GSE GUIDE SHARE EUROPE UK REGION

GSE UK Conference 2019
Dock into the Dark Side

IBM z15 Model T01 Hardware Overview Part 1

Parwez Hamid
Consultant
parwez_hamid@hotmail.com

7 November 2019
Session BO

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.

Acknowledgements and Notes:

- The content of this presentation was created by IBM Redbooks and the Washington Systems Centre. Modified by Parwez Hamid.

- It's assumed that the audience has knowledge of prior generations of IBM Z Servers.

- Always refer to the appropriate sources for the latest information e.g. IMPP (GC28-7002) Resource Link, Exception Letters, PSP Device buckets, Must Read section of the TDA-SA Guides etc.
- CAUTION: A number of the screen captures are from development and are included for illustration purposes ONLY. Announced versions of the product may be different.

- **Disclaimer: Any numbers quoted for performance/measurements were collected in a controlled environment running an IBM developed workload under z/OS. Individual customer results may vary. Results are workload dependent.**

Table of Content

z15 Announcement Dates

IBM z15 availability Dates – Driver Level 41

- **Announced – September 12, 2019**

- **General Availability – September 23, 2019**
  - New features and functions for the IBM z15 (Type number: 8561)

- **November 14, 2019 - *orders cannot be placed until November 14, 2019***
  - IBM Z Hardware Management Appliance (#0100): HMC/SE housed in one physical server inside CPC frame
  - Dynamic Partition Manager (DPM) on IBM z15 (#0016)
  - zTPF exploitation of System Recovery Boost

- **January 29, 2020**
  - All remaining MES orders for IBM z15 Model T01

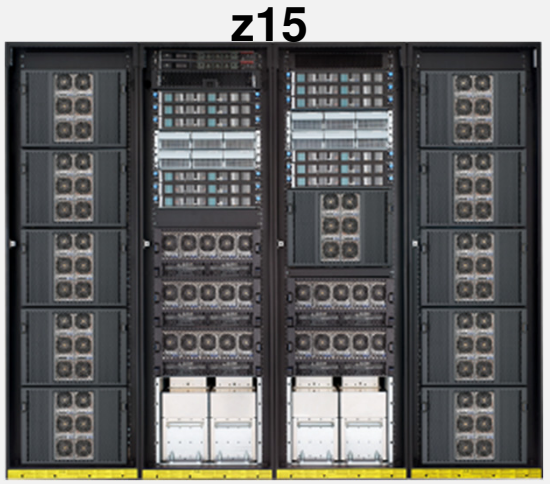

IBM z15 Redboooks

- **New Redbooks – Draft Versions**
  - IBM z15 Technical Introduction, SG24-8850
  - IBM z15 Technical Guide, SG24-8851
  - IBM Z Connectivity Handbook, SG24-5444
  - IBM Z Functional Matrix, REDP-5157
  - IBM z15 Configuration Setup, SG24-8860

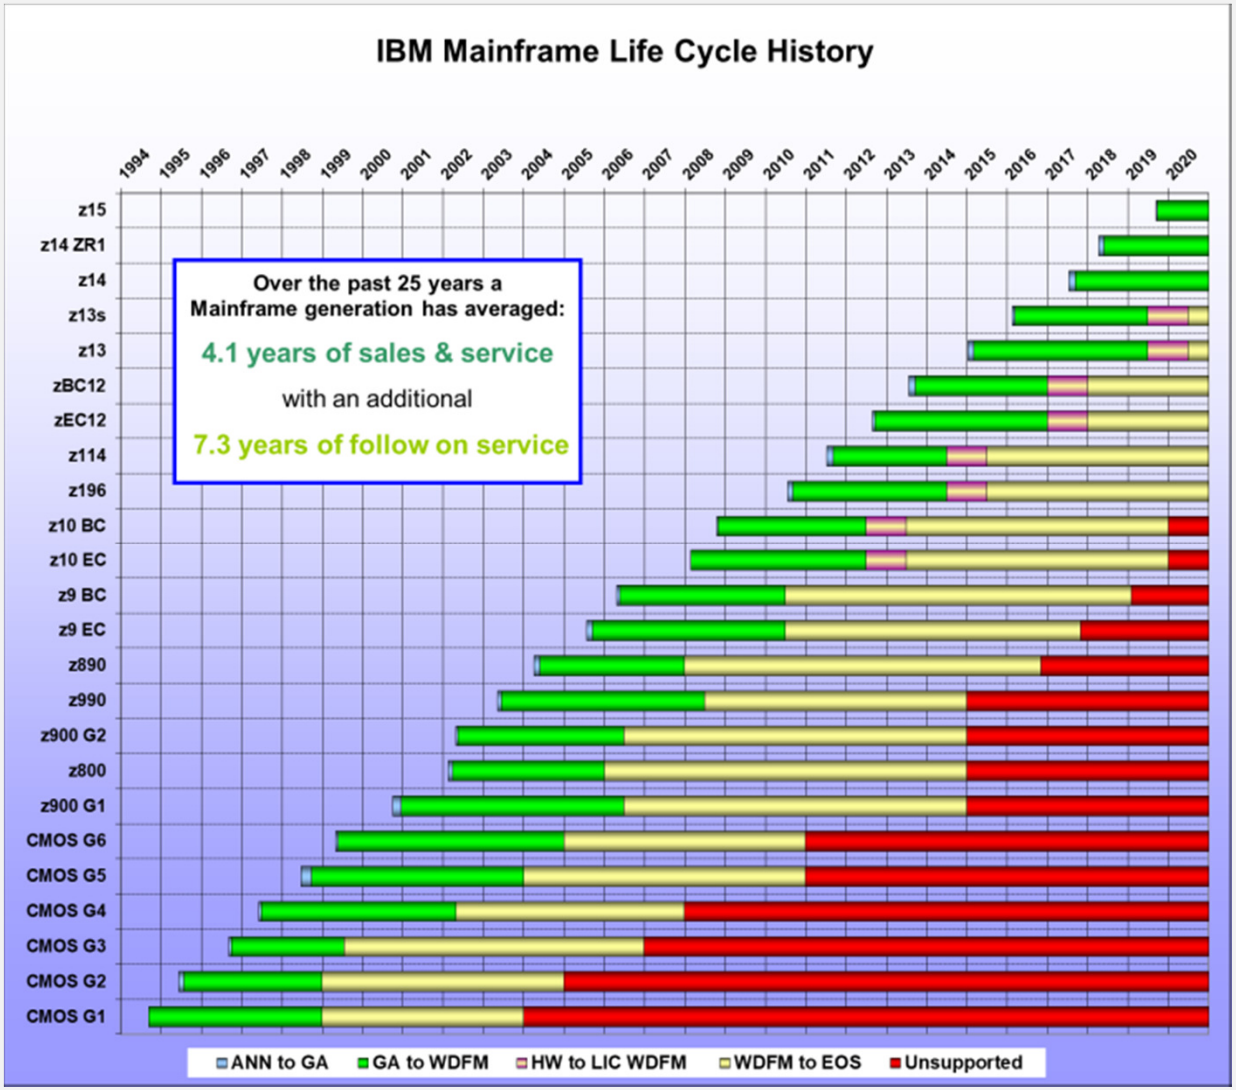

IBM Mainframe Life Cycle History

Machine			Dates					Years		
Type	Model	Family	ANN	GA	HW WDFM	LIC WDFM	EOS	ANN to GA	GA to HW WDFM	HW WDFM to EOS
8561	T01	z15	September 12, 2019	September 23, 2019				0.0		
3907	ZR1	z14 ZR1	April 10, 2018	May 31, 2018				0.1		
3906	M0n	z14	July 17, 2017	September 13, 2017				0.2		
2965	Nnn	z13s	February 16, 2016	March 10, 2016	June 30, 2019	June 30, 2020		0.1	3.3	
2964	Nnn	z13	January 14, 2015	March 9, 2015	June 30, 2019	June 30, 2020		0.1	4.3	
2828	Hnn	zBC12	July 23, 2013	September 20, 2013	December 31, 2016	December 31, 2017		0.2	3.3	
2827	Hnn	zEC12	August 28, 2012	September 19, 2012	December 31, 2016	December 31, 2017		0.1	4.3	
2818	Mnn	z114	July 12, 2011	September 9, 2011	June 30, 2014	June 30, 2015		0.2	2.8	
2817	Mnn	z196	July 22, 2010	September 10, 2010	June 30, 2014	June 30, 2015		0.1	3.8	
2098	E10	z10 BC	October 21, 2008	October 28, 2008	June 30, 2012	June 30, 2013	December 31, 2019	0.0	3.7	7.5
2097	Enn	z10 EC	February 26, 2008	February 26, 2008	June 30, 2012	June 30, 2013	December 31, 2019	0.0	4.3	7.5
2096	R07/S07	z9 BC	April 27, 2006	May 26, 2006	June 30, 2010		January 31, 2019	0.1	4.1	8.6
2094	Snn	z9 EC	July 26, 2005	September 16, 2005	June 30, 2010		October 31, 2017	0.1	4.8	7.3
2086	A04	z890	April 7, 2004	May 28, 2004	December 31, 2007		October 31, 2016	0.1	3.6	8.8
2084		z990	May 13, 2003	June 16, 2003	June 30, 2008		December 31, 2014	0.1	5.0	6.5
2064	2nn	z900 G2	April 30, 2002	May 15, 2002	June 30, 2006		December 31, 2014	0.0	4.1	8.5
2066		z800	February 19, 2002	March 29, 2002	December 31, 2005		December 31, 2014	0.1	3.8	9.0
2064	1nn	z900 G1	October 3, 2000	December 18, 2000	June 30, 2006		December 31, 2014	0.2	5.5	8.5
9672	nn7	CMOS G6	May 3, 1999	May 28, 1999	December 31, 2004		December 31, 2010	0.1	5.6	6.0
9672	nn6	CMOS G5	June 23, 1998	September 30, 1998	December 31, 2003		December 31, 2010	0.3	5.3	7.0
9672	Rn5	CMOS G4	June 9, 1997	June 30, 1997	April 30, 2002		December 31, 2007	0.1	4.8	5.7
9672	Rn4	CMOS G3	September 10, 1996	September 30, 1996	July 20, 1999		December 31, 2006	0.1	2.8	7.4
9672	Rn2/3	CMOS G2	June 12, 1995	July 31, 1995	December 29, 1998		December 31, 2004	0.1	3.4	6.0
9672	Rn1	CMOS G1	September 13, 1994	September 13, 1994	December 29, 1998		December 31, 2003	0.0	4.3	5.0
							Average	0.1	4.1	7.3

Legend:

ANN	Announcement of a new product
GA	General availability of a product
HW WDFM	Hardware withdrawal from marketing
LIC WDFM	Licensed Internal Code withdrawal from marketing.
EOS	End of service (unsupported)

ANN to GA - actual or announced number of years between ANN and GA

GA to HW WDFM - actual or announced number of years between GA and HW WDFM

HW WDFM to EOS - actual or announced number of years between HW WDFM and EOS

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

IBM z15 at a glance

System, Processor, Memory
One model: T01; Five features: Max34, Max71, Max108, Max145, Max190
12 core 5.2GHz 14nm PU SCM
1 - 190 PUs configurable as CPs, zIIPs, IFLs, ICFs (up to 215 PUs)
Increased Uniprocessor capacity
Up to 34 sub capacity CPs at capacity settings 4, 5, or 6
Enhanced Out-of-Order and new on-chip HW compression
Enhanced processor/cache design with 2x L3 on-chip and 1.4x L4 (SCM) cache sizes
Up to 40 TB DRAM, protected by Redundant Array of Independent Memory (RAIM)
Virtual Flash Memory (VFM) granularity – 0.5 TB / Feature , up to 12
256 GB HSA
Improved pipeline design and cache management

Announce: Sept. 12, 2019

I/O Subsystem, Parallel Sysplex, STP, Security
PCIe+ Gen3 I/O fanouts with 2 x 16 GBps Buses
6 CSS, 4 Subchannel sets per CSS
0 – 12 PCIe+ I/O Drawers (Gen3) – PDU Models
0 – 11 PCIe+ I/O Drawers (Gen3) – BPA Models
Next generation FICON Express16AS
25 and 10 GbE RoCE Express2
Integrated Coupling Adapter (ICA SR) and Coupling express LR for coupling links
Support for up to 384 coupling CHPIDs per CPC
CFCC Level 24 (HMC 2.15.0)
Crypto Express7S
STP configuration and usability enhancements (GUI)
IBM zHyperLink Express
OSA-Express7S
IBM Secure Service Container

RAS, simplification and others	
L3, L4 Cache Symbol ECC	Enhanced Dynamic Memory Relocation for EDA and CDR
N+2 pumps design for Air Cooled System	Coupling Facility Resiliency enhancements
ASHRAE Class A3 design	Enhanced SE and HMC Hardware (security)
Power – iPDU and BPA	TKE 9.2 LICC and new Smart cards
System Recovery Boost*	Simplified and enhanced functionality for STP configuration
Universal Spare SCMs (CP and SC)	Virtual Flash Memory

PR/SM
Up to 190 CPUs per partition
IBM Dynamic Partition Manager updates
Up to 85 LPARs
16 TB Memory per partition

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 (all features)
- FICON Express16SA
- 25GbE and 10GbE RoCE Express2.1
- IBM zHyperLink Express1.1
- Integrated Coupling Adapter SR1.1

