

---

**IBM Z**  
**Introduction**  
**May 2021**

**IBM z15**  
**And IBM z14**  
Frequently Asked Questions

Worldwide

**IBM**

---

## Table of Contents

IBM z15 Announcement and Common update information .....	3
IBM z15 Model T02 Hardware .....	8
IBM z15 Model T01 Hardware .....	15
Performance .....	23
z15 Warranty.....	27
z15 T02 Software Pricing.....	28
z15 T01 Software Pricing.....	32
Physical Planning for z15.....	36
I/O.....	41
IBM z15 Networking.....	46
IBM Virtual Flash Memory.....	47
Integrated Acceleration for zEDC and zEDC on z/OS.....	49
z15 Sustainability .....	51
z/OS .....	51
z/VSE.....	52
z/VM.....	53
KVM.....	54
Linux on IBM Z .....	55
IBM Dynamic Partition Manager.....	56
Academic Initiative and Skills.....	57

## IBM z15 Announcement and Common update information

### Tell me about the IBM z15™ (z15).

On September 12, 2019, IBM announced the IBM z15 - the next generation in the IBM Z® family. The new IBM z15 T01 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 customers need assurance that their data is safe, where services are always on and fulfilled instantly. The z15 can help customers adapt to planned or unplanned events while keeping services and operations running smoothly and continuously, whether on premise or in the cloud.

### Tell me about the IBM z15 T02.

On April 14, 2020, the newest member of the z15 family was announced – the IBM z15 Model T02. The z15 T02 is an air-cooled, single frame, efficient design, with a low entry cost. Like the other servers in the z15 family, it is designed to help enable cloud native development and deployment, achieve encryption everywhere, and provide cyber resiliency to ensure scalable isolation of workloads to protect from threats while ensuring continuous availability of services.

While it's built on the z15 processor chip, the z15 T02 is closer in comparison to the IBM z14® ZR1. The z15 T02 is housed in a single 19-inch frame, designed with redundant power, cooling, and line cords. It also has the option of single or three phase power to accommodate data center requirements. These factors can help reduce power costs, reduce footprint cost, and install in virtually any existing data center helping facility standardization.

The z15 T02 will support the new System Recovery Boost introduced on the z15 T01. Also, the server can be used in conjunction with IBM Data Privacy Passports which is designed to enforce security and privacy protections to eligible data not only on IBM Z, but across platforms.

The z15 T02 is designed with integrated accelerators to offload general processors and reduce system overhead CPU cycles for application workload processing. The IBM Integrated Accelerator for z Enterprise Data Compression is provided on each processor chip and uses industry-standard compression formats for file compression that can enable reduction in the size of data which can save storage space and increase data transfer rates. The Integrated Accelerator for Z Sort, standard on IBM z15, optimizes elapsed time for sort workloads, which typically occur during customer batch windows.

## What is the 16U Reserved space for DS8910F?

For customers utilizing single-phase power and fewer than 2 PCIe+ drawers, there is an option to hold 16U Reserved Space (#0151) in the z15 T02 rack. This 16U of space in the frame, available only on the z15 Model T02, can be populated by a customer with IBM DS8910F Model 993 storage in order to reduce the physical footprint in the data center. When the 16U Reserved Space feature is ordered, the z15 T02 order will provide additional power ports in redundant power distribution units (PDUs), along with weight ballast and air flow fillers as appropriate. Note that IBM would first complete installation of the z15 T02 and turn it over to customer operation, and then the storage can be added if the 16U Reserved Space (#0151) has been ordered. This new hardware configuration flexibility feature is perfect for customers who fit into smaller I/O configurations and need the additional ease of a single-footprint approach. The 16U Reserved space for DS8910F is only available on the IBM z15 Model T02. **Note:** Review the [z15 T02 and LinuxONE III LT2 Exception Letter on Resource Link](#)® to understand when this offering has completed its compliance certification in your geography. Failure to do so may result in a non-compliant machine configuration.

## Why does IBM Z use the terms single frame and multi frame to describe IBM z15 systems?

The IBM z15 Model T01 is a water-cooled multi frame system that scales from one frame to four 19" frames to accommodate large amounts of processor, memory and I/O. It is designed to deliver the ultimate in scalability and expandability for customers with enormous workloads and intensive growth.

The IBM z15 Model T02 is an air-cooled single frame system that delivers expandability within a single 19" footprint for workloads that have more moderate growth.

Customers need to consider current and future projected capacity requirements when selecting the server that fits their needs best.

## What is the new Secure Execution for Linux?

Secure Execution for Linux® on Z is a new capability introduced on April 14, 2020 for the IBM z15 family – available for both the z15 T01 and the z15 T02. Engineered to help protect against insider and outsider threats in multi-tenanted cloud environments, it ensures users, and even system administrators, are unable to access sensitive data in Linux-based virtual environments. Secure Execution for Linux protects the confidentiality and integrity of data, by isolating data at the virtual machine level, and ensuring that only the people within the organization that have a need-to-know have access to data in the clear.

For more information see the Secure Execution FAQ document at:

<https://www.ibm.com/downloads/cas/G1WLJDAY>

The solution brief: <https://www.ibm.com/downloads/cas/O158MBWG>

The customer brief at: <https://www.ibm.com/downloads/cas/GPLNZLE2>

## What is the new Hardware Management Appliance?

New for the z15 family, the new IBM Z Hardware Management Appliance feature provides the capability for the Hardware Management Console (HMC) and Support Element (SE) functionality to be contained within redundant physical servers inside the server frame. When you order the IBM Z Hardware Management Appliance feature, this will provide logically a Primary and Alternate Support Element and two peer Hardware Management Consoles on two physical servers in the CPC frame. This eliminates the need for having to manage a separate physical server or servers for one or more HMCs outside of the frame. For the User Interface experience, you must use remote browsing controls from your own workstation into HMC within the IBM Z Hardware Management Appliance.

If you have multiple systems, you don't need to order the Hardware Management Appliance feature for all servers. The recommendation is that you consider having the IBM Z Hardware Management Appliance features on one or two servers, but the rest of the servers don't need to include Hardware Management Appliance features. Although those servers would have redundant SEs. The IBM Z Hardware Management Appliance feature is optional.

### **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.

The System Recovery Boost GDPS® enhancements are a combination of changes to GDPS processing/scripting itself, and GDPS taking advantage of new functionality being offered by the z15 Support Element. Various APIs used by GDPS on the SE are being streamlined and improved, while GDPS itself will be doing more in parallel when possible. In addition, GDPS will be enhancing its scripting language to allow multiple systems to be addressed in parallel by some commands, as opposed to sequentially as before, and making changes to some control flows to improve efficiency.

In addition to the Speed Boost, zIIP Boost and GDPS enhancements capabilities that are included on the z15 at no additional charge, on the IBM z15 T01, you also have the option to build upon this base functionality with System Recovery Boost Upgrade, a new Capacity on Demand offering that lets you maximize performance and parallelism during the boost period by unlocking previously un-used “dark cores” for additional zIIP capacity when you need it the most.

On the z15 T01, new with the April 14, 2020 announcements, IBM has increased the term for System Recovery Boost Upgrade on IBM z15 Model T01 from one year to five years, available in one-year increments. New feature code 6799 will record how many subscription years you purchase. Feature code 6802, SRB Upgrade Record, will still enable temporary activation of additional physical zIIP processors to be used in conjunction with System Recovery Boost. Feature code 9930, SRB Upgrade Authorization, is still required to enable the ordering of SRB Upgrade Record (#6802). Feature code 9930 drives the necessary contracts to use the Boost Record. Feature code 6799 further qualifies the use of these existing feature codes.

As announced on November 26, 2019, no specific z15 feature codes need to be ordered to be able to use the base capabilities of System Recovery Boost. Feature codes 6802, 6799 and 9930 are required only to make use of the optional System Recovery Boost Upgrade Record, which allows the temporary activation of additional zIIP processor resources in conjunction with System Recovery Boost.

Note that the System Recovery Boost Upgrade capability is only available on the z15 T01.

For more Information check out the Systems Recovery Boost FAQs:

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

In September 2020, IBM will deliver new System Recovery Boost capabilities that will support a new class of boost, called recovery process boosts, that can be applied to a range of sysplex recovery processes, including sysplex partitioning, CF structure recovery, CF data sharing member recovery, and HyperSwap® recovery to minimize disruption of recovery processes and maximize sysplex resiliency. Now, you can use System Recovery Boost to mitigate the disruption caused by these recovery processes and maximize the benefits of your sysplex environment. The new short-duration recovery process

boosts take advantage of the same GP Speed Boost and zIIP Boost technologies used to accelerate IPLs, but are instead used to accelerate one or more of the sysplex recovery processes I mentioned before. Each individual recovery process boost can last up to five minutes, and you can use up to thirty minutes worth of these boosts in a consecutive 24-hour period.

