

IBM z17 and LinuxONE Emperor 5 Technical Overview

Kenny Stine
IBM Z Technical Specialist
Washington System Center
kjstine@us.ibm.com



The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a more complete list of IBM Trademarks, see <https://www.ibm.com/legal/copytrade> :

*BladeCenter®, CICS®, DataPower®, Db2®, e business (logo)®, ESCON, eServer, FICON®, IBM®, IBM (logo)®, IMS, MVS, OS/390®, POWER6®, POWER6+, POWER7®, Power Architecture®, PowerVM®, PureFlex, PureSystems, S/390®, ServerProven®, Sysplex Timer®, System p®, System p5, System x®, z Systems®, System z9®, System z10®, WebSphere®, X-Architecture®, z13™, z13s™, **IBM z14™**, **IBM z15™**, **IBM z16™**, **IBM z17™**, z Systems™, z9®, z10, z/Architecture®, z/OS®, z/VM®, z/VSE®, zEnterprise®, zSeries®, IBM Z®, IBM LinuxONE III™, IBM LinuxONE Emperor™, IBM LinuxONE Rockhopper™

The following are trademarks or registered trademarks of other companies.

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.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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.

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

*** All other products may be trademarks or registered trademarks of their respective companies.**

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput 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. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured Sync new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.




This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained Sync the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

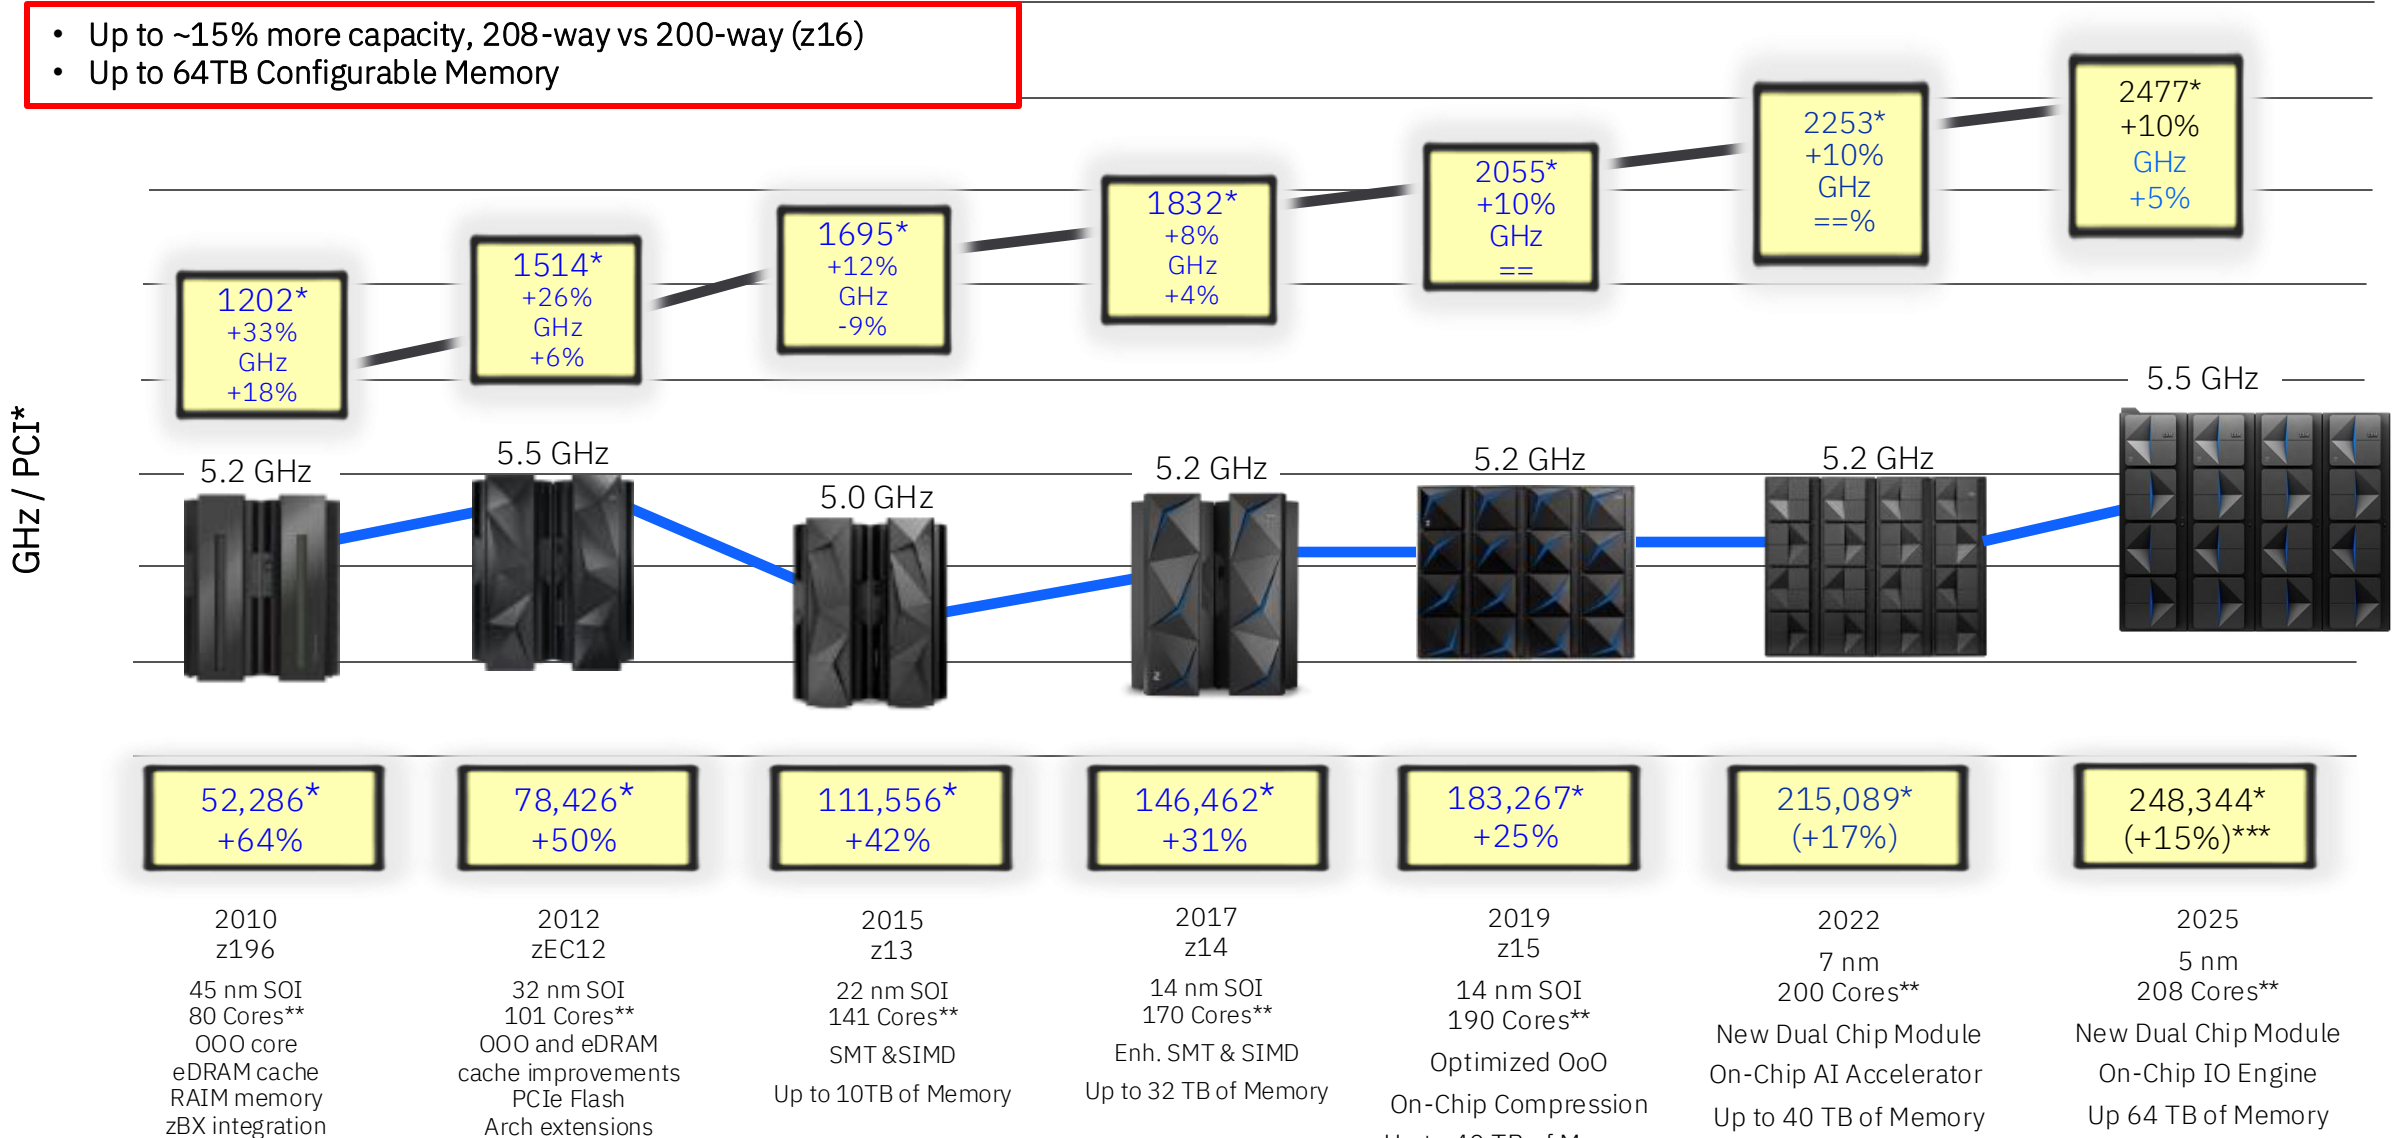
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

IBM z17 Multi-frame Technical Overview



z17 continues the CMOS

- Up to ~15% more capacity, 208-way vs 200-way (z16)
- Up to 64TB Configurable Memory



