IBM z15 Model T02 Hardware Overview

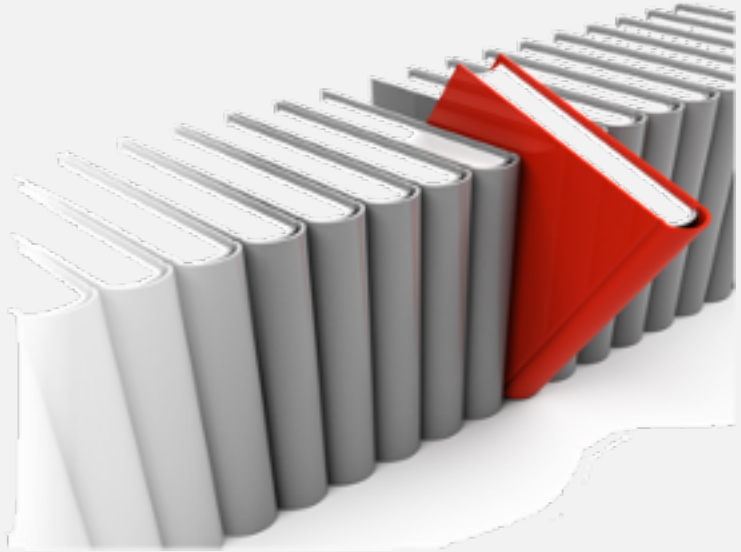

Kenny Stine
IBM Z Technical Specialist
Washington Systems Center
kjstine@us.ibm.com

Trademarks

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 www.ibm.com/legal/copytrade.shtml:

*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™, **z14™**, **z15™**, 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.

Table of Content

z15 Announcement Dates

z15 T02 and LinuxONE III LT2 key Dates 2020

- **May 15, 2020**
  - New Build Systems:
    - IBM z15 Model T02
    - IBM LinuxONE III LT2
  - Features and functions for the IBM z15 Model T02
    - Upgrades from ZR1: 3907 Model ZR1 upgrades to IBM z15 Model T02

z15 T02 and LinuxONE III LT2 key Dates 2020

- **May 15 2020**
  - MES orders for the IBM z15 Model T02 that are delivered solely through a modification to the machine's Licensed Internal Code (LIC)
  - HMC Table Top KMM (#0148) on IBM z15 Model T02
  - HMC Rack Mount KMM (#0154) on IBM z15 Model T02
  - Customer Supplied HMC KMM (#0188) on IBM z15 Model T02
  - HMC Tower (#0062) on IBM z15 Model T02
  - HMC Rack Mount (#0063) on IBM z15 Model T02
  - TKE Rack Mount (#0087) on IBM z15 Model T02
  - TKE (#0088) on IBM z15 Model T02
  - TKE 9.2 LIC (#0881) on IBM z15 Model T02
  - TKE Rack Mount KMM (#0156) on IBM z15 Model T02
  - TKE Table Top KMM (#0157) on IBM z15 Model T02
  - Customer Supplied TKE KMM (#0190) on IBM z15 Model T02
  - Smart Card Reader (#0891) on IBM z15 Model T02
  - Additional Smart Cards (#0900) on IBM z15 Model T02

z15 T02 and LinuxONE III LT2 key Dates 2020

- **September 17, 2020**
  - All remaining MES orders for IBM z15 Model T02/LT2

z15 Design Principles

Elevate your hybrid cloud with IBM z15

Service Level Excellence

Industry's highest level of business uptime to meet SLA and regulatory compliance

Data Protection & Privacy

Industry-first solution to protect sensitive data across your multicloud

Mission Critical Cloud

Integrate seamlessly into hybrid multicloud, blockchain and AI

Standardized & Flexible for the Cloud Data Center

Modular, scalable and proven cloud-ready infrastructure

z15 Overview

z15 T02 Functions and Features (Driver Level 41)

System, Processor, Memory
One model with one or two CPC drawers, five available sizes
12 core 14nm PU SCM (7, 8, 9, 10 or 11 active cores per PU SCM)
Up to 65 configurable PUs as CPs, zIIPs, IFL, ICFs, or optional SAPs (up to 6 CPs)
Up to 14% increased uni-processor capacity vs. z14 ZR1
156 CP capacity settings
19” Rack, ASHRAE class A3 (for Data Center requirements relief)
Integrated (on-chip) Accelerator for zEDC
Enhanced processor / cache design with bigger cache sizes
Up to 16 TB of Memory protected by Redundant Array of Independent Memory (RAIM)
Two 8 U Reserved (rack space) features
Up to 40 LPARs
IBM Dynamic Partition Manager
Secure Service Container
LPAR Group Absolute Capping
CPUMF sampling w/o PE Mode enablement

I/O Subsystem, Parallel Sysplex, STP, Security
Up to 4 PCIe+ Gen3 I/O fanouts (dual port)
PCIe+ I/O Drawer (up to 4 per system, up to 64 PCIe features),
3 LCSSs, 3 Subchannel Sets per LCSS
32K I/O Devices per channel for all FICON features
FICON Express16S+
zHyperLink Express 1.1
25 GbE RoCE Express2.1
Shared Memory Communications - Direct Memory Access over Internal Shared Memory (ISM) – SMC-D
Virtual Flash Memory (512 GB per feature, up to four features)
CFCC Level 24
Crypto Express7S and Crypto enhancements

RAS, Other Infrastructure Enhancements	
Hardware Management Appliance	Single Keyboard Mouse Monitor for SE (USB-C)
STP Enhancements – Introducing Precision Time Protocol	Rack-Mounted Support Elements
Key Locks for doors	Tower & Rack-mounted HMCs and TKEs
Support for ASHRAE Class A3 datacenter	TKE 9.2 LICC

z15 T02 System Design Changes (vs. z14 ZR1)

- 14 nm Processor with optimized Out-of-Order, new DEFLATE and Modulo Arithmetic
- 12 Cores per PU SCM design (7 – 11 active)
- 4 CP SCMs per Drawer, one or two CPC drawers
- Integrated I/O with PCIe+ Gen3
- Single System Controller Chip, 960MB L4
- Simplified CPC Drawer SMP Fabric

- Integrated (on-chip) compression
- Crypto Express7S (single/dual port)
- 25GbE and 10GbE RoCE Express2.1
- IBM zHyperLink Express1.1
- Integrated Coupling Adapter SR1.1

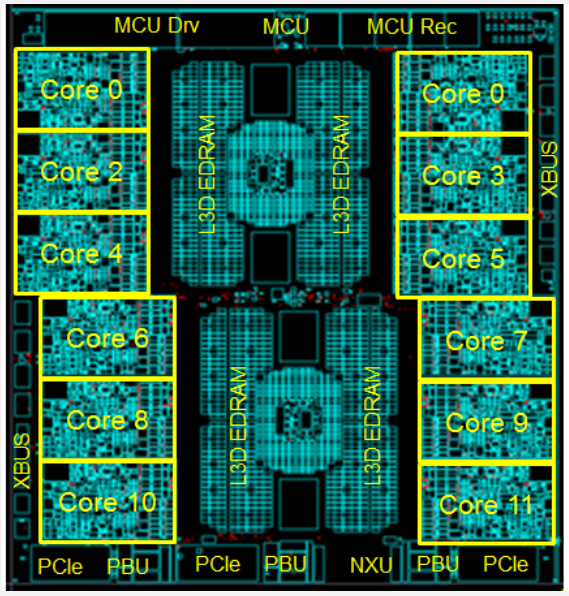

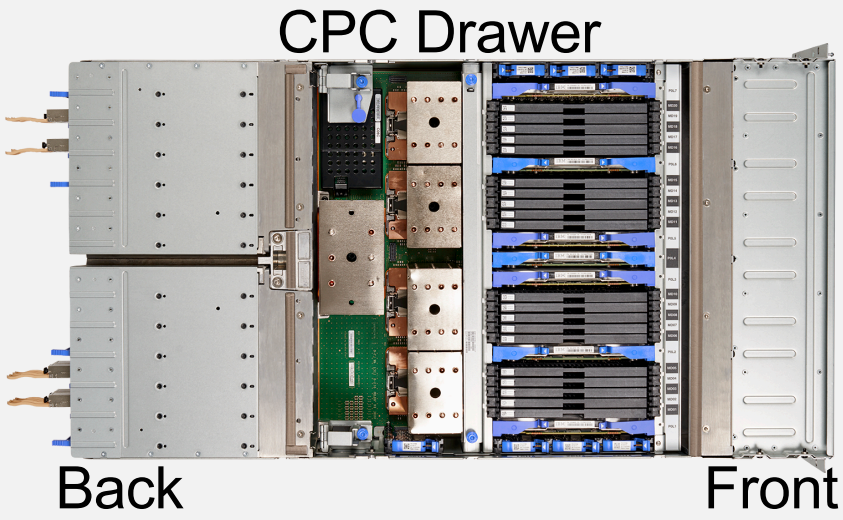

IBM Z – Processor Roadmap

45 nm
z196
9/2010

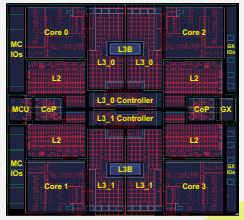

- Top Tier Single Thread Performance,System Capacity
- Accelerator Integration
- Out of Order Execution
- Water Cooling
- PCIe I/O Fabric
- RAIM
- Enhanced Energy Management

32 nm
zEC12
8/2012

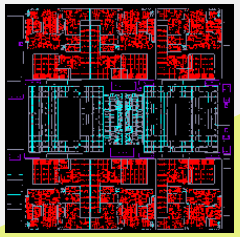

- Leadership Single Thread, Enhanced Throughput
- Improved out-of-order
- Transactional Memory
- Dynamic Optimization
- 2 GB page support
- Step Function in System Capacity

22 nm
z13
1/2015

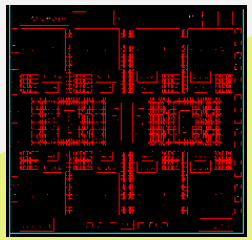

- Leadership System Capacity and Performance
- Modularity & Scalability
- Dynamic SMT
- Supports two instruction threads
- SIMD
- PCIe attached accelerators
- Business Analytics Optimized

14 nm
z14
7/2017

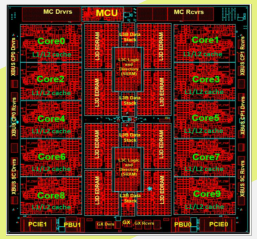

- Pervasive encryption
- Low latency I/O for acceleration of transaction processing for DB2 on z/OS
- Pause-less garbage collection for enterprise scale JAVA applications
- New SIMD instructions
- Optimized pipeline and enhanced SMT
- Virtual Flash Memory

14 nm
z15
9/2019

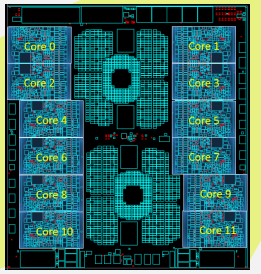

- Focus on power efficiency and new on-chip architectures
- Improved and enlarged caches
- Optimized Out-of-Order architecture
- New Secure Execution support
- On-chip compression support (DEFALTE)

*Capacity and performance ratios are 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 .

z15 Continues the CMOS Mainframe Heritage

PCI – Processor Capacity Index
Full capacity uniprocessor

MHz, PCI*

4500
4000
3500
3000
0

661*

z10 BC
65 nm SOI
3.5 GHz
5 CPs (10 IFLs)

2008

782*
+18%

z114
45 nm SOI
3.8 GHz (+8.6%)
5 CPs (10 IFLs)

2011

1064*
+36%

zBC12
32 nm SOI
4.2 GHz (+10.5%)
6 CPs (13 IFLs)

2013

1430*
+34%

z13s
22 nm SOI
4.3 GHz (+2.4%)
6 CPs (20 IFLs)
Up to 4 TB memory

Introduced:
SMT and SIMD

2016

1570*
+10%

z14 ZR1
14 nm SOI
4.5 GHz (+4.6%)
6 CPs (30 IFLs)
Up to 8 TB memory

19” Rack;
Enhanced
SMT and SIMD

2018

1761*
+12%

z15 T02
14 nm SOI
4.5 GHz (==)
6 CPs (65 IFLs)
Up to 16 TB memory

New on-core and
on-chip
architectures

2020

*NOTE: PCI (MIPS) Tables are NOT adequate for making comparisons of IBM Z processors in proposals

z15 T02 sub-capacity CP Granularity

- 26 CP capacity levels (26 x 6 = 156)
- Processor Value Unit for IFL = 100 (No change since z10 BC)

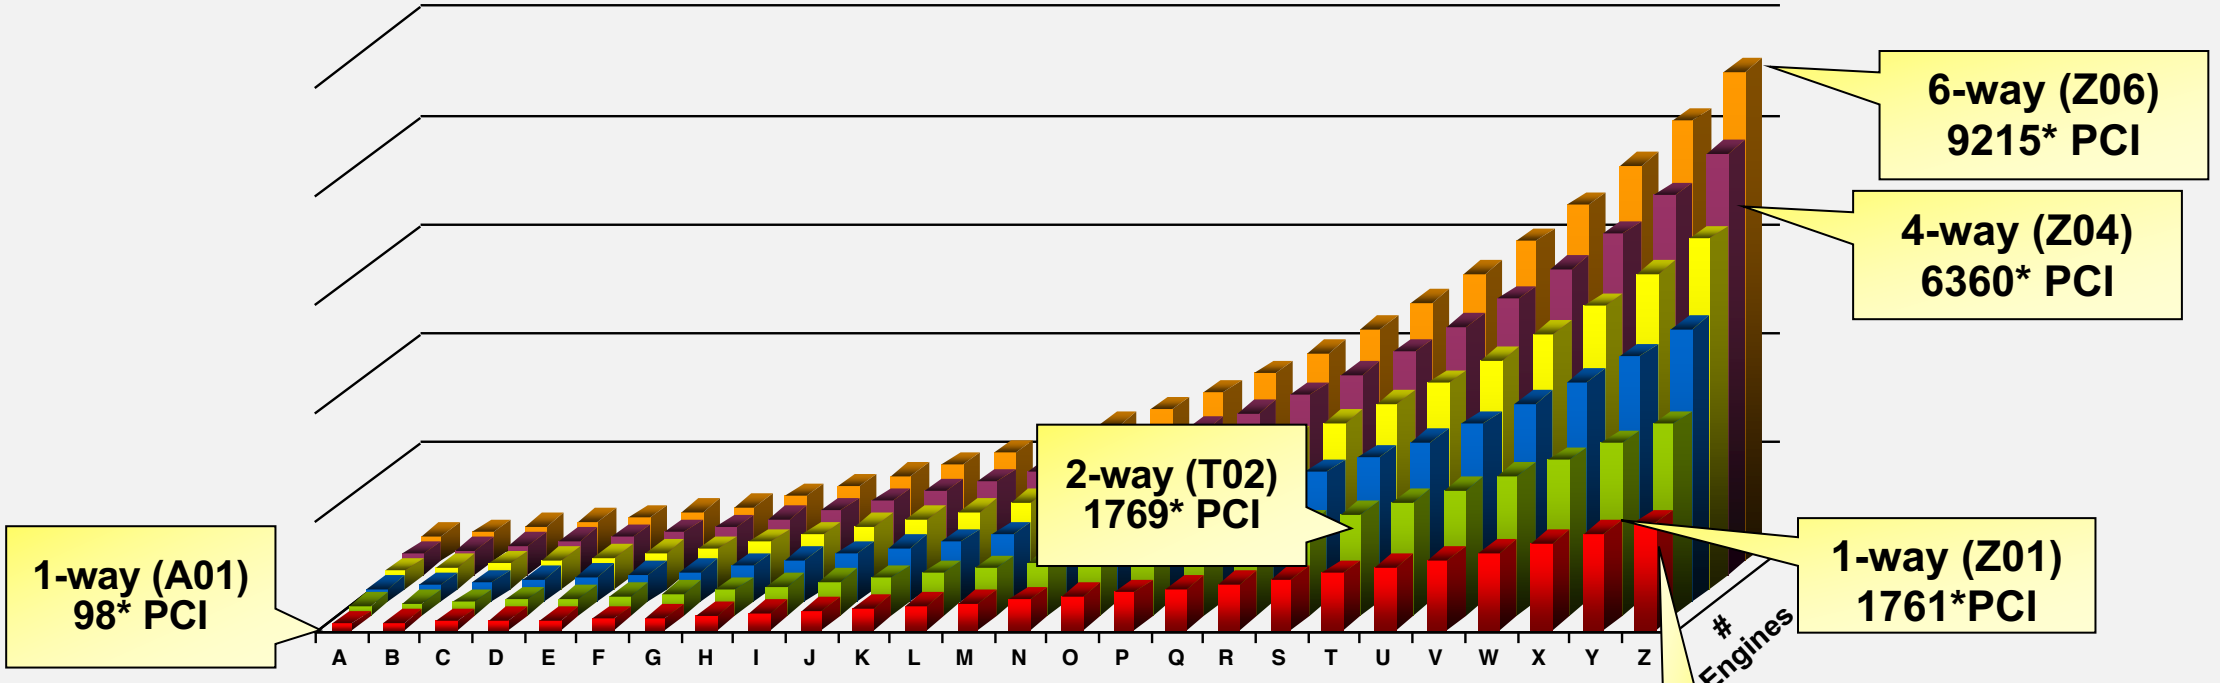

PCI – Processor Capacity Index

*Capacity and performance ratios are 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 .

Call to Action – Sizing done right – Best Practices
CPU Measurement Facility

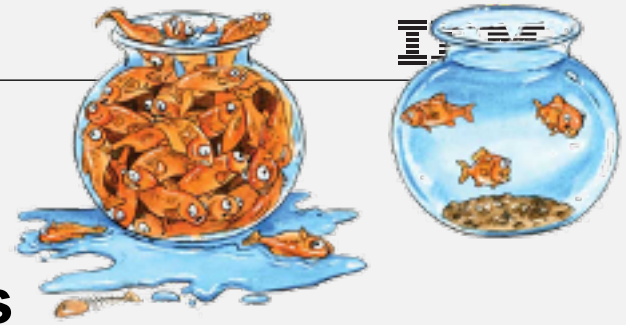

- **Ensure the CPU MF data is captured and kept for analysis**

- **Performance, Capacity Planning and Problem Determination**

- **Critical Migration Action for every IBM Z (z/OS and z/VM)**
  - **CPU MF Counters must be enabled on their current processor**
  - **CPU MF Counters must be enabled on their z15**

In z/OS there is a HIS started task.
This is run on each System/LPAR and writes SMF 113 records.
This should be set up and run on all partitions.

z/VM also gathers CPU MF Counters through new z/VM Monitor Records.

z15 Processor Design and Structure

z15 Processor Drawer (Top View)

- Each PU SCM:
  - 14nm
  - Four PU SCMs
  - One Memory Controller per PU Chip
  - Five DDR4 DIMM slots per Memory Controller
    - 20 DIMMs total per drawer

- Each drawer:
  - Two logical PU clusters (0 and 1)
  - Up to Four PU Chips per CPC Drawer
    - 12 PUs / SCMs, (7-11 active cores PNs)
    - 8,17,27,38 active PUs per drawer
    - 41 active PUs per drawer – Max4, Max13, Max21, Max31 and Max65
  - One SC Chip (960 MB L4 cache)
  - DIMM slots: 20 DIMMs to support up to 8 TB of addressable memory (10 TB RAIM)
  - Two Flexible Support Processors/ OSC Cards
  - 12 fanout slots for PCIe+ I/O drawer or PCIe coupling fanouts (ICA SR).

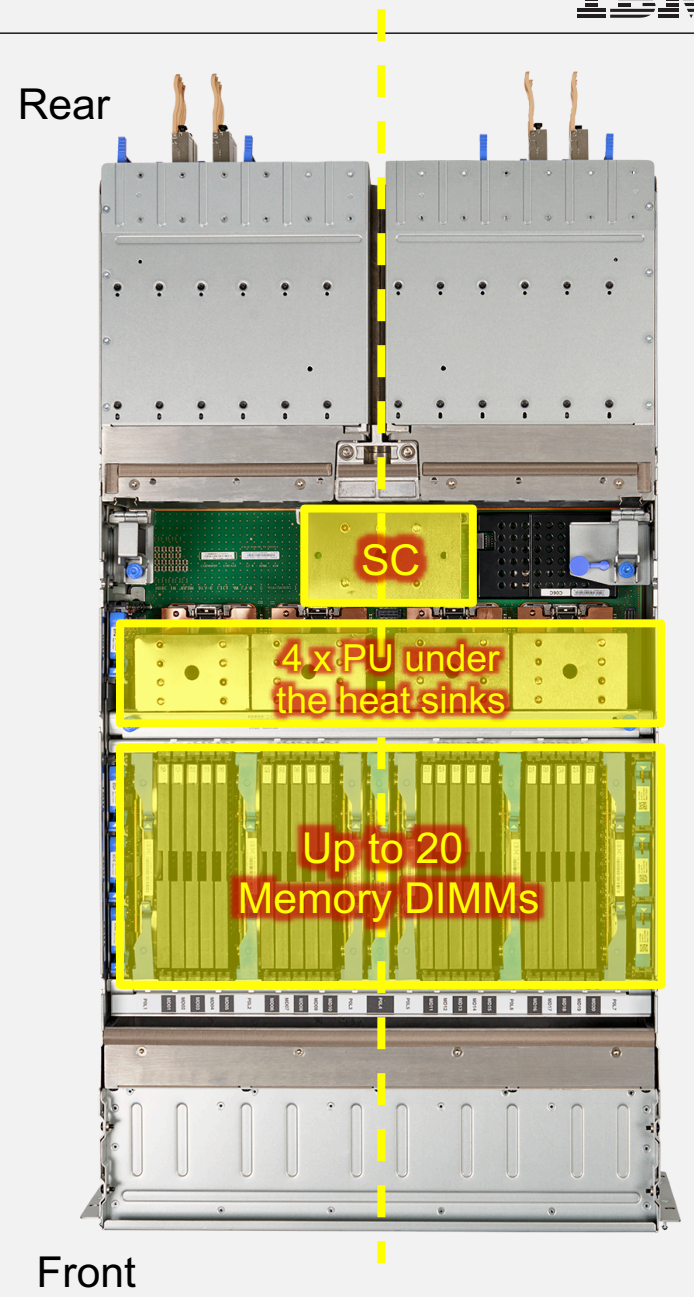

z15 Processor (CPC) Drawer Connectivity

- 12 PCIe fanout slots per z15 CPC drawer
  - Increase from 8 PCIe fanouts in z14 ZR1
- Integrated Coupling Adapter (ICA) SR1.1
  - Two ports @ 8 GBps* (PCIe Gen3) for short distance coupling
  - 150m fiber optic coupling link
- PCIe+ Fanout Gen3
  - Two ports @ 16GBps (PCIe Gen3)
  - Connects to the PCIe Interconnect Gen3 in the PCIe+ I/O drawers

