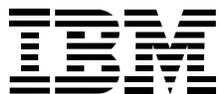

IBM System z
Introduction
July 2014

IBM zEnterprise EC12 (zEC12) and IBM zEnterprise BC12 (zBC12)

Frequently Asked Questions

Worldwide



ZSQ03058-USEN-16

Table of Contents

Announcements	3
IBM zEnterprise BC12 (zBC12) Hardware	4
zEC12 Hardware	10
Academic Initiative and Skills	17
zEDC Express and Compression Acceleration.....	18
10GbE RoCE Express and SMC-R	27
Transactional Execution	33
IBM zAware.....	35
Flash Express.....	42
zEC12 and zBC12 Warranty	49
IBM zEnterprise BladeCenter Extension (zBX)	50
POWER7 Blades for zBX	56
IBM System x blades for zBX.....	60
IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z).....	65
zBX Networking	68
zBX Availability	71
zBX and blade Warranty and Maintenance.....	74
IBM zEnterprise Unified Resource Manager.....	76
Ensemble	81
Hardware Management Console (HMC)	84
zBC12 Upgradeability	87
zEC12 Upgradeability	88
Software Pricing – zBC12.....	89
Software Pricing – zEC12	92
Performance.....	96
Capacity on Demand offerings.....	100
I/O	107
Networking	113
Older, but important, networking information.....	114
Water cooling option	117
Overhead Cabling and Power.....	121
zEC12 Power requirements (including High Voltage DC Power option)	122
Integrated Facility for Linux (IFL).....	124
zAAPs and zIIPs	125
zEnterprise and PureSystems	128
Security	130
IBM System Storage & zEnterprise	136
STP	140
Parallel Sysplex	142
Cloud and Mobile.....	146
GDPS	147
z/OS.....	148
Linux on System z.....	149
z/VM.....	153
z/VSE.....	155
Statements of Direction	157

Announcements

Question:

What is being announced for the IBM zEnterprise® System family of products on February 24, 2014?

Answer:

IBM is strengthening its commitment to support Smarter Computing by announcing enhancements to the IBM zEnterprise EC12 (zEC12), the IBM zEnterprise BC12 (zBC12), and the IBM zEnterprise BladeCenter® Extension (zBX). New support provided for the IBM WebSphere® DataPower® Integration Appliance XI50 for zEnterprise firmware V6.0 can help with the rapid data transformation for cloud and mobile applications.

A mobile application is now available, IBM Mobile Systems Remote, to allow you to display the configuration, status, and metrics information for systems managed by the IBM System z® Hardware Management Console (HMC). The application can be downloaded from the web to your iOS devices from the Apple App Store or, for Android devices, from the Google Play Store.

Additional Common Cryptographic Architecture enhancements are now being made available in support of Die Deutsche Kreditwirtschaft (DK) AES PIN, new message authentication codes (MAC) using the AESCMAC algorithm, and User Defined Extension (UDX) simplification for PKA Key Translate. These enhancements are available on the zEC12, zBC12, IBM zEnterprise 196 (z196), and IBM zEnterprise 114 (z114) servers.

Delivering on statements of direction, z/VM® V6.3 now supports guest exploitation for compression acceleration using the zEDC Express feature and Remote Direct Memory Access over Converged Ethernet using the 10GbE RoCE Express feature, which is designed to help improve data transmission and communications across servers. Two additional enhancements are being offered for z/VM V6.3 environments - CPU pooling and an environment information interface.

z/OS® V2.1 was made generally available September 30, 2013. To further enhance the groundwork for the next tier of mainframe computing, additional enhancements are now being offered for both z/OS V2.1 environments and for z/OS platform products.

An enhancement to CFCC Level 19 is designed for exploitation of the Flash Express feature to help improve the resiliency of the Coupling Facility with cost-effective standby capacity to handle overflow of WebSphere MQSeries® shared queues.

Through our continuous investment, the zEnterprise System is built to help you embrace IT and shifting market forces. It can help you reshape the value you deliver to your customers with a level of efficiency and economics unmatched in the industry.

IBM zEnterprise BC12 (zBC12) Hardware

Question:

Tell me about the processor chip used in the IBM zEnterprise BC12 (zBC12)?

Answer:

The zBC12 leverages the same microprocessor technology as the zEnterprise EC12 (zEC12) but is tuned and optimized for the processing requirements of the entry-level mainframe.

The core design combines all the best qualities of the latest mainframe design philosophy with a generous infusion of POWER7®-inspired computational muscle, a reflection of the teamwork between IBM's best processor designers.

The microprocessor chip has a higher-frequency superscalar design, improved cache structure, second generation “out of order” execution sequence and is optimized to provide software performance improvements for Java™, PL/I, compilers, DB2® and more.

The zBC12 is powered by up to 18 microprocessors running at 4.2 GHz and delivers up to 36 percent improvement in performance per core compared to its predecessor, the zEnterprise 114 (z114), and up to 58% more total traditional workload capacity.

Question:

What about the zBC12 makes it so adept at consolidation and infrastructure simplification?

Answer:

Today, most datacenters contain mainframe and distributed systems, running a diverse set of workloads or industry applications. The zEnterprise BC12 is designed to help address the complexity and inefficiencies of these siloed IT infrastructures.

As businesses look to transform their IT infrastructures, many choose to start this journey by addressing existing inefficiencies consolidating workloads to lower the cost of operations. IBM System z can do this on a grand scale. The extreme virtualization capabilities of the zBC12 can support an average of 40 distributed servers or more on a single core depending on the workload, or 520 in a single footprint, delivering a virtual Linux® server for as low as \$1.00 per day and providing an opportunity to help you to collapse infrastructures and drive greater data center efficiencies.

Unlike other proclaimed cloud solutions that are defined by a siloed architecture resource pool, the zBC12 leaps beyond the virtualization of a single platform and can integrate, virtualize and manage multiple server architectures for optimal application placement. You have the flexibility to deploy and manage applications across a virtualized pool of heterogeneous resources and operating environments including z/OS, z/VM, z/VSE®, Linux on System z, AIX®, Linux on IBM System x® and Microsoft® Windows®.

When configured with a IBM zEnterprise BladeCenter Extension (zBX), you can combine System z, UNIX® and Intel® server technologies into a single unified system—integrating workloads with affinity to mainframe applications and data—and with the zEnterprise Unified Resource Manager, manage it all with the same tools, techniques and resources for consistent, automated and reliable service delivery.

The zBC12 offers a range of scaling capabilities:

- Scale up—over 4900 general purpose MIPS in a single footprint
- Scale out—consolidate up to 520 distributed servers
- Scale within—specialty engines, cryptographic processors, hypervisors
- Scale beyond its traditional boundaries—when configured with the zBX—supports the integration of up to 112 distributed blade servers and workload optimizers

With improved processor performance, increased capacity, new hybrid computing capabilities, and significant power, space and cooling benefits, the zBC12 is now a genuine “data-center in a box” solution and a perfect fit for infrastructure simplification and true Cloud Computing.

Question:

What is the machine type of the zBC12?

Answer:

The zBC12 machine type is 2828.

Question:

What zBC12 models were announced?

Answer:

The zBC12 is available in two models; a single central processing drawer model, the H06 and a two drawer model, the H13 which offers the additional flexibility for I/O and coupling expansion and/or increased specialty engine capability. With up to 13 configurable cores, the model naming is indicative of how many total processor units are available for customer purchase and characterization.

The cores can be configured as general purpose processors (CPs), Integrated Facilities for Linux (IFLs), System z Application Assist Processors (zAAPs), System z Integrated Information Processors (zIIPs), Internal Coupling Facilities (ICFs), additional System Assist Processors (SAPs), or used as additional spares. And on the H13, you can have up to two “dedicated” spares.

With up to 6 CPs and 156 available capacity settings offered across either model, you have the freedom to choose the right capacity setting for your needs with the flexibility to scale on demand as workload requirements increase.

And we offer both IFL only and ICF only configurations where you can have up to 13 IFLs or ICFs in a standalone environment.

Question:

How much memory is available on the zBC12? What is the RAIM feature?

Answer:

The zBC12 supports up to 512 GB of real (usable) RAIM-protected memory—an industry exclusive currently available only on System z. Beyond the 496 GB of “customer purchasable memory,” there is an additional 16 GB of memory for the Hardware System Area (HSA) which holds the I/O configuration data for the server.

The minimum initial amount of memory that can be ordered is 8 GB on the Model H06 and 16 GB for the Model H13. Memory can be ordered in increments of 8 or 32 GB.

The zEnterprise servers – zBC12, zEC12, z114 and zEnterprise 196 (z196) – are the only high-end servers in the industry to offer the redundant array of independent memory (RAIM) technology. RAIM is similar to what is known in the disk storage industry as RAID. RAIM technology provides protection for the dynamic random access memory (DRAM), dual inline memory modules (DIMMs), and at the memory channel level, delivers the most resilient memory subsystem to date. RAIM is always active, and IBM has already factored internal RAIM overhead into all advertised memory configuration amounts.

Question:

To upgrade from an IBM System z10 Business Class™ (z10 BC™) to a zBC12, is it a field upgrade or a whole box replacement (push/pull)?

Answer:

There is a field upgrade from the z10 BC to the zBC12.

Question:

When is the zBC12 available?

Answer:

The general availability worldwide for the zBC12 is Sept 20, 2013.

Question:

Why have a two model packaging structure and what are the differences between the H06 and H13?

Answer:

We have spread the resources across two drawers on the zBC12 with each drawer separately containing two Single Chip Modules (SCMs), one storage control module and slots for ten DIMMs for memory. This allows us to avoid shipping and charging for unnecessary technology that some clients will never need or use.

It is important to select the appropriate model for your processing needs. Other than just for capacity or for raw engine count, there are other variables or reasons that may dictate the need for an H13 such as:

- When you require more than six total processor units (mix of CPs or specialty engines) or are on the verge today and want to be able to grow non-disruptively into the H13 (e.g. you are on a 5-way today and want the ability to add more than one specialty engine in the future).
- If you need more than 240 GB of memory offered on the H06 you can get up to 496 GB of customer usable memory on the H13.
- There is a limit to the number of hub slots and amount of bandwidth that is available with the one drawer system (H06) and if you exceed that bandwidth or need additional I/O capability or have increased coupling link requirements then the 2nd drawer (H13) will be needed.

Upgrading from the H06 to H13 will require a scheduled outage so please plan accordingly.

Important planning note: There are many variables that can influence the proper determination of model structure and capacity setting requirements. Do not use “one number” capacity comparisons! Work with IBM technical support for capacity planning! Customers can now use zPCR. The IBM Processor Capacity Reference (zPCR) is a free tool available for download that can be used to size your System z processors. <http://www.ibm.com/support/techdocs/atmsastr.nsf/WebIndex/PRS1381>

Question:

How many spare processing cores are on the zBC12?

Answer:

There are two ‘dedicated’ IBM provided spares on the model H13 only. Unused or unassigned engines on either the H06 or H13 can be used as spares.

Question:

Can either of these ‘dedicated’ IBM provided spare processing cores be used for other purposes?

Answer:

No, these spares are exclusively reserved to provide failover in the extremely unlikely event of a processor failure. The spare cores protect all processor types (CPs, IFLs, zAAPs, zIIPs, and ICFs). Any additional spares or unassigned cores can be activated for other purposes within the limits of the allowable configuration.

Question:

How many System Assist Processors (SAPs) are on the zBC12?

Answer:

Two System Assist Processors (SAPs) are standard on both the model H06 and H13. Additionally, more SAPs may be acquired from the pool of available processing units within the model.

Question:

What is the integrated firmware processor (IFP)?

Answer:

The integrated firmware processor (IFP) is allocated from the pool of processors for the whole system. Unlike the characterized processors, you don't pay for the IFP. It is solely used for infrastructure management of the PCIe features – 10GbE RoCE Express and zEDC Express. It is not customer usable or 'visible'.

Question:

What is the LPAR absolute hardware capacity setting?

Answer:

Customers, particularly our Linux customers where software pricing is based on a "per core" basis, need the ability to limit an LPAR to a specific amount of hardware processor capacity in a way that is unaffected by subsequent physical or logical configuration changes. Introduction of firmware and software appliances creates an even greater need for this capability.

The LPAR absolute hardware capacity setting will allow for specification of an absolute capacity limit for an individual logical partition. This is specified in absolute processor capacity (for example 2.5 processors). Use of this setting should allow more granular and flexible software pricing for non-z/OS systems.

Question:

Is the zBC12 physically bigger than the z114?

Answer:

No, the zBC12 has the same physical footprint as the z114 and can use the same floor cut outs. If you choose the top-exit cabling option, the clearance adds six inches to the width of the standard zBC12.

Below are a summary of the Physical Configuration requirements for both models of the zBC12.

Physical Configuration	H06	H13
Max Weight (Base/w Batteries/w Batteries & Overhead I/O cabling/w batteries, Overhead I/O & Balanced Power)	1802 lbs/2028 lbs/2123 lbs/ 2235 lbs	2064 lbs/2290 lbs/2385 lbs/2497 lbs)
Footprint	30 in. W x 50 in. D or 0.97 sq meters (10.42 sq ft)	30 in. W x 50 in. D or 0.97 sq meters (10.42 sq ft)
Service	36 in. W x 140 in. D or 3.16 sq meters (30.38 sq ft)	36 in. W x 140 in. D or 3.16 sq meters (30.38 sq ft)
Product Dimensions (W x D x H) ¹³	30.87 x 62 x 79.26 in. (784 x 1575 x 2013 mm)	30.87 x 62 x 79.26 in. (784 x 1575 x 2013 mm)

Question:

Does the zBX Model 002 attach to the zBC12?

Answer:

No, your zBX Model 002 will need to be upgraded to the zBX Model 003. This is a no charge upgrade and entitlements on the z114 will be moved to the zBC12 at no charge.

Question:

Can you help me understand the firmware implementation of zBC12?

Answer:

In general, zBC12 and zBX firmware is implemented in the traditional System z manner. What is different is that firmware (for example, the hypervisors) for blades installed in a zBX are supported from the attached zBC12 firmware, to help optimize the integration of the zBX with the zBC12, as well as treating the blade firmware as an "always there" component of the zEnterprise System. Also, there are some changes to licensing of certain firmware in a zEnterprise System. First, firmware for the zBC12 cryptography feature includes technology that requires license terms in addition to the standard license terms governing use of LIC, so an addendum to the LIC license will be included when a zBC12 is configured with a crypto card feature having elliptical curve cryptography capabilities. Second, certain zBX components and features include firmware that is licensed under non-IBM terms (called "separately licensed code" or "SLC"). IBM will deliver the license agreements governing use of SLC along with the IBM license agreement for LIC when a zBX is ordered.

Question:

What is the maximum allowable distance between the zBC12 frame and the zBX frame?

Answer:

The controlling zBC12 must be connected to the zBX with a 26 meter (85') cable. This requirement is for serviceability reasons.

Question:

If I order a zBC12 with the non-raised floor (NRF) option, can I also order a zBX with a NRF option?

Answer:

Yes. Both the zBC12 and zBX support NRF implementations.

Question:

What cooling options are available for the zBC12 and the zBX?

Answer:

The zBC12 is a single frame, air cooled system only. Water cooling is not an option for the zBC12, unlike the zEC12.

For the zBX there is an optional rear door heat exchanger available.

Question:

Is it possible to update driver code on the zBX separately from the zBC12, on a different schedule?

Answer:

No. Once the zBX is installed the firmware is common and updated as a unified zEnterprise system.

Question:

Can I buy a zBC12 that has only IFL or ICF processors without including a general-purpose processor (CP)?

Answer:

Yes. Similar to the System z114, you can order only IFLs or ICFs in a zBC12, using a model capacity identifier of A00 with 1 to 13 IFLs or ICFs.

Question:

Does the zBC12 offer the same hybrid capabilities as the zEC12?

Answer:

Yes. The zBC12 offers identical hybrid computing capabilities to the zEC12. The zBC12 simply represents a new entry point into the hybrid computing model for mid-sized clients.

When we are talking about hybrid computing and the introduction of the zBC12, it's important to realize that the central processing complex can be the zEC12, z196, z114 or now the zBC12. It's the only component that is different or interchangeable based on the size and capacity requirements of the mainframe footprint.

The entire hybrid environment is equally accessible from the zBC12 client and the zEC12. There are no limitations in terms of what you can put in a zBX or how Unified Resource Manager works when you have a zBC12 vs. a zEC12 in the mix. It is not handicapped in any way. It uses the same zBX, the same blades, optimizers and operating environment support, and the same Unified Resource Manager firmware. Everything we will talk about regarding hybrid computing applies equally to a zBC12, zEC12, z114 and z196 environment.

Question:

Is the zBC12 just a low end derivative of the zEC12?

Answer:

The zBC12 offering features many of the leadership capabilities of the zEC12 in a package that is designed for mid-sized businesses, bringing hybrid computing within the reach of a much broader set of clients. The zBC12 is much more than just a derivative offering. While it boasts similar qualities of service, what we built was unique for the mid-sized business – simpler packaging, more flexibility and improved price performance – are just a few of elements of the design.

Question:

Will the zBC12 support z/OS 2 GB pages?

Answer:

The zBC12 and z/OS will offer 2 GB page support designed to reduce memory management overhead and improve overall system performance by enabling middleware to use 2 GB pages. These improvements are expected due to improved effective translation lookaside buffer (TLB) coverage and a reduction in the number of steps the system must perform to translate a 2 GB page virtual address. Exploitation is provided for the IBM 31-bit SDK for z/OS, Java Technology Edition, V7.0.0 (5655-W43) and SDK7 IBM 64-bit SDK for z/OS, Java Technology Edition, V7.0.0 (5655-W44). Also, along with this support, z/OS is designed to make the pageable link pack area (PLPA) and common page data sets optional, used only for quick and warm start IPLs.

zEC12 Hardware

Question:

Tell me about the processor chip used in the zEC12.

Answer:

The zEC12 design is based on IBM's hexa-core 5.5 GHz application-tuned out-of-order superscalar chip. The microprocessor design delivers a boost to performance for all workloads.

The zEC12 will be the first general-purpose large-scale enterprise server with a Transactional Execution Facility, designed to help eliminate software locking overhead that can impact performance.

Transactional Execution will offer increased scalability and parallelism to drive higher transaction throughput. IBM's Java Runtime Environment is planned to exploit the Transactional Execution Facility in an upcoming maintenance roll-up (see Statement of general direction section).

The IBM Enterprise PL/I compiler is planned to exploit the Decimal-Floating-Point Zoned-Conversion Facility for increased efficiency to improve performance.

- Java exploits a significant set of the new instructions available on zEC12 servers. These changes to Java are designed to enhance array bounds checking, and help the system fetch data and instructions operate more efficiently.
- Support announced July 23, 2013 – The zEC12 and z/OS will offer 2 GB page support designed to reduce memory management overhead and improve overall system performance by enabling middleware to use 2 GB pages. These improvements are expected due to improved effective translation lookaside buffer (TLB) coverage and a reduction in the number of steps the system must perform to translate a 2 GB page virtual address. Exploitation is provided for the IBM 31-bit SDK for z/OS, Java Technology Edition, V7.0.0 (5655-W43) and SDK7 IBM 64-bit SDK for z/OS, Java Technology Edition, V7.0.0 (5655-W44). Also, along with this support, z/OS is designed to make the pageable link pack area (PLPA) and common page data sets optional, used only for quick and warm start IPLs.
- Support announced February 24, 2013 - Additional Common Cryptographic Architecture enhancements are made available in support of Die Deutsche Kreditwirtschaft (DK) AES PIN, new message authentication codes (MAC) using the AES CMAC algorithm, and User Defined Extension (UDX) simplification for PKA Key Translate. Delivering on our SOD from July 23, 2013, an enhancement to CFCC Level 19 is available, designed for exploitation of the Flash Express feature to help improve the resiliency of the Coupling Facility with cost-effective standby capacity to handle overflow of WebSphere MQ shared queues.

Question:

What other enhancements were announced for the zEC12 on July 23, 2013?

Answer:

Today we are introducing enhancements to the IBM zEnterprise EC12 (zEC12), the IBM zEnterprise BladeCenter Extension (zBX) Model 003, and the IBM zEnterprise Unified Resource Manager (zManager) to help strengthen the core attributes of zEnterprise - today's modern mainframe.

- Helping to energize your applications.
- zManager enhancements
- CPU management for the IBM BladeCenter HX5 (7873) blade in the zBX Model 003 for increased policy-driven workload management
- Availability management for monitoring and reporting availability of workloads spanning virtual servers in logical partitions, on HX5 blades, and on PS701 blades

- z/OS V2.1 Shared Memory Communications over Remote Direct Memory Access (SMC-R) and a new high-speed networking link, 10GbE RoCE Express, a great combination helping to transparently optimize z/OS server-to-server communications in a multiple-CPC environment, and z/OS LPAR-to-LPAR communications in a single-server environment, with reduced latency and lower CPU overhead than traditional TCP/IP communications.
- Technology refresh for Ethernet environments with OSA-Express5S.
- Helping you save money.
- LPAR absolute hardware capacity setting to limit an LPAR to a specific amount of hardware processor capacity, designed to offer improved granularity and flexibility for non-z/OS systems.
- z/OS V2.1 zEnterprise Data Compression (zEDC) acceleration capability and the zEDC Express feature are designed to support a new data compression function that provides high performance, low-latency compression without significant CPU overhead. Refer to the Planned Availability dates section.
- Increased number of subchannels per FICON® port; now up to 24k from 16k to support more base and alias devices.
- Coupling Thin Interrupts to help improve the efficiency of environments where shared engines are used as Coupling Facilities.
- LPAR absolute hardware capacity setting will allow for a more granular and flexible software pricing for non-z/OS systems.
- Securing it all
- Extended cryptographic algorithms for IBM Enterprise Public-Key Cryptography Standards (PKCS) #11 (EP11) to support higher-quality digital signatures
- New Trusted Key Entry (TKE) workstation with a setup wizard for simplification of startup TKE tasks and a full-function migration wizard for the quick and accurate deployment of Crypto Express4S features configured as EP11 coprocessors.
- EAL5+ Common Criteria certification for the zEC12.

Note: All new zEC12 enhancements for July 23, 2013, will require Driver 15, available on September 20, 2013.

Question:

What is the machine type of the zEC12?

Answer:

The zEC12 machine type is 2827.

Question:

Where did you get the name 'zEnterprise EC12 (zEC12) – did it have a meaning?

Answer:

The 'EC' stands for the fact this is an enterprise class machine. The '12' is that this is the twelfth generation of CMOS technology that started with the 9672 processor line we introduced in 1994.

Question:

Does the zEC12 offer more available subcapacity processors than the IBM zEnterprise 196 (z196)?

Answer:

Yes. The zEC12 can have up to 20 subcapacity Central Processors (CPs) per server while the z196 accommodates up to 15 subcapacity CPs. These subcapacity processors can physically reside in multiple books. These additional subcapacity options provide customers with the flexibility to match their infrastructure to their business volumes at every moment in time, to improve operational efficiency.

Question:

What zEC12 models were announced?

Answer:

IBM announced the following five zEC12 models. Please note that the last two digits of the model number indicate the maximum number of processor units (PUs) available for customer purchase on that model.

- A zEC12 H20 model can be a 1-way through 20-way – which means there are 20 orderable cores contained on one processor “book” (plus many supporting cores in every model, including a minimum of two spare cores).
- A zEC12 H43 model can be a 1-way through 43-way (43 orderable cores) contained on two books.
- A zEC12 H66 model can be a 1-way through 66-way (66 orderable cores) contained on three books.
- A z1EC12 H89 model can be a 1-way through 89-way (89 orderable cores) contained on four books.
- The enhanced capacity zEC12 HA1 model can be a 1-way through 101-way (101 orderable cores) contained on four books.

Customers that reach book limits can easily upgrade from smaller models to larger models (except the HA1) nondisruptively, i.e. without requiring a service interruption of the machine.

The cores can be configured as general purpose processors (CPs), Integrated Facilities for Linux (IFLs), System z Application Assist Processors (zAAPs), System z Integrated Information Processors (zIIPs), additional System Assist Processors (SAPs), Internal Coupling Facilities (ICFs) and/or used as additional spares.

Up to twenty subcapacity CPs can be active on the server regardless of hardware model. In other words, subcapacity configurations are available on any of the models (H20 to HA1) as long as the server is configured (not necessarily the same as purchased) with fifteen or fewer general-purpose processors.

Question:

How much memory is available on the zEC12? What is the RAIM feature you mentioned earlier? How much HSA is there on the zEC12?

Answer:

The zEC12 server supports up to 3 TB of real memory per server and 1.0 TB per LPAR, but the actual maximum physical memory sizes are related to the number of books in the system. The minimum initial amount of memory that can be ordered is 32 GB for all models.

The zEC12 is designed with redundant array of independent memory (RAIM) technology. RAIM is similar to what is known in the disk storage industry as RAID. RAIM technology provides protection for the dynamic random access memory (DRAM), dual inline memory modules (DIMMs), and at the memory channel level, delivers the most resilient memory subsystem to date.

Please note that the 3 TB maximum memory is customer-usable, RAIM-protected memory. RAIM is always active, and IBM has already factored internal RAIM overhead into all advertised memory configuration amounts.

On the zEC12, IBM provides 32 GB of HSA, independent of customer-purchased memory.

Question:

What is different about the Model HA1 compared to H89? What does the 'A' in HA1 stand for?

Answer:

The Model HA1 is an enhanced capacity model which contains a different configuration of MCMs than other models. The zEC12 is fully populated with four high density books and 101 orderable cores. You can configure the HA1 machine to be a 1 to 101-way. Like the other four-book model, the H89, the HA1 can be ordered with a minimum of 32 GB of memory and up to a maximum of 3 TB.

Upgrading from any other model of the zEC12 to a Model HA1 will require a planned outage of that machine, but you can still avoid application service interruptions if you exploit Parallel Sysplex® which can automatically and dynamically shift workloads to another machine.

The 'A' in HA1 represents the number '10'. So there are 101 configurable cores in the Model HA1.

Question:

When I look at the cores and your point about 120 total cores – I get 101 configurable cores (CP, IFL, zIIP, etc.), 2 spares, 16 SAPs. That adds up to 119, not 120. What is the unaccounted-for core?

Answer:

There is one core in all configurations which is "reserved" for the integrated firmware processor (IFP).

Question:

What is the integrated firmware processor (IFP)?

Answer:

The integrated firmware processor (IFP) is allocated from the pool of processors for the whole system. Unlike the characterized processors, you don't pay for the IFP. It is solely used for infrastructure management of the PCIe features –10GbE RoCE Express and zEDC Express. It is not customer usable or 'visible'.

Question:

How many spare processing cores are on the zEC12?

Answer:

IBM ships every zEC12 machine with a minimum of two spare processing cores. These spares can be shared across the books. The zEC12 offers core-level (engine-level) level sparing.

Question:

Can any of these spare processing cores be used for other purposes?

Answer:

No, the minimum allotment of spare cores is exclusively reserved to provide failover in the extremely unlikely event of a processor failure. The spare cores protect all processor types (CPs, IFLs, zAAPs, zIIPs, and ICFs). Any additional spare cores above the minimum allotment can be activated for other purposes.

Question:

How many System Assist Processors (SAPs) are on the zEC12?

Answer:

The answer depends on the model. The standard number of SAPs provided to the customer is as follows:

- The zEC12 Model H20 has four.
- The zEC12 Model H43 has eight.
- The zEC12 Model H66 has twelve.
- The zEC12 Model H89 has sixteen.
- The zEC12 Model HA1 has sixteen.

These are in addition to the "visible" customer configurable cores. For example, an HA1 has 16 SAPs, 2 mandatory minimum spare cores, and up to 101 customer configurable cores.

Additionally, more SAPs may be acquired from the pool of available processing units within the model.

Question:

Is the zEC12 physically bigger than the z196?

Answer:

The zEC12 has the same floor cutouts as the z196. The depth of the covers will increase by 64mm/ 2.52 inches in the front. If you choose optional water cooling or top-exit I/O cabling the clearance is different (4" additional depth for water and 12" width additional for top-exit I/O cabling) from the standard zEC12.

Question:

Can you help me understand the firmware implementation of zEC12?

Answer:

In general, zEC12 and IBM zEnterprise BladeCenter Extension (zBX) firmware is implemented in the traditional IBM System z manner and identical to what was done with the IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114). The firmware (for example, the hypervisors) for blades installed in a zBX are supported from the attached zEC12 firmware, to help optimize the integration of the zBX with the zEC12, as well as treating the blade firmware as an "always there" component of the zEnterprise System. Firmware for the zEC12 cryptography feature includes technology that requires license terms in addition to the standard license terms governing use of LIC, so an addendum to the LIC license will be included when a zEC12 is configured with a crypto card feature having elliptical curve cryptography capabilities. Second, certain zBX components and features include firmware that is licensed under non-IBM terms (called "separately licensed code" or "SLC"). IBM will deliver the license agreements governing use of SLC along with the IBM license agreement for LIC when a zBX is ordered.

Question:

What is the LPAR absolute hardware capacity setting?

Answer:

Customers, particularly our Linux customers where software pricing is based on a "per core" basis, need the ability to limit an LPAR to a specific amount of hardware processor capacity in a way that is unaffected by subsequent physical or logical configuration changes. Introduction of firmware and software appliances creates an even greater need for this capability.

The LPAR absolute hardware capacity setting will allow for specification of an absolute capacity limit for an individual logical partition. This is specified in absolute processor capacity (for example 2.5 processors). Use of this setting should allow more granular and flexible software pricing for non-z/OS systems.

Note: All new zEC12 enhancements for July 23, 2013, will require Driver 15, available on September 20, 2013.

Question:

What is the maximum allowable distance between the zEC12 frame and the zBX Model 003 frame?

Answer:

25 meters. The controlling zEC12 must be connected to the zBX with a 26 meter (85') cable. This requirement is for serviceability reasons.

Question:

What is the radiator-based cooling on the zEC12?

Answer:

The zEC12 is designed with an environmental focus to help improve data center efficiency. It has a radiator-based air-cooled system designed for more efficient cooling and improved maintenance. The system has redundancy of all active components with air backup if there is a rare plumbing problem. A fill and drain tool is required for freeze for install, some radiator service actions and system relocation if freezing temperatures could be encountered.

Question:

What cooling options are available for the zEC12 and the zBX?

Answer:

For the zEC12 there is optional water cooling available. For the zBX there is an optional rear door heat exchanger available.

Question:

Is it possible to update driver code on the zBX separately from the zEC12, on a different schedule?

Answer:

No. Once the zBX is installed the firmware is common and updated as a unified zEnterprise system.

Question:

Can I buy a zEC12 that has only ICF processors without including a general-purpose processor (CP)?

Answer:

Yes. Similar to the z196, you can order only ICFs in a zEC12, using a model capacity identifier of 400 with a maximum of 101 ICFs (this is a change from the z196 where the maximum was 16).

Question:

Can I buy a machine that is made up entirely of Integrated Facility for Linux (IFL) engines and no general-purpose processor (CP)?

Answer:

Yes. Similar to the z196, you can order an IFL only zEC12, using a capacity identifier of 400 with 1 to 101 IFLs. This dedicated IBM System z Linux server is named the IBM Enterprise Linux Server. For more information, check the [ELS home page \(ibm.com/systems/z/os/linux/els.html\)](http://ibm.com/systems/z/os/linux/els.html), or ask your IBM or Business Partner representative for more information.

Question:

Is there still an option for overhead cabling?

Answer:

Yes you can order overhead cabling and you can also order overhead power on the zEC12.

Question:

Tell me about the non-raised floor option on the zEC12?

Answer:

The non-raised floor option allows a zEC12 to be ordered and installed without a raised floor. This feature will be a great option for clients looking for 'cement' floors for disaster recovery centers, backup sites, lower- cost operations facilities. With a non-raised floor zEC12, you must order overhead power, overhead cabling, and radiator-based air-cooling.

Note that a zBX can be installed on a non-raised floor. On the zBX Model 002, and the zBX Model 003 prior to March 2013, the I/O and power cables will exit from the bottom of the machine. In March 2013, IBM announced the ability for top exit I/O and power for the zBX Model 003.

Question:

What about the zEC12 is perfect for consolidation of server farms?

Answer:

Scale is the first big factor making the zEC12 so compelling for consolidating distributed servers. The 5.5 GHz chip means more processing power per core, and there are lots more cores available in each frame – up to 101 “visible” cores per server but many more helping to keep those 101 focused on real work – like the System Assist Processor (SAP) for executing I/O. There is also benefit using internal networks (HiperSockets™) for communications.

The zEnterprise server virtualization capabilities can help to support more virtual Linux servers than any of its competitors, hundreds to thousands of virtual servers in a single footprint with unsurpassed qualities of service. The zEC12 offers benefits such as the concurrent add/delete of I/O, the redundant array of 3 TB of independent memory (RAIM), the multi-chip module design, the HiperSockets network traffic analyzer, and the Crypto Express4S cryptographic enhancements.

Plus the ability to run Linux workloads in an environment designed to run at sustained high processor utilization can help because of the designed off-loading of network and disk I/O traffic through the Server assist processors. There is also benefit to having built in processor failover capabilities in both memory and processor.

The use of integrated blades offers an added dimension for workload consolidation and optimization. And IBM zEnterprise Unified Resource Manager (zManager) governs Linux on IBM System z and blade resources for greater command and control.

Academic Initiative and Skills

Question:

I am an intern and have found out many of the languages used by the mainframe are no longer taught in school. How can a student like myself be sure to get exposure and learn the necessary skills to be ready to work on the mainframe?

Answer:

Welcome to the System z world! Actually I'd suggest checking out our Academic Initiative website. In reference to a similar question from the prior session on how to develop System z skills, there are over 1000 schools that are part of the System z academic initiative. These schools offer a variety of courses and curriculum across all areas of operating systems and application development. For more information see: <http://www.ibm.com/ibm/university/academic/pub/page/systemz>

Question:

What is the IBM global plan for the generation Y education toward the new System z platform?

Answer:

System z has a strong commitment to building skills in our industry. You can look at the Academic Initiative website: www.ibm.com/systems/z/education/academic/index.html.

You will find references to universities we are working with to help develop mainframe skills as well as see a broad spectrum of resources available to students. We also offer a resume site to help match students with prospective clients.

Question:

Is there a way to get hands on experience on the actual operating system outside of an existing IBM client?

Answer:

Yes, there are various options for getting access to mainframe systems. It depends on what system image is required - z/OS, z/VM, DB2, IMS™, etc and the purpose of the access.

IBM has mainframe systems located in universities and also at IBM depending on the purpose of the system access and what you are attempting to do.

The best way to help you is to schedule a call with Don Resnik, IBM System z Client Skills leader and discuss your requirement. Please send an e-mail to resnik@us.ibm.com and suggest a day and time for a call.

Question:

This information is all very good. In my region we are on-boarding new z/VM (Linux on System z) customers but much of our most knowledgeable IT staff has retired. How do we go about acquiring education about how to use z/VM in order to support our customers at an acceptable level?

Answer:

In reference to a similar question from the prior session on how to develop System z skills, there are over 1000 schools that are part of the System z academic initiative. These schools offer a variety of courses and curriculum across all areas of operating systems and application development. For more information see <http://www.ibm.com/ibm/university/academic/pub/page/systemz>

zEDC Express and Compression Acceleration

Question:

What is IBM Enterprise Data Compression (zEDC)?

Answer:

A new capability of z/OS V2.1, IBM zEnterprise Data Compression (zEDC) combined with a new hardware feature of the IBM zEnterprise EC12 (zEC12) or IBM zEnterprise BC12 (zBC12) called zEDC Express, offers a compression acceleration solution designed for high performance, low latency compression with little additional overhead.

Question:

Why is zEDC needed? Isn't there already compression on every System z processor chip?

Answer:

Today, there is explosive growth in the volume of data being generated. Clients must be able to store, process, protect, audit and use the data to create business value. Storing and handling data is costly in both disk space and CPU consumption. Data must be moved across systems and enterprises, where the size of the files directly affects how long it takes to exchange these files.

Data compression has been available since early in the life of the mainframe. Historically, IBM had introduced the S/390 Data Compression, right on the processor chip, which offered one-fifth the number of processor cycles of software-only compression. Implemented with the CMPSC instruction, it is an IBM proprietary compression using a dictionary-based approach. (This type of compression is heavily used by DB2, for example.) DFSMS™ also utilizes hardware compression with generic compression available for VSAM and non-VSAM extended format data sets, as well as tailored compression available for non-VSAM extended format data sets.

