



SHARE in Boston

August 11 - 16, 2013



Introducing the IBM zEnterprise® BC12 (zBC12) and EC12 (zEC12) GA2 Hardware: I/O Subsystem, Features, and Parallel Sysplex Coupling Function

Session ID: 13706

Wednesday, August 14, 2013, 9:30 AM
Room 312, Hynes Convention Center

Speaker: Harv Emery

System z Hardware,
Advanced Technical Sales



Please complete session evaluations
online at SHARE.org/BostonEval



Permission is granted to SHARE to publish this presentation in the SHARE Proceedings.
IBM retains its right to distribute copies of this presentation to whomever it chooses.



Trademarks

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

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®

* Registered trademarks of IBM Corporation

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.

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.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

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.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and

Linux is a registered trademark of Linus Torvalds 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.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](http://www.openstack.org).

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

Windows Server and the Windows logo are trademarks of the Microsoft group of countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

* Other product and service names might be trademarks of IBM or other 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.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g. zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

zBC12 and zEC12 GA2

July 23, 2013

Announcement Hardware

Highlights



IBM zEnterprise EC12 and zBC12 – September 2013 – Dr 15

IBM zEnterprise EC12 (2827)



- Announced 08/12
- 5 models – H20, H43, H66, H89, and HA1
 - Up to 101 CPs
- Granular Offerings for up to 20 CPs
- PU (Engine) Characterization
 - CP, IFL, ICF, zAAP, zIIP, SAP, IFP
- On Demand Capabilities
 - CoD: CIU, CBU, On/Off CoD, CPE
- Memory – up to 3 TB for Server, 1 TB per LPAR
 - 32 GB Fixed HSA
- Channels
 - PCIe bus
 - Four CSSs, up to 60 LPARs
 - 3 Subchannel Sets per CSS
 - FICON Express8 and 8S
 - zHPF
 - OSA 10 GbE, GbE, 1000BASE-T
 - InfiniBand Coupling Links
 - Flash Express
 - Compression Acceleration (zEDC)
 - RDMA over CE (RoCE)
- Configurable Crypto Express4S
- Parallel Sysplex clustering
- HiperSockets – up to 32
- Enhanced Book Availability
- IBM zAware
- Unified Resource Manager
- Operating Systems
 - z/OS, z/VM, z/VSE, z/TPF, Linux on System z

IBM zEnterprise Blade Extension (2458)



- First Announced 7/10
- Model 003 for zEC12 – 08/12
- zBX Racks (up to 4) with:
 - BladeCenter Chassis
 - N + 1 components
 - Blades
 - Top of Rack Switches
 - 8 Gb FC Switches
 - Power Units
 - Advanced Management Modules
- Up to 112 Blades
 - POWER7 Blades
 - IBM System x Blades
 - IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (M/T 2462-4BX)
- Operating Systems
 - AIX 5.3 and higher
 - Linux for Select IBM x Blades
 - Microsoft Windows for x Blades
- Hypervisors
 - PowerVM Enterprise Edition
 - Integrated Hypervisor for System x

IBM zEnterprise BC12 (2828)



- Announced 07/13
- 2 models – H06 and H13
 - Up to 6 CPs, up to 13 characterized PUs
- High levels of Granularity available
 - 156 Capacity Settings
- PU (Engine) Characterization
 - CP, IFL, ICF, zAAP, zIIP, SAP, IFP
- On Demand Capabilities
 - CoD: CIU, CBU, On/Off CoD, CPE
- Memory – up to 512 GB for Server
 - 16 GB Fixed HSA
- Channels
 - PCIe bus
 - Two CSSs, up to 30 LPARs
 - 2 Subchannel Sets per CSS
 - FICON Express8 and 8S
 - zHPF
 - OSA 10 GbE, GbE, 1000BASE-T
 - InfiniBand Coupling Links
 - Flash Express
 - Compression Acceleration (zEDC)
 - RDMA over CE (RoCE)
- Configurable Crypto Express 4S
- Parallel Sysplex clustering
- HiperSockets – up to 32
- IBM zAware
- Unified Resource Manager
- Operating Systems
 - z/OS, z/VM, z/VSE, z/TPF, Linux on System z

zBC12: Extending the capabilities of the modern mainframe

zBC12

Machine Type: 2828

2 Models: H06 and H13

Twice the capacity at the same entry IBM hardware cost as the zEnterprise 114

36% boost in per core processor performance

58% more total z/OS system capacity

Up to **6** CPs

Up to **13** IFLs

62% more total Linux capacity with **27%** price performance improvement

2x available memory

156 available capacity settings

Fully Upgradeable from the z10 BC and z114; and to the zEC12 H20

New technology with unprecedented performance

- New 4.2GHz core with improved cache designs and new hardware function designed to boost software performance
- Increased core counts and memory and SSI scale for additional flexibility, growth and economies of scale consolidation
- Increased granularity for right-sizing system to business needs
- Built to support future data center design, modernization and energy efficiencies requirements




Innovative Features bolster platform differentiators

- Storage Class Memory and integrated system health analytics to improve availability: Flash Express and IBM zAware
- Enhanced high speed, low latency networking fabric for z/OS to z/OS communications: 10GbE RoCE Express
- New compression acceleration adapter to improve system performance and reduce CPU and storage costs : zEDC Express
- Enhanced security with extended EP11 and enhanced IBM CCA support
- Hybrid-enabled to optimize workload integration and to provide System z governance: zBX and Unified Resource Manager

Improved Platform Economics

- Modular two drawer design for low cost of entry
- Improved price performance across the stack
- Second generation upgradeability for investment protection

IBM zEnterprise family – will be withdrawn from Marketing, June 30, 2014

IBM zEnterprise 196 (2817)	IBM zEnterprise BladeCenter Extension (2458)	IBM zEnterprise 114 (2818)
		
<ul style="list-style-type: none"> • Announced 7/10 – Server with up to 96 PU cores • 5 models – Up to 80-way • Granular Offerings for up to 15 CPs • PU (Engine) Characterization <ul style="list-style-type: none"> – CP, SAP, IFL, ICF, zAAP, zIIP • On Demand Capabilities <ul style="list-style-type: none"> – CoD, CIU, CBU, On/Off CoD, CPE • Memory – up to 3 TB for Server and up to 1 TB per LPAR <ul style="list-style-type: none"> – 16 GB Fixed HSA • Channels <ul style="list-style-type: none"> – PCIe bus – Four LCSSs – 3 Subchannel Sets – MIDAW facility – Up to 240 ESCON channels – Up to 288 FICON channels – FICON Express8 and 8S – zHPF – OSA 10 GbE, GbE, 1000BASE-T – InfiniBand Coupling Links • Configurable Crypto Express3 • Parallel Sysplex clustering • HiperSockets – up to 32 • Up to 60 logical partitions • Enhanced Availability • Unified Resource Manager • Operating Systems <ul style="list-style-type: none"> – z/OS, z/VM, z/VSE, z/TPF, Linux on System z 	<ul style="list-style-type: none"> • Announced 7/10 • Model 002 for z196 or z114 • zBX Racks with: <ul style="list-style-type: none"> – BladeCenter Chassis – N + 1 components – Blades – Top of Rack Switches – 8 Gb FC Switches – Power Units – Advance Management Modules • Up to 112 Blades <ul style="list-style-type: none"> – POWER7 Blades – IBM System x Blades – IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (M/T 2462-4BX) • Operating Systems <ul style="list-style-type: none"> – AIX 5.3 and higher – Linux for Select IBM x Blades – Microsoft Windows for x Blades • Hypervisors <ul style="list-style-type: none"> – PowerVM Enterprise Edition – Integrated Hypervisor for System x 	<ul style="list-style-type: none"> • Announced 07/11 – Server with up to 18 PU cores • 2 models – M05 and M10 <ul style="list-style-type: none"> – Up to 5 CPs • High levels of Granularity available <ul style="list-style-type: none"> – 130 Capacity Indicators • PU (Engine) Characterization <ul style="list-style-type: none"> – CP, SAP, IFL, ICF, zAAP, zIIP • On Demand Capabilities <ul style="list-style-type: none"> – CoD, CIU, CBU, On/Off CoD, CPE • Memory – up to 256 GB for Server <ul style="list-style-type: none"> – 8 GB Fixed HSA • Channels <ul style="list-style-type: none"> – PCIe bus – Two LCSSs – 2 Subchannel Sets – MIDAW facility – Up to 240 ESCON channels – Up to 128 FICON channels – FICON Express8 and 8S – zHPF – OSA 10 GbE, GbE, 1000BASE-T – InfiniBand Coupling Links • Configurable Crypto Express3 • Parallel Sysplex clustering • HiperSockets – up to 32 • Up to 30 logical partitions • Unified Resource Manager • Operating Systems <ul style="list-style-type: none"> – z/OS, z/VM, z/VSE, TPF, z/TPF, Linux on System z

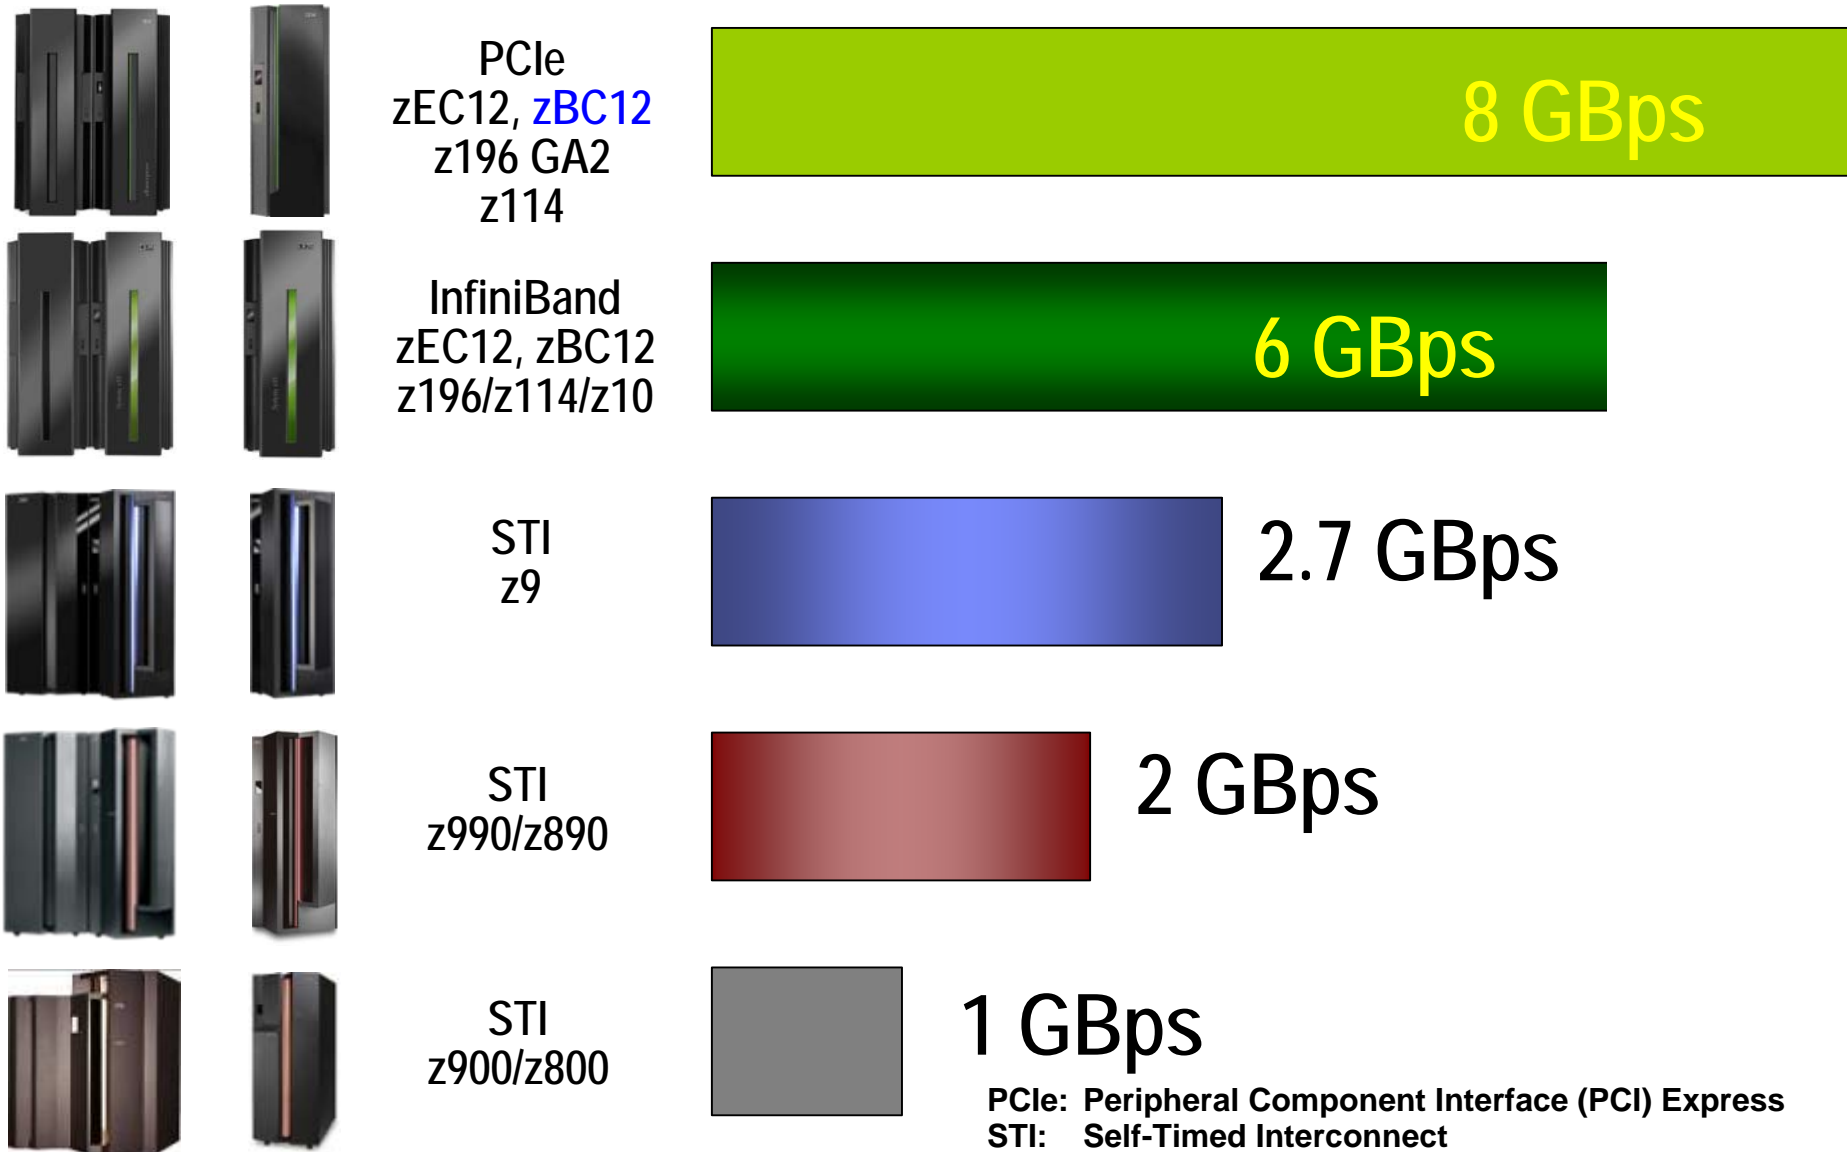
IBM zEnterprise Hardware Withdrawals from Marketing

