Introducing the new IBM z15 T02
Extending the family with the new air cooled, single frame model

IBM z15 Model T02

IBM LinuxONE III Model LT2

IBM LinuxONE III Model LT1

IBM z15 Model T01

| Built on IBM z15 chip technology | Cloud native development and deployment | Encryption everywhere protecting eligible data | Resiliency and availability in simplified package |

IBM z15 Model T02, IBM LinuxONE III Model LT2

IBM z15™

Machine Type: 8562
Model T02

Feature Based Sizing

CPC Drawer	Customer PUs	Max Memory
Max4	4	2 TB
Max13	13	4 TB
Max21	21	4 TB
Max31	31	8 TB
Max65	65	16 TB

IBM LinuxONE™ III

Machine Type: 8562
Model LT2

Feature Based Sizing

CPC Drawer	IFLs	Max Memory
Max4	4	2 TB
Max13	13	4 TB
Max21	21	4 TB
Max31	31	8 TB
Max65	65	16 TB

IBM z15 Model T02   **IBM LinuxONE III Model LT2**

TechU

The New IBM z15 A technical review of the Processor Design, New Features, I/O Cards, and Crypto 2020

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

2020 IBM Systems TechU

IBM

Table of Content

- z15 Overview
- z15 Processor Design and Structure
- z15 Memory
- z15 New Features and Functionality
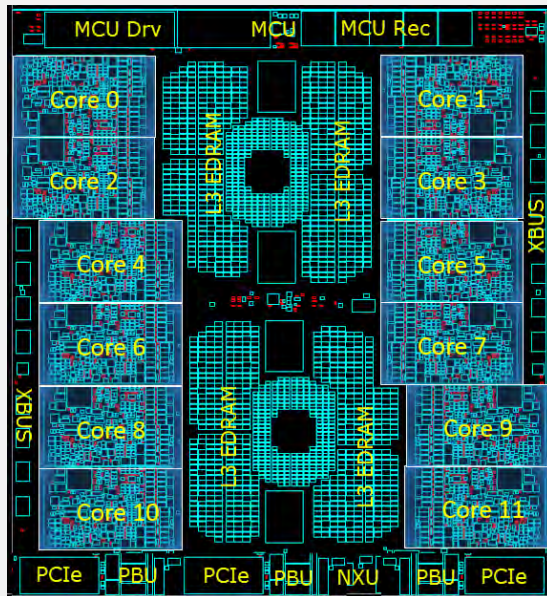- z15 I/O Infrastructure

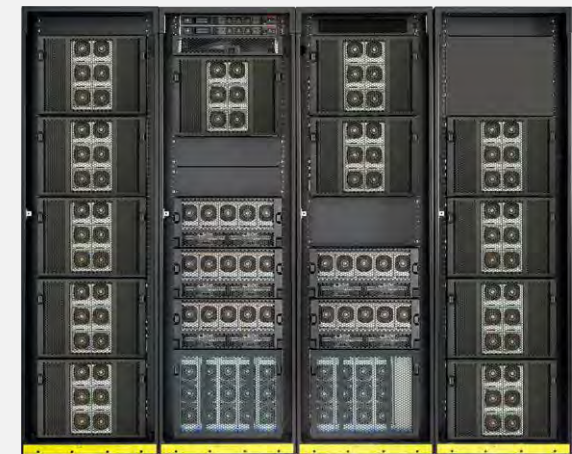# z15 Overview

z15 System Design Changes

- 14 nm Processor with optimized Out-of-Order, new DEFLATE and SORT

- 12 Cores per PU SCM design

- 4 CP SCMs per Drawer, up to five 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)

- OSA-Express7S 25GbE

- OSA-Express6S

- New FICON Express16SA

- 25GbE and 10GbE RoCE Express2

- IBM zHyperLink Express
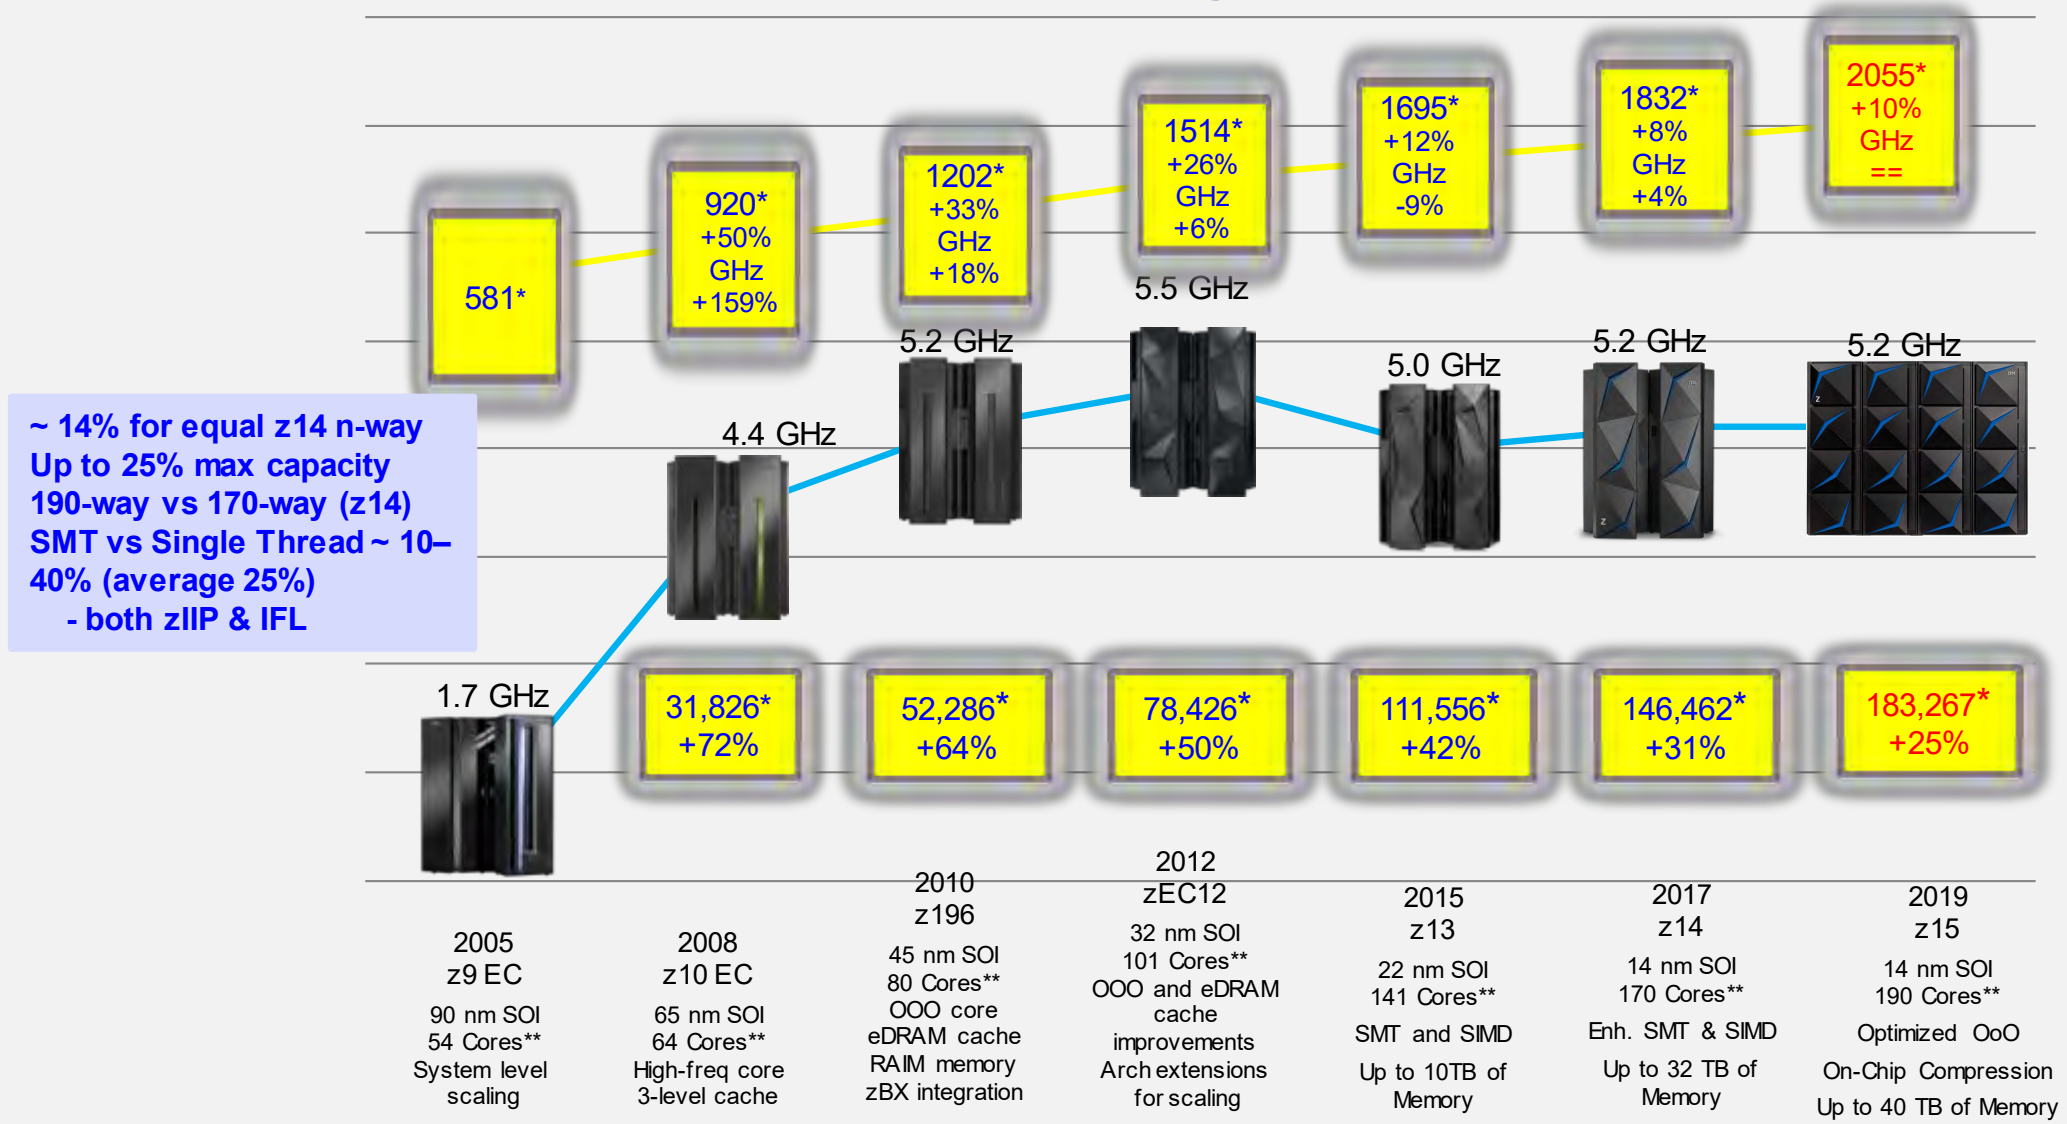
- Coupling Express Long Reach

Back   Front

z15 Continues the CMOS Mainframe Heritage

GHz / PCI*

~ 14% for equal z14 n-way
Up to 25% max capacity
190-way vs 170-way (z14)
SMT vs Single Thread ~ 10–40% (average 25%)
- both zIIP & IFL

