



zEnterprise System Introduction to Ensembles

GC27-2609-06

Level 06b





zEnterprise System
Introduction to Ensembles

GC27-2609-06

Note:

Before using this information and the product it supports, read the information in “Safety” on page vii, “Notices,” on page 71, and *IBM Systems Environmental Notices and User Guide*, Z125–5823.

This edition, GC27-2609-06, applies to the IBM zEnterprise System (zEnterprise). This edition replaces all versions of GC27-2609-05.

There might be a newer version of this document in a **PDF** file available on **Resource Link**.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 2010, 2014.**

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

Contents

Figures v

Safety vii

Safety notices	vii
World trade safety information.	vii
Laser safety information	vii
Laser compliance	vii

About this publication ix

Intended audience	ix
Related information.	ix
Accessibility	x
How to send your comments.	x

Summary of changes. xi

Chapter 1. The changing data center . . . 1

Common concerns	2
zEnterprise and ensembles.	5

Chapter 2. Introduction to zEnterprise ensembles 9

What is a zEnterprise ensemble?.	9
Ensemble infrastructure	13
Supported hardware, firmware, and software	20
Hardware	20
Hypervisors and hardware firmware	22
Software	24
Optimizations	29
Select IBM blades	30

Chapter 3. zEnterprise ensemble system management. 33

Hypervisor management	40
Ensemble membership management	42
Storage management	44
Virtual server management	46
Virtual network management	47
Performance management	49
Availability management	52
Energy management, monitoring, and controls	53
zEnterprise System Web Services API.	55

Chapter 4. Planning your system . . . 57

The Insurance Company: A sample implementation	58
The Insurance Company plan for a zEnterprise ensemble	58
The Insurance Company business goals and processes	61
The Insurance Business Application in a zEnterprise ensemble	63
The virtual servers supporting the Insurance Business Application	66

Appendix. Notices 71

Trademarks	72
Electronic emission notices	72

Glossary 77

Index 83

Figures

1. IBM timeline of speciality engines	2	19. IBM zEnterprise Unified Resource Manager API	55
2. Heterogeneous system infrastructure	3	20. The Insurance Company hardware configuration before zEnterprise	58
3. Increasing complexity driven by virtualization and infrastructure complexity	4	21. The Insurance Company hardware configuration with zEnterprise components	59
4. Infrastructure integration	5	22. A closer look at the z196 for the Insurance Company IT shop	60
5. Basic ensemble configuration with zBX.	9	23. A closer look at the zBX for the Insurance Company IT shop	61
6. Basic infrastructure of a two-node ensemble.	10	24. Architecture view of the Insurance Business Application	62
7. Hardware Management Consoles	11	25. Three-tier view of the Insurance Business Application	63
8. Simple view of intraensemble data network	13	26. Mapping the Insurance Business Application to the Insurance ensemble.	64
9. Networks in an ensemble with an IEDN vSwitch bridge to zCPC HyperSockets.	16	27. Software support for layers of the Insurance Business Application on the z196	65
10. Networks in an ensemble with an vSwitch Uplink to a customer-provided data network	16	28. Software support for layers of the Insurance Business Application on POWER and System x blades	66
11. Basic ensemble configuration without zBX	17	29. Software support for Annuity policies on the z196	67
12. Basic ensemble configuration with zBX	18	30. Software support for Annuity policies on blades in the zBX	68
13. Ensemble configuration with zEC12 or zBC12 and z196 or z114 systems, each with zBX attached.	19	31. Software support for Variable Life policies on the z196.	69
14. Mixed system environment	19	32. Software support for Variable Life policies on blades in the zBX	70
15. View of the PR/SM hypervisor layer	23		
16. View of z/VM instance running three virtual machines	25		
17. IBM zEnterprise Unified Resource Manager zBX Model 002 (HMC version 2.11.x) "Manage", "Advanced Manage", and "Automate" features.	34		
18. IBM zEnterprise Unified Resource Manager zBX Model 003 (HMC version 2.12.x) "Manage" and "Automate" features.	35		

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 System z® models can use I/O cards such as, ESCON®, 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

The information technology (IT) challenge is to simplify and optimize the IT infrastructure so that your business can efficiently respond in a cost-effective way.

The IBM zEnterprise® System (zEnterprise) ensemble offers a solution to improve IT infrastructure by tightly integrating different types of servers, storage, operating system software, and networking into a heterogeneous environment. The zEnterprise integrated platform management uses virtualization and firmware to provide the foundation on which the next generation data centers are being built.

Figures included in this document illustrate concepts and are not necessarily accurate in content, appearance, or specific behavior.

Intended audience

This publication is intended for executives, data processing managers, data processing technical staff, consultants, and vendors who want to explore the advantages of zEnterprise ensembles.

Related information

Important: Ensure that you are using the most recent version of all related documentation for the version of IBM zEnterprise Unified Resource Manager you are researching.

The following publications document the zEnterprise environment.

- For Hardware Management Console (HMC) or Support Element (SE) console version 2.12.1 user interface or task information, go to the console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmca/v2r12m1/index.jsp>.
- *zEnterprise System Hardware Management Console Operations Guide for Ensembles*, SC27-2615 Version 2.11.1 and SC27-2622 Version 2.12.0, contains an overview of the user interface and the tasks that you use to configure and manage a zEnterprise environment.
- *zEnterprise System SE Operations Guide*, SC28-6906 Version 2.11.1 and SC28-6920 Version 2.12.0, contains an overview of the user interface and the tasks that you use to configure and manage a zEnterprise environment.
- *zEnterprise System Ensemble Workload Resource Group Management Guide*, GC27-2621 (Unified Resource Manager version 2.12.1), helps you define, create, and manage the performance and availability goals of business application solutions deployed in a zEnterprise environment.
- *zEnterprise System Ensemble Performance Management Guide*, GC27-2607 (Unified Resource Manager version 2.12.0 or prior), helps you define, create, and manage the performance goals of business application solutions deployed in a zEnterprise environment.
- *zEnterprise System Ensemble Planning and Configuring Guide*, GC27-2608, leads you through the planning and implementation phases for moving business application solutions from a distributed environment to an integrated set of virtual servers running in a zEnterprise environment.
- *zEnterprise Unified Resource Manager*, SG24-7921, introduces the Unified Resource Manager, system concepts, planning, implementing, and managing solutions in an integrated zEnterprise environment.
- *System z Hardware Management Console Web Services API*, SC27-2626 Version 2.12.1, SC27-2617 Version 2.12.0, and SC27-2616 Version 2.11.1, defines the application programming interface (API) for zEnterprise. You can use the API to programmatically accomplish ensemble tasks that are available through the Hardware Management Console (HMC).

- *zEnterprise BC12 System Overview*, SA22-1089, *zEnterprise EC12 System Overview*, SA22-1088, *zEnterprise 196 System Overview*, SA22-1086, and *zEnterprise 114 System Overview*, SA22-1087, describe the design, components, functions, features, and capabilities of the IBM zEnterprise servers.
- *System z Advanced Workload Analysis Reporter (zAware) Guide*, SC27-2623, describes an integrated set of analytic applications that monitor software running on z/OS® and model normal system behavior.

For information on the IBM WebSphere® DataPower® Integration Appliance XI50 for zEnterprise, see <http://www.ibm.com/software/integration/datapower/library/documentation>.

Other IBM System z publications are available on Resource Link®. You can access these books at <http://www.ibm.com/servers/resourcelink>, under the Library section.

Accessibility

This publication is in Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience difficulties using this PDF file you can request a web-based format of this publication. 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 request, 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 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.

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. 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).

Summary of changes

Technical changes to the text are indicated by a vertical bar (|) to the left of the change. This edition contains the following additions or changes.

Summary of changes for GC27-2609-06b

This edition adds support for IBM WebSphere DataPower Integration Appliance version 6.0.

See topics:

- “DataPower XI50z” on page 29
- “Related information” on page ix

Summary of changes for GC27-2609-06a

This edition clarifies the levels of operating systems and hypervisors that run on the different models of zBX.

See topics:

- “POWER blade” on page 30
- “System x blade” on page 31

Summary of changes for GC27-2609-06

This edition adds IBM zEnterprise BC12 (zBC12), Unified Resource Manager version 2.12.1, Availability Management, and support for dynamic processor management on virtual servers deployed on IBM System x® blades based on policy goals.

See topics:

- “zEnterprise hardware” on page 20
- Added zBC12 to figures Figure 5 on page 9, Figure 12 on page 18, and Figure 13 on page 19.
- For Ensemble Availability Management, see topics “System z platform management concepts extended” on page 11, Chapter 3, “zEnterprise ensemble system management,” on page 33, and “Availability management” on page 52.
- zBX integrated hypervisor for IBM System x blades supports dynamic processor management on page 24.
- Chapter 3, “zEnterprise ensemble system management,” on page 33 enhanced the functions available with the Automate suite.
- z/VM® support changes:
 - z/VM V6.1 end of service is effective May 1, 2013.
 - Unified Resource Manager does not support z/VM V6.3 as a member. See topic “Unified Resource Manager, z/VM V6.3, and OpenStack” on page 28.
- “Related information” on page ix.
- Unified Resource Manager version 2.12.1 changes in various topics throughout the book.
- Editorial changes in various topics throughout the book.

Summary of changes for GC27-2609-05a

This edition adds support for IBM WebSphere DataPower Integration Appliance version 5.0.

See topics:

- “Related information” on page ix
- “DataPower XI50z” on page 29
- Editorial changes in various topics throughout the book.

Summary of changes for GC27-2609-05

This edition adds IBM zEnterprise EC12 (zEC12), zBX Model 003, planning for upgrades and zBX moves, and the term workload resource group. Also, removed IBM Smart Analytics Optimizer for DB2® for z/OS.

See topics:

- Chapter 3, “zEnterprise ensemble system management,” on page 33
- “Performance management” on page 49
- Added zEC12 to figures Figure 12 on page 18, and Figure 13 on page 19.
- “zEnterprise hardware” on page 20
- “IBM zEnterprise BladeCenter Extension (zBX)” on page 21
- Migration scenarios in Chapter 4, “Planning your system,” on page 57
- Editorial changes in various topics throughout the book.

Summary of changes for GC27-2609-04a

This edition adds virtual switch HiperSockets™ intraensemble data network (IEDN) network bridge support for z/VM and an increase in the number select System x blades supported in the IBM zEnterprise BladeCenter® Extension (zBX). See topics:

- HiperSockets availability for the IEDN on page 14
- “IBM zEnterprise BladeCenter Extension (zBX)” on page 21
- Editorial changes in various topics throughout the book.

Summary of changes for GC27-2609-04

This edition adds Microsoft Windows support for System x blades in an ensemble, Web Services Application Programming Interface (Web Services API), and some new tasks to the ensemble. See topics:

- HiperSockets availability for the IEDN on page 14
- “System x blade” on page 31
- Discover Storage Resources on page 45
- Network Monitors Dashboard on page “Network Monitors Dashboard” on page 48
- Load Balancing Report on page “Load Balancing Report” on page 52
- “zEnterprise System Web Services API” on page 55
- Editorial changes in various topics throughout the book.

Summary of changes for GC27-2609-03

This edition adds the IBM System x blade, Advanced Manage suite, IBM zEnterprise 114 (z114) and some new tasks to the ensemble. See topics:

- “zEnterprise hardware” on page 20

- zBX integrated hypervisor for IBM System x blades on page 24
- “Select IBM blades” on page 30
- Figures; Figure 6 on page 10, Figure 7 on page 11, Figure 13 on page 19, Figure 14 on page 19, and Figure 17 on page 34.
- Open Graphical Console on page 47
- Customize Network Settings on page 48
- Editorial changes in various topics throughout the book.

Summary of changes for GC27-2609-02

This edition adds the IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z) to the IBM zEnterprise BladeCenter Extension (zBX) . See topics:

- Figure 6 on page 10
- “IBM zEnterprise BladeCenter Extension (zBX)” on page 21
- “DataPower XI50z” on page 29
- IEDN Interfaces in zBX Blade Details on page 54

This edition includes enhancements to the **Manage Alternate HMC** task. See topic:

- “Ensemble membership management tasks” on page 42

This edition includes editorial changes in various topics throughout the book.

Summary of changes for GC27-2609-01

This edition adds the alternate Hardware Management Console (alternate HMC) to the ensemble. See topics:

- “What is a zEnterprise ensemble?” on page 9
- “Hardware Management Console” on page 10
- “Configuration examples” on page 17
- “Hypervisor management” on page 40
- “Ensemble membership management tasks” on page 42
- Editorial changes in various topics throughout the book.

Chapter 1. The changing data center

The goal of information technology (IT) is to create an infrastructure that is dynamic, scalable, and flexible to satisfy the needs of mission critical work, development, and deployment of new workloads. The IT model is shifting. Real time event-driven workloads, richer content, and modular technologies are altering the composition, deployment, and management of systems. The demand for raw processing power is growing at an exponential rate.

The traditional processing power growth vehicle of faster single-thread processors is giving way to “throughput computing,” multithread, multi-core, and massively parallel environments. A fresh view on how applications are deployed and heterogeneous resources are managed is needed.

Multiplatform solutions have become the norm for handling computational acceleration and specialized processing. As these heterogeneous systems grow, the end-to-end management can become a burden on resources and the IT budget.

Many of the IT assets have become islands of computing resources connected together with multiple networks. This infrastructure is driving cost and complexity to unsustainable levels, while hampering organizational maneuverability. There is a growing realization that our IT systems, network infrastructure, and IT asset management are reaching a breaking point.

Business trends are driving a new generation of applications that are diverse in their functions, processing requirements, and processing power requirements. Advanced analytics is fast becoming a staple of business processing. An intense focus on risk management is driving the need for real time intensive processing of vast amounts of data. A single solution might require multiple architectures (for example, transactional and analytical processing).

This shift to parallelism also opened the door to a wide class of specialty processors and appliances. They are designed to handle specific tasks better than general-purpose processors. The demand for more real-time analytics and queries is driving the manufacture of specialty engines that can accelerate algorithmic processing. Appliances have emerged to handle various networking algorithms and data transformations. They also represent a new hierarchy of devices that must be configured, managed, and maintained.

The illustration in Figure 1 on page 2 shows how IBM historically responded to the increasing demand for processing capacity. For example:

- IBM System z Application Assist Processor (zAAP) is a specialized processor that provides a Java™ execution environment. This environment enables the integration of Java-based web applications and service-oriented architecture (SOA) based technologies to be integrated with core z/OS business applications and backend database systems.
- IBM System z Integrated Information Processor (zIIP) is a specialized processor designed to help free general computing capacity. The zIIP lowers the overall total cost of computing, for select data and transaction processing workloads. The zIIP is also designed to help enterprise resource planning (ERP), customer relationship management (CRM), and select network encryption workloads on the server.




Figure 1. IBM timeline of speciality engines

The IBM zEnterprise System (zEnterprise) with the optional IBM zEnterprise BladeCenter Extension (zBX) continues this tradition with the introduction of the optimizers and select IBM blades.

Common concerns

As IT business solutions become ever more complex, so has the complexity of the data center infrastructure.

Infrastructure

The IT dilemma in managing the complexity that business solutions require today is the core motivation for providing virtualization and management capabilities across heterogeneous architectures. The illustration in Figure 2 on page 3 shows the core issues for providing virtualization and management capabilities across heterogeneous architectures. This illustration depicts the platform necessary to process a single transaction with all the components representing islands of computing. These islands of computing all use different languages and there is little coordination among them. The software does not take advantage of larger symmetric multiprocessing systems, workloads that require optimization to meet throughput, and processor utilization objectives. Virtualization is trying to get control of this optimization, but ultimately, having more islands of virtual resources to manage only complicates managing an end-to-end business solution. zEnterprise provides a better approach: The ability to manage IT infrastructure and business applications as an integrated whole from a single point of control.




Figure 2. Heterogeneous system infrastructure

Complexity and limitations include:

- Many tiers or nodes of independent resources connected over multiple corporate networks are less dynamic with minimal resource sharing
- Managing infrastructure islands affects human productivity
- Platform management information is not an end-to-end view
- Automation policies are limited to tier or node boundaries
- Redundancy is pervasive for operational staff, hardware, and software
- Software and policies are not able to communicate across architectures
- Managing IT complexity uses most of IT resources for maintaining the existing infrastructure rather than adding new capabilities
- Supply inefficiencies and demand spikes put pressure on environmental concerns dealing with increasing power, cooling, physical space, and proliferation of hardware
- A significant percentage of infrastructure capacity is under used
- A infrastructure with islands of virtual resources reduces the flexibility and responsiveness of IT to deliver end-to-end business solutions
- Systems connected over multiple corporate networks are facing a proliferation of new risks and threats

Virtualization

Virtualization has major benefits and introduces more complexity.

Creating highly virtualized resources demands a stronger, more integrated service management approach. Combining business and IT infrastructures can be a daunting task, if not handled in an integrated way. As you adopt higher levels of automation and new innovative methods of delivery, you must be confident all of it can be achieved securely. You want the applications required to be available, recover quickly, and adjust to any failures or outages. Demands for a better enterprise wide information infrastructure providing integrated information to users is on going challenge.

Enterprise-wide virtualization can deliver significant value but it also potentially reflects a new form of complexity, as shown in Figure 3 on page 4. Even with the most carefully constructed estimators, scripts, automation, and configuration tools it is difficult. Do-it-yourself planning, assembling, testing, configuration, and maintenance tasks are required. Flexibility is gained at the expense of complexity and loss of potentially differentiating business value.




Figure 3. Increasing complexity driven by virtualization and infrastructure complexity

Without attention to heterogeneity, IT continues to struggle with islands and fragmentation of the computing infrastructure.

Specialization management

The drive to specialization is also substantially increasing the inventory of processors and threads. Specialization is exacerbating the complexity of managing systems and applications. Specialization is making it more difficult to maintain service-level agreements, and it is diminishing the ability to completely secure user information. Monitoring, reporting, and balancing goal-oriented workloads in context with workload performance and availability policies in a hybrid environment is difficult.

