
IBM z Systems
Introduction
May 2017

IBM z BladeCenter Extension (zBX)
Frequently Asked Questions

Worldwide



ZSQ03098-USEN-01

Table of Contents

IBM z BladeCenter® Extension (zBX).....	3
POWER7 Blades for zBX.....	7
IBM System x blades for zBX.....	10
IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z)	13
zBX Networking.....	14
zBX Availability	16
zBX and blade Warranty and Maintenance.....	19
IBM z Unified Resource Manager.....	21
Ensemble.....	24

IBM z BladeCenter® Extension (zBX)

What is the hardware machine type of the zBX model 004? Is there a machine type for the optimizers and blades installed in the zBX?

The zBX model 004 is machine type 2458. The new Model 004 can be included in an ensemble with a IBM z13s™, IBM z13™, IBM zEnterprise® BC12 (zBC12), IBM zEnterprise EC12 (zEC12), IBM zEnterprise 114 (z114) or IBM zEnterprise 196 (z196) and is no longer connected to the processor via the internode management network (INMN). The Model 003 attaches to the zBC12 and zEC12. The Model 002 attaches to the z114 or z196.

IBM POWER7® and IBM System x® blades were ordered independently of the zBX (see the POWER7 and System x section of this document for details).

DataPower® XI50z is ordered as a feature of the zBX but will generate machine type 2462 Model 4BX.

What advantages will I have with the zBX Model 004?

Our clients have asked us to remove the tight maintenance network integration between the IBM z Systems™ server and the zBX. With the zBX Model 004 we are going to install two Support Elements (SE) in the zBX. Now, firmware updates and service activity will no longer be linked to the server. The SEs will be connected directly to the ensemble HMC. Breaking this connection, i.e. between the server and the zBX, will also remove the distance limitation of 25 meters from the server. For some clients this is important as they traditionally don't keep their distributed servers next to the z Systems in the data center.

There will still be a data connection to a server (the IEDN network). The IBM z Unified Resource Manager goals will not change, still providing hardware management, virtualization management, and workload management based on user defined priorities.

The zBX Model 004 update will include a refresh for DataPower XI50z with the delivery of DataPower firmware 7.0.

If I have four zBX frames (racks), is that considered one zBX or four zBXs?

A single configuration can include up to four frames (depending on the number of blades ordered).

If I have more than one frame in the zBX are they bolted together?

Yes IBM offers 'marriage kits' in the ship group if you have more than one frame. These bolt the frames together.

Is there anything that I will need to customize on the BladeCenter Chassis?

No. The zBX is a totally integrated system. The only thing that is different per installation is that a customer decides which blades they put into the zBX – a combination that can include DataPower XI50z, select POWER7, and System x blades.

How many z Systems servers can connect to a zBX?

By defining an ensemble, a combination of up to eight z Systems servers can connect an ensemble to one zBX.

What is a 'node' in terms associated with an ensemble?

A single zBX Model 004, z13s or z13, or any zBC12, zEC12, z196, or z114 with an optionally attached zBX. A node can be a member of only one ensemble.

Can I have an ensemble with a z13s or z13 node and a zBX Model 004 node and attach a zBC12 (or zEC12) with a zBX, and then also attach a z196 or z114 servers? Can I do that kind of mix and match?

Yes, workloads can be defined with resources on any of those nodes and be part of the ensemble. But remember that the zBX Model 003 can only be connected to the zBC12 or zEC12 and the zBX Model 004 does not have a direct management connection to any server.

What is the physical size of the zBX?

Customers can order one to four zBX racks depending upon the number of optimizers and blades required. The first rack will be a standard 42u 19" sized rack and the rest will be 42u 19" expansion racks. Customers can request a 36u height reduction option if needed to get the zBX model 003 into their data center.

How much time will it take to load the firmware on the blades using a zBX?

There will be some variability based on the configuration. Firmware load times can range from a few minutes to several hours.

What if I put a supported blade that already has data on it into the zBX?

When the supported blade is discovered and configured, firmware will be loaded which will replace the BIOS and remove any existing data – including any Machine Code that was on the blade.

If I have a zBX with four frames, are there top-of-rack (TOR) switches in each frame?

No, TOR switches for management and data are only in the first frame, so all Ethernet network connectivity flows through those switches. Each BladeCenter has a pair of SAN switches, and the switches need direct attachment to SAN external switches

If I order a server with a non-raised floor (NRF) option, can I also order a zBX with a NRF option?

A zBX Model 004 can be installed on a non-raised floor as it has the ability for top exit I/O and power.

I have a zBX Model 003 attached to my zBC12/zEC12. I'd like to upgrade the zBC12/zEC12 to a z13s/z13 and move the zBX Model 003 to different zEC12/zBC12. If I purchase new entitlements on the second zEC12/zBC12 is there any issue for me to move the zBX Model 003?

No. You can do this move.

Can I detach my Model 003 and let it sit standalone?

No. IBM doesn't support this action, as it could age to an unacceptable condition.

Why are there only two BladeCenter chassis per zBX frame?

Due to the BladeCenters' weight, a lift tool would be required, that doesn't exist. Additionally, as the zBX is shipped as a single unit, more than two BladeCenters would exceed shipment tilt and limitations.

What is the benefit of having a zBX compared to blades installed in a stand-alone BladeCenter H?

There are several benefits of using the zBX. One is the high speed 10 Gb layer 2 network that is configured for high availability and secured by the Unified Resource Manager. The connectivity can reduce latency and overhead and potentially remove the requirement for additional hardware (specifically firewalls). Another benefit is the Unified Resource Manager, which provides uniformity of management tasks, independent of server type or operating system. The system administrator uses Unified Resource Manager to set up new virtual servers in the same way, independently of the operating system and the hardware.