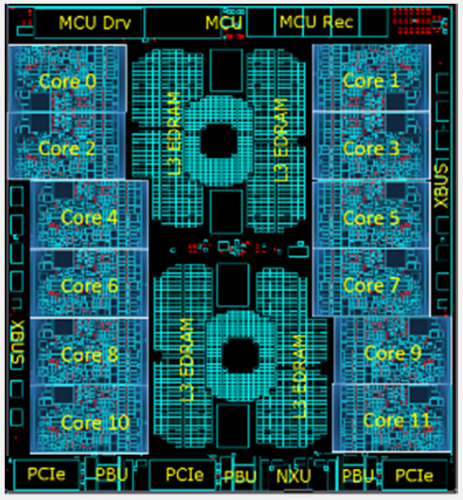

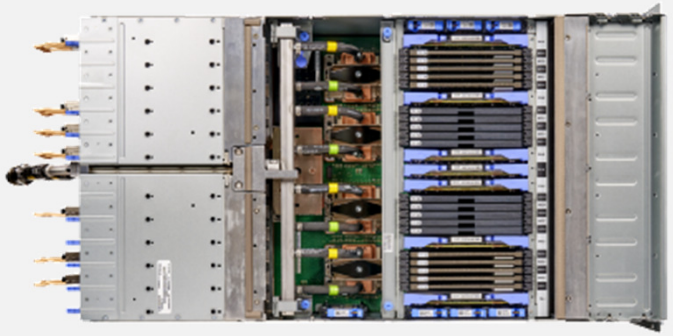

Back Front

IBM Z – Processor Roadmap

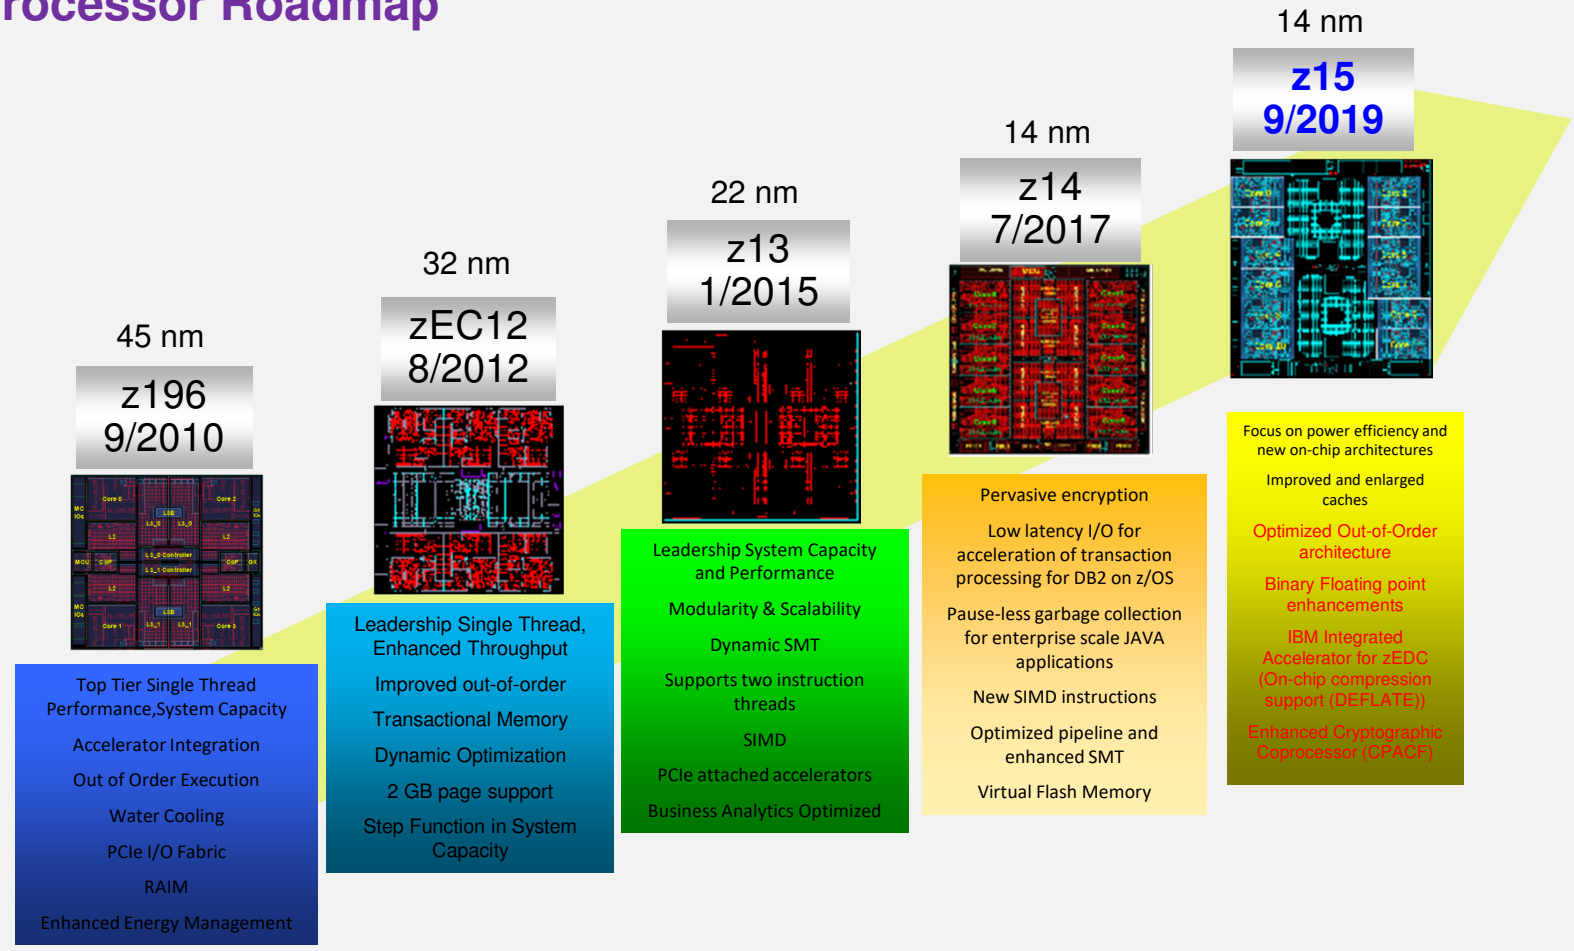

z15 Continues the CMOS Mainframe Heritage

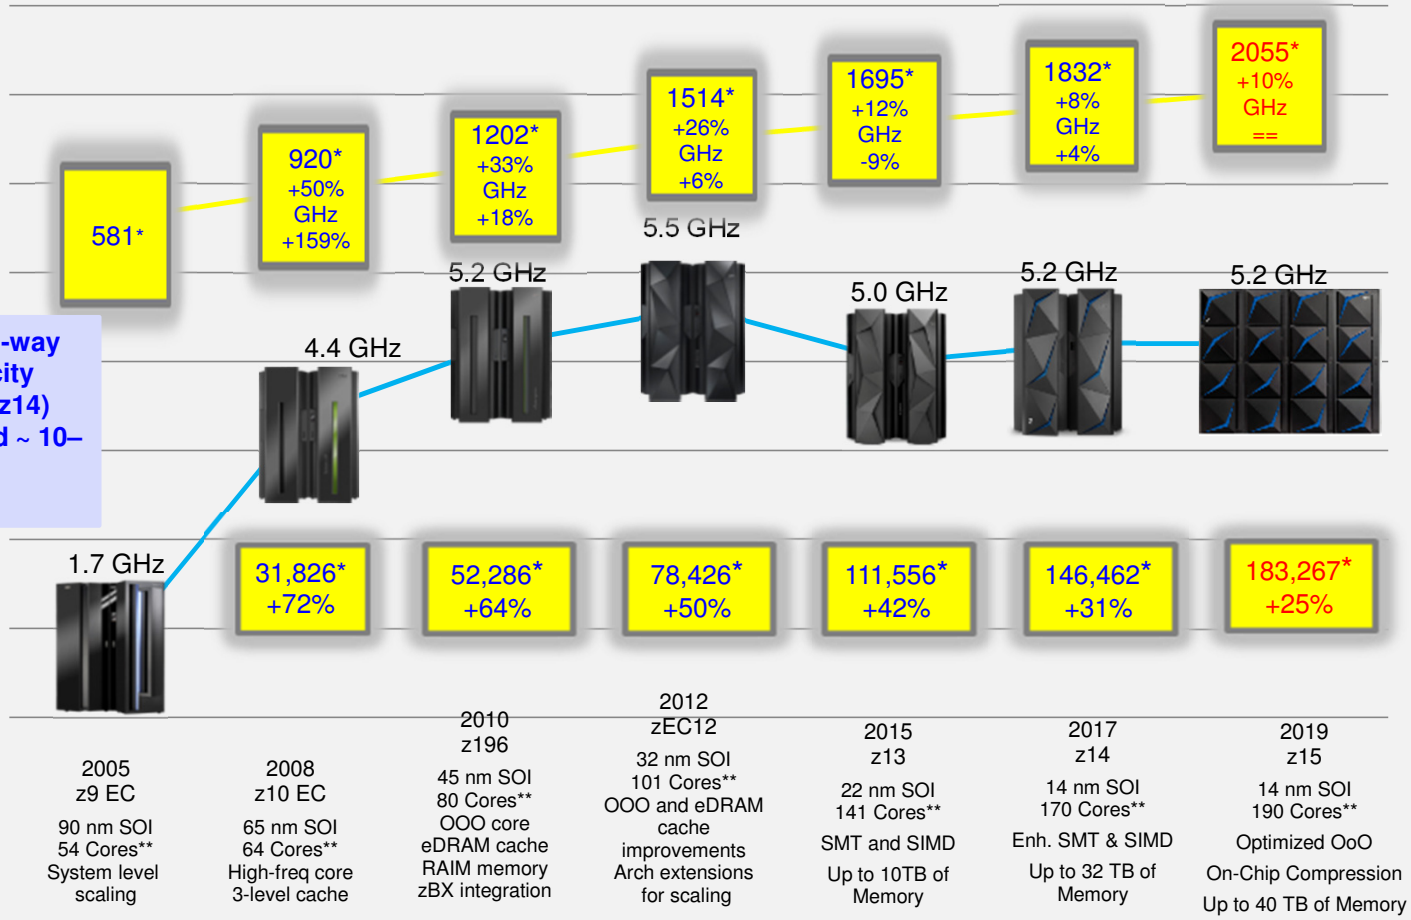

~ 14% for equal z14 n-way
Up to 25% max capacity
190-way vs 170-way (z14)
SMT vs Single Thread ~ 10–40% (average 25%)
- both zIIP & IFL

GHz / PCI*

* Processor Capacity Index (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).

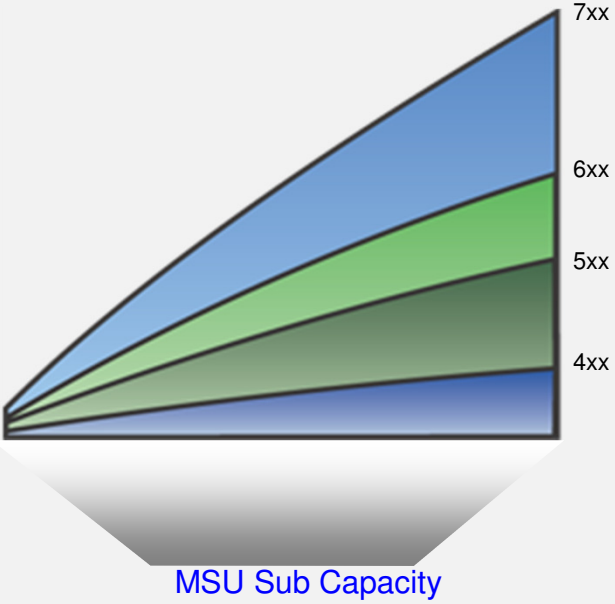

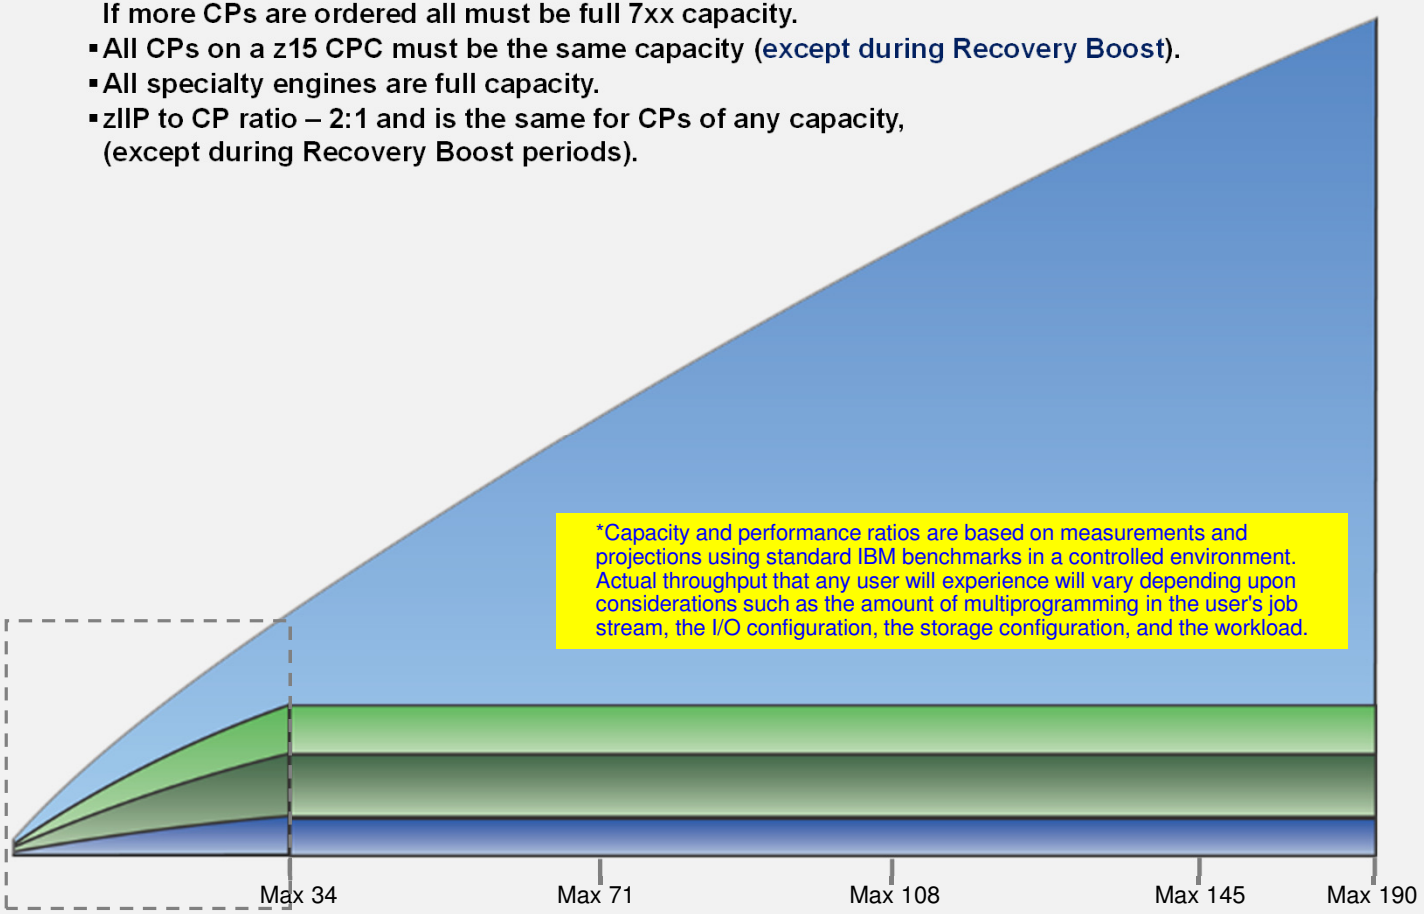

*Capacity and performance ratios are based on measurements and projections using standard IBM benchmarks in a controlled environment. Actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload.

Reminder – 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 Hardware Instrumentation Services (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)
  - 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)

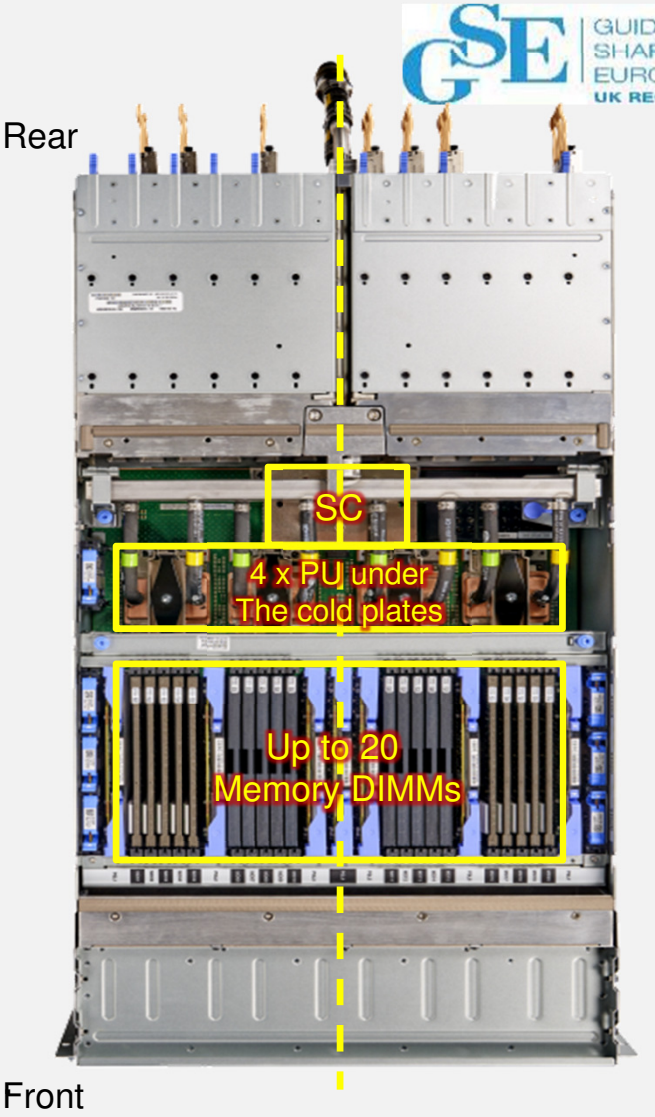

z15 Processor (CPC) Drawer Connectivity

- **12 PCIe fanout slots per z15 CPC drawer**
  - Increase from 10 PCIe fanouts in z14
- **Integrated Coupling Adapte (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

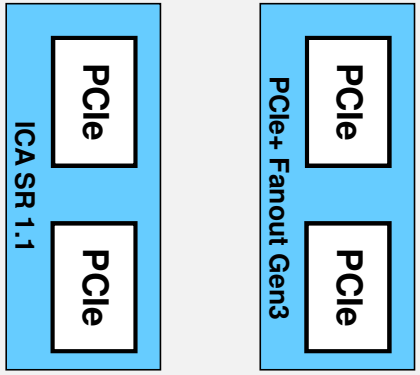

ICA SR1.1 Coupling Links

To PCIe+ I/O Drawers

* 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.

No InfiniBand fanouts

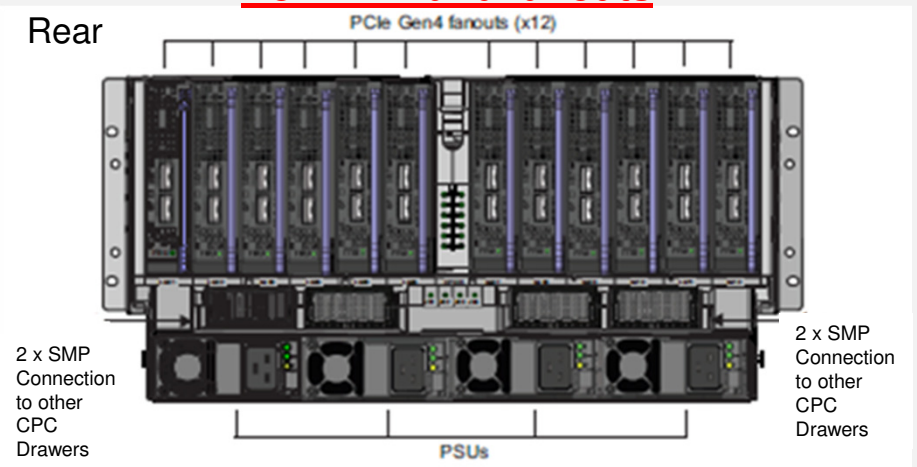

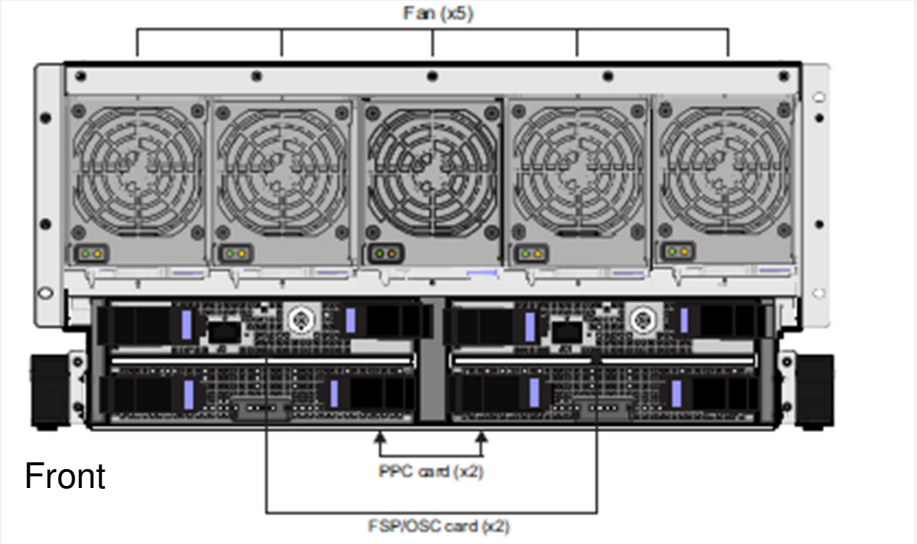

12-Core Processor Chip Detail

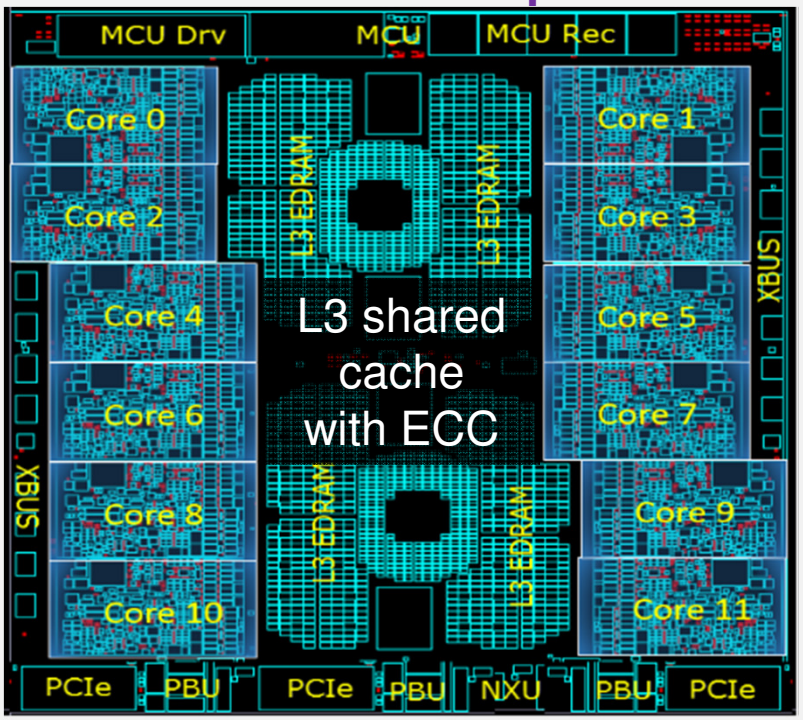

- **14nm SOI Technology**
  - 12 Cores
  - 17 layers of metal
  - 696 mm2 chip area
  - 9.2B transistors versus 6.2B on z14
- **20% reduction area**
- **20% reduction in power**

