
IBM Z
Introduction
October 2018

IBM z14
Frequently Asked Questions

Worldwide



ZSQ03115-USEN-07

Table of Contents

z14 Hardware.....	3
IBM Storage	13
Pervasive Encryption	14
Performance	16
z14 Warranty.....	20
Hardware Management Console (HMC)	21
HMC Mobile for IBM Z.....	24
Power (including HV-DC Power) and Overhead Cabling.....	29
Water cooling option for z14 Models M01-M05	31
IBM Secure Service Container	37
z/VSE Network Appliance	39
IBM Dynamic Partition Manager	40
I/O	42
IBM Virtual Flash Memory	48
Networking.....	50
zEDC Express and Compression Acceleration.....	50
Parallel Sysplex.....	55
STP	60
z/OS.....	61
z/VSE.....	62
z/VM	63
KVM.....	64
Linux on IBM Z	65
Compilers	66
Academic Initiative and Skills.....	67

z14 Hardware

Tell me more about the IBM z14™ from 2017 and the IBM z14 (z14) that was introduced on April 10, 2018?

In July 2017, IBM introduced the IBM z14 Models M01, M02, M03, M04 and M05. The z14 is designed to address the important issues of security and to be the infrastructure clients can trust in the digital economy. The z14 offers a new server for our traditional IBM Z® clients, but also provides a rapid deployment model for faster time to value, a highly secure and trusted environment to protect from breaches and an open software stack for the widest choice in applications.

On April 10, 2018, IBM has developed a new model for the z14 with a lower cost and point of entry to extend the IBM Z to clients of all sizes. The new IBM z14 Model ZR1 delivers secure capabilities in a smaller, industry-standard 19-inch frame, that can easily co-exist with other platforms in a cloud data center. It can be the base for an integrated hybrid cloud for mission critical core business workloads that which demand maximum security.

The introduction of the z14 ZR1 also demonstrates that IBM will no longer distinguish between ‘midrange’ and ‘high end’ systems. IBM Z will drive to merging the life cycles of our offerings, introducing a family of servers that spans both single frame and dual frame models.

What is the machine type of the z14?

The IBM z14 machine type for the models M01, M02, M03, M04 and M05 is 3906.

The IBM z14 Model ZR1 machine type is 3907.

How much capacity will the z14 deliver?

Powered by 170 of the world’s most powerful microprocessors IBM z14 Model M05 delivers up to 35% more total general system processing capacity than the IBM z13® (z13). The z14 Model ZR1 is expected to deliver 13% more total general system processing capacity than the IBM z13s® (z13s).

How many capacity options are offered on the z14 Model ZR1?

The z14 Model z14 offers the same granularity as the z13s – 26 capacity levels x 6 CPs equals

What is the available subcapacity granularity on the z14 dual frames?

The z14 dual frame models can have up to 33 subcapacity Central Processors (CPs) per server while the z13 accommodates up to 30 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 model was announced for the IBM z14 on April 10th 2018?

On April 10, 2018, IBM announced the IBM z14 Model ZR1, a single frame model of the IBM z14 product line. There are thirty configurable cores in the z14 ZR1. The z14 ZR1 is introducing new feature based sizing (4, 12, 24, 30) – done at configuration time. This will replace a model sizing for the different sizes of z14 ZR1 available and will be designated by use of a feature code.

The Max4 feature code (FC #0636) can be a 1-way through a 4-way – which means there are 4 orderable cores contained on one CPC drawer (plus supporting cores including a spare core). When ordering with Max4 feature code, there can be up to four general purpose processors (CPs) defined using the capacity setting feature – meaning there are 104 capacity settings with the Max4 feature code – no xx5 or xx6 capacity settings. Any unused cores can be configured as 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. There can be up to 2 TB of memory and up to 16 I/O features configured with the Max4 feature code. There can be up to 2 features of ICA SR and up to 16 features of Coupling Express LR on the Max4.

The Max12 feature code (FC #0637) can be a 1-way through a 12-way – which means there are 12 orderable cores contained on one CPC drawer (plus supporting cores including a spare core). When ordering with Max12 feature code, there can be up to six CPs defined using the capacity setting feature – the full 156 granular options. Any unused cores can be configured as IFLs, zIIPs, SAPs, ICFs and/or additional spares. There can be up to 4 TB of memory and up to 32 I/O features configured with the Max12 feature code. There can be up to 4 features of ICA SR and up to 16 features of Coupling Express LR on the Max12.

The Max24 feature code (FC #0638) can be a 1-way through a 24-way – which means there are 24 orderable cores contained on two CPC drawers (plus supporting cores including a spare core). When ordering with Max24 feature code, there can be up to six CPs defined using the capacity setting feature – the full 156 granular options. Any unused cores can be configured as IFLs, zIIPs, SAPs, ICFs and/or additional spares. There can be up to 8 TB of memory and up to 64 I/O features configured with the Max4 feature code. There can be up to 8 features of ICA SR and up to 16 features of Coupling Express LR on the Max24.

The Max30 feature code (FC #0639) can be a 1-way through a 30-way – which means there are 30 orderable cores contained on two CPC drawers (plus supporting cores including a spare core). When ordering with Max30 feature code, there can be up to six CPs defined using the capacity setting feature – the full 156 granular options. Any unused cores can be configured as IFLs, zIIPs, SAPs, ICFs and/or additional spares. There can be up to 8 TB of memory and up to 64 I/O features configured with the Max4 feature code. There can be up to 8 features of ICA SR and up to 16 features of Coupling Express LR on the Max30.

For additional information on specs refer to the [z14 Model ZR1 datasheet](#).

What dual frame models are available for the IBM z14?

On July 17, 2017, IBM announced five z14 dual frame models.

A z14 Model M01 model can be a 1-way through 33-way – which means there are 33 orderable cores contained on one CPC drawer (plus many supporting cores in every model, including a minimum of two spare cores).

A z14 Model M02 model can be a 1-way through 69-way (69 orderable cores) contained in two CPC drawers.

A z14 Model M03 model can be a 1-way through 105-way (105 orderable cores) contained in three CPC drawers.

A z14 Model M04 model can be a 1-way through 141-way (141 orderable cores) contained in four CPC drawers.

The enhanced capacity z14 Model M05 model can be a 1-way through 170-way (170 orderable cores) contained in four CPC drawers.

Customers that reach CPC drawer limits can easily upgrade from M01 models to M04 models non-disruptively, i.e. without requiring a service interruption of the machine. There is no upgrade option to get to a Model M05. The z14 Model M05 must be an initial order only.

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-three sub-capacity CPs can be active on the server regardless of hardware model. In other words, sub-capacity configurations are available on any of the models (M01 to M05) as long as the server is configured (not necessarily the same as purchased) with thirty-three or fewer general-purpose processors.

For more information refer to the [z14 Models M01-M05 datasheet](#).

What is different about the z14 Model M05 compared to z13 Model NE1?

The z14 Model M05 is an enhanced capacity model which contains a different configuration of SCMs than other models. The z14 Model M05 is fully populated with four high density CPC drawers and 170 orderable cores. You can configure the M05 machine to be a 1 to 170-way. The M05 can be ordered with a minimum of 256 GB of memory and up to a maximum of 32 TB.

What is new is that there is no upgrade path into the z14 M05 from other z14 models. Clients that are looking for the enhanced capacity model will need to make it the initial order.

What are some key innovations in the z14 processor chip?

The new 10-core z14 processor chip leverages the density and efficiency of 14nm silicon-on-insulator technology to deliver increased performance and capacity across a wide range of workloads. Much of that increase results from innovations in the microprocessor design, driven by tight collaboration across hardware, firmware, and software development. These innovations include:

- IBM Z has redesigned the cache architecture with 1.5x more on-chip cache per core compared to the z13/z13s and on-chip L3 cache was doubled. Bigger and faster caches help to avoid untimely swaps and memory waits while maximizing the throughput of concurrent workloads.
- The z14 consolidates the z13 two storage controller chips per drawer into one, eliminating a chip crossing for many accesses to memory and shared data, thus reducing the time processors spend waiting for the data they need.
- New instructions in the single instruction, multiple data (SIMD) facility offers a boost for traditional workloads using decimal operations (i.e. COBOL, PL/I) and new application like analytics (i.e. Open Data Analytics for z/OS®). Select PKCS#11 clear key operations also now support for SIMD for performance gains.

- The new Guarded Storage Facility (GSF) will deliver pause-less garbage collection to enable enterprise scale Java™ applications to run with fewer and shorter pauses for garbage collection on larger and larger heaps. Our performance work has determined that customers running z/OS Java SDK 8 applications can achieve up to 3x better throughput for response-time constrained Service Level Agreements (SLAs) using new IBM Java 8 SR5 Pause-less garbage collection feature on z14 compared to using Java 8 SR3 on z13.¹
- The z14 has optimized second generation Simultaneous Multithreading (SMT) for superior Linux performance with new efficiency enhancements – better thread/resource balancing, address translation redesigned to avoid stress from SMT for the translation buffer, Linux aware branch prediction tuning help Linux mapping for modules, libraries and binaries, performance optimized key virtualization for Linux (aka keyless guest), instruction execution protection facility adds security by preventing stack-overflow and similar attacks (malware attacks). SMT will also continue to support zIIPs and has been extended to support System Assist Processors (SAPs). SMT on the z14 improves throughput for an IFL or zIIP to benefit exploiters.
- Improved compression ratio (using Huffman coding) and order-preserving compression for the on-chip compression coprocessor which results in fewer CPU cycles to enable further compression of data, improving memory, transfer, and disk efficiency. In the future Db2® plans to enable new order-preserving compression for Db2 indices² which can also take advantage of the compression co-processor
- Every core has hardware accelerated encryption implementing the CP Assist for Cryptographic Function (CPACF) which provides cryptographic functions and hashing functions in support of clear-key operations. CPACF encryption rates for like modes and data sizes on z14 are up to six times faster than z13/z13s³. Use of encryption is a key element of the pervasive encryption capabilities of the z14 platform.
- The z14 ZR1 can have up to 8 TB of memory and the z14 Models M01-M05 can have up to 32 TB memory. The additional memory can support new workloads, data-in-memory applications, larger local buffer pools as well as efficiently process huge amounts of information for faster business insight.
- The z14 ZR1 offers 30 configurable cores for Linux – a 50% increase from the z13s. Up to 170 configurable cores the z14 M01-M05 offers up to 35% more total capacity in a single footprint compared to a z13. The z14 can scale to securely support the digital engagement that is accelerating as business executes on digital transformation. The amount of work generated by of mobile, analytics, vast amounts of data and customized user experiences require business to have greater scalability.

¹ Performance results based on IBM internal tests running Java Store Inventory and Point-of-Sale in COMPOSITE mode application benchmark on standalone dedicated IBM z14 and z13 machines using z/OS 2.2.2 with APAR OA51643 and no other workloads running in the LPAR under test. Both z14 and z13 were configured with 1CP and 8 SMT zIIPs with total 17 hardware threads. Capacity projections were done to estimate the benefit of moving from z13 z/OS 2.2 Java 8 SR3 to z14 z/OS 2.2 Java 8 SR5 with Pause-less garbage collection enabled by java option -Xgc:concurrentScavenge. The response-time constrained Service Level Agreements (SLAs) metric used for this claim was based on geometric mean of (throughput @ 10ms, 25ms, 50ms, 75ms and 100ms response time SLAs). Hardware instrumentation data was collected and analyzed on all benchmarks to verify performance results. IBM 64-bit SDK for z/OS Java Technology Edition, Version 8 SR3 on z13 was used as the baseline. IBM 64-bit SDK for z/OS Java Technology Edition, Version 8 SR5 is scheduled to GA September 2017.

² IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

³ Based on preliminary internal IBM lab measurements on a standalone dedicated system in a controlled environment and compared to the z13. Results may vary

Is it true the z14 ZR1 is housed in an industry standard 19” rack?

Yes, the z14 Model ZR1 is housed in an industry standard 19-inch rack form factor. The design will have power distribution unit (PDU)-based power along with redundant power, cooling and line cords. These factors let you lower power costs, reduce footprint cost and install in virtually any existing data center. By having the same footprint as other datacenter servers, you gain facility standardization.

