



# SHARE in Boston

August 11 - 16, 2013



## Introducing the IBM zEnterprise® BC12 (zBC12) and EC12 (zEC12) GA2 Hardware: Processor, Memory, System Structure, and Installation Planning

**Session ID: 13707**

Wednesday, August 14, 2013, 8:00 AM  
Room 312, Hynes Convention Center

**Speaker: Harv Emery**

System z Hardware,  
Advanced Technical Sales



Please complete session evaluations  
online at [SHARE.org/BostonEval](http://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](http://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](http://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

zEnterprise

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

# 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.**

# zBC12 and zEC12 GA2 Technical Introduction





## zBX and zManager – Proven hybrid computing environment

### **Support for multi-platform infrastructure with single resource management**

- Investment protection Model 002 to Model 003
  - zManager entitlements move at no-charge
- zBX opens **cloud opportunities** with its fit-for-purpose strategy

### »»»» **What's new?**

#### **For zManager**

- Performance policy managing and load balancing for HX5
- Monitor and reporting to an availability policy
- Support for Microsoft Windows Server 2012 on HX5 blades

#### **New networking connectivity option**

- Layer-2 support between external data network and zBX\*

#### **In the Future**

- With strategy for future including tightly integrated hardware delivery, new compute node form factors, improved virtualization capabilities \*\*

Note: zManager does not support z/VM V6.3 as a member of an ensemble

\* Some restrictions will apply

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



# Why mission critical business analytics on zEnterprise?



## Mission critical = zEnterprise

- C Level executives treat business analytics as mission critical
- Bullet proof security, reliability, disaster recovery, and availability at 99.999%
- High scalability to support mixed workload analytics, high volume of users and requests

## DB2 Analytics Accelerator for z/OS

- Complex business analytics queries will run extremely fast
- Reduces need for distributed data marts
- Storage savings by keeping DB2 z/OS historical data in the accelerator only

## Co-residency of data

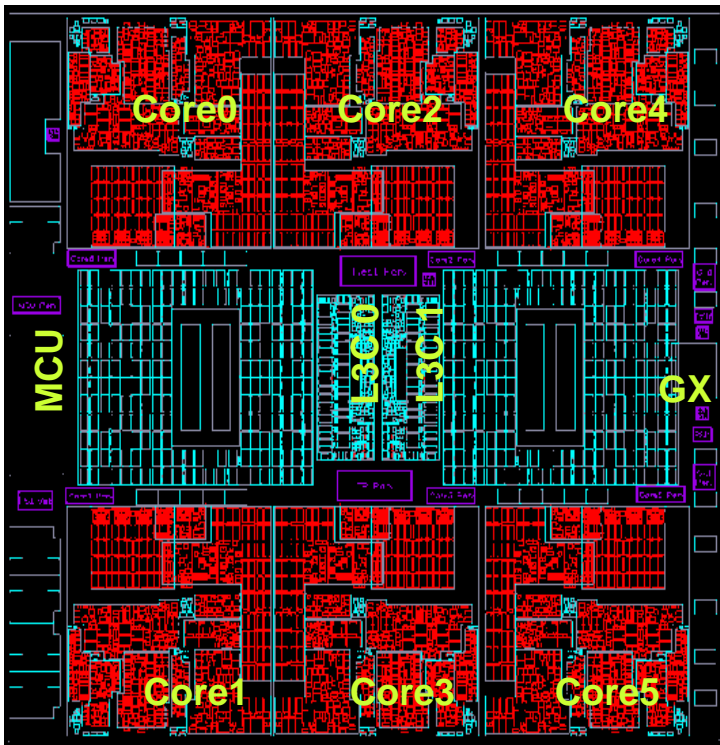
- Production data (70% originates on System z) and data in the warehouse both reside on the same platform
- Data in the warehouse is much more current
- Reduces costs and risk – no need to manage and copy production data into distributed data warehouses and data marts

## True hybrid computing platform

- High volume transaction processing, batch, short duration and complex analytics queries all running concurrently on the same hybrid platform
- Excellent platform for consolidation of costly distributed data warehouses, data marts and business analytics tools

**zEnterprise: The most optimal platform for mixed workloads, large user population, heavy concurrency and high volumes of requests**

# zEC12 and zBC12 Hex Core Processor Chip Detail

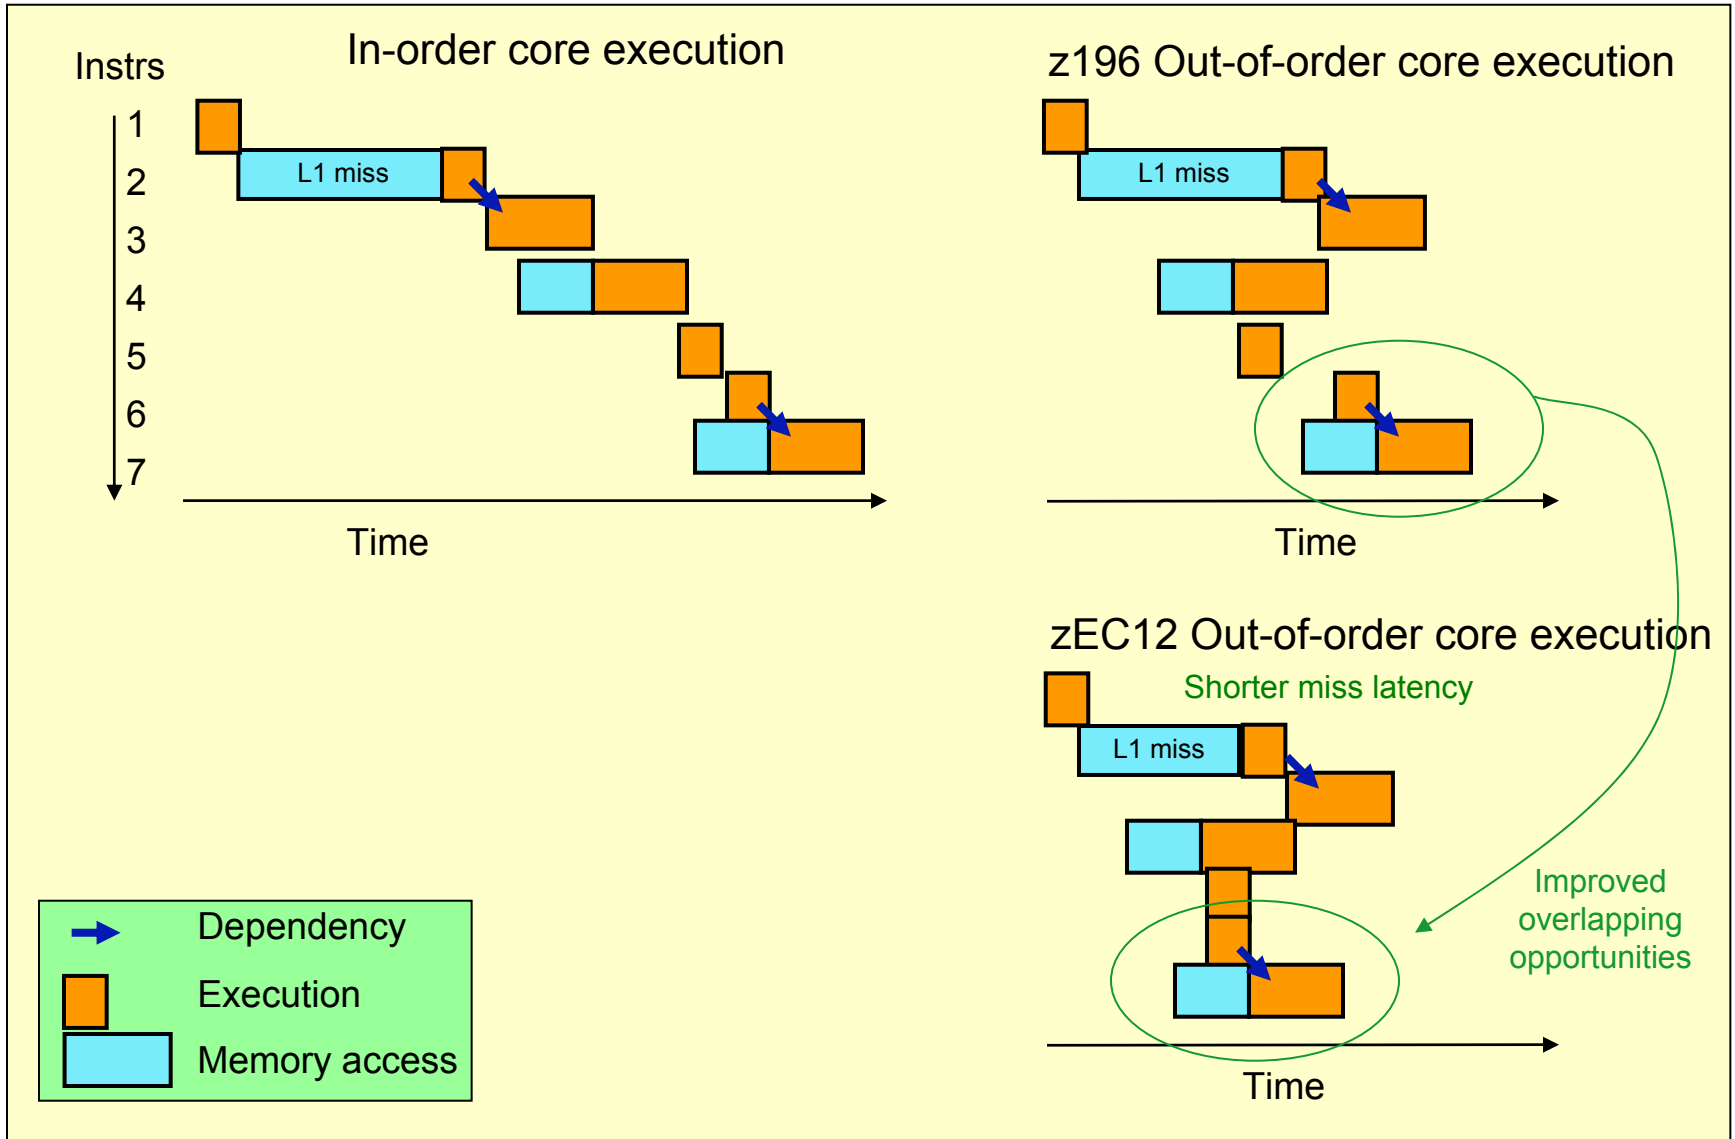


- **13S 32nm SOI Technology**
  - 15 layers of metal
  - 7.68 km wire
- **2.75 Billion Transistors**