### ***New with the May 4, 2021 announcements***

#### **What is the new addition, Hardware Consumption Solution, to the Tailored Fit Pricing Model for IBM Z family?**

There is a new addition in the Tailored Fit Pricing for IBM Z family, the Hardware Consumption Solution. The new offering provides a simple, flexible cloud-like solution for unpredictable, spiky, business critical workloads. It provides an always on fixed capacity corridor, that sits on top of purchased capacity. This capacity is charged based on usage when workloads spike above the base, purchased capacity and into the corridor.

The offering consists of three priced items: a subscription fee based on the size of the corridor, a usage fee, and a maintenance fee. Customers pay a monthly subscription and maintenance fee for this based on the amount of capacity. If the workload spikes into this corridor throughout the day, customers will pay a usage fee based on their hourly peak. Spikes into the corridor for more than 4 hours, and the usage charge, changes to the daily rate.

It's important to note that this new offering is for general purpose central processing units (CPs) only, and in order to take advantage of the new Hardware Consumption Solution, the customer should have Tailored Fit Pricing for their software installed; they must be on the current IBM z15 generation of hardware; and must have a Tailored Fit Pricing for hardware workload analysis done. The Tailored Fit Pricing for hardware analysis needs to be run on the customer's workload over the last few months to determine if this new Hardware Consumption Offering is appropriate and will help with workload spikes. This analysis helps determine if additional base capacity might be a better option and, if Tailored Fit Pricing – Hardware Consumption Solution is the right solution, how large the corridor should be sized.

### ***New with the May 4, 2021 announcements***

#### **What is the new Remote Code Load for IBM Z Firmware?**

Remote Code Load for IBM Z Firmware optimizes resilience and keeps your IBM Z system up to date with the latest features, fixes, and maintenance without requiring someone to be physically in the data center to install and monitor planned updates to your system. This optional feature provides secure, remote installation and monitoring by IBM for planned updates to your IBM Z system

#### **What is IBM Data Privacy Passports?**

Data Privacy Passports is a data centric security solution that enables eligible data to play an active role in its own protection. It lets you implement field level data protection to protect that data throughout its lifecycle. IBM Data Privacy Passports, in conjunction with IBM z15, and available via an IBM z15 only PID, is 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 protection that can span hybrid and multiparty computing environments. Data Privacy Passports now exploits the latest release of Hyper Protect Virtual Servers, which delivers improved exploitation of key IBM z15 features, with more exploitation continuously being introduced.

For more information about IBM Data Privacy Passports V1.0 announced on March 10, 2020 see Software Announcement 220-062.

The IBM Data Privacy Passports FAQ can be found here:  
<https://www.ibm.com/downloads/cas/L8EWKEP9>

## **What is the new IBM Enterprise Key Management Foundation Web Edition?**

As z/OS® data set encryption becomes more widely adopted, organizations also must contend with an ever-growing set of encryption keys. Effective management of these keys is essential to ensure both the availability and security of the encrypted information. Centralized management of keys is necessary to perform the complex tasks related to key generation, renewal, backup and recovery.

The IBM Enterprise Key Management Foundation Web Edition (EKMF Web) is a flexible and highly secure key management system for the enterprise.

The new IBM Key Management Foundation Web Edition (EKMF Web) offers real-time, centralized secure management of keys in an enterprise. It works with the IBM Crypto Express Hardware Security Module (HSM) for secure key generation. Use of EKMF Web requires Trusted Key Entry (TKE) and the key repository is Db2® based and hosted on one z/OS LPAR. EKMF Web provides policy based key generation, policy based key rotation, key usage tracking, and key backup & recovery.

For more information refer to:

<https://www.ibm.com/products/enterprise-key-management-foundation-web>

## **What are the new Sysplex Time Protocol (STP) enhancements for Precision Time Protocol (PTP)?**

The April 14, 2020 announcement for z15 T01 and z15 T02 introduced a new external time source option for STP, called Precision Time Protocol (PTP), which is the subject of the IEEE 1588 standard. PTP provides much better time accuracy for systems than was possible when using Network Time Protocol (NTP) alone as a time source. The use of PTP is most important for customers facing regulatory issues, such as in the financial industry, where tight synchronization to a Universal Coordinated Time source is demanded by new regulations.

IBM Z today can achieve excellent time accuracy where NTP together with Pulse Per Second are used as a time source, but IBM Z is likely just one part of your processing environment. Pulse Per Second is not an option for much of the other equipment in your data center, and it can be complicated to have to manage multiple timing protocols. The use of a PTP time source offers you an opportunity to begin to migrate to a single timing protocol that will meet all of your needs and help to meet the regulations. PTP does not alleviate the need to also use Pulse Per Second to meet the regulations for this particular generation of IBM Z. Pulse Per Second can be used in conjunction with PTP to achieve higher time accuracy than PTP alone. This limitation is planned to be removed with future systems, as noted in the Statement of Direction released with IBM z15 on September 12, 2019. This is a first step toward our goal of simplifying your data center time management needs.

## **What was announced on January 14, 2020?**

IBM fulfilled the September 2019 statement of direction for IBM z15 T01 and IBM LinuxONE III LT1 to deliver Fibre Channel Endpoint Security.

IBM Fibre Channel Endpoint Security is a new end-to-end solution for the IBM z15 T01 that is designed to provide a means to help ensure the integrity and confidentiality of all data flowing on Fibre Channel links between authorized server and storage devices, creating a trusted storage network that encrypts data in flight.

For more details check out the separate frequently asked questions document:

<https://www.ibm.com/downloads/cas/Y6E9KLA8>

## IBM z15 Model T02 Hardware

### How much capacity will the z15 T02 deliver?

The largest IBM z15 T02 is expected to provide approximately 15% more capacity than the largest IBM z14 Model ZR1 (z14 ZR1) with some variation based on workload and configuration. <sup>1</sup>

### What is the difference between a single frame z15 Model T01 and this new z15 Model T02 that is also a single frame?

The new z15 T02 is our entry model into the z15 family of server, so is based on the z15 processor chip. The z15 T02 is IBM's lower cost air-cooled single-frame system with an entry capacity of 98 Processor Capacity Index (PCI) and a maximum of six central processors (CPs) at 9,215 PCI. The z15 T01 is still the performance optimized, high frequency system. A single frame z15 T01 can have up to 5 drawers or up to 190 CP cores or significantly more capacity than the z15 T02 at over 21,000 PCI.

### What is the machine type of the z15 T02?

The IBM z15 T02 machine type is 8562.

### What is the sizing structure for the z15 T02?

The z15 T02 will follow the z15 T01 and the z14 ZR1 concept of feature-based sizing instead of models. There is one orderable model – the T02 – and there will be 5 new feature codes for ordering capacity – the Max4, Max13, Max21, Max31 and Max65.

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. An example is a customer requiring 6 Central Processors (CPs), 6 zIIPs and 3 IFLs – a total of 15 configured cores. The customer could choose a Max21, Max31 or Max65 depending on the growth they were expecting, or the possibility of wanting the enhanced availability of the Max65.

### What is the available subcapacity granularity on the z15 T02?

The z15 T02 can have up to six Central Processors (CPs) per server – the same as the z14 ZR1. There are up to 26 capacity settings (A0x-Z0x) available for those CPs – so a total of 156 capacity settings. For more information refer to the IBM z15 T02 datasheet:

<https://www.ibm.com/downloads/cas/0VZ0KE68>

The IBM z15 spec sheet: <https://www.ibm.com/downloads/cas/JD90D9QR>

### What is meant by the statement about “enhanced availability on the z15 T02 Max65”?

The z15 T02 design incorporates two Central Processor Complex (CPC) drawers for the Max65. The second drawer enhances availability by allowing Concurrent Drawer Repair (CDR) – also referred to as Enhanced Drawer Availability (EDA), a feature previously not available on single frame servers. CDR allows a single CPC drawer in a multi-drawer server, if there is sufficient memory installed, to be removed, serviced and reinstalled without bringing the system down.

---

<sup>1</sup> Based on preliminary internal measurements and projections and compared to the z14 ZR1. 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/lspindex>.



## **I have lots of spare racks in my datacenter, can I supply the rack for the IBM z15 T02?**

No, the z15 will still be fully assembled, and tested at an IBM plant prior to being packaged for shipping to our customers 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 T02?**

It is possible to upgrade any IBM z14 ZR1 to the IBM z15 T02. You can do a migration offering (similar to the upgrade but you are unable to retain the serial number) upgrade from any IBM z13s®.

There are some nuances to the z15 T02 to z15 T02 upgrades that you need to know.

- There are concurrent upgrades from Max31 to Max65 and other upgrades will be disruptive.
- There is no upgrade to the z15 T01
- There are no MES path into the LinuxONE III models.
- There is an upgrade path from the LinuxONE III Model LT2 to the z15 T02.

