
IBM Z
Introduction
September 2019

IBM z15
And IBM z14
Frequently Asked Questions

Worldwide



Table of Contents

IBM z15 Hardware	3
Performance.....	9
z15 Warranty	13
z15 Software Pricing.....	14
Physical Planning for z15	18
I/O	22
IBM z15 Networking.....	25
IBM Virtual Flash Memory	26
Integrated Acceleration for zEDC and zEDC on z/OS	28
z/OS.....	31
z/VSE.....	32
z/VM	33
KVM.....	34
Linux on IBM Z	35
IBM Dynamic Partition Manager	36
IBM z14 Hardware	38
z14 ZR1 Warranty – additional information.....	47
Water cooling option for z14 Models M01-M05	48
Academic Initiative and Skills	50

IBM z15 Hardware

Tell me more about the IBM z15™ (z15).

The IBM z15 is the next generation IBM Z. The z15 is focused on delivering a system designed for secure digital transformation and hybrid cloud.

The IBM z15 delivers the platform for a hybrid cloud that has the power, availability, agility, and the speed users demand. And incorporates the security users and regulators require. Wrapped in a packaging the brings along operational efficiencies to help curb the skills gap. With a comprehensive portfolio of IBM solutions, the z15 can deploy and support a secure cloud environment, as well as expand access to it.

IBM z15 delivers the data privacy and business continuity for mission critical cloud workloads. The z15 extends pervasive encryption beyond the border of the IBM Z environment with a goal of protecting data beyond the platform and into distributed and hybrid cloud environments.

IBM z15 is built for a secure, "always on" world because clients need assurance that their data is safe, where services are always on and fulfilled instantly. The z15 can help clients adapt to planned or unplanned events while keeping services and operations running smoothly and continuously, whether on premise or in the cloud.

What is this IBM System Recovery Boost?

Built into IBM z15, System Recovery Boost is an innovative solution that diminishes the impact of downtime, planned or unplanned, so you can restore service and recover workloads substantially faster than on previous IBM Z generations with zero increase in IBM software MSU consumption or cost.

With System Recovery Boost, you can unleash additional processing capacity during a temporary performance increase known as the boost period. By enabling general-purpose processors to run at full-capacity speed, and by allowing general-purpose workloads to run on zIIP processors, the boost period accelerates the entire recovery process in the image(s) being boosted. By decreasing the time it takes to shut down, restart and process backlog, System Recovery Boost enables you to optimize your planned maintenance strategy, ensure service-level attainment, and accelerate execution of DR testing and site switches.

For more Information check out the Systems Recovery Boost FAQs:

<https://www.ibm.com/downloads/cas/1NWEJKOX>

What is new with pervasive encryption on the z15?

With Pervasive Encryption in IBM z14™, IBM enabled the encryption of enterprise data on Z at scale without requiring any application changes. With pervasive encryption we were able to capture the imagination of clients, analysts, and the mainstream media, because we brought something really unique and different to the table.

Now, with z15, IBM is expanding our Pervasive Encryption story to not only protect data on the platform, but to deliver new value for our clients, by meeting one of the greatest growing demands in the market – privacy. We went one step further because part of the feedback we received was, “Well it’s great that you’re protecting the data, on the Z platform but what can you do to actually help me protect that data wherever it goes around my hybrid cloud?” Since doesn’t always just live on Z there is the need to not

only protect data on the platform, but also protect and enforce data as it moves throughout the enterprise.

Bottom line the z15 will offer protection of z data wherever it resides – We’re doubling down on our security leadership to ensure not only is data protected on your system of record, but wherever it is needed throughout the hybrid cloud environment. IBM is announcing IBM Data Privacy Passports, a consolidated data security solution that protects data after it leaves the system of record, minimizing the risk of security breach, potential noncompliance and financial liability.

What is IBM Data Privacy Passports?

IBM Z Data Privacy Passports, in conjunction with IBM z15 and available via an IBM z15 only PID, is being designed to enforce security and privacy protections to data not only on Z, but across platforms. It provides a data-centric security solution that enables data to play an active role in its own protection.

For more information about IBM Z Data Privacy Passports V1.0 beta program announced on September 12, 2019 refer to Software Announcement 219-452.

How much capacity will the z15 deliver?

The largest IBM z15 is expected to provide approximately 25% more capacity than the largest IBM z14 (z14) with some variation based on workload and configuration. ¹

Is the one frame IBM z15 a replacement for the IBM z14 ZR1?

No. The one frame system is not the replacement for z14 ZR1. It is still the performance optimized, high frequency system. A single frame z15 can have up to 3 drawers or up to 108 cores or significantly more capacity than the z14 ZR1.

What is the machine type of the z15?

The IBM z15 machine type is 8561.

How many models are they for the z15?

There is only one model type for the z15 – the Model T01.

The z15 will follow the z14 ZR1 concept of feature-based sizing instead of models. There will only be one Model – the T01 – and there will be 5 new feature codes for ordering capacity – the Max34, Max71, Max108, Max145 and Max190.

Similar to the Model structure we’ve had in the past, it will be possible to order less configurable CPs than the Max setting. Thus allowing configuration of additional cores such as IBM z Integrated Information Processors (zIIPs), IBM z Integrated Facility for Linux (IFL) processors, additional System Assist Processors (SAPs), Internal Coupling Facilities (ICFs) and/or used as additional spares.

For example: a client requiring 20 Central Processors (CPs), 40 zIIPs and 5 IFLs – a total of 65 configured cores. The client could choose a Max71 or Max 108 depending on the growth they were expecting.

¹ Based on preliminary internal measurements and projections and compared to the z14. Official performance data will be available upon announce. Results may vary by customer based on individual workload, configuration and software levels. Visit LSPR website for more details at: <https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex>.

