

IBM z Unified Resource Manager – table of features identified by seven management areas

The z Unified Resource Manager is firmware that executes on the Hardware Management Console (HMC) and Support Element (SE). There are seven management areas that make up Unified Resource Manager – Operational Controls, Hypervisor Management, Virtual Server Management, Performance Management, Availability Management, Network Management, Energy Management. This table lays out the seven management areas, provides a short explanation of the features and what they provide, and highlights if they are available in the Manage (green square ■), Advanced Manage (orange circle ●) or Automate (blue triangle ▲) suite of Unified Resource Manager for z/VM®, POWER7® Blades, and IBM WebSphere® DataPower® Integration Appliance XI50 for zEnterprise (DataPower XI50z).

Please be aware that the system is required to be in an ensemble (FC #0025) in order to enable these functions.

Members of an ensemble can be at different functional levels but the lowest of those levels will be what is supported (i.e. if one node is at Manage and another at Automate, the ensemble will be at Manage).

The z/VM 6.2 hypervisor, announced October 12, 2011, offers new system capabilities. The most significant is the VMSSI Feature, which provides the ability to group up to four z/VM systems in a Single System Image (SSI) cluster. This clustering technology is the backbone for Live Guest Relocation (LGR), which allows guest mobility between z/VM systems in the SSI cluster. See the table at the end of the feature/function charts to clarify how Unified Resource Manager can be used with z/VM V6.2. z/VM V6.2 is the last release of z/VM that can be managed by zManager and zEC12, BC12 are the last servers on which z/VM will be managed.

Operational Controls ¹		z/VM V6.2	POWER7 Blade	DataPower XI50z	IBM System x Blade
Change Management	View Firmware Information	N/A	■	■	■
	Retrieve Firmware Changes	N/A	■	■	■
	Change Firmware Levels	N/A	■	■	■
	Backup/Restore Critical Data	N/A	■	■	■
Problem Management	Automatic Error Logging and first-failure data capture (FFDC) data collection	N/A	■	■	■
	Problem analysis (hardware and LIC problem detection) and call home reporting	N/A	■	■	■
	View hardware messages	N/A	■	■	■
	View open problems	N/A	■	■	■
	Manual problem reporting and data collection	N/A	■	■	■
Guided repair and verification for a service action – repair activities are guided by the HMC and verified automatically		N/A	■	■	■
Configuration Management	Vital Product Data – fully automatic and coherent integrated resource discovery and inventory for all elements of the system without requiring user configuration, deployment of libraries or sensors, or user scheduling. Verifies resources are entitled to be powered on and managed.	■	■	■	■
	Edit Frame Layout – Frames, Switches, BladeCenters	N/A	■	■	■
	Edit Frame Layout – MES support (add/remove)	N/A	■	■	■
	Edit Frame Layout – Management Enablement	N/A	■	■	■
	Dynamic add of storage resources for blades	N/A	■	N/A	■ ²
Operations Management	Allows power on and power off via HMC	N/A	■	■	■
	Event Notification for support personnel (based on logged events or state changes)	■	■	■	■
	Scheduled Operations for functions such as firmware update, activate, deactivate)	■	■	■	■
	Time synchronization	N/A	■	■	■
	Launch Full Device Console	■	■	■	■

¹ Available on IBM z Systems Central Processor Complex (zCPC) – IBM z13s, IBM z13, IBM zEnterprise BC12, IBM zEnterprise EC12, IBM zEnterprise 196 or IBM zEnterprise 114 (z13s, z13, zBC12, zEC12, z196 z114) – and for zCPC only does not require ensemble (FC#0025)

² Available on IBM z13s, z13, zEC12 or zBC12 and IBM z BladeCenter® Extension (zBX) Model 004 and zBX Model 003 but not on z196/z114/zBX Model 002

Operational Controls ³		z/VM V6.2	POWER7 Blade	DataPower XI50z	IBM System x Blade
Performance Management	Monitors Dashboard	■	■	■	■
	Reporting of energy consumption and temperature monitoring	N/A	■	■	■
Business Management	User Management	■	■	■	■
	Auditing	■	■	■	■
	Device Status and Details	N/A	■	■	■
	Automatic Service Network Configuration	N/A	■	■	■
	Documentation	■	■	■	■

