



z Systems

# Hardware Management Console Web Services API Version 2.13.0

SC27-2627-00

Level 00a







z Systems

Hardware Management Console Web Services API  
Version 2.13.0

SC27-2627-00

**Note:**

Before using this information and the product it supports, read the information in “Safety” on page xix, Appendix E, “Notices,” on page 943, and *IBM Systems Environmental Notices and User Guide*, Z125-5823.

This edition, SC27-2627-00, applies to the IBM z Systems servers.

There might be a newer version of this document in a **PDF** file available on **Resource Link**. Go to <http://www.ibm.com/servers/resourcelink> and click **Library** on the navigation bar. A newer version is indicated by a lowercase, alphabetic letter following the form number suffix (for example: 00a, 00b, 01a, 01b).

© **Copyright IBM Corporation 2015.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Figures . . . . .</b>	<b>ix</b>	<b>Chapter 4. Asynchronous notification . . . . .</b>	<b>29</b>
<b>Tables . . . . .</b>	<b>xv</b>	JMS basics . . . . .	29
<b>Safety . . . . .</b>	<b>xix</b>	Connecting to the API message broker . . . . .	29
Safety notices . . . . .	xix	Per-session notification topics . . . . .	30
World trade safety information . . . . .	xix	Notification message formats . . . . .	31
Laser safety information . . . . .	xix	Common message characteristics . . . . .	31
Laser compliance . . . . .	xix	Status change notification. . . . .	33
<b>About this publication . . . . .</b>	<b>xxi</b>	Property change notification. . . . .	33
Related publications . . . . .	xxi	Inventory change notification . . . . .	35
Related HMC and SE console information . . . . .	xxi	Job completion notification . . . . .	35
Extending zBX connectivity options to Layer-2 . . . . .	xxi	Log entry notification . . . . .	36
Revision bars . . . . .	xxi	<b>Chapter 5. Data model definitions . . . . .</b>	<b>37</b>
Accessibility. . . . .	xxii	Data model concepts . . . . .	37
Accessibility features. . . . .	xxii	Objects in the data model. . . . .	37
Keyboard navigation. . . . .	xxii	Properties in the data model. . . . .	38
IBM and accessibility. . . . .	xxii	Shared data model schema elements . . . . .	39
How to send your comments . . . . .	xxii	Base managed object properties schema . . . . .	39
<b>Chapter 1. Introduction . . . . .</b>	<b>1</b>	<b>Chapter 6. General API services . . . . .</b>	<b>43</b>
Overview . . . . .	1	General API services operations summary . . . . .	43
Components of the API. . . . .	1	Session management services . . . . .	43
Enabling and accessing the API . . . . .	2	Query API Version . . . . .	44
Authentication and access control . . . . .	3	Logon . . . . .	45
Alternate HMC considerations . . . . .	3	Logoff . . . . .	48
Compatibility . . . . .	4	Get Notification Topics . . . . .	49
API versioning . . . . .	4	Asynchronous job processing . . . . .	51
Allowable changes within a major version . . . . .	4	Query Job Status . . . . .	52
Requirements on client applications. . . . .	5	Delete Completed Job Status. . . . .	54
Summary of API version updates . . . . .	5	<b>Chapter 7. Ensemble composition . . . . .</b>	<b>57</b>
<b>Chapter 2. Base definitions . . . . .</b>	<b>15</b>	Ensemble composition operations summary . . . . .	57
Data types. . . . .	15	Ensemble object . . . . .	58
Input and output representation . . . . .	16	Data Model . . . . .	58
Representing API data types in JSON. . . . .	16	Operations . . . . .	61
<b>Chapter 3. Invoking API operations . . . . .</b>	<b>19</b>	List Ensembles . . . . .	61
HTTP protocol standard . . . . .	19	Get Ensemble Properties . . . . .	63
Connecting to the API HTTP server . . . . .	19	Update Ensemble Properties. . . . .	65
HTTP header field usage . . . . .	19	List Ensemble Nodes . . . . .	67
Required request header fields . . . . .	20	Get Node Properties . . . . .	69
Optional request headers . . . . .	20	Add Node to Ensemble . . . . .	71
Standard response headers . . . . .	21	Remove Node from Ensemble . . . . .	74
Additional response headers. . . . .	22	Inventory service data. . . . .	75
Media types . . . . .	22	Usage notes . . . . .	76
HTTP status codes . . . . .	23	<b>Chapter 8. zBX infrastructure elements . . . . .</b>	<b>77</b>
Error response bodies . . . . .	24	zBX physical network overview . . . . .	77
Common request validation reason codes . . . . .	24	zBX infrastructure operations summary . . . . .	78
Common request processing reason codes . . . . .	26	zBX object. . . . .	79
Use of chunked response encoding . . . . .	26	Data model . . . . .	80
Filter query parameters . . . . .	27	Operations . . . . .	87
Regular expression syntax . . . . .	28	List zBXs of a CPC . . . . .	87
		List zBXs of an Ensemble. . . . .	89
		Get zBX Properties . . . . .	91

	Get EC/MCL Description of zBX (Node)	. . . . .	95		Data model	. . . . .	184
	Activate zBX (Node)	. . . . .	98		Operations	. . . . .	184
	Deactivate zBX (Node)	. . . . .	100		Set zBX (Node) Power Save	. . . . .	184
	Get zBX (Node) Audit Log	. . . . .	103		Set zBX (Node) Power Capping	. . . . .	186
	Get zBX (Node) Security Log	. . . . .	105		Energy Management for CPC object	. . . . .	189
	List zBX (Node) Hardware Messages	. . . . .	108		Data model	. . . . .	189
	Get zBX (Node) Hardware Message Properties	. . . . .	110		Operations	. . . . .	189
	Delete zBX (Node) Hardware Message	. . . . .	112		Set CPC Power Save	. . . . .	189
	Inventory service data	. . . . .	114		Set CPC Power Capping	. . . . .	191
	zBX Top-of-Rack switches	. . . . .	115		Set zCPC Power Save	. . . . .	194
	Data model	. . . . .	115		Set zCPC Power Capping	. . . . .	195
	Operations	. . . . .	117		Get CPC Energy Management Data	. . . . .	197
	List Top-of-Rack Switches of a zBX	. . . . .	117		Energy Management for BladeCenter object	. . . . .	199
	Get Top-of-Rack Switch Properties	. . . . .	119		Data model	. . . . .	199
	Get Top-of-Rack Switch Port Details	. . . . .	121		Operations	. . . . .	199
	Update Top-of-Rack Switch Port Properties	. . . . .	123		Set BladeCenter Power Save	. . . . .	199
	Add MAC Filters to Top-of-Rack Switch Port	. . . . .	125		Set BladeCenter Power Capping	. . . . .	201
	Remove MAC Filters from Top-of-Rack Switch Port	. . . . .	127		Energy Management for Blade object	. . . . .	204
	Add Top-of-Rack Switch Port to Virtual Networks	. . . . .	129		Data model	. . . . .	204
	Remove Top-of-Rack Switch Port from the Virtual Networks	. . . . .	131		Operations	. . . . .	204
	Rack object	. . . . .	133		Set Blade Power Save	. . . . .	204
	Data model	. . . . .	133		Set Blade Power Capping	. . . . .	206
	Operations	. . . . .	134		<b>Chapter 10. Virtualization management.</b>	. . . . .	<b>209</b>
	List Racks of a zBX	. . . . .	134		Virtualization host operations summary	. . . . .	209
	Get Rack Properties	. . . . .	136		Virtual server operations summary	. . . . .	210
	Inventory service data	. . . . .	137		Virtualization Host object	. . . . .	211
	BladeCenter object	. . . . .	138		Data model	. . . . .	211
	Data model	. . . . .	138		Operations	. . . . .	222
	Operations	. . . . .	141		List Virtualization Hosts of a zBX (Node)	. . . . .	222
	List BladeCenters in a Rack	. . . . .	141		List Virtualization Hosts of a Node	. . . . .	224
	List BladeCenters in a zBX	. . . . .	143		List Virtualization Hosts of an Ensemble	. . . . .	227
	Get BladeCenter Properties	. . . . .	145		List Virtualization Hosts of a CPC	. . . . .	229
	Activate BladeCenter	. . . . .	147		Get Virtualization Host Properties	. . . . .	232
	Deactivate BladeCenter	. . . . .	150		Update Virtualization Host Properties	. . . . .	237
	Inventory service data	. . . . .	152		List Virtual Switches	. . . . .	239
	Blade object	. . . . .	153		Get Virtual Switch Properties	. . . . .	240
	Data model	. . . . .	153		Create IEDN Virtual Switch	. . . . .	243
	Operations	. . . . .	158		Create QDIO Virtual Switch	. . . . .	246
	List Blades in a BladeCenter	. . . . .	158		Get Switch Controllers	. . . . .	249
	List Blades in a zBX	. . . . .	160		Update Virtual Switch	. . . . .	251
	Get Blade Properties	. . . . .	163		Delete Virtual Switch	. . . . .	255
	Activate a Blade	. . . . .	167		Activating a Virtualization Host	. . . . .	256
	Deactivate a Blade	. . . . .	169		Deactivating a Virtualization Host	. . . . .	257
	Create IEDN Interface for a DataPower XI50z Blade	. . . . .	171		SMAPI Error Response Body	. . . . .	257
	Delete IEDN Interface for a DataPower XI50z Blade	. . . . .	174		Inventory service data	. . . . .	258
	Inventory service data	. . . . .	175		Virtual Server Object	. . . . .	260
	<b>Chapter 9. Energy management</b>	. . . . .	<b>179</b>		Data Model	. . . . .	261
	Groups	. . . . .	180		Operations	. . . . .	286
	Special states	. . . . .	181		List Virtual Servers of a zBX (Node)	. . . . .	287
	Power saving	. . . . .	182		List Virtual Servers of a Node	. . . . .	289
	Group power saving	. . . . .	182		List Virtual Servers of an Ensemble	. . . . .	291
	Power capping	. . . . .	182		List Virtual Servers of a CPC	. . . . .	294
	Group capping	. . . . .	182		List Virtual Servers of a Virtualization Host	. . . . .	297
	Energy management operations summary	. . . . .	183		Create Virtual Server	. . . . .	299
	Energy Management for zBX (Node) object	. . . . .	184		Delete Virtual Server	. . . . .	304
					Get Virtual Server Properties	. . . . .	306
					Update Virtual Server Properties	. . . . .	314
					Create Network Adapter	. . . . .	319
					Delete Network Adapter	. . . . .	322

Get Network Adapter Properties . . . . .	323	Remove Virtualization Host Storage Resource from Virtualization Host Storage Group . . . . .	420
Update Network Adapter . . . . .	325	Notifications. . . . .	421
Reorder Network Adapter . . . . .	328	Inventory service data . . . . .	422
Create Virtual Disk . . . . .	330	Usage notes . . . . .	422
Delete Virtual Disk . . . . .	334		
Get Virtual Disk Properties . . . . .	336	<b>Chapter 12. Virtual network management. . . . .</b>	<b>423</b>
Update Virtual Disk Properties . . . . .	338	Virtual network management operations summary	423
Reorder Virtual Disks . . . . .	341	Virtual Network object . . . . .	423
Activate Virtual Server . . . . .	343	Data model . . . . .	424
Deactivate Virtual Server . . . . .	345	List Virtual Networks . . . . .	424
Mount Virtual Media . . . . .	348	Get Virtual Network Properties . . . . .	426
Mount Virtual Media Image . . . . .	350	Update Virtual Network Properties . . . . .	427
Unmount Virtual Media . . . . .	352	Create Virtual Network . . . . .	430
Migrate Virtual Server . . . . .	354	Delete Virtual Network . . . . .	432
Initiate Virtual Server Dump . . . . .	356	List Members of a Virtual Network . . . . .	434
Inventory service data . . . . .	357	Inventory service data . . . . .	436
		<b>Chapter 13. Workload resource group management. . . . .</b>	<b>439</b>
<b>Chapter 11. Storage Management. . . . .</b>	<b>363</b>	Overview. . . . .	439
Terms . . . . .	363	Workload resource group operations summary	440
Object model overview . . . . .	364	Workload Resource Group object . . . . .	443
Storage management operations summary. . . . .	364	Data model . . . . .	443
Storage Resource object . . . . .	366	List Workload Resource Groups of an Ensemble	448
Data model . . . . .	366	Get Workload Resource Group Properties . . . . .	450
Operations . . . . .	367	Create Workload Resource Group . . . . .	452
List Storage Resources . . . . .	367	Delete Workload Resource Group . . . . .	455
Get Storage Resource Properties . . . . .	370	Update Workload Resource Group . . . . .	456
Create Storage Resource . . . . .	371	List Virtual Servers of a Workload Resource Group . . . . .	458
Update Storage Resource Properties . . . . .	373	Add Virtual Server to a Workload Resource Group . . . . .	461
Delete Storage Resource . . . . .	375	Remove Virtual Server from a Workload Resource Group . . . . .	463
Export World Wide Port Names List. . . . .	377	List Groups of Virtual Servers of a Workload Resource Group . . . . .	465
Import Storage Access List . . . . .	379	Add Group of Virtual Servers to a Workload Resource Group . . . . .	467
List Virtualization Host Storage Resources of a Storage Resource . . . . .	381	Remove Group of Virtual Servers from a Workload Resource Group . . . . .	469
Inventory service data . . . . .	383	List Workload Element Groups of a Workload Resource Group . . . . .	471
Virtualization Host Storage Resource object . . . . .	384	Add Workload Element Group to a Workload Resource Group . . . . .	472
Data model . . . . .	384	Remove Workload Element Group from a Workload Resource Group . . . . .	474
Operations . . . . .	386	Performance Policy object . . . . .	476
List Virtualization Host HBA Ports . . . . .	386	Data model . . . . .	476
List Virtualization Host Storage Resources. . . . .	388	Notifications of property changes to performance policies . . . . .	480
Get Virtualization Host Storage Resource Properties . . . . .	391	List Performance Policies . . . . .	481
Create Virtualization Host Storage Resource . . . . .	395	Get Performance Policy Properties . . . . .	483
Delete Virtualization Host Storage Resource . . . . .	398	Create Performance Policy . . . . .	486
Add Virtualization Host Storage Resource Paths	400	Delete Performance Policy . . . . .	488
Remove Virtualization Host Storage Resource Paths . . . . .	403	Update Performance Policy. . . . .	490
Discover Virtualization Host Storage Resources	406	Activate Performance Policy . . . . .	493
List Virtual Disks of a Virtualization Host Storage Resource . . . . .	408	Import Performance Policy . . . . .	494
Notifications. . . . .	410	Export Performance Policy . . . . .	496
Inventory service data . . . . .	410		
Virtualization Host Storage Group object . . . . .	411		
Data model . . . . .	411		
Operations . . . . .	412		
List Virtualization Host Storage Groups . . . . .	412		
Get Virtualization Host Storage Group Properties . . . . .	414		
List Virtualization Host Storage Resources in a Virtualization Host Storage Group . . . . .	416		
Add Virtualization Host Storage Resource to Virtualization Host Storage Group . . . . .	418		

Performance management reports . . . . .	499		Password Rule operations summary . . . . .	609
Generate Workload Resource Groups Report			LDAP Server Definition operations summary	610
(Performance Management) . . . . .	500		Group operations summary . . . . .	610
Generate Workload Resource Group			CPC operations summary . . . . .	611
Performance Index Report . . . . .	504		Logical partitions operation summary . . . . .	612
Generate Workload Resource Group Resource			Activation profile operations summary . . . . .	613
Adjustments Report . . . . .	507		Capacity record operations summary . . . . .	614
Generate Virtual Servers Report . . . . .	511		Shared nested objects . . . . .	615
Generate Virtual Server CPU Utilization Report	515		Console object . . . . .	617
Generate Virtual Server Resource Adjustments			Data model . . . . .	618
Report. . . . .	518		Get Console Properties . . . . .	621
Generate Hypervisor Report . . . . .	523		Restart Console. . . . .	624
Generate Hypervisor Resource Adjustments			Make Console Primary . . . . .	626
Report. . . . .	531		Shutdown Console . . . . .	627
Generate Service Classes Report . . . . .	535		Reorder User Patterns . . . . .	628
Generate Service Class Resource Adjustments			Get Console Audit Log . . . . .	630
Report. . . . .	538		Get Console Security Log . . . . .	633
Generate Service Class Hops Report. . . . .	543		List Console Hardware Messages. . . . .	637
Generate Service Class Virtual Server Topology			Get Console Hardware Message Properties . . . . .	639
Report. . . . .	548		Delete Console Hardware Message . . . . .	640
Generate Load Balancing Report . . . . .	556		Inventory service data . . . . .	641
Workload Element Group object . . . . .	558		User-related-access permission. . . . .	645
Data model . . . . .	558		User object . . . . .	646
List Workload Element Groups of an Ensemble	561		Data model . . . . .	646
Get Workload Element Group Properties . . . . .	562		List Users . . . . .	650
Create Workload Element Group . . . . .	563		Get User Properties . . . . .	652
Delete Workload Element Group . . . . .	566		Update User Properties . . . . .	654
Update Workload Element Group . . . . .	567		Add User Role to User . . . . .	657
List Virtual Servers of a Workload Element			Remove User Role from User . . . . .	659
Group . . . . .	568		Create User . . . . .	661
Add Virtual Server to a Workload Element			Delete User . . . . .	664
Group . . . . .	569		Inventory service data . . . . .	666
Remove Virtual Server from a Workload			User Role object . . . . .	667
Element Group . . . . .	571		Data model . . . . .	668
Availability Policy object . . . . .	572		List User Roles . . . . .	670
Data model . . . . .	573		Get User Role Properties . . . . .	673
List Availability Policies . . . . .	575		Update User Role Properties . . . . .	675
Get Availability Policy Properties. . . . .	577		Add Permission to User Role . . . . .	677
Create Availability Policy . . . . .	578		Remove Permission from User Role . . . . .	680
Delete Availability Policy . . . . .	580		Create User Role . . . . .	682
Update Availability Policy . . . . .	581		Delete User Role . . . . .	684
Activate Availability Policy . . . . .	583		Inventory service data . . . . .	686
Ensemble Availability Management reports . . . . .	584		Task object . . . . .	687
Generate Workload Resource Groups Report			Data model . . . . .	687
(Ensemble Availability Management) . . . . .	586		List Tasks. . . . .	688
Generate Workload Resource Group Availability			Get Task Properties . . . . .	689
Status Report . . . . .	589		Inventory service data . . . . .	691
Generate Virtual Servers Report (Ensemble			User Pattern object . . . . .	691
Availability Management) . . . . .	594		Data model . . . . .	692
Generate Availability Status Report . . . . .	597		List User Patterns . . . . .	693
Get Performance Management Velocity Level			Get User Pattern Properties. . . . .	695
Range Mappings . . . . .	599		Update User Pattern Properties . . . . .	697
Inventory service data . . . . .	601		Create User Pattern . . . . .	699
			Delete User Pattern . . . . .	701
			Inventory service data . . . . .	702
<b>Chapter 14. Core z Systems resources 607</b>			Password Rule object. . . . .	703
Operations Summary. . . . .	607		Data model . . . . .	704
Console operations summary . . . . .	607		List Password Rules . . . . .	706
User operations summary . . . . .	608		Get Password Rule Properties . . . . .	708
User Role operations summary . . . . .	608		Update Password Rule Properties . . . . .	710
Task operations summary . . . . .	609		Create Password Rule . . . . .	712
User Pattern operations summary . . . . .	609			



	Delete Password Rule . . . . .	714	List Reset Activation Profiles . . . . .	834	
	Inventory service data . . . . .	715	Get Reset Activation Profile Properties . . . . .	836	
	LDAP Server Definition object . . . . .	717	Update Reset Activation Profile Properties . . . . .	837	
	Data model . . . . .	718		Inventory service data . . . . .	838
	List LDAP Server Definitions . . . . .	720	Image activation profile . . . . .	839	
	Get LDAP Server Definition Properties . . . . .	722	Data model . . . . .	839	
	Update LDAP Server Definition Properties . . . . .	724	List Image Activation Profiles . . . . .	854	
	Create LDAP Server Definition . . . . .	725	Get Image Activation Profile Properties . . . . .	856	
	Delete LDAP Server Definition . . . . .	728	Update Image Activation Profile Properties . . . . .	859	
	Inventory service data . . . . .	729		Inventory service data . . . . .	861
	Group Object . . . . .	730	Load activation profile . . . . .	861	
	Data model . . . . .	731	Data model . . . . .	861	
	List Custom Groups . . . . .	732	List Load Activation Profiles . . . . .	863	
	Get Custom Group Properties . . . . .	734	Get Load Activation Profile Properties . . . . .	865	
	Create Custom Group . . . . .	735	Update Load Activation Profile Properties . . . . .	867	
	Delete Custom Group . . . . .	737		Inventory service data . . . . .	868
	Add Member to Custom Group . . . . .	738	Group profile . . . . .	869	
	Remove Member from Custom Group . . . . .	740	Data model . . . . .	869	
	List Custom Group Members . . . . .	741	List Group Profiles . . . . .	869	
	Inventory service data . . . . .	743	Get Group Profile Properties . . . . .	871	
	CPC object . . . . .	743	Update Group Profile Properties . . . . .	873	
	Data model . . . . .	743		Inventory service data . . . . .	874
	List CPC Objects . . . . .	754	Capacity records . . . . .	875	
	List Ensemble CPC Objects . . . . .	756	Data model . . . . .	875	
	Get CPC Properties . . . . .	758	List Capacity Records . . . . .	877	
	Update CPC Properties . . . . .	764	Get Capacity Record Properties . . . . .	878	
	Activate CPC . . . . .	766		Inventory service data . . . . .	879
	Deactivate CPC . . . . .	768			
	Import Profiles . . . . .	770			
	Export Profiles . . . . .	771			
	Add Temporary Capacity . . . . .	772			
	Remove Temporary Capacity . . . . .	774			
	Swap Current Time Server . . . . .	776			
	Set STP Configuration . . . . .	777			
	Change STP-only Coordinated Timing Network . . . . .	778			
	Join STP-only Coordinated Timing Network . . . . .	780			
	Leave STP-only Coordinated Timing Network . . . . .	781			
	Get CPC Audit Log . . . . .	782			
	Get CPC Security Log . . . . .	784			
	List CPC Hardware Messages . . . . .	787			
	Get CPC Hardware Message Properties . . . . .	789			
	Delete CPC Hardware Message . . . . .	791			
	Inventory service data . . . . .	792			
	Logical Partition object . . . . .	793			
	Data model . . . . .	793			
	List Logical Partitions of CPC . . . . .	808			
	Get Logical Partition Properties . . . . .	810			
	Update Logical Partition Properties . . . . .	813			
	Activate Logical Partition . . . . .	814			
	Deactivate Logical Partition . . . . .	816			
	Reset Normal . . . . .	818			
	Reset Clear . . . . .	820			
	Load Logical Partition . . . . .	822			
	PSW Restart . . . . .	824			
	Start Logical Partition . . . . .	825			
	Stop Logical Partition . . . . .	827			
	SCSI Load . . . . .	828			
	SCSI Dump . . . . .	830			
	Inventory service data . . . . .	832			
	Reset activation profile . . . . .	833			
	Data model . . . . .	833			
			<b>Chapter 15. Inventory and metrics services. . . . .</b>	<b>881</b>	
			Inventory services operations summary . . . . .	881	
			Metrics service operations summary . . . . .	881	
			Inventory service . . . . .	882	
			Get Inventory . . . . .	882	
			Metrics service . . . . .	888	
			Create Metrics Context . . . . .	888	
			Get Metrics . . . . .	891	
			Delete Metrics Context . . . . .	895	
			<b>Chapter 16. zManager metric groups</b>	<b>897</b>	
			Monitors dashboard metric groups . . . . .	897	
			BladeCenter temperature and power metric group . . . . .	897	
			Blade power . . . . .	898	
			Channels . . . . .	898	
			CPC overview . . . . .	898	
			zBX (Node) overview . . . . .	899	
			Logical partitions . . . . .	900	
			zCPC environmentals and power . . . . .	900	
			zCPC processors . . . . .	901	
			Blade CPU and memory metric group . . . . .	902	
			Cryptos . . . . .	902	
			Flash memory adapters . . . . .	903	
			RoCE adapters . . . . .	903	
			Ensemble power . . . . .	903	
			Performance management metrics groups . . . . .	904	
			Virtual server CPU and memory metrics group . . . . .	904	
			Virtualization host CPU and memory metrics group . . . . .	905	
			Workload service class data metrics group . . . . .	907	

Network management metrics . . . . . 907  
    Virtualization host and virtual server metrics . . . . . 908  
    Optimizer network metrics . . . . . 915  
Physical switches . . . . . 918  
    Top-of-rack switch ports metrics . . . . . 919  
    ESM switch port metrics . . . . . 921

**Appendix A. XML document structure of a performance policy . . . . . 923**

XML structure of a `ServiceClasses` element . . . . . 923  
Sample XML document for a performance policy . . . . . 927

**| Appendix B. Enum values for a type of managed objects within User Roles 931**

**| Appendix C. Enum values for the User Role object . . . . . 933**

**| Appendix D. Enum values for the Task object. . . . . 935**

**Appendix E. Notices . . . . . 943**

Trademarks . . . . . 944

Class A Notices. . . . . 945

**Index . . . . . 949**

## Figures

1. Logon: Request . . . . .	48	49. Get Top-of-Rack Switch Properties: Request	120
2. Logon: Response . . . . .	48	50. Get Top-of-Rack Switch Properties: Response	121
3. Logoff: Request . . . . .	49	51. Get Top-of-Rack Switch Port Details: Request	123
4. Logoff: Response . . . . .	49	52. Get Top-of-Rack Switch Port Details: Response . . . . .	123
5. Get Notification Topics: Request . . . . .	51	53. Update Top-of-Rack Switch Port Properties: Request . . . . .	125
6. Get Notification Topics: Response . . . . .	51	54. Update Top-of-Rack Switch Port Properties: Response . . . . .	125
7. Query Job Status: Request . . . . .	54	55. Add MAC Filters to Top-of-Rack Switch Port: Request . . . . .	127
8. Query Job Status: Response . . . . .	54	56. Add MAC Filters to Top-of-Rack Switch Port: Response . . . . .	127
9. Delete Completed Job Status: Request . . . . .	56	57. Remove MAC Filters from Top-of-Rack Switch Port: Request . . . . .	129
10. Delete Completed Job Status: Response	56	58. Remove Mac Filters from Top-of-Rack Switch Port: Response . . . . .	129
11. List Ensembles: Request . . . . .	63	59. Add Top-of-Rack Switch Port to Virtual Networks: Request . . . . .	131
12. List Ensembles: Response . . . . .	63	60. Add Top-of-Rack Switch Port to Virtual Networks: Response . . . . .	131
13. Get Ensemble Properties: Request . . . . .	64	61. Remove Top-of-Rack Switch Port from the Virtual Networks: Request . . . . .	133
14. Get Ensemble Properties: Response . . . . .	65	62. Remove Top-of-Rack Switch Port from the Virtual Networks: Response . . . . .	133
15. Update Ensemble Properties: Request . . . . .	67	63. List Racks of a zBX: Request . . . . .	135
16. Update Ensemble Properties: Response	67	64. List Racks of a zBX: Response . . . . .	136
17. List Ensemble Nodes: Request . . . . .	69	65. Get Rack Properties: Request . . . . .	137
18. List Ensemble Nodes: Response . . . . .	69	66. Get Rack Properties: Response . . . . .	137
19. Get Node Properties: Request . . . . .	71	67. Rack object: Sample inventory data . . . . .	138
20. Get Node Properties: Response . . . . .	71	68. List BladeCenters in a Rack: Request	143
21. Ensemble object: Sample inventory data	76	69. List BladeCenters in a Rack: Response	143
22. List zBXs of a CPC: Request . . . . .	89	70. List BladeCenters in a zBX: Request . . . . .	145
23. List zBXs of a CPC: Response . . . . .	89	71. List BladeCenters in a zBX: Response	145
24. List zBXs of an Ensemble: Request . . . . .	91	72. Get BladeCenter Properties: Request . . . . .	146
25. List zBXs of an Ensemble: Response . . . . .	91	73. Get BladeCenter Properties: Response	147
26. Get zBX Properties: Request . . . . .	93	74. Activate BladeCenter: Request . . . . .	149
27. Get zBX Properties: Response (Part 1). . . . .	94	75. Activate BladeCenter: Response . . . . .	149
28. Get zBX Properties: Response (Part 2). . . . .	95	76. Deactivate BladeCenter: Request . . . . .	151
29. Get EC/MCL Description of zBX (Node): Request . . . . .	96	77. Deactivate BladeCenter: Response . . . . .	152
30. Get EC/MCL Description of zBX (Node): Response (Part 1) . . . . .	97	78. BladeCenter object: Sample inventory data	153
31. Get EC/MCL Description of zBX (Node): Response (Part 2) . . . . .	98	79. List Blades in a BladeCenter: Request	160
32. Activate zBX (Node): Request . . . . .	100	80. List Blades in a BladeCenter: Response	160
33. Activate zBX (Node): Response . . . . .	100	81. List Blades in a zBX: Request . . . . .	162
34. Deactivate zBX (Node): Request . . . . .	102	82. List Blades in a zBX: Response . . . . .	163
35. Deactivate zBX (Node): Response . . . . .	103	83. Get Blade Properties: Request . . . . .	164
36. Get zBX (Node) Audit Log: Request . . . . .	104	84. Get Blade Properties: Response for blade of type "system-x" (similar for type "power") . . . . .	165
37. Get zBX (Node) Audit Log: Response	105	85. Get Blade Properties: Response for blade of type "dpx150z": . . . . .	166
38. Get zBX (Node) Security Log: Request	107	86. Get Blade Properties: Response for blade of type "isaopt": . . . . .	167
39. Get zBX (Node) Security Log: Response	108	87. Activate a Blade: Request . . . . .	169
40. List zBX (Node) Hardware Messages: Request	110	88. Activate a Blade: Response . . . . .	169
41. List zBX (Node) Hardware Messages: Response . . . . .	110	89. Deactivate a Blade: Request . . . . .	171
42. Get zBX (Node) Hardware Message Properties: Request . . . . .	112	90. Deactivate a Blade: Response . . . . .	171
43. Get zBX (Node) Hardware Message Properties: Response . . . . .	112		
44. Delete zBX (Node) Hardware Message: Request . . . . .	114		
45. Delete zBX (Node) Hardware Message: Response . . . . .	114		
46. zBX object: Sample inventory data . . . . .	115		
47. List Top-of-Rack Switches: Request . . . . .	119		
48. List Top-of-Rack Switches: Response . . . . .	119		

91. Activate a Blade: Sample inventory data for a blade of type "power". . . . .	176	129. List Virtual Servers of a Virtualization Host: Request . . . . .	299
92. Activate a Blade: Sample inventory data for a blade of type "system-x". . . . .	177	130. List Virtual Servers of a Virtualization Host: Response . . . . .	299
93. Energy management as applied throughout layers of enterprise management . . . . .	179	131. Create Virtual Server: Request for a virtual server of type "power-vm" . . . . .	304
94. Example of a Example of a zBX node group that contains a BladeCenter with blades. . . . .	180	132. Create Virtual Server: Response for a virtual server of type "power-vm" . . . . .	304
95. Example of a CPC group that contains a zCPC . . . . .	180	133. Delete Virtual Server: Request . . . . .	306
96. Example of a CPC (Ensemble) group that contains a zCPC and a BladeCenter . . . . .	181	134. Delete Virtual Server: Response . . . . .	306
97. Example of a CPC (Ensemble) group that contains a zCPC . . . . .	181	135. Get Virtual Server Properties: Request . . . . .	308
98. List Virtualization Hosts of a zBX (Node): Request . . . . .	224	136. Get Virtual Server Properties: Response for virtual servers of "power-vm" (Part 1) . . . . .	309
99. List Virtualization Hosts of a zBX (Node): Response . . . . .	224	137. Get Virtual Server Properties: Response for virtual servers of "power-vm" (part 2) . . . . .	310
100. List Virtualization Hosts of a Node: Request . . . . .	226	138. Get Virtual Server Properties: Response for virtual servers of type "prsm" . . . . .	311
101. List Virtualization Hosts of a Node: Response . . . . .	226	139. Get Virtual Server Properties: Response for virtual servers of type "x-hyp" (Part 1) . . . . .	312
102. List Virtualization Hosts of an Ensemble: Request . . . . .	228	140. Get Virtual Server Properties: Response for virtual servers of type "x-hyp" (Part 2) . . . . .	313
103. List Virtualization Hosts of an Ensemble: Response . . . . .	229	141. Get Virtual Server Properties: Response for virtual servers of type "zvm" . . . . .	314
104. List Virtualization Hosts of a CPC: Request . . . . .	231	142. Update Virtual Server Properties: Request for a virtual server of type "x-hyp" . . . . .	319
105. List Virtualization Hosts of a CPC: Response . . . . .	232	143. Update Virtual Server Properties: Response for a virtual server of type "x-hyp" . . . . .	319
106. Get Virtualization Host Properties: Request . . . . .	233	144. Create Network Adapter: Request for a virtual server of type "x-hyp" . . . . .	321
107. Get Virtualization Host Properties: Response for virtualization host of type "prsm" . . . . .	234	145. Create Network Adapter: Response for a virtual server of type "x-hyp" . . . . .	322
108. Get Virtualization Host Properties: Response for virtualization host of type "power-vm" . . . . .	235	146. Delete Network Adapter: Request. . . . .	323
109. Get Virtualization Host Properties: Response for virtualization host of type "x-hyp" . . . . .	236	147. Delete Network Adapter: Response . . . . .	323
110. Get Virtualization Host Properties: Response for virtualization host of type "zvm" . . . . .	237	148. Get Network Adapter Properties: Request . . . . .	325
111. List Virtual Switches: Request . . . . .	240	149. Get Network Adapter Properties: Response . . . . .	325
112. List Virtual Switches: Response . . . . .	240	150. Update Network Adapter: Request for a virtual server of type "x-hyp" . . . . .	328
113. Get Virtual Switch Properties: Request . . . . .	242	151. Update Network Adapter: Response for a virtual server of type "x-hyp" . . . . .	328
114. Get Virtual Switch Properties: Response for virtual switch of type "iedn" . . . . .	242	152. Reorder Network Adapter: Request . . . . .	330
115. Get Virtual Switch Properties: Response for virtual switch of type "qdio" . . . . .	243	153. Reorder Network Adapter: Response . . . . .	330
116. Get Switch Controllers: Request . . . . .	250	154. Create Virtual Disk: Request for a virtual server of type "power-vm" . . . . .	333
117. Get Switch Controllers: Response . . . . .	251	155. Create Virtual Disk: Response for a virtual server of type "power-vm" . . . . .	333
118. Virtualization Host object: Sample inventory data for a virtualization host of type "power-vm" . . . . .	259	156. Create Virtual Disk: Request for a virtual server of type "x-hyp". . . . .	333
119. Virtualization Host object: Sample inventory data for a virtualization host of type "prsm" . . . . .	260	157. Create Virtual Disk: Response for a virtual server of type "x-hyp". . . . .	334
120. Virtualization Host object: Sample inventory data for a virtualization host of type "x-hyp" . . . . .	260	158. Delete Virtual Disk: Request . . . . .	335
121. List Virtual Servers of a zBX (Node): Request . . . . .	288	159. Delete Virtual Disk: Response . . . . .	336
122. List Virtual Servers of a zBX (Node): Response . . . . .	289	160. Get Virtual Disk Properties: Request for a virtual server of type "power-vm". . . . .	337
123. List Virtual Servers of a Node: Request . . . . .	291	161. Get Virtual Disk Properties: Response for a virtual server of type "power-vm". . . . .	338
124. List Virtual Servers of a Node: Response . . . . .	291	162. Get Virtual Disk Properties: Request for a virtual server of type "x-hyp" . . . . .	338
125. List Virtual Servers of an Ensemble: Request . . . . .	293	163. Get Virtual Disk Properties: Response for a virtual server of type "x-hyp" . . . . .	338
126. List Virtual Servers of an Ensemble: Response . . . . .	294	164. Update Virtual Disk Properties: Request for a virtual server of type "x-hyp" . . . . .	341
127. List Virtual Servers of a CPC: Request . . . . .	296		
128. List Virtual Servers of a CPC: Response . . . . .	297		

165. Update Virtual Disk Properties: Response for a virtual server of type "x-hyp" . . . . .	341	205. Create Virtualization Host Storage Resource: Request . . . . .	398
166. Reorder Virtual Disks: Request . . . . .	343	206. Create Virtualization Host Storage Resource: Response . . . . .	398
167. Reorder Virtual Disks: Response . . . . .	343	207. Delete Virtualization Host Storage Resource: Request . . . . .	400
168. Activate Virtual Server: Request . . . . .	345	208. Delete Virtualization Host Storage Resource: Response . . . . .	400
169. Activate Virtual Server: Response . . . . .	345	209. Add Virtualization Host Storage Resource Paths: Request . . . . .	403
170. Deactivate Virtual Server: Request . . . . .	347	210. Add Virtualization Host Storage Resource Paths: Response . . . . .	403
171. Deactivate Virtual Server: Response . . . . .	348	211. List Virtual Disks of a Virtualization Host Storage Resource: Request . . . . .	410
172. Mount Virtual Media: Request . . . . .	350	212. List Virtual Disks of a Virtualization Host Storage Resource: Response . . . . .	410
173. Mount Virtual Media: Response . . . . .	350	213. List Virtualization Host Storage Groups: Request . . . . .	414
174. Unmount Virtual Media: Request . . . . .	353	214. List Virtualization Host Storage Groups: Response . . . . .	414
175. Unmount Virtual Media: Response . . . . .	354	215. Get Virtualization Host Storage Group Properties: Request . . . . .	416
176. Initiate Virtual Server Dump: Request . . . . .	357	216. Get Virtualization Host Storage Group Properties: Response . . . . .	416
177. Initiate Virtual Server Dump: Response . . . . .	357	217. List Virtual Networks: Request . . . . .	426
178. Virtual Server object: Sample inventory data for a virtual server of type "power-vm" (Part 1) . . . . .	359	218. List Virtual Networks: Response . . . . .	426
179. Virtual Server object: Sample inventory data for a virtual server of type "power-vm" (Part 2) . . . . .	360	219. Get Virtual Network Properties: Request . . . . .	427
180. Virtual Server object: Sample inventory data for a virtual server of type "prsm" . . . . .	361	220. Get Virtual Network Properties: Response . . . . .	427
181. Virtual Server object: Sample inventory data for a virtual server of type "x-hyp" . . . . .	362	221. Update Virtual Network Properties: Request . . . . .	430
182. Object model. . . . .	364	222. Update Virtual Network Properties: Response . . . . .	430
183. List Storage Resources: Request . . . . .	369	223. Create Virtual Network: Request . . . . .	432
184. List Storage Resources: Response . . . . .	369	224. Create Virtual Network: Response . . . . .	432
185. Get Storage Resource Properties: Request . . . . .	371	225. Delete Virtual Network: Request . . . . .	434
186. Get Storage Resource Properties: Response . . . . .	371	226. Delete Virtual Network: Response . . . . .	434
187. Create Storage Resource: Request . . . . .	373	227. List Members of a Virtual Network: Request . . . . .	436
188. Create Storage Resource: Response . . . . .	373	228. List Members of a Virtual Network: Response . . . . .	436
189. Update Storage Resource Properties: Request . . . . .	375	229. Virtual Network object: Sample inventory data. . . . .	437
190. Update Storage Resource Properties: Response . . . . .	375	230. List Workload Resource Groups of an Ensemble: Request . . . . .	449
191. Delete Storage Resource: Request . . . . .	376	231. List Workload Resource Groups of an Ensemble: Response . . . . .	450
192. Delete Storage Resource: Response . . . . .	377	232. Get Workload Resource Group Properties: Request . . . . .	451
193. Export World Wide Port Names List: WWPN list: Request . . . . .	379	233. Get Workload Resource Group Properties: Response . . . . .	452
194. Export World Wide Port Names List: WWPN list: Response . . . . .	379	234. Create Workload Resource Group: Request . . . . .	454
195. List Virtualization Host Storage Resources of a Storage Resource: Request . . . . .	382	235. Create Workload Resource Group: Response . . . . .	455
196. List Virtualization Host Storage Resources of a Storage Resource: Response . . . . .	383	236. Delete Workload Resource Group: Request . . . . .	456
197. Storage Resource object: Sample inventory data. . . . .	384	237. Delete Workload Resource Group: Response . . . . .	456
198. List Virtualization Host HBA Ports: Request . . . . .	387	238. Update Workload Resource Group: Request . . . . .	458
199. List Virtualization Host HBA Ports: Response . . . . .	388	239. Update Workload Resource Group: Response . . . . .	458
200. List Virtualization Host Storage Resources: Request . . . . .	391	240. List Virtual Servers of a Workload Resource Group: Request . . . . .	460
201. List Virtualization Host Storage Resources: Response . . . . .	391	241. List Virtual Servers of a Workload Resource Group: Response . . . . .	461
202. Get Virtualization Host Storage Resource Properties: Request . . . . .	393	242. Add Virtual Server to a Workload Resource Group: Request . . . . .	463
203. Get Virtualization Host Storage Resource Properties: Response for Virtualization Host of type "power-vm" or "x-hyp". . . . .	394	243. Add Virtual Server to a Workload Resource Group: Response . . . . .	463
204. Get Virtualization Host Storage Resource Properties: Response for Virtualization Host of type "zvm" . . . . .	395		

244. Remove Virtual Server from a Workload Resource Group: Request . . . . .	465	283. Generate Workload Resource Group Availability Status Report: Request . . . . .	594
245. Remove Virtual Server from a Workload Resource Group: Response . . . . .	465	284. Generate Virtual Servers Report: Request	596
246. List Groups of Virtual Servers of a Workload Resource Group: Request . . . . .	467	285. Generate Virtual Servers Report: Request	599
247. List Groups of Virtual Servers of a Workload Resource Group: Response . . . . .	467	286. Get Performance Management Velocity Level Range Mappings: Request . . . . .	601
248. Add Group of Virtual Servers to a Workload Resource Group: Request . . . . .	469	287. Workload Resource Group: Sample inventory data (Part 1) . . . . .	602
249. Add Group of Virtual Servers to a Workload Resource Group: Response . . . . .	469	288. Workload Resource Group: Sample inventory data (Part 2) . . . . .	603
250. Remove Group of Virtual Servers from a Workload Resource Group: Request . . . . .	471	289. Workload Resource Group: Sample inventory data (Part 3) . . . . .	604
251. Remove Group of Virtual Servers from a Workload Resource Group: Response . . . . .	471	290. Workload Resource Group: Sample inventory data (Part 4) . . . . .	605
252. List Performance Policies: Request . . . . .	482	291. Get Console Properties: Request . . . . .	621
253. List Performance Policies: Response . . . . .	483	292. Get Console Properties: Response (Part 1)	622
254. Get Performance Policy Properties: Request	484	293. Get Console Properties: Response (Part 2)	623
255. Get Performance Policy Properties: Response (Part 1) . . . . .	485	294. Get Console Properties: Response (Part 3)	624
256. Get Performance Policy Properties: Response (Part 2) . . . . .	486	295. Shutdown Console: Request. . . . .	628
257. Create Performance Policy: Request . . . . .	488	296. Shutdown Console: Response . . . . .	628
258. Create Performance Policy: Response	488	297. Reorder User Patterns: Request . . . . .	630
259. Delete Performance Policy: Request . . . . .	490	298. Reorder User Patterns: Response . . . . .	630
260. Delete Performance Policy: Response	490	299. Get Console Audit Log: Request . . . . .	632
261. Update Performance Policy: Request	492	300. Get Console Audit Log: Response. . . . .	633
262. Update Performance Policy: Response	493	301. Get Console Security Log: Request . . . . .	635
263. Activate Performance Policy: Request	494	302. Get Console Security Log: Response . . . . .	636
264. Activate Performance Policy: Response	494	303. List Console Hardware Messages: Request	638
265. Export Performance Policy: Request . . . . .	497	304. List Console Hardware Messages: Response	639
266. Export Performance Policy: Response	498	305. Get Console Hardware Message Properties: Request . . . . .	640
267. Relationship between reports and the properties used . . . . .	500	306. Get Console Hardware Message Properties: Response . . . . .	640
268. Generate Workload Resource Groups Report: Request . . . . .	504	307. Delete Console Hardware Message: Request	641
269. Generate Workload Resource Group Performance Index Report: Request . . . . .	506	308. Delete Console Hardware Message: Response	641
270. Generate Workload Resource Group Resource Adjustments Report: Request . . . . .	511	309. Console object: Sample inventory data (Part 1) . . . . .	642
271. Generate Virtual Servers Report: Request	515	310. Console object: Sample inventory data (Part 2) . . . . .	643
272. Generate Virtual Server CPU Utilization Report: Request. . . . .	518	311. Console object: Sample inventory data (Part 3) . . . . .	644
273. Generate Virtual Server Resource Adjustments Report: Request . . . . .	522	312. Console object: Sample inventory data (Part 4) . . . . .	645
274. Generate Hypervisor Report: Request	531	313. List Users: Request. . . . .	652
275. Generate Hypervisor Resource Adjustments Report: Request . . . . .	535	314. List Users: Response . . . . .	652
276. Generate Service Classes Report: Request	538	315. Get User Properties: Request . . . . .	654
277. Generate Service Class Resource Adjustments Report: Request. . . . .	543	316. Get User Properties: Response . . . . .	654
278. Generate Service Class Hops Report: Request	548	317. Update User Properties: Request . . . . .	657
279. Generate Service Class Virtual Server Topology Report: Request . . . . .	556	318. Update User Properties: Response . . . . .	657
280. Generate Load Balancing Report: Request	558	319. Add User Role to User: Request . . . . .	659
281. Relationship between reports and the properties used . . . . .	585	320. Add User Role to User: Response . . . . .	659
282. Generate Workload Resource Groups Report (Ensemble Availability Management): Request	589	321. Remove User Role from User: Request	661
		322. Remove User Role from User: Response	661
		323. Create User: Request . . . . .	664
		324. Create User: Response . . . . .	664
		325. Delete User: Request . . . . .	665
		326. Delete User: Response . . . . .	666
		327. User object: Sample inventory data . . . . .	667
		328. List User Roles: Request . . . . .	672
		329. List User Roles: Response . . . . .	673
		330. Get User Role Properties: Request. . . . .	674
		331. Get User Role Properties: Response . . . . .	675

332. Update User Role Properties: Request	677	388. Delete Custom Group: Response	738
333. Update User Role Properties: Response	677	389. Add Member to Custom Group: Request	740
334. Add Permission to User Role: Request	679	390. Add Member to Custom Group: Response	740
335. Add Permission to User Role: Response	680	391. Remove Member from Custom Group: Request	741
336. Remove Permission from User Role: Request	682	392. Remove Member from Custom Group: Response	741
337. Remove Permission from User Role: Response	682	393. List Custom Group Members: Request	743
338. Create User Role: Request	684	394. List Custom Group Members: Response	743
339. Create User Role: Response	684	395. List CPC Objects: Request	756
340. Delete User Role: Request	685	396. List CPC Objects: Response	756
341. Delete User Role: Response	686	397. List Ensemble CPC Objects: Request	758
342. User Role object: Sample inventory data	687	398. List Ensemble CPC Objects: Response	758
343. List Tasks: Request	689	399. Get CPC Properties: Request	759
344. List Tasks: Response	689	400. Get CPC Properties: Response (Part 1)	760
345. Get Task Properties: Request	690	401. Get CPC Properties: Response (Part 2)	761
346. Get Task Properties: Response	691	402. Get CPC Properties: Response (Part 3)	762
347. Task object: Sample inventory data	691	403. Get CPC Properties: Response (Part 4)	763
348. List User Patterns: Request	695	404. Get CPC Properties: Response (Part 5)	764
349. List User Patterns: Response	695	405. Get CPC Audit Log: Request	783
350. Get User Pattern Properties: Request	696	406. Get CPC Audit Log: Response	784
351. Get User Pattern Properties: Response	697	407. Get CPC Security Log: Request	786
352. Update User Pattern Properties: Request	698	408. Get CPC Security Log: Response	787
353. Update User Pattern Properties: Response	699	409. List CPC Hardware Messages: Request	789
354. Create User Pattern: Request	701	410. List CPC Hardware Messages: Response	789
355. Create User Pattern: Response	701	411. Get CPC Hardware Message Properties: Request	791
356. Delete User Pattern: Request	702	412. Get CPC Hardware Message Properties: Response	791
357. Delete User Pattern: Response	702	413. Delete CPC Hardware Message: Request	792
358. User Pattern object: Sample inventory data	703	414. Delete CPC Hardware Message: Response	792
359. List Password Rules: Request	708	415. List Logical Partitions of CPC: Request	810
360. List Password Rules: Response	708	416. List Logical Partitions of CPC: Response	810
361. Get Password Rule Properties: Request	709	417. Get Logical Partition Properties: Request	811
362. Get Password Rule Properties: Response	710	418. Get Logical Partition Properties: Response (Part 1)	812
363. Update Password Rule Properties: Request	712	419. Get Logical Partition Properties: Response (Part 2)	813
364. Update Password Rule Properties: Response	712	420. List Reset Activation Profiles: Request	835
365. Create Password Rule: Request	714	421. List Reset Activation Profiles: Response	835
366. Create Password Rule: Response	714	422. Get Reset Activation Profile Properties: Request	837
367. Delete Password Rule: Request	715	423. Get Reset Activation Profile Properties: Response	837
368. Delete Password Rule: Response	715	424. List Image Activation Profiles: Request	856
369. Password Rule object: Sample inventory data	717	425. List Image Activation Profiles: Response	856
370. List LDAP Server Definitions: Request	721	426. Get Image Activation Profile Properties: Request	857
371. List LDAP Server Definitions: Response	722	427. Get Image Activation Profile Properties: Response (Part 1)	858
372. Get LDAP Server Definition Properties: Request	723	428. Get Image Activation Profile Properties: Response (Part 2)	859
373. Get LDAP Server Definition Properties: Response	723	429. List Load Activation Profiles: Request	864
374. Update LDAP Server Definition Properties: Request	725	430. List Load Activation Profiles: Response	865
375. Update LDAP Server Definition Properties: Response	725	431. Get Load Activation Profile Properties: Request	866
376. Create LDAP Server Definition: Request	727	432. Get Load Activation Profile Properties: Response	867
377. Create LDAP Server Definition: Response	727	433. List Group Profiles: Request	871
378. Delete LDAP Server Definition: Request	729	434. List Group Profiles: Response	871
379. Delete LDAP Server Definition: Response	729	435. Get Group Profile Properties: Request	872
380. LDAP Server Definition object: Sample inventory data	730		
381. List Custom Groups: Request	733		
382. List Custom Groups: Response	734		
383. Get Custom Group Properties: Request	735		
384. Get Custom Group Properties: Response	735		
385. Create Custom Group: Request	737		
386. Create Custom Group: Response	737		
387. Delete Custom Group: Request	738		

436. Get Group Profile Properties: Response	873	442. Get Metrics: Response	895
437. Get Inventory: Request	887	443. Delete Metrics Context: Request	896
438. Get Inventory: Response	887	444. Delete Metrics Context: Response	896
439. Create Metrics Context: Request	890	445. Policy XML example, Part 1.	928
440. Create Metrics Context: Response	891	446. Policy XML example, Part 2.	929
441. Get Metrics: Request	894		



## Tables

1. Summary of updates by API version number	5		37. Delete zBX (Node) Hardware Message: HTTP	
2. Primitive data types.	15		status and reason codes	113
3. Compound data types	15		38. zBX Top-of-Rack switches: properties	115
4. Primitive data types in JSON notation	16		39. zBX Top-of-Rack switches: tor-port-info	
5. Compound data types in JSON notation	17		nested object properties	116
6. General API services: operations summary	43		40. Rack object: base managed object properties	
7. General API services: URI variables	43		specializations	133
8. topic-info object	50		41. Rack object: class specific properties	134
9. Ensemble composition: operations summary	57		42. BladeCenter object: base managed object	
10. Ensemble composition: URI variables	58		properties specializations	138
11. Ensemble object: base managed object			43. BladeCenter object: class specific properties	139
properties specializations	58		44. BladeCenter object: energy management	
12. Ensemble object: class specific properties	59		related additional properties	139
13. Ensemble composition: energy management			45. Activate BladeCenter: HTTP status and	
related additional properties	60		reason codes	149
14. Ensemble composition: MAC address prefix			46. Activate BladeCenter: Job status and reason	
nested object related additional properties	60		codes	149
15. Ensemble composition: node properties	61		47. Activate BladeCenter: HTTP status and	
16. zBX infrastructure: operations summary	78		reason codes	151
17. zBX infrastructure: URI variables	79		48. Deactivate BladeCenter: Job status and reason	
18. zBX object: base managed object properties			codes	151
specializations	80		49. Blade object: base managed object properties	
19. zBX object: class specific properties	81		specializations	153
20. zBX object: class specific additional properties			50. Blade object: class specific properties	154
(for zBX node)	81		51. Blade object: energy management related	
21. zBX object: ec-mcl-description nested object			additional properties	155
properties (for zBX node)	84		52. Blade object: IEDN interface nested object	
22. zBX object: action nested object properties (for			properties.	157
zBX node)	84		53. Energy management: operations summary	183
23. zBX object: ec nested object properties (for zBX			54. Energy management: URI variables	183
node)	84		55. Virtualization management - virtualization	
24. zBX object:mcl nested object properties (for			host: operations summary	209
zBX node)	84		56. Virtualization management- virtualization	
25. zBX object: ipv6-info properties (for zBX node)	84		host: URI variables.	210
26. zBX object: hardware-message object properties			57. Virtualization management - virtual server:	
(for zBX node)	85		operations summary	210
27. zBX object: energy management related			58. Virtualization management - virtual server:	
additional properties (for zBX node)	85		URI variables	211
28. Get EC/MCL Description of zBX (Node):			59. Virtualization Host object: base managed	
HTTP status and reason codes	96		object properties specializations	212
29. Activate zBX (Node): HTTP status and reason			60. Virtualization Host object: class specific	
codes	100		additional properties	213
30. Activate zBX (Node): Job status and reason			61. iedn-virtual-switch object properties	216
codes	100		62. real-uplink object properties	218
31. Deactivate zBX (Node): HTTP status and			63. qdio-virtual-switch object properties	219
reason codes	102		64. Virtual Server object: base managed object	
32. Dectivate zBX (Node): Job status and reason			properties specializations	261
codes	102		65. Virtual Server object: class specific additional	
33. Get zBX (Node) Audit Log: HTTP status and			properties.	264
reason codes	104		66. Virtual Server object: virtual server	
34. Get zBX (Node) Security Log: HTTP status			availability status with reasons nested object	
and reason codes	107		properties.	278
35. List zBX (Node) Hardware Messages: HTTP			67. Virtual Server object: virtual server	
status and reason codes	109		performance policy nested object	279
36. Get zBX (Node) Hardware Message			68. Virtual Server object: virtual server	
Properties: HTTP status and reason codes	111		availability policy nested object	279

69. mac-prefix object properties . . . . .	280	108. Format of a perf-status-data-point object	502
70. network-adapter-power object properties	280	109. Format of a service-class-pi-data object	505
71. network-adapter-x-hyp object properties	280	110. Format of a pi-data-point object . . . . .	505
72. network-adapter-zvm object properties	281	111. Format of a report-hypervisor-details object	524
73. network-adapter-prsm object properties	283	112. Format of a PowerVM report-hypervisor- virtual-servers object . . . . .	525
74. Virtual disk object properties . . . . .	284	113. Format of an x Hyp report-hypervisor-virtual- servers object . . . . .	526
75. fullpack-virtual-disk object properties	285	114. Format of a z/VM report-hypervisor-virtual- servers object . . . . .	527
76. fullpack-virtual-disk-zvm object properties	285	115. Format of a PR/SM report-hypervisor-virtual- servers object . . . . .	529
77. storage-group-based-virtual-disk object properties . . . . .	285	116. Format of an equivalent-workload-service- class object . . . . .	545
78. linked-virtual-disk object properties . . . . .	286	117. Format of a hop-entry object . . . . .	545
79. Valid values for the access-modes property of a virtual disk object . . . . .	286	118. Format of a hops-report-statistics object	546
80. Activate Virtual Server: HTTP status and reason codes . . . . .	345	119. Format of a hop-application-env object	546
81. Storage management: ensemble-level storage operations . . . . .	364	120. Format of a hop-application-env-virtual- server object . . . . .	547
82. Storage management: virtualization host storage operations . . . . .	365	121. Format of an equivalent-workload-service- class object . . . . .	550
83. Storage management: storage group operations . . . . .	365	122. Format of a topo-hop object. . . . .	550
84. Storage management: URI variables . . . . .	366	123. Format of a topo-virtual-server-node object	550
85. Storage Resource object: base managed object properties specializations . . . . .	366	124. Format of an appl-env-vs-response-entry object . . . . .	552
86. Storage Resource object: class specific properties . . . . .	367	125. Format of an appl-env-vs-utilization-entry object . . . . .	553
87. Virtualization Host Storage Resource object properties . . . . .	384	126. Format of a child-virtual-server-node-link object . . . . .	554
88. Virtualization Host Storage Resource object: path-information-fcp object properties . . . . .	385	127. Workload object: base managed object properties specializations . . . . .	559
89. Virtualization Host Storage Resource object: path-information-eckd object properties . . . . .	386	128. Workload Element Group object: type-specific properties . . . . .	559
90. Virtualization Host Storage Group object properties . . . . .	411	129. Workload Element Group object: Workload element group availability status nested object properties . . . . .	560
91. Virtualization Host Storage Group object: free-space-information object properties . . . . .	411	130. Workload Element Group object: Workload element group availability status with reasons nested object properties . . . . .	561
92. Virtual network management: operations summary . . . . .	423	131. Availability Policy object: type-specific properties . . . . .	573
93. Virtual network management: URI variables	423	132. Availability Policy object: Workload element group override nested object properties . . . . .	575
94. Virtual Network object: base managed object properties specializations . . . . .	424	133. Format of a workload-report-entry object	587
95. Virtual Network object: class specific additional properties . . . . .	424	134. Format of nested avail-status-data-point object . . . . .	588
96. Workload resource group management: operations summary . . . . .	440	135. Format of a policy-activation-record object	590
97. Workload management: URI variables	443	136. Format of availability-status-record object	590
98. Workload object: base managed object properties specializations . . . . .	443	137. Format of a parent object . . . . .	593
99. Workload object: type-specific properties	444	138. Core z Systems resources - Console: operations summary . . . . .	607
100. perf-policy-summary-object . . . . .	447	139. Core z Systems resources - Console: URI variables . . . . .	608
101. avail-policy-summary-object . . . . .	447	140. Core z Systems resources - User: operations summary . . . . .	608
102. Performance Policy object: type-specific properties . . . . .	476	141. Core z Systems resources - User: URI variables . . . . .	608
103. Performance Policy object: Service class nested object properties . . . . .	478	142. Core z Systems resources - User Role: operations summary . . . . .	608
104. Performance Policy object: classification rule nested object properties . . . . .	479	143. Core z Systems resources - User Role: URI variables . . . . .	609
105. Performance Policy object: filter nested object properties . . . . .	480		
106. Format of a workload-report-entry object	501		
107. Format of a cpu-utilization-range object	502		

144. Core z Systems resources - Task: operations summary . . . . .	609	182. ipv6-info properties . . . . .	620
145. Core z Systems resources - Task: URI variables . . . . .	609	183. machine-info properties . . . . .	620
146. Core z Systems resources - User Pattern: operations summary . . . . .	609	184. hardware-message object properties . . . . .	620
147. Core z Systems resources - User Pattern: URI variables . . . . .	609	185. Reorder User Patterns: HTTP status and reason codes . . . . .	629
148. Core z Systems resources - Password Rule: operations summary . . . . .	609	186. log-entry-info object properties. . . . .	631
149. Core z Systems resources - Password Rule: URI variables . . . . .	610	187. event-details-info object properties . . . . .	631
150. Core z Systems resources - LDAP Server Definition: operations summary . . . . .	610	188. event-data-item-info object properties	631
151. Core z Systems resources - LDAP Server Definition: URI variables. . . . .	610	189. List Console Hardware Messages: HTTP status and reason codes . . . . .	638
152. Core z Systems resources - Group: operations summary . . . . .	610	190. User object: base managed object properties specializations . . . . .	646
153. Core z Systems resources - Groups: URI variables . . . . .	611	191. User object: class specific additional properties. . . . .	647
154. Core z Systems resources - CPC: operations summary . . . . .	611	192. List Users: HTTP status and reason codes	652
155. Core z Systems resources - CPC: URI variables . . . . .	612	193. Update User Properties: HTTP status and reason codes . . . . .	656
156. Core z Systems resources - Logical partitions: operations summary . . . . .	612	194. Add User Role to User: HTTP status and reason codes . . . . .	658
157. Core z Systems resources - Logical partitions: URI variables . . . . .	613	195. Remove User Role from User: HTTP status and reason codes . . . . .	660
158. Core z Systems resources - Reset activation profile: operations summary . . . . .	613	196. Create User: HTTP status and reason codes	663
159. Core z Systems resources - Image activation profile: operations summary . . . . .	613	197. Delete User: HTTP status and reason codes	665
160. Core z Systems resources - Load activation profile: operations summary . . . . .	613	198. User Role object: base managed object properties specializations . . . . .	668
161. Core z Systems resources - Group profile: operations summary . . . . .	614	199. User Role object: class specific additional properties. . . . .	668
162. Core z Systems resources - Activation profile: URI variables . . . . .	614	200. permission-info object properties . . . . .	670
163. Core z Systems resources - Capacity record: operations summary . . . . .	614	201. Update User Role Properties: HTTP status and reason codes . . . . .	676
164. Core z Systems resources - Capacity record: URI variables . . . . .	614	202. Add Permission to User Role: HTTP status and reason codes . . . . .	679
165. ec-mcl-description object . . . . .	615	203. Remove Permission from User Role: HTTP status and reason codes . . . . .	682
166. action object . . . . .	615	204. Create User Role: HTTP status and reason codes . . . . .	684
167. ec object . . . . .	615	205. Delete User Role: HTTP status and reason codes . . . . .	685
168. mcl object. . . . .	615	206. Task object: properties . . . . .	687
169. stp-config object . . . . .	616	207. User Pattern object: properties . . . . .	692
170. stp-node object . . . . .	616	208. Update User Pattern Properties: HTTP status and reason codes . . . . .	698
171. psw-description object . . . . .	616	209. Create User Pattern: HTTP status and reason codes . . . . .	700
172. zaware-network object . . . . .	616	210. Password Rule object: properties . . . . .	704
173. ip-info object. . . . .	617	211. character-rule object properties. . . . .	705
174. network-ip-info object. . . . .	617	212. custom-character-set object properties	706
175. absolute-capping object . . . . .	617	213. Update Password Rule Properties: HTTP status and reason codes . . . . .	711
176. Console object: base managed object properties specializations . . . . .	618	214. Create Password Rule: HTTP status and reason codes . . . . .	713
177. Console object: class specific additional properties. . . . .	618	215. LDAP Server Definition object: properties	718
178. network-info object properties . . . . .	619	216. Create LDAP Server Definition: HTTP status and reason codes . . . . .	727
179. detailed-network-info properties . . . . .	619	217. Group object: base managed object properties specializations . . . . .	731
180. detailed-network-info properties . . . . .	619	218. Group object: class specific additional properties. . . . .	731
181. ipv4-info properties . . . . .	620	219. match-info object properties. . . . .	731
		220. CPC object: base managed object properties specializations . . . . .	744

221. CPC object: class specific additional properties . . . . .	744	249. zCPC processors metric group . . . . .	901
222. ipv6-info object properties . . . . .	749	250. Blade CPU and memory metric group . . . . .	902
223. hardware-message object properties . . . . .	750	251. Crypto metric group . . . . .	902
224. CPC object: energy management related additional properties . . . . .	750	252. Flash memory adapters metric group . . . . .	903
225. Get CPC Audit Log: HTTP status and reason codes . . . . .	783	253. RoCE adapters metric group . . . . .	903
226. Get CPC Security Log: HTTP status and reason codes . . . . .	786	254. Ensemble power metric group . . . . .	904
227. List CPC Hardware Messages: HTTP status and reason codes . . . . .	788	255. Virtual server CPU and memory metric group . . . . .	904
228. Get CPC Hardware Message Properties: HTTP status and reason codes . . . . .	790	256. Virtualization host CPU and memory metric group . . . . .	906
229. Delete CPC Hardware Message: HTTP status and reason codes . . . . .	792	257. Workload metrics group - service class data metric group . . . . .	907
230. Logical Partition object: base managed object properties specializations . . . . .	793	258. Virtualization host (vSwitch) uplink metric group . . . . .	909
231. Logical Partition object: class specific additional properties . . . . .	794	259. Virtualization host (vSwitch) by virtual network metric group . . . . .	911
232. Reset activation profile: properties . . . . .	833	260. Attached virtual server network adapters metric group . . . . .	914
233. Image activation profile: properties . . . . .	839	261. Optimizer IEDN virtual network interface metric group . . . . .	916
234. Load activation profile: properties . . . . .	862	262. Optimizer IEDN physical network adapter metric group . . . . .	917
235. Group profile: properties . . . . .	869	263. Top-of-rack switch port metrics group . . . . .	919
236. Capacity records: type-specific properties . . . . .	875	264. Optimizer IEDN physical network adapter metric group . . . . .	921
237. caprec-proc-info object . . . . .	876	265. Performance policy XML elements . . . . .	923
238. caprec-target object . . . . .	876	266. Performance policy XML: Elements in a <b>ServiceClass</b> element . . . . .	924
239. Inventory service: operations summary . . . . .	881	267. Performance policy XML: Elements required for a <b>Velocity</b> element . . . . .	924
240. Metrics service: operations summary . . . . .	881	268. Performance policy XML: Elements in a <b>Filter</b> element . . . . .	926
241. Metrics service: URI variables . . . . .	881	269. Enum values for a class of managed objects . . . . .	931
242. BladeCenter temperature and power metric group . . . . .	897	270. Enum values for the <b>name</b> property of User Role objects with a <b>type</b> of " <b>system-defined</b> " . . . . .	933
243. Blade power metric group . . . . .	898	271. Enum values for the <b>name</b> property of Task objects . . . . .	935
244. Channels metric group . . . . .	898		
245. CPC overview metric group . . . . .	899		
246. zBX node overview metric group . . . . .	900		
247. Logical partitions metric group . . . . .	900		
248. zCPC environmentals and power metric group . . . . .	901		

---

## Safety

---

### Safety notices

Safety notices may be printed throughout this guide. **DANGER** notices warn you of conditions or procedures that can result in death or severe personal injury. **CAUTION** notices warn you of conditions or procedures that can cause personal injury that is neither lethal nor extremely hazardous. **Attention** notices warn you of conditions or procedures that can cause damage to machines, equipment, or programs.

### World trade safety information

Several countries require the safety information contained in product publications to be presented in their translation. If this requirement applies to your country, a safety information booklet is included in the publications package shipped with the product. The booklet contains the translated safety information with references to the US English source. Before using a US English publication to install, operate, or service this IBM® product, you must first become familiar with the related safety information in the *Systems Safety Notices*, G229-9054. You should also refer to the booklet any time you do not clearly understand any safety information in the US English publications.

---

### Laser safety information

All IBM z Systems™ (z Systems™) models can use I/O cards such as FICON®, Open Systems Adapter (OSA), InterSystem Channel-3 (ISC-3), or other I/O features which are fiber optic based and utilize lasers (short wavelength or long wavelength lasers).

### Laser compliance

All lasers are certified in the US to conform to the requirements of DHHS 21 CFR Subchapter J for Class 1 or Class 1M laser products. Outside the US, they are certified to be in compliance with IEC 60825 as a Class 1 or Class 1M laser product. Consult the label on each part for laser certification numbers and approval information.

**CAUTION: Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. (C027)**

**CAUTION: This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)**



---

## About this publication

This publication defines, for reference purposes, the external interface of the Hardware Management Console (HMC) Web Services Application Programming Interface (Web Services API) for IBM z Systems, (HMC version 2.13.0). This document specifies the capabilities, input and output formats, and behaviors of the Web Services API as viewed by an application external to the HMC that is leveraging that interface.

---

## Related publications

The following publications provide information which supplements the information found within this document:

- *z Systems Capacity On Demand User's Guide*, SC28-6943
- *z Systems Ensemble Workload Resource Group Management Guide*, GC27-2629
- *z Systems Ensemble Planning Guide*, GC27-2631
- *z Systems Processor Resource/Systems Manager Planning Guide*, SB10-7162
- *z/VM CP Planning and Administration Guide*, SC24-6178
- *z/VM CP Commands and Utility Reference*, SC24-6175
- *z13™ Installation Manual for Physical Planning*, GC28-6938
- *zEC12 Installation Manual for Physical Planning*, GC28-6914-01
- *zBC12 Installation Manual for Physical Planning*, GC28-6923-00
- *z BladeCenter Extension Model 004 Installation Manual for Physical Planning*, GC27-2630
- *zBX Model 003 Installation Manual for Physical Planning*, GC27-2619-01
- *zBX Model 002 Installation Manual for Physical Planning*, GC27-2611-03

---

## Related HMC and SE console information

- Hardware Management Console (HMC) and Support Element (SE) information can be found on the console help system, or on the IBM Knowledge Center at <http://www.ibm.com/support/knowledgecenter/> (Select **z Systems** on the navigation bar, and then select your server).

---

## Extending zBX connectivity options to Layer-2

Customer experience with zBX has led IBM to depart from its original requirement to exclusively support Layer-3 connectivity between the external data network and the intraensemble data network (IEDN) top-of-rack (TOR) switches in the zBX. A Redpaper™ is available, illustrating a set of pre-tested configuration examples in support of both Layer-2 and Layer-3 connectivity. The Redpaper, *IBM zEnterprise BladeCenter Extension: Network Connectivity Options* (REDP-5036), includes a description of limitations and trade-offs when deploying Layer-2 versus Layer-3 connectivity.

The Redpaper can be accessed at the following website: <http://www.redbooks.ibm.com/redpieces/abstracts/redp5036.html?Open>.

---

## Revision bars

A technical change to the text is indicated by a vertical line to the left of the change.

- New publications do not usually contain any revision bars. However, for the convenience of application programmers that have used the previous version of this book, *Hardware Management Console Web Services API Version 2.12.1*, SC27-2626-00a, revision bars have been added to illustrate the changes made since that document was published.

For more information about what has changed since the last publication, see “Summary of API version updates” on page 5.

---

## Accessibility

IBM strives to provide products with usable access for everyone, regardless of age or ability.

Accessible publications for this product are offered in HTML format and can be downloaded from Resource Link<sup>®</sup> at <http://www.ibm.com/servers/resourcelink>.

If you experience any difficulty with the accessibility of any z Systems information, go to Resource Link at <http://www.ibm.com/servers/resourcelink> and click **Feedback** from the navigation bar on the left. In the **Comments** input area, state your question or comment, the publication title and number, choose **General comment** as the category and click **Submit**. You can also send an email to [reslink@us.ibm.com](mailto:reslink@us.ibm.com) providing the same information.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

## Accessibility features

The following list includes the major accessibility features in z Systems documentation:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Customizable display attributes such as color, contrast, and font size
- Communication of information independent of color
- Interfaces commonly used by screen magnifiers
- Interfaces that are free of flashing lights that could induce seizures due to photo-sensitivity.

## Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

## IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

---

## How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. Send your comments by using Resource Link at <http://www.ibm.com/servers/resourcelink>. Click **Feedback** on the navigation bar on the left. You can also send an email to [reslink@us.ibm.com](mailto:reslink@us.ibm.com). Be sure to include the name of the book, the form number of the book, the version of the book, if applicable, and the specific location of the text you are commenting on (for example, a page number, table number, or a heading).



---

## Chapter 1. Introduction

- | This chapter provides an overview of IBM z Unified Resource Manager (zManager) APIs, how to enable and access them, and considerations for compatibility.

---

### Overview

- The IBM z Unified Resource Manager (zManager) is a collection of advanced hardware and virtualization management functions delivered as z Systems firmware. The functions of zManager are implemented as a cooperating set of components hosted on the Hardware Management Console (HMC), the Support Element (SE), the IBM z BladeCenter Extension (zBX) Model 004, the blades of an IBM zEnterprise® BladeCenter Extension (zBX) Model 003 or Model 002, and as extensions to z/VM®. It provides a uniform, integrated and workload-oriented administrative model for the heterogeneous computing configuration provided by a z Systems environment. The functions provided by zManager include:
- Hardware inventory, initialization, configuration, monitoring and problem analysis for the components of a z Systems CPC, including both the traditional z Systems computing resources as well as the zBX.
  - Firmware installation and update for the HMC, SE, traditional CPC components, zBX infrastructure, Intra-Ensemble Data Network (IEDN) elements, and zBX blades.
  - Operational control and energy management for these hardware elements.
  - Configuration of function-specialized workload accelerators available as blade optimizers in the zBX.
  - Provisioning, configuration, control and monitoring of virtualized computing systems (virtual servers) on the firmware-managed IBM blade and z/VM environments.
  - Secure management of the IEDN through the provisioning and control of IEDN virtual networks.
  - Automatic and workload-oriented performance optimization of the heterogeneous, virtualized environment.

The HMC serves as the administrative access point for zManager. In that capacity, the HMC provides a web-based, remote-able graphical user interface (UI) to make the zManager functions available to users. In addition, it hosts the implementation of zManager Web Service API (Web Services API) that is described in this document.

The Web Services API is a web-oriented programming interface that makes the underlying zManager capabilities available for use by higher level management applications, system automation functions, or custom scripting. The functions that are exposed through the API support several important usage scenarios in virtualization management, including resource inventory, provisioning, monitoring, automation and workload-based optimization among others.

### Components of the API

The Web Services API consists of two major components. Both components are accessed by client applications by establishing TCP/IP network connections with the HMC.

#### Web services interface

The web services interface is a request-and-response oriented programming interface by which client applications obtain information about the system resources managed by zManager, and by which those applications can perform provisioning, configuration or control actions on those resources.

As is the case for any web-oriented interface, client applications interact with this interface by means of the Hypertext Transfer Protocol (HTTP), an application protocol that flows over TCP/IP socket connections. Client applications request operations by forming and sending text-oriented request messages as defined by HTTP, and the Web Services API responds with text-oriented HTTP response messages. The use of HTTP makes the API client-programming-language neutral, and thus accessible to a

wide variety of client applications. Client applications can be developed in programming languages such as Java™, or in scripting languages such as Perl or Python that include extensive support for performing HTTP operations.

The design of the API's mapping to HTTP has been influenced by the Representational State Transfer (REST) style of interface design. The manageable resources of the system are associated with and identified by durable URIs, and the basic get, update, create and delete operations on those manageable resources are mapped directly to the HTTP GET, PUT, POST and DELETE methods. Request and response data is provided using JavaScript Object Notation (JSON), a simple, open and portable transfer representation. Mapping the functions of the API to HTTP in this way simplifies client application development and allows access to the API without the need for extensive client side tooling or libraries as is often the case in other approaches to web services interface design.

Broadly speaking, the web service interface provides two categories of operations:

- Resource (or object) oriented operations, in which a particular request is targeted at a single manageable resource instance and typically affects just that single resource instance. The majority of the API has this orientation, for example providing functions for interacting with the virtual servers, virtualization hosts, virtual networks and workloads of the system.
- Service oriented operations, in which a particular request operates across many or all manageable resources of the system. The service-oriented operations are provided to support usage scenarios that cannot be accomplished efficiently using an object-by-object sequence of individual requests. The operations provided by the Metrics and Inventory services of the API are examples of service-oriented operations.

### **Asynchronous notification facility**

The web services interface described above is useful to satisfy many usage scenarios, particularly those in which the client application's interest in and interaction with zManager is focused on performing a short-term task. In these kinds of applications (typical of automation or simple provisioning), the client application forms a request, gets a response, processes the response and then “forgets” about the zManager resource it interacted with. That is, the application does not attempt to retain (or cache) information about zManager resources long term and then keep that cache up to date.

However, more sophisticated management applications, including those for discovery, monitoring and advanced provisioning, are not single-request-and-forget with respect to their interest in zManager. Rather, such applications have a need to obtain and retain (i.e., cache) information about the inventory, configuration and status of many zManager resources, and to keep that cached information up to date.

In order to support these more sophisticated applications, the Web Services API provides an asynchronous notification facility by which zManager can inform interested client applications about changes to the resources managed by zManager.

The API's asynchronous notification facility is designed around the Java Message Service (JMS), an open, standard framework and API for sending messages between two or more applications.

---

## **Enabling and accessing the API**

The Web Services API is provided on an HMC that is running with firmware version 2.11.1 or later. The API can be used to query, configure and control Central Processing Complexes (CPCs) containing a Support Element (SE) that is running with firmware version 2.11.1 or higher.

By default, the Web Services API is disabled on the HMC. When disabled, the HMC internal firewall is configured to prohibit connections to any of the TCP/IP ports used by the API. When in this state, requests to connect to the API network ports are completely ignored by the HMC without a connection-refused response.

The Web Services API can be enabled and the scope of access to it configured using the **Customize API Settings** task in the HMC UI. This task, which previously provided configuration settings for the SNMP APIs, has been extended with new controls for the Web Services API as well.

The **Customize API Settings** task allows an installation to enable the API via an overall enabled/disabled setting. When enabled, the HMC internal firewall is reconfigured to allow access to the relevant network ports. When the API is enabled with default settings, the HMC allows connections to the API functions from client applications accessing the HMC from any TCP/IP address. For additional security, an installation can configure the HMC to permit connections to the API ports only from selected network addresses or subnets. These addresses or subnets are specified by the **Customize API Settings** task as well. If specified, these connection restrictions are enforced by the HMC internal firewall.

In addition to the overall enablement on/off control and the optional client network address filtering, access to the API is further secured by the requirement for per-user authorization.

The HMC **User Profiles** and **Manage Users Wizard** tasks define access and other characteristics of an HMC user. These tasks have been extended to provide a new property (**Allow access to management interfaces**) to indicate whether a particular HMC user is to be permitted to use the API or not. By default, this setting is disabled for an HMC user profile and thus attempts to establish an API session by that user are rejected. The installation can use the **Customize API Settings**, **User Profiles**, or **Manage Users Wizard** tasks of the HMC to set this property for one or more HMC users and thus allow those users to access the API.

Once a user is permitted to establish API sessions, its actions within those sessions are subject to the HMC's access control model, as is described in the section that follows.

## Authentication and access control

The HMC provides a built-in access control model in which an HMC user authenticates itself to the HMC to establish its identity, and then based on that identity is permitted or denied the ability to perform certain operations as specified by the access control configuration. These operations, and the objects on which they are permitted, are managed with object and task/action permissions that are grouped into roles that are assigned to HMC users. Roles are managed with the **Customize User Controls** task on the HMC. Roles are assigned to users with the **User Profiles** or **Manage Users Wizard** tasks.

Use of the Web Services API is subject to the same access control policy as is used for UI operations.

Establishing an API session with the HMC requires the initiating application to provide a valid HMC logon ID and corresponding password in order to authenticate and establish the identity under which its requests will be performed. (See “Logon” on page 45 for more information.) The API requires the use of SSL connections so that these login credentials can be flowed securely. The user credentials are validated by the HMC in the same way they are validated for a logon to the UI, either via the HMC's built-in user registry or by use of an LDAP directory server.

Once a client application has established an API session, its ability to access various managed object instances and the operations that can be performed on those instances is regulated based on the identity associated with the API session and the access control policy configured in the HMC for those managed object instances. Access control requirements vary based on the class of managed object and the operation for the managed object. These access control requirements for API actions mirror the requirements for corresponding tasks in the HMC user interface. Details on the authorization requirements for an operation are specified in the description of that operation.

## Alternate HMC considerations

- I A z Systems ensemble is managed by a pair of HMCs that operate in a primary/alternate configuration rather than by a single HMC. At any point in time, one HMC of the pair is designated as the primary HMC and has active management responsibility for the resources of the ensemble. The other HMC is

designated as the alternate HMC, and when in this role acts only as a standby for the primary HMC and mirrors configuration updates from the primary in case the primary fails, but does not otherwise perform active management.

The Web Services API can and should be enabled on both the primary and alternate HMCs of the pair.

However, client applications should generally connect only with the primary HMC for API purposes. Because the alternate HMC is not performing active management, it is unable to perform the management actions implied by API requests, and thus most API operations are designed to be rejected if directed to the alternate HMC (with HTTP status code 404, reason code 3). The API operations supported on the alternate HMC are limited to those that control the alternate HMC function itself. The description of an operation specifically indicates that it is supported on the alternate HMC if this is the case.

---

## Compatibility

The capabilities of the Web Services API will evolve as additional management functionality is added to zManager. Over time, this evolution could result in a mixture of HMC and client application versions coexisting in a customer environment. The principles and guidelines outlined in this section are intended to maximize the compatibility and interoperability among HMC and client applications in such a mixed environment.

### API versioning

Since the functionality of the Web Services API may evolve over time, each functional level of the API is identified by a version number. This version number is represented in major.minor form, with the initial version of the API designated as version 1.1.

The API version offered by an HMC can be determined before API logon by using the Query API Version operation (GET /api/version). The version number of the API is also provided in the response from the Logon operation.

Enhancements to the API specification that maintain compatibility with previous versions (see principles below) are indicated by incrementing the minor portion of the version number. So, for example, the first set of compatible changes to the API would be designated as version 1.2, following the initial 1.1 version.

Because the minor versions within a major version stream (e.g. the 1.x versions) are considered compatible, the HMC always offers and behaves according to the latest minor version of the API specification it supports. That means, for example, the API does not offer any facility by which a client can request version 1.1 behaviors on an HMC that offers version 1.2 level of functionality.

While reasonable effort will be made to preserve compatibility, it may become necessary to make changes to zManager (and thus the API) that do not maintain compatibility with the previous version. If this occurs, the introduction of this new (incompatible) behavior is indicated by incrementing the major part of the version number, and starting the minor part of the version number again at 1. The first such version would thus be identified as version 2.1.

### Allowable changes within a major version

The following kinds of changes to the API specification are allowable within a major version, and thus result in changes to the minor but not major parts of the API version number.

- Adding new object classes or new operations on existing object classes.
- Adding new properties to the data model of an existing object class.
- Changing existing properties of an existing object class from read-only or mutable to writable using the API.
- Adding new URIs and operations related to those URIs.

- Adding new optional query parameters to existing URIs where the behavior in the absence of this query parameter is unchanged.
- Adding new optional fields into input bodies where the behavior in the absence of these new fields is unchanged.
- Adding new fields to the response bodies of existing operations.
- Adding additional header or body fields to existing notification messages.
- Adding data for new classes of objects to the results provided by the Inventory service.
- Generating new types of notification messages.
- Generating property change notifications for new properties, or for existing properties that did not provide those notifications previously.
- Adding new enumeration values to enumeration-type fields returned in response bodies without removing or changing the meaning of any existing enumeration values.
- Adding new error status and reason codes.
- Adding new metric groups.
- Adding new metric fields to the end of existing metric groups.

## Requirements on client applications

In order for a client application to correctly interoperate with an HMC that may be offering a higher minor version of the API, client applications must be designed and developed following the simple principle of “ignore what you don’t understand” when interpreting responses or messages received from the HMC. This is necessary because the principles of allowable changes specified in the preceding section allow new fields to be added to preexisting responses or messages.

More specifically, a client application must:

- Ignore, without error, any field in a response body that is not recognized by the application.
- Ignore, without error, any header or body field in a notification message that is not recognized by the application.
- Ignore, without error, any notification message of an unrecognized type that may be received by the application.
- Ignore, without error, any object appearing in the response to an Inventory request that are of a class not recognized by the application.
- Tolerate receiving a value in a field with an enumeration data type that is an unexpected value. If the application is attempting to display this field, it might consider mapping the unrecognized enumeration value to some value indicating “other” or “unknown”.
- Ignore, without error, extra values provided in a row of metric group data that reside beyond the last field currently expected by the application.

These conditions can arise as a result of API extensions that are considered allowable within a given major version. Following the “ignore what you don’t understand” principle prepares a client application to tolerate these API additions should they occur.

---

## Summary of API version updates

The following functions were introduced in the respective API version:

*Table 1. Summary of updates by API version number*

API Ver	Description	HMC MCL	SE MCL
HMC/SE Version 2.11.1			

Table 1. Summary of updates by API version number (continued)

API Ver	Description	HMC MCL	SE MCL
1.1	Added reason codes 0, 105 and 108 as possible HTTP status code 400 (Conflict) error conditions reported by the <b>Migrate Virtual Server</b> operation.	N48180.278	N48168.275
1.1	Changed the <b>backing-virtualization-host-storage-resource</b> property of the Virtual Server data model (in the fullpack-virtual-disk nested object) from a read-only property to a writeable property.	N48180.287	None
1.1	Increased the maximum request body size for the <b>Import Storage Access List</b> operation to 1 MB, and increased the maximum request body size for most other operations to 64KB.	N48180.288	N48168.294
1.1	Added the <b>cores-per-processor</b> property to the Blade object data model (read-only).	N48180.296	None
1.1	Added inventory and property-change notification support for the Virtualization Host Storage Resource object.	N48180.308	N48168.315
1.1	Added the <b>inventory-error-details</b> field and related inventory-error-info nested object to the inventory-error document returned by the <b>Get Inventory</b> operation when error condition 5 is encountered.	N48180.314	N48168.321
1.1	<ul style="list-style-type: none"> <li>Added the <b>properties=common</b> query parameter to the <b>Get Virtual Server Properties</b> operation.</li> <li>Added the virtual-server-common category and power-vm-virtual-server-common, prsm-virtual-server-common, x-hyp-virtual-server-common and zvm-virtual-server-common classes to the <b>Get Inventory</b> operation.</li> </ul>	N48180.319	N48168.325
1.1	Changed the <b>resources</b> field in the request body for the <b>Get Inventory</b> operation from required to optional.	N48180.340	None
1.1	Corrected the range checking for the <b>load-address</b> field of the <b>Load Logical Partition</b> operation so that the operation correctly supports loading from an alternate subchannel.	N48180.360	None
1.2	<ul style="list-style-type: none"> <li>Added the <b>Mount Virtual Media Image</b> operation.</li> <li>Increased API version number from 1.1 to 1.2.</li> </ul>	N48180.361	None
1.2	Corrected the format of the URI returned by the <b>List Members of Virtual Network</b> operation for a zBX TOR port to reflect the correct canonical URI format for zBX TOR port elements.	N48180.363	None
1.2	Added the <b>power-vm-partition-id</b> property to the Virtual Server object data model for PowerVM® virtual servers (read-only)	N48180.363	N48168.378
1.2	Added HTTP status code 409 (Conflict) as a possible error condition for the List Virtualization Host Storage Resources and Get Virtualization Host Storage Resource Properties operations.	N48180.376	None
1.2	Added the <b>feature-list</b> property to the Virtualization Host object data model. This property is provided for virtualization hosts on all CPCs supported by the Web Services API, but the particular features provided by a given virtualization host will differ based on the release and MCL level of the CPC.	N48180.380	N48168.402
<b>HMC/SE Version 2.12.0</b>			

Table 1. Summary of updates by API version number (continued)

API Ver	Description	HMC MCL	SE MCL
1.3	<p>The following extensions are provided by the HMC Web Services API for HMCs at version 2.12.0, and apply to all CPCs supported by the Web Services API:</p> <ul style="list-style-type: none"> <li>• Increased API version number from 1.2 to 1.3.</li> <li>• Added the <b>power-saving-state</b> property to the for BladeCenter and Blade objects data models, and added the <b>cpc-power-saving-state</b> and <b>zpc-power-saving-state</b> properties to the CPC object data model.</li> <li>• Added <b>"not-supported"</b> as a possible enumeration value for the power-save-allowed and power-cap-allowed properties of BladeCenter and Blade objects, and added <b>"not-supported"</b> as a possible enumeration value of the <b>cpc-power-save-allowed</b>, <b>cpc-power-cap-allowed</b>, <b>zpc-power-save-allowed</b> and <b>zpc-power-cap-allowed</b> properties of the CPC object.</li> <li>• Added the <b>status</b> (read-only), <b>acceptable-status</b> (writeable), <b>perf-status</b> (read-only) and <b>compliant-perf-status</b> (writable) properties to the Workload Resource Group object data model.</li> <li>• Added <b>most-severe-perf-status</b> and <b>perf-status-data-points</b> fields, and related <b>perf-status-data-point</b> nested object to the response from the <b>Generate Workload Resource Groups Report</b> operation.</li> <li>• Added the <b>perf-policies</b> property to the Virtual Server object data model, and also added related virtual server performance policy nested object.</li> <li>• Added <b>"data-retrieval-error"</b> as a possible enumeration value for the status-detailed field in the response for the <b>Generate Load Balancing Report</b> operation.</li> <li>• Added an optional request body containing an optional <b>force</b> input field to the <b>Unmount Virtual Media</b> operation.</li> <li>• Changed the <b>Create Virtual Server</b> operation for a zVM virtual server to require the <b>password</b> field on input rather than allowing it to be optional. This change has been made to improve security.</li> <li>• Added HTTP status code 409 (Conflict) as a possible error response reported by the following Storage Management operations: <ul style="list-style-type: none"> <li>– <b>Import Storage Access List</b></li> <li>– <b>Create Virtualization Host Storage Resource</b></li> <li>– <b>Delete Virtualization Host Storage Resource</b></li> <li>– <b>Add Virtualization Host Storage Resource Paths</b></li> <li>– <b>Remove Virtualization Host Storage Resource Paths</b></li> <li>– <b>Discover Virtualization Host Storage Resources.</b></li> </ul> </li> </ul>	H09182.023	H09173.028

Table 1. Summary of updates by API version number (continued)

API Ver	Description	HMC MCL	SE MCL
1.3	<p>The following extensions are provided by the HMC Web Services API for HMCs at version 2.12.0, but apply only to CPCs with SE version 2.12.0:</p> <ul style="list-style-type: none"> <li>• Added <b>cp-cpu-consumption-percent</b>, <b>ifl-cpu-consumption-percent</b> and <b>other-cpu-consumption-percent</b> fields to the response from the <b>Generate Hypervisor Report</b> operation for zVM virtualization hosts. These new fields are provided for zVM virtualization hosts running at version 6.2 or greater.</li> <li>• Added the following in support of IBM zAware partitions. These changes apply only for partitions of the new "<b>zaware</b>" type: <ul style="list-style-type: none"> <li>– Added "<b>zaware</b>" as a possible value of the <b>activation-mode</b> property of Logical Partition and Image Activation Profile objects.</li> <li>– Added the <b>zaware-network</b>, <b>network-ip-info</b> and <b>ip-info</b> nested objects as common nested object definitions used for new properties of Logical Partition objects.</li> <li>– Added the <b>zaware-host-name</b>, <b>zaware-master-userid</b>, <b>zaware-master-pw</b>, <b>zaware-network-info</b>, <b>zaware-gateway-info</b> and <b>zaware-dns-info</b> properties to the Logical Partition and Image Activation Profile object data models.</li> <li>– Added HTTP status 400 reason code 306 as a possible error response from the Load Logical Partition, PSW Restart, Start Logical Partition, Stop Logical Partition, and Update Image Activation Profile Properties operations when these operations are attempted on an IBM zAware partition.</li> </ul> </li> <li>• Added the Cryptos and Flash Memory Adapters metric groups for CPC objects. Data entries are provided for a CPC in these metric groups if the CPC has one or more Cryptos or Flash Memory Adapters installed.</li> <li>• Added new <b>cp-cpu-time</b>, <b>ifl-cpu-time</b>, <b>zaap-cpu-time</b>, <b>ziip-cpu-time</b> and <b>icf-cpu-time</b> metrics to the Virtualization Host CPU and Memory metric group (for zVM virtualization hosts).</li> </ul>	H09182.023	H09173.028
1.3	Added HTTP status code 409 (Conflict) as a possible error condition for the List Virtualization Host Storage Resources and Get Virtualization Host Storage Resource Properties operations.	H09182.062	None
1.3	Added property change support for the <b>unique-device-id</b> property of the Storage Resource object.	H09182.102	None
1.3	Added the <b>feature-list</b> property to the Virtualization Host object data model. This property is provided for virtualization hosts on all CPCs supported by the Web Services API, but the particular features provided by a given virtualization host will differ based on the release and MCL level of the CPC.	H09182.119	H09173.149
<b>HMC/SE Version 2.12.1</b>			



Table 1. Summary of updates by API version number (continued)

API Ver	Description	HMC MCL	SE MCL
1.4	<p>The following extensions are provided by the HMC Web Services API for HMCs at version 2.12.1 or later, and apply to all CPCs supported by the Web Services API:</p> <ul style="list-style-type: none"> <li>• Increased API version number from 1.3 to 1.4.</li> <li>• Added the <b>Get Network Adapter Properties</b> operation for Virtual Server objects.</li> <li>• Added the <b>List Virtualization Host Storage Resources of a Storage Resource</b> operation for Storage Resource objects.</li> <li>• Added the <b>List Virtual Disks of a Virtualization Host Storage Resource</b> operation for Virtualization Host objects.</li> <li>• Added API support for absolute capping to the Logical Partition and Image Activation Profile objects, including the following API extensions: <ul style="list-style-type: none"> <li>– Added the <b>absolute-processing-capping</b>, <b>absolute-aap-capping</b>, <b>absolute-ifl-capping</b>, <b>absolute-ziip-capping</b> and <b>absolute-cf-capping</b> properties to the data model for Logical Partition objects.</li> <li>– Added the <b>absolute-general-purpose-capping</b>, <b>absolute-aap-capping</b>, <b>absolute-ifl-capping</b>, <b>absolute-ziip-capping</b> and <b>absolute-icf-capping</b> properties to the data model for Image Activation Profile elements of CPC objects.</li> </ul> </li> <li>• Added the <b>partition-identifier</b> property to the data model for Logical Partition objects.</li> </ul> <p>The following extensions are provided by the HMC Web Services API for HMCs at version 2.12.1 or later, but primarily apply only to CPCs with SE version 2.12.1 or later:</p>	H49574.020	H49564.021

Table 1. Summary of updates by API version number (continued)

API Ver	Description	HMC MCL	SE MCL
1.4	<ul style="list-style-type: none"> <li>• Added support for management of processor performance for virtual servers on x-hyp virtualization hosts, comprising the following API extensions:               <ul style="list-style-type: none"> <li>– Added the <b>cpu-perf-mgmt-enabled-x-hyp</b> property to the data model for Ensemble objects.</li> <li>– Added the <b>cpu-shares-supported</b> property to the data model for Virtualization Host objects.</li> <li>– Extended the existing <b>cpu-perf-mgmt-enabled</b>, <b>initial-shares</b>, <b>minimum-shares</b> and <b>maximum-shares</b> properties of Virtual Server objects to now also be applicable to virtual servers of type x-hyp.</li> <li>– Added the <b>workload-processor-mgmt-status</b>, <b>workload-processor-mgmt-status-reason</b>, <b>initial-shares</b>, <b>shares</b>, <b>min-shares</b>, and <b>max-shares</b> fields to the response for the <b>Generate Hypervisor Report</b> operation of Ensemble objects. when issued for x-hyp virtualization hosts. These fields were previously provided for PowerVM and/or z/VM virtualization hosts but not for x-hyp virtualization hosts</li> </ul> </li> <li>• Added support for ensemble availability management, comprising the following API extensions:               <ul style="list-style-type: none"> <li>– Added the <b>workload-element-groups</b>, <b>active-avail-policy</b>, <b>default-avail-policy</b>, <b>custom-avail-policies</b>, <b>avail-status</b>, and <b>compliant-avail-status</b> properties and related nested objects to the data model for Workload Resource Group objects.</li> <li>– Added the Availability Policy element of a Workload Resource Group and corresponding operations for elements of this class, including <b>List</b>, <b>Create</b>, <b>Delete</b>, <b>Get Properties</b>, <b>Update Properties</b> and <b>Activate</b> operations for Availability Policy elements of Workload Resource Group objects.</li> <li>– Added the Workload Element Group object and corresponding operations on objects of this class, including <b>List</b>, <b>Create</b>, <b>Delete</b>, <b>Get Properties</b> and <b>Update Properties</b> operations for Workload Element Group objects.</li> <li>– Added <b>Add To</b> and <b>Remove From</b> operations for managing the inclusion of Workload Element Groups within Workload Resource Group objects.</li> <li>– Added <b>List</b>, <b>Add To</b> and <b>Remove From</b> operations for managing the inclusion of Virtual Servers within Workload Element Group objects.</li> <li>– Added reporting operations for availability management, including the <b>Generate Workload Resource Groups Report (Ensemble Availability Management)</b>, <b>Generate Workload Resource Group Availability Status Report</b>, <b>Generate Virtual Server Report (Ensemble Availability Management)</b>, and <b>Generate Availability Status Report</b> operations.</li> <li>– Added an enumeration value of "<b>workload-element-group</b>" as a possible value of the <b>inclusion-type</b> field in the response for the <b>List Virtual Servers of a Workload Resource Group</b> operation, to specify the additional way in which virtual servers can become members of a Workload Resource Group.</li> <li>– Added the <b>avail-status</b>, <b>acceptable-avail-status</b>, <b>avail-policies</b> and <b>workload-element-group</b> properties and related nested objects to the data model for Virtual Server objects.</li> </ul> </li> <li>• Added the <b>heat-load</b>, <b>heat-load-forced-air</b> and <b>heat-load-water</b> metrics to the zCPC Environmentals and Power metric group.</li> <li>• Added RoCE Adapters and Ensemble Power Metric groups.</li> </ul>	H49574.020	H49564.021

Table 1. Summary of updates by API version number (continued)

API Ver	Description	HMC MCL	SE MCL
1.5	<p>Added the ability to specify per-virtual-server shutdown timeouts and to perform deactivate actions that either use or override the shutdown timeout specified in the virtual server or its hosting virtualization host configuration. This new capability applies to CPCs with SE version 2.12.1 or later that have the specified MCL installed. This new capability comprises the following detailed API extensions:</p> <ul style="list-style-type: none"> <li>• Increased API version number from 1.4 to 1.5.</li> <li>• Added the enumeration value "<b>virtual-server-shutdown-timeout-override-support</b>" as a possible feature identifier included in the <b>feature-list</b> property of the Virtualization Host object to indicate the availability of this new capability to virtual servers hosted on a Virtualization Host instance.</li> <li>• Added the <b>shutdown-timeout-source</b> and <b>shutdown-timeout</b> properties to the data model of the Virtual Server object to allow a customized default shutdown timeout to be configured for a particular virtual server.</li> <li>• Added the <b>shutdown-timeout</b> field as an optional request body field for the <b>Deactivate Virtual Server</b> operation to allow an individual deactivation action to be performed using a shutdown timeout that is different than the timeout configured as a default for the virtual server or its hosting virtualization host.</li> </ul>	H49574.075	H49564.070
<b>HMC/SE Version 2.13.0</b>			

Table 1. Summary of updates by API version number (continued)

API Ver	Description	HMC MCL	SE MCL
1.6	<p>The following extensions are provided by the HMC Web Services API for HMCs at version 2.13.0 and apply to all SE versions supported by the Web Services API:</p> <ul style="list-style-type: none"> <li>• Increased API version number from 1.5 to 1.6.</li> <li>• Added API support for managing HMC user and role definitions, comprising the following API extensions: <ul style="list-style-type: none"> <li>– Added the User object representing a defined HMC user, and associated List, Create, Delete, Get Properties, Update Properties, Add User Role, and Remove User Role operations for User objects.</li> <li>– Added the User Role object representing a security role for HMC users, and associated List, Create, Delete, Get Properties, Update Properties, Add Permission and Remove Permission operations for User Role objects.</li> <li>– Added the Task object representing the permission to invoke an HMC UI task or request an associated API operation and associated List and Get Properties operations for Task objects.</li> <li>– Added the User Pattern object representing a pattern string used to match user IDs during logon and associated List, Create, Delete, Get Properties and Update Properties operations for User Pattern objects.</li> <li>– Added the <b>Reorder User Patterns</b> operation for Console objects.</li> <li>– Added the Password Rule object representing a password format and expiration policy specification and associated List, Create, Delete, Get Properties and Update Properties operations for Password Rule objects.</li> <li>– Added the LDAP Server Definition object representing the configuration of an LDAP server used for HMC authentication and associated List, Create, Delete, Get Properties and Update Properties operations for LDAP Server Definition objects.</li> <li>– Added the <b>replication-override-possible</b> properties for Group objects.</li> <li>– Added the enumeration values "user" and "user-role" as possible object class values for the Inventory Service.</li> </ul> </li> <li>• Added API support for determining the characteristics of the virtual server network adapter used for GPMP purposes, comprising the following API extensions: <ul style="list-style-type: none"> <li>– Added the <b>gpmp-network-adapter</b> property to the data model for Virtual Server objects.</li> <li>– Added the <b>adapter-type</b> property to the network-adapter-power and network-adapter-x-hyp nested object for Virtual Server objects.</li> </ul> </li> <li>• Added the <b>gpmp-available-version</b> property to the data model for Virtualization Host objects.</li> <li>• Added the <b>management-enablement-level</b> property to the data model for CPC objects.</li> </ul> <p>The following extensions are provided by the HMC Web Services API for HMCs at version 2.13.0 but primarily apply only to managed system with SE version 2.13.0 or later:</p>	N98841	CPC: N98775 zBX: N98822

Table 1. Summary of updates by API version number (continued)

API Ver	Description	HMC MCL	SE MCL
1.6	<ul style="list-style-type: none"> <li>• Added API support for zBX Model 004 ensemble nodes, comprising the following API extensions:               <ul style="list-style-type: none"> <li>– Added the <b>type</b> property to the data model for a zBX object to indicate whether the zBX is a CPC-attached (zBX Model 002/003) or ensemble node (zBX Model 004) zBX. As related extensions, also added the <b>type</b> property to the response for the <b>List zBXs of Ensemble</b> operation and also added it as an optional filtering query parameter for that operation. Note that when a zBX object represents a zBX node, the value of its <b>parent</b> property contains the URI of the Ensemble of which it is a member rather than the URI of the CPC to which it is attached.</li> <li>– Added an optional new input format for the <b>Add Node to Ensemble</b> operation to allow a zBX Model 004 to be specified as the node to be added to the ensemble.</li> <li>– Added the enumeration value "<b>zbx</b>" as a possible value for the <b>type</b> property of a Node element of an Ensemble object.</li> <li>– Added the <b>max-nodes</b>, <b>max-cpc-nodes</b> and <b>max-zbx-nodes</b> properties to the data model for an Ensemble object.</li> <li>– Added many additional properties to the data model for a zBX object when that object represents a zBX node.</li> <li>– Added the Get EC/MCL Description, Activate, Deactivate, Set Power Save, Set Power Capping, List Virtualization Hosts and List Virtual Servers operations for zBX node objects.</li> <li>– Added the <b>Activate</b> and <b>Deactivate</b> operations for BladeCenter objects contained within a zBX node.</li> <li>– Added the <b>List Virtualization Hosts</b> and <b>List Virtual Servers</b> operations for Ensemble nodes in general.</li> <li>– Added the enumeration value "<b>node</b>" as a possible value of the availability status nested object of a Virtual Server object.</li> <li>– Added the zBX (Node) Overview metric group as a metric group that can be requested using the Metric Service.</li> </ul> </li> <li>• Added API support for manipulation of hardware messages, comprising the following API extensions:               <ul style="list-style-type: none"> <li>– Added the <b>hardware messages</b> container property and hardware message nested objects to the data model for zBX node and CPC objects and to the data model for the HMC Console object.</li> <li>– Added the List Hardware Messages, Get Hardware Message Properties and Delete Hardware Message operations for zBX node and CPC objects and the HMC Console object.</li> </ul> </li> <li>• Added API support for the retrieval of audit and security log information, comprising the following API extensions:               <ul style="list-style-type: none"> <li>– Added a notification topic and Log Entry notification messages to provide for asynchronous push-type delivery of new log entries for the HMC Console object to interested API clients.</li> <li>– Added the <b>Get Notification Topics</b> operation for Session objects to allow API clients a more general way of retrieving the names of notification topics for an API session.</li> <li>– Added the <b>Get Audit Log</b> and <b>Get Security Log</b> operations for zBX node and CPC objects and the HMC Console object.</li> </ul> </li> <li>• Added the <b>smt-usage</b>, <b>thread-0-usage</b> and <b>thread-1-usage</b> metrics in the zCPC Processors metric group.</li> <li>• Added the enumeration value "<b>virtio-scsi</b>" as a possible value for the emulation mode property of a virtual disk element of a Virtual Server object.</li> </ul>	N98841	CPC: N98775 zBX: N98822

**Note:** When an API extension indicates the addition of new properties to the data model for a specified object class, such an extension also includes standard changes to several related operations as well even though, for brevity, these related changes are not specifically mentioned in the table. In general, an extension to an object's data model will also include corresponding changes to the inputs to or responses from **Get Properties**, **Update Properties** and **Create** operations for that object class, as appropriate. In addition, the new properties are included in Inventory Service data for objects of the specified class.

## Chapter 2. Base definitions

This chapter provides basic definitions of data types, representation formats and other fundamental syntactic elements that apply across the Web Services API.

### Data types

The following data types are used in the definition of the management data model, input and output parameters and notification message formats in the Web Services API.

Table 2. Primitive data types

Data type	Description
Boolean	A logical truth value: either the value <b>true</b> or the value <b>false</b> .
Byte	An integer value in the range $-2^7$ to $(2^7)-1$ (the range of a signed 8-bit integer)
Float	An IEEE 754 floating point number in the range $+/-4.9E-324$ to $+/-3.4028235E+38$ . Note that, although IEEE 754 provides for representations of positive or negative Infinity and NaN, such values are not used within the API.
Long	An integer value in the range $-2^{63}$ to $(2^{63})-1$ (the range of a signed 64-bit integer)
Integer	An integer value in the range $-2^{31}$ to $(2^{31})-1$ (the range of a signed 32-bit integer)
Short	An integer value in the range $-2^{15}$ to $(2^{15})-1$ (the range of a signed 16-bit integer)
String	A sequence of Unicode characters. When the number of characters in the string is bounded, the length or length range is provided in parenthesis, for example String (16) for a 16 character string, or String (0-256) for a string that may range in length from 0 (empty) to 256 characters.
String Enum	A String enumeration, i.e. a String for which the possible values are constrained to be one of a specified set of choices.
String/URI	A String that contains a URI path used to designate object instances or operations within the API.
String/IPV4 Address	A String that contains an Internet Protocol Version 4 address presented in dotted-decimal notation.  Example: "127.0.0.1"
String/IPV6 Address	A string that contains an Internet Protocol Version 6 address presented in colon-separated-hexadecimal notation. Leading and consecutive groups of zeros may be omitted in the representation as is conventional for IPV6 addresses presented in this form.  Example: "2001:db8:85a3::8a2e:370:7334"
Timestamp	A non-negative Long integer quantity where the value represents a date and time expressed as the number of milliseconds since midnight on January 1, 1970 UTC, or the special value -1 to indicate special treatment of the timestamp field.

Table 3. Compound data types

Data type	Description
Array of <T>	A ordered sequence of zero or more elements each of data type <T>. An array may be empty, i.e. have no elements contained within it.
Object	A nested data structure providing a set of fields, each field having a name, data type and value. Object types do not formally have names. However, descriptions of these nested objects will often assign reference names to allow connections to be made in the documentation between points of use and definition for a given nested object.

## Input and output representation

Except for a few special cases, the operations provided by the Web Services API expect their input and provide their output using a representation known as JavaScript Object Notation, or JSON for short. The JSON representation is also used within the bodies of notification messages emitted by the API. Unless some different representation is specifically mentioned in the description of an operation or message, all operations and messages should be understood to use JSON notation.

JavaScript Object Notation (JSON) is a lightweight, text-based, language-independent data interchange format that defines a small set of formatting rules for the portable representation of structured data. JSON can represent four primitive types (strings, numbers, booleans, and the value null) and two structured types (objects and arrays) that together provide sufficient expressive power to represent the manageable resource configuration, state, inputs, and outputs that appear in this API.

A JSON string is a sequence of zero or more Unicode characters enclosed in quotes.

A JSON object is an unordered collection of zero or more name/value pairs (sometimes referred to in this document as fields or properties), where a name is a string and a value is a primitive type (string, number, boolean, or null), an array, or a nested object. Each name/value pair is represented in the form **"name": value** and is separated from the next name/value pair by a comma. The collection of name/value pairs comprising the object is enclosed by left and right braces e.g. { ... }.

An array is an ordered sequence of zero or more values separated from each other by commas and enclosed in left and right square brackets e.g. [10,20,30]). The values in the array can be primitive or structured types, i.e. arrays of objects or arrays of arrays are permitted.

The precise BNF syntax of JSON notation is not provided in this document, but can be found in the IETF information document RFC 4726, *The application/json Media Type for JavaScript Object Notation (JSON)*, July 2006. This RFC can be found in text format on the World Wide Web at:

<http://www.ietf.org/rfc/rfc4627.txt>

## Representing API data types in JSON

The following tables define the mapping between the API data types and their corresponding representation in JSON notation.

*Table 4. Primitive data types in JSON notation*

API data type	JSON representation
Boolean	A JSON boolean with keywords <b>true</b> and <b>false</b>
Byte, Integer, Long, Short	A JSON number with a sign and integer component, but no fraction or exponent part.
Float	A JSON number, possibly including fraction or exponent parts.  On output, values with a magnitude greater than or equal to $10^{-3}$ and less than $10^7$ are representation in floating-point format with a fraction part but not exponent part (e.g. 1.7, -32.467). Values with magnitudes outside that range are represented in scientific notation with both fraction and exponent parts (e.g. -4.23E127).
String, String Enum	Represented as a JSON string enclosed in quotes.
Timestamp	An unsigned JSON number with integer component, but no fraction or exponent part.



Table 5. Compound data types in JSON notation

Data type	Description
Array of <T>	A JSON square-bracket-enclosed array with elements represented according to the data type <T>.
Object	A JSON curly-brace-enclosed object, with the fields/properties of the nested object represented as name/value members of the object. The name of a property/field is used directly as the name part of the JSON object member, and the value of the field/property is provided as the value part of the member.

All strings in the JSON representation (object member names, and string values) are encoded in UTF-8.



---

## Chapter 3. Invoking API operations

The Web Services API provides an extensive set of operations that client applications can invoke to obtain information about the manageable resources of the system, to change those resources' characteristics, and to take action on them. Because the API is designed using a web services orientation, these operations are accessed by means of Hypertext Transport Protocol (HTTP) protocol messages flowing across TCP/IP network connections.

Most aspects of HTTP protocol usage required to invoke API operations or receive responses apply universally across all of the operations of the API. Rather than repeat these details in the description of each and every operation, this common information is instead provided in this chapter. The material in this chapter should be considered to apply to each and every operation of the API unless the operation-specific description indicates otherwise. Thus, the information in this chapter should be consulted in conjunction with the operation-specific descriptions elsewhere in this document when determining how to invoke a specific API operation.

---

### HTTP protocol standard

The Web Services API has been designed in accordance with the HTTP version 1.1 protocol, as defined in the W3C internet standards document *RFC 2616, Hypertext Transfer Protocol – HTTP/1.1, June 1999*. This RFC can be found in HTML format on the World Wide Web at: <http://www.w3.org/Protocols/rfc2616/rfc2616.html>

The API requires that all clients interact using the HTTP/1.1 protocol. The API does not support clients that use HTTP/1.0.

**Note:** While the API does not specifically assume or exclude any particular client user agent, its use and interpretation of HTTP elements has been designed presuming that the client application interacting with the API is a programmatic web application client or HTTP-capable scripting client rather than a standard browser-based application.

---

### Connecting to the API HTTP server

When the Web Services API is enabled, the HMC API HTTP server listens for SSL-based socket connections on TCP port 6794. The HMC is enabled for both the SSL version 3 and TLS version 1 protocols on this SSL port. It does not accept non-SSL connections. The set of cipher suites enabled for the HMC API HTTP server is controlled by the **Certificate Management** task on the HMC. Note, the default set of cipher suites may change with updates to the HMC if one or more of the cipher suites are found to be weak or vulnerable.

The listening port for the API HTTP server is a fixed port number and is not subject to customer reconfiguration. Thus, client applications can treat this as a well-known port number rather than requiring customer input when configuring the networking parameters the client will use to connect to the HMC.

---

### HTTP header field usage

HTTP request and response messages include elements known as header fields (often referred to simply as headers for short) that provide request metadata. Certain headers are required or provided in all HTTP messages, while others are present in selected messages depending on content.

This section describes the use of header fields by the Web Services API.

## Required request header fields

The following HTTP request headers are relevant to all request methods (GET, PUT, POST, DELETE) and are required on all API requests (except as indicated for the Logon and Query API Version operations).

HTTP header name	Rqd/Opt	Description
<b>Host</b>	Required	Specifies the Internet host and port number of the HMC to which the request is being directed, as obtained from the original URI given by the client application. The Web Services API enforces that this header is provided as required by the HTTP protocol, but does not check or use the value of the header in any way.
<b>X-API-Session</b>	Required <sup>1</sup>	An opaque string that provides a cryptographically strong identifier of the API session (known as a session id) under which this request is executed. This header is required on all requests that require authentication. The Login operation begins a new HMC session and includes credentials identifying the HMC user for the session. Upon successful authentication, the Login operation returns the value to be used in the <b>X-API-Session</b> header for all subsequent requests of the same session. Failure to include this header on a request requiring authentication results in status code 403 (Forbidden) with reason code 4. Specifying an invalid session id results in status code 403 (Forbidden) with reason code 5.
Note:		
1. Not required on requests to the Query API Version and Logon operations since these operations can be performed before an API session has been established.		

For requests made using the HTTP PUT or POST methods, the following additional request headers are required if a request body is being provided. If an operation being requested via POST method does not require a request body, these headers can be omitted.

HTTP header name	Rqd/Opt	Description
<b>Content-Length</b>	Required if request body present	When used in a request, specifies the length of the request body. If omitted, the request is presumed to not contain a body.  The API limits the size of request bodies in order to control usage of memory resources on the HMC. Unless a different limit is specified for a particular operation, in general the largest request body accepted by the API is 64KB. Requests with bodies that exceed this maximum are rejected with an HTTP status 413 (Request Entity Too Large) response.
<b>Content-Type</b>	Required if request body present	When used in a request, specifies the MIME media type of the request body contained in the request. This header is required if the <b>Content-Length</b> header is supplied and specifies a non-zero request body length, otherwise status code 400 (Bad Request) will result.

## Optional request headers

The following HTTP request headers are relevant to all request methods (GET, PUT, POST, DELETE) and may be specified on these method requests but are not required. If present, they are interpreted by the API in the indicated way.

HTTP header name	Rqd/Opt	Description
<b>Accept</b>	Optional	<p>Specifies the list of response MIME media types that the client application is prepared to accept for the response to the request. This header is provides for content negotiation between the client and the server in cases where the Web Services API supports multiple possible response media types for a given operation.</p> <p>In the current implementation, the Web Services API supports only a single response media type for each operation. For the majority of operations, that media type is JSON (<b>application/json</b>), but selected operations support a different media type (indicated in the descriptions of those special operations).</p> <p>If this header is omitted, the Web Services API responds using the (single) media type supported for the operation. If the header is included, it must allow for the single media type that the operation supports, otherwise the request will fail with HTTP status code 406 (Not Acceptable).</p> <p>If an operation is extended to support multiple media types, compatibility will be maintained for existing clients that request the operation without specifying an <b>Accept</b> header.</p>
<b>X-Audit-Id</b>	Optional	<p>A string that provides additional client identity information that is included in all audit records created for this request, in addition to the API user's HMC login identity. This header is intended to provide improved audit logging in the case of clients that make requests on behalf of multiple upstream users while logged into the API under a single HMC login identity. Such clients should provide the identity of their upstream user in this header so that the requests of different upstream users can be distinguished in the HMC audit logs. The HMC will use up to the first 64 characters of information from this header if present, and silently ignore the remainder of the header's value if it is longer than 64 characters.</p>
<b>X-Client-Correlator</b>	Optional	<p>A string that provides diagnostic information pertaining to this request that is of significance to the client, such as a client request number or the like. The HMC will record this information in selected diagnostic trace or log data it collects so as to allow better cross-correlation of this information with similar information maintained by the client. This data supplied in this header is intended to assist in product problem determination and does not otherwise affect the operation of the API. The HMC will use up to the first 64 characters of information from this header if present, and silently ignore the remainder of the header's value if it is longer than 64 characters.</p>

## Standard response headers

The following HTTP response headers are always provided in the response to all requests.

HTTP header name	Description
<b>Date</b>	<p>The date and time, from the perspective of the HMC's clock, at which the response message was generated. As required by the HTTP protocol specification, this date is an HTTP full date sent in the RFC 1123-defined fixed length format.</p> <p>Example: Sun, 08 Oct 1961 10:08:00 GMT</p>

The following HTTP response headers are provided in the response to all requests except those that result in a 204 (No Content) HTTP status code.

HTTP header name	Description
<b>Content-Length</b>	When used in a response, specifies the length of the response body. If omitted, the response does not contain a body.
<b>Content-Type</b>	When used in a response, specifies the MIME media type of the response body. This response header is provided any time the <b>Content-Type</b> header is provided and specifies a non-zero length.

## Additional response headers

Some operations may return additional response headers beyond those described in “Standard response headers” on page 21. The following table describes these possible additional response headers.

Operations that return these additional headers indicate that they do so in the operation description.

HTTP header name	Description
<b>Location</b>	The URI of the resources that was created by the operation.  This header is provided for operations that complete successfully with an HTTP status code of 201 (Created).
<b>X-API-Session</b>	An opaque string that provides a cryptographically strong identifier of the API session that was created for the client.  This header is provided in the response to a successful <b>Logon</b> operation.
<b>X-Request-Id</b>	A string that provides diagnostic information identifying the request from the perspective of the HMC. This same information is included in the API log entries that are recorded by the HMC for the request. If captured by the client from a response, a client application developer or support technician can use this information to locate the HMC API log entry corresponding to a particular request. The value of this header will be 64 characters or less.  This header is provided in the responses to all requests.

## Media types

The following media types are applicable to the use of the Web Services API, and thus may appear in the values of **Accept** or **Content-Type** header fields.

MIME media type	Description
application/json	JavaScript Object Notation (JSON), as described by RFC 4627. This media type is used by the Web Services API for both request and response representation for the majority of the operations in the API. The JSON text is encoded using the UTF-8 charset.
application/vnd.ibm-z-zmanager-metrics	Custom output format used for providing the results to the Get Metrics operation of the metrics service. The result text is encoded using the UTF-8 charset.
application/xml	Extensible Markup Language, used for the input and output formats for the Export Performance Policy and Import Performance Policy operations of the workload object. The XML text is encoded using the UTF-8 charset.
application/octet-stream	Binary data. This media type is used by the Web Services API for the request representation for the <b>Mount Virtual Media Image</b> operation of the Virtual Server object.

## HTTP status codes

The HMC API provides standard HTTP status codes in the response to requests to indicate the success or failure of the request. Unless stated otherwise in the description of an operation, the following general interpretations of the status code values apply.

HTTP status code	Description/Causes
200 (OK)	The request has succeeded completely. A response body is provided that contains the results of the request.
201 (Created)	The request has succeeded completely and resulted in the creation of a new managed resource/object. The URI for the newly created managed resource is provided in a <b>Location</b> header. (POST methods only)
202 (Accepted)	The request was successfully validated and has been accepted to be carried out asynchronously.
204 (No Content)	The request succeeded completely, and no additional response information is provided.
400 (Bad Request)	The request was missing required input, had errors in the provided input, or included extraneous input. Additional information regarding the error is provided in an error response body that includes a reason code with additional information.
403 (Forbidden)	Multiple error conditions result in this status code: <ul style="list-style-type: none"> <li>• The request requires authentication but no <b>X-API-Session</b> header was provided, or one was provided but the session ID was invalid.</li> <li>• The user under which the API request was authenticated is not authorized to perform the requested operation.</li> <li>• The ensemble is not operating at the management enablement level required to perform this operation.</li> </ul>
404 (Not Found)	Multiple error conditions result in this status code: <ul style="list-style-type: none"> <li>• The URI does not designate an extant resource, or designates a resource for which the API user does not have object-access permission.</li> <li>• The URI designates a resource or operation that is not supported by the HMC because it is currently the alternate HMC.</li> </ul>
405 (Method Not Allowed)	The request specifies an HTTP method that is not valid for the designated URI.
406 (Not Acceptable)	The <b>Accept</b> header for the request does not include at least one content representation supported by the Web Services API.
409 (Conflict)	The managed resource is in an incorrect state (status) for performing the requested operation. Additional information regarding the error is provided in an error response body that includes a reason code with additional information.
413 (Request Entity Too Large)	The request includes a request body that is too large.  Unless a different limit is specified for a particular operation, in general the largest request body accepted by the API is 64 KB.
415 (Unsupported Media Type)	The <b>Content-Type</b> header for the request specifies a representation that is not supported by the Web Services API.
500 (Server Error)	A server error occurred during processing of the request.
501 (Not Implemented)	The request specifies an HTTP method that is not recognized by the server (for any resource). <b>Note:</b> The response body that accompanies this error is not a JSON response body as defined in “Error response bodies” on page 24.
503 (Service Unavailable)	The request could not be carried out by the HMC due to some temporary condition.
505 (HTTP Version Not Supported)	The request specifies an HTTP protocol version that is not supported by the Web Services API.

## Error response bodies

For most 4xx and 5xx HTTP error status codes, additional diagnostic information beyond the HTTP status code is provided in the response body for the request. This information is provided in the form of a JSON object containing the following fields:

Field name	Type	Description
<b>http-status</b>	Integer	HTTP status code for the request.
<b>request-method</b>	String	The HTTP method (DELETE, GET, POST, PUT) that caused this error response.
<b>request-uri</b>	String	The URI that caused this error response.
<b>reason</b>	Integer	Numeric reason code providing more details as to the nature of the error) than is provided by the HTTP status code itself. This reason code is treated as a sub-code of the HTTP status code and thus must be used in conjunction with the HTTP status code to determine the error condition. Standard reason codes that apply across the entire API are described in "Common request validation reason codes." Additional operation-specific reason codes may also be documented in the description of the specific API operations.
<b>message</b>	String	Message describing the error. This message is not currently localized.
<b>stack</b>	String	Internal HMC diagnostic information for the error. This field is supplied only on selected 5xx HTTP status codes.
<b>error-details</b>	Object	A nested object that provides additional operation-specific error information. This field is provided by selected operations, and the format of the nested object is as described by that operation.

### Usage notes:

- The message provided in the **message** field is primarily intended as a convenience for use by developers when developing and testing client applications. Because it is not localized, it may not be appropriate for client applications to simply pass this message on to their clients when reporting errors to those upstream clients. Instead, client applications can use the value in the **reason** field as a key in obtaining a client-provided message that may be more appropriate to use.
- Because the reason code is treated as a sub-code of the HTTP status code, the same reason code value is often defined for multiple different HTTP status codes and has a different meaning in each case. For example, reason code 1 when considered for a 400 (Bad Request) status code has a different meaning than when considered for a 403 (Forbidden) status code. For this reason, client applications that make decisions based on the reason codes should always include checking the HTTP status code as part of the relevant logic (e.g. test for status code == 400 AND reason code == 1, not just reason code == 1 alone).

## Common request validation reason codes

The Web Services API performs request validation on each request it receives to ensure the request is correctly formed and appropriate before it begins processing the request. Many errors of basic request syntax can occur on all or a large number of the operations provided by the API. Validation for these kinds of errors is done in a common way across all of the operations and results in a common (not request-specific) reason code being reported if errors are detected. Other validation operation-specific by nature, and results in operation-specific reason codes when errors are detected.

The following table provides the HTTP status codes and reason codes for common request validation. These status and reason codes may be reported on any of the operations of the API.



HTTP status code	Reason code	Description
400 (Bad Request)	1	The request included an unrecognized or unsupported query parameter.
	2	A required request header is missing or invalid.
	3	A required request body is missing.
	4	A request body was specified when not expected.
	5	A required request body field is missing.
	6	The request body contains an unrecognized field (i.e. one that is not listed as either required or optional in the specification for the request body format for the operation).
	7	The data type of a field in the request body is not as expected, or its value is not in the range permitted.
	8	The value of a field does not provide a unique value for the corresponding data model property as required.
	9	The request body is not a well-formed JSON document.
	10	An unrecognized X-* header field was specified.
	11	The length of the supplied request body does not match the value specified in the <b>Content-Length</b> header.
	13	The maximum number of logged in user sessions for this user ID has been reached; no more are allowed.
	14	Query parameters on the request are malformed or specify a value that is invalid for this operation. Common causes include the inability to successfully decode a parameter element, the presented parameters are not in the expected key=value format, the value is not a valid regular expression, or the value is not a valid enum for the operation.
	15	The request body contains a field whose presence or value is inconsistent with the presence or value of another field in the request body. A prerequisite condition or dependency among request body fields is not met.
	403 (Forbidden)	1
3		The ensemble is not operating at the management enablement level required to perform this operation.
4		The request requires authentication but no X-API-Session-header was specified in the request.
5		An X-API-Session header was provided but the session id specified in that header is not valid.
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.
	2	A URI in the request body does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.
	3	The request URI designates a resource or operation that is not available on the Alternate HMC.
	4	The object designated by the request URI does not support the requested operation.

HTTP status code	Reason code	Description
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.
	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	3	The operation cannot be performed because the object designated by the request URI is currently locked to prevent disruptive changes from being made.
	8	The operation cannot be performed because the request would result in the object being placed into a state that is inconsistent with its data model or other requirements. The request body contains a field whose presence or value is inconsistent with the current state of the object or some aspect of the system, and thus a prerequisite condition or dependency would no longer be met.

## Common request processing reason codes

Certain common error conditions can be encountered during the processing of many of the operations of the API. When they are encountered they are reported using the same HTTP status and reason code by any operation of the API that may encounter them.

These common request processing reason codes are listed in the following table:

HTTP status code	Reason code	Description
500 (Server Error)	19	An Asynchronous operation was terminated while running because the host HMC was restarted, or a failover to the alternate HMC occurred.
	Other in range 0 - 39	An internal processing error has occurred and no additional details are documented.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	2	The request could not be processed because the SE is not currently communicating with an element of a zBX needed to perform the requested operation.
	3	This request would exceed the limit on the number of concurrent API requests allowed.

## Use of chunked response encoding

For most API operations, the size of the response data is modest and therefore standard HTTP response payload transfer encoding is used. In this encoding, the length of the entire payload of the response message is provided in the response before any of the contents of the response payload are written to the socket connection. But some operations, such as the Get Inventory operation of the Inventory service and the Get Metrics operation of the Metrics service, can produce very large responses. Use of standard transfer encoding for these kinds of operations is inefficient for the HMC because it requires the entire response be generated and buffered before any of it is sent in order to compute and send the total length of the response body before sending any of the contents of the response data.

To avoid the need for the buffering the entire response, and to instead allow the response to be transmitted in smaller segments as they are prepared, operations that return large responses use HTTP chunked response encoding instead. Chunked transfer encoding is an HTTP V1.1 data transfer feature that allows the payload of the response message to be split into a sequence of smaller parts known as chunks, with the size of each chunk transmitted as part of the chunk rather than requiring the transmission of the size of the full response payload.

Chunked transfer encoding is defined in the HTTP/1.1 protocol standard, RFC 2616, cited earlier in this section.

The HTTP protocol standard requires that all HTTP/1.1 applications (client or server) be capable of receiving and handling chunked transfer-encoded messages, so the use of this encoding by the API HTTP server is within the options allowed by the protocol standard. However, since this format may be unexpected to naively-written applications, its use is limited by the API HTTP server to the special circumstances that warrant its use to improve performance or efficiency. Therefore, a client application can safely assume that an operation will not use chunked transfer encoding for its responses unless the use of this encoding is specifically mentioned in the description of the operation.

---

## Filter query parameters

Some operations allow for the (optional) use of designated query parameters for conveying additional request parameters. Although query parameters can be used to convey various kinds of additional request information, most operations that make use of query parameters do so for the purpose of filtering the response entries to a subset of what would otherwise be returned. For example, this kind of filtering is typically provided on operations that are described as List operations (e.g. List Virtual Servers of Ensemble). This section describes the interpretation/handling of filter-type query parameters across all of the operations of the API.

As would be expected, if an operation is invoked without specifying any of its possible filter-type query parameters, the operation returns all of the result entries applicable to the request. For example, the List Virtual Servers of Ensemble operation invoked with no filtering query parameters returns all of the Virtual Server objects in the Ensemble to which the API user has access.

If one or more filter-type query parameters are specified, the combination of those parameters specifies a logical match expression that is evaluated against each entry that is a candidate for inclusion in the result to determine if the entry is included or not. Within that expression, there may be multiple occurrences of the same-named query parameter and/or there may be occurrences of differently-named query parameters. The query parameters are interpreted as a logical expression using the following rules:

- Multiple occurrences of the same-named query parameter are interpreted as a group that is connected by a logical OR operation among all of query parameters with the same name. An entry remains a candidate for inclusion in the result as long as it matches at least one of the values specified for this particular query parameter.
- Occurrences of differently-named query parameters are first organized into OR'ed groups as mentioned above, and then these groups are interpreted as being connected by logical AND operations. Thus an entry is included in the result only if it matches at least one value from each of the differently-named groups of parameters.
- As an example, a query string of “name=fee&type=fie&name=foe&type=fum” is interpreted as specifying the expression (name=fee OR name=foe) AND (type=fie OR type=fum). Note that the order in which the query parameters appear in the string is not important.

As a filter-type query parameter is applied against a candidate entry, it is determined to match or not as follows:

- If the query parameter is of data type String, the parameter's value is interpreted as a regular expression pattern and is considered to match if the corresponding String property of the candidate entry matches the pattern.
- If the query parameter is of data type String Enum, the parameter's value is compared against the corresponding Enum property of the candidate entry and is considered to match if they are exactly the same value.

## Regular expression syntax

The values of String-type filtering query parameters are interpreted as regular expressions. The regular expression syntax used is the same as that used by the Java programming language, as specified for the `java.util.regex.Pattern` class. Documentation on that syntax can be found at on the following web page: <http://download.oracle.com/javase/1.4.2/docs/api/java/util/regex/Pattern.html>

---

## Chapter 4. Asynchronous notification

The Web Services API includes an asynchronous notification facility by which client applications may subscribe to and receive notification messages regarding a set of predefined management events. These events include:

- Addition and removal of managed objects to/from the inventory of resources that are managed by the HMC.
- Changes to specified properties of managed object instances.
- Changes to the operational status of managed objects.
- Completion of asynchronously processed jobs.
- | • Addition of entries to the HMC's audit log.
- | • Addition of entries to the HMC's security log.

The zManager notification facility is based on the Java Message Service (JMS) architecture and API for exchanging messages among two or more applications.

---

### JMS basics

In the JMS model, message-based communication between producing and consuming applications is coordinated by an intermediate component known as a message broker that acts as the clearinghouse for message exchange. The message broker provides a registry of the available destinations to which messages can be posted, and a store for messages that have been posted but not yet consumed.

Applications acting in the role of message producer create messages and post them (via the broker) to the destination appropriate for the type of message. The messages are associated with the destination and retained by the broker until they have been consumed by interested applications.

Applications acting in the role of message consumer connect to a message destination (again, via the broker) in order to express interest in receiving messages posted to it. As messages are posted to the destination by producers, the broker makes the messages available to interested consumers which then receive and process the message.

In the Web Services API notification facility, the HMC acts both as the message broker and the message producer for API notification messages. API client applications act as message consumers.

For the broker function, the HMC includes an integrated JMS message broker implementation based on Apache ActiveMQ, an open, standards-based implementation of JMS. This integrated broker is configured to allow internal HMC function to act as message producers, and to allow external applications to connect as message consumers. However, external applications cannot produce and send messages using the HMC integrated broker.

---

### Connecting to the API message broker

As part of the Web Services API, the HMC provides an integrated JMS message broker based on Apache ActiveMQ Version 5.2.0. This message broker is active on the HMC whenever the Web Services API is enabled.

When active, the integrated broker listens for client connections using the following transports supported by ActiveMQ:

- OpenWire flowing over SSL connections, listening port: 61617.

- STOMP (Streaming Text Oriented Messaging Protocol) flowing over SSL connections, listening port: 61612.

The broker is enabled for the SSL version 3 and TLS version 1 protocols on these SSL ports.

The listening ports for the message broker are fixed port numbers and are not subject to customer reconfiguration. Thus, client applications can treat them as well-known port numbers rather than requiring customer input when configuring the networking parameters the client will use to connect to the HMC.

In order to connect to the integrated message broker, clients must provide a valid HMC user name and password in order to identify the HMC user making the connection. This information is validated using the standard HMC user authentication mechanisms before allowing the connection to succeed. The integrated message broker does not allow any anonymous or unauthenticated connections.

---

## Per-session notification topics

As part of its access control enforcement, the Web Services API limits the distribution of notification messages to clients that have object-access permission to the managed object for which the notification was generated.

In order to accomplish this, the API does not define a single (or fixed number of) notification topics to which all messages are posted and from which any or all clients can receive messages. Rather, the API uses per-session notification topics.

- | In this approach, each API session is associated with a set of JMS destinations that are created by the  
| HMC when the session is created or other actions are performed using the session, and are used for  
| providing notifications destined to that session. The names of the object notification and job completion  
| destinations are returned as part of the results from the Login operation. The names of all destinations  
| available to a session are returned by the **Get Notification Topics** operation. Each session has the  
| following per-session notification destinations:
- An object notification topic, used by the HMC to send notifications that pertain to the inventory and status of managed objects that this session has permission to access.
  - A job notification topic, used by the HMC to send notifications that pertain to the status of asynchronous operations that are initiated by this session.
  - An audit notification topic, used by the HMC to send notifications that pertain to the HMC's audit log, but only if the client has permission to the **Audit and Log Management** task. Without the required task permission, there is no such destination available to the session.
  - A security notification topic, used by the HMC to send notifications that pertain to the HMC's security log, but only if the client has permission to the **View Security Logs** task. Without the required task permission, there is no such destination available to the session.

| The session is also associated with an HMC user (identified during API session login) that in turn has a  
| set of object access permissions defined for it that determine the managed resources that it is authorized  
| to access. The HMC user also has a set of task permissions defined for it that determine the tasks that it  
| is authorized to perform.

| As notification messages are created for managed resource changes or other events, they are distributed  
| to sessions according to the object access permissions of those sessions. More specifically, when a  
| notification is generated pertaining to some managed resource, it is published to the object notification  
| topics of all sessions for which the related API user has object-access permission to that managed  
| resource, and is omitted from the object notification topics of sessions for which the user does not have  
| object-access permission. As a result, a particular API session will have access to all notifications for all  
| managed resources to which its API user has access permission, but will not have access to notifications  
| for managed resources that it does not.

Notification messages for job completion are sent only to the job notification topic of the API session that initiated the asynchronous processing represented by the job.

## Notification message formats

- I Five types of notification messages are provided by the API. The JMS messages created for all types of notifications share a common set of message characteristics and header fields, which are extended with additional header fields and message body formats that vary by the type of notification.

## Common message characteristics

The characteristics described in this section apply to all notification messages published by the Web Services API.

### Message format

The following JMS message characteristics apply to all notification messages sent by the Web Services API. These characteristics are echoed in the message themselves in the values of the standard JMS message header fields.

Characteristic	Description
Destination	The per-session object or job notification topic as indicated below for each type of notification.
Message type	Text message. The contents of the body vary by the type of notification.
Delivery mode	Nonpersistent.
Expiration	No expiration.
Priority	4 (highest normal priority)
Message ID	A unique message ID for the message.
Correlation ID	Not set by the API.
Timestamp	The time, from the HMC's perspective, that this message was sent.

In addition to the standard JMS message headers, the following additional message properties are provided in all notification messages to allow for message selection:

Message Property Name	Description
<b>notification-type</b>	The type of notification contained in this message, varies by notification type.
<b>session-sequence-nr</b>	The sequence number of this notification within the session. This number starts at 0 when the API session is created, and is incremented each time a notification message is published to this session.
<b>global-sequence-nr</b>	The sequence number of this notification from the HMC. This number starts at 0 when the HMC is started, and is incremented each time a notification message is published to any API session.
<b>object-uri</b>	The current value of the <b>object-uri</b> property (canonical URI) of the managed object for which the notification is being emitted.
<b>object-id</b>	The current value of the <b>object-id</b> property (durable unique identifier) of the managed object.
<b>element-uri</b>	The current value of the <b>element-uri</b> property of the element object for which the notification is being emitted. This message property is included only when the message pertains to an element object of a managed object.
<b>element-id</b>	The current value of the <b>element-id</b> property (local identifier) of the element object. This message property is included only when the message pertains to an element object of a managed object.
<b>class</b>	The current value of the <b>class</b> property (kind of object) of the object, i.e. the kind of object for which the notification is being emitted.

Message Property Name	Description
<b>name</b>	The current value of the <b>name</b> property (display name) for the object to which the notification pertains. <b>Note:</b> Note: In some circumstances the <b>name</b> property may be unavailable, in which case this field is set to an empty string. This may occur, for example, if a property change occurs and is to be reported on very shortly before (essentially concurrent with) the removal of that object from the inventory.

When a notification message pertains to an element object, the message includes **element-uri** and **element-id** fields in addition to **object-uri** and **object-id** fields. The element-\* fields identify the element object instance while the object-\* fields identify the containing managed object instance. In this case, the **class** field provides the class of the element object, and the **name** field provides the name of the element object.

When a notification message pertains to a managed object, the message contains **object-uri** and **object-id** fields but not the element-\* fields and the **class** field provides the class of the managed object and the **name** field provides the name of the managed object.

### Grouping of notifications

A particular managed resource may often experience a series of changes that occur in rapid succession. This might occur, for example, when a user uses an object's Details task in the HMC UI to make a set of changes to the object's properties and then selects the Finish button to complete the task. All of the pending property changes collected by the task will be made on the managed object in very quick succession in response to the Finish button being selected, rather than before as the user was interacting with the task.

In order to reduce the notification traffic in these situations, the notification messages have been designed to allow the HMC to report multiple changes of the same type (e.g. property changes, status changes) for the same managed resource in a single message rather than generating a distinct message for each change. In order to do such grouping, the HMC may delay generation of a notification message for a change for a brief period of time in order to allow coalescing of that change report with others that occur for the same managed resource within the coalescing time window. This optimization will be performed while maintaining the following characteristics:

- Grouping of notifications may be done for property change, status change, and log entry notifications, but will not be done for other notification types.
- Notifications will be buffered for a maximum of 1 second.
- The grouping of change reports will not obscure a client's ability to correctly observe the temporal ordering of the individual changes made to a particular object or between objects based on the messages sent to a session. That is, notification messages will always be generated so that the ordering of the messages as determined by their session sequence numbers reflects the temporal order in which the changes were actually made. All of the changes reported to a session in a message with a lower session sequence number will have occurred before any of the changes reported in a message with a higher session sequence number. Further, the ordering of change reports within a particular message reflects the order in which they occurred to that object as well.
- Coalescing of multiple changes into a single notification message will occur for at most a single pending notification message (thus, of a single type) at a time. If a need arises to report a change of a different type than is currently pending (for example, a need to report a status change when there is currently a set of pending property change reports), coalescing will end for that pending message and it will be posted to notification topics as appropriate. This is necessary in order to maintain the API client's ability to correctly observe temporal ordering.
- The grouping of change reports into notification messages occurs independently for each session, so that one session may receive a different distribution of change reports across notification messages than another session.



The degree to which message grouping occurs or not depends on the timing of changes and possibly other considerations and thus is to be strictly considered an optimization and not guaranteed behavior. Client applications should infer no particular semantic significance to change reports being delivered in a single message vs. a sequence of messages.

## Status change notification

A Status Change notification is emitted by the API to report changes to the **status** property of a managed object.

Characteristic	Description
Destination	The per-session object notification topic for each API session that is authorized to receive the notification.

In addition to the common JMS message headers described above, the following additional message properties are provided for this type of notification:

Message property name	Description
<b>notification-type</b>	Contains the value " <b>status-change</b> ".

The body of a Status Change notification message is a JSON representation of an object that contains the following fields and values:

Field name	Type	Description
<b>change-reports</b>	Array of objects	An array of nested change-report objects, the format of which is described in the next table. The order in which these objects appear in this array reflects the temporal order in which the changes occurred.

Each nested change-report object has the following fields and values:

Field name	Type	Description
<b>old-status</b>	String	The old (previous) value of the <b>status</b> property for the object. The value of this field will be one of the possible enumeration values for the <b>status</b> property as defined for this class of object.
<b>old-additional-status</b>	String	The old (previous) value of the <b>additional-status</b> property for the object. The value of this field will be one of the possible enumeration values for the <b>additional-status</b> property as defined for this class of object.
<b>new-status</b>	String	The new (current) value of the <b>status</b> property for the object. The value of this field will be one of the possible enumeration values for the <b>status</b> property as defined for this class of object.
<b>new-additional-status</b>	String	The new (current) value of the <b>additional-status</b> property for the object. The value of this field will be one of the possible enumeration values for the <b>additional-status</b> property as defined for this class of object.
<b>has-unacceptable-status</b>	Boolean	The value of the <b>has-unacceptable-status</b> property of the object, based on its new status. If true, the object is now considered to have unacceptable status because its current status is not one of the configured acceptable status values for this object.

## Property change notification

A Property Change notification is emitted by the API to report changes to the properties of a managed object where the data model description indicates that modification notification support (qualifier "pc") is available for that property.

Characteristic	Description
Destination	The per-session object notification topic for each API session that is authorized to receive the notification.

In addition to the common JMS message headers described above, the following additional message properties are provided to allow for message selection:

Message property name	Description
notification-type	Contains the value " <b>property-change</b> ".
property-names	Blank-separated list of the names of the properties for which change reports are provided in the body of this notification message. The list always includes a leading and trailing blank so that a property name can be specified as a blank-delimited word in a message selector to avoid matching unintended properties that have the desired property name as a substring.

The body of a Property Change notification message is a JSON representation of an object that contains the following fields and values:

Field name	Type	Description
change-reports	Array of objects	An array of nested change-report objects, the format of which is described in the next table. The order in which these objects appear in this array reflects the temporal order in which the changes occurred.

Each nested change-report object has the following fields and values:

Field name	Type	Description
property-name	String	The name of the property (as specified in the object's data model) that has changed.
old-value	Based on model	<p>If the property is not a container-type or write-only property, this field contains the old (previous) value of the property for the object. The value of this field will be of the data type indicated for this property in the object's data model.</p> <p>If the property is a container-type property (i.e. marked with the (c) qualifier), this field does not provide the complete previous value. Rather, it provides an array of entries that have been removed from the value of the container property. The value of these entries will be of the data type indicated for the property in the object's data model. If no entries have been removed, null is provided.</p> <p>If the property is a write-only property (i.e. marked with the (wo) qualifier), this field does not provide the value of the property. Rather, this field always contains null.</p>
new-value	Based on model	<p>If the property is not a container-type or write-only property, this field contains the new (current) value of the property for the object. The value of this field will be of the data type indicated for this property in the object's data model.</p> <p>If the property is a container-type property (i.e. marked with the (c) qualifier), this field does not provide the complete new value. Rather, it provides an array of entries that have been added to the value of the container property. The value of these entries will be of the data type indicated for the property in the object's data model. If no entries have been added, null is provided.</p> <p>If the property is a write-only property (i.e. marked with the (wo) qualifier), this field does not provide the value of the property. Rather, this field always contains null.</p>

## Inventory change notification

An Inventory Change notification is emitted by the API to report the addition or removal of a managed object to/from the current inventory of resources that are being managed by zManager. This occurs when managed resources are created or deleted, but also may occur in other situations, such as when zManager reestablishes its inventory of (already-existing) managed resources upon restart of the primary HMC.

For some kinds of managed objects, an Inventory Change notification is also emitted by the API to report the addition or removal of an element of a managed object. Such notifications do not occur for all elements, but rather only when specifically described in the documentation for a class of managed object.

Because an Inventory Change notification may be generated more than once for the same conceptual object, these notifications cannot be interpreted as designating a resource creation action.

Characteristic	Description
Destination	The per-session object notification topic for each API session that is authorized to receive the notification.

In addition to the common JMS message headers described above, the following additional message properties are provided to allow for message selection:

Message property name	Description
<b>notification-type</b>	Contains the value <b>"inventory-change"</b> .
<b>name</b>	Not provided for this notification. Always an empty string.
<b>action</b>	The value <b>"add"</b> when the object has been added to the inventory, or <b>"remove"</b> when it is being removed.

The body of an inventory change notification is null.

## Job completion notification

A Job Completion notification is emitted by the API to report that an operation that runs asynchronously to the client application has completed its processing.

Asynchronous operations are those that complete with an HTTP status code of 202 (Accepted) when requested by the client. A Job Completion Notification message is sent to the API session that initiated the job when such an operation completes, and provides to the client application the URI of the job that has completed so the client application can use the Query Job Status operation to obtain results for the job.

Characteristic	Description
Destination	The per-session job notification topic for each API session that is authorized to receive the notification.

In addition to the common JMS message headers described above, the following additional message properties are provided to allow for message selection:

Message property name	Description
<b>notification-type</b>	The value <b>"job-completion"</b> .
<b>job-uri</b>	The URI identifying the asynchronous job that has just completed execution.

The body of a job completion notification is null.

## Log entry notification

A Log Entry notification is emitted by the API to report the addition of a log entry to its corresponding console log.

Characteristic	Description
Destination	The audit notification topic or security notification topic for each API session that is authorized to receive the notification.

In addition to the common JMS message headers described above, the following additional message properties are provided to allow for message selection:

Message property name	Description
notification-type	Contains the value "log-entry".

The body of a Log Entry notification message is a JSON representation of an object that contains the following fields and values:

Field name	Type	Description
log-entries	Array of objects	An array of nested log-entry-info objects, the format of which is described in Table 186 on page 631. The order in which these objects appear in this array reflects the temporal order in which the log entries were created.

---

## Chapter 5. Data model definitions

This chapter covers data model concepts and shared data model schema elements.

---

### Data model concepts

zManager provides resource management and control functions for the various resources that comprise a z Systems ensemble environment. In performing these functions, zManager establishes a separation between those aspects of resource management that are handled entirely by system firmware, and the other aspects for which customer or installation visibility, configuration and control is appropriate.

In order to specify the external aspects in a succinct way, the web Services API is described in this document in terms of a conceptual data model that it offers for the resources that it manages. This data model is an information structuring technique that conceptually defines the kinds of resources that are managed by zManager and for each, the information that is available for and the operations that can be performed on resources of that kind. This data model is intended to provide the complete perspective that clients of the API can have regarding the logical resources of the system while insulating them from implementation details.

### Objects in the data model

The manageable resources of the environment are represented in the management system as entities referred to as objects. Each distinct manageable resource is represented by a separate object instance, and the life cycle of an instance corresponds with the lifecycle of the manageable resource it represents. For example, for physical entities, such as an IBM blade in a zBX, the object that represents it are created implicitly when the physical entity is attached to and configured to be (or entitled to be) part of the system. This object continues to exist so long as the IBM blade remains entitled on the system. For some logical entities, such as the virtualization management functions on an IBM blade, the object that represents it (virtualization host) is also implicitly rather than explicitly created. For other logical entities, such as a virtual server on an IBM blade, the creation of the management model object instance for it is, in fact, the mechanism that triggers the creation of the corresponding manageable resource in the system.

There are different kinds of manageable physical or logical resources in the system, and each kind manifests different observable characteristics. As a result, there are different classes of objects in the data model. Objects of the same class represent the same kind of resource and provide a defined set of *properties* that capture the attributes of that kind of resource that the Web Services API exposes.

### Managed objects and element objects

The object classes defined in the data model fall into one of two main categories: managed object classes (or simply managed objects), and element object classes (element objects).

These two categories are very similar in that they are both, ultimately, unordered collections of named properties that capture the key attributes of a resource instance. The categories differ primarily in how prominently they are handled in the API: the way that instances of them are designated to perform operations on them, and the degree to which API facilities such as inventory and change notification can be offered for objects in that category.

Managed objects are the first-class entities in the data model and the API. They represent the primary manageable resources of the system, such as ensembles, blades, virtual servers and workloads. These kinds of objects typically appear prominently in the main displays of the HMC user interface.

Instances of managed objects are registered and indexed in the zManager managed object registry, and thus can be directly referenced by URIs that form a relatively "flat" namespace. The URI of a managed

object designates its object instance based on its class and a unique, durable, UUID-based identifier called an object ID. For example, the URI of a virtual server is of the form `/api/virtual-servers/{vs-object-id}` where the identifier at the end of the URI is globally unique. Inventory change, property change, and status change notifications can be generated for managed objects.

In comparison, element objects represent the secondary or more-detailed aspects of the system. Examples include the virtual network adapters of a virtual server, or the performance policies of a workload. These kinds of entities do not generally appear in the main displays of the HMC user interface, but rather are displayed only within particular management tasks offered by the UI.

Instances of element objects are not directly registered in the zManager object registry, but rather are associated with or “attached to” some containing or related managed object instance. As a result, access to these elements is indirect, through the containing managed object. The URIs that designate element objects are hierarchical in nature, with the leftmost part of the URI identifying the managed object to which the element is attached. For example, the URI for a virtual disk of a virtual server is of the form `/api/virtual-servers/{vs-object-id}/virtual-disks/{disk-id}` in which the `{disk-id}` at the end is only necessarily unique within the context of the related virtual server. Inventory, property and status change notifications are not provided for element objects. Instead, when changes to elements are reported, those changes are done via property change notifications emitted for the associated managed object.

## Properties in the data model

Object in the management system contain, fundamentally, unordered collections of name/value pairs called properties that capture the key characteristics of the manageable resources they represent. The defined set of named properties that are maintained for a particular kind of resource constitutes the specification of the data model class for that kind of resource.

As a result, in the chapters that follow, the description of the management interfaces for a class of resource begins with a Data Model section that specifies the properties that are exposed by the API for that kind of resource.

Each property has a name, a data type, and a semantic description in prose.

The property name is the programmatic identifier of the property. This identifier is used within requests and responses to indicate that a field represents a particular property of the data model. It is the “name” part of the name/value pair that is the property.

The property data type indicates the kind of information that can be represented by the property, just as a variable’s data type indicates the kind of information that can be stored in a variable. The data type provides information on the nature of the “value” part of the name/value pair that is the property.

### Property characteristics

Properties are classified as being either writeable or read-only from the perspective of an API client application.

Writable properties are ones that can have their values read by `Get <class> Properties` or similar operations and can also have their values directly changed by `Update <class> Properties` operations.

- | Properties that are classified as write-only can have their values directly changed by `Update <class> Properties` operations, but cannot have their values read by using `Get <class> Properties` or similar operations.
- | Properties that are classified as read-only can have their values read by using `Get <class> Properties` or similar operations, but cannot have their values changed directly.

Although properties that are classified as read-only cannot have their values changed directly, their values may nonetheless be affected by other operations supported by a class of object. For example, a class of object might include an **is-enabled** property that is classified as read-only because the enabled state of the resource cannot be affected by a simple `Update <class> Properties` operation on that

**is-enabled** property. However, this object might also define a Change State or Enable operation that can perform this enabling, and as a side effect will alter the value of the **is-enabled** property.

In addition to the read-only vs. writeable classification, properties defined for managed object also can differ in whether changes to them result in property change or status change notifications being emitted for the managed or not. For properties that have property or status change support, these notifications are emitted asynchronously by the API any time the value of the property changes, whether that change was made via this API, the HMC UI, or implicitly by the system. Changes to the values of properties for which change notification support is not provided do not result in such notifications.

In the tables of properties that appear within this document, the characteristics of properties are indicated by qualifier annotations in parenthesis following the property name. The qualifiers have the following meanings:

Qualifier notation	Description
(w)	The property is a writeable property. Any property that lacks this qualifier and the (wo) qualifier is considered read-only and thus is not directly modifiable.
(wo)	The property is a write-only property. Any property that lacks this qualifier and the (w) qualifier is considered read-only and thus is not directly modifiable.
(ro)	Although this property is writable when present in other managed object classes, it is read only in this class. This qualifier is only used when a managed object class specializes the definition of a base managed object property and overrides the writeable characteristic of the base definition.
(pc)	Change to this property's value will result in Property Change notifications.
(c)	The property is a container-type property for which deltas (changes) are reported in Property Change notifications rather than complete old and new values.
(sc)	Change to this property will result in Status Change notifications.
(mg)	This property represents a performance or utilization metric of the object that is included in a metric group available via the Metrics Service of this API. The value of this property may change very frequently and, therefore, property change notifications are not emitted for changes to this property. Client applications interested in obtaining metric information frequently should obtain this information through use of the Metrics Service of this API.

---

## Shared data model schema elements

The data-model schema fragments in this section define groups of properties that are used in common ways in specifying the data models for the managed object classes defined in the API.

The description of the data model for a specific object class specifies the shared schema elements it is incorporating within the Data Model section of that description, if any. It will also include a description of the specializations that apply to that class's use of the shared schema, such as additional constraints on properties, class-specific values for properties, etc.

## Base managed object properties schema

This data-model fragment contains the basic properties that are present in the representation of many of the managed object types that represent manageable resources.

Name	Qualifier	Type	Description
<b>object-uri</b>	—	String (1-255)	The canonical URI path that designates this managed object instance and serves as the primary reference and retrieval key for this instance. The URI path is formed based on a unique and permanently-assigned object ID (see the <b>object-id</b> property in the next row of this table), and as result, an object's URI path will not change as a result of changes to properties of the object. Further, this canonical URI path is independent of the containment hierarchy and thus will not change if this object instance is moved within the hierarchy.
<b>object-id</b>	—	String (36)	The object identifier for the managed object instance. This value is unique in space and time, and is permanently associated with this instance while it is managed by this HMC or its associated alternate HMC. (If the instance is removed from this HMC and later managed by another HMC, it will have a new and different object identifier when managed by that other HMC.) It is generated by zManager and assigned when the managed resource is created or first discovered, and is immutable thereafter. As example, a managed object's object ID will not change as a result of changes to display name, changes in the location of this resource in the containment hierarchy, or across restarts of the HMC.
<b>parent</b>	—	String (1-255)	The canonical URI path of the managed object that is conceptually the parent of this object in the containment hierarchy. This property is null for objects that do not have a parent.
<b>class</b>	—	String (1-32)	The class of resource represented by this managed object. Each distinct class of resource has a different type name, while all instances of the same type share the same type name. The specific value used for a class of object is specified in the Data Model section for that object type. Example: " <b>virtual-server</b> ".
<b>name</b>	(w)(pc)	String (1-64)	The current display name of the managed object as defaulted or specified for the object. This is the simple name of this object, not qualified by containment hierarchy. This name must consist only of alphanumeric characters and the following special characters: period (.), hyphen (-), at sign (@), underscore (_), and space. It must not begin or end with a space. Some resource types do not support the setting of a user-assigned display name. For such objects, this property is not settable, and instead always provides a name assigned by the HMC or SE.
<b>description</b>	(w)(pc)	String (0-1024)	Arbitrary text providing additional descriptive information about this managed resource. This information is retained for the resource and may be shown as part of the object's details on the user interface, but is otherwise not generally used by zManager. This property may be null.
<b>is-locked</b>	(pc)	Boolean	The object is locked and thus disruptive actions or tasks cannot be performed on it.

## Operational status properties

Many (but not all) classes of managed objects support the concept of operational status. That is they maintain information about the current functional state (Not Communicating, Not Operating, etc.) of the managed resource and whether that current functional state is considered acceptable (not alert causing) or not. If a class of object supports the operational status concept, it provides the standard properties defined in the following table (referred to as the operational status properties) in addition to those defined earlier in this section.

Unless stated to the contrary, any object class data model that includes the base managed object properties schema should be understood to also provide these operational status properties as well. For object classes for which that is not the case, the data model description will specifically point out that operational status and thus these operational-status-related properties are not provided for that object.



The operational status properties are as follows:

Name	Qualifier	Type	Description
<b>status</b>	(sc)	String Enum	The current operational status of the managed resource. The possible status values vary by managed object class and are specified in the description of each managed object class that provides this property.
<b>additional-status</b>	(sc)	String Enum	A qualifier to the <b>status</b> property used by selected object classes to provide finer grained operational status tracking.
<b>acceptable-status</b>	(w)(pc)	Array of String Enum	The set of operational status values that the managed resource can be in and be considered to be in an acceptable (not alert causing) state.
<b>has-unacceptable-status</b>	(sc)	Boolean	If true, the current operational status of the managed resource is not one of the acceptable statuses for this resource.



## Chapter 6. General API services

This chapter describes the services that are provided by the Web Services API for creating and deleting API sessions and performing other general functions.

### General API services operations summary

Table 6. General API services: operations summary

Operation name	HTTP method and URI path
"Query API Version" on page 44	GET /api/version
"Logon" on page 45	POST /api/sessions
"Logoff" on page 48	DELETE /api/sessions/this-session
"Get Notification Topics" on page 49	GET /api/sessions/operations/get-notification-topics
"Query Job Status" on page 52	GET /api/jobs/{job-id}
"Delete Completed Job Status" on page 54	DELETE /api/jobs/{job-id}

Table 7. General API services: URI variables

Variable	Description
{job-id}	The identifier of an asynchronous job associated with this user, as returned in the response of the operation that initiated the job.

### Session management services

Almost all operations of the Web Services API are requested and carried out in the context of an API session that is used for determining the client's authority to access managed resources and perform requested operations, and is also used to scope the delivery of asynchronous notifications. An API session is an HMC concept that is independent of and layers on top of network-related considerations such as a TCP/IP socket connection. As a result, a single API session may span multiple TCP/IP socket connect/disconnect sequences from the same client.

Sessions are created upon request from a client by using the **Logon** operation, and may be explicitly terminated by a client using the **Logoff** operation. Sessions may also be terminated by the HMC due to inactivity when no requests are made using the session over a period of 6 hours. (This session timeout is not configurable.) However, termination of a session due to inactivity will not occur as long as a client application uses the API's notification facility to maintain a JMS subscription to one or more of the session's JMS notification topics. The existence of such a subscription is considered by the HMC to indicate that a client is still using the session and thus it is not terminated even if no requests are made using it.

The scope of a session is the particular HMC instance on which it was created via a **Logon** operation. That is, an API session created on one HMC of a primary/alternate HMC pair is not also available on the alternate HMC of that pair. Nor will that session (and its associated **session-id**) be available on the other

HMC should a primary/alternate role switch occur due to a failure of the primary HMC. After a primary/alternate role switch, client applications will need to establish new sessions by making Logon requests again.

Sessions are identified by clients using a **session-id**, which is a string of up to 64 characters in length that is returned to the client in the results from a successful **Logon** operation. This string is generated in a cryptographically-secure manner. A **session-id** string is a form of authentication credentials for a user equivalent in power to a user's user ID and password. Because of this, a **session-id** should be transmitted only within SSL connections.

In order to indicate that subsequent requests are to be performed in the context of a designated session, the client supplies the appropriate **session-id** to the HMC in each such subsequent request. This is done by supplying the **session-id** as the value of the **X-API-Session** HTTP header which is an application-specific header defined by and recognized by the HMC.

The **Logon** and **Query API Version** operations are the only two operations in the Web Services API that can be performed without an API session so requests for these operations do not need to provide the **X-API-Session** HTTP header. All other operations are valid only in the context of an API session and thus requests for all other operations must supply an **X-API-Session** header with a valid **session-id** in order to be successfully executed.

## Query API Version

The **Query API Version** operation returns information about the level of Web Services API supported by the HMC.

### HTTP method and URI

GET /api/version

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
api-major-version	Integer	The major-version part of the API version in effect for this session
api-minor-version	Integer	The minor-version part of the API version in effect for this session
hmc-version	String (5-8)	The version number of the HMC firmware. This is a string of the form <i>v.r.m</i> , where each of <i>v</i> , <i>r</i> and <i>m</i> can be one or two digits. Example: "2.11.1".
hmc-name	String (1-16)	The name assigned to the HMC

### Description

This operation returns name and version information for the HMC and the HMC API.

This operation can be requested without an API session being open, i.e. no **X-API-Session** header, and **session-id** is required on input.

This operation can be invoked on the alternate HMC.

For more information about the function included in each API version, see "Summary of API version updates" on page 5.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 44.

Under normal conditions, no error response codes are returned by this request. (HTTP Status code 500 could possibly result if internal HMC errors occur.)

## Example HTTP interaction

Request:

---

```
GET /api/version HTTP/1.1
```

---

Response:

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 20 Jul 2011 17:22:23 GMT
content-type: application/json;charset=UTF-8
content-length: 119
{
  "api-major-version": 1,
  "api-minor-version": 1,
  "hmc-name": "HMCR32PRI",
  "hmc-version": "2.11.1"
}
```

---

## Logon

The **Logon** operation establishes an API session with the Web Services API.

### HTTP method and URI

**POST** /api/sessions

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
userid	String	Required	The name of the HMC user to be associated with the new API session. This name may be of arbitrary length, i.e. the HMC does not have a defined maximum length.
password	String	Required	The password used to authenticate the HMC user identified by the <b>userid</b> field. The required length and valid characters are determined by the password policy in effect for the user ID.
new-password	String	Optional	A new password to be established for the user defined by the <b>userid</b> field. The required length and valid characters are determined by the password policy in effect for the user ID.

The largest request body accepted by this operation is 512 bytes. Requests with bodies that exceed this maximum are rejected with an HTTP status 413 (Request Entity Too Large) response.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
api-session	String (1-64)	The <b>session-id</b> of the newly created session. The client must specify this value in the <b>X-API-Session</b> header of all subsequent requests that are to be performed in the session.
notification-topic	String (1-128)	The name of the JMS topic the HMC will use to send object-related notification messages to this session.
job-notification-topic	String (1-128)	The name of the JMS topic the HMC will use to send job-related notification messages to this session.
api-major-version	Integer	The major-version part of the API version in effect for this session
api-minor-version	Integer	The minor-version part of the API version in effect for this session
password-expires	Integer	The time interval, in days, until the user's current password expires. A value of 0 indicates that the password will expire within the next 24 hours. A value of -1 indicates that the HMC does not enforce password expiration for this user, however, if this user is authenticated with an external authentication mechanism (e.g. LDAP) such expiration might be enforced by that mechanism.

## Description

This operation opens a new API session with the Web Services API. Authentication is performed as part of this process.

The characteristics and permissions of an HMC user are specified in an HMC User Profile or User Template. The user name provided in the **userid** field of the request body is used to select a corresponding User Profile/Template based on the name. If such a User Profile/Template is found, the client's authority to operate as this HMC user is authenticated by validating the password provided in the **password** field using the authentication method specified in the User Profile/Template. If the password authentication is successful, a new API session is created and the session-id for the new session is provided in the **api-session** field in the response from this operation. This same value is also provided by an **X-API-Session** HTTP header field in the response.

If the request specifies an **X-API-Session** HTTP header field on input (indicating that this operation be performed under some designed session), the logon request fails and status code 400 (Bad Request) is returned.

If an HMC User Profile/Template corresponding to the user ID field does not exist, or if the password validation fails, the logon request fails and status code 403 (Forbidden) is returned. There is no reason code to distinguish these reasons for the failure. If the User Profile/Template is marked as disabled or the associated password has expired, or if the or if the User Profile/Template is not configured to allow use of the API, the login request also fails with status code 403 (Forbidden) and a reason code identifying the specific cause.

If user authentication is successful and the request body contains the optional **new-password** field, the password associated with the User Profile/Template is changed to the specified new value as part of the **Logon** operation. If the new password does not meet the requirements of the password policy in effect for this User Profile/Template or if the User Profile/Template password is not changeable because it is managed by an external authentication mechanism, the request fails with status code 400 (Bad Request) and a reason code indicating the cause of the failure.

As part of establishing the new API session, JMS topics are created that will be used by the HMC to send object-related and job-related notification messages to this session and the names of these topics are provided in fields of the response body. The name of the topic used for object-related notifications is provided in the **notification-topic** field of the response, and the name of the topic used for job-related notifications is provided in the **job-notification-topic** field.

This operation can be invoked on the alternate HMC, however password changes cannot be requested (i.e. the **new-password** field cannot be provided) in this case.

## Authorization requirements

This operation has the following authorization requirement:

- The HMC User Profile or User Template selected by the **userid** field must be configured to allow use of the Web Services API.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 46.

The following HTTP status codes are returned for the indicated operation-specific errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	12	The request specified an <b>X-API-Session</b> header, which is interpreted as an attempt to unnecessarily logon again when already logged on.
	42	Password changes cannot be requested when logging on to the alternate HMC.
	43	The password for this user cannot be changed, for example because it is managed in an external authentication mechanism such as LDAP.
	44	The new password does not conform to the requirements of the password policy in effect for this user.
	45	The user’s password has expired and no new-password field was specified.
403 (Forbidden)	0	User authentication failed.
	40	The user is disabled.
	41	The user is not authorized to use the HMC Web Services interface.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- The **Logon** operation checks for and prevents requests that specify an **X-API-Session** header on input in order to detect client applications that unnecessarily log on again when already logged on. It is valid to have multiple sessions, but in order to more explicitly indicate that this is desired, the client application needs to request each logon without referencing any existing session.
- Some of the information returned by this operation is also present in the response body of a successful **Get Notification Topics** request. Specifically, the information contained in the **notification-topic** and

**job-notification-topic** fields is also included in the **Get Notification Topics** response. That operation identifies all JMS topics available to the API user, possibly including topics other than those identified in the **Logon** response.

## Example HTTP interaction

---

```
POST /api/session HTTP/1.1
content-type: application/json
content-length: 58
{
  "password": "12345678",
  "userid": "APIUSER"
}
```

---

Figure 1. Logon: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 02 Nov 2011 18:41:27 GMT
x-api-session: 4hy7c4nogldz4b59ajegzb1dulec641ziyv6uf73zs43205edv
content-type: application/json;charset=UTF-8
content-length: 207
{
  "api-major-version": 1,
  "api-minor-version": 1,
  "api-session": "4hy7c4nogldz4b59ajegzb1dulec641ziyv6uf73zs43205edv",
  "job-notification-topic": "APIUSER.229job",
  "notification-topic": "APIUSER.229",
  "password-expires": 29
}
```

---

Figure 2. Logon: Response

## Logoff

The **Logoff** operation closes an API session with the Web Services API.

### HTTP method and URI

**DELETE** /api/sessions/**this-session**

### Description

This operation closes an API session with the Web Services API.

The session to be closed is indicated by the **session-id** in the **X-API-Session** header of the request. If the **session-id** designates an open session, the API session is closed and status code 204 (No Content) is returned. Closing of the API session includes closing/deleting any Metrics Service retrieval contexts or JMS notification topics associated with the session. However, asynchronous actions initiated by the session continue to run.

Once a session is closed, its **session-id** is no longer valid for use in subsequent Web Services API requests. Attempts to do so will result in the same errors as any other attempt to use a session-requiring operation without providing a valid **session-id**.

This operation can be invoked on the alternate HMC.



## Authorization requirements

This operation has the following authorization requirement:

- No explicit authorization is required, however the client application must possess and present a valid **session-id** of the session to be closed.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned with no response body.

The following HTTP status codes are returned for the indicated operation-specific errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
DELETE /api/session/this-session HTTP/1.1
x-api-session: zksmpaxgtcasy5uixmtwaudqe8ha6fy0006bzm2bd8yo
```

---

*Figure 3. Logoff: Request*

---

```
204 No Content
date: Wed, 20 Jul 2011 18:33:56 GMT
x-request-id: Sx32 Rx0
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

---

*Figure 4. Logoff: Response*

## Get Notification Topics

The **Get Notification Topics** operation returns a structure that describes the JMS notification topics associated with the API session. These topics allow the user to receive various types of asynchronous notifications from the HMC.

### HTTP method and URI

**GET** /api/sessions/operations/get-notification-topics

## Response body contents

On successful completion, the response body is a JSON object with the following field:

Field name	Type	Description
topics	Array of objects	Array of nested topic-info objects as described in the next table

Each nested topic-info object contains the following fields:

Table 8. topic-info object

Field name	Type	Description
topic-type	String Enum	<p>The type of notification topic, which provides an indication of the type of data found on the topic. A given value will only be represented at most once within a single response. One of the following values:</p> <ul style="list-style-type: none"> <li>• <b>"object-notification"</b> - The object notification topic. This topic is consistent with the information returned in the <b>notification-topic</b> field in the response body of a successful Logon request. It is used by the HMC to send object-related notifications to this session.</li> <li>• <b>"job-notification"</b>- The job notification topic. This topic is consistent with the information returned in the <b>job-notification-topic</b> field in the response body of a successful Logon request. It is used by the HMC to send job-related notifications to this session.</li> <li>• <b>"audit-notification"</b> - the audit notification topic. This topic is used by the HMC to send audit-related events to this session.</li> <li>• <b>"security-notification"</b> - the security notification topic. This topic is used by the HMC to send security-related events to this session.</li> </ul>
topic-name	String (1-128)	The name of the notification topic. API users can connect using this name to receive notifications for the topic.

## Description

This operation returns a list of all JMS topics to which the API user is authorized to connect. As there exists at least one JMS topic available to any authenticated user, the returned JSON array will never be empty.

## Authorization Requirement

This operation has the following authorization requirement:

- No explicit authorization is required; however, the response to this request is limited to the topics to which the user is authorized to connect.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and an associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
GET /api/sessions/operations/get-notification-topics HTTP/1.1
x-api-session: 21tfe2c2q3ti2b2pwq1wfwuzifo4qymqa8ktzjep7dbyr110k
```

Figure 5. Get Notification Topics: Request

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Sat, 14 Sept 2013 18:03:00 GMT
content-type: application/json;charset=UTF-8
{ "topics" :
  [
    { "topic-type":"object-notification", "topic-name":"mikeuser.1" },
    { "topic-type":"job-notification", "topic-name":"mikeuser.1job" },
    { "topic-type":"audit-notification", "topic-name":"mikeuser.1aud" },
    { "topic-type":"security-notification", "topic-name":"mikeuser.1sec" }
  ]
}
```

Figure 6. Get Notification Topics: Response

### Usage notes

Some of the information returned by this operation is also present in the response body of a successful **Logon** request. This operation is intended to provide a superset and will contain all JMS topics available to the API user including the two topics indicated in the **Logon** response.

## Asynchronous job processing

Some of the operations provided in the Web Services API may take a significant amount of elapsed time to complete. In order to optimize the usage of HMC session resources and to allow the client application the opportunity to perform other processing, such long-running operations are structured to be executed asynchronously (rather than synchronously) from the perspective of the client application.

In a synchronous operation, the Web Services API does not respond to the client application's request until all of the processing associated with the request is complete (successfully or in error) and the API can provide a final result status for the operation. The client application thread is typically blocked (not running) during this time.

By contrast, in a function that operates asynchronously, the Web Services API performs just the minimal front-end validation and set up work needed to accept the request to perform the indicated operation, and then quickly returns an HTTP 202 (Accepted) result to the client indicating that the operation request

has been started but is not yet finished. Along with the HTTP 202 (Accepted) result, the client application is provided with a URI that represents the asynchronous job that is in progress. This URI is of the form `/api/job/{job-id}`.

At any point after receiving the HTTP 202 (Accepted) result, the client application can invoke the **Query Job Status** operation described in this section to determine if the job is complete or not. If the job is not yet complete, the **Query Job Status** request returns an indication that the job is still running. If the job is complete, the **Query Job Status** request returns an indication that the job is now complete along with the final status code, reason code and result data associated with the now-completed asynchronous processing. Once complete, job status is retained by the HMC for a minimum of 4 hours to allow the client application time to retrieve the results, but this status and results are not held indefinitely.

Since the major reason an API operation is structured to be asynchronous is that it will take significant time to complete, very frequent polling for completion via calls to **Query Job Status** can lead to significant unproductive use of client application and HMC resources. In order to eliminate the need to poll at all, the Web Services API also provides asynchronous notifications of job completion via its JMS notification capability. IBM recommends that client applications use this notification facility to determine when a job is complete rather than polling for completion. See "Job completion notification" on page 35 for details on this notification mechanism.

If it is not practical for a client application to use asynchronous notification of job completion, the application should introduce elapsed-time delays between successive **Query Job Status** requests to poll for job completion in order to reduce unproductive use of resources.

## Query Job Status

The **Query Job Status** operation returns the status associated with an asynchronous job.

### HTTP method and URI

**GET** `/api/jobs/{job-id}`

In this request, the URI variable `{job-id}` is the identifier of an asynchronous job associated with this user, as returned in the response of the operation that initiated the job.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
status	String Enum	An indication of the current disposition of the job. The possible values are as follows: <ul style="list-style-type: none"> <li><b>"running"</b> - indicates that the job was found and has not run to completion at the time of the query.</li> <li><b>"complete"</b> - indicates that the job was found and has completed running. The successful or error completion of the job is indicated by the <b>job-status-code</b> and <b>job-reason-code</b> fields.</li> </ul>

Field name	Type	Description
job-status-code	Integer; Field provided only if status is "complete"	<p>The job completion status code. This field is provided only if the <b>status</b> field is set to "complete".</p> <p>This field provides the major status code describing the success or failure completion of the asynchronous action represented by the job. It is expressed in terms of an HTTP status code (i.e. the HTTP status code that would have been returned for the operation had it been performed synchronously).</p> <p>The values provided here and their meaning depend on the particular action that is being performed asynchronously. The description of these values is provided as part of the description for the operation that initiated the asynchronous job.</p>
job-reason-code	Integer; Field provided only if status is "complete"	<p>The job completion reason code. This field is provided only if the <b>status</b> field is set to "complete" and only if the <b>job-status-code</b> field indicates an error completion (status code other than 2XX).</p> <p>When present, this field provides a more detailed reason code describing the success or failure completion of the asynchronous action represented by the job. It is expressed in terms of the API reason code as are returned in standard error response bodies provided by the API.</p> <p>The values provided here and their meaning depend on the particular action that is being performed asynchronously. The description of these values is provided as part of the description or the operation that initiated the asynchronous job.</p>
job-results	Object; Field provided only if status is "complete"	<p>A nested object that provides results for the job. This field is provided only if the <b>status</b> field is set to "complete" but is optional even for complete jobs. If the asynchronous operation has no result information (beyond the job status and reason codes) then this field is omitted.</p> <p>The structure of the nested object provided by this field and its meaning depends on the particular action that is being performed asynchronously. The description of this object's structure is provided as part of the description of the operation that initiated the asynchronous job.</p>

## Description

The **Query Job Status** operation returns the status associated with an asynchronous job. The results depend on whether the job is still running or is complete.

If the job designated by the URI is still running, the operation sets the **status** field in the response body to "running" and provides no other information about the job. The client application may repeat the query at a later time, but should avoid frequent polling since that can lead to unproductive use of client and HMC resources. In order to eliminate the need to poll at all, the client application can (and should) use the asynchronous notifications facility provided by the API to receive notification of job completion via a JMS-based message. See "Job completion notification" on page 35 for details on this notification mechanism.

If the job is complete, the operation sets the **status** field in the response body to "complete" and provides the other completion-related fields defined in the Response Body Contents section above to report the results to the client application. Once complete, job status is retained by the HMC for a minimum of 4 hours to allow the client application time to retrieve the results, but this status and results are not held indefinitely.

If the URI does not designate a job associated with the current API session HTTP status code 404 (Not Found) is returned to the client.

## Authorization requirements

This operation has the following authorization requirement:

- The user must be correctly authenticated.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 52.

The following HTTP status codes are returned for the indicated operation-specific errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The URI does not designate an asynchronous job associated with the API user.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/jobs/86e44546-107f-11e1-bde0-0010184c8334 HTTP/1.1
x-api-session: 21tfe2c2q3ti2b2pwq1wfwuzifoi4rymq8ktzjep7dbyr110k
```

---

*Figure 7. Query Job Status: Request*

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 16 Nov 2011 18:19:35 GMT
content-type: application/json;charset=UTF-8
content-length: 63
{
  "job-reason-code": 0,
  "job-status-code": 200,
  "status": "complete"
}
```

---

*Figure 8. Query Job Status: Response*

## Delete Completed Job Status

The **Delete Completed Job Status** operation deletes the job status and results associated with a job that has been completed.

### HTTP method and URI

```
DELETE /api/jobs/{job-id}
```

In this request, the URI variable *{job-id}* is the identifier of an asynchronous job associated with this session, as returned in the URI of the operation that initiated the job.

## Description

The **Delete Completed Job Status** operation deletes the job status and results associated with a job that has been completed.

If the job designated by the request URI is complete, its completion status and results are deleted from the HMC and status code 204 (No Content) is returned to the client.

If the job is still running, the operation fails and HTTP status code 409 (Conflict) is returned to the client.

If the URI does not designate a job associated with the current API user, or if the job's status has already been deleted (either explicitly, or due to expiration of the status retention interval), HTTP status code 404 (Not Found) is returned to the client.

## Authorization requirements

This operation has the following authorization requirement:

- The user must be correctly authenticated.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated operation-specific errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The URI does not designate an asynchronous job associate with this session.
409 (Conflict)	40	The URI designates an asynchronous job that is not yet complete.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- This operation is defined to operate only on jobs that have been completed, i.e. have a **status** field with value **"complete"**. As a result, this operation cannot be used to cancel an in-progress asynchronous operation. The ability to cancel an in-progress asynchronous operation is not provided by the API.
- Once an asynchronous job is complete, job status is retained by the HMC for a minimum of 4 hours to allow the client application time to retrieve the results, but this status and results are not held indefinitely. At the expiration of the retention interval, job status is deleted as if the Delete Complete Job Status operation were called.

## Example HTTP interaction

---

```
DELETE /api/jobs/86e44546-107f-11e1-bde0-0010184c8334 HTTP/1.1
x-api-session: 21tfe2c2q3ti2b2pwq1wfwuzifoi4rymqa8ktzjep7dbyr110k
```

---

*Figure 9. Delete Completed Job Status: Request*

---

```
204 No Content
date: Wed, 16 Nov 2011 18:19:35 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

<No response body>

---

*Figure 10. Delete Completed Job Status: Response*



## Chapter 7. Ensemble composition

A z Systems ensemble is a grouping of one or more computing systems, referred to as nodes, that are managed in a coordinated way for purposes of virtualization and workload management. Currently, the ensemble can be comprised of zEnterprise (or later) Central Processing Complex (CPC) nodes and their associated system resources or zBX nodes. Each CPC node consists of the traditional z Systems elements (processors, memory, I/O, LPARs) along with an optional CPC-attached processor feature termed an IBM zEnterprise BladeCenter Extension (zBX) (model 003 or earlier) that provides blade-based computing resources. Starting with IBM z BladeCenter Extension (zBX) Model 004, rather than being treated as part of a CPC node, a zBX instead appears as a distinct ensemble node of its own that provides those blade-based resources. In addition to the per-member resources, the ensemble also encompasses certain ensemble-wide resources that are shared by all of the members of the ensemble, including a secure, platform-managed data network.

- l An ensemble will typically consist of at least one node. However, when initially created, an ensemble starts out with no nodes, so an empty ensemble is a legitimate configuration that may sometimes exist.
- l CPCs or zBX nodes are explicitly added as nodes to, or removed from, an ensemble using operations provided in the API or using management tasks on the HMC UI.
- l Not all models of z Systems processors support ensembles and ensemble-based management. Platform support for ensembles and ensemble-based management was first delivered with the zEnterprise family.

For purposes of controlling the configuration and operational state of the elements of an ensemble, each ensemble is managed from a customer-designated primary/alternate pair of Hardware Management Consoles (HMC) that is referred to as the ensemble HMCs for that ensemble. At any one point in time, only one HMC of the pair is in an active primary role, with the second of the pair acting strictly as a passive backup or alternate. The primary ensemble HMC may manage at most one ensemble. The primary ensemble HMC is the component that provides the main administrative user interface for the ensemble, and it is the component that acts as the access point for the majority of the Web Services API. However, a selected set of API function is also available from the alternate HMC.

For purposes of this definition, it is usually sufficient to focus on the primary ensemble HMC only, except when specifically considering aspects that pertain to the setup of the primary/alternate pair or handling the failover mechanism. Therefore, use of the term “ensemble HMC” without further primary or alternate qualification should be interpreted as indicating the primary ensemble HMC.

**Note:** Properties related to the identity and addressing of the alternate HMC for the ensemble managed by the current HMC are provided as part of the data model for the Console object.

### Ensemble composition operations summary

The following tables provide an overview of the ensemble composition operations provided.

*Table 9. Ensemble composition: operations summary*

Operation name	HTTP method and URI path
“List Ensembles” on page 61	GET /api/ensembles
“Get Ensemble Properties” on page 63	GET /api/ensembles/{ensemble-id}
“Update Ensemble Properties” on page 65	POST /api/ensembles/{ensemble-id}

Table 9. Ensemble composition: operations summary (continued)

Operation name	HTTP method and URI path
"List Ensemble Nodes" on page 67	GET /api/ensembles/{ensemble-id}/nodes
"Get Node Properties" on page 69	GET /api/ensembles/{ensemble-id}/nodes/{node-id}
"Add Node to Ensemble" on page 71	POST /api/ensembles/{ensemble-id}/nodes
"Remove Node from Ensemble" on page 74	DELETE /api/ensembles/{ensemble-id}/nodes/{node-id}

Table 10. Ensemble composition: URI variables

Variable	Description
{ensemble-id}	Object ID (UUID) of an Ensemble object
{node-id}	Element ID of a node of an Ensemble.

## Ensemble object

- An ensemble object represents a single z Systems ensemble.

## Data Model

This object includes the properties defined in the "Base managed object properties schema" on page 39, including the operational-status properties, with the following class-specific specialization:

Table 11. Ensemble object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/ URI	The canonical URI path for an ensemble object is of the form /api/ensembles/{ensemble-id} where {ensemble-id} is the value of the <b>object-id</b> property of the ensemble object.
<b>parent</b>	—	String/ URI	An ensemble object is conceptually a root object and has no parent, so this property is always null.
<b>class</b>	—	String	The class of an ensemble object is "ensemble".
<b>name</b>	(w)(pc)	String (1-16)	The display name specified for the ensemble. Alphanumeric, space, or any of "+<=>%&*\\"(),_./:;?" are valid characters.
<b>description</b>	(w)(pc)	String (0-128)	Text describing the Ensemble. Alphanumeric, space, or any of "+<=>%&*\\"(),_./:;?" are valid characters.
<b>status</b>	(sc)	String Enum	The status of the ensemble representing the current communication status between the primary and alternate HMC: <ul style="list-style-type: none"> <li>"<b>alternate-communicating</b>" – The primary is communicating to the alternate</li> <li>"<b>alternate-not-communicating</b>" - The primary is not communicating to the alternate</li> </ul>
<b>additional-status</b>	—	String Enum	An ensemble object has no additional-status.

## Class specific additional properties

In addition to the properties defined in included schemas, this object includes the following additional class-specific properties:

Table 12. Ensemble object: class specific properties

Name	Qualifier	Type	Description
<b>management- enablement-level</b>	(pc)	String Enum	<p>The zManager management enablement level for this ensemble. The level determines which zManager advanced management functions are available for use. Values:</p> <ul style="list-style-type: none"> <li>• <b>"manage"</b>- Gives you the basics for managing an ensemble. It includes HMC operational controls for zBX, change management of firmware across the ensemble, energy monitoring, virtual networking with automated provisioning, virtual server management, and a base level of performance management.</li> <li>• <b>"automate"</b>- Provides more leverage from the ensemble by managing workloads and energy. This level of support includes power capping, power savings mode, and platform performance management.</li> </ul>
<b>cpu-perf-mgmt- enabled-x-hyp</b>	(w)(pc)	Boolean	<p>If true, management of processor performance is enabled for x Hyp virtualization hosts. Management of processor performance is also available for virtual servers.</p> <p>Performance management properties may be updated if the ensemble <b>management-enablement-level</b> is <b>"automate"</b>.</p>
<b>cpu-perf-mgmt- enabled-power-vm</b>	(w)(pc)	Boolean	<p>If true, management of processor performance is enabled for PowerVM virtualization hosts. Management of processor performance is also available for virtual servers.</p> <p>Performance management properties may be updated if the ensemble <b>management-enablement-level</b> is <b>"automate"</b>.</p>
<b>cpu-perf-mgmt- enabled-zvm</b>	(w)(pc)	Boolean	<p>If true, management of processor performance is enabled for z/VM virtualization hosts. Management of processor performance is also available for virtual servers.</p> <p>Performance management properties may be updated if the ensemble <b>management-enablement-level</b> is <b>"automate"</b>.</p>
<b>unique-local-unified- prefix</b>	(pc)	String	<p>Unique local address (ULA) prefix applied to addresses used for management communication between virtual servers and their virtualization hosts.</p> <p>The ULA prefix of the form fdXX:XXXX:XXXX:./48 is formed by substituting the X's with a pseudo-random 40-bit global ID using the algorithm defined in RFC 4193.</p> <p>The prefix may not be updated through the API.</p>
<b>load-balancing- enabled</b>	(w)(pc)	Boolean	<p>If true, Load Balancing is enabled for this ensemble.</p> <p>Load Balancing properties may be updated if the ensemble <b>management-enablement-level</b> is <b>"automate"</b>.</p>
<b>load-balancing-port</b>	(w)(pc)	Integer (1024- 65535)	<p>The Load Balancing port value in the range 1024-65534. The default port is 3860.</p> <p>Load Balancing properties may be updated if the ensemble <b>management-enablement-level</b> is <b>"automate"</b>.</p>

Table 12. Ensemble object: class specific properties (continued)

Name	Qualifier	Type	Description
<b>load-balancing-ip-addresses</b>	(w)(pc)	Array of String	<p>The IPV4 address or IPV6 addresses of Load Balancers allowed access to the Load Balancing function.</p> <p>The strings are in dotted-decimal form (“nnn.nnn.nnn.nnn”) for IPV4 addresses.</p> <p>The strings are in eight groups of four hexadecimal digits separated by colons (e.g. 2001:0db8:85a3:0000:0000:8a2e:0370:7334), for IPV6 addresses.</p> <p>Load Balancing properties may be updated if the ensemble <b>management-enablement-level</b> is <b>"automate"</b>.</p>
<b>mac-prefix</b>	(pc)	String (2)	<p>The Prefix Address <i>xx:00:00:00:00:00/8</i> is the ensemble mac prefix. All Mac addresses dynamically generated by zManager will be within the ensemble mac prefix. <i>xx</i> defaults to 02.</p> <p>The MAC prefix may not be updated through the API.</p>
<b>reserved-mac-address-prefixes</b>	(pc)	Array of objects	<p>The list of reserved MAC address prefixes. Each reserved MAC address prefix will be an object in the form of a MAC address prefix nested object, as described in “MAC address prefix nested object.”</p> <p>MAC address prefixes may not be updated through the API.</p>
<b>max-nodes</b>	—	Integer	The maximum number of nodes allowed in the ensemble.
<b>max-cpc-nodes</b>	—	Integer	The maximum number of CPC nodes allowed in the ensemble.
<b>max-zbx-nodes</b>	—	Integer	The maximum number of BXs (either CPC-attached or nodes) allowed in the ensemble.

**Energy management related additional properties:** In addition to the properties defined above, this object includes the following additional class-specific properties related to energy management. For further explanation of the various states involved, please see “Special states” on page 181,

Table 13. Ensemble composition: energy management related additional properties

Name	Qualifier	Type	Description
<b>power-rating</b>	—	Integer	Specifies the maximum power usage in watts (W) of this ensemble. This is a calculated value as indicated by the electrical rating labels or system rating plates of the ensemble components.
<b>power-consumption</b>	(mg)	Integer	Specifies the current power consumption in watts (W) for this ensemble.

**MAC address prefix nested object:** A MAC address prefix object is a nested object of an ensemble object. An ensemble may contain zero or more MAC address prefixes. MAC address prefix properties may not be updated through the API.

The following properties are supported:

Table 14. Ensemble composition: MAC address prefix nested object related additional properties

Field name	Type	Description
<b>mac-address</b>	String	The MAC address represented as 6 groups of two hexadecimal digits separated by colons, e.g. 01:23:45:67:89:ab. The MAC address uses the ensemble prefix.

Table 14. Ensemble composition: MAC address prefix nested object related additional properties (continued)

Field name	Type	Description
prefix-length	Integer	The bit length of the MAC address prefix. This is a 2-digit value with these parameters in the range 12-44.

## Node object

A node is an element of the ensemble object that represents a system that is currently a member of the ensemble. Each node contains the following properties:

Table 15. Ensemble composition: node properties

Field name	Type	Description
element-uri	String/ URI	The URI path for a node of an ensemble is of the form <code>/api/ensembles/{ensemble-id}/nodes/{node-id}</code> where <code>{ensemble-id}</code> is the value of the <b>object-id</b> property of the ensemble, and <code>{node-id}</code> is a locally unique element ID for the node.  For nodes of <b>type "cpc"</b> , <code>{node-id}</code> is the value of the <b>object-id</b> property of the CPC that is represented by the node.  For nodes of <b>type "zbx"</b> , <code>{node-id}</code> is the value of the <b>object-id</b> property of the zBX node that is represented by the node.
parent	String/ URI	The canonical URI path of the ensemble containing this node.
class	String (4)	The value <b>"node"</b> .
type	String Enum	This is the type of node. Nodes have the following types: <ul style="list-style-type: none"> <li>• <b>"cpc"</b></li> <li>• <b>"zbx"</b></li> </ul>
member	String/ URI	The canonical URI path of the system element that is represented as a member of the ensemble by this object.  For nodes of type <b>"cpc"</b> , this is the URI path of the CPC object.  For nodes of type <b>"zbx"</b> , this is the URI path of the zBX node object.

## Operations

### List Ensembles

The **List Ensembles** operation lists the ensembles managed by the HMC.

### HTTP method and URI

GET `/api/ensembles`

### Query parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) used to limit returned objects. If matches are found, the response will be an array with all ensembles that match. If no match is found, the response will be an empty array.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
ensembles	Array of objects	Array of nested ensemble-info objects (described in the next table). If the HMC is not a primary HMC, an empty array is provided.

Each nested ensemble-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the Ensemble object, in the form <code>/api/ensembles/{ensemble-id}</code> .
name	String	Display name of the Ensemble object.
status	String Enum	The <b>status</b> property of the Ensemble object.

## Description

This operation lists the ensembles that are managed by this HMC. The object URI, display name, and status are provided for each.

If the **name** query parameter is specified, the returned list is limited to those ensembles that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

An ensemble is included in the list only if the API user has object-access permission for that object. If an HMC is a manager of an ensemble but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

If the HMC does not manage any ensembles, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to any ensemble object to be included in the result.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

This operation, as well as other aspects of the Web Services API, has been structured to allow for the possibility that an HMC might be the manager of multiple ensembles. However, in the current implementation, an HMC can be the manager of at most one ensemble, so the resulting list of ensembles will contain either zero or one entries.

## Example HTTP interaction

---

```
GET /api/ensembles HTTP/1.1
x-api-session: 1f2g70m2e9b4sawt53ydp1oc9nzsz4sduc2wr2bzgmy5dhs7pz
```

---

Figure 11. List Ensembles: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 20 Jul 2011 18:41:03 GMT
content-type: application/json;charset=UTF-8
content-length: 207
{
  "ensembles": [
    {
      "name": "R32Ensemble",
      "object-uri": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026",
      "status": "alternate-communicating"
    }
  ]
}
```

---

Figure 12. List Ensembles: Response

## Get Ensemble Properties

The **Get Ensemble Properties** operation retrieves the properties of a single Ensemble object that is designated by its object ID.

### HTTP method and URI

```
GET /api/ensembles/{ensemble-id}
```

In this request, the URI variable *{ensemble-id}* is the object ID of the Ensemble object for which properties are to be obtained.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the ensemble object as defined in the Data Model section above. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

The operation returns the current properties for the ensemble object specified by *{ensemble-id}*.