- **Effective September 20, 2013, IBM will withdraw from the following features on IBM zEnterprise EC12 (zEC12):**
 - OSA-Express4S GbE LX (#0404) – Replaced by OSA-Express5S GbE LX (#0413)
 - OSA-Express4S GbE SX (#0405) – Replaced by OSA-Express5S GbE SX (#0414)
 - OSA-Express4S 10 GbE LR (#0406) – Replaced by OSA-Express5S 10 GbE LR (#0415)
 - OSA-Express4S 10 GbE SR (#0407) – Replaced by OSA-Express5S 10 GbE SR (#0416)
 - OSA-Express4S 1000BASE-T (#0408) – Replaced by OSA-Express5S 1000BASE-T (#0417)
 - HMC (#0091) – Replaced by HMC (#0092)
 - TKE workstation (#0841) – Replaced by TKE workstation (#0842)
 - TKE 7.2 LIC (#0850) - Replaced by TKE 7.3 LIC (#0872)
- **Effective June 30, 2014, IBM will withdraw from marketing:**
 - All models of the [IBM zEnterprise 196 \(z196\)](#) and all upgrades to the z196 from the IBM System z9 EC (z9 EC), or IBM System z10 EC (z10 EC).
 - All models of the [IBM zEnterprise 114 \(z114\)](#) and all upgrades to the z114 from the IBM System z9 BC (z9 BC), or IBM System z10 BC (z10 BC).
 - The [zEnterprise BladeCenter Extension \(zBX\) Model 002](#). Model conversions and hardware MES features applied to an existing z196 or z114 server or to an existing zBX Model 002.
 - [Field installed features and conversions that are delivered solely through a modification to the machine's Licensed Internal Code \(LIC\) will continue to be available until June 30, 2015.](#) After June 30, 2015, features and conversions that are delivered solely through a modification to the LIC will be withdrawn.
 - [The Capacity on Demand offerings that are configured prior to withdrawal are usable until the offering expiration date or termination date, as applicable.](#)
- **Adding Power and System x Blade Enablements for the zBX Model 002.**
 - **For the following features and machine types, MES orders may be accepted for after June 30, 2014, and before June 30, 2015, that do not require a zBX Configured Chassis Assembly (#0603) to complete the MES order. These are LIC CC only enablements.**
 - M/T 2817 Feature Codes #0041, #0042, #0044, #0045, #0046, #0052
 - M/T 2828 Feature Codes #0041, #0042, #0044, #0045, #0046, #0052
 - M/T 2458 Model 002 Feature Codes #0612, #0613
 - After June 30, 2014, this enables an MES order to be placed that will allow additional Power and System x Blade Servers to be installed in a zBX that had server slots available in their existing chassis assemblies. No zBX Model 002 orders will be accepted that include zBX Chassis Assembly (#0603) after June 30, 2014.
- **On or after the effective dates, these can no longer be ordered directly from IBM.**

I/O Subsystem Introduction



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



zEC12 and zBC12 support two I/O infrastructures for I/O feature cards

▪ PCI Express Generation 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

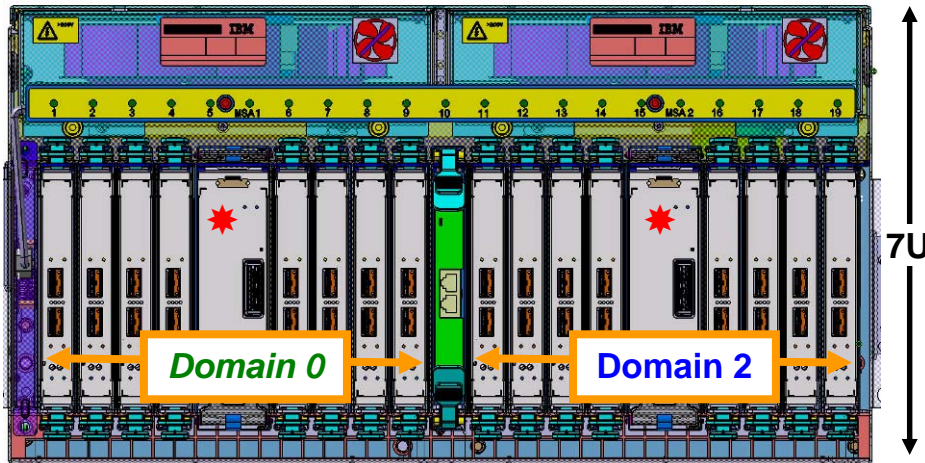
Lower energy consumption

▪ 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

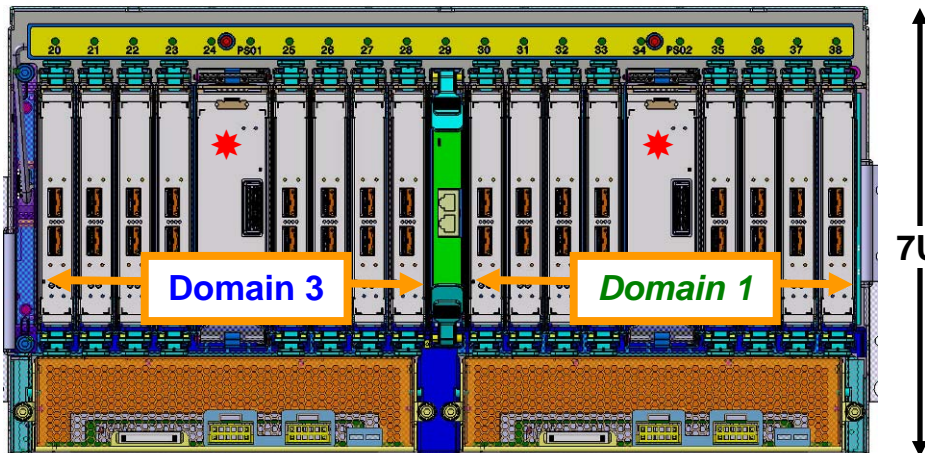
PCIe 32 I/O slot drawer (New build or Carry Forward)

Front



- Supports only PCIe I/O cards
 - zBC12: Up to two 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 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)

Rear



Introducing “Native” PCIe Technology



PCIe I/O Features – Introducing “Native” PCIe Flash Express, zEDC Express and 10GbE RoCE Express

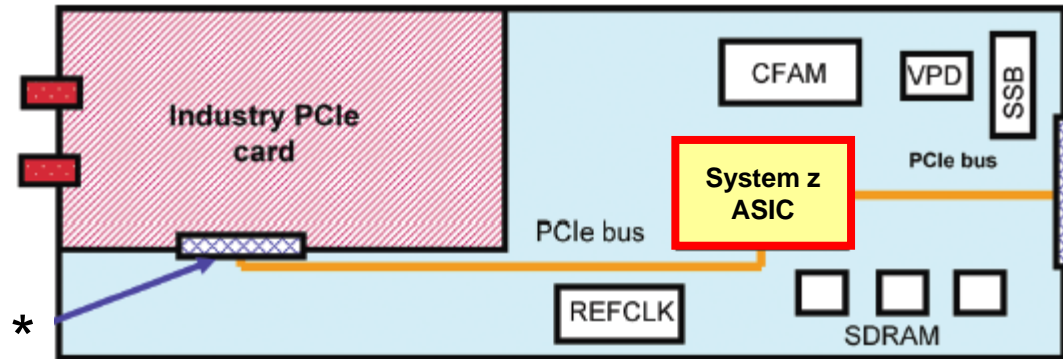
Traditional System z I/O PCIe Feature

- One System z 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 Express4S is not defined as a CHPID**
- Virtualization and support by Channel Subsystem LIC on System Assist Processors (SAPs)

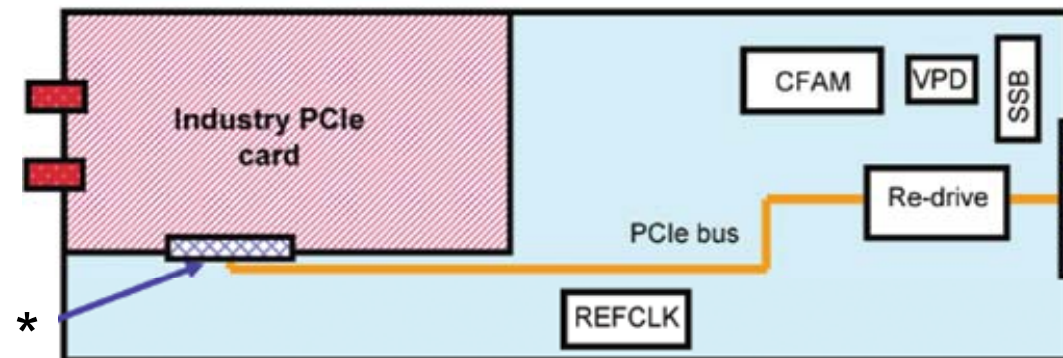
Native PCIe Features



- System z ASIC role moved to the the new System z I/O Controller (zIOC) in the PCIe I/O fanout
- 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 System z I/O PCIe Features: FICON Express8S, OSA-Express5S and 4S, **Crypto Express4S**



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) 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 **00 – FF**
 - **NOT** assigned to a Channel Subsystem so ANY LPAR can be defined to a 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. For Example, a zEDC Express feature supports up to 15 partitions.
 - 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.

Channel Subsystem and Subchannel Sets on zBC12

Two Logical Channel Subsystems

- 15 LPARs each, 30 total
- Two Subchannel Sets each
 - SS-0: 63.75 k subchannels primarily for base devices
 - SS-1: 64 k subchannels primarily for parallel access and other special usage

Channels

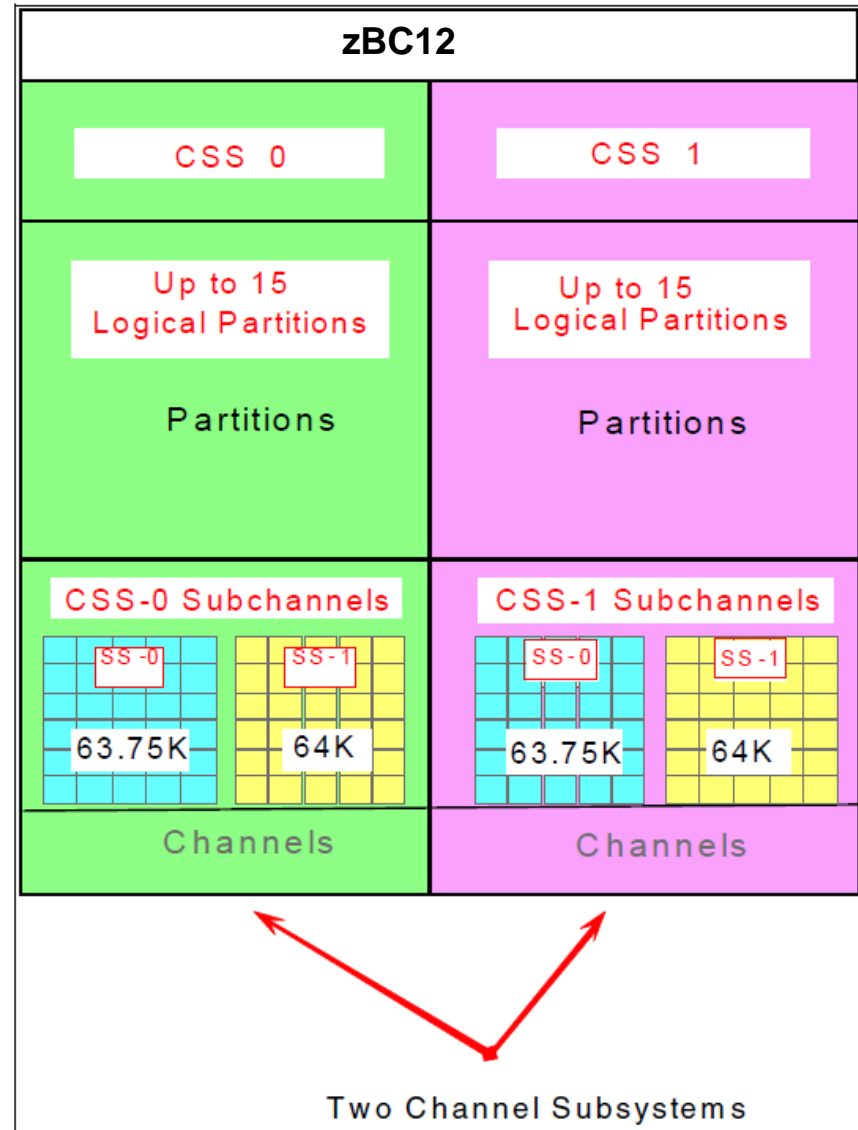
- Up to 256 per channel subsystem (CSS) with channel path identifier (CHPID) numbers in the range 00 – FF
- Can be defined as:
 - **Dedicated** to support one LPAR
 - **Reconfigurable** to support multiple LPARs, one at a time
 - **Shared** by multiple LPARs in one logical channel subsystem at the same time
 - **Spanned** to support LPARs in both logical channel subsystems

Functions (NEW)

- Up to 256 FUNCTIONS per server, independent of CSS, with Function Identifier (FID=) 00 - FF
- Access to 10GbE RoCE Express and zEDC Express
- Can be defined as:
 - **Dedicated** to support one LPAR
 - **Reconfigurable** to support multiple LPARs, one at a time