However, there is still a massive amount of data that is not being compressed. You need to consider the trade-offs in the processor resources to compress and decompress versus the savings in storage and wall-clock time.

zEDC can be more efficient. Compression is done on the new zEDC Express feature, having little effect on CPU resources. zEDC is compatible with industry standard, open zlib based compression – used today by Java and other applications. Also, zEDC is different – zEDC is dictionary-less. zEDC scans text and looks for the re-use of a phrase and puts in a 'back' reference. For large files, zEDC may provide a compression ratio that is better than existing BSAM/QSAM compression options, reducing I/O requirements which may provide a shorter batch window, while improving CPU time for compression.

EDC is complementary to existing System z compression technology. Smaller records and files best suited for hardware chip compression will still use hardware compression.

Instructions to get best benefit, and larger files that should need a different compression algorithm will be directed to zEDC.

Question:

Is there a difference between compaction and compression? Sometimes the term compaction is used (in DFSMS Data Class definitions) and sometimes the term compression is used.

Answer:

With regard to DFSMS DASD compression, there is no difference. When DFSMS compression for DASD was introduced about 20 years ago the COMPACTION option already existed to request compression for tape. Therefore, it made sense to enhance the existing compression-related option to include the new values for generic and tailored compression (which were only valid with DSNTYPE=EXT). However, since no existing COMPACTION parameter existed in IGDSMSxx member of SYS1.PARMLIB, DFSMS introduced a new COMPRESS keyword to allow users to request generic and tailored compression.

Question:

How does zEDC work with zlib, and will that help with the data transfer requirements for my enterprise?

Answer:

z/OS V2.1 provides the zlib library which supports the sending of compression and decompression requests to the zEDC Express. The z/OS-provided zlib library is provided as a UNIX archive file that can be statically linked into IBM, ISV, or customer applications that currently use zlib, enabling additional exploitation of compression through zEDC Express and expanding potential compression opportunities. This makes zEDC a great fit for interaction with services that are implemented on a variety of architectures and hosted on heterogeneous architectures. When zlib uses zEDC, there may be up to 118X reduction in CPU and up to 24X throughput improvement.¹

Some examples include:

You could compress a file on your z/OS v2.1 system using zEDC (such as using PKZIP), then move that file to a Linux for System z or Windows system, and the file could be decompressed on that system.

An FTP scenario from z/OS to a non-z/OS system could look like

- Compress a file with PKZIP, IBM Encryption Facility for z/OS, or IBM 31-bit and 64-bit SDK71 for z/OS Java Version 7 Release 1
- FTP the file to the other system
- Decompress the file on that other system

Similarly, a file could be compressed on a Windows system using an application such as PKZIP, and sent to z/OS where the decompression could be accelerated with zEDC.

Question:

And, what about the compression I have in my DS8000® disk flash control unit; will there be an impact to this if I implement zEDC?

Answer:

zEDC disk compression extends the value of your DS8000 disk flash purchase by using less flash to hold the same amount of data. zEDC will allow you to more fully utilize the flash space.

Question:

Can you give me examples of what functions and applications exploit zEDC?

Answer:

There are already quite a few exploitation use cases for zEDC, and the list will continue to grow. Exploitation includes:

¹ Based on projections and measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

- z/OS V2.1
 - SMF logger
 - DFSMS BSAM/QSAM extended format data sets
 - IBM 31-bit and 64-bit SDK71 for z/OS Java Version 7 Release 1 and higher, IBM 31-bit and 64-bit SDK7 for z/OS SR7 and higher
 - DFSMSdss™ and DFSMShsm™ plans to exploit zEDC by the end of the 3Q14
- IBM Encryption Facility for z/OS V1.2
- IBM Sterling Connect:Direct for z/OS V5.2
- IBM Security zSecure™ V2.1 (via BSAM/QSAM)
- IBM WebSphere MQ for z/OS V8 (COMPMSG(ZLIBFAST))

Many ISV applications are already exploiting zEDC or have plans to exploit. A few of the ISVs who have exploited zEDC or have plans to are: Alebra, ASE, PKWARE, and Software AG. Contact your product representatives for more information about their plans.

Question:

How does zEDC improve SMF logstream recording?

Answer:

zEDC with z/OS SMF Logstream recording and zEDC with QSAM/BSAM alleviate SMF constraints across the entire life cycle of a record. zEDC can help to improve SMF logger compression with less system overhead and storage usage. Programs that read SMF data from a zEDC compressed QSAM/BSAM data set can also see a reduction in elapsed time. In fact, our internal measurements have shown that using zEDC, up to 4X less data is stored in System Logger and reducing Logger CPU usage by up to 30%.² Corresponding support in the SMF dump program IFASMF DL is designed to support both hardware-based and software-based decompression, and decompression; support is available on z/OS V1.12 and z/OS V1.13 (5694-A01) with the PTF for APAR OA41156. When running the SMF dump job against an SMF logstream with compressed records on a system with zEDC, the elapsed time for the dump job may be reduced by up to 15% and the CPU time for the job may be reduced by up to 40%.¹ Note that this new function is expected to allow higher write rates for SMF data when hardware compression is enabled. RMF™ support for hardware compression includes SMF Type 74 subtype 9 records and a new Monitor I PCIE Activity report intended to provide information about compression activity on the system.

Question:

How does QSAM/BSAM benefit from zEDC?

Answer:

In general, z/OS access methods compress data by policy. zEDC accelerated compression will improve disk compression (requiring less space) for sequential data and reduce the CPU resources needed. QSAM/BSAM can save disk cost and in some cases shorten elapsed time, reducing batch windows. BSAM/QSAM extended format data, can realize benefits from zEDC to help reduce disk space up to 75% and improve effective bandwidth without significant CPU overhead.¹ Some will see up to 80% reduced CPU cost compared to tailored and generic compression options.¹

zEDC can be used by DFSMSdss/DFSMShsm to deliver efficient compression when backing up / restoring data; this support became available in 3Q/2014.

² Based on projections and measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

Question:

How does IBM Java for z/OS exploit zEDC?

Answer:

The IBM 31-bit and 64-bit SDK7 for z/OS Java Technology Edition, Version 7 Release 1 (5655-W43 and 5655-W44) and (IBM SDK7 for z/OS Java Technology Edition, Version 7 SR7 and higher) provides exploitation of the zEDC Express feature; z/OS Java can transparently exploit zEDC through standard APIs, such as java/util/zip, where buffer sizes greater than 4K will reflect greater benefit for CPU utilization and elapsed time. This can benefit use cases such as application business data exchange, HTTP responses for web services, and large object storage.

A zEDC java.util.zip.Deflater in memory test improved elapsed time up to 55x and CPU time up to 240x when compared to zlib software compression.¹ And, using java.util.zip.GZIPOutputStream class, there was up to 90% reduction in CPU time using zEDC hardware versus zlib software and up to 74% reduction in Elapsed time using zEDC hardware versus zlib software.³

Question:

How does IBM Encryption Facility for z/OS benefit from zEDC?

Answer:

The IBM Encryption Facility for z/OS, 5655-P97, with the PTF for APAR OA43869, exploits zEnterprise Data Compression (zEDC) for z/OS V2.1, running on zEC12 and zBC12 servers with the zEDC Express adapter with the appropriate levels of IBM SDK Java for z/OS Version 7 and Version 1 Release 1 (5655-W43 and 5655-W44). zEDC can provide IBM EF users reductions of up to 60% in elapsed time and up to 70% in CPU time for environments where compression is already in use.¹ For IBM EF users not already using compression, compression with zEDC can provide IBM EF users a reduction of up to 44% in elapsed time and up to 46% in CPU times.²

This will compliment the software compression support that exists today with Encryption Facility OpenPGP support.

Question:

How does IBM Sterling Connect:Direct for z/OS V5.2 benefit from zEDC?

Answer:

The IBM Sterling Connect:Direct for z/OS V5.2 uses zEDC to provide high performance data compression and decompression for managed file transfer. The file transfers can be z/OS to z/OS, or z/OS to distributed systems (UNIX, Linux, Linux on System z, and Windows). zEDC is fully compatible with the zlib compression used today. When IBM Sterling Connect:Direct for z/OS V5.2 uses zEDC, it may provide up to 80% reduction in elapsed time to transfer a file from z/OS to z/OS (sender TCB time and receiver TCB time), and with significant improvement in CPU time compared to the zlib software compression.²

If you are exchanging files with a distributed system, that system uses the software zlib capability to compress and decompress files, as we said, zEDC's zlib libraries are compatible with the industry standard. On z/OS systems with zEDC, zEDC would be used to perform the compress and decompress operations.

³ Based on projections and measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

Question:

How does IBM Security zSecure V2.1 benefit from zEDC?

Answer:

The audit capabilities of IBM zSecure V2.1 are improved as the Access Monitor data (data about actual usage of resource profiles and to remove unused profiles and unused authorizations in with the Access Monitor) and the CKFREEZE files can benefit with zEDC QSAM/BSAM data set compression. zEDC can reduce the output size of Access Monitor and UNLOAD files by up to 10X and CKFREEZE files reduced by up to 4X. ²

Question:

How does IBM WebSphere MQ for z/OS V8 benefit from zEDC?

Answer:

WebSphere MQSeries has always provided compression options for message data passed over MQ channels via the COMPMSG attribute.

The existing zlib options are:

- ZLIBFAST - Message data compression is performed using the zlib compression technique. A fast compression time is preferred.
- ZLIBHIGH - Message data compression is performed using the zlib compression technique. A high level of compression is preferred.

Starting with WebSphere MQ for z/OS V8 the COMPMSG (ZLIBFAST) attribute will now use zEDC when available to perform compression and decompression of message data.

Question:

If I have software compression and zEDC will my files be compressed twice?

Answer:

The ability to compress data is usually provided via a specific product, service or client application. If data flows between such software, it is possible that data is compressed twice without knowledge of the zEDC Express or existing compression technology. Any interactions should be examined to determine the correct location and technology to be used for the compression. When files are compressed a second time, there may be some growth in the size of the file, reflecting header information associated with the zlib compression format.

Question:

I am currently using DB2 SMF record compression. How does this interact with zEDC enabled SMF logstreams?

Answer:

The DB2 SMF compression ZPARM option enabled Run Length Encoding compression. This compression style reduces data size by elimination of repeated characters. The zEDC compression works across several SMF records and provides more advanced compression and can be used in coordination with the DB2 support.

Question:

What are the software and hardware requirements for z/EDC?

Answer:

Software requirements:

- Requires z/OS 2.1 (w/PTFs) and the new zEDC Express for z/OS feature
- z/OS V1.13 and V1.12 offer software decompression support only

Hardware requirements:

- zEC12 or zBC12
 - New zEDC Express feature for PCIe I/O drawer (FC#0420)
 - Each feature can be shared across up to 15 LPARs
 - Up to 8 features available on zEC12 or zBC12
 - The recommended high availability configuration per server is four features
 - This provides up to 4 GB/s of compression/decompression
 - Provides high availability during concurrent update (half devices unavailable during update)
 - Recommended minimum configuration per server is two features
 - To install zEDC Express in an existing zEC12/zBC12
 - Apply z/OS Service; Hot plug a zEDC Express adapter; update your IODF, and Dynamic Activate
- You should make sure that you have the latest software and hardware service installed.
-

Question:

How is the zEDC Express feature defined to the zEC12 or zBC12?

Answer:

The zEDC Express feature, along with the 10GbE RoCE Express feature, are not defined as a CHPID and do not consume a CHPID number.

They will be defined as part of an I/O configuration using the Hardware Configuration Definition (HCD) program. Each physical device has a PCHID (Physical Channel ID) which represents where the PCIe device is installed. Each PCHID is mapped to up to fifteen FIDs (Function ID) which each represent a Virtual Function. Each Virtual Function is assigned to a single z/OS LPAR.

Alternatively, zEDC can also be defined via an I/O Configuration Program (IOCP). HCD will likely be used to generate the IOCP deck in most cases.

Question:

So, I need to order two features both hardware and software? That is something new.

Answer:

To use zEDC, you need to order both the zEC12 or zBC12 H/W feature and the z/OS V2.1 priced feature. This is a new delivery model, that is designed to help with the overall TCA and designed to scale as you expand your use of zEDC. eConfig will notify you that there is a required S/W feature at the time you configure the zEDC Express feature for PCIe I/O drawer (FC#0420).

Question:

How do I order the zEDC Express for z/OS feature?

Answer:

zEDC Express for z/OS is similar to other z/OS priced features from an ordering and delivery perspective. It is a monthly license charge (MLC) feature similar to SDSF. You would license zEDC Express for z/OS to those servers that will be exploiting zEDC (those servers should have the zEDC Express feature (FC#420). By licensing the feature at the time you plan to start using zEDC will align the start of the MLC billing with your usage. Contact TechLine or your MLC pricing representative for details.

Question:

Can the zEDC Express feature be shared between multiple LPARs?

Answer:

Yes. The zEDC Express feature can be shared across up to 15 LPARs.

Question:

What are the considerations for data sharing if all my systems are not using zEDC, such as if I have z/OS V1.13 systems or z196 systems?

Answer:

It is highly recommended that you have both the zEDC Express feature and z/OS V2.1 on each 'end' of a connection, including backup systems. In an emergency, if you need to send a file to a server that doesn't have the feature, (such as a z196) or access the file from a release prior to z/OS V2.1, z/OS V1.12 and z/OS 1.13 will have toleration support. This means they'll be able to decompress data, but it will be very CPU intensive. For more information about the z/OS implementation of zlib, including implications for systems on toleration releases, see z/OS MVS™ Programming: Callable Services for High-Level Languages. You are also able to receive files from other platforms that have been compressed with the industry-standard zlib library and use zEDC to decompress those files.

Question:

How many zEDC Express features will I need to have on my system?

Answer:

The minimum purchase is one. IBM recommends at least two zEDC Express feature for high availability. Four zEDC Express features are highly recommended to aid with normal maintenance activities on the server. As the cards are installed in pairs, four cards would ensure that there is not a single point of failure when maintenance is performed or a potential problem with a card is encountered.

Question:

How do I set up my z/OS system to support compression?

Answer:

If you are running z/OS V2.1 (or higher) and have configured the hardware, then the process is similar to what you might be doing today to enable z/OS priced features. The first step is to notify IBM that you are starting to use the feature. Secondly, you update the IFAPRDxx PARMLIB member to specify that the z/OS zEDC software feature is ENABLED. Lastly, you follow the z/OS V2.1 documentation to customize the exploiting functions – SMF, BSAM/QSAM.

For SMF, you specify COMPRESS on the LSNAME parameter in your SMFPRMxx PARMLIB member. For BSAM/QSAM you set up a policy (DATACLASS or via the COMPRESS parameter in your IGDSMSxx PARMLIB member) to set up compression. There are no changes needed to the access method.

Question:

Is there a way to force the dump job (IFASMF DL) to fail if a compressed logstream is accessed from an LPAR that isn't set up to use zEDC Express?

Answer:

Yes. Do not specify SOFTINFLATE and the job will fail.

Question:

I've re-linked my application with the z/OS V2.1 provided zlib library; however, I do not appear to be using the zEDC Express.

Answer:

The first thing to do is to verify that the z/OS LPAR has zEDC Express devices installed. Once it is confirmed that there are zEDC Express devices assigned to this partition check the following:

- That the IFAPRDxx parmlib member has been updated to reflect the enablement of the zEDC Software License
 - That the application Address Space has UACC(READ) access to FPZ.ACCELERATOR.COMPRESSION.
 - Check that the application is providing adequately sized input buffers.
-

Question:

What planning tools are there? How can I determine if I will see benefit from using zEDC? And if so, how many cards will I need?

Answer:

The IBM Washington Systems Center (WSC) has created a sizing tool for use with zEDC, the IBM System z Batch Network Analyzer (zBNA). zBNA is a free, as-is tool that analyzes batch windows using SMF records to help determine if you have files that are candidates for zEDC. It can also help estimate the number of H/W features you will need. It is available from:

IBMers:

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5126>

Business Partners:

https://www.ibm.com/partnerworld/wps/servlet/mem/ContentHandler/tech_PRS5133

Customers:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132>

Question:

What data is provided to help me understand my system's use of zEDC?

Answer:

z/OS offers enhanced RMF™ reporting to report on vital statistics relevant for compression. The SMF 74 SubType 9 record has been updated with new PCIe and zEDC statistics. RMF Monitor III can be used to post-process these records and obtain statistics

Question:

Is zEDC supported on Linux for System z?

Answer:

Linux for System z does not currently support zEDC.

Question:

What did IBM announce on February 24, 2014 regarding z/VM and zEDC?

Answer:

Guest exploitation support for the zEDC Express feature (#0420) on the zEC12 and zBC12 servers, designed for compression acceleration for resource optimization. PTFs became available June 27, 2014.

Question:

What are the benefits of zEDC support for z/VM customers?

Answer:

zEDC is based on new PCIe technology that provides high-speed low-latency connectivity. z/VM support enables the hardware feature to be exploited by guests such as z/OS that have PCIe drivers.

Question:

Are any z/OS APARs required to exploit zEDC as a guest of z/VM?

Answer:

Yes, the PTF for APARs OA44482 and OA43256 are required when exploiting zEDC as a z/VM guest.

10GbE RoCE Express and SMC-R

Question:

How does the IBM announcement on February 24, 2014 regarding z/VM benefit customers who also have the 10GbE RoCE Express feature?

Answer:

z/VM 6.3 has been enhanced to provide guest exploitation of the 10GbE RoCE Express feature (#0411) on the zEC12 and zBC12 servers to allow guests to utilize Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) for optimized networking.

Specifically, any operating system that can exploit the RDMA capability of the 10GbE RoCE Express feature can now do so while running as a guest operating system under z/VM. For example, this new support allows z/OS 2.1 to be tested with RoCE on a z/VM platform in order to fully understand the use and capabilities of this feature before moving to a production environment. Customers who initially test z/OS as a guest under z/VM, and/or those that run z/OS 2.1 as an image on z/VM, can now leverage the 10GbE RoCE Express capability.

The z/OS PTF for APAR OA44482 is required when exploiting RoCE as a z/VM guest.

Question:

What is RDMA?

Answer:

Remote Direct Memory Access (RDMA) is a protocol defined by the Infiniband® Trade Association (IBTA). RDMA allows a host to write or read memory from a remote host without involving the remote host's CPU and Operating System (OS). It requires a specialized network adapter to transfer data as it bypasses operating system layers and many communications protocol layers that are otherwise required for communication between applications. RDMA thus reduces software overhead, providing for high-throughput, low-latency networking.

Question:

What is RoCE?

Answer:

RDMA over Converged Ethernet (RoCE) "is a standard that defines a new RDMA protocol over Ethernet. With advances in data center convergence over reliable Ethernet and Data Center Bridging (DCB), RoCE uses the efficient RDMA mechanism to provide lower CPU overhead and increase mainstream data center application performance at 10GigE and 40GigE link speeds and beyond." ⁴

Question:

I've heard of RoCE before. Is the one in the market today the same for zEC12 and zBC12?

Answer:

The 10GbE RoCE Express feature and its ports are equivalent to the RDMA network adapter, however the only System z exploitation of RoCE is by z/OS V2R1 using SMC-R communication protocol.

⁴ Quoted from Page 1 of "RDMA over Converged Ethernet (RoCE) Fact Sheet" at <https://cw.infinibandta.org/document/dl/7261>

Question:

What is the new Shared Memory Communications RDMA (SMC-R)?

Answer:

SMC-R is a new capability on z/OS V2R1, which works in conjunction with the feature 10GbE RoCE Express, a new PCIe feature on the zEC12 or zBC12, and RoCE. SMC-R is a socket over RDMA communication protocol that optimizes server to server networking for z/OS users.

If you have a z/OS environment that that uses TCP/IP and has applications, such as CICS® or DB2 talking to a database like DB2, in another LPAR or z/OS system, then once you've established the connection, SMC-R provides host-to-host direct memory access for communicating without the traditional TCP/IP processing overhead. This configuration can help to reduce both network latency and CPU resource consumption over traditional TCP/IP communications.

Best of all, SMC-R allows any existing sockets applications that exploit TCP to transparently benefit from RDMA for exchanging data over a RoCE network (no application change). Plus you can benefit from the RDMA InfiniBand technology by leveraging your existing 10GbE Ethernet infrastructure. And SMC-R provides for critical qualities of services that are required by your enterprise environment such as high availability and load balancing (when redundant paths network hardware paths are available), preservation of existing IP topology and IP security, minimal administrative and operational changes, dynamic discovery and setup of other RDMA connections over RoCE fabrics.

Question:

What is the value I will get from SMC-R and the 10GbE RoCE Express feature?

Answer:

The value of RDMA technology is about providing significantly improved network performance for server to server communications. Network improvement attributes are typically described as latency, throughput, CPU cost and scalability. Improvements in network performance can potentially improve (increase) application workload transaction rates while reducing your CPU cost.

The combination of the 10GbE RoCE Express feature and the SMC-R protocol provide for significant improvements in network performance. The network latency characteristics provided by this combined solution are compelling: network latency is typically expressed as “network round trip time.” This latency attribute can translate to an improved overall application transaction rate for z/OS to z/OS workloads. Workloads that are network intensive and transaction oriented (sometimes described as “request/response” workloads – that require multiple and even hundreds of network (“client/server”) flows to complete a single transaction will realize the most benefit. An example of such a workload is a WebSphere application server communicating with a remote database server such as DB2. Another is CICS to CICS communications with IPIC. IBM internal benchmarking has shown up to 48% reduction in response time and up to 10% CPU savings for CICS transactions using DPL (Distributed Program Link) to invoke programs in remote CICS regions in another z/OS system via CICS IP interconnectivity (IPIC) when using SMC-R vs. standard TCP/IP.⁵

Message sizes also impact the performance characteristics of a workload. Transactional workloads that exchange larger messages (e.g. web services such as WAS to DB2 or CICS) will see even greater benefit. IBM has seen a 40% reduction in overall transaction response time for WebSphere Application Server v8.2 Liberty profile TradeLite workload accessing z/OS DB2 in another system measured in internal benchmarks.⁶ Another benchmark we have done shows WebSphere MQSeries for z/OS realizes up to 3X increase in messages per second it can deliver across z/OS systems when using

⁵ Based on internal IBM benchmarks using a modeled CICS workload driving a CICS transaction that performs 5 DPL calls to a CICS region on a remote z/OS system, using 32K input/output containers. Response times and CPU savings measured on z/OS system initiating the DPL calls. The actual response times and CPU savings any user will experience will vary.

⁶ Based on projections and measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

SMC-R vs. standard TCP/IP.⁷

Finally, streaming (or bulk) application workloads (e.g. FTP) will also experience reduction in CPU consumption. Our internal measurements show up to 50% CPU savings for FTP binary file transfers across z/OS systems when using SMC-R vs. standard TCP/IP.⁸

Remember that ultimately the actual performance benefits that can be achieved by each user are very unique to each application workload and other unique system environmental factors (CPU and memory utilization, network bandwidth, network path and congestion, etc.).

Question:

What is the zEC12 and zBC12 10GbE RoCE Express feature?

Answer:

The 10GbE RoCE Express feature is a new PCIe based network adapter on System z that will provide access to RDMA technology such as RoCE and low latency and high network bandwidth with very efficient host CPU utilization.

Question:

How many ports are on the 10GbE RoCE Express feature?

Answer:

Two 10 GbE SFP+ ports are provided on each 10GbE RoCE Express feature, but z/OS SMC-R exploitation uses only one of the two ports (whichever one is configured first in the z/OS TCP profile). The ports are numbered “1” and “2”.

Question:

Does the SMC-R exploitation of RoCE require special Ethernet (RoCE-capable) switches?

Answer:

No, SMC-R does not require a RoCE-capable switch. SMC-R requires a standard 10GbE switch⁹. The Global pause frame (a standard Ethernet switch feature for Ethernet flow control described in the IEEE 802.3x standard) should be enabled on the switch.¹⁰

Question:

How far apart can different servers be to communicate with each other using [SMC-R/RoCE]?

Answer:

The distance from the 10GbE RoCE Express port to the 10GbE switch is limited to 300 meters with OM3 fiber cable. The latency advantages of RDMA are diminished when travelling long distances, and so RDMA performs best when used within datacenter distances. RoCE Express features connected to a single 10GbE switch are preferable. Therefore the distance between two different servers would not exceed 600 meters with a switch in the middle.

⁷ Based on internal IBM benchmarks using a modeled WebSphere MQ for z/OS workload driving non-persistent messages across z/OS systems in a request/response pattern. The benchmarks included various data sizes and number of channel pairs. The actual throughput and CPU savings users will experience may vary based on the user workload and configuration.

⁸ Based on internal IBM benchmarks in a controlled environment using z/OS V2R1 Communications Server FTP client and FTP server, transferring a 1.2GB binary file using SMC-R (10GbE RoCE Express feature) vs standard TCP/IP (10GbE OSA Express4 feature). The actual CPU savings any user will experience may vary.

⁹ SMC-R can also be exploited using direct RoCE Express to RoCE Express connectivity (without any switch). However, this type of direct physical connectivity forms a single physical point to point connection disallowing any other connectivity with other LPARs (e.g. any other SMC-R peers). While this is a viable option for test scenarios it is not practical (not recommended) for production deployment.

¹⁰ The 802.3x “Pause” feature is implemented differently with different switch vendors. Some implement the feature on a per-port basis; others on a global basis. Most switches have “Global Pause” enabled but you will want to verify with the switch vendor that this is the case on the switch that is part of the RoCE fabric.

Question:

I have two datacenters that are approximately 10km apart. Can we connect the zEC12s (one at each location) with RoCE if we use DWDM between the two datacenters?

Answer:

This is currently not a supported implementation but we understand the requirement.

Resource Link® contains all the latest qualification letters with the supported distances depending on your DWDM

vendor: <https://www.ibm.com/servers/resourcelink/lib03020.nsf/pages/systemzQualifiedWdmProductsForGdpSolutions?OpenDocument&pathID>

Question:

So will I still need OSA once I exploit SMC-R using the 10 GbE RoCE Express feature?

Answer:

Yes, the SMC-R protocol leverages your existing IP topology and the TCP/IP connection to control your SMC-R connectivity. You will still need to establish a standard TCP/IP connection over a QDIO OSA port and an Ethernet network to the peer. SMC-R uses this TCP connection to determine eligibility to exploit the RoCE fabric and then to build a point-to-point SMC-R link for the data flow over the RDMA connection path.

Question:

Will SMC-R and RoCE support SNA?

Answer:

No. SMC-R requires TCP.

Question:

Will SMC-R and RoCE support SNA with Enterprise Extender?

Answer:

No. SMC-R requires TCP, and Enterprise Extender relies on UDP.

Question:

What is the feature code for 10GbE RoCE Express?

Answer:

The 10GbE RoCE Express is FC# 0411 10 GbE SR.

Question:

Can the 10GbE RoCE Express feature be shared between multiple LPARs?

Answer:

No. Each 10GbE RoCE Express feature can only be used by a single LPAR.

Question:

How many 10GbE RoCE Express features will I need to have on my system?

Answer:

Each LPAR should be provisioned with 2 features (PFIDs) for redundancy. Up to 16 features are supported on the zEC12 or zBC12.

Question:

What operating systems support this new technology?

Answer:

z/OS V2R1, contains support for the SMC-R protocol, and is the only operating system that is announcing exploitation of the 10GbE RoCE Express feature on the zEC12 or zBC12.

Question:

What changes do I need to make to the applications to take advantage of RoCE and SMC-R?

Answer:

There are no changes required. The application is not involved in the decision to use SMC-R or not. This is all handled inside z/OS Communications Server code.

Question:

Can I enable SMC-R but exclude some applications from using it?

Answer:

Yes, for applications that act as the TCP server you can configure NOSMCR on the port definition.

Question:

Can I communicate memory to memory between my distributed system and my System z – such as a SAP application server and the z/OS DB2 database?

Answer:

No. The only operating system that has announced support for the SMC-R protocol is z/OS. So the offering is only for z/OS to z/OS communication.

Question:

I see that Sysplex Distributor supports SMC-R with RoCE. How does it work?

Answer:

If a z/OS Sysplex Distributor (SD) client has a RoCE connection to the selected z/OS Sysplex Distributor target, and VIPARoute has been enabled, then the connection is eligible to use SMC-R. In this case both outbound and inbound traffic will flow directly to and from the target and client over the RoCE connection. Thus even the inbound flow, which normally traverses the SD node, bypasses the SD route to provide better throughput directly to the target.

Question:

What security is there for the 10GbE RoCE Express feature? How do I know that data traveling across the RDMA fabric is protected?

Answer:

SMC-R preserves almost all the existing network security, IP topology, and network administrative and operational models available in z/OS Communications Server. For example, the data that travels across the RoCE fabric can be protected with encryption, data integrity controls, authentication, access controls and so on.

Question:

What differences exist between a HiperSockets network fabric and a RoCE communication fabric?

Answer:

HiperSockets represents a virtual connection technology that requires neither a network adapter card nor special cabling. RoCE requires a hardware slot in the PCIe drawer and imposes cabling and cooling requirements.

Both are essentially memory-to-memory operations, but each uses memory differently.

HiperSockets utilizes CPU and communication protocol layers of TCP/IP while transmitting traffic. Offload to a zIIP processor alleviates the cost of this CPU utilization. RoCE minimizes the use of CPU and TCP/IP communication protocol layers for the data traffic.

HiperSockets represents a System z proprietary link type; RoCE is a standardized link type; SMC-R is a draft RFC within the IETF database. Both RoCE and SMC-R could thus be applicable to platforms other than zEnterprise.

HiperSockets may be used between z/OS, Linux on System z, z/VM, and z/VSE, including as part of a DynamicXCF network on z/OS. RoCE and SMC-R may only be used between z/OS partners. HiperSockets supports all IP traffic between partners. RoCE and SMC-R only supports TCP traffic.

HiperSockets relies on standard TCP/IP routing and therefore automatic backup to OSA relies on TCP/IP dynamic routing. RoCE and SMC-R will automatically fall-back to OSA in OSD-mode when the RoCE and SMC-R path is not eligible for use.

Note:

All new zEC12 enhancements for July 23, 2013, will require Driver 15, available on September 20, 2013.

Transactional Execution

Question:

Can you tell me a little more about Transactional Execution?

Answer:

Transactional Execution is a computer science innovation that enables a sequence of computer instructions to update data in memory in an exclusive manner while running in a multi-threading operating environment. It ensures that when a given thread tries to update data in memory and other independent threads coincidentally update the target memory asynchronously, that the given thread is signaled of this condition. The concept has been around for almost 20 years and the first hardware implementation was prototyped in IBM's Blue Gene[®]/Q processors that have been designed to power the 20 petaflops Sequoia supercomputer being built for Lawrence Livermore National Labs (e.g. see <http://arstechnica.com/gadgets/2011/08/ibms-new-transactional-memory-make-or-break-time-for-multithreadedrevolution/>).

Transactional Execution does not mean the memory itself is 'transactional' but that there is a new hardware way of accessing memory that enables the memory to be used for transactional updates.

Question:

Can you refresh me on the programming use of 'locks'?

Answer:

If a computer program is operating on data in memory and that data gets updated by another independent program that's executing simultaneously, the data becomes corrupted from the point of view of the first program and is no longer accurate for the first program's subsequent computation. Data integrity has been compromised. Obviously if computer programs cannot rely on the integrity of the data they are operating on, they would be unreliable and useless.

The traditional way to prevent this from happening is the use of 'locks' which are programming concepts that protect the data being updated. All programs that want to access the same data in memory must get permission to access it through a lock that is either 'available' or 'unavailable'. When the lock is 'available' the requesting program gets temporary ownership of the lock, making it immediately unavailable for other requesting programs (on a first-come-first-served basis). Thus it is essential that the owning program 'releases' the lock as soon as it has successfully updated the data.

When the lock is 'unavailable' the requesting program must wait until the lock becomes available. In effect, multi-threaded processing becomes single-threaded as a common memory resource is exclusively accessed by an updating program. This system of holding locks and releasing them is ubiquitous in all computer systems from the biggest transaction processors to the smallest mobile devices and is usually implemented in software (operating systems and middleware).

Multi-threaded software that correctly implements memory updates without loss of data integrity is sometimes described as 'atomic' – the concept meaning that either data is completely and correctly updated, or it is not updated at all (the analog being that atoms are the smallest units and cannot be (easily) split further). All IBM software, including CICS, IMS[™], WAS, WebSphere MQ, TPF, z/OS, z/VM and z/VSE, all use this technique to ensure our customers' data never gets corrupted. The operating system itself may need to update fundamental data structures (chains of control blocks, queues, lists, tables, vectors ...) very frequently and always with atomicity (the penalty for loss of atomicity would likely be an operating system crash) so this sort of processing is critical.

Question:

So how does Transactional Execution work?

Answer:

Transactional Execution provides an alternative approach to using locks, and should, in time, allow the elimination (or elision) of locks. This is especially important because computer systems are becoming increasingly multi-threaded and thus spending more time processing locks is inefficient. Transactional execution is optimized to assume no inter-thread corruption of data in which case there is very little overhead, no locking is required and the transactional sequence completes with a 'zero condition code'. Eliminating the locking/unlocking saves about 160 instructions for every transactional update. So if most of the time the condition code is zero, there is potential to save a significant number of instructions throughout the stack and during the peaks of operational capacity (hence saving money).

What happens if the programming code has an error during transactional execution?

If the condition code returned is non-zero, this signals that another thread has 'trampled on the data' and the logic would need to be retried.

Question:

What software will be taking advantage of Transactional Execution?

Answer:

IBM's Java Runtime Environment is the first exploiter of Transactional Execution Facility. First exploiter of this was Java where we saw up to 2X improved scalability for Java multi threaded workloads. We expect in the future to have exploitation by the z/OS operating system and other compilers.

IBM zAware

Question:

Why did IBM introduce IBM zAware?

Answer:

Today's complex data centers often require a team of experts to monitor systems for abnormal behavior, and to diagnose anomalies before they cause visible impact to service levels. These tasks are difficult for many reasons, including the amount of change, complexity and volume of messages to examine. When an error occurs, it is often complex and reflective of multiple application changes, system configuration settings, patches, and so on. IBM zAware provides a solution for detecting and diagnosing anomalies in z/OS system's that help IT staff more quickly understand unusual data center situations, and diagnose anomalies so that they can improve problem determination.

Question:

Is IBM zAware an analytics application?

Answer:

IBM zAware consists of an integrated set of analytic applications, packaged as an integrated firmware stack used to analyze z/OS OPERLOG messages. As firmware, it is ordered as a feature of the zEC12 or zBC12.

IBM zAware uses a self learning approach to model normal system behavior based on 90 days of baseline system data, using pattern recognition to identify unexpected messages from monitored systems. IBM zAware provides near real-time detection of anomalies that can then be viewed and further diagnosed, such as:

- Rare events leading to problems
 - Problems due to incorrect system or application configuration settings, modifications or fixes
 - Complex incidents involving several different subsystems
 - The triggers and events leading up to a problem
-

Question:

What do you mean by 'IBM zAware monitored clients'? Can I receive OPERLOG messages from other machines in my data center?

Answer:

Good questions. Let's regroup on some of the terms. IBM zAware runs in a partition, which must be on a zEC12 or zBC12. This partition is called the **IBM zAware host partition**. The **IBM zAware host system** is the zEC12 or zBC12 that hosts the IBM zAware host partition. There may be instances that you have multiple IBM zAware host partitions on one zEC12 or zBC12.

IBM zAware can accept messages to be analyzed from other System z servers as long as they are running the supported z/OS operating system (z/OS V1.13 and required exploitation software or z/OS V2.1). The z/OS partition that sends the messages to be analyzed by IBM zAware is called the **IBM zAware monitored client**. Depending on the system configuration, there may be multiple IBM zAware monitored clients on a server.

The **IBM zAware environment** is the collection of the IBM zAware host system and the IBM zAware monitored clients that are sending information to the IBM zAware host system.

Question:

How is IBM zAware different than other tools?

Answer:

Existing problem management tools help determine whether certain specific events are related to system behavior using hard-coded predefined rules related to a known, defined set of conditions. IBM zAware is unique in its ability to detect events automatically using a self learning, self-tuning approach without requiring *any* codification of rules. As a result, IBM zAware can automatically and proactively identify many different types of anomalies and unusual events in your system.

Question:

How is problem determination easier with IBM zAware?

