The IBM z13 Part II: Crypto, I/O Design, Features and Functions, Parallel Sysplex and Implementation Planning (17435)

Harv Emery
Executive IT Specialist, z Systems Hardware
IBM Washington System Center

Please evaluate sessions:
www.share.org/Orlando-Eval

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®, GDPS®, IBM®, IBM (logo)®, IMS, MVS, OS/390®, POWER7®, Power Architecture®, PowerVM®, S/390®, ServerProven®, Sysplex Timer®, System p®, System x®, System z®, System z9®, System z10®, Tivoli®, WebSphere®, X-Architecture®, z Systems, z9®, z10, z13, z/Architecture®, z/OS®, z/VM®, z/VSE®, zEnterprise®, zSeries®

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.
Oracle®, Java and all Java-based trademarks are trademarks of Oracle Corporation and/or its affiliates in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured Sync new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained Sync the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

IBM z13 Launch

IBM z13 platform positioning

Platform Core Capabilities:

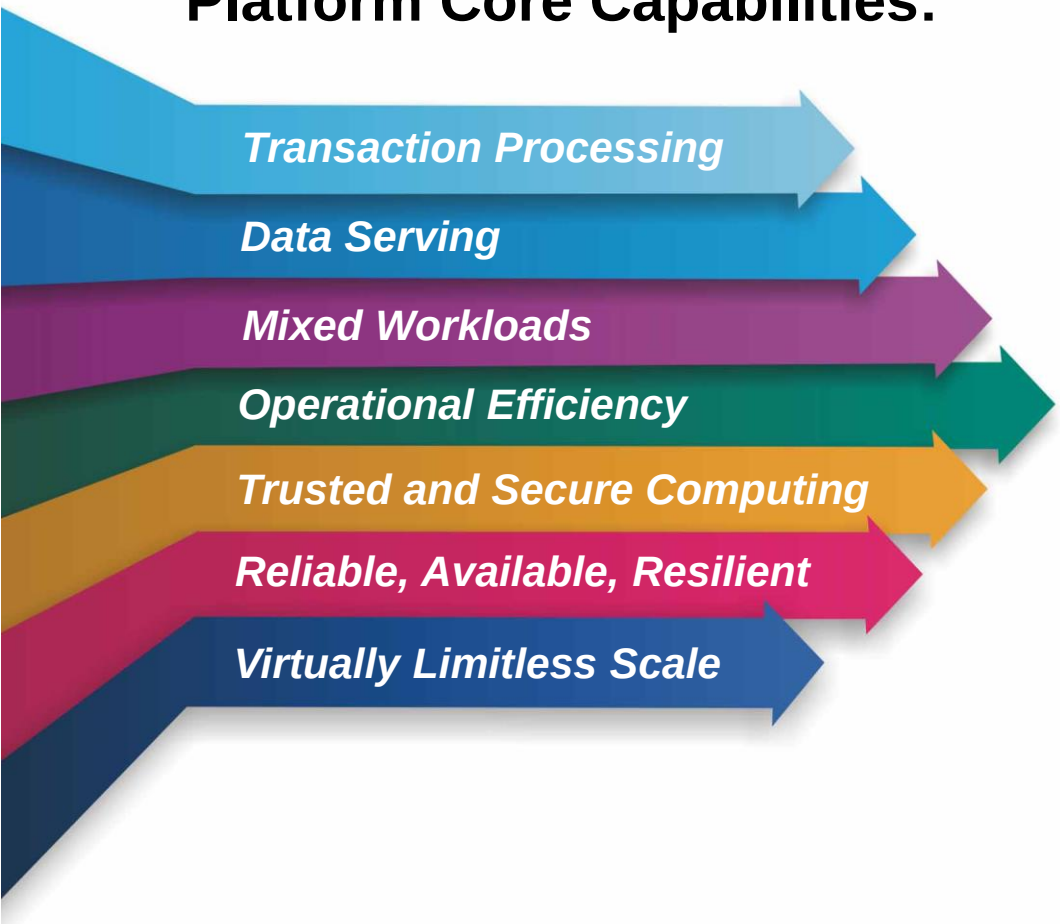

- *The world's premier transaction and data engine now enabled for the **mobile** generation*
- *The integrated transaction and **analytics** system for right-time insights at the point of impact*
- *The world's most efficient and trusted **cloud** system that transforms the economics of IT*

IBM z Systems Server Generations

z196 "N-2"

- Announced: 7/22/2010
- End of mkt: **6/30/2014 (LIC 2015)**
- Chip: 4 core, *5.2 GHz*
- Key z/Arch: *Out of order execution*
- Up to 80 client cores
- CP, IFL, ICF, zAAP, zIIP
- Single thread
- zIIP-zAAP to CP ratio 1x1
- Uni MIPS: 1,202
- Max MIPS: 52,286
- Max mem *3 TB (RAIM)*
- Max mem/LPAR: 1 TB
- LCSS: 4, LPARs: 60
- Subchannel Sets: 3/LCSS
- Max I/O Slots: 160*
- IO buses: 6 GBps (*8 GBps GA2*)
- Max FICON Channels: 320
- FICON Express8 (*8S GA2*)
- Max OSA Ports: 96
- OSA-Express4S (GA2)
- Crypto Express3
- Coupling: *128 CHPIDs*, ISC3, PSIFB DDR: 12x, 1x
- CTN: STP Mixed
- *zBX Model 2 (Attached)*

zE12 "N-1"

- Announced: 8/28/2012
- End of mkt: **Not yet announced**
- Chip: *6 core, 5.5 GHz*
- Key z/Arch: *Transactional Exec*
- Up to *101* client cores
- CP, IFL, ICF, zAAP, zIIP
- Single thread
- zIIP-zAAP to CP *ratio 2x1 (GA2)*
- Uni MIPS: 1,514
- Max MIPS: 78,426
- Max mem 3 TB (RAIM)
- Max mem/LPAR: 1 TB
- LCSS: 4, LPARs: 60
- Subchannel Sets: 3/LCSS
- Max I/O Slots: 160*
- I/O buses: 8 GBps
- Max FICON Channels: 320
- FICON Express8S
- Max OSA Ports: 96
- OSA-Express4S *(5S GA2)*
- *Crypto Express4S*
- Coupling: 128 CHPIDs, PSIFB DDR: 12x, 1x
- CTN: STP mixed
- *Flash Express*
- *Native PCIe: zEDC, RoCE (GA2)*
- *zBX Model 3 (Attached)*

z13 "N"

- Announced 1/14/2015
- End of mkt: Current generation
- *Chip: 8 core,* 5.0 GHz
- Key z/Arch*: SMT, Vector (SIMD)*
- Up to *141* client cores
- CP, IFL, ICF, zIIP *(No zAAPs)*
- *SMT Support: zIIP, IFL*
- zIIP to CP ratio 2x1
- Uni MIPS: *1,695*
- Max MIPS: *111,566*
- Max mem *10 TB* (RAIM)
- Max LPAR: *10 TB (z/OS: 4TB)*
- *LCSS: 6, LPARs: 85*
- Subchannel Sets: *4/LCSS*
- Max I/O Slots: 160*
- I/O buses: *16 GBps*
- Max FICON Channels: 320
- *FICON Express16S**
- Max OSA Ports: 96
- *OSA-Express5S**
- *Crypto Express5S*
- Coupling: *256 CHPIDs*, *ICA 8 GBps*, PSIFB DDR: 12x, 1x
- *CTN: STP Only*
- *Flash Express (refresh)*
- Native PCIe: zEDC, *RoCE**
- *zBX Model 4 (Independent)*

** Major functional enhancements*

IBM z13 Key Planning and Support Dates

- **January 14, 2015 – Announcement Day**
  - IBM United States Hardware Announcement Letter 115-001
- **March 9, 2015 – General Availability** ✓
- **June 26, 2015** ✓ **(Bundle S14, Driver 22. Released June 10)**
  - Field install of MES hardware features for z13 Models N30, N63, N96, NC9, and NE1
  - z/VM V6.3 support for Multi-VSwitch Link Aggregation
  - Support for 256 Coupling CHPIDs
  - HMC STP Panel Enhancements: Initialize Time, Set Date and Time, Time Zone, View-Only Mode
  - Fibre Channel Protocol (FCP) channel configuration discovery and debug
  - Improved High Performance FICON for z Systems (zHPF) I/O Execution at Distance
  - IBM zAware support for Linux on z Systems
- **September 25, 2015**
  - FICON Dynamic Routing
  - Forward Error Correction (FEC) for FICON Express16S
  - Storage Area Network (SAN) Fabric I/O Priority

IBM z13 Model Structure and Performance

z13 System Offering Overview

zEC12 zBX Model 3

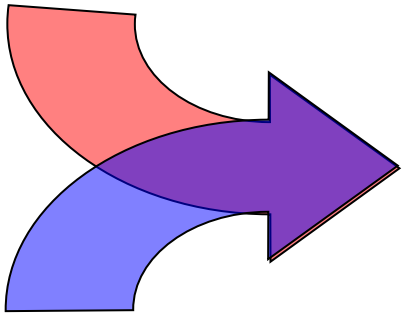

z196 zBX Model 2

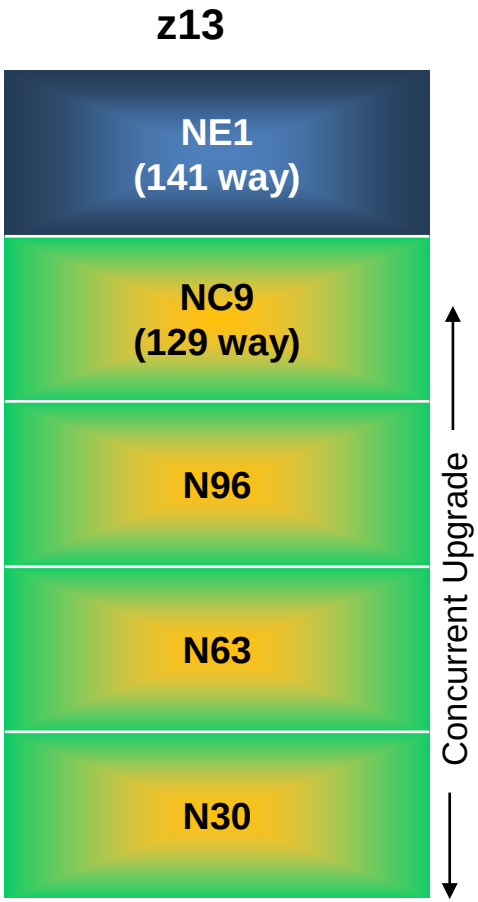

Machine Type for z13

- 2964

Processors

- 39 PUs per drawer (42 in NE1)
- Sub-capacity available up to 30 CPs
- 2 standard spare PUs per system

Memory

- System minimum = 64 GB with separate 96 GB HSA
- Maximum: ~10 TB / ~2.5TB per drawer
- RAIM memory design
- Purchase Increments – 32 to 512 GB

I/O

- Up to 14 fanouts per drawer
  - **Up to 10 PCIe Gen 3 fanouts: 1-port 16 GBps I/O or 2-port 8 GBps PCIe coupling**
  - Up to 4 IFB HCA fanouts: 2-port 6 GBps I/O, 2-port 12x PSIFB, or 4-port 1x PSIFB

On upgrade from zEC12 or z196

- Detach zBX Model 3 or 2 and upgrade it to zBX Model 4 (Option: Move zBX Model 3)
- Feature convert installed zAAPs to zIIPs (default) or another processor type
- For installed On Demand Records, change temporary zAAPs to zIIPs. Stage the record

z13 8-Core Processor Chip Detail

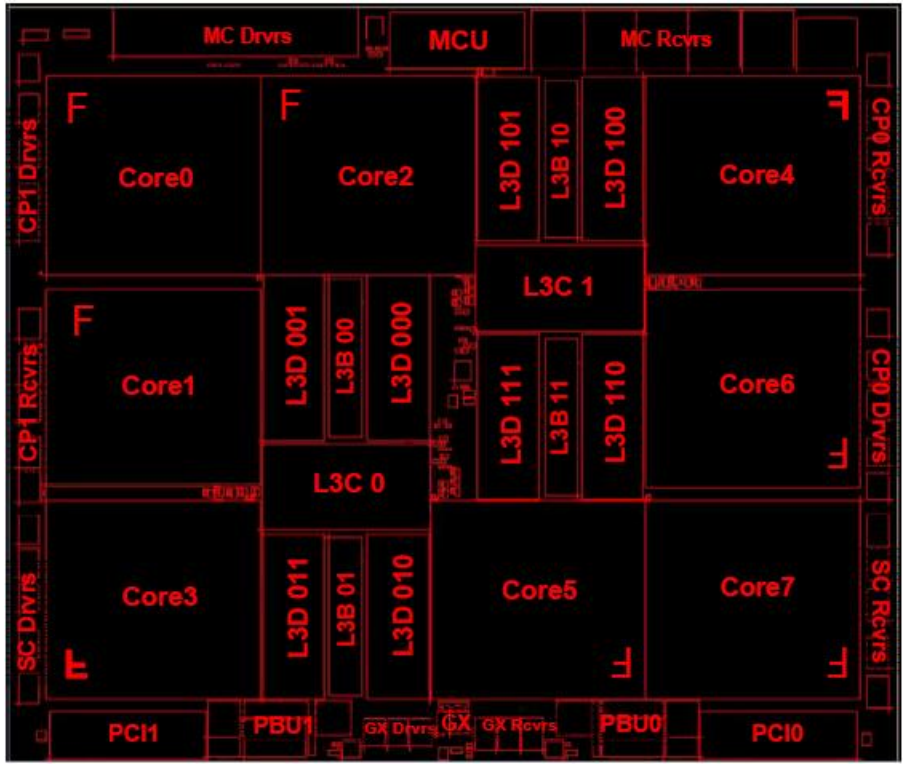

- **14S0 22nm SOI Technology**
  - 17 layers of metal
  - 3.99 Billion Transistors
  - 13.7 Miles copper wire
- **Chip Area**
  - 678.8 mm²
  - 28.4 x 23.9 mm
  - 17,773 power pins
  - 1,603 signal I/Os

- **Up to eight active cores (PUs) per chip**
  - **5.0 GHz (v5.5 GHz zEC12)**
  - L1 cache/ core
    - 96 KB I-cache
    - 128 KB D-cache
  - L2 cache/ core
    - 2M+2M Byte eDRAM split private L2 cache
- **Single Instruction/Multiple Data (SIMD)**
- **Single thread or 2-way simultaneous multithreading (SMT) operation**
- **Improved instruction execution bandwidth:**
  - Greatly improved branch prediction and instruction fetch to support SMT
  - Instruction decode, dispatch, complete increased to 6 instructions per cycle
  - Issue up to 10 instructions per cycle
  - Integer and floating point execution units
- **On chip 64 MB eDRAM L3 Cache**
  - Shared by all cores
- **I/O buses**
  - One GX++ I/O bus
  - Two PCIe I/O buses
- **Memory Controller (MCU)**
  - Interface to controller on memory DIMMs
  - Supports RAIM design

IBM z13
I/O Subsystem
Introduction

z13 Model NE1 or NC9 Radiator (Air) Cooled – Under the covers (Front View)

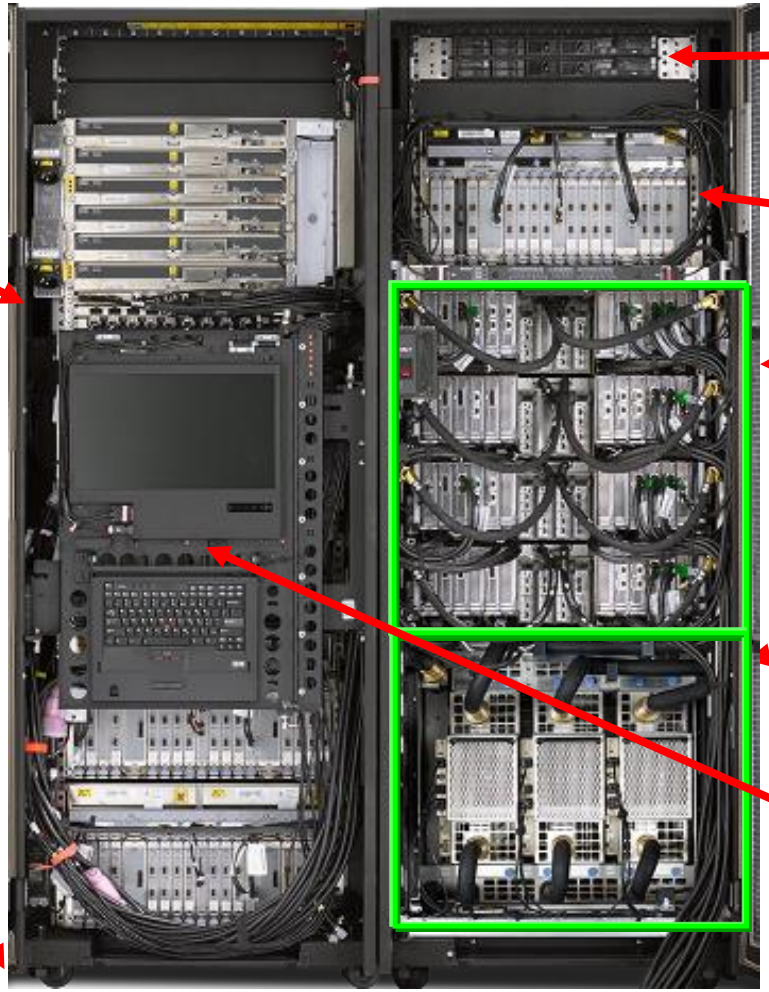

z13 Processor Drawer Connectivity for I/O and Coupling

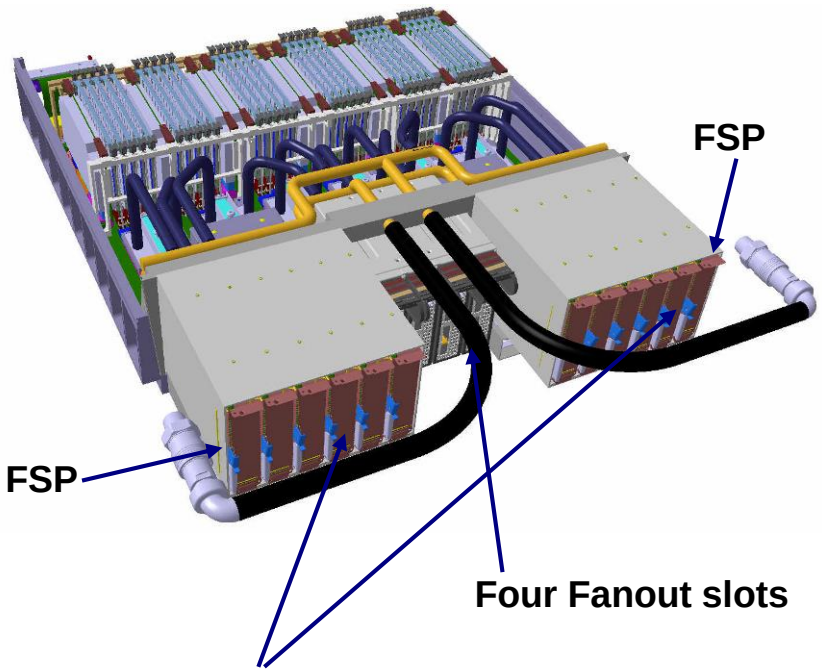

- **Ten PCIe fanout slots per drawer (40 maximum)**
  - **ICA (PCIe-O SR) two-port 8 GBps PCIe Gen3 fanout 150 meter fiber optic coupling link**