Firmware and software management

Firmware and software management of managed islands of computing is time-consuming and becoming more complex to manage. In a typical multi-tiered system, there might be a dozen or more firmware code loads and software applications to manage.

zEnterprise ensemble

A requirement has emerged for platform management of a "hybrid" architecture. Figure 4 illustrates how the platform management of a "hybrid" infrastructure compares to previous IT solutions. A new enterprise IT solution is described in the next topic.




Figure 4. Infrastructure integration

zEnterprise and ensembles

zEnterprise delivers value through industry leading security, availability, scalability, virtualization, and management capabilities. With the addition of IBM zEnterprise BladeCenter Extension (zBX) and ensembles, zEnterprise quality of service (QoS) and management capabilities are extended to a heterogeneous hybrid platform. This solution provides a single integrated firmware management, monitoring, and policy-based framework across multiple operating systems and hardware resources. The zEnterprise with a zBX provides a single logical virtualized system, to lower the cost and complexity of managing a multi-tiered IT infrastructure. These platform and system management capabilities provide operational advantages that a multi-system, potentially multi-vendor, or "piece parts" approach cannot match.

zEnterprise

zEnterprise is optimized to provide support for mission-critical data for mixed workloads. zEnterprise is engineered with a strong threading model and robust cache in support of the processing of data. With this strategy, zEnterprise delivers exceptional data management capabilities while also accommodating mixed solutions that have close affinity to the data.

The role of System z today includes the following strengths:

- Secure and resilient enterprise data hub

- Enterprise server for mission-critical applications requiring high levels of availability and security

- Highly efficient consolidation platform for exceptional cost savings

The role of System z tomorrow includes:

- A cost-efficient platform across broader enterprise workloads
- Better energy monitoring and management of system energy consumption
- Multi-tiered business application host for a wider range of critical applications
- Extension of System z QoS (RAS) and platform management to virtualized heterogeneous platforms and applications

Within an IT infrastructure, there are substantial pockets of heterogeneous resources surrounding and integrated with System z servers. It is not practical to put all this work on the System z servers. But, zEnterprise world-class virtualization and resource management has been extended to a selected set of hardware and software solutions. The zEnterprise with a zBX, creates an ensemble of heterogeneous resources as a single highly optimized virtual container that can be managed as a single enterprise.

zEnterprise and ensembles represent the state-of-the-art technologies to enhance IT platform management.

Ensembles

As the list of operational challenges facing many IT organizations continues to grow, it is becoming more apparent that a change is needed. Transformation to a state-of-the-art IT infrastructure for growth and success in the future must be initiated.

It is time to start thinking differently about infrastructure. Infrastructure must become more dynamic. It needs a foundation for breakthrough productivity, accelerated value creation, and increased velocity to achieve the faster pace demands of a smarter world. The evolutionary approach of zEnterprise and ensembles spans distributed physical and virtual assets and IT processes to address the business imperatives of today and tomorrow.

zEnterprise and ensembles deliver a logically integrated and managed view of the infrastructure resources. The ensemble management capabilities of zEnterprise improve the ability to integrate, monitor, and dynamically manage heterogeneous server resources as a single logical virtualized environment.

Value of zEnterprise and ensembles:

- **Reduced Cost**
 - Simplified infrastructure of system resources and operational policies
 - Maximized utilization
 - Savings on power consumption with better monitoring and planned usage during peak workloads
 - IT management reductions through intelligent automation of resource provisioning, management of all firmware updates from a single point of control, system monitoring, and problem resolution
 - Integration of computationally intensive functions to specialty processors
- **Managed Risk**
 - End-to-end customized governance and management of multi-tiered infrastructure
 - Extension of mainframe management to heterogeneous resources of a multi-tiered infrastructure
 - Improved security and automated problem resolution
 - Reduced human error for firmware maintenance
 - Dynamic policy-based workload management across a multi-tiered infrastructure
 - Ensured assignment of needed resources to highest priority workloads
 - Enterprise wide virtualization across systems, storage, networks, and applications
 - Controlled access to data and information, protect network and physical infrastructure against threats and vulnerabilities
- **Improved Service**
 - Integration, monitoring, and management of multiple platform resources as a single, logical, virtualized system
 - Intelligent, autonomic, and dynamic resource allocation to maximize cross-platform efficiencies
 - Integration of real-time analytics with operational data

- Integration, monitoring, and management of multiple through sharing large pools of compute resources
- Maintenance of continuous business operations 24x7

Chapter 2. Introduction to zEnterprise ensembles

A zEnterprise ensemble is a collection of one or more zEnterprise nodes (including any attached zBX) that are managed as a single logical virtualized system by the IBM zEnterprise Unified Resource Manager (Unified Resource Manager or zManager) on the Hardware Management Console (HMC).

What is a zEnterprise ensemble?

A pair of HMCs is required to configure, control, and monitor the hardware, software, and firmware of an ensemble, see Figure 5 and Figure 6 on page 10. The HMC used to create the ensemble is the only HMC that can manage member resources of the ensemble. This HMC is called the primary HMC. If the HMC managing the ensemble fails, the alternate HMC becomes the primary HMC of the ensemble. Attached storage (for example DS 8300 and 5020) is not a member of an ensemble.




Figure 5. Basic ensemble configuration with zBX

- An ensemble is managed as a single virtualized pool of server resources.
- Individual ensembles can be geographically dispersed.
- Nodes for a given ensemble are deployed within distances no greater than what is supported by the direct connection of nodes using 10 GB Ethernet fibers. See “Maximum distances between network connections” on page 20.
- Virtualized server resources might be native LPAR, z/VM virtual images, x Hyp (zBX integrated hypervisor for IBM System x blades) hosted virtual images or PowerVM® hosted virtual images.
- Unified Resource Manager is integrated firmware.
- System z10® can connect to the zEnterprise server using ISC-3 or PSIFB channels but cannot be part of the managed ensemble.
- The zBX provides support for:
 - IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z)

Combined hardware and software solution that simplifies, helps govern, and enhances the security of XML and IT services.

- Select IBM blades

Supports application integration with zEnterprise transaction processing, messaging, and data serving capabilities.

For a list of supported IBM blades or supported storage devices, go to <http://www.ibm.com/systems/z/hardware/zenterprise/zbx.html>.




Figure 6. Basic infrastructure of a two-node ensemble.

Hardware Management Console

The HMC performs system management tasks or performs both system management tasks and ensemble management tasks. The HMC provides a single point of control and single system image for those CPCs or nodes defined to it. (A single CPC, with or without an attached zBX, is called a node. The term node is most often used when referring to an ensemble.

When managing an ensemble, a pairing of HMCs is required - the primary HMC and the alternate HMC. The HMC managing the nodes in an ensemble is referred to as the primary HMC. The primary HMC can also manage CPCs that are not members of an ensemble. The alternate HMC is used as backup. If the primary HMC fails, the alternate HMC take over the role of the primary HMC.

An HMC, other than the primary HMC or the alternate HMC, can run most tasks for CPCs that are in an ensemble. However, it cannot perform any ensemble management tasks.

The HMC can manage up to 100 CPCs. Eight of these CPCs can be members of an ensemble. A CPC that is not a member of an ensemble can be managed by up to 32 HMCs. A single node can be a member of only one ensemble.

The HMC uses VLAN and an included PCI Express Ethernet adapter for handling both single and dual Ethernet configuration. The HMC is supplied with two Ethernet ports.

The physical location of the HMC hardware features (standard, optional, or both) is dictated by the specific PC. Some features can be mutually exclusive with other features depending on the PC model. Each CPC must be connected to at least one HMC on the same network as the Support Elements of the CPC.

Figure 7 illustrates multiple HMCs controlling multiple systems.




Figure 7. Hardware Management Consoles

Each HMC in Figure 7 control different systems.

- Primary HMC (A) can directly control any system including ensemble-specific functions.
- Alternate HMC (A1) cannot directly control any CPC at this time.
- Primary HMC (A) and alternate HMC (A1) must be located on the same network subnet.
- HMC (B) and HMC (C) can directly control zEnterprise(1), zEnterprise (2), and zEnterprise (3) systems, except ensemble-specific functions are unavailable for system zEnterprise (1) and zEnterprise (2).
- HMC (A1), HMC (B), and HMC (C) can use the **Remote Hardware Management Console** task to perform ensemble-specific tasks through HMC (A).

System z platform management concepts extended

Platform management involves many components, such as:

- Ensemble Membership Management
- Performance Management
- Availability Management (Unified Resource Manager version 2.12.1)
- Workload Membership Management
- Virtual Server Management
- Virtual Network Management
- Energy Management
- Entitlement Management

- zEnterprise System Web Services API

Together these platform components provide:

- Workload optimized resource management
 - Integration of advanced platform virtualization and management functions into the zEnterprise HMC and Support Element infrastructure
 - Management of servers, storage, and networks as single logical system
- A single delivery mechanism
 - Platform management components packaged and serviced as integrated zEnterprise firmware using existing firmware delivery, deployment, upgrade, and service mechanisms
 - Extensions to existing zEnterprise HMC and Support Element hardware appliances

Many of System z platform management tasks on the HMC have been extended to support the IT infrastructure for the heterogeneous resources integrated into zEnterprise:

- Change management enhancements for heterogeneous resources:
 - View firmware information
 - Retrieve firmware changes
 - Change firmware levels
 - Backup and restore critical data
- Configuration management enhancements for heterogeneous resources:
 - Configure physical hardware
 - Discover, validate, and inventory hardware
- Operations management enhancements for heterogeneous resources:
 - Power on and power off hardware
 - Perform event notification based on logged events or state changes
 - Support device console
- Performance management enhancements for heterogeneous resources:
 - Monitor and report performance goals. These functions provide system administrators and performance analysts with the data they need to understand whether performance goals are being met.
 - Manage virtual server processor capacity. You can use this function to manage processor resources across virtual servers based on a goal-oriented performance policy.
 - Load balancing recommendations control the distribution of incoming work. This function facilitates distributing incoming work to best achieve the goals of your workload performance policies.
 - Processor management. This feature provides Unified Resource Manager control of processor shares for the virtual servers.
- Availability management for heterogeneous resources:
 - Provides availability management for workload resource groups in an ensemble.
 - Provides consistent monitoring, analysis, and reporting of hard and soft failures of virtual servers and element groups, analyzed in the context of the workload resource group.
 - Highlights out of compliance conditions, based on the redundancy objectives that are set in the availability policy for a workload resource group.
 - Integrates with other key Unified Resource Manager disciplines such as performance management.
 - Audit and log capabilities
- Energy management enhancements for heterogeneous resources:
 - Enhanced **System Activity Display** task
 - Monitor energy consumption and temperature

- Report energy data to IBM using **Transmit System Availability Data** task
- Problem management enhancements for heterogeneous resources:
 - Log errors automatically
 - Analyze problems and use call-home reporting
 - View hardware messages
 - View open problems
 - Collect problem data
- Business management enhancements for heterogeneous resources:
 - Manage user roles
 - Audit security
 - Configure and manage networks

Ensemble infrastructure

ensemble introduces two new internal networks. These networks are the intraensemble data network (IEDN) and the intranode management network (INMN). Existing external customer-managed networks are also present. The requirements vary depending on your configuration. Considerations include whether nodes within an ensemble have an attached zBX and whether you plan to connect zEnterprise systems to existing System z networks.

Inside a node

Intraensemble data network (IEDN)

At a high level, the IEDN can be represented as shown in Figure 8.




Figure 8. Simple view of intraensemble data network

Physically, in a node, a zEnterprise central processor complex (CPC) connects to a zBX by using top-of-rack (TOR) switches in the zBX rack. An Open Systems Adapter-Express5S (OSA-Express5S), Open Systems Adapter-Express4S (OSA-Express4S) or Open Systems Adapter-Express3 (OSA-Express3) 10 Gigabit Ethernet (GbE) feature (CHPID type OSX) provides connectivity to the IEDN. An IEDN is required for application data communications between the managing zEnterprise CPC and attached zBX (single node). The IEDN is optional for application data communications between nodes.

Select IBM blades and optimizers (such as XI50z) within a zBX have dedicated interfaces that are connected to the IEDN. This connection is through Electronic Switch Modules (ESMs) located in the zBX BladeCenter. ESM uplinks also attach to the TOR switches.

An IEDN enables communication to flow along these network components and thus between the nodes that comprise an ensemble. Specifically, an IEDN provides the following functions:

- zEnterprise applications can communicate between operating system images to share data.
- zEnterprise applications can communicate to a zBX.
- zBX applications can communicate -to- another zBX within an ensemble.

A 10 GbE OSA data link (Layer 2) network supports each IEDN. Virtual servers and optimizers connect to the IEDN for communication in support of customer workloads. IBM provides and manages an IEDN as part of the ensemble. Specifically, an HMC manages the ensemble through its user interface. The HMC includes networking tasks that are collectively known as network virtualization function tasks.

An IEDN requires two OSA Express5S, OSA-Express4S or OSA-Express3 10 GbE adapters (thus ensuring redundancy). The IEDN supports running IPv4 or IPv6 protocols. The IEDN default is one virtual LAN, and you can define multiple VLAN subnets.

The IEDN also can support an internal LAN segment for the HiperSockets network. Each zCPC can have one IQD CHPID with IEDN access defined to enable IEDN HiperSockets (Internal Queued Direct I/O Extensions function IQDX) communication to the virtual servers. This enables the HiperSockets network to extend outside the CPC to the entire ensemble. The z/VM virtual switch bridge support is used to connect a virtual switch to a HiperSockets network, essentially bridging the HiperSockets network to the virtual switch's simulated network and the physical IEDN network. The IEDN IQDX HiperSockets network is a flat Layer 2 network. See, Figure 9 on page 16.

Customer experience with the IBM zEnterprise BladeCenter Extension (zBX) led IBM to depart from its original requirement to exclusively support Layer 3 connectivity between the external data network and the 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)*, available at website <http://www.redbooks.ibm.com/redpieces/abstracts/redp5036.html>, includes a description of limitations and trade-offs when deploying Layer 2 versus Layer 3 connectivity.

Communications through the IEDN have the benefit of being performed through a physical network within the machine and, therefore, are more secure than external network communications.

Intranode management network (INMN)

An INMN is required for platform management within a node. The network enables an SE to communicate with the select IBM blade hypervisors and optimizer blades that it controls. An ensemble must have an INMN.

An INMN is contained within a single node; it cannot span multiple nodes. IBM provides the INMN as a private 1000BASE-T Ethernet network. However, unlike an IEDN an INMN configuration (VLAN, IP addressing) is automatically deployed by the HMC and its network virtualization function tasks.

An Express5S, OSA-Express4S or OSA-Express3 (CHPID type OSM) provides zEnterprise connectivity to the INMN. An INMN requires two ports from two different Express5S, OSA-Express4S or OSA-Express3 1000BASE-T Ethernet adapters thus ensuring redundancy. The INMN supports only IPv6 protocols.

Customer-provided data network

A customer-provided data network provides network communication external to the IEDN.

Depending on factors such as network traffic patterns and load balancing requirements, you can connect the customer-provided data network to the IEDN in these ways, see Figure 9 on page 16:

- Connect through a TOR switch by a router or an OSA associated with a zEnterprise CPC. You would typically do this connection in a migration scenario.
- Connect through an OSA channel using CHPID type OSD, using a System z router.

You manage the customer-provided data networks. The HMC does not manage customer-provided data networks or the virtualization of the network.

The z/VM virtual switch support provides the option to connect the IQD HiperSockets channel to the customer-provided data network. The vSwitch bridge support is used to connect the virtual switch to a HiperSockets network, essentially bridging the HiperSockets network to the virtual switch's simulated network and the customer-provided data network. z/VM virtual switch bridge support provides access to the customer-provided data network for QDIO Enhanced Buffer State Management facility (QEBSM) guests. See, Figure 10 on page 16.

Customer-provided management network

You provide and manage the customer-provided management network (HMC network). The HMC does not manage the customer-provided management network.

The HMC uses the customer-provided management network to access a specific node in the ensemble. This network also provides HMC and Support Element cross-node communication within the ensemble.

Private system control network

The PSCN is an Internet Protocol version 4 (IPv4) network and is statically configured. If zBX units are present, the PSCN connects the zEnterprise CPC Bulk Power Hub to the TOR switches. The connection uses 26 meter 1000BASE-T Ethernet cables. IBM firmware provides and manages the PSCN. The HMC does not manage the PSCN.

A private system control network (PSCN) connects the Support Element to the hardware control elements within a node. This network supports problem and error detection and is essential for serviceability.

Figure 9 on page 16 and Figure 10 on page 16 diagram the constituent networks and supported vSwitch configurations in and ensemble. The internal PSCN is not visible in the figure.




Figure 9. Networks in an ensemble with an IEDN vSwitch bridge to zCPC HyperSockets




Figure 10. Networks in an ensemble with an vSwitch Uplink to a customer-provided data network

Configuration examples

Each ensemble configuration requires:

- A primary HMC, and alternate HMC, used to control and manage the ensemble.
 - Only one pair of HMCs can be included in an ensemble. One HMC (the primary) manages the ensemble, and an alternate HMC is also included.
 - Both HMCs are attached to the customer-provided management network and are located on the same network subnet.
 - The alternate HMC is a mirrored backup of the HMC managing the ensemble and it provides limit function when the primary HMC is functional.

Example: Basic ensemble configuration rules: without zBX

An ensemble node consisting of a zEnterprise without an attached zBX is supported, see Figure 11.