Answer:

Examination of the sheer volume of messages in an operations log (OPERLOG) can make the analysis task daunting — for larger clients over 25 million messages daily is not uncommon. There are over 40,000 unique message IDs defined for z/OS and the IBM software that runs on z/OS systems which makes analysis particularly complex. It is virtually impossible to identify problem messages given this volume without use of a tool. Also IBM zAware can recognize problems that occur anywhere on the sysplex.

Question:

What messages does IBM zAware analyze?

Answer:

IBM zAware is able to analyze large quantities of message log data, or OPERLOG. Using prior message log data and modeling, IBM zAware builds a model of normal system behavior, tailored to your system, to compare to current message log data from the connected z/OS system. IBM zAware monitors the z/OS OPERLOG, which contains all messages written to the z/OS console, including application-generated messages and suppressed messages. Through OPERLOG data, IBM zAware detects unusual message patterns that typical monitoring systems miss, as well as unique messages that might indicate system problems. Its ability to pinpoint deviations in normal system behavior improves real-time event diagnostics.

Question:

How does IBM zAware develop a baseline?

Answer:

IBM zAware needs to create a baseline model to perform basic analysis using an estimated 90 days worth of data for building a more accurate model. Your installation can modify this training period based on knowledge of the workloads running on z/OS. Also you can prime the server by transferring prior data for IBM zAware monitored clients, and request the server to build a model for each client from the transferred data.

Question:

How does IBM zAware score and display unusual messages?

Answer:

The IBM zAware server determines anomaly scores by comparing current data from each client to the model it has built. For each 2-minute interval, the server collects messages and reports on them in a 10 minute interval. To calculate the anomaly score, the IBM Aware server considers

- The number of times a message is issued within the interval
- The probability of occurrence for this message on this system
- If a message appears outside its normal context (group of messages)
- IBM expert knowledge of this message

IBM zAware analyzes the messages and displays scores using one color coded rectangle per interval using a web-based GUI. If the interval contains normal messages, the server assigns a low score and light blue color to the interval (rectangle). If the server detects anomalous message patterns, based on the criteria above, it assigns a higher score and darker color to the interval (rectangle).

Question:

What is the change to IBM zAware from the July 23, 2013 announcement?

Answer:

When a new workload is added to a system being monitored by IBM zAware, or moved to a different system, it often generates messages that are not recognized by IBM zAware. These messages are subsequently flagged as irregular and cause orange bars to appear on the IBM zAware analysis panel because they are not part of the baseline. When a new baseline is created, these messages will be marked appropriately.

While this is a valuable tool to see the impact of change, in the case where both acceptable anomalies appear in the same timeframe as an actual problem, it's beneficial to have some control over which messages continue to highlight intervals as usual.

The new enhancement will allow the user, using a graphical user interface (GUI) panel, to mark the desired messages as "ignore." An ignored message will not be a part of the IBM zAware analysis and scoring either forever or until the next time a baseline is created.

Question:

How are messages sent from the IBM zAware monitored client to the IBM zAware host?

Answer:

Messages from the IBM zAware monitored client machine are sent to the IBM zAware host across your IP network or HiperSocket, by the z/OS Logger component.

Question:

What configuration is required for the IBM zAware environment on a zEC12 or zBC12?

Answer:

The IBM zAware partition running on the zEC12 or the zBC12 host or DR system requires:

- The IBM zAware features need to be added to the zBC12 configuration (via econfig and the hardware order process)
- A shared or dedicated OSA port, with an IP address
- The system must be using OPERLOG as the hardcopy medium
- On either a zEC12 or zBC12, your installation can assign only one of two processor types for the IBM zAware partition: Integrated Facilities for Linux (IFLs) or central processors (CPs). The IFLs or CPs can be shared or dedicated. This is dependent on your configuration and volume. You will need a

little more capacity to perform the initial bulk load than when you've stabilized the volume of messages you send to IBM zAware.

- Use the following guidelines to determine the amount of processor resource that your installation requires. Keep in mind that IFLs run at full capacity, but CPs can run at full capacity or various subcapacity settings, depending on the zEC12 or zBC12 model. These guidelines are based on IFLs or CPs that are running at full capacity.
- For an IBM zAware partition on a zEC12:
 - For configurations with up to 10 monitored clients, for a total maximum rate of 500 messages per second:
 - Approximately 25% of one zEC12 IFL or 25% of one full-capacity zEC12 CP (for example, model 701) is required for initial priming and training.
 - Approximately 20% of one zEC12 IFL or 20% of one full-capacity zEC12 CP (for example, model 701) is required for analysis operations.
 - For configurations with up to 10 monitored clients, for a total maximum rate of 1500 messages per second:
 - Approximately 80% of one zEC12 IFL or 80% of one full-capacity zEC12 CP (for example, model 701) is required for initial priming and training.
 - Approximately 40% of one zEC12 IFL or 40% of one full-capacity zEC12 CP (for example, model 701) is required for analysis operations.

Note: If you connect more than 10 monitored images during a 15-minute interval when the maximum message rate per second is approximately 1500, the capacity of a single zEC12 IFL or a single full-capacity zEC12 CP might be overrun during the initial priming and training phase. To avoid this potential condition, you can configure a second zEC12 IFL or CP for IBM zAware to use.

- For an IBM zAware partition on a zBC12
 - For configurations with up to 10 monitored clients, for a total maximum rate of 500 messages per second:
 - Approximately 30% of one zBC12 IFL or 30% of one full-capacity zBC12 CP (for example, model z01) is required for initial priming and training.
 - Approximately 20% of one zBC12 IFL or 20% of one full-capacity zBC12 CP (for example, model z01) is required for analysis operations.
 - For configurations with up to 10 monitored clients, for a total maximum rate of 1500 messages per second:
 - Approximately 95% of one zBC12 IFL or 95% of one full-capacity zBC12 CP (for example, model z01) is required for initial priming and training.
 - Approximately 40% of one zBC12 IFL or 40% of one full-capacity zBC12 CP (for example, model z01) is required for analysis operations.
 - If you connect more than 10 monitored images during a 15-minute interval when the maximum message rate per second is approximately 1500, the capacity of a single zBC12 IFL or CP might be overrun during the initial priming and training phase. To avoid this potential condition, you can configure a second zBC12 IFL or CP for IBM zAware to use.
- Storage – estimated 500 GB
- 4 GB of base memory for up to 6 IBM zAware monitored clients; 256 MB additional memory per additional IBM zAware monitored clients after six
- IBM zAware monitored clients must be running z/OS V1.13 plus required PTFs
- Browsers:
 - Mozilla Firefox Extended Support Release (ESR) 10 with available updates
 - Microsoft Windows Internet Explorer 9
 - For further information please review the IBM Advanced Workload Analysis Reporter (IBM zAware): [www.ibm.com/servers/resourceink/lib03010.nsf/0/4F9114255D7D1F3285257A6A0077C2CA/\\$file/SC27-2623-00.pdf](http://www.ibm.com/servers/resourceink/lib03010.nsf/0/4F9114255D7D1F3285257A6A0077C2CA/$file/SC27-2623-00.pdf)

Question:

Which processors and operating systems are supported as IBM zAware monitored clients?

Answer:

These servers can be IBM zAware monitored clients as long as they run z/OS V1R13 with PTFs for system logger and bulk load client) or higher:

- An IBM zEnterprise EC12 (zEC12) server or IBM zEnterprise BC12 (zBC12)
 - IBM zEnterprise 196 (z196) or IBM zEnterprise 114 (z114)
 - IBM System z10® Enterprise Class (z10 EC™) or IBM System z10 Business Class (z10 BC)
-

Question:

Does IBM zAware support z/OS under z/VM as an IBM zAware monitored client?

Answer:

A z/OS can be an IBM zAware monitored client when running as a z/VM guest. As long as z/OS is Release 1.13 and has the appropriate z/OS service installed on z/OS.

Question:

How is the price of IBM zAware calculated for the zBC12?

Answer:

- Pricing is based on a calculation of the total number of CPs on the IBM zAware host server plus on the IBM zAware monitored client servers.
 - The CP count is the sum of the purchased CPs on the IBM zAware host *plus* the purchased CPs on IBM zAware monitored servers, not including specialty engines (no CBU, On Off CoD or specialty engines – zIIPs, zAAPs, IFLs, or CFs). The count must include high water marks CPs.
 - On the zBC12, the minimum block of CPs that you can purchase is a 2-pack (equal to a count of two purchased CPs). The next option is a 4-pack followed by a 6-pack. After you have more CPs than six, the pricing is similar to the zEC12. Round up the total purchased CPs to the nearest multiple of 10, which is then divided by 10 to yield the number of 10-packs you will need to purchase.
 - The IBM/BP CTS or seller will help determine the number of IBM zAware connections and use this when placing the order in eConfig.
-

Question:

How is the price of IBM zAware calculated for the zEC12?

Answer:

- Pricing is based on a calculation of the total number of CPs on the IBM zAware host server plus on the IBM zAware monitored client servers.
- The CP count is the sum of the purchased CPs on the IBM zAware host *plus* the purchased CPs on IBM zAware monitored servers, not including specialty engines (no CBU, On Off CoD or specialty engines – zIIPs, zAAPs, IFLs, or CFs). The count must include high water marks CPs.
- On the zEC12, the minimum block of CPs that you can purchase is a 10-pack (equal to a count of 10 purchased CPs). You need to round up the total up to the nearest multiple of 10, which is then divided by 10 to yield the number of 10-packs you will need to purchase.
- The IBM/BP CTS or seller will help determine the number of IBM zAware connections and use this when placing the order in eConfig.

Question:

Are there any other costs to be considered in the costs of running IBM zAware?

Answer:

You need to include the additional hardware costs that may be needed for running IBM zAware – such as the CP or IFL capacity, memory, storage or OSA ports.

Question:

What if I have a DR (disaster recovery) or CA (continuous availability) machine that needs to be able to run IBM zAware? Will I have to purchase those IBM zAware connections for the DR machine?

Answer:

A Disaster Recovery option is available to ship IBM zAware firmware to a zEC12 or zBC12 DR server to be used in case of emergency. The quantity of IBM zAware connection DR features (FC #0102) will be calculated the same way features are calculated for an IBM zAware host system but these will be no charge (zero priced) features.

Question:

When I am running IBM zAware, can I define more than one OSA connection be used for redundancy on the z/OS images?

Answer:

Yes. IBM zAware has the ability to support multiple OSA connections from a system being monitored - something you might do when creating a high availability design.

Question:

What feature codes are needed to order IBM zAware?

Answer:

IBM zAware is available with IBM zEC12 and zBC12 models.

Features for ordering IBM zAware will be located on the Processor usage panel of eConfig

- FC 0011 will represent that IBM zAware is installed.

For the zEC12:

- FC 0101 will represent the quantity of IBM zAware CP 10-packs on the zEC12
- FC 0102 will represent the quantity of IBM zAware CP 10-packs on the zEC12 DR server

For the zBC12:

- FC 0138 will represent the IBM zAware CP 2-pack on the zBC12
- FC 0140 will represent the IBM zAware CP 4-pack on the zBC12
- FC 0142 will represent the IBM zAware CP 6-pack on the zBC12
- FC 0150 will represent the quantity of IBM zAware CP 10-packs on the zBC12
- FC 0139 will represent the IBM zAware CP 2-pack on the zBC12 DR server
- FC 0141 will represent the IBM zAware CP 4-pack on the zBC12 DR server
- FC 0143 will represent the IBM zAware CP 6-pack on the zBC12 DR server
- FC 0151 will represent the quantity of IBM zAware CP 10-packs on the zBC12 DR server

Question:

What is the distance that IBM zAware will support between the monitored client and the IBM zAware host server?

Answer:

In March 2013, IBM announced that the maximum supported distance between the IBM zAware host system and monitored clients has been increased to 3500 km.

Question:

How does IBM zAware deal with messages that are produced from the IBM Healthchecker facility? If the Healthchecker produces a message indicating something in my environment is at odds with best practices, does IBM zAware read that message and does it assign an anomaly score to it?

Answer:

The strength of IBM zAware is that it does not require hard-coded rules which need updating but 'learns' what's normal for a system and dynamically highlights what's different. Due to our experiences, when we designed IBM zAware, we knew certain messages are always indicative of a problem, and so we built into the scoring of IBM zAware knowledge of a certain subset of messages, some of which happen to be some Healthchecker messages. In other words, IBM zAware uses many factors to impact both a message score, and a (10-minute) interval score, including rarity of a message, it's frequency, if it's within its normal context of messages, and finally, if it happens to contain a message we know is either critical, or possibly of concern. A Healthchecker message by itself may not cause the score to rise, and in some cases that's good because if Healthchecker is already catching it, you don't need IBM zAware to also do so.

Note:

All new zEC12 enhancements for July 23, 2013, will require Driver 15, available on September 20, 2013.

Flash Express

Question:

What is Flash Express?

Answer:

Flash Express is an optional feature on the zEC12 platform that exploits solid state drive (SSD) technology, providing a faster paging device compared to hard disk drive (HDD) storage – extending options for faster paging can improve the availability of your systems. For companies with demanding service level expectations, Flash Express can dramatically improve availability during transitions of workload processing where paging spikes might occur. For example, when your workloads *shift from batch to start of day processing, or when a new region is started.*

Question:

What was announced in the February 24, 2014 announcement concerning Flash Express supporting the Coupling Facility Control Code (CFCC)?

Answer:

IBM is fulfilling a statement of direction from July 2013.

The Flash Express feature is designed to offer improved availability and performance for key workloads at critical processing times. Flash memory is designed to help improve availability during workload transition periods and spikes such as start-of-day processing or batch online transitions.

With enhancements in CFCC Level 19, exploitation of the Flash Express feature is designed to help improve resiliency while providing cost-effective standby capacity to help handle the overflow of WebSphere MQ shared queues. You can now specify overflow areas for certain Coupling Facility list structures in the Storage Class Memory (SCM) provided by the Flash Express feature. This is designed to allow structure data to be migrated to Flash Express memory as needed and migrated back to real memory to be processed. When using WebSphere MQ for z/OS Version 7 (5655-R36), this new capability is expected to help provide significant buffering against enterprise messaging workload spikes and to help provide support for storing very large amounts of data in shared queue structures, potentially allowing several hours of data to be stored without causing interruptions in processing.

Question:

How can Flash Express improve the availability of my environment?

Answer:

Flash Express helps organizations improve availability and performance as often demanded by service level agreements, by improving paging, thus reducing specific paging delays and resultant system degradation. In particular, it helps reduce paging delays caused by page-ins of workloads during transition periods such as when starting new WebSphere servant regions, serving new workloads after a period of time where the system was processing other work (like overnight to first shift), or for handling page-ins needed for collection of runtime diagnostics like SVC dumps. These workloads and others could cause significant delays or stall conditions - Flash Express reduces these delays. In addition, z/OS V1.13 provide support for writing pageable (1M) large pages to Flash Express, which can improve CPU performance and flexibility compared to writing smaller 4K pages or using large pages that previously were fixed in main memory.

Question:

Which servers, operating systems and middleware support Flash Express?

Answer:

Flash Express is supported on zEC12 and zBC12 servers and minimally requires z/OS V1.13 with the z/OS V1R13 RSM Enablement Offering web deliverable available on <http://www.ibm.com/systems/z/os/zos/downloads/>

The Linux distributors from Red Hat and SUSE are supporting Flash Express when used for Linux for System z running natively on the server. The support is available with RHEL 6.4 and SES 11 SP3 today.

Other IBM software products plan to support pageable large pages in the future, including Java SDK7 5655-W43 and 5655-W44, Note that IBM SDK7 for z/OS Java is available for use by IBM middleware products running Java, such as IBM IMS 12 (5635-A03), IBM DB2 10 for z/OS (5605-DB2), and the Liberty profile of IBM WebSphere Application Server for z/OS v8.5 (5655-W65); and is planned for use by a future release of CICS Transaction Server for z/OS.

Question:

What are the plans to exploit Flash Express in a native System z Linux environment?

Answer:

Now that we have their support, IBM will be working with the Linux community to explore Flash Express exploitation opportunities.

Question:

How do I order Flash Express?

Answer:

Flash Express is ordered through eConfig as an optional feature of zEC12 and zBC12. Cards are ordered in pairs;

- Order in pairs, (where second card provides redundancy for the first card)
 - A pair of cards equals 2.8 TB total memory equating to 1.4TB customer usable storage
 - Maximum of 4 card pairs (total of 5.6 TB)
-

Question:

How much Flash Express capacity will I need to order?

Answer:

You should plan to configure roughly the same amount of Flash Express paging space as you configure disk paging space today. Refer to z/OS in the Tuning Reference Guide SA22-7591 for chapters that discuss Auxiliary Storage Management (ASM) and the paging datasets. Most customers will find one pair of Flash Express cards covers their paging needs. You should not need to order additional capacity above your current paging space.

Question:

Do I need to change my applications to use Flash Express?

Answer:

No you need not alter applications. Flash Express is automatically used for paging if defined to the partition unless PAGESCM=NONE is specified in IEASYSxxx.

Allocation across LPARs:

- All paging data can easily reside on Flash Express
- No definition required for placement of data on Flash Express
- At IPL, z/OS detects if Flash Express is assigned to the partition.
- z/OS automatically uses Flash Express for paging unless specified otherwise via PARMLIB

Question:

If I use Flash Express, will I require less real memory or less paging space on auxiliary storage?

Answer:

Your application performance will be negatively affected if real memory is reduced. You should not reduce real memory if you use Flash Express. Flash Express does not serve the same purpose nor is it used in the same way as real memory. Real memory can be used for delivering performance improvements to your system software and will not be depleted in the event, say, of an SVC dump, which would otherwise cause paging spikes and consume a significant amount of real memory, potentially detracting from other more productive memory use.

If you choose to stop using Flash Express or in the event of failover from Flash Express, you need to ensure that you have sufficient disk paging space available. Therefore, it is recommended that you retain the same amount of disk paging capacity that you have today even when using Flash Express.

Question:

Can I subsequently order additional Flash Express adapters, and subsequently MES a server that did not originally include Flash Express?

Answer:

Yes, you can place an MES order for additional card pairs to be installed at a later time.

Question:

Can Flash Express adapters be replaced or additional adapters added non-disruptively?

Answer:

Yes Flash Express adapters can be added concurrently. Furthermore, if an adapter fails, the system runs using the mirrored card, and a failing adapter can be replaced concurrently.

Question:

Where are Flash Express cards installed in the server?

Answer:

The Flash Express cards will be installed in the PCIe I/O drawer of the zEC12 or zBC12.

Question:

What hardware set-up and configuration is required in order to allocate Flash Express for use by an LPAR?

Answer:

Flash Express introduces a new class of memory, storage class memory. A profile specifies the amount of storage class memory initially brought online to the LPAR; additional Flash memory can be brought online or offline at a later point. Using the Flash Memory Allocation panel on the HMC or SE you can define an initial allocation for storage class memory and a maximum amount per partition. You might wish to allocate all of the storage class memory at once and later dynamically change the memory available to a particular logical partition as needed.

Question:

Is Flash Express automatically used by the system?

Answer:

At IPL z/OS detects if any SCM is assigned to the partition and will automatically use Flash Express for paging unless PAGESCM=NONE is specified in IEASYSxxx.

Question:

How do I change the Flash Express capacity assigned to an LPAR?

Answer:

The initial and maximum Flash Express available to a particular LPAR is specified through the SE or HMC through the Flash Memory Allocation panel. The panel allows users to increase the maximum allocation as long as the Flash Express capacity exists. A z/OS online command can also be used to bring additional flash capacity online after an IPL, as long as the Flash Express capacity exists.

Question:

What is dynamic reconfiguration for Flash Express?

Answer:

Dynamic Reconfiguration announced in 2012 and available early in 2013. It allows a user to dynamically configure the amount of flash memory to be moved into and out of the partition.

Question:

Describe the use of pageable large pages and how it relates to Flash Express.

Answer:

Pageable large pages require Flash Express. If you are using Pageable Large pages, you can gain additional performance (CPU) benefits. For instance Pageable Large Pages can help deliver performance improvements as paging is accomplished using larger aggregate page-ins, which results in faster and fewer paging operations and fewer Translation Look aside operations. Customers may see a resultant reduction in CPU cycles for such page-ins.

Question:

Is data stored on Flash Express persistent across IPLs? What about PLPA data?

Answer:

No, data on Flash Express is not persistent across IPLs. Customers will note that PLPA data will remain persistent across IPLs, because PLPA data resides on disk.

Question:

Is the data stored on the Flash Express card encrypted?

Answer:

Yes. The data stored on the System z Flash Express cards is protected – even if these cards are removed from System z. The Flash Express cards have self-encryption capabilities and will not release data until they are provisioned with a data encryption key served by the IBM System z Support Element (SE). Systems with Flash Express cards use smart cards that are plugged into the Support Elements (SE) to protect the keys that enable the Flash Express cards to release clear text data. In the event of a failure, or when migrating to another processor, removal of the smart cards causes a "cryptographic erasure" - the data encryption key can no longer be served to the Flash Express cards. This provides assurance that your data remains safe and helps guard against sensitive information from being recovered from the Flash Express cards. This is one of numerous safeguards that IBM System z provides to help keep your data secure.

Question:

Is there any special configuration needed for encryption support?

Answer:

No special configuration is required. The Flash Express cards have self-encryption capabilities and will not release data until they are provisioned with a data encryption key served by the IBM System z Support Element (SE).

Question:

Do the SSDs on the Flash Express card ever wear out or degrade over time?

Answer:

All Flash wears out over time, but the lifetime of Flash is designed to last many years, beyond the effective lifetime of current Flash technology.

Question:

Are there ways to find out about the remaining life left in my Flash Express card?

Answer:

The Service Element sends logs to IBM that report on the state of the card. The cards run as mirrored RAID 10 pairs, so that the failure of any one card does not affect the customer's operations. If a card experiences a failure, the card is automatically called out for service and replacement, while the system runs on the partner card.

Question:

Will I still require paging disk space if I have Flash Express cards installed?

Answer:

Yes. Paging space on disk serves as a backup. If the flash overflows, it would page to disk, or if there was a disaster where both Flash Express cards fail, you would default to paging to disk.

Question:

Will having additional pairs of Flash Express cards improve paging?

Answer:

Ordering additional card pairs will not improve performance. The Flash Express card pairs provide built in resiliency (card pairs), purchasing additional card pairs can provide additional resiliency in the unlikely event that there is a problem with both cards of a pair.

Generally, you should not need to order additional capacity above your current paging space. The amount of paging space is unchanged, so you will use the same paging space on Flash Express as on you use with spinning disk today.

Question:

While extremely rare, what happens if an error occurs causing *both* cards in one Flash Express pair to fail?

Answer:

We recommend that you continue to allocate disk paging space. In the case of a failure, paging would revert to paging on disk. If failover occurs to disk, additional resiliency can be derived from having an extra card pair.

Question:

Will having more Flash Express cards improve resiliency?

Answer:

If you want additional resiliency, a second card pair could be configured in a second PCIe I/O drawer.

Question:

Is IBM's recent acquisition of Texas Memory Systems (TMS) related to your Flash Express offering?

Answer:

No. Although TMS offers PCIe based cards similar to our Flash Express, our firmware is not designed to work with the TMS cards, nor have we done any development or test work with these cards. While we don't comment on future products, we will investigate the TMS product line and determine if a switch over is something to consider in the future.

Question:

Can you ever allocate too much flash memory across your LPARS – where you would get into a problem?

Answer:

No. Allocating too much flash to your LPARs should never cause a problem. The sum of all flash memory that is online to all the partitions on that CEC at any given instance can not exceed the amount of physical flash memory available on that CEC.

If you try to bring online more flash memory than what is physically available the online command would fail. If this happens during IPL a z/OS message is issued saying we could not bring online the amount of Flash requested and would only bring online what is available.

Question:

How long would it take to de-allocate the flash offline on z/OS?

Answer:

The answer is it depends, if the flash memory that is being configured offline is not being utilized (that is, it is not currently allocated to any z/OS address spaces) it can be taken offline very quickly. If the flash memory that is coming offline contains z/OS paging data then that data would have to be migrated either to other parts of flash memory or disk before the particular flash memory can come offline. How long this takes it would depend on how much flash memory is being configured offline and if there is space in other flash memory increments or disk to migrate the data that currently resides on the flash SCM increments that are coming offline.

Question:

Is the working memory set size change for z/OS with flash available?

Answer:

No.

zEC12 and zBC12 Warranty

Question:

What is proprietary diagnostic support?

Answer:

A zEC12 and zBC12 comes with proprietary diagnostic support. Proprietary diagnostic support gives the full maintenance package including call home support and repair and verify (R&V) procedures to assist the System z Service Support Representative (zSSR) in doing problem repairs.

Question:

How long is the machine entitled to proprietary support?

Answer:

The default is for the first year of warranty period and remains in effect if you purchase an IBM maintenance agreement.

Question:

What happens when the warranty ends and there is no maintenance agreement in place or maintenance is being provided by someone besides IBM?

Answer:

The machine reverts to a base service state. The non-IBM representative will not have access to IBM proprietary tools.

Question:

What is the difference between base and proprietary service state?

Answer:

Base support includes repair and verify help. The additional proprietary service state includes locating of IBM field replacement units, help indicators on replacement units, guided videos and specialized tools for zSSR use.

Question:

Can other service providers still fix the machine?

Answer:

Yes, but they will not have access to IBM intellectual property (i.e. proprietary diagnostic support).

Question:

If I am on base support, what level of support will I get if I upgrade (MES) my system?

Answer:

Base or Proprietary Service state does not change any of the MES warranty or maintenance policies presently in effect and the MES assumes the status of the machine.

Question:

Will I lose the call home capability if my zEC12 and zBC12 is no longer on an IBM maintenance agreement?

Answer:

Yes, if the machine is no longer on warranty or has a valid IBM maintenance agreement in place.

IBM zEnterprise BladeCenter Extension (zBX)

Question:

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

Answer:

The zBX machine type is 2458. The new Model 003 attaches to the zEC12 and zBC12. The Model 002 attaches to the z196 or z114. IBM POWER7® and IBM System x blades are ordered independently of the zBX (see the POWER7 and System x section of this document for details). DataPower XI50z is ordered as a feature of the zBX but will generate machine type 2462 Model 4BX.

Question:

What are the minimum levels of System z operating system software releases needed to support a zBX as part of a zEnterprise System?

Answer:

The following are the minimum levels that are required:

- z/OS
 - zBX Ensemble support: z/OS V1.11 or higher¹¹
 - z/VM
 - z/VM V6.2 for Enterprise Unified Resource Manager ensemble support¹²
 - z/VSE
 - IEDN support: z/VSE 5.1
 - z/VSE 4.3: IEDN access through z/VM VSWITCH (z/VM V6.2)
-

Question:

If I have four zBX frames (racks) connected to one zEnterprise (zEC12, zBC12, z196 or z114), is that considered one zBX or four zBXs?

Answer:

A single configuration can include up to four frames (depending on the number of blades ordered). The zEnterprise only supports the attachment of one optional zBX per server.

Question:

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

Answer:

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

Question:

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

Answer:

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

¹¹ Note that z/OS V1.10 is no longer a "supported" operating system and you will require a LifeCycle Extension contract V1.10 to participate. In September 2013, z/OS V1.11 will also fall under this requirement.

¹² See section 'Statement of Direction' for information on OpenStack, z/VM 6.3 and Unified Resource Manager

Question:

How many zEnterprise servers can connect to a zBX?

Answer:

By defining an ensemble, a combination of up to eight zEnterprise servers can connect over the intra-ensemble data network (IEDN) to one zBX, but only one zEC12, zBC12, z196 or z114 “owns” the management and support connection (intra-node management network - INMN).

Question:

Can I have an ensemble with a zEC12 (or zBC12) and a zBX Model 003, and attach z196 or z114 servers? Can I do that kind of mix and match?

Answer:

Yes, workloads can be defined with resources on both a z196 and the zBX Model 003, but the zBX Model 003 can only be connected using the INMN to the zEC12 or zBC12.

Question:

What is the physical size of the zBX?

Answer:

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

Question:

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

Answer:

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

Question:

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

Answer:

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

Question:

If I have a zBX Model 003 with four frames, are there top-of-rack (TOR) switches in each frame?

Answer:

No, TOR switches for management and data are only in the first frame, so all zBX-003 network connectivity flows through those switches.

Question:

If I order a zEC12 or zBC12 with a non-raised floor (NRF) option, can I also order a zBX with a NRF option?

Answer:

Yes a zBX can be installed on a non-raised floor. On the zBX Model 002, and the zBX Model 003 prior to March 2013, the I/O and power cables will exit from the bottom of the machine. In March 2013, IBM announced the ability for top exit I/O and power for the zBX Model 003.

Question:

I have a zBX Model 002 attached to my z196 (or z114). I'd like to upgrade the z196 to a zEC12 (or zBC12) and move the zBX Model 002 to different z196 (or z114). If I purchase new entitlements on the second z196 (or z114) is there any issue for me to move the zBX Model 002?

Answer:

With the announcement on August 28, 2012 this capability is available. You MUST attach the zBX Model 002 to another server.

Question:

I have a zBX Model 002 attached to my z196 (or z114). I'd like to upgrade the z196 to a zEC12 (or z114 to a zBC12) and upgrade the zBX Model 002 to a zBX Model 003 but attach it to a different zEC12 (or z114). If I purchase new entitlements on the second zEC12 (or zBC12) is there any issue for me to move the zBX Model 003?

Answer:

You are able to move the upgraded zBX Model 003 to a different zEC12 (or zBC12). You MUST attach the zBX Model 003 to another server.

Question:

Can I detach my Model 002 or Model 003 and let it sit standalone?

Answer:

No,

Question:

Why are there only two BladeCenter chassis per zBX frame?

Answer:

The maximum power that a rack can support is 240amps. Two BladeCenters, fully configured in accordance with the zBX infrastructure and the System z redundancy requirements, could reach this limit. Additionally, as the zBX is shipped as a single unit, more than two BladeCenters would exceed shipment tilt and weight limitations.

Question

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

Answer:

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

Question:

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

Answer:

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

Question:

How is the BladeCenter chassis configured in the zBX?

Answer:

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

The components of the BladeCenter chassis are also not managed independently. The operational controls provided on the zEC12 or zBC12 by the Unified Resource Manager enable the underlying hardware infrastructure components to be managed transparently.

Question:

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

Answer:

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

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

Question:

What workloads are good candidates to run on the zEnterprise System?

Answer:

The zEnterprise provides the opportunity to bring together most or all of the application and information components of important end-to-end business applications into one unified, simplified, and easily managed server environment. Strong candidates for zBX hosting include application components with affinities to System z application and information services. Such applications are found in every industry, including banking, insurance, retail, government, and manufacturing. Some of the candidate workloads include business intelligence, data warehousing, business analytics, ERP (including SAP applications), CRM, infrastructure services (such as monitoring, storage management, and security services), Web serving, and other multi-tier application architectures. The zEnterprise also offers new opportunities for recentralizing enterprise information (for improved governance and customer privacy protection), master data management, and enterprise reporting.

IBM has developed a Fit for Purpose methodology to help aid in business and application deployment decisions. When looking at the big picture with complex topology's that many clients have, it is easy to see why one size fits all is not the answer some think it is. IBM System z representatives can meet with prospective zEnterprise customers to perform this methodology and identify the strongest application candidates and determine the business value that zEnterprise delivers. Please contact an IBM representative for more information.

Question:

Today I have WebSphere Application Server running on distributed servers, and I've been considering moving at least some of those applications to the mainframe. Should I move my applications to WebSphere Application Server for z/OS, WebSphere Application Server for Linux on System z, or to a zBX blade?

Answer:

It depends. For example, if your applications have close affinity to z/OS-based information and/or applications, or require the highest service qualities, hosting them on WebSphere Application Server for z/OS is likely the best strategy. If z/OS affinity is less important, or if there are particular affinities with Linux or AIX (operational or otherwise), Linux on System z, System x blades or AIX blades may make the most sense. In some cases the best choice may be to run WebSphere Application Server on all three operating systems.

IBM offers "Fit for Purpose" workshops to review your IT architecture and make recommendations for server platform placement. Please ask your IBM representative for more information.

Question:

I am a user of a large ERP application and have been running it on System z for a long time. I have migrated everything in my shop to System z and now run DB2 for z/OS, several applications on z/OS, all my ERP application servers on Linux on System z, and some other applications on Linux on System z. The few exceptions to our general deployment strategy consist of some third party applications. The application vendors do not support their products on System z, but the applications are critical to our business. I think these applications would be good candidates for a zBX blade solution. Correct?

Answer:

Yes. The zBX's System x or POWER7 blades provide an excellent option for such applications, and customers can use them to consolidate applications that do not yet run on Linux on System z or on z/OS.

Customers should also ask their vendors for the solution attributes they need, including support for their preferred platforms. Through its PartnerWorld® program, IBM welcomes and assists vendors in bringing their applications to more operating systems, including Linux on System z and z/OS. In this way, vendors can expand their market opportunities, increase customer satisfaction and retention, and thus grow their revenues and profits. Typically the effort required is minimal, and IBM is happy to work with vendors to help them achieve these goals.

Question:

Are there any applications specific to zBX?

Answer:

There are no applications specific to zBX per se. The POWER7 blades, or System x blades, installed in the zBX are exactly the same POWER7 or System x blades installed in any other BladeCenter chassis and are interchangeable. Naturally, those POWER7 blades run the same AIX operating system and the System x blades run the same Linux or Windows operating system. There is no difference in the blades or in the core chassis, by design.

However, applications running on the zBX-based blades enjoy several benefits: easy-to-use management with the Unified Resource Manager, reduced infrastructure complexity, and closer proximity to the System z workload components.

Question:

Are new workloads such as virtual desktops (vdi) good candidates for the zBX?

Answer:

Yes. Over the past ten years we have had several customers do virtual desktop solutions on System z Linux. The most successful one right now is using open source project called eyeos (<http://eyeos.org/>). It is being run several places, including a University in the southeast US running 10,000 virtual desktops.

Question:

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

Answer:

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

Question:

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

Answer:

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

Question:

Does IBM System z plan to implement new blades and chassis in the zBX as they become available?

Answer:

Our intention is to integrate desired new technologies into the zBX after we have tested and confirmed support for the new devices.

POWER7 Blades for zBX

Question:

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

Answer:

Effective 1/15, the standalone PS701 Express blade will no longer be available for order. But, System z has made arrangements for zBX users of PS701 blades to be able to still order them for use in the zBX. When placing an order, through an IBM Business Partner or Distributor, an IBM sales representative, or through <http://www.ibm.com>, you must make sure that the representative (or techline) are aware you want the PS701 blade ordered using the RPQ 8A2139. You will also need to provide the zBX Model and serial number as the RPQ is only for ordering PS701 blades for zBX environments. You will still need to make sure that your Central Processing Complex (CPC) has entitlements on it to support the PS701.

You are required to order the hardware configuration specified by this zBX website document: http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&apname=STGE_ZS_USEN&htmlfid=ZSY03019USEN&attachment=ZSY03019USEN.PDF

Question:

How do I order the POWER7 blade?

Answer:

The PS701 Express blade with features and software may be obtained through an IBM Business Partner or Distributor, an IBM sales representative, or through <http://www.ibm.com>. It is recommended that you use the following document, posted and regularly updated on the zBX website, for the most current ordering

information. http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&apname=STGE_ZS_USEN&htmlfid=ZSY03019USEN&attachment=ZSY03019USEN.PDF

Question:

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

Answer:

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

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

Question:

Which version of PowerVM™ do I need?

Answer:

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

Question:

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

Answer:

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

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

The following is a link for entering the activation codes in the blades' existing BladeCenter environment: <http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7hch/iphchenablevios.htm?resultof=%22%61%63%74%69%76%61%74%69%6f%6e%22%20%22%61%63%74%69%76%22%20%22%63%6f%64%65%22%20>

Question:

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

Answer:

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

Question:

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

Answer:

The blade would be rejected by Unified Resource Manager.

Question:

What applications can run on the POWER7 blade in a zBX?

Answer:

Applications that are certified for POWER7 and PowerVM EE are compatible. The same software and hardware environment is provided on the POWER7 blades in the zBX.

Question:

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

Answer:

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

Question:

Are the Power blades managed by Systems Director?

Answer:

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

Question:

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

Answer:

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

Question:

What adapters can be ordered on the PS701 blade?

Answer:

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

Question:

What storage attaches to the POWER7 blades?