On successful execution, all of the current properties as defined by the Data Model for the ensemble object are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing ensemble object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the Ensemble object designated by *{ensemble-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is provided and the response body is as described in “Response body contents” on page 63.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object-id in the URI ( <i>{ensemble-id}</i> ) does not designate an existing ensemble object, or the API user does not have object-access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026 HTTP/1.1
x-api-session: 1f2g70m2e9b4sawt53ydp1oc9nzzs4sduc2wr2bzgmy5dhs7pz
```

---

Figure 13. Get Ensemble Properties: Request



---

```

200 OK
x-request-id: Sx3a Rx1
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 20 Jul 2011 18:41:03 GMT
content-type: application/json;charset=UTF-8
content-length: 959
{
  "acceptable-status": [
    "alternate-communicating"
  ],
  "class": "ensemble",
  "cpu-perf-mgmt-enabled-power-vm": true,
  "cpu-perf-mgmt-enabled-zvm": true,
  "description": "",
  "has-unacceptable-status": false,
  "is-locked": false,
  "load-balancing-enabled": true,
  "load-balancing-ip-addresses": [
    "1.1.1.1"
  ],
  "load-balancing-port": 9876,
  "mac-prefix": "02:00:00:00:00:00",
  "management-enablement-level": "automate",
  "name": "R32Ensemble",
  "object-id": "87d73ffc-75b2-11e0-9ba3-0010184c8026",
  "object-uri": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026",
  "parent": null,
  "power-consumption": 8311,
  "power-rating": 46522,
  "reserved-mac-address-prefixes": [
    {
      "mac-address": "02:c1:00:00:00:00",
      "prefix-length": 16
    }
  ],
  "status": "alternate-communicating",
  "unique-local-unified-prefix": "fd07:8c9:1ba3:0:0:0:0:0"
}

```

---

Figure 14. Get Ensemble Properties: Response

## Update Ensemble Properties

The **Update Ensemble Properties** operation modifies simple writable properties of an ensemble object.

### HTTP method and URI

**POST** /api/ensembles/{ensemble-id}

In this request, the URI variable {ensemble-id} is the object ID of the Ensemble object for which properties are to be updated.

### Request body contents

The request body is expected to contain a JSON object with one or more of the following fields. Only fields that are being changed are necessary to supply.

Field name	Type	Description
name	String (1-16)	The new name to give the ensemble, as described in “Data Model” on page 58.
description	String (0-128)	The new description to give the ensemble, as described in “Data Model” on page 58.
cpu-perf-mgmt-enabled-power-vm	Boolean	The PowerVM virtualization host processor performance management enablement setting, as described in “Data Model” on page 58.
cpu-perf-mgmt-enabled-zvm	Boolean	The z/VM virtualization host processor performance management enablement setting, as described in “Data Model” on page 58.
load-balancing-enabled	Boolean	The Load Balancing enablement setting for this ensemble, as described in “Data Model” on page 58.
load-balancing-port	Integer	The Load Balancing port for this ensemble, as described in “Data Model” on page 58.
load-balancing-ip-address	Array of Strings	The Load Balancing ip addresses, as described in “Data Model” on page 58.

## Description

This operation updates one or more simple writeable properties of the ensemble object identified by *{ensemble-id}*.

On successful execution, the ensemble object has been updated with the supplied property values and status code 204 (No Content) is returned without supplying a response body.

If the update changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

The URI path must designate an ensemble object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must also have permission to the Ensemble Details task as well, otherwise status code 403 (Forbidden) is returned.

The request body is validated against the schema described in the Request Body Contents section. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the Ensemble object designated by *{ensemble-id}*
- Action/task permission to the **Ensemble Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
	228	The ensemble's <b>management-enablement-level</b> property does not allow the updating of a property specified in the request body.
404 (Not Found)	1	The object ID in the URI ( <i>ensemble-id</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to the object.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	3	The operation cannot be performed because the object designated by the request URI is currently locked.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334 HTTP/1.1
x-api-session: 297n8iun1251svgcju9tvsai0rrew4ieawx97ykucbxy69bwr2
content-type: application/json
content-length: 34
{
  "name": "SS-Ensemble-1"
}
```

---

Figure 15. Update Ensemble Properties: Request

---

```
204 No Content
date: Wed, 07 Dec 2011 04:54:01 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 16. Update Ensemble Properties: Response

## List Ensemble Nodes

The **List Ensemble Nodes** operation lists the nodes of an ensemble managed by the HMC.

### HTTP method and URI

```
GET /api/ensembles/{ensemble-id}/nodes
```

In this request, the URI variable *{ensemble-id}* is the object ID of the Ensemble object for which nodes are to be listed.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
nodes	Array of objects	Array of nested node-info objects (described in the next table). If the ensemble has no nodes, an empty array is provided.

Each nested node-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	URI path of the ensemble node object, in the form <code>/api/ensembles/{ensemble-id}/nodes/{node-id}</code> .  For a node of <b>type "cpc"</b> , the node object represents a CPC as a member of the ensemble, and the <code>{node-id}</code> component of this URI path is the <b>object-id</b> property of the underlying CPC object.  For a node of <b>type "zbx"</b> , the node object represents a zBX node as a member of the ensemble, and the <code>{node-id}</code> component of this URI path is the <b>object-id</b> property of the underlying zBX node object.
type	String Enum	The <b>type</b> property of the underlying object that is represented by the node.
name	String	The <b>name</b> property of the underlying object that is represented by the node
status	String Enum	The <b>status</b> property of the underlying object that is represented by the node

## Description

This operation lists the nodes of an ensemble specified by its `{ensemble-id}`. The element URI and type are provided for each.

A node is included in the list only if the API user has object-access permission to the underlying object that is represented by the node. For a node of **type "cpc"**, the underlying object is a CPC. For a node of **type "zbx"**, the underlying object is a BX node. If an ensemble contains a node but the API user does not have permission to the related object, that object is simply omitted from the list but no error status code results.

If the ensemble is empty (has no nodes), an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the ensemble designated in the request URI
- Object access permission to the underlying object that is represented by a node included in the result.  
For nodes of **type "cpc"**, the underlying object is a CPC. For nodes of **type "zbx"**, the underlying object is a zBX node.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>ensemble-id</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes HTTP/1.1
x-api-session: 1f2g70m2e9b4sawt53ydp1oc9nzzs4sduc2wr2bzbzgm5dhs7pz
```

---

Figure 17. List Ensemble Nodes: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 20 Jul 2011 18:41:03 GMT
content-type: application/json;charset=UTF-8
content-length: 250
{
  "nodes": [
    {
      "element-uri": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
      "name": "R32",
      "status": "operating",
      "type": "cpc"
    }
  ]
}
```

---

Figure 18. List Ensemble Nodes: Response

### Get Node Properties

The **Get Node Properties** operation retrieves the properties of a single Node object.

#### HTTP method and URI

```
GET /api/ensembles/{ensemble-id}/nodes/{node-id}
```

#### URI variables

Variable	Description
<i>{ensemble-id}</i>	Object ID of the ensemble containing the node for which properties are to be obtained.
<i>{node-id}</i>	Element ID of the node which for which properties are to be obtained.

## Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the node object as defined in “Data Model” on page 58. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

The operation returns the current properties for the node object specified by the request URI.

On successful execution, all of the current properties as defined by the Data Model for the ensemble object are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing node object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the ensemble designated in the request URI
- Object access permission to the underlying object that is represented by a node included in the result.  
For nodes of type "cpc", the underlying object is a CPC.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is provided and the response body is as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to the object.
	225	The element-id in the URI ( <i>{node-id}</i> ) does not designate an existing node of the ensemble designed by <i>{ensemble-id}</i> .
	227	The element-id in the URI ( <i>{node-id}</i> ) does not designate an existing object, or the API user does not have access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
GET /api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-e79e49dd340 HTTP/1.1
x-api-session: 1f2g70m2e9b4sawt53ydp1oc9nzzs4sduc2wr2bzgmy5dhs7pz
```

Figure 19. Get Node Properties: Request

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 20 Jul 2011 18:41:03 GMT
content-type: application/json;charset=UTF-8
content-length: 295
{
  "class": "node",
  "element-uri": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-
    be79e49dd340",
  "member": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026",
  "type": "cpc"
}
```

Figure 20. Get Node Properties: Response

## Add Node to Ensemble

- The **Add Node to Ensemble** operation adds a CPC or zBX node to an ensemble, creating a new Node object to represent the membership as a result.

### HTTP method and URI

**POST** /api/ensembles/{ensemble-id}/nodes

- In this request, the URI variable {ensemble-id} is the object ID of the ensemble object to which the CPC or zBX node is to be added as a node.

### Request body contents

- The request body is a JSON object requiring exactly one of the following fields:

Field name	Type	Description
cpc	String/URI	The canonical URI path identifying the CPC to be added to the targeted ensemble as a new node. If this field is provided, the <b>member</b> field is not allowed.
member	String/URI	The canonical URI path of the system element that is represented as a member of the ensemble by this object.  If a CPC is to be added, this is the URI path of a CPC object. If a zBX node is to be added, this is the URI path of a zBX object that has a <b>type</b> property with value <b>"node"</b> . If this field is provided, the <b>cpc</b> field is not allowed.

### Response body contents

Field name	Type	Description
node-uri	String/URI	The canonical URI path identifying the node that was created in the targeted ensemble, in the form /api/ensembles/{ensemble-id}/nodes/{node-id}.

## Description

- | This operation adds a CPC or zBX node to the ensemble targeted by the request URI and creates a new
- | Node object to represent the membership. Refer to the *z Systems Ensemble Planning Guide* for details on
- | managing members of an ensemble.
- | All ensemble nodes must have the same **management-enablement-level**. A node with a
- | **management-enablement-level** of "automate" may be added to an ensemble with a **management-**
- | **enablement-level** of "manage", but the node (not the ensemble) will be downgraded to a
- | **management-enablement-level** of "manage". To later upgrade a downgraded node, all the other nodes in
- | the ensemble must have their **management-enablement-level** upgraded to "automate" (achieved by
- | installing the appropriate LICCC records on the nodes).

Upon successful completion, HTTP status code 201 (Created) is returned and the response body includes the URI of the node that was created to represent the membership. This URI is also provided as the value of the **Location** header in the response.

If the CPC is already a node of an ensemble (either the targeted ensemble or another one) HTTP status code 400 (Bad Request) is returned with associated reason code 224.

- | The URI path must designate an existing ensemble and the API user must have object-access permission
- | to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API
- | user must also have action/task permission to the Add Member to Ensemble task as well, otherwise
- | status code 403 (Forbidden) is returned. Additionally if the CPC or zBX node is ineligible to be added to
- | the ensemble a status code 409 (Conflict) is returned.

The request body is validated against the schema described in "Request body contents" on page 71. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the ensemble object passed in the request URI
- | • Object access permission to the CPC or zBX node object passed in the request body.
- Action/task permission to the Add Member to Ensemble task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in "Response body contents" on page 71.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.



HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	6	The request body contains an unrecognized field (i.e. one that is not listed as either required or optional in the specification for the request body format for the operation).
	221	The operation cannot be performed because the CPC or zBX node to be added as a member is not an eligible machine model.
	222	The operation cannot be performed because the CPC does not have the Ensemble without zBX Feature or the zBX Feature installed.
	224	The operation cannot be performed because the CPC or zBX node is a member of another ensemble or already a member of the ensemble targeted in the request URI.
	229	The operation cannot be performed because the maximum number of CPCs would be exceeded for the ensemble.
	230	A CPC or zBX node may be added to an Ensemble only if the LICCC QoS value of the CPC or zBX node is the same as or is higher than the aggregated QoS of the Ensemble.
	232	The operation cannot be performed because the maximum number of zBXs would be exceeded for the ensemble.
	233	The operation cannot be performed because the maximum number of nodes would be exceeded for the ensemble.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>ensemble-id</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to the object.
	227	The object-id in the URI of the CPC or zBX node object in the request body does not designate an existing object, or the API user does not have access permission to it.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	3	The operation cannot be performed because the object designated by the request URI is currently locked to prevent disruptive changes from being made.
	220	The operation cannot be performed because the CPC or zBX node to be added as a member is currently busy performing some other operation.
	223	The operation cannot be performed because the CPC or zBX node to be added as a member is not in the correct state to be added to the ensemble. The CPC cannot be added if the CPC or zBX node operational status is <b>"not-communicating"</b> .
	226	The operation cannot be performed because the CPC or zBX node object to be added as a member is currently locked to prevent disruptive changes from being made.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- This operation continues to support the "cpc" input field for compatibility when adding a CPC to an ensemble. However, you should use the "member" input field for a more general approach to add either a CPC or zBX node to an ensemble.
- When a node's **management-enablement-level** is downgraded to **"manage"**, management functions requiring a higher **management-enablement-level** can no longer be executed. Such functions, like Energy Management, will no longer be available.

## Remove Node from Ensemble

The **Remove Node from Ensemble** operation removes a node from an ensemble.

### HTTP method and URI

**DELETE** /api/ensembles/{ensemble-id}/nodes/{node-id}

#### URI variables

Variable	Description
{ensemble-id}	Object ID of the ensemble from which the targeted node is to be removed as a member.
{node-id}	Element ID of the node which is to be removed from the targeted ensemble.

### Description

This operation removes a specified node from the specified ensemble. The node is identified by the {node-id} variable in the URI, and the ensemble is identified by the {ensemble-id} variable in the request URI.

- Refer to the *z Systems Ensemble Planning Guide* for details on managing members of an ensemble.

Upon successfully removing the node as a member, HTTP status code 204 (No Content) is returned.

- The URI path must designate an existing ensemble object and the API user must have object-access permission to it. Furthermore, the URI path must designate an existing node element. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must also have Remove Member from Ensemble action/task permission as well, otherwise status code 403 (Forbidden) is returned. Additionally if the CPC or zBX node is ineligible to be removed from the ensemble a status code 409 (Conflict) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the ensemble object passed in the request URI
- Object access permission to the system element represented by the node element. For nodes of **type "cpc"**, this is a CPC object. For nodes of **type "zbx"**, this is a zBX node object.
- Action/task permission to the **Remove Member from Ensemble** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	225	The operation cannot be performed because the node ( <i>{node-id}</i> ) to be removed is not a member of the ensemble ( <i>{ensemble-id}</i> ) designated by the request URI.
	231	The operation cannot be performed because the node ( <i>{node-id}</i> ) to be removed has entitled blades.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to the object.
	227	The element-id in the URI ( <i>{node-id}</i> ) does not designate an existing object, or the API user does not have access permission to it.
409 (Conflict)	2	The operation cannot be performed because the object ( <i>{ensemble-id}</i> ) designated by the request URI is currently busy performing some other operation.
	3	The operation cannot be performed because the object ( <i>{ensemble-id}</i> ) designated by the request URI is currently locked to prevent disruptive changes from being made.
	220	The operation cannot be performed because the node ( <i>{node-id}</i> ) to be removed as a member is currently busy performing some other operation.
	223	The operation cannot be performed because the node ( <i>{node-id}</i> ) to be removed as a member is not in the correct state to be removed from the ensemble. The node cannot be removed if the operational status is <b>"not-communicating"</b> .
	226	The operation cannot be performed because the node to be removed as a member is currently locked to prevent disruptive changes from being made.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Inventory service data

Information about the Ensembles managed by the HMC and the associated nodes can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for ensemble and node objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "ensemble" are to be included. Information for a particular ensemble (and associated node) is included only if the API user has object-access permission to that object.

For each ensemble to be included, the inventory response array includes the following:

- An array entry for the ensemble object itself. This entry is a JSON object with the same contents as is specified in the Response Body Contents section for the **Get Ensemble Properties** operation. That is, the data provided is the same as would be provided if a **Get Ensemble Properties** operation were requested targeting this object.
- An array entry for each node associated with the ensemble. For each such node, an entry is included that is a JSON object with the same contents as specified in the Response Body Contents section of the

**Get Ensemble Properties** operation. As a result, the data provided is the same as would be obtained if a **Get Node Properties** operation were requested for each node listed by a **List Ensemble Nodes** operation targeting the ensemble.

The array entry for an ensemble object will appear in the results array before entries for associated nodes.

## Sample inventory data

The following fragment is an example of the JSON object that would be included in the Get Inventory response to describe an ensemble (named "R32Ensemble") with a single node as member. These objects would appear as a sequence of array entries in the response array:

---

```
{
  "acceptable-status": [
    "alternate-communicating"
  ],
  "class": "ensemble",
  "cpu-perf-mgmt-enabled-power-vm": false,
  "cpu-perf-mgmt-enabled-zvm": true,
  "description": "FVT Test",
  "has-unacceptable-status": false,
  "is-locked": false,
  "load-balancing-enabled": true,
  "load-balancing-ip-addresses": [
    "1.1.1.1"
  ],
  "load-balancing-port": 9876,
  "mac-prefix": "02:00:00:00:00:00",
  "management-enablement-level": "automate",
  "name": "R32Ensemble",
  "object-id": "87d73ffc-75b2-11e0-9ba3-0010184c8026",
  "object-uri": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026",
  "parent": null,
  "power-consumption": 8298,
  "power-rating": 46288,
  "reserved-mac-address-prefixes": [],
  "status": "alternate-communicating",
  "unique-local-unified-prefix": "fd07:8c9:1ba3:0:0:0:0:0"
},
{
  "class": "node",
  "element-uri": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "member": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026",
  "type": "cpc"
}
```

---

Figure 21. Ensemble object: Sample inventory data

## Usage notes

When configured as recommended by IBM, the process of recovering from the failure of the primary ensemble-management HMC by takeover by the alternate HMC includes movement of the IP address of the former primary to the new primary. When this occurs, explicit redirection of API requests to the newly designated primary HMC is not needed. However, the IP address swapping may not be possible in certain network configurations. The address of the alternate ensemble-management HMC for the current Ensemble is provided as properties of the Console object (representing the current HMC) to allow applications to explicitly redirect requests to the other HMC of the pair in these cases.

---

## Chapter 8. zBX infrastructure elements

- | A z Systems ensemble is a grouping of one or more CPC nodes and optional IBM z BladeCenter Extension (zBX) Model 004 nodes that are managed together in a coordinated way for purposes of virtualization and workload management. CPC members of the ensemble may also have optional CPC-attached IBM zEnterprise BladeCenter Extension (zBX) Model 002 or 003 features if supported by the processor model.
- | A zBX, in either Model 004 node form, or the Model 002 or 003 CPC-attached or node form, has as its key components one or more BladeCenter chassis, one or more Blades and VLAN-capable switches all mounted in dedicated racks. Each zBX is connected to other members of the ensemble through a dedicated and integrated data network. In addition, the components of a zBX are connected to their controlling Support Element (either in the associated CPC for a zBX Model 002 or 003, or within the zBX node itself for a zBX Model 004) through a dedicated and integrated management network.

Within the ensemble, advanced management capabilities are provided for the blades housed within the zBX. The zBX components are configured, managed and serviced in the same way as the other components of a z Systems CPC. Management of the zBX is provided only via the controlling Support Element when the CPC or zBX node is a member of an ensemble.

---

### zBX physical network overview

The zBX contains physical network switches that provide the connectivity between the blades and CPCs in the Intra-Ensemble Data Network (IEDN). There are two types of Ethernet switches within each zBX:

- Top-of-Rack Switches (TORs) – A pair of TORs reside in each zBX and act as a primary and backup. TORs connect the blades in the zBX to z Systems network interfaces, and to other external networking equipment, such as routers.
- Ethernet Switch Modules (ESMs) – A pair of ESMs reside in each BladeCenter chassis and connect the blades in the zBX to the IEDN and provide the links to the TORs.

The initial configuration and setup of the physical switches are provided by zManager. The ESMs are not accessible for configuration changes through the Web Services API, however, performance metrics are provided through the Metrics Service. Some TOR management will typically be required by an external administrator; therefore, when managing virtual networks, administrators must consider requirements for configuring the Top-of-Rack switch ports. Only certain types of TOR ports support configuration by the zManager user, these are:

- External - Ports that connect to a customer's external network.
- Internal - Ports are internal ports that extend to the ESMs which connect to the blades. The configuration of this type of port is intended only to support the ISAOPT Coordinator. In this case, VLAN- tagging is required to allow traffic from ISAOPT to z Systems to be tagged with the proper VLAN ID.

The configuration properties supported by these ports are:

- Virtual networks - The following port types can be configured to a virtual network:
  - External
  - Internal
- MAC Filters – The following port types support MAC filters:
  - External

## zBX infrastructure operations summary

The following tables provide an overview of the operations provided.

Table 16. zBX infrastructure: operations summary

Operation name	HTTP method and URI path
"List zBXs of a CPC" on page 87	GET /api/cpcs/{cpc-id}/zbx
"List zBXs of an Ensemble" on page 89	GET /api/ensembles/{ensemble-id}/zbx
"Get zBX Properties" on page 91	GET /api/zbx/{zbx-id}
"Get EC/MCL Description of zBX (Node)" on page 95	GET /api/zbx/{zbx-id}/operations/get-ec-mcl-description
"Activate zBX (Node)" on page 98	POST /api/zbx/{zbx-id}/operations/activate
"Deactivate zBX (Node)" on page 100	POST /api/zbx/{zbx-id}/operations/deactivate
"Get zBX (Node) Audit Log" on page 103	GET /api/zbx/{zbx-id}/operations/get-audit-log
"Get zBX (Node) Security Log" on page 105	GET /api/zbx/{zbx-id}/operations/get-security-log
"List zBX (Node) Hardware Messages" on page 108	GET /api/zbx/{zbx-id}/hardware-messages
"Get zBX (Node) Hardware Message Properties" on page 110	GET /api/zbx/{zbx-id}/hardware-messages/{hardware-message-id}
"Delete zBX (Node) Hardware Message" on page 112	DELETE /api/zbx/{zbx-id}/hardware-messages/{hardware-message-id}
"List Top-of-Rack Switches of a zBX" on page 117	GET /api/zbx/{zbx-id}/top-of-rack-switches
"Get Top-of-Rack Switch Properties" on page 119	GET /api/zbx/{zbx-id}/top-of-rack-switches/{tor-id}
"Get Top-of-Rack Switch Port Details" on page 121	GET /api/zbx/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}
"Update Top-of-Rack Switch Port Properties" on page 123	POST /api/zbx/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}
"Add MAC Filters to Top-of-Rack Switch Port" on page 125	POST /api/zbx/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}/operations/add-mac-filters
"Remove MAC Filters from Top-of-Rack Switch Port" on page 127	POST /api/zbx/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}/operations/remove-mac-filters
"Add Top-of-Rack Switch Port to Virtual Networks" on page 129	POST /api/zbx/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}/operations/add-virtual-networks
"Remove Top-of-Rack Switch Port from the Virtual Networks" on page 131	POST /api/zbx/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}/operations/remove-virtual-networks

Table 16. zBX infrastructure: operations summary (continued)

Operation name	HTTP method and URI path
“List Racks of a zBX” on page 134	GET /api/zbx/{zbx-id}/racks
“Get Rack Properties” on page 136	GET /api/racks/{rack-id}
“List BladeCenters in a Rack” on page 141	GET /api/racks/{rack-id}/bladecenters
“List BladeCenters in a zBX” on page 143	GET /api/zbx/{zbx-id}/bladecenters
“Get BladeCenter Properties” on page 145	GET /api/bladecenters/{bladecenter-id}
“Activate BladeCenter” on page 147	POST /api/bladecenters/{bladecenter-id}/operations/activate
“Deactivate BladeCenter” on page 150	POST /api/bladecenters/{bladecenter-id}/operations/deactivate
“List Blades in a BladeCenter” on page 158	GET /api/bladecenters/{bladecenter-id}/blades
“List Blades in a zBX” on page 160	GET /api/zbx/{zbx-id}/blades
“Get Blade Properties” on page 163	GET /api/blades/{blade-id}
“Activate a Blade” on page 167	POST /api/blades/{blade-id}/operations/activate
“Deactivate a Blade” on page 169	POST /api/blades/{blade-id}/operations/deactivate
“Create IEDN Interface for a DataPower XI50z Blade” on page 171	POST /api/blades/{blade-id}/iedn-interface
“Delete IEDN Interface for a DataPower XI50z Blade” on page 174	DELETE /api/blades/{blade-id}/iedn-interface/{iedn-interface-id}

Table 17. zBX infrastructure: URI variables

Variable	Description
{cpc-id}	Object ID of a CPC
{zbx-id}	Object ID of a zBX
{tor-id}	Element ID of a TOR object
{port-id}	Element ID of a TOR port
{rack-id}	Object ID of a Rack
{bladecenter-id}	Object ID of a BladeCenter
{blade-id}	Object ID of a Blade
{iedn-interface-id}	Element ID of an IEDN Interface

## zBX object

- | A zBX object represents a zBX. The zBX may be a zBX Model 004 ensemble node (referred to as a zBX node), or a CPC-attached zBX Model 003 or earlier.

## Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, including the operational-status properties when representing a zBX node, with the following class-specific specialization:

Table 18. zBX object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>name</b>	(ro)	String (1-64)	The name of the zBX. If the value of the <b>type</b> property is " <b>cpc-attached</b> ", the name is assigned by zManager as 21 characters of the form <machine type>-<machine model>-<machine serial number> of the zBX object. <sup>1</sup>  If the value of the <b>type</b> property is " <b>node</b> ", the name of the zBX is assigned during installation.
<b>description</b>	—	String	The description property of the zBX object.
<b>object-uri</b>	—	String/ URI	The canonical URI path for a zBX object is of the form /api/zbx/{zbx-id}, where {zbx-id} is the value of the <b>object-id</b> property of the zBX object.
<b>parent</b>	—	String/ URI	The URI path of the parent of this zBX.  If the value of the <b>type</b> property is " <b>cpc-attached</b> ", this property provides the canonical URI path of the parent CPC object.  If the value of the <b>type</b> property is " <b>node</b> ": <ul style="list-style-type: none"> <li>• If the zBX is a member of an ensemble, this property contains the canonical URI path for the Ensemble object.</li> <li>• Otherwise, the value of this property is null.</li> </ul>
<b>class</b>	—	String	The class of a zBX object has the value " <b>zbx</b> ".
<b>status</b>	(sc)	String Enum	The current operational status of the zBX object. One of: <ul style="list-style-type: none"> <li>• "<b>operating</b>" – all blades and all BladeCenter chassis of the zBX have a status of "<b>operating</b>". A blade with a status of "<b>no-power</b>" state results in zBX status of "<b>exceptions</b>". A blade with a status of "<b>service</b>" is excluded from summarized status for the zBX which would not affect the zBX having a status of "<b>operating</b>".</li> <li>• "<b>not-communicating</b>" - the HMC is not communicating with the Support Element of the zBX.</li> <li>• "<b>exceptions</b>" – at least one blade or BladeCenter chassis has a status not equal to 'operating'. Objects with a status of "<b>service</b>" are excluded.</li> <li>• "<b>status-check</b>" - the Primary SE is not communicating with the zBX (or all blades and all BladeCenter chassis have a status of "<b>status-check</b>").</li> <li>• "<b>service</b>" – the zBX has been placed in service mode.</li> <li>• "<b>no-power</b>" - all blades in all BladeCenter chassis of the zBX have a status of "<b>no-power</b>". The power state of each BladeCenter chassis is not part of this summary status.</li> <li>• "<b>service-required</b>" – service is required for the zBX. It is operating on the last redundant part of a particular type.</li> </ul> <p>This property is provided only if the value of the <b>type</b> property is "<b>node</b>", otherwise it is omitted.</p>
<b>acceptable-status</b>	(w)(pc)	Array of String Enum	The set of operational status values in which the zBX object can exist and be considered in an acceptable (not alert causing) state. One or more of the values listed for the <b>status</b> property.  This property is provided only if the value of the <b>type</b> property is " <b>node</b> ", otherwise it is omitted.



Table 18. zBX object: base managed object properties specializations (continued)

Name	Qualifier	Type	Description of specialization
<p><b>Note:</b> <sup>1</sup> This <b>name</b> property for a CPC-attached zBX is currently assigned by zManager and is not writeable. However, it is possible that the API could be extended to allow this property to be writeable, in which case an API or User-Interface user could change the name to contain arbitrary data. Therefore, API client applications should not rely on the contents and format of the <b>name</b> property always being in the form of the zManager-assigned name.</p>			

## Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties (regardless of the value of the **type** property):

Table 19. zBX object: class specific properties

Name	Qualifier	Type	Description
<b>type</b>	—	String Enum	The type of the zBX: <ul style="list-style-type: none"> <li>• "cpc-attached"</li> <li>• "node"</li> </ul>
<b>machine-type</b>	—	String (4)	The type of the zBX. For example, 2458.
<b>machine-model</b>	—	String (3)	The model of the zBX. For example, 004.
<b>machine-serial</b>	—	String (12)	The serial number of the zBX. For example, 000002001234.
<b>current-isaopt-entitlements</b>	—	Integer	Number of blades entitled as ISAOPT blades.
<b>max-isaopt-entitlements</b>	—	Integer	Maximum licensed to be entitled as ISAOPT blades.
<b>current-power-entitlements</b>	—	Integer	Number of blades entitled as POWER7 <sup>®</sup> blades.
<b>max-power-entitlements</b>	—	Integer	Maximum licensed to be entitled as POWER7 blades.
<b>current-systemx-entitlements</b>	—	Integer	Number of blades entitled as System x IBM blades.
<b>max-systemx-entitlements</b>	—	Integer	Maximum licensed to be entitled as System x blades.
<b>current-dpxi50z-entitlements</b>	—	Integer	Number of blades entitled as DataPower XI50z blades.
<b>max-dpxi50z-entitlements</b>	—	Integer	Maximum licensed to be entitled as DataPower XI50z blades.

In addition to the properties defined in the previous table, if the value of the **type** property of the zBX is "node", the zBX object also includes the following additional class-specific properties:

Table 20. zBX object: class specific additional properties (for zBX node)

Name	Qualifier	Type	Description
<b>se-version</b>	(pc)	(String 1-8)	The current release level of the primary SE licensed internal code. For example, "2.13.0". Note that the alternate SE is usually at the same level, except when installing new licensed internal code levels.
<b>has-hardware-messages</b>	(pc)	Boolean	The zBX node object has hardware messages (true), or does not have hardware messages (false).
<b>is-ensemble-member</b>	(pc)	Boolean	Whether the zBX node is currently part of an ensemble (true) or not (false).

Table 20. zBX object: class specific additional properties (for zBX node) (continued)

Name	Qualifier	Type	Description
is-service-required	—	Boolean	Whether the zBX node is operating using the last redundant part of a particular type (true) or not (false). If true, repairs should be made before additional parts fail that would make this zBX node non-operational.
ec-mcl-description	—	ec-mcl-description object	Describes the Engineering Change (EC) and MicroCode Level (MCL) for the zBX node. Refer to the description of ec-mcl-description for details.  It may take a significant amount of time for the HMC API to provide this information when requested. Because of this, and because of the specialized nature of this information, the EC/MCL information is omitted from the response to <b>Get zBX Properties</b> and from inventory data for the zBX, and instead the value of this property is null for those operations. Clients can use the <b>Get EC/MCL Description of zBX (Node)</b> operation to obtain the value of this property when required.
has-automatic-se-switch-enabled	—	Boolean	Automatic switching between primary and alternate Support Elements is enabled for the zBX node object (true), or is not enabled (false).  Support Element (Version 2.12.1 and newer) information can be found on the console help system, or on the IBM Knowledge Center at <a href="https://www.ibm.com/support/knowledgecenter">https://www.ibm.com/support/knowledgecenter</a> . (Select <b>z Systems</b> on the navigation bar, and then select your server). For information about earlier versions of the Support Element, see the <i>Support Element Operations Guide</i> .
lan-interface1-type	(pc)	String Enum	The adapter type of the Support Element's LAN interface 1. One of the following: <ul style="list-style-type: none"> <li>• "ethernet"</li> <li>• "token-ring"</li> <li>• "unknown"</li> </ul>
lan-interface1-address	(pc)	String (12)	The MAC address of the Support Element's LAN interface 1. The address is provided without any embedded colon separators (example: "f0def14b63af").
lan-interface2-type	(pc)	String Enum	The adapter type of the Support Element's LAN interface 2. One of the following: <ul style="list-style-type: none"> <li>• "ethernet"</li> <li>• "token-ring"</li> <li>• "unknown"</li> </ul>
lan-interface2-address	(pc)	String (12)	The MAC address of the Support Element's LAN interface 2. The address is provided without any embedded colon separators (example: "f0def14b63af").
network1-ipv4-mask	(pc)	String (1-15)	The network IP mask value.
network1-ipv4-pri-ipaddr	(pc)	String IPV4 address	The primary IPv4 address or a null object if not configured.
network1-ipv4-alt-ipaddr	(pc)	String IPV4 address	The alternate IPv4 address or a null object if not configured.
network1-ipv6-info	—	Array of ipv6-info objects	A list of objects describing the Support Element's IPv6 network connections. If no IPv6 connections are defined, an empty list is returned.
network2-ipv4-mask	(pc)	String (1-15)	The network IP mask value.

Table 20. zBX object: class specific additional properties (for zBX node) (continued)

Name	Qualifier	Type	Description
<b>network2-ipv4-pri-ipaddr</b>	(pc)	String IPV4 address	The primary IPv4 address or a null object if not configured.
<b>network2-ipv4-alt-ipaddr</b>	(pc)	String IPV4 address	The alternate IPv4 address or a null object if not configured.
<b>network2-ipv6-info</b>	—	Array of ipv6-info objects	A list of objects describing the Support Element's IPv6 network connections. If no IPv6 connections are defined, an empty list is returned.
<b>management-enablement-level</b>	(pc)	String Enum	<p>The zManager management enablement level for this zBX node. The <b>management-enablement-level</b> values of an ensemble's CPC and zBX node members of the ensemble determine the ensemble's <b>management-enablement-level</b>, which determines the zManager advanced management functions that are available for use.</p> <p>All ensemble nodes must have the same <b>management-enablement-level</b>. A CPC or zBX node with a <b>management-enablement-level</b> of "automate" may be added to an ensemble with a <b>management-enablement-level</b> of "manage", but this will downgrade the system's <b>management-enablement-level</b> until all ensemble nodes are upgraded to "automate".</p> <p>Values:</p> <ul style="list-style-type: none"> <li>• <b>"manage"</b> - Provides the basic capabilities for managing an ensemble. It includes HMC operational controls for change management of firmware across the ensemble, energy monitoring, virtual networking with automated provisioning, virtual server management, and a base level of performance management.</li> <li>• <b>"automate"</b> - Provides more leverage from the ensemble by managing workloads and energy. This level of support includes power capping, power savings mode, and platform performance management.</li> </ul>
<b>hardware-messages</b>	(c)(pc)	Array of hardware-message objects	<p>The complete list of all zBX node hardware messages, each identified by its URI. This list corresponds to the list provided by the <b>List zBX (Node) Hardware Messages</b> operation. If the zBX node has no hardware messages, then an empty array is provided.</p> <p>The list of returned hardware messages can change as a result of the new messages being dynamically added or removed by the infrastructure or due to hardware messages being deleted via the <b>Delete zBX (Node) Hardware Message</b> operation.</p> <p><b>Note:</b> This property is not returned by the <b>Get zBX Properties</b> operation or included in Inventory Service results for the zBX object, and only sessions associated with an HMC user with permission to the Hardware Messages task will receive a property-change notification for this property.</p>

Each ec-mcl-description nested object and its related action, ec and mcl nested objects have the properties shown in the following tables:

Table 21. zBX object: ec-mcl-description nested object properties (for zBX node)

Name	Type	Description
action	Array of action objects	An optional array of pending action objects. This field is only provided when the HMC is communicating with the zBX node's SE.
ec	Array of ec objects	An optional array of EC objects. This field is only provided when the HMC is communicating with the zBX node's SE.

Table 22. zBX object: action nested object properties (for zBX node)

Name	Type	Description
type	String Enum	Valid value is "zhybrid-blades-activation" - zHybrid accelerator blades are pending an activation.
activation	String Enum	One of: <ul style="list-style-type: none"> <li>"current" - The action is for the current activation.</li> <li>"next" - The action is for the next install and activation.</li> </ul>
pending	Boolean	Indicates whether the action is pending (true) or not pending (false).

Table 23. zBX object: ec nested object properties (for zBX node)

Name	Type	Description
number	String (1-6)	Engineering Change stream identifier.
part-number	String (1-8)	Engineering Change stream part number..
type	String (1-32)	Engineering Change stream name.
description	String (1-65)	Engineering Change stream descriptive text.
mcl	Array of mcl objects	The list of MicroCode Levels for this Engineering Change.

Table 24. zBX object:mcl nested object properties (for zBX node)

Name	Type	Description
type	String Enum	One of: <ul style="list-style-type: none"> <li>"retrieved" - A retrieved or staged level.</li> <li>"activated" - An activated or applied level.</li> <li>"accepted" - A committed level.</li> <li>"installable-concurrent" - A non-disruptive apply-able level.</li> <li>"removable-concurrent" - A non-disruptive reject-able level.</li> </ul>
level	String (1-3)	MicroCode Level.
last-update	Timestamp	Time stamp of the last update, in the number of milliseconds since midnight January 1, 1970 UTC. A null object is returned if no updates have occurred.

Each ipv6-info nested object has the following properties:

Table 25. zBX object: ipv6-info properties (for zBX node)

Name	Type	Description
type	String Enum	The IPv6 scope. One of the following values: <ul style="list-style-type: none"> <li>"link-local"</li> <li>"static"</li> <li>"auto"</li> </ul>
prefix	Integer	The number of leading bits of the IPv6 address that represent the network prefix.

Table 25. zBX object: ipv6-info properties (for zBX node) (continued)

Name	Type	Description
<b>pri-ip-address</b>	String IPv6 address	The primary IPv6 address.
<b>alt-ip-address</b>	String IPv6 address	The alternate IPv6 address or a null object if not configured.

Each hardware-message nested object has the following properties:

Table 26. zBX object: hardware-message object properties (for zBX node)

Name	Type	Description
<b>element-uri</b>	String/URI	The canonical URI path of the zBX node hardware message. The URI is in the following form: <code>/api/zbx/s/{zbx-id}/hardware-messages/{hardware-message-id}</code> , where <code>{hardware-message-id}</code> is the value of the <b>element-id</b> property of the hardware message.
<b>element-id</b>	String (36)	The unique identifier for the hardware message. The <b>element-id</b> is in the form of a UUID.
<b>parent</b>	String/URI	The parent of a zBX (Node) hardware message is the zBX Node object. The <b>parent</b> value is the canonical URI path for the zBX (Node).
<b>class</b>	String	The <b>class</b> of a hardware-message object is <b>"hardware-message"</b> .
<b>timestamp</b>	Timestamp	The <b>timestamp</b> represents the date and time when the hardware message was created.
<b>text</b>	String	The text of the hardware message.

## Energy Management Related Additional Properties

In addition to the class-specific additional properties defined in the previous section, if the value of the **type** property is **"node"**, the zBX object also includes the following additional class-specific properties related to energy management. For further explanation of the various states involved, see Chapter 9, "Energy management," on page 179.

Table 27. zBX object: energy management related additional properties (for zBX node)

Name	Qualifier	Type	Description
<b>power-rating</b>	—	Integer	Specifies the maximum power usage of this zBX node in watts (W). This value is a calculated value, as indicated by the electrical rating labels or system rating plates of the zBX node components.
<b>power-consumption</b>	(mg)	Integer	Specifies the current power consumption in watts (W) of this zBX node. The zBX power consumption includes the power consumption of the blades contained within all BladeCenter chassis and the shared infrastructure.

Table 27. zBX object: energy management related additional properties (for zBX node) (continued)

Name	Qualifier	Type	Description
<b>power-saving</b>	—	String Enum	<p>Specifies the current power saving setting of the zBX node. Power saving reduces the energy consumption of a system, and can be managed it using the Set Power Saving operation. The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - Specifies not reducing the power consumption and performance of the zBX node. This is the default setting.</li> <li>• <b>"low-power"</b> - Specifies low power consumption for all components of the zBX node enabled for power saving.</li> <li>• <b>"custom"</b> - Specifies that some, but not all, components of the zBX node are in the Low power setting.</li> <li>• <b>"not-supported"</b> - Specifies that power saving is not supported for this zBX node.</li> <li>• <b>"not-available"</b> - Specifies that <b>power-saving</b> property could not be read from this zBX node.</li> <li>• <b>"not-entitled"</b> - Specifies that the server is not entitled to power saving.</li> </ul>
<b>power-saving-state</b>	—	String Enum	<p>Specifies whether power saving is supported and entitled for the zBX node set by the user. Note that this property indicates the user setting and may not match the real state of the hardware compared to the <b>power-saving</b> property. For more information, see "Group power saving" on page 182. The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - Specifies not reducing the power consumption and performance of the zBX node. This setting will be applied to all children.</li> <li>• <b>"low-power"</b> - Specifies low power consumption for all components of the zBX node enabled for power saving.</li> <li>• <b>"custom"</b> - Specifies that the zBX node does not control the children. This is the default setting.</li> <li>• <b>"not-supported"</b> - Specifies that power saving is not supported for this zBX node.</li> <li>• <b>"not-entitled"</b> - Specifies that the server is not entitled to power saving.</li> </ul>
<b>power-save-allowed</b>	—	String Enum	<p>Should be used to determine if a call of the power save operation is currently allowed. If a value other than <b>"allowed"</b> is returned the caller may reckon that the power save operation will fail.</p> <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Alter power save setting is allowed for this zBX node</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power saving.</li> <li>• <b>"not-supported"</b> - Specifies that power saving is not supported for this zBX node.</li> </ul>

Table 27. zBX object: energy management related additional properties (for zBX node) (continued)

Name	Qualifier	Type	Description
<b>power-capping-state</b>	—	String Enum	Specifies the current power capping setting of the zBX node. Power capping limits peak power consumption of a system, and you can manage it by using the Set Power Cap operation. The possible settings include: <ul style="list-style-type: none"> <li>• <b>"disabled"</b> - Specifies not setting the power cap of the zBX node and not limiting the peak power consumption. This is the default setting.</li> <li>• <b>"enabled"</b> - Specifies capping all components of the zBX node available for power capping to limit the peak power consumption of the zBX node.</li> <li>• <b>"custom"</b> - The components of the zBX node can be individually configured for power capping.</li> <li>• <b>"not-supported"</b> - Specifies that power capping is not supported for this zBX node.</li> <li>• <b>"not-entitled"</b> - Specifies that the server is not entitled to power capping.</li> </ul>
<b>power-cap-minimum</b>	—	Integer	Specifies the minimum value for the zBX node cap value in watts (W). This is a sum of the component minimum cap values.
<b>power-cap-maximum</b>	—	Integer	Specifies the maximum value for the zBX node cap value in watts (W). This is a sum of the component maximum cap values.
<b>power-cap-current</b>	—	Integer	Specifies the current cap value for the zBX node in watts (W). The current cap value indicates the power budget for the zBX node and is the sum of the component Cap values.
<b>power-cap-allowed</b>	—	String Enum	Should be used to determine if a call of the power capping operation is currently allowed. If a value other than <b>"allowed"</b> is returned the caller may reckon that the power capping operation will fail. <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Alter power capping setting is allowed for this zBX node</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power capping.</li> <li>• <b>"not-supported"</b> - Specifies that power capping is not supported for this zBX node.</li> </ul>
<b>ambient-temperature</b>	(mg)	Float	Specifies the input air temperature in degrees Celsius (°C) as measured by the system.
<b>exhaust-temperature</b>	—	Float	Specifies the exhaust air temperature in degrees Celsius (°C) as calculated by the system. This is useful in determining potential hot spots in the data center.

## Operations

### List zBXs of a CPC

The **List zBXs of a CPC** operation lists the CPC-attached zBXs associated with the CPC. The CPC must be a member of an ensemble.

## HTTP method and URI

GET /api/cpcs/{cpc-id}/zbx

In this request, the URI variable *{cpc-id}* is the object ID of the CPC object whose zBXs are to be obtained.

### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) used to limit returned objects

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
zbx	Array of objects	Array of nested zbx-info objects, described in the next table. If no zBXs are returned, an empty array is provided.

Each nested zbx-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the zBX object in the form /api/zbx/{zbx-id}.
name	String	The <b>name</b> property of the zBX object (based on its machine type, machine model, and machine serial number).

## Description

- The **List zBXs of a CPC** operation lists the CPC-attached zBXs that are associated with this CPC. The object URI and name are provided for each zBX.

If the **name** query parameter is specified, then a zBX is included in the list only if the name pattern matches the **name** property of the object.

A zBX is included in the list only if the API user has object-access permission for the CPC with which the zBX is associated. If the HMC is a manager of a zBX, but the API user does not have permission to it, that object is omitted from the list, but no error status code results.

If the CPC does not have a zBX, an empty list is provided and the operation completes successfully. Note that if the CPC is not a member of an ensemble, it cannot have a zBX. Therefore, an empty list is provided.

## Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the CPC object specified by the request URI.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."



Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	241	CPC object ( <i>{cpc-id}</i> ) is not a member of an ensemble.
404 (Not Found)	1	The object ID <i>{cpc-id}</i> does not designate an existing CPC object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- This operation has been structured to allow for the possibility that multiple CPC-attached zBXs might be associated with a CPC. However, in the current implementation, each CPC can have at most one zBX. Therefore, the resulting list of zBXs will contain either zero or one entry.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/zbx HTTP/1.1
x-api-session: 5q07hx6jgp2ngn2cypq1zxot76sfwnzky0ih8nddd5hz6bpiue
```

---

Figure 22. List zBXs of a CPC: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 21 Jul 2011 17:48:58 GMT
content-type: application/json;charset=UTF-8
content-length: 160
{
  "zbx": [
    {
      "name": "2458-002-0000000ZBX26",
      "object-uri": "/api/zbx/54a9716c-a326-11e0-9469-001f163805d8"
    }
  ]
}
```

---

Figure 23. List zBXs of a CPC: Response

## List zBXs of an Ensemble

The **List zBXs of an Ensemble** operation lists the zBXs, which are associated with each CPC in the ensemble.

### HTTP method and URI

```
GET /api/ensembles/{ensemble-id}/zbx
```

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble object whose zBXs are to be obtained.

**Query parameters:**

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) used to limit returned objects
type	String Enum	Optional	Filter string used to limit returned objects to those that have a matching <b>type</b> property. Value must be a valid zBX <b>type</b> property value.

**Response body contents**

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
zbx	Array of objects	Array of nested zbx-info objects, described in the next table. If the Ensemble does not have any CPCs or zBXs associated with it, an empty array is provided.

Each nested zbx-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the zBX object in the form <code>/api/zbx/{zbx-id}</code> .
name	String	The <b>name</b> property of the zBX object.
type	String Enum	Type of the zBX object.

**Description**

The **List zBXs of an Ensemble** operation lists the zBXs that are part of the ensemble. The object URI, name, and type are provided for each zBX.

If the **name** query parameter is specified, then a zBX is included in the list only if the name pattern matches the **name** property of the object.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid zBX **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those zBXs that have a **type** property matching a specified type value. If the **type** parameter is omitted, this filtering is not done.

A CPC-attached zBX is included in the list only if the API user has object-access permission for the CPC object with which it is associated. A zBX node is included in the list only if the API user has object-access permission to that object. If the HMC is a manager of a zBX, but the API user does not have this access permission to it, that object is omitted from the list, but no error status code results.

If the HMC does not manage any zBXs, an empty list is provided and the operation completes successfully.

**Authorization requirements**

This operation has the following authorization requirements:

- Object-access permission to the Ensemble object passed in the request URI

- | • Object-access permission to the CPC objects with which a zBX is associated (for CPC-attached zBX objects).
- | • Object-access permission to the zBX object (for zBX node objects).

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 90.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{ensemble-id}</i> does not designate an existing ensemble object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/zbx HTTP/1.1
x-api-session: 5q07hx6jgp2ngn2cypq1zxot76sfwnzky0ih8nddd5hz6bpiue
```

---

Figure 24. List zBXs of an Ensemble: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 21 Jul 2011 17:48:58 GMT
content-type: application/json;charset=UTF-8
content-length: 160
{
  "zbx": [
    {
      "name": "2458-002-0000000ZBX26",
      "object-uri": "/api/zbx/54a9716c-a326-11e0-9469-001f163805d8"
    }
  ]
}
```

---

Figure 25. List zBXs of an Ensemble: Response

## Get zBX Properties

The **Get zBX Properties** operation retrieves the properties of a single zBX object that is designated by its object ID.

### HTTP method and URI

```
GET /api/zbx/{zbx-id}
```

In this request, the URI variable *{zbx-id}* is the object ID of the zBX object for which properties are to be obtained.

## Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the zBX object as defined in the “Data model” on page 80. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

## Description

The **Get zBX Properties** operation returns the current properties for the zBX object specified by *{zbx-id}*.

- | On successful execution, all of the current properties as defined in “Data model” on page 80 for the zBX object are provided in the response body, and HTTP status code 200 (OK) is returned.
- | The URI path must designate an existing zBX object. If the URI path designates a CPC-attached zBX, the API user must have object-access permission to the CPC with which the zBX is associated. If the URI path designates a zBX node, the API user must have object-access permission to that object. If these conditions are not met, status code 404 (Not Found) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- | • Object-access permission to the CPC objects with which the zBX is associated (for CPC-attached zBX objects).
- | • Object-access permission to the zBX object (for zBX node objects).

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{zbx-id}</i> does not designate an existing zBX object, or the API user does not have object access permission to the CPC with which the zBX is associated.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
| GET /api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408 HTTP/1.1  
| x-api-session: 65815bpxckw5ejih51h67bt5zygjwrh8v34u33s9o5dg4atnnz
```

---

*Figure 26. Get zBX Properties: Request*

```
| 200 OK
| server: zSeries management console API web server / 2.0
| cache-control: no-cache
| date: Mon, 08 Sep 2014 17:28:10 GMT
| content-type: application/json;charset=UTF-8
| content-length: 2066
| {
|   "acceptable-status": [
|     "operating"
|   ],
|   "class": "zbx",
|   "current-dpxi50z-entitlements": 0,
|   "current-isaopt-entitlements": 0,
|   "current-power-entitlements": 0,
|   "current-systemx-entitlements": 0,
|   "description": "Defined zBX Node",
|   "ec-mcl-description": null,
|   "has-automatic-se-switch-enabled": false,
|   "has-hardware-messages": true,
|   "has-unacceptable-status": true,
|   "is-ensemble-member": true,
|   "is-service-required": false,
|   "lan-interface1-address": "40f2e910664e",
|   "lan-interface1-type": "ethernet",
|   "lan-interface2-address": "40f2e910664f",
|   "lan-interface2-type": "ethernet",
|   "machine-model": "004",
|   "machine-serial": "000002560057",
|   "machine-type": "2458",
|   "management-enablement-level": "manage",
|   "max-dpxi50z-entitlements": 7,
|   "max-isaopt-entitlements": 0,
|   "max-power-entitlements": 0,
|   "max-systemx-entitlements": 0,
|   "name": "P2560057",
|   "network1-ipv4-alt-ipaddr": "9.60.14.99",
|   "network1-ipv4-mask": "255.255.255.0",
|   "network1-ipv4-pri-ipaddr": "9.60.14.98",
|   "network1-ipv6-info": [
|     {
|       "alt-ip-address": "2002:93c:ffb:1:42f2:e9ff:fe10:708f",
|       "prefix": 64,
|       "pri-ip-address": "2002:93c:ffb:1:42f2:e9ff:fe10:664f",
|       "type": "auto"
|     },
|     {
|       "alt-ip-address": "fdd8:673b:d89b:1:42f2:e9ff:fe10:708f",
|       "prefix": 64,
|       "pri-ip-address": "fdd8:673b:d89b:1:42f2:e9ff:fe10:664f",
|       "type": "auto"
|     },
|     {
|       "alt-ip-address": "fe80::42f2:e9ff:fe10:708f%eth0",
|       "prefix": 64,
|       "pri-ip-address": "fe80::42f2:e9ff:fe10:664f%eth0",
|       "type": "link-local"
|     }
|   ]
| },
| ]
```

Figure 27. Get zBX Properties: Response (Part 1)

```

"network2-ipv4-alt-ipaddr": "9.60.15.99",
"network2-ipv4-mask": "255.255.255.0",
"network2-ipv4-pri-ipaddr": "9.60.15.98",
"network2-ipv6-info": [
  {
    "alt-ip-address": "fe80::42f2:e9ff:fe10:708e%eth1",
    "prefix": 64,
    "pri-ip-address": "fe80::42f2:e9ff:fe10:664e%eth1",
    "type": "link-local"
  }
],
"object-id": "ace3f861-3421-33d5-bcf4-f61f9057d408",
"object-uri": "/api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408",
"parent": "/api/ensembles/2d6677f8-4b9e-11e3-96d7-42f2e90d2883",
"power-cap-allowed": "not-entitled",
"power-capping-state": "disabled",
"power-consumption": null,
"power-rating": null,
"power-save-allowed": "not-entitled",
"power-saving": "not-entitled",
"power-saving-state": "not-entitled",
"se-version": "2.13.0",
"status": "service",
"type": "node"
}

```

Figure 28. Get zBX Properties: Response (Part 2)

## Get EC/MCL Description of zBX (Node)

The **Get EC/MCL Description of zBX (Node)** operation retrieves the value of the **ec-mcl-description** property of a zBX node that is designated by its object ID.

### HTTP method and URI

**GET** `/api/zbx/{zbx-id}/operations/get-ec-mcl-description`

In this request, the URI variable `{zbx-id}` is the object ID of the zBX node for which EC/MCL information is to be retrieved.

### Response body contents

On successful completion, the response body is a JSON object that provides the current value of the **ec-mcl-description** property for the zBX object. The response body object contains a single field named **ec-mcl-description** with a value that is in the format described for the same-named property in the “Data model” on page 80.

### Description

The **Get EC/MCL Description of zBX (Node)** operation returns the current value of the **ec-mcl-description** property for the zBX node object specified by `{zbx-id}`.

On successful execution, the value of the property as defined in “Data model” on page 80 is provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing zBX node object and the API user must have object-access permission to that object.

### Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the zBX node object designated in the URI.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 95.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

Table 28. Get EC/MCL Description of zBX (Node): HTTP status and reason codes.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The object ID <i>{zbx-id}</i> designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of “node”).
404 (Not Found)	1	The object ID <i>{zbx-id}</i> does not designate an existing zBX node, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/zbxes/8dac6a58-935e-352c-996e-ded17dbf92c0/operations/get-ec-mcl-description HTTP/1.1
x-api-session: 470bbm7331p8nkpu5x1qasezo8apxyvclifznc1ktr07jej8x1
```

---

Figure 29. Get EC/MCL Description of zBX (Node): Request



```

| 200 OK
| server: zSeries management console API web server / 2.0
| cache-control: no-cache
| date: Thu, 02 Oct 2014 16:19:25 GMT
| content-type: application/json;charset=UTF-8
| content-length: 2473
| {
|   "ec-mcl-description": {
|     "action": [
|       {
|         "activation": "current",
|         "pending": true,
|         "type": "zhybrid-blades-activation"
|       },
|       {
|         "activation": "next",
|         "pending": false,
|         "type": "zhybrid-blades-activation"
|       }
|     ],
|     "ec": [
|       {
|         "description": "SE Framework",
|         "mcl": [
|           {
|             "last-update": null,
|             "level": "000",
|             "type": "retrieved"
|           },
|           {
|             "last-update": null,
|             "level": "000",
|             "type": "activated"
|           },
|           {
|             "last-update": null,
|             "level": "000",
|             "type": "accepted"
|           },
|           {
|             "last-update": null,
|             "level": "000",
|             "type": "installable-concurrent"
|           },
|           {
|             "last-update": null,
|             "level": "000",
|             "type": "removable-concurrent"
|           }
|         ],
|         "number": "N98205",
|         "part-number": "00LY716",
|         "type": "Base EC"
|       }
|     ],
|   },

```

Figure 30. Get EC/MCL Description of zBX (Node): Response (Part 1)

```

|         {
|           "description": "BladeCenter Components",
|           "mcl": [
|             {
|               "last-update": null,
|               "level": "000",
|               "type": "retrieved"
|             },
|             {
|               "last-update": null,
|               "level": "000",
|               "type": "activated"
|             },
|             {
|               "last-update": null,
|               "level": "000",
|               "type": "accepted"
|             },
|             {
|               "last-update": null,
|               "level": "000",
|               "type": "installable-concurrent"
|             },
|             {
|               "last-update": null,
|               "level": "000",
|               "type": "removable-concurrent"
|             }
|           ],
|           "number": "N98209",
|           "part-number": "00LY720",
|           "type": "Other Optional EC"
|         }
|       ]
|     }
|   }

```

Figure 31. Get EC/MCL Description of zBX (Node): Response (Part 2)

## Activate zBX (Node)

The **Activate zBX (Node)** operation activates all blades in all BladeCenter chassis in the zBX node designated by its object ID. This operation is asynchronous.

### HTTP method and URI

**POST** /api/zbx/{zbx-id}/operations/activate

In this request, the URI variable {zbx-id} is the object ID of the zBX node object to activate.

### Response body contents

Once the activation request is accepted, the response body contains a JSON object with the following field:

Field name	Type	Description
job-uri	String/ URI	URI that may be queried to retrieve activation status updates. The canonical URI path for zBX activation status updates is of the form /api/jobs/{job-id}.

## | Asynchronous result description

| Once the activation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Activate zBX (Node)** request.

| The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. (For more information, see “Query Job Status” on page 52. When the status of the job is **"complete"**, the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) that are set. (See the description that follows.) The **job-results** field is null for asynchronous zBX activation jobs.

## | Description

| The **Activate zBX (Node)** operation activates all blades in all chassis in the zBX node object specified by *{zbx-id}*. Activation brings blades into a state of **"operating"**. If a blade is already powered on when the activation operation is requested, the blade is powered off and then brought to a state of **"operating"**. The order in which the blades of the zBX are activated is unspecified. If the blade is a host to a virtualization application, then this application is activated also. See “Activating a Virtualization Host” on page 256 for more information.

| The URI path must designate an existing zBX object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. The zBX designated by the URI path must be a zBX node (that is, have a **type** of **"node"**) otherwise status code 400 (Bad Request) is returned. In addition to having object-access permission to the zBX, the API user must also have permission to the **Activate** task, otherwise status code 403 (Forbidden) is returned. The zBX must be in the correct state to perform the activate action, otherwise status code 409 (Conflict) is returned.

| When the zBX node activation job is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the activation job. See “Query Job Status” on page 52 for information on how to query job status. When the activate job has completed, an asynchronous result message is sent, with “Job status and reason codes” on page 100.

## | Authorization requirements

| This operation has the following authorization requirement:

- | • Object-access permission to the zBX node object designated by *{zbx-id}* and all its BladeCenter chassis and blades.
- | • Object-access permission to the **Activate** task.

## | HTTP status and reason codes

| On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 98.

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 29. Activate zBX (Node): HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The object ID <i>{zbx-id}</i> designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of “node”).
403 (Forbidden)	0	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID <i>{zbx-id}</i> does not designate an existing zBX object, or the API user does not have object access permission to it.
409 (Conflict)	1	The object ID <i>{zbx-id}</i> designates a zBX object that is not in the correct state (status) for performing the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Table 30. Activate zBX (Node): Job status and reason codes

HTTP error status code	Reason code	Description
200 (OK)	N/A	Activation completed successfully.
500 (Server Error)	100	zBX activation failed.
	101	The zBX activation job timed out.

## Example HTTP interaction

---

```
POST /api/zbx/5e5b9b5e-3c76-3fac-8f82-d9b0055773fc/operations/activate HTTP/1.1
x-api-session: 1gtnlpz3uk1d3es07mdnbsnsnrgss2xupmai32nkgvaysm5434
```

---

Figure 32. Activate zBX (Node): Request

---

```
202 Accepted
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Wed, 10 Sep 2014 18:14:40 GMT
content-type: application/json;charset=UTF-8
content-length: 60
{
  "job-uri": "/api/jobs/54e8e598-3916-11e4-8e38-02c00001303e"
}
```

---

Figure 33. Activate zBX (Node): Response

## Deactivate zBX (Node)

The **Deactivate zBX (Node)** operation deactivates all blades in all BladeCenter chassis in the zBX node specified by its object ID. This operation is asynchronous.

## HTTP method and URI

**POST** `/api/zbxes/{zbx-id}/operations/deactivate`

In this request, the URI variable `{zbx-id}` is the object ID of the zBX node object to deactivate.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String/ URI	URI that may be queried to retrieve deactivation status updates. The canonical URI path for zBX deactivation status updates is of the form <code>/api/jobs/{job-id}</code> .

## Asynchronous result description

Once the deactivation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Deactivate zBX (Node)** request.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. (For more information, see “Query Job Status” on page 52. When the status of the job is **"complete"**, the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) that are set. (See the description that follows.) The **job-results** field is null for asynchronous zBX deactivation jobs.

## Description

The **Deactivate zBX (Node)** operation deactivates all blades in all BladeCenter chassis in the zBX node object specified by `{zbx-id}`. Deactivation powers off all blades after an orderly shutdown of any hardware and software activity running on the blades. The order in which the blades of the zBX are deactivated is unspecified. If the blade is a host to a virtualization application, then this application is deactivated also. See “Deactivating a Virtualization Host” on page 257 for more information.

The URI path must designate an existing zBX object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. The zBX designated by the URI path must be a zBX node (that is, have a **type** of **"node"**) otherwise status code 400 (Bad Request) is returned. In addition to having object-access permission to the zBX, the API user must also have permission to the **Deactivate** task, otherwise status code 403 (Forbidden) is returned. The zBX must be in the correct state to perform the deactivate action, otherwise status code 409 (Conflict) is returned.

When the zBX node deactivation job is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the deactivation job. See “Query Job Status” on page 52 for information on how to query job status. When the deactivate job has completed, an asynchronous result message is sent, with “Job status and reason codes” on page 102.

## Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the zBX node object specified by `{zbx-id}` and all its BladeCenter chassis and blades.
- Object-access permission to the **Deactivate** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 101.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 31. Deactivate zBX (Node): HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The object ID <i>{zbx-id}</i> designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of “node”).
403 (Forbidden)	0	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID <i>{zbx-id}</i> does not designate an existing zBX object, or the API user does not have object access permission to it.
409 (Conflict)	1	The object ID <i>{zbx-id}</i> designates a zBX object that is not in the correct state (status) for performing the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Table 32. Deactivate zBX (Node): Job status and reason codes

HTTP error status code	Reason code	Description
200 (OK)	N/A	Deactivation completed successfully.
500 (Server Error)	100	zBX deactivation failed.
	101	The zBX deactivation job timed out.

## Example HTTP interaction

---

```
POST /api/zbx/5e5b9b5e-3c76-3fac-8f82-d9b0055773fc/operations/deactivate HTTP/1.1
x-api-session: 5upqikkhel1vhasloh25bnhelekvemvnmhutc4kfegav9571zx
```

---

Figure 34. Deactivate zBX (Node): Request

```

| 202 Accepted
| server: zSeries management console API web server / 2.0
| cache-control: no-cache
| date: Wed, 10 Sep 2014 18:12:50 GMT
| content-type: application/json;charset=UTF-8
| content-length: 60
| {
|   "job-uri": "/api/jobs/13308ef8-3916-11e4-bb52-02c00001303e"
| }

```

Figure 35. Deactivate zBX (Node): Response

## Get zBX (Node) Audit Log

The **Get zBX (Node) Audit Log** operation returns the zBX node's audit log, filtered according to the query parameters, if specified.

### HTTP method and URI

**GET** /api/zbxes/{zbx-id}/operations/get-audit-log

In this request, the URI variable {zbx-id} is the object ID of the target zBX node object.

### Query parameters:

Name	Type	Req/Opt	Description
begin-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created earlier than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.
end-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created later than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.

### Response body contents

On successful completion, the response body is a JSON array of JSON objects returned using HTTP chunked transfer encoding. Each array element is a log-entry-info object containing information about a single log entry. The array elements are in order of increasing timestamp. See Table 186 on page 631 for more information.

### Description

This operation returns the zBX node's audit log in increasing timestamp order, filtered according to the query parameters, if specified. Each log entry pertains to a specific event that occurred on or to a managed object or the zBX node itself. The log entries can be limited by specifying explicit filtering criteria on the request. If the **begin-time** query parameter is specified, then any entries earlier than that time are omitted. If the **end-time** query parameter is specified, then any entries later than that time are omitted.

The URI path must designate an existing zBX object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. The zBX must be a zBX node (that is, have a **type** of "node"); otherwise, status code 400 (Bad Request) is returned. In addition, the API user must have action/task permission to the **Audit and Log Management** task; otherwise, status code 403 (Forbidden) is returned.

On successful execution, the response body contains an array of filtered log entries. If the audit log is empty or there are no entries to be returned after filtering, then an empty array is provided. Each log entry contains the event ID, event name and event message. If there are data items included in the event message, they are available separately. The order and meaning of the substitution items for each event ID are documented in the IBM Knowledge Center, at <http://www.ibm.com/support/knowledgecenter/> (Select **z Systems** on the navigation bar, and then select your server). The information can be found in the HMC Introduction topic **Audit, Event, and Security Log Messages**.

### Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the zBX node object specified in the request URI.
- Action/task permission to the **Audit and Log Management** task.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 103.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

Table 33. Get zBX (Node) Audit Log: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The request URI designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of “ <b>node</b> ”).
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.
	4	The zBX node designated by the request URI does not support this operation.
500 (Server Error)	307	The request timed out while attempting to communicate with the SE or while attempting to get the log data.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
GET /api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408/operations/get-audit-log
  HTTP/1.1
x-api-session: bp66qdtmjnofktpyw0f6hvsipo7skdv22v1aov70ecngc2kvn
```

Figure 36. Get zBX (Node) Audit Log: Request



```

200 OK
server: zSeries management console API web server / 2.0
transfer-encoding: chunked
cache-control: no-cache
date: Fri, 03 Oct 2014 01:19:54 GMT
content-type: application/json;charset=ISO-8859-1
[
  {
    "event-data-items": [
      {
        "data-item-number": 0,
        "data-item-type": "string",
        "data-item-value": "B.2.01"
      }
    ],
    "event-details": [],
    "event-id": "3296",
    "event-message": "Stop of hypervisor B.2.01 was triggered.",
    "event-name": "HV-STOP",
    "event-time": 1411619111020,
    "user-uri": "/api/users/146c44a0-a3fd-11e3-bfca-df30ad17ee78",
    "userid": "HVADMIN"
  },
  {
    "event-data-items": [
      {
        "data-item-number": 0,
        "data-item-type": "string",
        "data-item-value": "B.2.01"
      }
    ],
    "event-details": [],
    "event-id": "3298",
    "event-message": "Stop of hypervisor B.2.01 was successful.",
    "event-name": "HV-STOP",
    "event-time": 1411619111970,
    "user-uri": "/api/users/146c44a0-a3fd-11e3-bfca-df30ad17ee78",
    "userid": "HVADMIN"
  }
]

```

Figure 37. Get zBX (Node) Audit Log: Response

## Get zBX (Node) Security Log

The **Get zBX (Node) Security Log** operation returns the zBX node's security log, filtered according to the query parameters, if specified.

### HTTP method and URI

**GET** /api/zbx/{zbx-id}/operations/get-security-log

In this request, the URI variable *{zbx-id}* is the object ID of the target zBX node object.

### Query parameters:

Name	Type	Req/Opt	Description
begin-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created earlier than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.

Name	Type	Req/Opt	Description
end-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created later than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.

## Response body contents

On successful completion, the response body is a JSON array of JSON objects returned using HTTP chunked transfer encoding. Each array element is a log-entry-info object containing information about a single log entry. The array elements are in order of increasing timestamp. See Table 186 on page 631 for more information.

## Description

This operation returns the zBX node's security log in increasing timestamp order, filtered according to the query parameters, if specified. Each log entry pertains to a specific event that occurred on or to a managed object or the zBX node itself. The log entries can be limited by specifying explicit filtering criteria on the request. If the **begin-time** query parameter is specified, then any entries earlier than that time are omitted. If the **end-time** query parameter is specified, then any entries later than that time are omitted.

The URI path must designate an existing zBX object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. The zBX must be a zBX node (that is, have a **type** of "node"); otherwise, status code 400 (Bad Request) is returned. In addition, the API user must have action/task permission to the **View Security Logs** task; otherwise, status code 403 (Forbidden) is returned.

On successful execution, the response body contains an array of filtered log entries. If the security log is empty or there are no entries to be returned after filtering, then an empty array is provided. Each log entry contains the event ID, event name and event message. If there are data items included in the event message, they are available separately. The order and meaning of the substitution items for each event ID are documented in the IBM Knowledge Center, at <http://www.ibm.com/support/knowledgecenter/> (Select **z Systems** on the navigation bar, and then select your server). The information can be found in the HMC Introduction topic **Audit, Event, and Security Log Messages**.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the zBX node object specified in the request URI.
- Action/task permission to the **View Security Logs** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

Table 34. Get zBX (Node) Security Log: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The request URI designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of “node”).
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.
	4	The zBX node designated by the request URI does not support this operation.
500 (Server Error)	307	The request timed out while attempting to communicate with the SE or while attempting to get the log data.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
GET /api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408/operations/get-security-log
HTTP/1.1
x-api-session: bp66qdtmjnofktpyw0f6hvsipo7skdv22v1aov70ecngc2kvn
```

Figure 38. Get zBX (Node) Security Log: Request

```

| 200 OK
| server: zSeries management console API web server / 2.0
| transfer-encoding: chunked
| cache-control: no-cache
| date: Fri, 03 Oct 2014 01:19:56 GMT
| content-type: application/json;charset=ISO-8859-1
| [
|   {
|     "event-data-items": [],
|     "event-details": [],
|     "event-id": "778",
|     "event-message": "Mirroring data from the primary Support Element to the
|       alternate Support Element started.",
|     "event-name": "ASEMIRRST",
|     "event-time": 1412265359190,
|     "user-uri": null,
|     "userid": null
|   },
|   {
|     "event-data-items": [
|       {
|         "data-item-number": 0,
|         "data-item-type": "string",
|         "data-item-value": "Service Status is active."
|       }
|     ],
|     "event-details": [],
|     "event-id": "779",
|     "event-message": "Mirroring data from the primary Support Element to the
|       alternate Support Element failed. Service Status is active.",
|     "event-name": "ASEMIRRNO",
|     "event-time": 1412265360460,
|     "user-uri": null,
|     "userid": null
|   }
| ]

```

Figure 39. Get zBX (Node) Security Log: Response

## List zBX (Node) Hardware Messages

The **List zBX (Node) Hardware Messages** operation lists the current set of hardware messages associated with the zBX (Node).

### HTTP method and URI

**GET** /api/zbx/{zbx-id}/hardware-messages

In this request, the URI variable *{zbx-id}* is the object ID of a zBX Node object for which hardware messages are to be listed.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
hardware-messages	Array of hardware-message-info objects	Array of nested hardware-message-info objects as defined in the next table.

Each nested hardware-message-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	The canonical URI path of the hardware message. The URI is in the following form: <code>/api/zbxs/{zbx-id}hardware-messages/{hardware-message-id}</code>
timestamp	Timestamp	The date and time the hardware message was created
text	String	The text of the hardware message.

## Description

This operation returns a set of zBX (Node) hardware messages in increasing timestamp order, filtered according to the query parameters, if specified. Each hardware message describes an event or notification that may require the operator's attention. The list of hardware messages can be limited by specifying explicit filtering criteria on the request.

If the **begin-time** query parameter is specified, then any entries earlier than that time are omitted. If the **end-time** query parameter is specified, then any entries later than that time are omitted.

If there are no hardware messages associated with the zBX (Node), or if no hardware messages are to be included in the results due to filtering, an empty array is returned and the operation completes successfully.

The URI path must designate an existing zBX (Node) and the API user must have object access permission to it; otherwise, status code 404 (Not Found) is returned. In addition, the API user must have Action/Task permission to the **Hardware Messages** task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the zBX Node object designated by `{zbx-id}`.
- Action/Task permission to the **Hardware Messages** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents" on page 108.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 35. List zBX (Node) Hardware Messages: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	7	The <b>begin-time</b> value is greater than the <b>end-time</b> value.
	240	The request URI designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of "node").
403 (Forbidden)	1	The API user does not have Action/Task permission for the <b>Hardware Messages</b> task.

Table 35. List zBX (Node) Hardware Messages: HTTP status and reason codes (continued)

HTTP error status code	Reason code	Description
404 (Not Found)	1	The object ID in the URI ( <i>{zbx-id}</i> ) does not designate an existing zBX Node object, or the API user does not have object access permission to the object.
	4	The zBX node designated by the request URI does not support this operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
GET /api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408/hardware-messages HTTP/1.1
x-api-session: 311hjzy0a4o0wrt3zmzn3d0zwx9z5pj8sxuy46xufut63n9
```

Figure 40. List zBX (Node) Hardware Messages: Request

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Mon, 08 Sep 2014 18:38:49 GMT
content-type: application/json;charset=UTF-8
content-length: 2421
{
  "hardware-messages": [
    {
      "element-uri": "/api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408/hardware-messages/0d3f86aa-3765-11e4-9a39-42f2e910664b",
      "text": "Licensed internal code has detected a problem. [Problem # 5]",
      "timestamp": 1410186788140
    },
    {
      "element-uri": "/api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408/hardware-messages/73f58498-3509-11e4-93e5-42f2e910664b",
      "text": "You can no longer modify the default users roles. Your changes have been saved and copied to a new user ID identified in the User Profiles task.",
      "timestamp": 1409927544430
    }
  ]
}
```

Figure 41. List zBX (Node) Hardware Messages: Response

## Get zBX (Node) Hardware Message Properties

The Get zBX (Node) Hardware Message Properties operation retrieves the properties of a single zBX (Node) hardware message.

### HTTP method and URI

```
GET /api/zbx/{zbx-id}/hardware-messages/{hardware-message-id}
```

### URI Variables:

Variable	Description
<i>{zbx-id}</i>	Object ID of the zBX Node object.
<i>{hardware-message-id}</i>	Element ID of the hardware message to retrieve.

## Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the zBX (Node) hardware message object as defined in “Data model” on page 80. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

## Description

This operation retrieves the properties of a single zBX (Node) hardware message specified by *{hardware-message-id}*.

The URI path must designate an existing zBX (Node), and the API user must have object access permission to it; otherwise, status code 404 (Not Found) is returned.

The URI path must designate an existing hardware message; otherwise status code 404 (Not Found) is returned. In addition, the API user must have Action/Task permission to the **Hardware Messages** task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the zBX (Node) object designated by *{zbx-id}*.
- Action/Task permission to the **Hardware Messages** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

*Table 36. Get zBX (Node) Hardware Message Properties: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The request URI designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of “node”).
403 (Forbidden)	1	The API user does not have Action/Task permission for the <b>Hardware Messages</b> task.

Table 36. Get zBX (Node) Hardware Message Properties: HTTP status and reason codes (continued)

HTTP error status code	Reason code	Description
404 (Not Found)	1	The object ID in the URI <i>{zbx-id}</i> does not designate an existing zBX Node object, or the API user does not have object access permission to the object.
	4	The zBX node designated by the request URI does not support this operation.
	322	The element ID in the URI <i>{hardware-message-id}</i> does not designate an existing zBX (Node) hardware message.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
GET /api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408/
    hardware-messages/0d3f86aa-3765-11e4-9a39-42f2e910664b HTTP/1.1
x-api-session: 31lhjzy0a4o0wrt3zmn3d0zxw9z5pj8sxuy46xufut63n9
```

Figure 42. Get zBX (Node) Hardware Message Properties: Request

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Mon, 08 Sep 2014 18:38:49 GMT
content-type: application/json;charset=UTF-8
content-length: 354
{
  "class": "hardware-message",
  "element-id": "0d3f86aa-3765-11e4-9a39-42f2e910664b",
  "element-uri": "/api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408/hardware-messages/
    0d3f86aa-3765-11e4-9a39-42f2e910664b",
  "parent": "/api/zbx/ace3f861-3421-33d5-bcf4-f61f9057d408",
  "text": "Licensed internal code has detected a problem. [Problem # 5]",
  "timestamp": 1410186788140
}
```

Figure 43. Get zBX (Node) Hardware Message Properties: Response

## Delete zBX (Node) Hardware Message

The **Delete zBX (Node) Hardware Message** operation deletes a single zBX (Node) hardware message

### HTTP method and URI

```
DELETE /api/zbx/{zbx-id}/hardware-messages/{hardware-message-id}
```

#### URI Variables:

Variable	Description
<i>{zbx-id}</i>	Object ID of the zBX Node object.
<i>{hardware-message-id}</i>	Element ID of the hardware message to delete.



## Description

This operation deletes a specific zBX (Node) hardware message. The hardware message to be deleted is identified by the *{hardware-message-id}* variable in the URI.

The URI path must designate an existing zBX (Node) and the API user must have object access permission to it; otherwise status code 404 (Not Found) is returned.

The URI path must designate an existing hardware message; otherwise, status code 404 (Not Found) is returned. In addition, the API user must have Action/Task permission to the **Hardware Messages** task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the zBX (Node) object designated by *{zbx-id}*.
- Action/Task permission to the **Hardware Messages** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned with no response body provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

*Table 37. Delete zBX (Node) Hardware Message: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The request URI designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of “node”).
403 (Forbidden)	1	The API user does not have Action/Task permission for the <b>Hardware Messages</b> task.
404 (Not Found)	1	The object ID in the URI <i>{zbx-id}</i> does not designate an existing zBX Node object, or the API user does not have Object-access permission to the object.
	4	The zBX node designated by the request URI does not support this operation.
	322	The element ID in the URI <i>{hardware-message-id}</i> does not designate an existing zBX (Node) hardware message.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
DELETE /api/zbxs/ace3f861-3421-33d5-bcf4-f61f9057d408/
      hardware-messages/0d3f86aa-3765-11e4-9a39-42f2e910664b HTTP/1.1
x-api-session: 311hjzy0a4o0wrt3zmzn3d0zxw9z5pj8sxuy46xufut63ncs9
```

---

*Figure 44. Delete zBX (Node) Hardware Message: Request*

---

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Mon, 08 Sep 2014 18:38:49 GMT
<No response body>
```

---

*Figure 45. Delete zBX (Node) Hardware Message: Response*

## Inventory service data

Information about the zBXs managed by the HMC can be optionally included in the inventory data provided by the Web Services API Inventory Service.

Inventory entries for zBX objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**zbx**" are to be included. An entry for a particular CPC-attached zBX is included only if the API user has object-access permission to the CPC with which that object is associated. An entry for a particular zBX node is included only if the API user has object-access permission for that zBX node object.

For each zBX object to be included, the inventory response array includes entry that is a JSON object with the same contents as is specified in the Response Body Contents section for "Get zBX Properties" on page 91. That is, the data provided is the same as would be provided if a **Get zBX Properties** operation were requested targeting this object.

### Sample inventory data

The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single zBX (that has a **type** of "**node**"). This object would appear as one array entry in the response array:

```

|   {
|     "class": "zbx",
|     "current-dpxi50z-entitlements": 2,
|     "current-isaopt-entitlements": 0,
|     "current-power-entitlements": 4,
|     "current-systemx-entitlements": 2,
|     "description": "Represents one zBX",
|     "is-locked": false,
|     "machine-model": "004",
|     "machine-serial": "0000020ZBX26",
|     "machine-type": "2458",
|     "max-dpxi50z-entitlements": 28,
|     "max-isaopt-entitlements": 0,
|     "max-power-entitlements": 28,
|     "max-systemx-entitlements": 28,
|     "name": "ZBX001",
|     "object-id": "28ba8930-7bc4-11e0-a905-001f163803de",
|     "object-uri": "/api/zbxes/28ba8930-7bc4-11e0-a905-001f163803de",
|     "parent": "/api/ensembles/37c6f8a9-8d5e-3e5d-8466-be79e49dd340"
|     "type": "node"
|   }

```

Figure 46. zBX object: Sample inventory data

## zBX Top-of-Rack switches

A Top-of-Rack Switch element of a zBX represents one of the VLAN-capable switches provided by the zBX to serve as the network infrastructure of the Intra-Ensemble Data Network (IEDN). (Note that the Top-of-Rack switches used for management purposes are considered internal components of zManager and are not surfaced to API clients as TOR objects.)

### Data model

The Top-of-Rack switch object provides the following properties:

Table 38. zBX Top-of-Rack switches: properties

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/ URI	The canonical URI path for a TOR object is of the form <code>/api/zbxes/{zbx-id}/top-of-rack-switches/{tor-id}</code> where <code>{tor-id}</code> is the value of the <b>element-id</b> property of the Top-of-Rack switch object.
<b>element-id</b>	—	String	The element ID of the TOR.
<b>parent</b>	—	String/ URI	The canonical URI path of the zBX that is the parent of the TOR.
<b>class</b>	—	String	The class of a TOR object is <b>"top-of-rack-switch"</b> .
<b>name</b>	—	String (32)	The name of the TOR. This name cannot be changed.
<b>tor-ports-list</b>	—	Array of tor-port- info objects	Array of nested tor-port-info objects (described in the next table), each entry of which describes one port on the TOR switch.

The following tor-port-info nested object describes the properties of a single TOR port:

Table 39. zBX Top-of-Rack switches: tor-port-info nested object properties

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/ URI	Canonical URI of the port for this TOR: /api/zbx/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id} when {port-id} is the value of the <b>port-num</b> property of the port.
<b>port-num</b>	—	String (3)	The TOR port number, as a string of 1 to 3 numeric digits.
<b>type</b>	—	String Enum	This is the type of port. TOR ports have the following types: <ul style="list-style-type: none"> <li>• <b>"internal"</b> - Provides connectivity to an ISAOPT coordinator blade.</li> <li>• <b>"external"</b> - Provides connectivity to networks external to the IEDN.</li> </ul>
<b>port-access</b>	(w)	String Enum	This describes the port access mode associated with this TOR port. Possible values: <ul style="list-style-type: none"> <li>• <b>"trunk"</b></li> <li>• <b>"access"</b></li> </ul>
<b>all-virtual-networks</b>	(w)	Boolean	If true then all defined virtual networks in the ensemble can be used with this TOR port. When true, the <b>virtual-networks-list</b> property value will be an empty array. <p>This property can have the value true only if the <b>type</b> property has the value <b>"internal"</b> and the <b>port-access-mode</b> property has the value <b>"trunk"</b>. That is, only internal ports configured in trunk mode can be configured to allow access to all defined virtual networks.</p> <p>If true, attempts to invoke the <b>Add Top-of-Rack Switch Port to Virtual Networks</b> or <b>Remove Top-of-Rack Switch Port from Virtual Networks</b> operations will result in errors.</p> <p>Changing this property from false to true will remove all virtual networks from the current <b>virtual-networks-list</b>.</p>
<b>virtual-networks-list</b>	—	Array of URI	This is the list of virtual network URIs that represent the virtual networks that can be used for this port. There can be more than one virtual network URI in this list only if the <b>port-access</b> property has the value <b>"trunk"</b> . <p>If the <b>all-virtual-networks</b> property is false, the <b>Add Top-of-Rack Switch Port to Virtual Networks</b> and <b>Remove Top-of-Rack Switch Port from Virtual Networks</b> operations may be used to modify this list.</p> <p>Changing the <b>all-virtual-networks</b> property from false to true will remove all virtual networks from this list.</p>

Table 39. zBX Top-of-Rack switches: tor-port-info nested object properties (continued)

Name	Qualifier	Type	Description
<b>allow-all-macs</b>	(w)	Boolean	<p>If true, all MAC addresses are allowed to access this port and the <b>mac-filter-list</b> array will be empty.</p> <p>If false, only those MAC addresses listed in the <b>mac-filter-list</b> property are permitted to access this port. Network traffic from all other MAC addresses is filtered out.</p> <p>If true, attempts to invoke the <b>Add MAC Filter to Top-of-Rack Switch Port</b> and <b>Remove MAC Filter from Top-of-Rack Switch Port</b> operations will result in errors.</p> <p>Changing the value from false to true will remove all MAC filters from the current <b>mac-filter-list</b> property.</p> <p>This property is writable only for TOR ports with a type value of <b>"external"</b>. The value of this property is always false for TOR ports with a <b>type</b> value of <b>"internal"</b>.</p>
<b>mac-filter-list</b>	—	Array of String (17)	<p>The list of MAC addresses to allow (permit) for this port. Each entry in the list is a MAC address. A MAC address is represented as a string of length 17 consisting of 6 groups of two lower-case hexadecimal digits separated by colons ( : ), e.g. "01:23:45:67:89:ab".</p> <p>If the <b>allow-all-macs</b> property is true, this list is empty. If the <b>allow-all-macs</b> property is false, the <b>Add MAC Filter to Top-of-Rack Switch Port</b> and <b>Remove MAC Filter from Top-of-Rack Switch Port</b> operations may be used to modify this list.</p> <p>Changing the <b>allow-all-macs</b> property from false to true will remove all MAC addresses from this list.</p> <p>This property is only applicable to TOR ports with a <b>type</b> value of <b>"external"</b>. The value of this property is always an empty array of TOR ports with a <b>type</b> value of <b>"internal"</b>.</p>

## Operations

### List Top-of-Rack Switches of a zBX

The **List Top-of-Rack Switches of a zBX** operation returns a list of the Top-of-Rack (TOR) switches for a zBX managed by the HMC.

#### HTTP method and URI

**GET** /api/zbx/{zbx-id}/top-of-rack-switches

In this request, the URI variable *{zbx-id}* is the object ID of a zBX object for which top-of-rack switches are to be listed.

#### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
tors-list	Array of objects	Array of nested objects where each element in the array is a TOR-info object, described in the next table.

Each TOR-info object contains the following fields:

Field name	Type	Description
element-uri	String/ URI	Canonical URI path of the TOR object.
element-id	String	Object ID of the TOR object.
name	String	Display name of the TOR object.

## Description

On successful completion, this operation obtains a list of URIs that represent the Top-of-Rack (TOR) switches for the zBX designated by *{zbx-id}*. If no TORs are present, then the response body is provided and contains an empty **tors-list** array.

- | A TOR switch within a CPC-attached zBX is included in the list only if the API user has object-access permission to the CPC associated with the zBX that contains the TOR switch. A TOR switch within a zBX node is included in the list only if the API user has object-access permission to the zBX node object that contains the TOR switch. If the ensemble contains a TOR switch but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

If the Ensemble does not contain any TOR switches, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- | • Object access permission to the CPC associated with the zBX containing TOR switches to be listed in the result (for CPC-attached zBX objects)
- | • Object access permission to the zBX node containing TOR switches to be listed in the result (for CPC-attached zBX objects)
- | • Action/task permission to **Configure Top-of-Rack Switch** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 117.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message:

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	0	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Configure Top of Rack Switch</b> task is required.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{zbx-id}</i> ) does not designate an existing resource, or designates a resource for which the API user does not have object-access permission.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches HTTP/1.1
x-api-session: sqernk0lqu0s49oul0drld6ifcb1cp1f902og86d9apesqjsl
```

---

Figure 47. List Top-of-Rack Switches: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 23 Nov 2011 20:57:48 GMT
content-type: application/json;charset=UTF-8
content-length: 409
{
  "tors-list": [
    {
      "element-id": "189ad58d-c19f-4cdc-8dc8-639f3ccd4f49",
      "element-uri": "/api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/189ad58d-c19f-4cdc-8dc8-639f3ccd4f49",
      "name": "GG0210487100"
    },
    {
      "element-id": "fea63433-e03d-4ea1-a5da-6a2fc7abe844",
      "element-uri": "/api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/fea63433-e03d-4ea1-a5da-6a2fc7abe844",
      "name": "GG0210487101"
    }
  ]
}
```

---

Figure 48. List Top-of-Rack Switches: Response

## Get Top-of-Rack Switch Properties

The **Get Top-of-Rack Switch Properties** operation obtains the properties of the TOR specified by the *{tor-id}* managed object.

### HTTP method and URI

```
GET /api/zbxs/{zbx-id}/top-of-rack-switches/{tor-id}
```

#### URI variables:

Variable	Description
<i>{zbx-id}</i>	Object ID of the zBX object containing the TOR switch.
<i>{tor-id}</i>	Element ID of the TOR object.

## Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the TOR object as defined in the “Data model” on page 115. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

## Description

On successful completion, the **Get Top-of-Rack Switch Properties** operation obtains the properties of the TOR specified by *{tor-id}*.

- | The URI path must designate an existing TOR switch object. If the TOR switch is within a CPC-attached zBX, the API user must have object-access permission to the CPC associated with the zBX. If the TOR switch is within a zBX node, the API user must have object-access permission to the zBX node. If these conditions are not met, status code 404 (Not Found) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- | • Object access permission to the CPC associated with the zBX containing the TOR switch (for CPC-attached zBX objects)
- | • Object-access permission to the zBX node containing the TOR switch (for zBX node objects)
- | • Action/task permission to **Configure Top-of-Rack Switch** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 119.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message:

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	0	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Configure Top of Rack Switch</b> task is required.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{zbx-id}</i> ) does not designate an existing resource, or designates a resource for which the API user does not have object-access permission.
	360	The TOR element ID in the URI ( <i>{tor-id}</i> ) does not designate an existing top of rack switch.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/zbx/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/
  fea63433-e03d-4ea1-a5da-6a2fc7abe844 HTTP/1.1
x-api-session: sqernk0lqu0s49ou10dr1d6ifcb1cp1f902og86d9apesqjs1
```

---

Figure 49. Get Top-of-Rack Switch Properties: Request



```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 23 Nov 2011 20:57:53 GMT
content-type: application/json;charset=UTF-8
content-length: 1314
{
  "class": "top-of-rack-switch",
  "element-id": "fea63433-e03d-4ea1-a5da-6a2fc7abe844",
  "element-uri": "/api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/fea63433-e03d-4ea1-a5da-6a2fc7abe844",
  "name": "GG0210487101",
  "parent": "/api/zbxs/da5d7720-a337-11e0-9555-00262df332b3",
  "tor-ports-list": [
    {
      "all-virtual-networks": false,
      "allow-all-macs": true,
      "element-uri": "/api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/fea63433-e03d-4ea1-a5da-6a2fc7abe844/ports/8",
      "mac-filter-list": [],
      "port-access": "access",
      "port-num": "8",
      "type": "internal",
      "virtual-networks-list": [
        "/api/virtual-networks/4ccb3c0c-a703-11df-a6fc-00215ef9b504"
      ]
    },
    {
      "all-virtual-networks": false,
      "allow-all-macs": true,
      "element-uri": "/api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/fea63433-e03d-4ea1-a5da-6a2fc7abe844/ports/34",
      "mac-filter-list": [],
      "port-access": "trunk",
      "port-num": "34",
      "type": "external",
      "virtual-networks-list": [
        "/api/virtual-networks/4ccb3c0c-a703-11df-a6fc-00215ef9b504"
      ]
    }
  ]
}

```

Figure 50. Get Top-of-Rack Switch Properties: Response

## Get Top-of-Rack Switch Port Details

The **Get Top-of-Rack Switch Port Details** operation obtains the properties of the designated TOR port specified by the *{port-id}* identifier.

### HTTP method and URI

**GET** /api/zbxs/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}

#### URI variables:

Variable	Description
<i>{zbx-id}</i>	Object ID of the zBX object containing the TOR switch.
<i>{tor-id}</i>	Element ID of the TOR object.
<i>{port-id}</i>	Port number of the TOR port.

## Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the designated TOR port defined in the “Data model” on page 115 by `tor-port-info`. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

## Description

On successful completion, the **Get Top-of-Rack Switch Port Details** operation obtains the properties of the TOR port specified by `{port-id}`.

## Authorization requirements

This operation has the following authorization requirements:

- | • Object access permission to the CPC associated with the zBX containing the TOR switch (for CPC-attached zBX objects)
- | • Object-access permission to the zBX node containing the TOR switch (for zBX node objects)
- | • Action/task permission to **Configure Top-of-Rack Switch** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message:

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	0	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Configure Top-of-Rack Switch</b> task is required.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <code>{zbx-id}</code> ) does not designate an existing resource, or designates a resource for which the API user does not have object-access permission.
	360	The TOR element ID in the URI ( <code>{tor-id}</code> ) does not designate an existing top of rack switch.
	361	The port ID the URI ( <code>{port-id}</code> ) does not designate an existing port.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/
    fea63433-e03d-4ea1-a5da-6a2fc7abe844/ports/34 HTTP/1.1
x-api-session: sqernk0lqu0s49oul0drld6ifcb1cp1f902og86d9apesqjs1
```

---

Figure 51. Get Top-of-Rack Switch Port Details: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 23 Nov 2011 20:58:19 GMT
content-type: application/json;charset=UTF-8
content-length: 427
{
  "all-virtual-networks": false,
  "allow-all-macs": false,
  "element-uri": "/api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/
    fea63433-e03d-4ea1-a5da-6a2fc7abe844/ports/34",
  "mac-filter-list": [
    "00:24:7e:e0:ea:9e"
  ],
  "port-access": "trunk",
  "port-num": "34",
  "type": "external",
  "virtual-networks-list": [
    "/api/virtual-networks/4ccb3c0c-a703-11df-a6fc-00215ef9b504",
    "/api/virtual-networks/cd42c1c0-1615-11e1-817f-00215e6a0c26"
  ]
}
```

---

Figure 52. Get Top-of-Rack Switch Port Details: Response

## Update Top-of-Rack Switch Port Properties

The **Update Top-of-Rack Switch Port Properties** operation updates selected properties of the specified TOR port.

### HTTP method and URI

```
POST /api/zbxs/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}
```

URI variables:

Variable	Description
{zbx-id}	Object ID of the zBX object containing the TOR switch.
{tor-id}	Element ID of the TOR object.
{port-id}	Port number of the TOR port.

### Request body contents

The request body is expected to contain a JSON object that provides the new values of any writeable property that is to be updated by this operation. Field names and data types in this JSON object are expected to match the corresponding property names and data types defined by the “Data model” on page 115. The JSON object can and should omit fields for properties whose values are not to be changed by this operation.

## Description

This operation updates writeable properties of the TOR port specified by *{port-id}*.

The request body does not need to specify a value for all writeable properties, but rather can and should contain fields for the properties to be updated. Object properties for which no input value is provided remain unchanged by this operation.

The request body is validated against the schema described in the “Request body contents” on page 123. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. In addition to occurring for common validation reasons, status code 400 is returned when the requested changes are not valid considering the port's type, or for inconsistencies in the request and the **port-access-mode** setting.

## Authorization requirements

This operation has the following authorization requirements:

- l • Object access permission to the CPC associated with the zBX containing the TOR switch (for CPC-attached zBX objects)
- l • Object-access permission to the zBX node containing the TOR switch (for zBX node objects)
- Action/task permission to **Configure Top-of-Rack Switch** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message:

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	133	The port's port-access setting is not valid for the update. A change from "trunk" to "access" mode is not allowed if more than one virtual network is defined to the port.
	134	The port's type does not support MAC filters. Only ports that have a type of "external" support MAC filters.
	139	The port's type does not support the <b>all-virtual-networks</b> setting.
403 (Forbidden)	0	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Configure Top-of-Rack Switch</b> task is required.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{zbx-id}</i> ) does not designate an existing resource, or designates a resource for which the API user does not have object-access permission.
	360	The TOR element ID in the URI ( <i>{tor-id}</i> ) does not designate an existing top of rack switch.
	361	The port ID the URI ( <i>{port-id}</i> ) does not designate an existing port.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/
    fea63433-e03d-4ea1-a5da-6a2fc7abe844/ports/34 HTTP/1.1
x-api-session: sqernk0lqu0s49oul0drld6ifcb1cplf902og86d9apesqjsl
content-type: application/json
content-length: 49
{
  "allow-all-macs": false,
  "port-access": "trunk"
}
```

---

Figure 53. Update Top-of-Rack Switch Port Properties: Request

---

```
204 No Content
date: Wed, 23 Nov 2011 20:57:59 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 54. Update Top-of-Rack Switch Port Properties: Response

## Add MAC Filters to Top-of-Rack Switch Port

The **Add MAC Filters to Top-of-Rack Switch Port** operation adds MAC address filters to the designated Top-of-Rack Switch port to permit a list of MAC addresses to connect to this TOR port.

### HTTP method and URI

**POST** /api/zbxs/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}/operations/add-mac-filters

#### URI variables:

Variable	Description
{zbx-id}	Object ID of the zBX object containing the TOR switch.
{tor-id}	Element ID of the TOR object.
{port-id}	Port number of the TOR port.

### Request body contents

The request body is expected to contain a JSON object that provides the following:

Field name	Type	Rqd/Opt	Description
mac-address-list	Array of String	Required	Array of MAC addresses to add to the <b>mac-filter-list</b> for the specified TOR port. If successful, these MAC addresses will be permitted to access the TOR port. This operation will fail if the TOR port's <b>allow-all-macs</b> is true.

## Description

On successful execution, this operation adds the MAC provided in the **mac-address-list** field to the TOR port's **mac-filter-list**. If the inputs contain addresses that are already in the list, these will be accepted without error. These MAC addresses will be permitted to access the specified TOR port.

The request body is validated against the schema described in “Request body contents” on page 125. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. Consider the following to prevent bad requests:

- Ensure the type of the target port supports the setting of MAC filters
- The **allow-all-macs** property must be false
- Ensure that the MAC address is valid.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC associated with the zBX containing the TOR switch (for CPC-attached zBX objects)
- Object-access permission to the zBX node containing the TOR switch (for zBX node objects)
- Action/task permission to **Configure Top-of-Rack Switch** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message:

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	134	The port's type does not support MAC filters. Only ports that have a <b>type</b> value of "external" support MAC filters.
	135	This operation cannot be performed when the <b>allow-all-macs</b> property of the TOR port is true.
	136	The MAC address is invalid.
403 (Forbidden)	0	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Configure Top-of-Rack-Switch</b> task is required.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{zbx-id}</i> ) does not designate an existing resource, or designates a resource for which the API user does not have object-access permission.
	360	The TOR element ID in the URI ( <i>{tor-id}</i> ) does not designate an existing top of rack switch.
	361	The port ID the URI ( <i>{port-id}</i> ) does not designate an existing port.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/
 63433-e03d-4ea1-a5da-6a2fc7abe844/ports/34/operations/add-mac-filters HTTP/1.1
x-api-session: sqernk01qu0s49oul0drld6ifcb1cplf902og86d9apesqjs1
content-type: application/json
content-length: 43
{
  "mac-address-list": [
    "00:24:7E:E0:EA:9E"
  ]
}
```

---

Figure 55. Add MAC Filters to Top-of-Rack Switch Port: Request

---

```
204 No Content
date: Wed, 23 Nov 2011 20:58:07 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 56. Add MAC Filters to Top-of-Rack Switch Port: Response

## Remove MAC Filters from Top-of-Rack Switch Port

The **Remove MAC Filters from Top-of-Rack Switch Port** operation removes MAC address filters from the designated Top-of-Rack Switch port.

### HTTP method and URI

```
POST /api/zbxs/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}/operations/remove-mac-filters
```

#### URI variables:

Variable	Description
{zbx-id}	Object ID of the zBX object containing the TOR switch.
{tor-id}	Element ID of the TOR object.
{port-id}	Port number of the TOR port.

### Request body contents

The request body is expected to contain a JSON object that provides the following:

Field name	Type	Rqd/Opt	Description
mac-address-list	Array of String	Required	This is an array of the MAC addresses to remove from the MAC mac-filter-list for the specified TOR port.

### Description

On successful execution, this operation removes the MAC provided in the **mac-address-list** field to the TOR port's **mac-filter-list**. If a MAC address input is not currently in the **mac-filter-list**, it will be ignored without resulting in an error.

The request body is validated against the schema described in “Request body contents” on page 127. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. Consider the following to prevent bad requests:

- Ensure the type of the target port supports the setting of MAC filters
- The **allow-all-macs** property must be false
- Ensure that the MAC address is valid.

## Authorization requirements

This operation has the following authorization requirements:

- | • Object access permission to the CPC associated with the zBX containing the TOR switch (for CPC-attached zBX objects)
- | • Object-access permission to the zBX node containing the TOR switch (for zBX node objects)
- | • Action/task permission to **Configure Top-of-Rack Switch** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message:

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	134	The port's type does not support MAC filters. Only ports that have a <b>type</b> value of "external" support MAC filters.
	135	This operation cannot be performed when the <b>allow-all-macs</b> property of the TOR port is true.
	136	The MAC address is invalid.
403 (Forbidden)	0	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Configure Top-of-Rack Switch</b> task is required.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{zbx-id}</i> ) does not designate an existing resource, or designates a resource for which the API user does not have object-access permission.
	360	The TOR element ID in the URI ( <i>{tor-id}</i> ) does not designate an existing top of rack switch.
	361	The port ID the URI ( <i>{port-id}</i> ) does not designate an existing port.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.



## Example HTTP interaction

---

```
POST /api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/
fea63433-e03d-4ea1-a5da-6a2fc7abe844/ports/34/operations/remove-mac-filters HTTP/1.1
x-api-session: sqernk01qu0s49oul0drld6ifcb1cplf902og86d9apesqjs1
content-type: application/json
content-length: 43
{
  "mac-address-list": [
    "00:24:7E:E0:EA:9E"
  ]
}
```

---

Figure 57. Remove MAC Filters from Top-of-Rack Switch Port: Request

---

```
204 No Content
date: Wed, 23 Nov 2011 20:58:40 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 58. Remove Mac Filters from Top-of-Rack Switch Port: Response

## Add Top-of-Rack Switch Port to Virtual Networks

The **Add Top-of-Rack Switch Port to Virtual Networks** operation adds Top-of-Rack switch port to the specified list of virtual networks.

### HTTP method and URI

```
POST /api/zbxs/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}/operations/add-virtual-networks
```

#### URI variables:

Variable	Description
{zbx-id}	Object ID of the zBX object containing the TOR switch.
{tor-id}	Element ID of the TOR object.
{port-id}	Port number of the TOR port.

### Request body contents

The request body is expected to contain a JSON object that provides the following:

Field name	Type	Rqd/Opt	Description
add-virtual-networks	Array of String/URI	Required	Array of virtual network URIs that define the virtual networks that can be used for the specified TOR port.

## Description

On successful execution, this operation adds the specified virtual networks to the TOR port's **virtual-networks-list** array. If the inputs contain virtual networks that are already in the list, these will be accepted without error. Traffic for VLAN IDs representing these virtual networks will now be able to be used on the specified TOR port.

The request body is validated against the schema described in "Request body contents" on page 129. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. Consider the following to prevent bad requests:

- Ensure the **type** of the target port supports the setting of virtual networks
- The target TOR port's **all-virtual-networks** property must be false when issuing this request
- If the port's **port-mode** property is "access", then it only supports one virtual network.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC associated with the zBX containing the TOR switch (for CPC-attached zBX objects)
- Object-access permission to the zBX node containing the TOR switch (for zBX node objects)
- Action/task permission to **Configure Top-of-Rack Switch** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message:

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	137	The <b>port-access</b> property is "access", therefore, only one virtual network is allowed in the TOR port's <b>virtual-networks-list</b> .
	138	This operation cannot be performed when the <b>all-virtual-networks</b> property of the TOR port is true.
403 (Forbidden)	0	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Configure Top-of-Rack Switch</b> task is required.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{zbx-id}</i> ) does not designate an existing resource, or designates a resource for which the API user does not have object-access permission.
	360	The TOR element ID in the URI ( <i>{tor-id}</i> ) does not designate an existing top of rack switch.
	361	The port ID the URI ( <i>{port-id}</i> ) does not designate an existing port.
	362	A URI in the <b>add-virtual-networks</b> field of the request body does not designate an existing virtual network.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/zbxs/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/
    fea63433-e03d-4ea1-a5da-6a2fc7abe844/ports/34/operations/add-virtual-networks HTTP/1.1
x-api-session: sqernk0lqu0s49oul0drld6ifcb1cp1f902og86d9apesqjsl
content-type: application/json
content-length: 88
{
  "add-virtual-networks": [
    "/api/virtual-networks/cd42c1c0-1615-11e1-817f-00215e6a0c26"
  ]
}
```

---

Figure 59. Add Top-of-Rack Switch Port to Virtual Networks: Request

---

```
204 No Content
date: Wed, 23 Nov 2011 20:58:18 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

<No response body>

---

Figure 60. Add Top-of-Rack Switch Port to Virtual Networks: Response

## Remove Top-of-Rack Switch Port from the Virtual Networks

The **Remove Top-of-Rack Switch Port from the Virtual Networks** operation removes the Top-of-Rack Switch port from the specified virtual networks.

### HTTP method and URI

**POST** /api/zbxs/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}/operations/remove-virtual-networks

#### URI variables:

Variable	Description
{zbx-id}	Object ID of the zBX object containing the TOR switch.
{tor-id}	Element ID of the TOR object.
{port-id}	Port number of the TOR port.

### Request body contents

The request body is expected to contain a JSON object that provides the following:

Field name	Type	Rqd/Opt	Description
remove-virtual-networks	Array of String/URI	Required	List of virtual network URIs to remove from the <b>virtual-network-list</b> array for the specified TOR port.

## Description

On successful execution, this operation removes the specified virtual networks from the TOR port's **virtual-networks-list** array. If a virtual network in the input is not currently in the **virtual-networks-list**, it will be ignored without resulting in an error.

The request body is validated against the schema described in “Request body contents” on page 131. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. Consider the following to prevent bad requests:

- Ensure the **type** of the target port supports the setting of virtual networks
- The target TOR port's **all-virtual-networks** property must be false when issuing this request.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC associated with the zBX containing the TOR switch (for CPC-attached zBX objects)
- Object-access permission to the zBX node containing the TOR switch (for zBX node objects)
- Action/task permission to **Configure Top-of-Rack Switch** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message:

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	138	This operation cannot be performed when the <b>all-virtual-networks</b> property of the TOR port is true.
403 (Forbidden)	0	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Configure Top-of-Rack Switch</b> task is required.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{zbx-id}</i> ) does not designate an existing resource, or designates a resource for which the API user does not have object-access permission.
	360	The TOR element ID in the URI ( <i>{tor-id}</i> ) does not designate an existing top of rack switch.
	361	The port ID the URI ( <i>{port-id}</i> ) does not designate an existing port.
	362	A URI in the <b>remove-virtual-networks</b> field of the request body does not designate an existing virtual network.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/zbx/da5d7720-a337-11e0-9555-00262df332b3/top-of-rack-switches/
fea63433-e03d-4ea1-a5da-6a2fc7abe844/ports/34/operations/remove-virtual-networks HTTP/1.1
x-api-session: sqernk01qu0s49oul0drld6ifcb1cplf902og86d9apesqjs1
content-type: application/json
content-length: 91
{
  "remove-virtual-networks": [
    "/api/virtual-networks/cd42c1c0-1615-11e1-817f-00215e6a0c26"
  ]
}
```

---

Figure 61. Remove Top-of-Rack Switch Port from the Virtual Networks: Request

---

```
204 No Content
date: Wed, 23 Nov 2011 20:58:32 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 62. Remove Top-of-Rack Switch Port from the Virtual Networks: Response

## Rack object

A Rack object represents a rack that houses the zBX components.

There is at least one Rack object for each zBX, designated as the rack in location B. This rack houses the top of rack (TOR) switches for the zBX, the Support Elements (for a zBX node), and one or two BladeCenter chassis. Rack B is the first rack to be populated when installing a zBX. Additional racks are added for larger configurations that contain more than two BladeCenter chassis in the zBX. The additional racks, each containing one to two chassis, are designated with consecutive location codes (C, D, E).

## Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, but does not provide the operational-status-related properties defined in that schema because it does not maintain the concept of an operational status.

The following class-specific specializations apply to the other base managed object properties:

Table 40. Rack object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>name</b>	(ro)	String (1-64)	The name of the object. Currently, this is assigned by zManager based on the rack's location (has the same value as the location property). <sup>1</sup>
<b>description</b>	—	String	This field is not provided.
<b>object-uri</b>	—	String/URI	The canonical URI path for a Rack object is of the form <code>/api/racks/{rack-id}</code> .
<b>parent</b>	—	String/URI	The canonical URI path of the parent zBX object, of the form <code>/api/zbx/{zbx-id}</code> .
<b>class</b>	—	String	The value "rack".

Table 40. Rack object: base managed object properties specializations (continued)

Name	Qualifier	Type	Description of specialization
<b>Note:</b>			
1. This <b>name</b> property is currently assigned based on the location of the rack and thus has the same value as the <b>location</b> property. However, it is possible that the API could be extended to allow this property to be writable, in which case an API or User-Interface user could change the name to contain arbitrary data. Therefore, API client applications that are interested in determining the location of the rack should not rely on the contents and format of the <b>name</b> property, but rather obtain location information from the location property.			

## Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties:

Table 41. Rack object: class specific properties

Name	Type	Description
serial-number	String	The 12 character serial number of the rack.
location	String	The zManager assigned location code for the rack, as an uppercase alphabetic character starting with "B" for the first rack of the zBX and with subsequent racks being designated by consecutive letters.

## Operations

### List Racks of a zBX

The **List Racks of a zBX** operation lists the racks where the zBX components are mounted.

#### HTTP method and URI

**GET** /api/zbx/{zbx-id}/racks

In this request, the URI variable {zbx-id} is the object ID of the zBX object whose racks are to be obtained.

#### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) used to limit returned objects

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
racks	Array of objects	Array of nested rack-info objects, described in the next table

Each nested rack-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the Rack object in the form /api/racks/{rack-id}
name	String	The <b>name</b> property of the Rack object

## Description

The **List Racks of a zBX** operation lists the racks where the zBX components are mounted. BladeCenters are plugged into the rack, and the blades are plugged into the BladeCenters. The object URI and name are provided for each rack.

If the **name** query parameter is specified, then a rack is included in the list only if the name pattern matches the **name** property of the object.

If the HMC does not manage any racks, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- | • Object access permission to the CPC with which the zBX and Rack objects are associated (for a CPC-attached zBX).
- | • Object-access permission to the zBX object (for a zBX node).

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 134.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{zbx-id}</i> does not designate a zBX object, or the API user does not have object access permission to the CPC with which it is associated.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/zbx/54a9716c-a326-11e0-9469-001f163805d8/racks HTTP/1.1
x-api-session: 5q07hx6jgp2ngn2cypq1zxot76sfwnzky0ih8nddd5hz6bpiue
```

---

Figure 63. List Racks of a zBX: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 21 Jul 2011 17:49:01 GMT
content-type: application/json;charset=UTF-8
content-length: 142
{
  "racks": [
    {
      "name": "B",
      "object-uri": "/api/racks/b434398a-a328-11e0-9b4a-001f163805d8"
    }
  ]
}
```

---

Figure 64. List Racks of a zBX: Response

## Get Rack Properties

The **Get Rack Properties** operation retrieves the properties of a single Rack object that is designated by its object ID.

### HTTP method and URI

**GET** /api/racks/{*rack-id*}

In this request, the URI variable *{rack-id}* is the object ID of the Rack object for which properties are to be obtained.

### Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Rack object as defined in the “Data model” on page 133. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

The **Get Rack Properties** operation returns the current properties for the Rack object specified by *{rack-id}*.

On successful execution, all of the current properties as defined in “Data model” on page 133 for the Rack object are provided in the response body, and HTTP status code 200 (OK) is returned.

- | The URI path must designate an existing Rack object. If the URI path designates a rack within a
- | CPC-attached zBX, the API user must have object-access permission to the CPC with which the zBX is
- | associated. If the URI path designates a rack within a zBX node, the API user must have object-access
- | permission to the zBX node object. If these conditions are not met, status code 404 (Not Found) is
- | returned.

### Authorization requirements

This operation has the following authorization requirements:

- | • Object-access permission to the CPC associated with the zBX within which the rack object designated
- | by *{rack-id}* is contained (for racks within a CPC-attached zBX).
- | • Object-access permission to the zBX within which the rack object designated by *{rack-id}* is contained
- | (for racks within a zBX node).



## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 136.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{rack-id}</i> does not designate an existing Rack object, or the API user does not have object access permission to the CPC with which the rack is associated.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/racks/04419e1e-a9db-11e0-8077-f0def1610d20 HTTP/1.1
x-api-session: 5wksmtktms30ajeohh0bn411fdzusmk1ld4jydd1des5t78aq
```

---

Figure 65. Get Rack Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 15 Nov 2011 14:22:52 GMT
content-type: application/json;charset=UTF-8
content-length: 272
{
  "class": "rack",
  "description": "Rack",
  "is-locked": false,
  "location": "C",
  "name": "C",
  "object-id": "04419e1e-a9db-11e0-8077-f0def1610d20",
  "object-uri": "/api/racks/04419e1e-a9db-11e0-8077-f0def1610d20",
  "parent": "/api/zbxs/291e385e-a9cd-11e0-8650-f0def1610d20",
  "serial-number": "12345"
}
```

---

Figure 66. Get Rack Properties: Response

## Inventory service data

Information about the zBX racks managed by the HMC can be optionally included in the inventory data provided by the Web Services API Inventory Service.

Inventory entries for Rack objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class **"rack"** are to be included. An entry for a particular Rack is included only if the API user has object-access permission to the CPC with which that object is associated.

For each Rack object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for “Get Rack Properties” on page 136. That is, the data provided is the same as would be provided if a **Get Rack Properties** operation were requested targeting this object.

## Sample inventory data

The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single rack. This object would appear as one array entry in the response array:

```
{
  "class": "rack",
  "description": "Rack",
  "is-locked": false,
  "location": "B",
  "name": "B",
  "object-id": "28bc03c8-7bc4-11e0-a905-001f163803de",
  "object-uri": "/api/racks/28bc03c8-7bc4-11e0-a905-001f163803de",
  "parent": "/api/zbx/28ba8930-7bc4-11e0-a905-001f163803de",
  "serial-number": "123456123456"
}
```

Figure 67. Rack object: Sample inventory data

## BladeCenter object

A BladeCenter object represents a single BladeCenter of the zBX.

### Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, including the operational-status properties, with the following class-specific specialization:

Table 42. BladeCenter object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>name</b>	(ro)	String (1-64)	The zManager-assigned name of the zBX BladeCenter. <sup>1</sup>
<b>description</b>	—	String	This property is not provided.
<b>object-uri</b>	—	String/URI	The canonical URI path for a zBX BladeCenter object is of the form <code>/api/bladecenters/{bladecenter-id}</code>
<b>parent</b>	—	String/URI	The canonical URI path of the parent Rack object, of the form <code>/api/racks/{rack-id}</code>
<b>class</b>	—	String	The value <b>"bladecenter"</b>
<b>status</b>	(sc)	String Enum	The status of the blade center. Values: <ul style="list-style-type: none"> <li><b>"no-power"</b></li> <li><b>"operating"</b></li> </ul>
<b>additional-status</b>	—	String Enum	This property is not provided.

Table 42. BladeCenter object: base managed object properties specializations (continued)

Name	Qualifier	Type	Description of specialization
<b>Note:</b>			
1. This <b>name</b> property is currently assigned based on the location of the BladeCenter and is of the form <b>RackName.BladecenterName</b> (e.g. B.1). However, it is possible that the API could be extended to allow this property to be writable, in which case an API or User-Interface user could change the name to contain arbitrary data. Therefore, API client applications that are interested in determining the location of the BladeCenter should not rely on the contents and format of the <b>name</b> property, but rather obtain location information from the <b>location</b> property.			

## Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties:

Table 43. BladeCenter object: class specific properties

Name	Qualifier	Type	Description
<b>machine-type</b>	—	String	4 characters
<b>machine-model</b>	—	String	3 characters
<b>machine-serial</b>	—	String	7 characters
<b>location</b>	—	String (4)	4 Characters (RxxB) RxxB – BladeCenter vertical position in rack R
<b>has-hardware-messages</b>	(pc)	Boolean	The BladeCenter has a hardware message (true) or the BladeCenter does not have a hardware message (false).

## Energy Management Related Additional Properties

In addition to the properties defined above, this object includes the following additional class-specific properties related to energy management. For further explanation of the various states involved, please see the "Overview" section in the "Energy Management" chapter of this document.

Table 44. BladeCenter object: energy management related additional properties

Name	Qualifier	Type	Description
<b>power-rating</b>	—	Integer	Specifies the maximum power usage of this BladeCenter in watts (W). This value is a calculated value, as indicated by the electrical rating labels or system rating plates of the BladeCenter components.
<b>power-consumption</b>	(mg)	Integer	Specifies the current power consumption in watts (W) of this BladeCenter. The BladeCenter power consumption includes the power consumption of the Blades contained within the BladeCenter and the shared infrastructure.

Table 44. BladeCenter object: energy management related additional properties (continued)

Name	Qualifier	Type	Description
<b>power-saving</b>	—	String Enum	<p>Specifies the current power saving setting of the BladeCenter. Power saving reduces the energy consumption of a system, and can be managed it using the Set Power Saving operation. The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - Specifies not reducing the power consumption and performance of the BladeCenter. This is the default setting.</li> <li>• <b>"low-power"</b> - Specifies low power consumption for all components of the BladeCenter enabled for power saving.</li> <li>• <b>"custom"</b> - Specifies that some, but not all, components of the BladeCenter are in the Low power setting.</li> <li>• <b>"not-supported"</b> - Specifies that power saving is not supported for this BladeCenter.</li> <li>• <b>"not-available"</b> - Specifies that <b>power-saving</b> property could not be read from this BladeCenter.</li> <li>• <b>"not-entitled"</b> - Specifies that the server is not entitled to power saving.</li> </ul>
<b>power-saving-state</b>	—	String Enum	<p>Specifies the power saving setting of the BladeCenter set by the user. Note that this property indicates the user setting and may not match the real state of the hardware compared to the <b>power-saving</b> property. For more information, see "Group power saving" on page 182. The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - Specifies not reducing the power consumption and performance of the BladeCenter. This setting will be applied to all children.</li> <li>• <b>"low-power"</b> - Specifies low power consumption for all components of the BladeCenter enabled for power saving.</li> <li>• <b>"custom"</b> - Specifies that the BladeCenter does not control the children. This is the default setting.</li> <li>• <b>"not-supported"</b> - Specifies that power saving is not supported for this BladeCenter.</li> <li>• <b>"not-entitled"</b> - Specifies that the server is not entitled to power saving.</li> </ul>
<b>power-save-allowed</b>	—	String Enum	<p>Should be used to determine if a call of the power save operation is currently allowed. If a value other that "allowed" is returned the caller may reckon that the power save operation will fail.</p> <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Alter power save setting is allowed for this BladeCenter</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power saving.</li> <li>• <b>"not-supported"</b> - Specifies that power saving is not supported for this BladeCenter.</li> <li>• <b>"under-group-control"</b> - The BladeCenter is under group control and cannot be individually altered.</li> </ul>

Table 44. BladeCenter object: energy management related additional properties (continued)

Name	Qualifier	Type	Description
<b>power-capping-state</b>	—	String Enum	Specifies the current power capping setting of the BladeCenter. Power capping limits peak power consumption of a system, and you can manage it by using the Set Power Cap operation. The possible settings include: <ul style="list-style-type: none"> <li>• <b>"disabled"</b> - Specifies not setting the power cap of the BladeCenter and not limiting the peak power consumption. This is the default setting.</li> <li>• <b>"enabled"</b> - Specifies capping all components of the BladeCenter available for power capping to limit the peak power consumption of the BladeCenter.</li> <li>• <b>"custom"</b> - Specifies individually configuring the components of the BladeCenter for power capping.</li> <li>• <b>"not-supported"</b> - Specifies that power capping is not supported for this BladeCenter.</li> <li>• <b>"not-entitled"</b> - Specifies that the server is not entitled to power capping.</li> </ul>
<b>power-cap-minimum</b>	—	Integer	Specifies the minimum value for the BladeCenter cap value in watts (W). This is a sum of the component minimum cap values.
<b>power-cap-maximum</b>	—	Integer	Specifies the maximum value for the BladeCenter cap value in watts (W). This is a sum of the component maximum cap values.
<b>power-cap-current</b>	—	Integer	Specifies the current cap value for the BladeCenter in watts (W). The current cap value indicates the power budget for the BladeCenter and is the sum of the component Cap values.
<b>power-cap-allowed</b>	—	String Enum	Should be used to determine if a call of the power capping operation is currently allowed. If a value other than <b>"allowed"</b> is returned the caller may reckon that the power capping operation will fail. <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Alter power capping setting is allowed for this BladeCenter</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power capping.</li> <li>• <b>"not-supported"</b> - Specifies that power capping is not supported for this BladeCenter.</li> <li>• <b>"under-group-control"</b> - The BladeCenter is under group control and cannot be individually altered.</li> </ul>
<b>ambient-temperature</b>	(mg)	Float	Specifies the input air temperature in degrees Celsius (°C) as measured by the system.
<b>exhaust-temperature</b>	—	Float	Specifies the exhaust air temperature in degrees Celsius (°C) as calculated by the system. This is useful in determining potential hot spots in the data center.

## Operations

### List BladeCenters in a Rack

The List BladeCenters in a Rack operation lists the BladeCenters that are mounted into the rack.

#### HTTP method and URI

GET /api/racks/{rack-id}/bladecenters

In this request, the URI variable *{rack-id}* is the object ID of the Rack object whose BladeCenters are to be obtained.

#### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) used to limit returned objects

#### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
blade-centers	Array of objects	Array of nested BladeCenter-info objects, described in the next table. If the rack does not have any BladeCenters mounted in it, an empty array is provided.

Each nested BladeCenter-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the BladeCenter object in the form <code>/api/bladecenters/{bladecenter-id}</code>
name	String	The <b>name</b> property of the BladeCenter object (for example, B.1)
status	String Enum	The <b>status</b> property of the BladeCenter object

#### Description

The **List BladeCenters in a Rack** operation lists the BladeCenters that are mounted in the rack. The object URI, name, and status are provided for each BladeCenter.

If the **name** query parameter is specified, then a BladeCenter is included in the list only if the name pattern matches the **name** property of the object.

A BladeCenter is included in the list only if the API user has object-access permission for that object. If the HMC is a manager of a BladeCenter, but the API user does not have permission to it, that object is omitted from the list, but no error status code results.

If the HMC does not manage any BladeCenters, an empty list is provided and the operation completes successfully.

#### Authorization requirements

This operation has the following authorization requirements:

- | • Object access permission to the CPC associated with the zBX in which the rack specified by the URI is contained (for a CPC-attached zBX).
- | • Object access permission to the zBX in which the rack specified by the URI is contained (for a zBX node).
- | • Object access permission to any BladeCenter object which is to be included in the result.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 142.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{rack-id}</i> does not designate a Rack object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/racks/04419e1e-a9db-11e0-8077-f0def1610d20/bladecenters HTTP/1.1
x-api-session: 5wksmtkmts30ajeohh0bn411fdzusmk1ld4jydd1des5t78aq
```

---

Figure 68. List BladeCenters in a Rack: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 15 Nov 2011 14:22:52 GMT
content-type: application/json;charset=UTF-8
content-length: 231
{
  "blade-centers": [
    {
      "name": "C.1",
      "object-uri": "/api/bladecenters/22e79848-3957-35dc-b88e-c661f9c8b680",
      "status": "operating"
    },
    {
      "name": "C.2",
      "object-uri": "/api/bladecenters/e3ee0adc-27c0-355e-93b9-ace8a3d2da15",
      "status": "operating"
    }
  ]
}
```

---

Figure 69. List BladeCenters in a Rack: Response

## List BladeCenters in a zBX

The **List BladeCenters in a zBX** operation lists the BladeCenters in a zBX.

### HTTP method and URI

```
GET /api/zbx/{zbx-id}/bladecenters
```

In this request, the URI variable *{zbx-id}* is the object ID of the zBX object whose BladeCenters are to be obtained.

#### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) used to limit returned objects

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
blade-centers	Array of objects	Array of nested BladeCenter-info objects, described in the next table. If the zBX does not have any BladeCenters, an empty array is provided.

Each nested BladeCenter-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the BladeCenter object in the form <code>/api/bladecenters/{bladecenter-id}</code>
name	String	The <b>name</b> property of the BladeCenter object (for example, B.1)
status	String Enum	The <b>status</b> property of the BladeCenter object

### Description

The **List BladeCenters in a zBX** operation lists the BladeCenters in the zBX. The object URI, name, and status are provided for each BladeCenter.

If the **name** query parameter is specified, then a BladeCenter is included in the list only if the name pattern matches the **name** property of the object.

A BladeCenter is included in the list only if the API user has object-access permission for that object. If the HMC is a manager of a BladeCenter, but the API user does not have permission to it, that object is omitted from the list, but no error status code results.

If the HMC does not manage any BladeCenters, an empty list is provided and the operation completes successfully.

### Authorization requirements

This operation has the following authorization requirement:

- | • Object access permission to the CPC to which the zBX specified by the URI is attached (for a CPC-attached zBX).
- | • Object access permission to the zBX specified by the URI (for a zBX node).
- | • Object access permission to any BladeCenter object which is to be included in the result.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”



Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{zbx-id}</i> does not designate a zBX object, or the API user does not have object access permission to the CPC with which it is associated.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/zbxes/291e385e-a9cd-11e0-8650-f0def1610d20/bladecenters HTTP/1.1
x-api-session: 5wksmtkmts30ajeohh0bn411fdzusmk1ld4jydd1des5t78aq
```

---

Figure 70. List BladeCenters in a zBX: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 15 Nov 2011 14:22:52 GMT
content-type: application/json;charset=UTF-8
content-length: 443
{
  "blade-centers": [
    {
      "name": "C.1",
      "object-uri": "/api/bladecenters/22e79848-3957-35dc-b88e-c661f9c8b680",
      "status": "operating"
    },
    {
      "name": "B.1",
      "object-uri": "/api/bladecenters/2ae200b3-fa8e-3db7-b34a-ec08780aac6",
      "status": "operating"
    },
    {
      "name": "C.2",
      "object-uri": "/api/bladecenters/e3ee0adc-27c0-355e-93b9-ace8a3d2da15",
      "status": "operating"
    },
    {
      "name": "B.2",
      "object-uri": "/api/bladecenters/f5bca837-bc0d-34e6-bcef-766411287439",
      "status": "operating"
    }
  ]
}
```

---

Figure 71. List BladeCenters in a zBX: Response

## Get BladeCenter Properties

The **Get BladeCenter Properties** operation retrieves the properties of a single BladeCenter object that is designated by its object ID.

## HTTP method and URI

**GET** /api/bladecenters/{*bladecenter-id*}

In this request, the URI variable *{bladecenter-id}* is the object ID of the BladeCenter object for which properties are to be obtained.

## Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Rack object as defined in the “Data model” on page 138. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

## Description

The **Get BladeCenter Properties** operation returns the current properties for the BladeCenter object specified by *{bladecenter-id}*.

On successful execution, all of the current properties as defined in “Data model” on page 138 for the BladeCenter object are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing BladeCenter object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the BladeCenter object specified by *{bladecenter-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{bladecenter-id}</i> does not designate an existing BladeCenter object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/bladecenters/ECEC1940F05B39EA9B3AEA5C1600AB1E HTTP/1.1
x-api-session: 5q07hx6jgp2ngn2cypq1zxot76sfwnzky0ih8nddd5hz6bpiue
```

---

Figure 72. Get BladeCenter Properties: Request

---

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 21 Jul 2011 17:49:02 GMT
content-type: application/json;charset=UTF-8
content-length: 948
{
  "acceptable-status": [
    "operating"
  ],
  "ambient-temperature": 18.0,
  "class": "bladecenter",
  "description": "Represents one BladeCenter",
  "exhaust-temperature": 25.5,
  "has-hardware-messages": false,
  "has-unacceptable-status": false,
  "is-locked": false,
  "location": "B01B",
  "machine-model": "HC1",
  "machine-serial": "KQZZXLF",
  "machine-type": "8852",
  "name": "B.2",
  "object-id": "ECEC1940F05B39EA9B3AEA5C1600AB1E",
  "object-uri": "/api/bladecenters/ECEC1940F05B39EA9B3AEA5C1600AB1E",
  "parent": "/api/racks/b434398a-a328-11e0-9b4a-001f163805d8",
  "power-cap-allowed": "allowed",
  "power-cap-current": 9561,
  "power-cap-maximum": 9561,
  "power-cap-minimum": 3473,
  "power-capping-state": "custom",
  "power-consumption": 1876,
  "power-rating": 9561,
  "power-save-allowed": "allowed",
  "power-saving": "high-performance",
  "status": "operating"
}

```

---

Figure 73. Get BladeCenter Properties: Response

## Activate BladeCenter

The **Activate BladeCenter** operation activates all blades in the BladeCenter chassis designated by its object ID. The BladeCenter chassis must be contained within a zBX node. This operation is asynchronous.

### HTTP method and URI

**POST** `/api/bladecenters/{bladecenter-id}/operations/activate`

In this request, the URI variable `{bladecenter-id}` is the object ID of the BladeCenter object to activate.

### Response body contents

Once the activation request is accepted, the response body contains a JSON object with the following field:

Field name	Type	Description
job-uri	String/ URI	URI that may be queried to retrieve activation status updates. The canonical URI path for BladeCenter activation status updates is of the form <code>/api/jobs/{job-id}</code> .

## | **Asynchronous result description**

| Once the activation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Activate BladeCenter** request.

| The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. (For more information, see “Query Job Status” on page 52. When the status of the job is “**complete**”, the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) that are set. (See the description that follows.) The **job-results** field is null for asynchronous BladeCenter activation jobs.

## | **Description**

| The **Activate BladeCenter** operation activates all blades in the BladeCenter object specified by *{bladecenter-id}*. Activation brings blades into a state of “**operating**”. If a blade is already powered on when the activation operation is requested, the blade is powered off and then brought to a state of “**operating**”. The order in which the blades of the BladeCenter are activated is unspecified. If the blade is a host to a virtualization application, then this application is activated also. See “Activating a Virtualization Host” on page 256 for more information.

| The URI path must designate an existing BladeCenter object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. The BladeCenter designated by the URI path must be contained within a zBX node (that is, a zBX that has a **type** of “**node**”) otherwise status code 400 (Bad Request) is returned. In addition to having object-access permission to the BladeCenter, the API user must also have permission to the **Activate** task, otherwise status code 403 (Forbidden) is returned. The BladeCenter must be in the correct state to perform the activate action, otherwise status code 409 (Conflict) is returned.

| When the BladeCenter activation job is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the activation job. See “Query Job Status” on page 52 for information on how to query job status. When the activate job has completed, an asynchronous result message is sent, with “Job status and reason codes” on page 149.

## | **Authorization requirements**

| This operation has the following authorization requirements:

- | • Object-access permission to the BladeCenter object designated by *{bladecenter-id}* and all its blades.
- | • Action/task permission to the **Activate** task.

## | **HTTP status and reason codes**

| On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 147.

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 45. Activate BladeCenter: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The object ID <i>{bladecenter-id}</i> designates a BladeCenter object that is not contained within a zBX node (that is, it contained within a zBX that does not have a <b>type</b> of “node”).
403 (Forbidden)	0	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID <i>{bladecenter-id}</i> does not designate an existing BladeCenter object, or the API user does not have object access permission to it.
409 (Conflict)	1	The object ID <i>{bladecenter-id}</i> designates a BladeCenter object that is not in the correct state (status) for performing the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Table 46. Activate BladeCenter: Job status and reason codes

HTTP error status code	Reason code	Description
200 (OK)	N/A	Activation completed successfully.
500 (Server Error)	100	BladeCenter activation failed.
	101	The BladeCenter activation job timed out.

## Example HTTP interaction

```
POST /api/bladecenters/bc000101-0000-0000-0000-000000000000/
  operations/activate HTTP/1.1
x-api-session: 2cdh1xvxjwh11wks636npd16nkvl dhij2ptku8une68o3ypep2
```

Figure 74. Activate BladeCenter: Request

```
202 Accepted
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Wed, 10 Sep 2014 19:20:02 GMT
content-type: application/json;charset=UTF-8
content-length: 60
{
  "job-uri": "/api/jobs/76c88c00-391f-11e4-8e38-02c00001303e"
}
```

Figure 75. Activate BladeCenter: Response

## Deactivate BladeCenter

The **Deactivate BladeCenter** operation deactivates a BladeCenter chassis specified by its object ID and all of its blades. The BladeCenter chassis must be contained within a zBX node. This operation is asynchronous.

### HTTP method and URI

**POST** `/api/bladecenters/{bladecenter-id}/operations/deactivate`

In this request, the URI variable `{bladecenter-id}` is the object ID of the BladeCenter chassis object to deactivate.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String/ URI	URI that may be queried to retrieve deactivation status updates. The canonical URI path for BladeCenter deactivation status updates is of the form <code>/api/jobs/{job-id}</code> .

### Asynchronous result description

Once the deactivation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Deactivate BladeCenter** request.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. (For more information, see “Query Job Status” on page 52. When the status of the job is “**complete**”, the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) that are set. (See the description that follows.) The **job-results** field is null for asynchronous zBX deactivation jobs.

### Description

The **Deactivate BladeCenter** operation deactivates the BladeCenter object specified by `{bladecenter-id}` and all of its blades. Deactivation powers off all blades after an orderly shutdown of any hardware and software activity running on the blades. The order in which the blades of the BladeCenter are deactivated is unspecified. If the blade is a host to a virtualization application, then this application is deactivated also. See “Deactivating a Virtualization Host” on page 257 for more information.

The URI path must designate an existing BladeCenter object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. The BladeCenter designated by the URI path must be contained within a zBX node (that is, a zBX that has a **type** of “**node**”) otherwise status code 400 (Bad Request) is returned. In addition to having object-access permission to the BladeCenter, the API user must also have permission to the **Deactivate** task, otherwise status code 403 (Forbidden) is returned. The BladeCenter must be in the correct state to perform the deactivate action, otherwise status code 409 (Conflict) is returned.

When the BladeCenter deactivation job is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the deactivation job. See “Query Job Status” on page 52 for information on how to query job status. When the deactivate job has completed, an asynchronous result message is sent, with “Job status and reason codes” on page 151.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the BladeCenter object specified by *{bladecenter-id}* and all its blades.
- Action/task permission to the **Deactivate** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 150.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 47. Activate BladeCenter: HTTP status and reason codes.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The object ID <i>{bladecenter-id}</i> designates a BladeCenter object that is not contained within a zBX node (that is, is contained within a zBX that does not have a <b>type</b> of “node”).
403 (Forbidden)	0	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID <i>{bladecenter-id}</i> does not designate an existing BladeCenter object, or the API user does not have object access permission to it.
409 (Conflict)	1	The object ID <i>{bladecenter-id}</i> designates a BladeCenter object that is not in the correct state (status) for performing the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Table 48. Deactivate BladeCenter: Job status and reason codes

HTTP error status code	Reason code	Description
200 (OK)	N/A	Deactivation completed successfully.
500 (Server Error)	100	BladeCenter deactivation failed.
	101	The BladeCenter deactivation job timed out.

## Example HTTP interaction

```
POST /api/bladecenters/bc000101-0000-0000-0000-000000000000/
  operations/deactivate HTTP/1.1
x-api-session: 48sppzb4o6vnnl1j4eguvj4vp8cm7z8oyyfdnz631oy6074vb9
```

Figure 76. Deactivate BladeCenter: Request

---

```
| 202 Accepted
| server: zSeries management console API web server / 2.0
| cache-control: no-cache
| date: Wed, 10 Sep 2014 19:18:22 GMT
| content-type: application/json;charset=UTF-8
| content-length: 60
| {
|   "job-uri": "/api/jobs/3b52ea4e-391f-11e4-bb52-02c00001303e"
| }
```

---

*Figure 77. Deactivate BladeCenter: Response*

## Inventory service data

Information about the BladeCenter chassis managed by the HMC can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for BladeCenter objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**bladecenter**" are to be included. An entry for a particular BladeCenter is included only if the API user has object-access permission to that object.

For each BladeCenter object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for "Get BladeCenter Properties" on page 145. That is, the data provided is the same as would be provided if a **Get BladeCenter Properties** operation were requested targeting this object.

## Sample inventory data

The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single BladeCenter. This object would appear as one array entry in the response array:



```

{
  "acceptable-status": [
    "operating"
  ],
  "additional-status": "unknown",
  "ambient-temperature": 18.5,
  "class": "bladecenter",
  "description": "Represents one BladeCenter",
  "exhaust-temperature": 24.5,
  "has-hardware-messages": false,
  "has-unacceptable-status": false,
  "is-locked": false,
  "location": "B10B",
  "machine-model": "HC1",
  "machine-serial": "KQZZXLF",
  "machine-type": "8852",
  "name": "B.1",
  "object-id": "ECEC1940F05B39EA9B3AEA5C1600AB1E",
  "object-uri": "/api/bladecenters/ECEC1940F05B39EA9B3AEA5C1600AB1E",
  "parent": "/api/racks/28bc03c8-7bc4-11e0-a905-001f163803de",
  "power-cap-allowed": "under-group-control",
  "power-cap-current": 9444,
  "power-cap-maximum": 9444,
  "power-cap-minimum": 4042,
  "power-capping-state": "disabled",
  "power-consumption": 1875,
  "power-rating": 9444,
  "power-save-allowed": "allowed",
  "power-saving": "high-performance",
  "status": "operating"
}

```

Figure 78. BladeCenter object: Sample inventory data

## Blade object

A Blade object represents a single blade in the zBX.

### Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, including the operational-status properties, with the following class-specific specialization:

Table 49. Blade object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>name</b>	(ro)	String (1-64)	The zManager assigned name of the blade <sup>1, 2</sup>
<b>description</b>	—	String	This property is not provided.
<b>object-uri</b>	—	String/URI	The canonical URI path for a zBX Blade object is of the form <code>/api/blades/{blade-id}</code> .
<b>parent</b>	—	String/URI	The canonical URI path for a parent BladeCenter object, in the form <code>/api/bladecenters/{bladecenter-id}</code> . <sup>2</sup>
<b>class</b>	—	String	The value <b>"blade"</b> .

Table 49. Blade object: base managed object properties specializations (continued)

Name	Qualifier	Type	Description of specialization
<b>status</b>	(sc)	String Enum	The status of the Blade object. Possible values: <ul style="list-style-type: none"> <li>• <b>"no-power"</b> - the blade is powered off</li> <li>• <b>"status-check"</b> - the blade and the console are not communicating</li> <li>• <b>"not-operating"</b> - the blade is powered on and communicating with the console but is not running work</li> <li>• <b>"stopped"</b> - operations on the blade are quiesced</li> <li>• <b>"definition-error"</b> - an error has occurred when loading the blade with the firmware</li> <li>• <b>"operating"</b> - blade is operating normally.</li> </ul>
<b>additional-status</b>	—	String Enum	This property is not provided.

**Notes:**

1. This **name** property is currently assigned based on the location of the Blade and is of the form **RackName.BladecenterName.BladeSlot** (e.g. B.1.01). However, it is possible that the API could be extended to allow this property to be writeable, in which case an API or User-Interface user could change the name to contain arbitrary data. Therefore, API client applications that are interested in determining the location of the blade should not rely on the contents and format of the **name** property, but rather obtain location information from the **location** property.
2. The location of a blade can be moved from slot to slot within a zBX. When a blade is moved to a different slot, the original URI of this blade is retained. Because the blade **name**, **parent** and **location** is based on the slot location of the blade, these three properties can change for a given URI when the blade is moved within the zBX. The relocation of a blade generates an inventory change notification to report the removal of the blade, then an inventory change notification to report the addition of the blade. Upon addition of the blade, expect the values of these properties to differ.

**Class specific additional properties**

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties:

Table 50. Blade object: class specific properties

Name	Qualifier	Type	Description	Supported "type" values
<b>type</b>	—	String Enum	Type of the blade. Values: <ul style="list-style-type: none"> <li>• <b>"power"</b> – the POWER7 blade</li> <li>• <b>"system-x"</b> – the System x blade</li> <li>• <b>"isaopt"</b> – the IBM Smart Analytics Optimizer blade</li> <li>• <b>"dpxi50z"</b> – the DataPower® XI50z blade.</li> </ul>	All
<b>processors</b>	—	Integer	number of zBX blade processors.	All
<b>cores-per-processor</b>	—	Integer	The number of processing cores provided by each processor of the zBX blade.	All
<b>memory-size</b>	—	Integer	memory size of the zBX blade specified in MB.	All
<b>machine-type</b>	—	String	4 characters.	All
<b>machine-model</b>	—	String	3 characters.	All
<b>machine-serial</b>	—	String	12 characters.	All
<b>location</b>	—	String	8 Characters (RxxBBSyy) <sup>1</sup> <ul style="list-style-type: none"> <li>• RxxB – BladeCenter vertical position in rack R</li> <li>• BSyy – physical slot of Blade Server</li> </ul>	All

Table 50. Blade object: class specific properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>isaopt-mode</b>	(pc)	String Enum	The mode of the ISAOPT blade. Values: <ul style="list-style-type: none"> <li>• "worker"</li> <li>• "coordinator"</li> </ul>	isaopt
<b>virtualization-host</b>	—	String/URI	The canonical URI path for the virtualization host being hosted by the blade.	power, system-x
<b>has-hardware-messages</b>	(pc)	Boolean	The blade has a hardware message (true) or the blade does not have a hardware message (false).	All
<b>licensed-features</b>	—	String	The Features that this DPXI50Z blade is licensed for. Each licensed feature in the string is delimited with commas (.). The blade must be Operating to retrieve this property; if it is not, null is returned instead.	dpxi50z
<b>iedn-interfaces</b>	(pc)	Array of iedn-interface Objects	Complex object defining the IEDN Interfaces configured to this DPXI50Z blade. The blade must be Operating to retrieve this property; if it is not, null is returned instead.	dpxi50z

**Note:** <sup>1</sup>The location of a blade can be moved from slot to slot within a zBX. When a blade is moved to a different slot, the original URI of this blade is retained. Because the blade **name**, **parent** and **location** is based on the slot location of the blade, these three properties can change for a given URI when the blade is moved within the zBX. The relocation of a blade generates an inventory change notification to report the removal of the blade, then an inventory change notification to report the addition of the blade. Upon addition of the blade, expect the values of these properties to differ.

**Energy management related additional properties:** In addition to the properties defined above, this object includes the following additional class-specific properties related to energy management. For further explanation of the various states involved, see Chapter 9, "Energy management," on page 179.

Table 51. Blade object: energy management related additional properties

Name	Qualifier	Type	Description
<b>power-rating</b>	—	Integer	Specifies the maximum power usage in watts (W) of this blade. This is a calculated value, as indicated by the electrical rating label or system rating plate of the blade.
<b>power-consumption</b>	(mg)	Integer	Specifies the current power consumption in watts (W) for this blade.
<b>power-saving</b>	—	String Enum	Specifies the current power saving setting of the blade. Power saving is used to reduce the energy consumption of a system and can be managed in the Set Power Saving operation. The possible settings include: <ul style="list-style-type: none"> <li>• "high-performance" - Specifies not reducing the power consumption of the blade. This is the default setting.</li> <li>• "low-power" - Specifies reducing the performance of the blade to allow for low power consumption.</li> <li>• "not-supported" - Specifies power saving is not supported for this blade.</li> <li>• "not-available" - Specifies <b>power-saving</b> property could not be read from this blade.</li> <li>• "not-entitled" - Specifies the server is not entitled to power saving.</li> </ul> Additional power savings modes may be introduced as zManager is extended to support additional power saving capabilities.

Table 51. Blade object: energy management related additional properties (continued)

Name	Qualifier	Type	Description
<b>power-saving-state</b>	—	String Enum	Specifies the power saving setting of the Blade set by the user. Please note that this property indicates the user setting and may not match the real state of the hardware compared to the <b>power-saving</b> property. The possible settings include: <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - Specifies not reducing the power consumption of the blade. This is the default setting.</li> <li>• <b>"low-power"</b> - Specifies low power consumption for all components of the blade enabled for power saving.</li> <li>• <b>"not-supported"</b> - Specifies power saving is not supported for this blade.</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power saving.</li> </ul>
<b>power-save-allowed</b>	—	String Enum	Should be used to determine if a call of the power save operation is currently allowed. If a value other than <b>"allowed"</b> is returned the caller may reckon that the power save operation will fail. <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Alter power save setting is allowed for this blade</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power saving.</li> <li>• <b>"not-supported"</b> - Specifies power saving is not supported for this blade.</li> <li>• <b>"under-group-control"</b> - The blade is under group control and cannot be individually altered.</li> </ul>
<b>power-capping-state</b>	—	String Enum	Specifies the current power capping setting of the blade. Power capping limits peak power consumption of a system, and you can manage it with the Set Power Cap operation. The possible settings include: <ul style="list-style-type: none"> <li>• <b>"disabled"</b> - Specifies not setting the power cap of the blade and not limiting the peak power consumption. This is the default setting.</li> <li>• <b>"enabled"</b> - Specifies limiting the peak power consumption of the blade to the current cap value.</li> <li>• <b>"not-supported"</b> - Specifies that power capping is not supported for this blade.</li> <li>• <b>"not-entitled"</b> - Specifies that the server is not entitled to power capping.</li> </ul>
<b>power-cap-minimum</b>	—	Integer	Specifies the minimum value for the blade cap value in watts (W).
<b>power-cap-maximum</b>	—	Integer	Specifies the maximum value for the blade cap value in watts (W).
<b>power-cap-current</b>	—	Integer	Specifies the current cap value for the blade in watts (W). The current cap value indicates the power budget for the blade.

Table 51. Blade object: energy management related additional properties (continued)

Name	Qualifier	Type	Description
<b>power-cap-allowed</b>	—	String Enum	<p>Should be used to determine if a call of the power capping operation is currently allowed. If a value other than <b>"allowed"</b> is returned the caller may reckon that the power capping operation will fail.</p> <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Alter power capping setting is allowed for this blade</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power capping.</li> <li>• <b>"not-supported"</b> - Specifies power capping is not supported for this blade.</li> <li>• <b>"under-group-control"</b> - The blade is under group control and cannot be individually altered.</li> </ul>

**IEDN interface nested object:** The IEDN (Intraensemble Data network) Interface object defines the VLAN sub-interface on the DPXI50Z blade:

Table 52. Blade object: IEDN interface nested object properties

Name	Type	Description
<b>name</b>	String (1-64)	Display name of the IEDN Interface. Valid characters are: a->z, A->Z, 0->9, underscore, dash and period. Periods can not be consecutive.
<b>is-active</b>	Boolean	Administrative State of the IEDN Interface configuration. Values: <ul style="list-style-type: none"> <li>• True – Enabled means that the Op-State of the IEDN interface will be up if there are no errors in the configuration.</li> <li>• False – Disabled means that the Op-State of the IEDN interface will always be down.</li> </ul>
<b>network-uri</b>	String/URI	<p>The canonical URI of the virtual network to which this IEDN interface is connected.</p> <p>The combination of the virtual network with the Ethernet Interface must be unique per blade. That is, each interface is associated with one virtual network, and the virtual network associated with one interface cannot also be associated with another.</p>
<b>ethernet-interface</b>	String Enum	<p>The physical Ethernet network interface, either <b>"eth7"</b> or <b>"eth9"</b>.</p> <p>The combination of the virtual network with the Ethernet Interface must be unique per blade.</p>
<b>ip-address</b>	String	<p>IP Address in either IPv4 or IPv6 format.</p> <p>No default value is provided.</p>
<b>net-mask</b>	String	<p>Network Mask associated with the IP Address in either IPv4 or IPv6 format.</p> <p>If an ipv4 ip-address is provided, valid values are 0-32, the default is 32.</p> <p>If an ipv6 ip-address is provided, valid values are 0-128, the default is 128.</p>
<b>secondary-ip-address</b>	Array of Strings	<p>List of secondary IP Addresses of either IPv4 or IPv6 format.</p> <p>No default value is provided.</p>

Table 52. Blade object: IEDN interface nested object properties (continued)

Name	Type	Description
<b>secondary-net-mask</b>	Array of Strings	List of secondary Network Masks associated with the Secondary IP Addresses of either IPv4 or IPv6 format.  If a secondary ipv4 ip-address is provided, valid values are 0-32, the default is 32.  If a secondary ipv6 ip-address is provided, valid values are 0-128, the default is 128.
<b>ipv4-gateway</b>	String	The IPv4 address to use for the default IPv4 gateway. No default value is provided.
<b>ipv6-gateway</b>	String	The IPv6 address to use for the default IPv6 gateway. No default value is provided.
<b>is-ipv6-auto-config-enabled</b>	Boolean	True if the IPv6 Auto configuration is enabled. Otherwise, false. When enabled, the interface is configured with a link-local secondary address. When disabled, uses the defined primary address.  Default is false.
<b>is-slaac</b>	Boolean	IPv6 Auto configuration must be enabled to utilize this option.  True if IPv6 Stateless Address is enabled. Otherwise, false. When enabled, the IPv6 address is obtained when connected to the network. When disabled, the interface uses the defined primary address.  Default is false.
<b>dad-transmit</b>	Integer	IPv6 Auto configuration must be enabled to utilize this option.  Specify the number of duplicate address detection (DAD) attempts to perform.  Default is 1.
<b>dad-transmit-delay</b>	Integer	IPv6 Auto configuration must be enabled to utilize this option.  When the number of duplicate address detection (DAD) attempts is greater than 1, specify the delay between attempts in milliseconds.  Default is 1000.
<b>mac-address</b>	String (17)	The MAC address represented as 6 groups of two lower-case hexadecimal digits separated by colons, e.g. "01:23:45:67:89:ab". Length is 17 characters. The MAC address uses the ensemble prefix.

## Operations

### List Blades in a BladeCenter

The List Blades in a BladeCenter operation lists all the blades in a BladeCenter.

#### HTTP method and URI

**GET** /api/bladecenters/{*bladecenter-id*}/blades

In this request, the URI variable *{bladecenter-id}* is the object ID of the BladeCenter object whose blades are to be obtained.

#### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) used to limit returned objects

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
blades	Array of objects	Array of nested blade-info objects, described in the next table. If the BladeCenter does not have any blades mounted in it, an empty array is provided.

Each nested blade-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the Blade object in the form <code>/api/blades/{blade-id}</code>
name	String	The <b>name</b> property of the Blade object (for example, B.1.01)
status	String Enum	The <b>status</b> property of the Blade object
type	String Enum	The type of the blade.

## Description

The **List Blades in a BladeCenter** operation lists the blades that are in the BladeCenter. The object URI, name, status, and type are provided for each blade.

If the **name** query parameter is specified, then a blade is included in the list only if the name pattern matches the **name** property of the object.

A blade is included in the list only if the API user has object-access permission for that object. If the HMC is a manager of a blade, but the API user does not have permission to it, that object is omitted from the list, but no error status code results.

If the HMC does not manage any blades, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission the BladeCenter object specified by the URI
- Object access permission to any Blade object are to be included in the result.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{bladecenter-id}</i> does not designate a BladeCenter object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/bladecenters/ECEC1940F05B39EA9B3AEA5C1600AB1E/blades HTTP/1.1
x-api-session: 5q07hx6jgp2ngn2cypq1zxot76sfwnzky0ih8nddd5hz6bpiue
```

---

Figure 79. List Blades in a BladeCenter: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 21 Jul 2011 17:49:02 GMT
content-type: application/json;charset=UTF-8
content-length: 386
{
  "blades": [
    {
      "name": "B.2.02",
      "object-uri": "/api/blades/938706AC3FF111D78B5600215EC0330E",
      "status": "operating",
      "type": "power"
    },
    {
      "name": "B.2.03",
      "object-uri": "/api/blades/B8210BC02D1E11E0AE81E41F13FE1430",
      "status": "operating",
      "type": "system-x"
    }
  ]
}
```

---

Figure 80. List Blades in a BladeCenter: Response

## List Blades in a zBX

The **List Blades in a zBX** operation lists all the blades in all the BladeCenters in a zBX. This operation has an optional parameter to specify the blade type to return in the list. If this parameter is omitted, all blades of all blade types are returned.

### HTTP method and URI

```
GET /api/zbx/{zbx-id}/blades
```

In this request, the URI variable *{zbx-id}* is the object ID of the zBX object whose blades are to be obtained.

### Query parameters:



Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) used to limit returned objects based on matching the objects <b>name</b> property
type	String	Optional	Filter string used to limit returned objects to those that have a matching <b>type</b> property. Value must be a valid blade <b>type</b> property value. To request that the results include blades of multiple types, specify this parameter multiple times. For example "type=power&type=isaopt".

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
blades	Array of objects	Array of nested blade-info objects, described the next table. If the zBX does not have any blades, an empty array is provided.

Each nested blade-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the Blade object in the form <code>/api/blades/{blade-id}</code>
name	String	The <b>name</b> property of the Blade object (for example, B.1.01)
status	String Enum	The <b>status</b> property of the Blade object
type	String Enum	The type of the blade

## Description

The **List Blades in a zBX** operation lists the blades in the zBX. The object URI, name, status, and type are provided for each blade.

A blade is included in the list only if the API user has object-access permission for that object. If the HMC is a manager of a blade, but the API user does not have permission to it, that object is omitted from the list, but no error status code results.

If the **name** query parameter is specified, the returned list is limited to those blades that have a **name** property matching at least one specified name filter pattern. If the **name** parameter is omitted, this filtering is not done.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid blade **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those blades that have a **type** property matching a specified type value. If the **type** parameter is omitted, this filtering is not done.

If both the **name** and **type** query parameters are specified, a blade is included in the list only if it passes both the name and type filtering criteria.

If the HMC does not manage any blades, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- | • Object access permission to the CPC to which the zBX specified by the URI is attached (for a CPC-attached zBX).
- | • Object access permission to the zBX specified by the URI (for a zBX node).
- | • Object access permission to any Blade object which is to be included in the result.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 161.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{zbx-id}</i> does not designate a zBX object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/zbx/54a9716c-a326-11e0-9469-001f163805d8/blades HTTP/1.1
x-api-session: 5q07hx6jgp2ngn2cypq1zxot76sfwnzky0ih8nddd5hz6bpiue
```

---

Figure 81. List Blades in a zBX: Request

---

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 21 Jul 2011 17:49:02 GMT
content-type: application/json;charset=UTF-8
content-length: 386
{
  "blades": [
    {
      "name": "B.2.02",
      "object-uri": "/api/blades/938706AC3FF111D78B5600215EC0330E",
      "status": "operating",
      "type": "power"
    },
    {
      "name": "B.2.03",
      "object-uri": "/api/blades/B8210BC02D1E11E0AE81E41F13FE1430",
      "status": "operating",
      "type": "system-x"
    }
  ]
}

```

---

Figure 82. List Blades in a zBX: Response

## Get Blade Properties

The **Get Blade Properties** operation retrieves the properties of a single Blade object that is designated by its object ID.

### HTTP method and URI

**GET** /api/blades/{blade-id}

In this request, the URI variable *{blade-id}* is the object ID of the Blade object for which properties are to be obtained.

### Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Rack object as defined in the “Data model” on page 138. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

The **Get Blade Properties** operation returns the current properties for the Blade object specified by *{blade-id}*.

On successful execution, all of the current properties as defined in “Data model” on page 153 for the Blade object are provided in the response body, and HTTP status code 200 (OK) is returned. If the blade is a DPXI50z blade and its current status is not **operating**, then null is returned as the value of the **licensed-features** and **iedn-interfaces** properties, but the operation otherwise succeeds.

The URI path must designate an existing Blade object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the Blade object specified by *{blade-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 163.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{blade-id}</i> does not designate an existing Blade object, the object ID designates a Blade object that is not of the correct type, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/blades/B8210BC02D1E11E0AE81E41F13FE1430 HTTP/1.1
x-api-session: 5q07hx6jgp2ngn2cypq1zxot76sfwnzky0ih8nddd5hz6bpiue
```

---

Figure 83. Get Blade Properties: Request

---

```
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 21 Jul 2011 17:49:03 GMT
content-type: application/json;charset=UTF-8
content-length: 989
{
  "acceptable-status": [
    "operating"
  ],
  "class": "blade",
  "cores-per-processor": 8,
  "has-hardware-messages": false,
  "has-unacceptable-status": false,
  "is-locked": false,
  "location": "B01BBS03",
  "machine-model": "AC1",
  "machine-serial": "06NL721",
  "machine-type": "7872",
  "memory-size": 131072,
  "name": "B.2.03",
  "object-id": "B8210BC02D1E11E0AE81E41F13FE1430",
  "object-uri": "/api/blades/B8210BC02D1E11E0AE81E41F13FE1430",
  "parent": "/api/bladecenters/ECEC1940F05B39EA9B3AEA5C1600AB1E",
  "power-cap-allowed": "allowed",
  "power-cap-current": 268,
  "power-cap-maximum": 500,
  "power-cap-minimum": 268,
  "power-capping-state": "enabled",
  "power-consumption": 241,
  "power-rating": 500,
  "power-save-allowed": "unknown",
  "power-saving": "not-supported",
  "processors": 2,
  "status": "operating",
  "type": "system-x",
  "virtualization-host": "/api/virtualization-hosts/931b25d6-82e1-11e0-b9e4-f0def10bff8d"
}
```

---

Figure 84. Get Blade Properties: Response for blade of type "system-x" (similar for type "power")

---

```
200 OK
x-request-id: Sx3 Rx13
x-client-correlator: 21
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 15 Nov 2011 14:22:55 GMT
content-type: application/json;charset=UTF-8
content-length: 899
{
  "acceptable-status": [
    "operating"
  ],
  "class": "blade",
  "cores-per-processor": 8,
  "has-hardware-messages": false,
  "has-unacceptable-status": false,
  "iedn-interfaces": [],
  "is-locked": false,
  "licensed-features": " MQ, TAM, DataGlue, JAXP-API, PKCS7-SMIME, SQL-ODBC,
    WebSphere-JMS, RaidVolume, iSCSI, LocateLED, AppOpt, zBX",
  "location": "C01BBS01",
  "machine-model": "4BX",
  "machine-serial": "6800442",
  "machine-type": "4195",
  "memory-size": 12288,
  "name": "C.2.01",
  "object-id": "eadb0be8-6fdb-11df-8f6a-e41f137a29e4",
  "object-uri": "/api/blades/eadb0be8-6fdb-11df-8f6a-e41f137a29e4",
  "parent": "/api/bladecenters/e3ee0adc-27c0-355e-93b9-ace8a3d2da15",
  "power-cap-allowed": "under-group-control",
  "power-cap-current": 444,
  "power-cap-maximum": 444,
  "power-cap-minimum": 144,
  "power-capping-state": "disabled",
  "power-consumption": 115,
  "power-rating": 444,
  "power-save-allowed": "unknown",
  "power-saving": "not-supported",
  "processors": 2,
  "status": "operating",
  "type": "dpxi50z"
}
```

---

Figure 85. Get Blade Properties: Response for blade of type "dpxi50z":

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 15 Nov 2011 14:22:55 GMT
content-type: application/json;charset=UTF-8
content-length: 773
{
  "acceptable-status": [
    "operating"
  ],
  "class": "blade",
  "cores-per-processor": 8,
  "has-hardware-messages": false,
  "has-unacceptable-status": false,
  "is-locked": false,
  "isaopt-mode": "coordinator",
  "location": "B10BBS01",
  "machine-model": "PEL",
  "machine-serial": "KQWZTNG",
  "machine-type": "7870",
  "memory-size": 49152,
  "name": "B.1.01",
  "object-id": "fa8d1eea-95ab-33cf-bf86-a03cc1346222",
  "object-uri": "/api/blades/fa8d1eea-95ab-33cf-bf86-a03cc1346222",
  "parent": "/api/bladecenters/2ae200b3-fa8e-3db7-b34a-ec08780aaac6",
  "power-cap-allowed": "under-group-control",
  "power-cap-current": 515,
  "power-cap-maximum": 500,
  "power-cap-minimum": 220,
  "power-capping-state": "disabled",
  "power-consumption": 181,
  "power-rating": 500,
  "power-save-allowed": "unknown",
  "power-saving": "not-supported",
  "processors": 2,
  "status": "operating",
  "type": "isaopt"
}

```

Figure 86. Get Blade Properties: Response for blade of type "isaopt":

## Activate a Blade

The **Activate a Blade** operation activates a Blade object designated by its object ID. This operation is asynchronous.

### HTTP method and URI

**POST** /api/blades/{blade-id}/operations/activate

In this request, the URI variable {blade-id} is the object ID of the Blade object to activate.

### Response body contents

Once the activation request is accepted, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String/ URI	URI that may be queried to retrieve activation status updates.

## Asynchronous result description

Once the activation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Activate a Blade** operation.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. (For more information, see “Query Job Status” on page 52.) When the status of the job is “**complete**”, the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) that are set. (See the description that follows.) The **job-results** field is null for asynchronous blade activation jobs.

## Description

The **Activate a Blade** operation activates the Blade object specified by *{blade-id}*. Activation brings the blade into a state of “operating”. If the blade is already powered on when the activation operation is requested, the blade is powered off and then brought to a state of “operating”. If the blade is a host to a virtualization application, then this application is activated also. See the “Activating a Virtualization Host” on page 256 for more information.

The URI path must designate an existing Blade object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Activate** task; otherwise, status code 403 (Forbidden) is returned.

When the blade activation job is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the activation job. See “Query Job Status” on page 52 for information on how to query job status. When the activate job has completed, an asynchronous result message is sent with job status and reason codes described in “Job status and reason codes” on page 169.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the Blade object specified by *{blade-id}*
- Action/task permission to the **Activate** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 167.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	0	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID <i>{blade-id}</i> does not designate an existing Blade object, or the API user does not have object access permission to it.
409 (Conflict)	1	The object ID <i>{blade-id}</i> designates a Blade object that is not in the correct state (status) for performing the requested operation.



Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status code	Reason code	Description
200 (OK)	N/A	Activation completed successfully.
500 (Server Error)	100	Blade activation failed.
500 (Server Error)	101	Blade activation job timed out.

## Example HTTP interaction

---

```
POST /api/blades/62f508a6-2d21-11e0-813b-e41f13fe15a8/operations/activate HTTP/1.1
x-api-session: 6c6s3b1p02v9x9az6brcic9q9dk34jjhxw1lw0squegu0ktia5k
```

---

Figure 87. Activate a Blade: Request

---

```
202 Accepted
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 05:45:45 GMT
content-type: application/json;charset=UTF-8
content-length: 60
{
  "job-uri": "/api/jobs/b8824300-1728-11e1-aea4-0010184c8334"
}
```

---

Figure 88. Activate a Blade: Response

## Deactivate a Blade

The **Deactivate a Blade** operation deactivates a Blade object designated by its object ID. This operation is asynchronous.

### HTTP method and URI

```
POST /api/blades/{blade-id}/operations/deactivate
```

In this request, the URI variable *{blade-id}* is the object ID of the Blade object to deactivate.

### Response body contents

Once the activation request is accepted, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String/ URI	URI that may be queried to retrieve deactivation status updates.

## Asynchronous result description

Once the deactivation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Deactivate a Blade** operation.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. (For more information, see “Query Job Status” on page 52). When the status of the job is “**complete**”, the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) that are set. (See the description that follows.) The **job-results** field is null for asynchronous blade activation jobs.

## Description

The **Deactivate a Blade** operation activates the Blade object specified by *{blade-id}*. Deactivation powers off the blade after an orderly shutdown of any hardware and software activity running on the blade. If the blade is a host to a virtualization application, then this application is deactivated also. See the “Deactivating a Virtualization Host” on page 257 for more information.

The URI path must designate an existing Blade object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Deactivate** task; otherwise, status code 403 (Forbidden) is returned. If the blade is not in the correct state to perform the deactivate, a status code of 409 (Conflict) is returned.

When the blade deactivation job is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the deactivation job. See “Query Job Status” on page 52 for information on how to query job status. When the activate job has completed, an asynchronous result message is sent with job status and reason codes described in “Job status and reason codes” on page 171.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Blade object specified by *{blade-id}*
- Action/task permission to the **Deactivate** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 169.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	0	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID <i>{blade-id}</i> does not designate an existing Blade object, or the API user does not have object access permission to it.
409 (Conflict)	1	The object ID <i>{blade-id}</i> designates a Blade object that is not in the correct state (status) for performing the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status code	Reason code	Description
200 (OK)	N/A	Deactivation completed successfully.
500 (Server Error)	100	Blade deactivation failed.
500 (Server Error)	101	Blade deactivation job timed out.

## Example HTTP interaction

---

```
POST /api/blades/62f508a6-2d21-11e0-813b-e41f13fe15a8/operations/deactivate HTTP/1.1
x-api-session: 6c6s3b1p02v9x9az6brcic9q9dk34jjhxw1lw0squegu0ktia5k
```

---

Figure 89. Deactivate a Blade: Request

---

```
202 Accepted
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 05:45:04 GMT
content-type: application/json;charset=UTF-8
content-length: 60
{
  "job-uri": "/api/jobs/a0247864-1728-11e1-aea4-0010184c8334"
}
```

---

Figure 90. Deactivate a Blade: Response

## Create IEDN Interface for a DataPower XI50z Blade

The **Create IEDN Interface for a DataPower XI50z Blade** operation adds an IEDN interface with the designated properties to the DataPower XI50z blade configuration.

### HTTP method and URI

```
POST /api/blades/{blade-id}/iedn-interface
```

In this request, the URI variable *{blade-id}* is the object ID of the DPXI50Z Blade object to which the IEDN interface is to be added.

### Request body contents

The request body contains the following writeable IEDN interface properties of the DPXI50Z Blade object that will be used to create the IEDN interface:

Field name	Type	Rqd/Opt	Description
name	String (1-64)	Required	The name of the IEDN interface. Valid characters are: a-z, A-Z, 0-9, underscore, dash and period. Periods cannot be consecutive.  Must be unique per blade.

Field name	Type	Rqd/Opt	Description
is-active	Boolean	Required	Administrative state of the IEDN interface configuration. Values: <ul style="list-style-type: none"> <li>• True – Enabled means that the Op-State of the IEDN interface will be up if there are no errors in the configuration</li> <li>• False – Disabled means that Op-State of the IEDN interface will always be down</li> </ul>
network-uri	String/URI	Required	The canonical URI of the virtual network to which this IEDN interface is connected.  The combination of the virtual network with the Ethernet interface must be unique per blade. That is, each interface is associated with one virtual network, and the virtual network associated with one interface cannot be associated with another.
ethernet-interface	String Enum	Required	The physical Ethernet network interface – either "eth7" or "eth9".  The combination of the virtual network with the Ethernet interface must be unique per blade.
ip-address	String	Optional	An IP address in either IPv4 or IPv6 format.  No default value is provided.
net-mask	String	Optional	The network mask associated with the IP address in either IPv4 or IPv6 format.  If an IPv4 <b>ip-address</b> is provided, valid values are 0-32; the default is 32.  If an IPv6 <b>ip-address</b> is provided, valid values are 0-128; the default is 128.
secondary-ip-address	List of Strings	Optional	A list of secondary IP addresses in either IPv4 or IPv6 format.  No default value is provided.
secondary-net-mask	List of Strings	Optional	A list of secondary network masks associated with the secondary IP addresses in either IPv4 or IPv6 format.  If a secondary IPv4 <b>ip-address</b> is provided, valid values are 0-32, the default is 32.  If a secondary IPv6 <b>ip-address</b> is provided, valid values are 0-128, the default is 128
ipv4-gateway	String	Optional	The IPv4 address to use for the default IPv4 gateway.  No default value is provided.
ipv6-gateway	String	Optional	The IPv6 address to use for the default IPv6 gateway  No default value is provided.
is-ipv6-auto-config-enabled	Boolean	Optional	True if the IPv6 auto configuration is enabled. Otherwise, false.  When enabled, the interface is configured with a link-local secondary address. When disabled, the interface uses the defined primary address.  Default is false.

Field name	Type	Rqd/Opt	Description
is-slaac	Boolean	Optional	IPv6 auto configuration must be enabled to utilize this option.  True if IPv6 stateless address is enabled. Otherwise, false.  When enabled, the IPv6 address is obtained when connected to the network. When disabled, the interface uses the defined primary address.  Default is false.
dad-transmit	Integer	Optional	IPv6 auto configuration must be enabled to utilize this option.  The number of duplicate address detection (DAD) attempts to perform.  Default is 1.
dad-transmit-delay	Integer	Optional	IPv6 auto configuration must be enabled to utilize this option.  When the number of duplicate address detection (DAD) attempts is greater than 1, the delay between attempts in milliseconds.  Default is 1000.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
element-uri	String/ URI	Canonical URI path of the IEDN interface object, in the form <code>/api/blades/{blade-id}/iedn-interfaces/{iedn-interface-id}</code> , where <code>{iedn-interface-id}</code> is the value of the <b>name</b> property.

## Description

The **Create IEDN Interface for a DataPower XI50z Blade** operation creates the IEDN interface for the DataPower XI50z blade as specified by the given properties. The DPXI50Z blade must be operating to perform this operation. Any properties identified as optional may be excluded from the request body. If an optional property is not found in the request body, its value will be set to its default value.

The add of the IEDN interface is permitted if the API user has object-access permission for that DataPower XI50z blade.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the DPXI50z Blade object specified by `{blade-id}`
- Action/task permission to the **zBX Blade Details** task.
- | • Object access permission to the CPC associated with the zBX containing the DPXI50z Blade object specified by the URI variable `{blade-id}` (for CPC-attached zBX objects).
- | • Object access permission to the zBX node containing the DPXI50z Blade object specified by the URI variable `{blade-id}` (for zBX node objects).
- | • Object access permission to the Virtual Network object specified by the **network-uri** field.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 173

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	0	The API user does not have task access authority to the <b>zBX Blade Details</b> task.
404 (Not Found)	1	The object ID <i>{blade-id}</i> does not designate an existing Blade object, or the API user does not have object access permission to it.
409 (Conflict)	1	The object ID <i>{blade-id}</i> designates a Blade object that is not in the correct state (status) for performing the requested operation. The blade must be operating to perform this operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Delete IEDN Interface for a DataPower XI50z Blade

The **Delete IEDN Interface for a DataPower XI50z Blade** operation removes the specified IEDN interface for a DataPower XI50z blade.

### HTTP method and URI

**DELETE** /api/blades/{blade-id}/iedn-interface/{iedn-interface-id}

#### URI variables

Variable	Description
<i>{iedn-interface-id}</i>	Element ID of the IEDN interface
<i>{blade-id}</i>	Object ID of the blade

### Description

The **Delete IEDN Interface for a DataPower XI50z Blade** operation deletes the IEDN interface that is defined for the DataPower XI50z blade. The DPXI50Z blade must be operating to perform this operation.

The IEDN interface is removed only if the API user has object-access permission for that DataPower XI50z blade.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the DPXI50z Blade object specified by *{blade-id}*
- Action/task permission to the **zBX Blade Details** task.
- Object access permission to the CPC associated with the zBX containing the DPXI50z Blade object specified by the URI variable *{blade-id}* (for CPC-attached zBX objects).

- Object access permission to the zBX node containing the DPXI50z Blade object specified by the URI variable *{blade-id}* (for zBX node objects).

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID <i>{blade-id}</i> does not designate an existing Blade object, or the API user does not have object access permission to it.
	2	The object ID <i>{iedn-interface-id}</i> does not exist on the HMC.
409 (Conflict)	1	The object ID <i>{blade-id}</i> designates a Blade object that is not in the correct state (status) for performing the requested operation. The blade must be operating to perform this operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Inventory service data

Information about the blades managed by the HMC can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for Blade objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by inventory class, implicitly via a containing category, or by default) that objects of the various blade type-specific inventory classes are to be included. An entry for a particular blade is included only if the API user has object-access permission to that blade and the applicable type-specific inventory class has been specified, as described in the following table:

Inventory class	Includes blades with “type” value
power-blade	power
system-x-blade	system-x
dpxi50z-blade	dpxi50z
isaopt-blade	isaopt

For each Blade object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for “Get Blade Properties” on page 163. That is, the data provided is the same as would be provided if a **Get Blade Properties** operation were requested targeting this object.

## Sample inventory data

The following fragments are examples of the JSON objects that would be included in the **Get Inventory** response to describe a single Blade object of a particular type. These objects would appear as array entries in the response array.

---

```
{
  "acceptable-status": [
    "operating"
  ],
  "class": "blade",
  "cores-per-processor": 8,
  "has-hardware-messages": false,
  "has-unacceptable-status": true,
  "is-locked": false,
  "location": "B10BBS13",
  "machine-model": "71Y",
  "machine-serial": "06C9FDA",
  "machine-type": "8406",
  "memory-size": 32768,
  "name": "B.1.13",
  "object-uri": "/api/blades/D5C5CB8A3F5511D78B5600215EC03866",
  "parent": "/api/bladecenters/ECEC1940F05B39EA9B3AEA5C1600AB1E",
  "power-cap-allowed": "under-group-control",
  "power-cap-current": 277,
  "power-cap-maximum": 382,
  "power-cap-minimum": 277,
  "power-capping-state": "disabled",
  "power-consumption": 151,
  "power-rating": 382,
  "power-save-allowed": "under-group-control",
  "power-saving": "high-performance",
  "processors": 8,
  "status": "status-check",
  "type": "power",
  "virtualization-host": "/api/virtualization-hosts/baa17718-2990-11e0-8d5b-001f163803de"
}
```

---

Figure 91. Activate a Blade: Sample inventory data for a blade of type "power"



---

```
{
  "acceptable-status": [
    "operating"
  ],
  "class": "blade",
  "cores-per-processor": 8,
  "has-hardware-messages": false,
  "has-unacceptable-status": true,
  "is-locked": false,
  "location": "B10BBS12",
  "machine-model": "AC1",
  "machine-serial": "06NL728",
  "machine-type": "7872",
  "memory-size": 131072,
  "name": "B.1.12",
  "object-uri": "/api/blades/62F508A62D2111E0813BE41F13FE15A8",
  "parent": "/api/bladecenters/ECEC1940F05B39EA9B3AEA5C1600AB1E",
  "power-cap-allowed": "under-group-control",
  "power-cap-current": 278,
  "power-cap-maximum": 519,
  "power-cap-minimum": 278,
  "power-capping-state": "disabled",
  "power-consumption": 249,
  "power-rating": 519,
  "power-save-allowed": "unknown",
  "power-saving": "not-supported",
  "processors": 2,
  "status": "status-check",
  "type": "system-x",
  "virtualization-host": "/api/virtualization-hosts/47d3a864-82e1-11e0-b9e4-f0def10bff8d"
}
```

---

Figure 92. Activate a Blade: Sample inventory data for a blade of type "system-x"



## Chapter 9. Energy management

Energy Management is a management task that is pervasive and spread across several components in the systems management stack. Each layer in the stack needs to implement two key functions:

- A set of management functions appropriate for this level at the stack. Energy management functions provided by lower layers can be used to implement these functions.
- Management interfaces are provided that allows management layers above to configure and control the energy management functions.

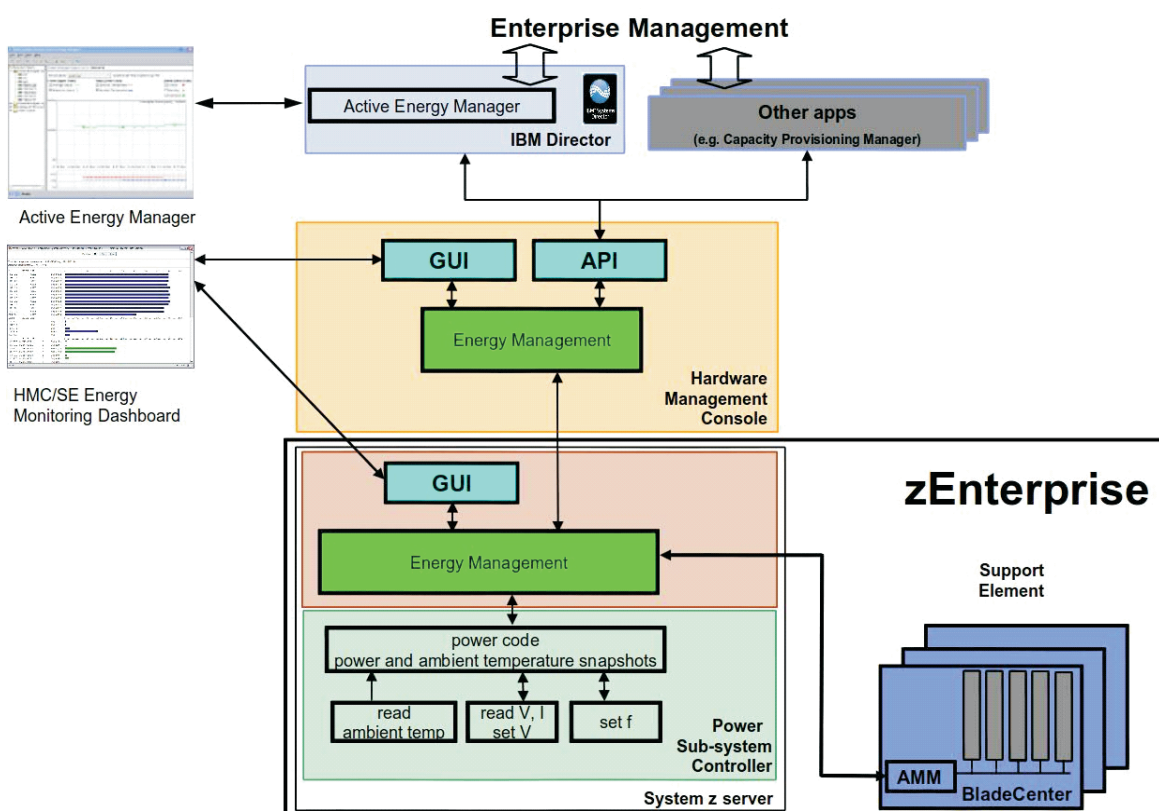


Figure 93. Energy management as applied throughout layers of enterprise management

To achieve this several pieces are needed:

### Power and thermal monitoring

"You can't improve what you don't measure" is a trivial engineering paradigm. Measuring energy consumption and the thermal environment is key for management. Energy Monitoring for z Systems was initially introduced with z9<sup>®</sup>. Starting with zEnterprise 196 the power consumption of attached BladeCenters is made available and factored into the presented system level power consumption.

### Energy control

Based on the measurement data - either for an individual system or aggregated for a group of servers or even a complete data center - analytics can be implemented. These can keep a watch on given limits or can identify optimizing potentials. At a system level energy control mechanisms will be provided to allow for changing energy consumption of a system. These energy controls can be categorized into two groups:

- **Power saving** - Power saving mechanisms are used to reduce the average energy consumption of a system. Through powering off components or reducing performance the power saving is typically achieved. For zEnterprise 196 the Static Power Savings Mode is implemented that reduces processor frequency and voltage for power saving purposes.
- **Power capping** - Power capping is a means to limit peak power consumption of a system. This is especially important in constraint data center environments. Today power and cooling allocation in data centers is usually done via the label power. This typically leads to a significant overprovisioning. Through power capping the power allocation for a system can be adjusted better to the real power consumption of a system and therefore more servers can be deployed within the same physical limits of their data center.

---

## Groups

A group is composed of an object that contains groups or another object and the object or objects it contains. For example, a group might be a zBX node that contains a BladeCenter with blades, as illustrated in Figure 94




Figure 94. Example of a zBX node group that contains a BladeCenter with blades

Figure 95 represents a CPC that contains a zCPC.




Figure 95. Example of a CPC group that contains a zCPC

Another example of a group might be a CPC (Ensemble) that contains a zCPC and BladeCenter, as shown in Figure 96.

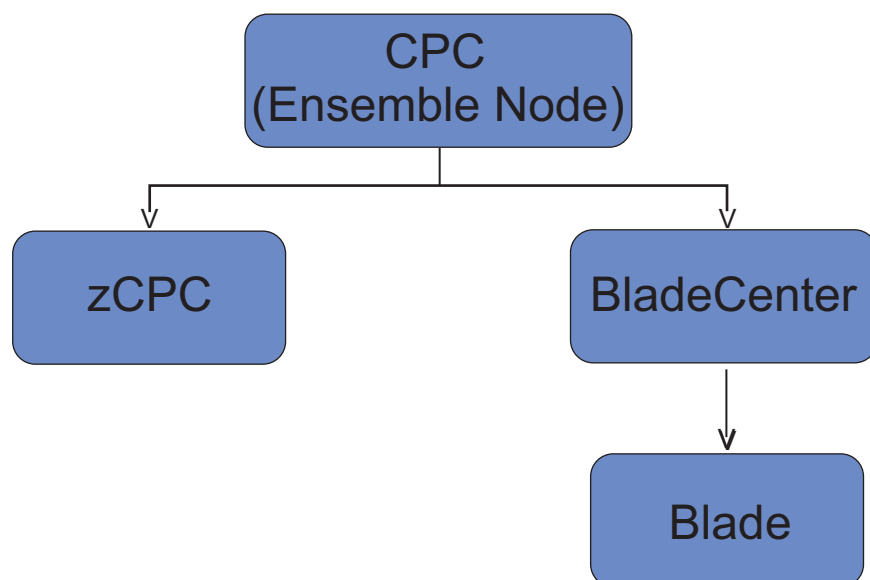


Figure 96. Example of a CPC (Ensemble) group that contains a zCPC and a BladeCenter

A fourth example is a CPC (Ensemble) that contains a zCPC as shown in Figure 97

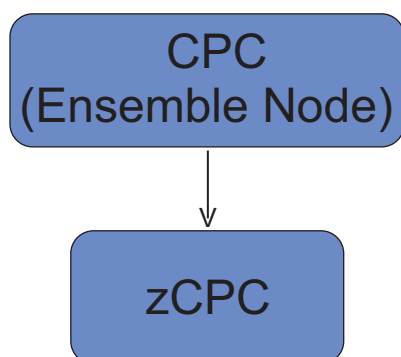


Figure 97. Example of a CPC (Ensemble) group that contains a zCPC

## Special states

In this chapter, the following states are used but the reasoning behind the states isn't always clear. So they are explained here in more detail:

### "custom"

Occurs only on groups and indicates that the group does not control the children. Clients are able to alter the children of a group individually.

### "under-group-control"

Occurs only on children and indicates that a group controls the state. When clients want to alter the state, the group must be set to "custom" first.

### "not-supported"

Indicates that the feature (either power saving or power capping) is currently not supported, possible reasons can be:

- Hardware does not support it → permanent

- Firmware level does not support it → can change after a firmware update
- The hardware is not powered on → can change after the device is powered on.

**"not-available"**

Couldn't read the state of the underlying hardware.

**"not-entitled"**

Indicates that the automate feature is not installed and so power saving and power capping is not allowed.

---

## Power saving

Power saving is a function that reduces the energy consumption of a system. Please note that power saving is only available if the Automate management enablement feature is installed. The possible settings include:

### High performance

The power consumption and performance of the object are not reduced. This is the default setting.

### Low power

The performance of the object is reduced to allow for low power consumption. When this setting is selected for CPC and BladeCenter objects, all components of the object enabled for power saving have reduced performance to allow for low power consumption. Use this setting to enable group power saving.

**Note:** You can only set the power saving setting of the zCPC to Low power one time per calendar day in an air cooled system. This power save property is set to Not Supported if the current zCPC power saving setting is High performance but the zCPC has already entered Low power once within the calendar day.

### Custom

Use Custom to disable group power saving and individually configure the components of the object for power saving.

**Note:** This setting is available only for CPC and BladeCenter objects.

## Group power saving

The following are important concepts regarding group power saving:

- Group power saving settings replace individual object settings--that is, the Power Saving setting of a CPC or BladeCenter supersedes the Power Saving setting of any object contained within the CPC or BladeCenter.
- You can enable group power saving by setting the Power Saving setting of the CPC or BladeCenter to Low power or High performance.
- You can change individual Power Saving settings only if the object is not under group power saving control.
- To disable group power saving without changing the individual Power Saving settings of the group members, change the Power Saving setting of the CPC or BladeCenter to Custom.

---

## Power capping

Please note that power capping is only available if the Automate management enablement feature is installed.

### Group capping

The following are important concepts regarding group power capping:

- Group caps replace individual object caps—that is, the Cap Value of a CPC or BladeCenter supersedes the power cap of any object contained within the CPC or BladeCenter.
- You can enable group capping by setting the Power Capping setting of the CPC or BladeCenter to Enabled.
- You can change individual Cap Values if the object is not under group capping control.
- If a CPC or BladeCenter contains an object that does not support power capping, the Power Rating is used in calculating the minimum power cap value for the group. The Power Rating can be found on the details window for an object.
- The maximum Cap Value for a group is the sum of the Power Rating of all Group objects.
- When a group component is powered off or removed, the group cap is redistributed to the remaining group components.
- To disable group capping without changing the individual power caps of the group members, change the Power Capping setting of the CPC or BladeCenter to Custom.

## Energy management operations summary

The following tables provide an overview of the operations provided. All POST operation were executed asynchronously and provide a query URI where the result of the request can be queried.

**Note:** The zCPC is not modeled as a full entity like CPC, BladeCenters or blades, because only energy management needs the zCPC to represent only z Systems hardware without zBX. That is the reason why all zCPC related operations are tied to the CPC.

Table 53. Energy management: operations summary

Operation name	HTTP method and URI path
“Set zBX (Node) Power Save” on page 184	POST /api/zbx/{zbx-id}/operations/set-power-save
“Set zBX (Node) Power Capping” on page 186	POST /api/zbx/{zbx-id}/operations/set-power-capping
“Set CPC Power Save” on page 189	POST /api/cpcs/{cpc-id}/operations/set-cpc-power-save
“Set CPC Power Capping” on page 191	POST /api/cpcs/{cpc-id}/operations/set-cpc-power-capping
“Get CPC Energy Management Data” on page 197	GET /api/cpcs/{cpc-id}/energy-management-data
“Set zCPC Power Save” on page 194	POST /api/cpcs/{cpc-id}/operations/set-zcpc-power-save
“Set zCPC Power Capping” on page 195	POST /api/cpcs/{cpc-id}/operations/set-zcpc-power-capping
“Set BladeCenter Power Save” on page 199	POST /api/bladecenters/{bladecenter-id}/operations/set-power-save
“Set BladeCenter Power Capping” on page 201	POST /api/bladecenters/{bladecenter-id}/operations/set-power-capping
“Set Blade Power Save” on page 204	POST /api/blades/{blade-id}/operations/set-power-save
“Set Blade Power Capping” on page 206	POST /api/blades/{blade-id}/operations/set-power-capping

Table 54. Energy management: URI variables

Variable	Description
{zbx-id}	Object ID of a zBX node

Table 54. Energy management: URI variables (continued)

Variable	Description
<i>{cpc-id}</i>	Object ID of a CPC
<i>{bladecenter-id}</i>	Object ID of a BladeCenter
<i>{blade-id}</i>	Object ID of a blade

## Energy Management for zBX (Node) object

### Data model

The data model for a zBX object includes some properties related to energy management when that zBX object represents a zBX node. These properties are described under the data model for the zBX object: "Energy Management Related Additional Properties" on page 85.

### Operations

#### Set zBX (Node) Power Save

Use the **Set zBX (Node) Power Save** operation to set the power save setting of a zBX node.

#### HTTP method and URI

**POST** `/api/zbx/{zbx-id}/operations/set-power-save`

In this request, the URI variable *{zbx-id}* is the object ID of the zBX node.

#### Request body contents

The request body is a JSON object with the following fields:

Name	Type	Rqd/Opt	Description
power-saving	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>"<b>high-performance</b>" - The power consumption and performance of the components of the zBX node are not reduced. This is the default setting.</li> <li>"<b>low-power</b>" - All components of the zBX node enabled for power saving.</li> <li>"<b>custom</b>" - Components may have their own settings changed individually. No component settings are actually changed when this mode is entered.</li> </ul>

#### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.



## Description

Use this operation to control the average energy consumption of a zBX node object designated by *{zbx-id}*, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated zBX node is under group control (see “Group capping” on page 182) or the **power-saving** property of the zBX node is set to **"not-supported"** or **"not-entitled"**. (See “Energy Management Related Additional Properties” on page 85 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-saving settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described in “HTTP status and reason codes.” After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to blade, BladeCenter, and zBX objects
- Action/task permission to the **Power Save** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 184.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The object ID in the URI <i>{zbx-id}</i> designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of <b>"node"</b> ).
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>{zbx-id}</i> ) does not designate an existing zBX object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the zBX hardware
	162	Communication error occurred while trying to access the zBX hardware
	163	An error occurred at one or more children

If the job reason code is 163, the **job-results** field provided by the **Query Job Status** operation will contain an object with the following fields:

Field name	Type	Description
errors	Object array	A list of error objects, containing detailed error information about errors occurred on children
at-least-one-operation-succeed	Boolean	<b>True</b> indicates that the operation was successful for at least one child.

Each error object has this structure:

Job status codes	Job reason code	Description
object-uri	String URI	The canonical URI path for a specific object where the error occurred
reason-code	Integer	Specify the specific error type, possible values are: <ul style="list-style-type: none"> <li>• 160 - A firmware error occurred while executing the operation</li> <li>• 161 - A hardware error occurred while performing the energy management operation</li> <li>• 162 - Communication error occurred while trying to access the hardware</li> </ul>
message	String	A non-localized message provided for development purposes only. Client applications should not display this message directly to the user.

## Set zBX (Node) Power Capping

Use the **Set zBX (Node) Power Capping** operation to set the power capping settings of a zBX node.

### HTTP method and URI

**POST** `/api/zbxs/{zbx-id}/operations/set-power-capping`

In this request, the URI variable `{zbx-id}` is the object ID of the zBX node.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
power-capping-state	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>• <b>"disabled"</b> - The power cap of the zBX node is not set and the peak power consumption is not limited. This is the default setting.</li> <li>• <b>"enabled"</b> - The peak power consumption of the zBX node is limited to the current cap value.</li> <li>• <b>"custom"</b> - Individually configure the components of the zBX node for power capping. No component settings are actually changed when this mode is entered.</li> </ul>
power-cap-current	Integer	Optional	Specifies the current cap value for the zBX node in watts (W). The current cap value indicates the power budget for the zBX node.  This field is only required if the <b>power-capping-state</b> field is set to <b>"enabled"</b> . The <b>power-cap-current</b> must be between <b>power-cap-minimum</b> and <b>power-cap-maximum</b> :  <b>power-cap-minimum</b> <= value <= <b>power-cap-maximum</b>

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.

## Description

Use this operation to limit the peak power consumption of a zBX node object designated by *{zbx-id}*, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated zBX is under group control (see “Group capping” on page 182) or the **power-capping-state** property of the zBX is set to **"not-supported"** or **"not-entitled"**. (See “Energy Management Related Additional Properties” on page 85 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-capping settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described in “HTTP status and reason codes” on page 188. After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to blade, BladeCenters, and zBX objects
- Action/task permission to the **Power Capping** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 187.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	240	The object ID in the URI ( <i>zbx-id</i> ) designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of “node”).
400 (Bad Request)	5	The <b>power-cap-current</b> field is not set, but <b>power-capping-state</b> field is set to “enabled”.
	7	The <b>power-cap-current</b> field contains a value that is not in the range <b>power-cap-minimum ... power-cap-maximum</b>
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>zbx-id</i> ) does not designate an existing zBX object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the zBX hardware
	162	Communication error occurred while trying to access the zBX hardware
	163	An error occurred at one or more children

If the job reason code is 163, the **job-results** field provided by the **Query Job Status** operation will contain an object with the following fields:

Field name	Type	Description
errors	Object array	A list of error objects, containing detailed error information about errors occurred on children
at-least-one-operation-succeed	Boolean	<b>True</b> indicates that the operation was successful for at least one child.

Each error object has this structure:

Job status codes	Job reason code	Description
object-uri	String URI	The canonical URI path for a specific object where the error occurred
reason-code	Integer	Specify the specific error type, possible values are: <ul style="list-style-type: none"> <li>• 160 - A firmware error occurred while executing the operation</li> <li>• 161 - A hardware error occurred while performing the energy management operation</li> <li>• 162 - Communication error occurred while trying to access the hardware</li> </ul>
message	String	A non localized message provided for development purposes only. Client applications should not display this message directly to the user.

## Energy Management for CPC object

The energy management for the CPC object represents all energy management for the CPC.

### Data model

The data model for a CPC object includes some properties related to energy management. These properties are described in “Energy management related additional properties” on page 750.

### Operations

#### Set CPC Power Save

Use the Set CPC Power Save operation to set the power save setting of a CPC.

#### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/set-cpc-power-save

In this request, the URI variable {cpc-id} is the object ID of the CPC.

#### Request body contents

The request body is a JSON object with the following fields:

Name	Type	Rqd/Opt	Description
power-saving	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - The power consumption and performance of the CPC are not reduced. This is the default setting.</li> <li>• <b>"low-power"</b> - Low power consumption for all components of the CPC enabled for power saving.</li> <li>• <b>"custom"</b> - Components may have their own settings changed individually. No component settings are actually changed when this mode is entered.</li> </ul>

#### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.

## Description

Use this operation to control the average energy consumption of a CPC object designated by *{cpc-id}*, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated CPC is under group control (see “Group capping” on page 182) or the **cpc-power-saving** property of the CPC is set to **"not-supported"** or **"not-entitled"**. (See “Energy management related additional properties” on page 750 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-saving settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described in “HTTP status and reason codes.” After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to all blade, BladeCenter, CPC and zCPC objects
- Action/task permission to the **Power Save** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 189.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully

Job status codes	Job reason code	Description
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the blade or z Systems hardware
	162	Communication error occurred while trying to access the blade or z Systems hardware
	163	An error occurred at one or more children

If the job reason code is 163, the **job-results** field provided by the **Query Job Status** operation will contain an object with the following fields:

Field name	Type	Description
errors	Object array	A list of error objects, containing detailed error information about errors occurred on children
at-least-one-operation-succeed	Boolean	<b>True</b> indicates that the operation was successful for at least one child.

Each error object has this structure:

Job status codes	Job reason code	Description
object-uri	String URI	The canonical URI path for a specific object where the error occurred
reason-code	Integer	Specify the specific error type, possible values are: <ul style="list-style-type: none"> <li>• 160 - A firmware error occurred while executing the operation</li> <li>• 161 - A hardware error occurred while performing the energy management operation</li> <li>• 162 - Communication error occurred while trying to access the hardware</li> </ul>
message	String	A non-localized message provided for development purposes only. Client applications should not display this message directly to the user.

## Set CPC Power Capping

Use the **Set CPC Power Capping** operation to set the power capping settings of a CPC.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/set-cpc-power-capping

In this request, the URI variable *{cpc-id}* is the object ID of the CPC.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
power-capping-state	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>• <b>"disabled"</b> - The power cap of the CPC is not set and the peak power consumption is not limited. This is the default setting.</li> <li>• <b>"enabled"</b> - The peak power consumption of the CPC is limited to the current cap value.</li> <li>• <b>"custom"</b> - Individually configure the components of the BladeCenter for power capping. No component settings are actually changed when this mode is entered.</li> </ul>
power-cap-current	Integer	Optional	Specifies the current cap value for the CPC in watts (W). The current cap value indicates the power budget for the CPC.  This field is only required if the <b>power-capping-state</b> field is set to <b>"enabled"</b> . The <b>power-cap-current</b> must be between <b>cpc-power-cap-minimum</b> and <b>cpc-power-cap-maximum</b> :  <b>cpc-power-cap-minimum</b> <= value <= <b>cpc-power-cap-maximum</b>

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.

## Description

Use this operation to limit the peak power consumption of a CPC object designated by *{cpc-id}*, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated CPC is under group control (see “Group capping” on page 182) or the **cpc-power-capping-state** property of the CPC is set to **"not-supported"** or **"not-entitled"**. (See “Energy management related additional properties” on page 750 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-capping settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described in “HTTP status and reason codes” on page 193. After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to all blade, BladeCenter, CPC and zCPC objects
- Action/task permission to the **Power Capping** task.



## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 192.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
400 (Bad Request)	7	The <b>power-cap-current</b> field contains a value that is not in the range <b>cpc-power-cap-minimum ... cpc-power-cap-maximum</b>
	5	The <b>power-cap-current</b> field is not set, but <b>power-capping-state</b> field is set to "enabled".
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the blade or z Systems hardware
	162	Communication error occurred while trying to access the blade or z Systems hardware
	163	An error occurred at one or more children

If the job reason code is 163, the **job-results** field provided by the **Query Job Status** operation will contain an object with the following fields:

Field name	Type	Description
errors	Object array	A list of error objects, containing detailed error information about errors occurred on children
at-least-one-operation-succeed	Boolean	<b>True</b> indicates that the operation was successful for at least one child.

Each error object has this structure:

Job status codes	Job reason code	Description
object-uri	String URI	The canonical URI path for a specific object where the error occurred
reason-code	Integer	Specify the specific error type, possible values are: <ul style="list-style-type: none"> <li>• 160 - A firmware error occurred while executing the operation</li> <li>• 161 - A hardware error occurred while performing the energy management operation</li> <li>• 162 - Communication error occurred while trying to access the hardware</li> </ul>
message	String	A non localized message provided for development purposes only. Client applications should not display this message directly to the user.

## Set zCPC Power Save

Use the **Set zCPC Power Save** operation to set the power save settings of the zCPC portion of a CPC.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/set-zcpc-power-save

In this request, the URI variable {cpc-id} is the object ID of the CPC.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
power-saving	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - The power consumption and performance of the zCPC are not reduced. This is the default setting.</li> <li>• <b>"low-power"</b> - Low power consumption for all components of the zCPC enabled for power saving.</li> </ul>

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.

## Description

Use this operation to control the average energy consumption of a zCPC portion of the CPC {cpc-id}, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated zCPC is under group control (see "Group capping" on page 182) or the **zcpc-power-saving** property of the zCPC is set to **"not-supported"** or **"not-entitled"**. (See "Energy management related additional properties" on page 750 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-saving settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body

includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described in “HTTP status and reason codes.” After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to all blade, BladeCenter, CPC and zCPC objects
- Action/task permission to the **Power Save** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 194.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the z Systems hardware
	162	Communication error occurred while trying to access the z Systems hardware

## Set zCPC Power Capping

Use the **Set zCPC Power Capping** operation to set the power capping settings of the zCPC portion of a CPC.

## HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/set-zcpc-power-capping

In this request, the URI variable {cpc-id} is the object ID of the CPC.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
power-capping-state	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>• <b>"disabled"</b> - The power cap of the zCPC is not set and the peak power consumption is not limited. This is the default setting.</li> <li>• <b>"enabled"</b> - The peak power consumption of the zCPC is limited to the current cap value.</li> </ul>
power-cap-current	Integer	Optional	Specifies the current cap value for the zCPC in watts (W). The current cap value indicates the power budget for the zCPC.  This field is only required if the <b>power-capping-state</b> field is set to <b>"enabled"</b> . The <b>power-cap-current</b> must be between <b>zcpc-power-cap-minimum</b> and <b>zcpc-power-cap-maximum</b> :  <b>zcpc-power-cap-minimum</b> <= value <= <b>zcpc-power-cap-maximum</b>

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.

## Description

Use this operation to limit the peak power consumption of a zCPC object designated by {cpc-id}, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated zCPC is under group control (see "Group capping" on page 182) or the **zcpc-power-capping-state** property of the zCPC is set to **"not-supported"** or **"not-entitled"**. (See "Energy management related additional properties" on page 750 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-capping settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described below. After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to all blade, BladeCenter, CPC and zCPC objects

- Action/task permission to the **Power Capping** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 196.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
400 (Bad Request)	7	The <b>power-cap-current</b> field contains a value that is not in the range <b>zpcp-power-cap-minimum ... zpcp-power-cap-maximum</b>
	5	The <b>power-cap-current</b> field is not set, but <b>power-capping-state</b> field is set to "enabled".
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the z Systems hardware
	162	Communication error occurred while trying to access the z Systems hardware

## Get CPC Energy Management Data

Use the **Get CPC Energy Management Data** operation to retrieve all energy management related data in one single call.

### HTTP method and URI