- **Chip Area**
  - 597 mm<sup>2</sup>
  - 23.7mm x 25.2mm
  - 10000+ Power pins
  - 1071 signal I/Os

- **Up to six active cores (PUs) per chip**
  - **5.5 GHz – zEC12, 4.2 GHz – zBC12**
  - L1 cache/ core
    - 64 KB I-cache
    - 96 KB D-cache
  - L2 cache/ core
    - 1M+1M Byte hybrid split private L2 cache
- **Dedicated Co-processor (COP) per core**
  - Crypto & compression accelerator
  - Includes 16KB cache
- **Improved Out of Order and Superscalar Instruction Execution**
- **Second Level Branch Prediction Table**
  - Supports 3.5 times more entries
- **On chip 48 MB eDRAM L3 Cache**
  - **Only 24 MB usable on zBC12**
  - Shared by all cores
- **Interface to SC chip / L4 cache**
  - 44 GB/sec to two SC chips (zEC12)
  - 16.8 GB/sec to one SC chip (zBC12)
- **Memory Controller (MCU)**
  - Interface to controller on memory DIMMs
  - Supports RAIM design

# Out of Order Execution – z10, z196 & z114 vs zEC12 & zBC12



## zEC12 & zBC12 Architecture Extensions

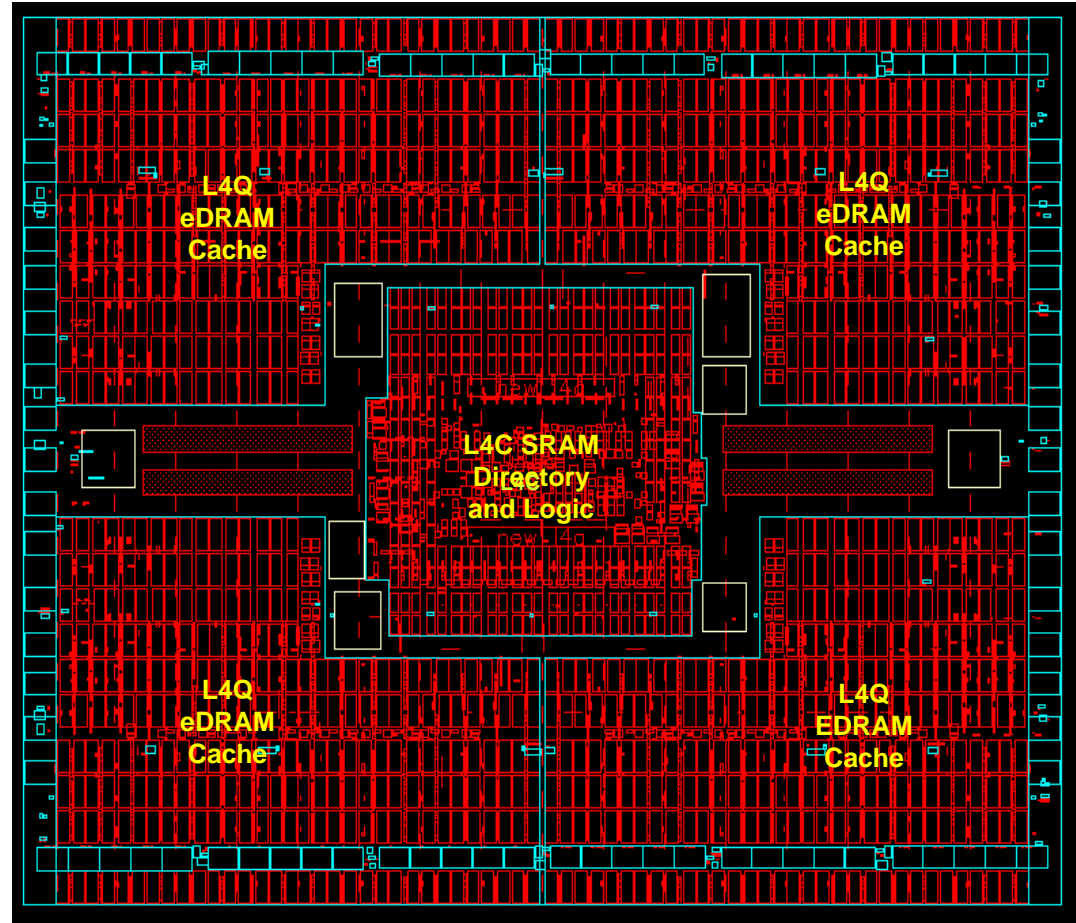
- **Transactional Execution (a/k/a Transactional Memory)**
  - Software-defined sequence treated by hardware as atomic “transaction”
  - Enables significantly more efficient software
    - Highly-parallelized applications
    - Speculative code generation
    - Lock elision
  - **Exploitation by Java 7 for z/OS and the IBM Enterprise COBOL Compiler for z/OS, V5.1** longer-term opportunity for DB2, z/OS V2.1, etc
- **2 GB page frames**
  - Increased efficiency for DB2 buffer pools, Java heap, other large structures
- **Software directives to improve hardware performance**
  - Data usage intent improves cache management
  - Branch pre-load improves branch prediction effectiveness
  - Block prefetch moves data closer to processor earlier, reducing access latency
- **New Decimal-Floating-Point Zoned-Conversion Facility that can help to improve performance applications compiled with the Enterprise PL/I for z/OS, V4.3.**



**Note: IBM Enterprise PL/I for z/OS, V4.3 and IBM Enterprise COBOL Compiler for z/OS, V5.1 (GA: June 21, 2013) can optionally exploit new z/Architecture instructions introduced from 2000 (z900) to 2012 (zEC12)**

# zEC12 & zBC12 Storage Control (SC) Chip Detail

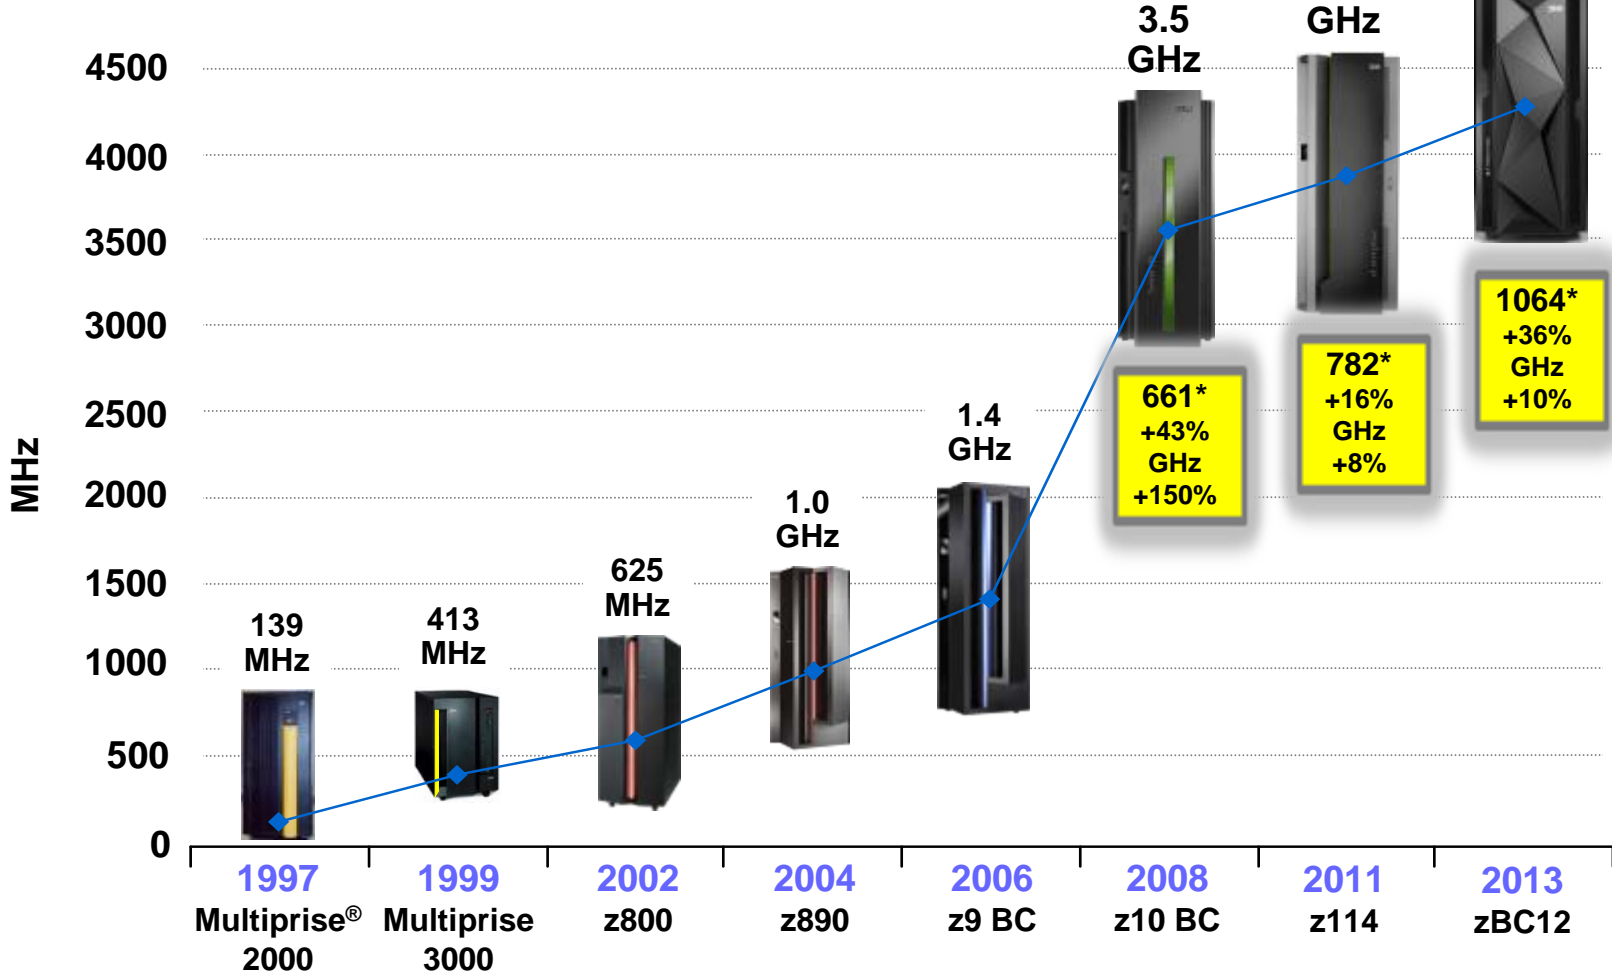
- **CMOS 13S 32nm SOI Technology**
  - 15 layers of metal
- **Chip Area –**
  - 526 mm<sup>2</sup>
  - 26.72mm x 19.67mm
  - 7311 Power Connectors
  - 1819 Signal Connectors
- **3.3 Billion Transistors**
  - 2.1 Billion eDRAM transistors
- **eDRAM Shared L4 Cache**
  - 192 MB per SC chip
  - 2 SCs = 384 MB per zEC12 book
  - **1 SC = 192 MB per zBC12 drawer**
- **6 CP chip interfaces**
- **Book fabric interfaces**
  - zEC12 Book Interconnect
  - zBC12 Drawer Interconnect
- **1 Clock domain**
- **4 Unique chip voltage supplies**



