zEnterprise.

A New Dimension in Computing

Vicom Infinity Inc.

Introduction to IBM System zEC12

Bob Lindsay

blindsay@vicominfinty.com

Trademarks

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

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

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

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*BladeCenter®, DB2®, e business(logo)®, DataPower®, ESCON, eServer, FICON®, IBM®, IBM (logo)®, MVS, OS/390®, POWER6®, POWER6+, POWER7®, Power Architecture®, PowerVM®, S/390®, ServerProven®, Sysplex Timer®, System p®, System p5, System x®, System z®, System z9®, System z10®, WebSphere®, X-Architecture®, z9®, z10, 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.
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 registered 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 from 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 from 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.

zEC12 Continues the CMOS Mainframe Heritage

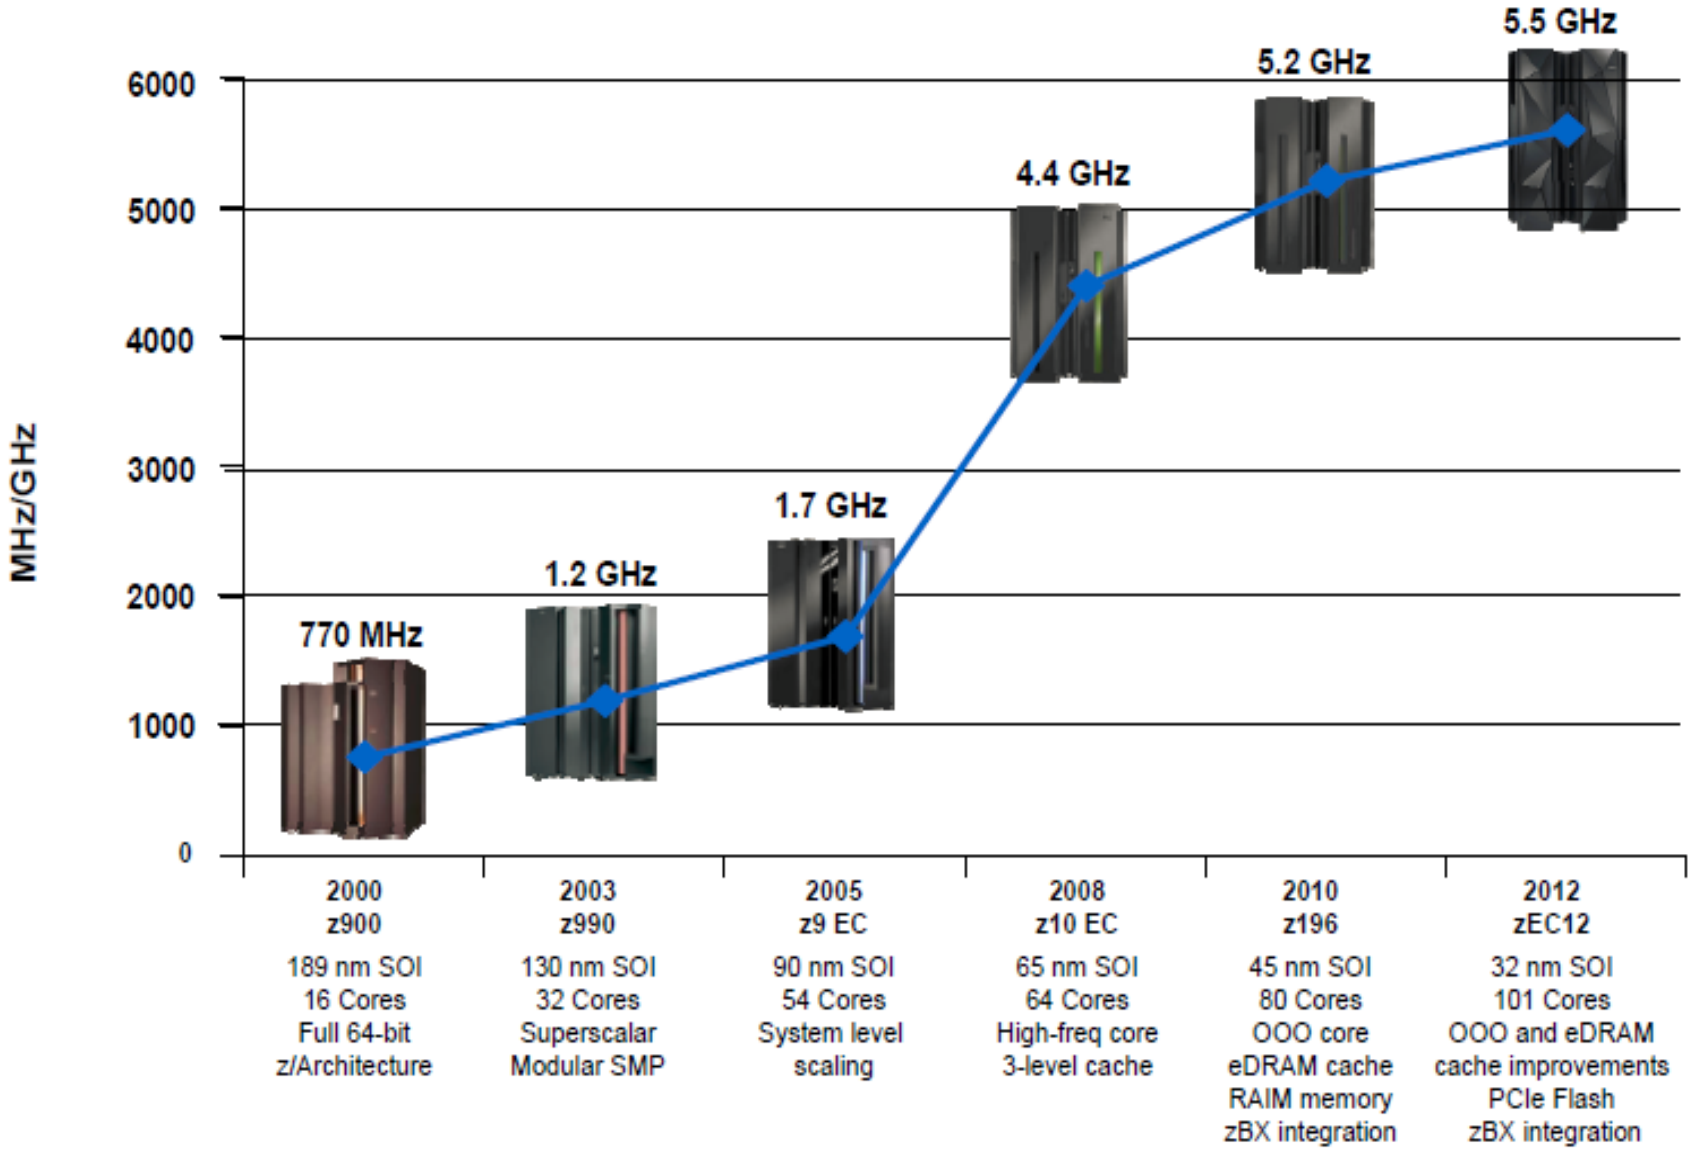

IBM System z: Design Comparison for High End Systems

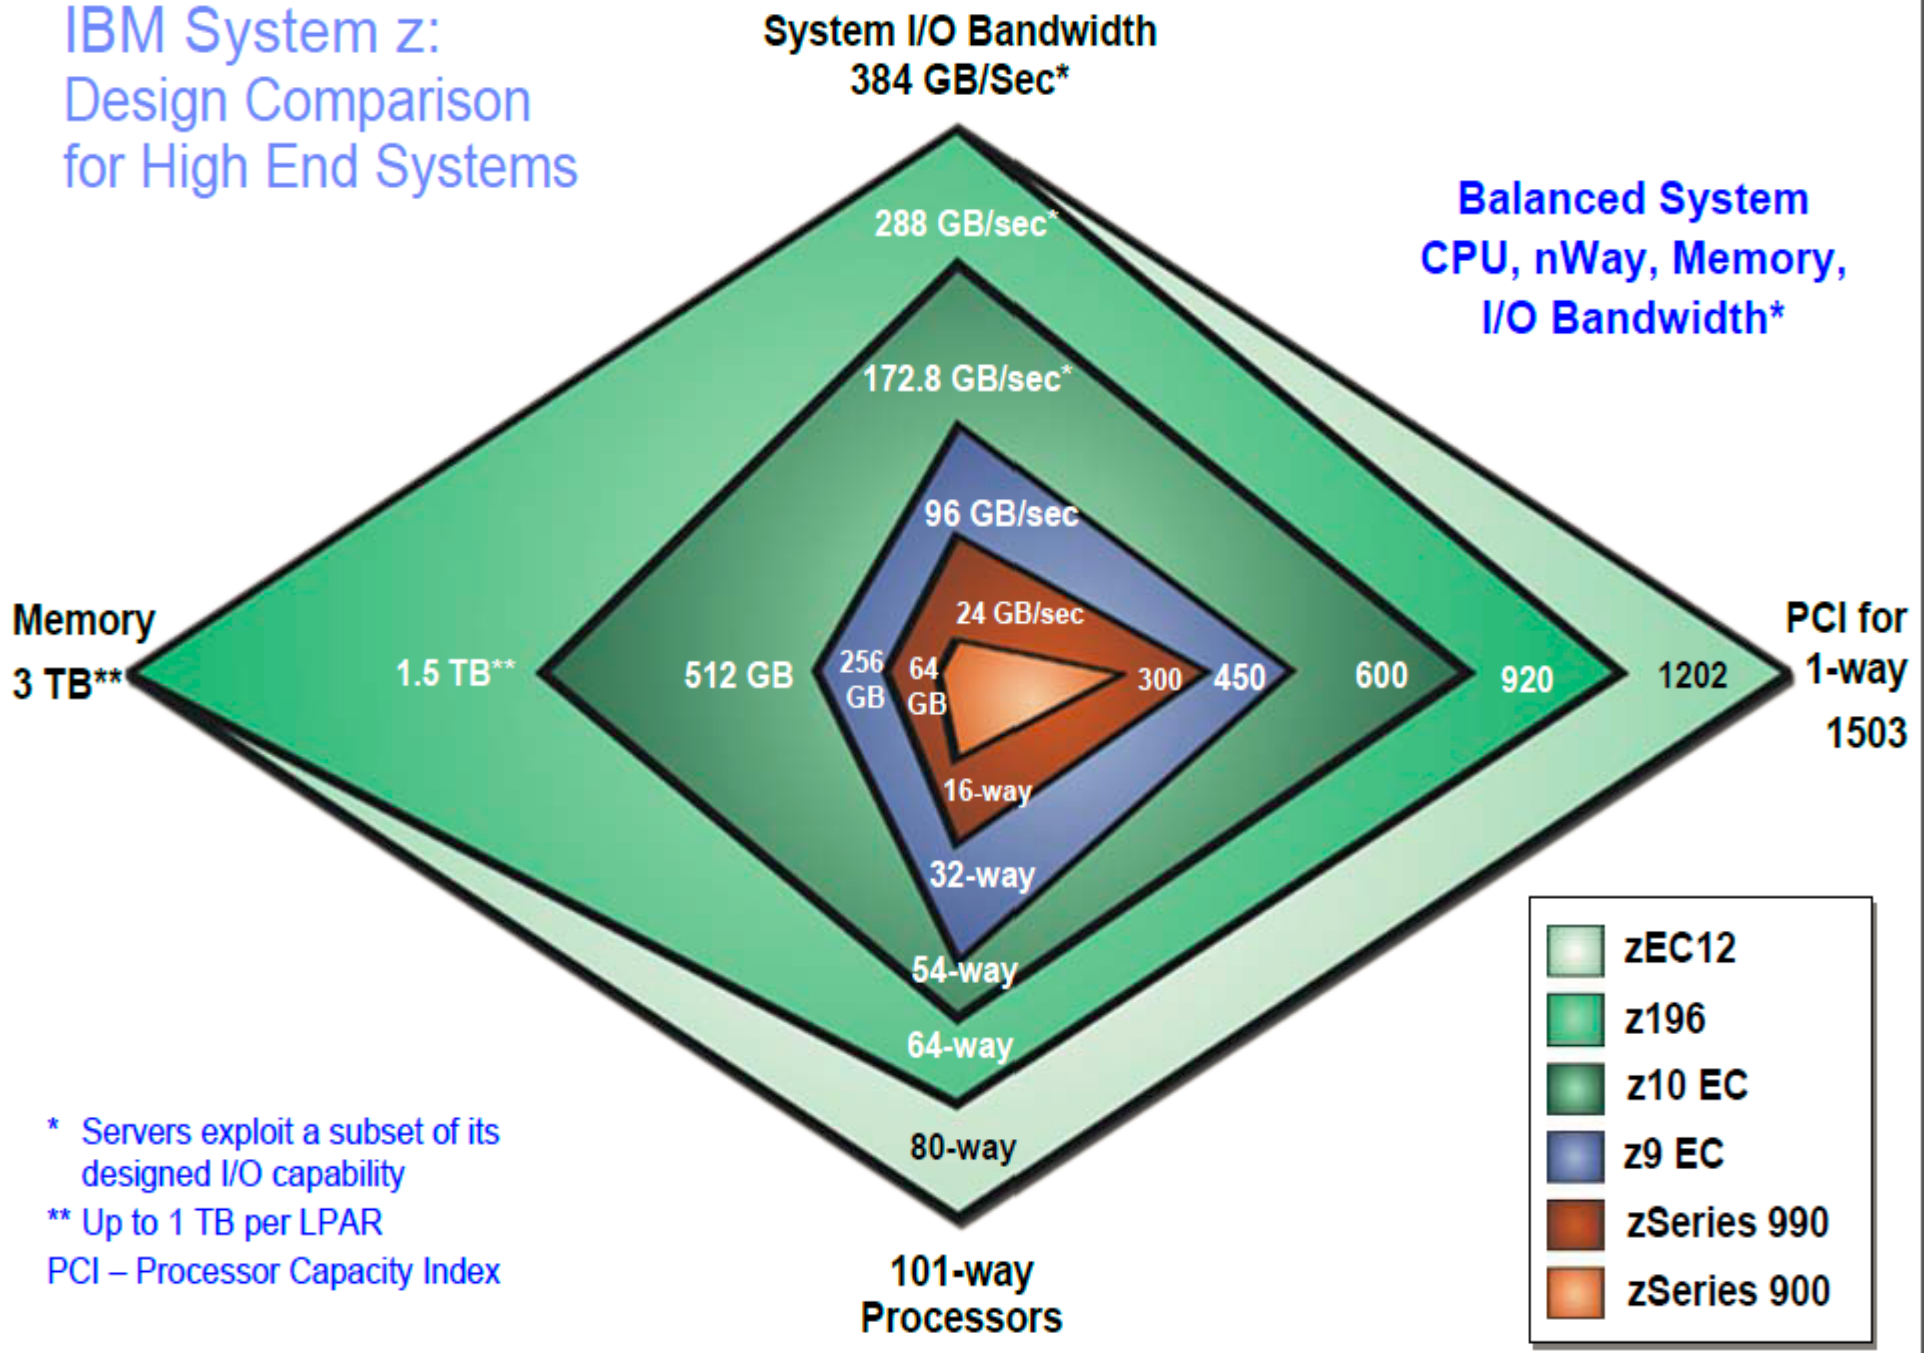

* Servers exploit a subset of its designed I/O capability

** Up to 1 TB per LPAR

PCI – Processor Capacity Index

IBM zEC12 Functions and Features

Processor, Memory, RAS

- Five Hardware Models
- Six core 32nm processor chip with 25% greater capacity per core than z196[1]
- Up to 101 configurable cores with 50% greater capacity than z196 model M80[1]
- 60 CP Subcapacity Settings
- Up to 3 TB memory
- z/Architecture Enhancements including 2 GB Pages, Transactional Execution and Runtime Instrumentation
- Flash memory resiliency improvements and pageable large page support
- Improved availability with IBM zAware

Parallel Sysplex, Security, I/O

- CFCC Level 18 Enhancements
- NTP security enhancements for STP
- OSA-Express4S 1000BASE-T
- Crypto Express4S (FIPS 140-2 cert and PKCS#11 support)
- Trusted Key Entry (TKE) 7.2

Upgradeable from z10 EC and z196

Environmental

- Optional Non Raised Floor
- Optional Overhead Power and Overhead I/O Cabling
- Improved N+1 Radiator-based Air Cooling
- Improved Optional Water Cooling with Air Cooled Backup
- Static Power Save Mode
- Optional High Voltage DC Power

Ensemble, Platform Management

- IBM zEnterprise BladeCenter Extension Model 3