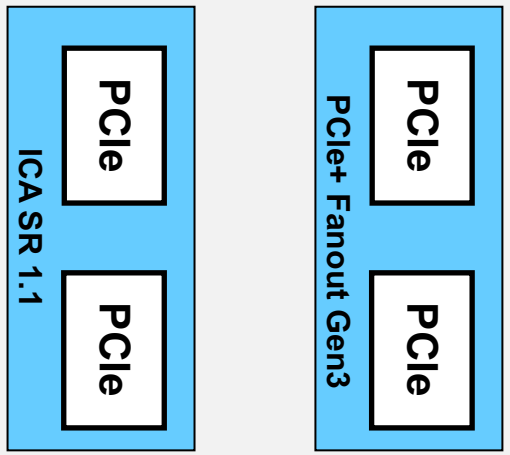

ICA SR1.1 Coupling Links

To PCIe+ I/O Drawers

No InfiniBand fanouts

Rear

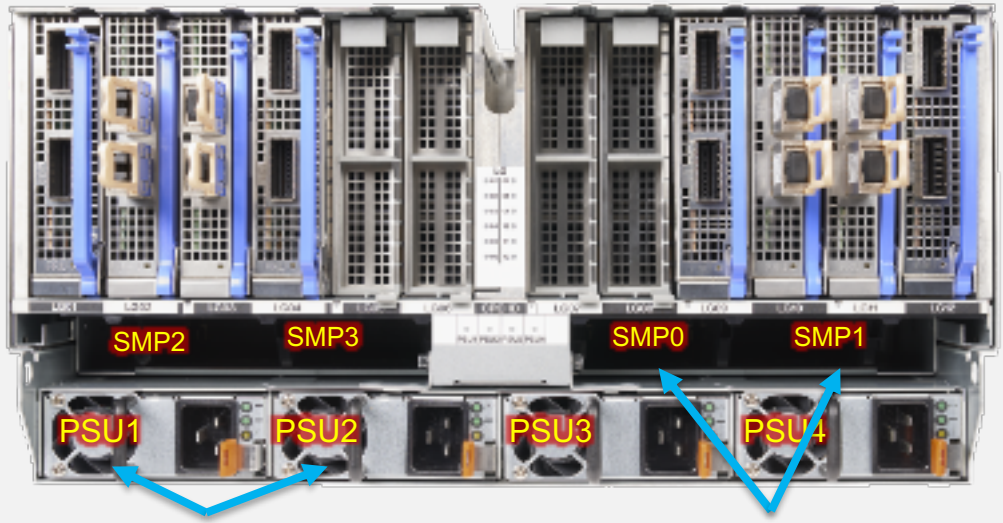

Power from iPDU/BPA

SMP Connection to other CPC Drawers

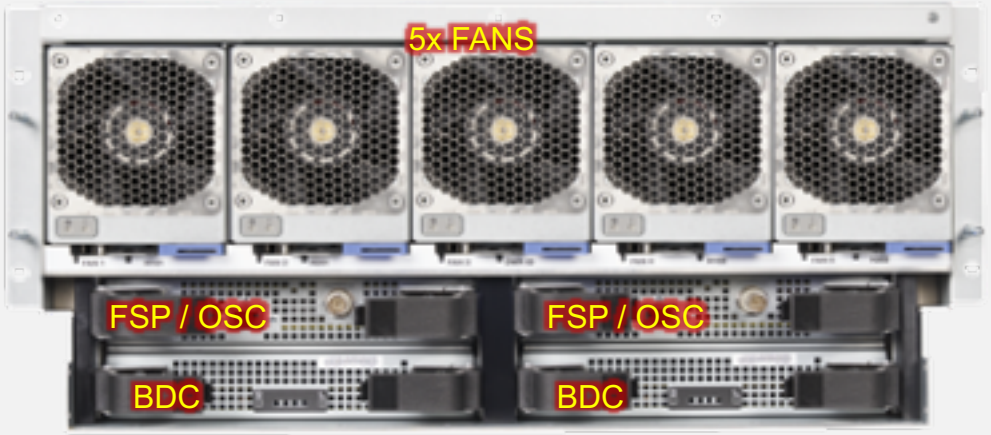

Front

* The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

12-Core Processor Chip Detail

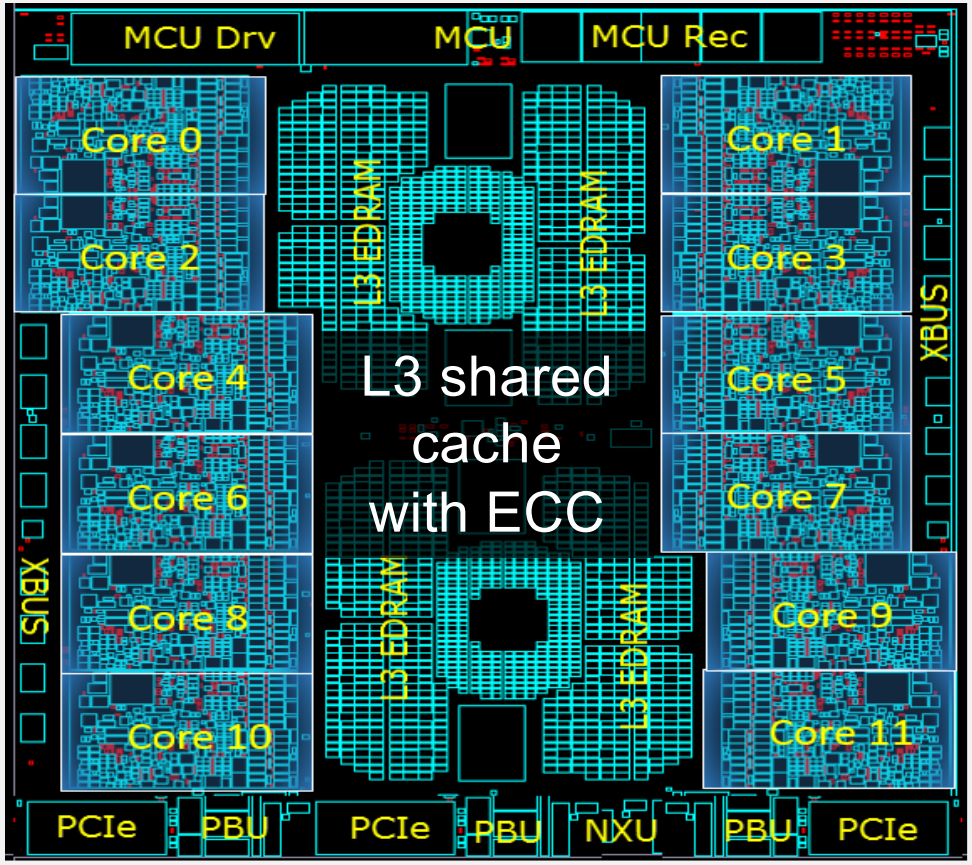

- **14nm SOI Technology**
  - 12 Cores
  - 17 layers of metal
  - 696 mm2 chip area
  - 9.2B transistors versus 6.2B on z14 ZR1
- **20% reduction area**
- **20% reduction in power**

- **4.5 GHz core frequency**
- **7, 8, 9, 10 or 11 active cores per chip**
- **IBM Integrated Accelerator for z Enterprise Data Compression (zEDC)**
  - **On-chip compression accelerator (NXU)**
- **On Core L1/L2 Cache**
  - L2-I from 2MB to 4MB per core
- **On chip L3 Cache**
  - Shared by all on-chip cores
  - Communicates with cores, memory, I/O and system controller single chip module.
  - L3 from 128MB to 256MB per chip
- **I/O buses**
  - Each CP chip will support up to 3 PCIe buses
    - PCIe+ I/O Drawer Fanout
    - ICA SR 1.1 Coupling Links

z15 Core Details (1 of 2)

- Each processor unit, or core, is a superscalar and out-of-order processor that supports 10 concurrent issues to execution units in a single CPU cycle*:
  - Fixed-point unit (FXU): The FXU handles fixed-point arithmetic
  - Load-store unit (LSU): The LSU contains the data cache. It is responsible for handling all types of operand accesses of all lengths, modes, and formats as defined in the z/Architecture
  - Instruction fetch and branch (IFB) (prediction) and Instruction cache & merge (ICM). These two sub units (IFB and ICM) contain the instruction cache, branch prediction logic, instruction fetching controls, and buffers. Its relative size is the result of the elaborate branch prediction
  - Instruction decode unit (IDU): The IDU is fed from the IFU buffers, and is responsible for parsing and decoding of all z/Architecture operation codes
  - Translation unit (XU): The XU has a large translation lookaside buffer (TLB) and the Dynamic Address Translation (DAT) function that handles the dynamic translation of logical to physical addresses
  - Instruction sequence unit (ISU): This unit enables the out-of-order (OoO) pipeline. It tracks register names, Out-of-Order instruction dependency, and handling of instruction resource dispatch

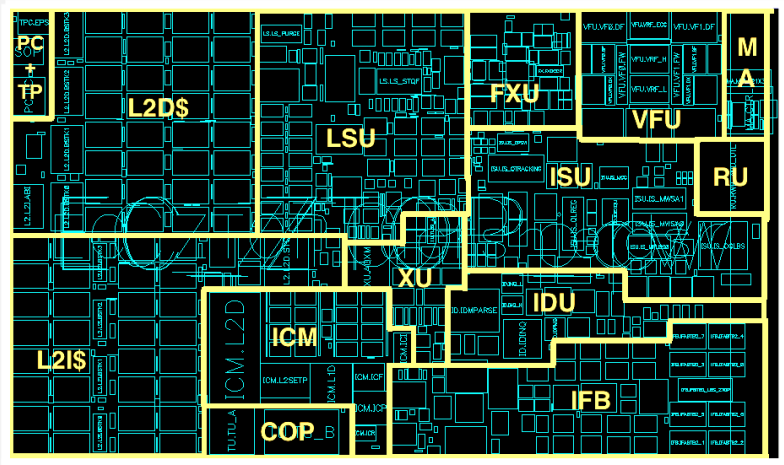

*See speaker notes

z15 Core Details (2 of 2)

- Instruction fetching unit (IFU) (prediction): These units contain the instruction cache, branch prediction logic, instruction fetching controls, and buffers. Its relative size is the result of the elaborate branch prediction design.
- Recovery unit (RU): The RU keeps a copy of the complete state of the system that includes all registers, collects hardware fault signals, and manages the hardware recovery actions.
- Dedicated Co-Processor (CoP): The dedicated coprocessor is responsible for data compression, encryption functions and Merge/SORT for each core
- Core pervasive unit (PC) for instrumentation, error collection
- Modulo arithmetic (MA) unit – support for Elliptic Curve Cryptography
- Vector and Floating point Units (VFU):
  - BFU - binary floating point unit,
  - DFU - decimal floating point unit,
  - DFx - decimal fixed point unit,
  - FPd - floating point divide unit,
  - VXx - vector fixed point unit unit,
  - VXs – vector string unit,
  - VXp - vector permute unit,
  - VXm - vector multiply unit
- L2I/L2D – Level 2 instruction/data cache

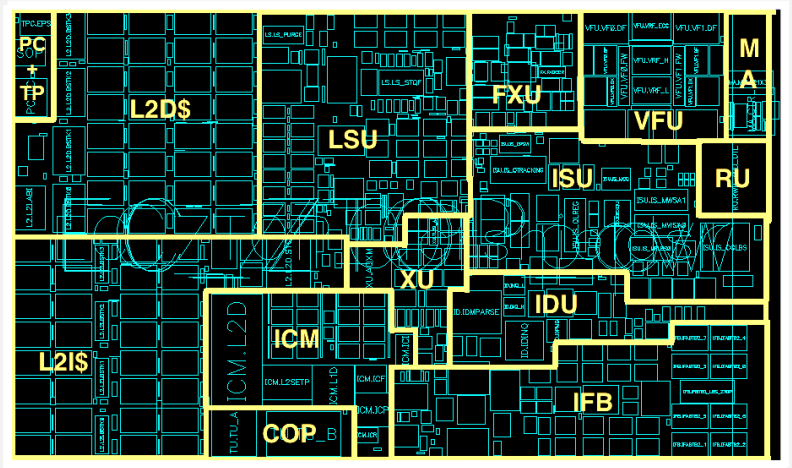

z15 Processor Enhancements Overview*

- Divide engine scheduler
- BFU „tandem“ operation:
  - 2x single precision bandwidth
- BFU latency reduction
- Double-bandwidth vector loads
- Mulitply/Divide speed ups
- Conversion speed ups
- New and enhanced vector instructions
  - Load/store reversed (to help with little endian conversion)
  - More vector shift operations
  - VECTOR STRING SEARCH, for fast string search, supporting different encodings
  - New vector FP converts
- New and enhanced Vector Packed Decimal instructions
- Enhanced Out-of-Order:
  - GCT size increase
  - Issue queue size increase
  - New Mapper design
- New TAgged Geometric Page History Table (PHT)
- Branch prediction – simplified and larger BTB
- Add special hardware to significantly accelerate frequently used functions:
  - Compression/Decompression (DEFLATE)
  - Sort
- Load/Store unit enhancements
- Operand Store Compare improvements

z15 Co-Processor (COP) Overview

- Co-Processor results (data) now stored direct via L1D cache
- Re-designed crypto/hashing/UTF-conversion/compression engines for increased throughput
- True Random Number Generator

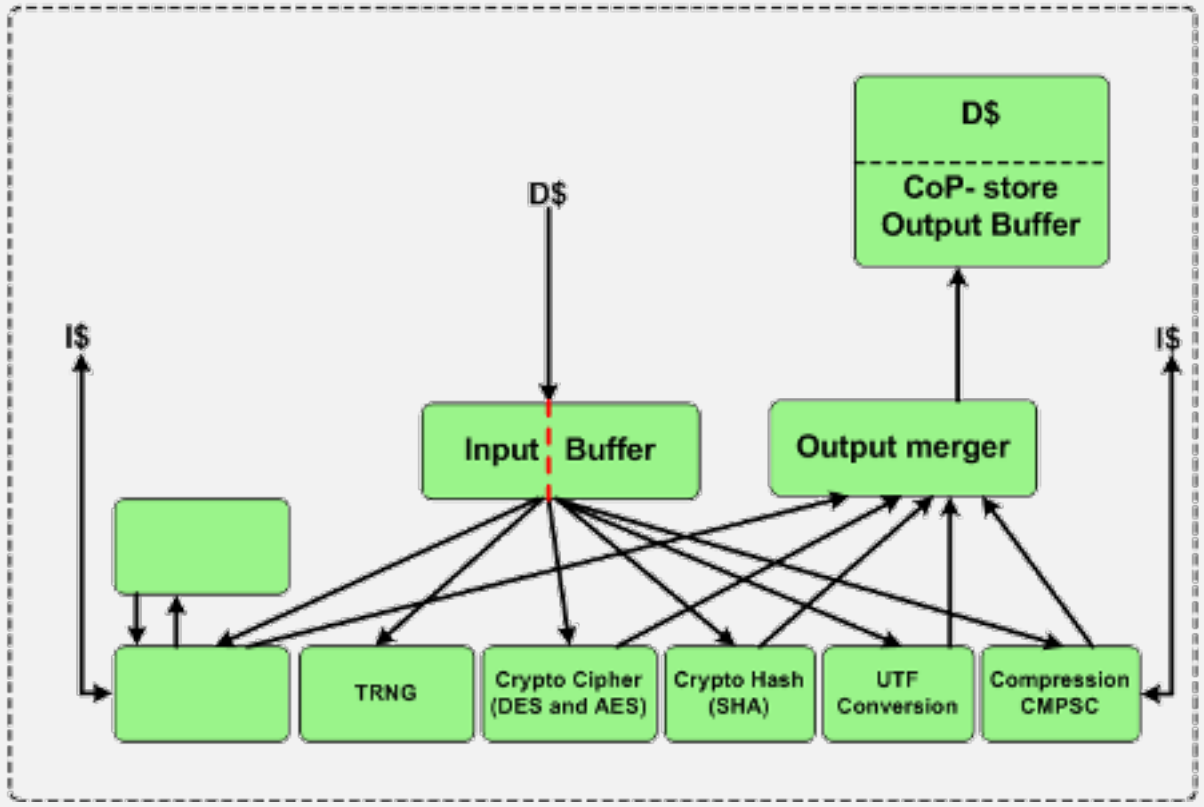

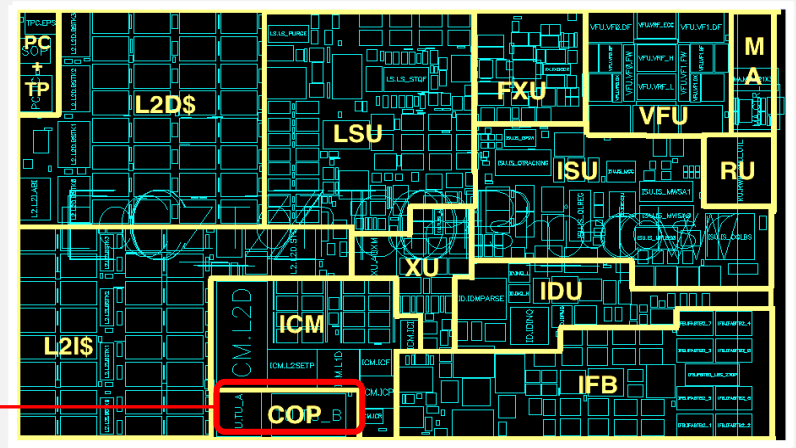

z15 SC Chip

- 14nm SOI technology
- 960 MB shared eDRAM L4 Cache
- System Interconnect
- System Coherency Manager
- X and A Bus - Support for:
  - 4 CPs using 4 x-buses
  - 2 drawers using 1 A-buses (point-to-point).

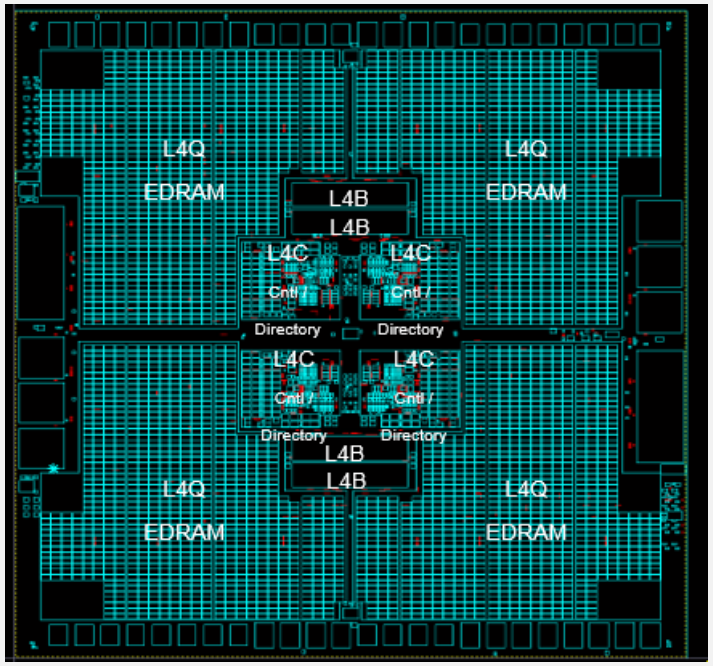

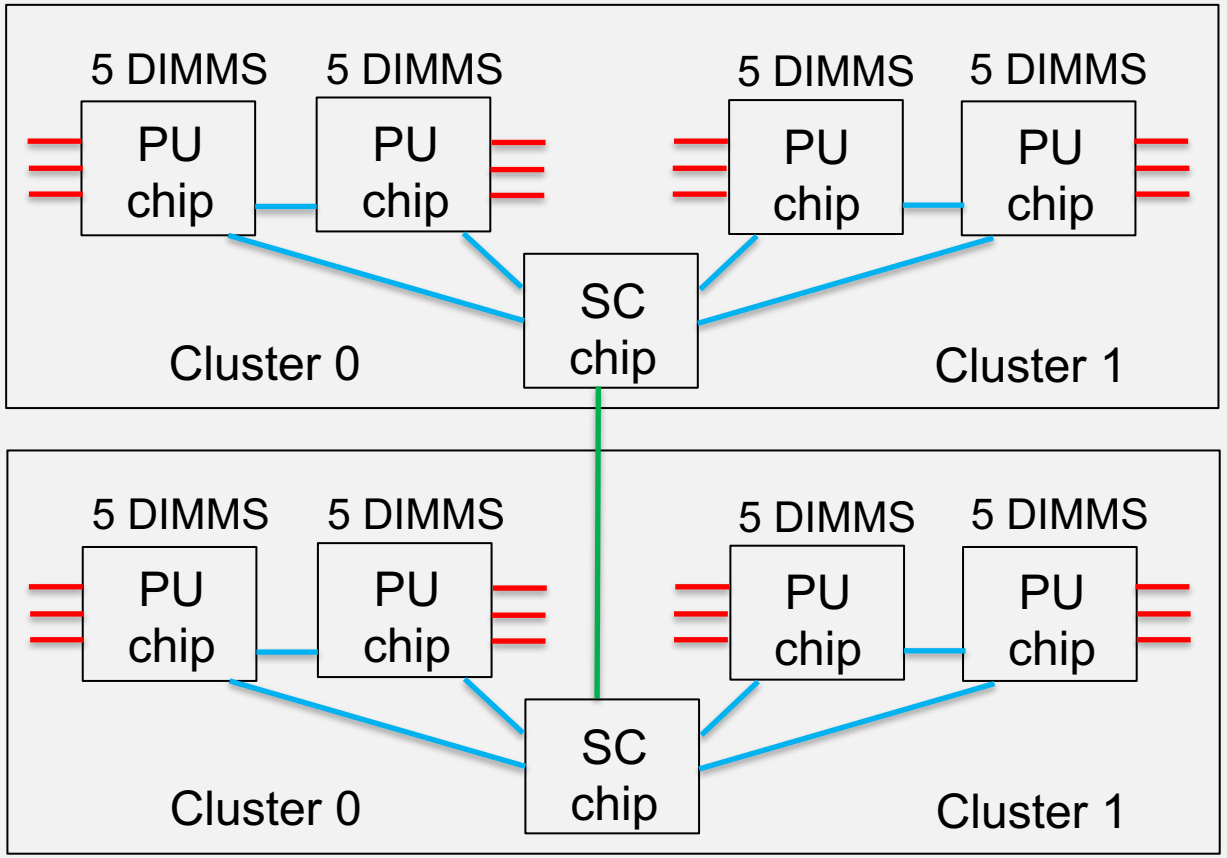

3 PCIe busses to I/O fanouts

Cache topology comparison: z15 T02 vs. z14 ZR1

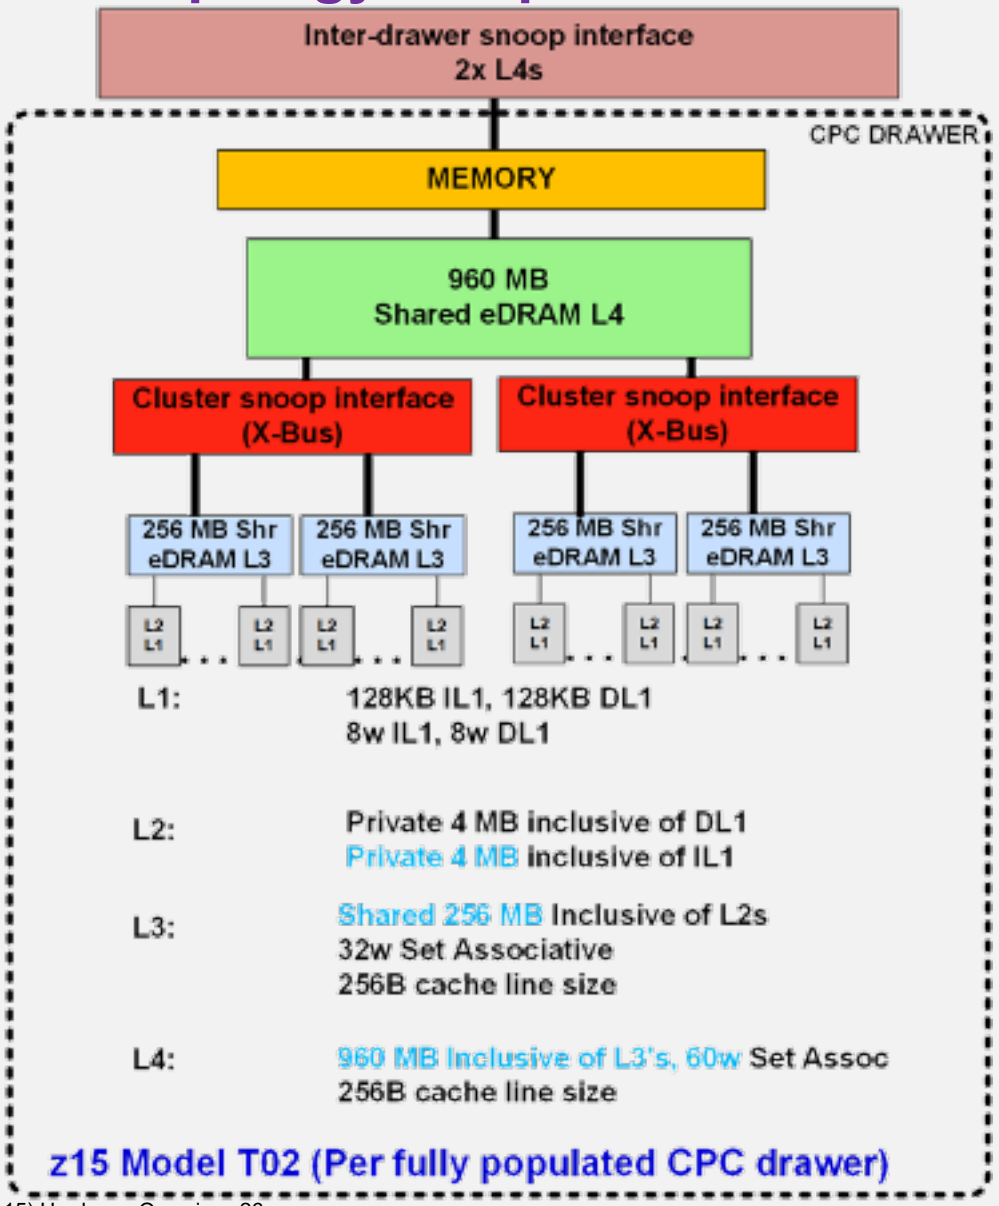

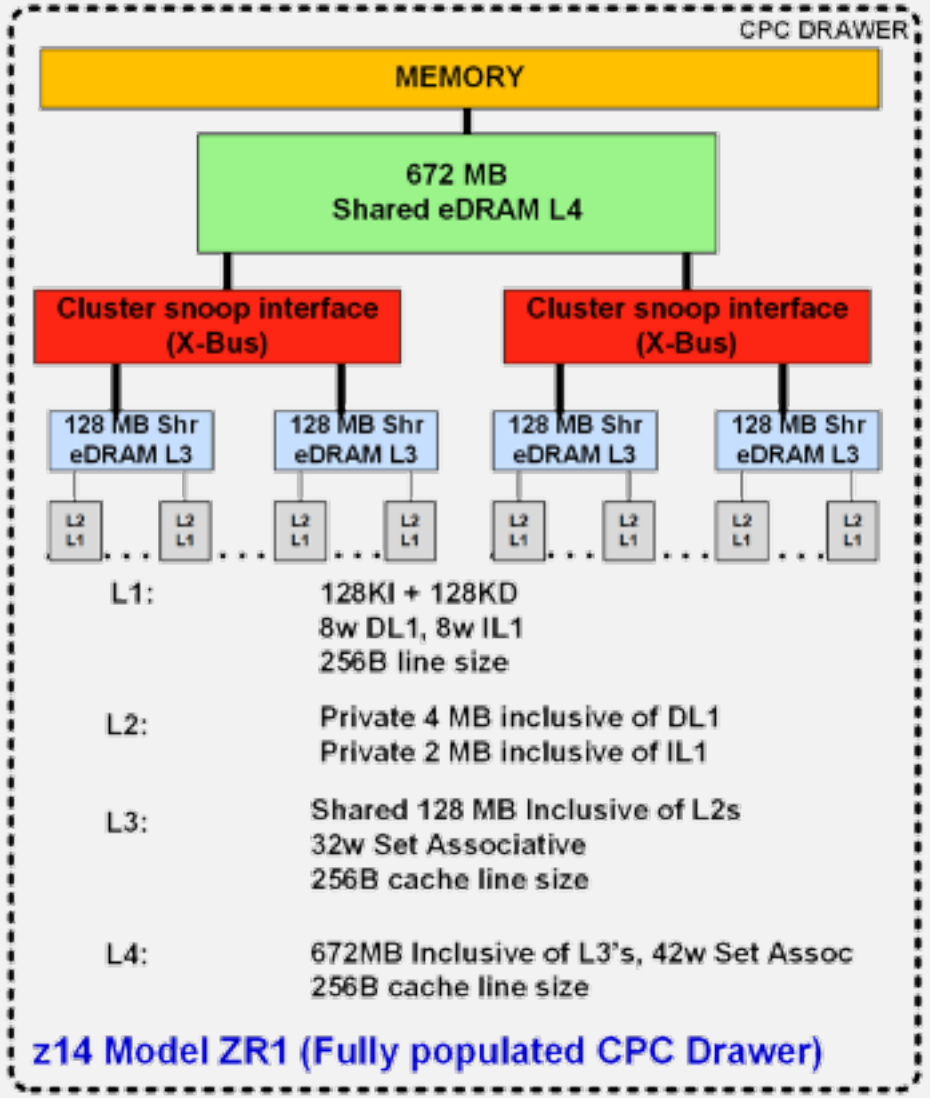

8562-T02 Processing Units

Model	Feature Code Description	Drawers/ Cores	Chips/ CPs	IFLs/ uIFLs	zIIPs	ICFs	Std SAPs	Optional SAPs	Std. Spares	IFP
T02 FC0505	0649 Max4	1/8	1 0-4	0-4 0-3	0-2	0-4	2	0-2	1	1
	0650 Max13	1/17	2 0-6	0-13 0-12	0-7	0-13	2	0-2	1	1
	0651 Max21	1/27	3 0-6	0-21 0-20	0-12	0-21	3	0-2	2	1
	0652 Max31	1/38	4 0-6	0-31 0-30	0-12	0-31	4	0-8	2	1
	0653 Max65	2/76	8 0-6	0-65 0-64	0-12	0-65	8	0-8	2	1

1. At least one CP, IFL, or ICF must be purchased in every machine.
2. Two zIIPs may be purchased for each CP purchased if cores are available. (2:1). This remains true for sub-capacity CPs and for “banked” CPs.
3. “uIFL” = Unassigned IFL
4. The IFP is conceptually an additional, special purpose SAP – used by PCIe I/O features and some other functions.