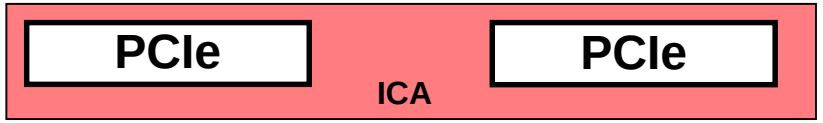

  - **PCIe Gen3 one-port 16 GBps PCIe fanout connects to a switch card for an 8-slot PCIe I/O domain (Plugs in pairs)**

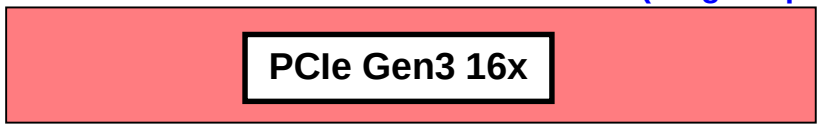

- **Four IFB HCA (GX++) fanout slots per drawer**
  (16 maximum on a four drawer system)
  - HCA2-C 2-port 6 GBps I/O drawer fanout (plugs in pairs)
  - HCA3-O 2-port 12x IFB Coupling Link fanout
  - HCA3-O LR 4-port 1x IFB Coupling Link fanout

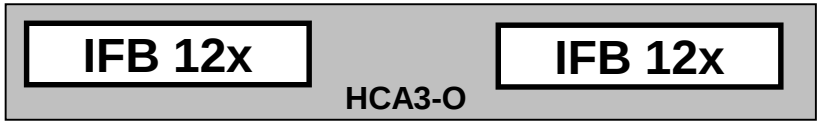

Carry forward
(One pair only)

Carry forward or
New Build

CPC Drawer I/O Fanout and FSP Locations

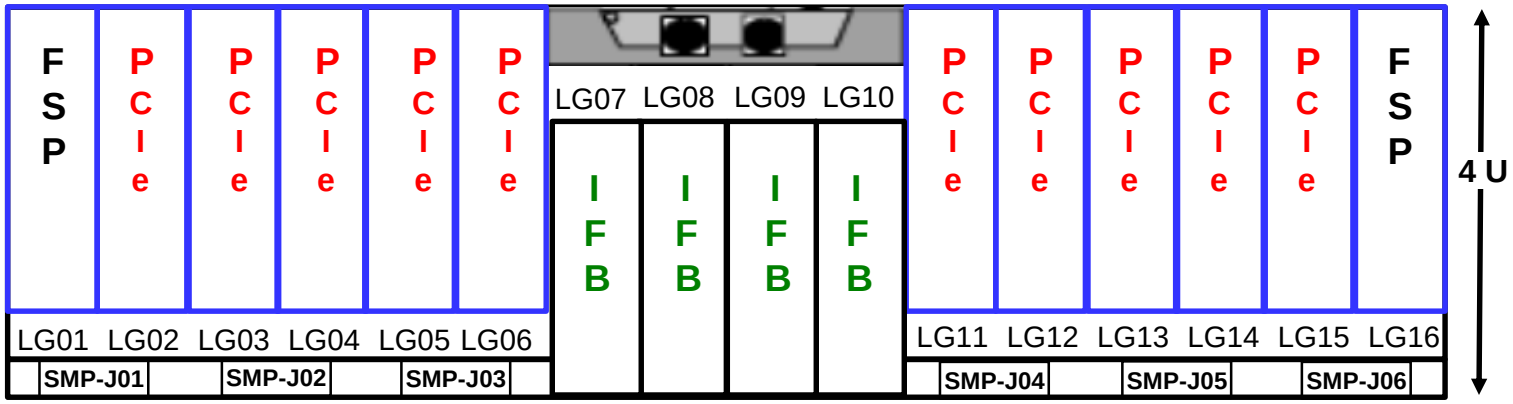

- PCIe Fanout Slots (Ten), slots LG02 – LG06 and LG11 – LG15, can support:
  - Up to 10 one-port PCIe 16 GBps I/O fanouts to support up to 10 domains in 32-slot PCIe I/O drawers
    Note: A zEC12 book with eight two-port 8 GBps PCIe fanouts supports up to 16 domains in 32-slot PCIe I/O drawers; but **a z13 CPC drawer supports double the bandwidth to each domain**
  - Up to 10 ICA (PCIe-SR) two-port coupling fanouts to support up to 20 8 GBps coupling links
- IFB Fanout Slots (Four), LG07 – LG10, can support:
  - Up to four HCA3-O 12x InfiniBand coupling fanouts, 8 12x 6 GBps links – Two per fanout
  - Up to four HCA3-O LR 1x InfiniBand coupling fanouts 16 1x 5 Gbps links – Four per fanout
    Note: A zEC12 book with 8 two-port HCA3-O 12x InfiniBand coupling fanouts can support 16 12x links
    A zEC12 book with 8 four-port HCA3-O LR 1x InfiniBand coupling fanouts can support 32 1x links
  - Up to two two-port HCA2-C 6GBps I/O fanouts (2 8-slot I/O drawers) with two slots left
- Slots LG01 and LG16 always have Flexible Support Processors (FSPs)
- SMP-J01 to J06 connectors are for A-Bus cables to nodes in other CPC drawers

z Systems I/O Subsystem Internal Bus Interconnect Speeds

PCIe Gen3
z13

16 GBps

PCIe Gen2
zEC12/zBC12/
z196/z114

8 GBps

InfiniBand
z10/z196/z114/
zEC12/zBC12

6 GBps

STI
z9

2.7 GBps

STI
z990/z890

2 GBps

STI: Self-Timed Interconnect

IBM z13: Optimized as a system of insight for digital business

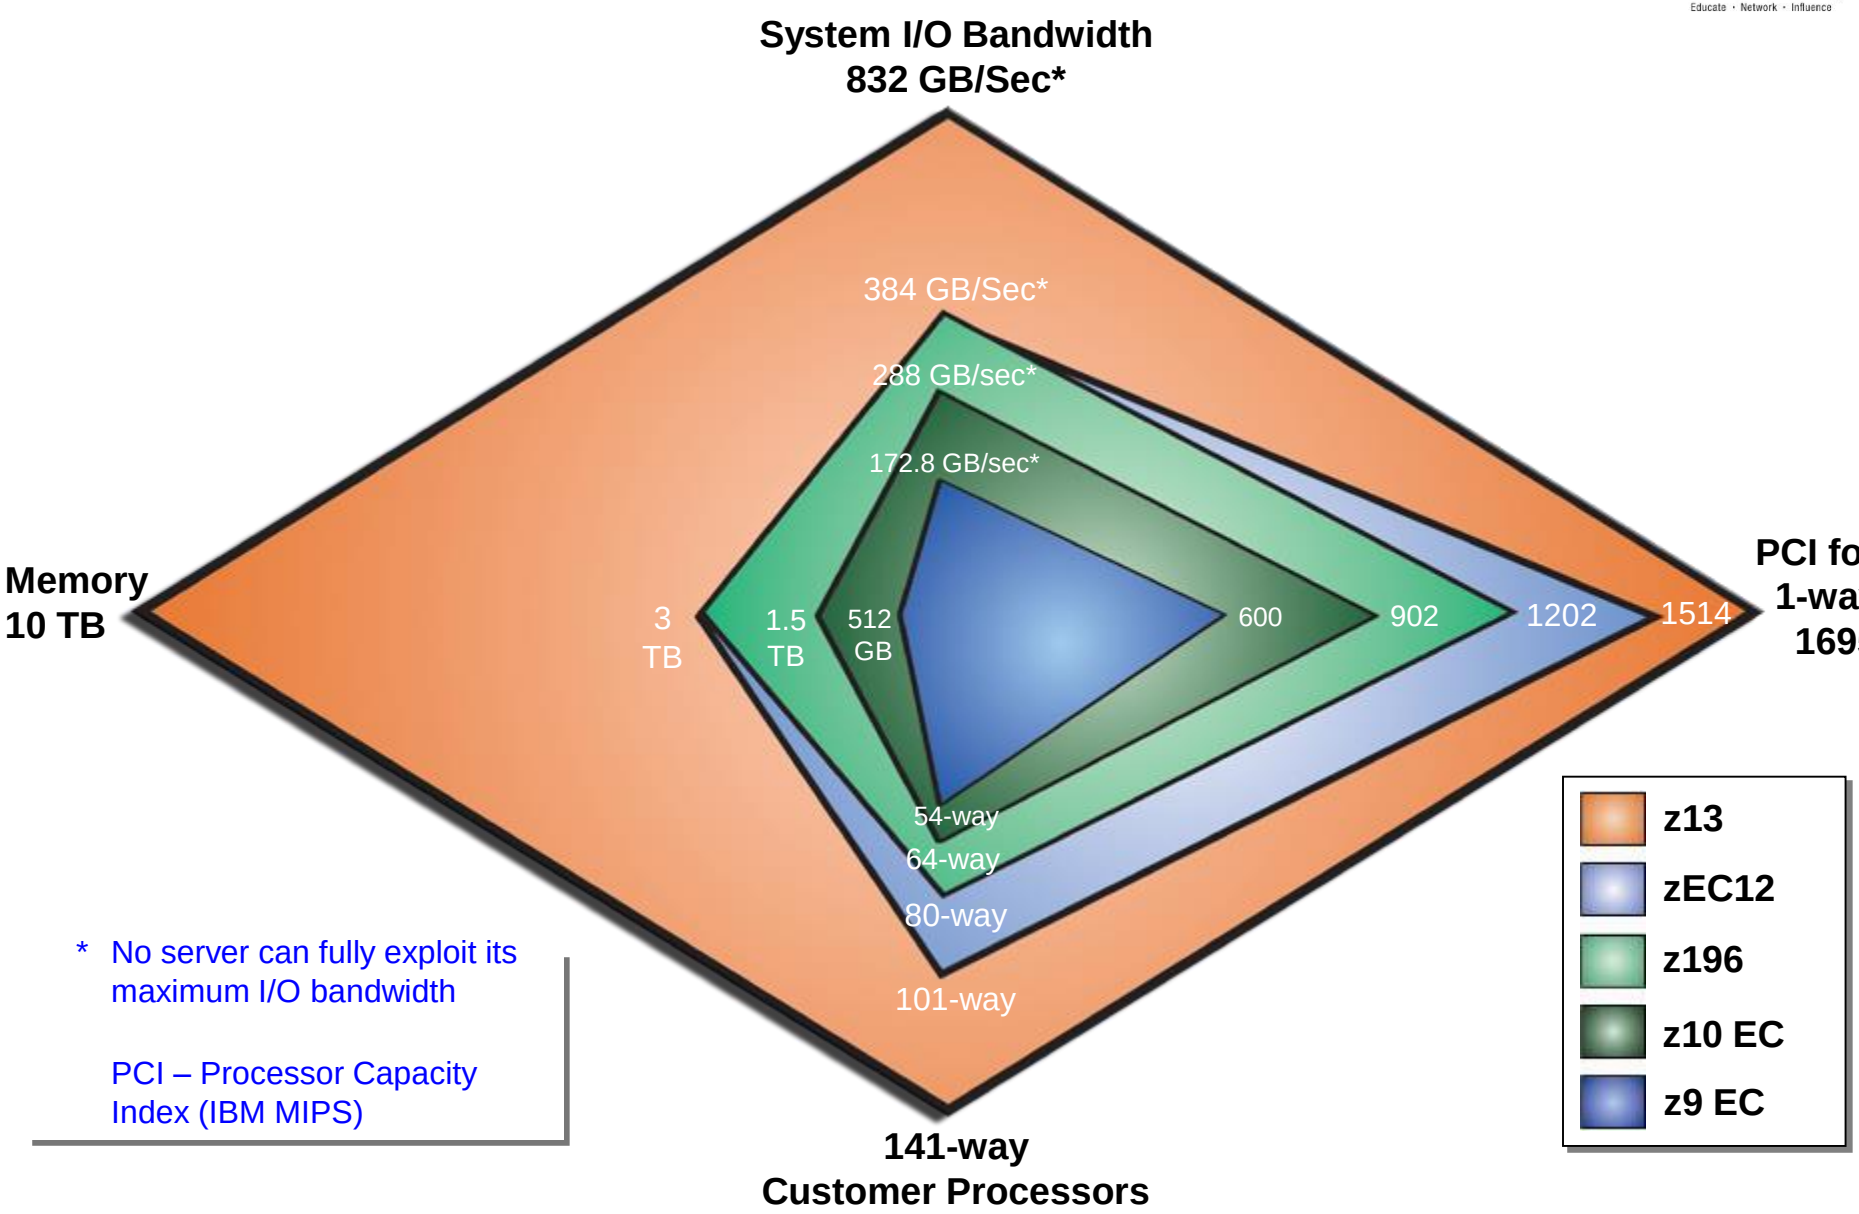

* No server can fully exploit its maximum I/O bandwidth

PCI – Processor Capacity Index (IBM MIPS)

PCIe 32 I/O slot drawer

Front

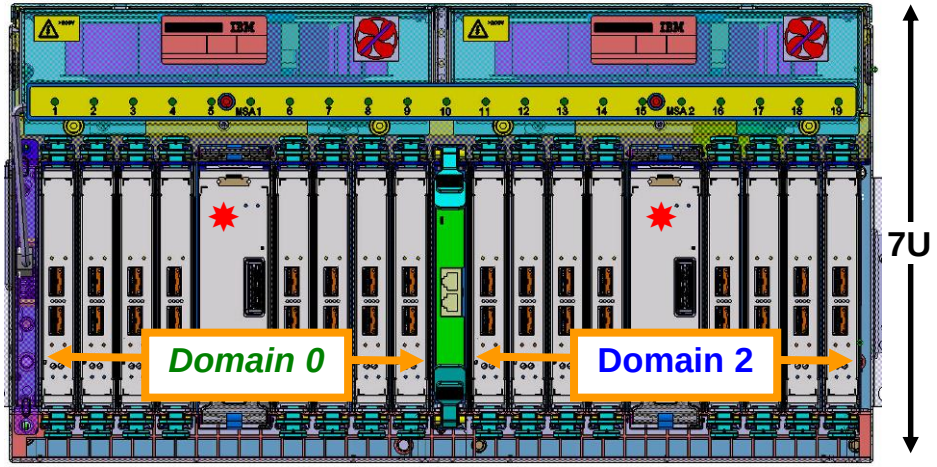

Rear

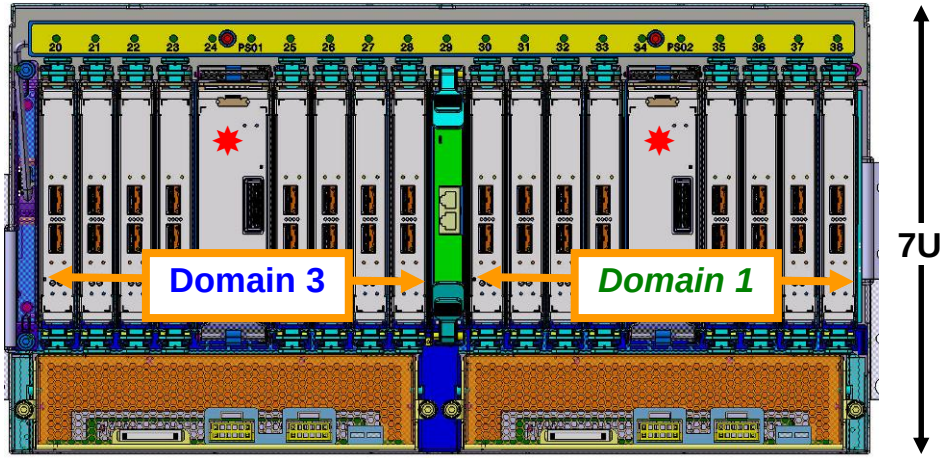

- Supports only PCIe I/O cards
  - z13: Up to five drawers
  - zEC12: Up to five drawers
- Supports 32 PCIe I/O cards, 16 front and 16 rear, vertical orientation, in four 8-card domains (shown as 0 to 3).
- Requires four **16 GBps** PCIe switch cards (✸), each connected to a **16 GBps** PCIe I/O interconnect to activate all four domains.
- To support Redundant I/O Interconnect (RII) between front to back domain pairs *0-1* and 2-3 the two interconnects to each pair will be from 2 different PCIe fanouts. (All four domains in one of these cages can be activated with two fanouts.)
- Concurrent field install and repair.
- Requires 7 EIA Units of space (12.25 inches ≈ 311 mm)