Can an IFL-only server with a zBX attached run Unified Resource Manager?

The Unified Resource Manager does not require a specific configuration or operating system, so an IFL-only server is compatible with the zBX.

How is the BladeCenter chassis configured in the zBX?

All the blades are virtualized, so the underlying hardware is never material for management and provisioning. The BladeCenter chassis is part of the configured hardware that comes with the zBX order. When an order is placed, depending on the number of blades specified in the configurator you plan to install, the required hardware is placed into the zBX. This also means that when you want to upgrade the zBX to add blades or optimizers, the z configurator (eConfig) will do that planning work again and the resulting configuration will have the necessary hardware.

The components of the BladeCenter chassis are also not managed independently. The operational controls on the Unified Resource Manager enable the underlying hardware infrastructure components to be managed transparently.

Is the Unified Resource Manager the mechanism for managing the BladeCenter chassis and its I/O modules?

Yes. You define virtual servers using the Unified Resource Manager. All aspects of a virtual server can be defined – CPU, memory, network, console, disk storage, and virtual DVD – taking into account differences in the underlying capabilities. (For example a virtual DVD is not supported for a z/VM-based virtual server). In addition, virtual guests can be listed, started and stopped, reconfigured, and deleted when no longer required. A virtual server definition can also be moved from one hypervisor to another of the same type.

You also use Unified Resource Manager to define workloads. A workload is a logical grouping of physical and virtual resources within the context of named business processes. When you set up a workload you will assign the virtual servers, storage and network (VLANs) resources that you want to make available to the workload.

If I have spare slots in the zBX can I install blades that are not controlled by Unified Resource Manager?

No. The current design point of the zBX, and a very important one from a control and management point of view, is that a blade in a zBX must be managed by the Unified Resource Manager. We do not have an option for a non-managed zBX blade.

What is the maximum number of blades I can put in the zBX?

The total capacity of a zBX cannot exceed 112 blades and can be any combination of types and capacities. The zBX will support up to 112 POWER7 blades, up to 56 System x blades, and up to 28 DataPower XI50z (doublewide) blades. The blades for POWER7, System x and DataPower XI50z can be shared in the same BladeCenter chassis. Remember that DataPower appliance blades are “doublewide” and use two slots.

I am an IBM z Systems client today using zBX. I am interested in migrating to the latest zBX Model 004. Are there resources to help me in this migration, and to do customization of my environment?

Yes, IBM STG Lab Services can assist with this. You can contact IBM STG Lab Based Services via the Internet at: <http://www.ibm.com/systems/services/labservices/> or send an email to stqls@us.ibm.com

POWER7 Blades for zBX

I know IBM Systems team has announced the end of marketing for the PS701 blade. What does that mean for the zBX? Does the announcement mean that I am no longer able to order PS701 blades for my zBX?

That is correct. Effective 1/15/14, the standalone PS701 Express blade was withdrawn. If you are able to obtain blades, you are required to meet the hardware configuration specified by this zBX website document:

http://www.ibm.com/common/ssi/cgibin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_US_EN&htmlfid=%0bZSY03019USEN&attachment=ZSY03019USEN.PDF

Will we have Linux available to run on the POWER7 blades?

No. The POWER7 blades can run AIX® 5.3 (Technology Level 12) or later in POWER6® and POWER6+™ compatibility mode, AIX 6.1 (Technology Level 5), and AIX 7.1.

From a Linux standpoint, the z13s and z13 supports both Linux on z Systems (available and supported on the server) and Linux for System x (available and supported on the IBM System x blades installed in the zBX).

Which version of PowerVM™ do I need?

You will need to purchase a license for PowerVM Enterprise Edition (EE). You do not need to obtain a physical copy of the PowerVM software as it will be loaded as z Systems firmware. Note that when the blade is installed in the zBX and “discovered” by the Unified Resource Manager some PowerVM functions, such as Live Partition Mobility, will not be exploited.

If I have a rack with POWER7 blades installed in it, such as an existing BladeCenter H, can I integrate those existing blades into a zBX so they can become part of an ensemble?

If the blades match the published model and configuration specifications for integration into the zBX, those blades are interchangeable and can be installed in the zBX through the standard defined processes.

NOTE: If the blades are not already licensed for PowerVM Enterprise, and do not have the PowerVM Enterprise activation codes entered, the licenses must be obtained, and the activation codes must be entered before the blades are removed from their current BladeCenter chassis. This is because Unified Resource Manager does not provide the User interface to enter activation codes, so you will have to enter them in their current BladeCenter environment if you don't already have them. Additionally, you cannot install an existing BladeCenter H chassis into the zBX.

The following is a link for entering the activation codes in the blades' existing BladeCenter environment:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7hch/iphchenablevios.htm?resultof=%22%61%63%74%69%76%61%74%69%6f%6e%22%20%22%61%63%74%69%76%22%20%22%63%6f%64%65%22%20>

How is the installation of AIX and the hypervisor managed? Today firmware on a POWER7 blade can only be updated via an OS installed on the blade or by booting a diagnostic CD. How will this be done in the zBX?

The hypervisors are managed by Unified Resource Manager and are loaded automatically when the blade is "discovered and configured" within the zBX. The hypervisor will not be accessible to the customer. Once a virtual server has been created the customer follows normal operating system installation (AIX) procedures to create a new image.

What if the POWER7 blade has more memory on it than the supported configuration?

The blade would be rejected by Unified Resource Manager.

Will there be separate certification and support statements required for the various middleware products run on the supported AIX?