# zBC12 Structure



# zBC12 continues the System z mainframe heritage Uniprocessor MIPS and GHz comparison



\* MIPS Tables are NOT adequate for making comparisons of System z processors




# zBC12 SCM Vs z114 SCM Comparison

## zBC12 Single Chip Modules (SCMs)


- **Processor Unit (PU) SCM**
  - 50mm x 50mm – fully assembled
  - PU Chip size 23.7 mm x 25.2 mm
  - Six core chip with 4 or 5 active cores (4.2 GHz)
- **Storage Control (SC) SCM**
  - 50mm x 50mm – fully assembled
  - SC Chip size 26.72 mm x 19.67 mm
  - 192 MB L4 cache per SCM
- **Processor Drawer**
  - Two PU SCMs (9 processors) per drawer
  - One SC SCM (192 MB of L4 cache) per drawer
  - H06: One Drawer, H13: Two Drawers



Single PU Chip without Module Thermal Cap



Single SC Chip without Module Thermal Cap




## z114 Single Chip Modules (SCMs)

- **Processor Unit (PU) SCM**
  - 50mm x 50mm – fully assembled
  - PU Chip size 23.50 mm x 21.80 mm
  - Four core chip with 3 or 4 active cores (3.8 GHz)
- **Storage Control (SC) SCM**
  - 50mm x 50mm – fully assembled
  - SC Chip size 24.43 mm x 19.60 mm
  - 96 MB L4 cache per SCM
- **Processor Drawer**
  - Two PU SCMs (7 processors) per drawer
  - One SC SCM (96 MB of L4 cache) per drawer
  - M05: One Drawer, M10: Two Drawers

Single PU Chip without Module Thermal Cap




Single SC Chip without Module Thermal Cap





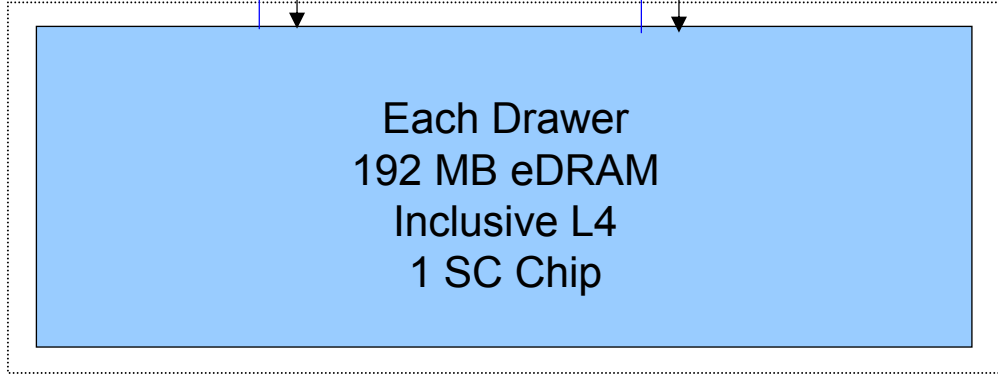
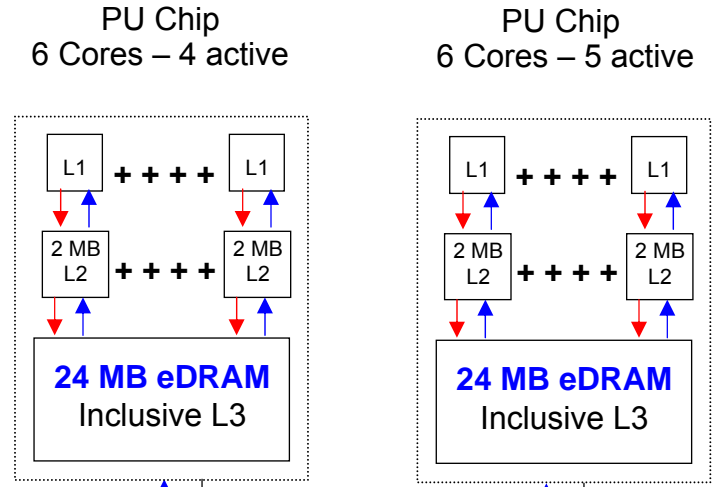
## zBC12 Processor Drawer (Top View)

- Two Processor SCMs (9 processors)
- One Storage Control SCM (192 MB L4 cache)
- Slots for 10 DIMMs (40 to 320 GB RAIM)





# zBC12 Drawer Level Cache Hierarchy



- + Cache for cores 1 to 6
- LRU Cast-Out
- CP Stores
- Data Fetch Return

Cache Comparison	
z114	zBC12
L1 – 64K/128K	L1 – 64K/96K
L2 – 1.5 M	L2 – 1M/1M
L3 – 12 M	L3 – 24 M
L4 – 96 M	L4 – 192 M



## Processor Drawers (Model H06 and H13)


### ■ Model H06 – One Processor Drawer

- Up to 6 client engines plus 2 SAPs and 1 IFP
  - No more than 6 CPs, IFLs or ICFs
  - zIIPs and zAAPs limited to **two** per CP
- 8 to 240 GB of client memory
- Up to 4 I/O Fanouts to support I/O drawers and coupling links



### ■ Model H13 – Two Processor Drawers

- **More processors, memory and I/O fanouts**
- Up to 13 client engines plus 2 SAPs, 1 IFP and 2 spares
  - No more than 6 CPs, up to 13 IFLs or ICFs
  - zIIPs and zAAPs limited to **two** per CP
- 8 to 496 GB of client memory
- Up to 8 I/O Fanouts to support I/O drawers and coupling links



**Planning Note: Unlike the zEC12 Books, add or remove of a zBC12 processor drawer requires a scheduled outage**

## zEnterprise BC12 Models H06 versus H13

- **M/T 2828 – Model H06**

- Air cooled
- Single Frame
- Non-raised floor option available
- 30 LPARs

- **Processor Units (PUs)**

- One Processor drawer
- 9 processors
  - 2 SAPs and 1 IFP standard
  - Up to 6 CPs
  - Up to 6 IFLs or ICFs
  - zIIP and zAAPs per ratio to CPs
  - 0 spares when fully configured

- **M/T 2828 – Model H13**

- Air cooled
- Single Frame
- Non-raised floor option available
- 30 LPARs

- **Processor Units (PUs)**

- **Two processor drawers**
- **18 processors**
  - 2 SAPs and 1 IFP standard
  - Up to 6 CPs
  - **Up to 13 IFLs or ICFs**
  - zIIPs and zAAPs per ratio to CPs
  - **2 dedicated spare processors**

- **Why the H13 (2<sup>nd</sup> processor drawer)?**

More than 6 Customer engines (processors)

More than 240 GB memory (Linux?)