Supported I/O Features

z13 "New Build" I/O and MES Features Supported

Note - "Plan Ahead" for I/O drawers is not offered on z13

New Build Features

- **Features – PCIe I/O drawer**
  - ***FICON Express16S*** *(SX and LX, 2 SFPs, 2 CHPIDs)*
  - **FICON Express8S** (SX and LX, 2 SFPs, 2 CHPIDs)
  - **OSA-Express5S**
    - **10 GbE LR and SR** (1 SFP, 1 CHPID)
    - **GbE SX, LX, and 1000BASE-T** (2 SFPs, 1 CHPID)
  - ***10 GbE RoCE Express*** *(2 supported SR ports)*
  - **zEDC Express**
  - ***Crypto Express5S***
  - ***Flash Express (Technology Refresh)***

PCIe I/O drawer

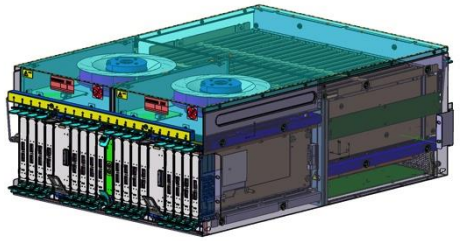

32 I/O slots

- **PCIe Coupling Link Feature (Fanout)**
  - ***ICA PCIe-O SR two 8GBps PCIe Gen3 Coupling Link***
- **InfiniBand Coupling Features (Fanouts)**
  - **HCA3-O two 12x 6GBps InfiniBand DDR Coupling Links**
  - **HCA3-O LR four 1x 5Gbps InfiniBand DDR or SDR Coupling Links**

z13 "Carry Forward" I/O Features Supported

Note – "Plan Ahead" for I/O drawers is not offered on z13

Carry Forward Features

- **Features – PCIe I/O drawer**
  - **FICON Express8S** (SX and LX, 2 SFPs, 2 CHPIDs)
  - **OSA-Express5S (All)**
  - **OSA-Express4S (All)**
  - ***10 GbE RoCE Express*** *(Both ports supported on z13)*
  - **zEDC Express**
  - **Flash Express**
  - **Not Supported: Crypto Express4S**

PCIe I/O drawer 32 I/O slots

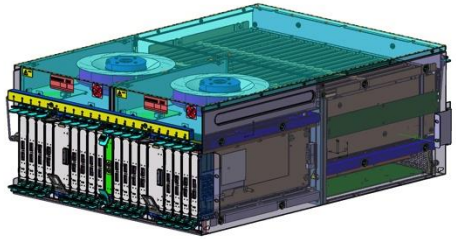

- **Features – I/O drawer (No MES adds)**
  - **FICON Express8** (SX and LX, 4 SFPs, 4 CHPIDs)
    SoD: IBM plans not to support FICON Express8 on the next high end z Systems server.
  - **Not Supported: ESCON, FICON Express4, OSA-Express3, ISC-3, and Crypto Express3**

I/O drawer 8 I/O slots

- **InfiniBand Coupling Features (Fanouts)**
  - **HCA3-O two 12x 6GBps InfiniBand DDR Coupling Links**
  - **HCA3-O LR four 1x 5Gbps InfiniBand DDR or SDR Coupling Links**
  - **NOT Supported: HCA2-O 12x, HCA2-O LR 1x InfiniBand Coupling Links**

z13 Carry Forward (Field Upgrade) Rules for I/O Features (All PCIe I/O Features Can be Carried Forward)

FICON Express8 Features Carried Forward	8-slot I/O Drawers Required	Maximum PCIe Drawers/ Slots (CF or Add)
0	0	5/160
1 to 8	1	4/128
9 to 16	2	3/96
17 or more	Not Supported!	

Empty slots in a carried forward drawer can NOT be filled by MES.

SoD: IBM plans not to support FICON Express8 on the next high end z Systems server.

Note: Large I/O configurations may require two or more CPC drawers.

z13 CPC Drawer and I/O Drawer Locations

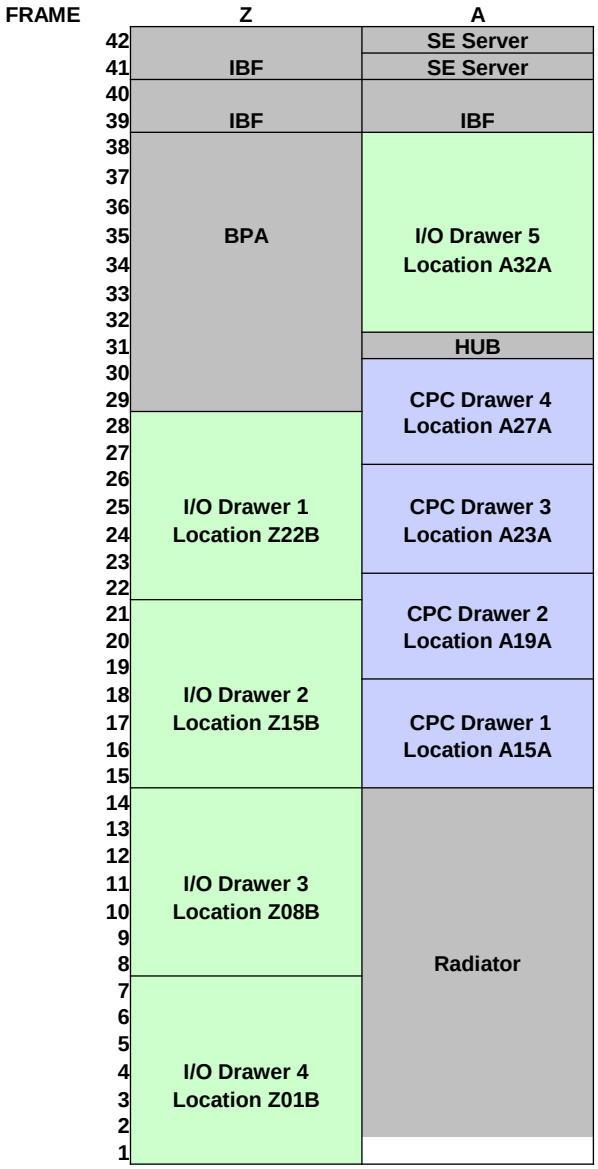

- Drawer locations are based on the front view of the machine: Frame A (right), Frame Z (left) and EIA Unit location of the lower left of drawer corner
- Locations are reported in eConfig “AO Data” reports along with PCHIDs for I/O definition
- CPC Drawers are populated from bottom to top
  - Drawer 1: A15A – N30, N63, N96, NC9 and NE1
  - Drawer 2: A19A – N63, N96, NC9 and NE1
  - Drawer 3: A23A – N96, NC9 and NE1
  - Drawer 4: A27A – NC9 and NE1
- Old technology 8-slot I/O drawers (if present) populate top down in Frame Z
  - Drawer 1: Z22B, Drawer 2: Z15B
- PCIe 32-slot I/O Drawers populate in remaining locations:
  - PCIe I/O Drawer 1: Z22B, Z15B or Z08B
  - PCIe I/O Drawer 2: Z15B, Z08B, or Z01B
  - PCIe I/O Drawer 3: Z08B, Z01B, or A32A
  - PCIe I/O Drawer 4: Z01B
  - PCIe I/O Drawer 5: A32A

z13 I/O Drawer Layout

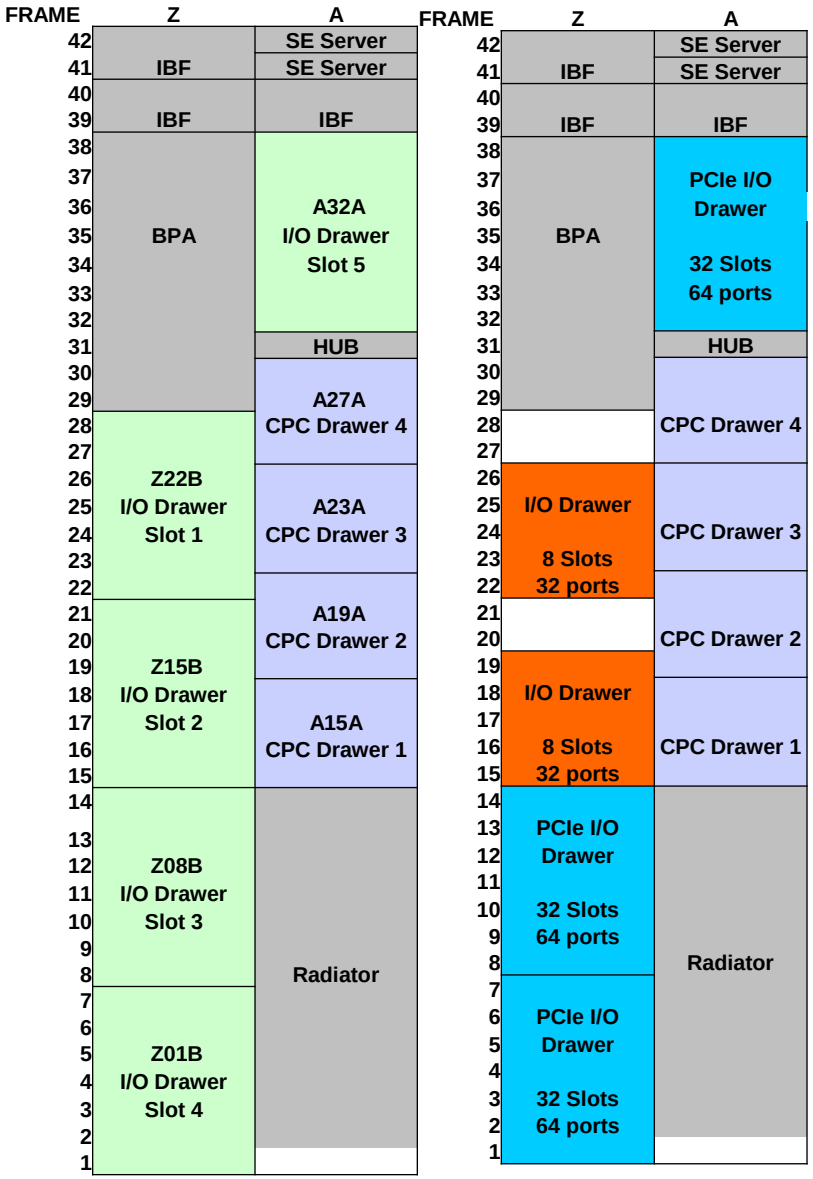

- An I/O drawer slot is a physical location in the A or Z frame for an I/O drawer or PCIe I/O drawer to be inserted = 7u
- A PCIe I/O drawer uses 1 I/O frame slot = 7u
  - 32 two port I/O slots = 64 ports each
  - 5 drawers maximum = 160 slots, 320 ports total
- An 8-slot I/O drawer uses 0.7 frame slot = 5u
  - 8 four port I/O slots = 32 ports total
  - 2 drawers carry forward **ONLY** maximum in I/O frame slots 1 and 2 only
- The 8-slot I/O drawers (if present) populate top down in the Z Frame
  - Drawer 1: Z22B, Drawer 2: Z15B
- PCIe 32-slot I/O Drawers populate in remaining locations, starting in the Z Frame:
  - PCIe I/O Drawer 1: Z22B, Z15B or Z08B
  - PCIe I/O Drawer 2: Z15B, Z08B, or Z01B
  - PCIe I/O Drawer 3: Z08B, Z01B, or A32A
  - PCIe I/O Drawer 4: Z01B
  - PCIe I/O Drawer 5: A32A

Channel Subsystems
Subchannel Sets
and Partitions

Logical channel subsystems (CSS), subchannel sets (SS), Function Definitions, and Logical Partitions on z13

- Six Logical Channel Subsystems (CSS) each with four subchannel sets (SS) and up to 256 channels
  - Maximum channel count includes channels spanned to more than one CSS
  - Total physical channels depend on I/O features configured
  - Up to 63.75k base IODEVICEs in SS 0 and 64 k alias IODEVICEs each in SS 1 to SS 3 per CSS
- FUNCTION definition support for virtualized RoCE and zEDC independent of CSS
- Up to 85 Logical Partitions: 15 each in CSS 0 – 4, 10 in CSS 5 (Partitions B – F Reserved)
  - Only channels and IODEVICEs defined in its CSS can be assigned to an LPAR
  - Any defined FUNCTION can be assigned to any LPAR

z13					
Function Definitions for up to16 RoCE features (31 LPARs each) and 8 zEDC features (15 LPARs each)					
CSS 0	CSS 1	CSS 2	CSS 3	CSS 4	CSS 5
Up to 15 Logical Partitions	**Up to 15 Logical Partitions**	**Up to 15 Logical Partitions**	**Up to 15 Logical Partitions**	**Up to 15 Logical Partitions**	**Up to 10 Logical Partitions**
Subchannel Sets: **SS 0 – 63.75 k** **SS 1 – 64 k** **SS 2 – 64 k** **SS 3 – 64 k**	**Subchannel Sets:** **SS 0 – 63.75 k** **SS 1 – 64 k** **SS 2 – 64 k** **SS 3 – 64 k**	**Subchannel Sets:** **SS 0 – 63.75 k** **SS 1 – 64 k** **SS 2 – 64 k** **SS 3 – 64 k**	**Subchannel Sets:** **SS 0 – 63.75 k** **SS 1 – 64 k** **SS 2 – 64 k** **SS 3 – 64 k**	**Subchannel Sets:** **SS 0 – 63.75 k** **SS 1 – 64 k** **SS 2 – 64 k** **SS 3 – 64 k**	**Subchannel Sets:** **SS 0 – 63.75 k** **SS 1 – 64 k** **SS 2 – 64 k** **SS 3 – 64 k**
Up to 256 Channels	Up to 256 Channels	Up to 256 Channels	Up to 256 Channels	Up to 256 Channels	Up to 256 Channels

Cryptography

Crypto Express5S (New Build)

- **Native PCIe card (FC #0890)**
  - Resides in the PCIe I/O drawer
  - Requires CPACF Enablement (FC #3863)
- **New Crypto Module**
  - Designed to more than double Crypto Express4S performance (Added L2 Cache, New Crypto ASIC and processor upgrade)
  - **Designed to support up to 85 domains for logical partitions or z/VM guests**
- **Designed to Meet Physical Security Standards**
  - FIPS 140-2 level 4
  - ANSI 9.97
  - Payment Card Industry (PCI) HSM
  - Deutsche Kreditwirtschaft (DK)

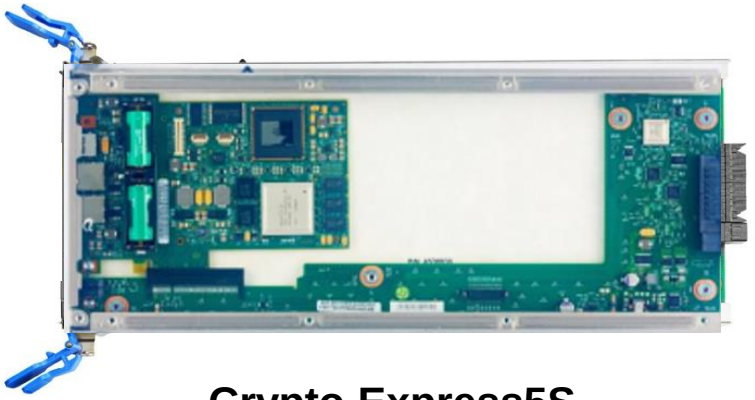

Crypto Express5S
(FC #0890)
(Placeholder picture)

- **New Functions, Standard and Compliance**
  - Drivers: NIST via FIPS standards and implementation guidance requirements; emerging banking standards: and strengthening of cryptographic standards for attack resistance
  - VISA Format Preserving Encryption (VFPE) for credit card numbers
  - Enhanced public key Elliptic Curve Cryptography (ECC) for users such a Chrome, Firefox, and Apple's iMessage
- **New Trusted Key Entry Workstation**
  - Workstation and LIC – FC #0847 with new crypto module and TKE LIC 8.0 is required
  - Required: EP11 (PKCS #11) Mode, Recommended: Common Cryptographic Architecture (CCA) Mode
  - Additional Smart Cards (FC #0892) – Support for stronger encryption than previous cards

FICON

FICON Express16S – SX and 10KM

- For FICON, zHPF, and FCP environments
  - CHPID types: FC and FCP
    - 2 PCHIDs/CHPIDs
- Auto-negotiates to 4, 8, or 16 Gbps
  - 2Gbps connectivity **NOT** supported
  - FICON Express8S will be available to order for 2Gbps connectivity
- Increased bandwidth compared to FICON Express8S
- 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
- 2 channels of LX or SX (no mix)
- Small form factor pluggable (SFP) optics
  - Concurrent repair/replace action for each SFP