If middleware or applications run on PowerVM on a POWER7 blade today, then existing support statements are sufficient. IBM does not require and does not plan to issue separate certification and support statements. IBM has also been working closely with its ISV partners and already has substantial feedback that they share this view.

Are the Power blades managed by Systems Director?

No. The zBX blades are managed by the Unified Resource Manager.

What are the current allowable configurations of the PS701 blades installed in the zBX?

There are three memory configurations of PS701 BladeCenter Express blades currently supported in the zBX: 32 GB DIMM, 64 GB DIMM and 128 GB DIMM.

What adapters can be ordered on the PS701 blade?

The standard 2port 10 GbE (FC# 8275) and the Fibre Channel Adapter (FC# 8242) must be ordered. They are the only adapters that are supported in the z Systems environment.

Is there a requirement to set up the VIOS? If so, who does that?

VIOS is considered part of the hypervisor, and is therefore managed and delivered by the z Systems firmware update process. VIOS and PowerVM will be automatically installed on each blade after it is entitled and activated. So, there is no install of VIOS by the customer. Additionally, any of the manual configurations of VIOS that is traditionally performed will no longer be required. To keep VIOS more stable, there is no external access to it. The Unified Resource Manager will affect the required configuration changes to VIOS in response to the POWER® LPAR definition activities that are performed on the HMC.

Will the POWER7 support PowerHA™?

PowerHA can be configured to perform automated service recovery for the applications running in virtual servers deployed in zBX. PowerHA automates application failover from one virtual server in a IBM System p® blade to another virtual server in a different System p blade with similar configuration.

Failover protects service (masks a service interruption) in case of unplanned or planned (scheduled) service interruption. During failover, clients may experience a short service unavailability, while resources are configured by PowerHA on the new virtual server.

Power HA configuration for the zBX environment is similar to standard Power environments, with the particularity that it uses only virtual I/O resources. Currently, PowerHA for zBX support is limited to failover inside the same ensemble and not for business continuity across ensembles.

Refer to TechDocs FLASH10750 for the most current information on supported releases:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/webindex/flash10750>

IBM System x blades for zBX

What select System x blades will be supported in the zBX?

The zBX will support IBM BladeCenter HX5 (7873 blades). The HX5 (7873) blade is built on the IBM BladeCenter family of products and designed using the IBM X-Architecture® technology with fifth-generation eX5 innovation, with capabilities that help promote greater system uptime and energy-smart design.

There are four HX5 single-wide two socket HX5 (7873) blade configurations that have been tested by IBM for inclusion in the zBX Model 003 and Model 004. The models are based on four memory configurations - 64 GB, 128 GB, 192 GB and 256 GB. ONLY those blade configurations defined by IBM z Systems will be supported and work in the zBX. The blades are optimized to achieve maximum performance for both the system and its virtual machines. The blades come under the management of the IBM z Unified Resource Manager.

MAX5 or other x86 blades are not currently supported in the zBX. If unsupported blades are installed in the xBX, they will be rejected by Unified Resource Manager.

What is the capacity of zBX in terms of the number of System x Blades that can be supported?

We support 56 HX5 blades per zBX. These are 16 core blades so a total of 896 cores will be supported at a level of 1 virtual server per 2 cores.

Who orders the blades? Is it part of the configuration process by my IBM team?

Similar to the PS701 blade for the zBX today, the System x blade will not be part of the zBX configuration. The following document should be used as a guide for the specific configuration required.

http://www.ibm.com/common/ssi/cgibin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_US_EN&htmlfid=ZSL03128USEN&attachment=ZSL03128USEN.PDF

How will the IBM System x Blades and supported Linux® and Microsoft® Windows® operating systems be priced?

The pricing for IBM HX5 Blades and the supported Linux and Windows operating system is provided through your standard qualified channels. z Systems does not price or ship these as part of the zBX order.

I currently have a 4-socket Linux subscription. Can I use this on the 2-socket blade installed in the zBX?

Yes. It should be no problem. Red Hat charges by socket pairs – so the subscription is “up to 4-sockets” and it should work on the 2-socket blade without a problem.

What would happen if the blade I try to install does not match the supported System x configuration?

The blade would be rejected by Unified Resource Manager.

What operating system software releases and hypervisor are supported on the System x blades in the zBX model 004?

The two Linux operating system configurations that are supported are Red Hat Enterprise Linux (RHEL) 5.5 and up, 6.0 and up, 7.0 and up. And SUSE Linux Enterprise Server (SLES) 10 (SP4) and up, SLES 11 (SP1) and up and SLES 12 and up.

The following levels of Microsoft Windows are supported: Microsoft Windows Server 2012 R1, Microsoft Server 2012 R2, Microsoft Windows Server 2008 R2 and Microsoft Windows Server 2008 (SP2) – any of these must be the 64-bit versions.¹ We recommend that you order Datacenter Edition as that supports an unlimited number of virtual guests, but that is not a requirement.

For operating system support on older zBX models see the zBX spec page:

http://www.ibm.com/systems/z/hardware/zenterprise/zbx_specs.html

When ordering the operating system for your HX5 (7873) blade, you should not have it preinstalled on the blade. This is primarily due to the fact that the Unified Resource Manager will install an integrated hypervisor (which is KVM based) on the IBM z BladeCenter HX5 (7873) when it finds the blade installed on the zBX. Second, you may already have a license for the operating system that can be used. Therefore, you may not have a requirement to order the operating system at the time the blade is ordered.

What is KVM in relation to the zBX and what is the difference between 'standard' KVM and the 'KVM based hypervisor' used by zBX for the System x blades?