- Unified Resource Manager enhancements for zEC12 and zBX Model 3
- Unified Resource Manager support for ensembles with zEC12, z196, z114, and zBX Models 2 and 3
- Doubled IEDN bandwidth internal to the zBX Model 3
- Upgraded POWER7 and System x General Purpose Blades Hypervisor Levels
- Continued incremental improvements

[1]Based on preliminary internal measurements and projections against a z196. Official performance data will be available upon announce and can be obtained online at LSPR (Large Systems Performance Reference) website at:

I/O Subsystem Internal Bus Interconnect Speeds (GigaBytes per second)

PCIe
zEC12
z196 GA2
z114

8 GBps

InfiniBand
zEC12
z10/z196 GA1

6 GBps

STI
z9

2.7 GBps

STI
z990/z890

2 GBps

STI
z900/z800

1 GBps

PCIe: Peripheral Component Interface (PCI) Express
STI: Self-Timed Interconnect

zEC12 still supports two I/O infrastructures for I/O feature cards

- **PCI Express 2 I/O infrastructure**
  - PCIe fanouts supporting 8 GBps PCIe I/O interconnect
  - PCIe switches with Redundant I/O interconnect in for I/O domains in the 7U, 32-slot, 4-domain PCIe I/O drawer introduced with z114 and z196 GA2
  - PCI Express I/O feature cards

    Based on selected industry standard PCIe I/O

    Designed to:

    Improve I/O port purchase granularity (fewer ports per card)

    Improve performance

    Increase I/O port density

    Lower energy consumption

- **InfiniBand I/O infrastructure (*Carry forward only, NO MES adds*)**
  - InfiniBand fanouts supporting a 6 GBps InfiniBand I/O interconnect
  - InfiniBand I/O card domain multiplexers with Redundant I/O interconnect in:

    The 14U, 28-slot, 7-domain I/O cage

    The 5U, 8-slot, 2-domain IO drawer
  - Selected legacy I/O feature cards

zEC12 Model H89 or HA1 Radiator (Air) Cooled – Under the covers

Front view

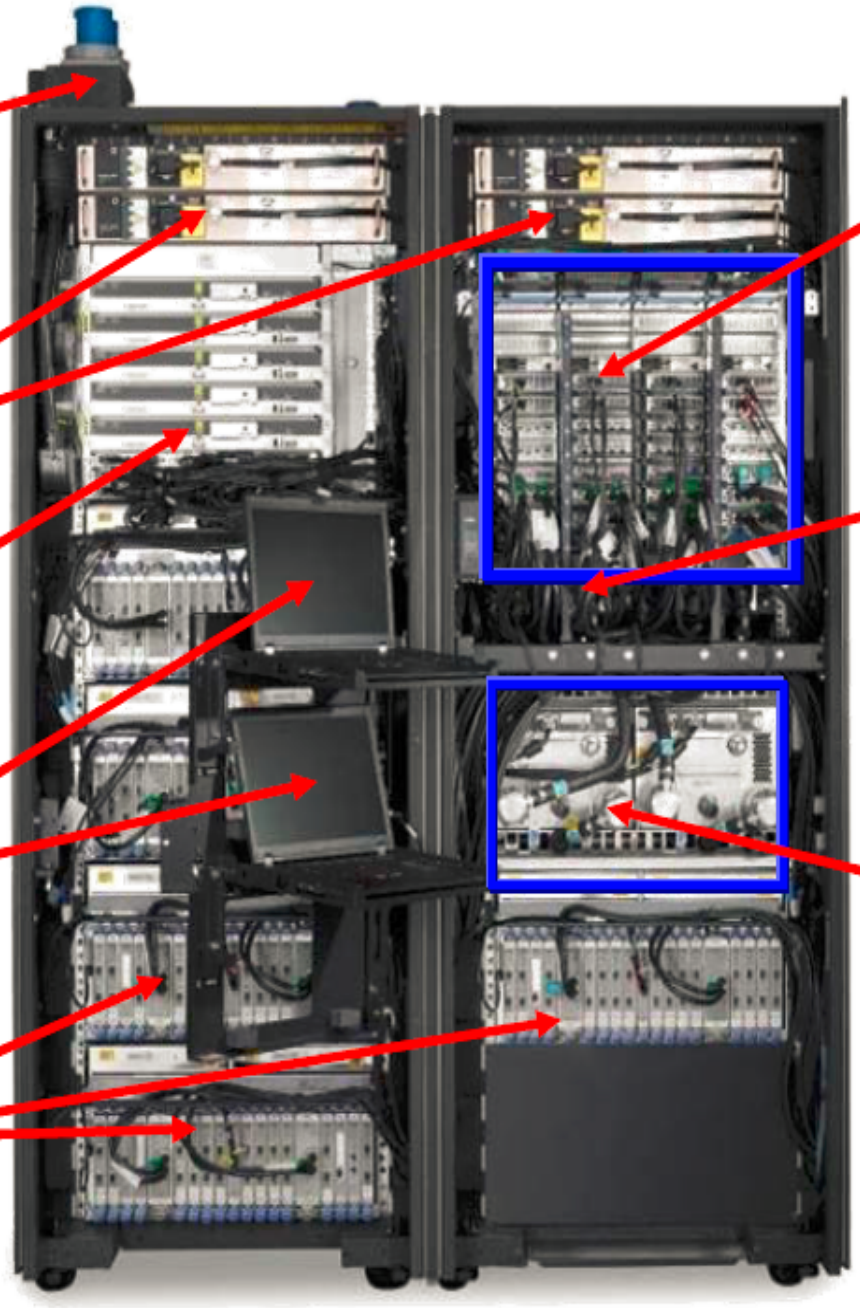

Overhead Power Cables (option)

Internal Batteries (option)

Power Supplies

2 x Support Elements

PCIe I/O drawers (Maximum 5 for zEC12)

Processor Books with Flexible Support Processors (FSPs), PCIe and HCA I/O fanouts

PCIe I/O interconnect cables and Ethernet cables for FSP cage controller cards

N+1 Radiator-based Air Cooling Unit

Optional FICON LX Fiber Quick Connect (FQC) not shown

28-slot I/O Cage (Introduced with z900, *Carry Forward only*)