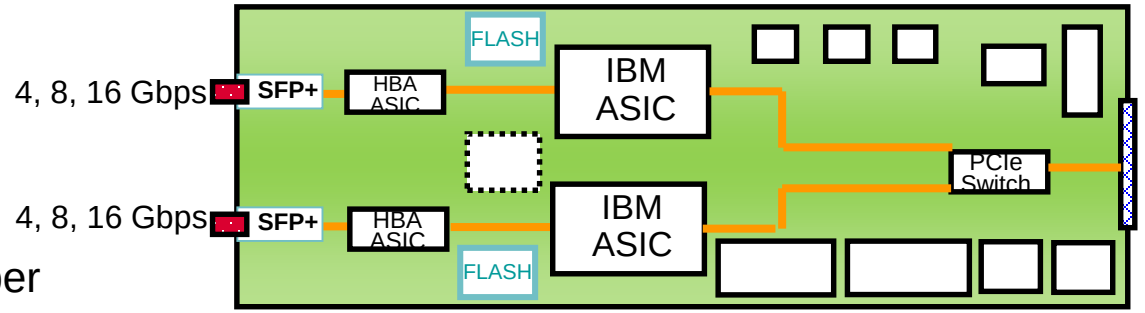

FC 0418 – 10KM LX, FC 0419 – SX

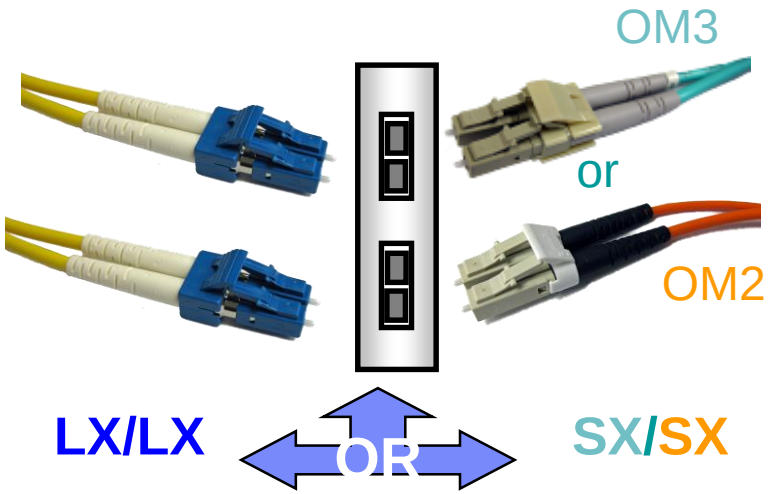

New FICON Function on z13

- FICON Express16S - 16 Gbps Link Speeds
  - Designed with the DS8870 to provide substantially improved DB2 transactional latency and up to 32% reduction in elapsed time for I/O bound batch jobs.
- 32K devices per FICON channel on all types of FICON channel
  - Up to 85 Logical Partitions: More flexibility for server consolidation
- Fourth subchannel set for each LCSS
  - Designed to eliminate single points of failure for storage after a disk failure by facilitating the exploitation of IBM DS8870 multi-target Metro Mirror storage replication with IBM Geographically Dispersed Parallel Sysplex™ (IBM GDPS®) and IBM Tivoli Storage Productivity Center for Replication HyperSwap
- Preserve Virtual WWPNs for NPIV configured FCP channels
  - Designed to simplify migration to a new-build z13
- Improved zHPF Performance at Extended Distance – GA
  - Can reduce the impact of distance on I/O response times by 50% for large data writes, providing significant response time improvements for multi-site IBM Parallel Sysplex® environments
- Forward Error Correction (FEC) on FICON Express16S – GA September 25, 2015
  - Designed to work with supporting storage capabilities of the Fibre Channel link protocol to enable operation at higher speeds, over longer distances, with reduced power and higher throughput, while retaining traditional FICON reliability and robustness
- FICON Dynamic Routing (EBR/OxID compatibility) – GA September 25, 2015
  - Designed to enable exploitation of SAN dynamic routing polices in the fabric to lower cost and improve performance for supporting I/O devices
- Mainframe SAN Fabric Priority – GA September 25, 2015
  - Mainframe SAN Fabric Priority, with exploiting storage products, extends the z/OS Work Load Manager (WLM) to the SAN infrastructure providing improved resilience and autonomic capabilities while enhancing the value of FICON Dynamic Routing

zHPF and FICON Performance* on z Systems

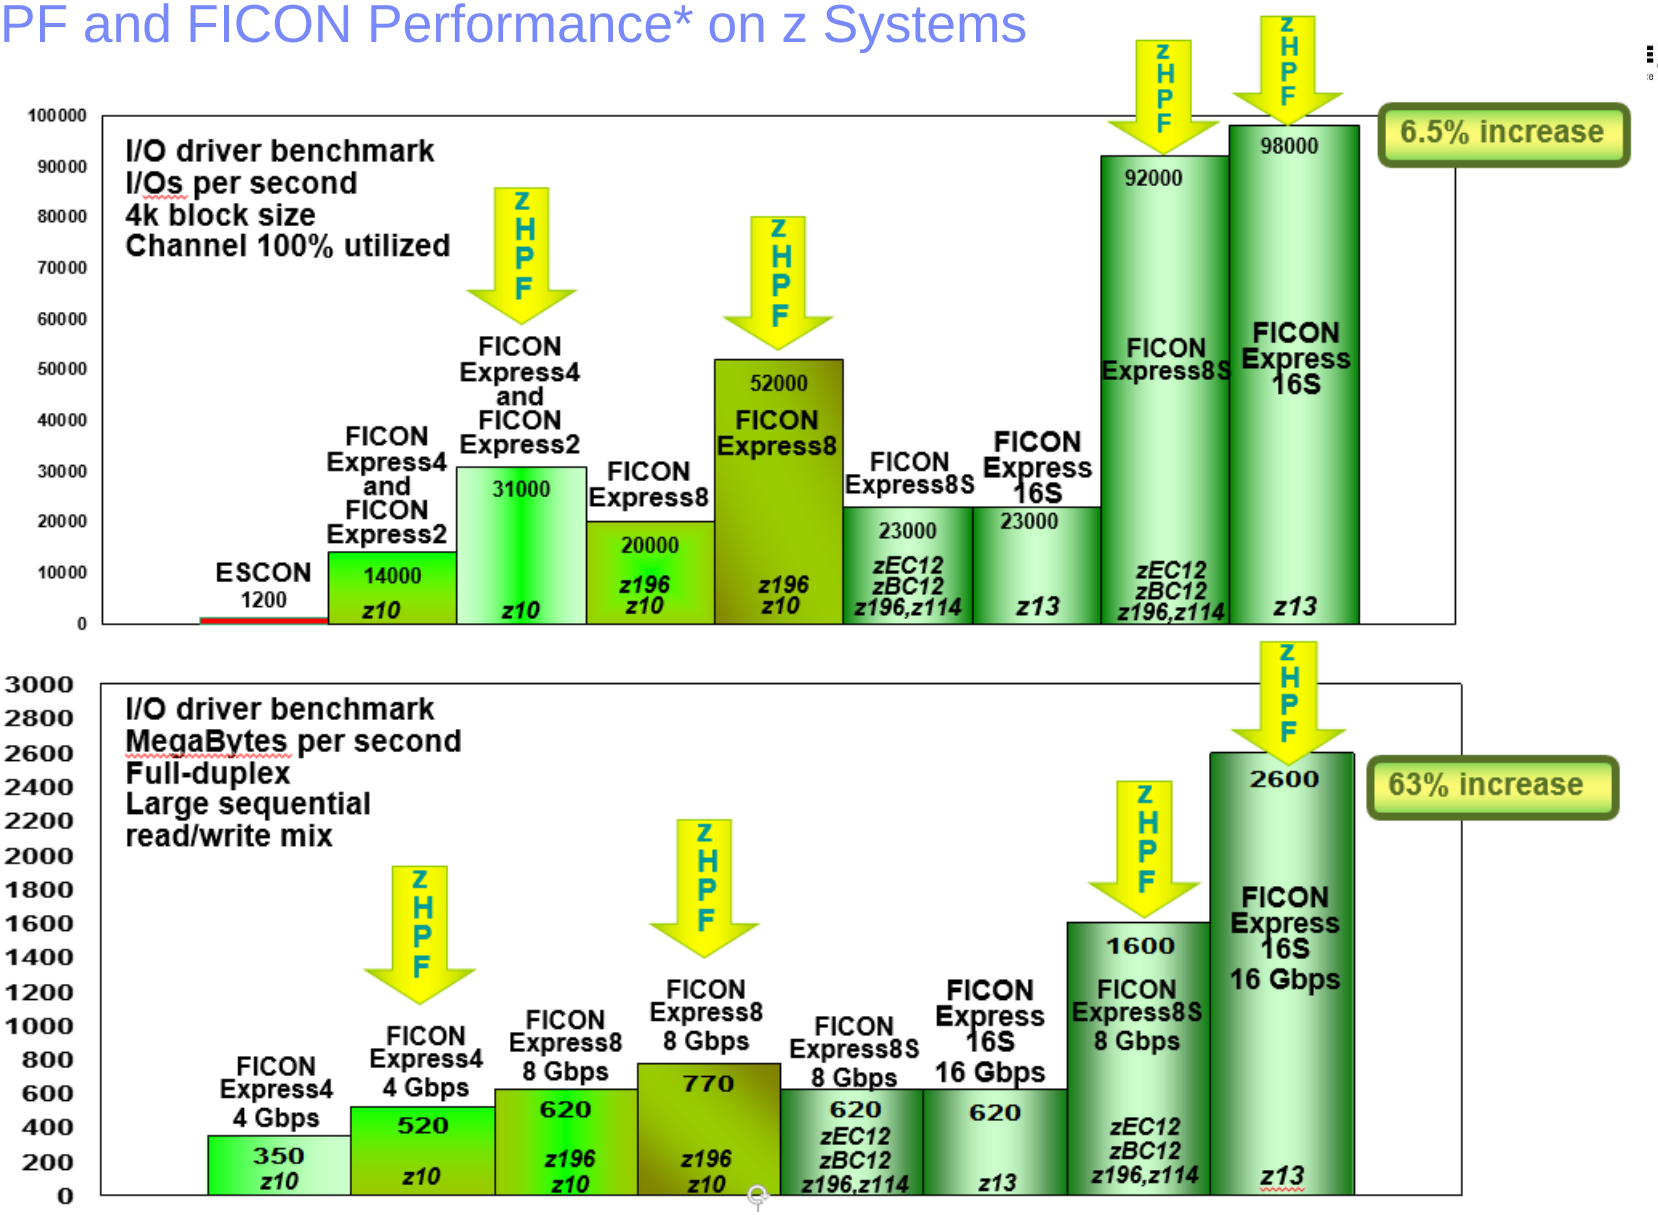

*This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration,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

FCP Performance* on z Systems

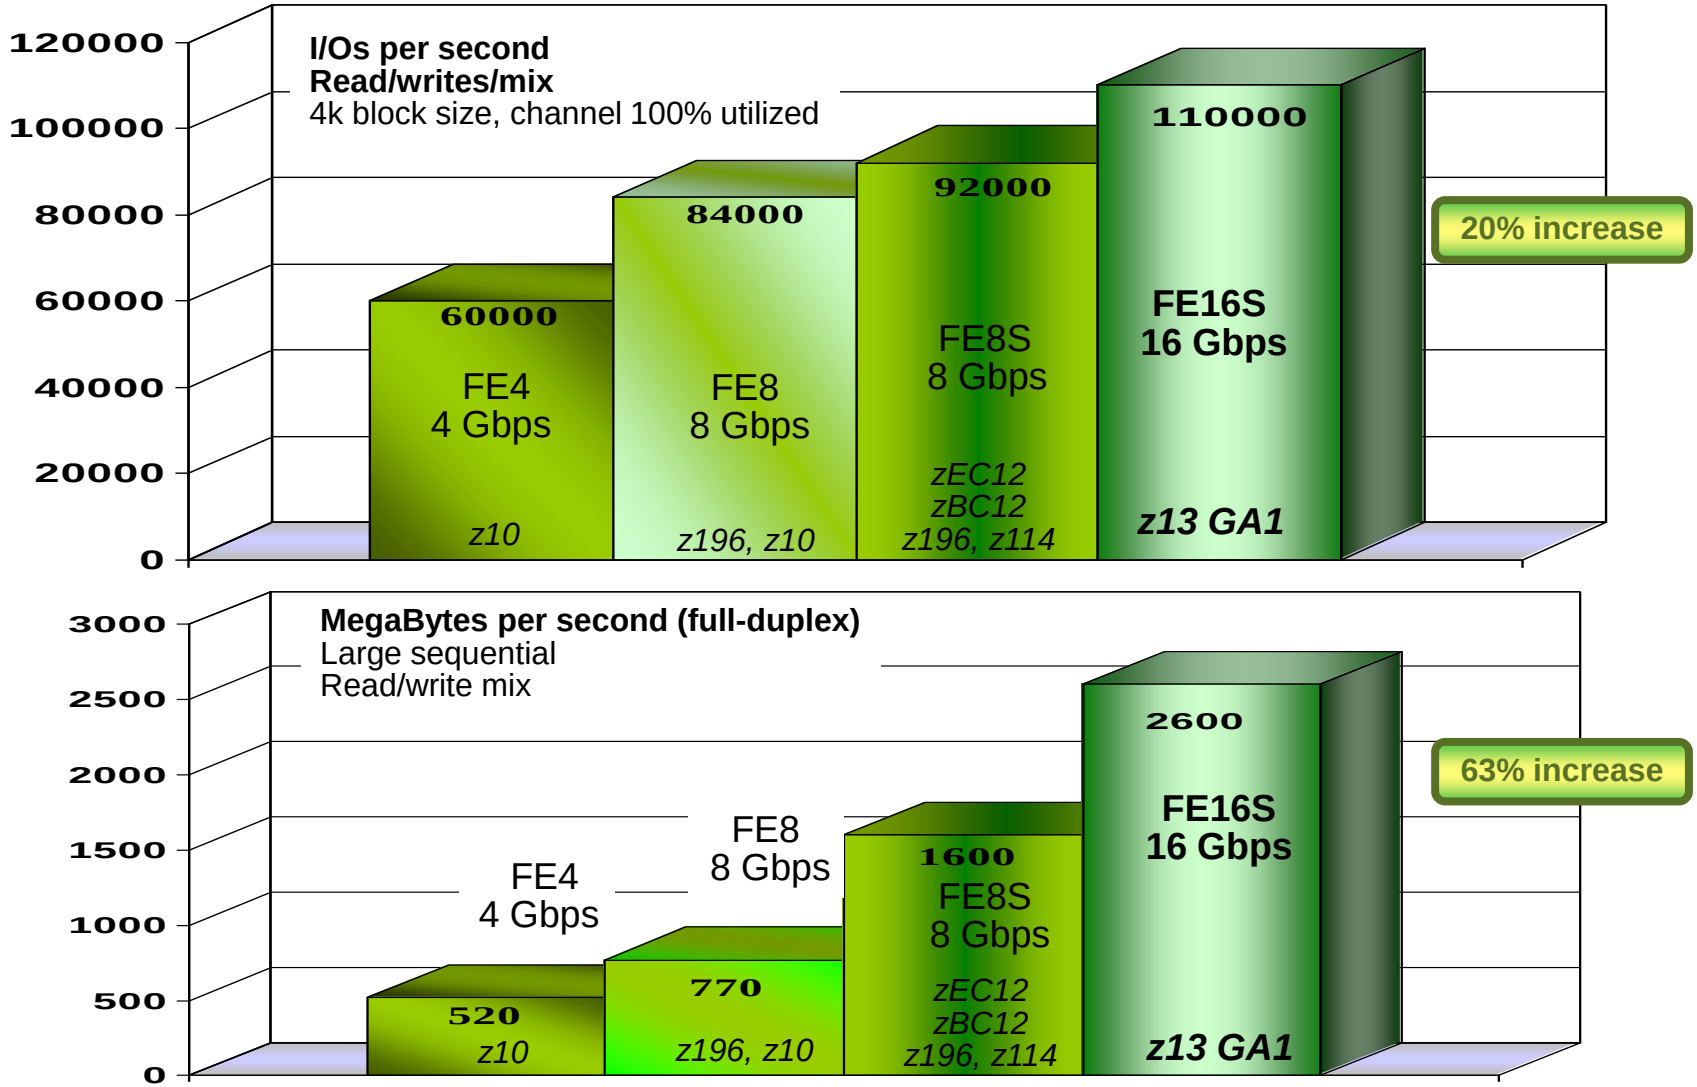

*This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration,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

Networking

New Networking Function on z13

- 10 GbE RoCE Express Virtualization Support
  - Designed to enable both ports on a RoCE Express feature and to allow sharing of each RoCE Express feature by up to 31 logical partitions
- Static VCHID Support for HiperSockets Channels
  - Designed to facilitate resource management by providing a consistent identifier for HiperSockets channels
- OSA OSD Channel Multi VSWITCH Link Aggregation (LAG) Support – June 26, 2015
  - Designed to improve z/VM V6.3 virtual networking capabilities and to permit sharing of supporting OSD channels among multiple z/VM V6.3 images

10 GbE RoCE Express

z13 - 10GbE RoCE Express Feature

- **Designed to support high performance system interconnect**
  - Shared Memory Communication (SMC) over Remote Direct Memory Access (RDMA) (SMC-R) Architecture exploits RDMA over Converged Ethernet (CE) - RoCE
  - Shares memory between peers
  - Read/write access to the same memory buffers without application changes
  - Designed to increase transaction rates greatly with low latency and reduced CPU cost
- **Configuration**
  - *z13 - Both 10 GbE SFP+ ports enabled*
  - *z13 - Support for up to 31 Logical Partitions*
  - A switched connection requires an enterprise-class 10 GbE switch with SR Optics, *Global Pause* enabled & *Priority Flow Control (PFC)* disabled
  - Point-to-point connection is supported
  - *Either connection supported to z13, zEC12 and zBC12*
  - Not defined as a CHPID and does not consume a CHPID number
  - Up to 16 features supported on a zBC12/zEC12
  - Link distance up to 300 meters over OM3 50 micron multimode fiber
- **Exploitation and Compatibility**
  - z/OS V2.1
  - IBM SDK for z/OS Java Technology Edition, Version 7.1 (February 24, 2014)
  - z/VM V6.3 support for z/OS V2.1 guest exploitation (June 27, 2014)
  - Linux on z Systems – IBM is working with Linux distribution partners to include support in future releases*