A node without a zBX *requires* the following:

- Intranode management network (INMN)
- Customer-provided management network (HMC network)

A node without a zBX *might require* the following, depending on the System z applications:

- Intraensemble data network (IEDN)
 - For an ensemble with a single zEnterprise system, an LC DUPLEX cable between 2 OSX CHPIDs is required.
 - For an ensemble with 2 - 8 zEnterprise CPCs, a single zBX rack is required.
- Customer-provided data network




Figure 11. Basic ensemble configuration without zBX

Example: Basic ensemble configuration rules: with zBX

Each node with zBX *requires* the following, see Figure 12:

- Intranode management network (INMN)
- Customer-provided management network (HMC network)
- Private system control network

Each node with zBX *might require* the following, depending on the System z applications:

- Intraensemble data network (IEDN)
- Customer-provided data network




Figure 12. Basic ensemble configuration with zBX

You can extend the IEDN (optional connection) to all nodes of an ensemble, as illustrated in Figure 13 on page 19.

Example: Ensemble configuration with zEC12 or zBC12 and z196 or z114 servers, each with zBX attached




Figure 13. Ensemble configuration with zEC12 or zBC12 and z196 or z114 systems, each with zBX attached

The ensemble infrastructure supports external ISC-3, PSIFB, and customer network connections, as illustrated in Figure 14. The systems that are connected by external ISC-3, PSIFB, and customer networks are not a member of the ensemble.

Example: Mixed system environment




Figure 14. Mixed system environment

A similar mixed system environment is possible by using zEC12, zBC12, and a zBX 2458-003.

Maximum distances between network connections

The following list summarizes maximum distances between network connections:

- Intraensemble data network
 - If multimode fiber (short range optics):
 - With 50 μ m fiber at 2000 MHz-km: 300 m (984 feet)
 - With 50 μ m fiber at 500 MHz-km: 82 m (269 feet)
 - With 62.5 μ m fiber at 200 MHz-km: 33 m (108 feet)
 - If single-mode fiber (long range optics): 10 km (6.2 mi)
- Intranode management network
 - zEnterprise server to the zBX TOR switches: 26 m (85.3 feet) or 100 m (328 feet)
IBM recommends that the zBX is placed near the zEnterprise server.
- Fibre channel Storage Area networks
 - If multimode fiber at 8 Gbps:
 - With 50 μ m fiber at 2000 MHz-km: 150 m (492 feet)
 - With 50 μ m fiber at 500 MHz-km: 50 m (164 feet)
 - With 62.5 μ m fiber at 200 MHz-km: 21 m (69 feet)
 - If multimode fiber at 4 Gbps:
 - With 50 μ m fiber at 2000 MHz-km: 380 m (1247 feet)
 - With 50 μ m fiber at 500 MHz-km: 150 m (492 feet)
 - With 62.5 μ m fiber at 200 MHz-km: 70 m (230 feet)
 - If multimode fiber at 2 Gbps:
 - With 50 μ m fiber at 2000 MHz-km: 500 m (1640 feet)
 - With 50 μ m fiber at 500 MHz-km: 300 m (984 feet)
 - With 62.5 μ m fiber at 200 MHz-km: 150 m (492 feet)
- Customer-provided management network (HMC network)
 - 100 m (328 feet)

Supported hardware, firmware, and software

The hybrid zEnterprise infrastructure supports the following hardware, firmware, and software elements.

Hardware

Supported hardware in an ensemble node.

zEnterprise hardware

Supported zEnterprise hardware in an ensemble node.

Table 1. Supported zEnterprise hardware.

	zBC12 (MT 2828)	zEC12 (MT 2827)	z114 (MT 2818)	z196 (MT 2817)
CPC				
books or drawers	1 or 2 drawers	1 - 4 books	1 or 2 drawers	1 - 4 books
CPs, zIIPs, zAAPs, IFLs, or ICFs	1 - 13	1 - 101	1 - 10	1 - 80
Maximum number of SAPs	4	32	4	32
Memory				

Table 1. Supported zEnterprise hardware. (continued)

	zBC12 (MT 2828)	zEC12 (MT 2827)	z114 (MT 2818)	z196 (MT 2817)
Maximum per book or drawer	240 GB first drawer	0.75 TB per book	120 GB first drawer	0.75 TB per book
Maximum	496 GB	3 TB	248 GB	3 TB
Per partition	up to 496 GB	up to 1 TB	up to 248 GB	up to 1 TB
I/O				
Cargo cages	0	0 - 1	0	0 - 2
I/O drawers	0 - 2	0 - 6	0 - 4	0 - 6
PCIe I/O drawers	0 - 2	0 - 6	0 - 2	0
Maximum number of channels		1024	240	1024
Power and cooling				
Number of frames	1	2	1	2
Refrigerant-cooled	No	No	No	Yes
Water-cooled	No	Yes	No	Yes
Radiator-cooled	No	Yes	No	No
Air-cooled	Yes	No	Yes	No
Power source	AC or DC	AC or DC	AC or DC	AC or DC
zBX (MT 2458)				
Model 002	No	No	Yes	Yes
Model 003	Yes	Yes	No	No

For additional information about zEnterprise servers:

On [ibm.com](http://www.ibm.com/systems/z/hardware/zenterprise/index.html)[®], go to <http://www.ibm.com/systems/z/hardware/zenterprise/index.html>

On **Resource Link**, go to <http://www.ibm.com/servers/resourcelink> and click **Library** on the navigation bar.

IBM zEnterprise BladeCenter Extension (zBX)

Supported IBM zEnterprise BladeCenter Extension (zBX) hardware in an ensemble node.

The machine type (MT) is 2458 and is composed of the following sub-systems:

- Has one to four racks
- Has private service, management, and data network switches
- Uses 1 - 2 BladeCenter chassis per rack
- Supports zBX blades types:
 - Select IBM POWER[®] blades, 0 - 112 blades
 - Select IBM System x blades, 0 - 56 blades
 - DataPower XI50z, 0 - 28 blades
- zBX is ordered by specifying the number of optimizers and IBM blades required by the managing zEnterprise CPC
- The zEnterprise order determines all required infrastructure (blade, BladeCenters, switches, and racks) for zBX.
- Model 002, all blades have up to a 10 GbE network interface bandwidth to the IEDN TOR switches.
- Model 003, all blades have up to a 20 GbE network interface bandwidth to the IEDN TOR switches.

- All OSX connections must be point to point, you can place no switches or routers in the path between the OSX and zBX.
- Only one zBX per managing zEnterprise CPC . The managing CPC must be a zEnterprise CPC.
- IBM recommends that the zBX be placed adjacent to the zEnterprise server.
- All other types of customer networks (for example, OSC, OSD, OSE, and OSN connections) are managed outside the ensemble.

Items procured from a source other than IBM System z:

- All select IBM blades.
- Storage devices (DASD)
- All Ethernet optical cables
- OSA to zBX cables
- zBX to zBX cables
- Private FC DASD (from supported list)
- All 8 GB FC optical cables

For additional information about zBX MT 2458 on **Resource Link**, go to <http://www.ibm.com/servers/resourcelink> and click **Library** on the navigation bar.

For a list of supported IBM blades or supported storage devices, go to <http://www.ibm.com/systems/z/hardware/zenterprise/zbx.html>.

Hypervisors and hardware firmware

Supported hypervisors and firmware in an ensemble node.

Hypervisors

Supported hypervisors in an ensemble.

z/VM The z/VM hypervisor is designed to help extend the business value of mainframe technology across the enterprise by integrating applications and data while providing exceptional levels of availability, security, and operational ease. With z/VM virtualization technology, you can run hundreds of thousands of Linux servers on a single server running with other zEnterprise operating systems, such as z/OS, or a large-scale Linux-only enterprise server solution. With z/VM V6.2 you can improve productivity by managing ensembles of integrated hardware and enterprise wide workloads with Unified Resource Manager.

PR/SM™

Processor Resource/System Manager (PR/SM) is integrated into all zEnterprise models. PR/SM runs directly on the hardware to control hardware resources and monitor guest virtual machines. It transforms physical resources into virtual resources, so that many logical partitions can share the same physical resources. The zEnterprise firmware manages the PR/SM hypervisor.

Using PR/SM you can divide physical system resources (dedicated or shared) into isolated logical partitions. Each logical partition operates like an independent system running its own operating environment. On the latest zEnterprise zEC12 models, you can create up to 80 logical partitions running z/VM, z/OS, z/OS.e, or Linux on System z, on a single system.

PR/SM enables each logical partition to have dedicated memory (which you can dynamically re-configure as needed) and dedicated or shared processors and I/O. Logical partitions with dedicated resources own the resources to which they are assigned. Logical partitions with shared resources appear to own the resources to which they are assigned, but many logical processors share the resources. In other words, PR/SM transforms physical resources into virtual resources so that many logical partitions can share the same physical resources.

The following figure shows a physical system running three logical partitions. Each logical partition is running an operating environment and applications independent of the other logical

partitions, while sharing the same physical resources.




Figure 15. View of the PR/SM hypervisor layer

PR/SM enables logical partitions to share processor and I/O resources. Multiple Image Facility (MIF) and multiple-channel subsystem (MCSS) are the root technologies that enable I/O resource sharing at the hardware level. These technologies enable I/O channel paths and subsets of I/O devices attached to those channel paths to be shared among logical partitions.

Use PR/SM to dynamically add and delete resources like processors, I/O, and memory across logical partitions while the partitions are actively in use. When you apply this dynamic resource allocation or dynamic LPAR you can dynamically redefine all available system resources to reach optimum capacity for each partition.

PowerVM Enterprise edition

PowerVM Enterprise edition on select POWER blades offers virtualization capabilities for AIX®. You can define servers down to 0.10 of a processor and flex server size with demand. Use PowerVM network support to share expensive disk and network resources, while minimizing any management and maintenance costs.

zEnterprise POWER blade LPARs can be either dedicated, which means they consume full physical PUs or shared, in which each virtual CPU has to be backed by at least a tenth of physical PU. Each blade always has the VIOS partition defined. This partition consumes a 10th of every physical processor on the system. Therefore, you can have up to $9 * n$ virtual servers per blade, "n" being the number of physical PUs. For blades that have 4 or 8 PUs the number of available virtual servers is 36 or 72. The number could be fewer, if the virtual servers are defined with more than one virtual PU. Another factor that could limit the number of virtual servers is the entitled capacity, which is independent of the number of virtual PUs. For example, a virtual server with one virtual PU and a minimum entitled capacity of 0.5 reserves 50% of a physical PU.

PowerVM provides a secure and scalable virtualization environment for AIX applications built upon the advanced RAS features and leading performance for the select POWER blades.

PowerVM can virtualize processor, memory, and I/O resources to increase asset utilization and reduce infrastructure costs. You can adjust processor capability to meet changing workload demands.

The PowerVM hypervisor is always active on a POWER blade. The PowerVM Hypervisor is responsible for dispatching the virtual server workload across the shared physical processors. The PowerVM Hypervisor also enforces security between virtual servers and can provide communication between virtual servers that enables the virtual network functions.

Combining these PowerVM features can help transform computing resources into the agile, responsive, and energy efficient infrastructure demanded by IT enterprises. zEnterprise firmware manages the PowerVM hypervisor installation and code updates.

zBX integrated hypervisor for IBM System x blades (using Kernel-based Virtual Machine)

zEnterprise firmware manages the integrated hypervisor installation and code updates. Utilizing the zBX integrated hypervisor for IBM System x blade, provides for logical device integration between zEnterprise and System x blades for multi-tiered applications.

Starting with version 2.12.1 of Unified Resource Manager, Unified Resource Manager provides dynamic processor management for virtual servers deployed on IBM System x blades based on policy goals.

The maximum number of active virtual servers for x Hyp is 99. Attempting to activate more than 99 virtual servers will result in the error: ERROR - cannot start virtual server, since LPAR limit is exceeded on hypervisor.

The zBX integrated hypervisor for IBM System x blades is abbreviated as INxHY or x Hyp.

Hardware firmware

Additional firmware managed by the zEnterprise ensemble solution.

- OSX
- AMM
- zBX switches

Software

Supported software in an ensemble node.

z/OS

z/OS is a 64-bit operating system designed to offer a stable, secure, continuously available, and scalable environment for applications running on a System z server.

z/OS changes for zEnterprise

z/OS V2.1, R10, R11, R12, R13; system exploitation with appropriate PTF, z/OS R9; system toleration

- OSA-Express QDIO interfaces for new CHPID types OSM and OSX:
 - OSM provides connectivity to the intranode management network
 - OSX provides connectivity to the intraensemble data network
- Connectivity to new CHPID types allowed only if the LPAR is a member of an ensemble
- Minimal configuration required for OSAs that are part of the IEDN
- OSM IPv6 link is local and automatically created
- You define the OSX connectivity with the INTERFACE statement in the TCP/IP profile
- With guest platform management provider (GPMP) installed, you can benefit from these advanced capabilities and additional performance data:
 - Cooperative performance management with z/OS workload manager, which enables z/OS WLM to manage the z/OS segments of multi-tiered workloads with knowledge of the overall workload goals.
 - Load balancing recommendations that help optimize network workload distribution within the ensemble.
 - Performance monitoring and data reports
 - Collection of application data from IBM middleware products that have been instrumented with the Open Group Application Response Measurement (ARM) 4.0 standard.
- z/VM guests running z/OS

- With z/OS V1.13 (or later) and z/OSMF V1.13 (or later), RMF™ CIM based performance data can be gathered for Linux on System z, Linux on System x®, and AIX® systems to provide a consistent monitoring solution for zEnterprise ensembles. The Resource Monitoring plug-in for the z/OS Management Facility, first made available with z/OSMF V1.12, combines performance metrics from those platforms and z/OS metrics in common graphic views.

For more information about establishing z/OS connectivity to the new CHPID types, see the section about using TCP/IP in an ensemble in the *z/OS Communications Server IP Configuration Guide*.

z/VM

z/VM is a 64-bit virtual machine hypervisor. The z/VM hypervisor runs other operating systems in the virtual machines it creates. IBM mainframe operating systems, such as z/OS, z/VM, and Linux on System z, can be run as guest systems in their own virtual machines. z/VM can run any combination of guest systems.

z/VM emulates and distributes system resources so that many virtual machines can share the same physical resources. The z/VM hypervisor is not zEnterprise firmware. However, the zEnterprise management firmware can optionally manage the z/VM hypervisor and its virtual resources.

With z/VM, you can create tens to thousands of virtual machines within one logical partition. Like logical partitions, each virtual machine operates as an independent system. z/VM emulates the existence of available physical resources (system functions, storage, and I/O) to each virtual machine as a shared resource pool. Thus, z/VM has great control over how and to what extent a specific virtual machine uses the physical resources.

The following figure shows a system with two logical partitions, each running z/VM. Both logical partitions are running independently of each other but sharing the same physical resources. Each z/VM instance is running three virtual machines. Each virtual machine is running an operating environment and applications independent of the other virtual machines while sharing the same physical resources.




Figure 16. View of z/VM instance running three virtual machines

The number of virtual machines you can create is limited only by the resources available.

z/VM helps to improve productivity by hosting workloads such as z/OS, z/OS.e, and Linux on System z on the same System z server. z/VM also supports z/VM as a guest operating system.

- Use the HMC to manage:
 - Guests and to install Linux on System z

- Installation of z/VM in a virtual machine from the HMC DVD drive without having to establish an external network connection.
- Expanded guest support for specialty processors, which permits all processor types (CPs, IFLs, zIIPs, zAAPs, and ICFs) to be defined in the same z/VM LPAR.
- When managed in an Ensemble, z/VM provides these advanced facilities:
 - Redistribution of CPU shares to z/VM virtual processors, which helps virtual machines to be managed more efficiently.
 - New function that can automatically start and stop virtual processors based on virtual processor utilization and workload characteristics.

z/VM changes for zEnterprise

z/VM is essentially the same operating system within or without an ensemble. When participating in the ensemble environment, z/VM is regarded as another hypervisor in the ensemble. z/VM automatically detects when it is running in an ensemble or when it has been removed from an ensemble. When a part of the ensemble, z/VM and its virtual servers must be managed from the HMC. From the HMC, you can manage the lifecycle of z/VM itself and of the virtual servers (provisioning, starting, stopping, and so on). Management of operating system images and application programs, however, are managed as done previously.

z/VM also participates in the enhanced monitoring facility and static power savings mode when those hardware components are present. CP automatically detects their availability and supports them.

The HMC communicates with z/VM using the SCLP (service call logical processor) interface with an internal management network called the intranode management network (INMN). Data communications within the ensemble are handled through another internal network called the intraensemble data network (IEDN).

The IEDN enables you to set up connections between the virtual servers running in the ensemble. From z/VM, the virtual servers attach to the IEDN through virtual switches. The virtual servers can also attach to the external network through the OSA.

For z/VM to participate in the ensemble, the following configuration tasks are required:

- Set up the OSM and OSX OSA channel paths
- Define controllers for the IEDN and INMN networks
- Configure DirMaint™ (or other user directory manager)
- Configure SMAPI servers

When participating in the ensemble environment managed from the HMC by the Unified Resource Manager, you cannot manage z/VM from a z/VM operator console. The HMC becomes the management focus for the following tasks:

- Create virtual servers on z/VM
- Manage network connections for virtual servers
- Manage storage resources for virtual servers
- Monitor performance of virtual servers and z/VM
- Manage hypervisor networking resources

The Unified Resource Manager network support enhances the virtual and real networking components of z/VM to provide integrated management of these resources within an ensemble.