Answer:

A list of supported storage products to support the POWER7 blades can be found on the web at: http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=PS&infotype=SA&appname=STGE_ZS_ZS_USE_N&htmlfid=ZSP03437USEN&attachment=ZSP03437USEN.PDF. No storage products can be installed in the zBX.

Question:

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

Answer:

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

Question:

Will the POWER7 support PowerHA™?

Answer:

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

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

Power HA configuration for the zBX environment is similar to standard Power environments, with the particularity that it uses only virtual I/O resources. Currently, PowerHA for zBX support is limited to

failover inside the same ensemble and not for business continuity across ensembles.

Refer to TechDocs FLASH10750 for the most current information on supported releases: <http://www.ibm.com/support/techdocs/atmastr.nsf/webindex/flash10750>

Question:

Are there any published case studies that justify the use of zBX with IBM pSeries® servers? And that include what workloads are a good fit for zBX.

Answer:

Case studies that justify the use of zBX with POWER7 blades are available for BG-Phoenixis (http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=AB&infotype=PM&appname=STGE_ZS_ZS_USEN&htmlfid=ZSC03118USEN&attachment=ZSC03118USEN.PDF) and

Marist College

(http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=AB&infotype=PM&appname=STGE_ZS_ZS_USEN&htmlfid=ZSC03162USEN&attachment=ZSC03162USEN.PDF).

A press release and video are available for

IZUMIYA <http://www.ibm.com/press/us/en/pressrelease/41326.wss> http://www.ibm.com/software/success/cssdb.nsf/CS/ARBN-98DKUK?OpenDocument&Site=eserverzseries&cty=en_us

Workloads that are a good fit for zBX are those that have an affinity to System z such as multi-tier applications where the data base tier is on the mainframe platform and the application servers are on distributed platforms. Examples include transaction processing, database management, business processing applications, analytics, web, collaboration, and infrastructure workloads. However, a fit for purpose study can be conducted to help develop a custom solution.

IBM System x blades for zBX

Question:

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

Answer:

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

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

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

Question:

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

Answer:

We support 56 HX5 blades per zBX. These are 16 core blades so a total of 896 cores will be supported at a level of 1 virtual server per 2 cores. Our intention is to work with our clients and understand their needs going forward including the need for additional scale for both physical and virtual server support and/or additional IBM System x Blade types.

Question:

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

Answer:

Similar to the PS701 blade for the zBX today, the System x blade will not be part of the zBX configuration. The System x blade with features and software may be obtained through an IBM Business Partner or Distributor, an IBM sales representative, or through the web. The following document should be used as a guide for specific ordering information. http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&appname=STGE_Z_S_ZS_USEN&htmlfid=ZSL03128USEN&attachment=ZSL03128USEN.PDF

Question:

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

Answer:

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

Question:

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

Answer:

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

Question:

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

Answer:

The blade would be rejected by Unified Resource Manager.

Question:

What type of storage will be available for the System x blade?

Answer:

For IBM open storage information you can use the IBM System Storage® Interoperation Center (SSIC) web site - <http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>. For information on support from other industry leaders you can use the IBM ServerProven® web site - <http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/>. Your IBM storage specialist (FTSS – Field Technical Support Specialist) or BP storage specialist can also assist you in finding an open storage product to support the IBM BladeCenter HX5 7873.

Question:

What operating system software releases and hypervisor are supported on the System x blades?

Answer:

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

Three versions of Microsoft Windows are supported. Microsoft Windows Server 2012, Microsoft Windows Server 2008 R2 and Microsoft Windows Server 2008 (SP2) – any of these must be the 64-bit versions.¹³ We recommend that you order Datacenter Edition as that supports an unlimited number of virtual guests, but that is not a requirement.

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

Question:

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

Answer:

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

¹³ All supported releases of Windows may not be available from Microsoft but continue to be available from IBM and other resellers

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

Question:

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

Answer:

Yes. zBX uses KVM as the hypervisor for x86 blades. The flavor of KVM that is used is: Red Hat Enterprise Virtualization - Hypervisor 6.4 (RHEV-H 6.4)¹⁴. The Microsoft SWP certification is inherited from Red Hat and can be found here on the Microsoft site:

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

Question:

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

Answer:

Yes. zBX uses KVM as the hypervisor for x86 blades. The flavor of KVM that is used in the zBX Model 003 is: Red Hat Enterprise Virtualization - Hypervisor 6.4 (RHEV-H 6.4)¹. Red Hat RHEL certifications are inherited from Red Hat and can be found on the Red Hat site:

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

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

Question:

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

Answer:

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

Question:

What applications can run on the System x blade in a zBX?

Answer:

Applications that are certified for the operating systems supported by the zBX can run on a System x blade.

Question:

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

Answer:

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

¹⁴ KVM on zBX Model 002 (z196 and z114 servers) is Red Hat Enterprise Virtualization – Hypervisor 6.1 (RHEV-H 6.1).

Question:

What workloads do you expect to execute on the System x blades? I am trying to figure out if my Windows application will be a good candidate.

Answer:

The IBM zEnterprise provides the ideal foundation and infrastructure required by our customers' increasingly complex workloads. With Windows support, the aperture of applications is widened, and applications that run on Windows will run on the zEnterprise. Best fit workloads that can take advantage of zEnterprise are workloads and applications that rely on System z for data serving and application components, whether it is DB2 for z/OS, IMS, CICS, WebSphere or even Oracle. In addition, they are workloads that require the strength of System z but often have application components on Power® or Intel that are required to complete the end to end business process.

If you have applications running in these two or three tier environments, they may be excellent candidates for migrating to a zEnterprise. Our recommendation is to work with your IBM or Business Partner Sales Representative to run a study – such as our Fit-for-Purpose workshop or Business Value Assessment - to help you identify the value zEnterprise can bring your environment.

Question:

Are the System x blades managed by Systems Director?

Answer:

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

Question:

I am an Independent Software Vendor (ISV) and I see that the zEnterprise supports Linux running on the BladeCenter Extension (zBX). Does that mean that I no longer have a need to enable my applications to run on Linux for System z?

Answer:

No, the need remains. Most applications should still be enabled and deployed on Linux for System z in the zEnterprise. Putting Linux applications on System z allows those applications to take direct advantage of the z/Architecture® that provides not only close proximity between applications and database, but it also includes the following:

- Using the IBM zEnterprise – the server family that can offer extreme scalability and performance for transaction processing and data serving
- The only hardware with L4 Cache to help with high throughput
- Internal private communications within the box at up to 6 GB/s
- Hardware encryption capabilities
- I/O and Network off-load built into the hardware that allows high throughput even when virtualized and running with greater than 90% CPU utilization.
- And with zBX the ability to leverage software/ hardware combinations to further enhance a complete solution
- Leveraging the internal architecture and new capabilities will allow most ISV code to run in a manner giving customers increased value in the code they are running.

While deploying an application on a zBX provides centralized management through the Unified Resource Manager, applications running on the zEnterprise also receive the well-known “qualities of service” that remain System z's competitive advantage. Although different applications have different needs, some of these qualities are reliability, performance, availability, accounting, scalability, I/O bandwidth, staff productivity, and so on. Applications desiring System z levels of these attributes should

choose deployment on Linux on System z. As well, the consolidation capabilities on the zEnterprise allow cost attractive large scale data center simplification and optimization. Those applications that have fit well on a blade in the past, and don't require the qualities of service of System z, will remain a good fit on the blades in a zBX environment.

Remember that in general, a mixture of applications running on Linux for System z will overall perform better as System z typically provides higher utilization levels when compared to a number of blades. Although your mileage may vary, distributed servers often run at average utilization levels in the range of 5 to 20%. On an IFL, one can achieve nearly 100% utilization nearly 100% of the time. The zEC12, the z196 and the z114 also have many performance enhancements that allow direct native Linux on System z applications to perform better than on earlier System z processors.

Finally, many ISVs already have their applications directly executing on Linux for System z and will want to continue to enhance and add to their portfolios of those native applications.

Question:

Is the zBX strategy to be able to consolidate customer developed Microsoft.NET applications onto the z platform? Is it possible to do this under Linux on System z?

Answer:

The zBX can be used to consolidate Microsoft.NET applications. Such applications may be supported on virtual servers running Microsoft Windows or they may be supported on virtual servers running Linux and using the open source Mono Project by Novell. There is a version of Mono for Linux on System z, but as far as I know it is only available for the Debian distribution. Sine Nomine Associates (www.sinenomine.net) provides support for Debian on System z and is responsible for the Mono port. I'm sure that Sine Nomine Associates would be happy to provide additional details.

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

Question:

Does zEnterprise still support the DataPower Appliance?

Answer:

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

Question:

What is the value to the DataPower XI50z of the new firmware level v6.0 that was announced February 24, 2013?

Answer:

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

Other enhancements include:

- Mobile web traffic security for IBM Worklight®: Easy-to-use authentication integration for Worklight platform
- Local response caching and integration with DataPower XC10 Caching Appliance: Optimized application delivery with local response caching on the appliance and seamless integration with elastic caching XC10 appliances
- Enhanced value for System z: New integration capabilities between DataPower and IMS that are designed to allow IMS transactions to more easily consume external web services and remote applications to more easily consume IMS data as a service

Question:

I have a zBX Model 002. Will I be able to take advantage of the new DataPower v6 firmware load?

Answer:

No. Support for DataPower firmware v6 will only be available on DataPower appliances in a zBX Model 003.

Question:

Where can I get more information on DataPower firmware v6?

Answer:

The release notes for DataPower 6.0 are available at: <http://pic.dhe.ibm.com/infocenter/wsdatap/v6r0m0/>

Select WebSphere DataPower Integration Appliance.

Question:

What was announced for DataPower on October 29, 2013?

Answer:

DataPower firmware v5.0 became available for the zBX Model 002 effective November 30, 2013.

Some new functions in firmware v5.0 include:

- New support for WebSphere Service Registry and Repository (WSRR) subscription can help to distinguish similar Saved Search Queries and support automatic synchronization and enforcement between WSRR and DataPower . This is designed to provide more consumable and centralized service level agreement (SLA) management. DataPower XI50z V5.0 appliances support the IETF Open Authorization (OAuth) 2.0 protocol. Using the OAuth protocol helps to decrease the need to share your credentials with third parties. It is designed to provide an authorization service separate and apart from the resource owner. OAuth is focused on the emerging Web 2.0 infrastructure and the popularity of APIs that exist to provide customizable access to an organization's applications. For example, eBay provides an API to provide enhanced shopping experiences by integrating with third-party applications; Twitter and Facebook provide APIs that extend their applications by providing content sharing capabilities. Each of these integrations requires focused attention on all aspects of security and the need to consider all access to be untrusted until proven otherwise.
- AAA expanded to support other service objects: Authentication, Authorization, Accounting (AAA - "triple A") for user or entity security is a framework within the WebSphere DataPower firmware. DataPower takes advantage of AAA extensively to support the OAuth 2.0 protocol. AAA is used to authenticate both the resource owner's and OAuth client's identities. It is also used for authorizing a request.
- In release 3.8.1, DataPower introduced form-based authentication, which is tied closely with web application firewall. As of firmware V5.0 AAA has been expanded to support other service objects. DataPower XI50z V5.0 is designed to act to protect access to resources when defined as a Policy Enforcement Point (PEP) for a resource server receiving and authorizing OAuth 2.0 requests. DataPower XI50z V5.0 firmware is also designed to improve processing power with extended memory support removing some of the limitations for large files.

For additional information, the release notes for DataPower 5.0 are available at: <http://www-01.ibm.com/support/docview.wss?uid=swg27027516> Select WebSphere DataPower Integration Appliance.

The DataPower XI50z V5.0 firmware was previously made available for the zBX Model 003.

Question:

What was the SOD that was announced July 23, 2013 for DataPower XI50z and the zBX?

Answer:

IBM intends to introduce a version of the WebSphere DataPower Integration Appliance XI52 Virtual Edition for use within the IBM zEnterprise BladeCenter Extension (zBX). IBM intends this virtual appliance to run on System x blades installed within a zBX. This program is designed to provide industry-leading security, integration, and optimization capabilities for System z, while offering similar management capabilities to the physical WebSphere DataPower appliance models.

Question:

Is the Application Optimizer (AO) included in the XI50z license?

Answer:

Yes. It is standard with XI50z.

Question:

Where can I get information on DataPower XI50z?

Answer:

There is a set of System z software FAQs with DataPower XI50 information: http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ST&infotype=SA&appname=STGE_ZS_ZS_USEN&htmlfid=ZSQ03052USEN&attachment=ZSQ03052USEN.PDF

zBX Networking

Question:

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

Answer:

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

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

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

Question:

What new capability is available for connectivity to a customer network?

Answer:

Prior to December 16, 2011, the zBX top of rack switches supported a 10 Gigabit Ethernet environment adjacent to the customer network. Some of our customers did not have a 10 Gbps network defined and therefore were unable to set up connectivity to the zBX. To help satisfy the requirements of those environments with 1 Gbps Ethernet, IBM made available to the external ports of the TORs optics for both single mode (LX) and multimode (SX) fiber optic cabling environments. These optics are referred to as 1000BASE-LX (feature 0634) and 1000BASE-SX (feature 0635). Note that if you use one of these features, you can only connect the IEDN to the external customer network and not to OSX CHPIDs.

Question:

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

Answer:

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

Question:

What features/functions are needed to attach a zEC12 or zBC12 to a zBX?

Answer:

To attach the zEC12 (or zBC12), to a zBX, you need two OSA-Express5S, OSA-Express4S (carry forward only) or OSA-Express3 (carry forward only) 1000BASE-T Ethernet features with one port on each defined as CHPID type OSM.

Question:

If I upgrade the zBX and it expands to a second rack, do I order additional OSA adapters to connect to it?

Answer

No, the connectivity you established during the installation of the first rack is sufficient.

Question:

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

Answer:

Yes, you can use externally facing ports on the TOR switches in the zBX and OSA adapters on the zEC12 to attach directly to your external network.

Question:

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

Answer:

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

Question:

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

Answer:

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

Question:

With the October 12, 2011 announcement, why did IBM announce going from four to twelve SAN connections per BladeCenter chassis within the zBX (FC#0615)?

Answer:

We increased the number of SAN connections per BladeCenter chassis to enable additional bandwidth to SAN-attached storage on a BladeCenter by BladeCenter basis where needed by customer applications. Use of the additional connections is optional.

Question:

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

Answer:

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

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

Question:

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

Answer:

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

zBX Availability

Question:

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

Answer:

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

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

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

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

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

Question:

What can you tell me about the GDPS enhancements that were in the July 23, 2013 announcement?

Answer:

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

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

This is designed to extend GDPS/PPRC continuous availability and disaster recovery capabilities to virtual servers in a zBX for planned and unplanned outages. The planned availability is fourth quarter 2013 with GDPS/PPRC v3.10 SPE.

Question

Will zBX frames have the same availability as the classic zEnterprise server?

Answer

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

Question:

Is there redundancy in the blade hardware itself?

Answer:

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

Question:

What availability connection exists between the zEC12 (or zBC12) and the zBX? Can the zBX continue to run if there is no host (for example, during a Power-on-Reset)?

Answer:

Yes the zBX can continue to run without the host (and, of course, vice versa.) However, the zBX may lose the management "guidance" provided by the zEC12 (or zBC12) for configuration, operations, performance and reporting support. Other zEnterprise servers in the ensemble that includes the zBX can access the zBX even if the host isn't available because their connection to the zBX is via the private data network (IEDN).

Question:

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

Answer:

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

Hardware redundancy at various levels

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

Question:

What are the availability characteristics of the networks that connect the zEC12 (or zBC12) and the zBX?

Answer:

There are two connections between a zEnterprise and a zBX. Both have excellent availability features.

- A private and physically isolated internal management network (the INMN) connects all zEnterprise System resources (CPCs, BladeCenters, blades, top-of-rack switches, power distribution units, and hypervisors) for management purposes. This network is factory pre-wired, internally switched, configured, and managed with full redundancy for high availability. These characteristics enable simplified and secure management with no effect on customer data traffic.
- A private and secure intra-ensemble data network (the IEDN) connects all elements of a zEnterprise System ensemble for normal application interconnection purposes. This network is access-controlled using integrated virtual local area network (VLAN) provisioning. This requires no external customer-managed switches or routers, which helps to reduce the need for firewalls and encryption, and simplifies network configuration and management, while providing full redundancy for high availability. The management of the network provides enforcement of strict access control across heterogeneous environments, further augmenting security and simplicity.

Question:

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

Answer:

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

zBX and blade Warranty and Maintenance

Question:

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

Answer:

The warranty for the zBX is a traditional IBM System z warranty - one year PD (problem determination) and 24x7 onsite FRU (field replaceable unit) service performed by a zSSR (System z Service Support Rep).

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

Question:

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

Answer:

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

Question:

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

Answer:

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

Question:

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

Answer:

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

Question:

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

Answer:

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

Question:

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

Answer:

You should provide your IBM sales representative with an order letter for the type and quantity of blades you plan to install in the zBX. It is important for your letter to specify the number of node enablement features on the zEnterprise and the number of zBX blade feature codes (the number will be the same).

Question:

I don't have IBM maintenance on my zEnterprise (zEC12, zBC12, z196 or z114). How can I order maintenance for the zBX?

Answer:

If you don't have IBM maintenance on the zEnterprise, or if you don't plan to have it on the zBX, you will need to purchase HMC Feature Code #0090 in order for the IBM zSSR to load updates to the zBX.

Question:

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

Answer:

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

Question:

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

Answer:

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

Question:

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

Answer:

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

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

IBM zEnterprise Unified Resource Manager

Question:

What is the new Ensemble Availability Management function that was announced July 23, 2013 for zEnterprise Unified Resource Manager?

Answer:

Ensemble Availability Management offers the ability to create user-defined availability policies for managing the availability virtual servers, along with monitoring and reporting functions to help ensure virtual servers are executing in line with those policies. Defining your own customer workloads (by name), their business priorities and the availability management capabilities allows differentiation between multiple workloads in an ensemble. Workload availability status represents whether all workload elements are operating and available. Element Groups may be configured to manage virtual servers with redundancy objectives, enabling availability and capacity management among redundant virtual servers that provide the same function.

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

Question:

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

Answer:

The Unified Resource Manager is designed to provide systems management capabilities across the multi-architecture environment of zEnterprise. Many of our clients are planning to exploit these system management capabilities to deploy a framework for a heterogeneous cloud environment, thereby providing an effective means to help deliver IT services. To accelerate the delivery of its cloud offerings, IBM recently announced plans to base all of its cloud services and software on open architecture and standards, including OpenStack. OpenStack is an infrastructure as a service (IaaS) cloud computing open source project. IBM joined the project in 2012 and in support of 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 System x hypervisor based on Kernel-based Virtual Machine (KVM). Looking ahead, IBM will continue to enable OpenStack to provide heterogeneous systems management across zEnterprise, z/VM, and distributed platforms, which in turn can be exploited by subsequent IBM SmartCloud offerings. These offerings are designed to provide you enablement for enterprise-wide cloud deployments and greater flexibility by removing the need to develop specific interfaces for different cloud services.

Question:

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

Answer:

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

Additional information can be found in the zEnterprise Unified Resource Manager Redbook (www.redbooks.ibm.com) or in the white books in the library available on Resource Link¹⁵ (www.ibm.com/servers/resourceLink).

Question:

What is Dynamic Storage Provisioning?

Answer:

Dynamic Storage Provisioning allows virtual disks to be added/allocated to a running virtual server without having to take the virtual server down. This ability is available for virtual servers running on the Power and x86 blades on both the zBX mod 002 (z196/z114) and mod 003 (zEC12/zBC12). Note that the guest operating system must support the dynamic addition of virtual based disks to use dynamic storage provisioning.

Question:

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

Answer:

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

Question:

What is the philosophy of the zEnterprise Unified Resource Manager?

Answer:

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

¹⁵ A userid and password are required to access Resource Link – if you don't have one then follow the instructions to set one up.

Question:

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

Answer:

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

Question:

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

Answer:

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

When Unified Resource Manager is executed, you have the ability to define a workload as the set of virtual resources required and defined to support a business process. The workload can include z/OS, z/VM, Linux on System z, and AIX on POWER7. New on July 23, 2013, is the ability to manage resource optimization for virtual servers running on IBM BladeCenter HX5 (7873) blades in the zBX. A performance policy can be attached to the workload that defines its service-level requirements. Similar to Workload Manager on z/OS, if goals are not being met the Unified Resource Manager will take appropriate corrective action.

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

Question:

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

Answer:

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

Question:

In the Unified Resource Manager I define service objectives, and the Unified Resource Manager monitors the systems and can "perform dynamic CPU movement when needed." Does this mean that a software task can migrate to a different processor?

Answer:

Initially, the Unified Resource Manager will dynamically adjust the allocation of processor capacity in Linux under z/VM and POWER hypervisors to balance the allocation of processing capacity across the virtual guests hosted by that hypervisors. This is not movement of the workload, but rather a movement of the processor capacity supporting the workloads on each system. In future, the Unified Resource Manager may exploit live guest relocation capabilities to actually move running virtual guests from one system to another. With the July 23, 2013 announcement, support for balancing processing capacity has been extended to the System x hypervisor (zBX Model 003 Driver level 15 only).

Question:

What is the basic functional difference between the Manage, Advanced Management and the Automate suites?

Answer:

Manage is the function of Unified Resource Manager that delivers core operational controls, installation support and configuration. Advanced Management for select System x blades allows setting up business workloads and associating resources with a business process, as well as monitoring and reporting performance and performing energy management. Automate, prior to the July 23, 2013 announcement, for select POWER7 blades and DataPower XI50z provided energy management. On those POWER7 blades, Automate also allowed setting up business workloads and associating resources with a business process, as well as monitoring and reporting performance and performing energy management in addition to setting up workload goals, dynamic resource provisioning, and optimization. For zEC12 processors at driver level 15, and all zBC12 processors, there is no 'Advanced Management' suite. Support for CPU management for System x blades was added to the Automate firmware in the July 23, 2013 announcement.

Question

Is the Unified Resource Manager included in the base price of the zEnterprise?

Answer:

No, the Unified Resource Manager functions are organized into three suites – Manage, Advanced Manage¹⁶ and Automate. Each node of resources (for example, the z196, an IFL, a POWER7 blade, System x blade or DataPower XI50z) will have an associated price for manage, advanced management or automate. There is additional information in the pricing section.

For the zBC12 there will only be two tiers of management – Manage and Automate. This is because the CPU management for System x blades is now part of the Automate tier for the zBC12. Advanced Manage did not support CPU management for System x blades.

Question:

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

Answer:

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

Question:

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

Answer:

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

¹⁶ System x blades on earlier zEC12 servers will not have full Automate capabilities until they are upgraded to driver level 15. Prior to that level they can only have Advanced Management functions including: automate capabilities for System x blades Wizard-driven set up of resources in accordance with specified business process; ability to monitor and report performance; Load balance recommendations

Question:

What is the guest platform management provider (GPMP)?

Answer:

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

Question:

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

Answer:

The GPMP does not run in a Linux native LPAR.

Question:

How does z/OS work with Unified Resource Manager?

Answer:

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

Question:

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

Answer:

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

Ensemble

Question:

What is an ensemble?

Answer:

An ensemble is a collection of one or more zEnterprise System nodes in which each node comprises a zEnterprise System (zEC12, zBC12, z196 or z114) and its optionally attached zBX (Model 003 for zEC12 and zBC12, and Model 002 for z196 or z114). An ensemble can consist of a single zEnterprise System with no zBX attached, or two to eight zEnterprise Systems where at least one of the zEnterprise Systems has a zBX attached. The resources of a zEnterprise System ensemble are managed by Unified Resource Manager, and virtualized as a single pool of resources integrating system and workload management across the multisystem, multitier, multi-architecture environment.

Question:

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

Answer:

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

Question:

Where is the code to set up an ensemble?

Answer:

You will need to order FC#0025 on the zEC12, zBC12, z196 or z114, and then use customer interfaces of the Unified Resource Manager on the HMC to create an ensemble.

Question:

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

Answer:

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

Question:

Is Tivoli® software mandatory when setting up an ensemble?

Answer:

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

Question:

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

Answer:

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

Question:

Must I have a zBX on each zEnterprise System in the ensemble?

Answer:

No, this is not a requirement. However, if there are multiple nodes in the ensemble, at least one of them must include a zBX.

Question:

Can I have an ensemble with a single zEnterprise System with no zBX attached?

Answer:

Yes. The single CEC must have two OSA-Express5S, OSA-Express4S (carry forward only), or OSA-Express3 (carry forward only) 1000BASE-T features (OSM) and two OSA-Express5S, OSA-Express4S (carry forward only) or OSA-Express 3 (carry forward only) 10 GbE features (OSX) (with loopback cables), in addition to Unified Resource Manager.

Question:

Can a System z10 be part of an ensemble?

Answer:

No.

Question:

If I have one zEnterprise System (zEC12, zBC12, z196 or z114) with a zBX in an ensemble and I want to allow another zEnterprise System to be able to access data on the zBX and not be in an ensemble, can I use my existing network to do that, without having to be in the ensemble?

Answer:

Yes, in fact that would be how a System z10 would access the zBX data, since a System z10 cannot be in an ensemble. So, if the second zEnterprise System is not in an ensemble then it can connect to the zBX using your existing network. If the second zEnterprise System is in a separate ensemble then it can connect to the zBX using your existing network but would not be able to define workloads that include virtual servers in the zBX. However, if the second zEnterprise System is in the same ensemble as the first, it should be connected to the zBX using the intra-ensemble data network.

Question:

Can I have an ensemble with a zEC12 (or zBC12) and a zBX Model 003, and attach z196 or z114 servers? Can I do that kind of mix and match?

Answer:

Yes, workloads can be defined with resources on both a z196 and the zBX Model 003, but the zBX Model 003 can only be connected to the zEC12 or zBC12.

Question:

Can IBM DB2 Analytics Accelerator for z/OS be a part of an ensemble?

Answer:

No, but it could be connected to a member of an ensemble. In that case, ensemble nodes with z/OS LPARs that are in the same DB2 data sharing group as the one with the attached IBM DB2 Analytics Accelerator can access it.

Question:

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

Answer:

Yes, z/VM V6.2 and its virtual servers can be part of a zEnterprise ensemble and can be managed by the Unified Resource Manager. But, z/VM V6.3 and its virtual servers cannot be managed by Unified Resource Manager and if you plan to use the VMSSI feature, and especially if you plan to use its Live Guest relocation capability, then the best practice is to not use Unified Resource Manager to manage z/VM virtual servers.

Question:

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

Answer:

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 System x hypervisor based on Kernel-based Virtual Machine (KVM).

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

Hardware Management Console (HMC)

Question:

What was announced July 23, 2013 about managing my OSA Express cards on the HMC?

Answer:

Support for OSA-Express5S and OSA-Express4S features is new in the OSA Advanced Facilities on the HMC and is exclusive to the zEC12 and the zBC12. The latest driver level is required. "OSA/SF on the HMC" provides configuration, validation, activation, and display support for selected functions of the OSA-Express5S and OSA-Express4S.

OSA Advanced Facilities on the HMC ("OSA/SF on the HMC") is required for the OSA-Express5S features operating as CHPID types of OSD, OSE, or OSN. Either OSA/SF on the HMC or the existing OSA/SF operating system component under z/OS, z/VM, and z/VSE can be used for the OSA-Express4S feature configured with these CHPID types. Note that OSA/SF on the HMC plays no role in the configuration and management of CHPID types OSC, OSM, or OSX.

Detailed information is available in the following manual available from ResourceLink: Open Systems Adapter/Support Facility on the Hardware Management Console (SC14-7580)

Question:

Are the primary HMC and backup HMC that manage the ensemble exclusive to zEC12 or zBC12 or can they also be used as 'regular' HMCs?

Answer:

The firmware to manage the ensemble is an application running on the HMC. That HMC can be used for other 'regular' HMC functions. The backup HMC has no function other than to mirror the ensemble state information from the primary. Accordingly, no other function can execute on the backup HMC.

Question:

If I have a primary ensemble HMC and an alternate ensemble HMC, and the primary fails does the system automatically migrate to the alternate? How am I notified – is there a 'call home' service call generated? Do I need to immediately assign another alternate ensemble HMC or is it alright to run without a backup until I get the primary fixed and put it back in use as my alternate?

Answer:

If the primary ensemble HMC fails, the alternate ensemble HMC that has been configured for the ensemble automatically takes over and assumes the primary ensemble HMC role. This is done to maintain availability of the functions performed by or through the primary ensemble HMC. A call-home service call is generated to notify IBM of this event and the operator is notified by a hardware message. The ensemble HMC that has taken over the primary role will continue to operate in that role without a functional alternate ensemble HMC, which is acceptable for a short period of time. It is not necessary to immediately assign a new alternate ensemble HMC. However, the environment is at increased risk of unavailability due to the lack of redundancy. Therefore IBM recommends that the failed HMC be repaired promptly and then brought back into use (now as the alternate ensemble HMC) so that redundancy is restored.

Question:

If I have to take my primary HMC out of service are there switch "roles" functions that I must initiate and how long does it take to complete?

Answer:

Yes, you can initiate a role switch to have the alternate ensemble HMC take over the primary role so that the (before the switch) primary can be taken out of service. The switch is initiated using a function of the Manage Alternate HMC task. The switch process is initiated very quickly but involves an (automatic) reboot of the alternate ensemble HMC and thus does not complete until that reboot has finished. The full process typically requires several minutes to complete.

Question:

Can a Primary ensemble HMC and Alternate ensemble HMC participate in mirroring with other HMCs at the same Driver level?

Answer:

Mirroring is a special synchronization function that copies all important configuration and state data from the primary ensemble HMC to the alternate so that the alternate has the full state it needs in order to take over the primary role if necessary. This mirroring occurs only between designed primary and alternate ensemble HMCs for an ensemble. No other HMCs (regardless of driver level) can participate in this mirroring. However, the HMC Data Replication feature can be used to replicate many categories of configuration data between HMCs, including between an ensemble HMC and other HMCs. Replication relationships between HMCs are set up using the Configure Data Replication task.

Question:

What is the distance permitted between the primary ensemble HMC and the alternate ensemble HMC? Can it be in an alternate location?

Answer:

There is no specific physical limit on the allowable distance between the primary and alternate ensemble HMCs for an ensemble. However, there is a requirement that the IP addresses of both the primary and alternate ensemble HMC be on the same TCP/IP subnet. This same-subnet requirement means that the scope of the subnet to which the ensemble HMCs are attached will, in practice, be a limiting consideration. While some customers might have configurations where the raised-floor network to which the ensemble HMCs are attached spans locations or sites, this is not typically the case. Thus, in most cases, the primary and alternate ensemble HMCs will both have to be in the same location/site.

Question:

Where will I see the humidity and altitude readings?

Answer:

The humidity and altitude readings are both shown on the SAD and CPC detail panels of the HMC.

Question:

Who will need access to the ensemble HMC console?

Answer

Due to the more significant management functions provided by the ensemble HMC you need to pay much more attention to the "defined roles" in the HMC. The HMC includes security facilities that can be used to isolate different roles (for example, storage manager, network manager, virtualization manager, systems programmer, operator) from each other.

Question:

Can the different functions of the ensemble HMC be secured by using LDAP services?

Answer

Yes, all functions can be secured by the use of an LDAP server. It is recommended that z/OS (and its IBM Tivoli Directory Server function) be used as the LDAP server for the ensemble HMC.

Question:

Do zBXs have their own HMC consoles?

Answer:

No, the IBM CE uses the ensemble HMC to apply firmware maintenance to the blades in the zBX nodes as well as to the zEnterprise elements. There is one exception. If your zEnterprise or zBX are not on an IBM Maintenance Agreement then you will need to order HMC feature 0092 in order for IBM to load updates to the zBX.

Question:

Can access to different functions of the HMC and ensemble HMC be audited?

Answer:

Yes, many log functions have been added to the HMC to allow a security administrator to audit every access to HMC functions.

Question:

What role does the Support Element play in the connections between the zEnterprise and the zBX?

Answer:

The Support Element (SE) is a management appliance that has been extended to provide resource management for the zBX that may be associated with a host zEnterprise. Customer workloads do not run in or through the SE. From a workload perspective, the blade elements are accessed from other systems within the ensemble over the IEDN network. The blades are installed in the zBX, which is connected to the zEnterprise through two networks: the intra-ensemble data network (IEDN) and the intra-node management network (INMN). Both networks are connected through a pair OSA Express5S cards to the zEnterprise.

Note:

All new zEC12 enhancements for July 23, 2013, will require Driver 15, available on September 20, 2013.

zBC12 Upgradeability

Question:

What are the upgrade paths from a z114 or a z10 BC to the zBC12?

Answer:

All models of the z114 and the z10 BC can upgrade directly to a zBC12.

Question:

Is there an upgrade path from a zBC12 Model H06 to Model H13?

Answer:

Yes. But careful planning should be done as it is a disruptive upgrade due to the additional drawer.

Question:

Is there an upgrade path from a zBC12 to zEC12?

Answer:

Yes. There is an upgrade path from the Model H13 to the zEC12 Model H20.

Question:

Will I be able to upgrade from a full capacity zBC12 to a subcapacity zBC12?

Answer:

Yes. Each of the first six general purpose processors on the zBC12 can be defined as full capacity or one of twenty-five subcapacity settings. This creates a 26 by 6 matrix of settings. As long as upgrades are positive capacity growth, you can move around anywhere within the matrix when adding capacity. Additional processors beyond these six, i.e. specialty engines, are always defined as full capacity.

zEC12 Upgradeability

Question:

What are the upgrade paths from a z196 or a z10 EC to the zEC12?

Answer:

All models of the z196 and the z10 EC can upgrade directly to a zEC12.

Question:

Is there an upgrade path from a zBC12 to zEC12?

Answer:

Yes. There is an upgrade path from the Model H13 to the zEC12 Model H20.

Question:

Can I upgrade any zEC12 to a Model HA1?

Answer:

Yes, but it is important to note that upgrading other zEC12 models to the HA1 will be disruptive due to the difference between MCMs on the HA1 and those on other zEC12 models.

Question:

Will I be able to upgrade from a full capacity zEC12 to a subcapacity zEC12?

Answer:

Yes. Each of the first twenty general purpose processors on the zEC12 can be defined as full capacity or one of three subcapacity settings. This creates a 20 by 4 matrix of settings. As long as upgrades are positive capacity growth, you can move around anywhere within the matrix when adding capacity. When your number of general-purpose processors exceeds twenty, then all of the general-purpose processors must be full capacity.

Question:

Can I have subcapacity processors if I have more than one book?

Answer:

Yes, remember that only twenty subcapacity processors can be active on the server no matter what model you have installed. If more than twenty CPs have been purchased on a server, a selection can be made to activate only 20 or fewer subcapacity features. This means that the subcapacity settings are available on any of the models as long as they are configured (not the same as purchased) with twenty or fewer general-purpose processors.

Software Pricing – zBC12

Question:

What software pricing is announced for a standalone zBC12?

Answer:

IBM announced a new Technology Transition Offering (TTO) called Technology Update Pricing for Advanced Entry Workload License Charges (AEWLC). Technology Update Pricing for AEWLC uses the reporting mechanisms and existing MSU tiers of the AEWLC price metric while increasing the software price performance improvements provided by AEWLC.

Technology Update Pricing for AEWLC applies only to eligible z/OS, z/TPF, and z/VSE operating systems and their associated middleware programs when running on a single, stand-alone zBC12 server.

Question:

What is the price reduction available with the Technology Update Pricing for AEWLC available for the zBC12?

Answer:

The percent reduction in the monthly AEWLC is based on the number of zBC12 full capacity MSUs. While the percent reduction is only based on full capacity MSUs, you remain eligible for sub-capacity pricing under business as usual terms. AEWLC pricing for a standalone zBC12 is reduced by the percentage in the table below.

<i>Schedule of AEWLC reductions for Technology Update Pricing for AEWLC</i> MSUs: Number of zBC12 Full Capacity MSUs for a stand-alone zBC12	Reduction in Monthly AEWLC
1 – 7 MSUs	5.0%
8 - 17 MSUs	5.0%
18 – 30 MSUs	5.0%
31 – 45 MSUs	5.0%
46 – 87 MSUs	4.0%
88 – 175 MSUs	4.0%
176 – 260 MSUs	4.0%
261 – 315 MSUs	4.0%
316 – 390 MSUs	4.0%
391 and more MSUs	4.0%

The number of MSUs used to determine the MSUs of a stand-alone zBC12 server are based on the announced IBM full capacity ratings that can be found on the Mainframe Exhibits section of the System z Software Contracts website at:

<http://ibm.com/systems/z/swprice/reference/exhibits/hardware.html>

Question:

If I upgrade from a System z10 BC to zBC12 will software pricing improvements be additive?