- **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
- **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 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

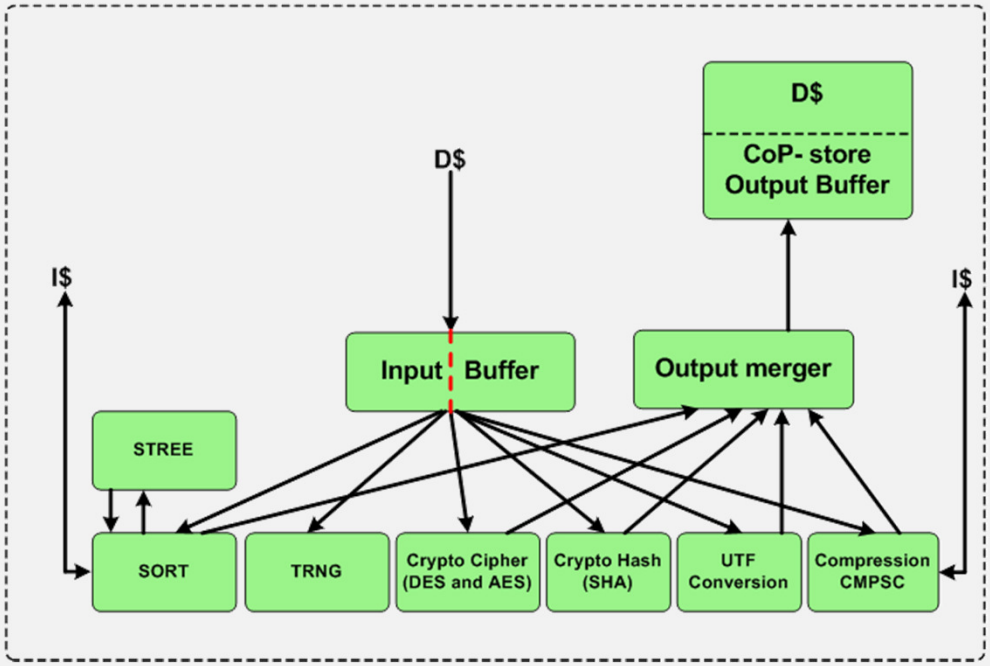

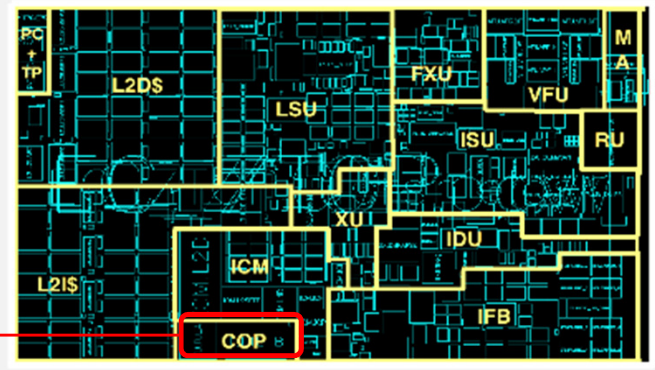

z15 SC Chip

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

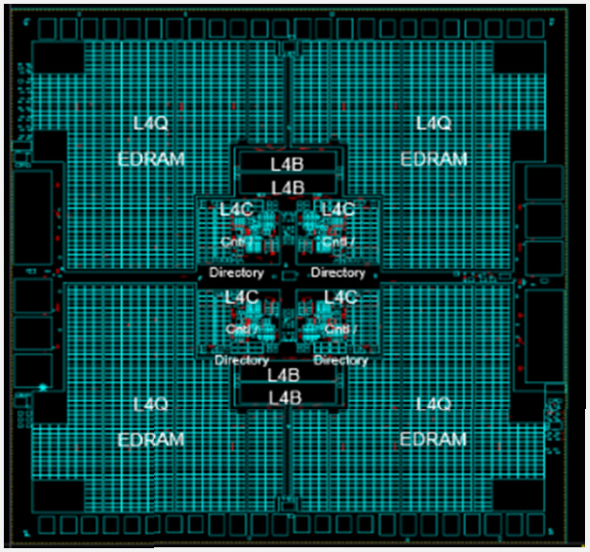

Fully Populated Drawer

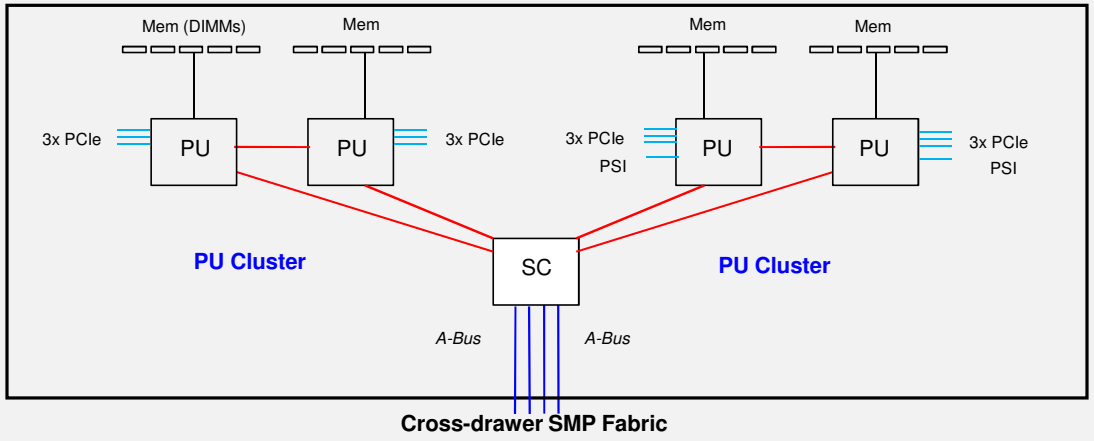

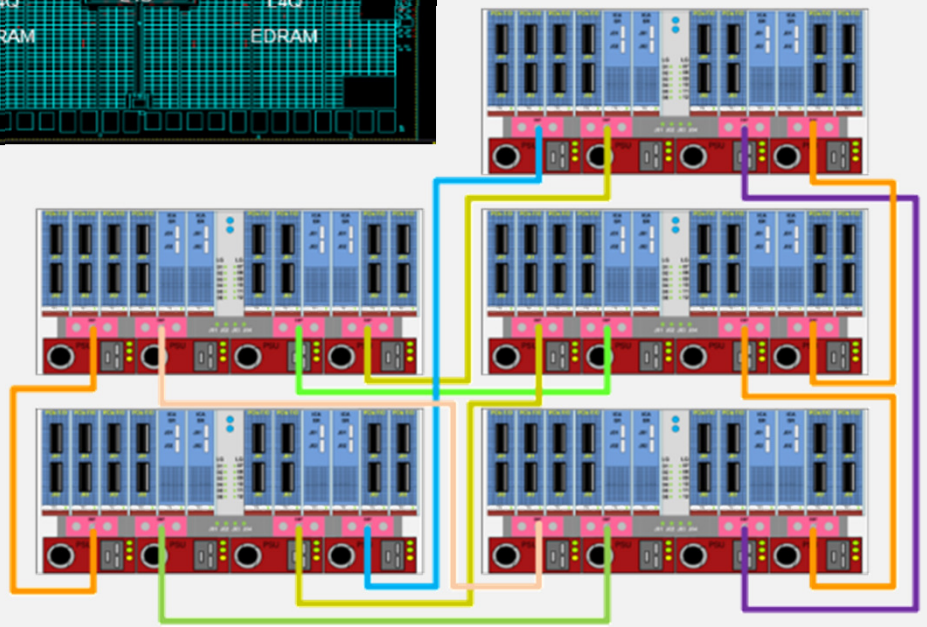

Cache topology comparison: z15 vs. z14

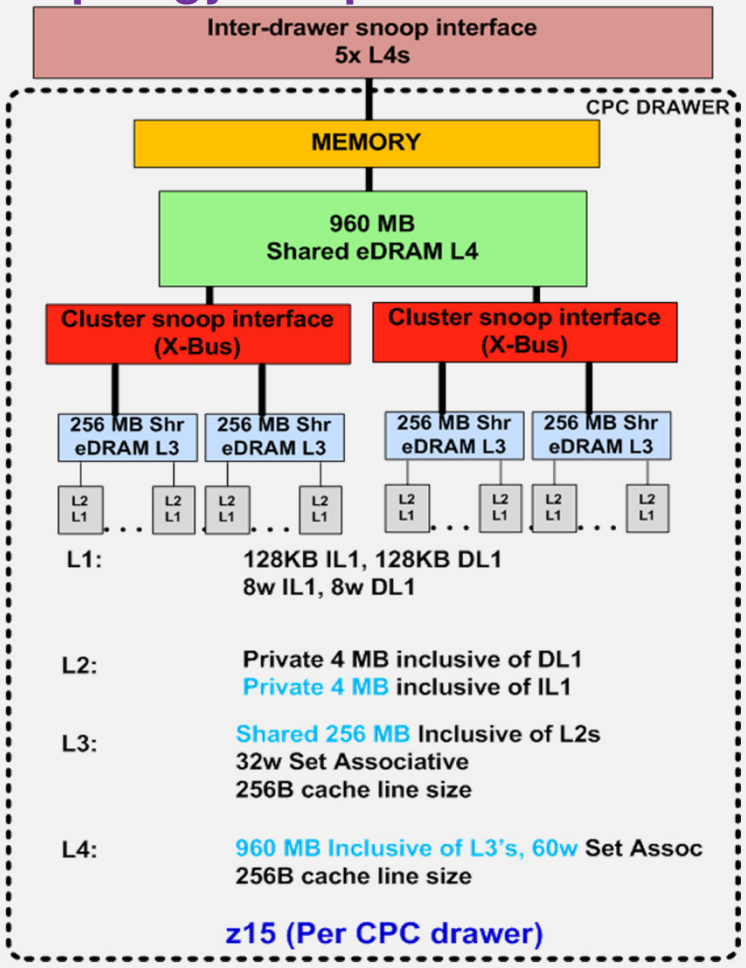

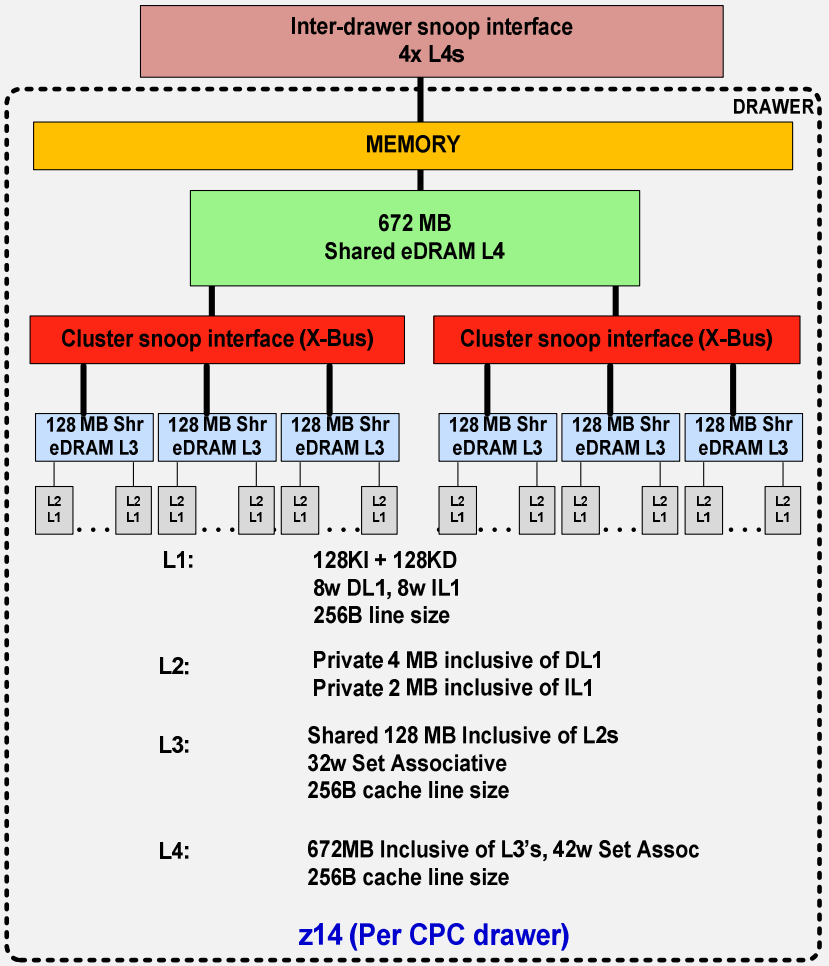

z15 Capacity Considerations

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

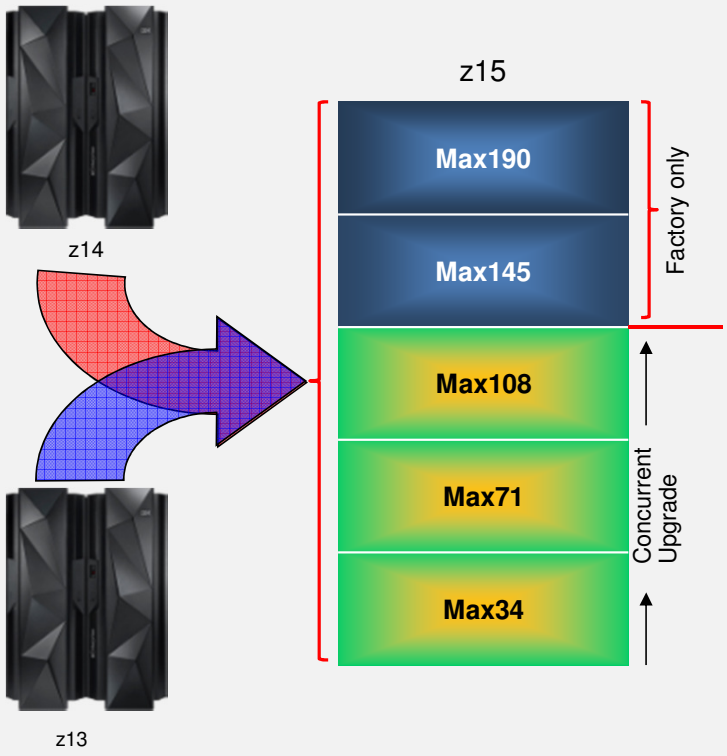

- **z15 to z15 upgrades**
  - z15 Concurrent upgrade from Max 34 to Max71 to Max108
    - Each max level adds a CPC drawer
    - No MES upgrade to Max145 or Max190
  - Additional I/O Drawers
    - Based on available space in current frames and/or I/O expansion frames
  - *No conversion available between 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 - PDU-based configuration (max. configuration example)