On May 3, 2011, IBM and Red Hat, Inc. announced an agreement to develop products and solutions jointly, based on Red Hat's KVM (Kernel-based Virtual Machine). KVM is open source software that provides a full virtualization solution for Linux on x86 architecture.

There is no difference between the standard KVM hypervisor code and our 'KVM-based hypervisor' code. KVM is KVM. The difference is that the KVM hypervisor used by zBX for x86 blades is encapsulated and only accessible through the Unified Resource Manager interface. This also means that only the functions Unified Resource Manager surfaces and exploits for KVM are available.

Is the Microsoft Windows 2012 and 2008 Server Operating System certified for the zBX virtualization environment?

Yes. zBX uses KVM as the hypervisor for x86 blades. The flavor of KVM that is used in the model 004 is: Red Hat Enterprise Virtualization - Hypervisor 6.5 (RHEV-H 6.5)². The Microsoft SWVP certification is inherited from Red Hat and can be found here on the Microsoft site:

<http://www.windowsservercatalog.com/results.aspx?&chtext=&cstext=&csttext=&chbtext=&bCatID=1521&cplID=17597&avc=34&ava=0&avq=0&OR=1&PGS=25&ready=0>

¹ All supported releases of Windows may not be available from Microsoft but continue to be available from IBM and other resellers

² For KVM levels on older zBX Models, see: http://www-03.ibm.com/systems/z/hardware/zenterprise/zbx_specs.html

Are the RHEL and SLES Linux Operating Systems certified for the zBX virtualization environment?

Yes. zBX uses KVM as the hypervisor for x86 blades. The flavor of KVM that is used in the zBX Model 004 is: Red Hat Enterprise Virtualization - Hypervisor 6.5 (RHEV-H 6.5)². Red Hat RHEL certifications are inherited from Red Hat and can be found on the Red Hat site:

<http://www.redhat.com/resourcelibrary/articles/enterprise-linux-virtualization-support>

For SuSE SLES, we have obtained a separate certification for the zBX, running the HX5 blade, which can be found on the Novell site: <http://developer.novell.com/yessearch/>

Does IBM intend to support virtualization on the System x blades? What about VMware - are there plans to make this available?

IBM requires the use of virtualization on System x blades, however, we do not have plans to support VMware or its hypervisor.

Will there be separate certification and support statements required for the various middleware products run on the supported Windows or Linux on z Systems?

If middleware or applications run on KVM on a HX5 blade today, then existing support statements are sufficient. IBM does not require and does not plan to issue separate certification and support statements. IBM has also been working closely with its ISV partners and already has substantial feedback that they share this view.

Are the System x blades managed by Systems Director?

No. The zBX blades are managed by the Unified Resource Manager.

IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z)

Does the zBX Model 004 still support the DataPower Appliance?

Yes. IBM DataPower Gateway Appliances extend industry-leading service-oriented architecture (SOA), security, control, optimization, and integration capabilities to web, mobile, and API workloads. The DataPower appliance can be inserted in the zBX and supported by Unified Resource Manager.

What is the value to the DataPower XI50z of the firmware level v7.0?

There are multiple functions in the v7.0 firmware level. Important to z Systems clients are new intuitive and easy-to-use tool for creating and deploying common DataPower configuration workloads; an embedded, on-demand router functionality for WebSphere® Application Server Network Deployment environments; and native JSON (JavaScript Object Notation) support which brings enhanced security & control for REST (Representational State Transfer) services.

Other enhancements include:

- GatewayScript: A JavaScript runtime that is secured, optimized and tuned for the gateway environment to simplify configuration for developers and provide an easier development paradigm for Mobile, Web, & API
- New Virtual Edition for Developers provides a low cost, per user pricing, and easy to use gateway for developers
- Support for Citrix XenServer hypervisor provides additional deployment flexibility on-premise & cloud deployments
 - WebSocket Proxy support enables full-duplex, bi-directional, & low-latency communication for Mobile & Web applications, Internet of Things
- Improved security & traffic control functionality in support of IBM API Management offering

Where can I get information on DataPower XI50z?

There is a set of z Systems software FAQs with DataPower XI50 information:

http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ST&infotype=SA&appname=STGE_ZS_ZS_USEN&htmlfid=ZSQ03052USEN&attachment=ZSQ03052USEN.PDF

zBX Networking

Can I use Layer-2 forwarding to connect my network (customer network) to the zBX in order to utilize resources available in the ensemble?

Customer experience with the IBM z BladeCenter Extension (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 now available, illustrating a set of pre-tested configuration examples in support of both Layer-2 and Layer-3 connectivity. The Redpaper, IBM z 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/>

We still strongly recommend Layer-3 routing to connect from the external network into the zBX. Utilizing a Layer-2 switch can cause unexpected results. Layer-3 routing can be achieved through the use of either a Central Processor Complex (CPC) operating system (e.g. z/OS® or Linux on z Systems) enabled as a router (possibly with additional IP filtering configured) or an external hardware IP router (and possibly an external load balancing solution). This second method would likely be used for traffic that is directed at the blades which may or may not access the CPC as the 2nd-tier (e.g. database server).

For either case, the 'best practice' for secure entry into the IEDN is to use a firewall at the entry point, providing for the isolation, logging and auditability that are typical security requirements when crossing security zones.

Is 1000BASE-T available for access to my network (customer network)?

No, for connection between the zBX and customer routers, you are only able to utilize fiber optic 1000BASE-SX (SX multimode fiber) or 1000BASE-LX (LX single mode fiber for longer distances).

Given that you have incorporated a private data network and a private management network in the zBX, can I still attach to my company's external network?