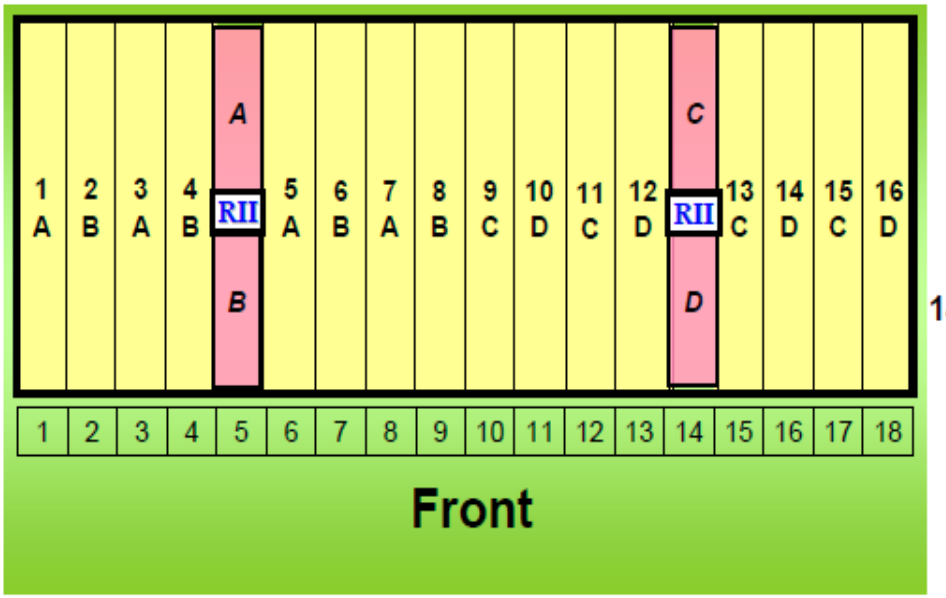

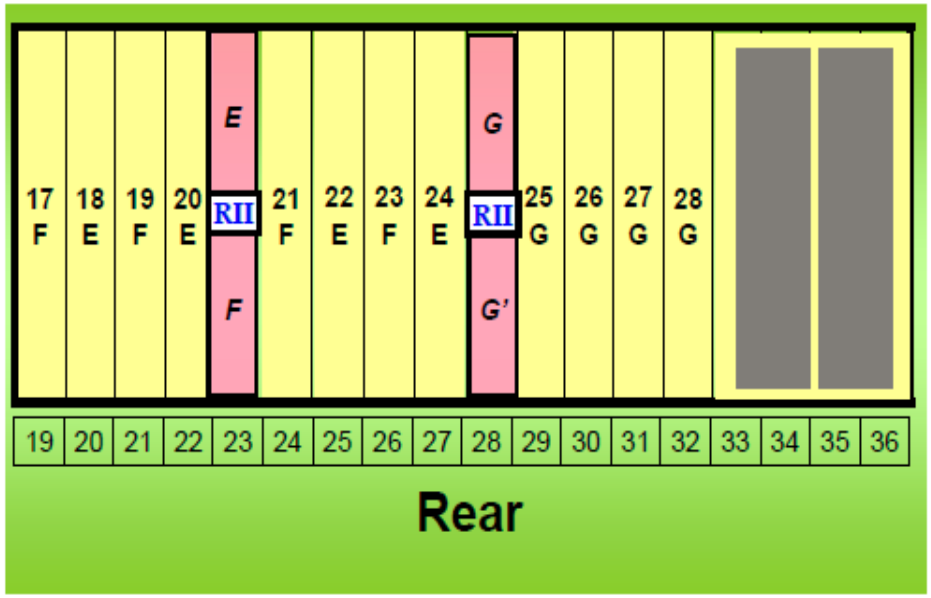

14 U

14 U

- Supports all carry forward non-PCIe I/O and Crypto Express3 cards
- Supports 28 I/O cards, 16 front and 12 rear, vertical orientation, in seven 4-card domains (shown as A to G)
- Requires eight IFB-MP daughter cards (A to G'), each connected to a 6 GBps InfiniBand I/O interconnect to activate all seven domains.
- To support Redundant I/O Interconnect (RII), the two interconnects to each domain pair (A-B, C-D, E-F, and G-G') must come from two different InfiniBand fanouts. (All seven domains in one of these cages can be activated with four fanouts.)
- Disruptive field install or remove
- Requires 14 EIA Units of space (24.5 inches ≈ 622 mm)

8-slot I/O Drawer (Introduced with z10 BC, *Carry Forward only*)

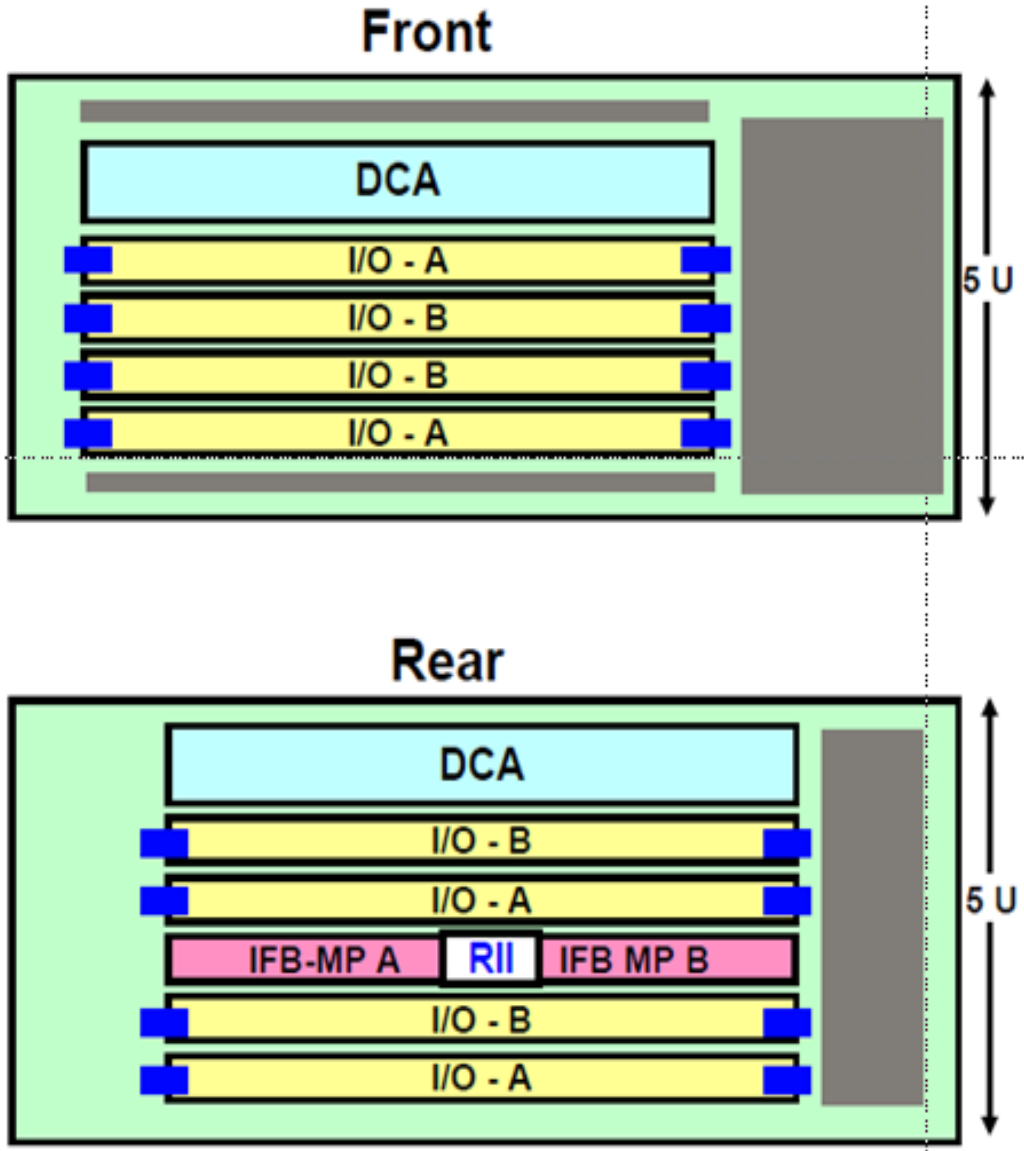

- Supports carry forward non-PCIe I/O and Crypto Express3 cards
- Supports 8 I/O cards, 4 front and 4 back horizontal orientation, in two 4-card dom (shown as A and B)
- Requires two IFB-MP daughter cards, ea connected to a 6 GBps InfiniBand interconnect to activate both domains.
- To support Redundant I/O Interconnect (between the two domains, the two interconnects must be from two different InfiniBand fanouts. (Two fanouts can sup two of these drawers.)
- Concurrent add, repair.
- Requires 5 EIA Units of space (8.75 inches ≈ 222 mm)

PCIe I/O drawer and PCIe I/O features

- **Increased infrastructure bandwidth**
  - PCI Express 2 x16 - ***8 GBps interconnect***
    (Compared to 6 GBps 12x InfiniBand DDR interconnect)
  - PCI Express 2 x8 - ***4 GBps*** available to PCIe I/O feature cards
    (Compared to 2 GBps or less available to I/O feature cards in the 8 slot drawer or 28 slot cage)
- **Compact**
  - Two 32-slot PCIe I/O drawers occupy the same space as one 28-slot I/O cage
  - Increases I/O port density 14% (Equivalent to an increase from 28 to 32 slots)
- **Improved I/O port purchase granularity**
  - "Half high" I/O feature cards compared to older I/O feature cards
  - Two ***FICON Express8S*** channels per feature (Four on FICON Express8)
  - One or two ***OSA-Express4S*** ports per feature (Two or four on OSA-Express3)
- **Reduced power consumption**
- **Designed for Improved Reliability, Availability, and Serviceability**
  - Concurrent field MES install and repair
  - Symmetrical, redundant cooling across all cards and power supplies
  - Temperature monitoring of critical ASICs

PCIe 32 I/O slot drawer (New build or Carry Forward)

Front

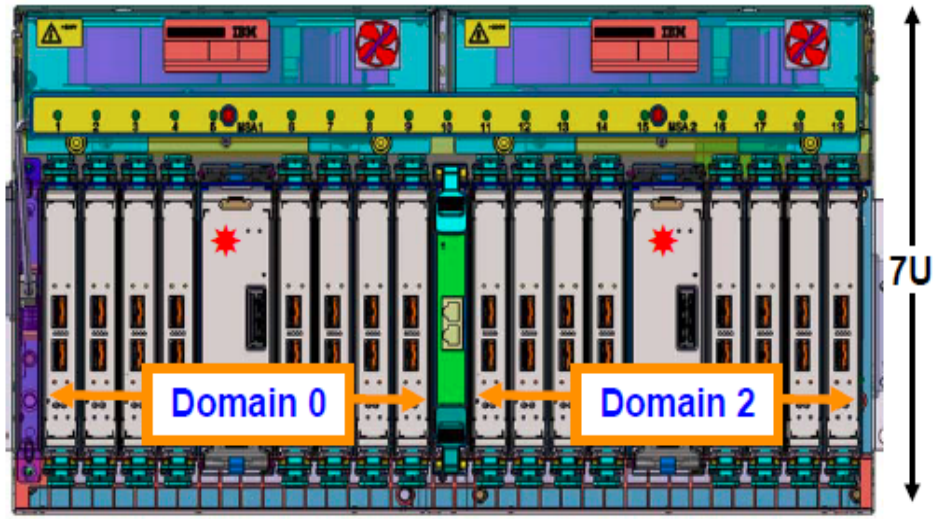

Rear

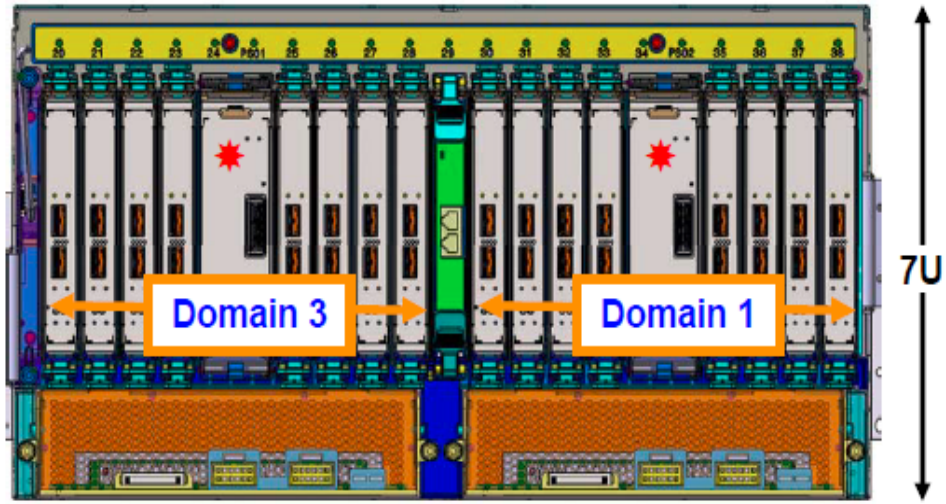

- Supports only new PCIe I/O cards.
- Supports 32 PCIe I/O cards, 16 front and 16 rear, vertical orientation, in four 8-card domains (shown as 0 to 3).
- Requires four PCIe switch cards (✸), each connected to an 8 GBps PCIe I/O interconnect to activate all four domains.
- To support Redundant I/O Interconnect (RII) between front to back domain pairs 0-1 and 2-3 the two interconnects to each pair must 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)

zEC12 I/O Features supported

Note - zEC12 does not offer "Plan Ahead" for I/O drawers or cages.