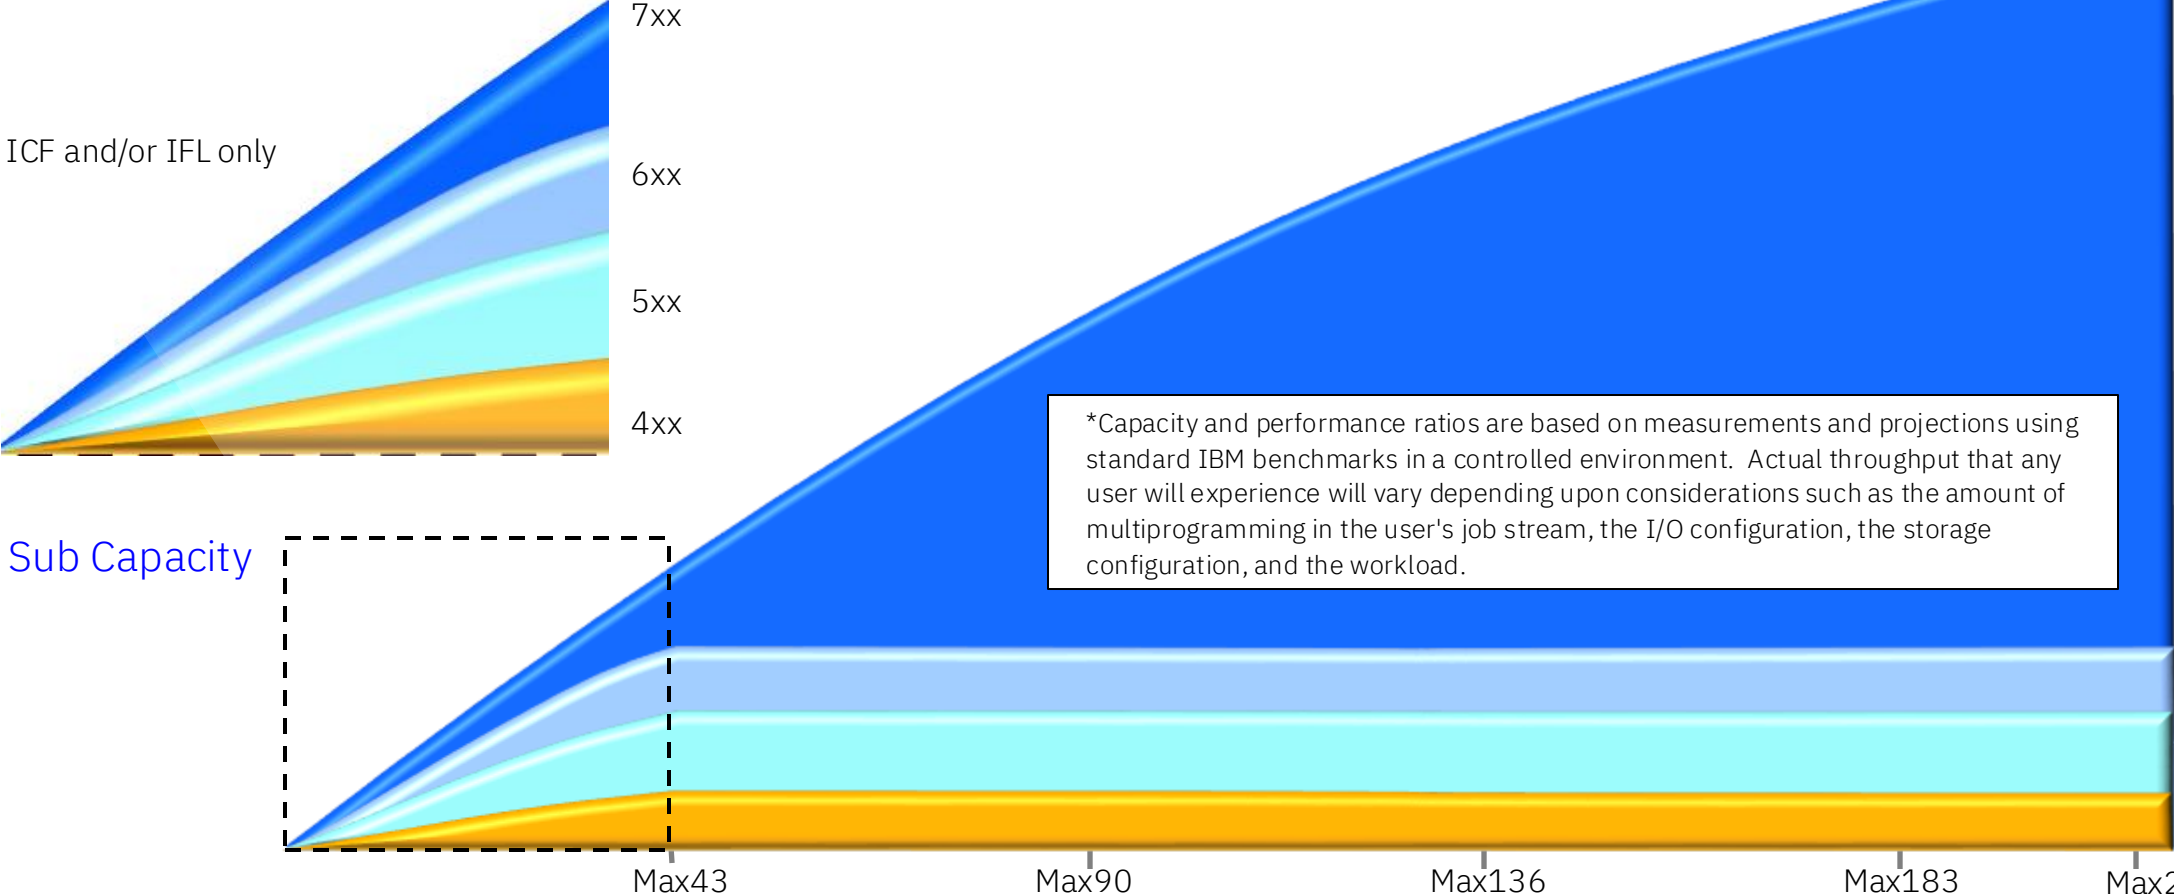
* MIPS Tables are NOT adequate for making comparisons of IBM Z processors. Additional capacity planning required
 ** Number of PU cores for customer use
 *** MIPS are currently estimates for z17. Actuals for z17 MIPS will be released at system announce

ME1 Full and Sub-Capacity CP Offerings

CP Capacity – Relative to Full Capacity Uni
 701 = 100% ≈ 2,477 PCI (IBM MIPS)
 601 ≈ 66% ≈ 1,642 PCI
 501 ≈ 41% ≈ 1,030 PCI
 401 ≈ 12% ≈ 296 PCI

- Subcapacity CPs, up to 43 may be ordered.
 If more CPs are ordered all must be full 7xx capacity.
- All CPs on an ME1 CPC must be the same capacity (except during Recovery Boost).
- All specialty engines are full capacity.

701-7K8
 601-643
 501-543
 401-443
 400 = 100% ICF and/or IFL only



IBM z17 Processor Design and Structure

CPC Drawer

CP Layout

- 8 CPs in 4 Dual Chip Modules (DCMs)
- Built on 5 nm technology
- Up to 2 PCIe Gen5 interfaces per CP
- Up to 1 memory controller per CP
- 1 Integrated Accelerator for AI per CP
- 1 zEDC Accelerator per CP
- 1 DPU per CP