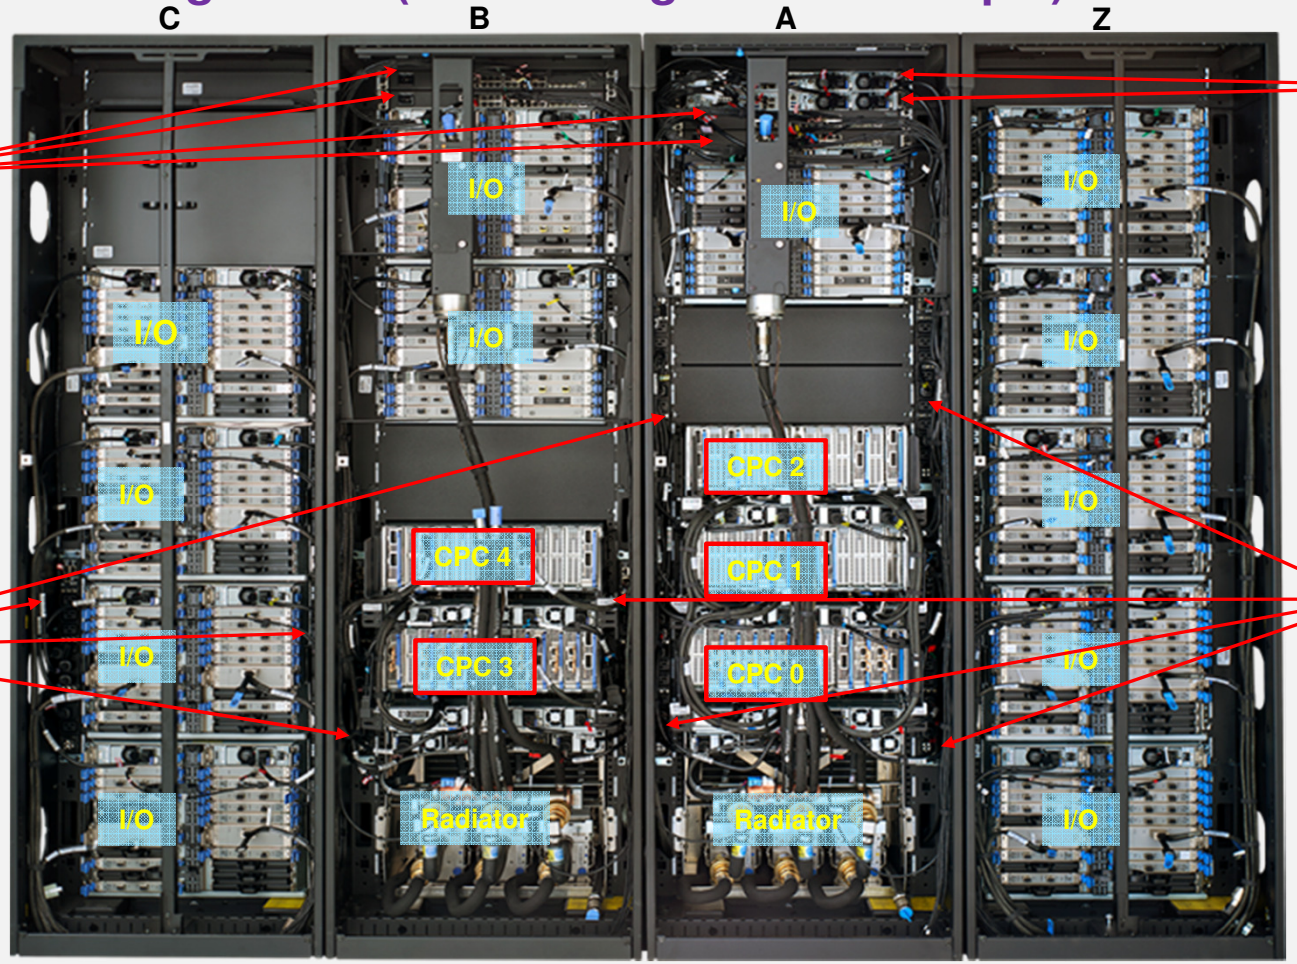

GbE Switches (4x or 2x)**

Support Elements (2x)

iPDUs (3 Phase)

iPDUs (3 Phase)

- 4 or 5 CPC drawer configurations are factory Only
- Concurrent add 2-3 CPC w/CPC Plan Ahead (FC 2271 and 2272)
- **MES Add of 4th or 5th CPC drawer not supported**

z15 BPA-based water cooled (max. configuration example)

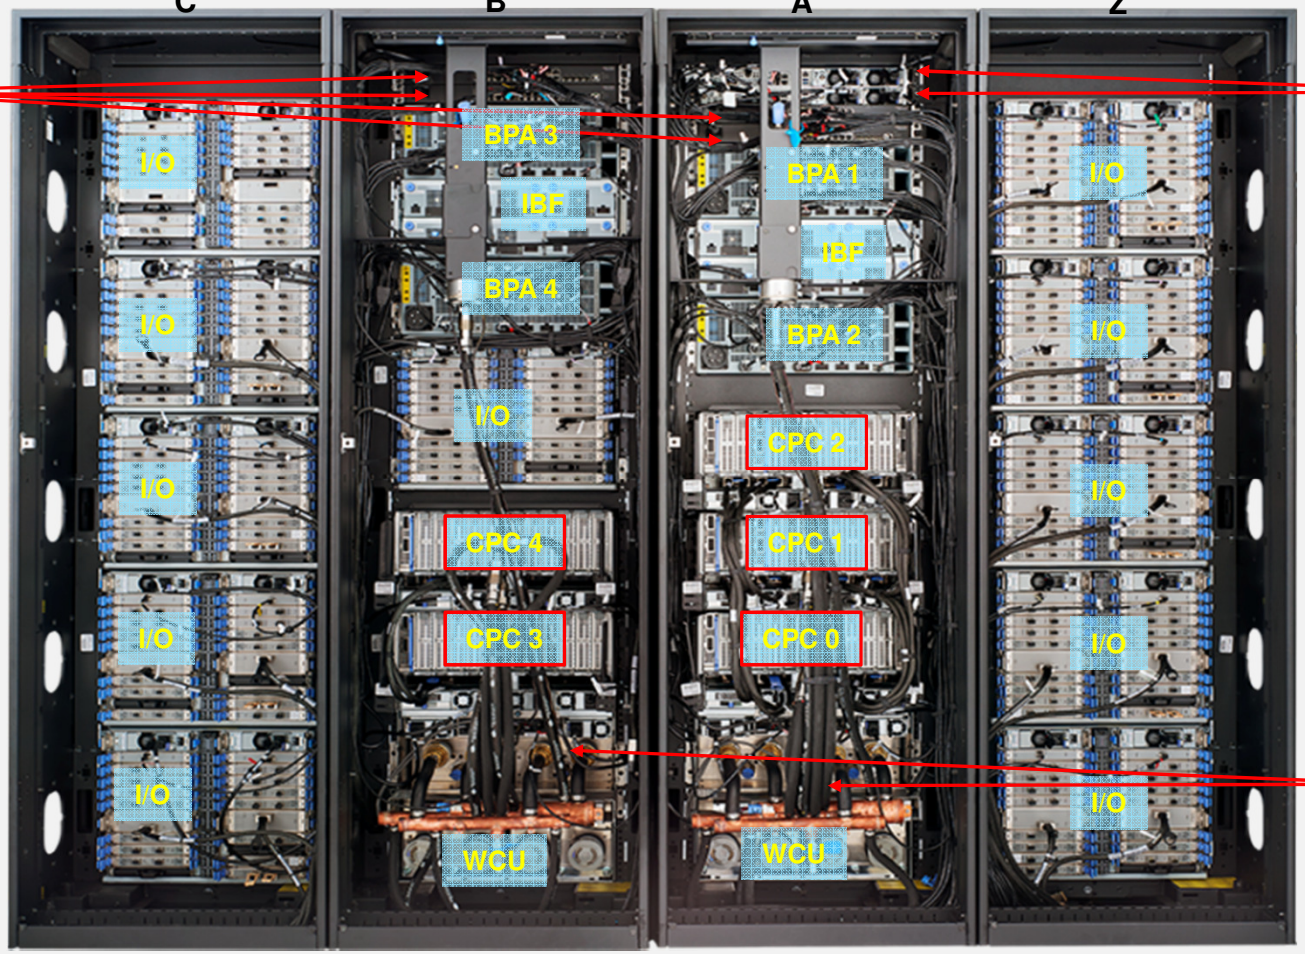

- 4 or 5 CPC drawer configurations are factory Only
- Concurrent add 2-3 CPC w/CPC Plan Ahead (FC 2271 and 2272)
- **MES Add of 4th or 5th CPC drawer not supported**

Integrated system design for z15

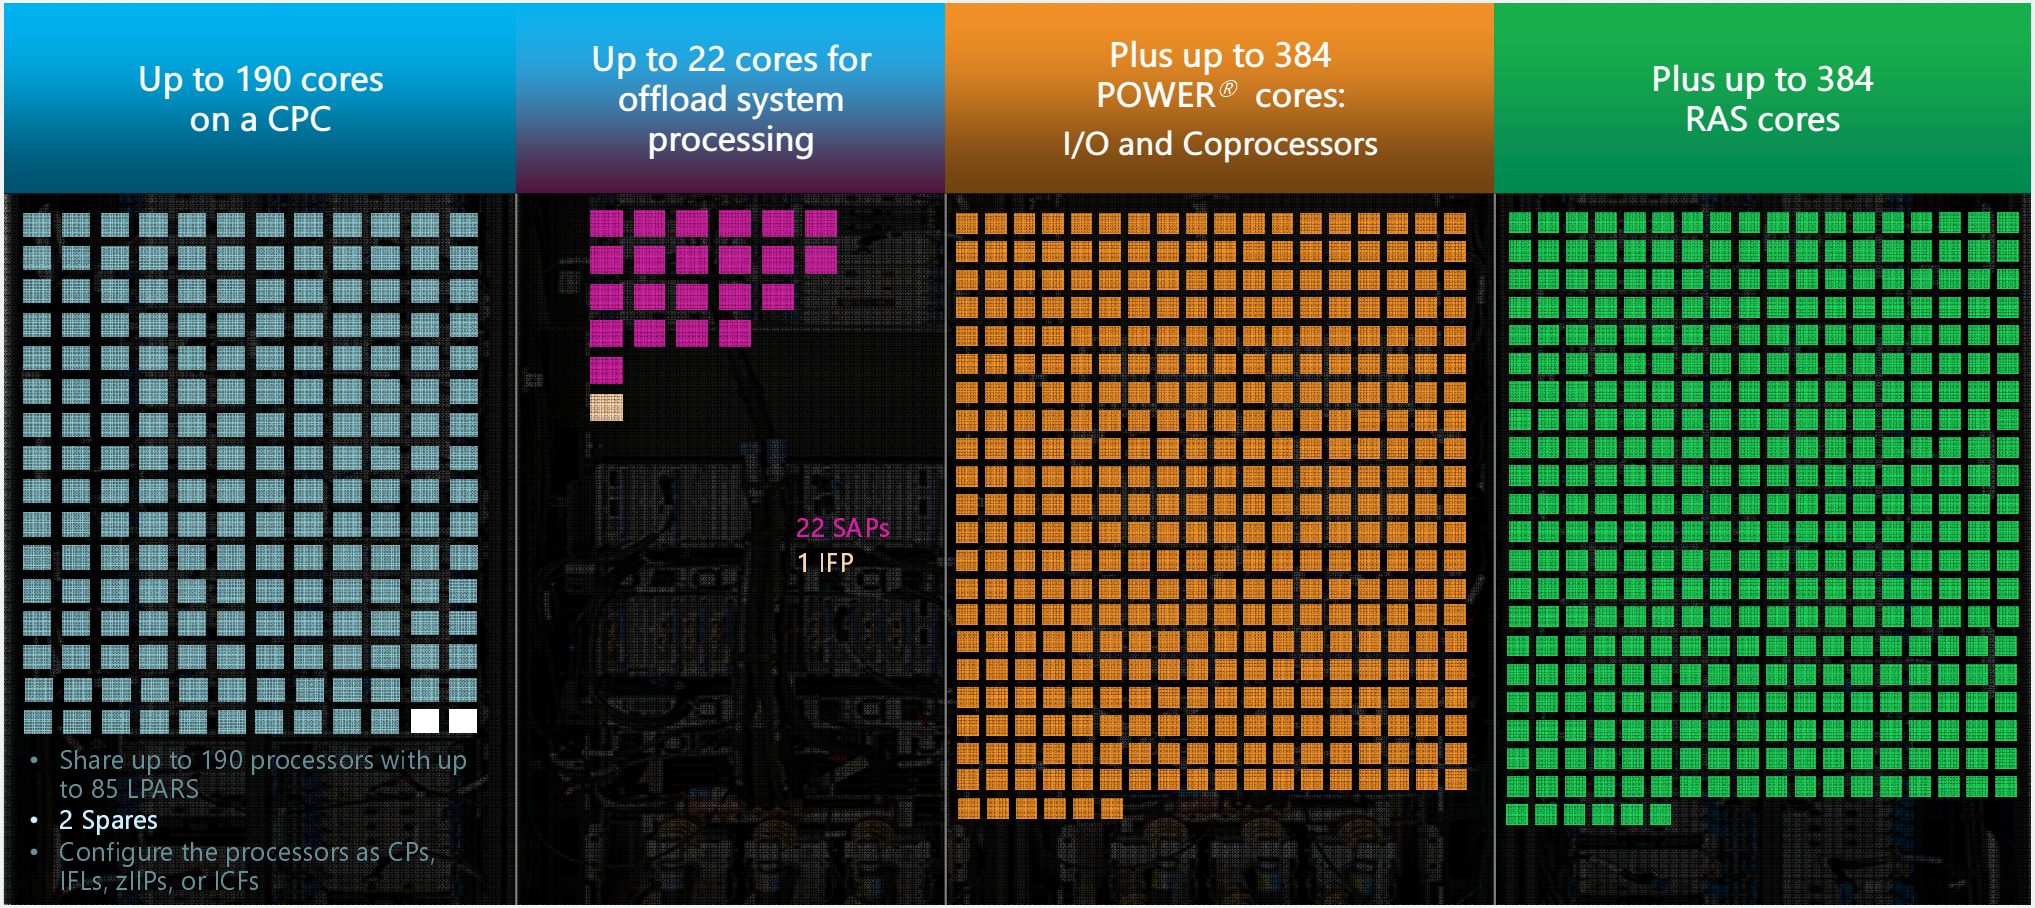

z15 Memory

z15 Purchased Memory Offering Ranges

Feature	Standard Memory GB (Min* – Max)	Flexible Memory GB
Max34	512 - 7936	NA
Max71	512 - 16128	512 - 7936
Max108	512 - 24320	512 - 16128
Max145	512 - 32512	512 - 24320
Max190	512 - 40704	512 - 32512

- **Purchased Memory – Memory available for assignment to LPARs**
- **Hardware System Area – Standard 256 GB of addressable memory for system use outside customer memory**
- **Standard Memory – Provides minimum physical memory required to hold customer purchase memory plus 256 GB HSA**
- **Flexible Memory – Provides additional physical memory needed to support activation base customer memory and HSA on a multiple CPC drawer z15 with one drawer out of service.**
- **No Plan Ahead Memory for new orders.**
  - Existing Plan Ahead memory can be carried forward during an upgrade.