The following command is new for this support:

- QUERY NIC

The following commands and/or their responses have been updated for this support:

- COUPLE
- DEFINE NIC
- DEFINE VSWITCH
- QUERY CONTROLLER
- QUERY LAN
- QUERY VIRTUAL NIC
- QUERY VMLAN
- QUERY VSWITCH
- SET LAN
- SET NIC
- SET VMLAN
- SET VSWITCH

For more information about managing z/VM and its virtual servers in an ensemble, see the hardware operations guides such as:

- *zEnterprise System Ensemble Planning and Configuring Guide*
- *zEnterprise System Ensemble Performance Management Guide* (Unified Resource Manager version 2.12.0 or prior)
- *zEnterprise System Hardware Management Console Operations Guide for Ensembles* (Unified Resource Manager version 2.12.0 or prior)
- For Hardware Management Console (HMC) 2.12.1 user interface or task information, go to the console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmc/v2r12m1/index.jsp>.
- *CP Planning and Administration* available from the z/VM bookshelf on website <http://www.vm.ibm.com/library/>
- *Systems Management Application Programming (SMAPI)* available from the z/VM bookshelf on website <http://www.vm.ibm.com/library/>
- The HMC help system

z/VM V6.1 with applicable PTFs and V6.2 with applicable PTFs support zEnterprise ensemble infrastructure.

Features in IBM z/VM V6.1 (end of support is effective May 1, 2013) are designed to provide:

- Full participation in the creation, deletion, and management of virtual server resources from the Unified Resource Manager
- Direct and secure access to the zEnterprise intranode management network (INMN) and intraensemble data network (IEDN)
- Support for the zEnterprise zEC12, z114, and z196
- Improved security with support for secure-key encryption capabilities of the CP Assist for Cryptographic Functions (CPACF) and support for the Federal Information Protection Standard (FIPS) 140-2
- Improved resiliency through z/VM Extended Remote Copy (XRC) time stamping, which z/VM uses to participate in long-distance environments that require synchronized I/O such as GDPS®, which has a planned availability date of December 2010.
- Increased guest and system throughput with improvements to z/VM memory management scalability
- Support for Crypto Express4S

The z/VM System Management Application Programming Interface (SMAPI) and the z/VM Directory Maintenance Facility (DirMaint) have been enhanced to enable the IBM Unified Resource Manager to manage the z/VM hypervisor and the virtual servers it supports. The SMAPI enhancements provide APIs and function in the following functional areas:

- System configuration file management
- User Directory management
- Disk management
- Network management
- Guest facilities management
- Directory parsing
- API function level interrogation

z/VM V6.1 with the PTF for APAR VM64822 provides z/VM SMAPI, Performance Toolkit, and DirMaint support for the Unified Resource Manager.

z/VM provides access to the zEnterprise INMN and the IEDN through both its real and virtual networking capabilities. As a result, the deployment and management of z/VM network topology is integrated into the zEnterprise environment and is managed by the Unified Resource Manager. After the zBX is connected to a zEnterprise CPC through the INMN, the resources of the zBX are a logical extension of the zEnterprise CPC. You can then control and manage these resources through the HMC and Support Element of the zEnterprise CPC. The INMN requires OSA-Express5S or Express3 1000BASE-T and the IEDN requires a port on two separate OSA-Express5S, OSA-Express4S, or OSA-Express3 10 GbE features.

z/VM V6.1 with the PTF for APAR VM64822 provides z/VM Networking and Performance Toolkit support for the zEnterprise INMN and IEDN networks.

z/VM V6.2 with PTFs provides support for HiperSockets over the IEDN. HiperSockets Integration with the IEDN extends the reach of HiperSockets network outside the zCPC to the entire ensemble. See Intraensemble data network (IEDN) on page 14.

z/VM V5.4, V6.2, or V6.3 is required to define, modify, and delete OSX (for IEDN) or OSM (for INMN) CHPID types when z/VM is the controlling LPAR for dynamic I/O.

For the latest information on ensemble management support for z/VM see website <http://www.vm.ibm.com/service/vmrequirm.html>.

Unified Resource Manager, z/VM V6.3, and OpenStack

Unified Resource Manager is designed to provide systems management capabilities across the multi-architecture environment of zEnterprise. Customers planning to exploit these system management capabilities might 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 System z is making contributions to the OpenStack project that are designed to enable z/VM V6.3 to be the first System z 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 zEnterprise integrated firmware hypervisors - PR/SM, PowerVM, and zBX integrated hypervisor for IBM System x blades (using Kernel-based Virtual Machine).

Instructions for migrating IEDN managed devices to native z/VM management can be found on website <http://www.vm.ibm.com/service/vmrequrm.html>.

z/VM V6.2 Single System Image (SSI) recommendation within a zEnterprise ensemble

If z/VM is configured to be ensemble managed, the z/VM image should not be part of an SSI cluster or CSE complex.

Linux on System z

Linux operating systems that have been compiled to use the Integrated Facility for Linux (IFL) engines on System z. The IFL is a processor dedicated to Linux workloads on IBM System z servers. The IFL is supported by the z/VM virtualization software and the Linux operating system.

Supported distributions are:

- Red Hat Enterprise Linux (RHEL) 5 and 6 or later.
- SUSE Linux Enterprise Consolidation Suite (SLES) 10 and 11 or later.

Optimizations

The zBX infrastructure works with the IBM zEnterprise CPC to enhance workload optimization by delivering an integrated hardware and software platform.

DataPower XI50z

The DataPower XI50z is a combined hardware and software solution that simplifies, helps govern, and enhances the security of XML and IT services.

The DataPower XI50z software architecture is simple. There is a customized, hardened, native-code operating system kernel that implements core functionality. The OS resides in the firmware and can be updated by applying updated firmware images.

On top of this layer is the functionality that is implemented in XSLT style sheets. These style sheets are read-only. The appliance uses these style sheets to implement specific functionality.

The next layer is the software stack that consists of configurations that you develop. These configurations are the application proxies and processing policies to process message traffic for your application. Configuration files and application artifacts can reside in the file system on the appliance or can be stored on remote servers and retrieved and cached at startup. Storage and retrieval from remote servers means that in a highly secure environment these files and artifacts never reside on the appliance.

The DataPower XI50z is ordered as a separate System z Machine Type (2462-4BX). The following 4 firmware loads are available:

- Base IBM includes: IBM WebSphere DataPower base, Tivoli® Access Manager, and Application Optimization
- Base IBM + TIBCO EMS
- Base IBM + SQL-ODBC
- Base IBM + TIBCO EMS + SQL-ODBC

For more information, see the WebSphere DataPower Integration Appliance Information Center at:

Version	Link
4.0.1	http://publib.boulder.ibm.com/infocenter/wsdatap/4mt/index.jsp?topic=%2Fcom.ibm.dp.xi.doc%2Fwelcome.htm
5.0.0	http://pic.dhe.ibm.com/infocenter/wsdatap/v5r0m0/index.jsp?topic=%2Fcom.ibm.dp.xi.doc%2Fwelcome.htm
6.0.0	http://pic.dhe.ibm.com/infocenter/wsdatap/v6r0m0/index.jsp?topic=%2Fcom.ibm.dp.xi.doc%2Fwelcome.html

Select IBM blades

The IBM blades can be installed in the zBX, enabling application integration with System z transaction processing, messaging, and data serving capabilities. The blades are managed as part of a single logical virtualized environment by Unified Resource Manager.

POWER blade

With PowerVM installed, the POWER blade offers the possibility of running an application certified to run on the AIX levels that are shown in Table 2. For the most up-to-date information, see the IBM zEnterprise BladeCenter Extension (zBX) website at http://www.ibm.com/systems/z/hardware/zenterprise/zbx_specs.html.

Table 2. Supported operating systems and hypervisors on select POWER blades

Operating Systems	Supported Levels on hardware:	
	IBM BladeCenter PS701 Express Blades installed in zBX Model 002	IBM BladeCenter PS701 Express Blades installed in zBX Model 003
	IBM zEnterprise 114 (z114) and IBM zEnterprise 196 (z196)	IBM zEnterprise EC12 (zEC12) and IBM zEnterprise BC12 (zBC12)
AIX Note: AIX Express, Standard, or Enterprise Edition can be ordered.	• AIX 5.3 TL12+ and up • AIX 6.1 TL5+ and up • AIX 7.1 and up	• AIX 5.3 TL12+ and up • AIX 6.1 TL5+ and up • AIX 7.1 and up
Hypervisors	PowerVM Enterprise Edition VIOS 2.2.1	PowerVM Enterprise Edition VIOS 2.2.2

IBM BladeCenter PS701 Express and the IBM zEnterprise System planning information can be found at http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_USEN&htmlfid=ZSY03019USEN&attachment=ZSY03019USEN.PDF.

System x blade

With zBX integrated hypervisor for IBM System x blades (using Kernel-based Virtual Machine) installed, the System x blade offers the possibility of running an application certified to run on 64-bit Windows and Linux.

Supported 64-bit Windows and Linux operating systems on select IBM System x blades are shown in Table 3. For the most up-to-date information, see the IBM zEnterprise BladeCenter Extension (zBX) website at http://www.ibm.com/systems/z/hardware/zenterprise/zbx_specs.html.

Table 3. Supported operating systems and hypervisors on select IBM System x blades

Operating Systems	Supported Levels on Hardware:	
	IBM BladeCenter HX5 Blades installed in zBX Model 002	IBM BladeCenter HX5 Blades installed in zBX Model 003
	IBM zEnterprise 114 (z114) and IBM zEnterprise 196 (z196)	IBM zEnterprise EC12 (zEC12) and IBM zEnterprise BC12 (zBC12)
Linux	- Red Hat Enterprise Linux (RHEL) 5.5 and up - Red Hat Enterprise Linux (RHEL) 6.0 and up	- Red Hat Enterprise Linux (RHEL) 5.5 and up - Red Hat Enterprise Linux (RHEL) 6.0 and up
	- SUSE Linux Enterprise Server (SLES) 10 (SP4) and up - SUSE Linux Enterprise Server (SLES) 11 (SP1) and up	- SUSE Linux Enterprise Server (SLES) 10 (SP4) and up - SUSE Linux Enterprise Server (SLES) 11 (SP1) and up
64-bit Microsoft Windows Note: IBM recommends the DataCenter Edition for all Windows operating system levels	- Microsoft Windows Server 2008 R2 - Microsoft Windows Server 2008 (SP2)	- Microsoft Windows Server 2012 - Microsoft Windows Server 2008 R2 - Microsoft Windows Server 2008 (SP2)
Hypervisors	Red Hat Enterprise Virtualization Hypervisor (RHEV-H) 6.1	Red Hat Enterprise Virtualization Hypervisor (RHEV-H) 6.4

IBM BladeCenter HX5 (7873) Planning Information for the IBM zEnterprise System can be found at http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_USEN&htmlfid=ZSL03128USEN&attachment=ZSL03128USEN.PDF.

Chapter 3. zEnterprise ensemble system management

zEnterprise with or without an attached zBX, benefits from resource management as an ensemble. Ensemble management tasks are tightly integrated into the firmware, providing for integration of workload, performance, energy, configuration, change, problem, operation, and security management in support of service-oriented workloads. In zEnterprise, the term *workload resource group* is defined as a collection of virtual servers that perform a customer-defined collective purpose. A workload resource group generally can be viewed as a multi-tiered application. Each workload resource group is associated with a set of policies that define performance goals. Resources are grouped and managed logically as a single system. Unified Resource Manager provides the support to manage the ensemble resources.

The Unified Resource Manager includes a combination of new and existing system management tasks. The tasks are designed to provide a single, cohesive management context that applies across all managed objects of the ensemble.

- **Hypervisor Management**

Provides tasks for managing hypervisor lifecycle, managing storage resources, performing RAS/FFDC and monitoring the supported hypervisors.

- **Ensemble membership management**

Provides tasks for creating an ensemble and controlling membership of the ensemble.

- **Storage Management**

Provides a common user interface for allocation and deallocation of physical and virtual storage resources for an ensemble.

- **Virtual server management**

Provides lifecycle management to create, delete, activate, deactivate, and modify definitions of virtual servers.

- **Virtual network management**

Provides for the lifecycle management of networking resources for an ensemble.

- **Performance management**

Provides a global performance view of all the virtual servers supporting workload resource groups deployed in an ensemble. The virtual server workload resource group performance goal is like a simplified z/OS WLM policy.

- You can define, monitor, report, and manage the performance of virtual servers based on workload resource group performance policies.
- Policies are associated to the workload resource group.
 - From the overall workload resource group “performance health” report, you can review contributions of individual virtual servers.
 - You can manage resources across virtual servers within a hypervisor instance.
 - Provides load balancing recommendations
- Availability management for heterogeneous resources:
 - Provides availability management for workload resource groups in an ensemble.
 - Provides consistent monitoring, analysis, and reporting of hard and soft failures of virtual servers and element groups, analyzed in the context of the workload resource group.
 - Highlights out of compliance conditions, based on the redundancy objectives that are set in the availability policy for a workload resource group.
 - Integrates with other key Unified Resource Manager disciplines such as performance management.
 - Audit and log capabilities

- **Energy management, monitoring, and controls**
 - Monitors energy usage and control power-saving settings, accessed through the new monitors dashboard
 - Monitoring of virtual server resources for CPU use and delays, with capability of creating a graphical trend report
- **zEnterprise System Web Services API**
 - Builds on existing SNMP/CIM function plus new Unified Resource Manager capabilities
 - Supports modern scripting languages (e.g., Perl, Python) that have HTTP supporting libraries
 - Enables programmatic access to the same functions used by the HMC UI
 - Enables management of a zEnterprise from external tools

Unified Resource Manager supports different levels of resource management. A feature determines the hardware management functions and operational controls that are available for a zEnterprise mainframe and any attached zBX. These named features are "Manage", "Advanced Manage", and "Automate".




Figure 17. IBM zEnterprise Unified Resource Manager zBX Model 002 (HMC version 2.11.x) "Manage", "Advanced Manage", and "Automate" features




Figure 18. IBM zEnterprise Unified Resource Manager zBX Model 003 (HMC version 2.12.x) "Manage" and "Automate" features

In HMC versions 2.11.x, the "Advanced Manage" suite is similar to the "Automate" suite, providing Unified Resource Manager functions for virtual servers on System x blades. The "Advanced Manage" suite provides the same workload and performance management functions except for dynamic processor resource adjustments. In HMC 2.12.1 or higher, the "Advanced Manage" and "Automate" suites merge together as "Automate".

Resource management functions provided through the "Manage", the "Automate", and the "Advanced Manage" suites do not apply for:

- IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise

The following tables provide a short explanation of the management functions, what they provide, and highlights whether they are available in the Manage (square ■), Advanced Manage (circle ●) or Automate (triangle ▲) suite for z/VM, POWER blades, System x blades, and the IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise.

Note: The zEnterprise system is required to be in an ensemble (which requires feature code 0025) in order to enable these functions.

For the latest comparison of the functions and what they provide, see http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=PS&infotype=SA&appname=STGE_ZS_ZS_USEN&htmlfid=ZSL03113USEN&attachment=ZSL03113USEN.PDF

Table 4. Operational controls ¹

Management function	Description	z/VM V6.2	POWER blades	System x blades	DataPower XI50z
Change management	View firmware information	n/a	■	■	■
	Retrieve firmware changes	n/a	■	■	■
	Change firmware levels	n/a	■	■	■
	Back up/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	■	■ ²	n/a
Operations management	Allows power on and power off through the 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	■	■	■	■
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	■	■	■	■

¹ Available on zEnterprise servers. For central processor complex (zCPC) does not require ensemble (FC 0025).² Available on zEC12 or zBC12/IBM zEnterprise BladeCenter Extension (zBX) model 003, not available on z196 or z114 with zBX model 002.

Table 5. Hypervisor management

Management function	Description	z/VM V6.2	POWER blades	System x blades	DataPower XI50z
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 the 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, and network metrics	■	■	■	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

Table 6. Virtual server lifecycle management. Enabling directed and dynamic virtual server provisioning spanning all hypervisors from a single, uniform point of control

Management function	Description	z/VM V6.2	POWER blades	System x blades	DataPower XI50z
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

Table 7. Workload awareness and platform performance management

Management function	Description	z/VM V6.2	POWER blades	System x blades ¹	DataPower XI50z
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

Table 7. Workload awareness and platform performance management (continued)

Management function	Description	z/VM V6.2	POWER blades	System x blades ¹	DataPower XI50z
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.	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
Note:					
Reporting does not replace tools that report on operating system resources and performance.	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

¹ System x blades on z114 and z196 will only have Advanced Management functions. Early zEC12 servers will not have full Automate capabilities until they are upgraded to Version 2.12.1. 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.

Table 8. Workload awareness and availability monitoring and reporting

Management function	Description	zCPC	z/VM V6.2	POWER blades	System x blades ¹	DataPower XI50z
	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 whether all workload elements are operating and available. Element groups can be configured to manage virtual servers according to 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

¹ System x blades on z114 and z196 will only have Advanced Management functions. Early zEC12 servers will not have full Automate capabilities until they are upgraded to Version 2.12.1. 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.

Table 9. Network management

Management function	Description	zCPC	z/VM V6.2	POWER blades	System x blades	DataPower XI50z
	Create virtual network: virtual networks allow virtual servers to be connected together. The virtual network has a name, an associated VLAN identifier, and a list of servers that are authorized to connect to it. Virtual servers are connected to a virtual network for which they are authorized via a virtual network interface card (vNIC) and are assigned a virtual media access control (vMAC) address on the network.	■	■	■	■	■
	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

Table 10. Energy management. Energy monitoring and management can help you better understand the power and cooling demands of the zEnterprise system.

Management function	Description	zCPC	POWER blades	System x blades	DataPower XI50z
Monitoring of energy consumption and key environmental parameters		■	■	■	■
Power saving mode ¹ : Power saving mode 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 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.		▲	▲	●▲	▲

¹ For zCPC only, does not require ensemble (FC#0025).

z/VM V6.2 Single System Image (SSI) recommendation within a zEnterprise ensemble

If z/VM is configured to be ensemble managed, the z/VM image should not be part of an SSI cluster or CSE complex.