MES/Model Considerations

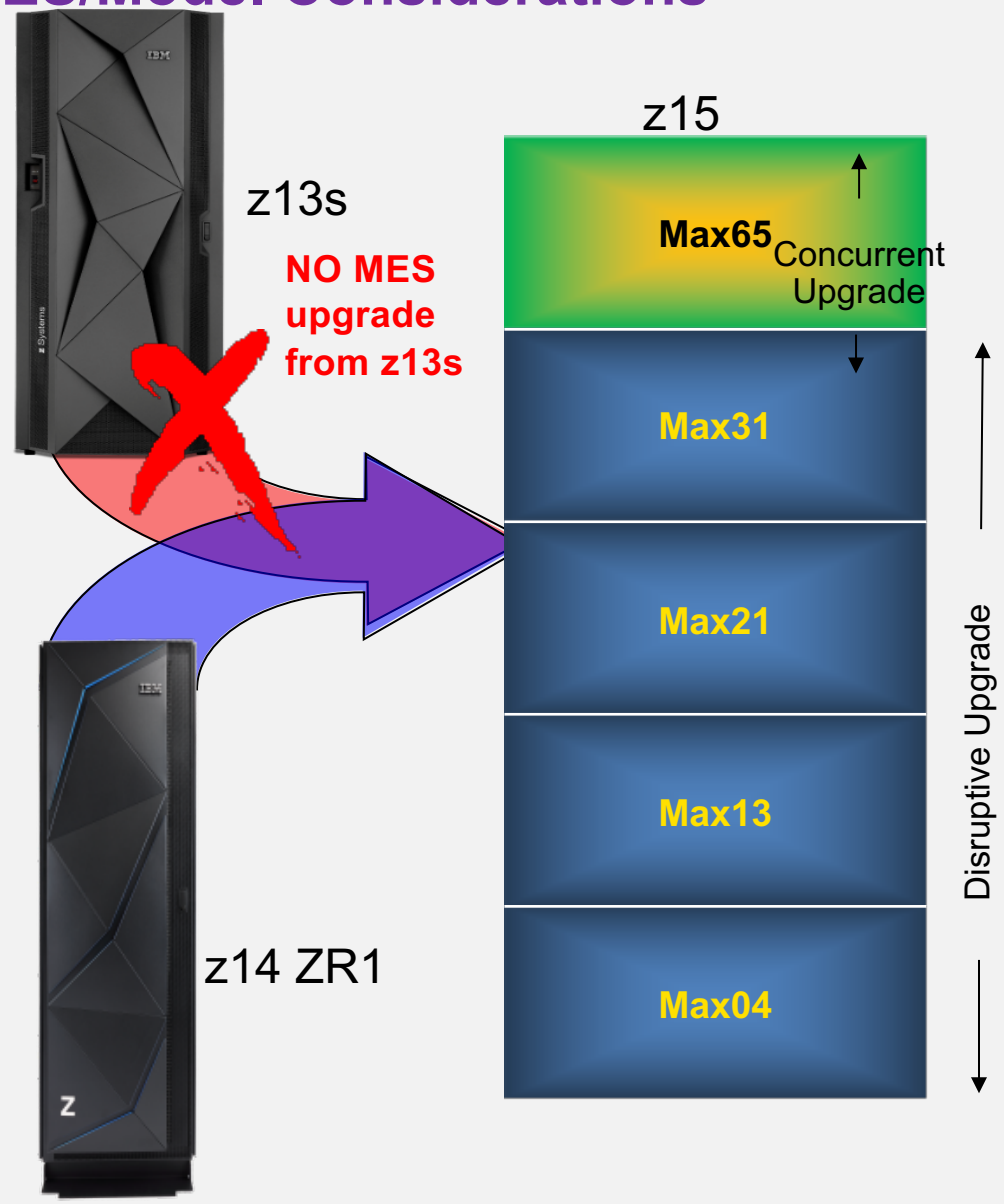

- **z15 T02 to z15 T02 upgrades**
  - Concurrent upgrade from Max31 to Max65
    - Adds a CEC drawer
  - Disruptive upgrade from Max04 – Max31
    - Adds additional PUs and Memory
  - Additional I/O Drawers
    - Based on available space in current frame and available fanouts
- **No frame roll MES from a z13s to a z15 T02**
  - Migration Offering available
- **Any z14 ZR1 to any z15 T02**
- **No Upgrade from z15 T02 to z15 T01**
- **No LinuxONE model conversions to LinuxONE III**
- **LinuxONE III LT2 to z15 T02 MES available**

Integrated system design for z15

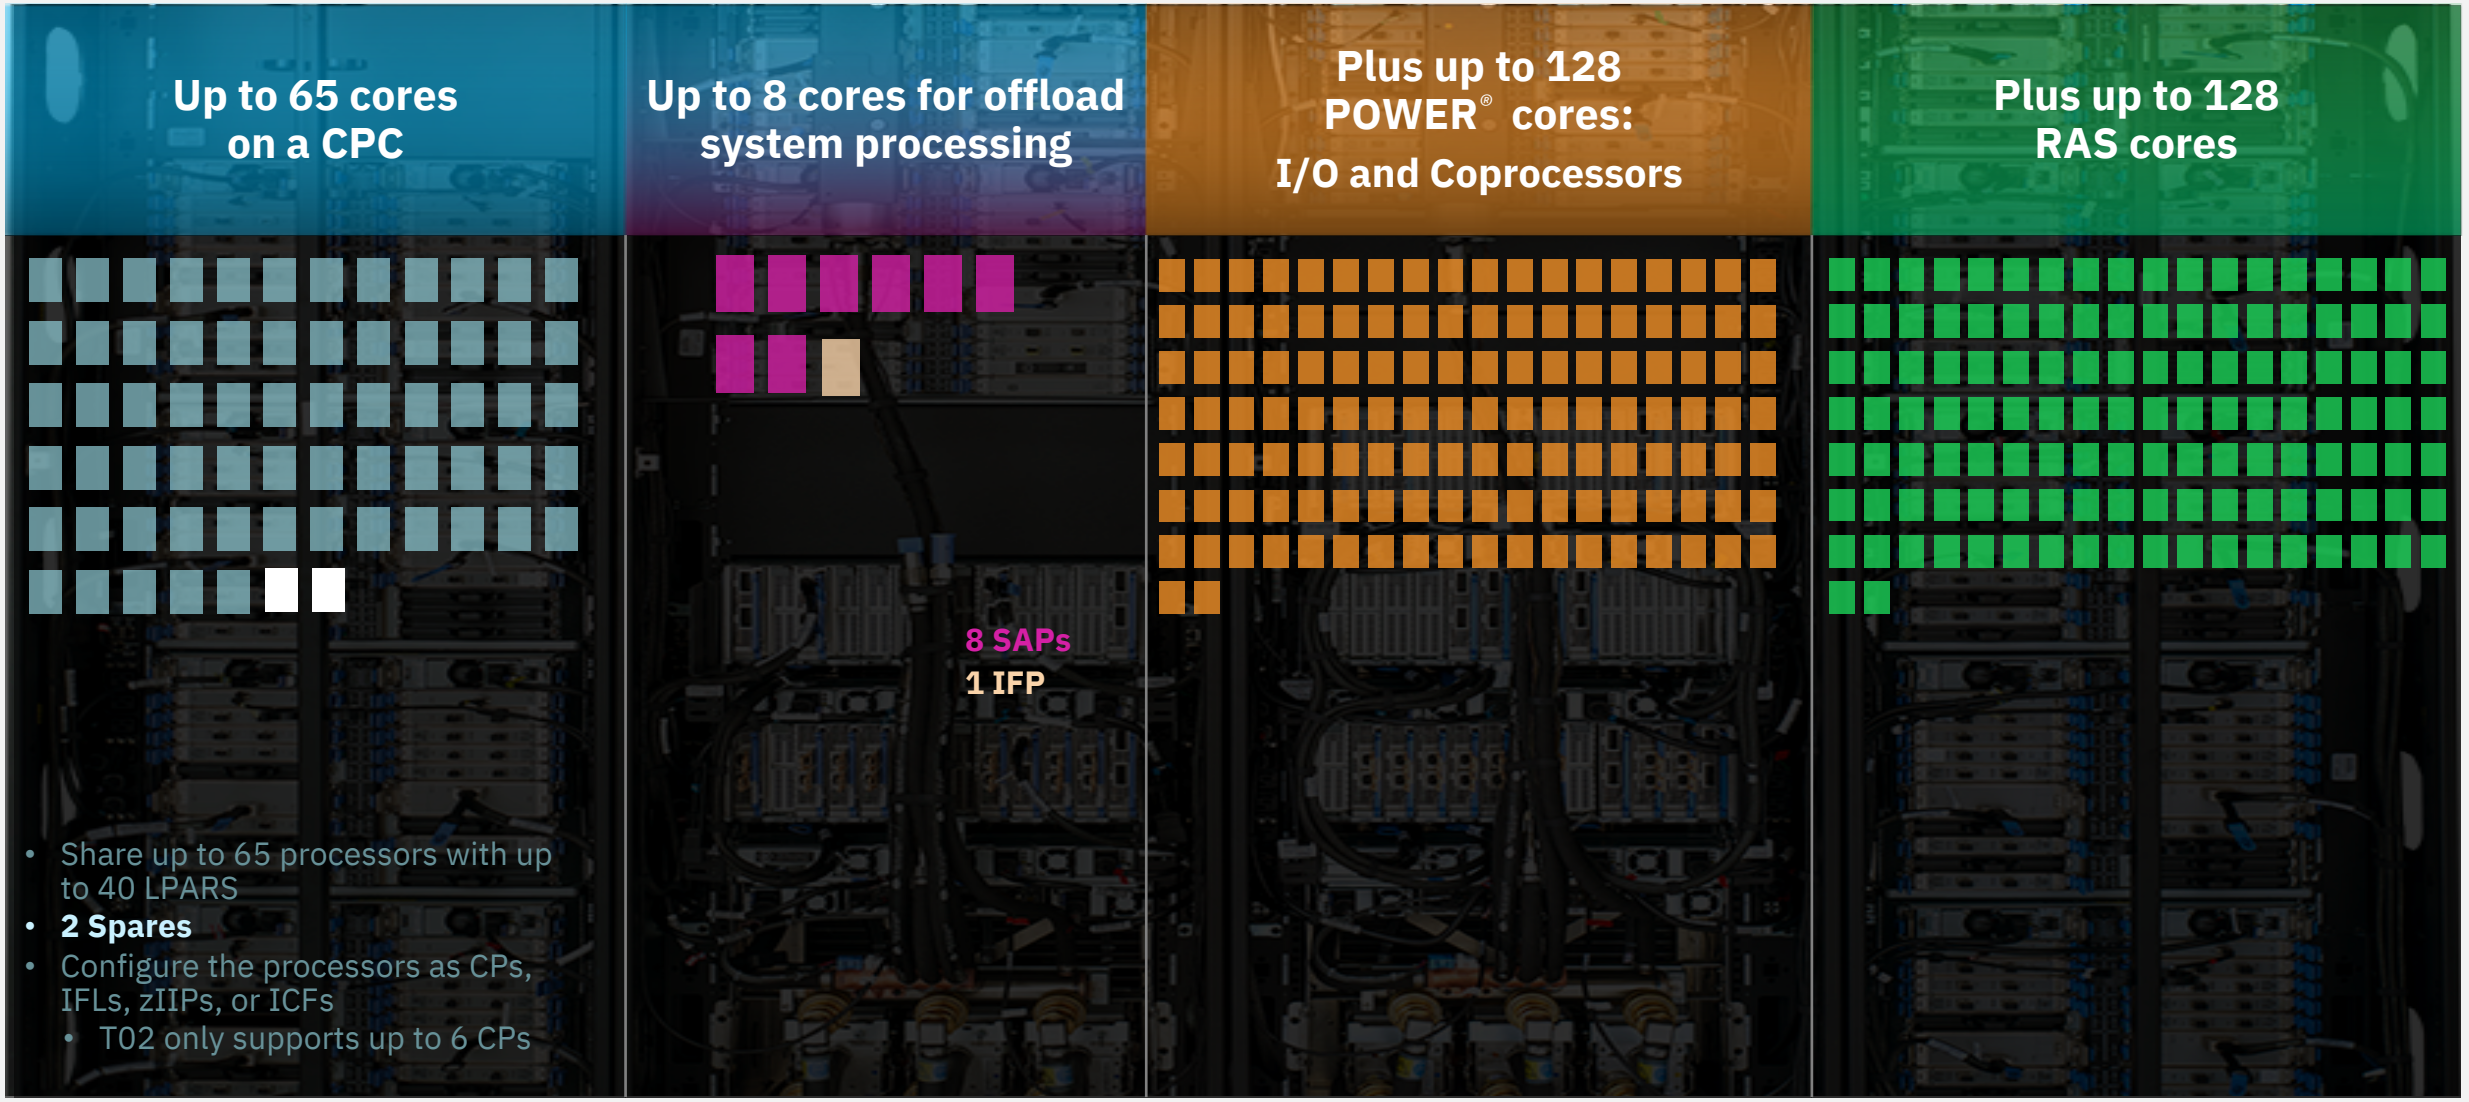

z15 Memory

Memory Considerations

Model	Feature	Min	Max
T02/LT2	Max4	64 GB	2 TB
T02/LT2	Max13	64 GB	4 TB
T02/LT2	Max21	64 GB	4 TB
T02/LT2	Max31	64 GB	8 TB
T02/LT2	Max65	64 GB	16 TB

Concurrent memory upgrades via licensed internal code (LICC) are available at several capacity levels.

- DDR4 Memory DIMMS (32, 64, 128, 256, 512 GB)
- An additional **160 GB** of memory is reserved above the customer purchase amount for the Hardware System Area (HSA).
- An additional 20% of memory is reserved above the customer purchase amount for Redundant Array of Independent Memory (RAIM).

Flexible Memory and Plan Ahead Memory is not offered.

IBM Virtual Flash Memory

- IBM Virtual Flash Memory
  - Replacement for IBM Flash Express I/O features – same use cases.
  - Saves at least two PCIe I/O Drawer Slots from z13s upgrade
  - Less power consumption
  - During z13s upgrade, Feature Conversion for IBM Flash Express
  - z14 ZR1 offered 4 VFM features at .5 TB per feature for a total of 2TB

- Increment Size
  - Up to four features/increments
    - 0.5 TB | 1.0 TB | 1.5 TB | 2.0 TB

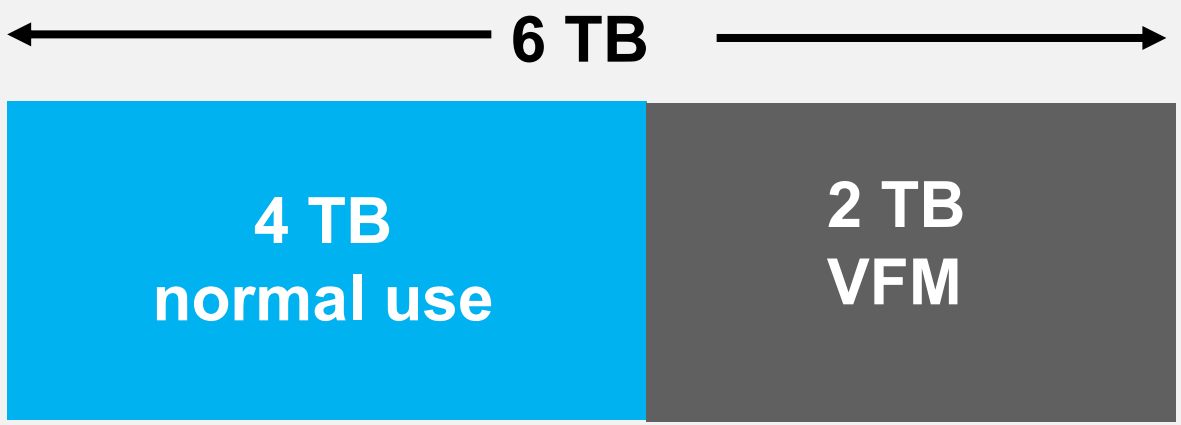

z15 New Features and Functionality

z15 Integrated Accelerator – Design Overview

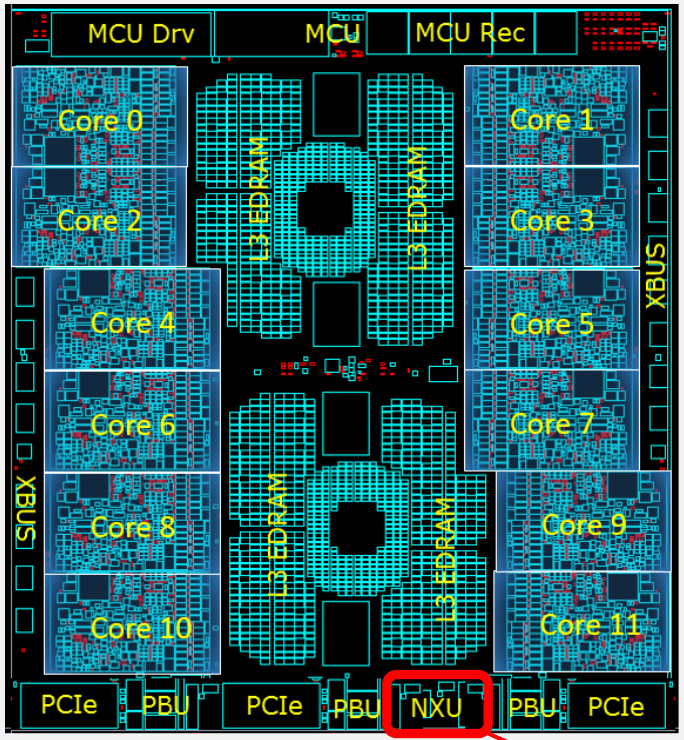

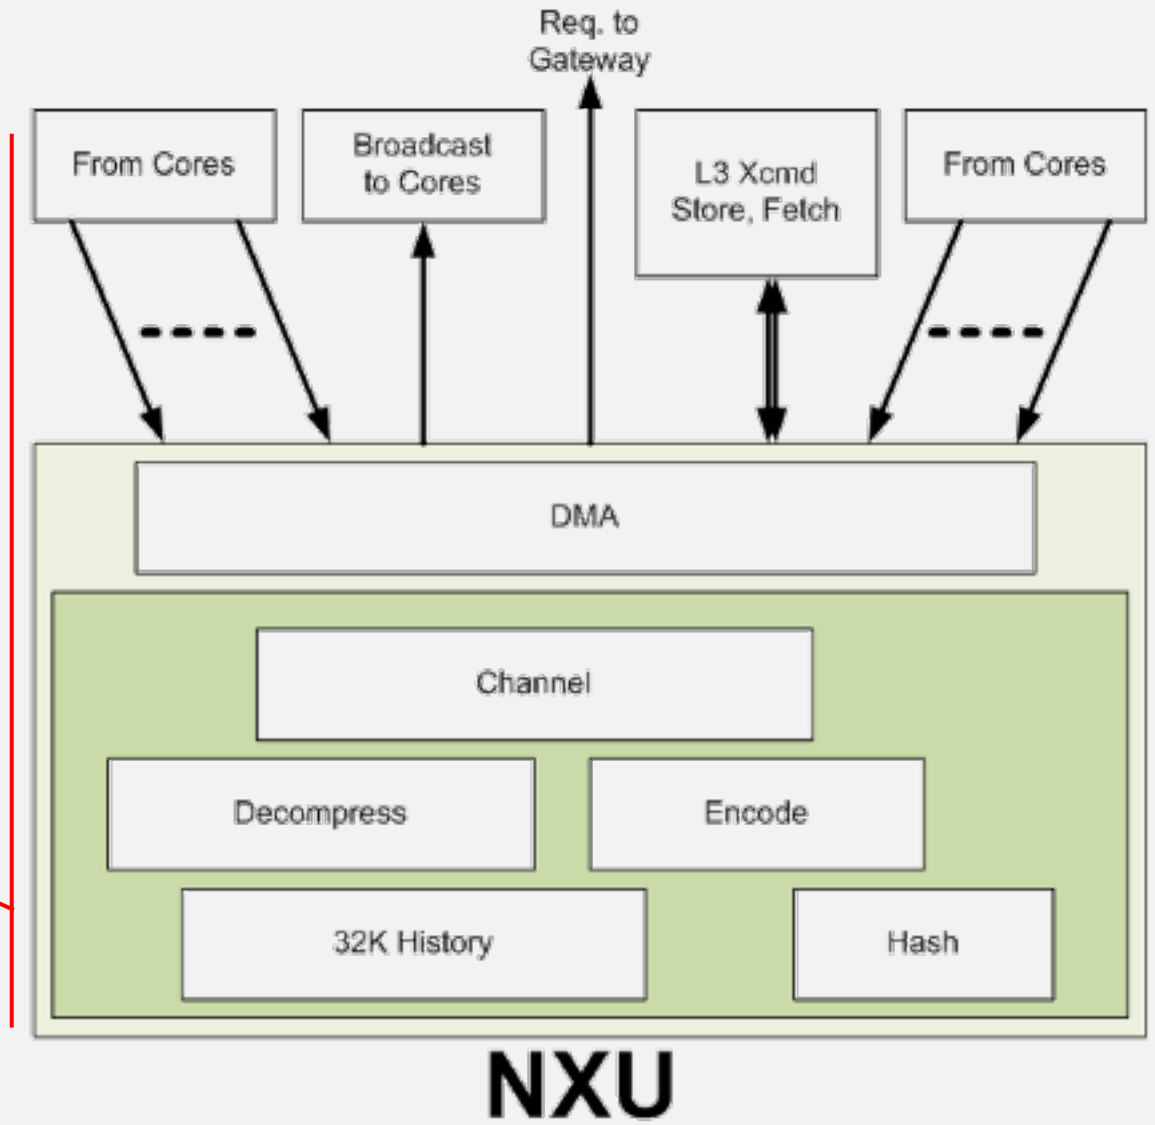

z15 Integrated Accelerator for zEDC

- Compression/Decompression implemented in Nest Accelerator Function
  - Replacement for existing zEDC Express adapter in I/O drawer
  - Nest accelerator unit per processor chip, shared by cores on this chip
  - Supports **DEFLATE compliant compression/decompression + GZIP CRC/ZLIB Adler**

- Brand new concept of sharing and operating an accelerator function in the nest
  - Low Latency
  - High Bandwidth
  - Problem State Execution
  - HW/FW Interlocks to ensure System Responsiveness

- Architected Instruction
  - Executed in Millicode
  - Operating shared HW accelerator on behalf of issuing core

- zEDC Express devices will NOT CARRY FORWARD

Integrated Accelerator for zEDC – z15 T02

	IBM z14 ZR1™ with zEDC Express	z15 T02 with Integrated Accelerator for zEDC
Application elapsed time	Application elapsed time is affected by the time required for the data to be offloaded to and retrieved from the zEDC adapter (PCIe infrastructure in the PCIe I/O Drawer))	Up to 8x faster application elapsed time with no additional CPU time using IBM z15 Integrated Accelerator for zEDC compared to z14 zEDC Express for both compression and decompression.*
Total CEC Throughput	Fully Configured z14 ZR1 – 8 GB/s	Compress up to 95 GB/sec with the Integrated Accelerator for zEDC on the largest IBM z15 T02 **
Virtualization	15 LPARs or VMs per adapter	All LPARs and VMs have 100% access
Capacity Planning	Clients run zEDC cards at 30-50% to handle LPAR consolidation for DR	Enable everything – More than enough throughput
Compatibility	Full compatible with z15 family	Fully compatible with zEDC
Sequential Data Sets	Selectively enabled by application	Enable everything – More than enough throughput
Migration to Tape or VTS	Balanced against data set compression	Enable everything – More than enough throughput
Network Traffic (e.g. Connect:Direct)	Enabled only if enough capacity available	Enable everything
Linux on Z Support	Limited client adoption, virtualization layer adding complexity and affecting throughput	Fully available for Open Source software – NO virtualization employed (on-chip engine) NEW DIFFERENTIATION AGAINST Linux on other platforms

Integrated Accelerator for zEDC support in Linux

Supported with

- Red Hat RHEL 8.1
- SUSE SLES 12 SP5
- Ubuntu 19.10

- On-Chip Compression provides value for existing and new compression users
- Less CPU consumption for compression
- Fully enabled in highly virtualized environments
- All LPARs and VMs (z/VM & KVM) have 100% access
- Avoids having to pick and choose which Linux guests may use accelerator
- Most beneficial for:
- Java (e.g. IBM WebSphere Application Server)
- Backup (e.g. for DBs)
- Network traffic (e.g. Apache, NGINX)
- IBM MQ

Unparalleled Business Continuity and Resiliency for the Digital Age

System Recovery Boost

*For the **ultimate** boost*

***UNLEASH* the power of IBM Z**

Restore normal service and catch up after both planned and unplanned events faster than on any previous Z machine…

- ✓ *Faster IPL and subsystem/middleware restart and recovery*
- ✓ *Faster recovery of production workload*
- ✓ *Boosted capacity to deliver on SLAs catch up on workload faster*
- ✓ *Faster GDPS reconfiguration for automated start-up, shut-down, site switches, etc.*
- ✓ *Easy set-up and management*
- ✓ *Avaliable for every IBM z15*
  - ✓ *Turn subcapacity processors to full capacity per LPAR*
  - ✓ *Allow GP work to be run on currently available zIIPs*

…all without introducing additional IBM software MSU consumption and cost!