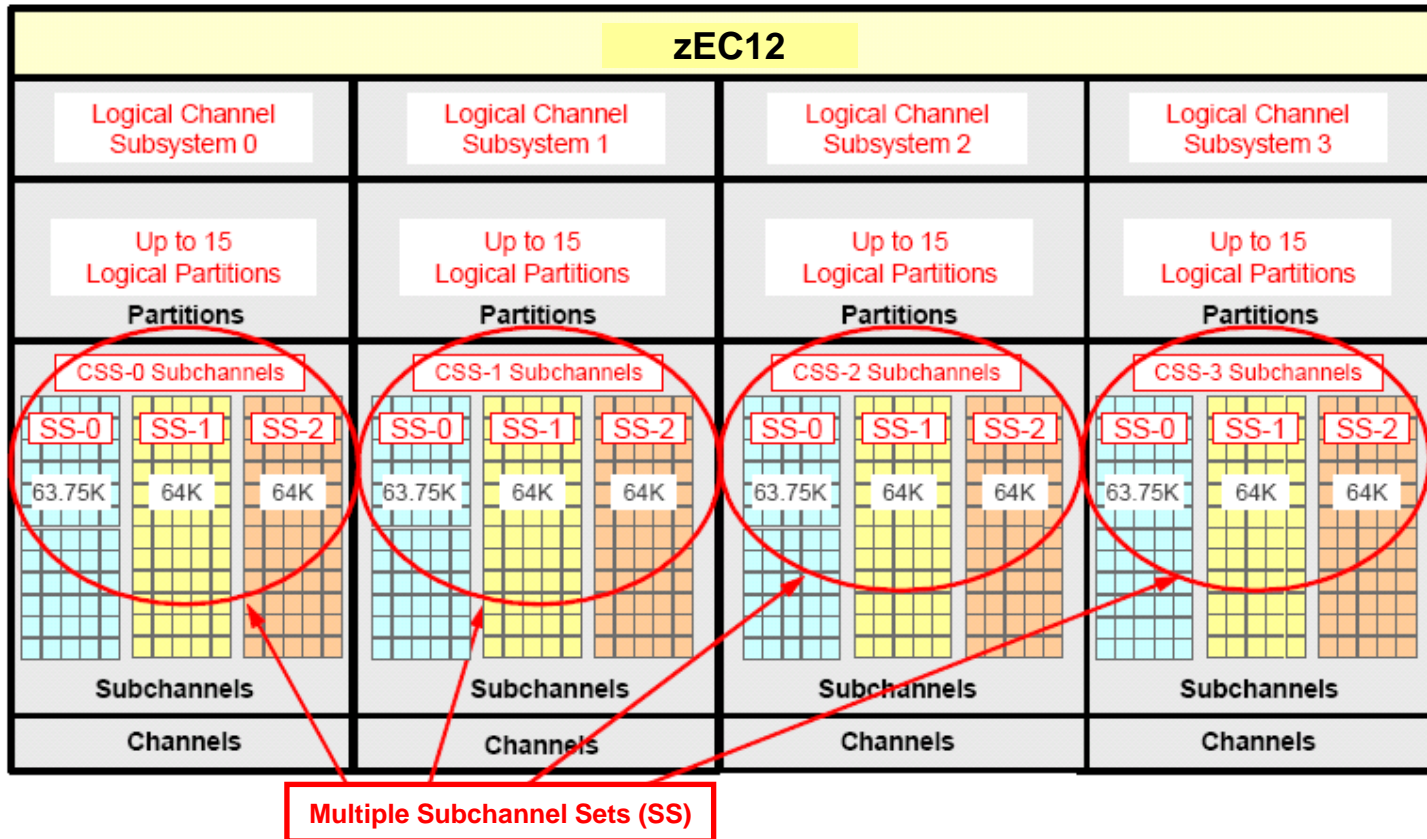


Note: Future System z server generations are likely to support more than 256 FUNCTIONS as the number of native PCIe features increases



Multiple subchannel sets on zEC12

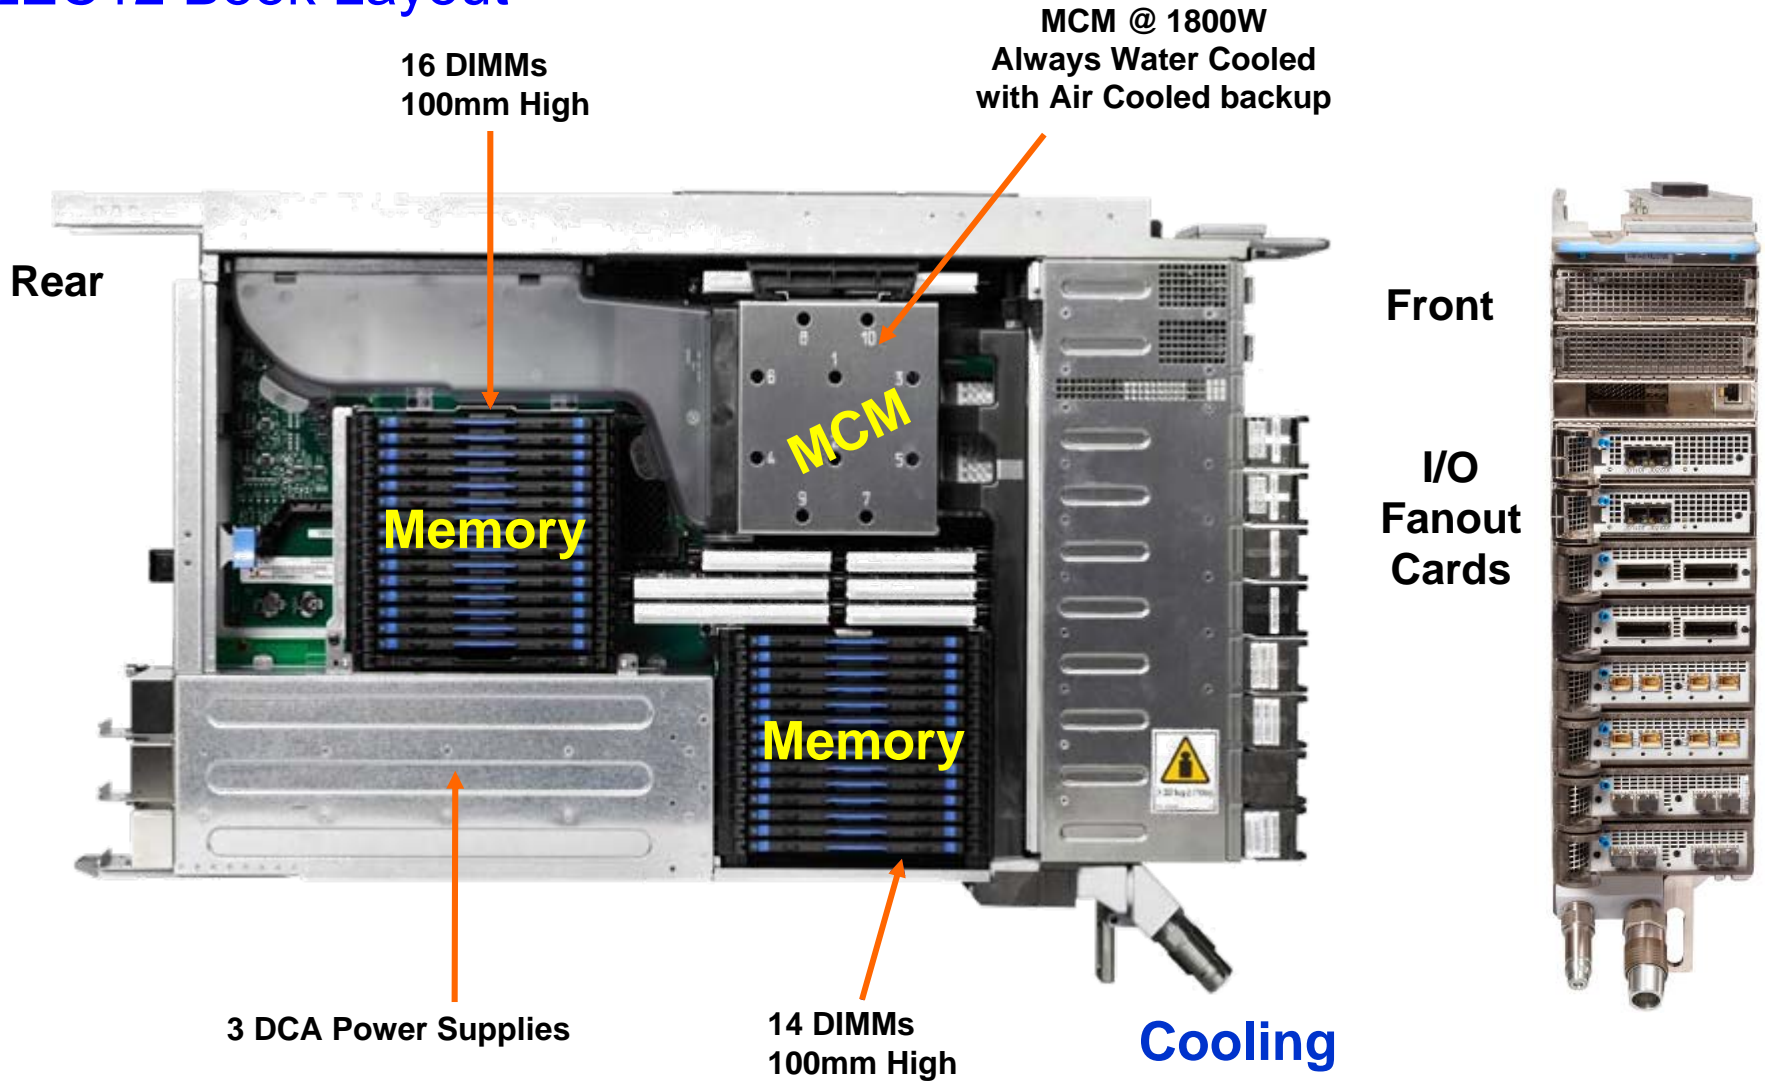
- Four Logical Channel Subsystems (CSS) each supporting up to 256 channels and three subchannel sets (SS)
 - Up to 63.75k base IODEVICEs in SS-0 and 128k alias IODEVICEs (SS-1 and SS-2) per CSS
- Up to 15 LPARs per CSS, up to 60 total
 - Each LPAR has access to the channels and IODEVICEs in its CSS
- **Up to 256 FUNCTION definitions per zEC12, independent of CSS, for native PCIe**



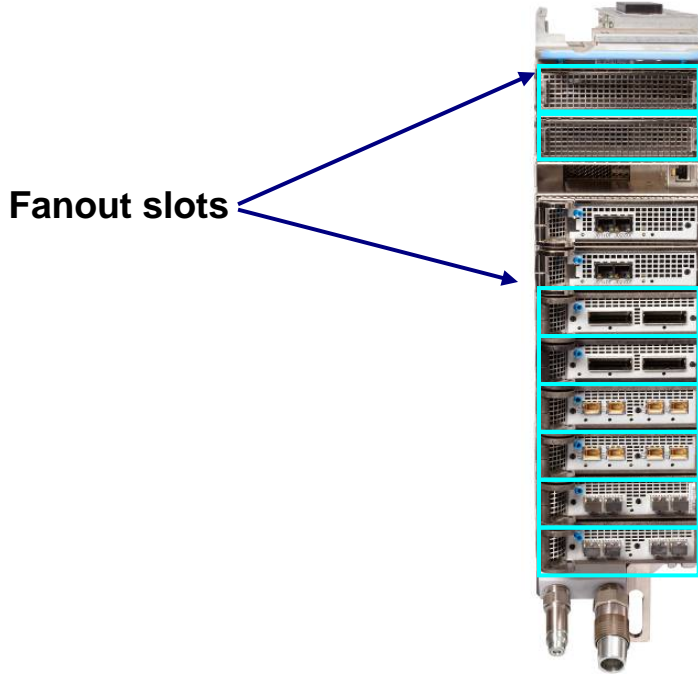
I/O Subsystem Structure



zEC12 Book Layout

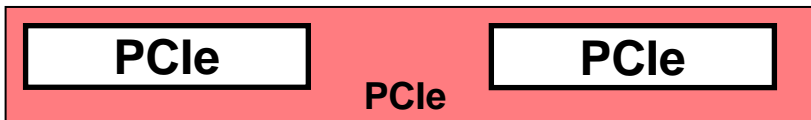


zEC12 Connectivity for I/O and Coupling

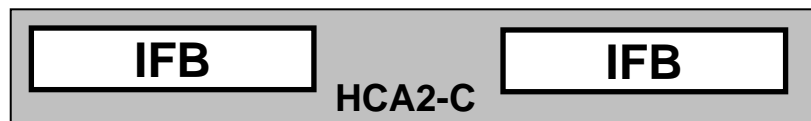


- **Up to 8 fanout cards per zEC12 book**
 - H20 (1 book) – up to 8
 - H43 (2 books) – up to 16
 - H66 (3 books) – up to 20
 - H89 and HA1 (4 books) – up to 24

- **I/O fanouts compete for fanout slots with the the InfiniBand HCA fanouts that support coupling:**
 - HCA2-O 12x two InfiniBand DDR links (CF only)
 - HCA2-O LR two 1x InfiniBand DDR links (CF only)
 - **HCA3-O two 12x InfiniBand DDR links**
 - **HCA3-O LR four 1x InfiniBand DDR links**
 -



- **PCIe fanout – PCIe I/O Interconnect links**
 Supports two copper cable PCIe 8 GBps interconnects to two 8-card PCIe I/O domain multiplexers.
Always plugged in pairs for redundancy.



- **HCA2-C fanout – InfiniBand I/O Interconnect**
 Supports two copper cable 12x InfiniBand DDR 6 GBps interconnects to two 4-card I/O domain multiplexers. **(Carry forward)**
Always plugged in pairs for redundancy.

zEC12 Book and I/O fanout plugging

H20

H43

- **I/O fanouts**
 - H20 up to 8
 - H43 up to 16
 - H66 up to 20
 - H89/HA1 up to 24

- **Cooling Design**
 - Fanout locations D1 and D2 are NOT used in the 2nd or 3rd Book for Model H66

- Fanout locations D1 and D2 are NOT used in any Book for Models H89 and HA1

Front views

H66

H89 and HA1

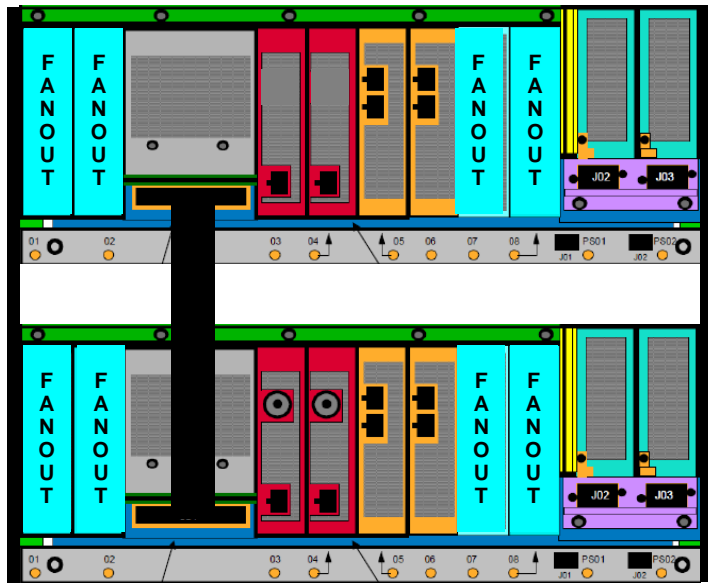
ECF	OSC	ECF	OSC
filler	D1 I/O	filler	filler
	D2 I/O		
	D3 FSP		
	D4 FSP		
	D5 I/O		
	D6 I/O		
	D7 I/O		
	D8 I/O		
	D9 I/O		
	DA I/O		

ECF	OSC	ECF	OSC
filler	D1 I/O	filler	filler
	D2 I/O		
	D3 FSP		
	D4 FSP		
	D5 I/O		
	D6 I/O		
	D7 I/O		
	D8 I/O		
	D9 I/O		
	DA I/O		

ECF	OSC	ECF	OSC		
filler	D1 I/O	filler	filler		
	D2 I/O				
	D3 FSP			D3 FSP	D3 FSP
	D4 FSP			D4 FSP	D4 FSP
	D5 I/O			D5 I/O	D5 I/O
	D6 I/O			D6 I/O	D6 I/O
	D7 I/O			D7 I/O	D7 I/O
	D8 I/O			D8 I/O	D8 I/O
	D9 I/O			D9 I/O	D9 I/O
	DA I/O			DA I/O	DA I/O

