IBM z Systems
Introduction
May 2017

IBM z13s and IBM z13

Frequently Asked Questions

Worldwide

ZSQ03076-USEN-15

Table of Contents

z13s Hardware

What is the machine type of the IBM z13s?

The z13s machine type is 2965.

What z13 models were announced?

The z13s is available in two models – the N10 and N20, the model naming indicating how many total processor units are available for client purchase and characterization.

The N10 can be configured with up to 10 configurable cores and up to 1 TB of memory. The N20 can be configured with up to 20 configurable cores and up to 4 TB of memory.

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

We also offer both IFL only and ICF only configurations where you can have up to ten or twenty IFLs or ICFs in a standalone environment.

How much capacity will the z13s deliver?

The z13s offers a maximum of 6 CPs with the range up to 7123 millions of instructions per second (MIPS)[1] in a single footprint or 1.4x the capacity of the zBC12. For Linux® the z13s contains up to 20 IFLs for Linux specialty engines – 2.1x that of the zBC12.

Where can I find the MIPS ratings for z13s?

IBM z Systems provides capacity comparisons among processors based on a variety of measured workloads which are published in the Large System Performance Reference (LSPR). We don't provide MIPS but we do provide Processor Capacity Index (PCI) values. They can be found at: https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex

[1] Based on preliminary internal measurements and projections. Official performance data will be available upon announce and can be obtained online at LSPR (Large Systems Performance Reference) website at: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument . Actual performance results may vary by customer based on individual workload, configuration and software levels.

Did IBM increase the number of logical partitions (LPARs) available on the z13s?

Yes. There are now up to 40 logical partitions (LPARs) available for use versus 30 on the IBM zEnterprise® BC12 (zBC12). This will allow clients to support more 'isolated/separate' workloads on a single footprint.

How many HiperSockets™ will be available on the z13s?

There will still be 30 HiperSockets available. Please review the new Shared Memory Communications – Direct Memory Access (SMC-D) in the I/O section of the FAQs for information about a new inter-system communications protocol and their relationship to HiperSockets.

Why are there are door locks on the new z13s?

Yes. The front and rear doors will arrive with the door lock installed. There have been client requests for security rules for servers, for cryptographic features and for external auditors. The decision was made to have them on all servers. The use of the door locks is up to the client choice.

Each door will come with two keys. If the doors are locked, then the client is responsible for unlocking the doors for IBM to perform service. IBM support personnel will not have 'master' keys.

How many spare processing cores are on the z13s?

There are two 'dedicated' IBM provided spares on the model N20 only. Unused or unassigned engines on either the N10 or N20 can be used as spares.

Can any of the spare processing cores on the z13s be used for other purposes?

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

What is a CPC drawer on the z13s?

The Central Processor Complex Drawer contains the physical collection of hardware that consists of main storage, one or more central processors, timers, and channels – installed in one unit called the CPC drawer. Earlier z Systems had "books" versus "drawers" to house the same collection of hardware.

There can up to two CPC drawers installed on the z13s. The N10 has one and the N20 will have up to two CPC drawers depending on the I/O and memory configuration.

What is new about the processor chip used in the z13s?

The z13s and z13 use the same processor chip technology.

With changes in the chip technology industry, microprocessor frequency is no longer the primary means to achieve incremental performance. Hence, the 22nm silicon-on-insulator 8-core processor chip achieves performance gains over zBC12 by the following micro-architecture innovations:

- Increased instruction parallelism through wider instruction decode bandwidth, increased execution bandwidth and a more aggressive out-of-order-execution.

- Economies of scale, using simultaneous multi-threading (SMT) to execute two instruction streams (or threads) on a processor core which delivers more throughput for Linux on z Systems and IBM z Integrated Information Processor (zIIP) eligible workloads.
- Single Instruction Multiple Data (SIMD), a vector processing model providing instruction level parallelism, to speed workloads such as analytics and mathematical modeling. For example, IBM Java®, XL C/C++ on z/OS, XL C/C++ib Linux on z, COBOL 5.2 and PL/I 4.5 exploit SIMD and enhanced floating point to deliver improved performance over and above that provided by the faster processor.
- On chip cryptographic and compression coprocessors receive a performance boost improving both general processors and Integrated Facility for Linux (IFL) cryptographic performance and allowing compression of more data, helping to save disk space and reducing data transfer time.
- Redesigned cache architecture, leveraging our leadership eDRAM technology to provide twice as much second, third and fourth level caches compared to the zBC12. Bigger and faster caches help to avoid untimely swaps and memory waits while maximizing the throughput of concurrent workloads.

How does SIMD on the z13s differ from the vector processing of past servers?

This is not the engineering/scientific vector processing used for weather, climate, etc. This vector processing can be used for analytics operations. It provides much needed extensions to aid in the performance of complex database operations and provides the ability for a single instruction stream to process multiple data streams simultaneously rather than one at a time.

With the z13s announcement, KVM has announced capability to support new analytics workloads with Single Instruction Multiple Data (SIMD) – available 1Q16. KVM will also support SIMD on the z13.

How is simultaneous multithreading implemented on the z13s?

We have implemented simultaneous multithreading on the z13s the same as we did on the z13. The incremental throughput is achieved partly because the new processor chip offers intelligently implemented 2-way simultaneous multithreading. Simultaneous multithreading (SMT) allows two active instruction threads per core, each dynamically sharing the core's execution resources. SMT will be available in z13s for workloads running on the IFL and the zIIP.

Each operating system / hypervisor has the ability to intelligently drive SMT in a way that is best for its unique requirements. In z/OS®, SMT management consistently drives the cores to high thread density, in an effort to reduce SMT variability and deliver repeatable performance across varying CPU utilization – thus providing more predictable SMT capacity. But remember that z/OS is only supporting SMT for workloads running on the zIIP. In z/VM®, SMT management optimizes throughput by spreading a workload over the available IFL cores until it demands the additional SMT capacity (z/VM only supports SMT for IFL cores).

With the z13s announcement, KVM has announced capability to deliver higher compute capacity with support for simultaneous multithreading (SMT) – available 1Q16. KVM will also support SMT on the z13.

What is the HSA and how much is there on the z13s?

The hardware system area (HSA) HSA is a non-addressable storage area that contains system microcode and configuration dependent control blocks. The HSA has a fixed size and is not part of the purchased memory that you order and install. The HSA has sufficient reserved space to allow for dynamic I/O reconfiguration changes to the maximum capability of the processor. On the z13s there is 40 GB of HSA, independent of client-purchased memory.

How much memory is available on the z13s?

The z13s server supports up to 4 TB of real memory per server (LPAR limits are dependent on the operating system), but the actual maximum physical memory sizes are related to the number of CPC drawers in the system. The minimum initial amount of memory that can be ordered is 64 GB for all models. The maximum memory on the N10 is 1 TB (minus 40 GB HSA). The maximum memory on the N20 with 1 CPC drawer is 2 TB (minus 40 GB HSA). It is required to have the 2 CPC drawer version of the z13s N20 to enable 4 TB (minus 40 GB HSA) of memory.

Please note that the 4 TB maximum memory is client-usable, RAIM-protected memory. RAIM is always active, and IBM has already factored in an additional 20% of RAIM memory supplying redundancy.

The hardware system area (HSA) HSA is a non-addressable storage area that contains system microcode and configuration dependent control blocks. The HSA has a fixed size and is not part of the purchased memory that you order and install. The HSA has sufficient reserved space to allow for dynamic I/O reconfiguration changes to the maximum capability of the processor. On the z13s there is 40 GB of HSA, independent of client-purchased memory.

What is the RAIM feature of the z13s?

The z13 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.

What is the value of 4 TB of memory on the z13s?

The z13s offers up to 4 TB of memory—8X more than the zBC12. This will advantage many types of users. Large memory can reduce latency and CPU cost, and thus improving operational efficiency, for WebSphere® Application Server and Java applications by allowing larger heaps without an increase in paging. Large memory for IBM MQ® V8 can help to cost effectively manage the increasing message volumes generated from today's mobile and cloud applications. Configuring more memory for use by IBM Cognos® Dynamic Cubes on z13s improves the affordability of keeping the entire application life cycle on z Systems. z/VM 6.3 and later supports 1 TB of real memory to help clients keep pace with increasing business demands and thus Linux application servers, database servers, analytic and cloud workloads running native or under z/VM, may see performance benefits when taking advantage of large shared, virtualized memory.

What is the integrated firmware processor (IFP) on a z13s?

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 client usable or 'visible'.

How many System Assist Processors (SAPs) are on the z13s?

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

The z13s Model N10 has two.

The z13s Model N20 has three.

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

What is the LPAR absolute hardware capacity setting on a z13s?

Users, particularly our Linux clients 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, introduced with the zBC12, allows 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.

What is Absolute Capping of an LPAR group on the z13s?

LPAR group capping is similar to the hardware "absolute" capping of individual partitions. Processor Resource/Systems Manager™ (PR/SM™) and the Hardware Management tool have been enhanced to support an option to limit the amount of physical processor capacity consumed by a group of logical partitions (LPARs).

Group capping allows customer to set a cap on the physical usage of processors across a group of partitions. Using the HMC, the customer defines the group of partitions such that they are not able to consume more than the defined number of processors-worth (CP or IFL) of capacity across the group. For example, LPARs 3, 4 and 7 form a group that is capped to not use more than seven Linux cores.

Within a defined group, individual LPARs can also have their own independent Absolute Cap.

This is done without z/OS workload management (WLM) so it is very useful for non-z/OS environments where subcapacity pricing support (group soft capping) is not applicable.

What is LPAR dynamic memory management on the z13s?

PR/SM has been enhanced to support more flexibility as to how additional physical memory is dynamically added to a logical partition.

Rather than attempting to fully populate a logical partition's reserved storage element when it is initially configured online, the operating system in the partition can request a single storage increment be attached (and subsequently can request additional increments if desired). This allows

a more gradual, flexible addition of memory to the partition as needed over time. For example, the Linux image in LPAR3 could request more memory, and PR/SM dynamic memory management can dynamically give incremental memory to LPAR3 without human intervention.

This is done without z/OS workload management (WLM) so it is very useful for non-z/OS environments where subcapacity pricing support (group soft capping) is not applicable.

What cooling options are available for the z13s?

The z13s is a single frame, air cooled system only - water cooling is not an option for the z13s.

What is ASHRAE?

Data centers continue to use higher amounts of energy with each generation of IT equipment[2]. The American Society of Heating, Refrigeration, and Air Conditioning Engineers (ASHRAE) is an organization that (among other things) created a set of environmental guidelines to help data center operators simplify the process of selecting IT equipment for their data center environment. Data center operators should use these classifications of IT equipment to determine their energy envelope, which can help them reduce overall energy use by improving energy efficiency.

What is the ASHRAE rating of the z13s?

The z13s is rated for the ASHRAE class A3 environment for all models and configurations. Class A3 allows a wider operating temperature range (including up to 40 degrees C inlet air), as well as a wider humidity range than the previous generation's class A2 rating. This provides significant additional system robustness and safety margin for data center cooling and humidity control failures. It allows installation in a wider range of data centers, including lower cost data centers. Energy savings are enabled in multiple ways. The increased temperature and humidity range provides the opportunity to use compressor-less cooling solutions in many locales. Additionally, the expanded temperature range allows for a data center to run at a higher temperature for some time periods with high external temperatures, reducing the required cooling infrastructure and energy usage.

For additional information please refer to
http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=WH&infotype=SA&htmlfid=XBW03029USEN&attachment=XBW03029USEN.PDF

Can I buy a z13s that has only ICF processors without including a general-purpose processor (CP)?

Yes. You can order only ICFs in a z13s, using a model capacity identifier of A00 with a maximum of 10 ICFs on the N10 and 20 ICFs on the N20. There is still a limit of 16 ICF engines for a single Coupling Facility LPAR.

---

[2] ASHRAE, Datacom Equipment Power Trends and Cooling Applications, 2nd Edition, ASHRAE Datacom Series 2

Can I buy a z13s that is made up entirely of Integrated Facility for Linux (IFL) engines and no general-purpose processor (CP)?

Yes. You can order an IFL-only z13s, using a capacity identifier of A00 with 1 to 10 IFLs on the N10 and 1 to 20 IFLs on the N20. For more information, ask your IBM or Business Partner representative for more information.

Is there still an option for overhead cabling on the z13s?

Yes, you can order overhead cabling and you can also order overhead power on the z13s.

What non-raised floor option is available on the z13s?

The non-raised floor option allows a z13s to be ordered and installed without needing a raised floor. This feature will be a great option for clients looking for 'cement' floors for disaster recovery centers, backup sites, or other lower- cost operations facilities. With a non-raised floor z13s, you must order overhead power and overhead cabling.

Why might I be interested in the new rack mounted HMC or new rack mounted Trusted Key Entry (TKE) console on the z13s?

The new optional rack mounted management console can help to save space when it is a premium in data centers or help meet data center best practices.

Will my order for the rack mounted HMC of the z13s or rack mounted TKE include the rack?

No. Most of our clients already have 19" racks with available space in them, so we will not provide a rack. If you need one you'll need to work with your sales rep to purchase one.

Is there still going to be a zBX available for the z13s?

Yes. Clients who have embarked on our hybrid computing strategy will see investment protection for their zBX. The z13s will support the IBM z BladeCenter® Extension (zBX) Model 004. The zBX Model 004 is only available as an upgrade from either a zBX Model 003 or zBX Model 002. The zBX Model 004 is designed for clients that desire to continue managing the integrated hybrid environment with IBM z Unified Resource Manager.

What availability enhancements were made to the z13s?

The z13s continues the drive for continuous reliable operation provided by its predecessors with the following Reliability, Availability and Serviceability (RAS) improvements:

Improved soft error resilience in the processor cores

Lane shadowing, hardware buffer retry, and independent channel recovery which are designed to improve the DIMM interface

Continued use of RAIM in the main memory to protect DRAM

Improved robustness in the level 3 and level 4 cache

Improved isolation with the addition of integrated time domain reflectometry logic to chip interfaces

Enhanced integrated sparing designed to reduce the complexity and number of repair actions

I am interested in making sure my IT personnel are well trained in all aspects of IBM z Systems. Are there resources to help me with client education?

Yes, IBM STG Lab Services and Training and our Global Training Providers can assist with this. You can contact IBM STG Lab Based Services via the Internet at: http://www.ibm.com/systems/services/labservices/ or send an email to stgls@us.ibm.com

z13 Hardware

What were the new enhancements to the z13 with the February 16, 2016 announcement?

Optimize internal communications between operating systems transparently with no changes to applications, using Shared Memory Communications - Direct Memory Access (SMC-D). *See SMC-D section for more details.*

Simplify z Systems hardware virtual infrastructure management for KVM for IBM z Systems including integrated dynamic I/O management with IBM Dynamic Partition Manager (DPM). *See IBM Dynamic Partition Manager section for more details.*

Improved Time-To-Value with faster deployment and implementation of software solutions delivered as virtual software appliances and firmware appliances with Secure Service Container. *See Secure Service Container section for more details.*

LPAR enhancement to provide group physical capacity limit enforcement - *See HMC section for more details*

LPAR enhancement for dynamic memory management

Trusted Key Entry (TKE) 8.1 License Internal Code (LIC) – *See Security section for more details*

Support to help facilitate SAN configuration setting and debugging – *See HMC section for more details*

Protect sensitive data by using Transport Layer Security support in the OSA-Integrated Console Controller (OSA-ICC) - *See HMC section for more details*

Export / import physical port WWPNs for FCP Channels – *See Networking section for more details*

Fibre Channel Read Diagnostic Parameter extended link service (ELS) support – *See Networking section for more details*

z/VSE® Network Appliance for z/VSE systems running in LPAR - *See z/VSE Network Appliance section for more details.*

Hardware Management Console (HMC) enhancements - *See HMC section for more details.*

What is LPAR dynamic memory management on z13?

PR/SM has been enhanced to support more flexibility as to how additional physical memory is dynamically added to a logical partition.

Rather than attempting to fully populate a logical partition's reserved storage element when it is initially configured online, the operating system in the partition can request a single storage increment be attached (and subsequently can request additional increments if desired). This allows a more gradual, flexible addition of memory to the partition as needed over time. For example, the Linux image in LPAR3 could request more memory, and PR/SM dynamic memory management can dynamically give incremental memory to LPAR3 without human intervention.

The following z13 FAQs are from the January 2015 FAQ document.

What is the machine type of the z13?

The z13 machine type is 2964.

How much capacity with the z13 deliver?

Powered by 141 of the world's most powerful microprocessors capable of z/OS execution of more than 111,000 millions of instructions per second (MIPS)[3], z13 delivers up to 40% more total general system processing capacity than the IBM zEnterprise EC12 (zEC12).

Does the z13 offer more available subcapacity processors than the zEC12?

Yes. The z13 can have up to 30 subcapacity Central Processors (CPs) per server while the zEC12 accommodates up to 20 subcapacity CPs. These subcapacity processors can physically reside in multiple CPC drawers. 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.

What is a CPC drawer on the z13?

The packaging on the z13 will transition from a book with an 8 chip Glass Ceramic Multi-chip Module (MCM) to the Single Chip Modules (SCM). Each CPC drawer will support 2 logical nodes. Each logical node contains 3 processor unit (PU) chips and one system controller (SC) chip.

The SCM packaging is used today on the zBC12. The design change allows the z13 to offer improved scalability, reliability and simplicity over the book design.

What z13 models were announced?

IBM announced the following five z13 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 z13 Model N30 model can be a 1-way through 30-way – which means there are 30 orderable cores contained on one CPC drawer (plus many supporting cores in every model, including a minimum of two spare cores).

A z13 Model N63 model can be a 1-way through 63-way (63 orderable cores) contained in two CPC drawers.

A z13 Model N96 model can be a 1-way through 96-way (96 orderable cores) contained in three CPC drawers.

A z13 Model NC9 model can be a 1-way through 129-way (129 orderable cores) contained in four CPC drawers.

---

[3] Based on preliminary internal measurements and projections. Official performance data will be available upon announce and can be obtained online at LSPR (Large Systems Performance Reference) website at: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument . Actual performance results may vary by customer based on individual workload, configuration and software levels.

The enhanced capacity z13 Model NE1 model can be a 1-way through 141-way (141 orderable cores) contained in four CPC drawers.

Customers that reach CPC drawer limits can easily upgrade from smaller models to larger models (except the NE1) 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), z Integrated Information Processors (zIIPs), additional System Assist Processors (SAPs), Internal Coupling Facilities (ICFs) and/or used as additional spares.