## **What was the SOD about the 16U reserved space available in the IBM z15 T02?**

The April 14, 2020 announcement for the z15 T02 includes a Statement of Direction (SOD) for a 16U of reserved space in the z15 T02. The SOD is: In the future, IBM plans to test a co-located DS8910F solution that can be utilized in the 16U Reserved space for single phase power z15 T02 model.

Customers must consider leaving enough room for the reserved space and staying with the single phase power option, if they would like to consider configuring this option and co-locating their storage in the future. Note: All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

## **What is 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.

The System Recovery Boost GDPS enhancements are a combination of changes to GDPS processing/scripting itself, and GDPS taking advantage of new functionality being offered by the z15 Support Element. Various APIs used by GDPS on the SE are being streamlined and improved, while GDPS itself will be doing more in parallel when possible. In addition, GDPS will be enhancing its scripting language to allow multiple systems to be addressed in parallel by some commands, as opposed to sequentially as before, and making changes to some control flows to improve efficiency.

In addition to the Speed Boost, zIIP Boost and GDPS enhancements capabilities that are included on the z15 at no additional charge, on the IBM z15 T01 only (not available on the z15 T02), you also have the option to build upon this base functionality with System Recovery Boost Upgrade, a Capacity on Demand offering that lets you unlock additional zIIP capacity when you need it.

Additional information on System Recovery Boost can be found in the FAQ document:

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

*New with the May 4, 2021 announcements*

### **What is the new addition, Hardware Consumption Solution, to the Tailored Fit Pricing Model for IBM Z family?**

There is a new addition in the Tailored Fit Pricing for IBM Z family, the Hardware Consumption Solution. The new offering provides a simple, flexible cloud-like solution for unpredictable, spiky, business critical workloads. It provides an always on fixed capacity corridor, that sits on top of purchased capacity. This capacity is charged based on usage when workloads spike above the base, purchased capacity and into the corridor.

The offering consists of three priced items: a subscription fee based on the size of the corridor, a usage fee, and a maintenance fee. Customers pay a monthly subscription and maintenance fee for this based on the amount of capacity. If the workload spikes into this corridor throughout the day, customers will pay a usage fee based on their hourly peak. Spikes into the corridor for more than 4 hours, and the usage charge, changes to the daily rate.

It's important to note that this new offering is for general purpose central processing units (CPUs) only, and in order to take advantage of the new Hardware Consumption Solution, the customer should have Tailored Fit Pricing for software installed; they must be on the current IBM z15 generation of hardware; and must have a Tailored Fit Pricing for hardware workload analysis done. The Tailored Fit Pricing for hardware analysis needs to be run on the customer's workload over the last few months to determine if this new Hardware Consumption Offering is appropriate and will help with workload spikes. This analysis helps determine if additional base capacity might be a better option and, if Tailored Fit Pricing – Hardware Consumption Solution is the right solution, how large the corridor should be sized.

*New with the May 4, 2021 announcements*

### **What is the new Remote Code Load for IBM Z Firmware?**

Remote Code Load for IBM Z Firmware optimizes resilience and keeps your IBM Z system up to date with the latest features, fixes, and maintenance without requiring someone to be physically in the data center to install and monitor planned updates to your system. This optional feature provides secure, remote installation and monitoring by IBM for planned updates to your IBM Z system

### **What is the compression accelerator on the z15 T02?**

The IBM z15 T02 integrates file compression capabilities with an on-chip compression coprocessor. The Integrated Accelerator for zEDC can reduce data storage and communications requirements and costs, as well as increase data transfer rates to boost throughput, without adversely impacting response times. The Integrated Accelerator for zEDC replaces the IBM zEnterprise Data Compress (zEDC) Express adapter on earlier Z servers. Best of all, combined with pervasive encryption principles, you have the processing capacity and the confidence to compress and encrypt everything.

### **What is the on-chip accelerated sort for IBM z15?**

The Integrated Accelerator for Z Sort, standard on IBM z15, can help optimize elapsed time for eligible sort workloads, which typically occur during customer batch windows. By providing one sort accelerator per core, it can help accelerate frequently used functions to speed up sorting, shorten batch windows, and improve select database functions, such as reorganization. A new SORTL instruction, exploited by DFSORT™ and the Db2 for z/OS Utilities Suite, leverages the new IBM Integrated Accelerator for Z Sort to help reduce CPU usage and can help improve elapsed time for eligible sort workloads. Review the latest DFSORT and Db2 for z/OS Utilities Suite PTFs for exploitation of the Integrated Accelerator for Z Sort.

## What is the new Secure Execution for Linux?

Secure Execution for Linux on Z is a new capability introduced on April 14, 2020 for the IBM z15 family – available for both the z15 T01 and the z15 T02. Engineered to help protect against insider and outsider threats in multi-tenanted cloud environments, it ensures users, and even system administrators, are unable to access sensitive data in Linux-based virtual environments. Secure Execution for Linux protects the confidentiality and integrity of data, by isolating data at the virtual machine level, and ensuring that only the people within the organization that have a need-to-know have access to data in the clear.

For more information see the Secure Execution solution brief:

<https://www.ibm.com/downloads/cas/O158MBWG>

And the Secure Execution customer brief at: <https://www.ibm.com/downloads/cas/GPLNZLE2>

## What was the increase to coupling capabilities in the z15 T02?

In the z15 T02 there is a huge increase in coupling capabilities relative to the z14 ZR1 -- largely enabled by the possibility of a second drawer. The z15 T02 footprint enables more processing capacity and more coupling connectivity:

- 3X as many ICA SR links - the maximum number of physical ICA SR coupling links (ports) is increased from a maximum of 16 per CPC on z14 ZR1 to 48 per CPC.
- 2X as many Internal Coupling Channels (ICPs) and 1.5 as many total coupling channels CHPIDs- The maximum number ICP is increased from 32 on z14 ZR1 to 64, and the maximum overall number of allowed coupling channels (CHPIDs) per CPC is increased from 256 to 384.
- 2X as many CE LR links - There is an increase in maximum number of CE LR coupling links from 32 to 64 (ZR1 to T02).
- 2X as many Internal Coupling Facility (ICF) processors – up to 65 compared to 30 on the z14 ZR1.

## What is the Hardware Management Appliance?

New for the z15, the IBM Z Hardware Management Appliance feature provides the capability for the Hardware Management Console (HMC) and Support Element (SE) functionality to be contained within redundant physical servers inside the server frame. When you order the IBM Z Hardware Management Appliance feature, this will provide logically a Primary and Alternate Support Element and two peer Hardware Management Consoles on two physical servers in the CPC frame. This eliminates the need for having to manage a separate physical server or servers for one or more HMCs outside of the frame. For the User Interface experience, you must use remote browsing controls from your own workstation into HMC within the IBM Z Hardware Management Appliance.

If you have multiple systems, you don't need to order the Hardware Management Appliance feature or the HMCs for all servers. The recommendation is that you consider having the IBM Z Hardware Management Appliance features on one or two servers, but the rest of the servers don't need to include Hardware Management Appliance features.

The IBM Z Hardware Management Appliance feature is optional and MUST be ordered initially. It cannot be added after installation has occurred.

## What is the HSA and how much is there on the z15 T02?

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 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 T02 there is 160 GB of HSA, independent of customer-purchased memory.

## How much memory is available on the z15 T02?

The z15 T02 server supports up to 16TB of real memory per server, 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 64 GB for all models. The z15 T02 is designed with redundant array of independent memory (RAIM) technology. RAIM is analogous to what is known in the disk storage industry as RAID.

Please note that the listed maximum memory is customer-usable, RAIM-protected memory. RAIM is always active, and IBM has already factored in additional RAIM memory supplying redundancy.

## Tell me about memory on the 2<sup>nd</sup> CPC drawer z15 T02?

The z15 T02 two CPC drawer system divides the memory between drawers. Memory may or may not be equal on each drawer. To support continuous availability through CPC drawer repairs, verify that the memory configuration on each drawer will support your LPAR configuration.

## What value might I experience if I purchase larger memory on the z15 T02?

Memory on the z15 T02 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 customers 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.

## How many System Assist Processors (SAPs) are on the z15 T02?

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

- The z15 T02 Max4 and Max13 have two SAPs.
- The z15 T02 Max21 has three SAPs.
- The z15 T02 Max31 has four SAPs.
- The z15 T02 Max65 has eight 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 T02?

IBM ships the z15 T02 Max4 and Max13 with one spare. The z15 T02 Max21, Max31 and Max65 have two spare processing cores, as well as any uncharacterized core that isn't currently in use can also be used as a spare.

## Can any of the z15 T02 spare processing cores be used for other purposes?

No, the minimum allotment of z15 T02 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 T02 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 T02, using a model capacity identifier of A00, with a maximum of 65 IFLs or ICFs. There is still a limit of 16 ICF engines for a single Coupling Facility LPAR. Besides using the capacity identifier of A00, IBM does offer Linux only servers based on z15 T02 technology – the IBM LinuxONE III Model LT2. The LinuxONE systems may optionally add one CP (at C01 capacity) but it can only be used for the GDPS Virtual Appliance.