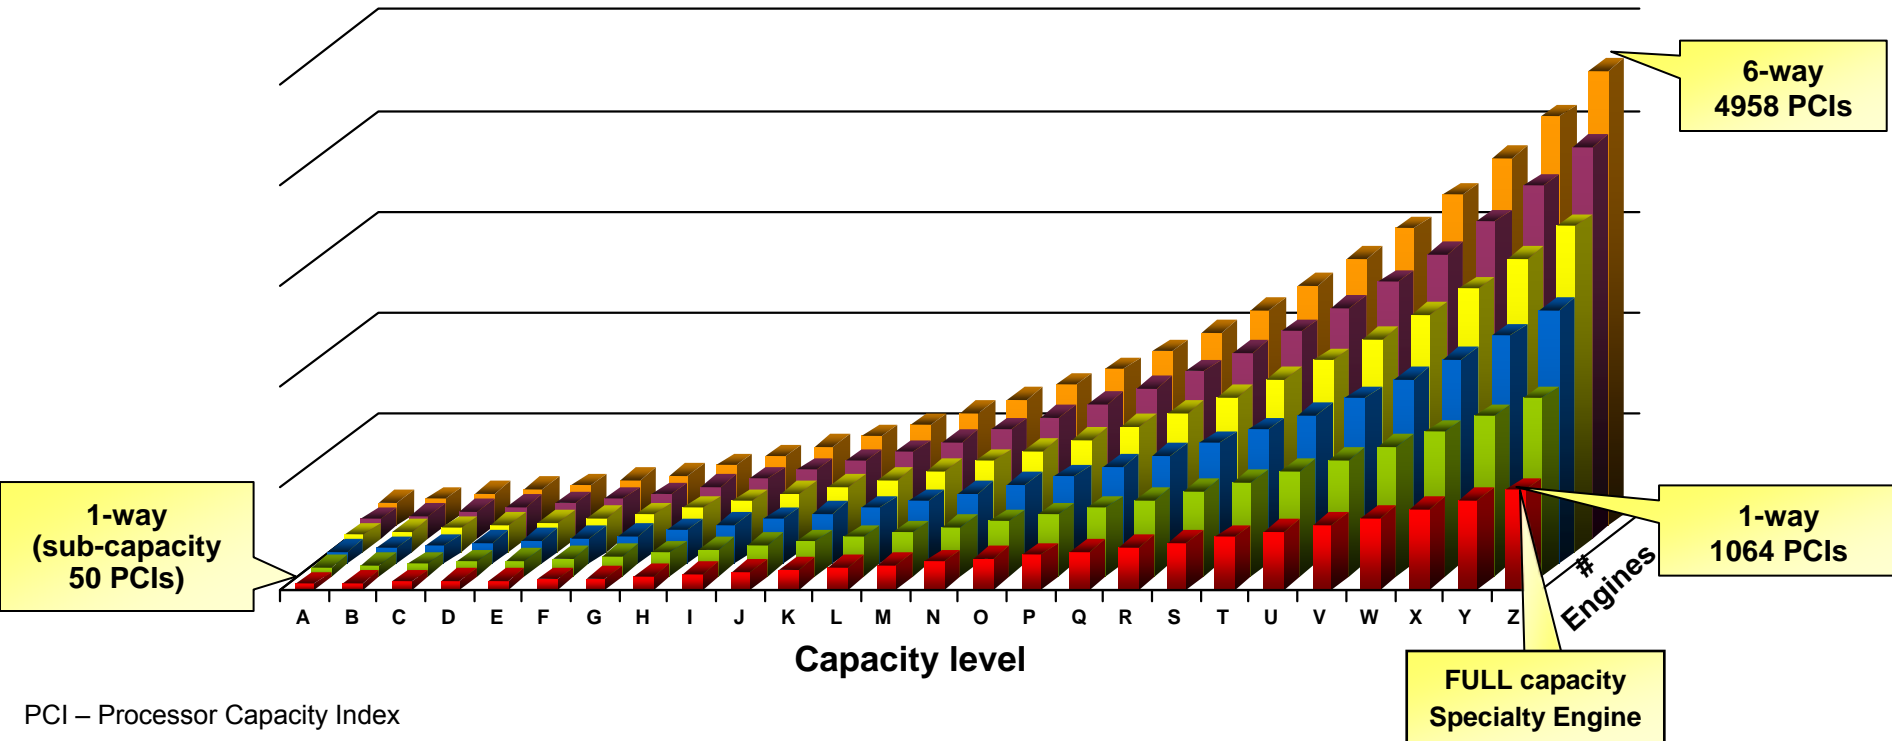
More than 4 I/O Fanouts for connectivity – especially PSIFB links

- Needs vary depending on I/O feature and PSIFB link requirements

## zBC12 Sub-capacity Processor Granularity

- The zBC12 has 26 CP capacity levels (26 x 6 = 156)
  - Up to 6 CPs at any capacity level
  - All CPs must be the same capacity level
- The **two** for one entitlement to purchase **two** zAAPs and/or **two** zIIPs for each CP purchased is the same for CPs of any capacity.
- All specialty engines run at full capacity
  - Processor Value Unit for IFL = 100 (Equal to z114)

Number of zBC12 CPs	Base Ratio	Ratio z114 To zBC12
1 CP	z114 Z01	1.36
2 CPs	z114 Z02	1.37
3 CPs	z114 Z03	1.37
4 CPs	z114 Z04	1.36
5 CPs	z114 Z05	1.36
6 CPs	z114 Z05	1.58



PCI – Processor Capacity Index

zBC12 Processor Features – New zAAP/zIIP to CP 2:1 ratio and new IFP



Model	Drawers/ PUs	CPs	IFLs uIFLs	<i>zAAPs</i>	<i>zIIPs</i>	ICFs	Std SAPs	Optional SAPs	Std. Spares	<i>IFP</i>
<b>H06</b>	<b>1/9</b>	<b>0-6</b>	<b>0-6 0-5</b>	<b>0-4</b>	<b>0-4</b>	<b>0-6</b>	<b>2</b>	<b>0-2</b>	<b>0</b>	<b>1</b>
<b>H13</b>	<b>2/18</b>	<b>0-6</b>	<b>0-13 0-12</b>	<b>0-8</b>	<b>0-8</b>	<b>0-13</b>	<b>2</b>	<b>0-2</b>	<b>2</b>	<b>1</b>

- ▶ **The integrated firmware processor (IFP) supports Resource Group (RG) Licensed Internal Code (LIC) to provide native PCIe I/O feature management and virtualization functions. This LIC is conceptually similar to the Channel Subsystem LIC that supports traditional System z I/O features.**
- ▶ Adding the second drawer to upgrade from the zBC12 Model 6 to Model 13 is disruptive.

Notes:

1. “PU” stands for a Processor Unit, a functional processor core that can be purchased by the client.
2. At least one CP, IFL, or ICF must be purchased in every machine
- 3. Up to two zAAPs and two zIIPs may be purchased for each CP purchased if PUs are available. This remains true for sub-capacity CPs and for “banked” CPs.**
4. “uIFL” stands for Unassigned IFL
5. Optional SAPs are almost never needed except if running the z/TPF operating system.


## IBM zIIP and zAAP Simplification (Update of the August 28, 2012 Statement of Direction)

- **IBM System z Integrated Information Processor (zIIP) and IBM System z Application Assist Processor (zAAP) Simplification<sup>1</sup>**
    - The IBM zEnterprise EC12 **and the zEnterprise BC12** are planned to be the last System z servers to offer support for zAAP specialty engine processors. IBM intends to continue support for running zAAP workloads on zIIP processors ("zAAP on zIIP"). This is intended to help simplify capacity planning and performance management, while still supporting all the currently eligible workloads.
    - In addition, IBM has provided a PTF for [APAR OA38829](#) on z/OS V1.12 and V1.13 to remove the restriction that prevents zAAP-eligible workloads from running on zIIP processors when a zAAP is installed on the server. **Note:** This PTF still does not allow zAAP-eligible work to run on a zIIP in any LPAR with a zAAP assigned to it.
- This works on any System z server that supports zIIPs and zAAPs, not just on zEC12 and zBC12.**

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.



# zBC12 H13 – Under the covers



Rear View



Front View

Internal Batteries (optional)