Shutdown/Startup/Recovery Boost

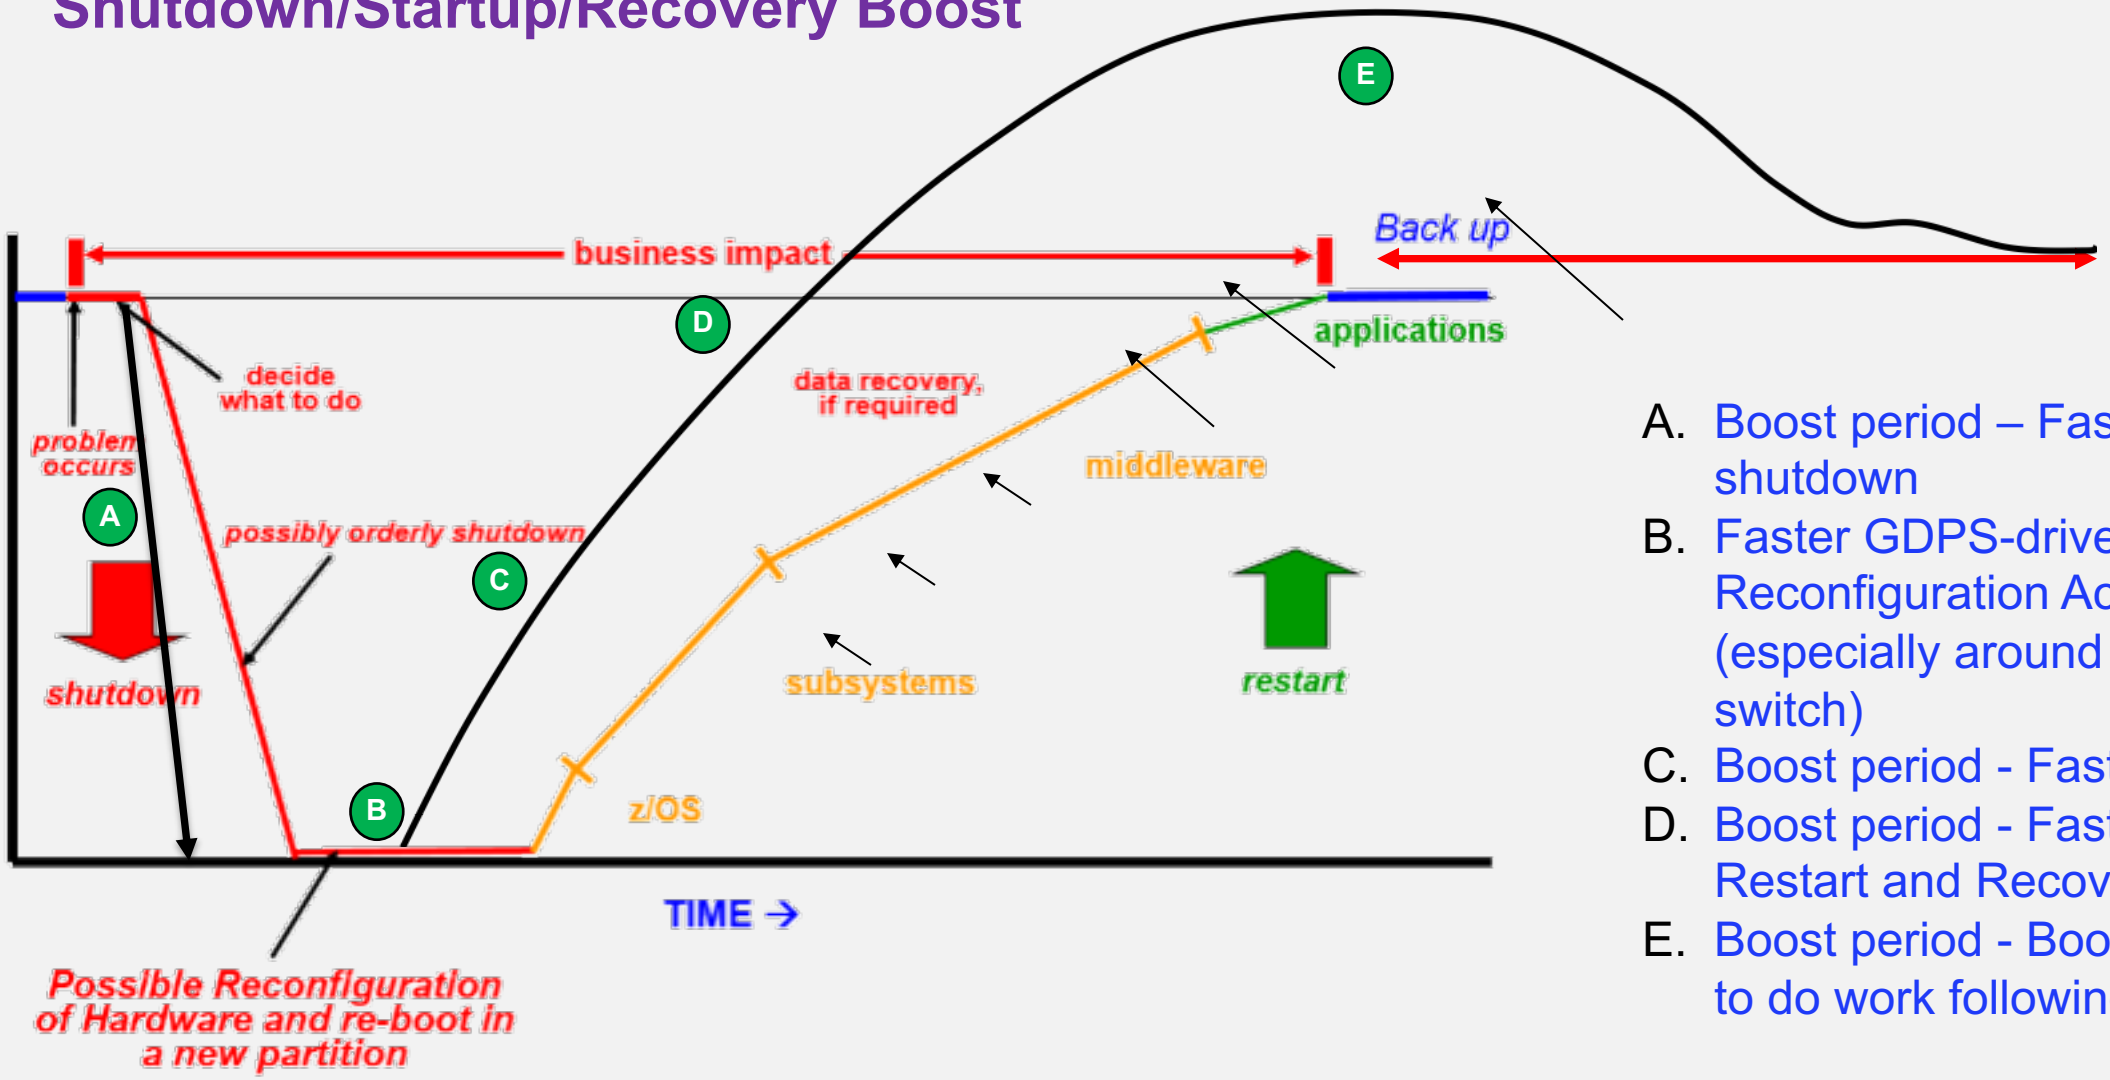

A. Boost period – Faster planned shutdown
B. Faster GDPS-driven HW Reconfiguration Activities (especially around DR and site switch)
C. Boost period - Faster IPL
D. Boost period - Faster Middleware Restart and Recovery
E. Boost period - Boosted Capacity to do work following IPL

No Increase in Client's IBM Software Costs!

Overview

- System Recovery Boost expedites planned system shutdown processing, system IPL, middleware/workload restart and recovery, and the client workload execution that follows, by

1. Processor Capacity Boost using zIIPs
   Using the client's already-entitled GPs and zIIPs

2. Speed Boost
   On sub-capacity machine models, providing a Boost in processor speed by running the GP processors at full-capacity speed, for the Boosting images, during the Boost period.

3. Expedited GDPS Reconfiguration
   Expediting and parallelizing GDPS reconfiguration actions that may be part of the client's restart, reconfiguration, and recovery process

System Recovery Boost Upgrade record not available for the z15 T02

- Boost features can be used together and are enabled on a per LPAR basis
  - Supported operating systems z/OS, z/VM, z/TPF*, and zVSE*
- 30 minutes for shutdown, and 60 minutes for IPL
- All this without increasing the 4HRA IBM software billing cost or MSU consumption costs

* Planned future support

System Recovery Boost

- **BOOST={SYSTEM | ZIIP | SPEED | NONE}**
- SYSTEM: activate whichever boosts the machine and environment accommodates.
- SPEED: Must be on a sub-capacity processor.
- zIIP: HIPERDISPATCH=YES on a non-dedicated partition.
- NONE: neither boost. NO can be used as a synonym of NONE.
- DISPLAY IPLINFO,BOOST can be used to display the value used.

Hardware Management Appliance (HMA)

- **The HMC code runs as an appliance on a higher performance model of the Support Element.**
- **Logon to HMC remotely from your browser.**
- **Same Driver 41 LIC as a physical HMC.**
- **Up to 10 physical HMCs plus the HMA can be ordered and used.**
- **Optional – FC0100**
- **PR/SM or DPM mode**
- **<u>New Build Only – No MES</u>**

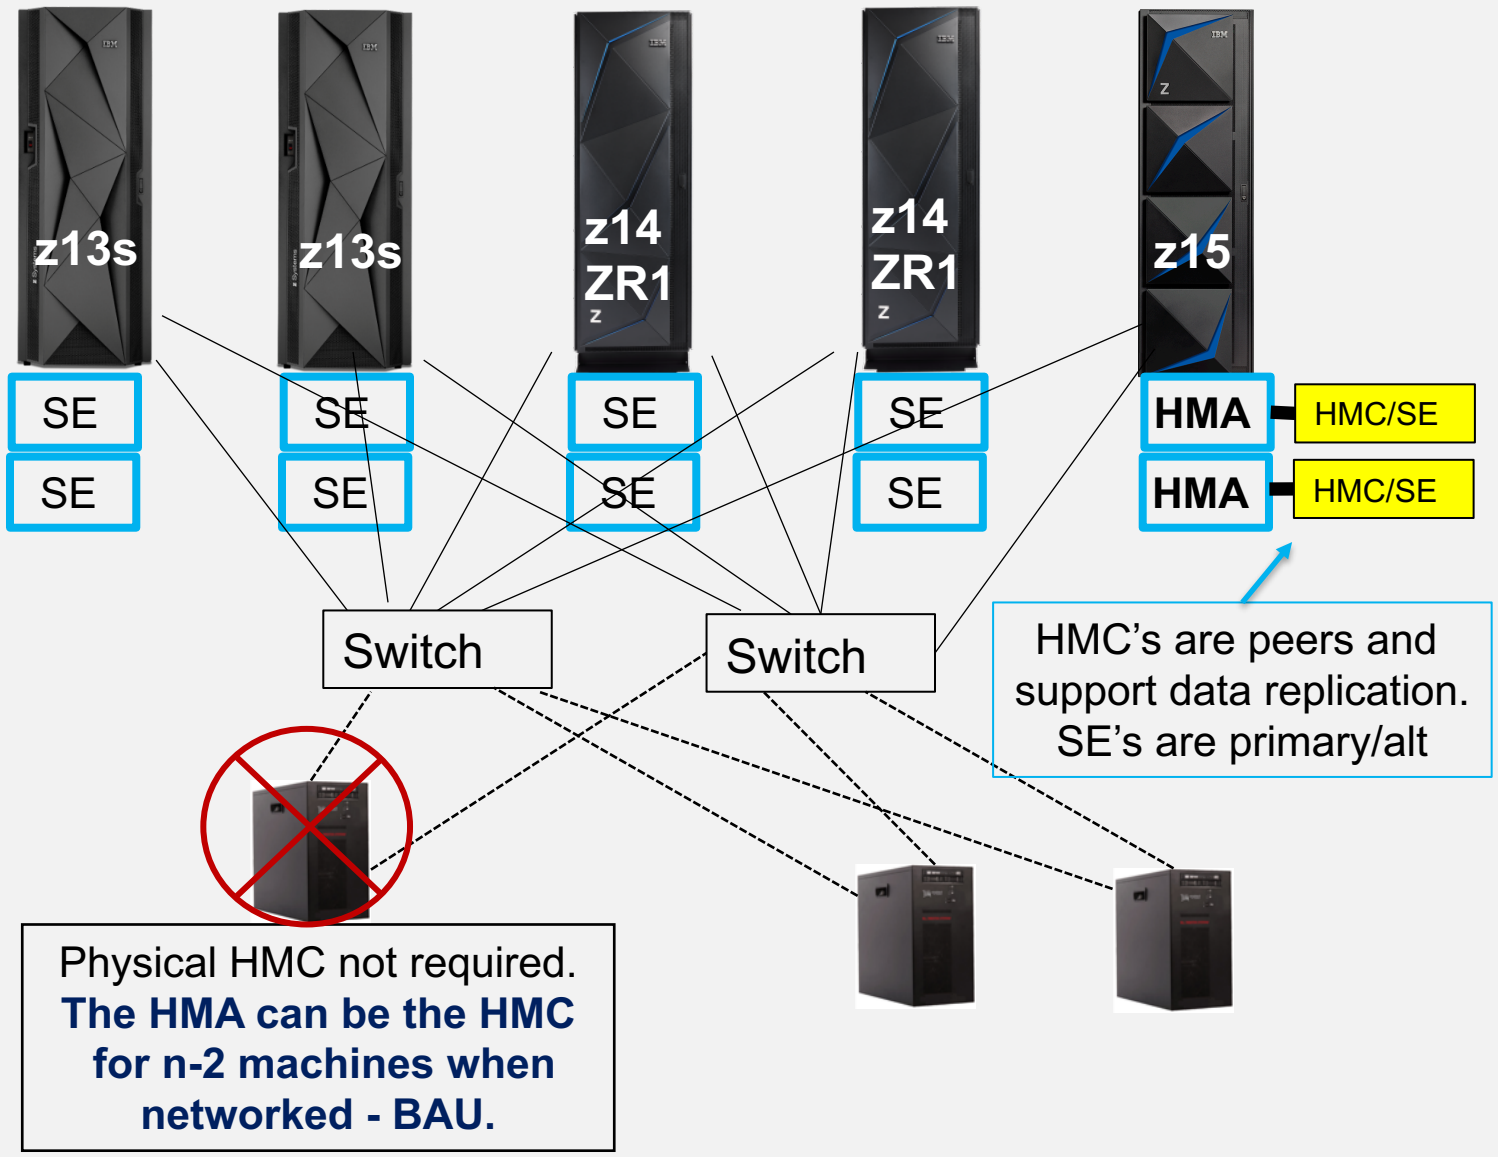

Expanding the z/OS Software Ecosystem

No z/OS skills required

Redbook – SG24-8457

New ***z/OS Container Extensions*** delivers unmodified Linux on Z Docker images running inside z/OS

Challenges

- Porting desired software to z/OS requires effort and presents time-to-value and currency concerns
- Requirement to deploy dependent software hosted on separate Linux servers leads to complicated z/OS operational procedures and hinders the ability to take full advantage of z/OS Qualities of Service.

Client Value

- Modernize z/OS workloads by providing flexibility for development and operations on Z
- Integration with other DevOps tools and Linux applications all in z/OS
- Maintain operational control and extend z/OS Qualities of Service to Linux software
- Make use of existing IT investments by employing Linux within the Z platform
- Requires FC 0104 Cloud Container Foundation

BONUS – workloads run on zCX are zIIP eligible

What is PTP and why it is introduced to STP?

- The PTP Standard has been originally approved in 2002, with update in 2008:
  - Provides more accurate timestamps to connected devices
  - Initially used for Power Distribution Systems, Telecommunications, and Laboratories
  - **Requires Customer Network Infrastructure to be PTP-capable**
  - Accuracy comparison* :
    - NTP – synchronize to within 100 milliseconds
    - NTP with Pulse Per Second – to within 10 microseconds
    - PTP – to sub-microsecond accuracy
- Regulatory requirements for time synchronization (to UTC):
  - Financial Industry Regulations
    - FINRA - 50 milliseconds
    - MiFID II - 100 microseconds
  - Payment Card Industry (PCI) Requirements and Security Assessment Procedures V3.2.1 (May 2018) requires an auditable, tightly synchronized system for credit card companies
- How will z15 use PTP?
  - New External Time Source (like NTP is used today)
  - Use of PTP is optional – customers can continue to use NTP
  - PTP will be provided via the Support Element

LPAR placement improvements

Drawer 1				Drawer 2			
Cluster 1		Cluster 2		Cluster 1		Cluster 2	
C1	C2	C1	C2	C1	C2	C1	C2

- **Chip level optimization**
- **Movement of IFL processors and ICF processors**
  - Isolation of GP+ZIIP, IFL and ICF by chip, to the degree possible to avoid cross sharing of cache
- **In the rare situation where a logical partition's processors span more than one drawer, the memory for the partition will be consolidated into the drawer with the maximum entitlement, if possible (for Dedicated and Hiperdispatch=YES partitions only)**
- **Memory movement within clusters of a drawer**
- **Global Activation**
  - LPAR Act/Deact, Sparing Event
- **Value – Optimal placement of partitions translates to better performance of partitions**

Check placement every 50 seconds

z15 I/O Infrastructure

Up to 40 Logical Partions

- Three Logical Channel Subsystems (LCSS) each with three subchannel sets (SS) and up to 256 channels
- Maximum channel count includes channels spanned to more than one LCSS
- Total physical channels depend on I/O features configured
- Up to 63.75k base IODEVICEs in SS 0 and 64 k alias IODEVICEs in SS 1 and SS 2 per LCSS

8562		
HSA fixed at 160 GB		
LCSS 0	LCSS 1	LCSS 2
Up to 15 Logical Partitions	Up to 15 Logical Partitions	Up to 10 Logical Partitions
Subchannel Sets: SS 0 – 63.75 k SS 1 – 64 k SS 2 – 64 k	Subchannel Sets: SS 0 – 63.75 k SS 1 – 64 k SS 2 – 64 k	Subchannel Sets: SS 0 – 63.75 k SS 1 – 64 k SS 2 – 64 k
Up to 256 Channels	Up to 256 Channels	Up to 256 Channels
Up to 768 channels per System		

PCIe+ I/O Drawer – 16 slots

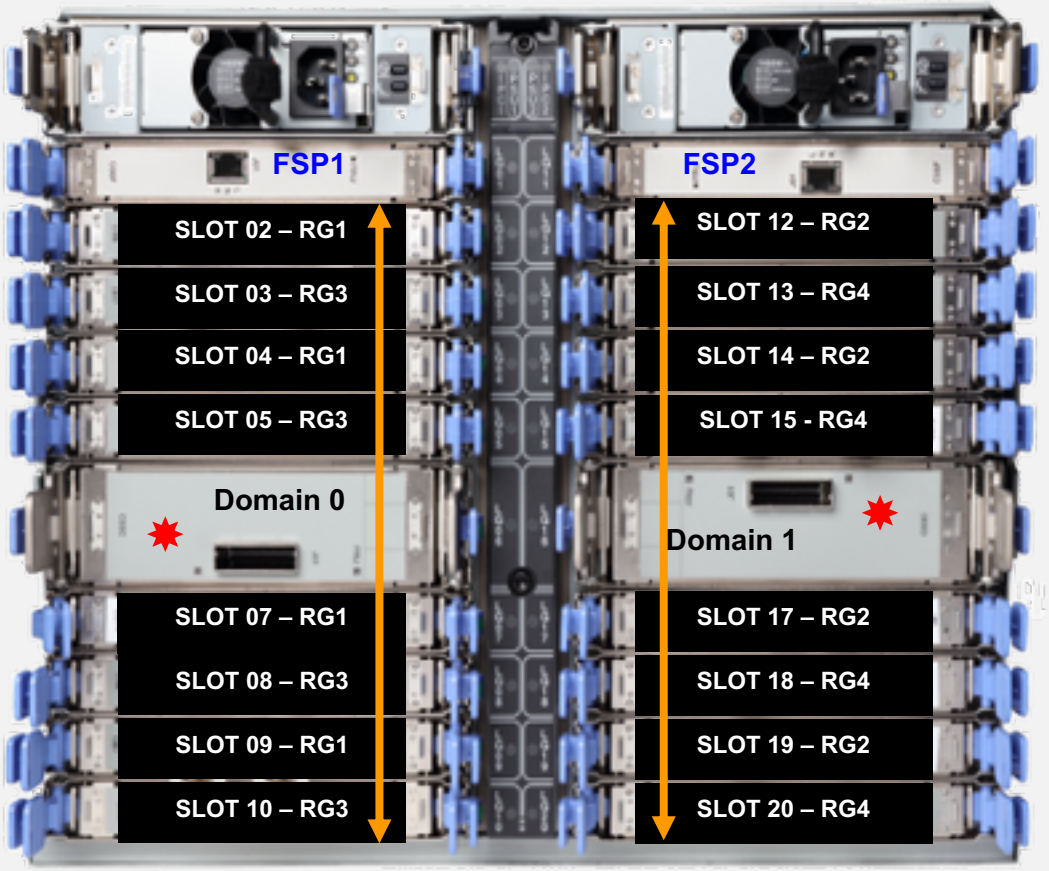

- **Supports PCIe I/O cards**
  - First introduced on the z14 ZR1/Rockhopper II.
- **Supports 16 PCIe I/O cards, horizontal orientation, in two 8-card domains.**
- **Requires two 16 GBps PCIe Interconnect cards (✸), each connected to a 16 GBps PCIe+ Fanout Gen3 to activate both domains.**
- **To support Redundant I/O Interconnect (RII) between domain pairs 0/1 the interconnects to each pair will be from 2 different PCIe+ Fanout Gen3.**
- **Concurrent repair of drawer & concurrent install of all I/O features (hot plug).**

CPC and PCIe+ I/O Drawer Structure and Interconnect

- Max PCIe+ I/O drawer count is power is determined by the CPC feature

CPC Drawer Feature	Max PCIe+ I/O Drawers	Max PCIe I/O Cards
FC 0649 (Max4) FC 0653 (Max65)	3	48
FC 0650 (Max13) FC 0651 (Max21) FC 0652 (Max31)	4	64

- I/O drawer location configuration dependent
  - Placement starts under the first CPC Drawer (z14 ZR1 placed the first I/O drawer above)
  - Second CPC Drawer is placed above the first CPC Drawer
    - The Max65 (the second CPC Drawer) can also be plugged here or FC 2271 reserving space to upgrade to the Max65
- Fanout Management
  - Plugging will be balanced across all CPC Drawers
    - Fanout cable lengths will account for potential future rebalancing.
  - New internal cabling routing system will ensure all fanout cables will be organized and out of way of any component needing service

A new Chip Design

What's Changed:

- IBM has introduced new componentry onto all new build I/O cards on for the z15 to address components at end of life

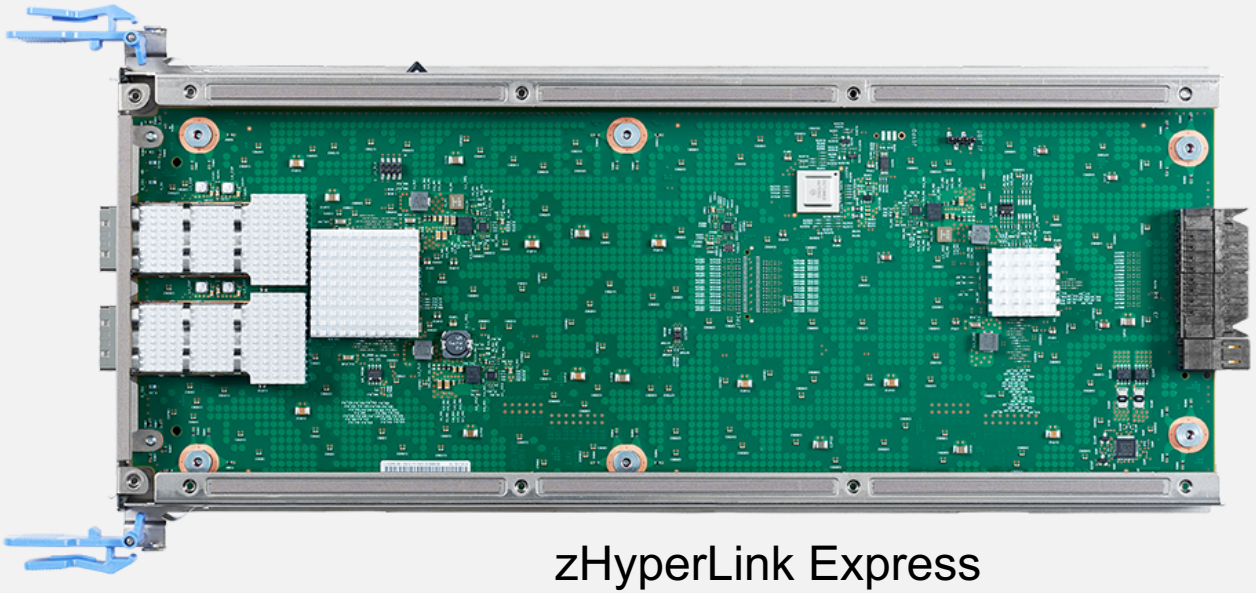

zHyperLink Express

Features Affected:

- All new build I/O Cards
  - Carry forward not affected

Changes to these features

- **No functional change to the card.**
- z15 new build I/O cards will have a different Feature Code from the previous generation.
  - Example
    - zHyperlink Express z14 FC 0431 / z15 FC 0451

z15 I/O Features (new build)

- FICON Express16S+
  - FC 0427, 0428
- zHyperLink Express1.1 
  - FC 0451
- OSA Express6S:
  - GbE (LX, SX) 10GbE (LR, SR) , 1000BASE-T; FC 0422, 0423, 0424, 0425, 0426
- OSA Express7S 25GbE SR;
  - FC 0449
- 10GbE RoCE Express2.1;
  - FC 0432
- 25GbE RoCE Express2.1;
  - FC 0450
- Crypto Express7S;
  - FC 0899, 0898
- Coupling Express LR; 
  - FC 0433
- Integrated Coupling Adapter (ICA) SR1.1; 
  - FC 0176
- IBM Adapter for NVMe1.1; 
  - FC0448
- FCP Express32S; 
  - FC 0438, 0439

 z15 Only

 LinuxONE III Only

No InfiniBand Coupling Links

z15 I/O Features – Carry Forward

- FICON Express16S
  - FC 0417, 0418
- FICON Express8S
  - FC 0409, 0410
- OSA-Express7S 25GbE SR
  - FC 0429
- OSA-Express5S;
  - FC 0413, 0414, 0415, 0416, 0417
- 10GbE RoCE Express;
  - FC 0411
- 10GbE RoCE Express2;
  - FC 0412
- 25GbE RoCE Express2;
  - FC 0430
- zHyperLink Express;
  - FC 0431
- Crypto Express6S;
  - FC 0893
- Crypto Express5S;
  - FC 0890
- Coupling Express LR;
  - FC 0433
- Integrated Coupling Adapter (ICA) SR;
  - FC 0172

There is no MES Carry forward for LinuxONE

FICON Express16S+

- **For FICON, zHPF, and FCP**
  - CHPID types: FC and FCP
  - ***Both ports must be same CHPID type***
    - 2 PCHIDs / CHPIDs
- **Auto-negotiates to 4, 8, or 16 Gbps**
  - 2 Gbps connectivity not supported
  - FICON Express8S will be available for 2Gbps (carry forward only)
- **Increased performance compared to FICON Express16S**
- **Small form factor pluggable (SFP) optics**
  - Concurrent repair/replace action for each SFP
  - 10KM LX - 9 micron single mode fiber
    - Unrepeated distance - 10 kilometers (6.2 miles)
  - SX - 50 or 62.5 micron multimode fiber
    - Distance variable with link data rate and fiber type
- **Two channels of LX or SX (no mix)**

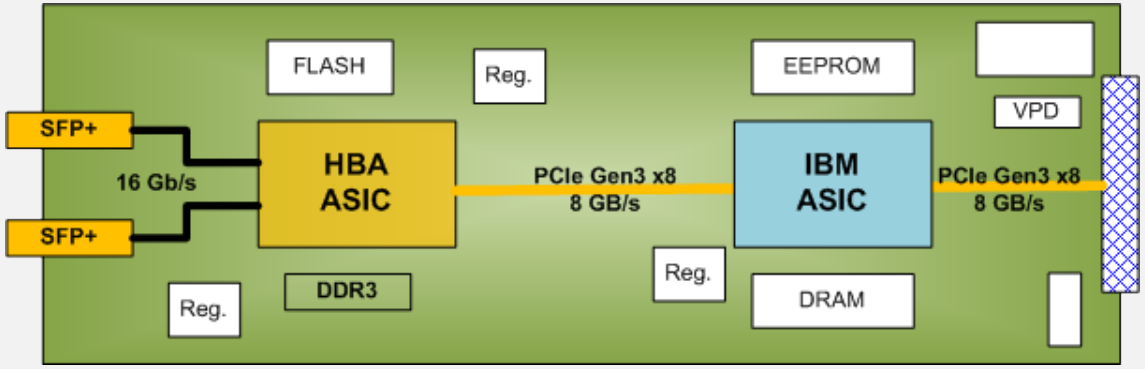

FC 0427 – 10KM LX, FC 0428 – SX

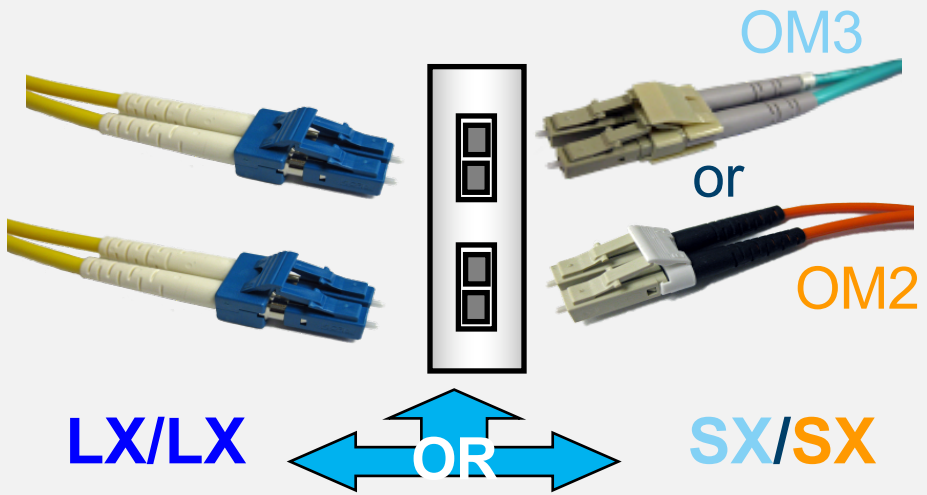

FICON Considerations

- FICON Express16S+ **auto-negotiates to 4, 8 or 16 Gbps**
  - 2 Gbps connectivity not supported for point to point connections
  - 2 Gbps supported through a switch with 8 or 16 Gb optics
- FICON Express16S+ cannot mix the port types

zHyperlink Express1.1 support for z15 T02

- Feature Code 0451 new build, FC 0431 Carry Forward
  - Min/Max, 0/16; Increment = one feature (two ports)
- DS8880 models 984, 985, 986 and 988*, DS8900
- For current support, check IBM Website*

- Db2 V11 and V12 Read Support (Batch and Transactions)
  - Synchronous Db2 database reads – 4K CI/page sizes only
  - For Db2, think of zHyperLink as part of the storage hierarchy in terms of latency that is between the DB2 buffer pools and DS8K cache
- Ability to turn zHyperLinks on and off at the Dataset level
  - Via Storage Class policy or via VARY SMS command
- Db2 Active Log Support
  - Simplex or Synchronous replication (metro mirror, 150 meters)
  - Multi-volume log data sets not supported
- VSAM Read Support (Batch and Transactions)
  - All VSAM record type data sets (i.e. KSDS, ESDS, RRDS, VRRDS) support zHyperLink, along with all types of access (NSR,LSR,GSR,RLS) except NSR Sequential.

Note: Data sets do not have to be extended format

OSA Express7S 25 GbE SR

- 25 Gigabit Ethernet (25 GbE)
  - New build and carry forward
  - CHPID types: OSD
  - Multimode (SR) fiber ONLY
  - One port SR
    - 1 PCHID/CHPID
  - Small form factor pluggable (SFP+) optics
  - LC Duplex
  - 25GbE does NOT auto-negotiate to a slower speed.
  - Up to 48 features per system

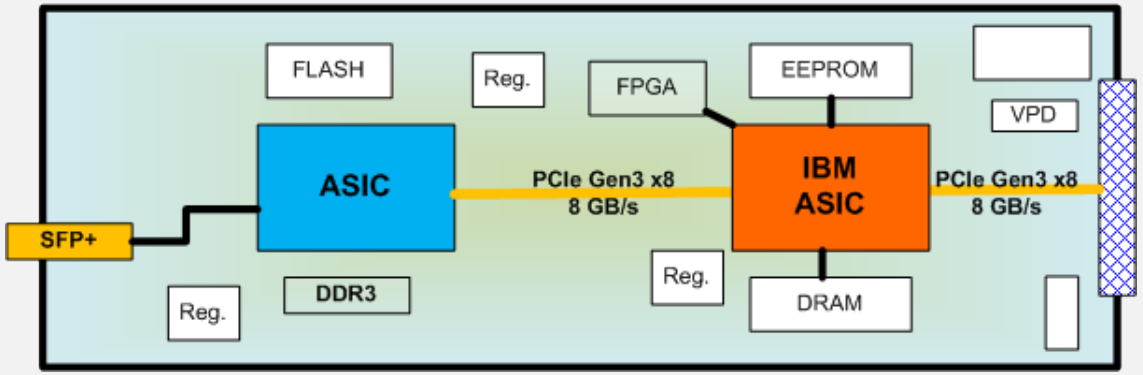

FC #0429 – 25GbE SR

OSA-Express6S Fiber Optic Features

- 10 Gigabit Ethernet (10 GbE)
  - CHPID types: OSD, OSX
  - Single mode (LR) or multimode (SR) fiber
  - One port of LR or one port of SR
    - 1 PCHID/CHPID
  - Small form factor pluggable (SFP+) optics
  - LC duplex
  - OSX CHPID NOT Supported

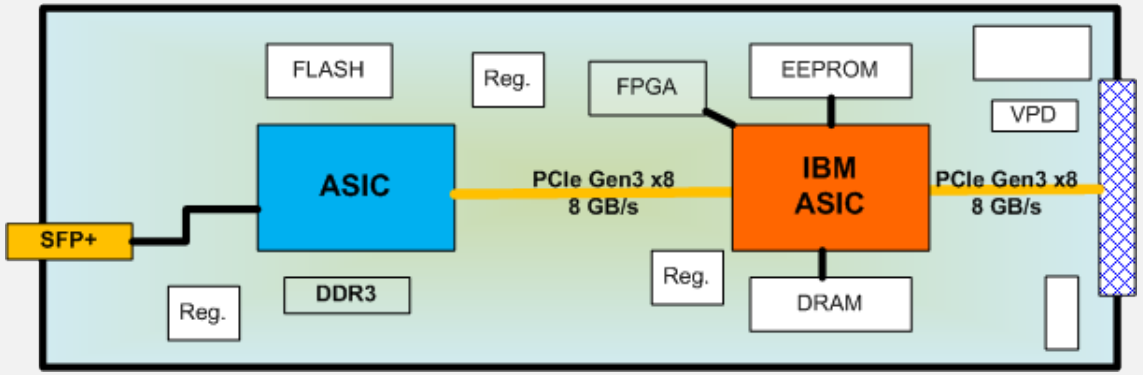

FC 0424 – 10 GbE LR, FC 0425 – 10 GbE SR

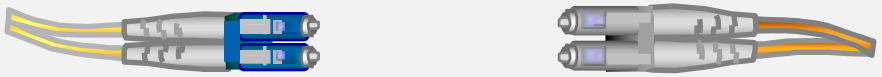

- Gigabit Ethernet (1 GbE)
  - CHPID types: OSD (OSN not supported)
  - Single mode (LX) or multimode (SX) fiber
  - Two ports of LX or two ports of SX
    - 1 PCHID/CHPID
  - Small form factor pluggable (SFP+) optics
    - Concurrent repair/replace action for each SFP
  - LC Duplex

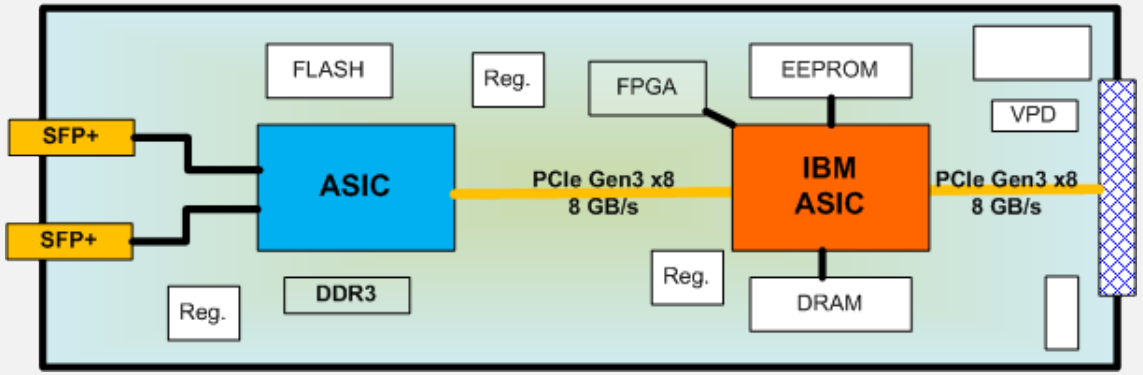

FC 0422 – GbE LX, FC 0423 – GbE SX

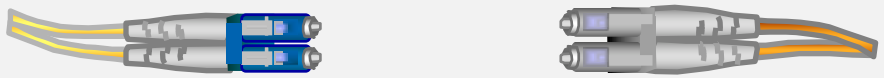

OSA-Express6S 1000BASE-T Ethernet feature

- PCIe form factor feature supported by PCIe I/O drawer
  - One two-port CHPID per feature
  - Half the density of the OSA-Express3 version
- Small form factor pluggable (SFP+) transceivers
  - Concurrent repair/replace action for each SFP
- Exclusively Supports: Auto-negotiation to 100* or 1000 Mbps and full duplex only on Category 5 or better copper
  - No 10Mbps
  - RJ-45 connector
  - Operates at “line speed”

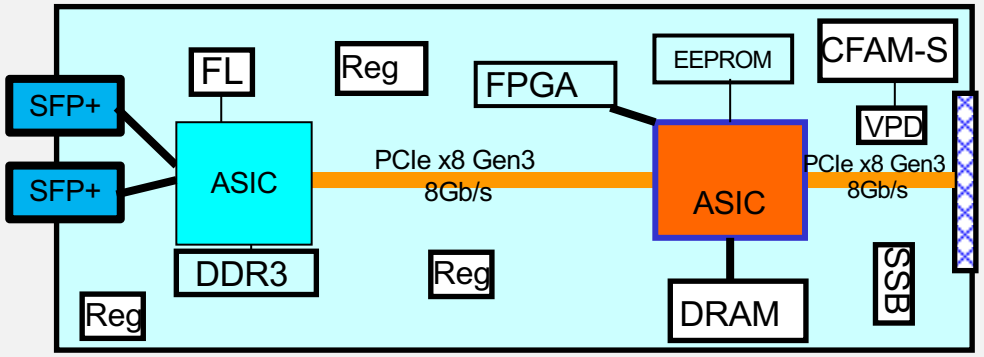

FC 0426

Connector = RJ-45

CHPID TYPE Support:

Operation Mode	TYPE	Description
OSA-ICC	OSC	TN3270E, non-SNA DFT, OS system console operations
QDIO	OSD	TCP/IP traffic when Layer 3, Protocol-independent when Layer 2
Non-QDIO	OSE	TCP/IP and/or SNA/APPN/HPR traffic
DPM Connection	OSM	Connectivity for DPM Mode

* OSA-Express6S 1000BASE-T adapters (#0426) will be the last generation of OSA 1000BASE-T adapters to support connections operating at 100 Mb/second link speed. Future OSA-Express 1000BASE-T adapter generations will support operation only at 1000 Mb/second (1Gb/s) link speed.

10 and 25 GbE RoCE Express2.1

Description	Feature Code	Ports	Max. Features per system (z15 T02)
25GbE RoCE Express2.1	0450	2	**8 (16 ports)**
10GbE RoCE Express2.1	0432	2	**8 (16 ports)**

NEW

Capabilities

- Card electronics update with 25GbE and 10GbE RoCE Express2.1 (compare to FC 0430 and FC 0412)
- Virtualization - 63 Virtual Functions per port (126 VFs per feature)
- Improved RAS - ECC double bit correction

Old 10GbE RoCE Express → FC0411 (2-Ports on z15/z14/z13/z13s, 1-Port on zEC12)

Shared Memory Communications-Remote (SMC-R)
Shared Memory Communications-Direct (SMC-D)

Linux to Linux support z15 and z14 GA2

https://linux-on-z.blogspot.com/p/smc-for-linux-on-ibm-z.html

No application changes required.

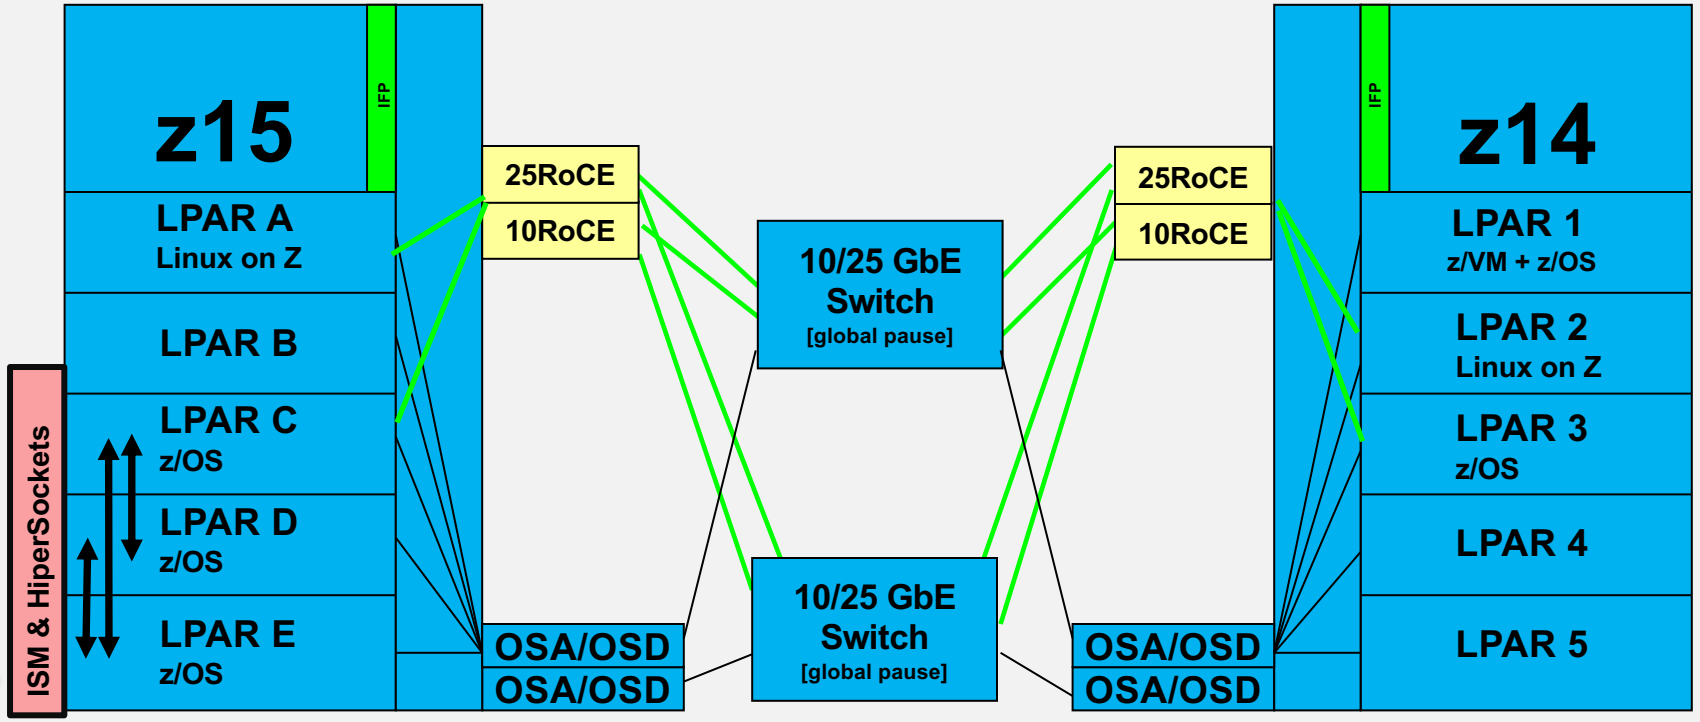

z13 supports 10GbE RoCE Express only.

ISM = Internal Shared Memory via SMC-D
LPAR to LPAR, HiperSockets or OSD connections are required to establish the SMC-D communications.
No additional hardware purchase required.

Central Processor Assist for Cryptographic Function (CPACF)

- Feature Code 3863, CFACF enablement
  No Charge
- Value = Lower latency for encryption operations & better performance
- Hardware accelerated encryption on every core designed to provide faster encryption and decryption than previous servers.
- **New - Elliptic Curve Cryptography clear key support in CPACF. No application changes.**
  - Value = better ECC performance & throughput.
- Support for new Algorithms
  - *EdDSA (Ed448, Ed25519), ECDSA (P-256, P-384, P-521), ECDH(P-256, P-384, P521, X25519, X448)*
  - *Support for protected key signature creation*

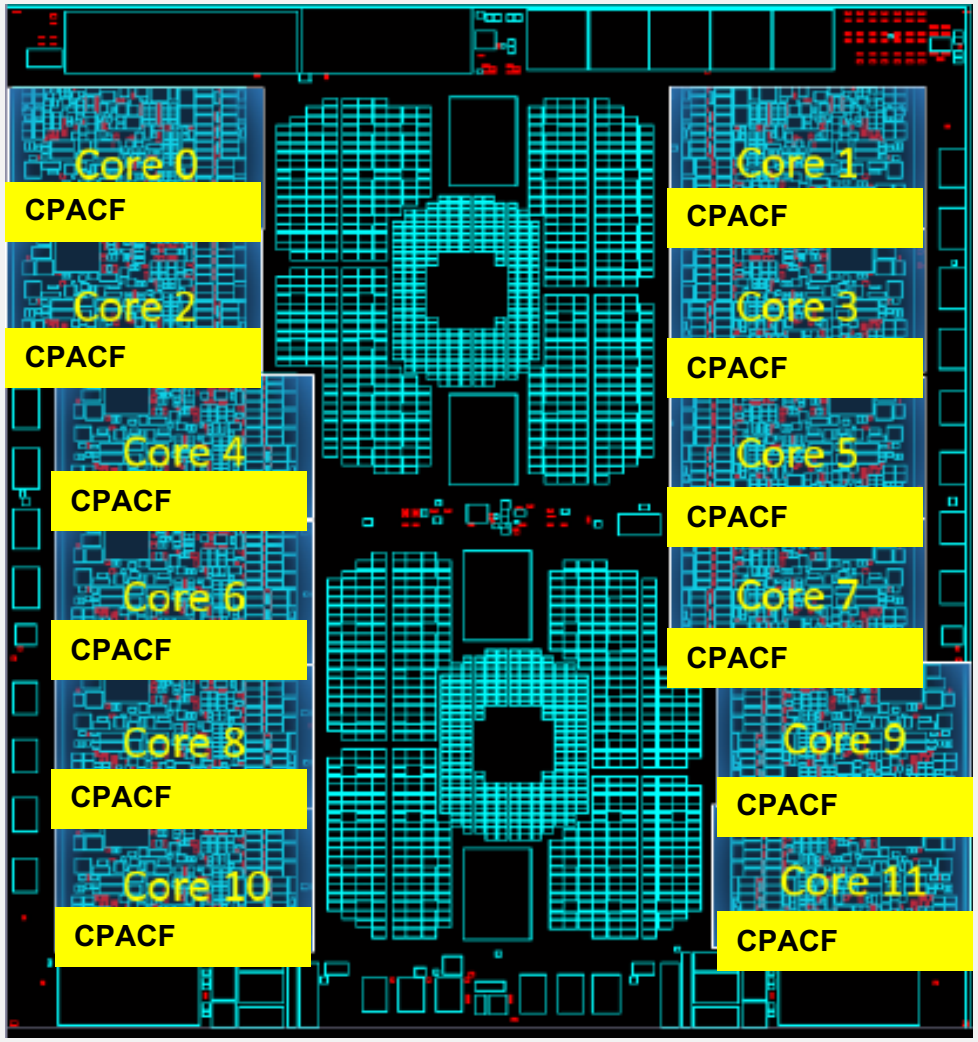

Crypto Express7S

Two new cards designed for z15

- FC 0899 Single HSM
  - Max 16 per server
- FC 0898 Dual HSM
  - Max 20 per server

A mix of Crypto cards can be ordered for both new build and carry forward

- Max combined total: 40 co-processors

New design and format driven by the adoption of blockchain and other highly secure applications

- Designed for 2X performance improvement
- Support for new Algorithms
  - *SHA3, SHA3* XOF modes, FFX, VAES3, BPS
  - **Dilithium (Quantum Safe)**