Up to thirty subcapacity CPs can be active on the server regardless of hardware model. In other words, subcapacity configurations are available on any of the models (N30 to NE1) as long as the server is configured (not necessarily the same as purchased) with fifteen or fewer general-purpose processors.

What can you tell me about the processor chip used in the z13?

- The z13 is available with up to 141 configurable processor units for performance and scaling advantages over prior generations of the mainframe, supporting up to 8,000 virtual servers in a single footprint.
- With changes in the chip technology industry, microprocessor frequency is no longer the primary means to achieve performance. Hence, the new 22nm silicon-on-insulator 8-core processor chip achieves performance gains over zEC12 by the following micro-architecture innovations:
- Increased instruction parallelism through wider instruction decode bandwidth, increased execution bandwidth and a more aggressive out-of-order-execution.
- Economies of scale, using simultaneous multi-threading (SMT) to execute two instruction streams (or threads) on a processor core which delivers more throughput for Linux on z Systems and IBM z Integrated Information Processor (zIIP) eligible workloads.
- Single Instruction Multiple Data (SIMD), a vector processing model providing instruction level parallelism, to speed workloads such as analytics and mathematical modeling. For example, IBM Java, COBOL 5.2 and PL/I 4.5 exploit SIMD and enhanced floating point to deliver improved performance over and above that provided by the faster processor.
- On chip cryptographic and compression coprocessors receive a performance boost improving both general processors and Integrated Facility for Linux (IFL) cryptographic performance and allowing compression of more data, helping to save disk space and reducing data transfer time.
- Redesigned cache architecture, leveraging our leadership eDRAM technology to provide twice as much second level cache and substantially more third and fourth level caches compared to the zEC12. Bigger and faster caches help to avoid untimely swaps and memory waits while maximizing the throughput of concurrent workloads.

How is simultaneous multithreading implemented on the z13?

Incremental throughput is achieved partly because the new processor chip offers intelligently implemented 2-way simultaneous multithreading. Simultaneous multithreading (SMT) allows two active instruction threads per core, each dynamically sharing the core's execution resources. SMT will be available in z13 for workloads running on the IFL and the zIIP.

Each operating system / hypervisor has the ability to intelligently drive SMT in a way that is best for its unique requirements. In z/OS, SMT management consistently drives the cores to high thread density, in an effort to reduce SMT variability and deliver repeatable performance across varying CPU utilization – thus providing more predictable SMT capacity. But remember that z/OS is only supporting SMT for workloads running on the zIIP. In z/VM, SMT management optimizes throughput by spreading a workload over the available IFL cores until it demands the additional SMT capacity (z/VM only supports SMT for IFL cores).

With the z13s announcement, KVM has announced capability to deliver higher compute capacity with support for Simultaneous Multithreading (SMT) on both the z13s and the z13 – available 1Q16.

How does SIMD on the z13 differ from the vector processing of past servers?

This is not the engineering/scientific vector processing used for weather, climate, etc. This vector processing can be used for analytics operations. It provides much needed extensions to aid in the performance of complex database operations and provides the ability for a single instruction stream to process multiple data streams simultaneously rather than one at a time.

With the z13s announcement, KVM has announced capability to support new analytics workloads with Single Instruction Multiple Data (SIMD) on both the z13s and the z13 – available 1Q16.

What is the HSA and how much is there on the z13?

The hardware system area (HSA) HSA is a non-addressable storage area that contains system microcode and configuration dependent control blocks. The HSA has a fixed size and is not part of the purchased memory that you order and install. The HSA has sufficient reserved space to allow for dynamic I/O reconfiguration changes to the maximum capability of the processor. On the z13 there is 96 GB of HSA, independent of client-purchased memory.

How much memory is available on the z13?

The z13 server supports up to 10 TB of real memory per server (LPAR limits are dependent on the operating system), but the actual maximum physical memory sizes are related to the number of CPC drawers in the system. The minimum initial amount of memory that can be ordered is 64 GB for all models.

What is the RAIM feature of the z13?

The z13 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 10 TB maximum memory is customer-usable, RAIM-protected memory. RAIM is always active, and IBM has already factored in an additional 20% of RAIM memory supplying redundancy.

What is the value of 10 TB memory on the z13?

The z13 offers up to 10 TB of memory—3X more than the zEC12. This will advantage many types of users. Large memory can reduce latency and CPU cost, and thus improving operational efficiency, for WebSphere Application Server and Java applications by allowing larger heaps without an

increase in paging. Large memory for IBM MQ V8 can help to cost effectively manage the increasing message volumes generated from today's mobile and cloud applications. Configuring more memory for use by IBM Cognos Dynamic Cubes on z13 improves the affordability of keeping the entire application life cycle on z Systems. z/VM 6.3 or later supports 1 TB of real memory to help clients keep pace with increasing business demands and thus Linux application servers, database servers, analytic and cloud workloads running native or under z/VM, may see performance benefits when taking advantage of large shared, virtualized memory.

What is different about the z13 Model NE1 compared to zEC12 Model HA1?

The Model NE1 is an enhanced capacity model which contains a different configuration of SCMs than other models. The z13 Model NE1 is fully populated with four high density CPC drawers and 141 orderable cores. You can configure the NE1 machine to be a 1 to 141-way. Like the other four-CPC drawer model, the NC9, the NE1 can be ordered with a minimum of 64 GB of memory and up to a maximum of 10 TB.

Upgrading from any other model of the z13 to a Model NE1 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.

What does the 'E' in Model NE1 of the z13 stand for?

The 'E' in NE1 represents the number '14'. So there are 141 configurable cores in the Model NE1.

What does the 'C' in Model NC9 of the z13 stand for?

Note: The 'C' in NC9 represents the number '12' so there are 129 configurable cores in the Model NC9.

What is the integrated firmware processor (IFP) on a z13?

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'.

How many spare processing cores are on the z13?

IBM ships every z13 machine with a minimum of two spare processing cores. These spares can be shared across the drawers. The z13 offers core-level (engine-level) sparing.

On the z13, can any of these spare processing cores be used for other purposes?

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

How many System Assist Processors (SAPs) are on the z13?

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

The z13 Model N30 has six.

The z13 Model N63 has twelve.

The z13 Model N96 has eighteen.

The z13 Model NC9 and NE1 have twenty four.

These are in addition to the "visible" customer configurable cores. For example, an NE1 has 24 SAPs, 2 mandatory minimum spare cores, and up to 141 customer configurable cores.

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

What is the LPAR absolute hardware capacity setting on a z13?

Clients, particularly our Linux clients 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 allows 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.

What is Absolute Capping of an LPAR group on a z13?

LPAR group capping is similar to the hardware “absolute” capping of individual partitions. PR/SM and the Hardware Management tool have been enhanced to support an option to limit the amount of physical processor capacity consumed by a group of logical partitions (LPARs).

Group capping allows customer to set a cap on the physical usage of processors across a group of partitions. Using the HMC, the customer defines the group of partitions such that they are not able to consume more than the defined number of processors-worth (CP or IFLs) of capacity across the group. For example, LPARs 3, 4 and 7 form a group that is capped to not use more than seven Linux cores

Within a defined group, individual LPARs can also have their own independent Absolute Cap.

What is the radiator-based cooling on the z13?

The z13 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. A fill and drain tool is required for install and some radiator service actions.

Is there a water cooled option available for the z13?

Yes optional water cooling is available. 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.

Can I buy a z13 that has only ICF processors without including a general-purpose processor (CP)?

Yes. You can order only ICFs in a z13, using a model capacity identifier of 400 with a maximum of 141 ICFs. There is still a limit of 16 ICF engines for a single Coupling Facility LPAR.

Can I buy a z13 that is made up entirely of Integrated Facility for Linux (IFL) engines and no general-purpose processor (CP)?

Yes. You can order an IFL only z13, using a capacity identifier of 400 with 1 to 141 IFLs. For more information, ask your IBM or Business Partner representative for more information.

Is there still an option for overhead cabling on the z13?

Yes, you can order overhead cabling and you can also order overhead power on the z13.

Tell me about the non-raised floor option on the z13?

The non-raised floor option allows a z13 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 z13, you must order overhead power, overhead cabling, and radiator-based air-cooling.

Why might I be interested in the z13 new rack mounted HMC or new rack mounted TKE?

The new optional rack mounted management console can help to save space when it is a premium in data centers or help meet data center best practices.

Will my order for the rack mounted HMC or rack mounted TKE on a z13 include the rack?

No. Most of our clients already had 19" racks with available space in them, so we will not provide a rack. If you need one you'll need to work with your sales rep to purchase one.

Is there still going to be a zBX available for the z13?

Yes. Clients who have embarked on our hybrid computing strategy will see investment protection for their zBX. The z13 will support the IBM z BladeCenter Extension (zBX) Model 004. The zBX Model 004 is only available as an upgrade from either a zBX Model 003 or zBX Model 002. The zBX Model 004 is designed for clients that desire to continue managing the integrated hybrid environment with IBM z Unified Resource Manager.

What availability enhancements were made to the z13 (January 2015 announcement)?

The z13 continues the drive for continuous reliable operation provided by its predecessors with the following Reliability, Availability and Serviceability (RAS) improvements:

Improved soft error resilience in the processor cores

Lane shadowing, hardware buffer retry, and independent channel recovery which are designed to improve the DIMM interface

Continued use of RAIM in the main memory to protect DRAM

Improved robustness in the level 3 and level 4 cache

Improved isolation with the addition of integrated time domain reflectometry logic to chip interfaces

RAS changes to adjust to the new CPC drawer structure

Enhanced integrated sparing designed to reduce the complexity and number of repair actions

I am interested in making sure my IT personnel are well trained in all aspects of IBM z Systems. Are there resources to help me with client education?

Yes, IBM STG Lab Services and Training and our Global Training Providers can assist with this. You can contact IBM STG Lab Based Services via the Internet at: http://www.ibm.com/systems/services/labservices/ or send an email to stgls@us.ibm.com

End of z13 FAQs from the January 2015 FAQ document.

Performance

What are the major changes to the z/OS V2R1 LSPR?

The LSPR ratios reflect the range of performance between z Systems 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/lsprindex

What is the multi-image table in the LSPR?

Typically, IBM z Systems 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 z Systems processors.

What multi-image configurations are used to produce the LSPR multi-image table?

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.

Can I use the LSPR multi-image table for capacity sizing?

For high-level sizing, the multi-image table may be used. However, the most accurate sizing requires 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 configuration reflected in the multi-image LSPR table.

What model is used as the "base" or "reference" processor in the z/OS V2R1 LSPR table?

The 2094-701 processor model is used as the base in the z/OS V2R1 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.

What "capacity scaling factors" are commonly used?

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 V2R1 studies, the capacity scaling factor commonly associated with the reference processor set to a 2094-701 is 593 which is unchanged from that used originally 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 V2R1 multi-image table the commonly used scaling factor is .944x593=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 V2R1 multi-image table by 559.792. The PCI (Processor Capacity Index) column in the z/OS V2R1 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.

How much variability in performance should I expect when moving a workload to a z13 or z13s processor?

As with the introduction of any new server, workloads with differing characteristics will see variation in performance when moved to a z13 or z13s. 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 z13 and z13s each have significant improvements in its micro-processor design (major pipeline restructure) and in its memory hierarchy (increased cache sizes at all levels). However, workloads with different characteristics will see varying performance values from these changes. This suggests that the range of variation in performance of workloads will be somewhat larger than that seen in recent processor generations.

Once my workload is up and running on a z13 or z13s, how much variability in performance will I see?

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 support for larger numbers of engines, the z13 and the z13s have the capability to deliver up to 1.4 times the capacity of the largest previous server in each family. 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 z13 and the z13s.

How do I get performance information for my TPF products running on a z13 or z13s?

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.

What is HiperDispatch on the z13 and z13s and how does it impact performance?

HiperDispatch is the z/OS exploitation of PR/SM's Vertical CPU Management (VCM) capabilities and is exclusive to z Systems processors since the IBM System z10®. Rather than dispatching 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 logical processors. 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 efficiency 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 IBM zEnterprise 196 (z196) and later processors due to their sensitivity to the chip-level shared cache topology.

A white paper is available concerning HiperDispatch at:
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101229

What is the performance improvement a z/VM customer might experience on the z13 or the z13s?

The performance ratios that a z/VM customer workload might experience when migrating to z13 or z13s from older processors will vary. For z/VM, a single workload was run for the LSPR which has characteristics similar to the AVERAGE 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.

Where can I read more about the performance of z/VM?

The z/VM Performance Resources Page, located at http://www.vm.ibm.com/perf/ contains information on z/VM Performance.

What is the performance improvement a z/VSE customer might experience on the z13 or z13s?

The performance ratios that a z/VSE customer workload might experience when migrating to a z13 or z13s is represented by the range of ratios for a comparable z/OS migration. For example, the published ratio in the LSPR between the zEC12 702 and the z13 702 is approximately 8% to 16%. z/VSE workloads should expect this same range of performance for this migration. Consult the LSPR for other examples of moves to z13 or z13s.

z13 Warranty

What is proprietary diagnostic support for z13 or z13s?

A z13 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 z Systems Service Support Representative (zSSR) in doing problem repairs.

How long is a z13s or z13 entitled to proprietary support?

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

What happens when the warranty on a z13s or z13 ends and there is no maintenance agreement in place or maintenance is being provided by someone besides IBM?

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

What is the difference between base and proprietary service state for a z13s or z13?

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.

Can other service providers still fix a z13s or z13?

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

If I am on base support on a z13s or z13, what level of support will I get if I upgrade (MES) my system?

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.

Will I lose the call home capability if my z13s or z13 is no longer on an IBM maintenance agreement?

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

Hardware Management Console (HMC)

What is the new Hardware Management Console (HMC) code for z13s and z13?

The z13s will only support an HMC with code 2.13.1. The initial offering of the z13 shipped with 2.13.0. Those purchased after March 10, 2016, will have 2.13.1.[4] Enhancements to 2.13.1 include Absolute Capping of a LPAR Group, removal of Java Applet HMC task implementations to help with remote HMC browsing, a security enhancement for data protection, a new ASCII translation for display/alter of storage data and SNMP/BCPii performance enhancements. Enhancements to 2.13.1 and 2.13.0 include no longer having a Systems Activity Display (SAD), support for a new user management dashboard, alternatives for the USB Flash Memory drive, changes to the HMC time source, data replication versioning (all HMCs must be at the same level to be able to perform data replication with their HMC).

What is the new Absolute Capping of an LPAR group for z13s or z13?

LPAR group capping is similar to the hardware "absolute" capping of individual partitions. PR/SM and the Hardware Management tool have been enhanced to support an option to limit the amount of physical processor capacity consumed by a group of logical partitions (LPARs).

Group capping allows customer to set a cap on the physical usage of processors across a group of partitions. Using the HMC, the customer defines the group of partitions such that they are not able to consume more than the defined number of processors-worth (CP or IFL) of capacity across the group. For example, LPARs 3, 4 and 7 form a group that is capped to not use more than seven Linux cores

This is done without z/OS workload management (WLM) so it is very useful for non-z/OS environments where subcapacity pricing support (group soft capping) is not applicable.

What enhancement to the HMC will help me protect sensitive data on a z13s or 13?

Currently, we advise all customers who need a secure connection to install a VPN device between their server and client infrastructure when using OSA-ICC channels. To avoid this requirement and provide additional customer value, IBM has added Transport Layer Security (TLS) / Secure Sockets Layer (SSL) with Certificate Authentication with the HMC/SE support providing the controls for the certificate management. The OSA Integrated Console Controller (ICC) will use the OSC Physical Channel ID (PCHID) to provide a secure and validated method for connecting clients to the z Systems host.

The HMC/SE support has a new Secure TCP port field added for each LAN port on the Edit Server Configuration panel. Additionally, the Card Specific Advanced Facilities will be updated to handle the certificate management.

---

[4] Earlier z13 servers can be upgraded to the new firmware after availability.

What new capability is available for remote browsing of a HMC – particularly in terms of Java Applets for z13s or z13?

In HMC 2.13.0, from the initial announcement of the z13, there were a few customer tasks which are implemented via Java Applets. These include Operating System Messages, Integrated 3270 Console, Integrated ASCII Console, and Text Console. That Java applet approach has created challenges for customers who use remote browsing to their HMC, and when the Java level changes on their remote workstations, it might cause those tasks to no longer work until some special action such as a custom Java policy file is created to allow them to continue to work.

In the HMC 2.13.1 release associated with the announcement of the z13s and updated z13, the implementations of the above specified customer tasks will no longer be Java Applet based, and a remote browser execution of those tasks will no longer be affected by changes of Java levels on the remote workstations. In some cases the look and feel of those tasks will change slightly. In addition, the Integrated 3270 Console has been enhanced. When multiple remote users are connected into the same HMC or different HMC for a specific LPAR, any remote user will be able to take control of the Integrated 3270 Console session.

The implementation changes to remove the Java Applet implementations for the Operating System Messages, Integrated 3270 Console, Integrated ASCII Console, and Text Console tasks will also be patched back to the HMC 2.12.1 code level around the time of the HMC 2.13.1 is released. This should be noted since customers will also get the advantages of this new implementation approach on the HMC 2.12.1 level, but also the appearance/controls of those tasks will be slightly different once those MCLs are applied on the HMC.

What is the enhancement to the HMC of a z13s or z13 that will be most helpful to Linux on z?

Today, Display/Alter displays an EBCDIC translation of storage to the user in addition to a hex display of the data. In the new HMC 2.13.1 (available for z13s and updated z13), the user will now also be able to view an ASCII translation of storage. This is very useful for operating systems such as Linux on z that handles text in ASCII.

What can you tell me about the FCP SAN discovery tool for z13s or z13?

The FCP SAN Explorer function on the HMC has been enhanced with additional functions to facilitate SAN configuration setting and debugging. This facility can now display the name of the active fabric zone set, a list of zones an initiator is a member of, and enhanced diagnostic data for the initiator and target fabric links.

What are the performance enhancements for SNMP/BCPii for a z13s or z13?

Performance enhancements for SNMP/BCPii will be realized by allowing query or set operations for multiple attributes and/or category based attributes via new HwmcaEnhancedGet and HwmcaEnhancedSet API interfaces. Previously, an individual API request was required for each attribute being acted upon. Additionally, a single interface attribute was limited to a single LPAR, but now a single attribute query or update API can be sent to multiple LPARs.