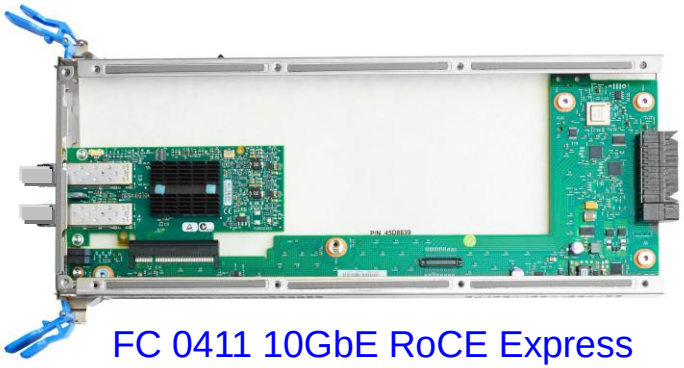

FC 0411 10GbE RoCE Express

OM3 fiber recommended

*Note: All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

z13: 10GbE RoCE Express Sample Configuration

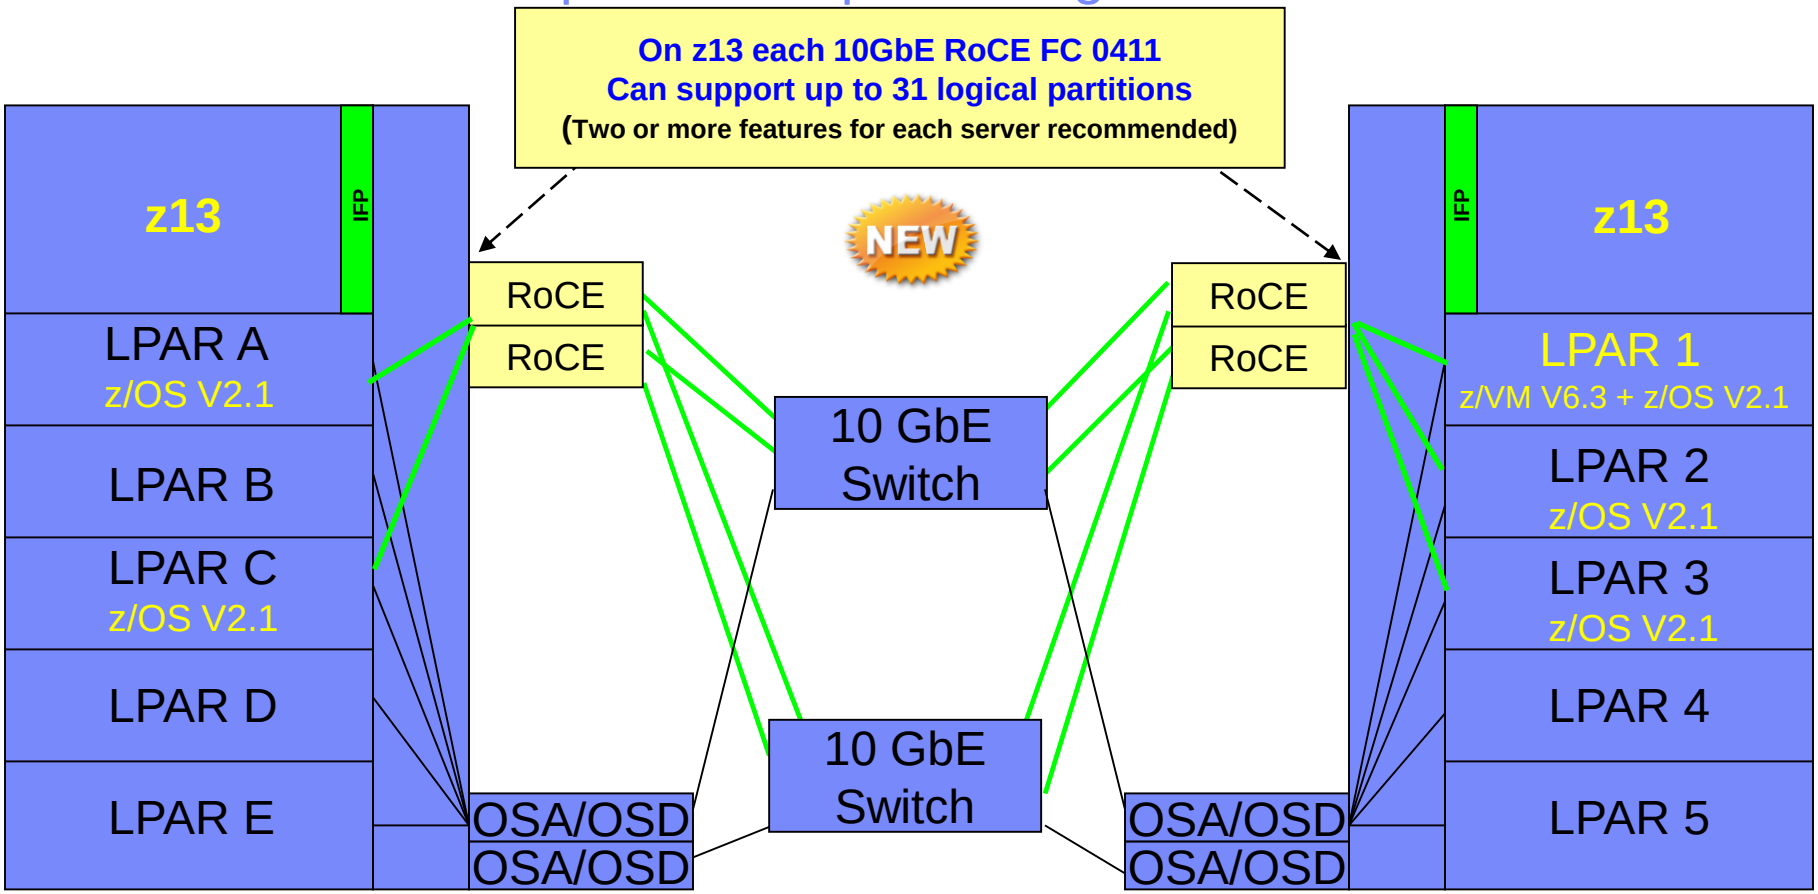

- This configuration allows redundant SMC-R connectivity among LPAR A, LPAR C, LPAR 1, LPAR 2, and LPAR 3
- Both 10 GbE
- LPAR to LPAR OSD connections are required to establish the SMC-R communications
  - 1 GbE OSD connections can be used instead of 10 GbE
  - OSD connections can flow through the same 10 GbE switches or different switches
  - z13 exclusive: Simultaneous use of both 10 GbE ports on 10 GbE RoCE Express features

z13 I/O Features, Channels, Ports, Domains, and Functions

Features	Offered As	Maximum # of features	Channels, Ports, Domains, Functions	Increments per Feature	Purchase increments
FICON – (Maximum of 160 features (320 channels) total only if all are FICON Express16S or 8S features.)					
FICON Express16S[1]	NB[2]	160	320[4] channels maximum	2 channels/feature	2 channels
FICON Express8S[3]	NB	160	320[4] channels maximum	2 channels/feature	2 channels
FICON Express8[4]	CF[2]	16	64 channels maximum	4 channels/feature	**CF Only**
Networking – (No more than 48 networking features total counting features of all types. One channel per feature)					
OSA-Express5S	NB	48	96 ports maximum	Ports: 2, 10 GbE 1	1 feature
OSA-Express4S	CF	48	96 ports maximum	Ports: 2, 10 GbE 1	**CF Only**
Crypto – (No more that 16 crypto features)					
Crypto Express5S[1]	NB	16	85 Domains/Adapter	1 PCIe Adapter	2, 3 - 16
Special purpose – These features provide Native PCI FUNCTIONs or Storage Class Memory (SCM)					
10GbE RoCE Express	NB	16	31 FUNCTIONs/Adapter	2 ports/Adapter	1 feature
Flash Express[1] (FC#0403)	NB	8 (4 Pairs)	1.4 TB SCM per pair	1 PCIe Adapter	2 (1 Pair)
Flash Express (FC#0402)	CF	8 (4 Pairs)	1.4 TB SCM per pair	1 PCIe Adapter	**CF Pairs Only**
zEDC Express	NB	8	15 FUNCTIONs/Adapter	1 PCIe Adapter	1 feature

Notes:
1. **Bold blue text indicates new features for z13**
2. **NB** = New Build, and if previously offered Carry Forward, **CF-** Carry Forward **ONLY**
3. FICON Express8S is offered on New Build to support point to point 2 Gbps attachment
4. Any 8-slot drawer limits maximum memory in any LPAR to 1 TB; One 8-slot drawer limits maximum FICON channels to 288, two 8-slot drawers limit maximum FICON channels to 256. (These numbers are REDUCED by 4 for each empty slot in an 8-slot drawer.)

Parallel Sysplex and Server Time Protocol

z13 Parallel Sysplex Enhancements (Introduced with CFCC 20)

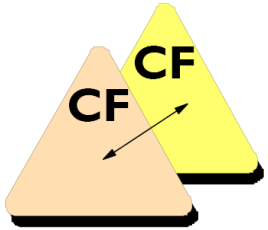

- **Support for up to 141 ICF processors**
  - The maximum number of logical processors in a CF LPAR remains at 16

- **Coupling Links Support**
  - PCIe-O SR 8 GBps 150 m
    - Up to 16 features (Up to 10 per drawer) = 32 ports
  - HCA3-O LR 1x 5 Gbps long distance links
    - Up to 16 features (4 per drawer) = 64 ports
  - HCA3-O 12x 150 m
    - Up to 16 features (Up to 4 per drawer) = 32 ports
  - Internal Coupling (Up to 32 ICP CHPIDs, 16 ICP-ICP links)
  - Coupling CHPID definitions
    - Up to 256 (Increased from 128) – June 26, 2015
    - The maximum defined to one CF partition remains at 128

- **PCIe-O SR 8 GBps 150 m links (2 ports per feature)**
  - Up to 4 Coupling CHPID TYPE=CS5 definitions per port, 8 per feature
  - Cable/point to point maximum distance options:
    - 150 Meters with 12-pair OM4 50/125 micron fiber (Recommended)
    - 100 Meters with 12-pair OM3 50/125 micron fiber
      (Note: InfiniBand 12x links use 12-pair OM3 cabling with **different** connectors)
  - Estimated Performance Approximately Equivalent to InfiniBand 12x

- **Improved Scalability and Support for Large CF Structures**

z13 Parallel Sysplex Coupling Connectivity

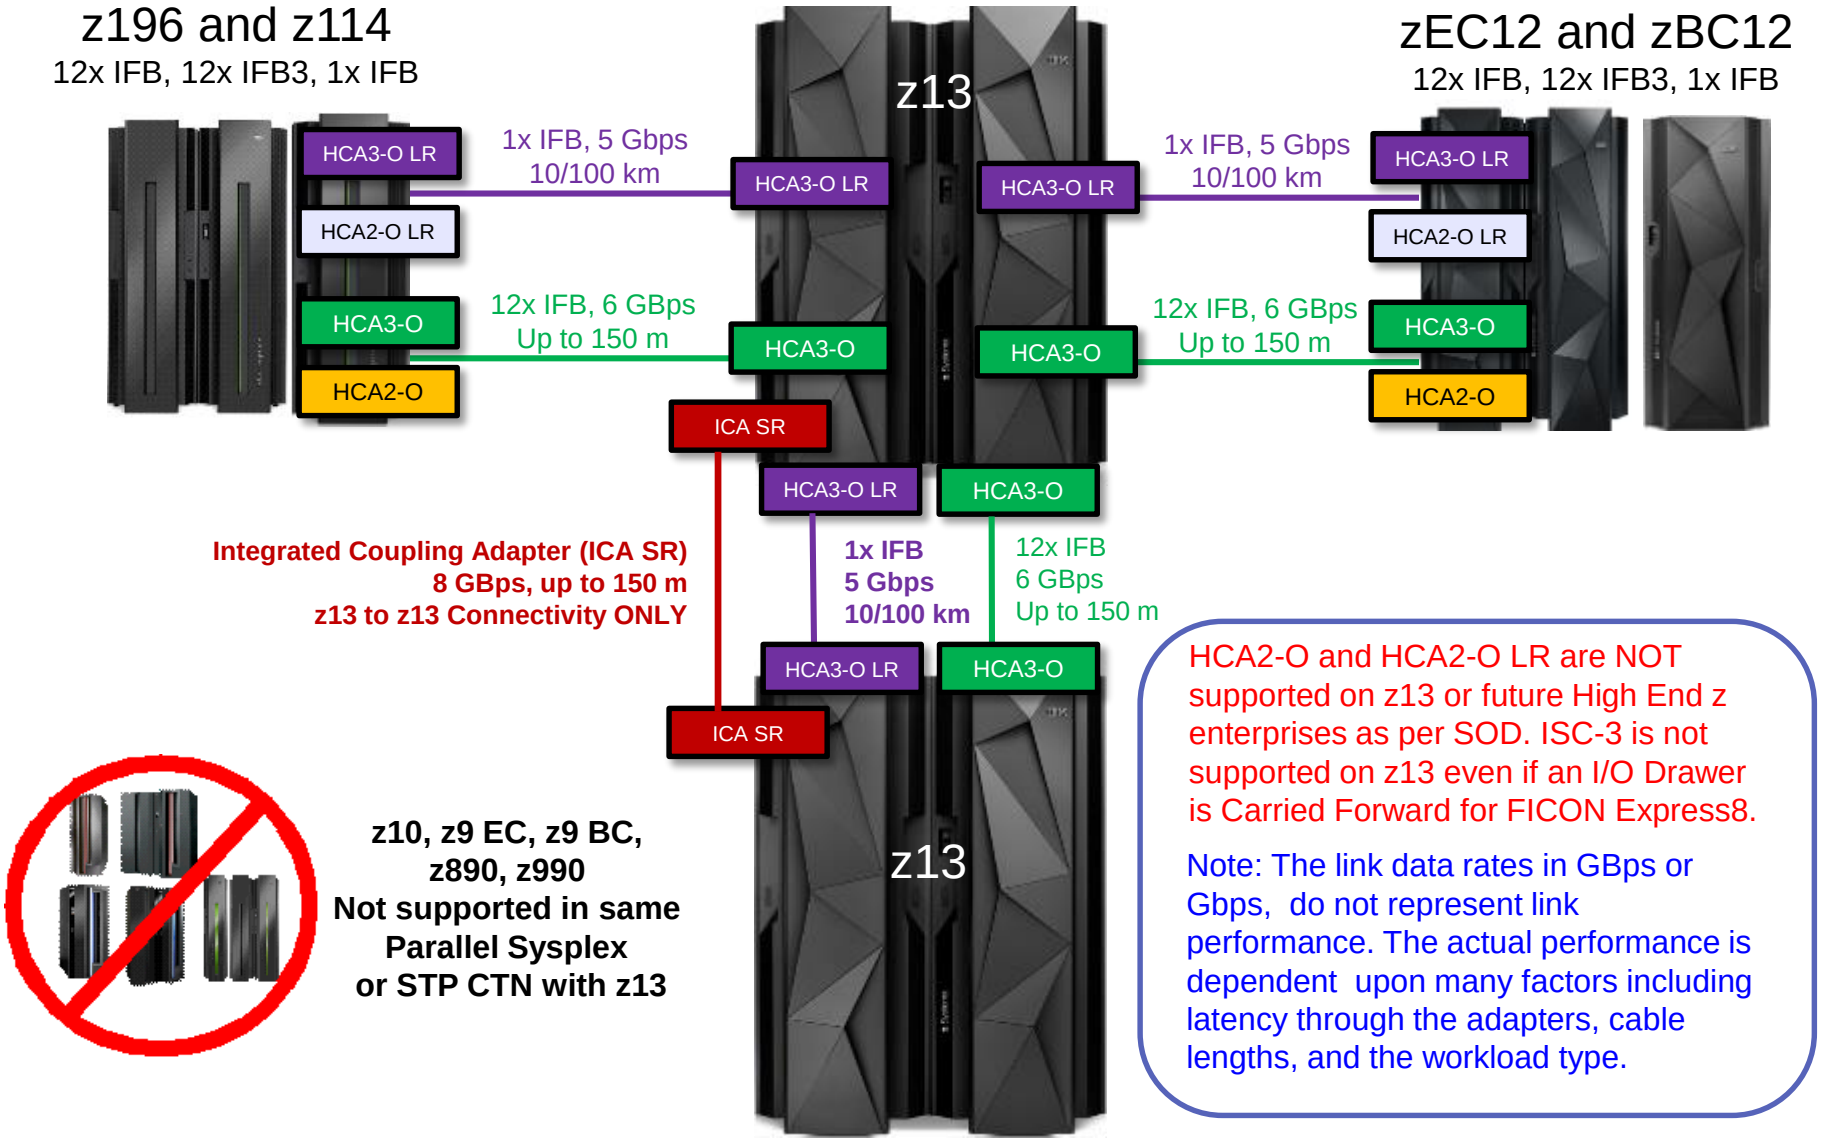

HCA2-O and HCA2-O LR are NOT supported on z13 or future High End z enterprises as per SOD. ISC-3 is not supported on z13 even if an I/O Drawer is Carried Forward for FICON Express8.