Fanouts

- 12 PCIe slots for both I/O fanouts and ICA 2.0 SR Coupling links

Memory

- 8 4U DDR4/DDR5 DDIMMs per Memory Controller
- Max 48 DIMMs per CPC Drawer

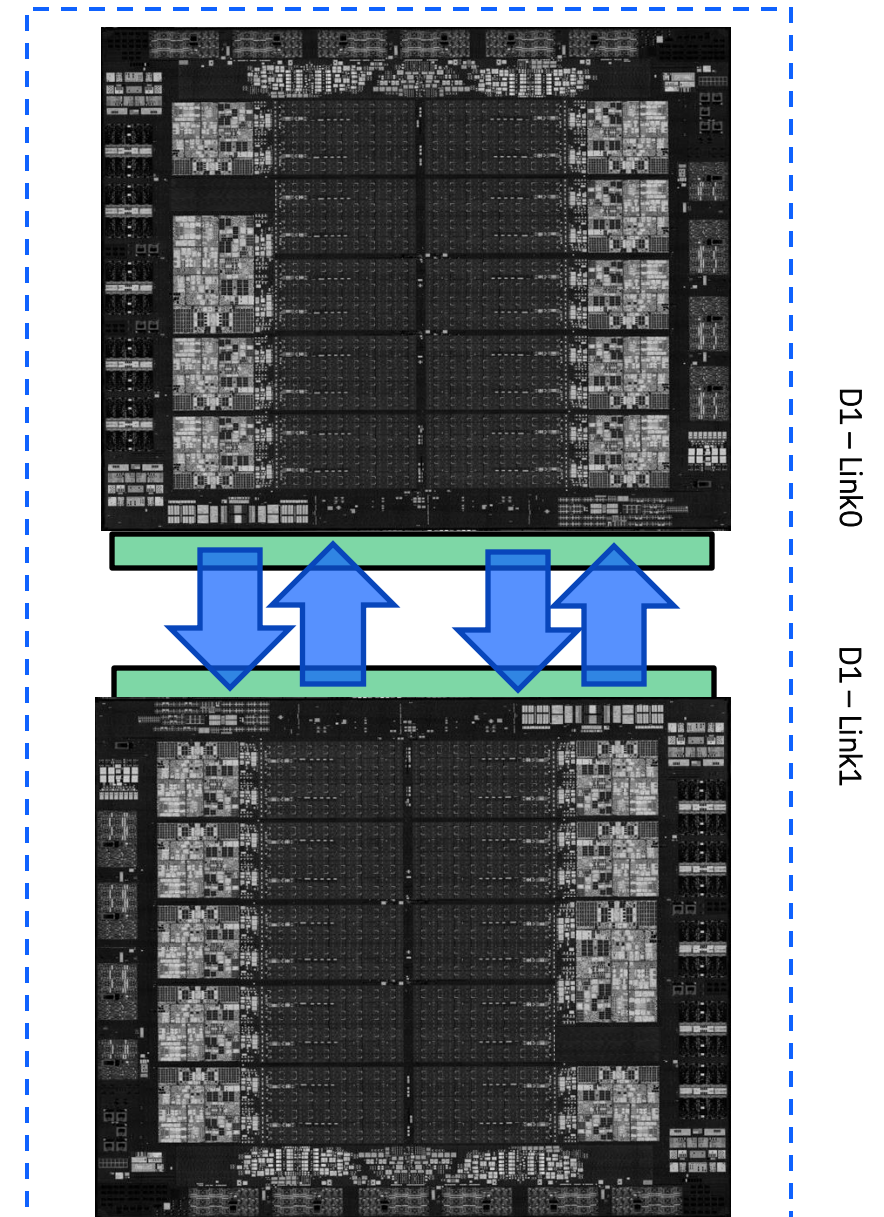
Rear



Front

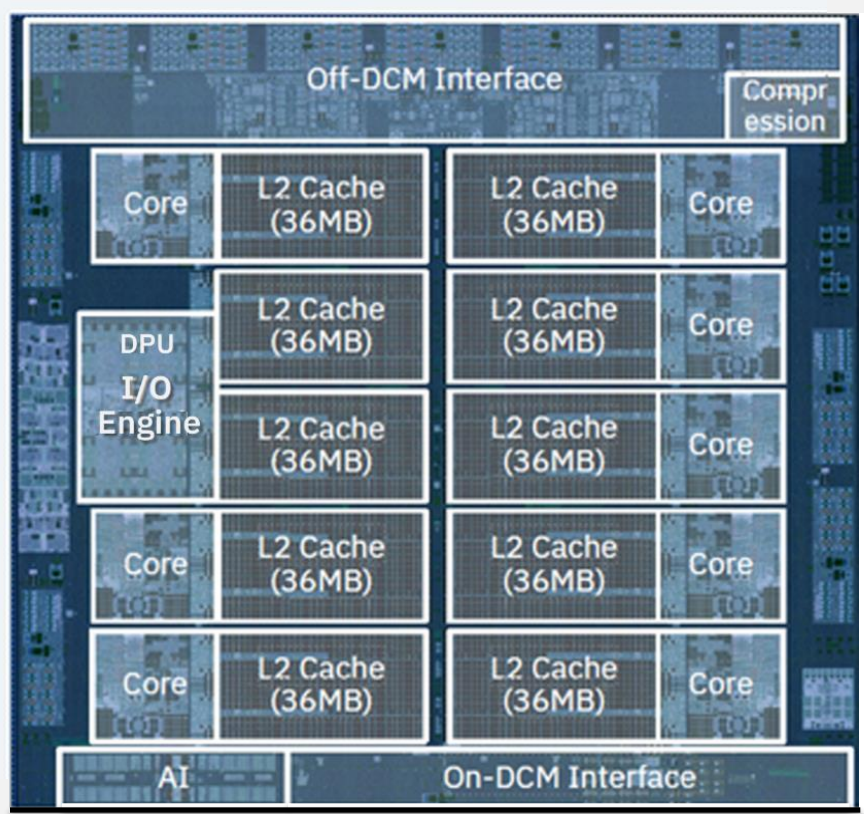
IBM z17 Dual Chip Module (DCM)

- 2 CP chips per DCM
- CP chips connected by highspeed MBus
 - 2 sets of 700 signals between CPs
 - **~119 GB/s** bandwidth in each set
- Each DCM has:
 - up to **15 active cores**
 - up to **16 DDIMS of memory**
 - 3 PCIe Adapter connections
 - Up to 1 SMP Connection (to other CPC Drawers)
 - **~66 GB/s** bandwidth between drawers
 - 4 XBus links
 - **~66 GB/s** bandwidth per link
 - **~264 GB/s** bandwidth between DCMs



*Speed and performance data based on running at max CP frequency in laboratory test

8-Core Telum II Processor Chip



5nm 5HPP Technology

- 8 cores per CP
- 1 IO Engine
- 24.1 Miles of wire per chip (5.3 more miles than z16)
- ~23 mm x 22 mm
- 31.0B transistors (7.5B more than IBM z16)

- 8 processor chips per CPC Drawer in 4 DCMs
- Up to 8 active cores per chip
- One DPU per chip
- One Integrated Accelerator for AI on every Telum II chip
- **On-Core L1 Cache**
 - Private 128K L1I and 128K L1D
- **On-Core/Chip L2 Cache**
 - Each core has access to a private 36 MB cache
 - Up to 18 MB of each cache can be used by other cores as virtual cache depending on the current activity
 - L2 cache of an inactive core becomes shared virtual L3 cache by the active cores of the chip
 - L2 cache of an inactive core of another CP can become virtual L4 cache
- **I/O buses**
 - Each CP chip will support up to 2 Gen-5 PCIe buses

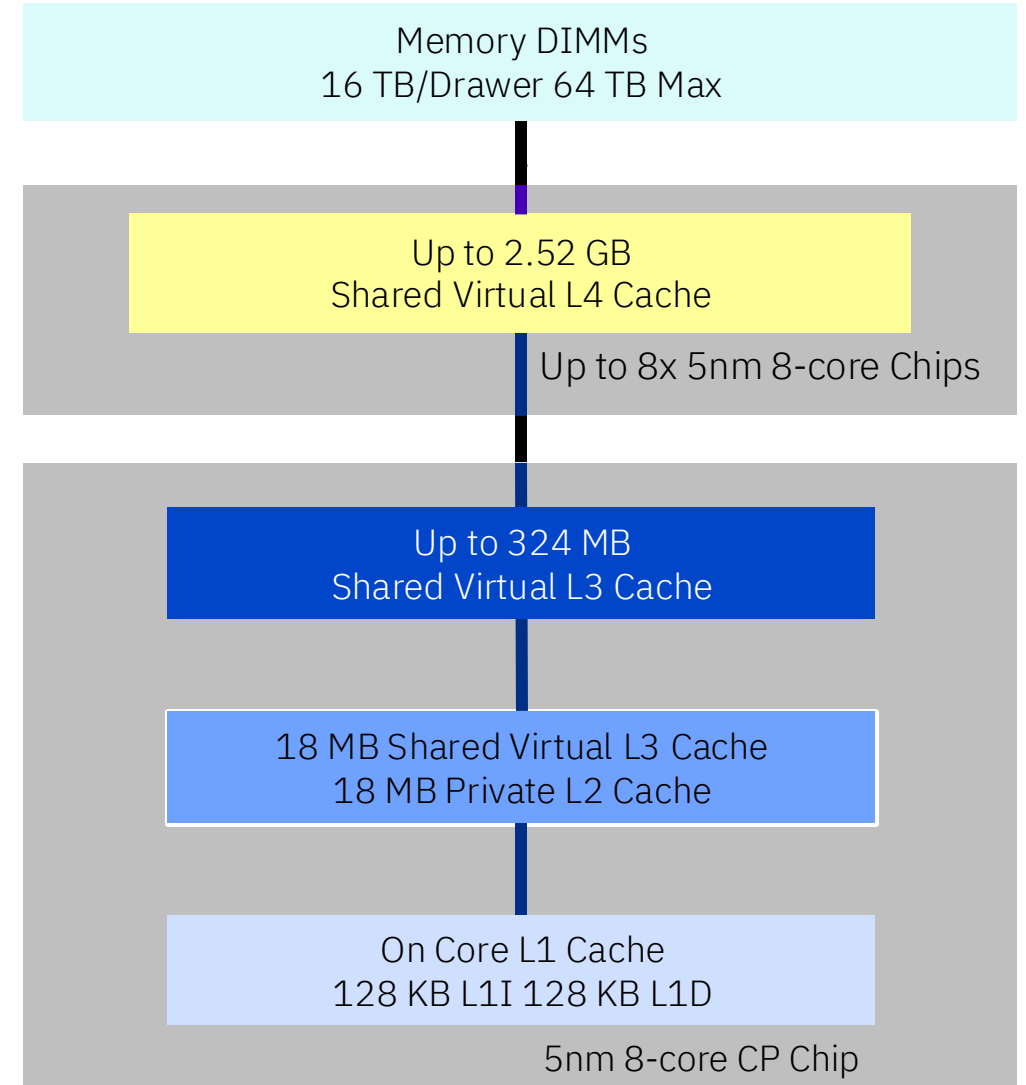
IBM z17 Multi-frame cache

New Cache Design on the CP chip

- On Core L1 Cache
 - Private 128KB L1I and 128KB L1D
- Each core has access to 36 MB Private L2 Cache
- Implementing virtual L3 and L4 Cache

How does it work

- There is no L3 physical cache present on the cores
- There is no physical L4 Cache
 - All CPs L2 are interconnected via buses
- Virtual cache on the same CP will be seen as additional virtual L3 cache to the core
- Virtual Cache on a different CP on the same drawer will be seen as L4 Cache



Dedicated low-latency Integrated Accelerator for AI


Centralized on-chip accelerator shared by all cores



With IBM z17, process up to 5 million inference operations per second with less than 1ms response time using a Credit Card Fraud Detection model



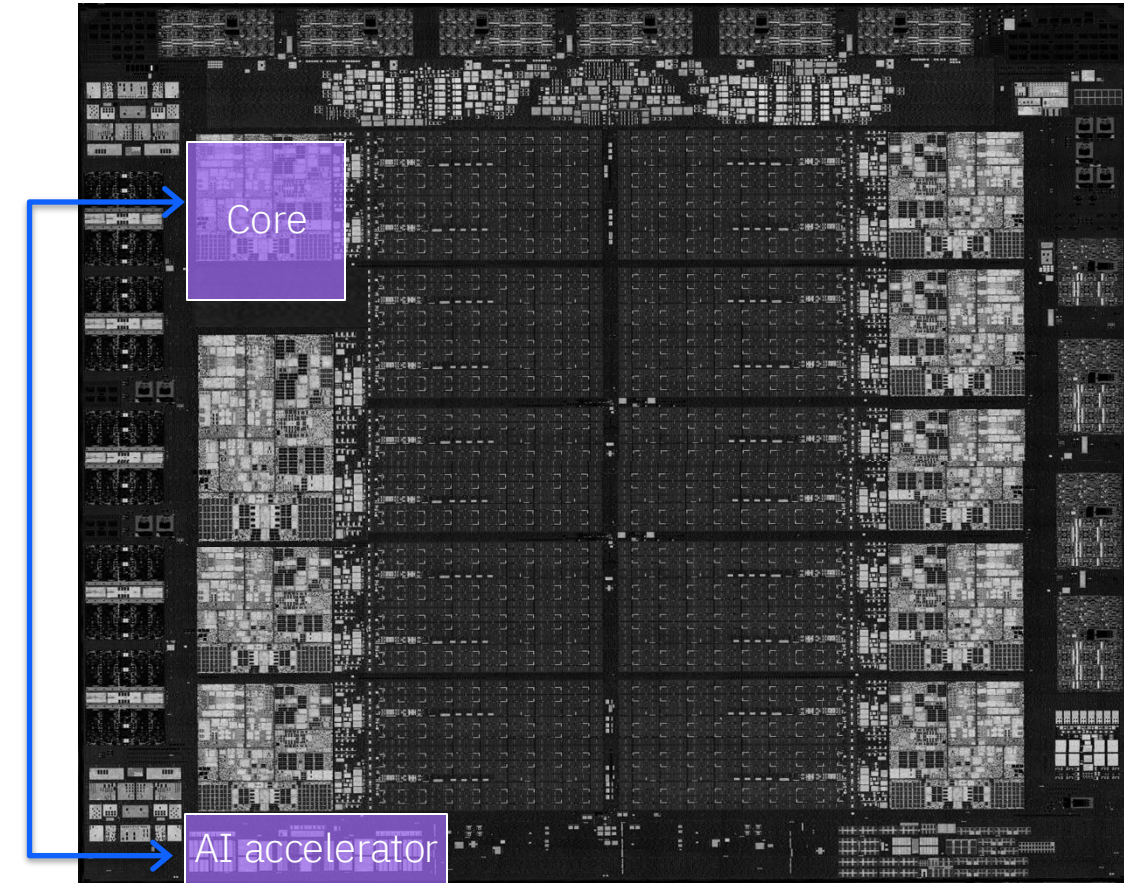
Very low and consistent inference latency and scalable capacity