What is the available subcapacity granularity on the z15?

The z15 can have up to 34 subcapacity Central Processors (CPs) per server while the z14 accommodates up to 33 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.

Up to 34 sub-capacity CPs can be active on the server regardless of feature code. In other words, sub-capacity configurations as long as the server is configured (not necessarily the same as purchased) with thirty-four or fewer general-purpose processors.

For more information refer to the IBM z15 datasheet: <https://www.ibm.com/downloads/cas/NN7GBPJ1> and the IBM z15 spec sheet: <https://www.ibm.com/downloads/cas/2GO7LRDV>

I have lots of spare racks in my datacenter, can I supply the rack for the IBM z15?

No, the z15 will still be fully assembled, and tested at an IBM plant prior to being packaged for shipping to our clients as completed, tested, individual frames requiring minimum assembly on site. The frame(s) are part of what makes up the complete system for regulatory compliance.

What upgrades are there for the IBM z15?

It is possible to upgrade any IBM z14 Model M01-M05 or IBM z13 (any model) to the IBM z15.

There are some nuances to the z15 to z15 upgrades that you need to know.

- There are concurrent upgrades from Max34 to Max 71 to Max108. ²
- Each Max level adds a CPC (Central Processor Complex) drawer.
- There are no upgrades into the Max 145 or Max190. These are factory build only.
- There is no upgrade from a water-radiator cooled z15 to a customer-water cooled z15 – or vice versa. You cannot change a customer-water cooled z15 to a water-radiator cooled z15.

Any IBM LinuxONE™ III can be upgraded to a radiator-based air cooled z15.

There are no upgrades between from an IBM z15 into the LinuxONE III models.

Can any customer supplied infrastructure be added to the IBM z15 (similar to the 16U Reserved feature code on the IBM z14 ZR1)?

No. There is no reserved space in the IBM z15 to support adding other features such as storage, switch, HMC, etc.

What is the HSA and how much is there on the z15?

The hardware system area (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

² Please work closely with your IBM or BP Z technical representative as you create your configuration. If you need to leave room for CPC upgrade growth, there are 'reservation' feature codes (FC2271 and FC2272) to 'hold' space in Frame A for future CPC growth. If that space has not been reserved, I/O Drawers will be placed in that area - potentially reducing the need for another frame for I/O Drawers.

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 z15 there is 256 GB of HSA, independent of client-purchased memory.

How much memory is available on the z15?

The z15 server supports up to 40TB 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 Max feature code of the system (thus the number of CPC drawers). The minimum initial amount of memory that can be ordered is 512 GB for all models.

What is RAIM?

The z15 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 experience if I purchase larger memory on the z15?

Memory on the z15 will benefit many types of users. Large memory can reduce latency and CPU cost, and thus improve 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 rates.

z/VM® 6.4 and above 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 on IFLs or under z/VM, may see performance benefits when taking advantage of large shared, virtualized memory.

What is the integrated firmware processor (IFP) in z15?

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 and Coupling Express LR. It is not customer usable or 'visible'.

How many SAPs are on the z15?

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

- The z15 Max34 has four SAPs.
- The z15 Max71 has eight SAPs.
- The z15 Max108 has twelve SAPs.
- The z15 Max145 has sixteen SAPs.
- The z15 Max190 has twenty-two SAPs.

Up to 8 more SAPs may be acquired from the pool of available processing units within the model.

How many spare processing cores are on the z15?

IBM ships every z15 machine with a minimum of two spare processing cores. These spares can be shared across the drawers. The z15 offers core-level (engine-level) sparing.

Can any of the z15 spare processing cores be used for other purposes?

No, the minimum allotment of z15 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 unallocated cores above the minimum allotment can be activated for other purposes.

Can I buy a z15 that has only IFL or ICF processors without including a general-purpose processor (CP)?

Yes. You can order only IFLs or ICFs in a z15, using a model capacity identifier of 400, with a maximum of 190 IFLs or ICFs. There is still a limit of 16 ICF engines for a single Coupling Facility LPAR.

Besides using the capacity identifier of 400, IBM does offer Linux only servers based on z15 technology – the IBM LinuxONE III. The LinuxONE™ systems may optionally add one CP but it can only be used for the GDPS® Appliance.

Tell me about the new frame on the z15.

The IBM z15 is built with a 19” form factor that flexibly scales from 1 – 4 frames depending on the system configuration. This means that many clients will actually see a significant reduction in the amount of floor space taken up within the data center. More importantly, this system can now fit within most data center aisles – offering new flexibility in data center layout, space efficiency, cabling efficiency and participation in new air-cooling containment systems.

Client feedback has suggested reducing system footprint, particularly system depth, from predecessor z systems to maximize space savings and flexibility while maintaining future system expansion capability was a win. The fact that z15 can deliver up to 34 high performance processors and 48 I/O slots, or 71 high performance processors and 32 I/O slots, or 108 high performance processors and 16 I/O slots in a single frame, is a perfect example of this. (See physical planning section for more information.)

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

How can I quickly understand what I need to do in order to migration to the new z15 system and begin taking advantage of the great new features and capabilities?

Maintaining currency is critical to obtaining support from software vendors, exploiting new innovative features of z15 exploit new hardware, and maintain compliance with government/industry regulation. IBM Systems Lab Services has been providing a Currency and Migration Services offering for numerous IBM z generations. The offering is typically customized to the needs of the individual client and can include a quick assessment of IBM z/OS® and IBM Subsystem Products that are within the supported migration releases. It can also include assessments for KVM, z/VM and Linux on IBM Z. You can contact Systems Lab Services directly via e-mail at ibmsls@us.ibm.com or contact your IBM or Business Partner Z Sales Representative.