The new 19-inch rack design opens a new opportunity for IBM Z. For the first time in IBM Z history, depending on the configuration, there is a potential for up to 16U of available frame space – called the 16U Reserved feature.

I have lots of spare racks in my datacenter, can I supply the rack for the IBM z14 ZR1?

No, the z14 ZR1 will still be assembled at an IBM plant prior to shipping to a client. The rack will be part of the standard IBM configuration of the server.

What is the new 16U Reserved feature code on the IBM z14 ZR1?

The z14 ZR1 is really an exciting introduction for IBM Z. The 19” rack offers great potential to be able to coexist with other platforms in the same data center. Another feature that will make this possible is the new 16U Reserved feature code (FC# 0617). For the first time in IBM Z history, depending on the configuration, there is a potential for up to 16U of available frame space, optionally available to be used to support other IBM or non-IBM equipment.

The 16U of space in the rack is specifically available for smaller I/O configurations – in place of a third or fourth PCIe+ drawer. Businesses might choose to install storage, servers, switches or other equipment within the frame when space is available. There are some restrictions that apply so see the z14 Installation planning guide for more details.

When the feature code is ordered, it may provide two additional Power Distribution Units, a cable management structure, air flow fillers, and weight ballast depending on the number of PCIe+ IO drawers.

Why might I be interested in the rack mounted HMC or rack mounted TKE on the z14?

The 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.

The rack mounted HMC or TKE can be installed in the 16U reserved feature (FC #0617) of the z14 ZR1 if that has been ordered and space is available. But 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.

What if I want to add a 3rd or 4th I/O drawer at a later point to the z14 ZR1 and remove the 16U feature - can the ballast and feature be removed?

The IBM z14 Model ZR1 introduced a new space for 16U Reserved (#0617) that could be used for qualifying client server-related entities. In the updated release level, we allow clients that have ordered the 16U Reserved space to be able to delete that feature if they will be adding I/O adapters that require the third I/O drawer slots. This means a more future-ready, flexible configuration for clients in data warehouses where floor space is a premium. In addition, the z14 Model ZR1 16U Reserved space can be utilized for IBM DS8882F Storage.

How is simultaneous multithreading (SMT) implemented on the z14?

Simultaneous multithreading (SMT) allows two active instruction threads per core, each dynamically sharing the core's execution resources. The implementation of SMT on z14 is similar to that on z13/z13s, with incremental improvements in the design. As on the z13/z13s, SMT will be available in z14 for workloads running on the IFL and the zIIP.

Each operating system / hypervisor may 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, this is done 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). z/VM 6.4 supports dynamically switching between one thread and two threads per core with no IPL required when enabled for SMT.

What is the HSA and how much is there on the z14?

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 z14 ZR1 there is 64 GB of HSA, and on the z14 M01-M05 there is 192 GB of HSA, independent of client-purchased memory.

How much memory is available on the z14?

The z14 ZR1 supports up to 8 TB of real memory. The z14 M01-M05 server supports up to 32 TB of real memory per server (LPAR limits are dependent on the operating system or 16 TB), 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 256 GB for all models.

What is RAIM?

The z14 is designed with redundant array of independent memory (RAIM) technology. RAIM is analogous 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 listed maximum memory is customer-usable, RAIM-protected memory. RAIM is always active, and IBM has already factored in an additional 25% of RAIM memory supplying redundancy.

What value might I see on if I purchase larger memory on the z14?

Memory on the z14 will advantage many types of users. Large memory can reduce latency and CPU cost, and thus improving operational efficiency. The additional memory can support new workloads, data-in-memory applications, efficiently process huge amounts of information for faster business insight. Large local buffer pools for Db2, VSAM and other data managers may see better CPU time, better response time, significant synchronous read IO reductions and much improved transaction rates⁴.

z/VM 6.4 supports 2 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) of a z14?

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 adapters –10GbE RoCE Express2, Coupling Express LR and zEDC Express. It is not customer usable or 'visible'.

How many spare processing cores are on the z14?

IBM ships every z14 ZR1 with a minimum of one spare and the z14 M01-M05 machines with a minimum of two spare processing cores. These spares can be shared across the drawers. The z14 offers core-level (engine-level) sparing.

Can any of the z14 spare processing cores be used for other purposes?

No, the minimum allotment of z14 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 SAPs are on the z14?

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

- The z14 Model ZR1 has two SAPs.
- The z14 Model M01 has five SAPs.
- The z14 Model M02 has ten SAPs.
- The z14 Model M03 has fifteen SAPs.
- The z14 Model M04 has twenty SAPs.
- The z14 Model M05 has twenty-three SAPs.

More SAPs may be acquired from the pool of available processing units within the model.

⁴ See IBM z13 and large memory blog: <http://mainframeinsights.com/ibm-z13-performance-benefits-large-memory/>

What benefit will I see in SSL handshaking with use of Crypto Express6S on the z14?

SSL handshake performance on z14 with Crypto Express6S is 2X the SSL handshake performance on z13/z13s with Crypto Express5S.⁵

What is the cooling on the z14?

The z14 ZR1 is a single frame, air cooled system only. Air flow is designed to cool front to back on the subsystems. If a client chooses to take advantage of the 16U reserved space filler components MUST be installed front to back to take advantage of the air flow as designed.

The z14 M01-M05 models are designed with an environmental focus to help improve data center efficiency. They have 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.

What is the water cooled option for the z14?

Optional water cooling is available for the z14 M01-M05 models. 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.

There is no water cooling for the z14 ZR1.

What is ASHRAE?

Data centers continue to use higher amounts of energy with each generation of IT equipment. 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 z14?

The z14 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>

⁵ Disclaimer: Based on preliminary internal IBM lab measurements on a standalone dedicated system in a controlled environment and compared to the z13. Results may vary.

Can I buy a z14 that has only IFL or ICF processors without including a general-purpose processor (CP)?

Yes. You can order only IFLs or ICFs in a z14, using a model capacity identifier of A00 for the z14 Model ZR1 with 1 to 30 IFLs or ICFs, and capacity identifier 400 for the z14 Models M01-M05 with a maximum of 170 IFLs or ICFs. There is still a limit of 16 ICF engines for a single Coupling Facility LPAR.

Besides using the A00 or 400, IBM does offer Linux only servers based on z14 technology – the IBM LinuxONE Emperor™ II and IBM LinuxONE Rockhopper™ II. The LinuxONE™ systems may optionally add one CP but it can only be used for the GDPS® Appliance.

Is there still an option for overhead cabling on z14?

Yes, you can order overhead cabling and you can also order overhead power on the z14. On the z14 ZR1 there are flexible options for top and bottom exit cabling.

Tell me about the non-raised floor option on the z14?

The non-raised floor option allows a z14 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 z14, you must order overhead power, overhead cabling, and radiator-based air-cooling.

What is the LPAR absolute hardware capacity setting on z14?

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 z14?

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 LPAR dynamic memory management on z14?

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.

This is supported in both PR/SM mode and IBM Dynamic Partition Manager (DPM) mode.

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

Yes, IBM Systems Lab Services and our Global Training Providers can assist with this. IBM Lab Services provides training through its technical events and private skills transfer engagements to clients and business partners. The Global Training Providers provide ongoing client digital and classroom based education. You can contact IBM Systems Lab Services via the Internet at:

<http://www.ibm.com/systems/services/labservices/> or send an email to ibmsls@us.ibm.com

Are there resources available to help me migrate to the new machine and to help with new workloads?

Yes, IBM Systems Lab Services can provide assistance to clients along a wide portfolio of options for clients. You can contact IBM Systems Lab Services via the Internet at:

<http://www.ibm.com/systems/services/labservices/> or send an email to ibmsls@us.ibm.com

IBM Storage

What is the Storage that IBM announced on August 21, 2018 concerning DS8880F Storage and IBM z14 Model ZR1?

IBM Storage delivered a high-performance, high availability, ECKD™/Distributed format flash storage that can be mounted in an IBM z14 ZR1 server rack within the 16U Reserved space.

The IBM DS8880F family now offers a new all-flash array model, DS8882F, to meet the demand for higher speed storage in a smaller footprint. All-flash arrays deliver higher IOPS and bandwidth with lower power consumption to reduce the total cost of ownership as compared to hybrid or HDD-based solutions.

DS8882F is a rack-mounted, all-flash storage system designed for customers with small capacity needs looking to boost performance. With 16U height, DS8882F is designed to take advantage of the space available in the IBM Z z14 Model ZR1 system frame, and can be integrated into existing 19-inch form factor racks.

Does the new DS8880F storage connect to the z14 ZR1 with external FICON® connections?

The storage and the server can be connected with a FICON Express16S+ and an IBM zHyperLink™ Express connection.

Can I install spinning disk drives in this new DS8880F product?

No, this new product only supports the IBM High Performance Flash Enclosure (HPFE).

Pervasive Encryption

What does IBM mean by pervasive encryption?

Many encryption policies have been based on selective encryption—manual selection of the most critical data. Development of a comprehensive approach to data protection requires a large investment in time and money. The decision on encryption had to involve where to execute, what is the impact to service level agreements, who is responsible, how extensive is the plan for encryption and how much will it cost.

The IBM z14 introduces a new approach. IBM z14 allows business to defend and protect critical assets with unrivaled encryption and intelligent data monitoring—without compromising transactional throughput or response times. Most importantly, there are no application changes.

We call this pervasive encryption.

By encrypting as much of your data and transactional pipeline as possible, you can reduce potential data breach risks and financial losses—and help comply with complex regulatory mandates. IBM z14 pervasive encryption capabilities give you a transparent and consumable approach to provide 100% coverage for in-flight and at-rest data with zero host CPU host. Using simple policy controls announced in z/OS 2.3, z14 pervasive computing streamlines data protection for mission critical datasets.

How is the z14 able to do pervasive encryption?

The performance enhancements on the IBM z14 make pervasive encryption achievable and that changes the security game. The Central Processor Assist for Cryptographic Function (CPACF), standard on every core, provides the dramatic improvements in hardware acceleration that make support of pervasive encryption affordable. CPACF will be used for encryption of data in-flight and data at-rest. CPACF on the z14 provides more than 6x faster encryption for like modes and data sizes with compared to z13/z13s.⁶ The new Crypto Express6S adapter offers 2x the SSL handshake performance on z14 with compared to z13/z13s with Crypto Express5S.⁷ New Crypto enhancements enables compliance with PCI, ANSI and other evolving standards, provides enhanced performance, simplified Trusted Key Entry (TKE) processes and a new smart card to meet expected encryption strengths required for compliance.

A new smart card was shipped with TKE 9.1. The blue smart card supports 521-bit Elliptic Curve (EC) cryptography and has more memory than the previous TKE smart card. Many new features in TKE 9.1 take advantage of, or require the new blue smart card.

There are important limitations you should know about the blue smart card

- TKE 9.1 is the minimum supported release for the blue smart card.
- The blue smart card is configured to run in FIPS mode. Therefore, the smart card hardware prohibits the generation of any 1024-bit RSA keys.

⁶Based on preliminary internal IBM lab measurements on a standalone dedicated system in a controlled environment and compared to the z13. Results may vary.

⁷Based on preliminary internal IBM lab measurements on a standalone dedicated system in a controlled environment and compared to the z13. Results may vary.

For z/OS, dataset encryption will protect z/OS data sets automatically throughout their life cycle. Linux file and volume encryption transparently exploits on chip cryptographic functions to leverage hardware performance gains and expects⁸ an extended level of security through protected key for encryption for data at-rest. This will allow encryption of complete disks (volumes) for example.

Coupling facility (CF) encryption is a key piece of pervasive encryption helping to protect z/OS CF data end to end, using encryption that's transparent to applications. Although no encryption happens on the CF itself, data is encrypted on a host in the sysplex using CPACF on a per-workload, per-structure basis, using established policies before being sent to the CF. The data written to the CF remains encrypted until it is read from the CF and decrypted by a host elsewhere in the sysplex. Thus, the encrypted data is safe in-flight as it flows to the External Coupling Facility, as well as when it is at-rest on the CF.