Supported features

- **Features – PCIe I/O drawer**
  - ***Crypto Express4S***
  - ***Flash Express***
  - ***FICON Express8S***
    SX and LX
  - ***OSA-Express4S***
    10 GbE LR and SR
    GbE SX, LX, and **1000BASE-T**

PCIe I/O drawer

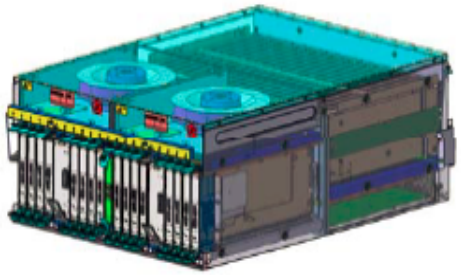

32 I/O slots

- **Features – I/O cage and I/O drawer (No MES adds)**
  - **Not Supported: ESCON, older FICON, FICON Express4 LX 4 km, OSA-Express2, PSC**
  - Crypto Express3 (Carry forward)
  - FICON Express8 (Carry forward)
  - FICON Express4 10 km LX and SX (Carry forward)
  - ISC-3 (Carry forward – **except RPQ 8P2602**)
  - OSA-Express3 (Carry forward)
    10 GbE, GbE, GbE 1000BASE-T

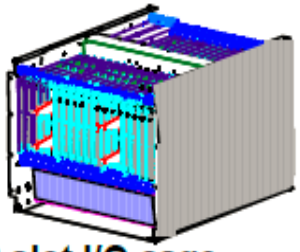

28 slot I/O cage

8 slot I/O drawer

Removal of Support for Crypto Express3 (August 28, 2012 Statement of Direction[1])

- **The IBM zEC12 is planned to be the last high-end System z server to offer support of the Crypto Express3 feature (#0864).**
  Crypto Express3 will not be supported on future high-end System z servers as carry forward on an upgrade. Enterprises should begin migrating from the Crypto Express3 feature to the Crypto Express4S feature (#0865).

Note 1: 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.

Removal of Support for FICON Express4 (August 28, 2012 Statement of Direction[1])

- **The IBM zEC12 is planned to be the last high-end System z server to offer support of the FICON Express4 features (#3321, #3322). FICON Express4 will not be supported on future high-end System z servers as carry forward on an upgrade.**
  Enterprises should continue migrating from the FICON Express4 features to the FICON Express8S features (#0409, #0410).

Note 1: 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.

Note: The August 28, 2012 announcement includes NO Statement of Direction about FICON Express8 support. Future "Carry Forward" support may or may not be possible.

Removal of Support for OSA-Express3 (August 28, 2012 Statement of Direction[1])

- **The IBM zEC12 is planned to be the last high-end System z server to offer support of the Open System Adapter-Express3 (OSA-Express3 #3362, #3363, #3367, #3370, #3371) family of features. OSA-Express3 will not be supported on future high-end System z servers as carry forward on an upgrade.** Enterprises should continue migrating from the OSA-Express3 features to the OSA-Express4S features (#0404, #0405, #0406, #0407, #0408).

Note 1: 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.

Removal of Support for ISC-3 Coupling Links (August 28, 2012 Statement of Direction[1])

- **The IBM zEC12 is intended to be the last high end System z server that will support ISC-3 coupling links.**
  Enterprises should begin migrating from ISC-3 features (#0217, #0218, #0219) to 12x InfiniBand (#0171 - HCA3-O) or 1x InfiniBand (#0170 - HCA3-O LR) coupling links.

Note 1: 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.

What happened to RPQ 8P2534? It is not available on zEC12.

- **RPQ 8P2534 History Lesson**
  - Introduced for z114 and z196 at GA2
  - Allows non-PCIe I/O cards to be added to "fill empty slots" if there is no way to add PCIe I/O cards
    (All PCIe slots full, no space to add a PCIe drawer)
  - Still available on z196 and z114

- **What do I do on zEC12 if more I/O features are needed and there is no way to add PCIe I/O?**
  - MES remove enough carry forward I/O to allow removal of an 8-slot I/O drawer (possibly concurrent)
    or a 28-slot I/O cage (definitely disruptive)
  - MES add PCIe I/O features and a 32-slot PCIe I/O drawer

- **Teaching points:**
  1. **Only a possible issue on "same serial number" push-pull MES upgrades from z10 or z196 to zEC12 because only those machines can carry forward non-PCIe I/O features.**
  2. **Plan ahead. Remove enough non-PCIe I/O on those upgrades to ensure that future PCIe I/O can be added.**

Removal of Support for ESCON (July 12, 2011 Statement of Direction) – FULFILLED

The IBM zEnterprise 196 and the IBM zEnterprise 114 ~~will be~~ are the last System z servers to support ESCON channels: IBM plans not to offer ESCON channels as an orderable feature on future System z servers. In addition, ESCON channels **cannot be carried forward** on an upgrade to such follow-on servers. This plan applies to channel path identifier (CHPID) types CNC, CTC, CVC, and CBY and to featured 2323 and 2324. System z customers should continue migrating from ESCON to FICON. Alternate solutions are available for connectivity to ESCON devices. IBM Global Technology Services offers an ESCON to FICON Migration solution, Offering ID #6948-97D, to help simplify and manage an all FICON environment with continued connectivity to ESCON devices if required.

Notes:

- For z196, this new Statement of Direction restates the SOD in Announcement letter 111-112 of **February 15, 2011**. It also confirms the SOD in Announcement letter 109-230 of **April 28, 2009** that "ESCON Channels will be phased out."

Optica PRIZM FICON Converter

http://www.opticatech.com/

Supports the elimination of ESCON channels on the host while maintaining ESCON and Bus/Tag-based devices and applications

What is PRIZM?

- A purpose-built appliance designed exclusively for IBM System z; enables ESCON devices to be connected to FICON channels or fabrics
- Allows ESCON devices to connect to FICON channels and FICON fabrics/networks
  - Prizm also supports attachment of parallel (bus/tag) devices to FICON channels via ESBT module
- Converts 1 or 2 FICON channels (CHPID type FC) into 4, 8 or 12 ESCON channels
  - Replace aging ESCON Directors with PRIZM (maintenance savings)
  - Achieve streamlined infrastructure and reduced Total Cost of Ownership
- Qualified by the IBM Vendor Solutions Lab in POK for all ESCON devices; qualified for connectivity to Brocade and Cisco FICON switching solutions
  - Refer to: http://www-03.ibm.com/systems/z/hardware/connectivity/index.html
    - Products -- > FICON / FCP Connectivity -- > Other supported devices
- PRIZM is available via IBM Global Technology Services: ESCON to FICON Migration offering (#6948-97D)

Front

Back

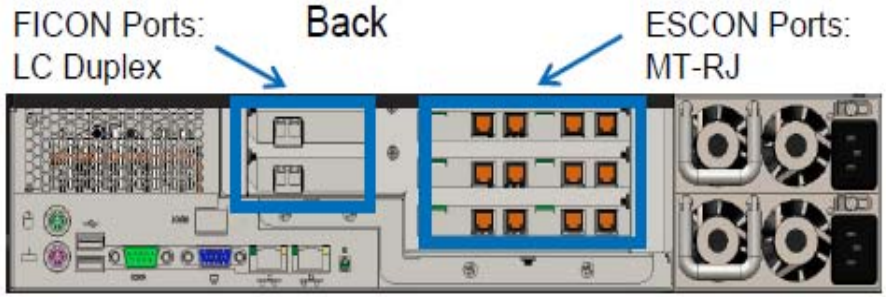

Where Does Prizm Fit in the Data Center?

Topologies supported by Prizm

- Local: direct attached or switched
- Remote: ISL (cascaded) or IP channel extended

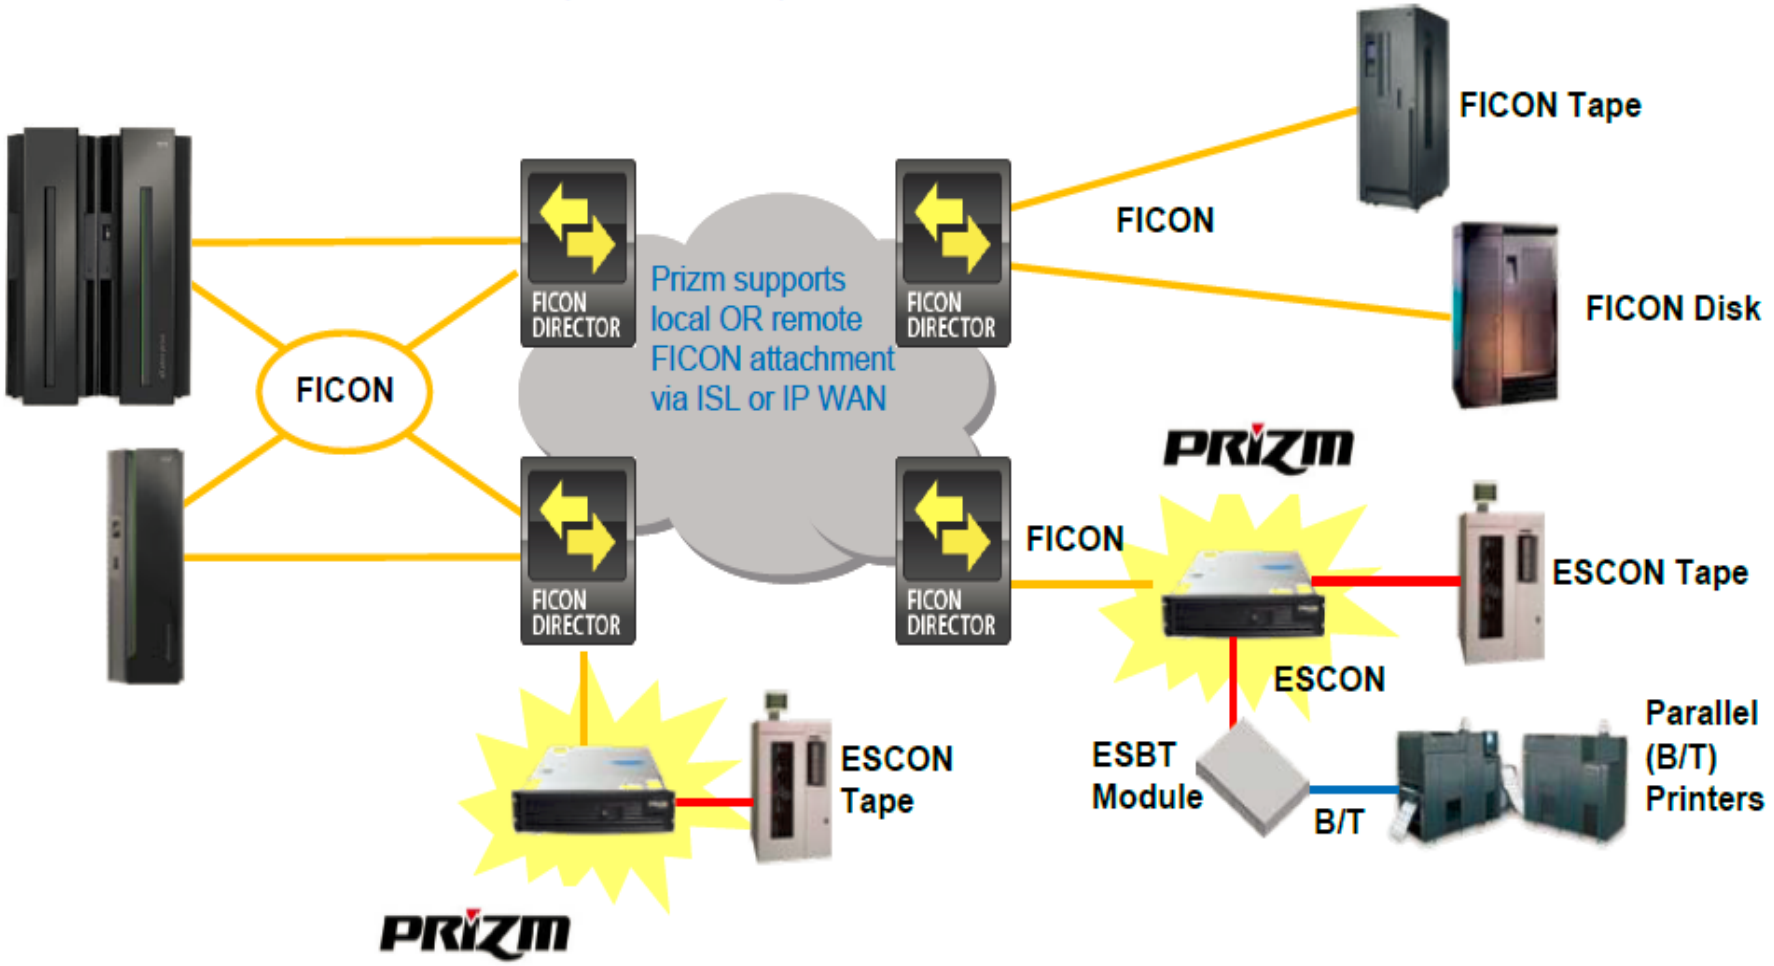

Flash Express

What Is Flash Express?