I recognize the importance of security and see the compelling value that IBM z Pervasive Encryption and the new IBM Data Privacy Passports can provide me, how can IBM help me get started and strengthen my company's data security position?

IBM Systems Lab Services provides global proven expertise in the IBM Z security segment and offers a portfolio of engagement offerings that help clients assess their current data security posture, plan, implement and exploit the strong IBM Z security capabilities including Pervasive Encryption, Data Privacy Passports as well as RACF® and base security z/OS and Linux functions. 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.

Performance

What are the major changes to the z/OS V2R3 LSPR?

The LSPR ratios reflect the range of performance between IBM Z mainframes as measured using a wide variety of application benchmarks. 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/lsprindex>

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 V2R3 LSPR table?

The 2094-701 processor model is used as the base in the z/OS V2R3 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 V2R3 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 V2R3 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 V2R3 multi-image table by 559.792. The PCI (Processor Capacity Index) column in the z/OS V2R3 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 z15 processor?

As with the introduction of any new server, workloads with differing characteristics will see variation in performance when moved to an IBM z15. 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 z15 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 z15, 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 z15 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 z15.

How do I get performance information for my TPF products running on an IBM z15?

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 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/atsmastr.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 z15?

The performance ratios a z/VM customer workload might experience when migrating to an IBM z15 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 z15 customer might experience?

The performance ratios a Linux customer workload might experience when migrating to an IBM z15 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 z15?

The performance ratios that a z/VSE® customer workload might experience when migrating to an IBM z15 are represented by the range of ratios for a comparable z/OS migration. For example, the published ratio in the LSPR between the z14 702 and the z15 702 is approximately 11% to 15%. z/VSE workloads should expect this same range of performance for this migration. Consult the LSPR for other examples of moves to an IBM z15.

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

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

z15 Warranty

What is proprietary diagnostic support for z15?

A z15 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 z15 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 z15 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 z15?

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 z15?

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

If I am on base support of a z15, 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 z15 is no longer on an IBM maintenance agreement?

Yes, if the machine is no longer on warranty or does not have a valid IBM maintenance agreement in place.

Is there a higher level warranty (or Service Level Agreement) available?

Yes, through an offering called IBM Z Warranty Service Upgrade (zWSU). Contact your local TSS seller or Z BP seller for additional information on pricing/content.

z15 Software Pricing

What software pricing is announced with the z15?

Four new or updated Technology Transition Offerings are being announced. They are Technology Update Pricing for the z15 and one new and two revised Transition Charges for Sysplexes or Multiplexes programs.

Technology Update Pricing for the z15 (TU7) uses the reporting mechanisms and existing MSU tiers of the AWLC or CMLC price metrics while increasing the software price performance improvements provided by AWLC or CMLC.

The four Technology Transition Offerings apply only to eligible z/OS and z/TPF operating systems and their associated middleware programs when running in the following configurations:

- On a stand-alone z15 server
- In an aggregated z/OS Parallel Sysplex® with a z15 server
- In an aggregated z/TPF Loosely Coupled Complex with a z15 server
- In a Multiplex with a z15 server

What software pricing is available on a standalone z15 or a z/OS Parallel Sysplex or a z/TPF Loosely Coupled Complex or a Multiplex with all z15 technology?

For a standalone z15, Sysplex, Complex, or Multiplex with all z15 servers, IBM will offer the Technology Update Pricing for the z15 (TU7) which leverages the existing AWLC or CMLC pricing metrics while offering price performance improvements for z15. The Technology Update Pricing for z15 provides reduced AWLC or CMLC pricing on the z15 as compared to AWLC or CMLC on the z14 with greater than 3 MSUs.

Note: If you choose to do so, you may continue to use PSLC pricing with the z15.

What is the price reduction available with the Technology Update Pricing for the z15?

The percent reduction in the monthly AWLC or CMLC is based on the number of z15 full capacity MSUs. While the percent reduction is only based on full capacity MSUs, you remain eligible for sub-capacity pricing under business as usual terms. AWLC or CMLC pricing for a standalone z15 or Parallel Sysplex or Loosely Coupled Complex or Multiplex of all z15 servers is reduced by the percentage in the table below.

Schedule of AWLC or CMLC reductions for Technology Update Pricing for the z15 (TU7)

Quantity of z15 Full Capacity MSUs for a stand-alone server, or the sum of Full Capacity MSUs in an actively coupled Parallel Sysplex, z/TPF Loosely Coupled Complex, or Multiplex made up entirely of z15 servers	Reduction in Monthly AWLC or CMLC
1 – 3 MSUs	0.0%
4 - 45 MSUs	8.0%
46 - 315 MSUs	16.0%
316 - 1315 MSUs	18.0%
1316 - 2676 MSUs	20.0%
2677 - 5476 MSUs	24.0%
5477 - more MSUs	28.0%

If I have Tailored Fit Pricing, will I benefit from the z15 Technology Transition Offerings?

With Tailored Fit Pricing, you will receive the benefit of the Technology Update Pricing or Transition Charges for Sysplexes, Complexes or Multiplexes for which your environment qualifies. For example, if you have all z15 servers with Tailored Fit Pricing, you will qualify for Technology Update Pricing for the z15 (TU7). If you have only z14 and z15 servers, you will qualify for the Transition Charges for Sysplexes, Complexes or Multiplexes (TC5.)

If I upgrade from a z14 to a z15 server will software pricing improvements be additive?

You will receive the pricing advantage of Technology Update Pricing for the z15 (TU7) for the z15 server. The Technology Update Pricing for the z15 (TU7) price reduction includes the benefits of the prior price reductions. Technology Update Pricing for the z14 (TU5) is available when upgrading to a z14.