ECF	OSC	ECF	OSC		
filler	D1 I/O	filler	filler		
	D2 I/O				
	D3 FSP			D3 FSP	D3 FSP
	D4 FSP			D4 FSP	D4 FSP
	D5 I/O			D5 I/O	D5 I/O
	D6 I/O			D6 I/O	D6 I/O
	D7 I/O			D7 I/O	D7 I/O
	D8 I/O			D8 I/O	D8 I/O
	D9 I/O			D9 I/O	D9 I/O
	DA I/O			DA I/O	DA I/O

zBC12 Connectivity for I/O and Coupling

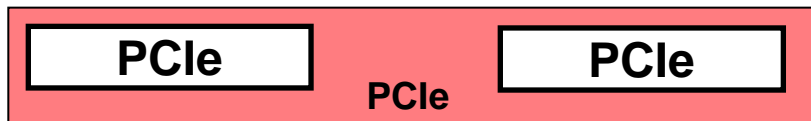


Drawer 2
H13 only

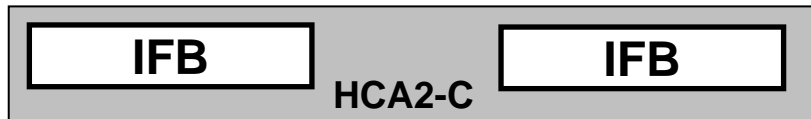
Drawer 1
H06 and H13

- **Up to 4 fanouts per zBC12 processor drawer**
 - H06 (one CEC drawer) – up to 4 fanouts
 - H13 (two CEC drawers) – up to 8 fanouts

- **I/O fanouts compete for fanout slots with the InfiniBand HCA fanouts that support coupling:**
 - HCA2-O 12x two InfiniBand DDR links (CF only)
 - HCA2-O LR two 1x InfiniBand DDR links (CF only)
 - **HCA3-O two 12x InfiniBand DDR links**
 - **HCA3-O LR four 1x InfiniBand DDR links**

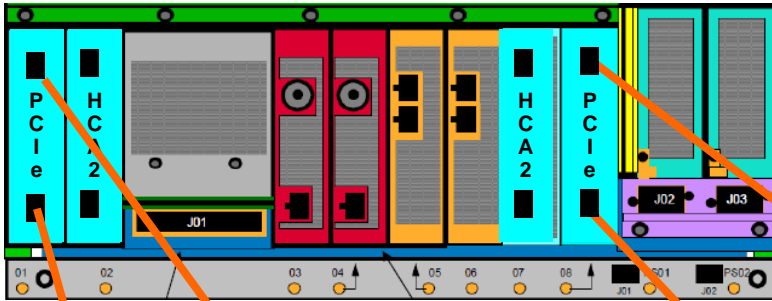


- **PCIe fanout – PCIe I/O Interconnect links**
Supports two PCIe 8 GBps interconnects on copper cables to two 8-card PCIe I/O domain switches.
Always plugged in pairs for redundancy.



- **HCA2-C fanout – InfiniBand I/O Interconnect (CF only)**
Supports two 12x InfiniBand DDR 6 GBps interconnects on copper cables to two 4-card I/O domain multiplexers.
Always plugged in pairs for redundancy.

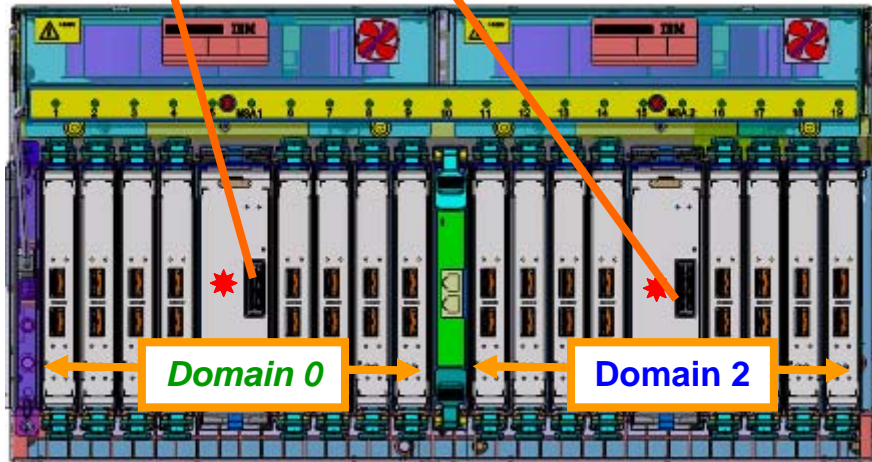
zBC12 Redundant I/O Interconnect Example – One PCIe I/O drawer



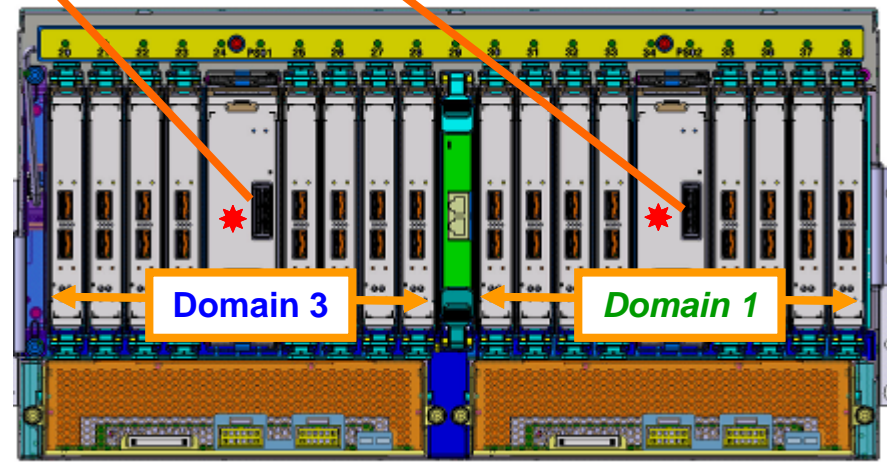
- Different PCIe Fanouts Support Domain Pairs:
 - 0 and 1
 - 2 and 3
- Normal operation each PCIe interconnect in a pair supports the eight I/O cards in its domain.
- Backup operation: One PCIe interconnect supports all 16 I/O cards in the domain pair.

PCIe switch cards (*)

Front



Rear



Supported I/O Features



zBC12 I/O Features supported



Supported features

PCIe I/O drawer (zBC12 does not offer Plan Ahead for PCIe I/O Drawers)

- The drawer itself does NOT carry forward but the cards do
- **Cards that Carry Forward**
 - OSA-Express4S 1 GbE LX and SX, OSA-Express4S 10 GbE LR and SR
 - FICON Express8S 10 km LX and SX
- **Cards New Build**
 - FICON Express8S 10 km LX and SX
 - Crypto Express4S (1 coprocessor)
 - **OSA-Express5S GbE LX and SX, OSA-Express5S 10 GbE LR and SR, OSA-Express5S 1000BASE-T**
 - IBM Flash Express
 - **10 GbE RoCE Express**
 - **zEDC Express**

32 slot PCIe I/O drawer



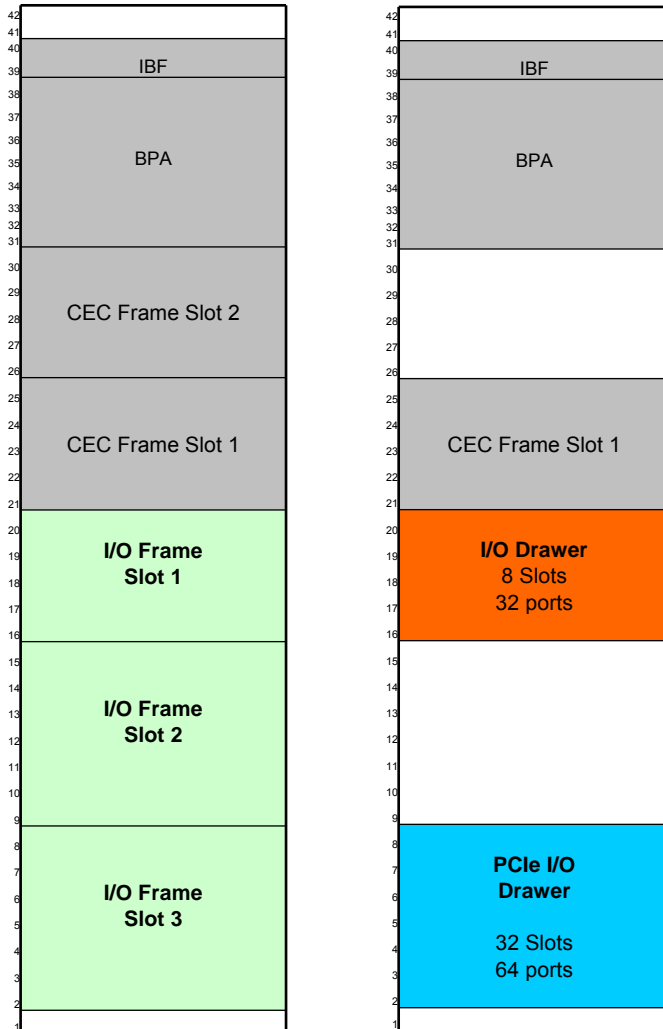
I/O drawer (Carry forward only. No MES adds.)

- The drawer itself can carry forward. It cannot be ordered on new build
- One I/O Drawer can be carried forward, two I/O drawers requires **RPQ 8P2733**
- **Cards that Carry Forward**
 - **Not Supported:** ESCON, older FICON, FICON Express4 LX 4 km (4-port or 2-port), OSA-Express2, PSC
 - OSA-Express3 Gigabit LX and SX, OSA-Express3 1000BASE-T, OSA-Express3-2P 1000BASE-T, OSA-Express3 10 Gigabit LR and SR, OSA-Express3-2P Gigabit SX
 - FICON Express8 10KM LX, FICON Express8 SX, FICON Express4 10KM LX (4-port only), FICON Express4 SX, **FICON Express4-2C SX**
 - Crypto Express3, **Crypto Express3-1P**
 - ISC3

8 slot I/O drawer



zBC12Frame Layout for I/O



- An **I/O frame slot** is a physical location in the frame for a legacy I/O drawer or PCIe I/O drawer
 - I/O Frame slot 1 = 5u
Supports 1 I/O drawer ONLY
 - I/O Frame slot 2 = 7u
Supports 1 I/O drawer or 1 PCIe I/O drawer
 - I/O Frame slot 3 = 7u
Supports 1 PCIe I/O drawer
- **PCIe I/O drawer**
 - Frame slots 2 and 3
 - 32 two port I/O slots = 64 ports
 - 2 maximum (1 with 2 I/O drawers - RPQ **8P2733**)
- **I/O drawer**
 - Frame slots 1 and 2
 - 8 four port I/O slots = 32 ports
 - 1 maximum (2 with RPQ **8P2733**)

zBC12 Carry Forward (Field Upgrade) Rules for I/O Features (All z114 PCIe I/O Features Can be Carried Forward)



Non-PCIe Feature Cards Carried Forward	8-slot I/O Drawers Required	Maximum PCIe Drawers/ Slots (CF or Add)
0	0	2/64
1 to 8	1	2/64
9 to 16 (RPQ 8P2733)	2	1/32
17 or more	Not Supported!	

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

zEC12 GA2 “New Build” I/O and MES Features Supported

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



New Build Features

■ Features – PCIe I/O drawer

- FICON Express8S (SX and LX)
- **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** (*1 supported SR port, NO CHPIDs*)
- **zEDC Express**
- **Crypto Express4S**
- **Flash Express**

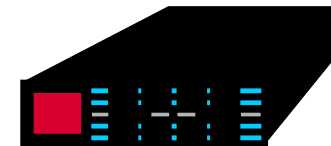
PCIe I/O drawer



32 I/O slots

■ Features – I/O drawer (No MES adds)

- ISC-3 (**RPQ 8P2602**)
Available ONLY on zEC12 Models H66, H89 or HA1
when 16 InfiniBand HCAs are included in the configuration.



I/O drawer

8 I/O slots

■ InfiniBand Coupling Features (Fanouts)

- HCA3-O 12x InfiniBand
- HCA3-O LR 1x InfiniBand

zEC12 GA2 “Carry Forward” I/O Features Supported

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

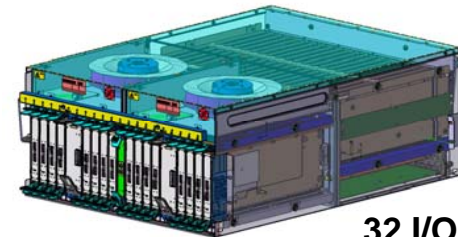


Carry Forward Features

■ Features – PCIe I/O drawer

- All PCIe features supported at z196 GA2
- Any number may be carried forward
FICON Express8S, OSA-Express4S

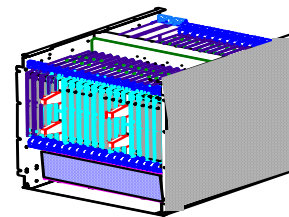
PCIe I/O drawer



32 I/O slots

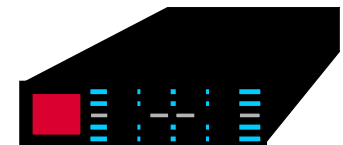
■ Features – I/O cage and I/O drawer (No MES adds)

- **Not Supported:** ESCON, older FICON, FICON Express4 LX 4 km, OSA-Express2, PSC
- **No more than 44 of these features may be carried forward:**
Crypto Express3, FICON Express8
FICON Express4 10 km LX and SX
OSA-Express3
ISC-3 (Mother + Daughters)



28 slot I/O cage

I/O drawer

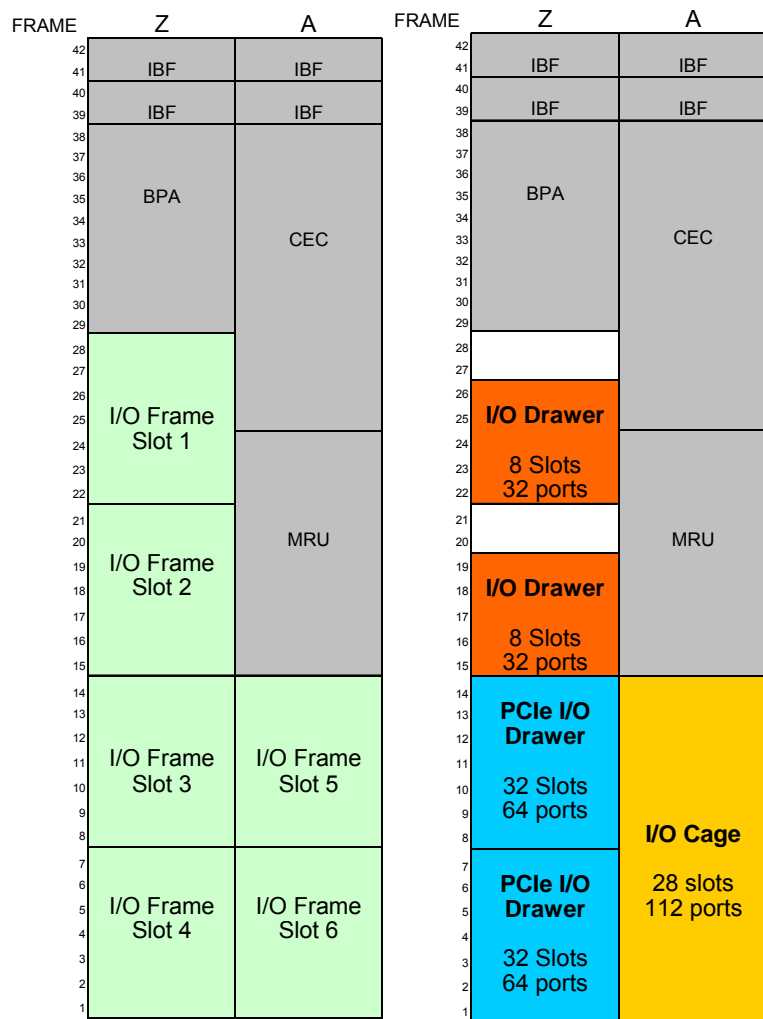


8 I/O slots

■ InfiniBand Coupling Features (Fanouts)