```
GET /api/cpcs/{cpc-id}/energy-management-data
```

In this request, the URI variable *{cpc-id}* is the object ID of the CPC.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
objects	Array of objects	An array of nested em-data objects containing the energy management data. The format of each nested object is given in the next table.

Each nested em-data object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The canonical URI path of the specific object to which this em-data object pertains.
object-id	String	Object-id property of the specific object to which this em-data object pertains.
class	String	The type of the specific object to which this em-data object pertains.
properties	Object	Nested object containing the energy management properties for the object identified by the object-uri field, as described in the data model section for objects of the type indicated by the class field.
error-occurred	Boolean	If true, indicates that an error occurred while querying the data for the object specified by the object-uri. As a consequence the property could be null or incomplete.

## Description

The **Get CPC Energy Management Data** is a convenience operation to allow a client to retrieve all energy management related data for a CPC in a single request rather than invoking several requests to retrieve this data.

Note that this operation returns data for a child object of the designated CPC only if the API user has object-access permission to that object. Children objects for which the API user does not have access are omitted from the response and no error is indicated.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by the request, and for any children objects for which data is to be returned.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.

HTTP error status code	Reason code	Description
404 (Not Found)	1	The object ID in the URI ( <i>cpc-id</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

---

## Energy Management for BladeCenter object

### Data model

The data model for a BladeCenter object includes some properties related to energy management. These properties are described in Chapter 8, “zBX infrastructure elements,” on page 77, under the data model for the BladeCenter object: “Energy Management Related Additional Properties” on page 139.

### Operations

#### Set BladeCenter Power Save

Use the **Set BladeCenter Power Save** operation to set the power save setting of a BladeCenter.

#### HTTP method and URI

**POST** /api/bladecenters/{*bladecenter-id*}/operations/set-power-save

In this request, the URI variable *bladecenter-id* is the object ID of the BladeCenter.

#### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
power-saving	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - The power consumption and performance of the BladeCenter are not reduced. This is the default setting.</li> <li>• <b>"low-power"</b> - Low power consumption for all components of the BladeCenter enabled for power saving.</li> <li>• <b>"custom"</b> - Components may have their own settings changed individually. No component settings are actually changed when this mode is entered.</li> </ul>

#### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.

## Description

Use this operation to control the average energy consumption of a BladeCenter object designated by *{bladecenter-id}*, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated BladeCenter is under group control (see “Group capping” on page 182) or the **power-saving** property of the BladeCenter is set to **"not-supported"** or **"not-entitled"**. (See “Energy Management Related Additional Properties” on page 139 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-saving settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described in “HTTP status and reason codes.” After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to all blade, BladeCenter, CPC and zCPC objects
- Action/task permission to the **Power Save** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 199.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>{bladecenter-id}</i> ) does not designate an existing BladeCenter object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.



## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the BladeCenter hardware
	162	Communication error occurred while trying to access the BladeCenter hardware
	163	An error occurred at one or more children

If the job reason code is 163, the **job-results** field provided by the **Query Job Status** operation will contain an object with the following fields:

Field name	Type	Description
errors	Object array	A list of error objects, containing detailed error information about errors occurred on children
at-least-one-operation-succeed	Boolean	<b>True</b> indicates that the operation was successful for at least one child.

Each error object has this structure:

Job status codes	Job reason code	Description
object-uri	String URI	The canonical URI path for a specific object where the error occurred
reason-code	Integer	Specify the specific error type, possible values are: <ul style="list-style-type: none"> <li>• 160 - A firmware error occurred while executing the operation</li> <li>• 161 - A hardware error occurred while performing the energy management operation</li> <li>• 162 - Communication error occurred while trying to access the hardware</li> </ul>
message	String	A non localized message provided for development purposes only. Client applications should not display this message directly to the user.

## Set BladeCenter Power Capping

Use the **Set BladeCenter Power Capping** operation to set the power capping settings of a BladeCenter.

### HTTP method and URI

**POST** `/api/bladecenters/{bladecenter-id}/operations/set-power-capping`

In this request, the URI variable `{bladecenter-id}` is the object ID of the BladeCenter.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
power-capping-state	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>• <b>"disabled"</b> - The power cap of the BladeCenter is not set and the peak power consumption is not limited. This is the default setting.</li> <li>• <b>"enabled"</b> - Capping all components of the BladeCenter available for power capping to limit the peak power consumption of the BladeCenter.</li> <li>• <b>"custom"</b> - Individually configure the components of the BladeCenter for power capping. No component settings are actually changed when this mode is entered.</li> </ul>
power-cap-current	Integer	Optional	Specifies the current cap value for the BladeCenter in watts (W). The current cap value indicates the power budget for the BladeCenter.  This field is only required if the <b>power-capping-state</b> field is set to <b>"enabled"</b> . The <b>power-cap-current</b> must be between <b>power-cap-minimum</b> and <b>power-cap-maximum</b> :  <b>power-cap-minimum</b> <= value <= <b>power-cap-maximum</b>

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.

## Description

Use this operation to limit the peak power consumption of a BladeCenter object designated by *{bladecenter-id}*, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated BladeCenter is under group control (see "Group capping" on page 182) or the **power-capping-state** property of the BladeCenter is set to **"not-supported"** or **"not-entitled"**. (See "Energy Management Related Additional Properties" on page 139 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-capping settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described in "HTTP status and reason codes" on page 203. After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to all blade, BladeCenter, CPC and zCPC objects

- Action/task permission to the **Power Capping** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 202.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
400 (Bad Request)	7	The <b>power-cap-current</b> field contains a value that is not in the range <b>power-cap-minimum ... power-cap-maximum</b>
	5	The <b>power-cap-current</b> field is not set, but <b>power-capping-state</b> field is set to "enabled".
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>{bladecenter-id}</i> ) does not designate an existing BladeCenter object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the BladeCenter hardware
	162	Communication error occurred while trying to access the BladeCenter hardware
	163	An error occurred at one or more children

If the job reason code is 163, the **job-results** field provided by the **Query Job Status** operation will contain an object with the following fields:

Field name	Type	Description
errors	Object array	A list of error objects, containing detailed error information about errors occurred on children
at-least-one-operation-succeed	Boolean	<b>True</b> indicates that the operation was successful for at least one child.

Each error object has this structure:

Job status codes	Job reason code	Description
object-uri	String URI	The canonical URI path for a specific object where the error occurred
reason-code	Integer	Specify the specific error type, possible values are: <ul style="list-style-type: none"> <li>• 160 - A firmware error occurred while executing the operation</li> <li>• 161 - A hardware error occurred while performing the energy management operation</li> <li>• 162 - Communication error occurred while trying to access the hardware</li> </ul>
message	String	A non localized message provided for development purposes only. Client applications should not display this message directly to the user.

## Energy Management for Blade object

### Data model

The data model for a Blade object includes some properties related to energy management. These properties are described in Chapter 8, “zBX infrastructure elements,” on page 77, under the data model for the Blade object: “Energy management related additional properties” on page 155.

### Operations

#### Set Blade Power Save

Use the **Set Blade Power Save** operation to set the power save setting of a blade.

#### HTTP method and URI

**POST** /api/blades/{blade-id}/operations/set-power-save

In this request, the URI variable {blade-id} is the object ID of the blade.

#### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
power-saving	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - The power consumption and performance of the blade are not reduced. This is the default setting.</li> <li>• <b>"low-power"</b> - Low power consumption for all components of the blade enabled for power saving.</li> </ul>

#### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.

## Description

Use this operation to control the average energy consumption of a Blade object designated by *{blade-id}*, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated blade is under group control (see “Group capping” on page 182) or the **power-saving** property of the blade is set to **"not-supported"** or **"not-entitled"**. (See “Energy management related additional properties” on page 155 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-saving settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described below. After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to all blade, BladeCenter, CPC and zCPC objects
- Action/task permission to the **Power Save** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 204.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>{blade-id}</i> ) does not designate an existing Blade object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully

Job status codes	Job reason code	Description
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the blade hardware
	162	Communication error occurred while trying to access the blade hardware

## Set Blade Power Capping

Use the **Set Blade Power Capping** operation to set the power capping settings of a blade.

### HTTP method and URI

**POST** `/api/blades/{blade-id}/operations/set-power-capping`

In this request, the URI variable `{blade-id}` is the object ID of the blade.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
power-capping-state	String Enum	Required	The possible settings are: <ul style="list-style-type: none"> <li>"disabled" - The power cap of the blade is not set and the peak power consumption is not limited. This is the default setting.</li> <li>"enabled" - Capping all components of the blade available for power capping to limit the peak power consumption of the blade.</li> </ul>
power-cap-current	Integer	Optional	Specifies the current cap value for the blade in watts (W). The current cap value indicates the power budget for the blade.  This field is only required if the <b>power-capping-state</b> field is set to "enabled". The <b>power-cap-current</b> must be between <b>power-cap-minimum</b> and <b>power-cap-maximum</b> :  <b>power-cap-minimum</b> <= value <= <b>power-cap-maximum</b>

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
job-uri	String	URI of the asynchronous job that may be queried to retrieve status updates for action initiated by this operation.

## Description

Use this operation to limit the peak power consumption of a Blade object designated by `{bladecenter-id}`, or to remove a power consumption limit for this object. You can closely manage power allocations within the physical limits of your data center.

This operation will always fail if the designated blade is under group control (see "Group capping" on page 182) or the **power-capping-state** property of the blade is set to "not-supported" or "not-entitled".

(See “Energy management related additional properties” on page 155 for details on this property.) In addition, this operation is only available if the ensemble is functioning at the Automate management enablement level.

The action to change the power-capping settings occurs asynchronously. If the request is accepted, an asynchronous job is initiated and an HTTP Status code of 202 (Accepted) is returned. The response body includes a URI that may be queried to retrieve the status of the asynchronous job. See the description of the **Query Job Status** operation for information on how to query job status. When the asynchronous job has completed, an asynchronous result message is sent, with Job status and reason codes described below. After completion, the **Query Job Status** operation may be used to retrieve the completion results.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to all blade, BladeCenter, CPC and zCPC objects
- Action/task permission to the **Power Capping** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 206.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
400 (Bad Request)	7	The <b>power-cap-current</b> field contains a value that is not in the range <b>power-cap-minimum ... power-cap-maximum</b>
	5	The <b>power-cap-current</b> field is not set, but <b>power-capping-state</b> field is set to "enabled".
403 (Forbidden)	1	The user is not authorized to access the object or perform this task.
	3	The server is not entitled to perform energy management.
404 (Not Found)	1	The object ID in the URI ( <i>blade-id</i> ) does not designate an existing Blade object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The operation cannot be performed because the object designated by the request URI is not in the correct state.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status codes	Job reason code	Description
200 (OK)	N/A	Operation executed successfully

<b>Job status codes</b>	<b>Job reason code</b>	<b>Description</b>
500 (Server Error)	160	A firmware error occurred while executing the operation
	161	A hardware error occurred while performing the operation on the blade hardware
	162	Communication error occurred while trying to access the blade hardware



## Chapter 10. Virtualization management

zManager provides facilities for running virtualized computing systems, termed Virtual Servers, on top of system-firmware-managed hosting environments that coordinate the physical system resources that the virtual servers share. These hosting environments upon which the virtual servers run are known as Virtualization Hosts.

The Virtualization Host APIs provide access to the set of virtualization hosts managed by an Ensemble. PowerVM, x Hyp, z/VM and PR/SM™ virtualization hosts are supported. APIs exist to query virtualization hosts, retrieve and update select properties of virtualization hosts, and perform operations on virtualization hosts, such as activate and deactivate.

Within z Systems, a virtual server can be described as a container for an operating system. It could be a logical partition on an IBM POWER7 Blade, a virtual machine on an IBM System x Blade, a virtual machine defined under z/VM, or a z Systems logical partition (LPAR). It is created in cooperation with the virtualization host on the hardware involved. The virtual server container includes the definition of processor resources, network interfaces and storage devices it will access.

### Virtualization host operations summary

The following tables provide an overview of the virtualization host operations provided.

Table 55. Virtualization management - virtualization host: operations summary

Operation name	HTTP method and URI path
"List Virtualization Hosts of a zBX (Node)" on page 222	GET /api/zbx/{zbx-id}/virtualization-hosts
"List Virtualization Hosts of a Node" on page 224	GET /api/ensembles/{ensemble-id}/nodes/{node-id}/virtualization-hosts
"List Virtualization Hosts of an Ensemble" on page 227	GET /api/ensembles/{ensemble-id}/virtualization-hosts
"List Virtualization Hosts of a CPC" on page 229	GET /api/cpcs/{cpc-id}/virtualization-hosts
"Get Virtualization Host Properties" on page 232	GET /api/virtualization-hosts/{virt-host-id}
"Update Virtualization Host Properties" on page 237	POST /api/virtualization-hosts/{virt-host-id}
"List Virtual Switches" on page 239	GET /api/virtualization-hosts/{virt-host-id}/virtual-switches
"Get Virtual Switch Properties" on page 240	GET /api/virtualization-host/{virt-host-id}/virtual-switches/{virtual-switch-id}
"Create IEDN Virtual Switch" on page 243	POST /api/virtualization-hosts/{virt-host-id}/virtual-switches/operations/add-iedn
"Create QDIO Virtual Switch" on page 246	POST /api/virtualization-hosts/{virt-host-id}/virtual-switches/operations/add-qdio
"Get Switch Controllers" on page 249	GET /api/virtualization-hosts/{virt-host-id}/operations/get-switch-controllers
"Update Virtual Switch" on page 251	POST /api/virtualization-hosts/{virt-host-id}/virtual-switches/{virtual-switch-id}

Table 55. Virtualization management - virtualization host: operations summary (continued)

Operation name	HTTP method and URI path
"Delete Virtual Switch" on page 255	DELETE /api/virtualization-hosts/{virt-host-id}/virtual-switches/{virtual-switch-id}

Table 56. Virtualization management- virtualization host: URI variables

Variable	Description
{zbx-id}	Object ID of a zBX object
{node-id}	Object ID of a Node object
{ensemble-id}	Object ID of an Ensemble object
{cpc-id}	Object ID of a CPC object
{virt-host-id}	Object ID of a Virtualization Host object
{virtual-switch-id}	Element ID of a virtual switch

## Virtual server operations summary

The following tables provide an overview of the virtual server operations provided.

Table 57. Virtualization management - virtual server: operations summary

Operation name	HTTP method and URI path
"List Virtual Servers of a zBX (Node)" on page 287	GET /api/zbxs/{zbx-id}/virtual-servers
"List Virtual Servers of a Node" on page 289	GET /api/ensembles/{z-ensemble-id}/nodes/{node-id}/virtual-servers
"List Virtual Servers of an Ensemble" on page 291	GET /api/ensembles/{ensemble-id}/virtual-servers
"List Virtual Servers of a CPC" on page 294	GET /api/cpcs/{cpc-id}/virtual-servers
"List Virtual Servers of a Virtualization Host" on page 297	GET /api/virtualization-hosts/{virt-host-id}/virtual-servers
"Create Virtual Server" on page 299	POST /api/virtualization-hosts/{virt-host-id}/virtual-servers
"Delete Virtual Server" on page 304	DELETE /api/virtual-servers/{virtual-server-id}
"Get Virtual Server Properties" on page 306	GET /api/virtual-servers/{virtual-server-id}
"Update Virtual Server Properties" on page 314	POST /api/virtual-servers/{virtual-server-id}
"Create Network Adapter" on page 319	POST /api/virtual-servers/{virtual-server-id}/network-adapters
"Delete Network Adapter" on page 322	DELETE /api/virtual-servers/{virtual-server-id}/network-adapters/{network-adapter-id}
"Get Network Adapter Properties" on page 323	GET /api/virtual-servers/{virtual-server-id}/network-adapters/{network-adapter-id}
"Update Network Adapter" on page 325	POST /api/virtual-servers/{virtual-server-id}/network-adapters/{network-adapter-id}

Table 57. Virtualization management - virtual server: operations summary (continued)

Operation name	HTTP method and URI path
“Reorder Network Adapter” on page 328	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/operations/reorder-network-adapters
“Create Virtual Disk” on page 330	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/virtual-disks
“Delete Virtual Disk” on page 334	DELETE /api/virtual-servers/{ <i>virtual-server-id</i> }/virtual-disks/{ <i>virtual-disk-id</i> }
“Get Virtual Disk Properties” on page 336	GET /api/virtual-servers/{ <i>virtual-server-id</i> }/virtual-disks/{ <i>virtual-disk-id</i> }
“Update Virtual Disk Properties” on page 338	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/virtual-disks/{ <i>virtual-disk-id</i> }
“Reorder Virtual Disks” on page 341	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/operations/reorder-virtual-disks
“Activate Virtual Server” on page 343	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/operations/activate
“Deactivate Virtual Server” on page 345	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/operations/deactivate
“Mount Virtual Media” on page 348	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/operations/mount-virtual-media
“Mount Virtual Media Image” on page 350	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/operations/mount-virtual-media-image
“Unmount Virtual Media” on page 352	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/operations/unmount-virtual-media
“Migrate Virtual Server” on page 354	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/operations/migrate
“Initiate Virtual Server Dump” on page 356	POST /api/virtual-servers/{ <i>virtual-server-id</i> }/operations/initiate-dump

Table 58. Virtualization management - virtual server: URI variables

Variable	Description
{ <i>zbx-id</i> }	Object ID of a zBX object
{ <i>node-id</i> }	Object ID of a Node object
{ <i>ensemble-id</i> }	Object ID of an ensemble object
{ <i>cpc-id</i> }	Object ID of a CPC object
{ <i>virt-host-id</i> }	Object ID of a Virtualization Host object
{ <i>virtual-server-id</i> }	Object ID of a Virtual Server object
{ <i>network-adapter-id</i> }	Element ID of a network adapter object
{ <i>virtual-disk-id</i> }	Element ID of a virtual disk object

## Virtualization Host object

- | A Virtualization Host object represents a single z Systems virtualization host.

### Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, with the following class-specific specialization:

Table 59. Virtualization Host object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/ URI	The canonical URI path for a Virtualization Host object is of the form <code>/api/virtualization-hosts/{object-id}</code> .
<b>name</b> <sup>1</sup>	(ro)	String (1-64)	The display name of the virtualization host. Not writeable. The name is derived from the name of the underlying hosting environment (e.g. blade).
<b>class</b>	—	String (19)	The class will always be <b>"virtualization-host"</b> .
<b>parent</b>	—	String/ URI	The canonical URI path of the node that manages the virtualization host.
<b>status</b>	(sc)	String Enum	<p>The current operational status of the virtualization host.</p> <p>Valid values common to all types:</p> <ul style="list-style-type: none"> <li>• <b>"operating"</b> - indicates virtualization host is running normally</li> <li>• <b>"not-operating"</b> - indicates the virtualization host is deactivated, in the process of loading, or has a error</li> <li>• <b>"status-check"</b> - indicates the support element is not communicating with the CPC or zBX node</li> <li>• <b>"not-communicating"</b> - indicates the HMC is not communicating with the support element.</li> </ul> <p>Valid values for type <b>"power-vm"</b>, <b>"x-hyp"</b>:</p> <ul style="list-style-type: none"> <li>• <b>"definition-error"</b> - indicates that the specified zBX Blade does not match the characteristics of the installed zBX Blade</li> <li>• <b>"no-power"</b> - indicates the zBX Blade powered off</li> <li>• <b>"stopped"</b> - indicates the Support Element with the zBX Blade is stopped.</li> </ul> <p>Valid values for type <b>"zvm"</b>:</p> <ul style="list-style-type: none"> <li>• <b>"exceptions"</b> - indicates at least one CP is operating, but at least one CP is not operating.</li> <li>• <b>"not-activated"</b> - indicates the image has not been activated.</li> </ul> <p>Valid values for type <b>"prsm"</b>:</p> <ul style="list-style-type: none"> <li>• <b>"degraded"</b> - indicates the CPC or zBX node is operating but some hardware is not available.</li> <li>• <b>"exceptions"</b> - indicates that at least one Central Processor (CP) is operating, but at least one CP is not operating.</li> <li>• <b>"no-power"</b> - indicates that the CPC or zBX node is powered off.</li> <li>• <b>"service"</b> - indicates that CPs are in service status.</li> <li>• <b>"service-required"</b> - indicates that the CPC or zBX node is still operating but is using the last redundant part of a particular type.</li> </ul>
<b>additional status</b>	—		This property is not provided.
<p><b>Note:</b> <sup>1</sup>The location of a blade can be moved from slot to slot within a zBX. When a blade is moved to a different slot, the original URI of this blade and Virtualization Host object is retained. Since the name of the blade and virtualization host is based on the slot location of the blade, the <b>name</b> property can change for a given URI when the blade is moved within the zBX. The relocation of a blade generates inventory change notifications to report the removal of the blade and the corresponding Virtualization Host object, then inventory change notifications to report the addition of the blade and the virtualization host. Upon addition of the blade and virtualization host, expect the value of the <b>name</b> property to differ.</p>			

## Class specific additional properties

In addition to the properties defined in included schemas, this object includes the following additional type-specific properties:

Table 60. Virtualization Host object: class specific additional properties

Name	Qualifier	Type	Description	Supported "type" values
<b>type</b>	—	String Enum	Type of the virtualization host.  Values: <ul style="list-style-type: none"> <li>"<b>power-vm</b>"- a virtualization host running on a POWER7 blade</li> <li>"<b>x-hyp</b>" - The canonical URI path for the System x blade that hosts the virtualization host.</li> <li>"<b>zvm</b>" - a z/VM operating system instance that is participating as an ensemble-managed virtualization host</li> <li>"<b>prsm</b>" - the virtualization host representation of a CPC or zBX node</li> </ul>	All
<b>hosting-environment</b>	—	String/URI	The hosting environment (cpc, image, blade) of the virtualization host.  Value based on type: <ul style="list-style-type: none"> <li>"<b>power-vm</b>" - The canonical URI path for the POWER7 blade that hosts the virtualization host.</li> <li>"<b>x-hyp</b>" - The canonical URI path for the System x blade that hosts the virtualization host.</li> <li>"<b>zvm</b>" - The canonical URI path for the PR/SM virtual server that is hosting the z/VM virtualization host.</li> <li>"<b>prsm</b>" - The canonical URI path for the CPC or zBX node that is the PR/SM virtualization host.</li> </ul>	All
<b>virtual-server-shutdown-timeout</b>	(w)(pc)	Integer	Amount of time, in seconds, to allow a virtual server to shut down. After the elapsed time has passed, the virtual server will be forcefully terminated.  The value may be -1 to indicate to wait "forever" or any integer value between 0 and 86400 to specify an exact wait time in seconds. If a value over 86400 is specified, it will be accepted and a value of 86400 will be set for this property.  This value will be not be used if the virtual server supports the <b>shutdown-timeout-source</b> property and it is set to any value other than " <b>virtualization-host</b> ".	power-vm, x-hyp
<b>auto-start-virtual-servers</b>	(w)(pc)	Boolean	Automatically start any virtual servers configured to start when the virtualization host is started.	power-vm, x-hyp
<b>total-memory</b>	(pc)	Integer	The total amount of memory the hypervisor has (in MB).	All
<b>memory-increment-in-megabytes</b>	—	Integer	The minimum "increment", in megabytes, that is required when setting the memory size for a virtual server.	power-vm, x-hyp, zvm

Table 60. Virtualization Host object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>minimum-memory-size-for-virtual-server</b>	—	Integer	The minimum memory size, in megabytes, that can be configured for a virtual server.	power-vm, x-hyp, zvm
<b>maximum-memory-size-for-virtual-servers</b>	—	Integer	The maximum memory size, in megabytes, that can be configured for a virtual server.	power-vm, x-hyp, zvm
<b>maximum-allowed-dedicated-processors</b>	—	Integer	The maximum number of dedicated processors that can be configured for a virtual server.	power-vm, zvm
<b>maximum-allowed-virtual-processors</b>	—	Integer	The maximum number of virtual processors that can be configured for a virtual server.	power-vm, x-hyp, zvm
<b>maximum-allowed-processing-units</b>	—	Float	The maximum number of processing units that can be configured for a virtual server. Processing units are a unit of measure for shared processing power across one or more virtual processors. One shared processing unit on one virtual processor accomplishes approximately the same work as one dedicated processor.	power-vm
<b>maximum-allowed-ide-devices</b>	—	Integer	The maximum number of IDE devices that can be configured for a virtual server	x-hyp
<b>mixed-mode-boot-restriction</b>	—	Boolean	Indicates if a virtual disk boot restriction exists for the virtualization host's virtual servers that would prevent booting from " <b>virtio</b> " virtual disks when both " <b>ide</b> " and " <b>virtio</b> " virtual disks are defined. In such a case, the virtual server will always boot from an " <b>ide</b> " virtual disk.	x-hyp
<b>gpmp-available-version</b>	—	String	The version of the Guest Platform Management Provider (GPMP) ISO that is available and can be mounted (using the Mount Virtual Media operation) on virtual servers hosted on this virtualization host. This zManager-provided ISO contains installation images that are used to install the GPMP into a virtual server. The value is null if the available GPMP version information cannot be obtained.	power-vm, x-hyp, zvm
<b>inband-monitoring-supported</b>	—	Boolean	If true, in-band monitoring is supported for the Virtualization Host, allowing it to gather performance metrics on its virtual servers.	power-vm, x-hyp
<b>supported-keyboard-languages</b>	—	Array of Strings	List of keyboard languages supported by the hypervisor for graphical console connections.  The value is <b>null</b> if no key map is supported or necessary or if graphical console is not supported.  See the Virtual Server object's <b>keyboard-language</b> property for the description of an array element value.	x-hyp

Table 60. Virtualization Host object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>is-bridge-capable</b>	—	Boolean	<p>If the virtualization host is capable of having virtual switches with bridge ports, the value is true. For true to be returned the following requirements must be met; otherwise, false is returned:</p> <ul style="list-style-type: none"> <li>• The <b>os-level</b> property of the PR/SM virtual server running the z/VM virtualization host is "621" or greater</li> <li>• The <b>se-version</b> property of the CPC or zBX node object is 2.11.1 or higher.</li> </ul>	zvm
<b>feature-list</b>	—	Array of String Enums	<p>Optional features or behaviors supported by this virtualization host. The presence of one of the Enum values described below indicates that this virtualization host provides the described feature. If the virtualization host has no such optional features, an empty array is provided.</p> <p>Possible <b>feature-list</b> values for an "x-hyp" virtualization host:</p> <ul style="list-style-type: none"> <li>• <b>"boot-sequence-mixed-disk-restriction"</b> - Indicates that a virtual disk boot restriction exists for the virtualization host's virtual servers that would prevent booting from <b>"virtio"</b> virtual disks when both <b>"ide"</b> and <b>"virtio"</b> virtual disks are defined. In such a case, the virtual server will always boot from an <b>"ide"</b> virtual disk.</li> <li>• <b>"boot-sequence-network-priority-restriction"</b> - Indicates that a virtual network adapter boot restriction exists for the virtualization host's virtual servers that would force <b>"network-adapter"</b> to be the first entry in the <b>boot-sequence</b> property.</li> <li>• <b>"boot-sequence-virtual-disk-enforcement"</b> - Indicates that a virtualization host's virtual server's <b>boot-sequence</b> property must contain an entry for <b>"virtual-disk"</b> if at least one virtual disk has been configured.</li> <li>• <b>"boot-sequence-network-adapter-enforcement"</b> - Indicates that a virtualization host's virtual server's <b>boot-sequence</b> property must contain an entry for <b>"network-adapter"</b> if at least one network adapter has been configured.</li> <li>• <b>"boot-sequence-virtual-media-enforcement"</b> - Indicates that a virtualization host's virtual server's <b>boot-sequence</b> property must contain an entry for <b>"virtual-media"</b>.</li> </ul> <p>Possible <b>feature-list</b> values for an "x-hyp" or "power" virtualization host:</p> <ul style="list-style-type: none"> <li>• <b>"virtual-server-shutdown-timeout-override-support"</b> - Indicates that a virtualization host's virtual servers support their own <b>shutdown-timeout-source</b> and <b>shutdown-timeout</b> property.</li> </ul>	All

Table 60. Virtualization Host object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
cpu-shares-supported	—	Boolean	If true, CPU share management is supported for the virtualization host, allowing it to manage the processor shares on its virtual servers.	x-hyp

## Virtual switch objects

A virtual switch is a special type of guest LAN that provides external LAN connectivity through an OSA-Express device without the need for a routing virtual machine.

**Note:** Some properties are only valid when mutable prerequisite properties have specific values. When such properties are not valid, their value is **null**. For instance a **iedn-virtual-switch** object's router property value is **null** when the **layer-mode** value is **"eth"**.

**iedn-virtual-switch object:** The **iedn-virtual-switch** object is a virtual-switch object that defines the IEDN virtual switch of a virtualization-host object of type **"zvm"**.

Table 61. *iedn-virtual-switch* object properties

Name	Qualifier	Type	Description
<b>element-id</b>	—	String	Unique ID for the virtual switch within the scope of the containing virtualization host. The element-id is actually the name of the virtual switch. Once the virtual switch is created, the name cannot be changed.
<b>element-uri</b>	—	String/URI	The canonical URI path for the virtual switch is of the form <code>/api/virtualization-hosts/{virt-host-id}/virtual-switches/{element-id}</code> , where <code>{virt-host-id}</code> is the object-id of the virtualization host.
<b>parent</b>	—	String/URI	The URI path of the z/VM Virtualization Host that hosts this virtual switch.
<b>class</b>	—	String (14)	Always <b>"virtual-switch"</b>
<b>type</b>	—	String Enum	The virtual switch type. Always <b>"iedn"</b>
<b>name</b>	—	String (1-8)	The unique name identifying the virtual switch.  1-8 alphanumeric uppercase characters or any of the following characters: "@# \$"
<b>switch-status</b>	—	String	Virtual switch status, read-only.
<b>layer-mode</b>	—	String Enum	Indicates the transport for the virtual switch is Ethernet or IP. Values: <ul style="list-style-type: none"> <li><b>"eth"</b> : This type is Data Link (Layer 2) based, where the Ethernet frame is used as the point of reference for source and destination Media Access Control (MAC) addresses in transporting Ethernet frames on the LAN.</li> <li><b>"ip"</b>: This type is Network (Layer 3) based, where the IP package is used as the point of reference for source and destination IP addresses in transporting IP packets on the LAN.</li> </ul>



Table 61. iedn-virtual-switch object properties (continued)

Name	Qualifier	Type	Description
<b>router</b>	—	String Enum	The router type.  Prerequisites: layer-mode is <b>"ip"</b> .  Values: <ul style="list-style-type: none"> <li>• <b>"none"</b> - Indicates that the OSA-Express device identified will not act as a router to the virtual switch. If a datagram is received at this device for an unknown IP address, the datagram will be discarded.</li> <li>• <b>"primary"</b> - Indicates that the OSA-Express device identified will act as a primary router to the virtual switch. If a datagram is received at this device for an unknown IP address, the datagram will be passed to the virtual switch. The only time to set PRIMARY for a virtual switch is if you have a guest attached to the virtual switch that is providing a routing function for systems attached to another network.</li> </ul>
<b>queue-size</b>	—	Integer	Queue storage in megabytes: decimal number between 1 and 8
<b>ip-timeout</b>	—	Integer	IP timeout in seconds: decimal number between 1 and 240
<b>is-connect-uplinks</b>	—	Boolean	If the switch is connected to uplinks, the value is true.
<b>real-uplinks</b>	—	Array of real-uplink objects	A list of zero to three real-uplink objects.
<b>is-use-any-available-controller</b>	—	Boolean	True if the switch uses any available controller.
<b>controllers</b>	—	Array of strings	A list of names of switch controllers. Prerequisites: <b>is-use-any-available-controller</b> is false.
<b>bridge-value-type</b>	(w)	String Enum	Whether the bridge port is a primary or secondary port. There can only be 1 primary and up to 4 secondary bridge ports.  Values: <ul style="list-style-type: none"> <li>• <b>"primary"</b> - Bridge port is the primary port</li> <li>• <b>"secondary"</b> - Bridge port is a secondary port.</li> </ul> Prerequisite: This field is applied to a vSwitch that has a valid <b>bridge-device-number</b> .  If not specified when the bridge port is defined, then the default is <b>"secondary"</b> . This field can be modified when the <b>bridge-connection-status</b> is <b>"disconnected"</b> .
<b>bridge-device-number</b>	(w)	String (1-4)	The bridge device number. The following values may be specified: <ul style="list-style-type: none"> <li>• The bridge real device number, 1~4 hex digits.</li> </ul> Prerequisites: The <b>is-bridge-capable</b> property of the virtualization host is true and the <b>layer-mode</b> property of the virtual switch is <b>"eth"</b> .  This property is required to define a device for a bridge port and must be a nonblank, valid device number.

Table 61. *iedn-virtual-switch object properties (continued)*

Name	Qualifier	Type	Description
<b>bridge-connection-status</b>	(w)	String Enum	Whether the bridge port is connected or in standby state.  Prerequisites: The <b>is-bridge-capable</b> property of the virtualization host is true and the virtual switch has a valid <b>bridge-device-number</b> .  Values: <ul style="list-style-type: none"> <li>• <b>"connected"</b> - Bridge port is connected</li> <li>• <b>"disconnected"</b> - Bridge port is disconnected.</li> <li>• <b>"standby"</b> - Bridge port is in standby state</li> </ul>
<b>mtu-size-enforcement</b>	(w)	String Enum	How the virtual switch MTU size is enforced.  Prerequisites: The <b>is-bridge-capable</b> property of the virtualization host is true.  Values: <ul style="list-style-type: none"> <li>• <b>"external"</b> - The MTU size will be set to the size used by the OSA adapter. This is the default value.</li> <li>• <b>"off"</b> - MTU enforcement is disabled</li> <li>• <b>"user-defined"</b> - The MTU size is specified in the <b>mtu-size</b> field</li> </ul>
<b>mtu-size</b>	(w)	Integer	The MTU specification represents the acceptable MTU size, in bytes, enforced by the virtual switch. MTU size is a decimal number between 512-65535.  Prerequisite: This field is applied to a virtual switch where the <b>is-bridge-capable</b> property of the virtualization host is true and the <b>mtu-size-enforcement</b> property of the virtual switch is <b>"user-defined"</b> .

**real-uplink object:** This is the embedded object definition for a virtual switch. For IEDN virtual switch, the **switch-uri** property is required. For QDIO virtual switch, the **switch-uri** property is ignored.

Table 62. *real-uplink object properties*

Name	Type	Description
<b>switch-uri</b>	String	The <b>element-uri</b> of the network-adapter-prsm object in use by the z/VM virtual switch to provide a connection between a virtual network and a physical network.  Prerequisites: <b>type</b> in the parent switch is <b>"iedn"</b> .
<b>device-number</b>	String (1-8)	The uplink real device number, 1-8 hex digits. The value could be either <b>"none"</b> or a string with the regular expression pattern <code>[0-9a-fA-F]{1,4}(.P[0-9][0-9])?</code> . If the optional <code>(.P[0-9][0-9])</code> part is not specified, a default string <code>".P00"</code> will be used.

## qdio-virtual-switch object

The qdio-virtual-switch object is a virtual-switch object that defines the QDIO virtual switch of a virtualization-host object of type **"zvm"**.

Table 63. *qdio-virtual-switch* object properties

Name	Qualifier	Type	Description
<b>element-id</b>	—	String	Unique ID for the virtual switch within the scope of the containing virtualization host. The element-id is actually the name of the virtual switch. Once the virtual switch is created, the name cannot be changed.
<b>element-uri</b>	—	String/URI	The canonical URI path for the virtual switch is of the form <code>/api/virtualization-hosts/{virt-host-id}/virtual-switches/{element-id}</code> , where <code>{virt-host-id}</code> is the <b>object-id</b> of the virtualization host.
<b>parent</b>	—	String/URI	The URI path of the z/VM Virtualization Host that hosts this virtual switch.
<b>class</b>	—	String (14)	Always <b>"virtual-switch"</b>
<b>type</b>	—	String Enum	Virtual switch type. Always <b>"qdio"</b>
<b>name</b>	—	String (1-8)	The unique name identifying the virtual switch.  1-8 alphanumeric uppercase characters, or any of the following characters: "@# \$"
<b>switch-status</b>	—	String	Virtual switch status, read-only.
<b>is-vlan-aware</b>	—	Boolean	True if not VLAN unaware. VLAN unaware is a classification for a networking device that indicates it does not support the IEEE 802.1Q VLAN specification for VLAN membership and VLAN frame formats. These devices ignore the additional fields within the Ethernet frame that carry VLAN specific semantics.
<b>vlan-id</b>	—	String (1-4)	Vlan ID : 1-4 decimal digits, maximum 4094. Prerequisites: <b>is-vlan-aware</b> is true.
<b>vlan-object-uri</b>	—	String/URI	The canonical URI path for the associated virtual network.  Prerequisites: <b>is-vlan-aware</b> is true.
<b>vlan-port-type</b>	—	String Enum	Indicates the port type of the simulated NIC. This setting is applied to the switch, not the single port on the switch.  Prerequisites: <b>is-vlan-aware</b> is true.  Values: <ul style="list-style-type: none"> <li><b>"trunk"</b>: When the switch port is configured in trunk mode, it will allow the flow of traffic from multiple virtual networks (i.e. VLANs). The port must be configured with those virtual networks.</li> <li><b>"access"</b>: When the switch port is configured in access mode, it will support a single virtual network. Traffic from the virtual server's network adapter will be tagged with the virtual network configured for this port, and traffic destined to the virtual server on this port will be verified that it is tagged with the configured virtual network.</li> </ul>
<b>vlan-native-id</b>	—	String	Native vlan ID : 1-4 decimal digits, maximum 4094.  Prerequisites: <b>is-vlan-aware</b> is true.  A native VLAN ID, (usually VLAN ID 0001) is deployed internally by the virtual switch to associate or flow untagged frames through the switching fabric. Only those guests that are configured for the native VLAN ID will receive or send untagged frames.

Table 63. *qdio-virtual-switch object properties (continued)*

Name	Qualifier	Type	Description
<b>is-gvrp-enabled</b>	—	Boolean	True if GVRP is enabled.  Prerequisites: <b>is-vlan-aware</b> is true.  Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) is an application defined in the IEEE802.1Q standard that allows for the control of VLANs. It runs only on 802.1Q trunk links. It prunes trunk links so that only active VLANs will be sent across trunk connections.
<b>layer-mode</b>	—	String Enum	Indicates the transport for the virtual switch is Ethernet or IP. Values: <ul style="list-style-type: none"> <li>• <b>"eth"</b>: This type is Data Link (Layer 2) based, where the Ethernet frame is used as the point of reference for source and destination Media Access Control (MAC) addresses in transporting Ethernet frames on the LAN.</li> <li>• <b>"ip"</b>: This type is Network (Layer 3) based, where the IP package is used as the point of reference for source and destination IP addresses in transporting IP packets on the LAN.</li> </ul>
<b>router</b>	—	String Enum	The router type.  Prerequisites: layer-mode is <b>"ip"</b> .  Values: <ul style="list-style-type: none"> <li>• <b>"none"</b> - Indicates that the OSA-Express device identified will not act as a router to the virtual switch. If a datagram is received at this device for an unknown IP address, the datagram will be discarded.</li> <li>• <b>"primary"</b> - Indicates that the OSA-Express device identified will act as a primary router to the virtual switch. If a datagram is received at this device for an unknown IP address, the datagram will be passed to the virtual switch. The only time to set PRIMARY for a virtual switch is if you have a guest attached to the virtual switch that is providing a routing function for systems attached to another network.</li> </ul>
<b>queue-size</b>	—	Integer	Indicates the upper limit of the amount of fixed storage CP and Queued Direct I/O Hardware Facility will use for buffers for each OSA-Express data device.  Queue storage in megabytes: decimal number between 1 and 8.
<b>ip-timeout</b>	—	Integer	IP timeout in seconds: decimal number between 1 and 240.
<b>is-connect-uplinks</b>		Boolean	True if connected to uplinks.
<b>real-uplinks</b>	—	Array of real-uplink objects	A list of zero to three real-uplink objects.
<b>is-use-any-available-controller</b>	—	Boolean	True if the switch uses any available controller.
<b>controllers</b>	—	Array of strings	A list of names of switch controllers.  Prerequisites: <b>is-use-any-available-controller</b> is false.

Table 63. *qdio-virtual-switch object properties (continued)*

Name	Qualifier	Type	Description
<b>bridge-value-type</b>	(w)	String Enum	<p>Whether the bridge port is a primary or secondary port. There can only be 1 primary and up to 4 secondary bridge ports.</p> <p>Values:</p> <ul style="list-style-type: none"> <li>• <b>"primary"</b> - Bridge port is the primary port</li> <li>• <b>"secondary"</b> - Bridge port is a secondary port.</li> </ul> <p>Prerequisite: This field is applied to a vSwitch that has a valid <b>bridge-device-number</b>.</p> <p>If not specified when the bridge port is defined, then the default is <b>"secondary"</b>. This field can be modified when the <b>bridge-connection-status</b> is <b>"disconnected"</b>.</p>
<b>bridge-device-number</b>	(w)	String (1-4)	<p>The bridge device number. The following values may be specified:</p> <ul style="list-style-type: none"> <li>• The bridge real device number, 1~4 hex digits.</li> </ul> <p>Prerequisites: The <b>is-bridge-capable</b> property of the virtualization host is true and the <b>layer-mode</b> property of the virtual switch is <b>"eth"</b>.</p> <p>This property is required to define a device for a bridge port and must be a nonblank, valid device number.</p>
<b>bridge-connection-status</b>	(w)	String Enum	<p>Whether the bridge port is connected, disconnected or in standby state.</p> <p>Prerequisites: The <b>is-bridge-capable</b> property of the virtualization host is true and the virtual switch has a valid <b>bridge-device-number</b>.</p> <p>Values:</p> <ul style="list-style-type: none"> <li>• <b>"connected"</b> - Bridge port is connected (default)</li> <li>• <b>"disconnected"</b> - Bridge port is disconnected</li> <li>• <b>"standby"</b> - Bridge port is in standby state.</li> </ul>
<b>mtu-size-enforcement</b>	(w)	String Enum	<p>How the virtual switch MTU size is enforced.</p> <p>Prerequisites: The <b>is-bridge-capable</b> property of the virtualization host is true.</p> <p>Values:</p> <ul style="list-style-type: none"> <li>• <b>"external"</b> - The MTU size will be set to the size used by the OSA adapter. This is the default value.</li> <li>• <b>"off"</b> - MTU enforcement is disabled</li> <li>• <b>"user-defined"</b> - The MTU size is specified in the <b>mtu-size</b> field</li> </ul>
<b>mtu-size</b>	(w)	Integer	<p>The MTU specification represents the acceptable MTU size, in bytes, enforced by the virtual switch. MTU size is a decimal number between 512-65535.</p> <p>Prerequisite: This field is applied to a virtual switch where the <b>is-bridge-capable</b> property of the virtualization host is true and the <b>mtu-size-enforcement</b> of the virtual switch is <b>"user-defined"</b>.</p>

## Operations

If a virtualization host operation accesses a z/VM virtualization host and encounters an error while communicating with the virtualization host via SMAPI, the response body is a SMAPI Error Response Body.

### List Virtualization Hosts of a zBX (Node)

The **List Virtualization Hosts of a zBX (Node)** operation lists the Virtualization Hosts managed by the zBX node with the given identifier.

#### HTTP method and URI

**GET** /api/zbx/{zbx-id}/virtualization-hosts

In this request, the URI variable *{zbx-id}* is the object ID of the zBX node object.

#### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property.  Value must be a valid Virtualization Host <b>type</b> property value.

#### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtualization-hosts	Array of objects	Array of virtualization-host-info objects, described in the next table. Returned array may be empty.

Each nested virtualization-host-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The canonical URI path of the Virtualization Host
name	String	Name of the Virtualization Host object
type	String Enum	Type of Virtualization Host (" <b>power-vm</b> " or " <b>x-hyp</b> ")
status	String Enum	Current status of the Virtualization Host

#### Description

This operation lists the Virtualization Hosts that are managed by the identified zBX node. The **object-uri**, **name**, **type**, and **status** are provided for each.

If the **name** query parameter is specified, the returned list is limited to those Virtualization Hosts that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid Virtualization Host **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those Virtualization Hosts that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not done.

If both **name** and **type** query parameters are specified, a Virtualization Host is included in the list only if it passes both the **name** and **type** filtering criteria.

The zBX specified by *{zbx-id}* must be a zBX node (that is, have a **type** of "**node**"), otherwise status code 400 (Bad Request) is returned.

A Virtualization Host is included in the list only if the API user has object-access permission to its hosting-environment. If an HMC is a manager of a Virtualization Host but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

If the zBX node does not manage any Virtualization Hosts or if no Virtualization Hosts are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the specified zBX node
- Object access permission to hosting-environment of the Virtualization Hosts managed by the specified zBX node.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is provided and the response body is as described in "Response body contents" on page 222.

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	80	A <b>type</b> query parameter defines an invalid value.
	89	The request URI designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of " <b>node</b> ").
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Example HTTP interaction

```
GET /api/zbx/8dac6a58-935e-352c-996e-ded17dbf92c0/virtualization-hosts HTTP/1.1
x-api-session: 579shzuc4ca3lwnly7xk2s64w6qmwy2v2zsikdrn4ky9sn9vv1
```

Figure 98. List Virtualization Hosts of a zBX (Node): Request

```
200 OK
server: zSeries management console API web server / 2.0
date: Mon, 09 Feb 2015 19:24:33 GMT
content-type: application/json;charset=UTF-8
content-length: 444
{
  "virtualization-hosts": [
    {
      "name": "B.1.01",
      "object-uri": "/api/virtualization-hosts/40a480b6-a7f1-11e4-ad46-42f2e9106e93",
      "status": "operating",
      "type": "power-vm"
    },
    {
      "name": "B.2.03",
      "object-uri": "/api/virtualization-hosts/88f0a20a-a7f1-11e4-ad46-42f2e9106e93",
      "status": "operating",
      "type": "x-hyp"
    }
  ]
}
```

Figure 99. List Virtualization Hosts of a zBX (Node): Response

## List Virtualization Hosts of a Node

The **List Virtualization Hosts of a Node** operation lists the Virtualization Hosts managed by the node with the given identifier.

### HTTP method and URI

```
GET /api/ensembles/{ensemble-id}/nodes/{node-id}/virtualization-hosts
```

#### URI variables:

Name	Description
{ensemble-id}	The object ID of the Ensemble object
{node-id}	The object ID of the Node object

#### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property.  Value must be a valid Virtualization Host <b>type</b> property value.



## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtualization-hosts	Array of objects	Array of virtualization-host-info objects, described in the next table. Returned array may be empty.

Each nested virtualization-host-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The canonical URI path of the Virtualization Host
name	String	Name of the Virtualization Host object
type	String Enum	Type of Virtualization Host. (" <b>power-vm</b> ", " <b>x-hyp</b> ", " <b>zvm</b> ", or " <b>prsm</b> ")
status	String Enum	Current status of the Virtualization Host

## Description

This operation lists the Virtualization Hosts that are managed by the specified node. The `object-uri`, `name`, `type`, and `status` are provided for each.

If the **name** query parameter is specified, the returned list is limited to those Virtualization Hosts that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid Virtualization Host **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those Virtualization Hosts that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not done.

If both **name** and **type** query parameters are specified, a Virtualization Host is included in the list only if it passes both the **name** and **type** filtering criteria.

A Virtualization Host is included in the list only if the API user has object-access permission to its hosting-environment. If an HMC is a manager of a Virtualization Host but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

If the Node does not manage any Virtualization Hosts or if no Virtualization Hosts are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the specified node
- Object access permission to hosting-environment of the Virtualization Hosts managed by the specified node.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is provided and the response body is as described in “Response body contents” on page 225.

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	80	A <b>type</b> query parameter defines an invalid value.
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
GET /api/ensembles/de7915f4-a127-11e2-8273-5cf3fcae8019/nodes/8dac6a58-
935e-352c-996e-ded17dbf92c0/virtualization-hosts HTTP/1.1
x-api-session: 579shzuc4ca3lwnly7xk2s64w6qmwy2v2zsjkdrn4ky9sn9vv1
```

Figure 100. List Virtualization Hosts of a Node: Request

```
200 OK
server: zSeries management console API web server / 2.0
date: Mon, 09 Feb 2015 19:24:27 GMT
content-type: application/json;charset=UTF-8
content-length: 444
{
  "virtualization-hosts": [
    {
      "name": "B.1.01",
      "object-uri": "/api/virtualization-hosts/40a480b6-a7f1-11e4-ad46-42f2e9106e93",
      "status": "operating",
      "type": "power-vm"
    },
    {
      "name": "B.2.03",
      "object-uri": "/api/virtualization-hosts/88f0a20a-a7f1-11e4-ad46-42f2e9106e93",
      "status": "operating",
      "type": "x-hyp"
    }
  ]
}
```

Figure 101. List Virtualization Hosts of a Node: Response

## List Virtualization Hosts of an Ensemble

The **List Virtualization Hosts of an Ensemble** operation lists the Virtualization Hosts managed by the ensemble with the given identifier.

### HTTP method and URI

GET /api/ensembles/{*ensemble-id*}/virtualization-hosts

In this request, the URI variable *{ensemble-id}* is the object ID of the Ensemble object.

### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property. Value must be a valid Virtualization Host <b>type</b> property value.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtualization-hosts	Array of objects	Array of virtualization-host-info objects, described in the next table. Returned array may be empty.

Each nested virtualization-host-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The canonical URI path of the Virtualization Host
name	String	Name of the Virtualization Host
type	String Enum	Type of Virtualization Host
status	String Enum	Current status of the Virtualization Host

### Description

This operation lists the Virtualization Hosts that are managed by the identified ensemble.

If the **name** query parameter is specified, the returned list is limited to those Virtualization Hosts that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid Virtualization Host **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those Virtualization Hosts that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not done.

A Virtualization Host is included in the list only if the API user has object-access permission to its hosting-environment. If an HMC is a manager of a Virtualization Host but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

If the ensemble does not manage any Virtualization Hosts or if no Virtualization Hosts are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the specified ensemble
- Object access permission to hosting-environment of the Virtualization Hosts managed by the specified ensemble.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is provided and the response body is as described in “Response body contents” on page 227.

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	80	A <b>type</b> query parameter defines an invalid value.
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/virtualization-hosts HTTP/1.1
x-api-session: 4oup923zgs27vmd1wpzvu47ikgzhxa8bwimjofpq6d3eq3j13q
```

---

Figure 102. List Virtualization Hosts of an Ensemble: Request

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 22 Jul 2011 15:02:37 GMT
content-type: application/json;charset=UTF-8
content-length: 824
{
  "virtualization-hosts": [
    {
      "name": "APIVM1",
      "object-uri": "/api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4",
      "status": "operating",
      "type": "zvm"
    },
    {
      "name": "B.2.02",
      "object-uri": "/api/virtualization-hosts/ba97ff30-2990-11e0-8d5b-001f163803de",
      "status": "operating",
      "type": "power-vm"
    },
    {
      "name": "B.2.03",
      "object-uri": "/api/virtualization-hosts/931b25d6-82e1-11e0-b9e4-f0def10bff8d",
      "status": "operating",
      "type": "x-hyp"
    },
    {
      "name": "R32",
      "object-uri": "/api/virtualization-hosts/bab76208-2990-11e0-8d5b-001f163803de",
      "status": "operating",
      "type": "prsm"
    }
  ]
}

```

Figure 103. List Virtualization Hosts of an Ensemble: Response

## List Virtualization Hosts of a CPC

The **List Virtualization Hosts of a CPC** operation lists the Virtualization Hosts managed by the CPC with the given identifier.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/virtualization-hosts

In this request, the URI variable {cpc-id} is Object ID of the CPC object.

### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property.  Value must be a valid Virtualization Host <b>type</b> property value.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtualization-hosts	Array of objects	Array of virtualization-host-info objects, described in the next table. Returned array may be empty.

Each nested virtualization-host-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The canonical URI path of the Virtualization Host
name	String	Name of the Virtualization Host object
type	String Enum	Type of Virtualization Host. (" <b>power-vm</b> ", " <b>x-hyp</b> ", " <b>zvm</b> ", or " <b>prsm</b> ")
status	String Enum	Current status of the Virtualization Host

## Description

This operation lists the Virtualization Hosts that are managed by the specified CPC. The object URI, object ID, display name, and display description are provided for each.

- | If the object-id *{cpc-id}* does not identify a CPC object to which the API user has object-access permission
- | or if the CPC is not a member of an ensemble, a 404 status code is returned.

If the **name** query parameter is specified, the returned list is limited to those Virtualization Hosts that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid Virtualization Host **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those Virtualization Hosts that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not done.

If both **name** and **type** query parameters are specified, a Virtualization Host is included in the list only if it passes both the **name** and **type** filtering criteria.

A Virtualization Host is included in the list only if the API user has object-access permission to its hosting-environment. If an HMC is a manager of a Virtualization Host but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

The list that is returned is never empty because a CPC always has a PR/SM Virtualization Host, and may have additional ones as well. If no Virtualization Hosts are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the specified CPC
- Object access permission to hosting-environment of the Virtualization Hosts managed by the specified CPC.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is provided and the response body is as described in “Response body contents” on page 230.

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	80	A <b>type</b> query parameter defines an invalid value.
404 (Not Found)	1	A CPC with object-id <i>{cpc-id}</i> does not exist on the HMC or API user does not have object-access permission for it.
	100	The CPC with object-id <i>{cpc-id}</i> is not a member of an ensemble.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/virtualization-hosts HTTP/1.1
x-api-session: 4oup923zgs27vmd1wpzvu47ikgzhxa8bwimjofpq6d3eq3j13q
```

---

Figure 104. List Virtualization Hosts of a CPC: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 22 Jul 2011 15:02:37 GMT
content-type: application/json;charset=UTF-8
content-length: 824
{
  "virtualization-hosts": [
    {
      "name": "APIVM1",
      "object-uri": "/api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4",
      "status": "operating",
      "type": "zvm"
    },
    {
      "name": "B.2.02",
      "object-uri": "/api/virtualization-hosts/ba97ff30-2990-11e0-8d5b-001f163803de",
      "status": "operating",
      "type": "power-vm"
    },
    {
      "name": "B.2.03",
      "object-uri": "/api/virtualization-hosts/931b25d6-82e1-11e0-b9e4-f0def10bff8d",
      "status": "operating",
      "type": "x-hyp"
    },
    {
      "name": "R32",
      "object-uri": "/api/virtualization-hosts/bab76208-2990-11e0-8d5b-001f163803de",
      "status": "operating",
      "type": "prsm"
    }
  ]
}
```

---

Figure 105. List Virtualization Hosts of a CPC: Response

## Get Virtualization Host Properties

The **Get Virtualization Host Properties** operation retrieves the properties of a single Virtualization Host object that is designated by its object-id.

### HTTP method and URI

**GET** /api/virtualization-hosts/{virt-host-id}

In this request, the URI variable {virt-host-id} is Object ID of the Virtualization Host object for which properties are to be obtained.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the Virtualization Host object as defined in the “Data model” on page 211. Field names and data types in the JSON object are the same as the property and data types defined in the data model.

### Description

Retrieve the current values for properties supported by the specified Virtualization Host.



If the object-id *{virt-host-id}* does not identify a Virtualization Host object for which the API user has object-access permission to its hosting-environment, a 404 status code is returned.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to hosting-environment of the Virtualization Host with object-id *{virt-host-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is provided and the response body is as described in “Response body contents” on page 232.

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/bab76208-2990-11e0-8d5b-001f163803de HTTP/1.1
x-api-session: 4oup923zgs27vmd1wpzv47ikgzhxa8bwimjofpq6d3eq3j13q
```

---

Figure 106. Get Virtualization Host Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 22 Jul 2011 15:02:37 GMT
content-type: application/json;charset=UTF-8
content-length: 640
{
  "acceptable-status": [
    "operating"
  ],
  "class": "virtualization-host",
  "description": "Initial description",
  "feature-list": [],
  "has-unacceptable-status": false,
  "hosting-environment": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "is-locked": false,
  "name": "R32",
  "object-id": "bab76208-2990-11e0-8d5b-001f163803de",
  "object-uri": "/api/virtualization-hosts/bab76208-2990-11e0-8d5b-001f163803de",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "status": "operating",
  "type": "prsm",
  "virtual-server-shutdown-timeout": 200
}
```

---

Figure 107. Get Virtualization Host Properties: Response for virtualization host of type "prsm"

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 22 Jul 2011 15:02:37 GMT
content-type: application/json;charset=UTF-8
content-length: 988
{
  "acceptable-status": [
    "operating"
  ],
  "auto-start-virtual-servers": true,
  "class": "virtualization-host",
  "description": "single update ABC ",
  "feature-list": [],
  "has-unacceptable-status": true,
  "hosting-environment": "/api/blades/938706AC3FF111D78B5600215EC0330E",
  "is-locked": false,
  "maximum-allowed-dedicated-processors": 57,
  "maximum-allowed-processing-units": 57.600000000000001,
  "maximum-allowed-virtual-processors": 64,
  "maximum-memory-size-for-virtual-server": 27648,
  "memory-increment-in-megabytes": 256,
  "minimum-memory-size-for-virtual-server": 256,
  "name": "B.2.02",
  "object-id": "ba97ff30-2990-11e0-8d5b-001f163803de",
  "object-uri": "/api/virtualization-hosts/ba97ff30-2990-11e0-8d5b-001f163803de",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "status": "operating",
  "type": "power-vm",
  "virtual-server-shutdown-timeout": 2147483647
}
```

---

Figure 108. Get Virtualization Host Properties: Response for virtualization host of type "power-vm"

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 22 Jul 2011 15:02:37 GMT
content-type: application/json;charset=UTF-8
content-length: 930
{
  "acceptable-status": [
    "operating"
  ],
  "auto-start-virtual-servers": true,
  "class": "virtualization-host",
  "description": "",
  "feature-list": [],
  "has-unacceptable-status": true,
  "hosting-environment": "/api/blades/B8210BC02D1E11E0AE81E41F13FE1430",
  "is-locked": false,
  "maximum-allowed-ide-devices": 3,
  "maximum-allowed-virtual-processors": 16.0,
  "maximum-memory-size-for-virtual-server": 125829,
  "memory-increment-in-megabytes": 1,
  "minimum-memory-size-for-virtual-server": 1,
  "mixed-mode-boot-restriction": true,
  "name": "B.2.03",
  "object-id": "931b25d6-82e1-11e0-b9e4-f0def10bff8d",
  "object-uri": "/api/virtualization-hosts/931b25d6-82e1-11e0-b9e4-f0def10bff8d",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "status": "operating",
  "type": "x-hyp",
  "virtual-server-shutdown-timeout": 302
}
```

---

Figure 109. Get Virtualization Host Properties: Response for virtualization host of type "x-hyp"

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 22 Jul 2011 15:02:37 GMT
content-type: application/json;charset=UTF-8
content-length: 893
{
  "acceptable-status": [
    "operating"
  ],
  "class": "virtualization-host",
  "cpu-shares-supported": true,
  "description": "An initial description",
  "feature-list": ["boot-sequence-network-priority-restriction"],
  "has-unacceptable-status": false,
  "hosting-environment": "/api/virtual-servers/4401b16c-ac9a-11e0-ade4-001f163805d8",
  "is-locked": false,
  "maximum-allowed-dedicated-processors": 64,
  "maximum-allowed-virtual-processors": 64,
  "maximum-memory-size-for-virtual-server": 1048576,
  "memory-increment-in-megabytes": 1,
  "minimum-memory-size-for-virtual-server": 64,
  "name": "APIVM1",
  "object-id": "342d80e0-65ff-11e0-acfd-f0def10c03f4",
  "object-uri": "/api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "status": "operating",
  "type": "zvm",
  "virtual-server-shutdown-timeout": 212
}

```

Figure 110. Get Virtualization Host Properties: Response for virtualization host of type "zvm"

## Update Virtualization Host Properties

The **Update Virtualization Host Properties** operation updates one or more of the writeable properties of a Virtualization Host.

### HTTP method and URI

**POST** /api/virtualization-hosts/{*virt-host-id*}

In this request, the URI variable {*virt-host-id*} is the object ID of the Virtualization Host object for which properties are to be updated.

### Request body contents

The request body contains a JSON object that provides the new values of the writeable properties of the Virtualization Host object as defined in the Data Model section above. Field names and data types in the JSON object are the same as the property or relationship names and data types defined in the data model.

Writeable properties are only valid if they are supported for a Virtualization Host whose **type** property matches the given **type** property value. For example, auto-start-virtual-servers is only a writeable property for type "**power-vm**" and "**x-hyp**" Virtualization Hosts.

## Description

This operation updates writeable properties of the Virtualization Host object specified by the request URI.

The request body contains an object with one or more fields with field names that correspond to the names of properties for this object. On successful execution, the value of each corresponding property of the object is updated with the value provided by the input field, and status code 204 (No Content) is returned without supplying any response body. The request body does not need to specify a value for all writeable properties, but rather can and should contain fields for the properties to be updated. Object properties for which no input value is provided remain unchanged by this operation.

If the update changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

The URI path must designate an existing Virtualization Host object and the API user must have object-access permission to its hosting-environment. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must also have task permission to the **System Details** task as well, otherwise status code 403 (Forbidden) is returned.

The request body is validated against the data model for this object type to ensure that it contains only writeable properties and the data types of those properties are as required. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Action/task permission to the **System Details** task
- Object access permission to hosting-environment of the Virtualization Host with object-id *{virt-host-id}*.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## List Virtual Switches

The **List Virtual Switches** operation lists the virtual switches managed by the z/VM Virtualization Host with the given identifier.

### HTTP method and URI

**GET** /api/virtualization-host/{*virt-host-id*}/virtual-switches

In this request, the URI variable *{virt-host-id}* is the object ID of the Virtualization Host.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtual-switches	Array of objects	Array of virtual-switch-info objects, described in the next table. Returned array may be empty.

Each nested virtual-switch-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	The <b>element-uri</b> property of virtual-switch element.
type	String Enum	The <b>type</b> property of virtual-switch element.
name	String	The <b>name</b> property of virtual-switch element.

### Description

This operation lists the virtual switches that are managed by the identified Virtualization Host.

If the Virtualization Host does not have any virtual network switches, an empty list is provided and the operation completes successfully.

### Authorization requirements

This operation has the following authorization requirements:

- Action/task role permission to the **Manage Virtual Switches** task
- Object access permission to hosting-environment of the Virtualization Host with object-id *{virt-host-id}*.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.

HTTP error status code	Reason code	Description
404 (Not Found)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMIPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMIPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4/virtual-switches HTTP/1.1
x-api-session: 4oup923zgs27vmd1wpzv47ikgzhxa8bwimjofpq6d3eq3j13q
```

---

Figure 111. List Virtual Switches: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 22 Jul 2011 15:02:46 GMT
content-type: application/json;charset=UTF-8
content-length: 411
{
  "virtual-switches": [
    {
      "element-uri": "/api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4/virtual-switches/S2777",
      "name": "S2777",
      "type": "iedn"
    },
    {
      "element-uri": "/api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4/virtual-switches/4179SW3",
      "name": "4179SW3",
      "type": "qdio"
    }
  ]
}
```

---

Figure 112. List Virtual Switches: Response

## Get Virtual Switch Properties

The **Get Virtual Switch Properties** operation retrieves the properties of a single virtual switch that is designated by the request URI.



## HTTP method and URI

GET /api/virtualization-host/{*virt-host-id*}/virtual-switches/{*virtual-switch-id*}

### URI variables

Variable	Type	Description
<i>virt-host-id</i>	String	Object ID of the Virtualization Host
<i>virtual-switch-id</i>	String	Element ID of the Virtual Switch

## Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the virtual switch object as defined in the “Data model” on page 211. Field names and data types in the JSON object are the same as the property or relationship names and data types defined in the data model.

## Description

This operation returns a JSON object that provides the current values of the properties for the virtual switch element object as defined in the “Data model” on page 211.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to hosting-environment of the Virtualization Host with object-id *{virt-host-id}*
- Task role permission to the **Manage Virtual Switches** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/57ab94c8-03e6-11e1-baf3-001f163805d8/virtual-switches/SSSS HTTP/1.1
x-api-session: 3ch9r8g2lavxw9st52brubgk4bpsqik3jcbg8hdwpkg5f1wpx
```

---

*Figure 113. Get Virtual Switch Properties: Request*

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 07 Dec 2011 06:01:18 GMT
content-type: application/json;charset=UTF-8
content-length: 575
{
  "bridge-connection-status": null,
  "bridge-device-number": null,
  "bridge-value-type": null,
  "class": "virtual-switch",
  "controllers": null,
  "element-id": "SSSS",
  "element-uri": "/api/virtualization-hosts/57ab94c8-03e6-11e1-baf3-001f163805d8/
  virtual-switches/SSSS",
  "ip-timeout": 5,
  "is-connect-uplinks": true,
  "is-use-any-available-controller": true,
  "layer-mode": "eth",
  "mtu-size": null,
  "mtu-size-enforcement": "external",
  "name": "SSSS",
  "parent": "/api/virtualization-hosts/57ab94c8-03e6-11e1-baf3-001f163805d8",
  "queue-size": 8,
  "real-uplinks": [],
  "router": null,
  "switch-status": "Defined",
  "type": "iedn"
}
```

---

*Figure 114. Get Virtual Switch Properties: Response for virtual switch of type "iedn"*

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 07 Dec 2011 06:01:38 GMT
content-type: application/json;charset=UTF-8
content-length: 718
{
  "bridge-connection-status": null,
  "bridge-device-number": null,
  "bridge-value-type": null,
  "class": "virtual-switch",
  "controllers": null,
  "element-id": "TESTQPOP",
  "element-uri": "/api/virtualization-hosts/57ab94c8-03e6-11e1-baf3-001f163805d8/
    virtual-switches/TESTQPOP",
  "ip-timeout": 5,
  "is-connect-uplinks": true,
  "is-gvrp-enabled": false,
  "is-use-any-available-controller": true,
  "is-vlan-aware": true,
  "layer-mode": "eth",
  "mtu-size": null,
  "mtu-size-enforcement": "external",
  "name": "TESTQPOP",
  "parent": "/api/virtualization-hosts/57ab94c8-03e6-11e1-baf3-001f163805d8",
  "queue-size": 8,
  "real-uplinks": [],
  "router": null,
  "switch-status": "Defined",
  "type": "qdio",
  "vlan-id": "AWARE",
  "vlan-native-id": "1",
  "vlan-object-uri": "",
  "vlan-port-type": "ACCESS"
}

```

Figure 115. Get Virtual Switch Properties: Response for virtual switch of type "qdio"

## Create IEDN Virtual Switch

The **Create IEDN Virtual Switch** operation creates an IEDN virtual network switch for the z/VM Virtualization Host.

### HTTP method and URI

**POST** /api/virtualization-host/{*virt-host-id*}/virtual-switches/operations/create-iedn

In this request, the URI variable *{virt-host-id}* is the object ID of the Virtualization Host.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The <b>name</b> property of iedn-virtual-switch object
layer-mode	String Enum	Optional	The <b>layer-mode</b> property of iedn-virtual-switch object. Default value is <b>"eth"</b> .
router	String Enum	Optional	The <b>router</b> property of iedn-virtual-switch object. Default value is <b>"none"</b> .

Field name	Type	Rqd/Opt	Description
queue-size	Integer	Optional	The <b>queue-size</b> property of iedn-virtual-switch object. Default value is 8.
ip-timeout	Integer	Optional	The <b>ip-timeout</b> property of iedn-virtual-switch object. Default value is 5.
is-connect-uplinks	Boolean	Optional	The <b>is-connect-uplinks</b> property of iedn-virtual-switch object. Default value is false.
real-uplinks	Array of real-uplink objects	Required if <b>is-connect-uplinks</b> is true	The <b>real-uplinks</b> property of iedn-virtual-switch object. Default value is an empty list. If <b>is-connect-uplinks</b> is true, the list is required and contains 1-3 real-uplink objects.
is-use-any-available-controller	Boolean	Optional	The <b>is-use-any-available-controller</b> property of iedn-virtual-switch object. Default value is true. If the value is true and virtual switch is successfully created and there is a controller available in the z/VM, this property will be set to false, and the available controller will be used and shown in the <b>controllers</b> property.
controllers	Array of Strings	Required if <b>is-use-any-available-controller</b> is false	The <b>controllers</b> property of iedn-virtual-switch object
bridge-value-type	String	Optional; Allowed only if <b>bridge-device-number</b> has been specified	The <b>bridge-value-type</b> property of the iedn-virtual-switch object. Default value is "secondary".
bridge-device-number	String	Optional; Allowed only if the <b>is-bridge-capable</b> property of the virtualization host is true and the <b>layer-mode</b> property of the virtual switch is "eth"	The <b>bridge-device-number</b> property of the iedn-virtual-switch object. Default value is null.
bridge-connection-status	String	Optional; Allowed only if <b>bridge-device-number</b> has been specified	Values: <ul style="list-style-type: none"> <li>• "connect" - Connect the bridge port</li> <li>• "disconnect" - Disconnect the bridge port</li> </ul> Default value is "connect".
mtu-size-enforcement	String	Optional; Allowed only if the <b>is-bridge-capable</b> property of the virtualization host is true	The <b>mtu-size-enforcement</b> property of the iedn-virtual-switch object. Default value is "external".
mtu-size	Integer	Optional; Allowed only if the <b>is-bridge-capable</b> property of the virtualization host is true and the <b>mtu-size-enforcement</b> property of the virtual switch is "user-defined"	The <b>mtu-size</b> property of the iedn-virtual-switch object. Required if <b>mtu-size-enforcement</b> is "user-defined".

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	The <b>element-uri</b> property of the created virtual-switch element.

## Description

This operation creates the IEDN virtual switch for the identified Virtualization Host and then returns its URI. The response also includes a **Location** header that provides this URI.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to hosting-environment of the Virtualization Host with object-id *{virt-host-id}*
- Action/task permission to the **Manage Virtual Switches** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents.”

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	81	A virtual switch with the <b>name</b> specified in the request body already exists on the Virtualization Host with object-id <i>{virt-host-id}</i> .
	87	The iedn-virtual-switch object <b>bridge-device-number</b> property specified is already in use.
	88	The iedn-virtual-switch object <b>bridge-device-number</b> property specified is not defined available to the z/VM Virtualization Host.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
409 (Conflict)	80	The Virtualization Host already has a vSwitch with a bridge where the iedn-virtual-switch object <b>bridge-value-type</b> property is " <b>primary</b> ".
	81	The Virtualization Host already has four vSwitches with a bridge where the iedn-virtual-switch object <b>bridge-value-type</b> property is " <b>secondary</b> ". Four is the maximum.

HTTP error status code	Reason code	Description
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Create QDIO Virtual Switch

The **Create QDIO Virtual Switch** operation creates a QDIO virtual network switch for the z/VM Virtualization Host.

### HTTP method and URI

**POST** /api/virtualization-hosts/{*virt-host-id*}/virtual-switches/operations/create-qdio

In this request, the URI variable *{virt-host-id}* is the object ID of the Virtualization Host.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The <b>name</b> property of qdio-virtual-switch object
is-vlan-aware	Boolean	Optional	The <b>is-vlan-aware</b> property of qdio-virtual-switch object. Default value is false.
vlan-id	String	Optional; Allowed only if <b>is-vlan-aware</b> is true	The <b>vlan-id</b> property of qdio-virtual-switch object. Default value is "".
vlan-port-type	String Enum	Optional; Allowed only if <b>is-vlan-aware</b> is true	The <b>vlan-port-type</b> property of qdio-virtual-switch object. Default value is "access".
vlan-native-id	String	Optional; Allowed only if <b>is-vlan-aware</b> is true	The <b>vlan-native-id</b> property of qdio-virtual-switch object. Default value is "1".
is-gvrp-enabled	Boolean	Optional; Allowed only if <b>is-vlan-aware</b> is true	The <b>is-gvrp-enabled</b> property of qdio-virtual-switch object. Default value is false.
layer-mode	String Enum	Optional	The <b>layer-mode</b> property of qdio-virtual-switch object. Default value is "eth".
router	String Enum	Optional	The <b>router</b> property of qdio-virtual-switch object. Default value is "none".
queue-size	Integer	Optional	The <b>queue-size</b> property of qdio-virtual-switch object. Default value is 8.

Field name	Type	Rqd/Opt	Description
ip-timeout	Integer	Optional	The <b>ip-timeout</b> property of qdio-virtual-switch object. Default value is 5.
is-connect-uplinks	Boolean	Optional	The <b>is-connect-uplinks</b> property of qdio-virtual-switch object. Default value is false.
real-uplinks	Array of real-uplink objects	Required if <b>is-connect-uplinks</b> is true	The <b>real-uplinks</b> property of qdio-virtual-switch object. Default value is an empty list. If <b>is-connect-uplinks</b> is true, the list is required and contains 1-3 real-uplink objects.
is-use-any-available-controller	Boolean	Optional	The <b>is-use-any-available-controller</b> property of qdio-virtual-switch object. Default value is true. If the value is true and virtual switch is successfully created and there is a controller available in the z/VM, this property will be set to false, and the available controller will be used and shown in the <b>controllers</b> property
controllers	Array of Strings	Required if <b>is-use-any-available-controller</b> is false	The <b>controllers</b> property of qdio-virtual-switch object
bridge-value-type	String	Optional; Allowed only if <b>bridge-device-number</b> has been specified	The <b>bridge-value-type</b> property of the qdio-virtual-switch object. Default value is "secondary".
bridge-device-number	String	Optional; Allowed only if <b>theis-bridge-capable</b> property of the virtualization host is true and <b>thelayer-mode</b> property of the virtual switch is "eth"	The <b>bridge-device-number</b> property of the qdio-virtual-switch object. Default value is null.
bridge-connection-status	String	Optional; Allowed only if <b>bridge-device-number</b> has been specified	Values: <ul style="list-style-type: none"> <li>• "connect" - Connect the bridge port. (Default value)</li> <li>• "disconnect" - Disconnect the bridge port</li> </ul> Default value is "connect".
mtu-size-enforcement	String	Optional; Allowed only if the <b>is-bridge-capable</b> property of the virtualization host is true	The <b>mtu-size-enforcement</b> property the qdio-virtual-switch object. Default value is "external".
mtu-size	Integer	Optional; Allowed only if the <b>is-bridge-capable</b> property of the virtualization host is true and the <b>mtu-size-enforcement</b> property of the virtual switch is "user-defined"	The <b>mtu-size</b> property the qdio-virtual-switch object. Required if <b>mtu-size-enforcement</b> is "user-defined".

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	The <b>element-uri</b> property of the created virtual-switch element.

## Description

This operation creates the QDIO virtual switch for the identified Virtualization Host and then returns its element URI. The response also includes a **Location** header that provides this URI.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to hosting-environment of the Virtualization Host with object-id *{virt-host-id}*
- Action/task permission to the **Manage Virtual Switches** task

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents.”

If errors occur, the following HTTP status codes are provided, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	81	A virtual switch with the <b>name</b> specified in the request body already exists on the Virtualization Host with object-id <i>{virt-host-id}</i> .
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
409 (Conflict)	80	The Virtualization Host already has a vSwitch with a bridge where the qdio-virtual-switch object <b>bridge-value-type</b> property is " <b>primary</b> ".
	81	The Virtualization Host already has four vSwitches with a bridge where the qdio-virtual-switch object <b>bridge-value-type</b> property is " <b>secondary</b> ". Four is the maximum.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.



## Get Switch Controllers

The **Get Switch Controllers** operation gets a list of controllers for the z/VM Virtualization Host. Controllers are z/VM TCP/IP virtual machines used to manage OSA-Express devices associated with virtual switches. For details, see the *z/VM CP Commands and Utility Reference*, SC24-6175.

### HTTP method and URI

**GET** /api/virtualization-hosts/{virt-host-id}/operations/get-switch-controllers

In this request, the URI variable {virt-host-id} is the object ID of the Virtualization Host.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
controllers	Array of controller objects	An array of controller object as defined in the next table.

Each nested controller object provides the properties for a single switch controller, and has the following format:

Field name	Type	Description
name	String (0-8)	Name for the controller
is-available	Boolean	True if the controller is available to control an additional set of OSA-Express devices associated with the virtual switch.
vdev-range	String (1-9)	Identifies the device range where the OSA-Express devices associated with a virtual switch can be attached. The value is two virtual device address values separated by a hyphen (e.g. "8800-88FF") or "*", which indicates that the virtual device address used to attach the OSA-Express devices is the same as the real device address identified by the virtual switch's real-uplinks.
is-ip	Boolean	True if the virtual switch controller can initialize an OSA-Express device in IP mode. See Virtual Switch layer mode.
is-eth	Boolean	True if the virtual switch controller can initialize an OSA-Express device in ETH mode. See Virtual Switch layer mode.
is-vlan-arp	Boolean	True if the virtual switch controller can register IP addresses on the OSA-Express with the proper VLAN groups (VLAN_ARP).
is-gvrp	Boolean	True if the virtual switch controller can register VLAN IDs in use on a virtual switch with GVRP-aware switches (GVRP).
is-linkagg	Boolean	True if the controller can control virtual switches that are using link aggregation. A Link Aggregation port group is two or more links that are grouped together to appear as a single logical link.
is-isolation	Boolean	True if the controller can control virtual switches that are using the isolation setting.

### Description

This operation lists the controllers for the identified Virtualization Host.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to hosting-environment of the Virtualization Host with object-id *{virt-host-id}*.
- Action/task permission to the **Manage Virtual Switches** task

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in the “Response body contents” on page 249.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/0e4a5d94-a8c1-11e0-9492-00262df332b3/operations/
  get-switch-controllers HTTP/1.1
x-api-session: 1kcjo7etEU67cygzoknoww8caid5jyck6yfupyqz3619t3r
```

---

Figure 116. Get Switch Controllers: Request

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 07 Dec 2011 05:49:49 GMT
content-type: application/json;charset=UTF-8
content-length: 653
{
  "controllers": [
    {
      "is-available": true,
      "is-eth": true,
      "is-gvrp": true,
      "is-ip": true,
      "is-isolation": true,
      "is-linkagg": true,
      "is-vlan-arp": true,
      "name": "DTCVSW1",
      "vdev-range": "0600-F000"
    },
    {
      "is-available": true,
      "is-eth": true,
      "is-gvrp": true,
      "is-ip": true,
      "is-isolation": true,
      "is-linkagg": true,
      "is-vlan-arp": true,
      "name": "DTCENS1",
      "vdev-range": "*"
    }
  ]
}

```

Figure 117. Get Switch Controllers: Response

## Update Virtual Switch

The **Update Virtual Switch** operation modifies an existing virtual network switch for the z/VM Virtualization Host.

### HTTP method and URI

**POST** /api/virtualization-host/{*virt-host-id*}/virtual-switches/{*virtual-switch-id*}

#### URI variables

Variable	Description
{ <i>virt-host-id</i> }	Object ID of the Virtualization Host
{ <i>virtual-switch-id</i> }	Element ID of the virtual switch

## Request body contents

For IEDN virtual switch:

Field name	Type	Rqd/Opt	Description
router	String Enum	Optional; Allowed only if <b>layer-mode</b> is " <b>ip</b> "	The <b>router</b> property of iedn-virtual-switch object
queue-size	Integer	Optional	The <b>queue-size</b> property of iedn-virtual-switch object.
ip-timeout	Integer	Optional	The <b>ip-timeout</b> property of iedn-virtual-switch object.
is-connect-uplinks	Boolean	Optional	The <b>is-connect-uplinks</b> property of iedn-virtual-switch object
is-use-any-available-controller	Boolean	Optional	The <b>is-use-any-available-controller</b> property of iedn-virtual-switch object. If the value is true and virtual switch is successfully updated and there is a controller available in the z/VM, this property will be set to false, and the available controller will be used and shown in the <b>controllers</b> property.
controllers	Array of Strings	Optional; Allowed only if <b>is-use-any-available-controller</b> is false	The <b>controllers</b> property of iedn-virtual-switch object.
real-uplinks	Array of real-uplink objects	Required if <b>is-connect-uplinks</b> is true	The <b>real-uplinks</b> property of iedn-virtual-switch object. If <b>is-connect-uplinks</b> is true, the list is required and contains 1-3 real-uplink objects.
bridge-value-type	String	Optional; Allowed only if <b>bridge-device-number</b> has been specified	The <b>bridge-value-type</b> property of the iedn-virtual-switch object
bridge-device-number	String	Optional; Allowed only if the <b>is-bridge-capable</b> property of the virtualization host is true and the <b>layer-mode</b> property of the virtual switch is " <b>eth</b> "	The <b>bridge-device-number</b> property of the iedn-virtual-switch object
bridge-connection-status	String	Optional; Allowed only if <b>bridge-device-number</b> has been specified	Values: <ul style="list-style-type: none"> <li>• "<b>connect</b>" - Connect the bridge port. (Default value)</li> <li>• "<b>disconnect</b>" - Disconnect the bridge port</li> </ul>
mtu-size-enforcement	String	Optional; Allowed only if <b>theis-bridge-capable</b> property of the virtualization host is true	The <b>mtu-size-enforcement</b> property of the iedn-virtual-switch object

Field name	Type	Rqd/Opt	Description
mtu-size	Integer	Optional;  Allowed only if the <b>is-bridge-capable</b> property of the virtualization host is true and the <b>mtu-size-enforcement</b> property of the virtual switch is "user-defined"	The <b>mtu-size</b> property of the iedn-virtual-switch object. Required if <b>mtu-size-enforcement</b> is "user-defined".

For QDIO virtual switch:

Field name	Type	Rqd/Opt	Description
router	String Enum	Optional;  Allowed only if <b>layer-mode</b> is "ip"	The <b>router</b> property of the qdio-virtual-switch object
queue-size	Integer	Optional	The <b>queue-size</b> property of qdio-virtual-switch object.
ip-timeout	Integer	Optional	The <b>ip-timeout</b> property of qdio-virtual-switch object.
is-connect-uplinks	Boolean	Optional	The <b>is-connect-uplinks</b> property of qdio-virtual-switch object
is-use-any-available-controller	Boolean	Optional	The <b>is-use-any-available-controller</b> property of qdio-virtual-switch object. If the value is true and virtual switch is successfully updated and there is a controller available in the z/VM, this property will be set to false, and the available controller will be used and shown in the <b>controllers</b> property.
controllers	Array of Strings	Optional;  Allowed only if <b>is-use-any-available-controller</b> is false	The <b>controllers</b> property of qdio-virtual-switch object.
real-uplinks	Array of real-uplink objects	Required if <b>is-connect-uplinks</b> is true	The <b>real-uplinks</b> property of qdio-virtual-switch object. If <b>is-connect-uplinks</b> is true, the list is required and contains 1-3 real-uplink objects.
bridge-value-type	String	Optional;  Allowed only if <b>bridge-device-number</b> has been specified	The <b>bridge-value-type</b> property of the qdio-virtual-switch object
bridge-device-number	String	Optional;  Allowed only if <b>theis-bridge-capable</b> property of the virtualization host is true and the <b>layer-mode</b> property of the virtual switch is "eth"	The <b>bridge-device-number</b> property of the qdio-virtual-switch object

Field name	Type	Rqd/Opt	Description
bridge-connection-status	String	Optional; Allowed only if <b>bridge-device-number</b> has been specified	Values: <ul style="list-style-type: none"> <li>"connect" - Connect the bridge port. (Default value)</li> <li>"disconnect" - Disconnect the bridge port</li> </ul>
mtu-size-enforcement	String	Optional; Allowed only if the <b>is-bridge-capable</b> property of the virtualization host is true	The <b>mtu-size-enforcement</b> property of the qdio-virtual-switch object
mtu-size	Integer	Optional; Allowed only if <b>theis-bridge-capable</b> property of the virtualization host is true and the <b>mtu-size-enforcement</b> property of the virtual switch is "user-defined"	The <b>mtu-size</b> property of the qdio-virtual-switch object. Required if <b>mtu-size-enforcement</b> is "user-defined".

## Description

This operation modifies the virtual switch for the identified virtual server.

## Authorization requirements

This operation has the following authorization requirements:

- Action/task role permission to the **Manage Virtual Switches** task
- Object access permission to hosting-environment of the Virtualization Host with object-id *{virt-host-id}*.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	87	The iedn-virtual-switch object <b>bridge-device-number</b> property specified is already in use.
	88	The iedn-virtual-switch object <b>bridge-device-number</b> property specified is not defined available to the z/VM Virtualization Host.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.

HTTP error status code	Reason code	Description
409 (Conflict)	80	The Virtualization Host already has a vSwitch with a bridge where the qdio-virtual-switch object <b>bridge-value-type</b> property is <b>"primary"</b> .
	81	The Virtualization Host already has four vSwitches with a bridge where the qdio-virtual-switch object <b>bridge-value-type</b> property is <b>"secondary"</b> . Four is the maximum.
	82	The <b>bridge-connection-status</b> property for an iedn-virtual-switch or qdio-virtual-switch object must be <b>"disconnected"</b> to update bridge properties.
	83	To change the <b>bridge-device-number</b> property for an iedn-virtual-switch or qdio-virtual-switch object, it must first be set to <b>"none"</b> . This removes the bridge device.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Delete Virtual Switch

The **Delete Virtual Switch** operation deletes an existing virtual network switch specified by an element identifier for the z/VM Virtualization Host with the given object identifier.

### HTTP method and URI

**DELETE** /api/virtualization-hosts/{*virt-host-id*}/virtual-switches/{*virtual-switch-id*}

#### URI variables

Variable	Description
{ <i>virt-host-id</i> }	Object ID of the Virtualization Host
{ <i>virtual-switch-id</i> }	Element ID of the virtual switch

### Description

This operation deletes the virtual switch for the identified Virtualization Host.

### Authorization requirements

This operation has the following authorization requirements:

- Action/task role permission to the **Manage Virtual Switches** task
- Object access permission to hosting-environment of the Virtualization Host with object-id {*virt-host-id*}.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Activating a Virtualization Host

This operation is not directly accessible; it will occur as a side effect of activating the hosting environment. If **auto-start-virtual-servers** is true, the Virtualization Host activation will also activate all virtual servers on the Virtualization Host whose **auto-start** property is true.

See the Activate sections of the hosting environment objects for operation details, including URI parameters, response body contents, authorization requirements, and HTTP status and reason codes.

## Asynchronous result description

Once the activation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the Activate Virtualization Host request.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "**complete**", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description in “Job Status and Reason Codes” on page 257. The results also include a job-results nested object that has the following form:

Field name	Type	Description
<b>failed-virtual-servers</b>	Array of String/URI	Array of the virtual servers whose activation failed. This field only exists if the virtualization host activated, but one or more virtual servers failed to activate, indicated by job-status-code 500 and job-reason-code 102.



## Job Status and Reason Codes

Job status code	Job reason code	Description
200 (OK)	N/A	Activation completed successfully.
500 (Server Error)	100	Virtualization host activation failed.
	101	Virtualization host activation job timed out.
	102	Virtualization host activation succeeded, but some virtual servers failed to activate, see the response body for a list of virtual servers that failed to activate.

## Deactivating a Virtualization Host

This operation is not directly accessible; it will occur as a side effect of deactivating the hosting environment.

See the Deactivate sections of the hosting environment objects for operation details, including URI parameters, response body contents, authorization requirements, and HTTP status and reason codes.

### Asynchronous result description

Once the deactivation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the Deactivate Virtualization Host request.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "**complete**", the results include a job completion status code and reason code (fields `job-status-code` and `job-reason-code`) which are set as indicated in operation description below. The results also include a job-results nested object that has the following form:

Field name	Type	Description
<code>failed-virtual-servers</code>	Array of String/URI	Array of the virtual servers whose deactivation failed. This field only exists if the virtualization host deactivated, but one or more virtual servers failed to deactivate, indicated by <code>job-status-code</code> 500 and <code>job-reason-code</code> 102.

## Job Status and Reason Codes

Job status code	Job reason code	Description
200 (OK)	N/A	Deactivation completed successfully.
500 (Server Error)	100	Virtualization host deactivation failed.
	101	Virtualization host deactivation job timed out.
	102	Virtualization host deactivation succeeded, but some virtual servers failed to deactivate, see the response body for a list of virtual servers that failed to deactivate.

## SMAPI Error Response Body

If an operation encounters an error while communicating with a z/VM Virtualization Host via SMAPI, a 503 (Service Unavailable) status code is returned with reason code 100. If an operation is able to

communicate with a z/VM Virtualization Host and execute a command via SMAPI, but the SMAPI command returns a non-zero return code, a 503 (Service Unavailable) status code is returned with reason code 101 and the standard error response body is extended with the following information about the SMAPI command failure:

Field name	Type	Description
<b>smapi-command</b>	String	The name of the SMAPI command that encountered the error
<b>smapi-return-code</b>	String	The return code returned when executing the SMAPI command
<b>smapi-reason-code</b>	String	The reason code returned when executing the SMAPI command
<b>smapi-message</b>	String	The error message returned when executing the SMAPI command or <b>null</b> if no error message was provided

## Inventory service data

Information about the Virtualization Hosts managed by the HMC can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for Virtualization Host objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by inventory class, implicitly via a containing category, or by default) that objects of the various Virtualization Host type-specific inventory classes are to be included. An entry for a particular Virtualization Host is included only if the API user has object-access permission to that object and the applicable type-specific inventory class has been specified, as described in the following table:

Inventory class	Includes Virtualization Hosts with "type" value
power-vm-virtualization-host	power-vm
prsm-virtualization-host	prsm
x-hyp-virtualization-host	x-hyp
zvm-virtualization-host	zvm

For each Virtualization Host object to be included, the inventory response array includes the following:

- An entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the **Get Virtualization Host Properties** operation. That is, the data provided is the same as would be provided if a **Get Virtualization Host Properties** operation were requested targeting this object.
- An array entry for each Virtual Switch object associated with the virtualization host. For each such resource, an entry is included that is a JSON object with the same contents as specified in the Response Body Contents section of the Get Virtual Switch Properties operation. As a result, the data provided is the same as would be obtained if a Get Virtual Switch Properties operation were requested for each resource listed by a List Virtual Switches operation targeting the virtualization host.
- An array entry for each Virtualization Host Storage Resource object associated with the virtualization host. For each such resource, an entry is included that is a JSON object with the same contents as specified in the Response Body Contents section of the **Get Virtualization Host Storage Resource Properties** operation, however storage path accessibility status is not provided. (More specifically, the accessible property of path-information-fcp and path-information-eckd nested objects will always null.) This data is described in Chapter 11, "Storage Management," on page 363. As a result, the data provided is the same as would be obtained if a **Get Virtualization Host Storage Resource Properties** operation were requested with the **include-path-accessibility** query parameter specified as false for each resource listed by a List Virtualization Host Storage Resources operation targeting the virtualization host.
- An array entry for each Virtualization Host Storage Group object associated with the virtualization host. For each such group, an entry is included that is a JSON object with the same contents as specified in the Response Body Contents section of the **Get Virtualization Host Storage Group**

**Properties** operation. This data is described in Chapter 11, “Storage Management,” on page 363. As a result, the data provided is the same as would be obtained if a **Get Virtualization Host Storage Group Properties** operation where requested for each group listed by a **List Virtualization Host Storage Groups** operation targeting the virtualization host.

The array entry for a Virtualization Host object will appear in the results array before entries for associated virtual switches, virtualization host storage resources, or groups.

### Sample inventory data

The following fragments are examples of the JSON objects that would be included in the Get Inventory response to describe a single Virtualization Host object of a particular type. These objects would appear as array entries in the response array.

---

```
{
  "acceptable-status": [
    "operating"
  ],
  "auto-start-virtual-servers": false,
  "class": "virtualization-host",
  "description": "",
  "feature-list": [],
  "has-unacceptable-status": "true",
  "minimum-memory-size-for-virtual-server": 256,
  "name": "B.1.14",
  "object-id": "baab1cd2-2990-11e0-8d5b-001f163803de",
  "object-uri": "/api/virtualization-hosts/baab1cd2-2990-11e0-8d5b-001f163803de",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "status": "operating",
  "type": "power-vm",
  "virtual-server-shutdown-timeout": 300
}
```

---

Figure 118. Virtualization Host object: Sample inventory data for a virtualization host of type "power-vm"

---

```

{
  "acceptable-status": [
    "operating",
    "channel-acceptable"
  ],
  "auto-start-virtual-servers": false,
  "class": "virtualization-host",
  "description": "Initial description",
  "feature-list": [],
  "has-unacceptable-status": "false",
  "hosting-environment": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "is-locked": false,
  "name": "R32",
  "object-id": "bab76208-2990-11e0-8d5b-001f163803de",
  "object-uri": "/api/virtualization-hosts/bab76208-2990-11e0-8d5b-001f163803de",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "status": "operating",
  "type": "prsm",
  "virtual-server-shutdown-timeout": 200
}

```

---

Figure 119. Virtualization Host object: Sample inventory data for a virtualization host of type "prsm"

---

```

{
  "acceptable-status": [
    "operating"
  ],
  "auto-start-virtual-servers": true,
  "class": "virtualization-host",
  "cpu-shares-supported": true,
  "description": "",
  "feature-list": ["boot-sequence-network-priority-restriction"],
  "has-unacceptable-status": "true",
  "hosting-environment": "/api/blades/b8210bc0-2d1e-11e0-ae81-e41f13fe1430",
  "is-locked": false,
  "maximum-allowed-ide-devices": 3,
  "maximum-allowed-virtual-processors": 16.0,
  "maximum-memory-size-for-virtual-server": 125829,
  "memory-increment-in-megabytes": 1,
  "minimum-memory-size-for-virtual-server": 1,
  "mixed-mode-boot-restriction": true,
  "name": "B.1.03",
  "object-id": "931b25d6-82e1-11e0-b9e4-f0def10bff8d",
  "object-uri": "/api/virtualization-hosts/931b25d6-82e1-11e0-b9e4-f0def10bff8d",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/nodes/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "status": "operating",
  "type": "x-hyp",
  "virtual-server-shutdown-timeout": 302
}

```

---

Figure 120. Virtualization Host object: Sample inventory data for a virtualization host of type "x-hyp"

---

## Virtual Server Object

- I A Virtual Server object represents a single z Systems virtual server.

## Data Model

This object includes the properties defined the “Base managed object properties schema” on page 39, with the following class-specific specialization:

Table 64. Virtual Server object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>name</b>	(w)(pc)	String (1-64)	<p>The display name of the virtual server. Names must be unique to other existing virtual servers on the virtualization host.</p> <p>The format of the name varies based on the virtual server type:</p> <ul style="list-style-type: none"> <li>• <b>"zvm"</b>: 1-8 characters, characters may be uppercase alphanumeric or any of the following characters: "@#\$.:"</li> <li>• <b>"power-vm"</b>, <b>"x-hyp"</b>: 1-64 characters, must begin with an alphabetic characters, other characters may be characters may be alphanumeric, a space, or any of the following characters: "!@#%&amp;^*()_+ =,.;'~"</li> <li>• <b>"prsm"</b>: 1-8 characters, alphanumeric</li> </ul> <p>For virtual servers of type <b>"prsm"</b>, this property is the LPAR name as defined in the active IOCDs and is immutable.</p> <p>This property is also immutable for virtual servers whose <b>type</b> property is <b>"zvm"</b>, though it is user-defined through the Create Virtual Server operation.</p> <p>For <b>"power-vm"</b> and <b>"x-hyp"</b> virtual servers, this property may only be modified when the virtual server's status is <b>"not-operating"</b>.</p>
<b>description</b>	(w)(pc)	String	Read-only for virtual servers of type <b>"prsm"</b> .
<b>object-uri</b>	—	String/ URI	The canonical URI path for a Virtual Server object is of the form <code>/api/virtual-servers/{object-id}</code> .
<b>parent</b>	—	String/ URI	The URI path of the virtualization host that hosts this virtual server.
<b>class</b>	—	String (14)	Always <b>"virtual-server"</b> .

Table 64. Virtual Server object: base managed object properties specializations (continued)

Name	Qualifier	Type	Description of specialization
<b>status</b>	(sc)	String Enum	<p>The current operational status of the managed resource.</p> <p><b>"prsm"</b> values:</p> <ul style="list-style-type: none"> <li>• <b>"not-activated"</b> - indicates the virtual server's LPAR image was not activated.</li> <li>• <b>"operating"</b> - indicates all CPs are operating</li> <li>• <b>"not-operating"</b> - indicates no CPs are operating</li> <li>• <b>"not-communicating"</b> - indicates the HMC is not communicating with the support element</li> <li>• <b>"exceptions"</b> - indicates the virtual server has a problem of some sort, such as not being able to access storage</li> <li>• <b>"status-check"</b> - indicates at least one CP is operating, but at least one CP is not operating. When a <b>"prsm"</b> virtual server has this status, its LPAR Image is not available.</li> </ul> <p><b>"power-vm"</b> and <b>"x-hyp"</b> values:</p> <ul style="list-style-type: none"> <li>• <b>"operating"</b> - indicates virtual server is in an activated and running state</li> <li>• <b>"not-operating"</b> - indicates the virtual server is deactivated, in the process of loading, or has a error</li> <li>• <b>"not-communicating"</b> - indicates the HMC is not communicating with the support element</li> <li>• <b>"exceptions"</b> - indicates the virtual server has a problem of some sort, such as not being able to access storage</li> <li>• <b>"status-check"</b> - indicates the virtual server is powered on, but the virtualization host is not communicating so there is no status available</li> <li>• <b>"migrating"</b> - indicates the virtual server is in the process of migrating to another virtualization host</li> <li>• <b>"starting"</b> - indicates the virtual server in the process of powering on</li> <li>• <b>"stopping"</b> - indicates the virtual server in the process of powering off.</li> </ul> <p><b>"zvm"</b> values:</p> <ul style="list-style-type: none"> <li>• <b>"operating"</b> – indicates the virtual machine is logged on and has no current failure conditions</li> <li>• <b>"not-operating"</b>– indicates the virtual machine is logged on, but currently has some sort of failure condition</li> <li>• <b>"not-activated"</b> – indicates the virtual machine is not logged on</li> <li>• <b>"not-communicating"</b> - indicates the HMC is not communicating with the support element</li> <li>• <b>"logoff-timeout-started"</b> – indicates a logoff timeout has been started for the virtual machine</li> <li>• <b>"storage-limit-exceeded"</b> – indicates the storage limit has been exceeded for the virtual machine</li> <li>• <b>"forced-sleep"</b> – indicates the virtual machine is logged on, but has been placed into a forced sleep state</li> <li>• <b>"unknown"</b> – indicates the status of the virtual machine reported by z/VM was not recognized.</li> </ul>

Table 64. Virtual Server object: base managed object properties specializations (continued)

Name	Qualifier	Type	Description of specialization
<b>acceptable-status</b>	(w)(pc)	Array of String Enum	For virtual servers of type " <b>prsm</b> ", this property is immutable.  For " <b>zvm</b> " virtual servers, the following status values may not be set as acceptable: " <b>not-communicating</b> " and " <b>unknown</b> ".
<b>additional-status</b>	(sc)	String Enum	The property is not provided for virtual servers of type " <b>power-vm</b> ", " <b>x-hyp</b> ", or " <b>prsm</b> ".  For " <b>zvm</b> " virtual servers, this property only applies when the status is " <b>not-operating</b> " and only if the status is " <b>not-operating</b> " because of the current state of its hosting-environment. In other cases for a " <b>zvm</b> " virtual server, the value of this property is null.  Values when non-null: <ul style="list-style-type: none"> <li>• "<b>host-env-not-activated</b>" - indicates the "<b>zvm</b>" virtual server is not operating because the hosting-environment is not activated</li> <li>• "<b>host-env-not-capable</b>" - indicates the "<b>zvm</b>" virtual server is not operating because the hosting-environment is not allowing communications. This can be due to a temporary condition at the support element for the associated CPC or can be due to SMAPI communication problems.</li> <li>• "<b>host-env-not-operating</b>" - indicates the "<b>zvm</b>" virtual server is not operating because the hosting-environment is not operating</li> </ul>

## Data model notes

For an x Hyp virtual server, the UUID assigned by zManager as the virtual server's object ID and thus provided as the virtual server's **object-id** property is also used as the virtual server's System Management BIOS (SMBIOS) UUID. As a result, this value is visible to software running within the virtual server as the UUID field of the SMBIOS System Identification (Type 1) structure. Guest operating system interfaces or utilities, such as the Linux **lspci** command, may be available to query this value. Because the object ID is available both from outside the guest as well as within, it can serve as a reliable correlating value for management application that have a need to associate data obtained from these two environments.

## Class specific additional properties

In addition to the properties defined in included schemas, this object includes the following additional class-specific properties.

**Note:** Many properties are only valid for virtual servers of specific "type". These value are only included in a Virtual Server object if the virtual server is of that type. For example a virtual server with type "**power-vm**" will define a **processing-mode** property ("**power-vm**" only) and **mac-prefix** property ("**power-vm**", "**zvm**") but not an **initial-share-mode** property ("**zvm**" only).

Other properties are only valid when mutable prerequisite properties have specific values. When such properties are not valid, their value is null. For instance a "**power-vm**" virtual server's **initial-virtual-processors** property value is null when the **processing-mode** value is "**dedicated**".

Table 65. Virtual Server object: class specific additional properties

Name	Qualifier	Type	Description	Supported "type" values
<b>type</b>	—	String Enum	Type of the virtual server. Values: <ul style="list-style-type: none"> <li>"<b>power-vm</b>" - a virtual server that has been defined on an IBM POWER7 blade</li> <li>"<b>x-hyp</b>" - a virtual server that has been defined on an IBM System x blade</li> <li>"<b>zvm</b>" - a virtual server that has been defined on an IBM z/VM operating system instance that is participating as an ensemble-managed Virtualization Host</li> <li>"<b>prsm</b>" - the virtual server representation of an LPAR image.</li> </ul>	All
<b>gpmp-status</b>	(pc)	String Enum	Status of the Guest Platform Management Provider (GPMP). Values: <ul style="list-style-type: none"> <li>"<b>unknown</b>" – Guest performance agent status could not be determined</li> <li>"<b>not-operating</b>" – Guest performance agent is not operating</li> <li>"<b>operating</b>" – Guest performance agent is operating</li> <li>"<b>status-check</b>" – There is a failure in communications between the virtual server and the guest performance agent.</li> </ul> <p>If virtual server type is "<b>prsm</b>", <b>gpmp-status</b> is only available if virtual server is running z/OS®.</p>	All
<b>cpu-perf-mgmt-enabled</b>	(w)(pc)	Boolean	If true, management of processor performance is enabled for this virtual server if management of processor performance is enabled at the ensemble level. <b>Note:</b> Management of processor performance at the ensemble level is enabled by virtual server type. See the ensemble object's <b>cpu-perf-mgmt-enabled-x-hyp</b> , <b>cpu-perf-mgmt-enabled-power-vm</b> , and <b>cpu-perf-mgmt-enabled-zvm</b> properties.  Prerequisite: <ul style="list-style-type: none"> <li>The parent virtualization host's <b>cpu-shares-supported</b> value is true.</li> </ul>	x-hyp, power-vm, zvm
<b>hostname</b>	(pc)	String	Virtual server host name. This data is only available if the Guest Platform Management Provider is running on the virtual server.	power-vm, x-hyp, zvm
<b>os-name</b>	(pc)	String	The name given to this system by its operating system. This data is only available if the Guest Platform Management Provider is running on the virtual server.	All
<b>os-type</b>	(pc)	String	The type of operating system that is running on the virtual server. This data is only available if the Guest Platform Management Provider is running on the virtual server.	All
<b>os-level</b>	(pc)	String	The release level of the operating system, as reported by the OS itself. This data is only available if the Guest Platform Management Provider is running on the virtual server.	All



Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>mac-prefix</b>	(pc)	mac-prefix Object	MAC address provides the means of identification that forwards frames within the LAN segment. This prefix is the first part of all MAC addresses assigned for the virtual server. The remaining bits are dynamically assigned when the virtualization host has to generate a MAC address for a network adapter.	power-vm, zvm
<b>processing-mode</b> <sup>1</sup>	(w)	String Enum	The manner in which virtual processors are associated with the physical processors available on the virtualization host. Values: <ul style="list-style-type: none"> <li>• <b>"shared"</b> - In shared mode, the virtual servers can use fractions of physical processors. The processor capacity is assigned in 0.1 units of physical processor, equivalent to 1.0 virtual processing unit. The virtual processor units can be shared among multiple virtual servers.</li> <li>• <b>"dedicated"</b> - In dedicated mode, the processors are assigned in whole units of physical processors. The processors that are dedicated to the virtual server cannot be used by other virtual servers. The virtual server cannot use any processors other than its own dedicated processors.</li> </ul>	power-vm
<b>minimum-dedicated-processors</b> <sup>1</sup>	(w)	Integer	Defines the minimum number of dedicated processors that the virtual server can use. <sup>3</sup>  Prerequisite: <b>processing-mode</b> is <b>"dedicated"</b> .  Limits: <ul style="list-style-type: none"> <li>• <math>\geq 1</math></li> <li>• <math>\leq</math> <b>maximum-allowed-dedicated-processors</b></li> </ul>	power-vm
<b>initial-dedicated-processors</b> <sup>1, 2</sup>	(w)(pc)	Integer	Defines the initial number of dedicated processors that the virtual server can use; the number of dedicated processors to be provided to the virtual server when it is next activated. <sup>3</sup>  Prerequisite: <b>processing-mode</b> is <b>"dedicated"</b> .  <b>"power-vm"</b> Limits: <ul style="list-style-type: none"> <li>• <math>\geq 1</math></li> <li>• <math>\leq</math> <b>maximum-allowed-dedicated-processors</b></li> <li>• <math>\geq</math> <b>minimum-dedicated-processors</b></li> </ul>	power-vm
<b>maximum-dedicated-processors</b> <sup>1</sup>	(w)	Integer	The upper limit for the number of dedicated processors to for the virtual server to consume. <sup>3</sup>  Prerequisite: <b>processing-mode</b> is <b>"dedicated"</b>  <b>"power-vm"</b> Limits: <ul style="list-style-type: none"> <li>• <math>\geq 1</math></li> <li>• <math>\leq</math> <b>maximum-allowed-dedicated-processors</b></li> <li>• <math>\geq</math> <b>minimum-dedicated-processors</b></li> <li>• <math>\geq</math> <b>initial-dedicated-processors</b></li> </ul>	power-vm

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>minimum-virtual-processors</b> <sup>1</sup>	(w)	Integer	<p>Defines the minimum number of virtual processors that the virtual server can use.<sup>3</sup></p> <p>Prerequisite: <b>processing-mode</b> is "shared".</p> <p>Limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq 1</math></li> <li>• <math>\leq 10 * \text{minimum-processing-units}</math></li> <li>• <math>\geq \text{minimum-processing-units}</math></li> </ul>	power-vm
<b>initial-virtual-processors</b> <sup>1,2</sup> ( <sup>1</sup> applies to PowerVM and x Hyp only, <sup>2</sup> applies only to PowerVM)	(w)	Integer	<p>Defines the initial number of virtual processors that the virtual server can use; the number of virtual processors to be provided to the virtual server when it is next activated.<sup>3</sup></p> <p>Prerequisites:</p> <ul style="list-style-type: none"> <li>• "x-hyp", "zvm": none</li> <li>• "power-vm": <b>processing-mode</b> is "shared"</li> </ul> <p>"power-vm" limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq 1</math></li> <li>• <math>\leq 10 * \text{initial-processing-units}</math></li> <li>• <math>\geq \text{initial-processing-units}</math></li> <li>• <math>\geq \text{minimum-virtual-processors}</math></li> </ul> <p>"x-hyp" limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq 1</math></li> <li>• <math>\leq \text{maximum-allowed-virtual-processors}</math></li> </ul> <p>"zvm" limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq 1</math></li> </ul>	power-vm, x-hyp, zvm
<b>maximum-virtual-processors</b> <sup>1</sup> ( <sup>1</sup> applies to PowerVM only)	(w)	Integer	<p>The upper limit for the number of virtual processors to for the virtual server to consume.<sup>3</sup></p> <p>Prerequisites:</p> <ul style="list-style-type: none"> <li>• "x-hyp", "zvm": none</li> <li>• "power-vm": <b>processing-mode</b> is "shared"</li> </ul> <p>"power-vm" limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq 1</math></li> <li>• <math>\leq 10 * \text{maximum-processing-units}</math></li> <li>• <math>\geq \text{maximum-processing-units}</math></li> <li>• <math>\geq \text{minimum-virtual-processors}</math></li> <li>• <math>\geq \text{initial-virtual-processors}</math></li> <li>• <math>\leq \text{maximum-allowed-virtual-processors}</math></li> </ul> <p>"zvm" limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq 1</math></li> <li>• <math>\geq \text{initial-virtual-processors}</math></li> <li>• <math>\leq \text{maximum-allowed-virtual-processors}</math></li> </ul>	power-vm, zvm

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>minimum-processing-units</b> <sup>1</sup>	(w)	Float	<p>The minimum number of processing units required for this virtual server to start running.</p> <p>Prerequisite: <b>processing-mode</b> is "shared".</p> <p>Limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq 0.1 * \text{minimum-virtual-processors}</math></li> <li>• <math>\leq \text{minimum-virtual-processors}</math></li> <li>• <math>\leq \text{maximum-allowed-processing-units}</math></li> <li>• Fixed to two decimal places</li> </ul>	power-vm
<b>initial-processing-units</b> <sup>1,2</sup>	(w)	Float	<p>The number of processing units representing the initial processor scheduling target for this virtual server. If resources are available, the virtual server may receive more than this amount.<sup>4</sup></p> <p>Prerequisite: <b>processing-mode</b> is "shared".</p> <p>Limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq .1 * \text{initial-virtual-processors}</math></li> <li>• <math>\leq \text{initial-virtual-processors}</math></li> <li>• <math>\leq \text{maximum-allowed-processing-units}</math></li> <li>• <math>\geq \text{minimum-processing-units}</math></li> <li>• Fixed to two decimal places</li> </ul>	power-vm
<b>maximum-processing-units</b> <sup>1</sup>	(w)	Float	<p>The maximum number of processing units that will be allocated to the virtual server (processor utilization capping).</p> <p>Prerequisite: <b>processing-mode</b> is "shared".</p> <p>If you have enabled processor management, Maximum processing units defines the upper limit for zManager. A virtual server with a capacity equal to this maximum value cannot receive resources from other virtual servers on the blade.</p> <p>Limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq 0.1 * \text{maximum-virtual-processors}</math></li> <li>• <math>\leq \text{maximum-virtual-processors}</math></li> <li>• <math>\leq \text{maximum-allowed-processing-units}</math></li> <li>• <math>\geq \text{minimum-processing-units}</math></li> <li>• <math>\geq \text{initial-processing-units}</math></li> <li>• Fixed to two decimal places</li> </ul>	power-vm

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>initial-share-mode</b>	(w)	String Enum	<p>Defines the virtual servers' initial share of system resources either relative to other virtual servers or absolutely. Values:</p> <ul style="list-style-type: none"> <li>"<b>relative</b>" - Grants all virtual servers different priorities for processor and I/O. A relative share allocates to a virtual server a portion of the total system resources minus those resources allocated to virtual servers with an absolute share. Also, a virtual server with a relative share receives access to system resources that is proportional with respect to other virtual servers with relative shares. For example, if a virtual server (VM1) has a relative share of 100, and a second virtual server (VM2) has a relative share of 200, VM2 receives twice as much access to system resources as VM1.</li> <li>"<b>absolute</b>" - Grants universal access and priority over all other virtual servers. An absolute share allocates to a virtual server an absolute percentage of all available system resources. For example, if you assign a virtual server an absolute share of 50%, CP allocates to that virtual server approximately 50% of all available resources (regardless of the number of other virtual servers running).</li> </ul>	zvm
<b>initial-shares</b>	(w)(pc)	Float/ Integer	<p>On x Hyp virtualization hosts:</p> <p>Defines the initial share value for the virtual server relative to other virtual servers on the virtualization host. This field can not be updated when the virtual server's <b>status</b> is "<b>operating</b>" and <b>cpu-perf-mgmt-enabled</b> is "<b>true</b>". This field expects an integer.</p> <p>Default value (used if <b>cpu-shares-supported</b> changes to true: 1024.</p> <p>Limits:</p> <ul style="list-style-type: none"> <li>&gt;= <b>minimum-shares</b></li> <li>&lt;= <b>maximum-shares</b></li> </ul> <p>Prerequisite:</p> <ul style="list-style-type: none"> <li>The parent virtualization host's <b>cpu-shares-supported</b> value is true.</li> </ul> <p>On z/VM virtualization hosts:</p> <p>If <b>initial-share-mode</b> is "<b>relative</b>", defines the initial share value for the virtual server relative to other virtual servers on the virtualization host. When <b>maximum-share-mode</b> is "<b>relative</b>", this value must be an Integer between 1 and min (maximum-shares, 10000).</p> <p>If <b>initial-share-mode</b> is "<b>absolute</b>", defines the initial share value for the virtual server absolutely. When <b>maximum-share-mode</b> is "<b>absolute</b>", this value must be a Number between 0.1 and min (maximum-shares, 100), fixed to one decimal place.</p>	x-hyp, zvm

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>share-limit</b>	(w)	String	<p>Define how the virtual server's share of system resources is limited. Values:</p> <ul style="list-style-type: none"> <li>• <b>"none"</b> - Specifies that a server share of processing resource is not limited.</li> <li>• <b>"soft"</b> - Specifies that the share of processing resource is limited but at times these servers can receive more than their limit, if no other server can use the available resources.</li> <li>• <b>"hard"</b> - Specifies that the share of processing resource is limited. These servers can not receive more than their limit.</li> </ul>	zvm
<b>maximum-share-mode</b>	(w)	String Enum	<p>Defines the virtual servers' maximum share of system resources either relative to other virtual servers or absolutely.</p> <p>Prerequisites: <b>share-limit</b> is <b>"soft"</b> or <b>"hard"</b>.</p> <p>Values:</p> <ul style="list-style-type: none"> <li>• <b>"relative"</b> - Grants all virtual servers different priorities for processor and I/O. A relative share allocates to a virtual server a portion of the total system resources minus those resources allocated to virtual servers with an absolute share. Also, a virtual server with a relative share receives access to system resources that is proportional with respect to other virtual servers with relative shares. For example, if a virtual server (VM1) has a relative share of 100, and a second virtual server (VM2) has a relative share of 200, VM2 receives twice as much access to system resources as VM1.</li> <li>• <b>"absolute"</b> - Grants universal access and priority over all other virtual servers. An absolute share allocates to a virtual server an absolute percentage of all available system resources. For example, if you assign a virtual server an absolute share of 50%, CP allocates to that virtual server approximately 50% of all available resources (regardless of the number of other virtual servers running).</li> </ul>	zvm
<b>minimum-shares</b>	(w)(pc)	Integer	<p>Defines the minimum share value for the virtual server on the virtualization host. This field can not be updated when the virtual server's <b>status</b> is <b>"operating"</b> and <b>cpu-perf-mgmt-enabled</b> is <b>"true"</b>. This field expects an integer.</p> <p>Limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq 2</math></li> <li>• <math>\leq</math> <b>initial-shares</b></li> </ul> <p>Prerequisite:</p> <ul style="list-style-type: none"> <li>• The parent virtualization host's <b>cpu-shares-supported</b> value is true.</li> </ul>	x-hyp

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>maximum-shares</b>	(w)(pc)	Float/ Integer	<p>On x Hyp virtualization hosts:</p> <p>Defines the maximum share value for the virtual server on the virtualization host. This field can not be updated when the virtual server's <b>status</b> is <b>"operating"</b> and <b>cpu-perf-mgmt-enabled</b> is <b>"true"</b>. This field expects an integer.</p> <p>Default value (used if <b>cpu-shares-supported</b> changes to true): 262144.</p> <p>Limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq</math> <b>initial-shares</b></li> <li>• <math>\leq</math> 262144</li> </ul> <p>Prerequisite:</p> <ul style="list-style-type: none"> <li>• The parent virtualization host's <b>cpu-shares-supported</b> value is true.</li> </ul> <p>On z/VM virtualization hosts:</p> <p>If <b>maximum-share-mode</b> is <b>"relative"</b>, defines the maximum share value for the virtual server relative to other virtual servers on the virtualization host. When <b>maximum-share-mode</b> is <b>"relative"</b>, this value must be an Integer between 1 and min (maximum-shares, 10000).</p> <p>If <b>maximum-share-mode</b> is <b>"absolute"</b>, defines the maximum share value for the virtual server absolutely. When <b>maximum-share-mode</b> is <b>"absolute"</b>, this value must be a Number between 0.1 and min (maximum-shares, 100), fixed to one decimal place.</p> <p>Prerequisites: <b>share-limit</b> is <b>"soft"</b> or <b>"hard"</b>.</p>	x-hyp, zvm
<b>minimum-memory<sup>1</sup></b>	(w)	Integer	<p>Minimum memory value for use by the virtual server, specified in megabytes (MB). Must be a multiple of virtualization host's <b>memory-increment-in-mega-bytes</b> value.</p> <p><b>"power-vm"</b> limits:</p> <ul style="list-style-type: none"> <li>• <math>\geq</math> <b>minimum-memory-size-for-virtual-server</b></li> <li>• multiple of <b>memory-increment-in-mega-bytes</b></li> </ul>	power-vm

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>initial-memory</b> <sup>1,2</sup> ( <sup>1</sup> applies to PowerVM and x Hyp only, <sup>2</sup> applies only to PowerVM)	(w)(pc)	Integer	Initial memory value for use by the virtual server, specified in MB. Must be a multiple of virtualization host's <b>memory-increment-in-mega-bytes</b> value.  <b>"power-vm"</b> limits: <ul style="list-style-type: none"> <li>• &gt;= <b>minimum-memory-size-for-virtual-server</b></li> <li>• multiple of <b>memory-increment-in-mega-bytes</b></li> <li>• &gt;= <b>minimum-memory</b></li> <li>• &lt;= <b>maximum-memory</b></li> <li>• &lt;= <b>maximum-memory-size-for-virtual-server</b></li> </ul> <b>"x-hyp"</b> limits: <ul style="list-style-type: none"> <li>• &gt;= <b>minimum-memory-size-for-virtual-server</b></li> <li>• multiple of <b>memory-increment-in-mega-bytes</b></li> <li>• &lt;= <b>maximum-memory-size-for-virtual-server</b></li> </ul> <b>"zvm"</b> limits: <ul style="list-style-type: none"> <li>• &gt;= <b>minimum-memory-size-for-virtual-server</b></li> <li>• multiple of <b>memory-increment-in-mega-bytes</b></li> <li>• &lt;= <b>maximum-memory</b></li> <li>• &lt;= <b>maximum-memory-size-for-virtual-server</b></li> </ul>	power-vm, x-hyp, zvm
<b>maximum-memory</b> <sup>1</sup> ( <sup>1</sup> applies to PowerVM only)	(w)	Integer	Maximum memory value for use by the virtual server, specified in MB. Must be a multiple of virtualization host's <b>memory-increment-in-mega-bytes</b> value.  <b>"power-vm"</b> limits: <ul style="list-style-type: none"> <li>• &gt;= <b>minimum-memory-size-for-virtual-server</b></li> <li>• multiple of <b>memory-increment-in-mega-bytes</b></li> <li>• &gt;= <b>minimum-memory</b></li> <li>• &gt;= <b>initial-memory</b></li> <li>• &lt;= <b>maximum-memory-size-for-virtual-server</b></li> </ul> <b>"zvm"</b> limits: <ul style="list-style-type: none"> <li>• &gt;= <b>minimum-memory-size-for-virtual-server</b></li> <li>• multiple of <b>memory-increment-in-mega-bytes</b></li> <li>• &lt;= <b>maximum-memory-size-for-virtual-server</b></li> </ul>	power-vm, zvm
<b>workloads</b>	(c)(pc)	Array of String/URI	The canonical URI path of each workload resource group to which the virtual server is assigned.	All
<b>network-adapters</b>	—	Array of objects	Array of nested Network Adapter objects defining the virtual server's network adapters. The nested objects are of type <b>network-adapter-power</b> , <b>network-adapter-x-hyp</b> , <b>network-adapter-zvm</b> or <b>network-adapter-prsm</b> depending on the type of the virtual server.	All
<b>virtual-disks</b>	—	Array of objects	Array of nested Virtual Disk objects defining the virtual server's virtual disks.	power-vm, x-hyp, zvm

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>mounted-media-name</b>	—	String (0-255)	The display name of the mounted ISO or null if no media is mounted.	power-vm, x-hyp
<b>boot-mode</b>	(w)(pc)	String Enum	The boot mode of the virtual server. Values: <ul style="list-style-type: none"> <li>• <b>"normal"</b> - The virtual server boots in normal mode.</li> <li>• <b>"sms"</b> - The virtual server boots to the System Management Services (SMS) menu.</li> <li>• <b>"diagnostic-default-boot-list"</b>- The sequence of devices read at startup.</li> <li>• <b>"diagnostic-stored-boot-list"</b> - The virtual server performs a service mode boot using the service mode boot list saved in NVRAM.</li> <li>• <b>"open-firmware-prompt"</b> - The virtual server boots to the open firmware prompt.</li> </ul>	power-vm



Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
boot-sequence	(w)(pc)	Array of String Enum	<p>List of boot sources. Values:</p> <ul style="list-style-type: none"> <li>• <b>"virtual-disk"</b> – virtual-disks will be tried in order</li> <li>• <b>"network-adapter"</b> – the first network-adapter will be used as a boot source</li> <li>• <b>"virtual-media"</b> – virtual media will be used as a boot source</li> </ul> <p>List must contain one and only one element for <b>type "power-vm"</b></p> <p><b>"virtual-disk"</b> is not a valid value if the virtual server's <b>virtual-disks</b> property is empty.</p> <p><b>"network-adapter"</b> is not a valid value if the virtual server's <b>network-adapter</b> property is empty.</p> <p><b>"power-vm"</b> network-adapter boot notes:</p> <ul style="list-style-type: none"> <li>• If boot-network-adapter-client-ip, boot-network-adapter-subnet-ip, boot-network-adapter-gateway-ip, and boot-network-adapter-server-ip are not defined, boot network settings defined via DHCP.</li> <li>• If boot-network-adapter-client-ip, boot-network-adapter-subnet-ip, and boot-network-adapter-server-ip are defined, but boot-network-adapter-gateway-ip is not defined, boot network settings are manually defined and netboot is limited to its own subnet.</li> <li>• If boot-network-adapter-client-ip, boot-network-adapter-subnet-ip, boot-network-adapter-gateway-ip, and boot-network-adapter-server-ip are all defined, boot network settings are manually defined and netboot is not limited to its own subnet.</li> </ul> <p><b>"x-hyp"</b> virtual-disk boot notes:</p> <ul style="list-style-type: none"> <li>• Booting is always enabled for Virtual DVDs (virtual media), virtual disks, and virtual network adapters if devices of these types are defined for the virtual server. Although it is possible to change the order in which booting is attempted, it is not possible to disable booting from devices of these types if they are present for the virtual server. See the text directly after this table for considerations that apply to the value of the boot-sequence property.</li> <li>• If the virtual server resides on a node that is being managed by Support Element Version 2.11.1, a virtual network interface (if present) must be either the first or last entry in the virtual server's boot sequence. That is, the value <b>"network-adapter"</b> must appear as either the first or last entry in the <b>boot-sequence</b> property.</li> </ul> <p><b>"x-hyp"</b> virtual-disk boot notes:</p> <ul style="list-style-type: none"> <li>• If the virtualization host's mixed-mode-boot-restriction value is true and the virtual-disks property defines both <b>"virtio"</b> and <b>"ide"</b> virtual disks, the virtual server will only boot from an <b>"ide"</b> virtual-disk when the <b>boot-sequence</b> value contains <b>"virtual-disk"</b>.</li> </ul>	power-vm, x-hyp

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>boot-network-adapter-client-ip</b>	(w)	String (0-39)	The IP address of the virtual server used during network boot in dotted-decimal form ("nnn.nnn.nnn.nnn").  Prerequisites: <b>boot-sequence</b> contains " <b>network-adapter</b> ".	power-vm
<b>boot-network-adapter-subnet-ip</b>	(w)	String (0-17)	The IPv4 subnet mask of the network used during network boot of the virtual server in dotted-decimal form ("nnn.nnn.nnn.nnn").  Prerequisites: <b>boot-sequence</b> contains " <b>network-adapter</b> ".	power-vm
<b>boot-network-adapter-gateway-ip</b>	(w)	String (0-39)	The IP address of the gateway system used during network boot of the virtual server in dotted-decimal form ("nnn.nnn.nnn.nnn").  Prerequisites: <b>boot-sequence</b> contains " <b>network-adapter</b> ".	power-vm
<b>boot-network-adapter-server-ip</b>	(w)	String (0-39)	The IP address of the boot server on which the disk boot source files reside in dotted-decimal form ("nnn.nnn.nnn.nnn").  Prerequisites: <b>boot-sequence</b> contains " <b>network-adapter</b> ".	power-vm
<b>keylock</b>	(w)	String Enum	Indicates the state of the system key lock at boot time. Value: <ul style="list-style-type: none"> <li>• <b>"normal"</b> – Enables a regular operating system boot of the virtual server including all services, under program control.</li> <li>• <b>"manual"</b> - Requires system operator to manually boot the virtual server. This setting is useful for service scenarios. The virtual server boots to the diagnostics menu, then you have to open a console to the virtual server and work with the displayed menu to choose the next steps.</li> </ul>	power-vm
<b>auto-start</b>	(w)(pc)  (w only for PowerVM and x Hyp)	Boolean	If true, this virtual server is automatically started when its hosting virtualization host is started.	All
<b>dlpar-enabled</b>	(w)	Boolean	If true, this virtual server is configured for Dynamic Logical Partitioning. Note: This setting does not enable any DLPAR prerequisites (such as the Resource Monitoring and Control network).	power-vm
<b>dlpar-active</b>	—	Boolean	If true, DLPAR support is enabled and specific virtual server properties may be updated when the virtual server status is <b>"operating"</b> .	power-vm
<b>gpmp-support-enabled</b>	(w)(pc)  (ro for PR/SM)	Boolean	If true, Guest Platform Management Provider support is enabled, allowing the GPMP to gather performance data for work running on the virtual server. <b>Note:</b> This only defines if support is enabled, not if it is active.	All

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>gpmp-version</b>	—	String	Guest Platform Management Provider version or keyword <b>"unavailable"</b> if not known.  If virtual server type is <b>"prsm"</b> , <b>"gpmp-version"</b> is only available if virtual server is running z/OS.	power-vm, x-hyp, zvm
<b>password</b>	(w)	String (1-8)	The password or passphrase to be used for authentication. Password must meet the guidelines defined by the VM system. 1-8 uppercase characters long.	zvm
<b>privilege-classes</b>	(w)	String (1-32)	String defining z/VM CP privilege classes denoted by the letters A through Z (uppercase), the numbers 1 through 6, and the word "ANY". 1-32 characters long	zvm
<b>ipl-device</b>	(w)(pc)	String (1-8)	Specifies the virtual device that you want to IPL. It is the virtual device number or sysname of the virtual device. <sup>5</sup>  A virtual device number is a four character hex number.  The sysname is a 1 - 8 uppercase alphanumeric-character name of the named saved system you want to IPL.	zvm
<b>ipl-load-parameters</b>	(w)(pc)	String (0-8)	The z/VM LOADPARM option value; used to pass a load parameter of up to 8 bytes of data to the operating system you are IPLing. <sup>5</sup>  0-8 characters long, character may be uppercase alphanumeric, a space, or a period	zvm
<b>ipl-parameters</b>	(w)(pc)	String (0-64)	String defining the z/VM IPL parameter list. <sup>5</sup>  0:64 characters.	zvm
<b>associated-logical-partition</b>	—	String/URI	The canonical URI path for the Logical Partition that is the PR/SM virtual server.	prsm
<b>inband-monitoring-enabled</b>	(w)(pc)	Boolean	If true, in-band monitoring support is enabled, allowing the hypervisor to gather performance metrics for the virtual server.  Prerequisites: parent virtualization host's <b>inband-monitoring-supported</b> value is true.	power-vm, x-hyp

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>keyboard-language</b>	(w)(pc)	String	<p>String describing the locale required by graphical consoles connecting to the virtual server.</p> <p>The value is <b>null</b> if the parent virtualization host's <b>supported-keyboard-languages</b> array is <b>null</b> or empty.</p> <p>String value is the locale's language, country, and variant values separated by underscores:</p> <ul style="list-style-type: none"> <li>• The language value is either an empty string or a lowercase ISO 639 language code (2-4 characters).</li> <li>• The country value is either an empty string or an uppercase ISO 3166 two-letter code.</li> <li>• The variant value is either an empty string or a string defining variations for a language/country pair. For example, Serbian locales have Cyrillic ("Cyril") and Latin ("Latn") variants.</li> </ul> <p>If the language is an empty string, the string will begin with an underscore. If both the language and country fields are empty strings, the value will be an empty string.</p> <p>Example values: "en", "en_GB", "en_US", "sr_RS", "sr_RS_Cyrl", and "sr_RS_Latn".</p> <p>Value must be defined in the parent virtualization host's <b>supported-keyboard-languages</b> value array.</p>	x-hyp
<b>avail-status</b>	(pc)	Array of objects	<p>The list of availability statuses of the virtual server, each with reasons explaining the status, in order of severity. The list will always contain at least one element in the form of a virtual server availability status with reasons nested object, as described in Table 66 on page 278.</p>	
<b>acceptable-avail-status</b>	(w) (pc)	Array of String Enum	<p>The set of <b>avail-status</b> values that the virtual server can be in and be considered to be in an acceptable (not alert causing) state. A virtual server <b>avail-status</b> value of <b>"not-supported"</b> has no impact on a virtual server's acceptable state.</p> <p>Prerequisites: The virtual server's <b>avail-status</b> is not <b>"not-supported"</b>.</p> <p>Values may not contain <b>"not-supported"</b>.</p> <p>Refer to the virtual server availability status with reasons nested object's <b>avail-status</b> property described in Table 66 on page 278 for possible values.</p>	
<b>workload-element-groups</b>	(pc)	Array of String/URI	<p>The canonical URI path of each workload element group to which the virtual server is assigned.</p>	
<b>perf-policies</b>	(pc)	Array of objects	<p>The list of workloads and performance policies that are active on the virtual server, each with the activation status of the policy on the virtual server. The list will contain one element for each workload with a performance policy that has been activated on the virtual server in the form of a virtual server performance policy nested object, as described in Table 67 on page 279.</p>	All

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>avail-policies</b>	(pc)	Array of objects	The list of workloads and availability policies that are active on the virtual server. The list will contain one element for each workload to which the virtual server is assigned in the form of a virtual server availability policy nested object, as described in Table 68 on page 279.	
<b>power-vm-partition-id</b>	—	Integer	The identifier of the PowerVM partition in which the virtual server is running. A virtual server is assigned to a PowerVM partition during the virtual server activation process, and may be assigned to a different partition each time it is activated. Therefore, this property has a value only when the virtual server's <b>status</b> property is <b>"operating"</b> . If the virtual server is in some other status, the value of this property is null.	power-vm
<b>shutdown-timeout-source</b>	(w)(pc)	String Enum	The source of the <b>shutdown-timeout</b> property. Values: <ul style="list-style-type: none"> <li>• <b>"virtualization-host"</b> - Use the <b>virtual-server-shutdown-timeout</b> property defined in the virtualization host that contains this virtual server.</li> <li>• <b>"virtual-server"</b> - Use the <b>shutdown-timeout</b> property defined for this virtual server.</li> </ul>	power-vm, x-hyp
<b>shutdown-timeout</b>	(w)(pc)	Integer	Amount of time, in seconds, to allow a virtual server to shut down. After the elapsed time has passed, the virtual server will be forcefully terminated.  The value may be -1 to indicate to wait "forever" or any integer value between 0 and 86400 to specify an exact wait time in seconds.  This value is only used when the <b>shutdown-timeout-source</b> property is set to <b>"virtual-server"</b> .	power-vm, x-hyp
<b>gpmp-network-adapter</b>	—	Object	Nested network adapter object defining the virtual server's GPMP network adapter. For a virtual server of type <b>"x-hyp"</b> the nested object is of type <b>network-adapter-x-hyp</b> . For a virtual server of type <b>"power-vm"</b> the nested object is of type <b>network-adapter-power</b> .	power-vm, x-hyp

Table 65. Virtual Server object: class specific additional properties (continued)

Name	Qualifier	Type	Description	Supported "type" values
<b>Notes:</b>				
<p><sup>1</sup> These properties represent configuration values for the virtual server that are used to establish the virtual server's initial configuration when it is activated.</p> <p><sup>2</sup> These properties update runtime values when the virtual server status is <b>"operating"</b> and DLPAR is active (dlpar-active value is true). If DLPAR is not active, updating these properties will not affect runtime values. DLPAR support is only available for <b>"power-vm"</b>.</p> <p><sup>3</sup> Note: Because zManager-managed <b>"power-vm"</b> virtual servers consume resources on the Power ASB only when activated, it is possible to overcommit and define a set of virtual servers on a virtualization host that together require more virtual processors than can be simultaneously supported. PowerVM resource limits are enforced at the time that a virtual server is activated.</p> <p><sup>4</sup> This value is a configuration value that is used to set the scheduling goal when the virtual server is activated. If CPU performance management is active for this virtual server, the run time value may be adjusted as the virtual server operates. This will not change the "initial" value and thus does not trigger property change event. The current processing unit goal for the virtual server will be available as a virtual server metric.</p> <p><sup>5</sup> The maximum length of the IPL Directory statement resulting from the ipl-device, ipl-load-parameters, and ipl-parameters values may not exceed 72 characters. Refer to the <i>z/VM CP Planning and Administration Guide</i> for details on the IPL Directory statement.</p>				

The following considerations apply to the value of the **boot-sequence** property:

- Because a virtual server always has a virtual DVD device, the value **"virtual-media"** must always appear somewhere within the value of the **boot-sequence** property.
- If the virtual server has one or more virtual disks defined, then the value **"virtual-disk"** must appear within **boot-sequence** property. When the first virtual disk is defined for a virtual server, the **boot-sequence** property is updated to add **"virtual-disk"** as the last entry in the list. When the last virtual disk is removed, the value **"virtual-disk"** is removed from the **boot-sequence** property.
- Similarly, if the virtual server has one or more virtual network interfaces defined, then the value **"network-adapter"** must also appear within the **boot-sequence** property. When the first virtual network interface is defined for a virtual server, the **boot-sequence** property is updated to add **"network-adapter"** as the last entry in the list. When the last virtual network interface is removed, the value **"network-adapter"** is removed from the **boot-sequence** property.

The virtual server availability status with reasons nested object is nested within a Virtual Server object to encapsulate the availability status with the reasons explaining why it is in that state.

Table 66. Virtual Server object: virtual server availability status with reasons nested object properties

Name	Type	Description
avail-status	String Enum	<p>The availability status of the virtual server, determined by its operational status, its GPMP, its CPC, and its OS. The possible values are as follows:</p> <ul style="list-style-type: none"> <li>• <b>"available"</b> - The virtual server is considered available.</li> <li>• <b>"exposed"</b> - The virtual server is considered exposed.</li> <li>• <b>"critical"</b> - The virtual server is considered critically exposed.</li> <li>• <b>"not-available"</b> - The virtual server is considered not available.</li> <li>• <b>"not-supported"</b> - The virtual server is from a CPC that does not support Ensemble Availability Management.</li> </ul>

Table 66. Virtual Server object: virtual server availability status with reasons nested object properties (continued)

Name	Type	Description
reasons	String Enum	<p>The list of reasons given for the availability status of the virtual server. The possible values are as follows:</p> <ul style="list-style-type: none"> <li>• <b>"vs"</b> - The virtual server's operational status is contributing to this availability status.</li> <li>• <b>"cpc"</b> - The virtual servers CPC's operational status is contributing to this availability status.</li> <li>• <b>"gpmp"</b> - The virtual server's GPMP's status is contributing to this availability status.</li> <li>• <b>"node"</b> - The operational status of the node is contributing to this availability status.</li> </ul> <p>The list may be empty if no reasons are given.</p>

The virtual server performance policy nested object is nested within a Virtual Server object to encapsulate the performance policies that are active on the virtual server and their activation status with respect to the virtual server.

Table 67. Virtual Server object: virtual server performance policy nested object

Name	Type	Description
workload-uri	String/URI	The canonical URI path of the workload resource group to which the performance policy identified by <b>policy-name</b> is assigned.
policy-name	String	The <b>name</b> property of the Performance Policy object.
activation-status	String Enum	<p>The status of activating the performance policy on the virtual server. Possible values are:</p> <ul style="list-style-type: none"> <li>• <b>"initializing"</b> - Status is initializing (not yet known).</li> <li>• <b>"successful"</b> - The policy has been successfully activated on the virtual server.</li> <li>• <b>"failed"</b> - Policy activation failed.</li> <li>• <b>"pending"</b> - Policy activation is pending (in progress).</li> </ul>

The virtual server availability policy nested object is nested within a Virtual Server object to encapsulate the availability policies that are active on the virtual server.

Table 68. Virtual Server object: virtual server availability policy nested object

Name	Type	Description
workload-uri	String/URI	The canonical URI path of the workload resource group to which the availability policy identified by <b>policy-uri</b> is assigned.
policy-uri	String/URI	Canonical URI path of the Availability Policy object, in the form <code>/api/workload-resource-groups/{workload-id}/availability-policies/{policy-id}</code> .
avail-status-impact-exclusion	Boolean	<p>The <b>avail-status</b> value of the virtual server is excluded from impacting the availability status of the workload identified by <code>{workload-uri}</code> by the <b>avail-status-impact-exclusion</b> property of the availability policy identified by <code>{policy-uri}</code>.</p> <p>Value is always false if the <b>class</b> is <b>"workload-resource-group"</b> or <b>"workload-element-group"</b> or if the <b>class</b> is <b>"virtual-server"</b> and the <b>parent-class</b> is <b>"workload-element-group"</b>.</p>

**mac-prefix object:** A mac-prefix object defines a "**power-vm**" or "**zvm**" virtual server's mac-prefix property value.

Table 69. mac-prefix object properties

Name	Type	Description
<b>mac-address</b>	String (17)	The MAC address represented as 6 groups of two lower-case hexadecimal digits separated by colons ( : ), e.g. "01:23:45:67:89:ab". Length is 17 characters. The MAC address uses the ensemble prefix.
<b>prefix-length</b>	Integer	The bit length of the MAC address prefix. This is a 2-digit value with these parameters in the range 12-44.

**network-adapter objects:** The bit length of the MAC address prefix. This is a 2-digit value with these parameters in the range 12-44.

The network-adapter-power object defines a network adapter of a virtual-server of type "**power-vm**".

Table 70. network-adapter-power object properties

Name	Qualifier	Type	Description
<b>element-id</b>	—	String	Unique ID for the virtual network adapter within the scope of the containing virtual server.  It is a randomly generated integer value when the object is created. Currently this value will change if this object is modified or the group of network adapters of the virtual server is reordered.
<b>element-uri</b>	—	String/ URI	The canonical URI path for the virtual network adapter is of the form /api/virtual-servers/{virtual-server-id}/network-adapters/{element-id}, where {virtual-server-id} is the <b>object-id</b> of the virtual server.
<b>network-uri</b>	(w)	String/ URI	The canonical URI path for the associated virtual network or <b>null</b> if the network adapter is not connected to a virtual network.
<b>mac-address</b>	—	String (17)	The MAC address of the network adapter represented as 6 groups of two lowercase hexadecimal digits separated by colons ( : ), e.g. "01:23:45:67:89:ab". Length is 17 characters. The MAC address uses the ensemble prefix.
<b>adapter-type</b>	—	String Enum	The type of the network adapter. The possible values are as follows: <ul style="list-style-type: none"> <li>• "<b>regular</b>" - The network adapter is the regular network adapter ( non-monitoring).</li> <li>• "<b>gpmp</b>" - The network adapter is the resource monitoring network adapter used for Guest Platform Management Provider (GPMP).</li> </ul>

The network-adapter-x-hyp object defines a network adapter of a virtual-server object of type "**x-hyp**".

Table 71. network-adapter-x-hyp object properties

Name	Qualifiers	Type	Description
<b>element-id</b>	—	String	Unique ID for the virtual network adapter within the scope of the containing virtual server.
<b>element-uri</b>	—	String/ URI	The canonical URI path for the virtual network adapter is of the form /api/virtual-servers/{virtual-server-id}/network-adapters/{element-id}, where {virtual-server-id} is the <b>object-id</b> of the virtual server.



Table 71. network-adapter-x-hyp object properties (continued)

Name	Qualifiers	Type	Description
<b>emulation-mode</b>	—	String Enum	The network adapter emulation mode. Values: <ul style="list-style-type: none"> <li>• "e1000" – Intel E1000</li> <li>• "rtl8139" – Realtek 8139</li> <li>• "virtio" – RedHat virtio</li> </ul>
<b>network-uri</b>	(w)	String/ URI	The canonical URI path for the associated virtual network or <b>null</b> if the network adapter is not connected to a virtual network.
<b>mac-address</b>	—	String (17)	The MAC address of the network adapter represented as 6 groups of two lowercase hexadecimal digits separated by colons (:), e.g. "01:23:45:67:89:ab". Length is 17 characters.
<b>adapter-type</b>	—	String Enum	The type of the network adapter. The possible values are as follows: <ul style="list-style-type: none"> <li>• "regular" - The network adapter is the regular network adapter ( non-monitoring).</li> <li>• "gpmp" - The network adapter is the resource monitoring network adapter used for Guest Platform Management Provider (GPMP).</li> </ul>

The network-adapter-zvm object defines a network adapter of a virtual-server object of type "zvm".

**Note:** Some properties are only valid for network adapters of specific "type". These value are only included in a network-adapter-zvm object if the network adapter is of that type. For example, a network adapter with type **rmc** will not define an **interface-type** property.

Other properties are only valid when mutable prerequisite properties have specific values. When such properties are not valid, their value is null. For instance a network-adapter-zvm's **real-device-address** property value is null when the **interface-type** value is "virtual-iedn".

Table 72. network-adapter-zvm object properties

Name	Qualifier	Type	Description
<b>element-id</b>	—	String (1-4)	Unique ID for the virtual network adapter within the scope of the containing virtual server. This element-id is actually the <b>virtual-device-address</b> listed in this table.  This element ID is not immutable. It will be changed if the <b>virtual-device-address</b> of this object is modified.
<b>element-uri</b>	—	String/ URI	The canonical URI path for the virtual network adapter is of the form /api/virtual-servers/{virtual-server-id}/network-adapters/{element-id}, where {virtual-server-id} is the <b>object-id</b> of the virtual server.  This URI is not immutable because the <b>element-id</b> component of it can change. See the description for the <b>element-id</b> property.
<b>virtual-device-address</b>	(w)	String (1-4)	Virtual device address 1-4 character hex string
<b>device-count</b>	—	Integer	The number of device addresses to reserve

Table 72. network-adapter-zvm object properties (continued)

Name	Qualifier	Type	Description
<b>type</b>	(ro)	String Enum	<p>Network adapter type. Values:</p> <ul style="list-style-type: none"> <li>"osx": OSX is a CHPID type. A network adapter that connects through an OSX CHPID to provide connectivity to a virtual network of the IEDN (Intra-Ensemble Data Network).</li> <li>"osd": OSD is a CHPID type. A network adapter that connects through an OSD CHPID to provide connectivity to external network.</li> <li>"iqd": IQD is a CHPID type. A network adapter that connects through an IQD CHPID to provide connectivity to a HiperSockets™ network that is not part of the IEDN.</li> <li>"iqdx": IQDX is a CHPID type that provides connectivity to the IEDN HiperSockets network.</li> <li>"rmc": Identifies the network adapter is the Remote Monitoring and Control network adapter.</li> </ul>
<b>interface-type</b>	(w)	String Enum	<p>Network adapter switch type. Legal values based on <b>type</b> value.</p> <p>osx values:</p> <ul style="list-style-type: none"> <li>"none"</li> <li>"virtual-iedn": Virtual IEDN switch</li> <li>"physical-iedn": Physical IEDN switch</li> </ul> <p>osd values:</p> <ul style="list-style-type: none"> <li>"none"</li> <li>"virtual-iedn": Virtual IEDN switch</li> <li>"virtual-qdio": Virtual QDIO switch</li> <li>"physical-qdio": Physical QDIO switch</li> </ul> <p>iqd values:</p> <ul style="list-style-type: none"> <li>"none"</li> <li>"physical-iqdn": Physical IQDN switch</li> </ul> <p>iqdx values:</p> <ul style="list-style-type: none"> <li>"none"</li> <li>"physical-iedn": Physical IQDX switch</li> </ul>
<b>real-device-address</b>	(w)	String (1-4)	<p>The device address that has been assigned to the port on the switch that the "dedicated" NIC is mapped to.</p> <p>Prerequisites: <b>interface-type</b> is "physical-iedn", "physical-qdio", or "physical-iqdn"</p> <p>1-4 character hex number (range 0-FFFF).</p>
<b>switch-uri</b>	(w)	String/URI	<p>The switch to which the network adapter is connected.</p> <p>Prerequisites: <b>interface-type</b> is not "none".</p> <ul style="list-style-type: none"> <li>If <b>interface-type</b> is "virtual-iedn" or "virtual-qdio", the value is the <b>element-uri</b> of the <b>virtual-switch</b> object in use by the z/VM network adapter.</li> <li>If <b>interface-type</b> is "physical-iedn" and type is "osx", the value is the <b>element-uri</b> of the <b>network-adapter-prsm</b> object in use by the z/VM network adapter.</li> </ul>

Table 72. network-adapter-zvm object properties (continued)

Name	Qualifier	Type	Description
<b>port-mode</b>	(w)	String Enum	<p>The port mode, or null if the virtual-switch to which the network adapter was connected no longer exists.</p> <p>Prerequisites: <b>interface-type</b> is "virtual-iedn" or "virtual-qdio". If <b>interface-type</b> is "virtual-qdio", the virtual switch identified by <b>switch-uri</b> must be VLAN aware (its <b>is-vlan-aware</b> property is true).</p> <p>Values:</p> <ul style="list-style-type: none"> <li>"trunk": When the switch port is configured in trunk mode, it will allow the flow of traffic from multiple virtual networks (i.e. VLANs). The port must be configured with those virtual networks.</li> <li>"access": When the switch port is configured in access mode, it will support a single virtual network. Traffic from the virtual server's network adapter will be tagged with the virtual network configured for this port, and traffic destined to the virtual server on this port will be verified that it is tagged with the configured virtual network.</li> </ul>
<b>vlan-ids</b>	(w)	String (0-19)	<p>A space-delimited String defining the VLAN IDs.</p> <p>Prerequisites: <b>interface-type</b> is "virtual-qdio"</p> <p>Value may contain one of the following:</p> <ul style="list-style-type: none"> <li>zero to four VLAN IDs</li> <li>zero to two ranges of VLAN IDs.</li> </ul> <p>A VLAN ID is a decimal integer from 1-4094.</p> <p>A range is defined by two VLAN IDs separated by a hyphen.</p> <p>If <b>port-mode</b> is "access" this string must define a single VLAN ID.</p> <p>Examples: "", "0 10", "0 10 100 2000", "0-10 2000-2000", "0-10 1000-2000"</p>
<b>network-uris</b>	(w)	Arrays of String/URI	<p>A list of associated Virtual Networks. Each item represents the canonical URI path for the associated virtual network.</p> <p>Prerequisites: <b>interface-type</b> is "virtual-iedn" or "physical-iedn".</p> <p>List must contain at least one network URI. If <b>port-mode</b> is "access", list must contain exactly one URI.</p>

The network-adapter-prsm object defines a network adapter of a virtual-server object of type "prsm".

Table 73. network-adapter-prsm object properties

Name	Qualifier	Type	Description
<b>element-id</b>	—	String	Unique ID for the virtual network adapter within the scope of the containing virtual server.
<b>element-uri</b>	—	String/URI	The canonical URI path for the virtual network adapter is of the form /api/virtual-servers/{virtual-server-id}/network-adapters/{element-id}, where {virtual-server-id} is the <b>object-id</b> of the virtual server.
<b>type</b>	—	String Enum	The physical switch type. Either "osx" or "idqx".

Table 73. *network-adapter-prsm* object properties (continued)

Name	Qualifier	Type	Description
css	—	String (1)	The channel subsystem ID
chpid	—	String (2)	The channel path ID
network-uris	(w)	Array of String/ URI	A list including the canonical URI path for each associated virtual network.

**Virtual disk objects:** A Virtual Disk is virtual storage space provided by a virtualization host to a guest virtual server. A Virtual Disk is based upon a storage resource, but may be further virtualized by a Virtualization Host.

Every virtual disk object contains the following base properties in addition to type-specific properties:

Table 74. *Virtual disk object properties*

Name	Qualifier	Type	Description
element-id	—	String (36)	The unique identifier for the virtual disk instance. This identifier is in the form of a UUID.
element-uri	—	String/URI	Canonical URI path of the virtual disk object, in the form <code>/api/virtual-servers/{virtual-server-id}/virtual-disks/{element-id}</code> , where <code>{virtual-server-id}</code> is the <b>object-id</b> of the virtual server.
name	(w)	String (1-64)	The name of the virtual disk. It must consist only of alphanumeric characters, spaces and the following special characters: “_.”, and it must begin with an alphabetic character.
description	(w)	String (0-1024)	The description for the virtual disk. It must consist only of alphanumeric characters, spaces and the following special characters: “_.”.
type	—	String Enum	The type of virtual disk. Values: <ul style="list-style-type: none"> <li>• <b>"fullpack"</b> - A virtual disk that is backed by an entire virtualization host storage resource</li> <li>• <b>"storage-group-based"</b> - A virtual disk that is backed by resources allocated from a virtualization host storage group. This type only applies to virtual disks of a virtual server whose <b>type</b> property is <b>"zvm"</b>.</li> <li>• <b>"linked"</b> - A virtual disk that is linked to a virtual disk owned by another virtual server. This type only applies to virtual disks of a virtual server whose <b>type</b> property is <b>"zvm"</b>.</li> </ul>
size	—	Long	The size of the virtual disk, in bytes.
owner	—	String/URI	Canonical URI path of the virtual server that owns this virtual disk.
emulation-mode	(w)	String Enum	The disk I/O emulation mode. Values: <ul style="list-style-type: none"> <li>• <b>"virtio"</b> – This virtual disk emulates VIRTIO</li> <li>• <b>"ide"</b> – This virtual disk emulates IDE</li> <li>• <b>"virtio-scsi"</b> – This virtual disk emulates VIRTIO-SCSI</li> </ul> <p>Only valid for virtual disks of a virtual server whose <b>type</b> property is <b>"x-hyp"</b>.</p>

A fullpack-virtual-disk object contains information about a fullpack virtual disk of a non-z/VM virtual server (i.e., the virtual server's **type** property is not **"zvm"**, and the virtual disk's **type** property is **"fullpack"**).

In addition to base properties, a fullpack-virtual-disk object contains the following properties:

Table 75. fullpack-virtual-disk object properties

Name	Qualifier	Type	Description
<b>backing-virtualization-host-storage-resource</b>	(w)	String/URI	Canonical URI path of the Virtualization Host Storage Resource object that backs this virtual disk.

A fullpack-virtual-disk-zvm object contains information about a z/VM virtual server's fullpack virtual disk (i.e., the virtual server's **type** property is "zvm", and the virtual disk's **type** property is "fullpack").

In addition to base properties, a fullpack-virtual-disk-zvm object contains the following properties:

Table 76. fullpack-virtual-disk-zvm object properties

Name	Qualifier	Type	Description
<b>device-address</b>	—	String (1-4)	The virtual device address. This is the device address by which the virtual server knows this virtual disk. The string form of a 1-4 digit hexadecimal number.
<b>access-mode</b>	(w)	String Enum	The access mode describing the virtual server's permission to read and/or write to this virtual disk. The values are defined in Table 79 on page 286.
<b>read-password</b>	(w)	String (0-8)	The read password for this virtual disk. Characters must be uppercase.
<b>write-password</b>	(w)	String (0-8)	The write password for this virtual disk. Characters must be uppercase.
<b>multi-password</b>	(w)	String (0-8)	The multiple-write password for this virtual disk. Characters must be uppercase.
<b>backing-virtualization-host-storage-resource</b>	—	String/URI	Canonical URI path of the Virtualization Host Storage Resource object that backs this virtual disk

A storage-group-based-virtual-disk object contains information about a virtual server's storage-group-based virtual disk (i.e., the virtual disk's **type** property is "storage-group-based").

In addition to base properties, a storage-group-based-virtual-disk object contains the following properties:

Table 77. storage-group-based-virtual-disk object properties

Name	Qualifier	Type	Description
<b>device-address</b>	—	String (1-4)	The virtual device address. This is the device address by which the virtual server knows this virtual disk. The string form of a 1-4 digit hexadecimal number.
<b>access-mode</b>	(w)	String Enum	The access mode describing the virtual server's permission to read and/or write to this virtual disk. The values are defined Table 79 on page 286.
<b>read-password</b>	(w)	String (0-8)	The read password for this virtual disk. Characters must be uppercase.
<b>write-password</b>	(w)	String (0-8)	The write password for this virtual disk. Characters must be uppercase.
<b>multi-password</b>	(w)	String (0-8)	The multiple-write password for this virtual disk. Characters must be uppercase.

Table 77. storage-group-based-virtual-disk object properties (continued)

Name	Qualifier	Type	Description
<b>backing-storage-group</b>	—	String/URI	Canonical URI path of the Virtualization Host Storage Group object that backs this virtual disk.

A linked-virtual-disk object contains information about a virtual server's linked virtual disk (i.e., the virtual disk's **type** property is **"linked"**).

In addition to base properties, a linked-virtual-disk object contains the following properties:

Table 78. linked-virtual-disk object properties

Name	Qualifier	Type	Description
<b>device-address</b>	—	String (1-4)	The virtual device address. This is the device address by which the virtual server knows this virtual disk. The string form of a 1-4 digit hexadecimal number.
<b>access-mode</b>	(w)	String Enum	The access mode describing the virtual server's permission to read and/or write to this virtual disk. The values are defined in Table 79.
<b>base-virtual-disk</b>	—	String/URI	Canonical URI path of the virtual disk to which this virtual disk is ultimately linked.
<b>source-virtual-disk</b>	—	String/URI	Canonical URI path of the virtual disk to which this virtual disk is directly linked.

The possible access modes for a virtual disk of a z/VM virtual server are listed in the following table. This table is effectively a String enumeration of the valid values for the **access-modes** property of a virtual disk object.

Table 79. Valid values for the access-modes property of a virtual disk object

Type	Description
String Enum	The virtual disk's access mode. These modes correspond exactly to the disk access modes in the z/VM user directory. Values: <ul style="list-style-type: none"> <li>• <b>"read-only"</b></li> <li>• <b>"read-write"</b></li> <li>• <b>"multi-write"</b></li> <li>• <b>"unsupported"</b> – any access mode other than the supported ones listed here. A virtual disk created by some means other than zManager-provided functions could be created with such an access mode.</li> </ul>

#### Usage notes for virtual disks:

- For virtual disk objects whose **type** property is **"linked"**, the “source” virtual disk is always the same as the “base” virtual disk. The base virtual disk property is provided for the case where a virtual disk is a link to a virtual disk which is itself a link to some other virtual disk. As such a configuration is not supported by zManager, base and source will be identical.

## Operations

If a virtual server operation accesses a z/VM virtualization host and encounters an error while communicating with the virtualization host via SMAPI, the response body is a SMAPI Error Response Body.

## List Virtual Servers of a zBX (Node)

The **List Virtual Servers of a zBX (Node)** operation lists the virtual servers managed by the zBX node with the given identifier.

### HTTP method and URI

**GET** /api/zbx/{zbx-id}/virtual-servers

In this request, the URI variable *{zbx-id}* is the object ID of the zBX node object.

### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property.  Value must be a valid virtual server <b>type</b> property value.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtual-servers	Array of objects	Array of virtual-server-info objects, described in the next table. Returned array may be empty.

Each nested virtual-server-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The canonical URI path of the Virtual Server object
name	String	The <b>name</b> property of the Virtual Server object
type	String Enum	The <b>type</b> property of the Virtual Server object
status	String Enum	The <b>status</b> property of the Virtual Server object

### Description

This operation lists the virtual servers that are managed by the identified zBX node. The **object-uri**, **name**, **type**, and **status** are provided for each.

If the object-id *{zbx-id}* does not identify a zBX object to which the API user has object-access permission, or if the zBX is not a member of an ensemble, a 404 status code is returned.

If the **name** query parameter is specified, the returned list is limited to those virtual servers that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid virtual server **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the

returned list is limited to those virtual servers that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not done.

The zBX specified by *{zbx-id}* must be a zBX node (that is, have a **type** of "**node**"), otherwise status code 400 (Bad Request) is returned.

A virtual server is included in the list only if the API user has object-access permission for that object. If an HMC is a manager of a virtual server but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

If the zBX node does not manage any virtual servers or if no virtual servers are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the zBX node whose object ID is *{zbx-id}*
- Object access permission to each Virtual Server object to be included in the result.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents" on page 287.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	105	A <b>type</b> query parameter defines an invalid value.
	112	The request URI designates a zBX object that is not a zBX node (that is, does not have a <b>type</b> of " <b>node</b> ").
404 (Not Found)	1	A zBX with object ID <i>{zbx-id}</i> does not exist on HMC or API user does not have object-access permission for it.
	100	The zBX with object ID <i>{zbx-id}</i> is not a member of an ensemble.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

### Example HTTP interaction

```
GET /api/zbx/8dac6a58-935e-352c-996e-ded17dbf92c0/virtual-servers HTTP/1.1
x-api-session: 579shzuc4ca3lwnly7xk2s64w6qmw2v2zsisikdrn4ky9sn9vv1
```

Figure 121. List Virtual Servers of a zBX (Node): Request



```

200 OK
server: zSeries management console API web server / 2.0
date: Mon, 09 Feb 2015 19:24:36 GMT
content-type: application/json;charset=UTF-8
content-length: 447
{
  "virtual-servers": [
    {
      "name": "B.2.03_SERVER_1",
      "object-uri": "/api/virtual-servers/bde7c552-a890-11e4-91e1-42f2e9106e93",
      "status": "operating",
      "type": "x-hyp"
    },
    {
      "name": "B.1.01_SERVER_2",
      "object-uri": "/api/virtual-servers/97cec09e-ae3c-11e4-a8fe-42f2e9106e93",
      "status": "operating",
      "type": "power-vm"
    }
  ]
}

```

Figure 122. List Virtual Servers of a zBX (Node): Response

## List Virtual Servers of a Node

The **List Virtual Servers of a Node** operation lists the virtual servers managed by the node with the given identifier.

### HTTP method and URI

**GET** /api/ensembles/{ensemble-id}/nodes/{node-id}/virtual-servers

#### URI variables:

Name	Description
{ensemble-id}	The object ID of the Ensemble object
{node-id}	The object ID of the Node object

#### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property. Value must be a valid virtual server <b>type</b> property value.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtual-servers	Array of objects	Array of nested virtual-server-info objects, described in the next table.

| Each nested virtual-server-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The canonical URI path of the Virtual Server object
name	String	The <b>name</b> property of the Virtual Server object
type	String Enum	The <b>type</b> property of the Virtual Server object
status	String Enum	The <b>status</b> property of the Virtual Server object

## | Description

| This operation lists the virtual servers that are managed by the identified Node. The **object-uri**, **name**, **type**, and **status** are provided for each.

| If the object ID *{node-id}* does not identify a Node object to which the API user has object-access permission or if the node is not a member of an ensemble, a 404 status code is returned.

| If the **name** query parameter is specified, the returned list is limited to those virtual servers that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

| If the **type** query parameter is specified, the parameter is validated to ensure it is a valid virtual server **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those virtual servers that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not done.

| A virtual server is included in the list only if the API user has object-access permission for that object. If an HMC is a manager of a virtual server but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

| If the Node does not manage any virtual servers or if no virtual servers are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.

## | Authorization requirements

| This operation has the following authorization requirements:

- | • Object access permission to the node whose object ID is *{node-id}*
- | • Object access permission to each Virtual Server object to be included in the result.

## | HTTP status and reason codes

| On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 289.

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	105	A <b>type</b> query parameter defines an invalid value.
404 (Not Found)	1	A node with object ID <i>{node-id}</i> does not exist on HMC or API user does not have object-access permission for it.
	100	The node with object ID <i>{node-id}</i> is not a member of an ensemble.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
GET /api/ensembles/de7915f4-a127-11e2-8273-5cf3fcae8019/nodes/
    8dac6a58-935e-352c-996e-ded17dbf92c0/virtual-servers HTTP/1.1
x-api-session: 579shzuc4ca3lwnly7xk2s64w6qmw2v2zsisikdrn4ky9sn9vv1
```

Figure 123. List Virtual Servers of a Node: Request

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Mon, 09 Feb 2015 19:24:33 GMT
content-type: application/json;charset=UTF-8
content-length: 447
{
  "virtual-servers": [
    {
      "name": "B.2.03_SERVER_1",
      "object-uri": "/api/virtual-servers/bde7c552-a890-11e4-91e1-42f2e9106e93",
      "status": "operating",
      "type": "x-hyp"
    },
    {
      "name": "B.1.01_SERVER_2",
      "object-uri": "/api/virtual-servers/97cec09e-ae3c-11e4-a8fe-42f2e9106e93",
      "status": "operating",
      "type": "power-vm"
    }
  ]
}
```

Figure 124. List Virtual Servers of a Node: Response

## List Virtual Servers of an Ensemble

The **List Virtual Servers of an Ensemble** operation lists the virtual servers managed by the ensemble with the given identifier.

### HTTP method and URI

```
GET /api/ensembles/{ensemble-id}/virtual-servers
```

In this request, the URI variable *{ensemble-id}* is the object ID of the Ensemble object.

**Query parameters:**

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property.  Value must be a valid virtual server <b>type</b> property value.

**Response body contents**

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtual-servers	Array of objects	Array of virtual-server-info objects, described in the next table. Returned array may be empty.

Each nested virtual-server-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The canonical URI path of the Virtual Server object
name	String	The <b>name</b> property of the Virtual Server object
type	String Enum	The <b>type</b> property of the Virtual Server object
status	String Enum	The <b>status</b> property of the Virtual Server object

**Description**

This operation lists the virtual servers that are managed by the identified ensemble. The **object-uri**, **object-id**, **name**, **type**, and **status** are provided for each.

If the object-id *{ensemble-id}* does not identify an ensemble object to which the API user has object-access permission, a 404 status code is returned.

If the **name** query parameter is specified, the returned list is limited to those virtual servers that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid virtual server **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those virtual servers that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not done.

A virtual server is included in the list only if the API user has object-access permission for that object. If an HMC is a manager of a virtual server but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

If the ensemble does not manage any virtual servers or if no virtual servers are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the ensemble whose object-id is *{ensemble-id}*
- Object access permission to each Virtual Server object to be included in the result.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 292.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	105	A <b>type</b> query parameter defines an invalid value.
404 (Not Found)	1	An ensemble with object-id <i>{ensemble-id}</i> does not exist on HMC or API user does not have object-access permission for it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334/virtual-servers HTTP/1.1
x-api-session: 64zdknfzh24fqw8f9v5g99h96rhrxhodmawtknx4iutmqrfrvf
```

---

Figure 125. List Virtual Servers of an Ensemble: Request

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Sat, 12 Nov 2011 21:09:23 GMT
content-type: application/json;charset=UTF-8
content-length: 1008
{
  "virtual-servers": [
    {
      "name": "APIVM2",
      "object-uri": "/api/virtual-servers/63911688-03f4-11e1-881f-001f163805d8",
      "status": "operating",
      "type": "prsm"
    },
    {
      "name": "SS-Web-Svr-1",
      "object-uri": "/api/virtual-servers/7ba96572-0d72-11e1-892f-f0def14b63af",
      "status": "not-operating",
      "type": "x-hyp"
    },
    {
      "name": "SSWSVR3",
      "object-uri": "/api/virtual-servers/86819708-0d72-11e1-bf89-f0def14b63af",
      "status": "not-activated",
      "type": "zvm"
    },
    {
      "name": "ZOS",
      "object-uri": "/api/virtual-servers/636768f6-03f4-11e1-881f-001f163805d8",
      "status": "operating",
      "type": "prsm"
    },
    {
      "name": "SS-web-Srv-2",
      "object-uri": "/api/virtual-servers/960529e2-0d3b-11e1-9f64-f0def14b63af",
      "status": "operating",
      "type": "power-vm"
    }
  ]
}

```

Figure 126. List Virtual Servers of an Ensemble: Response

## List Virtual Servers of a CPC

The **List Virtual Servers of a CPC** operation lists the virtual servers managed by the CPC with the given identifier.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/virtual-servers

In this request, the URI variable {cpc-id} is the object ID of the CPC object.

### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.

Name	Type	Rqd/Opt	Description
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property.  Value must be a valid virtual server <b>type</b> property value.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtual-servers	Array of objects	Array of nested virtual-server-info objects, described in the next table.

Each nested virtual-server-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The canonical URI path of the Virtual Server object
name	String	The <b>name</b> property of the Virtual Server object
type	String Enum	The <b>type</b> property of the Virtual Server object
status	String Enum	The <b>status</b> property of the Virtual Server object

## Description

This operation lists the virtual servers that are managed by the identified CPC. The **object-uri**, **name**, **type**, and **status** are provided for each.

If the object-id *{cpc-id}* does not identify a CPC object to which the API user has object-access permission or if the CPC is not a member of an ensemble, a 404 status code is returned.

If the **name** query parameter is specified, the returned list is limited to those virtual servers that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid virtual server **type** property value. If the value is not valid, a 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those virtual servers that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not done.

A virtual server is included in the list only if the API user has object-access permission for that object. If an HMC is a manager of a virtual server but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

If the CPC does not manage any virtual servers or if no virtual servers are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the cpc whose object ID is *{cpc-id}*
- Object access permission to each Virtual Server object to be included in the result.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 295.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	105	A <b>type</b> query parameter defines an invalid value.
404 (Not Found)	1	A CPC with object-id <i>{cpc-id}</i> does not exist on the HMC or API user does not have object-access permission for it.
	100	The CPC with object-id <i>{cpc-id}</i> is not a member of an ensemble.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/virtual-servers HTTP/1.1
x-api-session: 64zdknfzh24fqw8f9v5g99h96rhrxhodmawtknx4iutmgrfvf
```

---

Figure 127. List Virtual Servers of a CPC: Request



```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Sat, 12 Nov 2011 21:09:23 GMT
content-type: application/json;charset=UTF-8
content-length: 423
{
  "virtual-servers": [
    {
      "name": "ZOS",
      "object-uri": "/api/virtual-servers/636768f6-03f4-11e1-881f-001f163805d8",
      "status": "operating",
      "type": "prsm"
    },
    {
      "name": "SS-web-Srv-2",
      "object-uri": "/api/virtual-servers/960529e2-0d3b-11e1-9f64-f0def14b63af",
      "status": "operating",
      "type": "power-vm"
    }
  ]
}

```

Figure 128. List Virtual Servers of a CPC: Response

## List Virtual Servers of a Virtualization Host

The **List Virtual Servers of a Virtualization Host** operation lists the virtual servers managed by the virtualization host with the given identifier.

### HTTP method and URI

**GET** `/api/virtualization-hosts/{virt-host-id}/virtual-servers`

In this request, the URI variable `{virt-host-id}` is the object ID of the Virtualization Host object.

### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtual-servers	Array of objects	Array of nested virtual-server-info objects, described in the next table.

Each nested virtual-server-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	The <b>object-uri</b> property of the Virtual Server object.

Field name	Type	Description
name	String	The <b>name</b> property of the Virtual Server object
type	String Enum	The <b>type</b> property of the Virtual Server object
status	String Enum	The <b>status</b> property of the Virtual Server object

## Description

This operation lists the virtual servers that are managed by the identified virtualization host. The **object-uri**, **name**, **type**, and **status** are provided for each.

If the object-id *{virt-host-id}* does not identify a Virtualization Host object for which the API user has object-access permission to its hosting-environment, a 404 status code is returned.

If the **name** query parameter is specified, the returned list is limited to those virtual servers that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

A virtual server is included in the list only if the API user has object-access permission for that object. If an HMC is a manager of a virtual server but the API user does not have permission to it, that object is simply omitted from the list but no error status code results.

If the virtualization host does not manage any virtual servers or if no virtual servers are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to hosting-environment of the virtualization host with object-id *{virt-host-id}*
- Object access permission to each Virtual Server object to be included in the result.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 297.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	A Virtualization Host with object-id <i>{virt-host-id}</i> does not exist on HMC or API user does not have object-access permission for its hosting-environment.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/71822c16-0401-11e1-8eda-001f163805d8/virtual-servers HTTP/1.1
x-api-session: 64zdknfzh24fqw8f9v5g99h96hrxhodmawtknx4iutmqrfrvf
```

---

Figure 129. List Virtual Servers of a Virtualization Host: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Sat, 12 Nov 2011 21:08:38 GMT
content-type: application/json;charset=UTF-8
content-length: 434
{
  "virtual-servers": [
    {
      "name": "SS-Web-Svr-1",
      "object-uri": "/api/virtual-servers/7ba96572-0d72-11e1-892f-f0def14b63af",
      "status": "operating",
      "type": "x-hyp"
    },
    {
      "name": "SS-FTP-Svr-1",
      "object-uri": "/api/virtual-servers/5483e6b2-09fc-11e1-9f08-001f163805d8",
      "status": "not-operating",
      "type": "x-hyp"
    }
  ]
}
```

---

Figure 130. List Virtual Servers of a Virtualization Host: Response

## Create Virtual Server

The **Create Virtual Server** operation creates a **"power-vm"**, **"x-hyp"**, or **"zvm"** virtual server with the given properties on the identified Virtualization Host. This operation is not supported for **"prsm"** Virtualization Hosts/virtual servers.

### HTTP method and URI

```
POST /api/virtualization-hosts/{virt-host-id}/virtual-servers
```

In this request, the URI variable *{virt-host-id}* is the object ID of the Virtualization Host object.

### Request body contents

Properties are only valid if they are supported for a virtual server whose **type** property matches the given **type** property value. For instance a virtual server with type **"power-vm"** will define a **processing-mode** property (PowerVM only) and **mac-prefix** property (PowerVM, z/VM) but not an **initial-share-mode** property (z/VM only).

Properties may also only be valid if prerequisite properties have specific values. For example, a PowerVM virtual server's **initial-virtual-processors** is only valid when the **processing-mode** value is **"shared"**.

See the "Data Model" on page 261 for a complete definition of all properties including for what virtual server types they are valid and their prerequisites.

Field name	Type	Rqd/Opt	Description
type	String Enum	Required	The value to be set as the virtual server's <b>type</b> property. The value may not be " <b>prsm</b> ".
name	String	Required	The value to be set as the virtual server's <b>name</b> property.
description	String	Optional	The value to be set as the virtual server's <b>description</b> property. Default value: an empty string
cpu-perf-mgmt-enabled	Boolean	Optional	The value to be set as the virtual server's <b>cpu-perf-mgmt-enabled</b> property. Default value: true. Valid only when the ensemble's <b>management-enablement-level</b> is " <b>automate</b> ".
processing-mode	String Enum	Optional	The value to be set as the virtual server's <b>processing-mode</b> property. Default value: shared
minimum-virtual-processors	Integer	Optional	The value to be set as the virtual server's <b>minimum-virtual-processors</b> property. Default value: 1
initial-virtual-processors	Integer	Optional	The value to be set as the virtual server's <b>initial-virtual-processors</b> property. Default value: 1
maximum-virtual-processors	Integer	Optional	The value to be set as the virtual server's <b>maximum-virtual-processors</b> property. Default value based on <b>type</b> : <ul style="list-style-type: none"> <li>power-vm: max (<b>maximum-allowed-dedicated-processors</b>, <b>initial-virtual-processors</b>)</li> <li>zvm: 1</li> </ul>
minimum-dedicated-processors	Integer	Optional	The value to be set as the virtual server's <b>minimum-dedicated-processors</b> property. Default value: 1
initial-dedicated-processors	Integer	Optional	The value to be set as the virtual server's <b>initial-dedicated-processors</b> property. Default value: 1
maximum-dedicated-processors	Integer	Optional	The value to be set as the virtual server's <b>maximum-dedicated-processors</b> property. Default value: <b>maximum-allowed-dedicated-processors</b>
minimum-processing-units	Float	Optional	The value to be set as the virtual server's <b>minimum-processing-units</b> property. Default value: 0.1
initial-processing-units	Float	Optional	The value to be set as the virtual server's <b>initial-processing-units</b> property. Default value: 0.1
maximum-processing-units	Float	Optional	The value to be set as the virtual server's <b>maximum-processing-units</b> property. Default value: min ( <b>maximum-virtual-processors</b> , <b>maximum-allowed-processing-units</b> )
initial-share-mode	String Enum	Optional	The value to be set as the virtual server's <b>initial-share-mode</b> property. Default value: " <b>relative</b> "
initial-shares	Float	Optional	The value to be set as the virtual server's <b>initial-shares</b> property.  For x Hyp virtual servers: <ul style="list-style-type: none"> <li>Default: 1024</li> <li>Prerequisite: The parent virtualization host's <b>cpu-shares-supported</b> value is <b>true</b>.</li> </ul> For z/VM servers: <ul style="list-style-type: none"> <li>Default: Based on initial-share-mode: <ul style="list-style-type: none"> <li>relative: 100</li> <li>absolute: 10.0</li> </ul> </li> </ul>
share-limit	String Enum	Optional	The value to be set as the virtual server's <b>share-limit</b> property. Default value: " <b>none</b> "
maximum-share-mode	String Enum	Optional	The value to be set as the virtual server's <b>maximum-share-mode</b> property. Default value: " <b>relative</b> "

Field name	Type	Rqd/Opt	Description
minimum-shares	Float	Optional	The value to be set as the virtual server's <b>minimum-shares</b> property. Default value: 2
maximum-shares	Float	Optional	The value to be set as the virtual server's <b>maximum-shares</b> property. Default value for x Hyp virtual servers: 262144. Default value for z/VM virtual servers is based on <b>maximum-share-mode</b> : <ul style="list-style-type: none"> <li>• relative: 100</li> <li>• absolute: 10.0</li> </ul>
minimum-memory	Integer	Optional	The value to be set as the virtual server's <b>minimum-memory</b> property. Default value: max (256, <b>initial-memory</b> )
initial-memory	Integer	Optional	The value to be set as the virtual server's <b>initial-memory</b> property. Default value: 1024
maximum-memory	Integer	Optional	The value to be set as the virtual server's <b>maximum-memory</b> property. Default value: <b>initial-memory</b>
boot-mode	String Enum	Optional	The value to be set as the virtual server's <b>boot-mode</b> property. Default: normal
boot-network-adapter-client-ip	String	Optional	The value to be set as the virtual server's <b>boot-network-adapter-client-ip</b> property. Default value: an empty string
boot-network-adapter-subnet-ip	String	Optional	The value to be set as the virtual server's <b>boot-network-adapter-subnet-ip</b> property. Default value: an empty string
boot-network-adapter-gateway-ip	String	Optional	The value to be set as the virtual server's <b>boot-network-adapter-gateway-ip</b> property. Default value: an empty string
boot-network-adapter-server-ip	String	Optional	The value to be set as the virtual server's <b>boot-network-adapter-server-ip</b> property. Default value: an empty string
keylock	String Enum	Optional	The value to be set as the virtual server's <b>keylock</b> property. Default value: " <b>normal</b> "
auto-start	Boolean	Optional	The value to be set as the virtual server's <b>auto-start</b> property. Default value: false. Prerequisite: <b>type</b> is " <b>power-vm</b> " or " <b>x-hyp</b> ".
dlpar-enabled	Boolean	Optional	The value to be set as the virtual server's <b>dlpar-enabled</b> property. Default value: false
gpmp-support-enabled	Boolean	Optional	The value to be set as the virtual server's <b>gpmp-support-enabled</b> property. Default value: false. Prerequisite: <b>type</b> is " <b>power-vm</b> ", " <b>zvm</b> " or " <b>x-hyp</b> ".
rmc-virtual-device-address	String	Optional	The virtual device address of the resource monitoring and control network to create in order to support the Guest Platform Management Provider (GPMP).  Prerequisite: <b>type</b> is " <b>zvm</b> " and <b>gpmp-support-enabled</b> is true.  Required if <b>gpmp-support-enabled</b> is true.  1-4 character hex string
password	String	Required if <b>type</b> is " <b>zvm</b> "	The value to be set as the virtual server's <b>password</b> property. Prerequisite: <b>type</b> is " <b>zvm</b> "
privilege-classes	String	Optional	The value to be set as the virtual server's <b>privilege-classes</b> property. Default value "G"
ipl-device	String	Optional	The value to be set as the virtual server's <b>ipl-device</b> property. Default value "CMS"

Field name	Type	Rqd/Opt	Description
ipl-load-parameters	String	Optional	The value to be set as the virtual server's <b>ipl-load-parameters</b> property. Default value: an empty string
ipl-parameters	String	Optional	The value to be set as the virtual server's <b>ipl-parameters</b> property. Default value "AUTOOCR".
inband-monitoring-enabled	Boolean	Optional	The value to be set as the virtual server's <b>inband-monitoring-enabled</b> property. Default value: false
keyboard-language	String	Optional	The value to be set as the virtual server's <b>keyboard-language</b> property. Default value: the first value in the virtualization host's <b>supported-keyboard-languages</b> array or null if that array is null or empty.
shutdown-timeout-source	String Enum	Optional	The value to be set as the virtual server's <b>shutdown-timeout-source</b> property. Default value: "virtualization-host".  Prerequisite: <b>type</b> is "power-vm" or "x-hyp"
shutdown-timeout	Integer	Optional	The value to be set as the virtual server's <b>shutdown-timeout</b> property. Default value: 300  Prerequisite: <b>type</b> is "power-vm" or "x-hyp"

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
object-uri	String/ URI	The <b>object-uri</b> property of the created Virtual Server object.

## Description

This operation creates a virtual server with the values specified on the identified Virtualization Host and then returns its object-uri in the response body. The response also includes a **Location** header that provides this URI.

Any properties identified as "Required" must be included in the request body. Any properties identified as "Optional" may be excluded from the request body; if an optional property is not found in the request body, its value will be set to its default value.

If the API user does not have authority to perform the Create Virtual Server task or the zManager management enablement level for the ensemble does not allow the setting of a provided property, a 403 (Forbidden) status code is returned. A 404 (Not Found) status code is returned if the object-id *{virt-host-id}* does not identify a Virtualization Host object for which the API user has object-access permission to its hosting-environment. If the Virtualization Host is of **type "zvm"** and its status is not-activated, a 409 (Conflict) status code is returned.

If the Request Body Contents fail to validate, a 400 (Bad Request) status code is returned. This may occur because the document fails to define a required property or defines a property that is not supported for the given virtual server type. This may also occur if the document fails to define a single valid virtual server, for instance defining a property with an invalid value (e.g. an initial-memory value less than zero or a virtual server **type** of "zvm" when the Virtualization Host **type** is "power-vm"). A 400 (Bad Request) will also be returned if the request body contains a property that is not valid given the value of a prerequisite property (e.g. defining a value for **initial-dedicated-memory** when the value of **processing-mode** is "shared").

If the Request Body Contents are valid, the virtual server is created on the target Virtualization Host and its properties are defined to their corresponding request body content's properties' values. If a property is omitted from the request body, its default value is used when creating the virtual server. The newly created virtual server will have empty list values defined for its **storage-adapters**. For PowerVM and x Hyp virtual servers, the **network-adapters** property will contain the definition of the network-adapter for the Default network and it will have a **boot-sequence** array that contains only "**network-adapter**". For z/VM virtual servers, the **network-adapters** property value will be an empty list. The virtual server will also be a member of the default workload resource group.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Virtualization Host's hosting-environment
- Action/task permission to **Create Virtual Server** task.
- Action/task permission to the **Virtual Server Performance Details** task if the **gmp-support-enabled** or **cpu-perf-mgmt-enabled** field is specified in the request body.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in "Response body contents" on page 302.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	101	A virtual server with the <b>name</b> specified in the request body already exists on the Virtualization Host with object-id <i>{virt-host-id}</i> .
	102	The object-id <i>{virt-host-id}</i> identifies a Virtualization Host whose <b>type</b> property value differs from the <b>type</b> defined in the request body.
	103	A "shares" ( <b>initial-shares</b> or <b>maximum-shares</b> ) property value type is not correct based on the value of its corresponding "share-mode" ( <b>initial-share-mode</b> or <b>maximum-share-mode</b> ). For example, the <b>initial-share-mode</b> is defined as "relative" but the <b>initial-shares</b> value is not an Integer.
403 (Forbidden)	1	API user does not have action permission to the <b>Create Virtual Server</b> task.
	100	The ensemble's <b>management-enablement-level</b> property value does not allow the updating of a provided property.
404 (Not Found)	1	A Virtualization Host with object-id <i>{virt-host-id}</i> does not exist on HMC or API user does not have object-access permission for its hosting-environment.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	101	Parent Virtualization Host has a <b>status</b> value that is not valid to perform the operation (attempted to create a virtual server of type " <b>zvm</b> " and Virtualization Host is not active).

HTTP error status code	Reason code	Description
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
POST /api/virtualization-hosts/2f7bf364-03f8-11e1-8eda-001f163805d8/virtual-servers HTTP/1.1
x-api-session: 64zdknfzh24fqw8f9v5g99h96hrxhodmawtknx4iutmgrfvf
content-type: application/json
content-length: 161
{
  "description": "Spacely Sprockets Web Store Web Server #2",
  "initial-memory": 2048,
  "initial-virtual-processors": 2,
  "name": "SS-Web-Svr-2",
  "type": "power-vm"
}
```

---

Figure 131. Create Virtual Server: Request for a virtual server of type "power-vm"

---

```
201 Created
server: zSeries management console API web server / 1.0
location: /api/virtual-servers/7d298eb8-0d72-11e1-8c83-f0def14b63af
cache-control: no-cache
date: Sat, 12 Nov 2011 21:08:41 GMT
content-type: application/json;charset=UTF-8
content-length: 74
{
  "object-uri": "/api/virtual-servers/7d298eb8-0d72-11e1-8c83-f0def14b63af"
}
```

---

Figure 132. Create Virtual Server: Response for a virtual server of type "power-vm"

## Delete Virtual Server

The **Delete Virtual Server** operation deletes the identified virtual server. This operation is not supported for PR/SM virtual servers.

### HTTP method and URI

**DELETE** /api/virtual-servers/{*virtual-server-id*}

In this request, the URI variable {*virtual-server-id*} is the object ID of the Virtual Server object.



## Description

A 409 (Conflict) status code is returned if the virtual server has a **status** other than **"not-operating"** or **"not-activated"**. A 409 (Conflict) status code is also returned if the virtual server is busy for the duration of the request.