New Function with the z17

- Transform function between different data layouts/types
- Gaussian Error Linear Unit function
- NNPA instruction can be dispatched to any available accelerator in a drawer

Every CP chip has one Integrated Accelerator for AI built-in



9175 ME1 Processing Units

Offering	Feature Code Description	Chips/CPs	Dual Chip Modules	IFLs/uIFLs	zIIPs/uzIIPs	ICFs/uICFs	Std SAPs	Std. Spares	IFP
ME1	Max43	8 0-43	4	0-43 0-42	0-42 0-41	0-43 0-42	5	2	2
	Max90	16 0-90	8	0-90 0-89	0-89 0-88	0-90 0-89	10	2	2
	Max136	24 0-136	12	0-136 0-135	0-135 0-134	0-136 0-135	16	2	2
	Max183	32 0-183	16	0-183 0-182	0-182 0-181	0-183 0-182	21	2	2
	Max208	32 0-208	16	0-208 0-207	0-207 0-206	0-208 0-207	24	2	2

1. At least one CP, IFL, or ICF must be purchased in every machine.
2. The IFP is conceptually an additional, special purpose SAP – used by PCIe I/O features and some other functions.
3. Additional SAPs have been dropped



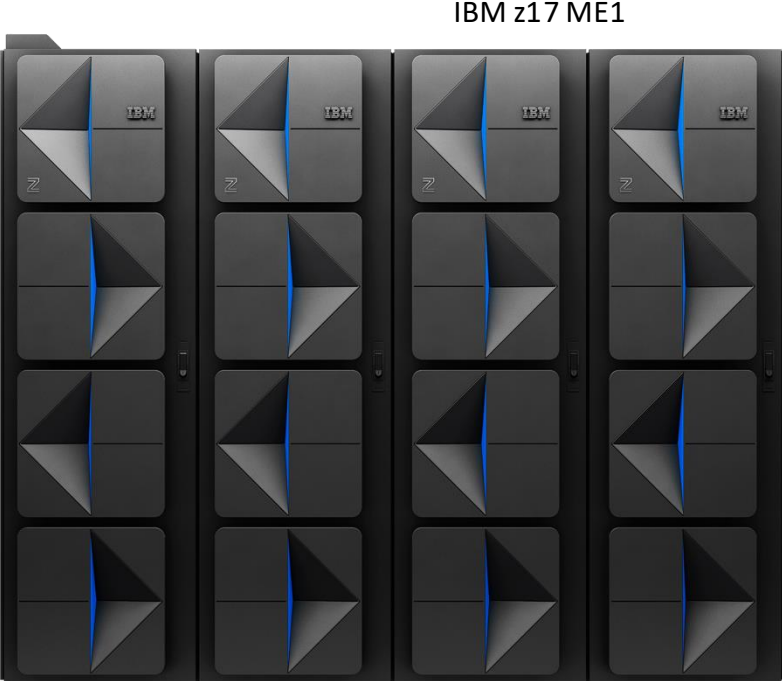

9175 ML1 Processing Units

Offering	Feature Code Description	Chips/CPs	Dual Chip Modules	IFLs/uIFLs	Std SAPs	Std. Spares	IFP
ML1	Max43	8 0-1	4	1-43 0-42	5	2	2
	Max90	16 0-1	8	1-90 0-89	10	2	2
	Max136	24 0-1	12	1-136 0-135	16	2	2
	Max183	32 0-1	16	1-183 0-182	21	2	2
	Max208	32 0-1	16	1-208 0-207	24	2	2

1. At least one IFL must be purchased in every machine.
2. The IFP is conceptually an additional, special purpose SAP – used by PCIe I/O features and some other functions.
3. Additional SAPs have been dropped



IBM z17 System Upgrades*



- IBM z17 to IBM z17 model upgrades
 - IBM z17 ME1 Max43 to Max90, Max136
 - No field upgrade to Max183 or Max208 (these features are Factory built and shipped only).
 - ***Additional I/O Drawers can be added based on available space in current frames and/or I/O expansion frames***
- Any IBM z15 T01 (8561, all models) to any IBM z17 (frame roll)
- Any IBM z16 A01 (3931, all models) to any IBM z17 (frame roll)

IBM z17 Memory

IBM z17 memory considerations

Offering	Feature	Min	Max
ME1/ML1	Max43	512 GB	16 TB
ME1/ML1	Max90	512 GB	32 TB
ME1/ML1	Max136	512 GB	48 TB
ME1/ML1	Max183	512 GB	64 TB
ME1/ML1	Max208	512 GB	64 TB

- Concurrent memory upgrades via licensed internal code (LICC) are available at several capacity levels.
- DDR4 Memory DIMMS (32, 64, 128 GB)
- DDR5 Memory DIMMS (32, 64, 128, 256, 512 GB)

- An additional **884 GB** of memory is reserved above the customer purchase amount for the Hardware System Area (HSA)
- DIMMs include **RAIM overhead**
- RAIM design implemented to improve RAS


IBM z17 New Features and Functions

IBM Spyre Accelerator PCIe attached card

- 75W PCIe gen5 x16 adapter
- 128GB of LPDDR5 memory
- 300+ TOPS
- Up to 8 cards per I/O drawer
- Generative AI: 8 cards form a logical cluster
- 1TB of aggregate memory
- 1.6TB per second aggregate memory bandwidth



Designed to handle **Large Language Models**
and **Generative AI** use cases



AI/Spyre Accelerators




Clients need the following infrastructure in place to manage Spyre adapters.

- **3 SSC LPARs required**
 - **One Appliance Control Center**
 - 2 logical IFL/CP Shared
 - 16 GB Memory
 - 50 GB DASD (ECKD™/SCSI)
 - Networking connections
 - **Two Spyre Support Appliances**
 - 2 logical IFL/CP Shared each
 - 50 GB Memory each
 - 50 Gb DASD (ECKD/SCSI) each
 - Networking Connections for each


- Each system that has spyre adapters will need **two Spyre Support Appliances** to manage the Spyre adapters
 - Recommended to split adapters evenly between the two SSCs
- Clients only need **one Appliance Control Center** in an environment as long as the SSC has network access to where they want to deploy Spyre Support Appliances
 - Acts as “App Store” to install the Spyre Support Appliance code
 - IBM requires at least 2 physical IFLs

IBM Z Flexible Capacity for Cyber Resiliency




Greater Flexibility

Dynamically shift workloads between IBM z16 and IBM z17



Complete Client Control

Flexibility to activate and deactivate capacity on participating machines




Improved Compliance for Disaster Recovery

Fully automatable using solutions like GDPS



Disaster Recovery & DR Testing

Activate required capacity at your DR site to continue running your business. Automate and test recovery procedures for unplanned outages, including cyber attacks, to provide near continuous availability and DR




Frictionless Compliance

Meet the ever-evolving requirements of global regulators, allowing a highly automated and fast process to demonstrate a production site swap



Facility Maintenance

Run your production workload from your alternate site while you perform maintenance at your primary site with the capacity, you need



Proactive Avoidance

Protect your critical business services from natural disasters. Avoid rolling power outages. Migrate your critical workloads to an alternate Site before your business gets impacted and stay there for up to one year

Connectivity Mapping Tool Overview

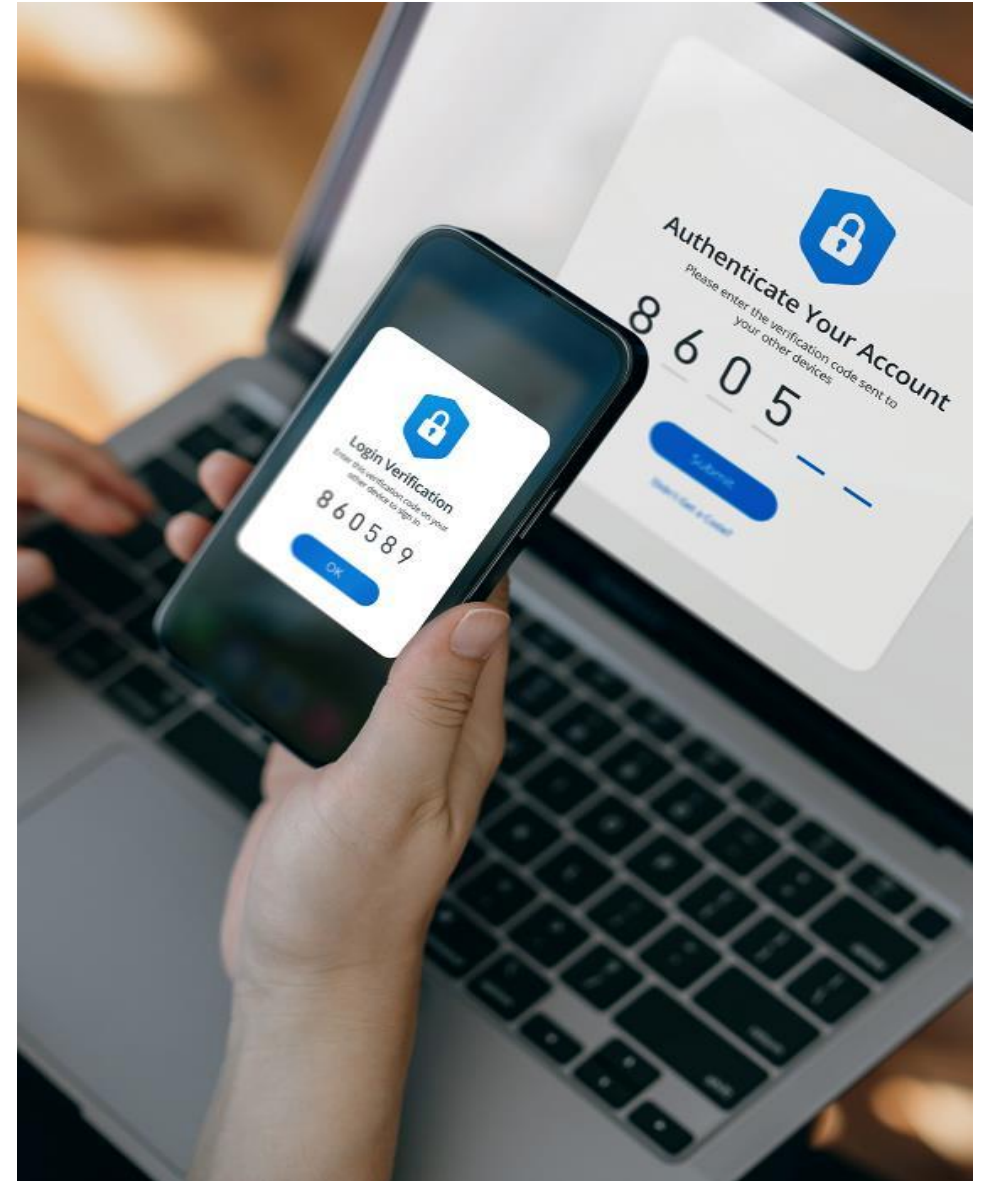
- The Connectivity Mapping Tool (CMT) maps the I/O ports of your hardware to the CHPID/FID definitions in your IOCP source statements.
- The tool guides you in making mappings that will maximize path availability by avoiding single points-of-failure.
- The Connectivity Mapping Tool can save time by minimizing or eliminating changes when adding new hardware or upgrading existing hardware.