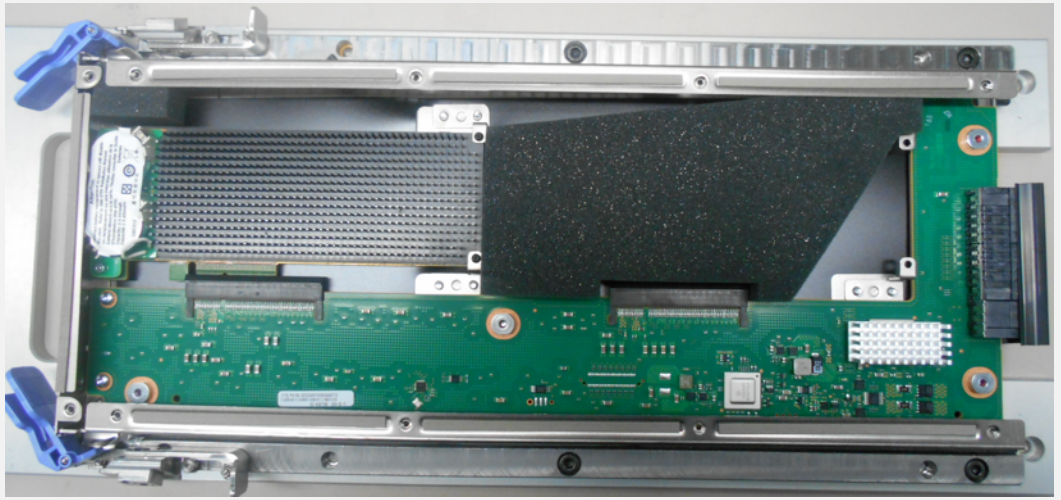

FC 0899

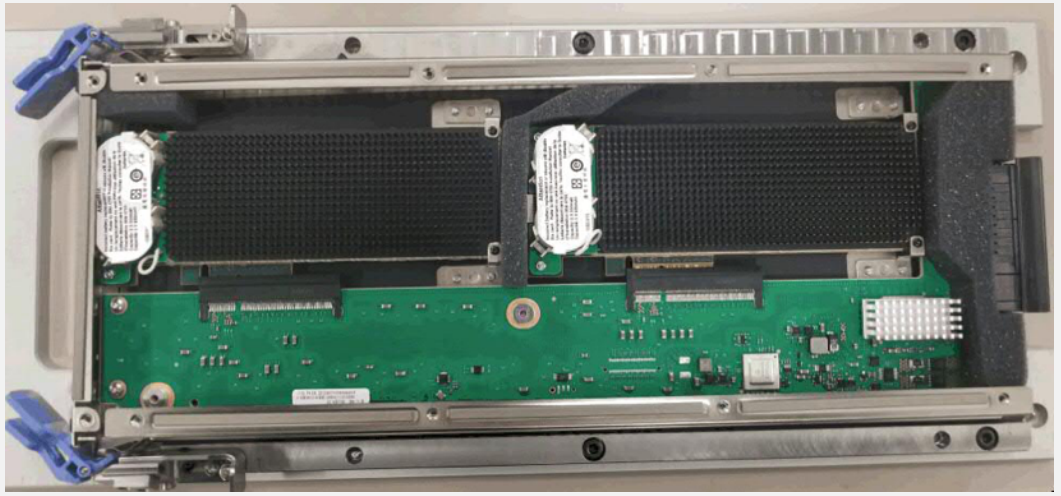

FC 0898

Crypto Express7S co-Processor

- **Half the length and height of the PCIe Standard (approx. 23mm x 23mm)**
- **Double number of public key cryptographic engines**
- **Double number of processors (PPC)**
- **Preprocessing and functionality offloading from main processor**
- **Embedded True Random Number Generator**
- **Designed to be FIPS 140-2 Level 4 compliant**
- **EP11 can now run with Protected Keys**

Three Crypto Express7S configuration options

- Only one configuration option can be chosen at any given time
- Switching between configuration modes will erase all card secrets
  - Exception: Switching from CCA to accelerator or vice versa

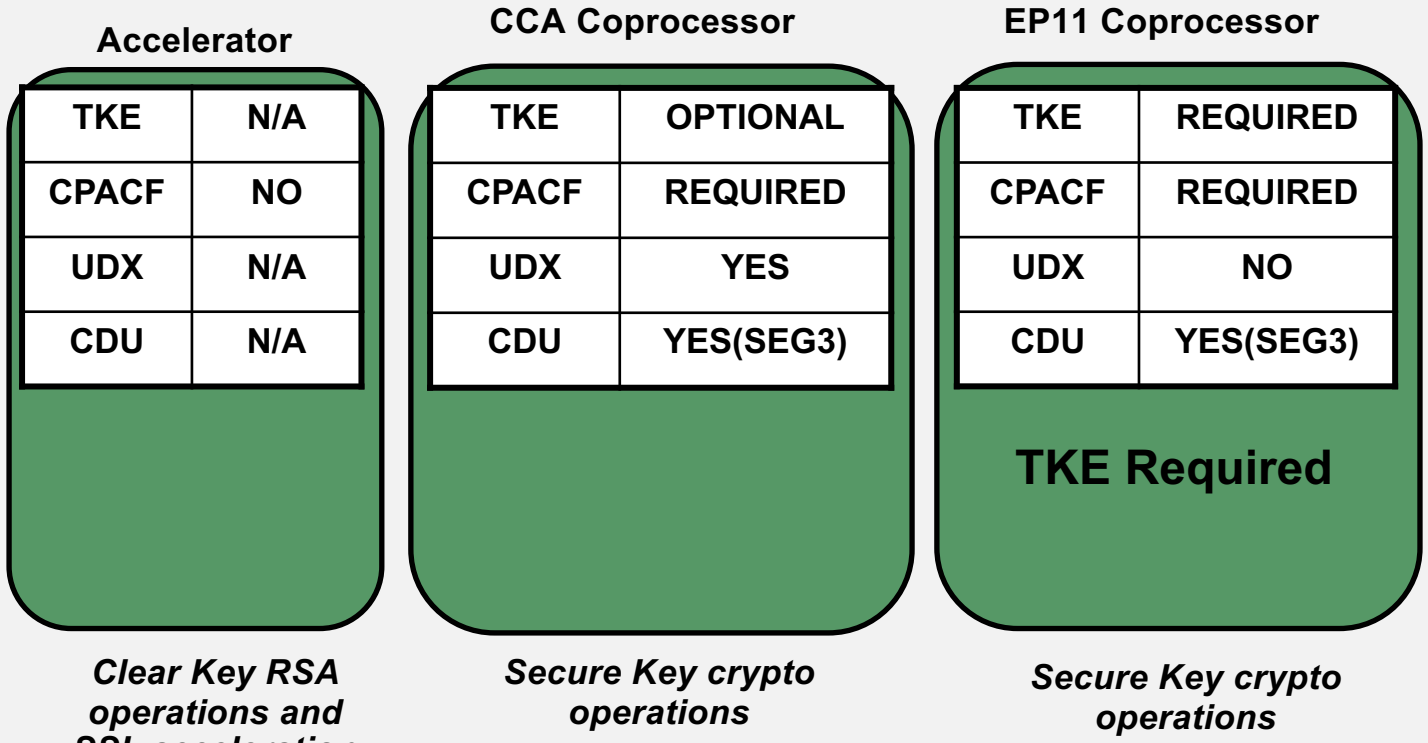

Trusted Key Entry and TKE LIC

- Although customers can carry forward their TKEs on IBM z15, new capabilities will only be delivered on newly manufactured TKEs.
- The IBM z15 environment can contain both TKEs that have been carried forward and newly manufactured TKEs.
- TKE 9.2 LIC is required if you choose to use the TKE to manage a Crypto Express7S.
- CCA in PCI-HSM mode and EP11 also require a smartcard reader plus FIPS certified smart cards.
- Smart card readers: Feature 0891
  - (2 Identiv readers and 20 00RY790 smart cards)
- Smart cards: FC0900
  - Smart card readers from FC0885 or FC0891 can be carried forward to any TKE 9.2 workstation.

Description	FC	z15 with LIC 9.2
TKE 9.2 LIC	0881	Crypto Express5S Crypto Express6S Crypto Express7S
Workstation	0087 Rack 0085 Rack	Yes
Workstation	0088 Tower 0086 Tower	Yes
Workstation	0841	No
Workstation	0842 Tower	No
Workstation	0847 Tower	No
Workstation	0849 Rack	No
Workstation	0097 Rack	Yes w/4768 (FC0844)
Workstation	0098 Tower	Yes w/4768 (FC0844)
Workstation	0080 Rack	Yes w/4768 (FC0844)
Workstation	0081 Tower	Yes w/4768 (FC0844)

LinuxOne Only

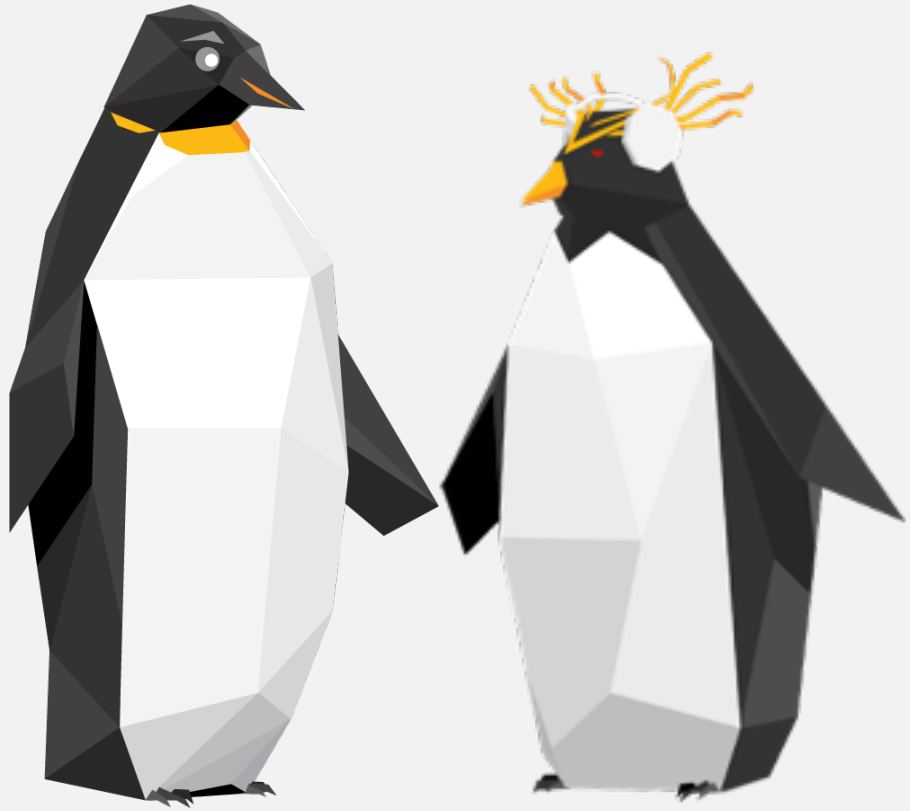

FCP Express32S

- For FCP Only
  - CHPID types: FCP
    - 2 PCHIDs/CHPIDs
- Auto-negotiates to 8, 16, or 32 Gbps
- Concurrent repair/replace of SFP optics
  - 10KM LX - 9 micron single mode fiber
    - Unrepeated distance - 10 kilometers (6.2 miles)
    - Receiving device must also be LX
  - SX - 50 or 62.5 micron multimode fiber
    - Distance variable with link data rate and fiber type
    - Receiving device must also be SX
- Maximum of 32 two port features

LinuxONE Only

IBM Adapter for NVMe1.1 – FC0448

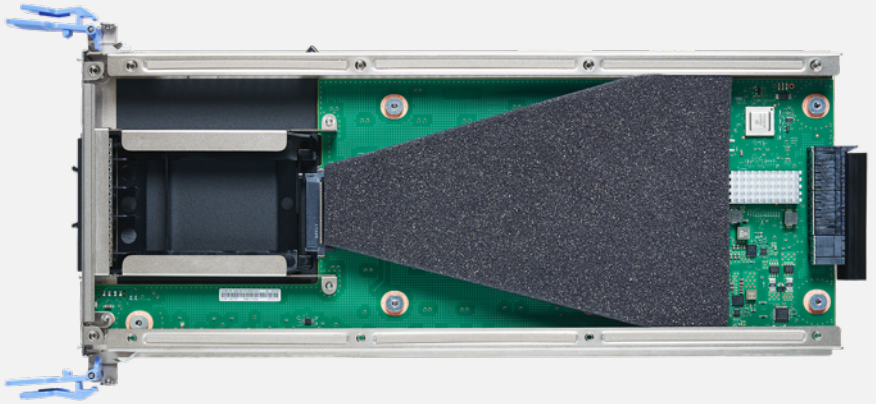

- "Built in" storage.
- In plan:
  - Dynamic Partition Manager (DPM) support for data volumes
  - Multiple NVMe's per LPAR
- Uses the normal LinuxONE PCIe EC Stream.
- Carrier Card
  - Zero ports
  - IBM provides a carrier card into which NVMe SSDs can be plugged.
  - IBM service will install the vendor SSD concurrently into the carrier card on-site. Hot/cold plug.

- Up to 16 features in increments of one.
- The vendor SSD card will be purchased by the client from a reseller or directly from the vendor. Verified in IBM Z
  - Intel PN SSDPE2KX010T701 (1TB) - Up to 16 TB
  - Intel PN SSDPE2KX040T701 (4TB) - Up to 64 TB
  - OTHER VENDORS BEING VERIFIED.
    - All can coexist on the same system and same I/O Drawer.

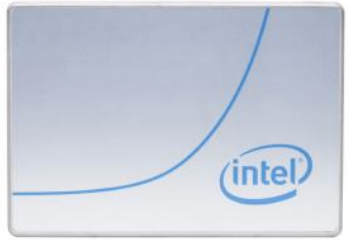

- Details can be found in the IMPP GC28-7011. Performance testing found here

Features	New Build (NB) Carry Forward (CF)	Maximum Features	Ports	Increment
FICON Express16S+	NB	64	2	1
FICON Express16S	CF	64	2	1
FICON Express8S	CF	64	2	1
FCP Express32S (LinuxONE) 	NB	64	2	1
OSA Express7S 25GbE SR1.1	NB/CF	48	2	1
OSA Express6S 1GbE / 10GbE	CF	48	2 / 1	1
OSA Express5S	CF	48	2	1
Crypto Express7S (2 HW Security Modules (HSMs))	NB	20	2 HSMs	2, 1 thereafter
Crypto Express7S (1 HSM)	NB	16	1 HSM	2, 1 thereafter
Crypto Express6S (1 HSM)	CF	16	1 HSM	2, 1 thereafter
Crypto Express5S (1 HSM)	CF	16	1 HSM	2, 1 thereafter
25GbE RoCE Express2.1	NB	8	2	1
10GbE RoCE Express2.1	NB	8	2	1
10GbE RoCE Express	CF	8	2	1
10GbE RoCE Express2	CF	8	2	1
25GbE RoCE Express2	CF	8	2	1
NVMe Express (LinuxONE) 	NB	16	n/a	1
zHyperlink Express1.1 	NB	16	2	1
zHyperlink Express	CF	16	2	1
Virtual Flash Memory	NB	4 Units	0.5 TB	1
Integrated Coupling SR1.1 	NB	**	2	1
Integrated Coupling SR	CF	**	2	1
Coupling Express2 LR	NB/CF	32	2	1

** ICA SR Features: MAX4=2, MAX13=4, MAX21=6, MAX31=8, MAX65=16

z15 Coupling

Parallel Sysplex and Server Time Protocol

- z15 CFCC Level 24
- The z15 can be connected with coupling links to the following servers with these MCL requirements.

- z14 (3906/3907) at Driver 36
  - CFCC Product Release 23 – Service Level 0.13
  - Bundle S13 / MCL P41419.003 (February 2019)

- z13/z13s (2964/2965) at Driver 27
  - CFCC Product Release 21 – Service Level 2.20
  - Bundle S82 / MCL P08416.008 (February 2019)

- Enhancements to coupling communication up to 4 virtual paths (CHPIDs) per link.
  - Max ICA SR1.1 increased from 8 to 24
  - Max coupling CHPIDs per system increased to 384
    - Currently 256 on the z14 ZR1, 50% more
    - Max 128 per LPAR same as z14 ZR1
  - Max ICP CHPIDs increased to 64
    - Currently 32 on the z14, 2X more

Integrated HCD for S/A Coupling Facility

Eliminate S/A CF planned outages for I/O Reconfigurations

The client needs a controlling HCD Instance, which is the User Interface

- **z14 GA2 or higher**
- **z/OS 2.1, 2.2, 2.3 with APARS:**
  - **OA53952 (IOS)**
  - **OA54912 (HCD)**
  - **AA55404 (IOCP)**
  - **IO25603 (HCM)**

Value = Keep the business running!
Easier migrations with no CF outage.

On z14 (GA2)**, enable functionality by adding the MCS_1 LPAR to the target system's configuration.**

Functionality & MCS_1 LPAR is already present in the z15.

z15 Coupling Links

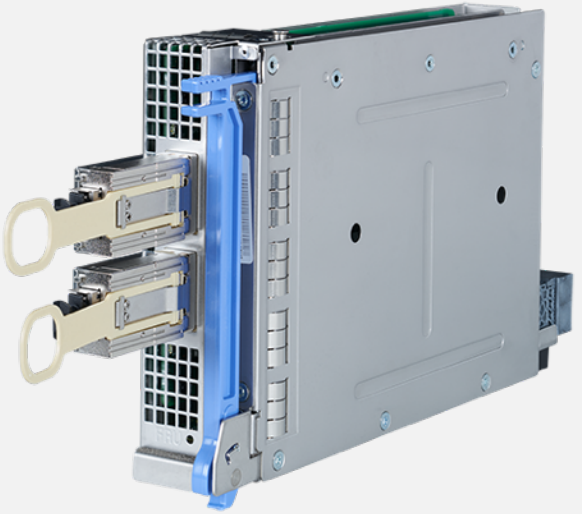

IBM Integrated Coupling Adapter SR1.1 (ICA SR1.1)

- *Coupling Connectivity into the Future (Short Distance)*
- Coupling CHPID CS5, Performance similar to Coupling over InfiniBand 12X IFB3 protocol
- PCIe+ Gen3, Fanout in the CPC drawer, 2-ports per fanout, 150m;
- z13 GA1 availability

Coupling Express2 LR (CE LR)

- *Coupling Connectivity into the Future (Long Distance)*
- Coupling CHPID CL5, Performance similar to Coupling over InfiniBand 1x
- PCIe+ I/O drawer required for CL5 adapter
- Adapter (2-port card): same adapter as 10GbE RoCE Express but with Coupling Optics and Firmware
- 10 Gbps, Up to 4 CHPIDs per port, 32 buffers (i.e. 32 subchannels) per CHPID
- Distance: 10 KM Unrepeated; up to 100 KM with qualified DWDM
- Point-to-Point
- Retrofitted on z13 GA2

The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

z15 Coupling Considerations

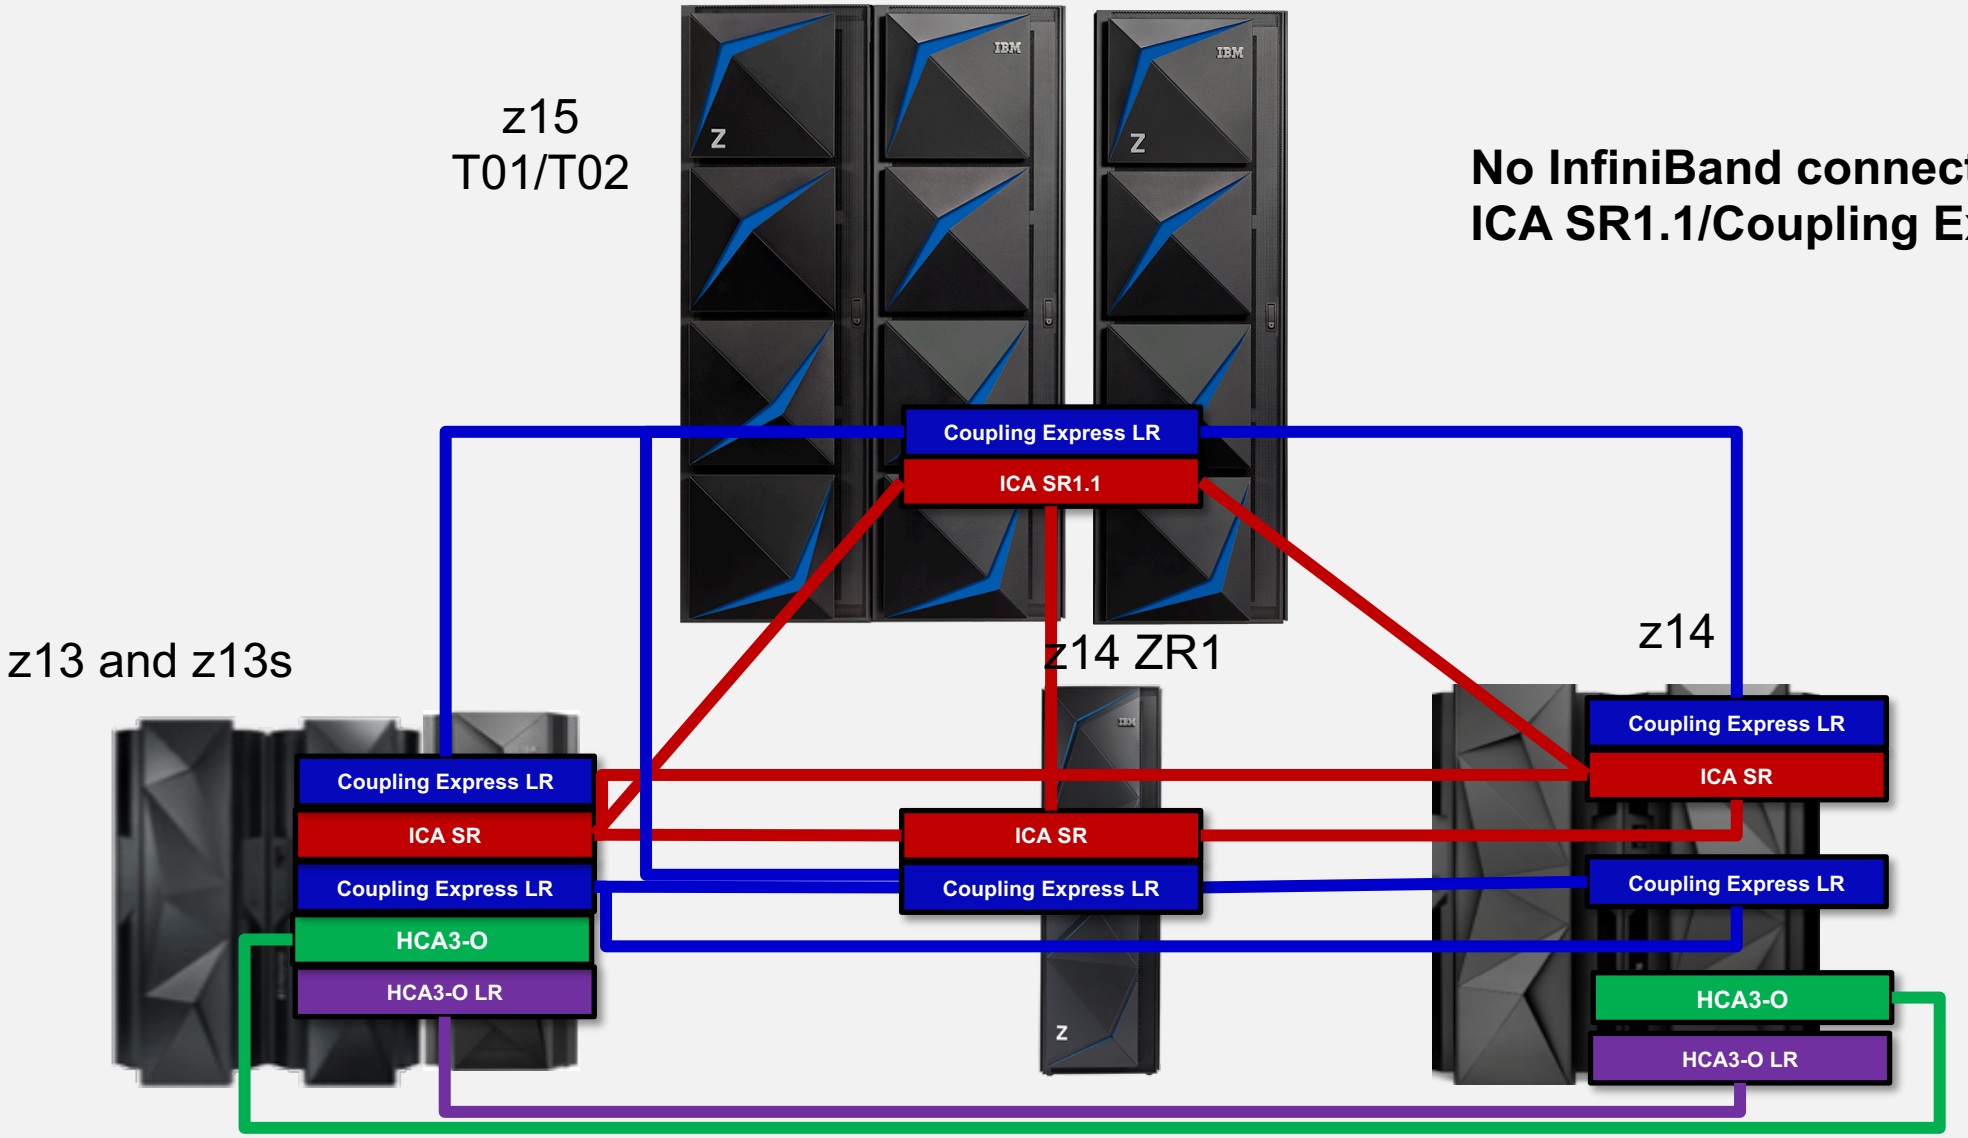

No InfiniBand connections
ICA SR1.1/Coupling Express LR Only

HCA3-O / HCA3-O LR Only communicate between the z13/z13s/z14 (M0x Models)