Where can I get more information on pervasive encryption?

There is a set of FAQs on pervasive encryption located at:

<https://www.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSQ03116USEN>

I am interested in ensuring I am fully utilizing the new pervasive encryption. Are there resources available to help me assess my shop in terms of my readiness and to assist in implementing any gaps for me?

Yes, IBM Systems Lab Services has a Pervasive Encryption Readiness Assessment, and has a full set of offerings to fully execute the various level of encryption and security. You can contact IBM Systems Lab Services via the Internet at: <https://www.ibm.com/it-infrastructure/services/lab-services> or send an email to ibmsls@us.ibm.com.

I am interested in ensuring I am fully utilizing the all of the security capabilities of the z14, but I am currently using non-IBM security software. Are there resources available to help me migrate to IBM supported security products?

Yes, IBM Systems Lab Services can help you through the migration from non-IBM security software, and to implement all of the hardware and software security capabilities that IBM has to offer. You can contact IBM Systems Lab Services via the Internet at: <https://www.ibm.com/it-infrastructure/services/lab-services> or send an email to ibmsls@us.ibm.com.

Are there any publications I can reference to get started with data set encryption on IBM Z which is a core component of pervasive encryption?

Yes please reference the Redbook, "Getting Started with z/OS Data Set Encryption" at www.redbooks.ibm.com/redpieces/abstracts/sg248410.html?Open

Also, for z/OS V2.2 clients, or clients with both z/OS V2.3 and V2.2 that want to do data set encryption, there's a publicly available data set encryption FAQ specific to z/OS V2.2:

<https://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FQ131494>

⁸ IBM is working with the Linux distribution partners to get the functionality included in their distribution for Linux on IBM Z.

Performance

What are the major changes to the z/OS V2R2 LSPR?

The LSPR ratios reflect the range of performance between IBM Z mainframes 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 IBM System 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/resourceink/lib03060.nsf/pages/lspindex>

What is the multi-image table in the LSPR?

Typically, IBM Z processors are configured with multiple images of z/OS. Thus, the LSPR continues to include a table of performance ratios based on average multi-image z/OS configurations for each processor model as determined from the profiling data. The multi-image table is used as the basis for setting MIPS and MSUs for IBM Z 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 five at low-end models to nine at the high end. Most systems were configured with two major images (those defined with >20% relative weight). On low- to mid-range 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 processor 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 V2R2 LSPR table?

The 2094-701 processor model is used as the base in the z/OS V2R2 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 V2R2 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 V2R2 multi-image table the commonly used scaling factor is $0.944 \times 593 = 559.792$. Note the 0.944 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 V2R2 multi-image table by 559.792. The PCI (Processor Capacity Index) column in the z/OS V2R2 multi-image table shows the result of this calculation. Note that the PCI column was 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 an IBM z14 processor?

As with the introduction of any new server, workloads with differing characteristics will see variation in performance when moved to an IBM z14. The performance ratings for a server are determined by the performance of a reference workload that represents what we understand to be the major components of our customers' production environments. While we feel the ratings provide good "middle-of-the-road" values, we also recognize some customers' workloads will differ somewhat from the reference workload we used. The IBM z14 has improvements in its microprocessor design and in its memory hierarchy. However, workloads with different characteristics will see varying performance values from these changes. It is expected that the range of variation in performance of workloads will be similar to that seen in recent processor generations.

Once my workload is up and running on an IBM z14, 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 microprocessor and memory hierarchy design and support for larger numbers of engines, the IBM z14 provides a significant increase in capacity over the largest previous server in each family. Continued enhancements to z/OS 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, some variation in performance might be seen, simply due to the expected larger and more complex LPAR configurations that can be supported by the IBM z14.

How do I get performance information for my TPF products running on an IBM z14?

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 z/OS HiperDispatch and how does it impact performance?

z/OS HiperDispatch is the z/OS exploitation of PR/SM’s Vertical CPU Management (VCM) capabilities and is exclusive to IBM Z processors since the IBM System

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 z/OS HiperDispatch at: <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101229>.

What is z/VM HiperDispatch and how does it impact performance?

z/VM HiperDispatch is the z/VM exploitation of PR/SM's Vertical CPU Management (VCM) capabilities. z/VM HiperDispatch improves CPU efficiency by causing the z/VM Control Program to run virtual servers in a manner that recognizes and exploits IBM Z machine topology to increase the effectiveness of physical machine memory cache. This includes: a) requesting PR/SM to handle the partition's logical processors in a manner that exploits physical machine topology, b) dispatching virtual servers in a manner that tends to reduce their movement within the partition's topology and c) dispatching multiprocessor virtual servers in a manner that tends to keep the server's virtual CPUs close to one other within the partition's topology. z/VM HiperDispatch can also improve performance by automatically tuning the LPAR's use of its logical CPUs to try to use only those logical CPUs to which it appears PR/SM will be able to deliver a full physical processor's worth of computing power. This includes: a) sensing and forecasting key indicators of workload intensity and b) automatically configuring the z/VM system not to use underpowered logical CPUs.

An article is available concerning z/VM HiperDispatch at: <http://www.vm.ibm.com/perf/tips/zvmhd.html>.

What is the performance improvement a z/VM customer might experience on an IBM z14?

The performance ratios a z/VM customer workload might experience when migrating to an IBM z14 from older processors will vary. For the z/VM LSPR curves, a single workload having characteristics similar to the AVERAGE relative nest intensity workload was used. 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 Linux on IBM z14 customer might experience?

The performance ratios a Linux customer workload might experience when migrating to an IBM z14 from older processors will vary. For the Linux on Z LSPR curves, a single workload having characteristics similar to the LOW relative nest intensity workload was used. 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.

What is the performance improvement a z/VSE customer might experience on an IBM z14?

The performance ratios that a z/VSE® customer workload might experience when migrating to an IBM z14 are represented by the range of ratios for a comparable z/OS migration. For example, the published ratio in the LSPR between the z13 702 and the z14 702 is approximately 6% to 11%. z/VSE workloads should expect this same range of performance for this migration. Consult the LSPR for other examples of moves to an IBM z14.

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

<https://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS1381>

z14 Warranty

Will my IBM service rep install components (IBM or non-IBM) into the 16U Reserved space of the z14 ZR1?

Installation of components in the 16U Reserved space of the z14 ZR1 is customer responsibility.

Who will do problem determination (PD) and problem source identification (PSI) on the hardware that is installed within the 16U Reserved space of the z14 ZR1?

PD and PSI of components in the 16U Reserved space is customer responsibility.

What is proprietary diagnostic support for z14?

A z14 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 IBM Z Service Support Representative (zSSR) in doing problem repairs.

How long is a z14 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 z14 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 on the z14?

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 z14?

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

If I am on base support of a z14, 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 z14 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 new about the Hardware Management Console (HMC) for z14?

Driving for simplification was a key goal in development of z14 which is demonstrated by the updated design of the IBM Hardware Management Console (HMC). There are simplification updates to improve workspace and manage system time. New security features include Multifactor Authentication and a new HMC Mobile application for monitor and recover action controls. HMC 2.14.1 Enhancements provide improvements in the areas of I/O, Security, System Concurrency, User Experience, and HMC Mobile.

What is new with the user interface on the HMC for the z14?

The z14 Hardware Management Console and Support Element will only support Tree Style user interface. Classic Style user interface is no longer supported. The following enhancements have been made to Tree Style user interface to aid with new users of Tree Style as well as address user pain points of existing users.

- A new Masthead with Favorites and Search controls will help users quickly find and launch tasks.
- Tasks will now open in tabs within the user interface (instead of separate browser windows) to make finding and managing running tasks easier.
- There is also the ability to have the task tab 'pop-out' into a separate window giving ability to have a similar parallel task display window as on previous Tree Style controls and allowing the task display(s) to be viewed on other physical display(s).

What is the new Manage System Time task for the HMC on the z14?

The Manage System Time task replaces the System (Sysplex) Time task on HMC and can be the single point of system time management for multiple systems. The new Manage System Time task provides a simplified workflow for system time management including improved help tools to compliment and improve system administrator skills and inline definition of technical terms. There is improved user experience with visual representation of configuration panels and guidance provided within the workflows. Topology displays of system time networks help administrators run a configuration. And errors are surfaced in visualization for easier problem resolution of setup errors.

What are the new security capabilities of HMC?

New security capabilities in the Hardware Management Console (HMC) are available with IBM z14 including multi-factor authentication, firmware integrity monitoring in support of NIST Standard 800-147B, crypto compliance levels, FTP through HMC, SNMP/BCPii API Security Controls, secure console to console communication enhancements, remote browser IP address limiting, and more.

Can you tell me a little more the new Multi-Factor Authentication on IBM Z?

The Hardware Management Console will now offer an optional control of Multifactor Authentication in addition to the userid/password controls provided today.

If Multi-factor option is selected for a given user, that user will now be required to enter a second authentication factor using a TOTP (Time-based One-Time Password). It will be implemented by freely available smartphone and web apps utilizing a secret key provided per HMC user.

What is the new HMC Mobile Interface on IBM Z?

Introducing IBM HMC Mobile for Z and LinuxONE, the app that allows you to monitor and manage your systems from your mobile device. Keep watch over all of your systems and partitions and receive alerts when hardware messages, OS messages, or status changes arise.

Key features:

- Access all of your systems, even if they're spread across multiple HMCs
- View system status, hardware messages, and configuration details
- View partition status, OS messages, and configuration details
- Start/activate, stop/deactivate and change the activation profiles of partitions
- Use personalized push notification settings that allow you to receive updates only for the systems and partitions that you're most interested in monitoring
- Take advantage of personalized security options such as multi-factor authentication, fingerprint ID, restricted object access and actions based on user role and IP, a "view only" mode, and more

What is the new firmware integrity monitoring on the Support Element (SE) and HMC for the z14?

The z14 will offer an enhancement on the Support Element that provides notification if tampering with booting of firmware on the server (CPC) is detected. This enhancement is designed to meet the BIOS Protection Guidelines recommended and published by the National Institute of Standards and Technology (NIST) in Special Publication 800-147B.

If tampering is detected, the Support Element will issue a customer alert via a warning or a lock of the Support Element, depending on the configuration. If "call home" support is enabled on a z14 Hardware Management Console managing the Support Element, additional analysis of the Support Element will be performed and displayed by IBM Resource Link®.

In addition to this support, the Hardware Management Console also has been enhanced to provide firmware integrity for attempted tamper monitoring and reporting. A newly manufactured Hardware Management Console directly ordered with z14, or at a later time, is required for this protection. Any detected events of attempted tampering will be logged and will issue a customer alert via a warning or a lock of the Hardware Management Console, depending on setup configuration. In addition, if "call home" support is enabled on the Hardware Management Console, supplementary analysis of events logged by the Hardware Management Console will be available on IBM Resource Link.

Although customers can carry forward their Hardware Management Consoles on z14, these tamper protection capabilities will be delivered only on newly manufactured Hardware Management Consoles. The z14 environment can contain both Hardware Management Consoles that have been carried forward and newly manufactured Hardware Management Consoles.

What is the FTP though HMC that was announced with z14?

To maximize security features, we recommend that customers keep their IBM Z on a dedicated network with one HMC network used for that network and the second HMC network used for outward communication (IBM Support Facility, remote browsing, automation). However, for systems prior to z14, this created a security challenge for FTP operations originating from the SE from working. A customer has had to either put their FTP server on the IBM Z dedicated network or put their IBM Z on their intranet network.

Starting with z14, all FTP operations originating from the SE will be proxied through a managing HMC. This now allows the FTP SE originated operations to follow our security recommendation. In addition, all HMC/SE tasks that support FTP will provide three options of FTP: FTP, FTPS and SFTP.

- FTPS is SSL based and uses certificates to authenticate servers.
- SFTP is SSH based and uses SSH keys to authenticate servers.
- Username and passwords are required for client authentication in all 3 protocols.

What is new about the Secure Console to Console Communications of the z14?