© Copyright IBM Corp. 2024, 2025



HMC Dual Control

- Dual Control is in response to feedback that a single HMC user is no longer adequate to support client security policy.
 - Protect against malicious insider attacks
 - Protect against user actions mistakes
- The HMC solution will address providing a second user authorization of a requested action of a different user and reviewing the selections actions for correctness.
- This optional Dual Control will be provided for disruptive tasks and Crypto related settings
 - User would make all selections for a given Dual Control task
 - Real time notification for Authorization User Review/Confirm request



Supported I/O features



New Build I/O Features

Description		Feature Code	Ports	Max Features	Comments - CHPID Types (FID Types)
ICA SR2.0	NEW	0216	2	48	CS5
Coupling Express3 LR 10GB	NEW	0498	2	64	CL5
Coupling Express3 LR 25GB	NEW	0499	2	64	CL6
Network Express SR 10G	NEW	0524	2	48	OSH, (NETH)
Network Express LR 10G	NEW	0525	2	48	OSH, (NETH)
Network Express SR 25G	NEW	0526	2	48	OSH, (NETH)
Network Express LR 25G	NEW	0527	2	48	OSH, (NETH)
OSA-Express7S 1.2 GbE SX		0455	2	48	OSD, OSC
OSA-Express7S 1.2 GbE LX		0454	2	48	OSD, OSC
OSA-Express7S 1.2 10GbE SR		0457	1	48	OSD
OSA-Express7S 1.2 10GbE LR		0456	1	48	OSD
OSA-Express7S 1.2 25GbE SR		0459	1	48	OSD
OSA-Express7S 1.2 25GbE LR		0460	1	48	OSD

New Build I/O Features

Description	Feature Code	Ports	Max Features	Comments- CHPID Types
zHyperLink Express2.0 NEW	0351	2	16	HYL
Crypto Express8S (1 HSM)	0909	N/A	16	
Crypto Express8S (2 HSM)	0908	N/A	30	
IBM Adapter for NVMe 1.1	0448	N/A	16	(NVMe) customer supplied NVMe
IDAA Internal Storage - 15TB NEW	0528	N/A	16	(NVMe) IBM supplied NVMe LinuxONE Only
Reserved Spyre (AI Adapter) NEW	0061	N/A	48	Sets of 8
FICON Express32-4P LX NEW	0387	4	96	FC, FCP
FICON Express32-4P SX NEW	0388	4	96	FC, FCP



I/O Port Limitations

FICON

- Combined total of all generations of FICON Port – 384

Networking

- Combined total of OSA and Networking Express Adapters – 48

Crypto

- Combined total of all generations of Crypto HSMs – 60

NVMe Adapter

- Combined total of NVMe adapters (including the IDAA Internal Storage - 15TB) – 16


Spyre

- Total number of reservable Spyre adapters in sets of 8 – 48

DPU

- Combined total of DPU managed I/O Ports (FICON Express32-4P, and Networking Express) – 384

FICON



Qualified 3rd-party FICON switches

Partner	FICON Switches Supported		Firmware Supported
	Broadcom Name	IBM Name	
Broadcom/ Brocade	X7-8	SAN512B-7	FOS 9.1.1d
	X7-4	SAN256B-7	
	G720	SAN64B-7	FOS 9.2.0c1
	7850	SAN42B-R7	
	Cisco Name	IBM Name	
Cisco	MDS 9710	SAN384C-6	NX-OS 9.4(1a)
	MDS 9706	SAN192C-6	
	MDS 9220i	SAN16C-R	

Note: for up-to-date hardware support and service dates, please visit the End-of-Life pages for [Cisco](#) and [Broadcom](#).



End-to-end solution for data-in-flight protection

IBM Fibre Channel Endpoint Security enables FICON or Fibre Channel Protocol (FCP) Links

Challenges

- Encrypt all data in-flight by corporate directive
- Protect the integrity and confidentiality of data in-flight

Customer Value

- Gain confidence that all data flowing within and across data centers is traveling between trusted entities
- **Ensure ability to provide auditable information** verifying that customer data is only accessed by trusted IBM Z and storage devices
- Use on all IBM Z operating systems




New Statement of Direction

- Given the increasing importance of providing the highest level of data protection to IBM Z clients, IBM intends to require the use of IBM Fibre Channel Endpoint Security for all FICON connected devices starting with the release of IBM z17+1. This direction will require investment by IBM Infrastructure teams, FICON storage vendors and IBM Z clients as an important step towards continuing to secure the most mission critical workloads. In support of this direction, all new FICON-connected storage systems introduced after December 31, 2024, will be required to support IFCES to connect to z17+1.

* Not Official Language

IBM z17

Networking Options



Networking updates and strategy

- The introduction of the Network Express adapter on IBM z17 will enable converging of legacy OSA and ROCE into one hardware offering
- The OSH CHPID type will support all legacy functions available with OSD, utilizing enhanced QDIO (EQDIO) architecture (CHPID type OSH) for OSA-style I/O.
 - OSH CHPID not supported in z/OS 2.4 and earlier releases nor z/VSE
- OSA Express7S 1.2 adapter is available on the IBM z17
 - Customers who require the legacy QDIO architecture (CHPID type OSD) must use OSA-Express7S 1.2 adapter.
 - CHPID type OSC requires a OSA Express7S 1.2 GbE SX/LX
- CHPID type OSE will not be supported on IBM z17 generation
- 1000Base-T copper SFP is no longer supported on IBM z17 generation.



Networking restrictions

- SNMP support for OSA (OSH) will not be available at z17 GA (targeted for Sept/2025)
- OSAENTA (Network Traffic Analyzer) trace support will not be available (Alternatively, a sniffer (wireshark) trace may be used)
- QDIOSYNC support will not be available (Automated mechanism to collect SW and HW traces concurrently)
- NETH PFIDs (RoCE) – Guideline
 - For clients using IBM Network Express 10G or 25G with z/OS CommServer 2.5 or 3.1, it is recommended to set up a maximum of 16 NETH FIDs on the same adapter when OSH is also enabled.
- **zOS guest deployed in zVM vSwitch environment must be Layer 3 QDIO (IPAQENET (OSD CHPID) OSA-Express)**
 - **A z/VM vSwitch attached to a Network Express card must operate in Layer-2 mode. This means that a z/OS guest cannot be deployed on this vSwitch, as the z/OS guest must use a QDIO interface (a current z/VM limitation) and z/OS does not support Layer-2 with QDIO. Linux supports both modes today, so Linux guests using QDIO can use this configuration.**



Networking Summary

How to determine how many Networking Express ports to order?

- Order at least the number of ports equal to the greater of:
 - Current number of OSA ports
 - Current number of RoCE Ports

z/OS 2.5 and Higher
z/VM 7.3 and Higher

For general Networking:

- Use Networking Express Adapters for 10/25 GbE communication
 - CHPID: OSH

For OSA-ICC communication:

- Use OSA Express 7S 1.2 GbE
 - CHPID: OSC

For SMC-R communication:

- Use Networking Express Adapters for 10/25 GbE communication
 - FID: NETH
 - z/OS only
 - Recommended to limit SMC-R to 16 VFs per port
 - Must use the same PCHID as the paired OSH

z/OS 2.4 or lower
zVSEⁿ

For general Networking:

- Use OSA Express 7S 1.2 10/25/GbE
 - CHPID: OSD

For OSA-ICC communication:

- Use OSA Express 7S 1.2 GbE
 - CHPID: OSC

For SMC-R communication:


- Use Networking Express Adapters for 10/25 GbE communication
 - FID: NETH
 - Recommended to limit SMC-R to 16 VFs per port

Linux

For general Networking:

- Use Networking Express Adapters for 10/25 GbE communication
 - FID: NETH
 - Limited to 127 VFs per port

IBM z17 Artificial Intelligence Enhancement Spyre Accelerator




z17 PCIe attached AI accelerator placements

Frame C
3 IO drawer

Frame B
3 IO drawer

Frame Z
5 IO drawer



- Minimum 8 adapters
- Adapters are spread out across the fewest IO drawers due to power and cooling impacts
- Max of qty (8) AI accelerators in a single IO drawer
- Max of 48 total AI accelerators
- Qty (48) accelerators could be split in as little as 6 IO drawers


IBM z17 NVMe



IDAA Internal Storage - 15TB – FC0528

NVMe - IBM z16

- IBM's only offering is a carrier card
- Clients are responsible for purchasing their own NVMe SSDs
- Clients are responsible for diagnosing and troubleshooting NVMe SSD problems
- Clients must purchase a separate contract for IBM SSRs to install and replace the SSDs