Power Supplies

Processor drawers: Memory and HCAs  
Top Drawer: Model H13 only

I/O drawer

PCIe I/O drawers

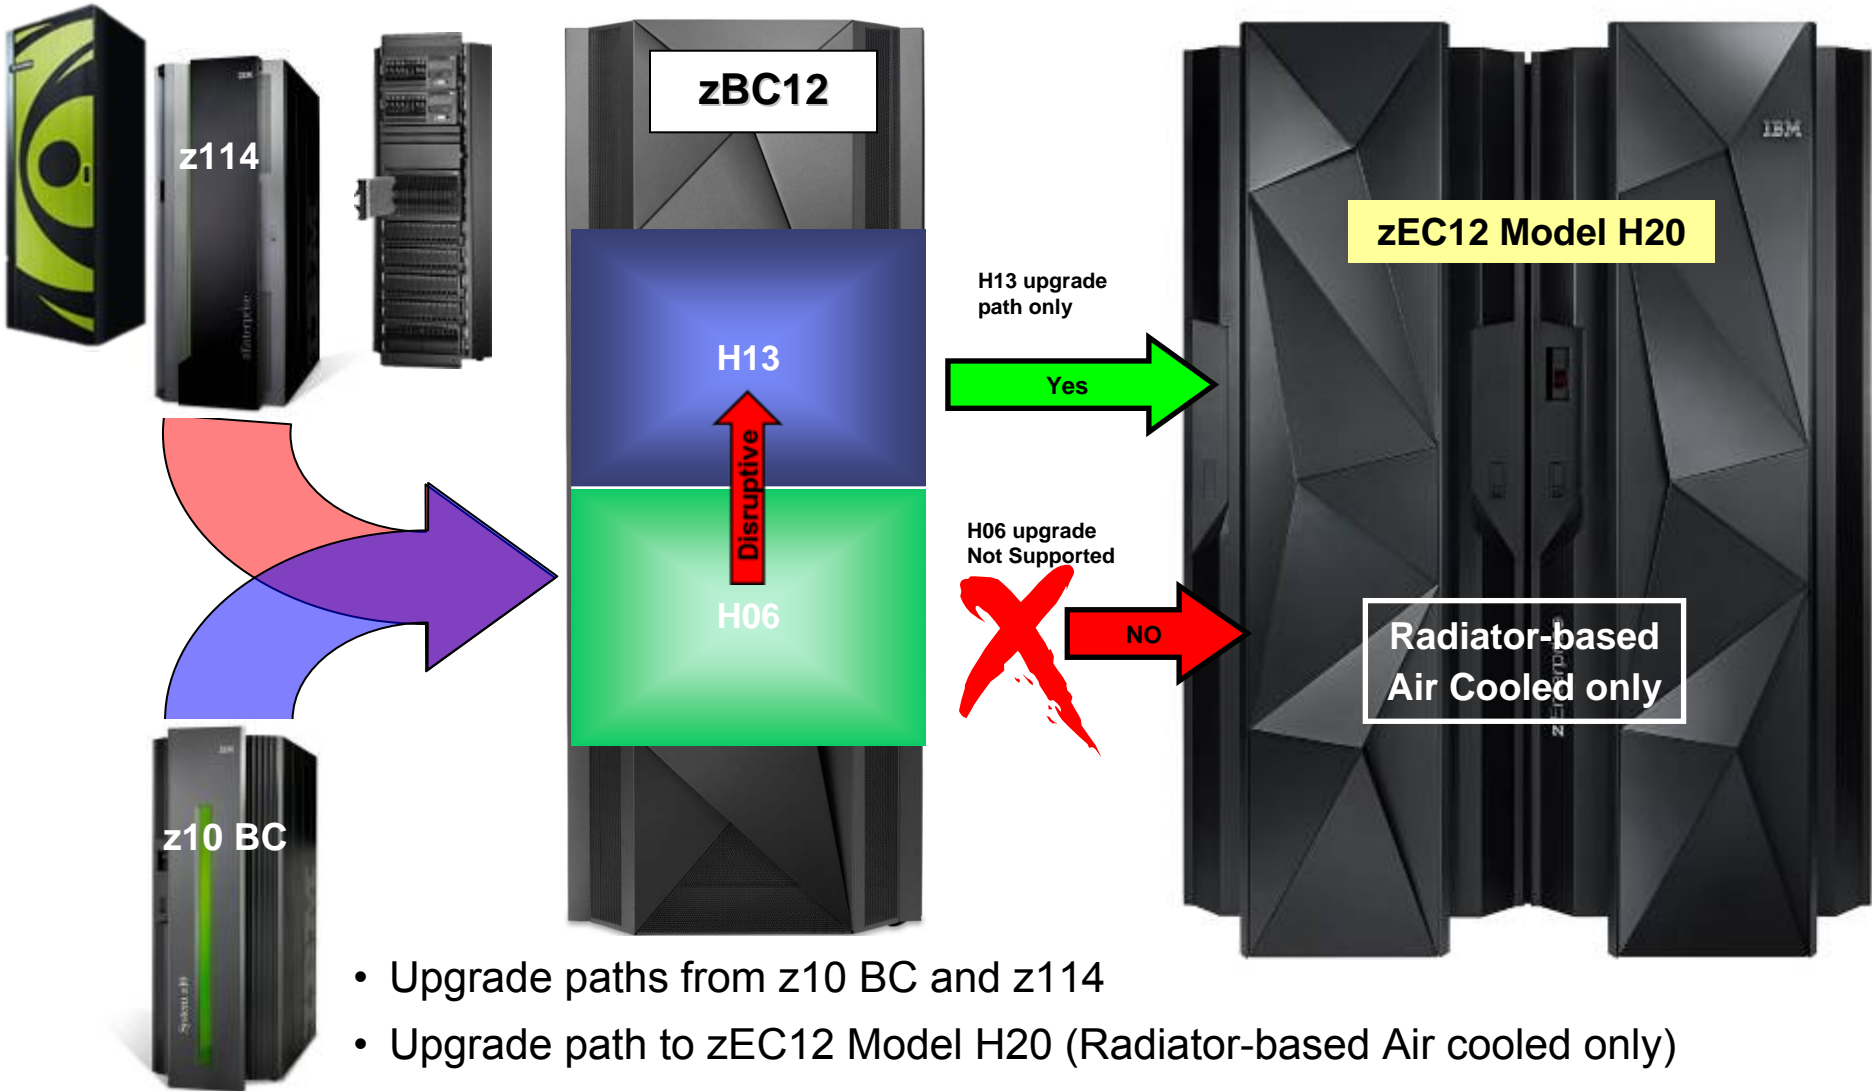
Note: One I/O drawer with two PCIe I/O drawers requires a total of six HCAs so it requires a Model H13.

Ethernet cables for internal System LAN connecting Flexible Service Processor (FSP) cage controller cards (not shown)

2 x Support Elements



## zBC12 Upgrade paths



- Upgrade paths from z10 BC and z114
- Upgrade path to zEC12 Model H20 (Radiator-based Air cooled only)
- Disruptive upgrade H06 to H13 and from H13 to zEC12 H20

# zEC12 Structure



## zBC12 SCM Vs zEC12 MCM Comparison

### zBC12 Single Chip Modules (SCMs)

- **Processor Unit (PU) SCM**

- 50mm x 50mm – fully assembled
- PU Chip size 23.7 mm x 25.2 mm
- Six core chip with 4 or 5 active cores




- **Storage Control (SC) SCM**

- 50mm x 50mm – fully assembled
- SC Chip size 26.72 mm x 19.67 mm
- 196 MB L4 cache per SCM


- **Processor Drawer**

- Two PU SCMs (9 processors) per drawer
- One SC SCM (192 MB of L4 cache) per drawer
- H06: One Drawer, H13: Two Drawers

Single PU Chip  
without  
Module Thermal Cap



Single SC Chip  
without  
Module Thermal Cap



### zEC12 Multi Chip Module (MCM)

- **Technology**

- 96mm x 96mm with 102 glass ceramic layers
- 7,356 LGA connections to 8 chip sites


- **Six 6-core Processor (PU) chips**

- Each with 4, 5 or 6 active cores
- 27 active processors per MCM (30 in Model HA1)
- PU Chip size 23.7 mm x 25.2 mm