## **What are the new Sysplex Time Protocol (STP) enhancements for Precision Time Protocol (PTP)?**

The April 14, 2020 announcement for z15 T01 and z15 T02 introduced a new external time source option for STP, called Precision Time Protocol (PTP), which is the subject of the IEEE 1588 standard. PTP provides much better time accuracy for systems than was possible when using Network Time Protocol (NTP) alone as a time source. The use of PTP is most important for customers facing regulatory issues, such as in the financial industry, where tight synchronization to a Universal Coordinated Time source is demanded by new regulations.

IBM Z today can achieve excellent time accuracy where NTP together with Pulse Per Second are used as a time source, but IBM Z is likely just one part of your processing environment. Pulse Per Second is not an option for much of the other equipment in your data center, and it can be complicated to have to manage multiple timing protocols. The use of a PTP time source offers you an opportunity to begin to migrate to a single timing protocol that will meet all of your needs and help to meet the regulations. PTP does not alleviate the need to also use Pulse Per Second to meet the regulations for this particular generation of IBM Z. Pulse Per Second can be used in conjunction with PTP to achieve higher time accuracy than PTP alone. This limitation is planned to be removed with future systems, as noted in the Statement of Direction released with IBM z15 on September 12, 2019. This is a first step toward our goal of simplifying your data center time management needs.

## **Tell me about the frame on the z15 T02 – is it similar to the z14 ZR1?**

The IBM z15 T02 is built with the same 19” form factor available on the z14 ZR1. For customers that are on z13s and older technology, they should see a reduction in the amount of floor space taken up within the data center. The system can now fit within most data center aisles – offering new flexibility in data center layout, space efficiency, cabling efficiency and participation in new data center air-cooling containment systems.

## **How is the single frame z15 T01 air-cooled different than the z15 T02 air-cooled?**

The z15 T01 air-cooled single frame has a radiator heat exchanger built into the system. The processors are cooled with an internal water loop to provide maximum cooling and performance. At the frame level the system is air-cooled. The z15 T02 processors are air-cooled using custom heat sinks providing a better price performance option.

## **What is the integrated firmware processor (IFP) in z15 T02?**

The integrated firmware processor (IFP) is allocated from the pool of processors for the entire system. Unlike other characterized PUs, the IFP is standard and are not defined by the customer. It is a single PU that is dedicated solely to supporting the native Peripheral Component Interconnect Express (PCIe) I/O feature management and virtualization functions and other functions such as Stand Alone Coupling Facility Dynamic I/O Reconfiguration and Resource Group management.

## **Tell me about the power options in the z15 T02?**

The z15 T02 uses an intelligent power distribution unit (iPDU). It is intelligent, fully managed and fully integrated within the system's power distribution, control, and maintenance support structure.

The z15 T02 supports either Single Phase or Three Phase Input power depending on configurations. Single Phase can be used for 1 CPC drawer and up to 4 I/O drawers. 3 Phase can be used for all configurations. If you select the feature code to reserve the 2<sup>nd</sup> CPC drawer, or if you order a Max65, you are required to use 3 Phase power.

### **What is the new IBM Enterprise Key Management Foundation Web Edition?**

As z/OS data set encryption becomes more widely adopted, organizations also must contend with an ever-growing set of encryption keys. Effective management of these keys is essential to ensure both the availability and security of the encrypted information. Centralized management of keys is necessary to perform the complex tasks related to key generation, renewal, backup and recovery.

The IBM Enterprise Key Management Foundation Web Edition (EKMF Web) is a flexible and highly secure key management system for the enterprise.

The new IBM Key Management Foundation Web Edition (EKMF Web) offers real-time, centralized secure management of keys in an enterprise. It works with the IBM Crypto Express Hardware Security Module (HSM) for secure key generation. Use of EKMF Web requires Trusted Key Entry (TKE) and the key repository is Db2 based and hosted on one z/OS LPAR. EKMF Web provides policy based key generation, policy based key rotation, key usage tracking, and key backup & recovery.

For more information refer to:

<https://www.ibm.com/products/enterprise-key-management-foundation-web>

## IBM z15 Model T01 Hardware

### How much capacity will the z15 T01 deliver?

The largest IBM z15 T01 is expected to provide approximately 25% more capacity than the largest IBM z14 (z14) with some variation based on workload and configuration. <sup>2</sup>

### Is the one frame IBM z15 T01 a replacement for the IBM z14 ZR1?

No. The one frame z15 T01 system is not the replacement for z14 ZR1. It is still the performance optimized, high frequency system. A single frame z15 T01 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 T01?

The IBM z15 T01 machine type is 8561.

### How many models are they for the z15?

There are two model types for the z15. The z15 T01 and the air-cooled single frame server the z15 T02 that was announced April 14, 2020.

This follows the models on the z14 with the use of feature-based sizing instead of models. For the z15 T01 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 z15 T01 customer requiring 20 Central Processors (CPs), 40 zIIPs and 5 IFLs – a total of 65 configured cores. The customer could choose a Max71 or Max 108 depending on the growth they were expecting.

### What is the available subcapacity granularity on the z15 T01?

The z15 T01 can have up to 34 subcapacity Central Processors (CPs) per server while the z14 M01-M05 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 T01 datasheet:

<https://www.ibm.com/downloads/cas/NN7GBPJ1>

and the IBM z15 T01 spec sheet: <https://www.ibm.com/downloads/cas/2GO7LRDV>

---

<sup>2</sup> 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/resourceink/lib03060.nsf/pages/lspindex>.

## **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 customers 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 z14 to z15 upgrades that you need to know.

- There are concurrent upgrades from Max34 to Max 71 to Max108.<sup>3</sup>
- 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 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 customer-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 additional RAIM memory supplying redundancy.

---

<sup>3</sup> 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.



### **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 customers 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 entire system. Unlike other characterized PUs, the IFP is standard and not defined by the customer. It is a single PU that is dedicated solely to supporting the native Peripheral Component Interconnect Express (PCIe) I/O feature management and virtualization functions.

### **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, as well as any uncharacterized core can be used as a spare. 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 customers 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.

Customer 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.)

These questions are from an earlier section in the FAQ and are repeated here. They highlight some of the new features of the z15 that were announced on April 14, 2020

## What is the new Secure Execution for Linux?

Secure Execution for Linux on Z is a new capability introduced on April 14, 2020 for the IBM z15 family – available for both the z15 T01 and the z15 T02. Engineered to help protect against insider and outsider threats in multi-tenanted cloud environments, it ensures users, and even system administrators, are unable to access sensitive data in Linux-based virtual environments. Secure Execution for Linux protects the confidentiality and integrity of data at scale, by isolating data at the virtual machine level, and ensuring that only the people within the organization that have a need-to-know have access to data in the clear.

For more information see the Secure Execution solution brief:

<https://www.ibm.com/downloads/cas/O158MBWG>

And the Secure Execution customer brief at: <https://www.ibm.com/downloads/cas/GPLNZLE2>

## What is the new Hardware Management Appliance?

New for the z15 family, the new IBM Z Hardware Management Appliance feature provides the capability for the Hardware Management Console (HMC) and Support Element (SE) functionality to be contained within redundant physical servers inside the server frame. When you order the IBM Z Hardware Management Appliance feature, this will provide logically a Primary and Alternate Support Element and two peer Hardware Management Consoles on two physical servers in the CPC frame. This eliminates the need for having to manage a separate physical server or servers for one or more HMCs outside of the frame. For the User Interface experience, you must use remote browsing controls from your own workstation into HMC within the IBM Z Hardware Management Appliance.

If you have multiple systems, you don't need to order the Hardware Management Appliance feature for all servers. The recommendation is that you consider having the IBM Z Hardware Management Appliance features on one or two servers, but the rest of the servers don't need to include Hardware Management Appliance features. Although those servers would have redundant SEs.

The IBM Z Hardware Management Appliance feature is optional.

## What is System Recovery Boost on z15?

Built into IBM z15, System Recovery Boost is an innovative solution that lets you 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 use your already-entitled Central Processors and zIIPs to unleash additional processing capacity on an LPAR-by-LPAR basis, providing a fixed-duration performance boost

to reduce the length and mitigate the impact of downtime. This 'boost' can be used twice per IPL (30-minute boost for shutdown and 60-minute boost for startup). The value for customers is:

- Faster shutdown - accelerate image shutdown to prepare for planned shutdown activities.
- Faster GDPS automation actions - drive faster and more efficient GDPS automation actions to rapidly reconfigure and recover your environment.
- Faster restart & recovery - expedite the restart and recovery of images, middleware environments and customer workloads to accelerate return to pre-shutdown SLAs.
- Faster elimination of workload backlog - utilize additional capacity for a fixed period during recovery, so you can process backlog faster after planned or unplanned downtime.