Hypervisor management

In a typical layered architecture, the layer that provides platform virtualization is called the hypervisor. Virtualization is the process of simulating the underlying physical hardware as isolated logical entities. The hypervisor provides to each virtual entity the illusion of full access to the physical hardware of the system. Multiple operating systems can use and share the logical entities. For each instance of a guest operating system managed by a hypervisor, the system resources are virtualized to appear as dedicated to each guest.

zEnterprise hypervisor management firmware provides tasks for managing deployment, hypervisor lifecycle, managing storage resources, performing RAS/FFDC and monitoring the supported hypervisors.

Deployment of the hypervisor components is integrated into the existing Support Element and HMC firmware, except for z/VM. The Support Element and HMC user interface firmware management tasks manage all updates and repairs to the hypervisors. The methods of handling updating the firmware for the PR/SM hypervisor and z/VM hypervisor remains intact and unchanged.

Lifecycle management includes the following tasks:

- Activate virtual server
- Deactivate virtual server
- Create virtual servers definitions
- Duplicate virtual server definitions
- Manage hypervisor installation and updates
- Configure and manage the IEDN

- Entitle and non-entitle the blades

Resource management

Configure and connect virtual switches, virtual networks, and virtual storage resources to virtual servers.

RAS/FFDC includes the following:

- If either the primary Support Element or primary HMC is temporarily unavailable, the system continues to operate. If the primary Support Element or primary HMC fails, the failing primary Support Element or primary HMC will automatically switch or can be manually switched to its alternate Support Element or alternate HMC, respectively.
- Hypervisor management resynchronizes the Support Element hypervisor objects and their shadow objects on the HMC when the primary and alternate HMCs switch roles.
- Repairs and replaces the hypervisors for hardware failures.
- Provides access to text consoles and graphical consoles to access guest console for each virtual server.
- Creates and captures FFDC data system logs. Support for call-home.

Monitoring includes system activities, activity profiles, workload reports, thermal data, and power consumption data for the ensemble resources. Monitored data is displayed using the Monitors Dashboard, Workloads Report, and Virtual Servers Report tasks on the HMC. The data can be used by the platform performance management component (PPM).

The hypervisor management requires access to persistent storage for definitions, states, and monitoring data for the managed objects. The firmware uses the existing Support Element and HMC save, restore, and mirroring mechanism to ensure that critical data is available either on media or on the alternate Support Element and alternate HMC. The exception is z/VM, as it is the authoritative source of the z/VM configuration. z/VM configuration information might be stored on the Support Element to account for communication disturbances to the z/VM SMAPI.

Hypervisor management tasks

Basic hypervisor management tasks are provided by Support Element and HMC user interface tasks. See the *zEnterprise System Hardware Management Console Operations Guide for Ensembles* (Unified Resource Manager version 2.12.0 or prior) for details about these tasks. For version 2.12.1, go to console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmca/v2r12m1/index.jsp>.

zBX Blade Details

Includes a brief description of the hypervisor, a virtual server shutdown option, and a hypervisor autorestart option for the targeted zBX Blade.

Image Details

Includes a brief description of the hypervisor, a virtual server shutdown option, and a hypervisor autorestart option for the targeted system image.

Activate

Starts the hypervisor.

Deactivate

Stops the hypervisor.

Manage Storage Resources

Provides a common interface for allocation and deallocation of physical and virtual storage resources for an ensemble, hypervisors, images, and virtual servers objects.

New Virtual Server

Guides you through the process of configuring a virtual server for the selected hypervisor.

Set Power Cap

Limits peak power consumption of your system resources.

Set Power Saving

Reduces the average energy consumption of your system resources.

Ensemble membership management

With the appropriate code level installed, the HMC provides tasks for creating an ensemble and controlling membership of the ensemble. The HMC that creates the ensemble is enabled to perform subsequent platform management functions on that ensemble. In addition, the HMC can still perform all the non-ensemble HMC functions on systems, regardless of whether the HMC is managing an ensemble.

The HMC managing the ensemble owns the platform management configuration and policy that spans all of the managed members in the ensemble. It also has an active role in system monitoring and adjustment. With these new responsibilities, the primary HMC is required to have an alternate HMC for redundancy. The alternate HMC has a limited set of tasks and is designated solely as a backup HMC in case the primary HMC fails.

After the ensemble is created, the ensemble-managed objects are available to the HMC user interface. The ensemble-specific managed objects are:

- Ensemble
- Ensemble Members
- Hypervisors
- Storage Resources
- Virtual Networks
- Virtual Servers
- Workloads
- zBX Blades
- zBX BladeCenters

Ensemble membership management tasks

Ensemble membership management tasks are described in detail in the *zEnterprise System Hardware Management Console Operations Guide for Ensembles* (Unified Resource Manager version 2.12.0 or prior) for details about these tasks. For version 2.12.1, go to console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmca/v2r12m1/index.jsp>.

Add Member to Ensemble

Adds the selected CPC to the ensemble to be managed as part of a single logical virtualized system. When the task is finished, the CPC is a member of the ensemble.

Alternate Details

Provides the ensemble name, machine type, machine model, serial number, and status on the alternate HMC.

CPC Details

Provides the following additional information for ensemble management:

Instance Information - Displays the name of the ensemble if an ensemble is managing the CPC.

zBX Information - Displays the characteristics and settings of the attached zBX, when installed.

Energy Management Information - Displays the power and thermal monitoring information, when the appropriate feature is installed.

Product Information - Displays the Unified Resource Manager suite level in the Management Enablement Level field.

Create Ensemble

Guides you through creating an ensemble and optionally adding members to the ensemble.

Customize User Controls

Defines and customizes user roles. A *user role* is a collection of authorizations. You can create a user role to define the set of tasks permitted for a given class of user (task roles). You can create a user role to define the set of managed objects that are manageable for a user (managed resource roles). Use this task to view and manage task roles and resource roles introduced for ensemble management. When you have defined or customized the user roles you can use the **User Profiles** task to create new users with their own permissions.

Delete Ensemble

Deletes the selected ensemble from the primary HMC. After the ensemble is deleted, the primary HMC title designation is removed and the alternate HMC is an unassigned alternate HMC.

Ensemble Details

Provides the name, description, management suite, status of the HMC, performance management and availability settings, energy monitoring information, network MAC address information, alternate HMC assigned, and option to lock out disruptive tasks for the ensemble.

Ensemble Management Guide

Provides an educational task that walks you through an example of creating an ensemble. The guide provides for adding members and using the ensemble for managing a business application. Use this guide to assist you with setting up an ensemble.

Hardware Messages

Displays consolidated hardware-related messages for all selected hardware in the processor cluster, including your Hardware Management Console. See the hardware messages section of the *Hardware Management Console Operations Guide* (Unified Resource Manager version 2.12.0 or prior) for basic information about this task. For version 2.12.1, go to console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmc/v2r12m1/index.jsp>.

Manage Alternate HMC

Are functions that vary depending on the Hardware Management Console (HMC) to which you are logged on.

From an HMC that is not the primary HMC or the alternate HMC, this task performs the following functions:

- Enables or disables the HMC pairing capability.

From the primary HMC, this task performs the following functions:

- Switches the roles of the primary HMC and the alternate HMC
- Queries the switchover capabilities
- Mirrors data from the primary HMC to the alternate HMC
- Resets the fenced alternate HMC due to an automatic switchover
- Enables or disables the automatic switchover capability
- Enables or disables service mode
- Pre-loads the alternate HMC with a new level of code.

From the alternate HMC, this task performs the following functions:

- Switches the roles of the primary HMC and the alternate HMC

- Queries the switchover capabilities.

Operating System Messages

Displays the consolidated operating system-generated messages for all CPC members of the selected ensemble.

Remove Member from Ensemble

Removes the selected CPC from the ensemble. The removed CPC is no longer managed as part of the ensemble.

Remove Object Definition

Removes a system (CPC) from the **Systems** group in the navigation pane for the tree style user interface or the **Defined CPCs** group in the classic style user interface.

User Profiles

Manages system users that log on to the Hardware Management or Support Element Console.

The predefined task roles for ensemble management include:

Version 2.12.1	Version 2.12.0 and prior
• Ensemble Administrator Tasks • Virtual Network Administrator Tasks • Virtual Server Administrator Tasks • Virtual Server Operator Tasks • Storage Resource Administrator Tasks • Workload Administrator Tasks • Policy Administrator Tasks • Policy Operator Tasks • Energy Management Administrator Tasks	• Ensemble Administrator Tasks • Virtual Network Administrator Tasks • Virtual Server Administrator Tasks • Virtual Server Operator Tasks • Storage Resource Administrator Tasks • Workload Administrator Tasks • Performance Management Administrator Tasks • Performance Management Operator Tasks • Energy Management Administrator Tasks

Storage management

Storage management provides a common interface for allocation and deallocation of physical and virtual storage resources for the virtual servers in an ensemble. The **Manage Storage Resources** task provides two different views of the storage resources. You can view storage resource by: Storage Resources or Virtual Disks.

The Manage Storage Resources task simplifies storage management by presenting a single point of control for all the storage resources associated with hypervisors in the ensemble. The only exceptions are virtual servers running in native LPARs. Your storage administrators can add capacity and distribute storage across different workload resource groups or applications through this task. The main features are:

- Is integrated into the system hardware and firmware
- Supports sharing of storage resources across the enterprise
- Supports a virtualized pool of storage from the storage systems attached to a SAN

The following storage resource views are available from this task:

Storage Resource

Lists each defined storage resource with its hypervisor, providing the owner status, resource type, storage size, group name, and description for each resource.

Virtual Disks

Lists each virtual disk with its virtual server, providing the disk size and description, hypervisor, storage resource (SR) name, SR size and description, and shared status for each resource.

The following storage resource actions are available from this task:

- **Add Storage Resource**

Adds new storage resources to the selected hypervisor from a list of available storage resources.

- **Add Storage Resource to Group**

Adds a storage resource to a z/VM Storage Group. Virtual devices are then provided in the form of mini disks.

- **Compare Access List**

Compares the storage connectivity of two hypervisors.

- **Discover Storage Resources**

Detects the storage resources available to one or more hypervisors.

- **Export World Wide Port Name (WWPN) List**

Exports the aggregated Fibre Channel host ports of each hypervisor in the form of a text file.

- **Import Storage Access List**

Imports a storage access list. Storage access lists have two purposes:

- Authorizes the ensemble to manage the storage resources contained in the list.
- Indicates which storage resources can be accessed by each hypervisor in the ensemble.

You can add storage resources at any time. The storage access list defines what storage resources can be accessed by each hypervisor. Multiple hypervisors can access the same storage resource. The list consists of a number of host worldwide port names (WWPNs) and entries for configured storage resource. The list includes properties, such as addressing information, or device type information.

- **Remove Storage Resource**

Removes a storage resource from the list of accessible storage resources of a hypervisor within the ensemble.

- **Remove Storage Resource from Group**

Removes a storage resource from a z/VM Storage Group.

- **Retest Communication with Storage Resources**

Tries to communicate, in real time, with selected resources.

- **Test Communication with Storage Resources**

Tries to communicate, in real time, with selected resources.

With appropriate user authority, attribute and resource settings (resource details) for a storage group, hypervisor, storage resource, or virtual server can be viewed or modified if you click the links provided in the table columns.

Storage device types include:

- Virtual DVD drive
- HMC DVD or USB drive
- Devices for PowerVM and INxHY are backed by FC SAN devices
- z/VM mini disks backed by either FC SAN devices or ECKD™ devices

FC SAN or FICON attached ECKD devices can be used as storage devices. You are responsible for managing your SAN and for providing FC targets in the right granularity. You can use Multi-pathing. A virtual server can boot from a storage device. Virtual disks are based upon regular or full-pack minidisk ECKD.

You can upload and delete image files through the hypervisor. Uploaded images are temporary data and do not persist anywhere except on the local storage device of the blade.

All hypervisors can use storage resource type of FC logical units (LUNs). Type ECKD volumes are for z/VM only. For FC on z/VM, only full logical units are assigned to virtual servers.

Virtual server management

At a high level, each virtual server requires a guest operating system, an image to boot, a configuration of IP addresses, a quantity of memory to use, storage access, and network access. The storage and network device access is commonly mapped into the physical storage and network infrastructure of the system.

Virtual Server Management tasks provide consistent basic functions for virtual servers hosted by z/VM and IBM blade hypervisors:

- Lifecycle management
- Resource management
- Consistent HMC interface.

Virtual Server Management Capabilities include:

- Lists the virtual servers
- Groups objects (security)
- Supports creating and deleting virtual servers
- Supports activating and deactivating virtual servers
- Provides task to assign a virtual storage device to the guest
- Provides task to assign a virtual network adapter to the guest
- Optionally adds virtual server to a workload resource group
- Installation of operating system on a virtual server
- Provides an operating system security model (access rights) for operating system consoles running on HMC or Support Element
- Provides ability to view and modify the details of the virtual server definition: name, description, processor capacity, memory utilization, attached network devices, attached storage devices, and workload resource group policies.

Virtual server management tasks

See the *zEnterprise System Hardware Management Console Operations Guide for Ensembles* (Unified Resource Manager version 2.12.0 or prior) for details about these tasks. For version 2.12.1, go to console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmca/v2r12m1/index.jsp>.

Activate

Starts the virtual server.

Choose z/VM Virtual Servers to Manage

Assigns z/VM virtual machines to be managed as virtual servers in the ensemble.

Customize Scheduled Operations

Schedules the times and dates for certain operations. The scheduled operations for an ensemble are: **Activate Availability Policy** (Unified Resource Manager version 2.12.1 or later), **Activate Performance Policy**, **Set Power Saving**, **Activate** virtual server, and **Deactivate** virtual server.

Deactivate

Stops one or more virtual servers gracefully.

Delete Virtual Server

Removes the selected virtual servers from the hosting hypervisor.

Initiate Hypervisor Dump

Sends a command to begin a memory dump of the hypervisor running on the target IBM blade.

Initiate Virtual Server Dump

Sends an operating system-specific command to begin a memory dump of the operating system running on the virtual server.

Initiate z/VM Management Guest Dump

Sends command to begin a memory dump of the target z/VM management guest to facilitate problem determination.

Manage Virtual Switches

Creates virtual switches for z/VM virtual servers to access IEDN HiperSockets (IQDX), IEDN network, and external OSD networks.

Migrate Virtual Server

Migrates the source virtual server definition to another hypervisor of the same type.

Monitors Dashboard

Provides a summary of system activities, activity profiles, workload reports for the ensemble. The summaries include information about processor usage, zBX blades usage, channel usage, memory usage, network I/O usage, storage usage, power consumption, and thermal measurements.

Mount Virtual Media

Mounts an ISO image to the virtual server as a virtual DVD and provides the option to boot the server with the virtual DVD.

New Virtual Server

Guides you through the process of adding a new virtual server definition to the selected hypervisor.

New Virtual Server Based On

Guides you through the process of creating virtual server definitions based on an existing virtual server definition.

Open Graphical Console

Opens a graphical console to the VGA output on the selected System x blade virtual server to install or administer the operating system.

Open Text Console

Opens a text-based console with the operating system on the selected IBM blade virtual server.

Restart z/VM Management Guest

Issues a restart request to the target z/VM management guest machine.

Virtual Server Details

Provides a summary list and status of the virtual server definition.

Virtual network management

Virtual network management provides tasks for the provisioning and managing of networking resources in an ensemble.

The network virtualization function has two key roles in ensemble function and management:

- The network virtualization function manages all connectivity for virtual servers and optimizers across the intraensemble data network (IEDN). Each virtual server and optimizer must belong to a virtual network to connect to the IEDN. You create virtual networks through the network virtualization function.
A virtual server or virtual NIC must have a VLAN ID for IEDN connection and communication. The network virtualization function controls this access. A VLAN ID is obtained only by association with a virtual network defined through the HMC user interface.
- The network virtualization function manages the MAC registry for the ensemble. You can use it to assign VMAC prefixes or full MAC addresses to hypervisors and virtual servers across the ensemble in

order to prevent VMAC collisions. These operations are performed automatically over the intranode management network (INMN). From the HMC, an administrator can reserve MAC prefixes for special situations.

Network management tasks support static migration of INxHY and PowerVM. For example, the virtual server must be stopped for migration to occur.

Virtual network management tasks

See the *zEnterprise System Hardware Management Console Operations Guide for Ensembles* (Unified Resource Manager version 2.12.0 or prior) for details about these tasks. For version 2.12.1, go to console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmc/v2r12m1/index.jsp>.

The following network tasks are available:

Add Hosts to Virtual Network

Adds hosts to the selected virtual network. Hosts include virtual servers and virtual NICs.

Configure Top-of-rack (TOR) Switch

Provides connectivity to the intraensemble data network (IEDN), supporting application data within an ensemble. Only the top-of-rack (TOR) switches connected to the IEDN in an ensemble are available for this task. Configuration actions cover switch ports, VLAN settings, and MAC address filtering.

Customize Network Settings

Opens the WebSphere DataPower XI50 user interface for configuring the selected DPXI50z zBX blade network settings.

Delete Virtual Network

Removes a virtual network from the ensemble.

Manage Ensemble MAC Prefixes

Configures Ensemble MAC Prefix

This function customizes the ensemble MAC prefix for cases where existing DHCP servers might not accept a MAC address that is administered locally. If a DHCP server serves multiple ensembles, you can use this task to ensure that the DHCP server does not encounter MAC collisions.

Reserves MAC Address Prefix

This function reserves a MAC prefix for rare situations when resources outside the ensemble are connected to a TOR switch within the ensemble. In some situations, having a reserved MAC prefix available avoids potential MAC collisions within the ensemble.

Manage Virtual Networks

Provides tasks for creating or deleting a network, adding hosts, removing hosts, viewing details, or editing details of a virtual network.

Network Monitors Dashboard

Provides tasks for monitoring network metrics and displaying statistics for the networking resources associated with the IEDN. The dashboard consists of three views, they are: Virtual Networks, Physical Interfaces, and Physical Switches.

New Virtual Network

Creates a new virtual network for use in the ensemble.