- All: HCA2-O 12x, HCA2-O LR 1x, HCA3-O 12x and HCA3-O 1x LR

zEC12 Frame Layout for Carry Forward I/O – Air Cooled*



- An **I/O frame slot** is a physical location in the A or Z frame for an I/O cage, I/O drawer or PCIe I/O drawer to be inserted = 7u
- **PCIe I/O drawer** uses 1 I/O frame slot = 7u
 - 32 two port I/O slots = 64 ports
 - 5 drawers maximum
- **I/O cage** uses 2 I/O frame slots = 14u
 - 28 four port I/O slots = 112 ports
 - 1 cage carry forward only maximum in I/O frame slots 5+6 only
- **I/O drawer** uses 0.7 frame slot = 5u
 - 8 four port I/O slots = 32 ports
 - Requires 2u of free space for future upgrade to the PCIe I/O drawer
 - 2 drawers carry forward only maximum in I/O frame slots 1 and 2 only

* Locations differ if water cooled; but the number of I/O frame slots is identical.

zEC12 Carry Forward (Field Upgrade) Rules for I/O Features (All PCIe I/O Features Can be Carried Forward)

Non-PCIe Feature Cards Carried Forward	8-slot I/O Drawers Required	28-slot I/O Cages Required	Maximum PCIe Drawers/ Slots (CF or Add)
0	0	0	5/160
1 to 8	1	0	5/160
9 to 16	2	0	4/128
17 to 28	0	1	4/128
29 to 36	1	1	3/96
37 to 44	2	1	2/64
45 or more	Not Supported!		

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

Note: Maximum I/O configurations require a zEC12 H43 (2 books) even without IFB links.

Supported I/O features to consider removing on an MES upgrade

I/O Feature Description	Reason to Consider Removal
FICON Express4	<p>*SOD: Not supported on future System z servers Lower performance than FICON Express8S or FICON Express8 Note: Carry forward FICON Express4 if 1 Gbps point-to-point connectivity is required</p>
ISC-3 Coupling Links	<p>*SOD: Not supported on future System z servers Less responsive and less bandwidth than InfiniBand 1x or 12x coupling links Does not support for multiple coupling CHPIDs per link</p>
OSA-Express3	<p>*SOD: Not supported on future System z servers 10GbE – Lower performance than 10 GbE OSA-Express4S or OSA-Express5S</p>
Crypto-Express3	<p>*SOD: Not supported on future System z servers Does not support the Public-Key Cryptography Standard (PKCS) EP11 mode</p>
FICON Express8	<p>Fewer I/O operations per second and lower bandwidth than FICON Express8S</p>
<p>Note: If any of the old technology I/O features above remains on a zBC12, two or more old technology I/O fanouts will be required. The fanout slots might be better used to support an additional PCIe I/O drawer or PSIFB coupling.</p>	
HCA2-O 12x PSIFB HCA2-O LR 1x PSIFB	<p>*SOD: Not supported on future System z servers 12x PSIFB – No support for the improved IFB-3 protocol supported by HCA3-O 1x PSIFB – Half the number of ports (only 2) compared to HCA3-O LR</p>

*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.

What about RPQ 8P2534? Not available on zEC12 or zBC12

- **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 or zBC12 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 EC or z196 to zEC12 or z10 BC or z114 to zBC12 because only those machines can have 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.”

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 System z Compression Call instruction (System z 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**

- **Exploitation and Compatibility**
 - Exclusive to zEC12 GA2 and zBC12
 - z/OS Support Planned:
 - z/OS V2.1 – Hardware exploitation for SMF log data in September 2013, for BSAM and QSAM SOD* for 1Q2014, DFSMSdss and DFSMSHsm SOD* for 3Q2014, and in a future update of IBM SDK for z/OS Java
 - z/OS V1.13 and V1.12 - Software support only, no hardware compression/decompression acceleration support
 - z/VM SOD* for guest support
 - Authorized APIs for ISV use are planned



zEDC Express
FC # 0420

*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 **



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 **



zEDC Express



z/VM 6.3 support for guests***



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.

10 GbE RoCE Express

For more detail on RoCE exploitation, be sure to get the handouts from
Gus Kassimis's Tuesday Sessions: 13627 and 13628



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 ****



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



z/OS V2.1 SMC-R



z/VM 6.3 support for guests*



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.

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**
 - One 10 GbE SFP+ port enabled per feature
 - Each feature must be dedicated to one LPAR
 - A switched connection requires an enterprise-class 10 GbE switch
 - SR Optics, *Global Pause* enabled & *Priority Flow Control (PFC)* disabled
 - Point-to-point connection is supported
 - 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**
 - Exclusive to zEC12 GA2 and zBC12
 - z/OS V2.1 Exploitation
 - SOD* - Future update of IBM SDK for z/OS Java
 - SOD* - Future z/VM support for guests to exploit the 10GbE RoCE Express
 - Linux on System z – IBM is working with Linux distribution partners to include support in future releases



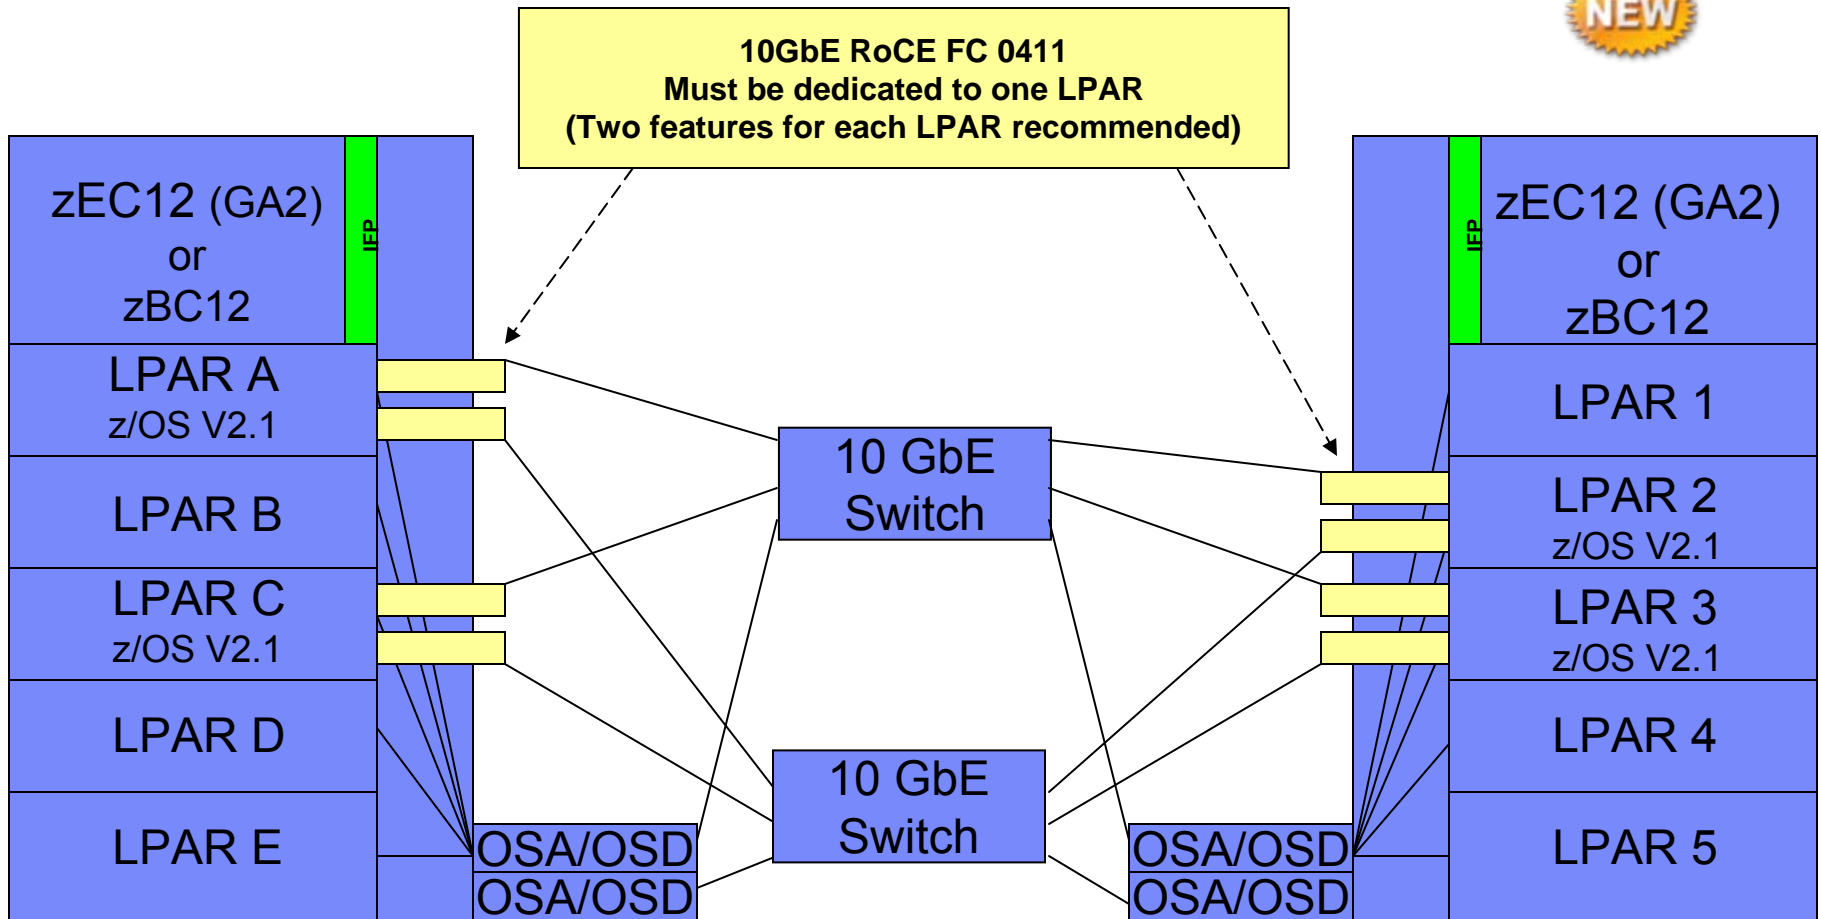
FC 0411 10GbE RoCE Express



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.

10GbE RoCE Express Sample Configuration



- This configuration allows redundant SMC-R connectivity among LPAR A, LPAR C, LPAR 2, and LPAR 3
- 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

Flash Express



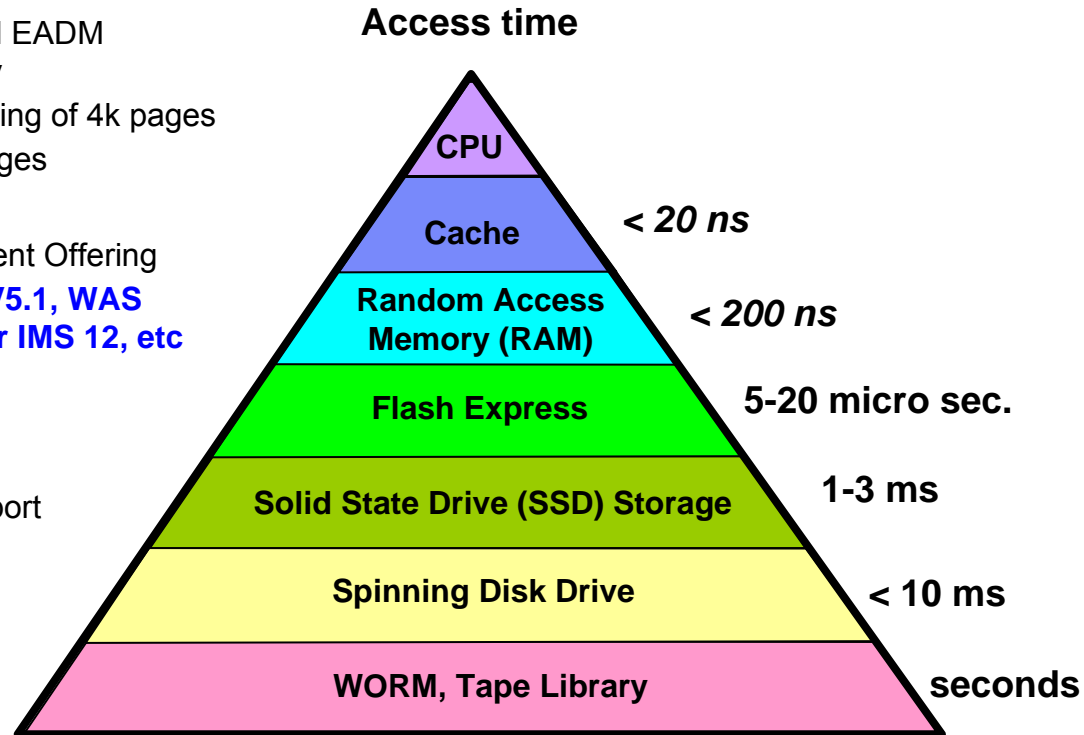
Why Flash Express on zBC12 and zEC12?