- Flash Express is internal storage implemented via NAND Flash SSDs (Solid State Drives) mounted in PCIe Flash Express feature cards
  - Which plug into PCIe I/O drawers in pairs
  - Data security provided on the feature cards
  - A pair provides 1.6 TB of storage
  - A maximum of 4 pairs are supported in a system

- Flash Express card "plugging rules"
  - Only two pairs of Flash Express cards per PCIe I/O drawer
  - Slots 1 and 14 on the front side and 25 and 33 on the rear are reserved for Flash Express
  - These will be the last slots used to plug other PCIe features

- Internal Flash is accessed using the new System z architected EADM (Extended Asynchronous Data Mover) Facility
  - An extension of the ADM architecture used in the past with expanded storage
  - Flash Express cards do NOT require definition in the IOCDS
  - Access is initiated with a Start Subchannel instruction
  - Flash Express does not reduce the number of available I/O device subchannels
  - Subchannels used were previously reserved

Flash Express - Twin-Cable Interconnects a RAID 10 Mirrored Pair

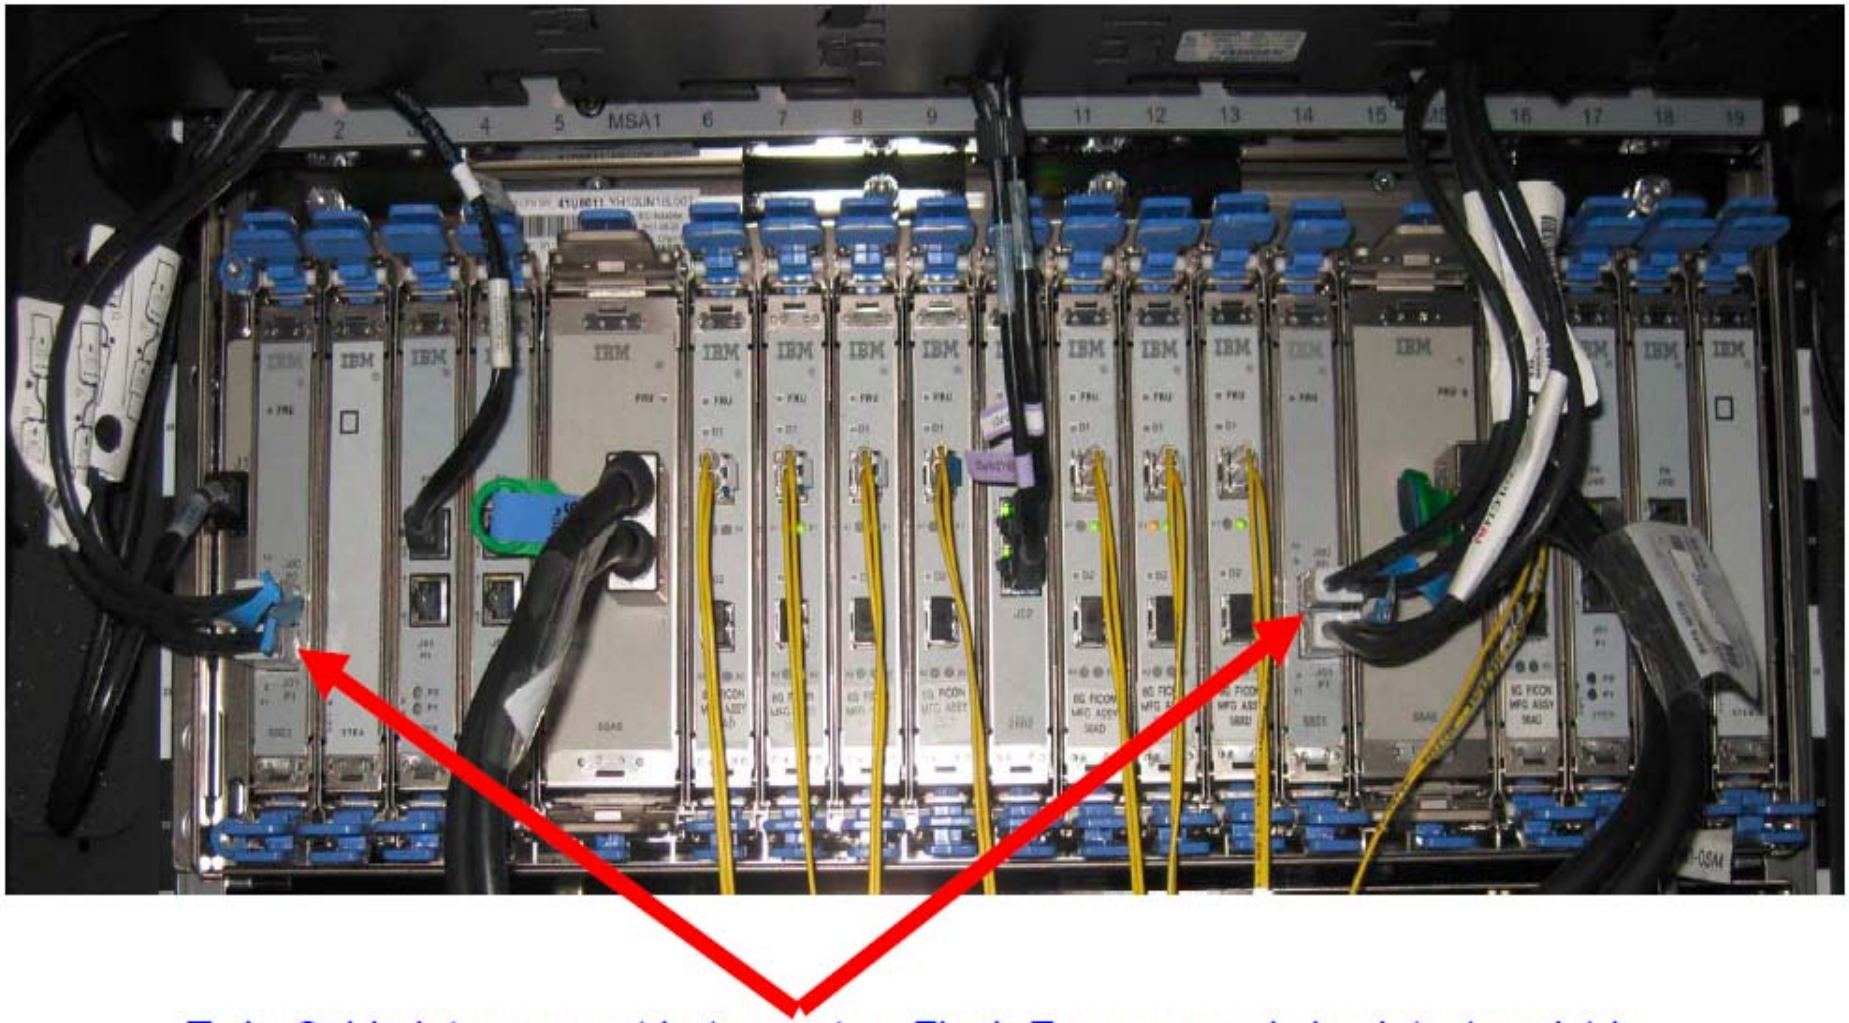

Twin-Cable interconnect between two Flash Express cards in slots 1 and 14

Crypto Express4S

- **One PCIe adapter per feature**
  - Initial order – two features
- **FIPS 140-2 Level 4**
- **Installed in the PCIe I/O drawer**
- **Up to 16 features per server**
- **Prerequisite: CPACF (#3863)**

Three configuration options for the PCIe adapter

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

- Accelerator
  - For SSL acceleration
  - Clear key RSA operations

- **Enhanced: Secure IBM CCA coprocessor (default)**
  - Optional: TKE workstation (#0841) for security-rich, flexible key entry or remote key management

- **New: IBM Enterprise PKCS #11 (EP11) coprocessor**
  - Designed for extended evaluations to meet public sector requirements
    - Both FIPS and Common Criteria certifications
  - ***Required:*** TKE workstation (#0841) for management of the Crypto Express4S when defined as an EP11 coprocessor

FICON Express8S – PCIe I/O drawer

- **For FICON, zHPF, and FCP environments**
  - CHPID types: FC and FCP
  - 2 PCHIDs/CHPIDs
- **Auto-negotiates to 2, 4, or 8 Gbps**
- **Increased performance compared to FICON Express8**

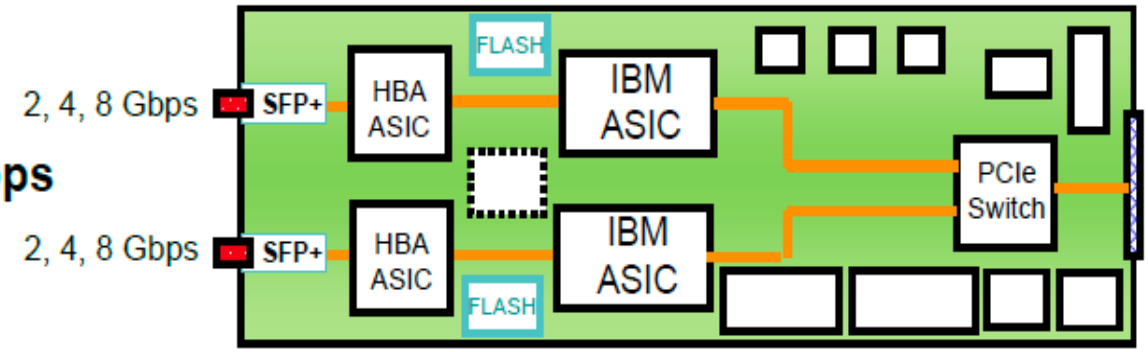

- **10KM LX - 9 micron single mode fiber**
  - Unrepeated distance - 10 kilometers (6.2 miles)
  - Receiving device must also be LX
  - **Note: Only LX has FQC support**
- **SX - 50 or 62.5 micron multimode fiber**
  - Distance variable with link data rate and fiber type
  - Receiving device must also be SX
- **2 channels of LX or SX (no mix)**
- **Small form factor pluggable (SFP) optics**
  - Concurrent repair/replace action for each SFP

0409 – 10KM LX, # 0410 – SX

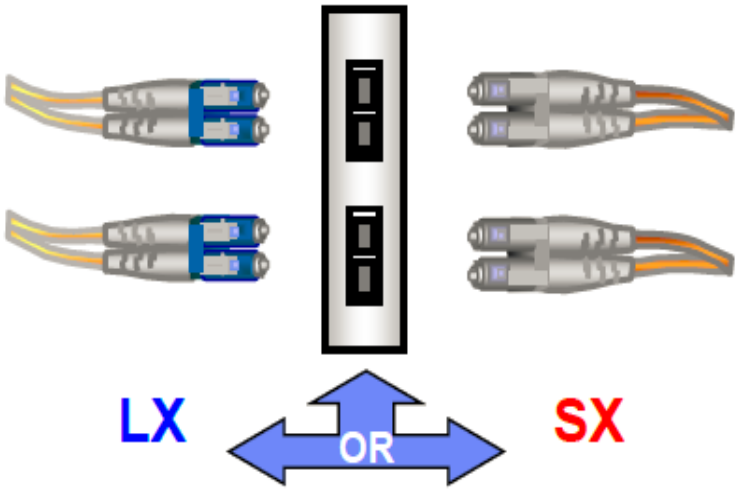

FICON Express8S – PCIe I/O drawer

- **For FICON, zHPF, and FCP environments**
  - CHPID types: FC and FCP
  - 2 PCHIDs/CHPIDs
- **Auto-negotiates to 2, 4, or 8 Gbps**
- **Increased performance compared to FICON Express8**