Answer:

Yes. You will receive the pricing advantage of AEWLC on the z114, and you will also receive the Technology Update Pricing for AEWLC for the zBC12.

Question:

What software pricing announcements were made for a zBC12 in a Parallel Sysplex?

Answer:

IBM announced revisions to the Technology Update Pricing for Advanced Workload License Charges (AWLC) offering and two revised Transition Charges for Sysplexes offerings.

The revisions to the Technology Update Pricing for AWLC offering and the two Transition Charges for Sysplexes offerings apply only to eligible z/OS and z/TPF operating systems and their associated middleware programs, when running in an aggregated Parallel Sysplex or a Loosely Coupled Complex, respectively.

The Technology Update Pricing for AWLC offers actively coupled Parallel Sysplexes and Loosely Coupled Complexes, that consist entirely of zBC12 servers and IBM zEnterprise EC12 (zEC12) servers, the benefit of aggregated Technology Update Pricing for AWLC.

The Transition Charges for Sysplexes offers actively coupled Parallel Sysplexes and Loosely Coupled Complexes with different generations of server hardware technology partial benefits of AWLC or partial benefits of Technology Upgrade Pricing for AWLC in proportion to the percentage of zBC12 server and other IBM zEnterprise server capacity in the sysplex or complex.

Question:

Will I still be able to recognize the benefits of sub-capacity pricing with AEWLC?

Answer:

Yes. AEWLC allows you to report and pay for software based on sub-capacity charges using the existing SCRT process. AEWLC also allows for full-capacity pricing based on the rated MSUs of your zBC12 server.

Question:

What contracts are required for Technology Update Pricing for AEWLC?

Answer:

In order to get Sub-Capacity Technology Update Pricing for AEWLC, the following contract is required to be executed:

ICA Attachment for IBM System z AEWLC, (Z125-8755).

The Supplement for Technology Transition Offerings to Attachment for IBM System z Advanced Workload License Charges and to Attachment for IBM System z Advanced Entry Workload License Charges and to IBM System z Machines Exhibit (Z125-8994) describes the Technology Transition Offerings and should be delivered to each zBC12 customer, but it does not get executed.

The System z Machines Exhibit (Z125-3901) provides the terms for Full Capacity AEWLC. The Exhibit and the Supplement should be delivered to each zBC12 customer, but they do not get executed.

Question:

What contracts are required for zBC12 machine participation in the two Transition Charges for Sysplexes offerings?

Answer:

In order for a machine to participate in either of the Transition Charges for Sysplexes offerings, the following contract is required to be executed:

ICA Attachment for IBM System z AWLC, (Z125-8538).

The Supplement for Technology Transition Offerings to Attachment for IBM System z Advanced Workload License Charges and to Attachment for IBM System z Advanced Entry Workload License Charges and to IBM System z Machines Exhibit (Z125-8994) documents the Technology Transition Offerings and should be delivered to each zBC12 customer, but it does not get executed.

Additional contracts are required for WLC if applicable.

Question:

What is the Processor Value Unit (PVU) rating used for Passport Advantage® software on the zBC12?

Answer:

The PVU rating will be 100 for both IFL engines and CP engines, the same as the z114.

Question:

How will the IBM POWER7 PS701 or BladeCenter HX5 (System x) blades and supported AIX, Linux and Windows operating systems be priced?

Answer:

The pricing for IBM POWER7 or IBM HX5 Blades and the supported AIX, Linux and Windows operating system are purchased (and licensed) through your standard qualified channels. System z does not price or ship these as part of the zBX order.

Question:

Where can I get more information on IBM software charges?

Answer:

Please refer to www.ibm.com/systems/z/swprice/

Software Pricing – zEC12

Question:

What software pricing is announced with the zEC12?

Answer:

Three Technology Transition Offerings are being announced. They are Technology Update Pricing for AWLC and two revised Transition Charges for Sysplexes programs.

Technology Update Pricing for AWLC uses the reporting mechanisms and existing MSU tiers of the AWLC price metric while increasing the software price performance improvements provided by AWLC.

The three Technology Transition Offerings apply only to eligible z/OS and zTPF operating systems and their associated middle ware programs when running on a zEC12 server or in an aggregated Parallel Sysplex with a zEC12.

Question:

What software pricing is available on a standalone zEC12 or a Parallel Sysplex with all zEC12 technology?

Answer:

For a standalone zEC12, or a Parallel Sysplex with all zEC12, IBM will offer the Technology Update Pricing for AWLC which leverages the existing AWLC pricing metric while offering price performance improvements for zEC12. The Technology Pricing for AWLC provides reduced AWLC pricing on the zEC12 as compared to AWLC on the z196 with greater than 3 MSUs. The Technology Update Pricing for AWLC is one of three Technology Transition Offerings announce with the zEC12.

Note: If choose to do so, you may continue to use PSLC pricing with the zEC12.

Question:

What is the price reduction available with the Technology Update Pricing for AWLC available for the zEC12?

Answer:

The percent reduction in the monthly AWLC is based on the number of zEC12 full capacity MSUs. While the percent reduction is only based on full capacity MSUs, you remain eligible for sub-capacity pricing under business as usual terms. AWLC pricing for a standalone zEC12 or Parallel Sysplex of all zEC12 servers is reduced by the percentage in the table below.

Schedule of AWLC reductions for Technology Update Pricing for AWLC (TU1)

MSUs: Number of zEC12 Full Capacity MSUs for a stand-alone zEC12, or the sum of the Full Capacity MSUs in an actively coupled Parallel Sysplex made up entirely of zEC12 servers	Reduction in Monthly AWLC
4 – 45 MSUs	2.0%
46 - 315 MSUs	4.0%
316 – 1315 MSUs	4.5%
1316 – 2676 MSUs	5.0%
2677 – 5476 MSUs	6.0%
5477 or more MSUs	7.0%

The number of MSUs used to determine either the MSUs of a stand-alone zEC12 server or the total

MSUs of a zEC12 Parallel Sysplex are based on the announced IBM full capacity ratings that can be found on the Mainframe Exhibits section of the System z Software Contracts website at <http://ibm.com/systems/z/swprice/reference/exhibits/hardware.html>

Question:

If I add a zEC12 in a Parallel Sysplex with one or more z196 and/or z114 servers, but no z10™ servers, can I benefit from the new Technology Update Pricing for AWLC?

Answer:

No. AWLC pricing may be used across a Parallel Sysplex when all the systems are zEC12, z196, or z114. The AWLC Transition Charges for Sysplexes (TC2) provides a reduction in AWLC pricing based on the percentage of the full capacity MSUs that are zEC12. When you migrate to all zEC12 servers, the Technology Update Pricing for AWLC may apply.

The AWLC Transition Charges for Sysplexes (TC2) is one of the three new Technology Transition Offerings announced with the zEC12.

Question:

What AWLC price reduction is provided by the Transition Charges for Sysplexes (TC2) program when I have a zEC12 in a Parallel Sysplex with z196 and z114 (and no z10 servers)?

Answer:

The AWLC price reduction provides a portion of the price performance benefit related to AWLC pricing based on the proportion of zEC12 MSU's in the Parallel Sysplex. The number of MSUs used to determine both the total MSUs for the Parallel Sysplex and zEC12 server proportion of MSUs are based on the announced IBM full-capacity ratings

at <http://ibm.com/systems/z/swprice/reference/exhibits/hardware.html>

The reduction schedule is below. For example, if the percentage of zEC12 hardware in the fully qualified Parallel Sysplex is 40% the AWLC reduction will be 1.5%. If the percentage of zEC12 hardware in the fully qualified Parallel Sysplex is 75% the AWLC reduction will be 4.0%.

Schedule of AWLC reductions for Transition Charges for Sysplexes (TC2)

Percentage of zEC12 Full Capacity MSUs in the Sysplex	Reduction in Monthly AWLC
> 0 - 20%	0.5%
21 - 40%	1.5%
41 - 60%	3.0%
61 - 80%	4.0%
81 - < 100%	4.5%

Question:

If I have a zEC12 in a Parallel Sysplex with a z10 server, can I use AWLC pricing?

Answer:

No, you can't use AWLC pricing if there is a z10 server in the Parallel Sysplex, but the AWLC Transition Charges for Sysplexes (TC1) provides partial AWLC benefit. The Transition Charges for Parallel Sysplexes (TC1) provides aggregated VWLC pricing as well as a reduction in VWLC. VWLC pricing is not available to the zEC12, z196 or z114 except when using the AWLC Transition Charges for Sysplexes (TC1). When you migrate off of z10 servers, one of the other Technology Transition Offerings with AWLC pricing may apply.

The AWLC Transition Charges for Sysplexes (TC1) is one of the three new Technology Transition Offerings announced with the zEC12, but it is a modification of the previously announced AWLC Transition Charges for Sysplexes.

Question:

How is the VWLC price reduction provided by the AWLC Transition Charges for Sysplexes (TC1) program calculated?

Answer:

The VWLC price reduction provides a portion of the price performance benefit related to AWLC pricing based on the proportion of zEnterprise (zEC12, z196 and z114) MSU's in the Parallel Sysplex. The number of MSUs used to determine both the total MSUs for the Parallel Sysplex and zEnterprise server proportion of MSUs are based on the announced IBM full-capacity ratings at <http://ibm.com/systems/z/swprice/reference/exhibits/hardware.html>

The reduction schedule is below. For example, if the percentage of zEnterprise hardware in the fully qualified Parallel Sysplex is 37% the VWLC reduction will be 1.5%. If the percentage of zEnterprise hardware in the fully qualified Parallel Sysplex is 80% the VWLC reduction will be 4.0%.

Schedule of VWLC reductions for Transition Charges for Sysplexes (TC1)

Percentage of zEnterprise Full Capacity MSUs in the Sysplex	Reduction in Monthly VWLC
> 0 - 20%	0.5%
21 - 40%	1.5%
41 - 60%	3.0%
61 - 80%	4.0%
81 - < 100%	4.5%

Question:

If I upgrade from a System z10 EC to zEC12 will software pricing improvements be additive?

Answer:

Yes. You will receive the pricing advantage of AWLC on the z196, and you will also receive the Technology Update Pricing for AWLC for the zEC12.

Question:

Will I still be able to recognize the benefits of sub-capacity pricing with AWLC?

Answer:

Yes. AWLC allows you to report and pay for software based on sub-capacity charges using the existing SCRT process. AWLC also allows for full-capacity pricing based on the rated MSUs of your zEC12 server. All three of the Technology Transition Offerings may be used in a sub-capacity environment.

Question:

What contracts are required for AWLC pricing and the three Technology Transition Offerings?

Answer:

In order to get AWLC pricing or any of the Technology Transition Offerings, the following contract is required to be executed:

- Attachment for IBM System z Advanced Workload License Charges (Z125-8538)

The additional benefits of the Technology Transition Offerings are provided in a supplement called the Supplement for Technology Transition Offerings to Attachment for IBM System z Advanced Workload License Charges (Z125-8994). This supplement does not require a signature.

Question:

What is the Processor Value Unit (PVU) rating used for Passport Advantage software on the zEC12 and zBC12?

Answer:

The PVU rating for the zEC12 will be 120 for both IFL engines and CP engines, the same as the z196 and the z10 EC. The PVU rating for the zBC12 will be 100 for both IFL engines and CP engines, the same as the z114 and z10 BC.

Question:

How will the IBM POWER7 PS701 or BladeCenter HX5 (System x) blades and supported AIX, Linux and Windows operating systems be priced?

Answer:

The pricing for IBM POWER7 or IBM HX5 Blades and the supported AIX, Linux and Windows operating system are purchased (and licensed) through your standard qualified channels. System z does not price or ship these as part of the zBX order.

Question:

Where can I get more information on IBM software charges?

Answer:

Please refer to: www.ibm.com/systems/z/swprice/

Performance

Question:

What are the major changes to the z/OS V1R13 LSPR

Answer:

The LSPR ratios reflect the range of performance between System z servers as measured using a wide variety of application benchmarks. The latest release of LSPR continues with the methodology introduced with the z/OS V1R11 LSPR. Prior to that version, workloads had been categorized by their application type or software characteristics (for example, CICS, OLTP-T, LoIO-mix). With the introduction of CPU MF (SMF 113) data starting with the z10 processor, insight into the underlying hardware characteristics that influence performance was made possible. The LSPR defines three workload categories, LOW, AVERAGE, HIGH, based on the metric called “Relative Nest Intensity (RNI)” which reflects a workload’s use of a processor’s memory hierarchy. For details on RNI and the workload categories, please reference the LSPR documentation or go to <https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lspindex>

Question:

What is the multi-image table in the LSPR?

Answer:

Typically, IBM System z processors are configured with multiple images of z/OS. Thus, the LSPR continues to include a table of performance ratios based on average multi-image z/OS configurations for each processor model as determined from the profiling data. The multi-image table is used as the basis for setting MIPS and MSUs for IBM System z processors.

Question:

What multi-image configurations are used to produce the LSPR multi-image table?

Answer:

A wide variety of multi-image configurations exist. The main variables in a configuration typically are: 1) number of images, 2) size of each image (number of logical engines), 3) relative weight of each image, 4) overall ratio of logical engines to physical engines, 5) the number of books and 6) the number of ICFs/IFLs. The configurations used for the LSPR multi-image table are based on the average values for these variables as observed across a processor family. It was found that the average number of images ranged from 5 at low-end models to 9 at the high end. Most systems were configured with 2 major images (those defined with >20% relative weight). On low- to midrange models, at least one of the major images tended to be configured with a number of logical engines close to the number of physical engines. On high-end boxes, the major images were generally configured with a number of logical engines well below the count of physical engines reflecting the more common use of these processors for consolidation. The overall ratio of logical to physical engines (often referred to as “the level of over-commitment” in a virtualized environment) averaged as high as 5:1 on the smallest models, hovered around 2:1 across the majority of models, and dropped to 1.3:1 on the largest models. The majority of models were configured with one book more than necessary to hold the enabled processing engines, and an average of 3 ICFs/IFLs were installed.

Question:

Can I use the LSPR multi-image table for capacity sizing?

Answer:

For high-level sizing, the multi-image table may be used. However, the most accurate sizings require using the zPCR tool's *LPAR Configuration Capacity Planning* function, which can be customized to exactly match a specific multi-image configuration rather than the average configurations reflected in the multi-image LSPR table.

Question:

What model is used as the “base” or “reference” processor in the z/OS V1R13 LSPR table?

Answer:

The 2094-701 processor model is used as the base in the z/OS V1R13 table. Thus, the ITRR for the 2094-701 appears as 1.00.

Note that in zPCR the reference processor may be set at the user's discretion.

Question:

What “capacity scaling factors” are commonly used?

Answer:

The LSPR provides capacity ratios among various processor families. It has become common practice to assign a capacity scaling value to processors as a high-level approximation of their capacities. The commonly used scaling factors can change based on the version of LSPR. For z/OS V1R13 studies, the capacity scaling factor commonly associated with the reference processor set to a 2094-701 is 593 which is unchanged from that used with z/OS V1R11. This value reflects a 2094-701 configured with a *single image* of z/OS - no complex LPAR configuration (i.e., multiple z/OS images) effects are included. For the z/OS V1R13 multi-image table the commonly used scaling factor is $.944 \times 593 = 559.792$. Note the .944 factor reflects the fact that the multi-image table has processors configured based on the average client LPAR configuration; on a 2094-701, the cost to run this complex configuration is approximately 5.6%. The commonly used capacity scaling values associated with each model of a processor may be approximated by multiplying the AVERAGE column of ITRRs in the LSPR z/OS V1R13 multi-image table by 559.792. The PCI (Processor Capacity Index) column in the z/OS V1R13 multi-image table shows the result of this calculation. Note that the PCI column was actually calculated using zPCR, thus the full precision of each ITRR is reflected in the values. Minor differences in the resulting PCI calculation may be observed when using the rounded values from the LSPR table.

Of course, using a table of values based on a capacity scaling factor only allows for a gross approximation of the relative capacities among the processor models. A more accurate analysis may be conducted by using zPCR to perform a detailed LPAR configuration assessment to develop the capacity ratio between a “before” and “after” configuration.

Question:

How much variability in performance should I expect when moving a workload to a zEnterprise EC12 (zEC12) processor?

Answer:

As with the introduction of any new server, workloads with differing characteristics will see some variation in performance when moved to a zEC12. The performance ratings for a server are determined by the performance of the average workload that represents what we understand to be the major components of our customers' production environments. While the ratings provide good "middle-of-the-road" values, they do represent an average, and by definition some workloads fall higher than the average and some workloads fall below. The zEC12 has fairly balanced improvements both in its micro-processor design (higher clock speed and improved out-of-order execution) and in its memory hierarchy (much larger on-chip shared cache and book-level shared cache). This suggests that the range of variation in performance of workloads with different hardware characteristics will be somewhat tighter than that seen in recent processor generations.

Question:

Once my workload is up and running on a zEnterprise EC12 (zEC12), how much variability in performance will I see?

Answer:

Minute-to-minute, hour-to-hour and day-to-day performance variability generally grows with the size (capacity) of the server and the complexity of the LPAR configuration. With its improved micro-processor and memory hierarchy design, and the capability to be configured with up to 101 engines, the zEC12 has the capability to deliver up to 1.5 times the capacity of the largest previous server. Continued enhancements to HiperDispatch have been made to help reduce the potential for increased performance variability. In the spirit of autonomic computing, PR/SM and the z/OS dispatcher cooperate to automatically place and dispatch logical partitions to help optimize the performance of the hardware, and minimize the interference of one partition to another. However, while the average performance of workloads is expected to remain reasonably consistent when viewed at small increments of time or by individual jobs or transactions, performance can potentially see some variation simply due to the expected larger and more complex LPAR configurations that can be supported by the zEC12.

Question:

What is HiperDispatch and how does it impact performance?

Answer:

HiperDispatch is the z/OS exploitation of PR/SM's Vertical CPU Management (VCM) capabilities and is exclusive to System z processors since the z10. Rather than dispatch tasks randomly across all logical processors in a partition, z/OS will tie tasks to small queues of logical processors, and dispatch work to a "high priority" subset of the logicals. PR/SM provides processor topology information and updates to z/OS, and ties the high priority logical processors to physical processors. HiperDispatch can lead to improved efficiencies in both the hardware and software in the following two manners: 1) work may be dispatched across fewer logical processors therefore reducing the "multi-processor (MP) effects" and lowering the interference among multiple partitions; 2) specific z/OS tasks may be dispatched to a small subset of logical processors which PR/SM will tie to the same physical processors thus improving the hardware cache re-use and locality of reference characteristics such as reducing the rate of cross-book communication. Note the value of HiperDispatch is higher on the z196 and zEC12 processors due to its sensitivity to the chip-level shared cache topology.

A white paper is available concerning HiperDispatch

at: <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101229>

Question:

What is the performance improvement a z/VM customer might experience on the zEnterprise EC12 (zEC12)?

Answer:

The performance ratios that a z/VM customer workload might experience when migrating to zEC12 from older processors will vary. For z/VM, a single workload was run for the LSPR which has characteristics similar to the HIGH relative nest intensity workload. However, customer workloads have been shown to cover the full range from LOW to HIGH RNI workloads. Thus, it is suggested that you consider the full range of LSPR workloads.

Question:

Where can I read more about the performance of z/VM?

Answer:

The z/VM Performance Resources Page, located at <http://www.vm.ibm.com/perf/>, contains information on z/VM Performance.

Question:

What is the performance improvement a z/VSE customer might experience on the zEC12?

Answer:

The performance ratios that a z/VSE customer workload might experience when migrating from a IBM eServer™ zSeries® server to the zEC12 is represented by the range of ratios for a comparable z/OS migration. For example, the published ratio in the LSPR between the z196 702 and the zEC12 702 is approximately 24% to 29%. z/VSE workloads should expect this same range of performance.

Question:

How do I get performance information for my TPF products running on a zEC12?

Answer:

TPF provides “Workload Specifics ITRRs” separately from the LSPR tables. For more information please contact your TPF Support Representative, or send a request to tpfqa@us.ibm.com.

Capacity on Demand offerings

Question:

Does the zEC12 and zBC12 come with the same Capacity on Demand architecture available on the other zEnterprise System servers (the IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114)?

Answer:

Yes. The provisioning architecture framework is available on all zEnterprise platforms. The architecture offers increased flexibility and capabilities over previous product lines, where only one temporary entitlement record (TER) could be active at a given time. With the latest Capacity on Demand architecture, up to eight different TERs may be installed and active at the same time. Only one On/Off Capacity on Demand TER can be activated at a time. Multiple Capacity Backup and Capacity for Planned Event TERs can be activated at the same time. In addition, the architecture allows concurrent permanent upgrades while temporary capacity is active.

Question:

What are the major differences from IBM System z10 capacity on demand offerings?

Answer:

IBM Resource Link will monitor all installed On/Off CoD records. Every 90 days, Resource Link will generate a replenishment record for each installed record that will move the expiration date out 180 days. The record must be "enabled" for auto-renewal. The next time the system connects through the IBM remote support facility to the IBM Support System a replenishment record is pushed to the system and installed. Once set, no customer renewal action is required. The auto-renewal capability may be cancelled or suspended at any time.

Ordering options were added to Resource Link to allow the purchase of banked CP (high water mark) or unassigned IFL capacity and to allow users to explore capacity needs and upgrade options before purchasing a permanent engine upgrade.

On/Off CoD Administrative Test supports standard order flow, including approval steps that allow for testing of processes and procedures without any resulting hardware or IBM program charges.

With the zEnterprise, up to four Capacity on Demand records can be ordered for pre-installation, including records for CBU and CPE; these records will come pre-installed instead of being staged on the Service Element. If more than four records are ordered with the system, none will be installed and all will be staged on the service element. IBM installation service representatives will assist with the installation of any staged records.

Question:

How does On/Off CoD work?

Answer:

On/Off CoD is available for up to twice the 'purchased' capacity of a given machine based on the LSPR mixed workload multi-image ITRR (Internal Throughput Rate Ratio). On/Off CoD upgrades are allowed for any processor configuration, up to the limit, as long as the number or capacity level of the processors is increased. Upgrades that decrease the number or the capacity level of processors are not allowed. CP capacity upgrades will be ordered based on a percentage increase over the currently purchased capacity. Specialty engines will be ordered in full-engine increments and will always be full-capacity processors.

Question:

How do resource tokens work on the zEnterprise?

Answer:

Management of temporary capacity through On/Off CoD is further enhanced through the introduction of resource tokens. For CP capacity, a resource token represents an amount of processing capacity that will result in one MSU of software cost for one day – an MSU-day. For specialty engines, a resource token represents activation of one engine of that type for one day – an IFL-day, a zIIP-day or a zAAP-day. The different resource tokens are contained in separate pools within the On/Off CoD record.

Using the Resource Link ordering process, the customer determines how many tokens to put into each pool. Once On/Off CoD resources are activated, tokens are removed from their pools every 24 hours. The number removed is based on the highest activation level for that engine type during the previous 24 hours.

Resource tokens are intended to help bound the hardware costs associated with using On/Off CoD. The use of resource tokens is optional and they are available on either a prepaid or post-paid basis. Prepaid resource tokens are priced based on the total number of resource tokens contained in the On/Off CoD record. When post-paid, the total billing against the On/Off CoD record is limited by the number of resource tokens contained in the record.

For more information, refer to the Capacity on Demand Users Guide, SC28-2605.

Question:

What is the change that allows me to alter my On/Off CoD capacity without having to order and download a new On/Off CoD record?

Answer:

With the Capacity on Demand record structure, instead of ordering a separate On/Off CoD record for each possible configuration, a single reusable record that identifies the maximum possible activated configuration can be ordered. CP capacity is ordered based on a percentage increase over the currently purchased capacity. Other engines are ordered in full-engine increments. When temporary capacity is needed, the desired target configuration is identified by the customer. Need more processing capacity? Select a new larger target configuration without a new activation. The amount of active temporary capacity can also be decreased by selecting a lower target. When temporary capacity is no longer required, the customer simply returns the machine to its base configuration. The On/Off CoD record is still available and can be used for additional capacity at any time. The maximum temporary upgrade is still restricted to a maximum of two times the machine's purchased capacity configuration. Customers will be charged for the additional activated capacity on a 24-hour basis, and if capacity is increased multiple times during a 24 hour period, the charge will be based on the greatest amount of capacity activated.

Question:

What is the API that is available for On/Off CoD?

Answer:

There is an API provided within the existing HMC SNMP interface that is designed to enable customers to use other automation code (which conforms to the API) to enable activation of On/Off CoD on the HMC without human intervention. This will allow flexibility of operation of the On/Off CoD function. A specification for API is available on the IBM Resource Link website. This API is exclusive to the zEC12, zBC12, z196 and z114.

Question:

What is banked or unassigned capacity?

Answer:

Banked or unassigned capacity can be either general-purpose or IFL capacity that has been purchased by the end user, but for business reasons has been unassigned (turned off) by the end user such that no workload can be executed on the unassigned capacity.

The customer is responsible for any additional maintenance or software charges that may result from activating unassigned capacity.

Question:

Can you tell me where to get more information on the 2014 announcement concerning On/Off Capacity on Demand maintenance?

Answer:

There are some FAQs on the announcement. They can be found at: http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ST&infotype=SA&appname=STGE_ZS_ZS_USEN&htmlfid=ZSQ03070USEN&attachment=ZSQ03070USEN.PDF

Question:

Can On/Off CoD use unassigned IFL or CP capacity for any purpose other than its original intended one?

Answer:

Yes. An unassigned processor may be activated as any temporary engine type. An unassigned CP may be activated as a CP, IFL, zIIP, zAAP, ICF, or SAP. Likewise, an unassigned IFL may be temporarily activated as a CP, IFL, zIIP, zAAP, ICF, or SAP. However, if you activate unassigned capacity for any purpose other than its original purchased purpose, that activation will be priced as if the capacity were unowned. So temporarily activating an unassigned CP as a CP engine results in no daily hardware charges. However, activating that same unassigned CP as an IFL, zIIP, zAAP, ICF, or SAP results in the same daily hardware charge as an unpurchased engine.

Question:

Do I need to follow a different process to use unassigned capacity with On/Off CoD?

Answer:

No. The order process is the same.

Question:

Will temporary capacity be available for specialty processors (IFLs, ICFs, zAAPs, zIIPs, SAPs) on the zEnterprise?

Answer:

Yes. All engine types are available for temporary upgrades.

Question:

Can I order On/Off CoD if I have subcapacity CP processors?

Answer:

Yes. On/Off CoD is available for up to twice the 'purchased' capacity for a given machine but the number or capacity setting of CPs cannot be decreased. With the full 'matrix' upgradeability of the subcapacity processors, a machine's capacity can be temporarily changed with processors of equal or greater capacity depending on the capacity requirements.

On the zEC12, when subcapacity settings are used for On/Off CoD upgrades, the number of general purpose processors cannot exceed 20. (NOTE – when using On/Off CoD to increase the speed of subcapacity processors on the zEC12 to full capacity, it is possible to have more than 20 processors active at one time.)

Question:

When can I begin placing On/Off CoD orders against my zEnterprise?

Answer:

On/Off CoD can be initiated as soon as the profile for the zEnterprise is established. The prerequisite of establishing a profile is the signing of the necessary contract supplements associated with ordering features 9900 and 9896.

Question:

Will I be able to do Capacity Backup Upgrade (CBU) capability on any engine type?

Answer:

Yes. The zEnterprise is able to activate all processor types as part of CBU upgrades: IFLs, zAAPs, zIIPs, ICFs, CPs and SAPs.

Question:

Can I order CBU processors if I have subcapacity processors?

Answer:

Yes. CBU is available for machines configured with subcapacity processors but a CBU Upgrade cannot decrease the number or the capacity level of installed CP processors.

On the zEC12 – remember that the configuration cannot exceed 20 subcapacity processors (CP or CBU). A customer no longer has to increase the quantity of CPs for CBU, but can just increase the capacity of the existing CPs count by ordering CBU CP features with greater capacity. When the quantity of CBU processors on the zEC12 exceeds 20, all CBU processors will be full capacity.

Question:

Can I add CBU capacity by selecting CBU engines that have more CP capacity than my permanent configurations and have fewer actual CPs than the CBU machine's base configuration?

Answer:

No, a CBU Upgrade cannot reduce the number of CPs configured for the machine.

Question:

Can I convert an active permanent engine to another engine type during CBU?

Answer:

No. All active permanent engines must remain as part of the CBU environment (although they may change in capacity) and cannot be converted to another type while the CBU Upgrade is active. Unassigned engines may be used as other engine types within a CBU upgrade's configuration.

Question:

What can I do during a CBU test?

Answer:

Customers may now execute productive workload on the capacity of a CBU upgrade during a CBU test provided that a) an amount of System z productive workload capacity equivalent to the CBU Upgrade is shut down or otherwise made unusable by the customer for the duration of the test, and b) the appropriate contracts are in place. All CBU contracts contain these new CBU Test terms. Existing CBU

customers will need to execute IBM Customer Agreement Amendment for IBM System z Capacity Backup Upgrade Tests, form number Z125-8145.

Question:

What Capacity on Demand features can I order on IBM Resource Link?

Answer:

Permanent processor and memory upgrades, On/Off CoD, CPE, and CBU records can be ordered from Resource Link.

Question:

Is FC 9898, the permanent upgrade via CIU feature, still a prerequisite to On/Off CoD?

Answer:

No. A machine may be enabled for On/Off CoD upgrades without having to enable the machine for permanent upgrades via CIU.

Question:

If I have Capacity on Demand features on my system now will they be lost if I upgrade to a zEnterprise?

Answer:

The CBU,CPE, Permanent Upgrades via CIU, On/Off Capacity on Demand, and On-Line CoD Buying enablement features can be brought forward to the zEC12 including any CBU, CPE and On/Off Capacity on Demand records installed or staged.

Question:

Has the way I order Capacity on Demand features changed?

Answer:

Permanent processor and memory upgrades, CBU and the Capacity for Planned Event may be ordered either through your IBM or BP sales representative, or Resource Link. On/Off Capacity on Demand orders must still be placed through Resource Link only.

Question:

Tell me about the Capacity Provisioning capabilities of the zEnterprise?

Answer:

A Capacity Provisioning Manager was introduced on z/OS V1.10 also available on z/OS V1.9 with a PTF. The Capacity Provisioning Manager can monitor z/OS systems on zEC12, zBC12, z196, z114 and System z10 servers. Activation and deactivation of temporary capacity can be suggested or performed automatically based on user-defined schedules and workload criteria.

The Capacity Provisioning Control Center is a tool for managing capacity provisioning for zEC12, zBC12, z196, z114 and System z10 servers. It is designed to manage provisioning policies and domain configurations. Provisioning policies specify the criteria for capacity increases and decreases, while domain configurations specify systems to be observed and servers to be managed. In z/OS V1.10, support was provided for a policy definition application which requires a workstation running Microsoft Windows XP. z/OS V1.12 added support for averaged rolling performance intervals, CICS and IMS transaction monitoring, and Microsoft Windows Vista.

Specifically, the Capacity Provisioning Control Center provides the following functions:

- Create and edit Capacity Provisioning policies
- Create and edit Capacity Provisioning domain configurations
- Connect to the Provisioning Manager; Display the status of the Provisioning Managers

- Install Capacity Provisioning policies and domain configurations into the Provisioning Manager
- With z/OS V1.11 the z/OS Capacity Provisioning exploits the new BCPii capability to remove the requirement for a TCP/IP connection for communication between the Capacity Provisioning Manager and the Support Element (SE) or Hardware Management Console (HMC). In addition, Capacity Provisioning provides improved logical processor management support. This support allows you to specify that Capacity Provisioning Manager should prompt the operator to configure logical processors online or offline as needed for capacity changes. Capacity Provisioning support for BCPii and logical processor management is also available on z/OS V1.10 with PTF UA47421.
- With z/OS V1.12 the z/OS Capacity Provisioning uses the delay data for transaction service classes provided by RMF to help determine whether a provisioning action is required for servers on which CICS and IMS are running. Monitoring delay data for CICS and IMS transaction classes is intended to help improve capacity provisioning decisions for servers with LPARs running CICS and IMS. This function is also available now for z/OS V1.10 and z/OS V1.11 with the PTF for APAR OA29641.
- For z/OSV1.13, Capacity Provisioning Manager now allows you to specify different quantities for obtaining the first capacity increment and subsequent increments, to help you add the right amount of capacity quickly. This is available for z/OS V1.11 and z/OS V1.12 with the PTF for APAR OA35284. In addition, support is provided for recurring time conditions, which can simplify defining the Capacity Provisioning policy for events that repeat and is also available for z/OS V1.11 and z/OS V1.12 with the PTF for APAR OA35284. Also in z/OS Capacity Provisioning Control Center support is added for the 32-bit and 64-bit versions of Microsoft Windows 7 Professional Edition.
- In z/OSMF V1.13, the Capacity Provisioning task is designed to support easier monitoring of z/OS Capacity Provisioning Manager status, which can reduce the time it takes to get capacity provisioning status. The Windows-based CPCC is still required for managing the z/OS Capacity Provisioning Manager
- IBM intends for z/OS V1R13 to be the final release for which SNMP as supported protocol for the communication to the HMC or Support Element is available. Customers currently using SNMP for communication, should migrate to BCPii. The migration includes enabling the communication through BCPii for the provisioning manager user and adding a new key to the Capacity Provisioning Manager parameter file.

Question:

What is Capacity for Planned Event (CPE)?

Answer:

CPE is temporary access to capacity intended to replace capacity lost within the enterprise due to a planned event such as a facility upgrade or system relocation. CPE is similar to CBU in that it is intended to replace lost capacity; however, it differs in its scope and intent. Where CBU addresses disaster recovery scenarios that can take up to three months to remedy, CPE is intended for short-duration events lasting up to three days, maximum. Each CPE record is ordered with the capacity needed to meet business needs.

Question:

How will CPE be priced?

Answer:

There is a fixed price for each CPE record (which is for a three day event). The price depends on the amount of capacity ordered. There are no additional IBM programs or IBM maintenance service charges for capacity activated by CPE. IBM, at its discretion, reserves the right to add maintenance charges at a later time.

Question:

How is CPE ordered?

Answer:

CPE records can be ordered by Resource Link or via your IBM or BP sales representative as an MES order. The system must be enabled for Capacity for Planned Event, Feature code #9912. The enablement process for each Capacity on Demand offering begins when you order the associated enablement feature code and sign the associated IBM contract.

- The following contracts must be signed one time within a given country before IBM will accept an order for the first instance of the CPE Enablement feature code (feature #9912):
 - IBM Customer Agreement Attachment for IBM System z Capacity on Demand Offerings (US form #Z125-7879)
 - IBM Customer Agreement Attachment for IBM System z Replacement Capacity Offerings (US form #Z125-7880)
 - IBM Customer Agreement Attachment for IBM System z Capacity for Planned Events (US form #Z125-7882).

I/O

Question:

Does the zEC12 and zBC12 use the same PCIe I/O drawer?

Answer:

Yes. The PCIe I/O drawer and the new form factor I/O cards were introduced to support a direct Peripheral Component Interconnect Express Generation 2 (PCIe Gen2) infrastructure with increased capacity, granularity, and infrastructure bandwidth, as well as increased reliability, availability, and serviceability. Hardware features have been introduced to support the PCIe I/O drawer for the SAN, the LAN, and security environments to support PCIe Gen2 - FICON Express8S, OSA-Express5S and the new Crypto Express4S.

With their compact design, two PCIe I/O drawers occupy the same space as one I/O cage while delivering 14% more capacity - up to 128 channels (64 I/O cards) are supported versus the 112 channels (28 I/O cards) offered with the I/O cage.

With the PCIe I/O drawer we are introducing enhanced granularity features for the storage area network (SAN) and the local area network (LAN).

The industry-standard Peripheral Component Interconnect Express Generation 2 (PCIe Gen2) infrastructure from the processor book to the I/O cards, supports an 8 gigabytes per second (8 GBps) bus (compared to the prior 6 GBps bus). This infrastructure offers enhanced bandwidth in the connection to the processor book as well as the distribution buses contained in the drawer.

The PCIe I/O drawer is designed with symmetrical, redundant cooling across all I/O cards and power supplies for improved RAS. Dual PCIe paths to the system processors from an I/O card provide enhanced resiliency. Also included are temperature monitoring of critical ASICs to optimize cooling and reliability.