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 System z architected EADM (Extended Asynchronous Data Mover) Facility
- Designed to enable extremely responsive paging of 4k pages
- Designed to enable 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, SoD* for IMS 12, etc
- Linux on System z
 - SLES 11 SP3 and RHEL 6.4
- **Statement of Direction***: CFCC Level 19 for WebSphere MQ Shared Queue overflow support



*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 cuts away at availability lapses

19%

Reduction in total dump time for 36 GB standalone dump

10x

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

Up to **28%**

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



Innovative PCI-e storage-class memory
Enables the use of large 1 MB pageable pages

Typical Client Use Cases:

Improve availability and performance during workload transition and spikes

Faster, less disruptive diagnostics



Cost effective, resilient solution for overflow of MQ shared queues in Coupling Facility (SOD) *

Supported on z/OS V1.13 or Higher **



Dynamic Reconfiguration for Flash Express


Flash Express feature for zEC12 and zBC12

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

** With the March 2013 RSM Enablement Offering Web deliverable.

Flash Express PCIe Adapter Card


Four 400 GB SSDs support
1.4 TB of Storage Class Memory
(AES encrypted)



→ Cable connections to form a RAID 10 Array
on a pair of Flash Express Cards

Exclusive to IBM zEC12 and zBC12
Up to four Flash Express pairs per system

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



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

Cryptography



Crypto Express4S (zEC12 GA Review)

- **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 (No change compared to z196)**
 - For SSL acceleration
 - Clear key RSA operations
- **Enhanced: Secure IBM CCA coprocessor (default)**
 - **Optional:** TKE workstation (#0841) with TKE LIC 7.2 (#850) 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) with TKE LIC 7.2 (#850) for management of the Crypto Express4S when defined as an EP11 coprocessor

zBC12 and zEC12 GA2 Enhancements



- **IBM Enterprise Public Key Cryptography Standards #11(EP11) LIC**
 - Support for PKCS #1 v2.1 PSS
 - Domain Parameter Generate
 - Support for Diffie-Hellman and Elliptic Curve Diffie-Hellman
- **IBM Common Cryptographic Architecture (CCA) 4.4 LIC**
 - EMV enhancement: Diversified Key Generation TDES CBC support
 - DUKPT Initial PIN Encrypting Key (IPEK) Support
 - Export/Import TDES key under AES transport key
 - Remote Key Export (RKX) key wrapping support
 - Integrate specific UDX functions into CCA:
 - Recover PIN from Offset (PFO)
 - Symmetric Key Export with Data verb (SXD)
 - Authentication Parameter Generate (APG)
- **Trusted Key Entry (TKE) 7.3 (FC #0872) LIC**
 - Requires TKE Workstation Hardware 7.3 (FC #0842) or 7.2 (FC #0841) (TKE 7.3 LIC ships installed on TKE Workstation Hardware 7.3)
 - EP11 full function migration wizard
 - Workstation setup wizard
 - Support for PCI-HSM compliance
 - Set Master Key from the TKE
 - Increased strength of TKE session keys

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



▪ Gigabit Ethernet (GbE)

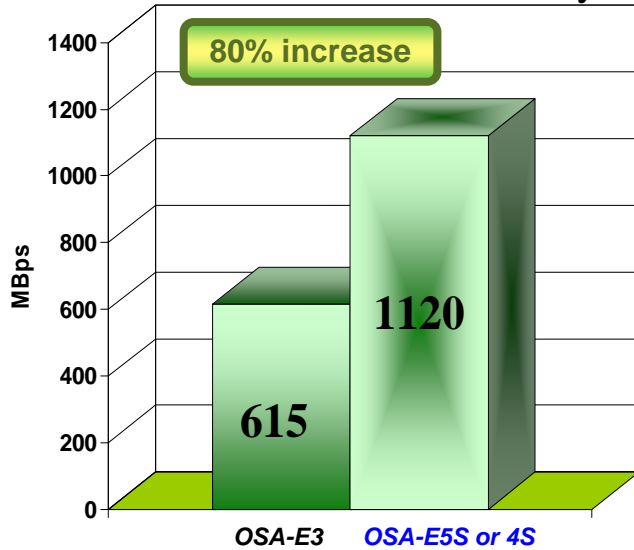
- CHPID types: OSD (**OSN not supported**)
- Single mode (LX) or multimode (SX) fiber
- Two LX or SX SFP+ (D1 top, D2 Bottom)
 - 1 PCHID/CHPID
- **Small form factor pluggable (SFP+) transceivers**
- LC Duplex



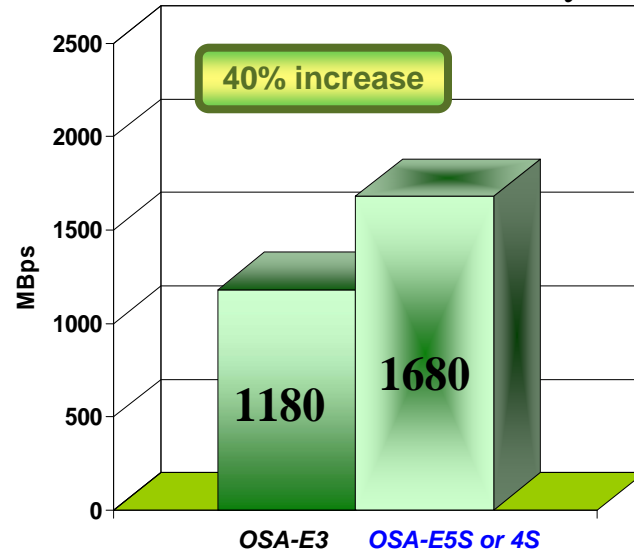
Note: OSA-Express5S features are designed to have the same performance as equivalent OSA-Express4S features.

OSA-Express5S and 4S 10 GbE performance (laboratory)

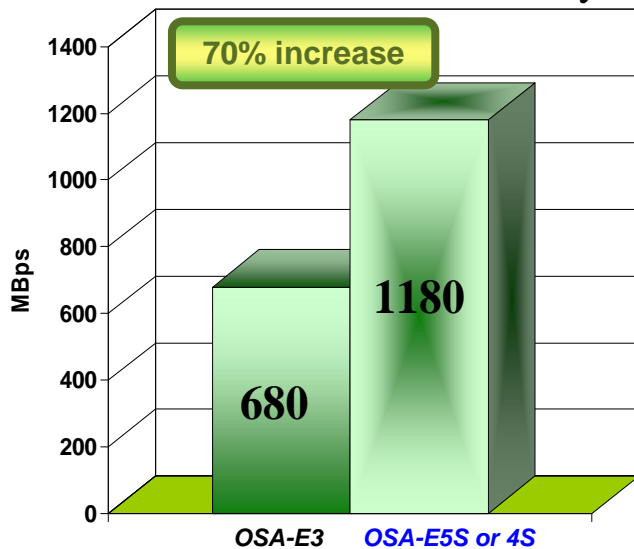
Inbound Streams – 1492 Byte MTU



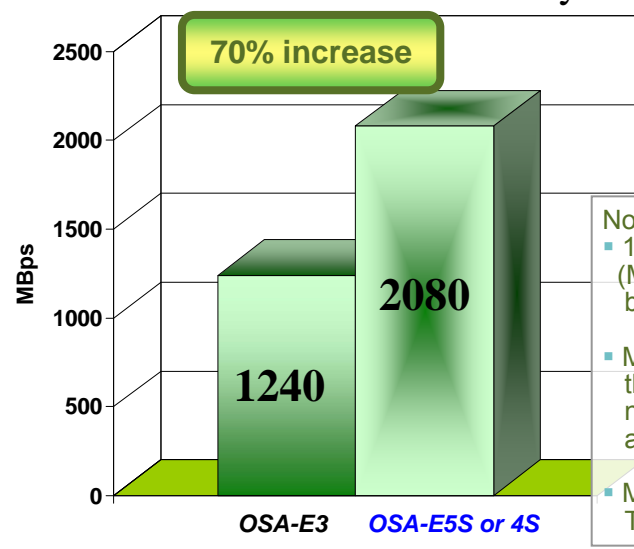
Mixed Streams – 1492 Byte MTU



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

FICON



zBC12 and zEC12 GA2 FICON Overview

▪ Features:

- FICON Express8S (**New build** or **Carry forward**) – 2, 4 or 8 Gbps
 - 10 km LX FC 0409 and SX FC 410: 2 Ports/CHPIDs
- FICON Express8 (**Carry forward only**) – 2, 4 or 8 Gbps
 - 10 km LX FC 3325 and SX FC 3326: 4 Ports/CHPIDs
- FICON Express4 (**Carry forward only**) – 1, 2 or 4 Gbps
 - 10 km LX FC 3321 and SX FC 3322: 4 Ports/CHPIDs; SX (2C) FC 3318 2 ports/CHPIDs

Notes: 1. **Carry forward FICON Express4 if 1 Gbps point-to-point connectivity is required**
 2. **SOD*: FICON Express4 will not be supported on future System z Servers**

- **Not supported:**

- zBC12:** 4 km LX FC 3324: 4 Ports/CHPIDs or 4 km LX (2C) FC 3323 2 ports/CHPIDs
- zEC12:** Same as above plus FICON Express4 SX (2C) FC 3318 2 ports/CHPIDs

▪ Connectors and optics on all FICON features:

- LC Duplex connectors, the same as z114, z10 BC and z9 BC
- Small form factor pluggable (SFP) optics for good serviceability

▪ CHPID TYPEs supported on all features: FCP or FC (selectable by port)

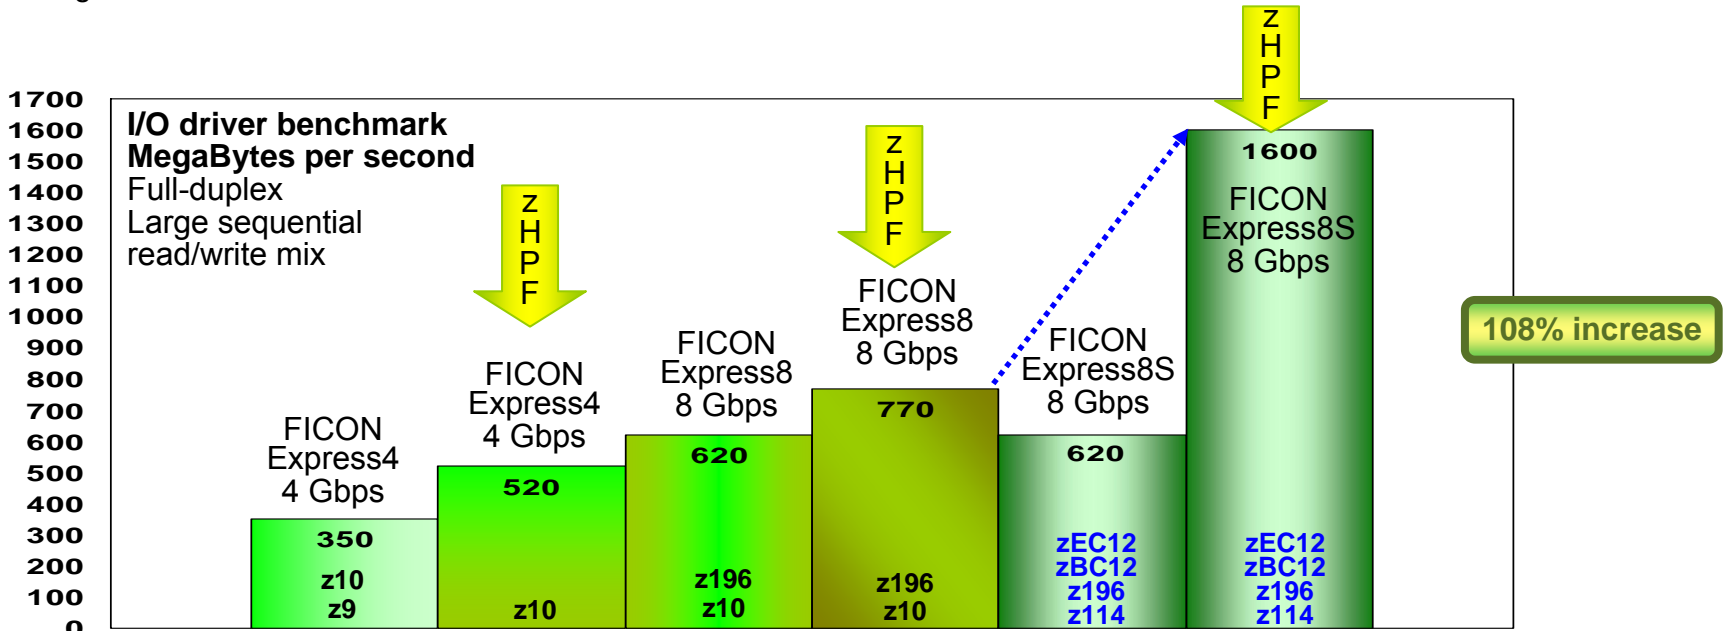
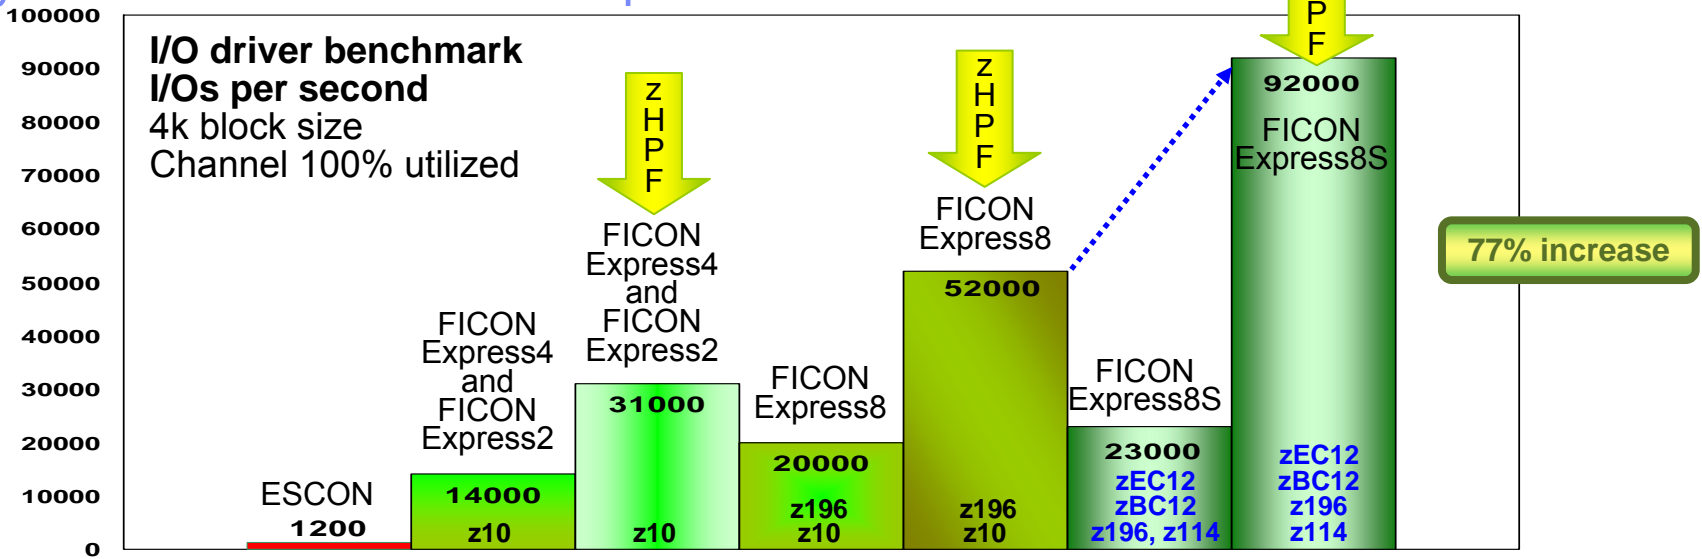
▪ Increased I/O Devices (subchannels) per channel for all FICON features:

- TYPE=FC: Increased to 24k from 16k to support more base and alias devices


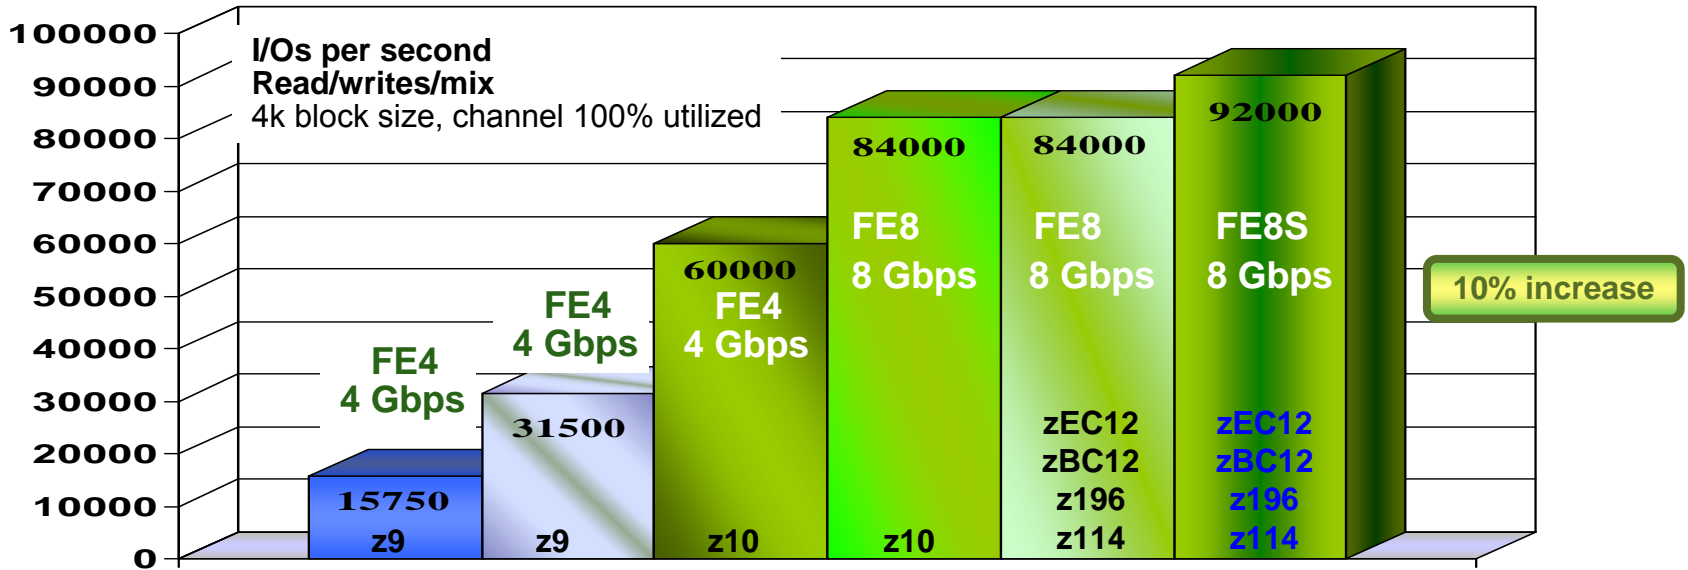


***Note:** All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

System z FICON and zHPF performance



FCP performance on System z

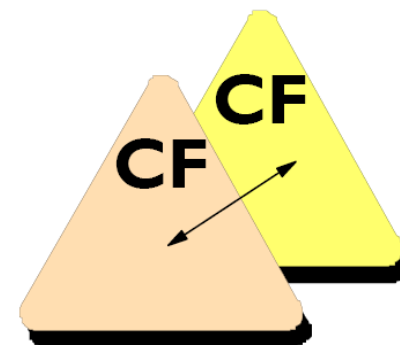


Parallel Sysplex and Server Time Protocol

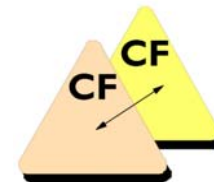


zEC12 GA2 and zBC12 Parallel Sysplex Scalability and Monitoring Enhancements (Introduced in CFCC 18 for zEC12, CFCC 19 for zBC12)

- **Support for up to 101 ICF processors on zEC12**
 - 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 on zEC12 up to 64 1x PSIFB Coupling Links**
 - The limit on previous Enterprise Class System z servers was 48
 - Provides additional link connectivity for STP and for Parallel Sysplex configurations with members of more than one 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
- **Performance/Throughput Enhancements for**
 - Dynamic cache structure size alteration
 - DB2 conditional GBP cache write around the cache to disk
 - Parallel cache castout processing
 - Storage class and castout class contention avoidance

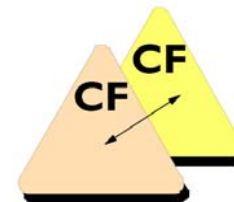


zBC12 and zEC12 GA2 CFCC Level 19: Coupling “Thin Interrupts” Support



- A z/Architecture enhancement allows a new class of “thin interrupts” to be generated when Parallel Sysplex coupling events such as the following occur:
 - On the CF side: **If DYNDISP=THININTERRUPT (new) is specified**
 - A CF command is received by a CF LPAR with shared logical processors (e.g. arrival of a primary CF command that needs to be processed)
 - A CF signal is received by a CF LPAR with shared logical processors (e.g. arrival of a secondary message duplexing signal that needs to be processed)
 - At completion of a secondary message sent by the CF (e.g. completion of a secondary message duplexing signal sent by the image)
 - On the z/OS side: **If enabled in COUPLExx or by SETXCF FUNCTIONS in z/OS 2.1 or z/OS 1.13 and 1.12 with PTFs**
 - A CF signal is received (e.g. arrival of a secondary message from another system)
 - An asynchronous CF operation completes on a z/OS image
- The coupling thin interrupt causes PR/SM to dispatch the partition with shared logical processors, if it is not already dispatched, to allow the request or signal to be processed
- Once dispatched, the CF or z/OS partition with shared logical processors:
 - Will use existing “poll for work” logic as-is to locate and process the work.
 - The partition will give up control:
 - When work is exhausted OR
 - When PR/SM takes the physical processor away from the shared logical processor
 - **The new thin interrupt simply expedites dispatching of the partition.**
- This MAY in some cases eliminate the need to dedicate processors to a CF; but, **Dedicated logical processors are still recommended for best CF performance**

zBC12 and zEC12 GA 2: Parallel Sysplex CFCC Level 19 (continued)



▪ **Statement of Direction*:** CFCC Level 19 exploitation of Flash Express:

- IBM intends to provide exploitation of the Flash Express feature (#0402) on zEC12 and zBC12 servers with Coupling Facility Control Code (CFCC) Level 19 for certain coupling facility list structures in the first half of 2014 for the z/OS V2.1 environment. This new function is designed to allow list structure data to be migrated to Flash Express memory as needed when the consumers of data do not keep pace with its creators for some reason, and migrate it back to real memory to be processed. When using WebSphere MQ for z/OS Version 7 (5655-R36), this new capability is expected to provide significant buffering against enterprise messaging workload spikes and provide support for storing very large amounts of data in shared queue structures, potentially allowing several hours' data to be stored without causing interruptions in processing.

▪ **Structure and CF Memory Sizing with CFCC Level 19**

- May increase memory requirements when moving from CFCC Level 18 (or below) to CFCC Level 19
- Use of the CFSizer Tool is recommended: <http://www.ibm.com/systems/z/cfsizer/>

▪ **Support**

- Exploitation: z/OS V1.12 and higher + PTFs (Note: Flash Express exploitation not yet announced)
- Compatibility: z/OS V1.11 or V1.10 with LifeCycle Extension (Ends for V1.10 in September, 2013)
- Virtual Coupling Facility Guest Support: z/VM V6.3. V6.2 or V5.4 plus PTFs

***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.

Server Time Protocol Enhancements



- **zEC12 GA2 and zBC12: Improved SE time accuracy.** Optionally, the SE can be configured to connect to an external time source periodically to maintain highly accurate time that can be used, if required, to initialize CTN time during POR.

- **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)

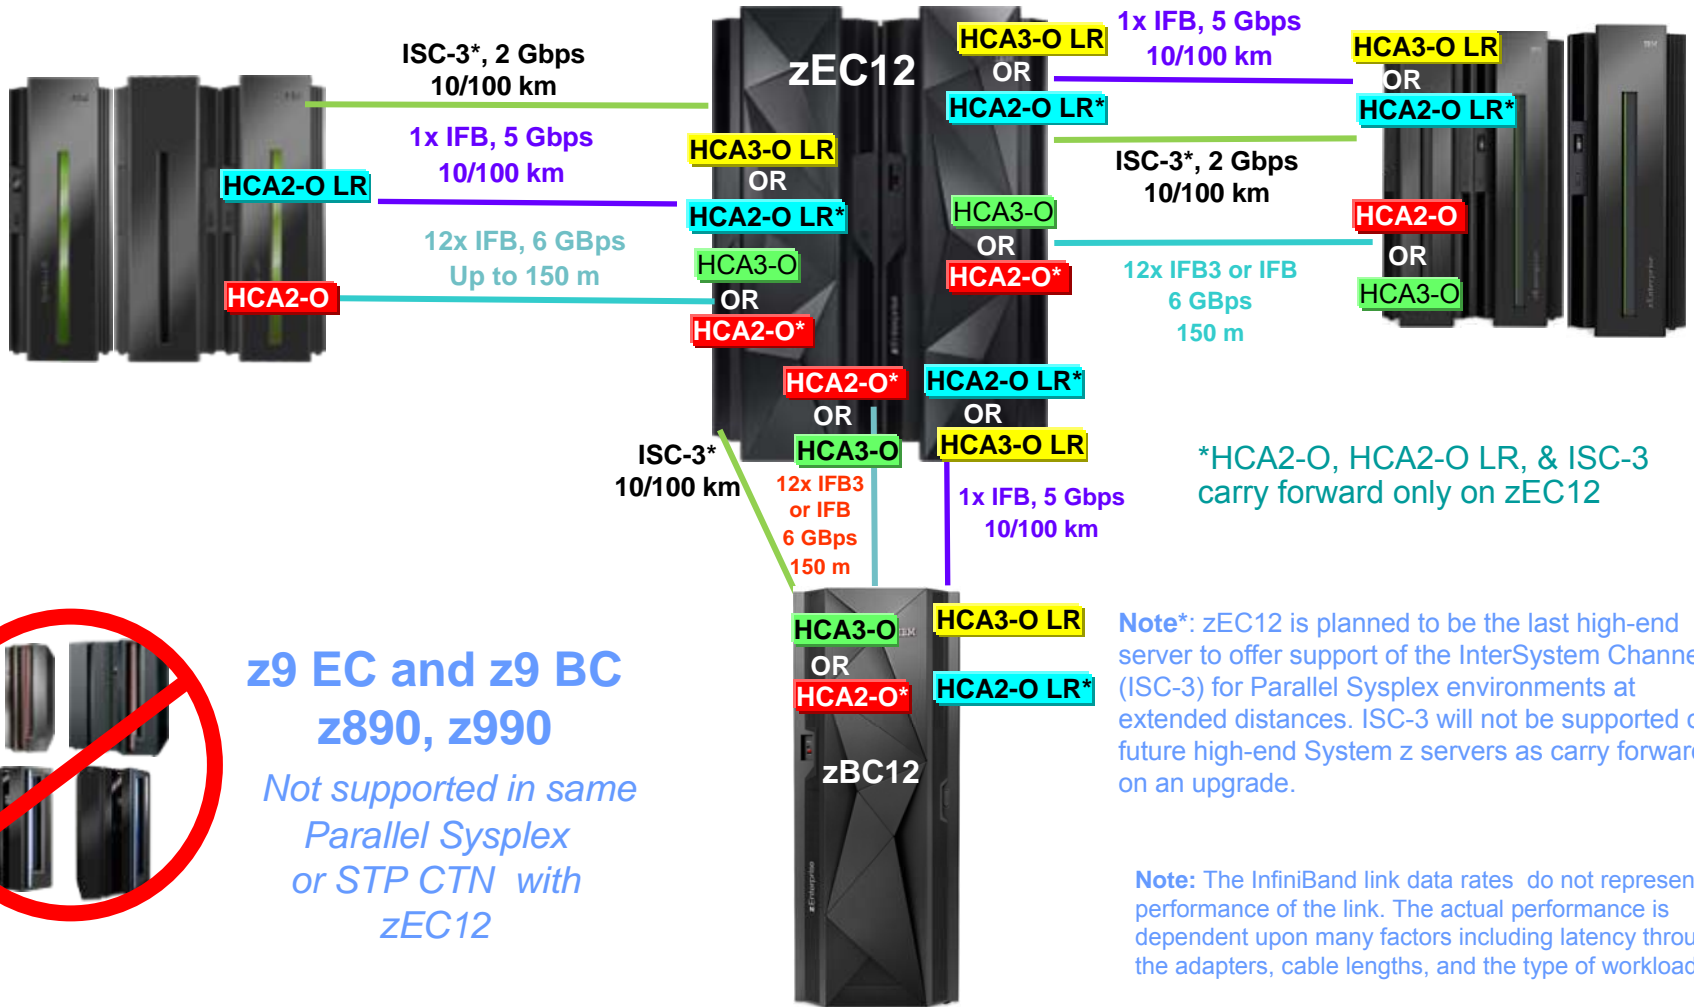
zEC12/zBC12 Parallel Sysplex Connectivity

z10 EC and z10 BC

12x IFB, 1x IFB & ISC-3

z196 and z114

12x IFB, 12x IFB3, 1x IFB, & ISC-3



z9 EC and z9 BC z890, z990

Not supported in same
Parallel Sysplex
or STP CTN with
zEC12

Note*: zEC12 is planned to be the last high-end server to offer support of the InterSystem Channel-3 (ISC-3) for Parallel Sysplex environments at extended distances. ISC-3 will not be supported on future high-end System z servers as carry forward on an upgrade.

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.