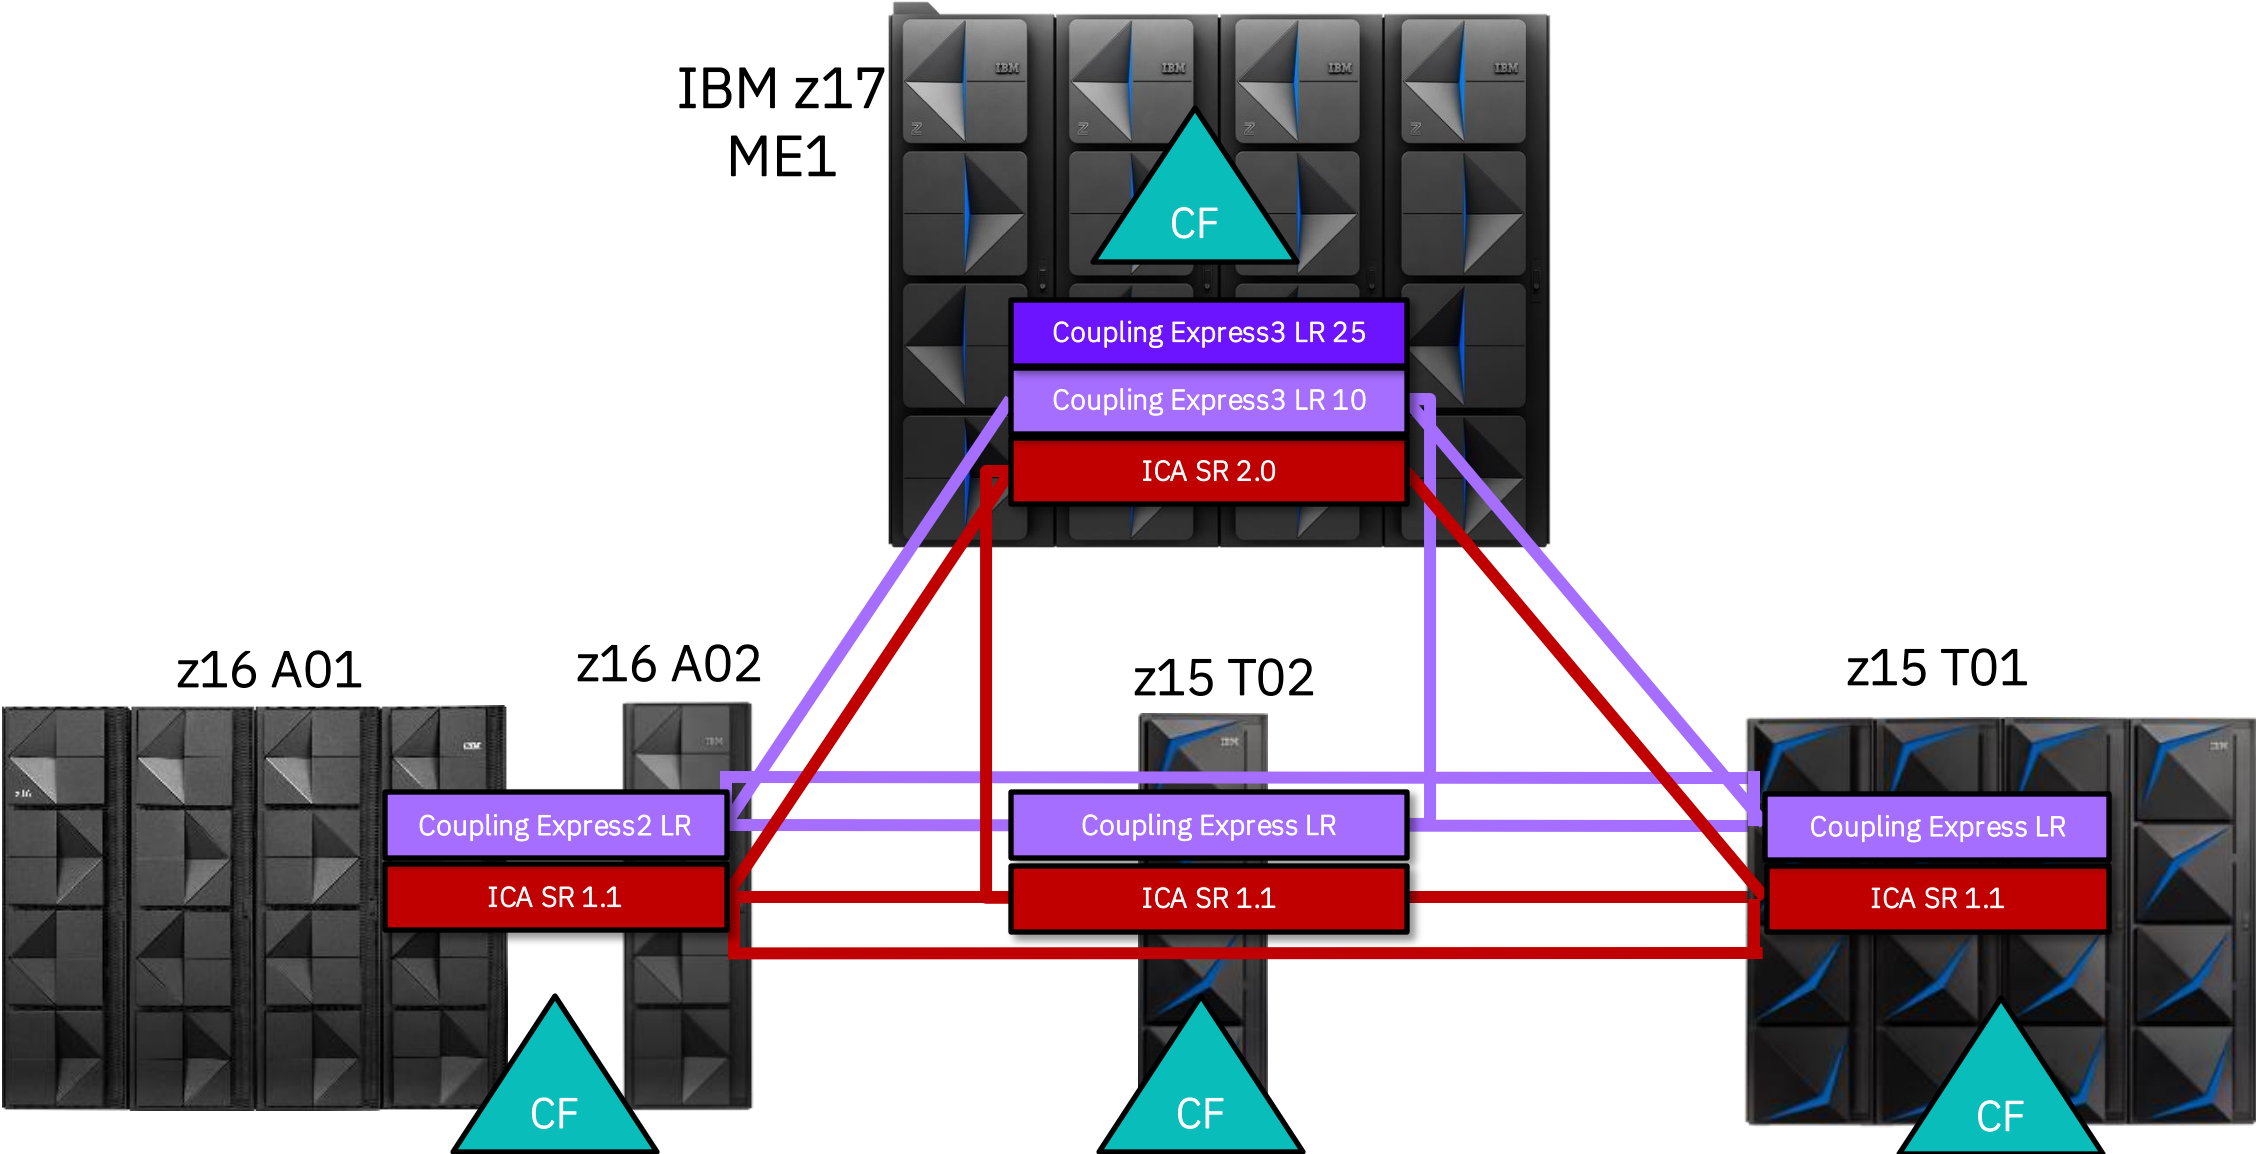


New with IBM z17: Robust NVMe

- New Feature Code (0528) for the Carrier and SSD drive
- A FRU will be created for the carrier+SSD
- Client must be running IDAA
- For LinuxONE only