581*

920*
+50% GHz
+159%

1202*
+33% GHz
+18%

1514*
+26% GHz
+6%

1695*
+12% GHz
-9%

1832*
+8% GHz
+4%

2055*
+10% GHz
==

1.7 GHz

4.4 GHz

5.2 GHz

5.5 GHz

5.0 GHz

5.2 GHz

5.2 GHz

31,826*
+72%

52,286*
+64%

78,426*
+50%

111,556*
+42%

146,462*
+31%

183,267*
+25%

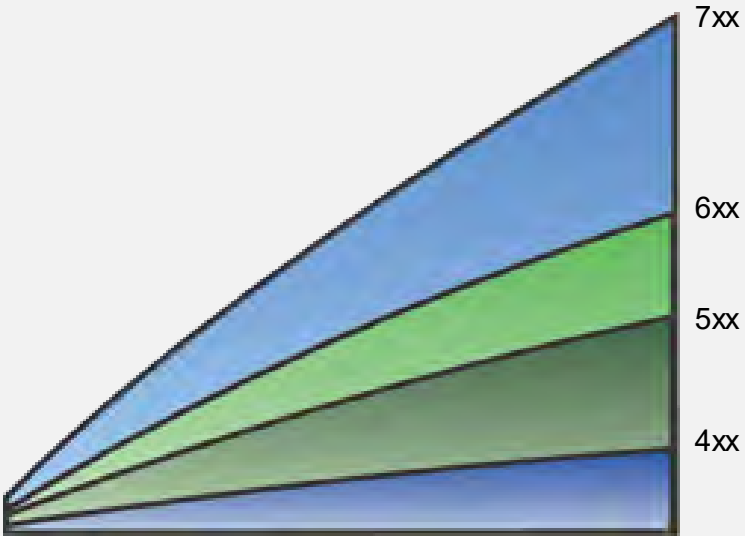| 2005 z9 EC	2008 z10 EC	2010 z196	2012 zEC12	2015 z13	2017 z14	2019 z15
90 nm SOI	65 nm SOI	45 nm SOI	32 nm SOI	22 nm SOI	14 nm SOI	14 nm SOI
54 Cores**	64 Cores**	80 Cores**	101 Cores**	141 Cores**	170 Cores**	190 Cores**
System level scaling	High-freq core 3-level cache	OOO core eDRAM cache RAIM memory zBX integration	OOO and eDRAM cache improvements Arch extensions for scaling	SMT and SIMD Up to 10TB of Memory	Enh. SMT & SIMD Up to 32 TB of Memory	Optimized OoO On-Chip Compression Up to 40 TB of Memory

*   PCI Tables are NOT adequate for making comparisons of IBM Z processors. Additional capacity planning required
** Number of PU cores for customer use
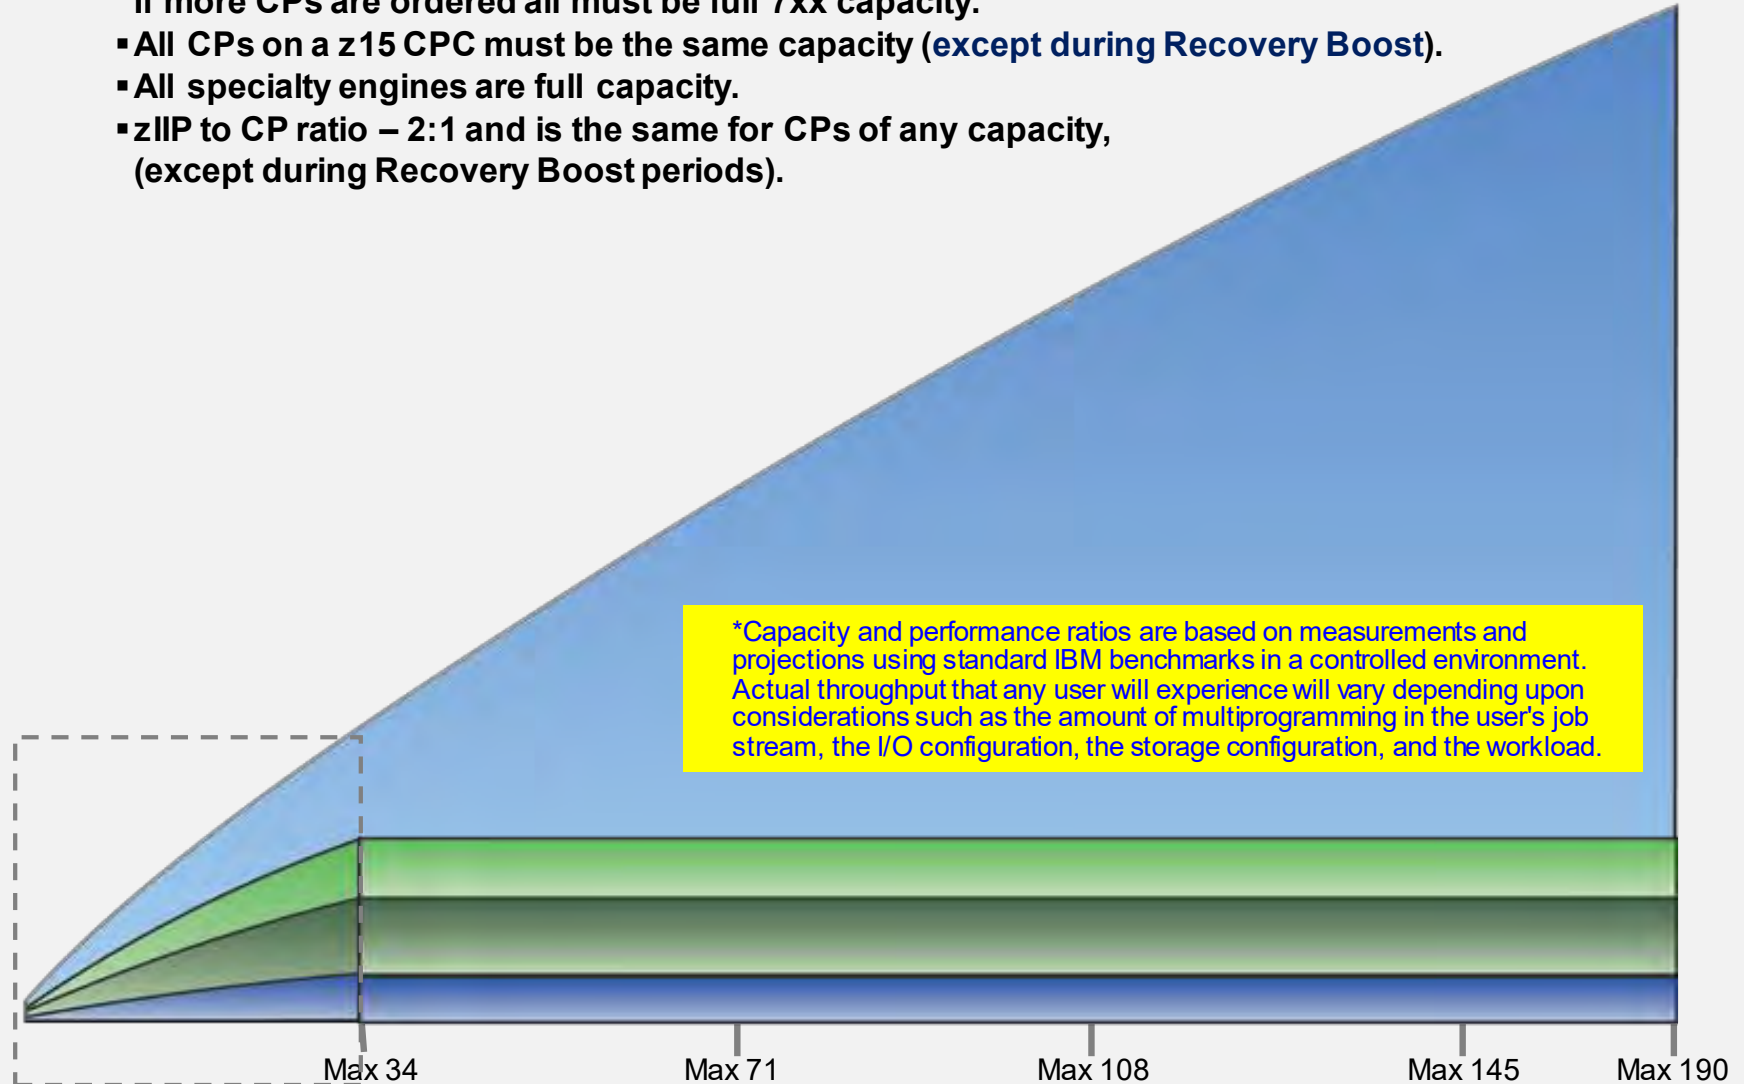
z15 Full and Sub-Capacity CP Offerings

CP Capacity – Relative to Full Capacity Uni
- 701 = 100% ≈ 2,055.3 PCI
- 601 ≈ 0.56% ≈ 1,151 PCI
- 501 ≈ 0.38% ≈ 781 PCI
- 401 ≈ 0.13% ≈ 267.2 PCI

- Subcapacity CPs, up to 34 may be ordered (286 capacity levels). If more CPs are ordered all must be full 7xx capacity.
- All CPs on a z15 CPC must be the same capacity (except during Recovery Boost).
- All specialty engines are full capacity.
- zIIP to CP ratio – 2:1 and is the same for CPs of any capacity, (except during Recovery Boost periods).

7xx

6xx

5xx

4xx

MSU Sub Capacity

*Capacity and performance ratios are based on measurements and projections using standard IBM benchmarks in a controlled environment. 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.

Max 34   Max 71   Max 108   Max 145   Max 190

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

- **Ensure the CPU MF data is captured and <u>kept</u> 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.**

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

z15 Processor Design and Structure

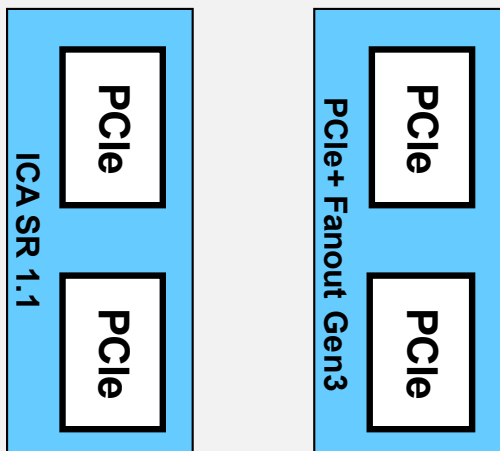# z15 Processor Drawer (Top View)

- 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)
  - Four PU Chips per CPC Drawer:
    - 41 active PUs per drawer - Max34, Max71, Max108 and Max145
    - 43 active PUs per drawer – Max190
  - One SC Chip (960 MB L4 cache)
  - DIMM slots: 20 DIMMs to support up to 8 TB of addressable memory (10 TB RAIM)
  - Water cooling for PU SCMs, air cooled SC SCM
  - Two Flexible Support Processors/ OSC Cards
  - 12 fanout slots for PCIe+ I/O drawer or PCIe coupling fanouts (ICA SR)