Question:

Tell me a little about the FICON Express8S feature to be used with the PCIe I/O drawer in a zEC12, zBC12, z196 or z114?

Answer:

IBM introduced a new generation of features for the storage area network (SAN) in support of the PCIe 8 GBps host bus and the PCIe I/O drawer. The FICON Express8S features for the multimode (SX) and single mode (LX) fiber optic cabling environments have path length reductions for High Performance FICON for System z (zHPF) and Fibre Channel Protocol (FCP), increased start I/Os, improved throughput for zHPF and FCP with the introduction of a hardware data router, and increased port granularity - two channels/ports per feature.

With a design for increased performance for zHPF and FCP: FICON Express8S contains an IBM ASIC which is designed to support the 8 GBps PCIe interface to the PCIe I/O drawer and increased start I/Os. In addition, a hardware data router has been added in support of the zHPF and FCP protocols for path length reduction and increased throughput. FICON Express8S supports a link data rate of 2, 4, or 8 Gbps auto-negotiated.

With these changes FICON Express8S, when supporting the zHPF or FCP protocols, has been designed to achieve full duplex line speed - 8 Gbps - in each direction.

Question:

In the past, I have noticed that when you announce a new link data rate for FICON and FCP channels, that you withdraw support for a lower link data rate. Will you continue to support the 2 Gbps link data rate on features after FICON Express8 and FICON Express8S, which support a link data rate of 2, 4, or 8 Gbps? These features do not support 1 Gbps link data rates.

Answer:

You are correct. The industry standard for 16 Gbps interfaces, is autonegotiation to 4, 8, or 16 Gbps. These interfaces will not support 2 Gbps link data rates. Migration to 4 Gbps interfaces, at a minimum, should be taken into consideration when planning future infrastructure purchases for fiber optic cabling, as well as purchases of switches, directors, and control units. It is important, also to remember that as the link data rate increases the unrepeated distance and link loss budget decreases, especially when using multimode fiber optic cabling. When purchasing multimode fiber optic cabling for your environment it is important to consider migrating to 50 micron fiber rated at 2000 MHz-km (OM3).

Question:

Can you give me some more details about the performance of the FICON Express8S features?

Answer:

In laboratory measurements, using FICON Express8S in a z196 with the zHPF protocol and small data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum of 92,000 IOs/sec, compared to the maximum of 52,000 IOs/sec achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 77% increase and applies to reads, writes, and a read/write mix. Results on a zEC12 are comparable.

In laboratory measurements, using FICON Express8S in a z196 with the zHPF protocol and a mix of large sequential read and write data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum throughput of 1600 MB/sec (reads + writes) compared to a maximum of 770 MB/sec (reads + writes) achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 108% increase. Results on a zEC12 are comparable.

This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

Increased performance for the FCP protocol: A FICON Express8S feature, when defined as CHPID type FCP, conforms to the Fibre Channel Protocol (FCP) standard to support attachment of SCSI devices, to complement the classical storage attachment supported by FICON and zHPF channels.

In laboratory measurements, using FICON Express8S in a z196 with the FCP protocol for small data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum of 92,000 IOs/sec, compared to the maximum of 84,000 IOs/sec achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 10% increase and applies to reads, writes, and a read/write mix. Results on a zEC12 are comparable.

In laboratory measurements, using FICON Express8S in a z196 with the FCP protocol and an internal driver supporting the hardware data router, executing a mix of large sequential read and write data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum throughput of 1600 MB/sec (reads + writes) compared to the maximum of 770 MB/sec (reads + writes) achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 108% increase. Results on a zEC12 are comparable.

The FCP protocol is supported by z/VM, z/VSE, and Linux on System z. Note: zHPF is not supported by z/VSE.

Question:

What is new in the area of T10-DIF?

Answer:

Recognizing that high reliability is important to maintaining the availability of business-critical applications, the System z Fibre Channel Protocol (FCP) has implemented support of the American National Standards Institute's (ANSI) T10 Data Integrity Field (DIF) standard. Data integrity protection fields are generated by the operating system and propagated through the storage area network (SAN). System z helps to provide added end-to-end data protection between the operating system and the storage device.

An extension to the standard, Data Integrity Extensions (DIX), provides checksum protection from the application layer through the host bus adapter (HBA), where cyclical redundancy checking (CRC) protection is implemented.

T10-DIF support is now available for Linux on System z environments as well as z/VM for guest exploitation. Data integrity protection fields are generated by the operating system and propagated through the storage area network (SAN). System z helps to provide added end-to-end data protection between the operating system and the storage device.

T10-DIF support by the FICON Express8S and FICON Express8 features, when defined as CHPID type FCP, is supported on the zEC12, zBC12, z196, and z114. Exploitation of the T10-DIF standard is supported by z/VM for guest exploitation and by Linux on System z. Exploitation is also required by the storage device.

Increased port granularity: The FICON Express8S 10KM LX and SX features for single mode and multimode fiber optic cabling environments each now have two channels/ports per feature versus the four channels per feature for the FICON Express8 features. This design helps facilitate purchasing the right number of ports to help satisfy your application requirements and to better optimize for redundancy.

The FICON Express8S features, supporting CHPID types FC (zHPF, FICON, channel-to-channel) and FCP, are exclusive to zEC12, zBC12, z196 and z114. They are for use exclusively in the PCIe I/O drawer and are supported by z/OS, z/VM, z/VSE, z/TPF, and Linux on System z. Note that zHPF is not supported by z/VSE.

Question:

What is the High Performance FICON for System z (zHPF) relating to QSAM, BPAM and BSAM access methods?

Answer:

zHPF was enhanced to support certain I/O transfers for workloads using QSAM, BPAM, and BSAM access methods. Significant I/O performance improvements are expected without the need for application changes. This builds upon existing zHPF support for VSAM, Extended Format sequential, zFS, and PDSE data sets and provides support for these QSAM, BPAM, and BSAM data set types when a new parameter is specified in the IGDSMSxx member of parmlib:

- Basic non-extended format physical sequential data sets
- Basic and large format sequential data sets

The zHPF enhancement for QSAM, BPAM, and BSAM access methods is exclusive to the zEnterprise servers (zEC12, zBC12, z196 and z114) and applies to all supported FICON features (CHPID type FC). It is supported by the IBM System Storage DS8700 and DS8800 series and z/OS.

Question:

What is the zHPF enhancement relating to format writes?

Answer:

zHPF has been enhanced to support format writes. This capability applies to all of the same data set types that were originally supported by the Modified Indirect Data Address Word (MIDAW) facility and zHPF, in addition to QSAM, BPAM, and BSAM data sets described above. The performance value of these enhancements are highest for small blocks, which are typically used for databases.

DB2 utilities use format writes. Those utilities that load or restore data into a table space or index with a 4K page size are expected to experience the most benefit.

The zHPF enhancement for format writes is exclusive to the zEnterprise servers and applies to all supported FICON features (CHPID type FC). It is supported by the IBM System Storage DS8700 and DS8800 series and z/OS.

Question:

What is the High Performance FICON for System z enhancement relating to DB2 list prefetch processing?

Answer:

zHPF has been enhanced to provide improvements for DB2 list prefetch processing.

zHPF list prefetch is supported by the FICON Express4 and FICON Express8 features. However, those features limit the number of discontinuities that a single zHPF channel program can have to 22. If the number exceeds 22, z/OS splits the operation into two I/Os. This applies to 4K pages, since DB2 for z/OS typically reads 32x4K pages. The FICON Express8S features remove this limitation. This results in fewer I/Os when using the FICON Express8S features.

The zHPF enhancement for DB2 list prefetch is exclusive to the zEnterprise servers and applies to the FICON Express8S features (CHPID type FC) exclusively. It is supported by the IBM System Storage DS8700 and DS8800 series and z/OS.

Question:

Can I continue to carry forward my PSCs to the next generation high end-mainframe?

Answer:

Yes, you cannot buy new but you can carry forward PSCs when you do an upgrade.

Question:

Tell me a little about the new channel subsystem enhancement regarding support for 24k subchannels per FICON Express channel (port) that was announced July 23, 2013.

Answer:

To help facilitate growth as well as continuing to enable server consolidation, we are now supporting up to 24k subchannels per FICON Express channel (channel path identifier - CHPID). Now you will be able to define more devices per FICON channel, which includes primary, secondary, and alias devices. The maximum number of subchannels across all device types addressable within an LPAR remains at 63.75k for subchannel set 0 and 64k-1 for subchannel sets 1 and higher. This support is exclusive to the zEC12 and the zBC12 and applies to the FICON Express8S, FICON Express8, and FICON Express4 features when defined as CHPID type FC. This is supported by z/OS, z/VM, and Linux on System z.

Question:

What is z/OS discovery and auto configuration for FICON channels for zEnterprise?

Answer:

With zEC12, zBC12, z196 or z114 and z/OS, z/OS discovery and auto configuration (zDAC), is designed to automatically perform a number of I/O configuration definition tasks for new and changed disk and tape controllers connected to a switch or director that is attached to a FICON channel. When new controllers are added to an I/O configuration or changes are made to existing controllers, the system is designed to discover them and propose configuration changes based on a policy you define in the hardware configuration dialog (HCD) or in the Hardware Configuration Manager (HCM, optional priced feature of z/OS). Your policy can include preferences for availability and bandwidth including parallel access volume (PAV) definitions, control unit numbers, and device number ranges.

zDAC is designed to perform discovery for all systems, including those in a Parallel Sysplex that support the function. The proposed configuration will incorporate the current contents of the I/O definition file (IODF) with additions for newly installed and changed control units and devices. zDAC is designed to help simplify I/O configuration on zEC12, zBC12, z196 or z114 running z/OS and reduce complexity and setup time.

zDAC applies to all FICON features supported on zEC12, zBC12, z196 or z114 when configured as CHPID type FC and is supported by z/OS. zDAC is exploited by the IBM System Storage DS8700. For more information refer to Hardware Announcement 110-168, dated July 20, 2010, IBM System Storage DS8700 (M/T 242x) delivers z/OS Distributed Data Backup and new functional capabilities.

Question:

What was the change in the Statements of Direction relating to ISC-3 your recent announcement?

Answer:

To help with future planning we have identified a number of features that cannot be ordered on future servers. Originally we had said that the z196 and z114 were to be the last servers to support InterSystem Channel-3 (ISC-3) coupling links, but we are amending that to state that the zEC12 and zBC12 will be the last ones to offer ordering of ISC-3.

Question:

I don't see anything about ESCON® in the zEC12 or zBC12 announcements. Why is that?

Answer:

Since April of 2008, System z has indicated that ESCON channels would be phased out. The z196 and z114 were the last servers to support ESCON channels. They cannot be ordered new nor carried forward on the zEC12 or the zBC12.

Question:

I currently have ESCON devices. Are there alternatives if I wish to retain my ESCON channels to attach to future System z servers?

Answer:

Alternate solutions are available for connectivity to ESCON devices. IBM Global Technology Services offers an ESCON to FICON Migration solution, Offering ID #6948-97D, to help facilitate migration from ESCON to FICON. This offering should help customers to simplify and manage a single physical and operational environment.

Question:

What Storage Area Network (SAN) products are currently qualified to operate at 8 Gbps with the zEnterprise System?

Answer:

The most current list of qualified SAN products are now available for review on Resource Link. <http://www.ibm.com/servers/resourcelink/>

“Sign In” with valid user ID and password. On the left, click on the "Library" link. Locate the listing of "Hardware products for servers" around the middle of the Web page. Click on the link “Switches and directors qualified for IBM System z FICON and FCP channels”

Question:

Can a 1 Gbps device be connected to a Storage Area Network (SAN) director that supports 8 Gbps?

Answer:

Based upon the Fibre Channel Physical Interface (FC-PI-4) standard (Rev. 8.00),

A 8 Gbps optics can autonegotiate to 2, 4, or 8 Gbps. A 4 Gbps optics can autonegotiate to 1, 2, or 4 Gbps. You will need 4 Gbps optics in the switch or director to connect downstream to a 1 Gbps device. If a FICON Express4 feature is carried forward into a zEC12, zBC12, z196 or z114, it can autonegotiate to 1 Gbps.

Question:

You have a tool for assigning worldwide port names (WWPN) for both virtual and physical Fibre Channel Protocol (FCP) ports. Can you tell me more about it?

Answer:

In May of 2009 support for assigning WWPNs to virtual FCP ports was made available on System z10. The tool assists with preplanning of the SAN environment - to allow you to set up your SAN in advance of the server installation. In September of 2010 support was expanded to include assigning WWPNs to physical FCP ports and is offered for all zEnterprise servers. In addition, you can retain your SAN configuration if a FICON feature is replaced instead of altering your SAN configuration based upon the FICON feature's “burned in” WWPN. The WWPN tool can be downloaded from the “Tools” section of Resource Link at: <http://www.ibm.com/servers/resourcelink/>

Note:

All new zEC12 enhancements for July 23, 2013, will require Driver 15, available on September 20, 2013.

Networking

Question:

Tell me about the new OSA-Express5S that was announced July 23, 2013:

Answer:

A new generation of Ethernet features is being introduced for use in the PCIe I/O drawer and continues to be supported by the 8 GBps PCIe Gen2 host bus. This is an introduction of the full family of features - 1000BASE-T Ethernet for copper environments, in addition to 10 Gigabit Ethernet (10 GbE) and Gigabit Ethernet (GbE) for single-mode and multimode fiber optic environments.

The performance characteristics are comparable to the OSA-Express4S features. They also retain the same form factor and port granularity - two ports per feature for the 1000BASE-T Ethernet and Gigabit Ethernet features, and one port per feature for the 10 Gigabit Ethernet features. And, the first time offered for networking, the OSA-Express5S features have small form factor pluggable+ (SFP+) transceivers.

The OSA-Express5S family of features, is exclusive to the zEC12 and to zBC12 and are supported by z/OS, z/VM, z/VSE, z/TPF, and Linux on System z.

Question:

What was announced July 23, 2013 about managing my OSA Express cards on the HMC?

Answer:

Support for OSA-Express5S and OSA-Express4S features is new in the OSA Advanced Facilities on the HMC and is exclusive to the zEC12 and the zBC12. The latest driver level is required. "OSA/SF on the HMC" provides configuration, validation, activation, and display support for selected functions of the OSA-Express5S and OSA-Express4S.

OSA Advanced Facilities on the HMC ("OSA/SF on the HMC") is required for the OSA-Express5S features operating as CHPID types of OSD, OSE, or OSN. Either OSA/SF on the HMC or the existing OSA/SF operating system component under z/OS, z/VM, and z/VSE can be used for the OSA-Express4S feature configured with these CHPID types. Note that OSA/SF on the HMC plays no role in the configuration and management of CHPID types OSC, OSM, or OSX.

Detailed information is available in the following manual available from ResourceLink: Open Systems Adapter/Support Facility on the Hardware Management Console (SC14-7580)

Note:

All new zEC12 enhancements for July 23, 2013, will require Driver 15, available on September 20, 2013.

Older, but important, networking information

Question:

What is checksum offload for z/OS environments and z/VM guest exploitation?

Answer:

Checksum offload for IPv6 packets is available for z/OS environments and for z/VM guest exploitation. When the checksum function is offloaded from the host, CPU cycles are reduced, improving performance. With the introduction of OSA-Express4S, the checksum offload function is now performed for IPv6 packets as well as IPv4 packets, whether the traffic goes out to the local area network (LAN), comes in from the LAN, or flows logical partition-to-logical partition through OSA-Express4S.

Checksum offload provides the capability of calculating the Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Protocol (IP) header checksums for Internet Protocol Version 4 (IPv4) packets and now, IPv6 packets.

When checksum offload was introduced in May of 2003, it was limited to IPv4 packets. Checksum offload for IPv6 packets is exclusive to OSA-Express4S features (CHPID types OSD and OSX) on zEnterprise. It is supported by z/OS and z/VM for guest exploitation. Checksum offload for IPv4 packets is currently available for all in-service releases of z/OS and Linux on System z.

Checksum offload for LPAR-to-LPAR traffic in the z/OS environment is included in the OSA-Express4S design for both IPv4 and IPv6 packets.

Question:

Tell me about large send for IPv6 packets.

Answer:

Large send for IPv6 packets is available for z/OS V1.13 and z/VM 5.4 environments: Large send (also referred to as TCP segmentation offload) is designed to improve performance by offloading outbound TCP segmentation processing from the host to an OSA-Express4S feature by employing a more efficient memory transfer into OSA-Express4S.

Large send support for IPv6 packets applies to the OSA-Express4S features (CHPID type OSD and OSX), and is exclusive to zEnterprise. It is supported by z/OS V1.13 and z/VM V5.4. (Large send for IPv4 packets is currently available for all in service releases of z/OS, Linux on System z, and z/VM for guest exploitation.)

Note: Large send is not supported for LPAR-to-LPAR packets.

In addition, the OSA-Express4S, z/OSV 1.13, and z/VM 5.4 (or guests) supports checksum offload for IPv6 packets, and for checksum offload for IPv6 and IPv4 for LPAR-to-LPAR traffic.

Question:

Can you tell me about Inbound Workload Queuing (IWQ)?

Answer:

Inbound workload queuing (IWQ) for the OSA-Express features is able to differentiate and separate inbound Enterprise Extender traffic to a new input queue, in addition to the current and existing interactive, streaming, and sysplex distributor queues. The Enterprise Extender separation and processing associated with the Enterprise Extender input queue provides improved scalability and performance for Enterprise Extender.

With each input queue representing a unique type of workload, each having unique service and processing requirements, the IWQ function allows z/OS to use appropriate processing resources for each input queue. This approach allows multiple concurrent z/OS processing threads to process each unique input queue to avoid traditional resource contention. In a heavily mixed workload environment, this "off the wire" network traffic separation provided by OSA-Express IWQ reduces the conventional z/OS processing required to identify and separate unique workloads.

Inbound workload queuing for Enterprise Extender is supported by the OSA-Express4S and OSA-Express3 features when defined as CHPID types OSD or OSX. It is exclusive to zEnterprise, and is supported by z/OS and by z/VM for guest exploitation.

Question:

Is there a command to help gather OSA configuration information?

Answer:

As additional complex functions have been added to OSA, the ability for the system administrator to display, monitor, and verify the specific current OSA configuration unique to each operating system has become more complex. With OSA-Express4S and OSA-Express3 features you have the capability for the operating system to directly query and display the current OSA configuration information (similar to OSA/SF). z/OS exploits this OSA capability by introducing a TCP/IP operator command called Display OSAINFO. Display OSAINFO allows the operator to monitor and verify the current OSA configuration, which helps to improve the overall management, serviceability, and usability of OSA-Express3.

Display OSAINFO is available starting with OSA-Express3 supporting CHPID types OSD, OSM, and OSX, the z/OS operating system, and z/VM for guest exploitation

Question:

Can you remind me about the HiperSockets enhancement that was made in 2010 for the Open Systems Adapters family of features supporting connectivity to the local area network (LAN)?

Answer:

A single logical partition can now connect to up to 32 HiperSockets, double the number of previously supported. With HiperSockets, you have independent, integrated, virtual local area networks (LANs). No physical cabling or external connections are required. HiperSockets supports Internet Protocol version 4 (IPv4) and IPv6. Up to 32 HiperSockets are supported by z/OS, z/VM, z/VSE, and Linux on System z.

An enhancement to the HiperSockets architecture, HiperSockets network traffic analyzer, is designed to help simplify problem isolation and resolution. You now have the capability to trace Layer 2 and Layer 3 HiperSockets network traffic. HS NTA allows Linux on System z to control the trace for the internal virtual LAN and to capture the records into host memory and storage (file systems). Linux on System z tools can be used to format, edit, and process the trace records for analysis by system programmers and network administrators. This function is available for zEnterprise and System z10 servers.

For network performance- inbound workload queuing: z/OS workloads are becoming more diverse; each type of workload may have different service level requirements. OSA-Express-3 introduces inbound workload queuing (IWQ), which creates multiple input queues and allows OSA to differentiate workloads "off the wire" and then assign work to a specific input queue (per device) to z/OS. With each input queue representing a unique type of workload, each having different service and processing requirements, the IWQ function allows z/OS to pre-assign appropriate processing resources for each input queue. This approach allows multiple concurrent z/OS processing threads to then process each input queue (workload), avoiding traditional resource contention. In a heavily mixed workload environment, this "off the wire" network traffic separation provided by OSA-Express3 IWQ reduces the conventional z/OS processing required to identify and separate workloads, which results in improved overall system performance and scalability. IWQ is supported on zEnterprise and System z10 and is exclusive to OSA-Express3 CHPID types OSD and OSX (CHPID type OSX is exclusive to zEnterprise).

It is estimated that networking throughput for interactive workloads can be improved by 30-50%. The interactive throughput measurements were obtained on System z10 EC, model 2097-E64 with OSA-Express3 Inbound Workload Queuing function. Actual benefits will depend on the amount of data being transferred, presence of bulk-data traffic in the mix, and whether communication is z/OS to z/OS or z/OS to distributed system. It is anticipated that IWQ will provide benefit for streaming workloads by reducing the number of costly network retransmissions due to out-of-order packets, and will benefit Sysplex Distributor traffic by providing a more streamlined processing path.

For network management - query and display OSA configuration: Previously, OSA-Express system architecture introduced the capability for operating systems to dynamically register the OSA configuration. This approach significantly improved OSA-Express usability by reducing the burden placed on the system administrator to manually configure OSA-Express for each unique operating system configuration.

Traditionally, the Open Systems Adapter Support Facility (OSA/SF) has provided the administrator with the ability to monitor the OSA configuration. As additional complex functions have been added to OSA, the ability for the system administrator to display, monitor, and verify the specific current OSA configuration unique to each operating system has become more complex. OSA-Express3 introduces the capability for the operating system to directly query and display the current OSA configuration information (similar to OSA/SF). z/OS exploits this OSA capability by introducing a TCP/IP operator command - Display OSAINFO. It allows the operator to monitor and verify the current OSA configuration, which helps to improve the overall management, serviceability, and usability of OSA-Express3. Display OSAINFO is exclusive to OSA-Express3 CHPID types OSD, OSM, and OSX, the z/OS operating system, and z/VM for guest exploitation.

Water cooling option

Question:

Will there be a water cooling option for the zBC12?

Answer:

No.

Question:

When should I consider installing water cooling on the zEC12?

Answer:

- If you have installed z196 servers with the water cooling option.
- If you have a problem with hot spots in your data center, water cooling will help eliminate them.
- If you are limited on power in your data center, a water-cooled system is a way to increase server capacity without increasing power requirements.

Your power savings will vary based on the server configuration in terms of the number of processor books and I/O cards, as well as on the power and cooling used in your data center. For a well-utilized (not maximum) four-book system, you can expect to see savings of up to about 6%.

- Simply said, if you have a data center that is bounded by limited power capacity or if you want to reduce server input power and the cost to remove server heat load, you should look at the water cooling option. In addition, you should explore capabilities such as high voltage DC input.
 - If you are building a new data center water cooling may be an important way to get a significant reduction in energy use. When considering water cooling it's important to look at your entire data center strategy. System z is one component but even more significant improvements in removal of heat load can be achieved by implementing water cooling across your other server platforms.
-

Question:

Which zEC12 models will benefit most from water cooling?

Answer:

Three- and four-book systems will see the most benefit from water cooling. However, some customers with an overall water cooling strategy in their data centers may want to consider water cooling on one- or two-book systems.

Question:

Can I tap into my building's chilled water?

Answer:

A large number of data centers use chilled water somewhere in their cooling infrastructure, for example to provide cooling to CRAH (Computer Room Air Handler) units. The chilled water requirements (temperature range, pH, hardness, contamination, particulates, etc.) for the zEC12 server should match the characteristics of the data center's available chilled water in most cases.

Question:

How long will it take me get my data center ready for a water cooled system?

Answer:

It depends on whether you have chilled water in your data center (most do), how it was installed, pipe locations, and on the condition of those components. In some cases, it will be quite straightforward to get ready for installation. In other cases it may require more time.

Question:

What changes will I need in my floor tile cut outs for water?

Answer:

There is no change to the floor cutouts for the zEC12 with the water cooling option – but the option will add 4 inches depth to the rear of the server. The floor tile cutouts for zEC12 are the same as those for the z196 and the z10 EC.

Question:

Is there any change to the height of the server for water cooling?

Answer:

No.

Question:

Is there any change to the weight of the server?

Answer:

With water cooling, the installed weight will increase approximately 145 pounds (66 kg).

Question:

How many Water Cooling Units (WCU) are on a water cooled zEC12?

Answer:

Normal operation of the system requires one WCU. However, there will be two units for high availability.

Question:

What would happen if chilled water input is lost?

Answer:

We are providing two levels of air cooling backup mode for maximum robustness – the system will continue to operate with at a small performance decrease even when chilled water input is lost to both WCUs. In this backup mode, all system heat is delivered to the system's exhaust air. Even in the extremely unlikely condition of the internal water loop not functioning there is an additional air cooling backup air flow that exhausts out the top of the processor frame like in the air cooled zEC12.

Question:

Will data center water be going through my server?

Answer:

No - the water on the server side of the WCU is contained in a closed-loop system, maintained by IBM service personnel. There is a fill and drain tool delivered with a water-cooled system that is used to fill the system initially and to service it if needed. If maintenance is required on the closed-loop water system for WCU or Radiator systems, backup air blowers at the top of the system will continue to cool the system with possible performance reduction.

Question:

Is there any basic maintenance required because of water cooling once the system is installed?

Answer:

No scheduled maintenance is required. If service is ever needed, the IBM engineer will have the appropriate tools to work with the water systems as needed.

Question:

What will happen if the humidity or temperature of my data center has a fluctuation – will I have a problem with condensation on internal server pipes?

Answer:

The system has triple-redundant humidity sensors and will regulate the temperature of the internal water cooling loop safely above the dew point. Even if due to some fault the data center humidity goes above specification, the system will prevent condensation from occurring.

Question:

Will the zEC12 contain a rear door heat exchanger like I can purchase today for a BladeCenter rack?

Answer:

No. There is an air-to-water heat exchanger (HX) contained within the rear of each rack (not in the rear door). This heat exchanger removes heat from the systems exhaust air and transfers it to the internal water loop. This is in addition to the heat load that is removed from the high-performance processor modules by direct-attached cold plates connected to the internal water loops.

Question:

Where in the server is my data center chilled water connected?

Answer:

The data center chilled water is connected to the WCU with two connectors on each WCU - one is the inflow (supply) and the other is the outflow (return). The WCU limits the flow rate of chilled water to what is required to cool the system based on chilled water temperature and system heat load. The four insulated chilled water hoses (two supply and two return) that attach to the server pass up through the floor tile cutout at the rear of the A-frame.

Question:

What requirements are there for the data center chilled water?

Answer:

In general, the requirement is for standard building chilled water (BCW) at 6-20 degrees C (43-68 degrees F). The total hardness of the water must not exceed 200 mg/L as calcium carbonate pH must be between seven and nine. Turbidity must be less than 10 NTUs (Nephelometric Turbidity Units). Bacteria must be less than 1000 CFUs (Colony Forming Units) / ml. The water should be as free of particulate matter as feasible.

Question:

What types of hoses and connections will I need to have?

Answer:

Only IBM hoses can be used to connect to the server. The hoses are 14 feet (4m) long. The system end of each hose has a dry-break quick-disconnect coupling, and the supply and return ends are cut to length to be attached to barbed fittings sized for the 1-inch ID hose. The hoses will be attached by the IBM installation team. Shutoff valves are required on the chilled water supply and return pipes on the customer side of each connection.

Question:

If I don't want water on the zEC12, what can I expect for air cooling?

Answer:

The zEC12 is designed with a new radiator-based air-cooled system designed for more efficient cooling and improved maintenance. It is fully redundant of all active components with air backup if there is a rare plumbing problem.

The systems are shipped drained, to protect against freezing. A fill and drain tool is required on installation to fill and prior to unistall to drain.

Question:

Can you give me more information on the specs of air cooling?

Answer:

All zEC12 servers are unchanged from the z196 in terms of air flow. It is still front to back airflow, with very similar air flow (cfm) requirements. At the system level we use internal liquid cooling for high reliability, high performance processor MCM. Air temperature, altitude and humidity requirements are the same as z196 (Class1 ASHRAE/Acoustic Category 1B).

Overhead Cabling and Power

Question:

How much additional room is needed for the overhead options?

Answer:

On the zEC12 and zBC12, overhead cabling will add 12" (~30 cm) to the width (not depth) of the combined frames on the zEC12 or single frame on the zBC12. Side-cable "chimneys" or raceways are attached to the four corners of the system frames.

Overhead power will add approximately 20" to accommodate the power cord bend radius (there is only ~7" to the top for the unmated power connector and ~12" to the top of the mated power connector).

The zEC12 system's optional water connections are bottom exit only.

Question:

Will overhead options add additional weight?

Answer:

Yes and no. Overhead cabling adds at maximum configuration approximately 200 lbs (90 kg) to the weight of the zEC12 – and will be less in the zBC12 as there is only one side panel update needed. Overhead power will add no additional weight of the zEC12 or zBC12.

The key is that the panels are added after the machine is placed in position, and therefore doesn't have to be moved into place, and the heavier weight is not significant for floor loading.

Question:

Will I need to have additional space in my floor cutouts because of these 'raceways'?

Answer:

Top-exit I/O cabling will add 12" of width to the zEC12 or zBC12, but the increases will only be above the floor to the sides of the system.

Question:

Any configuration things I should watch for when ordering overhead power or overhead I/O?

Answer:

Yes. When overhead power is selected in econfig, it will force overhead I/O to be added. The converse is not true – overhead I/O selection will not force overhead power. Also, selection of non raised floor option will force both – overhead power and overhead I/O.

Question:

If my zEC12 or zBC12 is installed off a raised floor, can some cables exit the bottom of the machine?

Answer:

No. On the zEC12 and zBC12, there is no tailgate to allow bottom exit of any power or signal cables when the machine is configured for installation off of a raised floor.

Question:

If my zEC12 and zBC12 is installed on a raised floor with top exit I/O specified, can I split up my I/O cables so that some exit out the top and some exit under the raised floor?

Answer:

Yes; but, note that if top exit power is also specified, power cables MUST exit the top.

zEC12 Power requirements (including High Voltage DC Power option)

Question:

What happens to power consumption of the zEC12 if I add a zBX and blades into the infrastructure?

Answer:

The zEC12 and zBX are two separate products so they will each have their own individual power use and independent power cords.

Question:

What is static power save mode?

Answer:

Static power save mode is a function that will be available with Unified Resource Manager's Automate suite of functions. You have the ability to reduce the power consumption of the zEC12 when full performance is not required. It can be switched on and off during runtime with no disruption to currently running workloads, aside from the change in performance. You can use power save mode for periods of lower utilization (for example weekends or third shift) or for capacity backup systems where you keep them "running" but with reduced energy consumption. And systems can quickly be brought back to full performance.

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

Static power save mode is executed via commands on the HMC.

The subcapacity values reported by SCRT are designed to reflect the correct MSU values for any time when products were running under Static Power Save Mode. A detailed explanation of SCRT's support for Static Power Save Mode for different conditions will be provided in the SCRT Users Guide when the support is available.

Question:

What is query maximum potential power?

Answer:

Query maximum potential power is a component of the 'Manage' suite of functions from Unified Resource Manager. The function is implemented in the Support Element (SE) and can be used to calculate the maximum potential power draw of the zEC12 based on the configuration, the altitude of the computer room, the room temperature and the highest single fault service scenario power condition for the configuration applying reasonable tolerances. This monitoring capability can facilitate operations personnel with no System z knowledge to query the maximum potential draw of the system in order to manage the overall energy draw of the data center. It is recommended that this function be used in conjunction with the Power Estimation Tool, available via Resource Link®.

Question:

What is line cord plan ahead (FC #3001) and Balanced Power Plan Ahead (FC#2000)?

Answer:

The line cord plan ahead option allows the ability to plan ahead for the second set of line cords. It must be the same feature selection as the initial set. The Plan Ahead function provides the means of ordering now, the hardware that will be required in the future to help avoid a disruptive hardware install in the future.

Phase currents are minimized when they are balanced among the three input phases. Balanced Power Plan Ahead is designed to allow you to order the full complement of bulk power regulators (BPRs) on any configuration, to help ensure that the configuration will be in a balanced power environment.

Question:

What line cord requirements will there be for the zEC12?

Answer:

Just like the z196, the zEC12 will have full 2N redundancy and can have two or four (3-phase) line cords (if AC powered) depending on system configuration. The name plate rating of each line cord is the same as the z196 which is the same or slightly less than the System z10 Enterprise Class (z10 EC).

Question:

How many line cords will be required for High Voltage DC Power option on the zEC12?

Answer:

The same number of line cords as are required for AC operation.

Question:

What nominal DC supply voltages will be supported?

Answer:

We'll support 380 VDC – 520 VDC nominal.

Question:

What level of savings will I see using High Voltage DC Power?

Answer:

High-voltage DC power saves approximately 3% of server input power. In North America today, the data center Universal Power Supply (UPS) and Power Distribution Units (PDU) consume 8-20% of total data center power. Our current estimates are that HV DC when implemented in a new data center could save on the order of 7 – 12% of server input power.

Integrated Facility for Linux (IFL)

Question:

What is the Integrated Facility for Linux (IFL)?

Answer:

The Integrated Facility for Linux (IFL) is a processor dedicated to Linux workloads on IBM System z servers. It is an optional feature, designed, attractively priced and supported by the z/VM virtualization software and the Linux operating system.

The inclusion of an IFL will not increase charges for System z software running on general purpose (standard) processors.

Question:

Can I buy a machine that is made up entirely of Integrated Facility for Linux (IFL) engines and no traditional engines?

Answer:

Yes. It is possible to configure a server with IFL engines only. A dedicated IBM System z Linux server is available. It is named The IBM Enterprise Linux Server.

For more information, please look at the Web site: ibm.com/systems/z/os/linux/els.html or ask your IBM or Business Partner representative for more information about the IBM Enterprise Linux Server.

Question:

Where can I get more information on the IFL?

Answer:

For more information, attributes and values, please look at the Web site: ibm.com/systems/z/os/linux/solutions/ifl.html

zAAPs and zIIPs

Question:

What was the IBM System z specialty engine capability, called “zAAPs on zIIPs” originally announced on August 18, 2009?

Answer:

z/OS V1.11 and later was enhanced with a capability called ‘zAAP on zIIP’ that can enable System z Application Assist Processor (zAAP) eligible workloads to run on System z Integrated Information Processors (zIIPs). This capability can allow you to run zIIP- and zAAP-eligible workloads together on one type of specialty engine – the zIIP.

This is ideal for customers without enough zAAP- or zIIP-eligible workload to justify a specialty engine today; the combined eligible workloads may make the acquisition of a zIIP more cost effective. With the “zAAP on zIIP” capability, the combined eligible TCB and enclave SRB workloads may now make the acquisition of a single zIIP cost effective.

This capability is also intended to provide more value for customers having only zIIP processors by making Java and XML-based workloads eligible to run on existing zIIPs.

The “zAAP on zIIP” capability can help optimize server resources and simplify systems management by reducing the need to plan for and manage multiple types of specialty engines.

This was designed so that if you have zAAPs on the server and “zAAP on zIIP” is enabled, then z/OS will not honor “zAAP on zIIP” and workloads will be dispatched as normal on zAAP and zIIP engines.

In addition, current IBM terms and conditions surrounding System z specialty engines still apply. For example all servers needed to comply with the 1:1 ratio on CP to zIIP and/or zAAP. Information on the July 23, 2013 change to the 2:1 ratio is further down in this section.

Question:

How do I turn on the “zAAP on zIIP” capability?

Answer:

The “zAAP on zIIP” ships default enabled. If you wish to disable the function for any reason, you must IPL with ZAAPZIIP=NO in the IEASYSxx Parmlib member.

Question:

Is there a special price, feature, or software charge for zAAP on zIIP capability?

Answer:

No special prices, codes, charges are needed. The zAAP on zIIP capability is available at no additional charge and is in the base of z/OS V1.11 and later.

Question:

Can I convert all my zAAPs to zIIPs and then take advantage of the “zAAP on zIIP” function?

Answer:

You can explore the conversion of zAAPs to zIIPs going forward. Customer planning and testing is recommended before eliminating all zAAPs.

Question:

If I run z/OS as guest under z/VM, how does “zAAP on zIIP” capability work?