Operating System Support



Operating System Support for zBC12

- Currency is key to operating system support and exploitation of future servers
- The following releases of operating systems are supported on zBC12
(refer to PSP buckets for any required maintenance):

Operating System	Supported levels
z/OS	<ul style="list-style-type: none"> ▪ z/OS V2.1 with PTFs (Exploitation plus zEDC Express and 10GbE RoCE Support) ▪ z/OS V1.13 and V1.12 with PTFs (Exploitation) ▪ z/OS V1.11 with PTFS (Toleration, Lifecycle extension required) ▪ z/OS V1.10 (Toleration, Lifecycle Extension required – Ends 9/2013)
Linux on System z	<ul style="list-style-type: none"> ▪ SUSE SLES 10 and SLES 11 ▪ Red Hat RHEL 5 and RHEL 6
z/VM	<ul style="list-style-type: none"> ▪ z/VM V6.3 – exploitation support ▪ z/VM V5.4 and V6.2 – toleration support. (Support for V6.1 ended in April, 2013)
z/VSE	<ul style="list-style-type: none"> ▪ z/VSE compatibility with PTFs, support for: <ul style="list-style-type: none"> • z/VSE 5.1 • z/VSE 4.3
z/TPF	<ul style="list-style-type: none"> ▪ V1.1

- Support for p Blades in zBX Model 003
 - AIX 5.3 Technology Level 12 or higher, AIX 6.1 Technology Level 5 or higher, AIX 7.. All with PowerVM™ Enterprise Edition
- Support for Linux and Windows* environments on select System x blades in zBX Model 003
 - 64 bit version support only
 - Red Hat RHEL 5.5 and higher, 6.0 and higher and SLES 10 (SP4), 11 SP1 and higher
 - Microsoft Windows Server 2008 R2 and Microsoft Windows Server 2008 (SP2) (Datacenter Edition recommended)

THANK YOU



Session 13706

Please complete session evaluations online at SHARE.org/BostonEval



Technical Backup Charts



zBC12 and zEC12 GA2 Summary



zEC12 I/O Feature Cards at GA2



Features	Offered As	Maximum # of features	Maximum channels/adapters	Increments per feature	Purchase increments
FICON					
FICON Express8S	NB	160	320 channels	2 ports/channels	2 channels
FICON Express8	CF ¹	44	176 channels	4 ports/channels	CF Only
FICON Express4 10km LX, SX	CF ¹	44	176 channels	4 or 2 ports/channels	CF Only
ISC-3 Coupling	CF ¹	12	48 links	1, 2, 3 or 4 links	CF Only
OSA-Express					
OSA-Express5S⁴	NB	48	96 ports	1 (10 GbE) or 2 ports	1 feature/channel
OSA-Express4S	CF ¹	48	96 ports	1 (10 GbE) or 2 ports	CF Only
OSA-Express3	CF ¹	24	96 ports	2 (10 GbE) / 4 ports	CF Only
Crypto					
Crypto Express4S⁴	NB	16	16 coprocessors	1 coprocessor	1 feature ²
Crypto Express3 ^{***}	CF ¹	8	16 coprocessors	2 coprocessors	CF Only
Special Purpose					
10 GbE RoCE⁴	NB	16	16 usable SR ports	1 usable SR port	1 feature ³
zEDC Express⁴	NB	8	8 coprocessors	1 coprocessor	1 feature ³
Flash Express⁴	NB	8	8 PCIe adapters	1 PCIe adapter	2 features

1. Carry forward ONLY

2. Two coprocessors initially, one thereafter

3. Purchase in pairs recommended

4. New on zEC12 and zBC12

NB = New Build

CF = Carry Forward

zBC12 I/O Feature Cards



Features	Offered As	Maximum # of features	Maximum channels/adapters	Increments per feature	Purchase increments
FICON					
FICON Express8S	NB	64	128 channels	2 ports/channels	2 channels
FICON Express8	CF ¹	8, 16 RPQ*	64 channels	4 ports/channels	CF Only
FICON Express4 10km LX, SX	CF ¹	8, 16 RPQ*	64 channels	4/2 ports/channels	CF Only
ISC-3 Coupling	CF ¹	8, 12 RPQ*	48 links	1, 2, 3 or 4 links	CF Only
OSA-Express					
OSA-Express5S⁴	NB	48	96 ports	1 (10 GbE) or 2 SFPs	1 feature/channel
OSA-Express4S	CF ¹	48	96 ports	1 (10 GbE) or 2 ports	CF Only
OSA-Express3	CF ¹	8, 16 RPQ*	64 ports	2 (10 GbE) / 4 ports	CF Only
Crypto					
Crypto Express4S⁴	NB	16	16 coprocessors	1 coprocessor	1 feature ²
Crypto Express3	CF ¹	8	16 coprocessors	2/1 coprocessors	CF Only
Special Purpose					
10 GbE RoCE⁴	NB	16	16 usable SR ports	1 usable SR port	1 feature ³
zEDC Express⁴	NB	8	8 coprocessors	1 coprocessor	1 feature ³
Flash Express⁴	NB	8	8 PCIe adapters	1 PCIe adapter	2 features

1. Carry forward ONLY

2. Two coprocessors initially, one thereafter

3. Purchase in pairs recommended

4. New on zEC12 and zBC12

* RPQ 8P2733 Limits the maximum number of PCIe I/O features to 32

NB = New Build

CF = Carry Forward

System z – Maximum Coupling Links and CHPIDs



Server	1x IFB (HCA3-O LR)	12x IFB-IFB3 (HCA3-O)	1x IFB (HCA2-O LR)	12x IFB (HCA2-O)	IC	ICB-4	ISC-3	Maximum External Links	Maximum Coupling CHPIDs
zEC12	64 H20 – 32* H43 – 64*	32 H20 – 16* H43 – 32*	32 ⁽⁴⁾ H20 – 16* H43 – 32*	32 ⁽⁴⁾ H20 – 16* H43 – 32*	32	N/A	48 ⁽⁴⁾	112 ⁽¹⁾ H20 – 72* ⁽²⁾ H43 – 104* ⁽¹⁾	128
zBC12	H13 – 32* H06 – 16*	H13 – 16* H06 – 8*	H13 – 16* H06 – 8*	H13 – 16* H06 – 8*	32	N/A	48 ⁽⁴⁾	H13 – 72* ⁽²⁾ H06 – 56* ⁽³⁾	128
z196	48 M15 – 32*	32 M15 – 16* M32 – 32*	32 M15 – 16* M32 – 32*	32 M15 – 16* M32 – 32*	32	N/A	48	104 ⁽¹⁾ M15 – 72* ⁽²⁾ M32 – 100* ⁽¹⁾	128
z114	M10 – 32* M05 – 16*	M10 – 16* M05 – 8*	M10 – 12* M05 – 8*	M10 – 16* M05 – 8*	32	N/A	48	M10 – 72* ⁽²⁾ M05 – 56* ⁽³⁾	128
z10 EC	N/A	N/A	32 E12 – 16*	32 E12 – 16*	32	16 (32/RPQ)	48	64	64
z10 BC	N/A	N/A	12	12	32	12	48	64	64

* Uses all available fanout slots. Allows no other I/O or coupling.

1. A zEC12 H66, H89 or HA1 supports a maximum 112 extended distance links (64 1x IFB and 48 ISC-3) with no 12x IFB links
A zEC12 H43 supports a maximum 104 extended distance links (56 1x IFB and 48 ISC-3) with no 12x IFB links or other I/O.
A zEC12 H20 or z196 M15 supports a maximum 72 extended distance links (24 1x IFB and 48 ISC-3) with no 12x IFB links or other I/O.
A z196 M49, M66 or M80 supports a maximum of 96 extended distance links (48 1x IFB and 48 ISC-3) with 8 12x IFB links
A z196 M32 supports a maximum of 96 extended distance links (48 1x IFB and 48 ISC-3) with 4 12x IFB links and no other I/O.
2. zEC12 H20, **zBC12 H13**, z196 M15 or z114 M10 support a maximum of 72 extended distance links (24 1x IFB and 48 ISC-3) with no 12x IFB links or I/O.
3. **zBC12 H06** or z114 M05 supports a maximum of 56 extended distance links (8 1x IFB and 48 ISC-3) with no 12x IFB links or I/O.
4. zEC12 H20 and H43 and **zBC12 H06 and H13** support ISC-3, HCA2-O and HCA2-O LR as carry forward only, not on new build
zEC12 H89 and HA1 (only) support ISC-3 as carry forward and on new-build by RPQ when 16 PSIFB fanout features are also configured

zEC12 GA Functional Comparison to z196

Processor / Memory

- Uniprocessor Performance
 - System Capacity
 - Processor Design
 - Models
 - Processing cores
 - Granular Capacity
 - Memory
 - Fixed HSA
- Up to 25% performance improvement over z196 uniprocessor
 - Up to 50% system capacity performance improvement over z196 80-way
 - New 5.5 GHz processor chip versus 5.2 GHz
 - Five models with up to 4 books (z196 had five models)
 - Up to 101 cores to configure, up to 80 on z196
 - Up to 161 capacity settings versus 125 on the z196
 - Up to 3 TB RAIM memory (same as z196)
 - Up to 32 GB fixed HSA versus z196 has 16 GB fixed HSA

Virtualization and Alternative Processors

- Virtualization
 - zEnterprise BladeCenter Extension (zBX)
- zEnterprise Unified Resource Manager provides virtualization management for blades installed in the zBX Mod 003.
 - zEnterprise Unified Resource Manager has “resource workload awareness” where hybrid resources can be managed and optimized across the zEnterprise.
 - zEnterprise System is a truly integrated hardware platform that is able to span and intelligently manage resources across mainframe and distributed technologies – including select POWER7 and IBM System x blades
 - Supported optimizer is IBM WebSphere DataPower XI50 in the zBX Mod 003.
 - zBX Model 003 (versus zBX Model 002 which attaches to z196)

Connectivity

- HiperSockets™
 - FICON
 - I/O subsystem
 - Internal I/O Bandwidth
 - Coupling
 - Cryptography
- Both zEC12 and z196 support of 32 HiperSockets
 - PCIe I/O infrastructure with FICON Express8S and OSA-Express4S adapters including new OSA-Express4S 1000BASE-T
 - Industry standard 8 GBps PCI Express for I/O to provide both high speed connectivity and high bandwidth
 - Parallel Sysplex Coupling with HCA3 DDR InfiniBand Coupling Links up to 6 GBps
 - Crypto Express4S enhanced with new FIPS 140-2 Level 4 cert and PKCS#11 support
 - Elliptic Curve Cryptography (ECC)

RAS

- RAS Focus
 - Availability
- New IBM zAware offers high speed analytics facilitates the ability to consume large quantities of message logs for smarter monitoring
 - zEC12 offers advanced memory enhancements (RAIM) and advanced power and thermal optimization and management that can help to control heat / improve RAS
 - New PCIe Flash Express on zEC12 to handle paging workload spikes and improve availability – not available on z196

Environmentals

- Energy
 - Cooling
- Power Save modes for the processor
 - New improved integrated cooling system
 - Optional Non Raised Floor and overhead cabling options for both I/O and (New!) Power
 - Optional water cooling and DC power

IBM System z Business Class Configuration Comparisons

	z10 BC™ E10	z114 M05	z114 M10	zBC12 H06	zBC12 H13
Uniprocessor Performance	673 MIPS	782 MIPS		1064 MIPS	
z/OS Capacity	26-2760 MIPS	26 - 3139 MIPS		50 – 4958 MIPS	
Total System Memory	248 GB	120 GB	248 GB	240 GB	496 GB
Configurable Engines	10	5	10	6	13
Configurable CPs	0-5	0-5		0 – 6	
LPARS/LCSS	30/2	30/2		30/2	
HiperSockets	16	32		32	
I/O drawers	Up to 4	Up to 3	Up to 3	Up to 3 ⁽¹⁾	Up to 3 ⁽¹⁾
I/O slots per I/O drawers/ PCIe I/O drawers	8	8/32		8/32 ⁽²⁾	
FICON® Channels	128	128		128 ⁽³⁾	
OSA Ports	96	96		96	
ESCON® Channels	480	240		0 ⁽⁴⁾	
IFB host bus Bandwidth PCIe Gen2 Bandwidth	6.0 GB/sec(IFB)	6.0 GB/sec (IFB) 8.0 GB/sec (PCIe)		6.0 GB/sec (IFB) 8.0 GB/sec (PCIe)	
ICB-4/ISC-3 ⁽⁸⁾ /PSIFB	12/48/12	0/48/8 -16	0/48/16 - 32	0 ⁽⁵⁾ /32/8 -16 ⁽⁶⁾	0 ⁽⁵⁾ /32/16 - 32 ⁽⁷⁾
zIIP/zAAP Maximum Qty	5	2	5	3	6
IFL Maximum Qty	10	5 (3139 MIPS)	10 (5390 MIPS)	6 (4958 MIPS)	13 (8733 MIPS)
ICF Maximum Qty	10	5	10	6	13
Capacity Settings	130	130	130	156	156
Upgradeable	Upgrade to z114 or zBC12	Upgrade to zBC12 H06, H13	Upgrade to zBC12 H06, H13	Upgrade H06 to H13, H13 to zEC12 Model H20 (Radiator-based air cooled only)	

Notes for Configuration comparisons chart

(1) Up to 3 drawers standard, a combination of I/O drawers and PCIe I/O drawers as defined

H06		H13	
I/O drawer	PCIe I/O drawer	I/O drawer	PCIe I/O drawer
0	0	0	0
0	1	0	1
0	2	0	2
1	0	1	0
1	1	1	1
		1	2
2*	0	2*	0
2*	1	2*	1

* 2nd I/O drawer offered via an RPQ

- (2) 28 slots per I/O cage, 8 card slots per I/O drawer, 32 per PCIe I/O drawer
- (3) FICON count is based on 2 PCIe I/O drawers (z114/zBC12 or 4 I/O drawers (z10 BC)
- (4) Quantity of 0 ESCON channels is consistent with Statement of Direction
- (5) Quantity of 0 ICB-4 links is consistent with Statements of Direction
- (6) 8 ports of 12x IFB, 16 ports of 1x IFB links available on model H06 based on 4 HCA
- (7) 16 ports of 12x IFB, 32 ports 1x IFB links available on model H13 based on 8 HCA
- (8) ISC-3s. Carry forward only for zBC12/zEC12. Not available for 'new' build or migration offerings

IBM System z Config Comparisons, zBC12 vs. zEC12 Model H20

	zBC12 H06	zBC12 H13	zEC12 Model H20
Uniprocessor Performance	1064 MIPS		1514 MIPS
z/OS Capacity	50 – 4958 MIPS		240 – 21380 MIPS
Maximum System Memory	240 GB	496 GB	704 GB
Configurable Engines	6	13	20
Configurable CPs	0 – 6		0 - 20
LPARS/CSS	30/2		60/4
HiperSockets	32		32
I/O Cages/ I/O drawers/ PCIe I/O drawers	0/2*/2	0/2*/2	1/2/5
I/O slots per Cage/ I/O drawers/ PCIe I/O drawers	0/8/32		28/8/32
FICON® Channels	128 ⁽³⁾		256
OSA Ports (10GbE/1GbE/1000BASE-T)	48/96/96		48/96/96
ESCON® Channels	0 ⁽⁴⁾		0 ⁽⁴⁾
IFB host bus Bandwidth PCIe Gen2 Bandwidth	6.0 GB/sec (IFB) 8.0 GB/sec (PCIe)		6.0 GB/sec 8.0 GB/sec
ICB-4/ISC-3 ⁽⁸⁾ /PSIFB	0 ⁽⁵⁾ /32/8 - 16 ⁽⁶⁾	0 ⁽⁵⁾ /32/16 - 32 ⁽⁷⁾	0 ⁽⁵⁾ /48/16 - 32
zIIP/zAAP Maximum Qty	3	6	10
IFL Maximum Qty	6 (4518 MIPS)	13 (8199 MIPS)	20 (21380 MIPS)
ICF Maximum Qty	6	13	20
Capacity Settings	156	156	80
Upgradeable	Upgrade H06 to H13, H13 to zEC12 Model H20 (Radiator cooled only)		zEC12 H43, H66, H89, HA1 Radiator and Water Cooled

See previous chart for foot notes

Last Slide