Note: The link data rates in GBps or Gbps, do not represent link performance. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the workload type.

z13 Coupling Link Details at GA

Features	Offered as	Maximum # of features	Maximum connections	Increments per feature	Purchase increments
PCIe-O ICA SR (GA1)	NB	16	32 links[1]	2 links	2 links
HCA3-O LR (1x)	NB/CF	16	64 links[2]	4 links	4 links
HCA3-O (12x)	NB/CF	16	32 links	2 links	2 links
ICP (Standard)	NB/CF	NA	32 ICP CHPIDs, 16 ICP-ICP Links		

Notes: [1] Same physical number of links as 12X PSIFB on zEC12
[2] Same physical number of links as 1X PSIFB on zEC12

NB = New build, Migration Offering, z Systems Exchange Program
CF = Carry Forward

Link type	GA1	Port Qty	Protocol	Link data rate	Fiber core	Fiber bandwidth	Fiber type	Light source	Cable	Connector	Maximum distance	Repeated Distance
Short distance												
HCA3-O fanout (12x IFB)	X	2	IFB	6 GBps	50 micron	2000 MHz-km 850 nm	OM3 Multimode	SX	24-fiber cable assembly	MTP (split) Tx & RX	150 meters	N/A
PCIe-O SR for Coupling (Fanout in CPC drawer)	X	2	PCIe Gen3	8 GBps	50 micron	4700 MHz-km 850 nm	**OM4 Multimode**	SX	24-fiber cable assembly	MTP (new)	**150 meters**	N/A
					50 micron	2000 MHz-km 850 nm	**OM3 Multimode**	SX	24-fiber cable assembly	MTP (new)	**100 meters**	N/A
Long distance												
HCA3-O LR fanout (1x IFB)	X	4	IFB	5 Gbps	9 micron	1310 nm	Single mode	LX	1 fiber pair	LC Duplex	10 km 20 km RPQ	100 km

Installation Planning for z13

z13 Hardware Management Console

- **HMC System Unit and LIC Support**
  - New Build: HMC FC 0092 desk side or 0094 rack mounted HMC (0 – 10 orderable per z13)
  - Carry Forward: HMCs FC 0091 or FC 0092 will be upgraded to control z13
  - zEnterprise Ensemble Primary and Alternate HMCs required to support z13
    - An identical pair is required (Two FC 0094, two FC 0092 or two FC 0091)
    - At Driver 22 – HMC LIC Application level 2.13.0
  - **No-charge ECAs** orderable by IBM service are available for **older** z Systems servers to upgrade their HMC FC 0092 and FC0091 HMCs to HMC Driver 22 LIC
    - FC 0091: ECA 348 (For FC 0091 with 8 GB memory, ECA 332 is also required to upgrade to 16 GB)
    - FC 0092: ECA 348
  - HMC application in Driver 22 supports z9 (N-4) and later only
- **HMC Display Support for HMC FC 0092**
  - 22 inch flat panel FC 6096 (No change from zEC12)
- ***New Backup Options***
  - Critical z13 HMC data: USB Storage and FTP/Secure FTP
  - Critical z13 SE data: SE to Alternate SE Hard Drive or and FTP/Secure FTP Server
    - Older machine HMC and SE USB storage only. New optional 32 GB USB "stick" offered if needed
- **HMC 1000BASE-T LAN Switches – No longer offered**
  - FC 0070 10/100/1000BASE-T switches – (Carry Forward Only)
  - Recommended Alternative: Compatible customer provided 1000BASE-T switches
- **See the z13 Library on Resource Link for the latest publications**
  - "Installation Manual for Physical Planning" for HMC FC 0091, 0092 and 0094 feature physical characteristics
  - "Integrating the HMC Broadband RSF into your Enterprise"
  - "Hardware Management Console Operations Guide" and "Support Element Operations Guide"

z13: HMC Feature Code #0094, Display and Keyboard

HMC 1U System Unit:

HMC Display and Keyboard: IBM 1U 18.5-inch Standard Console

Note: The System unit and tray must be mounted in a customer rack in two adjacent 1U locations in the "ergonomic zone" between 21U and 26U. Three C13 power receptacles are required, two for the System Unit and one for the Display and Keyboard.

z13 Physical Planning

- **Extend / Maintain zEC12 Datacenter Characteristics**
  - 2 frame base system (CEC, I/O, service system and PP&C)
  - No significant increase in weight
  - Maintain floor tile cutouts for raised floor system (same as z10 EC, z196, and zEC12)

- **Better control of energy usage and improved efficiency in your data center**
  - **Support for ASHRAE Class A2 datacenter (Up to 35º C and 80% relative humidity)**
  - **Upgraded radiator (air) cooling compared to zEC12 with N+2 pumps and blowers**
  - **Upgraded water cooling compared to zEC12 support for 24º C water (was 20º C, 15º on z196)**
  - Same number of power cords (2 or 4) as “equivalent” zEC12 configuration
  - Maintain 27.5 kW box max input power (same as z10 EC, z196, and zEC12)
  - Maintain DC input power capability, overhead I/O cabling option, and overhead power options

z13 New Installation Planning Consideration

- Rear Cover Adjustable Airflow
  - The IBM z13 has a new rear door design that includes reversible rear door panels that can be installed two different ways to allow exhaust airflow to be directed upward or downward.
  - This design addresses issues experienced by a few datacenters due to fixed downward exhaust airflow on older z Systems servers.
  - Action: Advise IBM prior to the install of the desired airflow direction.

Airflow UP

Airflow Down

- Locking Doors
  - In response to client requirements, IBM z13 has doors that include standard key locks compliant with industry standards. There are four locks, each provided with two keys. Locking the doors or leaving them unlocked is a client option.
  - Action: Advise IBM of whether or not the doors are to be locked. It is a client responsibility to maintain custody of the keys and, if the doors are to be locked to establish key control procedures, to ensure that the doors are unlocked promptly whenever required (24x7) for IBM service, and to ensure they are locked again after service is complete.

z13 Installation - Raised Floor options

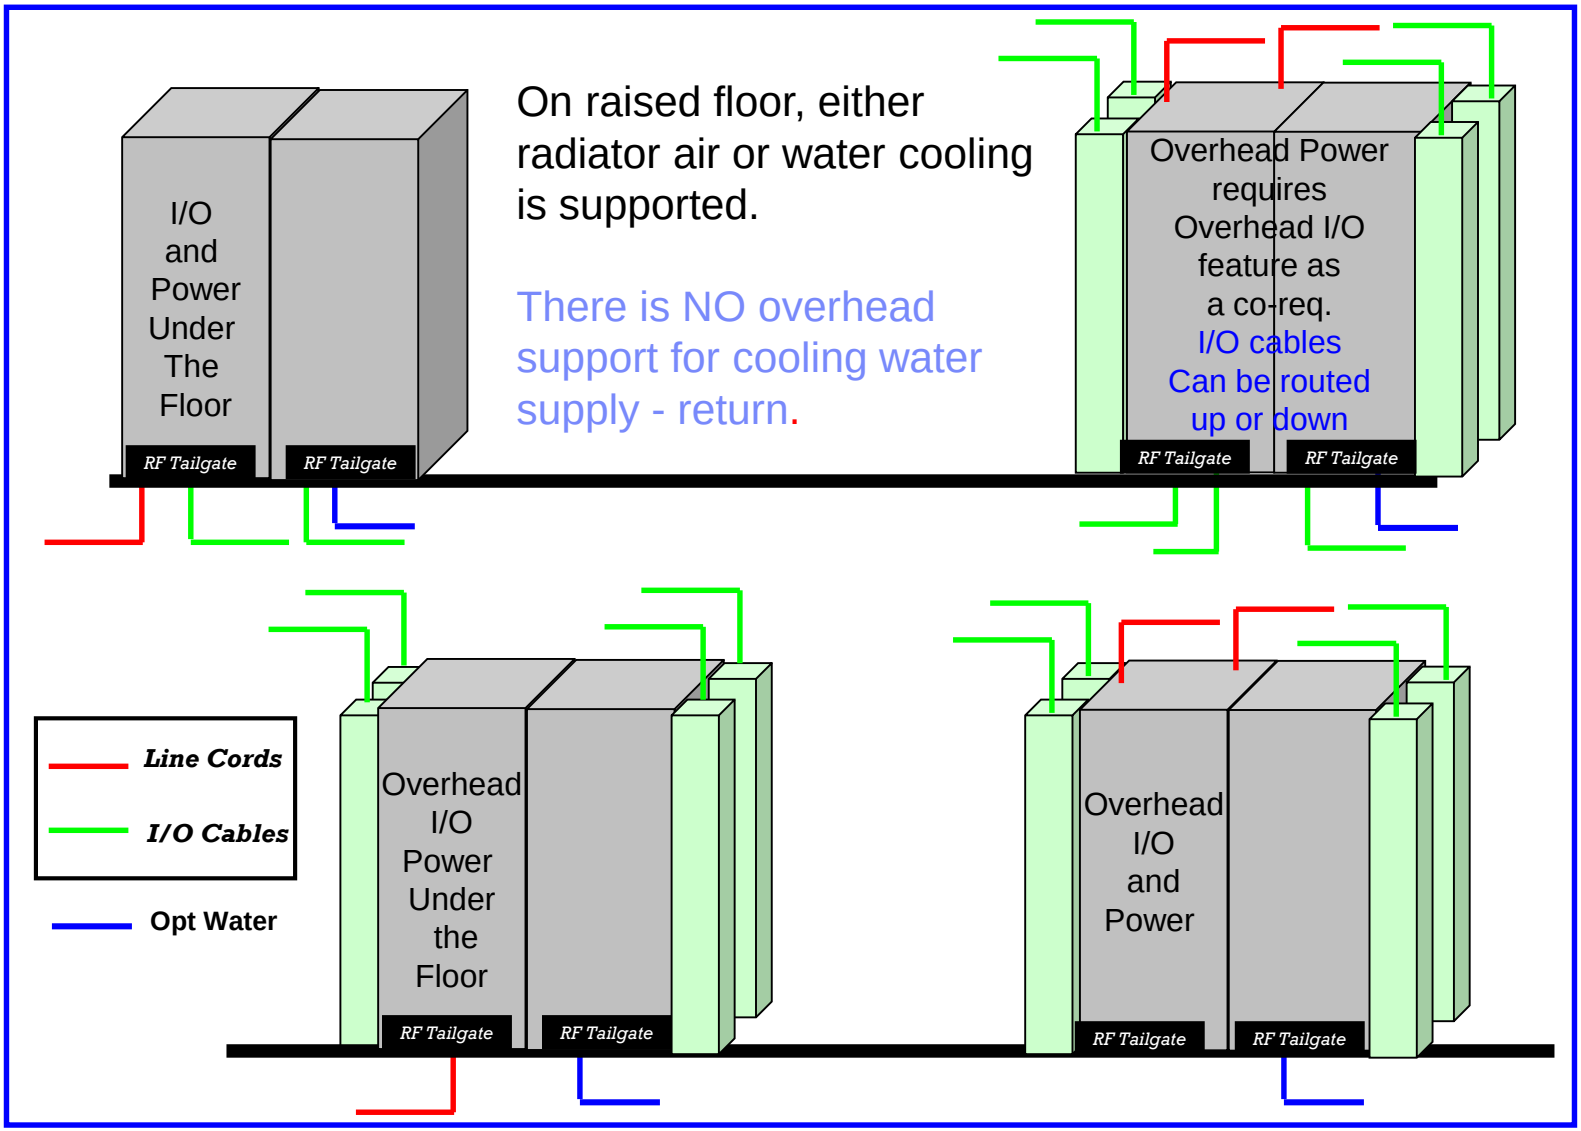

Top Exit Power option: When selected for a raised floor the Top Exit I/O feature is a coreq. Also the diagram for this configuration should depict the I/O routing up thru the I/O chimneys and also routing thru the bottom of the frame using the raised floor tailgates.

z13 Installation – Non-Raised Floor option

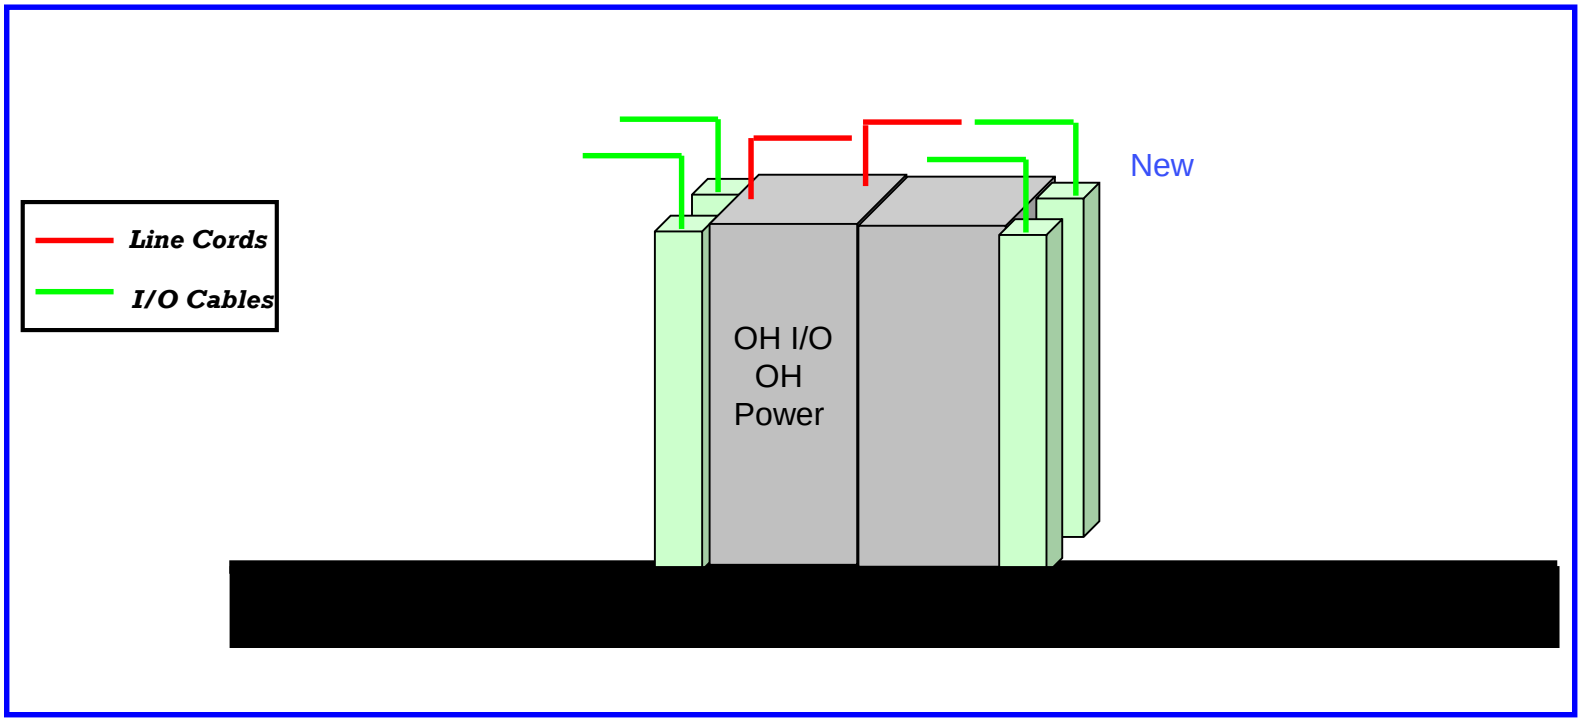

If z13 is NOT installed on a raised floor, overhead I/O, overhead power, and radiator (air) cooling options are required.

Water cooling is NOT supported. NO cables may exit at floor level.

z13 New Fill and Drain Tool (FDT) and Lift/Tool Ladder

Approximate FDT unit dimensions:

- 35 inches from floor to top of handle
- 30 inches long
- 22 inches wide

New FDT: FC #3380
Or order upgrade kit FC #3379 if a zEC12 FDT FC # 3378 will remain on site

New Universal Lift Tool/Ladder: FC #3105
Or order upgrade kit FC #3103 if a zEC12 Universal Lift Tool/Ladder FC #3359 will remain on site

System Fill Procedure

- Driven through Repair & Verify on SE
- 15-20 minute procedure
- Initial setup includes:
  - Starting R&V
  - Gathering FDT, adapter kit, and BTA water solution
  - Plugging FDT into bulk power port on system

Statements of Direction
I/O and HMC

Statements of Direction: I/O

- The IBM z13 will be the last z Systems server to support FICON Express8 channels: IBM z13 will be the last high-end server to support FICON Express8. Enterprises should begin migrating from FICON Express8 channel features (#3325, #3326) to FICON Express16S channel features (#0418, #0419). FICON Express8 will not be supported on future high-end z Systems servers as carry forward on an upgrade.

- The IBM z13 will be the last z Systems server to offer ordering of FICON Express8S channel features. Enterprises that have 2 Gb device connectivity requirements must carry forward these channels.

- The IBM z13 will be the last generation of z Systems hardware servers to support configuring OSN CHPID types. OSN CHPIDs are used to communicate between an operating system instance running in one logical partition and the IBM Communication Controller for Linux on z Systems (CCL) product in another logical partition on the same CPC. See announcement letter #914-227 dated 12/02/2014 for details regarding withdrawal from marketing for the CCL product.

- IBM intends to provide support for the Read Diagnostic Parameters Extended Link Service command for fiber channel SANs as defined in the T11.org FC-LS-3 draft standard. Support for the Read Diagnostic Parameters Extended Link Service command is intended to improve SAN reliability and fault isolation.

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Statements of Direction: HMC

- Removal of support for Classic Style User Interface on the Hardware Management Console and Support Element: The IBM z13 will be the last z Systems server to support Classic Style User Interface. In the future, user interface enhancements will be focused on the Tree Style User Interface.

- Removal of support for the Hardware Management Console Common Infrastructure Model (CIM) Management Interface: IBM z13 will be the last z Systems server to support the Hardware Console Common Infrastructure module (CIM) Management Interface. The Hardware Management Console Simple Network Management Protocol (SNMP), and Web Services Application Programming Interfaces (APIs) will continue to be supported. IBM intends to provide support for the Read Diagnostic Parameters Extended Link Service command for fiber channel SANs as defined in the T11.org FC-LS-3 draft standard. Support for the Read Diagnostic Parameters Extended Link Service command is intended to improve SAN reliability and fault isolation.