Remember that the zBX Model 004 does not have a private management network. But the answer is yes, you can use externally facing ports on the TOR switches in the zBX and OSA adapters on the server to attach directly to your external network.

Since I have the IEDN, why might I want connectivity to my external network?

The IEDN provides connectivity between different elements of the nodes that form an ensemble. Your external network provides access to the ensemble resources from nodes residing outside the ensemble - for example, from the Internet or from servers in your intranet. To protect the ensemble resources from users and systems on the external network while allowing traffic on the secured IEDN to flow in the clear, some customers might use physical or virtual firewalls, encryption, and other technologies. In summary, many network configuration choices are available to balance throughput, latency, data, and security requirements.

What resources can use the secure data network (IEDN)?

The data network (IEDN) is controlled by the Unified Resource Manager, which defines the virtual networks that can be used for communication. When you create virtual servers, you can define network interfaces for them and associate them with these virtual networks, enabling the virtual servers on the same virtual network to communicate with one another.

With twelve SAN connections per BladeCenter chassis within the zBX, is it mandatory that I use all twelve connections?

No, there is no need to use all the SAN connections. We STRONGLY recommend that you connect at least four. There are two ports on each blade and two switches in each BladeCenter, so we believe four provides the minimum amount of redundancy required by best practices.

When a blade attempts to log in to the SAN fabric, the switch forwards that connection over the least-loaded of the ports connected to the SAN. So if four ports are plugged, it will load balance over four. If six are connected, it will load balance over six, and so on. This is only done at fabric login time. Rebalancing will occur when the blade is (re)booted or if the connection fails (for whatever reason) and is re-established. So if the customer starts off with four ports, activates all the blades, then plugs more cables into the rest of the ports, only the first four ports will be used until another fabric login occurs.

I know the Fibre Channel switch modules in the BladeCenters (installed in the zBX - FC# 0606) are the ones used to attach to the disk storage I supply. Do those switch models auto-negotiate (like the FICON Express8/8S feature) to whatever speed the device is at the other end?

Yes, we recommend that your SAN ports be set to allow to them to auto-negotiate with the fibre channel ports in the zBX modules to achieve the highest speed the SAN port to which they connect supports (2, 4 or 8 Gbps).

zBX Availability

Can you discuss zBX and ensemble availability a little? I'd like to be sure I understand my responsibilities in this area.

To ensure maximum availability for workloads in an ensemble with zBX, configure multiple nodes with sufficient zBX capacity to accommodate the work of a failing zBX. This can be accomplished in two phases: develop a single node with a zBX and clone the next ones.

The zBX is designed and delivered with high levels of hardware and firmware redundancy, so losing one is unlikely. Nevertheless it is a single point of failure and must be redundantly configured on another ensemble node for high availability.

BladeCenters and blades also contain single points of failure, so virtual servers and disks needing high availability must be redundantly configured on failure isolated components (e.g. back up blades in a different BladeCenter in a different frame or zBX.)

Because of eConfig tool constraints, extra blade entitlements may need to be purchased to protect from a BladeCenter failure. This is because the zBX does not support multiple sparsely populated BladeCenter Chassis. High availability of virtual servers running in zBX can be achieved through automated failover and recovery to/from redundant blades in the same or different zBX.

High availability of external data can be achieved by means of synchronous mirroring and automated switching between primary and secondary disks, using capabilities like Metro Mirror and TPC-R HyperSwap® and OpenSwap as well as GDPS® HyperSwap, xDR and Distributed Cluster Management and Veritas® Cluster Server.

How can GDPS be used to provide availability for my zBX?

GDPS products currently support the start, stop, and switching of applications executing on virtual servers in a zBX; however, this requires that the target virtual servers be active. For example, when a planned site switch from site 1 to site 2 is performed, GDPS will stop the applications executing on virtual servers in a site 1 zBX, reverse disk replication and online the disk to the target active virtual servers in site 2 zBX, and start the workloads on the target virtual servers. With this announcement, GDPS/PPRC will exploit Unified Resource Manager APIs and no longer require the target virtual servers be active. For example, when a planned site switch from site 1 to site 2 is performed, GDPS will stop the applications executing on virtual servers in a site 1 zBX, deactivate the virtual servers in site 1, reverse disk replication, activate the target virtual servers in site 2 zBX, and start the workloads.

GDPS/PPRC is designed to provide a metro distance continuous availability/disaster recover capability. For greater system resiliency GDPS/PPRC is adding failover support for Unified Resource Manager, or Hardware Management Console (HMC). If the primary HMC becomes unreachable from the GDPS/PPRC controlling system then GDPS will communicate with the alternate HMC to request that it takes over as the primary. After the alternate HMC assumes the primary role, GDPS will begin communicating with it. GDPS/PPRC supports starting and stopping zBX virtual servers that boot from IBM System Storage® DS8000® series.

This is designed to extend GDPS/PPRC continuous availability and disaster recovery capabilities to virtual servers in a zBX for planned and unplanned outages. Availability starts with GDPS/PPRC v3.10 SPE.

Will zBX frames have the same availability as the classic z Systems server?

No, z Systems servers offer superior business service availability that results from closer collaboration among different components including system hardware, firmware, operating systems (z/OS, z/VM®, z/VSE®, and z/TPF), storage and middleware. Blades do not offer the same qualities of service merely because they are attached to a z Systems machine. However, the blades benefit from the capabilities of the robust z Systems management environment such as first-failure data capture, “call home”, and guided repair, as well as from the redundancy built into the zBX hardware.

Is there redundancy in the blade hardware itself?