Rear

SC

4 x PU under
The cold plates

Up to 20
Memory DIMMs

Front
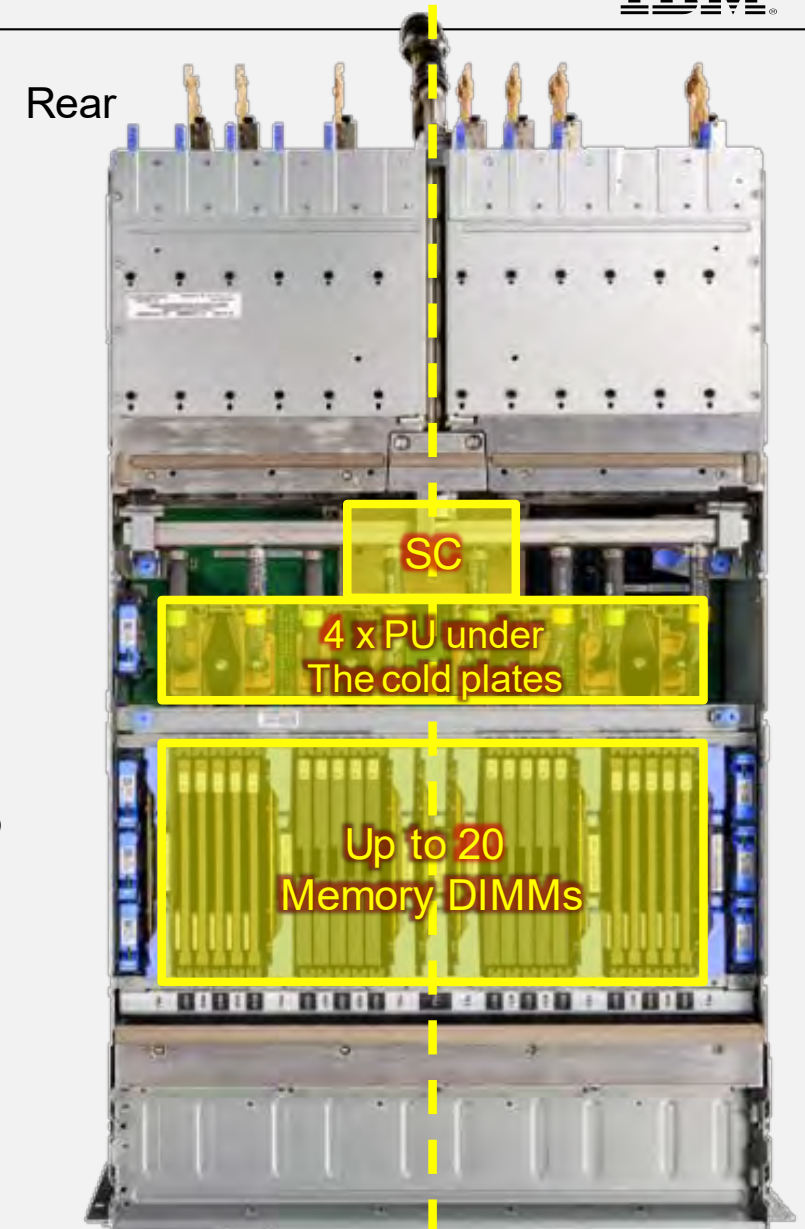
z15 Processor (CPC) Drawer Connectivity

No InfiniBand fanouts

Rear

- 12 PCIe fanout slots per z15 CPC drawer
  - Increase from 10 PCIe fanouts in z14

- 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)
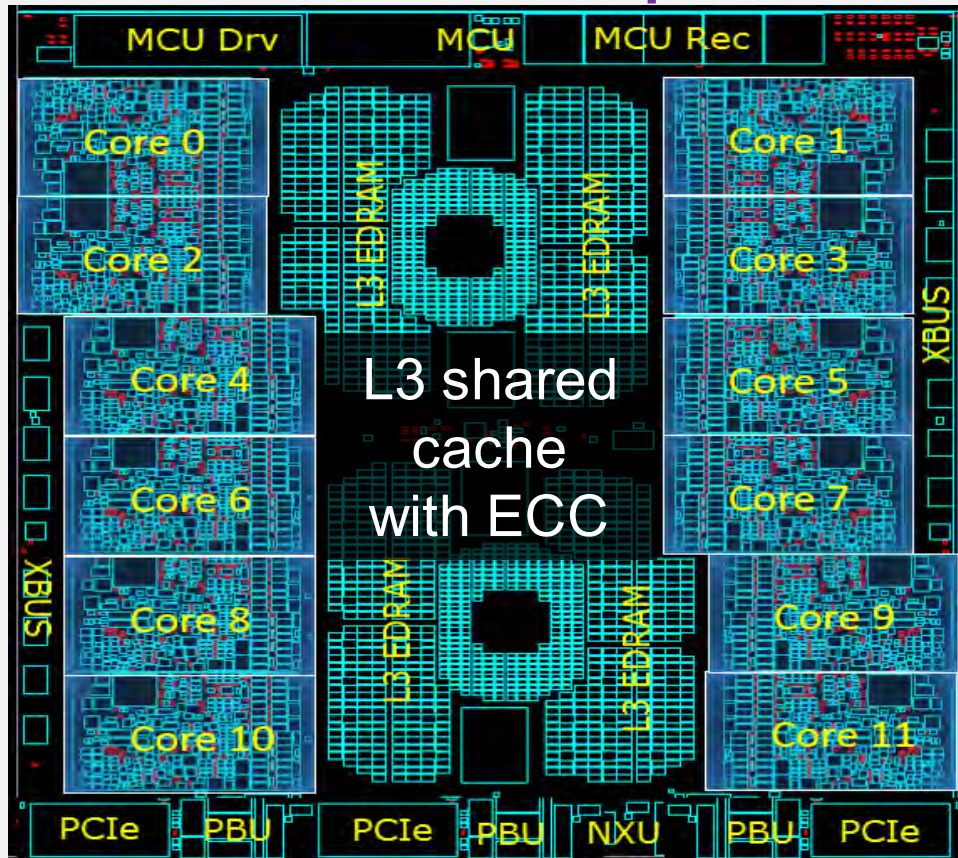 - Connects to the PCIe Interconnect Gen3 in the PCIe+ I/O drawers

SMP2   SMP3   SMP0   SMP1

PSU1   PSU2   PSU3   PSU4

Power from iPDU/BPA

SMP Connection to other CPC Drawers

5x FANS

FSP / OSC   FSP / OSC

BDC   BDC

Front

ICA SR 1.1

PCIe

PCIe

ICA SR1.1 Coupling Links

PCIe+ Fanout Gen3

PCIe

PCIe

To PCIe+ I/O Drawers

* 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.
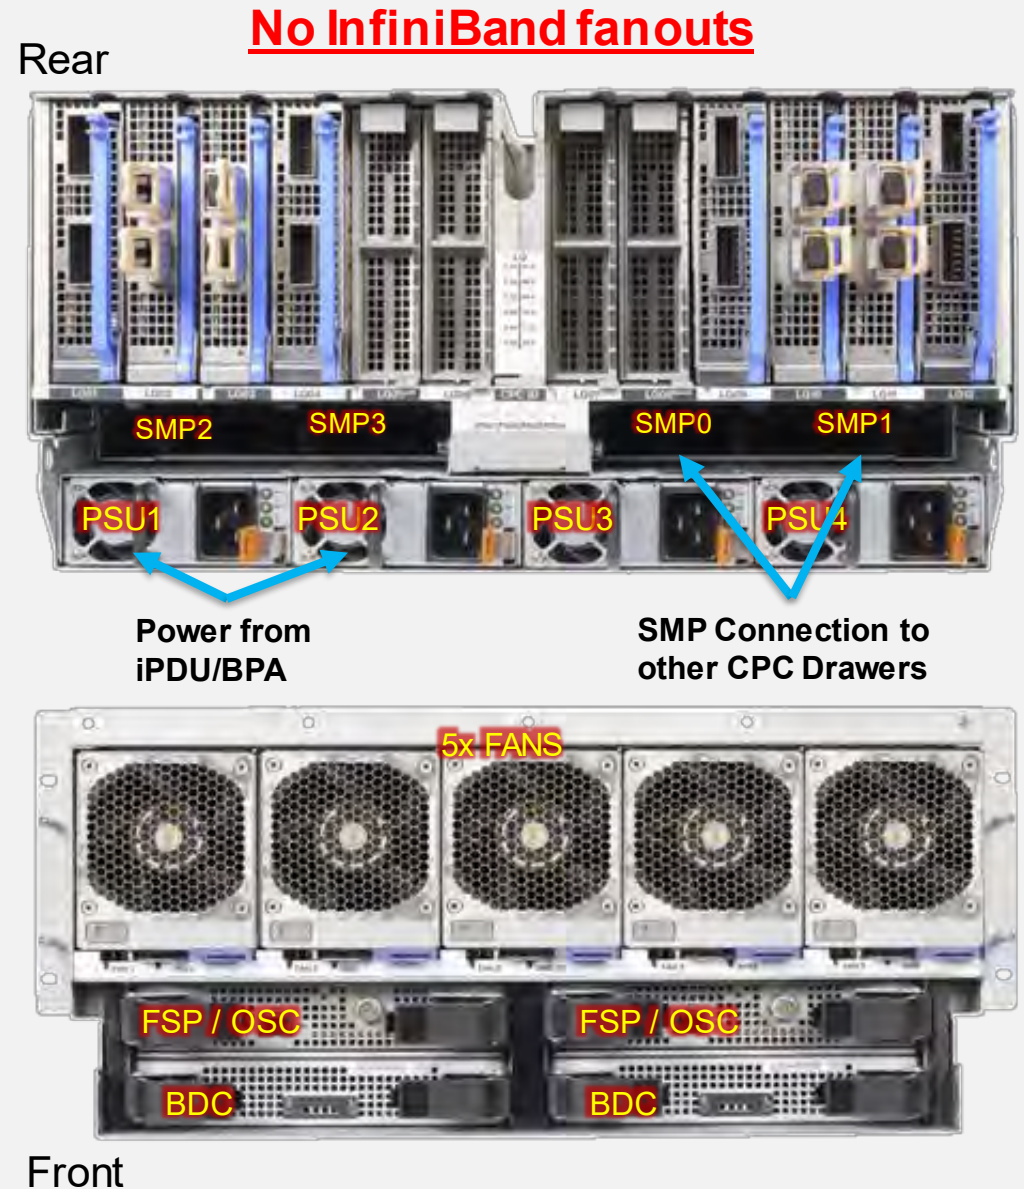
12-Core Processor Chip Detail

- **5.2 GHz core frequency**
- **8, 9, 10, 11 or 12 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

- **14nm SOI Technology**
  - 12 Cores
  - 17 layers of metal
  - 696 mm2 chip area
  - 9.2B transistors versus 6.2B on z14