Answer:

When z/OS runs as a guest under z/VM, zAAP on zIIP is not available if zAAPs are defined in the virtual machine for z/OS or in the partition where z/VM is running.

Question:

What was announced with respect to zAAPs on August 28th, 2012 and July 23, 2013?

Answer:

The zEC12 and zBC12 will be the last System z servers to offering zAAP specialty engines – in the future we will only offer zIIP specialty engines. This is intended to provide greater simplification for customers in that they will need to plan to test and plan for fewer specialty engines while still making use of the savings achievable through specialty engines.

IBM has released PTF for APAR OA38829 on z/OS V1R12 and V1R13. This PTF allows zAAP-eligible workloads to be dispatched on zIIPs even when there are active zAAPs installed - but not in the LPAR running ‘zAAP on zIIP’.¹⁷

Current IBM terms and conditions surrounding System z specialty engines still apply.

Question:

Are there changes to the workloads that are eligible to run on zAAPs and zIIPs?

Answer:

The workload eligibility rules have not changed. See the authorized use table at http://public.dhe.ibm.com/systems/support/warranty/pdfs/aut/Authorized_Use_Table_01-2012_en_US.pdf

Question:

What did IBM announce July 23, 2013 with respect to the 1:1 zIIP/zAAP terms and conditions ratio?

Answer:

IBM announced a modification to the zIIP/zAAP ratio for zEC12 and zBC12 only servers.

The ratio will now be 2:1. For the zEC12 and the zBC12, you may purchase up to two zIIP and/or up to two zAAP processors for every general purpose processor you purchase on the server.

For z196 and z114 servers and earlier, you must still follow the 1:1 ratio.

Question:

With the change in the 2:1 ratio, how will that effect how many zAAPs and zIIPs I can have on my processor?

Answer:

Previously with the 2:1 ratio you were allowed to have one zIIP and one zAAP for every general purpose CP. As of July 23, if you have a zEC12 or zBC12 with one CP, you may have up to two zIIPs and two zAAPs on that server. If you have two CP's on a zEC12 or zBC12, you may have up to four zIIPs and four zAAPs, and so forth.

¹⁷ This APAR will be updated in 4Q13 to support the 2:1 ratio.

Question:

If I want to have a zIIP, with the 2:1 ratio, does this change mean I have to buy two zIIPs (or zAAPs) for every CP I purchase?

Answer:

No. You can purchase what you need, the change just gives the option for additional zIIP capacity should your workload need it.

Question:

Will the new 2:1 ratio apply to the z196 or z114?

Answer:

No. The change in the ratio is only for the zEC12 or zBC12.

Question:

Will the new 2:1 ratio apply to CBUs as well?

Answer:

Yes. During configuration of CBU you will be able to apply the 2:1 ratio when ordering zIIP/zAAP CBUs.

Question:

I am interested in doing "zAAP on zIIP" and today I have two CPs, two zAAPs and two zIIPs (following the 1:1 ratio). Can I start moving zAAP workload over to the zIIPs and go to two CPs and four zIIPs?

Answer:

"Technically, if you are on a zEC12 or a zBC12 it is fine as running with two CPs and four zIIPs will follow in the 2:1 ratio of zIIPs:CPs. Remember that you can do zAAPs on zIIPs even if zAAPs are installed on the server for the purpose of testing zAAP on zIIP.

With the new 2:1 ratio you might be better off just migrating your work and modifying the zAAPs to be zIIPs so you end up with two CPs and four zIIPs and no zAAPs.

Question:

Will IBM's configurators be available to make these 2:1 modifications on July 23rd?

Answer:

No. The configurators, including those on IBM Resource Link, will not be updated until 4Q13. On orders to take advantage of the 2:1 ratio, prior to the 4Q13 availability, work with your IBM sales rep or BP.

Question:

Why did IBM change the zIIP and / or zAAP to CP ratio?

Answer:

We changed the ratio to 2:1 because we have announced our intention to replace zAAP functionality with zIIPs and that we are discontinuing zAAPs on future processors. This change will allow for easier migration.

Question:

Why is doubling the ratio good for me?

Answer:

Now that zIIPs and zAAPs are merged in to one engine, and you can have twice as many as before, more workload can be offloaded. The zIIPs serve double duty (they can be assigned to Java execution when there is major demand, and DRDA processing when that is important - or both, in workload-dependent ratios) and may achieve higher utilization rates.

zEnterprise and PureSystems

Question:

Can IBM PureSystems™ and zEnterprise be complementary?

Answer:

Yes they are complimentary offerings. PureSystems can be connected to a zEnterprise through a 10 GbE connection. However it is important to understand the attached PureSystems resources will not be under the control of zManager and PureSystems. FSM cannot manage System z resources. Federation of system management can be done at a higher level though smart cloud technologies. For clients who want to tightly integrate and centrally manage distributed resources that have affinity to System z data or core applications, would be best served to host those distributed resources on a zBX and have them managed by zManager.

Question:

What was announced on **October 9, 2012** with regards to IBM PureData™ Systems?

Answer:

IBM announced a new member to its IBM PureSystems Family. In addition the recently announced IBM PureFlex™ and IBM PureApplication™ Systems they are offering a new expert integrated PureData System for transactional and Analytic Services. The PureData Systems come in three flavors:

PureData Systems for Transaction Processing - Optimized for delivering highly reliable and scalable data services for transactional applications.

PureData Systems for Analytics - A data warehouse system powered by Netezza® technology for high performance analytics and reporting on large volumes of data

PureData Systems for Operational Analytics - An operational warehouse system that balances the demands of delivering analytics to real-time decision making in business operations. It handles continuous loading of data, complex data analysis, and up to 1000 or more concurrent operational queries.

Question:

Does the introduction of PureSystems change the positioning of PureSystems and Enterprise Systems?

Answer:

No, the PureSystem portfolio continues to be focused on clients who want the greatest integration possible to simplify and accelerate their targeted solution deployment and management. In contrast the Enterprise Systems are for clients that want the highest levels of service for their critical business computing, requiring unique tailored solutions that are optimized for their business and provide economies at the enterprise scale.

Question:

Are there corresponding offerings to PureData Systems that run on System z?

Answer:

Answer: Yes. For each of the three new PureData Solutions system z offers corresponding solutions that add the characteristics associated with Enterprise Class Systems. The following table summarizes the corresponding offerings:

PureData offering	System z corresponding offering
PureData System for Transaction Processing	DB2 on z/OS with DataSharing
PureData System for Analytics	IBM DB2 Analytics Accelerator or as part of a zEnterprise Analytics System 9700
PureData System for Operational Analytics	zEnterprise Analytics System 9700 with IBM DB2 Analytics Accelerator

Question:

When might I choose a PureData System versus a Solution on System z?

Answer:

The basics start with the overall positioning of PureSystem versus Enterprise Systems. The PureSystem portfolio continues to be focused on clients who want the greatest integration possible to simplify and accelerate their targeted solution deployment and management. Enterprise Systems are for clients that demand the highest levels of service for their business computing, find incremental value in building unique tailored solutions that are optimized for their business providing economies at enterprise scale. PureData Systems enable clients who want to capitalize on an expert integrated, fit for purpose, system for transactional and analytic coupled with PureFlex and PureApplication platforms. These modular systems are designed and optimized to build from basic building blocks to robust systems. PureData also offers clients an improvement in security and availability over competitive UNIX and x86 offerings. PureData System provides an out-of-the-box transaction processing and analytic services solution for clients focused on the quickest time to value.

For comparison, the System z solutions are for clients that want a single centrally managed hybrid technology platform that is workload optimized. Their requirements demand a system that provides the ultimate security, industry-leading resiliency, and multi-platform governance that dynamically delivers resources to workloads the business finds most critical. System z transaction processing and analytic solutions are for clients that require a next generation system that integrates workload-optimized OLTP and DW/BI for operational analytics. In addition, these clients also need the economies of data mart consolidation coupled with application servers either co-located with data, or connected in zBX. The clients also are looking for solutions that provide the scalability to support 1000s of users with high volume transactional and analytic environments and want to deliver these large scale environments with the lowest cost per transaction or query possible. System z solutions provide a single footprint solution to today's need for tighter integration of all applications with minimal data movement or network traffic.

Security

Question:

What is the security technology rating?

Answer:

IBM Processor Resource/Systems Manager (PR/SM) on the IBM zEnterprise EC12, with Driver Level D12K, was certified effective February 19, 2013 at a Common Criteria Evaluation Assurance Level (EAL) 5+. Common Criteria (CC) certification at EAL5+ provides the assurance that the process of specification, implementation, and evaluation of a computer security product has been conducted in a rigorous and standardized fashion against a set of internationally recognized criteria. This certification is indicative of the scrutiny that an independent third party evaluates IBM System z. This helps to assure that many disparate applications running in different operating environments in separate logical partitions on one IBM zEnterprise EC12 server will be isolated and distinct from each other. The zEnterprise EC12 servers now join the IBM zEnterprise 196 and IBM zEnterprise 114 servers, and previous IBM mainframes, as the world's only servers with the highest level of hardware security certification, Common Criteria Evaluation Assurance Level 5+.

The zBC12 server is designed to meet the EAL 5+ criteria for Common Criteria certification.

Question:

Tell me about the Crypto Express4S feature that was announced August 28, 2012.

Answer:

The Crypto Express4S feature, exclusive to zEC12 and zBC12, is IBM's latest tamper-resistant cryptographic coprocessor that is installed in a PCIe I/O drawer. This feature provides a secure programming and hardware environment wherein crypto processes are performed. It is suited to applications requiring high-speed security-sensitive cryptographic operations for data encryption and digital signing, and secure management and use of cryptographic keys.

The Crypto Express4S feature can be configured in one of three ways: as a secure IBM Common Cryptographic Architecture (CCA) coprocessor, as a secure IBM Enterprise PKCS #11 (EP11) coprocessor or as an accelerator.

Question:

What is the IBM Enterprise PKCS #11 coprocessor?

Answer:

The IBM Enterprise Public-Key Cryptography Standards (PKCS) #11 (EP11) Licensed Internal Code (LIC) is designed to provide open industry-standard cryptographic services. Based on PKCS #11 specification v2.20 and more recent amendments, EP11 leverages the IBM Crypto Express4S feature to provide enhanced firmware capabilities. This firmware is designed to meet the rigorous FIPS 140-2 Level 4 and Common Criteria EAL 4+ certifications. The Crypto Express4S configuration option is designed to meet public sector and European Union requirements where standardized crypto services and certifications are needed.

EP11 supports secure PKCS #11 keys – keys that never leave the secure boundary of the coprocessor unencrypted. The prior PKCS #11 implementation, which supported only clear keys, was provided by z/OS. Key protection was accomplished solely by Resource Access Control Facility (RACF®) dataset protection. With EP11, keys can now be generated and securely wrapped under the EP11 Master Key, all within the bounds of the coprocessor. Thus, EP11 provides enhanced security qualities when using PKCS #11 functions. EP11 is exclusive to IBM zEnterprise EC12 and IBM zEnterprise BC12 and is supported by z/OS and z/VM (for guest exploitation).

The TKE workstation is required to manage a Crypto Express4S feature that is configured as an EP11 coprocessor. The TKE smart card reader and smart cards (FC #0885) are also mandatory to manage the EP11 card (and optional when the card is configured as a CCA coprocessor).

Enhancements to PKCS#11 for July 23, 2013 exclusive to the Crypto Express4S feature and are available on zBC12 and select zEC12 servers. They are supported by z/OS and by z/VM for guest exploitation.

PKCS #11 v2.1 Probabilistic Signature Scheme (PSS): EP11 now supports the latest algorithm that is used in digital signature applications, offering enhanced security characteristics over prior digital signature algorithms.

EP11 Key agreement algorithms supported:

- Diffie-Hellman: 1024-bit, 2048-bit
- Elliptic Curve Diffie-Hellman
- National Institute of Standards and Technology (NIST): 192-bit, 224-bit, 256-bit, 384-bit, 521-bit
- Brainpool: 160-bit, 192-bit, 224-bit, 256-bit, 320-bit, 384-bit, 512-bit

Offload Generation of Domain Parameters: The generation of a PKCS #11 secret key (via the CSFPGSK/CSFPGSK6 APIs) requires a domain parameter, which is generated as part of the process based on the key type attribute. With the enhancements to the CEX4S the generation of the domain parameter can be done on the P11 coprocessor instead of in software, reducing consumption of CPU resources.

Question:

What are the Common Cryptographic Architecture (CCA) enhancements?

Answer:

When the Crypto Express4S feature is configured as a CCA coprocessor, the CCA functions from Crypto Express3 are still available and supported.

New functions from February 24, 2014. These CCA enhancements are available on zBC12 and select zEC12, z196, and z114 servers. They are supported by z/OS and by z/VM for guest exploitation:

Die Deutsche Kreditwirtschaft (DK) AES PIN support: The German banking industry organization, DK, has defined a new set of Personal Identification Number (PIN) processing functions to be used on the internal systems of the banks and servers. CCA is designed to support these functions that are essential to those parts of the German banking industry governed by DK requirements. The functionality includes key management support for new AES key types, AES key derivation support, and several DK specific PIN and administrative functions. The intellectual property rights regarding the methods and specification of this support belong to the German Banking Industry Committee. DK is an association of the German banking industry. The German Banking Industry Committee is the hybrid term, in English, for Die Deutsche Kreditwirtschaft. Prior to August 2011, DK was named ZKA for Zentraler Kreditausschuss, or Central Credit Committee.

New Message Authentication Code (MAC) support: CCA now supports new message authentication codes using the Advanced Encryption Standard Cipher-based MAC (AES-CMAC) algorithm (Note available March 31, 2014)

User Defined Extension (UDX) simplification for PKA Key Translate: The Integrated Cryptographic Service Facility (ICSF) and CCA are designed to allow businesses to create extensions to the base CCA services. All UDX services are provided under contract to specific customers by IBM Global Business Services. With this announcement, ICSF and CCA add support for a UDX to the base CCA services. Support is added for translating an external RSA CRT key into new formats. These new formats use tags to identify key components. Depending on which new rule array keyword is used with the PKA Key Translate callable service, the service TDES encrypts those components in either CBC or ECB mode. (Note available March 31, 2014)

Functions from July 23, 2013. These CCA enhancements are available on zBC12 and select zEC12, z196, and z114 servers. They are supported by z/OS and by z/VM for guest exploitation:

Export Triple Data Encryption Standard (TDES) key under Advanced Encryption Standard (AES) transport keys. The AES encryption algorithm has greater security than the TDES encryption algorithm. CCA has added the ability to use AES key-encrypting keys (KEKs) to wrap your TDES keys to help begin moving to AES for key management. All of the TDES key wrapping functions are still available, but a parallel set of AES wrapping functions are now available for use.

Diversified Key Generation Cipher Block Chaining (CBC) support. During the Europay, Mastercard and Visa (EMV) smart card personalization process, session keys are derived and then used to secure messages to the EMV cards. Some EMV card personalization specifications require the use of TDES CBC mode to derive these session keys. This enhancement adds that capability to the existing key derivation options in CCA.

Initial PIN Encrypting Key (IPEK) support. The IPEK is the initial key that is loaded into a point-of-sale (POS) terminal before it is deployed for use, when that terminal will use the Derived Unique Key Per Transaction (DUKPT) key protocol. CCA has added a function that allows the Hardware Security Module (HSM) to securely derive an IPEK and return it to the application program in an encrypted key token, which can then be securely installed in a POS terminal.

Remote Key Export (RKX) key wrapping method support. In a previous release, CCA added the capability to wrap keys using a proprietary enhanced mode algorithm. This included the ability to set a default preference for the wrapping method to be used, as well as options to override that default in most CCA functions. The RKX function now supports that ability as well.

Integration of User Defined Extensions (UDX) into CCA. A UDX is designed to allow you to add custom functions to the CCA application programming interface (API) running in the Hardware Security Module (HSM). CCA has included the following three UDXs in the standard CCA APIs, avoiding the requirement for a UDX: Recover PIN from Offset, Symmetric Key Export with Data, and Authentication Parameter Generate. This means these functions can now be invoked using the standard APIs, avoiding the requirement to install and manage a UDX.

These are functions announced in August 2012:

Improved wrapping key strength. In order to comply with cryptographic standards, including ANSI X9.24 Part 1 and PCI-HSM, a key must not be wrapped with a key weaker than itself. With this release, CCA allows you to configure the coprocessor to ensure that your system meets these key wrapping requirements. This Crypto function is exclusive to zEC12 and the zBC12 and is supported by z/OS and z/VM.

DUKPT for Message Authentication Code (MAC) and encryption keys. Derived Unique Key Per Transaction (DUKPT) is defined in the ANSI X9.24 Part 1 standard. It provides a method in which a separate key is used for each transaction or other message sent from a device. This makes it so that an attacker who is able to discover the value of a key would only be able to gain information about a single transaction and not about any that preceded it or that follow it. The original definition of DUKPT only allowed derivation of keys to be used in encryption of personal identification number (PIN) blocks. The purpose was to protect PINs that were entered at a point-of-sale (POS) device and then sent to a host system for verification. Recent versions of X9.24 Part 1 expanded this so that DUKPT can also be used to derive keys for MAC generation and verification, and for data encryption and decryption. Three separate variations of the DUKPT key derivation process are used so that there is key separation between the keys derived for PIN, MAC, and encryption purposes. This function is exclusive to zEC12 and zBC12 and is supported by z/OS and z/VM.

Secure Cipher Text Translate2 (CTT2). CTT2 is a new data encryption service that takes as input data encrypted with one key and returns the same data encrypted under a different key. This service has the advantage that it provides the ability to securely change the encryption key for cipher text without exposing the intermediate plain text. This Crypto function is exclusive to zEC12 and zBC12 and is supported by z/OS and z/VM. The Cipher Text Translated2 API also provides a way to migrate secure key ciphertext from DES/TDES to AES encryption.

Compliance with new random number generation standards. The standards defining acceptable methods for generating random numbers have been enhanced to include improved security properties. Now, random number generation in the Crypto Express4S feature when defined as a coprocessor conforms to the Deterministic Random Bit Generator (DRBG) requirements defined in NIST Special Publication 800-90/90A, using the SHA-256 based DRBG mechanism. The methods in these NIST standards supersede those previously defined in NIST FIPS 186-2, ANS (X9.31), and ANS (X9.62). With these improvements, client applications can help to meet the timeline outlined in Chapter 4 of NIST SP800-131 for switching to the new methods and ceasing use of the old methods. This Crypto function is exclusive to zEC12 and zBC12 and is supported by z/OS and z/VM.

CCA Enhancements for applications supporting American Express EMV cards. Two changes have been made to the CCA application programming interface (API) to help improve support of payment card applications for American Express EMV cards. The Transaction_Validation service is used to generate and verify American Express card security codes (CSCs). This release adds support for the American Express CSC version 2.0 algorithm. The PIN_Change/Unblock verb is used for PIN maintenance. It prepares an encrypted message portion for communicating an original or replacement PIN for an EMV smart card. The verb embeds the PINs in an encrypted PIN block using information supplied. With this CCA enhancement, PIN_Change/Unblock adds support for the message. This crypto function is available on zEC12, zBC12 and select z196, z114, System z10, and IBM System z9 servers and is supported by z/OS and z/VM.

Question:

And what about the accelerator?

Answer:

When the Crypto Express4S feature is configured as an accelerator, it is optimized for Secure Sockets Layer (SSL) acceleration and clear key RSA operations.

Question:

What are the new enhancements for Trusted Key Entry (TKE) 7.3?

Answer:

Full-function migration wizard for EP11: The full-function migration wizard is designed to provide the ability to quickly and accurately collect and apply data to the Crypto Express features configured as EP11 coprocessors. This wizard previously supported CCA, and has now been enhanced to also support EP11. (Please note that the zEC12 requires Driver 15 for this support.)

Workstation setup wizard: The setup wizard performs the most common TKE workstation initialization functions, ensuring speed and accuracy of new TKE hardware deployment. It simplifies the process while greatly reducing errors. The wizard can also be run to verify the TKE workstation has been configured correctly.

Allow Set Master Key from the TKE workstation: Initially setting or changing any type of master key on a Crypto Express feature must be done carefully. If a master key is set or changed when key stores have not been properly prepared for the new master key, the keys in the store will become unusable. In a recovery situation, quickly loading the correct master key quickly is critical. The TKE workstation will allow you to set any master key from the TKE workstation without having to perform an additional step from a TSO session.

Restricted PIN support: The latest CCA enhancements are designed to allow users to prevent the automatic generation of certain PIN values, or the replacement of existing PINs with certain PIN values. The TKE 7.3 LIC includes a new tab for specifying restricted PIN values. This enhancement is exclusive to the TKE 7.3 LIC.

New AES operational keys: Five new AES operational keys can be managed from the TKE 7.3 workstation. The key types are MAC, PINCALC, PINPROT, PINPRW, and DKYGENKY.

Close Host and Unload Authority Signature Key: The Close Host enhancement is designed to allow you to explicitly sign off a host. The Unload Authority Signature Key enhancement allows you to explicitly remove the current authority signature key without ending the TKE application. When you have many users with different roles, users no longer have to end the TKE application before the TKE workstation is utilized by another user.

New access control for managing host list entries: The TKE workstation profile role has a new access control point to allow you to create, change, or delete a host list entry. This is designed to provide stronger separation of duties between users of a host list entry and users that manage the entries.

Domain Group changes:

- When a user is creating or changing a domain group, a domain can only be included in the group once. This ensures that domain commands are only sent to a domain once.
- If you manage a host crypto module role from a domain group, the user must explicitly select which Domain Access Control Points are to be set. The user specifies that either every domain access control point is selected for every crypto module in the group, or only the domain access control points for the domains in the group are selected. This enhancement allows you to manage a 'module-scoped role' from inside a domain group.

User-defined CCA and EP11 Domain Control lists: When managing CCA or EP11 Domain Control Points, the user can save the settings to a file which can then later be applied to other domains. This enhancement allows for fast and accurate deployment of new or recovered domains.

Increased session key strength: When using the latest version of smart cards on a TKE 7.3 workstation, a 256-bit AES session key will be used for all smart card operations. Refer to the TKE Workstation User's Guide, TKE Version 7.3, SC14-7511, in the Library, Hardware products for servers, TKE workstation section of Resource Link for further information. Access Resource Link at: www.ibm.com/servers/resourcelink

Note: The TKE 7.3 LIC (#0872) is supported on zBC12, zEC12, z196, and z114.

Question:

What is the new Elliptic Curve Cryptography Digital Signature Algorithm (ECDSA)?

Answer:

Elliptic Curve Cryptography Digital Signature Algorithm (ECDSA) is an emerging public-key algorithm. ECC is based on a computationally harder mathematical problem than RSA. It provides the same or better security (per-key-bit) with much shorter key lengths than RSA keys. This makes ECC appropriate in resource constrained environments such as smart cards, that may have limited space for the storage of key material. ECC is capable of providing digital signature functions and key agreement functions. The new IBM Common Cryptographic Architecture (CCA) provides ECC key generation, digital signature generation and verification functions that are compliant with the ECDSA method described in ANSI X9.62 "Public Key Cryptography for the Financial Services Industry". With RSA impractical at key lengths with strength-per-key-bit equivalent to AES-192 and AES-256, the strength-per-key-bit is substantially greater with an algorithm that uses elliptic curves.

Question:

What restrictions or limitations, if any, are associated with the use of Elliptic Curve Cryptography?

Answer:

First, the user of zEnterprise hardware cryptography is not authorized to use an Elliptic Curve Cryptography Feature's Machine Code as a commercial certificate authority.

Second, upon transfer of the zEnterprise hardware cryptography feature (or machine containing the feature) to another party, the customer must provide the other party with a copy of Licensed Internal Code (LIC) license and the license's Addendum for Elliptic Curve Cryptography.

Question:

Where can I get more information on IBM's encryption offerings?

Answer:

The best place to get started is the Enterprise security website at: www.ibm.com/systems/z/solutions/security.html. There are good links to the solutions as well as available papers and case studies.

Note:

All new zEC12 enhancements for July 23, 2013, will require Driver 15, available on September 20, 2013.

IBM System Storage & zEnterprise

Question:

How has the DS8870 been optimized for System z?

Answer:

IBM System Storage DS8870, IBM's flagship storage system, has had a unique relationship with the IBM mainframe that dates back almost half of a century. IBM architects, designs, builds, and tests System z and DS8000 systems together, and this long history of close collaboration and tight integration has been providing unique benefits for our clients for decades.

This unparalleled integration includes advanced support of the System z channel architecture and data-access methods, as well as self-optimization features such as intelligent caching algorithms, automated drive tiering and integrated QoS management. In addition to performance optimization, DS8000 and System z also boast the highest levels of business continuity with advanced mirroring and failover capabilities.

Here is a sample of the benefits from the unmatched integration between DS8870 and System z:

- Run operational analytics solutions with the fastest query response times on the market
 - DB2 query performance can improve up to 8X by exploiting DS8000 with SSDs, the List Prefetch Optimizer caching algorithm and High Performance FICON (zHPF)
- Improve the efficiency of high performance workloads
 - IBM DS8870 with the new High Performance Flash Enclosure (HPFE) accelerates System z database performance by up to 3.2x
 - DS8870 elevates performance to a new level with an upgraded IBM POWER7 controllers and up to 1 terabyte of system cache.
 - DS8870 offers hybrid (multi-tier) and all-flash configurations for the ultimate in flash optimization flexibility.
 - DS8870 with Easy Tier™ offers up to 4x automatic increase for multi-tier configurations, 3x automatic increase for single-tier configurations, and up to 10x automatic increase for the new Easy Tier Application API
 - The DS8870 with HPFE for 4x faster flash performance in 50% less space than existing flash options
- Designed to support near-continuous operations the highest levels of availability
 - Together, DS8870 Metro Mirror and GDPS HyperSwap are designed to provide over six-nines system availability for the most critical applications. In fact, over 75% of the worlds leading banks rely on this solution for their most important applications.
- Protect and secure critical enterprise information from external and internal threats
 - DS8870 ships every system with innovative self-encrypting drives
- Enable cost-effective storage consolidation
 - DS8870 series supports an extensive variety of server platforms, drive tiers, and application workloads
 - Support for the new z/OS zEnterprise Data Compression and zEDC Express makes more efficient use of your DS/8000 capacity – both flash and traditional spinning drives.

These are just a few of the capabilities the IBM System Storage DS8870 series offers that enable it to support the highest levels of enterprise class performance, reliability and continuous operations for System z environments.

Question:

How does IBM Smarter Storage support System z workload consolidation and optimization?

Answer:

The IBM Smarter Storage portfolio offers a variety of storage systems that can support the flexibility and scalability requirements for workload consolidation and optimization.

The DS8870 can scale up to over 2 petabytes of data and supports a wide range of server platforms and operating systems, including z/OS, z/VSE, z/VM, z/TPF, Linux on System z. It can support a combination of traditional spinning drives of varying rotational speeds, SSDs and ultra-fast flash cards (within the High Performance Flash Enclosure) in both hybrid (multi-tier) and all-flash configurations. In addition, to help simplify management and accommodate changing application requirements, administrators can easily and nondisruptively add and expand storage capacity with Dynamic Volume Expansion. Thin provisioning is another feature that can significantly reduce the time administrators spend provisioning new storage while keeping applications online.

The DS8870 also offers advanced self-optimizing performance capabilities, such as automated tiering, dynamic workload prioritization and intelligent caching algorithms that can maintain the right performance and QoS for consolidated applications. DS8870 also includes a fully-redundant design and sophisticated data integrity, as well as a variety of high-availability features designed to keep mission-critical applications online.

IBM has a variety of enterprise tape and virtual tape offerings that can also provide capabilities to help clients support their datacenter consolidation initiatives.

Question:

Besides the DS8870, what IBM Storage alternatives are available to support System z operating systems, such as Linux on System z, z/VM, z/VSE or z/TPF?

Answer:

Linux on System z, z/VSE and z/VM are supported on the XIV[®] Storage System, Storwize[®] V7000 and SAN Volume Controller (SVC). As clients consider migrating workloads from distributed servers to the integrated System zEnterprise BladeCenter Extension (zBX) server blades, many will prefer to use their existing storage since they're familiar with the management techniques, tooling, etc. For those considering new storage for their distributed data, DS8870 is an ideal choice for consolidating all their System z operating systems. Alternatively, the XIV Storage System, Storwize V7000 and SAN Volume Controller (SVC) also support Linux on System z, z/VSE and z/VM.

Question:

What IBM Storage offerings are available to support the midsize zBC12 or z114?

Answer:

DS8870 offers a single-frame Business Class configuration, as well as a new Entry-Class option that are perfectly suited for zBC12 or z114 environments with streamlined, lower cost configurations than the standard Enterprise Class DS8870 configuration. These entry-level configuration options support all System z operating systems and reduce the number of device adapters and I/O enclosures, yet increases the number of drives available in the frame. This effectively lowers the price point for zBC12 and z114 environments.

Question:

What does IBM Storage offer for long-term data protection for zEC12 and zBC12 environments?

Answer:

To help organizations ensure business continuity and cost-effectively protect vital data and business processes, IBM offers the Virtualization Engine TS7700, a mainframe virtual tape solution with advanced management features to optimize tape processing. The TS7700 Virtualization Engine enables clients to implement a fully-integrated tiered storage hierarchy of disk and tape and leverage the benefits of tape virtualization. To ensure high availability and rapid disaster recovery, the TS7700 Virtualization Engine can be deployed in a variety of grid configurations, designed to help eliminate downtime during planned and unplanned outages, upgrades and routine maintenance.

In addition, the IBM System Storage TS3500 Tape Library is designed to provide a highly-scalable, automated tape library for both mainframe and open systems backup and archive that can scale to support the largest System z Enterprise Class environments. The TS3500 can scale up to 240 frames and up to 225,000 cartridges. It's has a multi-path architecture that is designed to support redundant control paths for high availability. It's self-encrypting and provides WORM (Write Once Read Many) technology so the data cannot be overwritten, which helps support regulatory compliance needs. It also supports open servers, as well as z/OS, z/Linux, z/VM and z/VSE.

Question:

What does IBM Storage offer in terms of securing stored data?

Answer:

IBM Storage offers innovative self-encrypting disk and tape offerings that can secure clients' data at rest. These offerings include self-encrypting drives that are designed to encrypt data as it's written to the drive and decrypt it as it's read from the drive. Unlike competitive solutions, each drive has an embedded encryption engine that enables performance to match that of traditional non-encrypting drives. This drive-level encryption approach reduces the risk that information could be compromised when storage media are physically removed from the storage systems for archiving.

IBM Security Key Lifecycle Manager for z/OS manages the encryption keys for storage, simplifying deployment and maintaining availability to data at rest natively on the System z mainframe environment. Security Key Lifecycle Manager for z/OS simplifies key management and compliance reporting for privacy of data and compliance with security regulations. It is designed to help manage the growing volume of encryption keys across an organization with simplified deployment, configuration and administration of key generation, as well as key life cycle management.

IBM storage encryption solutions are designed to be transparent to the operating system, applications, databases, system administrators and users, making deployment extremely simple.

Question:

What is the IBM SAN Volume Controller and IBM FlashSystem bundle and it's relationship to Linux?

Answer:

IBM has recently improved the IBM FlashSystem bundle, creating a new product offering, FlashSystem V840. The speed and MicroLatency of Flash, and the flexibility and software defined storage advanced capabilities have been combined into a single offering that scales in both performance (up to 2.5M IOPS) and capacity (from 2TB up to 320TB), while serving IO to popular servers and their operating system. Linux, including zLinux is listed in the SSIC (insert URL hot spot here). The FlashSystem V840 is an SSR installed process, so an IBM Customer Engineer performs the racking etc to get the FlashSystem V840 up and operational. Lab Based Services also provides a set of services to provide rapid time to value, so that you can realize your investment faster.

Using the Fiber Channel Protocol (FCP) adapters the SVC had originally been qualified on Red Hat Enterprise Linux RHEL4 and SUSE SLES9 SP3. It continues to be supported on all releases above the first supported level.

Question:

What services are part of the IBM FlashSystem and SVC Integration bundle?

Answer:

The STG Lab Services System Storage team will provide implementation and virtualization services for IBM FlashSystem with IBM SVC.

Answer: In a typical basic V840 implementation, the STG Lab Services System Storage team will provide implementation and virtualization services for IBM FlashSystem V840. This service focuses on configuring the control and storage enclosure to create a single cluster configuration and consists of

- Ensuring all code updates are at the latest and necessary levels
- Assisting with the connecting and provisioning the V840
- Configuring the V840 using best practices
- Creation of Vdisk mirrors for dual flash implementations and backup
- Providing skills transfer to designated technical staff

A range of additional services can be provided based on your specific needs. Ask your IBM or BP Storage Specialist.

Question:

If I have an existing SVC can I just purchase a FlashSystem and use it with my Linux on System z

Answer:

The IBM FlashSystem 710, 720, 810 and 820 are all supported when used in conjunction with the IBM San Volume Controller at maintenance level 6.4.1.5 and above (including stream 7.1). For a full description of IBM's firmware support levels and a full listing of IBM storage interoperability the following URL's should assist you in finding the best match for your environment.

Device support matrix can be found at: <http://www.ibm.com/support/docview.wss?uid=ssg1S1004392> and <http://www.ibm.com/support/docview.wss?uid=ssg1S1004366>

Question:

How do I order the IBM SVC and IBM FlashSystem bundle?

Answer:

Please see your IBM Storage sales team or IBM business partner to order the IBM FlashSystem products or FlashSystem bundle.

Services offerings are available from IBM Systems Lab Services and Training. To learn more about IBM FlashSystem services and other related products and services, contact stgls@us.ibm.com or visit: ibm.com/systems/services/labservices

Question:

Where do I go for more information?

Answer:

ibm.com/storage/flash

STP

Question:

What has changed regarding how the zEnterprise Systems (zEC12, zBC12, z196 and z114) are time synchronized with other System z servers?

Answer:

The zEnterprise does not support attachment to the Sysplex Timer®. However, the zEnterprise can use Server Time Protocol (STP), which has been generally available since January 2007. Server Time Protocol (STP) is designed to allow events occurring in different servers to be properly sequenced in time. STP is designed for servers that have been configured in a Parallel Sysplex or a basic sysplex (without a coupling facility), as well as servers that are not in a sysplex but need time synchronization.

Question:

How can the zEnterprise be time synchronized to an existing External Time Reference (ETR) network of System z10 servers?

Answer:

First, migrate the existing ETR network to at least an STP Mixed Coordinated Timing Network (CTN). Once this is done, a zEnterprise can join the Mixed CTN as a Stratum 2 server. The zEnterprise can use STP to synchronize with one of the System z10 servers still synchronized to the Sysplex Timer. The servers synchronized to the Sysplex Timer are Stratum 1 servers. It is highly recommended that there are at least two Stratum 1 servers in a Mixed CTN, so as to not have a single point of failure. The STP feature is required on the Stratum 1 servers (the System z10 servers in this example).

Question:

Can I enter an ETR ID on the Server Time Protocol (STP) panels such that a zEC12 machine can participate in a Mixed Coordinated Timing Network (CTN)?

Answer:

Yes, this is possible and required in order to join a mixed CTN. A zEC12 can be a Stratum 2 or Stratum 3 server in an STP Mixed CTN. However a zEC12, zBC12, z196 or z114 machine cannot be a Stratum 1 server in a Mixed CTN since it cannot attach to an IBM Sysplex Timer. Time synchronization can only be established by using STP. For migration purposes it is possible to enter the ETR ID on the zEnterprise HMC to allow it to join a Mixed CTN that has System z10 servers attached to the Sysplex Timer. Note that any new high end server announced following the zEC12 will NOT be able to join a Mixed CTN.

Question:

Can a zEnterprise be configured in a Mixed CTN or STP-only CTN with IBM eServer zSeries 990 or IBM eServer zSeries 890 (z990,z890) servers?

Answer:

No. The zEnterprise servers are designed to coexist in the same CTN with (n-2) server families. This allows zEnterprise servers to participate in the same CTN with System z10 servers, but not with System z9 servers.

Question:

Can I use the same External Time Source options for STP on the zEC12?

Answer:

Yes. STP on the zEC12 still offers two options for external Time Source - Network Time Protocol (NTP) server and NTP server with Pulse per second (PPS). The IBM zEC12 or zBC12 do not support dial-up modems for use with the Remote Support Facility (RSF), or the External Time Source (ETS) option of Server Time Protocol (STP). The currently available Network Time Protocol (NTP) server option for ETS as well as Internet time services available using broadband connections can be used to provide the same degree of accuracy as dial-up time services.