No. If a blade failed due to (for example) a hard drive problem, you would install a new blade and reload it. You might lose some ISO images that you’ve uploaded to that blade.

The zBX is populated with standard IBM-certified components like BladeCenter Chassis, top-of-rack (TOR) switches, power supplies and optimizers or blades. Most of these parts are based on the standard BladeCenter and blade hardware offerings. So what kind of availability can I expect with the zBX?

Much of the existing product’s reliability is inherited by the zBX. But to enhance the availability of the zBX, new and enhanced RAS capabilities have been incorporated into infrastructure. These include:

Hardware redundancy at various levels

- Power infrastructure
- Power and switch units in the BladeCenter chassis
- Cabling for management of zBX and data connections
- Concurrent to system operations
- Install and growth
- Hardware repair
- Firmware fixes and driver upgrades
- First failure data capture and problem analysis
- Automated call home for hardware/firmware problems

What are the availability characteristics of the networks that connect the z13s/z13 and the zBX Model 004?

There is only one connection between a z13s/z13 and a zBX Model 004.

- A private and secure intra-ensemble data network (the IEDN) connects all elements of an ensemble for normal application interconnection purposes. This network is access-controlled using integrated virtual local area network (VLAN) provisioning. This requires no external customer-managed switches or routers, which helps to reduce the need for firewalls and encryption, and simplifies network configuration and management, while providing full redundancy for high availability. The management of the network provides enforcement of strict access control across heterogeneous environments, further augmenting security and simplicity.

Can I put storage or other things in the zBX rack?

No. The zBX is managed by the Unified Resource Manager, which ensures that blades in the zBX are of supported types. Unsupported blades in the zBX such as storage blades would not be powered on or configured.

zBX and blade Warranty and Maintenance

What is the warranty for the zBX? How is the warranty for the IBM BladeCenter PS701 Express blade or IBM BladeCenter HX5 7873 different than the zBX warranty?

The warranty for the zBX is a traditional IBM z Systems warranty - one year PD (problem determination) and 24x7 onsite FRU (field replaceable unit) service performed by a zSSR (z Systems Service Support Rep). Once a Blade is accepted by the zBX, it inherits the zBX's warranty. If the Blade is un-entitled and removed from the zBX, it can be serviced by IBM by the machine type under which it was purchased.

The IBM BladeCenter PS701 Express or IBM BladeCenter HX5 blade comes with a 3-year warranty (limited) 9 hours per day, Monday through Friday (excluding holidays), next business day at no additional cost; on-site for selected components; CRU (customer replaceable unit) for all other units (varies by country).

If I have a problem on a blade, how do I get maintenance support?

Hardware, firmware or software calls/problems can be initiated by the machine, by you or by your IBM z Systems Service Representative, via a PMV report on the HMC. Calls will be routed to appropriate hardware or software product support group. Details on this process are documented in the HMC Operations Guide (V2.11.1 Doc # SC28-6905)

Once I install a supported blade into the zBX is it still (CRU) customer replaceable unit maintenance?

No, the supported blade installed in the zBX will take on the same z Systems maintenance as is on the zBX, so it will now be 24x7 onsite FRU by a zSSR.

When I order a POWER7 or System x blade for the zBX should I order post-warranty maintenance?

No. Do not purchase post-warranty maintenance service. Post-warranty maintenance service for the blade servers in the zBX is provided by the IBM maintenance on the 2458 zBX with the appropriate feature codes, either #0612 (Power Blade Enablement) or #0613 (IBM System x Blade Enablement).

Who physically installs the POWER7 or System x blades in the zBX?

You do, or you can negotiate a GTS contract to have IBM install them for you.

I have a supported POWER7 or System x blade and I'd like to install it into the zBX and make it accessible to the zBX. What is the process to make the zBX ready to use the blade?

If you have any entitled BladeCenter slots, not occupied, you can use them. If you don't you should provide your IBM sales representative with an order letter for the type and quantity of blades you plan to install in the zBX. It is important for your letter to specify the number of enablement features on the zBX and the number of zBX blade feature codes (the number will be the same).

I don't have IBM maintenance on my IBM z Systems hardware (z13s/z13, zBC12/zEC12, z196, z114). How can I order maintenance for the zBX?

You will need to purchase HMC Feature Code #0091, 0092 or 0094, depending on the zBX, in order for the IBM zSSR to load updates to the zBX.

Can I remove a blade from my zBX and install it a BladeCenter H? What happens to the warranty and maintenance on the blade?

If you remove a blade from the zBX, you no longer have 24x7 Field Replaceable Unit (FRU) maintenance by the zSSR. The blade reverts to any remaining warranty that remains on the blade - blade warranty is 9am-5pm Monday through Friday Customer Replaceable Unit (CRU).

Can I install supported blades that I purchased prior to zBX availability in the zBX?

You will need to verify the blades meet the zBX requirements and you will need order entitlement support for the quantity of blades on the zBX.

What additional POWER7 blade Predictive Failure Analysis® (PFA) capabilities will z Systems add over existing BladeCenter PFA capabilities?

In addition to the PFA capabilities on the IBM POWER7 BladeCenter PS701 Express blade, the Unified Resource Manager's Problem Management is extended to the blades to enable automatic logging of error and first-failure data capture information. Blade-related problems are analyzed to determine their severity and are reported to IBM through the call-home mechanism for potential proactive service action. Problems that are system-detected can be managed through these mechanisms, as can customer-initiated problems.

When a service action is required for the blade, repair activities are guided by the HMC and the effectiveness of the repair is verified automatically.

IBM z Unified Resource Manager

What is the Ensemble Availability Management function for Unified Resource Manager?