In addition, a new SE task will be provided to better understand what BCPii requests are being made to the SE. For BCPii there could be multiple sources of applications making BCPii requests to the SE, and this task will provide a summarized view of historical data for the customer to better understand the source, frequency, and interactions of the BCPii initiated requests. This will allow

the customer to potentially tune the application requests of the multiple BCPii requesters to work towards a better performance model for their BCPii request environment.

The following z13 FAQs are from the January 2015 FAQ document and are valid for both the z13 and the new z13s.

Why might I be interested in the new rack mounted HMC on the z13s or z13?

The new optional rack mounted management console can help to save space when it is a premium in data centers or help meet data center best practices.

Will my z13s or z13 order for the rack mounted HMC include the rack?

No. Most of our clients already had 19" racks with available space in them, so we will not provide a rack. If you need one you'll need to work with your IBM or BP sales rep to purchase one.

What is the new minimum Hardware Management Console (HMC) code for z13s and z13?

Enhancements to 2.13.0, and later, include no longer having a Systems Activity Display (SAD), support for a new user management dashboard, alternatives for the USB Flash Memory drive, changes to the HMC time source, data replication versioning (all HMCs must be at the same level to be able to perform data replication with their HMC), and validation for all licensed internal code delivery and updates onto itself with a digitally signed firmware process.

What is replacing the Systems Activity Display (SAD) on the z13s and z13?

The System Activity Display task is no longer supported for z13s and z13 systems. The Monitors Dashboard task (available since z196) provides equivalent functionality. The System Activity display functionality is available for legacy systems (zEC12/zBC12 or earlier) which are managed by the Hardware Management Console 2.13.0 or later.

For the z13s and z13,what can you tell me about the User Management Dashboard of the HMC?

The HMC has reengineered user related tasks by establishing a User Management Dashboard task which replaces the following tasks: User Profiles, Customer User Controls, Password Profiles, Manage Enterprise Directory Server Definitions, User Templates, and User ID Patterns. This User Management Dashboard task provides additional functionality such as more granular user management controls and inheritance controls for objects added to groups. It is recommended to look at the HMC *What's New* section in the online help or *IBM Knowledge Center* to view detailed descriptions for this new task including getting started tutorials for different usage scenarios.

My company's security policies aren't keen on having a USB Flash Memory drive available on the HMC or the z13s or z13, what are my alternatives?

With the new code for the z13s and z13, the HMC will provide alternative options for each task that currently has an option to utilize a USB Flash memory drive (note that the USB Flash memory drive is still supported). The alternatives are FTP Servers and Remote Browser from a workstation and they will be documented in *IBM Knowledge Center*. If you prefer no USB Flash Memory Drive

usage, there is the Feature Code 0845 'Read-Only Media Option' that can be selected during ordering.

Note that these options are only available for managed z13s and z13 systems. The legacy managed systems (zEC12/zBC12 or earlier) still have a requirement for the USB Flash Memory Drive.

Note that the USB Flash drive cannot be used to back up critical Support Element (SE) data on the z13s/z13. It can back up the HMC. For the SE there are two choices: Backup to the SE hard drive or backup to an FTP server.

What changed with the HMC time source for z13s and z13?

The HMC will no longer define its time source using the Add Object Definition. The time source is now defined on the Customize Console Date/Time task. This will provide a clearer identification of all defined time sources including validation of Server Time Protocol (STP) Coordinated Timing Network IDs (CTN IDs).

Where are humidity and altitude readings on an HMC for z13s and z13?

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

What is the LPAR absolute hardware capacity setting on the z13s and z13?

Clients, particularly our Linux clients 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, introduced with the zBC12, allows 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.

This is the end of the January 2015 FAQs for HMC.

Power requirements (including High Voltage DC Power option)

What is static power save mode for the z13?

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 z13 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 z13 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.

Static power save mode is not available for subcapacity processors, and thus this capability is not available on the z13s.

What is query maximum potential power on the z13s and z13?

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 z13s and z13 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 z Systems 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®.

What is line cord plan ahead (FC #2000) and Balanced Power Plan Ahead (FC#3003) on the z13?

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.

How many line cords will be required for High Voltage DC Power option on the z13?

The same number of line cords as are required for AC operation.

What nominal DC supply voltages will be supported for the z13s and z13?

We'll support 380 VDC – 520 VDC nominal.

Overhead Cabling and Power

How much additional room is needed for the overhead cabling options on the z13s or z13?

On the z13s, overhead cabling will add 6" (~15 cm) to the width (not depth) of the frame on the z13s. Side-cable "chimneys" are attached on the left side of the server.

On the z13, overhead cabling will add 12" (~30 cm) to the width (not depth) of the combined frames on the z13. 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 z13 system's optional water connections are bottom exit only.

Will overhead options on the z13s or z13 add additional weight?

Yes and no. At maximum configuration on the z13s, overhead cabling adds approximately 78 lbs (~35 kg) to the weight. At maximum configuration on the z13, overhead cabling adds approximately 158 lbs (~70 kg) to the weight. Overhead power will add no additional weight to the z13s or the z13.

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.

Will I need to have additional space in my floor cutouts for the z13s or z13 because of these 'raceways'?

Top-exit I/O cabling will add 6" of width to the z13s and 12" of width to the z13, but the increases will only be above the floor to the sides of the system.

What benefit is there to add overhead power to the z13s or z13 after installation?

Overhead power is shipped separately and installed on-site to allow for door clearance.

Any configuration things I should watch for when ordering overhead power or overhead I/O on a z13s or z13?

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.

z13 Water cooling option

When should I consider installing water cooling on the z13?

A few examples are:

If you have installed zEC12 or 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 up to about 6.5%.

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. The z Systems 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.

Which z13 models will benefit most from water cooling?

Three- and four-processor drawer 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-processor drawer systems.

Can I tap into my building's chilled water to add the water cooling option on the z13?

A large number of data centers use chilled water somewhere in their cooling infrastructure, for example to provide cooling to CRAC (Computer Room Air Conditioner) or CRAH (Computer Room Air Handler) units. The chilled water requirements (temperature range, pH, hardness, contamination, particulates, etc.) for the z13 server should match the characteristics of the data center's available chilled water in most cases.

How long will it take me get my data center ready for a water cooled system when installing a z13 that has it?

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.

What changes will I need in my floor tile cut outs for water cooling on a z13?

There is no change to the floor cutouts for the z13 with the water cooling option – but the option will add 4 inches depth to the rear of the server. The floor tile cutouts for z13 are the same as those for the zEC12 and z196.

Is there any change to the height of the z13 for water cooling?

No.

Is there any change to the weight of the z13?

With water cooling, the installed weight will increase approximately 75 pounds (23 kg).

How many Water Cooling Units (WCU) are on a water cooled z13?

The z13 has two Water Cooled Units (WCU), in a N+1 design.

Will data center water be going through my z13?

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 the system that is used to fill the system initially and to service it if needed.

Is there any basic maintenance required because of water cooling on the z13 once the system is installed?

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.

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 on the z13?

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.

If my z13s or z13 is installed off a raised floor, can some cables exit the bottom of the machine?

No. 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.

If my z13s or z13 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?

Yes; but, note that if top exit power is also specified, power cables MUST exit the top.

Secure Service Container

What is Secure Service Container on a z13s or z13?

The Secure Service Container is a new partition type which, along with an appliance installer, enables the secure deployment of firmware and software virtual appliances. Secure Service Container will be delivered as part of the base code on each z13s or z13 with updated 2016-level microcode.

Secure Service Container also provides a framework which the appliance software solutions use as base – which includes the operating system and middleware.

What is the value of having the Secure Service Container on the z13s or z13?

Secure Service Container will shorten the deployment and implementation of select firmware and software virtual appliances. It supplies a common appliance installer to ease installation of the software virtual appliance.

Secure Service Container virtual appliances are protected against tampering and can be encrypted so that admins and people with root authority are not able to view the appliance data or code.

Will I only be able to install a software virtual appliance when I do a new server or update my firmware on the z13s or z13?

No, software virtual appliance solutions can be released independent from hardware releases – you should expect to see them delivered as software products are delivered today – either from shopz or via a web download.

Will I need to install a specific operating system to support a software virtual appliance on the z13s or z13?

No, Secure Service Container provides the framework for the software virtual appliances, which includes an embedded operating system, the necessary middleware and the solution software. You get a full image that contains the whole solution so that you can deploy on a Secure Service Container LPAR without having to install any OS nor middleware before.

What differentiates the Secure Service Container LPAR partition type on a z13s or z13?

The Secure Service Container LPAR type provides a substantial security envelope for appliances, checked by the Secure Service Container bootloader. It protects the appliance image from unauthorized modifications and it enables security related certifications of appliances, even from system administrators.

The protected resources include delivery media (including updates), installed code (including the bootloaders), any data that is produced by the appliance and the memory on the server.

The image is kept encrypted and signed on the disk at all times.

Will you allow the Secure Service Container to be on IBM zEnteprise BC12 (zBC12), IBM zEnterprise EC12 (zEC12), IBM zEnterprise 114 (z114) or IBM zEnterprise 196 (z196)?

No, Secure Service Container is only available on the z13s and the z13 at firmware level 27. Only selected appliances that do not depend on the security features listed above might be deployed on general Linux LPARs on older servers.

What type of appliances will be available for use in a Secure Service Container partition on the z13s or z13?

There are several appliances that we are looking at – in the areas of integrated analytics and security. Look for announcements in the future. We wanted the z13s and z13 to have the needed infrastructure to support the software virtual appliances as soon as they were available.

Can I include the software virtual appliance in my existing ELA (Enterprise License Agreement) for the z13s or z13?

IBM cannot comment on unannounced products. You will need to await the announcement of the software virtual appliance to get that answer.

Can ISVs use the Secure Service Container framework to host their software virtual appliances on the z13s or z13?

Secure Service Container does not support ISV software virtual appliances today, but this sounds like a good opportunity that IBM will consider for the future.

z/VSE Network Appliance

What is the z/VSE Network Appliance?

z/VSE includes a function called Linux Fast Pass (LFP). It allows selected TCP/IP applications to communicate via the TCP/IP stack on Linux on z System without requiring a TCP/IP stack on z/VSE.

On z13 or z13s with the correct firmware level the z/VSE Network Appliance is an integrated solution providing TCP/IP stack functionality. TCP/IP applications running on z/VSE in LPAR may use LFP to communicate via the z/VSE Network Appliance.

What are the benefits of having the z/VSE Network Appliance?

Using the z/VSE Network Appliance instead of a TCP/IP stack on z/VSE may offer up to 3x faster throughput due to a shorter path length and less overhead. The z/VSE Network Appliance may reduce the MSU consumption in z/VSE by up to 20% compared to a TCP/IP stack within z/VSE.[5] The benefits depend on your workload.

For heavy loaded environments, the zVSE Network Appliance may free up z/VSE resources bypassing the TCP/IP stack on z/VSE.

The LFP function is provided with the z/VSE products at no additional cost. For selected TCP/IP applications no TCP/IP product is required on z/VSE. The z/VSE Network Appliance removes the need for a Linux distribution for communication with LFP.

Does this replace z/VSE Linux Fast Path (LFP)?

No. LFP will continue to be available for environments, where z/VSE runs under z/VM through the z/VSE z/VM IP Assist (VIA) or where z/VSE is connected to a Linux on z Systems server in an LPAR or under z/VM.

If I don't have a z13s or z13 with the correct firmware, can I still use the z/VSE Network Appliance?

No. However, you may use LFP for environments where z/VSE runs under z/VM through the z/VSE z/VM IP Assist (VIA) or where z/VSE is connected to a Linux on z Systems server in an LPAR or under z/VM.

What level of z/VSE is required to support z/VSE Network Appliance?

The z/VSE Network Appliance is supported on z/VSE 6.1, 5.2, and 5.1.

Where do I get a copy of z/VSE Network Appliance?

When it is available you can download it from the web.

---

[5] Based on IBM measurements with high FTP workload

Can I use z/VSE Network Appliance if I'm running z/VSE under z/VM?

No, the z/VSE Network Appliance cannot be used under z/VM. z/VSE systems running under z/VM can either use LFP with the z/VSE z/VM IP Assist (VIA) or a Linux on z Systems server under z/VM.

IBM Dynamic Partition Manager

What is IBM Dynamic Partition Manager?

IBM Dynamic Partition Manager is designed to perform simplified configuration for Linux users. It allows partitions to be quickly configured, along with the management of system resources including integrated dynamic I/O management, as easily as other virtualized environments. It was developed for new-to-z users working on servers with z/VM 6.4, KVM and/or Linux as a partition-hosted operating system.

What are the benefits of having IBM Dynamic Partition Manager?

IBM Dynamic Partition Manager allows administrators that are new to the environment to be able to:

- Quickly create a new partition, including the I/O configuration, from a single management end-point
- Modify system resources without disrupting running workloads
- Monitor sources of system failure incidents and conditions or events which might lead to workload degradation
- Create alarms for events, conditions, and state changes
- Update individual partition resources to adjust capacity, redundancy, availability, or isolation.

Does IBM Dynamic Partition Manager replace PR/SM?

No, IBM Dynamic Partition Manager is not a replacement for PR/SM – it IS PR/SM. More specifically, it is a new administrative mode of PR/SM that allows for simplified configuration of partitions, associated resources and I/O. To use the new mode, rather than the 'classic' mode of PR/SM, the machine needs to be IML'd in IBM Dynamic Partition Manager mode.

What operating systems does IBM Dynamic Partition Manager support?

IBM Dynamic Partition Manager supports Linux running virtualized with z/VM V6.4, KVM or Linux running as a partition-hosted operating system. IBM Dynamic Partition Manager does not support z/OS, z/VSE or z/TPF.

Are there storage/disk restrictions with IBM Dynamic Partition Manager on the z13s or z13?

IBM Dynamic Partition Manager supports FCP storage devices only. You cannot use it with ECKD™ devices.

What servers support IBM Dynamic Partition Manager?

IBM Dynamic Partition Manager is supported by IBM z13s, IBM z13, IBM LinuxONE Rockhopper (2965) and IBM LinuxONE Emperor™ (2964).

Anything I should be aware of when I switch my server to IBM Dynamic Partition Manager?

IBM Dynamic Partition Manager must be activated with an IML.

Can I run both z/OS (or z/VSE or z/TPF) and a Linux server in IBM Dynamic Partition Manager on a single z13s or z13?

No, z/OS (or z/VSE or z/TPF) requires ‘classic’ mode of PR/SM, and z Linux server in IBM Dynamic Partion Manager, nees the z13s or z13 machine to be IML’d in IBM Dynamic Partition Manager mode.

I/O

What can you tell me about FICON Express16S?

With the introduction of FICON Express16S on the z13s and z13, you now have additional growth opportunities for your storage area network (SAN). FICON Express16S supports a link data rate of 16 gigabits per second (Gbps) and auto-negotiation to 4 or 8 Gbps for synergy with existing switches, directors, and storage devices. With support for native FICON®, High Performance FICON for z Systems (zHPF), and Fibre Channel Protocol (FCP), the z13s and z13 servers enables you to position your SAN for even higher performance – helping you to prepare for an end-to-end 16 Gbps infrastructure to meet the lower latency and increased bandwidth demands of your applications.

The new FICON Express16S channel will work with your existing fiber optic cabling environment, both single mode and multimode optical cables.

What kind of increased performance can I expect to get with zHPF using FICON Express16S?

Using FICON Express16S on a z13s or z13, large data transfer I/O operations with zHPF (reads + writes) can achieve 2600 MB/sec a 63% increase in throughput over FICON Express8S on a zBC12 or IBM zEnterprise EC12 (zEC12).

Using FICON Express16S on a z13s or z13, small data transfer I/O operations with zHPF can achieve 93,000 IO/sec an improvement over FICON Express8S on a zBC12 or zEC12.

Disclaimer: 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.

What kind of increased performance can I expect to get for the FCP protocol using FICON Express16S?

Using FICON Express16S on a z13s or z13, large data transfer I/O operations with FCP (reads + writes) can achieve 2560 MB/sec a 60% increase in throughput over FICON Express8S on a zBC12 or zEC12.

Using FICON Express16S on a z13s or z13, small data transfer I/O operations with FCP can achieve 110,000 IO/sec a 20% increase in throughput over FICON Express8S on a zBC12 or zEC12.

Disclaimer: 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.

What will Forward Error Correction (FEC) do for the z13s or z13?

The z13 was the first system to use a standards based approach for enabling Forward Error Correction (FEC) codes to improve resilience by reducing I/O errors. The improvement provided by FEC is the same improvement that would occur if the optical signal strength would be doubled. This technology will allow z Systems I/O to operate at higher speeds, over longer distances, with

reduced power and higher throughput, while retaining the same reliability and robustness that FICON has traditionally been known for. This capability will also be available for the z13s.

On the IBM DS8870, FEC helps to preserve data reliability for Peer to Peer Remote Copies. IBM announced the support of this new I/O infrastructure for DS8870 Storage Systems on May 11, 2015.

What is FICON Dynamic Routing?

With the z13s and the z13 server, FICON channels are no longer restricted to the use of static Storage Area Network (SAN) routing policies for Inter-Switch Links (ISLs) for cascaded FICON directors. The z Systems feature that supports dynamic routing in the Storage Area Network (SAN) is called FICON Dynamic Routing (FIDR). It is designed to support the dynamic routing policies provided by the FICON Director manufacturers, for example, Brocade's Exchange Based Routing (EBR) and Cisco's Open Exchange ID Routing (OxID). Please check with the switch provider for their support statement.

FICON Dynamic Routing can help clients reduce costs. It does this with the ability to share SANs between their FICON and FCP traffic (e.g. IBM's DS8870 Metro Mirror technology). It can help to improve performance due to SAN dynamic routing policies better exploiting all the available ISL bandwidth through higher utilization of the ISLs. And it can help to simplify management of their SAN fabrics due to static routing policies assigning different ISL routes with each power-on-reset which makes the SAN fabric performance difficult to predict. Clients will need to ensure that all devices in their FICON SAN support FICON Dynamic Routing before they implement this feature.

FICON Dynamic Routing on the IBM DS8870 enables clients to use Brocade Exchange Based Routing (EBR) or CISCO OxID routing across cascaded FICON Directors to simplify configuration and capacity planning, provide persistent and repeatable performance and higher resiliency. Sharing of switches is simplified and hardware costs can be reduced in Peer to Peer Remote Copy configurations by allowing FICON and FCP to share the same switch infrastructure. As IBM's Metro Mirror technology uses FCP as the transport, FICON and Metro Mirror, will flow over the same Inter Switch Links (ISLs) and be managed with consistent fabric priority. IBM announced the support of this new I/O infrastructure for DS8870 Storage Systems on May 11, 2015.