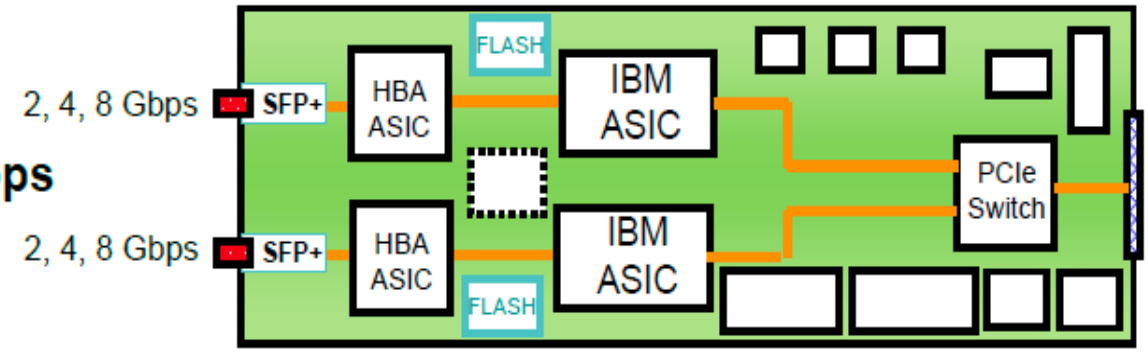

- **10KM LX - 9 micron single mode fiber**
  - Unrepeated distance - 10 kilometers (6.2 miles)
  - Receiving device must also be LX
  - **Note: Only LX has FQC support**
- **SX - 50 or 62.5 micron multimode fiber**
  - Distance variable with link data rate and fiber type
  - Receiving device must also be SX
- **2 channels of LX or SX (no mix)**
- **Small form factor pluggable (SFP) optics**
  - Concurrent repair/replace action for each SFP

0409 – 10KM LX, # 0410 – SX

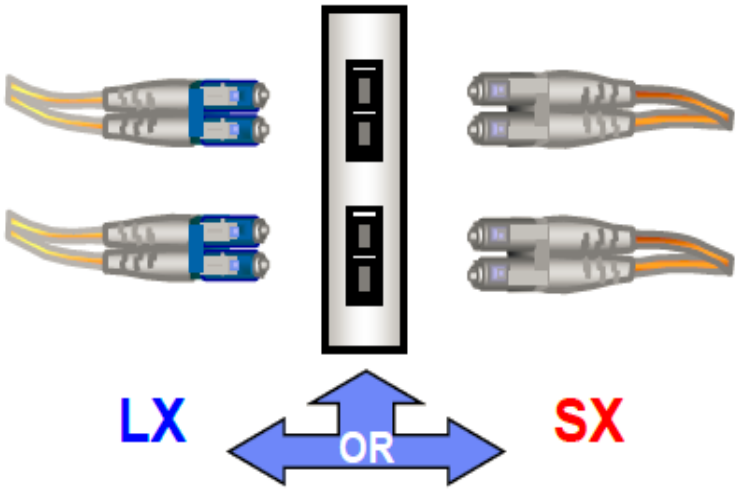

FCP performance on System z

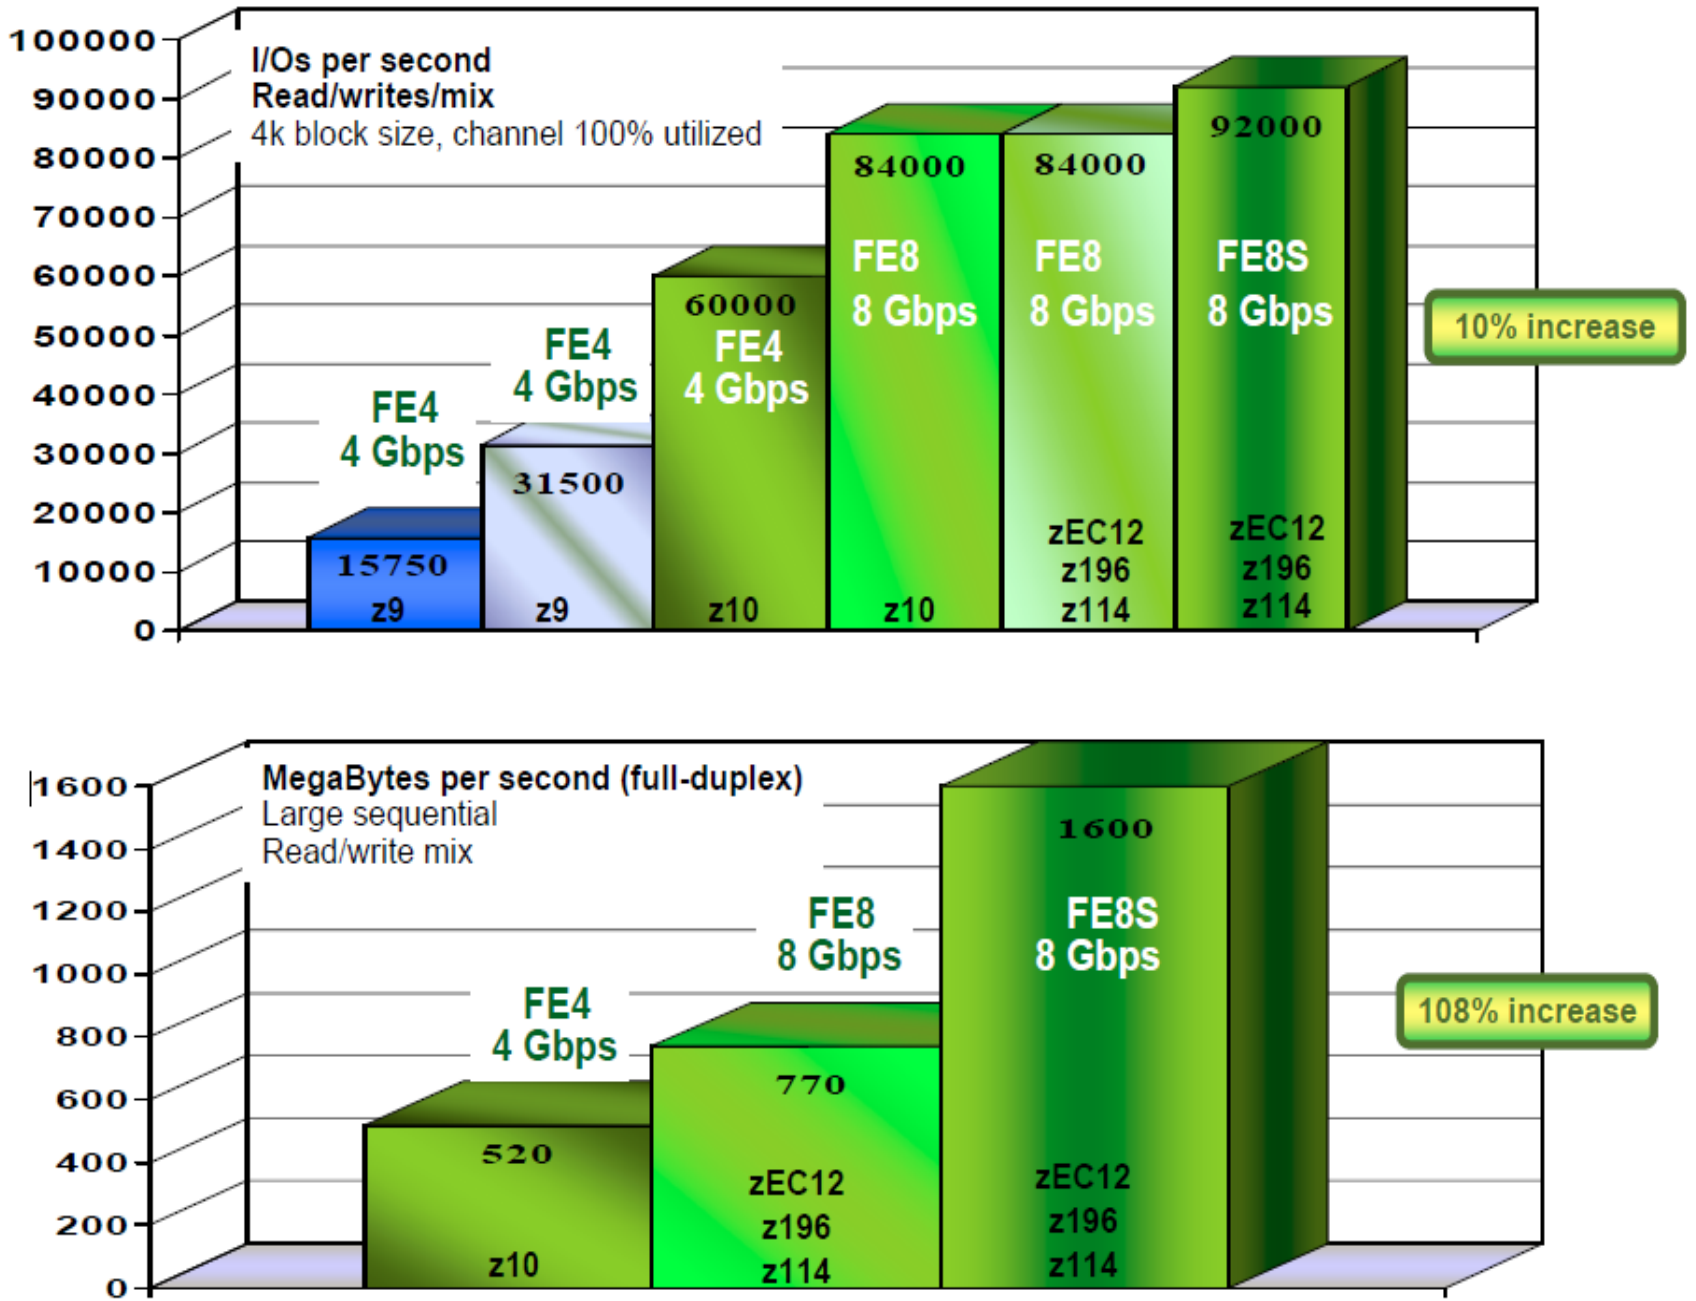

OSA-Express4S 1000BASE-T Ethernet Feature - PCIe I/O Drawer

- **PCI-e form factor card supported by PCIe I/O drawer**
  - One two-port CHPID per card
    Half the density of the OSA-Express3 version
- **Auto-negotiation to 10, 100, 1000 Mbps over Category 5 or better copper**
- **RJ-45 connector**
- **Operates at "line speed"**
- **CHPID TYPE Support:**

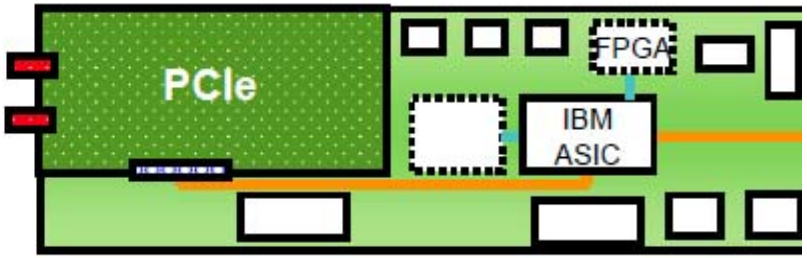

FC 0408

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

zEC12 and zBX Model 3 Ensemble Connectivity

- Networks – Redundant zBX-003 "B" frame top of rack switches
  - Up to 8 nodes, zEC12, z196, and z114 in a common Ensemble
  - Intranode management network (**INMN**) – 1000BASE-T (OSM)
  - Intraensemble data network (**IEDN**) –
    10 GbE LR or SR (OSX)
  - Corporate network to zBX - 10 GbE LR or SR, GbE LX or SX
    Note: zEC12 may also connect via the corporate network (OSD)
- Storage – Redundant 8 Gbps SW Fibre Channel
  - Local SAN storage for IBM blades