z15 Coupling Migration without HCA3-O connections

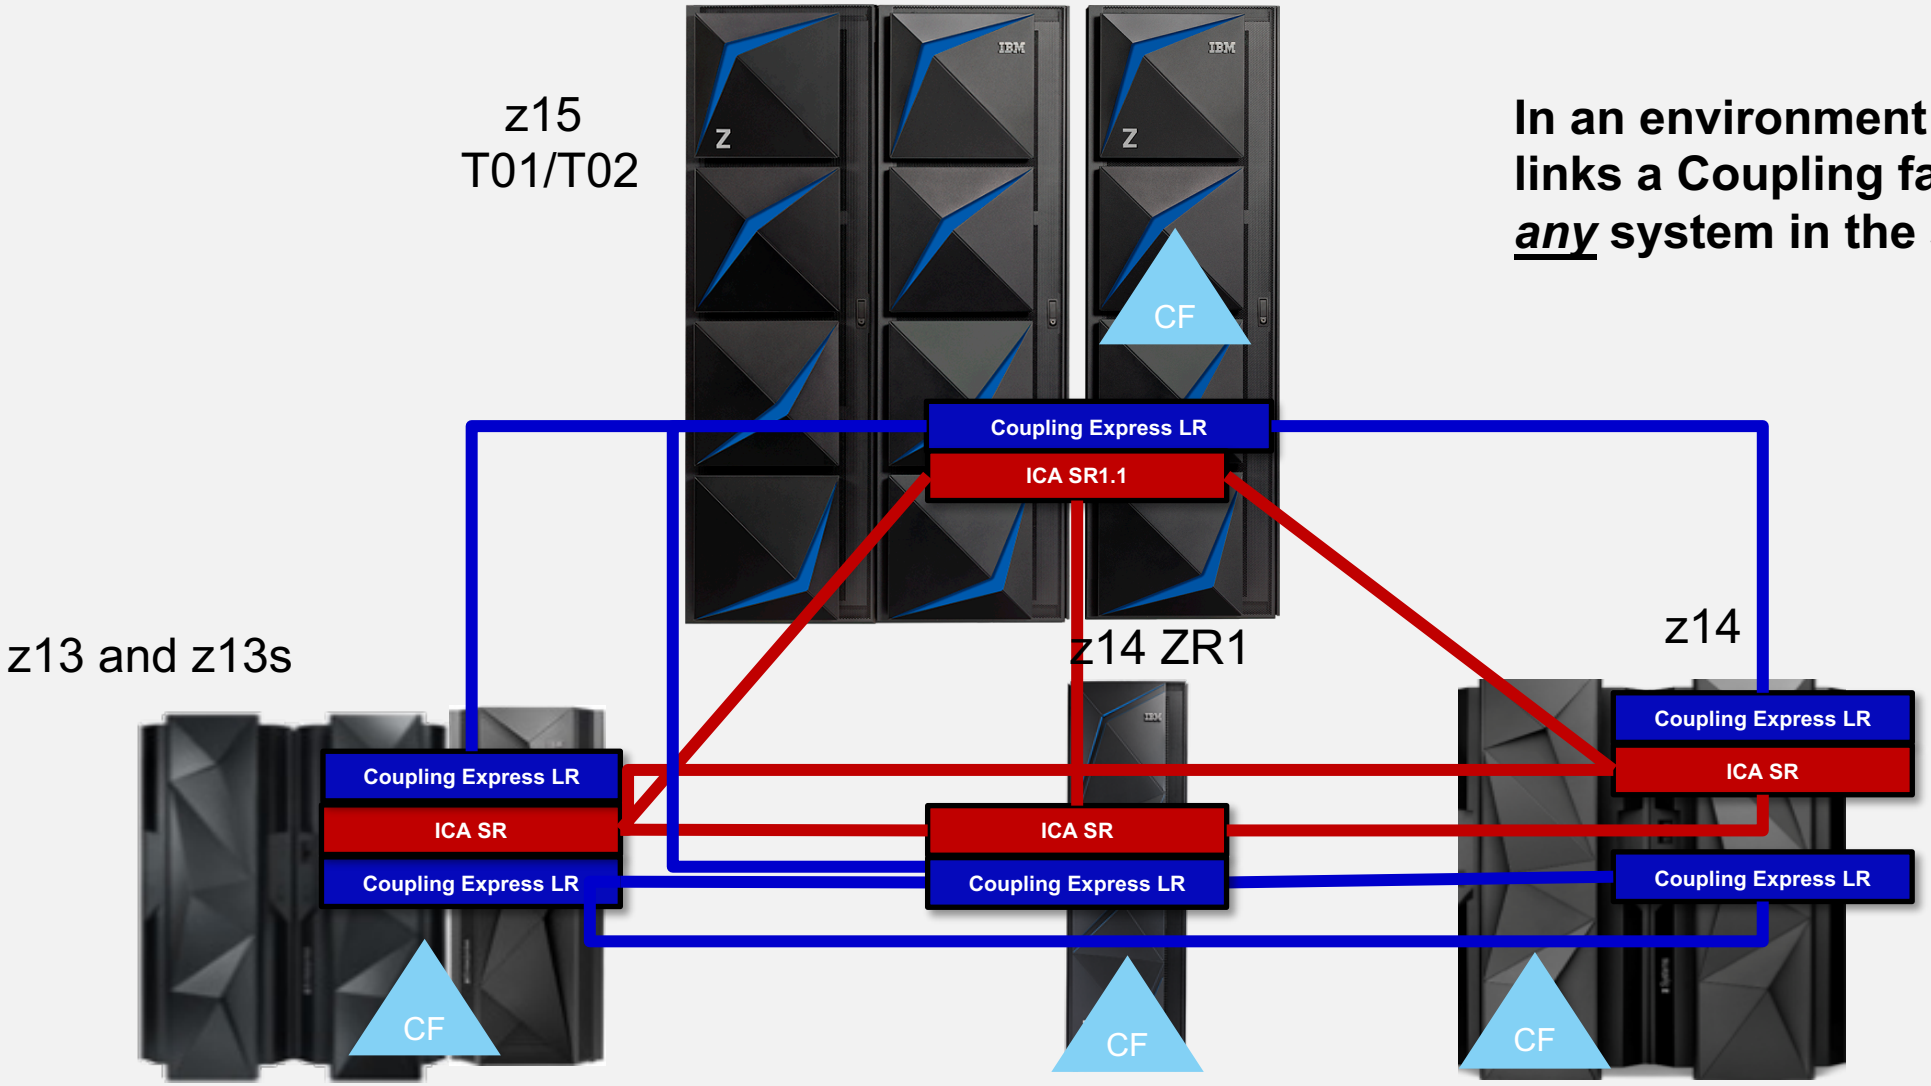

In an environment with no HCA3-O links a Coupling facility can reside on *any* system in the sysplex

z15 Coupling Migration with HCA3-O environment

z15
T01/T02

z13/z13s withdrawn from marketing June 2019

In an environment with a z13 or z13s that only has HCA3-O coupling optics the Coupling Facility *MUST* reside on a z14 with coupling connections to all other systems

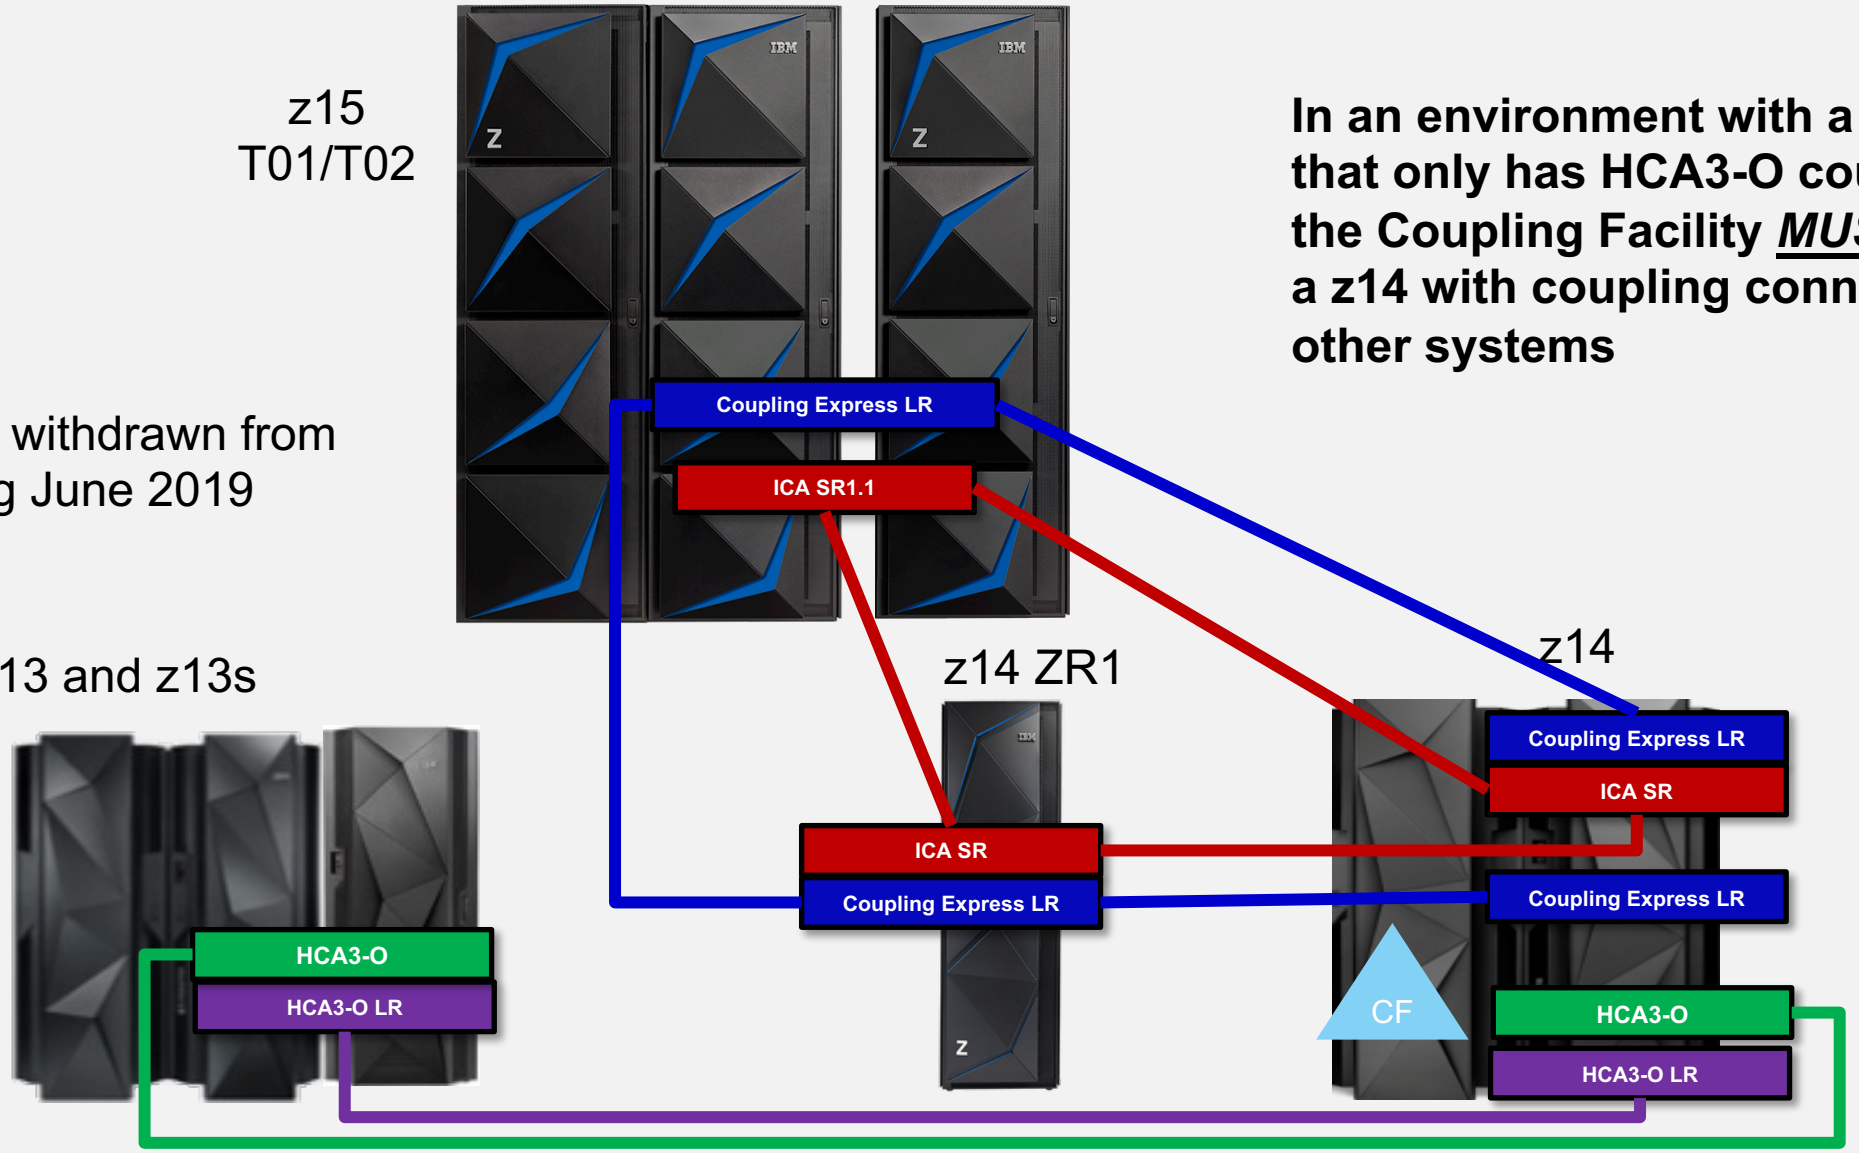

HCA3-O / HCA3-O LR Only communicate between the z13/z13s/z14 (M0x Models)

z15 STP

Why Precision Time Protocol (PTP)?

Prepare for prerequisite regulations and audits.

Synchronization Regulations

- SEC - within 50ms of UTC
- FINRA (investor fairness) - within 50ms of UTC
- MiFID (European – like FINRA) - within 100µs of UTC
- Other regulating agencies are developing standards ie. PCI

Single Point of Failure reporting

- Another European regulation
- If the STP Preferred Time Server & Current Time Server (PTS/CT) is taken down or fails, the Backup Time Server (BTS) becomes the Current Time Server. In this mode, it could be seen as a single point of failure.

PTP History

- Standard designed for labs, cell phone networks, radar, robots
- Finding its way into computer networking
- Capable of synchronization to less than a microsecond

Precision Time Protocol (PTP) - IEEE 1588

The Precision Time Protocol (PTP) is a protocol used to synchronize clocks throughout a computer network. Other vendor platforms may also exploit PTP.

- **PTP is designed for local systems requiring accuracies beyond those attainable using NTP. [Wikipedia]**
- **The first IBM Z implementation will provide the use of an external time source (ETS) capable of supporting PTP.**
  - **Accuracy today with**
    - **NTP – within about 100ms**
    - **NTP with PPS – within about 10µs**
    - **PTP with PPS – within about 10µs**
    - **PTP future without PPS – 100ns or less.**

z15 provides a stepping stone to further PTP support.

- **Sets the stage for exciting future enhancements.**

PTP ETS exploitation is optional on z15. No feature code required.

Implementation details

- PTP requires network infrastructure support
  - For z15, the PTP Ethernet cable must plug directly into the SE and HMC/SE virtual appliance*
  - Due to SE to CPC interface, for current implementation, Pulse Per Second still required for PTP to meet Financial Regulations

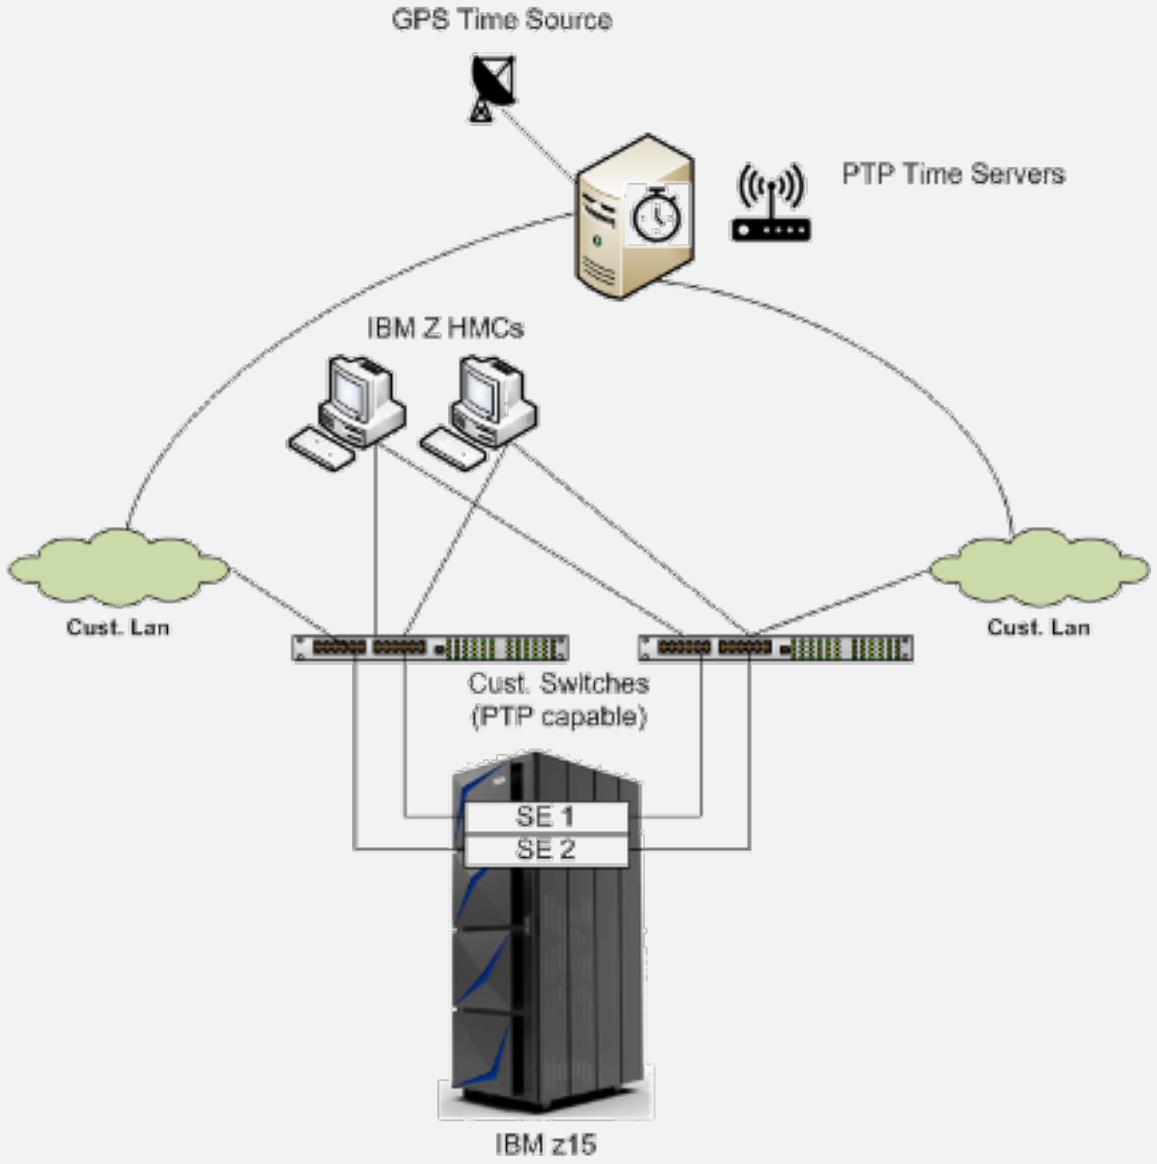

z15 Power and cooling

Power considerations for z15

- z15 Model T02 max vs. z14 ZR1:
  - Core frequency: 4.5 GHz (same as 3907)
  - 76 physical cores (active) on max system vs. 35
  - max memory +100% = 16TB vs 8TB

- Power choice: Single Phase* or Three Phase Power Distribution Units (PDU)
  - Estimated maximum power for z15 T02:
    - Max65 feature with three PCIe+ I/O drawers – 8.1 kW

- One Model T02 and five features (one or two CPC drawers)
  - Max4, Max13, Max21, Max31 – One CPC Drawer
  - Max65 – Two CPC Drawers
  - Air cooled CPC Drawers (air cooled PUs, SCs) and PCIe+ I/O Drawers
  - Single frame system, 19” industry format
  - Front to back air flow

Power considerations for z15 T02

- Three-phase power:
  - Low voltage 4 wire *Delta (“*Δ”) (FC 0629)
  - High voltage 5 wire *Wye* (“Y”) (FC 0630)

- Single phase:
  - Single phase system has only one CPC Drawer (up to Max31)
  - No concurrent upgrade to Max65 (and no CPC1 reserve feature FC 2271). Upgrade requires an outage if going from 1ph to 3-ph.
  - FC 0630 only with single phase line cord

- No Bulk Power (BPA), and no Internal Battery (IBF)
- No EPO (Emergency Power Off) option
- No High Voltage DC option (~~HVDC~~)

z15 T02 Power Cords Considerations

- System has been designed in a single 19” frame format (similar to z14 ZR1)

- Number of power cords (and PDUs) depends on the system configuration:

CPC \ I/O	0	1	2	3	4
1	2	2	4	4	4
1 & FC2271	2	2	4	4	4
2	4	4	4	4	4

Blue -> single or three phase
Purple -> three phase only

- A conversion between single phase power to three phase power will require a CEC wide outage
- For dimensions, weights, and utility power consumption see the manual:
  - ***IBM 8562 Installation Manual for Physical Planning, GC28-7011***

Power Line Cords

Feature Code	Description	Comments	Co-Req Feature Code
FC 7892	30A/208V 14ft w/TwistLock	US, Japan, all CCN, all LA 1-Phase w/1 CPC.	FC 0630
FC 7893	30A/208V 14ft w/Russelstoll	US, Japan, all CCN, all LA. 1-Phase w/1 CPC.	FC 0630
FC 7894	32A/250V Cord EMEA & AP	EMEA & AP. 1-Phase w/1 CPC.	FC 0630
FC 7895	32A/250V Cord Aus & NZ	Aus & NZ. 1-Phase w/1 CPC.	FC 0630
FC 7896	32A/250V Cord Korea	Korea. 1-Phase w/1 CPC.	FC 0630
FC 7897	32A/250V LSZH Cord	WW except US, Japan, Australia, Korea, India, China & CCN. 1-Phase w/1 CPC.	FC 0630
FC 7946	30A/400V 3Ph Wye w/Hubbell	US, CCN, JAPAN.	FC 0630
FC 7947	32A/380-415V 3Ph Wye	World Trade.	FC 0630
FC 7948	32A/380-415V 3Ph Wye LSZH	World Trade.	FC 0630
FC 7952	30A/250V 3Ph w/Hubbell	US, CCN, JAPAN.	FC 0629
FC 7953	30A/250V 3Ph w/Cut End	World Trade.	FC 0629

CCN = Canada, Australia, New Zealand and some Caribbean islands
LSZH = Low smoke zero halogen

Environmental Classes and IBM Z

z13	A1	Temp range 15C (59F) to 32C (89.6F) RH range 20% to 80% Max dew point 17C (62.6F) (note 1) Alt 10,000 ft.	zEC12
	A2	Temp range 10C (50F) to 35C (95F) RH range 20% to 80% Max dew point 21C (69.8F) (note 1) Alt 10,000 ft.	zBC12
z15, z14	A3	Temp range 5C (41F) to 40C (104F) RH range -12C (10.4F) min dewpoint; 8% to 85% Max dew point 24C (75.2F) (note 1) Alt 10,000 ft.	z15 T02 z14 ZR1, z13s
Recommended (all)		Temp range 18C (64.4F) to 27C (80.6F) RH range 5.5C (41.9F) min dewpoint; up to 60% Max dew point 15C (59F) (note 1)	

Note 1: Actual inlet air moisture content range (grams moisture/Kg dry air)

A1:	2 - 12
A2:	1.5 - 16
A3:	1.5 - 19
Recommended:	6.2 - <10

Hot and cold aisle cooling containment

- Partitions used in the aisle cooling containment solution should be self-supporting and not attach to the frame for structural integrity.
- Service clearances must be maintained when the system is installed in the cooling containment solution.
- Airflow requirements of the z15 must meet the intended and abnormal operating conditions to assure sufficient airflow can be provided, particularly in cold aisle cooling containment systems.
- Ensure racks across the aisle will not negatively impact the z15 operation or the z15 impacts other racks.

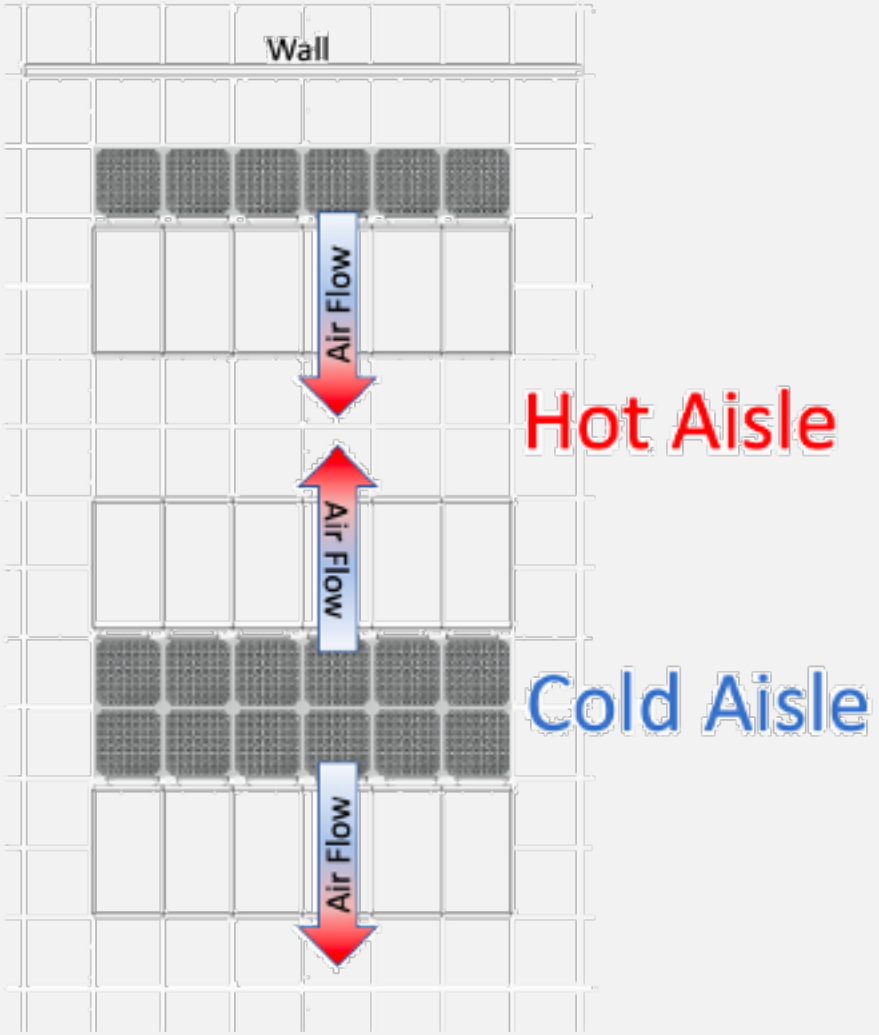

z15 Physical Configuration

z15 T02/LT2 System Overview

- Designed to align with data center trends, to optimize cost, density, & flexibility
- Robust 19” frame package
  - Same frame as the z14 ZR1
  - Based on the design principals of the z15 T01
  - “True 19 inch” frame (depth, width, cabling, airflow, power)
    - 42U high (40U while being moved with height reduction option)
    - 600 mm width frame measured without side covers
    - 1.20 meter depth with acoustic covers

- This change allows us to become a more strategic part of a clients’ datacenter infrastructure
- Designed to fit most datacenters hot/cold aisle containment systems
- Can be integrated into existing aisles
- All rear exit cabling
- Can be on Raised/Non-raised Floors

Package Shipping Specifications

- **z15 shipment:**
  - Systems are shipped protected with an anti-static poly bag with desiccant inside. This is placed in a wooden container and mounted on pallets requiring commercial lift transportation. This packaging is used for all world-wide shipments.