Hypervisor Management		z/VM V6.2	POWER7 Blade	DataPower XI50z	IBM System x Blade
Virtualization - manage hypervisors and support virtual servers for application deployment.		■	■	N/A	■
Hypervisors managed as firmware	Booted automatically at power on reset and isolated on the internal Unified Resource Manager network	N/A	■	N/A	■
	Managed through HMC	■	■	N/A	■
Basic hypervisor management tasks	Deploy and initialize hypervisor	N/A	■	N/A	■
	Start, stop and query/list hypervisors	■	■	N/A	■
	Update and repair hypervisors	N/A	■	N/A	■
	Monitor hypervisors and their resource use – CPU, memory, consumption	■	■	N/A	■
	Manage ISO images	N/A	■	N/A	■
	Create virtual networks	■	■	N/A	■
Allow agents in virtual server operating systems to communicate with a manager running in the hypervisor or the hypervisor management stack		■	■	N/A	■

Virtual Server Lifecycle Management – enabling directed and dynamic virtual server provisioning across all hypervisors from a single, uniform point of control		z/VM V6.2	POWER7 Blade	DataPower XI50z	IBM System x Blade
Create (Provision)	CPU	■	■	N/A	■
	Memory	■	■	N/A	■
	Network	■	■	N/A	■
	Console	■	■	N/A	■
	Storage	■	■	N/A	■
	Virtual DVD	N/A	■	N/A	■
List	■	■	N/A	■	
Start/Stop	■	■	N/A	■	
View/Modify configuration	■	■	N/A	■	
Migrate definition	■	■	N/A	■	
Delete	■	■	N/A	■	

³ Available on z Systems Central Processor Complex (zCPC) – zBC12, zEC12, z196 or z114 – and for zCPC only does not require ensemble (FC#0025)

Workload Awareness and Platform Performance Management – management of CPU resource across virtual servers hosted in the same hypervisor instance to achieve workload performance policy objectives		z/VM V6.2	POWER 7 Blade	DataPower XI50z	IBM System x Blade ⁴
Workload definition and workload-based performance definition - Defining your own customer workloads (by name) and the performance management capabilities allows differentiation between multiple workloads in an ensemble.		▲	▲	N/A	● ▲
Platform resource monitoring based on performance policy - performance monitoring and reporting functions primary objective is to provide data to check whether performance objectives are being met. Note: reporting will not replace tools that report on operating system resources and performance	Workload – provides a high level performance status and objective achievement information for all the workloads in the ensemble.	▲	▲	N/A	● ▲
	Virtual Server - Provides data for virtual servers associated with the workload. The data includes allocated resources, resource utilization data and delay statistics.	▲	▲	N/A	● ▲
	Service Class – provides a list of all the service classes defined in the workload performance policy and includes active performance policy objectives and importance definitions, indication when service class performance violated defined threshold.	▲	▲	N/A	● ▲
	Resource adjustment report – provides adjustment actions taken over report interval	▲	▲	N/A	▲
	Hypervisor Report – provides details about the virtual servers that are running in the same hypervisor instance, and how these virtual servers are competing for shared resources	▲	▲	N/A	● ▲
	Load Balancing Report - provides details about the load balancing groups to which external routers are distributing incoming work requests	▲	▲	N/A	● ▲
Dynamic, goal-oriented resource management - Manage CPU across virtual servers within a hypervisor instance		▲	▲	N/A	▲
Load balancing recommendation to external routers - Through SASP communication with Unified Resource Manager, external routers receive weight recommendations that enable the routers to distribute incoming work more effectively among virtual servers in the load balancing group.		▲	▲	N/A	● ▲

⁴ System x blades on z114, z196 will only have Advanced Management functions. Early zEC12 servers will not have full Automate capabilities until they are upgraded to driver level 15. Prior to that level they can only have Advanced Management functions including: automate capabilities for System x blades Wizard-driven set up of resources in accordance with specified business process; ability to monitor and report performance ; Load balance recommendations