HMC consoles have used anonymous cipher suites for some inter-console communication purposes. These cipher suites, while providing encryption and integrity validation, do not provide cryptographic authentication. Network security scanners can detect this, and anonymous cipher suites may not be acceptable to some customers' security policies. Starting with the IBM z14 HMC/SE, the secure console to console communications solution will no longer use anonymous cipher suites and will begin using an industry standard based password driven cryptography system. This system provides cryptographic encryption, integrity validation and authentication.

HMC Mobile for IBM Z

What are the minimum requirements to connect to a Hardware Management Console (HMC) on an IBM Z using HMC Mobile?

The HMC on the IBM z14, IBM z13 or IBM z13s must be at least version 2.14.0 and have MCL P42675.104 applied.

How do I connect to an HMC on an IBM Z using HMC Mobile on my smartphone?

You will first need to download the HMC Mobile app to your smartphone. Then you need to:

- Navigate to the **Home Screen**.
- Open Settings.
- Select Manage HMCs.
- Select the + icon.

How do I logout of an HMC on an IBM Z when I'm done using the HMC Mobile app on my smartphone?

To exit the HMC Mobile app you need to:

- Navigate to the **Home Screen**.
- Select the **Hardware Manage Consoles** button at the bottom of the screen.
- Select Logout.

What is 'system' referring to with an IBM Z?

A 'system' in IBM Z is a physical collection of hardware that is also referred to as a Central Processor Complex (CPC).

How do I change which of IBM Z environments I see with the HMC Mobile app on my smartphone?

From your smartphone you need to:

- Navigate to the **Home Screen**.
- Open **Settings**.
- Select **Manage HMCs**.
- Toggle systems between visible or hidden.

I'm using the HMC Mobile app for IBM Z, why is there a blue dot next to my system?

When using the HMC Mobile app, a blue dot next to a system indicates that it contains at least one unread hardware message on the system or OS message on one of the system's partitions.

On the HMC Mobile app for IBM Z, what does the check or 'x' icon next to each system and partition mean?

When using the HMC Mobile app, a check indicates that the status of the IBM Z or a partition on the Z is in an acceptable state. An 'x' indicates that the status of the system or the partition is in an unacceptable state.

On the HMC Mobile app for IBM Z, how often is the data for each system and partition refreshed?

The data on the HMC Mobile app for IBM Z is refreshed every 30 seconds.

What is a logical partition on an IBM Z?

Among the system control functions is the capability to partition the IBM Z into several logical partitions (LPARs). An LPAR is a subset of the processor hardware that is defined to support an operating system. An LPAR contains resources (processors, memory, and input/output devices) and operates as an independent system. Multiple logical partitions can exist within a mainframe hardware system. The installation planners may elect to share I/O devices across several LPARs, but that is a customer decision.

Why is there a blue dot next to a partition when I'm using the HMC Mobile app for IBM Z?

On the HMC Mobile app, a blue dot indicates that the partition contains at least one unread OS message.

Why can't I activate/start, deactivate/stop, or change the activation profile of a partition using the HMC Mobile app for IBM Z?

In order to perform an action on a partition, you must be granted permission. Identify and access management for HMC Mobile is controlled in the **HMC Mobile Settings** on the physical HMC.

Why is there a blue dot next to my hardware or operating system message when I'm on the IBM Z using HMC Mobile?

A blue dot next to a hardware or operation system message indicates that it is unread. After it is opened, the blue dot no longer appears. You can also mark all messages as read by selecting the **'mark all as read' icon** in the navigation bar.

What is a priority operating system message for IBM Z?

A priority operating system message on IBM Z indicates that the operation system requires immediate attention.

What is a held operating system message for IBM Z?

A held operating system message indicates that the operating system requires a response.

When using the HMC Mobile app on an IBM Z, how do I specify which events trigger a push notification?

When you are in the HMC Mobile app you need to:

- On the **Home Screen**, select the system from which you want to receive push notifications.
- Scroll down to Notification Settings.
- Select the events that you want to trigger a push notification.

When using the HMC Mobile app on an IBM Z, how do I specify which events trigger a push notification on a partition?

When you are using the HMC Mobile app you need to:

- On the **Home Screen**, select the system within which the partition is located.
- Push the **Partitions** tab.
- Select the partition from which you want to receive push notifications.
- Scroll down to Notifications Settings.
- Select the events that you want to trigger a push notification.

To receive push notifications for all of the partitions within a system, perform the following:

- On the **Home Screen**, select the system within which the partition is located.
- Scroll down to Notifications Settings.
- Select the events that you want to trigger a push notification.

On the HMC Mobile app for IBM Z, how do I view a list of all the push notifications I've received?

When using the HMC Mobile app, select the **'bell'** icon in the navigation bar on the **Home Screen**.

I'm not receiving push notifications on the HMC Mobile app for IBM Z – why not?

There are several reasons you might not be receiving push notifications:

- A system administrator has disabled push notifications for the HMC that the system is connected to. You can verify your security policies by navigating to **Home Screen > Settings > View security policies**. Identity and access management for HMC Mobile is controlled in the **HMC Mobile Settings** task on the HMC.
- Notifications for all of your systems and partitions are turned off. See "How do I specify which events trigger a push notification for a system?" and "How do I specify which events trigger a push notification for a partition?"
- Your phone settings might not allow HMC Mobile to send you push notifications.

With the HMC Mobile app for IBM Z, what is an HMC Mobile password?

An HMC Mobile password is a unique password that is separate from your HMC User ID password. It is used to unlock the app.

Using the HMC Mobile app for IBM Z, how do I set the HMC Mobile password?

To set a HMC Mobile password using the HMC Mobile app:

- Navigate to the **Home Screen**.
- **Open Settings**.
- Open **Configure HMC Mobile password**.
- Enable **HMC Mobile password**. If you have never set up an HMC Mobile password, you will be prompted to create one.

Using the HMC Mobile app for IBM Z, how do I enable or disable the HMC Mobile password?

To enable or disable the HMC Mobile password using the HMC Mobile app:

- Navigate to the **Home Screen**.
- **Open Settings**.
- Open **Configure HMC Mobile password**.
- Enable or disable **HMC Mobile password**.

Using the HMC Mobile app for IBM Z, how do I change my HMC Mobile password?

To change a HMC Mobile password using the HMC Mobile app:

- Navigate to the **Home Screen**.
- **Open Settings**.
- Open **Configure HMC Mobile password**.
- Select **Change HMC Mobile password**. Passwords can only be changed if **HMC Mobile password** is enabled.

Using the HMC Mobile app for IBM Z, can I use my TouchID, FaceID, or fingerprint as an HMC Mobile password?

To set up the use of TouchID, FaceID or fingerprint using the HMC Mobile app:

- Navigate to the **Home Screen**.
- **Open Settings**.
- Open **Configure HMC Mobile password**.
- Select **Use TouchID, Use FaceID, or Use fingerprint**.

While using the HMC Mobile app for IBM Z I'm being asked to enter an authorization code - why?

Your account has been enabled for multi-factor authentication which requires you to enter an authentication code every time you login. If you haven't setup multi-factor authentication, you must login to the HMC console, after which you will be prompted.

Where do I setup the HMC Mobile app for IBM Z security policies on the HMC?

Mobile security policies are defined in the **HMC Mobile Settings** task.

I am using HMC Mobile app for IBM Z, how do I view the security policies that were setup on my HMC?

To view the security policies that were set up on your HMC for the HMC Mobile app:

- Navigate to the **Home Screen**.
- Open **Settings**.
- Open **View security policies**.

What is User ID password secure storage in the HMC Mobile app for IBM Z?

User ID password secure storage allows your HMC User ID password to be securely stored on your mobile device. User ID password secure storage allows you to view the systems and partitions you have access to, but you will still need to login to make changes.

What are 'actions' in the HMC Mobile app for IBM Z?

All management activities are considered actions. For the current release of HMC Mobile, this is limited to starting/activating, stopping/deactivating, or changing the activation profile of partitions.

How do I access the new HMC Mobile Interface available?

Visit the website ibm.biz/hmc-mobile or download from the respective iOS (ibm.biz/hmc-mobile-ios) and Google play (ibm.biz/hmc-mobile-android) stores.

Power (including HV-DC Power) and Overhead Cabling

Will there be 3-phase power on the z14 Model ZR1?

No. Only single phase power will be available.

Why are there no Internal Battery Features (IBF) on the z14 ZR1?

The redesign of the z14 ZR1 for the 19" rack eliminated the use of the current 24" rack technology (i.e. IBM z13s) in the new system. Our research indicated that clients are more dependent on external PDUs than internal features so we decided not to design IBF in the new server.

On a picture of the server, I saw what looked like extension outlets in the z14 ZR1 frame. Is that what they are?

No. Those are there to power functional units (I/O, CPs) within the system. They are not available for client use.

How do I power the new equipment I install in my z14 ZR1 frame in the 16U Reserved space?

There are 10 outlets of the 12 in each of the lower pair of PDUs set aside for external equipment when the flex feature is ordered. Absent this feature, none of the outlets are available for client use.

Is there any limit on the power draw of the new equipment I install in my z14 ZR1 frame in the 16U reserved space?

Yes, there are power limits. No single outlet can exceed 2,000 watts. No adjacent pair of outlets can exceed a total of 3200 watts. The total power on all available outlets cannot exceed 3400 watts. For equipment which does not have dual utility power inputs, the limit of 3400 watts can be separately applied to each PDU of the lower pair.

What is line cord plan ahead (FC #2000) and Balanced Power Plan Ahead (FC#3003) on a z14?

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.

There is no line cord plan ahead or Balanced Power Plan Ahead on the z14 ZR1.

How many line cords will be required for High Voltage DC Power option on a z14?

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

The z14 ZR1 does not support High Voltage DC Power.

What nominal DC supply voltages will be supported for the z14?

We'll support 380 VDC – 520 VDC nominal.

The z14 ZR1 does not support High Voltage DC Power.

Will I need to have additional space in my floor cutouts because of these 'raceways' on the z14?

Top-exit I/O cabling will add 12" of width to the z14, but the increases will only be above the floor to the sides of the system.

What benefit is it that that overhead power is added after installation on the z14?

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 the z14?

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.

How much additional room is needed for the overhead options on a z14?

On the new single frame z14 ZR1, when top exit cabling is selected on an order, the required infrastructure is provided. With the z14 ZR1, there is a top hat that is attached to the top of the frame in order to manage overhead cabling. The top hat can be configured to exit the cables either from the top front or the top rear of the rack. This same feature can be used for overhead power cabling, or for overhead IO cabling. There are no longer any cable chimneys that would need attached to the sides of the racks. Overhead cabling will add 1" width and approximately 5" depth to the frame.

On the dual frame z14 models, overhead cabling will add 12" (~30 cm) to the width (not depth) of the combined frames on the z14. 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 z14 system's optional water connections are bottom exit only.

Water cooling option for z14 Models M01-M05

When should I consider installing water cooling on the z14 Models M01-M05?

A few examples are:

If you have installed z13, 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 IBM Z is one component but even more significant improvements in removal of heat load can be achieved by implementing water cooling across your other server platforms.

Which z14 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.

The z14 Model ZR1 will not be available as a water cooler server.

Can I tap into my building's chilled water for z14 water cooling?

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 z14 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 z14 system?

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 z14?

There is no change to the floor cutouts for the z14 with the water cooling option – but the option will add 4 inches depth to the rear of the server. The floor tile cutouts for z14 are the same as those for the z13 and zEC12.

Is there any change to the height of the server for water cooling on the z14?

No.

Is there any change to the weight of the server on the z14 with water cooling?

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

How many Water Cooling Units (WCU) are on a water cooled z14?

The z14 has two Water Cooled Units (WCU), in a N+1 design.

Will data center water be going through my z14 server with water cooling?

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 z14 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 z14?

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 z14 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 z14 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.

z14 Model ZR1 - Software Pricing

What software pricing is announced with the z14 Model ZR1?

A new Technology Transition Offering (TTO) called Technology Update Pricing for the z14 Model ZR1 (TU6) is being announced along with revisions to the Technology Update Pricing for the z14 offering and three revised Transition Charges for Sysplexes or Multiplexes offerings.