Question:

What STP enhancement was announced on August 28, 2012?

Answer:

Enhanced broadband security is provided with Network Time Protocol (NTP) authentication support, when an NTP server is accessed to get accurate time for the STP Coordinated Timing Network (CTN). This helps ensure the time packets are not tampered with. New panels were added to the HMC to define and manage Symmetric Key and Autokey Trusted Certificates.

Additional support allows one to issue and receive NTP commands from the HMC panels.

Question:

What new recovery enhancements were introduced for STP and the zEnterprise System in September 2011?

Answer:

In September 2011, STP was enhanced to consistently block a potential disruptive action on any server that has an STP role (PTS, BTS or Arbiter). With specific MCLs installed on machines with these roles, STP forces the role to be removed or reassigned prior to proceeding with a disruptive action.

See <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102019> for details.

An STP enhancement was made available to the z10, z196 and z114. This also applies to the zEC12 and zBC12. The enhancement affects how recovery works and can protect the environment from an entire Coordinated Timing Network (CTN) outage.

See <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102037> for details.

Question:

What recovery enhancement was introduced for STP on July 12, 2011?

Answer:

The new generation of host channel adapters (HCA3-O or HCA3-O LR), introduced for coupling, have been designed to send a reliable unambiguous "going away signal" (GOSIG) to indicate that the server on which the HCA3 is running is about to enter a failed (check stopped) state. When the GOSIG sent by the Current Time Server (CTS) in an STP-only Coordinated Timing Network (CTN) is received by the Backup Time Server (BTS), the BTS can safely take over as the CTS without relying on the previous recovery methods of Offline Signal (OLS) in a two-server CTN or the Arbiter in a CTN with three or more servers.

This enhancement is exclusive to the zEnterprise (zEC12, zBC12, z196 and z114) and is available only if you have an HCA3-O or HCA3-O LR on the CTS communicating with an HCA3-O or HCA3-O LR on the BTS. Note that the already available STP recovery design is still available for the cases when a GOSIG is not received or for other failures besides a server failure.

Parallel Sysplex

Question:

What was announced in the February 24, 2014 announcement concerning Flash Express supporting the Coupling Facility Control Code (CFCC)?

Answer:

IBM is fulfilling a statement of direction from July 2013.

The Flash Express feature is designed to offer improved availability and performance for key workloads at critical processing times. Flash memory is designed to help improve availability during workload transition periods and spikes such as start-of-day processing or batch online transitions.

With enhancements in CFCC Level 19, exploitation of the Flash Express feature is designed to help improve resiliency while providing cost-effective standby capacity to help handle the overflow of WebSphere MQ shared queues. You can now specify overflow areas for certain Coupling Facility list structures in the Storage Class Memory (SCM) provided by the Flash Express feature. This is designed to allow structure data to be migrated to Flash Express memory as needed and migrated back to real memory to be processed. When using WebSphere MQ for z/OS Version 7 (5655-R36), this new capability is expected to help provide significant buffering against enterprise messaging workload spikes and to help provide support for storing very large amounts of data in shared queue structures, potentially allowing several hours of data to be stored without causing interruptions in processing.

Question:

What is Coupling Thin Interrupts?

Answer:

Coupling Thin Interrupts is designed to improve the efficiency of environments where shared engines are used as Coupling Facilities. While dedicated engines continue to be recommended to obtain the best Coupling Facility performance, Coupling Thin Interrupts may help to facilitate the use of a shared pool of engines, helping to lower hardware acquisition costs.

You may now experience CF response time improvements or more consistent CF response times when using Coupling Facilities with shared engines. This may also allow more environments with multiple CF images to coexist in a server, and share CF engines with reasonably good performance. The response time for asynchronous CF requests may also be improved as a result of using Coupling Thin Interrupts on the z/OS host system, regardless of whether the CF is using shared or dedicated engines.

Coupling Thin Interrupts is exclusive to the zEC12 and zBC12, and to z/OS V1.12 and above (with PTFs).

Question:

What was announced about Flash Express supporting the Coupling Facility?

Answer:

IBM announced a statement of direction. It intends to provide exploitation of the Flash Express feature (#0402) on zEC12 and zBC12 servers with Coupling Facility Control Code (CFCC) Level 19 for certain Coupling Facility list structures in the first half of 2014. This new function is designed to allow list structure data to be migrated to Flash Express memory as needed when the consumers of data do not keep pace with its creators for some reason, and migrate it back to real memory to be processed.

When using WebSphere MQ this new capability is expected to provide significant buffering against enterprise messaging workload spikes and provide support for storing very large amounts of data in shared queue structures, potentially allowing several hours' data to be stored without causing interruptions in processing.¹⁸

Question:

What was introduced on July 12, 2011 in relation to coupling links?

Answer:

IBM introduced a means to simplify your Parallel Sysplex connectivity requirements and do more with less, with the latest generation of host channel adapters (HCA3s) for coupling - the HCA3-O fanout for 12x InfiniBand (12x IFB) with improved service times using the 12x InfiniBand3 (12x IFB3) protocol and the HCA3-O LR fanout for 1x InfiniBand (1x IFB) with four ports of connectivity and optional additional subchannels for extended-distance solutions.

12x InfiniBand coupling links support a link data rate of 6 gigabytes per second (GBps) at distances up to 150 meters (492 feet). 1x IFB coupling links support a link data rate of 5 gigabits per second (Gbps) server-to-server or 2.5 or 5 Gbps when using dense wavelength division multiplexers (DWDMs). The 1x IFB coupling links support distances up to 10 km (6.2 miles) unrepeated. Greater distances are possible with an RPQ or when using qualified DWDM solutions.¹⁹

Both 12x and 1x InfiniBand coupling links also allow you to share physical links by defining multiple logical links (channel path identifiers - CHPIDs) - up to 16 CHPIDs across two ports for 12x IFB, and up to 16 CHPIDs across four ports for 1x IFB.

Note: The 12x InfiniBand HCA3-O fanout can connect to the 12x InfiniBand HCA2-O fanout on a z114, z196, or z10, but an HCA3-O fanout cannot connect to an HCA1-O fanout on a z9.

Question:

What benefits will I receive converting to the InfiniBand protocol and HCA3 adapters?

Answer:

The 12x IFB3 protocol is designed to deliver improved service times. When no more than four CHPIDs are defined per HCA3-O port, the 12x IFB3 protocol is used. When using the 12x IFB3 protocol, synchronous service times are designed to be 40% faster than when using the 12x IFB protocol.

The 1x InfiniBand coupling links are designed to offer improved physical connectivity: The HCA3-O long reach (LR) fanout for 1x InfiniBand coupling links has been designed with four ports per feature to satisfy requirements for more physical connectivity. The added connectivity will be helpful as clients migrate from InterSystem Channel-3 (ISC-3) to 1x InfiniBand coupling links. Up to 32 subchannels (devices) per CHPID for 1x InfiniBand coupling links: To improve link utilization and coupling throughput at increased distances between the Coupling Facility (CF) and the operating system or between CFs separated by increased distances, we are now optionally supporting 32 subchannels (devices) per CHPID, versus the current 7 devices per CHPID. The additional device definition is available whether using an HCA3-O LR or HCA2-O LR feature for 1x IFB. This increase is designed to help eliminate the need to add CHPIDs or physical links in order to help improve coupling throughput at extended distance.

The HCA3-O feature (#0171) for 12x IFB and the HCA3-O LR feature (#0170) for 1x IFB are available starting with the z196 and z114, utilize CHPID type CIB, and are supported in the z/OS environment.

¹⁸ All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these statements of general direction is at the relying party's sole risk and will not create liability or obligation for IBM.

¹⁹ The InfiniBand link data rates do not represent the performance of the link. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

Question:

Were any default settings changed?

Answer:

Yes. In HCD definitions, the default number of devices (subchannels) for IFB links was changed from 32 back to 7. This is consistent with the maximum number of devices allowed in 12x IFB3 links. Customers with 1x IFB3 links can choose to enter 32 in the channel definition.

Question:

What connectivity capabilities were added to the zEC12 and zBC12 to support larger data sharing environments?

Answer:

As data sharing workloads continue to grow, the Parallel Sysplex infrastructure needs to anticipate the increased requirements for coupling resources. To do this, we have added connectivity capabilities to support larger data sharing environments.

Connectivity improvements with up to 64 1x IFB links: There is a maximum of up to 16 IFB fanout cards available in the zEC12 and zBC12, so each card can be used to support up to four 1x IFB links or two 12x IFB3 or IFB links. IBM now allows all 16 fanout cards to be configured for 1x IFB to support up to 64 links. This is an increase from a maximum of 48 1x IFB links possible with the IBM z196.

Greater than 16 ICFs per server: Although we are still keeping a maximum of 16 ICFs per LPAR, we now allow all the cores to be configured as ICFs across multiple LPARs. In the zEC12, one can configure up to 101 ICFs, requiring at least seven CF images. This allows a server to support multiple coupling facilities that may span multiple Parallel Sysplexes.

Question:

Do the zEC12 and zBC12 support ISC-3 coupling links?

Answer:

Yes, ISC-3 links are supported on the EC12 and zBC12. However, ISC-3 links will not be orderable on the zEC12 or zBC12, they can only be carried over from the z196, z114, or System z10. IBM plans on not supporting the ISC-3 link on the server after the zEC12 and zBC12.

Question:

What other generations of servers can connect to a zEC12 server in a Parallel Sysplex environment?

Answer:

The EC12 and zBC12 servers are designed to coexist in the same Parallel Sysplex environment with (n-2) server families. This allows it to coexist with other zEC12s, zBC12s, the z196 and z114 servers and the System z10 servers. Connectivity to the System z9 is not supported. You should also ensure that your z196, z114, and System z10 servers are at the latest driver level.

Question:

What servers will support CFCC Level 19?

Answer:

CFCC Level 19 comes with zEC12 (driver level 15) or zBC12. It is not rolled back to previous systems.

Question:

What System Management changes are added with the announcement?

Answer:

Information is now available to RMF and operator display commands to provide detailed performance and configuration information on the different Parallel Sysplex InfiniBand links, estimated distance, PCHID, HCA, and Port information, and running degraded or not. This is done for 12x or 1x links.

Question:

Are there any performance improvements?

Answer:

The coupling facility can now better manage the DB2 Group Buffer Pool (GBP) for DB2 batch jobs that do a lot of update activity. It can tell which pages are not part of the normal working set to request DB2 to write these directly to disk, helping the buffer hit ratio.

Other performance improvements help when making dynamic alters or changing element/entry ratios, as well as improving internal management of storage classes and castout classes.

Question:

Are there any RAS improvements?

Answer:

There are now more triggers for non-disruptive dumping of coupling control information, more trace points can be collected, and the ability to have larger trace tables, and specific trace tables set up for individual trace points. This helps ensure the required trace information is collected.

Note:

All new zEC12 enhancements for July 23, 2013, will require Driver 15, available on September 20, 2013.

Cloud and Mobile

Question:

Which improvements were done to support OpenStack?

Answer:

The z/VM Systems Management API (SMAPI) was enhanced to provide additional functionality required for xCAT and OpenStack support. Separately, IBM made a contribution to the OpenStack open source project that leverages xCAT and the enhanced SMAPI, allowing z/VM to be integrated into the OpenStack environment.

Question:

Are you able to provision mainframe CICS regions onto the cloud?

Answer:

As you know CICS v5.1 was made available in 4Q2012. In this version, CICS provides many "cloud-style" capabilities including those that can automate the provisioning of CICS regions. For more information go to: <http://www.ibm.com/software/products/us/en/cics-tserver-zos51/>

Question:

Have "cloud-based" or "cloud ready" systems been FedRAMP certified?

Answer:

Our understanding is that FedRAMP is a mandate for Federal government supported services datacenters. It is a security model that the providers need to comply with in order to work with the government. It involves end to end data center processes, policies, audibility, monitoring etc. We have found no indication that it applies directly to HW systems though there may be sourcing restrictions.

Question:

Which are the most important directions of IBM about Cloud Computing using Linux on System z?

Answer:

IBM's direction with Cloud computing is the use of OpenStack with our cloud offerings and software. With the announcement we made today, we announced that z/VM 6.3 has been enabled for OpenStack. This is in direct support for our Linux on System z environment.

Question:

What is System z's compelling value to support those new business imperatives like Mobile?

Answer:

These new business imperatives, in general, require that you have the necessary capabilities such as security, reliability, and scalability in place to be able to fully develop, deliver, and deploy these new business imperatives. For example, mobile requires that you have a complete secure and reliable infrastructure to connect critical and sensitive data to mobile devices. The zEnterprise system is the only system in the marketplace today that can provide the security required to connect this data and has 99.999% reliability.

GDPS

Question:

What can you tell me about the GDPS announcement on July 23, 2013?

Answer:

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

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

This is designed to extend GDPS/PPRC continuous availability and disaster recover capabilities to virtual servers in a zBX for planned and unplanned outages. The planned availability is fourth quarter 2013 with GDPS/PPRC v3.10 SPE.

Question:

What is the recommended disaster recovery (DR) capability for zEnterprise?

Answer:

GDPS is the recommended DR solution for zEnterprise clients. GDPS supports the process of aligning CEC configurations across the primary and DR sites, and coordinates storage replication as well as failover / failback for the zEnterprise. Specific to the zBX configuration, clients must ensure that the 2 sites have symmetric zBX configurations and must manually coordinate / replicate configuration changes from primary to DR site as necessary. The API support provided in GA2 may be used in conjunction with scripting to query the configurations and replicate them using automation. GDPS will manage and coordinate storage replication functions for both z/OS storage as well as distributed storage. For the zEnterprise, GDPS will leverage Distributed Cluster Management (DCM) to coordinate with TSA Application Manager (AppMan) in order to manage application components running on the zBX during a planned site switch (e.g., production site maintenance) or unplanned site switch (e.g., DR situation). Further information on GDPS can be found in IBM Redbooks®: SG24-6374-07 <http://publib-b.boulder.ibm.com/abstracts/sg246374.html?Open>

Question:

Where can I find an electronic version (including updates) of the GDPS frequently asked questions list?

Answer:

You can go to the GDPS Website at:
<http://www.ibm.com/systems/z/advantages/gdps/faqs.html>

z/OS

Question:

Where can I find an electronic version (including updates) of the z/OS and/ or z/OS Management Facility frequently asked questions?

Answer:

Please visit:

http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ST&infotype=SA&appname=STGE_ZS_ZS_USEN&htmlfid=ZSQ03064USEN&attachment=ZSQ03064USEN.PDF

Question:

Where can I find more information about z/OS?

Answer:

Visit the z/OS website: www.ibm.com/systems/z/os/zos/

Linux on System z

Question:

What are the unique advantages of running Linux on IBM System z?

Answer:

In short, Linux on System z can benefit you by freeing your business from IT complexity and improving the responsiveness of your systems and your people. The Linux environment helps you extract the maximum value from your IT budget through software savings, operational efficiency, power and space savings, as well as providing highest levels of security and virus-resistance while delivering legendary system reliability.

In detail, the Linux on System z provides unique advantages through the exploitation of the leading capabilities of the System z server hardware and the z/VM virtualization software:

Single-server simplicity for small to large-scale deployments

- Run from tens to hundreds of virtual Linux servers on a single physical server
- Saving opportunities in software costs through less licenses
- Saving opportunities in operational management costs
- Low power, cooling, and floor space requirements
- Very fast internal communication capabilities allow to reduce the physical network
- Minimize hardware needed for business continuance and disaster recovery
- Workload deployment on a “scale up” machine means fewer cables, fewer components, fewer operational effort.

Advanced resource utilization and dynamic allocation

- Efficiency at scale - high flexibility, scalability and resource utilization
- Share-everything system architecture
- Resources can be dynamically shared and reconfigured
- Over-commitment of system resources at high levels
- Achieve nearly 100% utilization of system resources nearly 100% of the time
- Create and deploy virtual Linux servers in minutes

Massive workload consolidation capability

- Server capabilities allow to start small and grow within the server, up to hundreds parallel workloads
- **Non-disruptive growth within one physical server**
- Proven horizontally and vertically scalability
- Large-scale virtual growth instead of physical expansion on x86 or RISC servers
- Run workloads with very large resource requirements on one physical server

Rock-solid system security, ensured isolation and reliability

- Take advantage of certified security
- Virus and intrusion resistance
- Virtual Linux images are totally isolated and protected
- Ultimate security with high-speed cryptography integrated as part of the chip
- Highest hardware security certification: Common Criteria Evaluation Assurance Level 5 (EAL5)
- Live Guest Relocation for non-disruptive maintenance
- Broad set of disaster recovery and continuous availability solutions

These advantages position the Enterprise Linux Server to run and manage many and mixed Linux workloads concurrently – a good fit for Linux workloads, including cloud computing.

You can “do more with less”.

Question:

What is Linux on System z?

Answer:

Linux on System z refers to the open source operating system environment running on the following IBM System z servers:

- IBM zEnterprise EC12 (zEC12) and IBM zEnterprise BC12 (zBC12)
- IBM zEnterprise 114 (z114) and IBM zEnterprise 196 (z196)
- IBM System z10 Enterprise Class (z10 EC) and IBM System z10 Business Class (z10 BC),
- IBM System z9 Enterprise Class (z9[®] EC) and IBM System z9 Business Class (z9 BC),
- IBM eServer zSeries 990, 890, 900, 800, (z990, z890, z900, z800) and S/390[®] (9672 G5, G6 and Multiprise[®] 3000).

Linux has been ported and is supporting these architectures. In addition, the ports are supporting some devices that are specific to these environments.

Please note that latest Linux on System z distributions may have dropped support of older System z and zSeries servers in their current releases (see next question).

Linux on System z exploits the strengths and reliability features of the outstanding z/VM virtualization and System z hardware technologies, while preserving the openness and stability of Linux.

Question:

What are the IBM tested and supported Linux distributions for Linux on System z?

Answer:

The latest information can be found on the tested platforms Web page (ibm.com/systems/z/os/linux/resources/testedplatforms.html).

Question:

Where can I get commercial Linux for System z distributions?

Answer:

Commercial Distributions are available from IBM Technical Support Services, or directly from our Linux Distribution Partners Red Hat (www.redhat.com) and SUSE (www.suse.com). Commercially available Linux distributions are obtained by purchasing the Red Hat and SUSE subscription service, and can include the Linux operating system enabled for specific hardware platforms, an assortment of device drivers, routines for installation, and value add programs like web servers and shells. The subscription service will also include access to product updates and fixes.

IBM suggests that these distributions be your first choice for production environments because of their availability of service and support from the Linux distributors or IBM Support Line. All System z Linux production environments should be covered by a Linux support agreement and maintained with a current Linux subscription service.

Question:

What IBM software products are available for Linux on System z?

Answer:

IBM offers a large software portfolio and new IBM software products are added constantly.

The latest information about IBM software available for the “Red Hat Enterprise Linux (RHEL)” versions and “SUSE Enterprise Linux Server (SLES)” versions for System z are available via the [IBM Software Product Availability tool](#)

(<http://publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/productsOnOs.html>).

Question:

Where can I get the latest and more information on Linux on IBM System z?

Answer:

The most current information about Linux on System z is available at: ibm.com/systems/z/linux

A separate FAQ document for download exists as well: [FAQ – Linux on System z](#)

Question:

What is the IBM Enterprise Linux Server?

Answer:

The IBM Enterprise Linux Server provides a virtualized infrastructure with the power and flexibility you need for your enterprise Linux applications. It is a proven platform for Linux workloads and consolidation, and its life-cycle management costs can be considerably less expensive than other Linux platforms.

The Enterprise Linux Server can provide significant advantages. It has the capability to run a large number of parallel workloads—up to hundreds—in a highly efficient and economical way. It can share system resources at extreme levels of utilization and provision flexible IT services. It provides load-balancing and efficient systems management helping to achieve superior levels of service and improved operational efficiency.

The IBM Enterprise Linux Server is designed from the ground up for enterprise-class workloads, can help on lower software license costs, can secure data from the growing threats and helps to use less of the environmental resources.

The Enterprise Linux Server provides an infrastructure that is data ready, mobile ready, security ready and cloud ready. It is a simple way to run an efficient and economic infrastructure.

Question:

What components constitute the Enterprise Linux Server?

Answer:

The Enterprise Linux Server includes IBM System z hardware, hardware maintenance, IBM virtualization software components and software support & subscription with a three to five year solution pricing.

- **Hardware options:**

You can choose from the IBM System z family of servers, the zEC12, zBC12, z196 or a z114 server. The Enterprise Linux Server includes one of the above System z servers with IFL processors, memory and I/O connectivity. Maintenance for three years is included for the server and all hardware components as well (one year warranty and two years pre-paid).

- **Virtualization software:**

z/VM Version 6 and the z/VM features: z/VM Directory Maintenance Facility, z/VM Resource Access Control Facility, Performance Toolkit for VM™, z/VM RSCS Feature. Three to five years Subscription and Support (S&S) for these software products is included as well.

Linux is not included. Linux can be ordered from the Linux distributors Red Hat or SUSE. Both Linux distribution partners provide specific offerings for the Enterprise Linux Server.

Question:

Where do I find more information about Enterprise Linux Server (ELS)?

Answer:

The most current information is available at: [ELS home page \(ibm.com/systems/z/os/linux/els.html\)](http://ibm.com/systems/z/os/linux/els.html).

A separate FAQ document for download exists as well: [ELS - FAQ](#)

Question:

What is the IBM SAN Volume Controller and IBM FlashSystem bundle and it's relationship to Linux?

Answer:

IBM has recently improved the IBM FlashSystem bundle, creating a new product offering, FlashSystem V840. The speed and MicroLatency of Flash, and the flexibility and software defined storage advanced capabilities have been combined into a single offering that scales in both performance (up to 2.5M IOPS) and capacity (from 2TB up to 320TB), while serving IO to popular servers and their operating system. Linux, including zLinux is listed in the SSIC (insert URL hot spot here). The FlashSystem V840 is an SSR installed process, so an IBM Customer Engineer performs the racking etc to get the FlashSystem V840 up and operational. Lab Based Services also provides a set of services to provide rapid time to value, so that you can realize your investment faster.

Question:

If I have an existing SVC can I just purchase a FlashSystem and use it with my Linux on System z

Answer:

The IBM FlashSystem 710, 720, 810 and 820 are all supported when used in conjunction with the IBM San Volume Controller at maintenance level 6.4.1.5 and above (including stream 7.1). For a full description of IBM's firmware support levels and a full listing of IBM storage interoperability the following URL 's should assist you in finding the best match for your environment.

Device support matrix can be found at: <http://www.ibm.com/support/docview.wss?uid=ssg1S1004392> and <http://www.ibm.com/support/docview.wss?uid=ssg1S1004366>

Question:

How do I order the IBM SVC and IBM FlashSystem bundle?

Answer:

Please see your IBM Storage sales team or IBM business partner to order the IBM FlashSystem products or FlashSystem bundle.

Services offerings are available from IBM Systems Lab Services and Training. To learn more about IBM FlashSystem services and other related products and services, contact stgls@us.ibm.com or visit: ibm.com/systems/services/labservices

Question

Where do I go for more information?

Answer:

ibm.com/storage/flash

z/VM

Question:

What did IBM announce on February 24, 2014 regarding z/VM?

Answer:

The following z/VM 6.3 enhancements were announced:

- CPU pooling support, designed to allow you to cap CPU utilization for a group of guests, to better balance resource usage
- Environment Information interface, designed to allow software to gather information about the environment in which it is operating
- Guest exploitation support for the zEDC Express feature (#0420) on the zEC12 and zBC12 servers, designed for compression acceleration for resource optimization
- Guest exploitation support for the 10GbE RoCE Express feature (#0411) on the zEC12 and zBC12 servers, designed to allow guests to utilize Remote Direct memory Access over Converged Ethernet (RoCE) for optimized networking

PTFs were available on June 27, 2014.

Question:

What are the benefits of zEDC and RoCE support for z/VM customers?

Answer:

Both zEDC and RoCE are based on new PCIe technology that provides high-speed low-latency connectivity. z/VM support enables these hardware features to be exploited by guests such as z/OS that have PCIe drivers.

Question:

Are any z/OS APARs required to exploit RoCE and zEDC as a guest of z/VM?

Answer:

Yes, the PTF for APARs OA44482 and OA43256 are required when exploiting zEDC as a z/VM guest. The PTF for APAR OA44482 is required when exploiting RoCE as a z/VM guest.

Question:

What did IBM announce on August 28, 2012 regarding z/VM?

Answer:

z/VM V5.4, z/VM V6.1, z/VM V6.2 and z/VM 6.3 provides compatibility support for the zEC12 and the zBC12. For more about z/VM and required compatibility support, see <http://www.vm.ibm.com/>

Question:

Today I'm using Unified Resource Manager to manage z/VM and virtual servers running on z/VM. When I upgrade to z/VM 6.2 can I continue to use the Unified Resource Manager?

Answer:

Yes, you can continue to use the Unified Resource Manager to manage z/VM and its virtual servers. If you plan to use the z/VM Single System Image feature (VMSSI) then the suggested best practice is that you should no longer use Unified Resource Manager to manage z/VM virtual servers as some VMSSI functions, such as Live Guest Relocation, will not be exploited.

Question:

So, if I am using VMSSI and am not managing my z/VM V6.2 virtual machines with Unified Resource Manager can I still use Unified Resource Manager to manage my POWER7 and System x blades?

Answer:

Yes.

Question:

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

Answer:

Yes, z/VM V6.2 and its virtual servers can be part of a zEnterprise ensemble and can be managed by the Unified Resource Manager. But, if you plan to use the VMSSI feature, especially if you plan to use virtual server mobility, then the best practice is that you should no longer use Unified Resource Manager to manage z/VM virtual servers.

Question:

I plan to use Unified Resource Manager for management of my z/VM systems and I do not plan to use the VMSSI feature of z/VM Version 6 Release 2. Are there any differences in configuring z/VM to be ensemble managed between z/VM Version 6 Release 1 and z/VM Version 6 Release 2?

Answer:

Yes, the process is somewhat different between the two releases. The z/VM Version 6 Release 1 support for Unified Resource Manager delivered as service, while the z/VM 6.2 support is in the base of the product with some configuration already defined. As a result, most people find it easier to configure z/VM to be used in an ensemble environment on z/VM 6.2.

Question:

I am interested in deploying z/VM V6.2 VMSSI. Are their resources to help me understand, plan and deploy a z/VM Single System Image?

Answer:

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

Question:

Where can I find a full set of z/VM frequently asked questions?

Answer:

Please visit: <http://www.vm.ibm.com/faq>

Question:

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

Answer:

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 System x hypervisor based on Kernel-based Virtual Machine (KVM).

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

z/VSE

Question:

I am a z/VSE user. What's in the zBC12 announcement for me?

Answer:

For all z/VSE users:

- Latest hardware technology
- Higher single engine (PU) performance for single threaded workloads, benefitting batch times of specific jobs, or other non-parallel workloads
- Lower hardware maintenance cost
- Substantially lower cost per MSU for small (less than 10 MSU) systems - largest savings for 3 MSU clients.
- All Linux on System z benefits if you run Linux together with z/VSE
- Growth capacity - processors and memory (important if you are using Linux on System z)

For z/VSE users on zBC10 and older:

- Software (pricing) savings (AEWLC advantage)
 - Significant software license savings if you are still on zSeries (z890, z800, z900, z990) or older hardware
-

Question:

Does z/VSE 5.1 provide VLAN support?

Answer:

Yes. z/VSE 5.1 provides (Global) VLAN support for OSA Express (CHPID type OSD and OSX) and HiperSockets (CHPID IQD) devices. Global VLAN support can be used with in a Layer 3 configuration. z/VSE 5.1 provides VLAN support for OSA Express (CHPID type OSD and OSX) devices in a Layer 2 configuration that carries IPv6 traffic. VLAN in a Layer 2 configuration requires the IPv6/VSE product licensed from Barnard Software, Inc.

VLAN support is required when z/VSE participates in an IEDN using OSA Express for zBX (CHPID OSX) devices.

Question:

Does z/VSE 5.1 provide Layer 2 support?

Answer:

Yes. z/VSE 5.1 adds Layer 2 support for OSA Express devices (CHPID type OSD or OSX) when used in an IPv6 link. Layer 2 requires the IPv6/VSE product licensed from Barnard Software, Inc. With Layer 2 support you can include a z/VM VSWITCH in your z/VSE IPv6 configuration and, when accessing the IEDN through a z/VM VSWITCH, you can use IPv6 links in addition to IPv4 links.

Question:

What did IBM announce on October 12, 2011 and April 2, 2013 with regard to z/VSE?

Answer:

IBM announced in 2011 z/VSE Version 5 Release 1. This new release once again demonstrates IBM's commitment to z/VSE clients. This new version is another significant step in the ongoing evolution of z/VSE, by offering 64-bit virtual addressing for workload growth and new applications as well as exploitation of innovative zEnterprise System technology.

For more information please see the Announcement letter

at: http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS211-118&open&cm_mmc=5733- -n- -vrm_newsletter- -10207_134134&cmibm_em=dm:0:12046832

In 2013 IBM announced that z/VSE allows 64-bit Input/Output (I/O) processing for applications.

For more information please see the Announcement letter

at: <http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS213-090>

Question:

Where can I get the latest and more information on z/VSE?

Answer:

The most current information about z/VSE is available at: ibm.com/zvse

Statements of Direction

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these statements of general direction is at the relying party's sole risk and will not create liability or obligation for IBM.

From the July 23, 2013 announcement letter (Note – since July 23, 2013, when new function is delivered, the SOD on this:

WebSphere DataPower: IBM WebSphere DataPower Integration Appliance XI52 Virtual Edition for use within the IBM zEnterprise BladeCenter Extension provides flexible deployment options: IBM intends to introduce a version of the WebSphere DataPower Integration Appliance XI52 Virtual Edition for use within the IBM zEnterprise BladeCenter Extension (zBX). IBM intends for this virtual appliance to run on System x blades installed within a zBX and is designed to provide industry-leading security, integration, and optimization capabilities for System z similar to the physical WebSphere DataPower appliance models.

zEDC Express exploitation by z/OS for DFSMSdss and DFSMSHsm: IBM intends to provide exploitation of the zEDC Express feature for DFSMSdss and DFSMSHsm by the end of third quarter 2014.

Removal of support for 3-in-1 Bolt Down Kits: The IBM zEnterprise EC12 and IBM zEnterprise BC12 are planned to be the last System z servers to offer ordering of the 3-in-1 Bolt Down Kits for raised-floor and non-raised-floor environments. This applies to the zEC12 features #8000, #8001, #8002 and to the zBC12 features #8016 and #8017. Alternate solutions are available.

Removal of support for the HCA2-O fanouts for 12x IFB and 1x IFB InfiniBand coupling links: The IBM zEnterprise EC12 and the IBM zEnterprise BC12 are planned to be the last System z servers to support the following features as carry forward on an upgrade: HCA2-O fanout for 12x IFB coupling links (#0163) and HCA2-O LR fanout for 1x IFB coupling links (#0168). Enterprises should continue migrating to the HCA3-O fanout for 12x IFB (#0171) and the HCA3-O LR fanout for 1x IFB (#0170).

Removal of support for connections to an STP Mixed CTN: The IBM zEnterprise EC12 and the IBM zEnterprise BC12 are the last System z servers to support connections to an STP Mixed CTN. This includes the Sysplex Timer (9037). After the zEC12 and the zEnterprise BC12, servers that require Time synchronization, such as to support a base or Parallel Sysplex, will Require Server Time Protocol (STP), and all servers in that network must be configured in STP-only mode.

Removal of ISC-3 support on System z: The IBM zEnterprise EC12 and the IBM zEnterprise BC12 are planned to be the last System z servers to offer support of the InterSystem Channel-3 (ISC-3) for Parallel Sysplex environments at extended distances. ISC-3 will not be supported on future System z servers as carry forward on an upgrade. Previously we announced that the IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114) servers were the last to offer ordering of ISC-3. Enterprises should continue migrating from ISC-3 features (#0217, #0218, #0219) to 12x InfiniBand (#0171 - HCA3-O fanout) or 1x InfiniBand (#0170 - HCA3-O LR fanout) coupling links.

Removal of OSA-Express3 support on System z: The IBM zEnterprise EC12 and the IBM zEnterprise BC12 are planned to be the last System z servers to offer support of the Open System Adapter-Express3 (OSA-Express3 #3362, #3363, #3367, #3370, #3371) family of features. OSA-Express3 will not be supported on future System z servers as carry forward on an upgrade. Enterprises should continue migrating from the OSA-Express3 features to the OSA-Express4S (#0404, #0405, #0406, #0407, #0408) and OSA-Express5S (#0413, #0414, #0415, #0416, #0417) features.

Updated 08/09/13 - Removal of support for IEEE 802.3 Ethernet frame types: The IBM zEnterprise EC12 and IBM zEnterprise BC12 are planned to be the last System z servers to support IEEE 802.3 Ethernet frame types on OSA-Express QDIO interfaces in Layer 3 mode. This statement applies to CHPID types OSD and OSX when they are used in Layer 3 mode. These OSA-Express CHPID types in

Layer 3 mode are planned to support Ethernet DIX Version 2 (DIX V2) exclusively on future System z servers.

OSA-Express non-QDIO (CHPID type OSE) supporting SNA/APPN/HPR with Link Station Architecture (LSA), TCP/IP passthru environments with LAN Channel Station (LCS), and QDIO CHPID types OSD and OSX running in Layer 2 mode are not affected by this change.

Removal of FICON Express4 support on System z: The IBM zEnterprise EC12 and the IBM zEnterprise BC12 are planned to be the last System z servers to offer support of the FICON Express4 features (#3321, #3322). FICON Express4 will not be supported on future System z servers as carry forward on an upgrade. Enterprises should continue migrating from the FICON Express4 features to the FICON Express8S features (#0409, #0410).

Removal of Crypto Express3 support on System z: The IBM zEnterprise EC12 and the IBM zEnterprise BC12 are planned to be the last System z servers to offer support of the Crypto Express3 features (#0864 and #0871 - zEnterprise BC12). Crypto Express3 will not be supported on future System z servers as carry forward on an upgrade. Enterprises should continue migrating from the Crypto Express3 features to the Crypto Express4S feature (#0865).

IBM System z Integrated Information Processor (zIIP) and IBM System z Application Assist Processor (zAAP) simplification: The IBM zEnterprise EC12 and the IBM zEnterprise BC12 are planned to be the last System z servers to offer support for zAAP specialty engine processors. IBM intends to continue support for running zAAP workloads on zIIP processors ("zAAP on zIIP"). This is intended to help simplify capacity planning and performance management, while still supporting all the currently eligible workloads.

Stabilization of z/VM V5.4 support: The IBM zEnterprise EC12 and IBM zEnterprise BC12 are planned to be the last System z servers supported by z/VM V5.4 and the last 1 System z servers that will support z/VM V5.4 running as a guest (second level). z/VM V5.4 will continue to be supported until December 31, 2014, or until the IBM System z9 EC and IBM System z9 BC are withdrawn from support, whichever is later. Refer to Announcement 912-144, (RFA56762), Software withdrawal/discontinuance of service: IBM System z selected products, dated August 7, 2012.

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.



©Copyright IBM Corporation 2014
IBM Systems and Technology Group
Route 100
Somers, New York 10589
U.S.A.

Produced in the United States of America,
07/2014

IBM, IBM eServer, IBM logo, AIX, BladeCenter, Blue Gene/Q, CICS, DataPower, DB2, DFSMS, DFSMDdss, DFSMShsm, DS8000, Easy Tier, ESCON, FICON, GDPS, HyperSockets, HyperSwap, IMS, InfiniBand, MQSeries, Multiprise, MVS, Parallel Sysplex, PartnerWorld, Passport Advantage, Performance Toolkit for VM, Power, POWER, POWER6, POWER6+, POWER7, PowerHA, PowerVM, Predictive Failure Analysis, PR/SM, pSeries, PureApplication, PureData, PureFlex, PureSystems, RACF, Redbooks, Resource Link, RMF, S/390, ServerProven, Storwize, Sysplex Timer, System p, System Storage, System x, System z, System z9, System z10, System z10 Business Class, Tivoli, Veritas, WebSphere, x-Architecture, XIV, z9, z10, z10 BC, z10 EC, z/Architecture, zEnterprise, z/OS, zSeries, z/VM and z/VSE are trademarks or registered trademarks of the International Business Machines Corporation.

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

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

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

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

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

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

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

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

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

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

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

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

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

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