- **20% reduction area**
- **20% reduction in power**

- **I/O buses**
  - Each CP chip will support up to 3 PCIe buses
    - PCIe+ I/O Drawer Fanout
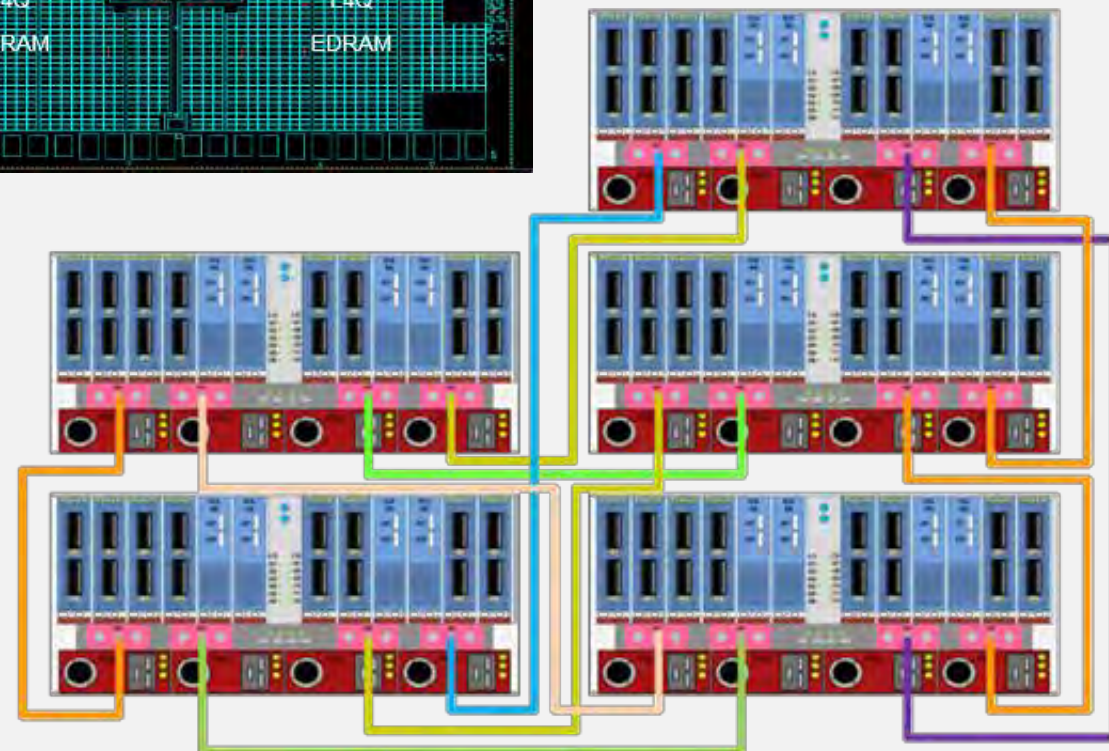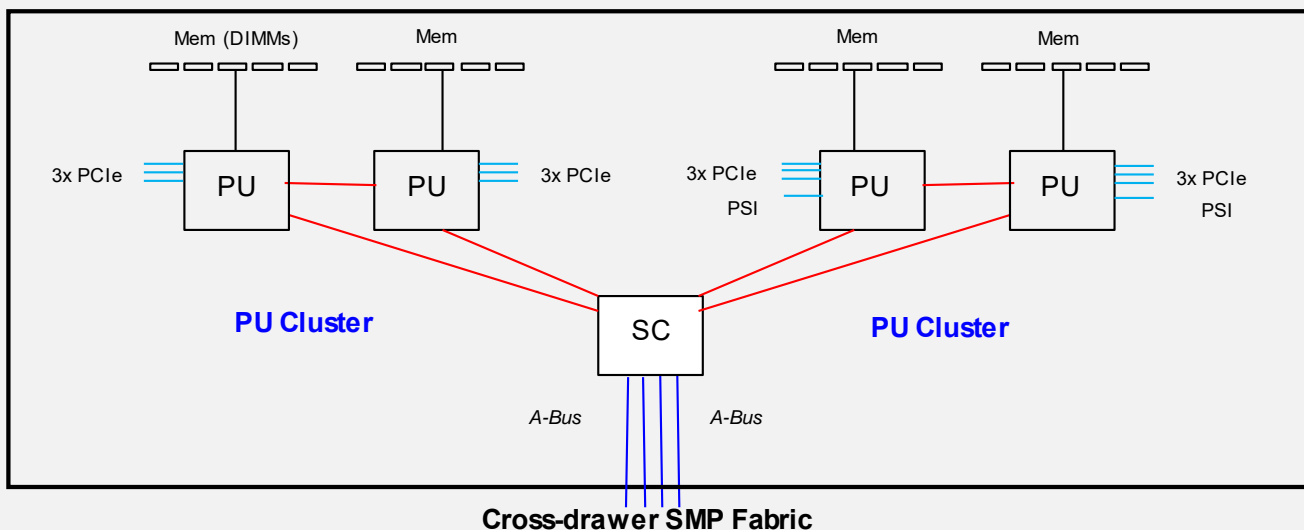   - ICA SR 1.1 Coupling Links

z15 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
  - 5 drawers using 4 A-buses (point-to-point).
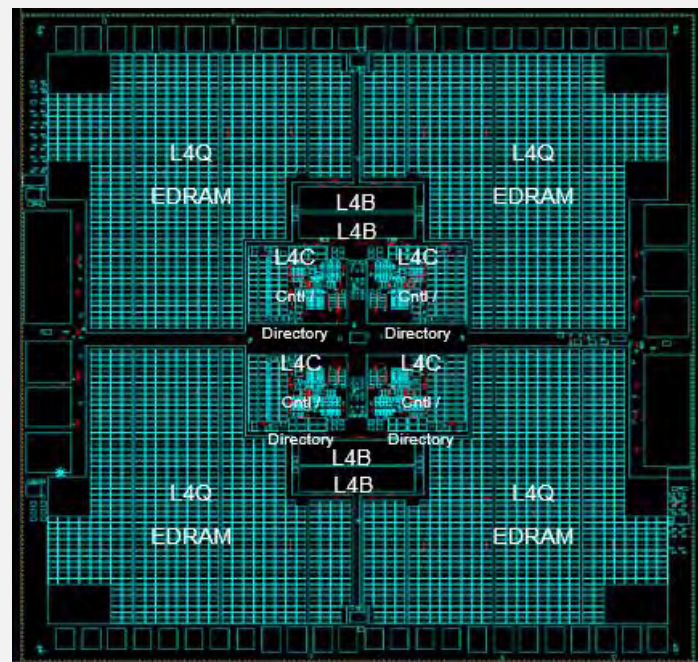

Fully Populated Drawer

Cross-drawer SMP Fabric

z15 Capacity Considerations

Feature	Feature Code	Drawers/ Cores	CPs	IFLs uIFLs	zIIPs	ICFs	Std SAPs	Optional SAPs	Std. Spares	IFP	Memory
Max34	0655	1/41	0-34	0-34 0-33	0-22	0-34	4	0-8	2	1	8 TB
Max71	0656	2/82	0-71	0-71 0-70	0-46	0-71	8	0-8	2	1	16 TB
Max108	0657	3/123	0-108	0-108 0-107	0-70	0-108	12	0-8	2	1	24 TB
Max145	0658	4/164	0-145	0-145 0-144	0-96	0-145	16	0-8	2	1	32 TB
Max190	0659	5/215	0-190	0-190 0-189	0-126	0-190	22	0-8	2	1	40 TB

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 Dynamic I/O for Standalone Coupling Facility.**
5. **The Max142 and Max 190 is factory build only**

MES/Model Considerations

z13

z15

Max190
Max145
Max108
Max71
Max34

Factory Only

Concurrent Upgrade

z14
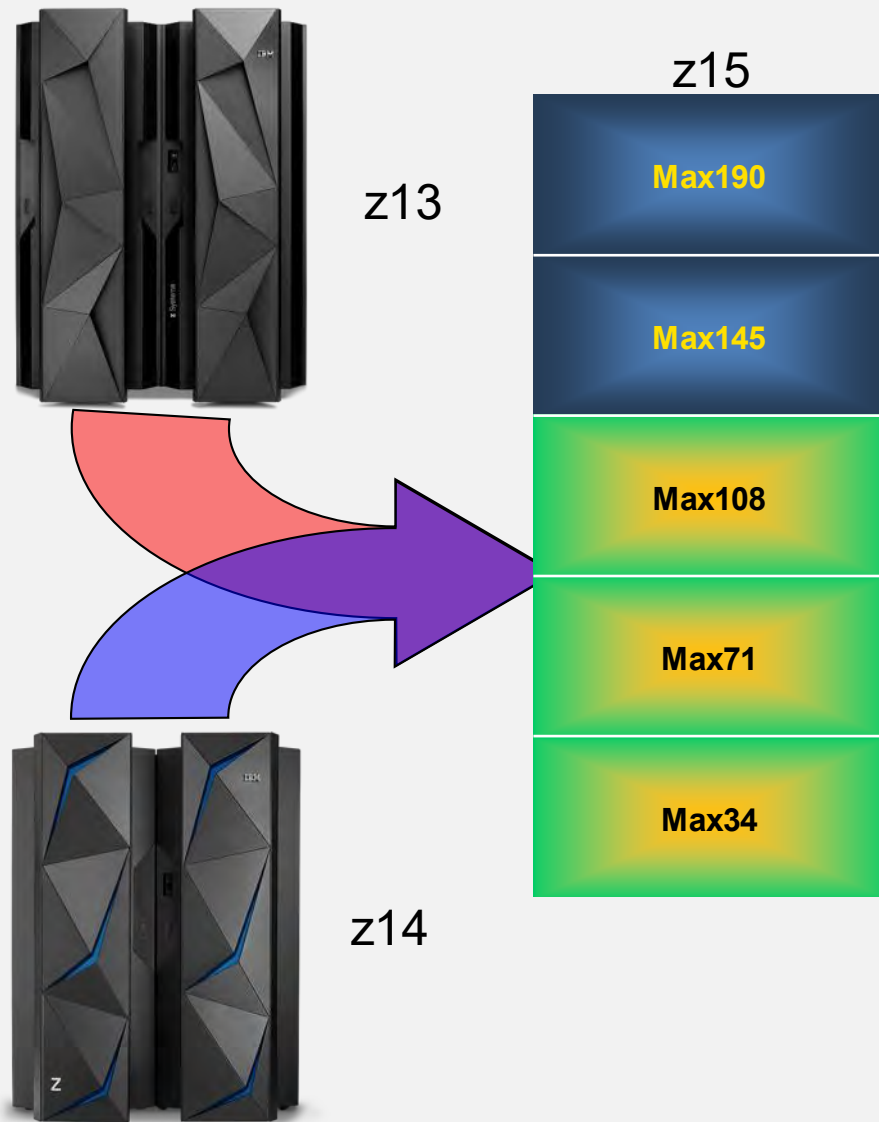
- **z15 to z15 upgrades**
  - z15 Concurrent upgrade from Max 34 to Max71 to Max108
    - Each max level adds a CEC drawer
    - No MES upgrade to Max145 or Max190
  - Additional I/O Drawers
    - Based on available space in current frames and/or I/O expansion frames
  - *No conversion available between power types*
- **Any z13 to any z15**
- **Any z14 M01- M05 to any z15**
- **No LinuxONE model conversions to LinuxONE III**
- **LinuxONE III to z15 MES available**

z15 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 z13 upgrade

- Less power consumption

- During z13 upgrade, Feature Conversion for IBM Flash Express

- z14 offered 4 VFM features at 1.5 TB per feature for a total of 6TB

- **Increment Size**

    - Up to twelve features/increments

        - **0.5 TB | 1.0 TB | 1.5 TB | 2.0 TB**

        - **2.5 TB | 3.0 TB | 3.5 TB | 4.0 TB**

        - **4.5 TB | 5.0 TB | 5.5 TB | 6.0 TB**

12 TB

| 10 TB normal use | 2 TB VFM |
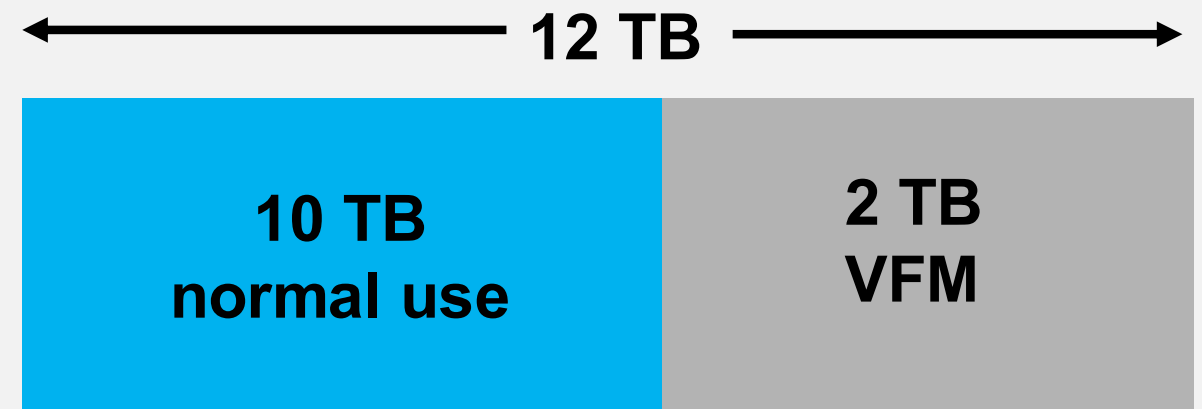
z15 New Features and Functionality

Integrated Accelerator for zEDC – z15

	IBM z14™ with zEDC Express	z15 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 – 16 GB/s	Compress up to 260 GB/sec with the Integrated Accelerator for zEDC on the largest IBM z15.**
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	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

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
   a) Using the client's already-entitled GPs and zIIPs
   b) System Recovery Boost Record
   - CoD like record activating unused/available processor resources on the machine to provide additional zIIP processors and capacity. (Priced Feature)

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

- 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, 60 minutes for IPL, and 6 Hours for System Recovery Boost Record.