MSUs: Quantity of Full Capacity MSUs for a stand-alone server, or the sum of the Full Capacity MSUs in an actively coupled Parallel Sysplex or z/TPF Loosely Coupled Complex or Multiplex made up entirely of one generation of servers	<u>z14 – TU5</u> Reduction in Monthly AWLC or CMLC	<u>z15 – TU7</u> Reduction in Monthly AWLC or CMLC
4 – 45 MSUs	6.0%	8.0%
46 - 315 MSUs	12.0%	16.0%
316 – 1315 MSUs	14.0%	18.0%
1316 – 2676 MSUs	15.0%	20.0%
2677 – 5476 MSUs	18.0%	24.0%
5477 or more MSUs	21.0%	28.0%

If I add a z15 into a Sysplex or Complex or Multiplex with one or more z14 and/or z14 Model ZR1 servers, but no older technology servers, can I benefit from the new Technology Update Pricing for the z15 (TU7)?

No. The Transition Charges for Sysplexes or Multiplexes (TC5) provides a reduction in AWLC or CMLC pricing based upon the number of MSUs of z15, z14, and z14 Model ZR1 servers. When you migrate to all z15 servers, the Technology Update Pricing for the z15 (TU7) may apply.

Schedule of AWLC or CMLC reductions for Transition Charges for Sysplexes or Multiplexes (TC5)

Quantity of z15 and z14 Full Capacity MSUs in an actively coupled Parallel Sysplex, Loosely Coupled Complex, or Multiplex	Reduction in Monthly AWLC
1 – 3 MSUs	0.0%
4 - 45 MSUs	7.0%
46 - 315 MSUs	14.0%
316 - 1315 MSUs	16.0%
1316 - 2676 MSUs	17.0%
2677 - 5476 MSUs	21.0%
5477 - more MSUs	24.0%

If I add a z15 into a Sysplex or Complex or Multiplex with one or more z14, z14 Model ZR1, IBM z13® and/or IBM z13s® servers, but no older technology servers, can I benefit from the new Technology Update Pricing for the z15?

No. The previously announced Transition Charges for Sysplexes or Multiplexes (TC4) is updated to include the z15. The Transition Charges for Sysplexes or Multiplexes (TC4) provides a reduction in AWLC or CMLC pricing based the number of MSUs of z15, z14, z14 Model ZR1, z13, and z13s servers. When you migrate to all z15 servers, the Technology Update Pricing for the z15 (TU7) may apply.

If I add a z15 into a Multiplex with one or more IBM zEnterprise® EC12 (zEC12) and/or IBM zEnterprise BC12 (zBC12) servers, can I benefit from the new Technology Update Pricing for the z15?

No. The previously announced Transition Charges for Multiplexes (TC3) is updated to include the z15. The Transition Charges for Multiplexes (TC3) provides a reduction in CMLC pricing based on the percentage of the full capacity MSUs that are z15, z14, z13, z13s, zEC12, and zBC12. When you migrate to all z15 servers, the Technology Update Pricing for the z15 (TU7) may apply.

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

Yes. AWLC and CMLC allow you to report and pay for software based on sub-capacity charges using the existing SCRT process. AWLC and CMLC also allow for full-capacity pricing based on the rated MSUs of your z15 server. All of the Technology Transition Offerings may be used in a sub-capacity environment.

What contracts are required for AWLC or CMLC pricing and the Technology Transition Offerings?

In order to get AWLC pricing with any of the Technology Transition Offerings, the following contract is required to be executed:

- Attachment for IBM z Advanced Workload License Charges (Z125-8538)

In order to get CMLC pricing with any of the Technology Transition Offerings, the following contract is required to be executed:

- Attachment for Country Multiplex Pricing (Z126-6965)

The additional benefits of the Technology Transition Offerings are provided in a supplement called the Supplement for Technology Transition Offerings (Z125-8994). This supplement does not require a signature.

What is the Processor Value Unit (PVU) rating used for Passport Advantage® software on the z15?

The PVU rating for the z15 will be 120 for both IFL engines and CP engines, the same as the z14, z13, zEC12, z196 and the z10 EC. The PVU rating for the z14 Model ZR1, z13s, zBC12, z114 and z10 BC remains 100 for both IFL engines and CP engines.

The number of MSUs used to determine either the MSUs of a stand-alone z15 server or the total MSUs of a z15 Sysplex, Complex, or Multiplex are based on the announced IBM full capacity ratings that can be found on the Mainframe Exhibits section of the z Systems Software Contracts website at <http://ibm.com/systems/z/swprice/reference/exhibits/hardware.html>

Where can I get more information on IBM software charges?

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

Physical Planning for z15

What are a few changes I'll see in the new frame configuration on the IBM z15?

This new configuration enables significant floor space reduction for most clients. For the new z15 package:

- Whether on raised floor or not, each frame sits on a space that is one data center floor tile wide (0.6 m or 24 inches) by two floor tiles deep (1.2 m or 48 inches). This enables z15 participation in almost all existing data center equipment aisles, as well as the latest more efficient hot and cold air containment cooling designs.
- All high-performance coupling connections, I/O cabling and power cords are located at the back of the frame. Additional cable management technology separates internal system cabling from I/O cabling allowing better cable management within the system frames.
- Side covers are truly optional and only applied on the exposed sides of a system for appearance only.
- There are two power options - intelligent power distribution unit (iPDU) and Bulk Power Assembly (BPA).
- The doors are designed for acoustics and optimized for air flow. Since the system has its exhaust air more widely distributed and exiting the system at lower velocity than in past Z systems, there is no need for the louvered rear doors used on z13 and z14 to direct exhaust air up or down.
- Maximum individual frame powers are typically 5 - 10 kW, which matches the capability of the majority of data centers
- Frame weight is better distributed so there is no need for weight distribution bars. Other than at the location of floor cutouts, most data center raised floors will require no additional reinforcement.
- Top exit I/O no longer increases system width as in the predecessor Z systems.
- There is a new, faster fill and drain tool for the new footprint. There is a new 19" lift tool introduced with z14 ZR1