Technology Update Pricing for the z14 Model ZR1 (TU6) 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 z14 Model ZR1 (TU6) 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 z14 Model ZR1 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 z14 (TU5) offering and the three 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 z14 Model ZR1?

The percent reduction in the monthly AEWLC is based on the number of z14 Model ZR1 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 z14 Model ZR1 is reduced by the percentage in the table below.

Schedule of AEWLC reductions for Technology Update Pricing for the z14 Model ZR1 (TU6)

MSUs: Quantity of z14 Model ZR1 Full Capacity MSUs for a stand-alone server	Reduction in Monthly AEWLC
1-11 MSUs	18.0%
12-17 MSUs	18.0%
18-30 MSUs	18.0%
31-45 MSUs	15.0%
46-87 MSUs	14.0%
88-175 MSUs	14.0%
176-260 MSUs	13.0%
261-315 MSUs	13.0%
316-390 MSUs	13.0%
391 - more MSUs	13.0%

The number of MSUs used to determine the MSUs of a stand-alone z14 Model ZR1 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 z13s to a z14 Model ZR1 will software pricing improvements be additive?

You will receive the pricing advantage of Technology Update Pricing for the z14 Model ZR1 (TU6) for the z14 Model ZR1. The Technology Update Pricing for the z14 Model ZR1 (TU6) price reduction includes the benefits of the prior price reductions. Technology Update Pricing for the z14 Model ZR1 (TU6) is available when upgrading from a z13s to a z14 Model ZR1.

Schedule of AEWLC reductions for Technology Update Pricing for the z14 Model ZR1 (TU6)

MSUs: Quantity of z14 ZR1 Full Capacity MSUs for a stand-alone server	Reduction in Monthly AEWLC for z14 ZR1 (TU6)	Reduction in Monthly AEWLC for z13s (TU4)
1-11 MSUs *	18.0%	13.0%
12-17 MSUs *	18.0%	13.0%
18-30 MSUs	18.0%	13.0%
31-45 MSUs	15.0%	10.0%
46-87 MSUs	14.0%	9.0%
88-175 MSUs	14.0%	9.0%
176-260 MSUs	13.0%	9.0%
261-315 MSUs	13.0%	9.0%
316-390 MSUs	13.0%	9.0%
391 - more MSUs	13.0%	9.0%

*The MSUs for TU4 are 1-10 and 11-17 respectively

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 IBM Z Software Contracts website at:

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

What software pricing announcements were made for a z14 Model ZR1 in a Parallel Sysplex or Multiplex?

IBM announced revisions to the Technology Update Pricing for z14 (TU5) offering and three revised Transition Charges for Sysplexes offerings.

The revisions to the Technology Update Pricing for z14 offering and the three 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 z14 offers actively coupled Parallel Sysplexes and Loosely Coupled Complexes that consist entirely of z14 (Models M01 – M05 and Model ZR1) servers the benefit of aggregated Technology Update Pricing for z14 when priced with AWLC, or in a Multiplex priced with Country Multiplex License Charges (CMLC).

Transition Charges for Sysplexes or Multiplexes (TC4): When two or more servers that consist of one or more z14 (all models) servers with one or more z13 or z13s 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 that is related to Technology Update Pricing (TU3), which is based on the amount of server capacity within the TTO-eligible environment.

Transition Charges for Sysplexes or Multiplexes (TC3): When two or more servers that consist of one or more z14 (all models) servers with one or more z13, z13s, 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 that is related to Technology Update Pricing (TU3). This is based on the amount of z14 (all models), z13 and z13s server capacity within the TTO-eligible environment.

Transition Charges for Multiplexes (TC2): For existing qualified multiplex only, when two or more servers consisting of one or more z14 (all models), z13, z13s, zEC12, or zBC12 servers with one or more z196 or z114 servers that comprise a TTO-eligible environment, those servers receive a reduction to CMLC pricing across the TTO-eligible environment. This reduction provides a portion of the price-performance benefit that is related to Technology Update Pricing (TU1) for AWLC. This is based on the proportion of z14 (all models), z13, z13s, zEC12, or zBC12 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 z14 (TU5) or PSLC, but not both, for your IBM software programs within the z/OS or z/TPF operating system families on z14 (all models) servers in a Parallel Sysplex or Loosely Coupled Complex.

Will I still be able to recognize the benefits of sub-capacity pricing with AEWLC on the z14 Model ZR1?

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 z14 Model ZR1?

In order to get Sub-Capacity Technology Update Pricing for the z14 Model ZR1, the following contract is required to be executed:

ICA Attachment for IBM 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 Z Advanced Workload License Charges, Attachment for IBM Z Advanced Entry Workload License Charges, Attachment for Country Multiplex Pricing, and to IBM Z Machines Exhibit (Z125-8994) describes the Technology Transition Offerings and should be delivered to each z14 Model ZR1 customer, but it does not get executed.

The Z Machines Exhibit (Z125-3901) provides the terms for Full Capacity AEWLC. The Exhibit and the Supplement should be delivered to each z14 Model ZR1 customer, but they do not get executed.

What is the Processor Value Unit (PVU) rating used for Passport Advantage® software on the z14 Model ZR1?

The PVU rating for the z14 Model ZR1 will be 100 for both IFL engines and CP engines, the same as the Rockhopper II, z13s, Rockhopper, zBC12, z114 and the z10 BC. The PVU rating for the z14 Models M01 – M05, Emperor II, 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 software charges?

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

IBM Secure Service Container

What is IBM Secure Service Container for z14?

IBM Secure Service Container consists of two parts. First it provides a framework which the appliance software solutions use as base – which includes the operating system and middleware.

The IBM Secure Service Container is a partition type which, along with an appliance installer, enables the secure deployment of firmware and software virtual appliances. IBM Secure Service Container will be delivered as part of the base code on the z14.

What is the value of having the IBM Secure Service Container on a z14?

An IBM exclusive, IBM Secure Service Container builds upon the industry leading isolation of IBM Z logical partitions. The reason—even with the highest levels of peer isolation, many organizations realize they also need vertical isolation to protect sensitive data from administrative staff and contractors who manage the infrastructure. By completely restricting system administrator access to the container, IBM Secure Service Container protects against the misuse of privileged user credentials.

IBM 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.

Will I need to install a specific operating system to support a software virtual appliance on the z14?

No, IBM 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 an IBM Secure Service Container LPAR without having to install any OS nor middleware beforehand.

What differentiates the IBM Secure Service Container LPAR partition type on the z14?

The IBM Secure Service Container LPAR type provides a substantial security envelope for appliances, checked by the IBM 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.

If I use a pre-z13 server for disaster recovery, can I run software appliances on the IBM Secure Service Container on older servers?

No, IBM Secure Service Container is only available on the z14, z13s and z13. Selected appliances that do not depend on the security features listed above may be deployed on general Linux LPARs on older servers.

What type of appliances will be available for use in an IBM Secure Service Container partition on the z14?

IBM® Secure Service Container now integrates with Docker and open container technologies to provide a best of both worlds experience for your developers and operations teams, to simplify and standardize the deployment of cloud-native and cloud-ready container-based applications on the z14.

Plus, by leveraging the Secure Service Container technology, ISVs will be able to create enriched services and product offerings built on a secure container foundation.

There are several IBM appliances available today that use the IBM Secure Service Container. 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 IBM Secure Service Container, the IBM zAware V3.1 code packaged with Operations Analytics for IBM Z can be thought of as a software appliance versus a firmware appliance.

IBM Blockchain Platform is a fully managed blockchain service running on IBM LinuxONE in the IBM cloud delivering a secure, isolated compute environment ideally suited for workloads with sensitive data. IBM Blockchain Platform is running in an IBM Secure Service Container.

The z/VSE Network Appliance is designed for faster internal LPAR to LPAR communications between z/VSE and Linux environments. It runs within an IBM Secure Service Container.

Can I migrate my z13 appliances to z14?

Yes. A migration is possible with the IBM Secure Service Container Installer. With the Migrate feature in the Installer it is possible to attach the disks of the appliance to the new system if available. Please review the migration information for specific appliances.

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

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

What are the new enhancements for IBM Secure Service Container?

The IBM Secure Service Container has been expanded to integrate with the IBM Cloud Private platform for hybrid and private cloud deployments on IBM Z and LinuxONE. With IBM Cloud Private, clients can deploy containerized IBM Middleware applications as well as utilize common management tooling for deploying homegrown or other third party Docker and Kubernetes based applications. The integration of Secure Service Containers and IBM Cloud Private will provide differentiated security to containerized middleware and applications deployed in hybrid and private clouds. In addition, clients will be able to directly use the environment to manage and help to securely deploy their own Docker and Kubernetes based mission-critical applications for cloud-native, cloud-ready deployment.

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 IBM Z without requiring a TCP/IP stack on z/VSE.

The z/VSE Network Appliance is designed for faster LPAR to LPAR (or z/VM to LPAR) communications between z/VSE and Linux on IBM Z environments. It runs within an IBM Secure Service Container.

On z14 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 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.

For heavy loaded environments, the z/VSE Network Appliance may free up z/VSE resources by 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 the z/VSE Network Appliance 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 IBM Z server in an LPAR or under z/VM.

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

The z/VSE Network Appliance is supported on z/VSE V5.2, or later.

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

Contact IBM at email zvse@de.ibm.com in order to receive instructions to download the appliance.

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

Yes, starting with z/VSE V6.2. Older z/VSE systems running under z/VM can either use LFP with the z/VSE z/VM IP Assist (VIA) or a Linux on IBM Z server under z/VM.

IBM Dynamic Partition Manager

What is IBM Dynamic Partition Manager (DPM)?

IBM Dynamic Partition Manager is designed to perform simplified configuration of hardware resources 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 KVM on Z and/or Linux as a partition-hosted operating system.

What are the benefits of having IBM Dynamic Partition Manager (DPM)?

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 endpoint
- 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 (DPM) 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 (DPM) support?

IBM Dynamic Partition Manager supports Linux running virtualized with KVM or z/VM 6.4, 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 (DPM)?

Previously, IBM Dynamic Partition Manager only supported FCP storage devices. With the z14 announcement, we are announcing support for FICON ECKD storage devices.

What IBM Z mainframes support IBM Dynamic Partition Manager (DPM)?

IBM Dynamic Partition Manager is supported by IBM z14, IBM z13s, IBM z13, IBM LinuxONE Rockhopper (2965) and IBM LinuxONE Emperor (2964).

Anything I should be aware of when I switch z14 my server to IBM Dynamic Partition Manager (DPM)?

IBM Dynamic Partition Manager must be activated with an IML. Also you cannot have both IBM Dynamic Partition Manager and “classic” PR/SM mode IMLed on the same server.

What is new on IBM Dynamic Partition Manager (DPM) with the z14 announcement?

New support with IBM Dynamic Partition Manager is available for auto configuration of devices to simplify Linux Operating System Installation, where Linux distro installers exploit the function.

Also, Dynamic Partition Manager will support secure FTP through HMC for booting and installing an Operating system via FTP. This is really a capability of the new HMC code. For maximum security, we recommend that customers keep their IBM Z on a dedicated network with one HMC network used for that network and the second HMC network used for outward communication (IBM Support Facility, remote browsing, automation). However, for systems prior to z14, this created a security challenge for FTP operations originating from the SE from working. A customer can either put their FTP server on the IBM Z dedicated network or put their IBM Z on their intranet network. Starting with z14, all FTP operations originating from the SE will proxy through a managing HMC. This now allows the FTP SE originated operations to follow our security recommendation.

The new DPM 3.2 provides improvements to the workspace and managing system time, new security features that include Multifactor Authentication, and a new HMC Mobile application for monitor and recover action controls. Enhancements to DPM simplify the installation of the Linux operating system, and support additional hardware adapters. It provides support of import of FICON based ECKD DASD configuration data from an existing machine or data center, IOD enables a quick setup of the Storage Configuration when installing a new Linux machine.

I/O

Tell me about the new FICON Express16S+ for z14.

FICON Express16S+ is designed with a boost in I/O rates and a reduction in single stream latency to help absorb large application and transaction spikes driven by large unpredictable analytic and mobile workloads.