- All this without increasing the 4HRA IBM software billing cost or MSU consumption costs

* Planned future support

IBM®

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

IBM Fibre Channel Endpoint Security enables FICON® or
FCP Links from the IBM z15 to the IBM DS8900F storage family to be encrypted and protected

Challenges	Client value

Challenges

- Encrypt all data in-flight by corporate directive.

- Protect the integrity and confidentiality of data in-flight.

Client value

- Gain confidence that all data flowing within and across datacenters 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.

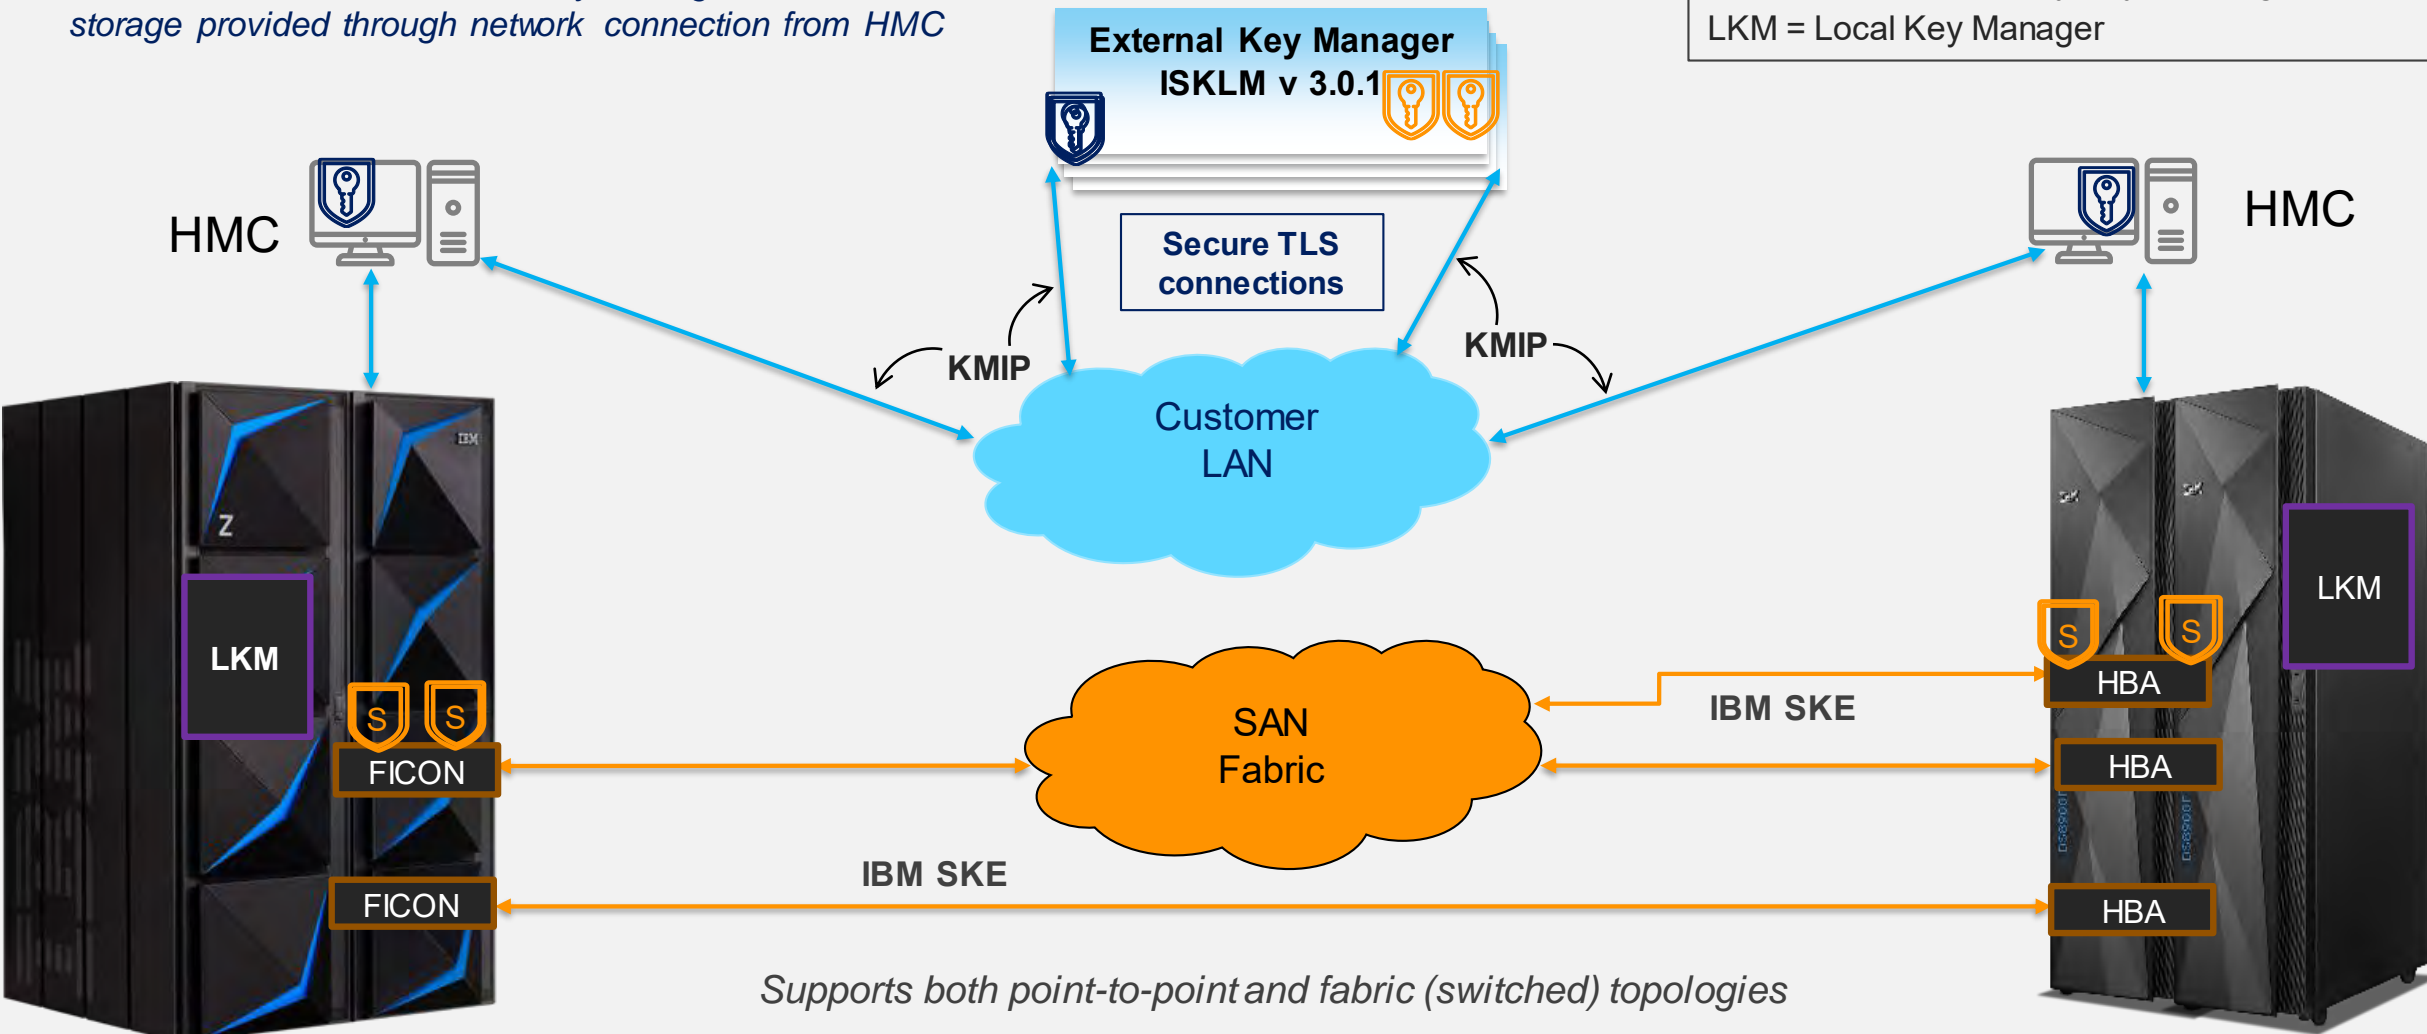- Use on all IBM Z operating systems.

  z/OS   Linux   z/TPF   z/VM   z/VSE

- Reduce and eliminate insider threats of unauthorized access to data in-flight.

High level diagram and solution components

EKM in a Multi-Master group of 2-4 ISKLM instances

Communication with External Key Manager on both Z and storage provided through network connection from HMC

External Key Manager ISKLM v 3.0.1

KMIP = Key Management Interoperability Protocol

IBM SKE = IBM Security Key Exchange

LKM = Local Key Manager

HMC

Secure TLS connections

HMC

KMIP

KMIP

Customer LAN

LKM

LKM

SAN Fabric

IBM SKE

FICON

HBA

HBA

IBM SKE

FICON

HBA

Supports both point-to-point and fabric (switched) topologies

IBM z15

IBM DS8900F

IBM Z Hardware Management Appliance

HMCs

- z15 introduces Hardware Management Appliance
  - HMC & SE packaged redundantly inside A frame
  - Eliminates the need for managing separate HMC boxes outside of CPC package
  - No change in general user experience (most use cases involve remote browser access)
  - Can be used to manage N-2 systems (z13/z13s, z14), not just z15
  - **NO Standalone HMCs are needed for this environment ***