Still supported are:

- Both raised and non-raised floor installation as well as top and bottom exit I/O, and top and bottom exit power.
- Front to back airflow
- A water-cooling option for which there are more details in other FAQs

With the internal configuration changes implemented for the 19" frame, the amount of I/O that can be attached is dependent on the type of power that is selected. There is a maximum of twelve PCIe+ I/O drawers (16 I/O slots each) when iPDU power is selected, or a maximum of eleven PCIe+ I/O drawers when BPA is selected.

With iPDU, you can fit more IO capacity without needing extra frames, save significant power (10 - 30%), and move to a standardized power approach within your data center (4 wire, 3 phase 200 - 240 VAC, or 5 wire, 3 phase 380 - 415 VAC) . BPA enables IBF (internal Battery Feature, an internal short term UPS), typically has fewer line cords, and is required for the water cooling option.

What is the cooling on the z15?

The z15 is designed with an environmental focus on improving data center efficiency. There are two options for cooling on the z15 – radiator or customer water cooling. Both options require a fill and drain tool (FDT) for initial fill of the no maintenance, closed internal cooling loop for the high performance, high reliability, high efficiency Z processor modules. The FDT is also needed to drain the system before shipping, MES processor drawer adds and very rarely for service. The new tool was designed to reduce the complexity of operation and improve efficiency.

The radiator-based air-cooled system is designed for more efficient cooling and improved maintenance. A radiator-cooled, iPDU power system is the most energy efficient z15 system configuration. Customer-water cooled is available for the z15. It is the solution if your data center cannot handle the full heat load of the z15 system going to data center air. This could be due to air cooling limitations at the z15 install location (sometimes referred to as a "hot spot") or an overall limited air cooling capacity in the data center.

What power options are on the IBM z15?

There are two power options - intelligent power distribution unit (iPDU) and Bulk Power Assembly (BPA). Both are 2N, that is designed with twice the minimum needed line cords so that normally all the power can be dropped to one power side and the system keeps running. Both are intelligent, fully managed and fully integrated within the system's power distribution, control, and maintenance support structure. The use of the iPDU power reduces power consumption 10 - 30% versus BPA. iPDU often requires one less frame for a system configuration, is more energy efficient, and supports more maximum I/O slots. BPA is needed for IBF (Internal Battery Feature), or water cooling and often requires 2 fewer line cords to the system.

What makes the iPDU intelligent?

"Intelligent" means that the PDU has communication interfaces to the systems internal control and monitoring system and an internal microprocessor that allows internal parameters (mainly voltages and currents) to be read. It also allows individual outlets to be turned on and off by command. The device's firmware allows the unit to perform certain higher level functions, for example, a total system power reset on command.

A standard PDU is simply a grouping of AC outlets protected by circuit breakers that distribute input AC power. It does nothing but distribute power.

What is ASHRAE?

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.

What is the ASHRAE rating of the z15?

The z15 is rated ASHRAE Class A3 (4th Edition, 2015) for all models and configurations. Like z14, this supports the wider operating temperature range (including up to 40 degrees C inlet air), as well as a wider humidity range than the z13's Class A2 rating, but adds the support of very low humidity for the recommended operating range. This will allow many data centers to avoid the energy and water use required for added humidification. As for z14, the ASHRAE A3 extended temperature and humidity ranges are meant mainly for robustness reasons, to assure z15 system operation even with data center cooling or environmental control failures. It can also allow installation in a wider range of data centers,

including lower cost data centers. Energy savings can be 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 short time periods with high external temperatures, reducing the required cooling infrastructure and energy usage.

Will there be an Internal Battery Features (IBF) on the z15?

Yes. But the IBF is only available on BPA power z15 servers. It is Lithium ion battery technology based, not VRLA (valve-regulated lead-acid) like predecessor Z systems.

Are there changes to the IBF from prior IBM Z servers?

The Internal Battery Feature (IBF) available with z15 contains Lithium ion batteries greater than 300 Wh. As such they are fully regulated Dangerous Goods so care will need to be taken for any repair actions, relocations or machine returns. You will be responsible to handle the batteries in accordance with all local, legal and environmental requirements.

The need for the IBF will be redundant if your location already provides uninterruptible power. IBF may no longer be a business requirement. In client locations without an uninterruptible power supply, there are resiliency advantages to using an IBF feature when a power failure occurs, for example, for preserving modified data in Coupling Facility images until power is restored.

On a picture of a rear view of an iPDU power z15 systems' internals, I saw what looked like some available c19 receptacles in the frame. Can these be used to power components outside the z15?

No. Those are the AC power ports of the iPDU system that are designed, controlled, and monitored to only power internal system hardware (CPC drawers, I/O drawers, processor cooling, SEs, etc.). Nothing else should be plugged in to these locations.

What is Balanced Power Plan Ahead (FC#3003) on a z15?