What availability enhancements will I receive with zHPF Extended Distance II?

High Performance FICON for z Systems (zHPF) has been enhanced to allow all large write operations (> 64 KB) at distances up to 100 km to be executed in a single round trip to the control unit thereby not elongating the I/O service time for these write operations at extended distances.

zHPF Extended Distance II on IBM DS8870 or DS8880 allows clients to achieve service level agreements after a disaster or when a storage control unit failure causes a HyperSwap® event. This capability is required especially for GDPS® HyperSwap configurations where the secondary DASD subsystem is in another site.

IBM announced the support of this new I/O infrastructure for DS8870 Storage Systems on May 11, 2015.

What enhancements are on the z13s and z13 for increasing scalability?

The z13 has channel subsystem (CSS) scalability with support for six logical channel subsystems (LCSSs) which are required to support the eighty-five LPARs for z13, four subchannel sets (to support more devices per logical channel subsystem), and 32K devices per FCP channel. Additionally, a fourth subchannel set for each logical channel subsystem (LCSS) is provided to facilitate elimination of single points of failure for storage after a disk failure by simplifying the exploitation of IBM's DS8870 Multi-target Metro Mirror storage replication with TPC-R HyperSwap.

The z13s channel subsystem (CSS) scalability will support up to three logical channel subsystems (LCSSs) which are required to support the forty LPARs for z13s, three subchannel sets (to support more devices per logical channel subsystem), and 32K devices per FCP channel. Additionally, a third subchannel set for each logical channel subsystem (LCSS) is provided to facilitate elimination of single points of failure for storage after a disk failure by simplifying the exploitation of IBM's DS8880 Multi-target Metro Mirror storage replication with TPC-R HyperSwap.

Does the z13s and z13 still use the PCIe I/O drawer for the I/O infrastructure?

Yes. The PCIe I/O drawer and the form factor I/O cards support a direct Peripheral Component Interconnect Express Generation 3 (PCIe Gen3) infrastructure with increased capacity, granularity, and infrastructure bandwidth, as well as increased reliability, availability, and serviceability. PCIe Gen 3 supported hardware features include FICON Express16S, FICON Express8S, OSA-Express5S and the Crypto Express5S.

With their compact design, two PCIe I/O drawers occupy the same space as one I/O cage while delivering more capacity.

The industry-standard PCIe Gen3 infrastructure from the processor drawer 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.

Will I be able to carry forward an I/O drawer from my zBC12, zEC12, z114 and z196, and also which of my PCIe I/O features can I continue to ONLY carry forward to the z13s and z13?

A z13s or z13 new build will only support PCIe I/O drawers, but on an upgrade you can carry forward a maximum of two I/O drawers from the zBC12, zEC12, z113 and z196. This means that a maximum of 16 FICON Express8 features (FC 3325 & 3326) for z13, and 8 FICON Express8 features for the z13s are the only features that can be carried forward and are not available with a new build z13s or z13.

Why isn't there anything about ESCON® in the z13s or z13 announcement?

Since April of 2008, z Systems 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 z13s or z13.

What Storage Area Network (SAN) products are currently qualified to operate at 16 Gbps with z Systems?

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 z Systems FICON and FCP channels”.

IBM zHyperLink for IBM Storage and IBM z Systems

What are zHyperLinks?

The zHyperLink technology is a new mainframe attach link. It is the result of collaboration between DB2 for z/OS, the z/OS operating System, IBM z Systems® processors and DS8880 storage to deliver the extreme low latency I/O access for DB2 for z/OS Applications.

What is the expected value of low latency I/O for DB2 on z/OS applications using zHyperLinks?

zHyperLinks are expected to save client's money, improve the scalability of DB2 for z/OS, enhance system resilience and provide industry leading I/O latency for transaction processing.

Do zHyperLinks replace FICON® technology?

No, zHyperLink technology is intended to complement FICON technology to accelerate those I/O requests that are typically used for transaction processing. These links are point-to-point connections between the CEC and the storage system and are limited to 150 meter distances. These links do not take away from the z Architecture 8 channel path limit.

Will my current DS8880 support IBM zHyperlink technology?

Yes. Clients will be able to upgrade their DS8880 storage devices to support IBM zHyperLink.

Flash Express

What is Flash Express on z13s and z13?

Flash Express is an optional feature 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.*

How can Flash Express improve the availability of my z13s and z13 environment?

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.

Which servers, operating systems and middleware support Flash Express?

The Flash Express feature is supported on z13s, z13, zBC12 and zEC12 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 z Systems running natively on the server. The support is available with RHEL 6.4 and above and SES 11 SP3 and above today. There are no use cases for Linux at this time.

KVM does not support Flash Express at this time.

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.

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

IBM is working with the Linux community to explore Flash Express exploitation opportunities.

How do I order Flash Express?

Flash Express is ordered through eConfig as an optional feature (FC#0403[6]) of z13s and z13. Cards are ordered 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)

How much Flash Express capacity will I need to order?

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.

Does Flash Express use the Hardware System Area (HSA)?

On the z13s and the z13, there is now read/write cache located in the system's Hardware System Area (HSA). The amount of flash cache on the z13s is .5 GB and, on the z13 is 4 GB. The benefit to HSA cache will be an advantage during diagnostic collection.

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

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

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

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

---

[6] Flash Express FC#0402 can be carried forward only

you retain the same amount of disk paging capacity that you have today even when using Flash Express.

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

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

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

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.

Where are Flash Express cards installed in the server?

The Flash Express cards will be installed in the PCIe I/O drawer.

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

Flash Express introduces storage class memory for the mainframe. 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.

Is Flash Express automatically used by the system?

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.

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

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.

What is dynamic reconfiguration for Flash Express?

Dynamic Reconfiguration allows a user to dynamically configure the amount of flash memory to be moved into and out of the partition.

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

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.

Is data stored on Flash Express persistent across IPLs or PLPA data?

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.

How can I use Flash Express to support Coupling Facility Control Code (CFCC)?

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) or later, this 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.

Is the data stored on the Flash Express card encrypted?

Yes. The data stored on the Flash Express cards is protected – even if these cards are removed from z Systems. 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 Support Element (SE). Systems with Flash Express cards use smart cards that are plugged into the 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 z Systems provides to help keep your data secure.

Is there any special configuration needed for encryption support with Flash Express?

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 SE.

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

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.

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

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.

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

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.

Will having additional pairs of Flash Express cards improve paging?

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.

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

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.

Will having more Flash Express cards improve resiliency?

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

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

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 cannot 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.

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

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.

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

No.

Networking

What can you tell me about the FCP SAN discovery tool?

The FCP SAN Explorer function on the HMC has been enhanced with additional functions to facilitate SAN configuration setting and debugging. This facility can now display the name of the active fabric zone set, a list of zones an initiator is a member of, and enhanced diagnostic data for the initiator and target fabric links.

What is the usability update for worldwide port naming (WWPN) assignments?

A z System server automatically assigns worldwide port names (WWPNs) to the physical ports of an FCP channel based upon the Physical Channel ID (PCHID). When an FCP channel is moved to a different physical slot position this WWPN assignment changes. The z13 and z13s will now allow for the modification of these default assignments, allowing FCP channels to keep previously assigned WWPNs, even after being moved to a different slot position. This capability can eliminate the need for reconfiguration of the SAN in many situations, and is especially helpful on a system upgrade as you can now import/export the naming.

What is the extended link service command to improve SAN diagnostics?

A new extended link service (ELS) command called Read Diagnostic Parameter (RDP) has been added to the Fibre Channel T11 standard to allow z Systems to obtain additional diagnostic data from the Small Form Factor Pluggable (SFP) optics located throughout the storage area network (SAN) fabric. The ultimate goal of the RDP will be to improve the accuracy of identifying a failed or failing component without unnecessarily replacing additional components in the SAN fabric (FICON cards, optics, cables, and so on).

z Systems FICON and FCP channels and the IOP code will provide a means to read this additional diagnostic data for all the ports accessed in the I/O configuration and make the data available to a z Systems partition. For z/OS clients utilizing FICON channels, z/OS will display the data with a new message and display command. For Linux on z Systems, z/VM, z/VSE, and KVM for IBM z clients utilizing FCP channels, this diagnostic data will be available in a new panel in the SAN Explorer tool.

The initial RDP support will be for data collection and display only. In the future, we will investigate the possibility of doing analysis on this collected data.

Tell me about OSA-Express5S:

The OSA-Express5S is for use in the PCIe I/O drawer and is supported by the 8 GBps PCIe Gen3 host bus. The OSA-Express5S features have small form factor pluggable+ (SFP+) transceivers.

The OSA-Express5S Gigabit Ethernet (GbE) short wavelength (SX) feature and the OSA-Express5S Gigabit Ethernet (GbE) long wavelength (LX) feature each has one PCIe adapter and two ports. The two ports on a feature share one channel path identifier (CHPID type OSD exclusively – Layer 3 TCP/IP or Layer 2 traffic). The ports support attachment to a one Gigabit per second (Gbps) Ethernet Local Area Network (LAN). Each port can be defined as a spanned channel and can be shared among logical partitions and across logical channel subsystems.

The OSA-Express5S 10 Gigabit Ethernet (GbE) short reach (SR) feature and the OSA-Express5S 10 Gigabit Ethernet (GbE) long reach (LR) feature each has one PCIe adapter and one port per

feature. Each port supports channel path identifier (CHPID) types OSD (Layer 3 TCP/IP or Layer 2 traffic), or OSX (intraensemble data network - IEDN). The 10 GbE SR feature is designed to support attachment to a multimode fiber 10 Gigabits per second (Gbps) Ethernet Local Area Network (LAN) or Ethernet switch capable of 10 Gbps. The 10 GbE LR feature is designed to support attachment to a single mode fiber 10 Gigabit per second (Gbps) Ethernet Local Area Network (LAN) or Ethernet switch capable of 10 Gbps. Each port can be defined as a spanned channel and can be shared among logical partitions within and across logical channel subsystems.

The OSA-Express5S 1000BASE-T Ethernet feature for copper environments has one PCIe adapter and two ports. The two ports share one channel path identifier (CHPID). The following CHPIDs are supported:

OSC (TN3270E and non-SNA DFT)

OSD (TCP/IP, Layer 2, Layer 3 traffic)

OSE (SNA/APPN/HPR traffic and TCP/IP passthru)

OSM (intranode management network - INMN)

OSN (OSA-Express for NCP)

10GbE RoCE Express and Shared Memory Communications – RDMA (SMC-R)

What is the LPAR sharing capability for 10GbE RoCE Express?

No, however the good news is that, similar to the z13, a 10GbE RoCE Express feature can be shared across LPARs in the new z13s. Single Root I/O Virtualization (SR-IOV) for 10GbE RoCE will be introduced to allow a PCIe device to appear to be multiple separate physical PCIe devices. The SR-IOV specification is an industry standard specification to help promote interoperability. This provides the ability for that PCIe device to present multiple instances of itself up to the Operating System instance or hypervisor. RoCE continues to support SMC-R for z/OS to z/OS communication or for Linux for z Systems to Linux for z Systems communication (z/OS LPAR or 2nd level guest running under z/VM).

What is RDMA?

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.

What is RoCE?

RDMA over Converged Ethernet (RoCE) "is a standard that defines a 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." [7]

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

SMC-R is a capability on z/OS V2R1 and later, which works in conjunction with the feature 10GbE RoCE Express, a PCIe feature on z Systems 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

---

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

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.

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

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.[8]

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 v8.5 Liberty profile TradeLite workload accessing z/OS DB2 in another system measured in internal benchmarks.[9] Another benchmark we have done shows WebSphere MQ for z/OS realizes up to 3X increase in messages per second it can deliver across z/OS systems when using SMC-R vs. standard TCP/IP.[10]

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.[11]

---

[8] 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.
[9] Based on projections and measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.
[10] 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.
[11] 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 (10 GbE RoCE Express feature) vs standard TCP/IP (10GbE OSA Express4 feature). The actual CPU savings any user will experience may vary.

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.).

What is the 10GbE RoCE Express feature?

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

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

Two 10GbE SFP+ ports are provided on each 10GbE RoCE Express feature. The z13s and z13 allows z/OS SMC-R exploitation to take advantage of both ports.

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

No, SMC-R does not require a RoCE-capable switch. SMC-R requires a standard 10GbE switch[12]. Auto negotiate to 1GbE is not supported. 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.[13]

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

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.

In support of multiple site (i.e. disaster recovery) configurations IBM has conducted SMC-R performance testing with 10GbE RoCE Express feature at longer distances (10km and 100km). In each case the support for the extended distances was provided by the switch optics (10km) and DWDM technology (100km). For additional information on these performance benchmarks refer to the SMC-R Reference Information "SMC-R / RoCE at Distance": http://www.ibm.com/software/network/commserver/SMCR

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

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/systemzQualifiedWdmProductsForGdpsSolutions?OpenDocument&pathID

---

[12] 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.
[13] 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.

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

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.

Will SMC-R and RoCE support SNA?

No. SMC-R requires TCP.

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

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

What is the feature code for 10GbE RoCE Express?

The 10 GbE RoCE Express is FC# 0411.

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

On the z13s and z13, the 10 GbE RoCE Express feature can support up to 31 logical partitions.

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

Each LPAR should be provisioned with 2 features (PFIDs) for redundancy. Up to 16 features are supported on the z13s and z13.

What operating systems support 10GbE RoCE Express technology?

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

z/VM 6.3 and later has been enhanced to provide guest exploitation of the 10GbE RoCE Express feature (#0411) 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 do so while running as a guest operating system under z/VM. For example, z/OS 2.1 can 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.

KVM does not have support for 10GbE RoCE Express feature at this time.

The 10GbE RoCE Express feature is supported by SLES12 SP1 and RHEL 7.2, but there are some issues being worked and we will notify our clients when support is available.

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

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.

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

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

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

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.

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

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.

How can I determine what level of benefit I might expect from SMC-R and the 10GbE RoCE Express feature?

IBM has provided a tool called the SMC Applicability Tool (SMCAT). The SMCAT has been provided for z/OS V1R13 and V2R1 (via PTFs). The tool does not have any dependency on RoCE, SMC-R, or SMC-D. Instead SMCAT will monitor your TCP/IP traffic for a specified IP address, group of IP addresses or IP subnets and then produce a summary report describing the network traffic associated with the monitored addresses. The summary report will provide information about how much traffic (percentage) to/from those IP addresses is eligible for and well suited for SMC-R and SMC-D.

SMCAT Overview document is available on the SMC-R reference material web site:

http://www.ibm.com/software/network/commserver/SMCR/
ftp://public.dhe.ibm.com/software/os/systemz/pdf/SMC_Applicability_Tool_Overview_3-03-15.pdf

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

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.

Shared Memory Communications – Direct (SMC-D)

What is Shared Memory Communications – Direct Memory Access?

Direct Memory Access (SMC-D) leverages the existing Shared Memory Communications protocol used over RDMA (SMC-R) to provide highly optimized inter-system operating system communications. Instead of using RoCE, SMC-D uses Internal Shared Memory (ISM) technology within the system. ISM provides adapter virtualization (virtual functions) to facilitate the internal communications. SMC-D does not require any additional physical hardware (provided with firmware within z13s/z13) – no adapters, card slots, switches, fabric management, PCI infrastructure – so there are cost savings from using the capability. It can be enabled in z/OS 2.2 with single TCP/IP profile keyword[14].

Does SMC-D replace HiperSockets?

SMC-D (like SMC-R) still requires an IP network, a TCP/IP connection and SMC-D only applies to connections using TCP sockets. The IP network required for SMC-D can be provided with an external LAN (e.g. OSA-Express) or provided by HiperSockets. SMC-D is only supported by z/OS. For these reasons, HiperSockets will still be needed (i.e. for Linux on System z environments, other network protocols, etc.). The SMC protocol bypasses TCP/IP processing (related to exchanging user data) for providing a direct optimized form of communications. The performance benefits of SMC-D is superior to (in some cases significantly better than) HiperSockets.

Some examples are:

Up to 61% CPU savings for FTP file transfers across z/OS systems versus HiperSockets[15]

Up to 9x improvement in throughput with more than a 88% decrease in CPU consumption and a 90% decrease in response time for streaming workloads versus using HiperSockets[16]

Up to 91% improvement in throughput and up to 48% improvement in response time for interactive workloads versus using HiperSockets[17]

What is the benefit of Internal Shared Memory (ISM)?

Internal Shared Memory (ISM) provides adapter virtualization with high scalability that allows logical partitions (or guest virtual machines) to logically share virtual memory.

ISM is based on PCIe architecture, which leverages the existing PCIe eco-system, assets such as HCD, IOCDS, IOCP and memory I/O virtualization and translation. ISM is provided by a VCHID (similar to HiperSockets). There can be up to 32 ISM VCHIDs per CPC and up to 255 virtual functions (VFs) per VCHID. That means there can be up to 8160 VFs per CPC.

What is the value I will get from SMC-D?

The value of SMC-D is about providing significantly improved performance for inter-system communications. The value is related to co-located workloads. These improvement attributes are

---

[14] ISM VCHIP and PFIDs must be defined in HCD (IOCDS)
[15] All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.
[16] All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.
[17] All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

typically described as latency, throughput and CPU cost. Improvements can potentially improve (increase) application workload transaction rates while reducing your CPU cost. The type and amount of improvement is based on the workload characteristics (i.e. OLTP vs streaming).

SMC-D is transparent to socket applications, require no IP topology changes and preserves connection level security. SMC-D does not expose data to the external network. It requires no additional hardware (such as adapters, card slots, switches, fabric management). It provides higher scalability, bandwidth and virtualization.

Can I use SMC-D to talk to another z13s or z13 system?

No, SMC-D is for intra-CPC communications only (two LPARS or z/VM guests on a single CPC).

For communications between CPCs SMC-R and RoCE is an option.

Can I use both SMC-D and SMC-R at the same time?

Yes, both SMC-R and SMC-D can be used at the same time. When the TCP connection is established the SMC protocol will dynamically determine which variations are supported by both hosts. When both peer hosts are enabled for SMC-D, then the protocol will determine if both peers are eligible for SMC-D connectivity (i.e. both hosts are on the same CPC, have access to the same ISM VCHID and VLAN (when applicable)) and then select SMC-D when possible. When SMC-D connectivity is not possible, then SMC-R eligibility is evaluated.

How can I determine if SMC-D applies to my environment and what level of benefit I might expect?