Note: z15 HMC and Hardware Management Appliance support ONLY N-2 Servers

Client LAN

Remote Browser (HMC access)

| SE | | | SE | | | SE | | | SE | HMC VA |
| SE | | | SE | | | SE | | | SE | HMC VA |

z13

z14

z15

z15
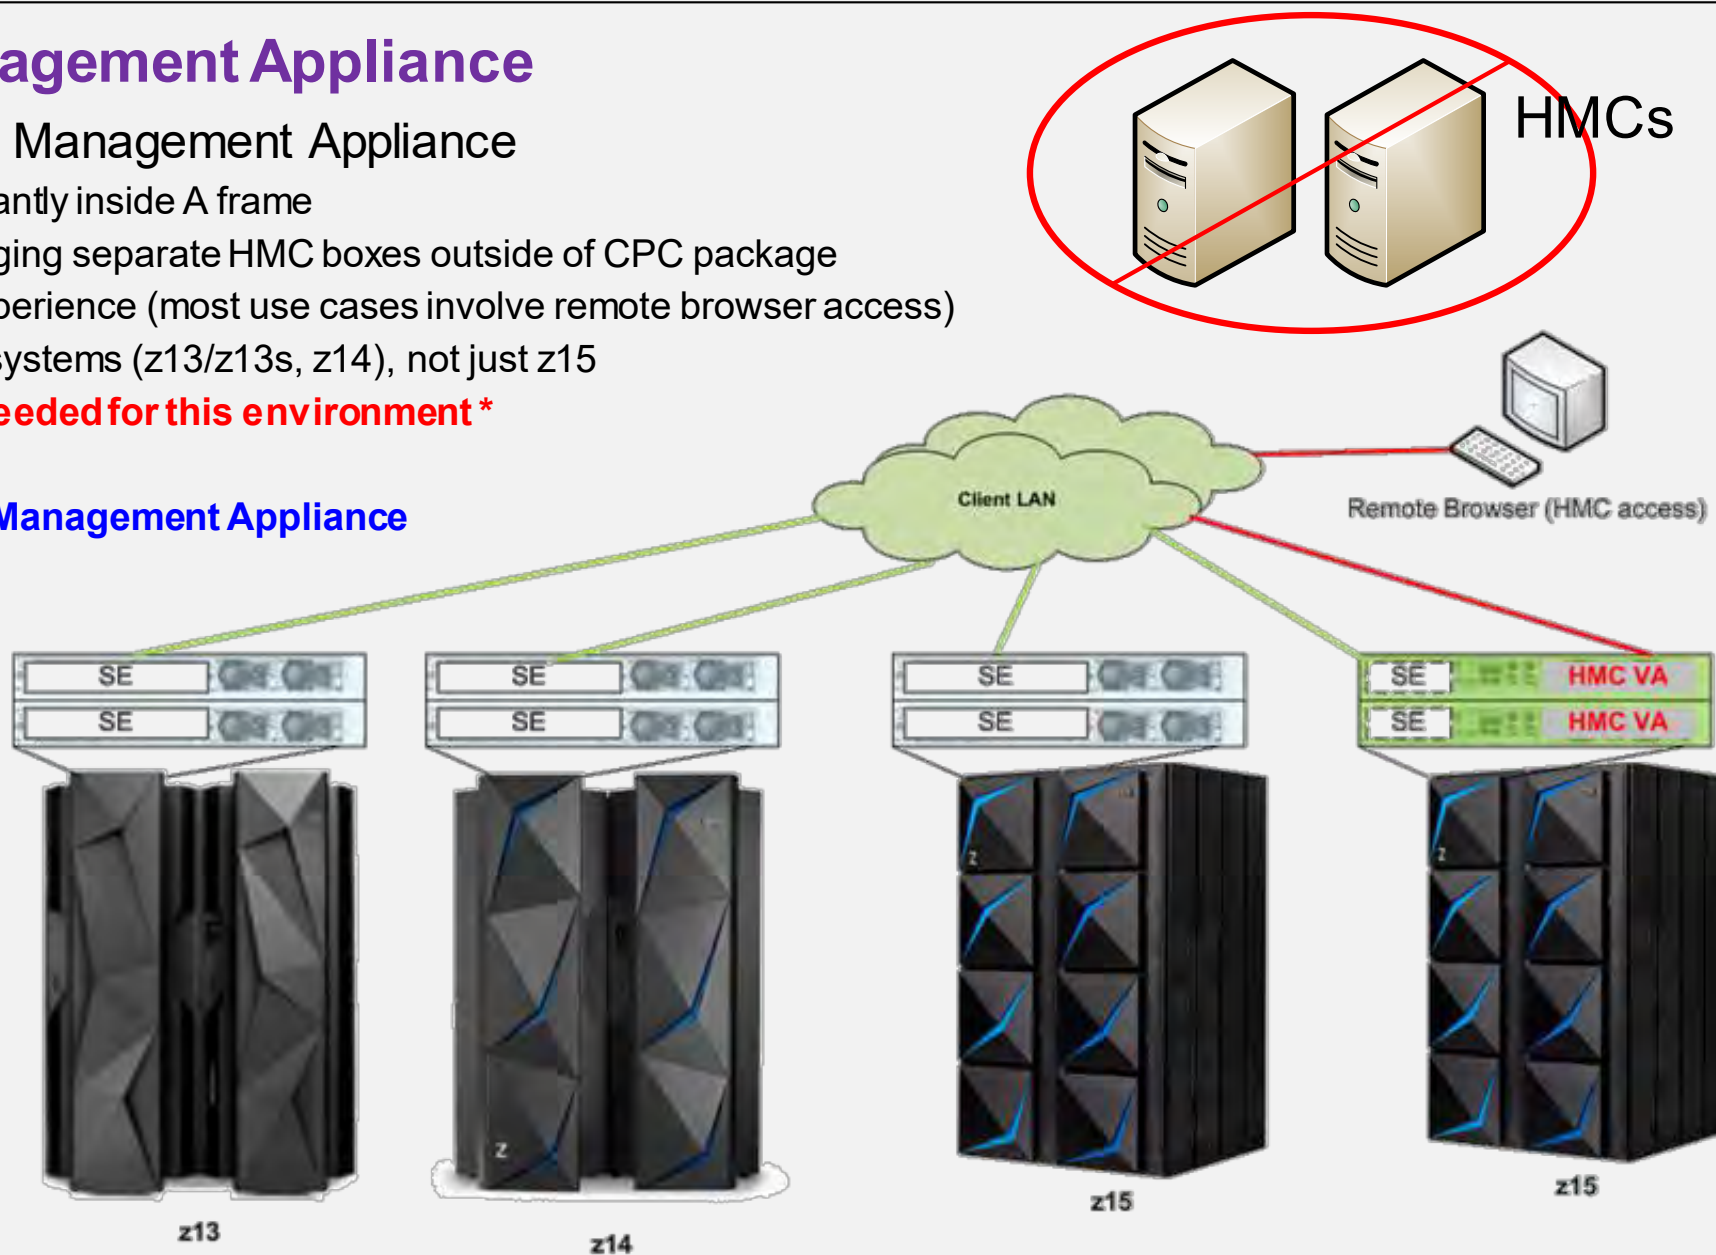
Protect sensitive data in dumps

Today:

In sending data in dumps to vendors, clients risk accidentally sharing sensitive data—putting themselves at risk in more ways than one.

Organizations are forced to make a choice between regulatory compliance and serviceability.

With z15:

With peace of mind, knowing that data in dumps will be appropriately protected, clients can more easily collaborate with vendors to fix major issues.

This solution would ensure that open problems can be addressed without fear of sensitive data exposure.

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
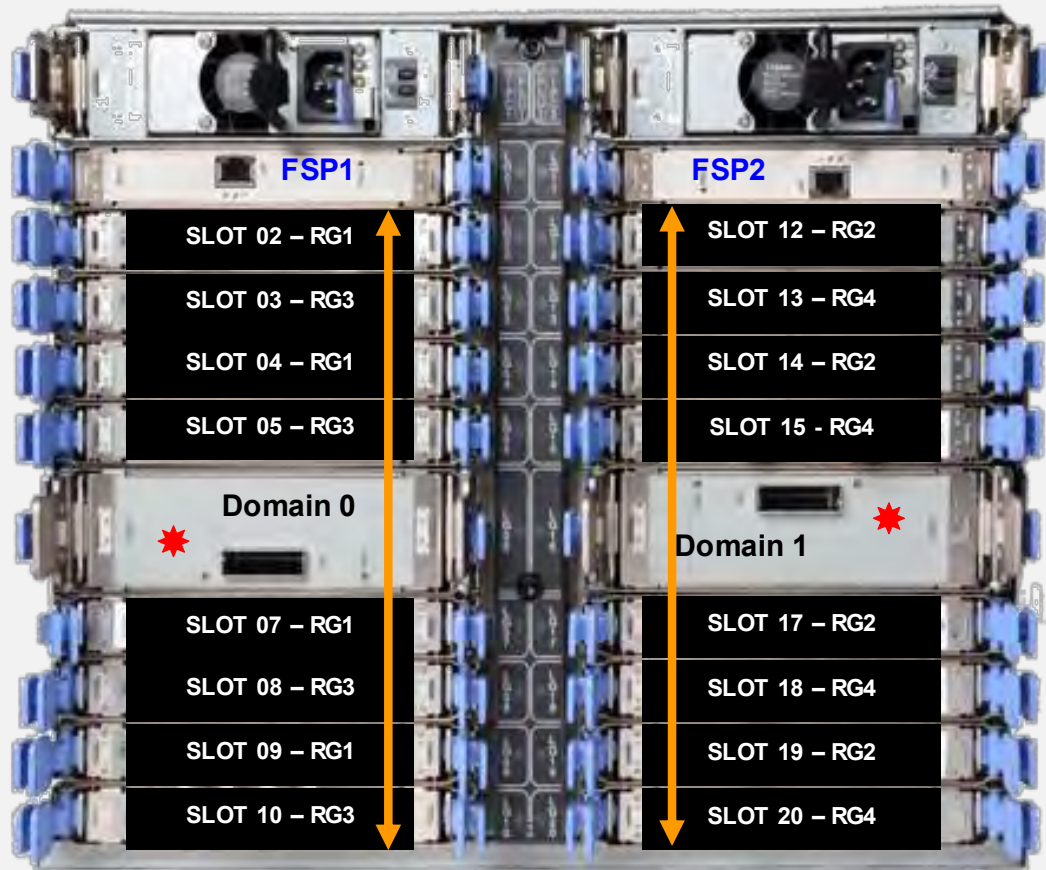 - PTP will be provided via the Support Element

POLL

What new feature are you most interesting in learning more?
a)   Integrated Accelerator for zEDC
b)   System Recovery Boost
c)   IBM Fibre Channel Endpoint Security
d)   Precision Time Protocol
e)   Data Privacy for Diagnostics

z15 I/O Infrastructure

PCIe+ I/O Drawer – 16 slots

FSP1

FSP2

SLOT 02 – RG1
SLOT 03 – RG3
SLOT 04 – RG1
SLOT 05 – RG3
Domain 0
SLOT 07 – RG1
SLOT 08 – RG3
SLOT 09 – RG1
SLOT 10 – RG3

SLOT 12 – RG2
SLOT 13 – RG4
SLOT 14 – RG2
SLOT 15 - RG4
Domain 1
SLOT 17 – RG2
SLOT 18 – RG4
SLOT 19 – RG2
SLOT 20 – RG4

8U

- **Supports PCIe I/O cards**
  - First introduced on the z14 ZR1/Rockhopper II
  - PCIe+ I/O drawers locations are dependent on power type (BPA or iPDU) and CPC drawer count.