THANK YOU

Complete Statements of Direction and Backup

Statements of Direction

- IBM plans to accept for review certification requests from cryptography providers by the end of 2015, and intends to support the use of cryptography algorithms and equipment from providers meeting IBM's certification requirements in conjunction with z/OS and z Systems processors in specific countries. This is expected to make it easier for customers to meet the cryptography requirements of local governments.
- KVM offering for IBM z Systems: In addition to the continued investment in z/VM, IBM intends to support a Kernel-based Virtual Machine (KVM) offering for z Systems that will host Linux on z Systems guest virtual machines. The KVM offering will be software that can be installed on z Systems processors like an operating system and can co-exist with z/VM virtualization environments, z/OS, Linux on z Systems, z/VSE and z/TPF. The KVM offering will be optimized for z Systems architecture and will provide standard Linux and KVM interfaces for operational control of the environment, as well as providing the required technical enablement for OpenStack for virtualization management, allowing enterprises to easily integrate Linux servers into their existing infrastructure and cloud offerings.
- In the first half of 2015, IBM intends to deliver a GDPS/Peer to Peer Remote Copy (GDPS/PPRC) multiplatform resiliency capability for customers who do not run the z/OS operating system in their environment. This solution is intended to provide IBM z Systems customers who run z/VM and their associated guests, for instance, Linux on z Systems, with similar high availability and disaster recovery benefits to those who run on z/OS. This solution will be applicable for any IBM z Systems announced after and including the zBC12 and zEC12

Statements of Direction

- Enhanced RACF password encryption algorithm for z/VM: In a future deliverable an enhanced RACF/VM password encryption algorithm is planned. This support will be designed to provide improved cryptographic strength using AES-based encryption in RACF/VM password algorithm processing. This planned design is intended to provide better protection for encrypted RACF password data in the event that a copy of RACF database becomes inadvertently accessible.

- IBM intends that a future release of IBM CICS Transaction Server for z/OS will support 64-bit SDK for z/OS, Java Technology Edition, Version 8 (Java 8). This support will enable the use of new facilities delivered by IBM z13 which are exploited by Java 8, including Single Instruction Multiple Data (SIMD) instructions for vector operations and simultaneous multithreading (SMT).

- z/VM support for Single Instruction Multiple Data (SIMD): In a future deliverable IBM intends to deliver support to enable z/VM guests to exploit the Vector Facility for z/Architecture (SIMD).

- Removal of support for Expanded Storage (XSTORE): z/VM V6.3 is the last z/VM release that will support Expanded Storage (XSTORE) for either host or guest usage. The IBM z13 server family will be the last z Systems server to support Expanded Storage (XSTORE).

Statements of Direction

- The IBM z13 will be the last z Systems server to support running an operating system in ESA/390 architecture mode; all future systems will only support operating systems running in z/Architecture mode. This applies to operating systems running native on PR/SM as well as operating systems running as second level guests. IBM operating systems that run in ESA/390 mode are either no longer in service or only currently available with extended service contracts, and they will not be usable on systems beyond IBM z13. However, all 24-bit and 31-bit problem-state application programs originally written to run on the ESA/390 architecture will be unaffected by this change.

- Stabilization of z/VM V6.2 support: The IBM z13 server family is planned to be the last z Systems server supported by z/VM V6.2 and the last z systems server that will be supported where z/VM V6.2 is running as a guest (second level). This is in conjunction with the statement of direction that the IBM z13 server family will be the last to support ESA/390 architecture mode, which z/VM V6.2 requires. z/VM V6.2 will continue to be supported until December 31, 2016, as announced in announcement letter # 914-012.

- Product Delivery of z/VM on DVD/Electronic only: z/VM V6.3 will be the last release of z/VM that will be available on tape. Subsequent releases will be available on DVD or electronically.

- Removal of support for Classic Style User Interface on the Hardware Management Console and Support Element: The IBM z13 will be the last z Systems server to support Classic Style User Interface. In the future, user interface enhancements will be focused on the Tree Style User Interface.

Statements of Direction

- Removal of support for the Hardware Management Console Common Infrastructure Model (CIM) Management Interface: IBM z13 will be the last z Systems server to support the Hardware Console Common Infrastructure module (CIM) Management Interface. The Hardware Management Console Simple Network Management Protocol (SNMP), and Web Services Application Programming Interfaces (APIs) will continue to be supported.

- The IBM z13 will be the last z Systems server to support FICON Express8 channels: IBM z13 will be the last high-end server to support FICON Express8. Enterprises should begin migrating from FICON Express8 channel features (#3325, #3326) to FICON Express16S channel features (#0418, #0419). FICON Express8 will not be supported on future high-end z Systems servers as carry forward on an upgrade.

- The IBM z13 will be the last z Systems server to offer ordering of FICON Express8S channel features. Enterprises that have 2 Gb device connectivity requirements must carry forward these channels.

- Removal of an option for the way shared logical processors are managed under PR/SM LPAR: The IBM z13 will be the last high-end server to support selection of the option to "Do not end the timeslice if a partition enters a wait state" when the option to set a processor run time value has been previously selected in the CPC RESET profile. The CPC RESET profile applies to all shared logical partitions on the machine, and is not selectable by logical partition.

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Statements of Direction

- IBM intends to provide support for the Read Diagnostic Parameters Extended Link Service command for fiber channel SANs as defined in the T11.org FC-LS-3 draft standard. Support for the Read Diagnostic Parameters Extended Link Service command is intended to improve SAN reliability and fault isolation.

- The IBM z13 will be the last generation of z Systems hardware servers to support configuring OSN CHPID types. OSN CHPIDs are used to communicate between an operating system instance running in one logical partition and the IBM Communication Controller for Linux on z Systems (CCL) product in another logical partition on the same CPC. See announcement letter #914-227 dated 12/02/2014 for details regarding withdrawal from marketing for the CCL product.

IBM z13 and zBX Model 004

IBM z13 (2964)

- Available – March 9, 2015
- 5 models – NE1, NC9, N96, N63, N30
  - Up to 141 customer configurable engines
- Sub-capacity Offerings for up to 30 CPs
- PU (Engine) Characterization
  - CP, IFL, ICF, zIIP, SAP, IFP (No zAAPs)
- SIMD instructions, SMT for IFL and zIIP
- On Demand Capabilities
  - CoD: CIU, CBU, On/Off CoD, CPE
- Memory – up to 10 TB
  - Up to 10 TB per LPAR (if no FICON Express8)
  - 96 GB Fixed HSA
- Channels
  - PCIe Gen3 16 GBps channel buses
  - Six CSSs, up to 85 LPARs
  - 4 Subchannel Sets per CSS
  - FICON Express16S or 8S (8 Carry forward)
  - OSA Express5S (4S carry forward)
  - HiperSockets – up to 32
  - Flash Express
  - zEnterprise Data Compression
  - RDMA over CE (RoCE) with SR-IOV Support
- Crypto Express5S
- Parallel Sysplex clustering, PCIe Coupling, Internal Coupling and InfiniBand Coupling
- IBM zAware: z/OS and Linux on z Systems
- Operating Systems
  - z/OS, z/VM, z/VSE, z/TPF, Linux on z Systems

IBM zBX Model 4 (2458-004)

- Available – March 9, 2015
- Upgrade ONLY stand alone Ensemble node converted from an installed zBX Model 2 or 3
- Doesn't require a 'owning' CPC
- Management – Unified Resource Manager
- zBX Racks (up to 4) with:
  - Dual 1U Support Elements, Dual INMN and IEDN TOR switches in the 1st rack
  - HMC LAN attached (no CPC BPH attachment)
  - 2 or 4 PDUs per rack
- Up to 8 BladeCenter H Chassis
  - Space for 14 blades each
  - 10 GbE and 8 Gbps FC connectivity
  - Advanced Management Modules
  - Redundant connectivity, power, and cooling
- Up to 112 single wide IBM blades
  - IBM BladeCenter PS701 Express
  - IBM BladeCenter HX5 7873
  - IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (M/T 2462-4BX)
  - IBM WebSphere DataPower® Integration Appliance XI50z with Firmware 7.0
- Operating Systems
  - AIX 5.3 and higher
  - Linux on System x
  - Microsoft Windows Server on System x
- Hypervisors
  - KVM Hypervisor on System x
  - PowerVM Enterprise Edition

"Native" PCIe Technology

PCIe I/O Features – "Native" (AKA "Direct Attach") PCIe

Flash Express, zEDC Express and 10GbE RoCE Express

- **Traditional z Systems I/O PCIe Feature**
  - One z Systems ASIC per Channel/PCHID
  - Definition and LPAR Assignment
    - HCD/IOCP CHPID definition **or**
    - Firmware definition outside HCD/IOCP is possible for some. For example: ***Crypto Express5S*** *is not defined as a CHPID*
  - Virtualization and support by Channel Subsystem LIC on System Assist Processors (SAPs)

NEW

- **Native PCIe Features**
  - z Systems ASIC role moved to the new z Systems I/O Controller (zIOC) in the PCIe I/O fanout or the processor
  - Definition and LPAR Assignment
    - HCD/IOCP FUNCTION definition similar to CHPID definition but with different rules **or**
    - Firmware definition outside HCD/IOCP is possible for some. For example: ***Flash Express*** *is not defined with FUNCTIONs*
  - Virtualization and support by the zIOC and Redundancy Group LIC running on the Integrated Firmware Processor (IFP) (***Note: NOT applicable to Flash Express***)

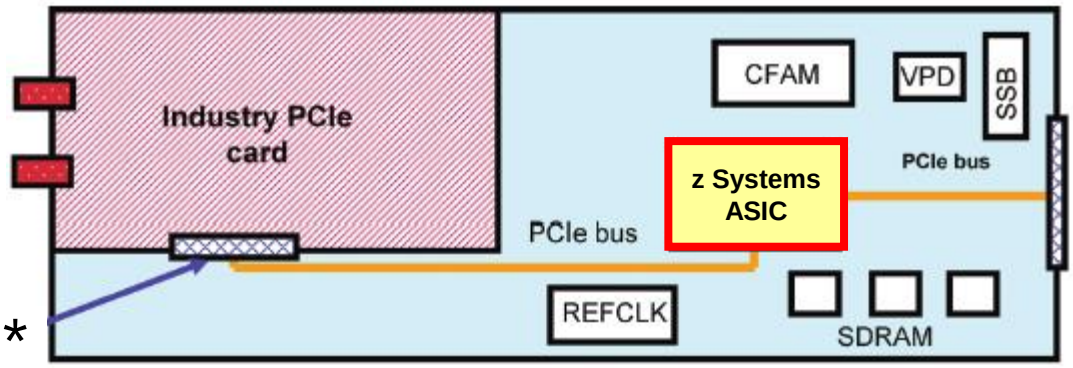

Traditional z Systems I/O PCIe Features: FICON Express16S and 8S, OSA-Express5S and 4S, *Crypto Express5S*

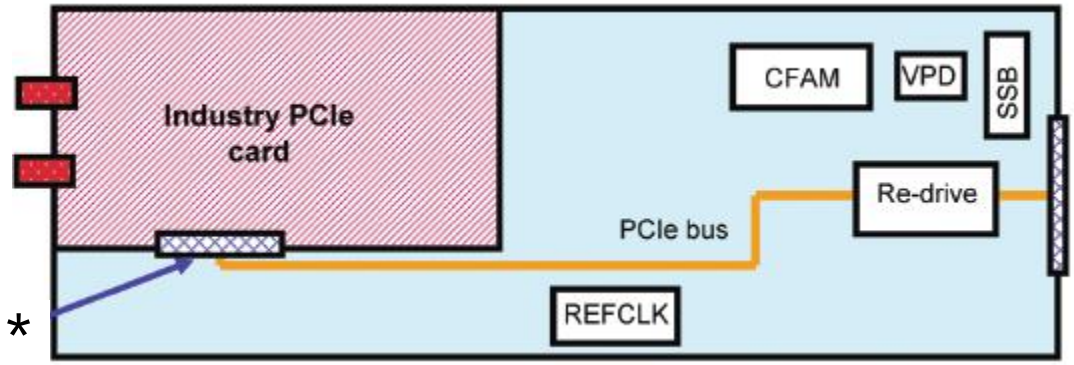

Native PCIe Feature: zEDC Express, 10GbE RoCE Express, and *Flash Express*

***PCIe Adapter Connector**

"Native PCIe" FUNCTION definition, assignment and mapping

- **Conceptually similar to channel (CHPID) or I/O device definition with different rules**

- **FUNCTION Definition in HCD or HCM to create IOCP input**
  - Uniquely identified by a hexadecimal **FUNCTION Identifier** (**FID**) in the range **000 – FFF**
  - **NOT** assigned to a Channel Subsystem so ANY LPAR can be assigned any FUNCTION.
  - Has a **PARTITION** parameter that dedicates it to **ONE** LPAR or allows reconfiguration among a group of LPARs. (**A FUNCTION can NOT be defined as shared**.)
  - If the intended PCIe hardware supports multiple partitions, has a decimal **Virtual Function Identifier** (**VF=**) in the range 1 – n, where n is the maximum number of partitions the PCIe feature supports. Examples: A RoCE feature supports up to 31 partitions, a zEDC Express feature supports up to 15
  - May have other parameters specific to the PCIe feature.
    For Example, 10GbE RoCE Express requires a **Physical Network Identifier (PNETID=)**.

- **FUNCTION Mapping to hardware**
  - Assign a Physical Channel Identifier (**PCHID=**) to identify the hardware feature in a specific PCIe I/O drawer and slot to be used for the defined FUNCTION.
  - Methods:
    - Manually using the configurator (eCONFIG) "AO Data" report
    - With assistance using the CHPID Mapping tool with eConfig Configuration Report File (CFR) input
  - **Note: Unlike CHPIDs, multiple FUNCTIONs can be mapped to the SAME PCHID.** This is conceptually similar to mapping multiple InfiniBand coupling CHPIDs to the same adapter and port.

Networking

OSA-Express5S 1000BASE-T Ethernet Feature - PCIe I/O Drawer

- **PCI-e form factor card supported by PCIe I/O drawer**
  - One two-port PCHID/CHPID per card
    Half the density of the OSA-Express3 version
- **Two small form factor pluggable (SFP+) transceivers (D1 top, D2 bottom)**
- **Auto-negotiates to 100 Mbps or 1 Gbps full duplex only**
- **RJ-45 connector on Cat 5 or better copper cable**
- **Operates at "line speed"**
- **CHPID TYPE Support:**

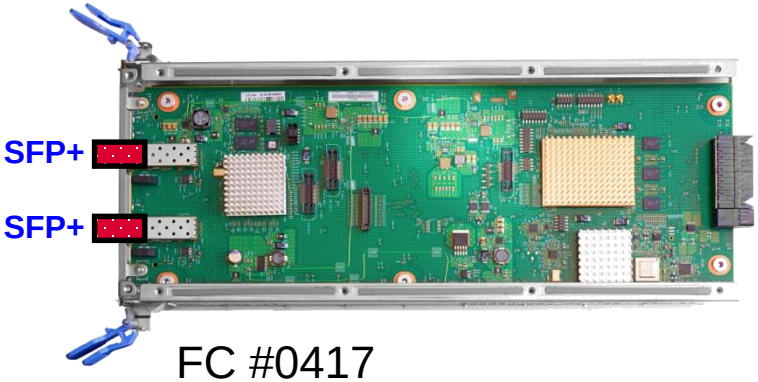

FC #0417

Connector = RJ-45

Mode	TYPE	Description
OSA-ICC	OSC	TN3270E, non-SNA DFT, OS system console operations
QDIO	OSD	TCP/IP traffic when Layer 3, Protocol-independent when Layer 2
Non-QDIO	OSE	TCP/IP and/or SNA/APPN/HPR traffic
Unified Resource Manager	OSM	Connectivity to intranode management network (INMN)
OSA for NCP (LP-to-LP)	OSN	NCPs running under IBM Communication Controller for Linux (CCL)

Note: OSA-Express5S feature are designed to have the same performance and to require the same software support as equivalent OSA-Express4S features.

OSA-Express5S fiber optic – PCIe I/O drawer

- **10 Gigabit Ethernet (10 GbE)**
  - CHPID types: OSD, OSX
  - Single mode (LR) or multimode (SR) fiber
  - One LR or SR SFP+ (D1 top)
    - 1 PCHID/CHPID
  - **Small form factor pluggable (SFP+) transceiver**
  - LC duplex

0415 – 10 GbE LR, # 0416 – 10 GbE SR

- **Gigabit Ethernet (GbE)**
  - CHPID types: OSD
  - Single mode (LX) or multimode (SX) fiber
  - Two LX or SX SFP+ (D1 top, D2 Bottom)
    - 1 PCHID/CHPID
  - **Small form factor pluggable (SFP+) transceivers**
  - LC Duplex

0413 – GbE LX, # 0414 – GbE SX

Note: OSA-Express5S features are designed to have the same performance as equivalent OSA-Express4S features.