Ensemble Availability Management offers the ability to create user-defined availability policies for managing the availability virtual servers, along with monitoring and reporting functions to help ensure virtual servers are executing in line with those policies. 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 represents whether all workload elements are operating and available. Element Groups may be configured to manage virtual servers with redundancy objectives, enabling availability and capacity management among redundant virtual servers that provide the same function.

This enhanced function is available for virtual servers in logical partitions and on HX5 blades and PS701 blades in the zBX model 003 and zBX model 004.

What is the relationship between Unified Resource Manager and OpenStack®?

The Unified Resource Manager is designed to provide systems management capabilities across the multi-architecture environment of an ensemble. Many of our clients are planning to exploit these system management capabilities to deploy a framework for a heterogeneous cloud environment, thereby providing an effective means to help deliver IT services. To accelerate the delivery of its cloud offerings, IBM recently announced plans to base all of its cloud services and software on open architecture and standards, including OpenStack. OpenStack is an infrastructure as a service (IaaS) cloud computing open source project. IBM joined the project in 2012 and in support of z Systems is making contributions to the OpenStack project that are designed to enable z/VM V6.3 to be the first z Systems operating environment that can be managed with these open cloud architecture-based interfaces.

The management of z/VM environments in Unified Resource Manager is now stabilized as part of the evolution of the IBM cloud strategy and adoption of OpenStack. Accordingly, **Unified Resource Manager does not provide systems management support for z/VM V6.3.** However, Unified Resource Manager will continue to play a distinct and significant role in the management of virtualized environments created by integrated firmware hypervisors - PR/SM™, PowerVM, and System x hypervisor based on Kernel-based Virtual Machine (KVM). Looking ahead, IBM will continue to enable OpenStack to provide heterogeneous systems management across the z13s/z13, z/VM, and distributed platforms, which in turn can be exploited by subsequent IBM SmartCloud offerings. These offerings are designed to provide you enablement for enterprise-wide cloud deployments and greater flexibility by removing the need to develop specific interfaces for different cloud services.

What are the APIs for Unified Resource Manager and where can I find more information about them?

Application program interfaces (APIs) are designed to provide access to the same underlying functions that support the Unified Resource Manager user interface capabilities for inventory, provisioning, configuration, operational control, monitoring, and workload optimization of the heterogeneous physical and logical resources of a z Systems environment. The capabilities available using the APIs are consistent with their counterparts provided by the Hardware

Management Console (HMC) user interface and may be exploited to enable discovery, monitoring, and provisioning use cases.

Additional information can be found in the Unified Resource Manager Redbook (www.redbooks.ibm.com) or in the white books in the library available on Resource Link³ (www.ibm.com/servers/resourcelink).

What is Dynamic Storage Provisioning?

Dynamic Storage Provisioning allows virtual disks to be added/allocated to a running virtual server without having to take the virtual server down. This ability is available for virtual servers running on the Power and x86 blades on the zBX. Note that the guest operating system must support the dynamic addition of virtual based disks to use dynamic storage provisioning.

Can you share some information about the Server/Application State Protocol (SASP) support?

Switches/Routers (e.g., F5 Big Switch) that are enabled for SASP can take advantage of the Unified Resource Manager SASP load balancing feature to load balance traffic across Linux (for blades) / AIX / Windows types of virtual servers in the zBX.

What is the philosophy of the Unified Resource Manager?

The Unified Resource Manager is firmware that manages the integration of multiple platform resources as a single virtualized system and provides a single point of control for the ensemble. The Unified Resource Manager allows customers to inspect, report, and manage all connected resources and automate their deployment to deliver the highest service level agreements (SLAs). Similar to the z/OS Workload Manager, the Unified Resource Manager allows the automation to be aligned with business goals and priorities for workloads on the zBX servers and Linux for z Systems server images. z/OS Workload Manager continues to manage z/OS workloads.

How do I use the Unified Resource Manager to set up resources for an application workload?

The Unified Resource Manager is accessible from the Hardware Management Console (HMC), either directly or through a Web browser. From the HMC you define the virtual servers that are needed for the workload. Once the virtual servers are selected, you are led through the definition of workload policies by a wizard.

Is Unified Resource Manager using the Workload Manager capabilities of z/OS to manage the resources that it has defined?

Unified Resource Manager is not using z/OS Workload Manager (WLM); it is firmware that comes with the z13s/z13 and runs on the HMC and it works side-by-side with the z/OS WLM. The Unified Resource Manager can monitor workloads on the zBX, Linux on z Systems, and z/OS workloads. It can manage workloads and resources on the zBX or Linux on z Systems.

When Unified Resource Manager is executed, you have the ability to define a workload as the set of virtual resources required and defined to support a business process. The workload can include z/OS, z/VM, Linux on z Systems, and AIX on POWER7. Virtual servers running on IBM BladeCenter HX5 (7873) blades in the zBX have the ability to manage resource optimization. A performance

³ A userid and password are required to access Resource Link – if you don't have one then follow the instructions to set one up.

policy can be attached to the workload that defines its service-level requirements. Similar to Workload Manager on z/OS, if goals are not being met the Unified Resource Manager will take appropriate corrective action.

IBM understands the requirement to establish complementary policies for both z/OS WLM and Unified Resource Manager to target the same objectives.

Does the Unified Resource Manager support live migration of virtual servers?

IBM understands the requirement for Unified Resource Manager to provide live migration of virtual machines (whether on POWER7 or System x blades), but the capability is not yet available.

Does IBM's Software Group (SWG) sell Unified Resource Manager?

No, Unified Resource Manager is part of the firmware of the z13s/z13, sold and delivered as a feature of the z13s/z13 by the z Systems hardware sellers (STG).