- **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).**

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

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

zHyperLink Express

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
    - OSA-Express7s 25 GbE   z14 FC 0429 / z15 FC 0449
   - zHyperlink Express   z14 FC 0431 / z15 FC 0451

z15 I/O Features (new build)

- **FICON Express16SA**
  - FC 0436, 0437
- **zHyperLink Express1.1** 🚫🐧
  - FC 0451
- **OSA Express7S:**
  - 25 GbE SR1.1, GbE (LX, SX) 10GbE (LR, SR) , 1000BASE-T; FC 0442, 0443, 0444, 0445, 0449, 0446
- **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
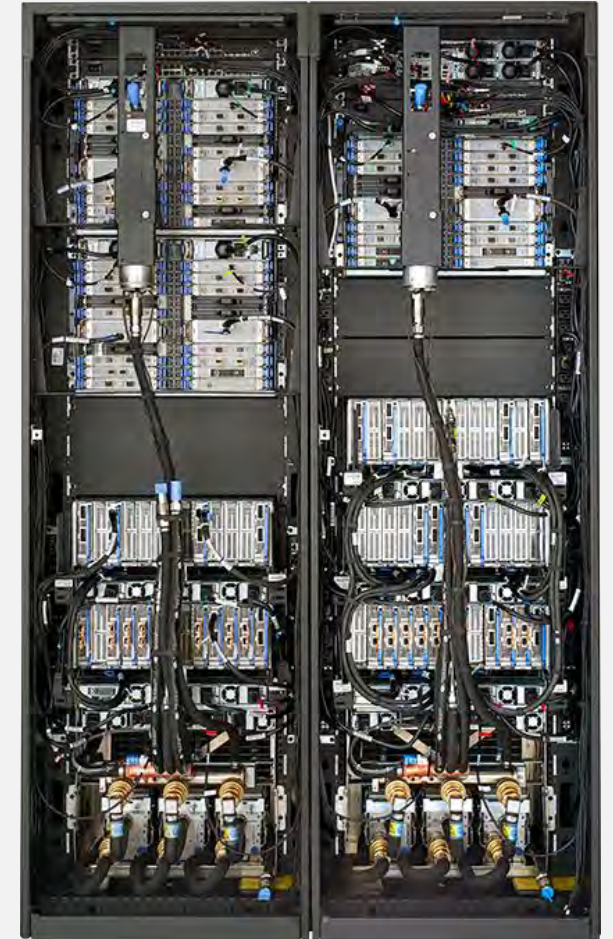
🚫🐧 z15 Only

🐧 LinuxONE III Only

No InfiniBand Coupling Links

z15 I/O Features – Carry Forward

- **FICON Express16S+**
  - **FC 0427, 428**
- **FICON Express16S**
  - **FC 0417, 0418**
- **FICON Express8S**
  - **FC 0409, 0410**
- **OSA-Express7S 25GbE SR**
  - **FC 0429**
- **OSA-Express6S**
  - **FC 0422, 0423, 0424, 0425, 0426**
- **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 0433**
- **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

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*
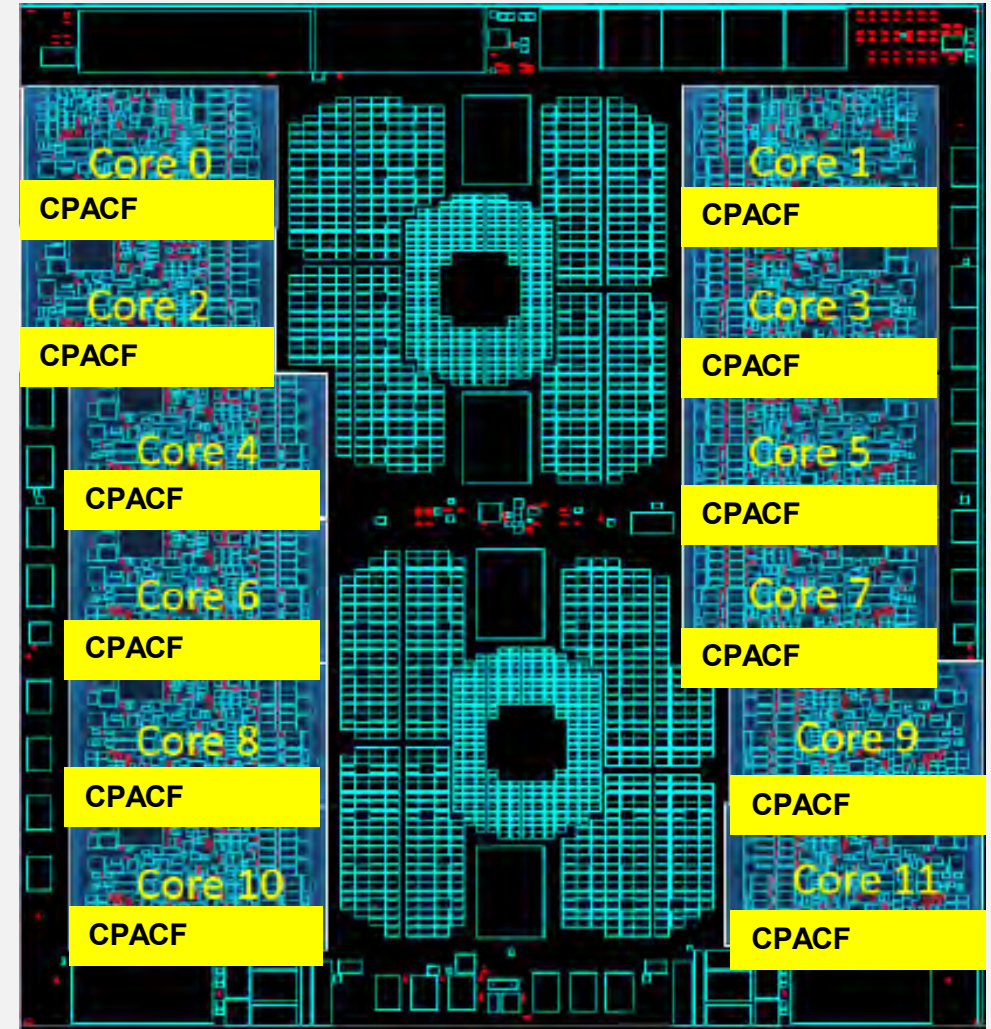
Crypto Express7S

Two new cards designed for z15

- FC 0899 Single HSM

  – Max 16 per server

- FC 0898 Dual HSM

  – Max 8 per server

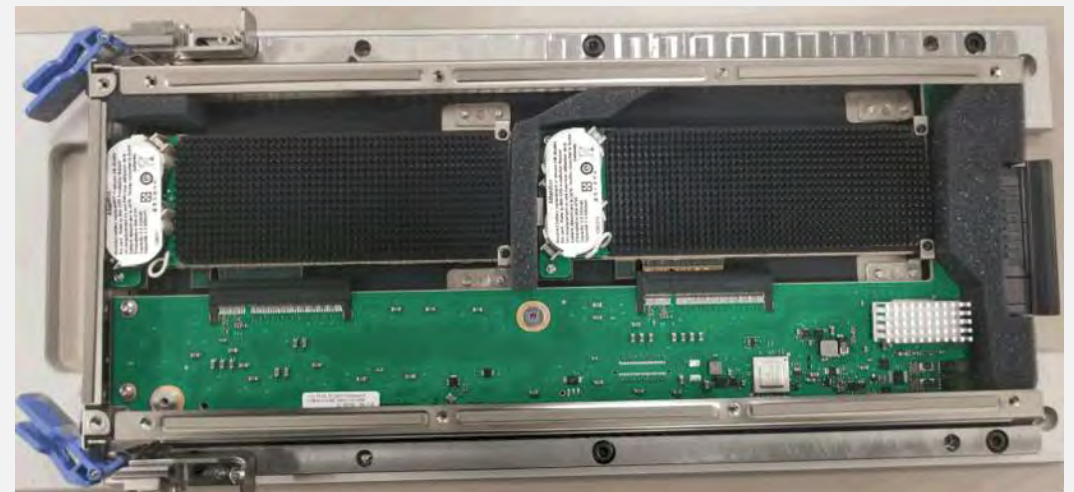## A mix of Crypto cards can be ordered for both new build and carry forward

- Max combined total: 16 HSMs

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
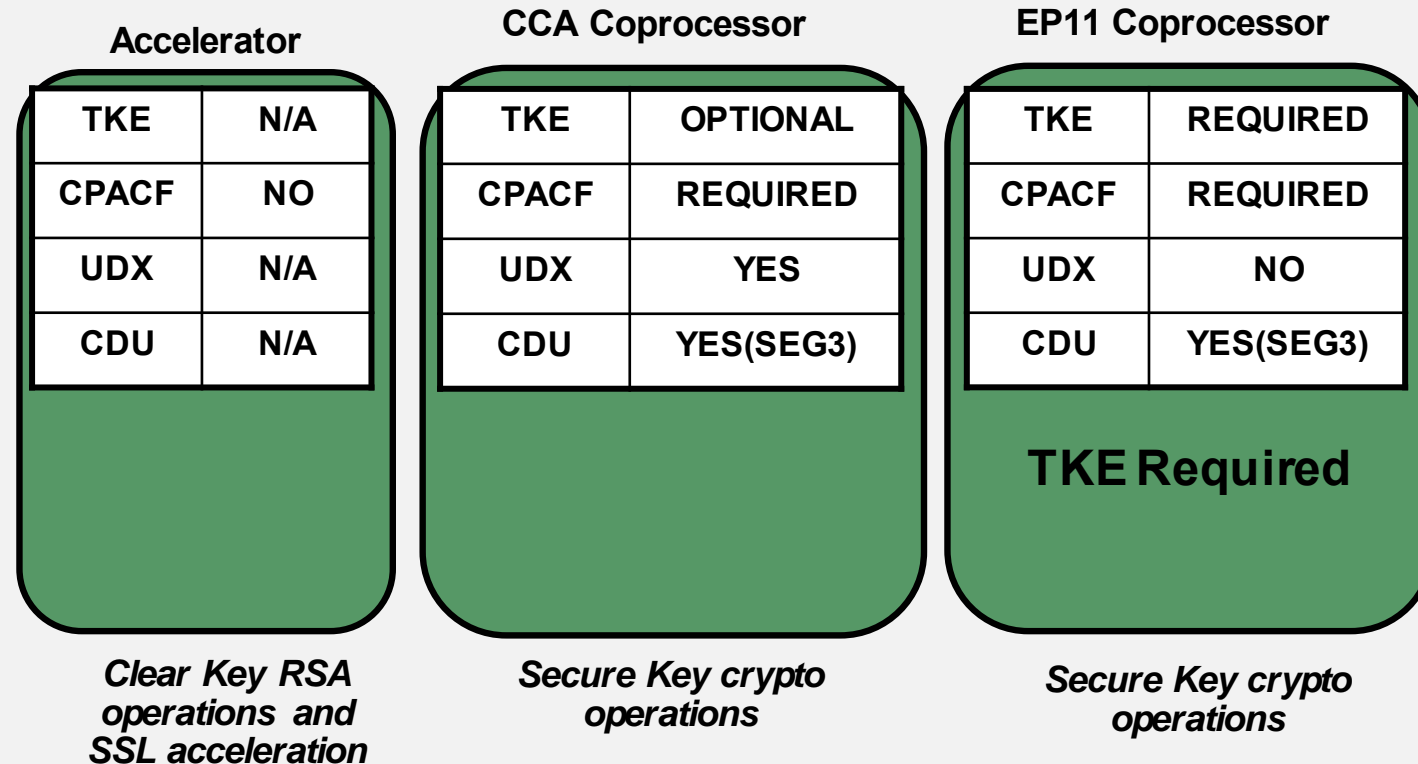
  - **Dilithium (Quantum Safe)**

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

Accelerator

TKE	N/A
CPACF	NO
UDX	N/A
CDU	N/A

Clear Key RSA operations and SSL acceleration

CCA Coprocessor

TKE	OPTIONAL
CPACF	REQUIRED
UDX	YES
CDU	YES(SEG3)

Secure Key crypto operations

EP11 Coprocessor

TKE	REQUIRED
CPACF	REQUIRED
UDX	NO
CDU	YES(SEG3)

TKE Required

Secure Key crypto operations

z15 Operating Systems

z15 operating system support

z/OS

- z/OS 2.4 with PTFs
- 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

- RHEL 7.6
- RHEL 8.0
- SLES 12.4
- SLES 15.1
- Ubuntu 16.04 LTS
- Ubuntu 18.04 LTS

- **Minimum Distributions**
- SLES 15.1
- SLES 12.4
- SLES 11.4
- RHEL 8.0
- RHEL 7.6
- RHEL 6.10
- Ubuntu 16.04 LTS
- Ubuntu 18.04 LTS

IBM cannot legally discuss z15 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.