IBM has provided a tool called the SMC Applicability Tool (SMCAT). The SMCAT has been provided for z/OS V1R13 and V2R1 (via PTFs). The tool does not have any dependency on RoCE, SMC-R, or SMC-D. Instead SMCAT will monitor your TCP/IP traffic for a specified IP address, group of IP addresses or IP subnets and then produce a summary report describing the network traffic associated with the monitored addresses. The summary report will provide information about how much traffic (percentage) to/from those IP addresses is eligible for and well suited for SMC-R and SMC-D.

SMCAT Overview document is available on the SMC-R reference material web site:
http://www.ibm.com/software/network/commserver/SMCR/

Do I need to have a 10GbE RoCE Express feature?

No, there is no requirement for RoCE (RDMA over Converged Ethernet). A TCP/IP connection over QDIO OSA or HiperSockets (internal QDIO (iQDIO)) is required.

Do I need to have an OSA connection to take advantage of SMC-D?

Yes, a TCP/IP connection over QDIO OSA or HiperSockets (internal QDIO (iQDIO)) is required. The OSA connection may be a 1 GbE or 10 GbE connection but it must be defined as CHPID type of OSD (QDIO).

What is the feature code for SMC-D?

There is no orderable feature code. SMC-D will come standard with the z13s and with new updated microcode for the z13.

What operating systems support this technology?

z/OS 2.2 (plus PTFs), support for the SMC-D protocol, and is the only operating system that is announcing exploitation.

z/VM V6.3 (plus PTFs planned to be available by March 10, 2016) and above, will provide support for guest exploitation within the CPC.

IBM is working with its Linux distribution partners to include support in future distribution releases.

What changes do I need to make to the applications to take advantage of SMC-D?

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

What security is there for SMC-D? How do I know that data traveling across the memory is protected?

SMC-D 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 memory between operating systems can be protected with encryption, data integrity controls, authentication, access controls and so on.

Note: Additional technical details are available at http://www.ibm.com/support/techdocs

IBM zAware V2.0

With the announcement of IBM® Operations Analytics for z Systems®, v3.1, that includes IBM zAware function, what should I buy to get IBM zAware function?

Our clients have expressed concerns about the flexibility of the IBM zAware V2.0 firmware offering as it was difficult to add enhancements that bring new capabilities and functions to the product outside of the hardware release cycle or to do service updates on their own schedule.

IBM has included IBM zAware function as part of the IBM software product IBM Operations Analytics for z Systems™ v3.1 - see announcement letter 216-373 dated September 13, 2016. When installed in the Secure Service Container (formerly known as z Appliance Container Infrastructure), the IBM zAware V3.1 code packaged with Operations Analytics for IBM z Systems can be thought of as a software appliance versus a firmware appliance.

By using a software offering, our clients can add or upgrade via web downloads, apply maintenance as its available, and have more flexibility in managing the offering. Our suggestion is to purchase Operations Analytics for IBM z Systems and quit using IBM zAware.

What can you tell me about the Operations Analytics for z Systems software offering?

Operations Analytics for IBM z Systems enables users to search, visualize and analyze the vast amounts of structured and unstructured operational data across IBM z Systems operating environments, including log, event and service request data, as well as performance metrics. IBM Operations Analytics for z Systems is made up of three capabilities:

- Enhancements to Operations Analytics for z Systems including: Problem Insights dashboard eliminates the need for tedious searching of volumes of operational data; Expanded WebSphere Application Server Insights provides improved understanding of WAS request activities and CPU usage; Expert Advice capability when analytics server is deployed in isolation from the internet.
- IBM System z Advanced Workload Analysis Reporter, V3.1 (IBM zAware) provides an integrated, self-learning analytics solution that helps identify unusual behaviors of workloads running on z/OS logical partitions (LPARs) and Linux on z Systems.
- IBM Common Data Provider for z Systems, V1.1 provides access to z/OS operations data in near real-time. With its support for all standard SMF record types and a broad set of log types, it is the single point of data collection you need to gain deep insights into the operations of your z/OS-based data center.

For more information, go to the web page for IBM Operations Analytics for z Systems at:
http://www.ibm.com/software/products/en/ibm-operations-analytics-for-z-systems

If I have IBM zAware, as part of Operations Analytics for z Systems, can I monitor both z/OS and/or Linux messages?

Yes. Operations Analytics for z Systems, which includes IBM zAware V3.1, monitors both z/OS and/or Linux messages.

When I ordered the IBM zAware firmware feature of the hardware, I had the option to put an IBM zAware disaster recovery (DR) feature on machines that would run the IBM zAware host code if there was a disaster. Do I need a DR software license for Operations Analytics for System z?

No. The Operations Analytics for z Systems does not have a DR license. If you need a copy of the software for DR please work with your IBM or BP z Software representative.

Where can I get information about migrating my IBM zAware firmware offering to IBM zAware V3.1 that is part of Operations Analytics for z Systems.

There is an IBM zAware V3.1 Guide that has that information. It can be found at:
http://www.ibm.com/support/knowledgecenter/SS55JD_3.1.0/com.ibm.zosla.doc/zaware.pdf

What should I do if I have IBM zAware firmware on my z System today?

If you are currently using IBM zAware, you can continue to use the firmware. But IBM is strongly encouraging the z/OS® and mixed environment clients to move to the new software product in order to get more function and flexibility.

If you are a Linux only client on an IBM z System® or a IBM LinuxONE™ server, you should continue to use the firmware offering at this time.

If you haven't implemented or even installed IBM zAware yet, you should contact your IBM seller to discuss what options make the most sense for you.

Will there be refunds or credits if I have IBM zAware today?

There will be no refunds for IBM zAware past purchases.

What did IBM announce concerning the cut off dates for ordering new features for IBM zAware?

IBM has announced cut off dates for ordering new features for IBM zAware - see announcement letter 916-201 dated November 1, 2016.

The last date to add IBM zAware to Systems servers, will be December 30, 2016. After December 30, 2016, the features to add IBM zAware to a z System will no longer be available. This includes feature number 0011. After December 30, 2016 clients who do not already have IBM zAware and want to add IBM zAware functionality to their machines will need to order Operations Analytics for z Systems.

Clients who already have IBM zAware will be able to order additional IBM zAware feature packs until June 29, 2017. This includes feature numbers 0101, 0102, 1010, 1011, 1012, 1013, 1014, 1015, 1017, 1018, 1019, 1020, 0138, 0139, 0140, 0141, 0142, 0143, 0150, and 0151.

For customers migrating from zEC12 or zBC12 to z13 or z13s, the last date to order these upgrades that will allow the carry forward of the IBM zAware firmware feature numbers will be June 29, 2017.

Current z Systems servers with the IBM zAware feature are able to bring forward current IBM zAware definitions through a technical migration path to the new release of Operations Analytics for z Systems that includes the IBM zAware functionality.

According to the terms of the IBM zAware contract, if the machine is discontinued/sold or if the customer no longer needs IBM zAware, they need to remove the IBM zAware features from the machine (MES). There will be no refunds for IBM zAware past purchases.

Starting with the IBM z13s™ (z13s), and z13 at the firmware load from March 2016, IBM zAware is installed in a Secure Service Container (formerly known as zACI). Is that the same for the IBM zAware that comes with Operations Analytics for z Systems?

Yes, it is preferred that IBM zAware V3.1 is installed in a Secure Service Container, but let's look again at the three pieces of Operations Analytics for z Systems V3.1.

On IBM z13 (at GA2) and IBM z13s servers, IBM zAware V3.1 installs in its own Secure Service Container LPAR. If that isn't an option (such as having a zEC12 or other reason), IBM zAware V3.1 can install in a Linux on z Systems LPAR. There is no Linux operating system needed with no Linux operating system installed.

Operations Analytics for z Systems V3.1 installs in a Linux (native or guest) image and consumes data from each z/OS and Linux image. Alternatively, Operations Analytics for z Systems can run in Linux on x.

Lastly, Common Data Provider for z Systems installs in each z/OS LPAR and sends data to destinations of choice.

Will my customer need to remove the IBM zAware feature codes from the z Systems server if they are no longer using IBM zAware?

The short answer is 'Yes'. According to the terms of the IBM zAware contract, if the machine is discontinued/sold or if the customer no longer needs IBM zAware, they need to remove the IBM zAware features from the machine (MES).

What Statement of Direction about IBM zAware was announced on November 1, 2016?

The Statement of Direction from the November 1, 2016 announcement is that the last generation of z Systems servers that will support IBM zAware as a firmware feature are z13 and z13s.

The IBM zAware firmware feature is being replaced with Operations Analytics for z Systems, but it's still important to understand the IBM zAware offering and so these questions and answers are still included.

Why did IBM introduce IBM zAware?

Today's 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 can be difficult for many reasons, including the volume of system data and messages to examine. When an error occurs, it can be complex to diagnose and reflective of changes in application, system configuration settings, patches, and so on. IBM zAware provides an analytics solution for detecting and diagnosing anomalies in z/OS and Linux for z Systems that help IT staff more quickly understand unusual data center situations, detect anomalies and diagnose problems faster so that they can improve system health and availability.

Is IBM zAware an analytics application?

IBM zAware is an analytics application used to analyze z/OS OPERLOG messages and Linux on z Systems syslog messages.

IBM zAware is delivered as an integrated firmware appliance, and ordered as a feature of the server, available on the z13s, z13, zBC12 and zEC12.

IBM zAware uses a self-learning approach to model normal system behavior based on historical 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

How is IBM zAware different than other tools?

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.

What do you mean by 'IBM zAware monitored clients'? Can I receive OPERLOG and Linux on z Systems syslog messages from other machines in my data center?

Good questions. Let's regroup on some of the terms. IBM zAware runs in a partition, which must be on the z13s or z13. This partition is called the IBM zAware host partition. The IBM zAware host system is the z13s or z13 that hosts the IBM zAware host partition. There may be instances that you have multiple IBM zAware host partitions on one z13s or z13.

IBM zAware can accept messages to be analyzed from other z Systems 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) and all supported distributions for Linux for z Systems. The z/OS partition or Linux systems 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.

How is problem determination easier with IBM zAware?

Examination of the sheer volume of messages in an operations log (OPERLOG) or Linux on z Systems syslog message logs can make the analysis task daunting. For larger z/OS clients, over 25 million messages daily is not uncommon. There are over 40,000 unique message IDs defined for z/OS.

For the z13s and z13, IBM zAware is also able to process Linux messages, which have no message Ids. It is virtually impossible to identify problem messages given this volume without use of a tool. Also IBM zAware can identify anomalies on any of the IBM zAware monitored clients within a parallel Sysplex, to help pinpoint the origin and show the scope of impact.

How does IBM zAware develop a baseline?

IBM zAware needs to create a baseline model to perform basic analysis using historical data for building a more accurate model. Your installation can modify this training period based on knowledge of the workloads running on z/OS or Linux on z. 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.

Can I use IBM zAware to mark 'ignore' for some messages?

Yes for z/OS. 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.

Users can use 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.

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

Messages from the IBM zAware monitored client machine are sent to the IBM zAware host across your IP network or HiperSockets, by the z/OS Logger component or for Linux on z, the existing syslog daemon is used to send data to the IBM zAware host.

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

The maximum supported distance between the IBM zAware host system and monitored clients is 3500 km.

What configuration is required for the IBM zAware environment on a z13s, z13, zEC12 or zBC12?

The place to get information on sizing is the IBM z Advanced Workload Analysis Reporter (IBM zAware) Guide SC27-2623 available on ResourceLink http://www.ibm.com/servers/resourcelink

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

For z/OS monitoring, z/OS V1R13 with PTFs for system logger and bulk load client) or higher is required. For Linux on z Systems, monitored clients that use the existing syslog daemon support available in either SUSE Linux Enterprise Server (SLES) 10 or later or Red Hat Enterprise Linux (RHEL) 6 or later can be monitored.

These servers can be IBM zAware monitor clients:

- IBM z13s (z13s) or IBM z13 (z13)
- 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)

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

A z/OS can be an IBM zAware monitored client when running as a z/VM guest.

A Linux server can be an IBM zAware monitored client when running as a z/VM guest. Linux on z System can also be monitored when running under KVM for z.

Likewise, a native Linux server is a supported monitored client.

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?

A Disaster Recovery option is available to ship IBM zAware firmware to a z13s or z13 to be used in case of emergency. The quantity of IBM zAware connection DR features (FC #1011 on the z13 and FC #1017, #1018, #1019 or #1020 on the z13s) will be calculated the same way features are calculated for an IBM zAware host system but these will be no charge (zero priced) features.

How is the price of IBM zAware calculated for the z13s or z13?

- Pricing is based on a calculation of the total number of monitored z/OS CPs and monitored IFLs on the IBM zAware host server plus on the IBM zAware monitored client servers.
- For z/OS only - 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.
- For Linux on z only – The CP count is the sum of the purchased IFLs on the IBM zAware host *plus* the purchased IFLs on IBM zAware monitored servers not including CBU IFLs or On/Off CoD IFLs. The count must include high water marks IFLs.
- For both z/OS and Linux on z monitoring - The CP count is the sum of the purchased CPs and IFLs on the IBM zAware host *plus* the purchased CPs and IFLs on IBM zAware

monitored servers, not including no CBU, On Off CoD or zIIPs, zAAPs or CFs. The count must include high water marks CPs and IFLs.

- On the z13s, 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 then six, the pricing is similar to the z13. 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. On the z13, the minimum block of CP/IFLs that you can purchase is a 10-pack (equal to a count of 10 purchased CP/ILFs). 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.

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

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.

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

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.

What feature codes are needed to order IBM zAware on z13s or z13?

Features for ordering IBM zAware on the z13s will be located on the Processor usage panel of eConfig

- FC 1012 will represent the quantity of IBM zAware CP/IFL 2-packs (FC 1017 for DR)
- FC 1013 will represent the quantity of IBM zAware CP/IFL 4-packs (FC 1018 for DR)
- FC 1014 will represent the quantity of IBM zAware CP/IFL 6-packs (FC 1019 for DR)
- FC 1015 will represent the quantity of IBM zAware CP/IFL 10-packs (FC 1020 for DR)

Features for ordering IBM zAware on the z13 will be located on the Processor usage panel of eConfig

- FC 1010 will represent the quantity of IBM zAware CP/IFL 10-packs
- FC 1011 will represent the quantity of IBM zAware CP/IFL 10-packs DR server

zEDC Express and Compression Acceleration

What is IBM Enterprise Data Compression (zEDC)?

A capability of z/OS V2.1 (and beyond), IBM zEnterprise Data Compression (zEDC) combined with the hardware feature called zEDC Express, offers a compression acceleration solution designed for high performance, industry standard, low latency compression with little additional overhead.

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

zEDC and CPU compression solve two different requirements.

The CPU compression in every z server using the CMPSC instruction is much faster than software compression and is optimized for short DB2 rows.

zEDC uses an industry standard compression format optimized for compression of large sequential data. It is best for sequential dataset where it can store the same data in half the space as CPU compression and at a lower CPU cost.

zEDC is DEFLATE RFC compliant and is fully compatible with the open source zlib library. This allows zEDC to accelerate middleware that handles standard compression including IBM Content Manager OnDemand, IBM WebSphere MQ channel compression, Java stream classes and managed file transfer software like the IBM Sterling Connect:Direct product. It also allows for easy adoption by Independent Software Vendors as well, for which many z Systems ISVs have enabled their products with zEDC support.

Systems deploying zEDC have access to both CPU compression and zEDC compression allowing use of the best compression technology for the job.

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.

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.

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

z/OS V2.1 and later z/OS releases provide the zlib library which supports the sending of compression and decompression requests to the zEDC Express. The z/OS-provided zlib library is 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.[18]

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 z Systems 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.

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

Exploitation includes:

- z/OS V2.1 and later
  - SMF logger
  - DFSMS BSAM/QSAM extended format data sets
  - IBM 31-bit and 64-bit SDK for z/OS Java Version 8, Java Version 7 Release 1 , IBM 31-bit and 64-bit SDK7 for z/OS SR7 and
  - DFSMSdss™ and DFSMShsm™
- 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)
- IHS Server 8.5.5.4 (PI24424 and later (APAR PI24423)
- IBM Content Manager OnDemand V9.5.0.3 (APAR PI41677)

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.

How does zEDC improve SMF logstream recording?

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

---

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

CPU usage by up to 30%.[19] Corresponding support in the SMF dump program IFASMFDL 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.

What support is there in tools to get information on hardware compression for z Systems?

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.

How does QSAM/BSAM benefit from zEDC?

QSAM/BSAM can save disk cost and in some cases shorten elapsed time, reducing batch windows, reduce replication costs compared to uncompress or CPU compressed data, and reduce CPU cost compared to CPU based compression.

What support is there in tools to get information on hardware compression for z Systems?

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

How does IBM Java for z/OS exploit zEDC?

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. This support is also available with IBM 31-bit and 64-bit SDK7 for z/OS Java Technology Edition, Version 8 (5655-DGG and 5655-DGH)

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] 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.[20]

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

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

---

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

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. [19]

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

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

The IBM Sterling Connect:Direct for z/OS V5.2 uses zEDC to reduce file transfer times. IBM measurements show IBM Sterling Connect:Direct for z/OS V5.2 using zEDC 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. Connect:Direct using zEDC for compression is fully compatible with existing Connect:Direct clients that support software compression.

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

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. QSAM/BSAM using zEDC can reduce the output size of Access Monitor and UNLOAD files by up to 10X and CKFREEZE files reduced by up to 4X. [19]

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

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.

What happens if data that is already compressed gets compressed again by zEDC?

Compressing data that is already compressed data is common. For example, when you send a file compressed by PKZIP using Connect:Direct over a compressed connection you are compressing data that is already compressed. In some cases zEDC can make compressed data even smaller. In other cases compressing already compressed data results in no size change or a slight increase in data size. In all cases compressing already compressed data worked correctly and should not be a concern.

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

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 will make the resulting SMF data smaller even if it was previously compressed with the DB2 ZPARM option. zEDC can be used in coordination with the DB2 support or it can replace the use of ZPARM.

What are the software and hardware requirements for zEDC Express?

Software requirements:

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

Hardware requirements:

- z13s, z13, zEC12 or zBC12
- 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 z13s, z13, zEC12, zBC12
- The recommended high availability configuration per server is a minimum of 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 z13s, z13, zEC12 or zBC12
  - Apply z/OS Service if necessary; 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.

How is the zEDC Express feature defined to the z13s or z13?

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.

For z/OS to use zEDC I need to order two features both hardware and software? That is something new.

To use zEDC, you need to order both the hardware zEDC Express feature and the z/OS priced feature. When an order is placed, eConfig provide a notification that there is a required software feature for z/OS at the time of configuration of the zEDC Express feature for PCIe I/O drawer (FC#0420).