This operation deletes the identified virtual server, which includes the following:

- The virtual server's network adapters are disconnected and deleted.
- The virtual server's virtual disks are removed from the virtual server.
- Virtual servers of type **"zvm"** are also removed from the z/VM Directory.
- The virtual server is removed from the Virtualization Host.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server
- Object access permission to the virtual server's Virtualization Host's hosting-environment.
- Action/task permission to **Delete Virtual Server** task

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
403 (Forbidden)	1	API user does not have action permission to the <b>Delete Virtual Server</b> task.
404 (Not Found)	1	A virtual server with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it or its virtualization host's hosting-environment.
409 (Conflict)	1	Virtual server status is not valid to perform the operation (must be either <b>"not-operating"</b> or <b>"not-activated"</b> ).
	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Example HTTP interaction

---

```
DELETE /api/virtual-servers/7ba96572-0d72-11e1-892f-f0def14b63af HTTP/1.1
x-api-session: 64zdknfzh24fqw8f9v5g99h96hrxhodmawtknx4iutmqrfrvf
```

---

Figure 133. Delete Virtual Server: Request

---

```
204 No Content
date: Sat, 12 Nov 2011 21:09:24 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

<No response body>

---

Figure 134. Delete Virtual Server: Response

## Get Virtual Server Properties

The **Get Virtual Server Properties** operation retrieves the properties of a single Virtual Server object that is designated by its object-id.

### HTTP method and URI

```
GET /api/virtual-servers/{virtual-server-id}
```

In this request, the URI variable *{virtual-server-id}* is the object ID of the Virtual Server object.

### Query parameters

Name	Type	Req/Opt	Description
properties	String	Optional	Identifies the properties to be returned for the Virtual Server object. The only supported value is " <b>common</b> ", which indicates that only a specific subset of properties that are quickly available and common to all types of virtual server should be returned. This properties included in this subset is specified in "Response body contents." If this query parameter is omitted, then all properties for the virtual server as defined in "Data Model" on page 261 are returned.

## Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Virtual Server object.

If the **properties** query parameter is not specified, the response body provides all of the properties for the virtual server as defined in the Data Model section above. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

If the **properties=common** query parameter is specified, the response body provides the values for only the following properties: **acceptable-status**, **auto-start**, **class**, **description**, **has-unacceptable-status**, **hostname**, **is-locked**, **name**, **object-id**, **object-uri**, **os-level**, **os-name**, **os-type**, **parent**, **status**, **type**, and **workloads**. Field names and data types in the JSON object are the same as the property names and data types defined in the data model. These properties are common to all types of virtual servers, and can be provided more quickly than the entire set of properties for the virtual server.

## Description

Returns the current values of the properties for the Virtual Server object as defined in the “Data Model” on page 261.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the Virtual Server object designated by *{virtual-server-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 306.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	A Virtualization Host with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMIPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMIPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

Retrieval of the full property set for a virtual server will generally require an interaction with the SE, and in some cases (for example, for z/VM virtual servers) may also require an interaction with the virtualization host. Some API applications may choose to represent virtual servers in a generic, rather than highly type-specialized, way and thus not need the full set of properties for each virtual server. The **properties=common** query parameter is provided to allow for improved performance for such applications by avoiding the processing that would be incurred in retrieving unneeded properties. The properties specified by **properties=common** are available for all virtual server types, and are generally cached on the HMC so that they can be retrieved quickly and without requiring an interaction with the SE or virtualization host.

## Example HTTP interaction

---

```
GET /api/virtual-servers/7ba96572-0d72-11e1-892f-f0def14b63af HTTP/1.1  
x-api-session: 64zdknfzh24fqw8f9v5g99h96rhrxhodmawtknx4iutmgqrfvf
```

---

*Figure 135. Get Virtual Server Properties: Request*

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 15 Nov 2011 17:11:03 GMT
content-type: application/json;charset=UTF-8
content-length: 3175
{
  "acceptable-status": [
    "operating"
  ],
  "auto-start": false,
  "boot-mode": "normal",
  "boot-network-adapter-client-ip": "192.168.1.22",
  "boot-network-adapter-gateway-ip": "192.168.1.1",
  "boot-network-adapter-server-ip": "192.168.1.254",
  "boot-network-adapter-subnet-ip": "255.255.255.0",
  "boot-sequence": [
    "network-adapter"
  ],
  "class": "virtual-server",
  "cpu-perf-mgmt-enabled": true,
  "description": "Spacely Sprockets Web Store Web Server #2",
  "dlpar-active": false,
  "dlpar-enabled": false,
  "gpmp-network-adapter":
|
| {
|   "adapter-type": "gpmp"
|   "element-id": "a3cbdcea-0fac-11e1-902b-f0def14b63af",
|   "element-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/
|     network-adapters/a3cbbdcea-0fac-11e1-902b-f0def14b63af",
|   "mac-address": "02:ff:a6:b3:3c:c6",
|   "network-uri": null
| }
  "gpmp-status": "not-operating",
  "gpmp-support-enabled": false,
  "gpmp-version": "unavailable",
  "has-unacceptable-status": true,
  "hostname": null,
  "inband-monitoring-enabled": false,
  "initial-dedicated-processors": null,
  "initial-memory": 2048,
  "initial-processing-units": 0.20000000000000001,
  "initial-virtual-processors": 2,
  "is-locked": false,
  "keylock": "normal",
  "mac-prefix": {
    "mac-address": "02:d0:19:aa:76:00",
    "prefix-length": 40
  },
  "maximum-dedicated-processors": null,
  "maximum-memory": 2048,
  "maximum-processing-units": 7.0,
  "maximum-virtual-processors": 7,
  "minimum-dedicated-processors": null,
  "minimum-memory": 2048,
  "minimum-processing-units": 0.10000000000000001,
  "minimum-virtual-processors": 1,
  "mounted-media-name": "gpa.iso",
  "name": "SS-Web-Svr-2",

```

Figure 136. Get Virtual Server Properties: Response for virtual servers of "power-vm" (Part 1)

```

"network-adapters": [
  {
    "adapter-type": "regular"
    "element-id": "596dd87c-0db7-11e1-9251-f0def14b63af",
    "element-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/
network-adapters/596dd87c-0db7-11e1-9251-f0def14b63af",
    "mac-address": null,
    "network-uri": "/api/virtual-networks/f920171e-03f2-11e1-8e8e-0010184c8334"
  },
  {
    "adapter-type": "regular"
    "element-id": "c4bbdcea-0fac-11e1-903e-f0def14b63af",
    "element-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/
network-adapters/c4bbdcea-0fac-11e1-903e-f0def14b63af",
    "mac-address": null,
    "network-uri": "/api/virtual-networks/f3aece54-0db8-11e1-9e2c-00215e69dea0"
  }
],
"object-id": "588d8c18-0db7-11e1-b1f1-f0def14b63af",
"object-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af",
"os-level": null,
"os-name": null,
"os-type": null,
"parent": "/api/virtualization-hosts/2f7bf364-03f8-11e1-8eda-001f163805d8",
"processing-mode": "shared",
"status": "not-operating",
"type": "power-vm",
"virtual-disks": [
  {
    "backing-virtualization-host-storage-resource": "/api/virtualization-hosts/
2f7bf364-03f8-11e1-8eda-001f163805d8/virtualization-host-storage-resources/
37699380-0fa9-11e1-b69e-f0def14b63af",
    "description": "Boot filesystem",
    "element-id": "c547fb4e-0fac-11e1-842b-f0def14b63af",
    "element-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/
virtual-disks/c547fb4e-0fac-11e1-842b-f0def14b63af",
    "name": "boot",
    "owner": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af",
    "size": 17179869184,
    "type": "fullpack"
  },
  {
    "backing-virtualization-host-storage-resource": "/api/virtualization-hosts/
2f7bf364-03f8-11e1-8eda-001f163805d8/virtualization-host-storage-resources/
b0e6c160-0fa9-11e1-b69e-f0def14b63af",
    "description": "Physical vol for sysvg",
    "element-id": "c568a13c-0fac-11e1-903e-f0def14b63af",
    "element-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/
virtual-disks/c568a13c-0fac-11e1-903e-f0def14b63af",
    "name": "pv01a",
    "owner": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af",
    "size": 8589934592,
    "type": "fullpack"
  }
],
"workloads": [
  "/api/workload-resource-groups/f9fbe5fa-03f2-11e1-8e8e-0010184c8334"
]
}

```

Figure 137. Get Virtual Server Properties: Response for virtual servers of "power-vm" (part 2)

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 15 Nov 2011 18:23:31 GMT
content-type: application/json;charset=UTF-8
content-length: 1121
{
  "acceptable-status": [
    "operating"
  ],
  "associated-logical-partition": "/api/logical-partitions/8394537b-1cc3-3607-88ad-06845b263439",
  "class": "virtual-server",
  "cpu-perf-mgmt-enabled": false,
  "description": "",
  "gpmp-status": "not-operating",
  "gpmp-support-enabled": false,
  "gpmp-version": "unavailable",
  "has-unacceptable-status": true,
  "is-locked": false,
  "name": "LP03",
  "network-adapters": [
    {
      "chpid": "42",
      "css": "0",
      "element-id": "OSX 0.42",
      "element-uri": "/api/virtual-servers/916f73c8-de39-11e0-a58e-f0def161133a/network-adapters/OSX%200.42",
      "network-uris": null,
      "type": "osx"
    },
    {
      "chpid": "43",
      "css": "0",
      "element-id": "OSX 0.43",
      "element-uri": "/api/virtual-servers/916f73c8-de39-11e0-a58e-f0def161133a/network-adapters/OSX%200.43",
      "network-uris": null,
      "type": "osx"
    }
  ],
  "object-id": "916f73c8-de39-11e0-a58e-f0def161133a",
  "object-uri": "/api/virtual-servers/916f73c8-de39-11e0-a58e-f0def161133a",
  "os-level": "6.2.0 - 1101",
  "os-name": "LP03",
  "os-type": "z/VM",
  "parent": "/api/virtualization-hosts/8e9cad8c-de39-11e0-a58e-f0def161133a",
  "status": "operating",
  "type": "prsm",
  "workloads": [
    "/api/workload-resource-groups/1fe39a72-de39-11e0-90a6-00215e67351a"
  ]
}

```

Figure 138. Get Virtual Server Properties: Response for virtual servers of type "prsm"

---

```

| 200 OK
| server: zSeries management console API web server / 1.0
| cache-control: no-cache
| date: Tue, 15 Nov 2014 16:56:19 GMT
| 'content-type': 'application/json'
| {
|   "os-name":null,
|   "shutdown-timeout-source":"virtualization-host",
|   "shutdown-timeout":300,
|   "boot-sequence":[
|     "virtual-media",
|     "virtual-disk",
|     "network-adapter"
|   ],
|   "gmp-status":"not-operating",
|   "avail-policies":[
|     {
|       "avail-status-impact-exclusion":false,
|       "policy-uri":"/api/workload-resource-groups/a40e2b96-9adb-11e1-8868-00215e67ad51/availability-policies/d0b6fecc-191e-11e2-8561-00215e67ad51",
|       "workload-uri":"/api/workload-resource-groups/a40e2b96-9adb-11e1-8868-00215e67ad51"
|     }
|   ],
|   "hostname":null,
|   "keyboard-language":"en_US",
|   "auto-start":false,
|   "os-type":null,
|   "avail-status":[
|     {
|       "reasons":[
|         "vs"
|       ],
|       "avail-status":"not-available"
|     }
|   ],
|   "gmp-network-adapter":
|     {
|       "adapter-type": "gmp"
|       "element-id": "a3cbdcea-0fac-11e1-902b-f0def14b63af",
|       "element-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/network-adapters/a3cbbdcea-0fac-11e1-902b-f0def14b63af",
|       "emulation-mode":"e1000"
|       "mac-address": "02:ff:a6:b3:3c:c6",
|       "network-uri": null
|     }
|   "gmp-version":"unavailable",
|   "is-locked":false,
|   "type":"x-hyp",
|   "status":"not-operating",
|   "maximum-shares":262144,
|   "mounted-media-name":null,
|   "description":"",
|   "parent":"/api/virtualization-hosts/6c54bdcc-0b3d-11e3-b37e-f0def1b30158",
|   "os-level":null,
|   "object-id":"f6755c24-a457-11e3-8a23-f0def1b30158",
|   "network-adapters":[
|     {

```

---

Figure 139. Get Virtual Server Properties: Response for virtual servers of type "x-hyp" (Part 1)



```

| "adapter-type": "regular"
  "network-uri": "/api/virtual-networks/a1b736ee-9adb-11e1-8868-00215e67ad51",
  "emulation-mode": "e1000",
  "element-id": "f6a47892-a457-11e3-8195-f0def1b30158",
  "mac-address": "02:ff:42:41:69:93",
  "element-uri": "/api/virtual-servers/f6755c24-a457-11e3-8a23-f0def1b30158/network-adapters/
    f6a47892-a457-11e3-8195-f0def1b30158"
  },
  ],
  "minimum-shares": 2,
  "virtual-disks": [
    {
      "description": "",
      "element-uri": "/api/virtual-servers/f6755c24-a457-11e3-8a23-f0def1b30158/virtual-disks/
        f6bdf2ae-a457-11e3-a23c-f0def1b30158",
      "backing-virtualization-host-storage-resource": "/api/virtualization-hosts/
        6c54bdcc-0b3d-11e3-b37e-f0def1b30158/virtualization-host-storage-resources/
        9e7d4072-a457-11e3-8195-f0def1b30158",
      "element-id": "f6bdf2ae-a457-11e3-a23c-f0def1b30158",
      "emulation-mode": "ide",
      "owner": "/api/virtual-servers/f6755c24-a457-11e3-8a23-f0def1b30158",
      "size": 1073741824,
      "type": "fullpack",
      "name": "testvdisk1"
    },
    {
      "description": "",
      "element-uri": "/api/virtual-servers/f6755c24-a457-11e3-8a23-f0def1b30158/virtual-disks/
        f6c29840-a457-11e3-a23c-f0def1b30158",
      "backing-virtualization-host-storage-resource": "/api/virtualization-hosts/
        6c54bdcc-0b3d-11e3-b37e-f0def1b30158/virtualization-host-storage-resources/
        be507bb2-a457-11e3-8a23-f0def1b30158",
      "element-id": "f6c29840-a457-11e3-a23c-f0def1b30158",
      "emulation-mode": "virtio-scsi",
      "owner": "/api/virtual-servers/f6755c24-a457-11e3-8a23-f0def1b30158",
      "size": 1073741824,
      "type": "fullpack",
      "name": "testvdisk_2"
    }
  ],
  "class": "virtual-server",
  "workload-element-groups": [],
  "name": "testbootsequence_1",
  "initial-memory": 1024,
  "perf-policies": [],
  "cpu-perf-mgmt-enabled": true,
  "acceptable-avail-status": [
    "available"
  ],
  "acceptable-status": [
    "operating"
  ],
  "initial-shares": 1024,
  "gmp-support-enabled": false,
  "initial-virtual-processors": 1,
  "object-uri": "/api/virtual-servers/f6755c24-a457-11e3-8a23-f0def1b30158",
  "workloads": [
    "/api/workload-resource-groups/a40e2b96-9adb-11e1-8868-00215e67ad51"
  ],
  "has-unacceptable-status": true,
  "inband-monitoring-enabled": false
}

```

| Figure 140. Get Virtual Server Properties: Response for virtual servers of type "x-hyp" (Part 2)

---

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 07 Dec 2011 05:31:36 GMT
content-type: application/json;charset=UTF-8
content-length: 1133
{
  "acceptable-status": [
    "operating"
  ],
  "additional-status": null,
  "auto-start": false,
  "class": "virtual-server",
  "cpu-perf-mgmt-enabled": false,
  "description": "",
  "gpmp-status": "not-operating",
  "gpmp-support-enabled": false,
  "gpmp-version": "unavailable",
  "has-unacceptable-status": false,
  "hostname": null,
  "initial-memory": 32,
  "initial-share-mode": "relative",
  "initial-shares": 100.0,
  "initial-virtual-processors": 1,
  "ipl-device": "CMS",
  "ipl-load-parameters": "",
  "ipl-parameters": "AUTOOCR",
  "is-locked": false,
  "mac-prefix": {
    "mac-address": "02:ca:0d:00:00:00",
    "prefix-length": 24
  },
  "maximum-memory": 0,
  "maximum-share-mode": null,
  "maximum-shares": null,
  "maximum-virtual-processors": 64,
  "name": "MEV1A",
  "network-adapters": [],
  "object-id": "1738f85e-1b6f-11e1-a8ab-f0def165da96",
  "object-uri": "/api/virtual-servers/1738f85e-1b6f-11e1-a8ab-f0def165da96",
  "os-level": null,
  "os-name": null,
  "os-type": null,
  "parent": "/api/virtualization-hosts/911f4808-de39-11e0-a58e-f0def161133a",
  "password": "999999",
  "privilege-classes": "BDEG",
  "share-limit": "none",
  "status": "operating",
  "type": "zvm",
  "virtual-disks": [],
  "workloads": [
    "/api/workload-resource-groups/1fe39a72-de39-11e0-90a6-00215e67351a"
  ]
}

```

---

Figure 141. Get Virtual Server Properties: Response for virtual servers of type "zvm"

## Update Virtual Server Properties

The **Update Virtual Server Properties** operation updates one or more of the writeable properties of a Virtual Server. This operation is not supported for PR/SM virtual servers.

## HTTP method and URI

POST /api/virtual-servers/{*virtual-server-id*}

In this request, the URI variable {*virtual-server-id*} is the object ID of the Virtual Server object.

### Request body contents

Fields are only valid if they are supported for a virtual server of the targeted type. For instance a virtual server with **type** property "**power-vm**" may define a processing-mode property (PowerVM only) and **mac-prefix** property (PowerVM, z/VM) but not an **initial-share-mode** property (z/VM only).

Properties may also only be valid if prerequisite properties have specific values. For example a PowerVM virtual server's **initial-virtual-processors** is only valid when the **processing-mode** value is "**shared**".

All fields in the following table are optional. If a field is not included in the request body contents, its value will not be updated (unless a prerequisite field is changed, as noted in the table).

See the "Data Model" on page 261 for a complete definition of all properties including for what virtual server types they are valid and their prerequisites.

Field name	Type	Description
name	String	The value to be set as the virtual server's <b>name</b> property. Prerequisite: virtual server <b>type</b> is not " <b>zvm</b> ".
description	String	The value to be set as the virtual server's <b>description</b> property.
acceptable-status	Array of String Enum	The value to be set as the virtual server's <b>acceptable-status</b> property.
cpu-perf-mgmt-enabled	Boolean	The value to be set as the virtual server's <b>cpu-perf-mgmt-enabled</b> property. Default value: true. Valid only when the ensemble's <b>management-enablement-level</b> is " <b>automate</b> ".
processing-mode	String Enum	The value to be set as the virtual server's <b>processing-mode</b> property.
minimum-virtual-processors	Integer	The value to be set as the virtual server's <b>minimum-virtual-processors</b> property. Default value (used if <b>processing-mode</b> changes to " <b>shared</b> "): 1
initial-virtual-processors	Integer	The value to be set as the virtual server's <b>initial-virtual-processors</b> property. Default value (used if <b>processing-mode</b> changes to " <b>shared</b> "): 1
maximum-virtual-processors	Integer	The value to be set as the virtual server's <b>maximum-virtual-processors</b> property. Default value for power-vm (used if <b>processing-mode</b> changes to " <b>shared</b> "): max ( <b>maximum-allowed-dedicated-processors</b> , <b>initial-virtual-processors</b> )
minimum-dedicated-processors	Integer	The value to be set as the virtual server's <b>minimum-dedicated-processors</b> property. Default value (used if <b>processing-mode</b> changes to " <b>dedicated</b> "): 1
initial-dedicated-processors	Integer	The value to be set as the virtual server's <b>initial-dedicated-processors</b> property. Default value (used if <b>processing-mode</b> changes to " <b>dedicated</b> "): 1
maximum-dedicated-processors	Integer	The value to be set as the virtual server's <b>maximum-dedicated-processors</b> property. Default value: <b>maximum-allowed-dedicated-processors</b> Default value (used if <b>processing-mode</b> changes to " <b>dedicated</b> "): <b>maximum-allowed-dedicated-processors</b>
minimum-processing-units	Float	The value to be set as the virtual server's <b>minimum-processing-units</b> property. Default value (used if <b>processing-mode</b> changes to " <b>shared</b> "): 0.1
initial-processing-units	Float	The value to be set as the virtual server's <b>initial-processing-units</b> property. Default value (used if <b>processing-mode</b> changes to " <b>shared</b> "): 0.1

Field name	Type	Description
maximum-processing-units	Float	The value to be set as the virtual server's <b>maximum-processing-units</b> property. Default value (used if <b>processing-mode</b> changes to "shared"): min ( <b>maximum-virtual-processors</b> , <b>maximum-allowed-processing-units</b> )
initial-share-mode	String Enum	The value to be set as the virtual server's <b>initial-share-mode</b> property.
initial-shares	Float	The value to be set as the virtual server's <b>initial-shares</b> property. Default value based on <b>initial-share-mode</b> (used if <b>initial-share-mode</b> changes): <ul style="list-style-type: none"> <li>• relative: 100</li> <li>• absolute: 10.0</li> </ul>
share-limit	String Enum	The value to be set as the virtual server's <b>share-limit</b> property.
maximum-share-mode	String Enum	The value to be set as the virtual server's <b>maximum-share-mode</b> property. Default value (used if <b>share-limit</b> changes from "none" to "soft" or "hard" ): relative
maximum-shares	Float	The value to be set as the virtual server's <b>maximum-shares</b> property. Default value based on <b>maximum-share-mode</b> (used if <b>maximum-share-mode</b> changes): <ul style="list-style-type: none"> <li>• relative: 100</li> <li>• absolute: 10.0</li> </ul>
minimum-memory	Integer	The value to be set as the virtual server's <b>minimum-memory</b> property.
initial-memory	Integer	The value to be set as the virtual server's <b>initial-memory</b> property.
maximum-memory	Integer	The value to be set as the virtual server's <b>maximum-memory</b> property.
boot-mode	String Enum	The value to be set as the virtual server's <b>boot-mode</b> property.
boot-sequence	Array of String Enum	The value to be set as the virtual server's <b>boot-sequence</b> property.
boot-network-adapter-client-ip	String	The value to be set as the virtual server's <b>boot-network-adapter-client-ip</b> property. Default value (used if <b>boot-sequence</b> changes to ["network-adapter"]): an empty string
boot-network-adapter-subnet-ip	String	The value to be set as the virtual server's <b>boot-network-adapter-subnet-ip</b> property. Default value (used if <b>boot-sequence</b> changes to ["network-adapter"]): an empty string
boot-network-adapter-gateway-ip	String	The value to be set as the virtual server's <b>boot-network-adapter-gateway-ip</b> property. Default value (used if <b>boot-sequence</b> changes to "network-adapter"): an empty string
boot-network-adapter-server-ip	String	The value to be set as the virtual server's <b>boot-network-adapter-server-ip</b> property. Default value (used if <b>boot-sequence</b> changes to "network-adapter"): an empty string
keylock	String Enum	The value to be set as the virtual server's <b>keylock</b> property.
auto-start	Boolean	The value to be set as the virtual server's <b>auto-start</b> property.
dlpar-enabled	Boolean	The value to be set as the virtual server's <b>dlpar-enabled</b> property.
gpmp-support-enabled	Boolean	The value to be set as the virtual server's <b>gpmp-support-enabled</b> property.

Field name	Type	Description
rmc-virtual-device-address	String	The virtual device address of the resource monitoring and control network to create in order to support the Guest Platform Management Provider (GPMP).  Prerequisite: <b>type</b> is "zvm" and <b>gpmp-support-enabled</b> is true.  Required if <b>gpmp-support-enabled</b> changes to true.  1-4 character hex string
password	String	The value to be set as the virtual server's <b>password</b> property.
privilege-classes	String	The value to be set as the virtual server's <b>privilege-classes</b> property.
ipl-device	String	The value to be set as the virtual server's <b>ipl-device</b> property.
ipl-load-parameters	String	The value to be set as the virtual server's <b>ipl-load-parameters</b> property.
ipl-parameters	String	The value to be set as the virtual server's <b>ipl-parameters</b> property.
inband-monitoring-enabled	Boolean	The value to be set as the virtual server's <b>inband-monitoring-enabled</b> property.
keyboard-language	String	The value to be set as the virtual server's <b>keyboard-language</b> property.
acceptable-avail-status	Array of String Enum	The value to be set as the virtual server's <b>acceptable-avail-status</b> property.

## Description

This operation updates a virtual server's properties with the values specified on the identified Virtualization Host.

If the API user does not have access action permission to Virtual Server Details or the zManager management enablement level for the ensemble does not allow the setting of a provided property, a 403 (Forbidden) status code is returned. A 404 (Not Found) status code is also returned if the object-id *{virt-host-id}* does not identify a Virtualization Host object to which the API user has object-access permission. If the Virtualization Host is of type "zvm" and its *status* is "not-activated", a 409 (Conflict) status code is returned.

If the Request Body Contents fail to validate, a 400 (Bad Request) status code is returned. This may occur because the document defines a field that is not supported for the given virtual server type or includes a field that is not supported because a prerequisite is not met (e.g. attempting to set **maximum-share-mode** when **share-limit** is "none").

If the Request Body Contents are valid, the virtual server's properties are updated to their corresponding request body content's field's values. All fields are optional and may be excluded from the request body; if an optional field is not found in the request body, its property's value will not be modified. As indicated, a property's value is set to its default value if the field is not included in the request body and a prerequisite or other linked field is changed (e.g. if **initial-share-mode** is changed from "relative" to "absolute", and **initial-shares** is not defined in the request body, initial-shares will be defaulted to 10.0).

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Virtual Server object designated by *{virtual-server-id}*
- Action/task permission to the **Virtual Server Details** task.

- Action/task permission to the **Virtual Server Performance Details** task if the **gpmp-support-enabled** or **cpu-perf-mgmt-enabled** field is to be updated, otherwise action/task permission to the **Virtual Server Details** task for all other fields.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	100	The operation does not support a virtual server of the given type.
	101	A virtual server with the name specified in the request body already exists on the Virtualization Host with object-id <i>{virt-host-id}</i> .
	103	A “shares” ( <b>initial-shares</b> or <b>maximum-shares</b> ) property value type is not correct based on the value of its corresponding “share-mode” ( <b>initial-share-mode</b> or <b>maximum-share-mode</b> ). For example, the <b>initial-share-mode</b> is defined as “relative” but the <b>initial-shares</b> value is not an integer.
403 (Forbidden)	1	API user does not have action permission to the <b>Virtual Server Details</b> task.
	100	The ensemble's <b>management-enablement-level</b> property value does not allow the updating of a provided property.
404 (Not Found)	1	A virtual-server with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it.
409 (Conflict)	1	Virtual server status is not valid to perform the operation (does not allow the updating of a specified virtual server property).
	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	2	The request could not be processed because the HMC is not currently communicating with an element of a zBX needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/virtual-servers/7ba96572-0d72-11e1-892f-f0def14b63af HTTP/1.1
x-api-session: 64zdknfzh24fqw8f9v5g99h96rhrxhodmawtknx4iutmqrfrvf
content-type: application/json
content-length: 169
{
  "auto-start": true,
  "description": "Spacely Sprockets Web Store Web Server #1A",
  "initial-virtual-processors": 2,
  "keyboard-language": "en_US",
  "name": "SS-Web-Svr-1A"
}
```

---

Figure 142. Update Virtual Server Properties: Request for a virtual server of type "x-hyp"

---

```
204 No Content
date: Sat, 12 Nov 2011 21:08:38 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 143. Update Virtual Server Properties: Response for a virtual server of type "x-hyp"

## Create Network Adapter

The **Create Network Adapter** operation creates a virtual network adapter for the PowerVM, x Hyp, or z/VM virtual server with the given identifier. This operation is not supported for PR/SM virtual servers.

### HTTP method and URI

**POST** /api/virtual-servers/{*virtual-server-id*}/network-adapters

In this request, the URI variable {*virtual-server-id*} is the object ID of the Virtual Server object.

### Request body contents

Request body contents vary based on virtual server type.

For "power-vm":

Field name	Type	Rqd/Opt	Description
network-uri	String/ URI	Optional	The <b>network-uri</b> property of the network-adapter-power object. Default: <b>null</b>

For "x-hyp":

Field name	Type	Rqd/Opt	Description
network-uri	String/ URI	Optional	The <b>network-uri</b> property of the network-adapter-x-hyp object. Default: <b>null</b>
emulation-mode	String Enum	Required	The <b>emulation-mode</b> property of the network-adapter-x-hyp object.

For "zvm":

Field name	Type	Rqd/Opt	Description
virtual-device-address	String	Required	The <b>virtual-device-address</b> property of network-adapter-zvm object
type	String	Required	The <b>type</b> property of network-adapter-zvm object. The value may not be "rmc".
interface-type	String Enum	Required	The <b>interface-type</b> property of network-adapter-zvm object
real-device-address	String	Optional	The <b>real-device-address</b> property of network-adapter-zvm object
switch-uri	String/ URI	Required if <b>interface-type</b> is any value except <b>none</b>	If the <b>interface-type</b> is virtual-iedn, or virtual-qdio, specify the <b>switch-uri</b> property of the of the network-adapter-zvm object. If the <b>interface-type</b> is physical-iedn, physical-qdio, or physical-iqdn, specify the <b>element-uri</b> of a network-adapter-prsm object in the list of network adapters from the <b>network-adapters</b> property of the z/VM virtual server's <b>parent</b> property's Virtual Server object.
port-mode	String Enum	Required if <b>interface-type</b> is virtual-iedn or virtual-qdio	The <b>port-mode</b> property of network-adapter-zvm object
vlan-ids	String	Required if <b>interface-type</b> is virtual-qdio	The <b>vlan-ids</b> property of network-adapter-zvm object
network-uris	Array of String/ URI	Required if <b>interface-type</b> is virtual-iedn or physical-iedn	The <b>network-uris</b> property of network-adapter-zvm object

## Response body contents

On successful completion, the response body contains the URI of the created network-adapter object.

Field name	Type	Description
element-uri	String/URI	The <b>element-uri</b> property of the created network adapter object.

## Description

This operation creates the network adapters for the identified virtual server and then returns the URI of the created object. The response also includes a **Location** header that provides this URI.

If the virtual server is of type "power-vm" or "x-hyp" and its **status** is neither "not-operating" nor "not-activated", a 409 (Conflict) status code is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Virtual Server object designated by *{virtual-server-id}*
- Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in "Response body contents."



The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
409 (Conflict)	1	Virtual server status is not valid to perform the operation.
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
	109	The switch-uri conflicts with the network adapter type
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
POST /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/network-adapters HTTP/1.1
x-api-session: 3vkw2ncjqhmpzo6bsyjojllfgzd0kjoqvbhskz2m6z7esmjrr
content-type: application/json
content-length: 104
{
  "emulation-mode": "e1000",
  "network-uri": "/api/virtual-networks/f920171e-03f2-11e1-8e8e-0010184c8334"
}
```

---

Figure 144. Create Network Adapter: Request for a virtual server of type "x-hyp"

---

```

201 Created
server: zSeries management console API web server / 1.0
location: /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/network-adapters/
4a5073e6-0f9b-11e1-94dc-f0def14b63af
cache-control: no-cache
date: Tue, 15 Nov 2011 15:05:48 GMT
content-type: application/json;charset=UTF-8
content-length: 129
{
  "element-uri": "/api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/
network-adapters/4a5073e6-0f9b-11e1-94dc-f0def14b63af"
}

```

---

Figure 145. Create Network Adapter: Response for a virtual server of type "x-hyp"

## Delete Network Adapter

The **Delete Network Adapter** operation deletes an existing virtual network adapter specified by an object identifier for the PowerVM, x Hyp, or z/VM virtual server with the given identifier. This operation is not supported for PR/SM virtual servers.

### HTTP method and URI

**DELETE** /api/virtual-servers/{*virtual-server-id*}/network-adapters/{*network-adapter-id*}

#### URI variables

Variable	Description
{ <i>virtual-server-id</i> }	Object ID of the Virtual Server object
{ <i>network-adapter-id</i> }	Element ID of the network adapter

### Description

This operation deletes the network adapters for the identified virtual server.

If virtual server is of type "zvm" and the target network adapter is the Remote Monitoring and Control network adapter (type "rmc"), a 400 (Bad Request) is returned.

If the virtual server is of type "power-vm" or "x-hyp" and its status is neither "not-operating" nor "not-activated", a 409 (Conflict) status code is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server
- Action/task role permission to the **Virtual Server Details** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	106	The targeted network adapter is the Resource Monitoring and Control network adapter.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
409 (Conflict)	1	Virtual server status is not valid to perform the operation.
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
DELETE /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/network-adapters/
4a5073e6-0f9b-11e1-94dc-f0def14b63af HTTP/1.1
x-api-session: 3vkw2ncjqhmpz06bsyjoj1fgzd0kjoqvbvhszk2m6z7esmjr
```

---

Figure 146. Delete Network Adapter: Request

---

```
204 No Content
date: Tue, 15 Nov 2011 15:05:50 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

<No response body>

---

Figure 147. Delete Network Adapter: Response

## Get Network Adapter Properties

The **Get Network Adapter Properties** operation retrieves the properties of a single network adapter object that is designated by its element ID and the object ID of the owning virtual server.

### HTTP method and URI

```
GET /api/virtual-servers/{virtual-server-id}/network-adapters/{network-adapter-id}
```

#### URI variables

Variable	Description
<i>{virtual-server-id}</i>	Object ID of the virtual server that owns the virtual disk.
<i>{network-adapter-id}</i>	Element ID of the network adapter object for which properties are to be obtained.

## Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the network adapter object as defined in the “Data Model” on page 261. Field names and data types in the JSON object are the same as the property names and data types that are defined in the data model.

## Description

This operation returns the current properties for the network adapter object that is specified by *{network-adapter-id}*.

On successful execution, all of the current properties as defined by the Data Model for the network adapter object are provided in the response body and HTTP status code 200 (OK) is returned.

The URI path must designate an existing Virtual Server object and you must have object-access permission to it. Furthermore, the URI path must designate an existing network adapter object. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, you must have action/task permission to the **Virtual Server Details** task, otherwise status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Virtual Server object designated by *{virtual-server-id}*
- Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in the “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code that is indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission for this operation.
404 (Not Found)	1	A object-id in the URI <i>{virtual-server-id}</i> does not designate an existing Virtual Server object, or the API user does not have object access permission to it.
	2	The element-id in the URI <i>{network-adapter-id}</i> does not designate an existing network adapter object.

HTTP error status code	Reason code	Description
409 (Conflict)	1	Virtual server status is not valid to perform the operation.
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
	108	The virtual switch to which the network adapter was connected is no longer defined in z/VM.
	109	The virtual switch specified by the <b>switch-uri</b> property conflicts with the network adapter type in the request.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/virtual-servers/c967fcb2-f950-11e2-856c-3c970e47dddb/network-adapters/
    7db9e036-f951-11e2-9276-3c970e47dddb HTTP/1.1
x-api-session: v3motdpyqmgz4oe2hq5vtxhvoc80uea9efd4xjbc8emb2uung
```

---

Figure 148. Get Network Adapter Properties: Request

---

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Tue, 30 Jul 2013 19:50:10 GMT
content-type: application/json;charset=UTF-8
content-length: 316
{
  "adapter-type": "regular"
  "element-id": "7db9e036-f951-11e2-9276-3c970e47dddb",
  "element-uri": "/api/virtual-servers/c967fcb2-f950-11e2-856c-3c970e47dddb/network-adapters/
    7db9e036-f951-11e2-9276-3c970e47dddb",
  "emulation-mode": "virtio",
  "mac-address": "02:ff:7d:2d:b2:43",
  "network-uri": "/api/virtual-networks/43ea31fe-f950-11e2-b884-00215e67ac16"
}
```

---

Figure 149. Get Network Adapter Properties: Response

### Update Network Adapter

The **Update Network Adapter** operation modifies an existing virtual network adapter specified by an object identifier for the PowerVM, x Hyp, z/VM, or PR/SM virtual server with the given identifier.

## HTTP method and URI

POST /api/virtual-servers/{*virtual-server-id*}/network-adapters/{*network-adapter-id*}

### URI variables

Variable	Description
{ <i>virtual-server-id</i> }	Object ID of the Virtual Server object
{ <i>network-adapter-id</i> }	Element ID of the network adapter

## Request body contents

Request body contents vary based on virtual server type.

For "prsm":

Field name	Type	Rqd/Opt	Description
network-uris	Array of String/ URI	Optional	The <b>network-uris</b> property of the network-adapter-prsm object.

For "power-vm":

Field name	Type	Rqd/Opt	Description
network-uri	String/ URI	Optional	The <b>network-uri</b> property of the network-adapter-power object. Specify a value of <b>null</b> to indicate that the network adapter should not be connected to a virtual network.

For "x-hyp":

Field name	Type	Rqd/Opt	Description
network-uri	String/ URI	Optional	The <b>network-uri</b> property of the network-adapter-x-hyp object. Specify a value of <b>null</b> to indicate that the network adapter should not be connected to a virtual network.
emulation-mode	String Enum	Optional	The <b>emulation-mode</b> property of the network-adapter-x-hyp object.

For "zvm":

Field name	Type	Rqd/Opt	Description
virtual-device-address	String	Optional	The <b>virtual-device-address</b> property of network-adapter-zvm object
interface-type	String Enum	Optional	The <b>interface-type</b> property of network-adapter-zvm object
real-device-address	String	Optional	The <b>real-device-address</b> property of network-adapter-zvm object
switch-uri	String/ URI	Optional; Allowed only if <b>interface-type</b> is virtual-iedn or virtual-qdio	The <b>switch-uri</b> property of network-adapter-zvm object.

Field name	Type	Rqd/Opt	Description
port-mode	String Enum	Optional; Allowed only if <b>interface-type</b> is virtual-iedn or virtual-qdio	The <b>port-mode</b> property of network-adapter-zvm object
vlan-ids	String	Optional; Allowed only if <b>interface-type</b> is virtual-qdio	The <b>vlan-ids</b> property of network-adapter-zvm object
network-uris	Array of String/ URI	Optional; Allowed only if <b>interface-type</b> is virtual-iedn or physical-iedn	The <b>network-uris</b> property of network-adapter-zvm object

## Description

This operation modifies the network adapters for the identified virtual server and then returns its properties.

If virtual server is of type "**zvm**" and the target network adapter is the Remote Monitoring and Control network adapter (type "**rmc**"), a 400 (Bad Request) is returned.

If the virtual server is of type "**power-vm**" or "**x-hyp**" and its **status** is neither "**not-operating**" nor "**not-activated**", a 409 (Conflict) status code is returned. If the virtual server is of type "**zvm**", the network adapter **interface-type** is "**virtual-iedn**" or "**virtual-qdio**", and the virtual-switch to which the network adapter was connected no longer exists, a 409 (Conflict) status code is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server.
- Object access permission to the Virtual Network object(s)
- Action/task permission to the **Virtual Server Details** task

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	106	The targeted network adapter is the Resource Monitoring and Control network adapter.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.

HTTP error status code	Reason code	Description
409 (Conflict)	1	Virtual server status is not valid to perform the operation.
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
	108	The virtual switch to which the network adapter was connected is no longer defined in z/VM.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/network-adapters/
4a5073e6-0f9b-11e1-94dc-f0def14b63af HTTP/1.1
x-api-session: 3vkw2ncjqhmpzo6bsyjoj1fgzd0kjoqvbvhsz2m6z7esmjrr
content-type: application/json
content-length: 105
{
  "emulation-mode": "virtio",
  "network-uri": "/api/virtual-networks/f3aece54-0db8-11e1-9e2c-00215e69dea0"
}
```

---

Figure 150. Update Network Adapter: Request for a virtual server of type "x-hyp"

---

```
204 No Content
date: Tue, 15 Nov 2011 15:05:48 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 151. Update Network Adapter: Response for a virtual server of type "x-hyp"

## Reorder Network Adapter

The **Reorder Network Adapter** operation reorders the virtual network adapter defined for the PowerVM or x Hyp virtual server with the given identifier. This operation is not supported for z/VM and PR/SM virtual servers.

### HTTP method and URI

```
POST /api/virtual-servers/{virtual-server-id}/operations/reorder-network-adapters
```

In this request, the URI variable *{virtual-server-id}* is the object ID of the Virtual Server object.



## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
network-adapter-uris	Array of String/URI	Required	Ordered list of element-uris for the network adapter object. The order of this array is significant.

## Description

This operation reorders the network adapters for the identified virtual server.

If the virtual server is of type **"power-vm"** or **"x-hyp"** and its **status** is neither **"not-operating"** nor **"not-activated"**, a 409 (Conflict) status code is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server
- Object access permission to the Virtual Network object(s)
- Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
409 (Conflict)	1	Virtual server status is not valid to perform the operation.
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Example HTTP interaction

---

```
POST /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/operations/
  reorder-network-adapters HTTP/1.1
x-api-session: 3vkw2ncjqhmpzo6bsyjojllfgzd0kjoqvhskz2m6z7esmjrr
content-type: application/json
content-length: 256
{
  "network-adapter-uris": [
    "/api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/network-adapters/
4a5073e6-0f9b-11e1-94dc-f0def14b63af",
    "/api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/network-adapters/
582b5d0e-0db7-11e1-b1f1-f0def14b63af"
  ]
}
```

---

Figure 152. Reorder Network Adapter: Request

---

```
204 No Content
date: Tue, 15 Nov 2011 15:05:48 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 153. Reorder Network Adapter: Response

## Create Virtual Disk

The **Create Virtual Disk** operation adds a virtual disk to the specified virtual server. This operation is not supported for PR/SM virtual servers.

### HTTP method and URI

**POST** /api/virtual-servers/{*virtual-server-id*}/virtual-disks

In this request, the URI variable {*virtual-server-id*} is the object ID of the virtual server that will own the new virtual disk.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The <b>name</b> property of the virtual disk object.
description	String	Optional	The <b>description</b> property of the virtual disk object.
size	Long	Required	The <b>size</b> property of the virtual disk object. Required and only valid for virtual disks with a <b>type</b> of <b>"storage-group-based"</b> .
type	String Enum	Required	The <b>type</b> property of the virtual disk object.
emulation-mode	String Enum	Required	The <b>emulation-mode</b> property of the virtual disk object. Required and only valid for all virtual disks for a virtual server with <b>type</b> of <b>"x-hyp"</b> .
access-mode	String Enum	Required	The <b>access-mode</b> property of the virtual disk object. Required and only valid for all virtual disks for a virtual server with <b>type</b> of <b>"zvm"</b> . Value may not be <b>"unsupported"</b> .

Field name	Type	Rqd/Opt	Description
device-address	String (1-4)	Required	The <b>device-address</b> property of the virtual disk object. Required for all virtual disks for a virtual server with <b>type</b> of "zvm".
read-password	String (0-8)	Optional	The <b>read-password</b> property of the virtual disk. For all virtual disks for a virtual server with <b>type</b> "zvm" and a virtual disk <b>type</b> of either "fullpack" or "storage-group-based".
write-password	String (0-8)	Optional	The <b>write-password</b> property of the virtual disk. For all virtual disks for a virtual server with <b>type</b> "zvm" and a virtual disk <b>type</b> of either "fullpack" or "storage-group-based".
multi-password	String (0-8)	Optional	The <b>multiple-write</b> password property of the virtual disk. For all virtual disks for a virtual server with <b>type</b> "zvm" and a virtual disk <b>type</b> of either "fullpack" or "storage-group-based".
password	String (0-8)	Optional	The password for the linked virtual disk. For virtual disks with a <b>type</b> of "linked".
backing- virtualization- host-storage- resource	String/ URI	Required	The <b>backing-virtualization-host-storage-resource</b> property of the virtual disk. Required and only valid for virtual disks with a <b>type</b> of "fullpack".
source-virtual-disk	String/ URI	Required	The <b>source-virtual-disk</b> property of the virtual disk. Required and only valid for virtual disks with a <b>type</b> of "linked".
backing-storage- group	String/ URI	Required	The <b>backing-storage-group</b> property of the virtual disk. Required and only valid for virtual disks with a <b>type</b> of "storage-group-based".

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	The <b>element-uri</b> of the newly created Virtual Disk object.

## Description

The Create Virtual Disk operation adds a virtual disk to the virtual server specified by the *{virtual-server-id}* portion of the request URI.

Upon successful completion, the **element-uri** field of the Response Body and the value of the **Location** response header identify the new virtual disk.

If this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

The URI path must designate an existing Virtual Server object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must also have action/task permission to the **Virtual Server Details** action as well, otherwise status code 403 (Forbidden) is returned.

If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

If the virtual server is of **type** "power-vm" or "x-hyp" and its **status** is not "operating" , "not-operating" or "not-activated", a 409 (Conflict) status code is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Virtual Server object designated by *{virtual-server-id}*
- Object access permission to the storage resources to be used by the virtual disk:
  - If virtual disk **type** is "fullpack", object access permission to the storage resource used by the Virtualization Host Storage Resource object designated by **backing-virtualization-host-storage-resource**.
  - If virtual disk **type** is "storage-group-based", object access permission to the storage resource used by every virtualization host storage resource owned by the Virtualization Host Storage Group object designated by **backing-storage-group**.
  - If virtual disk **type** is "linked", object access permission to the storage resources used by the virtual disk designated by **source-virtual-disk** and to the virtual server that owns it.
- Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and both the response body and the **Location** response header contain the URI of the newly created object.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	7	Another Virtual server instance running in this hypervisor with SCSI passthrough enabled.
	100	The operation does not support a virtual server of the given type.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	107	The <b>backing-virtualization-host-storage-resource</b> property does not designate an existing Virtualization Host Storage Resource object, or the API user does not have object access permission to it.
	108	The <b>backing-storage-group</b> property does not designate an existing Virtualization Host Storage Group object, or the API user does not have object access permission to it.
	109	The <b>source-virtual-disk</b> property does not designate an existing virtual disk object, or the API user does not have object access permission to it.
409 (Conflict)	1	Virtual server status is not valid to perform the operation.
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/virtual-disks HTTP/1.1
x-api-session: 29chofb3vqxpijctnd2kf5fiwnf118z62mq3yomtqjyz8bx2gn
content-type: application/json
content-length: 259
{
  "backing-virtualization-host-storage-resource": "/api/virtualization-hosts/
    2f7bf364-03f8-11e1-8eda-001f163805d8/virtualization-host-storage-resources/
    37699380-0fa9-11e1-b69e-f0def14b63af",
  "description": "Boot filesystem",
  "name": "boot",
  "type": "fullpack"
}
```

---

Figure 154. Create Virtual Disk: Request for a virtual server of type “power-vm”

---

```
201 Created
server: zSeries management console API web server / 1.0
location: /api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/virtual-disks/
    c547fb4e-0fac-11e1-842b-f0def14b63af
cache-control: no-cache
date: Tue, 15 Nov 2011 17:10:57 GMT
content-type: application/json;charset=UTF-8
content-length: 126
{
  "element-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/virtual-disks/
    c547fb4e-0fac-11e1-842b-f0def14b63af"
}
```

---

Figure 155. Create Virtual Disk: Response for a virtual server of type “power-vm”

---

```
| POST /api/virtual-servers/7fa1b30c-7d30-11e3-aaaa-00262dfc0ec9/virtual-disks/
|   817ceb06-7d30-11e3-b0c0-5cf3fc260857
| {
|   "emulation-mode": "virtio-scsi",
|   "backing-virtualization-host-storage-resource": "/api/virtualization-hosts/332de83c-c15c-11e2-
|     9510-00262dfc0ec9/virtualization-host-storage-resources/7ef9f73e-7d30-11e3-aaaa-
|     00262dfc0ec9",
|   "type": "fullpack",
|   "name": "fullpack_virtual_disk_3600",
|   "description": "fullpack virtual disk for API FVT testcase 3600"
| }
|
|
```

---

Figure 156. Create Virtual Disk: Request for a virtual server of type “x-hyp”

---

```

| 201 Created
| server: zSeries management console API web server / 2.0,
| location: /api/virtual-servers/7fa1b30c-7d30-11e3-aaaa-00262dfc0ec9/virtual-disks/
|   817ceb06-7d30-11e3-b0c0-5cf3fc260857,
| cache-control: no-cache,
| date: Tue, 14 Jan 2014 15:28:23 GMT,
| content-type: application/json; charset=UTF-8
| content-length: 126
| {
|   element-uri: "/api/virtual-servers/7fa1b30c-7d30-11e3-aaaa-00262dfc0ec9/virtual-disks/
|     817ceb06-7d30-11e3-b0c0-5cf3fc260857"
| }

```

---

Figure 157. Create Virtual Disk: Response for a virtual server of type "x-hyp"

## Delete Virtual Disk

The **Delete Virtual Disk** operation removes a specified virtual disk from the specified virtual server. This operation is not supported for PR/SM virtual servers.

### HTTP method and URI

**DELETE** /api/virtual-servers/{*virtual-server-id*}/virtual-disks/{*virtual-disk-id*}

#### URI variables

Variable	Description
{ <i>virtual-server-id</i> }	Object ID of the Virtual Server object
{ <i>virtual-disk-id</i> }	Element ID of the virtual disk object

### Description

The **Delete Virtual Disk** operation removes a specified virtual disk from the specified virtual server. The virtual disk is identified by the {*virtual-disk-id*} variable in the URI, and the virtual server is identified by the {*virtual-server-id*} variable in the URI.

Upon successfully removing the virtual disk, HTTP status code 204 (No Content) is returned and no response body is provided.

The URI path must designate an existing Virtual Server object and the API user must have object-access permission to it. Furthermore, the URI path must designate an existing virtual disk object. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must also have action/task permission to the Virtual Server Details action as well, otherwise status code 403 (Forbidden) is returned.

If the virtual server is of type **"power-vm"** or **"x-hyp"** and its **status** is neither **"not-operating"** nor **"not-activated"**, a 409 (Conflict) status code is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Virtual Server object designated by {*virtual-server-id*}.
- If virtual disk **type** is **"fullpack"**, object access permission to the Virtualization Host Storage Resource object designated by **backing-virtualization-host-storage-resource**

- If virtual disk **type** is "storage-group-based", object access permission to all virtualization host storage resources owned by the virtualization host Group object designated by **backing-storage-group**.
- If virtual disk **type** is "linked", object access permission to virtual disk designated by **source-virtual-disk** and the virtual server that owns it.
- Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	100	The operation does not support a virtual server of the given type.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI <i>{virtual-server-id}</i> does not designate an existing Virtual Server object, or the API user does not have object access permission to it.
	2	The object ID in the URI <i>{virtual-disk-id}</i> does not designate an existing virtual disk object.
409 (Conflict)	1	Virtual server status is not valid to perform the operation.
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Example HTTP interaction

---

```
DELETE /api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/virtual-disks/
c547fb4e-0fac-11e1-842b-f0def14b63af HTTP/1.1
x-api-session: 29chofb3vqxpijctnd2kf5fiwnf118z62mq3yomtqjyz8bx2gn
```

---

Figure 158. Delete Virtual Disk: Request

---

```
204 No Content
date: Tue, 15 Nov 2011 17:11:03 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

```
<No response body>
```

---

Figure 159. Delete Virtual Disk: Response

## Get Virtual Disk Properties

The **Get Virtual Disk Properties** operation retrieves the properties of a single virtual disk object that is designated by its object ID and the object ID of the owning virtual server. This operation is not supported for PR/SM virtual servers.

### HTTP method and URI

```
GET /api/virtual-servers/{virtual-server-id}/virtual-disks/{virtual-disk-id}
```

#### URI variables

Variable	Description
<i>{virtual-server-id}</i>	Object ID of the virtual server that will own the virtual disk.
<i>{virtual-disk-id}</i>	Element ID of the virtual disk object for which properties are to be obtained.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the virtual disk object as defined in the Data Model section above. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

This operation returns the current properties for the virtual disk object specified by *{virtual-disk-id}*.

On successful execution, all of the current properties as defined by the “Data Model” on page 261 are provided in the response body and HTTP status code 200 (OK) is returned.

The URI path must designate an existing Virtual Server object and the API user must have object-access permission to it. Furthermore, the URI path must designate an existing virtual disk object. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action/task permission to the **Virtual Server Details** action, otherwise status code 403 (Forbidden) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Virtual Server object designated by *{virtual-server-id}*
- If virtual disk **type** is **"fullpack"**, object access permission to the Virtualization Host Storage Resource object designated by **backing-virtualization-host-storage-resource**



- If virtual disk **type** is **"storage-group-based"**, object access permission to all virtualization host storage resources owned by the Virtualization Host Storage Group object designated by **backing-storage-group**.
- If virtual disk **type** is **"linked"**, object access permission to virtual disk designated by **source-virtual-disk** and the virtual server that owns it.
- Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	100	The operation does not support a virtual server of the given type.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI <i>{virtual-server-id}</i> does not designate an existing Virtual Server object, or the API user does not have object access permission to it.
	2	The object ID in the URI <i>{virtual-disk-id}</i> does not designate an existing virtual disk object.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMIPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMIPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/virtual-disks/
    c547fb4e-0fac-11e1-842b-f0def14b63af HTTP/1.1
x-api-session: 29chofb3vqxpijctnd2kf5fiwnf118z62mq3yomtqjyz8bx2gn
```

---

Figure 160. Get Virtual Disk Properties: Request for a virtual server of type "power-vm"

---

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 15 Nov 2011 17:10:57 GMT
content-type: application/json;charset=UTF-8
content-length: 516
{
  "backing-virtualization-host-storage-resource": "/api/virtualization-hosts/
  2f7bf364-03f8-11e1-8eda-001f163805d8/virtualization-host-storage-resources/
  37699380-0fa9-11e1-b69e-f0def14b63af",
  "description": "Boot filesystem",
  "element-id": "c547fb4e-0fac-11e1-842b-f0def14b63af",
  "element-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/virtual-disks/
  c547fb4e-0fac-11e1-842b-f0def14b63af",
  "name": "boot",
  "owner": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af",
  "size": 17179869184,
  "type": "fullpack"
}

```

---

Figure 161. Get Virtual Disk Properties: Response for a virtual server of type "power-vm"

---

```

| GET:/api/virtual-servers/7fa1b30c-7d30-11e3-aaaa-00262dfc0ec9/virtual-disks/817ceb06-7d30-
| 11e3-b0c0-5cf3fc260857
| x-api-session: 3rprukcrfwl0wfb57u0rw4e0epi5579ck253zbc02plfwypu0z
|
|

```

---

Figure 162. Get Virtual Disk Properties: Request for a virtual server of type "x-hyp"

---

```

| 200 OK
| server: zSeries management console API web server / 2.0,
| cache-control: no-cache,
| date: Tue, 14 Jan 2014 15:28:37 GMT
| content-type: application/json;charset=UTF-8,
| content-length: 605
| {
|   "description": "fullpack virtual disk for API FVT testcase 3600_updated",
|   "element-uri": "/api/virtual-servers/7fa1b30c-7d30-11e3-aaaa-00262dfc0ec9/virtual-disks/
| 817ceb06-7d30-11e3-b0c0-5cf3fc260857",
|   "backing-virtualization-host-storage-resource": "/api/virtualization-hosts/
| 332de83c-c15c-11e2-9510-00262dfc0ec9/virtualization-host-storage-resources/
| 7ef9f73e-7d30-11e3-aaaa-00262dfc0ec9",
|   "element-id": "817ceb06-7d30-11e3-b0c0-5cf3fc260857",
|   "emulation-mode": "virtio-scsi",
|   "owner": "/api/virtual-servers/7fa1b30c-7d30-11e3-aaaa-00262dfc0ec9",
|   "size": 1234,
|   "type": "fullpack",
|   "name": "fullpack_virtual_disk_3600_updated"
| }
|
|

```

---

Figure 163. Get Virtual Disk Properties: Response for a virtual server of type "x-hyp"

## Update Virtual Disk Properties

The **Update Virtual Disk Properties** operation updates one or more of the writeable properties of a virtual disk object. This operation is not supported for PR/SM virtual servers.

## HTTP method and URI

POST /api/virtual-servers/{*virtual-server-id*}/virtual-disks/{*virtual-disk-id*}

### URI variables

Variable	Description
{ <i>virtual-server-id</i> }	Object ID of the virtual server that owns the virtual disk.
{ <i>virtual-disk-id</i> }	Element ID of the virtual disk object for which properties are to be updated.

## Request body contents

The request body is expected to contain a JSON object that provides the new values of any writeable property that is to be updated by this operation. Field names and data types in this JSON object are expected to match the corresponding property names and data types defined in the data model for this object type. The JSON object can and should omit fields for properties whose values are not to be changed by this operation. All such fields are optional.

The JSON object may also contain the following non-data-model field:

Field name	Type	Rqd/Opt	Description
password	String (0-8)	Optional	The new password for a linked virtual disk. For virtual disks with a <b>type</b> of "linked"

Request body **access-mode** property value may not be "unsupported".

## Description

This operation updates writeable properties of the virtual disk object specified by {*virtual-disk-id*}.

The request body contains an object with one or more fields with field names that correspond to the names of properties for this object. On successful execution, the value of each corresponding property of the object is updated with the value provided by the input field, and status code 204 (No Content) is returned without supplying any response body. The request body does not need to specify a value for all writeable properties, but rather can and should contain only fields for the properties to be updated. Object properties for which no input value is provided remain unchanged by this operation.

If the update changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

The URI path must designate an existing virtual disk object and the API user must have object-access permission to it. Furthermore, the URI path must designate an existing Virtual Server object. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must also have action/task permission to the **Virtual Server Details** action as well, otherwise status code 403 (Forbidden) is returned.

The request body is validated against the data model for this object type to ensure that it contains only writeable properties and the data types of those properties are as required. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

If the virtual server is of **type** "power-vm" or "x-hyp" and its **status** is neither "not-operating" nor "not-activated", a 409 (Conflict) status code is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Virtual Server object designated by *{virtual-server-id}*
- If virtual disk **type** is "fullpack", object access permission to the Virtualization Host Storage Resource object designated by **backing-virtualization-host-storage-resource**
- If virtual disk **type** is "storage-group-based", object access permission to all virtualization host storage resources owned by the Virtualization Host Storage Group object designated by **backing-storage-group**.
- If virtual disk **type** is "linked", object access permission to virtual disk designated by **source-virtual-disk** and the virtual server that owns it.
- Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	100	The operation does not support a virtual server of the given type.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI <i>{virtual-server-id}</i> does not designate an existing Virtual Server object, or the API user does not have object access permission to it.
	2	The object ID in the URI <i>{virtual-disk-id}</i> does not designate an existing virtual disk object.
409 (Conflict)	1	Virtual server status is not valid to perform the operation.
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```

| POST /api/virtual-servers/7falb30c-7d30-11e3-aaaa-00262dfc0ec9/virtual-disks/817ceb06-
| 7d30-11e3-b0c0-5cf3fc260857
| x-api-session: 3rprukcrfwl0wfb57u0rw4e0epi5579ck253zbc02p1fwypu0z,
| content-type: application/json
| {
|   "emulation-mode": "virtio-scsi",
|   "name": "fullpack_virtual_disk_3600_updated",
|   "description": "fullpack virtual disk for API FVT testcase 3600_updated"
| }

```

Figure 164. Update Virtual Disk Properties: Request for a virtual server of type "x-hyp"

```

| 204 No Content
| date: Tue, 14 Jan 2014 15:28:36 GMT,
| cache-control: no-cache,
| server : zSeries management console API web server / 2.0
| < No Response Body>

```

Figure 165. Update Virtual Disk Properties: Response for a virtual server of type "x-hyp"

## Reorder Virtual Disks

The **Reorder Virtual Disks** operation reorders the virtual disks defined for the PowerVM or x Hyp virtual server with the given identifier. This operation is not supported for z/VM and PR/SM virtual servers.

### HTTP method and URI

**POST** /api/virtual-servers/{*virtual-server-id*}/operations/reorder-virtual-disks

In this request, the URI variable {*virtual-server-id*} is the object ID of the Virtual Server object.

### Request body contents

The JSON object may also contain the following field:

Field name	Type	Rqd/Opt	Description
virtual-disk-uris	Array of String/ URI	Required	Ordered list of virtual disk object <b>element-uri</b> property values. The order of URIs in the array defines the new order of the virtual disks.  The array size must equal the number of virtual disks for the virtual server and the array must include the <b>element-uri</b> of each virtual disk in the virtual server exactly once.

## Description

This operation reorders the virtual disks for the identified virtual server.

If the URI path does not designate an existing Virtual Server object or the API user does not have object access permission to that virtual server and the storage-resource objects backing its virtual disks, a 404 (Not Found) status code is returned. If the virtual server is of **type "zvm"** or **"prsm"**, a 400 (Bad Request) status code is returned. If the API user does not have action/task permission to the **Virtual Server Details** action, a 403 (Forbidden) status code is returned.

If **virtual-disk-uris** array size is not equal to the number of virtual disks in the virtual server or the array does not include the **element-uri** of each virtual disk in the virtual server, a 400 (Bad Request) status code is returned.

If the virtual server is of **type "power-vm"** or **"x-hyp"** and its status is neither **"not-operating"** nor **"not-activated"**, a 409 (Conflict) status code is returned.

If the request body contents are valid, the virtual server's virtual disks are reordered to match the order of their **element-uri** properties in the **virtual-disk-uris** array.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Virtual Server object designated by *{virtual-server-id}*
- Object access permission to the storage resources used by each virtual disk in the virtual server:
  - If virtual disk **type** is **"fullpack"**, object access permission to the storage resource used by the Virtualization Host Storage Resource object designated by **backing-virtualization-host-storage-resource**.
  - If virtual disk **type** is **"storage-group-based"**, object access permission to the storage resource used by every virtualization host storage resource owned by the Virtualization Host Storage Group object designated by **backing-storage-group**.
  - If virtual disk **type** is **"linked"**, object access permission to the storage resources used by the virtual disk designated by **source-virtual-disk** and the virtual server that owns it.
- Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	7	The <b>virtual-disk-uris</b> array does not designate a virtual disk of the virtual server identified in the request URI.
	100	The operation does not support a virtual server of the given type.
	106	The <b>virtual-disk-uris</b> array contains an invalid number of values.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI <i>{virtual-server-id}</i> does not designate an existing Virtual Server object, or the API user does not have object access permission to it.
409 (Conflict)	1	Virtual server status is not valid to perform the operation.
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.

HTTP error status code	Reason code	Description
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM Virtualization Host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM Virtualization Host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/operations/reorder-virtual-disks HTTP/1.1
x-api-session: 29chofb3vqxpijctnd2kf5fiwnf118z62mq3yomtqjyz8bx2gn
content-type: application/json
content-length: 247
{
  "virtual-disk-uris": [
    "/api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/virtual-disks/
c1418fba-0fac-11e1-903e-f0def14b63af",
    "/api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/virtual-disks/
c1149ce4-0fac-11e1-903e-f0def14b63af"
  ]
}
```

---

Figure 166. Reorder Virtual Disks: Request

---

```
204 No Content
date: Tue, 15 Nov 2011 17:10:50 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 167. Reorder Virtual Disks: Response

## Activate Virtual Server

The **Activate Virtual Server** operation accepts a request to asynchronously activate the identified virtual server.

### HTTP method and URI

**POST** /api/virtual-servers/{*virtual-server-id*}/operations/activate

In this request, the URI variable {*virtual-server-id*} is the object ID of the Virtual Server object.

### Response body contents

Once the activation request is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve activation status updates.

## Asynchronous result description

Once the activation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Activate Virtual Server** request.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "**complete**", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in the operation description. The **job-results** field is null for asynchronous virtual server activation jobs.

### Description

This operation asynchronously activates the identified virtual server. If the URI does not identify a valid virtual server a 404 (Not Found) status code is returned. If the user does not have authority to perform the Activate action, a 403 (Forbidden) status code is returned.

For specific hypervisors, the Unified Resource Manager limits the number of virtual servers that can be active at any given time. A virtual server is active when it is in Starting, Operating or Stopping state.

- The maximum number of active virtual servers for PowerVM is 32.
- The maximum number of active virtual servers for x Hyp is 99.

The virtual server activation job is then initiated and a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the activation job. See "Query Job Status" on page 52 for information on how to query job status. When the activate job has completed, an asynchronous result message is sent, including "Job status and reason codes" on page 345.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server
- Action/task permission to the **Activate** task.

### HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in "Response body contents" on page 343.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.



Table 80. Activate Virtual Server: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	7	Another virtual server instance is running in this hypervisor with SCSI passthrough enabled.
403 (Forbidden)	1	API user does not have access action permission to <b>Activate</b> .
404 (Not Found)	1	A virtual server with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Job status and reason codes

HTTP error status code	Reason code	Description
200 (OK)	N/A	Activation completed successfully
500 (Server Error)	100	Virtual server activation failed
	101	Virtual server activation job timed out

### Example HTTP interaction

---

```
POST /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/operations/activate HTTP/1.1
x-api-session: 1f6hv807yexwhfyk1p8ygrc8876y48adda2dfpvuz4t9iqeo9k
```

---

Figure 168. Activate Virtual Server: Request

---

```
202 Accepted
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 16 Nov 2011 18:49:07 GMT
content-type: application/json;charset=UTF-8
content-length: 60
{
  "job-uri": "/api/jobs/aa24c23e-1083-11e1-87b6-0010184c8334"
}
```

---

Figure 169. Activate Virtual Server: Response

### Deactivate Virtual Server

The **Deactivate Virtual Server** operation accepts a request to asynchronously deactivate the identified virtual server. If there is already a deactivate request active for the identified virtual server, it will be canceled and the action requested by this subsequent invocation will be performed instead.

## HTTP method and URI

POST /api/virtual-servers/{*virtual-server-id*}/operations/deactivate

In this request, the URI variable {*virtual-server-id*} is the object ID of the Virtual Server object.

## Request body contents

An optional request body can be specified as a JSON object with the following field:

Field name	Type	Rqd/Opt	Description
shutdown-timeout	Integer	Optional	<p>Amount of time, in seconds, to allow a virtual server to shutdown. After the elapsed time has passed, the virtual server will be forcefully terminated.</p> <p>The value may be -1 to indicate to wait “forever” or any integer value between 0 and 86400 to specify an exact wait time in seconds.</p> <p>Specifying this property will cause the deactivate operation to use this value, regardless of the setting of the <b>use-hypervisor-shutdown-timeout</b> property of the virtual server.</p>

## Response body contents

Once the deactivation request is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve deactivation status updates.

## Asynchronous result description

Once the deactivation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Deactivate Virtual Server** request.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "**complete**", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description below. The **job-results** field is null for asynchronous virtual server deactivation jobs.

## Description

This operation asynchronously deactivates the identified virtual server. If the URI does not identify a valid virtual server a 404 (Not Found) status code is returned. If the user does not have authority to perform the Deactivate action, a 403 (Forbidden) status code is returned.

The virtual server deactivation job is then initiated and a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the deactivation job. See “Query Job Status” on page 52 for information on how to query job status. When the deactivate job has completed, an asynchronous result message is sent, including “Job status and reason codes” on page 347.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server
- Action/task permission to the **Deactivate** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 346.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have access action permission to <b>Activate</b> .
404 (Not Found)	1	A virtual server with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

HTTP error status code	Reason code	Description
200 (OK)	N/A	Deactivation completed successfully
500 (Server Error)	100	Virtual server deactivation failed
	101	Virtual server deactivation job timed out

## Example HTTP interaction

---

```
POST /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/operations/deactivate HTTP/1.1
x-api-session: 1f6hv807yexwhfyk1p8ygrc8876y48adda2dfpvuz4t9iqeo9k
```

---

Figure 170. Deactivate Virtual Server: Request

---

```

202 Accepted
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 16 Nov 2011 18:50:49 GMT
content-type: application/json;charset=UTF-8
content-length: 60
{
  "job-uri": "/api/jobs/e6eb3504-1083-11e1-87b6-0010184c8334"
}

```

---

Figure 171. Deactivate Virtual Server: Response

## Mount Virtual Media

The **Mount Virtual Media** operation starts an asynchronous operation to mount the specified zManager-provided ISO to the identified PowerVM or x Hyp virtual server. This operation is not supported for PR/SM and z/VM virtual servers.

### HTTP method and URI

**POST** /api/virtual-servers/{*virtual-server-id*}/operations/mount-virtual-media

In this request, the URI variable {*virtual-server-id*} is the object ID of the Virtual Server object.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
iso	String Enum	The ID of the zManager-provided ISO to mount to the virtual server. Values: <ul style="list-style-type: none"> <li>"gpmp" – Guest Platform Performance Management ISO</li> <li>"virtio" – RedHat VirtIO driver CD</li> </ul>

### Asynchronous result description

Once the mount job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Mount Virtual Media** request.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "**complete**", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description below. The **job-results** field is null for asynchronous mount virtual media jobs.

### Description

The Mount Virtual Media operation starts an asynchronous operation to mount the specified zManager-provided ISO to the identified PowerVM or x Hyp virtual server. This operation is not supported for PR/SM and z/VM virtual servers.

Virtual media is a virtual DVD drive that can be attached to a virtual server. Because virtual media is allocated from IBM blade storage instead of on a real media, it may be faster to use than a physical real device.

Preloaded ISO images come in two variants. The first is an ISO image that is packaged on the SE, such as the GPMP Installation Image. When mounting this type of ISO image to the virtual server, the image will be uploaded from the SE to the hypervisor. The second type is an ISO image preloaded onto the hypervisor firmware, such as the Red Hat VirtIO Driver Image. When mounting this type of ISO image, no uploading will occur.

The response from this operation, success or failure, will be presented asynchronously. Once the mount operation job has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body from the **Mount Virtual Media** request.

- | If the API user does not have action permission for the **Virtual Server Details** task, a 403 (Forbidden) status code is returned. A 404 (Not Found) status code is returned if the object-id *{virtual-server-id}* does not identify a Virtual Server object to which the API user has object-access permission.

A mount operation cannot be performed when a virtual server or its Virtualization Host is busied by another operation or when the Virtualization Host **status** is not **"operating"**. Under these conditions a 409 (Conflict) status code is returned.

The mount virtual media job is then initiated and a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the mount job. See "Query Job Status" on page 52 for information on how to query job status. When the mount job has completed, an asynchronous result message is sent, including "Job status and reason codes" on page 350.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server
- | • Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	100	The operation does not support a virtual server of the given type.
403 (Forbidden)	1	API user does not have access action permission to <b>Virtual Server Details</b> task.
404 (Not Found)	1	A virtual server with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it.
409 (Conflict)	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
	101	Parent Virtualization Host has a status value that is not valid to perform the operation.
	105	Parent Virtualization Host object was locked/busy and the request timed out.

HTTP error status code	Reason code	Description
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Job status and reason codes

HTTP error status code	Reason code	Description
200 (OK)	N/A	Mount completed successfully
500 (Server Error)	100	Mount failed

### Usage notes

- The mount operation may take some time as the zManager-provided ISO may need to be internally uploaded to the hosting virtualization host from other elements in the management hierarchy.
- Upload and mount of user-provided ISOs is supported by the **Mount Virtual Media Image** operation available in version 1.2 of the Web Services API.

### Example HTTP interaction

---

```
POST /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/operations/mount-virtual-media HTTP/1.1
x-api-session: 29chofb3vqxpijctnd2kf5fiwnf118z62mq3yomtqjyz8bx2gn
content-type: application/json
content-length: 15
{
  "iso": "gpmp"
}
```

---

Figure 172. Mount Virtual Media: Request

---

```
202 Accepted
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 15 Nov 2011 17:10:50 GMT
content-type: application/json;charset=UTF-8
content-length: 60
{
  "job-uri": "/api/jobs/c4e59922-0fac-11e1-bfcd-00215e69dea0"
}
```

---

Figure 173. Mount Virtual Media: Response

## Mount Virtual Media Image

The **Mount Virtual Media Image** operation starts a synchronous operation to mount a user-provided ISO image to the identified PowerVM or x Hyp virtual server. This operation is not supported for PR/SM and z/VM virtual servers. Virtual media is a virtual DVD drive that can be attached to a virtual server. Because virtual media is allocated from IBM blade storage instead of on a real media, it may be faster to use than a physical real device. The contents of the ISO image is specified as binary data in the body of the POST request.

## HTTP method and URI

POST /api/virtual-servers/{*virtual-server-id*}/operations/mount-virtual-media-image

In this request, the URI variable *{virtual-server-id}* is the object ID of the Virtual Server object.

### Query parameters:

Name	Type	Rqd/Opt	Description
image-name	String	Optional	If specified, this will be used as the displayable name and returned as the mounted-media-name in the virtual server properties. If omitted, the service will generate a name for this use.

## Request body contents

The request body is the binary contents of an ISO image file. A MIME media type of `application/octet-stream` should be specified as the content-type on the request.

## Description

This operation mounts the provided ISO image content to the identified virtual server.

- | If the API user does not have action permission for the **Virtual Server Details** task, a 403 (Forbidden) status code is returned. A 404 (Not Found) status code is returned if the object-id *{virtual-server-id}* does not identify a Virtual Server object to which the API user has object-access permission.

A mount operation cannot be performed when a virtual server or its virtualization host is busied by another operation or when the virtualization host status is not **"operating"**. Under these conditions a 409 (Conflict) status code is returned.

If the mount completes successfully a 204 (No Content) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server.
- | • Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	100	The operation does not support a virtual server of the given type.
403 (Forbidden)	1	API user does not have action permission to the <b>Virtual Server Details</b> task.
404 (Not Found)	1	A virtual server with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it.

HTTP error status code	Reason code	Description
409 (Conflict)	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
	101	Parent Virtualization Host has a status value that is not valid to perform the operation.
	105	Parent Virtualization Host object was locked/busy and the request timed out.
500 (Server Error)	100	Mount failed.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Usage notes

- Because the ISO image is transmitted synchronously as part of the request body of this operation, the time to execute this operation can be expected to be, at least in part, proportional to the size of the ISO image being uploaded.
- Mount of preloaded ISO images is supported by the **Mount Virtual Media** operation of the Web Services API.

## Unmount Virtual Media

The **Unmount Virtual Media** operation unmounts the currently mounted ISO from the identified PowerVM or x Hyp virtual server. This operation is not supported for PR/SM and z/VM virtual servers.

### HTTP method and URI

**POST** `/api/virtual-servers/{virtual-server-id}/operations/unmount-virtual-media`

In this request, the URI variable *{virtual-server-id}* is the object ID of the Virtual Server object.

### Request body contents

An optional request body can be specified as a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
force	Boolean	Optional	If true, a forced unmount is requested. Otherwise, a normal unmount is requested. Default is false.

### Description

This operation unmounts the currently mounted ISO from the identified virtual server. Note that there are not separate unmount operations associated with the distinct **Mount Virtual Media** and **Mount Virtual Media Image** operations; this single **Unmount Virtual Media** operation may be used regardless of the mount operation used.

- 1 If the API user does not have action/task permission for the **Virtual Server Details** task, a 403 (Forbidden) status code is returned. A 404 (Not Found) status code is returned if the object-id *{virtual-server-id}* does not identify a Virtual Server object to which the API user has object-access permission.



If a request body is provided with the **force** request field specified as true, a request is issued to the virtualization host to attempt to force the unmount even if the media is locked by a guest OS. Otherwise, a normal unmount operation is performed.

An unmount operation cannot be performed when a virtual server or its virtualization host is busied by another operation or when the virtualization host **status** is not **"operating"**. Under these conditions a 409 (Conflict) status code is returned.

The virtual media is then unmounted and a 204 (No Content) status code is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server.
- Action/task permission to the **Virtual Server Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	100	The operation does not support a virtual server of the given type.
403 (Forbidden)	1	API user does not have action permission to the <b>Virtual Server Details</b> task.
404 (Not Found)	1	A virtual server with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it.
409 (Conflict)	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
	101	Parent Virtualization Host has a status value that is not valid to perform the operation.
	105	Parent Virtualization Host object was locked/busy and the request timed out.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/operations/unmount-virtual-media HTTP/1.1
x-api-session: 29chofb3vqxpijctnd2kf5fiwnf118z62mq3yomtqjyz8bx2gn
```

---

Figure 174. Unmount Virtual Media: Request

---

```
204 No Content
date: Tue, 15 Nov 2011 17:10:55 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

```
<No response body>
```

---

Figure 175. Unmount Virtual Media: Response

## Migrate Virtual Server

The **Migrate Virtual Server** operation moves the identified PowerVM or x Hyp virtual server to the target virtualization host. This operation is not supported for PR/SM and z/VM virtual servers.

### HTTP method and URI

**POST** /api/virtual-servers/{*virtual-server-id*}/operations/migrate

In this request, the URI variable {*virtual-server-id*} is the object ID of the Virtual Server object.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
virt-host-id	String	Object ID of the target Virtualization Host object.

### Response body contents

On successful completion, a response body is provided such that the caller may further target the migrated virtual server.

Field name	Type	Description
object-uri	String/URI	The <b>object-uri</b> property of the migrated Virtual Server object on the new virtualization host.

## Description

This operation moves the identified virtual server to a new virtualization host.

Virtual server migration is a composite action consisting of the steps of creating a new virtual server on the new virtualization host, configuring the new virtual server to correspond to the identified (original) virtual server, and then deleting the identified virtual server. Upon completion, this operation returns the **object-uri** of the new virtual server in the response body. The response also includes a **Location** header that provides this URI.

The steps of the composite migration action may be observable through inventory and property change notifications.

If the API user does not have action permission for the Migrate Virtual Server task, a 403 (Forbidden) status code is returned. A 404 (Not Found) status code is returned if the object-id {*virtual-server-id*} does not identify a Virtual Server object to which the API user has object-access permission. A 400 (Bad

Request) status code is returned if the request body fails to validate (e.g. the target virtualization host is of a different type than the current virtualization host).

A migrate operation cannot be performed when the virtual server, its storage-resources, or virtual-networks are busied by another operation. Migrate is also not possible when the virtual server's **status** is something other than **"not-operating"** or **"not-activated"** or when its current virtualization host or the target virtualization host **status** is **"not-communicating"** or **"status-check"**. Migration is also not possible when the virtual server has media mounted. Under any of these conditions a 409 (Conflict) status code is returned.

Virtual server migration is then attempted. Migration may fail if the user does not have object-access to the storage resources used by the virtual server's virtual disks. Migration may also fail if the virtual server has virtual disks and the target virtualization host does not have virtualization host storage resources for them. If these or other failure occurs, migration fails, the original virtual server is restored, and a 503 (Service Unavailable) status code is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server
- Object access permission to the virtual server's virtualization host's hosting-environment
- Object access permission to the target virtualization host's hosting-environment
- Object access permission to the storage resources used by the virtual server's virtual disks
- Action/task permission to the **Migrate Virtual Server** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in "Response body contents" on page 354.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	100	The operation does not support a virtual server of the given type.
	104	Target virtualization host is invalid because it is not the same type as the current virtualization host.
403 (Forbidden)	1	API user does not have action permission to the <b>Migrate Virtual Server</b> task.
404 (Not Found)	1	A Virtual Server object with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it or its virtualization host's hosting-environment.
	1	A virtualization host with object-id <i>{virt-host-id}</i> does not exist on HMC or API user does not have object-access permission for its hosting-environment.

HTTP error status code	Reason code	Description
409 (Conflict)	0	Migration is not possible for a reason other than those indicated by the other 409 (Conflict) reason codes.
	1	Virtual server status is not valid to perform the operation (must be either <b>"not-operating"</b> or <b>"not-activated"</b> ).
	2	Virtual Server object with ID <i>{virtual-server-id}</i> was busy and request timed out.
	101	Parent Virtualization Host has a status value that is not valid to perform the operation.
	103	Virtual server is an invalid target because it has virtual media is mounted.
	105	The target virtualization host already has a virtual server with the same name as the virtual server to be migrated.
	106	One or more of the virtual server's storage resources was busy and the request timed out.
	107	One or more of the virtual server's Virtual Network objects was busy and the request timed out.
	108	One or more storage resources used by the virtual server are not available on the target virtualization host.
503 (Service Unavailable)	0	Migration failed.
	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Initiate Virtual Server Dump

The **Initiate Virtual Server Dump** operation disruptively stops the identified PowerVM or x Hyp virtual server and sends an operating system-specific command to begin a memory dump of the operating system running on the virtual server. This operation is not supported for PR/SM or z/VM virtual servers.

### HTTP method and URI

**POST** `/api/virtual-servers/{virtual-server-id}/operations/initiate-dump`

In this request, the URI variable *{virtual-server-id}* is the object ID of the Virtual Server object.

### Description

The Initiate Virtual Server Dump operation disruptively stops the virtual server and initiates a dump.

Virtual server **status** must not be **"not-operating"**.

This operation disruptively stops the identified virtual server and sends an virtual-server-type-specific signal to begin a memory dump of the operating system running on the virtual server.

If the API user does not have action permission for the Initiate Virtual Server Dump task, a 403 (Forbidden) status code is returned. A 404 (Not Found) status code is returned if the object-id *{virtual-server-id}* does not identify a Virtual Server object to which the API user has object-access permission. A 400 (Bad Request) status code is returned if the request body fails to validate or if the operation is not supported for the given type of virtual server.

A 409 (Conflict) status code is returned if the virtual server status is **"not-operating"**.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the virtual server
- Action/task permission to the **Initiate Virtual Server Dump** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	100	The operation does not support a virtual server of the given type.
403 (Forbidden)	1	API user does not have action permission to the <b>Initiate Virtual Server Dump</b> task.
404 (Not Found)	1	A Virtual Server object with object-id <i>{virtual-server-id}</i> does not exist on HMC or API user does not have object-access permission for it.
409 (Conflict)	1	Virtual server status is not valid to perform the operation (status is <b>"not-operating"</b> ).
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/operations/initiate-dump HTTP/1.1
x-api-session: 6cq5qn64f6zv388z20hu6nzipx1ukaqr14ekri6sgezaxcwgztz0
```

---

Figure 176. Initiate Virtual Server Dump: Request

---

```
204 No Content
date: Wed, 07 Dec 2011 05:07:39 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

<No response body>

---

Figure 177. Initiate Virtual Server Dump: Response

## Inventory service data

Information about the virtual servers managed by the HMC can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for Virtual Server objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by inventory class, implicitly via a containing category, or by default) that objects of the various virtual server type-specific inventory classes are to be included. An entry for a particular Virtual Server is included only if the API user has object-access permission to that object and the applicable type-specific inventory class has been specified, as described in the following table:

Inventory class	Includes virtual servers with "type" value
power-vm-virtual-server power-vm-virtual-server-common	power-vm
prsm-virtual-server prsm-virtual-server-common	prsm
x-hyp-virtual-server x-hyp-virtual-server-common	x-hyp
zvm-virtual-server zvm-virtual-server-common	zvm

For each Virtual Server object to be included, the inventory response array includes entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the **Get Virtual Server Properties** operation. That is, the data provided is the same as would be provided if a Get Virtual Server Properties operation were requested targeting this object. For inventory class names that end with **-common** (e.g. **power-vm-virtual-server-common**), the data is the same as would be provided for a Get Virtual Server Properties operation with the **properties=common** query parameter specified. For the inventory class names that do not end in **-common** (e.g. **power-vm-virtual-server**), the results are the same as would be provided for a Get Virtual Server Properties operation with no **properties** query parameter specified.

## Sample inventory data

---

```

{
  "acceptable-status": [
    "operating"
  ],
  "auto-start": false,
  "boot-mode": "normal",
  "boot-network-adapter-client-ip": null,
  "boot-network-adapter-gateway-ip": null,
  "boot-network-adapter-server-ip": null,
  "boot-network-adapter-subnet-ip": null,
  "boot-sequence": [
    "virtual-disk"
  ],
  "class": "virtual-server",
  "cpu-perf-mgmt-enabled": true,
  "description": "Order processing for the Shimmer floor wax/dessert topping product",
  "dlpar-active": false,
  "dlpar-enabled": false,
  "gpmp-status": "not-operating",
  "gpmp-support-enabled": false,
  "gpmp-version": "unavailable",
  "has-unacceptable-status": true,
  "hostname": null,
  "initial-dedicated-processors": null,
  "initial-memory": 1024,
  "initial-processing-units": 0.10000000000000001,
  "initial-virtual-processors": 1,
  "is-locked": false,
  "keylock": "normal",
  "mac-prefix": {
    "mac-address": "02:ee:6b:6c:70:00",
    "prefix-length": 40
  },
  "maximum-dedicated-processors": null,
  "maximum-memory": 1024,
  "maximum-processing-units": 7.0,
  "maximum-virtual-processors": 7,
  "minimum-dedicated-processors": null,
  "minimum-memory": 1024,
  "minimum-processing-units": 0.10000000000000001,
  "minimum-virtual-processors": 1,
  "mounted-media-name": null,
  "name": "Shimmer orders server",
  "network-adapters": [
    {

```

---

Figure 178. Virtual Server object: Sample inventory data for a virtual server of type "power-vm" (Part 1)

---

```
    "element-id": "0000",
    "element-uri": "/api/virtual-servers/42eecf00-7b47-11e0-bc9e-001f163803de/
      network-adapters/0000",
    "network-uri": "/api/virtual-networks/9b9fe4f8-75b2-11e0-9219-0010184c8026"
  },
  {
    "element-id": "0010",
    "element-uri": "/api/virtual-servers/42eecf00-7b47-11e0-bc9e-001f163803de/
      network-adapters/0010",
    "network-uri": null
  }
],
"object-id": "119338a2-4081-11e0-9e7c-f0def10bff8d",
"object-uri": "/api/virtual-servers/119338a2-4081-11e0-9e7c-f0def10bff8d",
"parent": "/api/virtualization-hosts/baab1cd2-2990-11e0-8d5b-001f163803de",
"processing-mode": "shared",
"status": "not-operating",
"type": "power-vm",
"virtual-disks": [
  {
    "backing-virtualization-host-storage-resource": "/api/virtualization-hosts/
      baab1cd2-2990-11e0-8d5b-001f163803de/virtualization-host-storage-resources/
      abfce790-4080-11e0-8486-f0def10bff8d",
    "description": "",
    "element-id": "11b5b0a8-4081-11e0-8486-f0def10bff8d",
    "element-uri": "/api/virtual-servers/119338a2-4081-11e0-9e7c-f0def10bff8d/virtual-disks/
11b5b0a8-4081-11e0-8486-f0def10bff8d",
    "name": "Superman5",
    "owner": "/api/virtual-servers/119338a2-4081-11e0-9e7c-f0def10bff8d",
    "size": 5242880,
    "type": "fullpack"
  }
],
"workloads": [
  "/api/workload-resource-groups/a4019e06-7685-11e0-8fca-0010184c8026"
]
}
```

---

Figure 179. Virtual Server object: Sample inventory data for a virtual server of type "power-vm" (Part 2)



```

"acceptable-status": [
  "operating"
],
"associated-logical-partition": "/api/logical-partitions/0b239aa3-fea1-32e0-a38f-c632e7ee3b0c",
"class": "virtual-server",
"cpu-perf-mgmt-enabled": false,
"description": "",
"gpmp-status": "not-operating",
"gpmp-support-enabled": false,
"gpmp-version": "unavailable",
"has-unacceptable-status": true,
"is-locked": false,
"name": "VMALT2 ",
"network-adapters": [
  {
    "chpid": "F0",
    "css": "0",
    "element-id": "OSX 0.F0",
    "element-uri": "/api/virtual-servers/d44575cc-40ea-11e0-9814-001f163803de/network-adapters/OSX%200.F0",
    "name": "OSX 0.F0",
    "type": "osx"
  },
  {
    "chpid": "A0",
    "css": "0",
    "element-id": "OSX 0.A0",
    "element-uri": "/api/virtual-servers/d44575cc-40ea-11e0-9814-001f163803de/network-adapters/OSX%200.A0",
    "name": "OSX 0.A0",
    "network-uris": [
      "/api/virtual-networks/a9b6f8ce-771f-11e0-b1da-0010184c8026"
    ],
    "type": "osx"
  }
],
"object-id": "d44575cc-40ea-11e0-9814-001f163803de",
"object-uri": "/api/virtual-servers/d44575cc-40ea-11e0-9814-001f163803de",
"parent": "/api/virtualization-hosts/bab76208-2990-11e0-8d5b-001f163803de",
"status": "operating",
"type": "prsm",
"workloads": [
  "/api/workload-resource-groups/a4019e06-7685-11e0-8fca-0010184c8026"
]
}

```

Figure 180. Virtual Server object: Sample inventory data for a virtual server of type "prsm"

---

```
{
  "acceptable-status": [
    "operating"
  ],
  "auto-start": false,
  "boot-sequence": [
    "virtual-media"
  ],
  "class": "virtual-server",
  "description": "",
  "gpmp-status": "not-operating",
  "gpmp-support-enabled": false,
  "gpmp-version": "unavailable",
  "has-unacceptable-status": true,
  "hostname": null,
  "initial-memory": 4096,
  "initial-virtual-processors": 4,
  "is-locked": false,
  "mounted-media-name": "ubuntu-11.04-server-i386.iso",
  "name": "XVS1",
  "network-adapters": [],
  "object-id": "a4588932-8648-11e0-bbc1-f0def10bff8d",
  "object-uri": "/api/virtual-servers/a4588932-8648-11e0-bbc1-f0def10bff8d",
  "parent": "/api/virtualization-hosts/931b25d6-82e1-11e0-b9e4-f0def10bff8d",
  "status": "stopping",
  "type": "x-hyp",
  "virtual-disks": [],
  "workloads": [
    "/api/workload-resource-groups/a4019e06-7685-11e0-8fca-0010184c8026"
  ]
}
```

---

Figure 181. Virtual Server object: Sample inventory data for a virtual server of type "x-hyp"

---

## Chapter 11. Storage Management

zManager provides a common interface across the different Virtualization Host types and storage types that it supports. It allows a system administrator to create virtualized storage resources and attach them to virtual servers. The basic flow for all supported types of Virtualization Hosts and storage resources is as follows:

The server administrator defines his requirements to the storage administrator (e.g., the number of storage resources, their type and size information).

The server administrator uses zManager storage management interfaces to export the virtualization-host-specific information which is required to allow the Storage Area Network (SAN) administrator to setup the SAN accordingly. This information consists primarily of the Host WWPN List.

When the SAN administrator has finished configuring storage resources, the server administrator can add these new storage resources to the ensemble for management by performing one or more of the following actions:

- Triggering discovery of newly detected storage resources for a Virtualization Host
- Compiling a Storage Access List (a file with information about the storage resources, such as unique name and addressing information) and importing this list into zManager
- Manually adding each storage resource.

zManager offers interfaces to work with all ensemble-managed storage resources. For example, there are interfaces to:

- List the various storage-related entities (storage resources, Virtualization Host storage resources, Virtualization Host storage groups and virtual disks)
- List details of the various storage-related entities
- Identify storage resources that are to be managed by zManager
- Grant Virtualization Hosts access to storage resources
- Assign storage resources to a Virtualization Host storage group
- Assign storage resources to virtual servers by creating virtual disks on them.

---

### Terms

#### Host World Wide Port Name (WWPN) List

The Host WWPN List consists of a list of WWPNs of the Fibre-Channel host ports of each virtualization host. It can be exported for one or multiple virtualization hosts through a zManager function. The WWPN list is useful when the system administrator requires additional storage resources to be configured by the storage administrator. The storage administrator must enter these WWPNs into the SAN switches and storage controllers in order to allow these specific WWPNs to access the storage controllers / Logical Units.

#### Storage Access List (SAL)

The Storage Access List is provided by a storage administrator to a system administrator after configuring storage resources (e.g., FCP Logical Units). The Storage Access List consists of a number of host port WWPNs and entries for a configured storage resource with its properties, such as addressing information, or device type information, in the form of a Comma-Separated Values (CSV) file. Importing a Storage Access List offers a convenient way for letting zManager know which hypervisor has access to which storage resource, avoiding the cumbersome and error-prone process of adding storage resources (and their associated properties) manually.

**Storage Resource**

An addressable storage entity, allowing a virtualization host to write data to and read data from. A storage resource may be one of the following: a SCSI Logical Unit, attached via FCP, a file, a Volume attached via ESCON/FICON.

**Virtualization Host Storage Resource**

The representation of a storage resource from the perspective of a virtualization host. It is a storage resource to which the virtualization host has access.

**Virtualization Host Storage Group**

Representation of a z/VM Storage Group in zManager. It consists of homogeneous storage resources to which a z/VM virtualization host has access.

**Virtual Disk**

Virtual storage space provided by a virtualization host to a guest virtual server. A Virtual Disk is based upon a storage resource, but may be further virtualized by a virtualization host.

**Object model overview**

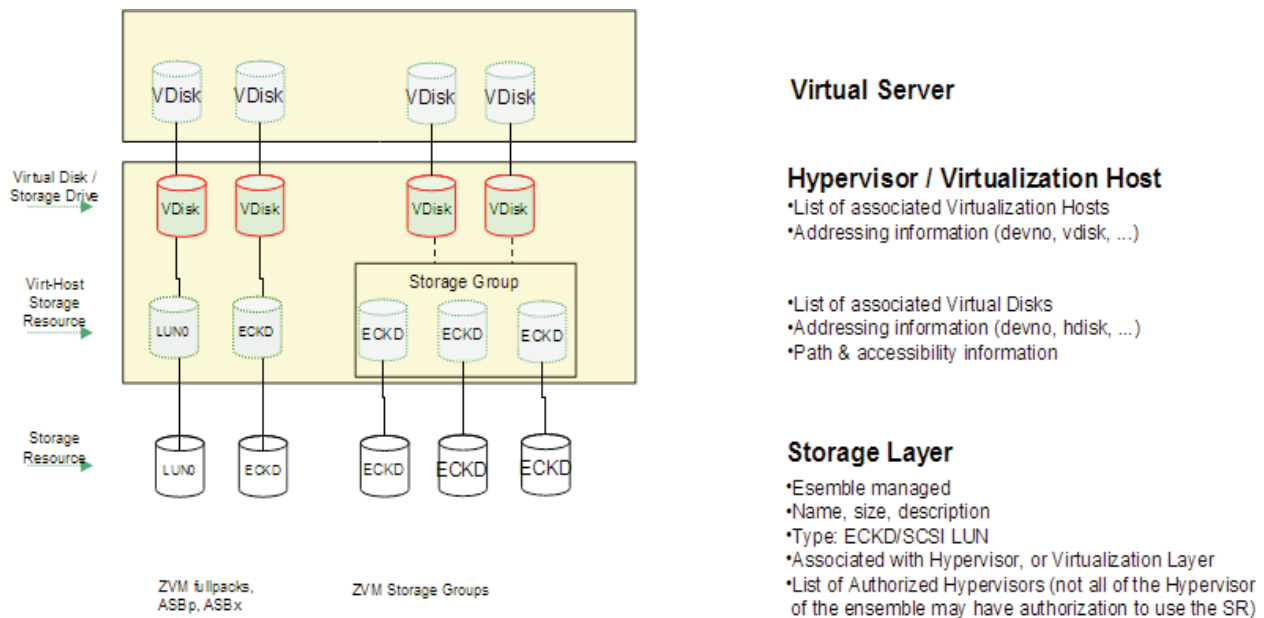


Figure 182. Object model

**Storage management operations summary**

The following tables provide an overview of the operations provided. The tables are organized according to the scope of the operations listed.

Table 81. Storage management: ensemble-level storage operations

Operation name	HTTP method and URI path
"List Storage Resources" on page 367	GET /api/ensembles/{ensemble-id}/storage-resources
"Get Storage Resource Properties" on page 370	GET /api/storage-resources/{storage-resource-id}
"Create Storage Resource" on page 371	POST /api/ensembles/{ensemble-id}/storage-resources

Table 81. Storage management: ensemble-level storage operations (continued)

Operation name	HTTP method and URI path
“Update Storage Resource Properties” on page 373	POST /api/storage-resources/{storage-resource-id}
“Delete Storage Resource” on page 375	DELETE /api/storage-resources/{storage-resource-id}
“Export World Wide Port Names List” on page 377	POST /api/ensembles/{ensemble-id}/operations/export-port-names
“Import Storage Access List” on page 379	POST /api/ensembles/{ensemble-id}/operations/import-storage-access-list
“List Virtualization Host Storage Resources of a Storage Resource” on page 381	POST /api/storage-resources/{storage-resource-id}/operations/list-virtualization-host-storage-resources

Table 82. Storage management: virtualization host storage operations

Operation name	HTTP method and URI path
“List Virtualization Host HBA Ports” on page 386	GET /api/virtualization-hosts/{virt-host-id}/hba-ports
“List Virtualization Host Storage Resources” on page 388	GET /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-resources
“Get Virtualization Host Storage Resource Properties” on page 391	GET /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-resources/{virt-host-storage-resource-id}
“Create Virtualization Host Storage Resource” on page 395	POST /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-resources
“Delete Virtualization Host Storage Resource” on page 398	DELETE /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-resources/{virt-host-storage-resource-id}
“Add Virtualization Host Storage Resource Paths” on page 400	POST /api/virtualization-hosts/{virt-host-id}/operations/add-paths
“Remove Virtualization Host Storage Resource Paths” on page 403	POST /api/virtualization-hosts/{virt-host-id}/operations/remove-paths
“Discover Virtualization Host Storage Resources” on page 406	POST /api/virtualization-hosts/{virt-host-id}/operations/discover-virtualization-host-storage-resources
“List Virtual Disks of a Virtualization Host Storage Resource” on page 408	POST /api/virtualization-hosts/{virt-host-id}/operations/list-virtual-disks-of-virtualization-host-storage-resource

Table 83. Storage management: storage group operations

Operation name	HTTP method and URI path
“List Virtualization Host Storage Groups” on page 412	GET /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-groups
“Get Virtualization Host Storage Group Properties” on page 414	GET /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-groups/{virt-host-storage-group-id}
“List Virtualization Host Storage Resources in a Virtualization Host Storage Group” on page 416	POST /api/virtualization-hosts/{virt-host-id}/operations/list-virtualization-host-storage-resources-in-group
“Add Virtualization Host Storage Resource to Virtualization Host Storage Group” on page 418	POST /api/virtualization-hosts/{virt-host-id}/operations/add-virtualization-host-storage-resource-to-group

Table 83. Storage management: storage group operations (continued)

Operation name	HTTP method and URI path
“Remove Virtualization Host Storage Resource from Virtualization Host Storage Group” on page 420	POST /api/virtualization-hosts/{virt-host-id}/operations/remove-virtualization-host-storage-resource-from-group

Table 84. Storage management: URI variables

Variable	Description
{ensemble-id}	Object ID of an ensemble object
{virt-host-id}	Object ID of a Virtualization Host object
{virtual-server-id}	Object ID of a Virtual Server object
{storage-resource-id}	Object ID of a Storage Resource object
{virt-host-storage-resource-id}	Element ID of a Virtualization Host Storage Resource object
{virt-host-storage-group-id}	Element ID of a Virtualization Host Storage Group object

**Note:** Although virtual disk operations are also storage related, they have a closer affinity to virtualization management. Thus, they are included within the specification for the “Virtual Server Object” on page 260 in Chapter 10, “Virtualization management,” on page 209.

## Storage Resource object

- | A Storage Resource object represents a single physical storage resource available to one or more z
- | Systems Virtualization Hosts in an ensemble.

## Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, but does not provide the operational-status-related properties defined in that schema because it does not maintain the concept of an operational status.

The following type-specific specializations apply to the other base managed object properties:

Table 85. Storage Resource object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/ URI	The canonical URI path for a Storage Resource object is of the form /api/storage-resources/{storage-resource-id} where {storage-resource-id} is the value of the object-id property of the Storage Resource object.
<b>parent</b>	—	String/ URI	The parent object of a Storage Resource object is an ensemble object.
<b>class</b>	—	String	The class of a Storage Resource object is <b>"storage-resource"</b> .
<b>description</b>	(w)(pc)	String (0-256)	The optional user-supplied description for the storage resource. This is the description that will be displayed on the user interface. It must consist only of alphanumeric characters, spaces and the following special characters: “._-”.
<b>name</b>	(w)(pc)	String (1-64)	The user-supplied name of the storage resource. This is the name that will be displayed on the user interface. It must consist only of alphanumeric characters, spaces and the following special characters: “._-”, and it must begin with an alphabetic character. This name must be unique within the Ensemble.

## Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties:

Table 86. Storage Resource object: class specific properties

Name	Qualifier	Type	Description
<b>type</b>	—	String Enum	The type of the storage resource. Values: <ul style="list-style-type: none"> <li>"<b>eckd</b>" – An Extended Count Key Data storage resource</li> <li>"<b>fcp</b>" – A Fibre-channel attached storage resource</li> <li>"<b>zvm-fcp</b>" – A Fibre-channel attached storage resource for use by z/VM virtualization hosts</li> </ul>
<b>size</b>	(w)(pc)	Long	The <b>size</b> of the storage resource. The units for this property are specified by the <b>allocation-units</b> property.
<b>allocation-units</b>	(w)	String Enum	The units for the size property. Values: <ul style="list-style-type: none"> <li>"<b>bytes</b>"– used only for storage resources whose <b>type</b> property is "<b>fcp</b>" or "<b>zvm-fcp</b>".</li> <li>"<b>cylinders</b>" – used only for storage resources whose <b>type</b> property is "<b>eckd</b>".</li> <li>"<b>unknown</b>"</li> </ul>
<b>allocation-status</b>	(pc)	String Enum	The status of the storage resource in terms of its current allocation. A storage resource is considered to be allocated for use if there is a virtual disk backed by this storage resource or the storage resource backs a hypervisor storage resource that is a member of a virtualization host storage group. <p>Values:</p> <ul style="list-style-type: none"> <li>"<b>free</b>" – the storage resource is not currently allocated for use.</li> <li>"<b>used</b>" – the storage resource is currently allocated for use.</li> </ul>
<b>unique-device-id</b>	(pc)	String	The unique device identifier assigned to this storage resource. This is a worldwide unique identifier based on information about the device. zManager creates a unique device identifier for each storage resource whose <b>type</b> property is " <b>fcp</b> ". A unique device identifier is not created for storage resources with a <b>type</b> property of " <b>eckd</b> " or " <b>zvm-fcp</b> "; for such storage resources, the value of this property is always null. <p>Note that even for storage resources with a <b>type</b> property of "<b>fcp</b>", the value of the property may be null or an empty string. This is typically the case if zManager has not yet accessed the storage resource and thus the unique device identifier is not yet known.</p>

## Operations

If a storage resource operation accesses a z/VM Virtualization Host and encounters an error while communicating with the Virtualization Host via SMAPI, the response body is as described in “SMAPI Error Response Body” on page 257.

## List Storage Resources

The **List Storage Resources** operation lists the storage resources in the ensemble.

### HTTP method and URI

**GET** /api/ensembles/{*ensemble-id*}/storage-resources

In this request, the URI variable {*ensemble-id*} is the object ID of the ensemble object.

**Query parameters:**

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property

**Response body contents**

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
storage-resources	Array of objects	Array of storage-resource-basic-info objects, described in the next table. If no storage resources are to be returned, an empty array is provided.

Each storage-resource-basic-info object contains the following fields:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the Storage Resource object
name	String	The <b>name</b> property of the Storage Resource object
type	String Enum	The <b>type</b> property of the Storage Resource object

**Description**

The **List Storage Resources** operation lists the storage resources in the ensemble. The object URI and other basic properties are provided for each storage resource.

If the **name** query parameter is specified, the returned list is limited to those storage resources that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

A set of basic properties is returned for each storage resource. See the storage-resource-basic-info object definition.

A storage resource is included in the list only if the API user has object-access permission for that object. If the ensemble contains a storage resource to which the API user does not have permission, that object is omitted from the list, but no error status code results. Note that this could result in an empty list.

The URI path must designate an existing ensemble object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Manage Storage Resources** task; otherwise, status code 403 (Forbidden) is returned.

If there are no storage resources in the ensemble, an empty list is provided and the operation completes successfully.

**Authorization requirements**

This operation has the following authorization requirements:

- Object-access permission to the ensemble object specified in the request URI
- Object-access permission to the Storage Resource objects passed in the response body
- Action/task permission to the **Manage Storage Resources** task.



## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 368.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{ensemble-id}</i> does not designate an existing ensemble object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026/storage-resources HTTP/1.1
x-api-session: 1rmnds0imna61i3110eu7drk7jsec93mvc1fbuqdb7xspk2fm5
```

---

Figure 183. List Storage Resources: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 04 Aug 2011 13:30:11 GMT
content-type: application/json;charset=UTF-8
content-length: 524
{
  "storage-resources": [
    {
      "name": "erictest",
      "object-uri": "/api/storage-resources/a23a998c-9693-11e0-aace-00215e69e3f5",
      "type": "fcp"
    },
    {
      "name": "Test2_ECKD",
      "object-uri": "/api/storage-resources/8d7da556-6598-11e0-8946-00215e69e3f5",
      "type": "eckd"
    },
    {
      "name": "FCP5618",
      "object-uri": "/api/storage-resources/b0be5b6e-5ada-11e0-b462-00215e69e3f5",
      "type": "zvm-fcp"
    }
  ]
}
```

---

Figure 184. List Storage Resources: Response

## Get Storage Resource Properties

The **Get Storage Resource Properties** operation retrieves the properties of a single Storage Resource object that is designated by its object ID.

### HTTP method and URI

**GET** /api/storage-resources/{*storage-resource-id*}

In this request, the URI variable {*storage-resource-id*} is the object ID of the Storage Resource object for which properties are to be obtained.

### Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Storage Resource object as defined in the “Data model” on page 366. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

The **Get Storage Resource Properties** operation returns the current properties for the Storage Resource object specified by {*storage-resource-id*}.

On successful execution, all of the current properties as defined in “Data model” on page 366 for the Storage Resource object are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing Storage Resource object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Manage Storage Resources** task; otherwise, status code 403 (Forbidden) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the Storage Resource object specified in the request URI
- Action/task permission to the **Manage Storage Resources** task.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID { <i>storage-resource-id</i> } does not designate an existing Storage Resource object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/storage-resources/6967f806-2023-11e1-9c1e-0010184c8334 HTTP/1.1
x-api-session: 1tcd8u2o682d6diyft8q9aafhfx125d8m87150yl6osfcje7k
```

---

Figure 185. Get Storage Resource Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 06 Dec 2011 16:00:26 GMT
content-type: application/json;charset=UTF-8
content-length: 402
{
  "allocation-status": "free",
  "allocation-units": "bytes",
  "class": "storage-resource",
  "description": "SS Ensemble volume V001",
  "is-locked": false,
  "name": "SS-V0001",
  "object-id": "6967f806-2023-11e1-9c1e-0010184c8334",
  "object-uri": "/api/storage-resources/6967f806-2023-11e1-9c1e-0010184c8334",
  "parent": "/api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334",
  "size": 8589934592,
  "type": "fcp",
  "unique-device-id": null
}
```

---

Figure 186. Get Storage Resource Properties: Response

## Create Storage Resource

The **Create Storage Resource** operation adds a storage resource to the specified ensemble.

### HTTP method and URI

**POST** /api/ensembles/{ensemble-id}/storage-resources

In this request, the URI variable {ensemble-id} is the object ID of the ensemble to which the new storage resource is to be added.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The <b>name</b> property of the new storage resource
description	String	Optional	The <b>description</b> property of the new storage resource
size	Long	Required	The <b>size</b> property of the storage resource
type	String Enum	Required	The <b>type</b> property of the storage resource

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the Storage Resource object, in the form <code>/api/storage-resources/{storage-resource-id}</code>

## Description

The **Create Storage Resource** operation identifies a new storage resource to be added to the ensemble specified by the `{ensemble-id}` portion of the request URI. Once added to the ensemble, the storage resource can be managed using the various storage-related zManager functions.

On successful execution, the **object-uri** field of the response body and the **Location** response header identify the new storage resource.

If this operation changes the value of any property for which property-change notifications are due, those notifications are issued asynchronously to this operation. Upon success, an Inventory Change notification is issued.

The URI path must designate an existing ensemble object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must also have action access permission to the **Add Storage Resource** task; otherwise, status code 403 (Forbidden) is returned.

If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object specified in the request URI
- Action/task permission to the **Add Storage Resource** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and both the response body and the **Location** response header contain the URI of the newly created object.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	8	The name specified for the new storage resource is not unique within the ensemble.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <code>{ensemble-id}</code> does not designate an existing ensemble object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

The **Create Storage Resource** operation performs a portion of the function provided by the Add Storage Resource subtask of the **Manage Storage Resources** task. The Add Storage Resource subtask performs some or all of the functions provided by the following Web Services APIs:

- **Create Storage Resource**  
Creates a Storage Resource object in zManager, representing storage resources (FCP LUNs, or ECKD™ volumes). Each Storage Resource object has a unique name. zManager reports additional storage resource specific information, such as a unique-device-id, after paths have been added and zManager was able to access the storage resource.
- **Create Virtualization Host Storage Resource**  
Creates an object in zManager, representing the virtualization host's view on the storage resource.
- **Add Virtualization Host Storage Resource Path**  
Adds a path between a virtualization host and a storage resource.

**Note:** a path to an FCP storage resource is identified by a host-port-wwpn, target-port-wwpn, and lun. See Table 88 on page 385. A path to an ECKD resource is identified by the device number of the ECKD volume. See Table 89 on page 386.

## Example HTTP interaction

---

```
POST /api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334/storage-resources HTTP/1.1
x-api-session: 1tcd8u2o682d6diyft8q9aafhfx125d8m87150y16osfcje7k
content-type: application/json
content-length: 109
{
  "description": "New Storage Resource",
  "name": "SS-New-Storage-Resource",
  "size": 8589934592,
  "type": "fcp"
}
```

---

Figure 187. Create Storage Resource: Request

---

```
201 Created
server: zSeries management console API web server / 1.0
location: /api/storage-resources/6967f806-2023-11e1-9c1e-0010184c8334
cache-control: no-cache
date: Tue, 06 Dec 2011 16:00:26 GMT
content-type: application/json;charset=UTF-8
content-length: 76
{
  "object-uri": "/api/storage-resources/6967f806-2023-11e1-9c1e-0010184c8334"
}
```

---

Figure 188. Create Storage Resource: Response

## Update Storage Resource Properties

The **Update Storage Resource Properties** operation updates one or more of the writeable properties of a Storage Resource object.

## HTTP method and URI

**POST** /api/storage-resources/{*storage-resource-id*}

In this request, the URI variable *{storage-resource-id}* is the object ID of the Storage Resource object for which properties are to be updated.

## Request body contents

The request body is expected to contain a JSON object that provides the new values of any writeable property that is to be updated by this operation. Field names and data types in this JSON object are expected to match the corresponding property names and data types defined in the data model for this object type. The JSON object can and should omit fields for properties whose values are not to be changed by this operation.

## Description

The **Update Storage Resource Properties** operation updates writeable properties of the Storage Resource object specified by *{storage-resource-id}*.

The request body contains an object with one or more fields with field names that correspond to the names of properties for this object. On successful execution, the value of each corresponding property of the object is updated with the value provided by the input field, and status code 204 (No Content) is returned without supplying any response body. The request body does not need to specify a value for all writeable properties, but rather can and should contain fields for the properties to be updated. Object properties for which no input value is provided remain unchanged by this operation.

If the update changes the value of any property for which property-change notifications are due, those notifications are issued asynchronously to this operation.

The URI path must designate an existing Storage Resource object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Storage Resources Details** task; otherwise, status code 403 (Forbidden) is returned.

The request body is validated against the data model for this object type to ensure that it contains only writeable properties and the data types of those properties are as required. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the Storage Resource object specified in the request URI
- Action/task permission to the **Storage Resources Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	8	The new name specified for the new storage resource is not unique within the ensemble.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{storage-resource-id}</i> does not designate an existing Storage Resource object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
POST /api/storage-resources/6967f806-2023-11e1-9c1e-0010184c8334 HTTP/1.1
x-api-session: 1tcd8u2o682d6diyft8q9aafhfx125d8m87150y16osfcje7k
content-type: application/json
content-length: 62
{
  "description": "SS Ensemble volume V001",
  "name": "SS-V0001"
}
```

---

Figure 189. Update Storage Resource Properties: Request

---

```
204 No Content
date: Tue, 06 Dec 2011 16:00:26 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 190. Update Storage Resource Properties: Response

## Delete Storage Resource

The **Delete Storage Resource** operation deletes the specified storage resource from the ensemble.

### HTTP method and URI

```
DELETE /api/storage-resources/{storage-resource-id}
```

In this request, the URI variable *{storage-resource-id}* is the object ID of the Storage Resource object to be deleted.

### Description

The **Delete Storage Resource** operation removes a specified storage resource from the ensemble. The storage resource is identified by the *{storage-resource-id}* variable in the URI. There must be no virtualization host storage resources associated with the storage resource to be deleted.

Upon successfully removing the storage resource, HTTP status code 204 (No Content) is returned and no response body is provided. An inventory change event is issued asynchronously.

The URI path must designate an existing Storage Resource object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Remove Storage Resource** task; otherwise, status code 403 (Forbidden) is returned. If there are any virtualization host storage resources associated with the storage resource, status code 409 (Conflict) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the Storage Resource object specified in the request URI
- Action/task permission to the **Remove Storage Resource** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{storage-resource-id}</i> does not designate an existing Storage Resource object, or the API user does not have object access permission to it.
409 (Conflict)	143	The object cannot be deleted at this time. There is a virtualization host storage resource associated with the storage resource.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- Before a storage resource can be deleted from the ensemble, any virtualization host storage resources associated with it must first be deleted.

## Example HTTP interaction

---

```
DELETE /api/storage-resources/6967f806-2023-11e1-9c1e-0010184c8334 HTTP/1.1
x-api-session: 1tcd8u2o682d6diyft8q9aafhx125d8m87150y16osfcje7k
```

---

Figure 191. Delete Storage Resource: Request



---

```
204 No Content
date: Tue, 06 Dec 2011 16:00:26 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

```
<No response body>
```

---

Figure 192. Delete Storage Resource: Response

## Export World Wide Port Names List

The **Export World Wide Port Names List** operation exports the world wide port names (WWPNs) of the fibre channel host ports of the specified virtualization hosts.

### HTTP method and URI

**POST** /api/ensembles/{*ensemble-id*}/operations/export-port-names

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble that contains the virtualization hosts specified in the request body.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
virtualization-hosts	Array of String/URI	Required	Array of canonical URI paths, one for each virtualization host whose WWPN list is to be exported

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
wwpn-list	String	The WWPN list in Comma-Separated Values (CSV) format

### Description

The **Export World Wide Port Names List** operation returns the list of host port WWPNs of the virtualization hosts specified by the **virtualization-hosts** field of the request body. These virtualization hosts must be part of the ensemble specified by *{ensemble-id}*. The list is provided in a JSON object as a single string in Comma-Separated Values (CSV) format. It will be of the format described in the **Import Storage Access List** operation, with only the statement type, **Location**, and **HostWwpn** fields filled in. This result can be used as the basis for a Storage Access List to be supplied as input to the **Import Storage Access List** operation.

On successful execution, the WWPN list for the specified virtualization hosts is provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing ensemble object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. If the array of virtualization host URIs is empty, status code 400 (Bad Request) is returned. The URIs in the request body must designate existing Virtualization Host objects and the API user must have

object-access permission to them; otherwise, status code 404 (Not Found) is returned. The virtualization hosts must be part of the specified ensemble; otherwise, status code 400 (Bad Request) is returned. In addition, the API user must have action access permission to the **Export WWPNS** task; otherwise, status code 403 (Forbidden) is returned.

The request body is validated against the schema described in “Request body contents” on page 377. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object specified in the request URI
- Object-access permission to the hosting object of the virtualization hosts specified in the request body
- Action/task permission to the **Export WWPNS** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 377.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	146	A virtualization host specified in the request body is not a member of the ensemble specified in the request URI.
	149	The array of virtualization host URIs in the request body is empty.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{ensemble-id}</i> does not designate an existing ensemble object, or the API user does not have object access permission to it.
	2	A URI specified in the request body does not identify a virtualization host.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	2	The request could not be processed because the HMC is not currently communicating with an element of a zBX needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- This operation creates the same file as the **Export Host Port WWPNS** task. This file can be used as the basis for a Storage Access List. The Comma-Separated Value format has been chosen because it allows customers to use spreadsheet applications to display, sort, add or modify data.

- The **List Virtualization Host HBA Ports** operation provides similar information in JSON format.

## Example HTTP interaction

---

```
POST /api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334/operations/export-port-names HTTP/1.1
x-api-session: 5mkvfjvxt6guptdr5omvc11et4tvazeb684stvvhq4c1aa7w4x
content-type: application/json
content-length: 158
{
  "virtualization-hosts": [
    "/api/virtualization-hosts/71822c16-0401-11e1-8eda-001f163805d8",
    "/api/virtualization-hosts/2f676d90-03f8-11e1-8eda-001f163805d8"
  ]
}
```

---

Figure 193. Export World Wide Port Names List: WWPN list: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 07 Dec 2011 06:47:31 GMT
content-type: application/json;charset=UTF-8
content-length: 394
{
  "wwpn-list":
    "#Version: 1
    #FCP_DEF:,Name,Size,Description,Location,HostWwpn,TargetWwpn,Lun
    #ECKD_DEF:,Name,Size,Description,Location,Devno,Volser
    #ZVM_FCP_DEF:,Name,Size,Description,Location,Devno,Volser,HostWwpn,TargetWwpn,Lun
    FCP,,R32:B.2.12,21000024ff24df01
    FCP,,R32:B.2.12,21000024ff24df00
    FCP,,R32:B.2.02,21000024ff2b47cb
    FCP,,R32:B.2.02,21000024ff2b47ca
    "
}
```

---

Figure 194. Export World Wide Port Names List: WWPN list: Response

## Import Storage Access List

The **Import Storage Access List** operation imports information about storage resources and the virtualization hosts that have access to them. The Storage Access List (SAL) contains information, such as host port WWPNS and properties, such as addressing and device type information for configured storage resources.

### HTTP method and URI

**POST** /api/ensembles/{*ensemble-id*}/operations/import-storage-access-list

In this request, the URI variable {*ensemble-id*} is the object ID of the ensemble object on which the Storage Access List is to be imported.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
sal	String	Required	Storage Access List in Comma-Separated Values (CSV) format

The largest request body accepted by this operation is 1 MB. Requests with bodies that exceed this maximum are rejected with an HTTP status 413 (Request Entity Too Large) response.

## Description

The **Import Storage Access List** operation imports the provided Storage Access List into the ensemble specified by *{ensemble-id}*. The Storage Access List specifies paths between host ports and configured storage resources. It is a convenient way to specify which virtualization hosts have access to each storage resource.

The Storage Access List logically consists of lines of text, each one being a statement that identifies a storage resource, a virtualization host that has access to that storage resource and a path for the virtualization host to use when accessing the storage resource. For the full definition of the Storage Access List, see the storage access list worksheet described in the appendix in the *z Systems Ensemble Planning Guide*.

On successful execution, status code 204 (No Content) is returned without supplying a response body.

The URI path must designate an existing ensemble object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Import SAL** task; otherwise, status code 403 (Forbidden) is returned. All virtualization hosts designated in the Storage Access List are marked busy for the duration of this request. If any of those virtualization hosts is already marked busy due to some other operation, then status code 409 (Conflict) is returned.

The request body is validated against the schema described in “Request body contents” on page 379. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. In addition, the CSV-formatted SAL designated by the **sal** field must be syntactically and semantically correct; otherwise, status code 400 (Bad Request) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object specified in the request URI
- Action/task permission to the **Import SAL** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and no response body is provided.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	140	The CSV-formatted Storage Access List designated by the <b>sal</b> field is not syntactically and semantically correct, or an error was encountered while processing a Storage Access List entry. The response body contains a message with more details.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{ensemble-id}</i> does not designate an existing ensemble object, or the API user does not have object access permission to it.

HTTP error status code	Reason code	Description
409 (Conflict)	2	The operation cannot be performed because a virtualization host designated by the Storage Access List is currently busy performing some other operation.
	150	The operation cannot be performed because a virtualization host designated by the Storage Access List is currently busy due to a zBX Move operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- This operation accepts the same data/file as the **Import Storage Access List** task. This file contains information about storage resources, virtualization host storage resources, and path information that is to be added to zManager.

The Comma-Separated Value format has been chosen because it allows customers to use spreadsheet applications in order to display, sort, add or modify data.

- The **Create Storage Resource**, **Create Virtualization Host Storage Resource**, and **Add Virtualization Host Storage Resource Paths** operations can be used to provide the same information to zManager in JSON format.

## List Virtualization Host Storage Resources of a Storage Resource

The **List Virtualization Host Storage Resources of a Storage Resource** operation lists the virtualization host storage resources that are backed by the storage resource identified in the request URI.

### HTTP method and URI

**POST** `/api/storage-resources/{storage-resource-id}/operations/list-virtualization-host-storage-resources`

In this request, the URI variable `{storage-resource-id}` is the object ID of the storage resource.

### Response body contents

On successful completion, the response body is a JSON object with the following field:

Field name	Type	Description
virtualization-host-storage-resources	Array of objects	Array of virtualization-host-storage-resource-basic-info objects, as described in the next table. If no virtualization host storage resources are backed by the storage resource, an empty array is provided.

Each virtualization-host-storage-resource-basic-info object contains the following fields:

Field name	Type	Description
element-uri	String/ URI	Canonical URI path of the Virtualization Host Storage Resource object.
name	String	The <b>name</b> property of the associated Storage Resource object.
type	String Enum	The <b>type</b> property of the associated Storage Resource object.

## Description

The **List Virtualization Host Storage Resources of a Storage Resource** operation returns basic information about each virtualization host storage resource that is backed by the storage resource specified by *{storage-resource-id}*. The element URI and other basic properties are provided for each virtualization host storage resource. See the *virtualization-host-storage-resource-basic-info* object definition.

The URI path must designate an existing Storage Resource object and you must have object-access permission to it. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, you must have action access permission to the Manage Storage Resources action; otherwise, status code 403 (Forbidden) is returned.

If there are no virtualization host storage resources backed by the storage resource, an empty array is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the Storage Resource object specified in the request URI
- Action/task permission to the **Manage Storage Resources** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 381.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{storage-resource-id}</i> does not designate an existing Storage Resource object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/storage-resources/ca10557a-bf20-11e2-8d27-5cf3fcaf5679/operations/
list-virtualization-host-storage-resources HTTP/1.1
x-api-session: 66c13hnchkn6pzirewro4gxi9vwm6do8r6slo6cqf5ibnkw8q2
```

---

Figure 195. List Virtualization Host Storage Resources of a Storage Resource: Request

---

```

200 OK
content-length:445,
server:zSeries management console API web server / 2.0,
cache-control:no-cache,
date:Wed, 26 Jun 2013 01:54:02 GMT,
content-type:application/json;charset=UTF-8

{
  "virtualization-host-storage-resources":[
    {
      "element-uri":"/api/virtualization-hosts/580f90c0-836d-11e1-b3b0-f0def1cb6750/
        virtualization-host-storage-resources/c9ce6462-bf20-11e2-9ebe-f0def1cb66b0",
      "name":"p93_b1_06_V7000_000004",
      "type":"fcp"
    },
    {
      "element-uri":"/api/virtualization-hosts/83ec3aa8-8cc9-11e1-b149-f0def1cb6750/
        virtualization-host-storage-resources/8283c9b0-dda4-11e2-b713-f0def1cb6750",
      "name":"p93_b1_06_V7000_000004",
      "type":"fcp"
    }
  ]
}

```

---

Figure 196. List Virtualization Host Storage Resources of a Storage Resource: Response

## Inventory service data

Information about the storage resources managed by the HMC can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for Storage Resource objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class **"storage-resource"** are to be included. An entry for a particular storage resource is included only if the API user has object-access permission to that object.

For each Storage Resource object to be included, the inventory response array includes entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the Get Storage Resource Properties operation. That is, the data provided is the same as would be provided if a Get Storage Resource Properties operation were requested targeting this object.

## Sample inventory data

The following fragment is an example of the JSON object that would be included in the Get Inventory response to describe a single storage resource. This object would appear as one array entry in the response array:

```

{
  "allocation-status": "used",
  "allocation-units": "bytes",
  "class": "storage-resource",
  "description": "2024-0080-e518-4ac0\r\n0000-0000-0000-0000",
  "name": "B1010000A",
  "object-id": "2d588ee2-25a2-11e0-94a7-0010184c8334",
  "object-uri": "/api/storage-resources/2d588ee2-25a2-11e0-94a7-0010184c8334",
  "parent": "/api/ensembles/87d73ffc-75b2-11e0-9ba3-0010184c8026",
  "size": 10737418240,
  "type": "fcp",
  "unique-device-id": "3E21360080E5000184AC00000441B4D07449B0F1814 FASSt03IBMfcp"
}

```

Figure 197. Storage Resource object: Sample inventory data

## Virtualization Host Storage Resource object

A virtualization host storage resource represents a storage resource to which the virtualization host has been granted access. It is the representation of a storage resource from the perspective of a virtualization host.

### Data model

This object includes the following properties:

Table 87. Virtualization Host Storage Resource object properties

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/ URI	Canonical URI path of the Virtualization Host Storage Resource object, in the form <code>/api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-resources/{virt-host-storage-resource-id}</code> where <code>{virt-host-storage-resource-id}</code> is the value of the <b>element-id</b> property of this object.
<b>element-id</b>	—	String (36)	The unique identifier for the virtualization host storage resource instance. This identifier is in the form of a UUID.
<b>parent</b>	—	String/URI	The parent object of a Virtualization Host Storage Resource object is a Virtualization Host object.
<b>class</b>	—	String	The class of a Virtualization Host Storage Resource object is <b>"virtualization-host-storage-resource"</b> .
<b>name</b>	(pc)	String	The name of the storage resource, as defined in the Storage Resource object's Data Model section.
<b>description</b>	(pc)	String	The description for the storage resource, as defined in the Storage Resource object's Data Model section.
<b>size</b>	(pc)	Long	The size of the storage resource, as defined in the Storage Resource object's Data Model section.
<b>allocation-units</b>	—	String Enum	The units for the <b>size</b> property, as defined in the Storage Resource object's "Data model" on page 366.
<b>type</b>	—	String Enum	The type of the storage resource, as defined in the Storage Resource object's Data Model section.
<b>storage-resource</b>	—	String/ URI	Canonical URI path of the Storage Resource object associated with this virtualization host storage resource.



Table 87. Virtualization Host Storage Resource object properties (continued)

Name	Qualifier	Type	Description
<b>paths</b>	(w)	Array of objects	Information about the paths by which the storage resource is accessible to this virtualization host. It is an array of path-information-fcp or path-information-eckd objects. If the virtualization host has no paths to the storage resource, an empty array is provided. Note that there can be at most one path for a virtualization host storage resource whose <b>type</b> property is "eckd".
<b>unique-device-id</b>	(pc)	String	The unique device identifier of the storage resource, as defined in the Storage Resource object data model. See "Class specific additional properties" on page 367.
<b>volume-serial-number</b>	(pc)	String (1-6)	The volume serial for this virtualization host storage resource. Only present if the <b>type</b> property is "eckd" or "zvm-fcp".
<b>device-number</b>	—	String (1-4)	The device number that the virtualization host uses to access an ECKD storage resource, or the EDEV (emulated volume) that is created for FCP storage resources when they get allocated to a Virtual Server, or added to a Virtualization Host Storage Group. The string form of a 1-4 digit hexadecimal number. Only present if the <b>type</b> property is "eckd" or "zvm-fcp".
<b>virtualization-host-storage-group</b>	(pc)	String/URI	Canonical URI path of the virtualization host storage group of which this virtualization host storage resource is a member or null if this virtualization host storage resource is not in a virtualization host storage group. Only present if the <b>type</b> property is "eckd" or "zvm-fcp".

A path-information-fcp object contains information about a single path by which an FCP storage resource is accessible to a virtualization host. This object describes a path for a virtualization host storage resource whose **type** property is "fcp" or "zvm-fcp".

Table 88. Virtualization Host Storage Resource object: path-information-fcp object properties

Name	Type	Description
<b>host-port-wwpn</b>	String (16)	The WWPN of the host port. The string form of a 16-digit hexadecimal number.
<b>controller-port-wwpn</b>	String (16)	The WWPN of the storage controller port. The string form of a 16-digit hexadecimal number.
<b>lun</b>	String (16)	The Logical Unit Number (LUN) of the storage resource. The string form of a 16-digit hexadecimal number.
<b>accessible</b>	Boolean	Indicates whether the storage resource is currently accessible to the virtualization host via this path.  Because path accessibility status can be time consuming to determine, by default such status is omitted when Virtualization Host Storage Resource properties are returned and instead the value provided for this property is <b>null</b> . Operations that provide path accessibility status data will specifically indicate the conditions under which they do so.

A path-information-eckd object contains information about a single path by which an ECKD storage resource is accessible to a virtualization host. This object describes a path for a virtualization host storage resource whose **type** property is "eckd".

**Note:** the path specified through URM is the device number specified in the system I/O Configuration for the ECKD volume. It is not to be confused with the path (CHPID) between the system and the Control Unit.

Table 89. Virtualization Host Storage Resource object: path-information-eckd object properties

Name	Type	Description
device-number	String (1-4)	The device number of the storage resource. The string form of a 1-4 digit hexadecimal number.
accessible	Boolean	Indicates whether the storage resource is currently accessible to the virtualization host via this path.  Because path accessibility status can be time consuming to determine, by default such status is omitted when Virtualization Host Storage Resource properties are returned and instead the value provided for this property is <b>null</b> . Operations that provide path accessibility status data will specifically indicate the conditions under which they do so.

## Operations

If a virtualization host storage resource operation accesses a z/VM virtualization host and encounters an error while communicating with the virtualization host via SMAPI, the response body is as described in “SMAPI Error Response Body” on page 257.

## List Virtualization Host HBA Ports

The **List Virtualization Host HBA Ports** operation lists information about Fibre-Channel HBA (Host Bus Adapter) ports for a virtualization host.

### HTTP method and URI

**GET** /api/virtualization-hosts/{virt-host-id}/hba-ports

In this request, the URI variable {virt-host-id} is the object ID of the virtualization host.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
hba-ports	Array of objects	Information about the virtualization host's HBA ports. It is an array of hba-port-information objects, described in the next table. If no virtualization host HBA ports are to be returned, an empty array is provided.

Each hba-port-information object contains the following fields:

Field name	Type	Description
wwpn	String	World Wide Port Name (WWPN) of the port.
identifier	String	Identifier for the port. It contains the device number of the subchannel for a virtualization host whose <b>type</b> property is "zvm". For Power ASB, an example is fscsi2. For Intel based ASBs, the identifier starts with "/dev/...".

## Description

The **List Virtualization Host HBA Ports** operation lists a virtualization host's HBA ports. All properties are provided for each port.

The URI path must designate an existing Virtualization Host object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Manage Storage Resources** task; otherwise, status code 403 (Forbidden) is returned.

If there are no HBA ports for the virtualization host, an empty list is provided and the operation completes successfully with HTTP status code 200 (OK).

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Export WWPNS** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 386.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{virt-host-id}</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	2	The request could not be processed because the HMC is not currently communicating with an element of a zBX needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/ba97ff30-2990-11e0-8d5b-001f163803de/hba-ports HTTP/1.1
x-api-session: 1rmnds0imna61i3110eu7drk7jsec93mvlc1fbuqdb7xspk2fm5
```

---

Figure 198. List Virtualization Host HBA Ports: Request

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 04 Aug 2011 13:30:12 GMT
content-type: application/json;charset=UTF-8
content-length: 210
{
  "hba-ports": [
    {
      "identifier": "fscsi3",
      "wwpn": "21000024ff2b47cb"
    },
    {
      "identifier": "fscsi2",
      "wwpn": "21000024ff2b47ca"
    }
  ]
}

```

Figure 199. List Virtualization Host HBA Ports: Response

## List Virtualization Host Storage Resources

The **List Virtualization Host Storage Resources** operation lists the virtualization host storage resources for a virtualization host.

### HTTP method and URI

**GET** /api/virtualization-hosts/{*virt-host-id*}/virtualization-host-storage-resources

In this request, the URI variable *{virt-host-id}* is the object ID of the virtualization host.

### Query parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property
properties	String	Optional	Identifies the properties of each virtualization host storage resource to be returned. The only supported value is <b>"all"</b> , which results in all properties being returned, including path accessibility status properties. If this query parameter is omitted, a set of basic properties is returned.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
virtualization-host-storage-resources	Array of objects	If the <b>properties=all</b> query parameter is specified, an array is provided whose elements are the set of virtualization host storage resource properties that would be returned on a <b>Get Virtualization Host Storage Resource Properties</b> request with the <b>include-path-accessibility</b> query parameter specified as true. If the <b>properties</b> query parameter is omitted, an array of virtualization-host-storage-resource-basic-info objects is returned, described in the next table. If no virtualization host storage resources are to be returned, an empty array is provided.

Each virtualization-host-storage-resource-basic-info object contains the following fields:

Field name	Type	Description
element-URI	String/ URI	Canonical URI path of the Virtualization Host Storage Resource object
name	String	The <b>name</b> property of the associated Storage Resource object
type	String Enum	The <b>type</b> property of the associated Storage Resource object

## Description

The **List Virtualization Host Storage Resources** operation lists a virtualization host's virtualization host storage resources. The object URI and other basic properties are provided for each virtualization host storage resource.

If the **name** query parameter is specified, the returned list is limited to those virtualization host storage resources that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

If the **properties** query parameter is specified, it controls the set of properties returned. A value of **"all"** results in all properties being returned, in exactly the same format as would be provided on a **Get Virtualization Host Storage Resource Properties** request with the **include-path-accessibility** query parameter specified as true. If the **properties** query parameter is omitted, a set of basic properties is returned for each virtualization host storage resource. See the virtualization-host-storage-resource-basic-info object definition. Any value other than **all** is not valid and results in an HTTP status code 400 (Bad Request).

The URI path must designate an existing Virtualization Host object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Manage Storage Resources** task; otherwise, status code 403 (Forbidden) is returned.

The virtualization host designated by the URI path must have a **status** of **"operating"**, otherwise status code 409 (Conflict) is returned.

If there are no virtualization host storage resources for the virtualization host, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Manage Storage Resources** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents" on page 388.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	145	A value other than "all" was specified for the <b>properties</b> query parameter, or this query parameter was specified more than once.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{virt-host-id}</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it.
409 (Conflict)	1	Virtualization host has a <b>status</b> that is not valid for this operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	2	The request could not be processed because the HMC is not currently communicating with an element of a zBX needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- The information provided about a virtualization host storage resource can optionally include path accessibility information, i.e. information on whether a path to the storage resources is currently operational or not. However, determining path accessibility status can be an expensive action, especially when a virtualization host has one or more non-functional paths as SAN rediscovery is implicitly triggered in this case. For this reason, an API application should obtain path accessibility information only if it is required to satisfy the function of the application. If not required, the application should obtain virtualization host storage resource information using techniques that omit path accessibility information in order to avoid unnecessary delays.

If path-accessibility information is required for all or many virtualization host storage resources, it may be significantly faster to obtain that information by making a single request to this operation with the **properties=all** parameter specified rather than by making a series of requests to the **Get Virtualization Host Storage Resource Properties** operation (with the **include-path-accessibility=true** parameter specified) to obtain that information one resource at a time. This is because overhead due to SAN rediscovery would be incurred at most one time by using this operation, but might be incurred on each and every iterated **Get Virtualization Host Storage Resource Properties** request.

On the other hand, if path-accessibility information is not required, using this operation with the **properties=all** parameter specified may incur unnecessary application delays. Instead, the application can bypass the determination of path-accessibility status by using this operation with **properties=all** omitted to obtain the URIs of the virtualization host's storage resources, and then iterating over those URIs and making a requests to the **Get Virtualization Host Storage Resource Properties** operation with **include-path-accessibility=false** specified (or defaulted) for each. If the application requires virtualization host storage resource information for all or many virtualization hosts in the ensemble, obtaining this information via the **Get Inventory** operation of the Inventory Service (for the virtualization host inventory categories) may provide the best performance. This service can provide data across all virtualization hosts in the ensemble in a single request, and bypasses the potentially costly determination of path-accessibility status when obtaining storage resource information.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/ba97ff30-2990-11e0-8d5b-001f163803de/virtualization-host-
storage-resources HTTP/1.1
x-api-session: 1rmnds0imna61i3110eu7drk7jsec93mvlc1fbuqdb7xspk2fm5
```

---

Figure 200. List Virtualization Host Storage Resources: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 04 Aug 2011 13:30:12 GMT
content-type: application/json;charset=UTF-8
content-length: 729
{
"virtualization-host-storage-resources": [
{
"element-uri": "/api/virtualization-hosts/ba97ff30-2990-11e0-8d5b-001f163803de/virtualization-
host-storage-resources/dfbf23f0-b7b7-11e0-a46d-f0def152d359",
"name": "B2L0012",
"type": "fcp"
},
{
"element-uri": "/api/virtualization-hosts/ba97ff30-2990-11e0-8d5b-001f163803de/virtualization-
host-storage-resources/ff83dea8-64d1-11e0-b579-f0def10c03f4",
"name": "test1234",
"type": "fcp"
}
]
}
```

---

Figure 201. List Virtualization Host Storage Resources: Response

## Get Virtualization Host Storage Resource Properties

The **Get Virtualization Host Storage Resource Properties** operation retrieves the properties of a single Virtualization Host Storage Resource object.

### HTTP method and URI

```
GET /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-resources/
{virt-host-storage-resource-id}
```

#### URI variables

Variable	Description
{virt-host-id}	Object ID of the virtualization host
{virt-host-storage-resource-id}	Element ID of the Virtualization Host Storage Resource object for which properties are to be obtained

#### Query parameters

Name	Type	Rqd/Opt	Description
include-path-accessibility	Boolean	Optional	If specified as true, the accessibility status of the paths for this resource is determined and reported as values of the <b>accessible</b> property. If specified as false or omitted, this status is not determined and instead the value of the <b>accessible</b> property is always <b>null</b> .

## Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Virtualization Host Storage Resource object as defined in “Data model” on page 384. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

## Description

The **Get Virtualization Host Storage Resource Properties** operation returns the current properties for the Virtualization Host Storage Resource object that is specified by its object ID *{virt-host-id}* and the object ID of the owning virtualization host *{virt-host-storage-resource-id}*.

On successful execution, all of the current properties as defined in “Data model” on page 384 for the Virtualization Host Storage Resource object are provided in the response body and HTTP status code 200 (OK) is returned. If the **include-path-accessibility** query parameter is specified as true, these properties include the current path accessibility status of each of the paths for the resource (as the **accessible** property). If **include-path-accessibility** is false (which is the default), this path accessibility status is not provided and instead the **accessible** property is always null.

The URI path must designate an existing Virtualization Host object and the API user must have object-access permission to it. Furthermore, the URI path must designate an existing Virtualization Host Storage Resource object. If any of these conditions are not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Manage Storage Resources** task; otherwise, status code 403 (Forbidden) is returned.

The virtualization host designated by the URI path must have a **status** of **"operating"**, otherwise status code 409 (Conflict) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Manage Storage Resources** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.



HTTP error status code	Reason code	Description
404 (Not Found)	1	The object ID <i>{virt-host-id}</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it.
	147	The object ID <i>{virt-host-storage-resource-id}</i> does not designate an existing Virtualization Host Storage Resource object for the specified virtualization host.
	148	There is no Storage Resource object associated with the Virtualization Host Storage Resource object identified by the object ID <i>{virt-host-storage-resource-id}</i> . This is most likely a temporary condition due to a delete operation in progress on the Storage Resource object.
409 (Conflict)	1	Virtualization host has a <b>status</b> that is not valid for this operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	2	The request could not be processed because the HMC is not currently communicating with an element of a zBX needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- Determining path accessibility status can be an expensive action, especially when a virtualization host has one or more non-functional paths as SAN rediscovery is implicitly triggered in this case. For this reason, this operation omits path accessibility status information (the **accessible** property) by default. Use the **include-path-accessibility** query parameter to request that this status be determined and reported when specifically needed by the application.
- If an application requires path accessibility status information for all or many of the storage resources of a virtualization host, IBM recommends using the **List Virtualization Host Storage Resources** operation with the **properties=all** query parameter to obtain information for all resources in a single request as a better performing approach than iteratively using **Get Virtualization Host Storage Resources** (with **include-path-accessibility=true**) one resource at a time. Repeated use of **Get Virtualization Host Storage Resource** may incur SAN rediscovery overhead once per request, but such overhead would be incurred at most once in the single **List Virtualization Host of Storage Resources** request.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/75ca2d2e-e854-11df-811c-00262df32766/
    virtualization-host-storage-resources/c0261be8-ec51-11df-85fe-00262df32766 HTTP/1.1
x-api-session: 3gcd77g1emvwq81dlmwxc8i4fwm4udx1by6i2auls4r6g529p1
```

---

Figure 202. Get Virtualization Host Storage Resource Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 06 Dec 2011 14:34:00 GMT
content-type: application/json;charset=UTF-8
content-length: 1043
{
  "allocation-units": "bytes",
  "class": "virtualization-host-storage-resource",
  "description": "xiv-41cb",
  "element-id": "c0261be8-ec51-11df-85fe-00262df32766",
  "element-uri": "/api/virtualization-hosts/75ca2d2e-e854-11df-811c-00262df32766/
  virtualization-host-storage-resources/c0261be8-ec51-11df-85fe-00262df32766",
  "name": "r93c1_12_hdisk4",
  "parent": "/api/virtualization-hosts/75ca2d2e-e854-11df-811c-00262df32766",
  "paths": [
    {
      "accessible": null,
      "controller-port-wwpn": "500173800aa50180",
      "host-port-wwpn": "2101001b32bf37e3",
      "lun": "41cb000000000000"
    },
    {
      "accessible": null,
      "controller-port-wwpn": "500173800aa50142",
      "host-port-wwpn": "2100001b329f37e3",
      "lun": "41cb000000000000"
    }
  ],
  "size": 34359738368,
  "storage-resource": "/api/storage-resources/c04f6f70-ec51-11df-a5bc-00215e6a0c27",
  "type": "fcp",
  "unique-device-id": "26112001738000AA5010B072810XIV03IBMfcp"
}
```

---

Figure 203. Get Virtualization Host Storage Resource Properties: Response for Virtualization Host of type "power-vm" or "x-hyp"

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 06 Dec 2011 14:34:01 GMT
content-type: application/json;charset=UTF-8
content-length: 651
{
  "allocation-units": "cylinders",
  "class": "virtualization-host-storage-resource",
  "description": "B71C",
  "device-number": "B71C",
  "element-id": "e85e91c2-ee02-11e0-a0eb-00262df332b3",
  "element-uri": "/api/virtualization-hosts/0e4a5d94-a8c1-11e0-9492-00262df332b3/
  virtualization-host-storage-resources/e85e91c2-ee02-11e0-a0eb-00262df332b3",
  "name": "B71C",
  "parent": "/api/virtualization-hosts/0e4a5d94-a8c1-11e0-9492-00262df332b3",
  "paths": [
    {
      "accessible": null,
      "device-number": "B71C"
    }
  ],
  "size": 60000,
  "storage-resource": "/api/storage-resources/ec9c3852-ee02-11e0-bc09-00215e6a0c27",
  "type": "eckd",
  "virtualization-host-storage-group": null,
  "volume-serial-number": "NNB7BC"
}

```

Figure 204. Get Virtualization Host Storage Resource Properties: Response for Virtualization Host of type "zvm"

## Create Virtualization Host Storage Resource

The **Create Virtualization Host Storage Resource** operation creates a new virtualization host storage resource for the virtualization host.

### HTTP method and URI

**POST** /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-resources

In this request, the URI variable {virt-host-id} is the object ID of the virtualization host that owns the new/modified virtualization host storage resource.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
storage-resource	String/URI	Required	Canonical URI path of the Storage Resource object to be associated with this virtualization host storage resource.
paths	Array of objects	Optional	The path information for the new virtualization host storage resource. It is either an array of new-path-fcp or new-path-eckd objects, as described in the <b>Add Virtualization Host Storage Resource Path</b> operation. Note that there can be at most one path for a virtualization host storage resource whose <b>type</b> property is "eckd".

Field name	Type	Rqd/Opt	Description
volume-serial-number	String (1-6)	Required when creating a VHSR on z/VM	The <b>volume-serial-number</b> property of the Virtualization Host Storage Resource object. This field is required when this operation is to create a new virtualization host storage resource on z/VM; otherwise, it is optional. This field only applies to storage resources whose <b>type</b> property is "eckd" or "zvm-fcp". The volume serial number is written to the ECKD storage resource or FCP storage resources when they get added to a virtualization host storage group. The same volume serial number is to be used when creating multiple virtualization host storage resources for a storage resource.
device-number	String (1-4)	Required when creating an FCP VHSR on z/VM	The <b>device-number</b> property of the Virtualization Host Storage Resource object. This field is required when this operation is to create a new virtualization host storage resource; otherwise, it is optional. This field only applies to storage resources whose <b>type</b> property is "eckd" or "zvm-fcp". The device number is used to access an ECKD storage resource, or assigned to the EDEV (emulated volume) that is created for FCP storage resources when they get allocated to a virtual server, or added to a virtualization host storage group.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Virtualization Host Storage Resource object, in the form <code>/api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-resources/{virt-host-storage-resource-id}</code>

## Description

The **Create Virtualization Host Storage Resource** operation creates a new virtualization host storage resource for the virtualization host specified by the `{virt-host-id}` portion of the request URI.

Upon successful completion, the **element-uri** field of the response body and the **Location** response header identify the new virtualization host storage resource. An inventory change event is emitted asynchronously. See "Notifications" on page 410 for more information.

The URI path must designate an existing Virtualization Host object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Add Storage Resource** task; otherwise, status code 403 (Forbidden) is returned. The virtualization host is marked busy for the duration of this request. If it is already marked busy due to some other operation, then status code 409 (Conflict) is returned.

If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Add Storage Resource** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and both the response body and the **Location** response header contain the URI of the newly created object.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	7	The Storage Resource object URI ( <b>storage-resource</b> ) in the request body designates a storage resource that is not compatible with the virtualization host designated in the request URI ( <i>virt-host-id</i> ).
	141	The virtualization host storage resource already has the maximum allowed number of paths.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>virt-host-id</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it.
	2	The storage resource URI specified in the request body does not identify a storage resource.
409 (Conflict)	2	The operation cannot be performed because a virtualization host designated by the request URI is currently busy performing some other operation.
	150	The operation cannot be performed because a virtualization host designated by the request URI is currently busy due to a zBX Move operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

The **Create Virtualization Host Storage Resource** operation performs a portion of the function provided by the Add Storage Resource subtask of the **Manage Storage Resources** task. The Add Storage Resource subtask performs some or all of the functions provided by the following Web Services APIs:

- **Create Storage Resource**

Creates a Storage Resource object in zManager, representing storage resources (FCP LUNs, or ECKD volumes). Each Storage Resource object has a unique name. zManager reports additional storage resource specific information, such as a unique-device-id, after paths have been added and zManager was able to access the storage resource.

- **Create Virtualization Host Storage Resource**

Creates an object in zManager, representing the hypervisor's view on the storage resource.

- **Add Virtualization Host Storage Resource Path**

Adds a path between a virtualization host and a storage resource.

## Example HTTP interaction

---

```
POST /api/virtualization-hosts/2f029af0-03f8-11e1-8eda-001f163805d8/
  virtualization-host-storage-resources HTTP/1.1
x-api-session: 68ps5rqvq1177xtcqz9rnrmls29mgllluq7ni0qlgnwjhup01c
content-type: application/json
content-length: 205
{
  "paths": [
    {
      "controller-port-wwpn": "20240080e5184ac0",
      "host-port-wwpn": "21000024ff2b4602",
      "lun": "1234001000000000"
    }
  ],
  "storage-resource": "/api/storage-resources/17556bdc-2034-11e1-83b8-0010184c8334"
}
```

---

Figure 205. Create Virtualization Host Storage Resource: Request

---

```
201 Created
server: zSeries management console API web server / 1.0
location: /api/virtualization-hosts/2f029af0-03f8-11e1-8eda-001f163805d8/
  virtualization-host-storage-resources/19625098-2034-11e1-b4a5-001f163805d8
cache-control: no-cache
date: Tue, 06 Dec 2011 18:00:24 GMT
content-type: application/json;charset=UTF-8
content-length: 155
{
  "element-uri": "/api/virtualization-hosts/2f029af0-03f8-11e1-8eda-001f163805d8/
  virtualization-host-storage-resources/19625098-2034-11e1-b4a5-001f163805d8"
}
```

---

Figure 206. Create Virtualization Host Storage Resource: Response

## Delete Virtualization Host Storage Resource

The **Delete Virtualization Host Storage Resource** operation removes a specified virtualization host storage resource from the specified virtualization host.

### HTTP method and URI

```
DELETE /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-resources/
  {virt-host-storage-resource-id}
```

### URI variables

Variable	Description
{virt-host-id}	Object ID of the virtualization host
{virt-host-storage-resource-id}	Element ID of the Virtualization Host Storage Resource object to be deleted

### Description

The **Delete Virtualization Host Storage Resource** operation removes a specified virtualization host storage resource from the specified virtualization host, removing all related path information at the same time. The virtualization host storage resource is identified by {virt-host-storage-resource-id} in the URI, and

the virtualization host is identified by *{virt-host-id}* in the URI. The virtualization host storage resource must not be part of a virtualization host storage group, and there must be no virtual disks backed by the virtualization host storage resource.

Upon successfully removing the virtualization host storage resource, HTTP status code 204 (No Content) is returned and no response body is provided. An inventory change event is issued asynchronously. See “Notifications” on page 410 for more information.

The URI path must designate an existing virtualization host and the API user must have object-access permission to it. Furthermore, the URI path must designate an existing Virtualization Host Storage Resource object. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must also have action access permission to the **Remove Storage Resource** task as well; otherwise, status code 403 (Forbidden) is returned. If the virtualization host storage resource is part of a virtualization host storage group or there is a virtual disk backed by the virtualization host storage resource, then status code 409 (Conflict) is returned. The virtualization host is marked busy for the duration of the request. If it is already marked busy due to some other operation, then status code 409 (Conflict) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Remove Storage Resource** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{virt-host-id}</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it.
	147	The object ID in the URI <i>{virt-host-storage-resource-id}</i> does not designate an existing Virtualization Host Storage Resource object for the specified virtualization host.
409 (Conflict)	2	The operation cannot be performed because the virtualization host designated by the request URI is currently busy performing some other operation.
	144	The object cannot be deleted at this time. Either the virtualization host storage resource is part of a virtualization host storage group or a virtual disk is backed by the virtualization host storage resource.
	150	The operation cannot be performed because the virtualization host designated by the request URI is currently busy due to a zBX Move operation.

HTTP error status code	Reason code	Description
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Usage note

- While this operation does delete all path information associated with this virtualization host storage resource, it does not delete the associated storage resource, even if there is no longer a virtualization host storage resource associated with the storage resource. If the storage resource is no longer needed, it can be deleted using the **Delete Storage Resource** operation.

### Example HTTP interaction

---

```
DELETE /api/virtualization-hosts/2f029af0-03f8-11e1-8eda-001f163805d8/
  virtualization-host-storage-resources/19625098-2034-11e1-b4a5-001f163805d8 HTTP/1.1
x-api-session: 68ps5rqvq1177xtcqz9rnrblms29mglluq7ni0qlgnwjhup01c
```

---

Figure 207. Delete Virtualization Host Storage Resource: Request

---

```
204 No Content
date: Tue, 06 Dec 2011 18:00:53 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

<No response body>

---

Figure 208. Delete Virtualization Host Storage Resource: Response

## Add Virtualization Host Storage Resource Paths

The **Add Virtualization Host Storage Resource Paths** operation adds a path definition to the virtualization host storage resource. The path is defined by its two endpoints – the virtualization host is at one end, and the associated storage resource is at the other end.

### HTTP method and URI

```
POST /api/virtualization-hosts/{virt-host-id}/operations/add-paths
```

In this request, the URI variable *{virt-host-id}* is the object ID of the Virtualization Host object.



## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
virtualization-host-storage-resource-element-uri	String/URI	Required	Canonical URI path of the Virtualization Host storage Resource object.
paths	Array of objects	Required	Either an array of new-path-fcp or new-path-eckd objects, described in the next table. Note that there can be at most one path for a virtualization host storage resource whose <b>type</b> property is "eckd".

A new-path-fcp object contains information about a single path by which an FCP storage resource is accessible to a virtualization host. This object describes a path for a virtualization host storage resource whose **type** property is either "fcp" or "zvm-fcp".

Name	Type	Rqd/Opt	Description
host-port-wwpn	String (16)	Required	The WWPN of the host port. The string form of a 16-digit hexadecimal number.
controller-port-wwpn	String (16)	Required	The WWPN of the storage controller port. The string form of a 16-digit hexadecimal number.
lun	String (16)	Required	The Logical Unit Number (LUN) of the storage resource. The string form of a 16-digit hexadecimal number.

A new-path-eckd object contains information about a single path by which an ECKD storage resource is accessible to a virtualization host. This object describes a path for a virtualization host storage resource whose **type** property is "eckd".

Name	Type	Rqd/Opt	Description
device-number	String (1-4)	Required	The device number of the storage resource. The string form of a 1-4 digit hexadecimal number.

## Description

The **Add Virtualization Host Storage Resource Paths** operation adds path definitions to the specified virtualization host storage resource. The **controller-port-wwpn**, **lun**, or **device-number**, as appropriate, along with the **host-port-wwpn**, must identify a configured path of the storage resource associated with this virtualization host storage resource. If that condition is not met, HTTP status code 400 (Bad Request) is returned.

Only a single path may be configured for an ECKD virtualization host storage resource. An attempt to define more than one path for such a resource will result in HTTP status code 400 (Bad Request) being returned.

The URI path must designate an existing Virtualization Host object, and the API user must have object-access permission to it. Furthermore, the request body must designate an existing Virtualization Host storage Resource object. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Add Storage Resource** task; otherwise, status code 403 (Forbidden) is returned. The virtualization host is marked busy for the duration of this request. If it is already marked busy due to some other operation, then status code 409 (Conflict) is returned.

Upon success, HTTP status code 204 (No Content) is returned and no response body is provided. If this operation changes the value of the **path** property, a property-change notification is issued asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Add Storage Resource** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	141	The virtualization host storage resource already has the maximum allowed number of paths.
	142	The specified path information does not identify a configured path available for the virtualization host to use to access the specified storage resource.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{virt-host-id}</i> ) does not designate an existing Virtualization Host object, or the API user does not have object access permission to it.
	147	The URI in the request body ( <i>{virtualization-host-storage-resource}</i> ) does not designate an existing Virtualization Host storage Resource object for the specified virtualization host.
409 (Conflict)	2	The operation cannot be performed because the virtualization host designated by the request URI is currently busy performing some other operation.
	150	The operation cannot be performed because the virtualization host designated by the request URI is currently busy due to a zBX Move operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	2	The request could not be processed because the HMC is not currently communicating with an element of a zBX needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- The **Add Virtualization Host Storage Resource Paths** operation performs a portion of the function provided by the Add Storage Resource subtask of the **Manage Storage Resources** task. The Add Storage Resource subtask performs some or all of the functions provided by the following Web Services APIs:
  - **Create Storage Resource**  
Creates a Storage Resource object in zManager, representing storage resources (FCP LUNs, or ECKD volumes). Each Storage Resource object has a unique name. zManager reports additional storage resource specific information, such as a unique-device-id, after paths have been added and zManager was able to access the storage resource.
  - **Create Virtualization Host Storage Resource**  
Creates an object in zManager, representing the hypervisor's view on the storage resource.
  - **Add Virtualization Host Storage Resource Path**  
Adds a path between a virtualization host and a storage resource.
- There is no request to remove a virtualization host storage resource path. In order to remove a path, the virtualization host storage resource must be deleted and recreated with only the desired path(s).

## Example HTTP interaction

---

```
POST /api/virtualization-hosts/2f029af0-03f8-11e1-8eda-001f163805d8/operations/add-paths HTTP/1.1
x-api-session: 68ps5rqvq1177xtcqz9rnrblms29mglluq7ni0qlgnwjhup01c
content-type: application/json
content-length: 315
{
  "paths": [
    {
      "controller-port-wwpn": "20240080e5184ac0",
      "host-port-wwpn": "21000024ff2b4603",
      "lun": "1234001000000000"
    }
  ],
  "virtualization-host-storage-resource-element-uri": "/api/virtualization-hosts/2f029af0-03f8-11e1-8eda-001f163805d8/virtualization-host-storage-resources/19625098-2034-11e1-b4a5-001f163805d8"
}
```

---

Figure 209. Add Virtualization Host Storage Resource Paths: Request

---

```
204 No Content
date: Tue, 06 Dec 2011 18:00:53 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 210. Add Virtualization Host Storage Resource Paths: Response

## Remove Virtualization Host Storage Resource Paths

The **Remove Virtualization Host Storage Resource Paths** operation removes a path definition from the virtualization host storage resource. The path is defined by its two endpoints – the virtualization host is at one end, and the associated storage resource is at the other end.

## HTTP method and URI

POST /api/virtualization-hosts/{*virt-host-id*}/operations/remove-paths

In this request, the URI variable *{virt-host-id}* is the object ID of the Virtualization Host object.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
virtualization-host-storage-resource-element-uri	String/URI	Required	Canonical URI path of the Virtualization Host Storage Resource object.
paths	Array of objects	Required	Information about the paths to be removed. It is either an array of path-fcp or path-eckd objects, as described in the next tables. Note that there can be at most one path for a virtualization host storage resource whose <b>type</b> property is "eckd".

A path-fcp object contains information about a single path by which an FCP storage resource is accessible to a virtualization host. This object describes a path for a virtualization host storage resource whose **type** property is either "fcp" or "zvm-fcp".

Name	Type	Rqd/Opt	Description
host-port-wwpn	String (16)	Required	The WWPN of the host port. The string form of a 16-digit hexadecimal number.
controller-port-wwpn	String (16)	Required	The WWPN of the storage controller port. The string form of a 16-digit hexadecimal number.
lun	String (16)	Required	The Logical Unit Number (LUN) of the storage resource. The string form of a 16-digit hexadecimal number.

A path-eckd object contains information about a single path by which an ECKD storage resource is accessible to a virtualization host. This object describes a path for a virtualization host storage resource whose **type** property is "eckd". A path-fcp object contains information about a single path by which an FCP storage resource is accessible to a virtualization host. This object describes a path for a virtualization host storage resource whose **type** property is either "fcp" or "zvm-fcp".

Name	Type	Rqd/Opt	Description
device-number	String (1-4)	Required	The device number of the storage resource. The string form of a 1-4 digit hexadecimal number.

## Description

The **Remove Virtualization Host Storage Resource Paths** operation removes path definitions from the specified virtualization host storage resource. The **controller-port-wwpn**, **lun**, or **device-number**, as appropriate, along with the **host-port-wwpn**, must identify a configured path of the storage resource associated with this virtualization host storage resource. If that condition is not met, HTTP status code 400 (Bad Request) is returned.

The URI path must designate an existing Virtualization Host object, and the API user must have object-access permission to it. Furthermore, the request body must designate an existing Virtualization Host Storage Resource object. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Remove Storage Resources**

task; otherwise, status code 403 (Forbidden) is returned. The virtualization host is marked busy for the duration of this request. If it is already marked busy due to some other operation, then status code 409 (Conflict) is returned.

Upon success, HTTP status code 204 (No Content) is returned and no response body is provided. If this operation changes the value of the **path** property, a property-change notification is issued asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Manage Storage Resources** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	142	The specified path information does not identify a configured path available for the virtualization host to use to access the specified storage resource.
403 (Forbidden)	1	API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{virt-host-id}</i> ) does not designate an existing Virtualization Host object, or the API user does not have object access permission to it, or it is not of a type for which this operation is supported.
	147	The URI in the request body ( <i>{virtualization-host-storage-resource}</i> ) does not designate an existing Virtualization Host Storage Resource object for the specified virtualization host.
409 (Conflict)	2	The operation cannot be performed because the virtualization host designated by the request URI is currently busy performing some other operation.
	150	The operation cannot be performed because the virtualization host designated by the request URI is currently busy due to a zBX Move operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	2	The request could not be processed because the HMC is not currently communicating with an element of a zBX needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- The **Remove Virtualization Host Storage Resource Paths** operation performs a portion of the function (remove path) provided by the Details for Storage Resource subtask of the **Manage Storage Resources** task. The operation performs some or all of the functions provided by the following Web Services APIs:
  - **Remove Virtualization Host Storage Resource Path**  
Removes a path between a virtualization host and a storage resource.
  - Gives a warning/error message in case an incorrect path or no path is selected.

## Discover Virtualization Host Storage Resources

The **Discover Virtualization Host Storage Resources** operation discovers the virtualization host storage resources for a virtualization host.

### HTTP method and URI

POST `/api/virtualization-hosts/{virt-host-id}/operations/discover-virtualization-host-storage-resources`

In this request, the URI variable `{virt-host-id}` is the object ID of the Virtualization Host object.

### Query parameters

Name	Type	Rqd/Opt	Description
prefix	String (1-50)	Optional	Prefix to use when constructing the value of the <b>name</b> property of each newly discovered virtualization host storage resource. It must consist only of alphanumeric characters and the following special characters: “._-”, and it must begin with an alphanumeric character.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
virtualization-host-storage-resources	Array of objects	On successful completion, the response body contains a JSON object that provides an array whose elements are discovered-virtualization-host-storage-resource objects (described in the next table) including path accessibility status. Field names and data types in the discovered-virtualization-host-storage-resource are the same as the property names and data types defined in the data model.  If no virtualization host storage resources are to be returned, an empty array is provided.

Each discovered-virtualization-host-storage-resource object contains the following fields:

Field name	Type	Description
parent	String/URI	The <b>parent</b> property of the Virtualization Host Storage Resource object.
class	String	The <b>class</b> property of the Virtualization Host Storage Resource object.
name	String	The <b>name</b> property of the Virtualization Host Storage Resource object.
size	Long	The <b>size</b> property of the Virtualization Host Storage Resource object.
allocation-units	String Enum	The <b>allocation-units</b> property of the Virtualization Host Storage Resource object.
type	String Enum	The <b>type</b> property of the Virtualization Host Storage Resource object.

Field name	Type	Description
storage-resource	String/URI	The <b>storage-resource</b> property of the Virtualization Host Storage Resource object. This will be null unless this virtualization host storage resource defines an additional path for the virtualization host to access an existing storage resource.
paths	Array of objects	The <b>paths</b> property of the Virtualization Host Storage Resource object.
unique-device-id	String	The <b>unique-device-id</b> property of the Virtualization Host Storage Resource object.
volume-serial-number	String (1-6)	The volume serial for this virtualization host storage resource. Only present if the <b>type</b> property is "eckd" or "zvm-fcp".
device-number	String (1-4)	The device number that the virtualization host uses to access an ECKD storage resource, or the EDEV (emulated volume) that is created for FCP storage resources when they get allocated to a Virtual Server, or added to a Virtualization Host Storage Group. The string form of a 1-4 digit hexadecimal number. Only present if the <b>type</b> property is "eckd" or "zvm-fcp".

## Description

The **Discover Virtualization Host Storage Resources** operation discovers a virtualization host's virtualization host storage resources. The current values of the properties for the discovered Virtualization Host Storage Resource objects, including path accessibility status, are returned as defined in the "Data model" on page 384.

If the **prefix** parameter is omitted, the ensemble's default prefix will be used. The default prefix consists of the name of the ensemble concatenated with "\_SR".

If there are no discovered virtualization host storage resources for the specified virtualization host, an empty array is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Discover Storage Resources** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents" on page 406.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{virt-host-id}</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it.

HTTP error status code	Reason code	Description
409 (Conflict)	2	The operation cannot be performed because the virtualization host designated by the request URI is currently busy performing some other operation.
	150	The operation cannot be performed because the virtualization host designated by the request URI is currently busy due to a zBX Move operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## List Virtual Disks of a Virtualization Host Storage Resource

The **List Virtual Disks of a Virtualization Host Storage Resource** operation lists the virtual disks that are backed by the virtualization host storage resource identified in the request body.

### HTTP method and URI

```
POST /api/virtualization-hosts/{virt-host-id}/operations/
list-virtual-disks-of-virtualization-host-storage-resource
```

In this request, the URI variable *{virt-host-id}* is the object ID of the virtualization host.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
virtualization-host-storage-resource	String/ URI	Required	Canonical URI path of the Virtualization Host Storage Resource object

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
virtual-disks	Array of objects	Array of virtual-disk-basic-info objects, as described in the next table. If no virtual disks are backed by the virtualization host storage resource, an empty array is provided.

Each virtual-disk-basic-info object contains the following fields:

Field name	Type	Description
element-URI	String/ URI	Canonical URI path of the virtual disk object
name	String	The <b>name</b> property of the virtual disk object
type	String Enum	The <b>type</b> property of the virtual disk object



## Description

The **List Virtual Disks of a Virtualization Host Storage Resource** operation returns basic information about each virtual disk that is backed by the virtualization host storage resource identified in the request body. The element URI and other basic properties are provided for each virtual disk. See the `virtual-disk-basic-info` object definition.

The URI path must designate an existing virtualization host and the API user must have object-access permission to its hosting object. Furthermore, the request body must designate an existing Virtualization Host Storage Resource object of that virtualization host. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Manage Storage Resources** action; otherwise, status code 403 (Forbidden) is returned.

If there are no virtual disks backed by the virtualization host storage resource, an empty array is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Manage Storage Resources** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 408.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have the required permission to this operation.
404 (Not Found)	1	The object ID <i>{virt-host-id}</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it.
	147	The URI in the request body ( <i>{virtualization-host-storage-resource}</i> ) does not designate an existing Virtualization Host Storage Resource object for the specified virtualization host.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/virtualization-hosts/83ec3aa8-8cc9-11e1-b149-f0def1cb6750/operations
  /list-virtual-disks-of-virtualization-host-storage-resource HTTP/1.1
x-api-session: 1tmwwkw0b80axgzun0mwfc7p5uxdv4gn1tbsul8e1sep2u4o2r

{
  "virtualization-host-storage-resource": "/api/virtualization-hosts/
    83ec3aa8-8cc9-11e1-b149-f0def1cb6750/virtualization-host-storage-resources/
    11c32616-a1f5-11e2-893e-f0def1cb6750"
}
```

---

Figure 211. List Virtual Disks of a Virtualization Host Storage Resource: Request

---

```
"content-length": "180",
"server": "zSeries management console API web server / 2.0",
"cache-control": "no-cache",
"date": "Wed, 26 Jun 2013 01:55:06 GMT",
"content-type": "application/json; charset=UTF-8"

{
  "virtual-disks": [
    {
      "element-uri": "/api/virtual-servers/5b529726-983c-11e1-baab-f0def1cb66b0/
        virtual-disks/8aa31a04-a205-11e2-a0fd-f0def1cb6750",
      "name": "hdisk0",
      "type": "fullpack"
    }
  ]
}
```

---

Figure 212. List Virtual Disks of a Virtualization Host Storage Resource: Response

## Notifications

Changes to properties of a virtualization host storage resource are reflected via a property change notification designating the virtualization host as the managed object and the virtualization host storage resource as the element object. The standard property change notification fields (old value, new value and property name) for each changed property are provided in the notification.

The creation or deletion of a virtualization host storage resource is reflected via an inventory change notification designating the virtualization host as the managed object and the virtualization host storage resource as the element object.

## Inventory service data

Information about the storage resources associated with a virtualization host can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for Virtualization Host Storage Resource objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**virtualization-host**" are to be included. An entry for a particular virtualization host storage resource is included only if the API user has object-access permission to the owning virtualization host.

## Virtualization Host Storage Group object

A virtualization host storage group is a representation of a z/VM Storage Group. It consists of homogeneous storage resources to which a z/VM virtualization host has access, and whose corresponding virtualization host storage resources have been placed into a storage group.

### Data model

This object includes the following properties:

Table 90. Virtualization Host Storage Group object properties

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/URI	Canonical URI path of the Virtualization Host Storage Group object, in the form <code>/api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-groups/{virt-host-storage-group-id}</code> where <code>{virt-host-storage-group-id}</code> is the value of the <b>element-id</b> property of the virtualization host storage group.
<b>element-id</b>	—	String (36)	The unique identifier for the virtualization host storage group instance. This identifier is in the form of a UUID.
<b>parent</b>	—	String/URI	The parent object of a Virtualization Host Storage Group object is a Virtualization Host object.
<b>class</b>	—	String	The class of a Virtualization Host Storage Group object is <b>"virtualization-host-storage-group"</b> .
<b>name</b>	—	String (1-64)	The name of the virtualization host storage group. It must consist only of alphanumeric characters, spaces and the following special characters: <code>“._-\$”</code> , and it must begin with an alphanumeric character or <code>“\$”</code> .
<b>description</b>	—	String (0-256)	The description for the virtualization host storage group. It must consist only of alphanumeric characters, spaces and the following special characters: <code>“._-\$”</code> .
<b>free-space</b>	—	Array of object	Information about the free space in the virtualization host storage group. It is an array of free-space-information objects, each of which describes an area of free storage in the virtualization host storage group. If there is no free space in the virtualization host storage group, an empty array is provided.
<b>virtualization-host</b>	—	String/URI	Canonical URI path of the virtualization host that owns this virtualization host storage group.

A free-space-information object contains information about an area of storage in the virtualization host storage group that is currently not in use.

Table 91. Virtualization Host Storage Group object: free-space-information object properties

Name	Type	Description
<b>device-number</b>	String (1-4)	The device number of the virtualization host storage resource which has free space on it. This is the string form of a 1-4 digit hexadecimal number.
<b>size</b>	Long	The size of the free storage area. The units for this property are specified by the <b>allocation-units</b> property.
<b>allocation-units</b>	String Enum	The units for the size property. Values: <ul style="list-style-type: none"> <li>• <b>"bytes"</b></li> <li>• <b>"cylinders"</b></li> </ul>

## Operations

If a virtualization host storage group operation accesses a z/VM virtualization host and encounters an error while communicating with the virtualization host via SMAPI, the response body is as described in “SMAPI Error Response Body” on page 257.

## List Virtualization Host Storage Groups

The **List Virtualization Host Storage Groups** operation lists the virtualization host storage groups for a virtualization host.

### HTTP method and URI

**GET** /api/virtualization-hosts/{*virt-host-id*}/virtualization-host-storage-groups

In this request, the URI variable {*virt-host-id*} is the object ID of the Virtualization Host object.

### Query parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
properties	String	Optional	Identifies the properties of each virtualization host storage group to be returned. The only supported value is " <b>all</b> ", which results in all properties being returned. If this query parameter is omitted, a set of basic properties considered to be of general interest is returned.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
virtualization-host-storage-groups	Array of objects	If the <b>properties=all</b> query parameter is specified, an array is provided whose elements are the set of virtualization host storage group properties that would be returned on a <b>Get Virtualization Host Storage Group Properties</b> operation. If the <b>properties</b> query parameter is omitted, an array of virtualization-host-storage-group-basic-info objects is returned, described in the next table. If no virtualization host storage groups are to be returned, an empty array is provided.

Each virtualization-host-storage-group-basic-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Virtualization Host Storage Group object
name	String	The <b>name</b> property of the Virtualization Host Storage Group object

### Description

The **List Virtualization Host Storage Groups** operation lists a virtualization host's virtualization host storage groups. The element URI and name are provided for each virtualization host storage group.

If the **name** query parameter is specified, the returned list is limited to those virtualization host storage groups that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not performed.

If the **properties** query parameter is specified, it controls the set of properties returned. A value of "all" results in all properties being returned, in exactly the same format as would be provided on a **Get Virtualization Host Storage Group Properties** operation. If the **properties** query parameter is omitted, a set of basic properties is returned for each virtualization host storage group. See the virtualization-host-storage-group-basic-info object definition. Any value other than "all" is not valid and results in an HTTP status code 400 (Bad Request).

The URI path must designate an existing Virtualization Host object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Manage Storage Resources** task; otherwise, status code 403 (Forbidden) is returned.

If the virtualization host has no virtualization host storage groups, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Manage Storage Resources** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents" on page 412.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	145	A value other than "all" was specified for the <b>properties</b> query parameter, or this query parameter was specified more than once.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{virt-host-id}</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it, or it is not of a type for which this operation is supported.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4/
    virtualization-host-storage-groups HTTP/1.1
x-api-session: 1rmnds0imna61i3110eu7drk7jsec93mvlc1fbuqdb7xspk2fm5
```

---

Figure 213. List Virtualization Host Storage Groups: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Thu, 04 Aug 2011 13:30:20 GMT
content-type: application/json;charset=UTF-8
content-length: 683
{
  "virtualization-host-storage-groups": [
    {
      "element-uri": "/api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4/
        virtualization-host-storage-groups/92560990-6aa4-11e0-b3b9-f0def10c03f4",
      "name": "$3390$"
    },
    {
      "element-uri": "/api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4/
        virtualization-host-storage-groups/92559f6e-6aa4-11e0-b3b9-f0def10c03f4",
      "name": "$3380$"
    },
    {
      "element-uri": "/api/virtualization-hosts/342d80e0-65ff-11e0-acfd-f0def10c03f4/
        virtualization-host-storage-groups/925337e2-6aa4-11e0-b3b9-f0def10c03f4",
      "name": "$FCP$"
    }
  ]
}
```

---

Figure 214. List Virtualization Host Storage Groups: Response

## Get Virtualization Host Storage Group Properties

The **Get Virtualization Host Storage Group Properties** operation retrieves the properties of a single Virtualization Host Storage Group object that is designated by its object ID and the object ID of the owning virtualization host.

### HTTP method and URI

```
GET /api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-groups/
    {virt-host-storage-group-id}
```

### URI variables

Variable	Description
{virt-host-id}	Object ID of the virtualization host
{virt-host-storage-group-id}	Element ID of the Virtualization Host Storage Group object for which properties are to be obtained

## Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Virtualization Host Storage Group object as defined in “Data model” on page 411. Field

names and data types in the JSON object are the same as the property names and data types defined in the data model.

## Description

The **Get Virtualization Host Storage Group Properties** operation returns the current properties for the Virtualization Host Storage Group object specified by *{virt-host-storage-group-id}* for the virtualization host specified by *{virt-host-id}*.

On successful execution, all of the current properties as defined by the data model for the Virtualization Host Storage Group object are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing Virtualization Host object and the API user must have object-access permission to it. Furthermore, the URI path must designate an existing Virtualization Host Storage Group object. If any of these conditions are not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Manage Storage Resources** task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Manage Storage Resources** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 414.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{virt-host-id}</i> ) does not designate an existing Virtualization Host object, or the API user does not have object access permission to it, or it is not of a type for which this operation is supported. The object ID in the URI ( <i>{virt-host-storage-group-id}</i> ) does not designate an existing Virtualization Host Storage Group object for the specified virtualization host.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/virtualization-hosts/0e4a5d94-a8c1-11e0-9492-00262df332b3/
    virtualization-host-storage-groups/305cf1fe-a8cf-11e0-9a45-00262df332b3 HTTP/1.1
x-api-session: 3gcd77g1emvq81dlmwx8i4fwm4udx1by6i2auls4r6g529p1
```

---

Figure 215. Get Virtualization Host Storage Group Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Tue, 06 Dec 2011 14:34:05 GMT
content-type: application/json;charset=UTF-8
content-length: 466
{
  "class": "virtualization-host-storage-group",
  "description": "$FCP$",
  "element-id": "305cf1fe-a8cf-11e0-9a45-00262df332b3",
  "element-uri": "/api/virtualization-hosts/0e4a5d94-a8c1-11e0-9492-00262df332b3/
    virtualization-host-storage-groups/305cf1fe-a8cf-11e0-9a45-00262df332b3",
  "free-space": [],
  "name": "$FCP$",
  "parent": "/api/virtualization-hosts/0e4a5d94-a8c1-11e0-9492-00262df332b3",
  "virtualization-host": "/api/virtualization-hosts/0e4a5d94-a8c1-11e0-9492-00262df332b3"
}
```

---

Figure 216. Get Virtualization Host Storage Group Properties: Response

## List Virtualization Host Storage Resources in a Virtualization Host Storage Group

The **List Virtualization Host Storage Resources in a Virtualization Host Storage Group** operation lists the virtualization host storage resources that are members of the specified virtualization host storage group.

### HTTP method and URI

**POST** /api/virtualization-hosts/{*virt-host-id*}/operations/list-virtualization-host-storage-resources-in-group

In this request, the URI variable {*virt-host-id*} is the object ID of the virtualization host that owns the virtualization host storage group.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
virtualization-host-storage-group	String/URI	Optional	Canonical URI path of the virtualization host storage group whose members are to be listed.  If none is specified, the operation returns information for all virtualization host storage groups.



## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
virtualization-host-storage-resources	Array of objects	Array of virtualization-host-storage-resource-basic-info objects is returned, described in the next table. If no virtualization host storage resources are to be returned, an empty array is provided.

Each virtualization-host-storage-resource-basic-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Virtualization Host Storage Resource object
name	String	The <b>name</b> property of the associated Storage Resource object
type	String Enum	The <b>type</b> property of the associated Storage Resource object

## Description

The **List Virtualization Host Storage Resources in a Virtualization Host Storage Group** operation returns basic information about each virtualization host storage resource in the virtualization host storage group specified by *{virtualization-host-storage-group}* for the virtualization host specified by *{virt-host-id}*. If *{virtualization-host-storage-group}* is omitted, information from all of the virtualization host's virtualization host storage group is returned. The object URI and other basic properties are provided for each virtualization host storage resource.

A set of basic properties is returned for each virtualization host storage resource. See the virtualization-host-storage-resource-basic-info object definition.

The URI path must designate an existing Virtualization Host object and the API user must have object-access permission to it. If any of these conditions is not met, status code 404 (Not Found) is returned. If a URI is specified in the request body and it does not identify a virtualization host storage group for the specified virtualization host, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Manage Storage Resources** action; otherwise, status code 403 (Forbidden) is returned.

If there are no virtualization host storage resources in the target virtualization host storage group(s), an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Manage Storage Resources** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission to this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{virt-host-id}</i> ) does not designate an existing Virtualization Host object, or the API user does not have object access permission to it, or it is not of a type for which this operation is supported.
	2	The URI specified in the request body does not identify a virtualization host storage group for the specified virtualization host.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Add Virtualization Host Storage Resource to Virtualization Host Storage Group

The **Add Virtualization Host Storage Resource to Virtualization Host Storage Group** operation adds a specified virtualization host storage resource to the appropriate virtualization host storage group.

### HTTP method and URI

**POST** `/api/virtualization-hosts/{virt-host-id}/operations/add-virtualization-host-storage-resource-to-group`

In this request, the URI variable *{virt-host-id}* is the object ID of the virtualization host that owns the virtualization host storage group.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
virtualization-host-storage-resource	String/URI	Optional	Canonical URI path of the Virtualization Host Storage Resource object to be added

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the virtualization host storage group to which the virtualization host storage resource was added in the form <code>/api/virtualization-hosts/{virt-host-id}/virtualization-host-storage-groups/{virt-host-storage-group-id}</code> .

Field name	Type	Description
element-id	String	Element ID of the Virtualization Host Storage Group object. This is the <i>{virt-host-storage-group-id}</i> portion of the URI path provided by the <b>element-uri</b> field.

## Description

The **Add Virtualization Host Storage Resource to Virtualization Host Storage Group** operation adds a specified virtualization host storage resource to the appropriate virtualization host storage group for the virtualization host specified by *{virt-host-id}*.

On successful execution, the virtualization host storage resource identified in “Request body contents” on page 418 has been added to the appropriate virtualization host storage group, and HTTP status code 200 (OK) is returned. “Response body contents” on page 418 identifies the virtualization host storage group to which the virtualization host storage resource was added.

On successful execution, a property-change notification is issued asynchronously to this operation. See “Notifications” on page 421 for more information.

The URI path must designate an existing Virtualization Host object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. If the URI in the request body does not identify a virtualization host storage resource for the specified virtualization host, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Add Storage Resource to Group** task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Add Storage Resource to Group** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID <i>{virt-host-id}</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it, or it is not of a type for which this operation is supported.
	147	The URI in the request body ( <i>{virtualization-host-storage-resource}</i> ) does not designate an existing Virtualization Host Storage Resource object for the specified virtualization host.

HTTP error status code	Reason code	Description
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Remove Virtualization Host Storage Resource from Virtualization Host Storage Group

The **Remove Virtualization Host Storage Resource from Virtualization Host Storage Group** operation removes a specified virtualization host storage resource from the appropriate virtualization host storage group.

### HTTP method and URI

**POST** `/api/virtualization-hosts/{virt-host-id}/operations/remove-virtualization-host-storage-resource-from-group`

In this request, the URI variable `{virt-host-id}` is the object ID of the virtualization host that owns the virtualization host storage resource.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
virtualization-host-storage-resource	String/URI	Required	Canonical URI path of the Virtualization Host Storage Resource object to be removed

### Description

The **Remove Virtualization Host Storage Resource from Virtualization Host Storage Group** operation removes a specified virtualization host storage resource from the appropriate virtualization host storage group for the virtualization host specified by `{virt-host-id}`.

On successful execution, the virtualization host storage resource identified in “Request body contents” has been removed from the appropriate virtualization host storage group, and HTTP status code 204 (No Content) is returned and no response body is provided.

On successful execution, a property-change notification is issued asynchronously to this operation. See “Notifications” on page 421 for more information.

The URI path must designate an existing Virtualization Host object, and the API user must have object-access permission to it. Furthermore, the request body must designate an existing Virtualization Host Storage Resource object for the specified virtualization host. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action access permission to the **Remove Storage Resource from Group** task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the hosting object of the virtualization host specified in the request URI
- Action/task permission to the **Remove Storage Resource from Group** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	API user does not have action permission to this operation.
404 (Not Found)	1	The object ID in the URI <i>{virt-host-id}</i> does not designate an existing Virtualization Host object, or the API user does not have object access permission to it, or it is not of a type for which this operation is supported.
	147	The URI in the request body ( <i>{virtualization-host-storage-resource}</i> ) does not designate an existing Virtualization Host Storage Resource object for the specified virtualization host.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.
	100	The request could not be processed because the HMC is unable to communicate with a z/VM virtualization host via SMAPI.
	101	The request could not be processed because a command executed on a z/VM virtualization host via SMAPI failed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Notifications

Changes to properties of a virtualization host storage group are reflected via a property change notification designating the virtualization host as the managed object and the virtualization host storage group as the element object. The standard property change notification fields (old value, new value and property name) for each changed property are provided in the notification.

The addition or removal of a virtualization host storage resource to/from a virtualization host storage group is reflected via a property change notification designating the virtualization host as the managed object and the virtualization host storage group as the element object. The changed property identified in the notification is the **virtualization-host-storage-resources** property. Note that this is not included as a property of a virtualization host or a virtualization host storage group in their respective data models; this property name is only used in the context of these property change notifications. When the notification is due to the addition of a virtualization host storage resource, the new value in the notification is the URI path of the added virtualization host storage resource and the old value is null. When the notification is due to the removal of a virtualization host storage resource, the old value in the notification is the URI path of the removed virtualization host storage resource and the new value is **null**. Thus, unlike other property change notifications, these notifications only identify the delta to the set of

members of the virtualization host storage group rather than the complete new and old membership sets. In that sense, they are more like inventory change notifications.

## Inventory service data

Information about the storage groups associated with a virtualization host can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for Virtualization Host Storage Group objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**virtualization-host**" are to be included. An entry for a particular virtualization host storage group is included only if the API user has object-access permission to the owning virtualization host.

## Usage notes

- The group to which a virtualization host storage resource is added is determined by the type of the resource. Specifically, FCP resources are added to a group named "\$FCP\$", 3380 resources are added to the "\$3380\$" group, and 3390 resources are added to the "\$3390\$" group. The user may not specify any attributes of these groups (e.g., name or description). Thus, there are no writeable properties of a group.
- It is not possible for a virtual disk to span multiple physical storage devices. Thus, the largest virtual disk that may be created in the virtualization host storage group is limited to the largest free space on any single virtualization host storage resource in the group. The information describing all such areas of free space may be useful when defining, or planning to define, multiple virtual disks in the group.

## Chapter 12. Virtual network management

- In a z Systems ensemble, the Intra-Ensemble Data Network (IEDN) provides the physical, layer 2 network to which all z Systems CECs and the blades in the zBX are attached. The zManager Manage Virtual Network tasks are used to provision the IEDN into virtual networks, and to manage those virtual networks. In an ensemble, a virtual network is defined by a name, VLAN ID, and zero or more associated network host interfaces of the following types: virtual server, optimizers, and Top-of-Rack (TOR) switch ports that attach ISAOPT and external routers. The ensemble has a default virtual network. This virtual network cannot be deleted, although its VLAN ID can be changed.

To communicate on the IEDN, a network host's network interfaces must be associated with a virtual network that is managed by zManager. Note that the different types of network hosts have different requirements for defining and configuring network interfaces. For example, part of the process of defining a virtual server is to create its virtual network interfaces and associate them with a virtual network. The same is true for most optimizers. In the case of ISAOPT, the association to a virtual network is a special case. For ISAOPT, the TOR port that attaches to the coordinator blade's access switch provides the virtual network interface attachment. In cases where the TOR's external ports are attached to external networking equipment, such as a router, these TOR ports are required to join one or more virtual networks defined by zManager.

### Virtual network management operations summary

The following tables provide an overview of the operations provided.

Table 92. Virtual network management: operations summary

Operation name	HTTP method and URI
"List Virtual Networks" on page 424	GET /api/ensembles/{ensemble-id}/virtual-networks
"Get Virtual Network Properties" on page 426	GET /api/virtual-networks/{virtual-network-id}
"Update Virtual Network Properties" on page 427	POST /api/virtual-networks/{virtual-network-id}
"Create Virtual Network" on page 430	POST /api/ensembles/{ensemble-id}/virtual-networks
"Delete Virtual Network" on page 432	DELETE /api/virtual-networks/{virtual-network-id}
"List Members of a Virtual Network" on page 434	GET /api/virtual-networks/{virtual-network-id}/host-vnics

Table 93. Virtual network management: URI variables

Variable	Description
{ensemble-id}	Object ID of an Ensemble object
{virtual-network-id}	Object ID of a Virtual Network

### Virtual Network object

- A Virtual Network object represents a single z Systems virtual network.

## Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, with the following type-specific specialization:

Table 94. Virtual Network object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/URI	The canonical URI path for a Virtual Network object is of the form <code>/api/virtual-networks/{virtual-network-id}</code> where <code>{virtual-network-id}</code> is the value of the <b>object-id</b> property of the Virtual Network object.
<b>parent</b>	—	String/URI	The parent of a Virtual Network object is an ensemble object.
<b>class</b>	—	String	The class of a Virtual Network object is " <b>virtual-network</b> ".
<b>name</b>	(w)(pc)	String (1-32)	The name of the Virtual Network object as known by zManager. This name must be unique across virtual networks. If the virtual network is the default virtual network, then the name cannot be changed. The name provided must be between 1 and 32 characters.

## Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties:

Table 95. Virtual Network object: class specific additional properties

Name	Qualifier	Type	Description
<b>vlan-id</b>	(w)(pc)	Integer	The VLAN ID assigned to this virtual network. Valid VLAN IDs are in the range of 10-1030.
<b>is-default</b>	—	Boolean	This value is true when the virtual network is the default virtual network.
<b>has-members</b>	(pc)	Boolean	This value is true when the virtual network has member network host interfaces attached to it.

## List Virtual Networks

Use the **List Virtual Networks** operation lists the virtual networks managed by the HMC.

### HTTP method and URI

**GET** `/api/ensembles/{ensemble-id}/virtual-networks`

In this request, the URI variable `{ensemble-id}` is the object ID of an Ensemble object.

### Query parameters

Name	Type	Rqd/Opt	Description
<code>is-default</code>	Boolean	Optional	Filter pattern to limit returned objects to those that have a matching <code>is-default</code> property. There is only one default virtual network, and it is always defined, so when <code>is-default=true</code> is specified, the response will be an array with one element.
<code>name</code>	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <code>name</code> property. If a match is found, the response will be an array with one element. If no match is found, the response will be an empty array.



## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
virtual-networks	Array of objects	Array of nested virtual-network-info objects as described in the next table.

Each nested virtual-network-info object contains the following:

Field name	Type	Description
object-uri	String/ URI	Canonical URI path of the Virtual Network object, in the form <code>/api/virtual-networks/{virtual-network-id}</code> .
name	String	Display name of the Virtual Network object.

## Description

This operation lists the virtual networks that are in the specified ensemble. The response body content is an array of one or more URI's that represent each virtual network. Each ensemble always has a single default virtual network with the **name** of "Default".

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the target Ensemble object
- Action/task permission to the **Manage Virtual Networks** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	1	The request included an unrecognized or unsupported query parameter.
403 (Forbidden)	1	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the Manage Virtual Networks task is required.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type, or designates a resource for which the API user does not have object-access permission.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Example HTTP interaction

---

```
GET /api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334/virtual-networks HTTP/1.1
x-api-session: 5tjfpd6i7dfria4i52czubs5ptly4fqig2gm3mkluph6ebga
```

---

Figure 217. List Virtual Networks: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 05:11:15 GMT
content-type: application/json;charset=UTF-8
content-length: 316
{
  "virtual-networks": [
    {
      "name": "Default",
      "object-uri": "/api/virtual-networks/f920171e-03f2-11e1-8e8e-0010184c8334"
    },
    {
      "name": "SS-Web-Store-Network",
      "object-uri": "/api/virtual-networks/e58564e0-1723-11e1-aea4-0010184c8334"
    }
  ]
}
```

---

Figure 218. List Virtual Networks: Response

## Get Virtual Network Properties

Use the **Get Virtual Network Properties** operation retrieves the properties of a single Virtual Network object that is designated by the *{virtual-network-id}*.

### HTTP method and URI

```
GET /api/virtual-networks/{virtual-network-id}
```

In this request, the URI variable *{virtual-network-id}* is the object ID of the Virtual Network object for which properties are to be obtained.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the Virtual Network object as defined in the Data Model section above. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

This operation returns the current properties for the Virtual Network object. In an ensemble, there is at least one virtual network. This is the default virtual network with the **name "Default"**.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the target Virtual Network
- Action/task permission to the **Manage Virtual Networks** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 426.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Manage Virtual Networks</b> task is required.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type, or designates a resource for which the API user does not have object-access permission.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/virtual-networks/e58564e0-1723-11e1-aea4-0010184c8334 HTTP/1.1
x-api-session: 5tjfp6ld6i7dfria4i52czubs5ptly4fqig2gm3mkluph6ebga
```

---

*Figure 219. Get Virtual Network Properties: Request*

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 05:11:15 GMT
content-type: application/json;charset=UTF-8
content-length: 368
{
  "class": "virtual-network",
  "description": "Spacely Sprockets Web Store Network",
  "has-members": true,
  "is-default": false,
  "is-locked": false,
  "name": "SS-Web-Store-Network",
  "object-id": "e58564e0-1723-11e1-aea4-0010184c8334",
  "object-uri": "/api/virtual-networks/e58564e0-1723-11e1-aea4-0010184c8334",
  "parent": "/api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334",
  "vlan-id": 1030
}
```

---

*Figure 220. Get Virtual Network Properties: Response*

## Update Virtual Network Properties

Use the **Update Virtual Network Properties** operation updates one or more of the writeable properties of a Virtual Network object.

## HTTP method and URI

**POST** /api/virtual-networks/{*virtual-network-id*}

In this request, the URI variable *{virtual-network-id}* is the object ID of the Virtual Network object for which properties are to be updated.

## Request body contents

The request body is expected to contain a JSON object that provides the new values of any writeable property that is to be updated by this operation. Field names and data types in this JSON object are expected to match the corresponding property names and data types defined by the data model for this object type. The JSON object can and should omit fields for properties whose values are not to be changed by this operation.

## Description

This operation updates writeable properties of the Virtual Network object specified by *{virtual-network-id}*.

If the **vlan-id** property is changed, and there are member attached hosts are not in the proper operating status to accept a virtual network change, then this request will fail with HTTP status code 409, and the response body will contain the list of URI's of the network hosts that were unable to accept a virtual network change. Upon failure, the virtual network's **vlan-id** is not changed, although requests for updates to other virtual network properties may have successfully occurred. Upon successful complete, the virtual network of the attached network hosts virtual network interfaces will be changed.

The request body is validated against the schema described in "Request body contents." If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

The request body does not need to specify a value for all writeable properties, but rather can and should contain fields for the properties to be updated. Object properties for which no input value is provided remain unchanged by this operation.

If the update changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the target Virtual Network object
- Action/task permission to the **Edit Virtual Network Properties** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	129	The default virtual network name of " <b>Default</b> " cannot be changed.
	131	The specified name or VLAN ID for a virtual network already exists. This includes specification of the name " <b>Default</b> ".
403 (Forbidden)	1	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Manage Virtual Networks</b> task is required.
404 (Not Found)	1	The object-id in the URI <i>{virtual-network-id}</i> does not designate an existing resource of the correct type, or designates a resource for which the API user does not have object-access permission.
409 (Conflict)	130	One or more of the network attached hosts that are members of this virtual network do not have the proper operating status to support changes to the virtual network at this time. The URIs of network hosts that were unable to accept the virtual network change are returned in the response body.  Retry the request.  Currently PowerVM and x Hyp servers have operating status restrictions on changes to a virtual network. See the network-adapter Data Model for details for virtual server status that allows network-adapter changes.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.
	2	The request could not be processed because the SE is not currently communicating with an element of a zBX needed to perform the requested operation.

On completion where the HTTP status code is 409 (Conflict), the standard error response body contains an error-details field that provides a list of network hosts that were unable to accept a virtual network change. The value of the error-details field is a nested object with the following fields:

Field name	Type	Description
<b>hosts-list</b>	Array of string/URI	Array of URIs that identify the network hosts that were unable to accept a virtual network change. The URIs for the network hosts are in the following forms:  Virtual server: <code>/api/virtual-servers/{virtual-server-id}</code>  Optimizer: <code>/api/blades/{blade-id}</code>

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/virtual-networks/e58564e0-1723-11e1-aea4-0010184c8334 HTTP/1.1
x-api-session: 5tjfp9ld6i7dfria4i52czubs5ptly4fqig2gm3mkluph6ebga
content-type: application/json
content-length: 103
{
  "description": "Spacely Sprockets Web Store Network",
  "name": "SS-Web-Store-Network",
  "vlan-id": 1030
}
```

---

Figure 221. Update Virtual Network Properties: Request

---

```
204 No Content
date: Fri, 25 Nov 2011 05:11:13 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 222. Update Virtual Network Properties: Response

## Create Virtual Network

Use the **Create Virtual Network** operation creates a new virtual network in the ensemble.

### HTTP method and URI

**POST** /api/ensembles/{ensemble-id}/virtual-networks

In this request, the URI variable {ensemble-id} is the object ID of the Ensemble object to which the Virtual Network is to be added.

### Request body contents

The request body is expected to contain a JSON object with writeable network properties of the Virtual Network object that will be used to create the virtual network:

Field name	Type	Rqd/Opt	Description
name	String	Required	The name of the Virtual Network object as known by zManager. This name must be unique across all virtual networks, and cannot be the value <b>"Default"</b> as this name is reserved as the zManager-defined name of the default virtual network. The name provided must be between 1 and 32 characters.
description	String	Optional	Display description of the Virtual Network object.
vlan-id	Integer	Required	The VLAN ID assigned to this virtual network. Valid VLAN IDs are in the range of 10-1030. This value must be unique across all virtual networks.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the Virtual Network object, in the form <code>/api/virtual-networks/{virtual-network-id}</code> .

## Description

This operation creates a new virtual network in the ensemble specified by *{ensemble-id}*. A virtual network with the same **vlan-id** or **name** must not already exist in the ensemble.

On successful execution of this operation the virtual network is created using the inputs as specified by the input fields of the request body. The URI of the new virtual network is provided in the response body and in a **Location** response header as well.

The request body is validated against the schema described in the Request Body Contents section. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the target Ensemble object
- Action/task permission to the **New Virtual Networks** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	0	This error indicates maximum number of virtual networks exceeded or unavailable resources.
	131	The specified name or VLAN ID for a virtual network already exists. This includes specification of the name <b>"Default"</b> .
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform the requested action.
404 (Not Found)	1	The object-id in the URI <i>{ensemble-id}</i> does not designate an existing resource of the correct type, or designates a resource for which the API user does not have object-access permission.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334/virtual-networks HTTP/1.1
x-api-session: 5tjfpd6i7dfria4i52czubs5ptly4fqig2gm3mkluph6ebga
content-type: application/json
content-length: 71
{
  "description": "New Network",
  "name": "SS-New-Network",
  "vlan-id": 11
}
```

---

Figure 223. Create Virtual Network: Request

---

```
201 Created
server: zSeries management console API web server / 1.0
location: /api/virtual-networks/e58564e0-1723-11e1-aea4-0010184c8334
cache-control: no-cache
date: Fri, 25 Nov 2011 05:11:13 GMT
content-type: application/json;charset=UTF-8
content-length: 75
{
  "object-uri": "/api/virtual-networks/e58564e0-1723-11e1-aea4-0010184c8334"
}
```

---

Figure 224. Create Virtual Network: Response

## Delete Virtual Network

Use the **Delete Virtual Network** operation removes a virtual network from an ensemble.

### HTTP method and URI

```
DELETE /api/virtual-networks/{virtual-network-id}
```

In this request, the URI variable *{virtual-network-id}* is the object ID of the Virtual Network object to be deleted.

### Description

On successful completion, this operation removes the virtual network specified by the URI from the ensemble. Any network attached hosts that are members of this virtual network and have the proper operating status to accept a virtual network change will be removed from the virtual network. Currently only PowerVM and x Hyp servers require the proper operating status to accept a virtual network change. If attached hosts cannot be removed from the virtual network, then this request will fail with an HTTP status code of 409, and the response body will contain the list of the URIs of the network hosts that were unable to accept a virtual network change. Upon failure, no network attached hosts are removed from the virtual network, and the virtual network is not deleted.

Note that the default virtual network, **"Default"**, cannot be deleted from the ensemble.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the target Ensemble object
- Action/task permission to the **New Virtual Networks** task.



## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	132	The default virtual network cannot be deleted.
404 (Not Found)	1	The object-id in the URI does not designate an existing resource of the correct type, or designates a resource for which the API user does not have object-access permission.
409 (Conflict)	130	One or more of the network attached hosts that are members of this virtual network do not have the proper operating status to support a change to the virtual network at this time. The URIs of network hosts that were unable to accept the virtual network change are returned in the response body.  Retry the request.  Currently PowerVM and x Hyp type servers have operating status restrictions on changes to a virtual network. See the network-adapter Data Model for details for virtual server status that allows network-adapter changes.
503 (Service Unavailable)	1	Communication between the HMC and SE is unavailable. Please retry the request.
	2	The request could not be processed because the SE is not currently communicating with an element of a zBX needed to perform the requested operation.