Additional information on System Recovery Boost can be found in the FAQ document:  
<https://www.ibm.com/downloads/cas/1NWEJKOX>

### What is IBM System Recovery Boost Upgrade?

System Recovery Boost Upgrade allows you to build upon System Recovery Boost base functionality with an optional Capacity on Demand offering. The option that can help to maximize performance and parallelism during the boost period by unlocking up to 20 additional zIIP engines for up to 6 hours to build upon the base functionality. **The System Recovery Boost Upgrade option is only available on the z15 T01.**

### What was announced on April 14<sup>th</sup> 2020 concerning System Recovery Boost Upgrade?

On the z15 T01, new with the April 14, 2020 announcements, IBM has increased the term for System Recovery Boost Upgrade on IBM z15 Model T01 from one year to five years, available in one-year increments. New feature code 6799 will record how many subscription years you purchase. Feature code 6802, SRB Upgrade Record, will still enable temporary activation of additional physical zIIP processors to be used in conjunction with System Recovery Boost. Feature code 9930, SRB Upgrade Authorization, is still required to enable the ordering of SRB Upgrade Record (#6802). Feature code 9930 drives the necessary contracts to use the Boost Record. Feature code 6799 further qualifies the use of these existing feature codes.

As announced on November 26, 2019, no specific z15 feature codes need to be ordered to be able to use the base capabilities of System Recovery Boost. Feature codes 6802, 6799 and 9930 are required only to make use of the optional System Recovery Boost Upgrade Record, which allows the temporary activation of additional zIIP processor resources in conjunction with System Recovery Boost. Note that the System Recovery Boost Upgrade capability is only available on the z15 T01.

### What was announced on August 4<sup>th</sup> 2020 regarding expedited Parallel Sysplex recovery?

In September 2020, IBM will deliver new System Recovery Boost capabilities that will support a new class of boost, called recovery process boosts, that can be applied to a range of sysplex recovery processes, including sysplex partitioning, CF structure recovery, CF data sharing member recovery, and HyperSwap recovery to minimize disruption of recovery processes and maximize sysplex resiliency. Now, you can use System Recovery Boost to mitigate the disruption caused by these recovery processes and maximize the benefits of your sysplex environment. The new short-duration recovery process boosts take advantage of the same GP Speed Boost and zIIP Boost technologies used to accelerate IPLs, but are instead used to accelerate one or more of the sysplex recovery processes I mentioned before. Each individual recovery process boost can last up to five minutes, and you can use up to thirty minutes worth of these boosts in a consecutive 24-hour period.

***New with the May 4, 2021 announcements***

## **What is the new addition, Hardware Consumption Solution, to the Tailored Fit Pricing Model for IBM Z family?**

There is a new addition in the Tailored Fit Pricing for IBM Z family, the Hardware Consumption Solution. The new offering provides a simple, flexible cloud-like solution for unpredictable, spiky, business critical workloads. It provides an always on fixed capacity corridor, that sits on top of purchased capacity. This capacity is charged based on usage when workloads spike above the base, purchased capacity and into the corridor.

The offering consists of three priced items: a subscription fee based on the size of the corridor, a usage fee, and a maintenance fee. Customers pay a monthly subscription and maintenance fee for this based on the amount of capacity. If the workload spikes into this corridor throughout the day, customers will pay a usage fee based on their hourly peak. Spikes into the corridor for more than 4 hours, and the usage charge, changes to the daily rate.

It's important to note that this new offering is for general purpose central processing units (CPUs) only, and in order to take advantage of the new Hardware Consumption Solution, the customer should have Tailored Fit Pricing for software installed; they must be on the current IBM z15 generation of hardware; and must have a Tailored Fit Pricing for hardware workload analysis done. The Tailored Fit Pricing for hardware analysis needs to be run on the customer's workload over the last few months to determine if this new Hardware Consumption Offering is appropriate and will help with workload spikes. This analysis helps determine if additional base capacity might be a better option and, if Tailored Fit Pricing – Hardware Consumption Solution is the right solution, how large the corridor should be sized.

**New with the May 4, 2021 announcements,**

## **What is the new Remote Code Load for IBM Z Firmware?**

Remote Code Load for IBM Z Firmware optimizes resilience and keeps your IBM Z system up to date with the latest features, fixes, and maintenance without requiring someone to be physically in the data center to install and monitor planned updates to your system. This optional feature provides secure, remote installation and monitoring by IBM for planned updates to your IBM Z system

## **What is the new IBM Enterprise Key Management Foundation Web Edition?**

As z/OS data set encryption becomes more widely adopted, organizations also must contend with an ever-growing set of encryption keys. Effective management of these keys is essential to ensure both the availability and security of the encrypted information. Centralized management of keys is necessary to perform the complex tasks related to key generation, renewal, backup and recovery.

The IBM Enterprise Key Management Foundation Web Edition (EKMF Web) is a flexible and highly secure key management system for the enterprise.

The new IBM Key Management Foundation Web Edition (EKMF Web) offers real-time, centralized secure management of keys in an enterprise. It works with the IBM Crypto Express Hardware Security Module (HSM) for secure key generation. Use of EKMF Web requires Trusted Key Entry (TKE) and the key repository is Db2 based and hosted on one z/OS LPAR. EKMF Web provides policy based key generation, policy based key rotation, key usage tracking, and key backup & recovery.

For more information refer to:

<https://www.ibm.com/products/enterprise-key-management-foundation-web>

## **What are the new Sysplex Time Protocol (STP) enhancements for Precision Time Protocol (PTP)?**

The April 14, 2020 announcement for z15 T01 and z15 T02 introduced a new external time source option for STP, called Precision Time Protocol (PTP), which is the subject of the IEEE 1588 standard. PTP provides much better time accuracy for systems than was possible when using Network Time Protocol (NTP) alone as a time source. The use of PTP is most important for customers facing regulatory

issues, such as in the financial industry, where tight synchronization to a Universal Coordinated Time source is demanded by new regulations.

IBM Z today can achieve excellent time accuracy where NTP together with Pulse Per Second are used as a time source, but IBM Z is likely just one part of your processing environment. Pulse Per Second is not an option for much of the other equipment in your data center, and it can be complicated to have to manage multiple timing protocols. The use of a PTP time source offers you an opportunity to begin to migrate to a single timing protocol that will meet all of your needs and help to meet the regulations. PTP does not alleviate the need to also use Pulse Per Second to meet the regulations for this particular generation of IBM Z. Pulse Per Second can be used in conjunction with PTP to achieve higher time accuracy than PTP alone. This limitation is planned to be removed with future systems, as noted in the Statement of Direction released with IBM z15 on September 12, 2019. This is a first step toward our goal of simplifying your data center time management needs.

## 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/resourcelink/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 customer 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](mailto: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, which are workload dependent. Consult the LSPR for 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 T02 Software Pricing**

### **What software pricing is announced with the z15 Model T02?**

A new Technology Transition Offering (TTO) called Technology Update Pricing for the z15 Model T02 (TU8) is being announced along with revisions to the Technology Update Pricing for the z15 T02 offering and three revised Transition Charges for Sysplexes or Multiplexes offerings.

Technology Update Pricing for the z15 Model T02 (TU8) 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 z15 Model T02 (TU8) 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 z15 Model T02 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 z15 (TU7) 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 customers have implemented Country Multiplex Pricing.

### **What is the price reduction available with the Technology Update Pricing for z15 Model T02?**

The percent reduction in the monthly AEWLC is based on the number of z15 Model T02 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 z15 Model T02 is reduced by the percentage in the table below.

Schedule of AEWLC reductions for Technology Update Pricing for the z15 Model T02 (TU8)

<b>MSUs: Quantity of z15 Model T02 Full Capacity MSUs for a stand-alone server</b>	<b>Reduction in Monthly AEWLC</b>
1-12 MSUs	23.0%
13-17 MSUs	23.0%
18-30 MSUs	23.0%
31-45 MSUs	20.0%
46-87 MSUs	19.0%
88-175 MSUs	19.0%
176-260 MSUs	17.0%
261-315 MSUs	17.0%
316-390 MSUs	17.0%
391 - more MSUs	17.0%

The number of MSUs used to determine the MSUs of a stand-alone z15 Model T02 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 z14 Model ZR1 to a z15 Model T02 will software pricing improvements be additive?**

You will receive the pricing advantage of Technology Update Pricing for the z15 Model T02 (TU8) for the z15 Model T02. The Technology Update Pricing for the z15 Model T02 (TU8) price reduction includes the benefits of the prior price reductions. Technology Update Pricing for the z15 Model T02 (TU8) is available when upgrading from a z14 Model ZR1 to a z15 Model T02.