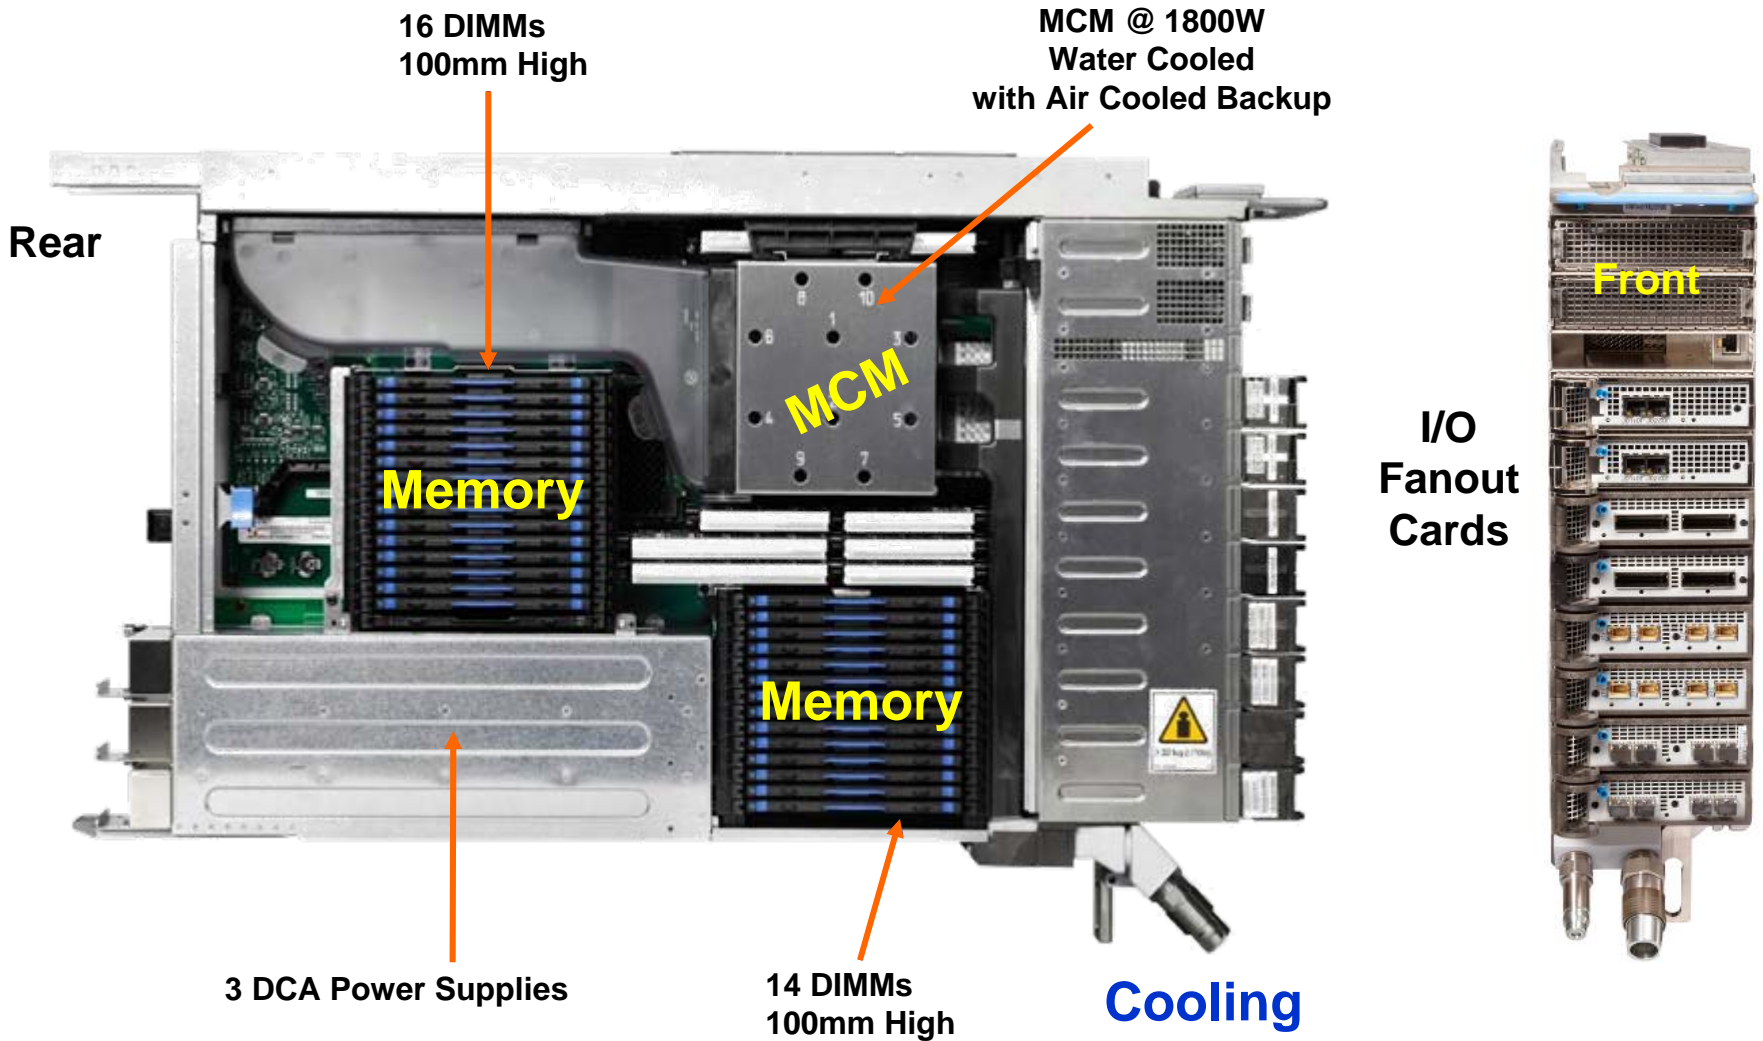
- **Two Storage Control (SC) chips per MCM**

- 192 MB L4 cache per SC, 384 MB per MCM
- SC Chip size 26.72 mm x 19.67 mm

- **One MCM per book, up to 4 books per System**

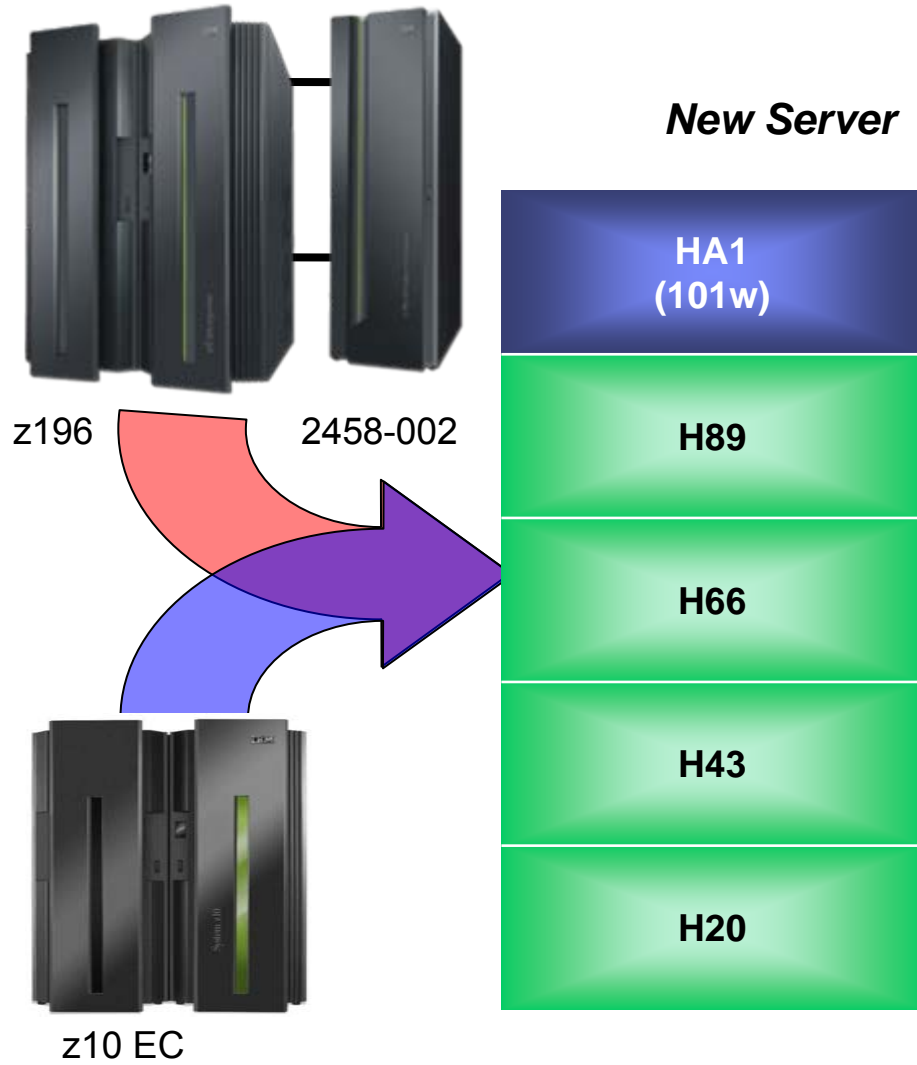


# zEC12 Book Layout



Note: Unlike the z196, zEC12 Books are the same for the Air and Water cooled Systems

# System Offering Overview



## Machine Type for zEC12

- ▶ 2827

## Processors

- ▶ 27 / 30 PUs per book
- ▶ Sub-capacity available up to 20 CPs
- ▶ 2 spares designated per system

## Memory

- ▶ System minimum = 32 GB with separate 32 GB HSA
- ▶ Maximum 3TB / 768GB per book
- ▶ RAIM memory design
- ▶ Purchase Increments – 32, 64, 96, 112, 128, 240, 256, 512 GB

## I/O

- ▶ Up to 16 connections per book (Up to 8 fanouts, 2 per fanout)
- ▶ PCIe connections 8 GB/sec
- ▶ InfiniBand 6 GB/sec