On completion where the HTTP status code is 409 (Conflict), the standard error response body contains an error-details field that provides a list of network hosts that were unable to accept a virtual network change. The value of the error-details field is a nested object with the following fields:

Field name	Type	Description
<b>hosts-list</b>	Array of string/URI	Array of URIs that identify the network hosts that were unable to accept a virtual network change. The URIs for the network hosts are in the following forms:  Virtual server: <code>/api/virtual-servers/{virtual-server-id}</code>  Optimizer: <code>/api/blades/{blade-id}</code>

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
DELETE /api/virtual-networks/e58564e0-1723-11e1-aea4-0010184c8334 HTTP/1.1
x-api-session: 5tjfpd6i7dfria4i52czubs5ptly4fqig2gm3mkluph6ebga
```

---

Figure 225. Delete Virtual Network: Request

---

```
204 No Content
date: Fri, 25 Nov 2011 05:11:15 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

```
<No response body>
```

---

Figure 226. Delete Virtual Network: Response

## List Members of a Virtual Network

Use the **List Members of a Virtual Network** operation lists the members of a Virtual Network.

### HTTP method and URI

```
GET /api/virtual-networks/{virtual-network-id}/host-vnics
```

In this request, the URI variable *{virtual-network-id}* is the object ID of the Virtual Network for which members are to be listed.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
host-vnics	Array of objects	Array of nested member-info objects, each element of which identifies an attached network host's network interface that is a member of this virtual network. The nested member-info object is described in the next table. If the virtual network has no network interfaces associated with it then an empty array is returned.

Each nested member-info object contains the following:

Field name	Type	Description
object-uri	String/URI	<p>A URI of an object or element that identifies an attached network host's network interface that is a member of this virtual network. A network interface can be one of the following: A network adapter of a virtual server or optimizer, or a Top-of-Rack-Switch (TOR) port that is attached to an ISAOPT optimizer or to an external router. The URI's for the network hosts/interfaces are:</p> <p>Virtual server:</p> <pre>/api/virtual-servers/{virtual-server-id}/network-adapters/{network-adapter-id}</pre> <p>Optimizer:</p> <pre>/api/blades/{blade-id}/iedn-interfaces/{iedn-interface-id}</pre> <p>TOR:</p> <pre>/api/zbxes/{zbx-id}/top-of-rack-switches/{tor-id}/ports/{port-id}</pre>

## Description

The List Members of a Virtual Network operation lists the members of a Virtual Network. The members of a virtual network are network attached hosts' network interfaces which include the following types: virtual servers' virtual network interfaces, optimizer virtual network interfaces, and Top-of-Rack (TOR) switch network interfaces (ports).

On successful execution, this operation returns an array of nested objects that identify the members of the Virtual Network specified by the request URI. This array may be empty if the virtual network has no associated network hosts. In some cases, a network host may have multiple interfaces associated with this virtual network.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the target Virtual Network object
- Object access to the specific member attached network host object
- Action/task permission to the **Manage Virtual Networks** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents" on page 434.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The user under which the API request was authenticated is not authorized to perform the requested operation. Permission to the <b>Manage Virtual Networks</b> task is required.

HTTP error status code	Reason code	Description
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type, or designates a resource for which the API user does not have object-access permission.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/virtual-networks/e58564e0-1723-11e1-aea4-0010184c8334/host-vnics HTTP/1.1
x-api-session: 5tjfpgd6i7dfria4i52czubs5ptly4fqig2gm3mkluph6ebga
```

---

Figure 227. List Members of a Virtual Network: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 05:11:15 GMT
content-type: application/json;charset=UTF-8
content-length: 274
{
  "host-vnics": [
    {
      "object-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af/
network-adapters/596dd87c-0db7-11e1-9251-f0def14b63af"
    },
    {
      "object-uri": "/api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af/
network-adapters/582b5d0e-0db7-11e1-b1f1-f0def14b63af"
    }
  ]
}
```

---

Figure 228. List Members of a Virtual Network: Response

## Inventory service data

Information about the Virtual Networks managed by the HMC can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for Virtual Network objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**virtual-network**" are to be included. An entry for a particular virtual network is included only if the API user has object-access permission to that object.

For each Virtual Network object to be included, the inventory response array includes entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the Get Virtual Network Properties operation. That is, the data provided is the same as would be provided if a Get Virtual Network Properties operation were requested targeting this object.

## Sample inventory data

The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single virtual network. This object would appear as one array entry in the response array:

---

```
{
  "class": "virtual-network",
  "description": "Network for servers of the Shimmer order processing application.",
  "has-members": false,
  "is-default": false,
  "name": "Shimmer Order Processing Network",
  "object-id": "e3a366f4-95e6-11e0-85c7-000c29bb873c",
  "object-uri": "/api/virtual-networks/e3a366f4-95e6-11e0-85c7-000c29bb873c",
  "parent": "/api/ensembles/890b6df2-93a4-11e0-887c-000c29bb873c",
  "vlan-id": 1008
}
```

---

*Figure 229. Virtual Network object: Sample inventory data*



---

## Chapter 13. Workload resource group management

This chapter describes workload resource group management APIs, and provides a data model and operations for each of the following: workload objects, Performance Policy objects, and performance management reports.

---

### Overview

Beginning with the IBM z Enterprise System, the z Systems family extends performance management capability to both traditional z Systems and BladeCenter hardware environments. For multi-tier applications that span z Systems hardware and BladeCenter hardware, this extended capability enables dynamic adjustments to processor (CPU) allocations to ensure that those applications are provided with sufficient resources. To manage hardware resources in the z Systems environment, you group these resources into workload resource groups and define performance policies for them. zManager uses these policy definitions to manage the resources for each workload resource group in an ensemble.

In computing terminology, the usual definition of the term *workload* is the amount of application processing that a computer performs at a given time, and this processing usually includes a specific number of connected users interacting with the running applications. This amount of work often serves as a benchmark for computer performance.

In contrast, a z Systems workload resource group, which is sometimes referred to more simply as a “workload”, is a customer-defined collection to be tracked, managed, and reported as a unit that reflects a business goal or function. This collection consists of virtualized hardware rather than software resources. Throughout the z Systems documentation, the term *workload resource group* means a collection of logical constructs called *virtual servers* that perform a customer-defined collective purpose.

Initially, after z Systems hardware is installed and the ensemble is configured, all virtual servers are associated with the default workload resource group and its default performance policy. Workload administrators use the HMC to create and name the new workload resource groups, and to define which virtual servers are associated with this workload resource group. A workload's virtual servers can then be further divided into workload element groups, ensuring availability of each business function.

After a custom workload resource group is defined, an administrator can use ensemble tasks in the HMC to create performance policies that describe the workload resource group performance objectives and importance, and to create service classes that set priority for and classify work within a policy. These values govern performance management. Each workload resource group can have one or more performance policies that describe the performance objectives and importance. Each performance policy has service classes that set the priority of and classify resources within the policy. Only one performance policy within a specific workload resource group can be active at any given time; having multiple policies with different defined goals gives an administrator the ability to quickly change priorities by merely activating a different policy. zManager uses the active performance policy and its service classes to manage how physical resources are applied to the virtual servers associated with the workload resource group.

The Web Services API provides the same workload resource group management capabilities as those provided through HMC tasks, including:

- Creating, updating and deleting workload resource groups, performance policies, or service classes.
- Activating a performance policy.

In addition, zManager also continuously collects performance-management data that is available through a variety of reports that you view through the HMC. By setting a time interval for any report, you can query historical performance data as it applied over that time period. zManager keeps only the last 36

hours of collected performance data, so that limit defines the maximum duration for reports. Through HMC tasks, these various reports present performance data in tables, charts or graphs. In many cases, you can request one report and select data within that report to request more detailed performance data that is available in another report type. For example, through the HMC you can select the **Workloads Report** task to list all workload resource groups defined in the ensemble. From that list, you can select a specific workload resource group to view a **Virtual Server Report** that lists all virtual servers defined to that specific workload resource group.

The Web Services API provide the same performance reports as those provided through HMC tasks, returning the same data and including the capability to query time intervals. With the zManager API, the caller requests a specific report through an operations call to generate and return the specific type of report. Data returned for items in one report can be used in requests for subsequent drill-down report generation requests for other report types.

For more information about workload resource group and performance management, including sample reports as they are returned through the HMC, see the *z Systems Ensemble Workload Resource Group Management Guide, GC27-2629*.

## Workload resource group operations summary

The following tables list the workload resource group operations that are provided through the Web Services API.

Table 96. Workload resource group management: operations summary

Operation name	HTTP method and URI path
"List Workload Resource Groups of an Ensemble" on page 448	GET /api/ensembles/{ensemble-id}/workload-resource-groups
"Get Workload Resource Group Properties" on page 450	GET /api/workload-resource-groups/{workload-id}
"Create Workload Resource Group" on page 452	POST /api/ensembles/{ensemble-id}/workload-resource-groups
"Delete Workload Resource Group" on page 455	DELETE /api/workload-resource-groups/{workload-id}
"Update Workload Resource Group" on page 456	POST /api/workload-resource-groups/{workload-id}
"List Virtual Servers of a Workload Resource Group" on page 458	GET /api/workload-resource-groups/{workload-id}/virtual-servers
"Add Virtual Server to a Workload Resource Group" on page 461	POST /api/workload-resource-groups/{workload-id}/operations/add-virtual-server
"Remove Virtual Server from a Workload Resource Group" on page 463	POST /api/workload-resource-groups/{workload-id}/operations/remove-virtual-server
"List Groups of Virtual Servers of a Workload Resource Group" on page 465	GET /api/workload-resource-groups/{workload-id}/virtual-server-groups
"Add Group of Virtual Servers to a Workload Resource Group" on page 467	POST /api/workload-resource-groups/{workload-id}/operations/add-virtual-server-group
"Remove Group of Virtual Servers from a Workload Resource Group" on page 469	POST /api/workload-resource-groups/{workload-id}/operations/remove-virtual-server-group
"List Workload Element Groups of a Workload Resource Group" on page 471	GET /api/workload-resource-groups/{workload-id}/workload-element-groups



Table 96. Workload resource group management: operations summary (continued)

Operation name	HTTP method and URI path
“Add Workload Element Group to a Workload Resource Group” on page 472	POST /api/workload-resource-groups/{workload-id}/operations/add-workload-element-group
“Remove Workload Element Group from a Workload Resource Group” on page 474	POST /api/workload-resource-groups/{workload-id}/operations/remove-workload-element-group
“List Performance Policies” on page 481	GET /api/workload-resource-groups/{workload-id}/performance-policies
“Get Performance Policy Properties” on page 483	GET /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}
“Create Performance Policy” on page 486	POST /api/workload-resource-groups/{workload-id}/performance-policies
“Delete Performance Policy” on page 488	DELETE /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}
“Update Performance Policy” on page 490	POST /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}
“Activate Performance Policy” on page 493	POST /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}/operations/activate
“Import Performance Policy” on page 494	POST /api/workload-resource-groups/{workload-id}/operations/import-performance-policy
“Export Performance Policy” on page 496	POST /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}/operations/export
“Generate Workload Resource Groups Report (Performance Management)” on page 500	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-workload-resource-groups-report
“Generate Workload Resource Group Performance Index Report” on page 504	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-workload-resource-group-performance-index-report
“Generate Workload Resource Group Resource Adjustments Report” on page 507	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-workload-resource-group-resource-adjustments-report
“Generate Virtual Servers Report” on page 511	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-virtual-servers-report
“Generate Virtual Server CPU Utilization Report” on page 515	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-virtual-server-cpu-utilization-report
“Generate Virtual Server Resource Adjustments Report” on page 518	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-virtual-server-resource-adjustments-report
“Generate Hypervisor Report” on page 523	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-hypervisor-report
“Generate Hypervisor Resource Adjustments Report” on page 531	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-hypervisor-resource-adjustments-report
“Generate Service Classes Report” on page 535	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-service-classes-report
“Generate Service Class Resource Adjustments Report” on page 538	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-service-class-resource-adjustments-report
“Generate Service Class Hops Report” on page 543	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-service-class-hops-report

Table 96. Workload resource group management: operations summary (continued)

Operation name	HTTP method and URI path
“Generate Service Class Virtual Server Topology Report” on page 548	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-service-class-virtual-server-topology-report
“Generate Load Balancing Report” on page 556	POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-load-balancing-report
“List Workload Element Groups of an Ensemble” on page 561	GET /api/ensembles/{ensemble-id}/workload-element-groups
“Get Workload Element Group Properties” on page 562	GET /api/workload-element-groups/{group-id}
“Create Workload Element Group” on page 563	POST /api/ensembles/{ensemble-id}/workload-element-groups
“Delete Workload Element Group” on page 566	DELETE /api/workload-element-groups/{group-id}
“Update Workload Element Group” on page 567	POST /api/workload-element-groups/{group-id}
“List Virtual Servers of a Workload Element Group” on page 568	GET /api/workload-element-groups/{group-id}/virtual-servers
“Add Virtual Server to a Workload Element Group” on page 569	POST /api/workload-element-groups/{group-id}/operations/add-virtual-server
“Remove Virtual Server from a Workload Element Group” on page 571	POST /api/workload-element-groups/{group-id}/operations/remove-virtual-server
“List Availability Policies” on page 575	GET /api/workload-resource-groups/{workload-id}/availability-policies
“Get Availability Policy Properties” on page 577	GET /api/workload-resource-groups/{workload-id}/availability-policies/{policy-id}
“Create Availability Policy” on page 578	POST /api/workload-resource-groups/{workload-id}/availability-policies
“Delete Availability Policy” on page 580	DELETE /api/workload-resource-groups/{workload-id}/availability-policies/{policy-id}
“Update Availability Policy” on page 581	POST /api/workload-resource-groups/{workload-id}/availability-policies/{policy-id}
“Activate Availability Policy” on page 583	POST /api/workload-resource-groups/{workload-id}/availability-policies/{policy-id}/operations/activate
“Generate Workload Resource Groups Report (Ensemble Availability Management)” on page 586	POST /api/ensembles/{ensemble-id}/availability-management/operations/generate-workload-resource-groups-report
“Generate Workload Resource Group Availability Status Report” on page 589	POST /api/ensembles/{ensemble-id}/availability-management/operations/generate-workload-resource-group-availability-status-report
“Generate Virtual Servers Report (Ensemble Availability Management)” on page 594	POST /api/ensembles/{ensemble-id}/availability-management/operations/generate-virtual-servers-report
“Generate Availability Status Report” on page 597	POST /api/ensembles/{ensemble-id}/availability-management/operations/generate-availability-status-report
“Get Performance Management Velocity Level Range Mappings” on page 599	GET /api/ensembles/{ensemble-id}/performance-management/velocity-level-range-mappings

Table 97. Workload management: URI variables

Variable	Description
<i>{ensemble-id}</i>	Object ID (UUID) of an ensemble object
<i>{workload-id}</i>	Object ID (UUID) of a workload object
<i>{policy-id}</i>	ID (UUID) of a Performance Policy object within a workload resource group
<i>{group-id}</i>	Object ID (UUID) of a Workload Element Group object.

## Workload Resource Group object

A *workload object* is a managed object representing a group of virtual servers in an ensemble. Virtual servers can be managed in workload resource groups either by adding or removing them directly to or from the workload resource group, or through user-defined managed object groups or workload element groups. By adding a group to a workload resource group, all of the virtual servers in that group are implicitly added to the workload resource group. By adding or removing a virtual server to or from a group, that virtual server is implicitly added to or removed from the workload resource group. A virtual server can belong to more than one custom workload resource group. Any virtual server that does not belong to a custom workload resource group is placed in the default workload resource group.

A workload resource group contains a default performance policy as well as any custom performance policies created for it. Exactly one performance policy is active for a specific workload resource group at any given time, and is applied to the virtual servers in that workload resource group. The default workload resource group also has a default performance policy that is always active.

## Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, with the type-specific specialization described in the following tables. Note that this object does not maintain the concept of an operational status, and therefore does not provide the operational-status-related properties.

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

Table 98. Workload object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/URI	The canonical URI path for a workload object is of the form <code>/api/workload-resource-groups/{workload-id}</code> where <i>{workload-id}</i> is the value of the <b>object-id</b> property of the workload object.
<b>object-id</b>	—	String (36)	The unique identifier for the workload resource group instance. This identifier is in the form of a UUID.
<b>parent</b>	—	String/URI	The parent of a workload resource group is conceptually its owning ensemble, so the parent value is the canonical URI path for the ensemble.
<b>class</b>	—	String	The class of a workload object is <b>"workload-resource-group"</b> .
<b>name</b>	(w) (pc)	String (1-64)	The display name specified for the workload resource group, which can be up to 64 characters made up of alphanumeric characters, blanks, periods, underscores, or dashes. Names must start with an alphabetic character and end with an alphabetic character, numeric character, underscore, period, or dash. Names must be unique to other existing workload resource groups in the ensemble.

Table 98. Workload object: base managed object properties specializations (continued)

Name	Qualifier	Type	Description of specialization
<b>description</b>	(w) (pc)	String (0-256)	Arbitrary text describing the workload resource group in up to 256 characters.
<b>status</b>	(sc)	String Enum	The current status of the workload resource group, which must be one of the following values:  <b>"compliant"</b> The contributors to workload compliance are all currently compliant.  <b>"not-compliant"</b> One or more of the contributors to the workload compliance are currently not-compliant.
<b>additional-status</b>	(sc)	String Enum	This property is not supported for workload resource groups and will not be included in its data model.
<b>acceptable-status</b>	(w)(pc)	Array of String Enum	The set of status values that the workload resource group can be in and be considered to be in an acceptable (not alert causing) state. By default, this is <b>"compliant"</b> for workloads, and <b>"not-compliant"</b> , <b>"compliant"</b> for the default workload.

In addition to the properties defined through included schemas, this object includes the additional type-specific properties in Table 99.

Table 99. Workload object: type-specific properties

Name	Qualifier	Type	Description
<b>is-default</b>	—	Boolean	This value identifies the default workload object. It is true for the default workload resource group and false for all other (custom) workload resource group in the ensemble.
<b>category</b>	(w) (pc)	String (0-32)	Up to 32 characters used to categorize the workload object, often with other workload resource groups within the ensemble.
<b>virtual-servers</b>	(c) (pc)	Array of String/URI	The complete list of all virtual servers in the workload, each identified by its URI, including those directly placed as well as those placed due to membership in a group. (This list corresponds to the list provided by the <b>List Virtual Servers of a Workload Resource Group</b> operation.) If the workload contains no virtual servers, then an empty array is provided. The virtual servers provided in this list can change as a result of the Add and Remove Virtual Server or the Add and Remove Groups operations on the workload. <b>Note:</b> This property is not returned by the <b>Get Workload Resource Group Properties</b> operation.
<b>directly-added-virtual-servers</b>	(c) (pc)	Array of String/URI	The list of virtual servers that have been directly placed in the workload, each identified by its URI. If the workload contains no directly placed virtual servers, then an empty array is provided. This value can be modified through the Add and Remove Virtual Server operations on the workload. <b>Note:</b> This property is not returned by the <b>Get Workload Resource Group Properties</b> operation.

Table 99. Workload object: type-specific properties (continued)

Name	Qualifier	Type	Description
<b>virtual-server-groups</b>	(c) (pc)	Array of String/URI	The list of the user-defined managed object groups, each identified by its URI, in the workload whose child virtual servers automatically become members of the workload. If the workload has no groups, an empty array is provided. This value can be modified through the Add and Remove Group operations on the workload. <b>Note:</b> This property is not returned by the <b>Get Workload Resource Group Properties</b> operation.
<b>workload-element-groups</b>	(c) (pc)	Array of String/URI	The list of the workload element groups, each identified by its URI, in the workload whose child virtual servers automatically become members of the workload. If the workload has no element groups, an empty array is provided. This value can be modified through the <b>Add Workload Element Group to a Workload Resource Group</b> and <b>Remove Workload Element Group from a Workload Resource Group</b> operations on the workload as well as the <b>CreateWorkload Element Group</b> and <b>Delete Workload Element Group</b> operations.
<b>active-perf-policy</b>	(pc)	Object	A perf-policy-summary object, as described in Table 100 on page 447, of the active performance policy in the workload resource group. This value can be modified through the operation "Activate Performance Policy" on page 493.
<b>perf-activation-node-count</b>	(pc)	Integer (0-n)	The number of nodes (ensemble members) that have successfully activated the active performance policy for the workload resource group. This value is between 0 and the total number of ensemble members.
<b>perf-activation-status</b>	(pc)	String Enum	The status of the last performance policy activation. The possible values are: <b>"initializing"</b> The workload resource group performance function is being initialized and status is not yet known. <b>"in-progress"</b> Performance policy activation is in progress and any new activation request is rejected. <b>"active"</b> Performance policy activation is complete.
<b>default-perf-policy</b>	—	Object	A perf-policy-summary object, as described in Table 100 on page 447, of the default performance policy in the workload resource group.
<b>custom-perf-policies</b>	(c) (pc)	Array of objects	A list of perf-policy-summary objects, as described in Table 100 on page 447, of the user-defined performance policies in the workload resource group. If no custom policies exist, an empty array is provided. This value can be modified through the operations "Create Performance Policy" on page 486 and "Delete Performance Policy" on page 488.

Table 99. Workload object: type-specific properties (continued)

Name	Qualifier	Type	Description
<b>perf-status</b>	(pc)	String Enum	The performance status of the workload, determined by the importance and PI value of the active performance policy's service classes. Possible values are: <b>"goals-met"</b> The PI values of all service classes in the active policy are less than or equal to one. <b>"exposed"</b> The PI value of a service class in the active policy is greater than one and its importance is <b>"low"</b> or <b>"lowest"</b> . <b>"severe"</b> The PI value of a service class in the active policy is greater than one and its importance is <b>"medium"</b> . <b>"critical"</b> The PI value of a service class in the active policy is greater than one and its importance is <b>"high"</b> or <b>"highest"</b> . <b>"no-status"</b> Performance status of the workload cannot be calculated.
<b>compliant-perf-status</b>	(w)(pc)	Array of String Enum	An array of values that define what <b>perf-status</b> values cause the workload's status to be compliant. Possible values are: <b>"goals-met"</b> The workload's performance status is <b>"goals-met"</b> . <b>"exposed"</b> The workload's performance status is <b>"exposed"</b> . <b>"severe"</b> The workload's performance status is <b>"severe"</b> . <b>"critical"</b> The workload's performance status is <b>"critical"</b> . <b>"no-status"</b> The workload's performance status is <b>"no-status"</b> .  By default, a workload's compliant performance statuses include <b>"goals-met"</b> .
<b>active-avail-policy</b>	(pc)	Object	An avail-policy-summary object, as described in Table 101 on page 447, of the active availability policy in the workload. This value can be modified with the <b>Activate Availability Policy</b> operation.
<b>default-avail-policy</b>	—	Object	An avail-policy-summary object, as described in Table 101 on page 447, of the default availability policy in the workload.
<b>custom-avail-policies</b>	(c) (pc)	Array of objects	A list of avail-policy-summary objects, as described in Table 101 on page 447, of the user-defined availability policies in the workload. If no custom policies exist, then an empty array is provided. This value can be modified through the <b>Create Availability Policy</b> and <b>Delete Availability Policy</b> operations on the workload.

Table 99. Workload object: type-specific properties (continued)

Name	Qualifier	Type	Description
<b>avail-status</b>	(pc)	String Enum	The availability status of the workload, determined by the active availability policy and the availability status of the workload's element groups and virtual servers. The possible values are as follows: <b>"available"</b> The workload is considered <b>"available"</b> . <b>"exposed"</b> The workload is considered <b>"exposed"</b> . <b>"critical"</b> The workload is considered critically exposed. <b>"not-available"</b> The workload is considered <b>"not-available"</b> . <b>"not-supported"</b> The workload is considered <b>"not-supported"</b> .
<b>compliant-avail-status</b>	(w) (pc)	Array of String Enum	An array of values that define what <b>avail-status</b> values causes the workload's status to be compliant. The possible values are: <b>"available"</b> The workload's availability is <b>"available"</b> . <b>"exposed"</b> The workload's availability status is <b>"exposed"</b> . <b>"critical"</b> The workload's availability status is <b>"critical"</b> . <b>"not-available"</b> The workload's availability status is <b>"not-available"</b> .  By default, a workload's compliant availability status includes <b>"available"</b> .

Table 100. perf-policy-summary-object

Field name	Type	Description
<b>element-uri</b>	String/URI	Canonical URI path of the Performance Policy object.
<b>element-id</b>	String	The <b>element-id</b> of the Performance Policy object.
<b>name</b>	String	Display name of the Performance Policy object.
<b>activation-status</b>	String Enum	The activation status of the Performance Policy object. The value is one of the following: <b>"not-active"</b> The performance policy is not currently the active policy in the workload resource group. <b>"in-progress"</b> The performance policy is currently being activated in the workload resource group. <b>"active"</b> The performance policy is currently the active policy in the workload resource group and its activation has completed.

Table 101. avail-policy-summary-object

Field name	Type	Description
<b>element-uri</b>	String/URI	Canonical URI path of the Availability Policy object.
<b>element-id</b>	String	The <b>element-id</b> of the Availability Policy object.
<b>name</b>	String	Display name of the Availability Policy object.

Table 101. *avail-policy-summary-object* (continued)

Field name	Type	Description
activation-status	String Enum	The activation status of the Availability Policy object. The value is one of the following:  <b>"not-active"</b> The availability policy is not currently the active policy in the workload resource group.  <b>"active"</b> The availability policy is currently the active policy in the workload resource group.

## List Workload Resource Groups of an Ensemble

Use the **List Workload Resource Groups of an Ensemble** operation to list the workload resource groups within the target ensemble.

### HTTP method and URI

**GET** /api/ensembles/{ensemble-id}/workload-resource-groups

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble object for which you are requesting a list of workload resource groups.

### Query parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	A filter pattern (regular expression) that limits returned objects to those that have a matching <b>name</b> property. If matches are found, the response is an array with all workload resource groups that match. If a match is not found, the response is an empty array.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
workloads	Array of objects	Array of nested workload-info objects as described in the next table.

### workload-info object

Field name	Type	Description
object-uri	String/URI	The canonical URI path of the workload object, in the form /api/workload-resource-groups/{workload-id}.
object-id	String	The <b>object-uri</b> of the workload object. This string is the <i>{workload-id}</i> portion of the URI path provided by the <b>object-uri</b> field.
name	String	The display name of the workload object.
status	String Enum	The status property of the Workload Resource Group object.



## Description

The **List Workload Resource Groups of an Ensemble** operation lists the workload resource groups that belong to the ensemble targeted by the request URI. The object URI, **status**, object ID, and display name are provided for each.

If the **name** query parameter is specified, the returned list is limited to the workload resource group in the ensemble that has a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

A workload resource group is included in the list only if the API user has object-access permission for that object. An error response is returned if the primary HMC does not manage the target ensemble or if you do not have the requirements listed in “Authorization requirements.”

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Object-access permission to all workload objects to be included in the result.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 448.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>ensemble-id</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/ensembles/b58f0846-8ef7-11df-bb72-00215ef9b504/workload-resource-groups HTTP/1.1
x-api-session: 466ske8jt8waeavxc5rc26gsfjqjfs6aji8qbj7jgne6yrdbq
```

---

Figure 230. List Workload Resource Groups of an Ensemble: Request

---

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Mon, 28 Nov 2011 03:52:05 GMT
content-type: application/json;charset=UTF-8
content-length: 452
{
  "workloads": [
    {
      "name": "Default",
      "object-id": "b65a3066-8ef7-11df-95c2-00215ef9b504",
      "object-uri": "/api/workload-resource-groups/b65a3066-8ef7-11df-95c2-00215ef9b504"
    },
    {
      "name": "SS-Web-Store-Workload",
      "object-id": "546a3398-1974-11e1-999c-00215e6a0c26",
      "object-uri": "/api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26"
    }
  ]
}

```

---

Figure 231. List Workload Resource Groups of an Ensemble: Response

## Get Workload Resource Group Properties

Use the **Get Workload Resource Group Properties** operation to retrieve the properties of a single workload object that is designated by its **object-id**. This operation returns the properties of a single workload resource group in the ensemble.

### HTTP method and URI

**GET** /api/workload-resource-groups/{*workload-id*}

In this request, the URI variable {*workload-id*} is the object ID of the workload object for which you are requesting properties.

### Response body contents

On successful completion, the response body is a JSON object provides the current values of the properties for the workload object as defined in “Data model” on page 443. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

Note that the **virtual-servers**, **virtual-server-groups**, and **workload-element-groups** properties are omitted from this response schema. They are available through their own operations: “List Virtual Servers of a Workload Resource Group” on page 458, “List Groups of Virtual Servers of a Workload Resource Group” on page 465, and “List Workload Element Groups of a Workload Resource Group” on page 471.

### Description

The **Get Workload Resource Group Properties** operation returns the current properties for the workload object specified by {*workload-id*}, as defined in “Response body contents.”

On successful execution, all of the current properties as defined by the Data Model for the Workload Resource Group object are provided in the response body, excluding **virtual-servers**, **virtual-server-groups**, and **workload-element-groups** and HTTP status code 200 (OK) is returned.

An error response is returned if the targeted workload resource group does not exist or if you do not have the requirements listed in “Authorization requirements.”

## Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the workload object passed in the request URI.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 450.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26 HTTP/1.1
x-api-session: 67prscbokwxz601bn1q3feysece2q4275agf27uupjnvr98lse
```

---

Figure 232. Get Workload Resource Group Properties: Request

---

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Mon, 28 Nov 2011 05:09:12 GMT
content-type: application/json;charset=UTF-8
content-length: 1244
{
  "active-perf-policy": {
    "activation-status": "active",
    "element-id": "160c563e-197f-11e1-8914-00215e6a0c26",
    "element-uri": "/api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/
performance-policies/160c563e-197f-11e1-8914-00215e6a0c26",
    "name": "Prime Shift"
  },
  "category": "Production",
  "class": "workload-resource-group",
  "custom-perf-policies": [
    {
      "activation-status": "active",
      "element-id": "160c563e-197f-11e1-8914-00215e6a0c26",
      "element-uri": "/api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/
performance-policies/160c563e-197f-11e1-8914-00215e6a0c26",
      "name": "Prime Shift"
    }
  ],
  "default-perf-policy": {
    "activation-status": "not-active",
    "element-id": "13ec9170-197f-11e1-8914-00215e6a0c26",
    "element-uri": "/api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/
performance-policies/13ec9170-197f-11e1-8914-00215e6a0c26",
    "name": "Default"
  },
  "description": "Spacely Sprockets Web Store Workload",
  "is-default": false,
  "is-locked": false,
  "name": "SS-Web-Store-Workload",
  "object-id": "13de1bfe-197f-11e1-8914-00215e6a0c26",
  "object-uri": "/api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26",
  "parent": "/api/ensembles/b58f0846-8ef7-11df-bb72-00215ef9b504",
  "perf-activation-node-count": 2,
  "perf-activation-status": "active"
}

```

---

Figure 233. Get Workload Resource Group Properties: Response

## Create Workload Resource Group

Use the **Create Workload Resource Group** operation to create a new custom workload resource group in an ensemble. The new workload resource group must be uniquely named within the ensemble.

### HTTP method and URI

**POST** `/api/ensembles/{ensemble-id}/workload-resource-groups`

In this request, the URI variable `{ensemble-id}` is the object ID of the ensemble object for which you want to create a new workload resource group.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String (1-64)	Required	The name to give the new workload resource group, as described in “Data model” on page 443. The passed name must be unique among all other workload resource groups currently in the ensemble.
description	String (0-256)	Optional	The description to give the new workload resource group, as described in “Data model” on page 443.
category	String (0-32)	Optional	The category to give the new workload resource group, as described in “Data model” on page 443.

## Response body contents

On successful completion, the response body is a JSON object with the following field:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the workload object, in the form <code>/api/workload-resource-groups/{workload-id}</code> .

## Description

The **Create Workload Resource Group** operation creates a new custom workload resource group in the ensemble identified by *{ensemble-id}*. On successful execution, the workload resource group is created and added to the ensemble and status code 201 is returned with a response body containing a reference to the new workload resource group.

The **Create Workload Resource Group** operation returns an empty workload resource group that has a default performance policy but does not contain any virtual servers or user-defined groups. To create a fully functional workload resource group, you also need to use these additional operations, in sequence, to minimize the number of policy activations on the virtual servers:

1. “Create Performance Policy” on page 486– Use this operation to create all custom performance policies desired for the new workload resource group. You can omit this step if you do not want to use any custom performance policies.
2. “Activate Performance Policy” on page 493– Use this operation to set the active policy for the new workload resource group. You can omit this step if you want the default performance policy to be the active policy.
3. “Add Virtual Server to a Workload Resource Group” on page 461 or “Add Group of Virtual Servers to a Workload Resource Group” on page 467– Use one or both of these operations to add all of the virtual servers and groups that you want to be members of the new workload resource group. As the virtual servers or groups are added, the active performance policy of the workload resource group is applied to them.

An error response is returned:

- If the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements” on page 454.
- If the request body is not valid. The request body is validated against the schema described in “Request body contents.”

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **New Workload Resource Group** task.

In addition, the target ensemble must be at the Automate entitlement level.

## HTTP status and reason codes

On successful completion, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents” on page 453.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	65	The name of the new workload resource group is not unique.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
	3	The ensemble is not operating at the management entitlement level required to perform this operation (Automate).
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/b58f0846-8ef7-11df-bb72-00215ef9b504/workload-resource-groups HTTP/1.1
x-api-session: 466ske8jt8waeavxc5rc26gsfjqjfs6aji8qbj7jgne6yrdbq
content-type: application/json
content-length: 73
{
  "description": "New Workload Resource Group",
  "name": "SS-New-Workload"
}
```

---

Figure 234. Create Workload Resource Group: Request

---

```

201 Created
server: zSeries management console API web server / 1.0
location: /api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26
cache-control: no-cache
date: Mon, 28 Nov 2011 03:52:01 GMT
content-type: application/json;charset=UTF-8
content-length: 83
{
  "object-uri": "/api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26"
}

```

---

Figure 235. Create Workload Resource Group: Response

## Delete Workload Resource Group

Use the **Delete Workload Resource Group** operation to remove a workload resource group from an ensemble, and deactivate its active policies against its workload element groups and virtual servers. Any virtual server that is not contained by another custom workload resource group directly or through a workload element group is moved into the default. A workload resource group may not be deleted if it contains a workload element group that belongs only to the target workload resource group; to delete the workload resource group you must first delete the workload element group..

### HTTP method and URI

**DELETE** /api/workload-resource-groups/{*workload-id*}

In this request, the URI variable {*workload-id*} is the object ID of the workload object that you want to remove.

### Description

The **Delete Workload Resource Group** operation deletes the workload object specified by {*workload-id*} from its ensemble.

On successful completion, all workload element groups and virtual servers in the workload resource group are removed from it, the workload resource group is removed from the ensemble, and status code 204 (No Content) is returned without supplying a response body. See the *z Systems Ensemble Workload Resource Group Management Guide*, GC27-2629, for details about managing virtual servers in workload resource groups.

An error response is returned if the targeted workload object does not exist or if you do not have the requirements listed in "Authorization requirements." Targeting the default workload resource group also returns an error response because the default workload resource group cannot be deleted.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the ensemble object that owns the workload.
- Object-access permission to the workload object passed in the request URI.
- Action/task permission to the **Delete Workload Resource Group** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned. Otherwise, the following HTTP status codes are returned for the indicated errors.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
	64	A workload element group in the workload resource group does not belong to any other workload resource groups. The workload element group must be deleted before you may delete the workload resource group.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to the object.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	61	The operation cannot be performed because one or more virtual servers in the workload resource group designated by the request URI are currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
DELETE /api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26 HTTP/1.1
x-api-session: 466ske8jt8waeavxcs5rc26gsfjqjfs6aji8qbj7jgne6yrdbq
```

---

Figure 236. Delete Workload Resource Group: Request

---

```
204 No Content
date: Mon, 28 Nov 2011 03:52:06 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 237. Delete Workload Resource Group: Response

## Update Workload Resource Group

Use the **Update Workload Resource Group** operation to modify simple writeable properties of a workload object.

### HTTP method and URI

```
POST /api/workload-resource-groups/{workload-id}
```



In this request, the URI variable *{workload-id}* is the object ID of the workload object that you are modifying.

## Request body contents

The request body is a JSON object with one or more of the following fields. You are required to supply only those fields that you want to change.

Field name	Type	Rqd/Opt	Description
name	String (1-64)	Optional	The new name to give the workload resource group, as described in “Data model” on page 443. The passed name must be unique among all other workload resource groups currently in the ensemble.
description	String (0-256)	Optional	The new description to give the workload resource group, as described in “Data model” on page 443.
category	String (0-32)	Optional	The new category to give the workload resource group, as described in “Data model” on page 443.
acceptable-status	Array of String Enum	Optional	An array of values that define what causes the workload's status to be compliant as described in “Data model” on page 443.
compliant-perf-status	Array of String Enum	Optional	An array of values that define what <b>perf-status</b> values cause the workload's status to be compliant as described in “Data model” on page 443.

## Description

The **Update Workload Resource Group** operation updates one or more simple writeable properties of the workload object identified by *{workload-id}*. These properties are those that are not modified through other specific operations, as noted in “Request body contents.”

The request body is validated against the schema described in “Request body contents.” On successful execution, the workload object is updated with the supplied property values and status code 204 (No Content) is returned without supplying a response body. Notifications for these property changes are sent asynchronously to this operation.

An error response is returned if the targeted workload resource group does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default workload resource group also invokes an error response because its name, description, and category cannot be modified.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Action/task permission to the **Workload Resource Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned. Otherwise, the following HTTP status codes are returned for the indicated errors.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
	65	The new name given to the workload resource group is not unique.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object id in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to the object.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
POST /api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26 HTTP/1.1
x-api-session: 466ske8jt8waeavxcs5rc26gsfjqjfs6aji8qbj7jgne6yrdbq
content-type: application/json
content-length: 114
{
  "category": "Production",
  "description": "Spacely Sprockets Web Store Workload",
  "name": "SS-Web-Store-Workload"
}
```

---

Figure 238. Update Workload Resource Group: Request

---

```
204 No Content
date: Mon, 28 Nov 2011 03:52:01 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 239. Update Workload Resource Group: Response

## List Virtual Servers of a Workload Resource Group

The **List Virtual Servers of a Workload Resource Group** operation lists the virtual servers in the passed workload. This is the way to get the information corresponding to the **virtual-servers** and **directly-added-virtual-servers** properties of a workload, as it is omitted from the standard GET operation.

### HTTP method and URI

**GET** /api/workload-resource-groups/{*workload-id*}/virtual-servers

In this request, the URI variable *{workload-id}* is the object ID of the workload object for which you are requesting a list of virtual servers.

**Query parameters:**

Field name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property. If matches are found, the response will be an array with information for all virtual servers that match. If no match is found, the response will be an empty array.

**Response body contents**

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
virtual-servers	Array of objects	Array of nested virtualserver-info objects as described in the next table.

Each nested virtualserver-info object contains the following fields:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the Virtual Server object.
name	String	The current value of the <b>name</b> property of the Virtual Server object.
status	String Enum	The current value of the <b>status</b> property of the Virtual Server object.
type	String Enum	The value of the <b>type</b> property of the Virtual Server object.
inclusion-type	Array of String Enum	The reason(s) this virtual server is included in the workload resource group. The possible values are as follows: <ul style="list-style-type: none"> <li>• <b>"direct"</b>: The virtual server was added directly to the workload resource group.</li> <li>• <b>"custom-group"</b>: The virtual server was added to the workload resource group due to its membership in one or more user-defined groups which belong to the workload resource group.</li> <li>• <b>"workload-element-group"</b>: The virtual server was added directly to the workload resource group due to its membership in one or more element groups which belong to the workload resource group.</li> </ul>

**Description**

The **List Virtual Servers of a Workload Resource Group** operation lists the virtual servers that belong to the workload resource group targeted by the request URI. The object URI, display name, status, type, and inclusion-type information are provided for each. The inclusion-type information details how the virtual server was added to the workload resource group, as described in “Response body contents.”

An error response is returned if the ensemble does not contain the targeted workload or if you do not have the requirements listed in “Authorization requirements.”

**Authorization requirements**

This operation has the following authorization requirement:

- Object-access permission to the workload object passed in the request URI.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 459.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/workload-resource-groups/29354e42-20db-11e1-9e8d-00215e6a0c26/virtual-servers HTTP/1.1
x-api-session: 4sz56y4z445abfdwva0qsee8wceszpk34chfjfpckfg0i5szv
```

---

Figure 240. List Virtual Servers of a Workload Resource Group: Request

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 07 Dec 2011 13:55:59 GMT
content-type: application/json;charset=UTF-8
content-length: 537
{
  "virtual-servers": [
    {
      "inclusion-type": [
        "custom-group"
      ],
      "name": "SS-Web-Svr-3",
      "object-uri": "/api/virtual-servers/1cc18bee-20db-11e1-b49e-00262df332b3",
      "status": "not-operating",
      "type": "x-hyp"
    },
    {
      "inclusion-type": [
        "direct"
      ],
      "name": "SS-Web-Svr-2",
      "object-uri": "/api/virtual-servers/1a4f490a-20db-11e1-b953-00262df332b3",
      "status": "not-operating",
      "type": "x-hyp"
    },
    {
      "inclusion-type": [
        "direct",
        "custom-group"
      ],
      "name": "SS-Web-Svr-1",
      "object-uri": "/api/virtual-servers/180b51de-20db-11e1-b357-00262df332b3",
      "status": "not-operating",
      "type": "x-hyp"
    }
  ]
}

```

Figure 241. List Virtual Servers of a Workload Resource Group: Response

## Add Virtual Server to a Workload Resource Group

Use the **Add Virtual Server to a Workload Resource Group** operation to add a virtual server to the target workload resource group.

### HTTP method and URI

**POST** `/api/workload-resource-groups/{workload-id}/operations/add-virtual-server`

In this request, the URI variable `{workload-id}` is the object ID of the workload object to which you want to add a virtual server.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
virtual-server	String/URI	The canonical URI identifying the Virtual Server object to be added to the target workload resource group.

## Description

The **Add Virtual Server to a Workload Resource Group** operation adds a virtual server directly to the workload resource group targeted by the request URI. See the *z Systems Ensemble Workload Resource Group Management Guide, GC27-2629*, for details about managing virtual servers in workload resource groups.

A virtual server can become a member of a workload resource group through two ways: by adding the virtual server individually or by adding a custom group that contains the virtual server. The way in which the virtual server was added has an effect on the outcome of this operation:

- If the virtual server is not already a member of the workload resource group, the virtual server is added and a property notification is sent asynchronously to this request. If the virtual server was previously in the default workload resource group, the virtual server is removed automatically from the default workload resource group.
- If the virtual server is already directly defined to the workload resource group (that is, the virtual server was added individually at an earlier time), the request is ignored with a successful return code.
- If the virtual server is already a member of the workload resource group through its membership in a group that already belongs to the workload resource group, the **virtual-servers** property does not change. However, it is now directly defined to the workload resource group such that if the virtual server is removed from the group later, it would remain in the workload resource group.

An error response is returned if the targeted workload resource group or virtual server does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default workload resource group also invokes an error response because you cannot directly modify the default workload resource group.

The request body is validated against the schema described in “Request body contents” on page 461. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the workload object passed in the request URI.
- Object access permission to the Virtual Server object passed in the request body.
- Action/task permission to the **Workload Resource Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
403 (Forbidden)	1	The API user does not have the required permission for this operation.

HTTP error status code	Reason code	Description
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to it.
	2	The <b>object-id</b> in the URI of the Virtual Server object in the request body does not designate an existing object, or the API user does not have access permission to it.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	61	The operation cannot be performed because the virtual server to be added is currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26/operations/
  add-virtual-server HTTP/1.1
x-api-session: 466ske8jt8waeavxcs5rc26gsfjqjfs6aji8qbj7jgne6yrdbq
content-type: application/json
content-length: 79
{
  "virtual-server": "/api/virtual-servers/ecd06dfc-1972-11e1-8879-00262df332b3"
}
```

---

Figure 242. Add Virtual Server to a Workload Resource Group: Request

---

```
204 No Content
date: Mon, 28 Nov 2011 03:52:04 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 243. Add Virtual Server to a Workload Resource Group: Response

## Remove Virtual Server from a Workload Resource Group

Use the **Remove Virtual Server from a Workload Resource Group** operation to remove a virtual server from the target workload resource group.

### HTTP method and URI

```
POST /api/workload-resource-groups/{workload-id}/operations/remove-virtual-server
```

In this request, the URI variable *{workload-id}* is the object ID of the workload object from which you want to remove the virtual server.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
virtual-server	String/URI	The canonical URI identifying the Virtual Server object to be removed from the targeted workload resource group.

## Description

The **Remove Virtual Server from a Workload Resource Group** operation removes a virtual server from being directly defined to the workload resource group targeted by the request URI. See the *Systems Ensemble Workload Resource Group Management Guide, GC27-2629*, for details about managing virtual servers in workload resource groups.

A virtual server can become a member of a workload resource group through two ways: by adding the virtual server individually or by adding a custom group that contains the virtual server. In some cases, a virtual server can be defined twice to the same workload resource group. The way in which the virtual server was added to the workload resource group has an effect on the outcome of this operation:

- If the virtual server was added individually to the workload resource group and is not also a member of a custom group that is already defined to this workload resource group, the virtual server is removed. A property change notification is sent asynchronously to this request.
- If the virtual server was added individually and as a member of a custom group in the workload resource group, the virtual server remains in the workload resource group only as a member of the custom group. In this case, a property change notification is not sent.
- If the virtual server is not already a member of the workload resource group (either individually or through a custom group), the request is ignored with a successful return code.

An error response is returned if the targeted workload resource group or virtual server does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default workload resource group also invokes an error response because you cannot directly modify the default workload resource group.

The request body is validated against the schema described in “Request body contents.” If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Object access permission to the Virtual Server object passed in the request body.
- Action/task permission to the **Workload Resource Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned.

Otherwise, the following HTTP status codes are returned for the indicated errors.



HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to it.
	2	The <b>object-id</b> in the URI of the Virtual Server object in the request body does not designate an existing object, or the API user does not have access permission to it.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	61	The operation cannot be performed because the virtual server to be removed is currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
POST /api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26/operations/
  remove-virtual-server HTTP/1.1
x-api-session: 466ske8jt8waeavxcs5rc26gsfjqjfs6aji8qbj7jgne6yrdbq
content-type: application/json
content-length: 79
{
  "virtual-server": "/api/virtual-servers/f1f581f4-196e-11e1-a344-00262df332b3"
}
```

---

Figure 244. Remove Virtual Server from a Workload Resource Group: Request

---

```
204 No Content
date: Mon, 28 Nov 2011 03:52:06 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 245. Remove Virtual Server from a Workload Resource Group: Response

## List Groups of Virtual Servers of a Workload Resource Group

The **List Groups of Virtual Servers of a Workload Resource Group** operation lists the user-defined managed object groups in the passed workload. This is a way to just get the information maintained in the **virtual-server-groups** property of a workload.

### HTTP method and URI

```
GET /api/workload-resource-groups/{workload-id}/virtual-server-groups
```

In this request, the URI variable *{workload-id}* is the object ID of the workload object for which you are requesting a list of groups.

**Query parameters:**

Field name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern to limit returned objects to those that have a matching <b>name</b> property. If matches are found, the response will be an array with information for all groups that match. If no match is found, the response will be an empty array.

**Response body contents**

On successful completion, the response body is a JSON object containing the following fields:

Field name	Type	Description
virtual-server-groups	Array of String/URI	An array of user-defined managed object groups in the workload resource group, each identified by its URI.

**Description**

The **List Groups of Virtual Servers of a Workload Resource Group** operation lists the user-defined managed object groups that belong to the workload resource group targeted by the request URI. The Group object URI is provided for each.

A group is included in the list only if you meet the requirements listed in “Authorization requirements”; otherwise, an error is returned. An error also is returned if the ensemble does not contain the targeted workload resource group.

**Authorization requirements**

This operation has the following authorization requirement:

- Object-access permission to the workload object passed in the request URI.

**HTTP status and reason codes**

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object-access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26/virtual-server-groups HTTP/1.1
x-api-session: 466ske8jt8waeavxcs5rc26gsfjqjfs6aji8qbj7jgne6yrdbq
```

---

Figure 246. List Groups of Virtual Servers of a Workload Resource Group: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Mon, 28 Nov 2011 03:52:05 GMT
content-type: application/json;charset=UTF-8
content-length: 78
{
  "virtual-server-groups": [
    "/api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341"
  ]
}
```

---

Figure 247. List Groups of Virtual Servers of a Workload Resource Group: Response

## Add Group of Virtual Servers to a Workload Resource Group

Use the **Add Group of Virtual Servers to a Workload Resource Group** operation to add a user-defined managed object group to the target workload resource group. Any virtual servers that the group contains are added as virtual servers of the workload resource group.

### HTTP method and URI

**POST** /api/workload-resource-groups/{workload-id}/operations/add-virtual-server-group

In this request, the URI variable *{workload-id}* is the object ID of the workload object for which you are requesting a list of groups.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
virtual-server-group	String/URI	The canonical URI identifying the user-defined Group object to be added to the targeted workload resource group.

### Description

The **Add Group of Virtual Servers to a Workload Resource Group** operation adds a user-defined group to the workload resource group targeted by the request URI. Through this operation, the virtual servers that belong to the group are added automatically to the workload. See the *z Systems Ensemble Workload Resource Group Management Guide*, GC27-2629, for details about managing virtual servers in workloads.

If the group is already defined to the workload, the request is ignored with a successful return code. If the group is not already defined to the workload resource group, the group is added and a property change notification for the **virtual-server-groups** property is sent asynchronously. All virtual servers in the group are added to the workload resource group, and another property change notification is sent asynchronously for the **virtual-servers** property including the virtual servers that are entirely new to the

workload. If any of these virtual servers previously belonged to the default workload resource group, they are removed automatically from the default workload resource group.

An error response is returned if the targeted workload resource group or custom group does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default workload resource group also results in an error because the default workload is not directly mutable. Pattern match groups are not supported so specifying them in an **Add Group of Virtual Servers to a Workload** operation results in an error.

The request body is validated against the schema described in “Request body contents” on page 467. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Object-access permission to the Group object passed in the request body.
- Object-access permission to all virtual servers contained by the group.
- Action/task permission to the **Workload Resource Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
	61	The <b>object-id</b> of the group to be added in the request body is that of a pattern match group.
	62	The group to be added contains virtual servers to which the API user does not have access.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to it.
	2	The <b>object-id</b> in the URI of the Group object in the request body does not designate an existing object, or the API user does not have access permission to it.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	62	The operation cannot be performed because the group to be added contains one or more virtual servers that are currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26/operations/
  add-virtual-server-group HTTP/1.1
x-api-session: 466ske8jt8waeavxc5rc26gsfjqjfs6aji8qbj7jgne6yrdbq
content-type: application/json
content-length: 76
{
  "virtual-server-group": "/api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341"
}
```

---

Figure 248. Add Group of Virtual Servers to a Workload Resource Group: Request

---

```
204 No Content
date: Mon, 28 Nov 2011 03:52:05 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

```
<No response body>
```

---

Figure 249. Add Group of Virtual Servers to a Workload Resource Group: Response

## Remove Group of Virtual Servers from a Workload Resource Group

Use the **Remove Group of Virtual Servers from a Workload Resource Group** operation to remove a user-defined managed object group from the target workload resource group. All virtual servers in the group are removed from the workload resource group, unless the virtual server belongs to another group that is associated with the same workload resource group, or is directly added.

### HTTP method and URI

```
POST /api/workload-resource-groups/{workload-id}/operations/remove-virtual-server-group
```

In this request, the URI variable *{workload-id}* is the object ID of the workload object from which you are removing the group.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
virtual-server-group	String/URI	The canonical URI identifying the user-defined Group object to be removed from the targeted workload resource group.

### Description

The **Remove Group of Virtual Servers from a Workload Resource Group** operation removes a user-defined group from the workload resource group targeted by the request URI. Through this operation, any virtual servers in the group are removed as well, unless they also belong to another group that is associated with this workload resource group. See the *z Systems Ensemble Workload Resource Group Management Guide*, GC27-2629, for details about managing virtual servers in workload resource groups.

If the group is not defined to the workload resource group, then the request is ignored with a successful return code. If the group is in the workload resource group, then it is removed and a property change notification for the **virtual-server-groups** property is sent asynchronously. Virtual servers in the group that are not also members of another group in the workload resource group, or directly added, are

removed from the workload resource group, and another property change notification is sent asynchronously for the **virtual-servers** property. If any of those virtual servers do not belong to another custom workload resource group, they are placed in the default workload resource group.

An error response is returned if the targeted workload resource group or group does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default workload resource group also invokes an error response because you cannot directly modify the default workload resource group.

The request body is validated against the schema described in “Request body contents” on page 469. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the workload object passed in the request URI.
- Object access permission to the Group object passed in the request body.
- Object access permission to all virtual servers contained by the group.
- Action/task permission to the **Workload Resource Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
	62	The group to be removed contains virtual servers to which the API user does not have access.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to it.
	2	The <b>object-id</b> in the URI of the Group object in the request body does not designate an existing object, or the API user does not have access permission to it.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	62	The operation cannot be performed because the group to be removed contains one or more virtual servers that are currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/workload-resource-groups/546a3398-1974-11e1-999c-00215e6a0c26/operations/
  remove-virtual-server-group HTTP/1.1
x-api-session: 466ske8jt8waeavxcs5rc26gsfjqjfs6aji8qbj7jgne6yrdbq
content-type: application/json
content-length: 76
{
  "virtual-server-group": "/api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341"
}
```

---

Figure 250. Remove Group of Virtual Servers from a Workload Resource Group: Request

---

```
204 No Content
date: Mon, 28 Nov 2011 03:52:06 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

<No response body>

---

Figure 251. Remove Group of Virtual Servers from a Workload Resource Group: Response

## List Workload Element Groups of a Workload Resource Group

The **List Workload Element Groups of a Workload Resource Group** operation lists the workload element groups in the specified workload.

### HTTP method and URI

**GET** /api/workload-resource-groups/{workload-id}/workload-element-groups

In this request, the URI variable {workload-id} is the object ID of the workload object for which you are requesting a list of workload element groups.

### Query parameters:

Field name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern to limit returned objects to those that have a matching <b>name</b> property. If matches are found, the response will be an array with information for all workload element groups that match. If no match is found, the response will be an empty array.

### Response body contents

On successful completion, the response body is a JSON object containing the following fields:

Field name	Type	Description
workload-element-groups	Array of String/URI	An array of workload element groups in the workload resource group, each identified by its URI.

### Description

The **List Workload Element Groups of a Workload Resource Group** operation lists the workload element groups that belong to the workload resource group targeted by the request URI. The Workload Element Group object URI is provided for each.

A workload element group is included in the list only if you meet the requirements listed in “Authorization requirements”; otherwise, an error is returned. An error also is returned if the ensemble does not contain the targeted workload resource group.

## Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the workload object passed in the request URI and to all Workload Element Group objects to be included in the result.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 471.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object-access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Add Workload Element Group to a Workload Resource Group

Use the **Add Workload Element Group to a Workload Resource Group** operation to add a workload element group to the target workload resource group. Any virtual servers that it contains are added as virtual servers of the workload resource group.

### HTTP method and URI

**POST** /api/workload-resource-groups/{workload-id}/operations/add-workload-element-group

In this request, the URI variable *{workload-id}* is the object ID of the workload object for which you are requesting a workload element group.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
workload-element-group	String/URI	The canonical URI identifying the Workload Element Group object to be added to the targeted workload resource group.

### Description

The **Add Workload Element Group to a Workload Resource Group** operation adds a workload element group to the workload resource group targeted by the request URI. A workload element group in a



- | workload has its child virtual servers automatically added to the workload. See the *z Systems Ensemble Workload Resource Group Management Guide*, GC27-2629, for details about managing workload element groups and virtual servers in workloads.

If the workload element group is already defined to the workload, the request is ignored with a successful return code. If the workload element group is not already defined to the workload resource group, then it is added and a property change notification for the **workload-element-groups** property is sent asynchronously. All virtual servers in the workload element group are added to the workload resource group, and another property change notification is sent asynchronously for the **virtual-servers** property including the virtual servers that are entirely new to the workload. If any of these virtual servers previously belonged to the default workload resource group, they are removed automatically from it.

An error response is returned if the targeted workload resource group or workload element group does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default workload resource group also results in an error because the default workload is not directly mutable. A workload element group may only be added if you have authority to access all virtual servers contained by it at the time of the request.

The request body is validated against the schema described in the “Request body contents” on page 472. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Object-access permission to the Workload Element Group object passed in the request body.
- Object-access permission to all virtual servers contained by the workload element group.
- Action/task permission to the **Workload Resource Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
	62	The workload element group to be added contains virtual servers to which the API user does not have access.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to it.
	2	The <b>object-id</b> in the URI of the Workload Element Group object in the request body does not designate an existing object, or the API user does not have access permission to it.

HTTP error status code	Reason code	Description
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	62	The operation cannot be performed because the workload element group to be added contains one or more virtual servers that are currently busy performing some other operation.
	65	The operation cannot be performed because the workload element group to be added is currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Remove Workload Element Group from a Workload Resource Group

Use the **Remove Workload Element Group from a Workload Resource Group** operation to remove a workload element group from the target workload resource group. All virtual servers in the group are removed from the workload resource group, unless the virtual server belongs to another group that is associated with the same workload resource group.

### HTTP method and URI

**POST** `/api/workload-resource-groups/{workload-id}/operations/remove-workload-element-group`

In this request, the URI variable `{workload-id}` is the object ID of the workload object from which you are removing the workload element group.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
workload-element-group	String/URI	The canonical URI identifying the Workload Element Group object to be removed from the targeted workload resource group.

### Description

The **Remove Workload Element Group from a Workload Resource Group** operation removes a workload element group from the workload resource group targeted by the request URI. Through this operation, any virtual servers in the workload element group are removed as well, unless they also belong to another group that is associated with this workload resource group. See the *z Systems Ensemble Workload Resource Group Management Guide, GC27-2629*, for details about managing virtual servers in workload resource groups.

If the workload element group is not defined to the workload resource group, then the request is ignored with a successful return code. If the workload element group is in the workload resource group, then it is removed and a property change notification for the **workload-element-groups** property is sent asynchronously. Virtual servers in the workload element group that are not also members of another group in the workload resource group, or directly added, are removed from the workload resource group, and another property change notification is sent asynchronously for the **virtual-servers** property. If any of those virtual servers do not belong to another custom workload resource group, they are placed in the default workload resource group.

An error response is returned if the targeted workload resource group or workload element group does not exist or if you do not have the requirements listed in “Authorization requirements” on page 475.

default workload resource group also invokes an error response because it is not directly mutable. A workload element group may only be removed if you have authority to access all virtual servers contained by it at the time of the request. A workload element group may not be removed from the only workload to which it belongs; to remove the workload element group from such a workload, the workload element group must be deleted.

The request body is validated against the schema described in “Request body contents” on page 474. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Workload Element Group object passed in the request URI.
- Object access permission to all virtual servers contained by the group.
- Action/task permission to the **Workload Resource Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
	62	The group to be removed contains virtual servers to which the API user does not have access.
	64	The workload element group to be removed does not belong to any other workloads. The workload element group must be deleted to remove it from the workload.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to it.
	2	The <b>object-id</b> in the URI of the Virtual Server object in the request body does not designate an existing object, or the API user does not have access permission to it.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	62	The operation cannot be performed because the group to be removed contains one or more virtual servers that are currently busy performing some other operation.
	65	The operation cannot be performed because the workload element group to be removed is currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Performance Policy object

Performance Policy objects are elements that define performance goals for virtual servers in a workload resource group. A performance policy consists mainly of an importance level and an ordered list of service classes that define the goals. See “Service class nested object” on page 478 for more details on service classes.

Any number of custom Performance Policy objects can be defined for a workload resource group. However, exactly one is active at a time. Every workload resource group contains an immutable default performance policy that is active if a custom workload policy is not activated. Note that changing any property of an active performance policy automatically causes its reactivation. Because activating a performance policy can fail, it follows then that updating properties of an active performance policy can fail. See “Activate Performance Policy” on page 493 for further details.

## Data model

This element includes the following type-specific properties. For definitions of the qualifier abbreviations in the following table, see “Property characteristics” on page 38.

Table 102. Performance Policy object: type-specific properties

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/URI	The canonical URI path for a performance policy is qualified by its workload object and is of the form <code>/api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}</code> where <code>{workload-id}</code> is the value of the <b>object-id</b> property of the parent workload object and <code>{policy-id}</code> is the value of the <b>element-id</b> property of the policy
<b>element-id</b>	—	String (36)	The unique identifier for the performance policy instance. The unique identifier is in the form of a UUID.
<b>parent</b>	—	String/URI	The parent of a performance policy is its owning workload resource group, and so the <b>parent</b> value is the canonical URI path for the workload resource group.
<b>class</b>	—	String	The class of a Performance Policy object is “performance-policy”.
<b>name</b>	(w) (pc)	String (1-64)	The display name specified for the performance policy, which can be up to 64 characters made up of alphanumeric characters, blanks, periods, underscores, or dashes. Names must start with an alphabetic character and end with an alphabetic character, numeric character, underscore, period, or dash. Names must be unique to other existing performance policies in the workload resource group.
<b>description</b>	(w) (pc)	String (0-256)	Arbitrary text describing the performance policy in up to 256 characters.
<b>is-default</b>	—	Boolean	This value is used to identify the default Performance Policy object. It is true for the default policy and false for all other (custom) policies in the workload resource group.

Table 102. Performance Policy object: type-specific properties (continued)

Name	Qualifier	Type	Description
<b>importance</b>	(w) (pc)	String Enum	<p>The importance value assigned to the performance policy, which is one of the following:</p> <ul style="list-style-type: none"> <li>• <b>"highest"</b></li> <li>• <b>"high"</b></li> <li>• <b>"medium"</b></li> <li>• <b>"low"</b></li> <li>• <b>"lowest"</b></li> </ul> <p>See the <i>z Systems Ensemble Workload Resource Group Management Guide</i>, GC27-2629, for detailed information on the <b>importance</b> property values.</p>
<b>custom-service-classes</b>	(w) (pc)	Array of objects (0-99)	<p>The ordered list of custom service classes in the performance policy. Each service class will be an object in the form of a service class object, as described in "Service class nested object" on page 478. This list includes only custom service classes and not the default service class for the policy. This array can contain from 0 to 99 entries.</p>
<b>default-service-class</b>	—	Object	<p>The object describing the default service class for the performance policy, in the form of the service class object, as described in "Service class nested object" on page 478.</p>
<b>revision</b>	(pc)	Integer (1-n)	<p>The revision count of the performance policy, which is the number of times the policy has been modified. The initial value for a new performance policy is 1.</p>
<b>activation-status</b>	(pc)	String Enum	<p>The activation status of the performance policy, which is one of the following values:</p> <ul style="list-style-type: none"> <li>• <b>"not-active"</b> – the performance policy is not currently the active policy for the workload resource group</li> <li>• <b>"in-progress"</b> – the performance policy is currently being activated for the workload resource group</li> <li>• <b>"active"</b> – the performance policy is currently the active policy for the workload resource group, and its activation has completed</li> </ul>
<b>last-activation-requested-date</b>	(pc)	Timestamp	<p>The standard date/time value indicating the time of the last activation request of the policy. If the policy has never been active or is the default policy of the default workload resource group, this value is negative.</p> <p>The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.</p>
<b>last-activation-completed-date</b>	(pc)	Timestamp	<p>The standard date/time value indicating the time the policy last completed activation. If the policy has never been active or is the default policy of the default workload resource group, this value is negative.</p> <p>The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.</p>

Table 102. Performance Policy object: type-specific properties (continued)

Name	Qualifier	Type	Description
<b>last-activated-by</b>	(pc)	String	The user name used for the last activation request. This value is the empty string if the policy has never been activated or is the default policy of the default workload resource group.
<b>last-modified-date</b>	(pc)	Timestamp	The standard date/time value indicating the time the policy was last modified. If the policy has never been modified, this property is equal to the <b>created-date</b> property. The value is negative for all default policies.  The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
<b>last-modified-by</b>	(pc)	String	The user name used to last modify the policy. If the policy has never been modified, this property is equal to the <b>created-by</b> property. The value is the empty string for all default policies.
<b>created-date</b>	—	Timestamp	The standard date/time value indicating the time the policy was created. The value is negative for all default policies.  The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
<b>created-by</b>	—	String	The user name used to create the policy. The value is the empty string for all default policies.

### Service class nested object

A service class object is a nested object of a Performance Policy object. A performance policy can contain zero or more ordered service classes, and contains a default service class. Service classes use classification rules to classify and apply performance goals to virtual servers in the workload resource group. Service classes are positional; during classification, zManager searches service classes from low-ordered to high-ordered service class to find a match for a virtual server. The default service class applies for any virtual server that does not match a custom service class.

You cannot directly change a service class object. When any property of a service class must be changed, zManager creates a new service class object to replace the existing one in the policy object. For this reason, service classes are not supported as objects in the standard sense, but rather as simple nested objects of a performance policy. Because of this fact, direct operations and notifications are not supported for service class objects.

The following properties are supported:

Table 103. Performance Policy object: Service class nested object properties

Field name	Type	Rqd/Opt	Description
<b>name</b>	String (1-64)	Required	The display name specified for the service class, which can be up to 64 characters made up of alphanumeric characters, blanks, periods, underscores, or dashes. Names must start with an alphabetic character and end with an alphabetic character, numeric character, underscore, period, or dash. Names must be unique to other existing service classes in the performance policy.
<b>description</b>	String (0-256)	Optional	Arbitrary text describing the performance policy in up to 256 characters.

Table 103. Performance Policy object: Service class nested object properties (continued)

Field name	Type	Rqd/Opt	Description
<b>type</b>	String Enum	Required	This identifies the type of the service class object: <ul style="list-style-type: none"> <li>• <b>"server"</b> – the service class type targeting specific virtual servers</li> </ul>
<b>goal-type</b>	String Enum	Required	This identifies the type of performance goal for the service class, which must be one of the following: <ul style="list-style-type: none"> <li>• <b>"velocity"</b> – the service class goal is based on a velocity value as given in the velocity property</li> <li>• <b>"discretionary"</b> – the service class goal is up to the discretion of zManager</li> </ul>
<b>business-importance</b>	String Enum	Optional*	This field identifies the business importance level assigned to the service class, which must be one of the following: <ul style="list-style-type: none"> <li>• <b>"highest"</b></li> <li>• <b>"high"</b></li> <li>• <b>"medium"</b></li> <li>• <b>"low"</b></li> <li>• <b>"lowest"</b></li> </ul> <p>* - A business importance is required for the velocity goal type, but is not used for a discretionary goal type.</p>
<b>velocity</b>	String Enum	Optional*	This field identifies the velocity goal value of the service class, which must be one of the following: <ul style="list-style-type: none"> <li>• <b>"fastest"</b></li> <li>• <b>"fast"</b></li> <li>• <b>"moderate"</b></li> <li>• <b>"slow"</b></li> <li>• <b>"slowest"</b></li> </ul> <p>* - A velocity value is only required for the velocity goal type.</p>
<b>classification-rule</b>	Object	Required	The rule used to filter the virtual servers for which the service class goal applies. The value must be a JSON object, as described in "Classification rule nested object."

### Classification rule nested object

A classification rule object is a nested object within a service class object. It is a recursive object used to define logical filter statements to be applied to the virtual servers of the workload resource group. If a virtual server passes the classification rule, the service class goal is applied.

Table 104. Performance Policy object: classification rule nested object properties

Field name	Type	Rqd/Opt	Description
<b>type</b>	String Enum	Required	This field identifies the type of classification rule object, which must be one of the following: <ul style="list-style-type: none"> <li>• <b>"and"</b> – for this rule to be true, both of the two rules referenced by this rule must be true</li> <li>• <b>"or"</b> – for this rule to be true, only one of the two rules referenced by this rule must be true</li> <li>• <b>"rule"</b> – the rule defines a filter that resolves to true or false based on its filter pattern against a virtual server</li> </ul> <p>If you specify <b>"and"</b> or <b>"or"</b>, exactly two classification rule objects must be nested inside this classification rule object so they can be logically compared. If you specify <b>"rule"</b>, exactly one filter object must be nested inside this object.</p>

Table 104. Performance Policy object: classification rule nested object properties (continued)

Field name	Type	Rqd/Opt	Description
<b>classification-rule-1</b>	Object	Required*	The first of two classification rule objects that must be nested inside this classification rule object if they are to be logically compared as defined by the <b>type</b> property. If the <b>type</b> is "rule", this property is not supported.  * - The <b>type</b> property value determines whether this field is required.
<b>classification-rule-2</b>	Object	Required*	The second of two classification rule objects that must be nested inside this classification rule object if they are to be logically compared as defined by the <b>type</b> property. If the <b>type</b> is "rule", this property is not supported.  * - The <b>type</b> property value determines whether this field is required.
<b>filter</b>	Object	Required*	A filter object defining the filter statement if the classification rule type is "rule". If the <b>type</b> is "and" or "or", this property is not supported.  * - The <b>type</b> property value determines whether this field is required.

### Filter nested object

A filter object is a nested object within a classification rule object whose type is "rule". The filter object defines a logical statement used to filter on properties of a virtual server.

Table 105. Performance Policy object: filter nested object properties

Field name	Type	Rqd/Opt	Description
<b>type</b>	String Enum	Required	This field identifies the type of the filter object, which must be one of the following: <ul style="list-style-type: none"> <li>"hostname" – the filter value is matched against the hostname of the virtual server</li> <li>"virtual-server-name" – the filter value is matched against the name of the virtual server</li> <li>"os-type" – the filter value is matched against the operating system type of the virtual server</li> <li>"os-level" – the filter value is matched against the operating system level of the virtual server</li> <li>"os-name" – the filter value is matched against the operating system name of the virtual server</li> </ul>
<b>operation</b>	String Enum	Required	This field identifies the logical filter operation, which must be one of the following: <ul style="list-style-type: none"> <li>"string-match" – the filter value must match the property defined by the filter type</li> <li>"string-not-match" – the filter value must not match the property defined by the filter type</li> </ul>
<b>value</b>	String (1-255)	Required	A string to match the property against. Only patterns "." (to match any character) and ".*" (to match any character zero or more times) are currently supported and ".*" may only be used at the end of the filter value string. The following characters must be escaped to be used directly: +-()[]\$^.*\.

## Notifications of property changes to performance policies

Notifications of property changes to Performance Policy objects are similar to notifications of changes to workload objects.

Because Performance Policy objects are sub-objects of workload objects, creating or deleting a performance policy does not result in an inventory notification. Instead, these operations result in the



appropriate property notification for the target workload object. However, updating performance policy properties results in property notifications for the policy itself. In these cases, the notification header contains an **element uri** and an **element id** to indicate the performance policy that changed.

## List Performance Policies

Use the **List Performance Policies** operation to list the performance policies within the target workload resource group.

### HTTP method and URI

**GET** `/api/workload-resource-groups/{workload-id}/performance-policies`

In this request, the URI variable `{workload-id}` is the object ID of the workload object for which you are requesting a list of performance policies.

### Query parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property. If a match is found, the response is an array with all policies that match. If no match is found, the response is an empty array.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
perf-policies	Array of objects	Array of nested perf-policy-info objects as the following table.

Each nested perf-policy-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Performance Policy object, in the form <code>/api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}</code> .
element-id	String	The <b>element-id</b> of the Performance Policy object. This field value is the <code>{policy-id}</code> portion of the URI path provided by the <b>element-uri</b> field.
name	String	Display name of the Performance Policy object.
activation-status	String Enum	The status of the Performance Policy object, which must be one of the following values: <ul style="list-style-type: none"> <li><b>"not-active"</b> – the performance policy is not currently the active policy for the workload resource group</li> <li><b>"in-progress"</b> – the performance policy is currently being activated for the workload resource group</li> <li><b>"active"</b> – the performance policy is currently the active policy for the workload resource group, and its activation has completed</li> </ul>

## Description

The **List Performance Policies** operation lists the performance policies that belong to the workload resource group targeted by the request URI. The element URI, element ID, display name, and status are provided for each.

If the **name** query parameter is specified, the returned list is limited to the policies in the workload resource group that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

An error response is returned if the HMC does not manage the targeted workload resource group or if you do not have the requirements listed in “Authorization requirements.”

## Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the workload object passed in the request URI. A policy is included in the list only if you also have object-access permission for that policy object.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 481.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/
    performance-policies HTTP/1.1
x-api-session: 67prschbokwxz6o1bn1q3feysece2q4275agf27uupjnvr981se
```

---

Figure 252. List Performance Policies: Request

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Mon, 28 Nov 2011 05:09:12 GMT
content-type: application/json;charset=UTF-8
content-length: 509
{
  "perf-policies": [
    {
      "activation-status": "not-active",
      "element-id": "13ec9170-197f-11e1-8914-00215e6a0c26",
      "element-uri": "/api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/performance-policies/13ec9170-197f-11e1-8914-00215e6a0c26",
      "name": "Default"
    },
    {
      "activation-status": "active",
      "element-id": "160c563e-197f-11e1-8914-00215e6a0c26",
      "element-uri": "/api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/performance-policies/160c563e-197f-11e1-8914-00215e6a0c26",
      "name": "Prime Shift"
    }
  ]
}

```

Figure 253. List Performance Policies: Response

## Get Performance Policy Properties

Use the **Get Performance Policy Properties** operation to retrieve the properties of a single performance policy.

### HTTP method and URI

**GET** /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}

#### URI variables

Variable	Description
{workload-id}	Object ID of the workload object for the workload resource group to which the performance policy is defined.
{policy-id}	Element ID of the Performance Policy object for which properties are to be obtained.

### Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Performance Policy object as defined in “Data model” on page 476. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

The **Get Performance Policy Properties** operation returns the current properties for the Performance Policy object specified by {policy-id}.

An error response is returned if the targeted workload resource group or policy does not exist or if you do not have the requirements listed in “Authorization requirements.”

## Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the workload object passed in the request URI.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 483.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The <b>object id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to the object.
	62	The <b>element id</b> in the URI ( <i>{policy-id}</i> ) does not designate an existing Performance Policy object in the workload resource group.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/
    performance-policies/160c563e-197f-11e1-8914-00215e6a0c26 HTTP/1.1
x-api-session: 67prscbokwxz601bn1q3feysece2q4275agf27uupjnvr98lse
```

---

Figure 254. Get Performance Policy Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Mon, 28 Nov 2011 05:09:12 GMT
content-type: application/json;charset=UTF-8
content-length: 1586
{
  "activation-status": "active",
  "class": "performance-policy",
  "created-by": "ENSADMIN",
  "created-date": 1322456941929,
  "custom-service-classes": [
    {
      "business-importance": "highest",
      "classification-rule": {
        "filter": {
          "operation": "string-match",
          "type": "virtual-server-name",
          "value": "SS\\-Premium\\-Web\\-Svr\\-.*"
        },
        "type": "rule"
      },
      "description": "",
      "goal-type": "velocity",
      "name": "Premium Class",
      "type": "server",
      "velocity": "fastest"
    },
    {
      "business-importance": "high",
      "classification-rule": {
        "filter": {
          "operation": "string-match",
          "type": "virtual-server-name",
          "value": "SS\\-Web\\-Svr\\-.*"
        },
        "type": "rule"
      },
      "description": "",
      "goal-type": "velocity",
      "name": "Regular Class",
      "type": "server",
      "velocity": "moderate"
    }
  ],
}
```

---

Figure 255. Get Performance Policy Properties: Response (Part 1)

```

"default-service-class": {
  "business-importance": "medium",
  "classification-rule": {
    "filter": {
      "operation": "string-match",
      "type": "(\\*)",
      "value": "(\\*)"
    },
    "type": "rule"
  },
  "description": "The default workload performance policy service class.",
  "goal-type": "velocity",
  "name": "Default",
  "type": "server",
  "velocity": "moderate"
},
"description": "Performance policy for prime shift",
"element-id": "160c563e-197f-11e1-8914-00215e6a0c26",
"element-uri": "/api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/
performance-policies/160c563e-197f-11e1-8914-00215e6a0c26",
"importance": "highest",
"is-default": false,
"last-activated-by": "ENSADMIN",
"last-activation-completed-date": 1322456944022,
"last-activation-requested-date": 1322456942144,
"last-modified-by": "ENSADMIN",
"last-modified-date": 1322456942090,
"name": "Prime Shift",
"parent": "/api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26",
"revision": 2
}

```

Figure 256. Get Performance Policy Properties: Response (Part 2)

## Create Performance Policy

Use the **Create Performance Policy** operation to create a new custom performance policy for a workload resource group. The new policy must be uniquely named within the workload resource group. The policy is inactive until the **Activate Performance Policy** operation is applied to it.

### HTTP method and URI

**POST** /api/workload-resource-groups/{workload-id}/performance-policies

In this request, the URI variable {workload-id} is the object ID of the workload object for which you are creating a new performance policy.

### Request body contents

The request body contains properties used to define the new performance policy, which are the writeable properties of a performance policy. Some properties are optional.

Field name	Type	Rqd/Opt	Description
name	String (1-64)	Required	The name to give the new performance policy, as described in “Data model” on page 476. The passed name must be unique among all other policies currently in the workload resource group.
description	String (0-256)	Optional	The description to give the new performance policy, as described in “Data model” on page 476.

Field name	Type	Rqd/Opt	Description
importance	String Enum	Required	The importance value to give the new performance policy, as described in “Data model” on page 476.
custom-service-classes	Array of objects (0-99)	Optional	The ordered list of custom service classes in the new performance policy. Each service class is an object in the form of a service class object, as described in “Service class nested object” on page 478. This array can contain from 0 to 99 entries.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Performance Policy object, in the form <code>/api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}</code>

## Description

The **Create Performance Policy** operation creates a new custom performance policy in a workload resource group, identified by *{workload-id}*.

On successful execution, the performance policy is created and added to the workload resource group and status code 201 is returned with a response body containing a reference to the new Performance Policy object. Note that the new policy is not active until the **Activate Performance Policy** operation is applied to it.

An error response is returned if the targeted workload resource group does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default workload object also results in an error because the default workload resource group cannot contain any performance policy other than the default performance policy.

The request body is validated against the data model for this object type to ensure that it contains only writable properties and that the data types of those properties are specified as required. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Action/task permission to the **New Performance Policy** task.

## HTTP status and reason codes

On successful completion, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents.”

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
	65	The name of the performance policy is not unique within its workload resource group, or one or more of the defined service class names is reserved or not unique.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to the object.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
415 (Unsupported Media Type)	0	The request does not include a <b>Content-Type</b> header that specifies the request body is of media type <b>application/xml</b> .

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
POST /api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/performance-policies HTTP/1.1
x-api-session: 67prscbokwxz6o1bn1q3feysece2q4275agf27uupjnv981se
content-type: application/json
content-length: 101
{
  "description": "Performance policy for prime shift",
  "importance": "highest",
  "name": "Prime Shift"
}
```

---

Figure 257. Create Performance Policy: Request

---

```
201 Created
server: zSeries management console API web server / 1.0
location: /api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/
performance-policies/160c563e-197f-11e1-8914-00215e6a0c26
cache-control: no-cache
date: Mon, 28 Nov 2011 05:09:01 GMT
content-type: application/json;charset=UTF-8
content-length: 142
{
  "element-uri": "/api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/
performance-policies/160c563e-197f-11e1-8914-00215e6a0c26"
}
```

---

Figure 258. Create Performance Policy: Response

## Delete Performance Policy

Use the **Delete Performance Policy** operation to remove a performance policy from a workload resource group. This operation cannot be performed on an active or default performance policy.



## HTTP method and URI

DELETE /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}

### URI variables

Variable	Description
{workload-id}	Object ID of the workload object whose Performance Policy object is to be deleted.
{policy-id}	Element ID of the Performance Policy object that is to be deleted.

## Description

The **Delete Performance Policy** operation deletes the Performance Policy object specified by {policy-id} from its workload resource group specified by {workload-id}.

On successful execution, the Performance Policy object is removed from the workload resource group and status code 204 (No Content) is returned without a response body.

An error response is returned if the targeted policy object does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting a default performance policy also results in an error because the default performance policy cannot be deleted. You also cannot delete an active performance policy without first activating another performance policy through the **Activate Performance Policy** operation.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Action/task permission to the **Delete Performance Policy** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned without a response body.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	63	The targeted performance policy is a default Performance Policy object.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object id</b> in the URI ({workload-id}) does not designate an existing workload object, or the API user does not have object access permission to the object.
	62	The <b>element id</b> in the URI ({policy-id}) does not designate an existing Performance Policy object in the workload resource group.

HTTP error status code	Reason code	Description
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	60	The performance policy to be deleted is currently active or in the progress of being activated. The operation can be retried after a different policy has been activated, which causes the <b>activation-status</b> of this policy to be <b>"not-active"</b> .

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
DELETE /api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/
performance-policies/160c563e-197f-11e1-8914-00215e6a0c26 HTTP/1.1
x-api-session: 67prscbokwxz6o1bn1q3feysece2q4275agf27uupjnv98l9e
```

---

Figure 259. Delete Performance Policy: Request

---

```
204 No Content
date: Mon, 28 Nov 2011 05:09:22 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 260. Delete Performance Policy: Response

## Update Performance Policy

Use the **Update Performance Policy** operation to modify one or more writeable properties of a Performance Policy object. Note that updating an active performance policy causes its reactivation.

### HTTP method and URI

**POST** /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}

#### URI variables

Variable	Description
{workload-id}	Object ID of the workload object to which the target performance policy is defined.
{policy-id}	Element ID of the Performance Policy object that is to be modified.

### Request body contents

The request body is a JSON object containing one or more of the writeable fields for a performance policy, as described in “Data model” on page 476. You need to supply only those fields that you want to modify.

## Description

The **Update Performance Policy** operation updates one or more writeable properties of the Performance Policy object identified by *{policy-id}*.

On successful execution, the Performance Policy object is updated with the supplied property values and status code 204 (No Content) is returned without a response body. Notifications for these property changes are sent asynchronously to this operation.

If the performance policy is active at the time of this request, a reactivation of the policy is submitted asynchronously to this operation. Because activation of a performance policy is rejected if another activation request is in progress for the target workload resource group, an update to an active performance policy also can be rejected. In this case, a 409 (Conflict) status code is returned and you can retry the update operation after the first activation completes.

An error response is returned if the targeted workload resource group or performance policy does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default performance policy also results in an error because the default performance policy cannot be directly modified.

The request body is validated against the schema described in “Request body contents” on page 490. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Action/task permission to the **Performance Policy Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned without a response body.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	63	The targeted performance policy is a default Performance Policy object.
	65	The new name given to the performance policy is not unique within its workload resource group, or one or more of the defined service class names are reserved or not unique.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to the object.
	62	The <b>element id</b> in the URI ( <i>{policy-id}</i> ) does not designate an existing Performance Policy object in the workload resource group.

HTTP error status code	Reason code	Description
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	60	The performance policy to be updated is currently in the progress of being activated. The operation can be retried after activation has completed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```

POST /api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/performance-policies/
160c563e-197f-11e1-8914-00215e6a0c26 HTTP/1.1
x-api-session: 67prscbokwxz601bn1q3feysece2q4275agf27uupjnv981se
content-type: application/json
content-length: 580
{
  "custom-service-classes": [
    {
      "business-importance": "highest",
      "classification-rule": {
        "filter": {
          "operation": "string-match",
          "type": "virtual-server-name",
          "value": "SS\\-Premium\\-Web\\-Svr\\-.*"
        },
        "type": "rule"
      },
      "goal-type": "velocity",
      "name": "Premium Class",
      "type": "server",
      "velocity": "fastest"
    },
    {
      "business-importance": "high",
      "classification-rule": {
        "filter": {
          "operation": "string-match",
          "type": "virtual-server-name",
          "value": "SS\\-Web\\-Svr\\-.*"
        },
        "type": "rule"
      },
      "goal-type": "velocity",
      "name": "Regular Class",
      "type": "server",
      "velocity": "moderate"
    }
  ]
}

```

Figure 261. Update Performance Policy: Request

---

```
204 No Content
date: Mon, 28 Nov 2011 05:09:01 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

```
<No response body>
```

---

Figure 262. Update Performance Policy: Response

## Activate Performance Policy

Use the **Activate Performance Policy** operation to activate a performance policy for a workload resource group. You can activate any performance policy defined to the workload resource group, including the default and currently active policy.

### HTTP method and URI

```
POST /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}/operations/activate
```

### URI variables

Variable	Description
{workload-id}	Object ID of the workload object to which the target performance policy is defined.
{policy-id}	Element ID of the Performance Policy object that is to be activated.

### Description

The **Activate Performance Policy** operation activates or reactivates a performance policy for a specific workload resource group.

On successful execution, the target policy is active and status code 204 (No Content) is returned without a response body. Notifications for ensuing property changes are sent asynchronously to this operation.

An activation request is not accepted if another activation request is in progress for this workload resource group. In this case, status code 409 (Conflict) is returned and you can retry this operation after the first activation completes.

An error response is returned if the targeted workload resource group or performance policy does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default workload resource group also results in an error because its default performance policy is permanently active.

### Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Action/task permission to the **Activate Performance Policy** task.

### HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned without a response body.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to the object.
	62	The <b>element id</b> in the URI ( <i>{policy-id}</i> ) does not designate an existing Performance Policy object in the workload resource group.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	60	Performance policy activation is currently in progress on the workload resource group. The operation can be retried after activation has completed.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/workload-resource-groups/13de1bfe-197f-11e1-8914-00215e6a0c26/performance-policies/
160c563e-197f-11e1-8914-00215e6a0c26/operations/activate HTTP/1.1
x-api-session: 67prscbokwxz6olbn1q3feysece2q4275agf27uupjnvr98lse
```

---

Figure 263. Activate Performance Policy: Request

---

```
204 No Content
date: Mon, 28 Nov 2011 05:09:01 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache
```

```
<No response body>
```

---

Figure 264. Activate Performance Policy: Response

## Import Performance Policy

Use the **Import Performance Policy** operation to create a new custom performance policy for a workload resource group. This operation is equivalent to the **Create Performance Policy** operation, except the request body is an XML document defining the configuration of the performance policy to be created. The new policy must be uniquely named within the workload resource group. The policy is inactive until the **Activate Performance Policy** operation is applied to it.

### HTTP method and URI

```
POST /api/workload-resource-groups/{workload-id}/operations/import-performance-policy
```

In this request, the URI variable *{workload-id}* is the object ID of the workload object for which you are importing a performance policy.

## Request body contents

The request body must contain an XML document that describes the performance policy to be created. The XML must conform to the schema described in Appendix A, “XML document structure of a performance policy,” on page 923. Note that the same rules apply as for the Create Performance Policy operation, that the name must be unique within the workload.

Because the request body is expected to be in XML format, the request should specify MIME type **application/xml** as the value of the HTTP **Content-Type** header for the request.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Performance Policy object, in the form <code>/api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}</code>

## Description

The **Import Performance Policy** operation imports a performance policy from an XML document and creates a Performance Policy object, adding it to the workload resource group identified by *{workload-id}*.

On successful execution, the Performance Policy object is created with the supplied property values, added to the collection of custom performance policies for the workload resource group. Status code 201 (Created) is returned with a response body describing the location of the new Performance Policy object. Note that the new policy is not active until the **Activate Performance Policy** operation is applied to it.

An error response is returned if the targeted workload does not exist or if you do not have the requirements listed in “Authorization requirements.” Targeting the default workload object also results in an error because the default workload resource group cannot contain any performance policy other than the default performance policy.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Action/task permission to the **Import Performance Policy** task.

## HTTP status and reason codes

On successful completion, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents.”

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload resource group is the default workload object.
	65	The name of the performance policy is not unique within its workload resource group, or one or more of the defined service class names is reserved or not unique.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to the object.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
415 (Unsupported Media Type)	0	The request does not include a <b>Content-Type</b> header that specifies the request body is of media type <b>application/xml</b> .

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Export Performance Policy

Use the **Export Performance Policy** operation to return the configuration of an existing performance policy in the form of an XML document.

### HTTP method and URI

**POST** /api/workload-resource-groups/{workload-id}/performance-policies/{policy-id}/operations/export

#### URI variables

Variable	Description
<i>{workload-id}</i>	Object ID of the workload object for the workload resource group to which the performance policy is defined.
<i>{policy-id}</i>	Element ID of the Performance Policy object to be exported.

## Response body contents

On successful completion, the performance policy is returned in the response body as an XML document that conforms to the schema described in Appendix A, “XML document structure of a performance policy,” on page 923.

### Description

The **Export Performance Policy** operation returns the configurable properties of an existing performance policy identified by *{policy-id}* in a workload resource group identified by *{workload-id}*. The policy configuration is returned in the form of an XML document. You can save this XML as a backup copy of the policy configuration or modify and import it through the **Import Performance Policy** operation to create a new custom policy.

On successful execution, the Performance Policy object returned as a XML document with status code 200 (OK).



An error response is returned if the targeted workload resource group or policy does not exist or if you do not have the requirements listed in “Authorization requirements.”

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Action/task permission to the **Export Performance Policy** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 496.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to the object.
	62	The <b>element id</b> in the URI ( <i>{policy-id}</i> ) does not designate an existing Performance Policy object in the workload resource group.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/workload-resource-groups/e39812b2-209d-11e1-8bbc-0010184c8334/performance-policies/
e44af260-209d-11e1-8bbc-0010184c8334/operations/export HTTP/1.1
x-api-session: 2ggnk3nbcxsi8w8qbjpn14pqhwgkvdrdu5hw8pjw1o171qu2r
```

---

Figure 265. Export Performance Policy: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 07 Dec 2011 06:37:21 GMT
content-type: application/xml;charset=utf-8
content-length: 1515
<?xml version="1.0" encoding="UTF-8"?>

<PerformancePolicy xmlns="http://www.ibm.com/PPM/WorkloadPerformancePolicy-1.0">
  <Name>Prime Shift</Name>
  <Description>Performance policy for prime shift</Description>
  <Version>3.00.00</Version>
  <UI>PPM Editor</UI>
  <WorkloadImportance>Highest</WorkloadImportance>

  <ServiceClasses>

    <ServiceClass>
      <Name>Premium Class</Name>
      <Type>Server</Type>
      <RuleBuilderElement>
        <RuleBuilderElementType>Rule</RuleBuilderElementType>
        <Filter>
          <FilterType>Virtual Server Name</FilterType>
          <FilterOperation>stringMatch</FilterOperation>
          <FilterValue>SS\\-Premium\\-Web\\-Svr\\-\\.\\*</FilterValue>
        </Filter>
      </RuleBuilderElement>
      <Goal>
        <Velocity>
          <Importance>Highest</Importance>
          <Level>Fastest</Level>
        </Velocity>
      </Goal>
    </ServiceClass>

    <ServiceClass>
      <Name>Regular Class</Name>
      <Type>Server</Type>
      <RuleBuilderElement>
        <RuleBuilderElementType>Rule</RuleBuilderElementType>
        <Filter>
          <FilterType>Virtual Server Name</FilterType>
          <FilterOperation>stringMatch</FilterOperation>
          <FilterValue>SS\\-Web\\-Svr\\-\\.\\*</FilterValue>
        </Filter>
      </RuleBuilderElement>
      <Goal>
        <Velocity>
          <Importance>High</Importance>
          <Level>Moderate</Level>
        </Velocity>
      </Goal>
    </ServiceClass>

  </ServiceClasses>

</PerformancePolicy>
```

---

Figure 266. Export Performance Policy: Response

---

## Performance management reports

Through specific APIs described in this topic, you can request zManager to generate historical reports that contain performance data related to specific performance management objects. You can request data for a specific time interval, up to 36 hours before the current time. Because the actual performance management objects might have changed since the time interval you request, or actually might not exist now, the performance management report APIs are not object-based.

Instead, all of the report APIs are implemented as performance-management specific operations that request an on-demand report to be generated and returned to the caller. They accept query parameters in a request block, and query the historical reporting database to generate a custom report for the caller.

The performance management report APIs are interrelated because you can use certain properties retrieved from one report as input to generate other reports, just as you can drill down for more information through the report tasks in the HMC UI.

Figure 267 on page 500 shows how the performance management reports are related, along with an indication of which properties you can use from one report to generate the next one through the APIs. Once you receive these properties in the response to an API request for one report, you can repeatedly reuse them as input properties for subsequent reports for the same performance management object, regardless of the time interval for the report.

For example, suppose that you use the API to request a Virtual Servers Report for the ensemble and the response body includes a **hypervisor-report-id** property. At any time in the future, you can use the **hypervisor-report-id** property directly to request a Hypervisor Report, rather than having to drill down to it again through the Virtual Servers Report. The **\*-report-id** property values do not change across invocations of the reporting APIs; instead, the values remain constant over time.

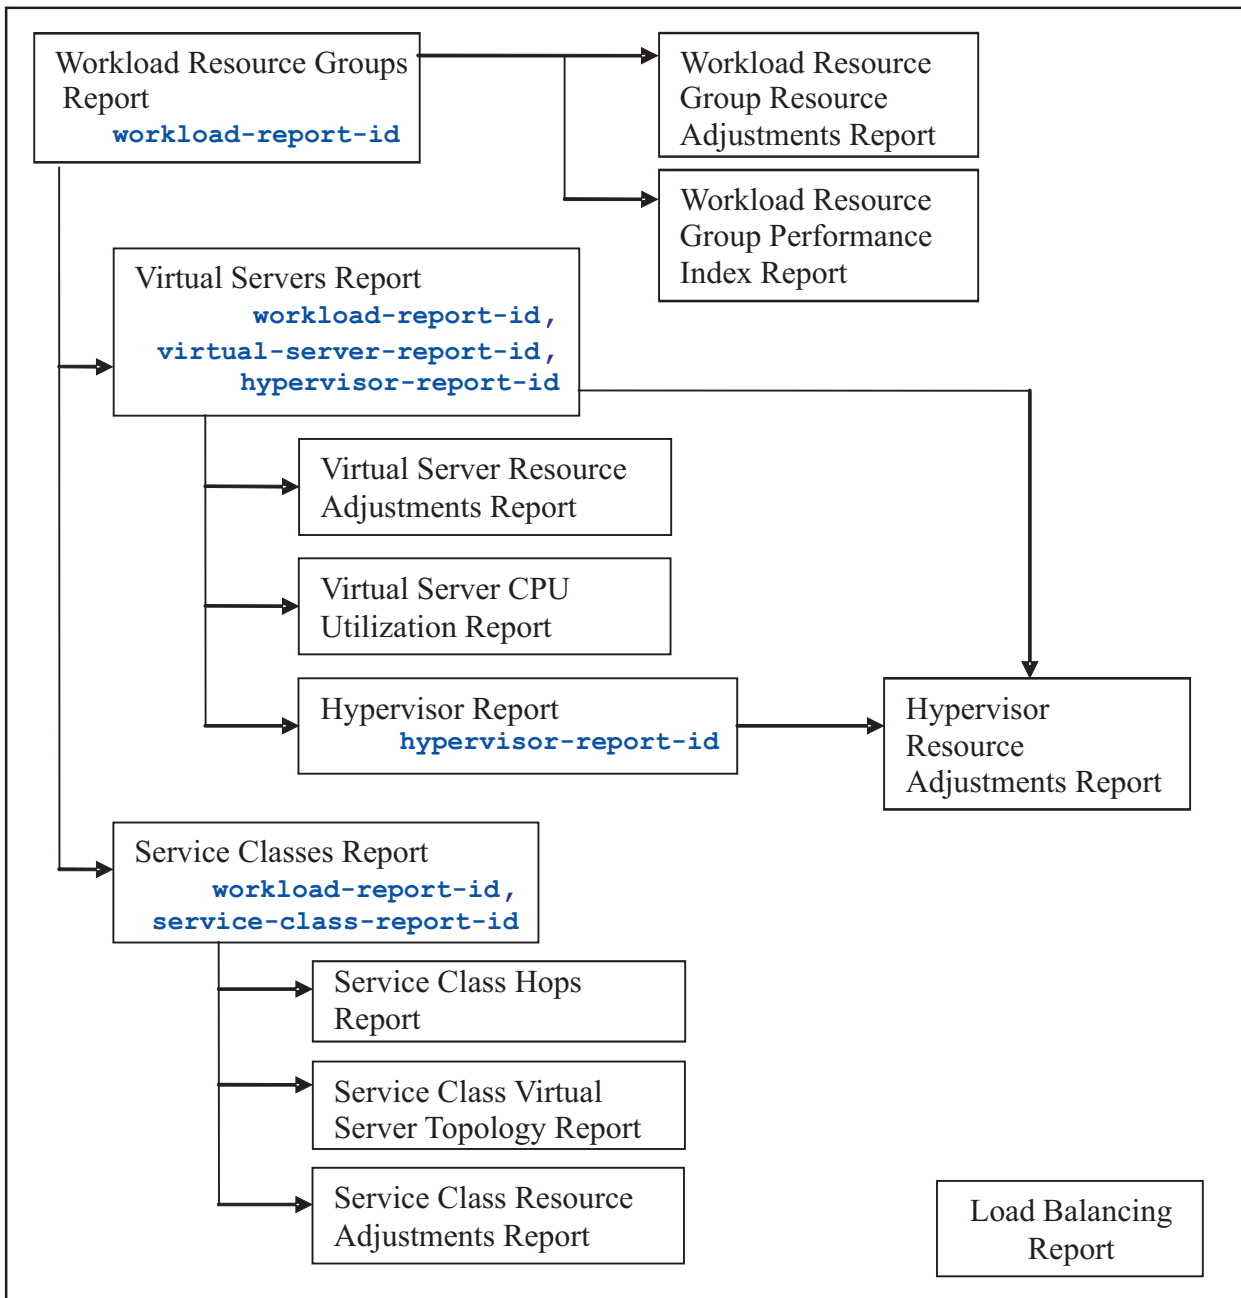


Figure 267. Relationship between reports and the properties used

## Generate Workload Resource Groups Report (Performance Management)

Use the **Generate Workload Resource Groups Report** operation to create a custom on-demand report that shows all workload resource groups in a specific ensemble over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

```
POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-workload-resource-groups-report
```

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble object for which you are requesting a workload resource groups report.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
report-workloads	Array of objects	Array of nested workload-report-entry objects described in Table 106.

Each nested workload-report-entry object contains the following fields:

*Table 106. Format of a workload-report-entry object*

Field name	Type	Description
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the Workload Resource Groups Report for this same interval.
performance-policy-name	String	The displayable name of the performance policy that was active during this reporting interval.

Table 106. Format of a workload-report-entry object (continued)

Field name	Type	Description
performance-policy-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific performance policy named in the <b>performance-policy-name</b> field. <b>Note:</b> If more than one performance policy was activated over the span of this reporting interval, only the name of the last one activated is returned in this field. The <b>multi-policy-activations</b> field indicates whether more than one policy was activated during this interval.
largest-pi-service-class-name	String	The displayable name of the performance management service class that had the largest PI (performance index) value over this reporting interval. <b>Optional:</b> This field is not returned if the performance index (PI) value could not be calculated.
largest-pi-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in <b>largest-pi-service-class-name</b> field. <b>Optional:</b> This field is not returned if the performance index (PI) value could not be calculated.
largest-pi	Float	The largest PI (performance index) value over this reporting interval. <b>Optional:</b> This field is not returned if the performance index (PI) value could not be calculated.
multi-policy-activations	Boolean	Indicates whether more than one performance policy was activated over the span of this reporting interval. The value is <b>"true"</b> if more than one policy or <b>"false"</b> if only one policy was activated during this time interval.
cpu-utilization-distribution	Array of objects	Array of nested cpu-utilization-range objects described in Table 107.
most-severe-perf-status	String Enum	The most severe <b>perf-status</b> value recorded for the workload over this reporting interval. See "Data model" on page 443 for more details about the valid values of this property.
perf-status-data-points	Array of objects	Array of nested perf-status-data-point objects described in Table 108.

Each nested cpu-utilization-range object contains the following fields:

Table 107. Format of a cpu-utilization-range object

Field name	Type	Description
low-boundary	Float	This value defines the low boundary of the CPU utilization range that this object is covering.
high-boundary	Float	This value defines the high boundary of the CPU utilization range that this object is covering.
virtual-server-count	Integer	This value is the number of virtual servers whose average CPU utilization was within the CPU utilization range during the reporting interval.

Each nested perf-status-data-point object contains the following fields:

Table 108. Format of a perf-status-data-point object

Field name	Type	Description
time	Integer	Standard date/time value indicating the date and time that this performance status value was recorded.  The standard time value is defined as the number of elapsed milliseconds after midnight January 1, 1970.

Table 108. Format of a perf-status-data-point object (continued)

Field name	Type	Description
perf-status	Float	The <b>perf-status</b> value of the workload at the recorded time. See the “Data model” on page 443 for more details about the valid values of this property.

## Description

The **Generate Workload Resource Groups Report** operation generates a report that contains the following information for the requested time interval:

- A list of all workload resource groups for which historical reporting information is available
- For each workload resource group:
  - The active performance policy at that time
  - A performance health indication (the performance index or PI) over that interval.

“Response body contents” on page 501 describes the full list of data that is returned in this report for each workload resource group. To request more detailed performance reporting data, you can use the returned **workload-report-id** property for a specific workload resource group as input for additional report operations.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

The request body is validated against the schema described in “Request body contents” on page 501. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Workloads Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 501.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-workload-resource-groups-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 30
}
```

---

Figure 268. Generate Workload Resource Groups Report: Request

## Generate Workload Resource Group Performance Index Report

Use the **Generate Workload Resource Group Performance Index Report** operation to create a custom on-demand performance index report for a specific workload resource group over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

```
POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-workload-resource-group-
performance-index-report
```

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble containing the workload resource group for which you want to receive a performance index report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
workload-report-id	String	Required	The identifier used by the performance management reporting to keep track of reporting data for the specific workload resource group for which the report is being requested. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the Workload Resource Groups Report for this same interval.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.



Field name	Type	Description
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field.
pi-data-by-service-class	Array of objects	Array of nested service-class-pi-data objects described in Table 109.

Each nested service-class-pi-data object contains the following fields:

*Table 109. Format of a service-class-pi-data object*

Field name	Type	Description
service-class-name	String	The displayable name of the service class.
service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>service-class-name</b> field.
pi-data-points	Array of objects	Array of nested pi-data-point objects described in Table 110.

Each nested pi-data-point object contains the following fields:

*Table 110. Format of a pi-data-point object*

Field name	Type	Description
pi-time	Integer	Standard date/time value indicating the requested starting date and time that this data point was taken. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
pi-value	Float	The specific performance index (PI) value recorded at the date/time in the <b>pi-time</b> field.

## Description

The **Generate Workload Resource Group Performance Index Report** operation generates a report that contains the following information for a specific workload resource group over the requested time interval:

- A list of the services classes within the requested workload resource group for which performance index data is available
- For each services class, a list of all of the individual PI data points that were recorded over that interval. Each PI data point contains both the actual PI value and the date/time that it was recorded.

“Response body contents” on page 504 describes the full list of data that is returned in this report for each workload resource group. To request more detailed performance reporting data, you can use the returned **service-class-report-id** property for a specific service class as input for additional report operations.

Note that when you request performance index information for a workload resource group through the HMC UI, the resulting report is a line graph display that plots these PI values with plot points and a line over time for each service class.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

The request body is validated against the schema described in “Request body contents” on page 504. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Workloads Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 504.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-workload-resource-group-performance-index-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group"
}
```

---

Figure 269. Generate Workload Resource Group Performance Index Report: Request

## Generate Workload Resource Group Resource Adjustments Report

Use the **Generate Workload Resource Group Resource Adjustments Report** operation to create a custom on-demand resource adjustments report for a specific workload resource group over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

**POST** /api/ensembles/{ensemble-id}/performance-management/operations/generate-workload-resource-group-resource-adjustments-report

In this request, the URI variable {ensemble-id} is the object ID of the ensemble containing the workload resource group for which you want to receive a resource adjustments report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Req/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
workload-report-id	String	Required	The identifier used by the performance management reporting to keep track of reporting data for the specific workload resource group for which the report is being requested. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the Workload Resource Groups Report for this same interval.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field.

Field name	Type	Description
successful-resource-adjustments	Array of objects	Array of nested successful-resource-adjustment objects. If no successful adjustments occurred during the requested time interval, an empty array is returned.
failed-resource-adjustments	Array of objects	Array of nested failed-resource-adjustment objects. If no failed adjustments occurred during the requested time interval, an empty array is returned.

Each nested successful-resource-adjustment object contains the following fields:

Field name	Type	Description
hypervisor-type	String Enum	The value of the <b>type</b> property of the virtualization host. See the virtualization host "Data model" on page 211 for more details about the valid values of this property.
adjustment-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
receiver-virtual-server-name	String	The name of the virtual server that was given additional resources.
receiver-workload-name	String	The displayable name of the workload resource group that was given additional resources.
receiver-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>receiver-workload-name</b> field.
receiver-service-class-name	String	The displayable name of the service class within the workload resource group that was given additional resources.
receiver-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>receiver-service-class-name</b> field.
receiver-processing-units-after	Float	The total number of processing units the receiver had after the additional resources were given to it.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
receiver-processing-units-before	Float	The total number of processing units the receiver had before the additional resources were given to it.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
adjustment-donors	Array of objects	Array of nested adjustment-donors objects.

Each nested adjustment-donors object contains the following fields:

Field name	Type	Description
donor-virtual-server-name	String	The name of a virtual server that gave up resources as part of this adjustment.

Field name	Type	Description
donor-processing-units-after	Float	The total number of processing units this donor had after it gave up resources.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
donor-processing-units-before	Float	The total number of processing units this donor had before it gave up resources.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
donor-workload-names	Array of strings	The displayable names of all workload resource groups that donated resources when this donor virtual server gave up resources.

Each nested failed-resource-adjustment object contains the following fields:

Field name	Type	Description
hypervisor-type	String Enum	The value of the <b>type</b> property of the virtualization host. See the virtualization host "Data model" on page 211 for more details about the valid values of this property.
adjustment-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
receiver-virtual-server-name	String	The name of the virtual server that needed but did not receive additional resources.
receiver-workload-name	String	The displayable name of the workload resource group that needed but did not receive additional resources.
receiver-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>receiver-workload-name</b> field.
receiver-service-class-name	String	The displayable name of the service class within the workload resource group that needed but did not receive additional resources.
receiver-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>receiver-service-class-name</b> field.
adjustment-fail-reason	String Enum	The reason why the adjustment failed. Values are: <ul style="list-style-type: none"> <li>• <b>"not-enough-capacity"</b> — The sum of the capacity that could be given up by all the available donors in the group was not enough to meet the capacity increase request of the intended receiver.</li> <li>• <b>"no-potential-donors"</b> — The group did not contain any potential donors so no assessments for increased capacity could be sent.</li> <li>• <b>"entitled-capacity-not-achievable"</b> — The extra capacity required to be added to the entitled capacity would have pushed the total beyond the maximum entitled capacity.</li> <li>• <b>"processor-not-fully-utilized"</b> — The virtual server is not fully using the processors it already has, so adding more will not have any effect.</li> <li>• <b>"requested-more-shares-than-max"</b> — More shares than the maximum allowed for this virtual server were requested.</li> <li>• <b>"not-enough-virtual-cpus"</b> — Not enough virtual CPUs were available for the projected required total consumed capacity to be achieved.</li> <li>• <b>"unknown"</b> — Unknown uncategorized failure reason.</li> </ul>

## Description

The **Generate Workload Resource Group Resource Adjustments Report** operation generates a report that contains the following information for a specific workload resource group over the requested time interval:

- A list of resource adjustments that zManager successfully made to maintain specified performance goals. For successful adjustments, the report includes:
  - A list of workload resource groups and the virtual servers that received additional resources (receivers)
  - A list of workload resource groups and the virtual servers that donated the additional resources (donors)
- A list of resource adjustments that zManager attempted but failed to make, along with a reason for the failure.

“Response body contents” on page 507 describes the full list of data that is returned in the report for this workload resource group.

Additional types of resource adjustment reports are available. This particular generate-workload-resource-group-resource-adjustments-report operation generates a resource adjustments report for a specific workload over the requested time period. What this means is that it will contain entries for all adjustments that affected the particular workload submitted; meaning all of the entries returned will involve that specific workload either having received additional resources, or having been forced to give up (donate) some of its resources.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

The request body is validated against the schema described in “Request body contents” on page 507. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Workloads Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 507.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-workload-resource-group-resource-adjustments-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group"
}
```

---

Figure 270. Generate Workload Resource Group Resource Adjustments Report: Request

## Generate Virtual Servers Report

Use the **Generate Virtual Servers Report** operation to create a custom on-demand report that shows all virtual servers that were members of a particular workload resource group over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

**POST** /api/ensembles/{*ensemble-id*}/performance-management/operations/generate-virtual-servers-report

In this request, the URI variable {*ensemble-id*} is the object ID of the ensemble object for which you are requesting a virtual server report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Req/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
workload-report-id	String	Required	The identifier used by the performance management reporting to keep track of reporting data for the specific workload resource group for which the report is being requested. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the Workload Resource Groups Report for this same interval.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field.
report-virtual-servers	Array of objects	Array of nested virtual-server-report-entry objects.

Each nested virtual-server-report-entry object contains the following fields:

Field name	Type	Description
virtual-server-name	String	The displayable name of the virtual server.
virtual-server-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific virtual server named in the <b>virtual-server-name</b> field.
hypervisor-type	String	The value of the <b>type</b> property of the virtualization host. See the virtualization host “Data model” on page 211 for more details about the valid values of this property.
hypervisor-name	String	The displayable name of the hypervisor under which this virtual server was running.
hypervisor-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific hypervisor under which this virtual server was running. (Hypervisor names returned in the <b>hypervisor-name</b> field are not unique).
defined-cpc-name	String	The name of the central processor complex (CPC) with which this virtual server was associated.
was-active	Boolean	Indicates whether this virtual server was actually active during this reporting interval. It will be true if it was active and false if it was not active. If this field is false, none of the fields in this next table this one can have data; so in that case none of the fields below here will be returned.
os-name	String	The name of the operating system image that is running on the virtual server as known to its operating system. <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.



Field name	Type	Description
os-type	String	The type of operating system that is running on the virtual server. <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
os-level	String	The release level of the operating system that is running on the virtual server. <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
hostname-ipaddr	String	The host name (or IP address) of the virtual server. <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
virtual-processors	String	The number of virtual processors associated with this virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
allocated-memory	String	For the x Hyp, PowerVM, and PR/SM hypervisor types, this field contains the allocated memory (in MB) configured for the virtual server. For z/VM, this field contains the amount of virtual memory currently resident in real memory for the guest. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
physical-cpu-utilization-percent	Float	The physical processor utilization percentage (%) of the virtual server. This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
hypervisor-cpu-delay-percent	Float	The hypervisor processing unit delay in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisors only: z/VM and PowerVM. In the case of z/VM, a value is available for this field only if sampling is turned on for the guest. <b>Optional:</b> This field is returned only if the virtual server was active during this interval and if the hypervisor type supports this information.
idle-time-percent	Float	The idle time percentage (%) of the virtual server. It is the percentage of total time that the virtual server had no work (that is, no application processes or internal hypervisor specific process states). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active during this interval and if the hypervisor type supports this information. For z/VM, this data is available only when sampling is enabled and started. For PowerVM and PR/SM, idle time data is not available.
other-time-percent	Float	The percentage (%) of total time that miscellaneous hypervisor specific internal process states had control for this virtual server. In other words, the percentage of time that the virtual server was not idle but also was not in a state of active CPU utilization or hypervisor CPU delay. This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active during this interval and if the hypervisor type supports this information. For z/VM, this data is available only when sampling is enabled and started. For PowerVM and PR/SM, other time data is not available.
service-class-name	String	The displayable name of the service class. <b>Optional:</b> This field is returned only if the virtual server was active and if the virtual server was associated with a specific service class within the requested workload resource group during this interval.
service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>service-class-name</b> field. <b>Optional:</b> This field is returned only if the <b>service-class-name</b> field is returned.

Field name	Type	Description
pi-value	Float	The average performance index (PI) value calculated over this reporting interval for the service class returned in the <b>service-class-name</b> field. This field is not returned if the virtual server was not active or if a PI value could not be calculated.
os-cpu-using-samples-percent	Float	The percentage of CPU using samples from among the total samples. For example, if there are 10 CPU using samples out of a total of 10 samples, then CPU using samples is 100% (because out of the total samples, all are CPU using samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
os-cpu-delay-samples-percent	Float	The percentage of CPU delay samples from among the total samples. For example, if there are 10 CPU delay samples and 10 samples that are not CPU delay samples, then CPU delay samples is 50% (because out of the total samples half are CPU delay samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
os-io-delay-samples-percent	Float	The percentage of I/O delay samples from among the total samples. The percent I/O delay is the percent of samples taken when work was delayed for non-paging DASD I/O. The I/O delay includes IOS queue, subchannel pending, and control unit queue delays. For example, if there are 10 I/O delay samples and 10 samples that are not I/O delay samples, then I/O delay samples is 50% (because out of the total samples half are I/O delay samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
os-page-delay-samples-percent	Float	The percentage of page delay samples from among the total samples. The percent page delay is the percent of samples when the address space experienced page faults in cross-memory access, and the page faults were resolved from auxiliary storage. For example, if there are 10 page delay samples and 10 samples that are not page delay samples, then page delay samples is 50% (because out of the total samples half are page delay samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.

## Description

The **Generate Virtual Servers Report** operation generates a report that contains the following information for a specific workload resource group over the requested time interval:

- A list of all virtual servers that were members in that specific workload resource group for which historical reporting information is available
- For each virtual server:
  - The unique performance reporting identifier for the virtual server
  - The name and unique performance reporting identifier of the hypervisor under which the virtual server was running
  - An indication of whether the virtual server was active over the reporting interval.

“Response body contents” on page 512 describes the full list of data that is returned in this report for each virtual server. To request more detailed performance reporting data, you can use the returned **virtual-server-report-id**, **hypervisor-report-id**, and **service-class-report-id** properties as input for additional report operations.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

The request body is validated against the schema described in “Request body contents” on page 511. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Virtual Servers Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 512.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-virtual-servers-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group"
}
```

---

Figure 271. Generate Virtual Servers Report: Request

## Generate Virtual Server CPU Utilization Report

Use the **Generate Virtual Server CPU Utilization Report** operation to create a custom on-demand report that shows the processor (CPU) utilization for a specific virtual server over a requested time period. The report is based on historical performance management data that was retained over that time period.

## HTTP method and URI

**POST** /api/ensembles/{*ensemble-id*}/performance-management/operations/generate-virtual-server-cpu-utilization-report

In this request, the URI variable {*ensemble-id*} is the object ID of the ensemble object for which you are requesting a virtual server CPU utilization report.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
virtual-server-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific virtual server named in the <b>virtual-server-name</b> field. This value is the same as one of the <b>virtual-server-report-id</b> values that are returned in the Virtual Servers Report. Alternatively, you can supply the <b>object-id</b> property value of an existing Virtual Server object.
hypervisor-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific hypervisor under which this virtual server was running. (Hypervisor names returned in the <b>hypervisor-name</b> field are not unique). This value is the same as the <b>hypervisor-report-id</b> value that is returned for this virtual server in the Virtual Servers Report. Alternatively, you can supply the <b>object-id</b> property value of an existing Virtualization Host object.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
virtual-server-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific virtual server named in the <b>virtual-server-name</b> field.
virtual-server-name	String	The displayable name of the virtual server.

Field name	Type	Description
hypervisor-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific hypervisor under which this virtual server was running. (Hypervisor names returned in the <b>hypervisor-name</b> field are not unique).
hypervisor-name	String	The displayable name of the hypervisor under which this virtual server was running.
cpu-utilization-data-points	Array of objects	Array of nested cpu-utilization-data-points objects.

Each nested cpu-utilization-data-points object contains the following fields:

Field name	Type	Description
cpu-utilization-time	Integer	Standard date/time value indicating the requested starting date and time that this data point was taken. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
cpu-utilization-value	Float	The specific CPU utilization value recorded at the date/time in the <b>cpu-utilization-time</b> field. This value is in fractional form representing a percentage (between 0 and 1 inclusive).

## Description

The **Generate Virtual Server CPU Utilization Report** operation generates a report that contains the following information for a specific virtual server over the requested time interval:

- All of the individual CPU utilization data points for that virtual server that were recorded over the requested reporting interval. Each CPU utilization data point contains both the actual CPU utilization value and the date/time that it was recorded.

“Response body contents” on page 516 describes the full list of data that is returned in this report for each virtual server.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

The request body is validated against the schema described in “Request body contents” on page 516. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Virtual Servers Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 516.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-virtual-server-cpu-utilization-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "virtual-server-report-id": "vs241510-5678900000-xxxxx",
  "hypervisor-report-id": "phyp241510-5678900000-xxxxx"
}
```

---

Figure 272. Generate Virtual Server CPU Utilization Report: Request

## Generate Virtual Server Resource Adjustments Report

Use the **Generate Virtual Server Resource Adjustments Report** operation to create a custom on-demand resource adjustments report for a specific virtual server over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

```
POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-virtual-server-
resource-adjustments-report
```

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble containing the workload resource group for which you want to receive a virtual server resource adjustments report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.

Field name	Type	Rqd/Opt	Description
workload-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the Workload Resource Groups Report for this same interval.
virtual-server-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific virtual server named in the <b>virtual-server-name</b> field. This value is the same as one of the <b>virtual-server-report-id</b> values that are returned in the Virtual Servers Report. Alternatively, you can supply the <b>object-id</b> property value of an existing Virtual Server object.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field.
virtual-server-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific virtual server named in the <b>virtual-server-name</b> field.
successful-resource-adjustments	Array of objects	Array of nested <b>successful-resource-adjustment</b> objects. If no successful adjustments occurred during the requested time interval, an empty array is returned.
failed-resource-adjustments	Array of objects	Array of nested <b>failed-resource-adjustment</b> objects. If no failed adjustments occurred during the requested time interval, an empty array is returned.

Each nested **successful-resource-adjustment** object contains the following fields:

Field name	Type	Description
hypervisor-type	String Enum	The value of the <b>type</b> property of the virtualization host. See the virtualization host "Data model" on page 211 for more details about the valid values of this property.
adjustment-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.

Field name	Type	Description
receiver-virtual-server-name	String	The name of the virtual server that was given additional resources.
receiver-workload-name	String	The displayable name of the workload resource group that was given additional resources.
receiver-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>receiver-workload-name</b> field.
receiver-service-class-name	String	The displayable name of the service class within the workload resource group that was given additional resources.
receiver-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>receiver-service-class-name</b> field.
receiver-processing-units-after	Float	The total number of processing units the receiver had after the additional resources were given to it.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
receiver-processing-units-before	Float	The total number of processing units the receiver had before the additional resources were given to it.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
adjustment-donors	Array of objects	Array of nested adjustment-donors objects.

Each nested adjustment-donors object contains the following fields:

Field name	Type	Description
donor-virtual-server-name	String	The name of a virtual server that gave up resources as part of this adjustment.
donor-processing-units-after	Float	The total number of processing units this donor had after it gave up resources.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
donor-processing-units-before	Float	The total number of processing units this donor had before it gave up resources.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
donor-workload-names	Array of strings	The displayable names of all workload resource groups that donated resources when this donor virtual server gave up resources.

Each nested failed-resource-adjustment object contains the following fields:



Field name	Type	Description
hypervisor-type	String Enum	The value of the <b>type</b> property of the virtualization host. See the virtualization host “Data model” on page 211 for more details about the valid values of this property.
adjustment-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
receiver-virtual-server-name	String	The name of the virtual server that needed but did not receive additional resources.
receiver-workload-name	String	The displayable name of the workload resource group that needed but did not receive additional resources.
receiver-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>receiver-workload-name</b> field.
receiver-service-class-name	String	The displayable name of the service class within the workload resource group that needed but did not receive additional resources.
receiver-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>receiver-service-class-name</b> field.
adjustment-fail-reason	String Enum	The reason why the adjustment failed. Values are: <ul style="list-style-type: none"> <li>• <b>"not-enough-capacity"</b> — The sum of the capacity that could be given up by all the available donors in the group was not enough to meet the capacity increase request of the intended receiver.</li> <li>• <b>"no-potential-donors"</b> — The group did not contain any potential donors so no assessments for increased capacity could be sent.</li> <li>• <b>"entitled-capacity-not-achievable"</b> — The extra capacity required to be added to the entitled capacity would have pushed the total beyond the maximum entitled capacity.</li> <li>• <b>"processor-not-fully-utilized"</b> — The virtual server is not fully using the processors it already has, so adding more will not have any effect.</li> <li>• <b>"requested-more-shares-than-max"</b> — More shares than the maximum allowed for this virtual server were requested.</li> <li>• <b>"not-enough-virtual-cpus"</b> — Not enough virtual CPUs were available for the projected required total consumed capacity to be achieved.</li> <li>• <b>"unknown"</b> — Unknown uncategorized failure reason.</li> </ul>

## Description

The **Generate Virtual Server Resource Adjustments Report** operation generates a report that contains the following information for a specific virtual server over the requested time interval:

- A list of resource adjustments that zManager successfully made to maintain specified performance goals. For successful adjustments, the report includes:
  - A list of workload resource groups and the virtual servers that received additional resources (receivers)
  - A list of workload resource groups and the virtual servers that donated the additional resources (donors)
- A list of resource adjustments that zManager attempted but failed to make, along with a reason for the failure.

“Response body contents” on page 519 describes the full list of data that is returned in the report for this workload resource group.

Additional types of resource adjustment reports are available. This particular generate-virtual-server-resource-adjustments-report operation generates a resource adjustments report for a specific virtual server within a specific workload over the requested time period. What this means is that it will contain entries

for all adjustments that affected the particular virtual server submitted; meaning all of the entries returned will involve that specific virtual server either having received additional resources, or having been forced to give up (donate) some of its resources.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

The request body is validated against the schema described in “Request body contents” on page 518. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Virtual Server Resource Adjustments Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 519.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-virtual-server-resource-adjustments-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group",
  "virtual-server-report-id": "vs241510-5678900000-xxxxx"
}
```

---

Figure 273. Generate Virtual Server Resource Adjustments Report: Request

## Generate Hypervisor Report

Use the **Generate Hypervisor Report** operation to create a custom on-demand report that shows information about a specific hypervisor and all virtual servers that were running under that hypervisor over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

**POST** /api/ensembles/{ensemble-id}/performance-management/operations/generate-hypervisor-report

In this request, the URI variable {ensemble-id} is the object ID of the ensemble object for which you are requesting a hypervisor report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
hypervisor-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific hypervisor under which this virtual server was running. (Hypervisor names returned in the <b>hypervisor-name</b> field are not unique).

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
hypervisor-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific hypervisor under which this virtual server was running. (Hypervisor names returned in the <b>hypervisor-name</b> field are not unique).
hypervisor-name	String	The displayable name of the hypervisor under which this virtual server was running.

Field name	Type	Description
hypervisor-type	String Enum	The value of the <b>type</b> property of the virtualization host. See the virtualization host "Data model" on page 211 for more details about the valid values of this property.
report-hypervisor-details	Object	An object that contains report details for the hypervisor itself. The format of the returned object is described in Table 111.
report-hypervisor-virtual-servers	Array of objects	An array of nested objects, each representing a virtual server that was under the control of this hypervisor during this reporting interval. The format of the returned object depends on the value returned in the hypervisor-type field: <ul style="list-style-type: none"> <li>• For "<b>power-vm</b>", see Table 112 on page 525</li> <li>• For "<b>x-hyp</b>", see Table 113 on page 526</li> <li>• For "<b>zvm</b>", see Table 114 on page 527</li> <li>• For "<b>prsm</b>", see Table 115 on page 529</li> </ul>
resource-adjustment-report	Object	If a hypervisor-level resource adjustment report is available for this same reporting interval, this object is returned. The format of the resource-adjustment-report object is the same as that described in "Response body contents" on page 532 for the <b>Generate Hypervisor Resource Adjustments Report</b> operation.

The report-hypervisor-details object for a hypervisor contains the fields in Table 111.

*Table 111. Format of a report-hypervisor-details object*

Field name	Type	Description
allocated-processing-units	Float	The total number of processing units that were allocated. This field is included in the object only for the PowerVM hypervisor type.
processor-count	Integer	The total number of physical processors.
cpu-consumption-percent	Float	The total consumption (utilization) percentage (%) of hypervisor CPUs. This percentage is in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisors only: z/VM, z Hyp, and PowerVM.
shared-cp-consumption-percent	Float	The total consumption (utilization) percentage (%) of shared processors. This percentage is in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisor only: PR/SM.
memory-used	Integer	The total amount of the hypervisor memory that was in use (in MB).
total-memory	Integer	The total amount of memory (in MB) that was that was available to the hypervisor.
cp-cpu-consumption-percent	Float	The total consumption (utilization) percentage (%) of the hypervisor's general purpose CPUs. It is in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisor only: z/VM.
ifl-cpu-consumption-percent	Float	The total consumption (utilization) percentage (%) of the hypervisor's IFL CPUs. It is in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisor only: z/VM.
other-cpu-consumption-percent	Float	The total consumption (utilization) percentage (%) of the hypervisor's CPU types other than CPs and IFLs. It is in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisor only: z/VM.

Each nested report-hypervisor-virtual-servers object for a PowerVM hypervisor contains the fields in Table 112 on page 525.

Table 112. Format of a PowerVM report-hypervisor-virtual-servers object

Field name	Type	Description
virtual-server-name	String	The displayable name of the virtual server.
workload-processor-mgmt-status	String Enum	The processor management status, which is one of the following values: <ul style="list-style-type: none"> <li>• "active"</li> <li>• "not-active"</li> </ul>
workload-processor-mgmt-status-reason	String Enum	A further explanation of the reason for processor management status returned in the <b>workload-processor-mgmt-status</b> field. Values are: <ul style="list-style-type: none"> <li>• "none"– the value of <b>workload-processor-mgmt-status</b> is "active"</li> <li>• "mgmt-disabled-global"– the option to enable processor performance management at the hypervisor type level was not set for this hypervisor type</li> <li>• "mgmt-disabled-vs"– the option to enable processor performance management at the virtual server level was not set for this virtual server</li> <li>• "dedicated-proc-mode"– the virtual server is running in dedicated processor mode, and only shared mode is supported for processor management</li> <li>• "internal-error"– An internal error has occurred.</li> <li>• "network-connection-failure"– the virtual server has no connectivity to its hypervisor.</li> <li>• "virtual-server-not-active"– the virtual server itself was not up and active during this reporting interval</li> </ul>
was-active	Boolean	Indicates whether this virtual server was active during this reporting interval. The value is "true" if the virtual server was active or "false" if it was not active.  If the value for this field is false, none of the remaining fields in this object can have data so they are not returned.
virtual-processor-count	Integer	This value is the number of virtual processors that were associated with this virtual server during the reporting interval. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
min-virtual-processors	Integer	The minimum number of virtual processors allowed for this virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
max-virtual-processors	Integer	The maximum number of virtual processors allowed for this virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
consumed-processors	Float	The consumed processors metric for this virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
was-dedicated	Boolean	Indicates whether this virtual server was running in dedicated mode (as opposed to shared mode) during this reporting interval. The value is "true" if the virtual server was in dedicated mode or "false" if it was in shared mode. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
was-capped	Boolean	Indicates whether this virtual server was capped during this reporting interval. The value is "true" if the virtual server was capped or "false" if it was not capped. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
hypervisor-cpu-delay-percent	Float	The hypervisor processing unit delay in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisors only: z/VM and PowerVM. In the case of z/VM, a value is available for this field only if sampling is turned on for the guest. This field is returned only if the virtual server was active and running in shared mode (as opposed to dedicated mode) during this interval.

Table 112. Format of a PowerVM report-hypervisor-virtual-servers object (continued)

Field name	Type	Description
processing-units	Float	The number of processing units that were assigned to this virtual server during the reporting interval. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
initial-processing-units	Float	The number of processing units that were initially assigned to this virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
min-processing-units	Float	The minimum number of processing units that this virtual server could use. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
max-processing-units	Float	The maximum number of processing units that this virtual server could use. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
min-memory	Integer	The minimum amount of memory (in MB) that this virtual server could use. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
max-memory	Integer	The maximum amount of memory (in MB) that this virtual server could use. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.

Each nested report-hypervisor-virtual-servers object for an x Hyp hypervisor contains the fields in Table 113.

Table 113. Format of an x Hyp report-hypervisor-virtual-servers object

Field name	Type	Description
virtual-server-name	String	The displayable name of the virtual server.
was-active	Boolean	Indicates whether this virtual server was active during this reporting interval. The value is <b>"true"</b> if the virtual server was active or <b>"false"</b> if it was not active.  If the value for this field is false, none of the remaining fields in this object can have data so they are not returned.
virtual-processor-count	Integer	This value is the number of virtual processors that were associated with this virtual server during the reporting interval. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
consumed-processors	Float	The consumed processors metric for this virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
hypervisor-cpu-delay-percent	Float	The hypervisor processing unit delay in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisors only: z/VM and PowerVM. In the case of z/VM, a value is available for this field only if sampling is turned on for the guest. <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
memory-in-use	Integer	The amount of memory (in MB) that was actually in use by this virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
allocated-memory	Integer	For the x Hyp, PowerVM, and PR/SM hypervisor types, this field contains the allocated memory (in MB) configured for the virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.

Table 113. Format of an x Hyp report-hypervisor-virtual-servers object (continued)

Field name	Type	Description
workload-processor-mgmt-status	String Enum	The status of automatic workload processor performance management during this reporting interval.  Values: <ul style="list-style-type: none"> <li>• "active"</li> <li>• "not-active"</li> </ul>
workload-processor-mgmt-status-reason	String Enum	A further explanation of the reason for workload processor performance management status returned in <b>workload-mgmt-status</b> .  Values: <ul style="list-style-type: none"> <li>• "none" - No further reason available (for example, would be "none" when <b>workload-processor-mgmt-status</b> is "active").</li> <li>• "mgmt-disabled-global" - The option to enable processor performance management at the hypervisor type level was not set for this hypervisor type.</li> <li>• "mgmt-disabled-vs" - The option to enable processor performance management at the virtual server level was not set for this virtual server.</li> <li>• "virtual-server-not-active" - The virtual server itself was not up and active during this reporting interval.</li> <li>• "hypervisor-not-correct-level" - The hypervisor controlling this virtual server is not the correct level to support automatic performance management.</li> <li>• "internal-error" - An internal error has been encountered (contact IBM support).</li> </ul>
initial-shares	Integer	The number of processor shares that were initially associated to this virtual server during this reporting interval. <b>Optional:</b> This data may not always be available (for example, if <b>was-active</b> is false). In any case, if it is not available this field will not be returned.
shares	Integer	The number of processor shares that were currently associated to this virtual server during this reporting interval. <b>Optional:</b> This data may not always be available (for example, if <b>was-active</b> is false). In any case, if it is not available this field will not be returned.
min-shares	Integer	The minimum number of processor shares that could be dynamically assigned to the virtual server. <b>Optional:</b> This data may not always be available (for example, if <b>was-active</b> is false). In any case, if it is not available this field will not be returned.
max-shares	Integer	The maximum number of processor shares that could be dynamically assigned to the virtual server. <b>Optional:</b> This data may not always be available (for example, if <b>was-active</b> is false). In any case, if it is not available this field will not be returned.

Each nested report-hypervisor-virtual-servers object for a z/VM hypervisor contains the fields in Table 114.

Table 114. Format of a z/VM report-hypervisor-virtual-servers object

Field name	Type	Description
virtual-server-name	String	The displayable name of the virtual server.
workload-processor-mgmt-status	String Enum	The processor management status, which is one of the following values: <ul style="list-style-type: none"> <li>• "active"</li> <li>• "not-active"</li> </ul>

Table 114. Format of a z/VM report-hypervisor-virtual-servers object (continued)

Field name	Type	Description
workload-processor-mgmt-status-reason	String Enum	A further explanation of the reason for processor management status returned in the <b>workload-processor-mgmt-status</b> field. Values are: <ul style="list-style-type: none"> <li>• <b>"none"</b>– the value of <b>workload-processor-mgmt-status</b> is <b>"active"</b></li> <li>• <b>"mgmt-disabled-global"</b>– the option to enable processor performance management at the hypervisor type level was not set for this hypervisor type</li> <li>• <b>"mgmt-disabled-vs"</b>– the option to enable processor performance management at the virtual server level was not set for this virtual server</li> <li>• <b>"absolute-share-mode"</b>– the virtual server is running in absolute share mode, and only relative share mode is supported for processor management</li> <li>• <b>"sampling-disabled"</b>– z/VM sampling was not enabled. Sampling must be enabled for processor management under z/VM.</li> <li>• <b>"internal-error"</b>– An internal error has occurred.</li> <li>• <b>"network-connection-failure"</b>– the virtual server has no connectivity to its hypervisor.</li> <li>• <b>"virtual-server-not-active"</b>– the virtual server itself was not up and active during this reporting interval</li> </ul>
was-active	Boolean	Indicates whether this virtual server was active during this reporting interval. The value is <b>"true"</b> if the virtual server was active or <b>"false"</b> if it was not active.  If the value for this field is false, none of the remaining fields in this object can have data so they are not returned.
virtual-processor-count	Integer	This value is the number of virtual processors that were associated with this virtual server during the reporting interval. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
consumed-processors	Float	The consumed processors metric for this virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
hypervisor-cpu-delay-percent	Float	The hypervisor processing unit delay in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisors only: z/VM and PowerVM. In the case of z/VM, a value is available for this field only if sampling is turned on for the guest. This field is returned only if the virtual server was active and sampling is enabled during this interval.
share-mode	String Enum	The processor share mode configured for this virtual server during this reporting interval. Valid values are <b>"absolute"</b> and <b>"relative"</b> . <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
share-limit	String Enum	The processor share limit configured for this virtual server during this reporting interval. Valid values are <b>"soft"</b> , <b>"hard"</b> and <b>"none"</b> . <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
shares	Integer	The number of processor shares that were associated with this virtual server during this reporting interval. <b>Optional:</b> This field is returned only if the virtual server was active and running in relative share mode (as opposed to absolute share mode) during this interval.
min-shares	Integer	The minimum number of processor shares that could be dynamically assigned to the virtual server during this reporting interval. <b>Optional:</b> This field is returned only if the virtual server was active and running in relative share mode (as opposed to absolute share mode) during this interval.
max-shares	Integer	The maximum number of processor shares that could be dynamically assigned to the virtual server during this reporting interval. <b>Optional:</b> This field is returned only if the virtual server was active and running in relative share mode (as opposed to absolute share mode) during this interval.



Table 114. Format of a z/VM report-hypervisor-virtual-servers object (continued)

Field name	Type	Description
initial-shares	Integer	The number of processor shares that were initially associated to this virtual server during this reporting interval. <b>Optional:</b> This data may not always be available (for example, if <b>was-active</b> is false). In any case, if it is not available this field will not be returned.
memory-used	Integer	The total amount of the hypervisor memory that was in use (in MB). <b>Optional:</b> This field is returned only if the virtual server was active during this interval.

Each nested report-hypervisor-virtual-servers object for a PR/SM hypervisor contains the fields in Table 115.

Table 115. Format of a PR/SM report-hypervisor-virtual-servers object

Field name	Type	Description
virtual-server-name	String	The displayable name of the virtual server.
was-active	Boolean	Indicates whether this virtual server was active during this reporting interval. The value is <b>"true"</b> if the virtual server was active or <b>"false"</b> if it was not active.  If the value for this field is false, none of the remaining fields in this object can have data so they are not returned.
logical-processor-count	Integer	The number of logical processors that were associated with this virtual server during this reporting interval. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
consumed-processors	Float	The consumed processors metric for this virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
allocated-memory	Integer	For the x Hyp, PowerVM, and PR/SM hypervisor types, this field contains the allocated memory (in MB) configured for the virtual server. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
was-dedicated	Boolean	Indicates whether this virtual server was running in dedicated mode (as opposed to shared mode) during this reporting interval. The value is <b>"true"</b> if the virtual server was in dedicated mode or <b>"false"</b> if it was in shared mode. <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
cpu-weight	Integer	The CPU weight that was associated to this virtual server during this reporting interval. This field is returned only if the virtual server was active and running in shared mode (as opposed to dedicated mode) during this interval.
min-cpu-weight	Integer	The minimum CPU weight that this virtual server could use. This field is returned only if the virtual server was active and running in shared mode (as opposed to dedicated mode) during this interval.
max-cpu-weight	Integer	The maximum CPU weight that this virtual server could use. This field is returned only if the virtual server was active and running in shared mode (as opposed to dedicated mode) during this interval.

## Description

The **Generate Hypervisor Report** operation generates a report that contains the following information for a specific hypervisor over the requested time interval:

- Information about the hypervisor itself, including the total number of physical processors and total consumption percentage of those processors

- A list of all virtual servers that were running under this specific hypervisor for which historical reporting information is available
- For each virtual server:
  - An indication of whether the virtual server was active over the reporting interval
  - Additional information such as processor counts and memory statistics.

“Response body contents” on page 523 describes the full list of data that is returned in this report for each virtual server. The information available for virtual servers varies depending on the hypervisor type.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

The request body is validated against the schema described in “Request body contents” on page 523. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Hypervisor Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 523.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-hypervisor-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "hypervisor-report-id": "phyp241510-5678900000-xxxxx"
}
```

---

Figure 274. Generate Hypervisor Report: Request

## Generate Hypervisor Resource Adjustments Report

Use the **Generate Hypervisor Resource Adjustments Report** operation to create a custom on-demand resource adjustments report for a specific hypervisor over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

**POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-hypervisor-resource-adjustments-report**

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble containing the workload resource group for which you want to receive a hypervisor resource adjustments report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
hypervisor-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific hypervisor under which this virtual server was running. (Hypervisor names returned in the <b>hypervisor-name</b> field are not unique). This value is the same as the <b>hypervisor-report-id</b> value that is returned for this virtual server in the Virtual Servers Report. Alternatively, you can supply the <b>object-id</b> property value of an existing Virtualization Host object.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
hypervisor-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific hypervisor under which this virtual server was running. (Hypervisor names returned in the <b>hypervisor-name</b> field are not unique).
successful-resource-adjustments	Array of objects	Array of nested successful-resource-adjustment objects. If no successful adjustments occurred during the requested time interval, an empty array is returned.
failed-resource-adjustments	Array of objects	Array of nested failed-resource-adjustment objects. If no failed adjustments occurred during the requested time interval, an empty array is returned.

Each nested successful-resource-adjustment object contains the following fields:

Field name	Type	Description
hypervisor-type	String Enum	The value of the <b>type</b> property of the virtualization host. See the virtualization host "Data model" on page 211 for more details about the valid values of this property.
adjustment-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
receiver-virtual-server-name	String	The name of the virtual server that was given additional resources.
receiver-workload-name	String	The displayable name of the workload resource group that was given additional resources.
receiver-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>receiver-workload-name</b> field.
receiver-service-class-name	String	The displayable name of the service class within the workload resource group that was given additional resources.
receiver-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>receiver-service-class-name</b> field.

Field name	Type	Description
receiver-processing-units-after	Float	The total number of processing units the receiver had after the additional resources were given to it.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
receiver-processing-units-before	Float	The total number of processing units the receiver had before the additional resources were given to it.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
adjustment-donors	Array of objects	Array of nested adjustment-donors objects.

Each nested adjustment-donors object contains the following fields:

Field name	Type	Description
donor-virtual-server-name	String	The name of a virtual server that gave up resources as part of this adjustment.
donor-processing-units-after	Float	The total number of processing units this donor had after it gave up resources.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
donor-processing-units-before	Float	The total number of processing units this donor had before it gave up resources.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
donor-workload-names	Array of strings	The displayable names of all workload resource groups that donated resources when this donor virtual server gave up resources.

Each nested failed-resource-adjustment object contains the following fields:

Field name	Type	Description
hypervisor-type	String Enum	The value of the <b>type</b> property of the virtualization host. See the virtualization host "Data model" on page 211 for more details about the valid values of this property.
adjustment-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
receiver-virtual-server-name	String	The name of the virtual server that needed but did not receive additional resources.
receiver-workload-name	String	The displayable name of the workload resource group that needed but did not receive additional resources.
receiver-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>receiver-workload-name</b> field.

Field name	Type	Description
receiver-service-class-name	String	The displayable name of the service class within the workload resource group that needed but did not receive additional resources.
receiver-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>receiver-service-class-name</b> field.
adjustment-fail-reason	String Enum	The reason why the adjustment failed. Values are: <ul style="list-style-type: none"> <li>• <b>"not-enough-capacity"</b> — The sum of the capacity that could be given up by all the available donors in the group was not enough to meet the capacity increase request of the intended receiver.</li> <li>• <b>"no-potential-donors"</b> — The group did not contain any potential donors so no assessments for increased capacity could be sent.</li> <li>• <b>"entitled-capacity-not-achievable"</b> — The extra capacity required to be added to the entitled capacity would have pushed the total beyond the maximum entitled capacity.</li> <li>• <b>"processor-not-fully-utilized"</b> — The virtual server is not fully using the processors it already has, so adding more will not have any effect.</li> <li>• <b>"requested-more-shares-than-max"</b> — More shares than the maximum allowed for this virtual server were requested.</li> <li>• <b>"not-enough-virtual-cpus"</b> — Not enough virtual CPUs were available for the projected required total consumed capacity to be achieved.</li> <li>• <b>"unknown"</b> — Unknown uncategorized failure reason.</li> </ul>

## Description

The **Generate Hypervisor Resource Adjustments Report** operation generates a report that contains the following information for a specific hypervisor over the requested time interval:

- A list of resource adjustments that zManager successfully made to maintain specified performance goals. For successful adjustments, the report includes:
  - A list of workload resource groups and the virtual servers that received additional resources (receivers)
  - A list of workload resource groups and the virtual servers that donated the additional resources (donors)
- A list of resource adjustments that zManager attempted but failed to make, along with a reason for the failure.

“Response body contents” on page 532 describes the full list of data that is returned in the report for this workload resource group.

Additional types of resource adjustment reports are available. This particular generate-hypervisor-resource-adjustments-report operation generates a resource adjustments report for a specific hypervisor over the requested time period. What this means is that it will contain entries for all adjustments that affected virtual servers within the particular hypervisor submitted; meaning all of the entries returned will involve a virtual server under the control of that specific hypervisor either having received additional resources, or having been forced to give up (donate) some of its resources.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements” on page 535.

The request body is validated against the schema described in “Request body contents” on page 531. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Hypervisor Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 532.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-hypervisor-resource-adjustments-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "hypervisor-report-id": "phyp241510-5678900000-xxxxx"
}
```

---

Figure 275. Generate Hypervisor Resource Adjustments Report: Request

## Generate Service Classes Report

Use the **Generate Service Classes Report** operation to create a custom on-demand report that shows all service classes within a particular workload resource group over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

```
POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-service-classes-report
```

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble object for which you are requesting a service classes report.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Req/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
workload-report-id	String	Required	The identifier used by the performance management reporting to keep track of reporting data for the specific workload resource group for which the report is being requested. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the Workload Resource Groups Report for this same interval.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field.
performance-policy-name	String	The displayable name of the performance policy that contains this service class.
performance-policy-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific performance policy named in the <b>performance-policy-name</b> field.
report-service-classes	Array of objects	Array of nested service-class-report-entry objects.

Each nested service-class-report-entry object contains the following fields:



Field name	Type	Description
service-class-name	String	The displayable name of the service class. <b>Optional:</b> This field is returned only if the virtual server was active and if the virtual server was associated with a specific service class within the requested workload resource group during this interval.
service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>service-class-name</b> field. <b>Optional:</b> This field is returned only if the <b>service-class-name</b> field is returned.
has-hop-data	Boolean	Indicates whether this service class has hop level data available for this reporting interval. The value is <b>"true"</b> if hop data is available or <b>"false"</b> if it is not available.
goal-type	String Enum	The type of performance goal that was defined for this service class. Values are <b>"velocity"</b> or <b>"discretionary"</b> .
goal-performance-level	String Enum	The performance goal of the service class as defined in performance policy. Possible values are: <b>"fastest"</b> , <b>"fast"</b> , <b>"moderate"</b> , <b>"slow"</b> , and <b>"slowest"</b> . This field is returned only if a velocity performance goal was defined for this service class.
business-importance	String Enum	The business importance level that was defined for this service class. Possible values are: <b>"highest"</b> , <b>"high"</b> , <b>"medium"</b> , <b>"low"</b> , and <b>"lowest"</b> . This field is returned only if a velocity performance goal was defined for this service class.
pi-value	Float	The average performance index (PI) value calculated over this reporting interval for the service class returned in the <b>service-class-name</b> field. This field is not returned if the virtual server was not active or if a PI value could not be calculated.
actual-performance-level	String Enum	The average level of performance that was actually measured over this reporting interval. Possible values are: <b>"fastest"</b> , <b>"fast"</b> , <b>"moderate"</b> , <b>"slow"</b> , and <b>"slowest"</b> . This field is returned only if a velocity performance goal was defined for this service class and if the performance index value could be measured.

## Description

The **Generate Service Classes Report** operation generates a report that contains the following information for a specific workload resource group over the requested time interval:

- A list of all service classes within that specific workload resource group for which historical reporting information is available
- For each service class:
  - The name and unique reporting ID of the performance policy that contains the service class
  - An indication of whether the service class has available hop data for this reporting interval
  - Details about the service class definition, such as the type of performance goal
  - The actual performance level achieved during this reporting interval.

“Response body contents” on page 536 describes the full list of data that is returned in this report for each service class. If hop data is available for this service class (check the **has-hop-data** field), you can use the **service-class-report-id** property as input for additional report operations related to service classes:

- “Generate Service Class Hops Report” on page 543
- “Generate Service Class Virtual Server Topology Report” on page 548

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements” on page 538.

The request body is validated against the schema described in “Request body contents” on page 536. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Service Classes Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 536.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-service-classes-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group"
}
```

---

Figure 276. Generate Service Classes Report: Request

## Generate Service Class Resource Adjustments Report

Use the **Generate Service Class Resource Adjustments Report** operation to create a custom on-demand resource adjustments report for a specific service class within a workload resource group over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

```
POST /api/ensembles/{ensemble-id}/performance-management/operations/
generate-service-class-resource-adjustments-report
```

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble containing the workload resource group for which you want to receive a service class resource adjustments report.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
workload-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the Workload Resource Groups Report for this same interval.
service-class-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>service-class-name</b> field. <b>Note:</b> This value is the same as the <b>service-class-report-id</b> value that is returned in one of the service class entries in the Service Classes Report for this same interval.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field.
service-class-name	String	The displayable name of the service class.
service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>service-class-name</b> field.

Field name	Type	Description
successful-resource-adjustments	Array of objects	Array of nested successful-resource-adjustment objects. If no successful adjustments occurred during the requested time interval, an empty array is returned.
failed-resource-adjustments	Array of objects	Array of nested failed-resource-adjustment objects. If no failed adjustments occurred during the requested time interval, an empty array is returned.

Each nested successful-resource-adjustment object contains the following fields:

Field name	Type	Description
hypervisor-type	String Enum	The value of the <b>type</b> property of the virtualization host. See the virtualization host "Data model" on page 211 for more details about the valid values of this property.
adjustment-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
receiver-virtual-server-name	String	The name of the virtual server that was given additional resources.
receiver-workload-name	String	The displayable name of the workload resource group that was given additional resources.
receiver-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>receiver-workload-name</b> field.
receiver-service-class-name	String	The displayable name of the service class within the workload resource group that was given additional resources.
receiver-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>receiver-service-class-name</b> field.
receiver-processing-units-after	Float	The total number of processing units the receiver had after the additional resources were given to it.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
receiver-processing-units-before	Float	The total number of processing units the receiver had before the additional resources were given to it.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
adjustment-donors	Array of objects	Array of nested adjustment-donors objects.

Each nested adjustment-donors object contains the following fields:

Field name	Type	Description
donor-virtual-server-name	String	The name of a virtual server that gave up resources as part of this adjustment.

Field name	Type	Description
donor-processing-units-after	Float	The total number of processing units this donor had after it gave up resources.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
donor-processing-units-before	Float	The total number of processing units this donor had before it gave up resources.  Processing units are returned in the units that are normal for the specific hypervisor type of the virtual server. The units are entitled capacity for PowerVM, and CPU shares for all other hypervisor types.
donor-workload-names	Array of strings	The displayable names of all workload resource groups that donated resources when this donor virtual server gave up resources.

Each nested failed-resource-adjustment object contains the following fields:

Field name	Type	Description
hypervisor-type	String Enum	The value of the <b>type</b> property of the virtualization host. See the virtualization host "Data model" on page 211 for more details about the valid values of this property.
adjustment-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
receiver-virtual-server-name	String	The name of the virtual server that needed but did not receive additional resources.
receiver-workload-name	String	The displayable name of the workload resource group that needed but did not receive additional resources.
receiver-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>receiver-workload-name</b> field.
receiver-service-class-name	String	The displayable name of the service class within the workload resource group that needed but did not receive additional resources.
receiver-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>receiver-service-class-name</b> field.
adjustment-fail-reason	String Enum	The reason why the adjustment failed. Values are: <ul style="list-style-type: none"> <li>• <b>"not-enough-capacity"</b> — The sum of the capacity that could be given up by all the available donors in the group was not enough to meet the capacity increase request of the intended receiver.</li> <li>• <b>"no-potential-donors"</b> — The group did not contain any potential donors so no assessments for increased capacity could be sent.</li> <li>• <b>"entitled-capacity-not-achievable"</b> — The extra capacity required to be added to the entitled capacity would have pushed the total beyond the maximum entitled capacity.</li> <li>• <b>"processor-not-fully-utilized"</b> — The virtual server is not fully using the processors it already has, so adding more will not have any effect.</li> <li>• <b>"requested-more-shares-than-max"</b> — More shares than the maximum allowed for this virtual server were requested.</li> <li>• <b>"not-enough-virtual-cpus"</b> — Not enough virtual CPUs were available for the projected required total consumed capacity to be achieved.</li> <li>• <b>"unknown"</b> — Unknown uncategorized failure reason.</li> </ul>

## Description

The **Generate Service Class Resource Adjustments Report** operation generates a report that contains the following information for a specific service class within a workload resource group over the requested time interval:

- A list of resource adjustments that zManager successfully made to maintain specified performance goals. For successful adjustments, the report includes:
  - A list of workload resource groups and the virtual servers that received additional resources (receivers)
  - A list of workload resource groups and the virtual servers that donated the additional resources (donors)
- A list of resource adjustments that zManager attempted but failed to make, along with a reason for the failure.

“Response body contents” on page 539 describes the full list of data that is returned in the report for this workload resource group.

There are multiple types of resource adjustments reports that can be requested, and each has its own API listed in this Performance Management Reports section. This particular generate-service-class-resource-adjustments-report operation generates a resource adjustments report for a specific service class within a specific workload over the requested time period. What this means is that it will contain entries for all adjustments that affected the particular service class submitted; meaning all of the entries returned will involve that specific service class either having received additional resources, or having been forced to give up (donate) some of its resources.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

The request body is validated against the schema described in “Request body contents” on page 539. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Service Class Resource Adjustments Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 539.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-service-class-resource-adjustments-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group",
  "service-class-report-id": "Batch Service"
}
```

---

Figure 277. Generate Service Class Resource Adjustments Report: Request

## Generate Service Class Hops Report

Use the **Generate Service Class Hops Report** operation to create a custom on-demand hops report for a service class within a particular workload resource group over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

Typically, a transactional service class has multiple hops, each hop corresponding to a tier in the transactional flow. Each hop can have one or more application environments associated with it; an application environment includes software and the server or network infrastructure that supports it. An application environment, in turn, might have multiple application environment servers.

**POST** /api/ensembles/{*ensemble-id*}/performance-management/operations/generate-service-class-hops-report

In this request, the URI variable {*ensemble-id*} is the object ID of the ensemble object for which you are requesting a hops report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
workload-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the Workload Resource Groups Report for this same interval.

Field name	Type	Rqd/Opt	Description
service-class-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>service-class-name</b> field. <b>Note:</b> This value is the same as the <b>service-class-report-id</b> value that is returned in one of the service class entries in the Service Classes Report for this same interval.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field.
service-class-name	String	The displayable name of the service class.
service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>service-class-name</b> field.
performance-policy-name	String	The displayable name of the performance policy that contains this service class.
performance-policy-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific performance policy named in the <b>performance-policy-name</b> field.
goal-type	String Enum	The type of performance goal that was defined for this service class. Values are <b>"velocity"</b> or <b>"discretionary"</b> .
goal-performance-level	String Enum	The performance goal of the service class as defined in performance policy. Possible values are: <b>"fastest"</b> , <b>"fast"</b> , <b>"moderate"</b> , <b>"slow"</b> , and <b>"slowest"</b> . This field is returned only if a velocity performance goal was defined for this service class.
business-importance	String Enum	The business importance level that was defined for this service class. Possible values are: <b>"highest"</b> , <b>"high"</b> , <b>"medium"</b> , <b>"low"</b> , and <b>"lowest"</b> . This field is returned only if a velocity performance goal was defined for this service class.
pi-value	Float	The average performance index (PI) value calculated over this reporting interval for the service class returned in the <b>service-class-name</b> field. This field is not returned if the virtual server was not active or if a PI value could not be calculated.



Field name	Type	Description
actual-performance-level	String Enum	The average level of performance that was actually measured over this reporting interval. Possible values are: " <b>fastest</b> ", " <b>fast</b> ", " <b>moderate</b> ", " <b>slow</b> ", and " <b>slowest</b> ". This field is returned only if a velocity performance goal was defined for this service class and if the performance index value could be measured.
equivalent-workload-service-classes	Array of objects	Because a specific virtual server can belong to more than one workload, the hops report might contain reflect combined values for multiple workload resource groups rather than those for the requested workload resource group and service class. In this case, this field contains an array of nested equivalent-workload-service-class objects, each of which identifies another workload resource group or service class, if any, that have virtual servers also contained within the reported hops. The format of the returned object is described in Table 116. Otherwise, this field contains an empty array.
hops	Array of objects	An array of nested hop-entry objects. The format of the returned object is described in Table 117.

Each nested equivalent-workload-service-class object contains the following fields:

*Table 116. Format of an equivalent-workload-service-class object*

Field name	Type	Description
equiv-workload-name	String	The displayable name of a workload resource group that is part of an equivalent service class that was included in this report.
equiv-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>equiv-workload-name</b> field.
equiv-service-class-name	String	The displayable name of the equivalent service class (within the workload resource group defined by the <b>equiv-workload-name</b> ) that was included in this report.
equiv-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>equiv-service-class-name</b> field.

Each nested hop-entry object contains the following fields:

*Table 117. Format of a hop-entry object*

Field name	Type	Description
hop-number	Integer	The number of this hop. The hop numbers reflect the relative order of the flow of a work request (transaction) from one application environment to the next. The first hop is 0, the next is hop 1, the next is hop 2, and so on.
hop-name	String	A text name generated for this hop number.
hop-statistics	Object	A hop-report-statistics object that reflects statistics for this hop as a whole (all application environment groups and their virtual servers that are contained in this hop). The format of the returned object is described in Table 118 on page 546.
hop-application-environments	Array of objects	An array of nested hop-application-env objects. The format of the returned object is described in Table 119 on page 546.

Each nested hops-report-statistics object contains the following fields:

Table 118. Format of a hops-report-statistics object

Field name	Type	Description
successful-transactions	Integer	The total number of transactions that completed successfully.
failed-transactions	Integer	The total number of transactions that failed.
stopped-transactions	Integer	Total number of transactions that stopped before completing. These transactions did not fail or complete successfully; a transaction can enter the stopped state if it encounters an error with the application or server that is processing the transaction. For example, if an application detects that its caller or client terminates the request before the transaction instance completes, the application can stop processing for the transaction instance and report it as stopped, rather than failed or successful.
inflight-transactions	Integer	The total number of transactions that had started, but not yet completed by the end of the requested reporting interval.  Optional: This data may not always be available. For example, when Linux is involved, it is not usually possible to obtain this statistic. If it is not available, this field will not be returned.
queue-time	Integer	The average amount of time (in microseconds) from the time a transaction is received until processing of the transaction begins.  Optional: This data may not always be available. For example, when Linux is involved, it is not usually possible to obtain this statistic. If it is not available, this field will not be returned.
execution-time	Integer	The average amount of time (in microseconds) that transactions took to execute.  Optional: This data may not always be available. For example, when Linux is involved, it is not usually possible to obtain this statistic. If it is not available, this field will not be returned.
successful-avg-response-time	Integer	Average response time (in microseconds) of all successful transactions.
inflight-avg-response-time	Integer	Average amount of time (in microseconds) spent toward response time for inflight transactions.  Optional: This data may not always be available. For example, when Linux is involved, it is not usually possible to obtain this statistic. If it is not available, this field will not be returned.

Each nested hop-application-env object contains the following fields:

Table 119. Format of a hop-application-env object

Field name	Type	Description
appl-env-name	String	The name of the application environment. An application environment is the environment that includes the software and the server or network infrastructure that supports it. As defined by The Open Group application response measurement (ARM) standard, the name is no more than 128 characters in length.
group-name	String	The name of the application environment group with which this application environment is associated. This value cannot be more than 128 characters long but is zero length when the application environment does not belong to a group.

Table 119. Format of a hop-application-env object (continued)

Field name	Type	Description
appl-env-hop-statistics	Object	A hop-statistics object that reflects statistics for this application environment as a whole (all virtual servers that are contained in this application environment). The format of the returned object is described in Table 118 on page 546.
hop-application-env-virtual-servers	Array of objects	An array of nested hop-application-env-virtual-server objects. The format of the returned object is described in Table 120.

Each nested hop-application-env-virtual-server object contains the following fields:

Table 120. Format of a hop-application-env-virtual-server object

Field name	Type	Description
virtual-server-name	String	The displayable name of the virtual server. This field identifies the virtual server within the application environment.
virtual-server-hop-statistics	Object	A hop-report-statistics object that reflects statistics for this specific virtual server within the application environment. The format of the returned object is described in Table 118 on page 546.

## Description

The **Generate Service Class Hops Report** operation generates a report that contains the following information for a specific workload resource group and service class over the requested time interval:

- A list of application-level hops associated with this specific workload resource group and service class for which historical reporting information is available. A hop corresponds to a tier in the transactional flow, and each hop can have one or more application environments associated with it.

A service class can have multiple hops associated with it, and each hop can have one or more associated application environments. These application environments might have multiple virtual servers.

- For each hop associated with the service class:
  - A common set of statistics
  - Statistics for the application environments and virtual servers within each hop.

“Response body contents” on page 544 describes the full list of data that is returned in this report.

For more information about transactional service classes, hops in the transaction flow, and application environments, see the topic “Enhanced monitoring and management through guest platform management providers” in the *z Systems Ensemble Workload Resource Group Management Guide*, GC27-2629.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. The hops report for the service class can be empty when one of the following circumstances is true for all of the virtual servers associated with the service class:

- A guest platform management provider is not installed or running on the operating system.
- No ARM-instrumented applications are running on the operating system.
- None of the ARM-instrumented applications running on the operating system are hop 0 application environments.

An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements” on page 548.

The request body is validated against the schema described in “Request body contents” on page 543. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Hops Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 544.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-service-class-hops-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group",
  "service-class-report-id": "Batch Service"
}
```

---

Figure 278. Generate Service Class Hops Report: Request

## Generate Service Class Virtual Server Topology Report

Use the **Generate Service Class Virtual Server Topology Report** operation to create a custom on-demand topology report for a specific service class within a workload resource group over a requested time period. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

```
POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-service-class-
virtual-server-topology-report
```

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble containing the workload resource group for which you want to receive a virtual server topology report.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
workload-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the Workload Resource Groups Report for this same interval.
service-class-report-id	String	Required	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>service-class-name</b> field. <b>Note:</b> This value is the same as the <b>service-class-report-id</b> value that is returned in one of the service class entries in the Service Classes Report for this same interval.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>workload-name</b> field.
service-class-name	String	The displayable name of the service class.
service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>service-class-name</b> field.

Field name	Type	Description
equivalent-workload-service-classes	Array of objects	Because a specific virtual server can belong to more than one workload, the hops report might contain reflect combined values for multiple workload resource groups rather than those for the requested workload resource group and service class. In this case, this field contains an array of nested equivalent-workload-service-class objects, each of which identifies another workload resource group or service class, if any, that have virtual servers also contained within the reported hops. The format of the returned object is described in Table 121. Otherwise, this field contains an empty array.
topo-hop-count	Integer	The total number of hops represented in this topology report.
topo-hops	Array of objects	An array of nested topo-hop objects. The format of the returned object is described in Table 122.

Each nested equivalent-workload-service-class object contains the following fields:

*Table 121. Format of an equivalent-workload-service-class object*

Field name	Type	Description
equiv-workload-name	String	The displayable name of a workload resource group that is part of an equivalent service class that was included in this report.
equiv-workload-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific workload resource group named in the <b>equiv-workload-name</b> field.
equiv-service-class-name	String	The displayable name of the equivalent service class (within the workload resource group defined by the <b>equiv-workload-name</b> ) that was included in this report.
equiv-service-class-report-id	String	An identifier used by the performance management reporting structure to keep track of reporting data for the specific service class named in the <b>equiv-service-class-name</b> field.

Each nested topo-hop object contains the following fields:

*Table 122. Format of a topo-hop object*

Field name	Type	Description
hop-number	Integer	The number of this hop. The hop numbers reflect the relative order of the flow of a work request (transaction) from one application environment to the next. The first hop is 0, the next is hop 1, the next is hop 2, and so on.
topo-virtual-server-nodes	Array of objects	An array of nested topo-virtual-server-node objects that represent virtual servers that are part of this hop. The format of the returned object is described in Table 123.

Each nested topo-virtual-server-node object contains the following fields:

*Table 123. Format of a topo-virtual-server-node object*

Field name	Type	Description
node-identifier	String	A topology node identifier that uniquely identifies this specific node within this entire topology report. For example, a virtual server that shows up in multiple hops of the transaction flow has the same virtual server name in all of those nodes but each node identifier is different.
virtual-server-name	String	The displayable name of the virtual server.

Table 123. Format of a topo-virtual-server-node object (continued)

Field name	Type	Description
hypervisor-type	String	The value of the <b>type</b> property of the virtualization host. See the virtualization host "Data model" on page 211 for more details about the valid values of this property.
hypervisor-name	String	The displayable name of the hypervisor under which this virtual server was running.
was-active	Boolean	Indicates whether this virtual server was active during this reporting interval. The value is <b>"true"</b> if the virtual server was active or <b>"false"</b> if it was not active. If the value of this field is false, none of the remaining fields in this table have data, so none of these fields are returned.
physical-cpu-utilization-percent	Float	The physical processor utilization percentage (%) of the virtual server. This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active during this interval.
hypervisor-cpu-delay-percent	Float	The hypervisor processing unit delay in fractional form (between 0 and 1 inclusive). This field is returned for the following hypervisors only: z/VM and PowerVM. In the case of z/VM, a value is available for this field only if sampling is turned on for the guest. <b>Optional:</b> This field is returned only if the virtual server was active during this interval and if the hypervisor type supports this information.
idle-time-percent	Float	The idle time percentage (%) of the virtual server. It is the percentage of total time that the virtual server had no work (that is, no application processes or internal hypervisor specific process states). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active during this interval and if the hypervisor type supports this information. For z/VM, this data is available only when sampling is enabled and started. For PowerVM and PR/SM, idle time data is not available.
other-time-percent	Float	The percentage (%) of total time that miscellaneous hypervisor specific internal process states had control for this virtual server. In other words, the percentage of time that the virtual server was not idle but also was not in a state of active CPU utilization or hypervisor CPU delay. This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active during this interval and if the hypervisor type supports this information. For z/VM, this data is available only when sampling is enabled and started. For PowerVM and PR/SM, other time data is not available.
os-cpu-using-samples-percent	Float	The percentage of CPU using samples from among the total samples. For example, if there are 10 CPU using samples out of a total of 10 samples, then CPU using samples is 100% (because out of the total samples, all are CPU using samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
os-cpu-delay-samples-percent	Float	The percentage of CPU delay samples from among the total samples. For example, if there are 10 CPU delay samples and 10 samples that are not CPU delay samples, then CPU delay samples is 50% (because out of the total samples half are CPU delay samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.

Table 123. Format of a topo-virtual-server-node object (continued)

Field name	Type	Description
os-io-delay-samples-percent	Float	The percentage of I/O delay samples from among the total samples. The percent I/O delay is the percent of samples taken when work was delayed for non-paging DASD I/O. The I/O delay includes IOS queue, subchannel pending, and control unit queue delays. For example, if there are 10 I/O delay samples and 10 samples that are not I/O delay samples, then I/O delay samples is 50% (because out of the total samples half are I/O delay samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
os-page-delay-samples-percent	Float	The percentage of page delay samples from among the total samples. The percent page delay is the percent of samples when the address space experienced page faults in cross-memory access, and the page faults were resolved from auxiliary storage. For example, if there are 10 page delay samples and 10 samples that are not page delay samples, then page delay samples is 50% (because out of the total samples half are page delay samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
appl-env-vs-response-data	Array of objects	An array of nested appl-env-vs-response-entry objects that contain response time data for each of the application environments with which the virtual server is associated. If no response data is available, this field contains an empty array. The format of the returned object is described in Table 124.
appl-env-vs-utilization-data	Array of objects	An array of nested appl-env-vs-utilization-data objects that contain utilization data for each of the application environments with which the virtual server is associated. If no response data is available, this field contains an empty array. The format of the returned object is described in Table 125 on page 553.
child-virtual-server-node-links	Array of objects	An array of nested child-virtual-server-node-link objects that contain a list of links to other virtual server nodes in the topology, if this particular virtual server sent transactions to one or more virtual server nodes in the next hop. These link objects provide the information required to identify those next-level nodes (or "child" virtual servers). If this virtual server node has no child nodes, this field contains an empty array. The format of the returned object is described in Table 126 on page 554.

Each nested appl-env-vs-response-entry object contains the following fields:

Table 124. Format of an appl-env-vs-response-entry object

Field name	Type	Description
appl-env-name	String	The name of the application environment. An application environment is the environment that includes the software and the server or network infrastructure that supports it.
group-name	String	The name of the application environment group with which this application environment is associated. This value cannot be more than 128 characters long but is zero length when the application environment does not belong to a group.
successful-transactions	Integer	The total number of transactions that completed successfully.
failed-transactions	Integer	The total number of transactions that failed.



Table 124. Format of an *appl-env-vs-response-entry* object (continued)

Field name	Type	Description
stopped-transactions	Integer	Total number of transactions that stopped before completing. These transactions did not fail or complete successfully; a transaction can enter the stopped state if it encounters an error with the application or server that is processing the transaction. For example, if an application detects that its caller or client terminates the request before the transaction instance completes, the application can stop processing for the transaction instance and report it as stopped, rather than failed or successful.
inflight-transactions	Integer	The total number of transactions that had started, but not yet completed by the end of the requested reporting interval.  Optional: This data may not always be available. For example, when Linux is involved, it is not usually possible to obtain this statistic. If it is not available, this field will not be returned.
queue-time	Integer	The average amount of time (in microseconds) from the time a transaction is received until processing of the transaction begins.  Optional: This data may not always be available. For example, when Linux is involved, it is not usually possible to obtain this statistic. If it is not available, this field will not be returned.
execution-time	Integer	The average amount of time (in microseconds) that transactions took to execute.  Optional: This data may not always be available. For example, when Linux is involved, it is not usually possible to obtain this statistic. If it is not available, this field will not be returned.
successful-avg-response-time	Integer	Average response time (in microseconds) of all successful transactions.
inflight-avg-response-time	Integer	Average amount of time (in microseconds) spent toward response time for inflight transactions.  Optional: This data may not always be available. For example, when Linux is involved, it is not usually possible to obtain this statistic. If it is not available, this field will not be returned.

Each nested *appl-env-vs-utilization-entry* object contains the following fields:

Table 125. Format of an *appl-env-vs-utilization-entry* object

Field name	Type	Description
appl-env-name	String	The name of the application environment. An application environment is the environment that includes the software and the server or network infrastructure that supports it.
group-name	String	The name of the application environment group with which this application environment is associated. This value cannot be more than 128 characters long but is zero length when the application environment does not belong to a group.
cpu-time	Integer	The processor (CPU) time, in microseconds, that was consumed on this virtual server as part of this application environment.

Table 125. Format of an appl-env-vs-utilization-entry object (continued)

Field name	Type	Description
os-cpu-using-samples-percent	Float	The percentage of CPU using samples from among the total samples. For example, if there are 10 CPU using samples out of a total of 10 samples, then CPU using samples is 100% (because out of the total samples, all are CPU using samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
os-cpu-delay-samples-percent	Float	The percentage of CPU delay samples from among the total samples. For example, if there are 10 CPU delay samples and 10 samples that are not CPU delay samples, then CPU delay samples is 50% (because out of the total samples half are CPU delay samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
os-io-delay-samples-percent	Float	The percentage of I/O delay samples from among the total samples. The percent I/O delay is the percent of samples taken when work was delayed for non-paging DASD I/O. The I/O delay includes IOS queue, subchannel pending, and control unit queue delays. For example, if there are 10 I/O delay samples and 10 samples that are not I/O delay samples, then I/O delay samples is 50% (because out of the total samples half are I/O delay samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.
os-page-delay-samples-percent	Float	The percentage of page delay samples from among the total samples. The percent page delay is the percent of samples when the address space experienced page faults in cross-memory access, and the page faults were resolved from auxiliary storage. For example, if there are 10 page delay samples and 10 samples that are not page delay samples, then page delay samples is 50% (because out of the total samples half are page delay samples). This percentage is in fractional form (between 0 and 1 inclusive). <b>Optional:</b> This field is returned only if the virtual server was active and a guest platform management provider was running on the virtual server during this interval.

Each nested child-virtual-server-node-link object contains the following fields:

Table 126. Format of a child-virtual-server-node-link object

Field name	Type	Description
child-node-identifier	String	The unique identifier assigned within this topology report to this child virtual server node.
child-hop-number	Integer	The number of the hop that contains this child virtual server node.
transaction-count	Integer	The number of transactions that flowed from the parent virtual server node to this specific child virtual server node over this reporting period. If no transaction data is available, this field is not returned.

## Description

The **Generate Service Class Virtual Server Topology Report** operation generates a report that contains the following information for a specific service class within a workload resource group over the requested time interval:

- A list of application-level hops associated with this specific workload resource group and service class for which historical reporting information is available. A hop corresponds to a tier in the transactional flow, and each hop can have one or more application environments associated with it.

A service class can have multiple hops associated with it, and each hop can have one or more associated application environments. These application environments might have multiple virtual servers.

- For each hop associated with the service class:
  - A common set of statistics
  - Statistics for the application environments and virtual servers within each hop.

This report returns information that is similar to the output of the **Generate Service Class Hops Report** operation; the virtual server nodes are organized by hop and the reports contain some similar statistics for each virtual server. This topology report, however, also returns information about the relationship between the virtual servers from one hop to the next. This type of information relationships is not always available, depending on the actual configuration of the application environments.

For an example of a graphical representation of the virtual server topology, use the **Virtual Server Topology Report** that is available through the HMC or see the topic about that report in the *z Systems Ensemble Workload Resource Group Management Guide*, GC27-2629. The **Virtual Server Topology Report** task graphically depicts the relationships between virtual servers that are running the workload resource group and providing the resources to complete the work. The framed image in that report displays graphical representations of virtual servers. Each is represented by an icon that displays the name and status of the virtual server; by clicking the icon, you can view statistics for that virtual server node. The **Generate Service Class Virtual Server Topology Report** operation generates the same information that you can access through the **Virtual Server Topology Report** and the icons in its display.

“Response body contents” on page 549 describes the full list of data that is returned in the report for this workload resource group.

If reporting data is not available for the requested time interval, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

The request body is validated against the schema described in “Request body contents” on page 549. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Virtual Server Topology Report** and the **View Statistics** tasks.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 549.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-service-class-virtual-server-topology-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group",
  "service-class-report-id": "Batch Service"
}
```

---

Figure 279. Generate Service Class Virtual Server Topology Report: Request

## Generate Load Balancing Report

Use the **Generate Load Balancing Report** operation to list information about network load balancing activity in an ensemble. Unlike other performance management reports, which are based on historical performance management data that was retained over a specific time period, this report contains data that is available when the **Generate Load Balancing Report** API is called.

### HTTP method and URI

**POST /api/ensembles/{ensemble-id}/performance-management/operations/generate-load-balancing-report**

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble for which you want to receive a load balancing report.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
load-balancing-groups	Array of objects	An array of nested load-balancing-Group objects. If no load balancing groups have been defined for this ensemble, an empty array is returned for this field.

Each nested load-balancing-Group object contains the following fields:

Field name	Type	Description
load-balancer-group-name	String	The name given to this load balancer group. The returned string in this field conforms to Server/Application State Protocol (SASP) standards, which allow a group name to consist of 1 through 255 UTF8 characters.
load-balancer-group-members	Array of objects	An array of nested load-balancer-group-member objects. If a group does not have any members associated with it, an empty array is returned for this field.

Each nested load-balancer-group-member object contains the following fields:

Field name	Type	Description
virtual-server-name	String	The name of a virtual server defined to this load balancing group.
hostname-ipaddr	String	The IP address of the virtual server.
weight	Float	The relative weight recommendation that zManager assigned to this member.
status	String Enum	The status of or environmental conditions for each virtual server in the selected load balancing group. Possible values are: <ul style="list-style-type: none"> <li>• "active"</li> <li>• "down"</li> <li>• "inactive"</li> <li>• "failure-recovery"</li> <li>• "quiesced"</li> <li>• "unknown"</li> </ul>
status-detailed	String Enum	A further explanation for the specific status condition "down" or "inactive"; for other status field values, this field is not returned. Possible values are: <ul style="list-style-type: none"> <li>• "matching-ip-address-not-found"</li> <li>• "gpmp-downlevel"</li> <li>• "lsof-not-available"</li> <li>• "no-process-listening"</li> <li>• "not-initialized"</li> <li>• "vs-not-operating"</li> <li>• "load-balance-data-initializing"</li> <li>• "unsupported-hypervisor-config"</li> <li>• "hypervisor-details-not-available"</li> <li>• "data-retrieval-error"</li> </ul>
port-number	Integer	The port number on the virtual server that is being used for load balancing purposes. This field is not returned if the port number is not available.
protocol-type	String Enum	The protocol in use for this virtual server; possible values are: <ul style="list-style-type: none"> <li>• "tcp" for Transmission Control Protocol, or</li> <li>• "udp" for User Datagram Protocol</li> </ul> This field is not returned if the protocol type is not available.

## Description

The **Generate Load Balancing Report** operation generates a report that contains a list of all of the load-balancing groups that zManager is monitoring within a specific ensemble. For each load-balancing group, the report also contains information about each individual virtual server that is a member of the group. "Response body contents" on page 556 describes the full list of data that is returned in this report for each load-balancing group and its members.

The generated load-balancing report contains data that is available when the **Generate Load Balancing Report** API is called. Load-balancing data for an ensemble is refreshed every 30 seconds. For the most efficient use of the **Generate Load Balancing Report** API, space additional API calls at least 30 seconds apart.

If no load-balancing groups are defined for this ensemble, an empty response object is provided and the operation completes successfully. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Load Balancing Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 556.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/operations/
generate-load-balancing-report
```

---

Figure 280. Generate Load Balancing Report: Request

## Workload Element Group object

A Workload Element Group object is a container object of redundant elements of one or more workloads. The group is used to maintain redundancy for a function within the workload business unit. It must always be in one or more user-defined workloads.

## Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, with the type-specific specialization described in the following tables. Note that this object does not maintain the concept of an operational status, and therefore does not provide the operational-status-related properties.

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

Table 127. Workload object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/URI	The canonical URI path for a Workload Element Group object is of the form <code>/api/workload-element-groups/{group-id}</code> where <code>{group-id}</code> is the value of the <b>object-id</b> property of the Workload Element Group object.
<b>object-id</b>	—	String (36)	The unique identifier of the Workload Element Group object, in the form of a UUID.
<b>parent</b>	—	String	The <b>object-uri</b> property of the owning ensemble object.
<b>class</b>	—	String	The class of a Workload Element Group object is <b>"workload-element-group"</b> .
<b>name</b>	(w) (pc)	String (1-64)	The user-specified workload element group name, which can be up to 64 characters made up of alphanumeric characters, blanks, periods, underscores, or dashes. Names must start with an alphabetic character and end with an alphabetic character, numeric character, underscore, period, or dash. Names must be unique to other existing workload element groups in the ensemble.
<b>description</b>	(w) (pc)	String (0-256)	The user-specified workload element group description, up to 256 characters, or if none was provided, an empty string.
<b>is-locked</b>	—	Boolean	This property is always false for a Workload Element Group object.

In addition to the properties defined through included schemas, this object includes the additional class-specific properties:

Table 128. Workload Element Group object: type-specific properties

Name	Qualifier	Type	Description
<b>category</b>	(w) (pc)	String (0-32)	Up to 32 characters used to categorize the workload object, often with other workloads within the ensemble.
<b>workloads</b>	(c) (pc)	Array of String/URI	An array of URIs of the workloads in which the workload element group belongs.  Value may not be null or an empty list.
<b>minimum-available-elements</b>	(w) (pc)	Integer	The number of workload elements that should minimally be available in the Workload Element Group object, by default, unless overridden by an Availability Policy object.  Value must be greater than or equal to zero.
<b>preferred-available-elements</b>	(w) (pc)	Integer	The number of workload elements that are preferred to be available in the Workload Element Group object, by default, unless overridden by an Availability Policy object.  Value must be greater than or equal to <b>minimum-available-elements</b> .

Table 128. Workload Element Group object: type-specific properties (continued)

Name	Qualifier	Type	Description
<b>virtual-servers</b>	(c) (pc)	Array of String/URI	This is a placeholder property for which container-type events will be fired to indicate the delta of the URIs of virtual servers of the workload element group which were added or removed. The actual list of virtual servers can be retrieved through the <b>List Virtual Servers of a Workload Element Group</b> operation.  Note that only virtual servers that support Ensemble Availability Management may be added to the list. Virtual servers which do not support Ensemble Availability Management will have an <b>avail-status</b> of “ <b>not-supported</b> ”.
<b>avail-statuses</b>	(pc)	Array of objects	The availability statuses of the workload element group for each workload it belongs to, each in the form of a workload element group availability status object, as described below. The availability status of a workload element group can vary by workload context due to the workload element group configuration overrides a workload’s availability policy can define.

### Workload element group availability status nested object

The workload element group availability status object is nested within a workload element group to provide the availability status for the workload element group within the context of each workload to which it belongs.

The following properties are supported:

Table 129. Workload Element Group object: Workload element group availability status nested object properties

Field name	Type	Rqd/Opt	Description
<b>workload-uri</b>	String/URI		The canonical URI path for a workload object that the workload element group availability status is in the context of, in the form <code>/api/workload-resource-groups/{workload-id}</code> where <code>{workload-id}</code> is the value of the <b>object-id</b> property of the workload object.
<b>avail-status-with-reasons</b>	Array of objects		The list of availability statuses of the workload element group, each with reasons explaining the status, in order of severity. The list will always contain at least one workload element in the form of a workload element group availability status with reasons nested object, as described in the next table.

**Workload element group availability status with reasons nested object:** The workload element group availability status with reasons nested object is nested within a workload element group availability status object for a particular workload to encapsulate the availability status with the reasons explaining why it is in that state.



Table 130. Workload Element Group object: Workload element group availability status with reasons nested object properties

Field name	Type	Rqd/Opt	Description
<b>avail-status</b>	String Enum		<p>The availability status of the workload element group, determined by the active availability policy in the workload and the availability status of the workload element group's virtual servers. The possible values are as follows:</p> <ul style="list-style-type: none"> <li>• <b>"available"</b>: The workload element group is considered available.</li> <li>• <b>"exposed"</b>: The workload element group is considered exposed.</li> <li>• <b>"critical"</b>: The workload element group is considered critically exposed.</li> <li>• <b>"not-available"</b>: The workload element group is considered not available.</li> </ul>
<b>reasons</b>	Array of String Enum		<p>The list of reasons given for the availability status of the workload element group. The possible values are as follows:</p> <ul style="list-style-type: none"> <li>• <b>"min-avail"</b>: The number of available workload elements in the workload element group is below the configured minimum.</li> <li>• <b>"pref-avail"</b>: The number of available workload elements in the workload element group is below the configured preference.</li> <li>• <b>"zero-avail"</b>: The number of available workload elements in the workload element group is zero and the <b>"min-avail"</b> is &gt; 0.</li> </ul> <p>The list may be empty if no reasons are given.</p>

## List Workload Element Groups of an Ensemble

Use the **List Workload Element Groups of an Ensemble** operation to list the workload element groups within the target ensemble.

### HTTP method and URI

GET /api/ensembles/{ensemble-id}/workload-element-groups

In this request, the URI variable {ensemble-id} is the object ID of the ensemble for which you are requesting a list of workload element groups.

### Query parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern to limit returned objects to those that have a matching <b>name</b> property. If matches are found, the response is an array with all workload element groups that match. If no match is found, the response is an empty array.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
workload-element-groups	Array of objects	Array of nested element-group-info objects as described in the next table.

Each nested element-group-info object contains the following fields:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the Workload Element Group object, in the form <code>/api/workload-element-groups/{group-id}</code> .
name	String	Display name of the Workload Element Group object.

## Description

The **List Workload Element Groups of an Ensemble** operation lists the workload element groups that belong to the ensemble targeted by the request URI. The object URI and display name are provided for each.

If the **name** query parameter is specified, the returned list is limited to the workload element group in the ensemble that has a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

A workload element group is included in the list only if you have object-access permission for that object. Additionally, you must have access to the ensemble object, or an error response will be returned.

If the HMC does not manage the targeted ensemble, then an error response occurs.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the ensemble object passed in the request URI and all Workload Element Group objects to be included in the result.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in the “Response body contents” on page 561.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{ensemble-id}</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Get Workload Element Group Properties

Use the **Get Workload Element Group Properties** operation to retrieve the properties of a single Workload Element Group object that is designated by its object ID.

### HTTP method and URI

**GET** `/api/workload-element-groups/{group-id}`

### URI variables

Variable	Description
<i>{group-id}</i>	Object ID of the Workload Element Group object for which properties are to be obtained.

## Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Workload Element Group object as defined in “Data model” on page 558. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

## Description

The **Get Workload Element Group Properties** operation returns the current properties for the Workload Element Group object specified by *{group-id}*.

On successful execution, all of the current properties as defined by the “Data model” on page 558 are provided in the response body and HTTP status code 200 (OK) is returned.

A URI path must designate an existing Workload Element Group object and you must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned.

## Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the Workload Element Group object designated by *{group-id}*.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{group-id}</i> ) does not designate an existing Workload Element Group object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Create Workload Element Group

Use the **Create Workload Element Group** operation to create a new custom workload element group in an ensemble. The new workload element group must be uniquely named within the ensemble.

## HTTP method and URI

**POST** `/api/ensembles/{ensemble-id}/workload-element-groups`

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble object for which a workload element group is to be created.

## Request body contents

The request body contains properties used to define the new workload element group, which are a subset of the properties of a workload element group. Some properties are optional:

Field name	Type	Rqd/Opt	Description
name	String (1-64)	Required	The value to be set as the workload element group's <b>name</b> property.
description	String (0-256)	Optional	The value to be set as the workload element group's <b>description</b> property.
category	String (0-32)	Optional	The value to be set as the workload element group's <b>category</b> property.
workloads	Array of String/URI	Required	The value to be set as the workload element group's <b>workloads</b> property.
minimum-available-elements	Integer	Optional	The value to be set as the workload element group's <b>minimum-available-elements</b> property. Default value: 2
preferred-available-elements	Integer	Optional	The value to be set as the workload element group's <b>preferred-available-elements</b> property. Default value: <b>minimum-available-elements</b>
virtual-servers	Array of String/URI	Optional	The value to be set as the workload element group's <b>virtual-servers</b> property. Default value: an empty list

## Response body contents

On successful completion, the response body is a JSON object with the following field:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the Workload Element Group object, in the form <code>/api/workload-element-groups/{group-id}</code> .

## Description

The **Create Workload Element Group** operation creates a new custom workload element group in the ensemble identified by *{ensemble-id}*.

The request body contains an object with one or more fields with field names that correspond to the names of properties for this object. The only required properties are the workload element group **name**, which must be unique among existing workload element groups in the ensemble, and the **workloads** in which the workload element group belongs. On successful execution, the workload element group is created, added to the identified workloads, and added to the ensemble and status code 201 is returned with a response body containing a reference to the new workload element group.

The URI path must designate an existing ensemble object and the API user must have object-access permission to it. The **workloads** property must designate a list of existing workload objects and the API user must have object-access permission to them. If provided, the **virtual-servers** property must designate a list of existing Virtual Server objects and the API user must have object-access permission to them. If

any of these conditions is not met, status code 404 (Not Found) is returned. If any of the passed virtual servers do not support Ensemble Availability Management, status code 400 (Bad Request) is returned. In addition, you must have permission to the **New Element Group** task and the ensemble must be at the AUTOMATE entitlement level, otherwise status code 403 (Forbidden) is returned.

The request body is validated against the “Data model” on page 558 to ensure that it contains only writeable properties and the data types of those properties are as required. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Object access permission to all Workload objects designated by **workloads**.
- Object access permission to all Virtual Server objects designated by **virtual-servers**.
- Action/task permission to the **New Element Group** task.

## HTTP status and reason codes

On successful completion, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents” on page 564.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	65	The name of the new workload element group is not unique.
	66	One or more of the virtual servers in the request body do not support Ensemble Availability Management.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
	3	The ensemble is not operating at the management entitlement level required to perform this operation (Automate).
404 (Not Found)	1	The object ID in the URI ( <i>ensemble-id</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to the object.
	2	A URI in <b>workloads</b> or <b>virtual-servers</b> does not designate an existing workload or Virtual Server object, or the API user does not have object access permission to the object.
409 (Conflict)	61	The operation cannot be performed because one or more virtual servers designated by the <b>virtual-servers</b> URI are currently busy performing some other operation.
	64	The operation cannot be performed because one or more workload resource groups servers designated by the <b>workloads</b> URI are currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Delete Workload Element Group

Use the **Delete Workload Element Group** operation to remove a workload element group from an ensemble and the workloads to which it belongs.

### HTTP method and URI

**DELETE** /api/workload-element-groups/{group-id}

In this request, the URI variable {group-id} is the object ID of the Workload Element Group object that you want to remove.

### Description

The **Delete Workload Element Group** operation deletes the Workload Element Group object specified by {group-id} from its ensemble.

On successful completion, the workload element group is removed from all workloads (potentially removing its virtual servers from the workloads), the workload element group is removed from the ensemble, and the status code 204 (No Content) is returned without supplying a response body. See the *z Systems Ensemble Workload Resource Group Management Guide*, GC27-2629, for details about managing workload element groups.

The URI path must designate an existing Workload Element Group object and the API user must have object-access permission to it. If this condition is not met, status code 404 (Not Found) is returned. Additionally, you must have permission to the **Delete Element Group** task. If this condition is not met, status code 403 (Forbidden) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the ensemble object that owns the workload element group.
- Object-access permission to the Workload Element Group object designated by {group-id}.
- Action/task permission to the **Delete Element Group** task.

### HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	62	The workload element group to be removed contains virtual servers to which the API user does not have access.
	68	The workload element group to be removed is assigned to one or more workload resource groups servers to which the API user does not have access.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ({group-id}) does not designate an existing Workload Element Group object, or the API user does not have object access permission to the object.

HTTP error status code	Reason code	Description
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	61	The operation cannot be performed because one or more virtual servers in the workload element group designated by the request URI are currently busy performing some other operation.
	64	The operation cannot be performed because one or more workload resource groups to which the workload element group is assigned are currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Update Workload Element Group

Use the **Update Workload Element Group** operation to modify one or more writeable properties of a Workload Element Group object.

### HTTP method and URI

**POST** /api/workload-element-groups/{group-id}

In this request, the URI variable *{group-id}* is the object ID of the Workload Element Group object that you are modifying.

### Request body contents

The request body is a JSON object with one or more of the writeable fields for a workload element group as described in the “Data model” on page 558. You are required to supply only those fields that you want to change.

### Description

The **Update Workload Element Group** updates one or more simple writeable properties of the Workload Element Group object identified by *{group-id}*.

On successful execution, the Workload Element Group object is updated with the supplied property values and status code 204 (No Content) is returned without supplying a response body. Notifications for these property changes are sent asynchronously to this operation.

The URI path must designate an existing Workload Element Group object and you must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, you must also have permission to the **Element Group Details** task as well, otherwise status code 403 (Forbidden) is returned.

The request body is validated against the schema described in “Request body contents.” If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

### Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the Workload Element Group object designated by *{group-id}*.
- Action/task permission to the **Element Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned. Otherwise, the following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	65	The new name given to the workload element group is not unique.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{group-id}</i> ) does not designate an existing Workload Element Group object, or the API user does not have object access permission to the object.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## List Virtual Servers of a Workload Element Group

The **List Virtual Servers of a Workload Element Group** operation lists the virtual servers in the passed workload element group. This is the way to get the information corresponding to the **virtual-servers** property of a workload element group, as it is omitted from the standard GET operation.

### HTTP method and URI

**GET** /api/workload-element-groups/{group-id}/virtual-servers

In this request, the URI variable *{group-id}* is the object ID of the Workload Element Group object for which you are requesting a list of virtual servers.

#### Query parameters:

Field name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern to limit returned objects to those that have a matching <b>name</b> property. If matches are found, the response will be an array with information for all virtual servers that match. If no match is found, the response will be an empty array.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
virtual-servers	Array of objects	Array of nested virtual-server-info objects as described in the next table.

Each nested virtual-server-info object contains the following fields:



Field name	Type	Description
object-uri	String/URI	Canonical URI path of the Virtual Server object.
name	String	The current value of the <b>name</b> property of the Virtual Server object.
status	String Enum	The current value of the <b>status</b> property of the Virtual Server object.
type	String Enum	The value of the <b>type</b> property of the Virtual Server object.

## Description

The **List Virtual Servers of a Workload Element Group** operation lists the virtual servers that belong to the workload element group targeted by the request URI. The object URI, display name, status, and type are provided for each.

A virtual server is included in the list only if you have object-access permission for that object. Additionally, you must have access to the Workload Element Group object, or an error response will be returned.

If the HMC's Ensemble does not contain the targeted workload element group, then an error response occurs.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the Workload Element Group object passed in the request URI and all Virtual Server objects to be included in the result.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents" on page 568.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{group-id}</i> ) does not designate an existing workload object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Add Virtual Server to a Workload Element Group

Use the **Add Virtual Server to a Workload Element Group** operation to add a virtual server to the target workload element group.

### HTTP method and URI

**POST** /api/workload-element-groups/{group-id}/operations/add-virtual-server

In this request, the URI variable *{group-id}* is the object ID of the Workload Element Group object to which you want to add a virtual server.

## Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
virtual-server	String/URI	The canonical URI identifying the Virtual Server object to be added to the target workload element group.

## Description

The **Add Virtual Server to a Workload Element Group** operation adds a virtual server directly to the workload element group targeted by the request URI. See the *z Systems Ensemble Workload Resource Group Management Guide, GC27-2629*, for details on managing virtual servers in workload element groups.

If the virtual server is already defined to the workload element group, then the request is ignored with a successful return code. If the virtual server is new to the workload element group, then it is added and a property notification is sent asynchronously to the request. Note that adding a virtual server may impact the availability status of the workload element group and its workload(s), depending on the workload element objectives of the workload element group and availability status of the virtual server.

An error response will be returned if the targeted workload element group or virtual server does not exist or you do not have object-access to either of them. If the virtual server being added does not support Ensemble Availability Management, status code 400 (Bad Request) will be returned. Additionally, you must have access to the **Element Group Details** task, or an error response will be returned.

The request body is validated against the schema described in the “Request body contents.” If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Workload Element Group object passed in the request URI.
- Object access permission to the Virtual Server object passed in the request body.
- Action/task permission to the **Element Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	66	The virtual server in the request body does not support Ensemble Availability Management.
403 (Forbidden)	1	The API user does not have the required permission for this operation.

HTTP error status code	Reason code	Description
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{group-id}</i> ) does not designate an existing Workload Element Group object, or the API user does not have object access permission to it.
	2	The <b>object-id</b> in the URI of the Virtual Server object in the request body does not designate an existing object, or the API user does not have access permission to it.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	61	The operation cannot be performed because the virtual server to be added is currently busy performing some other operation.
	64	The operation cannot be performed because one or more workload resource groups to which the workload element group is assigned are currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Remove Virtual Server from a Workload Element Group

Use the **Remove Virtual Server from a Workload Element Group** operation to remove a virtual server from the target workload element group.