z15 I/O Infrastructure

IBM Z I/O Subsystem *Internal* Bus Interconnect Speeds

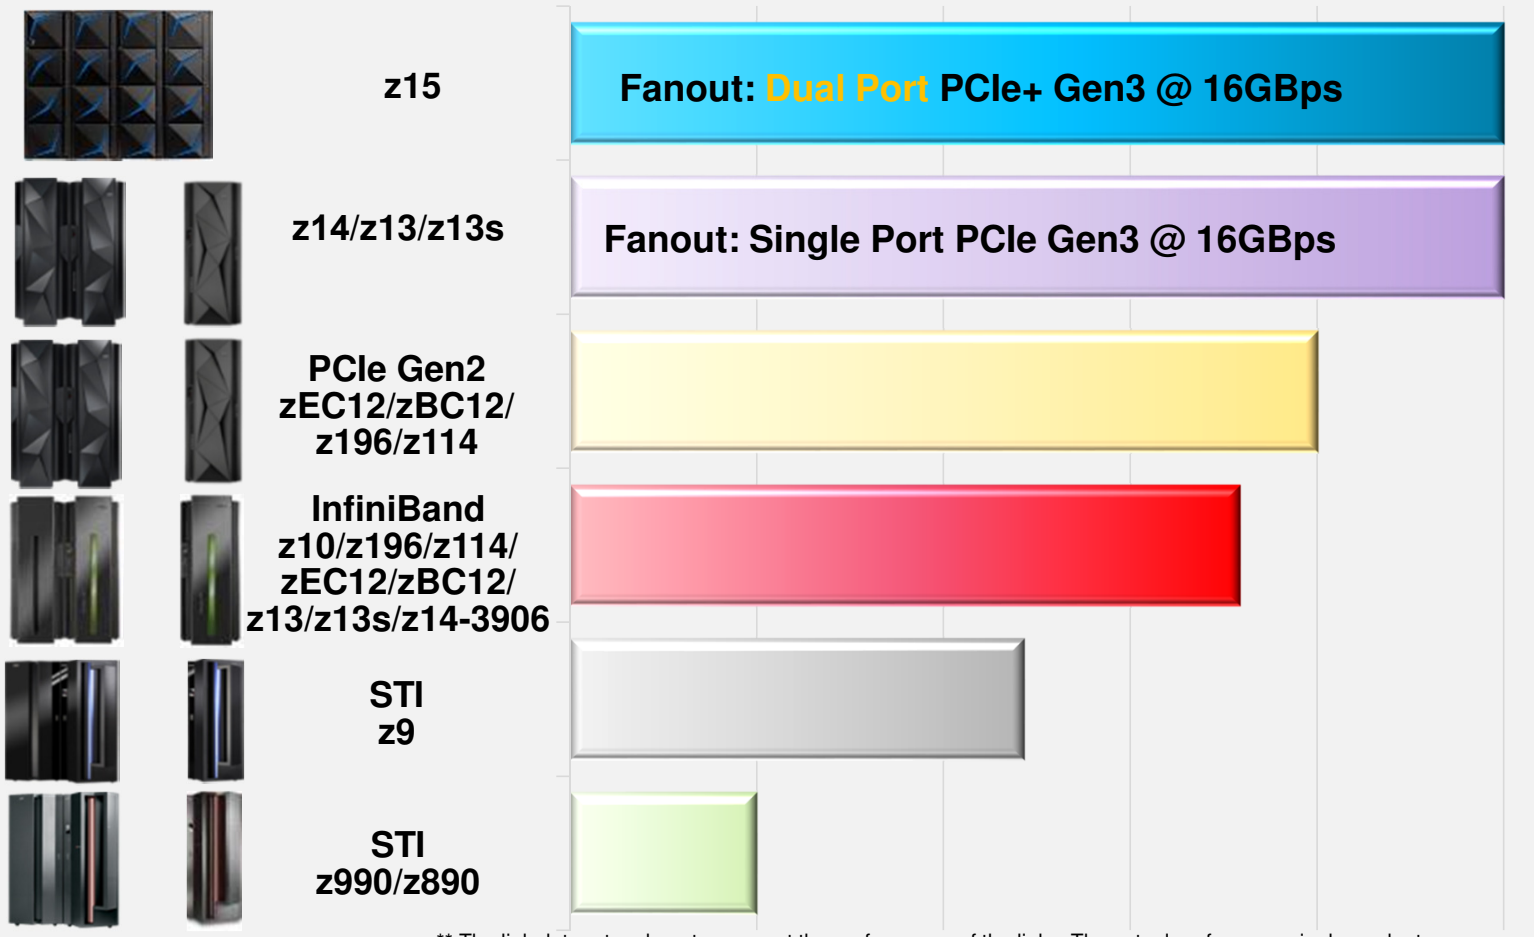

** The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

z15 I/O Infrastructure

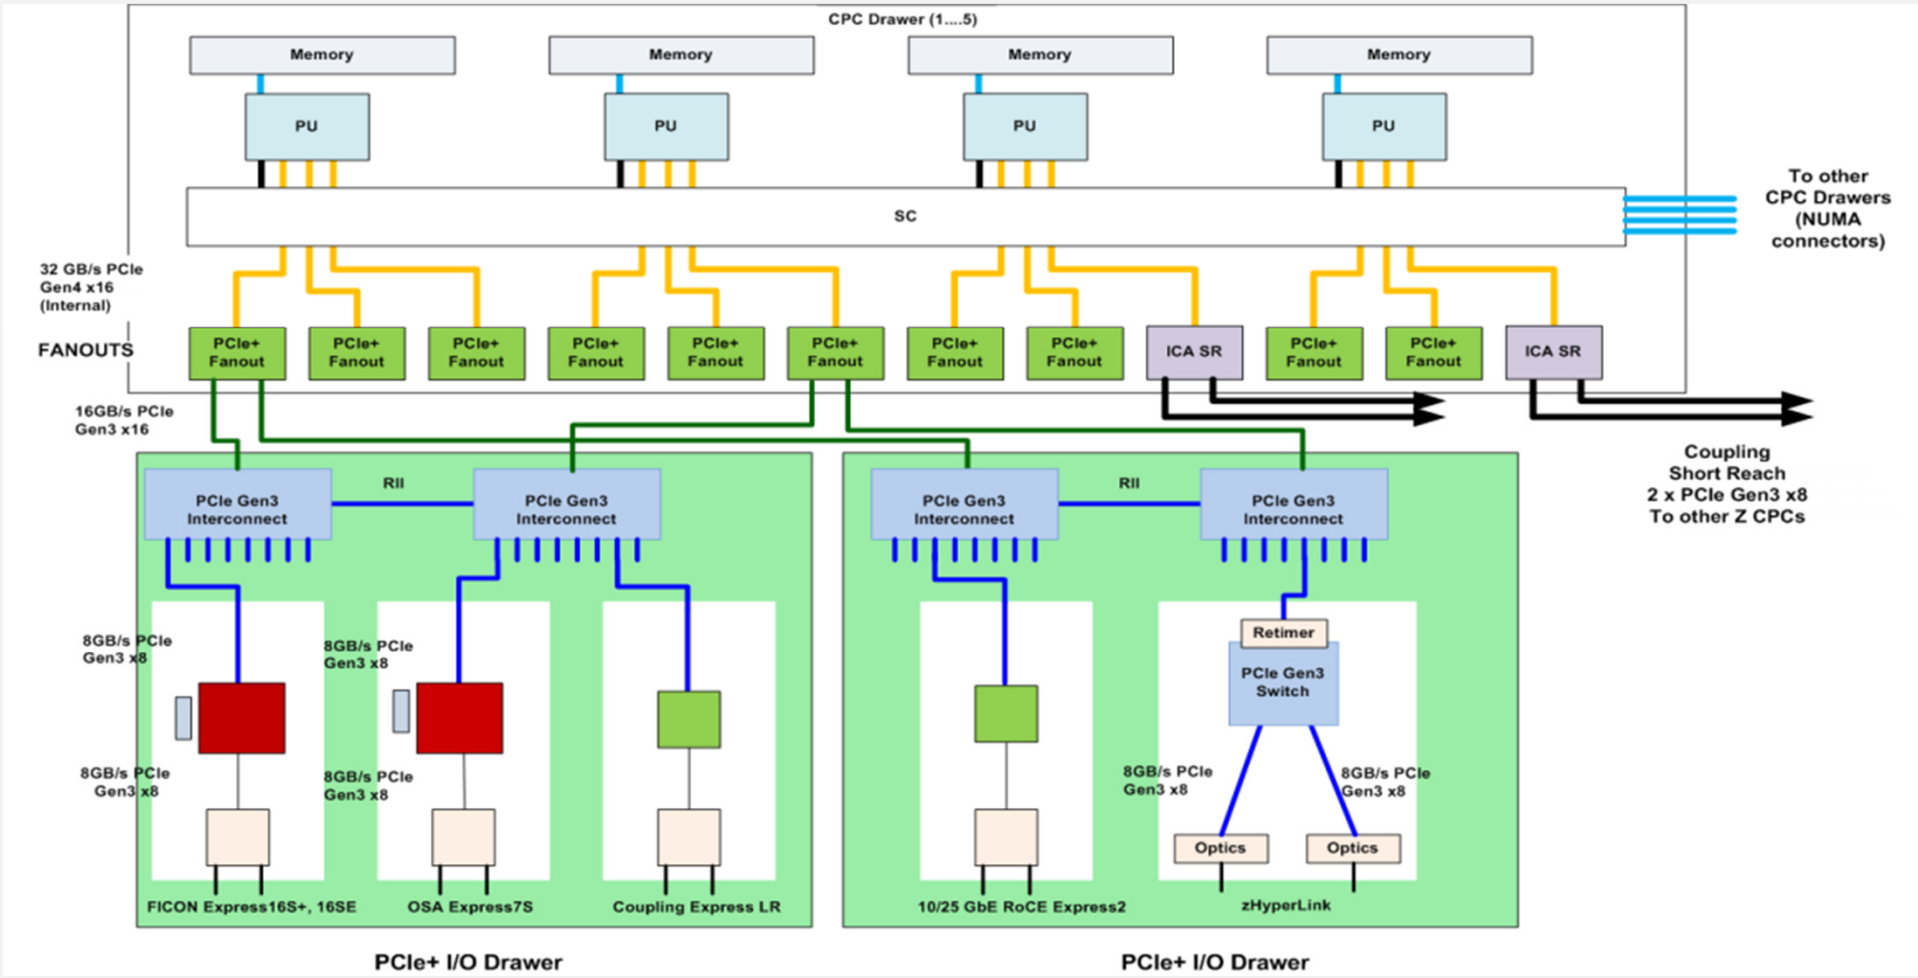

CPC & PCIe+ I/O Drawer Structure and Interconnect

- **Max drawer count power dependent**

Power Type	Max PCIe+ I/O Drawers	Max PCIe I/O Cards
iPDU (Power Distribution Unit)	12	192
BPA (Bulk Power Assembly)	11	176

- **I/O drawer location configuration dependent**
  - Placement starts above the CPC Drawers in the A Frame and expands to the Z frame then the C Frame
  - B Frame is ”Factory build only”, included if Max145 or Max190 are ordered

- **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

I/O Max - PDU

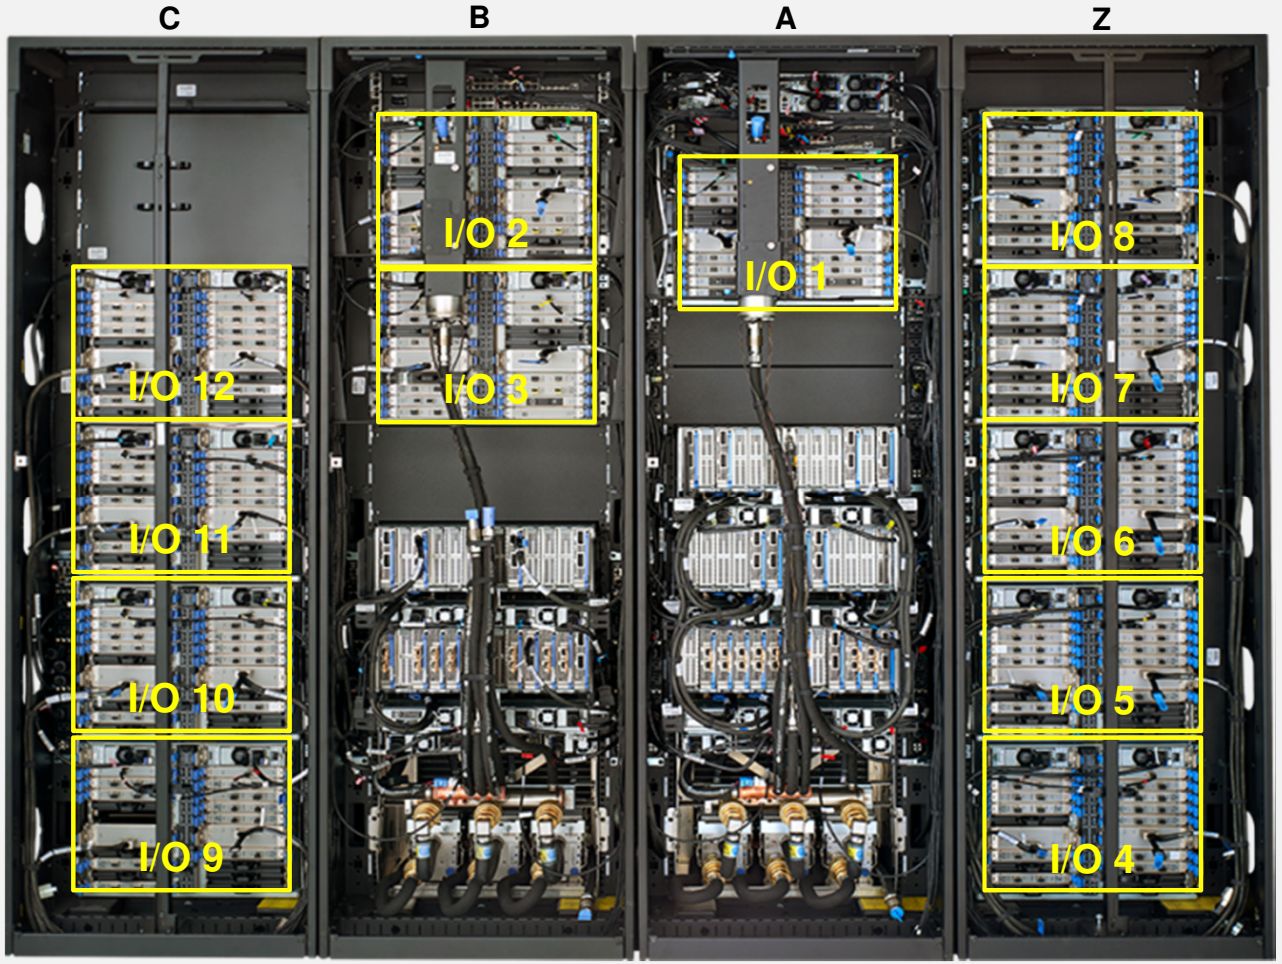

I/O Max - BPA

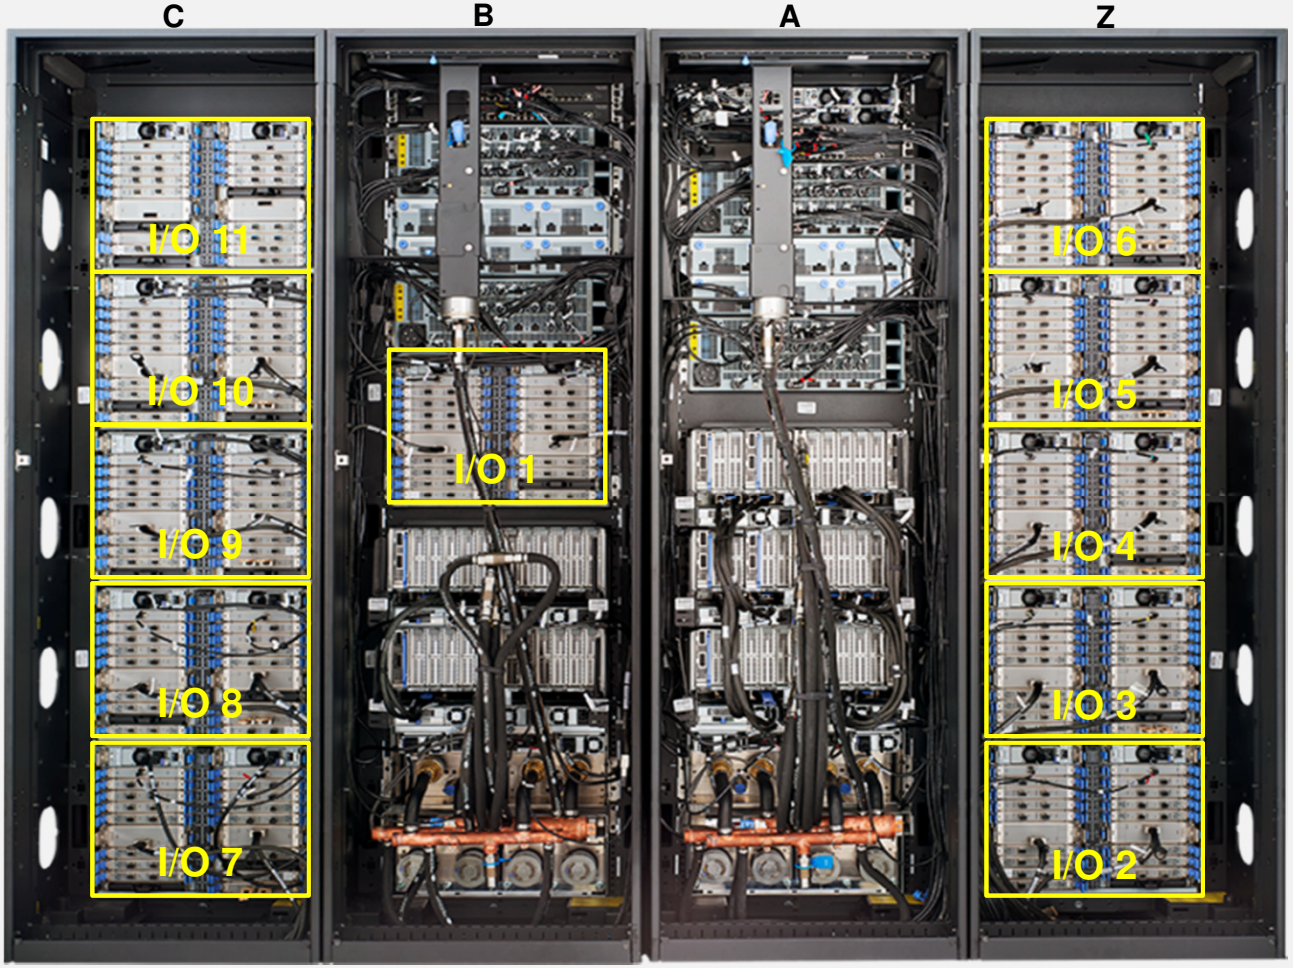

Technology Refresh for I/O Features

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

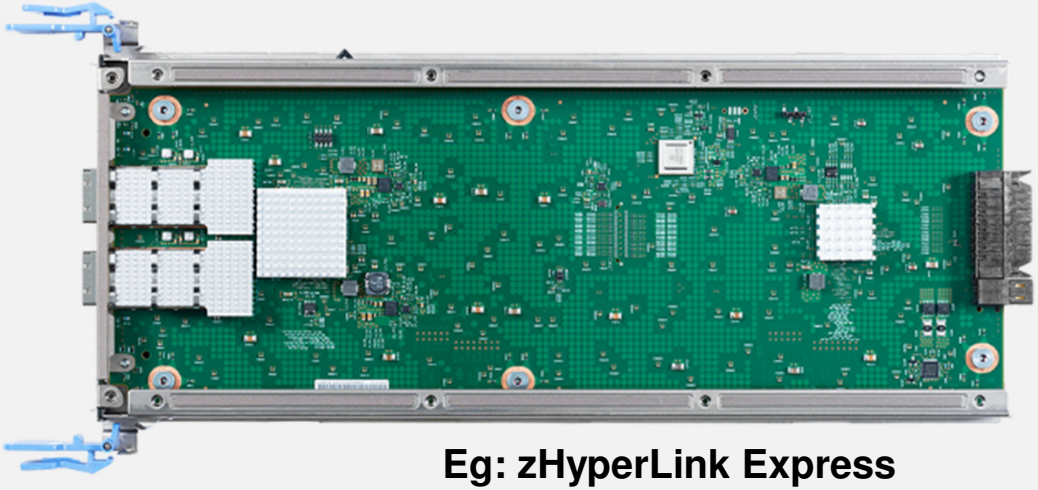

Eg: zHyperLink Express

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
    - 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

PCIe+ I/O Drawer – 16 slots

- **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).**

“Native” PCIe feature Plugging and Resource Groups (RGs)