Workload Awareness and Availability Monitoring and Reporting		zCPC	z/VM V6.2	POWER7 Blade	DataPower XI50z	IBM System x Blade⁵
Workload definition and workload-based availability definition - Defining your own customer workloads (by name), their business priorities and the availability management capabilities allows differentiation between multiple workloads in an ensemble. Workload availability status indicates the whether all workload elements are operating and available. Element Groups can be configured to manage virtual servers with redundancy objectives, for availability and capacity management among redundant virtual servers providing the same function.		▲	N/A	▲	N/A	▲
Platform resource monitoring based on availability policy - availability monitoring and reporting functions. Primary objective is to provide data to check whether availability objectives are being met.	Workload Report – shows high level availability status information, including active availability policy for all workloads in the ensemble, over a 36 hour period.	▲	N/A	▲	N/A	▲
	Virtual Server Report - shows availability status data for selected virtual servers associated with a workload	▲	N/A	▲	N/A	▲
	Availability Status Event Report – shows all availability status changes for the workload, its element groups, and virtual servers	▲	N/A	▲	N/A	▲

Network Management	zCPC	z/VM V6.2	POWER7 Blade	DataPower XI50z	IBM System x Blade
Create virtual network - Virtual networks allow virtual servers and appliances to be connected together.	■	■	■	■	■
Network Monitoring – provides monitoring and collection of metrics from the Layer 2 networking resources associated with the IEDN.	■	■	■	■	■
HiperSockets™ Integration with the IEDN – extends the reach of HiperSockets network outside the zCPC to the entire ensemble.	■	■	N/A	N/A	N/A

⁵ System x blades on z114, z196 will only have Advanced Management functions. Early zEC12 servers will not have full Automate capabilities until they are upgraded to driver level 15. Prior to that level they can only have Advanced Management functions including: automate capabilities for System x blades Wizard-driven set up of resources in accordance with specified business process; ability to monitor and report performance ; Load balance recommendations





Power save mode is available on the IBM z13, zEnterprise EC12 (zEC12) and zEnterprise 196 (z196) with Automate when full performance of the server is not required. Power save can be switched on and off during runtime with no disruption to currently running workloads, aside from the change in performance. You can use power save mode for periods of lower utilization (for example weekends or third shift) or for capacity backup systems where you keep them “running” but with reduced energy consumption.

With static power save mode you do not turn off the z13, zEC12 or z196 engines, but it slows down the clock speed and thus reduces power. Once the clocks are slowed, it also reduces the supply voltage to get the maximum savings possible for a given configuration.

Static Power Save is available on the z13, zEC12 and the z196. By nature of IBM z13s, zEnterprise BC12 (zBC12) and zEnterprise 114 (z114) having a slower clock speed already, they can not take advantage of Power Save. Power save operations on the zCPC can only be performed once a day, however this restriction only applies to air cooled machines.

Energy Management – energy monitoring and management can help better understand the power and cooling demand of the system Description	zCPC	POWER7 Blade	DataPower XI50z	IBM System x Blade ⁶
Monitoring of energy consumption and key environmental parameters	■	■	■	■
Power saving mode ⁷ - Power Save is a feature implemented with a combination of hardware and firmware functions to minimize the power consumption of a component. This reduction is typically achieved by reducing the performance of the component. The amount of power reductions is limited by the hardware components abilities to save power and is controlled by the component itself. However power save is providing a guaranteed reduction of power consumption.	▲	▲	N/A	N/A
Power capping - Power Capping is a feature that limits the peak power of a component and therefore is used for power and cooling capacity planning. Power can be allocated to components or systems and the power capping guarantees that the power consumption and the related heat dissipation will never exceed this specified limit.	▲	▲	▲	● ▲

Can z/VM 6.2 be a member of an ensemble and a Single System Image (SSI) or Cross System Extension (CSE) cluster?

z/VM is a member of an ensemble?	z/VM is a member of an SSI or CSE cluster?	
	Yes	No
Yes	Not recommended 	
No		

⁶ System x blades on z114, z196 will only have Advanced Management functions. Early zEC12 servers will not have full Automate capabilities until they are upgraded to driver level 15. Prior to that level they can only have Advanced Management functions including: automate capabilities for System x blades Wizard-driven set up of resources in accordance with specified business process; ability to monitor and report performance ; Load balance recommendations

⁷ For zCPC only does not require ensemble (FC#0025)