Remove Hosts from Virtual Network

Removes hosts from the selected virtual network.

Repair Virtual Network

Verifies that all hosts are configured as wanted and detects problems that might exist.

Virtual Network Details

Provides the name, description, VLAN ID, status, and a list of connected hosts for the network.

Performance management

The objective of performance management is to provide an understanding of whether performance goals are being achieved. To help isolate the source of the performance problems, the data is monitored at the virtual server level.

Platform management

The platform management has a global performance view of all the virtual servers supporting workload resource groups deployed in a zEnterprise ensemble. This type of knowledge is useful to optimize network workload resource group balancing decisions.

The platform management firmware component is responsible for goal-oriented resource monitoring, defining and implementing goal-oriented policies, and reporting across the ensemble. The platform management tasks are integrated into the HMC user interface and packaged as zEnterprise firmware. The performance management goals and policy administration goals run on a trusted internal platform-managed network.

A number of new concepts are introduced with platform management functions.

- Compatible networked systems, virtualization, and management functions are integrated as a single system. Regardless of the ensemble size, the complexity and costs are managed as if it was a single system.
- The platform workload resource group provides the context for managing platform resources based on the requirements of a deployed business application. The virtual servers that host a business application support the workload resource group. The workload resource group provides a grouping mechanism and a “management view” of these virtual servers. A given virtual server might be exclusively associated with a single workload resource group or shared across multiple workload resource groups. Associated with a workload resource group is a set of policies that specify goals controlling how platform resources are applied to the virtual servers. You define these goals in the context of the requirements of the business application that the workload resource group represents. Policies that define goals for different management disciplines are planned. Examples include performance, energy, and availability.
- Performance policies specify performance goals for the work running in the workload resource group. Performance policies control how platform resources are dynamically provided to the virtual servers making up the workload resource group. The performance policy is organized into service classes that assign performance goals.
- A workload resource group performance policy might contain one or more service classes. A service class is a group of work for which you have similar business goals or business requirements. Each service class contains a business performance goal and a classification rule.
- Each service class has one classification rule that identifies how one or more fragments of a workload resource group are associated with service class. For example, if you have a batch process for a “Payroll” workload resource group and you want to manage it to a specific business goal, you must define a service class (Payroll Batch) and specify a classification rule that identifies the servers on which the batch process runs.

Workload resource group and performance management tasks

See the *zEnterprise System Hardware Management Console Operations Guide for Ensembles* (Unified Resource Manager version 2.12.0 or prior) for details about these tasks. For version 2.12.1, go to console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmca/v2r12m1/index.jsp>.

Activate Policies

Activates a performance or availability policy for the target workload resource group. (Unified Resource Manager version 2.12.1)

Activate Performance Policy

Activates a performance policy for the target workload resource group. (Unified Resource Manager version 2.12.0 and prior)

Add Custom Groups

Adds one or more custom groups of virtual servers to the table in the **Virtual Servers** tab of the **Workload Resource Group Details** task.

Add Virtual Servers

Adds one or more virtual servers to the table in the **Virtual Servers** tab of the **Workload Resource Group Details** task.

Customize Scheduled Operations

Schedules the times and dates for certain operations. The scheduled operations for an ensemble are: **Activate Availability Policy** (Unified Resource Manager version 2.12.1 or later), **Activate Performance Policy**, **Set Power Saving**, **Activate** virtual server, and **Deactivate** virtual server.

Delete Workload Resource Group

Removes a target workload resource group and all the performance policies associated with it.

Element Group Details

Use to view or edit the properties of a target element group.

Ensemble Details

Provides the name, description, management suite, status of the HMC, performance and availability management settings, energy monitoring information, network MAC address information, alternate HMC assigned, and option to lock out disruptive tasks for the ensemble.

Hops Report

Illustrates information for each hop involved with handling a specific service class within a specific workload resource group during the requested interval time frame.

Hypervisor Report

Provides a report on virtual server consumption, successful adjustments, and failed adjustments of shared resources in the same hypervisor.

Load Balancing Report

Provides a report on what load balancers are currently configured to get load balancing recommendations from the platform performance manager. Also, provides details about the recommended weights for each load balance group.

Manage Workload Element Groups

Use to manage element groups that are defined in the ensemble or a specific workload resource group.

Monitor System Events

Provides for the creation and management of event monitors. Notifies users by email if an event passes through all the filters of the event monitor.

Monitors Dashboard

Provides a summary of system activities, activity profiles, workload reports for the ensemble. The summaries include information about processor usage, zBX blades usage, channel usage, memory usage, network I/O usage, storage usage, power consumption, and thermal measurements.

New Element Group

Use the New Element Group wizard to guide you through the process of creating a new element group. An element group is a group of virtual servers which are functional peers organized for redundancy on behalf of one or more workload resource groups.

New Performance Policy

Guides you through the process of creating a performance policy and any service classes associated with it.

New Service Class

Guides you through the process of creating a service class or an existing service class to create a service class. Service class is a group of work that has the same service goals or performance objectives, resource requirements, or availability requirements.

New Workload Resource Group

Guides you through the process of creating a workload resource group. This wizard also provides options for you to select virtual servers, create performance policies, and schedule activation of policies, for the new workload resource group.

Performance Policy Details

Provides the name, description, business importance, workload resource group status, and service classes for the performance policy.

Performance Policy Print/View

Provides a printable version of the properties of a target performance policy.

Performance Policy Revisions

Provides a view of the performance policy revision log file. Every time a policy is changed, its previous revision is archived. This function displays all revisions up to the current policy.

Resource Adjustment Report

Lists successful and failed resource adjustments for a workload resource group, virtual server, or service class over the selected reporting interval. This function presents information in two tables, one for the successful adjustments and one for failed adjustments.

Service Class Details

Provides the name, description, workload resource group, performance policy, business importance, performance goal, and classification rule for the service class.

Service Class Print/View

Provides a printable version of the properties of a target service class.

Service Classes Report

Provides a list of all the service classes defined for the workload resource group for a specified time interval. A line graph of performance index versus time is available for the service classes.

View Statistics

Provides a list of operating system processes and application environment statistics for the selected Virtual Server Topology Report node.

Virtual Server Topology Report

Provides a view of the relationships between virtual servers that are running the workload resource group and providing the resources to complete the work.

Virtual Server Report

Provides data for each virtual server associated with the workload resource group or a service class. The data includes allocated resources, resource utilization data, and delay statistics.

Workload Resource Group Details

Provides the name, description, category, virtual servers, and performance policies for the workload resource group.

Workload Resource Group Print/View

Provides a printable version of the properties of a target workload resource group.

Workload Report

Lists all associated workload resource groups and their active performance policies and health for a specified interval. This report provides high-level performance status and goal achievement

information for all the workload resource groups associated with active performance policies. Two workload charts are available with the **Workload Report** task: **CPU Utilization** and the **Performance Index**.

Load Balancing Report

Lists the load balancing groups to which external routers are distributing incoming work requests. Through communication with the Unified Resource Manager, these external routers receive weight recommendations that enable the routers to distribute incoming work more effectively among virtual servers in the ensemble.

Availability management

The primary purpose of IBM zEnterprise Unified Resource Manager availability monitoring and reporting functions is to provide system administrators with the data they need to understand whether availability goals are being met. If these goals are not being achieved, detailed data helps identify the source of availability problems.

With the default availability policy in effect, the Unified Resource Manager provides only availability status for individual virtual servers within a workload resource group. Through custom availability policies, you can benefit from more advanced availability management functions:

- Using element groups to eliminate single points of failure. An element group is a set of virtual servers that are functional peers, organized for redundancy on behalf of one or more workload resource groups.
- Assigning different availability goals to different workload resource groups.

Availability policies are integrated with other key Unified Resource Manager disciplines such as performance management, energy management. Availability policies specify availability and redundancy requirements for the virtual servers and element groups performing the work in the workload resource group.

Although the use of guest platform management providers is not a requirement for availability management functions, IBM recommends that you install one on each virtual server associated with an availability policy. Doing so can result in more accurate and timely availability status reports.

Availability management tasks

For information about availability management tasks, go to console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmca/v2r12m1/index.jsp>.

The following availability tasks are available:

Activate Policies - Activate availability policy

Use to activate an availability policy for the target workload resource group.

Availability Status Event Report

Lists all availability status changes for the workload resource group, its element groups, and its virtual servers over the selected reporting interval. The report is presented in table format.

Workload Resource Group Details: Availability Policy Details

Use to view or edit the properties of a given availability policy. (The default availability policy cannot be edited).

Workload Resource Group Details: Availability Policy Print/View

Use to view a printable version of the properties of a target availability policy. If you are connected remotely to the console, you can also print the information.

Workload Resource Group Details: Availability Policy Revisions

Displays the contents of the revision log file of the availability policy. Every time a policy is changed, its previous revision is archived. These revisions are all displayed, up to the current policy.

New Availability Policy

Creates a new availability policy. Defines availability goals for the elements in a workload resource group. This action cannot be performed on a default workload resource group target.

Manage Workload Element Groups: New Element Group and Element Group Details

Use to provide availability by eliminating single points of failure or for identifying functional capacity.

Energy management, monitoring, and controls

zEnterprise provides energy management functions to monitor, manage, and customize power allocations within the physical limits of your ensemble.

You can use monitoring tasks and related resource **Details** tasks to monitor power and thermal properties of your ensemble, CPCs, IBM blades, optimizers, and BladeCenters. Power capping and power saving are two configuration tasks introduced in zEnterprise. Use the **Set Power Cap** task to limit peak power consumption and **Set Power Saving** task to reduce average energy consumption of your system resources.

All zEnterprise new monitoring is done on the new re-engineered SAD display. You can use the SNMP API to query or alter these controls for the IBM Systems Directory plug-in, Active Energy Manager (AEM), to access the power and thermal data. The SNMP API provides for these controls to be queried and altered. In addition, the **Activation Reset Profile Option** task is updated so you can set up **Static Power Savings** mode settings at activation time.

Energy management includes:

- Monitors energy and thermal data using the HMC user interface
- Monitors energy and thermal data for zEnterprise using AEM
- Triggers **Power Savings Mode** for zEnterprise CPCs using the HMC or AEM
- Triggers **Power Savings Mode** for IBM blades using the HMC or AEM
- Provides **Set Power Cap** for BladeCenter (Group Capping) using the HMC or AEM
- Provides **Set Power Cap** for individual zBX blades using the HMC or AEM
- Provides **Set Power Cap** (query maximum potential power) for zEnterprise CPCs using the HMC or AEM

Energy management tasks

See the *zEnterprise System Hardware Management Console Operations Guide for Ensembles* (Unified Resource Manager version 2.12.0 or prior) for details about these tasks. For version 2.12.1, go to console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmca/v2r12m1/index.jsp>.

Activity

Indicates the system activity for CPCs or a group of CPCs. System activity includes channel and physical processing activity defined in the system profiles that are stored in selected CPCs.

CPC Details

Supports the following additional information for ensemble management:

Instance Information

Displays the name of the ensemble if an ensemble is managing the CPC.

zBX Information

Displays the characteristics and settings of the attached zBX, when installed.

Energy Management Information

Displays the power and thermal monitoring information, when the appropriate feature is installed.

Customize Scheduled Operations

Schedules the times and dates for certain operations. The scheduled operations for an ensemble are: **Activate Policies** (Unified Resource Manager version 2.12.1), **Activate Performance Policy** (Unified Resource Manager version 2.12.0 and prior), **Set Power Saving**, **Activate** virtual server, and **Deactivate** virtual server.

Customize/Delete Activation Profiles

Creates activation profiles, customizes existing profiles, or deletes unwanted profiles that are stored in the Support Element. The parameter for an ensemble that is in the activation profile is: **Set Power Saving**.

Ensemble Details

Provides the name, description, management suite, status of the HMC, performance and availability management settings, energy monitoring information, network MAC address information, alternate HMC assigned, and option to lock out disruptive tasks for the ensemble.

Monitors Dashboard

Provides a summary of system activities, activity profiles, workload reports for the ensemble. The summaries include information about processor usage, zBX blades usage, channel usage, memory usage, network I/O usage, storage usage, power consumption, and thermal measurements.

Set Power Cap

Limits the peak power consumption of a system resource or group of resources. You can closely manage power allocations within the physical limits of your data center.

Set Power Saving

Reduces the average energy consumption of a system component or group of components. You can closely manage power allocations within the physical limits of your data center.

zBX Blade Details

Provides the current instance information, hypervisor information, and acceptable status settings for the selected zBX Blade:

- Instance Information includes the status of the zBX Blade and information about operating conditions, characteristics, and settings of the zBX Blade.
- Product Information, is assigned to zBX Blades when they are manufactured and is primarily for identification.
- Acceptable Status settings determine which zBX Blade statuses are acceptable and which statuses are unacceptable. The Support Element console reports when the zBX Blade status becomes unacceptable.
- Energy Management Information specifies power and thermal monitoring information.
- Hypervisor Information, includes a brief description of the hypervisor, a virtual server shutdown option, and a hypervisor autorestart option for the targeted zBX Blade.
- IEDN Interfaces, displays the intraensemble data network (IEDN) configuration settings for the targeted IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z) blade. You can add an interface, view the details of an existing interface, or delete an interface. This tab is only available for DataPower XI50z blades.

zBX BladeCenter Details

Provides the following information for the selected BladeCenter:

- Acceptable Status settings determine which BladeCenter statuses are acceptable. The console reports when the BladeCenter status becomes unacceptable

- Instance Information includes the status of the BladeCenter and information about operating conditions, characteristics, and settings of the BladeCenter.
- Product Information is assigned to the BladeCenters when they are manufactured and is primarily for identification.
- Energy Management Information, specifies power and thermal monitoring information.

zEnterprise System Web Services API

The Web Services API is a web-oriented programming interface that makes the underlying Unified Resource Manager 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. The capabilities available using the APIs are consistent with their counterparts provided by the Hardware Management Console (HMC) user interface and might be used to enable discovery, monitoring, and provisioning use cases.




Figure 19. IBM zEnterprise Unified Resource Manager API

Some API characteristics are:

- TCP/IP Sockets/HTTP is the underlying network support with SSL for connection security.
- Governed by the functions they involve such as 'Manage' or 'Automate'.
- Architecture/design is based on current industry design practices. Requests and responses are structured as web services.
- Can be disabled for installations not using it.
- Connection to API requires authentication under an HMC application login identity.
- Individual requests by the API are authorized using the HMC's controls.
- Security and event log entries are created at same points as for UI invoked requests.

The APIs are fully documented and supported for customer and third-party use. See the *System z Hardware Management Console Web Services API*, SC27-2626 Version 2.12.1, SC27-2617 Version 2.12.0, and SC27-2616 Version 2.11.1, for details about API support.

Web Services API tasks

Basic API tasks are provided by HMC user interface. See the *zEnterprise System Hardware Management Console Operations Guide for Ensembles* (Unified Resource Manager version 2.12.0 or prior) for details about these tasks. For version 2.12.1, go to console online help or the **System z HMC and SE (Version 2.12.1) Information Center** at <http://pic.dhe.ibm.com/infocenter/hwmca/v2r12m1/index.jsp>.

Customize API Settings (The WEB Services tab)

Enables and customizes the Web Services Application Programming Interface (API).

User Profiles and Manage Users Wizard (Allow access to management interfaces)

Enables or disables a particular Hardware Management Console (HMC) user profile permission to use the Application Programming Interfaces (APIs).

Manage Web Services API Logs

Provides a view of the web services API logs or option to save the logs.

Chapter 4. Planning your system

To help plan your system, a sample implementation introduces how to move a business solution from a distributed environment to an ensemble.

The sample implementation shows a configuration for a distributed business solution and how that same solution is configured on an ensemble. The distributed business solution is an interconnected set of applications; some applications run in System z, and some run on System p®. A key value of an ensemble is that it brings all these applications together so that you can manage them as a single unit.

Topics and scenarios illustrating planning and configuration tasks use the sample implementation solution for the following publications:

- *zEnterprise System Ensemble Performance Management Guide*, GC27-2607 (Unified Resource Manager version 2.12.0 or prior), helps you define, create, and manage the performance goals of business application solutions deployed in a zEnterprise environment.
- *zEnterprise System Ensemble Workload Resource Group Management Guide*, GC27-2621 (Unified Resource Manager version 2.12.1), helps you define, create, and manage the performance and availability goals of business application solutions deployed in a zEnterprise environment.
- *zEnterprise System Ensemble Planning and Configuring Guide*, GC27-2608, leads you through the planning and implementation phases for moving business application solutions from a distributed environment to an integrated set of virtual servers running in a zEnterprise environment.

In addition, the *zEnterprise System Ensemble Planning and Configuring Guide*, GC27-2608 includes a few migration scenarios. The *Migration scenarios* and *Planning for upgrades and zBX moves* topics document the steps or considerations to move from existing distributed environments to an ensemble, upgrading zBX, moving zBX between systems, or moving DataPower XI50z blades between zBXs.

Scenarios include:

- Replacing one z10™ in a two CPC sysplex with z/VM
The migration action is to replace one z10 running part of the z/OS sysplex and z/VM with a z196 and enable that zCPC to be a member of new ensemble. The new ensemble does not have a zBX.
- Migrating from a single CPC sysplex and IBM POWER servers
The migration action is to replace a z10 with a z196, enable that z196 as a node in its own ensemble, and then migrate from external IBM POWER servers to POWER blades in the zBX.
- Planning for upgrades from z196 with zBX Model 002 to zEC12 with zBX Model 003
- Planning for upgrades of zBX Model 002 to zBX Model 003
- Planning for moving a zBX Model 002 from one z196 to another z196
- Planning for moving a zBX Model 003 from one zEC12 to another zEC12
- Planning for move of DataPower XI50z blades from one zBX to another zBX
- Planning for upgrades to the primary HMC and alternate HMC

The **Ensemble Management Guide** task located on the Hardware Management Console can help familiarize you with the various tasks to create and manage aspects of an ensemble.