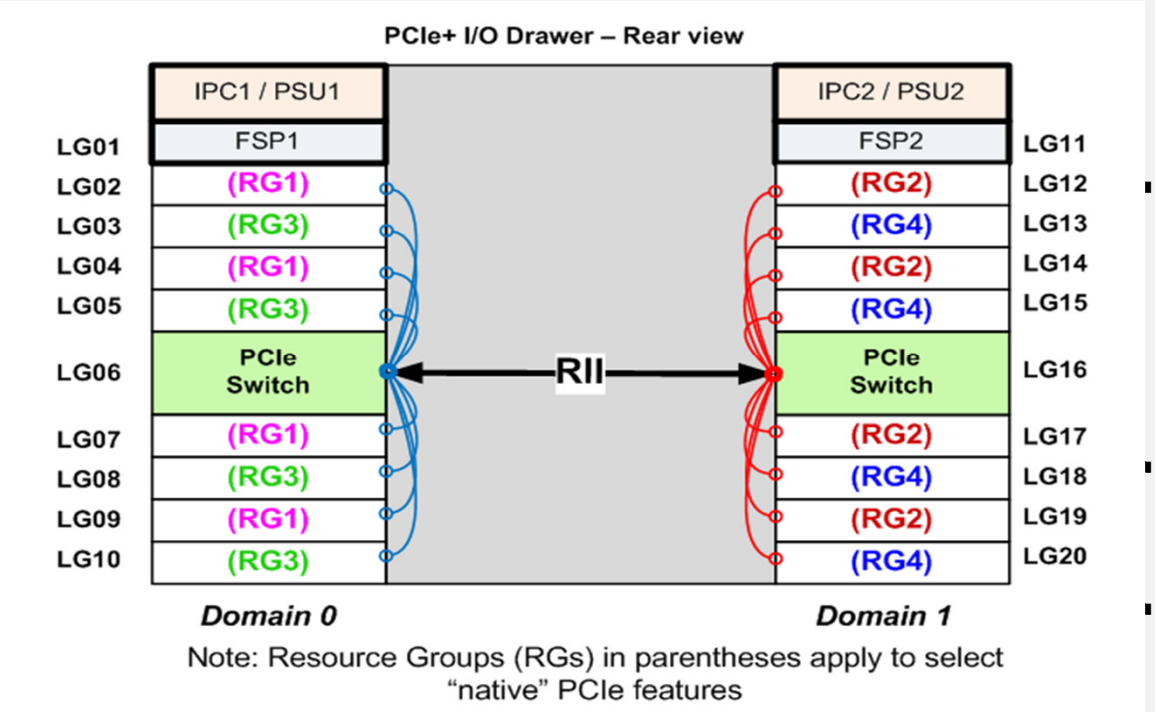

Note: Resource Groups (RGs) in parentheses apply to select “native” PCIe features

- **Configurator (eConfig) card placement:**
  - Places features of the same type in I/O slots to balance the number in I/O Domains and Resource Groups
  - Reports RG assigned in “AO Data”
- **Each slot in a drawer is assigned to an RG***
  - Note: These rules are NOT the same as z14 M0x.
  - *Domain Pair 0 and 1*
    RG1: slots 2, 4, 7 and 9
    RG2: slots 12, 14, 17 and 19
    RG3: slots 3, 5, 8 and 10
    RG4: slots 13, 15, 18 and 20
- **Each feature’s physical channel ID (PCHID) is assigned based on its slot and the PCIe I/O drawer location.**
- **For high availability: Note the RG assignment and PCHID of each native PCIe feature*. Balance assignment of each PCIe feature type to every LPAR between the four RGs.**

*** See speaker notes**

Review of the Integrated Firmware Processor (IFP)

- **Integrated firmware processor (IFP)**
  - **The IFP is allocated from the pool of non-client cores available for the whole system**
    - Unlike other characterized cores, the customer doesn't pay for the IFP
    - It's a single core dedicated solely for the purpose of supporting the native PCIe features and is initialized at POR if these features are present.
    - The z15 has four Resource Groups (RGs) which have firmware for:
      - **10GbE and 25GbE RoCE Express2.1**
      - **10GbE and 25GbE RoCE Express2**
      - **10GbE RoCE Express**
      - **Coupling Express LR**

IFP and Resource Groups – Basic Configuration

- **Resource Groups (RG)**
  - **Each Resource Group will handle 25% of the native PCIe features based on the plugging rules and purchases made in pairs of features**
    - During firmware updates, error conditions, etc. that affects one RG, ALL the features attached to that RG will be unavailable across all PCIe+ I/O Drawers
    - MCL update to Resource Group requires a RG outage of a few minutes

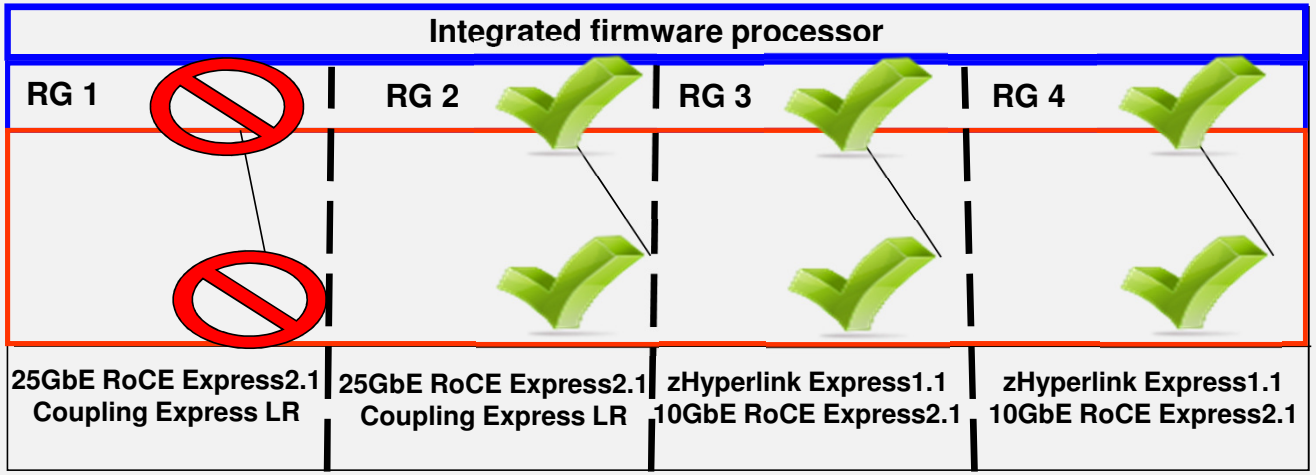

What is IBM zHyperLink™?

- **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 which was introduced with z14.**
- **z15 Introduces an updated feature, the IBM zHyperLink Express1.1**
- **IBM zHyperLink™ reduces latency by interconnecting the Z CPC directly to the I/O Bay of the DS8880**
- **zHyperLink improves application response time, cutting I/O sensitive workload response time in half without significant application changes.**

zHyperLink Express® at a Glance

- **Feature Code 0451 (FC 0431 carry forward)**
  - Two ports per feature
  - Maximum of 16 features (32 ports)
  - Function ID Type = HYL
  - Up to 127 Virtual Functions (VFs) per port (254 per feature)
  - Point to point connection using PCIe Gen3
  - Maximum distance: 150 meters
- **DS8880 models 984, 985, 986 and 988*, DS8900**

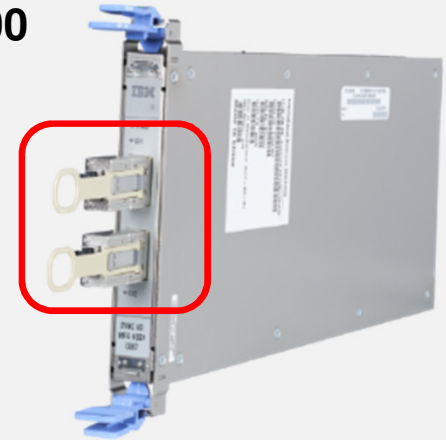

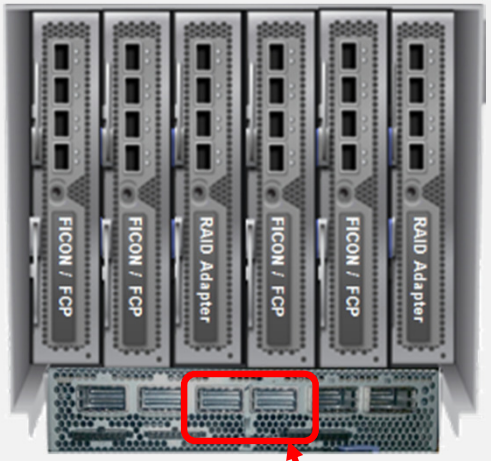

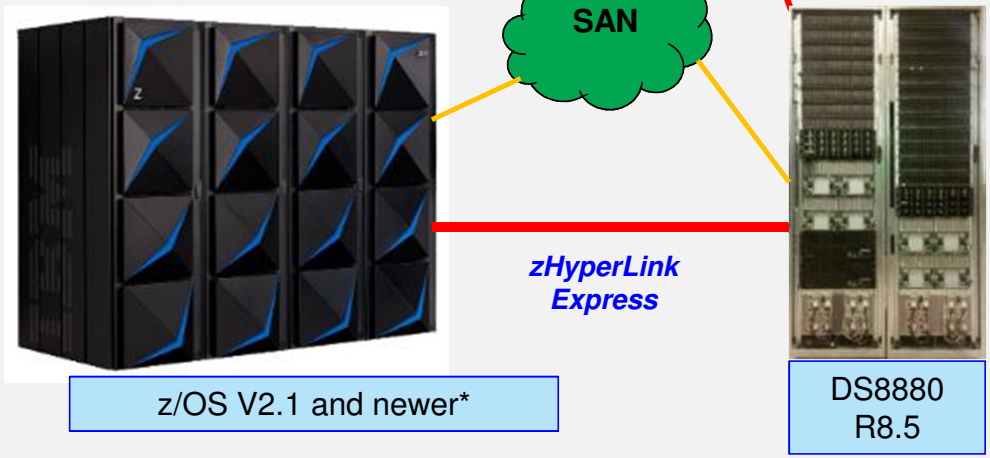

FICON Express16SA

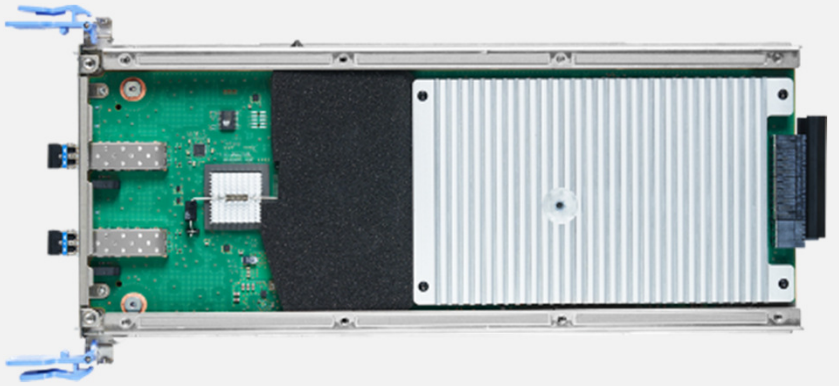

- **For FICON, zHPF, and FCP**
  - FC 0436 (LX) & 0437 (SX)
  - CHPID types: FC and FCP
    - Two PCHIDs/CHPIDs
    - NO mixed CHPIDs for same card – only FC or FCP
- **Same performance as FICON Express16S+**
- **Auto-negotiates to 8 or 16 Gbps**
  - Negotiation to 4 Gbps NOT supported
  - 2 and 4 Gbps 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

FICON Considerations

- **FICON Express16SA auto-negotiates to 8 or 16 Gbps**
  - 2 and 4 Gbps connectivity not supported for point to point connections
  - 2 and 4 Gbps supported through a switch with 8 or 16 Gb optics
- **FICON Express16SA cannot mix the port types**

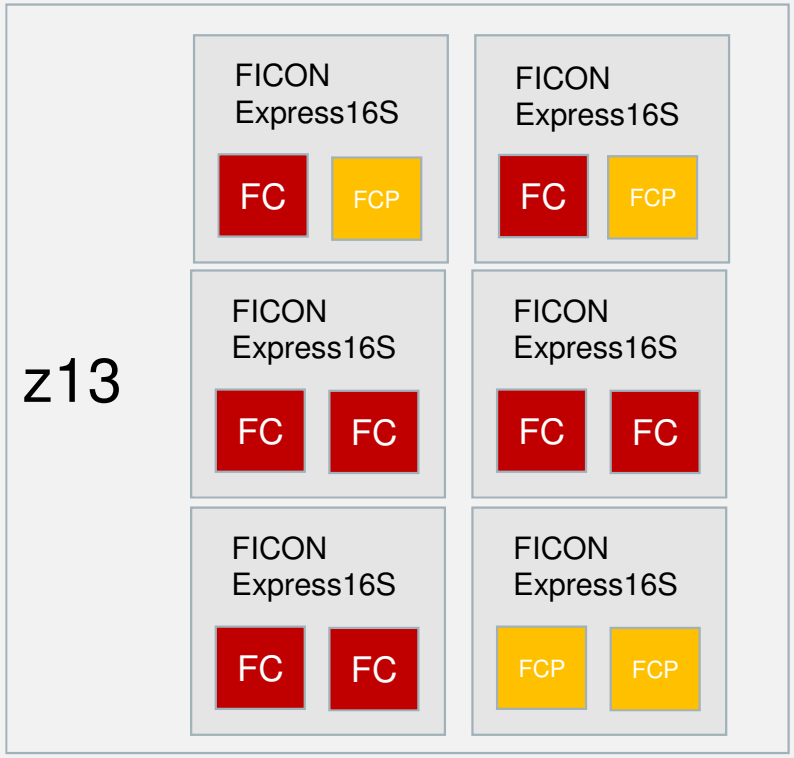

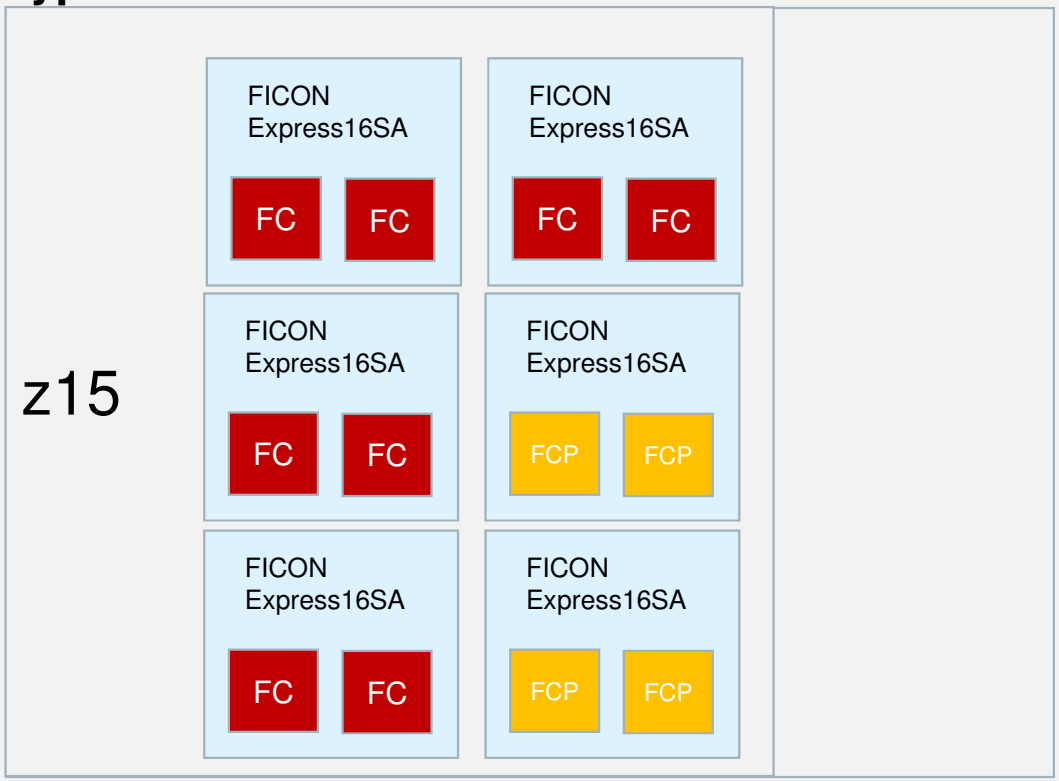

OSA Express7S 25 GbE SR1.1 – FC0449

- **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)

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

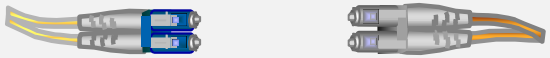

FC 0442 – GbE LX | FC 0443 – GbE SX

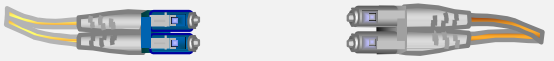

OSA-Express7S 1000BASE-T

- **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)

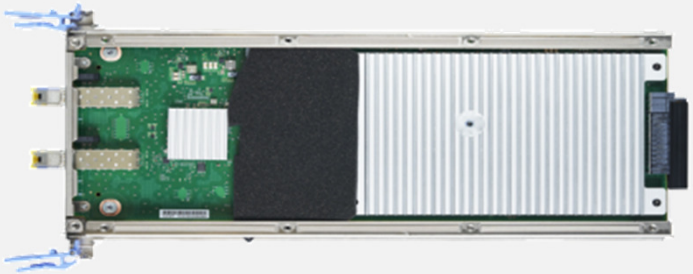

Connector = RJ-45

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)**

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)

No application changes required.