Schedule of AEWLC reductions for Technology Update Pricing for the z15 Model T02 (TU8)

<b>MSUs: Quantity of z15 T02 Full Capacity MSUs for a stand-alone server</b>	<b>Reduction in Monthly AEWLC for z15 T02 (TU8)</b>	<b>Reduction in Monthly AEWLC for z14 ZR1 (TU6)</b>
1-12 MSUs *	23.0%	18.0%
13-17 MSUs *	23.0%	18.0%
18-30 MSUs	23.0%	18.0%
31-45 MSUs	20.0%	15.0%
46-87 MSUs	19.0%	14.0%
88-175 MSUs	19.0%	14.0%
176-260 MSUs	17.0%	13.0%
261-315 MSUs	17.0%	13.0%
316-390 MSUs	17.0%	13.0%
391 - more MSUs	17.0%	13.0%

\*The MSUs for TU6 are 1-11 and 12-17 respectively

The number of MSUs used to determine the MSUs of a stand-alone z15 Model T02 server are based on the announced IBM full capacity ratings that can be found on the Mainframe Exhibits section of the System z Software Contracts website at:

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

### What software pricing announcements were made for a z15 Model T02 in a Parallel Sysplex or Multiplex?

IBM announced revisions to the Technology Update Pricing for z15 (TU7) offering and three revised Transition Charges for Sysplexes offerings.

The revisions to the Technology Update Pricing for z15 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 z15 offers actively coupled Parallel Sysplexes and Loosely Coupled Complexes that consist entirely of z15 (Model T01 and Model T02) servers the benefit of aggregated Technology Update Pricing for z15 when priced with AWLC, or in a Multiplex priced with Country Multiplex License Charges (CMLC).

Transition Charges for Sysplexes or Multiplexes (TC5): When two or more servers that consist of one or more z15 (all models) servers with one or more z14 (all models) 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 (TU7), which is based on the amount of server capacity within the TTO-eligible environment.

Transition Charges for Sysplexes or Multiplexes (TC4): When two or more servers that consist of one or more z15 (all models) servers with one or more z14 (all models), 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 (TU5), which is based on the amount of server capacity within the TTO-eligible environment.

Transition Charges for Multiplexes (TC3): For existing qualified multiplex only, when two or more servers consisting of one or more z15 (all models), z14 (all models), IBM z13®, z13s, IBM zEnterprise® zEC12, or IBM zEnterprise zBC12 servers 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 (TU3), which is based on the amount of server capacity within the TTO eligible environment.

If available for your specific configuration, you have a choice of selecting either Technology Update Pricing for the z15 (TU7) or PSLC, but not both, for your IBM software programs within the z/OS or z/TPF operating system families on z15 (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 z15 Model T02?**

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 z15 T02 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 z15 Model T02?**

In order to get Sub-Capacity Technology Update Pricing for the z15 Model T02, the following contract is required to be executed:

ICA Attachment for IBM System z AEWLC, (Z125-8755), or, if applicable, ICA Attachment for IBM Country Multiplex Pricing, (Z126-6965).

The Supplement for Technology Transition Offerings to Attachment for IBM System z Advanced Workload License Charges, Attachment for IBM System z Advanced Entry Workload License Charges, Attachment for Country Multiplex Pricing, and to IBM System z Machines Exhibit (Z125-8994) describes the Technology Transition Offerings and should be delivered to each z15 Model T02 customer, but it does not get executed.

The System z Machines Exhibit (Z125-3901) provides the terms for Full Capacity AEWLC. The Exhibit and the Supplement should be delivered to each z15 Model T02 customer, but they do not get executed.

### **What is the Processor Value Unit (PVU) rating used for Passport Advantage® software on the z15 Model T02?**

The PVU rating for the z15 Model T02 will be 100 for both IFL engines and CP engines, the same as the z14 Model ZR1, Rockhopper II, z13s, Rockhopper, zBC12, z114 and the z10 BC. The PVU rating for the z15, z14 Models M01 – M05, Emperor II, z13, Emperor, zEC12, z196 and z10 EC remains 120 for both IFL engines and CP engines.

### **What is the Tailored Fit Pricing for IBM Z pricing model?**

IBM introduced Tailored Fit Pricing as a simpler pricing model to allow customers to better leverage their existing platform investments as their business demands, and in a more cost competitive way. IBM's Tailored Fit Pricing model can address key concerns: Complexity of Sub-capacity pricing model leading to Z being managed as a cost center; difficulty in establishing cost of new workload deployment and the impact on cost of existing workloads; investment in tools and resources to manage sub-capacity that can inflate costs; lack of development and test resources; buying hardware for peak capacity to just to handle short term spikes, With simplicity as the core principle, this pricing model provides customers with an opportunity to more fully leverage Z capabilities and increase business value.

*New with the May 4, 2021 announcements*

## **What is the new addition, Hardware Consumption Solution, to the Tailored Fit Pricing Model for IBM Z family?**

There is a new addition in the Tailored Fit Pricing for IBM Z family, the Hardware Consumption Solution. The new offering provides a simple, flexible cloud-like solution for unpredictable, spiky, business critical workloads. It provides an always on fixed capacity corridor, that sits on top of purchased capacity. This capacity is charged based on usage when workloads spike above the base, purchased capacity and into the corridor.

The offering consists of three priced items: a subscription fee based on the size of the corridor, a usage fee, and a maintenance fee. Customers pay a monthly subscription and maintenance fee for this based on the amount of capacity. If the workload spikes into this corridor throughout the day, customers will pay a usage fee based on their hourly peak. Spikes into the corridor for more than 4 hours, and the usage charge, changes to the daily rate.

It's important to note that this new offering is for general purpose central processing units (CPUs) only, and in order to take advantage of the new Hardware Consumption Solution, the customer should have Tailored Fit Pricing for software installed; they must be on the current IBM z15 generation of hardware; and must have a Tailored Fit Pricing for hardware workload analysis done. The Tailored Fit Pricing for hardware analysis needs to be run on the customer's workload over the last few months to determine if this new Hardware Consumption Offering is appropriate and will help with workload spikes. This analysis helps determine if additional base capacity might be a better option and, if Tailored Fit Pricing – Hardware Consumption Solution is the right solution, how large the corridor should be sized.

## **Where can I get more information on IBM Z software charges?**

Please refer to: [www.ibm.com/systems/z/swprice/](http://www.ibm.com/systems/z/swprice/)

## z15 T01 Software Pricing

### What software pricing is announced with the z15 T01?

Four new or updated Technology Transition Offerings are being announced. They are Technology Update Pricing for the z15 T01 and one new and two revised Transition Charges for Sysplexes or Multiplexes programs.

Technology Update Pricing for the z15 T01 (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 T01 server
- In an aggregated z/OS Parallel Sysplex® with a z15 T01 server
- In an aggregated z/TPF Loosely Coupled Complex with a z15 T01 server
- In a Multiplex with a z15 T01 server

### What software pricing is available on a standalone z15 T01 or a z/OS Parallel Sysplex or a z/TPF Loosely Coupled Complex or a Multiplex with all z15 T01 technology?

For a standalone z15 T01, 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 T01. The Technology Update Pricing for z15 T01 provides reduced AWLC or CMLC pricing on the z15 T01 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 T01?

The percent reduction in the monthly AWLC or CMLC is based on the number of z15 T01 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 T01 or Parallel Sysplex or Loosely Coupled Complex or Multiplex of all z15 T01 servers is reduced by the percentage in the table below.

Schedule of AWLC or CMLC reductions for Technology Update Pricing for the z15 T01 (TU7)

Quantity of z15 T01 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 T01 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 T01 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 T01 servers with Tailored Fit Pricing, you will qualify for Technology Update Pricing for the z15 T01 (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 T01 (TU7) for the z15 T01 server. The Technology Update Pricing for the z15 T01 (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.

<b>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</b>	<b><u>z14 – TU5</u></b> Reduction in Monthly AWLC or CMLC	<b><u>z15 – TU7</u></b> 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 T01 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)

<b>Quantity of z15 and z14 Full Capacity MSUs in an actively coupled Parallel Sysplex, Loosely Coupled Complex, or Multiplex</b>	<b>Reduction in Monthly AWLC</b>
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 T01 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 T01 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 T01?**

The PVU rating for the z15 T01 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>

## **What is the Tailored Fit Pricing for IBM Z pricing model?**

IBM introduced Tailored Fit Pricing as a simpler pricing model to allow customers to better leverage their existing platform investments as their business demands, and in a more cost competitive way. IBM's Tailored Fit Pricing model can address key concerns: Complexity of Sub-capacity pricing model leading to Z being managed as a cost center; difficulty in establishing cost of new workload deployment and the impact on cost of existing workloads; investment in tools and resources to manage sub-capacity that can inflate costs; lack of development and test resources; buying hardware for peak capacity to just to handle short term spikes, With simplicity as the core principle, this pricing model provides customers with an opportunity to more fully leverage Z capabilities and increase business value.