This feature code requires BPA power. It populates all BPA with 3 BPR (Bulk Power Regulator) so that the system can through active internal control draw the same currents on all 3 phases. This is called balanced power, which with everything else being equal is a desirable condition. It may not be the right answer for most z15 clients if they are choosing BPA power just to get this feature. It depends on the specific system configuration. Due to the much lower energy use of iPDU powered z15 systems, in most cases this feature may provide no significant advantage versus iPDU, but will consume significantly more power. The purpose of balanced power is to reduce the highest aggregate phase current by distributing load equally across all 3 phases, particularly in the N power feeds case. The reason that for many z15 configurations there is no real advantage for balanced power is that due to its efficiency the iPDU power is much lower, so that even though not truly balanced, all of its phase currents are equal to or less than that for BPA with balanced phase currents.

Will there be High Voltage DC Power on the IBM z15?

No, the z15 does not support High Voltage DC Power.

What cabling options exist on the z15?

You can order top or bottom exit cabling options on the z15. For both the top and bottom exit cabling there is rear exit only. The z15 is designed with a vertical cable management system that includes a

vertical cable retention bar (aka 'spine') and tool-less cable retention clips. These spines assist in both top and bottom exit cabling, keeping external cables orderly and alleviating frame side cabling congestion.

Tell me about the non-raised floor option on the z15?

The z15 can be configured for installation without a raised floor. This is needed for clients looking for installation 'on slab', that is directly on a concrete floor. For a non-raised floor z15, you must order top exit I/O cabling (this also handles the line cords which must top exit), and radiator air-cooling.

What is the service lift tool on the z15?

The lift tool attaches to the front of a z15 frame and provides a platform to lift/lower heavy drawers out of the front of the system for MES adds or very rare service actions. It was introduced on the 19 inch z14 ZR1.

Where can I get more physical planning information for the z15?

The best place for information is the Installation Manual for Physical Planning (IMPP) which is available here:

<https://www.ibm.com/servers/resourcelink/lib03010.nsf/pages/8561InstallationManualForPhysicalPlanning?OpenDocument>

Tell me about the new FICON Express16SA for z15.

FICON Express16SA is designed to help absorb large application and transaction spikes like those driven by large unpredictable cloud, analytic, and mobile workloads.

The FICON Express16SA on the z15 supports a link data rate of 16 gigabits per second (Gbps) and auto-negotiation to 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 z15 servers enables you to position your SAN to meet the lower latency and increased bandwidth demands of your applications.

The new FICON Express16SA channel will work with your existing fiber optic cabling environment, both single mode and multimode optical cables.

The FICON Express16SA adapter is exclusive to the z15.

What kind of increased performance can I expect to get with zHPF using FICON Express16SA on z15?

The use of FICON Express16SA on an IBM z15 with the zHPF protocol and a mix of large sequential read and write data transfer I/O operation, or with small data transfer I/O operation, will achieve a throughput equal to that of the FICON Express16S+.

What kind of increased performance can I expect to get on commercial batch elapsed time on the z15 with FICON Express16SA?

The use of FICON Express16SA on an IBM z15 running I/O intensive batch workloads will achieve throughput equal to that of the FICON Express16S+.

What kind of increased performance can I expect to get for the FCP protocol using FICON Express16SA on z15?

A FICON Express16SA 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 Express16SA in an IBM z15 with the FCP protocol for small data transfer I/O operations will achieve throughput equal to that of FICON Express16S+.

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

How many channel subsystems (CSS) are supported on the z15?

The z15 has channel subsystem (CSS) scalability with support for six logical channel subsystems (LCSSs) which are required to support the eighty-five LPARs for z15, 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 DS8900 Multi-target Metro Mirror storage replication with TPC-R HyperSwap®.

Which of my PCIe I/O adapters can I continue to carry forward to the z15?

You can carry forward to the z15 the FICON Express16S+, FICON-Express16S, FICON Express8S, OSA-Express7S 25 GbE SR, OSA-Express6S, OSA-Express5S, 10GbE RoCE Express, 10 GbE RoCE Express2, 25 GbE RoCE Express2, Crypto Express6S, Coupling Express LR, and Integrated Coupling Adapter SR (ICA SR).

What is the 25GbE RoCE Express2.1 adapter?

The 25GbE RoCE Express2.1 is a name change from the original 25 GbE RoCE Express2 that was introduced on the 14. The adapter 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.

How many 25 GbE RoCE Express2.1, 25 GbE RoCE Express2, 10GbE RoCE Express2.1 or RoCE Express2 adapters does the z15 support?

The z15 supports up to 16 adapters are supported (32 ports). Each LPAR should be provisioned with 2 PFIDs for redundancy, ideally with each PFID defined on two separate ports across two separate adapters.

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”.

What is the worldwide port naming (WWPN) assignments that is on the z15?

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 z15 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 is zHyperLinks on the IBM Z and DS8900?

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 Z and DS8900F storage to deliver the extreme low latency I/O access for Db2 for z/OS Applications.

What is the expected value of low latency I/O for Db2 on z/OS Applications using 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.

IBM z15 Networking

What is the OSA-Express7S 25 Gigabit Ethernet Short Reach 1.1 (OSA-Express7S 25 GbE SR1.1) adapter for z15?

The OSA-Express7S 25 Gigabit Ethernet (GbE) Short Reach (SR) 1.1 feature provides improved performance from prior level OSA-Express6S 10GbE adapters. The 1.1 adapter is only a name change and has the same functionality as the OSA-Express7S 25 GbE SR that became available on the z14. 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.

What are the new OSA-Express7S adapters for the z15?

The z15 introduces an Ethernet technical refresh to complete the full family of the OSA-Express7S generation of adapters. There is the OSA-Express7S 1000BASE-T Ethernet for copper environments, in addition to OSA-Express7S 10 Gigabit Ethernet (10 GbE) and OSA-Express7S Gigabit Ethernet (GbE) for single-mode and multimode fiber optic environments. The performance characteristics are comparable to the OSA-Express6S 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. These 3 new OSA-Express7S family of adapters is exclusive to the z15.