Replay availability!

Kenny Stine

IBM Z Technical Specialist

Washington Systems Center

kjstine@us.ibm.com

ibm.com

- **Please feel free to send Questions you were unable to ask in the live call to my email**

- **The replay, presentation and Q&A transcript will be available on the TechU Talks page within 48 hours. https://www.ibm.com/services/learning/events/techutalks**

- **To follow TechU on Linked in, join the Linkedin Group : https://www.linkedin.com/groups/8518569/**

Notices and disclaimers

— © 2019 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.

— **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**

— Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed "as is" without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.

— IBM products are manufactured from new parts or new and used parts.
In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply."

— **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**

— Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those

— customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

— References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

— Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

— It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Notices and disclaimers continued

— Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of 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. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**

— The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

— IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml

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.

Thanks

FICON Express16SA

- **For FICON, zHPF, and FCP**
  - FC 0436 (LX) & 0437 (SX)
  - CHPID types: FC and FCP
    - Two PCHIDs/CHPIDs
    - NO mixed CHPIDs for same card – only FC or FCP

- **Same performance as FICON Express16S+**

- **Auto-negotiates to 8 or 16 Gbps**
  - Negotiation to 4 Gbps NOT supported
  - 2 and 4 Gbps supported through a switch with 8 or 16 Gbps optics

- Max. 192 features per system

- **Concurrent repair/replace of small form factor pluggable (SFP) optics**
  - Port components can be replaced instead of the entire adapter.
  - 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
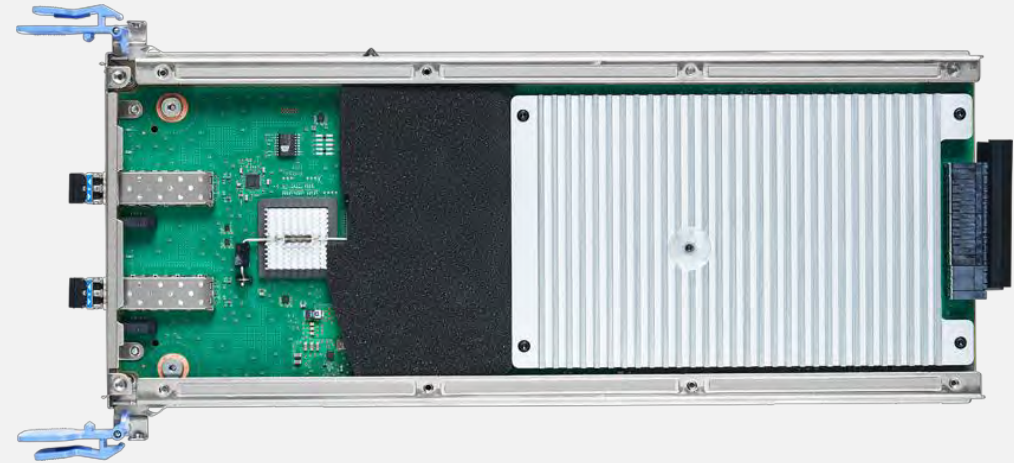
What is IBM zHyperLink™?

- zHyperLink Express is a direct connect short distance
  IBM Z I/O feature designed to work in conjunction with a FICON or High
  Performance FICON SAN infrastructure

- zHyperLink improves application response time, cutting I/O sensitive
  workload response time in half without significant application changes.

- Current Support
  - Db2 V11/V12 Read Support
  - Db2 V12 Active Log Support
  - VSAM Read Support
  - Db2 log writes with Metro Mirror support
    - The zHyperlink Write capability supports Metro Mirror and HyperSwap®
      within 150 meter distances.
    - Requires PTFs for z/OS and Db2®.
    - Not supported asynchronous replication (Global Mirror, zGlobal Mirror (XRC),
      Global Copy, Safeguarded Copy), and DS8882/F

IBM z15 CPC

Up to 5x reduction In latency

IBM DS8880

OSA Express7S 25 GbE SR1.1 – FC0449

- 25 Gigabit Ethernet (25 GbE)
  - 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

# OSA-Express7S Fiber Optic Features

- 10 Gigabit Ethernet (10 GbE)
  - CHPID types: OSD
  - 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
  - 10GbE does NOT auto-negotiate to a slower speed.
  - Up to 48 features per system (48 ports)

FC 0444 – 10 GbE LR | FC 0445 – 10 GbE SR

- Gigabit Ethernet (GbE)
  - CHPID types: OSD
  - 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
  - LC Duplex
 - Up to 48 features per system (96 ports)
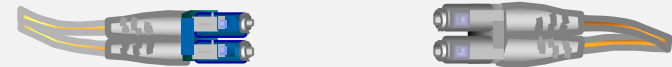
FC 0442 – GbE LX | FC 0443 – GbE SX
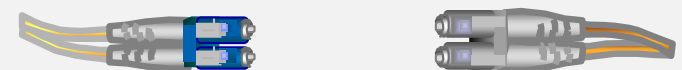
OSA-Express7S 1000BASE-T

- ## 1000BASE-T Ethernet (1 GbE)

  - Copper Wiring
  - Two ports with RJ-45 connector
    - 1 PCHID/CHPID
  - Small form factor pluggable (SFP+) transceivers
  - Concurrent repair/replace for each SFP transceiver
  - 1000 Mbps (1 gbps full duplex) – NO negotiation to lower speeds
  - Up to 48 features per system (96 ports)
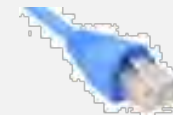
Connector = RJ-45

Operation Mode	CHPID 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
Dynamic Partition Manager	OSM	DPM Management

10 and 25 GbE RoCE Express2.1

Description	Feature Code	Ports	Max. Features per system (z15)
25GbE RoCE Express2.1	0450	2	**16 (32 ports)**
10GbE RoCE Express2.1	0432	2	**16 (32 ports)**

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)

LinuxOne Only

FCP Express32S

- ## For FCP Only
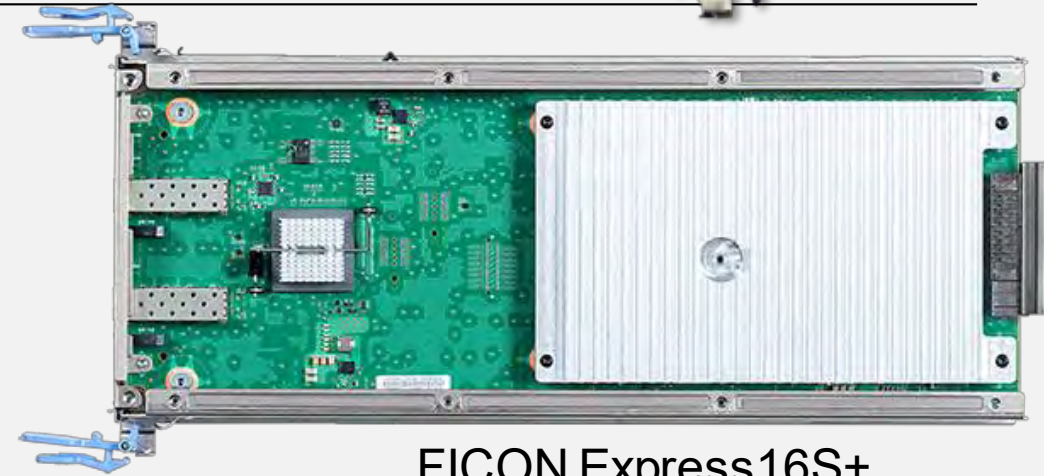 - FC 0438 LX & 0439 SX
  - CHPID types: FCP
    - 2 PCHIDs/CHPIDs

- ## Auto-negotiates to 8, 16, or 32 Gbps

FICON Express16S+

- ## Concurrent repair/replace of small form factor pluggable (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

- ## These features will only be available on the LinuxONE machines.

IBM Adapter for NVMe1.1 – FC0448

- **LinuxONE only**

- **"Built in" storage. No boot support initially.**
- **Uses the normal INext PCIe EC Stream.**
- **Carrier Card**
  - Zero ports
  - IBM provides a <u>carrier card</u> 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.**

- **Tested in IBM Z – we will not make support statements just testing statements**
  - Intel PN SSDPE2KX010T701 (1TB) – Up to 16 TB
  - Intel PN SSDPE2KX040T701 (4TB) - Up to 64 TB
    - Both can coexist on the same system and same I/O Drawer.

- **Details can be found in the IMPP GC28-7002. Performance testing found <u>here</u>**
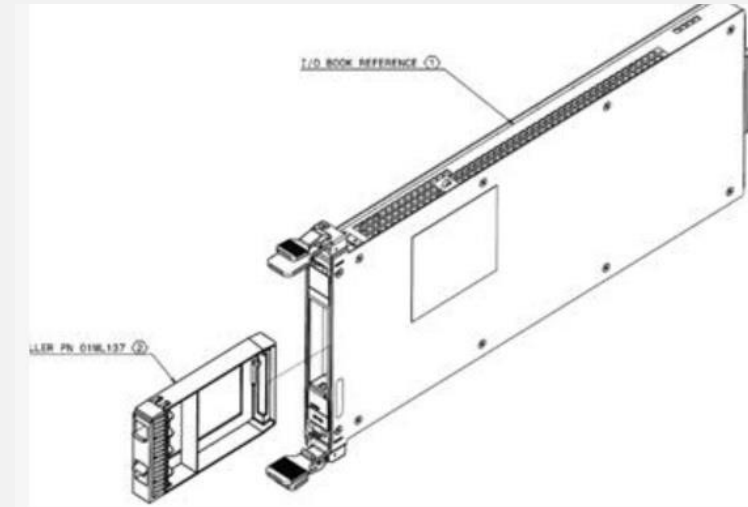
Features	New Build (NB) Carry Forward (CF)	Maximum Features	Ports	Increment
FICON Express16SA-LX and -SX	NB	192	2	1
FICON Express16S+	CF	192	2	1
FICON Express16S	CF	192	2	1
FICON Express8S	CF	192	2	1
FCP Express32S	NB	192	2	1
OSA Express7S	NB	48	2/1	1
OSA Express7S 25GbE SR1.1	NB/CF	48	1	1
OSA Express6S	CF	48	2/1	1
OSA Express5S	CF	48	2/1	1
Crypto Express7S (2 adapters)	NB	8	2 Adapters	2, 1 thereafter
Crypto Express7S (1 adapter)	NB	16	1 adapter	2, 1 thereafter
Crypto Express6S	CF	16	1 adapter	2, 1 thereafter
Crypto Express5S	CF	16	1 adapter	2, 1 thereafter
25GbE RoCE Express2.1	NB/CF	16	2	1
10GbE RoCE Express2.1	NB/CF	16	2	1
10GbE RoCE Express	CF	16	2	1
NVMe Express	NB	16	1	1
zHyperlink Express1.1	NB/CF	16	1	1
Virtual Flash Memory	NB	12 Units	0.5 TB	1
Integrated Coupling SR1.1	NB/CF	48	2	1
Coupling Express2 LR	NB/CF	32	2	1