IBM 9175 coupling connectivity

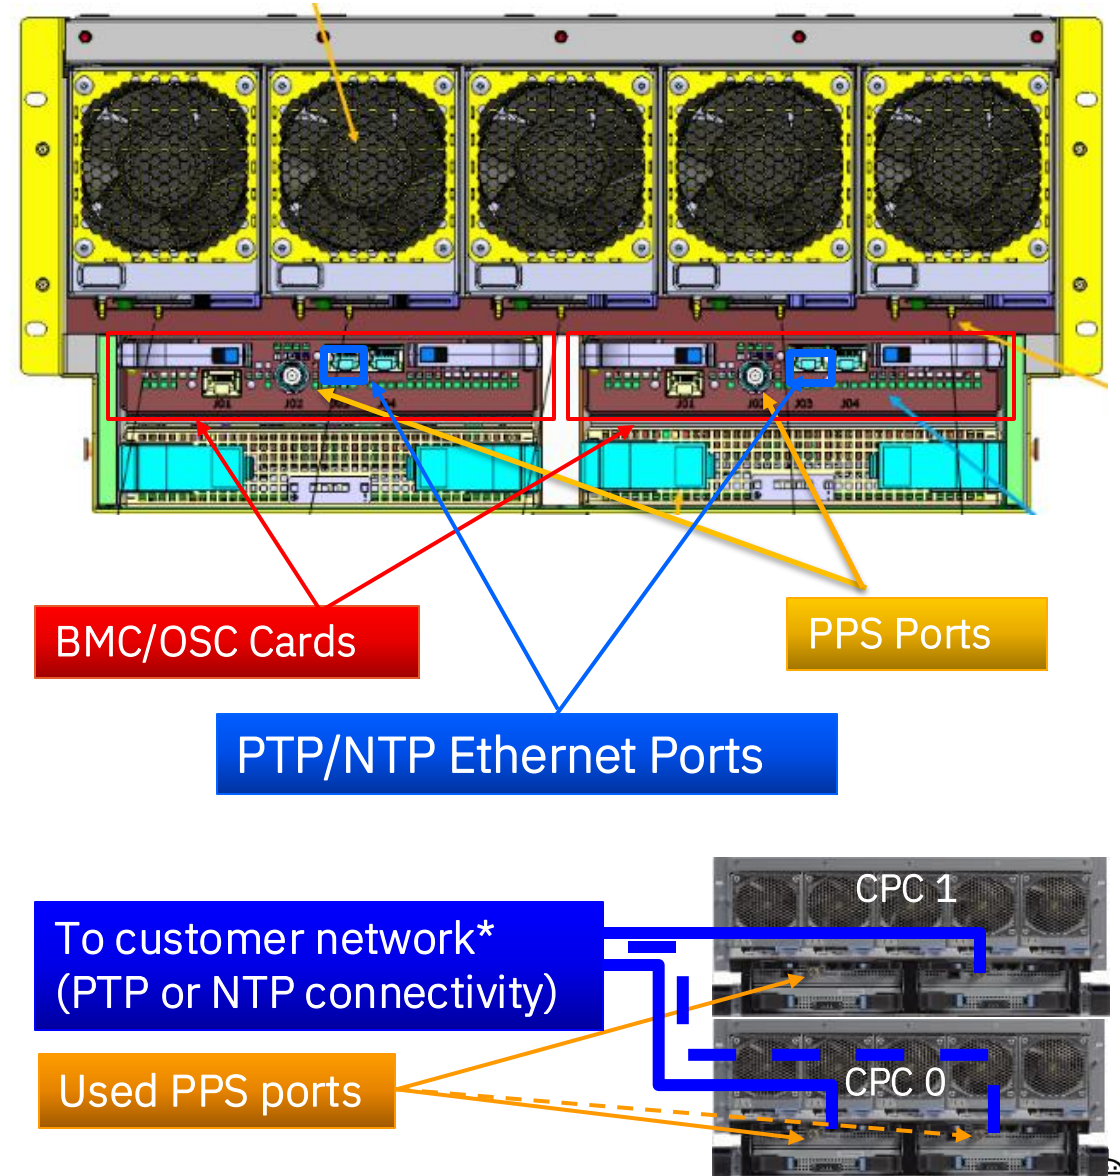


Server Time Protocol (STP)
Precision Time Protocol
(PTP)
Network Time Protocol
(NTP)



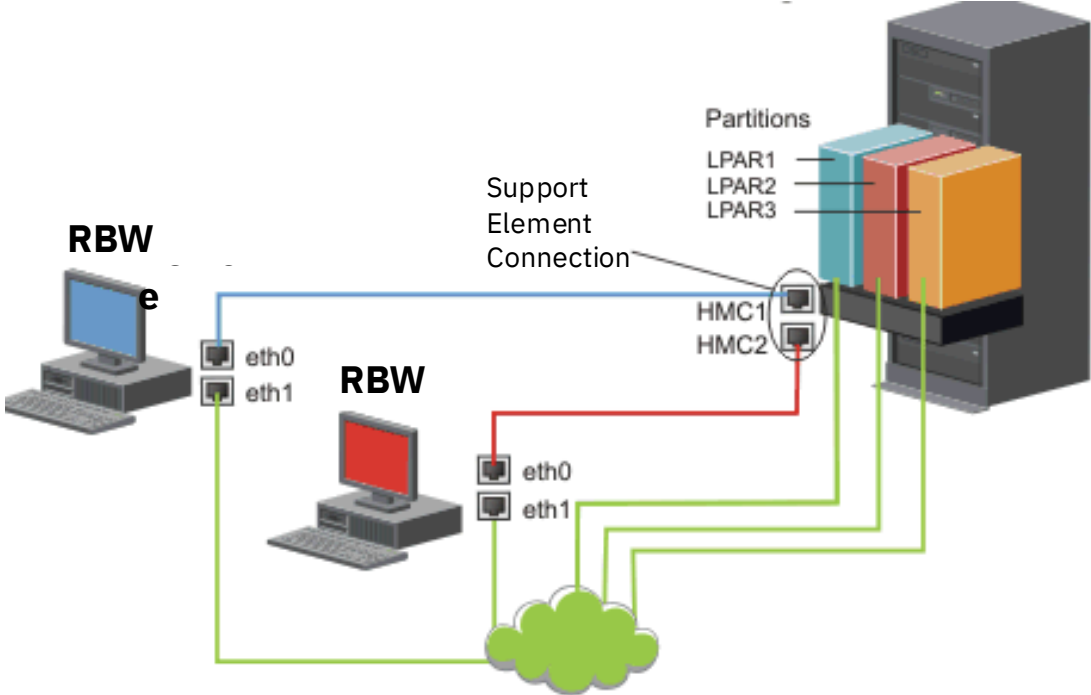
Oscillator cards

- IBM z17 configurations can have one to four CPC drawers
- Each CPC drawer has two combined BMC/OSC cards, each with one PPS port and one ETS port (RJ45 Ethernet, for both PTP and NTP)
- For timing signal redundancy, two links must be used
 - For a single CPC drawer system, both ports must be connected and configured for timing and/or PPS.
 - For a system with two CPC drawers, only the first ports in the first and second CPC drawer can be used, and only two ports (one in CPC 0 and one in CPC 1)



IBM z17 Multi-frame

HMC/SE



HMC/SE Changes

- **Single Sign On**
 - Added support for SSO
 - The logon will be transparent to the HMC other than the validation confirmation
- **Network Time Security**
 - Support NTS on NTP
- **HMC/OSA-ICC 2048+ Key Length Certificates**
- **TLS Cipher Suite Config UI Enhancements**
- **HMC Mobile Enhancements**
- **HMA to SE MES**
 - Standalone HMCs are no longer supported.
 - All client users will remote browse into the HMC
- **Import/Export from Remote File System**
 - Can use Export/Import of SA HMC data and HMC Data Replication to HMA HMCs
 - Adding import/export to remote browsing workstation HDD (in addition to USB & 3 FTP options)
- **BCPii Enhancements**
 - BCPii HWIREST API/v2 support for HMC Targets
 - Asynchronous Notification support for z/OS

Power and Cooling

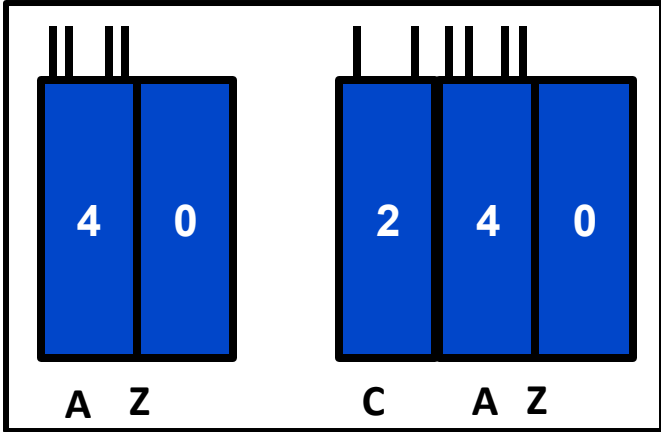
Summary of line cords (rear view)

iPDU Based Power

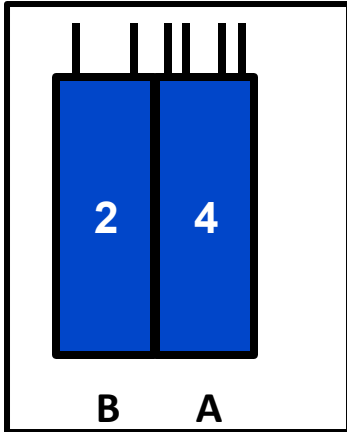


Frame **A**

I/O Expansion




CP Expansion **Factory Build Only**



CP Expansion **Factory Build Only**

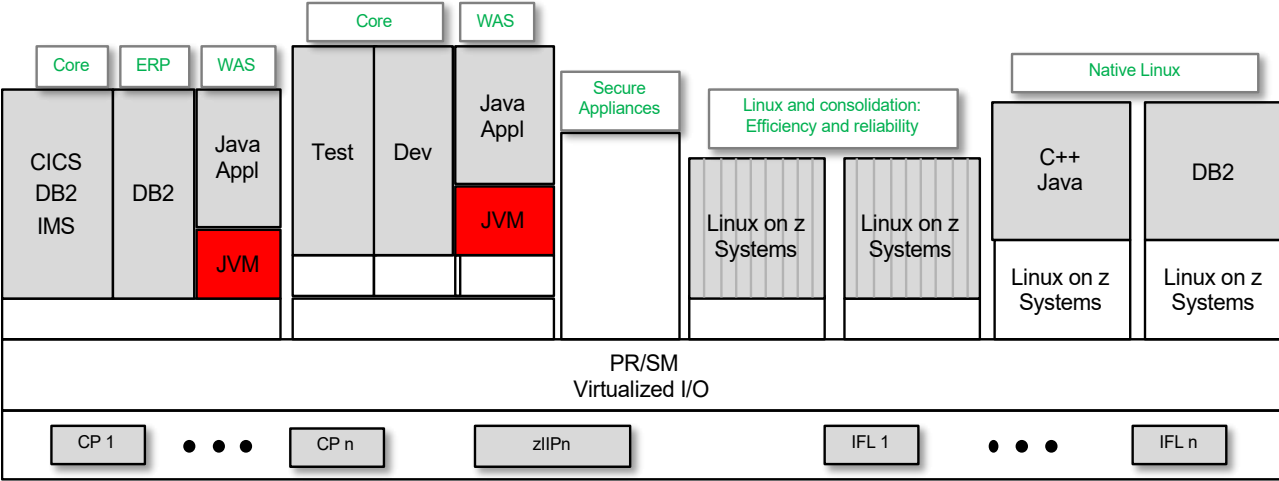
CP and I/O Expansion **Factory Build Only**



CP and I/O Expansion **Factory Build Only**

*4 line cords required in the A frame if more than 1 CPC drawer or 3 I/O drawers are present

IBM z17 Multi-frame operating systems



IBM z17 operating system support

z/OS


- z/OS V3.2 with PTFs
- z/OS V3.1 with PTFs
- z/OS V2.5 with PTFs
- z/OS V2.4 with PTFs for toleration only, Extended Support Required after Sept 2024

z/VM

- z/VM® 7.4 with PTFs
- z/VM 7.3 with PTFs

VSEⁿ

- [VSEⁿ V6.3 – 21st Century Software](#)



z/TPF

- z/TPF 1.1 with PTFs

Linux® on IBM Z

Minimum Distributions:

- SUSE SLES 15 SP6
- SUSE SLES 12 SP5
- Red Hat RHEL 9.4
- Red Hat RHEL 8.10
- Canonical Ubuntu 24.04 LTS
- Canonical Ubuntu 22.04 LTS

IBM cannot legally discuss IBM z17 multi frame exploitation prior to GA from distributors.