IBM Virtual Flash Memory

What is Virtual Flash Memory on the IBM z15?

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 z15 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 IBM Z servers support Virtual Flash Memory?

The Virtual Flash Memory is only supported on the z15 and z14. Prior to the z14, storage class memory was provided by the Flash Express adapter.

How do I order Virtual Flash Memory for the z15?

Virtual Flash Memory is a feature of IBM Z and is located within Random Array of Independent Memory (RAIM) on the IBM z15 (z15).

Granularity of order is improved on the z15 over the z14 where there was only 4 ordering options. On the z15 it is orderable in twelve .5 TB sizes between .5 TB and 6 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 z15?

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 PARMLIB

What is the correlation between pageable large pages on the z15 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.

Integrated Acceleration for zEDC and zEDC on z/OS

What is IBM z Enterprise Data Compression (zEDC) for IBM Z?

A capability of z/OS V2.1 (and beyond), IBM zEnterprise Data Compression (zEDC) combined with hardware acceleration (zEDC Express adapter for z14 and select earlier Z servers and Integrated Accelerator for zEDC for z15), offers a compression acceleration solution designed for high performance, industry standard, low latency compression with little additional overhead.

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.

What's new on zEDC compression for the z15?

There is now an Integrated Accelerator for zEDC on the IBM z15 processor chip. Clients no longer need to purchase zEDC Express adapters for their servers.

What software exploits the Integrated Accelerator for zEDC on the z15?

Software that takes advantage of the Integrated Accelerator and zEDC on z/OS include:

- Those that offer storage savings - z/OS SMF Logstream, z/OS QSAM/BSAM, z/OS DFHSM/DFDSS
- Those that offer storage compression: Db2 for z/OS V12 LOBs, z/FS in V2R3

What if I don't enable the zEDC feature of z/OS – is there any software when your running z/OS that can exploit the Integrated Accelerator for zEDC on the z15?

Yes. There are use cases where no z/OS feature is required on z15. The following z/OS offerings can exploit the Integrated Accelerator for zEDC:

- IBM Java
- For Db2 storage compression - IBM Content Manager OnDemand, IBM Encryption Facility
- For data transfer: IBM Sterling Connect:Direct®, z/OS HTTP Server, OpenSSH, IBM MQ® for z/OS

Is zEDC only for z/OS or can it be exploited by Linux on Z?

No. The zEDC Express adapter exploitation with Linux on Z requires one of the following Linux distributions for IBM Z: Canonical Ubuntu 16.04.03 LTS and later, Red Hat Enterprise Linux (RHEL) 7.3

and later, or SUSE Linux Enterprise Server (SLES) 12 SP3 and later. The new Integrated Accelerator for z15 requires RHEL 8.1.1³, SLES 12 SP5³, and Ubuntu 19.10³.

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, openssh, openssl, subversion, git, or PNG.

For use with z14 and zEDC Express, these are the sweet spot products - IBM Db2, IBM Java, IBM MQ, IBM WebSphere Application Server, as well as open source products such as Apache Kafka. On the z15 with Integrated Accelerator for zEDC, products like Java 8.0.6³, WAS 8.5.5³, and WAS 9.0.5³, IBM is working with Linux distributors for support.

What happens if data that is already compressed gets compressed again by zEDC on the z15?

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.

How do I order the zEDC z/OS feature?

zEDC 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 for z/OS to those servers that will be exploiting zEDC. Contact TechLine or your MLC pricing representative for details.

Is the Integrated Accelerator for zEDC shared between multiple LPARs on z15?

Yes.

What planning tools are there? How can I determine if I will see benefit from using zEDC?

The IBM Washington Systems Center (WSC) has 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/atsmastr.nsf/WebIndex/PRS5132>

³ IBM is working with the Linux distribution partners to get the functionality included in their distribution for Linux on Z.

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 the Integrated Adapter for zEDC on the z15?

z/VM guest exploitation support for the Integrated Adapter for zEDC

What are the benefits of zEDC for z/VM customers on the z15?

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

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.

If I use zEDC Express today for compression on z14 and earlier what changes will I need to make for z/OS?

The new Integrated Accelerator for zEDC requires no changes for z/OS users exploiting zEDC compression on z14 or earlier hardware.

Do I need to purchase additional hardware to take advantage of zEDC on z15?

No additional hardware is needed to utilize the benefits on zEDC on z15.

Will I need to make any application or I/O changes to utilize zEDC on z15?

No changes to applications or I/O are needed to utilize the Integrated accelerator for zEDC.

What is the difference between synchronous and asynchronous compression?

Synchronous deals more with compression at the application level, such as zlib. In this form of compression user applications invoke instructions in a virtual address space. Asynchronous compression is usually for a specific use case and deals with larger operations for authorized applications.

Will compression on my z14 with zEDC be compatible with z15?

Yes, z14 and z15 machines with zEDC will coexist and all data interchange will remain compatible. Data compressed and written with zEDC Express adapters will be read and decompressed by the Integrated Accelerator for zEDC.

Can I still use my current zlib?

Both the current zlib and the new zlib will function on both z14 and below as well as z15 hardware.

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

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

How can I learn more about the new z/CX Container Extensions and get some help getting started with this new capability?

IBM Systems Lab Services consultants get early access to new functions such as z/CX Container Extensions and based on that early experience have developed a new planning and implementation offering for zCX Container Extensions. As with many of our planning and implementation offerings, we customize the exact services to meet individual client needs. Our z/CX offering will help clients plan and use the z/CX capabilities while helping them learn more about the capability. You can contact Systems Lab Services directly via e-mail at: ibmsls@us.ibm.com or contact your IBM Sales Representative.

z/VSE

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

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 higher), 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)?