- **Packaged Shipping Specifications:**

	Width mm (in)	Depth mm (in)	Height mm (in)	Max. Weight kg (lb)
Crated Frame w/o covers	940 mm (37.0 in)	1334 mm (52.5 in)	2286 mm (90.0 in)	TBD
Cover Set with Packaging	457 mm (18.0 in)	2134 mm (84.0 in)	1016 mm (40.0 in)	TBD

- **Height Reduction (FC 9975):**
  - If during transit to or entry to the final installation location contains openings less than 2032 mm (80.0 in) high, FC 9975 should be ordered.
  - This feature reduces the frame height to 1898 mm (74.7 in), with the frame's top hat and side covers shipped in a separate carton.
  - Note: FC 9976 shall be ordered for "ship to return" packaging

**Excludes front and rear doors (called covers here), but does include side covers*

Guide For Raised Floor Preparation

- **Things to consider:**
  - Rear exit I/O and power only, both for bottom and top exit cabling
  - New floor cutouts may be necessary depending on placement
  - Raised floor stanchions locations may need to be reevaluated
  - As with previous system the z15 should be placed with adequate service clearance areas on both the front and back of the system.

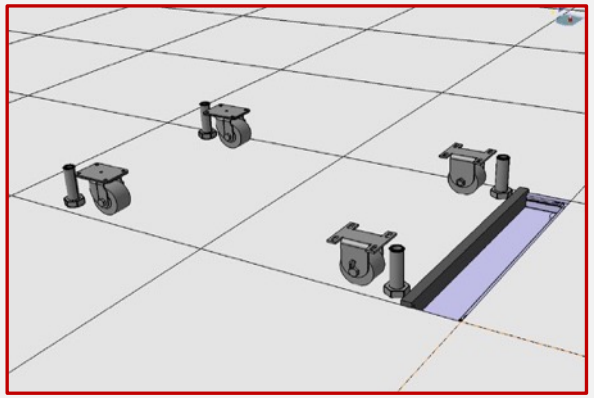

Max Floor Tile Cut-out

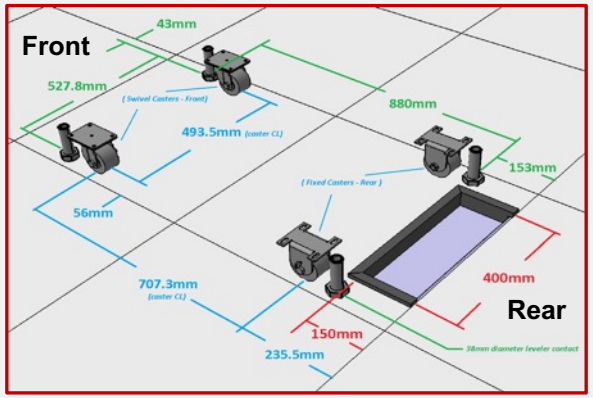

Typical Floor Tile Cut-out

Recommended Stanchion Locations

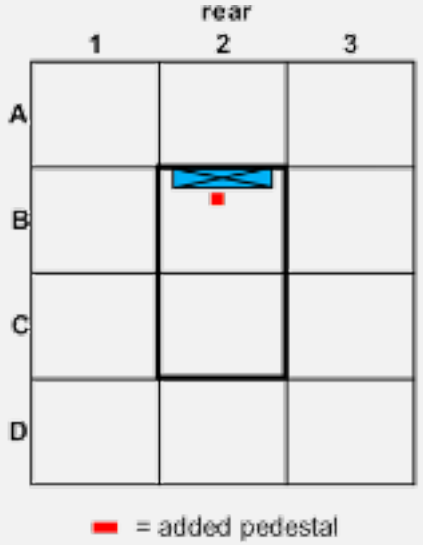

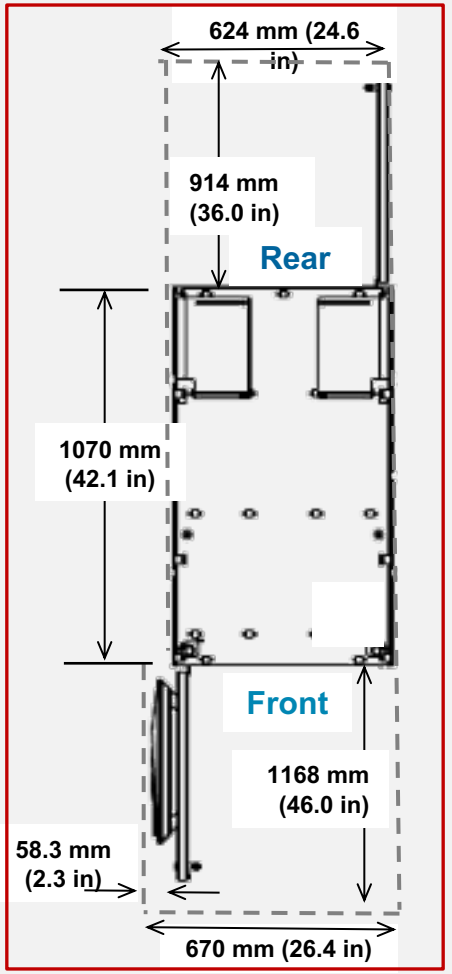

Service Clearance Area

Consult the z15 IMPP for additional information IMPP GC28-7011

Cabling Design and Management

- Top exit and bottom exit options (applicable to both power and I/O cabling)
  - Three feature codes available for both power and I/O cable management
  - FC 7899 Bottom Exit Cabling
    - Necessary hardware to allow cable egress out the bottom
    - Will still allow a client to exit out the top
  - FC 7898 Top Exit Cabling with Top hat
    - Includes an additional top hat hardware to use the trunking cables/slack management
    - Should be ordered if using fiber trunking services and you want to exit out the top
  - FC 7928 Top Exit Cabling without Top Hat
    - Should be order for non-raised floor or if client only wants top exit cabling
    - Comes with a cover to seal the bottom of the frame that limits access as well as a top plate on the top that will slide open for cables to egress.

FC 7898 Top hat

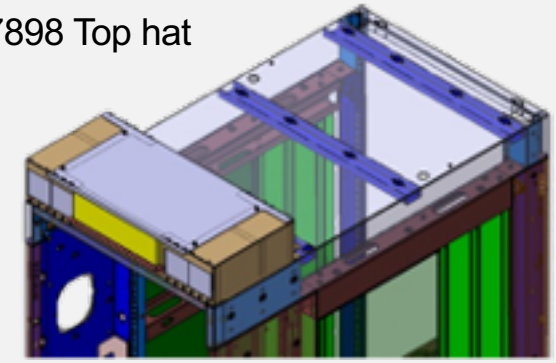

top plate

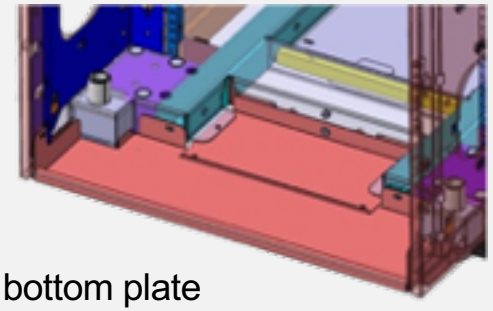

bottom plate

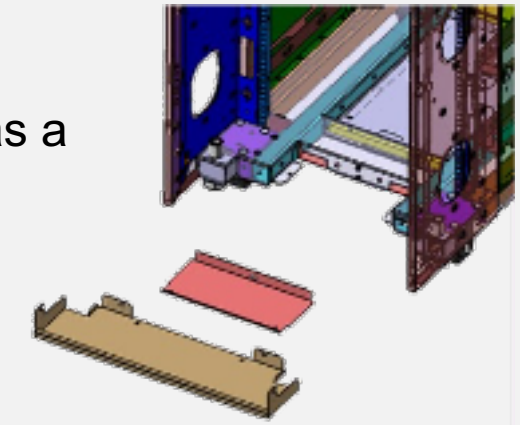

bottom plate

Cabling Spine

The full length spine is used for cable management when any of the following is ordered:

1. 2nd CPC Drawer (Max65)
2. 3rd PCIe+ I/O Drawer

z15 T02 Frame Overview

- Systems can reserve space for the additional CPC Drawer that supports the Max65
  - 2271 – Yellow box

- Removing these F/C can allow for one additional I/O drawer to be plugged in place of the CPC drawer

- **Clients *will not* be able to upgrade to the Max 65 if they override reserve features and plug I/O drawers.**

Frame Layout

Here are some sales examples you may run into

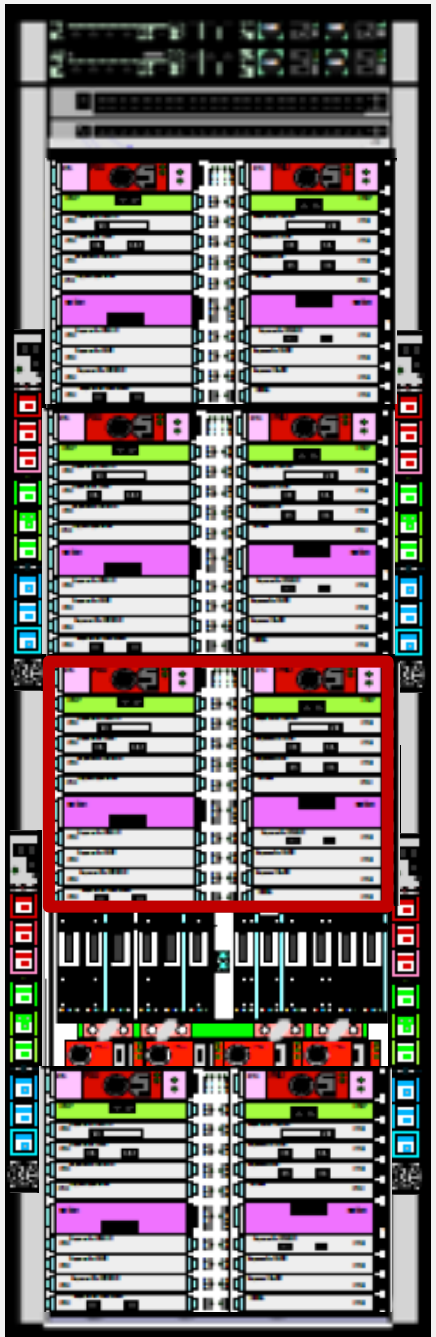

z15 HMC

IBM z15 Hardware Management Console

You can order up to 10 HMCs total for a system

Available HMCs for z15

Feature Code	Description	CF/NB
0062	HMC Tower	NB
0063	HMC Rack Mount	NB
0082	HMC Tower	CF
0083	HMC Rack Mount	CF
0095	HMC Tower	CF
0096	HMC Rack Mount	CF
0100	HMA	NB

Initial order only. No MES add later.

Available peripherals

- HMC Table top KMM FC 0148
- HMC Rack Mount KMM FC 0154

HMC Support

- Seamless hard drive encryption
- Will only support n-2 Systems
- DVD drives removed

Machine Family	Machine Type	Firmware Driver	SE Version
z15 T02	8562	41	2.15.0
Z15 T01	8561	41	2.15.0
z14 M0x	3906	36	2.14.1
z14 ZR1	3907	36	2.14.1
z13	2964	27	2.13.1
z13s	2965	27	2.13.1

HMC Support

- What's different?
  - No DVD/CD Drives on new HMCs, SEs & TKEs
  - Two options for service and functional operations
    - USB
      - IBM supplied USBs preformatted and loaded with the needed code
    - Electronic
      - FTP – File Transfer Protocol, SFTP – SSH File Transfer Protocol, SSL/TLS FTP - FTPS
      - Z Remote Support Facility (RSF), IBM Resource Link, and FTP/SFTP/FTPS Server connections from the HMC. Note that for an electronic-only delivery environment, there is a requirement that there are two HMCs on every unique network subnet where a Hardware Management Console, Support Element, or Trusted Key Entry workstation is connected.
        Two units required for dead HMC Recovery Scenario.
  - Both options provide:
    - Firmware loading
    - eBoD (eBusiness on Demand) Records via ResourceLink
    - Operating System code

SE, HMC & TKE – USB Summary

- FC0843
  - 32GB USB Load Media
    Used to load and install purchased software (operating systems, etc).
- FC0846
  - No physical media (previous slide)
- FC0848
  - 32GB USB Backup Media
    HMC or SE Critical Data Backup

Hardware Management Appliance details (1 of 2)

- The Support Element runs as a guest of the HMC.
  - Hardware: Two 1U Rack Mounted servers running HMCs first level.
  - SEs maintain primary/alternate relationship
  - HMCs do not have a primary/alternate relationship, though data replication can be set up
- Manage the SE using the new “Virtual Support Element Management” task
  - Allows start/stop of the SE
    - These operations should be seldom used since:
  - The SE is automatically started when the HMC boots
  - The SE “Shutdown” task should be used to gracefully shutdown the SE. The stop function should only be needed in the event of an SE hang
  - Allows user to open the graphical console of the SE

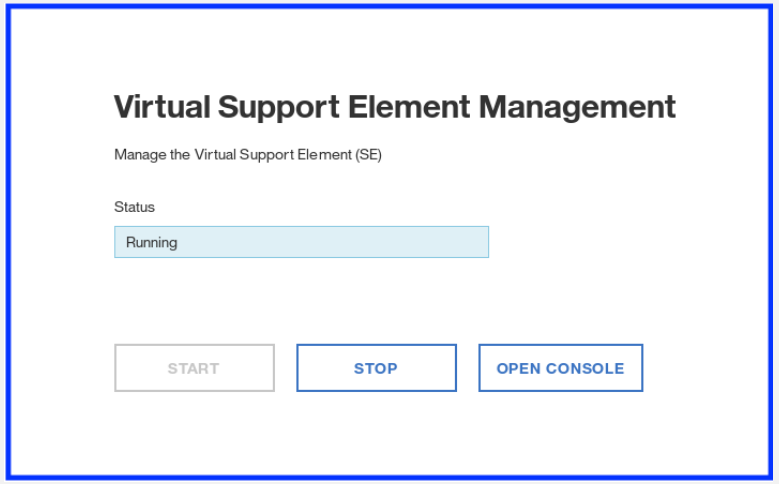

Hardware Management Appliance details (2 of 2)

- Accessing the HMC and SE.
  - You can access the HMC through a remote web browser as usual
  - If you walk up to the physical KMM display, you will see the HMC UI
  - The SE can be accessed via “Single Object Operations” from the HMC as usual
  - You can also use “Virtual Support Element Management” -> “Open Console” to view the full SE desktop

- Things to be aware of
  - Since the SE is running as a guest under the HMC, any reboot or shutdown will be disruptive to the SE running on that HMC
    - This means that operations like HMC code updates must be performed on the HMC hosting the Alternate SE first, followed by a Primary/Alternate switch before proceeding to update the other HMC.
  - ”Restart application” on the HMC will *not* disrupt operation of the SE.
  - Vast majority of SE and HMC tasks work the same way they do on standalone consoles

Mini KMM for SE

- SE Keyboard/Display intended for Service use
  - Non critical hardware for functional & monitoring of CPC, Images, & OSes
  - Customer physical access intended use => HMC only (no SE physical access)
    - Understand that some clients may physically access SE
    - New Human Interface for SE => clients must utilize service solution
    - For Virtual Appliance Feature
      - Intended HMC/SE Physical access via remote browser workstation
      - If clients access VA directly => must accept service solution
    - Recommendation to clients: Don’t physically access Mini KMM in IBM Z frame

- Mini (or Compact) KMM (Keyboard Monitor Mouse) solution
  - Single KMM
  - Attaches via USB-C cable connection
    - Connections via front/back of frame
    - Interface Adapter to switch between Primary/Alternate SEs

Mini KMM for SE

Open Cubby => KMM Stored

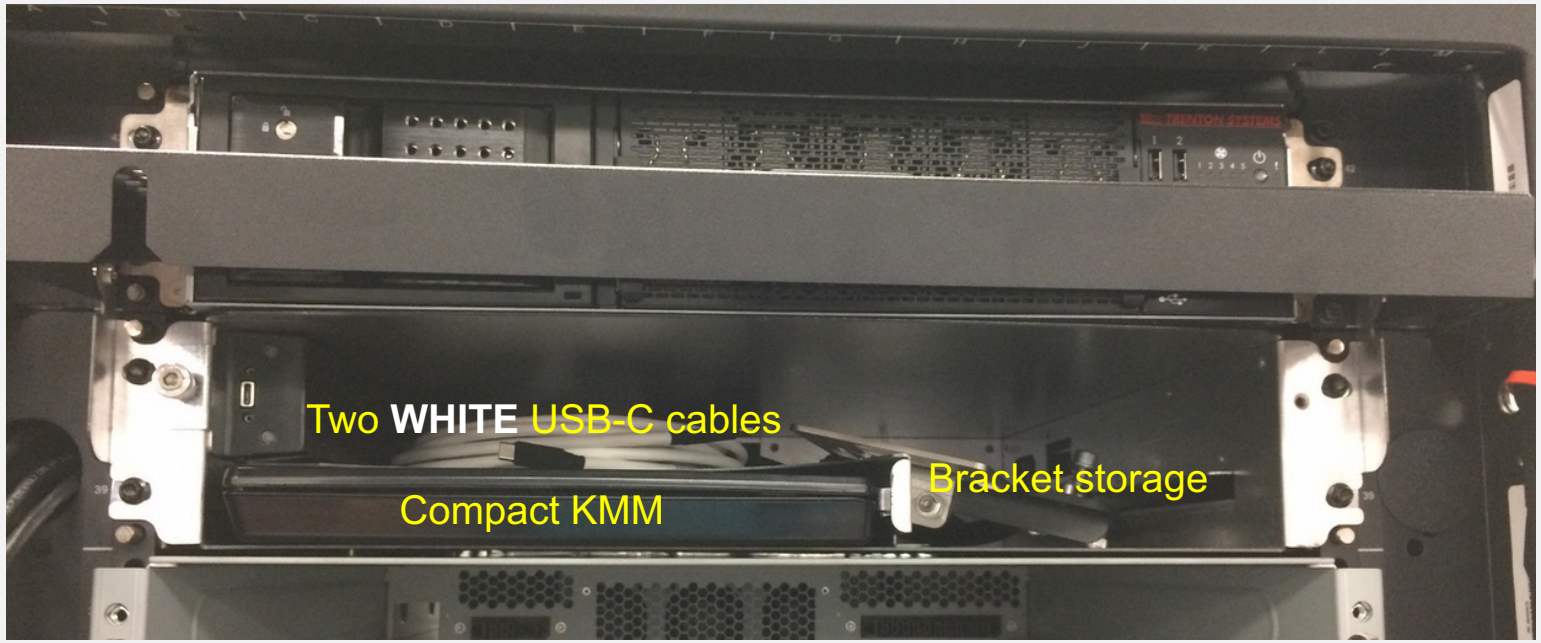

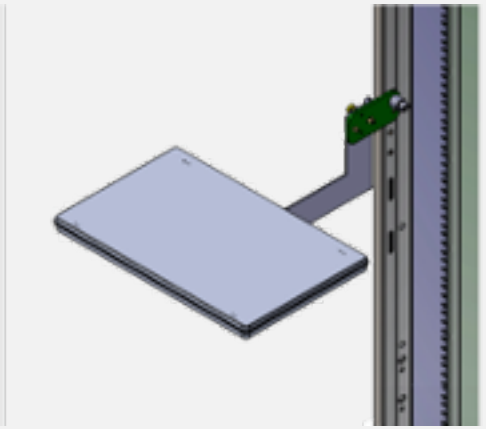

Mini KMM
Frame Mounting
Opens OUT of the drawer space

Mounts to either side,
left or right,
front or back

Can mount on an adjacent rack

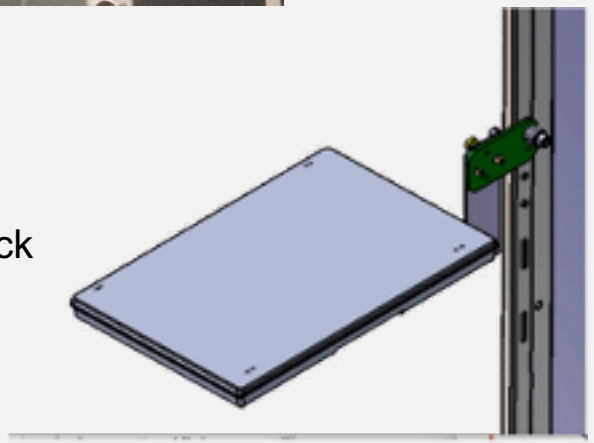

z15 Site Tools

Lift Tool and Extension Ladder

- Lift Tool FC 3100
  - Same as the z14 ZR1
  - 5 separate components that attach to the front of the system
  - Each site requires at least one for all z15 systems

–

- Extension Ladder FC 3101
  - Same as the z14 ZR1
  - Each datacenter requires at least on for all z15 Systems

z15 Operating Systems

IBM z15 T02 operating system support

z/OS
- **z/OS 2.4**
- **z/OS 2.3 with PTFs**
- **z/OS 2.2 with PTFs**
- **z/OS 2.1 (compatibility only)**
  - IBM Software Support Services purchase
  - September 2018, EoS

z/VM
- **z/VM 7.1 with PTFs**
- **z/VM 6.4 with PTFs**

z/VSE
- **z/VSE 6.2**

KVM Hypervisor
- **Red Hat RHEL 8.0 + higher**
- **SUSE SLES 15 SP1 maintweb + higher**
- **SUSE SLES 12 SP4 maintweb + higher**
- **Ubuntu 18.04 LTS**
- **Ubuntu 16.04 LTS**

Linux on IBM Z
- **Minimum Distributions**
- **RHEL 6.10**
- **RHEL 7.7.z alt**
- **RHEL 8.0**
- **SLES 12 SP4**
- **SLES 15 SP1**
- **Ubuntu 16.04 LTS**
- **Ubuntu 18.04 LTS**

IBM cannot legally discuss zNext exploitation prior to GA from distributors.

Officially Tested:
www.ibm.com/systems/z/os/linux/resources/testedplatforms.html

z/TPF
- **z/TPF 1.1 with PTFs**

z15 Statements of Direction

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remain at our sole discretion.

Statements of General Direction

- **Prepaid OOCoD tokens:** Beginning with IBM z15, new prepaid OOCoD tokens purchased will not carry forward to future systems.

- **TLS 1.0 for OSA, HMC, and SE:** IBM z15 will be the last IBM Z server to support the use of the Transport Layer Security protocol version 1.0 (TLS 1.0) for establishing secure connections to the Support Element (SE), Hardware Management Console (HMC), and OSA-Integrated Console Controller (channel path type OSC).

- **A new use of System Recovery Boost:** In the future, IBM plans to introduce a new use of System Recovery Boost that will focus on a limited number of short-duration boosts. These boosts are mediated by the operating system and designed to improve system resiliency during specific focused recovery actions.

- **Prepaid token expiration:** Beginning with IBM z15 Model T02, prepaid tokens for On/Off Capacity on Demand (On/Off CoD) will expire 5 years after LICCC expiration date.

- **Water Cooling:** IBM z15 is planned to be the last IBM Z server to offer customer water cooling .

- **Future HMC Hardware:** IBM z15 is planned to be the last server to offer the ability to order stand alone Hardware Management Console (HMC) hardware. For future systems, new HMC hardware can only be ordered in the form of the Hardware Management Appliance feature (#0100) which was introduced on IBM z15. The Hardware management Appliance feature provides redundant HMCs and Support Elements (SEs) that reside inside the Central Processor Complex (CPC) frame, and the ability to eliminate stand alone HMC hardware (tower or rack mounted) outside the CPC frame. Stand alone HMC hardware (tower or rack mounted) can still be ordered and used with IBM z15.

Statements of General Direction

- **Operational Data Generation and Analytics**: In the future IBM intends to deliver z/OS and Middleware interdependency data generation, and automated z/OS cross stack analytics to reduce skill requirements level and amount of time required to perform problem definition.
- **z/VSE exploitation of System Recovery Boost**: In the future, IBM intends to deliver native z/VSE exploitation of System Recovery Boost, which is expected to enable restoration of service from, and catch up after, both planned and unplanned outages faster than on any prior Z machine.
- **The conditional-SSKE facility:** IBM z15 will be the last high end server to support the conditional-SSKE facility.
- **Reserved space for DS8910F:** 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. Clients 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.

Thanks