*New with the May 4, 2021 announcements,*

## **What is the new addition, Hardware Consumption Solution, to the Tailored Fit Pricing Model for IBM Z family?**

There is a new addition in the Tailored Fit Pricing for IBM Z family, the Hardware Consumption Solution. The new offering provides a simple, flexible cloud-like solution for unpredictable, spiky, business critical workloads. It provides an always on fixed capacity corridor, that sits on top of purchased capacity. This capacity is charged based on usage when workloads spike above the base, purchased capacity and into the corridor.

The offering consists of three priced items: a subscription fee based on the size of the corridor, a usage fee, and a maintenance fee. Customers pay a monthly subscription and maintenance fee for this based on the amount of capacity. If the workload spikes into this corridor throughout the day, customers will pay a usage fee based on their hourly peak. Spikes into the corridor for more than 4 hours, and the usage charge, changes to the daily rate.

It's important to note that this new offering is for general purpose central processing units (CPs) only, and in order to take advantage of the new Hardware Consumption Solution, the customer should have Tailored Fit Pricing for software installed; they must be on the current IBM z15 generation of hardware; and must have a Tailored Fit Pricing for hardware workload analysis done. The Tailored Fit Pricing for hardware analysis needs to be run on the customer's workload over the last few months to determine if this new Hardware Consumption Offering is appropriate and will help with workload spikes. This analysis helps determine if additional base capacity might be a better option and, if Tailored Fit Pricing – Hardware Consumption Solution is the right solution, how large the corridor should be sized.

## **Where can I get more information on IBM software charges?**

Please refer to: [www.ibm.com/systems/z/swprice/](http://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 T02?

This new configuration will be identical to the IBM z14 ZR1 and will enable floor space reduction for IBM z13s and older customers. For the new z15 T02 package:

- Whether on raised floor or not, the 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 T02 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.
- The z15 T02 is shipped with side covers but if the system is installed in a row with other servers then one or both side covers can be removed. Front and rear covers must be installed.
- There is one power option - intelligent power distribution unit (iPDU).
- The doors are designed for acoustics and optimized for air flow.
- 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.
- The top exit I/O feature code no longer increases system width as in the predecessor Z systems.
- The 19" lift tool was introduced with z14 ZR1 and will be new for z13s and older customers.

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

### What are a few changes I'll see in the new frame configuration on the IBM z15 T01?

This new configuration enables significant floor space reduction for most customers. For the new z15 T01 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 T01 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.
- The z15 T01 is shipped with side covers but if the system is installed in a row with other servers then one or both side covers can be removed. Front and rear covers should always remain installed.
- There are two power options - intelligent power distribution unit (iPDU) and Bulk Power Assembly (BPA). Note that DC utility power is not supported.
- 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 is available for the z15 T01 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 T02?**

The z15 T02 is designed with an environmental focus on improving data center efficiency and uses forced air cooling without an internal water loop. The z15 T02 air-cooled single frame has a radiator heat exchanger built into the system. The processors are cooled with an internal water loop to provide maximum cooling and performance. At the frame level the system is air-cooled. The z15 T02 processors are air-cooled using custom heat sinks providing a better price performance option. The fill and drain tool (FDT) is not required for the z15 T02.

### **What is the cooling on the z15 T01?**

The z15 T01 is designed with an environmental focus on improving data center efficiency. There are two options for cooling on the z15 T01 – 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 T01 system configuration. Customer-water cooled is available for the z15 T01. It is the solution if your data center cannot handle the full heat load of the z15 T01 system going to data center air. This could be due to air cooling limitations at the z15 T01 install location (sometimes referred to as a "hot spot") or an overall limited air cooling capacity in the data center.

### **What is the power option on the IBM z15 T02?**

The z15 T02 uses an intelligent power distribution unit (iPDU). It is 2N, and 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. It is intelligent, fully managed and fully integrated within the system's power distribution, control, and maintenance support structure.

The z15 T02 supports either Single Phase or Three Phase Input power depending on configurations. Single Phase can be used for 1 CPC drawer and up to 4 I/O drawers. Three Phase can be used for all configurations. If you select the feature code to reserve the 2<sup>nd</sup> CPC drawer, you will want to use 3 Phase power.

### **What power options are on the IBM z15 T01?**

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.

### **Tell me about the number of line cords required in the z15 T02?**

The z15 T02 is powered by 2 or 4 line cords 30/32A line cords. One Central Processor Complex (CPC) and one PCIe+ I/O drawer require 2 line cords, all other configurations require 4.

### **What makes the iPDU intelligent?**

An Intelligent Power Distribution Unit (iPDU) is a networked power distribution unit that are used to optimize data center efficiency via the real-time, remote, power monitoring, environmental monitoring, and overall data center infrastructure integration. A standard PDU is simply a grouping of AC outlets protected by circuit breakers that distribute input AC power and its function is limited to distributing power. IBM Z uses an iPDU and focuses its monitoring capability to enable enhanced system diagnostics and serviceability, supported by the ability for individual outlets to be turned on and off by command.

### **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 T01 servers. It is Lithium ion battery technology based, not VRLA (valve-regulated lead-acid) like predecessor Z systems. The IBF is not available on the z15 T02.

## **Will there be an Internal Battery Feature (IBF) on future models of IBM Z?**

On August 4th announced a statement of direction regarding the removal of IBF support from future IBM Z systems. IBM z15 is planned to be the last IBM Z server to offer an integrated battery feature (IBF). For systems employing server time protocol (STP), a new function will sense N-mode power for coordinating timing network PTS/BTS. The new function design covers BPA, PDU, single frame and multi-frame models. There is no dependency on any particular power source or UPS. As customer data centers continue to improve power stability and uninterruptible power supply (UPS) coordination, IBM Z continues to innovate to help customers take advantage of common power efficiency and monitoring across their ecosystems.

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

The Internal Battery Feature (IBF) available with z15 T01 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 customer 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 T01?**

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 T01?**

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 customers 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.

## **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 I/O access in the z15 T02?**

The z15 T02, like the z14 ZR1, has moved all I/O to the rear of the system. Pulse Per Second does still plug into the front of the CPC drawer. For raised floor installations, this system will require new floor tile cutouts. See the Installation Manual for Physical Planning for details (hot links provided in last question of this section).

### **Tell me about top exit support in the z15 T02?**

The z15 T02 provides 2 different Top Exit options for I/O and Power. FC 7898 provides a cable management Top Hat to help manage cable entrance into the system. FC 7928 provides Top Exit support through sliding access panels at the top of the frame. No additional management support is provided. Note: Bottom exit cabling is FC 7899.

### **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 customers 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" z14 ZR1.

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

The best place for information on the z15 T02 is the Installation Manual for Physical Planning (IMPP) which is available here:

<https://www.ibm.com/servers/resourcelink/lib03010.nsf/pages/8562InstallationManualForPhysicalPlanning?OpenDocument>

The best place for information on the z15 T01 is the Installation Manual for Physical Planning (IMPP) which is available here:

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



**Unless specifically stated, assume that the use of the family name IBM z15 applies to both the z15 T01 and the z15 T02.**

### **Tell me about the new I/O Fanout on z15?**

The z15 Central Processor Complex (CPC) drawer has 12 new I/O Fanout card that provides 2 interfaces (ports) to the PCIe+ I/O drawer. Each PCIe+ I/O drawer requires 2 interfaces back to the CPC drawers. The system will come configured for high availability with a minimum of 2 I/O Fanout cards and each I/O drawer will be cabled to different fanout cards. This provides redundant paths to each I/O adapter.

### **Tell me about the new Crypto Express7S on z15?**

Both models of the z15 support two new Crypto adapters. In the past, the Crypto adapter has supported a 1port Hardware Security Module (HSM). To increase the cryptographic capabilities, Crypto Express7S comes in either 1port HSM or 2port HSM configurations. The IBM z15 Model T01 now supports up to 60 crypto HSMs, supporting 85 domains, which provides over 5100 virtual HSMs for ultimate scalability. The IBM z15 T02 now supports up to 40 crypto HSMs, supporting 40 domains which provides over 1600 virtual HSMs. The crypto Express 7S continues to enable workloads by adding functionality to the IBM Common Cryptographic Architecture (CCA) and Secure IBM Enterprise PKCS #11 (EP11).

### **Tell me about FICON Express16S+ for the z15 T02.**

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

With FICON Express16S+ on the z15 T02, you have growth opportunities for your storage area network (SAN). FICON Express16S+ supports a link data rate of 16 gigabits per second (Gbps) and auto-negotiation to 4 or 8 Gbps for synergy with existing switches, directors, and storage devices. With support for native FICON®, High Performance FICON for Z (zHPF), and Fibre Channel Protocol (FCP), the z14 servers enables you to position your SAN for even higher performance – helping you to prepare for an end-to-end 16 Gbps infrastructure to meet the lower latency and increased bandwidth demands of your applications.