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

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#0420). 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.

Can the zEDC Express feature be shared between multiple LPARs?

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

What are the considerations for data sharing and disaster recovery if all my systems are not using zEDC, such as if I have z/OS V1.13 systems or z196/z114 systems?

It is highly recommended that your z/OS is configured for zEDC accelerated compression on each z/OS LPAR that may use interact with zEDC compressed data. This includes disaster recovery sites which may need to assume responsibly for reading compressed data sets. If zEDC hardware acceleration is not available small amounts of data may be transparently decompressed by z/OS software however it is likely that the volume of zEDC compressed data processed during normal batch processing is not sustainable without zEDC hardware acceleration.

If you need to access QSAM/BSAM data sets or DFDSS/DFHSM data on 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 other applications that leverage the industry standard zlib library, servers without zEDC will be fully compatible but will use the CPU for compression and decompression tasks.

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

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.

How do I configure my z/OS system for the zEDC compression feature?

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.

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

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

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

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.

What planning tools are there to help determine if I will see benefit from using zEDC?

The IBM Washington Systems Center (WSC) has a created a sizing tool for use with zEDC, the IBM z Systems 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

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

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.

If a dump is captured to DASD and is compressed using the zEDC Express, can it be FTP'ed compressed or is it decompressed before the FTP?

The data on a compressed BSAM dataset is decompressed as it is read, so you would not FTP compressed data by pointing FTP to a compressed BSAM dataset. Instead you would use a zEDC enabled tool like PKZIP or 'Encryption Facility for z/OS' to create a zip file (RFC 4880 compliant

package) then FTP the zip file. All servers support opening zip files whether they have zEDC or not. zEDC just makes it run faster. People using the dump transfer tool would not typically zip the file before calling the tool but they could.

If a compressed dataset can be FTP'ed, and it is FTP'ed to a z System that does not support that level of compression, is there a programmatic way to decompress it?

The data on a compressed BSAM dataset is decompressed as it is read, so you would not FTP compressed data by pointing to a compressed BSAM dataset. Instead you would use a zEDC enabled tool like PKZIP or 'Encryption Facility for z/OS' to create a zip file (RFC 4880 compliant package) then FTP the zip file. All servers support opening zip files whether they have zEDC or not. zEDC just makes it run faster.

Does z/VM support zEDC?

z/VM guest exploitation support for the zEDC Express feature is available on the z13s, z13, zEC12 and zBC12 servers, with up to 15 z/VM guests sharing the zEDC Express feature. The PTFs for APARs OA44482 and OA43256 are required when exploiting zEDC as a z/VM guest.

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

z/VM does not directly benefit from zEDC. The z/VM support gives z/OS guest operating systems access to zEDC acceleration.

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.

Is the zEDC Express feature supported on Linux for z Systems?

zEDC Express support is included in SLES12 SP1 and in RHEL 7.2 as "Tech-Preview" line item. The outlook for lifting this Tech-Preview Flag is during the next weeks, sometime in 1Q2016 via a fix from the distributors.

Is the zEDC Express feature supported by KVM?

Not at this time.

IBM z Systems Specialty Engine Update

What was announced on with IBM z13s and IBM z13 concerning simultaneous multithreading (SMT)?

With the IBM z13s and z13, simultaneous multithreading (SMT) support was announced for zIIP processors. This expected to offer throughput improvements you can use to address the growing volume of zIIP-eligible work, such as Java®-based IBM WebSphere® Application Server and CICS® Transaction Server Java-based transactions, in addition to XML parsing.

Additionally, customers should plan on running zAAPs on zIIPs, as zAAPs are no longer orderable on the z13s or z13.

Beside zIIP engines, there are the Integrated Facility for Linux (IFL) and Internal Coupling Facility (ICF). IFLs are used to run Linux; only Linux for System z and z/VM virtualization technology and designated software products can run on IFLs. Running z/OS® and a high virtualized Linux environment on IBM z Systems™ servers provides a great opportunity *to* integrate consolidated servers with the strength of z/OS applications and data.

An ICF is used to run the Coupling Facility Control Code for Parallel Sysplex® environments; only Coupling Facility Control Code can run on ICF engines.

What is the Integrated Facility for Linux (IFL)?

The Integrated Facility for Linux (IFL) is a processor dedicated to running Linux workloads on IBM z Systems servers. It is an optional feature, designed, attractively priced and supported by z/VM virtualization and IBM Wave for z/VM virtualization management, KVM for IBM z Systems and the Linux operating system.

An IFL has the same functionality as a general purpose processor (CP), but with z13s and z13, the IFL provides an optional multithreading technology, designed to improve the performance capacity.

Since an IFL can run many virtual servers, a consolidation onto an IFL can result in less IT costs, especially when the software licensing costs are priced per core.

The inclusion of an IFL will not increase charges for z Systems software running on general purpose (standard) processors.

Can I buy a z System that is made up entirely of Integrated Facility for Linux (IFL) engines and has no CPs installed?

Yes. It is possible to configure a server with IFL engines only. You could also purchase a IBM LinuxONE Rockhopper or LinuxONE Emperor that ONLY has IFLs installed. For more information visit our website: www.ibm.com/LinuxONE .

What types and portions of workloads are eligible/authorized to execute on IFLs?

IBM authorizes customers to use IBM specialty engines (e.g. zIIPs and IFLs) only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines". The table can be found at: www.ibm.com/systems/support/machine_warranties/machine_code/aut.html.

Where can I get more information on the IFL?

For more information, attributes and values, please look at the Web site:
ibm.com/systems/z/os/linux/solutions/ifl.html

What is the "zAAP on zIIP" capability?

In 2009, IBM System z specialty engine related announcements included the following:

z/OS offers a capability that can enable IBM zEnterprise Application Assist Processor (zAAP) eligible workloads to run on IBM 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 called "zAAP on zIIP" and can help with migration to zIIPs.

In z/OS V1.11 and later IBM zEnterprise Application Assist Processor (zAAP) eligible workloads can run on IBM z Integrated Information Processors (zIIPs). This function can enable you to run zIIP- and zAAP-eligible workloads on the zIIP. This is ideal for customers without enough zAAP- or zIIP-eligible workload to justify a specialty engine; 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 more cost effective.

Second, some customers have investments in zIIP processors, but have very little zAAP eligible work to justify the addition of a zAAP. With the "zAAP on zIIP" capability, the addition of eligible TCB to the eligible enclave SRB workloads can maximize the investment of the zIIP.

Customers having only zIIP processors can now have additional value with 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.

As the z13s and z13 no longer support zAAPS, zAAP on zIIP can be used to help you migrate your zAAP-eligible workload on zIIPs.

What are the prerequisites for "zAAP on zIIP"?

The "zAAP on zIIP" capability is available with z/OS V1.11 and later and on all IBM System z9® and System z10 servers and later; some additional restrictions apply. If you wish to disable the function for any reason, you must IPL with ZAAPZIIP=NO in the IEASYSxx Parmlib member.

zAAP on zIIP is available at no additional charge.

What is the ratio in place for zIIP specialty engine terms and conditions for z13s and z13 servers?

You must purchase and maintain one or more general purpose processors for every two zIIP processors on the server. The ratio is 2:1 for zIIPs to CPs.

I am running on an older z Systems server that supports zAAP specialty engines, after my zAAP workloads are consolidated onto zIIPs my 1:1 zIIP to CP ratio is exceeded; what should I do?

Customers who are migrating have either met the required 1:1 ratio for zAAPs and/or the 1:1 ratio for zIIPs in the past. Customers will not be considered in violation of the 1:1 rule once they migrate. Customers who migrate zAAPs onto zIIPs, and find they now need more zIIPs than CPs will not be considered in violation of the 1:1 rule if they met the 1:1 requirements in the past and consolidation forces this exception. If you have further questions speak to your client rep.

How will consolidation of zAAP workloads onto zIIPs help customers?

Customers using zIIPs will find it simpler from a management and capacity planning perspective to plan for only one type of specialty engine. With one resource pool of zIIPs, customers need not plan for two specialty engine types. In addition, if there is "white space" left on the zIIPs, a consolidation onto zIIPs makes sense as the zIIP engines can be more efficiently utilized as opposed to having both zAAP and zIIP engines running at less than full capacity.

Are there changes to the workloads eligible to run on zAAPs and zIIPs?

The workload eligibility rules have not changed. Please see the authorized use table at http://www-947.ibm.com/systems/support/machine_warranties/machine_code/aut.html

What is CIM? How is CIM used on z/OS?

The Common Information Model (CIM) is a standard maintained by Distributed Management Task Force for defining and sharing information between devices and programs in a standard way. The CIM specification provides for a framework within which CIM client applications can request CIM servers to return information about system resources.

On z/OS, there is a CIM client programming interface that enables z/OS-based Java applications to request information on or perform management actions to z/OS system resources which are managed by CIM servers. With support for the CIM server on z/OS, applications can gain access to z/OS resources using an extensible, industry-standard model.

On z/OS information providers include components such as RMF™, WLM, DFSMS™, and BCP; systems management exploiters include z/OS Management Facility.

What CIM workloads are eligible for zIIP specialty engines?

z/OS CIM server processing is authorized to run on the zIIP. Specifically, the CIM server and the CIM provider workloads are eligible for zIIP. CIM providers are dynamically loaded into the CIM Server address space and become part of the CIM Server processing.

CIM client applications can be Java-based and are enabled to execute on the zAAP

Portions of the z/OS Management Facility are Java-based CIM client applications which access the z/OS CIM server; parts of these workloads are eligible for zAAP and zIIP, respectively.

z/OS CIM server eligibility for the zIIP requires z/OS V1.11 or later.

Please provide more detail on DB2® utilities use of zIIP.

As a continuation of DB2 exploitation of System z specialty engines, z/OS V1.11 and later, together with DB2 for z/OS Version 8 or DB2 9 DB2 utilities and later, offers additional capabilities for exploiting zIIPs. DB2 utilities sorting fixed-length records using IBM's memory object sorting technique can have a portion of the workload redirected to a zIIP when one is available.

DB2 for z/OS and DB2 Utilities for both DB2 Version 8 and DB2 9 and later take advantage of zIIP processors. SQL statements accessing DB2 for z/OS via DRDA® over a TCP/IP connection, DB2 utilities index maintenance, and complex star schema parallel queries in BI applications are all eligible for zIIP processing.

In addition, DB2 for z/OS utility functions used to maintain index maintenance structures are eligible for zIIP processing such as the BUILD portion of LOAD, REORG, and REBUILD functions. Portions of DB2 Utility sorts of fixed-length records in the memory object sorting path are eligible for zIIP.

Where can I get more information on DB2 for z/OS utility-driven sort zIIP enablement?

For more information, see Utilities Management Solution Page:
ibm.com/software/data/db2imstools/solutions/utilities-mgmt.html

See the Utilities Page:
ibm.com/software/data/db2imstools/db2tools/db2utilsuite/

How can I perform capacity planning to better understand zIIP capacity?

The zPCR (Processor Capacity Reference for IBM System z) tool is updated to reflect this function as well as the configuration rules concerning the zIIP capabilities. Support for projecting zIIP capacity, when logical CPs are associated with one or more LPARs continues unchanged.

zPCR is a Microsoft® Windows®-based productivity tool, designed to provide capacity planning insight for IBM System z processors running various workload environments under z/OS, z/VM®, and Linux® for System z. Capacity results are based on IBM's LSPR (Large Systems Performance Reference) data.

zPCR functions include:

- LSPR Processor Capacity Tables: Displays processor capacity ratios for up to 10 z/OS workload environments as well as z/VM and Linux.
- LPAR Configuration Capacity Planning: Any legitimate LPAR configuration can be defined. Capacity projections are generated for each partition as well as the LPAR host as a whole. The LPAR host processor can be configured with General purpose CPs, zIIPs, IFLs, and ICFs where valid. LPARs can be defined by type (General Purpose, IFL, or ICF), operating system/workload, configuration (dedicated/shared with number of logical CPs), and weight/capping assignments.

Note that #2 is the only way that zIIP capacity projections are provided.

How will PROJECTCPU measure eligible workloads if "zAAP on zIIP" capability is engaged?

Once the "zAAP on zIIP" capability is engaged (i.e. ZAAPZIIP=YES, zIIPs present, and no zAAPs present), the PROJECTCPU characteristic for zAAP will go to '0.00' as all the workload that was once zAAP-eligible will then be zIIP-eligible and measured under the zIIP characteristic.

PROJECTCPU will not distinguish what was once zAAP-eligible work – from the standpoint of PROJECTCPU, it will all be zIIP-eligible work.

Where can I get more information on the zPCR (Processor Capacity Reference for z) tool?

https://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381

Where can I find more information?

The zAAP on zIIP capability and the ZAAPZIIP (and ZZ) system parameter is described in z/OS V1R11.0 MVS™ Initialization and Tuning Reference (SA22-7592-18).

Parallel Sysplex

What is Asynchronous Coupling Facility (CF) Duplexing for Lock Structures?

IBM announced support for a new synchronous form of CF duplexing for lock structures - Asynchronous CF Duplexing for Lock Structures on October 4, 2016.. Additional information can be found in the announcement letter: 116-093.

What performance issues are addressed by using Coupling Facility (CF) to replicate my data if the distance is very far?

Yes. The new capability helps with reducing the overhead of duplexing at any distance.

Do I need anything new on my storage devices to implement CF Duplexing?

No.

Is CF Duplexing available for z/OS under VM?

Yes, z/VM V6.4 with PTFs can be used for z/OS exploitation of the guest coupling environment.

What exploiters can take advantage of CF Duplexing?

DB2® V12 with PTFs for APAR PI66689 and the DB2 internal resource lock manager (IRLM) 2.3 with PTFs for APAR PI68378.

How do I get the new code for CF Duplexing?

It requires CFCC Level 21 with service level 02.16, z/OS V2.2 SPE with PTFs for APARs OA47796 and OA49148, CF to CF connectivity via coupling links, and exploitation. Asynchronous CF Duplexing for Lock Structures is designed to be a general-purpose interface for any transactional CF lock structure exploiters.

Any additional information out there on Asynchronous Coupling Facility for Duplexing for Lock Structures?

Yes. There is a good blog written by Dave Surman, and IBM Distinguished Engineer specializing in Parallel Sysplex.
http://mainframeinsights.com/asynchronous-coupling-facility-cf-duplexing-for-lock-structures/

What scalability benefits are available for Parallel Sysplex with the z13s and z13?

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 increased configuration limits to support larger data sharing environments. In the Coupling Facility, we have added an availability enhancement that is a scalability enabler for large cache structures.

The z13 and z13s will support up to 256 Coupling CHPIDs, twice the 128 coupling CHPIDs supported on zEC12 and zBC12. This provides enhanced connectivity and scalability for a growing number of coupling channel types and facilitates consolidation of multiple Sysplexes into the same set of physical servers. Note that each CF image will continue to support a maximum of 128 coupling CHPIDs.

On the z13, Up to 141 ICF engines can be ordered on a single server across multiple Coupling Facility LPARs. This helps environments that use a server hosting multiple Coupling Facilities to support multiple Parallel Sysplexes. There is still a limit of 16 ICF engines for a single Coupling Facility LPAR.

On the z13s, Up to 20 ICF engines can be ordered on a single server across multiple Coupling Facility LPARs. This helps environments that use a server hosting multiple Coupling Facilities to support multiple Parallel Sysplexes. There is still a limit of 16 ICF engines for a single Coupling Facility LPAR.

CFCC Level 21 and CFCC Level 20 support the Coupling Facility use of Large Memory to improve availability for larger CF cache structures and data sharing performance with larger DB2 Group Buffer Pools (GBP). This support removes inhibitors to using large CF cache structures, enabling use of Large Memory to appropriately scale to larger DB2 Local Buffer Pools (LBP) and Group Buffer Pools (GBP) in data sharing environments.[21]

What coupling link options do I have with the z13s and the z13?

The z13 can couple using InfiniBand or the new IBM Integrated Coupling Adapter (ICA SR) coupling links.

What is the IBM Integrated Coupling Adapter (ICA SR) fanout that is on z13s and z13?

The ICA SR, introduced on the z13, and now also available on the z13s, offers greater short reach coupling connectivity than existing link technologies and enables greater overall coupling connectivity per CPC footprint than prior server generations. The ICA SR is a 2 port fanout used for short distance coupling connectivity and utilizes a new coupling channel type: CS5. It utilizes PCIe Gen3 technology, with x16 lanes that are bifurcated into 2 x8 lanes for coupling. The ICA SR is designed to drive distances up to 150m and support a link data rate of 8 GigaBytes per second (GBps).[22]

The ICA SR fanout resides in the PCIe fanout slot on the z13s or z13 CPC drawer. The ICA SR can only be used for coupling connectivity between z13s and/or z13 enterprises. The ICA SR is designed to support up to 4 CHPIDs per port and 7 buffers (i.e. subchannels) per CHPID. The ICA SR is only compatible with another ICA SR and thus cannot be connected to HCA3-O or HCA3-O LR coupling fanouts. The ICA SR fanout requires new cabling. The z13 supports 10 PCIe fanout slots per CPC drawer which can be used for ICA SR coupling or connectivity to the PCIe I/O drawer in the I/O subsystem. The z13s supports 8 PCIe fanout slots per CPC drawer which can be used for ICA SR coupling or connectivity to the PCIe I/O drawer in the I/O subsystem.

Starting with z13s and z13, clients are strongly encouraged to begin migrating to the ICA SR to ensure short distance coupling connectivity to future server generations.

---

[21] To learn more about the performance benefits of large DB2 structures, reference "IBM zEnterprise System(R): Performance Report on Exploiting Large Memory for DB2 Buffer Pools with SAP(R)" at http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102461

[22] Note: A link data rate of 8 GBps does not represent the actual performance of the link.

Can you review the Infiniband options on the z13s and z13?

On the z13s and z13, there are two highspeed InfiniBand coupling links – up to 6 GBps for 12x InfiniBand and up to 5 Gbps for 1x InfiniBand.[23]

The 12x InfiniBand can be used for short distances - up to 150 meters (492 feet). 12x IFB links support up to 7 subchannels (devices) per CHPID. 12x InfiniBand coupling links utilize the Host Channel Adapter 3 optical (HCA3-O) fanout. The HCA3-O fanout has two ports/links and is compatible with the HCA2-O fanout on zEC12, zBC12, z196, or z114 machines. HCA2-O is not supported on z13s or z13.