Officially Tested list [here](#).

Statement of Direction

Fiber Channel Endpoint Security: There is an increasing focus on the vital importance of cybersecurity from all sides, including regulatory groups and governments, and a particularly important area of cybersecurity is the protection of critical infrastructure. As many IBM Z® clients run the most mission critical applications, operate in highly regulated industries, and have an increasing amount of sensitive data, IBM must provide tools for securing client data and have a strong technology roadmap to continue to do so.

As data is being moved within and across data centers, authentication of the identities exchanging data and transparent encryption of the data in flight are required to strengthen security of the data. IBM Fibre Channel Endpoint Security (IFCES) is an end-to-end solution 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.

Given the increasing importance of providing the highest level of data protection to IBM Z clients, IBM intends to require the use of IBM Fibre Channel Endpoint Security for all FICON connected devices starting with the release of IBM z17+1. This direction will require investment by IBM Infrastructure teams, FICON storage vendors and IBM Z clients as an important step towards continuing to secure the most mission critical workloads. In support of this direction, all new FICON-connected storage systems introduced after December 31, 2024, will be required to support IFCES to connect to z17+1. IBM will continue its practice of licensing the FICON/IFCES specifications to vendors, and IBM will continue to actively support any vendor who chooses to implement these technologies into their products.

Clients are encouraged to begin or expedite their education on the requirements to enable adoption of IFCES. IBM has the tools and technology needed to assist clients in rapidly migrating to this more secure environment.

New Statement of Direction

IBM Spyre AI Accelerator: IBM intends to deliver the IBM Spyre AI Accelerator on Z, which is a PCIe-attached AI accelerator designed to support products, assistants, and services that leverage large language models on the IBM Z platform. The IBM Spyre AI Accelerator on IBM z17 is designed to bring generative AI to mainframes, empowering enterprises to scale-up AI capabilities and workloads. Clients will be able to deploy large language models (LLMs) including encoder, decoder, encoder-decoder architecture LLMs on the IBM Z platform, and benefit from the robust security, performance, and reliability it offers. IBM Z software products that will depend on Spyre Accelerators for full on-prem deployment include watsonx Code Assistant for IBM Z and watsonx Assistant for IBM Z. A hybrid deployment model is available for watsonx Assistant for IBM Z before Spyre Accelerator availability, as [announced here](#).

Other AI on IBM Z products that are also planned for Spyre AI Accelerator exploitation upon availability include Machine Learning for IBM z/OS, and components of AI Toolkit for IBM Z & LinuxONE like IBM Z Accelerated for Nvidia Triton Inference Server and IBM Z Accelerated for PyTorch. IBM Spyre AI Accelerator is planned to be available starting in 4Q 2025, in accordance with applicable import/export regulations

New Statement of Direction

AI Toolkit for IBM Z & LinuxONE with Spyre: IBM intends to deliver capabilities to exploit IBM Z Spyre AI accelerator with components of the AI Toolkit for IBM Z & LinuxONE. AI Toolkit for IBM Z & LinuxONE intends to deliver support for encoder large language models (LLMs) like BERT (Bidirectional Encoder Representation from Transformers) which clients will be able to infuse into their IBM Z applications to leverage AI at scale. Components of the AI Toolkit for IBM Z & LinuxONE that will exploit these capabilities using the Spyre AI Accelerator upon availability includes IBM Z Accelerated for Nvidia Triton Inference Server and IBM Z Accelerated for PyTorch.

Machine Learning for IBM z/OS with Spyre: IBM intends to deliver capabilities to exploit IBM Z Spyre AI accelerator with Machine Learning for IBM z/OS Enterprise Edition. Machine Learning for IBM z/OS intends to deliver support for encoder large language models (LLMs) like BERT (Bidirectional Encoder Representation from Transformers) which clients will be able to infuse into their IBM Z applications to leverage AI at scale. Machine learning for IBM z/OS will exploit these capabilities using the Spyre AI Accelerator upon availability.

Unlocking deeper insights for Db2 z/OS with AI on IBM Z: IBM intends to provide advanced vector database capabilities that exploit the IBM z17 hardware AI acceleration features with Db2 for z/OS. We're bringing together capabilities to enable clients to unlock deeper insights from their structured and unstructured data.

New Statement of Direction

Changes to TKE: IBM z17 is the last machine family where the tower hardware server form factor of the Trusted Key Entry (TKE) will be supported. The 1U TKE is available on the z17, and the 1U TKE will be the only supported hardware in the future. For console room use of the 1U TKE, IBM recommends that there is consideration to mount the 1U TKEs in a mini rack with acoustical noise protection.

Removal of support for White Smart Cards: Due to changes in technology, IBM is withdrawing support for white smart cards (parts 45D3398, 74Y0551, and 00JA710). White smart cards will not work on any release of TKE past release TKE 10.1. To prepare for this, clients must purchase new smart cards through Feature Code 0889 or 0886, initialize a set of smart cards that corresponds to each of your existing white smart cards, and copy the content from each old to new smart card.

Future Direction for IBM Threat Detection for z/OS: IBM intends to add network anomaly detection to its IBM Threat Detection for z/OS product.

New Statement of Direction

IBM intends to deliver anti-malware for IBM z/OS: IBM® plans to provide a software solution that introduces cyber anomaly detection and notification for the z/OS® platform to mitigate the potential risk of malicious software. IBM plans to provide the option of quarantine functionality that further extends existing remediation options. It is the intent for these combined functions, per NIST guidelines, to be used by the client to satisfy compliance regulations requiring anti-malware coverage for z/OS. This intent includes standards such as the Payment Card Industry Data Security Standard (PCI DSS) version 4.0.

To further defend against the potential risk of malicious software, IBM plans to enhance the IBM z/OS Authorized Code Scanner to provide static scanning of authorized code, adding to the IBM z/OS Authorized Code Scanner feature's existing collection of its dynamic scanning for development and test environments and its authorized code monitor for production systems.

IBM also plans to provide a software solution that simplifies z/OS data set encryption, encrypting and re-encrypting data at scale for both key rotation and initial encryption, and leveraging analytics to minimize application downtime. This is designed to simplify adherence to expanded compliance regulations such as PCI DSS v4.0.

New Statement of Direction

Future direction for tagging sensitive data on IBM Z: IBM intends to leverage the Telum II Processor, artificial intelligence, and machine learning technology to identify and classify sensitive data in z/OS data sets. This will provide an automatic function to discover, classify and tag data which will support organizations in implementing more targeted and effective security measures to protect that information.

Future direction for IBM Terraform - IBM intends in the near future to support IBM Terraform for Z and IBM® LinuxONE to enable the configuration and management of Z resources through Infrastructure as Code.

New support for BCPii: BCPii v1 and HMC/Support Element (SE) Simple Network Management Protocol (SNMP) are being deprecated. No future feature enhancements will be added to BCPii v1 or SNMP, but there are no plans to remove existing support for BCPii v1 or HMC/SE SNMP. All future automation functional enhancements will only be done for BCPii HWIREST/v2 and HMC WebServices APIs.

T10 Data Integrity Field (T10 DIF): As per a [March 2025 Statement of Direction announcement](#) IBM z16 is intended to be the last IBM Z generation to support T10 Data Integrity Field (T10 DIF)

New Statement of Direction

Removal of Support for Constrained Transactional Execution: IBM z17 is planned to be the last IBM Z hardware generation to support Constrained Transactional Execution. On subsequent IBM Z hardware generations, the TBEGINC instruction that is used to start a Constrained Transaction will receive an operation exception, and no Constrained Transactional Execution will be initiated or supported.

IBM recommends that programs currently using Constrained Transactional execution should dual path their usage, so that when running on an IBM Z hardware generation that supports Constrained Transactional Execution, Constrained Transactional Execution continues to be used; but when running on an IBM Z hardware generation that does not support Constrained Transactional Execution, alternative mechanisms are used instead.

IBM z17 provides several new Perform Locked Operation (PLO) functions that are intended to provide a set of alternative capabilities that can be used in lieu of Constrained Transactional Execution. Note that many other existing serialization mechanisms can also be used to provide the necessary atomicity as an alternative to the use of Constrained Transactional Execution.

Note that support for Non-Constrained Transactional Execution is also reduced starting with z17. Starting with z17, the TBEGIN instruction used to start a Non-Constrained Transaction will always complete with with Condition Code 1 and with no Transactional Execution initiated, requiring the program to use its mandatory “fallback path” for execution. It is planned that this reduced support for Non-Constrained Transactional Execution will continue on subsequent IBM Z hardware generations.

New Statement of Direction

Removal of Support for 10 Gb Coupling Express3 LR Coupling Links (CL5): z17 introduces new support for Coupling Express3 LR 25G Coupling Link adapters to provide long-distance coupling link connectivity with 25 gigabit bandwidth (CL6), in addition to supporting the existing 10 gigabit bandwidth using Coupling Express3 LR 10G (CL5).

The machine after z17 is planned to be the last IBM Z hardware generation which will support the 10 gigabit bandwidth (CL5) coupling links. The subsequent hardware generation after that is planned to support only 25 gigabit bandwidth CL6 coupling links for long-distance coupling link connectivity.

This approach supports IBM's usual n/n-2 sysplex connectivity policy regarding the transition off of 10 gigabit bandwidth CL5 coupling links and onto 25 gigabit bandwidth CL6 coupling links for long-distance connectivity. clients should plan to complete this transition within their sysplex environments before introducing a machine of the future IBM Z hardware generation that only supports 25 gigabit bandwidth CL6 coupling links for long-distance coupling link connectivity; all earlier machines that will need to connect to such a machine must also have already migrated to the use of 25 gigabit bandwidth CL6 coupling links by that time.

New Statement of Direction

Future Linux Distribution Support: The following Linux distributions plan to support IBM z17 and IBM® LinuxONE Next generation models for a final time, and do not support future generation IBM Z and IBM® LinuxONE generation models. RHEL 8

The upcoming SUSE Linux Enterprise Server 16 will support IBM z17, and SUSE is actively collaborating with IBM and the communities to extend support for next generations of the IBM Z platform.

Red Hat and IBM intend to deliver Red Hat OpenShift AI on IBM Z and IBM® LinuxONE.