Does the Unified Resource Manager increase what I will be charged for use of my z Systems software?

No, the Unified Resource Manager runs completely on the z Systems HMC as firmware so it does not contribute to your z Systems software bill.

What is the guest platform management provider (GPMP)?

GPMP is an optional application that may be installed in specific z/OS, Linux, AIX and Windows operating system images to support platform management functions. For example, the GPMP collects and aggregates performance data for virtual servers and workloads. Users view these reports through the Hardware Management Console that manages an ensemble.

Does the GPMP Linux monitoring agent run in a native Linux LPAR or only under z/VM?

The GPMP does not run in a Linux native LPAR.

How does z/OS work with Unified Resource Manager?

z/OS integrates with the Unified Resource Manager environment using a GPMP agent. For z/OS V1.12 (and later) the GPMP sends basic z/OS WLM data (such as system resource utilization, system delays, and paging delays) back to the Unified Resource Manager over the INMN. GPMP is designed to add workload relationships from the ensemble components to your z/OS workload; for example, linking a transaction that may have started on the zBX back to DB2 on z/OS data.

Do I need to support and manage the software implementation or upgrades of Unified Resource Manager and its agents?

No, it is completely serviced via firmware upgrades and fixes. However you may need to install an upgraded GPMP agent.

Ensemble

What is an ensemble?

An ensemble is a collection of one or more nodes (a node is a single zBX Model 004, z13s or z13, or any zBC12, zEC12, z196, or z114 with an optionally attached zBX). An ensemble can consist of 1 to 8 nodes that are or include a server, and up to 8 additional Model 004 nodes. The resources of an ensemble are managed by Unified Resource Manager, and virtualized as a single pool of resources integrating system and workload management across the multisystem, multitier, multi-architecture environment.

So is an ensemble a virtualized container managed by the Unified Resource Manager?

An ensemble is a virtualized platform cluster. It defines the scope of the Unified Resource Manager's management domain and binds its scope of management responsibility. Think of the ensemble as a cluster of virtualized system resources managed by a single Unified Resource Manager. So, conceptually it could be considered a logical container of virtualized resources pooled to form a single logical system.

Do all the management suites of function for Unified Resource Manager need to be at the same level within an ensemble?

When adding a machine into an ensemble, the machine or node should be at the same level as the ensemble it is entering. This is due to the fact that while machines or nodes may exist within a current ensemble with different levels, they will operate only at the level of the 'lowest denominator'. This means that if any node is at the Manage level, the entire ensemble will operate at the Manage level. If one or more machines in the ensemble are at a higher level, such as Automate, they will not get to take advantage of the additional function and will only operate at the Manage level. Once all nodes have been upgraded to Automate level, Automate function will become available. Additionally, upgrading from Manage to Automate is fully concurrent.

Is Tivoli® software mandatory when setting up an ensemble?

No. Tivoli software is not mandatory and will be dictated by what functions of the z13s/z13 and how much of them you want to manage from Tivoli.

Is an ensemble required to be a member of a Parallel Sysplex®?

No. A Parallel Sysplex is a cluster of z/OS workloads used for z/OS system availability, scalability, as well as workload and software management. An ensemble can help you manage cross-platform workloads. Logical partitions on a z13s, z13, zEC12, zBC12, z196 or z114 may, but do not have to be, members of a Parallel Sysplex.

Must I have a zBX on each z Systems server in the ensemble?

No, this is not a requirement. However, if there are multiple nodes in the ensemble, at least one of them must include a zBX (an attached Model 002/Model 003 or a Model 004) in order to use the IEDN. An ensemble that has multiple nodes and no zBX, cannot use the IEDN.

Can a System z10 be part of an ensemble?

No.

Can z/VM V6.2 virtual servers be part of an ensemble?

Yes, z/VM V6.2 and its virtual servers can be part of an ensemble and can be managed by the Unified Resource Manager. But, z/VM V6.3 and its virtual servers cannot be managed by Unified Resource Manager and if you plan to use the VMSSI feature, and especially if you plan to use its Live Guest relocation capability, then the best practice is to not use Unified Resource Manager to manage z/VM virtual servers. Also note that the zEC12/BC12 are the last servers that can manage z/VM6.2 virtual servers through Unified Resource Manager. This capability is not on the z13s/z13.

What has changed about z/VM V6 Release 3 and Unified Resource Manager?

The management of z/VM environments in Unified Resource Manager is now stabilized as part of the evolution of the IBM cloud strategy and adoption of OpenStack. Accordingly, Unified Resource Manager does not provide systems management support for z/VM V6.3. However, Unified Resource Manager will continue to play a distinct and significant role in the management of virtualized environments created by integrated firmware hypervisors - PR/SM, PowerVM, and System x hypervisor based on Kernel-based Virtual Machine (KVM).

This means that z/VM 6.3 cannot be a member of an ensemble.



©Copyright IBM Corporation 2017

IBM Corporation
New Orchard Road
Armonk, NY 10504
U.S.A.

Produced in the United States of America,
05/2017

IBM, IBM logo, AIX, BladeCenter, DataPower, DS8000, GDPS, HyperSwap, Parallel Sysplex, Power6, Power6+, Power7, Power, PowerHA, PowerVM, Predictive Failure Analysis, PR/SM, System p, System Storage, System x, Tivoli, Veritas, WebSphere, x-Architecture, z13, z13s, zEnterprise, z/VM and z/VSE are trademarks or registered trademarks of the International Business Machines Corporation.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

InfiniBand and InfiniBand Trade Association are registered trademarks of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.