### HTTP method and URI

**POST** /api/workload-element-groups/{*group-id*}/operations/remove-virtual-server

In this request, the URI variable *{group-id}* is the object ID of the Workload Element Group object from which you are removing the virtual server.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Description
virtual-server	String/URI	The canonical URI identifying the Virtual Server object to be removed from the targeted workload element group.

### Description

The **Remove Virtual Server from a Workload Element Group** operation removes a virtual server from being defined to the workload element group targeted by the request URI. See the *z Systems Ensemble Workload Resource Group Management Guide, GC27-2629*, for details on managing virtual servers in workload element groups.

If the virtual server is not defined to the workload element group, then the request is ignored with a successful return code. If the virtual server is defined to the workload element group, then it is removed and a property change notification will be sent asynchronously to this request. If the virtual server is no longer contained by a user-defined workload for any reason, then it will be placed in the default workload. Note that removing a virtual server may impact the availability status of the workload element group and its workload(s), depending on the redundancy objectives of the workload element group and availability status of the virtual server.

An error response is returned if the targeted workload element group or virtual server does not exist or if you do not have object-access to either of them. Additionally, you must have access to the **Element Group Details** task, or an error response will be returned.

The request body is validated against the schema described in “Request body contents” on page 571. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Workload Element Group object passed in the request URI.
- Object access permission to the Virtual Server object passed in the request body.
- Action/task permission to the **Element Group Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{group-id}</i> ) does not designate an existing Workload Element Group object, or the API user does not have object access permission to it.
	2	The <b>object-id</b> in the URI of the Virtual Server object in the request body does not designate an existing object, or the API user does not have access permission to it.
409 (Conflict)	2	The operation cannot be performed because the virtual server to be removed is currently busy performing some other operation.
	61	The operation cannot be performed because the group to be removed contains one or more virtual servers that are currently busy performing some other operation.
	64	The operation cannot be performed because one or more workload resource groups to which the workload element group is assigned are currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

---

## Availability Policy object

Performance Policy objects are elements of a workload which define availability objectives for its workloads elements and how the availability of those elements impacts the availability of the workload. An availability policy consists mainly of a set importance level, availability status impact statements, and workload elements group configuration override statements.

Any number of custom Availability Policy objects can be defined for a workload. However, exactly one will be active at a time. Every workload contains an immutable default availability policy which will be active if a custom availability policy is not activated.

## Data model

This element includes the following type-specific properties. For definitions of the qualifier abbreviations in the following table, see “Property characteristics” on page 38.

Table 131. Availability Policy object: type-specific properties

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/URI	The canonical URI path for an availability policy is qualified by its workload object and is of the form <code>/api/workloads/{workload-id}/availability-policies/{policy-id}</code> where <code>{workload-id}</code> is the value of the <b>object-id</b> property of the parent workload object and <code>{policy-id}</code> is the value of the <b>element-id</b> property of the policy.
<b>element-id</b>	—	String (36)	The unique identifier for the availability policy instance. The policy's unique identifier of the policy is in the form of a UUID.
<b>parent</b>	—	String/URI	The parent of an availability policy is its owning workload, and so the <b>parent</b> value is the canonical URI path for the workload resource group.
<b>class</b>	—	String	The class of an Availability Policy object is <b>"availability-policy"</b> .
<b>name</b>	(w) (pc)	String (1-64)	The display name specified for the availability policy, which can be up to 64 characters made up of alphanumeric characters, blanks, periods, underscores, or dashes. Names must start with an alphabetic character and end with an alphabetic character, numeric character, underscore, period, or dash. Names must be unique to other existing availability policies in the workload.
<b>description</b>	(w) (pc)	String (0-256)	Arbitrary text describing the availability policy in up to 256 characters.
<b>is-default</b>	—	Boolean	This value is used to identify the default Availability Policy object. It is true for the default policy and false for all other (custom) policies in the workload.
<b>importance</b>	(w) (pc)	String Enum	The importance value assigned to the availability policy, which is one of the following: <ul style="list-style-type: none"> <li>• <b>"highest"</b></li> <li>• <b>"high"</b></li> <li>• <b>"medium"</b></li> <li>• <b>"low"</b></li> <li>• <b>"lowest"</b></li> </ul> <p>See the <i>z Systems Ensemble Workload Resource Group Management Guide</i>, GC27-2629, for detailed information on the <b>importance</b> property values.</p>

Table 131. Availability Policy object: type-specific properties (continued)

Name	Qualifier	Type	Description
<b>vs-exclusions</b>	(w) (pc)	Array of String/URI	<p>The list of canonical URI paths of Virtual Server objects for which availability status will not impact the workload availability status.</p> <p>Note that only virtual servers that support Ensemble Availability Management may be added to the list. Virtual servers which do not support Ensemble Availability Management will have an <b>avail-status</b> of <b>"not-supported"</b>.</p> <p>Note that only virtual servers whose workload <b>inclusion-type</b> is <b>"direct"</b> or <b>"custom-group"</b> may be added to the list.</p>
<b>eg-overrides</b>	(w) (pc)	Array of objects	<p>The list of Workload Element Group objectives override values in place against workload element groups in the workload when the availability policy is active. Each workload element group override value will be in the form of a workload element group override object, as described in the next table.</p> <p>If no workload element group overrides are defined, an empty array will be returned.</p>
<b>revision</b>	(pc)	Integer (1-n)	The revision count of the availability policy, which is the number of times the policy has been modified. The initial value for a new availability policy is 1.
<b>activation-status</b>	(pc)	String Enum	<p>The activation status of the availability policy, which will have one of the following values:</p> <ul style="list-style-type: none"> <li>• <b>"not-active"</b>: The availability policy is not currently the active policy in the workload.</li> <li>• <b>"active"</b>: The availability policy is currently the active policy in the workload.</li> </ul>
<b>last-activation-date</b>	(pc)	Timestamp	<p>The standard date/time value indicating the time of the last activation of the policy. If the policy has never been active or is the default policy of the default workload, this will be negative.</p> <p>The standard time value is defined as the number of elapsed milliseconds after midnight January 1, 1970.</p>
<b>last-activated-by</b>	(pc)	String	The user name used for the last activation. This is an empty string if the policy has never been activated or is the default policy of the default workload.
<b>last-modified-date</b>	(pc)	Timestamp	<p>The standard date/time value indicating the time the policy was last modified. If the policy has never been modified, this property will be equal to the <b>created-date</b> property. The value will be negative for all default policies.</p> <p>The standard time value is defined as the number of elapsed milliseconds after midnight January 1, 1970.</p>

Table 131. Availability Policy object: type-specific properties (continued)

Name	Qualifier	Type	Description
<b>last-modified-by</b>	(pc)	String	The user name used to last modify the policy. If the policy has never been modified, this property will be equal to the <b>created-by</b> property. The value will be an empty string for all default policies.
<b>created-date</b>	—	Timestamp	The standard date/time value indicating the time the policy was created. The value will be negative for all default policies.  The standard time value is defined as the number of elapsed milliseconds after midnight January 1, 1970.
<b>created-by</b>	—	String	The user name used to create the policy. The value will be an empty string for all default policies.

### Workload element group override nested object

The element group override object is nested within a workload's availability policy. It is used to override the objectives values on an element group in the workload.

Table 132. Availability Policy object: Workload element group override nested object properties

Field name	Type	Rqd/Opt	Description
<b>eg-uri</b>	String/URI		The canonical URI path of the Workload Element Group object whose values to override.
<b>minimum-available-elements</b>	Integer		The minimum number of elements that should be available in the Workload Element Group object.  Value must be greater than or equal to zero.
<b>preferred-available-elements</b>	Integer		The number of elements that are preferred to be available in the Workload Element Group object.  Value must be greater than or equal to <b>minimum-available-elements</b> .

## List Availability Policies

Use the **List Availability Policies** operation to list the availability policies within the target workload.

### HTTP method and URI

**GET** /api/workload-resource-groups/{workload-id}/availability-policies

In this request, the URI variable *{workload-id}* is the object ID of the workload object for which you are requesting a list of availability policies.

### Query parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern to limit returned objects to those that have a matching <b>name</b> property. If a match is found, the response is an array with all policies that match. If no match is found, the response is an empty array.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
avail-policies	Array of objects	Array of nested avail-policy-info objects as described in the following table.

Each nested avail-policy-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Availability Policy object, in the form <code>/api/workload-resource-groups/{workload-id}/availability-policies/{policy-id}</code> .
element-id	String	The <b>element-id</b> of the Availability Policy object. This field value is the <code>{policy-id}</code> portion of the URI path provided by the <b>element-uri</b> field.
name	String	Display name of the Availability Policy object.
activation-status	String Enum	The status of the Availability Policy object, which must be one of the following values: <ul style="list-style-type: none"> <li>• <b>"not-active"</b> – the availability policy is not currently the active policy for the workload.</li> <li>• <b>"active"</b> – the availability policy is currently the active policy for the workload, and its activation has completed.</li> </ul>

## Description

The **List Availability Policies** operation lists the availability policies that belong to the workload targeted by the request URI. The element URI, element ID, display name, and status are provided for each.

If the **name** query parameter is specified, the returned list is limited to the policy in the workload that has a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

A policy is included in the list only if you have object-access permission for that object. Additionally, you must have access to the workload object, or an error response is returned.

If the HMC does not manage the targeted workload, then an error response occurs.

## Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the Workload Resource Group object designated by `{workload-id}`.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.



HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>{workload-id}</i> ) does not designate an existing Workload Resource Group object, or the API user does not have object access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Get Availability Policy Properties

Use the **Get Availability Policy Properties** operation to retrieve the properties of a single Availability Policy object, designated by its element ID.

### HTTP method and URI

**GET** /api/workload-resource-groups/{workload-id}/availability-policies/{policy-id}

#### URI variables

Variable	Description
<i>{workload-id}</i>	Object ID of the Workload Resource Group object for which the policy's properties are to be obtained.
<i>{policy-id}</i>	Element ID of the Availability Policy object for which properties are to be obtained.

### Response body contents

On successful completion, the response body is a JSON object that provides the current values of the properties for the Availability Policy object as defined in “Data model” on page 573. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

The **Get Availability Policy Properties** operation returns the current properties for the Availability Policy object specified by *{policy-id}*.

On successful execution, all of the current properties as defined by the “Data model” on page 573 are provided in the response body and HTTP status code 200 (OK) is returned.

The URI path must designate an existing Workload Resource Group and Availability Policy object and you must have object-access permission to the Workload Resource Group object. If either of these conditions is not met, status code 404 (Not Found) is returned.

### Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the Workload Resource Group object designated by *{workload-id}*.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 577.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{workload-id}</i> ) does not designate an existing Workload Resource Group object, or the API user does not have object access permission to the object.
	63	The element ID in the URI ( <i>{policy-id}</i> ) does not designate an existing Availability Policy object in the workload.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Create Availability Policy

Use the **Create Availability Policy** operation to create a new custom availability policy for a workload resource group. The new policy must be uniquely named within the workload. The policy is inactive until the **Activate Availability Policy** operation is applied to it.

### HTTP method and URI

**POST** /api/workload-resource-groups/{workload-id}/availability-policies

In this request, the URI variable *{workload-id}* is the object ID of the Workload Resource Group object for which you are creating a new availability policy.

### Request body contents

The request body contains properties used to define the new performance policy, which are the writeable properties of a performance policy. Some properties are optional.

Field name	Type	Rqd/Opt	Description
name	String (1-64)	Required	The name to give the new availability policy, as described in the “Data model” on page 573. The passed name must be unique among all other availability policies currently in the workload.
description	String (0-256)	Optional	The description to give the new availability policy, as described in the “Data model” on page 573.
importance	String Enum	Required	The importance value to give the new availability policy, as described in the “Data model” on page 573.
vs-exclusions	Array of String/URI	Optional	The <b>vs-exclusions</b> value to give the new availability policy, as described in the “Data model” on page 573.
eg-overrides	Array of objects	Optional	The <b>eg-overrides</b> value to give to the new availability policy, as described in the “Data model” on page 573.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Availability Policy object, in the form <code>/api/workload-resource-groups/{workload-id}/availability-policies/{policy-id}</code>

## Description

The **Create Availability Policy** operation creates a new custom availability policy in a workload, identified by `{workload-id}`.

The request body contains an object with field names that correspond to the names of properties for this object. The only required properties are the policy name, which must be unique among existing policies in the workload, and the new policy's importance. Supplying a description of the policy is optional. Also optional are the list of non-redundant virtual servers that should not impact the workload's availability status and Workload Element Group objectives overrides. On successful execution, the availability policy is created and added to the workload and status code 201 is returned with a response body containing a reference to the new policy. Note that the new policy is not active until the **Activate Availability Policy** operation is applied to it.

The URI path must designate an existing Workload Resource Group object and you must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, you must also have permission to the **New Availability Policy** task as well, otherwise status code 403 (Forbidden) is returned. Targeting the default Workload Resource Group object will incur a 400 (Bad Request) response, since the default workload may only ever contain its default availability policy.

The request body is validated against the "Data model" on page 573 to ensure that it contains only writeable properties and the data types of those properties are as required. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object passed in the request URI.
- Action/task permission to the **New Availability Policy** task.

## HTTP status and reason codes

On successful completion, HTTP status code 201 (Created) is returned and the response body is provided as described in "Response body contents."

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload is the default Workload Resource Group object.
	65	The name of the availability policy is not unique within its workload.
	66	A virtual server in the request body’s <b>vs-exclusions</b> list does not support Ensemble Availability Management.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{workload-id}</i> ) does not designate an existing Workload Resource Group object, or the API user does not have object access permission to the object.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Delete Availability Policy

Use the **Delete Availability Policy** operation to remove an availability policy from a workload. Note that this operation cannot be performed on an active or default availability policy.

### HTTP method and URI

**DELETE** /api/workload-resource-groups/*{workload-id}*/availability-policies/*{policy-id}*

#### URI variables

Variable	Description
<i>{workload-id}</i>	Object ID of the Workload Resource Group object whose Availability Policy object is to be deleted.
<i>{policy-id}</i>	Element ID of the Availability Policy object that is to be deleted.

## Description

The **Delete Availability Policy** operation deletes the Availability Policy object specified by *{policy-id}* from its workload specified by *{workload-id}*.

On successful execution, the Availability Policy object is removed from the workload and status code 204 (No Content) is returned without a response body.

The URI path must designate an existing Availability Policy object and you must have object-access permission to its workload. If this condition is not met, status code 404 (Not Found) is returned. Additionally, you must have permission to the **Delete Availability Policy** task. If this condition is not met, status code 403 (Forbidden) is returned. Targeting a default availability policy will also invoke an error response, since it cannot be deleted. An active availability policy cannot be deleted as well. To delete an active policy, first target another availability policy in the workload with the Activate Availability Policy operation, and then run the **Delete Availability Policy** operation.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object designated by *{workload-id}*.
- Action/task permission to the **Delete Availability Policy** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned without a response body.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	67	The targeted availability policy is a default Availability Policy object.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{workload-id}</i> ) does not designate an existing Workload Resource Group object, or the API user does not have object access permission to the object.
	63	The element ID in the URI ( <i>{policy-id}</i> ) does not designate an existing Availability Policy object in the workload.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.
	63	The availability policy to be deleted is currently active or in the progress of being activated. The operation can be retried after a different policy has been activated, which causes the <b>activation-status</b> of this policy to be <b>"not-active"</b> .

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Update Availability Policy

Use the **Update Availability Policy** operation to modify one or more writeable properties of an Availability Policy object

### HTTP method and URI

**POST** /api/workload-resource-groups/*{workload-id}*/availability-policies/*{policy-id}*

#### URI variables

Variable	Description
<i>{workload-id}</i>	Object ID of the workload object to which the target availability policy is defined.
<i>{policy-id}</i>	Element ID of the Availability Policy object that is to be modified.

## Request body contents

The request body is a JSON object containing one or more of the writeable fields for an availability policy, as described in the “Data model” on page 573. You need to supply only those fields that you want to modify.

## Description

The **Update Availability Policy** operation updates one or more writeable properties of the Availability Policy object identified by *{policy-id}*.

On successful execution, the Availability Policy object is updated with the supplied property values and status code 204 (No Content) is returned without a response body. Notifications for these property changes are sent asynchronously to this operation.

If the availability policy is active at the time of this request, the Ensemble Availability Management of the workload’s virtual servers to be dynamically updated without a reactivation of the policy.

The URI path must designate existing workload and Availability Policy objects and you must have object-access permission to the workload. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, you must also have permission to the **Availability Policy Details** task as well, otherwise status code 403 (Forbidden) is returned. Targeting a default availability policy will also invoke an error response, since it is not directly mutable.

The request body is validated against the schema described in the “Request body contents.” If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object designated by *{workload-id}*.
- Action/task permission to the **Availability Policy Details** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned without a response body.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	65	The new name given to the availability policy is not unique within its workload.
	67	The targeted performance policy is a default Availability Policy object.
403 (Forbidden)	1	The API user does not have the required permission for this operation.

HTTP error status code	Reason code	Description
404 (Not Found)	1	The object ID in the URI ( <i>{workload-id}</i> ) does not designate an existing Workload Resource Group object, or the API user does not have object access permission to the object.
	63	The element ID in the URI ( <i>{policy-id}</i> ) does not designate an existing Availability Policy object in the workload.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Activate Availability Policy

Use the **Activate Availability Policy** operation to activate a workload’s availability policy. This may be any availability policy in the workload, including the default, whose **activation-status** is **"not-active"**.

### HTTP method and URI

**POST** `/api/workload-resource-groups/{workload-id}/availability-policies/{policy-id}/operations/activate`

#### URI variables

Variable	Description
<i>{workload-id}</i>	Object ID of the workload object to which the target availability policy is defined.
<i>{policy-id}</i>	Element ID of the Availability Policy object that is to be activated.

### Description

The **Activate Availability Policy** operation activates an availability policy identified by *{policy-id}* in its workload identified by *{workload-id}*.

On successful execution, the target policy is active and status code 204 (No Content) is returned without a response body. Notifications for ensuing property changes are sent asynchronously to this operation.

On successful execution, the activation request completes and status code 204 (No Content) is returned without supplying a response body. Notifications for ensuing property changes are sent asynchronously to this operation.

The URI path must designate existing workload and Availability Policy objects and you must have object-access permission to the workload. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, you must also have permission to the **Activate Availability Policy** task as well, otherwise status code 403 (Forbidden) is returned. If the target availability policy is already active, a 409 (Conflict) is returned. Targeting the default workload will also invoke an error response, since its default availability policy is permanently active.

### Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the workload object designated by *{workload-id}*.
- Action/task permission to the **Activate Availability Policy** task.

## HTTP status and reason codes

On successful completion, HTTP status code 204 (No Content) is returned without a response body.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	60	The targeted workload is the default Workload Resource Group object.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{workload-id}</i> ) does not designate an existing Workload Resource Group object, or the API user does not have object access permission to the object.
	63	The element ID in the URI ( <i>{policy-id}</i> ) does not designate an existing Availability Policy object in the workload.
409 (Conflict)	2	The operation cannot be performed because the object designated by the request URI is currently busy performing some other operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

---

## Ensemble Availability Management reports

Each ensemble, with its set of workloads, contains Ensemble Availability Management components. The APIs related to actual creation, activation, deletion, etc. of Ensemble Availability Management components within the object model are contained in the previous sections where the objects are defined. This subsection describes a special category of Ensemble Availability Management APIs related to reporting.

Ensemble Availability Management has the ability to generate historical reports, on demand, which are based on reporting data that it has been saved over time which is related to the Ensemble Availability Management objects as they existed at those points in time. But since the actual objects involved may have changed since then, or actually may not even exist now, these reports are being treated differently than the normal object based APIs.

**Note:** The Ensemble Availability Management historical reporting data currently goes back up to 36 hours from the current time.

Instead of being object based, all of the Ensemble Availability Management report APIs will be implemented as availability-management specific operations, that request an on-demand report be generated and returned to the caller. They will take in the query parameters in a request block, and essentially do the queries into the historical reporting database to generate a custom report for that caller.

The Ensemble Availability Management report APIs are interrelated, in that certain properties retrieved from one report can be used as input to generate other reports. One could visualize drilling down for more detailed information about a specific item found within a previous report; just as one might traverse these same reports when using the Hardware Management Console user interface.

The following chart shows how the reports are related, along with an indication of which **-report-id** properties (explained in more detail below the chart) would be used from one report to generate the next



one via the APIs. The actual descriptions of these reports, how to generate them via the APIs, and the API properties, are fully described following this introduction. Therefore, it is suggested this chart be referred to again after reading through and understanding each of them individually.

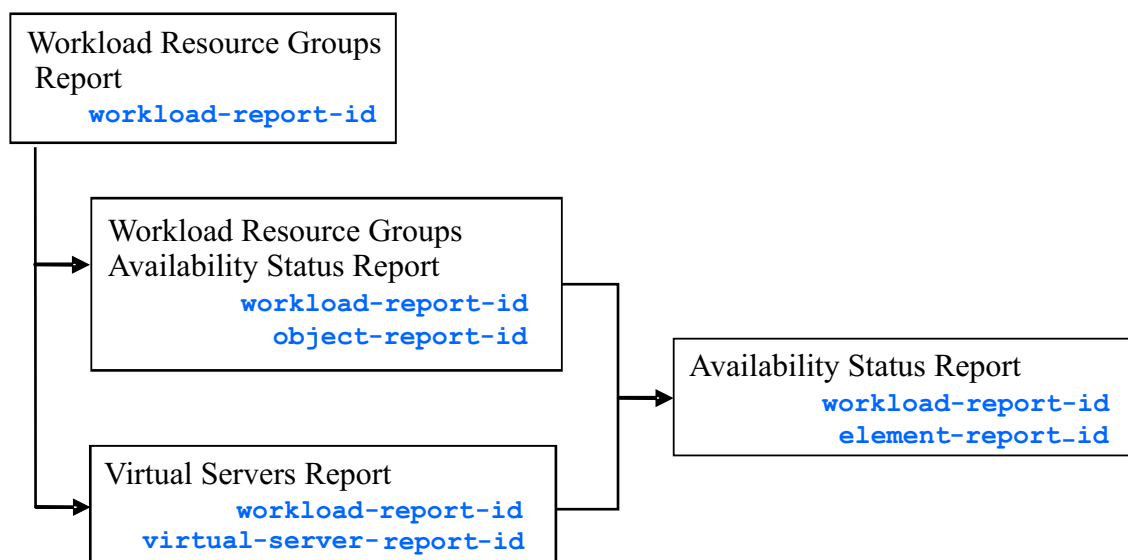


Figure 281. Relationship between reports and the properties used

The special properties, all containing names ending in **-report-id**, are the actual query keys into the internal reporting system. If following the drill down approach, one would never have to actually know what those key values are. Once they are returned as response properties from one report, the values in them can just be used again as-is for input request properties to subsequent reports concerning those same items (i.e. for the same virtual server or workload or hypervisor or service class).

Furthermore, they should be considered durable in that, for the same items, they do not change across invocations of the reporting APIs. For example, once the **workload-report-id** for a particular workload is obtained, it can be kept around and used over and over again when trying to generate subsequent reports involving that particular workload, regardless of the time interval of the report. In other words, once known for a particular item, they will remain constant, and so it is not necessary to start drilling again from the top every time to obtain a particular sub-report that requires any of those values as input request block properties.

**Note:** As a general rule, the reporting system internal keys are intentionally not documented as to what specific information is actually used to index the data for internal retrieval; so that there would be no client impact should the type of data used to internally index the data change in the future. But in the case of **virtual-server-report-id** and **object-report-id** (Virtual Servers Report), where the currently in use key types are unlikely to change in the future, the keys are being documented in order to provide an additional convenient way to obtain those key values directly from currently existing Workload objects. This allows a shortcut for calling certain select sub-reports directly, without having to obtain the key values via other calls to higher level reports.

- **virtual-server-report-id:** the value of the **object-id** property of the Virtual Server object.
- **object-report-id:** the value of the **object-id** property of the workload element group or Virtual Server object.

For example, if there is a Workload in the ensemble for which an Availability Status Report is desired, the API caller would be able to get it directly by providing the value from its **object-id** property as the input request **workload-report-id** property (as opposed to drilling through the Workload Resource Groups Report to get the **workload-report-id** from the reporting system).

As mentioned above concerning the **-report-id** values in general, though, the **workload-report-id** values obtained from the reports are durable over time, so if the calling application has a design such that it has already drilled down and saved these values for future use anyway, a direct call to the sub-report could be made using it from there as well.

The **-report-id** values are consistent between the performance management reports and Ensemble Availability Management. This allows for performance management and Ensemble Availability Management report data to be combined using **-report-id** values to generate combined workload performance and Ensemble Availability Management reports.

## Generate Workload Resource Groups Report (Ensemble Availability Management)

Use the **Generate Workload Resource Groups Report** of Ensemble Availability Management operation to create a custom on-demand Workload Resource Groups Report for a requested time period, based on historical Ensemble Availability Management data that was retained over that time period.

### HTTP method and URI

**POST** /api/ensembles/{ensemble-id}/availability-management/operations/generate-workload-resource-groups-report

In this request, the URI variable {ensemble-id} is the object ID of the ensemble object for which you are requesting a workload resource groups report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.

Field name	Type	Description
report-workloads	Array of objects	Array of nested workload-report-entry objects described in the next table.

Each nested workload-report-entry object contains the following fields:

*Table 133. Format of a workload-report-entry object*

Field name	Type	Description
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	The <b>object-id</b> property of the workload object represented by this entry. Because the report contains historical data, the workload object with this object ID may no longer exist.
workload-object-uri	String/URI	The <b>object-uri</b> property of the workload object represented by this entry. Because the report contains historical data, the workload object with this object URI may no longer exist.
workload-object-id	String	The <b>object-id</b> property of the workload object represented by this entry. Because the report contains historical data, the workload object with this object ID may no longer exist.
availability-policy-name	String	The name (displayable) of the availability policy that was active during this reporting interval. <b>Note:</b> If more than one availability policy was activated over the span of this reporting interval, only the name of the last one activated will be returned here. See the <b>multi-policy-activations</b> field to determine whether there was more than one.
availability-policy-element-uri	String/URI	The <b>element-uri</b> of the availability policy that was active during this reporting interval. Because the report contains historical data, the workload object with this <b>element-uri</b> may no longer exist.
availability-policy-element-id	String	The <b>element-id</b> property of the Availability Policy object represented by this entry. Because the report contains historical data, the workload object with this <b>element-id</b> may no longer exist.
most-severe-availability-status	String/Enum	The most severe <b>avail-status</b> value recorded for the workload over this reporting interval.
multi-policy-activations	Boolean	Indicates whether more than one availability policy was activated over the span of this reporting interval (true if more than one, false if not more than one).
avail-status-data-points	Array of objects	Array of nested avail-status-data-point objects, described in the next table.

Each nested avail-status-data-point object contains the following fields:

Table 134. Format of nested avail-status-data-point object

Field name	Type	Description
time	Integer	Standard date/time value indicating the date and time that this availability status value was recorded.  The standard time value is defined as the number of elapsed milliseconds after midnight January 1, 1970.
avail-status	String/ Enum	The <b>avail-status</b> value of the workload at the recorded time. See the “Data model” on page 443 for more details about the valid values of this property.

## Description

This operation generates an Ensemble Availability Management Workload Resource Groups Report. For a requested interval, designated by a start date/time and duration, it will generate and return a list of all workloads for which there is historical Ensemble Availability Management reporting information available. Additionally, for each workload, it also returns the active availability policy at that time and an availability health indication (availability status) over that interval. Please refer to the list of returned fields in the previous tables for a full list of all data that is returned in this report for each workload.

It should be noted that the returned **workload-object-id** for each returned workload can be used as input in the request bodies for further availability reports defined in this section that can be generated to dig deeper into other Ensemble Availability Management reporting data for a specific workload.

If there is no reporting data available for the requested time interval, an empty response object will be provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Workloads Resource Groups Report** task.

## HTTP status and reason codes

The following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP Status Code	Reason Code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See section “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform the requested operation (Action/task permission failure).
404 (Not Found)	1	An object with <b>object-id</b> <i>{ensemble-id}</i> does not exist on HMC or API user does not have object-access permission for the identified ensemble.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-123456789000/availability-management/operations/generate-workload-resource-groups-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 30
}
```

---

Figure 282. Generate Workload Resource Groups Report (Ensemble Availability Management): Request

## Generate Workload Resource Group Availability Status Report

Use the **Generate Workload Resource Group Availability Status Report** of an Ensemble Availability Management operation to create a custom on-demand availability status report for a specific workload over a requested time period. It is based on historical Ensemble Availability Management data that was retained over that time period.

### HTTP method and URI

```
POST /api/ensembles/{ensemble-id}/availability-management/operations/generate-workload-resource-group-availability-status-report
```

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble object for which you are requesting a report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated.  The standard time value is defined as the number of elapsed milliseconds after midnight January 1, 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report should cover (from the report-interval-start-time). Value must be greater than 0.
workload-report-id	String	Required	An identifier used by the Ensemble Availability Management reporting structure to keep track of reporting data for the specific workload this entry represents.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.

Field name	Type	Description
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The name (displayable) of the workload this report covers.
workload-report-id	String	An identifier used by the Ensemble Availability Management reporting structure to keep track of reporting data for the specific workload, named in workload-name, this report covers.
workload-object-uri	String/URI	The <b>object-uri</b> property of the workload object represented by this entry. Because the report contains historical data, the workload object with this <b>object-uri</b> may no longer exist.
policy-activation-records	Array of objects	Array of nested policy-activation-record objects described in Table 135. If none, an empty array will be returned for this field.
availability-status-records	Array of objects	Array of nested availability-status-record objects described in Table 136 If none, an empty array will be returned for this field.

Each nested policy-activation-record object contains the following fields:

*Table 135. Format of a policy-activation-record object*

Field name	Type	Description
time	Integer	Standard date/time value indicating the date and time that the availability policy was activated. The standard time value is defined as the number of elapsed milliseconds after midnight January 1, 1970.
availability-policy-name	String	The name (displayable) of the availability policy that was activated.
availability-policy-element-uri	String/URI	The <b>element-uri</b> property of the availability policy that was activated. Because the report contains historical data, the Availability Policy object with this <b>element-uri</b> may no longer exist.
availability-policy-revision	Integer (1-n)	The revision property of the availability policy that was activated.

Each availability-status-record object contains the following fields:

*Table 136. Format of availability-status-record object*

Field name	Type	Description
class	String Enum	The class of the object. Values: <ul style="list-style-type: none"> <li>• <b>"workload-resource-group"</b>: Entity is a workload resource group.</li> <li>• <b>"workload-element-group"</b>: Entity is a workload element group.</li> <li>• <b>"virtual-server"</b>: Entity is a virtual server.</li> </ul>
entity-name	String Enum	The name of the workload, workload element group, virtual server, or availability policy represented by the entry.

Table 136. Format of availability-status-record object (continued)

Field name	Type	Description
object-uri	String/URI	The <b>object-uri</b> property of the workload, workload element group, or Virtual Server object represented by this entry.  Because the report contains historical data, the object with this <b>object-uri</b> may no longer exist.
object-report-id	String	An identifier used by the Ensemble Availability Management reporting structure to keep track of reporting data for the specific workload, workload element group, or virtual server.
time	Integer	Standard date/time value indicating the date and time that this availability status change took place.  The standard time value is defined as the number of elapsed milliseconds after midnight January 1, 1970.
availability-policy-name	String	The name (displayable) of the availability policy that was active when this availability status change was reported.  Value is always null for entities whose <b>class</b> is " <b>virtual-server</b> ".
availability-policy-element-uri	String/URI	The <b>element-uri</b> property of the availability policy that was active when this availability status change was reported.  Because the report contains historical data, the Availability Policy object with this element-uri may no longer exist.  Value is always null for entities whose <b>class</b> is " <b>virtual-server</b> ".
availability-policy-revision	Integer (1-n)	The <b>revision</b> property of the availability policy that was active when this availability status change was reported.  Value is always null for entities whose <b>class</b> is " <b>virtual-server</b> ".
parents	Array of objects	Array of nested parent objects, described in the next table, listing the parent objects of this entity. These parents may be affected by the availability status change described by this record (and therefore have their own availability status change record in the report).  If none, an empty array will be returned for this field.  Value is always null for entities whose <b>class</b> is " <b>workload-resource-group</b> ".
old-avail-status	String Enum	The <b>avail-status</b> value of the object before this change.
new-avail-status	String Enum	The <b>avail-status</b> value of the object after this change.
avail-status-impact-exclusion	Boolean	The avail-status value of the object was excluded from impacting the workload availability status. This condition occurs when an object is selected to have its workload availability status impact ignored in an availability policy.  Value is always false if the <b>class</b> is " <b>workload-resource-group</b> " or " <b>workload-element-group</b> " or if the <b>class</b> is " <b>virtual-server</b> " and the <b>parent-class</b> is " <b>workload-element-group</b> ".

Table 136. Format of availability-status-record object (continued)

Field name	Type	Description
reason	String Enum	<p>The reason for the new availability status. Values:</p> <ul style="list-style-type: none"> <li>• <b>"at-least-preferred-available"</b>: The number of available virtual servers in the workload element group is greater than or equal to its preferred available objective.</li> <li>• <b>"fewer-than-preferred-available"</b>: The number of available virtual servers in the workload element group is fewer than its preferred available objective.</li> <li>• <b>"at-least-minimum-available"</b>: The number of available virtual servers in the workload element group is greater than or equal to its minimum available objective.</li> <li>• <b>"fewer-than-minimum-available"</b>: The number of available virtual servers in the workload element group is fewer than its minimum available objective.</li> <li>• <b>"status-is-exposed"</b>: Virtual server operating status indicates the element is exposed.</li> <li>• <b>"status-is-critical"</b>: Virtual server operating status indicates the element is critical.</li> <li>• <b>"status-is-notavailable"</b>: Virtual server operating status indicates the element is not available.</li> <li>• <b>"status-is-available"</b>: Virtual server operating status indicates the element is available.</li> <li>• <b>"gmp-is-exposed"</b>: GPMP operating status indicates the element is exposed.</li> <li>• <b>"gmp-is-available"</b>: GPMP operating status indicates the element is available.</li> <li>• <b>"member-is-exposed"</b>: Member operating status indicates the element is exposed.</li> <li>• <b>"member-is-critical"</b>: GPMP operating status indicates the element is critical.</li> <li>• <b>"member-is-available"</b>: GPMP operating status indicates the element is available.</li> <li>• <b>"entity-was-created"</b>: The entity was created.</li> <li>• <b>"entity-was-deleted"</b>: The entity was deleted.</li> <li>• <b>"entity-was-updated"</b>: The entity was updated.</li> <li>• <b>"virtual-server-deactivated"</b>: The virtual server was deactivated.</li> <li>• <b>"virtual-server-status-changed"</b>: The virtual server's operating status changed.</li> <li>• <b>"virtual-server-not-supported"</b>: The virtual server is not supported by Ensemble Availability Management.</li> <li>• <b>"policy-activated"</b>: A new policy was activated for the workload.</li> </ul>

Each nested policy-activation-record object contains the following fields:



Table 137. Format of a parent object

Field name	Type	Description
name	String Enum	<p>The name of the parent entity whose availability status may be affected by this entity's availability status change.</p> <p>A parent entity is defined by the name, URI, and <b>class</b> properties. The parent object's URI property and the <b>time</b> property of the <b>availability-status-record</b> object that contain the parent can be used to find the <b>availability-status-record</b> for the parent (the record whose <b>object-uri</b> matches this parent object's URI and whose time matches the record's time).</p> <p>If the record is of class <b>"virtual-server"</b>, the parent entity may be the workload in which the virtual server is a direct member (the target of the report) or the workload element group in which the virtual server is a member.</p> <p>If the record is of class <b>"workload-element-group"</b>, the parent entity is the workload in which the workload element group is a member (the target of the report).</p>
object-uri	String/URI	The <b>object-uri</b> property of the parent entity.
class	String Enum	<p>The <b>class</b> property of the parent entity. Values:</p> <ul style="list-style-type: none"> <li><b>"workload-resource-group"</b>: Entity is a workload resource group.</li> <li><b>"workload-element-group"</b>: Entity is a workload element group.</li> </ul>

## Description

In general, an availability status report shows information about the availability status changes for the workload and its elements. For each availability status change within the workload, the report lists the object to which the change occurred, why the change occurred, and what parent object(s) may be affected by the availability status change.

Multiple availability status updates resulting from a workload element availability status change will result in all related records being recorded with the same time. This time link and the parent properties allow one to link records together for reporting. For example:

time	object-uri	class	parent-uri	old-availability-status	new-availability-status
1316101902	<wl1-uri>	workload	null	critical	available
1316101902	<eg1-uri>	workload-element-group	<wl1-uri>	critical	exposed
1316101902	<vs1-uri>	virtual-server	<eg2-uri>	exposed	available
1316101902	<vs1-uri>	virtual-server	<eg2-uri>	exposed	available

If there is no reporting data available for the requested time interval, an empty response object will be provided and the operation completes successfully. If there are no availability status changes to report, then the **availability-status-records** array will be returned as an empty array.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble.
- Action/task permission to the **Availability Status Event Report** task.

## HTTP status and reason codes

The following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP Status Code	Reason Code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See section “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform the requested operation (Action/task permission failure).
404 (Not Found)	1	An object with <b>object-id</b> <i>{ensemble-id}</i> does not exist on HMC or API user does not have object-access permission for the identified ensemble.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/availability-management/operations/generate-workload-resource-group-availability-status-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group"
}
```

---

Figure 283. Generate Workload Resource Group Availability Status Report: Request

## Generate Virtual Servers Report (Ensemble Availability Management)

Use the **Generate Virtual Servers Report** of an Ensemble Availability Management operation to create a custom on-demand report that shows all virtual servers that were members of a particular workload over a requested time period for which reporting data is available. The report is based on historical performance management data that was retained over that time period.

### HTTP method and URI

**POST** `/api/ensembles/{ensemble-id}/availability-management/operations/generate-virtual-servers-report`

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble object for which you are requesting a virtual server report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Req/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.

Field name	Type	Req/Opt	Description
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
workload-report-id	String	Required	The identifier used by the Ensemble Availability Management reporting to keep track of reporting data for the specific workload for which the report is being requested. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the <b>Workload Resource Groups Report</b> for this same interval.

## Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the Ensemble Availability Management reporting structure to keep track of reporting data for the specific workload, named in the <b>workload-name</b> field.
report-virtual-servers	Array of objects	Array of nested virtual-server-report-entry objects, described in the next table.

Each nested virtual-server-report-entry object contains the following fields:

Field name	Type	Description
virtual-server-name	String	The displayable name of the virtual server.
virtual-server-report-id	String	An identifier used by the Ensemble Availability Management reporting structure to keep track of reporting data for the specific virtual server, because virtual server names by themselves are not unique.
most-severe-availability-status	String Enum	The most severe <b>avail-status</b> value recorded for the virtual server over the specified reporting interval. See the “Data model” on page 443 for more details about the valid values of this property.
avail-status-impact-exclusion	Boolean	The <b>most-severe-availability-status</b> value was excluded from impacting the workload availability status. This condition occurs when a virtual server is selected to have its workload availability status impact ignored in an availability policy.

## Description

The **Generate Virtual Servers Report** operation generates an Ensemble Availability Management report that contains the following information for a specific workload over the requested time interval:

- A list of all virtual servers that were members in that specific workload for which historical Ensemble Availability Management reporting information is available
- For each virtual server:
  - The unique availability reporting identifier for the virtual server
  - The name and unique availability reporting identifier of the hypervisor under which the virtual server was running
  - An indication of whether the virtual server was active over the reporting interval.

If there is no reporting data available for the requested time interval, an empty response object will be provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Virtual Servers Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 595.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group"
}
```

---

Figure 284. Generate Virtual Servers Report: Request

## Generate Availability Status Report

Use the **Generate Availability Status Report** of an Ensemble Availability Management operation to create a custom on-demand report showing the availability status of a workload, workload virtual server, or workload element group over a requested time period for which reporting data is available. It is based on historical Ensemble Availability Management data that was retained over that time period.

### HTTP method and URI

**POST** /api/ensembles/{ensemble-id}/availability-management/operations/generate-availability-status-report

In this request, the URI variable {ensemble-id} is the object ID of the ensemble object for which you are requesting an availability status report.

### Request body contents

The request body is a JSON object with the following fields:

Field name	Type	Req/Opt	Description
report-interval-start-time	Integer	Required	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
report-interval-duration	Integer	Required	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0.
workload-report-id	String	Required	The identifier used by the Ensemble Availability Management reporting to keep track of reporting data for the specific workload for which the report is being requested. <b>Note:</b> This value is the same as one of the <b>workload-report-id</b> values that are returned in the <b>Workload Resource Groups Report</b> , <b>Virtual Servers Report</b> , or <b>Workload Resource Groups Availability Status Report</b> for this same interval.
element-report-id	String	Required	The identifier used by the Ensemble Availability Management reporting to keep track of reporting data for the specific workload element for which the report is being requested or null if the report covers a workload and not a workload element. <b>Note:</b> This value is the same as one of the <b>virtual-server-report-id</b> values that are returned in the <b>Virtual Servers Report</b> , or the <b>object-report-id</b> returned in the <b>Workload Resource Groups Availability Status Report</b> for this same interval.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
report-interval-start-time	Integer	Standard date/time value indicating the requested starting date and time of the report to be generated. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970. <b>Note:</b> The returned value might differ from the requested time if no report data was available at the requested time, but data was available starting later within the requested interval.

Field name	Type	Description
report-interval-duration	Integer	The length, in minutes, of the interval this report covers, starting from the value specified for the <b>report-interval-start-time</b> field. This value must be greater than 0. <b>Note:</b> The returned value might differ from the requested duration if report data was not available for the entire requested interval, but data was available for a shorter duration within that same requested interval.
workload-name	String	The displayable name of the workload resource group.
workload-report-id	String	An identifier used by the Ensemble Availability Management reporting structure to keep track of reporting data for the specific workload, named in the <b>workload-name</b> field.
element-name	String	The displayable name of the workload element this report covers or null if the report covers a workload and not a workload element.
element-report-id	String	An identifier used by the Ensemble Availability Management reporting structure to keep track of reporting data for the specific workload element, specified by <b>element-name</b> , or null if the report covers a workload and not a workload element.
avail-status-data-points	Array of objects	Array of nested avail-status-data-point objects, described in the next table.

Each nested avail-status-data-point object contains the following fields:

Field name	Type	Description
time	Integer	Standard date/time value indicating the requested starting date and time that this value was recorded. The standard time value is defined as the number of elapsed milliseconds after midnight on 1 January 1970.
avail-status	String Enum	The <b>avail-status</b> value of the workload at the recorded time. See the “Data model” on page 443 for more details about the valid values of this property.

## Description

The **Generate Availability Status Report** operation generates an Ensemble Availability Management report that contains the following information for a specific workload over the requested time interval:

- A list of all virtual servers that were members in that specific workload for which historical Ensemble Availability Management reporting information is available
- For each virtual server:
  - The unique availability reporting identifier for the virtual server
  - The name and unique availability reporting identifier of the hypervisor under which it was running
  - An indication of whether the hypervisor was active over the reporting interval.

If there is no reporting data available for the requested time interval, an empty response object will be provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the ensemble object passed in the request URI.
- Action/task permission to the **Virtual Servers Report** task.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 597.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing Ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
POST /api/ensembles/12345678-1234-1234-1234-123456789000/availability-management/operations/generate-virtual-servers-report
{
  "report-interval-start-time": 1296149252662,
  "report-interval-duration": 60,
  "workload-report-id": "Payroll Workload Resource Group"
}
```

---

Figure 285. Generate Virtual Servers Report: Request

## Get Performance Management Velocity Level Range Mappings

Use the **Get Performance Management Velocity Level Range Mappings** operation to retrieve the numeric ranges that zManager uses to calculate the actual performance velocity for a service class. Each range of numbers is mapped to one of the valid descriptive values that are used in performance management reports for service classes that have a velocity rather than discretionary performance goal.

### HTTP method and URI

**GET** /api/ensembles/{*ensemble-id*}/performance-management/velocity-level-range-mappings

In this request, the URI variable *{ensemble-id}* is the object ID of the ensemble object.

### Response body contents

On successful completion, the response body is a JSON object with the following fields:

Field name	Type	Description
performance-management-velocity-level-range-mappings	Array of objects	An array of nested velocity-level-range-mapping-entry objects.

Each nested velocity-level-range-mapping-entry object contains the following fields:

Field name	Type	Description
low-boundary	Integer	This value defines the low boundary (inclusive) that defines this specific range of calculated velocity numbers.
high-boundary	Integer	This value defines the upper boundary (inclusive) that defines this specific range of calculated velocity numbers.
velocity-level	String Enum	The descriptive velocity level that this numeric range represents. Possible values are: <b>"fastest"</b> , <b>"fast"</b> , <b>"moderate"</b> , <b>"slow"</b> , and <b>"slowest"</b> .

## Description

The **Get Performance Management Velocity Level Range Mappings** operation retrieves the list of numeric ranges that zManager uses to calculate the actual performance for a service class that has a performance goal of the **"velocity"** (rather than **"discretionary"**) type. Each range of numbers is mapped to one of the valid descriptive values that are used in performance management reports; these descriptive values are: **"fastest"**, **"fast"**, **"moderate"**, **"slow"**, and **"slowest"**.

For service classes with a velocity performance goal, zManager:

1. Calculates a numeric value for actual performance
2. Compares the calculated numeric value to a predefined set of ranges, one range for each of the descriptive velocity-level values
3. Substitutes the appropriate descriptive value to present in performance management reports.

This substitution occurs because the actual performance is not significantly different between numbers that are close together; it is more meaningful to define a smaller set of velocity levels that describe the performance.

Through the **Get Performance Management Velocity Level Range Mappings** API and the service class metrics described in “Workload service class data metrics group” on page 907, you can calculate raw velocity numbers and compare and substitute the performance level using the same descriptive velocity levels that zManager uses in its performance management reports.

“Response body contents” on page 599 describes the full list of data that is returned for the **Get Performance Management Velocity Level Range Mappings** operation. An error response is returned if the targeted ensemble does not exist or if you do not have the requirements listed in “Authorization requirements.”

## Authorization requirements

This operation has the following authorization requirement:

- object-access permission to the ensemble object passed in the request URI.

## HTTP status and reason codes

On successful completion, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 599.

Otherwise, the following HTTP status codes are returned for the indicated errors. The response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.



HTTP error status code	Reason code	Description
404 (Not Found)	1	The object ID in the URI ( <i>{ensemble-id}</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/ensembles/12345678-1234-1234-1234-123456789000/performance-management/velocity-level-range-mappings
```

---

Figure 286. Get Performance Management Velocity Level Range Mappings: Request

## Inventory service data

Information about the Workload Resource Groups managed by the HMC and the associated Performance Policies can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for workload and associated policy objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "workload-resource-group" are to be included. Information for a particular workload (and associated policies) is included only if the API user has object-access permission to that object.

For each workload to be included, the inventory response array includes the following:

- An array entry for the workload object itself. This entry is a JSON object with the same contents as is specified in the Response Body Contents section for “Get Workload Resource Group Properties” on page 450. That is, the data provided is the same as would be provided if a **Get Workload Resource Group Properties** operation were requested targeting this object.
- An array entry for each policy associated with the workload. For each such policy, an entry is included that is a JSON object with the same contents as specified in the Response Body Contents section of “Get Performance Policy Properties” on page 483. As a result, the data provided is the same as would be obtained if a **Get Performance Policy Properties** operation were requested for each policy listed by a **List Performance Policies** operation targeting the workload.

The array entry for a workload object will appear in the results array before entries for associated performance policies.

## Sample inventory data

The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single workload (named “Online Ordering”) that defines a custom performance policy (“Prime shift”) in addition to the default policy. These objects would appear as a sequence of array entries in the response array:

---

```
{
  "active-perf-policy": {
    "activation-status": "active",
    "element-id": "cc79445a-95e4-11e0-b6ef-000c29bb873c",
    "element-uri": "/api/workload-resource-groups/cc79fa6c-95e4-11e0-b6ef-000c29bb873c
/performance-policies/cc79445a-95e4-11e0-b6ef-000c29bb873c",
    "name": "Prime shift"
  },
  "category": "Retail Operations",
  "class": "workload-resource-group",
  "custom-perf-policies": [
    {
      "activation-status": "active",
      "element-id": "cc79445a-95e4-11e0-b6ef-000c29bb873c",
      "element-uri": "/api/workload-resource-groups/cc79fa6c-95e4-11e0-b6ef-000c29bb873c
/performance-policies/cc79445a-95e4-11e0-b6ef-000c29bb873c",
      "name": "Prime shift"
    }
  ],
  "default-perf-policy": {
    "activation-status": "not-active",
    "element-id": "cc7b38dc-95e4-11e0-b6ef-000c29bb873c",
    "element-uri": "/api/workload-resource-groups/cc79fa6c-95e4-11e0-b6ef-000c29bb873c
/performance-policies/cc7b38dc-95e4-11e0-b6ef-000c29bb873c",
    "name": "Default"
  },
  "description": "Workload Resource Group for managing the online ordering application.",
  "is-default": false,
  "name": "Online Ordering",
  "object-id": "cc79fa6c-95e4-11e0-b6ef-000c29bb873c",
  "object-uri": "/api/workload-resource-groups/cc79fa6c-95e4-11e0-b6ef-000c29bb873c",
  "parent": "/api/ensembles/890b6df2-93a4-11e0-887c-000c29bb873c",
  "perf-activation-node-count": 0,
  "perf-activation-status": "active"
},
{
  "activation-status": "active",
  "class": "performance-policy",
  "created-by": "ENSADMIN",
  "created-date": 1307987073673,
```

---

Figure 287. Workload Resource Group: Sample inventory data (Part 1)

---

```
"custom-service-classes": [  
  {  
    "business-importance": "highest",  
    "classification-rule": {  
      "filter": {  
        "operation": "string-match",  
        "type": "virtual-server-name",  
        "value": "00WS(\\*)"  
      },  
      "type": "rule"  
    },  
    "description": "",  
    "goal-type": "velocity",  
    "name": "Web servers",  
    "type": "server",  
    "velocity": "fast"  
  },  
  {  
    "business-importance": "highest",  
    "classification-rule": {  
      "filter": {  
        "operation": "string-match",  
        "type": "virtual-server-name",  
        "value": "00DB(\\*)"  
      },  
      "type": "rule"  
    },  
    "description": "",  
    "goal-type": "velocity",  
    "name": "DB Servers",  
    "type": "server",  
    "velocity": "fastest"  
  }  
],
```

---

Figure 288. Workload Resource Group: Sample inventory data (Part 2)

---

```
"default-service-class": {
  "business-importance": "medium",
  "classification-rule": {
    "filter": {
      "operation": "string-match",
      "type": "(\\*)",
      "value": "(\\*)"
    },
    "type": "rule"
  },
  "description": "The default workload performance policy service class.",
  "goal-type": "velocity",
  "name": "Default",
  "type": "server",
  "velocity": "moderate"
},
"description": "Performance policy for prime shift operatoin",
"element-id": "cc79445a-95e4-11e0-b6ef-000c29bb873c",
"element-uri": "/api/workload-resource-groups/cc79fa6c-95e4-11e0-b6ef-000c29bb873c/performance-policies/cc79445a-95e4-11e0-b6ef-000c29bb873c",
"importance": "highest",
"is-default": false,
"last-activated-by": "PEDEBUG",
"last-activation-completed-date": 1307991376227,
"last-activation-requested-date": 1307991376226,
"last-modified-by": "ENSADMIN",
"last-modified-date": 1307987073673,
"name": "Prime shift",
"parent": "/api/workload-resource-groups/cc79fa6c-95e4-11e0-b6ef-000c29bb873c",
"revision": 1
},
{
  "activation-status": "not-active",
  "class": "performance-policy",
  "created-by": "",
  "created-date": -1,
  "custom-service-classes": [],
```

---

Figure 289. Workload Resource Group: Sample inventory data (Part 3)

---

```

"default-service-class": {
  "business-importance": "medium",
  "classification-rule": {
    "filter": {
      "operation": "string-match",
      "type": "(\\*)",
      "value": "(\\*)"
    },
    "type": "rule"
  },
  "description": "The default workload performance policy service class.",
  "goal-type": "velocity",
  "name": "Default",
  "type": "server",
  "velocity": "moderate"
},
"description": "The default workload performance policy",
"element-id": "cc7b38dc-95e4-11e0-b6ef-000c29bb873c",
"element-uri": "/api/workload-resource-groups/cc79fa6c-95e4-11e0-b6ef-000c29bb873c/performance-policies/cc7b38dc-95e4-11e0-b6ef-000c29bb873c",
"importance": "medium",
"is-default": true,
"last-activated-by": "PEDEBUG",
"last-activation-completed-date": 1307991361463,
"last-activation-requested-date": 1307991361463,
"last-modified-by": "",
"last-modified-date": -1,
"name": "Default",
"parent": "/api/workload-resource-groups/cc79fa6c-95e4-11e0-b6ef-000c29bb873c",
"revision": 1
}

```

---

Figure 290. Workload Resource Group: Sample inventory data (Part 4)



## Chapter 14. Core z Systems resources

These APIs provide access to and control of the following HMC/SE objects:

- Console
- Group
- CPC
- Logical Partition

In addition, these APIs provide access to the following CPC-related data items:

- Reset Activation Profiles
- Image Activation Profiles
- Load Activation Profiles
- Group Profiles
- Capacity Records

### Operations Summary

Following are the operations summaries for each of the core z Systems objects.

#### Console operations summary

The following table provides an overview of the operations provided for Console objects.

Table 138. Core z Systems resources - Console: operations summary

Operation name	HTTP method and URI path
"Get Console Properties" on page 621	GET /api/console
"Restart Console" on page 624	POST /api/console/operations/restart
"Make Console Primary" on page 626	POST /api/console/operations/make-primary
"Shutdown Console" on page 627	POST /api/console/operations/shutdown
"Reorder User Patterns" on page 628	POST /api/console/operations/reorder-user-patterns
"Get Console Audit Log" on page 630	GET /api/console/operations/get-audit-log
"Get Console Security Log" on page 633	GET /api/console/operations/get-security-log
"List Console Hardware Messages" on page 637	GET /api/console/hardware-messages
"Get Console Hardware Message Properties" on page 639	GET /api/console/hardware-messages/{hardware-message-id}
"Delete Console Hardware Message" on page 640	DELETE /api/console/hardware-messages/{hardware-message-id}

Table 139. Core z Systems resources - Console: URI variables

URI variable	Description
{hardware-message-id}	Element ID of the hardware message object

## User operations summary

The following table provides an overview of the operations provided for User objects.

Table 140. Core z Systems resources - User: operations summary

Operation name	HTTP method and URI path
"List Users" on page 650	GET /api/console/users
"Get User Properties" on page 652	GET /api/users/{user-id}
"Update User Properties" on page 654	POST /api/users/{user-id}
"Add User Role to User" on page 657	POST /api/users/{user-id}/operations/add-user-role
"Remove User Role from User" on page 659	POST /api/users/{user-id}/operations/remove-user-role
"Create User" on page 661	POST /api/console/users
"Delete User" on page 664	DELETE /api/users/{user-id}

Table 141. Core z Systems resources - User: URI variables

URI variable	Description
{user-id}	Object ID of the User object

## User Role operations summary

The following table provides an overview of the operations provided for User Role objects.

Table 142. Core z Systems resources - User Role: operations summary

Operation name	HTTP method and URI path
"List User Roles" on page 670	GET /api/console/user-roles
"Get User Role Properties" on page 673	GET /api/user-roles/{user-role-id}
"Update User Role Properties" on page 675	POST /api/user-roles/{user-role-id}
"Add Permission to User Role" on page 677	POST /api/user-roles/{user-role-id}/operations/add-permission
"Remove Permission from User Role" on page 680	POST /api/user-roles/{user-role-id}/operations/remove-permission
"Create User Role" on page 682	POST/api/console/user-roles
"Delete User Role" on page 684	DELETE /api/user-roles/{user-role-id}



Table 143. Core z Systems resources - User Role: URI variables

URI variable	Description
{user-role-id}	Object ID of the User Role object

## Task operations summary

The following table provides an overview of the operations provided for Task objects.

Table 144. Core z Systems resources - Task: operations summary

Operation name	HTTP method and URI path
"List Tasks" on page 688	GET /api/console/tasks
"Get Task Properties" on page 689	GET /api/console/tasks/{task-id}

Table 145. Core z Systems resources - Task: URI variables

URI variable	Description
{task-id}	Element ID of the Task object

## User Pattern operations summary

The following table provides an overview of the operations provided for User Pattern objects.

Table 146. Core z Systems resources - User Pattern: operations summary

Operation name	HTTP method and URI path
"List User Patterns" on page 693	GET /api/console/user-patterns
"Get User Pattern Properties" on page 695	GET /api/console/user-patterns/{user-pattern-id}
"Update User Pattern Properties" on page 697	POST /api/console/user-patterns/{user-pattern-id}
"Create User Pattern" on page 699	POST /api/console/user-patterns
"Delete User Pattern" on page 701	DELETE /api/console/user-patterns/{user-pattern-id}

Table 147. Core z Systems resources - User Pattern: URI variables

URI variable	Description
{user-pattern-id}	Element ID of the User Pattern object

## Password Rule operations summary

The following table provides an overview of the operations provided for Password Rule objects.

Table 148. Core z Systems resources - Password Rule: operations summary

Operation name	HTTP method and URI path
"List Password Rules" on page 706	GET /api/console/password-rules
"Get Password Rule Properties" on page 708	GET /api/console/password-rules/{password-rule-id}

Table 148. Core z Systems resources - Password Rule: operations summary (continued)

Operation name	HTTP method and URI path
“Update Password Rule Properties” on page 710	POST /api/console/password-rules/{password-rule-id}
“Create Password Rule” on page 712	POST /api/console/password-rules
“Delete Password Rule” on page 714	DELETE /api/console/password-rules/{password-rule-id}

Table 149. Core z Systems resources - Password Rule: URI variables

URI variable	Description
{password-rule-id}	Element ID of the Password Rule object

## LDAP Server Definition operations summary

The following table provides an overview of the operations provided for LDAP Server Definition objects.

Table 150. Core z Systems resources - LDAP Server Definition: operations summary

Operation name	HTTP method and URI path
“List LDAP Server Definitions” on page 720	GET /api/console/ldap-server-definitions
“Get LDAP Server Definition Properties” on page 722	GET /api/console/ldap-server-definitions/{ldap-server-definition-id}
“Update LDAP Server Definition Properties” on page 724	POST /api/console/ldap-server-definitions/{ldap-server-definition-id}
“Create LDAP Server Definition” on page 725	POST /api/console/ldap-server-definitions
“Delete LDAP Server Definition” on page 728	DELETE /api/console/ldap-server-definitions/{ldap-server-definition-id}

Table 151. Core z Systems resources - LDAP Server Definition: URI variables

URI variable	Description
{ldap-server-definition-id}	Element ID of the LDAP Server Definition object

## Group operations summary

The following table provides an overview of the operations provided for Group objects.

Table 152. Core z Systems resources - Group: operations summary

Operation name	HTTP method and URI path
“List Custom Groups” on page 732	GET /api/groups
“Get Custom Group Properties” on page 734	GET /api/groups/{group-id}
“Create Custom Group” on page 735	POST /api/groups

| Table 152. Core z Systems resources - Group: operations summary (continued)

Operation name	HTTP method and URI path
“Delete Custom Group” on page 737	DELETE /api/groups/{group-id}
“List Custom Group Members” on page 741	GET /api/groups/{group-id}/members
“Add Member to Custom Group” on page 738	POST /api/groups/{group-id}/operations/add-member
“Remove Member from Custom Group” on page 740	POST /api/groups/{group-id}/operations/remove-member

| Table 153. Core z Systems resources - Groups: URI variables

URI variable	Description
{group-id}	Object ID of a Group object

## CPC operations summary

The following tables provide an overview of the operations provided for CPC objects.

| Table 154. Core z Systems resources - CPC: operations summary

Operation name	HTTP method and URI path
“List CPC Objects” on page 754	GET /api/cpcs
“List Ensemble CPC Objects” on page 756	GET /api/ensembles/{ensemble-id}/cpcs
“Get CPC Properties” on page 758	GET /api/cpcs/{cpc-id}
“Update CPC Properties” on page 764	POST /api/cpcs/{cpc-id}
“Activate CPC” on page 766	POST /api/cpcs/{cpc-id}/operations/activate
“Deactivate CPC” on page 768	POST /api/cpcs/{cpc-id}/operations/deactivate
“Import Profiles” on page 770	POST /api/cpcs/{cpc-id}/operations/import-profiles
“Export Profiles” on page 771	POST /api/cpcs/{cpc-id}/operations/export-profiles
“Add Temporary Capacity” on page 772	POST /api/cpcs/{cpc-id}/operations/add-temporary-capacity
“Remove Temporary Capacity” on page 774	POST /api/cpcs/{cpc-id}/operations/remove-temporary-capacity
“Swap Current Time Server” on page 776	POST /api/cpcs/{cpc-id}/operations/swap-cts
“Set STP Configuration” on page 777	POST /api/cpcs/{cpc-id}/operations/set-stp-config
“Change STP-only Coordinated Timing Network” on page 778	POST /api/cpcs/{cpc-id}/operations/change-stponly-ctn

| Table 154. Core z Systems resources - CPC: operations summary (continued)

Operation name	HTTP method and URI path
“Join STP-only Coordinated Timing Network” on page 780	POST /api/cpcs/{cpc-id}/operations/join-stponly-ctn
“Leave STP-only Coordinated Timing Network” on page 781	POST /api/cpcs/{cpc-id}/operations/leave-stponly-ctn
“Get CPC Audit Log” on page 782	GET /api/cpcs/{cpc-id}/operations/get-audit-log
“Get CPC Security Log” on page 784	GET /api/cpcs/{cpc-id}/operations/get-security-log
“List CPC Hardware Messages” on page 787	GET /api/cpcs/{cpc-id}/hardware-messages
“Get CPC Hardware Message Properties” on page 789	GET /api/cpcs/{cpc-id}/hardware-messages/{hardware-message-id}
“Delete CPC Hardware Message” on page 791	DELETE /api/cpcs/{cpc-id}/hardware-messages/{hardware-message-id}

| Table 155. Core z Systems resources - CPC: URI variables

URI variable	Description
{cpc-id}	Object ID of a CPC object
{ensemble-id}	Object ID of an ensemble object
{hardware-message-id}	Element ID of the hardware message object

## Logical partitions operation summary

The following tables provide an overview of the operations provided for Logical Partition objects.

| Table 156. Core z Systems resources - Logical partitions: operations summary

Operation name	HTTP method and URI path
“List Logical Partitions of CPC” on page 808	GET /api/cpcs/{cpc-id}/logical-partitions
“Get Logical Partition Properties” on page 810	GET /api/logical-partitions/{logical-partition-id}
“Update Logical Partition Properties” on page 813	POST /api/logical-partitions/{logical-partition-id}
“Activate Logical Partition” on page 814	POST /api/logical-partitions/{logical-partition-id}/operations/activate
“Deactivate Logical Partition” on page 816	POST /api/logical-partitions/{logical-partition-id}/operations/deactivate
“Reset Normal” on page 818	POST /api/logical-partitions/{logical-partition-id}/operations/reset-normal
“Reset Clear” on page 820	POST /api/logical-partitions/{logical-partition-id}/operations/reset-clear
“Load Logical Partition” on page 822	POST /api/logical-partitions/{logical-partition-id}/operations/load
“PSW Restart” on page 824	POST /api/logical-partitions/{logical-partition-id}/operations/psw-restart

| Table 156. Core z Systems resources - Logical partitions: operations summary (continued)

Operation name	HTTP method and URI path
“Start Logical Partition” on page 825	POST /api/logical-partitions/{ <i>logical-partition-id</i> }/operations/start
“Stop Logical Partition” on page 827	POST /api/logical-partitions/{ <i>logical-partition-id</i> }/operations/stop
“SCSI Load” on page 828	POST /api/logical-partitions/{ <i>logical-partition-id</i> }/operations/scsi-load
“SCSI Dump” on page 830	POST /api/logical-partitions/{ <i>logical-partition-id</i> }/operations/scsi-dump

| Table 157. Core z Systems resources - Logical partitions: URI variables

URI variable	Description
{ <i>cpc-id</i> }	Object ID of a CPC object
{ <i>logical-partition-id</i> }	Object ID of a Logical Partition object

## Activation profile operations summary

The following tables provide an overview of the operations provided for the various types of Activation Profile objects.

| Table 158. Core z Systems resources - Reset activation profile: operations summary

Operation name	HTTP method and URI path
“List Reset Activation Profiles” on page 834	GET /api/cpcs/{ <i>cpc-id</i> }/reset-activation-profiles
“Get Reset Activation Profile Properties” on page 836	GET /api/cpcs/{ <i>cpc-id</i> }/reset-activation-profiles/{ <i>reset-activation-profile-name</i> }
“Update Reset Activation Profile Properties” on page 837	POST /api/cpcs/{ <i>cpc-id</i> }/reset-activation-profiles/{ <i>reset-activation-profile-name</i> }

| Table 159. Core z Systems resources - Image activation profile: operations summary

Operation name	HTTP method and URI path
“List Image Activation Profiles” on page 854	GET /api/cpcs/{ <i>cpc-id</i> }/image-activation-profiles
“Get Image Activation Profile Properties” on page 856	GET /api/cpcs/{ <i>cpc-id</i> }/image-activation-profiles/{ <i>image-activation-profile-name</i> }
“Update Image Activation Profile Properties” on page 859	POST /api/cpcs/{ <i>cpc-id</i> }/image-activation-profiles/{ <i>image-activation-profile-name</i> }

| Table 160. Core z Systems resources - Load activation profile: operations summary

Operation name	HTTP method and URI path
“List Load Activation Profiles” on page 863	GET /api/cpcs/{ <i>cpc-id</i> }/load-activation-profiles
“Get Load Activation Profile Properties” on page 865	GET /api/cpcs/{ <i>cpc-id</i> }/load-activation-profiles/{ <i>load-activation-profile-name</i> }

| Table 160. Core z Systems resources - Load activation profile: operations summary (continued)

Operation name	HTTP method and URI path
“Update Load Activation Profile Properties” on page 867	POST /api/cpcs/{cpc-id}/load-activation-profiles/{load-activation-profile-name}

| Table 161. Core z Systems resources - Group profile: operations summary

Operation name	HTTP method and URI path
“List Group Profiles” on page 869	GET /api/cpcs/{cpc-id}/group-profiles
“Get Group Profile Properties” on page 871	GET /api/cpcs/{cpc-id}/group-profiles/{group-profile-name}
“Update Group Profile Properties” on page 873	POST /api/cpcs/{cpc-id}/group-profiles/{group-profile-name}

| Table 162. Core z Systems resources - Activation profile: URI variables

URI variable	Description
{cpc-id}	Object ID of a CPC object
{group-profile-name}	Group profile name
{image-activation-profile-name}	Image activation profile name
{load-activation-profile-name}	Load activation profile name
{reset-activation-profile-name}	Reset activation profile name

## Capacity record operations summary

The following tables provide an overview of the operations provided for Capacity Record objects.

| Table 163. Core z Systems resources - Capacity record: operations summary

Operation name	HTTP method and URI path
“List Capacity Records” on page 877	GET /api/cpcs/{cpc-id}/capacity-records
“Get Capacity Record Properties” on page 878	GET /api/cpcs/{cpc-id}/capacity-records/{capacity-record-id}

| Table 164. Core z Systems resources - Capacity record: URI variables

URI variable	Description
{cpc-id}	Object ID of a CPC object
{capacity-record-id}	Capacity record identifier

## Shared nested objects

Some of the Core API objects share common nested objects and are documented here for ease of reference.

Table 165. *ec-mcl-description object*

Field name	Type	Description
<b>actions</b>	Array of action objects	An optional array of pending action objects. This field is only provided when the HMC is communicating with the CPC's SE.
<b>ec</b>	Array of ec objects	An optional array of EC objects. This field is only provided when the HMC is communicating with the CPC's SE.

Table 166. *action object*

Field name	Type	Description
<b>type</b>	String Enum	One of: <ul style="list-style-type: none"> <li>"channel-config" - channels pending a config on/off</li> <li>"coupling-facility-reactivation" - at least one coupling facility pending reactivation</li> <li>"power-on-reset-tracking" - there is a need for a power-on-reset</li> <li>"zhybrid-blades-activation" - (zHybrid accelerator blades are pending an activation.</li> </ul>
<b>activation</b>	String Enum	One of: <ul style="list-style-type: none"> <li>"current" - the action is for the current activation</li> <li>"next" - the action is for the next install and activation.</li> </ul>
<b>pending</b>	Boolean	Is the action pending (true) or not pending (false)

Table 167. *ec object*

Field name	Type	Description
<b>number</b>	String (1-6)	Engineering Change stream identifier.
<b>part-number</b>	String (1-8)	Engineering Change stream part number.
<b>type</b>	String (1-32)	Engineering Change stream name.
<b>description</b>	String (1-65)	Engineering Change stream descriptive text.
<b>mcl</b>	Array of mcl objects	The list of MicroCode Levels for this Engineering Change.

Table 168. *mcl object*

Field name	Type	Description
<b>type</b>	String Enum	One of: <ul style="list-style-type: none"> <li>"retrieved" - a retrieved or staged level</li> <li>"activated" - an activated or applied level</li> <li>"accepted" - a committed level</li> <li>"installable-concurrent" - a non-disruptive apply-able level</li> <li>"removable-concurrent" - a non-disruptive reject-able level.</li> </ul>
<b>level</b>	String (1-3)	Microcode level.
<b>last-update</b>	Timestamp	Time stamp of the last update, in the number of milliseconds since midnight January 1, 1970 UTC. A null object is returned if no updates have occurred.

Table 169. *stp-config object*

Field name	Type	Description
<b>stp-id</b>	String (1-8)	If in STP-only or Mixed CTN, the STP identifier. Otherwise, an empty string. Valid characters are 0-9, a-z, A-Z, underscore(_) and dash(-).
<b>etr-id</b>	Integer (0-31)	ETR Identifier, if in ETR mode. If not in ETR mode, a null object is returned.
<b>preferred-time-server</b>	stp-node-object	Describes the Preferred Timer Server. This property is optional, only returned on a Get request when the information is set.
<b>backup-time-server</b>	stp-node-object	Describes the Backup Timer Server. This property is optional, only returned on a Get request when the information is set.
<b>arbiter</b>	stp-node-object	Describes the arbiter of the CTN. This property is optional, only returned on a Get request when the information is set.
<b>current-time-server</b>	String Enum	Describes the CPC's role in the CTN. One of: <ul style="list-style-type: none"> <li>• <b>"preferred"</b> - CPC is the Preferred Time Server</li> <li>• <b>"backup"</b> - CPC is the Backup Time Server.</li> </ul>

This object is used to identify a CPC to the STP services. When used as input on the Set STP Configuration operation, if the **object-uri** field is not provided, all other fields are required. If the **object-uri** field is provided, all other fields are optional. If all fields are provided, the **object-uri** field is ignored.

Table 170. *stp-node object*

Field name	Type	Description
<b>object-uri</b>	String URI	If the CPC is known to the HMC, contains the CPC's object-uri. Otherwise, contains a null object
<b>type</b>	String (0-6)	The CPC machine type, right justified and left padded with zeros or a empty string
<b>model</b>	String (0-3)	The CPC machine model or a empty string
<b>manuf</b>	String (0-3)	The CPC manufacturer or a empty string
<b>po-manuf</b>	String (0-2)	The CPC plant of manufacturer or a empty string
<b>seq-num</b>	String (0-12)	The CPC sequence number or a empty string

Table 171. *psw-description object*

Field name	Type	Description
<b>psw</b>	String	Program Status Word (PSW) information for a single processor.
<b>cpid</b>	String (2)	The hexadecimal processor identifier, right justified and left padded with zeros.

Table 172. *zaware-network object*

Field name	Type	Description
<b>chpid</b>	String (2)	The required network adapter channel path identifier, in hexadecimal characters 0-9,a-f,A-F. The format is two hexadecimal digits (00-FF).
<b>ipaddr-type</b>	String Enum	Indicates how this network adapter's IP address is obtained. One of the following values: <ul style="list-style-type: none"> <li>• <b>"dhcp"</b> - obtains an IP address via DHCP</li> <li>• <b>"link-local"</b> - obtains a link-local IP address</li> <li>• <b>"static"</b> - uses the specified IP address information.</li> </ul>



Table 172. *zaware-network object (continued)*

Field name	Type	Description
<b>vlan-id</b>	Integer (0-65535)	If this network adapter is attached to a Virtual LAN, this field contains the VLAN identifier. Otherwise, a null object indicating that this network adapter is not attached to a Virtual LAN.
<b>static-ip-info</b>	network- ip-info object	When <b>ipaddr-type</b> is "static", contains the static IP information. Otherwise, contains a null object indicating that a static IP address is not used.

Table 173. *ip-info object*

Field name	Type	Description
<b>type</b>	String Enum	The type of IP address being provided. One of the following values: <ul style="list-style-type: none"> <li>• "ipv4" - an IPv4 address is provided</li> <li>• "ipv6" - an IPv6 address is provided.</li> </ul>
<b>ip-address</b>	String/ IPV4 address or String/ IPV6 address	The IP address to be used. The format of the string (IPv4 or IPv6) must be as indicated by the <b>type</b> field.

Table 174. *network-ip-info object*

Field name	Type	Description
<b>type</b>	String Enum	The type of IP address being provided. One of the following values: <ul style="list-style-type: none"> <li>• "ipv4" - an IPv4 address is provided</li> <li>• "ipv6" - an IPv6 address is provided.</li> </ul>
<b>ip-address</b>	String/ IPV4 address or String/ IPV6 address	The IP address to be used. The format of the string (IPv4 or IPv6) must be as indicated by the <b>type</b> field.
<b>prefix</b>	Integer	The number of leading bits of ip-address that represent the network prefix. <ul style="list-style-type: none"> <li>• When <b>type</b> is "ipv4" - valid values are 0-32</li> <li>• When <b>type</b> is "ipv6" - valid values are 0-128.</li> </ul>

Table 175. *absolute-capping object*

Field name	Type	Description
<b>type</b>	String Enum	The type of absolute capping. One of the following values: <ul style="list-style-type: none"> <li>• "none" - no absolute capping</li> <li>• "processors" - processor type absolute capping.</li> </ul>
<b>value</b>	Float	When <b>type</b> is "none", value is not specified. When <b>type</b> is "processors", value is in the range .01...255.00 in increments of .01 and is the limit of usage independent of priority.

## Console object

- | The Console object represents the single z Systems Hardware Management Console (HMC) application. The Console object offers a heterogeneous set of services and capabilities, from basic HMC control operations to general HMC information.

Object access to the single Console object representing the local HMC is automatic for all authenticated users. A Console may (but need not) participate in a { Primary, Alternate } pairing with another Console. The Console object helps facilitate the management of a HMC's role in such a pairing.

## Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, but does not provide the operational-status-related properties defined in that schema because it does not maintain the concept of an operational status.

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

The following class-specific specializations apply to the other base managed object properties:

Table 176. Console object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/URI	The canonical URI path of the Console object, of the form <code>/api/console</code>
<b>parent</b>	—	String/URI	A Console object has no parent, so this property is always a null object.
<b>class</b>	—	String	The class of a Console object is " <b>console</b> ".
<b>name</b>	(ro)	String	The installation assigned name
<b>description</b>	(ro)	String	This property is not supported, and returned as null.

## Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties:

Table 177. Console object: class specific additional properties

Name	Qualifier	Type	Description
<b>version</b>	—	String (1-8)	The version number for the Console object.
<b>paired-role</b>	—	String Enum	An indication (" <b>primary</b> " or " <b>alternate</b> ") of how the Console object functions in a pairing, or null if this Console object is not paired.
<b>ec-mcl-description</b>	—	ec-mcl-description object	A nested object that describes the EC (Engineering Change) and MCL (Microcode Level) for the Console. Refer to the description of the ec-mcl-description object for details.
<b>is-auto-switch-enabled</b>	—	Boolean	Automatic switching between primary and alternate Hardware Management Consoles is enabled (true), or is not enabled (false).
<b>network-info</b>	—	network-info object	A nested object describing the network information for this Hardware Management Console and for any other Hardware Management console with which this Hardware Management Console may be paired.
<b>machine-info</b>	—	machine-info object	A nested object describing the machine's BIOS characteristics.
<b>ip-swapping-available</b>	—	Boolean	Whether the current primary HMC IP addresses will be moved when making the alternate HMC the new primary (true) or not (false).

Table 177. Console object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>hardware-messages</b>	(c)(pc)	Array of hardware-message objects	<p>The complete list of all Console hardware messages, each identified by its URI. This list corresponds to the list provided by the <b>List Console Hardware Messages</b> operation. If the console has no hardware messages, then an empty array is provided.</p> <p>The list of returned hardware messages can change as the result of the new messages being dynamically added or removed by the infrastructure or due to hardware messages being deleted via the <b>Delete Console Hardware Message</b> operation.</p> <p><b>Note:</b> This property is not returned by the <b>Get Console Properties</b> operation, and only sessions associated with an HMC user with permission to the Hardware Messages task will receive a property-change notification for this property.</p>

Table 178. network-info object properties

Name	Type	Description
<b>this-hmc</b>	Array of detailed-network-info objects	The collection of network information for the local Hardware Management Console. The number of objects returned is a function of the machine model and type on which the Hardware Management Console is executing. This information is available in the <b>machine-info</b> property.
<b>paired-hmc</b>	paired-ip-info object	Describes the IP information for the paired Hardware Management Console known to the local Hardware Management Console

Table 179. detailed-network-info properties

Name	Type	Description
<b>hmc-name</b>	String (1-16)	The Hardware Management Console name
<b>interface-name</b>	String	The network interface name
<b>domain-name</b>	String (1-255)	The domain name configured for this network interface
<b>is-private</b>	Boolean	Whether the interface is private (true) or public (false).
<b>mac</b>	String (1-12)	The MAC address of this network interface.
<b>ipv4-address</b>	Array of ipv4-info objects	A collection of nested objects which describe the IPv4 addresses for this network interface.
<b>ipv6-address</b>	Array of ipv6-info objects	A collection of nested objects which describe the IPv6 addresses for this network interface.

Table 180. detailed-network-info properties

Name	Type	Description
<b>hmc-name</b>	String (1-16)	The Hardware Management Console name
<b>interface-name</b>	String	The network interface name
<b>domain-name</b>	String (1-255)	The domain name configured for this network interface
<b>is-private</b>	Boolean	Whether the interface is private (true) or public (false).
<b>mac</b>	String (1-12)	The MAC address of this network interface.

Table 180. detailed-network-info properties (continued)

Name	Type	Description
<b>ipv4-address</b>	Array of ipv4-info objects	A collection of nested objects which describe the IPv4 addresses for this network interface.
<b>ipv6-address</b>	Array of ipv6-info objects	A collection of nested objects which describe the IPv6 addresses for this network interface.

Table 181. ipv4-info properties

Name	Qualifier	Type	Description
<b>subnet-mask</b>	(pc)	String (1-15)	The IP mask value
<b>ip-address</b>	(pc)	String IPV4 address	The IPv4 address

Table 182. ipv6-info properties

Name	Qualifier	Type	Description
<b>prefix-length</b>	(pc)	Integer	The number of leading bits of the IPv6 address that represent the network prefix.
<b>ip-address</b>	(pc)	String IPV6 address	The IPv6 address

Table 183. machine-info properties

Name	Type	Description
<b>machine-type</b>	String (1-4)	The type of machine on which the Hardware Management Console is executing.
<b>machine-model</b>	String (1-3)	The model of machine on which the Hardware Management Console is executing.
<b>machine-serial</b>	String (1-10)	The serial number of machine on which the Hardware Management Console is executing.

Table 184. hardware-message object properties

Name	Type	Description
<b>element-uri</b>	String/URI	The canonical URI path of the Console hardware message. The URI is in the following form: /api/console/hardware-messages/{ <i>hardware-message-id</i> }, where { <i>hardware-message-id</i> } is the value of the <b>element-id</b> property of the hardware message.
<b>element-id</b>	String (36)	The unique identifier for the hardware message. The <b>element-id</b> is in the form of a UUID.
<b>parent</b>	String/URI	The parent of a console hardware message is the Console object. The <b>parent</b> value is the canonical URI path for the console.
<b>class</b>	String	The <b>class</b> of a hardware message object is " <b>hardware-message</b> ".
<b>timestamp</b>	Timestamp	The <b>timestamp</b> represents the date and time when the hardware message was created.
<b>text</b>	String	The text of the hardware message.

## Get Console Properties

The **Get Console Properties** operation retrieves the properties of the Console object.

### HTTP method and URI

**GET** /api/console

### Response body contents

On successful completion, the response body contains an object that provides the current values of the properties for the Console object as defined in “Data model” on page 618. Field names and data types in the object are the same as the property names and data types defined in the data model.

### Description

This operation returns the current properties for the Console object.

This operation may be targeted to an alternate Hardware Management Console.

On successful execution, HTTP status code 200 (OK) is returned and all of the current properties as defined by the Data Model for the Console object are provided in the response body.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/console HTTP/1.1
x-api-session: 6djngfmjus22whw3b5oj670c09rb2izzaafr4t3i1iw60ujmkd
```

---

*Figure 291. Get Console Properties: Request*

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 16:04:31 GMT
content-type: application/json;charset=UTF-8
content-length: 4250
{
  "class": "console",
  "description": "R32 Primary/Alternate HMC",
  "ec-mcl-description": {
    "ec": [
      {
        "description": "Hardware Management Console Framework",
        "mcl": [
          {
            "last-update": "2011-11-09T15:06:34Z",
            "level": "205",
            "type": "retrieved"
          },
          {
            "last-update": "2011-11-09T15:09:07Z",
            "level": "205",
            "type": "activated"
          },
          {
            "last-update": "2011-11-09T15:05:51Z",
            "level": "167",
            "type": "accepted"
          },
          {
            "level": "205",
            "type": "installable-concurrent"
          },
          {
            "level": "168",
            "type": "removable-concurrent"
          }
        ]
      },
      {
        "number": "N48180",
        "part-number": "45D8928",
        "type": "SYSTEM"
      }
    ],
  }
},
```

---

Figure 292. Get Console Properties: Response (Part 1)

```

    {
      "description": "Embedded Operating System",
      "mcl": [
        {
          "level": "000",
          "type": "retrieved"
        },
        {
          "level": "000",
          "type": "activated"
        },
        {
          "level": "000",
          "type": "accepted"
        },
        {
          "level": "000",
          "type": "installable-concurrent"
        },
        {
          "level": "000",
          "type": "removable-concurrent"
        }
      ],
      "number": "N48198",
      "part-number": "45D8929",
      "type": "OS"
    }
  ]
},
"ip-swapping-available": true,
"is-auto-switch-enabled": true,
"is-locked": false,
"machine-info": {
  "machine-model": "PAA",
  "machine-serial": "KQZBLKD",
  "machine-type": "7327"
},
"name": "HMCR32PRI",
"network-info": {
  "paired-hmc": {
    "hmc-name": "HMCR32ALT",
    "ipv4-address": [
      "9.60.15.47",
      "9.60.14.47"
    ],
    "ipv6-address": [
      "fdd8:673b:d89b:1:221:5eff:fe69:e3f5",
      "2002:93c:ffb:1:221:5eff:fe69:e3f5",
      "fe80:0:0:0:221:5eff:fe69:e3f5",
      "fe80:0:0:0:210:18ff:fe4c:8026"
    ]
  }
},
},

```

Figure 293. Get Console Properties: Response (Part 2)

---

```

"this-hmc": [
  {
    "domain-name": "endicott.ibm.com",
    "hmc-name": "HMCR32PRI",
    "interface-name": "eth0",
    "ipv4-address": [
      {
        "ip-address": "9.60.15.48",
        "subnet-mask": "255.255.255.0"
      }
    ],
    "ipv6-address": [
      {
        "ip-address": "fdd8:673b:d89b:1:221:5eff:fe69:dea0",
        "prefix-length": 64
      },
      {
        "ip-address": "2002:93c:ffb:1:221:5eff:fe69:dea0",
        "prefix-length": 64
      },
      {
        "ip-address": "fe80:0:0:0:221:5eff:fe69:dea0",
        "prefix-length": 64
      }
    ],
    "is-private": false,
    "mac": "00215E69DEA0"
  },
  {
    "domain-name": "endicott.ibm.com",
    "hmc-name": "HMCR32PRI",
    "interface-name": "eth1",
    "ipv4-address": [
      {
        "ip-address": "9.60.14.48",
        "subnet-mask": "255.255.255.0"
      }
    ],
    "ipv6-address": [
      {
        "ip-address": "fe80:0:0:0:210:18ff:fe4c:8334",
        "prefix-length": 64
      }
    ],
    "is-private": false,
    "mac": "0010184C8334"
  }
],
"object-id": "1e8b3137-85b0-3a06-9269-0b25fc170d44",
"object-uri": "/api/console",
"paired-role": "primary",
"parent": null,
"version": "2.11.1"
}

```

---

Figure 294. Get Console Properties: Response (Part 3)

## Restart Console

The **Restart Console** operation restarts the Hardware Management Console.



## HTTP method and URI

POST /api/console/operations/restart

## Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
force	Boolean	Optional	Whether the restart operation is processed when users are connected (true) or not (false). The default is false.

## Description

The Hardware Management Console is restarted. This operation may be targeted to an alternate Hardware Management Console.

By default, the restart does not occur if one or more users are currently connected to the Hardware Management Console. This can be overridden by use of the **force** field in the request body.

On success, HTTP status code 202 (Accepted) is returned.

## Authorization

To use **Restart Console**, you must have the following:

- Action/Task permission to the **Shutdown/Restart** task
- Remote Restart must be enabled on the Hardware Management Console.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	267	The operation is rejected, due to the presence of HMC users. Either wait until all HMC users have logged off or retry the request with the <b>force</b> field set to <b>"true"</b> .
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform this operation.
	269	This operation is currently blocked. The error message will contain information on the blocking application..
	270	The remote restart operation is not enabled on the HMC.
500 (Server Error)	273	An unexpected error occurred during the operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Make Console Primary

**Make Console Primary** initiates a primary/alternate role switch when directed at a Hardware Management Console that is currently operating in the alternate role. It is only supported for a Hardware Management Console participating in a {primary, alternate} pairing.

### HTTP method and URI

POST /api/console/operations/make-primary

### Description

This operation initiates a primary/alternate role switch when directed at a Hardware Management Console that is currently operating in the alternate role. As this command will cause both consoles to be rebooted, any active sessions (GUI based or API based) will be terminated.

If this operation is directed at a Hardware Management Console which does not participate in a {primary, alternate} pairing, it returns with HTTP status code 400 (Bad Request).

If this operation is directed at a Hardware Management Console whose **pair-role** is **"alternate"**, it initiates the role-switch process, and returns with HTTP status code 204 (No Content).

If this operation is directed at a Hardware Management Console whose **pair-role** is already **"primary"**, it has no effect, and returns with HTTP status code 204 (No Content).

### Authorization

To use **Make Console Primary**, you must have the following:

- Action/task permission to the **Manage Alternate HMC** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	265	The operation was directed at a Hardware Management Console which is not participating in a {primary, alternate} pairing, which is not supported.
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform this operation.
500 (Server Error)	273	An unexpected error occurred during the operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

### Usage notes

When configured as recommended by IBM, the process of recovering from the failure of the primary HMC by takeover by the alternate HMC includes movement of the IP address of the former primary HMC to the new primary HMC. When this occurs, explicit redirection of API requests to the newly

designated primary HMC is not needed. However, the IP address swapping may not be possible in certain network configurations. The address of the alternate HMC is provided to allow applications to explicitly redirect requests to the other HMC of the pair in these cases.

## Shutdown Console

**Shutdown Console** powers off the Hardware Management Console.

### HTTP method and URI

POST /api/console/operations/shutdown

### Request body contents

A request body must be specified. It has the following fields:

Field name	Type	Rqd/Opt	Description
force	Boolean	Optional	Whether the shutdown operation is processed when users are connected (true) or not (false). The default is false.

### Description

The Hardware Management Console is powered off.

This operation may be targeted to an alternate Hardware Management Console.

By default, the shutdown does not occur if one or more users are currently connected to the Hardware Management Console. This can be overridden by use of the force field in the request body.

The action to shutdown the Hardware Management Console occurs asynchronously. If the request is accepted, HTTP status code 202 (Accepted) is returned to indicate that the request has been initiated. However, because this action results in the targeted Hardware Management Console becoming inactive and powered off at completion, it is not possible to track the completion of this request. Thus no response body containing an asynchronous job URI is provided, nor is a job completion notification generated upon completion.

### Authorization

To use **Shutdown Console**, you must have the following:

- Action/task permission to the **Shutdown/Restart** task.
- Remote Shutdown must be enabled on the Hardware Management Console.

### HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned but no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	267	The operation is rejected, due to the presence of HMC users. Either wait until all HMC users have logged off or retry the request with the force field set to true.
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform this operation.
	270	The remote restart operation is not enabled on the HMC.
	304	This operation is currently blocked. The error message will contain information on the blocking application.
500 (Server Error)	273	An unexpected error occurred during the operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/console/operations/shutdown HTTP/1.1/
x-api-session: 5dul8zvlwa5s83eobcukaf1vug3s3kgidkyk9e5c5acsekabs1
content-type: application/json
content-length: 16
{
  "force": false
}
```

---

Figure 295. Shutdown Console: Request

---

```
202 Accepted
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Fri, 01 Mar 2013 19:38:25 GMT

<No response body>
```

---

Figure 296. Shutdown Console: Response

## | Reorder User Patterns

| The **Reorder User Patterns** operation changes the search order of the console's User Patterns used when a user logs on to the console.

### | HTTP method and URI

| **POST /api/console/operations/reorder-user-patterns**

## Request body contents

The request body is a JSON object with the following fields:

Name	Type	Rqd/Opt	Description
user-pattern-uris	Array of String/URI	Required	Ordered list of User Pattern object <b>element-uri</b> property values. The order of these URIs in the array defines the new order of the User Patterns.

## Description

This operation reorders the console's User Patterns.

On successful execution of this operation the User Patterns are reordered to match the order of their **element-uri** properties in the **user-pattern-uris** array in the request body.

The request body is validated against the schema described in “Request body contents.” If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. If a URI in the request body does not designate an existing User Pattern object, status code 404 (Not Found) is returned. The array in the request body must include each of the console's currently defined User Patterns and no others. In addition, the API user must have action/task permission to the Manage User Patterns task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirement

This operation has the following authorization requirements:

- Action/task permission to the **Manage User Patterns** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 185. Reorder User Patterns: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	7	The array in the request body is missing an existing User Pattern, or it contains an entry that designates a User Pattern that is not one of the console's current User Patterns.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	2	A URI in the request body does not designate an existing resource of the correct type.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
POST /api/console/operations/reorder-user-patterns HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 301
{
  "user-pattern-uris":[
    "/api/console/user-patterns/497bf4ec-1dbf-11e4-8ceb-1c6f65065a91",
    "/api/console/user-patterns/6d897292-3ceb-11e4-9e36-1c6f65065a91",
    "/api/console/user-patterns/e40b9ba6-48e0-11e4-82a1-1c6f65065a91",
    "/api/console/user-patterns/ec5b012a-4a7a-11e4-8777-1c6f65065a91"
  ]
}
```

Figure 297. Reorder User Patterns: Request

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT

<No response body>
```

Figure 298. Reorder User Patterns: Response

## Get Console Audit Log

The **Get Console Audit Log** operation returns the console audit log, filtered according to the query parameters, if specified.

### HTTP method and URI

**GET** /api/console/operations/get-audit-log

### Query Parameters

Name	Type	Rqd/Opt	Description
begin-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created earlier than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.
end-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created later than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.

## Response body contents

On successful completion, the response body is a JSON array of JSON objects returned using HTTP chunked transfer encoding. Each array element is a log-entry-info object containing information about a single log entry. The array elements are in order of increasing timestamp. See Table 186 on page 631 for more information.

A log-entry-info object contains information about a single log entry and the event which caused the entry. Each log-entry-info object contains the following fields:

Table 186. log-entry-info object properties

Name	Type	Description
<b>event-time</b>	Timestamp	The time when the event occurred.
<b>event-id</b>	String	The ID number for the event.
<b>event-name</b>	String (1-12)	The name for the event.
<b>userid</b>	String	The user ID of the HMC/SE user associated with the event, or null if there is no user associated with the event.
<b>user-uri</b>	String/URI	The canonical URI path of the HMC/SE user associated with the event, or null if there is no user associated with the event.
<b>event-message</b>	String	The complete, formatted message for the event.
<b>event-data-items</b>	Array of objects	An array of event-data-item-info objects, one for each item of event data associated with the <b>event-message</b> for the event. If there are no event data items, an empty array is provided. The order of items in this array is semantically significant and matches the documentation for this <b>event-id</b> .
<b>event-details</b>	Array of objects	An array of event-details-info objects, one for each event detail associated with the event. If there are no event details, an empty array is provided.

An event-details-info object contains information about a single event detail. Each event-details-info object contains the following fields:

Table 187. event-details-info object properties

Name	Type	Description
<b>event-details-message</b>	String	The complete, formatted details message.
<b>event-details-data-items</b>	Array of objects	An array of event-data-item-info objects, one for each item of event details data associated with the <b>event-details-message</b> for the event. If there are no event details data items, an empty array is provided.

An event-data-item-info object contains information about a single item of event data. Each event-data-item-info object contains the following fields:

Table 188. event-data-item-info object properties

Name	Type	Description
<b>data-item-number</b>	Integer	The number for this data item. This is the 0-based index of this event-data-item-info object in the event-data-items or event-details-data-items array in which it is contained. This number identifies the substitution variable to which this data item corresponds in the documentation for the event identified by <b>event-id</b> .
<b>data-item-type</b>	String Enum	Identifies the data type of the data item in the event-data-item field. Possible values are: <ul style="list-style-type: none"> <li>• "long"</li> <li>• "float"</li> <li>• "string"</li> </ul>
<b>data-item-value</b>	Varies. See <b>data-item-type</b> description	The data item.

## Description

This operation returns the console's audit log in increasing timestamp order, filtered according to the query parameters, if specified. Each log entry pertains to a specific event that occurred on or to a managed object or the console itself. The log entries can be limited by specifying explicit filtering criteria on the request. If the begin-time query parameter is specified, then any entries earlier than that time are omitted. If the end-time query parameter is specified, then any entries later than that time are omitted.

The API user must have action/task permission to the **Audit and Log Management** task; otherwise, status code 403 (Forbidden) is returned.

On successful execution, the response body contains an array of filtered log entries. If the audit log is empty or there are no entries to be returned after filtering, then an empty array is provided. Each log entry contains the event ID, event name and event message. If there are data items included in the event message, they are available separately. The order and meaning of the substitution items for each event ID are documented in the IBM Knowledge Center, at <http://www.ibm.com/support/knowledgecenter/> (Select **z Systems** on the navigation bar, and then select your server). The information can be found in the **HMC Introduction** topic **Audit, Event, and Security Log Messages**.

## Authorization requirement

This operation has the following authorization requirements:

- Action/task permission to the **Audit and Log Management** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described "Response body contents" on page 630.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Example HTTP interaction

```
GET /api/console/operations/get-audit-log HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

Figure 299. Get Console Audit Log: Request



```

200 OK
server: zSeries management console API web server / 2.0
transfer-encoding: chunked
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:34 GMT
content-type: application/json;charset=ISO-8859-1
[
  {
    "event-data-items":[
      {
        "data-item-number":0,
        "data-item-type":"string",
        "data-item-value":"GinkgoEnsemble"
      }
    ],
    "event-details":[],
    "event-id":"1988",
    "event-message":"An ensemble was created. The ensemble name is GinkgoEnsemble.",
    "event-name":"EnsembleMem",
    "event-time":1411405571970,
    "user-uri":"/api/users/ae72e844-3dc0-11e4-8dd1-1c6f65065a91",
    "userid":"ENSADMIN"
  },
  {
    "event-data-items":[
      {
        "data-item-number":0,
        "data-item-type":"string",
        "data-item-value":"Default"
      },
      {
        "data-item-number":1,
        "data-item-type":"string",
        "data-item-value":"Default"
      },
      {
        "data-item-number":2,
        "data-item-type":"string",
        "data-item-value":"OPERATOR"
      }
    ],
    "event-details":[],
    "event-id":"1982",
    "event-message":"Availability policy \"Default\" was activated in workload
    \"Default\" by user OPERATOR as a result of creating the workload.",
    "event-name":"AvlPolActive",
    "event-time":1411405588850,
    "user-uri":null,
    "userid":null
  }
]

```

Figure 300. Get Console Audit Log: Response

## Get Console Security Log

The **Get Console Security Log** operation returns the console security log, filtered according to the query parameters, if specified.

### HTTP method and URI

**GET** /api/console/operations/get-security-log

### Query Parameters

Name	Type	Rqd/Opt	Description
begin-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created earlier than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.
end-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created later than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.

## Response body contents

On successful completion, the response body is a JSON array of JSON objects returned using HTTP chunked transfer encoding. Each array element is a log-entry-info object containing information about a single log entry. The array elements are in order of increasing timestamp. See Table 186 on page 631 for more information.

## Description

This operation returns the console's security log in increasing timestamp order, filtered according to the query parameters, if specified. Each log entry pertains to a specific event that occurred on or to a managed object or the console itself. The log entries can be limited by specifying explicit filtering criteria on the request. If the begin-time query parameter is specified, then any entries earlier than that time are omitted. If the end-time query parameter is specified, then any entries later than that time are omitted.

The API user must have action/task permission to the **View Security Logs** task; otherwise, status code 403 (Forbidden) is returned.

On successful execution, the response body contains an array of filtered log entries. If the security log is empty or there are no entries to be returned after filtering, then an empty array is provided. Each log entry contains the event ID, event name and event message. If there are data items included in the event message, they are available separately. The order and meaning of the substitution items for each event ID are documented in the IBM Knowledge Center, at <http://www.ibm.com/support/knowledgecenter/> (Select **z Systems** on the navigation bar, and then select your server). The information can be found in the **HMC Introduction** topic **Audit, Event, and Security Log Messages**.

## Authorization requirement

This operation has the following authorization requirement:

- Action/task permission to the **View Security Logs** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/console/operations/get-security-log HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

*Figure 301. Get Console Security Log: Request*

```
200 OK
server: zSeries management console API web server / 2.0
transfer-encoding: chunked
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:34 GMT
content-type: application/json;charset=ISO-8859-1
[
  {
    "event-data-items":[
      {
        "data-item-number":0,
        "data-item-type":"string",
        "data-item-value":"acsadmin"
      },
      {
        "data-item-number":1,
        "data-item-type":"string",
        "data-item-value":"Sx149"
      },
      {
        "data-item-number":2,
        "data-item-type":"string",
        "data-item-value":"admin.my.company.com [1.2.3.4]"
      }
    ],
    "event-details":[],
    "event-id":"1941",
    "event-message":"User acsadmin has logged on to Web Services API session
      Sx149 from location admin.my.company.com [1.2.3.4]",
    "event-name":"WSA Logon",
    "event-time":1412285249660,
    "user-uri":"/api/users/ae8aed68-3dc0-11e4-8dd1-1c6f65065a91",
    "userid":"ACSADMIN"
  },
  {
    "event-data-items":[
      {
        "data-item-number":0,
        "data-item-type":"string",
        "data-item-value":"Dept Admin"
      }
    ],
    "event-details":[
      {
        "event-details-data-items":[],
        "event-details-message":"Task Role: Dept Admin --Based on role is
          null. Permitted tasks:"
      }
    ],
    "event-id":"1272",
    "event-message":"The task role Dept Admin has been created.",
    "event-name":"LOGTROLEADD",
    "event-time":1412285252280,
    "user-uri":"/api/users/ae8aed68-3dc0-11e4-8dd1-1c6f65065a91",
    "userid":"ACSADMIN"
  },
  {
    "event-data-items":[
      {
        "data-item-number":0,
        "data-item-type":"string",
        "data-item-value":"Dept Admin"
      }
    ],
    "event-details":[
      {
        "event-details-data-items":[],
        "event-details-message":"Task Role: Dept Admin --Based on role is
          null. Permitted tasks: ClassId=XVirtualServer"
      }
    ],
    "event-id":"1273",

```

## List Console Hardware Messages

The **List Console Hardware Messages** operation lists the current set of hardware messages associated with the console.

### HTTP method and URI

**GET** /api/console/hardware-messages

### Query Parameters

Name	Type	Rqd/Opt	Description
begin-time	Timestamp	Optional	A timestamp used to filter hardware messages. Messages created earlier than this time are omitted from the results. The value is specified as the number of milliseconds since the epoch and must be greater than or equal to 0.  If not specified, then no such filtering is performed.
end-time	Timestamp	Optional	A timestamp used to filter hardware messages. Messages created later than this time are omitted from the results. The value is specified as the number of milliseconds since the epoch and must be greater than or equal to 0.  If not specified, then no such filtering is performed.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
hardware-messages	Array of hardware-message-info objects	Array of nested hardware-message-info objects as defined in the next table.

Each nested hardware-message-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	The canonical URI path of the hardware message. The URI is in the following form: /api/console/hardware-messages/{hardware-message-id}
timestamp	Timestamp	The date and time the hardware message was created
text	String	The text of the hardware message.

### Description

This operation returns a set of console hardware messages in increasing timestamp order, filtered according to the query parameters, if specified.

If the **begin-time** query parameter is specified, then any entries earlier than that time are omitted. If the **end-time** query parameter is specified, then any entries later than that time are omitted.

If there are no hardware messages associated with the console, or if no hardware messages are to be included in the results due to filtering, an empty array is returned and the operation completes successfully.

| The API user must have Action/Task permission to the Hardware Messages task; otherwise, status code 403 (Forbidden) is returned.

### | **Authorization requirements**

| This operation has the following authorization requirement:

- | • Action/Task permission to the Hardware Messages task.

### | **HTTP status and reason codes**

| On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 637.

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

| *Table 189. List Console Hardware Messages: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	7	The <b>begin-time</b> value is greater than the <b>end-time</b> value.
403 (Forbidden)	1	The API user does not have Action/Task permission for the Hardware Messages task.

| Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### | **Example HTTP interaction**

---

```
GET /api/console/hardware-messages HTTP/1.1
x-api-session: 4pw14919jtzcwohdfe8s9gw5zzv7v73yksomswrg50t7ni4q8r
```

---

| *Figure 303. List Console Hardware Messages: Request*

---

```

200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Mon, 06 Oct 2014 17:02:37 GMT
content-type: application/json;charset=UTF-8
content-length: 206
{
  "hardware-messages": [
    {
      "element-uri": "/api/console/hardware-messages/11c5f16a-4d58-11e4-ba8d-
        02215e673710",
      "text": "Licensed internal code has detected a problem.  [Problem # 3]",
      "timestamp": 1412600137860
    }
  ]
}

```

---

Figure 304. List Console Hardware Messages: Response

## Get Console Hardware Message Properties

The **Get Console Hardware Message Properties** operation retrieves the properties of a single console hardware message.

### HTTP method and URI

**GET** /api/console/hardware-messages/{*hardware-message-id*}

In this request, the URI variable {*hardware-message-id*} is the unique identifier of the hardware message to be retrieved.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the console hardware message object as defined in “Data model” on page 618. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

### Description

This operation retrieves the properties of a single console hardware message specified by {*hardware-message-id*}.

The URI path must designate an existing hardware message; otherwise, status code 404 (Not Found) is returned. In addition, the API user must have Action/Task permission to the **Hardware Messages** task; otherwise, status code 403 (Forbidden) is returned.

### Authorization requirements

This operation has the following authorization requirement:

- Action/Task permission to the **Hardware Messages** task.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have Action/Task permission for the <b>Hardware Messages</b> task.
404 (Not Found)	322	The element ID in the URI <i>{hardware-message-id}</i> does not designate an existing console hardware message.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
GET /api/console/hardware-messages/11c5f16a-4d58-11e4-ba8d-02215e673710 HTTP/1.1
x-api-session: 1p5uuz63tskpoczrpjux2ni122ciqeoqrchza1t30jvb9ug09t
```

Figure 305. Get Console Hardware Message Properties: Request

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Mon, 06 Oct 2014 17:04:54 GMT
content-type: application/json;charset=UTF-8
content-length: 286
{
  "class": "hardware-message",
  "element-id": "11c5f16a-4d58-11e4-ba8d-02215e673710",
  "element-uri": "/api/console/hardware-messages/11c5f16a-4d58-11e4-ba8d-02215e673710",
  "parent": "/api/console",
  "text": "Licensed internal code has detected a problem. [Problem # 3]",
  "timestamp": 1412600137860
}
```

Figure 306. Get Console Hardware Message Properties: Response

## Delete Console Hardware Message

The **Delete Console Hardware Message** operation deletes a single console hardware message

### HTTP method and URI

```
DELETE /api/console/hardware-messages/{hardware-message-id}
```

In this request, the URI variable *{hardware-message-id}* is the unique identifier of the hardware message to be deleted.

### Description

This operation deletes a specific console hardware message. The hardware message to be deleted is identified by the *{hardware-message-id}* variable in the URI.



| The URI path must designate an existing hardware message; otherwise, status code 404 (Not Found) is returned. In addition, the API user must have Action/Task permission to the **Hardware Messages** task; otherwise, status code 403 (Forbidden) is returned.

### | **Authorization requirements**

| This operation has the following authorization requirement:

- | • Action/Task permission to the **Hardware Messages** task.

### | **HTTP status and reason codes**

| On success, HTTP status code 204 (No Content) is returned with no response body provided.

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have Action/Task permission for the <b>Hardware Messages</b> task.
404 (Not Found)	322	The element ID in the URI <i>{hardware-message-id}</i> does not designate an existing console hardware message.

| Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

### | **Example HTTP interaction**

---

```
DELETE /api/console/hardware-messages/6b2d61a4-a1ac-11e4-87ee-5ef3fcae8020 HTTP/1.1
x-api-session: c8un3odpy8yyp150o3poz1ud4gwyfodlwq495327bpyn2p0z
```

---

*Figure 307. Delete Console Hardware Message: Request*

---

```
204 No Content
date: Mon, 09 Feb 2015 20:07:31 GMT
server: zSeries management console API web server / 2.0
```

<No response body>

---

*Figure 308. Delete Console Hardware Message: Response*

### | **Inventory service data**

| Information about the console can be optionally included in the inventory data provided by the Inventory Service.

| Inventory entries for the Console objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**console**" are to be included.

For each Console object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for “Get Console Properties” on page 621. That is, the data provided is the same as would be provided if a **Get Console Properties** operation were requested targeting this object.

### Sample inventory data

The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single user. This object would appear as one array entry in the response array:

---

```

{
  "class": "console",
  "description": "Endicott Test HMC",
  "ec-mcl-description": {
    "ec": [
      {
        "description": "Hardware Management Console Framework",
        "mcl": [
          {
            "last-update": 1422026097000,
            "level": "39",
            "type": "retrieved"
          },
          {
            "last-update": 1423074018000,
            "level": "39",
            "type": "activated"
          },
          {
            "last-update": null,
            "level": "000",
            "type": "accepted"
          },
          {
            "last-update": null,
            "level": "39",
            "type": "installable-concurrent"
          },
          {
            "last-update": null,
            "level": "1",
            "type": "removable-concurrent"
          }
        ]
      },
      {
        "number": "N98841",
        "part-number": "00LY737",
        "type": "SYSTEM"
      }
    ]
  }
},

```

---

Figure 309. Console object: Sample inventory data (Part 1)

```
{
  "description": "Enablement of new features  ",
  "mcl": [
    {
      "last-update": null,
      "level": "000",
      "type": "retrieved"
    },
    {
      "last-update": null,
      "level": "000",
      "type": "activated"
    },
    {
      "last-update": null,
      "level": "000",
      "type": "accepted"
    },
    {
      "last-update": null,
      "level": "000",
      "type": "installable-concurrent"
    },
    {
      "last-update": null,
      "level": "000",
      "type": "removable-concurrent"
    }
  ],
  "number": "N98844",
  "part-number": "00LY740",
  "type": "ENABLE1"
},
]
```

Figure 310. Console object: Sample inventory data (Part 2)

```
| "ip-swapping-available": true,
| "is-auto-switch-enabled": true,
| "is-locked": false,
| "machine-info": {
|   "machine-model": "PBC",
|   "machine-serial": "KQ0N5RF",
|   "machine-type": "7382"
| },
| "name": "ZFXHMC2",
| "network-info": {
|   "paired-hmc": {
|     "hmc-name": "ZFXHMC1",
|     "ipv4-address": [
|       "9.60.15.110",
|       "9.60.14.110"
|     ],
|     "ipv6-address": [
|       "2002:93c:ffb:1:5ef3:fcff:feaf:deb1",
|       "fdd8:673b:d89b:1:5ef3:fcff:feaf:deb1",
|       "fe80:0:0:0:5ef3:fcff:feaf:deb1"
|     ]
|   },
|   "this-hmc": [
|     {
|       "domain-name": "",
|       "hmc-name": "ZFXHMC2",
|       "interface-name": "eth0",
|       "ipv4-address": [
|         {
|           "ip-address": "9.60.15.111",
|           "subnet-mask": "255.255.255.0"
|         }
|       ],
|       "ipv6-address": [
|         {
|           "ip-address": "2002:93c:ffb:1:5ef3:fcff:feae:8019",
|           "prefix-length": 64
|         },
|         {
|           "ip-address": "fdd8:673b:d89b:1:5ef3:fcff:feae:8019",
|           "prefix-length": 64
|         },
|         {
|           "ip-address": "fe80:0:0:0:5ef3:fcff:feae:8019",
|           "prefix-length": 64
|         }
|       ],
|       "is-private": false,
|       "mac": "5CF3FCAE8019"
|     }
|   ],
| },
```

Figure 311. Console object: Sample inventory data (Part 3)

```

|     {
|         "domain-name": "",
|         "hmc-name": "ZFXHMC2",
|         "interface-name": "eth1",
|         "ipv4-address": [
|             {
|                 "ip-address": "9.60.14.111",
|                 "subnet-mask": "255.255.255.0"
|             }
|         ],
|         "ipv6-address": [
|             {
|                 "ip-address": "fe80:0:0:0:5ef3:fcff:feae:801a",
|                 "prefix-length": 64
|             }
|         ],
|         "is-private": false,
|         "mac": "5CF3FCAE801A"
|     }
| ]
| },
| "object-id": "ec982d6c-bcc1-3ae8-b39c-a2efd14734b4",
| "object-uri": "/api/console",
| "paired-role": "primary",
| "parent": null,
| "version": "2.13.0"
| }

```

Figure 312. Console object: Sample inventory data (Part 4)

## User-related-access permission

An HMC user has access to certain information about their own user account. This information is contained in their User object and related objects. An API user's access permission to their own User object and specific related objects is known as user-related-access permission. Through such permission, those objects will be included in a List <class> operation, unless otherwise filtered out, and the API user is permitted to issue a Get <class> Properties operation on them, unless specifically prohibited.

The object types included in user-related-access permission are:

- User
- User Role
- User Pattern
- Password Rule
- LDAP Server Definition

User-related-access permission includes the following:

- Permission for an object of the above types to be included in the response body of a List <class> operation.
- Permission to view properties of objects of the following types through the **Get <class> Properties** and **Get Inventory** operations:
  - User
  - User Role
  - User Pattern
  - Password Rule
- Permission to update certain properties of the User object. See the “Data model” on page 646 or the **Update User Properties** operation for details.

## User object

A User object represents a single Hardware Management Console user. There are different types of console users. A typical customer-defined user is known as a standard user. A user template defines certain attributes of a group of users whose user IDs match the expression in a User Pattern; these definitions are known as template users. When a user logs on with a user ID that matches the expression in a User Pattern, a pattern-based user is created. There are certain user definitions supplied by the system; they are known as system-defined users.

All API users are permitted to see their own User object in a **List Users** response, issue **Get User Properties** for their own User object and, with the exception of pattern-based users, issue **Update User Properties** to alter certain properties of their own User object. An API user with action/task permission to the **Manage Users** task is permitted to view and change any standard or system-defined User object. An API user with action/task permission to the **Manage User Templates** task is permitted to view and change any template User object.

User objects may be replicated to this HMC via its Data Replication facility. If that is the case, the **Update User Properties**, **Add User Role to User**, **Remove User Role from User** and **Delete User** operations should be used with care as they will prevent further replication of the modified or deleted object to this HMC.

## System-defined users

In most respects, system-defined users are indistinguishable from standard users. They can be modified or even deleted. Most properties of system-defined users may be changed, but certain others are immutable; the immutable properties are denoted as such in the Data Model section that follows. While system-defined users can be deleted and their name reused for a standard user definition, that practice is discouraged due to the likely confusion such a situation would cause. The typical system-defined users include the following:

- ACSADMIN
- ADVANCED
- ENSADMIN
- ENSOPERATOR
- OPERATOR
- SERVICE
- SYSPROG

## Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, but does not provide the operational-status-related properties defined in that schema because it does not maintain the concept of an operational status.

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

The following class-specific specializations apply to the other base managed object properties:

Table 190. User object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
object-uri	—	String/URI	The canonical URI path of the User object is of the form <code>/api/users/{user-id}</code> , where <code>{user-id}</code> is the value of the <b>object-id</b> property (not the <b>name</b> property) of the User object.
parent	—	String/URI	The canonical URI path of the User object.

Table 190. User object: base managed object properties specializations (continued)

Name	Qualifier	Type	Description of specialization
<b>class</b>	—	String	The class of a User object is "user".
<b>name</b>	(ro)	String	<p>If type is not "template": the name (console user ID) of the User object. It must be 4-320 characters in length and consist only of alphanumeric characters, spaces and the following special characters: "@&lt;+:#="&amp;*()-/\,%_&gt;.?". This name must be unique among all users whose type is not "template" defined on the console.</p> <p>If type is "template": the name of the template. While preexisting template names are virtually unrestricted in terms of length and characters, new template names must conform to the length and character requirement of the <b>name</b> property described in the "Base managed object properties schema" on page 39. This name must be unique among all template definitions on the console.</p> <p>For the purpose of verifying uniqueness only, this name is treated in a case-insensitive fashion when used to create a new User object of any type.</p>
<b>description</b>	(w)(pc)	String (0-1024)	<p>The description of the User object.</p> <p>Default: an empty string</p>

### Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties:

**Note:** Some properties are only valid for users of a specific type. Such properties are only included in the User object if the user is of that type, as indicated by its **type** property. For example, a user with a **type** of "standard" includes the **disabled** property but not the **user-pattern-uri** property.

Certain properties are only valid when mutable prerequisite properties have specific values. When such properties are not valid, their value is **null**. For instance the **password-rule-uri** is **null** when the **authentication-type** value is "ldap".

Table 191. User object: class specific additional properties

Name	Qualifier	Type	Description
<b>type</b>	—	String Enum	<p>The type of user. Supported values are:</p> <ul style="list-style-type: none"> <li>"standard" - a standard, normal user.</li> <li>"template" - a user template.</li> <li>"pattern-based" - a user created dynamically from a User Pattern and its associated template.</li> <li>"system-defined" - a user supplied by the system. Certain properties of system-defined users are immutable.</li> </ul>
<b>user-pattern-uri</b>	—	String/URI	<p>The canonical URI path of the User Pattern object upon which this user is based.</p> <p>Prerequisite: <b>type</b> is "pattern-based"</p>
<b>disabled</b>	(w)(pc)	Boolean	<p>Indicates whether the user is currently disabled. When disabled, the user is prevented from logging on to the console via either the UI or the Web Services APIs.</p> <p>Prerequisite: <b>type</b> is not "template".</p> <p>Default: <b>false</b></p>

Table 191. User object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>authentication-type</b>	(w)(pc)	String Enum	<p>The type of user ID and password authentication used for this user, which must be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>"local"</b> - the console performs the authentication.</li> <li>• <b>"ldap"</b> - authentication is delegated to the LDAP server identified in the <b>ldap-server-definition-uri</b> property.</li> </ul> <p>If <b>type</b> is <b>"template"</b>, this must be <b>"ldap"</b>.</p> <p><b>Note:</b> The value of this property is a prerequisite for certain other properties. Changing this value requires certain properties to be included in the same request; see the <b>Update User Properties</b> operation for details.</p>
<b>password-rule-uri</b>	(w)(pc)	String/URI	<p>The canonical URI path of the Password Rule for this user.</p> <p>Prerequisite: <b>authentication-type</b> is <b>"local"</b>.</p>
<b>password</b>	(wo)(pc)	String	<p>The console logon password for this user. The specific length, character and other requirements on this password are controlled by the authentication type and Password Rule assigned to this user.</p> <p>Note the (wo) qualifier; this field may be altered via an API, but it is not included in the response when this object's properties are retrieved via an API.</p>
<b>password-expires</b>	—	Integer	<p>The time interval, in days, until the user's current password expires. A value of 0 indicates that the password will expire within the next 24 hours. A value of -1 indicates that the HMC does not enforce password expiration for this user; however, if this user is authenticated with an external authentication mechanism (e.g. LDAP) such expiration might be enforced by that mechanism.</p>
<b>force-password-change</b>	(w)(pc)	Boolean	<p>Indicates whether the user should be forced to change their console logon password the next time they log in.</p> <p>Prerequisite: <b>authentication-type</b> is <b>"local"</b></p> <p>Default: <b>true</b></p>
<b>ldap-server-definition-uri</b>	(w)(pc)	String/URI	<p>The canonical URI path of the configuration object for the LDAP server used for authentication of this user.</p> <p>Prerequisite: <b>authentication-type</b> is <b>"ldap"</b>.</p>
<b>userid-on-ldap-server</b>	(w)(pc)	String (0-32)	<p>The user ID for this user on the LDAP server identified in <b>ldap-server-definition-uri</b>, or <b>null</b> if the user's console user ID (value of the <b>name</b> property) should be used. See the LDAP Server Definition object for more information on how this property is used.</p> <p>Prerequisite: <b>authentication-type</b> is <b>"ldap"</b> and <b>type</b> is not <b>"template"</b>.</p> <p>Default: an empty string</p>
<b>session-timeout</b>	(w)(pc)	Integer (0-525600)	<p>The session timeout in minutes for this user. This is the interval over which a user's UI session can run before being prompted for identity verification. 0 indicates no timeout.</p> <p>Default: 0</p>



Table 191. User object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>verify-timeout</b>	(w)(pc)	Integer (0-525600)	The verification timeout in minutes for this user. This is the amount of time allowed for the user to re-enter their password after being prompted due to a session timeout (see the <b>session-timeout</b> property). 0 indicates no timeout.  Default: 15
<b>idle-timeout</b>	(w)(pc)	Integer (0-525600)	The idle timeout in minutes for this user. This is the amount of time the user's UI session can be idle before it is disconnected. 0 indicates no timeout.  Default: 0
<b>min-pw-change-time</b>	(w)(pc)	Integer (0-525600)	The minimum password change time in minutes for this user. This is the minimum amount of time that must elapse between changes to this user's password. 0 indicates no minimum; that is, the password can be changed immediately after it has just been changed.  Prerequisite: <b>authentication-type</b> is "local".  Default: 0
<b>max-failed-logins</b>	(w)(pc)	Integer (0-525600)	The maximum number of failed login attempts for this user. This is maximum number of consecutive failed login attempts before the user is temporarily disabled for the amount of time specified in the <b>disable-delay</b> property. 0 indicates that the user is never disabled due to failed login attempts.  Default: 3
<b>disable-delay</b>	(w)(pc)	Integer (0-525600)	The time in minutes that the user is disabled after exceeding the maximum number of failed login attempts specified in the <b>max-failed-logins</b> property. 0 indicates that the user is not disabled for any period of time after reaching the maximum number of invalid login attempts.  Default: 1
<b>inactivity-timeout</b>	(w)(pc)	Integer (0-525600)	The inactivity timeout in days for this user. This is the maximum number of days of inactivity (consecutive days with no login) before the user is disabled. 0 indicates no timeout.  Default: 0
<b>disruptive-pw-required</b>	(w)(pc)	Boolean	Indicates whether the user's password is required to perform disruptive actions via the UI.  Default: true
<b>disruptive-text-required</b>	(w)(pc)	Boolean	Indicates whether text input is required to perform disruptive actions via the UI.  Default: false
<b>allow-remote-access</b>	(w)(pc)	Boolean	Indicates whether the user is allowed to access the HMC via its remote web server interface  Default: false

Table 191. User object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>allow-management-interfaces</b>	(w)(pc)	Boolean	Indicates whether the user is allowed access to management interfaces. This includes access to the Web Services APIs.  Default: false
<b>max-web-services-api-sessions</b>	(w)(pc)	Integer (0-9999)	The maximum number of simultaneous Web Services API sessions the user is permitted to have.  Default: 100
<b>web-services-api-session-idle-timeout</b>	(w)(pc)	Integer (1-360)	The idle timeout in minutes for Web Services API sessions created by this user. This is the amount of time a Web Services API session can be idle before it is terminated.  Default: 360
<b>user-roles</b>	(c)(pc)	Array of String/URI	The list of user roles defined for this user. Each element in this array is a canonical URI path for a User Role object. The roles provided in this list can change as a result of the <b>Add User Role to User</b> and <b>Remove User Role from User</b> operations.  This property is immutable if <b>type</b> is "system-defined".
<b>default-group-uri</b>	(w)(pc)	String/URI	The canonical URI path of the user's default group or <b>null</b> if the user has no default group. Managed objects created by this user automatically become members of this group. The user must have object-access permission to this group. This must be a user-defined group to which the user has object-access permission.  API users are permitted to change their own default group designation via the <b>Update User Properties</b> operation.  Default: null
<b>replication-overwrite-possible</b>	—	Boolean	Indicates whether this object is customizable data that is replicated to this HMC from an HMC configured as a Data Source in the Data Replication service.

## List Users

The **List Users** operation lists standard, template and system-defined users defined to the console. With one very specific exception, pattern-based users are never included in the response to the **List Users** operation.

## HTTP method and URI

**GET** /api/console/users

## Query Parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property. Value must be a valid <b>type</b> property value, with the exception of "pattern-based".

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
users	Array of objects	Array of nested user-info objects as described in the next table.

Each nested user-info object contains the following fields:

Field name	Type	Description
object-uri	String/URI	The <b>object-uri</b> property of the User object.
name	String	The <b>name</b> property of the User object.
type	String Enum	The <b>type</b> property of the User object.

## Description

The **List Users** operation lists users defined to the console. Some basic properties are provided for each user.

If the **name** query parameter is specified, the returned list is limited to those users that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not performed.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid user **type** property value for this operation. If the value is not valid, status code 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those users that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not performed.

A user is included in the list only if the API user has sufficient access permission to that object. The access permission requirements for User objects vary depending on the type of User object. All API users have user-related-access permission to their own User object. Action/task permission to the **Manage Users** task includes access permission to all non-template Users, and action/task permission to the **Manage User Templates** task includes access permission to all template users. If there is a User object to which the API user does not have permission, that object is omitted from the list, but no error status code results.

A pattern-based user is only included in the response if it meets the filtering criteria and is the User object for the API user.

If there are no users defined to the console or if no users are to be included in the results due to filtering or access permissions, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirement:

- User-related-access permission to the User object included in the response body, or, depending on the type of User object, action/task permission to the **Manage Users** task or the **Manage User Templates** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 651.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 192. List Users: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/console/users?type=system-defined&name=.*ADMIN HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

Figure 313. List Users: Request

---

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:30 GMT
content-type: application/json;charset=UTF-8
content-length: 225
{
  "users": [
    {
      "name": "ACADMIN",
      "object-uri": "/api/users/ae8aed68-3dc0-11e4-8dd1-1c6f65065a91",
      "type": "system-defined"
    },
    {
      "name": "ENSADMIN",
      "object-uri": "/api/users/ae6bf048-3dc0-11e4-8dd1-1c6f65065a91",
      "type": "system-defined"
    }
  ]
}
```

---

Figure 314. List Users: Response

## Get User Properties

The **Get User Properties** operation retrieves the properties of a single User object that is designated by its object ID. With one very specific exception, this operation does not support pattern-based users.

## HTTP method and URI

**GET** /api/users/{user-id}

In this request, the URI variable *{user-id}* is either the object ID of the User object whose properties are to be returned or the keyword value **"this-user"** which designates the API user that issued the request.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the User object as defined in the Data Model section. Field names and data types in the JSON object are the same as the property names and data types defined in the “Data model” on page 646.

### Description

This operation returns the current properties of a single User object that is designated by *{user-id}*.

On successful execution, all of the current properties as defined in the Data Model for the User object, except those designated as write-only properties, are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing User object and the API user must have access permission to it. All API users have user-related-access permission to their own User object. Action/task permission to the **Manage Users** task includes access permission to all non-template Users, and action/task permission to the **Manage User Templates** task includes access permission to all template users. If the URI path does not designate an existing User object or the API user does not have access permission to it, status code 404 (Not Found) is returned.

This operation does not support pattern-based users, unless the target of the operation is the User object for the API user that issued the request.

### Authorization requirements

This operation has the following authorization requirement:

- User-related-access permission to the User object specified in the request URI, or, depending on the type of User object specified in the request URI, action/task permission to the **Manage Users** task or the **Manage User Templates** task.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type, or designates a resource for which the API user does not have the required authorization.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
GET /api/users/e9e8d20a-4a7a-11e4-91ee-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

Figure 315. Get User Properties: Request

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:30 GMT
content-type: application/json;charset=UTF-8
content-length: 1100
{
  "allow-management-interfaces":false,
  "allow-remote-access":false,
  "authentication-type":"local",
  "class":"user",
  "default-group-uri":null,
  "description":"Gabby McRosie - company president",
  "disable-delay":1,
  "disabled":false,
  "disruptive-pw-required":true,
  "disruptive-text-required":false,
  "force-password-change":true,
  "idle-timeout":0,
  "inactivity-timeout":0,
  "is-locked":false,
  "ldap-server-definition-uri":null,
  "max-failed-logins":3,
  "max-web-services-api-sessions":100,
  "min-pw-change-time":0,
  "name":"Gabby",
  "object-id":"e9e8d20a-4a7a-11e4-91ee-1c6f65065a91",
  "object-uri":"/api/users/e9e8d20a-4a7a-11e4-91ee-1c6f65065a91",
  "parent":"/api/console",
  "password-expires":-1,
  "password-rule-uri":"/api/console/password-rules/4a790766-3dbf-11e4-980d-1c6f65065a91",
  "replication-overwrite-possible":false,
  "session-timeout":0,
  "type":"standard",
  "user-roles":[
    "/api/user-roles/ea6f9b14-4a7a-11e4-affa-1c6f65065a91",
    "/api/user-roles/ea41a664-4a7a-11e4-91ee-1c6f65065a91",
    "/api/user-roles/ea094df0-4a7a-11e4-8777-1c6f65065a91"
  ],
  "userid-on-ldap-server":null,
  "verify-timeout":15,
  "web-services-api-session-idle-timeout":360
}
```

Figure 316. Get User Properties: Response

## Update User Properties

The **Update User Properties** operation updates the properties of a single User object that is designated by its object ID. This operation is not valid for pattern-based users.

### HTTP method and URI

**POST** /api/users/{user-id}

In this request, the URI variable *{user-id}* is the object ID of the User object whose properties are to be updated or the special keyword value **"this-user"** which designates the API user that issued the request.

### Request body contents

The request body is expected to contain a JSON object that provides the new values of any writeable property that is to be updated by this operation. Field names and data types in this JSON object are expected to match the corresponding property names and data types defined by the data model for this object type. The JSON object can and should omit fields for properties whose values are not to be changed by this operation.

### Description

This operation updates writeable properties of the User object specified by *{user-id}*.

The URI path must designate an existing User object and the API user must have permission to update it. All API users have user-related-access permission to their own User object, which includes permission to update certain properties. Action/task permission to the **Manage Users** task includes update permission to all non-template users, and action/task permission to the **Manage User Templates** task includes update permission to all template users. If the URI path does not designate an existing User object, status code 404 (Not Found) is returned. If the API user does not have user-related-access permission to the designated User object or action/task permission to the **Manage Users** or **Manage User Templates** task, whichever is appropriate, status code 404 (Not Found) is returned. If the API user does not have action/task permission to the **Manage Users** or **Manage User Templates** task, whichever is appropriate, an attempt to update any field other than the user's own **password** or **default-group-uri** field results in status code 403 (Forbidden), and specifying a group to which the API user does not have object-access permission in **default-group-uri** results in status code 404 (Not Found).

The request body is validated against the schema described in the Request Body Contents section. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. The request body validation will fail if it contains a property that is not valid because a prerequisite is not met (e.g., attempting to set **password-rule-uri** when the **authentication-type** value is **"ldap"**). An attempt to update a pattern-based user is not valid and fails with status code 400 (Bad Request).

The request body does not need to specify a value for all writeable properties, but rather can and should contain fields only for the properties to be updated. Object properties for which no input value is provided and no prerequisite property is changed remain unchanged by this operation. A property's value is set to its default value if the field is not included in the request body and a prerequisite field is changed such that the prerequisite condition becomes satisfied (e.g., if **authentication-type** is changed from **"ldap"** to **"local"**, and **force-password-change** is not defined in the request body, **force-password-change** will be defaulted to false). Note however that certain fields must be included in the request body if the request alters the **authentication-type** property, because they are required in order to perform the newly specified type of logon authentication. Specifically, changing **authentication-type** to **"local"** requires that the following properties also be specified in the same request: **password**, **password-rule-uri**. Changing **authentication-type** to **"ldap"** requires that the following property also be specified: **ldap-server-definition-uri**.

If the update changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

### Authorization requirements

This operation has the following authorization requirements:

- For a user to update their own **password** or **default-group-uri** property, user-related-access permission to the User object specified in the request URI or action/task permission to the **Manage Users** task is required.
- An API user with action/task permission to the **Manage Users** task or the **Manage User Templates** task, depending on the type of User object specified in the request URI, may update any writeable property of that User object.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

*Table 193. Update User Properties: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	311	The new password does not conform to the requirements of the password policy in effect for this user.
	314	This operation is not supported for an object of this type. Pattern-based users may not be updated.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type, or designates a resource for which the API user does not have the required authorization.
	323	The <b>password-rule-uri</b> field in the request body does not designate an existing Password Rule object.
	324	The <b>ldap-server-definition-uri</b> field in the request body does not designate an existing LDAP Server Definition object.
	325	The <b>default-group-uri</b> field in the request body does not designate an existing resource of the expected type, or designates a resource for which the user identified by the request URI does not have object-access permission.
409 (Conflict)	321	The user's <b>authentication-type</b> property cannot be changed at this time. The specified user is currently the only locally-authenticated user with permission to the tasks for managing users and user roles.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.



## Example HTTP interaction

```
POST /api/users/e9e8d20a-4a7a-11e4-91ee-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 145
{
  "allow-management-interfaces":true,
  "description":"A new and improved description of this User",
  "web-services-api-session-idle-timeout":240
}
```

Figure 317. Update User Properties: Request

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:30 GMT

<No response body>
```

Figure 318. Update User Properties: Response

## Add User Role to User

The **Add User Role to User** operation adds a specified User Role to a specified user. This operation is not valid for system-defined or pattern-based users.

### HTTP method and URI

**POST** /api/users/{user-id}/operations/add-user-role

In this request, the URI variable {user-id} is the object ID of the user to which a User Role is to be added or the special keyword value **"this-user"** which designates the API user that issued the request.

### Request body contents

The request body is expected to contain a JSON object with the following field:

Field name	Type	Rqd/Opt	Description
user-role-uri	String/ URI	Required	The canonical URI path of the User Role to be added.

### Description

This operation adds a User Role to a user.

On successful execution of this operation the User Role specified in the request body has been added to the user identified in the request URI.

The request body is validated against the schema described in "Request body contents." If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. If the request URI does not designate an existing User object, status code 404 (Not Found) is returned. If the API user does not have user-related-access permission to the designated User object or action/task permission to the **Manage Users** or **Manage User Templates** task, whichever is appropriate, status code 404 (Not Found) is returned. If the API user has user-related-access permission to

the designated User object but not action/task permission to the **Manage Users** or **Manage User Templates** task, whichever is appropriate, status code 403 (Forbidden) is returned. If the request body does not designate an existing User Role, status code 404 (Not Found) is returned. If the specified object is already in the collection of the user's User Roles, status code 409 (Conflict) is returned. An attempt to update the User Role collection of a system-defined or pattern-based user is not valid and fails with status code 400 (Bad Request).

If this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

### Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage Users** task to modify a standard user or the **Manage User Templates** task to modify a template user.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 194. Add User Role to User: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	314	This operation is not supported for an object of this type. The User Role collection of system-defined and pattern-based users may not be altered.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
	2	A URI in the request body does not designate an existing resource of the correct type.
409 (Conflict)	315	The object designated by the URI in the request body is already in the user's collection of User Roles.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
POST /api/users/e9e8d20a-4a7a-11e4-91ee-1c6f65065a91/operations/add-user-role HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 73
{
  "user-role-uri":"/api/user-roles/eaecdf34-4a7a-11e4-8777-1c6f65065a91"
}
```

Figure 319. Add User Role to User: Request

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:31 GMT

<No response body>
```

Figure 320. Add User Role to User: Response

## Remove User Role from User

The **Remove User Role from User** operation removes a specified User Role from a specified user. This operation is not valid for system-defined or pattern-based users.

### HTTP method and URI

**POST** /api/users/{user-id}/operations/remove-user-role

In this request, the URI variable {user-id} is the object ID of the user from which a User Role is to be removed or the special keyword value **"this-user"** which designates the API user that issued the request.

### Request body contents

The request body is expected to contain a JSON object with the following field:

Field name	Type	Rqd/Opt	Description
user-role-uri	String/ URI	Required	The canonical URI path of the User Role to be removed.

### Description

This operation removes a User Role from a user.

On successful execution of this operation the User Role specified in the request body has been removed from the user identified in the request URI.

The request body is validated against the schema described in "Request body contents." If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. If the request URI does not designate an existing User object, status code 404 (Not Found) is returned. If the API user does not have user-related-access permission to the designated User object or action/task permission to the **Manage Users** or **Manage User Templates** task, whichever is appropriate, status code 404 (Not Found) is returned. If the API user has user-related-access permission to the designated User object but not action/task permission to the **Manage Users** or **Manage User**

If the **Templates** task, whichever is appropriate, status code 403 (Forbidden) is returned. If the specified object is not in the collection of the user's User Roles, status code 409 (Conflict) is returned. An attempt to update the User Role collection of a system-defined or pattern-based user is not valid and fails with status code 400 (Bad Request).

If this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

### Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage Users** task to modify a standard user or the **Manage User Templates** task to modify a template user.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 195. Remove User Role from User: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	314	This operation is not supported for an object of this type. The User Role collection of system-defined and pattern-based users may not be altered.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
	2	A URI in the request body does not designate an existing resource of the correct type.
409 (Conflict)	316	The object designated by the URI in the request body is not in the user's collection of User Roles.
	321	The User Role cannot be removed at this time. The user is currently the only locally-authenticated user with permission to the tasks for managing users and user roles.
	328	The User Role cannot be removed at this time, because doing so would leave the user without object-access permission to their default group.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
POST /api/users/e9e8d20a-4a7a-11e4-91ee-1c6f65065a91/operations/remove-user-role HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 73
{
  "user-role-uri":"/api/user-roles/eaecdf34-4a7a-11e4-8777-1c6f65065a91"
}
```

Figure 321. Remove User Role from User: Request

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:31 GMT

<No response body>
```

Figure 322. Remove User Role from User: Response

## Create User

The **Create User** operation creates a standard or template User object with the given properties. This operation is not valid for system-defined or pattern-based users.

### HTTP method and URI

**POST** /api/console/users

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The value to be set as the user's <b>name</b> property.  Note that the length and character requirements that apply to this field are dependent on the value of the <b>type</b> field.
description	String	Optional	The value to be set as the user's <b>description</b> property.
type	String Enum	Required	The value to be set as the user's <b>type</b> property. Must be <b>"standard"</b> or <b>"template"</b> .
disabled	Boolean	Optional	The value to be set as the user's <b>disabled</b> property.
authentication-type	String Enum	Required	The value to be set as the user's <b>authentication-type</b> property.
password-rule-uri	String/URI	Required if <b>authentication-type</b> is <b>"local"</b>	The value to be set as the user's <b>password-rule-uri</b> property.
password	String	Required if <b>authentication-type</b> is <b>"local"</b>	The value to be set as the user's <b>password</b> property.
force-password-change	Boolean	Optional	The value to be set as the user's <b>force-password-change</b> property.

Field name	Type	Rqd/Opt	Description
ldap-server-definition-uri	String/URI	Required if <b>authentication-type</b> is "ldap"	The value to be set as the user's <b>ldap-server-definition-uri</b> property.
userid-on-ldap-server	String	Optional	The value to be set as the user's <b>userid-on-ldap-server</b> property.
session-timeout	Integer	Optional	The value to be set as the user's <b>session-timeout</b> property.
verify-timeout	Integer	Optional	The value to be set as the user's <b>verify-timeout</b> property.
idle-timeout	Integer	Optional	The value to be set as the user's <b>idle-timeout</b> property.
min-pw-change-time	Integer	Optional	The value to be set as the user's <b>min-pw-change-time</b> property.
max-failed-logins	Integer	Optional	The value to be set as the user's <b>max-failed-logins</b> property.
disable-delay	Integer	Optional	The value to be set as the user's <b>disable-delay</b> property.
inactivity-timeout	Integer	Optional	The value to be set as the user's <b>inactivity-timeout</b> property.
disruptive-pw-required	Boolean	Optional	The value to be set as the users <b>disruptive-pw-required</b> property.
disruptive-text-required	Boolean	Optional	The value to be set as the user's <b>disruptive-text-required</b> property.
allow-remote-access	Boolean	Optional	The value to be set as the user's <b>allow-remote-access</b> property.
allow-management-interfaces	Boolean	Optional	The value to be set as the user's <b>allow-management-interfaces</b> property.
max-web-services-api-sessions	Integer	Optional	The value to be set as the user's <b>max-web-services-api-sessions</b> property.
web-services-api-session-idle-timeout	Integer	Optional	The value to be set as the user's <b>web-services-api-session-idle-timeout</b> property.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the new User object.

## Description

This operation creates a new console user.

On successful execution of this operation the user is created using the inputs as specified by the request body. The URI of the new user is provided in the response body and in a **Location** response header as well. An Inventory Change notification is emitted asynchronously.

The request body is validated against the schema described in the "Request body contents" on page 661. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. The request body validation will fail if it contains a property that is not valid because a prerequisite is not met (e.g., specifying **password-rule-uri** when the **authentication-type** value is "ldap") or the specified name is not unique. If a URI in the request body does not designate an

| existing resource of the appropriate type, status code 404 (Not Found) is returned. In addition, the API user must have action/task permission to the **Manage Users** task to create a standard user or the **Manage User Templates** task to create a template user; otherwise, status code 403 (Forbidden) is returned.

| Certain user names are used internally by the Hardware Management Console and are therefore not available for use when creating a new user. An attempt to create a user with one of these names results in status code 400 (Bad Request) indicating that there is already a user with that name. The list of such names is case-insensitive and includes the following:

- | • SOOACSADMIN
- | • SOOSERVICE
- | • SOOSYSPROG
- | • SOOADVANCED
- | • SOOOPERATOR
- | • SOOENSADMIN
- | • SOOENSOPERATOR
- | • PEDEBUG

### | **Authorization requirements**

| This operation has the following authorization requirement:

- | • Action/task permission to the **Manage Users** task to create a standard user or the **Manage User Templates** task to create a template user.

### | **HTTP status and reason codes**

| On success, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents” on page 662, and the **Location** response header contains the URI of the newly created object.

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

| *Table 196. Create User: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	8	A user with the name specified in the request body already exists.
	311	The password does not conform to the requirements of the password policy in effect for this user.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	323	The <b>password-rule-uri</b> field in the request body does not designate an existing Password Rule object.
	324	The <b>ldap-server-definition-uri</b> field in the request body does not designate an existing LDAP Server Definition object.

| Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/console/users HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 234
{
  "authentication-type":"local",
  "description":"Gabby McRosie - company president",
  "name":"Gabby",
  "password":"abc123pw",
  "password-rule-uri":"/api/console/password-rules/4a790766-3dbf-11e4-980d-1c6f65065a91",
  "type":"standard"
}
```

---

Figure 323. Create User: Request

---

```
201 Created
server: zSeries management console API web server / 2.0
location: /api/users/e9e8d20a-4a7a-11e4-91ee-1c6f65065a91
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:29 GMT
content-type: application/json;charset=UTF-8
content-length: 64
{
  "object-uri":"/api/users/e9e8d20a-4a7a-11e4-91ee-1c6f65065a91"
}
```

---

Figure 324. Create User: Response

## Delete User

The **Delete User** operation deletes a User object designated by its object ID. This operation is not valid for pattern-based users.

### HTTP method and URI

**DELETE** /api/users/{*user-id*}

In this request, the URI variable {*user-id*} is the object ID of the User object to be deleted.

### Description

This operation removes a specified user from the console. The user is identified by the {*user-id*} variable in the URI.

Upon successfully removing the user, HTTP status code 204 (No Content) is returned and no response body is provided. An Inventory Change notification is emitted asynchronously.

The URI path must designate an existing User object; otherwise, status code 404 (Not Found) is returned. If the API user does not have user-related-access permission to the designated User object or action/task permission to the **Manage Users** or **Manage User Templates** task, whichever is appropriate, status code 404 (Not Found) is returned. If the API user has user-related-access permission to the designated User object but not action/task permission to the **Manage Users** or **Manage User Templates** task, whichever is appropriate, status code 403 (Forbidden) is returned. It is an error for an API user to attempt to delete his own User object; any attempt to do so results in status code 400 (Bad Request). If the request URI identifies a template user, and a user or a User Pattern refers to that template user, the request fails and



| status code 409 (Conflict) is returned. An attempt to delete a pattern-based user is not valid and fails with status code 400 (Bad Request).

### | **Authorization requirements**

| This operation has the following authorization requirement:

- | • Action/task permission to the **Manage Users** task to delete a non-template user, or the **Manage User Templates** task to delete a template user.

### | **HTTP status and reason codes**

| On success, HTTP status code 204 (No Content) is returned no response body is provided.

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

| *Table 197. Delete User: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	312	This operation is not supported for an object of this type. Pattern-based users may not be deleted.
	313	The request URI designates the API user's own User object.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
409 (Conflict)	317	The object cannot be deleted at this time. One or more users or User Patterns refer to this template user.
	320	The object cannot be deleted at this time. It is currently identified as the Automatic Logon ID for the Hardware Management Console.
	321	The object cannot be deleted at this time. It is currently the only locally-authenticated user with permission to the tasks for managing users and user roles.

| Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### | **Example HTTP interaction**

---

```
DELETE /api/users/e9e8d20a-4a7a-11e4-91ee-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

| *Figure 325. Delete User: Request*

---

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:31 GMT
```

```
<No response body>
```

---

*Figure 326. Delete User: Response*

## | **Inventory service data**

| Information about the users managed by the console can be optionally included in the inventory data provided by the Inventory Service.

| Inventory entries for User objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**user**" are to be included. An entry for a particular user is included only if the API user has access permission to that object as described in the **Get User Properties** operation.

| For each User object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the **Get User Properties** operation. That is, the data provided is the same as would be provided if a **Get User Properties** operation were requested targeting this object.

## | **Sample inventory data**

| The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single user. This object would appear as one array entry in the response array:

|

```

|   {
|     "allow-management-interfaces": false,
|     "allow-remote-access": false,
|     "authentication-type": "ldap",
|     "class": "user",
|     "default-group-uri": "/api/groups/fc400fc6-54e3-335d-854f-a0a6d10f000b",
|     "description": "System administrator",
|     "disable-delay": 1,
|     "disabled": false,
|     "disruptive-pw-required": true,
|     "disruptive-text-required": false,
|     "force-password-change": null,
|     "idle-timeout": 0,
|     "inactivity-timeout": 0,
|     "is-locked": false,
|     "ldap-server-definition-uri": "/api/console/ldap-server-definitions/
|       4927787e-34c4-11e4-alea-5ef3fcae8020",
|     "max-failed-logins": 3,
|     "max-web-services-api-sessions": 100,
|     "min-pw-change-time": null,
|     "name": "sysadmin",
|     "object-id": "9069f7a6-34c5-11e4-af4d-5ef3fcae8020",
|     "object-uri": "/api/users/9069f7a6-34c5-11e4-af4d-5ef3fcae8020",
|     "parent": "/api/console",
|     "password-expires": -1,
|     "password-rule-uri": null,
|     "replication-overwrite-possible": false,
|     "session-timeout": 0,
|     "type": "standard",
|     "user-roles": [
|       "/api/user-roles/b39afb87-d915-4070-a22f-91b158c6c01e"
|     ],
|     "userid-on-ldap-server": "sysadmin",
|     "verify-timeout": 15,
|     "web-services-api-session-idle-timeout": 15
|   }

```

Figure 327. User object: Sample inventory data

## User Role object

A User Role object represents an authority role which can be assigned to one or more console users. A role may allow access to specific managed objects, classes of managed objects, groups and/or tasks. There are two types of User Roles: user-defined and system-defined. User-defined User Roles are created by a console user, whereas the system-defined User Roles are pre-defined, standard User Roles supplied with the console.

User role objects may be replicated to this HMC via its Data Replication facility. If that is the case, the **Update User Role Properties**, **Add Permission to User Role**, **Remove Permission from User Role** and **Delete User Role** operations should be used with care as they will prevent further replication of the modified or deleted object to this HMC.

Through user-related-access permission described in “User-related-access permission” on page 645, API users are permitted to see certain User Role objects in a **List User Roles** response and issue **Get User Role Properties** for those User Role objects. An API user with action/task permission to the **Manage User Roles** task is permitted to view any User Role object and change any user-defined User Role object.

## | System-defined user roles

| There are many system-defined User Roles supplied with the console. System-defined roles may not be modified or deleted. The names assigned to the system-defined roles are enumerated in Appendix C, "Enum values for the User Role object," on page 933.

## | Data model

| This object includes the properties defined in the "Base managed object properties schema" on page 39, but does not provide the operational-status-related properties defined in that schema because it does not maintain the concept of an operational status.

| For definitions of the qualifier abbreviations in the following tables, see "Property characteristics" on page 38.

| The following class-specific specializations apply to the other base managed object properties:

| *Table 198. User Role object: base managed object properties specializations*

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/URI	The canonical URI path of the User Role object is of the form <code>/api/user-roles/{user-role-id}</code> , where <code>{user-role-id}</code> is the value of the <b>object-id</b> property of the User Role object.
<b>parent</b>	—	String/URI	The canonical URI path of the console object.
<b>class</b>	—	String	The class of a User Role object is <b>"user-role"</b> .
<b>name</b>	(ro)	String	The name of the User Role object. This name must be unique among all of the console's User Roles with the same <b>type</b> value. While pre-existing User Role names are virtually unrestricted in terms of length and characters, new User Role names must conform to the length and character requirements of the <b>name</b> property described in the "Base managed object properties schema" on page 39.  For the purpose of verifying uniqueness, this name is treated in a case-insensitive fashion when used to create a new User Role object.  The names of system-defined User Roles are listed in Appendix C, "Enum values for the User Role object," on page 933.
<b>description</b>	(w)(pc)	String (0-1024)	The description of the User Role object.  Default: an empty string

## | Class specific additional properties

| In addition to the properties defined via included schemas, this object includes the following additional class-specific properties:

| *Table 199. User Role object: class specific additional properties*

Name	Qualifier	Type	Description
<b>type</b>	—	String Enum	The type of User Role. Supported values are: <ul style="list-style-type: none"> <li>• <b>"user-defined"</b></li> <li>• <b>"system-defined"</b></li> </ul>

Table 199. User Role object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>associated-system-defined-user-role-uri</b>	(w)(pc)	String/URI	<p>Canonical URI path of the system-defined User Role associated with this User Role. The User Roles which are valid associated User Roles are identified in Appendix C, “Enum values for the User Role object,” on page 933. If this is a system-defined User Role, this property is null, because system-defined User Roles do not have an associated User Role.</p> <p>Default: the URI path of the Operator Tasks system-defined User Role.</p>
<b>permissions</b>	(c)(pc)	Array of objects	<p>The list of permissions included in this User Role. These permissions are identified by permission-info objects, as defined in the next table. The members provided in this list can change as a result of the <b>Add Permission to User Role</b> and <b>Remove Permission from User Role</b> operations. If there are no permissions in this User Role, an empty array is provided.</p> <p>Default: an empty array</p>
<b>is-inheritance-enabled</b>	(w)(pc)	Boolean	<p>Indicates whether User Role inheritance is enabled. When <b>true</b>, if this User Role permits access to a parent managed object, then all managed objects that are hosted by the parent managed object are also permitted by this role. When <b>false</b>, no such inherited permissions exist.</p> <p>The supported inheritance relationships are:</p> <p><b>Parent managed object:</b>  <b>Hosted managed object:</b>  <b>POWER® processor-based blade</b>  Virtual server</p> <p><b>System x blade</b>  Virtual server</p> <p><b>zBX BladeCenter</b></p> <ul style="list-style-type: none"> <li>• POWER processor-based blade</li> <li>• System x blade</li> <li>• IBM WebSphere® Datapower Integration Appliance XI50</li> </ul> <p>Default: <b>false</b>.</p>
<b>replication-overwrite-possible</b>	—	Boolean	<p>Indicates whether this object is customizable data that is replicated to this HMC from an HMC configured as a Data Source in the Data Replication service.</p>

Each nested permission-info object contains the following fields:

Table 200. permission-info object properties

Name	Type	Description
<b>permitted-object</b>	String/URI or String Enum	<p>Canonical URI path or String Enum which identifies the object(s) or task to which permission has been granted. Granting permission to a Task object gives the user action/task permission to the console task associated with that Task object, and the user is permitted to target any of their authorized objects with that task. If this identifies a class of managed objects, a specific managed object or a Group object, then permission is granted to view the details of the object(s) and target the object(s) with any authorized task for which it is an appropriate target. This field is one of the following:</p> <ul style="list-style-type: none"> <li>• The identifier for a class of managed objects. This identifier is the String Enum value for the class from Appendix B, "Enum values for a type of managed objects within User Roles," on page 931.</li> <li>• The URI of a specific managed object.</li> <li>• The URI of a Group object.</li> <li>• The URI of a Task object.</li> <li>• The well-known URI <code>"/api/system-manual-definition"</code>, which denotes a specific managed object known on the HMC UI as "System Manual Definition"<sup>1</sup></li> </ul> <p>The type of object(s) is indicated by the <b>permitted-object-type</b> property.</p>
<b>permitted-object-type</b>	String Enum	<p>Identifies the type of object(s) identified by the <b>permitted-object</b> property.</p> <ul style="list-style-type: none"> <li>• <b>"object"</b> - <b>permitted-object</b> contains a URI that identifies one of the following: <ul style="list-style-type: none"> <li>– A specific managed object</li> <li>– A Group object</li> <li>– A Task object.</li> </ul> </li> <li>• <b>"object-class"</b> - <b>permitted-object</b> contains a String Enum value from Appendix B, "Enum values for a type of managed objects within User Roles," on page 931.</li> </ul>
<b>include-members</b>	Boolean	<p>Indicates whether the members of a group are included in the User Role, or if only the group itself is included. True if members are included; false otherwise.</p> <p>Prerequisite: the <b>permitted-object</b> field identifies a Group object that does not use pattern-matching.</p>
<b>view-only-mode</b>	Boolean	<p>Indicates whether it is a task's view-only version that is in this User Role. Only certain tasks support a view-only mode. This field is only provided if the <b>permitted-object</b> field identifies such a Task object. A User Role cannot have both the view-only version and the non-view-only version in its set of permissions.</p> <p>Prerequisite: the <b>permitted-object</b> field identifies a Task object that supports a view-only mode.</p>

<sup>1</sup>The object identified by this special URI does not have full API support. It is only valid as a permitted object identifier in a User Role object. Thus it may be included in the **Get User Role Properties** response body, and it is valid in the request body of the **Add Permission to User Role** and **Remove Permission from User Role** operations.

## List User Roles

The **List User Roles** operation lists User Roles defined to the console.

## HTTP method and URI

GET /api/console/user-roles

### Query Parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property. Value must be a valid <b>type</b> property value.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
user-roles	Array of objects	Array of nested user-role-info objects as described in the next table.

Each nested user-role-info object contains the following fields:

Field name	Type	Description
object-uri	String/URI	The <b>object-uri</b> property of the User Role object.
name	String	The <b>name</b> property of the User Role object.
type	String Enum	The <b>type</b> property of the User Role object.

### Description

The **List User Roles** operation lists User Roles defined to the console. Some basic properties are provided for each user role.

If the **name** query parameter is specified, the returned list is limited to those User Roles that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not performed.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid User Role **type** property value. If the value is not valid, status code 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those User Roles that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not performed.

A User Role is included in the list only if the API user has user-related-access permission to that object or action/task permission to the **Manage User Roles** task. If there is a User Role to which the API user does not have permission, that object is omitted from the list, but no error status code results.

If there are no User Roles defined to the console or if no User Roles are to be included in the results due to filtering or access permissions, an empty list is provided and the operation completes successfully.

### Authorization requirements

This operation has the following authorization requirement:

- User-related-access permission to the User Role objects included in the response body or action/task permission to the **Manage User Roles** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 671.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage Notes

While it is intended that system-defined User Roles have a name that begins with “hmc-”, there is no such guarantee. API clients are cautioned to use the **type** property rather than **name** to reliably distinguish between system-defined and user-defined User Roles.

## Example HTTP interaction

---

```
GET /api/console/user-roles?name=.*ensemble.*task.* HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

Figure 328. List User Roles: Request



---

```

200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:32 GMT
content-type: application/json;charset=UTF-8
content-length: 589
{
  "user-roles":[
    {
      "name":"hmc-ensemble-management-facility-administrator-tasks",
      "object-uri":"/api/user-roles/ecfae496-40e7-4dd9-aa70-0a4fb299069a",
      "type":"system-defined"
    },
    {
      "name":"hmc-ensemble-administrator-tasks",
      "object-uri":"/api/user-roles/7d5a8716-6aad-4462-ad47-52cc3f898702",
      "type":"system-defined"
    },
    {
      "name":"Company ensemble task subset",
      "object-uri":"/api/user-roles/46ad96f0-4a6d-11e4-a1e7-1c6f65065a91",
      "type":"user-defined"
    },
    {
      "name":"hmc-ensemble-management-facility-operator-tasks",
      "object-uri":"/api/user-roles/84220072-d8a3-4af0-87f6-33eefbda498c",
      "type":"system-defined"
    }
  ]
}

```

---

Figure 329. List User Roles: Response

## Get User Role Properties

The **Get User Role Properties** operation retrieves the properties of a single User Role object that is designated by its object ID.

### HTTP method and URI

**GET** `/api/user-roles/{user-role-id}`

In this request, the URI variable `{user-role-id}` is the object ID of the User Role object whose properties are to be retrieved.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the User Role object as defined in the Data Model section. Field names and data types in the JSON object are the same as the property names and data types defined in the “Data model” on page 668.

### Description

This operation returns the current properties of a single User Role object that is designated by `{user-role-id}`.

On successful execution, all of the current properties as defined in the Data Model for the User Role object are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing User Role object and the API user must have user-related-access permission to it or action/task permission to the **Manage User Roles** task. If these conditions are not met, status code 404 (Not Found) is returned.

### Authorization requirements

This operation has the following authorization requirement:

- User-related-access permission to the User Role object specified in the request URI, or action/task permission to the **Manage User Roles** task.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 673.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type, or designates a resource for which the API user does not have the required authorization.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Usage Notes

While it is intended that system-defined User Roles have a name that begins with “hmc-”, there is no such guarantee. API clients are cautioned to use the **type** property rather than **name** to reliably distinguish between system-defined and user-defined User Roles.

### Example HTTP interaction

---

```
GET /api/user-roles/eb53f840-4a7a-11e4-affa-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

Figure 330. Get User Role Properties: Request

```

200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:32 GMT
content-type: application/json;charset=UTF-8
content-length: 853
{
  "associated-se-user-role-uri":null,
  "class":"user-role",
  "description":"Role for managing department business",
  "is-inheritance-enabled":false,
  "is-locked":false,
  "name":"Dept Admin",
  "object-id":"eb53f840-4a7a-11e4-affa-1c6f65065a91",
  "object-uri":"/api/user-roles/eb53f840-4a7a-11e4-affa-1c6f65065a91",
  "parent":"/api/console",
  "permissions":[
    {
      "permitted-object":"/api/console/tasks/4d5a39f0-c1df-4a2e-9a46-5bf6f6a759f3",
      "permitted-object-type":"object",
      "view-only-mode":false
    },
    {
      "include-members":true,
      "permitted-object":"/api/groups/cafb3a9b-6a34-4475-938f-98c5d60868a5",
      "permitted-object-type":"object"
    },
    {
      "permitted-object":"x-virtual-server",
      "permitted-object-type":"object-class"
    },
    {
      "permitted-object":"/api/console/tasks/45042443-7202-40b3-8630-b7a563a21d8d",
      "permitted-object-type":"object"
    }
  ],
  "replication-overwrite-possible":false,
  "type":"user-defined"
}

```

Figure 331. Get User Role Properties: Response

## Update User Role Properties

The **Update User Role Properties** operation updates the properties of a single user-defined User Role object that is designated by its object ID. System-defined User Roles are immutable; therefore, this operation is not valid for system-defined User Roles.

### HTTP method and URI

**POST** /api/user-roles/{user-role-id}

In this request, the URI variable {user-role-id} is the object ID of the User Role object whose properties are to be updated.

### Request body contents

The request body is expected to contain a JSON object that provides the new values of any writeable property that is to be updated by this operation. Field names and data types in this JSON object are expected to match the corresponding property names and data types defined by the data model for this object type. The JSON object can and should omit fields for properties whose values are not to be changed by this operation

## Description

- | This operation updates writeable properties of the User Role object specified by *{user-role-id}* .
- | The URI path must designate an existing User Role object; otherwise, status code 404 (Not Found) is returned. If the user does not have user-related-access permission to the designated User Role object or action/task permission to the **Manage User Roles** task, status code 404 (Not Found) is returned. If the user has user-related-access permission to the designated User Role object but not action/task permission to the **Manage User Roles** task, status code 403 (Forbidden) is returned.
- | The request body is validated against the schema described in the Request Body Contents section. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. An attempt to update a system-defined User Role is not valid and fails with status code 400 (Bad Request).
- | The request body does not need to specify a value for all writeable properties, but rather can and should contain fields only for the properties to be updated. Object properties for which no input value is provided remain unchanged by this operation.
- | If the update changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

- | This operation has the following authorization requirement:
  - Action/task permission to the **Manage User Roles** task.

## HTTP status and reason codes

- | On success, HTTP status code 204 (No Content) is returned and no response body is provided.
- | The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

| *Table 201. Update User Role Properties: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	314	This operation is not supported for an object of this type. System-defined User Roles may not be updated.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
	2	A URI in the request body does not designate an existing resource of the correct type.

- | Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
POST /api/user-roles/eb53f840-4a7a-11e4-affa-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 99
{
  "description":"A new and improved description of this User Role",
  "is-inheritance-enabled":true
}
```

Figure 332. Update User Role Properties: Request

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:32 GMT

<No response body>
```

Figure 333. Update User Role Properties: Response

## Add Permission to User Role

The **Add Permission to User Role** operation adds a specified permission to a specified user-defined User Role thereby granting that permission to all users that have that User Role. System-defined User Roles are immutable; therefore, this operation is not valid for system-defined User Roles.

### HTTP method and URI

**POST /api/user-roles/{user-role-id}/operations/add-permission**

In this request, the URI variable *{user-role-id}* is the object ID of the User Role to which a permission is to be added.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
permitted-object	String/ URI or String Enum	Required	<p>Canonical URI path or String Enum which identifies the object(s) or task to which permission is to be granted. See the Data Model for more information about permissions. This field is one of the following:</p> <ul style="list-style-type: none"> <li>The identifier for a class of managed objects. This identifier is the String Enum value for the class from Appendix B, "Enum values for a type of managed objects within User Roles," on page 931.</li> <li>The URI of a specific managed object.</li> <li>The URI of a Group object.</li> <li>The URI of a Task object.</li> </ul> <p>The type of object(s) is indicated by the <b>permitted-object-type</b> field.</p>

Field name	Type	Rqd/Opt	Description
permitted-object-type	String Enum	Required	Identifies the type of object(s) identified by the <b>permitted-object</b> field. Supported values are: <ul style="list-style-type: none"> <li>• <b>"object"</b> - <b>permitted-object</b> contains a URI that identifies one of the following: <ul style="list-style-type: none"> <li>– A specific managed object</li> <li>– A Group object</li> <li>– A Task object.</li> <li>– The well-known URI <code>"/api/system-manual-definition"</code>, which denotes a specific managed object known on the HMC UI as "System Manual Definition"</li> </ul> </li> <li>• <b>"object-class"</b> - <b>permitted-object</b> contains a String Enum value from Appendix B, "Enum values for a type of managed objects within User Roles," on page 931.</li> </ul>
include-members	Boolean	Optional	Indicates whether the members of the group are included in the User Role, or if only the group itself is included. True if members are included; false otherwise.  Prerequisite: the <b>permitted-object</b> field identifies a Group object that does not use pattern-matching.  Default: false
view-only-mode	Boolean	Optional	Indicates whether it is the task's view-only version that is being added to this User Role. Only certain tasks support a view-only mode. This field is only allowed if the <b>permitted-object</b> field identifies such a Task object.  Prerequisite: the <b>permitted-object</b> field identifies a Task object that supports a view-only mode.  Default: true

## Description

This operation adds a permission to a User Role

On successful execution of this operation the permission specified in the request body has been added to the User Role identified in the request URI.

The request body is validated against the schema described in the "Request body contents" on page 677. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. If the request URI does not designate an existing User Role object, status code 404 (Not Found) is returned. If the user does not have user-related-access permission to the designated User Role object or action/task permission to the **Manage User Roles** task, status code 404 (Not Found) is returned. If the user has user-related-access permission to the designated User Role object but not action/task permission to the **Manage User Roles** task, status code 403 (Forbidden) is returned. If the URI in the request body does not designate an existing resource, status code 404 (Not Found) is returned. An attempt to alter a system-defined User Role is not valid and fails with status code 400 (Bad Request).

If the specified permission is already in the User Role, or an attempt is made to have both the view-only and non-view-only versions of a Task in the User Role, status code 409 (Conflict) is returned.

If this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage User Roles** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 202. Add Permission to User Role: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	314	This operation is not supported for an object of this type. System-defined User Roles may not be altered.
	315	The permission identified in the request body is already in the User Role.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
	2	A URI in the request body does not designate an existing resource of the correct type.
409 (Conflict)	315	The permission identified in the request body is already in the User Role.
	327	The view-only and non-view-only versions of a Task cannot be in the same User Role. An attempt was made to add one version when the other was already in the User Role's set of permitted objects.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
POST /api/user-roles/eb53f840-4a7a-11e4-affa-1c6f65065a91/operations/
  add-permission HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 114
{
  "permitted-object":"/api/console/tasks/5d8c9f60-a2b3-4327-9fd4-791df8a60dcc",
  "permitted-object-type":"object"
}
```

Figure 334. Add Permission to User Role: Request

```

204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:32 GMT

```

<No response body>

Figure 335. Add Permission to User Role: Response

## Remove Permission from User Role

The **Remove Permission from User Role** operation removes a specified permission from a specified user-defined User Role. System-defined User Roles are immutable; therefore, this operation is not valid for system-defined User Roles.

### HTTP method and URI

**POST** /api/user-roles/{user-role-id}/operations/remove-permission

In this request, the URI variable {user-role-id} is the object ID of the User Role from which a permission is to be removed.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
permitted-object	String/ URI or String Enum	Required	<p>Canonical URI path or String Enum which identifies the object(s) or task to which permission is to be removed. See the Data Model for more information about permissions. This field is one of the following:</p> <ul style="list-style-type: none"> <li>The identifier for a class of managed object. This identifier is the String Enum value for the class from Appendix B, "Enum values for a type of managed objects within User Roles," on page 931.</li> <li>The URI of a specific managed object.</li> <li>The URI of a Group object.</li> <li>The URI of a Task object.</li> </ul> <p>The type of object(s) is indicated by the <b>permitted-object-type</b> field.</p>
permitted-object-type	String Enum	Required	<p>Identifies the type of object(s) identified by the <b>permitted-object</b> field. Supported values are:</p> <ul style="list-style-type: none"> <li><b>"object"</b> - <b>permitted-object</b> contains a URI that identifies one of the following: <ul style="list-style-type: none"> <li>A specific managed object.</li> <li>A Group object.</li> <li>A Task object.</li> <li>The well-known URI "/api/system-manual-definition", which denotes a specific managed object known on the HMC UI as "System Manual Definition"</li> </ul> </li> <li><b>"object-class"</b> - <b>permitted-object</b> contains a String Enum value from Appendix B, "Enum values for a type of managed objects within User Roles," on page 931.</li> </ul>
include-members	Boolean	Optional	<p>Indicates whether the members of the group are included in this operation, or if only the group itself is included. True if members are included; false otherwise.</p> <p>Prerequisite: the <b>permitted-object</b> field identifies a Group object.</p> <p>Default: false</p>



Field name	Type	Rqd/Opt	Description
view-only-mode	Boolean	Optional	<p>Indicates whether it is permission to a task's view-only version that is to be removed from this User Role. Only certain tasks support a view-only mode. This field is only allowed if the <b>permitted-object</b> field identifies such a Task object.</p> <p>Prerequisite: the <b>permitted-object</b> field identifies a Task object that supports a view-only mode.</p> <p>Default: true</p>

## Description

This operation removes a permission from a User Role

On successful execution of this operation the permission specified in the request body has been removed from the User Role identified in the request URI.

The request body is validated against the schema described in “Request body contents” on page 680. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. If the request URI does not designate an existing User Role object, status code 404 (Not Found) is returned. If the user does not have user-related-access permission to the designated User Role object or action/task permission to the **Manage User Roles** task, status code 404 (Not Found) is returned. If the user has user-related-access permission to the designated User Role object but not action/task permission to the **Manage User Roles** task, status code 403 (Forbidden) is returned. If the URI in the request body does not designate an existing resource, status code 404 (Not Found) is returned. An attempt to alter a system-defined User Role is not valid and fails with status code 400 (Bad Request).

If the specified permission is not in the User Role, or if removing it would leave one or more users without object-access permission to their default group, status code 409 (Conflict) is returned.

If this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage User Roles** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 203. Remove Permission from User Role: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	314	This operation is not supported for an object of this type. System-defined User Roles may not be altered.
	316	The permission identified in the request body is not in the User Role.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
	2	A URI in the request body does not designate an existing resource of the correct type.
409 (Conflict)	316	The permission identified in the request body is not in the User Role.
	328	The permission cannot be removed at this time, because doing so would leave one or more users without object-access permission to their default group.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
POST /api/user-roles/eb53f840-4a7a-11e4-affa-1c6f65065a91/operations/
  remove-permission HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 114
{
  "permitted-object":"/api/console/tasks/5d8c9f60-a2b3-4327-9fd4-791df8a60dcc",
  "permitted-object-type":"object"
}
```

Figure 336. Remove Permission from User Role: Request

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:32 GMT

<No response body>
```

Figure 337. Remove Permission from User Role: Response

## Create User Role

The **Create User Role** operation creates a user-defined User Role object with the given properties on the console. This operation is not valid for system-defined User Roles.

### HTTP method and URI

**POST** /api/console/user-roles

## Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The value to be set as the User Role's <b>name</b> property.
description	String	Optional	The value to be set as the User Role's <b>description</b> property.
associated-system-defined-user-role-uri	String/URI	Optional	The value to be set as the User Role's <b>associated-system-defined-user-role-uri</b> property.
is-inheritance-enabled	Boolean	Optional	The value to be set as the User Role's <b>is-inheritance-enabled</b> property.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the new User Role object.

## Description

This operation creates a new, empty, user-defined User Role.

On successful execution of this operation the User Role is created using the inputs as specified by the request body. The URI of the new User Role is provided in the response body and in a **Location** response header as well. An Inventory Change notification is emitted asynchronously. The **Add Permission to User Role** operation can then be used to add permissions to the new User Role.

The request body is validated against the schema described in “Request body contents.” If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. If the specified **name** is not unique, status code 400 (Bad Request) is returned. In addition, the API user must have action/task permission to the **Manage User Roles** task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage User Roles** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents,” and the **Location** response header contains the URI of the newly created object.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 204. Create User Role: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	8	A User Role with the name specified in the request body already exists.
403 (Forbidden)	1	The API user does not have the required permission for this operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19 and the Location response header contains the URI of the newly created object.

### Example HTTP interaction

```
POST /api/console/user-roles HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 78
{
  "description": "Role for managing department business",
  "name": "Dept Admin"
}
```

Figure 338. Create User Role: Request

```
201 Created
server: zSeries management console API web server / 2.0
location: /api/user-roles/eb53f840-4a7a-11e4-affa-1c6f65065a91
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:31 GMT
content-type: application/json;charset=UTF-8
content-length: 69
{
  "object-uri": "/api/user-roles/eb53f840-4a7a-11e4-affa-1c6f65065a91"
}
```

Figure 339. Create User Role: Response

### Delete User Role

The **Delete User Role** operation deletes a user-defined User Role object designated by its object ID. This operation is not valid for system-defined User Roles.

#### HTTP method and URI

```
DELETE /api/user-roles/{user-role-id}
```

In this request, the URI variable *{user-role-id}* is the object ID of the User Role object to be deleted.

#### Description

This operation removes a specified User Role from the console. The User Role is identified by the *{user-role-id}* variable in the URI.

Upon successfully removing the User Role, HTTP status code 204 (No Content) is returned and no response body is provided. An Inventory Change notification is emitted asynchronously.

The URI path must designate an existing User Role object; otherwise, status code 404 (Not Found) is returned. If the user does not have user-related-access permission to the designated User Role object or action/task permission to the **Manage User Roles** task, status code 404 (Not Found) is returned. If the user has user-related-access permission to the designated User Role object but not action/task permission to the **Manage User Roles** task, status code 403 (Forbidden) is returned. An attempt to delete a system-defined User Role is not valid and fails with status code 400 (Bad Request). If any user has the specified User Role, the request fails and status code 409 (Conflict) is returned.

### Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage User Roles** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 205. Delete User Role: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	312	This operation is not supported for an object of this type. System-defined User Roles may not be deleted.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
409 (Conflict)	317	The object cannot be deleted at this time. One or more users have this User Role.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
DELETE /api/user-roles/eb53f840-4a7a-11e4-affa-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

Figure 340. Delete User Role: Request

---

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:32 GMT
```

```
<No response body>
```

---

*Figure 341. Delete User Role: Response*

## | **Inventory service data**

| Information about the User Roles managed by the console can be optionally included in the inventory data provided by the Inventory Service.

| Inventory entries for User Role objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**user-role**" are to be included. An entry for a particular User Role is included only if the API user has access permission to that object as described in the **Get User Role Properties** operation.

| For each User Role object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the **Get User Role Properties** operation. That is, the data provided is the same as would be provided if a **Get User Role Properties** operation were requested targeting this object.

## | **Sample inventory data**

| The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single User Role. This object would appear as one array entry in the response array:

```

| {
|   "associated-system-defined-user-role-uri": "/api/user-roles/b39afb87-d915-
|     4070-a22f-91b158c6c01e",
|   "class": "user-role",
|   "description": "Role that allows management of users",
|   "is-inheritance-enabled": false,
|   "is-locked": false,
|   "name": "user_roles_9",
|   "object-id": "db7f9448-3737-11e4-a5fc-5ef3fcae8020",
|   "object-uri": "/api/user-roles/db7f9448-3737-11e4-a5fc-5ef3fcae8020",
|   "parent": "/api/console",
|   "permissions": [
|     {
|       "permitted-object": "/api/console/tasks/f8d653f4-eab2-4547-97c0-
|         a26f762218ba",
|       "permitted-object-type": "object"
|     },
|     {
|       "permitted-object": "/api/console/tasks/36e32fb4-7b60-4677-b462-
|         e786f337ea0f",
|       "permitted-object-type": "object"
|     },
|     {
|       "permitted-object": "/api/console/tasks/8ef2b7ca-c2d2-4a5b-9d52-
|         b4d1a28ccb15",
|       "permitted-object-type": "object"
|     }
|   ],
|   "replication-overwrite-possible": false,
|   "type": "user-defined"
| }

```

Figure 342. User Role object: Sample inventory data

## Task object

A Task object is an element of the console object and represents an action that a console user with appropriate authority can perform. These actions could be available via the console's graphical user interface, the Web Services APIs or both. Tasks are predefined by the console and cannot be created, modified or deleted.

All API users have object-access permission to all Tasks and thus are permitted to issue **List Tasks** and **Get Task Properties** for any Task object. Note that this object-access permission to a Task object does not give an API user action/task permission to the task represented by the Task object.

## Data model

The Task object contains the following properties.

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

Table 206. Task object: properties

Name	Qualifier	Type	Description of specialization
<b>element-uri</b>	—	String/URI	The canonical URI path of the Task object is of the form <code>/api/console/tasks/{task-id}</code> , where <code>{task-id}</code> is the value of the <b>element-id</b> property of the Task object.
<b>element-id</b>	—	String (36)	The unique identifier for this object.
<b>parent</b>	—	String/URI	The canonical URI path of the console object.

Table 206. Task object: properties (continued)

Name	Qualifier	Type	Description of specialization
class	—	String	The class of a Task object is "task".
name	(ro)	String Enum	The name of the Task object. The task names are documented in Appendix D, "Enum values for the Task object," on page 935.
description	—	String (0-1024)	The description of the Task object. Default: an empty string
view-only-mode-supported	—	Boolean	Indicates whether this task supports a view-only mode.

## List Tasks

The **List Tasks** operation lists Tasks defined to the console.

### HTTP method and URI

GET /api/console/tasks

### Query Parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
tasks	Array of objects	Array of nested task-info objects as described in the next table.

Each nested task-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	The <b>element-uri</b> property of the Task object.
name	String	The <b>name</b> property of the Task object.

## Description

This operation lists tasks defined to the console. Some basic properties are provided for each task.

If the **name** query parameter is specified the returned list is limited to those tasks that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not performed.

If no tasks are to be included in the results due to filtering, an empty list is provided and the operation completes successfully.



## Authorization requirements

This operation has no explicit authorization requirements.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 688.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/console/tasks?name=.*blade.* HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvvqx18c4r066ge9kcyzr4c
```

---

Figure 343. List Tasks: Request

---

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:32 GMT
content-type: application/json; charset=UTF-8
content-length: 219
{
  "tasks": [
    {
      "element-uri": "/api/console/tasks/94bc3309-b2f5-4a56-a1bc-399c60fd9fed",
      "name": "zbx-blade-details"
    },
    {
      "element-uri": "/api/console/tasks/99934f93-ab89-49fa-a01d-74412f48ab5d",
      "name": "zbx-bladecenter-details"
    }
  ]
}
```

---

Figure 344. List Tasks: Response

## Get Task Properties

The **Get Task Properties** operation retrieves the properties of a single Task object that is designated by its element ID.

### HTTP method and URI

```
GET /api/console/tasks/{task-id}
```

In this request, the URI variable *{task-id}* is the element ID of the Task object whose properties are to be returned.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the Task object as defined in the Data Model section. Field names and data types in the JSON object are the same as the property names and data types defined in the “Data model” on page 687.

### Description

This operation returns the current properties of a single Task object specified by *{task-id}*.

On successful execution, all of the current properties as defined in the Data Model for the Task object are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing Task object; otherwise, status code 404 (Not Found) is returned.

### Authorization requirements

This operation has no explicit authorization requirements.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/console/tasks/94bc3309-b2f5-4a56-a1bc-399c60fd9fed HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

Figure 345. Get Task Properties: Request

---

```

200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:32 GMT
content-type: application/json;charset=UTF-8
content-length: 259
{
  "class":"task",
  "description":"Details for a Blade",
  "element-id":"94bc3309-b2f5-4a56-a1bc-399c60fd9fed",
  "element-uri":"/api/console/tasks/94bc3309-b2f5-4a56-a1bc-399c60fd9fed",
  "name":"zbx-blade-details",
  "parent":"/api/console"
}

```

---

Figure 346. Get Task Properties: Response

## Inventory service data

Information about the tasks managed by the console can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for Task objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class **"console"** are to be included.

For each Task object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the **Get Task Properties** operation. That is, the data provided is the same as would be provided if a **Get Task Properties** operation were requested targeting this object.

## Sample inventory data

The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single Task. This object would appear as one array entry in the response array:

---

```

{
  "class": "task",
  "description": "Customize the Application Programming Interface for the
  console",
  "element-id": "ee76dca1-9aee-4530-bf66-667ae728ed10",
  "element-uri": "/api/console/tasks/ee76dca1-9aee-4530-bf66-667ae728ed10",
  "name": "customize-api-settings",
  "parent": "/api/console",
  "view-only-mode-supported": false
}

```

---

Figure 347. Task object: Sample inventory data

## User Pattern object

User Patterns and user templates allow a system administrator to define a group of console users at once whose user IDs all match a certain pattern (for example, a regular expression) and who have a certain set of attributes. A User Pattern object is an element of the console object and defines a pattern for user IDs that are not defined to the console but can be verified by an LDAP server for user authentication. Each pattern identifies a template User object which defines many characteristics of such users. A successful logon with a user ID that matches a User Pattern results in the creation of a pattern-based user, with

many of its attributes coming from the associated template. User Patterns are searched in a defined order during logon processing. That order can be customized through the Console object operation “Reorder User Patterns” on page 628.

Through user-related-access permission described in “User-related-access permission” on page 645, API users are permitted to see certain User Pattern objects in a **List User Patterns** response and issue **Get User Pattern Properties** for those User Pattern objects. An API user with action/task permission to the **Manage User Patterns** task is permitted to view and change any User Pattern object and change the pattern search order.

User Pattern objects may be replicated to this HMC via its Data Replication facility. If that is the case, the **Update User Pattern Properties** and **Delete User Pattern** operations should be used with care as they will prevent further replication of the modified or deleted object to this HMC.

## Data model

This object contains the following properties. Certain properties are only valid when mutable prerequisite properties have specific values. When such properties are not valid, their value is **null**. For instance the **template-name-override** is **null** when the **ldap-server-definition-uri** value is **null**.

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

Table 207. User Pattern object: properties

Name	Qualifier	Type	Description of specialization
<b>element-uri</b>	—	String/URI	The canonical URI path of the User Pattern object is of the form <code>/api/console/user-patterns/{user-pattern-id}</code> , where <code>{user-pattern-id}</code> is the value of the <b>element-id</b> property of the User Pattern object.
<b>element-id</b>	—	String (36)	The unique identifier for this object.
<b>parent</b>	—	String/URI	The canonical URI path of the console object.
<b>class</b>	—	String	The class of a User Pattern object is <b>"user-pattern"</b> .
<b>name</b>	(w)	String	The name of the User Pattern object. This name must be unique among all User Patterns on the console. The length and character requirements on this property are the same as those of the <b>name</b> property described in the “Base managed object properties schema” on page 39.  For the purpose of verifying uniqueness, this name is treated in a case-insensitive fashion when used to create a new User Pattern object.
<b>description</b>	(w)	String (0-1024)	The description of the User Pattern object.  Default: an empty string
<b>pattern</b>	(w)	String	The actual User Pattern expression. The combination of <b>pattern</b> and <b>type</b> must be unique on the console. Must not be an empty string.
<b>type</b>	(w)	String Enum	The style in which this pattern is expressed, which is one of the following values: <ul style="list-style-type: none"> <li><b>"glob-like"</b> - The pattern may include limited special characters: an asterisk in the pattern matches zero or more characters in the user ID, and a question mark in the pattern matches any single character in the user ID.</li> <li><b>"regular-expression"</b> - The pattern is expressed as a UNIX regular expression.</li> </ul>

Table 207. User Pattern object: properties (continued)

Name	Qualifier	Type	Description of specialization
<b>search-order-index</b>	—	Integer	A zero-based index position of this User Pattern in the pattern search order used during logon processing.
<b>retention-time</b>	(w)	Integer (0-2147483647)	The time in days that the user data for pattern-based users created based on this pattern will be retained. A value of 0 indicates not to retain the settings.
<b>user-template-uri</b>	(w)	String/URI	The canonical URI path of the User object containing the template definition that applies when a successful match occurs using this User Pattern.
<b>ldap-server-definition-uri</b>	(w)	String/URI	The canonical URI path of the LDAP Server Definition object which identifies the LDAP server to be used to validate the user ID when processing a logon for a user ID that matches this pattern, or null if that validation is to be performed by the console itself.
<b>template-name-override</b>	(w)	String	The name of the LDAP attribute that contains the name of the user template definition for the user, or <b>null</b> if there is no such attribute. When not <b>null</b> , this property overrides the user template identified in the <b>user-template-uri</b> field, and it must not be an empty string.  Prerequisite: <b>ldap-server-definition-uri</b> is not <b>null</b>  Default: <b>null</b>
<b>domain-name-restrictions</b>	(w)	String	The name of the LDAP attribute that contains the information about which consoles the user is allowed to log onto, or <b>null</b> if there is no such attribute. Must not be an empty string.  Prerequisite: <b>ldap-server-definition-uri</b> is not <b>null</b>  Default: <b>null</b>
<b>replication-override-possible</b>	—	Boolean	Indicates whether this object is customizable data that is replicated to this HMC from an HMC configured as a Data Source in the Data Replication service.

## List User Patterns

The **List User Patterns** operation lists User Patterns defined to the console.

### HTTP method and URI

**GET** /api/console/user-patterns

### Query Parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property. Value must be a valid <b>type</b> property value.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
user-patterns	Array of objects	Array of nested user-pattern-info objects as described in the next table.

Each nested user-pattern-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	The <b>element-uri</b> property of the User Pattern object.
name	String	The <b>name</b> property of the User Pattern object.
type	String Enum	The <b>type</b> property of the User Pattern object.

## Description

This operation lists User Patterns defined to the console. Some basic properties are provided for each User Pattern.

A User Pattern is included in the list only if the API user has user-related-access permission to that object or action/task permission to the **Manage User Patterns** task. If there is a User Pattern to which the API user does not have permission, that object is omitted from the list, but no error status code results.

If there are no User Patterns defined to the console or if no User Patterns are to be included in the results due to filtering or access permissions, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirement:

- User-related-access permission to the User Pattern objects included in the response body or action/task permission to the **Manage User Patterns** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage Notes

User Patterns are searched in a defined order during logon processing. That order can be customized through the Console object operation “Reorder User Patterns” on page 628. The **List User Patterns** operation does not specify the order in which the User Pattern URIs appear in the response body, and there is no guarantee that the order in the response will not change in subsequent invocations. Use the **search-order-index** property to determine a pattern’s position in the search order.

## Example HTTP interaction

---

```
GET /api/console/user-patterns?type=regular-expression HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

Figure 348. List User Patterns: Request

---

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT
content-type: application/json;charset=UTF-8
content-length: 314
{
  "user-patterns":[
    {
      "element-uri":"/api/console/user-patterns/497bf4ec-1dbf-11e4-8ceb-1c6f65065a91",
      "name":"IBM Intranet User Pattern",
      "type":"regular-expression"
    },
    {
      "element-uri":"/api/console/user-patterns/ec5b012a-4a7a-11e4-8777-1c6f65065a91",
      "name":"Company email pattern",
      "type":"regular-expression"
    }
  ]
}
```

---

Figure 349. List User Patterns: Response

## Get User Pattern Properties

The **Get User Pattern Properties** operation retrieves the properties of a single User Pattern object that is designated by its element ID.

### HTTP method and URI

```
GET /api/console/user-patterns/{user-pattern-id}
```

In this request, the URI variable *{user-pattern-id}* is the element ID of the User Pattern object whose properties are to be returned.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the User Pattern object as defined in the “Data model” on page 692. Field names and data types in the JSON object are the same as the property names and data types defined in the Data Model.

## Description

- | This operation returns the current properties of a single User Pattern object specified by *{user-pattern-id}*.
- | On successful execution, all of the current properties as defined in the Data Model for the User Pattern object are provided in the response body, and HTTP status code 200 (OK) is returned.
- | The URI path must designate an existing User Pattern object and the API user must have user-related-access permission to it or action/task permission to the **Manage User Patterns** task. If these conditions are not met, status code 404 (Not Found) is returned.

## Authorization requirements

- | This operation has the following authorization requirement:
  - | • User-related-access permission to the User Pattern object specified in the request URI or action/task permission to the **Manage User Patterns** task

## HTTP status and reason codes

- | On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 695.
- | The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type, or designates a resource for which the API user does not have the required authorization.

- | Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/console/user-patterns/ec5b012a-4a7a-11e4-8777-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

*Figure 350. Get User Pattern Properties: Request*



---

```

200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT
content-type: application/json;charset=UTF-8
content-length: 593
{
  "class":"user-pattern",
  "description":"User Pattern based on company email addresses",
  "domain-name-restrictions":null,
  "element-id":"ec5b012a-4a7a-11e4-8777-1c6f65065a91",
  "element-uri":"/api/console/user-pattern
s/ec5b012a-4a7a-11e4-8777-1c6f65065a91",
  "ldap-server-definition-uri":null,
  "name":"Company email pattern",
  "parent":"/api/console",
  "pattern":".*@our\\.company\\.com",
  "replication-overwrite-possible":false,
  "retention-time":8,
  "search-order-index":3,
  "template-name-override":null,
  "type":"regular-expression",
  "user-template-uri":"/api/users/ec473e56-4a7a-11e4-91ee-1c6f65065a91"
}

```

---

Figure 351. Get User Pattern Properties: Response

## Update User Pattern Properties

The **Update User Pattern Properties** operation updates the properties of a single User Pattern object that is designated by its element ID.

### HTTP method and URI

**POST** `/api/console/user-patterns/{user-pattern-id}`

In this request, the URI variable `{user-pattern-id}` is the element ID of the User Pattern object whose properties are to be updated.

### Request body contents

The request body is expected to contain a JSON object that provides the new values of any writeable property that is to be updated by this operation. Field names and data types in this JSON object are expected to match the corresponding property names and data types defined by the data model for this object type. The JSON object can and should omit fields for properties whose values are not to be changed by this operation.

### Description

This operation updates writeable properties of the User Pattern object specified by `{user-pattern-id}`.

The URI path must designate an existing User Pattern object; otherwise, status code 404 (Not Found) is returned. If the user does not have user-related-access permission to the designated User Pattern object or action/task permission to the **Manage User Patterns** task, status code 404 (Not Found) is returned. If the user has user-related-access permission to the designated User Pattern object but not action/task permission to the **Manage User Patterns** task, status code 403 (Forbidden) is returned.

The request body is validated against the schema described in the Request Body Contents section. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the

| validation error encountered. The request body validation will fail if it contains a property that is not valid because a prerequisite is not met (e.g., attempting to set **template-name-override** when the **ldap-server-definition-uri** value is **null**).

| The request body does not need to specify a value for all writeable properties, but rather can and should contain fields only for the properties to be updated. Object properties for which no input value is provided and no prerequisite property is changed remain unchanged by this operation. A property's value is set to its default value if the field is not included in the request body and a prerequisite field is changed such that the prerequisite condition becomes satisfied (e.g., if **ldap-server-definition-uri** is changed from non-null to **null**, and **template-name-override** is not defined in the request body, **template-name-override** will be defaulted to **null**).

### | Authorization requirements

| This operation has the following authorization requirement:

- | • Action/task permission to the **Manage User Patterns** task.

### | HTTP status and reason codes

| On success, HTTP status code 204 (No Content) is returned and no response body is provided.

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

| *Table 208. Update User Pattern Properties: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
	324	The <b>ldap-server-definition-uri</b> field in the request body does not designate an existing LDAP Server Definition object.
	326	The <b>user-template-uri</b> field in the request body does not designate an existing template type User object.

| Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### | Example HTTP interaction

---

```
POST /api/console/user-patterns/ec5b012a-4a7a-11e4-8777-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 95
{
  "description":"A new and improved description of this User Pattern",
  "retention-time":30
}
```

---

*Figure 352. Update User Pattern Properties: Request*

```

204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT

```

<No response body>

Figure 353. Update User Pattern Properties: Response

## Create User Pattern

The **Create User Pattern** operation creates a User Pattern object with the given properties on the console.

### HTTP method and URI

**POST** /api/console/user-patterns

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The value to be set as the User Pattern's <b>name</b> property.
description	String	Optional	The value to be set as the User Pattern's <b>description</b> property.
pattern	String	Required	The value to be set as the User Pattern's <b>pattern</b> property.
type	String Enum	Required	The value to be set as the User Pattern's <b>type</b> property.
retention-time	Integer	Required	The value to be set as the User Pattern's <b>retention-time</b> property.
user-template-uri	String/URI	Required	The value to be set as the User Pattern's <b>user-template-uri</b> property.
ldap-server-definition-uri	String/URI	Optional	The value to be set as the User Pattern's <b>ldap-server-definition-uri</b> property.
template-name-override	String	Optional	The value to be set as the User Pattern's <b>template-name-override</b> property.
domain-name-restrictions	String	Optional	The value to be set as the User Pattern's <b>domain-name-restrictions</b> property.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the new User Pattern object.

### Description

This operation creates a new User Pattern.

On successful execution of this operation the User Pattern is created using the inputs as specified by the request body. The new User Pattern is placed at the end of the pattern search order used during logon processing. The URI of the new User Pattern is provided in the response body and in a **Location** response header as well.

The request body is validated against the schema described in the “Request body contents” on page 699. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. The request body validation will fail if it contains a property that is not valid because a prerequisite is not met (e.g., specifying **template-name-override** when the **ldap-server-definition-uri** value is **null**) or the specified pattern is not unique. If a URI in the request body does not designate an existing resource of the appropriate type, status code 404 (Not Found) is returned. In addition, the API user must have action/task permission to the **Manage User Patterns** task; otherwise, status code 403 (Forbidden) is returned.

### Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage User Patterns** task.

### HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents” on page 699, and the **Location** response header contains the URI of the newly created object.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 209. Create User Pattern: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	8	A User Pattern with the pattern specified in the request body already exists.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	324	The <b>ldap-server-definition-uri</b> field in the request body does not designate an existing LDAP Server Definition object..
	326	The <b>user-template-uri</b> field in the request body does not designate an existing template type User object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Usage Notes

User Patterns are searched in a defined order during logon processing. That order can be customized through the Console object's **Reorder User Patterns** operation.

## Example HTTP interaction

---

```
POST /api/console/user-patterns HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 260
{
  "description":"User Pattern based on company email addresses",
  "name":"Company email pattern",
  "pattern":".*@our\\.company\\.com",
  "retention-time":8,
  "type":"regular-expression",
  "user-template-uri":"/api/users/ec473e56-4a7a-11e4-91ee-1c6f65065a91"
}
```

---

Figure 354. Create User Pattern: Request

---

```
201 Created
server: zSeries management console API web server / 2.0
location: /api/console/user-patterns/ec5b012a-4a7a-11e4-8777-1c6f65065a91
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT
content-type: application/json;charset=UTF-8
content-length: 83
{
  "element-uri":"/api/console/user-patterns/ec5b012a-4a7a-11e4-8777-1c6f65065a91"
}
```

---

Figure 355. Create User Pattern: Response

## Delete User Pattern

The **Delete User Pattern** operation deletes a User Pattern object designated by its element ID.

### HTTP method and URI

**DELETE** /api/console/user-patterns/{*user-pattern-id*}

In this request, the URI variable {*user-pattern-id*} is the element ID of the User Pattern object to be deleted.

### Description

This operation removes a specified User Pattern from the console. The User Pattern is identified by the {*user-pattern-id*} variable in the URI.

Upon successfully removing the User Pattern, HTTP status code 204 (No Content) is returned and no response body is provided.

The URI path must designate an existing User Pattern object; otherwise, status code 404 (Not Found) is returned. If the user does not have user-related-access permission to the designated User Pattern object or action/task permission to the **Manage User Patterns** task, status code 404 (Not Found) is returned. If the user has user-related-access permission to the designated User Pattern object but not action/task permission to the **Manage User Patterns** task, status code 403 (Forbidden) is returned.

### Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage User Patterns** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage Note

It is permitted to delete a User Pattern even if there is a pattern-based user based on that pattern logged onto the HMC at the time of the deletion. Note that this will cause certain operations issued with the value of the **user-pattern-uri** property in that user's User object to fail, most likely with a 404 (Not Found) status code.

## Example HTTP interaction

---

```
DELETE /api/console/user-patterns/ec5b012a-4a7a-11e4-8777-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

Figure 356. Delete User Pattern: Request

---

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT
```

<No response body>

---

Figure 357. Delete User Pattern: Response

## Inventory service data

Information about the User Patterns managed by the console can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for User Pattern objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**console**" are to be included. An entry for a particular User Pattern is included only if the API user has access permission to that object as described in the **Get User Pattern Properties** operation.

For each User Pattern object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the **Get User Pattern Properties** operation. That is, the data provided is the same as would be provided if a **Get User Pattern Properties** operation were requested targeting this object.

### Sample inventory data

The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single LDAP Server Definition. This object would appear as one array entry in the response array:

---

```

{
  "class": "user-pattern",
  "description": "",
  "domain-name-restrictions": null,
  "element-id": "8c1ae17a-34c4-11e4-a30c-5ef3fcae8020",
  "element-uri": "/api/console/user-patterns/8c1ae17a-34c4-11e4-a30c-5ef3fcae8020",
  "ldap-server-definition-uri": "/api/console/ldap-server-definitions/4927787e-34c4-11e4-a1ea-5ef3fcae8020",
  "name": "testpattern_inventory",
  "parent": "/api/console",
  "pattern": "user-*",
  "replication-overwrite-possible": false,
  "retention-time": 0,
  "search-order-index": 0,
  "template-name-override": null,
  "type": "glob-like",
  "user-template-uri": "/api/users/641c7d96-34c4-11e4-b2c2-5ef3fcae8020"
}

```

---

Figure 358. User Pattern object: Sample inventory data

## Password Rule object

A Password Rule object is an element of the console object and represents a rule which a console user(s) must follow when creating a console logon password. Each console user using local authentication is assigned a password rule. There are certain system-defined password rules available for use.

Through user-related-access permission described in “User-related-access permission” on page 645, API users are permitted to see certain Password Rule objects in a **List Password Rules** response and issue **Get Password Rule Properties** for those Password Rule objects. An API user with action/task permission to the **Manage Password Rules** task is permitted to view any Password Rule object and change any user-defined Password Rule object.

Password rule objects may be replicated to this HMC via its Data Replication facility. If that is the case, the **Update Password Rule Properties** and **Delete Password Rule** operations should be used with care as they will prevent further replication of the modified or deleted object to this HMC.

### System-defined password rules

Unlike user-defined password rules, the system-defined password rules may not be modified. While system-defined password rules can be deleted and their name reused for a user-defined password rule, that practice is discouraged due to the likely confusion such a situation would cause. The names of the typical system-defined password rules include:

- Basic
- Standard

- Strict

## Password rule parts

Password rule parts are optional requirements to be applied to individual parts of a password. These requirements are applied, in order, to the password, from left to right. Each of these requirements must be met by some part of the password in order for the password to meet all of the requirements of the Password Rule.

For example, to require a password to consist of 1-3 letters followed by a 4 or 5 digit number, two rule parts are defined. The first rule part requires from 1 to 3 characters, each of which must be alphabetic; the second rule part requires from 4 to 5 characters, each of which must be numeric. Passwords such as “pa1600” and “Hey90210” meet the requirements of both of those rule parts.

## Data model

The Password Rule object contains the following properties.

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

Table 210. Password Rule object: properties

Name	Qualifier	Type	Description of specialization
<b>element-uri</b>	—	String/URI	The canonical URI path for a Password Rule object is of the form <code>/api/console/password-rules/{password-rule-id}</code> , where <code>{password-rule-id}</code> is the value of the <b>element-id</b> property of the Password Rule object.
<b>element-id</b>	—	String (36)	The unique identifier for this object.
<b>parent</b>	—	String/URI	The canonical URI path of the console object.
<b>class</b>	—	String	The class of a Password Rule object is " <b>password-rule</b> ".
<b>name</b>	(ro)	String	The name of the Password Rule object. This name must be unique among all password rules on the console. While preexisting Password Rule names are virtually unrestricted in terms of length and characters, new Password Rule names must conform to the length and character requirements of the <b>name</b> property described in the “Base managed object properties schema” on page 39.  For the purpose of verifying uniqueness, this name is treated in a case-insensitive fashion when used to create a new Password Rule object.
<b>description</b>	(w)	String (0-1024)	The description of the Password Rule object.  Default: an empty string
<b>type</b>	—	String Enum	Identifies the type of password rule. It must be one of the following values: <ul style="list-style-type: none"> <li>• "<b>system-defined</b>" - A password rule defined by the system. System-defined rules may not be modified.</li> <li>• "<b>user-defined</b>" - A password rule defined by a user.</li> </ul>
<b>expiration</b>	(w)	Integer	The total number of days a password is valid before it expires. A value of 0 indicates that the password never expires.  Default: 0
<b>min-length</b>	(w)	Integer (1-256)	The minimum required length of the password. Cannot be greater than <b>max-length</b> .  Default: 8



Table 210. Password Rule object: properties (continued)

Name	Qualifier	Type	Description of specialization
<b>max-length</b>	(w)	Integer (1-256)	The maximum allowed length of the password. Cannot be less than <b>min-length</b> .  Default: 256
<b>consecutive-characters</b>	(w)	Integer	The maximum number of characters that are allowed to be repeated in a row. A value of 0 indicates that there is no such limit.  Default: 0
<b>similarity-count</b>	(w)	Integer	The maximum number of consecutive characters in the current password that can match consecutive characters in the previous password. A value of 0 indicates that there is no such limit.  Default: 0
<b>history-count</b>	(w)	Integer	The number of previous passwords to which a new password is compared for uniqueness. A value of 0 indicates that there is no such comparison.  Default: 0
<b>case-sensitive</b>	(w)	Boolean	Indicates whether the password is case sensitive.  Default: <b>false</b>
<b>character-rules</b>	(w)	Array of objects	Optional rules to be applied to individual parts of the password. These rules are applied, in order, to the password, from left to right. Each of these rules must be met by some part of the password in order for the password to meet the requirements of this Password Rule. This property is an array of nested character-rule objects as described in the next table. If there are no rule parts, an empty array is provided.  Default: <empty array>
<b>replication-override-possible</b>	—	Boolean	Indicates whether this object is customizable data that is replicated to this HMC from an HMC configured as a Data Source in the Data Replication service.

Each nested password-rule-part object contains the following fields:

Table 211. character-rule object properties

Name	Type	Description
<b>min-characters</b>	Integer	The minimum number of characters required by this password rule part. Must be at least 1, and cannot be greater than <b>max-characters</b> .
<b>max-characters</b>	Integer	The maximum number of characters allowed by this password rule part. Must be at least 1, and cannot be less than <b>min-characters</b> .
<b>alphabetic</b>	String Enum	This field determines the inclusion of alphabetic characters within this part of the password. It must be one of the following values: <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - There can be alphabetic characters.</li> <li>• <b>"not-allowed"</b> - There cannot be alphabetic characters.</li> <li>• <b>"required"</b> - There must be alphabetic characters.</li> </ul>
<b>numeric</b>	String Enum	This field determines the inclusion of numeric characters within this part of the password. It must be one of the following values: <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - There can be numeric characters.</li> <li>• <b>"not-allowed"</b> - There cannot be numeric characters.</li> <li>• <b>"required"</b> - There must be numeric characters.</li> </ul>

Table 211. *character-rule object properties (continued)*

Name	Type	Description
<b>special</b>	String Enum	<p>This field determines the inclusion of special characters within this part of the password. It must be one of the following values:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - There can be special characters.</li> <li>• <b>"not-allowed"</b> - There cannot be special characters.</li> <li>• <b>"required"</b> - There must be special characters.</li> </ul> <p>Special characters include: greater than (&gt;), less than (&lt;), tilde (~), exclamation mark (!), at sign (@), number sign (#), question mark (?), dollar sign (\$), vertical bar ( ), percent sign (%), caret (^), ampersand (&amp;), asterisk (*), left and right parentheses ( ), underscore (_), plus sign (+), hyphen (-), equals sign (=), left and right curly braces ( { } ), left and right square brackets ( [ ] ), back slash (\), forward slash (/), period (.), comma (,), colon (:), accent (^), quotation mark ("), semicolon (;), and apostrophe (').</p>
<b>custom-character-sets</b>	Array of objects	Optional specific character requirements for this password part, as specified in an array of nested custom-character-set objects defined in the next table. This allows the specification of custom character sets and their inclusion requirement. There can be up to 2 custom character sets for a rule part. If none are defined, an empty array is provided.

Each nested specific-property object contains the following fields:

Table 212. *custom-character-set object properties*

Name	Type	Description
<b>character-set</b>	String	A string consisting of the characters that comprise this custom character set.
<b>inclusion</b>	String Enum	<p>This field determines the inclusion of characters in this character set within part of the password. It must be one of the following values:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Characters can be included.</li> <li>• <b>"not-allowed"</b> - Characters cannot be included.</li> <li>• <b>"required"</b> - At least one character must be included.</li> </ul>

## List Password Rules

The **List Password Rules** operation lists Password Rules defined to the console.

### HTTP method and URI

**GET** /api/console/password-rules

#### Query Parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.
type	String Enum	Optional	Filter string to limit returned objects to those that have a matching <b>type</b> property. Value must be a valid <b>type</b> property value.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
password-rules	Array of objects	Array of nested password-rule-info objects as described in the next table. If no Password Rules are to be returned, an empty array is provided.

Each nested password-rule-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	The <b>element-uri</b> property of the Password Rule object.
name	String	The <b>name</b> property of the Password Rule object.
type	String Enum	The <b>type</b> property of the Password Rule object.

## Description

This operation lists Password Rules defined to the console. Some basic properties are provided for each Password Rule.

If the **name** query parameter is specified the returned list is limited to those Password Rules that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not performed.

If the **type** query parameter is specified, the parameter is validated to ensure it is a valid Password Rule **type** property value. If the value is not valid, status code 400 (Bad Request) is returned. If the value is valid, the returned list is limited to those Password Rules that have a **type** property matching the specified value. If the **type** parameter is omitted, this filtering is not performed.

A Password Rule is included in the list only if the API user has user-related-access permission to that object or action/task permission to the **Manage Password Rules** task. If there is a console Password Rule to which the API user does not have permission, that object is omitted from the list, but no error status code results.

If there are no Password Rules defined to the console or if no Password Rules are to be included in the results due to filtering or access permissions, an empty list is provided and the operation completes successfully.

## Authorization requirements

This operation has the following authorization requirements:

- User-related-access permission to the Password Rules objects included in the response body or action/task permission to the **Manage Password Rules** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
GET /api/console/password-rules?type=system-defined HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

Figure 359. List Password Rules: Request

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT
content-type: application/json;charset=UTF-8
content-length: 390
{
  "password-rules":[
    {
      "element-uri":"/api/console/password-rules/4a790766-3dbf-11e4-980d-1c6f65065a91",
      "name":"Basic",
      "type":"system-defined"
    },
    {
      "element-uri":"/api/console/password-rules/4a79360a-3dbf-11e4-980d-1c6f65065a91",
      "name":"Standard",
      "type":"system-defined"
    },
    {
      "element-uri":"/api/console/password-rules/4a792b24-3dbf-11e4-980d-1c6f65065a91",
      "name":"Strict",
      "type":"system-defined"
    }
  ]
}
```

Figure 360. List Password Rules: Response

### Get Password Rule Properties

The **Get Password Rule Properties** operation retrieves the properties of a single Password Rule object that is designated by its element ID.