1x InfiniBand are used for longer distances - up to 10 km 6.2 miles) unrepeated. 1x IFB links support up to 32 subchannels (devices) per CHPID. 1x InfiniBand coupling links utilize the Host Channel Adapter 3 optical long reach (HCA3-O LR) fanout. The HCA3-O LR fanout has four ports/links and is compatible with the HCA2-O LR fanout on zEC12, zBC12, z196, or z114 machines, which has two ports/links. HCA2-O LR is not supported on z13s or z13.

For z13, the maximum number of all HCA3 fanout features is limited to 16 per system. For z13s, the maximum number of all HCA3 fanout features is limited to 8 per system. Internal coupling links (ICs) can also be used for internal communication between Coupling Facilities (CFs) and z/OS images on the same server.

Does the z13s or z13 support ISC-3 coupling links?

No.

Where can I order the new cabling required for IBM Integrated Coupling Adapter (ICA SR)?

IBM strongly recommends clients order cabling for the new IBM Integrated Coupling Adapter (ICA SR) through Anixter or IBM Global Technology Services to get IBM qualified cables.
For more information, refer to *IBM z Systems Planning for Fiber Optic Links (FICON/FCP, Coupling Links, and Open System Adapters),* GA23-1407 which can be found in the Library section of Resource Link: http://www.ibm.com/servers/resourcelink/svc03100.nsf?OpenDatabase .

What is the value that Flash Express can play in supporting the Coupling Facility Control Code (CFCC)?

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[24], 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

---

[23] Note: The ICA SR and InfiniBand (PSIFB) 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. Systems Lab Services can assist your migration to ICA SR and PSIFB coupling links by providing services to assess the impact of the migration or to assist with the implementation of the migration.

[24] Note that the z13 will be released with CFC level 20

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.

What is the value of Coupling Thin Interrupts?

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.

Note: Coupling Think Interrupts was first released on the zEC12 GA2 and zBC12 and is available on CFCC LEVEL 19 or later. Note that z13 is at CFCC LEVEL 20. The z13s is a CFCC LEVEL xx.

What other generations of servers can connect to a z13s or z13 server in a Parallel Sysplex environment?

The z13s and z13 servers are designed to coexist in the same Parallel Sysplex environment with (n-2) server families. This allows it to coexist with other z13, z13s, zEC12, zBC12, as well as the z196 and z114 servers. Connectivity to the System z10 is not supported. You should also ensure that your zEC12, zBC12, z196, and z114 servers are at the latest driver level.

Are there resources to help with a migration over to PCIe coupling links?

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

STP

How is the z13s and z13 time synchronized with other z Systems servers?

The z13s and z13 do not support attachment to the Sysplex Timer®. However, the z13s or z13 can use Server Time Protocol (STP). 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.

How can the z13s or z13 be time synchronized to an existing External Time Reference (ETR) network of System z10 servers?

Neither the z13s or z13 can be time synchronized to z10 servers nor can it be time synchronized in a mixed Coordinated Timing Network (CTN). So, first, migrate the existing ETR network or mixed CTN to an STP only Coordinated Timing Network (CTN). Once this is done, a z13s or z13 can join the STP only CTN.

Can I enter an ETR ID on the Server Time Protocol (STP) panels such that a z13s or z13 machine can participate in a Mixed Coordinated Timing Network (CTN)?

No, neither a z13s nor z13 can join a mixed CTN.

Can a z13s or z13 be configured in a mixed CTN or STP-only CTN with IBM System z10 or older servers?

No. The z13s and z13 can only be configured in an STP-only CTN. The z13s and z13 servers are designed to coexist in the same CTN with (n-2) server families. This allows the z13s or z13 to participate in the same CTN with z196 or z114 servers, but not with System z10 or older servers.

Can I use the same External Time Source options for STP on the z13s or z13?

Yes. STP on the z13s and z13 still offers two options for external Time Source - Network Time Protocol (NTP) server and NTP server with Pulse per second (PPS). The IBM z13s and z13 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.

Security

What is the security technology rating for the z13s and z13?

IBM Processor Resource/Systems Manager (PR/SM) on the z13s and z13, is designed to be certified 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 z Systems. This helps to assure that many disparate applications running in different operating environments in separate logical partitions on the z13s or on the z13 will be isolated and distinct from each other.

What is new with the CP Assist for Cryptographic Function (CPACF)?

The CP Assist for Cryptographic Function (*CPACF*) has been redesigned for performance in the z13s and z13. Encryption algorithms such as TDES and AES perform up to 2 times and hashing functions such as SHA384 and SHA512 perform up to 3.5 times faster than the prior generation machine.

What can you tell me about the new Crypto Express5S?

The new Crypto Express5S feature is optimized for performance. Similar to the Crypto Express4S feature, it is installed in the PCIe I/O drawer configuration and offers a state of the art tamper resistant packaging. The new feature delivers greater performance in all three modes of operation - accelerator (SSL), secure CCA (Common Cryptographic Architecture) and secure PKCS#11. These enhancements will allow more data to be securely transferred across the internet. Updates to the z13s and z13 will allow the Crypto Express5S to be shared across more than 16 domains (up to the maximum number of LPARs on the system).

Will the Crypto Express4S feature be supported on the z13s or z13?

No, clients will need to use the new Crypto Express5S with their z13s or z13.

What is the new Common Cryptographic Architecture (CCA) enhancement Visa Data Secure Platform (DSP) - Point to Point Encryption (P2PE) announced with the z13s and z13?

With the z13, IBM introduced HSM support for Visa Data Secure Platform P2PE. Visa Data Secure Platform is a suite of products designed to help merchants prevent sensitive account holder data from being compromised. The z13 support included CCA-based callable services which support Visa's Standard Encryption method and the Visa Format Preserving Encryption method. This support relies on the Crypto Express5S coprocessor. The IBM z13 and IBM z13s will extend the Visa DSP support by adding a callable service that translates PIN blocks whose primary account number information has been encrypted using Visa DSP P2PE with static TDES or DUKPT keys. Support is provided for ISO-0 and ISO-3 PIN blocks.

What is the new Common Cryptographic Architecture (CCA) enhancement Secure AES GCM encryption mode announced with the z13s and z13?

Galois Counter Mode (GCM) is an authenticated encryption algorithm designed to provide both data authenticity (integrity) and confidentiality in a single operation. The algorithm can also serve as

a stand-alone MAC function. With the Crypto Express5S and IBM z13 or IBM z13s, CCA adds this encryption mode support to the existing Symmetric Algorithm Encipher (CSNBSAE) and Symmetric Algorithm Decipher (CSNBSAD) callable services. Previously AES GCM was supported with clear keys through ICSF. This feature will provide a secure key option.

What is the new Common Cryptographic Architecture (CCA) enhancement Interoperable ECC key derivation algorithm announced with the z13s and z13?

The EC Diffie-Hellman callable service is used to establish symmetric keys using a pair of ECC keys using the Elliptic Curve Diffie-Hellman (ECDH) protocol. Crypto Express5S coprocessor has been enhanced to add an additional key derivation scheme that allows the establishment of keys with non-CCA entities, without compromising security. The EC_Diffie-Hellman callable service creates a shared symmetric key with a pair of ECC (Elliptic Curve Cryptography) keys using the ECDH (Elliptic Curve Diffie-Hellman) protocol and the ANSX9.63 protocol static unified model key-agreement scheme.

What is the new Common Cryptographic Architecture (CCA) enhancement Addition of new Key Check Value algorithm announced with the z13s and z13?

The Key Test2 callable service generates or verifies the value of a clear or encrypted key or key part contained in an external or internal variable-length symmetric key-token, a DES key token or a key wrapped in an external TR-31 key block. With the Crypto Express5S, CCA supports a new CMAC-based Key Check Value algorithm for use with AES and TDES keys.

The following CCA enhancements were introduced with the IBM z13 and are now available on the IBM z13s.

What is the new Common Cryptographic Architecture (CCA) enhancement Visa Data Secure Platform Support with Point to Point Encryption including Visa format preserving encryption (VFPE) announced with the z13s and z13?

The z13s and z13 offers Visa format preserving encryption (VFPE) for protecting payment card data and can help provide additional security by enabling legacy databases and applications to contain encrypted data of sensitive fields without having to undertake a restructure of the database or applications. VFPE is a valuable tool for payment card applications that helps to maintain the character length between input clear text and resulting cipher text.

What is the new Common Cryptographic Architecture (CCA) enhancement Greater than 16 Domain support announced with the z13s and z13?

Greater than 16 Domain support will allow a cryptographic coprocessor to be shared across more than 16 domains, up to the maximum number of LPARs on the system. This support relies on enhanced firmware available with a minimum microcode level for the Crypto Express5S coprocessors. With the adjunct processor (AP) extended addressing (APXA) facility installed, the z Systems crypto architecture can support greater than 16 domains in an AP. Customers will have the flexibility of mapping individual LPARs to unique crypto domains or continuing to share crypto domains across LPARs.

What new capability was announced for Elliptic Curve Cryptography (ECC)[25] for z13s and z13?

The Elliptic curve based algorithms are gaining popularity, and the pace of adoption is accelerating. IBM z Systems first started supporting it on the zEC12/zBC12. With the z13s and z13 we enhanced public key support for constrained digital environments using Elliptic Curve Cryptography (ECC) by providing HW based ECC support via the CryptoExpress5S to improve performance. Some examples of the use of ECC include: the US government uses it to protect internal communications, it is the mechanism used to prove ownership of bitcoins, it provides signatures in Apple's iMessage service, it is used to encrypt DNS information with DNSCurve, and it is the preferred method for authentication for secure Web browsing over SSL/TLS. Chrome or Firefox are using it in establishing a secure connection.

What is the domain cloning enhancement for Trusted Key Entry (TKE) 8.1?

Domain cloning is planned to support a Configuration Migration Tasks application that gives administrators the ability to collect settings from one domain and apply the settings to any set of new domains. This feature significantly simplifies the process and reduces the time it takes to deploy new domains on a Host Crypto Module.

What is the launch coordinated master key role enhancement for Trusted Key Entry (TKE) 8.1?

Launch coordinated master key role from the TKE is planned to provide the ability to set master keys from the TKE in a way that the Key Data Sets (KDSs) are encrypted under the new master key. IBM plans to provide the ability to invoke the ICSF coordinated master key role function and now allow clients to do all master key management from the TKE.

What are the guided create features for roles and authority indexes enhancements for Trusted Key Entry (TKE) 8.1?

Guided create features are intended to provide new guided create functions for both Roles and Authorities tabs for Host Cryptographic Modules. This support steps the client through the process of creating these items and greatly clarifies and simplifies the creation process.

What are the two Certificate Authority wizard enhancements for Trusted Key Entry (TKE) 8.1?

The two certificate authority wizards are intended to take administrators through the process of creating all the smart cards used in a TKE zone and also take administrators through the process of creating all the smart cards needed for the Configuration Migrations Task wizards.

What are the display crypto module setting enhancement for Trusted Key Entry (TKE) 8.1?

The display crypto module setting enhancement is, a new TKE feature, that enables you to create a report that shows the current configuration of a Host Cryptographic Module.

---

[25] 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.

What is the HMAC key enhancement for Trusted Key Entry (TKE) 8.1?

The HMC key will enable TKE to support generation and loading of HMAC operational keys. This support will be limited to keys that are 16, 24, and 32 bytes in length.

What is the Save/Restore Customized Data enhancement for Trusted Key Entry (TKE) 8.1?

The Save/Restore Customized data feature will provide TKE with a new Privileged Mode Access ADMIN task that allows the client to save and restore client data on the TKE. This feature can be used to move client data from one TKE workstation to another or to restore client data on a TKE.

What is the Password Protect Console enhancement for Trusted Key Entry (TKE) 8.1?

Password protect console will support the ability to password protect the TKE console. When you are signed on in Privileged Mode Access with the ADMIN or SERVICE user ID you can disable automatic logon for the default console user and specify the required password. This password will then be required when the TKE console first starts, or when the launch application link is clicked on the pre-login page.

What is the Binary Key Part File utility enhancement for Trusted Key Entry (TKE) 8.1?

The Binary Key Part File utility will provide the ability to copy key material from a binary file and place it on a smart card. This provides a simplified and faster process compared to using secure key load from the smart card PIN pad.

What is the ACP Usage Information enhancement for Trusted Key Entry (TKE) 8.1?

The ACP Usage utility allows you to track which Domain Controls (Access Control Points) were actually "checked" within a domain. If tracking is activated and an authority check is done for an ACP, the "checked" indicator is turned on. This indicator stays on until the tracking data is cleared. The TKE is used to activate, deactivate, and clear the tracking data. In addition, TKE provides a feature for displaying the tracking information for each domain.

What is the Require Enhanced Host Password protection enhancement for Trusted Key Entry (TKE) 8.1?

Require enhanced host password protection will support the strongest possible encryption. When using HCR 77B0 or later with TKE 8.0 or later, the strongest possible encryption is used to protect the host password. The TKE uses the best encryption available for protecting the host sign-on password. However, with a new TKE 8.1 enhancement, you can configure the TKE so that it will allow a sign-on only to a system that uses the strongest encryption. IBM recommends that you install HCR 77B0 or later and configure the TKE to allow the sign-on only to systems that support the strongest host password protection.

What is Operational Key Option on Domain Groups enhancement for Trusted Key Entry (TKE) 8.1?

Operational key option on domain groups lets users specify where key commands are sent. Currently, when an operation key command is run from inside a domain group, the command is sent only to the master domain. IBM plans to add support to the domain group to add an attribute that indicates if operational key commands are to be sent to the master domain or all the domains in the group. The new group attribute value can be managed from the Create or Change Group features.

Is the TKE Server available as rack mounted?

Yes, with the z13s, the Trusted Key Entry (TKE) introduces a new 1U workstation that is a rack-mounted system. It is a combination of hardware and software, network-connected to one or more servers, and designed to provide a secure, flexible method for master and operational key entry as well as local and remote management of the cryptographic coprocessor features (Crypto Express5S when defined as a coprocessor). It provides the same functionality as the TKE workstation but as a rack-mounted server. If Trusted Key Entry 1U workstation is required on an IBM z13 or IBM z13s, then a TKE workstation with TKE 8.1 LIC (#0878) must be used. TKE workstations with TKE 8.1 LIC can also be used to control zEC12, zBC12, z196, and z114 servers.

Where can I get more information on IBM's encryption offerings?

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.

What is the IBM Enterprise PKCS #11 coprocessor?

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.

The TKE workstation is required to manage a Crypto Express5S or 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).

z/OS

Where can I find a full set of z/OS frequently asked questions?

Please visit:

www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ST&infotype=SA&htmlfid=ZSQ03081USEN&attachment=ZSQ03081USEN.PDF

z/VSE

I am a z/VSE user. What's in the z13s announcement for me?

For all z/VSE users:

- Latest hardware technology
- Higher single engine (PU) performance for z/VSE workloads, benefitting batch and online.
- All Linux on z Systems benefits if you run Linux together with z/VSE

What are the z/VSE software (pricing) savings using a z13s or z13?

Customers may choose Advanced Entry Workload License Changes (AEWLC), subject to all applicable terms and conditions.

See http://ibm.com/systems/z/resources/swprice/

Migration Pricing Option (MPO) is now available and allows for a longer transition period for migration to the new z/VSE V6.1. See z/VSE V6.1 Announcement for details: http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&appname=iSource&supplier=897&letternum=ENUS215-202

Does z/VSE support the Crypto Express5S adapter?

Yes:

z/VSE V6.1, V5.2, V5.1 with PTFs provide support for the Crypto Express5S adapter

Can z/VSE use Crypto Express5S support of greater than 16 domains on the new z13s and the z13?

Yes:

z/VSE V6.1, V5.2, V5.1 with PTFs

Does z/VSE support the FICON Express16S adapter?

Yes

Where can I get the latest and more information on z/VSE?

The most current information about z/VSE is available at: ibm.com/zvse

Also you can follow the z/VSE Twitter account at: twitter.com/IBMzVSE

Who can I contact if I need more information or have additional questions?

Send an e-mail to zvse@de.ibm.com

z/VM

Where can I find a full set of z/VM frequently asked questions?

Please visit:
http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ST&infotype=PM&appname=STGE_ZS_ZS_USEN&htmlfid=ZSQ03027USEN&attachment=ZSQ03027USEN.PDF

KVM

Where can I find a full set of KVM on the mainframe frequently asked questions?

Please visit:
http://w3-01.ibm.com/sales/ssi/cgi-bin/ssialias?subtype=ST&infotype=SA&htmlfid=ZSQ03090USEN&attachment=ZSQ03090USEN.PDF

Linux on z Systems

What are the unique advantages of running Linux on IBM z Systems?

Linux on z Systems can provide you an IT infrastructure that is freeing your business from IT complexity and improving the responsiveness of your systems and your people.

The Linux on z Systems 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.

Linux on IBM z workloads can be easily integrated with existing solutions running on z/OS and/or z/VSE—thus benefiting from co-location of data and applications, very fast internal network communication, and an all-encompassing high availability and disaster recovery solution.

Keeping it simple, “do more with less”.

What makes the Linux on z Systems an enterprise grade Linux solution?

While “Linux is Linux", which is true for the Linux kernel and other components, and which often allows to run Linux applications after a simple re-compile on another Linux platform, it’s the complete environment – the Linux platform – that provides the overall solution.

An enterprise grade Linux solution, in our understanding, has defined characteristics. It provides:

IT simplicity, allowing to run up to hundreds of different workloads in parallel on one server

Easy workload integration of new and existing data and applications

Flexible server provisioning, simple to manage

High productivity, based on efficient systems and life cycle management

Highest resource utilization levels

High levels of quality of service – security, availability, reliability

Linux on z Systems can offer these characteristics by exploiting the capabilities and benefits of the industry-leading IBM z Systems, the outstanding IBM z/VM virtualization and the empowered virtualization management of the IBM Wave for z/VM (IBM Wave), IBM zAware, Elastic Storage (based on IBM GPFS™) and IBM GDPS/PPRC resiliency technologies, to name a few facts.

What are the potential savings of when using Linux on IBM z Systems?

Linux on z Systems can help lower your IT costs in several areas. These areas are mainly based on the massive workload deployment capabilities and the high quality of service characteristics.

The potential areas of cost savings are:

Operational management – a single physical z Systems server can mean less effort for the systems and operational management. Think of all servers, cables, switches, routers that won’t be required with a single server environment. In addition, when using an intelligent virtualization management for the virtual servers, this can be very efficient, helping you on your productivity.