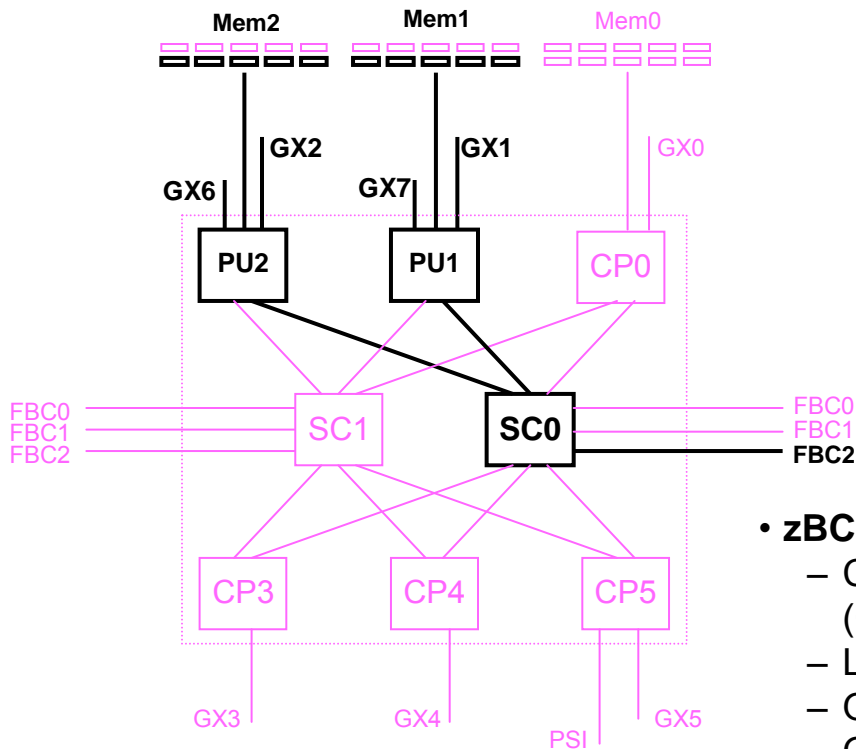
## Machine Type and Model for zBX

- ▶ 2458-003

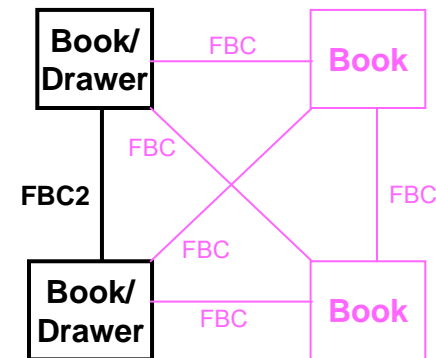


# zBC12 Internal System Structure – Compared to the zEC12

## Book / Drawer Structure



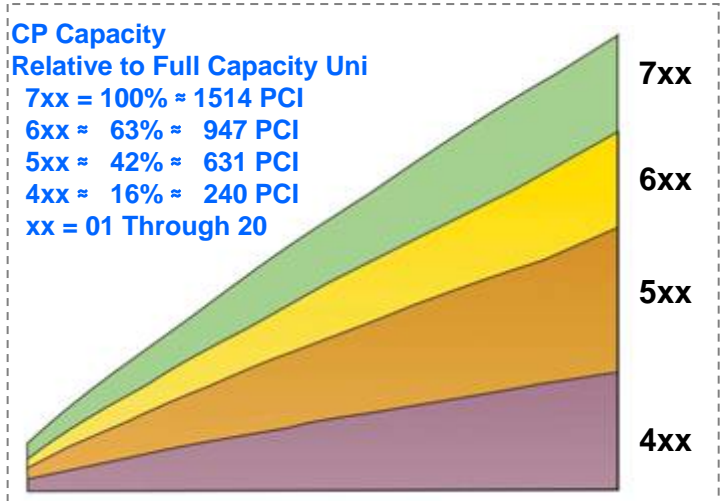
## Book / Drawer Interconnect



- **zBC12: A subset of the zEC12 design and chipset**
  - Cache: Half L3 on chip usable and L4 drawer cache (Only one SC chip per drawer instead of two)
  - Lower frequency: Reduced interconnect bandwidth
  - One fewer memory controllers per drawer
  - Only one DIMM per memory channel instead of two

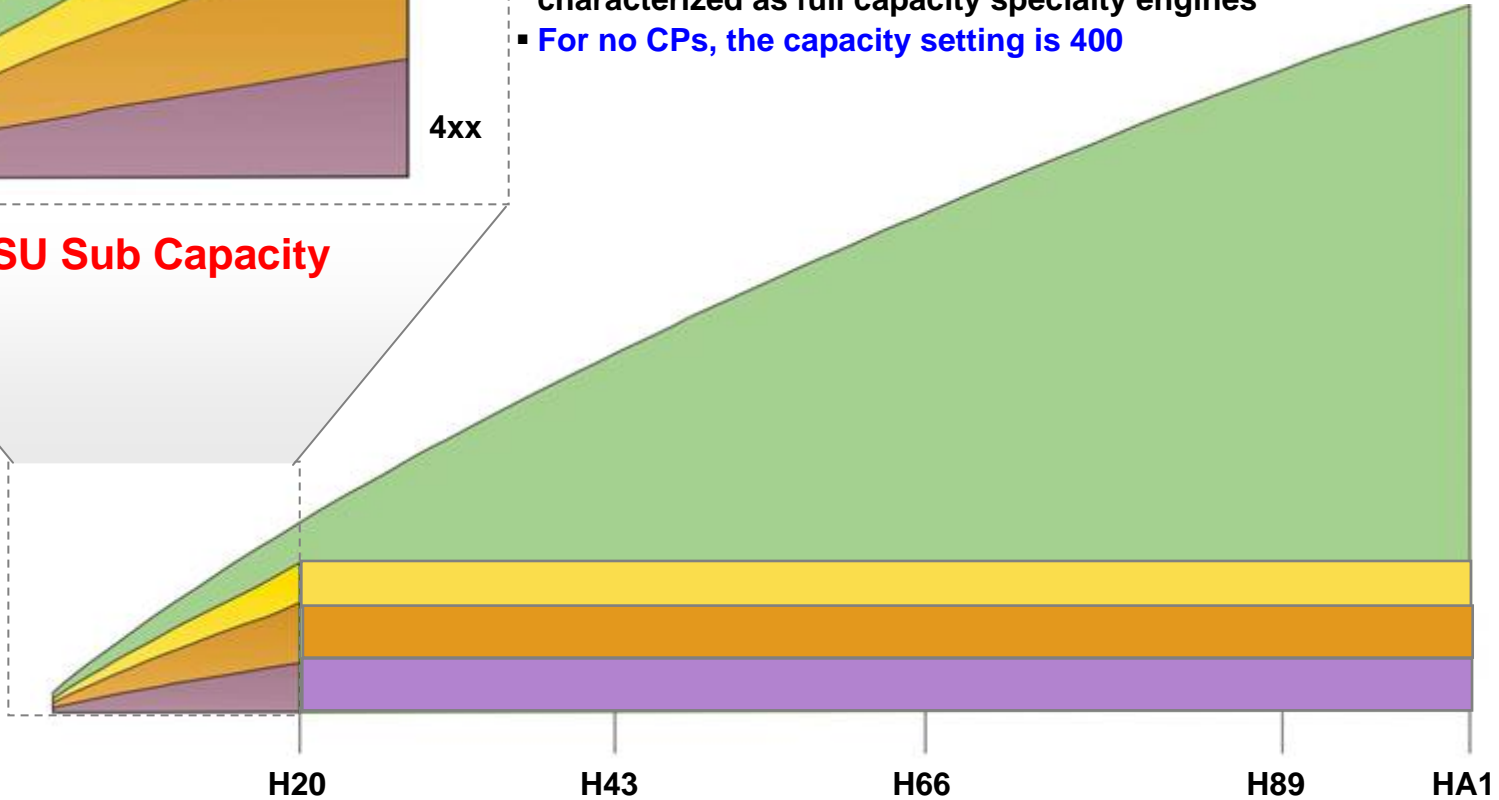
- **Diagrams color code**
  - zEC12 and zBC12
  - zEC12 only

# zEC12 Full and Sub-Capacity CP Offerings



- Subcapacity CPs, up to 20, may be ordered on ANY zEC12 model. If 21 or more CPs are ordered all must be full 7xx capacity
- All CPs on a zEC12 must be the same capacity
- All specialty engines run at full capacity. The **two** for one entitlement to purchase **two** zAAPs and **two** zIIPs for each CP purchased is the same for CPs of any capacity.
- Only 20 CPs can have granular capacity but other PU cores may be characterized as full capacity specialty engines
- For no CPs, the capacity setting is 400