#### HTTP method and URI

```
GET /api/console/password-rules/{password-rule-id}
```

In this request, the URI variable *{password-rule-id}* is the element ID of the Password Rule object whose properties are to be returned.

## Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the Password Rule object as defined in the “Data model” on page 704. Field names and data types in the JSON object are the same as the property names and data types defined in the Data Model.

## Description

This operation returns the current properties of a single Password Rule object that is designated by *{password-rule-id}*.

On successful execution, all of the current properties as defined in the Data Model for the Password Rule object are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing Password Rule object and the API user must have user-related-access permission to it or action/task permission to the **Manage Password Rules** task. If these conditions are not met, status code 404 (Not Found) is returned.

## Authorization requirements

This operation has the following authorization requirement:

- User-related-access permission to the Password Rule object specified in the request URI or action/task permission to the **Manage Password Rules** task

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type, or designates a resource for which the API user does not have the required authorization.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/console/password-rules/ecb26fb4-4a7a-11e4-affa-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

Figure 361. Get Password Rule Properties: Request

---

```

200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT
content-type: application/json;charset=UTF-8
content-length: 656
{
  "case-sensitive":true,
  "character-rules":[
    {
      "alphanumeric":"not-allowed",
      "custom-character-sets":[
        {
          "character-set":"*!^",
          "inclusion":"not-allowed"
        }
      ],
      "max-characters":256,
      "min-characters":8,
      "numeric":"not-allowed",
      "special":"required"
    }
  ],
  "class":"password-rule",
  "consecutive-characters":0,
  "description":"Password must be very special",
  "element-id":"ecb26fb4-4a7a-11e4-affa-1c6f65065a91",
  "element-uri":"/api/console/password-rules/ecb26fb4-4a7a-11e4-affa-1c6f65065a91",
  "expiration":0,
  "history-count":0,
  "max-length":256,
  "min-length":8,
  "name":"All specials",
  "parent":"/api/console",
  "replication-overwrite-possible":false,
  "similarity-count":0,
  "type":"user-defined"
}

```

---

Figure 362. Get Password Rule Properties: Response

## | Update Password Rule Properties

| The **Update Password Rule Properties** operation updates the properties of a single user-defined Password Rule object that is designated by its element ID.

### | HTTP method and URI

| **POST** /api/console/password-rules/{password-rule-id}

| In this request, the URI variable {password-rule-id} is the element ID of the Password Rule object whose properties are to be updated.

### | Request body contents

| The request body is expected to contain a JSON object that provides the new values of any writeable property that is to be updated by this operation. Field names and data types in this JSON object are expected to match the corresponding property names and data types defined by the data model for this object type. The JSON object can and should omit fields for properties whose values are not to be changed by this operation.

## Description

- | This operation updates writeable properties of the Password Rule object specified by *{password-rule-id}* .
- | The URI path must designate an existing Password Rule object; otherwise, status code 404 (Not Found) is returned. If the user does not have user-related-access permission to the designated Password Rule object or action/task permission to the **Manage Password Rules** task, status code 404 (Not Found) is returned.
- | If the user has user-related-access permission to the designated Password Rule object but not action/task permission to the **Manage Password Rules** task, status code 403 (Forbidden) is returned
- | The request body is validated against the schema described in the Request Body Contents section. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. An attempt to update a system-defined password rule is not valid and fails with status code 400 (Bad Request).
- | The request body does not need to specify a value for all writeable properties, but rather can and should contain fields only for the properties to be updated. Object properties for which no input value is provided remain unchanged by this operation.

## Authorization requirements

- | This operation has the following authorization requirement:
  - Action/task permission to the **Manage Password Rules** task.

## HTTP status and reason codes

- | On success, HTTP status code 204 (No Content) is returned and no response body is provided.
- | The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

*Table 213. Update Password Rule Properties: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	314	This operation is not supported for an object of this type. System-defined password rules may not be updated.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.

- | Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
POST /api/console/password-rules/ecb26fb4-4a7a-11e4-affa-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 109
{
  "description":"A new and improved description of this Password Rule",
  "expiration":90,
  "history-count":5
}
```

Figure 363. Update Password Rule Properties: Request

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT

<No response body>
```

Figure 364. Update Password Rule Properties: Response

## Create Password Rule

The **Create Password Rule** operation creates a user-defined Password Rule object with the given properties.

### HTTP method and URI

**POST** /api/console/password-rules

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The value to be set as the Password Rule's <b>name</b> property.
description	String	Optional	The value to be set as the Password Rule's <b>description</b> property.
expiration	Integer	Optional	The value to be set as the Password Rule's <b>expiration</b> property.
min-length	Integer	Optional	The value to be set as the Password Rule's <b>min-length</b> property.
max-length	Integer	Optional	The value to be set as the Password Rule's <b>max-length</b> property.
consecutive-characters	Integer	Optional	The value to be set as the Password Rule's <b>consecutive-characters</b> property.
similarity-count	Integer	Optional	The value to be set as the Password Rule's <b>similarity-count</b> property.
history-count	Integer	Optional	The value to be set as the Password Rule's <b>history-count</b> property.
case-sensitive	Boolean	Optional	The value to be set as the Password Rule's <b>case-sensitive</b> property.



Field name	Type	Rqd/Opt	Description
character-rules	Array of objects	Optional	The value to be set as the Password Rule's <b>character-rules</b> property.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the new Password Rule object.

## Description

This operation creates a new user-defined Password Rule.

On successful execution of this operation the Password Rule is created using the inputs as specified by the request body. The URI of the new Password Rule is provided in the response body and in a **Location** response header as well.

The request body is validated against the schema described in the “Request body contents” on page 712. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. If the specified name is not unique, status code 400 (Bad Request) is returned. In addition, the API user must have action/task permission to the Manage Password Rules task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage Password Rules** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents” and the Location response header contains the URI of the newly created object.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 214. Create Password Rule: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	8	A password rule with the name specified in the request body already exists.
403 (Forbidden)	1	The API user does not have the required permission for this operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```

POST /api/console/password-rules HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 298
{
  "character-rules":[
    {
      "alphabetic":"not-allowed",
      "custom-character-sets":[
        {
          "character-set":"*!^",
          "inclusion":"not-allowed"
        }
      ],
      "max-characters":256,
      "min-characters":8,
      "numeric":"not-allowed",
      "special":"required"
    }
  ],
  "description":"Password must be very special",
  "name":"All specials"
}

```

Figure 365. Create Password Rule: Request

```

201 Created
server: zSeries management console API web server / 2.0
location: /api/console/password-rules/ecb26fb4-4a7a-11e4-affa-1c6f65065a91
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT
content-type: application/json;charset=UTF-8
content-length: 82
{
  "element-uri":"/api/console/password-rules/ecb26fb4-4a7a-11e4-affa-1c6f65065a91"
}

```

Figure 366. Create Password Rule: Response

## Delete Password Rule

The **Delete Password Rule** operation deletes a Password Rule object designated by its element ID.

### HTTP method and URI

**DELETE** /api/console/password-rules/{password-rule-id}

In this request, the URI variable {password-rule-id} is the element ID of the Password Rule object to be deleted.

### Description

This operation removes a specified Password Rule from the console. The Password Rule is identified by the {password-rule-id} variable in the URI.

Upon successfully removing the Password Rule, HTTP status code 204 (No Content) is returned and no response body is provided.

The URI path must designate an existing Password Rule object; otherwise, status code 404 (Not Found) is returned. If the user does not have user-related-access permission to the designated Password Rule object or action/task permission to the **Manage Password Rules** task, status code 404 (Not Found) is returned. If the user has user-related-access permission to the designated Password Rule object but not action/task permission to the **Manage Password Rules** task, status code 403 (Forbidden) is returned. If any user has the specified Password Rule, the request fails and status code 409 (Conflict) is returned.

### Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage Password Rules** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
409 (Conflict)	317	The object cannot be deleted at this time. One or more users have this Password Rule.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
DELETE /api/console/password-rules/ecb26fb4-4a7a-11e4-affa-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

*Figure 367. Delete Password Rule: Request*

---

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT
```

```
<No response body>
```

---

*Figure 368. Delete Password Rule: Response*

### Inventory service data

Information about the Password Rules managed by the console can be optionally included in the inventory data provided by the Inventory Service.

| Inventory entries for Password Rule objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**console**" are to be included. An entry for a particular Password Rule is included only if the API user has access permission to that object as described in the **Get Password Rule Properties** operation.

| For each Password Rule object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the **Get Password Rule Properties** operation. That is, the data provided is the same as would be provided if a **Get Password Rule Properties** operation were requested targeting this object.

### | **Sample inventory data**

| The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single Password Rule. This object would appear as one array entry in the response array:

|

```

|   {
|     "case-sensitive": false,
|     "character-rules": [
|       {
|         "alphanumeric": "required",
|         "custom-character-sets": [],
|         "max-characters": 1,
|         "min-characters": 1,
|         "numeric": "not-allowed",
|         "special": "not-allowed"
|       },
|       {
|         "alphanumeric": "allowed",
|         "custom-character-sets": [],
|         "max-characters": 6,
|         "min-characters": 4,
|         "numeric": "required",
|         "special": "not-allowed"
|       },
|       {
|         "alphanumeric": "required",
|         "custom-character-sets": [],
|         "max-characters": 1,
|         "min-characters": 1,
|         "numeric": "not-allowed",
|         "special": "not-allowed"
|       }
|     ],
|     "class": "password-rule",
|     "consecutive-characters": 2,
|     "description": "",
|     "element-id": "56d11882-eaff-11e2-9ec7-5cf3fcae8019",
|     "element-uri": "/api/console/password-rules/56d11882-eaff-11e2-9ec7-5cf3fcae8019",
|     "expiration": 180,
|     "history-count": 0,
|     "max-length": 8,
|     "min-length": 6,
|     "name": "Strict",
|     "parent": "/api/console",
|     "replication-overwrite-possible": false,
|     "similarity-count": 0,
|     "type": "system-defined"
|   }

```

Figure 369. Password Rule object: Sample inventory data

## LDAP Server Definition object

An LDAP Server Definition object is an element of the console object and contains information about an LDAP server that may be used for console user authorization purposes. LDAP servers are sometimes referred to as Enterprise Directory Servers on the zManager user interface and publications.

All API users are permitted to issue **List LDAP Server Definitions** and retrieve very basic information for the LDAP Server Definition, if any, that applies to them. That is, they can use that operation to retrieve very basic information for the LDAP Server Definition identified in their User object. An API user with action/task permission to the **Manage LDAP Server Definitions** task is permitted to view and change any LDAP Server Definition object.

LDAP Server Definition objects may be replicated to this HMC via its Data Replication facility. If that is the case, the **Update LDAP Server Definition Properties** and **Delete LDAP Server Definition** operations should be used with care as they will prevent further replication of the modified or deleted object to this HMC.

## Data model

This object contains the following properties. Certain properties are only valid when mutable prerequisite properties have specific values. When such properties are not valid, their value is **null**. For instance the **search-filter** is **null** when the **location-method** value is **"pattern"**.

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

Table 215. LDAP Server Definition object: properties

Name	Qualifier	Type	Description of specialization
<b>element-uri</b>	—	String/URI	The canonical URI path for an LDAP Server Definition object is of the form <code>/api/console/ldap-server-definitions/{ldap-server-definition-id}</code> , where <code>{ldap-server-definition-id}</code> is the value of the <b>element-id</b> property of the LDAP Server Definition object.
<b>element-id</b>	—	String (36)	The unique identifier for this object.
<b>parent</b>	—	String/URI	The canonical URI path of the console object.
<b>class</b>	—	String	The class of an LDAP Server Definition object is <b>"ldap-server-definition"</b> .
<b>name</b>	(ro)	String	The name of the LDAP Server Definition object. This name must be unique among all LDAP Server Definition objects defined to the console. While preexisting LDAP Server Definition names are virtually unrestricted in terms of length and characters, new LDAP Server Definition names must conform to the length and character requirements of the <b>name</b> property described in the “Base managed object properties schema” on page 39.  For the purpose of verifying uniqueness, this name is treated in a case-insensitive fashion when used to create a new LDAP Server Definition object.
<b>description</b>	(w)	String (0-1024)	The description of the LDAP Server Definition object.  Default: an empty string
<b>primary-hostname-ipaddr</b>	(w)	String, String/IPV4 address, or String/IPV6 address	The host name or IP address of the primary LDAP server. It must contain at least 1 non-whitespace character.
<b>connection-port</b>	(w)	Integer	The TCP port number, which must be greater than 0, on which the server accepts connections, or <b>null</b> if the server uses the standard LDAP port appropriate for the value of the <b>use-ssl</b> property. The standard LDAP port values are 636 if <b>use-ssl</b> is <b>true</b> and 389 if <b>use-ssl</b> is <b>false</b>  Default: <b>null</b>

Table 215. LDAP Server Definition object: properties (continued)

Name	Qualifier	Type	Description of specialization
<b>backup-hostname-ipaddr</b>	(w)	String, String/IPV4 address, or String/IPV6 address	The host name or IP address of the backup LDAP server (which must contain at least 1 non-whitespace character), or <b>null</b> if there is none.  Default: <b>null</b>
<b>use-ssl</b>	(w)	Boolean	Indicates whether the server uses SSL for incoming connections.  Default: <b>false</b>
<b>tolerate-untrusted-certificates</b>	(w)	Boolean	Indicates whether the server should tolerate self-signed or otherwise untrusted certificates.  Prerequisite: <b>use-ssl</b> is <b>true</b>  Default: <b>false</b>
<b>bind-distinguished-name</b>	(w)	String	Part of the bind information for the initial connection, if needed: the distinguished name, which must contain at least 1 non-whitespace character; otherwise, <b>null</b> .  Default: <b>null</b>
<b>bind-password</b>	(wo)	String	Part of the bind information for the initial connection, if needed: the password; otherwise, <b>null</b>  Default: <b>null</b>
<b>location-method</b>	(w)	String Enum	The method this server uses to locate a user's directory entry. Must be one of the following: <ul style="list-style-type: none"> <li>• <b>"pattern"</b> - Use a distinguished name pattern.</li> <li>• <b>"subtree"</b> - Use a distinguished name subtree.</li> </ul>
<b>search-distinguished-name</b>	(w)	String	The distinguished name to use when searching for a user's directory entry.  If <b>location-method</b> is <b>"pattern"</b> , this is the distinguished name pattern to use when searching for a user's directory entry. It must include the string "{0}", which indicates where in the pattern the user ID is to be substituted. The user ID is the value of the <b>userid-on-ldap-server</b> property of the user's User object, unless it is <b>null</b> , in which case the <b>name</b> property of the User object is used.  If <b>location-method</b> is <b>"subtree"</b> , this is the distinguished name of the subtree to search for a user's directory entry.
<b>search-scope</b>	(w)	String Enum	Indicates how much of the subtree should be searched when searching for a user's directory entry in a subtree. The filter is specified in the <b>search-filter</b> property. Must be one of the following: <ul style="list-style-type: none"> <li>• <b>"all"</b> - Search the entire subtree.</li> <li>• <b>"one-level"</b> - Search one level only.</li> </ul> Prerequisite: <b>location-method</b> is <b>"subtree"</b>  Default: <b>"all"</b>

Table 215. LDAP Server Definition object: properties (continued)

Name	Qualifier	Type	Description of specialization
search-filter	(w)	String	The LDAP search filter to use when searching for a user's directory entry in a subtree. It must include the string "{0}", which indicates where in the filter the user ID is to be substituted. The user ID is the value of the <b>userid-on-ldap-server</b> property of the user's User object, unless it is <b>null</b> , in which case the <b>name</b> property of the User object is used. The <b>search-scope</b> property specifies how this filter is used during the search.  Prerequisite: <b>location-method</b> is "subtree"
replication-overwrite-possible	—	Boolean	Indicates whether this object is customizable data that is replicated to this HMC from an HMC configured as a Data Source in the Data Replication service.

## List LDAP Server Definitions

The List LDAP Server Definitions operation lists LDAP Server Definitions defined to the console.

### HTTP method and URI

GET /api/console/ldap-server-definitions

### Query Parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	Filter pattern (regular expression) to limit returned objects to those that have a matching <b>name</b> property.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
ldap-server-definitions	Array of objects	Array of nested ldap-server-definition-info objects as described in the next table.

Each nested ldap-server-definition-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	The <b>element-uri</b> property of the LDAP Server Definition object.
name	String	The <b>name</b> property of the LDAP Server Definition object.

### Description

This operation lists LDAP Server Definitions defined to the console. Some basic properties are provided for each LDAP Server Definition.



If the **name** query parameter is specified the returned list is limited to those LDAP Server Definitions that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not performed.

An LDAP Server Definition is included in the list only if the API user has user-related-access permission to that object or action/task permission to the **Manage LDAP Server Definitions** task. If there is an LDAP Server Definition to which the API user does not have permission, that object is omitted from the list, but no error status code results.

If there are no LDAP Server Definitions defined to the console or if no LDAP Server Definitions are to be included in the results due to filtering or access permissions, an empty list is provided and the operation completes successfully.

### Authorization requirements

This operation has the following authorization requirement:

- User-related-access permission to the LDAP Server Definition objects included in the response body or action/task permission to the **Manage LDAP Server Definitions** task

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 720.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/console/ldap-server-definitions?name=IBM.*&name=Company.* HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

Figure 370. List LDAP Server Definitions: Request

---

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:34 GMT
content-type: application/json;charset=UTF-8
content-length: 276
{
  "ldap-server-definitions":[
    {
      "element-uri":"/api/console/ldap-server-definitions/3ac6550e-1dbb-11e4-9aa4-1c6f65065a91",
      "name":"IBM LDAP server"
    },
    {
      "element-uri":"/api/console/ldap-server-definitions/ece481ca-4a7a-11e4-8777-1c6f65065a91",
      "name":"Company LDAP server"
    }
  ]
}
```

---

Figure 371. List LDAP Server Definitions: Response

## Get LDAP Server Definition Properties

The **Get LDAP Server Definition Properties** operation retrieves the properties of a single LDAP Server Definition object that is designated by its element ID.

### HTTP method and URI

**GET** `/api/console/ldap-server-definitions/{ldap-server-definition-id}`

In this request, the URI variable `{ldap-server-definition-id}` is the element ID of the LDAP Server Definition object whose properties are to be returned.

### Response body contents

On successful completion, the response body contains a JSON object that provides the current values of the properties for the Password Rule object as defined in the “Data model” on page 718. Field names and data types in the JSON object are the same as the property names and data types defined in the Data Model.

### Description

This operation returns the current properties of a single LDAP Server Definition object that is designated by `{ldap-server-definition-id}`.

On successful execution, all of the current properties as defined in the Data Model for the LDAP Server Definition object, except those designated as write-only properties, are provided in the response body, and HTTP status code 200 (OK) is returned.

The URI path must designate an existing LDAP Server Definition object; otherwise, status code 404 (Not Found) is returned. In addition, the API user must have action/task permission to the **Manage LDAP Server Definitions** task; otherwise, status code 403 (Forbidden) is returned.

### Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage LDAP Server Definitions** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 722.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
GET /api/console/ldap-server-definitions/ece481ca-4a7a-11e4-8777-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

Figure 372. Get LDAP Server Definition Properties: Request

```
00 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:34 GMT
content-type: application/json;charset=UTF-8
content-length: 640
{
  "backup-hostname-ipaddr":null,
  "bind-distinguished-name":null,
  "class":"ldap-server-definition",
  "connection-port":null,
  "description":"Directory server for the company",
  "element-id":"ece481ca-4a7a-11e4-8777-1c6f65065a91",
  "element-uri":"/api/console/ldap-server-definitions/ece481ca-4a7a-11e4-8777-1c6f65065a91",
  "location-method":"subtree",
  "name":"Company LDAP server",
  "parent":"/api/console",
  "primary-hostname-ipaddr":"ldap1.our.company.com",
  "replication-overwrite-possible":false,
  "search-distinguished-name":"o=our,ou=company.com",
  "search-filter":"email={0}",
  "search-scope":"all",
  "tolerate-untrusted-certificates":false,
  "use-ssl":true
}
```

Figure 373. Get LDAP Server Definition Properties: Response

## | Update LDAP Server Definition Properties

| The **Update LDAP Server Definition Properties** operation updates the properties of a single LDAP Server Definition object that is designated by its element ID.

### | HTTP method and URI

| **POST** `/api/console/ldap-server-definitions/{ldap-server-definition-id}`

| In this request, the URI variable `{ldap-server-definition-id}` is the element ID of the LDAP Server Definition object whose properties are to be updated.

### | Request body contents

| The request body is expected to contain a JSON object that provides the new values of any writeable property that is to be updated by this operation. Field names and data types in this JSON object are expected to match the corresponding property names and data types defined by the data model for this object type. The JSON object can and should omit fields for properties whose values are not to be changed by this operation.

### | Description

| This operation updates writeable properties of the LDAP Server Definition object specified by `{ldap-server-definition-id}`.

| The URI path must designate an existing LDAP Server Definition object; otherwise, status code 404 (Not Found) is returned. In addition, the API user must have action/task permission to the **Manage LDAP Server Definitions** task; otherwise, status code 403 (Forbidden) is returned.

| The request body is validated against the schema described in the Request Body Contents section. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. The request body validation will fail if it contains a property that is not valid because a prerequisite is not met (e.g., attempting to set **search-filter** when the **location-method** value is **"pattern"**).

| The request body does not need to specify a value for all writeable properties, but rather can and should contain fields only for the properties to be updated. Object properties for which no input value is provided and no prerequisite property is changed remain unchanged by this operation. A property's value is set to its default value if the field is not included in the request body and a prerequisite field is changed such that the prerequisite condition becomes satisfied (e.g., if **location-method** is changed from **"pattern"** to **"subtree"**, and **search-scope** is not defined in the request body, **search-scope** will be defaulted to **all**).

### | Authorization requirements

| This operation has the following authorization requirement:

- Action/task permission to the **Manage LDAP Server Definitions** task.

### | HTTP status and reason codes

| On success, HTTP status code 204 (No Content) is returned and no response body is provided.

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
POST /api/console/ldap-server-definitions/ece481ca-4a7a-11e4-8777-1c6f65065a91 HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 133
{
  "backup-hostname-ipaddr": "ldap2.my.company.com",
  "description": "A new and improved description of this LDAP Server Definition"
}
```

Figure 374. Update LDAP Server Definition Properties: Request

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:34 GMT

<No response body>
```

Figure 375. Update LDAP Server Definition Properties: Response

## Create LDAP Server Definition

The **Create LDAP Server Definition** operation creates an LDAP Server Definition object with the given properties.

### HTTP method and URI

**POST /api/console/ldap-server-definitions**

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The value to be set as the LDAP Server Definition's <b>name</b> property.
description	String	Optional	The value to be set as the LDAP Server Definition's <b>description</b> property.
primary-hostname-ipaddr	String	Required	The value to be set as the LDAP Server Definition's <b>primary-hostname-ipaddr</b> property.

Field name	Type	Rqd/Opt	Description
connection-port	Integer	Optional	The value to be set as the LDAP Server Definition's <b>connection-port</b> property.
backup-hostname-ipaddr	String	Optional	The value to be set as the LDAP Server Definition's <b>backup-hostname-ipaddr</b> property.
use-ssl	Boolean	Optional	The value to be set as the LDAP Server Definition's <b>use-ssl</b> property.
tolerate-untrusted-certificates	Boolean	Optional	The value to be set as the LDAP Server Definition's <b>tolerate-untrusted-certificates</b> property.
bind-distinguished-name	String	Optional	The value to be set as the LDAP Server Definition's <b>bind-distinguished-name</b> property.
bind-password	String	Optional	The value to be set as the LDAP Server Definition's <b>bind-password</b> property.
location-method	String Enum	Optional	The value to be set as the LDAP Server Definition's <b>location-method</b> property.
search-distinguished-name	String	Required	The value to be set as the LDAP Server Definition's <b>search-distinguished-name</b> property.
search-scope	String Enum	Optional	The value to be set as the LDAP Server Definition's <b>search-scope</b> property.
search-filter	String	Required if <b>location-method</b> is "subtree"	The value to be set as the LDAP Server Definition's <b>search-filter</b> property.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the new LDAP Server Definition object.

## Description

This operation creates a new LDAP Server Definition.

On successful execution of this operation the LDAP Server Definition is created using the inputs as specified by the request body. The URI of the new LDAP Server Definition is provided in the response body and in a **Location** response header as well

The request body is validated against the schema described in the "Request body contents" on page 725. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered. The request body validation will fail if it contains a property that is not valid because a prerequisite is not met (e.g., specifying **search-filter** when the **location-method** value is "**pattern**") or the specified name is not unique. In addition, the API user must have action/task permission to the **Manage LDAP Server Definitions** task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage LDAP Server Definitions** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents” on page 726, and the **Location** response header contains the URI of the newly created object.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 216. Create LDAP Server Definition: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	8	An LDAP Server Definition with the name specified in the request body already exists.
403 (Forbidden)	1	The API user does not have the required permission for this operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/console/ldap-server-definitions HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
content-type: application/json
content-length: 264
{
  "description":"Directory server for the company",
  "location-method":"subtree",
  "name":"Company LDAP server",
  "primary-hostname-ipaddr":"ldap1.our.company.com",
  "search-distinguished-name":"o=our,ou=company.com",
  "search-filter":"email={0}",
  "use-ssl":true
}
```

---

Figure 376. Create LDAP Server Definition: Request

---

```
201 Created
server: zSeries management console API web server / 2.0
location: /api/console/ldap-server-definitions/ece481ca-4a7a-11e4-8777-1c6f65065a91
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:33 GMT
content-type: application/json;charset=UTF-8
content-length: 94
{
  "element-uri":"/api/console/ldap-server-definitions/ece481ca-4a7a-11e4-8777-1c6f65065a91"
}
```

---

Figure 377. Create LDAP Server Definition: Response

## Delete LDAP Server Definition

The **Delete LDAP Server Definition** operation deletes an LDAP Server Definition object designated by its element ID.

### HTTP method and URI

**DELETE** /api/console/ldap-server-definitions/{*ldap-server-definition-id*}

In this request, the URI variable *{ldap-server-definition-id}* is the element ID of the LDAP Server Definition object to be deleted.

### Description

This operation removes a specified LDAP Server Definition from the console. The LDAP Server Definition is identified by the *{ldap-server-definition-id}* variable in the URI.

Upon successfully removing the LDAP Server Definition, HTTP status code 204 (No Content) is returned and no response body is provided.

The URI path must designate an existing LDAP Server Definition object; otherwise, status code 404 (Not Found) is returned. In addition, the API user must have action/task permission to the **Manage LDAP Server Definitions** task; otherwise, status code 403 (Forbidden) is returned. If any user is defined to use the specified LDAP Server Definition, the request fails and HTTP status code 409 (Conflict) is returned.

### Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Manage LDAP Server Definitions** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the correct type or the API user has no access permission to it.
409 (Conflict)	317	The object cannot be deleted at this time. One or more users have this LDAP Server Definition.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.



## Example HTTP interaction

---

```
DELETE /api/console/ldap-server-definitions/ece481ca-4a7a-11e4-8777-1c6f65065a91
HTTP/1.1
x-api-session: 2t4ixcf8nplr7yersi8i9b953fgxvqx18c4r066ge9kcyzr4c
```

---

*Figure 378. Delete LDAP Server Definition: Request*

---

```
204 No Content
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Thu, 02 Oct 2014 21:27:34 GMT

<No response body>
```

---

*Figure 379. Delete LDAP Server Definition: Response*

## Inventory service data

Information about the LDAP Server Definitions managed by the console can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for LDAP Server Definition objects are included in the response to the Inventory Service's **Get Inventory** operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**console**" are to be included. An entry for a particular LDAP Server Definition is included only if the API user has access permission to that object as described in the **Get LDAP Server Definition Properties** operation.

For each LDAP Server Definition object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for the **Get LDAP Server Definition Properties** operation. That is, the data provided is the same as would be provided if a **Get LDAP Server Definition Properties** operation were requested targeting this object.

## Sample inventory data

The following fragment is an example of the JSON object that would be included in the **Get Inventory** response to describe a single LDAP Server Definition. This object would appear as one array entry in the response array:

---

```

|   {
|     "backup-hostname-ipaddr": null,
|     "bind-distinguished-name": null,
|     "class": "ldap-server-definition",
|     "connection-port": null,
|     "description": "",
|     "element-id": "ffbf71f4-370d-11e4-a5fc-5ef3fcae8020",
|     "element-uri": "/api/console/ldap-server-definitions/ffbf71f4-370d-11e4-
|       a5fc-5ef3fcae8020",
|     "location-method": "pattern",
|     "name": "Temp_LDAP_13_56ba2f43-98c0-4848-9af8-cdb45b56f082",
|     "parent": "/api/console",
|     "primary-hostname-ipaddr": "bluepages.ibm.com",
|     "replication-overwrite-possible": false,
|     "search-distinguished-name": "uid={0}744,c=in,ou=bluepages,o=ibm.com",
|     "search-filter": null,
|     "search-scope": null,
|     "tolerate-untrusted-certificates": null,
|     "use-ssl": false
|   }

```

---

Figure 380. LDAP Server Definition object: Sample inventory data

---

## Group Object

The Hardware Management Console provides a fixed set of system-defined groups to which managed objects of certain types automatically belong, as members. For example, defined CPCs are automatically members of the CPC group. By their nature, the members of the system-defined groups are obtainable via list operations of the appropriate API. For example, all the CPCs managed by a Hardware Management Console can be obtained via a List CPCs operation. Therefore, list operations for system-defined groups are unnecessary. By their nature, the existence of a system-defined group and its content (members) is implicit. Therefore, create/delete operations for system-defined groups are both unnecessary and inappropriate.

These system-defined groups are distinct from user-defined (“custom”) groups. The latter are explicitly created by users for their own purposes: for example, it may be convenient for management purposes to take some proper subset of the members of the system-defined CPC group as a user-defined group of CPCs. User-defined groups may be homogeneous (all members of the same managed object type, as in this previous example), but need not be.

A Group object represents one or more managed objects which are called group members. Each member is of some object type: CPC, Logical Partition, etc. Note that groups may be heterogeneous (with member objects of differing types), and may even have other groups as members.

Users may define groups in one of two ways:

1. by use of a pattern-match expression to implicitly define membership (pattern-matching group)
2. by explicitly choosing members.

This API can be used to view/manage custom groups, and membership within these groups. The latter is subject to restrictions, based on which of the two fundamentally different means of definition the user employed:

- If pattern-matching was specified, then group membership is “implicit”. In this case, operations to add/remove a member are unnecessary (simply create/delete the managed object, itself, using the appropriate API operation). Accordingly, member-management operations are not supported for groups using pattern-matching.
- If pattern-matching was not specified, then group membership is “explicit”, and in this case operations to add/remove group members are both useful and appropriate. Accordingly, for custom groups not

based on pattern-matching, member-management operations are supported. Note that such operations do not affect the member object itself, only its group membership status.

- When groups are defined using pattern-matching, the types of managed objects to which pattern-matching is applied must be explicitly specified. Regardless of the POSIX regular expression specified as the match pattern, managed objects whose names match the pattern but who are not of the specified object type(s) are not considered to be members of the group.
- Groups are not intrinsically ordered in any way, nor are members within a given group. List-oriented operations therefore do not return ordered results.

## Data model

This object includes the properties defined in the “Base managed object properties schema” on page 39, but does not provide the operational-status-related properties defined in that schema because it does not maintain the concept of an operational status.

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38

The following class-specific specializations apply to the other base managed object properties:

Table 217. Group object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/ URI	The canonical URI path of the Group object, of the form <code>/api/groups/{group-id}</code> where <code>{group-id}</code> is the value of the <b>object-id</b> property of the Group object.
<b>parent</b>	—	String	This property is always a null object.
<b>class</b>	—	String	The class of a Group object is " <b>group</b> ".
<b>name</b>	(ro)	String	The group name specified by the user when the group was created
<b>description</b>	(ro)	String	The description specified by the user when the group was created, or if none was provided, the IBM provided caption text.
<b>replication-overwrite-possible</b>	—	Boolean	Indicates whether this object is replicated to this HMC from an HMC configured as a Data Source in the Data Replication service.

## Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties:

Table 218. Group object: class specific additional properties

Name	Type	Description
<b>match-info</b>	match-info object	A nested object which pertains to pattern-matching groups only, as described in the next table. An empty value is returned for groups which do not use pattern-matching.

The match-info object contains the following fields:

Table 219. match-info object properties

Name	Type	Description
<b>pattern</b>	String	A regular expression used to define membership for pattern-matching groups. This field has no length limitations.

Table 219. match-info object properties (continued)

Name	Type	Description
types	Array of String Enum	Specifies the type(s) of objects that are eligible for membership in pattern-matching groups. One or more of the following: <ul style="list-style-type: none"> <li>• "custom-groups" - zManager API objects of class "group"</li> <li>• "defined-cpc" - zManager API objects of class "cpc"</li> <li>• "director-timer-console" - ESCON director and/or Sysplex timer consoles</li> <li>• "ibm-fiber-saver" - IBM 2029 fiber optic data transports</li> <li>• "logical-partition" - zManager API objects of class "logical-partition"</li> <li>• "zvm-virtual-machines" - zManager API objects of class "virtual-server", type of "zvm"</li> <li>• "blade-center" - zManager API objects of class "bladecenter"</li> <li>• "data-power-xi50z-blades" - zManager API objects of class "blade", type of "dpxi50z"</li> <li>• "ibm-smart-analytics-optimizer-blades" - zManager API objects of class "blade", type of "isaopt"</li> <li>• "power-blade" - zManager API objects of class "blade", type of "power"</li> <li>• "system-x-blade" - zManager API objects of class "blade", type of "x-hyp"</li> <li>• "power-vm-virtual-server" - zManager API objects of class "virtual-server", type of "power-vm"</li> <li>• "system-x-virtual-server" - zManager API objects of class "virtual-server", type of "x-hyp"</li> <li>• "workload" - zManager API objects of class "workload-resource-group".</li> <li>• "defined-zbx-node" - zManager API objects of class "zbx", type of "node".</li> </ul>

## List Custom Groups

The **List Custom Groups** operation lists the custom groups which are visible to the API user.

### HTTP method and URI

GET /api/groups

### Query Parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	A regular expression used to limit returned objects to those that have a matching name property. If matches are found, the response will be an array with all objects that match. If no match is found, the response will be an empty array.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
groups	Array of group-info objects	Array of nested objects which identify groups that are visible to the API user.

Each nested group-info object contains the following fields:

Field name	Type	Description
object-uri	String/URI	The value of the Group object's <b>object-uri</b> property.
name	String	The value of the Group object's <b>name</b> property.

## Description

This operation lists the Group objects which are visible to the API user. Only groups to which the caller has authorization will be returned.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in the Response Body Contents section. If no groups exist, or if no groups are visible to the API user, HTTP status code 200 (OK) is returned, along with an empty response body.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the Group object.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described “Response body contents” on page 732.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	299	A query parameter has an invalid syntax.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/groups HTTP/1.1
x-api-session: 4ipkcbjpy5koce1t65213dvv85gi81iqy5bz8yrpt6vtrt8ks
```

---

Figure 381. List Custom Groups: Request

---

```
| 200 OK
| server: zSeries management console API web server / 1.0
| cache-control: no-cache
| date: Fri, 25 Nov 2011 16:02:42 GMT
| content-type: application/json;charset=UTF-8
| content-length: 283
| {
|   "groups": [
|     {
|       "name": "Finance department CPCs",
|       "object-uri": "/api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341"
|     },
|     {
|       "name": "Test Group",
|       "object-uri": "/api/groups/febde5ab-a4a6-35bf-9e01-83aae59d7e52"
|     }
|   ]
| }
```

---

Figure 382. List Custom Groups: Response

## Get Custom Group Properties

The **Get Custom Group Properties** operation retrieves the properties of a single Group object that is designated by the *{group-id}*.

### HTTP method and URI

**GET** /api/groups/{group-id}

In this request, the URI variable *{group-id}* is the object ID of the group.

### Response body contents

On success, HTTP status code 200 (OK) is returned and the response body contains an object that provides the current values of the properties for the Group object as defined in “Data model” on page 731. Field names and data types in the object are the same as the property names and data types defined in the data model.

### Description

This operation returns the current properties for the Group object designated by *{group-id}*.

The URI path *{group-id}* must designate an existing Group object.

### Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the Group object designated by *{group-id}*.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

On error, appropriate HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The <b>object-id</b> in the URI ( <i>group-id</i> ) does not designate an existing group, or the API user does not have sufficient access (as described above).

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341 HTTP/1.1
x-api-session: 42r6t4chltpvd614l61wi3111tf7fv2hes80hjqs3invt7cp
```

---

Figure 383. Get Custom Group Properties: Request

---

```
| 200 OK
| server: zSeries management console API web server / 1.0
| cache-control: no-cache
| date: Fri, 25 Nov 2011 16:45:45 GMT
| content-type: application/json;charset=UTF-8
| content-length: 250
| {
|   "class": "group",
|   "description": "Spacely Sprockets Web Servers",
|   "replication-overwrite-possible": false,
|   "is-locked": false,
|   "match-info": {},
|   "name": "SS-Web-Servers",
|   "object-id": "ee2782af-dd98-3ec0-bc2d-cfe2e9154341",
|   "object-uri": "/api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341",
|   "parent": null
| }
```

---

Figure 384. Get Custom Group Properties: Response

## Create Custom Group

Use the **Create Custom Group** operation to create a custom group.

### HTTP method and URI

**POST** /api/groups

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
name	String	Required	The name for the new custom Group object
description	String	Optional	The description for the new custom Group object

Field name	Type	Rqd/Opt	Description
match-info	match-info object	Optional	A nested object describing the pattern match. If not provided, this is not a pattern-match custom group. Refer to "Class specific additional properties" on page 731 for details.

## Response body contents

Field name	Type	Description
object-uri	String	The object URI of the new custom group.

## Description

Group objects are programmatically identified by object-id and not by name. To avoid the confusion which might result from allowing redundant names, the **name** property is required for this operation, and the (case-sensitive) value supplied for the name property must be distinct from that of all currently-existing Group objects. In keeping with restrictions imposed by the Hardware Management Console's Graphical User Interface (GUI), the following set of names is also not allowed:

- the current name of the Hardware Management Console
- the GUI View names {"Groups", "Exceptions", "Active Tasks", "Console Actions", "Task List", "Books", "Help", "Ensemble"}

On success, a custom group managed object is created reflecting the Request Body contents and HTTP status code 201 (Created) is returned.

## Authorization requirements

This operation has the following authorization requirement:

- Action/task permission to the **Grouping** task.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in "Response body contents." In addition, the **Location** response header contains the URI of the newly created object.

On error, appropriate HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	261	One of the following errors was detected: <ul style="list-style-type: none"> <li>• The pattern string specified in <b>match-info</b> is not valid. This must be a non-empty string which is a valid regular expression.</li> <li>• One or more of the types specified in <b>match-info</b> is invalid. At least one type must be specified, and all must be values as documented for the <b>match-info types</b> property.</li> </ul>
	290	The requested name is either reserved or already in use.
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform this operation.
500 (Server Error)	273	An unexpected error occurred during the operation.



Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/groups HTTP/1.1
x-api-session: 42r6t4chltpvd6l4l61wi3111tf7fv2hes80hjqs3inv7cp
content-type: application/json
content-length: 74
{
  "description": "Spacely Sprockets Web Servers",
  "name": "SS-Web-Servers"
}
```

---

Figure 385. Create Custom Group: Request

---

```
201 Created
server: zSeries management console API web server / 1.0
location: /api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341
cache-control: no-cache
date: Fri, 25 Nov 2011 16:45:44 GMT
content-type: application/json;charset=UTF-8
content-length: 65
{
  "object-uri": "/api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341"
}
```

---

Figure 386. Create Custom Group: Response

## Delete Custom Group

Use the **Delete Custom Group** operation to delete a custom group.

### HTTP method and URI

```
DELETE /api/groups/{group-id}
```

In this request, the URI variable *{group-id}* is the object ID of the group.

### Description

If successful, the custom group managed object designated by *{group-id}* is deleted.

If *{group-id}* does not identify an existing custom group, status code 404 (Not Found) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object-access permission to the custom group designated by *{group-id}*
- Action/task permission for the **Grouping** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

On error, appropriate HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform this operation.
404 (Not Found)	1	The URI's <i>{group-id}</i> does not designate an existing custom Group object, or the API user does not have object access to the group.
500 (Server Error)	273	An unexpected error occurred during the operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
DELETE /api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341 HTTP/1.1
x-api-session: 42r6t4chltpvd614161wi3111tf7fv2hes80hjqs3invt7cp
```

---

Figure 387. Delete Custom Group: Request

---

```
204 No Content
date: Fri, 25 Nov 2011 16:45:45 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 388. Delete Custom Group: Response

## Add Member to Custom Group

Use the **Add Member to Custom Group** operation to add a member to a custom group.

### HTTP method and URI

**POST** /api/groups/{group-id}/operations/add-member

In this request, the URI variable *{group-id}* is the object ID of the group.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
object-uri	String	Required	The object URI of the object to be added to the group

## Description

If successful, the managed object designated in the request body attains membership in the custom group identified by *{group-id}*.

The operation is subject to the following restrictions:

- The designated managed object must exist and must not already be a member of the group identified by *{group-id}*
- The group identified by *{group-id}* must be a custom group defined without a pattern-matching specification.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the custom Group object designated by *{group-id}*
- Object access permission to the object designated by the request body
- Action/task permission for the **Grouping** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

On error, appropriate HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	291	The designated managed object is already a member of the custom group identified by <i>{group-id}</i> .
	294	The group identified by <i>{group-id}</i> was defined using pattern-matching.
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform this operation.
	293	The addition of the member to the custom group designated by the URI ( <i>{group-id}</i> ) would introduce a circular reference, which is not permitted.
404 (Not Found)	1	The URI's <i>{group-id}</i> does not designate an existing custom Group object, or the API user does not have object access to the group.
	2	The request body does not designate an existing managed object, or the API user does not have sufficient access to the managed object.
500 (Server Error)	273	An unexpected error occurred during the operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341/operations/add-member HTTP/1.1
x-api-session: 42r6t4chltpvd614161wi3111tf7fv2hes80hjqs3invt7cp
content-type: application/json
content-length: 75
{
  "object-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af"
}
```

---

Figure 389. Add Member to Custom Group: Request

---

```
204 No Content
date: Fri, 25 Nov 2011 16:45:44 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 390. Add Member to Custom Group: Response

## Remove Member from Custom Group

Use the **Remove Member from Custom Group** operation to remove a member from a custom group.

### HTTP method and URI

**POST** /api/groups/{group-id}/operations/remove-member

In this request, the URI variable *{group-id}* is the object ID of the group

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
object-uri	String	Required	The object URI of the object to be removed from the group

### Description

The managed object designated in the request body relinquishes its membership in the custom group identified by *{group-id}*.

The operation is subject to the following restrictions:

- The managed object designated in the request body must currently be a member of the group identified by *{group-id}*.
- The group identified by *{group-id}* must be a custom group defined without a pattern-matching specification

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the custom Group object designated by *{group-id}*
- Object access permission to the object designated by the request body
- Action/task permission for the **Grouping** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

On error, appropriate HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	291	The designated managed object is not a member of the custom group identified by <i>{group-id}</i> .
	294	The group identified by <i>{group-id}</i> was defined using pattern-matching.
403 (Forbidden)	1	The user under which the API request was authenticated does not have the required authority to perform this operation.
404 (Not Found)	1	The URI's <i>{group-id}</i> does not designate an existing custom Group object, or the API user does not have object access to the group.
	2	The request body does not designate an existing managed object, or the API user does not have sufficient access to the managed object.
500 (Server Error)	273	An unexpected error occurred during the operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
POST /api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341/operations/remove-member HTTP/1.1
x-api-session: 42r6t4chltpvd614l61wi3111tf7fv2hes80hjqs3invt7cp
content-type: application/json
content-length: 75
{
  "object-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af"
}
```

---

Figure 391. Remove Member from Custom Group: Request

---

```
204 No Content
date: Fri, 25 Nov 2011 16:45:45 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

Figure 392. Remove Member from Custom Group: Response

## List Custom Group Members

Use the **List Custom Group Members** operation to list custom group members.

## HTTP method and URI

GET /api/groups/{group-id}/members

In this request, the URI variable *{group-id}* is the object ID of the group.

## Response body contents

Field name	Type	Description
members	Array of nested objects	Array of nested objects which identify members of the group designated by <i>{group-id}</i> .

Each nested member object contains the following fields:

Field name	Type	Description
object-uri	String/URI	The value of the member object's <b>object-uri</b> property.
name	String	The value of the member object's <b>name</b> property.

## Description

This operation lists the members of the Group object designated by *{group-id}*. The results of this operation only include references to member objects for which the API user has object access authority.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in the Response Body Contents section. If the group currently has no members, HTTP status code 200 (OK) is returned, along with an empty response body.

On error, appropriate HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the Group object.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

On error, appropriate HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The group designated by the URI ( <i>{group-id}</i> ) does not exist, or the API user does not have sufficient access (as described above).

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Example HTTP interaction

---

```
GET /api/groups/ee2782af-dd98-3ec0-bc2d-cfe2e9154341/members HTTP/1.1
x-api-session: 42r6t4chltipvd614161wi3111tf7fv2hes80hjqs3inv7cp
```

---

Figure 393. List Custom Group Members: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 16:45:45 GMT
content-type: application/json;charset=UTF-8
content-length: 207
{
  "members": [
    {
      "name": "SS-Web-Svr-1",
      "object-uri": "/api/virtual-servers/576569dc-0db7-11e1-b1f1-f0def14b63af"
    },
    {
      "name": "SS-Web-Svr-2",
      "object-uri": "/api/virtual-servers/588d8c18-0db7-11e1-b1f1-f0def14b63af"
    }
  ]
}
```

---

Figure 394. List Custom Group Members: Response

## Inventory service data

Information about custom groups can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for the Group objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of the class **"group"** are to be included. An entry for a particular group is included only if the API user has access permission to that object as described in the **Get Custom Group Properties** operation.

For each Group object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for "Get Custom Group Properties" on page 734. That is, the data provided is the same as would be provided if a **Get Custom Group Properties** operation were requested targeting this object.

---

## CPC object

A CPC object represents a single z Systems Central Processor Complex (CPC) that is being managed by an HMC.

## Data model

For definitions of the qualifier abbreviations in the following tables, see "Property characteristics" on page 38.

This object includes the properties defined in the "Base managed object properties schema" on page 39, with the following class-specific specialization:

Table 220. CPC object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/ URI	The canonical URI path of the CPC object, of the form <code>/api/cpcs/{cpc-id}</code> where <code>{cpc-id}</code> is the value of the <b>object-id</b> property of the CPC object.
<b>parent</b>	(pc)	String/ URI	If the CPC is a member of an ensemble, the parent is the canonical URI path for the ensemble object. Otherwise, this property contains a null object.
<b>class</b>	—	String	The class of a CPC object is <b>"cpc"</b> .
<b>name</b>	(ro)	String (1-8)	The CPC name
<b>description</b>	(ro)	String (0-1024)	The descriptive text associated with this CPC object.
<b>status</b>	(sc)	String Enum	The current operational status of the CPC object. One of: <ul style="list-style-type: none"> <li>• <b>"operating"</b> - the CPC is operational</li> <li>• <b>"not-communicating"</b> - the CPC is not communicating with the HMC</li> <li>• <b>"exceptions"</b> - the CPC has one or more unusual conditions</li> <li>• <b>"status-check"</b> - the SE is not communicating with the CPC</li> <li>• <b>"service"</b> - the CPC has been placed in service mode</li> <li>• <b>"not-operating"</b> - the CPC is not operational</li> <li>• <b>"no power"</b> - the CPC has no power</li> <li>• <b>"service-required"</b> - the CPC is operating on the last redundant part of a particular type</li> <li>• <b>"degraded"</b> - one or more of the CPC elements are degraded.</li> </ul>
<b>acceptable-status</b>	(w)(pc)	Array of String Enum	The set of operational <b>status</b> values in which the CPC object can exist and be considered in an acceptable (not alert causing) state. One or more of the values listed for the <b>status</b> property.

## Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties.

Table 221. CPC object: class specific additional properties

Name	Qualifier	Type	Description
<b>se-version</b>	(pc)	String (1-8)	The current release level of the primary SE internal code. For example, "2.11.1". Note that the alternate SE is normally at the same level, except when installing new internal code levels.
<b>has-hardware-messages</b>	(pc)	Boolean	The CPC object has hardware messages (true), or does not have hardware messages (false).
<b>is-ensemble-member</b>	(pc)	Boolean	Whether the CPC is currently part of an ensemble (true) or not (false).
<b>iml-mode</b>	(pc)	String Enum	The Initial Microcode Load (IML) mode type of the CPC object. One of: <ul style="list-style-type: none"> <li>• <b>"not-set"</b> - the CPC is not IMLed</li> <li>• <b>"esa390"</b> - the CPC is in ESA/390 mode</li> <li>• <b>"lpar"</b> - the CPC is in logical partition mode</li> <li>• <b>"esa390-tpf"</b> - the CPC is in ESA/390 TPF mode</li> </ul>
<b>next-activation-profile-name</b>	(w)(pc)	String (1-16)	The name of the activation profile to be used on the next activation of the CPC.
<b>last-used-activation-profile</b>	(pc)	String (0-16)	The name of the activation profile used on the last activation of the CPC or a null string.



Table 221. CPC object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>machine-model</b>	(pc)	String (1-3)	The model of the machine containing this CPC. For example, "M15".
<b>machine-type</b>	(pc)	String (1-4)	The type of the machine containing this CPC. For example, "2817".
<b>machine-serial-number</b>	(pc)	String (1-12)	The serial number of the machine containing this CPC. For example, "00 - SP1D92B".
<b>cpc-serial-number</b>	—	String (1-12)	The serial number of the CPC. For example, "00000SP1D92B".
<b>cpc-node-descriptor</b>	(pc)	String (2)	The hexadecimal number mapped to the device location of the CPC. This property identifies the CPC's relative order among other CPCs, if any, in the machine. For example, "00".
<b>is-cbu-installed</b>	—	Boolean	The Capacity Backup Upgrade (CBU) facility is installed (true), or not installed (false). Note: if status is "not-communicating", a null object is returned. Refer to the <i>Capacity On Demand User's Guide</i> for details.
<b>is-cbu-enabled</b>	—	Boolean	CBU is enabled (true), or is not enabled (false). Note: if status is "not-communicating", a null object is returned. Refer to the <i>Capacity On Demand User's Guide</i> for details.
<b>is-cbu-activated</b>	—	Boolean	CBU is activated (true), or is not activated (false). Note: if status is "not-communicating", a null object is returned. Refer to the <i>Capacity On Demand User's Guide</i> for details.
<b>is-real-cbu-available</b>	—	Boolean	Real CBU is available (true), or not available (false). Note: if status is "not-communicating", a null object is returned. Refer to the <i>Capacity On Demand User's Guide</i> for details.
<b>cbu-activation-date</b>	—	Timestamp	The date of CBU activation. Note: if status is "not-communicating", a null object is returned. Refer to the <i>Capacity On Demand User's Guide</i> for details.
<b>cbu-expiration-date</b>	—	Timestamp	The date of CBU expiration. Note: if status is "not-communicating", a null object is returned. Refer to the <i>Capacity On Demand User's Guide</i> for details.
<b>cbu-number-of-tests-left</b>	—	Integer	The number of CBU tests left. Note: if status is "not-communicating", a null object is returned. Refer to the <i>Capacity On Demand User's Guide</i> for details.
<b>is-service-required</b>	—	Boolean	Whether the CPC is operating using the last redundant part of a particular type (true) or not (false). If true, repairs should be made before additional parts fail that would make this CPC non-operational.  Support Element (Version 2.12.1 and newer) information can be found on the console help system, or on the IBM Knowledge Center at <a href="https://www.ibm.com/support/knowledgecenter">https://www.ibm.com/support/knowledgecenter</a> . (Select <b>z Systems</b> on the navigation bar, and then select your server). For information about earlier versions of the Support Element, see the <i>Support Element Operations Guide</i> .

Table 221. CPC object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>degraded-status</b>	(pc)	Array of String Enum	The set of degraded status values. If the CPC is not degraded, this property contains “not-degraded” as the only value. Otherwise, this property contains one or more of the following: <ul style="list-style-type: none"> <li>• <b>"memory"</b> - loss of memory</li> <li>• <b>"io"</b> - loss of I/O channels</li> <li>• <b>"node"</b> - one or more books are no longer functioning</li> <li>• <b>"ring"</b> - the ring connecting the book is open</li> <li>• <b>"cbu"</b> - CBU resources have expired</li> <li>• <b>"mru"</b> - cycle time reduction due to an MRU problem</li> <li>• <b>"ambient-temp"</b> - cycle time reduction due to a temperature problem</li> <li>• <b>"iml"</b> - CPC was IMLed during cycle time reduction.</li> </ul>
<b>processor-running-time-type</b>	(w)	String Enum	Denotes how the processor-running-time property value was determined. One of: <ul style="list-style-type: none"> <li>• <b>"system-determined"</b> - the processor running time is dynamically determined by the system</li> <li>• <b>"user-determined"</b> - the processor running time is set to a constant value.</li> </ul> <p>Note: if <b>iml-mode</b> is not <b>"lpar"</b>, a null object is returned.</p>
<b>processor-running-time</b>	(w)	Integer	The amount of continuous time, in milliseconds, allowed for logical processors to perform jobs on shared processors for the CPC object. Note: a null object is returned if the <b>iml-mode</b> is not <b>"lpar"</b> or processor-running-time-type is <b>"system-determined"</b> .
<b>does-wait-state-end-time-slice</b>	(w)	Boolean	Logical Partitions of the CPC should lose their share of running time when they enter a wait state (true), or should not lose their share of running time when they enter a wait state (false). Note: a null object is returned if the <b>iml-mode</b> is not <b>"lpar"</b> or processor-running-time-type is <b>"system-determined"</b> .
<b>is-on-off-cod-installed</b>	—	Boolean	On/Off Capacity on Demand is installed for the CPC object (true), or is not installed (false). Note: if <b>status</b> is “not-communicating”, a null object is returned. Refer to the <i>Capacity On Demand User’s Guide</i> for details.
<b>is-on-off-cod-enabled</b>	—	Boolean	On/Off CoD is enabled (true), or is not enabled (false). Note: if <b>status</b> is “not-communicating”, a null object is returned. Refer to the <i>Capacity On Demand User’s Guide</i> for details.
<b>is-on-off-cod-activated</b>	—	Boolean	On/Off CoD is currently activated for the CPC object (true), or is not currently activated (false). Note: if <b>status</b> is “not-communicating”, a null object is returned. Refer to the <i>Capacity On Demand User’s Guide</i> for details.
<b>on-off-cod-activation-date</b>	—	Timestamp	The date when On/Off CoD was activated. Note: if <b>status</b> is “not-communicating”, a null object is returned. Refer to the <i>Capacity On Demand User’s Guide</i> for details.
<b>software-model-permanent</b>	(pc)	String (1-3)	The software model based on the permanent processors. For example, “700”.
<b>software-model-permanent-plus-billable</b>	(pc)	String (1-3)	The software model based on the permanent plus billable processors. For example, “700”.

Table 221. CPC object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>software-model-permanent-plus-temporary</b>	(pc)	String (1-3)	The software model based on the permanent plus all temporary processors. For example, "700".
<b>msu-permanent</b>	—	Integer	The MSU value associated with the software model based on the permanent processors.
<b>msu-permanent-plus-billable</b>	—	Integer	The MSU value associated with the software model based on the permanent plus billable processors.
<b>msu-permanent-plus-temporary</b>	—	Integer	The MSU value associated with the software model based on the permanent plus all temporary processors.
<b>processor-count-general-purpose</b>	—	Integer	The count of active general purpose processors.
<b>processor-count-service-assist</b>	—	Integer	The count of active service assist processors.
<b>processor-count-aap</b>	—	Integer	The count of active IBM zEnterprise Application Assist Processor (zAAP) processors.
<b>processor-count-ifl</b>	—	Integer	The count of active Integrated Facility for Linux (IFL) processors.
<b>processor-count-icf</b>	—	Integer	The count of active Internal Coupling Facility (ICF) processors.
<b>processor-count-iip</b>	—	Integer	The count of active IBM z Integrated Information Processor (zIIP) processors.
<b>processor-count-defective</b>	—	Integer	The count of defective processors. Includes all processor types.
<b>processor-count-spare</b>	—	Integer	The count of spare processors. Includes all processor types.
<b>processor-count-pending</b>	—	Integer	The combined number of processors that will become active, when more processors are made available by adding new hardware or by deactivating capacity records.
<b>processor-count-pending-general-purpose</b>	—	Integer	The number of general purpose processors that will become active, when more processors are made available by adding new hardware or by deactivating capacity records. Note: if status is "not-communicating", a null object is returned.
<b>processor-count-pending-service-assist</b>	—	Integer	The number of service assist processors that will become active, when more processors are made available by adding new hardware or by deactivating capacity records. Note: if status is "not-communicating", a null object is returned.
<b>processor-count-pending-aap</b>	—	Integer	The number of Application Assist processors that will become active, when more processors are made available by adding new hardware or by deactivating capacity records. Note: if status is "not-communicating", a null object is returned.
<b>processor-count-pending-ifl</b>	—	Integer	The number of Integrated Facility for Linux processors that will become active, when more processors are made available by adding new hardware or by deactivating capacity records. Note: if status is "not-communicating", a null object is returned.
<b>processor-count-pending-icf</b>	—	Integer	The number of Integrated Coupling Facility processors that will become active, when more processors are made available by adding new hardware or by deactivating capacity records. Note: if status is "not-communicating", a null object is returned.

Table 221. CPC object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>processor-count-pending-iip</b>	—	Integer	The number of Integrated Information processors that will become active, when more processors are made available by adding new hardware or by deactivating capacity records. Note: if status is "not-communicating", a null object is returned.
<b>has-temporary-capacity-change-allowed</b>	—	Boolean	Whether API applications are allowed to make changes to temporary capacity (true), or not (false).
<b>ec-mcl-description</b>	—	ec-mcl-description object	Describes the Engineering Change (EC) and MicroCode Level (MCL) for the CPC object. An empty object is returned if the information is unavailable from the SE. Refer to the description of the ec-mcl-description object for details.
<b>has-automatic-switch-enabled</b>	—	Boolean	Automatic switching between primary and alternate Support Elements is enabled for the CPC object (true), or is not enabled (false).  Support Element (Version 2.12.1 and newer) information can be found on the console help system, or on the IBM Knowledge Center at <a href="https://www.ibm.com/support/knowledgecenter">https://www.ibm.com/support/knowledgecenter</a> . (Select <b>z Systems</b> on the navigation bar, and then select your server). For information about earlier versions of the Support Element, see the <i>Support Element Operations Guide</i> .
<b>stp-configuration</b>	—	stp-config object	Describes the Server Time Protocol (STP) configuration. Refer to the description of the stp-config object for details. Note: if the required feature(s) are not installed, the property is not returned.
<b>lan-interface1-type</b>	(pc)	String Enum	The adapter type of the Support Element's LAN interface 1. One of the following: <ul style="list-style-type: none"> <li>• "ethernet"</li> <li>• "token-ring"</li> <li>• "unknown"</li> </ul>
<b>lan-interface1-address</b>	(pc)	String (1-12)	The MAC address of the Support Element's LAN interface 1.
<b>lan-interface2-type</b>	(pc)	String Enum	The adapter type of the Support Element's LAN interface 2. One of the following: <ul style="list-style-type: none"> <li>• "ethernet"</li> <li>• "token-ring"</li> <li>• "unknown"</li> </ul>
<b>lan-interface2-address</b>	(pc)	String (1-12)	The MAC address of the Support Element's LAN interface 2.
<b>network1-ipv4-mask</b>	(pc)	String (1-15)	The network IP mask value.
<b>network1-ipv4-pri-ipaddr</b>	(pc)	String IPV4 address	The primary IPv4 address or a null object if not configured.
<b>network1-ipv4-alt-ipaddr</b>	(pc)	String IPV4 address	The alternate IPv4 address or a null object if not configured.
<b>network1-ipv6-info</b>	—	Array of ipv6-info objects	A list of objects describing the Support Element's IPv6 network connections. If no IPv6 connections are defined, an empty list is returned.
<b>network2-ipv4-mask</b>	(pc)	String (1-15)	The network IP mask value.
<b>network2-ipv4-pri-ipaddr</b>	(pc)	String IPV4 address	The primary IPv4 address or a null object if not configured.

Table 221. CPC object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>network2-ipv4-alt- ipaddr</b>	(pc)	String IPV4 address	The alternate IPv4 address or a null object if not configured.
<b>network2-ipv6-info</b>	—	Array of ipv6-info objects	A list of objects describing the Support Element's IPv6 network connections. If no IPv6 connections are defined, an empty list is returned.
<b>hardware-messages</b>	(c)(pc)	Array of hardware- message objects	<p>The complete list of all CPC hardware messages, each identified by its URI. This list corresponds to the list provided by the <b>List CPC Hardware Messages</b> operation. If the CPC has no hardware messages, then an empty array is provided.</p> <p>The list of returned hardware messages can change as a result of new messages being dynamically added or removed by the infrastructure or due to hardware messages being deleted via the <b>Delete CPC Hardware Message</b> operation.</p> <p><b>Note:</b> This property is not returned by the <b>Get CPC Properties</b> operation, and only sessions associated with an HMC user with permission to the Hardware Messages task will receive a property-change notification for this property.</p>
<b>management- enablement-level</b>	(pc)	String Enum	<p>The zManager management enablement level for this system. The <b>management-enablement-level</b> values of an ensemble's nodes determine the ensemble's <b>management-enablement-level</b>, which determines the zManager advanced management functions that are available for use.</p> <p>All ensemble nodes must have the same <b>management-enablement-level</b>. A CPC with a <b>management-enablement-level</b> of "automate" may be added to an ensemble with a <b>management-enablement-level</b> of "manage", but this will downgrade the system's <b>management-enablement-level</b> until all ensemble nodes are upgraded to "automate".</p> <p>Values:</p> <ul style="list-style-type: none"> <li>• "manage" - provides the controls for managing an ensemble. It includes HMC operational controls for zBX, change management of firmware across the ensemble, energy monitoring, virtual networking with automated provisioning, virtual server management, and a base level of performance management.</li> <li>• "automate" - Provides more leverage from the ensemble by managing workloads and energy. This level of support includes power capping, power savings mode, and platform performance management.</li> </ul>

Table 222. ipv6-info object properties

Name	Type	Description
<b>type</b>	String Enum	The IPv6 scope. One of the following values: <ul style="list-style-type: none"> <li>• "link-local"</li> <li>• "static"</li> <li>• "auto"</li> </ul>
<b>prefix</b>	Integer	The number of leading bits of the IPv6 address that represent the network prefix
<b>pri-ip-address</b>	String IPV6 address	The primary IPv6 address

Table 222. *ipv6-info* object properties (continued)

Name	Type	Description
<b>alt-ip-address</b>	String IPV6 address	The alternate IPv6 address or a null object if not configured

Table 223. *hardware-message* object properties

Name	Type	Description
<b>element-uri</b>	String/URI	The canonical URI path of the CPC hardware message. The URI is in the following form: <code>/api/cpcs/{cpc-id}/hardware-messages/{hardware-message-id}</code> , where <code>{hardware-message-id}</code> is the value of the <b>element-id</b> property of the hardware message.
<b>element-id</b>	String (36)	The unique identifier for the hardware message. The <b>element-id</b> is in the form of a UUID.
<b>parent</b>	String/URI	The parent of a CPC hardware message is the CPC object. The <b>parent</b> value is the canonical URI path for the CPC.
<b>class</b>	String	The <b>class</b> of a hardware message object is " <b>hardware-message</b> ".
<b>timestamp</b>	Timestamp	The <b>timestamp</b> represents the date and time when the hardware message was created.
<b>text</b>	String	The text of the hardware message.

### Energy management related additional properties

In addition to the properties defined above, this object includes the following additional class-specific properties related to energy management. For further explanation of the various states involved, please see "Special states" on page 181.

For definitions of the qualifier abbreviations in the following tables, see "Property characteristics" on page 38.

Table 224. *CPC* object: energy management related additional properties

Name	Qualifier	Type	Description
<b>cpc-power-rating</b>	—	Integer	Specifies the maximum power draw in watts (W) of this CPC. This is a calculated value as indicated by the electrical rating labels or system rating plates of the CPC components.
<b>cpc-power-consumption</b>	(mg)	Integer	Specifies the current power consumption in watts (W) for this CPC. The CPC power consumption includes the power consumption of the zCPC and BladeCenters. The BladeCenter power consumption includes the power consumption of the blades contained within the BladeCenter. If the system does not include a BladeCenter, the CPC power consumption will be equal to the zCPC power consumption.

Table 224. CPC object: energy management related additional properties (continued)

Name	Qualifier	Type	Description
<b>cpc-power-saving</b>	—	String Enum	<p>Specifies the current power saving setting of the CPC. Power saving is used to reduce the energy consumption of a system and can be managed in the Set Power Saving operation. The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - The power consumption and performance of the CPC are not reduced. This is the default setting.</li> <li>• <b>"low-power"</b> - All components of the CPC enabled for power saving will have reduced performance to allow for low power consumption.</li> <li>• <b>"custom"</b> - Custom mode indicates that some, but not all, components of the CPC are in the Low power setting.</li> <li>• <b>"not-supported"</b> - Power saving is not supported for this CPC.</li> <li>• <b>"not-available"</b> - Specifies that <b>cpc-power-saving</b> property could not be read from this CPC.</li> <li>• <b>"not-entitled"</b> - The server is not entitled for Power saving.</li> </ul>
<b>cpc-power-saving-state</b>	—	String Enum	<p>Specifies the power saving setting of the CPC set by the user. Please note that this property indicates the user setting and may not match the real state of the hardware compared to the <b>cpc-power-saving</b> property. For more information, see "Group power saving" on page 182. The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - Specifies not reducing the power consumption and performance of the CPC.</li> <li>• <b>"low-power"</b> - Specifies low power consumption for all components of the CPC enabled for power saving.</li> <li>• <b>"custom"</b> - Specifies that the CPC does not control the children. This is the default setting.</li> <li>• <b>"not-supported"</b> - Specifies that power saving is not supported for this CPC.</li> <li>• <b>"not-entitled"</b> - Specifies that the server is not entitled to power saving.</li> </ul>
<b>cpc-power-save-allowed</b>	—	String Enum	<p>Should be used to determine if a call of the power save operation is currently allowed. If a value other than <b>"allowed"</b> is returned the caller may reckon that the power save operation will fail.</p> <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Alter power save setting is allowed for this CPC</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-supported"</b> - Power saving is not supported for this CPC.</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power capping.</li> </ul>

Table 224. CPC object: energy management related additional properties (continued)

Name	Qualifier	Type	Description
<b>cpc-power-capping-state</b>	—	String Enum	Specifies the current power capping setting of the CPC. Power capping is used to limit peak power consumption of a system and can be managed in the Set Power Cap operation. The possible settings include: <ul style="list-style-type: none"> <li>• <b>"disabled"</b> - The power cap of the CPC is not set and the peak power consumption is not limited. This is the default setting.</li> <li>• <b>"enabled"</b> - All components of the CPC available for power capping will be capped to limit the peak power consumption of the CPC.</li> <li>• <b>"custom"</b> - The components of the CPC can be individually configured for power capping.</li> <li>• <b>"not-supported"</b> - Power capping is not supported for this CPC.</li> <li>• <b>"not-entitled"</b> - The server is not entitled for Power capping.</li> </ul>
<b>cpc-power-cap-minimum</b>	—	Integer	Specifies the minimum value for the CPC cap value in watts (W). This is a sum of the component minimum cap values.
<b>cpc-power-cap-maximum</b>	—	Integer	Specifies the maximum value for the CPC cap value in watts (W). This is a sum of the component maximum cap values.
<b>cpc-power-cap-current</b>	—	Integer	Specifies the current cap value for the CPC in watts (W). The current cap value indicates the power budget for the CPC and is the sum of the component Cap values.
<b>cpc-power-cap-allowed</b>	—	String Enum	Should be used to determine if a call of the power capping operation is currently allowed. If a value other than <b>"allowed"</b> is returned the caller may reckon that the power capping operation will fail. <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b>- Alter power capping setting is allowed for this CPC</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-supported"</b> - Power capping is not supported for this CPC.</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power capping.</li> </ul>
<b>zpcp-power-rating</b>	—	Integer	Specifies the maximum power draw in watts (W) of this zCPC. This is a calculated value as indicated by the electrical rating labels or system rating plates of the zCPC components.
<b>zpcp-power-consumption</b>	(mg)	Integer	Specifies the current power consumption of the zCPC in watts (W).
<b>zpcp-power-saving</b>	—	String Enum	Specifies the current power saving setting of the zCPC. Power saving is used to reduce the energy consumption of a system and can be managed in the Set Power Saving operation. The possible settings include: <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - The power consumption and performance of the zCPC are not reduced. This is the default setting.</li> <li>• <b>"low-power"</b> - The performance of the zCPC is reduced to allow for low power consumption.</li> <li>• <b>"not-supported"</b> - Power saving is not supported for this zCPC.</li> <li>• <b>"not-available"</b> - Specifies that <b>zpcp-power-saving</b> property could not be read for this zCPC.</li> <li>• <b>"not-entitled"</b> - The server is not entitled for Power saving.</li> </ul>



Table 224. CPC object: energy management related additional properties (continued)

Name	Qualifier	Type	Description
<b>zcpc-power-saving-state</b>	—	String Enum	Specifies the power saving setting of the zCPC set by the user. Please note that this property indicates the user setting and may not match the real state of the hardware compared to the <b>zcpc-power-saving</b> property. For more information, see “Group power saving” on page 182. The possible settings include: <ul style="list-style-type: none"> <li>• <b>"high-performance"</b> - Specifies not reducing the power consumption and performance of the zCPC. This is the default setting.</li> <li>• <b>"low-power"</b> - Specifies low power consumption for all components of the zCPC enabled for power saving.</li> <li>• <b>"not-supported"</b> - Specifies that power saving is not supported for this zCPC.</li> <li>• <b>"not-entitled"</b> - Specifies that the server is not entitled to power saving.</li> </ul>
<b>zcpc-power-save-allowed</b>	—	String Enum	Should be used to determine if a call of the power save operation is currently allowed. If a value other than <b>"allowed"</b> is returned the caller may reckon that the power save operation will fail. <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Alter power save is allowed for this zCPC</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power save.</li> <li>• <b>"under-group-control"</b> - The zCPC is under group control and cannot be individually altered.</li> <li>• <b>"not-supported"</b> - Power saving is not supported for this zCPC.</li> <li>• <b>"once-a-day-exceeded"</b> - Power saving mode has been entered at some point during the day and will not be allowed again until the next calendar day.</li> </ul>
<b>zcpc-power-capping-state</b>		String Enum	Specifies the current power capping setting of the zCPC. Power capping is used to limit peak power consumption of a system and can be managed in the Set Power Cap operation. The possible settings include: <ul style="list-style-type: none"> <li>• <b>"disabled"</b> - The power cap of the zCPC is not set and the peak power consumption is not limited. This is the default setting.</li> <li>• <b>"enabled"</b> - The peak power consumption of the zCPC is limited to the Current cap value.</li> <li>• <b>"custom"</b> - The components of the CPC can be individually configured for power capping.</li> <li>• <b>"not-supported"</b> - Power capping is not supported for this zCPC.</li> <li>• <b>"not-entitled"</b> - The server is not entitled for Power capping.</li> </ul>
<b>zcpc-power-cap-minimum</b>	—	Integer	Specifies the minimum value for the zCPC cap value in watts (W).
<b>zcpc-power-cap-maximum</b>	—	Integer	Specifies the maximum value for the zCPC cap value in watts (W).
<b>zcpc-power-cap-current</b>	—	Integer	Specifies the current cap value for the CPC in watts (W). The current cap value indicates the power budget for the zCPC.

Table 224. CPC object: energy management related additional properties (continued)

Name	Qualifier	Type	Description
zpcp-power-cap-allowed	—	String Enum	<p>Should be used to determine if a call of the power capping operation is currently allowed. If a value other than <b>"allowed"</b> is returned the caller may reckon that the power capping operation will fail.</p> <p>The possible settings include:</p> <ul style="list-style-type: none"> <li>• <b>"allowed"</b> - Alter power capping is allowed for this zCPC</li> <li>• <b>"unknown"</b> - Unknown reason</li> <li>• <b>"not-entitled"</b> - Specifies the server is not entitled to power save.</li> <li>• <b>"not-supported"</b> - Power capping is not supported for this zCPC.</li> <li>• <b>"under-group-control"</b> - Power capping is under group control</li> </ul>
zpcp-ambient-temperature	(mg)	Float	Specifies the input air temperature in degrees Celsius (°C) as measured by the system.
zpcp-exhaust-temperature	—	Float	Specifies the exhaust air temperature in degrees Celsius (°C) as calculated by the system. This is useful in determining potential hot spots in the data center.
zpcp-humidity	(mg)	Integer	<p>Specifies the amount of water vapor in the air as measured by the system. The humidity sensor gives a reading of the relative humidity of the air entering the system. The recommended long-term relative humidity for a system with an altitude from sea level to 900 meters (2953 feet) is 60%. The range of acceptable relative humidity is 8% - 80%.</p> <p>For more information, see the chapter related to environmental specifications in the <i>Installation Manual for Physical Planning</i>.</p>
zpcp-dew-point	(mg)	Float	<p>Specifies the air temperature in degrees Celsius (°C) at which water vapor will condense into water. This is a calculated value based on the current temperature and relative humidity. Cooling the server to the dew point can result in condensation on critical internal parts, leading to equipment failure, unless the computer room environment is adequately maintained to prevent it.</p> <p>For more information, see the chapter related to environmental specifications in the <i>Installation Manual for Physical Planning</i>.</p>
zpcp-heat-load	(mg)	Integer	Specifies the amount of heat in Btu/hr. removed from the system.
zpcp-heat-load-forced-air	—	Integer	Specifies the amount of heat in Btu/hr. removed from the system by forced-air.
zpcp-heat-load-water	—	Integer	Specifies the amount of heat in Btu/hr. removed from the system by chilled water. The value is always 0 on an air cooled system.
zpcp-maximum-potential-power	—	Integer	Specifies the maximum potential power consumption of a system in watts (W). This value is based on the configuration of the system and can be used for power and cooling planning.
zpcp-maximum-potential-heat-load	—	Integer	Specifies the maximum potential heat load of a system in Btu/hr. This value is based on the configuration of the system and can be used for power and cooling planning.

## List CPC Objects

The **List CPC Objects** operation returns a list of the zManager Web Services API capable CPCs managed by an HMC.

## HTTP method and URI

GET /api/cpcs

### Query Parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	A regular expression used to limit returned objects to those that have a matching name property. If matches are found, the response will be an array with all objects that match. If no match is found, the response will be an empty array.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
cpcs	Array of cpc-info objects	Array of nested cpc-info objects (described in the next table). If no matching CPC objects are found, an empty array is returned.

Each nested cpc-info object contains the following fields:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the CPC object
name	String	The name of the CPC object
status	String Enum	The current status of the CPC object

### Description

This operation lists the zManager Web Services API capable CPC objects that are managed by this HMC. The object URI, object ID and display name are provided for each CPC returned. CPCs that are not zManager Web Services API capable are not returned.

If the name query parameter is specified, the returned list is limited to those CPC objects that have a name property matching the specified filter pattern. If the name parameter is omitted, this filtering is not done.

An object is only included in the list if the API user has object-access permission for that object.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in the Response Body Contents section.

### Authorization requirements

This operation has the following authorization requirement:

- Object access permission to any CPC object to be included in the result.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	299	A query parameter has an invalid syntax.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/cpcs HTTP/1.1
x-api-session: 2jm2h7j25d1e1g5wbygmfriyjiiit8tp4iqiw8h09j8kz68i0k6
```

---

Figure 395. List CPC Objects: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 07:18:42 GMT
content-type: application/json;charset=UTF-8
content-length: 492
{
  "cpcs": [
    {
      "name": "P0LXSMOZ",
      "object-uri": "/api/cpcs/e8753ff5-8ea6-35d9-b047-83c2624ba8da",
      "status": "not-operating"
    },
    {
      "name": "R32",
      "object-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
      "status": "operating"
    },
    {
      "name": "ICHABOD",
      "object-uri": "/api/cpcs/ac15c987-90c6-3526-854e-4c612939260d",
      "status": "not-operating"
    }
  ]
}
```

---

Figure 396. List CPC Objects: Response

## List Ensemble CPC Objects

The **List Ensemble CPC Objects** operation returns a list of the CPCs associated with an ensemble.

### HTTP method and URI

```
GET /api/ensembles/{ensemble-id}/cpcs
```

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

**Query parameters:**

Name	Type	Rqd/Opt	Description
name	String	Optional	A regular expression used to limit returned objects to those that have a matching name property. If matches are found, the response will be an array with all objects that match. If no match is found, the response will be an empty array.

**Response body contents**

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
cpcs	Array of cpc-info objects	Array of nested cpc-info objects (described in the next table). If no matching CPC objects are found, an empty array is returned.

Each nested cpc-info object contains the following fields:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the CPC object
name	String	The name of the CPC object
status	String Enum	The current status of the CPC object

**Description**

This operation lists the CPC objects that are associated with the ensemble identified in the request URI. The object URI, object ID and display name are provided for each CPC returned. In contrast, the List Ensemble Nodes operation lists objects of any type that are associated with an ensemble.

If the **name** query parameter is specified, the returned list is limited to those CPC objects that have a **name** property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

A CPC object is only included in the list if the API user has object-access permission for that CPC object.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in the Response Body Contents section.

**Authorization requirements**

This operation has the following authorization requirements:

- Object access permission to the ensemble object designated by *{ensemble-id}*
- Object access permission to any CPC object to be included in the result.

**HTTP status and reason codes**

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>ensemble-id</i> ) does not designate an existing ensemble object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334/cpcs HTTP/1.1
x-api-session: 2jm2h7j25d1e1g5wbygmfrijjiit8tp4iqiw8h09j8kz68i0k6
```

---

Figure 397. List Ensemble CPC Objects: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 07:18:42 GMT
content-type: application/json;charset=UTF-8
content-length: 214
{
  "cpcs": [
    {
      "name": "R32",
      "object-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
      "status": "operating"
    },
    {
      "name": "ICHABOD",
      "object-uri": "/api/cpcs/ac15c987-90c6-3526-854e-4c612939260d",
      "status": "not-operating"
    }
  ]
}
```

---

Figure 398. List Ensemble CPC Objects: Response

## Get CPC Properties

The **Get CPC Properties** operation retrieves the properties of a single CPC object designated by *{cpc-id}*.

### HTTP method and URI

```
GET /api/cpcs/{cpc-id}
```

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Response body contents

On successful completion, the response body provides the current values of the properties for the CPC object as defined in “Data model” on page 743.

## Description

Some CPC properties are only available if the HMC is communicating with the SE, and are returned as null objects if the HMC is not communicating with the SE.

On successful execution, HTTP status code 200 (OK) is returned and the response body contains all of the current properties as defined “Data model” on page 743.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 758.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	272	Unable to obtain Server Time Protocol (STP) configuration. Retry the request later.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340 HTTP/1.1
x-api-session: 65aw2jahugn1wop51hsq0c6aldkx773dz9ulirrvvg2z853m4u
```

---

Figure 399. Get CPC Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 16:58:36 GMT
content-type: application/json;charset=UTF-8
content-length: 9001
{
  "acceptable-status": [
    "operating"
  ],
  "additional-status": "",
  "cbu-activation-date": 0,
  "cbu-expiration-date": 0,
  "cbu-number-of-tests-left": 0,
  "class": "cpc",
  "cpc-node-descriptor": "00",
  "cpc-power-cap-allowed": "allowed",
  "cpc-power-cap-current": 90000,
  "cpc-power-cap-maximum": 106936,
  "cpc-power-cap-minimum": 16019,
  "cpc-power-capping-state": "enabled",
  "cpc-power-consumption": 8160,
  "cpc-power-rating": 46522,
  "cpc-power-save-allowed": "allowed",
  "cpc-power-saving": "high-performance",
  "cpc-serial-number": "000020076D25",
  "degraded-status": [
    "not-degraded"
  ],
  "description": "Central Processing Complex (CPC)",
  "does-wait-state-end-time-slice": null,
  "ec-mcl-description": {
    "action": [
      {
        "activation": "current",
        "pending": false,
        "type": "channel-config"
      },
      {
        "activation": "current",
        "pending": false,
        "type": "coupling-facility-reactivation"
      },
      {
        "activation": "current",
        "pending": false,
        "type": "power-on-reset-tracking"
      },
      {
        "activation": "current",
        "pending": true,
        "type": "zhybrid-blades-activation"
      },
      {
        "activation": "next",
        "pending": false,
        "type": "channel-config"
      },
      {
```

---

Figure 400. Get CPC Properties: Response (Part 1)



```

    "activation": "next",
    "pending": false,
    "type": "coupling-facility-reactivation"
  },
  {
    "activation": "next",
    "pending": false,
    "type": "power-on-reset-tracking"
  },
  {
    "activation": "next",
    "pending": false,
    "type": "zhybrid-blades-activation"
  }
],
"ec": [
  {
    "description": "SE Framework",
    "mcl": [
      {
        "last-update": 1321019308748,
        "level": "216",
        "type": "retrieved"
      },
      {
        "last-update": 1321019499749,
        "level": "216",
        "type": "activated"
      },
      {
        "last-update": 1320675688749,
        "level": "179",
        "type": "accepted"
      },
      {
        "last-update": 943938000748,
        "level": "216",
        "type": "installable-concurrent"
      },
      {
        "last-update": 943938000749,
        "level": "180",
        "type": "removable-concurrent"
      }
    ],
    "number": "N48168",
    "part-number": "45D8918",
    "type": "Base EC"
  },
  {
    "description": "IBM x86 Blade Concurrent Components",
    "mcl": [
      {
        "last-update": 1321019296651,
        "level": "022",
        "type": "retrieved"
      },
      {

```

Figure 401. Get CPC Properties: Response (Part 2)

---

```
        "last-update": 1320675685652,
        "level": "009",
        "type": "accepted"
      },
      {
        "last-update": 943938000652,
        "level": "022",
        "type": "installable-concurrent"
      },
      {
        "last-update": 943938000652,
        "level": "010",
        "type": "removable-concurrent"
      }
    ],
    "number": "N48140",
    "part-number": "41U8008",
    "type": "Other Optional EC"
  },
  {
    "description": "Embedded Operating System T5xx Series",
    "mcl": [
      {
        "last-update": 1321018125779,
        "level": "001",
        "type": "retrieved"
      },
      {
        "last-update": 1321018167779,
        "level": "001",
        "type": "activated"
      },
      {
        "last-update": null,
        "level": "000",
        "type": "accepted"
      },
      {
        "last-update": 943938000779,
        "level": "001",
        "type": "installable-concurrent"
      },
      {
        "last-update": 943938000779,
        "level": "001",
        "type": "removable-concurrent"
      }
    ],
    "number": "N48197",
    "part-number": "45D8919",
    "type": "Base EC"
  }
],
},
```

---

Figure 402. Get CPC Properties: Response (Part 3)

```

"has-automatic-se-switch-enabled": true,
"has-hardware-messages": true,
"has-temporary-capacity-change-allowed": false,
"has-unacceptable-status": false,
"iml-mode": "lpar",
"is-cbu-activated": false,
"is-cbu-enabled": true,
"is-cbu-installed": false,
"is-ensemble-member": true,
"is-locked": false,
"is-on-off-cod-activated": false,
"is-on-off-cod-enabled": true,
"is-on-off-cod-installed": false,
"is-real-cbu-available": false,
"is-service-required": false,
"lan-interfacel-address": "f0def14b63af",
"lan-interfacel-type": "ethernet",
"lan-interface2-address": "f0def14b63af",
"lan-interface2-type": "ethernet",
"last-used-activation-profile": "DEFAULT",
"machine-model": "M15",
"machine-serial-number": "000020076D25",
"machine-type": "2817",
"management-enablement-level": "automate",

"msu-permanent": 1091,
"msu-permanent-plus-billable": 1091,
"msu-permanent-plus-temporary": 1091,
"name": "R32",
"network1-ipv4-alt-ipaddr": "9.60.15.6",
"network1-ipv4-mask": "255.255.255.0",
"network1-ipv4-pri-ipaddr": "9.60.15.5",
"network1-ipv6-info": [
  {
    "alt-ip-address": "fdd8:673b:d89b:1:f2de:f1ff:fe52:d359",
    "prefix": 64,
    "pri-ip-address": "fdd8:673b:d89b:1:f2de:f1ff:fe4b:63af",
    "type": "auto"
  },
  {
    "alt-ip-address": "fe80::f2de:f1ff:fe52:d359%eth0",
    "prefix": 64,
    "pri-ip-address": "fe80::f2de:f1ff:fe4b:63af%eth0",
    "type": "link-local"
  }
],
"network2-ipv4-alt-ipaddr": "9.60.14.6",
"network2-ipv4-mask": "255.255.255.0",
"network2-ipv4-pri-ipaddr": "9.60.14.5",
"network2-ipv6-info": [
  {
    "alt-ip-address": "fe80::f2de:f1ff:fe52:d359%eth1",
    "prefix": 64,
    "pri-ip-address": "fe80::f2de:f1ff:fe4b:63af%eth1",
    "type": "link-local"
  }
],

```

Figure 403. Get CPC Properties: Response (Part 4)

---

```

"next-activation-profile-name": "TESTCDU",
"object-id": "37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
"object-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
"on-off-cod-activation-date": 0,
"parent": "/api/ensembles/f8fc3a9c-03f2-11e1-ba83-0010184c8334",
"processor-count-aap": 1,
"processor-count-defective": 0,
"processor-count-general-purpose": 9,
"processor-count-icf": 2,
"processor-count-ifl": 2,
"processor-count-iip": 1,
"processor-count-pending": 0,
"processor-count-pending-aap": 0,
"processor-count-pending-general-purpose": 0,
"processor-count-pending-icf": 0,
"processor-count-pending-ifl": 0,
"processor-count-pending-iip": 0,
"processor-count-pending-service-assist": 0,
"processor-count-service-assist": 3,
"processor-count-spare": 0,
"processor-running-time": null,
"processor-running-time-type": "system-determined",
"se-version": "2.11.1",
"software-model-permanent": "709",
"software-model-permanent-plus-billable": "709",
"software-model-permanent-plus-temporary": "709",
"status": "operating",
"stp-configuration": {
  "current-time-server": "preferred",
  "etr-id": null,
  "stp-id": "12345678"
},
"zpcp-ambient-temperature": 25.399999618530273,
"zpcp-dew-point": 2.4000000953674316,
"zpcp-exhaust-temperature": 34.0,
"zpcp-heat-load": 20293,
"zpcp-heat-load-forced-air": 20293,
"zpcp-heat-load-water": 0,
"zpcp-humidity": 22,
"zpcp-maximum-potential-heat-load": 26571,
"zpcp-maximum-potential-power": 7782,
"zpcp-power-cap-allowed": "under-group-control",
"zpcp-power-cap-current": 23745,
"zpcp-power-cap-maximum": 27400,
"zpcp-power-cap-minimum": 7782,
"zpcp-power-capping-state": "enabled",
"zpcp-power-consumption": 5943,
"zpcp-power-rating": 27400,
"zpcp-power-save-allowed": "under-group-control",
"zpcp-power-saving": "high-performance"
}

```

---

Figure 404. Get CPC Properties: Response (Part 5)

## Update CPC Properties

The **Update CPC Properties** operation updates one or more writeable properties of the CPC object designated by *{cpc-id}*.

## HTTP method and URI

POST /api/cpcs/{cpc-id}

In this request, the URI variable {cpc-id} is the object ID of the target CPC object.

## Request body contents

The request body is expected to contain one or more field names representing writable CPC properties, along with the new values for those fields.

The request body can and should omit fields for properties whose values are not to be changed by this operation. Properties for which no input value is provided remain unchanged by this operation.

## Description

The request body object is validated against the data model for the CPC object type to ensure that the request body contains only writeable properties and the data types of those properties are as required. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

On successful execution, the value of each corresponding property of the object is updated with the value provided by the input field, and status code 204 (No Content) is returned.

When this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by {cpc-id}
- Action/task permission for the **CPC Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	268	The requested update requires that the <b>processor-running-time-type</b> property already contain <b>"user-determined"</b> or that the request body also requests an update of the <b>processor-running-time-type</b> property to <b>"user-determined"</b> .
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ({cpc-id}) does not designate an existing CPC object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Activate CPC

The **Activate CPC** operation activates the CPC object designated by *{cpc-id}*.

### HTTP method and URI

**POST** `/api/cpcs/{cpc-id}/operations/activate`

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
activation-profile-name	String (1-16)	Optional	The name of the activation profile to be used for the request. If not provided, the request uses the profile name specified in the activation-profile-name property for the CPC object.
force	Boolean	Optional	Whether this operation is permitted when the CPC is in <b>"operating"</b> status (true) or not (false). The default is false.

### Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve activation status updates.

### Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is **"complete"**, the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in "Description." The **job-results** field is null for this operation.

### Description

Activation is a process that makes a CPC operational, which means either:

- The CPC is ready to have a control program or operating system loaded, or
- The CPC has loaded and is running a control program or operating system.

Activation makes a CPC operational by:

- Using predefined information, referred to as an activation profile, to set the operational capabilities and characteristics of the CPC
- Checking the current status of the CPC, and then performing only the operations necessary to make it operational as specified in the activation profile.

So, using activation is not limited to starting the system. Using activation is recommended whenever you want to make the CPC or its logical partitions operational.

A complete activation activates the CPC and its logical partitions completely in a single step. The result of a complete activation is an operational CPC with logical partitions loaded and running operating systems. The current status of the CPC and its logical partitions determines which operations are performed during activation to make them operational. Activation may include:

1. Turning CPC power on.
2. Performing a power-on reset, this includes allocating system resources to the CPC.
3. Then activating logical partitions to support multiple images. Activating each logical partition includes:
  - a. Initializing it.
  - b. Allocating system resources to it.
  - c. Loading it with a control program or operating system.

Because the status of the CPC and its logical partitions determines which operations must be performed during activation to make them operational, one or more operations listed above may not be performed during activation. For example:

- Activating the CPC does not perform a power-on reset if the CPC has already been power-on reset and the applicable settings in its assigned activation profile, such as the operating mode and active input/output configuration data set (IOCDs), are already in effect.
- Activating the CPC does not perform any operations if the CPC is already operational and all settings in its assigned activation profile are already in effect.

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See “Query Job Status” on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent, with Job Status and Reason Codes described in “Job status and reason codes” on page 768.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission for the **Activate** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 766.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Job status and reason codes

Job status code	Job reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed or was rejected due to the current CPC status and use of the <code>force=false</code> parameter. If rejected due to <code>force=false</code> , the CPC status is unchanged. If the operation failed, the CPC status is unknown. Refer to the message parameter in the error response body for details.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Deactivate CPC

The **Deactivate CPC** operation deactivates the CPC object designated by `{cpc-id}`.

#### HTTP method and URI

**POST** `/api/cpcs/{cpc-id}/operations/deactivate`

In this request, the URI variable `{cpc-id}` is the object ID of the target CPC object.

#### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
<code>force</code>	Boolean	Optional	Whether this operation is permitted when the CPC is in <b>"operating"</b> status (true) or not (false). The default is false.

#### Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
<code>job-uri</code>	String/URI	URI that may be queried to retrieve activation status updates.

#### Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is **"complete"**, the results include a job completion status code and reason code (fields `job-status-code` and `job-reason-code`) which are set as indicated in “Job status and reason codes” on page 769. The `job-results` field is null for this operation.



## Description

Deactivation is an orderly process for shutting down and turning off the CPC.

Shutting down and turning off the CPC, referred to also as deactivating the CPC, includes:

- Ending hardware and software activity
- Clearing, releasing, and de-allocating hardware resources
- Turning off power.

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See “Query Job Status” on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent, with Job Status and Reason Codes seen in “Job status and reason codes.”

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission for the **Deactivate** task.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 768.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

Job status code	Job reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed or was rejected due to the current CPC status and use of the <code>force=false</code> parameter. If rejected due to <code>force=false</code> , the CPC status is unchanged. If the operation failed, the CPC status is unknown. Refer to the message parameter in the error response body for details.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Import Profiles

The **Import Profiles** operation imports activation profiles and/or system activity profiles for the CPC from the SE hard drive into the CPC object designated by *{cpc-id}*.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/import-profiles

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
profile-area	Integer (1-4)	Required	The numbered hard drive area from which the profiles are imported. Use the profile-area value specified on the prior <b>Export Profiles</b> operation.

### Description

The Support Element provides four reusable areas on its hard drive from which the data save by a prior Export Profiles can be read.

Exporting and importing profiles is necessary only when you intend to have your current system and Support Element replaced with a new system and Support Element. Support Element (Version 2.12.1 and newer) information can be found on the console help system, or on the IBM Knowledge Center at <https://www.ibm.com/support/knowledgecenter>. (Select **z Systems** on the navigation bar, and then select your server). For information about earlier versions of the Support Element, see the *Support Element Operations Guide*.

On success, HTTP status code 204 (No Content) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission to the CIM Actions ExportSettingsData task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.

HTTP error status code	Reason code	Description
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
500 (Server Error)	279	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Export Profiles

The **Export Profiles** operation exports activation profiles and/or system activity profiles from the CPC object designated by *{cpc-id}* to the SE hard drive.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/export-profiles

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
profile-area	Integer (1-4)	Required	The numbered hard drive area to which the profiles are exported. Any existing data is overwritten.

### Description

The Support Element provides four reusable areas on its hard drive that can be used as temporary save areas. The choice of save area is up to the caller.

Exporting and importing profiles is necessary only when you intend to have your current system and Support Element replaced with a new system and Support Element. Support Element (Version 2.12.1 and newer) information can be found on the console help system, or on the IBM Knowledge Center at <https://www.ibm.com/support/knowledgecenter>. (Select **z Systems** on the navigation bar, and then select your server). For information about earlier versions of the Support Element, see the *Support Element Operations Guide*.

On success, HTTP status code 204 (No Content) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission to the CIM Actions ExportSettingsData task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
500 (Server Error)	279	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Add Temporary Capacity

The **Add Temporary Capacity** operation adds temporary processors or increases temporary model capacity to the CPC object designated by *{cpc-id}*.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/add-temp-capacity

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
record-id	String (1-8)	Required	Identifies the capacity record to be used for this request
software-model	String (1-3)	Optional	The target software model. This value must be one of the software models defined within the capacity record. Implicit in this value is the number of general processors added. If not provided, the current software model is not changed.
processor-info	Array of processor-info objects	Optional	A nested object that defines the number of specialty processors to be added. If not provided, the number of specialty processors is not changed.
force	Boolean	Optional	Whether the operation proceeds if not enough processors are available (true) or not (false). The default is false.
test	Boolean	Required	Whether the request should activate real or test resources for the capacity record. Set true if test or set to false if real. This is mainly used for Capacity Backup Upgrade (CBU) activations. See the <i>Capacity on Demand User's Guide</i> , SC28-6943. For most records, field should be set to false.

### processor-info object

Field name	Type	Rqd/Opt	Description
processor-type	String Enum	Required	Identifies the type of specialty processors to be affected. One of: <ul style="list-style-type: none"> <li>• "aap" - Application Assist Processor</li> <li>• "ifl" - Integrated Facility for Linux processor</li> <li>• "icf" - Internal Coupling Facility processor</li> <li>• "iip" - Integrated Information Processors</li> <li>• "sap" - System Assist Processor</li> </ul>
num-processor-steps	Integer	Optional	The delta to the current number of processors. If not provided, the number of processors is not changed.

## Description

Removal of these temporary resources can be performed manually via the Remove Temporary Capacity operation or automatically upon expiration of the capacity record.

Refer to the *Capacity on Demand User's Guide* for details on temporary capacity changes.

On success, HTTP status code 204 (No Content) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission to the CIM Actions Activate task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	271	A duplicate processor-type entry was found in the processor-info array, remove the duplicate entry.
	275	The test value does not match the value stored in the capacity record.
	276	Either the request specifies more resources than available or the requested software model specifies fewer resources than the current software model.
	277	A temporary capacity record is already active. It must be deactivated before a new capacity record can be activated.
	278	The software-model value was not found in the capacity record. Only software models as defined in the target capacity record can be specified.
	298	The operation parameters conflict with the capacity record type: <ul style="list-style-type: none"> <li>• force=true is permitted only for CBU, CPE and loaner capacity records.</li> </ul>
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	274	The requested capacity record does not exist.

HTTP error status code	Reason code	Description
409 (Conflict)	1	The operation is unavailable in the current CPC state: <ul style="list-style-type: none"> <li>The SE is not configured to allow temporary capacity changes via an API</li> <li>The CPC <b>status</b> property is not <b>"operating"</b> or the CPC <b>iml-mode</b> property is not <b>"Ipar"</b></li> <li>No physical processors are operating</li> <li>The CPC is IMLed in a test or debug mode</li> <li>An IML is required</li> </ul>
	2	The operation was rejected by the Support Element (SE), because the SE is currently performing processing that requires exclusive control of the SE. Retry the operation at a later time.
	297	Some, but not all, of the requested resources were added.
500 (Server Error)	275	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Remove Temporary Capacity

The **Remove Temporary Capacity** operation removes temporary processors or decreases temporary model capacity from the CPC object designated by *{cpc-id}*.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/remove-temp-capacity

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
record-id	String (1-8)	Required	Identifies the capacity record to be used for this request
software-model	String (1-3)	Optional	The target software model. This value must be one of the software models defined within the capacity record. Implicit in this value is the number of general processors removed. If not provided, the current software model is not changed.
processor-info	Array of processor-info objects	Optional	A nested object that defines the number of specialty processors to be removed. If not provided, the number of specialty processors is not changed. Refer to "Request body contents" on page 772 of the <b>Add Temporary Capacity</b> operation for details.

### Description

When you are finished using all or part of a capacity upgrade, you can remove processors or decrease model capacity using this operation. You can only remove activated resources for the specific offering. You cannot remove dedicated engines or the last processor of a processor type.

If you remove resources back to the base configuration, the capacity record activation is completed. That is, if you remove the last temporary processor, your capacity record is deactivated. For a CBU and On/Off CoD record, to add resources again, you must use another **Add Temporary Capacity** operation. For an On/Off CoD test or CPE record, once the record is deactivated, it is no longer available for use. You can then delete the record.

After removal of the resources, the capacity record remains as an installed record. If you want a record deleted, you must manually select the record on the Installed Records page and click Delete.

Refer to the *Capacity on Demand User's Guide* for details on temporary capacity changes.

On success, HTTP status code 204 (No Content) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission to the CIM Actions Deactivate task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	271	A duplicate processor-type entry was found in the processor-info array, remove the duplicate entry.
	276	Either the request specifies more resources than are currently active or the requested software model specifies more resources than the current software model.
	278	The software-model value was not found in the capacity record. Only software models as defined in the target capacity record can be specified.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	274	The requested capacity record does not exist.
409 (Conflict)	1	The operation is unavailable in the current CPC state: <ul style="list-style-type: none"> <li>• The SE is not configured to allow temporary capacity changes via an API</li> <li>• The CPC <b>status</b> property is not <b>"operating"</b> or the CPC <b>iml-mode</b> property is not <b>"lpar"</b>.</li> </ul>
500 (Server Error)	275	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Swap Current Time Server

The **Swap Current Time Server** operation changes the role of the CPC object designated by *{cpc-id}* to the Current Time Server (CTS).

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/swap-cts

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
stp-id	String (1-8)	Required	Identifies the STP. Can contain 0-9, a-z, A-Z, underline (_) and dash (-).

### Description

This operation changes the role of the CPC object designated by *{cpc-id}* to the Current Time Server (CTS).

On success, HTTP status code 204 (No Content) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission to the **System (Sysplex) Time** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	285	The CPC targeted by this operation is already the Preferred Time Server.
	286	The operation was rejected for one of the following reasons: <ul style="list-style-type: none"> <li>• The stp-id field value does not match the current CTN identifier</li> <li>• The operation is not permitted for a mixed CTN.</li> </ul>
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The requested operation cannot be performed, due to the state of the object: <ul style="list-style-type: none"> <li>• Server Time Protocol is not enabled on this CPC</li> <li>• an ETR reverse migration is in progress</li> <li>• no alternate is active</li> <li>• this CPC is not the backup time server</li> </ul>



HTTP error status code	Reason code	Description
500 (Server Error)	272	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Set STP Configuration

The **Set STP Configuration** operation updates the configuration for an STP-only Coordinated Timing Network.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/set-stp-config

In this request, the URI variable {cpc-id} is the object ID of the target CPC object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
stp-id	String (1-8)	Required	The current STP identifier for the CTN, used to verify that the CPC is a member of the correct CTN. Can contain 0-9, a-z, A-Z, underline (_) and dash (-).
new-stp-id	String (1-8)	Optional	If provided, the new STP identifier for the CTN, Can contain 0-9, a-z, A-Z, underline (_) and dash (-).
force	Boolean	Required	Whether a disruptive operation is allowed (true) or rejected (false)
preferred-time-server	stp-node object	Required	Identifies the CPC object to be the Preferred Time Server. Refer to Table 170 on page 616 for details.
backup-time-server	stp-node object	Optional	Identifies the CPC object to be the Backup Time Server. If not provided, the STP has no Backup Time Server. Refer to Table 170 on page 616 for details.
arbiter	stp-node object	Optional	Identifies the CPC object to be the Arbiter for the CTN. If not provided, the STP has no Arbiter. Refer to Table 170 on page 616 for details.
current-time-server	String Enum	Required	Identifies the role of the Current Time Server (CTS). One of: <ul style="list-style-type: none"> <li>• <b>"preferred"</b> - the Preferred Time Server is the CTS</li> <li>• <b>"backup"</b> - the Backup Time Server is the CTS.</li> </ul>

### Description

The CPC object designated by {cpc-id} must be the system that becomes the Current Time Server (CTS).

On success, HTTP status code 204 (No Content) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by {cpc-id}

- Action/task permission to the **System (Sysplex) Time** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	282	The operation was rejected, due to an incomplete preferred, backup or arbiter nested object specification. Refer to Table 170 on page 616 for details.
	284	This operation does not target the Current Time Server CPC.
	285	The operation was rejected, due to one of the following configuration errors: <ul style="list-style-type: none"> <li>• A backup-time-server object is required when providing an arbiter object</li> <li>• A backup-time-server object is required when current-time-server is backup</li> <li>• The preferred-time-server, backup-time-server and arbiter objects do not reference different CPCs</li> </ul>
	286	The operation was rejected for one of the following reasons: <ul style="list-style-type: none"> <li>• The stp-id field value does not match the current CTN identifier</li> <li>• The operation is not permitted for a mixed CTN</li> </ul>
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	2	An object URI in one of the stp-node objects in the request body does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The requested operation cannot be performed, due to the state of the object: <ul style="list-style-type: none"> <li>• Server Time Protocol is not enabled on this CPC</li> <li>• an ETR reverse migration is in progress</li> <li>• no alternate is active</li> <li>• the operation is not permitted for a mixed-CTN.</li> </ul>
	287	The provided configuration can only be set by specifying force=true.
	288	No communication path between preferred-time-server and backup-time-server.
	289	No communication path to the arbiter.
500 (Server Error)	272	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Change STP-only Coordinated Timing Network

The **Change STP-only Coordinated Timing Network** operation, sent to the CPC object designated by *{cpc-id}* with the role of Current Time Server (CTS) in an STP-only Coordinated Timing Network (CTN), changes the STP ID portion of the CTN ID for the entire STP-only CTN.

## HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/change-stponly-ctn

In this request, the URI variable {cpc-id} is the object ID of the target CPC object.

## Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
stp-id	String (1-8)	Required	The new STP identifier. Can contain 0-9, a-z, A-Z, underline (_) and dash (-).

## Description

This operation, sent to the CPC object designated by {cpc-id} with the role of Current Time Server (CTS) in an STP-only Coordinated Timing Network (CTN), changes the STP ID portion of the CTN ID for the entire STP-only CTN.

On success, HTTP status code 204 (No Content) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by {cpc-id}
- Action/task permission to the **System (Sysplex) Time** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	284	The CPC targeted by the operation is not the Current Time Server. Retry the operation using the object-uri for the Current Time Server CPC.
	286	The operation is not permitted for a mixed CTN.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ({cpc-id}) does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The requested operation cannot be performed, due to the state of the object: <ul style="list-style-type: none"> <li>• Server Time Protocol is not enabled on this CPC</li> <li>• an ETR reverse migration is in progress</li> </ul>
500 (Server Error)	272	An unexpected error occurred during processing of the Server Time Protocol operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Join STP-only Coordinated Timing Network

The **Join STP-only Coordinated Timing Network** operation allows a CPC object designated by *{cpc-id}* to join an STP-only Coordinated Timing Network (CTN).

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/join-stponly-ctn

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
stp-id	String (1-8)	Required	Identifies the STP to be joined. Can contain 0-9, a-z, A-Z, underline (_) and dash (-).

### Description

If the CPC object is already participating in a different STP-only CTN and is the Current Time Server (CTS), the operation is rejected. Otherwise, the CPC object is removed from its current CTN and joins the specified CTN.

If the CPC object has an ETR ID, the ETR ID is removed.

On success, HTTP status code 204 (No Content) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission to the **System (Sysplex) Time** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The requested operation cannot be performed, due to the state of the object: <ul style="list-style-type: none"> <li>• Server Time Protocol is not enabled on this CPC</li> </ul>

HTTP error status code	Reason code	Description
500 (Server Error)	272	An unexpected error occurred during processing of the Server Time Protocol operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Leave STP-only Coordinated Timing Network

The **Leave STP-only Coordinated Timing Network** operation allows a CPC object designated by *{cpc-id}* to leave the STP-only Coordinated Timing Network (CTN) in which it currently participates.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/operations/leave-stponly-ctn

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Description

The CPC object cannot be the Current Time Server (CTS) in the CTN in which it is currently participating.

On success, HTTP status code 204 (No Content) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission to the **System (Sysplex) Time** task.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	286	The operation is not permitted for a mixed CTN.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
409 (Conflict)	1	The requested operation cannot be performed, due to the state of the object: <ul style="list-style-type: none"> <li>• Server Time Protocol is not enabled on this CPC</li> <li>• this CPC is not a member of a CTN</li> <li>• this CPC is the Current Time Server.</li> </ul>

HTTP error status code	Reason code	Description
500 (Server Error)	272	An unexpected error occurred during processing of the Server Time Protocol operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Get CPC Audit Log

The **Get CPC Audit Log** operation returns the CPC's audit log, filtered according to the query parameters, if specified.

### HTTP method and URI

**GET** `/api/cpcs/{cpc-id}/operations/get-audit-log`

In this request, the URI variable `{cpc-id}` is the object ID of the target CPC object.

### Query parameters:

Name	Type	Rqd/Opt	Description
begin-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created earlier than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.
end-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created later than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.

## Response body contents

On successful completion, the response body is a JSON array of JSON objects returned using HTTP chunked transfer encoding. Each array element is a log-entry-info object containing information about a single log entry. The array elements are in order of increasing timestamp. See Table 186 on page 631 for more information.

## Description

This operation returns the CPC's audit log in increasing timestamp order, filtered according to the query parameters, if specified. Each log entry pertains to a specific event that occurred on or to a managed object or the CPC itself. If the begin-time query parameter is specified, then any entries earlier than that time are omitted. If the end-time query parameter is specified, then any entries later than that time are omitted.

The URI path must designate an existing CPC object, the API user must have object-access permission to it, and that CPC must be at a release level that supports this operation. If any of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action/task permission to the **Audit and Log Management** task; otherwise, status code 403 (Forbidden) is returned.

On successful execution, the response body contains an array of filtered log entries. If the audit log is empty or there are no entries to be returned after filtering, then an empty array is provided. Each log entry contains the event ID, event name and event message. If there are data items included in the event message, they are available separately. The order and meaning of the substitution items for each event ID are documented in the IBM Knowledge Center, at <http://www.ibm.com/support/knowledgecenter/> (Select **z Systems** on the navigation bar, and then select your server). The information can be found in the HMC Introduction topic **Audit, Event, and Security Log Messages**.

### Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the CPC object specified in the request URI
- Action/task permission to the **Audit and Log Management** task.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 782.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 225. Get CPC Audit Log: HTTP status and reason codes.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.
	4	The CPC designated by the request URI does not support this operation.
500 (Server Error)	307	The request timed out while attempting to communicate with the SE or while attempting to get the log data.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/cpcs/1946e00f-401b-3aa6-84a3-5e49614743ec/operations/get-audit-log
  HTTP/1.1
x-api-session: lui4gmb59aunfkk8of69nrpks5mntq5xjtc613rdxjq53iv0e1j
```

---

Figure 405. Get CPC Audit Log: Request

---

```

200 OK
server: zSeries management console API web server / 2.0
transfer-encoding: chunked
cache-control: no-cache
date: Fri, 03 Oct 2014 00:10:35 GMT
content-type: application/json;charset=ISO-8859-1
[
  {
    "event-data-items": [],
    "event-details": [
      {
        "event-details-data-items": [],
        "event-details-message": "S32 - Disabled"
      }
    ],
    "event-id": "1809",
    "event-message": "Power Cap settings have changed.",
    "event-name": "POWERCAP",
    "event-time": 1412202863030,
    "user-uri": null,
    "userid": null
  },
  {
    "event-data-items": [
      {
        "data-item-number": 0,
        "data-item-type": "string",
        "data-item-value": "HMCCHGM(1.2.3.5)"
      },
      {
        "data-item-number": 1,
        "data-item-type": "string",
        "data-item-value": "S32"
      }
    ],
    "event-details": [],
    "event-id": "734",
    "event-message": "Remote support call generated on S32 is being handled by
      call-home server HMCCHGM(1.2.3.5).",
    "event-name": "TRSF_OFFER",
    "event-time": 1412262459470,
    "user-uri": null,
    "userid": null
  }
]

```

---

Figure 406. Get CPC Audit Log: Response

## Get CPC Security Log

The **Get CPC Security Log** operation returns the CPC's security log, filtered according to the query parameters, if specified.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/operations/get-security-log

In this request, the URI variable {cpc-id} is the object ID of the target CPC object.

### Query parameters:



Name	Type	Rqd/Opt	Description
begin-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created earlier than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.
end-time	Timestamp	Optional	A timestamp used to filter log entries. Entries created later than this time are omitted from the results. This is specified as the number of milliseconds since the epoch and must be greater than or equal to 0. If not specified, then no such filtering is performed.

## Response body contents

On successful completion, the response body is a JSON array of JSON objects returned using HTTP chunked transfer encoding. Each array element is a log-entry-info object containing information about a single log entry. The array elements are in order of increasing timestamp. See Table 186 on page 631 for more information.

## Description

This operation returns the CPC's security log in increasing timestamp order, filtered according to the query parameters, if specified. Each log entry pertains to a specific event that occurred on or to a managed object or the CPC itself. If the begin-time query parameter is specified, then any entries earlier than that time are omitted. If the end-time query parameter is specified, then any entries later than that time are omitted.

The URI path must designate an existing CPC object, the API user must have object-access permission to it, and that CPC must be at a release level that supports this operation. If either of these conditions is not met, status code 404 (Not Found) is returned. In addition, the API user must have action/task permission to the **View Security Logs** task; otherwise, status code 403 (Forbidden) is returned.

On successful execution, the response body contains an array of filtered log entries. If the security log is empty or there are no entries to be returned after filtering, then an empty array is provided. Each log entry contains the event ID, event name and event message. If there are data items included in the event message, they are available separately. The order and meaning of the substitution items for each event ID are documented in the IBM Knowledge Center, at <http://www.ibm.com/support/knowledgecenter/> (Select **z Systems** on the navigation bar, and then select your server). The information can be found in the HMC Introduction topic **Audit, Event, and Security Log Messages**.

## Authorization requirements

This operation has the following authorization requirement:

- Object-access permission to the CPC object specified in the request URI
- Action/task permission to the **View Security Logs** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

| Table 226. Get CPC Security Log: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The request URI does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.
	4	The CPC designated by the request URI does not support this operation.
500 (Server Error)	307	The request timed out while attempting to communicate with the SE or while attempting to get the log data.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not currently communicating with an SE needed to perform the requested operation.

| Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### | Example HTTP interaction

---

```
GET /api/cpcs/1946e00f-401b-3aa6-84a3-5e49614743ec/operations/get-security-log
  HTTP/1.1
x-api-session: 1ui4gmb59aunfkk8of69nrpks5mtmq5xjc613rdxjq53iv0e1j
```

---

| Figure 407. Get CPC Security Log: Request

```

200 OK
server: zSeries management console API web server / 2.0
transfer-encoding: chunked
cache-control: no-cache
date: Fri, 03 Oct 2014 00:10:35 GMT
content-type: application/json;charset=ISO-8859-1
[
  {
    "event-data-items": [],
    "event-details": [],
    "event-id": "778",
    "event-message": "Mirroring data from the primary Support Element to the
      alternate Support Element started.",
    "event-name": "ASEMIRRST",
    "event-time": 1412258403480,
    "user-uri": null,
    "userid": null
  },
  {
    "event-data-items": [
      {
        "data-item-number": 0,
        "data-item-type": "string",
        "data-item-value": "Communications to the Alternate SE was not
          active."
      }
    ],
    "event-details": [],
    "event-id": "779",
    "event-message": "Mirroring data from the primary Support Element to the
      alternate Support Element failed. Communications to the Alternate SE was not
      active.",
    "event-name": "ASEMIRRNO",
    "event-time": 1412258414110,
    "user-uri": null,
    "userid": null
  }
]

```

Figure 408. Get CPC Security Log: Response

## List CPC Hardware Messages

The **List CPC Hardware Messages** operation lists the current set of hardware messages associated with the CPC.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/hardware-messages

In this request, the URI variable {cpc-id} is the object ID of a CPC object for which hardware messages are to be listed.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
hardware-messages	Array of hardware-message-info objects	Array of nested hardware-message-info objects as defined in the next table.

Each nested hardware-message-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	The canonical URI path of the hardware message. The URI is in the following form: <code>/api/cpcs/{cpc-id}hardware-messages/{hardware-message-id}</code>
timestamp	Timestamp	The date and time the hardware message was created
text	String	The text of the hardware message.

## Description

This operation returns a set of CPC hardware messages in increasing timestamp order, filtered according to the query parameters, if specified. Each hardware message describes an event or notification that may require the operator's attention. The list of hardware messages can be limited by specifying explicit filtering criteria on the request.

If the **begin-time** query parameter is specified, then any entries earlier than that time are omitted. If the **end-time** query parameter is specified, then any entries later than that time are omitted.

If there are no hardware messages associated with the CPC, or if no hardware messages are to be included in the results due to filtering, an empty array is returned and the operation completes successfully.

The URI path must designate an existing CPC and the API user must have object access permission to it; otherwise, status code 404 (Not Found) is returned. In addition, the API user must have Action/Task permission to the **Hardware Messages** task; otherwise, status code 403 (Forbidden) is returned.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by `{cpc-id}`.
- Action/Task permission to the **Hardware Messages** task.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 787.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

Table 227. List CPC Hardware Messages: HTTP status and reason codes

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	7	The <b>begin-time</b> value is greater than the <b>end-time</b> value.
403 (Forbidden)	1	The API user does not have Action/Task permission for the <b>Hardware Messages</b> task.

Table 227. List CPC Hardware Messages: HTTP status and reason codes (continued)

HTTP error status code	Reason code	Description
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	4	The CPC designated by the request URI does not support this operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

```
GET /api/cpcs/1946e00f-401b-3aa6-84a3-5e49614743ec/hardware-messages HTTP/1.1
x-api-session: 35rr3x40qbnou8zwmx8ad801dp8koes4f2abcl6m1fg5jm4tug
```

Figure 409. List CPC Hardware Messages: Request

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Mon, 06 Oct 2014 17:08:58 GMT
content-type: application/json;charset=UTF-8
content-length: 242
{
  "hardware-messages": [
    {
      "element-uri": "/api/cpcs/1946e00f-401b-3aa6-84a3-5e49614743ec/hardware-messages/
        43a2922a-4d6b-11e4-972f-42f2e9ccd169",
      "text": "Licensed internal code has detected a problem. [Problem # 49]",
      "timestamp": 1412608381950
    }
  ]
}
```

Figure 410. List CPC Hardware Messages: Response

### Get CPC Hardware Message Properties

The **Get CPC Hardware Message Properties** operation retrieves the properties of a single CPC hardware message.

#### HTTP method and URI

```
GET /api/cpcs/{cpc-id}/hardware-messages/{hardware-message-id}
```

#### URI Variables:

Variable	Description
<i>{cpc-id}</i>	Object ID of the CPC object.
<i>{hardware-message-id}</i>	Element ID of the hardware message to retrieve.

## | Response body contents

| On successful completion, the response body contains a JSON object that provides the current values of the properties for the CPC hardware message object as defined in “Data model” on page 743. Field names and data types in the JSON object are the same as the property names and data types defined in the data model.

## | Description

| This operation retrieves the properties of a single CPC hardware message specified by *{hardware-message-id}*.

| The URI path must designate an existing CPC, and the API user must have object access permission to it; otherwise, status code 404 (Not Found) is returned.

| The URI path must designate an existing hardware message; otherwise status code 404 (Not Found) is returned. In addition, the API user must have Action/Task permission to the **Hardware Messages** task; otherwise, status code 403 (Forbidden) is returned.

## | Authorization requirements

| This operation has the following authorization requirements:

- | • Object access permission to the CPC object designated by *{cpc-id}*.
- | • Action/Task permission to the **Hardware Messages** task.

## | HTTP status and reason codes

| On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

| The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

| *Table 228. Get CPC Hardware Message Properties: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have Action/Task permission for the <b>Hardware Messages</b> task.
404 (Not Found)	1	The object ID in the URI <i>{cpc-id}</i> does not designate an existing CPC object, or the API user does not have object access permission to the object.
	4	The CPC designated by the request URI does not support this operation.
	322	The element ID in the URI <i>{hardware-message-id}</i> does not designate an existing CPC hardware message.

| Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
GET /api/cpcs/1946e00f-401b-3aa6-84a3-5e49614743ec/hardware-messages/
  43a2922a-4d6b-11e4-972f-42f2e9ccd169 HTTP/1.1
x-api-session: 5w3ci8h0ptx0l2jy3gdeiuzolzkxsdpfzs77dic5wi0ejzikh3
```

Figure 411. Get CPC Hardware Message Properties: Request

```
200 OK
server: zSeries management console API web server / 2.0
cache-control: no-cache
date: Mon, 06 Oct 2014 17:10:17 GMT
content-type: application/json;charset=UTF-8
content-length: 355
{
  "class": "hardware-message",
  "element-id": "43a2922a-4d6b-11e4-972f-42f2e9ccd169",
  "element-uri": "/api/cpcs/1946e00f-401b-3aa6-84a3-5e49614743ec/hardware-messages/
    43a2922a-4d6b-11e4-972f-42f2e9ccd169",
  "parent": "/api/cpcs/1946e00f-401b-3aa6-84a3-5e49614743ec",
  "text": "Licensed internal code has detected a problem. [Problem # 49]",
  "timestamp": 1412608381950
}
```

Figure 412. Get CPC Hardware Message Properties: Response

## Delete CPC Hardware Message

The **Delete CPC Hardware Message** operation deletes a single CPC hardware message

### HTTP method and URI

```
DELETE /api/cpcs/{cpc-id}/hardware-messages/{hardware-message-id}
```

### URI Variables:

Variable	Description
{cpc-id}	Object ID of the CPC object.
{hardware-message-id}	Element ID of the hardware message to delete.

### Description

This operation deletes a specific CPC hardware message. The hardware message to be deleted is identified by the {hardware-message-id} variable in the URI.

The URI path must designate an existing CPC and the API user must have object access permission to it; otherwise status code 404 (Not Found) is returned.

The URI path must designate an existing hardware message; otherwise, status code 404 (Not Found) is returned. In addition, the API user must have Action/Task permission to the **Hardware Messages** task; otherwise, status code 403 (Forbidden) is returned.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by {cpc-id}.

- Action/Task permission to the **Hardware Messages** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned with no response body provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

*Table 229. Delete CPC Hardware Message: HTTP status and reason codes*

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have Action/Task permission for the <b>Hardware Messages</b> task.
404 (Not Found)	1	The object ID in the URI <i>{cpc-id}</i> does not designate an existing CPC object, or the API user does not have Object-access permission to the object.
	4	The CPC designated by the request URI does not support this operation.
	322	The element ID in the URI <i>{hardware-message-id}</i> does not designate an existing CPC hardware message.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

```
DELETE /api/cpcs/0fdde999-5957-3129-99aa-b6f4bfbbc071/hardware-messages/
d5591a80-43f8-11e4-ac52-42f2e910664b HTTP/1.1
x-api-session: c8un3odpy8yyp150o3poz1ud4gwyfod1wyq495327bpyn2p0z
```

*Figure 413. Delete CPC Hardware Message: Request*

```
204 No Content
date: Mon, 09 Feb 2015 20:07:31 GMT
server: zSeries management console API web server / 2.0

<No response body>
```

*Figure 414. Delete CPC Hardware Message: Response*

## Inventory service data

Information about CPCs can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for the CPC objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "cpc" are to be included. An entry for a particular CPC is included only if the API user has access permission to that object as described in the **Get CPC Properties** operation.



- For each CPC object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for “Get CPC Properties” on page 758. That is, the data provided is the same as would be provided if a **Get CPC Properties** operation were requested targeting this object.

## Logical Partition object

The Processor Resource/Systems Manager™ (PR/SM) is a feature of IBM mainframes that enables logical partitioning of the CEC. A logical partition (LPAR) is a virtual machine at the hardware level. Each LPAR operates as an independent server running its own operating environment. Each LPAR runs its own operating system, which can be any mainframe operating system.

## Data model

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

This object includes the properties defined in “Base managed object properties schema” on page 39, with the following class-specific specialization:

Table 230. Logical Partition object: base managed object properties specializations

Name	Qualifier	Type	Description of specialization
<b>object-uri</b>	—	String/ URI	The canonical URI path of the Logical Partition object, of the form <code>/api/logical-partitions/{<i>logical-partition-id</i>}</code> where <code>{<i>logical-partition-id</i>}</code> is the value of the <b>object-id</b> property of the Logical Partition object.
<b>parent</b>	—	String/ URI	The canonical URI path of the associated CPC object.
<b>class</b>	—	String	The class of a Logical Partition object is <b>"logical-partition"</b> .
<b>name</b>	(ro)	String (1-8)	The name of the logical partition
<b>description</b>	(ro)	String (0-1024)	The descriptive text associated with this object.
<b>status</b>	(sc)	String Enum	One of the following values: <ul style="list-style-type: none"> <li><b>"operating"</b> - the logical partition has a active control program</li> <li><b>"not-operating"</b> - the logical partition's CPC is non operational</li> <li><b>"not-activated"</b> - the logical partition does not have an active control program</li> <li><b>"exceptions"</b> - the logical partition's CPC has one or more unusual conditions</li> </ul>
<b>acceptable-status</b>	(w)(pc)	Array of String Enum	An array of one or more status strings that determine an acceptable status for a logical partition. When a logical partition's <b>status</b> property contains one of the specified acceptable-status values, the <b>has-unacceptable-status</b> property contains false.

## Class specific additional properties

In addition to the properties defined via included schemas, this object includes the following additional class-specific properties. Refer to the *Processor Resource/Systems Manager Planning Guide* for more detailed explanations of the various properties.

There are additional notes throughout the table. Please refer to the note list at the end of the table.

Table 231. Logical Partition object: class specific additional properties

Name	Qualifier	Type	Description
<b>os-name</b> <sup>1</sup>	(pc)	String (0-8)	An operating system provided value, used to identify the operating system instance. The format of the value is operating system dependant. If not provided by the operating system, an empty string is returned.
<b>os-type</b> <sup>1</sup>	(pc)	String (0-8)	A human readable form of the operating system provided value for the type of the operating system active in this logical partition. If not provided, an empty string is returned.
<b>os-level</b> <sup>1</sup>	(pc)	String (0-32)	A human readable form of the operating system provided value for the level of the operating system active in this logical partition. If not provided, an empty string is returned.
<b>sysplex-name</b> <sup>1</sup>	(pc)	String (1-8)	Applicable only for z/OS, the name of the sysplex to which this logical partition is a member. Otherwise a null object is returned.
<b>has-operating-system-messages</b> <sup>1</sup>	—	Boolean	If true, object has operating system messages. If false, object does not have operating system messages.
<b>activation-mode</b>	—	String Enum	One of the following values: <ul style="list-style-type: none"> <li>• <b>"esa390"</b> - the logical partition is in ESA/390 mode</li> <li>• <b>"esa390tpf"</b> - the logical partition is in ESA/390 TPF mode</li> <li>• <b>"coupling-facility"</b> - the logical partition is running as a coupling facility</li> <li>• <b>"linux"</b> - the logical partition is in Linux mode</li> <li>• <b>"zvm"</b> - the logical partition is in z/VM mode</li> <li>• <b>"zaware"</b> - the logical partition is in IBM zAware mode.</li> <li>• <b>"not-set"</b> - the logical partition is not activated.</li> </ul>
<b>next-activation-profile-name</b>	(w)(pc)	String (1-16)	Image activation profile name or load activation profile name to be used on the next activate.
<b>last-used-activation-profile</b>	(pc)	String (0-16)	The last used activation profile name or a null string.
<b>initial-processing-weight</b> <sup>1, 2, 3</sup>	(w)	Integer	The relative amount of shared general purpose processor resources allocated to the logical partition:  Get:  <b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared general purpose processor.  <b>1-999</b> Represents the relative amount of shared general purpose processor resources allocated to the logical partition.  Update:  <b>1-999</b> Defines the relative amount of shared general purpose processor resources allocated to the logical partition.

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>initial-processing-weight-capped</b> <sup>1, 2, 3, 4</sup>	(w)	Boolean	<p>Whether the initial processing weight for general purpose processors is a limit or a target.</p> <p><b>True</b> Indicates that the initial general purpose processor processing weight for the logical partition is capped. It represents the logical partition's maximum share of general purpose processor resources.</p> <p><b>False</b> Indicates that the initial general purpose processor processing weight for the logical partition is not capped. It represents the share of general purpose processor resources guaranteed to a logical partition when all general purpose processor resources are in use. Otherwise, when excess general purpose processor resources are available, the logical partition can use them if necessary.</p>
<b>minimum-processing-weight</b> <sup>1, 2, 3</sup>	(w)	Integer	<p>The minimum relative amount of shared general purpose processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared general purpose processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared general purpose processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>0</b> There is no minimum value for the processing weight.</p> <p><b>1-999</b> Define the minimum relative amount of shared general purpose processor resources allocated to the logical partition. The value must be less than or equal to the <b>initial-processing-weight</b> property.</p>
<b>maximum-processing-weight</b> <sup>1, 2, 3</sup>	(w)	Integer	<p>The maximum relative amount of shared general purpose processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared general purpose processor.</p> <p><b>1-999</b> Represents the maximum relative amount of shared general purpose processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Defines the maximum relative amount of shared general purpose processor resources allocated to the logical partition. The value must be greater than or equal to the <b>initial-processing-weight</b> property.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>current-processing-weight</b> <sup>1,3</sup>	—	Integer	The relative amount of shared general purpose processor resources currently allocated to the logical partition.  <b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared general purpose processor.  <b>1-999</b> Represents the relative amount of shared general purpose processor resources currently allocated to the logical partition.
<b>current-processing-weight-capped</b> <sup>1,2,3</sup>	—	Boolean	Whether the current general purpose processing weight is a limit or a target.  <b>True</b> Indicates that the current general purpose processing weight for the logical partition is capped. It represents the logical partition's maximum share of resources, regardless of the availability of excess processor resources.  <b>False</b> Indicates that the current general purpose processing weight for the logical partition is not capped. It represents the share of resources guaranteed to a logical partition when all processor resources are in use. Otherwise, when excess processor resources are available, the logical partition can use them if necessary.
<b>workload-manager-enabled</b> <sup>1,5</sup>	(w)	Boolean	Whether or not z/OS Workload Manager is allowed to change processing weight related properties.  <b>True</b> Indicates that z/OS Workload Manager is allowed to change processing weight related properties for this logical partition.  <b>False</b> Indicates that z/OS Workload Manager is not allowed to change processing weight related properties for this logical partition.
<b>absolute-processing-capping</b>	(w)	absolute-capping object	The amount of absolute capping applied to the general purpose processor. <b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.
<b>defined-capacity</b> <sup>1</sup>	(w)	Integer	The defined capacity expressed in terms of Millions of Service Units (MSU)s per hour. MSU is a measure of processor resource consumption. The amount of MSUs a logical partition consumes is dependent on the model, the number of logical processors available to the partition, and the amount of time the logical partition is dispatched. The defined capacity value specifies how much capacity the logical partition is to be managed to by z/OS Workload Manager for the purpose of software pricing.  <b>0</b> No defined capacity is specified for this logical partition.  <b>1-nnnn</b> Represents the amount of defined capacity specified for this logical partition.

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>cluster-name</b> <sup>1</sup>	(pc)	String (0-8)	LPAR cluster name, which identifies membership in a group of logical partitions that are members of the same z/OS Parallel Sysplex®.
<b>partition-number</b> <sup>1</sup>	(pc)	String (2)	The partition number for the logical partition, in hexadecimal.
<b>partition-identifier</b> <sup>1</sup>	(ro)	String (2)	The partition identifier for the logical partition, in hexadecimal.
<b>initial-aap-processing-weight</b> <sup>1, 3, 6</sup>	(w)	Integer	<p>The relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Application Assist Processor (zAAP) processor.</p> <p><b>1-999</b> Represents the relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p>
<b>initial-aap-processing-weight-capped</b> <sup>1, 3, 4, 6, 7</sup>	(w)	Boolean	<p>Whether the initial processing weight for Application Assist Processor (zAAP) processors is a limit or a target.</p> <p><b>True</b> Indicates that the initial Application Assist Processor (zAAP) processor processing weight for the logical partition is capped. It represents the logical partition's maximum share of Application Assist Processor (zAAP) processor resources, regardless of the availability of excess Application Assist Processor (zAAP) processor resources.</p> <p><b>False</b> Indicates that the initial Application Assist Processor (zAAP) processor processing weight for the logical partition is not capped. It represents the share of Application Assist Processor (zAAP) processor resources guaranteed to a logical partition when all Application Assist Processor (zAAP) processor resources are in use. Otherwise, when excess Application Assist Processor (zAAP) processor resources are available, the logical partition can use them if necessary.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>minimum-aap-processing-weight</b> <sup>1, 3, 6</sup>	(w)	Integer	<p>The minimum relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Application Assist Processor (zAAP) processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>0</b> No minimum value for the processing weight.</p> <p><b>1-999</b> Define the minimum relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p>
<b>maximum-aap-processing-weight</b> <sup>1, 3, 6</sup>	(w)	Integer	<p>The maximum relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Application Assist Processor (zAAP) processor.</p> <p><b>1-999</b> Represents the maximum relative amount of shared Application Assist Processor (zAAP) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the maximum relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p>
<b>current-aap-processing-weight</b> <sup>1, 3</sup>	—	Integer	<p>The current relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Application Assist Processor (zAAP) processor.</p> <p><b>1-999</b> Represents the current relative amount of shared Application Assist Processor (zAAP) processor resources initially allocated to the logical partition.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>current-aap-processing-weight-capped</b> <sup>1, 3, 6, 7</sup>	—	Boolean	<p>Whether the current Application Assist Processor (zAAP) processing weight is a limit or a target.</p> <p><b>True</b> Indicates that the current Application Assist Processor (zAAP) processing weight for the logical partition is capped. It represents the logical partition's maximum share of resources, regardless of the availability of excess processor resources.</p> <p><b>False</b> Indicates that the current Application Assist Processor (zAAP) processing weight for the logical partition is not capped. It represents the share of resources guaranteed to a logical partition when all processor resources are in use. Otherwise, when excess processor resources are available, the logical partition can use them if necessary.</p>
<b>absolute-aap-capping</b>	(w)	absolute-capping object	<p>The amount of absolute capping applied to the Application Assist Processor (zAAP).</p> <p><b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.</p>
<b>initial-ifl-processing-weight</b> <sup>1, 3, 8</sup>	(w)	Integer	<p>The relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Integrated Facility for Linux (IFL) processor.</p> <p><b>1-999</b> Represents the relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>initial-ifl-processing-weight-capped</b> <sup>1, 3, 4, 8, 9</sup>	(w)	Boolean	<p>Whether the initial processing weight for Integrated Facility for Linux (IFL) processors is a limit or a target.</p> <p><b>True</b> Indicates that the initial Integrated Facility for Linux (IFL) processor processing weight for the logical partition is capped. It represents the logical partition's maximum share of Integrated Facility for Linux (IFL) processor resources, regardless of the availability of excess Integrated Facility for Linux (IFL) processor resources</p> <p><b>False</b> Indicates that the initial Integrated Facility for Linux (IFL) processor processing weight for the logical partition is not capped. It represents the share of Integrated Facility for Linux (IFL) processor resources guaranteed to a logical partition when all Integrated Facility for Linux (IFL) processor resources are in use. Otherwise, when excess Integrated Facility for Linux (IFL) processor resources are available, the logical partition can use them if necessary.</p>
<b>minimum-ifl-processing-weight</b> <sup>1, 3, 8</sup>	(w)	Integer	<p>The minimum relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Integrated Facility for Linux (IFL) processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>0</b> There is no minimum value for the processing weight.</p> <p><b>1-999</b> Define the minimum relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p>



Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>maximum-ift-processing-weight</b> <sup>1, 3, 8</sup>	(w)	Integer	<p>The maximum relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Integrated Facility for Linux (IFL) processor.</p> <p><b>1-999</b> Represents the maximum relative amount of shared Integrated Facility for Linux (IFL) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the maximum relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p>
<b>current-ift-processing-weight</b> <sup>1, 3</sup>	—	Integer	<p>The current relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Integrated Facility for Linux (IFL) processor.</p> <p><b>1-999</b> Represents the current relative amount of shared Integrated Facility for Linux (IFL) processor resources initially allocated to the logical partition.</p>
<b>current-ift-processing-weight-capped</b> <sup>1, 3, 8, 9</sup>	—	Boolean	<p>Whether the current Integrated Facility for Linux (IFL) processing weight is a limit or a target.</p> <p><b>True</b> Indicates that the current Integrated Facility for Linux (IFL) processing weight for the logical partition is capped. It represents the logical partition's maximum share of resources, regardless of the availability of excess processor resources.</p> <p><b>False</b> Indicates that the current Integrated Facility for Linux (IFL) processing weight for the logical partition is not capped. It represents the share of resources guaranteed to a logical partition when all processor resources are in use. Otherwise, when excess processor resources are available, the logical partition can use them if necessary.</p>
<b>absolute-ift-capping</b>	(w)	absolute-capping object	<p>The amount of absolute capping applied to the Integrated Facility for Linux (IFL) processor.</p> <p><b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>initial-ziip-processing-weight</b> <sup>1, 3, 10</sup>	(w)	Integer	<p>The relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Integrated Information Processors (zIIP) processor.</p> <p><b>1-999</b> Represents the relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p>
<b>initial-ziip-processing-weight-capped</b> <sup>1, 3, 4, 10, 11</sup>	(w)	Boolean	<p>Whether the initial processing weight for Integrated Information Processors (zIIP) processors is a limit or a target.</p> <p><b>True</b> Indicates that the initial Integrated Information Processors (zIIP) processor processing weight for the logical partition is capped. It represents the logical partition's maximum share of Integrated Information Processors (zIIP) processor resources, regardless of the availability of excess Integrated Information Processors (zIIP) processor resources.</p> <p><b>False</b> Indicates that the initial Integrated Information Processors (zIIP) processor processing weight for the logical partition is not capped. It represents the share of Integrated Information Processors (zIIP) processor resources guaranteed to a logical partition when all Integrated Information Processors (zIIP) processor resources are in use. Otherwise, when excess Integrated Information Processors (zIIP) processor resources are available, the logical partition can use them if necessary.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>minimum-ziip-processing-weight</b> <sup>1, 3, 10</sup>	(w)	Integer	<p>The minimum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Integrated Information Processors (zIIP) processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>0</b> There is no minimum value for the processing weight.</p> <p><b>1-999</b> Define the minimum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p>
<b>maximum-ziip-processing-weight</b> <sup>1, 3, 10</sup>	(w)	Integer	<p>The maximum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Integrated Information Processors (zIIP) processor.</p> <p><b>1-999</b> Represents the maximum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the maximum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p>
<b>current-ziip-processing-weight</b> <sup>1, 3</sup>	—	Integer	<p>The current relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Integrated Information Processors (zIIP) processor.</p> <p><b>1-999</b> Represents the current relative amount of shared Integrated Information Processors (zIIP) processor resources initially allocated to the logical partition.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>current-ziip-processing-weight-capped</b> <sup>1, 3, 10, 11</sup>	—	Boolean	<p>Whether the current Integrated Information Processors (zIIP) processing weight is a limit or a target.</p> <p><b>True</b> Indicates that the current Integrated Information Processors (zIIP) processing weight for the logical partition is capped. It represents the logical partition's maximum share of resources, regardless of the availability of excess processor resources.</p> <p><b>False</b> Indicates that the current Integrated Information Processors (zIIP) processing weight for the logical partition is not capped. It represents the share of resources guaranteed to a logical partition when all processor resources are in use. Otherwise, when excess processor resources are available, the logical partition can use them if necessary.</p>
<b>absolute-ziip-capping</b>	(w)	absolute-capping object	<p>The amount of absolute capping applied to the Integrated Information Processors (zIIP) processor.</p> <p><b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.</p>
<b>initial-cf-processing-weight</b> <sup>1, 3, 12</sup>	(w)	Integer	<p>The relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Internal Coupling Facility (ICF) processor.</p> <p><b>1-999</b> Represents the relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p> <p>Update:</p> <p><b>1-999</b> The relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>initial-cf-processing-weight-capped</b> <sup>1, 3, 4, 12, 13</sup>	(w)	Boolean	<p>Indicates whether the initial processing weight for Internal Coupling Facility (ICF) processors is a limit or a target.</p> <p><b>True</b> Indicates that the initial Internal Coupling Facility (ICF) processor processing weight for the Logical Partition object is capped. It represents the logical partition's maximum share of Internal Coupling Facility (ICF) processor resources, regardless of the availability of excess Internal Coupling Facility (ICF) processor resources.</p> <p><b>False</b> Indicates that the initial Internal Coupling Facility (ICF) processor processing weight for the Logical Partition is not capped. It represents the share of Internal Coupling Facility (ICF) processor resources guaranteed to a logical partition when all Internal Coupling Facility (ICF) processor resources are in use. Otherwise, when excess Internal Coupling Facility (ICF) processor resources are available, the logical partition can use them if necessary.</p>
<b>minimum-cf-processing-weight</b> <sup>1, 3, 12</sup>	(w)	Integer	<p>The minimum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Internal Coupling Facility (ICF) processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p> <p>Update:</p> <p><b>1-999</b> The minimum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>maximum-cf-processing-weight</b> <sup>1, 3, 12</sup>	(w)	Integer	<p>The maximum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p> <p>Get:</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Internal Coupling Facility (ICF) processor.</p> <p><b>1-999</b> Represents the maximum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p> <p>Update:</p> <p><b>1-999</b> Define the maximum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p>
<b>current-cf-processing-weight</b> <sup>1, 3</sup>	—	Integer	<p>The current relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p> <p><b>0</b> The <b>object-id</b> does not represent a logical partition with at least one shared Internal Coupling Facility (ICF) processor.</p> <p><b>1-999</b> Represents the current relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the Logical Partition object.</p>
<b>current-cf-processing-weight-capped</b> <sup>1, 3, 12, 13</sup>	—	Boolean	<p>Indicates whether the current Internal Coupling Facility (ICF) processing weight is a limit or a target.</p> <p><b>True</b> Indicates that the current Internal Coupling Facility (ICF) processing weight for the Logical Partition object is capped. It represents the logical partition's maximum share of resources, regardless of the availability of excess processor resources.</p> <p><b>False</b> Indicates that the current Internal Coupling Facility (ICF) processing weight for the Logical Partition is not capped. It represents the share of resources guaranteed to a logical partition when all processor resources are in use. Otherwise, when excess processor resources are available, the logical partition can use them if necessary.</p>
<b>absolute-cf-capping</b>	(w)	absolute-capping object	<p>The amount of absolute capping applied to the Internal Coupling Facility (ICF) processor.</p> <p><b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.</p>

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>program-status-word-information<sup>1</sup></b>	—	Array of psw-description objects	Describes the current PSW information for each CP associated with the logical partition. The information is obtained on each Get Logical Partition Properties request and is not cached. Refer to the description of the psw-description object for details.
<b>os-ipl-token<sup>1</sup></b>	(pc)	String (1-16)	Applicable only to z/OS, a value provided when z/OS is IPLed that uniquely identifies the instance of the operating system. Used by z/OS to obtain knowledge about the status of another system in the sysplex, and upon the demise of the system, potentially partition the system out of the sysplex immediately and reset the demised system. The value is a string of hexadecimal characters (0-9,A-Z), left justified.
<b>group-profile-capacity<sup>1</sup></b>	(w)	Integer	The current capacity value of the Group Profile with which the logical partition is associated. A null object is returned if the logical partition is not assigned to an LPAR group.
<b>group-profile-uri<sup>1</sup></b>	—	String/URI	The canonical URI of the Group Profile associated with the logical partition. A null object is returned if the logical partition is not assigned to an LPAR group.
<b>zaware-host-name<sup>14</sup></b>	(w)	String (1-64)	The IBM zAware host name. Valid characters are: a-z,A-Z,0-9, period(.), minus(-) and colon(:)
<b>zaware-master-userid<sup>14</sup></b>	(w)	String (1-32)	The IBM zAware master userid. Valid characters are: a-z,A-Z,0-9, period(.), minus(-) and underscore (_)
<b>zaware-master-pw<sup>14</sup></b>	(w)	String 98-256)	The IBM zAware master password. Valid characters are: a-z,A-Z,0-9 and !@#%&^&*( )_+{}   <>?=-  This property is not returned on a Get request, it can only be specified on an Update request.
<b>zaware-network-info<sup>14</sup></b>	(w)	Array of zaware-network objects	The set of networks available to IBM zAware. A minimum of 1 network and a maximum of 100 networks are permitted.  On an Update request, this property fully replaces the existing set.
<b>zaware-gateway-info<sup>14</sup></b>	(w)	ip-info object	The default gateway IP address information.
<b>zaware-dns-info<sup>14</sup></b>	(w)	Array of ip-info objects	The DNS IP address information. A minimum of 0 entries and a maximum of 2 entries are permitted.  On an Update request, this property fully replaces the existing set.

Table 231. Logical Partition object: class specific additional properties (continued)

Name	Qualifier	Type	Description
<b>Notes:</b>			
<ol style="list-style-type: none"> <li>1. If the logical partition <b>status</b> property is "<b>not-activated</b>", a null object is returned instead of the documented field type.</li> <li>2. An Update of this property is only valid for an <b>object-id</b> that represents a logical partition with at least one not dedicated general purpose processor.</li> <li>3. The value returned from a Get request is a null object for an <b>object-id</b> that does not represent a logical partition with at least one not dedicated general purpose processor.</li> <li>4. This property and the <b>workload-manager-enabled</b> property are mutually exclusive and cannot both be enabled at the same time. Therefore in order to enable this property it might be necessary to first disable the <b>workload-manager-enabled</b> property.</li> <li>5. This property and the various capping properties are mutually exclusive and cannot be enabled at the same time. Therefore in order to enable this property it may be necessary to first disable any capping property that is currently enabled.</li> <li>6. An Update of this property is only valid for an <b>object-id</b> that represents a logical partition with at least one not dedicated Application Assist Processor (zAAP) processor.</li> <li>7. The value returned from a Get request is always false for an <b>object-id</b> that does not represent a logical partition with at least one not dedicated Application Assist Processor (zAAP) processor.</li> <li>8. An Update of this property is only valid for an <b>object-id</b> that represents a logical partition with at least one not dedicated Integrated Facility for Linux (IFL) processor.</li> <li>9. The value returned from a Get request is always false for an <b>object-id</b> that does not represent a logical partition with at least one not dedicated Integrated Facility for Linux (IFL) processor.</li> <li>10. An Update of this property is only valid for an <b>object-id</b> that represents a logical partition with at least one not dedicated Integrated Information Processors (zIIP) processor.</li> <li>11. The value returned from a Get request is always false for an <b>object-id</b> that does not represent a logical partition with at least one not dedicated Integrated Information Processors (zIIP) processor.</li> <li>12. An Update of this property is only valid for an <b>object-id</b> that represents a logical partition with at least one not dedicated Internal Coupling Facility (ICF) processor.</li> <li>13. The value returned for a Get request is always false when the <b>object-id</b> does not represent a logical partition with at least one not dedicated Internal Coupling Facility (ICF) processor.</li> <li>14. On a Get request, this property is returned only when <b>activation-mode</b> is "<b>zaware</b>". On an Update request, this property can be updated only when <b>activation-mode</b> is "<b>zaware</b>".</li> </ol>			

## List Logical Partitions of CPC

The List Logical Partitions of CPC operation lists the logical partitions of a CPC.

### HTTP method and URI

GET /api/cpcs/{cpc-id}/logical-partitions

In this request, the URI variable {cpc-id} is the object ID of the target CPC.

### Query Parameters

Name	Type	Rqd/Opt	Description
name	String	Optional	A regular expression used to limit returned objects to those that have a matching name property. If matches are found, the response will be an array with all objects that match. If no match is found, the response will be an empty array.



## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
logical-partition	Array of logical-partition-info objects	Array of nested logical-partition-info objects (described in the next table)

Each nested logical-partition-info object contains the following fields:

Field name	Type	Description
object-uri	String/URI	Canonical URI path of the Logical Partition object
name	String	The name of the Logical Partition object
status	String Enum	The current status of the Logical Partition object

## Description

This operation lists the Logical Partition objects that belong to a CPC. The object URI, object ID and display name are provided for each.

If the **name** query parameter is specified, the returned list is limited to those Logical Partition objects that have a name property matching the specified filter pattern. If the name parameter is omitted, this filtering is not done.

An object is only included in the list if the API user has object-access permission for that object.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Object access permission to any Logical Partition object to be included in the result.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	299	A query parameter has an invalid syntax.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/logical-partitions HTTP/1.1
x-api-session: 65aw2jahugn1wop51hsq0c6aldkx773dz9ulirrvvg2z853m4u
```

---

Figure 415. List Logical Partitions of CPC: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 16:58:36 GMT
content-type: application/json;charset=UTF-8
content-length: 374
{
  "logical-partitions": [
    {
      "name": "APIVM1",
      "object-uri": "/api/logical-partitions/c7eb8134-826e-3a71-8d1a-00d706c874e9",
      "status": "operating"
    },
    {
      "name": "ZOS",
      "object-uri": "/api/logical-partitions/458e44e1-b0c2-391b-83ff-ecfd847295bd",
      "status": "not-operating"
    }
  ]
}
```

---

Figure 416. List Logical Partitions of CPC: Response

## Get Logical Partition Properties

The **Get Logical Partition Properties** operation retrieves the properties of a single Logical Partition object designated by *{logical-partition-id}*.

### HTTP method and URI

```
GET /api/logical-partitions/{logical-partition-id}
```

In this request, the URI variable *{logical-partition-id}* is the object ID of the target Logical Partition object.

### Response body contents

On successful completion, HTTP status code 200 (OK) is returned and the response body provides the current values of the properties for the Logical Partition object as defined in “Data model” on page 793.

### Description

The URI path must designate an existing Logical Partition object and the API user must have object-access permission to it. If either of these conditions is not met, status code 404 (Not Found) is returned.

On successful execution, HTTP status code 200 (OK) is returned and the response body contains all of the current properties as defined in “Data model” on page 793.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*.
- Object access permission to the logical partition's parent CPC object.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 810.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{logical-partition-id}</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/logical-partitions/c7eb8134-826e-3a71-8d1a-00d706c874e9 HTTP/1.1
x-api-session: 5obf0hwsfv1sg9kr5f93cph3zt6o5cptb61c1538wuyebdyzu4
```

---

Figure 417. Get Logical Partition Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 17:16:16 GMT
content-type: application/json;charset=UTF-8
content-length: 2366
{
  "acceptable-status": [
    "operating"
  ],
  "activation-mode": "esa390",
  "additional-status": "",
  "class": "logical-partition",
  "cluster-name": "",
  "current-aap-processing-weight": null,
  "current-aap-processing-weight-capped": null,
  "current-cf-processing-weight": null,
  "current-cf-processing-weight-capped": null,
  "current-ifl-processing-weight": null,
  "current-ifl-processing-weight-capped": null,
  "current-processing-weight": 100,
  "current-processing-weight-capped": false,
  "current-ziip-processing-weight": null,
  "current-ziip-processing-weight-capped": null,
  "defined-capacity": 0,
  "description": "LPAR Image",
  "group-profile-capacity": null,
  "group-profile-uri": null,
  "has-operating-system-messages": false,
  "has-unacceptable-status": false,
  "initial-aap-processing-weight": null,
  "initial-aap-processing-weight-capped": null,
  "initial-cf-processing-weight": null,
  "initial-cf-processing-weight-capped": null,
  "initial-ifl-processing-weight": null,
  "initial-ifl-processing-weight-capped": null,
  "initial-processing-weight": 100,
  "initial-processing-weight-capped": false,
  "initial-ziip-processing-weight": null,
  "initial-ziip-processing-weight-capped": null,
  "is-locked": false,
  "last-used-activation-profile": "APIVM1",
  "maximum-aap-processing-weight": null,
  "maximum-cf-processing-weight": null,
  "maximum-ifl-processing-weight": null,
  "maximum-processing-weight": 200,
  "maximum-ziip-processing-weight": null,
  "absolute-aap-capping":{"type": "none"},
  "absolute-cf-capping":{"value": 88.52, "type": "processors"},
  "absolute-ifl-capping":{"type": "none"}
  "absolute-processing-capping":{"value": 0.01, "type": "processors"},
  "absolute-ziip-capping":{"value": 2.01, "type": "processors"},
}
```

---

Figure 418. Get Logical Partition Properties: Response (Part 1)

```

"minimum-aap-processing-weight": null,
"minimum-cf-processing-weight": null,
"minimum-ifl-processing-weight": null,
"minimum-processing-weight": 50,
"minimum-ziip-processing-weight": null,
"name": "APIVM1",
"next-activation-profile-name": "APIVM1",
"object-id": "c7eb8134-826e-3a71-8d1a-00d706c874e9",
"object-uri": "/api/logical-partitions/c7eb8134-826e-3a71-8d1a-00d706c874e9",
"os-ipl-token": "0000000000000000",
"os-level": "6.2.0",
"os-name": "APIVM1",
"os-type": "z/VM",
"parent": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
"partition-number": 1,
"program-status-word-information": [
  {
    "cpid": "00",
    "psw": "07064000800000000000000000000000"
  },
  {
    "cpid": "01",
    "psw": "07064000800000000000000000000000"
  }
],
"status": "operating",
"sysplex-name": "SSICAPI1",
"workload-manager-enabled": true
}

```

Figure 419. Get Logical Partition Properties: Response (Part 2)

## Update Logical Partition Properties

The **Update Logical Partition Properties** operation updates one or more writeable properties of the Logical Partition object designated by *{logical-partition-id}*.

### HTTP method and URI

**POST** /api/logical-partitions/{*logical-partition-id*}

In this request, the URI variable *{logical-partition-id}* is the object ID of the target Logical Partition object.

### Request body contents

The request body is expected to contain one or more field names representing writable logical partition properties, along with the new values for those fields.

The request body can and should omit fields for properties whose values are not to be changed by this operation. Properties for which no input value is provided remain unchanged by this operation.

### Description

The request body object is validated against the data model for the Logical Partition object type to ensure that the request body contains only writeable properties and the data types of those properties are as required. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

On successful execution, the value of each corresponding property of the object is updated with the value provided by the input field, and status code 204 (No Content) is returned.

When this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **Change Object Definition** task.
- Object access permission to the logical partition's parent CPC object.
- For a logical partition whose **activation-mode** is "zaware", action/task permission for the **Firmware Details** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
403 (Forbidden)	1	The API user does not have the required permission for this operation.
404 (Not Found)	1	The object ID in the URI ( <i>{logical-partition-id}</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Activate Logical Partition

The **Activate Logical Partition** operation activates the Logical Partition object designated by *{logical-partition-id}*.

### HTTP method and URI

**POST** `/api/logical-partitions/{logical-partition-id}/operations/activate`

In this request, the URI variable *{logical-partition-id}* is the object ID of the target Logical Partition object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
activation-profile-name	String (1-16)	Optional	The name of the activation profile to be used for the request. If not provided, the request uses the profile name specified in the <b>next-activation-profile-name</b> property for the Logical Partition object.

Field name	Type	Rqd/Opt	Description
force	Boolean	Optional	Whether this operation is permitted when the logical partition is in "operating" status (true) or not (false). The default is false.

## Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

## Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "complete", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

## Description

Activation is a process that makes a logical partition operational, which means either:

- The logical partition is ready to have a control program or operating system loaded, or
- The logical partition has loaded and is running a control program or operating system.

Activating a logical partition includes:

- Initializing the logical partition
- Allocating system resources to the logical partition
- Loading the logical partition with a control program or operating system.

Since the status of the logical partition determines which operations must be performed during activation to make the logical partition operational, one or more operations listed above may not be performed during activation.

If planning to load the z/VM operating system in this logical partition, refer to Chapter 10, "Virtualization management," on page 209 for details on the activation of virtual servers.

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes" on page 816.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **Activate** task.
- Object access permission to the logical partition's parent CPC object.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 815.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>logical-partition-id</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed or was rejected due to the current logical partition status and use of the force=false parameter. If rejected due to force=false, the logical partition status is unchanged. If the operation failed, the logical partition status is unknown. Refer to the message parameter in the error response body for details.

## Deactivate Logical Partition

The **Deactivate Logical Partition** operation deactivates the Logical Partition object designated by *logical-partition-id*.

### HTTP method and URI

**POST** /api/logical-partitions/{*logical-partition-id*}/operations/deactivate

In this request, the URI variable *logical-partition-id* is the object ID of the target Logical Partition object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
force	Boolean	Optional	Whether this operation is permitted when the logical partition is in "operating" status (true) or not (false). The default is false.



## Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

## Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "complete", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

## Description

Deactivation is an orderly process for terminating a logical partition.

Deactivating a logical partition includes:

- Unloading the logical partition's control program or operating system
- Freeing system resources allocated to the logical partition.

After the logical partition is deactivated, the logical partition is no longer operational

If this logical partition currently has the z/VM operating system loaded, refer to Chapter 10, "Virtualization management," on page 209 for the details on deactivation of virtual servers.

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes" on page 818.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **Deactivate** task.
- Object access permission to the logical partition's parent CPC object.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>logical-partition-id</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed or was rejected due to the current logical partition status and use of the <code>force=false</code> parameter. If rejected due to <code>force=false</code> , the logical partition status is unchanged. If the operation failed, the logical partition status is unknown. Refer to the message parameter in the error response body for details.

## Reset Normal

The **Reset Normal** operation initializes a system or logical partition by clearing its pending interruptions, resetting its channel subsystem and resetting its processors. A reset prepares a system or logical partition for loading it with an operating system.

### HTTP method and URI

**POST** `/api/logical-partitions/{logical-partition-id}/operations/reset-normal`

In this request, the URI variable *{logical-partition-id}* is the object ID of the target Logical Partition object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
<code>force</code>	Boolean	Optional	Whether this operation is permitted when the logical partition is in " <b>operating</b> " status (true) or not (false). The default is false.
<code>os-ipl-token</code>	String (1-16)	Optional	Applicable only to z/OS, this parameter requests that this operation only be performed if the provided value matches the current value of the <b>os-ipl-token</b> property. This ensures that this operation is targeting the same IPL instance as when the <b>os-ipl-token</b> property was retrieved. IBM recommends that this parameter only be provided by callers that fully understand how the <b>os-ipl-token</b> parameter is managed by z/OS. The value is a string of hexadecimal characters (0-9, A-Z), left justified.

## Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

## Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "complete", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

## Description

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes" on page 820.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **Reset Normal** task.
- Object access permission to the logical partition's parent CPC object.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	264	The specified IPL Token value does not match the current IPL Token value.
404 (Not Found)	1	The object ID in the URI ( <i>{logical-partition-id}</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed or was rejected due to the current logical partition status and use of the force=false parameter. If rejected due to force=false, the logical partition status is unchanged. If the operation failed, the logical partition status is unknown. Refer to the message parameter in the error response body for details.

## Reset Clear

The **Reset Clear** operation initializes system or logical partition by clearing its pending interruptions, resetting its channel subsystem and resetting its processors. A reset prepares a system or logical partition for loading it with an operating system and clears main memory of the system or logical partition.

### HTTP method and URI

**POST** /api/logical-partitions/{*logical-partition-id*}/operations/reset-clear

In this request, the URI variable {*logical-partition-id*} is the object ID of the target Logical Partition object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
force	Boolean	Optional	Whether this operation is permitted when the logical partition is in " <b>operating</b> " status (true) or not (false). The default is false.
os-ipl-token	String (1-16)	Optional	Applicable only to z/OS, this parameter requests that this operation only be performed if the provided value matches the current value of the <b>os-ipl-token</b> property. This ensures that this operation is targeting the same IPL instance as when the <b>os-ipl-token</b> property was retrieved. IBM recommends that this parameter only be provided by callers that fully understand how the <b>os-ipl-token</b> parameter is managed by z/OS. The value is a string of hexadecimal characters (0-9,A-Z), left justified.

### Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

### Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "complete", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

## Description

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes."

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **Reset Clear** task.
- Object access permission to the logical partition's parent CPC object.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in "Response body contents" on page 820.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	264	The specified IPL Token value does not match the current IPL Token value.
404 (Not Found)	1	The object ID in the URI ( <i>{logical-partition-id}</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed or was rejected due to the current logical partition status and use of the force=false parameter. If rejected due to force=false, the logical partition status is unchanged. If the operation failed, the logical partition status is unknown. Refer to the message parameter in the error response body for details.

## Load Logical Partition

The **Load Logical Partition** operation resets a logical partition, to prepare it for loading an operating system, and loads the operating system.

### HTTP method and URI

**POST** /api/logical-partitions/{*logical-partition-id*}/operations/load

In this request, the URI variable {*logical-partition-id*} is the object ID of the target Logical Partition object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
load-address	String (1-5)	Required	The hexadecimal address of an I/O device that provides access to the control program to be loaded. The input value is right justified and padded with zeros to 5 characters.  Valid values are in the range "00000" to "nFFFF" where "n" is the number of subchannel sets provided by the CPC minus 1. So, for example, on a CPC that provides 3 subchannel sets, the valid range is "00000" to "2FFFF".
load-parameter	String (1-8)	Optional	Some control programs support the use of this property to provide additional control over the outcome of a <b>Load</b> operation. Refer to the configuration documentation for the control program to be loaded to see if this parameter is supported and if so, what values and format is supported. Omitting this field indicates that the value for this field is to be retrieved from the current IOCDs. Valid characters are 0-9, A-Z, blank and period.
clear-indicator	Boolean	Optional	Whether memory should be cleared before performing the <b>Load</b> (true) or not cleared (false). The default value is true.
timeout	Integer (60-600)	Optional	Amount of time, in seconds, to wait for the <b>Load</b> to complete. The default timeout value is 60 seconds.
store-status-indicator	Boolean	Optional	Whether status should be stored before performing the <b>Load</b> (true) or not stored (false). The default is false.
force	Boolean	Optional	Whether this operation is permitted when the logical partition is in " <b>operating</b> " status (true) or not (false). The default is false.
os-ipl-token	String (1-16)	Optional	Applicable only to z/OS, this parameter requests that this operation only be performed if the provided value matches the current value of the <b>os-ipl-token</b> property. This ensures that this operation is targeting the same IPL instance as when the <b>os-ipl-token</b> property was retrieved. IBM recommends that this parameter only be provided by callers that fully understand how the <b>os-ipl-token</b> parameter is managed by z/OS. The value is a string of hexadecimal characters (0-9, A-Z), left justified.

### Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

## Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "complete", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

## Description

This operation is not permitted for a logical partition whose **activation-mode** property is "zaware".

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes" on page 824.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **Load** task.
- Object access permission to the logical partition's parent CPC object.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in "Response body contents" on page 822.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	264	The specified IPL Token value does not match the current IPL Token value.
	306	This operation is not valid in the current activation mode.
404 (Not Found)	1	The object ID in the URI ( <i>{logical-partition-id}</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed or was rejected due to the current logical partition status and use of the force=false parameter. If rejected due to force=false, the logical partition status is unchanged. If the operation failed, the logical partition status is unknown. Refer to the message parameter in the error response body for details.

## PSW Restart

The **PSW Restart** operation restarts the first available processor of the Logical Partition object designated by *{logical-partition-id}*.

### HTTP method and URI

**POST** /api/logical-partitions/{*logical-partition-id*}/operations/psw-restart

In this request, the URI variable *{logical-partition-id}* is the object ID of the target Logical Partition object.

### Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

## Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "**complete**", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

## Description

This operation is not permitted for a logical partition whose **activation-mode** property is "**zaware**".

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes" on page 825.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **PSW Restart** task.
- Object access permission to the logical partition's parent CPC object.



## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 824.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	306	This operation is not valid in the current activation mode.
404 (Not Found)	1	The object ID in the URI ( <i>logical-partition-id</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed.

## Start Logical Partition

The **Start Logical Partition** operation starts the processors to process instructions of the Logical Partition object designated by *logical-partition-id*.

### HTTP method and URI

**POST** /api/logical-partitions/{*logical-partition-id*}/operations/start

In this request, the URI variable *logical-partition-id* is the object ID of the target Logical Partition object.

### Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

### Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "complete", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

## Description

This operation is not permitted for a logical partition whose **activation-mode** property is "zaware".

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes."

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **Start** task.
- Object access permission to the logical partition's parent CPC object.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in "Response body contents" on page 825.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	306	This operation is not valid in the current activation mode.
404 (Not Found)	1	The object ID in the URI ( <i>{logical-partition-id}</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed.

## Stop Logical Partition

The **Stop Logical Partition** operation stops the processors from processing instructions of the Logical Partition object designated by *{logical-partition-id}*.

### HTTP method and URI

**POST** /api/logical-partitions/{*logical-partition-id*}/operations/stop

In this request, the URI variable *{logical-partition-id}* is the object ID of the target Logical Partition object.

### Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

### Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "**complete**", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

### Description

This operation is not permitted for a logical partition whose **activation-mode** property is "**zaware**".

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes" on page 828.

### Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **Stop** task.
- Object access permission to the logical partition's parent CPC object.

### HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	306	This operation is not valid in the current activation mode.
404 (Not Found)	1	The object ID in the URI ( <i>logical-partition-id</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed.

## SCSI Load

The **SCSI Load** operation clears main storage, to prepare the logical partition for loading an operating system, and loads the operating system from the designated SCSI device.

### HTTP method and URI

**POST** /api/logical-partitions/{*logical-partition-id*}/operations/scsi-load

In this request, the URI variable *logical-partition-id* is the object ID of the target Logical Partition object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
load-address	String (1-5)	Required	The hexadecimal address of an I/O device that provides access to the control program to be loaded. The input value is right justified and padded with zeros to 5 characters.  Valid values are in the range "00000" to "nFFFF" where "n" is the number of subchannel sets provided by the CPC minus 1. So, for example, on a CPC that provides 3 subchannel sets, the valid range is "00000" to "2FFFF".
load-parameter	String (1-8)	Optional	Some control programs support the use of this property to provide additional control over the outcome of a <b>Load</b> operation. Refer to the configuration documentation for the control program to be loaded to see if this parameter is supported and if so, what values and format is supported. Omitting this field indicates that the value for this field is to be retrieved from the current IOCDs. Valid characters are 0-9, A-Z, blank and period

Field name	Type	Rqd/Opt	Description
world-wide-port-name	String (1-16)	Required	The worldwide port name (WWPN) of the target SCSI device to be used for this operation, in hexadecimal.
logical-unit-number	String (1-16)	Required	The hexadecimal logical unit number (LUN) to be used for the <b>SCSI Load</b> .
disk-partition-id	Integer (0-30)	Optional	The disk-partition-id (also called the boot program selector) to be used for the <b>SCSI Load</b> . The default value is 0.
operating-system-specific-load-parameters	String (1-256)	Optional	The operating system specific load parameters to be used for the <b>SCSI Load</b> . The default value is an empty string.
boot-record-logical-block-address	String (1-16)	Optional	The hexadecimal boot record logical block address to be used for the <b>SCSI Load</b> . The default value is hex zeros.
force	Boolean	Optional	Whether this operation is permitted when the logical partition is in <b>"operating"</b> status (true) or not (false). The default is false.
os-ipl-token	String (1-16)	Optional	Applicable only to z/OS, this parameter requests that this operation only be performed if the provided value matches the current value of the <b>os-ipl-token</b> property. This ensures that this operation is targeting the same IPL instance as when the <b>os-ipl-token</b> property was retrieved. IBM recommends that this parameter only be provided by callers that fully understand how the <b>os-ipl-token</b> parameter is managed by z/OS. The value is a string of hexadecimal characters (0-9, A-F), left justified.

## Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

## Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is **"complete"**, the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

## Description

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes" on page 830.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*

- Action/task permission for the **Load** task.
- Object access permission to the logical partition's parent CPC object.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in “Response body contents” on page 829.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>logical-partition-id</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed.

## SCSI Dump

The **SCSI Dump** operation loads a standalone dump program from a designated SCSI device.

### HTTP method and URI

**POST** /api/logical-partitions/{*logical-partition-id*}/operations/scsi-dump

In this request, the URI variable *logical-partition-id* is the object ID of the target Logical Partition object.

### Request body contents

The request body is expected to contain a JSON object with the following fields:

Field name	Type	Rqd/Opt	Description
load-address	String (1-5)	Required	<p>The hexadecimal address of an I/O device that provides access to the control program to be loaded. The input value is right justified and padded with zeros to 5 characters.</p> <p>Valid values are in the range "00000" to "nFFFF" where "n" is the number of subchannel sets provided by the CPC minus 1. So, for example, on a CPC that provides 3 subchannel sets, the valid range is "00000" to "2FFFF".</p>

Field name	Type	Rqd/Opt	Description
load-parameter	String (1-8)	Optional	Some control programs support the use of this property to provide additional control over the outcome of a <b>Load</b> operation. Refer to the configuration documentation for the control program to be loaded to see if this parameter is supported and if so, what values and format is supported. Omitting this field indicates that the value for this field is to be retrieved from the current IOCDs. Valid characters are 0-9, A-Z, blank and period.
world-wide-port-name	String (1-16)	Required	The worldwide port name (WWPN) of the target SCSI device to be used for this operation, in hexadecimal.
logical-unit-number	String (1-16)	Required	The hexadecimal logical unit number (LUN) to be used for the <b>SCSI Load</b> .
disk-partition-id	Integer (0-30)	Optional	The disk-partition-id (also called the boot program selector) to be used for the <b>SCSI Load</b> . The default value is 0.
operating-system-specific-load-parameters	String (1-256)	Optional	The operating system specific load parameters to be used for the <b>SCSI Load</b> . The default value is an empty string.
boot-record-logical-block-address	String (1-16)	Optional	The hexadecimal boot record logical block address to be used for the <b>SCSI Load</b> . The default value is hex zeros.
os-ipl-token	String (1-16)	Optional	Applicable only to z/OS, this parameter requests that this operation only be performed if the provided value matches the current value of the <b>os-ipl-token</b> property. This ensures that this operation is targeting the same IPL instance as when the <b>os-ipl-token</b> property was retrieved. IBM recommends that this parameter only be provided by callers that fully understand how the <b>os-ipl-token</b> parameter is managed by z/OS. The value is a string of hexadecimal characters (0-9, A-Z), left justified.

## Response body contents

Once the operation is accepted, the response body contains a JSON object with the following fields:

Field name	Type	Description
job-uri	String/URI	URI that may be queried to retrieve status updates.

## Asynchronous result description

Once the operation has completed, a job-completion notification is sent and results are available for the asynchronous portion of this operation. These results are retrieved using the **Query Job Status** operation directed at the job URI provided in the response body.

The result document returned by the **Query Job Status** operation is specified in the description for the **Query Job Status** operation. When the status of the job is "**complete**", the results include a job completion status code and reason code (fields **job-status-code** and **job-reason-code**) which are set as indicated in operation description. The **job-results** field is null for this operation.

## Description

When the operation is initiated, a 202 (Accepted) status code is returned. The response body includes a URI that may be queried to retrieve the status of the operation. See "Query Job Status" on page 52 for information on how to query job status. When the operation has completed, an asynchronous result message is sent. See "Job status and reason codes" on page 832.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the Logical Partition object designated by *{logical-partition-id}*
- Action/task permission for the **SCSI Dump** task.
- Object access permission to the logical partition's parent CPC object.

## HTTP status and reason codes

On success, HTTP status code 202 (Accepted) is returned and the response body is provided as described in "Response body contents" on page 831.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{logical-partition-id}</i> ) does not designate an existing Logical Partition object, or the API user does not have object access permission to the object.
500 (Server Error)	280	An IO exception occurred during the scheduling of the asynchronous request.

Additional standard status and reason codes can be returned, as described in Chapter 3, "Invoking API operations," on page 19.

## Job status and reason codes

HTTP error status code	Reason code	Description
204 (No Content)	N/A	Operation completed successfully.
500 (Server Error)	263	Operation failed.

## | Inventory service data

| Information about logical partitions can be optionally included in the inventory data provided by the Inventory Service.

| Inventory entries for the Logical Partition objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**logical-partition**" are to be included. An entry for a particular logical partition is included only if the API user has access permission to that object as described in the **Get Logical Partition Properties** operation.

| For each Logical Partition object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for "Get Logical Partition Properties" on page 810. That is, the data provided is the same as would be provided if a **Get Logical Partition Properties** operation were requested targeting this object.



## Reset activation profile

A Reset activation profile is used by a CPC Activate operation to control the activation of a CPC and, if properly configured with one or more image activation profiles, a set of Logical Partition(s).

An activation profile can only be created or deleted from the Hardware Management Console or the Support Element.

- | For information on customizing activation profiles, Support Element (Version 2.12.1 and newer)
- | information can be found on the console help system, or on the IBM Knowledge Center at
- | <https://www.ibm.com/support/knowledgecenter>. (Select **z Systems** on the navigation bar, and then
- | select your server). For information from earlier versions of the Support Element, see the *Support Element*
- | *Operations Guide*.

## Data model

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

This element includes the following properties.

Table 232. Reset activation profile: properties

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/ URI	The canonical URI path of the Reset Activation Profile object, of the form <code>/api/cpcs/{cpc-id}/reset-activation-profiles/{reset-activation-profile-name}</code> where <code>{reset-activation-profile-name}</code> is the value of the name property (Reset Activation Profile name).
<b>parent</b>	—	String/ URI	The canonical URI path of the associated CPC object.
<b>class</b>	—	String	The class of a Reset Activation Profile object is <b>"reset-activation-profile"</b> .
<b>name</b>	—	String (1-16)	The activation profile name, which uniquely identifies this profile within the set of activation profiles for the CPC object designated by <code>{cpc-id}</code> .
<b>description</b>	(w)	String (1-50)	The reset profile description
<b>iocds-name</b>	(w)	String (0-2)	The Input/Output Configuration Data Set name, in hexadecimal. An empty string indicates that the currently active IOCDS will be used. The active IOCDS is the one from the most recent power-on-reset of the CPC or, if using dynamic I/O configuration, the one last activated.
<b>processor-running-time-type</b>	(w)	String Enum	Defines whether the processor running time is determined dynamically or set manually for the CPC (see processor-running-time in this table). One of: <ul style="list-style-type: none"> <li>• <b>"system-determined"</b></li> <li>• <b>"user-determined"</b></li> </ul>
<b>processor-running-time</b>	(w)	Integer (0-100)	Amount of continuous time, in milliseconds, for logical processors to perform jobs on shared processors for the CPC, if <b>processor-running-time-type</b> is set to <b>"user-determined"</b> . If <b>processor-running-time-type</b> is <b>"system-determined"</b> , this property's value will always be returned as 0.
<b>end-timeslice-on-wait</b>	(w)	Boolean	If true and if <b>processor-running-time-type</b> is set to <b>"user-determined"</b> , CPC Logical Partitions lose their share of running time when they enter a wait state. If <b>processor-running-time-type</b> is <b>"system-determined"</b> , this property's value will always be returned as false.

## List Reset Activation Profiles

The **List Reset Activation Profiles** operation lists the Reset Activation Profiles associated with a particular CPC.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/reset-activation-profiles

In this request, the URI variable {cpc-id} is the object ID of the target CPC object.

#### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	A regular expression used to limit returned objects to those that have a matching name property. If matches are found, the response will be an array with all objects that match. If no match is found, the response will be an empty array.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
reset-activation-profiles	Array of reset-actprof-info objects	Array of nested objects (described in the following table).

Each reset-actprof-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Reset Activation Profile object.
name	String	The name of the Reset Activation Profile.

### Description

This operation lists the Reset Activation Profiles associated with a particular CPC.

If the **name** query parameter is specified, the returned list is limited to those Reset Activation Profiles that have a name property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

### Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by {cpc-id}.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 834.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	299	A query parameter has an invalid syntax.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
500 (Server Error)	281	An unexpected error occurred during the collection of the list of activation profiles.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/reset-activation-profiles HTTP/1.1
x-api-session: 5obf0hwsfv1sg9kr5f93cph3zt6o5cptb61cl538wuyebdyzu4
```

---

Figure 420. List Reset Activation Profiles: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 17:16:16 GMT
content-type: application/json;charset=UTF-8
content-length: 372
{
  "reset-activation-profiles": [
    {
      "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/reset-activation-profiles/
DEFAULT",
      "name": "DEFAULT"
    },
    {
      "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/reset-activation-profiles/
POWER_ON_RESET",
      "name": "POWER_ON_RESET"
    }
  ]
}
```

---

Figure 421. List Reset Activation Profiles: Response

## Get Reset Activation Profile Properties

The **Get Reset Activation Profile Properties** operation retrieves the properties of a single Reset Activation Profile designated by *{reset-activation-profile-name}*.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/reset-activation-profiles/{reset-activation-profile-name}

#### URI variables:

Variable	Description
<i>{cpc-id}</i>	Object ID of the target CPC object.
<i>{reset-activation-profile-name}</i>	Reset Activation Profile name

### Response body contents

On successful completion, the response body provides the current values of the properties for the Reset Activation Profile as defined in the “Data model” on page 833.

### Description

The URI path must designate an existing Reset Activation Profile and the API user must have object-access permission to the CPC. If either of these conditions is not met, status code 404 (Not Found) is returned.

On successful execution, HTTP status code 200 (OK) is returned and the response body contains all of the current properties as defined by the Data Model for the Reset Activation Profile object.

### Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	260	The activation profile name in the URI ( <i>{reset-activation-profile-name}</i> ) does not designate an existing activation profile.
500 (Server Error)	281	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/reset-activation-profiles/DEFAULT HTTP/1.1
x-api-session: 5obf0hwsfv1sg9kr5f93cph3zt6o5cptb61c1538wuyebdyzu4
```