The Insurance Company: A sample implementation

The Insurance Company recognizes that it must increase the efficiency of its business processes, while cutting costs, to remain competitive in the industry. Transitioning to an IBM zEnterprise System (zEnterprise) environment is part of the Insurance Company's plan to streamline business processes, increase speed to market, and provide better and faster services for customers and employees.

The Insurance Company recently adopted a fully automated insurance business solution to streamline business processes, using interactive electronic forms and signatures to expedite the insurance application process and other forms required for policy changes. Instead of requiring two to six visits and up to 60 days to process applications from new customers, insurance salespeople can now sell new policies in as little as one sales visit and process them in 45 minutes or less.

This business solution, called the Insurance Business Application, integrates the functions of multiple independent software vendors (ISVs) and IBM strategic software. The Insurance Business Application increases speed and accuracy, ensures compliant and secure processes, and applies real-time business rules for all processes. The Insurance Business Application currently is deployed in a mainframe-centric, heterogeneous infrastructure that consists of multiple platforms and operating system types.

In the future, the Insurance Company expects to launch new and improve existing business processes even more quickly and efficiently. To meet these business goals, the Insurance Company plans to transition the Insurance Business Application to a zEnterprise environment.

The Insurance Company plan for a zEnterprise ensemble

The Insurance Business Application is deployed in a heterogeneous environment, shown in Figure 20, that includes the following hardware:

- One IBM System z10 Enterprise Class (z10 EC)
- One IBM System p Power® 595
- One IBM BladeCenter containing BladeCenter HS20 blade servers
- One IBM System Storage® DS8300.



Figure 20. The Insurance Company hardware configuration before zEnterprise

The Insurance Company plans to change this hardware configuration by installing and configuring a one-node ensemble as shown in Figure 21 on page 59.

To determine which hardware to install, the Insurance Company information technology (IT) personnel analyzed their current hardware and the architecture of their insurance business application and projected the required capacity. They mapped the z10 EC resources one-to-one to the IBM zEnterprise 196 (z196) and, with the help of IBM representatives and capacity planning tools, determined the number and type of blades required for the IBM zEnterprise BladeCenter Extension (zBX).




Figure 21. The Insurance Company hardware configuration with zEnterprise components

To transition to this zEnterprise ensemble configuration, the Insurance Company plans to complete the following tasks:

- Install the z196 and move all business processes, services, and components already running on the z10 EC to the z196.
- Install one zBX, containing two chassis and the following blades:
 - 14 POWER blades
 - 7 System x (HX5) blades

The top-of-rack switch on the top left of the zBX in Figure 21 illustrates the connection to the z196 through two private networks:

 - The intraensemble data network (IEDN)
 - The intranode management network (INMN).
- Order the following suites to take advantage of the IBM zEnterprise Unified Resource Manager functions:
 - The **Automate** suite for select POWER blades
 - The **Advanced Management** suite for select HX5 blades
- Move policy-related applications and processes from the System p Power 595 and the BladeCenter HS20 blade servers to the zBX, to be monitored together in an ensemble workload resource group.
- Retain the DS8300, with 25 terabytes of virtualized disk, as the storage device for Insurance Company data.

The System p Power 595 remains in the IT shop to run applications that are not required to support insurance policy business processes. The DS8300 and the Power 595 are connected to the ensemble node through customer-managed data and management networks.

Figure 22 on page 60 provides a conceptual view of the z196 and its operating system configuration.




Figure 22. A closer look at the z196 for the Insurance Company IT shop

The z196 is a 2817 Model M80 with:

- Two terabytes of memory
- Twenty-eight logical partitions (LPAR)s managed by the Processor Resource/System Manager (PR/SM) hypervisor. The LPARs are used for the following purposes:
 - Six LPARs constitute a Parallel Sysplex® for production work, running z/OS V1R12 systems. The z/OS LPARs have DB2 9.1 and WebSphere Application Server installed.
 - Three LPARs form another Parallel Sysplex, also running z/OS V1R13, for testing purposes.
 - One LPAR runs z/VM 6.1 with Linux guests.

Figure 23 provides a conceptual view of the zBX configuration.




Figure 23. A closer look at the zBX for the Insurance Company IT shop

The zBX is a 2458 Model 002, with two BladeCenter chassis, containing the following blade types:

- 14 POWER blade PS701 blades. Each blade has a PowerVM Enterprise Edition hypervisor.
- 7 System x HX5 7873 blades. Each blade has a zBX integrated hypervisor for IBM System x blades (abbreviated as INxHY or x Hyp).

Some POWER and System x blades are reserved for production work, while others are reserved for testing and development purposes.

Note that Figure 23 does not illustrate the actual requirements of a physical configuration. For purposes of illustration in this chapter, the POWER and System x blades are grouped by type and the specific allotment of blades within a chassis is not shown.

With a plan and requirements in place for the hardware and operating systems, the Insurance Company IT personnel need to plan and prepare to move the Insurance Business Application to the new zEnterprise environment. This planning task requires a thorough understanding of the Insurance Company's business goals, and of the processes that comprise the Insurance Business Application.

The Insurance Company business goals and processes

The Insurance Company developed the Insurance Business Application to improve the speed, efficiency, flexibility, and security of business processes and to be able to apply real-time business rules for all processes. Through the Insurance Business Application, the Insurance Company already realized an improvement in the time required to process new applications to 45 minutes or less. The Insurance Company wants to achieve similar improvements for other insurance business processes. Through the application of business rules, the Insurance Company plans to temporarily modify business processes for such events as limited-time offers and marketing campaigns. The architecture and componentry of the Insurance Business Application provide the flexibility to achieve these goals from a software perspective.

The Insurance Business Application follows the style of a service-oriented architecture (SOA) in which an application is composed of the orchestration of a set of services supplied by service providers. Figure 24 on page 62

on page 62 illustrates an architectural view of the multiple layers of this application.




Figure 24. Architecture view of the Insurance Business Application

- The User layer illustrates the individuals or computer systems that need to interact with business services of the Insurance Company. These participants include insurance agents and regional agencies, insurance policy underwriters, policy owners and other clients, and interested parties such as medical insurers or providers.
- The Access Interface layer provides the view into the IT systems for different users. This enables the development of different channels and workplaces, which provide and control access to business services as appropriate for each participant.
- The Processes layer defines what tasks, in what order, and under what conditions, need to be accomplished by what resource (user or system). This layer enables procedural control and monitoring.
- The Business rules layer codifies the business-level controls that are separate from the procedural flow. This layer enables the Insurance Company to deploy variants to a business process without deploying a new process. For example, to support a limited-time marketing campaign for annuity policies, the Insurance Company can apply a set of business rules to modify their policy and underwriting processes.
- The Services layer separates what needs to be done from its implementation. This layer enables the aggregation of existing infrastructure with new infrastructure, speeds deployment of new implementation, and standardizes interfaces to functions. The Insurance Business Application uses business services that include:
 - The Policy Service, which provides functions related to insurance policies
 - The Underwriting Service (U/W), which provides functions that are related to performing life insurance underwriting
 - The Documentation Service, which provides functions that are related to creating, updating, handling, storing, and managing documents within the Insurance Company
 - Various administrative services, such as Billing, Claims Processing, and Marketing
- The Component layer provides the implementation of services, and controls access and changes to the Data layer. For this layer, the Insurance Business Application uses software from multiple ISVs and

IBM strategic software. For example, the applications or components involved to create and record a new insurance policy include a 'Policy Administration' vendor software package and DB2 for storing customer data.

- The Data layer represents all of the data used in the Insurance Company organization.

Figure 25 illustrates these layers in a three-tier view to indicate flows between users, processes, services, and components in the Insurance Business Application. This view is not a physical mapping of software onto hardware or operating systems; each process, service, and component can have several instances deployed across the hardware and software. “The Insurance Business Application in a zEnterprise ensemble” provides a general mapping of the application on the hardware and software environment.




Figure 25. Three-tier view of the Insurance Business Application

- Tier 1 consists of software elements that handle user interfaces and interactions, incoming work requests, and security checking for users and work requests. WebSphere Portal Server is one of the major software elements that supports these functions.
- Tier 2 consists of software elements that provide business logic and process orchestration. The Insurance Company uses the IBM WebSphere Enterprise Service Bus to integrate the applications and services that constitute the Insurance Business Application. WebSphere Application Server, WebSphere Process Server, and DB2 also contribute to the functions provided for this tier.
- Tier 3 consists of business services such as Policy, Claims, and Underwriting (U/W), and data storage services. As for Tier 2, the WebSphere products and DB2 contribute to the functions provided for Tier 3.

The Insurance Business Application in a zEnterprise ensemble

Figure 26 on page 64 provides a general illustration of the Insurance Business Application architecture physically mapped onto the z196 and the zBX in the Insurance ensemble.




Figure 26. Mapping the Insurance Business Application to the Insurance ensemble

- The User layer is represented on the left, with two methods of access: Intranet access for employees and Internet access through a firewall for Insurance Company clients and business partners.
- The Data layer is represented in each type of virtual server in the zEnterprise environment. The Insurance Company plans to keep its data warehouse on z/OS, shown by the large databases in the topmost Parallel Sysplex. Information in this data warehouse is copied out to smaller databases on the other virtual servers, for faster access to portions of the data that are most relevant to the applications that run on those virtual servers. The Insurance Company uses various DB2 products to provide this support.
- The remaining architectural layers of the Insurance Business Application are supported on various types of virtual server. Each user interface, process, service, and component can be composed of various software elements, and can have several instances of those elements deployed across the hardware and software.

For example, the Insurance Company uses several IBM WebSphere products to integrate and orchestrate business processes for the Insurance Business Application. Most of these products require WebSphere Application Server to run on the same system. So instances of WebSphere Application Server run on POWER blades and System x blades in the zBX, z/OS LPARs, and z/VM guests on a combination of operating systems: z/OS V1R13 and V1R12; z/VM V6R1; and SLES 10 Linux for System z.

The Insurance Business Application on the z196

On virtual servers in the z196, the Insurance Company plans to install a combination of IBM and ISV applications that handle mission-critical business processes and data. Figure 27 on page 65 provides a conceptual illustration of these applications, the Insurance Business Application architecture layers that they support, and the services and components that also run on virtual servers in the z196.




Figure 27. Software support for layers of the Insurance Business Application on the z196

On the z/OS LPARs for production and test, the combination of IBM and ISV applications includes:

- WebSphere Portal Server, which is a high-performance business process automation engine that helps form processes that meet business goals
- WebSphere Enterprise Service Bus, which is a flexible connectivity infrastructure for integrating applications and services
- DB2 Content Manager Enterprise Edition, which manages all types of content for multiple platforms, databases, and applications
- Vendor software that provides the functions to support the Policy, Underwriting, Documentation, and other services for the Insurance Business Application.
- DB2 for z/OS, the relational database management system (DBMS) for the z/OS operating system
- WebSphere Application Server, which provides the foundation for building, reusing, running, integrating, and managing SOA applications and services
- Lightweight Directory Access Protocol (LDAP) on z/OS, which provides the foundation for system and application security

On the z/VM guests, a similar combination of IBM and ISV applications runs on the Linux operating system. These virtual servers also have Apache HTTP Server, which provides fully integrated support for secure HTTPS transactions.

The Insurance Business Application on the zBX

On virtual servers on the POWER blades and System x blades in the zBX, the Insurance Company plans to install a combination of IBM and ISV applications that primarily secure, generate, and present online forms for Insurance Company employees and clients. Figure 28 on page 66 provides a conceptual illustration of these applications, the Insurance Business Application architecture layers that they support,

and the services and components that also run on the POWER blades and System x blades.




Figure 28. Software support for layers of the Insurance Business Application on POWER and System x blades

The combination of IBM and ISV applications includes:

- Vendor software that provides the functions to support the Policy, Underwriting, and Documentation services for the Insurance Business Application.
- IBM Lotus® Forms and IBM FileNet® P8 Platform products for document management and storage
- DB2 for Linux, UNIX and Windows, the relational database management system (DBMS) for the AIX operating system
- WebSphere Application Server, which provides the foundation for building, reusing, running, integrating, and managing SOA applications and services

These applications and components run on the following operating systems:

- On AIX on POWER blades with the PowerVM hypervisor
- On Microsoft Windows or Linux on System x blades with the zBX integrated hypervisor for IBM System x blades (x Hyp).

The virtual servers supporting the Insurance Business Application

The following topics provide conceptual illustrations of the virtual servers that the Insurance Company administrators plan to configure in the Insurance ensemble. On virtual servers in the z196, the Insurance Company plans to install a combination of IBM and ISV applications that handle mission-critical business processes and data. On virtual servers on the POWER blades and System x blades in the zBX, the Insurance Company plans to install a combination of IBM and ISV applications that primarily secure, generate, and present online forms for Insurance Company employees and clients.

Annuity server software configuration on the z196

Figure 29 on page 67 illustrates the virtual servers on the z196 that support the Annuity policy processes and software. For each virtual server, the figure indicates the installed operating system and a partial list of installed software. These virtual servers support both production and test work.




Figure 29. Software support for Annuity policies on the z196

Annuity server software configuration on the zBX

Figure 30 on page 68 illustrates a subset of the virtual servers in the zBX that support the Annuity policy processes and software. The figure shows virtual servers on only select POWER blades and one System x blade, with the installed operating system and a partial list of installed software. The set of virtual servers on the B.1.03 blade support test work, while those on blades B.1.11 and B.1.14 support production work only.




Figure 30. Software support for Annuity policies on blades in the zBX

Note that each virtual server on the POWER blades has its own AIX image; the virtual servers on the test blades run AIX V6.1, while the virtual servers on the production blades run AIX V5.3.

The full Annuity server software configuration for the System x blades in the zBX is similar.

- Some System x blades are reserved for test and development, while others are used for production. Figure 30 shows only part of the configuration: one set of virtual servers that support production work.
- Each virtual server on a System x blade has its own Linux on System x or Windows Server 2008 operating system image.
- Each virtual server supports a combination of IBM and ISV software applications that handle insurance policy processing.

Variable Life server software configuration on the z196

Figure 31 illustrates the virtual servers on the z196 that support the Variable Life policy processes and software. For each virtual server, the figure indicates the installed operating system and a partial list of installed software. These virtual servers support both production and test work.




Figure 31. Software support for Variable Life policies on the z196

Variable Life server software configuration on the zBX

Figure 32 illustrates a subset of the virtual servers on the zBX that support the Variable Life policy processes and software. The figure shows virtual servers on only select POWER blades and one System x blade, with the installed operating system and a partial list of installed software. The set of virtual servers on the B.1.03 blade support test work, while the virtual servers on blade B.1.07 and B.1.13 support production work only.




Figure 32. Software support for Variable Life policies on blades in the zBX

The Variable Life server software configuration for the System x blades in the zBX is similar:

- Some System x blades are reserved for test and development, while others are used for production. Figure 32 shows only part of the configuration: one set of virtual servers that support production work.
- Each virtual server on a System x blade has its own Linux on System x or Windows Server 2008 operating system image.
- Each virtual server supports a combination of IBM and ISV software applications that handle insurance policy processing.

Appendix. 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 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>.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

Microsoft and Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

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

Other product and service names might be trademarks of IBM or other companies.

Electronic emission notices

The following 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 instructions contained in the installation 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, by installation or use of this equipment other than as specified in

the installation manual, or by any other 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.

Canadian Department of Communications Compliance Statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité aux normes du ministère des Communications du Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

European Union (EU) Electromagnetic Compatibility Directive

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 equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

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.

European Community contact:

IBM Deutschland GmbH
Technical Regulations, Department M372
IBM-Allee 1, 71139 Ehningen, Germany
Telephone: 0049 (0) 7032 15-2941
email: lugi@de.ibm.com

EC Declaration of Conformity (In German)

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 89/336/EWG 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 der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der 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 89/336/EWG in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) vom 18. September 1998 (bzw. der EMC EG Richtlinie 89/336) für Geräte der Klasse A.

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Konformitätserklärung nach Paragraph 5 des EMVG ist die IBM Deutschland GmbH, 70548 Stuttgart.

Informationen in Hinsicht EMVG Paragraph 4 Abs. (1) 4:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

update: 2004/12/07

People's Republic of China Class A Compliance Statement

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 actions.

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。

Japan Class A Compliance Statement

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.

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

Korean Class A Compliance Statement

이 기기는 업무용(A급)으로 전자파적합등록을 한 기기이오니
판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의
지역에서 사용하는 것을 목적으로 합니다.

Taiwan Class A Compliance Statement

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.

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。

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

Glossary

A.

advanced management module (AMM). A hardware unit that provides system-management functions for all the blade servers in a BladeCenter chassis.

alternate HMC. A System z Hardware Management Console (HMC) that is paired with the primary HMC to provide redundancy.

See also [primary HMC](#).

AMM. See [advanced management module](#).

appliance. A software device that provides a narrow range of functions and generally runs on a hardware platform.

application environment. The environment that includes the software and the server or network infrastructure that supports it.

ARM-instrumented application. An application in which application response measurement (ARM) calls are added to the source code so that management systems can monitor the performance of the application. ARM is an Open Group standard.

Automate suite (Automate). The second of two suites of functionality associated with the IBM zEnterprise Unified Resource Manager. The Automate suite includes goal-oriented monitoring and management of resources and energy management.

See also [Manage suite](#).

availability policy. A description of the availability objectives for the virtual servers and element groups that support a workload resource group. Availability objectives consist of a workload impact setting and redundancy objectives for an element group.

B.

blade. A hardware unit that provides application-specific services and components. The consistent size and shape (or form factor) of each blade allows it to fit in a BladeCenter chassis.

BladeCenter chassis. A modular chassis that can contain multiple blades, allowing the individual blades to share resources such as the management, switch, power, and blower modules.