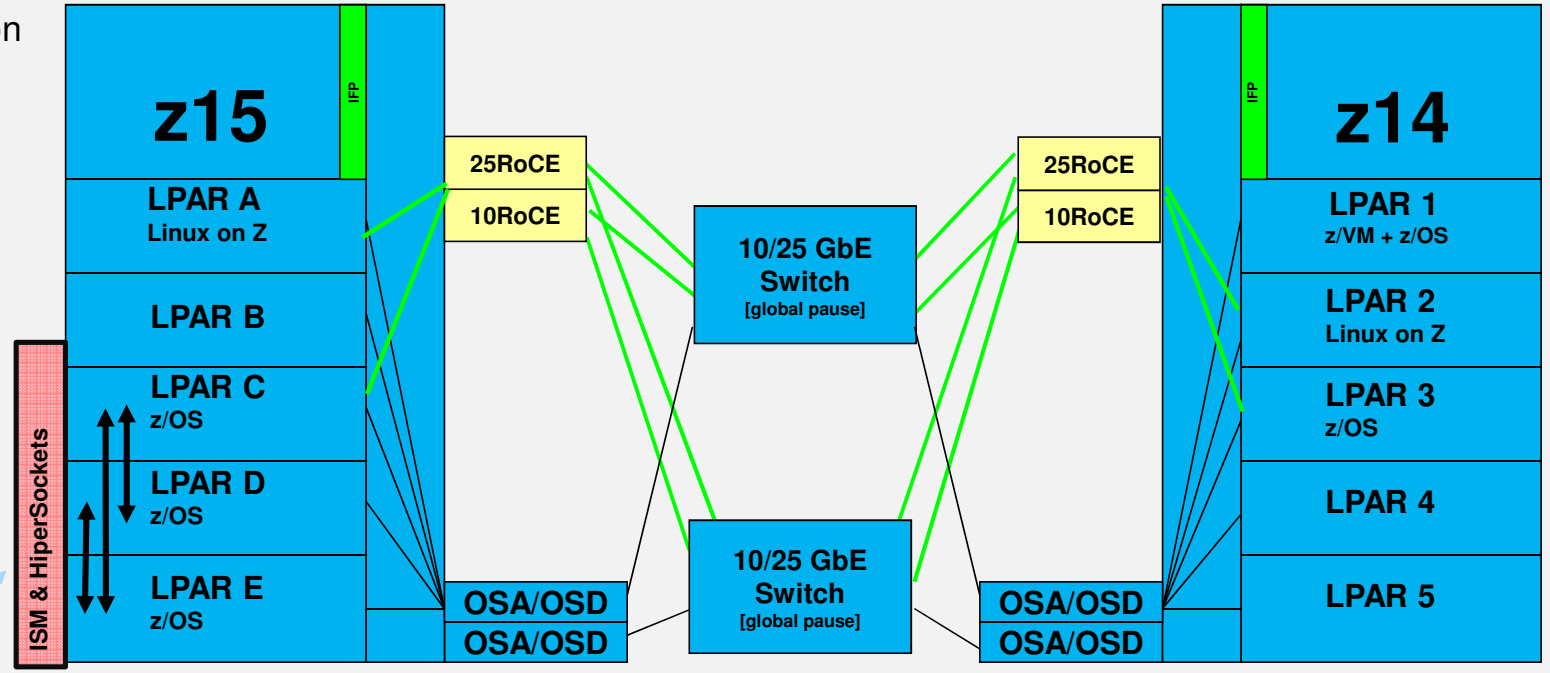

z13 supports 10GbE RoCE Express only.

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 required. SMC-D bypasses the IP stack , sparing CPU cycles.

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*

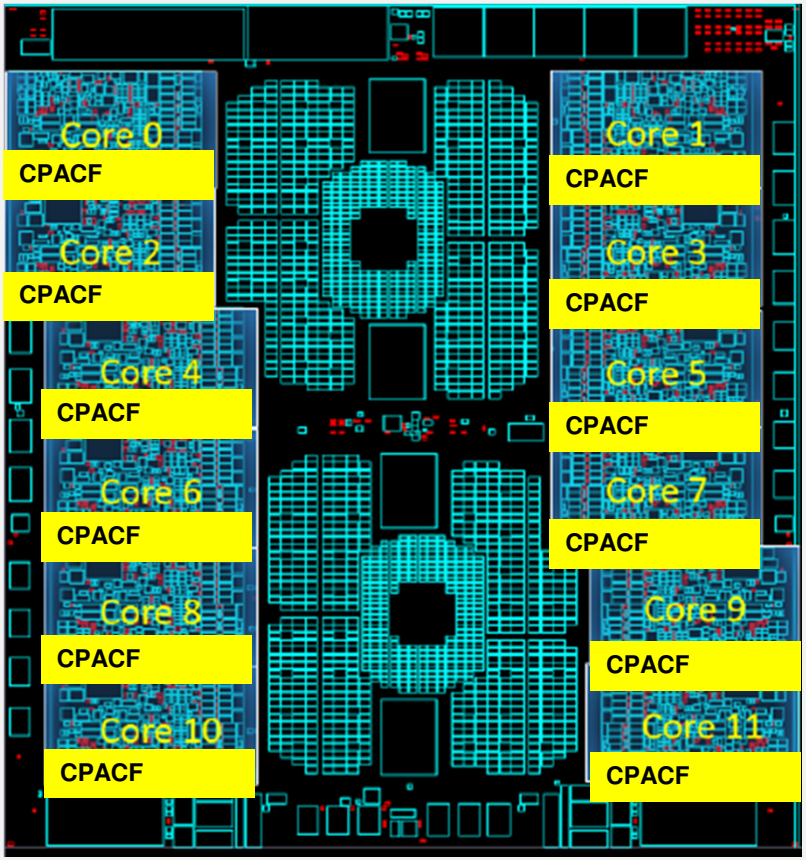

Crypto Express7S

Two new cards designed for z15

- FC 0899 One co-processor
  - Max 16 per server
- FC 0898 Two co-processors
  - Max 8 per server

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

- Max combined total: 16 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)**

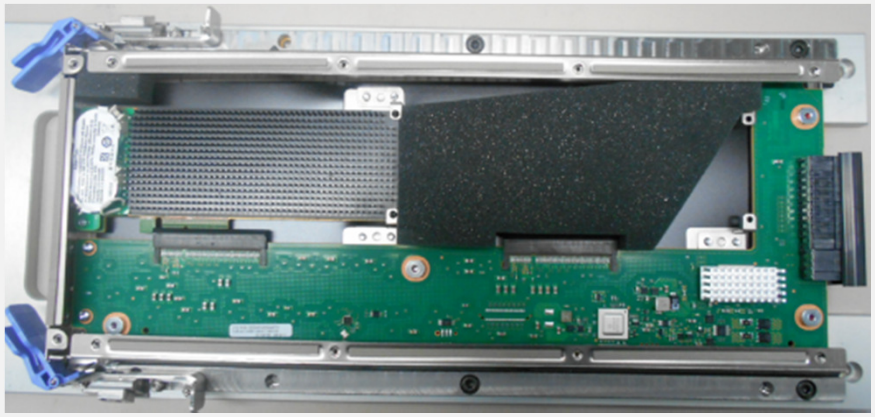

FC 0899

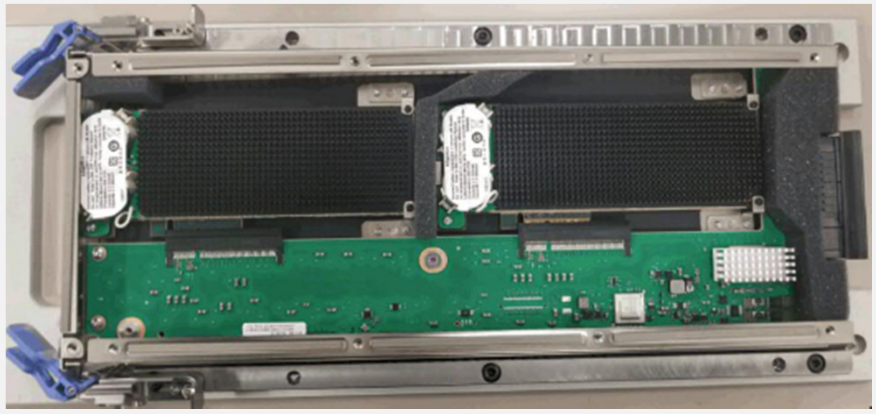

FC 0898

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

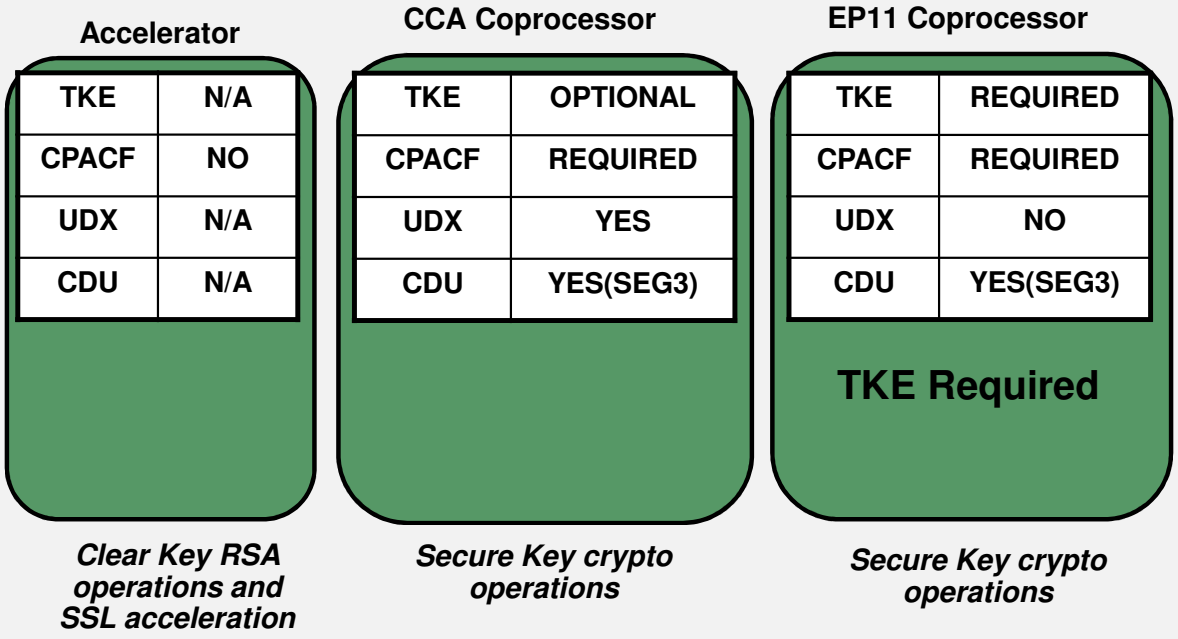

Trusted Key Entry and TKE LIC

- **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

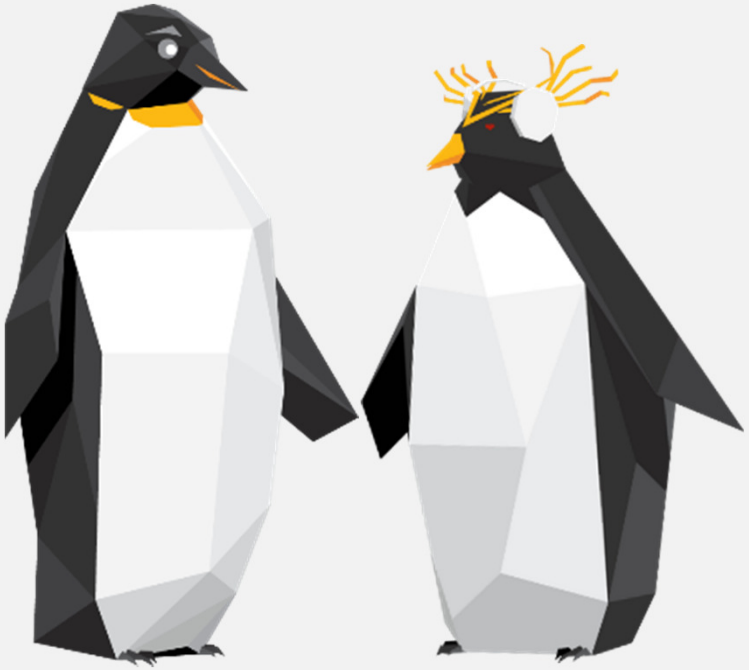

FCP Express32S

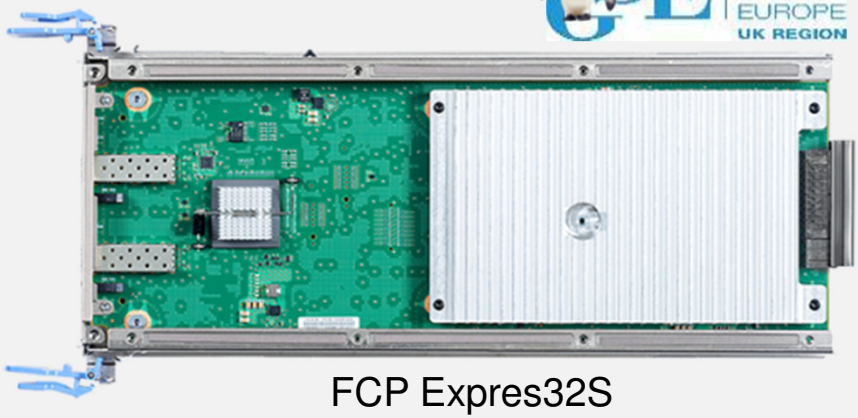

FCP Expres32S

- **For FCP Only**
  - FC 0438 LX & 0439 SX
  - CHPID types: FCP
    - 2 PCHIDs/CHPIDs
- **Auto-negotiates to 8, 16, or 32 Gbps**
- **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 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.

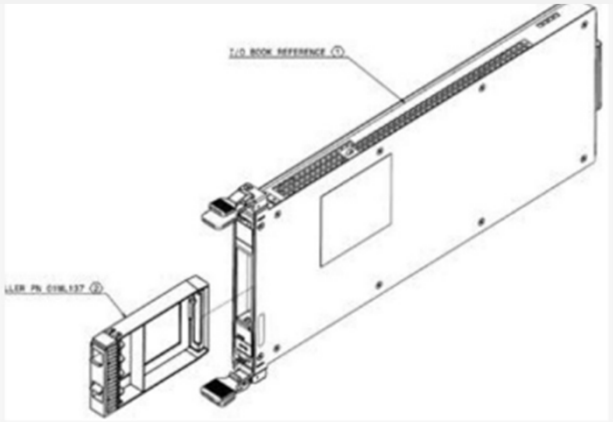

- **Up to 16 features in increments of one.**

- **The vendor SSD card will be purchased by the client from a reseller or directly from the vendor.**

- **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.**

z15 New and Interesting

z15 IBM Integrated Accelerator for zEnterprise Data Compression (zEDC)
(On-chip compression accelerator - Nest Acceleration Unit)

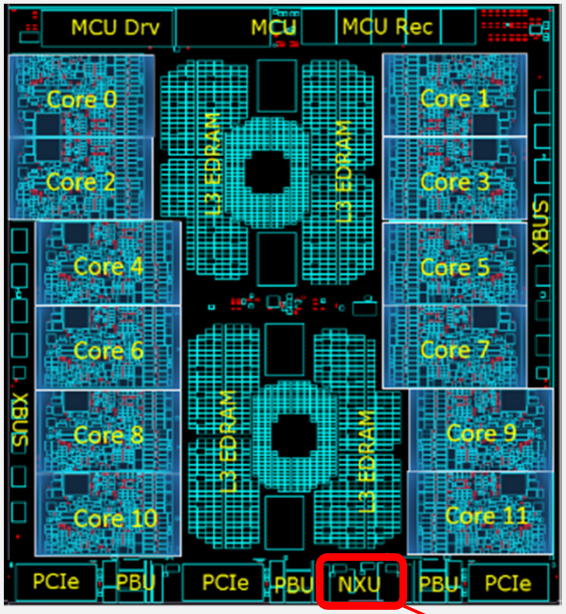

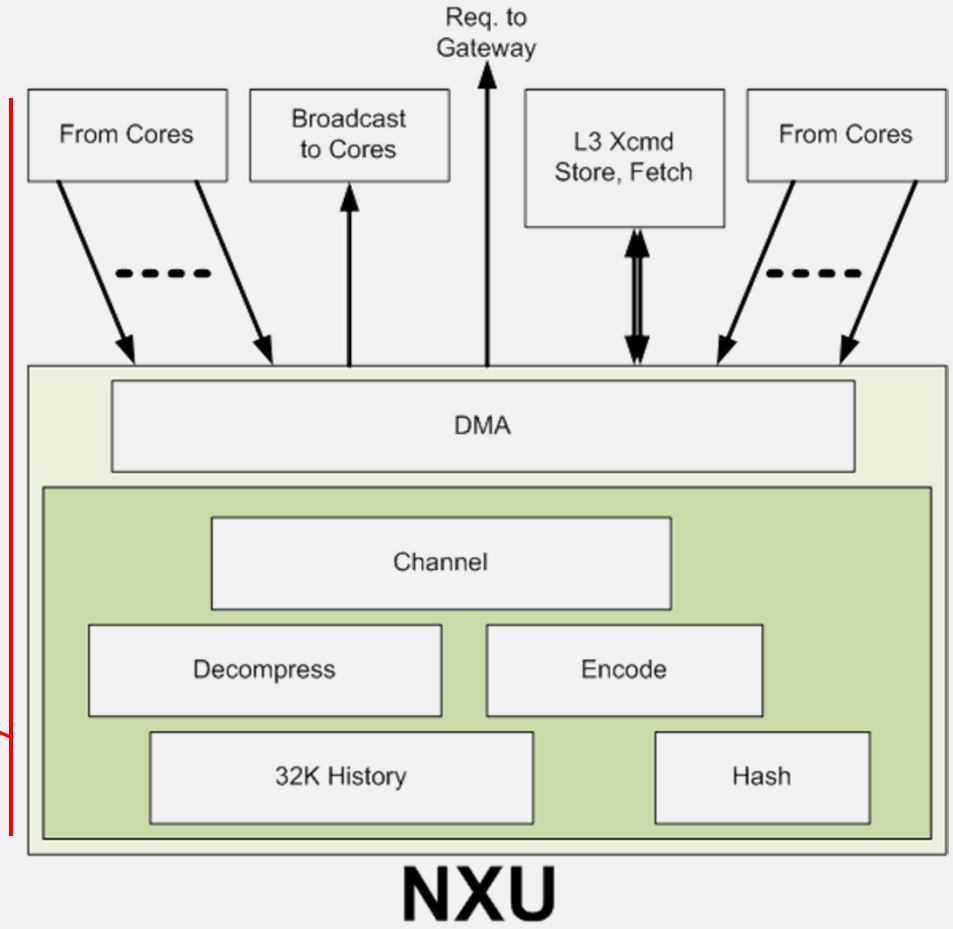

z15 Integrated Accelerator for zEDC

- 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 features will NOT CARRY FORWARD

z15 Compression VALUE

What changes moving to z15 from a z14 with zEDC Express?