With the introduction of FICON Express16S+ on the z14, 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 (zHPF), and Fibre Channel Protocol (FCP), the z14 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.

The FICON Express16S+ family of adapters is exclusive to the z14.

What kind of increased performance can I expect to get on commercial batch elapsed time on the z14 with FICON Express16S+?

Batch Elapsed time improves 17% when using FICON Express16S+ on z14 running I/O intensive batch workloads compared to running the same workload using FICON Express16S on a z13.

Batch Elapsed time improves 35% when using FICON Express16S+ on z14 running I/O intensive batch workloads compared to running the same workload using FICON Express8S on a z13.

Disclaimer: This performance data was measured in a controlled environment. 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 with zHPF using FICON Express16S+ on z14?

The use of FICON Express16S+ on an IBM z14 with the zHPF protocol and a mix of large sequential read and write data transfer I/O operations, and operating at 16 Gbps achieved a maximum throughput of 3200 MB/sec (reads + writes) compared to a maximum of 2560 MB/sec (reads + writes) achieved with FICON Express16S operating at 16 Gbps. This represents an approximately 25% increase over the prior generation FICON Express16S on z13/z13s.

The use of FICON Express16S+ in an IBM z14 with the zHPF protocol and small data transfer I/O operations, and operating at 16 Gbps achieved a maximum of 314,000 IOs/sec. This represents 3x improvement over the prior generation FICON Express16S on z13/z13s.

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 on z14?

A FICON Express16S+ adapters, when defined as CHPID type FCP, conforms to the Fibre Channel Protocol (FCP) standard to support attachment of SCSI devices, to complement the classical storage attachment supported by FICON and zHPF channels.

The use of FICON Express16S+ in an IBM z14 with the FCP protocol for small data transfer I/O operations, operating at 16 Gbps achieved a maximum of 380,000 IOs/sec, representing a greater than 3x improvement over the maximum of 110,000 IOs/sec achieved with FICON Express16S operating at 16 Gbps.

The use of FICON Express16S+ in an IBM z14 with the FCP protocol and operating at 16 Gbps, FICON Express16S+ achieved a maximum throughput of 3200 MB/sec (reads + writes). This represents approximately a 25% increase in throughput over the prior generation FICON Express16S on z13/z13s.

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.

The FCP protocol is supported by z/VM, z/VSE, and Linux on IBM Z.

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.

Does the FICON Express16S+ continue to support Forward Error Correction (FEC)?

The FICON Express16S+ will continue to support the Forward Error Correction (FEC) feature that was introduced on the FICON Express16S on z13. The feature is a Fibre Channel 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 IBM Z 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.

On the IBM DS8880, FEC helps to preserve data reliability for Peer to Peer Remote Copies. IBM announced support for FEC starting with the DS8870 Storage Systems on May 11, 2015.

Does the z14 still support FICON Dynamic Routing?

With the z14 (and the z13s and z13), 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 IBM Z 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 product 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 DS8880 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 originally announced the support for FICON Dynamic Routing for DS8870 Storage Systems on May 11, 2015.

NOTE: FICON Dynamic Routing is not supported in multi-hop configurations.

What are the availability advantages on the z14 with zHPF Extended Distance II?

High Performance FICON for Z (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 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 originally announced the support of zHPF Extended Distance II for DS8870 Storage Systems on May 11, 2015.

How many channel subsystems (CSS) are supported on the z14?

The z14 Model ZR1 has channel subsystem (CSS) support for three logical channel subsystems (LCSSs), with support for forty LPARs, and three subchannel sets.

The z14 Models M01-M05 have channel subsystem (CSS) scalability with support for six logical channel subsystems (LCSSs) which are required to support the eighty-five LPARs for z14, four subchannel sets (to support more devices per logical channel subsystem), and 32K devices per FICON 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 DS8880 Multi-target Metro Mirror storage replication with TPC-R HyperSwap.

Does the z14 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 adapters include FICON Express16S+, FICON Express16S, 10GbE RoCE Express2, 25GbE RoCE Express2, OSA-Express7S 25 GbE SR, OSA-Express6S, OSA-Express5S, Coupling Express LR, Crypto Express6S and zEDC Express for compression.

When will the OSA-Express7S GbE SR card be available?

The OSA-Express7S GbE SR card has a planned availability of April 9, 2019.

Which of my PCIe I/O adapters can I continue to carry forward to the z14?

You can carry forward to the z14 the FICON Express16S, FICON Express8S, OSA-Express5, OSA-Express4S 1000BASE-T, 10GbE RoCE Express, zEDC Express, Crypto Express5S and Integrated Coupling Adapter SR (ICA SR). Clients that install the new Coupling Express LR can carry those forward from a z13/z13s to the z14.

What is the 25GbE RoCE Express2 adapter?

The 25GbE RoCE Express2 with z14 GA2 code now allows for exploitation of routable RoCE and better adapter sharing. This could provide flexibility for clients to consolidate adapters or increase overall speed, while still seeing high reliability and performance established on prior adapters. There is one PCIe adapter, one PCHID, and two ports per feature. The RoCE (Remote Direct Memory Access (RDMA) over Converged Ethernet) IBM offerings include also the (#0412) 10GbE RoCE Express2, and clients should take care that the benefit of the (#0430) 25GbE RoCE Express2 is only recognized when connected to another (#0430) 25GbE RoCE Express2.

How many 10GbE RoCE Express2 adapters, 25GbE RoCE Express2 (or 10GbE RoCE Express adapters) does the z14 support?

On the z14 Model ZR1, up to 4 adapters are supported (8 ports). On the z14 Models M01-M05, up to 8 adapters are supported (16 ports). Each LPAR should be provisioned with 2 PFIDs for redundancy, ideally with each PFID defined on two separate ports across two separate adapters.

Can the 10 GbE RoCE Express2 or the 25GbE RoCE Express2 adapter be shared between multiple LPARs on the z14?

The RoCE Express2 adapter provides increased virtualization (sharing capability) allowing RoCE to be extended to more workloads. The 10 GbE Express2 adapter or the 25GbE RoCE Express2 on the z14 Model ZR1 supports 31 virtual functions (VF) per port. The adapter on the z14 Models M01-M05 can support 63 VFs per port times 2 ports for a total of 126 VFs.

Can I mix 10GbE RoCE Express, 25GbE RoCE Express2 and 10GbE RoCE Express2 adapters in the same z14?

The 10GbE RoCE Express can be carried forward on a z14. You can mix 10GbE RoCE Express and 25GbE RoCE Express2 adapters with the new 10GbE Express2 adapter but you cannot exceed a total of 4 adapters on the z14 Model ZR1 and a total of 8 adapters on the z14 Models M01-M05.

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

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 FICON and FCP channels”.

Does the z14 continue to support the FCP SAN discovery tool?

Yes. 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 worldwide port naming (WWPN) assignments that is on the z14?

An IBM Z 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 z14 allows 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 by allowing the import/export the naming.

What are zHyperLinks on the z14 and DS8880?

The zHyperLink technology is a new mainframe attached link. It is the result of collaboration between Db2 for z/OS, the z/OS operating System, IBM z14 and DS8880 storage to deliver the extreme low latency I/O access for Db2 for z/OS Applications. zHyperLink Writes is intended to accelerate Db2 log writes to help meet clients most stringent requirements and deliver superior service levels by processing high volume Db2 transactions.

What is the expected value of low latency I/O for Db2 on z/OS Applications using IBM 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 CPC 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. However, the DS8882F rack mounted storage solution that can be integrated into into IBM z14 ZR1 or IBM LinuxONE Rockhopper II systems is not supported at this time.

What kind of results will I see with zHyperLink on a z14?

zHyperLink improves application response time, cutting I/O sensitive workload response time by up to 50% without requiring application changes.⁹

⁹ This response time estimate is based on IBM internal measurements and projections that assume 75% or more of the workload response time is associated with read DASD I/O and the storage system random read cache hit ratio is above 80%. The actual performance that any user will experience may vary.

Are there any considerations/guidance that you could provide for zHyperLink migration?

IBM zHyperLink requires z/OS V2.1 or later and Db2 V12. Software maintenance (APARs) required for z/OS and Db2 will be listed in PSP information.

IBM Virtual Flash Memory

What is Virtual Flash Memory on the IBM z14?

Virtual Flash Memory is replacement technology for Flash Express providing faster paging 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, Virtual Flash Memory 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 Virtual Flash Memory on the z14 improve the availability of my environment?

Virtual Flash Memory 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 – Virtual Flash Memory reduces these delays. In addition, z/OS support for writing pageable (1M) large pages to Virtual Flash Memory, can improve CPU performance and flexibility.

Which z14 servers support Virtual Flash Memory?

The Virtual Flash Memory is only supported on the z14. Prior to the z14, storage class memory was provided by the Flash Express adapter.

How do I order Virtual Flash Memory for the z14?

Virtual Flash Memory is a feature of IBM Z and is located within Random Array of Independent Memory (RAIM) on the IBM z14 (z14).

On the z14 ZR1 it is orderable in 4 sizes - .5 TB, 1 TB, 1.5 TB or 2 TB.

On the z14 M01-M05 models it is orderable in 4 sizes - 1.5 TB, 3.0 TB, 4.5 TB, 6.0 TB.

Note: Virtual Flash Memory does not require PCIe I/O slots so users of Virtual Flash Memory can free up I/O slots for other users such as accelerators, channels, network, etc.

Do I need to change my applications to use Virtual Flash Memory on the z14?

No you need not alter applications. Virtual Flash Memory 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 Virtual Flash Memory
- No definition required for placement of data on Virtual Flash Memory
- At IPL, z/OS detects if Virtual Flash Memory is assigned to the partition.
- z/OS automatically uses Virtual Flash Memory for paging unless specified otherwise via PARM LIB

What is the correlation between pageable large pages on the z14 and Virtual Flash Memory?

Pageable large pages require Virtual Flash Memory. 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.

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

The Virtual Flash Memory is designed to offer improved availability and performance for key workloads at critical processing times. Virtual 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 Virtual Flash Memory 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 Virtual Flash Memory. This is designed to allow structure data to be migrated to Virtual Flash 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.

Can I concurrently add Virtual Flash Memory to the z14 after the initial shipment?

Yes – but it depends on the existing z14 configuration.

To upgrade the z14 to add Virtual Flash Memory, two hardware steps need to happen. First the memory for the VFM needs to be ordered and installed. If there are enough processor drawers, this could be done concurrently, otherwise there will be an outage. Once the physical memory is on the z14 then VFM can be added concurrently and assigned to a z/OS partition concurrently. Once VFM is on the z system then VFM can be concurrently added to a running z/OS partition as long as the LPAR activation record is set up correctly.

Networking

What is the OSA-Express7S 25 Gigabit Ethernet (GbE short reach (SR) adapter for z14?

The OSA-Express7S 25 Gigabit Ethernet (GbE) short reach (SR) feature adds significant speed and a new NIC module unit, and benefits from prior level OSA-Express6S 10GbE performance and reliability. There is one PCIe adapter and one port per feature. This adapter allows for some potential consolidation of prior level adapters. 25 GbE can provide increased bandwidth for workloads. The OSA-Express7S family of adapters is exclusive to the z14.

What is the OSA-Express6S for the z14?

OSA-Express6S is an Ethernet adapters being used in the PCIe I/O drawer and continues to be supported by the 16 GBps PCIe Gen3 host bus. This includes an introduction of the full family of adapters - 1000BASE-T Ethernet for copper environments, in addition to 10 Gigabit Ethernet (10 GbE) and Gigabit Ethernet (GbE) for single-mode and multimode fiber optic environments. The performance characteristics are comparable to the OSA-Express5S adapters. They also retain the same form factor and port granularity - two ports per adapter for the 1000BASE-T Ethernet and Gigabit Ethernet adapters, and one port per adapter for the 10 Gigabit Ethernet adapter. The OSA-Express6S family of adapters is exclusive to the z14.

zEDC Express and Compression Acceleration

What is IBM Enterprise Data Compression (zEDC) for z14?

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

The zEDC exploitation with Linux requires one of the following Linux distributions for IBM Z: Canonical Ubuntu 16.04.03 LTS and later, Red Hat Enterprise Linux 7.3 and later, or SUSE Linux Enterprise Server 12 SP3 and later.