C.

central processor complex (CPC). A physical collection of hardware that consists of main storage, one or more central processors, timers, and channels. In

the zEnterprise environment, the CPC consists of a zEnterprise mainframe and any attached IBM zEnterprise BladeCenter Extension (zBX).

See also [node](#) and [zCPC](#).

classification rule. A rule used by System z workload resource group manager firmware and software to assign a service class.

CPC. See [central processor complex](#).

D.

DataPower XI50z. See [IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise](#).

discretionary goal. A service class performance goal assigned to low priority work that does not have any specific performance goal. Work is run when system resources are available.

E.

element group. A collection of functionally equivalent, redundant virtual servers that can belong to more than one workload resource group.

ensemble. A collection of one or more zEnterprise nodes (including any attached zBX) that are managed as a single logical virtualized system by the Unified Resource Manager, through the Hardware Management Console.

ensemble member. A zEnterprise node that has been added to an ensemble.

See also [node](#).

F.

firmware. Licensed Internal Code (LIC) that is shipped with hardware. Firmware is considered an integral part of the system and is loaded and run at power on. Firmware is not open for customer configuration and is expected to run without any customer setup.

G.

GPMP. See [guest platform management provider](#).

guest platform management provider (GPMP). An optional suite of applications that is installed in specific z/OS, Linux, and AIX operating system images to support platform management functions. For example, the guest platform management provider collects and aggregates performance data for virtual servers and workload resource groups.

H.

Hardware Management Console (HMC). A user interface through which data center personnel configure, control, monitor, and manage System z hardware and software resources. The HMC communicates with each central processor complex (CPC) through the Support Element. On an IBM zEnterprise 196 (z196), using the Unified Resource Manager on the HMCs or Support Elements, personnel can also create and manage an ensemble.

See also primary HMC and alternate HMC.

HMC. See Hardware Management Console.

hypervisor. A program that allows multiple instances of operating systems or virtual servers to run simultaneously on the same hardware device. A hypervisor can run directly on the hardware, can run within an operating system, or can be imbedded in platform firmware. Examples of hypervisors include PR/SM, z/VM, and PowerVM Enterprise Edition.

I.

IBM blade. A customer-acquired, customer-installed select blade to be managed by IBM zEnterprise Unified Resource Manager. One example of an IBM blade is a POWER7® blade.

IBM System z Advanced Workload Analysis Reporter (IBM zAware). Firmware consisting of an integrated set of applications that monitor software running on z/OS and model normal system behavior. IBM zAware pattern recognition techniques identify unexpected messages, providing rapid diagnosis of problems caused by system changes. Operational controls and views of analytical data are available through the IBM zAware graphical user interface (GUI).

IBM System z Application Assist Processor (zAAP). A specialized processor that provides a Java execution environment, which enables Java-based web applications to be integrated with core z/OS business applications and backend database systems.

IBM System z Integrated Information Processor (zIIP). A specialized processor that provides computing capacity for selected data and transaction processing workloads and for selected network encryption workloads.

IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z). A purpose-built appliance that simplifies, helps secure, and optimizes XML and Web services processing.

IBM zAware. See IBM System z Advanced Workload Analysis Reporter (IBM zAware).

IBM zAware disaster recovery environment. An IBM zAware environment that is created expressly for disaster recovery purposes.

IBM zAware environment. A configuration that consists of an IBM zAware partition and the IBM zAware monitored clients that are sending information to the IBM zAware server that is running on the partition. The IBM zAware monitored clients do not have to run in the same IBM zAware host system that contains the partition.

IBM zAware host system. The zEC12 central processor complex (CPC) that contains the logical partition (LPAR) in which the IBM System z Advanced Workload Analysis Reporter (IBM zAware) runs.

IBM zAware model. A description of normal behavior that an IBM zAware server generates for a specific monitored z/OS system. Initially, this model is based on prior data (system and application messages) from the operations log (OPERLOG) for the z/OS system. The model is updated periodically and can be modified to include or exclude specific days of system operation. The IBM zAware server uses this model to detect system problems that are indicated in current data that the server receives from the specific z/OS system.

IBM zAware monitored client. A z/OS partition that sends OPERLOG logstream data to the IBM System z Advanced Workload Analysis Reporter (IBM zAware) for analysis. To detect problems, IBM zAware compares the system and application messages in these log files to a model of normal behavior for this z/OS system, and highlights anomalous results through the IBM zAware graphical user interface (GUI).

IBM zAware partition. The logical partition (LPAR) in the zEC12 central processor complex (CPC) in which only an IBM zAware server runs. The IBM zAware graphical user interface (GUI) provides operational controls and views of analytical data for IBM zAware monitored clients.

IBM zAware server. An instance of the IBM System z Advanced Workload Analysis Reporter (IBM zAware) that is receiving data from monitored clients.

IBM zEnterprise 114 (z114). The newest generation of the entry System z family of servers built on a new processor chip, featuring a 14-way core design with enhanced memory function and capacity, security, and on demand enhancements to support existing mainframe workloads and consolidation.

IBM zEnterprise 196 (z196). The previous generation of the System z high end family of servers built on a new processor chip, featuring a 96-way core design with enhanced memory function and capacity, security, and on demand enhancements to support existing mainframe workloads and large scale consolidation.

IBM zEnterprise BladeCenter Extension (zBX). A heterogeneous hardware infrastructure that consists of a BladeCenter chassis attached to a zEC12, z196, or z114. A BladeCenter chassis can contain IBM blades or optimizers.

IBM zEnterprise BladeCenter Extension (zBX) blade. Generic name for all blade types supported in an IBM zEnterprise BladeCenter Extension (zBX). This term includes IBM blades and optimizers.

IBM zEnterprise EC12 (zEC12). The newest generation of the System z high end family of servers built on a new processor chip, featuring a 120-way core design with enhanced memory function and capacity, security, and on demand enhancements to support existing mainframe workloads and large scale consolidation.

IBM zEnterprise System (zEnterprise). A heterogeneous hardware infrastructure that can consist of a zEC12, z196, or z114 and an attached IBM zEnterprise BladeCenter Extension (zBX), managed as a single logical virtualized system by the Unified Resource Manager.

IBM zEnterprise Unified Resource Manager. Licensed Internal Code (LIC), also known as firmware, that is part of the Hardware Management Console. The Unified Resource Manager provides energy monitoring and management, goal-oriented policy management, increased security, virtual networking, and data management for the physical and logical resources of a given ensemble.

IEDN. See intraensemble data network (IEDN).

IEDN TOR switch. See intraensemble data network (IEDN) TOR switch.

INMN. See intranode management network (INMN).

intraensemble data network (IEDN). A private high-speed network for application data communications within an ensemble. Data communications for workload resource groups can flow over the IEDN within and between nodes of an ensemble. The Unified Resource Manager configures, provisions, and manages all of the physical and logical resources of the IEDN.

intraensemble data network (IEDN) TOR switch. A top-of-rack switch that provides connectivity to the intraensemble data network (IEDN), supporting application data within an ensemble.

intranode management network (INMN). A private service network that the Unified Resource Manager uses to manage the resources within a single zEnterprise node. The INMN connects the Support Element to the zEC12, z196, or z114 and to any attached IBM zEnterprise BladeCenter Extension (zBX).

M.

Manage suite (Manage). The first suite of functionality associated with the IBM zEnterprise Unified Resource Manager. The Manage suite includes

core operational controls, installation, and configuration management, and energy monitoring.

management TOR switch. A top-of-rack switch that provides a private network connection between a zEC12, z196, or z114 Support Element and an IBM zEnterprise BladeCenter Extension (zBX).

member. See ensemble member.

N.

network interface card (NIC). A printed circuit board that plugs into a server. It controls the exchange of data over a network and provides the electronic functions for the data link protocol or access method, such as token ring or Ethernet.

NIC. See network interface card.

node. A single zEC12, z196, or z114 and any optionally attached IBM zEnterprise BladeCenter Extension (zBX). A node can be a member of only one ensemble.

See also central processor complex.

O.

optimizer. A special-purpose hardware component or appliance that can perform a limited set of specific functions with optimized performance when compared to a general-purpose processor. Because of its limited set of functions, an optimizer is an integrated part of a processing environment, rather than a standalone unit.

One example of an optimizer is the IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise.

out-of-band monitoring solution. A type of monitoring solution that runs on a dedicated server rather than relying on the use of a monitoring agent installed in the operating system. For example, the IBM System z Advanced Workload Analysis Reporter (IBM zAware) provides out-of-band monitoring because it runs in a dedicated PR/SM partition and monitors clients that run in other partitions in System z servers.

OSM. An OSA-Express channel path identifier (CHPID) type that provides connectivity to the intranode management network (INMN).

OSX. An OSA-Express channel path identifier (CHPID) type that provides connectivity to the intraensemble data network (IEDN).

P.

performance index. A number that indicates whether the performance goal for a service class was achieved, exceeded, or missed.

performance policy. A description of the performance objectives and importance of a workload resource group.

platform management. The subset of systems management focused on hardware and virtualization management.

PowerVM. See [PowerVM Enterprise Edition](#).

PowerVM Enterprise Edition (PowerVM). A hypervisor that provides a set of comprehensive systems technologies and services designed to enable aggregation and management of IBM POWER blade resources through a consolidated, logical view.

primary HMC. The System z Hardware Management Console (HMC) through which data personnel create and manage an ensemble. This HMC owns configuration and policy information that the Unified Resource Manager uses to monitor, manage, and adjust resources for all members of this ensemble.

See also [alternate HMC](#).

private system control network (PSCN). The private subsystem of the System z servers that is controlled by a fully redundant dual-Ethernet communications network. This network provides communication to all field-replaceable units (FRUs) and hierarchic control through a mirrored system of control cards and IP addresses. The PSCN provides a means for subsystems to communicate and control the dynamic parameters of system operation. The PSCN also supports error reporting, failure data collection and recovery detection, and correction of both the internal hardware and firmware of the System z servers.

PSCN. See [private system control network](#).

R.

rack. A free-standing structure or frame that can hold multiple servers and expansion units, such as BladeCenter blades.

redundancy objectives. An availability goal that identifies the minimum number and the preferred number of virtual servers in an element group that are required to support specific function for a workload resource group. Base redundancy objectives are set in the definition of the element group, but can be temporarily modified to suit a specific workload through an optional setting in an availability policy.

response time goal. A service class performance goal that defines end-to-end response time of work requests.

S.

service class. A collection of work that has the same service goals or performance objectives, resource requirements, or availability requirements.

static power save mode. A zEC12, z196, or z114 function used for periods of low utilization or potentially when a CBU system is sitting idle waiting to take over in the event of a failure. The server uses

frequency and voltage reduction to reduce energy consumption of the system. The customer initiates static power save mode by using the HMC or Support Element or Active Energy Manager.

T.

top-of-rack (TOR) switch. A network switch that is located in the first rack of an IBM zEnterprise BladeCenter Extension (zBX).

TOR switch. See [intraensemble data network \(IEDN\) TOR switch](#) and [management TOR switch](#).

transaction. A unit of processing consisting of one or more application programs, affecting one or more objects, that is initiated by a single request.

U.

Unified Resource Manager. See [IBM zEnterprise Unified Resource Manager](#).

V.

velocity goal. A service class performance goal that defines the acceptable amount of delay for work when work is ready to run. Velocity is the measure of how fast work should run when ready, without being delayed by contention for managed resources.

virtual appliance. A prepackaged software application that provides some well-defined business workflow, making it easier to deploy a solution with minimal configuration. Many tiers of operating system and applications can be packaged as a single virtual appliance. These tiers can depend on the hardware resources of different architectures.

See also [virtual server collection](#) and [virtual server image](#).

virtual server. A logical construct that appears to comprise processor, memory, and I/O resources conforming to a particular architecture. A virtual server can support an operating system, associated middleware, and applications. A hypervisor creates and manages virtual servers.

virtual server collection. A set of virtual servers that supports a workload resource group. This set is not necessarily static. The constituents of the collection at any given point are determined by the virtual servers involved in supporting the workload resource group at that time.

See also [virtual appliance](#) and [virtual server image](#).

virtual server image. A package containing metadata that describes the system requirements, virtual storage drives, and any goals and constraints for the virtual machine (for example, isolation and availability). The Open Virtual Machine Format (OVF) is a Distributed Management Task Force (DMTF) standard that describes a packaging format for virtual server images.

See also [virtual appliance](#) and [virtual server collection](#).

virtual server image capture. The ability to store metadata and disk images of an existing virtual server. The metadata describes the virtual server storage, network needs, goals, and constraints. The captured information is stored as a virtual server image that can be referenced and used to create and deploy other similar images.

virtual server image clone. The ability to create an identical copy (clone) of a virtual server image that can be used to create a new similar virtual server.

W.

workload. The amount of application processing that a computer performs at a given time. In z/OS WLM, a workload is a customer-defined collection of work to be tracked, managed, and reported as a unit. For zEnterprise, see [workload resource group](#).

workload impact setting. A list of the virtual servers that the Unified Resource Manager is to ignore when calculating the availability status of a workload resource group. A workload impact setting is an optional part of an availability policy.

workload resource group. A collection of virtual servers that perform a customer-defined collective purpose. A workload resource group generally can be viewed as a multi-tiered application. Each workload resource group is associated with a set of policies that define performance goals.

Z.

z/VM single system image (SSI) cluster. A collection of z/VM systems (called members) that can be managed, serviced, and administered as one system within which workloads can be deployed. An SSI cluster is intended to share a set of resources among all members.

z114. See [IBM zEnterprise 114 \(z114\)](#).

z196. See [IBM zEnterprise 196 \(z196\)](#).

zEC12. See [IBM zEnterprise EC12 \(zEC12\)](#).

zAAP. See [IBM System z Application Assist Processor](#).

zBX. See [IBM zEnterprise BladeCenter Extension \(zBX\)](#).

zBX blade. See [IBM zEnterprise BladeCenter Extension \(zBX\) blade](#).

zCPC. The physical collection of main storage, central processors, timers, and channels within a zEnterprise mainframe. Although this collection of hardware resources is part of the larger zEnterprise central processor complex, you can apply energy management

policies to the zCPC that are different from those that you apply to any attached IBM zEnterprise BladeCenter Extension (zBX) or blades.

See also [central processor complex](#).

zIIP. See [IBM System z Integrated Information Processor](#).

zEnterprise. See [IBM zEnterprise System \(zEnterprise\)](#).

Unified Resource Manager. See [IBM zEnterprise Unified Resource Manager](#).

Index

A

application programming interface (API)
 zEnterprise system web services
 api 55
 availability management 52

C

changing data center 1
 common concerns 2
 speciality engines 2
 zEnterprise and ensembles 5
 common concerns
 firmware and software
 management 4
 infrastructure 2
 specialization management 4
 virtualization 3
 configuration examples 17
 mixed system environment 19
 with zBX 18
 without zBX 17
 zEnterprise systems, each with a zBX
 attached 18

D

DataPower XI50z
 software 29

E

energy management 53
 ensemble 6, 9
 basic configuration 9
 configuration examples 17
 mixed system environment 19
 with zBX 18
 without zBX 17
 zEnterprise systems, each with a
 zBX attached 18
 customer-provided data network 14,
 15
 firmware 20, 24
 hardware 20
 IBM zEnterprise BladeCenter
 Extension 21
 zEnterprise 20
 Hardware Management Console 10
 hypervisors 22
 IBM System x blades 24
 PowerVM 23
 PR/SM 22
 z/VM LPAR 22
 hypervisors and firmware 22
 infrastructure 13
 intraensemble data network
 (IEDN) 13

ensemble (*continued*)
 intranode management network
 (INMN) 14
 network connection distances 20
 optimizations 29
 private system control network 15
 select IBM blade software 30
 software 20, 24
 DataPower XI50z 29
 Linux on System z 29
 POWER blade 30
 System x blade 31
 z/OS 24
 z/VM 25
 ensemble membership management 42
 ensemble system management 33

H

hardware
 IBM zEnterprise BladeCenter
 Extension 21
 zEnterprise 20
 Hardware Management Console 10
 HiperSockets 14, 15
 HiperSockets network 14
 hypervisor management 40
 hypervisors
 IBM System x blades 24
 PowerVM 23
 PR/SM 22
 z/VM LPAR 22

I

IBM System x blades
 hypervisor 24
 IBM Websphere DataPower Integration
 Appliance 29
 IBM zEnterprise Unified Resource
 Manager (Unified Resource
 Manager) 33
 IEDN enabled HiperSockets 14
 introduction 1
 IQD linked HiperSockets 15

L

Linux on System z 29

N

network
 customer-provided data network 14,
 15
 HiperSockets 14
 intraensemble data network
 (IEDN) 13

network (*continued*)
 intranode management network
 (INMN) 14
 maximum distances 20
 private system control network 15
 networks
 IEDN enabled HiperSockets 14
 IQD linked HiperSockets 15

O

optimizers
 DataPower XI50z software 29

P

planning 57
 ensemble management guide 57
 migration scenarios 57
 sample implementation 58
 POWER blade
 software 30
 PowerVM
 hypervisor 23
 PR/SM
 hypervisor 22

S

software
 DataPower XI50z 29
 Linux on System z 29
 POWER blade 30
 System x blade 31
 z/OS 24
 z/VM 25
 storage management 44
 System x blade
 software 31

T

tasks
 availability management 52
 energy management 53
 ensemble management guide 57
 ensemble membership
 management 42
 hypervisor management tasks 41
 manage storage resources 44
 virtual network management 48
 virtual server management 46
 web services api tasks 56
 workload and availability
 management 52
 workload and performance
 management 49

V

virtual network management 47
virtual server management 46

W

workload and availability
management 52
workload and performance
management 49

Y

your system 57

Z

z/OS 24
z/VM 25
z/VM LPAR
hypervisor 22
zEnterprise system web services api
api characteristics 55
web services api tasks 55



Printed in USA

GC27-2609-06