Improved Performance

Improved compression ratio means data gets smaller which increase the value of compression on z15

Better latency for synchronous callers of DFLTCC with minimal change in CPU cost

Simplification

Total system throughput increase relieves clients from managing capacity for zEDC Express cards

Easier path forward to enable more applications for compression

Increased Exploitation

New architected instruction opens up potential new use-cases for Linux on Z and z/TPF

zEDC Express to z15 Migration Considerations for z/OS

- **All z/OS configuration stay the same**
  - Nothing required to change when z/OS is moved from a z14 to z15
- **Fail-over and DR should be re-considered**
  - Ensure enough zEDC Express capacity on z13 and z14 systems
- **Performance Metrics**
  - No more RMF PCIE reporting for zEDC
  - Synchronous executions are not recorded (just an instruction invocation)
  - Asynchronous execution are recorded
    - SMF30 information captured for asynchronous usage
    - RMF EADM reporting enhanced (RMF 74.10) with information
    - SAP utilization updated to include time spent compressing + decompressing
- **zlib compatibility**
  - Industry standard interfaces provided by Open Source zlib library shipped in z/OS
  - Products using zlib for zEDC Express on z14 transparently migrate to Integrated Accelerator for zEDC

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 CPC 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

* Disclaimer: Measurements were collected in a controlled environment running an IBM developed workload under z/OS comprised of an equal mix of compression and decompression. Individual results may vary. Results are workload dependent

Fibre Channel Endpoint Security*

Clients who desire to ensure business and customer data is accessed only by trusted servers and storage devices within and across datacenters can leverage the IBM Z Fibre Channel Endpoint Security …

without application, operating system, or file system changes and without consuming host CPU cycles.

- **Enabled automatically between host and storage endpoints that are 'security-capable'**
- **Each established link must 'prove' its identity as a trusted component**
- **Trusted connections are identified; visible to both OS and HMC**
- **Policy can be established to enforce that only trusted connections can be made**
- **Each time a link goes down/up re-authentication/negotiation of encryptions keys occurs**
- **Zero host CPU cost**
- **Integrated key management**

***Statement of Direction**

IBM Z Hardware Management Appliance

- 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

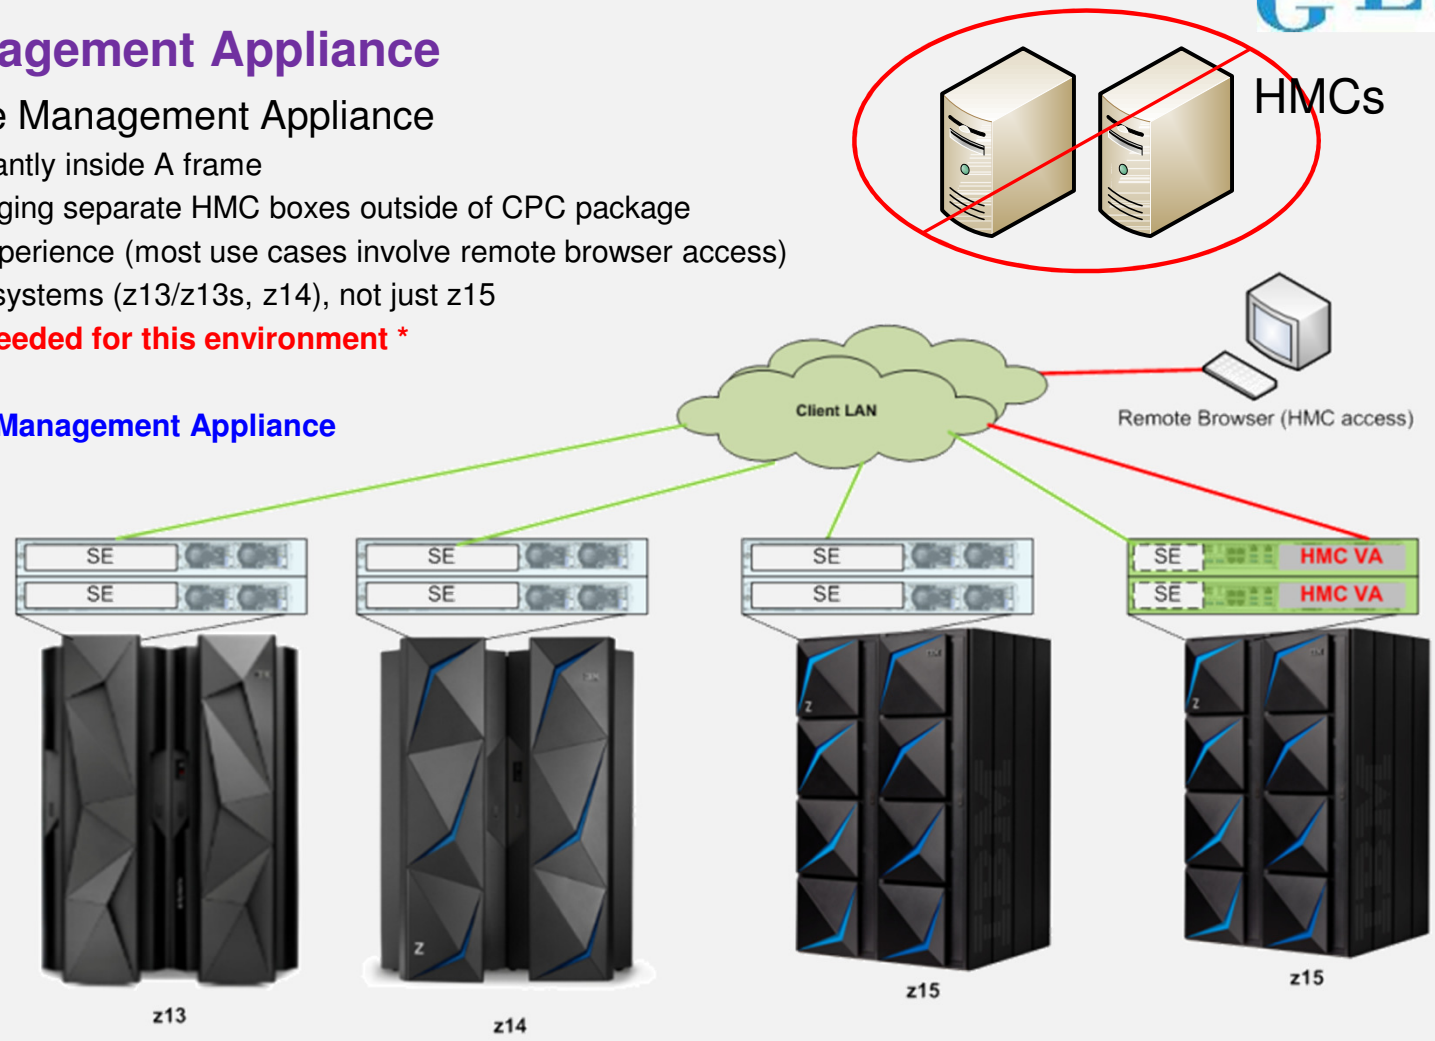

Protect sensitive data in dumps

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.

Protect sensitive data in diagnostic dumps – how to?

Must identify sensitive data in order to secure it

- Tag via z/OS APIs
  - Sensitive=yes, if user data
  - Sensitive=no, if meta data

Detect via Machine Learning (ML)

- Interrogate untagged data in a complete dump
- Tag if sensitive data is discovered

Ways of protecting sensitive data

- Redacting or sanitizing
  - Enforced data (irreversible)
- Encrypting (not MVP)
  - Protected data (reversible)

What is Precision Time Protocol (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				Drawer 3				Drawer 4				Drawer 5			
Cluster 1		Cluster 2		Cluster 1		Cluster 2		Cluster 1		Cluster 2		Cluster 1		Cluster 2		Cluster 1		Cluster 2	
C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2

- **Chip level optimization**
- **Movement of IFL processors and ICF processors**
  - Isolation of CP+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)**

Check placement every 50 seconds

- **Default processor assignments by POR, MES adds, and On Demand activation**:
  - Assign IFLs and ICFs to cores on chips in "high" drawers working down
  - Assign CPs and zIIP in low drawers working up.
  - Objective: Keep "Linux Only" and "Coupling Facility" using IFLs and ICFs "away" from partitions running z/OS on CPs and zIIPs and, in different drawers, if possible.
- ***New in z15: All processor types can be also dynamically reassigned, except IFPs***

Logical Partition Isolation by Container

Drawer 1				Drawer 2				Drawer 3				Drawer 4				Drawer 5			
Cluster 1		Cluster 2		Cluster 1		Cluster 2		Cluster 1		Cluster 2		Cluster 1		Cluster 2		Cluster 1		Cluster 2	
C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2	C1	C2

- **PR/SM makes optimum available memory and logical processor assignment at activation**
  - Logical Processors specified in the Image Profile, are assigned a core if Dedicated or a “home” drawer, node and chip if Shared. Later, if it becomes a HiperDispatch “Vertical High”, a Shared Logical Processor is assigned a specific core.
  - Ideally assign all memory in one drawer with the processors if everything “fits”
  - With memory striped across drawers with processors if memory or processors must be split
- **PR/SM optimizes resource assignment dynamically**
  - Examines partitions in priority order by the size of their “processor entitlement” (dedicated processor count or shared processor pool allocation by weight) to determine priority for optimization
  - Changes logical processor “home” drawer/node/chip assignment
  - Moves processors to different chips, nodes, drawers (LPAR Dynamic PU Reassignment)
  - Relocates partition memory to active memory in a different drawer or drawers using the newly optimized Dynamic Memory Relocation (DMR), also exploited by Enhanced Drawer Availability (EDA).
  - If available but inactive memory hardware is present (e.g. hardware driven by Flexible Memory) in a drawer where more active memory would help: activate it, reassign active partition memory to it, and deactivate the source memory hardware, again using DRM.
    (PR/SM can use all memory hardware but concurrently enables no more memory than the client has paid to use.)

PR/SM Partition Logical Processor and Memory Assignment

- **z15**
  - ***Improved memory affinity algorithm and improved logical partition placement algorithms based on z14 experience***
  - ***PR/SM re-optimization is the process of identifying "homes" for the partitions. Home of a partition is the node or set of nodes(s)/ drawer or set of drawers from which the partition's physical resources (processors and memory) will come from.***
  - ***Memory Allocation Goal***: Assign all memory in one drawer (single SC SCM, shared L4 cache). Improved performance due to lower latency memory access in drawer.
    - ***New for z15: Consolidation of storage onto drawers with the most processor entitlement***
  - ***Processor Allocation Goal***: Assign all logical processors to one CPC drawer; packed into chips of that drawer. Cooperate with operating system use of HiperDispatch
    - ***New for z15: All processor types can be dynamically reassigned except IFPs***
- **z14**
  - ***Memory Allocation Goal***: Assign all memory in one drawer (single SC SCM, shared L4 cache). Improved performance due to lower latency memory access in drawer (S- Bus has been eliminated)
  - ***Processor Allocation Goal***: Assign all logical processors to one CPC drawer; packed into chips of that drawer. Cooperate with operating system use of HiperDispatch
  - ***Optimization complexity*** for multiple logical partitions ***relieved*** due to increased memory support per drawer (8 TB).
- **z13**
  - ***Memory Allocation Goal***: Assign all memory in one drawer striped across the two nodes.
    Advantage: Lower latency memory access in drawer; smooth performance variability across nodes in the drawer
  - ***Processor Allocation Goal***: Assign all logical processors to one CPC drawer; packed into chips of that drawer. Cooperate with operating system use of HiperDispatch
  - ***Reality***: Easy for any given partition. Complex optimization for multiple logical partitions because some need to be split among drawers.

Capacity on Demand

Capacity on Demand (CoD) Offerings

- **On-line Permanent Upgrade**
  - Permanent upgrade performed by customer
- **Capacity Backup (CBU)**
  - For disaster recovery
  - Concurrently add CPs, IFLs, ICFs, zIIPs, SAPs
  - Pre-paid
- **Capacity for Planned Event (CPE)**
  - To replace capacity for short term lost within the enterprise due to a planned event such as a facility upgrade or system relocation
  - Predefined capacity for a fixed period of time (3 days)
  - Pre-paid
- **On/Off Capacity on Demand (On/Off CoD**
  - Production Capacity
  - Supported through software offering – Capacity Provisioning Manager (CPM)
  - Payment:
    - Post-paid or Pre-paid by purchase of capacity tokens
    - Post-paid with unlimited capacity usage
    - On/Off CoD records and capacity tokens configured on Resource Link
- **Customer Initiated Upgrade (CIU)**
  - Process/tool for ordering temporary and permanent upgrades via Resource Link
  - Permanent upgrade support:
    - Un-assignment of currently active capacity
    - Reactivation of unassigned capacity
    - Purchase of all PU types physically available but not already characterized
    - Purchase of installed but not owned memory
- **New with z15: System Recovery Boost record:**
  - Pre-paid temporary zIIP capacity for boost events – shutdown (30 min/event) , restart (IPL) (60 minutes/event):
  - Contract enablement (boost authorization): 9930
  - 20 zIIPs * 6 Hours for **one year**: FC 6802

z15 Basics of Capacity on Demand

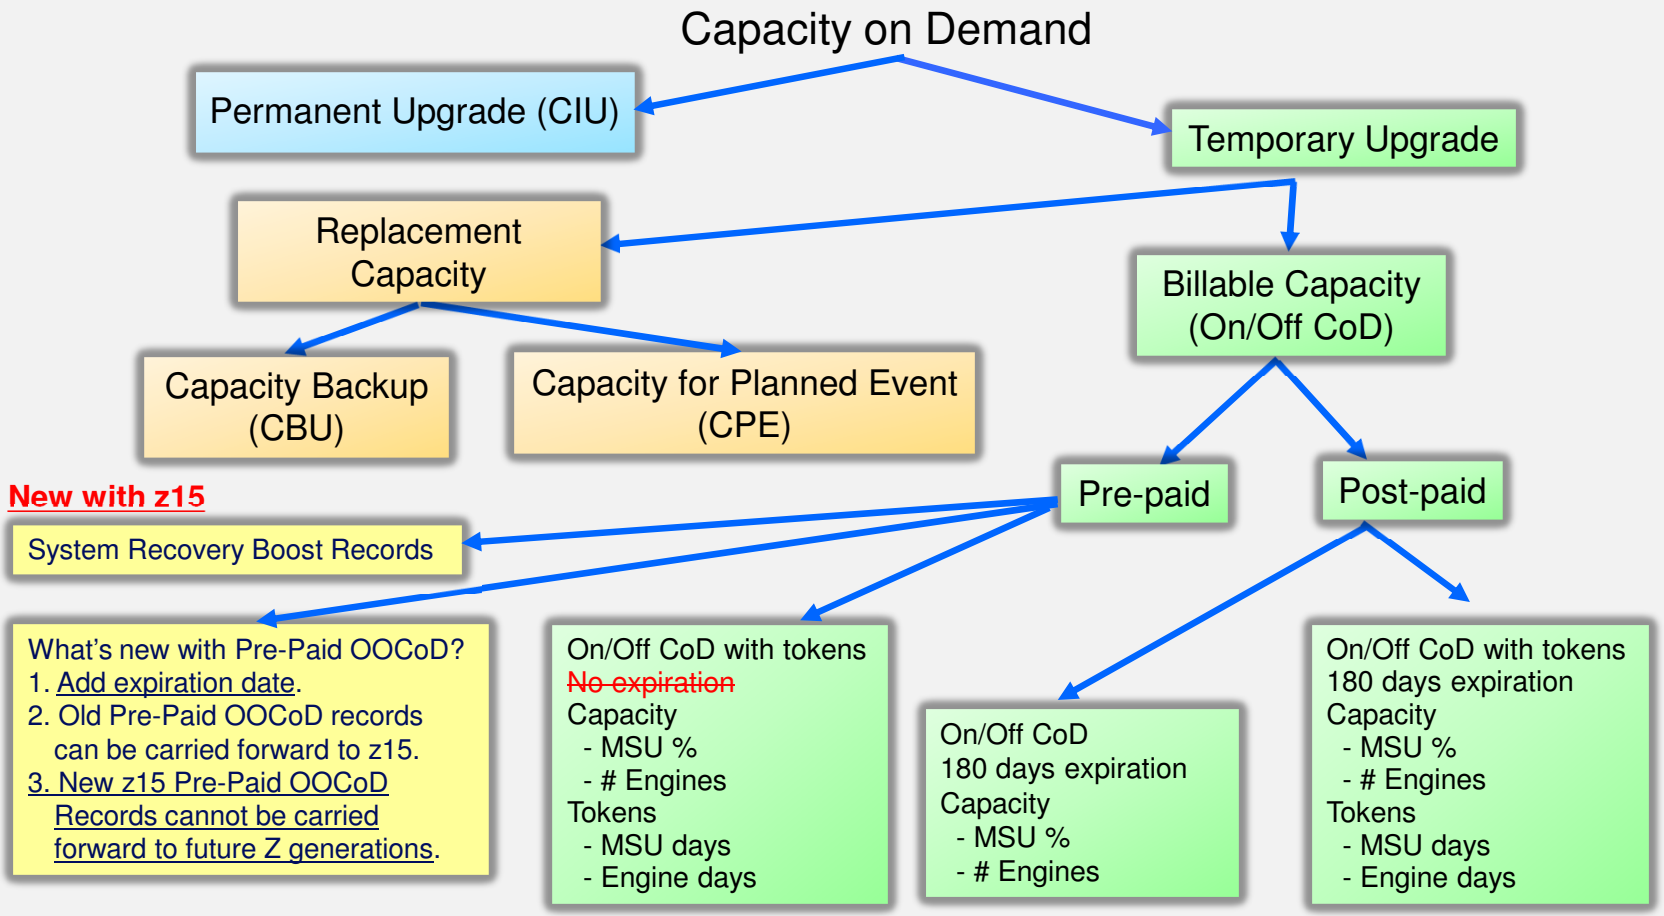

Any Questions?

Thanks!

Please submit your session feedback!

- **Do it online at** http://conferences.gse.org.uk/2019/feedback/nn
- **This session is** BO