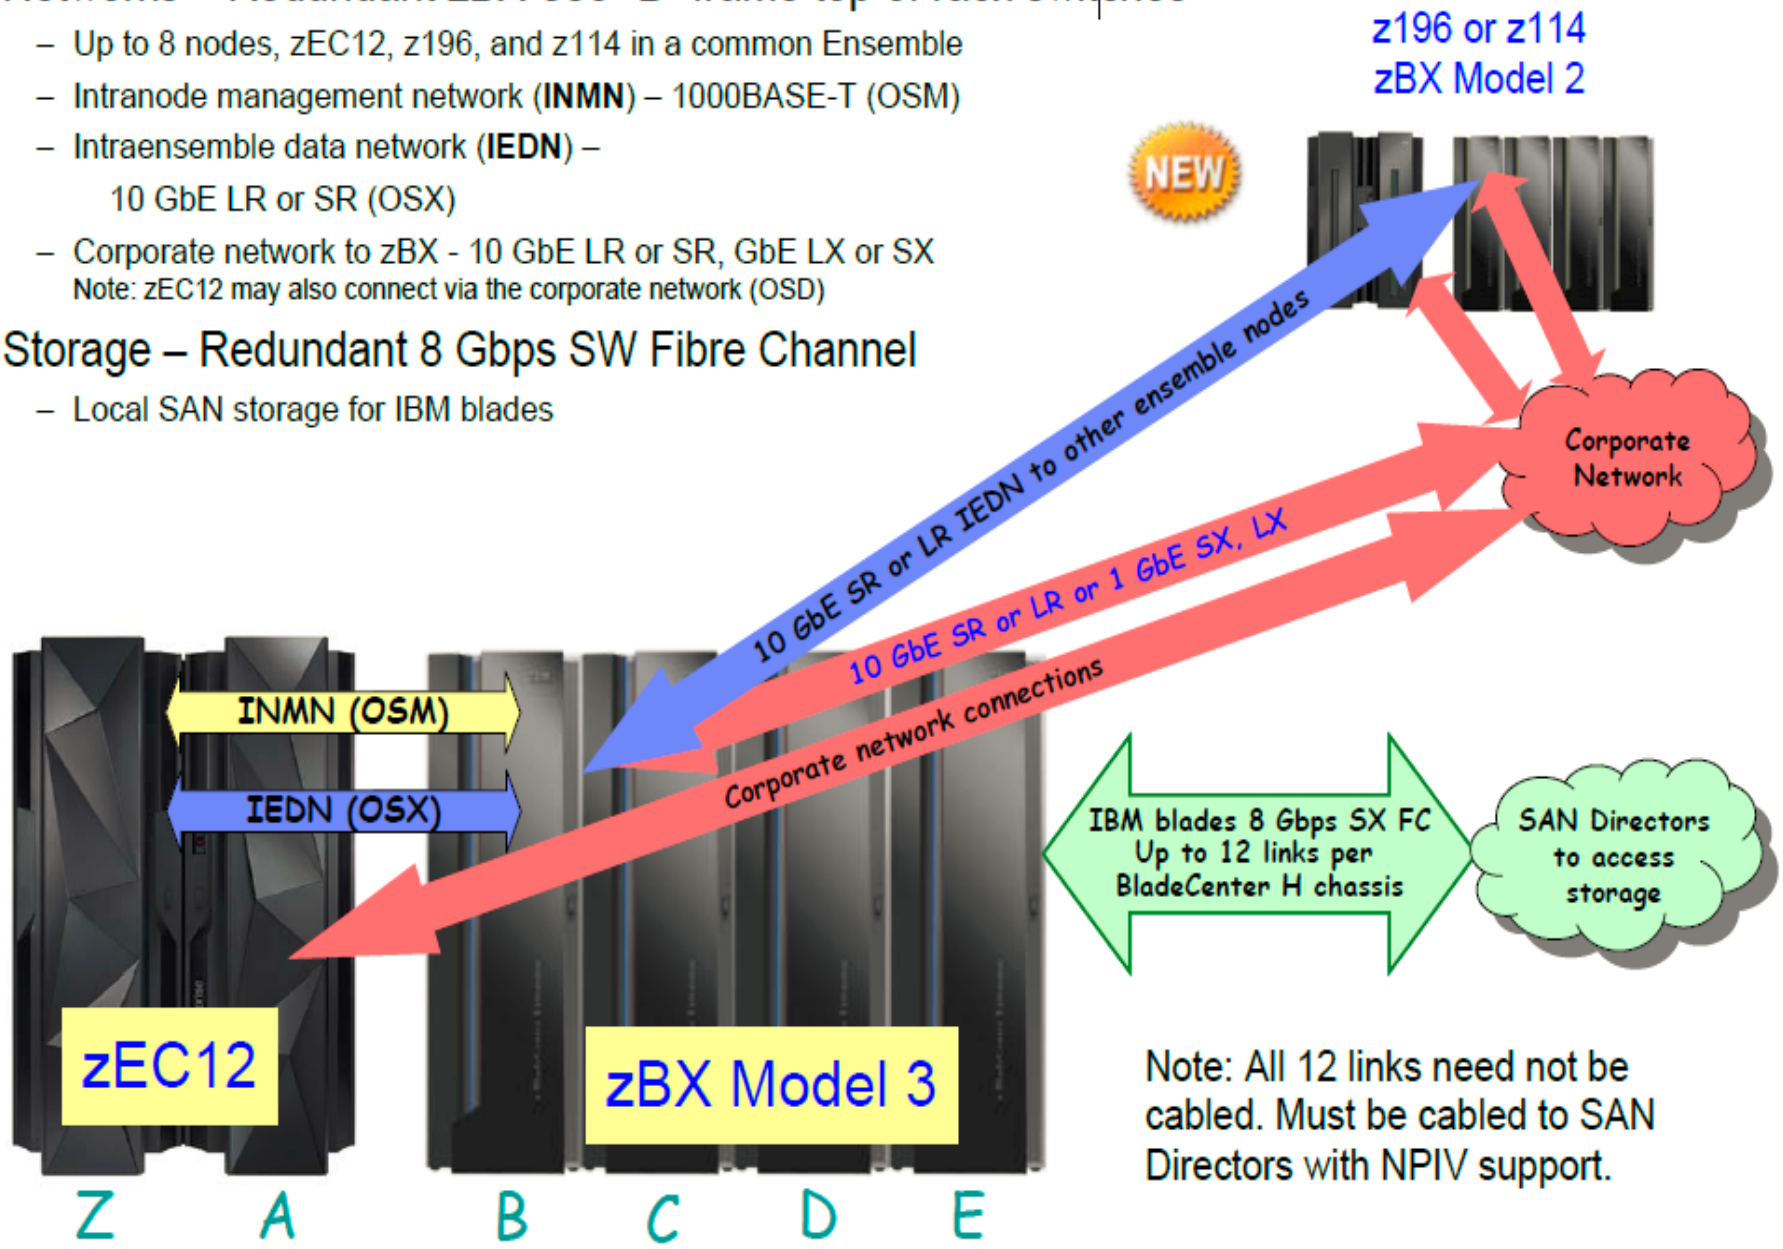

Note: All 12 links need not be cabled. Must be cabled to SAN Directors with NPIV support.

Removal of support for Ethernet half-duplex operation and 10 Mbps link data rate (August 28, 2012 Statement of Direction[1])

- **The OSA-Express4S 1000BASE-T Ethernet feature is planned to be the last copper Ethernet feature to support half-duplex operation and a 10 Mbps link data rate. The IBM zEC12 is planned to be the last Enterprise Class System z server to support half-duplex operation and a 10 Mbps link data rate for copper Ethernet environments.**

  Any future 1000BASE-T Ethernet feature will support full-duplex operation and auto-negotiation to 100 or 1000 Mbps exclusively.

  - **This restates the statement of direction in the March 6, 2012 announcement letter ENUS112-026 of March 6, 2012.**

Note 1: 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.

OSA-Express4S Fiber Optic Ethernet Features – PCIe I/O drawer

- **10 Gigabit Ethernet (10 GbE)**
  - CHPID types: OSD, OSX
  - LR - Single mode or SR - Multimode fiber
  - One port of LR or one port of SR
  - **One one-port CHPID**

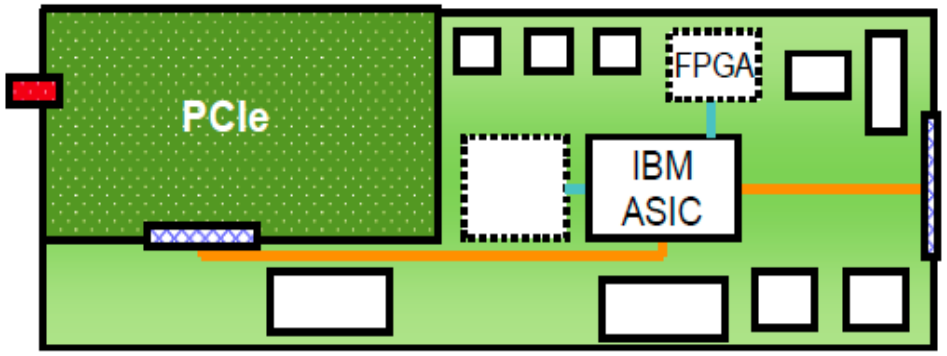

FC#0406 – 10 GbE LR, FC# 0407 – 10 GbE SR

- **Small form factor optics – LC Duplex**

- **Gigabit Ethernet (GbE)**
  - CHPID types: OSD (OSN not supported)
  - LX - Single mode or SX - Multimode fiber
  - Two ports of LX or two ports of SX
  - **One two-port CHPID**

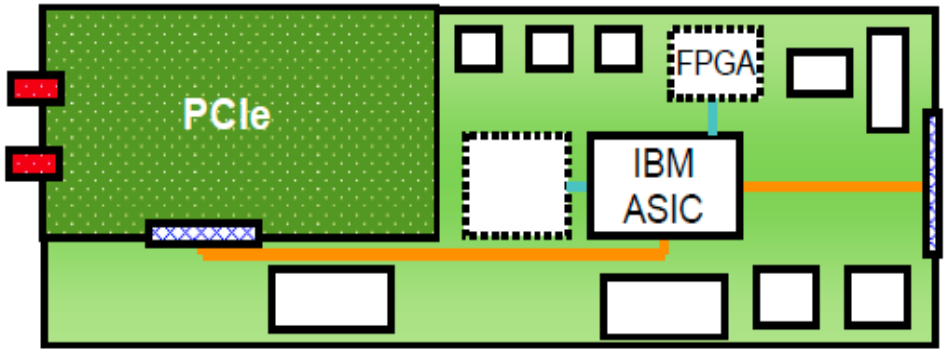

FC# 0404 – GbE LX, FC# 0405 – GbE SX

- **Small form factor optics – LC Duplex**

OSA-Express4S 10 GbE performance (laboratory)