OSA-Express5S and 4S 10 GbE Performance*

Inbound Streams – 1492 Byte MTU

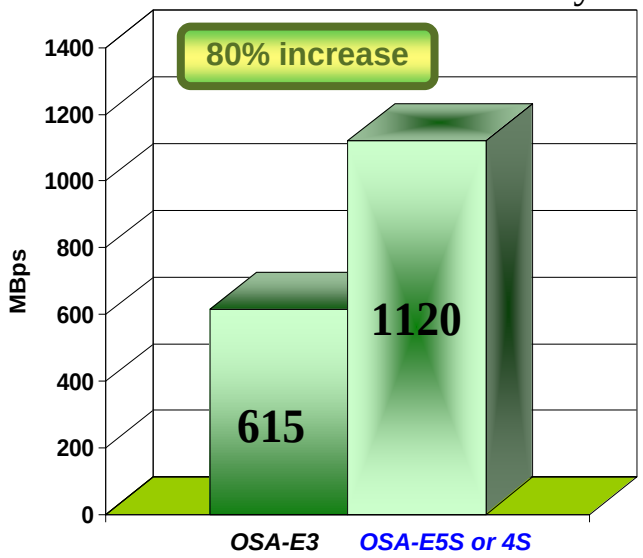

Mixed Streams – 1492 Byte MTU

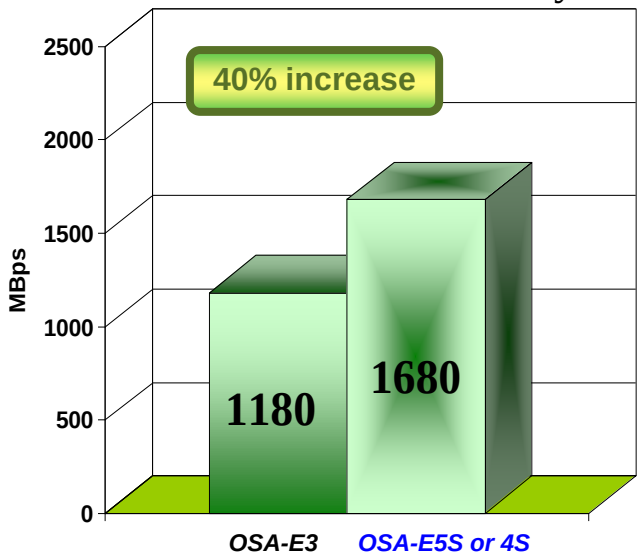

Inbound Streams – 8000 Byte MTU

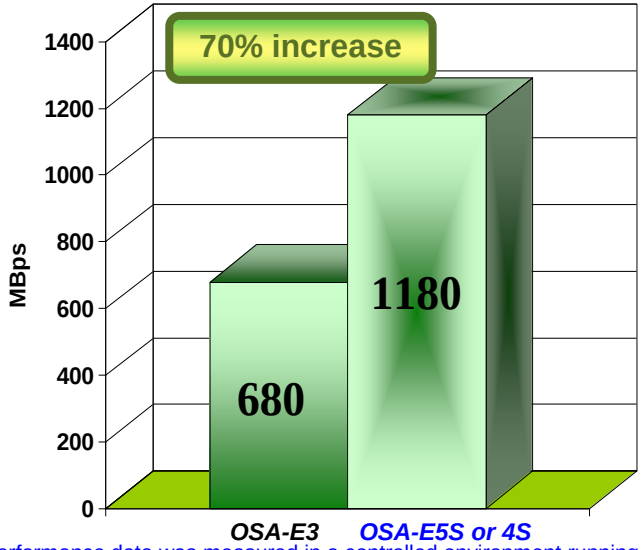

Mixed Streams – 8000 Byte MTU

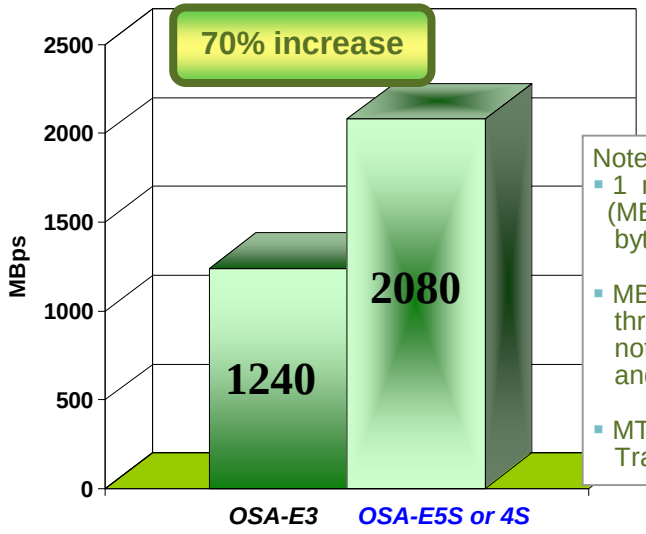

Notes:
- 1 megabyte per second (MBps) is 1,048,576 bytes per second
- MBps represents payload throughput (does not count packet and frame headers)
- MTU = Maximum Transmission Unit

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

Optimize server to server networking – transparently

"HiperSockets™-like" capability across systems

Up to 50% CPU savings for FTP file transfers across z/OS systems versus standard TCP/IP **

Up to 48% reduction in response time and 10% CPU savings for a sample CICS workload exploiting IPIC using SMC-R versus TCP/IP ***

Up to 40% reduction in overall transaction response time for WAS workload accessing z/OS DB2 ****

Up to 3X increase in WebSphere MQ messages delivered across z/OS systems ****

Shared Memory Communications (SMC-R):

Exploit RDMA over Converged Ethernet (RoCE) with qualities of service support for dynamic failover to redundant hardware

Typical Client Use Cases:

Help to reduce both latency and CPU resource consumption over traditional TCP/IP for communications across z/OS systems

Any z/OS TCP sockets based workload can **seamlessly** use SMC-R without requiring any application changes

NEW **z/OS V2.1 SMC-R**

NEW **z/VM 6.3 support for guests***

NEW **10GbE RoCE Express**

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

** Based on internal IBM benchmarks in a controlled environment using z/OS V2R1 Communications Server FTP client and FTP server, transferring a 1.2GB binary file using SMC-R (10GbE RoCE Express feature) vs standard TCP/IP (10GbE OSA Express4 feature). The actual CPU savings any user will experience may vary.

*** Based on internal IBM benchmarks using a modeled CICS workload driving a CICS transaction that performs 5 DPL (Distributed Program Link) calls to a CICS region on a remote z/OS system via CICS IP interconnectivity (IPIC), using 32K input/output containers. Response times and CPU savings measured on z/OS system initiating the DPL calls. The actual response times and CPU savings any user will experience will vary.

**** Based on projections and measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

Flash Express

Why Flash Express on z13?

- **Provides Storage Class Memory**
  - Implemented via NAND Flash SSDs (Solid State Drives) mounted in PCIe Flash Express features
  - Protected by strong AES Encryption done on the features
  - Not defined as I/O devices or with PCIe FUNCTIONs
  - Assigned to partitions similarly to Main Memory; but, not in the partition Image Profile. Reconfigurable.
  - Accessed using the new z Systems architected EADM (Extended Asynchronous Data Mover) Facility
  - Designed to enable extremely responsive paging of 4k pages to improve z/OS availability
  - Enables pageable large (1 MB) pages
- **Flash Express Exploitation**
  - z/OS V2.1, V1.13 + PTFs and RSM Enablement Offering
    - With z/OS Java SDK 7 SR3: CICS TS V5.1, WAS Liberty Profile V8.5, DB2 V11, IMS 12 and higher, SOD: Traditional WAS 8.0.0x*
    - CFCC Level 19 with WebSphere MQ for z/OS Version 7 MQ Shared Queue overflow support (March 31, 2014)
  - Linux on z Systems
    - SLES 11 SP3 and RHEL 6.4

10x

Faster response time and 37% increase in throughput compared to disk for morning transition

28%

Improvement in DB2 throughput leveraging Flash Express with Pageable Large Pages (PLP)

19%

Reduction in total dump time for a 36 GB standalone dump

~25%

Reduction in SVC dump elapsed time

*Note: All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Flash Express PCIe Adapter Card (Technology Refresh)

Four 400 GByte (G=10^{9}) SSDs support
1.4 TBytes (T=2^{40}) of Storage Class Memory
(AES encrypted)

Cable connections to form a RAID 10 Array across a pair of Flash Express Cards.

IBM zEnterprise Data Compression (zEDC)

zEDC Express feature

- **Designed to support high performance data serving by providing:**
  - A tenfold increase in data compression rates with much lower CP consumption than using software compression, including software compression that exploits the z Systems Compression Call instruction (z Systems hardware data compression)
  - A reduction in storage capacity required (creation of storage “white space”) that in turn reduces the cost of storage acquisition, deployment, operation, and management

- **Configuration:**
  - One compression accelerator per PCIe I/O feature card
  - Supports concurrent requests from up to 15 LPARs
  - Sustained aggregate 1 GBps compression rate when given large block inputs
  - Up to 8 features supported by zBC12 or zEC12
  - **Minimum two feature configuration recommended**

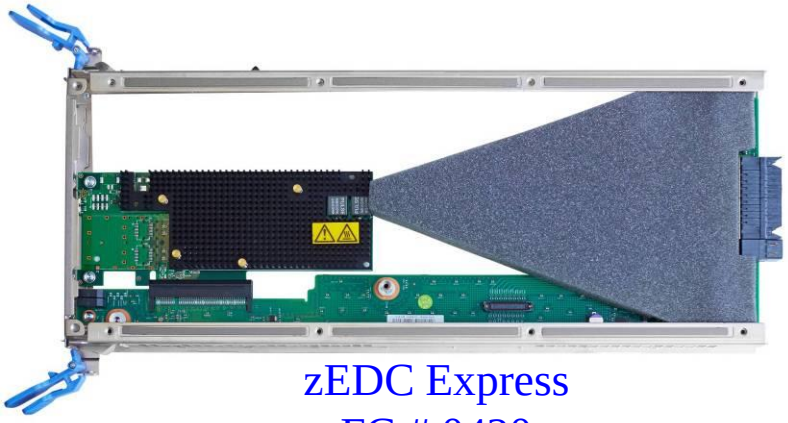

zEDC Express
FC # 0420

- **Exploitation and Compatibility**
  - Exclusive to zEC12 GA2 and zBC12
  - z/OS Support:
    - z/OS V2.1 – Hardware exploitation for SMF log data in September 2013, for IBM SDK for z/OS Java Technology Edition Version 7 Release 1 (5655-W43 and 5655-W44) with APAR OA43869 for zip and zlib compression, for BSAM and QSAM in 1Q2014 in PTFs for APAR OA42195, and for DFSMSdss and DFSMShsm SOD* for 3Q2014
    - z/OS V1.13 and V1.12 - Software support for decompression only, no hardware compression/decompression acceleration support

- **z/VM V6.3** support for z/OS V2.1 guest: June 27, 2014

*Note: All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

New hardware data compression accelerator can reduce CPU and storage

Every day 2.5 quintillion bytes of data are created

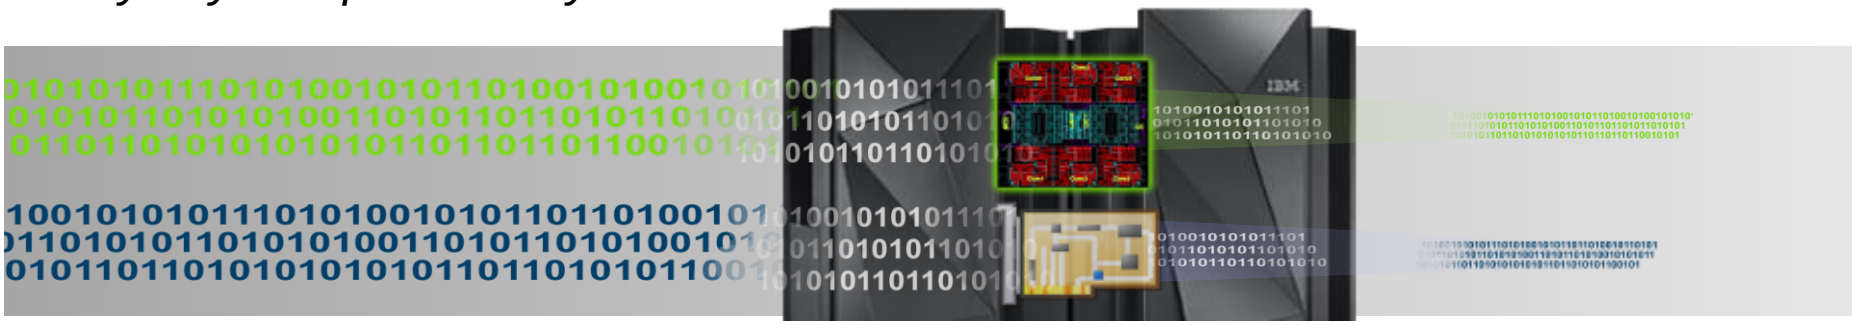

Compress your data *4X** *(efficient system data compression)*

*Up to 118X reduction in CPU and up to 24X throughput improvement when zlib uses zEDC ***

Data Ready

Efficiently compress active data by providing a low CPU, high performance, dedicated compression accelerator

Industry standard compliance compression for cross platform data distribution **

Typical Client Use Cases:

Significant disk savings with trivial CPU cost for large BSAM/QSAM sequential files

More efficiently store audit data in application logs

Reduce the amount of data needed for data migration and backup/restore **

Transparent acceleration of Java compressed applications **

NEW **zEDC Express**

NEW **z/VM 6.3 support for guests*****

NEW **z/OS V2.1 zEDC**

* The amount of data sent to an SMF logstream can be reduced by up to 75% using zEDC compression – reducing logger overhead

** These results are based on projections and measurements completed in a controlled environment. Results may vary by customer based on specific workload, configuration and software levels

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

Parallel Sysplex and Server Time Protocol

24x PCIe Gen3 Cable OM3/OM4 50/125 µm MM Cabling

- 24x PCIe Gen3 Cable required for new IBM Integrated Coupling Adapter (ICA SR)

- IBM qualified cables (Part numbers next chart) can be ordered from Anixter or IBM Global Technology
  - **Cable Distributor:**
    - Anixter ibmcabling@anixter.com or 877-747-2830
  - **Cable Suppliers:**
    - Computer Crafts http://www.computer-crafts.com/
    - TE Connectivity http://www.te.com/
    - Fujikura RBFiber@fujikura.com
- Fiber Core – 50 / 125 µm MM
- Connector – Single 24 fiber MPO – MPO
- Light Source – SX Laser
- Fiber bandwidth @ wavelength (OM4 Recommended)
  - 4700 MHz-km @850 nm OM4 for 150 m Max Length (Strongly Recommended)
  - 2000 MHz-km @850 nm OM3 for 100 m Max Length
- For more information, refer to
  - ***IBM z Systems Planning for Fiber Optic Links (FICON/FCP, Coupling Links, and Open System Adapters),* GA23-1407,** available in the Library section of Resource Link at http://www.ibm.com/servers/resourcelink/svc03100.nsf?OpenDatabase

24x PCIe Gen3 Cable Lengths OM3/OM4 50/125 µm MM Cabling

- IBM P/Ns for OM3,OM4 24-fiber cable assembly lengths (for ICA SR)

Item Description	IBM Cable P/N	Cable Length (m)	Cable Type	Connector Type
Fiber Optics – MPO / 24 OM4 (E1)				
Single 24-fiber cable assembly	00JA687	8.0m	**OM4**	MPO-MPO
Single 24-fiber cable assembly	00LU282	10.0m	**OM4**	MPO-MPO
Single 24-fiber cable assembly	00LU283	13.0m	**OM4**	MPO-MPO
Single 24-fiber cable assembly	00JA688	15.0m	**OM4**	MPO-MPO
Single 24-fiber cable assembly	00JA689	20.0m	**OM4**	MPO-MPO
Single 24-fiber cable assembly	00LU284	40.0m	**OM4**	MPO-MPO
Single 24-fiber cable assembly	00LU285	80.0m	**OM4**	MPO-MPO
Single 24-fiber cable assembly	00LU286	120.0m	**OM4**	MPO-MPO
Single 24-fiber cable assembly	00LU287	150.0m	**OM4**	MPO-MPO
Single 24-fiber cable assembly	00LU288	**Custom Length < 150.0m**	**OM4**	MPO-MPO
Fiber Optics – MPO / 24 OM4 (E1)				
Single 24-fiber cable assembly	00JJ548	8.0m	**OM3**	MPO-MPO
Single 24-fiber cable assembly	00LU290	10.0m	**OM3**	MPO-MPO
Single 24-fiber cable assembly	00LU291	13.0m	**OM3**	MPO-MPO
Single 24-fiber cable assembly	00JJ549	15.0m	**OM3**	MPO-MPO
Single 24-fiber cable assembly	00JJ550	20.0m	**OM3**	MPO-MPO
Single 24-fiber cable assembly	00LU292	40.0m	**OM3**	MPO-MPO
Single 24-fiber cable assembly	00LU293	80.0m	**OM3**	MPO-MPO
Single 24-fiber cable assembly	00LU294	100.0m	**OM3**	MPO-MPO
Single 24-fiber cable assembly	00LU295	**Custom Length < 100.0m**	**OM3**	MPO-MPO

Installation Planning for z13

z13 Requirements for Participation in a zEnterprise Ensemble

- Ensemble and Quality of Service (QoS) Features
  - Ensemble Feature: FC 0025 (always required for ensemble participation)
  - QoS selection: FC 0019, Manage, level or both FC 0019 and FC 0020, Automate, level
    Note: 1. All nodes in the same ensemble MUST have the same QoS feature level
    2. Priced Ensemble Blade and IFL Manage/Automate Features no longer exist on z13

- Intra-Node Management network (INMN) connectivity (Always required)
  - Two OSA-Express 1000BASE-T features to support two required OSM CHPIDs
    (Two OSA-Express5S FC 0417 or OSA-Express4S FC 0408 – CF only)
  - Two TYPE=OSM CHPIDs on the above, each cabled to a z13 internal System Control Hub (SCH)

- Intra-Ensemble Data Network (IEDN) connectivity with OSX (Optional)
  (Recommended for zBX connectivity, but OSD can be used)
  - One or more pairs of OSA-Express 10GbE features to support pairs OSX CHPIDs
    (OSA-Express5S 10 GbE LR FC 0415 or OSA-Express4S 10 GbE LR FC 0406 – CF only)
    (OSA-Express5S 10 GbE SR FC 0416 or OSA-Express4S 10 GbE SR FC 0407 – CF only)
  - Ordered to match LR or SR SFP optics features ordered for zBX
  - Cabled to the matching optics in the IEDN TOR switches in zBX

- Ensemble Primary and Alternate HMCs at Driver Level **22**
  - Identical hardware for both: Two HMC FCs 0091 or 0092 (deskside) or 0094 (rack mount)
  - Note: At this driver level, the Ensemble HMCs will also support nodes including zEC12, zBC12, z196, and z114 with or without managed zBX Model 3 or Model 2

Last Slide