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

### **What kind of performance can I expect to get with zHPF using FICON Express16S+ on z15 T02?**

The use of FICON Express16S+ on an IBM z15 T02 with the zHPF protocol and a mix of large sequential read and write data transfer I/O operations, and operating at 16 Gbps achieved a maximum throughput of 3200 MB/sec (reads + writes) compared to a maximum of 2560 MB/sec (reads + writes) achieved with FICON Express16S operating at 16 Gbps. This represents an approximately 25% increase over the prior generation FICON Express16S on z13s.

The use of FICON Express16S+ in an IBM z15 T02 with the zHPF protocol and small data transfer I/O operations, and operating at 16 Gbps achieved a maximum of 314,000 IOs/sec. This represents 3x improvement over the prior generation FICON Express16S on z13s.

Disclaimer: This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

## **What kind of performance can I expect to get on commercial batch elapsed time on the z15 T02 with FICON Express16S+?**

Batch Elapsed time improves 17% when using FICON Express16S+ on z15 T02 running I/O intensive batch workloads compared to running the same workload using FICON Express16S on a z13s.

Batch Elapsed time improves 35% when using FICON Express16S+ on z15 T02 running I/O intensive batch workloads compared to running the same workload using FICON Express8S on a z13s.

Disclaimer: This performance data was measured in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

## **What kind of performance can I expect to get for the FCP protocol using FICON Express16S+ on z15 T02?**

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

The use of FICON Express16S+ in an IBM z15 T02 with the FCP protocol for small data transfer I/O operations, operating at 16 Gbps achieved a maximum of 380,000 IOs/sec, representing a greater than 3x improvement over the maximum of 110,000 IOs/sec achieved with FICON Express16S operating at 16 Gbps.

The use of FICON Express16S+ in an IBM z15 T02 with the FCP protocol and operating at 16 Gbps, FICON Express16S+ achieved a maximum throughput of 3200 MB/sec (reads + writes). This represents approximately a 25% increase in throughput over the prior generation FICON Express16S on z13s.

The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

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

Disclaimer: This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

## Tell me about the new FICON Express16SA for z15 T01.

FICON Express16SA is designed to help absorb large application and transaction spikes like those driven by large unpredictable cloud, analytic, and mobile workloads. FICON Express16SA also supports IBM Fibre Channel Endpoint Security, a new end-to-end solution for the IBM z15 T01 that is designed to provide a means to help ensure the integrity and confidentiality of all data flowing on Fibre Channel links between authorized server and storage devices, creating a trusted storage network that encrypts data in flight. **NOTE: Fibre Channel Endpoint Security and the FICON Express16SA are not supported on the z15 T02.**

The FICON Express16SA on the z15 T01 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.

## What kind of increased performance can I expect to get with zHPF using FICON Express16SA on z15 T01?

The use of FICON Express16SA on an IBM z15 T01 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 T01 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 T01?

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.

## Where can I get more information on IBM Fibre Channel Endpoint Security for the z15 T01?

Please refer to the Fibre Channel Endpoint Security FAQ document located :  
<https://www.ibm.com/downloads/cas/Y6E9KLA8>

## How many channel subsystems (CSS) are supported on the z15 T02?

The z15 T02 has channel subsystem (CSS) support for three logical channel subsystems (LCSSs) with support for forty LPARs, and three subchannel sets.

The z15 T01 has channel subsystem (CSS) scalability with support for six logical channel subsystems (LCSSs) which are required to support the eighty-five LPARs for z15 T01, 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, 25 GbE RoCE Express2, Crypto Express6S, Coupling Express LR, and Integrated Coupling Adapter SR (ICA SR), IBM zHyperLink Express.

### **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 customers 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 T02 support?**

The z15 T02 supports up to 8 adapters are supported (16 ports). Each LPAR should be provisioned with 2 PFIDs for redundancy, ideally with each PFID defined on two separate ports across two separate adapters.

The z15 T01 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 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 customer'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 T01?**

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. The z15 T02 will only support the OSA-Express7S 25 GbE SR.

### **What are the new OSA-Express7S adapters for the z15 T01?**

The z15 T01 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 T01.

## 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 (VFM) 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. VFM on the z15 T01, is orderable in twelve .5TB sizes between .5TB and 6TB. VMF on the T02 is available in four .5TB sizes between .5TB and 2TB.

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. Customers 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

### 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.

### **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.

## z15 Sustainability

### Does IBM have a corporate policy for protecting the environment?

IBM has had long-standing corporate policies aimed at protecting the environment and conserving energy and natural resources. The IBM Product Design for the Environment program was established in 1991 to focus on product environmental design.

### How does the IBM z15 help support data center sustainability initiatives?

The IBM z15 is a transformed system that was designed for overall data center sustainability. IBM Z provides several tools, available with IBM z15, to help support data center sustainability initiatives.

- Single-system performance reduces energy consumption by 50% compared to the x86 systems running a similar workload.\*
- IBM z15 can be integrated into modern data center infrastructure management system through a set of security enabled, REST-based Web Services APIs. The Web Service API metrics groups allows IBM z15 to report key environmental and power consumption data, including ambient temperature and humidity, heat load, exhaust temperature, system power consumption, and power on each line cord phase. Available on Github, the zhmccustomer package is a customer library written in pure Python that interacts with the Web Services API of the HMC on Z machines. This package is designed to make the HMC Web Services API more easily consumable for Python programmers.
- IBM Z Energy Optimization Advisor provides insights into the operation of the IBM Z system and recommends data center actions that can be taken to reduce the overall system power. The advisor provides insight to data center operators, accessed through the user interface on the SE and HMC or through the Web Services API endpoints, about what actions they can take to save power.

## z/OS

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

Please visit: [ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSQ03081USEN](http://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 z/CX Container Extensions. As with many of our planning and implementation offerings, we customize the exact services to meet individual customer needs. Our z/CX offering will help customers 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](mailto: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](http://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](mailto:zvse@de.ibm.com).

## z/VM

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

Please visit: [IBM z/VM – Frequently Asked Questions \(ibm.com/downloads/cas/GDVOX61O\)](https://ibm.com/downloads/cas/GDVOX61O)

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

Please visit: [IBM Wave - Frequently Asked Questions \(ibm.com/downloads/cas/NOLXG3JM\)](https://ibm.com/downloads/cas/NOLXG3JM)

## KVM

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

KVM on IBM Z can and is 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 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](https://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: [Linux on IBM Z - Frequently Asked Question \(ibm.com/downloads/cas/O8W6EJ0P\)](https://ibm.com/downloads/cas/O8W6EJ0P)

### **What are the IBM tested and Partner certified Linux environments on IBM Z?**

Canonical, Red Hat and SUSE support z15 with their Linux distributions. The IBM tested and Partner certified Linux environments are shown at the “Tested platforms” web page, please visit [ibm.com/it-infrastructure/z/os/linux-tested-platforms](https://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, KVM on LinuxONE, z/VM, 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.

### Can I use DPM to set up Secure Execution?

Yes, IBM Dynamic Partition Manager R4.1 will support Secure Execution.



## 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! Here some options for you to start learning today!

- IBM hosts a Master the Mainframe contest that students can enter from September to December every year. 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. When the competition is not running you can enjoy the experience using the Master the Mainframe Learning Management System. Learn more at [ibm.biz/mastertheframe](http://ibm.biz/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.
- For those students looking to obtain professional certification, Coursera offer the [IBM z/OS Mainframe Practitioner Professional Certificate](#).
- 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 customer?

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 "Resources" tab of the IBM [Z Academic Initiative website](#) to register for the z/OS Academic Cloud or 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?

Here are some options:

- 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](#) LINK. NOT WORKING 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 2021  
IBM Corporation  
New Orchard Road  
Armonk, NY 10504  
U.S.A.  
05/21

IBM, IBM logo, IBM Z, Db2, DFSORT, FICON, GDPS, HyperSwap, IBM Sterling Connect:Direct, IMS, 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.

A current list of IBM trademarks is available on the Web at <https://www.ibm.com/legal/us/en/copytrade.shtml>, and select third party trademarks that might be referenced in this document is available at [https://www.ibm.com/legal/us/en/copytrade.shtml#section\\_4](https://www.ibm.com/legal/us/en/copytrade.shtml#section_4).

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

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

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

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

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

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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](#).

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

RStudio®, the RStudio logo and Shiny® are registered trademarks of RStudio, Inc.

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.

Zowe™, the Zowe™ logo and the Open Mainframe Project™ are trademarks of The Linux Foundation.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this documentation is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this documentation, it is provided "as is" without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this documentation or any other documentation. Nothing contained in this documentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in these materials to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

Product release dates and/or capabilities referenced in these materials may change at any time at IBM's sole discretion based on market opportunities or other factors and are not intended to be a commitment to future product or feature availability in any way.

ZSQ03115-USEN-14