IBM Dynamic Partition Manager supports FCP and FICON ECKD™ storage devices.

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

IBM Dynamic Partition Manager is supported by IBM z15, IBM z14, IBM z13s, IBM z13, IBM LinuxONE Rockhopper (2965), and IBM LinuxONE Emperor (2964).

Anything I should be aware of when I switch z15 my server to IBM Dynamic Partition Manager (DPM)?

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

What is available on IBM Dynamic Partition Manager (DPM) 3.2?

IBM Dynamic Partition Manager supports auto configuration of devices to simplify Linux Operating System Installation, where Linux distro installers exploit the function.

Dynamic Partition Manager supports 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 allowing the FTP SE originated operations to follow our security recommendation.

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.

IBM z14 Hardware

Tell me more about the IBM z14™.

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.

Tell me about the IBM z14 Model ZR1 announcement?

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 up 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 up 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](#).

Is it true the z14 ZR1 is housed in a 19” frame?

Yes, the z14 Model ZR1 is housed in a 19-inch frame 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 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.

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 much capacity will the z14 deliver?

Powered by 170 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).

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](#).

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 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.

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.

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.

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.

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/

z14 ZR1 Warranty – additional information

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.

Water cooling option for z14 Models M01-M05

When should I consider installing water cooling on the IBM z15?

The water-cooling option is the solution when your data center cannot handle the full heat load of the z15 system going to data center air. This could be due to air cooling limitations at the z15 install location (sometimes referred to as a "hot spot") or an overall limited air-cooling capacity in the data center. The water-cooling option typically takes the greatest majority to all of the z15 heat load to the supplied data center water. To understand what the heat loads to water and air are, use the Power and Weight Estimation Tool for 8561 on IBM Resource Link. As opposed to previous Z systems, the water cooled z15 is not the overall most energy efficient configuration. The most energy efficient configuration is always air cooled with iPDU power.

Can I tap into my building's chilled water for z15 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.

What changes will I need in my floor tile cut outs for water cooling on z14?

The water-cooled option can only be installed on a raised floor and the water lines must come up from under the raised floor. This means that there must be a large enough floor tile cutouts to the rear of the frames that contain the water cooling units to fit the insulated water lines plus any I/O or power cables. This will be the A frame, and if four or five CPC drawers, the B frame. The water cooling option adds 4 inches depth to the rear of the system.

Is there any change to the height of the server for water cooling on the z15?

No.

How many Water Cooling Units (WCU) are on a water cooled z15?

The z15 has two to four Water Cooled Units (WCU), in a 2N design. One to three CPC drawers require two feeds and two returns to the rear of the A frame for its two WCUs. Over three CPC drawers requires an additional two feeds and two returns to the rear of the B frame for its two WCUs.

Will data center water be going through my z15 server with water cooling?

No - the water on the internal z15 side of the WCU is contained in a closed-loop system, maintained by IBM service personnel. There is a fill and drain tool that is used to fill the system initially and to service it if needed, just as for the air cooled system.

Is there any basic maintenance required because of water cooling on the z14 once the system is installed?

No scheduled maintenance is required.

What will happen if the humidity in my data center fluctuates to a high level – will I have a problem with condensation on the internal water lines in the z15?

No. 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.

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 at [Mastertheframe.com](https://www.ibm.com/academic/mastertheframe).

If you'd like to take a class at a local school, 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 [IBM Z Skills website](#).

Students can also access [56 hours of free e-learning](#) as part of the Z Enterprise Computing Kickstart offering that's available on the "Software tab" of the IBM Z Academic Initiative website. Students can earn Open Mainframe Digital badges upon completion of the COBOL and Db2 courses that are a part of this offering.

Finally, you can connect to the [Z community](#) to keep up-to-date on all of the latest news, events, and offerings around Z. Join the community today!

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

Absolutely! We provide Z system access for faculty and students to use in the classroom and we provide community systems for developers to use just for playing around with for a free 120-day trial period. Faculty should visit the "Courseware" tab of the IBM [Z Academic Initiative website](#) to register for the z/OS Academic Cloud or visit the "Resources" tab to request access to Linux on Z as part of the 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 [IBM Z Skills Gateway](#) 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 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.

Visit the [IBM Z Skills website \(https://www.ibm.com/it-infrastructure/z/education/skills\)](https://www.ibm.com/it-infrastructure/z/education/skills) to access two additional helpful resources:

1. [IBM Z Skills Employer Resource Center \(http://ibm.biz/zemployerresources\)](http://ibm.biz/zemployerresources) - a one stop shop that provides resources to help with each phase of the talent lifecycle (developing a talent strategy, finding and attracting talent, training talent, growing and retaining talent, and transitioning talent and creating succession plans). The resource center includes manager testimonials, tips from recent hires, examples of proven tactics with examples, recommended training paths, videos and so much more. You need an IBM ID to access the materials, but it's free to get an ID and to view the content.
2. [Executive Summary \(https://www.ibm.com/downloads/cas/OVGDA3XO\)](https://www.ibm.com/downloads/cas/OVGDA3XO)- All of the top, key resources from the Employer Resource Center are highlighted, and links to the assets are provided on the second chart of this 2-page executive summary. This is an excellent high-level overview showcasing our top assets for Z Skills topics, including training.



© Copyright IBM Corporation 2018

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

IBM, IBM logo, IBM Z, Db2, ECKD, FICON, GDPS, HyperSwap, IMS, LinuxONE, LinuxONE Emperor, LinuxONE Rockhopper, MQ, Parallel Sysplex, PR/SM, RACF, Resource Link, RMF, System z10, WebSphere, z13, z13s, z14, z15, 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.