- Workload integration - integrating Linux workloads with solutions running on z/OS or z/VSE can benefit from co-location of data and applications, very fast internal network communication, resource sharing, and an all-encompassing high availability solution with unique arrangements for administration, security, backup and disaster recovery.

Maintenance – the maintenance costs of a single z Systems server can replace the maintenance costs of many x86 and UNIX® servers. But as mentioned above, it's not only the server maintenance, the maintenance of the other hardware components which aren't required in a single server environment, is omitted as well.

Software licensing – Linux on z Systems can support the consolidation of x86 and UNIX cores onto a single z Systems core (IFL processor). This can result in a dramatic cost reduction, since software licensing costs for Linux are usually priced per core.

Energy and space consumption – Linux on a z Systems server can allow for fewer servers and network components, requires less cooling, which help on savings in physical space and electricity consumption.

Business continuity – Linux z Systems provides trusted operations, it can scale up and out to meet spikes in server activity, helping to minimize costly transaction delays and potentially devastating system crashes. IBM z Systems are the most secure commercial servers available, have powerful encryption, protected your data 24/7. The suite of the z Systems built-in features, e.g. IBM zAware, can rapidly respond to, or even anticipate, threats to system health, helping to prevent costly system downtime. As well, high availability and disaster recovery solutions are available, such as IBM Spectrum Scale™ (based on GPFS technology) or the IBM GDPS high availability and disaster recovery capabilities.

What are the new advantages of Linux running on an IBM z13 (z13)?

To develop an *enterprise grade Linux* platform was a key design point for the z13s and z13. Existing server strengths have been further enhanced and new capabilities have been added:

New consolidation savings with increased total capacity in the same footprint more, higher capacity processors, more memory, more logical partition (LPAR) support, and improved sharing capabilities for networking and cryptographic features.

Better availability and more efficient use of critical data with the greatly expanded available redundant array of independent real memory.

Economies of scale with simultaneous multi-threading (SMT) processor design delivering more throughput for Linux workloads.

Improved performance of complex mathematical models, perfect for analytics processing, with Single Instruction Multiple Data (SIMD).

Improved ability to meet Service Level Agreements with new processor chip technology re-designed, larger cache and enhanced accelerators for cryptography.

Stronger and faster protection and integrity of data with the new Crypto Express5S cryptographic adapter.

Faster insight into the health of the Linux system with IBM zAware, helping you quickly pinpoint problems and therefore avoiding service disruptions.

Where can I get information on clients running Linux on z Systems in production?

Client success stories, sorted by industry, can be found at: ibm.com/systems/z/linux/stories.html

What is Linux on z Systems?

Linux on z Systems is the synonym for Linux that has been ported to the IBM z Systems architectures listed below. Linux is Linux, and Linux for z Systems has the identical look and feel as Linux on other platforms.

In addition, Linux for z Systems is supporting devices that are specific to the z Systems environment it is optimized to exploit the z Systems architecture for scalability, performance, security and resiliency.

Linux has been ported to the following IBM z Systems:

- IBM z13s (z13s) and IBM z13 (z13)
- IBM zEnterprise BC12 (zBC12) and IBM zEnterprise EC12 (zEC12)
- IBM zEnterprise 114 (z114) and IBM zEnterprise 196 (z196)
- IBM System z10 Enterprise Class (z10 EC) and BM 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 on z Systems inherits the advantages and can fully exploit the capabilities; that distinguishes Linux on z Systems from Linux on other platforms.

What are the IBM tested and supported Linux distributions for Linux on z Systems?

The latest information can be found on the tested platforms Web page (ibm.com/systems/z/os/linux/resources/testedplatforms.html).

Where can I get commercial Linux distributions for z Systems?

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). Bringing more flexibility and choice, Ubuntu will join Red Hat and SUSE as the premier Linux distributions on z Systems.

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 z Systems Linux production environments should be covered by a Linux support agreement and maintained with a current Linux subscription service.

What IBM software products are available for Linux on z Systems?

IBM offers a large software portfolio for "Red Hat Enterprise Linux" and "SUSE Enterprise Linux Server" distributions for z Systems, and new IBM software products are added constantly.

The information about IBM software available for these Linux distributions are available via the IBM Software Product Availability tool (hardware selection must be "z Systems" or "System z")

(http://publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/productsOnOs.html).

In addition, Linux on IBM z supports many non-IBM software solutions, including SAP and Oracle (IBM Global Solutions Directory: www-304.ibm.com/partnerworld/gsd/homepage, or check the distributor web pages), and there is a growing ecosystem of open source software available such as Docker, MariaDB, MongoDB, PostgreSQL, Spark, Apache Tomcat, or Node.js (ibm.com/systems/z/os/linux/open-source-software.html).

Where can I get the latest and more information on Linux on IBM z Systems?

The most current information about Linux on z Systems is available at: ibm.com/systems/z/linux

An additional information resource is the Blog: Linux and Mainframe (http://linuxmain.blogspot.de/), and a FAQ document for download exists as well: FAQ – Linux on z Systems

z13s - Software Pricing

What software pricing is announced with the z13s?

A new Technology Transition Offering (TTO) called Technology Update Pricing for the z13s is being announced along with revisions to the Technology Update Pricing for the z13 offering and two revised Transition Charges for Sysplexes or Multiplexes offerings.

Technology Update Pricing for the z13s uses the reporting mechanisms and existing Millions of Service Units per hour (MSU) tiers of the Advanced Entry Workload License Charges (AEWLC) pricing metric while extending the software price-performance provided by AEWLC.

Technology Update Pricing for the z13s 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 z13s server. It also applies to all these operating systems and their associated middleware programs when running in a z/VM guest environment.

The revisions to the Technology Update Pricing for the z13 offering and the two Transition Charges for Sysplexes or Multiplexes 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, or when clients have implemented Country Multiplex Pricing

What is the price reduction available with the Technology Update Pricing for z13s?

The percent reduction in the monthly AEWLC is based on the number of z13s 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 z13s is reduced by the percentage in the table below.

Schedule of AEWLC reductions for Technology Update Pricing for the z13s (TU4)

MSUs: Quantity of z13s Full Capacity MSUs for a stand-alone server	Reduction in Monthly AEWLC
1-10 MSUs	13.0%
11-17 MSUs	13.0%
18-30 MSUs	13.0%
31-45 MSUs	10.0%
46-87 MSUs	9.0%
88-175 MSUs	9.0%
176-260 MSUs	9.0%
261-315 MSUs	9.0%
316-390 MSUs	9.0%
391 - more MSUs	9.0%

The number of MSUs used to determine the MSUs of a stand-alone z13s 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

If I upgrade from a z114 to a z13s will software pricing improvements be additive?

You will receive the pricing advantage of Technology Update Pricing for the z13s (TU4) for the z13s. The Technology Update Pricing for the z13s (TU4) price reduction includes the benefits of the prior price reductions. Technology Update Pricing for AEWLC (TU2) is available when upgrading from a z114 to a zBC12.

Schedule of AEWLC reductions for Technology Update Pricing for the z13s (TU4)

MSUs: Quantity of z13s Full Capacity MSUs for a stand-alone server	Reduction in Monthly AEWLC for z13s (TU4)	Reduction in Monthly AEWLC for zBC12 (TU2)
1-10 MSUs	13.0%	5.0%
11-17 MSUs	13.0%	5.0%
18-30 MSUs	13.0%	5.0%
31-45 MSUs	10.0%	5.0%
46-87 MSUs	9.0%	4.0%
88-175 MSUs	9.0%	4.0%
176-260 MSUs	9.0%	4.0%
261-315 MSUs	9.0%	4.0%
316-390 MSUs	9.0%	4.0%
391 - more MSUs	9.0%	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

What software pricing announcements were made for a z13s in a Parallel Sysplex?

IBM announced revisions to the Technology Update Pricing for z13 offering and two revised Transition Charges for Sysplexes offerings.

The revisions to the Technology Update Pricing for z13 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, and priced with AWLC, or in a Multiplex priced with Country Multiplex License Charges (CMLC).

The Technology Update Pricing for z13 offers actively coupled Parallel Sysplexes and Loosely Coupled Complexes that consist entirely of z13s servers and z13 servers the benefit of aggregated Technology Update Pricing for z13 when priced with AWLC, or in a Multiplex priced with Country Multiplex License Charges (CMLC).

Transition Charges for Sysplexes or Multiplexes (TC2): When two or more servers consisting of one or more z13, z13s, zEC12, or zBC12 servers with one or more z196 or z114 servers comprise a

TTO-eligible environment, those servers receive a reduction to AWLC or CMLC pricing across the TTO-eligible environment. This reduction provides a portion of the price-performance benefit related to Technology Update Pricing for AWLC (TU1) based on the proportion of z13, z13s, zEC12, or zBC12 server capacity within the TTO-eligible environment.

Transition Charges for Sysplexes or Multiplexes (TC3): When two or more servers consisting of one or more z13 or z13s servers with one or more zEC12 or zBC12 servers comprise a TTO-eligible environment, those servers receive a reduction to AWLC or CMLC pricing across the TTO-eligible environment. This reduction provides a portion of the price-performance benefit related to Technology Update Pricing for the z13 (TU3) based on the amount of server capacity within the TTO-eligible environment.

If available for your specific configuration, you have a choice of selecting either Technology Update Pricing for the z13 or PSLC, but not both, for your IBM software programs within the z/OS or z/TPF operating system families on a z13s server.

Will I still be able to recognize the benefits of sub-capacity pricing with AEWLC on the z13s?

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 z13s server. All of the Technology Transition Offerings may be used in a sub-capacity environment.

What contracts are required for Technology Update Pricing for the z13s?

In order to get Sub-Capacity Technology Update Pricing for the z13s, the following contract is required to be executed:

ICA Attachment for IBM System z AEWLC, (Z125-8755), or, if applicable, ICA Attachment for IBM Country Multiplex Pricing, (Z126-6965).

The Supplement for Technology Transition Offerings to Attachment for IBM System z Advanced Workload License Charges, Attachment for IBM System z Advanced Entry Workload License Charges, Attachment for Country Multiplex Pricing, and to IBM System z Machines Exhibit (Z125-8994) describes the Technology Transition Offerings and should be delivered to each z13s 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 z13s customer, but they do not get executed.

What is the Processor Value Unit (PVU) rating used for Passport Advantage® software on the z13s?

The PVU rating for the z13s will be 100 for both IFL engines and CP engines, the same as the Rockhopper, zBC12, z114 and the z10 BC. The PVU rating for the z13, Emperor, zEC12, z196 and z10 EC remains 120 for both IFL engines and CP engines.

Where can I get more information on IBM z Systems software charges?

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

z13 Software Pricing

What software pricing is announced with the z13?

Three new or updated Technology Transition Offerings are being announced. They are Technology Update Pricing for the z13 and one new and one revised Transition Charges for Sysplexes programs.

Technology Update Pricing for the z13 (TU3) 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 z13 server or in an aggregated z/OS Parallel Sysplex or z/TPF Loosely Coupled Complex with a z13.

What software pricing is available on a standalone z13 or a z/OS Parallel Sysplex or a z/TPF Loosely Coupled Complex with all z13 technology?

For a standalone z13, or a Sysplex or a Complex with all z13 servers, IBM will offer the Technology Update Pricing for the z13 (TU3) which leverages the existing AWLC pricing metric while offering price performance improvements for z13. The Technology Update Pricing for z13 provides reduced AWLC pricing on the z13 as compared to AWLC on the zEC12 and the z196 with greater than 3 MSUs.

Note: If you choose to do so, you may continue to use PSLC pricing with the z13.

What is the price reduction available with the Technology Update Pricing for the z13 available for the z13?

The percent reduction in the monthly AWLC is based on the number of z13 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 z13 or Parallel Sysplex or Loosely Coupled Complex of all z13 servers is reduced by the percentage in the table below.

Schedule of AWLC reductions for Technology Update Pricing for the z13 (TU3)

MSUs: Quantity of z13 Full Capacity MSUs for a stand-alone server, or the sum of the Full Capacity MSUs in an actively coupled Parallel Sysplex or z/TPF Loosely Coupled Complexes made up entirely of z13 servers	Reduction in Monthly AWLC
4 – 45 MSUs	4.0%
46 - 315 MSUs	8.0%
316 – 1315 MSUs	9.0%
1316 – 2676 MSUs	10.0%
2677 – 5476 MSUs	12.0%
5477 or more MSUs	14.0%

The number of MSUs used to determine either the MSUs of a stand-alone z13 server or the total MSUs of a z13 Sysplex or Complex are based on the announced IBM full capacity ratings that can be found on the Mainframe Exhibits section of the z Systems Software Contracts website at http://ibm.com/systems/z/swprice/reference/exhibits/hardware.html

If I upgrade from a z196 to a z13 will software pricing improvements be additive?

You will receive the pricing advantage of Technology Update Pricing for the z13 (TU3) for the z13. The Technology Update Pricing for the z13 (TU3) price reduction includes the benefits of the prior price reductions. Technology Update Pricing for AWLC (TU2) is available when upgrading from a z196 to a zEC12.

MSUs: Quantity of Full Capacity MSUs for a stand-alone server, or the sum of the Full Capacity MSUs in an actively coupled Parallel Sysplex or z/TPF Loosely Coupled Complexes made up entirely of z13 or zEC12 servers	z13 - TU3 Reduction in Monthly AWLC	zEC12 – TU2 Reduction in Monthly AWLC
4 – 45 MSUs	4.0%	2.0%
46 - 315 MSUs	8.0%	4.0%
316 – 1315 MSUs	9.0%	4.5%
1316 – 2676 MSUs	10.0%	5.0%
2677 – 5476 MSUs	12.0%	6.0%
5477 or more MSUs	14.0%	7.0%

If I add a z13 into a Sysplex or Complex with one or more zEC12 and/or zBC12 servers, and but no z196 or z114 servers, can I benefit from the new Technology Update Pricing for the z13?

No. AWLC pricing may be used across a Sysplex or Complex when all the systems are z13, zEC12, or zBC12. The AWLC Transition Charges for Sysplexes (TC3) provides a reduction in AWLC pricing based the number of MSUs of z13, zEC12 and zBC12. When you migrate to all z13 servers, the Technology Update Pricing for the z13 (TU3) may apply.

MSUs: Quantity of z13, zEC12, and zBC12 Full Capacity MSUs in an actively coupled Parallel Sysplex or Loosely Coupled Complex	Reduction in Monthly AWLC
4 – 45 MSUs	2.8%
46 - 315 MSUs	5.6%
316 – 1315 MSUs	6.3%
1316 – 2676 MSUs	7.0%
2677 – 5476 MSUs	8.4%
5477 or more MSUs	9.8%

If I add a z13, zEC12 and/or zBC12 in a Sysplex or Complex with one or more z196 and/or z114 servers, can I benefit from the new Technology Update Pricing for the z13?

No. AWLC pricing may be used across a Sysplex or Complex when all the systems are z13, zEC12, zBC12, 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 z13 servers, the Technology Update Pricing for the z13 (TU3) may apply.

The previously announced AWLC Transition Charges for Sysplexes (TC2) is updated to include the z13.

Will I still be able to recognize the benefits of sub-capacity pricing with AWLC on a z13?

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 of the Technology Transition Offerings may be used in a sub-capacity environment.

What contracts are required for AWLC pricing and the Technology Transition Offerings for z13s or z13?

In order to get AWLC pricing or any of the Technology Transition Offerings, the following contract is required to be executed:

- Attachment for IBM 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 z Advanced Workload License Charges (Z125-8994). This supplement does not require a signature.

What is the Processor Value Unit (PVU) rating used for Passport Advantage® software on the z13?

The PVU rating for the z13 will be 120 for both IFL engines and CP engines, the same as the zEC12, z196 and the z10 EC. The PVU rating for the zBC12, z114 and z10BC remains 100 for both IFL engines and CP engines.

Where can I get more information on IBM z Systems software charges?

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

Academic Initiative and Skills

I am a student and the mainframe platform is new to me. How can I learn more about it? Are there ways for students to build skills in this area?

Students can absolutely learn more about the mainframe, and have fun while doing it too! IBM hosts a Master the Mainframe contest that students can enter. The contest is free to join and no previous mainframe knowledge is required. Participants earn prizes as they advance through three stages of the competition. Learn more about this exciting competition on our Academic Initiative website.

You can also build enterprise mainframe skills through our vast global network of schools that are a part of the z Systems Academic Initiative. These schools offer a variety of courses and curriculum across all areas of operating systems and application development. For more information see our school web page.

IBM also provides no fee community resources that are available through the course material webpage.

What is the IBM global plan for building the next generation of skills on the z Systems platform?

There is a strong commitment in the z Systems community for building skills in our industry. Watch students participating in our computing competitions on our Master the Mainframe YouTube channel.

You will find a list of universities IBM and our community are working with to help develop mainframe skills as well as see a broad spectrum of resources available to students. For more information see our school web page.

IBM also offers a no free job board to connect employers with job seekers. You can visit the site at zSystemsJobs.com.

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

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 send a note to zSkills@us.ibm.com with your request.

We're used to using VMWare for virtualization, so z/VM is new to me. How do I learn more about z/VM and Linux on z?

If you are used to using VMWare then learning z/VM should be an easy skill to acquire. Classes on z/VM and Linux on z are available through IBM's Global Training Providers. Visit the z/VM and Linux on z training roadmaps for a complete listing of courses.

Marist College also offers an instructor led online course: Running Linux Systems in a z/VM environment.

I have new hires joining my company and need to get them trained on z, what resources are available to help?

Several training options exist for building z skills:

Explore the z Systems training path tool to select the courses you need, view their course descriptions, and then enroll in the classes that suit your needs.

Leverage the Mainframe Contest Learning System. No mainframe skills are required to start learning z with this system. Build mainframe skills as you progress through all three stages of the self-paced learning.

Attend a one week, no-charge "z/OS Introduction and Workshop" either in person at the Dallas Innovation Center, or remotely. The first class is scheduled for June 27 – July 1, and a second class will be added later in the year. Send an email to zskills@us.ibm.com for additional information.

Marist College offers z certificate programs in z/OS Systems Programming (Associate, Professional and Expert), COBOL Application Programming, DB2 Application Programming, IMS Application Programming, and Assembler Language Application Programming.

zBX

Where can I find a full set of zBX frequently asked questions?

Please visit:
http://www.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ST&infotype=SA&htmlfid=ZSQ03098USEN&attachment=ZSQ03098USEN.PDF

IBM Corporation
New Orchard Road
Armonk, NY 10504
U.S.A.
Produced in the United States of America,
05/2017

ZSQ03076-USEN-15