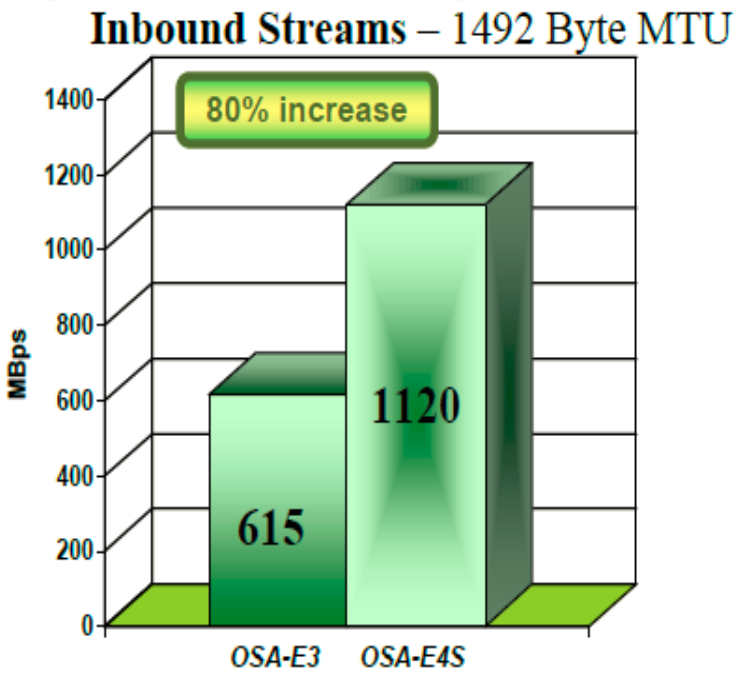

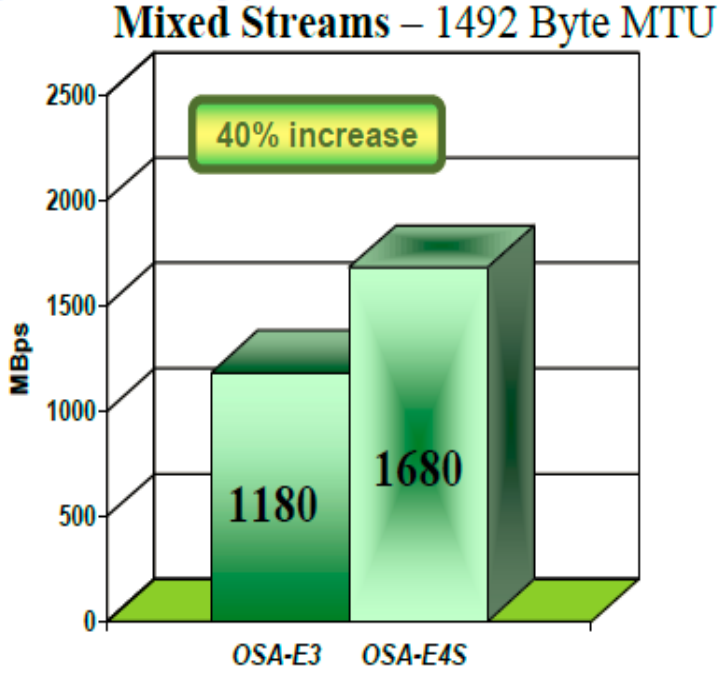

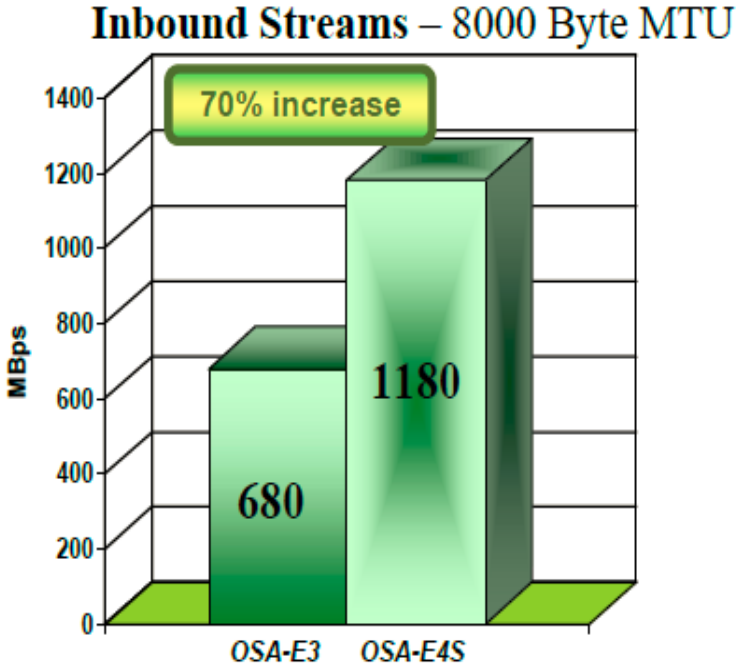

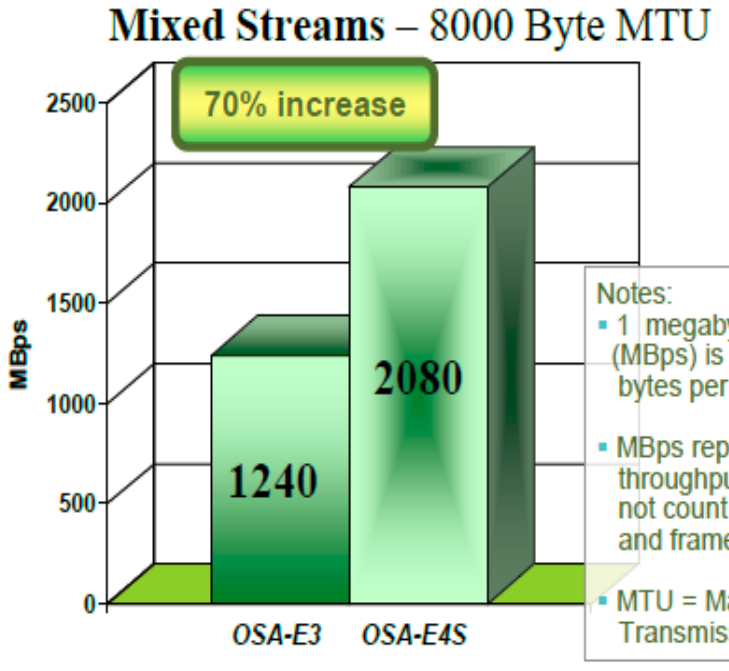

Notes:

- 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

Parallel Sysplex CFCC Level 18

- **Performance Improvements**
  - Elapsed time improvements when dynamically altering the size of a cache structure
  - DB2 conditional write to a group buffer pool (GBP) cache which allows selected entries be written around the cache to disk to reduce overhead
  - Performance improvements for coupling facility cache structures to avoid flooding the coupling facility cache with changed data and avoid excessive delays and backlogs for cast-out processing.
  - Performance throughput enhancements for parallel cache castout processing by extending the number of RCC cursors beyond 512.
  - Storage class and castout class contention avoidance by breaking up individual storage class and castout class queues to reduce storage class and castout class latch contention
- **Resiliency Improvements**
  - Enhanced capabilities to non-disruptively capture and collect extended diagnostic structure data from Coupling Facility structures that have encountered an error.
  - Verification of local cache controls for a Coupling Facility cache structure connector during registration of connection interest in a data item against lost lost cross-invalidation signals. .
- **Structure and CF Storage Sizing with CFCC level 18**
  - Moving to CFCC Level 18 may increase storage requirements
  - Use of the CF Sizer Tool is recommended: http://www.ibm.com/systems/z/cfsizer/

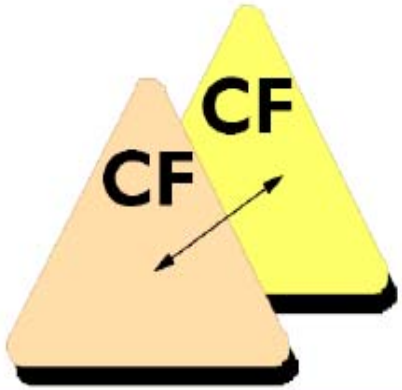

Sysplex Scalability and Monitoring Enhancements

- **Support for up to 101 Integrated Coupling Facility (ICF) processors**
  - The limit on previous System z servers was 16 ICFs
  - The maximum number of logical processors in a Coupling Facility LPAR remains at 16
- **Support for up to 64 1x InfiniBand Coupling Links**
  - The limit on previous System z servers was 48
  - Provides additional link connectivity for STP and for Parallel Sysplex configurations with members of more than on Sysplex on the same zEC12
  - Facilitates migration from ISC-3 to 1x InfiniBand Links
- **Coupling link characteristics reporting to z/OS**
  - Identifies underlying InfiniBand hardware characteristics for CIB CHPIDs to help with Sysplex monitoring and tuning
  - Enables RMF Monitor III to report additional information
    - InfiniBand Link type and protocol: 12x IFB, 12x IFB3 and 1x IFB
    - CHPID mapping to physical links – HCA IDs and port numbers
    - Calculated fiber optic link length
    - Fully functional or degraded status

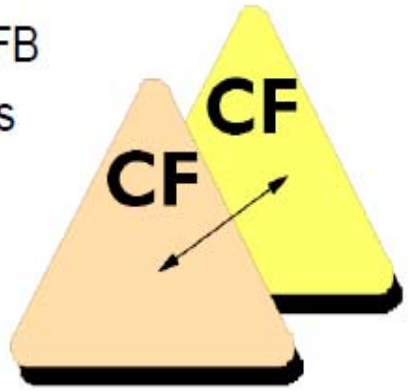

zEC12 Server Time Protocol Enhancements

- **Broadband Security Improvements for STP**
  - Authenticates NTP servers when accessed by the HMC client through a firewall
  - Authenticates NTP clients when the HMC is acting as an NTP server
  - Provides symmetric key (NTP V3-V4) and autokey (NTP V4) authentication (Autokey is not supported if Network Address Translation is used)
  - This is the highest level of NTP security available
- **Improved NTP Commands panel on HMC/SE**
  - Shows command response details
- **Telephone modem dial out to an STP time source is no longer supported**
  - All STP dial functions are still supported by broadband connectivity
  - zEC12 HMC LIC no longer supports dial modems (Fulfills the Statement of Direction in Letter 111-167, dated October 12, 2011)

Removal of Support for STP Mixed CTNs (August 28, 2012 Statement of Direction[1])

- **The zEC12 is intended to be the last high-end System z server to support connection to a Server Time Protocol (STP) Mixed Coordinated Timing Network (CTN) that includes a Sysplex Timer.**

  Beginning with the high-end System z server after zEC12, System z servers that require time synchronization to will require STP and must be connected to an STP-only CTN.

  - Enterprises should continue migration away from STP mixed CTNs to STP-only CTNs.

zEC12 1x InfiniBand Coupling Links

Multiple CHPIDs per link, 7 or 32* subchannels per CHPID (HCA3-O LR and HCA2-O LR)

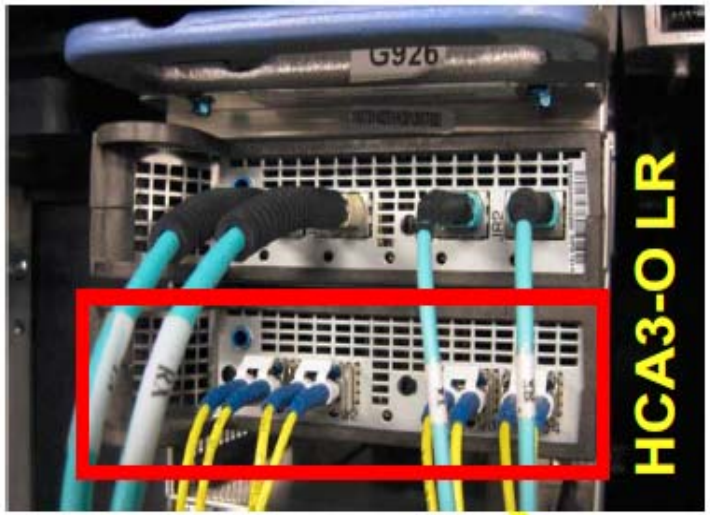

- Up to 16 CHPIDs using same physical links
  - More subchannels per physical link
  - Link sharing by different Sysplexes
- **Now** more subchannels per CHPID
  - **32 subchannels per CHPID**
  - **Option to define 7 or 32* subchannels**
  - **zEC12, z196 GA2, and z114 on HCA3-O LR or HCA2-O LR**

NEW

32 subchannels per CHPID (7* default)

Up to 16 CHPIDs per HCA3-O LR

---

512 subchannels per HCA3-O LR

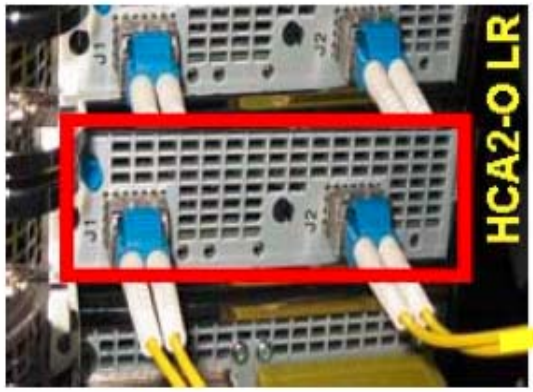

For Example:

***HCD with APAR 0A36617 fix changes the default back to 7 subchannels for CIB CHPIDs. Specify 32 when needed. (Recommended for extended distance.)**

CHPID FF
32 subchannels

CHPID FE
32 subchannels

One 1x IFB link

64 subchannels

zEC12 InfiniBand HCA3 Fanouts

Up to 16 CHPIDs – across 2 ports*

- **Exclusive to zEC12 and zEnterprise 196 and 114**

– HCA3-O fanout for 12x InfiniBand coupling links

- CHPID type – CIB
  - *Improved service times with 12x IFB3 protocol for HCA3-O to HCA3-O links*
  - Two ports per feature
  - Fiber optic cabling – 150 meters
  - Supports connectivity to HCA2-O
  - Link data rate of 6 GBps

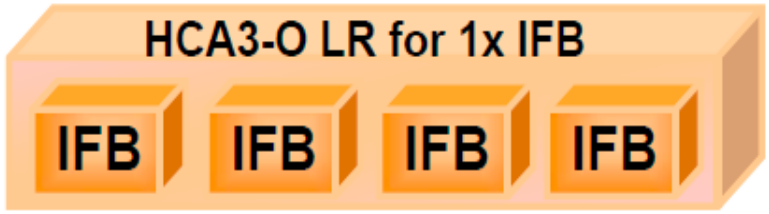

Up to 16 CHPIDs – across 4 ports*

– HCA3-O LR fanout for 1x InfiniBand coupling links

- CHPID type – CIB
  - *Four ports per feature*
  - Fiber optic cabling
    - 10 km unrepeated, 100 km repeated
  - Supports connectivity to HCA2-O LR
  - Link data rate server-to-server 5 Gbps
  - Link data rate with WDM; 2.5 or 5 Gbps

*** Performance considerations may reduce the number of CHPIDs per port.**

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

zEC12 Parallel Sysplex Coupling Connectivity