Isn't there already compression on every IBM Z processor chip – why is zEDC needed?

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.

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

How does zEDC work with zlib on my z14, 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.

The supporting z/OS and Linux releases provide the zlib library for lossless data-compression in a data format that in itself is portable across platforms. If the zEDC Express feature is available, the zlib library will exploit the ability to offload data compression and decompression tasks to the zEDC Express feature. This can accelerate the compression and decompression process significantly, and is fully transparent to the requesting application. Use of compression means that not only will clients see a substantial relief in CPU cycles otherwise used for those operations executed in software, but also see a performance gain in tasks heavily geared towards data (de)compression needs.

zEDC Express a great fit not only for system local use, but also for interaction with services that are implemented on a variety of architectures other than IBM Z and a multiplicity of different application environments hosted on Z or elsewhere.

Can you give me examples of z/OS and z/OS based functions and applications that exploit zEDC on the z14?

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 Security zSecure™ V2.3 (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 and later

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.

Can you give me examples of Linux for IBM Z based functions and applications that exploit zEDC on the z14?

On Linux for IBM Z, sweet spot workloads are large request sizes and products using gzip/deflate standard compression; zlib is also used by some network protocols and applications such as http/1.1, openssl, subversion, git, or PNG.

Sweet spot products include IBM Db2, IBM Java, IBM MQ, IBM WebSphere Application Server, as well as open source products such as Apache Kafka.

How does zEDC improve z/OS SMF logstream recording on the z14?

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.

How does QSAM/BSAM benefit from zEDC on the z14?

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.

How does IBM Encryption Facility for z/OS benefit from zEDC on the z14?

The IBM Encryption Facility for z/OS, 5655-P97, exploits zEnterprise Data Compression (zEDC) for z/OS, running on z14 with the zEDC Express adapter with the appropriate levels of IBM SDK Java for z/OS. zEDC can provide IBM EF users reductions in CPU time for environments where compression is already in use.

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

How does IBM Security zSecure V2.3 benefit from zEDC on the z14?

The audit capabilities of IBM zSecure V2.3 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.

How does IBM WebSphere MQ for z/OS V8 benefit from zEDC on z14?

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 on the z14?

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.

For z/OS to use zEDC on the z14, I need to order both hardware and software - why?

To use zEDC, you need to order both the hardware zEDC Express adapter 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 adapter for PCIe I/O drawer.

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. 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 adapter be shared between multiple LPARs on z14?

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

How many zEDC Express adapters will I need to have on my z14?

The minimum purchase is one. IBM recommends at least two zEDC Express adapters for high availability. Four zEDC Express adapters 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.

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

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 hardware adapters you will need. It is available from:

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

What z/OS data is provided to help me understand my system's use of zEDC on z14?

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.

Does z/VM support zEDC on the z14?

z/VM guest exploitation support for the zEDC Express adapter was added by z/VM 6.3 and is available on the z14, z13s, z13, zEC12 and zBC12 servers, with up to 15 z/VM guests sharing the zEDC Express adapter. 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 on z14?

z/VM does not directly benefit from zEDC. The z/VM support gives z/OS and Linux 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 adapter to be exploited by guests that have PCIe drivers.

Is the zEDC Express adapter supported by Linux for IBM Z on the z14?

zEDC Express support is included in SLES12 SP2, RHEL 7.3, Ubuntu 16.04, and their later Linux-Distro versions.

Parallel Sysplex

What long range coupling link options do I have with the z14?

For long range coupling connectivity, the z14 introduces the new Coupling Express Long Range (CE LR) links. This CE LR link is concurrently available on, and compatible with z13 and z13s systems, which allows for sysplex coexistence back to these systems. Unlike the 1x IFB Infiniband® coupling links, which were plugged into the processor drawer, the CE LR link is plugged within a PCIe I/O drawer slot, taking advantage of more industry standard I/O technology, such as converged Ethernet.

The location inside the PCIe I/O drawer allows users to fan links out across multiple 2-port CE LR cards, up to a maximum of 32 adapters (64 links maximum) for z14 and z13 servers, and up to a maximum of 16 adapters (32 links maximum) for z13s, servers in 2 link increments. Compared to the HCAO-3 IFB Infiniband LR with 4-port and 4-link increments, the CE LR link allows for more granularity when scaling up or completing maintenance. Link performance is similar to the InfiniBand 1x coupling link and uses identical Single Mode fiber. The CE LR link provides point-to-point coupling connectivity at distances of 10 km unrepeated and 100 km with a qualified dense wavelength division multiplexing (DWDM) device.

The z14 Models M01-M05 also support the LR InfiniBand coupling link and will be the last IBM Z mainframe models to support HCA3-O LR 1x IFB InfiniBand coupling links. Clients are strongly encouraged to begin migrating to the CE LR coupling links to ensure long distance coupling connectivity to future server generations.

The z14 Model ZR1 does not support any InfiniBand links.

What short range coupling link options are available for the z14?

The IBM Integrated Coupling Adapter SR (ICA SR) was first introduced on the z13, and is now also available on the z14. The ICA SR 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).¹⁰

The ICA SR fanout resides in the PCIe fanout slot on the z14 CPC drawer. The ICA SR can only be used for coupling connectivity between z14 and z13s and/or z13 servers. The ICA SR is designed to support up to 4 CHPIDs per port and 8 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 z14 Model ZR1 supports two PCIe fanout slots per CP chip. The number of PCIe fanouts supported is feature dependent. For the Max4 feature there will be two fanout slots. The Max12 will have 4 fanouts. The Max24 and Max30 will have 8 PCIe fanouts.

The z14 Models M01 – M05 support 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. Those models also support the

¹⁰ Note: A link data rate of 8 GBps does not represent the actual performance of the link.

InfiniBand coupling links and will be the last IBM Z mainframe models to support HCA3-O 12x IFB InfiniBand coupling links. Clients are strongly encouraged to begin migrating to the ICA SR to ensure short distance coupling connectivity to future server generations. There is no InfiniBand support on the z14 Model ZR1.

Can you review the Infiniband options on the z14 Models M01-M05?

On the z14 Models M01-M05, there are two InfiniBand coupling links.

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, also available on the z13 and z13s. The HCA3-O fanout has two ports/links and is compatible with the HCA2-O fanout on zEC12 and zBC12.

1x InfiniBand is 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 also available on the z13 and z13s. The HCA3-O LR fanout has four ports/links and is compatible with the HCA2-O LR fanout on zEC12 and zBC12 machines, which has two ports/links.

For z14 Models M01-M05, the maximum number of all HCA3 fanouts are limited to 16 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.

Where can I order cables for the z14?

IBM strongly recommends clients order cabling 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 z/OS Parallel Sysplex support for Asynchronous Coupling Facility (CF) Duplexing for Lock Structures?

A z/OS Parallel Sysplex is a tightly coupled cluster of operating system instances that cooperate to share data and process in parallel against a common client workload. CF lock structures provide a sysplex-wide point of serialization for controlling concurrent access to shared data in the cluster, which is essential to preserve data consistency and integrity across the cluster. In October 2016, IBM announced the optimization of the duplexing protocol between the CFs to make it asynchronous. Now Parallel Sysplex applications can obtain or release locks in one CF structure, without having to wait for synchronized updates to occur in the other CF structure.

Additional information can be found in the announcement letter: [116-093](#).

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

Yes. There is a good blog written by Dave Surman, an IBM Distinguished Engineer specializing in Parallel Sysplex.

<http://mainframeinsights.com/asynchronous-coupling-facility-cf-duplexing-for-lock-structures/>

What does Dynamic I/O for Standalone Coupling Facility provide?

Dynamic I/O for Standalone Coupling Facility capability Coupling Facilities (CFs) provide locking, caching, and list services between coupling-capable z/OS processors, and are a significant component of highly-available Parallel Sysplex configurations. Standalone CFs (which are, by definition, Coupling Facility images that reside on a CEC without a co-resident z/OS image) have, until now, been unable to participate in Dynamic I/O configuration changes that affect the standalone CF CEC, requiring the CEC to be re-IMLed or PORed in order to make such changes. This causes disruption for the Parallel Sysplex that is using the Standalone CF, as the CF image contents must be relocated other CF images.

With z14 GA2, standalone CF CECs can now seamlessly make hardware-only dynamic I/ O configuration changes on behalf of the CF partitions that reside there without requiring a disruptive re-IML or POR. This capability both improves client workload availability and minimizes the risks associated with relocation of CF structures and disruptive actions such as a CEC IML or POR. This enhancement requires z14 GA2 firmware support for the standalone CF CEC and it requires an enablement that is done during the IML of the standalone CF CEC. This enhancement also requires z14 GA2 firmware support both on the CF CEC and the CEC where the driving z/OS HCD system is running, however, no enablement IML is required on the driving system.

Note that standalone CFs generally do not have any external connectivity other than the coupling links used for Parallel Sysplex CF request traffic. This enhancement does not place any additional connectivity requirements on the standalone CF CEC.

What are the new scalability enhancements on the z14 for Coupling Facility processors?

There is better utilization of Coupling Facility (CF) processors with new scalability improvements on the z14. In the past, CF images have not scaled up well. CF work management and dispatching changes allow improved efficiency and scalability for coupling facility images.

These changes include ordered work queues have been eliminated from the CF in favor of first-in / first-out queues, avoiding the overhead of maintaining ordered queues. Protocols for system-managed duplexing have been simplified to avoid the potential for latching deadlocks between duplexed structures. And the CF image is now able to use its processors to perform specific work management functions when the number of processors in the CF image exceeds a threshold.

Asynchronous Cache Cross Invalidation (XI) is a new sysplex capability for performance/ scalability and improved cross-site operation. This new function allows the cache coherency messages that flow around the sysplex to maintain data integrity to be performed in an asynchronous fashion rather than synchronously, with exploitation from the data manager (Db2). The data manager must provide support to sync up with the asynchronous cross-invalidate messages at critical points in its processing, such as at transaction commit. The asynchronous protocol is expected to reduce cache services times and sysplex coupling overhead, particularly in sysplex environments which involve multiple sites with significant cross- site distances involved. The asynchronous protocol avoids distance latencies associated with the communication of XI messages across inter-site distance. This enhancement requires PTFs for z/OS and Db2.

STP CTN split/merge is a new sysplex timing capability for availability that allows two distinct timing

networks to be merged into one, or to split one timing network into two, nondisruptively. Previously, these timing network reconfigurations and transitions were disruptive to the running sysplex(es). This feature is especially helpful when working to combine or redistribute servers within a corporate structure.

What are the new CF list structure notifications on the z14?

There are significant enhancements in z13 to CF notifications that inform users about the status of shared objects within a Coupling Facility. The combination of these notification enhancements provides flexibility to accommodate notification preferences among various CF exploiters, and yields more consistent, timely notifications.

What is new for CF diagnostics on the z14 to assist in faster resolution of problems with the CF?

CF diagnostics will include new information about both successful and unsuccessful requests helping to improve diagnostic information and serviceability. The information will be written to z/OS logs and traces. This new information will include CF CHPID user for operation, as well as additional information describing errors such as interface control checks and channel control checks. The information should help to facilitate faster problem resolution.

How did the z14 provide Coupling Link constraint relief?

The IBM z14 Models M01-M05 provide additional physical and logical coupling link connectivity compared to z13. There is a 2X increase in physical ICA SR coupling links (ports) -- increased from 40 per CPC to 80 per CPC. These higher limits on z14 models can help in coupling link technology migration.

What is the value that Virtual Flash Memory can play in supporting the Coupling Facility Control Code (CFCC) on the z14?

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

With enhancements in CFCC, exploitation of Virtual Flash Memory 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 Virtual Flash Memory. This is designed to allow structure data to be migrated to Virtual Flash Memory as needed and migrated back to real memory to be processed. When using WebSphere MQ for z/OS Version 7 (5655-R36), this new capability is expected to help provide significant buffering against enterprise messaging workload spikes and to help provide support for storing very large amounts of data in shared queue structures, potentially allowing several hours of data to be stored without causing interruptions in processing.