---

Figure 422. Get Reset Activation Profile Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 17:16:18 GMT
content-type: application/json;charset=UTF-8
content-length: 384
{
  "class": "reset-activation-profile",
  "description": "This is the default Reset profile.",
  "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/reset-activation-profiles/
  DEFAULT",
  "end-timeslice-on-wait": false,
  "iocds-name": "a0",
  "name": "DEFAULT",
  "parent": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "processor-running-time": 0,
  "processor-running-time-type": "system-determined"
}
```

---

Figure 423. Get Reset Activation Profile Properties: Response

## Update Reset Activation Profile Properties

The **Update Reset Activation Profile Properties** operation updates one or more writeable properties of the Reset Activation Profile designated by *{reset-activation-profile-name}*.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/reset-activation-profiles/{reset-activation-profile-name}

#### URI variables:

Variable	Description
<i>{cpc-id}</i>	Object ID of the target CPC object.
<i>{reset-activation-profile-name}</i>	Reset Activation Profile name

### Request body contents

The request body is expected to contain one or more field names representing writable Reset Activation Profile properties, along with the new values for those fields.

The request body can and should omit fields for properties whose values are not to be changed by this operation. Properties for which no input value is provided remain unchanged by this operation.

## Description

The request body object is validated against the data model for the Reset Activation Profile to ensure that the request body contains only writeable properties and the data types of those properties are as required. If the request body is not valid, status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

On successful execution, the value of each corresponding property of the Reset Activation Profile is updated with the value provided by the input field, and status code 204 (No Content) is returned.

When this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission for the **Customize/Delete Activation Profiles** task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	300	The provided update values would result in an illegal state. Verify that the values are both internally consistent and consistent with the current state of the profile.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	260	The activation profile name in the URI ( <i>{reset-activation-profile-name}</i> ) does not designate an existing activation profile.
409 (Conflict)	2	The operation was rejected by the Support Element (SE), because the SE is currently performing processing that requires exclusive control of the SE. Retry the operation at a later time.
500 (Server Error)	281	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## | Inventory service data

| Information about reset activation profiles can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for the Reset Activation Profile objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "cpc" are to be included. An entry for a particular reset activation profile is included only if the API user has access permission to that object as described in the **Get Reset Activation Profile Properties** operation.

For each Reset Activation Profile object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for "Get Reset Activation Profile Properties" on page 836. That is, the data provided is the same as would be provided if a **Get Reset Activation Profile Properties** operation were requested targeting this object.

---

## Image activation profile

An Image activation profile is used by an Activate operation to activate a logical partition of a previously activated CPC.

An activation profile can only be created or deleted from the Hardware Management Console or the Support Element.

For information on customizing activation profiles, Support Element (Version 2.12.1 and newer) information can be found on the console help system, or on the IBM Knowledge Center at <https://www.ibm.com/support/knowledgecenter>. (Select **z Systems** on the navigation bar, and then select your server). For information from earlier versions of the Support Element, see the *Support Element Operations Guide*.

## Data model

For definitions of the qualifier abbreviations in the following tables, see "Property characteristics" on page 38.

This element includes the following properties. Some properties have additional notes associated with them. Refer to the table notes at the end of this table.

Table 233. Image activation profile: properties.

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/URI	The canonical URI path of the Image Activation Profile object, of the form <code>/api/cpcs/{cpc-id}/image-activation-profiles/{image-activation-profile-name}</code> where <code>{image-activation-profile-name}</code> is the value of the name property (Image Activation Profile name).
<b>parent</b>	—	String/URI	The canonical URI path of the associated CPC object.
<b>class</b>	—	String	The class of an Image Activation Profile object is " <b>image-activation-profile</b> ".
<b>name</b>	—	String (0-16)	The activation profile name, which uniquely identifies this profile within the set of activation profiles for the CPC object designated by <code>{cpc-id}</code> .
<b>description</b>	(w)	String (0-50)	The activation profile description

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>ipl-address</b> <sup>15</sup>	(w)	String (0-5)	<p>The hexadecimal address of an I/O device that provides access to the control program to be loaded. The input value will be right justified and padded with zeros to 5 characters. An empty string indicates that the value for this property is to be retrieved from the IOCDS used during a subsequent Load operation.</p> <p>Valid values are in the range "00000" to "nFFFF" where "n" is the number of subchannel sets provided by the CPC minus 1. So, for example, on a CPC that provides 3 subchannel sets, the valid range is "00000" to "2FFFF".</p>
<b>ipl-parameter</b>	(w)	String (0-8)	<p>Some control programs support the use of this property to provide additional control over the outcome of a Load operation. Refer to the configuration documentation for the control program to be loaded to see if this parameter is supported and if so, what values and format is supported. An empty string indicates that the value for this property is to be retrieved from the IOCDS used during a subsequent Load operation. Valid characters are 0-9, A-Z, blank and period. On an Update, a non-empty string is left justified and right padded with blanks to 8 characters.</p>
<b>initial-processing-weight</b> <sup>1</sup>	(w)	Integer	<p>The relative amount of shared general purpose processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared general purpose processor.</p> <p><b>1-999</b> Represents the relative amount of shared general purpose processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the relative amount of shared general purpose processor resources allocated to the logical partition.</p>
<b>initial-processing-weight-capped</b> <sup>1, 2, 3</sup>	(w)	Boolean	<p>Whether the initial processing weight for general purpose processors is a limit or a target.</p> <p><b>True:</b> Indicates that the initial general purpose processor processing weight for the logical partition is capped. It represents the logical partition's maximum share of general purpose processor resources.</p> <p><b>False:</b> Indicates that the initial general purpose processor processing weight for the logical partition is not capped. It represents the share of general purpose processor resources guaranteed to a logical partition when all general purpose processor resources are in use. Otherwise, when excess general purpose processor resources are available, the logical partition can use them if necessary.</p>



Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>minimum-processing-weight<sup>1</sup></b>	(w)	Integer	<p>The minimum relative amount of shared general purpose processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared general purpose processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared general purpose processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>0</b> There is no minimum value for the processing weight.</p> <p><b>1-999</b> Define the minimum relative amount of shared general purpose processor resources allocated to the logical partition. The value must be less than or equal to the initial-processing-weight property.</p>
<b>maximum-processing-weight<sup>1</sup></b>	(w)	Integer	<p>The maximum relative amount of shared general purpose processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared general purpose processor.</p> <p><b>1-999</b> Represents the maximum relative amount of shared general purpose processor resources allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Defines the maximum relative amount of shared general purpose processor resources allocated to the logical partition. Must be greater than or equal to the initial-processing-weight property.</p>
<b>absolute-general-purpose-capping</b>	(w)	absolute-capping object	<p>The amount of absolute capping applied to the general purpose processor.</p> <p><b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.</p>
<b>workload-manager-enabled<sup>4</sup></b>	(w)	Boolean	<p>Whether or not z/OS Workload Manager is allowed to change processing weight related properties.</p> <p><b>True:</b> Indicates that z/OS Workload Manager is allowed to change processing weight related properties for this logical partition.</p> <p><b>False:</b> Indicates that z/OS Workload Manager is not allowed to change processing weight related properties for this logical partition.</p>

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>defined-capacity</b>	(w)	Integer	The defined capacity expressed in terms of Millions of Service Units (MSU)s per hour. MSU is a measure of processor resource consumption. The amount of MSUs a logical partition consumes is dependent on the model, the number of logical processors available to the partition, and the amount of time the logical partition is dispatched. The defined capacity value specifies how much capacity the logical partition is to be managed to by z/OS Workload Manager for the purpose of software pricing.  <b>0:</b> No defined capacity is specified for this logical partition.  <b>1-nnnn:</b> Represents the amount of defined capacity specified for this logical partition
<b>ipl-type</b>	(w)	String Enum	One of: <ul style="list-style-type: none"> <li>• <b>"ipltype-standard"</b> - This image activation profile is used to perform a standard load.</li> <li>• <b>"ipltype-scsi"</b> - This image activation profile is used to perform a SCSI load.</li> <li>• <b>"ipltype-scsidump"</b> - This image activation profile is used to perform a SCSI dump.</li> </ul>
<b>worldwide-port-name<sup>15</sup></b>	(w)	String (1-16)	Worldwide port name of the target SCSI device, used for a SCSI load or SCSI dump, in hexadecimal.
<b>disk-partition-id</b>	(w)	Integer (0-30)	The disk partition number (also called the boot program selector) for the activation profile, used for a SCSI load or SCSI dump.
<b>logical-unit-number<sup>15</sup></b>	(w)	String (1-16)	Logical unit number value for the activation profile, used for a SCSI load or SCSI dump, in hexadecimal.
<b>boot-record-lba<sup>15</sup></b>	(w)	String (1-16)	Boot record logical block address for the activation profile, used for a SCSI load or SCSI dump, in hexadecimal.
<b>os-specific-load-parameters</b>	(w)	String (0-256)	Operating system-specific load parameters for the activation profile, used for a SCSI load or SCSI dump. On an Update, value is left justified and right padded with blanks to 256 characters.
<b>initial-aap-processing-weight<sup>5</sup></b>	(w)	Integer	The relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition at activation.  Get:  <b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Application Assist Processor (zAAP) processor.  <b>1-999</b> Represents the relative amount of shared Application Assist Processor (zAAP) processor resources initially allocated to the logical partition.  Update:  <b>1-999</b> Define the relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>initial-aap-processing-weight-capped</b> <sup>3, 5, 6</sup>	(w)	Boolean	<p>Whether the initial processing weight for Application Assist Processor (zAAP) processors is a limit or a target.</p> <p><b>True:</b> Indicates that the initial Application Assist Processor (zAAP) processor processing weight for the logical partition is capped. It represents the logical partition's maximum share of Application Assist Processor (zAAP) processor resources, regardless of the availability of excess Application Assist Processor (zAAP) processor resources.</p> <p><b>False:</b> Indicates that the initial Application Assist Processor (zAAP) processor processing weight for the logical partition is not capped. It represents the share of Application Assist Processor (zAAP) processor resources guaranteed to a logical partition when all Application Assist Processor (zAAP) processor resources are in use. Otherwise, when excess Application Assist Processor (zAAP) processor resources are available, the logical partition can use them if necessary.</p>
<b>minimum-aap-processing-weight</b> <sup>5</sup>	(w)	Integer	<p>The minimum relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Application Assist Processor (zAAP) processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared Application Assist Processor (zAAP) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>0</b> No minimum value for the processing weight.</p> <p><b>1-999</b> Define the minimum relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p>

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>maximum-aap-processing-weight<sup>5</sup></b>	(w)	Integer	<p>The maximum relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Application Assist Processor (zAAP) processor.</p> <p><b>1-999</b> Represents the maximum relative amount of shared Application Assist Processor (zAAP) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the maximum relative amount of shared Application Assist Processor (zAAP) processor resources allocated to the logical partition.</p>
<b>absolute-aap-capping</b>	(w)	absolute-capping object	<p>The amount of absolute capping applied to the Application Assist Processor (zAAP).</p> <p><b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.</p>
<b>initial-ifl-processing-weight<sup>7</sup></b>	(w)	Integer	<p>The relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition at activation.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Integrated Facility for Linux (IFL) processor.</p> <p><b>1-999</b> Represents the relative amount of shared Integrated Facility for Linux (IFL) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p>

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>initial-ift-processing-weight-capped</b> <sup>3, 7, 8</sup>	(w)	Boolean	<p>Whether the initial processing weight for Integrated Facility for Linux (IFL) processors is a limit or a target.</p> <p><b>True:</b> Indicates that the initial Integrated Facility for Linux (IFL) processor processing weight for the logical partition is capped. It represents the logical partition's maximum share of Integrated Facility for Linux (IFL) processor resources, regardless of the availability of excess Integrated Facility for Linux (IFL) processor resources.</p> <p><b>False:</b> Indicates that the initial Integrated Facility for Linux (IFL) processor processing weight for the logical partition is not capped. It represents the share of Integrated Facility for Linux (IFL) processor resources guaranteed to a logical partition when all Integrated Facility for Linux (IFL) processor resources are in use. Otherwise, when excess Integrated Facility for Linux (IFL) processor resources are available, the logical partition can use them if necessary.</p>
<b>minimum-ift-processing-weight</b> <sup>7</sup>	(w)	Integer	<p>The minimum relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Integrated Facility for Linux (IFL) processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared Integrated Facility for Linux (IFL) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>0</b> There is no minimum value for the processing weight.</p> <p><b>1-999</b> Define the minimum relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p>

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>maximum-ifl-processing-weight<sup>7</sup></b>	(w)	Integer	<p>The maximum relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation profile does not represent a logical partition with at least one shared Integrated Facility for Linux (IFL) processor.</p> <p><b>1-999</b> Represents the maximum relative amount of shared Integrated Facility for Linux (IFL) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the maximum relative amount of shared Integrated Facility for Linux (IFL) processor resources allocated to the logical partition.</p>
<b>absolute-ifl-capping</b>	(w)	absolute-capping object	<p>The amount of absolute capping applied to the Integrated Facility for Linux (IFL) processor.</p> <p><b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.</p>
<b>initial-internal-cf-processing-weight<sup>9</sup></b>	(w)	Integer	<p>The relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the logical partition at activation.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Internal Coupling Facility (ICF) processor.</p> <p><b>1-999</b> Represents the relative amount of shared Internal Coupling Facility (ICF) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the logical partition.</p>

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>initial-internal-cf-processing-weight-capped</b> <sup>3, 9, 10</sup>	(w)	Boolean	<p>Indicates whether the initial processing weight for Internal Coupling Facility (ICF) processors is a limit or a target.</p> <p><b>True:</b> Indicates that the initial Internal Coupling Facility (ICF) processor processing weight for the Image associated with the logical partition is capped. It represents the logical partition's maximum share of Internal Coupling Facility (ICF) processor resources, regardless of the availability of excess Internal Coupling Facility (ICF) processor resources.</p> <p><b>False:</b> Indicates that the initial Internal Coupling Facility (ICF) processor processing weight for the logical partition is not capped. It represents the share of Internal Coupling Facility (ICF) processor resources guaranteed to a logical partition when all Internal Coupling Facility (ICF) processor resources are in use. Otherwise, when excess Internal Coupling Facility (ICF) processor resources are available, the logical partition can use them if necessary.</p>
<b>minimum-internal-cf-processing-weight</b> <sup>9</sup>	(w)	Integer	<p>The minimum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the logical partition at activation.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Internal Coupling Facility (ICF) processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared Internal Coupling Facility (ICF) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the minimum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the logical partition.</p>

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>maximum-internal-cf-processing-weight<sup>9</sup></b>	(w)	Integer	<p>The maximum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Internal Coupling Facility (ICF) processor</p> <p><b>1-999</b> Represents the maximum relative amount of shared Internal Coupling Facility (ICF) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the maximum relative amount of shared Internal Coupling Facility (ICF) processor resources allocated to the logical partition.</p>
<b>absolute-icf-capping</b>	(w)	absolute-capping object	<p>The amount of absolute capping applied to the Internal Coupling Facility (ICF) processor.</p> <p><b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.</p>
<b>initial-ziip-processing-weight<sup>11</sup></b>	(w)	Integer	<p>The relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition at activation.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Integrated Information Processors (zIIP) processor.</p> <p><b>1-999</b> Represents the relative amount of shared Integrated Information Processors (zIIP) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p>



Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>initial-ziip-processing-weight-capped</b> <sup>3, 11, 12</sup>	(w)	Boolean	<p>Whether the initial processing weight for Integrated Information Processors (zIIP) processors is a limit or a target.</p> <p><b>True:</b> Indicates that the initial Integrated Information Processors (zIIP) processor processing weight for the logical partition is capped. It represents the logical partition's maximum share of Integrated Information Processors (zIIP) processor resources, regardless of the availability of excess Integrated Information Processors (zIIP) processor resources.</p> <p><b>False:</b> Indicates that the initial Integrated Information Processors (zIIP) processor processing weight for the logical partition is not capped. It represents the share of Integrated Information Processors (zIIP) processor resources guaranteed to a logical partition when all Integrated Information Processors (zIIP) processor resources are in use. Otherwise, when excess Integrated Information Processors (zIIP) processor resources are available, the logical partition can use them if necessary.</p>
<b>minimum-ziip-processing-weight</b> <sup>11</sup>	(w)	Integer	<p>The minimum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Integrated Information Processors (zIIP) processor.</p> <p><b>1-999</b> Represents the minimum relative amount of shared Integrated Information Processors (zIIP) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>0</b> There is no minimum value for the processing weight.</p> <p><b>1-999</b> Define the minimum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p>

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>maximum-ziip-processing-weight<sup>11</sup></b>	(w)	Integer	<p>The maximum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p> <p>Get:</p> <p><b>0</b> The Image Activation Profile does not represent a logical partition with at least one shared Integrated Information Processors (zIIP) processor.</p> <p><b>1-999</b> Represents the maximum relative amount of shared Integrated Information Processors (zIIP) processor resources initially allocated to the logical partition.</p> <p>Update:</p> <p><b>1-999</b> Define the maximum relative amount of shared Integrated Information Processors (zIIP) processor resources allocated to the logical partition.</p>
<b>absolute-ziip-capping</b>	(w)	absolute-capping object	<p>The amount of absolute capping applied to the Integrated Information Processors (zIIP) processor.</p> <p><b>Note:</b> Absolute capping does not apply to image profiles where the processors are dedicated to the partition. Absolute capping only applies to partitions using shared processors.</p>
<b>group-profile-uri</b>	(w)	String/ URI	<p>The canonical URI of the Group profile to be used for the logical partition activated by this profile, which provides the group capacity value. On a Get, a null object is returned if no Group profile is associated with this activation profile. On an Update, a null object indicates that no Group profile is to be associated with this activation profile.</p>
<b>load-at-activation</b>	(w)	Boolean	<p>If true, the logical partition will be loaded at the end of the activation.</p>
<b>central-storage</b>	(w)	Integer	<p>Defines the amount of central storage, measured in megabytes (MB), to be allocated for the logical partition's exclusive use at activation. This value must be a multiple of the storage granularity value.</p>
<b>reserved-central-storage</b>	(w)	Integer	<p>Defines the amount of central storage, measured in megabytes (MB), dynamically reconfigurable to the logical partition after activation. This value must be a multiple of the storage granularity value.</p>
<b>expanded-storage</b>	(w)	Integer	<p>Defines the amount of expanded storage, measured in megabytes (MB), to be allocated for the logical partition's exclusive use at activation. This value must be a multiple of the storage granularity value.</p>
<b>reserved-expanded-storage</b>	(w)	Integer	<p>Defines the amount of expanded storage, measured in megabytes (MB), dynamically reconfigurable to the logical partition after activation. This value must be a multiple of the storage granularity value.</p>

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>processor-usage</b>	(w)	String Enum	Defines how processors are allocated to the logical partition. One of the following values: <ul style="list-style-type: none"> <li>• <b>"dedicated"</b> - all processor types in the logical partition are to be exclusively available to this specific logical partition.</li> <li>• <b>"shared"</b> - all processors types in the logical partition are to be shareable across logical partitions.</li> </ul>
<b>number-dedicated-general-purpose-processors<sup>13</sup></b>	(w)	Integer	Defines the number of general purpose processors to be allocated for the logical partition's exclusive use at activation.
<b>number-reserved-dedicated-general-purpose-processors<sup>13</sup></b>	(w)	Integer	Defines the number of dedicated general purpose processors to be reserved for the logical partition, which can be dynamically configured after activation.
<b>number-dedicated-aap-processors<sup>13</sup></b>	(w)	Integer	Defines the number of Application Assist Processor (zAAP) processors to be allocated for the logical partition partition's exclusive use at activation.
<b>number-reserved-dedicated-aap-processors<sup>13</sup></b>	(w)	Integer	Defines the number of dedicated Application Assist Processor (zAAP) processors to be reserved for the logical partition, which can be dynamically configured after activation.
<b>number-dedicated-ifl-processors<sup>13</sup></b>	(w)	Integer	Defines the number of Integrated Facility for Linux (IFL) processors to be allocated for the logical partition partition's exclusive use at activation.
<b>number-reserved-dedicated-ifl-processors<sup>13</sup></b>	(w)	Integer	Defines the number of dedicated Integrated Facility for Linux (IFL) processors to be reserved for the logical partition, which can be dynamically configured after activation.
<b>number-dedicated-icf-processors<sup>13</sup></b>	(w)	Integer	Defines the number of Integrated Coupling Facility (ICF) processors to be allocated for the logical partition partition's exclusive use at activation.
<b>number-reserved-dedicated-icf-processors<sup>13</sup></b>	(w)	Integer	Defines the number of dedicated Integrated Coupling Facility (ICF) processors to be reserved for the logical partition, which can be dynamically configured after activation.
<b>number-dedicated-ziip-processors<sup>13</sup></b>	(w)	Integer	Defines the number of Integrated Information Processors (zIIP) processors to be allocated for the logical partition partition's exclusive use at activation.
<b>number-reserved-dedicated-ziip-processors<sup>13</sup></b>	(w)	Integer	Defines the number of dedicated Integrated Information Processors (zIIP) processors to be reserved for the logical partition, which can be dynamically configured after activation.
<b>number-shared-general-purpose-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared general purpose processors to be allocated for the logical partition at activation.
<b>number-reserved-shared-general-purpose-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared general purpose processors to be reserved for the logical partition, which can be dynamically configured after activation.
<b>number-shared-aap-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared Application Assist Processor (zAAP) processors to be allocated for the logical partition at activation.
<b>number-reserved-shared-aap-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared Application Assist Processor (zAAP) processors to be reserved for the logical partition, which can be dynamically configured after activation.

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>number-shared-ifl-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared Integrated Facility for Linux (IFL) processors to be allocated for the logical partition at activation.
<b>number-reserved-shared-ifl-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared Integrated Facility for Linux (IFL) processors to be reserved for the logical partition, which can be dynamically configured after activation.
<b>number-shared-icf-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared Integrated Coupling Facility (ICF) processors to be allocated for the logical partition at activation.
<b>number-reserved-shared-icf-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared Integrated Coupling Facility (ICF) processors to be reserved for the logical partition, which can be dynamically configured after activation.
<b>number-shared-ziip-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared Integrated Information Processors (ziIP) processors to be allocated for the logical partition at activation.
<b>number-reserved-shared-ziip-processors<sup>14</sup></b>	(w)	Integer	Defines the number of shared Integrated Information Processors (ziIP) processors to be reserved for the logical partition, which can be dynamically configured after activation.
<b>basic-cpu-counter-authorization-control</b>	(w)	Boolean	If true, the basic CPU counter facility for the logical partition is enabled.
<b>problem-state-cpu-counter-authorization-control</b>	(w)	Boolean	If true, the problem state CPU counter facility for the logical partition is enabled.
<b>crypto-activity-cpu-counter-authorization-control</b>	(w)	Boolean	If true, the crypto activity CPU counter facility for the logical partition is enabled.
<b>extended-cpu-counter-authorization-control</b>	(w)	Boolean	If true, the extended CPU counter facility for the logical partition is enabled.
<b>coprocessor-cpu-counter-authorization-control</b>	(w)	Boolean	If true, the coprocessor group CPU counter facility for the logical partition is enabled.
<b>basic-cpu-sampling-authorization-control</b>	(w)	Boolean	If true, the basic CPU sampling facility for the logical partition is enabled.
<b>permit-aes-key-import-functions</b>	(w)	Boolean	If true, importing of AES keys for the logical partition is enabled.
<b>permit-des-key-import-functions</b>	(w)	Boolean	If true, importing of DES keys for the logical partition is enabled.
<b>liccc-validation-enabled</b>	(w)	Boolean	If true, ensure that the image profile data conforms to the current maximum Licensed Internal Code Configuration Control (LICCC) configuration.
<b>global-performance-data-authorization-control</b>	(w)	Boolean	If true, the logical partition can be used to view the processing unit activity data for all other logical partitions activated on the same CPC.
<b>io-configuration-authorization-control</b>	(w)	Boolean	If true, the logical partition can be used to read and write any Input/Output Configuration Data Set (IOCDs) in the configuration.
<b>cross-partition-authority-authorization-control</b>	—	Boolean	If true, the logical partition can be used to issue control program instructions that reset or deactivate other logical partitions.

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>logical-partition-isolation-control</b>	(w)	Boolean	If true, reconfigurable channel paths assigned to the logical partition are reserved for its exclusive use.
<b>operating-mode</b>	(w)	String Enum	The operating mode for the logical partition: <ul style="list-style-type: none"> <li>• "esa390"</li> <li>• "esa390-tpf"</li> <li>• "coupling-facility"</li> <li>• "linux-only"</li> <li>• "zvm"</li> <li>• "zaware"</li> </ul>
<b>clock-type</b>	(w)	String Enum	One of: <ul style="list-style-type: none"> <li>• "standard" – Set the logical partition's clock is set to the same time set for the CPC's time source.</li> <li>• "lpar" - Set the logical partition's clock using an offset from the External Time Source's time of day.</li> </ul>
<b>time-offset-days</b>	(w)	Integer (0-999)	The number of days the logical partition's clock is to be offset from the External Time Source's time of day.
<b>time-offset-hours</b>	(w)	Integer (0-23)	The number of hours the logical partition's clock is to be offset from the External Time Source's time of day.
<b>time-offset-minutes</b>	(w)	Integer Enum	The number of minutes the logical partition's clock is to be offset from the External Time Source's time of day. Allowable values are 0, 15, 30 or 45.
<b>time-offset-increase-decrease</b>	(w)	String Enum	One of: <ul style="list-style-type: none"> <li>• "increase" – Set the logical partition's clock ahead of the External Time Source's time of day.</li> <li>• "decrease" – Set the logical partition's clock back from the External Time Source's time of day.</li> </ul>
<b>zaware-host-name<sup>16</sup></b>	(w)	String (1-64)	The IBM zAware host name. Valid characters are: a-z,A-Z,0-9, period(.), minus(-) and colon(:)
<b>zaware-master-userid<sup>16</sup></b>	(w)	String (1-32)	The IBM zAware master userid. Valid characters are: a-z,A-Z,0-9, period(.), minus(-) and underscore (_)
<b>zaware-master-pw<sup>16</sup></b>	(w)	String (8-256)	The IBM zAware master password. Valid characters are: a-z,A-Z,0-9 and !@#\$%^&*()_+{ <>?=-  This property is not returned on a Get request, it can only be specified on an Update request.
<b>zaware-network-info<sup>16</sup></b>	(w)	Array of zaware-network objects	The set of networks available to IBM zAware. A maximum of 100 networks are permitted.  On an Update request, this property fully replaces the existing set.
<b>zaware-gateway-info<sup>16</sup></b>	(w)	ip-info object	The default gateway IP address information.
<b>zaware-dns-info<sup>16</sup></b>	(w)	Array of ip-info objects	The DNS IP address information. A minimum of 0 entries and a maximum of 2 entries are permitted.

Table 233. Image activation profile: properties (continued).

Name	Qualifier	Type	Description
<b>Notes:</b>			
1. An Update of this property is only valid for an Image Activation Profile that represents a logical partition with at least one shared general purpose processor.			
2. The value returned for a Get request is always false when the Image Activation Profile does not represent a logical partition or the Image Activation Profile does not represent a logical partition with at least one shared general purpose processor.			
3. This property and the <b>workload-manager-enabled</b> property are mutually exclusive and cannot both be enabled at the same time. Therefore in order to enable this property it might be necessary to first disable the <b>workload-manager-enabled</b> property.			
4. This property and the various capping properties are mutually exclusive and cannot be enabled at the same time. Therefore in order to enable this property it may be necessary to first disable any capping property that is currently enabled.			
5. An Update of this property is only valid for an Image Activation Profile that represents a logical partition with at least one shared Application Assist Processor (zAAP) processor.			
6. The value returned for a Get request is always false when the Image Activation Profile does not represent a logical partition with at least one shared Application Assist Processor (zAAP) processor.			
7. An Update of this property is only valid for an Image Activation Profile that represents a logical partition with at least one shared Integrated Facility for Linux (IFL) processor.			
8. The value returned for a Get request is always false when the Image Activation Profile does not represent a logical partition with at least one shared Integrated Facility for Linux (IFL) processor.			
9. An Update of this property is only valid for an Image Activation Profile that represents a logical partition with at least one shared Internal Coupling Facility (ICF) processor.			
10. The value returned for a Get request is always false when the Image Activation Profile does not represent a logical partition with at least one shared Internal Coupling Facility (ICF) processor.			
11. An Update of this property is only valid for an Image Activation Profile that represents a logical partition with at least one shared Integrated Information Processors (zIIP) processor.			
12. The value returned for a GET request is always false when the Image Activation profile does not represent a logical partition with at least one shared Integrated Information Processors (zIIP) processor.			
13. The value of this property is a null object if the processor-usage property is " <b>shared</b> "			
14. The value of this property is a null object if the processor-usage property is " <b>dedicated</b> "			
15. An Update request accepts any mixture of [a-f,A-F,0-9], however the original string value is not saved and a subsequent Get request may not return the exact same set of lower/upper case letters.			
16. On a Get request, this property is returned only when <b>activation-mode</b> is " <b>zaware</b> ". On an Update request, this property can be specified only when <b>activation-mode</b> is " <b>zaware</b> ".			

## List Image Activation Profiles

The **List Image Activation Profiles** operation lists the Image Activation Profiles for the associated CPC object.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/image-activation-profiles

In this request, the URI variable {cpc-id} is the object ID of the target CPC object.

### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	A regular expression used to limit returned objects to those that have a matching name property. If matches are found, the response will be an array with all objects that match. If no match is found, the response will be an empty array.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
image-activation-profiles	Array of image-actprof-info objects	Array of nested objects (described in the following table).

Each image-actprof-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Image Activation Profile object.
name	String	The name of the Image Activation Profile.

## Description

This operation lists the Image Activation Profiles associated with a particular CPC.

If the **name** query parameter is specified, the returned list is limited to those Image Activation Profiles that have a name property matching the specified filter pattern. If the **name** parameter is omitted, this filtering is not done.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	299	A query parameter has an invalid syntax.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
500 (Server Error)	281	An unexpected error occurred during the collection of the list of activation profiles.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/image-activation-profiles HTTP/1.1
x-api-session: 5obf0hwsfv1sg9kr5f93cph3zt6o5cptb61c1538wuyebdyzu4
```

---

Figure 424. List Image Activation Profiles: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 17:16:18 GMT
content-type: application/json;charset=UTF-8
content-length: 506
{
  "image-activation-profiles": [
    {
      "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/image-activation-profiles/
      APIVM1",
      "name": "APIVM1"
    },
    {
      "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/image-activation-profiles/
      DEFAULT",
      "name": "DEFAULT"
    },
    {
      "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/image-activation-profiles/
      ZOS1",
      "name": "ZOS1"
    }
  ]
}
```

---

Figure 425. List Image Activation Profiles: Response

## Get Image Activation Profile Properties

The **Get Image Activation Profile Properties** operation retrieves the properties of a single Image Activation Profile designated by *{image-activation-profile-name}*.

### HTTP method and URI

```
GET /api/cpcs/{cpc-id}/image-activation-profiles/{image-activation-profile-name}
```

#### URI variables:

Variable	Description
<i>{cpc-id}</i>	Object ID of the target CPC object.
<i>{image-activation-profile-name}</i>	Image Activation Profile name



## Response body contents

On successful completion, the response body provides the current values of the properties for the Image Activation Profile as defined in the “Data model” on page 839.

### Description

The URI path must designate an existing Image Activation Profile and the API user must have object-access permission to the associated CPC object. If either of these conditions is not met, HTTP status code 404 (Not Found) is returned.

On successful execution, HTTP status code 200 (OK) is returned and the response body contains all of the current properties as defined by the Data Model for the Image Activation Profile object.

### Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	260	The activation profile name in the URI ( <i>{image-activation-profile-name}</i> ) does not designate an existing activation profile.
500 (Server Error)	281	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/image-activation-profiles/Z0S HTTP/1.1
x-api-session: 5obf0hwsfv1sg9kr5f93cph3zt6o5cptb61cl538wuyebdyzu4
```

---

Figure 426. Get Image Activation Profile Properties: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 17:16:18 GMT
content-type: application/json;charset=UTF-8
content-length: 3352
{
  "basic-cpu-counter-authorization-control": false,
  "basic-cpu-sampling-authorization-control": false,
  "boot-record-lba": "0",
  "central-storage": 4096,
  "class": "image-activation-profile",
  "clock-type": "standard",
  "coprocessor-cpu-counter-authorization-control": false,
  "cross-partition-authority-authorization-control": false,
  "crypto-activity-cpu-counter-authorization-control": false,
  "defined-capacity": 0,
  "description": "This is the ZOS Image profile.",
  "disk-partition-id": 0,
  "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/image-activation-profiles/ZOS",
  "expanded-storage": 0,
  "extended-cpu-counter-authorization-control": false,
  "global-performance-data-authorization-control": true,
  "group-profile-uri": null,
  "initial-aap-processing-weight": 0,
  "initial-aap-processing-weight-capped": false,
  "initial-ifl-processing-weight": 0,
  "initial-ifl-processing-weight-capped": false,
  "initial-internal-cf-processing-weight": 0,
  "initial-internal-cf-processing-weight-capped": false,
  "initial-processing-weight": 44,
  "initial-processing-weight-capped": false,
  "initial-ziip-processing-weight": 0,
  "initial-ziip-processing-weight-capped": false,
  "io-configuration-authorization-control": true,
  "ipl-address": "00000",
  "ipl-parameter": "    ",
  "ipl-type": "ipltype-standard",
  "liccc-validation-enabled": true,
  "load-at-activation": false,
  "logical-partition-isolation-control": false,
  "logical-unit-number": "0",
  "maximum-aap-processing-weight": 0,
  "maximum-ifl-processing-weight": 0,
  "maximum-internal-cf-processing-weight": 0,
  "maximum-processing-weight": 44,
  "absolute-aap-capping": {"type": "processors", "value": 67.95},
  "absolute-icf-capping": {"type": "none"},
  "absolute-ifl-capping": {"type": "none"},
  "absolute-general-purpose-capping": {"type": "processors", "value": 3.03},
  "absolute-ziip-capping": {"type": "processors", "value": 95.95},
```

---

Figure 427. Get Image Activation Profile Properties: Response (Part 1)

```

"maximum-ziip-processing-weight": 0,
"minimum-aap-processing-weight": 0,
"minimum-ifl-processing-weight": 0,
"minimum-internal-cf-processing-weight": 0,
"minimum-processing-weight": 44,
"minimum-ziip-processing-weight": 0,
"name": "ZOS",
"number-dedicated-aap-processors": null,
"number-dedicated-general-purpose-processors": null,
"number-dedicated-icf-processors": null,
"number-dedicated-ifl-processors": null,
"number-dedicated-ziip-processors": null,
"number-reserved-dedicated-aap-processors": null,
"number-reserved-dedicated-general-purpose-processors": null,
"number-reserved-dedicated-icf-processors": null,
"number-reserved-dedicated-ifl-processors": null,
"number-reserved-dedicated-ziip-processors": null,
"number-reserved-shared-aap-processors": 0,
"number-reserved-shared-general-purpose-processors": 0,
"number-reserved-shared-icf-processors": 0,
"number-reserved-shared-ifl-processors": 0,
"number-reserved-shared-ziip-processors": 0,
"number-shared-aap-processors": 0,
"number-shared-general-purpose-processors": 1,
"number-shared-icf-processors": 0,
"number-shared-ifl-processors": 0,
"number-shared-ziip-processors": 0,
"operating-mode": "esa390",
"os-specific-load-parameters": "
",
"parent": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
"permit-aes-key-import-functions": true,
"permit-des-key-import-functions": true,
"problem-state-cpu-counter-authorization-control": false,
"processor-usage": "shared",
"reserved-central-storage": 0,
"reserved-expanded-storage": 0,
"time-offset-days": 0,
"time-offset-hours": 0,
"time-offset-increase-decrease": "decrease",
"time-offset-minutes": 0,
"workload-manager-enabled": false,
"worldwide-port-name": "0"
}

```

Figure 428. Get Image Activation Profile Properties: Response (Part 2)

## Update Image Activation Profile Properties

The **Update Image Activation Profile Properties** operation updates one or more writeable properties of the Image Activation Profile designated by *{image-activation-profile-name}*.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/image-activation-profiles/{image-activation-profile-name}

URI variables:

Variable	Description
<i>{cpc-id}</i>	Object ID of the target CPC object.

Variable	Description
<i>{image-activation-profile-name}</i>	Image Activation Profile name

## Request body contents

The request body is expected to contain one or more field names representing writable Image Activation Profile properties, along with the new values for those fields.

The request body can and should omit fields for properties whose values are not to be changed by this operation. Properties for which no input value is provided remain unchanged by this operation.

## Description

The request body object is validated against the data model for the Image Activation Profile to ensure that the request body contains only writeable properties and the data types of those properties are as required. If the request body is not valid, HTTP status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

On successful execution, the value of each corresponding property of the Image Activation Profile is updated with the value provided by the input field, and HTTP status code 204 (No Content) is returned.

When this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission for the Customize/Delete Activation Profiles task.

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See "Common request validation reason codes" on page 24 for a list of the possible reason codes.
	300	The provided update values would result in an illegal state. Verify that the values are both internally consistent and consistent with the current state of the profile.
	306	The provided update values are not valid for the current operating-mode, the target image activation profile's operating-mode is not "zaware" or was not specified as "zaware" in the request body.

HTTP error status code	Reason code	Description
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	2	A URI in the request body does not designate an existing resource of the expected type, or designates a resource for which the API user does not have object-access permission.
	260	The activation profile name in the URI ( <i>{image-activation-profile-name}</i> ) does not designate an existing activation profile.
409 (Conflict)	2	The operation was rejected by the Support Element (SE), because the SE is currently performing processing that requires exclusive control of the SE. Retry the operation at a later time.
500 (Server Error)	281	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Inventory service data

Information about image activation profiles can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for the Image Activation Profile objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "cpc" are to be included. An entry for a particular image activation profile is included only if the API user has access permission to that object as described in the **Get Image Activation Profile Properties** operation.

For each Image Activation Profile object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for “Get Image Activation Profile Properties” on page 856. That is, the data provided is the same as would be provided if a **Get Image Activation Profile Properties** operation were requested targeting this object.

---

## Load activation profile

A Load activation profile is used to load a previously activated logical partition with a control program or operating system.

An activation profile can only be created or deleted from the Hardware Management Console or the Support Element.

For information on customizing activation profiles, Support Element (Version 2.12.1 and newer) information can be found on the console help system, or on the IBM Knowledge Center at <https://www.ibm.com/support/knowledgecenter>. (Select **z Systems** on the navigation bar, and then select your server). For information from earlier versions of the Support Element, see the *Support Element Operations Guide*.

## Data model

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

This element includes the following properties.

Table 234. Load activation profile: properties

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/ URI	The canonical URI path of the Load Activation Profile object, of the form <code>/api/cpcs/{cpc-id}/load-activation-profiles/{load-activation-profile-name}</code> where <code>{load-activation-profile-name}</code> is the value of the name property (Load Activation Profile name).
<b>parent</b>	—	String/ URI	The canonical URI path of the associated CPC object
<b>class</b>	—	String	The class of a Load Activation Profile object is <b>"load-activation-profile"</b> .
<b>name</b>	—	String (1-16)	The activation profile name, which uniquely identifies this profile within the set of activation profiles for the CPC object designated by <code>{cpc-id}</code> .
<b>description</b>	(w)	String (1-50)	The load profile description
<b>ipl-address</b>	(w)	String (0-5)	The hexadecimal address of an I/O device that provides access to the control program to be loaded. The input value will be right justified and padded with zeros to 5 characters. An empty string indicates that the value for this property is to be retrieved from the IOCDS used during a subsequent Load operation.  Valid values are in the range "00000" to "nFFFF" where "n" is the number of subchannel sets provided by the CPC minus 1. So, for example, on a CPC that provides 3 subchannel sets, the valid range is "00000" to "2FFFF".
<b>ipl-parameter<sup>1</sup></b>	(w)	String (0-8)	Some control programs support the use of this property to provide additional control over the outcome of a Load operation. Refer to the configuration documentation for the control program to be loaded to see if this parameter is supported and if so, what values and format is supported. An empty string indicates that the value for this property is to be retrieved from the IOCDS used during a subsequent Load operation. On an Update, a non-empty string is left justified and right padded with blanks to 8 characters.
<b>ipl-type</b>	(w)	String Enum	One of: <ul style="list-style-type: none"> <li>• <b>"ipltype-standard"</b> - Associated image activation profile is used to perform a standard load.</li> <li>• <b>"ipltype-scsi"</b> - Associated image activation profile is used to perform a SCSI load.</li> <li>• <b>"ipltype-scsidump"</b> - Associated image activation profile is used to perform a SCSI dump.</li> </ul>
<b>worldwide-port-name<sup>1</sup></b>	(w)	String (1-16)	Worldwide port name value for the activation profile, used for SCIS load and SCSI dump, in hexadecimal.
<b>disk-partition-id</b>	(w)	Integer (0-30)	Disk partition number (also called the boot program selector) for the activation profile, used for SCIS load and SCSI dump.
<b>logical-unit-number<sup>1</sup></b>	(w)	String (1-16)	Logical unit number value for the activation profile, used for SCIS load and SCSI dump, in hexadecimal.
<b>boot-record-lba<sup>1</sup></b>	(w)	String (1-16)	Boot record logical block address for the activation profile, used for SCIS load and SCSI dump, in hexadecimal.
<b>os-specific-load-parameters</b>	(w)	String (0-256)	Operating system-specific load parameters for the activation profile, used for SCIS load and SCSI dump.

Table 234. Load activation profile: properties (continued)

Name	Qualifier	Type	Description
<b>clear-indicator</b>	(w)	Boolean	Whether memory should be cleared before performing the Load (true) or not cleared (false). The default is to clear memory before performing the Load. This property cannot be set to false when the ipl-type is SCSI load or SCSI dump.
<b>store-status-indicator</b>	(w)	Boolean	Whether the store status function should be invoked before performing the Load (true) or not (false). The default is not to store status before performing the Load. The store status function stores the current values of the processing unit timer, the clock comparator, the program status word, and the contents of the processor registers in their assigned absolute storage locations. This property cannot be set to true when the ipl-type is SCSI load or SCSI dump, or when the ipl-type is " <b>ipltype-standard</b> " and clear-indicator is true.

1. An Update request accepts any mixture of [a-f, A-F, 0-9], however the original string value is not saved and a subsequent Get request may not return the exact same set of lower/upper case letters.

## List Load Activation Profiles

The **List Load Activation Profiles** operation lists the Load Activation Profiles for the associated CPC object.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/load-activation-profiles

In this request, the URI variable {cpc-id} is the object ID of the target CPC object.

#### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	A regular expression used to limit returned objects to those that have a matching name property. If matches are found, the response will be an array with all objects that match. If no match is found, the response will be an empty array.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
load-activation-profiles	Array of load-actprof-info objects	Array of nested objects (described in the next table).

Each load-actprof-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Load Activation Profile object.
name	String	The name of the Load Activation Profile.

## Description

This operation lists the Load Activation Profiles for the associated CPC object.

If the name query parameter is specified, the returned list is limited to those Load Activation Profiles that have a name property matching the specified filter pattern. If the name parameter is omitted, this filtering is not done.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 863.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 863.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	299	A query parameter has an invalid syntax.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
500 (Server Error)	281	An unexpected error occurred during the collection of the list of activation profiles.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/load-activation-profiles HTTP/1.1
x-api-session: 5obf0hwsfv1sg9kr5f93cph3zt6o5cptb61cl538wuyebdyzu4
```

---

Figure 429. List Load Activation Profiles: Request



---

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 17:16:19 GMT
content-type: application/json;charset=UTF-8
content-length: 363
{
  "load-activation-profiles": [
    {
      "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/load-activation-profiles/
DEFAULTLOAD",
      "name": "DEFAULTLOAD"
    },
    {
      "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/load-activation-profiles/
MODIFYL",
      "name": "MODIFYL"
    }
  ]
}

```

---

Figure 430. List Load Activation Profiles: Response

## Get Load Activation Profile Properties

The **Get Load Activation Profile Properties** operation retrieves the properties of a single Load Activation Profile designated by *{load-activation-profile-name}*.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/load-activation-profiles/{load-activation-profile-name}

#### URI variables

Variable	Description
<i>{cpc-id}</i>	Object ID of the target CPC object.
<i>{load-activation-profile-name}</i>	Load Activation Profile name.

### Response body contents

On successful completion, the response body provides the current values of the properties for the Load Activation Profile as defined in the “Data model” on page 861.

### Description

The URI path must designate an existing Load Activation Profile and the API user must have object-access permission to the associated CPC object. If either of these conditions is not met, status code 404 (Not Found) is returned.

On successful execution, HTTP status code 200 (OK) is returned and the response body contains all of the current properties as defined by the Data Model for Load Activation Profiles.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 865.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	260	The activation profile name in the URI ( <i>{load-activation-profile-name}</i> ) does not designate an existing activation profile.
500 (Server Error)	281	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/load-activation-profiles/DEFAULTLOAD HTTP/1.1
x-api-session: 5obf0hwsfv1sg9kr5f93cph3zt6o5cptb61cl538wuyebdyzu4
```

---

Figure 431. Get Load Activation Profile Properties: Request

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 17:16:19 GMT
content-type: application/json;charset=UTF-8
content-length: 793
{
  "boot-record-lba": "abcdef0123456789",
  "class": "load-activation-profile",
  "clear-indicator": true,
  "description": "This is the default Load profile.",
  "disk-partition-id": 0,
  "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/load-activation-profiles/
  DEFAULTLOAD",
  "ipl-address": "00D00",
  "ipl-parameter": " ",
  "ipl-type": "ipltype-scsi",
  "logical-unit-number": "0",
  "name": "DEFAULTLOAD",
  "os-specific-load-parameters": " ",
  "parent": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340",
  "store-status-indicator": false,
  "worldwide-port-name": "0"
}

```

Figure 432. Get Load Activation Profile Properties: Response

## Update Load Activation Profile Properties

The **Update Load Activation Profile Properties** operation updates one or more writeable properties of the Load Activation Profile designated by *{load-activation-profile-name}*.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/load-activation-profiles/{load-activation-profile-name}

#### URI variables

Variable	Description
<i>{cpc-id}</i>	Object ID of the target CPC object.
<i>{load-activation-profile-name}</i>	Load Activation Profile name.

### Request body contents

The request body is expected to contain one or more field names representing writable Load Activation Profile properties, along with the new values for those fields.

The request body can and should omit fields for properties whose values are not to be changed by this operation. Properties for which no input value is provided remain unchanged by this operation.

## Description

The request body object is validated against the data model for the Load Activation Profile to ensure that the request body contains only writeable properties and the data types of those properties are as required. If the request body is not valid, HTTP status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

On successful execution, the value of each corresponding property of the Load Activation Profile is updated with the value provided by the input field, and HTTP status code 204 (No Content) is returned.

When this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission for the **Customize/Delete Activation Profiles** task

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	300	The provided update values would result in an illegal state. Verify that the values are both internally consistent and consistent with the current state of the profile.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	260	The activation profile name in the URI ( <i>{load-activation-profile-name}</i> ) does not designate an existing activation profile.
409 (Conflict)	2	The operation was rejected by the Support Element (SE), because the SE is currently performing processing that requires exclusive control of the SE. Retry the operation at a later time.
500 (Server Error)	281	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## | Inventory service data

| Information about load activation profiles can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for the Load Activation Profile objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "cpc" are to be included. An entry for a particular load activation profile is included only if the API user has access permission to that object as described in the **Get Load Activation Profile Properties** operation.

For each Load Activation Profile object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for "Get Load Activation Profile Properties" on page 865. That is, the data provided is the same as would be provided if a **Get Load Activation Profile Properties** operation were requested targeting this object.

---

## Group profile

A Group profile is used to define the group capacity value for all logical partitions belonging to that group.

A logical partition becomes a member of a group profile by placing the group profile's URI in the image activation profile used to activate the logical partition.

A group profile can only be created or deleted from the Hardware Management Console or the Support Element.

For information on customizing activation profiles, Support Element (Version 2.12.1 and newer) information can be found on the console help system, or on the IBM Knowledge Center at <https://www.ibm.com/support/knowledgecenter>. (Select **z Systems** on the navigation bar, and then select your server). For information from earlier versions of the Support Element, see the *Support Element Operations Guide*.

## Data model

For definitions of the qualifier abbreviations in the following tables, see "Property characteristics" on page 38.

This element includes the following properties.

Table 235. Group profile: properties

Name	Qualifier	Type	Description
<b>element-uri</b>	—	String/ URI	The canonical URI path of the group profile object, of the form <code>/api/cpcs/{cpc-id}/group-profiles/{group-profile-name}</code> where <code>{group-profile-name}</code> is the value of the name property (group profile name).
<b>parent</b>	—	String/ URI	The canonical URI path of the associated CPC object
<b>class</b>	—	String	The class of a Group Profile object is " <b>group-profile</b> ".
<b>name</b>	—	String (1-16)	The group profile name, which uniquely identifies this profile within the set of activation profiles for the CPC object designated by <code>{cpc-id}</code>
<b>description</b>	(w)	String (1-50)	The group profile description
<b>capacity</b>	(w)	Integer	The upper bound, in MSUs, beyond which the rolling 4-hour average CPU utilization cannot exceed for the group. A value of 0 indicates the setting is unused.

## List Group Profiles

The **List Group Profiles** operation lists the Group Profiles for the associated CPC object.

## HTTP method and URI

GET /api/cpcs/{cpc-id}/group-profiles

In this request, the URI variable *{cpc-id}* is the object ID of the target CPC object.

### Query parameters:

Name	Type	Rqd/Opt	Description
name	String	Optional	A regular expression used to limit returned objects to those that have a matching name property. If matches are found, the response will be an array with all objects that match. If no match is found, the response will be an empty array.

## Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
group-profiles	Array of group-actprof-info objects	Array of nested objects (described in the next table).

Each group-actprof-info object contains the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Group Profile object.
name	String	The name of the Group Profile.

## Description

This operation lists the Group Profiles for the associated CPC object.

If the name query parameter is specified, the returned list is limited to those Group Profiles that have a name property matching the specified filter pattern. If the name parameter is omitted, this filtering is not done.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in the Response Body Contents section.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	299	A query parameter has an invalid syntax.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
500 (Server Error)	281	An unexpected error occurred during the collection of the list of group profiles.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/group-profiles HTTP/1.1
x-api-session: 5obf0hwsfv1sg9kr5f93cph3zt6o5cptb61c1538wuyebdyzu4
```

---

Figure 433. List Group Profiles: Request

---

```
200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 17:16:19 GMT
content-type: application/json;charset=UTF-8
content-length: 182
{
  "group-profiles": [
    {
      "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/group-profiles/DEFAULT",
      "name": "DEFAULT"
    }
  ]
}
```

---

Figure 434. List Group Profiles: Response

## Get Group Profile Properties

The **Get Group Profile Properties** operation retrieves the properties of a single Group Profile designated by *{group-profile-name}*.

### HTTP method and URI

```
GET /api/cpcs/{cpc-id}/group-profiles/{group-profile-name}
```

#### URI variables

Variable	Description
<i>{cpc-id}</i>	Object ID of the target CPC object.
<i>{group-profile-name}</i>	Group Profile name.

## Response body contents

On successful completion, the response body provides the current values of the properties for the Group Profile as defined in the “Data model” on page 869.

## Description

The URI path must designate an existing Group Profile and the API user must have object-access permission to the associated CPC object. If either of these conditions is not met, HTTP status code 404 (Not Found) is returned.

On successful execution, HTTP status code 200 (OK) is returned and the response body contains all of the current properties as defined by the Data Model for the Group Profile.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents.”

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	260	The group profile name in the URI ( <i>{group-profile-name}</i> ) does not designate an existing group profile.
500 (Server Error)	281	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

---

```
GET /api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/group-profiles/DEFAULT HTTP/1.1
x-api-session: 5obf0hwsfv1sg9kr5f93cph3zt6o5cptb61c1538wuyebdyzu4
```

---

Figure 435. Get Group Profile Properties: Request



---

```

200 OK
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Fri, 25 Nov 2011 17:16:20 GMT
content-type: application/json;charset=UTF-8
content-length: 250
{
  "capacity": 0,
  "class": "group-profile",
  "description": "This is the default Group profile.",
  "element-uri": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340/group-profiles/DEFAULT",
  "name": "DEFAULT",
  "parent": "/api/cpcs/37c6f8a9-8d5e-3e5d-8466-be79e49dd340"
}

```

---

Figure 436. Get Group Profile Properties: Response

## Update Group Profile Properties

The **Update Group Profile Properties** operation updates one or more writeable properties of the Group Profile object designated by *{group-profile-name}*.

### HTTP method and URI

**POST** /api/cpcs/{cpc-id}/group-profiles/{group-profile-name}

#### URI variables

Variable	Description
<i>{cpc-id}</i>	Object ID of the target CPC object.
<i>{group-profile-name}</i>	Group Profile name.

### Request body contents

The request body is expected to contain one or more field names representing writable Group Profile properties, along with the new values for those fields.

The request body can and should omit fields for properties whose values are not to be changed by this operation. Properties for which no input value is provided remain unchanged by this operation.

### Description

The request body object is validated against the data model for the Group Profile to ensure that the request body contains only writeable properties and the data types of those properties are as required. If the request body is not valid, HTTP status code 400 (Bad Request) is returned with a reason code indicating the validation error encountered.

On successful execution, the value of each corresponding property of the Group Profile is updated with the value provided by the input field, and HTTP status code 204 (No Content) is returned.

When this operation changes the value of any property for which property-change notifications are due, those notifications are emitted asynchronously to this operation.

## Authorization requirements

This operation has the following authorization requirements:

- Object access permission to the CPC object designated by *{cpc-id}*
- Action/task permission for the **Customize/Delete Activation Profiles** task

## HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	300	The provided update values would result in an illegal state. Verify that the values are both internally consistent and consistent with the current state of the profile.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	260	The activation profile name in the URI ( <i>{group-profile-name}</i> ) does not designate an existing group profile.
409 (Conflict)	2	The operation was rejected by the Support Element (SE), because the SE is currently performing processing that requires exclusive control of the SE. Retry the operation at a later time.
500 (Server Error)	281	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Inventory service data

Information group profiles can be optionally included in the inventory data provided by the Inventory Service.

Inventory entries for the Group Profile objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "**cpc**" are to be included. An entry for a particular group profile is included only if the API user has access permission to that object as described in the **Get Group Profile Properties** operation.

For each Group Profile object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for “Get Group Profile Properties” on page 871. That is, the data provided is the same as would be provided if a **Get Group Profile Properties** operation were requested targeting this object.

## Capacity records

A capacity record represents a temporary upgrade that can be applied to a CPC.

These upgrades are provided through the following offerings:

- **On/Off Capacity on Demand (On/Off CoD)** - This offering allows you to temporarily add additional capacity or specialty engines due to seasonal activities, period-end requirements, peaks in workload, or application testing.
- **Capacity Backup (CBU)** - This offering allows you to replace model capacity or specialty engines to a backup server in the event of an unforeseen loss of server capacity because of an emergency.
- **Capacity for Planned Events (CPE)** - This offering allows you to replace model capacity or specialty engines due to a relocation of workload during system migrations or a data center move.

## Data model

For definitions of the qualifier abbreviations in the following tables, see “Property characteristics” on page 38.

This element includes the following type-specific properties.

Table 236. Capacity records: type-specific properties

Name	Type	Description
<b>element-uri</b>	String/ URI	The canonical URI path of the capacity record object, of the form <code>/api/cpcs/{cpc-id}/capacity-records/{capacity-record-id}</code> where <code>{cpc-id}</code> is the value of the <b>object-id</b> property of the CPC object and <code>{capacity-record-id}</code> is the value of the <b>record-identifier</b> property of the Capacity Record object.
<b>parent</b>	String/ URI	The canonical URI path for the associated CPC object
<b>class</b>	String	The class of a capacity record object is " <b>capacity-record</b> ".
<b>record-identifier</b>	String (1-8)	The identifier for the capacity record.
<b>record-type</b>	String Enum	The type of capacity record. One of: <ul style="list-style-type: none"> <li>• "<b>unknown</b>" - the record does not specify a record-type</li> <li>• "<b>cbu</b>" - a Capacity Backup Upgrade record</li> <li>• "<b>oocod</b>" - an On/Off Capacity on Demand record</li> <li>• "<b>planned-event</b>" - a Capacity for Planned Events record</li> <li>• "<b>loaner</b>" - resources loaned to the installation.</li> </ul>
<b>activation-status</b>	String Enum	An indication if any of the resources defined for the record are currently activated. One of: <ul style="list-style-type: none"> <li>• "<b>unknown</b>" - the activation status of the record is not known</li> <li>• "<b>not-activated</b>" - the record is not currently active</li> <li>• "<b>real</b>" - the record is either active or pending activation, via an Add Temporary Capacity operation with a <code>test=false</code> option</li> <li>• "<b>test</b>" - the record is either active or pending activation via an Add Temporary Capacity operation with a <code>test=true</code> option</li> <li>• "<b>can-be-activated</b>" - the record is available for activation, but not currently active.</li> </ul>
<b>activation-date</b>	Timestamp	Defines the time stamp for when additional capacity for the record was activated.
<b>record-expiration-date</b>	Timestamp	Defines the time stamp for when the capacity record will expire.
<b>activation-expiration-date</b>	Timestamp	Defines the time stamp for when the additional capacity activated for the record will expire and no longer be active.

Table 236. Capacity records: type-specific properties (continued)

Name	Type	Description
<b>maximum-real-days</b>	Integer	Defines the maximum days that real additional capacity can be activated for the record. A value of -1 indicates that the number of days is unlimited.
<b>maximum-test-days</b>	Integer	Defines the maximum days that test additional capacity can be activated for the record. A value of -1 indicates that the number of days is unlimited.
<b>remaining-real-days</b>	Integer	Defines the remaining number of days that additional real capacity can be active for the record. A value of -1 indicates that the number of days is unlimited.
<b>remaining-test-days</b>	Integer	Defines the remaining number of days that additional test capacity can be active for the record. A value of -1 indicates that the number of days is unlimited.
<b>remaining-number-of-real-activations</b>	Integer	Defines the number of times that real additional capacity can be activated for the record. A value of -1 indicates that activation count is unlimited.
<b>remaining-number-of-test-activations</b>	Integer	Defines the number of times that test additional capacity can be activated for the record. A value of -1 indicates that activation count is unlimited.
<b>processor-info</b>	Array of caprec-proc-info objects	A nested object describing the processor capacities available with this capacity record.
<b>available-targets</b>	Array of caprec-target objects	A nested object describing the set of possible activation and deactivation targets contained within this capacity record. One of these targets is chosen via the software-model request body field on the <b>Add Temporary Capacity</b> or <b>Remove Temporary Capacity</b> operations.

Table 237. caprec-proc-info object

Name	Type	Description
<b>type</b>	String Enum	Identifies the type of specialty processor represented. One of: <ul style="list-style-type: none"> <li>• "cp" - central processor</li> <li>• "aap" - Application Assist Processor</li> <li>• "ifl" - Integrated Facility for Linux processor</li> <li>• "icf" - Internal Coupling Facility processor</li> <li>• "iip" - Integrated Information Processors processor</li> <li>• "sap" - System Assist Processor.</li> </ul>
<b>processor-step</b>	Integer	The number of processors steps available
<b>speed-step</b>	Integer	The CP processor speed activation step. A null object is returned for all other processor types.
<b>max-number-processors</b>	Integer	The maximum number of processors available
<b>remaining-processor-days</b>	Integer	The remaining processor days for this processor type. A -1 indicates an unlimited number of days.
<b>remaining-msu-days</b>	Integer	The remaining MSU days for this processor type. A -1 indicates an unlimited number of days. A null object is returned for processor types where this field is not meaningful.

Table 238. caprec-target object

Name	Type	Description
<b>processor-step</b>	Integer	The CPU processor step. This is the incremental delta CPUs compared to the current activation level. The returned value may be negative.

Table 238. *caprec-target object (continued)*

Name	Type	Description
speed-step	Integer	The CPU processor speed activation step. This is the incremental delta speed steps compared to current activation level. The returned value may be negative.
software-model	String (1-3)	The software model that this target represents
billable-msu-cost	Integer	The overall billable MSU cost for this target
billable-msu-delta	Integer	The change in billable MSU cost by activating this target. The value may be negative.

## List Capacity Records

The **List Capacity Records** operation lists the capacity record for a given CPC that are managed by this HMC.

### HTTP method and URI

**GET** /api/cpcs/{cpc-id}/capacity-records

In this request, the URI variable {cpc-id} is the object ID of the target CPC object.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
capacity-record	Array of objects	Array of nested objects (described in the next table).

Each nested object contains the following fields:

Field name	Type	Description
element-uri	String/URI	Canonical URI path of the Capacity Record object.
record-identifier	String	The record identifier of the Capacity Record

### Description

This operation lists the capacity record for a given CPC that are managed by this HMC.

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in the Response Body Contents section.

### Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by {cpc-id}.

### HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in "Response body contents."

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
500 (Server Error)	275	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Get Capacity Record Properties

The **Get Capacity Record Properties** operation retrieves the properties of a single Capacity Record designated by *{capacity-record-id}* from the CPC object designated by *{cpc-id}*.

### HTTP method and URI

**GET** /api/cpcs/{*cpc-id*}/capacity-records/{*capacity-record-id*}

#### URI variables

Variable	Description
<i>{cpc-id}</i>	Object ID of the target CPC object.
<i>{capacity-record-id}</i>	Capacity Record identifier, returned by a previous List Capacity Records operation

### Response body contents

On successful completion, HTTP status code 200 (OK) is returned and the response body provides the current values of the properties for the Capacity Record as defined in “Data model” on page 875.

### Description

The URI path must designate an existing Capacity Record and the API user must have object-access permission to the associated CPC object. If either of these conditions is not met, status code 404 (Not Found) is returned.

On successful execution, HTTP status code 200 (OK) is returned and the response body contains all of the current properties as defined by the Data Model for the Capacity Record object.

### Authorization requirements

This operation has the following authorization requirement:

- Object access permission to the CPC object designated by *{cpc-id}*.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in the Response Body Contents section.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
	276	The capacity record has expired, it can be deleted from the SE.
	302	The capacity record identifier in the URI ( <i>{capacity-record-id}</i> ) must be 1 to 8 characters.
404 (Not Found)	1	The object ID in the URI ( <i>{cpc-id}</i> ) does not designate an existing CPC object, or the API user does not have object access permission to the object.
	274	The capacity record identifier in the URI ( <i>{capacity-record-id}</i> ) does not designate an existing capacity record.
500 (Server Error)	275	An unexpected error occurred during the operation.
503 (Service Unavailable)	1	The request could not be processed because the HMC is not communicating with the SE needed to perform the requested operation.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Inventory service data

- | Information about capacity records can be optionally included in the inventory data provided by the Inventory Service.
- | Inventory entries for the Capacity Record objects are included in the response to the Inventory Service's Get Inventory operation when the request specifies (explicitly by class, implicitly via a containing category, or by default) that objects of class "cpc" are to be included. An entry for a particular capacity record is included only if the API user has access permission to that object as described in the **Get Capacity Record Properties** operation.
- | For each Capacity Record object to be included, the inventory response array includes an entry that is a JSON object with the same contents as is specified in the Response Body Contents section for “Get Capacity Record Properties” on page 878. That is, the data provided is the same as would be provided if a **Get Capacity Record Properties** operation were requested targeting this object.





## Chapter 15. Inventory and metrics services

The functions described in this chapter are termed “services” because unlike the interfaces described in many of the other chapters of this document, the functions described here are service-oriented rather than object-oriented in nature. That is, the functions of these services operate across multiple instances of managed objects rather than being directed at particular managed object instances.

The Inventory Service provides an efficient mechanism for retrieving identify and configuration information about all of the manageable resource instances that are managed by zManager. It provides this information in bulk form via a single request, and thus is expected to be a much more efficient means of determining this information than walking the entire resource tree one object at a time. It is anticipated that this service supports the requirements of a “discovery” phase of a client application that is interested in configuration information about all resources managed by zManager.

The Metrics Service provides a mechanism to retrieve performance metric data for resources that are managed by zManager. This data is captured periodically and buffered on the HMC. The data may include snapshots of performance data at a moment in time, or accumulated performance data, or both, as appropriate. This service is designed to support client applications that provide monitoring function for zManager managed resources.

### Inventory services operations summary

The following operation is provided by the Inventory service:

*Table 239. Inventory service: operations summary*

Operation name	HTTP method and URI path
“Get Inventory” on page 882	POST /api/services/inventory

### Metrics service operations summary

The following operations are provided by the Metrics service:

*Table 240. Metrics service: operations summary*

Operation name	HTTP method and URI path
“Create Metrics Context” on page 888	POST /api/services/metrics/context
“Get Metrics” on page 891	GET /api/services/metrics/context/{metrics-context-id}
“Delete Metrics Context” on page 895	DELETE /api/services/metrics/context/{metrics-context-id}

*Table 241. Metrics service: URI variables*

Variable	Description
{metrics-context-id}	Identifier of a metrics context object. Metrics contexts are associated with API sessions. Thus, this identifier is assigned by the metrics service so that it is unique within an API session and has a lifetime scoped to that session.

## Inventory service

The Inventory Service is an API which allows the client application to fetch a list of ensemble resources and their properties.

This service is intended to support clients that need to determine the inventory and properties of all of the resources of the ensemble (or at least a large portion of those resources). Retrieving this information in bulk form using this service is expected to be much more efficient than walking the resource tree one object at a time using the object-oriented operations of the Web Services API.

The ability to filter the results to only certain classes of resources is provided.

A response to an inventory request is a series of JSON objects returned using HTTP chunked transfer encoding. These objects will be in a format specified in the corresponding resource class's Inventory Data Model sections.

Resources returned are those to which the API client has object-level authorization.

### Get Inventory

The **Get Inventory** operation fetches ensemble resources and associated properties.

#### HTTP method and URI

**POST** /api/services/inventory

#### Request body contents

The request body can include a specification of the classes of resources that should be returned. It contains the following field:

Field name	Type	Description
resources	Array of String Enum	<p>An array of String values. Each element specifies a category or class of resource that should be returned. A category is simply a grouping of classes, so specifying a category is functionally equivalent to specifying all of its member classes. The request may include both categories and classes.</p> <p>Omitting the resources field, or providing an empty array, is equivalent to specifying an array listing all of the supported classes.</p> <p>Categories and associated class values:</p> <ul style="list-style-type: none"> <li>• Category: "zvm-resources" <ul style="list-style-type: none"> <li>– Class: "zvm-virtualization-host"</li> <li>– Class: "zvm-virtual-server"</li> </ul> </li> <li>• Category: "power-vm-resources" <ul style="list-style-type: none"> <li>– Class: "power-vm-virtualization-host"</li> <li>– Class: "power-vm-virtual-server"</li> </ul> </li> <li>• Category: "x-hyp-resources" <ul style="list-style-type: none"> <li>– Class: "x-hyp-virtualization-host"</li> <li>– Class: "x-hyp-virtual-server"</li> </ul> </li> <li>• Category: "prsm-resources" <ul style="list-style-type: none"> <li>– Class: "prsm-virtualization-host"</li> <li>– Class: "prsm-virtual-server"</li> </ul> </li> <li>• Category: "virtual-server-common" <ul style="list-style-type: none"> <li>– Class: "power-vm-virtual-server-common"</li> <li>– Class: "prsm-virtual-server-common"</li> <li>– Class: "x-hyp-virtual-server-common"</li> <li>– Class: "zvm-virtual-server-common"</li> </ul> </li> <li>• Category: "zbx-resources" <ul style="list-style-type: none"> <li>– Class: "zbx"</li> <li>– Class: "rack"</li> <li>– Class: "power-blade"</li> <li>– Class: "system-x-blade"</li> <li>– Class: "isaopt-blade"</li> <li>– Class: "dpxi50z-blade"</li> <li>– Class: "bladecenter"</li> </ul> </li> <li>• Category: "ensemble-wide-resources" <ul style="list-style-type: none"> <li>– Class: "ensemble"</li> <li>– Class: "workload-resource-group"</li> <li>– Class: "virtual-network"</li> <li>– Class: "storage-resource"</li> </ul> </li> <li>• Category: "core-resources" <ul style="list-style-type: none"> <li>– Class: "cpc"</li> <li>– Class: "logical-partition"</li> </ul> </li> <li>• Category: "console-resources" <ul style="list-style-type: none"> <li>– Class: "console"</li> <li>– Class: "custom-group"</li> <li>– Class: "user"</li> <li>– Class: "user-role"</li> </ul> </li> </ul>

Field name	Type	Description
<b>Notes:</b>		
<ul style="list-style-type: none"> <li>The various classes of virtual server, virtualization host, blade, and group above (for example, "x-hyp-virtual-server", "prsm-virtualization-host", "power-blade", and "custom-group") are actually type-specific variations of the object classes "virtual-server", "virtualization-host", "blade" and "group". They are specified with type qualifiers in the names here to allow distinguishing these types on inventory queries. The objects returned in the inventory response will be of the actual object classes ("virtual-server", "virtualization-host", "blade" or "group"), and will have appropriate type values as defined in the data models for those classes.</li> <li>The various classes within the "virtual-server-common" category above represent inventory queries that return a common subset of the virtual server's properties rather than the entire set of properties defined in the Data Model. They correspond to <b>Get Virtual Server Properties</b> requests with the <b>properties=common</b> query parameter specified. Refer to the documentation for the <b>Get Virtual Servers Properties</b> operation and the discussion of inventory service data for Virtual Server objects for specific information on the properties provided.</li> </ul>		

## Response body contents

On successful completion, the response body is a JSON array of JSON objects sent using HTTP chunked transfer encoding. The order in which these objects are returned is unspecified.

The array element documents are of 2 types:

- For resources that were successfully inventoried, the document will be as specified in the corresponding resource's inventory service data.
- For resources that were found but not successfully fully inventoried (i.e. the Object URI can be determined but not the properties), an inventory error document will be returned. Note that, even if one or more of these inventory error documents is contained in the response, an HTTP status code of 200 (OK) is still returned. The fields in the inventory error document are:

Field name	Type	Description
uri	String/URI	Canonical URI of the resource which could not be fully inventoried.
class	String	The class for these error documents is " <b>inventory-error</b> ".
inventory-error-code	Integer	<p>A reason code for the inventory failure. Note that all of these reasons indicate success in locating a resource, but some sort of failure in gathering its properties during inventory collection. A subsequent call to get the properties for the URI in this error document may succeed.</p> <ul style="list-style-type: none"> <li>1: Resource not found on target. Although the resource's URI was located on the HMC, its properties were subsequently not located on the HMC or SE on which the property data for the managed object is to be gathered.</li> <li>2: Communication problem. Communication problems were experienced with the SE on which the property data for the managed object is to be gathered.</li> <li>3: Plugin load error. The code responsible for capturing the properties of a resource class experienced an unexpected problem loading.</li> <li>4: Unknown plugin error. The code responsible for capturing the properties of a resource returned an unrecognized error.</li> <li>5: Unexpected plugin error. The code responsible for capturing the properties of a resource returned an unexpected error.</li> <li>6: Timeout error. The code responsible for capturing the properties of a resource did not respond within the time allocated to it.</li> </ul>
inventory-error-text	String	An error description for the inventory failure.

Field name	Type	Description
inventory-error-details	inventory-error-info Object	A nested inventory-error-info object that provides additional diagnostic information for unexpected inventory plugin errors. This field is provided if the <b>inventory-error-code</b> field is 5 (indicating unexpected plugin error). It is not provided for other <b>inventory-error-code</b> values. The format of the inventory-error-info object is defined in the next table.

The inventory-error-info object contains the following fields:

Field name	Type	Description
http-status	Integer	HTTP status code for the request.
request-uri	String	The URI that caused this error response.
reason	Integer	Numeric reason code providing more details as to the nature of the error) than is provided by the HTTP status code itself. This reason code is treated as a sub-code of the HTTP status code and thus must be used in conjunction with the HTTP status code to determine the error condition. Standard reason codes that apply across the entire API are described in "Common request validation reason codes" on page 24. Additional operation-specific reason codes may also be documented in the description of the specific API operations.
message	String	Message describing the error. This message is not currently localized.
stack	String	Internal HMC diagnostic information for the error. This field is supplied only on selected 5xx HTTP status codes.
error-details	Object	A nested object that provides additional operation-specific error information. This field is provided by selected operations, and the format of the nested object is as described by that operation.

## Description

The Get Inventory operation returns information on ensemble resources and associated properties.

A resource is included in the response if it matches any one of the list of resource classes in the request body. Specifying a category is equivalent to specifying its member classes. If a class is repeated on the request, either explicitly or effectively via categories, the operation will behave as if the class were only specified once. If no resources are specified in the request body, all resources are returned.

Furthermore, a resource is included in the response only if the API user has object-access permission for that resource. If an HMC is a manager of a resource but the API user does not have permission to it, that resource is simply omitted from the response. A success status code is still returned.

If the HMC does not manage any resources to which the user has access, or if no resources are found that match the request body specification, an empty response is returned with a 204 (No Content) status code.

In addition to objects for inventoried resources, the response may include objects for resources whose URIs could be determined, but whose properties could not, for some reason, be obtained. Rather than treat these resources as completely non-inventoried and omit them, the URI and an error reason are returned.

The order in which the objects are returned is unspecified.

The Get Inventory implementation may choose to limit the number of simultaneous in-process inventory requests. If such a limit is reached, further requests will return an HTTP 503 (Service Unavailable) error status code until previous requests complete and the number of in-process inventory requests falls back below the limit.

## Authorization requirements

This operation has the following authorization requirement:

- Object access permission to any object to be included in the result.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 884. If there are no resources to provide, HTTP status code 204 (No Content) is returned, along with an empty response body.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
503 (Service Unavailable)	200	The request could not be processed because of the number of currently pending inventory requests. The request can be reissued at a later time.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

The **Get Inventory** results represent a snapshot of inventory results as viewed from the HMC. The actual inventory can change, even as the results are being streamed back to the API client. Therefore, if the client wishes to stay informed about changes to the inventory and not risk missing any inventory changes, it should use the API event mechanisms to subscribe to inventory-related events before even issuing a **Get Inventory** request.

The **Get Inventory** results do not reflect all properties at a single moment in time. During the overall inventory collection process multiple resource's states and other properties may change. Therefore, states (or other properties) among two or more resources that might normally be expected to match (or have some other expected relationship) at one moment in time may instead return apparently inconsistent results in the **Get Inventory** response.

## Example HTTP interaction

The following example illustrates a typical response for a **Get Inventory** request for the ensemble class of resources. Responses for other classes will differ significantly from this because the data differs on a class by class basis. The format of the data returned by the Inventory Service for each class of object is described in a section entitled “Inventory service data” within the documentation for that object class.

---

```

POST /api/services/inventory HTTP/1.1
x-api-session: 38gu0i9so1y0vjemqytpcadnwq0esu32pjd3ub4jy61xadp72
content-type: application/json
content-length: 27
{
  "resources": [
    "ensemble"
  ]
}

```

---

Figure 437. Get Inventory: Request

---

```

200 OK
transfer-encoding: chunked
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 07 Dec 2011 03:42:15 GMT
content-type: application/json;charset=UTF-8
[
  {
    "acceptable-status": [
      "alternate-communicating"
    ],
    "class": "ensemble",
    "cpu-perf-mgmt-enabled-power-vm": false,
    "cpu-perf-mgmt-enabled-zvm": false,
    "description": "long ensemble name",
    "has-unacceptable-status": false,
    "is-locked": false,
    "load-balancing-enabled": false,
    "load-balancing-ip-addresses": [],
    "load-balancing-port": 3860,
    "mac-prefix": "02:00:00:00:00:00",
    "management-enablement-level": "automate",
    "name": "HMC_R74_ENSEMBLE",
    "object-id": "1f7ffb02-de39-11e0-88bd-00215e67351a",
    "object-uri": "/api/ensembles/1f7ffb02-de39-11e0-88bd-00215e67351a",
    "parent": null,
    "power-consumption": 24474,
    "power-rating": 65644,
    "reserved-mac-address-prefixes": [],
    "status": "alternate-communicating",
    "unique-local-unified-prefix": "fd2c:34be:df2:0:0:0:0:0"
  },
  {
    "class": "node",
    "element-uri": "/api/ensembles/1f7ffb02-de39-11e0-88bd-00215e67351a/nodes/9ba3b1aa-693a-3408-80ae-9d0808147ffa",
    "member": "/api/cpcs/9ba3b1aa-693a-3408-80ae-9d0808147ffa",
    "parent": "/api/ensembles/1f7ffb02-de39-11e0-88bd-00215e67351a",
    "type": "cpc"
  }
]

```

---

Figure 438. Get Inventory: Response

## Metrics service

The zEnterprise (or later) Ensembles, Central Processing Complexes (CPCs), and their associated system resources are instrumented at key points to collect performance and utilization data. The data is forwarded by the metric data providers to a buffer on the HMC where it is made available to consumers of this API.

The data collection instrumentation organizes and formalizes the data it collects into a series of named metric groups. The Metrics Service API allows specification of the metric groups the client wishes to query. The API returns some information about the format of the metrics that are being fetched. Specifically, a structure called a metrics context is associated with any metrics retrieval, and that structure includes metric group names, individual metric field names, and the associated individual metric data types.

The full metric group documentation, however, including descriptions of the data collected and the frequency of collection, can be found in Chapter 16, “zManager metric groups,” on page 897.

## Create Metrics Context

The **Create Metrics Context** operation creates a context under which metrics can be repeatedly retrieved. This context will be associated with the API session under which it was created.

### HTTP method and URI

**POST** /api/services/metrics/context

### Request body contents

A request body must be specified. It has the following fields:

Field name	Type	Description
anticipated-frequency-seconds	Integer	The number of seconds the client anticipates will elapse between <b>Get Metrics</b> calls against this context. The minimum accepted value is 15.
metric-groups	Array of Strings	Optional. Array of metric group names. If specified, then results from future <b>Get Metrics</b> requests associated with this context will be limited to only metrics with group names matching one of the specified values. If not specified, or if an empty array is specified, then results will not be limited with respect to metric group names.

### Response body contents

On successful completion, the response body contains a JSON object with the following fields:

Field name	Type	Description
metrics-context-uri	String/URI	Canonical URI path of the metrics context object created by this operation. This includes the metrics-context-id. E.g. “/api/services/metrics/context/1”, where “1” is the metrics-context-id.
metric-group-infos	Array of objects	Array of metric-group-info objects (described in the next table) that describe the data format for each metric group that may be returned by future GETs associated with this metric context.

Each nested metric-group-info object contains the following fields:



Field name	Type	Description
group-name	String	The name of the metric group for which we are providing descriptive information.
metric-infos	Array	Array of metric-info objects (described in the next table). These describe each metric for the group in the order that they will appear in future GETs associated with this context.

Each nested metric-info object contains the following fields:

Field name	Type	Description
metric-name	String	The name of the metric.
metric-type	String Enum	One of the following, describing the type of the metric:  "boolean-metric", "byte-metric", "double-metric", "long-metric", "integer-metric", "short-metric", "string-metric"  See the Get Metrics "Response body contents" on page 892 for further information on these metric types.

## Description

This operation establishes a context for future **Get Metrics** operations that is valid for the current API session. Because of the high frequency of invocation and large volume of data expected, the metrics service interface has been structured to optimize the transmission of data on each **Get Metrics** request. Thus, rather than use a self-describing representation for the results returned by each Get Metrics, the metrics service instead provides the descriptive metadata as results from this **Create Metrics Context** operation. It then returns the metric data in a compact format each time **Get Metrics** is invoked.

At a high level, the **Create Metrics Context** response communicates two primary pieces of information back to the client. One is the metrics-context-uri, which includes the ID of the metrics context that must be referenced on future GETs to associate them with this context. The other is the metric-groups description data. That data provides the metric type and metric name information for each metric group whose metrics may be returned by this context. This may be useful to the client for determining how to parse future **Get Metrics** responses for this context, although the full documentation on metric group formats is found in Chapter 16, "zManager metric groups," on page 897.

The anticipated-frequency-seconds specification which is required on the request body tells the metrics service how frequently the client anticipates issuing **Get Metrics** requests against this context. The metrics service may take no action based on this frequency, but reserves the right to invalidate and delete the metrics context if 4 times the specified frequency passes without receipt of an associated **Get Metrics** operation.

Optional result filtering is provided by field metric-groups on the request body. If a non-empty metric-groups arrays is specified, then future **Get Metrics** operations associated with this context will return only groups with names listed there.

Additionally, if a metric-groups array of group names is specified on the **Create Metrics Context** request, then the response JSON document will include only matching metric-group-info fields. If one or more names in the metric-groups array does not represent a metric group registered on the HMC, then HTTP error status code 400 (Bad Request) will be returned and the context will not be established.

Although the POST response fully describes and guarantees the ordering of metric-infos within a metric group for that context, as a matter of policy the HMC will further guarantee that, for a given metric group, any additions of new metrics to the group will be to the end of the list for the group.

## Authorization requirements

There are no authorization restrictions on creating a metrics context. However any future metric results returned by **Get Metrics** queries against that context will be restricted to managed objects accessible according to the permissions associated with the API session under which the metrics context was established.

Note that there is no indication via an HTTP status or reason code that future results may be restricted due to authorization restrictions. Rather, success is indicated and future **Get Metrics** responses behave just as if any restricted objects did not exist.

## HTTP status and reason codes

On success, HTTP status code 201 (Created) is returned and the response body is provided as described in “Response body contents” on page 888. The URI for the newly created context is also provided in the **Location** header of the response.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Example HTTP interaction

**Note:** These are example rather than actual metrics group names.

---

```
POST /api/services/metrics/context HTTP/1.1
x-api-session: 6a9oz3ymut6rvjijrft0loqhfzgp0rnu4mjishwh6d39jh31q
content-type: application/json
content-length: 96
{
  "anticipated-frequency-seconds": 45,
  "metric-groups": [
    "virtualization-host-cpu-memory-usage"
  ]
}
```

---

Figure 439. Create Metrics Context: Request

---

```

201 Created
transfer-encoding: chunked
server: zSeries management console API web server / 1.0
location: /api/services/metrics/context/1
cache-control: no-cache
date: Wed, 07 Dec 2011 04:01:59 GMT
content-type: application/json;charset=UTF-8
{
  "metric-group-infos": [
    {
      "group-name": "virtualization-host-cpu-memory-usage",
      "metric-infos": [
        {
          "metric-name": "processor-usage",
          "metric-type": "integer-metric"
        },
        {
          "metric-name": "memory-usage",
          "metric-type": "integer-metric"
        },
        {
          "metric-name": "network-usage",
          "metric-type": "integer-metric"
        },
        {
          "metric-name": "storage-rate",
          "metric-type": "integer-metric"
        },
        {
          "metric-name": "physical-cpu-time",
          "metric-type": "long-metric"
        },
        {
          "metric-name": "memory-used",
          "metric-type": "integer-metric"
        },
        {
          "metric-name": "virt-host-management-cpu-time-used",
          "metric-type": "long-metric"
        },
        {
          "metric-name": "virt-host-page-ins",
          "metric-type": "long-metric"
        },
        {
          "metric-name": "virt-host-page-outs",
          "metric-type": "long-metric"
        }
      ]
    }
  ],
  "metrics-context-uri": "/api/services/metrics/context/1"
}

```

---

Figure 440. Create Metrics Context: Response

## Get Metrics

The **Get Metrics** operation retrieves the current set of metrics associated with an established metrics context.

## HTTP method and URI

**GET** /api/services/metrics/context/{*metrics-context-id*}

In this request, the URI variable *{metrics-context-id}* is the identifier of the metrics context object for which metrics are to be obtained.

## Response body contents

On successful completion, the response body contains the set of metrics associated with the metrics context. The response is sent using HTTP chunked transfer encoding and UTF-8 character encoding. A MIME media type of **application/vnd.ibm-z-zmanager-metrics** is used and is specified in the **Content-Type** header on the response.

Because performance and scalability are a major concern for the metrics service, the response body is in a terse custom format, rather than being presented as a JSON object. The data type, name, and order information required to parse and interpret the response is provided in a previous **Create Metrics Context** response.

Data in this format will be delimited by newlines and commas.

Using a partial Extended Backus-Naur Form, where a comma (,) indicates concatenation and curly braces ({} ) indicate 0 or more repetitions, we can express the format this way:

```
MetricsResponse = {MetricsGroup},NL
MetricsGroup = MetricsGroupName,NL,{ObjectValues},NL
MetricsGroupName = StringValue
NL = "\n"
ObjectValues = ObjectURI,NL,Timestamp,NL,ValueRows,NL
Timestamp = LongValue
ObjectURI = StringValue
ValueRows = ValueRow,{ValueRow}
ValueRow = Value,{"",Value},NL
Value = BooleanValue | ByteValue | DoubleValue | LongValue | IntegerValue | ShortValue | StringValue
```

The *MetricsGroupName* is the name of the metrics group, as a *StringValue* as defined below.

The *Timestamp* is the time when the associated values were buffered (i.e. “cached”) on the HMC. It is expressed as an “epoch” timestamp: a *LongValue* giving the milliseconds since January 1, 1970, 00:00:00 GMT (just as is expected, for example, by the constructor of a *java.util.Date* object).

The *ObjectURI* is the canonical URI of the object, as a *StringValue* as defined below.

*NL* is a single newline character (Unicode U+000A).

All the varieties of *Value* will be represented as strings according to the following rules and limits:

- *BooleanValue*
  - Either the string true or the string false.
- *ByteValue*
  - A string representation of a signed decimal integer in the range -128 to 127 (i.e. the range of a signed 8 bit integer).
- *DoubleValue*

- A string representation of a 64 bit IEEE 754 floating point number in the range +/-4.9E-324 to +/-3.4028235E+38. Note that, although IEEE 754 provides for representations of positive or negative Infinity and NaN, such values are not allowed in the metric data feed and thus will not appear in a metrics service result. For results with a magnitude greater than or equal to 10<sup>-3</sup> and less than 10<sup>7</sup>, the string representation will be a dotted decimal (e.g. 1.7, -32.467). For results with magnitudes outside that range, the string representation will be computerized scientific notation (e.g. -4.23E127).
- LongValue
  - A string representation of a signed decimal integer in the range -9223372036854775808 to 9223372036854775807 (i.e. the range of a signed 64 bit integer).
- IntegerValue
  - A string representation of a signed decimal integer in the range -2147483648 to 2147483647 (i.e. the range of a signed 32 bit integer).
- ShortValue
  - A string representation of a signed decimal integer in the range -32768 to 32767 (i.e. the range of a signed 16 bit integer).
- StringValue
  - A string starting with a double-quote, ending with a double-quote, and with any embedded double-quotes or backslashes escaped with a preceding backslash (i.e. escaped as \" and \\).

## Description

On successful execution status code 200 (OK) is returned, with a response body as described above.

The URI path on the request must designate an existing metrics context for the current API session. If the URI designates an unrecognized context for the API session, then status code 404 (Not Found) is returned.

Note that under some circumstances the metrics service may delete a metrics context, requiring the client to establish a new context in order to resume metric retrievals. For example, the metrics service may choose to delete a given context if the time since the last associated **Get Metrics** request has exceeded 4 times the anticipated frequency specified when the context was created. In addition, the client may explicitly delete a metrics context with the **Delete Metrics Context** operation. If the URI designates a context that was once valid for the current API session, but no longer is, then status code 409 (Conflict) is returned.

## Authorization requirements

Only metrics referencing managed objects accessible according to the permissions associated with the API session under which the Get Metrics is being issued will be returned. Note that there is no indication via an HTTP status or reason code that results may have been restricted due to authorization restrictions. Rather, success is indicated and the responses are just as if any restricted objects did not exist.

## HTTP status and reason codes

On success, HTTP status code 200 (OK) is returned and the response body is provided as described in “Response body contents” on page 892.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.

HTTP error status code	Reason code	Description
404 (Not Found)	1	The metrics context ID in the URI ( <i>{metrics-context-id}</i> ) does not designate a metrics context for the associated API session.
409 (Conflict)	1	The metrics context ID in the URI ( <i>{metrics-context-id}</i> ) designates a metrics context for the associated API session that is no longer valid.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

## Usage notes

- Repeated metrics retrievals, even consecutive retrievals against a common metrics context, will not necessarily yield metrics for the exact same set of objects. Objects can come and go from the system's inventory due to various circumstances unrelated to the metrics data. The metrics feed for a particular metric group may stop for some reason, and the metrics data may therefore expire from the HMC's buffer (i.e. the metrics cache). The access permissions of a user associated with a metrics context may change, giving the user access to a smaller or larger set of objects (and therefore, perhaps, a smaller or larger set of metrics data).
- It is possible that there may be no metrics to return for one or more metric groups associated with the specified metrics context. For example, data for a metric group may not be buffered on the HMC at the time of the Get Metrics invocation, or authorization restrictions related to objects in a requested metric group may prevent any metrics for that group from being returned. If there is no metric data to be returned for a metric group name, then that group name does not appear in the response body. Furthermore, if there is no metric data to return for any metric group name associated with a context, the response body format above specifies that the body will consist only of a single newline character. This is nonetheless considered a successful response. In other words, an HTTP status code 200 (OK) will still be returned with such a response.

## Example HTTP interaction

---

```
GET /api/services/metrics/context/1 HTTP/1.1
x-api-session: 4g33uc2vith3v42vmce8wjr5q7x58x5ybh6hwd4vjpwr7j14sz
```

---

Figure 441. Get Metrics: Request

---

```

200 OK
transfer-encoding: chunked
server: zSeries management console API web server / 1.0
cache-control: no-cache
date: Wed, 07 Dec 2011 04:38:20 GMT
content-type: application/vnd.ibm-z-zmanager-metrics;charset=UTF-8
"virtualization-host-cpu-memory-usage"
"/api/virtualization-hosts/2f7bf364-03f8-11e1-8eda-001f163805d8"
1323232689283
0,14,0,3,2942386000,4608,2942386000,0,0

"/api/virtualization-hosts/c14cde64-03e6-11e1-baf3-001f163805d8"
1323232689283
0,64,-1,-1,76028000,1311,-1,-1,-1

"/api/virtualization-hosts/5f805d28-03e6-11e1-baf3-001f163805d8"
1323232689283
3,72,-1,-1,1731046000,1485,-1,-1,-1

"/api/virtualization-hosts/634fa694-03f4-11e1-881f-001f163805d8"
1323232689283
4,36,-1,-1,55878116000,17920,-1,-1,-1
<3 blank lines here (consecutive new lines)>

```

---

Figure 442. Get Metrics: Response

## Delete Metrics Context

The **Delete Metrics Context** operation deletes a metrics context.

### HTTP method and URI

```
DELETE /api/services/metrics/context/{metrics-context-id}
```

In this request, the URI variable *{metrics-context-id}* is the identifier of the metrics context object for which metrics are to be obtained.

### Description

This operation deletes the metrics context ID. That is, it disassociates it from the API session and cleans up any data associated with it. Further **Get Metrics** requests against this context will result in status code 409 (Conflict).

The URI path must designate an existing valid metrics context for the current API session. If the URI path represents an already invalidated metrics context for the current API session, status code 409 (Conflict) is returned. If the URI path does not represent a recognized metrics context for the current API session, status code 404 (Not Found) is returned.

### Authorization requirements

There are no authorization requirements for deleting a metrics context. The association with the API session is implicit, so there is no possibility of deleting a context that was created by a different API session. In other words, only the session which created a metrics context can delete it.

### HTTP status and reason codes

On success, HTTP status code 204 (No Content) is returned and no response body is provided.

The following HTTP status codes are returned for the indicated errors, and the response body is a standard error response body providing the reason code indicated and associated error message.

HTTP error status code	Reason code	Description
400 (Bad Request)	Various	Errors were detected during common request validation. See “Common request validation reason codes” on page 24 for a list of the possible reason codes.
404 (Not Found)	1	The metrics context ID in the URI ( <i>{metrics-context-id}</i> ) does not designate a metrics context for the associated API session.
409 (Conflict)	1	The metrics context ID in the URI ( <i>{metrics-context-id}</i> ) designates a metrics context for the associated API session that is no longer valid.

Additional standard status and reason codes can be returned, as described in Chapter 3, “Invoking API operations,” on page 19.

### Example HTTP interaction

---

```
DELETE /api/services/metrics/context/1 HTTP/1.1
x-api-session: 6a9oz3ymut6rvjijrft0loqhfzgp0rnu4mjishwh6d39jh31q
```

---

*Figure 443. Delete Metrics Context: Request*

---

```
204 No Content
date: Wed, 07 Dec 2011 04:01:59 GMT
server: zSeries management console API web server / 1.0
cache-control: no-cache

<No response body>
```

---

*Figure 444. Delete Metrics Context: Response*



## Chapter 16. zManager metric groups

This chapter provides a description of the metric groups that can be retrieved using the Metrics Service. For each metric group provided by the HMC, the material in this chapter describes the purpose and general characteristics of the metric group, and then defines the content of the metric group via a table that specifies the metric fields provided by the group. The order in which metric fields appear within these tables corresponds to the order in which the data items appear in a value row returned by the **Get Metrics** operation. For example, the metric field appearing in the first row of a metric group table (and identified in a table below as being in position 1) will be the first data item provided in a value row for that metric group; the metric field appearing in the second row (position 2) will be the next data item in a value row, and so on. Thus, the order in which metric fields are documented here is considered semantically significant and can be relied upon by client applications in order to simplify parsing of the data retrieved using the **Get Metrics** operation.

The contents of metric groups may be extended in future versions of this API. If a metric group is extended, new metric fields will be added to the end so as to not alter the relative positions of any of the existing fields. Such new fields would not be understood by a client application designed for an earlier version of the API. Therefore, applications must be developed using the philosophy of "ignore what you don't understand/expect" when processing metric group data order to tolerate such possible future extensions. See "Compatibility" on page 4 for more discussion on API compatibility principles.

### Monitors dashboard metric groups

The Monitors Dashboard task is the current system monitoring interface on the HMC. It gives a dashboard display to monitor system resources, such as power consumption, environmental data, processor usage, etc.

In order to provide programmatic access to this same data, the utilization and environment data that is displayed on the user interface is also provided via the Metrics Service in the following metric groups.

### BladeCenter temperature and power metric group

This metric group reports the ambient temperature and power consumption for each BladeCenter on the system.

**Metric Group Name**

"bladecenter-temperature-and-power"

**Collection Interval**

15 seconds

**Applicable Managed Object Class**

"bladecenter"

The following metrics are provided in each entry of this metric group:

*Table 242. BladeCenter temperature and power metric group*

Pos	Metric field name	Type	Units	Description
1	temperature-celsius	Double	Degrees Celsius	The ambient temperature
2	power-consumption-watts	Integer	Watts	The power consumption

## Blade power

This metric group reports power consumption for each blade on the system.

### Metric Group Name

"blade-power-consumption"

### Collection Interval

15 seconds

### Applicable Managed Object Class

"blade"

The following metrics are provided in each entry of this metric group:

Table 243. Blade power metric group

Pos	Metric field name	Type	Units	Description
1	power-consumption-watts	Integer	Watts	The power consumption

## Channels

This metric group reports the channel usage for each channel on the system. An instance of this metric group is created for each channel of a CPC.

### Metric Group Name

"channel-usage"

### Collection Interval

15 seconds

### Applicable Managed Object Class

"cpc"

The following metrics are provided in each entry of this metric group:

Table 244. Channels metric group

Pos	Metric field name	Type	Units	Description
1	channel-name	String	—	The name of the channel in the form CSS.Chpid
2	shared-channel	Boolean	—	True if the channel is shared among logical partitions, and false if it is not
3	logical-partition-name	String	—	For channel types for which logical partition names are available, the name of the owning logical partition or the value "shared" if the channel is shared. For other channel types for which the name is not available (for example, coupling channels), the value is always an empty string.
4	channel-usage	Integer	%	The channel percent usage (0 – 100%)

## CPC overview

This metric group reports the aggregated processor usage and channel usage, the ambient temperature, and total system power consumption for each system. The cpc-processor-usage is the average of the percentages of processing capacity for the physical processor lines in the active System Activity profile for the CPC. The channel-usage is the average of the percentages of I/O capacity for the channel lines in the active System Activity profile for the CPC.

### Metric Group Name

"cpc-usage-overview"

**Collection Interval**

15 seconds

**Applicable Managed Object Class****"cpc"**

The following metrics are provided in each entry of this metric group:

*Table 245. CPC overview metric group*

Pos	Metric field name	Type	Units	Description
1	<b>cpc-processor-usage</b>	Integer	%	The processor percent usage.
2	<b>channel-usage</b>	Integer	%	The channel percent usage.
3	<b>power-consumption-watts</b>	Integer	Watts	The total system power consumption.
4	<b>temperature-celsius</b>	Double	Degrees Celsius	The ambient temperature.
5	<b>cp-shared-processor-usage</b>	Integer	%	The processor percent usage for all CP shared processors. Set to -1 if there are no processors of this type.
6	<b>cp-dedicated-processor-usage</b>	Integer	%	The processor percent usage for all CP dedicated processors. Set to -1 if there are no processors of this type.
7	<b>ifl-shared-processor-usage</b>	Integer	%	The processor percent usage for all IFL shared processors. Set to -1 if there are no processors of this type.
8	<b>ifl-dedicated-processor-usage</b>	Integer	%	The processor percent usage for all IFL dedicated processors. Set to -1 if there are no processors of this type.
9	<b>icf-shared-processor-usage</b>	Integer	%	The processor percent usage for all ICF shared processors. Set to -1 if there are no processors of this type.
10	<b>icf-dedicated-processor-usage</b>	Integer	%	The processor percent usage for all ICF dedicated processors. Set to -1 if there are no processors of this type.
11	<b>iip-shared-processor-usage</b>	Integer	%	The processor percent usage for all zIIP shared processors. Set to -1 if there are no processors of this type.
12	<b>iip-dedicated-processor-usage</b>	Integer	%	The processor percent usage for all zIIP dedicated processors. Set to -1 if there are no processors of this type.
13	<b>aap-shared-processor-usage</b>	Integer	%	The processor percent usage for all zAAP shared processors. Set to -1 if there are no processors of this type.
14	<b>aap-dedicated-processor-usage</b>	Integer	%	The processor percent usage for all zAAP dedicated processors. Set to -1 if there are no processors of this type.
15	<b>all-shared-processor-usage</b>	Integer	%	The processor percent usage for all shared processors. Set to -1 if there are no processors of this type.
16	<b>all-dedicated-processor-usage</b>	Integer	%	The processor percent usage for all dedicated processors. Set to -1 if there are no processors of this type.

## | **zBX (Node) overview**

- | This metric group reports the average processor usage and network I/O usage, ambient temperature, and total system power for each zBX node. The processor usage is the average percent usage of all processors.
- | The network I/O usage is the average percent usage of all the blades that support network I/O. The ambient temperature value is the average of the ambient temperature values for all BladeCenter chassis.
- | An instance of this metric group is created for each zBX node.

### | **Metric Group Name**

| **"zbx-node-usage-overview"**

### | **Collection Interval**

| 15 seconds

| **Applicable Managed Object Class**  
| "zbx" (when type is "node")

| The following metrics are provided in each entry of this metric group:

| *Table 246. zBX node overview metric group*

Pos	Metric field name	Type	Units	Description
1	processor-usage	Integer	%	The average percent usage of all processors.
2	io-usage	Integer	%	The average I/O percent usage for all blades that support network I/O.
3	power-consumption-watts	Integer	Watts	The total system power consumption. This includes the zBX node and all its blades and BladeCenter chassis.
4	temperature-celsius	Double	Degrees Celsius	The average ambient temperature of all BladeCenter chassis.

## | Logical partitions

This metric group reports the processor usage and z/VM paging rate for each active logical partition (aka Image, LPAR Image, Zone, PR/SM virtual server) on the system.

**Metric Group Name**  
"logical-partition-usage"

**Collection Interval**  
15 seconds

**Applicable Managed Object Class**  
"logical-partition"

The following metrics are provided in each entry of this metric group:

*Table 247. Logical partitions metric group*

Pos	Metric field name	Type	Units	Description
1	processor-usage	Integer	%	The processor percent usage.
2	zvm-paging-rate	Integer	Pages Per Second	The z/VM paging rate. Only returned for logical partitions running z/VM level 6.1 or higher that have the appropriate agent running in them.
3	cp-processor-usage	Integer	%	The processor percent usage for all CP processors. Set to -1 if there are no processors of this type.
4	ifl-processor-usage	Integer	%	The processor percent usage for all IFL processors. Set to -1 if there are no processors of this type.
5	icf-processor-usage	Integer	%	The processor percent usage for all ICF processors. Set to -1 if there are no processors of this type.
6	iip-processor-usage	Integer	%	The processor percent usage for all IIP processors. Set to -1 if there are no processors of this type.

## zCPC environmentals and power

This metric group reports environmental data and power consumption for the zCPC, which is the legacy part of the system; i.e. without blades.

**Metric Group Name**  
"zpc-environmentals-and-power"

**Collection Interval**

15 seconds

**Applicable Managed Object Class****"cpc"**

The following metrics are provided in each entry of this metric group:

*Table 248. zCPC environmental and power metric group*

Pos	Metric field name	Type	Units	Description
1	<b>temperature-celsius</b>	Double	Degrees Celsius	The ambient temperature
2	<b>humidity</b>	Integer	%	The relative humidity
3	<b>dew-point-celsius</b>	Double	Degrees Celsius	The dew point
4	<b>power-consumption-watts</b>	Integer	Watts	The power consumption in watts
5	<b>heat-load</b>	Integer	Btu/hour	The total heat load of the system (heat load forced-air + heat load water)
6	<b>heat-load-forced-air</b>	Integer	Btu/hour	The heat load covered by forced-air
7	<b>heat-load-water</b>	Integer	Btu/hour	The heat load covered by water

**zCPC processors**

This metric group reports the processor usage for each physical zCPC processor on the system. This includes the System Assist Processors (SAPs) and does not include blades. An instance of this metric group is created for each processor of a CPC.

**Metric Group Name****"zpc-processor-usage"****Collection Interval**

15 seconds

**Applicable Managed Object Class****"cpc"**

The following metrics are provided in each entry of this metric group:

*Table 249. zCPC processors metric group*

Pos	Metric field name	Type	Units	Description
1	<b>processor-name</b>	String		The name of the zpc processor in the form processor-type + processor ID. For example, IFL01.
2	<b>processor-type</b>	String		The type of zpc processor. The valid types are: <b>"cp"</b> , <b>"ifl"</b> , <b>"icf"</b> , <b>"iip"</b> , <b>"aap"</b> , <b>"sap"</b> .
3	<b>processor-usage</b>	Integer	%	The processor percent usage.
4	<b>smt-usage</b>	Integer	%	The percentage of time the processor is running in simultaneous multithreading (SMT) mode. Set to -1 when SMT mode is not supported by the CPC.
5	<b>thread-0-usage</b>	Integer	%	The percent usage of thread 0 when the processor is running in simultaneous multithreading (SMT) mode. Set to -1 when SMT mode is not supported by the CPC.

Table 249. zCPC processors metric group (continued)

Pos	Metric field name	Type	Units	Description
6	thread-1-usage	Integer	%	The percent usage of thread 1 when the processor is running in simultaneous multithreading (SMT) mode. Set to -1 when SMT mode is not supported by the CPC.

## Blade CPU and memory metric group

This metric group reports CPU and memory utilization for each of the blades in the ensemble. This group provides data for all types of blades.

### Metric Group Name

"blade-cpu-memory-usage"

### Collection Interval

15 seconds

### Applicable Managed Object Class

"blade"

The following metrics are provided in each entry of this metric group:

Table 250. Blade CPU and memory metric group

Pos	Metric field name	Type	Units	Description
1	processor-usage	Integer	%	Processor utilization percentage (0-100%). Value for current interval. If the value is not currently available, for example if the blade is powered off, -1 is provided.
2	memory-usage	Integer	%	Memory utilization percentage (0-100%). Value for current interval. If the value is not currently available, for example if the blade is powered off, -1 is provided.

## Cryptos

This metric group reports the adapter usage for each crypto on the system. An instance of this metric group is created for each crypto of a CPC. If a CPC has no crypto adapters, then no data will appear in this metric group for that CPC.

### Metric Group Name

"crypto-usage"

### Collection Interval

15 seconds

### Applicable Managed Object Class

"cpc"

The following metrics are provided in each entry of this metric group:

Table 251. Crypto metric group

Pos	Metric field name	Type	Units	Description
1	channel-id	String		The physical channel identifier of the crypto
2	crypto-id	String		The crypto identifier of the crypto, decimal 0-15
3	adapter-usage	Integer	%	The adapter percent usage (0-100%)

## Flash memory adapters

This metric group reports the adapter usage for each Flash memory (Flash Express) adapter on the system. An instance of this metric group is created for each Flash memory adapter of the CPC. If a CPC has no flash memory adapters, then no data will appear in this metric group for that CPC.

**Note:** Flash Express has a planned exploitation of December 2012.

**Metric Group Name**  
"flash-memory-usage"

**Collection Interval**  
15 seconds

**Applicable Managed Object Class**  
"cpc"

The following metrics are provided in each entry of this metric group:

*Table 252. Flash memory adapters metric group*

Pos	Metric field name	Type	Units	Description
1	channel-id	String		The physical channel identifier of the Flash memory adapter
2	adapter-usage	Integer	%	The adapter percent usage (0-100%)

## RoCE adapters

This metric group reports the adapter usage for each RoCE (10GbE RoCE) adapter on the system. An instance of this metric group is created for each RoCE adapter of the CPC.

**Metric Group Name**  
"roce-usage"

**Collection Interval**  
15 seconds

**Applicable Managed Object Class**  
"cpc"

The following metrics are provided in each entry of this metric group:

*Table 253. RoCE adapters metric group*

Pos	Metric field name	Type	Units	Description
1	channel-id	String		The physical channel identifier of the RoCE adapter
2	adapter-usage	Integer	%	The adapter percent usage (0-100%)

## Ensemble power

This metric group reports the consumption for the Ensemble. An instance of this metric group is created for the Ensemble, if one exists on the HMC.

**Metric Group Name**  
"ensemble-power"

**Collection Interval**  
15 seconds

**Applicable Managed Object Class**  
"ensemble"

The following metrics are provided in each entry of this metric group:

Table 254. Ensemble power metric group

Pos	Metric field name	Type	Units	Description
1	<b>power-consumption-watts</b>	Integer	Watts	The total power consumption of all members of the Ensemble.

## Performance management metrics groups

Following are the performance management metrics groups.

### Virtual server CPU and memory metrics group

This metric group is collected for all types of virtual servers: PowerVM, PR/SM, x Hyp, and z/VM.

#### Metric Group Name

"virtual-server-cpu-memory-usage"

#### Collection Interval

15 seconds for PR/SM, x Hyp and z/VM virtual servers

30 seconds for PowerVM virtual servers

#### Applicable Managed Object Class

"virtual-server"

The following metrics are provided in each entry of this metric group:

Table 255. Virtual server CPU and memory metric group

Pos	Metric field name	Type	Units	Description
1	<b>processor-usage</b>	Integer	%	Physical utilization percentage (0-100%). Values for current interval.
2	<b>memory-usage</b>	Integer	%	Memory usage percentage (0-100%). Value for current interval.
3	<b>up-time</b>	Long	Microseconds	Time since the virtual server was started. Not supported for PR/SM, will be reported as -1.
4	<b>physical-cpu-time</b>	Long	Microseconds	Physical CPU time used by virtual server. This time accumulates until the virtualization host or the support element is restarted.
5	<b>virt-host-cpu-delay-time</b>	Long	Microseconds	The virtual processors were ready to run but not dispatched on physical processors due to competition with other virtual servers. This time accumulates until the virtualization host or the support element is restarted. Not supported for PR/SM, will be reported as -1.
6	<b>elapsed-time</b>	Long	Microseconds	Elapsed time over which physical-cpu-time and virt-host-cpu-delay-time have accumulated.
7	<b>cpu-allocation</b>	Integer	Share for z/VM and x Hyp, physical CPU equivalent for PowerVM, and processing weight for PR/SM	Amount of CPU resource allocated to virtual server. zManager adjusts this value when virtual server management is enabled. For PowerVM, this is the processing units setting. For z/VM and x Hyp, this is the CPU share value. For PS/SM this is the general purpose processing weight setting. Not supported for virtual servers with dedicated virtual processors, will be reported as -1.



Table 255. Virtual server CPU and memory metric group (continued)

Pos	Metric field name	Type	Units	Description
8	<b>current-physical-memory-used</b>	Integer	Megabytes	Amount of physical memory currently used by virtual server.
9	<b>os-total-cpu-using-samples<sup>1</sup></b>	Long	Count	Count of samples where virtual CPUs were seen in use. Will be 0 if GPMP not running on virtual server.
10	<b>os-total-cpu-delay-samples<sup>1</sup></b>	Long	Count	Count of samples where threads were delayed waiting for virtual CPUs. Will be 0 if GPMP not running on virtual server.
11	<b>os-total-paging-delay-samples<sup>1</sup></b>	Long	Count	Count of samples where threads were delayed waiting for page faults to be resolved. Will be 0 if GPMP not running on virtual server.
12	<b>os-total-io-delay-sample<sup>1</sup></b>	Long	Count	Count of samples where threads were delayed waiting for I/O to complete. Will be 0 if GPMP not running on virtual server.
13	<b>os-sampling-rate</b>	Integer	Count	Number of times per second OS samples are collected. Will be 0 if GPMP not running on virtual server.
14	<b>os-total-other-samples<sup>1</sup></b>	Long	Count	Count of samples where processes were in a state not included in the other sample counts.
15	<b>gpmp-active</b>	Boolean		True if the GPMP was active on the virtual server at the end of the interval. False otherwise.
16	<b>other-time</b>	Long	Microseconds	Sum of time any of the virtual processors of the virtual server were in a state other than dispatched on physical processors, delayed, or idle.
17	<b>idle-time</b>	Long	Microseconds	Sum of the time any of the virtual processors of the virtual server were idle.
18	<b>virtual-cpu-count</b>	Integer	Count	Number of virtual processors defined for the virtual server, totaled across all types (dedicated/shared, CP/IFL/zIIP/zAAP)
	<p><b>Note:</b></p> <p>1. On an interval basis the GPMP samples the state of OS threads. Each sampling interval the GPMP counts the number of threads in various states. Each of these states represents a sample type. The sample types are:</p> <ul style="list-style-type: none"> <li>• CPU Using: Each sample represents a thread found actively running on a virtual processor</li> <li>• CPU delay: Each sample represents a thread waiting to be dispatched on a virtual processor</li> <li>• Page Delay: Each sample represents a thread waiting for a page fault to be resolved</li> <li>• I/O Delay: Each sample represents a threads waiting for I/O to complete</li> <li>• Other: Each sample represents processes that had no threads in one of the above states.</li> </ul> <p>The GPMP keeps a running total for each sample type. These running totals are returned in the sample metrics that are part of the virtual-server-cpu-memory-usage metrics group. The sampling interval is returned in the os-sampling-rate metric.</p>			

## Virtualization host CPU and memory metrics group

### Metric Group Name

"virtualization-host-cpu-memory-usage"

### Collection Interval

15 seconds for PR/SM, x Hyp and z/VM virtualization hosts

30 seconds for PowerVM virtualization hosts

**Applicable Managed Object Class**  
**"virtualization-host"**

The following metrics are provided in each entry of this metric group:

*Table 256. Virtualization host CPU and memory metric group*

Pos	Metric field name	Type	Units	Description
1	<b>processor-usage</b>	Integer	%	Overall CPU utilization percentage for the virtualization host (0-100%). Value for current interval. For PR/SM only includes general purpose processors.
2	<b>memory-usage</b>	Integer	%	Memory usage percentage for the virtualization host (0-100%). Value for current interval.
3	<b>network-usage</b>	Integer	%	Network utilization percentage for the virtualization host (0-100%). Value for current interval. Not available for z/VM or PR/SM, will be reported as -1.
4	<b>storage-rate</b>	Integer	Kbytes per sec	Average Kbytes transferred to disk over interval. Value for current interval. Not available for z/VM or PR/SM, will be reported as -1.
5	<b>physical-cpu-time</b>	Long	Microseconds	Physical CPU time used. Cumulative value. For PR/SM only includes general purpose processors.
6	<b>memory-used</b>	Integer	Mbytes	Current memory in use by the virtualization host.
7	<b>virt-host-management-cpu-time-used</b>	Long	Microseconds	CPU time used for virtualization host management. Cumulative value. Note for PowerVM this is the CPU used by the VIOS partition. For x Hyp this is the Linux system mode CPU time. Currently not supported for z/VM or PR/SM, will be reported as -1.
8	<b>virt-host-page-ins</b>	Long	Count	Paging activity by the virtualization host to support hypervisor management. Page reads from paging space. Cumulative value. For PowerVM this represents VIOS paging. For x Hyp it is the base Linux paging. Currently not supported for z/VM or PR/SM, will be reported as -1.
9	<b>virt-host-page-outs</b>	Long	Count	Paging activity by the virtualization host to support hypervisor management. Page written to paging space. Cumulative value. For PowerVM this represents VIOS paging. For x Hyp it is the base Linux paging. Currently not supported for z/VM or PR/SM, will be reported as -1.
10	<b>cp-cpu-time<sup>1</sup></b>	Long	Microsecond	CPU time accumulated by the general purpose processors owned by the virtualization host. Only supported for z/VM, will be reported as -1 for other hypervisor types.
11	<b>ifl-cpu-time<sup>1</sup></b>	Long	Microsecond	CPU time accumulated by the IFL processors owned by the virtualization host. Only supported for z/VM, will be reported as -1 for other virtualization host types.

Table 256. Virtualization host CPU and memory metric group (continued)

Pos	Metric field name	Type	Units	Description
12	<b>zaap-cpu-time</b> <sup>1</sup>	Long	Microsecond	CPU time accumulated by the zAAP processors owned by the virtualization host. Only supported for z/VM, will be reported as -1 for other virtualization host types.
13	<b>ziip-cpu-time</b> <sup>1</sup>	Long	Microsecond	CPU time accumulated by the zIIP processors owned by the virtualization host. Only supported for z/VM, will be reported as -1 for other virtualization host types.
14	<b>icf-cpu-time</b> <sup>1</sup>	Long	Microsecond	CPU time accumulated by the ICF processors owned by the virtualization host. Only supported for z/VM, will be reported as -1 for other virtualization host types.

**Table Notes:**

1. Provided for virtualization hosts on CPCs with SE version 2.12.0 or later; metric not present for CPCs with earlier SE versions.

## Workload service class data metrics group

This metric group reports workload performance data on a per-service-class basis. At each collection interval for a given workload, one instance of this metric group is added to the metric cache for each service class of the active policy for that workload.

### Metric Group Name

"workload-service-class"

### Collection Interval

15 seconds

### Applicable Managed Object Class

"workload"

The following metrics are provided in each entry of this metric group:

Table 257. Workload metrics group - service class data metric group

Pos	Metric field name	Type	Units	Description
1	<b>policy-activation-time</b>	Long	Timestamp	Time of the last policy activation for this workload. This is the <b>last-activation-requested-date</b> property of the currently active policy.
2	<b>service-class-name</b>	String (1-64)	—	Name of service class
3	<b>velocity-numerator</b>	Long	Microseconds	Time value used for numerator of velocity calculation. Cumulative value since last policy activation.
4	<b>velocity-denominator</b>	Long	Microseconds	Time value used for denominator of velocity calculation. Cumulative value since last policy activation.

## Network management metrics

Following are the network management metric groups.

## Virtualization host and virtual server metrics

- 1 The zManager Metrics Service provides network metrics for z Systems network resources. These metrics are collected at the virtualization host (hypervisor) network layer. The virtualization host provides a virtual LAN for network communications between the virtual servers it is hosting, and transparently virtualizes the attached physical network interfaces providing server access to the physical network. In many cases, this network virtualization function within the virtualization host is commonly referred to as the “virtual switch” or “vSwitch”. In zManager, although only certain zManager platforms allow for explicit external exposure and management of the vSwitch, the network metrics collected at this layer will be exposed for each platform. In some cases, such as **"zvm"** and **"prsm"**, the virtualization host supports multiple vSwitches, and the vSwitches are externally identified within zManager. For example, in zManager, OSX and IQD-X chpids are referred to as vSwitches. The term “vSwitch” is also used to generically describe the network functions of the virtualization host. For **"x-hyp"**, **"power-vm"**, and **"prsm"** virtualization host types, a specific vSwitch resource is not explicitly externalized for management within zManager; therefore, it is the virtualization host itself that is associated with the metrics, and, where the metric group provides a vSwitch name, this value will be “N/A”. In these cases, the virtualization host implicitly represents a single vSwitch.

The network metrics provided at the virtualization host are the following:

- Virtualization host uplink metrics: Virtual network (i.e. per VLAN ID) metrics are not provided for uplinks. Metrics collected are provided on an interval. These metrics can be collected from the Virtualization Host (vSwitch) Uplink Metric Group. Metrics for two types of uplinks are provided:
  - Real Uplinks: These are metrics which are captured between the virtualization host or virtualization host's vSwitch and the attached physical network interfaces. These metrics can show the bandwidth that the virtualization host is contributing to the IEDN.
  - Virtual Uplinks: These metrics are captured by these vSwitch ports which are not directly attached to a physical vSwitch that attaches to the IEDN. The following describe the supported virtual uplinks:
    - An IQD-X chpid that is attached to a z/VM LPAR's vSwitch bridge port. The z/VM vSwitch bridge port provides the uplink from the IQD-X vSwitch.
    - A z/VM vSwitch virtual uplink. A virtual uplink connects a vSwitch to a z/VM virtual server. In this case, traffic is forwarded to this server, which may be useful for packet collection for debugging or analysis.
- Virtualization host by vSwitch by virtual network metrics: These metrics are captured between the virtualization host and its virtual servers. In cases, such as z/VM, where a virtualization host has multiple vSwitches, the metrics are captured by virtualization host by vSwitch by virtual network. These metrics can be collected from the Virtualization Host (vSwitch) by Virtual Network Metric Group.
- Network metrics from virtualization host for each attached virtual server's virtual network adapter by virtual network (VLAN ID): These metrics can be collected from the Attached Virtual Servers Network Adapter Metric Group.

In general, for the Virtualization Host by Virtual Network Metrics Group and the Attached Virtual Servers Network Adapter Metric Group, the metrics are collected at the virtualization host level for the attached virtual server virtual network adapter by virtual network (VLAN ID). The network metrics collected at this level provide a view of the performance between the virtual switch and the virtual server, with metrics such as bytes sent and bytes received. Other metrics such as packets dropped or discarded can help to determine if problems are occurring. Metrics are collected and provided on an interval, and each metric provided is the total cumulative value, and not a delta.

Providing metrics at the virtualization host's virtual network adapters provides a level of granularity that allows for the consumer to aggregate these metrics. Use of these metrics along with configuration data provided through the zManager external API allow the client application to determine resource utilization relationships.

## Virtualization host (vSwitch) uplink metric group

This metric group provides virtualization host (hypervisor)-based network metrics for the uplink ports. The metrics provided by this group represent the uplink metrics between the vSwitch and the physical network interface. These metrics are from the perspective of the vSwitch sending to and receiving from the physical network interfaces.

In the case of z/VM, there may be multiple uplink interfaces; therefore, multiple instances of this metric group may be provided for a single z/VM virtualization host. This is also true for a PR/SM virtualization host. For a PR/SM virtualization host, there may be multiple OSX's for which the uplink information provides the metrics between the OSX and the physical network.

Metrics are collected and provided on an interval, and each metric provided is the total cumulative value, and not a delta.

### Metric Group Name

"network-virtualization-host-uplink"

### Collection Interval

30 seconds

### Applicable Managed Object Class

"virtualization-host"

The following metrics are provided in each entry of this metric group:

Table 258. Virtualization host (vSwitch) uplink metric group

Pos	Metric field name	Type	Units	Description
1	uplink-id	String		<p>Name of the uplink interface.</p> <p>The uplink names are described as follows and the naming convention will be unique based upon the <b>type</b> property of the virtualization-host object:</p> <ul style="list-style-type: none"> <li>• <b>"power-vm"</b> - The names will be the platform-defined names of the physical network interfaces, for example "entx" or "enty". Where <i>x</i> and <i>y</i> are numeric values.</li> <li>• <b>"x-hyp"</b> - The names will be the platform-defined names of the physical network interfaces, for example "entx" or "enty". Where <i>x</i> and <i>y</i> are numeric values.</li> <li>• <b>"prsm"</b> - where the uplink is an OSX, the uplink is the pchid. The name will be in the format of "OSX pchid.port"</li> <li>• <b>"prsm"</b> - where the uplink is an IQDX chpid that is connected to a z/VM vSwitch bridge port. The name will be in the form "IQDX css.chpid:zvm lpar name.vswitch name."</li> <li>• <b>"zvm"</b> - If the uplink is OSX, then the uplink will identify the OSA css and chpid, the z/VM LPAR name and virtual device number of the OSX. The name will be in the format of "OSX css.chpid:zvm lpar name.vdev number".</li> </ul>
2	uplink-type	String Enum		<p>The uplink type is:</p> <ul style="list-style-type: none"> <li>• <b>"real"</b></li> <li>• <b>"virtual"</b></li> </ul>

Table 258. Virtualization host (vSwitch) uplink metric group (continued)

Pos	Metric field name	Type	Units	Description
3	<b>vswitch-name</b>	String		Name of the vSwitch. In the case where the vSwitch is not uniquely identified for a virtual host then this will be "N/A".  "N/A" is returned for all virtualization-host objects except where the <b>type</b> property of the virtualization-host object is " <b>zvm</b> ".
4	<b>bytes-sent</b>	Long	Count	Number of bytes sent from the uplink interface to the physical network.
5	<b>bytes-received</b>	Long	Count	Number of bytes received by the uplink interface from the physical network.
6	<b>packets-sent</b>	Long	Count	Number of packets sent from the uplink interface to the physical network.
7	<b>packets-received</b>	Long	Count	Number of packets received by this uplink interface from the physical network.
8	<b>packets-sent-dropped</b>	Long	Count	Number of packets that were dropped when sending from this uplink interface to the physical network.  Packets may be dropped due to conditions related to resource constraints such as a buffer shortage.
9	<b>packets-received-dropped</b>	Long	Count	Number of packets received by this uplink interface from the physical network that were dropped.  Packets may be dropped due to conditions related to resource constraints such as a buffer shortage.
10	<b>packets-sent-discarded</b>	Long	Count	Number of packets that were discarded when sending from the uplink interface to the physical network.  Packets may be discarded due to errors associated with the packet, such as malformed packets.
11	<b>packets-received-discarded</b>	Long	Count	Number of packets received by this uplink interface that were discarded.  Packets may be discarded due to errors associated with the packet, such as malformed packets.
12	<b>multicast-packets-sent</b>	Long	Count	Number of multicast packets sent from the uplink interface to the physical network.
13	<b>multicast-packets-received</b>	Long	Count	Number of multicast packets received by the uplink interface from the physical network.
14	<b>broadcast-packets-sent</b>	Long	Count	Number of broadcast packets sent from the uplink interface to the physical network.
15	<b>broadcast-packets-received</b>	Long	Count	Number of broadcast packets received by this uplink interface from the physical network.
16	<b>interval-bytes-sent</b>	Long	Bytes	Number of bytes sent by this uplink interface to the physical network over the collection interval.
17	<b>interval-bytes-received</b>	Long	Bytes	Number of bytes received by this uplink interface from the physical network over the collection interval.
18	<b>bytes-per-second-sent</b>	Long	Bytes per Second	Number of bytes sent per second by this uplink interface to the physical network over the collection interval.

Table 258. Virtualization host (vSwitch) uplink metric group (continued)

Pos	Metric field name	Type	Units	Description
19	<b>bytes-per-second-received</b>	Long	Bytes per Second	Number of bytes received per second by this uplink interface from the physical network over the collection interval.
20	<b>mac-address</b>	String		The MAC address of this uplink, if known. If it is not known then "N/A".
21	<b>flags</b>	Long		Flags indicating the types of metrics that are reported by this uplink. The value of this field should be interpreted as a bitmask. The meaning of each bit is as follows: <ul style="list-style-type: none"> <li>• 0x02 - Byte counts are supported</li> <li>• 0x04 - Packet counts are supported</li> <li>• 0x08 - Drop counts are supported</li> <li>• 0x10 - Discard counts are supported</li> <li>• 0x20 - Multicast counts are supported</li> <li>• 0x40 - Broadcast counts are supported</li> <li>• 0x80 - Interval bytes sent and received are supported</li> </ul>

### Virtualization host (vSwitch) by virtual network metric group

The virtualization host (vSwitch) Virtual Network metric collection group provides virtual network metrics by virtualization host by vSwitch by virtual network. These metrics are collected at the virtualization host's vSwitch port level by virtual network and aggregated to provide metrics for each vSwitch by virtual network. For each virtualization host there will be an instance of these metrics for each vSwitch (in cases where multiple vSwitches are associated with the virtualization host) for each virtual network (vlan-id); therefore, there may be multiple instances within the group. In cases where a "vSwitch" is not externalized for the virtualization host, the name of the vSwitch is "N/A" and the metrics are associated with the virtualization host. These metrics are essentially an aggregation of the metrics from the "Attached virtual server network adapters metric group" on page 913.

Metrics are collected and provided on an interval, and each metric provided is the total cumulative value, and not a delta.

#### Metric Group Name

"network-vswitch-by-virtual-network"

#### Collection Interval

30 seconds

#### Applicable Managed Object Class

"virtualization-host"

The following metrics are provided in each entry of this metric group:

Table 259. Virtualization host (vSwitch) by virtual network metric group

Pos	Metric field name	Type	Units	Description
1	<b>vswitch-name</b>	String		Name of the vSwitch. In the case where the vSwitch is not uniquely identified for a virtual host then this will be "N/A".  "N/A" is returned for all virtualization-host objects except where the <b>type</b> property of the virtualization-host object is "zvm".

Table 259. Virtualization host (vSwitch) by virtual network metric group (continued)

Pos	Metric field name	Type	Units	Description
2	<b>vlan-id</b>	Integer		VLAN ID of the virtual network for which metrics are being provided. This value corresponds to the <b>vlan-id</b> property of the related Virtual Network object.  There may be cases where this field is 0 when metrics are reported and the VLAN ID is unable to be determined.
3	<b>bytes-sent</b>	Long	Count	Number of bytes sent from this virtualization host or virtualization host's vSwitch to the attached virtual servers.
4	<b>bytes-received</b>	Long	Count	Number of bytes received by this virtualization host or virtualization host's vSwitch from the attached virtual servers.
5	<b>packets-sent</b>	Long	Count	Number of packets sent from this virtualization host or virtualization host's vSwitch to the attached virtual servers.
6	<b>packets-received</b>	Long	Count	Number of packets received by this virtualization host or virtualization host's vSwitch from the attached virtual servers.
7	<b>packets-sent-dropped</b>	Long	Count	Number of packets that were dropped when sending from this virtualization host or virtualization host's vSwitch to the attached virtual servers.  Packets may be dropped due to conditions related to resource constraints such as a buffer shortage.
8	<b>packets-received-dropped</b>	Long	Count	Number of packets received by this virtualization host or virtualization host's vSwitch from the attached virtual servers that were dropped.  Packets may be dropped due to conditions related to resource constraints such as a buffer shortage.
9	<b>packets-sent-discarded</b>	Long	Count	Number of packets that were discarded when sending from this virtualization host or virtualization host's vSwitch to the attached virtual servers.  Packets may be discarded due to errors associated with the packet, such as malformed packets.
10	<b>packets-received-discarded</b>	Long	Count	Number of packets received by this virtualization host or virtualization host's vSwitch from the virtual servers that were discarded.  Packets may be discarded due to errors associated with the packet, such as malformed packets.
11	<b>multicast-packets-sent</b>	Long	Count	Number of multicast packets sent from this virtualization host or virtualization host's vSwitch to the attached virtual servers.
12	<b>multicast-packets-received</b>	Long	Count	Number of multicast packets received by this virtualization host or virtualization host's vSwitch from the attached virtual servers.
13	<b>broadcast-packets-sent</b>	Long	Count	Number of broadcast packets sent from this virtualization host or virtualization host's vSwitch to the attached virtual servers.



Table 259. Virtualization host (vSwitch) by virtual network metric group (continued)

Pos	Metric field name	Type	Units	Description
14	<b>broadcast-packets-received</b>	Long	Count	Number of broadcast packets received by this virtualization host or virtualization host's vSwitch from the attached virtual servers.
15	<b>interval-bytes-sent</b>	Long	Bytes	Number of bytes sent by this virtual switch, for the VLAN specified by <b>vlan-id</b> , over the collection interval.
16	<b>interval-bytes-received</b>	Long	Bytes	Number of bytes received by this virtual switch, for the VLAN specified by <b>vlan-id</b> , over the collection interval.
17	<b>bytes-per-second-sent</b>	Long	Bytes per Second	Number of bytes sent per second by this virtual switch, for the VLAN specified by <b>vlan-id</b> , over the collection interval.
18	<b>bytes-per-second-received</b>	Long	Bytes per Second	Number of bytes received by this virtual switch, for the VLAN specified by <b>vlan-id</b> , over the collection interval.
19	<b>flags</b>	Long		Flags indicating the types of metrics that are supported. The value of this field should be interpreted as a bitmask. The meaning of each bit is as follows: <ul style="list-style-type: none"> <li>• 0x02 - Byte counts are supported</li> <li>• 0x04 - Packet counts are supported</li> <li>• 0x08 - Drop counts are supported</li> <li>• 0x10 - Discard counts are supported</li> <li>• 0x20 - Multicast counts are supported</li> <li>• 0x40 - Broadcast counts are supported</li> <li>• 0x80 - Interval bytes sent and received are supported</li> </ul>

### Attached virtual server network adapters metric group

This metric group provides metrics for each virtual server's virtual network adapter by virtual network. Each virtual network adapter may be associated with multiple virtual networks; therefore, there may be multiple instances of these metrics within the group. These metrics are collected at the virtualization host's (vSwitch) port level for the attached virtual server virtual network adapter by virtual network. These metrics are provided from the perspective of the virtualization host or virtualization host's vSwitch sending data to and receiving data from the virtual server's network adapter. However, there is an exception to this in the case of the PowerVM virtualization host. In this case, the metrics are collected at the virtual server's guest O/S level. The metrics are included in this list for consistency, allowing for a good approximation of the data flowing over the interfaces. Therefore, the metric counters are reversed from what is provided for the metric in this group. Counters like drops and discards are from the perspective of the guest O/S for this interface and not the virtualization host.

The resource ID of the associated virtual server's network adapter is provided to allow the client application to correlate the metrics with a particular virtual server.

Metrics are collected and provided on an interval, and each metric provided is the total cumulative value, and not a delta.

#### Metric Group Name

"network-virtual-server-attached-network-adapter"

#### Collection Interval

30 seconds

#### Applicable Managed Object Class

"virtual-server"

The following metrics are provided in each entry of this metric group:

Table 260. Attached virtual server network adapters metric group

Pos	Metric field name	Type	Units	Description
1	<b>network-adapter-id</b>	String		Element identifier of the virtual servers' network adapter. This value corresponds to the <i>{network-adapter-id}</i> portion of the URI for the network adapter. The full URI can be constructed using this value together with the URI of the parent virtual server.
2	<b>vlan-id</b>	Integer		VLAN ID of the virtual network for which metrics are being provided. This value corresponds to the <b>vlan-id</b> property of the related Virtual Network object.  There may be cases where this field is 0 when zManager is unable to determine the per virtual network metrics for this interface.
3	<b>mac-address</b>	String		MAC address of this interface.
4	<b>bytes-sent</b>	Long	Bytes	Number of bytes sent to this Virtual Server network interface for the virtual network.
5	<b>bytes-received</b>	Long	Bytes	Number of bytes received from this Virtual Server network interface for the virtual network.
6	<b>packets-sent</b>	Long	Count	Number of packets sent on this interface to the Virtual Server for the virtual network.
7	<b>packets-received</b>	Long	Count	Number of packets received from this Virtual Server network interface for the virtual network.
8	<b>packets-sent-dropped</b>	Long	Count	Number of packets that were dropped when sending to this Virtual Server Network interface for the virtual network.  Packets may be dropped due to conditions related to resource constraints such as a buffer shortage.
9	<b>packets-received-dropped</b>	Long	Count	Number of packets received from this Virtual Server network interface and for this virtual network that were dropped.  Packets may be dropped due to conditions related to resource constraints such as a buffer shortage.
10	<b>packets-sent-discarded</b>	Long	Count	Number of packets that were discarded when sending to this Virtual Server network interface for this virtual network.  Packets may be discarded due to errors associated with the packet, such as malformed packets.
11	<b>packets-received-discarded</b>	Long	Count	Number of packets received from this virtual server network interface for this virtual network that were discarded.  Packets may be discarded due to errors associated with the packet, such as malformed packets
12	<b>multicast-packets-sent</b>	Long	Count	Number of multicast packets sent to this virtual server network interface for this virtual network.
13	<b>multicast-packets-received</b>	Long	Count	Number of multicast packets received from this virtual server network interface for this virtual network.
14	<b>broadcast-packets-sent</b>	Long	Count	Number of broadcast packets sent to this virtual server virtual network interface for this virtual network.

Table 260. Attached virtual server network adapters metric group (continued)

Pos	Metric field name	Type	Units	Description
15	<b>broadcast-packets-received</b>	Long	Count	Number of broadcast packets received from this virtual server virtual network interface for this virtual network.
16	<b>interval-bytes-sent</b>	Long	Bytes	Number of bytes sent by this network adapter over the collection interval.
17	<b>interval-bytes-received</b>	Long	Bytes	Number of bytes received by this network adapter over the collection interval.
18	<b>bytes-per-second-sent</b>	Long	Bytes per Second	Number of bytes sent per second by this network adapter over the collection interval.
19	<b>bytes-per-second-received</b>	Long	Bytes per Second	Number of bytes received per second by this network adapter over the collection interval.
20	<b>flags</b>	Long		Flags indicating the types of metrics that are supported by this interface. The value of this field should be interpreted as a bitmask. The meaning of each bit is as follows: <ul style="list-style-type: none"> <li>• 0x02 - Byte counts are supported</li> <li>• 0x04 - Packet counts are supported</li> <li>• 0x08 - Drop counts are supported</li> <li>• 0x10 - Discard counts are supported</li> <li>• 0x20 - Multicast counts are supported</li> <li>• 0x40 - Broadcast counts are supported</li> <li>• 0x80 - Interval bytes sent and received are supported</li> </ul>

## Optimizer network metrics

Network metrics are provided for optimizer blades. The following metric groups are provided:

- Optimizer IEDN Virtual Network Interface Metric Group
- Optimizer IEDN Physical Network Adapter Metric Group

### Optimizer IEDN virtual network interface metric group

This metric group provides metrics for an optimizer's IEDN attached network interfaces by virtual network. Currently the following optimizers are supported for this metric group: Datapower.

Metrics are collected and provided on an interval, and each metric provided is the total cumulative value, and not a delta.

#### Metric Group Name

**"network-optimizer-attached-iedn-interface"**

#### Collection Interval

30 seconds

#### Applicable Managed Object Class

**"blade"**

The following metrics are provided in each entry of this metric group:

Table 261. Optimizer IEDN virtual network interface metric group

Pos	Metric field name	Type	Units	Description
1	<b>iedn-interface-id</b>	String		Element identifier of the optimizer's IEDN interface. This value corresponds to the <i>{iedn-interface-id}</i> portion of the URI for the interface. The full URI can be constructed using this value together with the URI of the parent blade.
2	<b>vlan-id</b>	Integer		VLAN ID of the virtual network for which metrics are being provided. This value corresponds to the <b>vlan-id</b> property of the related Virtual Network object.  There may be cases where this field is 0 when zManager is unable to determine the per virtual network metrics for this interface.
3	<b>mac-address</b>	String		MAC address of this interface.
4	<b>bytes-sent</b>	Long	Count	Number of bytes sent from this interface.
5	<b>bytes-received</b>	Long	Count	Number of bytes received by this interface.
6	<b>packets-sent</b>	Long	Count	Number of packets sent from this interface.
7	<b>packets-received</b>	Long	Count	Number of packets received by this interface.
8	<b>packets-sent-dropped</b>	Long	Count	Number of packets that were dropped when sending to this Virtual Server Network interface for the virtual network.  Packets may be dropped due to conditions related to resource constraints such as a buffer shortage.
9	<b>packets-received-dropped</b>	Long	Count	Number of packets received by this interface that were dropped.  Packets may be dropped due to conditions related to resource constraints such as a buffer shortage.
10	<b>packets-sent-discarded</b>	Long	Count	Number of packets that were discarded when sending from this interface.  Packets may be discarded due to errors associated with the packet, such as malformed packets.
11	<b>packets-received-discarded</b>	Long	Count	Number of packets received by this interface that were discarded.  Packets may be discarded due to errors associated with the packet, such as malformed packets.
12	<b>multicast-packets-sent</b>	Long	Count	Number of multicast packets sent from this interface.
13	<b>multicast-packets-received</b>	Long	Count	Number of multicast packets received by this interface
14	<b>broadcast-packets-sent</b>	Long	Count	Number of broadcast packets sent from this interface.
15	<b>broadcast-packets-received</b>	Long	Count	Number of broadcast packets received to this interface.
16	<b>interval-bytes-sent</b>	Long	Bytes	Number of bytes sent by this interface over the collection interval.
17	<b>interval-bytes-received</b>	Long	Bytes	Number of bytes received by this interface over the collection interval

Table 261. Optimizer IEDN virtual network interface metric group (continued)

Pos	Metric field name	Type	Units	Description
18	<b>bytes-per-second-sent</b>	Long	Bytes per Second	Number of bytes sent per second by this interface over the collection interval.
19	<b>bytes-per-second-received</b>	Long	Bytes per Second	Number of bytes received per second by this interface over the collection interval.
20	<b>flags</b>	Long		Flags indicating each types of metrics are supported. The value of this field should be interpreted as a bitmask. The meaning of each bit is as follows: <ul style="list-style-type: none"> <li>• 0x02 - Byte counts are supported</li> <li>• 0x04 - Packet counts are supported</li> <li>• 0x08 - Drop counts are supported</li> <li>• 0x10 - Discard counts are supported</li> <li>• 0x20 - Multicast counts are supported</li> <li>• 0x40 - Broadcast counts are supported</li> <li>• 0x80 - Interval bytes sent and received are supported</li> </ul>

### Optimizer IEDN physical network adapter metric group

This metric group provides metrics for an optimizer's IEDN attached network interfaces. Currently the following optimizers are supported for this metric group: Datapower.

Metrics are collected and provided on an interval, and each metric provided is the total cumulative value, and not a delta.

#### Metric Group Name

"optimizer-physical-network-adapter"

#### Collection Interval

30 seconds

#### Applicable Managed Object Class

"blade"

This metric collection provides metrics for an optimizer's physical network adapters.

Table 262. Optimizer IEDN physical network adapter metric group

Pos	Metric field name	Type	Units	Description
1	<b>network-adapter-id</b>	String		The network-adapter-id is the name of the optimizer's physical network adapter.  For the DataPower optimizer, there are two physical network interfaces. For example, names are commonly "eth7" and "eth9".
2	<b>mac-address</b>	String		MAC address of this interface
3	<b>bytes-sent</b>	Long	Count	Number of bytes sent from this interface.
4	<b>bytes-received</b>	Long	Count	Number of bytes received by this interface
5	<b>packets-sent</b>	Long	Count	Number of packets sent from this interface
6	<b>packets-received</b>	Long	Count	Number of packets received by this interface.
7	<b>packets-sent-dropped</b>	Long	Count	Number of packets that were dropped when sending from this interface.
8	<b>packets-received-dropped</b>	Long	Count	Number of packets received by this interface that were dropped.

Table 262. Optimizer IEDN physical network adapter metric group (continued)

Pos	Metric field name	Type	Units	Description
9	<b>packets-sent-discarded</b>	Long	Count	Number of packets that were discarded when sending from this interface.  Packets may be discarded due to errors associated with the packet, such as malformed packets.
10	<b>packets-received-discarded</b>	Long	Count	Number of packets received by this interface that were discarded.  Packets may be discarded by the virtual or physical due to errors associated with the packet, such as malformed packets.
11	<b>multicast-packets-sent</b>	Long	Count	Number of multicast packets sent from this interface.
12	<b>multicast-packets-received</b>	Long	Count	Number of multicast packets received by this interface.
13	<b>broadcast-packets-sent</b>	Long	Count	Number of broadcast packets sent from this interface.
14	<b>broadcast-packets-received</b>	Long	Count	Number of broadcast packets received by this interface.
15	<b>interval-bytes-sent</b>	Long	Bytes	Number of bytes sent by this network adapter over the collection interval.
16	<b>interval-bytes-received</b>	Long	Bytes	Number of bytes received by this network adapter over the collection interval.
17	<b>bytes-per-second-sent</b>	Long	Bytes per Second	Number of bytes received per second by this network adapter over the collection interval.
18	<b>bytes-per-second-received</b>	Long	Bytes per Second	Number of bytes sent per second by this network adapter over the collection interval.
19	<b>flags</b>	Long		Flags indicating the types of metrics that are reported by this uplink. The value of this field should be interpreted as a bitmask. The meaning of each bit is as follows: <ul style="list-style-type: none"> <li>• 0x02 - Byte counts are supported</li> <li>• 0x04 - Packet counts are supported</li> <li>• 0x08 - Drop counts are supported</li> <li>• 0x10 - Discard counts are supported</li> <li>• 0x20 - Multicast counts are supported</li> <li>• 0x40 - Broadcast counts are supported</li> <li>• 0x80 - Interval bytes sent and received are supported</li> </ul>

## Physical switches

The physical switches provide the connectivity between the blades and CPCs in the intraensemble data network (IEDN).

There are two types of Ethernet switches within each zBX:

- Top-of-Rack Switches (TORs) – A pair of TORs reside in each zBX and act as a primary and backup. TORs connect the blades in the zBX to z Systems network interfaces, and to other external networking equipment, such as routers.
- Ethernet Switch Modules (ESMs) - These switches connect the blades in the zBX to the IEDN and link the blades to the TORs.

The initial configuration and setup of the physical switches are provided by zManager. Some TOR ports can be managed by the user from the zManager UI. The ESMs are not accessible for configuration changes through zManager's external management interfaces.

Metrics are provided for the following TOR port types:

- External - These ports that connect to customer's external network
- IEDN Host – These ports connect to z Systems IEDN network adapters, such as OSX.
- ISAOPT - These ports connect to ISAOPT Coordinator.
- BladeCenter ESM Ports – These ports connect to the ESM switches (other than ISAOPT).
- zBX to zBX- These ports connect the TOR in one zBX to the TOR in another zBX.

Metrics are provided for the following ESM port types:

- Uplinks to TOR – These ports connect the ESM to the TOR.
- Blade Ports – These ports connect the ESMs to the Blade network adapters.

Monitoring the physical switches can allow for determining the health and performance of the switches. For example, metrics such as dropped and discarded packets can affect the overall performance of workloads flowing through these switches. Bytes transferred metrics for ports provide the ability to determine bandwidth utilization.

## Top-of-rack switch ports metrics

This metric collection group provides metrics for the Top-of-Rack (TOR) switch ports for the TORs in each zBX.

**Metric Group Name**

"top-of-rack-switch-ports"

**Collection Interval**

120 seconds

**Applicable Managed Object Class**

"zbx"

This metric collection provides metrics for an optimizer's physical network adapters.

*Table 263. Top-of-rack switch port metrics group*

Pos	Metric field name	Type	Units	Description
1	switch-location-info	String		The location of the switch in the zBX that was set by zManager.  The switch-location-info is a 4 character cage location identifier. The first character identifies the zBX rack, and the remaining characters identify the vertical location within the rack
2	port-num	Integer		Switch port number
3	type	String Enum		<ul style="list-style-type: none"> <li>• "B" (BladeCenter ESM port): attaches to an ESM switch.</li> <li>• "H" (Host port): attaches to z Systems network adapters for the IEDN</li> <li>• "E" (External port): attaches to external networking equipment</li> <li>• "I" (ISAOPT port): for ISAOPT</li> <li>• "Z" (zBX port): attaches the zBX to another zBX</li> <li>• "T" (TOR port): attaches to the other TOR switch in the same zBX</li> </ul>

Table 263. Top-of-rack switch port metrics group (continued)

Pos	Metric field name	Type	Units	Description
4	<b>remote-partner-info</b>	String		The remote partner depends upon the port type: <ul style="list-style-type: none"> <li>• If type is "Z": attached to a TOR in another zBX, this will be the location of the top-of-rack-switch.</li> <li>• If type is "E": "N/A"</li> <li>• If type is "H": This is the <i>cpc id.pchid</i> of the attached OSX.</li> <li>• If type is "B": The format of the remote-partner-info is <i>bladecenter chassis-id_ esm location-id</i> The first 4 characters identify the blade center chassis within this zBX. The last 4 characters identify the ESM location ID</li> <li>• If type is "I": The name of the blade center chassis containing the ISAOPT blades</li> <li>• If type is "T": The name of the other TOR switch.</li> </ul>
5	<b>bytes-sent</b>	Long	Count	Number of bytes sent from this port.
6	<b>bytes-received</b>	Long	Count	Number of bytes received to this port.
7	<b>packets-sent</b>	Long	Count	Number of packets sent from this port.
8	<b>packets-received</b>	Long	Count	Number of packets received to this port.
9	<b>packets-sent-dropped</b>	Long	Count	Number of packets that were dropped when sending from this port.
10	<b>packets-received-dropped</b>	Long	Count	Number of packets received to this port that were dropped.
11	<b>packets-sent-discarded</b>	Long	Count	Number of packets that were discarded when sending from this port. Packets may be discarded due to errors associated with the packet, such as malformed packets.
12	<b>packets-received-discarded</b>	Long	Count	Number of packets received by these ports that were discarded. Packets may be discarded by the virtual or physical due to errors associated with the packet, such as malformed packets
13	<b>multicast-packets-sent</b>	Long	Count	Number of multicast packets sent from this port.
14	<b>multicast-packets-received</b>	Long	Count	Number of multicast packets received to this port.
15	<b>broadcast-packets-sent</b>	Long	Count	Number of broadcast packets sent from this port.
16	<b>broadcast-packets-received</b>	Long	Count	Number of broadcast packets received to this port.
17	<b>interval-bytes-sent</b>	Long	Bytes	Number of bytes sent by this switch port over the collection interval.
18	<b>interval-bytes-received</b>	Long	Bytes	Number of bytes received by this switch port over the collection interval.
19	<b>bytes-per-second-sent</b>	Long	Bytes per Second	Number of bytes sent per second by this switch port over the collection interval.
20	<b>bytes-per-second-received</b>	Long	Bytes per Second	Number of bytes received per second by this switch port over the collection interval.



Table 263. Top-of-rack switch port metrics group (continued)

Pos	Metric field name	Type	Units	Description
21	<b>flags</b>	Long		Flags indicating the types of metrics that are reported by this uplink. The value of this field should be interpreted as a bitmask. The meaning of each bit is as follows: <ul style="list-style-type: none"> <li>• 0x02 - Byte counts are supported</li> <li>• 0x04 - Packet counts are supported</li> <li>• 0x08 - Drop counts are supported</li> <li>• 0x10 - Discard counts are supported</li> <li>• 0x20 - Multicast counts are supported</li> <li>• 0x40 - Broadcast counts are supported</li> <li>• 0x80 - Interval bytes sent and received are supported</li> </ul>

## ESM switch port metrics

This metric group provides metrics for the ESM switch ports for each ESM in each zBX.

### Metric Group Name

"ethernet-switch-module-ports"

### Collection Interval

120 seconds

### Applicable Managed Object Class

"zbx"

This metric collection provides metrics for an optimizer's physical network adapters.

Table 264. Optimizer IEDN physical network adapter metric group

Pos	Metric field name	Type	Units	Description
1	<b>switch-location-info</b>	String		The location of the switch in the zBX that was set by zManager.  The switch-location-info is a 4 character cage location identifier. The first character identifies the zBX rack, and the remaining characters identify the vertical location within the rack.
2	<b>port-num</b>	Integer		Switch port number
3	<b>type</b>	String Enum		<ul style="list-style-type: none"> <li>• "I" (Internal port): This port is connected to a Blade.</li> <li>• "E" (External port): This port is connected to a TOR.</li> </ul>
4	<b>remote-partner-info</b>	String		Information about the remote element connected to the port. The value depends on the <b>type</b> of the port. In the case where the port type is: <ul style="list-style-type: none"> <li>• If type is "E": This field is the name of the top-of-rack-switch connected to this ESM port. This TOR resides in the same zBX as the ESM.</li> <li>• If type is "I": This field is the name of the attached blade. The blade resides in the same zBX and blade center chassis as the ESM.</li> </ul>
5	<b>bytes-sent</b>	Long	Count	Number of bytes sent from this port.
6	<b>bytes-received</b>	Long	Count	Number of bytes received to this port.
7	<b>packets-sent</b>	Long	Count	Number of packets sent from this port.

Table 264. Optimizer IEDN physical network adapter metric group (continued)

Pos	Metric field name	Type	Units	Description
8	<b>packets-received</b>	Long	Count	Number of packets received to this port.
9	<b>packets-sent-dropped</b>	Long	Count	Number of packets that were dropped when sending from this port.
10	<b>packets-received-dropped</b>	Long	Count	Number of packets received to this port that were dropped.
11	<b>packets-sent-discarded</b>	Long	Count	Number of packets that were discarded when sending from this port. Packets may be discarded due to errors associated with the packet, such as malformed packets.
12	<b>packets-received-discarded</b>	Long	Count	Number of packets received by these ports that were discarded. Packets may be discarded by the virtual or physical due to errors associated with the packet, such as malformed packets.
13	<b>multicast-packets-sent</b>	Long	Count	Number of multicast packets sent from this port.
14	<b>multicast-packets-received</b>	Long	Count	Number of multicast packets received to this port.
15	<b>broadcast-packets-sent</b>	Long	Count	Number of broadcast packets sent from this port.
16	<b>broadcast-packets-received</b>	Long	Count	Number of broadcast packets received to this port.
17	<b>interval-bytes-sent</b>	Long	Bytes	Number of bytes sent by this switch port over the collection interval.
18	<b>interval-bytes-received</b>	Long	Bytes	Number of bytes received by this switch port over the collection interval.
19	<b>bytes-per-second-sent</b>	Long	Bytes per Second	Number of bytes sent per second by this switch port over the collection interval.
20	<b>bytes-per-second-received</b>	Long	Bytes per Second	Number of bytes received per second by this switch port over the collection interval.
21	<b>flags</b>	Long		<p>Flags indicating the types of metrics that are reported by this uplink. The value of this field should be interpreted as a bitmask. The meaning of each bit is as follows:</p> <ul style="list-style-type: none"> <li>• 0x02 - Byte counts are supported</li> <li>• 0x04 - Packet counts are supported</li> <li>• 0x08 - Drop counts are supported</li> <li>• 0x10 - Discard counts are supported</li> <li>• 0x20 - Multicast counts are supported</li> <li>• 0x40 - Broadcast counts are supported</li> <li>• 0x80 - Interval bytes sent and received are supported</li> </ul>

## Appendix A. XML document structure of a performance policy

To import a performance policy for a workload resource group through the HMC **Workload Resource Group Details** task or the **Import Performance Policy** operation, you must first create an XML document that defines the policy elements. Use this topic to create a properly structured XML document to import.

The XML document starts with the **WorkloadPerformancePolicy** element, which contains all the elements of a workload resource group performance policy. The following sample shows the correct structure of the major elements in the **WorkloadPerformancePolicy** element:

```
<WorkloadPerformancePolicy
xmlns="http://www.ibm.com/PPM/WorkloadPerformancePolicy">
  <Name> SampleWorkloadPerformancePolicy </Name>
  <Description> Sample performance policy for a workload </Description>
  <Version> 3.00.00 </Version>
  <UI> PPM Editor </UI>
  <WorkloadImportance> High </WorkloadImportance>
  <ServiceClasses> ... </ServiceClasses>
</WorkloadPerformancePolicy>
```

The following table describes the major elements in the **WorkloadPerformancePolicy** element.

Table 265. Performance policy XML elements

Element name	Rqd/Opt	Description
<b>Name</b>	Required	The display name specified for the performance policy. All <b>Name</b> elements must be up to 64 characters long, consisting of alphanumeric characters, blanks, periods, underscores, or dashes. Names must start with an alphabetic character and end with an alphabetic character, numeric character, underscore, period, or dash. Names must be unique to other existing performance policies in the workload resource group.
<b>Description</b>	Optional	Arbitrary text describing the performance policy in up to 256 characters.
<b>Version</b>	Required	A version number, a release number, and a level number, which are each separated by a period. The version number must be between 1 and 99. The release and level numbers must be between 0 and 99.
<b>UI</b>	Required	The name of the product that last edited the XML document. <b>UI</b> has the same length and content restrictions as the <b>Name</b> element.
<b>WorkloadImportance</b>	Required	The importance value assigned to the performance policy, which is one of the following: <b>highest</b> , <b>high</b> , <b>medium</b> , <b>low</b> , or <b>lowest</b>
<b>ServiceClasses</b>	Required	The <b>ServiceClasses</b> element contains one or more service classes, as defined in "XML structure of a <b>ServiceClasses</b> element."

### XML structure of a ServiceClasses element

The **ServiceClasses** element contains one or more **ServiceClass** elements that set the priority of and classify resources within a performance policy. A service class is a group of virtual servers that has the same service goals or performance objectives, resource requirements, or availability requirements.

The following sample shows the structure of a **ServiceClass** element within the **ServiceClasses** element:

```
<ServiceClasses>
  <ServiceClass>
    <Name> Web Hot </Name>
    <Description> Critical web transactions </Description>
    <Type> Server </Type>
```

Level 00a

```
<Goal> ... </Goal>  
  
<RuleBuilderElement> ... </RuleBuilderElement>  
  
</ServiceClasses>  
</ServiceClasses>
```

**ServiceClass** elements are positional. After you import and activate the policy, zManager searches service classes from low-ordered to high-ordered service class during classification to find a match.

The following table describes the elements in the **ServiceClass** element.

Table 266. Performance policy XML: Elements in a **ServiceClass** element

Element name	Rqd/Opt	Description
<b>Name</b>	Required	The name specified for the service class. The <b>Name</b> element must be a valid identifier up to 64 characters long, consisting of alphanumeric characters, blanks, periods, underscores, or dashes. Names must start with an alphabetic character and end with an alphabetic character, numeric character, underscore, period, or dash. Names must be unique to other existing service classes in the performance policy.
<b>Description</b>	Optional	Arbitrary text describing the service class in up to 256 characters.
<b>Type</b>	Required	This element identifies the resource associated with the service class. The <b>Type</b> value must be server to target specific virtual servers.
<b>Goal</b>	Required	The <b>Goal</b> element identifies the type of performance goal required for this service class, as described in "XML structure of a Goal element."
<b>RuleBuilderElement</b>	Required	The <b>RuleBuilderElement</b> contains the classification attributes for the service class, which are described in "XML structure of a RuleBuilderElement " on page 925.

## XML structure of a Goal element

The **Goal** element identifies the type of performance goal, which is either a velocity goal or a discretionary goal. If you code the **Velocity** element, you also need to code the **Importance** element. The following samples show how to code the XML for each type of goal:

```
<Goal>  
  <Discretionary/>  
</Goal>
```

or

```
<Goal>  
  <Velocity>  
    <Importance> Medium </Importance>  
    <Level> Fast </Level>  
  </Velocity>  
</Goal>
```

For every **Goal** element, you must code either the **Discretionary** element or the **Velocity** element. If you code the **Velocity** element, you must code the **Importance** and **Level** elements as described in the following table.

Table 267. Performance policy XML: Elements required for a **Velocity** element

Element name	Rqd/Opt	Description
<b>Importance</b>	Required	This field identifies the business importance level assigned to the service class, which must be one of the following: <b>highest</b> , <b>high</b> , <b>medium</b> , <b>low</b> , or <b>lowest</b>

Table 267. Performance policy XML: Elements required for a **Velocity** element (continued)

Element name	Rqd/Opt	Description
Level	Required	This field identifies the velocity goal value of the service class, which must be one of the following: <b>fastest</b> , <b>fast</b> , <b>moderate</b> , <b>slow</b> , or <b>slowest</b>

## XML structure of a RuleBuilderElement

The **RuleBuilderElement** represents a classification rule for the service class. Each service class within a performance policy has one or more classification rules that identify how hardware or software elements of a workload resource group are associated with the service class. So each **ServiceClass** element can have one or more **RuleBuilderElements**.

Each classification rule consists of one or more conditions that enable zManager to associate a service class to incoming work. A condition is a group containing a filter type, filter operator and filter value. To fully represent a classification rule, a **RuleBuilderElement** consists of one or more **Filter** elements, each containing one **FilterType**, one **FilterOperation**, and one **FilterValue** element.

Classification rules within a service class are positional. Assign any **RuleBuilderElement** containing a wildcard filter value to a high-order service class, and assign any **RuleBuilderElement** containing a specific filter value to a low-order service class.

The following sample shows the XML structure of a simple **RuleBuilderElement**:

```
<RuleBuilderElement>
  <RuleBuilderElementType> Rule </RuleBuilderElementType>
  <Filter>
    <FilterType> Virtual Server Name </FilterType>
    <FilterOperation> stringMatch </FilterOperation>
    <FilterValue> WebSales* </FilterValue>
  </Filter>
</RuleBuilderElement>
```

This **RuleBuilderElement** classifies all virtual servers with the string “WebSales” in the name as resources to be managed according to the goals set for the service class.

All **RuleBuilderElements** must begin with **RuleBuilderElementType**, which identifies the type of classification rule as one of the following:

### Rule

Defines a simple filter that resolves to true or false based on its filter pattern compared to a specified virtual server attribute.

### And

Defines a complex set of two rules; during classification, both of the two rules imbedded within this **RuleBuilderElement** must be true for the virtual server to be managed according to the goals for this service class

**Or** Defines a complex set of two rules; during classification, only one of the two rules imbedded within this **RuleBuilderElement** must be true for the virtual server to be managed according to the goals for this service class

If you code a **RuleBuilderElementType** value of **and** or **or**, exactly two **RuleBuilderElements** must be nested inside this **RuleBuilderElement** object so they can be logically compared. The following sample shows nested **RuleBuilderElements** for a complex set of rules that are equivalent to the expression Rule 1 **and** (Rule 2 **or** Rule 3):

```
<RuleBuilderElement>
  <RuleBuilderElementType>And</RuleBuilderElementType>
```

```

<RuleBuilderElement>
  <RuleBuilderElementType>Rule</RuleBuilderElementType>
  <Filter> ... filters for rule 1 ... </Filter>
</RuleBuilderElement>

<RuleBuilderElement>
  <RuleBuilderElementType>0r</RuleBuilderElementType>

  <RuleBuilderElement>
    <RuleBuilderElementType>Rule</RuleBuilderElementType>
    <Filter> ... filters for rule 2 ... </Filter>
  </RuleBuilderElement>

  <RuleBuilderElement>
    <RuleBuilderElementType>Rule</RuleBuilderElementType>
    <Filter> ... filters for rule 3 ... </Filter>
  </RuleBuilderElement>

</RuleBuilderElement>

</RuleBuilderElement>

```

Within every **Filter** element, you must code one **FilterType**, one **FilterOperation**, and one **FilterValue** element. These elements are described in the following table.

Table 268. Performance policy XML: Elements in a **Filter** element

Element name	Rqd/Opt	Description
<b>FilterType</b>	Required	<p>This element identifies the virtual server attribute to be used during classification. Valid values are:</p> <p><b>Hostname</b> The host name of the virtual server. To use this value, a guest platform management provider must be running on the operating system on the virtual server.</p> <p><b>Virtual Server Name</b> The virtual server ID for z/VM and x Hyp or the logical partition (LPAR) ID for PowerVM and PR/SM. This ID is the same as the name used on the <b>Virtual Servers</b> tab on the <b>Ensemble Management</b> window in the primary HMC.</p> <p><b>OS Type</b> The type of operating system, such as AIX® or Linux, that is running on the virtual server. You can select this type only if the virtual server definition contains the OS type. To use this value, a guest platform management provider must be running on the operating system on the virtual server.</p> <p><b>OS Level</b> The release level of the operating system running on the virtual server. To use this value, a guest platform management provider must be running on the operating system on the virtual server.</p> <p><b>OS Name</b> The name of the operating system image running on the virtual server as known to its operating system. To use this value, a guest platform management provider must be running on the operating system on the virtual server.</p>

Table 268. Performance policy XML: Elements in a **Filter** element (continued)

Element name	Rqd/Opt	Description
<b>FilterOperation</b>	Required	<p>This element identifies the logical filter operation, which must be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>stringMatch</b> – the filter value must match the property defined by the filter type</li> <li>• <b>stringNotMatch</b> – the filter value must not match the property defined by the filter type</li> </ul>
<b>FilterValue</b>	Required	<p>The value to be used for classification. Possible values vary based on the value specified for the <b>FilterType</b> element. A filter value cannot be longer than 255 characters. A filter value must be a regular expression that can include one period (.) as a substitute for one character and one wildcard (*) as a substitute for multiple characters. This wildcard can be used only at the end of an expression.</p> <ul style="list-style-type: none"> <li>• For filter type <b>Hostname</b>, provide a fully qualified host name that consists of the host name and the TCP domain name. The TCP domain name specifies a group of systems that share a common suffix (domain name). For example, given a fully qualified host name of <code>server1.us.ibm.com</code>: <ul style="list-style-type: none"> <li>– <code>server1</code> is the host name</li> <li>– <code>us.ibm.com</code> is the TCP domain name</li> </ul> <p>The fully qualified host name <code>server1.us.ibm.com</code> illustrates a character-based format, but you can use a numerical-based name. The numerical-based host name is not the same as the system's IP address.</p> </li> <li>• For filter type <b>OS Level</b>, provide the release level of the operating system.</li> <li>• For filter type <b>OS Name</b>, provide the LPAR name or virtual machine ID.</li> <li>• For filter type <b>OS Type</b>, provide the name of the operating system.</li> <li>• For filter type <b>Virtual Server Name</b>, provide the virtual server name shown in the <b>Virtual Server Details</b> window in the HMC.</li> </ul>

## Sample XML document for a performance policy

The following sample illustrates a correctly structured XML document for a workload resource group performance policy.

This performance policy consists of two service classes:

### SC1

This service class sets a performance goal of slow velocity and medium importance for virtual servers with attributes that match both of the following classification rules:

- The virtual server name is "Server1", and
- The hostname is "abc.pok.ibm.com"

### SC2

This service class sets a performance goal of slow velocity and medium importance for virtual servers with a virtual server name that matches "Lpar1".

---

```
<?xml version="1.0" encoding="UTF-8"?>
<WorkloadPerformancePolicy
  xmlns="http://www.ibm.com/PPM/WorkloadPerformancePolicy">

  <Name>Workload Policy name</Name>
  <Description>Workload performance policy for sample workload</Description>
  <Version>1.00.00</Version>
  <UI>PPM Editor</UI>
  <WorkloadImportance>High</WorkloadImportance>

  <ServiceClasses>

    <ServiceClass>
      <Name>SC1</Name>
      <Description>SC1 Description</Description>
      <Type>Server</Type>

      <RuleBuilderElement>
        <RuleBuilderElementType>And</RuleBuilderElementType>
        <RuleBuilderElement>
          <RuleBuilderElementType>Rule</RuleBuilderElementType>
          <Filter>
            <FilterType>Virtual Server Name</FilterType>
            <FilterOperation>stringMatch</FilterOperation>
            <FilterValue>Server1</FilterValue>
          </Filter>
        </RuleBuilderElement>
        <RuleBuilderElement>
          <RuleBuilderElementType>Rule</RuleBuilderElementType>
          <Filter>
            <FilterType>Hostname</FilterType>
            <FilterOperation>stringMatch</FilterOperation>
            <FilterValue>abc.pok.ibm.com</FilterValue>
          </Filter>
        </RuleBuilderElement>
      </RuleBuilderElement>

      <Goal>
        <Importance>Medium</Importance>
        <Velocity>Slow</Velocity>
      </Goal>
    </ServiceClass>

  </ServiceClasses>
</WorkloadPerformancePolicy>
```

---

Figure 445. Policy XML example, Part 1



---

```
<ServiceClass>
  <Name>SC2</Name>
  <Description>Service Class for virtual server</Description>
  <Type>Server</Type>

  <RuleBuilderElement>
    <RuleBuilderElementType>Rule</RuleBuilderElementType>
    <Filter>
      <FilterType>Virtual Server Name</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>Lpar1</FilterValue>
    </Filter>
  </RuleBuilderElement>

  <Goal>
    <Importance>Medium</Importance>
    <Velocity>Slow</Velocity>
  </Goal>
</ServiceClass>

</ServiceClasses>
</WorkloadPerformancePolicy>
```

---

Figure 446. Policy XML example, Part 2

|



## Appendix B. Enum values for a type of managed objects within User Roles

The valid Enum values that can be used to specify a class of managed objects within User Role objects are listed in the following table:

Table 269. Enum values for a class of managed objects

Enum for type name	Type
alternate-console	A Hardware Management Console (HMC) paired with the primary HMC to provide redundancy in an ensemble
blade-center	BladeCenter chassis which contains blades (DataPower XI50z, POWER, and/or System x)
cpc	CPC defined to the console via Add Object Definition
datapower-device	IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z) blade used to optimize XML and Web services processing
default-workload	System defined workload resource group used as a default in an ensemble
eckd-storage-resource	ECKD Storage Resource
ensemble	Ensemble object representing a collection of one or more nodes (CPC or zBX Node) managed as a single logical virtualized system
fcp-storage-resource	FCP Storage Resource
lpar-image	Partition where an operating system or coupling facility control code (CFCC) is run
p-blade	Select IBM POWER blade running AIX applications in an ensemble
p-virtual-server	Virtual server defining resources (processor, memory, and I/O) on a POWER blade
pattern-match-group	Group containing objects whose name matches a specified pattern
system-manual-definition	Template used to manually define a system object to the Hardware Management Console
undefined-cpc	CPC that is discovered in the console domain but not defined to the console
undefined-zbx-node	zBX Node that is discovered in the console domain but not defined to the console

Table 269. Enum values for a class of managed objects (continued)

Enum for type name	Type
user-defined-group	A group of objects defined by a user for organizing objects with a similar purpose or to assign different activation profiles
virtual-network	A group of virtual servers network interfaces that are associated with a VLAN ID
workload-element-group	Group of virtual servers organized for redundancy in a workload resource group
workload-resource-group	A collection of virtual servers associated with a set of policies that define performance and availability goals
x-blade	Select System x blade running 64-bit Windows or Linux applications in an ensemble
x-virtual-server	Virtual server defining resources (processor, memory, and I/O) on a System x blade
zbx-node	zBX Node defined to the console via Add Object Definition
zvm-fcp-storage-resource	z/VM FCP Storage Resource
zvm-virtual-machine	Virtualized processor, communications, storage, networking, and I/O resources running on the z/VM hypervisor

## Appendix C. Enum values for the User Role object

The valid Enum values for the **name** property of User Role objects with a **type** of "system-defined" are listed in the following table:

Table 270. Enum values for the **name** property of User Role objects with a **type** of "system-defined"

Enum for system-defined User Role name	System-defined User Role	Valid as an associated User Role
hmc-access-administrator-tasks	Access Administrator Tasks	Yes
hmc-advanced-operator-tasks	Advanced Operator Tasks	Yes
hmc-all-resources	All Resources	No
hmc-all-system-managed-objects	All System Managed Objects	No
hmc-bladecenter-objects	BladeCenter Objects	No
hmc-cim-actions	CIM Actions	Yes
hmc-defined-system-managed-objects	Defined System Managed Objects	No
hmc-dpxi50z-blade-objects	DPXI50z Blade Objects	No
hmc-energy-administrator-tasks	Energy Administrator Tasks	Yes
hmc-ensemble-administrator-tasks	Ensemble Administrator Tasks	Yes
hmc-ensemble-object	Ensemble Object	No
hmc-ibm-blade-objects	IBM Blade Objects	No
hmc-ibm-blade-virtual-server-objects	IBM Blade Virtual Server Objects	No
hmc-operator-tasks	Operator Tasks	Yes
hmc-policy-administrator-tasks	Policy Administrator Tasks	Yes
hmc-policy-operator-tasks	Policy Operator Tasks	Yes
hmc-service-representative-tasks	Service Representative Tasks	Yes
hmc-storage-resource-administrator-tasks	Storage Resource Administrator Tasks	Yes
hmc-storage-resource-objects	Storage Resource Objects	No
hmc-system-programmer-tasks	System Programmer Tasks	Yes
hmc-virtual-network-administrator-tasks	Virtual Network Administrator Tasks	Yes
hmc-virtual-network-objects	Virtual Network Objects	No
hmc-virtual-server-administrator-tasks	Virtual Server Administrator Tasks	Yes
hmc-virtual-server-operator-tasks	Virtual Server Operator Tasks	Yes
hmc-workload-administrator-tasks	Workload Administrator Tasks	Yes
hmc-workload-objects	Workload Objects	No
hmc-z-vm-virtual-machine-objects	z/VM Virtual Machine Objects	No
hmc-z-vm-virtual-machine-tasks	z/VM Virtual Machine Tasks	Yes



## Appendix D. Enum values for the Task object

The valid Enum values for the **name** property of Task objects are listed in the following table:

Table 271. Enum values for the **name** property of Task objects

Enum for task name	Task	view-only-mode supported
access-removable-media	Access Removable Media	false
activate-availability-policy	Activate Availability Policy	false
activate-base	Activate	false
activate-cim	Activate (CIM)	false
activate-performance-policy	Activate Performance Policy	false
activate-policies	Activate Policies	false
add-hosts-to-virtual-network	Add Hosts to Virtual Network	false
add-member-to-ensemble	Add Member to Ensemble	false
add-object-definition-z	Add Object Definition	false
add-storage-resource	Add Storage Resource	false
add-storage-resource-to-group	Add Storage Resource to Group	false
addmember-cim	AddMember (CIM)	false
alternate-details	Alternate Details	false
alternate-support-element	Alternate Support Element	false
alternate-support-element-engineering-changes	Alternate Support Element Engineering Changes (ECs)	false
analyze-console-internal-code	Analyze Console Internal Code	false
applybootconfig-cim	ApplyBootConfig (CIM)	false
archive-security-logs-base	Archive Security Logs	false
archive-security-logs-z	Archive Security Logs	false
audit-and-log-management	Audit and Log Management	false
authorize-internal-code-changes	Authorize Internal Code Changes	false
automatic-activation	Automatic Activation	false
availability-policy-details	Availability Policy Details	false
availability-status-event-report	Availability status Event Report	false
availability-status-monitor	Availability Status Monitor	false
backup-critical-console-data	Backup Critical Console Data	false
backup-critical-data	Backup Critical Data	false
block-automatic-licensed-internal-code-change-installation	Block Automatic Licensed Internal Code Change Installation	false
certificate-management	Certificate Management	false
change-console-internal-code	Change Console Internal Code	false
change-internal-code	Change Internal Code	false
change-lpar-controls	Change LPAR Controls	false
change-lpar-group-controls	Change LPAR Group Controls	false

Table 271. Enum values for the **name** property of Task objects (continued)

Enum for task name	Task	view-only-mode supported
change-lpar-i-o-priority-queuing	Change LPAR I/O Priority Queuing	false
change-object-definition	Change Object Definition	false
change-object-options	Change Object Options	false
choose-z-vm-virtual-machines-to-manage	Choose z/VM Virtual Machines to Manage	false
choose-z-vm-virtual-servers-to-manage	Choose z/VM Virtual Servers to Manage	false
clearlog-cim	ClearLog (CIM)	false
compare-access-lists	Compare Access Lists	false
concurrent-upgrade-engineering-changes	Concurrent Upgrade Engineering Changes (ECs)	false
configure-3270-emulators	Configure 3270 Emulators	false
configure-backup-settings	Configure Backup Settings	false
configure-channel-path-on-off	Configure Channel Path On/Off	true
configure-data-replication	Configure Data Replication	false
configure-top-of-rack-tor-switch	Configure Top-of-rack (TOR) Switch	false
configuretimeservers-cim	ConfigureTimeServers (CIM)	false
console-default-user-settings	Console Default User Settings	false
console-messenger	Console Messenger	false
copy-console-logs-to-media	Copy Console Logs to Media	false
create-ensemble	Create Ensemble	false
create-a-new-virtual-network	Create a New Virtual Network	false
create-welcome-text	Create Welcome Text	false
customer-information	Customer Information	false
customize-activity-profiles	Customize Activity Profiles	false
customize-api-settings	Customize API Settings	false
customize-automatic-logon	Customize Automatic Logon	false
customize-console-date-time	Customize Console Date/Time	false
customize-console-services	Customize Console Services	false
customize-customer-information	Customize Customer Information	false
customize-delete-activation-profiles	Customize/Delete Activation Profiles	false
customize-network-settings-base	Customize Network Settings	false
customize-network-settings-z	Customize Network Settings	false
customize-outbound-connectivity	Customize Outbound Connectivity	false
customize-product-engineering-access	Customize Product Engineering Access	false
customize-remote-service	Customize Remote Service	false
customize-scheduled-operations	Customize Scheduled Operations	false
customize-scheduled-operations-c	Customize Scheduled Operations	false
customize-support-element-date-time	Customize Support Element Date/Time	false
deactivate-base	Deactivate	false



Table 271. Enum values for the **name** property of Task objects (continued)

Enum for task name	Task	view-only-mode supported
deactivate-cim	Deactivate (CIM)	false
delete-availability-policy	Delete Availability Policy	false
delete-element-group	Delete Element Group	false
delete-ensemble	Delete Ensemble	false
delete-performance-policy	Delete Performance Policy	false
delete-virtual-network	Delete a Virtual Network	false
delete-virtual-server	Delete Virtual Server	false
delete-workload-resource-group	Delete Workload Resource Group	false
discover-storage-resources	Discover Storage Resources	false
domain-security	Domain Security	false
edit-frame-layout	Edit Frame Layout	false
edit-virtual-network-properties	Edit Virtual Network Properties	false
edit-the-vmrm-active-configuration-file	Edit the VMRM Active Configuration File	false
element-group-details	Element Group Details	false
enable-electronic-service-agent	Enable Electronic Service Agent™	false
enable-ftp-access-to-mass-storage-media	Enable FTP Access to Mass Storage Media	false
enable-i-o-priority-queuing	Enable I/O Priority Queuing	false
engineering-changes	Engineering Changes (ECs)	false
ensemble-details	Ensemble Details	false
ensemble-details-performance-management	Ensemble Details Performance Management	false
ensemble-management-guide	Ensemble Management Guide	false
environmental-efficiency-statistics	Environmental Efficiency Statistics	false
export-performance-policy	Export Performance Policy	false
export-world-wide-port-name-list	Export World Wide Port Name List	false
exportsettingdata-cim	ExportSettingData (CIM)	false
externalinterrupt-cim	ExternalInterrupt (CIM)	false
failover-cim	Failover (CIM)	false
fibre-channel-analyzer	Fibre Channel Analyzer	false
file-dialog	File Dialog	false
firmware-details	Firmware Details	false
format-media	Format Media	false
grouping	Grouping	false
hardware-messages	Hardware Messages	true
hops-report	Hops Report	false
hypervisor-report	Hypervisor Report	false
image-details	Image Details	false
import-performance-policy	Import Performance Policy	false

Table 271. Enum values for the **name** property of Task objects (continued)

Enum for task name	Task	view-only-mode supported
import-storage-access-list	Import Storage Access List	false
importsettingdata-cim	ImportSettingData (CIM)	false
initiate-hypervisor-dump	Initiate Hypervisor Dump	false
initiate-virtual-server-dump	Initiate Virtual Server Dump	false
initiate-zvm-management-guest-dump	Initiate zVM Management Guest Dump	false
input-output-configuration-save-and-restore	Input/output (I/O) Configuration Save and Restore	false
installation-complete-report	Installation Complete Report	false
integrated-3270-console	Integrated 3270 Console	false
integrated-ascii-console	Integrated ASCII Console	false
load	Load	false
load-balancing-report	Load Balancing Report	false
load-cim	Load (CIM)	false
load-from-removable-media-or-server	Load from Removable Media or Server	false
logical-processor-add	Logical Processor Add	false
maintain-z-vm-profiles	Maintain z/VM Profiles	false
maintain-z-vm-prototypes	Maintain z/VM Prototypes	false
maintain-z-vm-virtual-machines	Maintain z/VM Virtual Machines	false
maintain-z-vm-volume-space	Maintain z/VM Volume Space	false
manage-alternate-hmc	Manage Alternate HMC	false
manage-datapower-xi50z	Manage DataPower XI50z	false
manage-flash-allocation	Manage Flash allocation	false
manage-ldap-server-definitions	Manage LDAP Server Definitions	false
manage-password-rules	Manage Password Rules	false
manage-print-screen-files	Manage Print Screen Files	false
manage-remote-connections	Manage Remote Connections	false
manage-remote-support-requests	Manage Remote Support Requests	false
manage-ssh-keys	Manage SSH Keys	false
manage-storage-resources	Manage Storage Resources	false
manage-user-patterns	Manage User Patterns	false
manage-user-roles	Manage User Roles	false
manage-user-templates	Manage User Templates	false
manage-users	Manage Users	false
manage-virtual-networks	Manage Virtual Networks	false
manage-virtual-switches	Manage Virtual Switches	false
manage-web-services-api-logs	Manage Web Services API Logs	false
manage-workload-element-groups	Manage Workload Element Groups	false
manage-zbx-move	Manage zBX Move	false
migrate-virtual-server	Migrate Virtual Server	false

Table 271. Enum values for the **name** property of Task objects (continued)

Enum for task name	Task	view-only-mode supported
monitor-system-events	Monitor System Events	false
monitor-system-events-service-class-pi-monitor	Monitor System Events - Service Class PI Monitor	false
monitor-system-events-virtual-server-cpu-utilization-monitor	Monitor System Events - Virtual Server CPU Utilization Monitor	false
monitors-dashboard	Monitors Dashboard	false
mount-virtual-media	Mount Virtual Media	false
network-diagnostic-information	Network Diagnostic Information	false
network-monitors-dashboard	Network Monitors Dashboard	false
new-availability-policy	New Availability Policy	false
new-element-group	New Element Group	false
new-performance-policy	New Performance Policy	false
new-role-wizard	New User Role Wizard	false
new-template-wizard	New User Template Wizard	false
new-user-wizard	New User Wizard	false
new-virtual-server	New Virtual Server	false
new-virtual-server-based-on	New Virtual Server Based n	false
new-workload-resource-group	New Workload Resource Group	false
object-locking-settings	Object Locking Settings	false
offload-virtual-retain-data-to-removable-media	Offload Virtual RETAIN Data to Removable Media	false
open-graphical-console-novnc	Open Graphical Console	false
open-text-console	Open Text Console	false
operating-system-messages	Operating System Messages	true
osa-advanced-facilities	OSA Advanced Facilities	true
perform-a-console-repair-action	Perform a Console Repair Action	false
perform-problem-analysis	Perform Problem Analysis	false
perform-transfer-rate-test	Perform Transfer Rate Test	false
performance-policy-details	Performance Policy Details	false
product-engineering-directed-changes	Product Engineering Directed Changes	false
psw-restart	PSW Restart	false
pswrestart-cim	PSWRestart (CIM)	false
reassign-hardware-management-console	Reassign Hardware Management Console	false
reassign-i-o-path	Reassign I/O Path	false
reboot-support-element	Reboot Support Element	false
rebuild-vital-product-data	Rebuild Vital Product Data	false
remote-hardware-management-console	Remote Hardware Management Console	false
remote-service	Remote Service	false
remove-hosts-from-virtual-network	Remove Hosts from Virtual Network	false

Table 271. Enum values for the **name** property of Task objects (continued)

Enum for task name	Task	view-only-mode supported
remove-member-from-ensemble	Remove Member from Ensemble	false
remove-object-definition-z	Remove Object Definition	false
remove-storage-resource	Remove Storage Resource	false
remove-storage-resource-from-group	Remove Storage Resource from Group	false
removemember-cim	RemoveMember (CIM)	false
repair-moved-hosts	Repair Moved Hosts	false
repair-virtual-network	Repair Virtual Network	false
report-a-problem-base	Report a Problem	false
report-a-problem-z	Report a Problem	false
requeststatechange-cim	RequestStateChange (CIM)	false
reset-clear	Reset Clear	false
reset-normal	Reset Normal	false
resetclear-cim	ResetClear (CIM)	false
resetnormal-cim	ResetNormal (CIM)	false
restart-zvm-management-guest	Restart zVM Management Guest	false
retrieve-backup-file-from-ftp	Retrieve Backup File from FTP	false
retrieve-internal-code	Retrieve Internal Code	false
save-legacy-upgrade-data	Save Legacy Upgrade Data	false
save-restore-customizable-console-data	Save/Restore Customizable Console Data	false
save-upgrade-data	Save Upgrade Data	false
scsidump-cim	SCSIDump (CIM)	false
scsiload-cim	SCSILoad (CIM)	false
sendoscommand-cim	SentOSCommand (CIM)	false
service-class-resource-adjustments-report	Service Class Resource Adjustments Report	false
service-classes-report	Service Classes Report	false
service-status	Service Status	false
service-status	Service Status	false
set-power-cap	Set Power Cap	false
set-power-saving	Set Power Saving	false
setbootconfigrole-cim	SetBootConfigRole (CIM)	false
setbootconfigusage-cim	SetBootConfigUsage (CIM)	false
shutdown-or-restart	Shutdown or Restart	false
single-object-operations	Single Object Operations	false
single-step-console-internal-code	Single Step Console Internal Code	false
single-step-internal-code-changes	Single Step Internal Code Changes	false
special-code-load	Special Code Load	false
start-all	Start All	false
start-cim	Start (CIM)	false

Table 271. Enum values for the **name** property of Task objects (continued)

Enum for task name	Task	view-only-mode supported
stop-all	Stop All	false
stop-cim	Stop (CIM)	false
storage-resource-details	Storage Resource Details	false
system-activity-display	Activity	false
system-details	System Details	false
system-details-nodes	System Details	false
system-information	System Information	false
system-input-output-configuration-analyzer	System Input/Output Configuration Analyzer	false
system-sysplex-time	System (Sysplex) Time	true
test-communications-with-storage-resources	Test Communications with Storage Resources	false
tip-of-the-day	Tip of the Day	false
toggle-lock	Toggle Lock	false
transmit-console-service-data	Transmit Console Service Data	false
transmit-service-data	Transmit Service Data	false
transmit-vital-product-data-base	Transmit Vital Product Data	false
transmit-vital-product-data-z	Transmit Vital Product Data	false
undefine-z-vm-virtual-machines-for-management	Undefine z/VM Virtual Machines for Management	false
update-hmc-configuration-data	Update HMC Configuration Data	false
user-settings	User Settings	false
users-and-tasks	Users and Tasks	false
view-activation-profiles	View Activation Profiles	false
view-console-events	View Console Events	false
view-console-information	View Console Information	false
view-console-service-history	View Console Service History	false
view-console-tasks-performed	View Console Tasks Performed	false
view-frame-layout	View Frame Layout	false
view-licenses	View Licenses	false
view-pmv-records-base	View PMV Records	false
view-pmv-records-z	View PMV Records	false
view-security-logs	View Security Logs	false
view-service-history	View Service History	false
view-statistics	View Statistics	false
view-the-vmmr-measurement-data	View the VMMR Measurement Data	false
virtual-server-details	Virtual Server Details	false
virtual-server-performance-details	Virtual Server Performance Details	false
virtual-server-resource-adjustments-report	Virtual Server Resource Adjustments Report	false

| Table 271. Enum values for the **name** property of Task objects (continued)

Enum for task name	Task	view-only-mode supported
virtual-server-topology-report	Virtual Server Topology Report	false
virtual-servers-report	Virtual Servers Report	false
whats-new	What's New	false
workload-resource-adjustment-report	Workload resource Adjustments Report	false
workload-resource-group-details	Workload resource Group Details	false
workloads-report	Workloads Report	false
xnode-details	System Details	false
z-vm-virtual-machine-details	z/VM Virtual Machine Details	false
z-vm-virtual-network-information	z/VM Virtual Network Information	false
zbx-blade-details	Blade Details	false
zbx-bladecenter-details	BladeCenter Details	false

---

## Appendix E. Notices

This information was developed for products and services offered in the USA.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785 USA*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

---

## Trademarks

IBM, the IBM logo, and `ibm.com`<sup>®</sup> are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Intel, and Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a trademark of Linux Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Unix is a trademark of The Open Group in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.



---

## Class A Notices

The following Class A statements apply to this IBM product. The statement for other IBM products intended for use with this product will appear in their accompanying manuals.

### Federal Communications Commission (FCC) Statement

**Note:** This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

### Industry Canada Compliance Statement

This Class A digital apparatus complies with Canadian ICES-003.

### Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

### European Community Compliance Statement

This product is in conformity with the protection requirements of EU Council Directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

European Community contact:  
IBM Deutschland GmbH  
Technical Regulations, Department M372  
IBM-Allee 1, 71139 Ehningen, Germany  
Tele: +49 (0) 800 225 5423 or +49 (0) 180 331 3233  
email: halloibm@de.ibm.com

**Warning:** This is a Class A product. In a domestic environment, this product may cause radio interference, in which case the user may be required to take adequate measures.

## VCCI Statement - Japan

この装置は、クラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。 VCCI-A

The following is a summary of the VCCI Japanese statement in the box above:

This is a Class A product based on the standard of the VCCI Council. If this equipment is used in a domestic environment, radio interference may occur, in which case the user may be required to take corrective actions.

### Japanese Electronics and Information Technology Industries Association (JEITA) Confirmed Harmonics Guideline (products less than or equal to 20 A per phase)

高調波ガイドライン適合品

### Japanese Electronics and Information Technology Industries Association (JEITA) Confirmed Harmonics Guideline with Modifications (products greater than 20 A per phase)

高調波ガイドライン準用品

## Electromagnetic Interference (EMI) Statement - People's Republic of China

### 声 明

此为 A 级产品,在生活环境  
中,该产品可能会造成无线电干  
扰。在这种情况下,可能需要用  
户对其干扰采取切实可行的措  
施。

**Declaration:** This is a Class A product. In a domestic environment, this product may cause radio interference, in which case the user may need to perform practical action.

## Electromagnetic Interference (EMI) Statement - Taiwan

警告使用者：  
這是甲類的資訊產品，在  
居住的環境中使用時，可  
能會造成射頻干擾，在這  
種情況下，使用者會被要  
求採取某些適當的對策。

The following is a summary of the EMI Taiwan statement above.

**Warning:** This is a Class A product. In a domestic environment, this product may cause radio interference, in which case the user will be required to take adequate measures.

**IBM Taiwan Contact Information:**

台灣IBM 產品服務聯絡方式：  
台灣國際商業機器股份有限公司  
台北市松仁路7號3樓  
電話：0800-016-888

**Electromagnetic Interference (EMI) Statement - Korea**

이 기기는 업무용(A급)으로 전자파적합등록을 한 기기이오니  
판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의  
지역에서 사용하는 것을 목적으로 합니다.

**Germany Compliance Statement**

**Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit**

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

**Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten**

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 2004/108/EG in der Bundesrepublik Deutschland.

**Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC EG Richtlinie 2004/108/EG) für Geräte der Klasse A**

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:  
International Business Machines Corp.  
New Orchard Road  
Armonk, New York 10504  
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:  
IBM Deutschland GmbH  
Technical Regulations, Abteilung M372  
IBM-Allee 1, 71139 Ehningen, Germany  
Tel: +49 (0) 800 225 5423 or +49 (0) 180 331 3233  
email: halloibm@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

### **Electromagnetic Interference (EMI) Statement - Russia**

**ВНИМАНИЕ!** Настоящее изделие относится к классу А.  
В жилых помещениях оно может создавать радиопомехи, для  
снижения которых необходимы дополнительные меры

# Index

## A

accessibility xxii  
 contact IBM xxii  
 features xxii  
 Activate a Blade 167, 169  
 Activate Availability Policy 583  
 Activate BladeCenter 147  
 Activate CPC 766  
 Activate Logical Partition 814  
 Activate Performance Policy 493  
 Activate Virtual Server 343  
 Activate zBX (Node) 98  
 Activating a Virtualization Host 256  
 Add Group of Virtual Servers to a Workload Resource Group 467  
 Add MAC Filters to Top-of-Rack Switch Port 125  
 Add Member to Custom Group 738  
 Add Node to Ensemble 71  
 Add Permission to User Role 677  
 Add Temporary Capacity 772  
 Add Top-of-Rack Switch Port to Virtual Networks 129  
 Add User Role to User 657  
 Add Virtual Server to a Workload Element Group 569  
 Add Virtual Server to a Workload Resource Group 461  
 Add Virtualization Host Storage Resource Paths 400  
 Add Virtualization Host Storage Resource to Virtualization Host Storage Group 418  
 Add Workload Element Group to a Workload Resource Group 472  
 API version number, function included in 5  
 assistive technologies xxii

## C

Change STP-only Coordinated Timing Network 779  
 Create Availability Policy 578  
 Create Custom Group 735  
 Create IEDN Interface for a DataPower XI50z Blade 171  
 Create IEDN Virtual Switch 243  
 Create LDAP Server Definition 725  
 Create Metrics Context 888  
 Create Network Adapter 319  
 Create Password Rule 712  
 Create Performance Policy 486  
 Create QDIO Virtual Switch 246  
 Create Storage Resource 371  
 Create User 661  
 Create User Pattern 699  
 Create User Role 682  
 Create Virtual Disk 330  
 Create Virtual Network 430

Create Virtual Server 299  
 Create Virtualization Host Storage Resource 395  
 Create Workload Element Group 563  
 Create Workload Resource Group 452

## D

Deactivate BladeCenter 150  
 Deactivate CPC 768  
 Deactivate Logical Partition 816  
 Deactivate Virtual Server 346  
 Deactivate zBX (Node) 101  
 Deactivating a Virtualization Host 257  
 Delete Availability Policy 580  
 Delete Completed Job Status 54  
 Delete Console Hardware Message 640  
 Delete CPC Hardware Message 791  
 Delete Custom Group 737  
 Delete IEDN Interface for a DataPower XI50z Blade 174  
 Delete LDAP Server Definition 728  
 Delete Metrics Context 895  
 Delete Network Adapter 322  
 Delete Password Rule 714  
 Delete Performance Policy 489  
 Delete Storage Resource 375  
 Delete User 664  
 Delete User Pattern 701  
 Delete User Role 684  
 Delete Virtual Disk 334  
 Delete Virtual Network 432  
 Delete Virtual Server 304  
 Delete Virtual Switch 255  
 Delete Virtualization Host Storage Resource 398  
 Delete Workload Element Group 566  
 Delete Workload Resource Group 455  
 Delete zBX (Node) Hardware Message 112  
 Discover Virtualization Host Storage Resources 406

## E

Ensemble power metric group 903  
 Enum values for a class of managed objects 931  
 Enum values for Task objects 935  
 Enum values for User Role objects 933  
 Export Performance Policy 496  
 Export Profiles 771  
 Export World Wide Port Names List 377

## F

Flash memory adapters 903

## G

Generate Availability Status Report 597  
 Generate Hypervisor Report 523  
 Generate Hypervisor Resource Adjustments Report 531  
 Generate Load Balancing Report 556  
 Generate Service Class Hops Report 543  
 Generate Service Class Resource Adjustments Report 538  
 Generate Service Class Virtual Server Topology Report 548  
 Generate Service Classes Report 535  
 Generate Virtual Server CPU Utilization Report 516  
 Generate Virtual Server Resource Adjustments Report 518  
 Generate Virtual Servers Report 511  
 Generate Virtual Servers Report (Ensemble Availability Management) 594  
 Generate Workload Resource Group Availability Status Report 589  
 Generate Workload Resource Group Performance Index Report 504  
 Generate Workload Resource Group Resource Adjustments Report 507  
 Generate Workload Resource Groups Report (Ensemble Availability Management) 586  
 Generate Workload Resource Groups Report (Performance Management) 500  
 Get Availability Policy Properties 577  
 Get Blade Properties 163  
 Get BladeCenter Properties 146  
 Get Capacity Record Properties 878  
 Get Console Audit Log 630  
 Get Console Hardware Message Properties 639  
 Get Console Properties 621  
 Get Console Security Log 633  
 Get CPC Audit Log 782  
 Get CPC Energy Management Data 197  
 Get CPC Hardware Message Properties 789  
 Get CPC Properties 758  
 Get CPC Security Log 784  
 Get Custom Group Properties 734  
 Get EC/MCL Description of zBX (Node) 95  
 Get Ensemble Properties 63  
 Get Group Profile Properties 871  
 Get Image Activation Profile Properties 856  
 Get Inventory 882  
 Get LDAP Server Definition Properties 722  
 Get Load Activation Profile Properties 865  
 Get Logical Partition Properties 810  
 Get Metrics 892

Get Network Adapter Properties 323  
 Get Node Properties 69  
 Get Notification Topics 49  
 Get Password Rule Properties 708  
 Get Performance Management Velocity Level Range Mappings 599  
 Get Performance Policy Properties 483  
 Get Rack Properties 136  
 Get Reset Activation Profile Properties 836  
 Get Storage Resource Properties 370  
 Get Switch Controllers 249  
 Get Task Properties 689  
 Get Top-of-Rack Switch Port Details 121  
 Get Top-of-Rack Switch Properties 119  
 Get User Pattern Properties 695  
 Get User Properties 652  
 Get User Role Properties 673  
 Get Virtual Disk Properties 336  
 Get Virtual Network Properties 426  
 Get Virtual Server Properties 306  
 Get Virtual Switch Properties 241  
 Get Virtualization Host Properties 232  
 Get Virtualization Host Storage Group Properties 414  
 Get Virtualization Host Storage Resource Properties 391  
 Get Workload Element Group Properties 562  
 Get Workload Resource Group Properties 450  
 Get zBX (Node) Audit Log 103  
 Get zBX (Node) Hardware Message Properties 110  
 Get zBX (Node) Security Log 105  
 Get zBX Properties 91

**I**

IBM z Unified Resource Manager 1  
 Import Performance Policy 494, 923  
 Import Profiles 770  
 Import Storage Access List 379  
 Initiate Virtual Server Dump 356  
 Inventory Service Data 258

**J**

Join STP-only Coordinated Timing Network 780

**K**

keyboard  
 navigation xxii

**L**

Leave STP-only Coordinated Timing Network 781  
 List Availability Policies 575  
 List BladeCenters in a Rack 141, 143  
 List Blades in a BladeCenter 158  
 List Blades in a zBX 160  
 List Capacity Records 877

List Console Hardware Messages 637  
 List CPC Hardware Messages 787  
 List CPC Objects 755  
 List Custom Group Members 742  
 List Custom Groups 732  
 List Ensemble CPC Objects 756  
 List Ensemble Nodes 67  
 List Ensembles 61  
 List Group Profiles 870  
 List Groups of Virtual Servers of a Workload Resource Group 465  
 List Image Activation Profiles 854  
 List LDAP Server Definitions 720  
 List Load Activation Profiles 863  
 List Logical Partitions of CPC 808  
 List Members of a Virtual Network 434  
 List Password Rules 706  
 List Performance Policies 481  
 List Racks of a zBX 134  
 List Racks of a zBX of a zBX 117  
 List Reset Activation Profiles 834  
 List Storage Resources 367  
 List Tasks 688  
 List User Patterns 693  
 List User Roles 671  
 List Users 650  
 List Virtual Disks of a Virtualization Host Storage Resource 408  
 List Virtual Networks 424  
 List Virtual Servers of a CPC 294  
 List Virtual Servers of a Node 289  
 List Virtual Servers of a Virtualization Host 297  
 List Virtual Servers of a Workload Element Group 568  
 List Virtual Servers of a Workload Resource Group 458  
 List Virtual Servers of a zBX (Node) 287  
 List Virtual Servers of an Ensemble 291  
 List Virtual Switches 239  
 List Virtualization Host HBA Ports 386  
 List Virtualization Host Storage Groups 412  
 List Virtualization Host Storage Resources 388  
 List Virtualization Host Storage Resources in a Virtualization Host Storage Group 416  
 List Virtualization Host Storage Resources of a Storage Resource 381  
 List Virtualization Hosts of a CPC 229  
 List Virtualization Hosts of a Node 224  
 List Virtualization Hosts of a zBX (Node) 222  
 List Virtualization Hosts of an Ensemble 227  
 List Workload Element Groups of a Workload Resource Group 471  
 List Workload Element Groups of an Ensemble 561  
 List Workload Resource Groups of an Ensemble 448  
 List zBX (Node) Hardware Messages 108  
 List zBXs of a CPC 88  
 List zBXs of an Ensemble 89  
 Load Logical Partition 822

Logoff 48  
 Logon 45

**M**

Make Console Primary 626  
 managed objects, valid enum values 931  
 Migrate Virtual Server 354  
 Mount Virtual Media 348  
 Mount Virtual Media Image 351

**N**

navigation  
 keyboard xxii

**P**

policy XML document  
 description 923  
 example 927  
 PSW Restart 824

**Q**

Query API Version 44  
 Query Job Status 52

**R**

Remove Group of Virtual Servers from a Workload Resource Group 469  
 Remove MAC Filters from Top-of-Rack Switch Port 127  
 Remove Member from Custom Group 740  
 Remove Node from Ensemble 74  
 Remove Permission from User Role 680  
 Remove Temporary Capacity 774  
 Remove Top-of-Rack Switch Port from the Virtual Networks 131  
 Remove User Role from User 659  
 Remove Virtual Server from a Workload Element Group 571  
 Remove Virtual Server from a Workload Resource Group 463  
 Remove Virtualization Host Storage Resource from Virtualization Host Storage Group 420  
 Remove Virtualization Host Storage Resource Paths 404  
 Remove Workload Element Group from a Workload Resource Group 474  
 Reorder Network Adapter 328  
 Reorder User Patterns 628  
 Reorder Virtual Disks 341  
 Reset Clear 820  
 Reset Normal 818  
 Restart Console 625  
 RoCE adapters 903

**S**

SCSI Dump 830  
 SCSI Load 828  
 Set Blade Power Capping 206  
 Set Blade Power Save 204  
 Set BladeCenter Power Capping 201  
 Set BladeCenter Power Save 199  
 Set CPC Power Capping 191  
 Set CPC Power Save 189  
 Set STP Configuration 777  
 Set zBX (Node) Power Capping 186  
 Set zBX (Node) Power Save 184  
 Set zCPC Power Capping 196  
 Set zCPC Power Save 194  
 shortcut keys xxii  
 Shutdown Console 627  
 SMAPI Error Response Body 257  
 Start Logical Partition 825  
 Stop Logical Partition 827  
 Summary of updates by API version  
 number 5  
 Swap Current Time Server 776

**T**

Task objects, valid enum values 935

**U**

Unified Resource Manager 1  
 Unmount Virtual Media 352  
 Update Availability Policy 581  
 Update CPC Properties 765  
 Update Ensemble Properties 65  
 Update Group Profile Properties 873  
 Update Image Activation Profile  
 Properties 859  
 Update LDAP Server Definition  
 Properties 724  
 Update Load Activation Profile  
 Properties 867  
 Update Logical Partition Properties 813  
 Update Network Adapter 326  
 Update Password Rule Properties 710  
 Update Performance Policy 490  
 Update Reset Activation Profile  
 Properties 837  
 Update Storage Resource Properties 374  
 Update Top-of-Rack Switch Port  
 Properties 123  
 Update User Pattern Properties 697  
 Update User Properties 654  
 Update User Role Properties 675  
 Update Virtual Disk Properties 339  
 Update Virtual Network Properties 428  
 Update Virtual Server Properties 315  
 Update Virtual Switch 251  
 Update Virtualization Host  
 Properties 237  
 Update Workload Element Group 567  
 Update Workload Resource Group 456  
 User Role objects, valid enum  
 values 933

**W**

Web Services API 1  
 Workload Resource Group Details  
 task 923

**X**

XML document  
 description for performance  
 policy 923  
 example for performance policy 927

**Z**

zManager 1









Printed in USA

SC27-2627-00