What is Coupling Thin Interrupts for the z14?

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.

What other generations of servers can connect to a z14 server in a Parallel Sysplex environment?

The z14 servers are designed to coexist in the same Parallel Sysplex environment with (n-2) server families. This allows it to coexist with other z14, z13, z13s, zEC12 and zBC12 servers. You should also ensure that your z13, z13s, zEC12 and zBC12 servers are at the latest driver level.

I have a Parallel Sysplex and would like some assistance migrating to the CE LR coupling links, are there resources to help me through this migration?

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 does the z14 keep time synchronized with other IBM Z mainframes?

The z14 does not support attachment to the Sysplex Timer[®]. However, if you require time accuracy either for one or more IBM Z mainframes or you require the same time across heterogeneous platforms (IBM Z, UNIX, AIX[®], etc.) you can use Server Time Protocol (STP). STP is designed to allow events occurring in different servers to be properly sequenced in time.

What STP enhancements were announced on the z14?

There are a few enhancements being announced for STP.

The z14 will introduce an additional stratum level 4 for Server Time Protocol (STP) synchronization. With the additional stratum level, STP can synchronize systems up to 3 steps away from the Current Time Server (CTS). Prior systems only allowed synchronization up to level 3, or up to 2 steps from the CTS. This additional stratum level is not intended for long-term use; rather, it is specifically intended for short-term use during configuration changes for large timing networks, to avoid some of the cost and complexity caused by being constrained to a 3- stratum timing configuration.

The z14 will also introduce a new Graphical User Display for the STP network and configuration. The new user interface has been revamped for a quick, intuitive view of the various pieces of the STP relationship network map, including the status of the components of the timing network. The new z14 support allows the new level of HMC to manage older systems using the same new interface.

I have a Parallel Sysplex and would like some assistance migrating to the CE LR coupling links, are there resources to help me through this migration?

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

z/OS

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

Please visit: ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSQ03081USEN

z/VSE

What's in the z14 announcement for z/VSE users?

- Exploitation of latest hardware technology with z/VSE V6.2
- Opportunity to extend your workloads with Linux on IBM Z

What are my z/VSE software (pricing) savings on z14 Model ZR1?

Coinciding with the announcement of the z14 Model ZR1 server, IBM is making available a new Technology Transition Offering (TTO) called Technology Update Pricing for the z14 Model ZR1. Technology Update Pricing for the z14 Model ZR1 offers price-performance advantages for z14 Model ZR1 servers. In addition, zSeries Entry License Charge (zELC) pricing applies to the z14 Model ZR1 capacity setting A01, the entry z14 Model ZR1 server.

Technology Update Pricing for the z14 Model ZR1 extends the software price-performance provided by Advanced Entry Workload License Charges (AEWLC) for stand-alone z14 Model ZR1 servers, and applies to eligible z/VSE releases and their associated middleware programs, subject to all applicable terms and conditions.

For more information, refer to <https://ibm.com/it-infrastructure/z/software/pricing>.

Does z/VSE support the Crypto Express6S adapter on the z14?

Yes. Crypto Express6S on the z14 is supported on z/VSE V6.2, V6.1 and V5.2 with PTFs.

Can z/VSE use Crypto Express6S support of greater than 16 domains on the z14?

Yes. Crypto Express6S can support greater than 16 domains on the z14 with z/VSE V6.2, V6.1 and V5.2 with PTFs.

Does z/VSE support the FICON Express16S+ adapter on the z14?

Yes. FICON Express16S+ on the z14 is supported by z/VSE.

Does z/VSE support High Performance FICON for IBM Z (zHPF)?

Yes. z/HPF is supported by z/VSE V6.2 with PTFs.

Does z/VSE support the OSA-Express6S adapter on the z14?

Yes. OSA-Express6S on the z14 is supported by z/VSE V6.2, V6.1 and V5.2

Does z/VSE support the Vector Facility for z/Architecture® of a z14 (SIMD)?

Yes. SIMD on the z14 is supported by z/VSE V6.2 with PTFs.

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

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: ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=LUQ12358USEN

Where can I find information about the preview of z/VM V7.1?

Please visit: ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=LUQ12358USEN

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

Please visit: ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSQ03067USEN

KVM

Where can I find information about KVM running on IBM Z?

KVM on IBM Z can be offered by the Linux distribution partners, integrated into their Linux distribution. Please contact the distribution partners for detailed information about their KVM implementation.

IBM is committed to the KVM hypervisor and is responsible for the architecture and exploitation of the IBM Z hardware in the Linux upstream code using the same approach as for Linux.

Information about the KVM base technology for the IBM Z is available at:

ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_kvm_base.html

Linux on IBM Z

Where can I find a full set of Linux on IBM Z frequently asked questions?

Please visit: ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSQ03011USEN

What are the IBM tested and supported Linux distributions for Linux on Z?

Canonical, Red Hat and SUSE will support z14 with their Linux distributions. The IBM tested and supported Linux distributions are shown at the “Tested platforms” web page, please visit ibm.com/it-infrastructure/z/os/linux-tested-platforms

Compilers

What advantage does the IBM z14 have for COBOL?

IBM Enterprise COBOL for z/OS V6.2 reduces CPU usage of decimal and floating-point compute intensive applications by up to 38%, and on average by 19%, on z14 over the same applications built with Enterprise COBOL for z/OS V6.1 on z14¹¹.

IBM Enterprise COBOL for z/OS V6.2 reduces CPU usage of compute intensive applications on average by 10% on z14 over the same applications built with Enterprise COBOL for z/OS V6.1 on z14.

IBM Enterprise COBOL for z/OS V6.2 reduces CPU usage of decimal and floating-point compute intensive applications by up to 67%, and on average by 30%, on z14 over the same applications built with Enterprise COBOL for z/OS V5.2 on z14¹².

IBM Enterprise COBOL for z/OS V6.2 reduces CPU usage of compute intensive applications on average by 16% on z14 over the same applications built with Enterprise COBOL for z/OS V5.2 on z14.

All of these significant reductions in CPU usage are achieved through improvements in optimization and compiler exploitation of the new Vector Packed Decimal facility. The Vector Packed Decimal Facility in z14 allows the dominant COBOL data types, packed and zoned decimal, to be handled in wide 16 byte vector registers instead of in memory, resulting in significant improvements in COBOL computational performance.

What are the C/C++ Compiler benefits of the z14?

z/OS V2R3 XL C/C++ compiler demonstrates on average 13% reduction in CPU time for floating point intensive applications on z14 over the same applications built with z/OS V2R2 XL C/C++ compiler on z14. z/OS V2R3 XL C/C++ compiler demonstrate on average 8% reduction in CPU time for compute intensive applications on z14 over the same applications built with z/OS V2R2 XL C/C++ compiler on z14.¹³

¹¹ The performance improvements are based on internal IBM lab measurements on a z14 701 with 8GB of central storage. All benchmarks compiled with Enterprise COBOL for z/OS V6.2 used the options OPT(2), STGOPT, AFP(NOVOLATILE), HGPR(NOPRESERVE), ARCH(12), LIST. All benchmarks compiled with Enterprise COBOL for z/OS V6.1 used the options OPT(2), STGOPT, AFP(NOVOLATILE), HGPR(NOPRESERVE), ARCH(11), LIST. Performance results for specific applications will vary, depending on the source code, the compiler options specified, and other factors.

¹² The performance improvements are based on internal IBM lab measurements on a z14 701 with 8GB of central storage. All benchmarks compiled with Enterprise COBOL for z/OS V6.2 used the options OPT(2), STGOPT, AFP(NOVOLATILE), HGPR(NOPRESERVE), ARCH(12), LIST. All benchmarks compiled with Enterprise COBOL for z/OS V4.2 used the options LIB, OPT(FULL), LIST. Performance results for specific applications will vary, depending on the source code, the compiler options specified, and other factors.

¹³ The performance improvements are based on internal IBM lab measurements. All CPU intensive integer and floating-point benchmarks were compiled in 31-bit addressing mode and built using XPLINK, HGPR, O3, HOT and IPA LEVEL(2) with PDF compiler options. The benchmarks compiled with the z/OS V2R2 XL C/C++ compiler were built using the ARCH(11) TUNE(11) options; the benchmarks compiled with the z/OS V2R3 XL C/C++ compiler were built using the ARCH(12) TUNE(12) options; All benchmarks were executed on a z/OS V2R2 dedicated LPAR with 1 CP and 8GB Central Storage on z14. Performance results for specific applications will vary, depending on the source code, the compiler options specified, and other factors. Disclaimer: The performance improvements are based on internal IBM lab measurements. All CPU intensive integer and floating-point benchmarks were compiled in 31-bit addressing mode and built using XPLINK, HGPR, O3, HOT and IPA LEVEL(2) with PDF compiler options. The benchmarks compiled with the z/OS V2R2 XL C/C++ compiler were built using the ARCH(11) TUNE(11) options; the benchmarks compiled with the z/OS V2R3 XL C/C++ compiler were built using the ARCH(12) TUNE(12) options; All benchmarks were executed on a z/OS V2R2 dedicated LPAR with 1 CP and 8GB Central Storage on z14. Performance results for specific applications will vary, depending on the source code, the compiler options specified, and other factors.

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](#).

Or perhaps you'd like to take a class at a local school. If so, you should check out our vast [global network of schools](#) that are a part of the IBM Z Academic Initiative. These schools offer a variety of courses and curriculum across all areas of operating systems and application development. You can find the list on our Academic Initiative website.

IBM also provides no fee community resources that are available through the On The Hub., including [44 hours of free e-learning](#) where students can earn a Z Kickstart Open Badge.

Finally, you can connect to the growing community of mainframers young and old through the Advocacy Hub. There you will find information, games, challenges and education while connecting to the broader community. Go to ibm.biz/zstudent to get started.

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

Absolutely! We have hub systems for faculty and students to use in the classroom and community systems for developers to use just for playing around for a free trial period. Students and faculty should go to [On The Hub](#) to register for the z/OS or Linux system you need. Anyone can take advantage of the 120 day trial of Linux on z at LinuxONE Community Cloud.

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](#).

If you just want to play around, check out the free 120 day trial at the [LinuxONE Community Cloud](#).

I have new hires joining my company and need to get them trained on z, what resources are available to help?

IBM Offers the following resources to help train new hires on IBM Z.

- The [IBM Skills Gateway](#). Explore [Learning Journey roadmaps](#) for training on z/OS, Linux on Z, z/VM, KVM and our Z hardware.
- [Master the Mainframe Learning Management System](#). No mainframe skills are required to start learning z with this fun and engaging training. Build mainframe skills as you progress through three stages of self-paced learning.
- [z/OS Introduction and Workshop](#). This no-cost, five day hands-on workshop is offered quarterly and teaches z/OS Fundamentals. Check the website regularly for newly scheduled classes.
- [z/OS System Programming Certificates](#) offered by Marist College's Institute for Dace Center Professionals. Earn certificates in z/OS Systems Programming (Associate, Professional, and Expert), COBOL Application Programming, Db2 Application Programming, IMS™ Application Programming, and Assembler Language Application Programming.

For a comprehensive listing of resources available to clients to help attract, train, and retain z talent, visit our [Employer's Reference Guide](#).



© Copyright IBM Corporation 2018

IBM Corporation
New Orchard Road
Armonk, NY 10504
U.S.A.
10/2018

IBM, IBM logo, IBM Z, AIX, CICS, Db2, DFSMS, DFSMDdss, DFSMSHsm, ECKD, FICON, GDPS, HiperSockets, HyperSwap, IMS, InfiniBand, LinuxONE, LinuxONE Emperor, LinuxONE Rockhopper, MQ, Parallel Sysplex, PR/SM, Resource Link, RMF, Sysplex Timer, System z10, WebSphere, z13, z13s, z14, z/Architecture, zEnterprise, zHyperlink, z/OS, zSecure, z Systems, z/VM and z/VSE are trademarks or registered trademarks of the International Business Machines Corporation.

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

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

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

IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

ITIL is a Registered Trade Mark of AXELOS Limited.

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

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

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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

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

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

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

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

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

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

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

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