## MSU Sub Capacity





zEC12 Processor Features – New zAAP/zIIP to CP ratio and new IFB

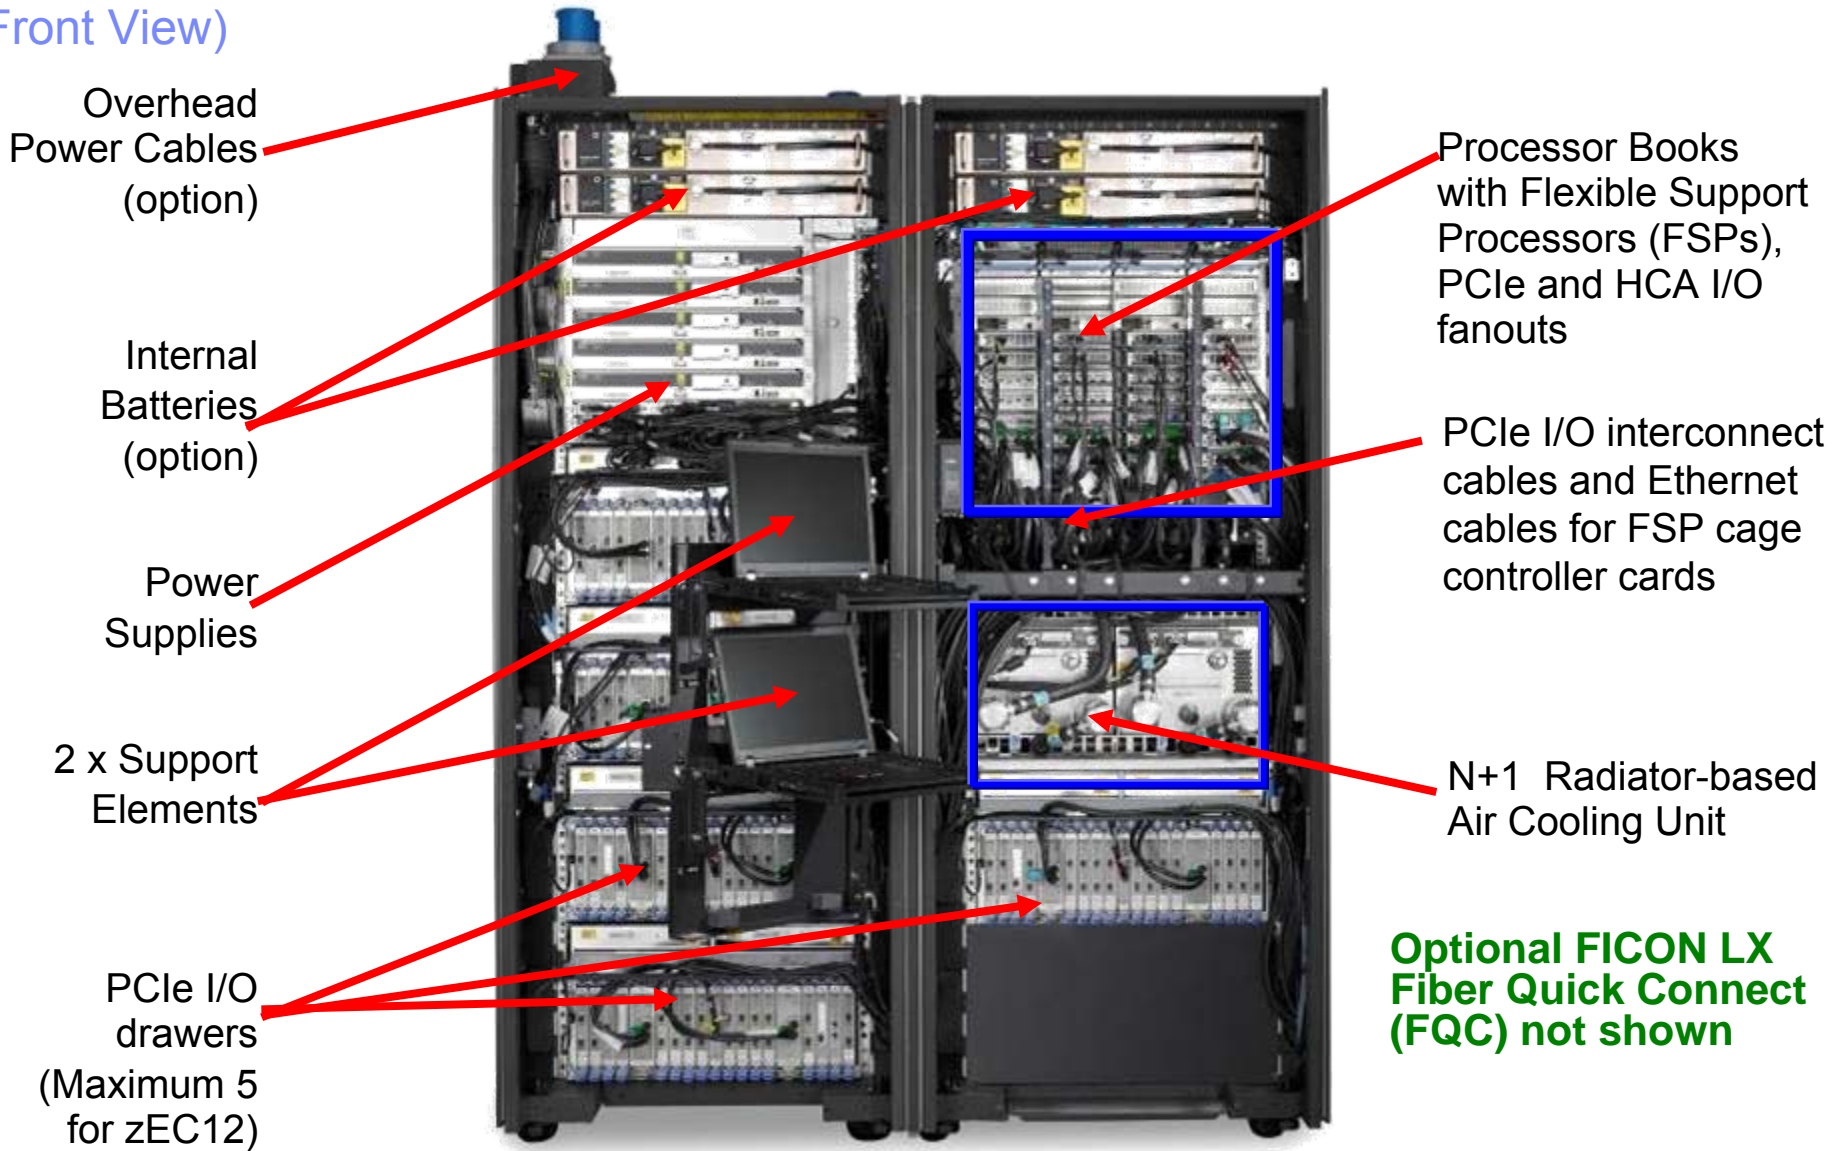


Model	Books/ PUs	CPs	IFLs uIFLs	<i>zAAPs</i>	<i>zIIPs</i>	ICFs	Std SAPs	Optional SAPs	Std. Spares	<i>IFP</i>
H20	1/27	0-20	0-20 0-19	0-13	0-13	0-20	4	0-4	2	1
H43	2/54	0-43	0-43 0-42	0-28	0-28	0-43	8	0-8	2	1
H66	3/81	0-66	0-66 0-65	0-44	0-44	0-66	12	0-12	2	1
H89	4/108	0-89	0-89 0-88	0-59	0-59	0-89	16	0-16	2	1
HA1	4/120	0-101	0-101 0-100	0-67	0-67	0-101	16	0-16	2	1

- ▶ zEC12 Models H20 to H89 use books with 27 core MCMs. The Model HA1 has 4 books with 30 core MCMs
- ▶ The maximum number of logical ICFs or logical CPs supported in a CF logical partition is 16
- ▶ **The integrated firmware processor (IFP), formerly reserved, is used for PCIe I/O support functions**
- ▶ Concurrent Book Add is available to upgrade in steps from model H20 to model H89

- Notes:
1. At least one CP, IFL, or ICF must be purchased in every machine
  2. Two zAAPs **and** two zIIPs may be purchased for each CP purchased if PUs are available.  
This remains true for sub-capacity CPs and for “banked” CPs.
  3. “uIFL” stands for Unassigned IFL
  4. The IFP is conceptually an additional, special purpose SAP

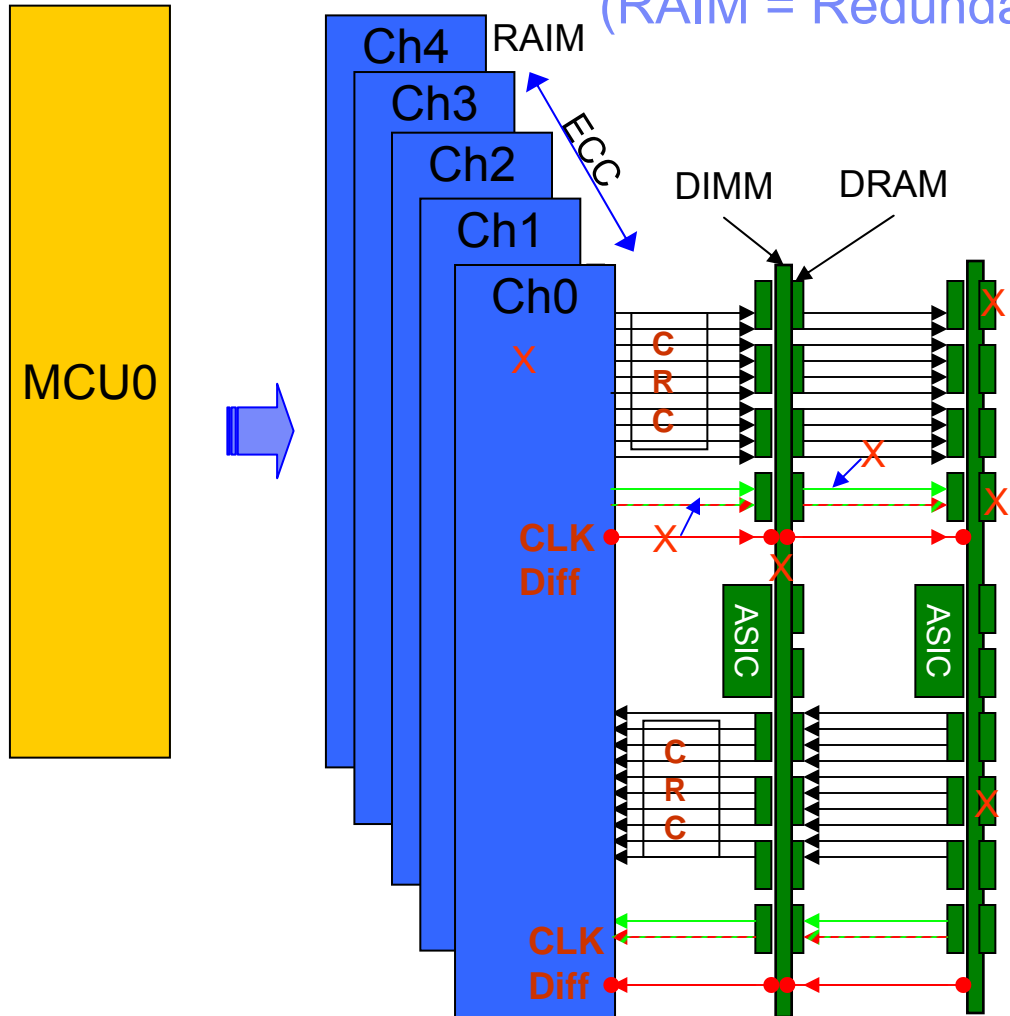
# zEC12 Model H89 or HA1 Radiator (Air) Cooled – Under the covers (Front View)



# zBC12 and zEC12 GA2 Memory



# zBC12 and zEC12 5-Channel RAIM Memory Controller Overview (RAIM = Redundant Array of Independent Memory)



## Layers of Memory Recovery

### ECC

- Powerful 90B/64B Reed Solomon code

### DRAM Failure

- Marking technology; no half sparing needed
- 2 DRAM can be marked
- Call for replacement on third DRAM

### Lane Failure

- CRC with Retry
- Data – lane sparing
- CLK – RAIM with lane sparing

### DIMM Failure (discrete components, VTT Reg.)

- CRC with Retry
- Data – lane sparing
- CLK – RAIM with lane sparing

### DIMM Controller ASIC Failure

- RAIM Recovery

### Channel Failure

- RAIM Recovery

**zEC12: Each channel supports one DIMM or a two deep DIMM cascade**

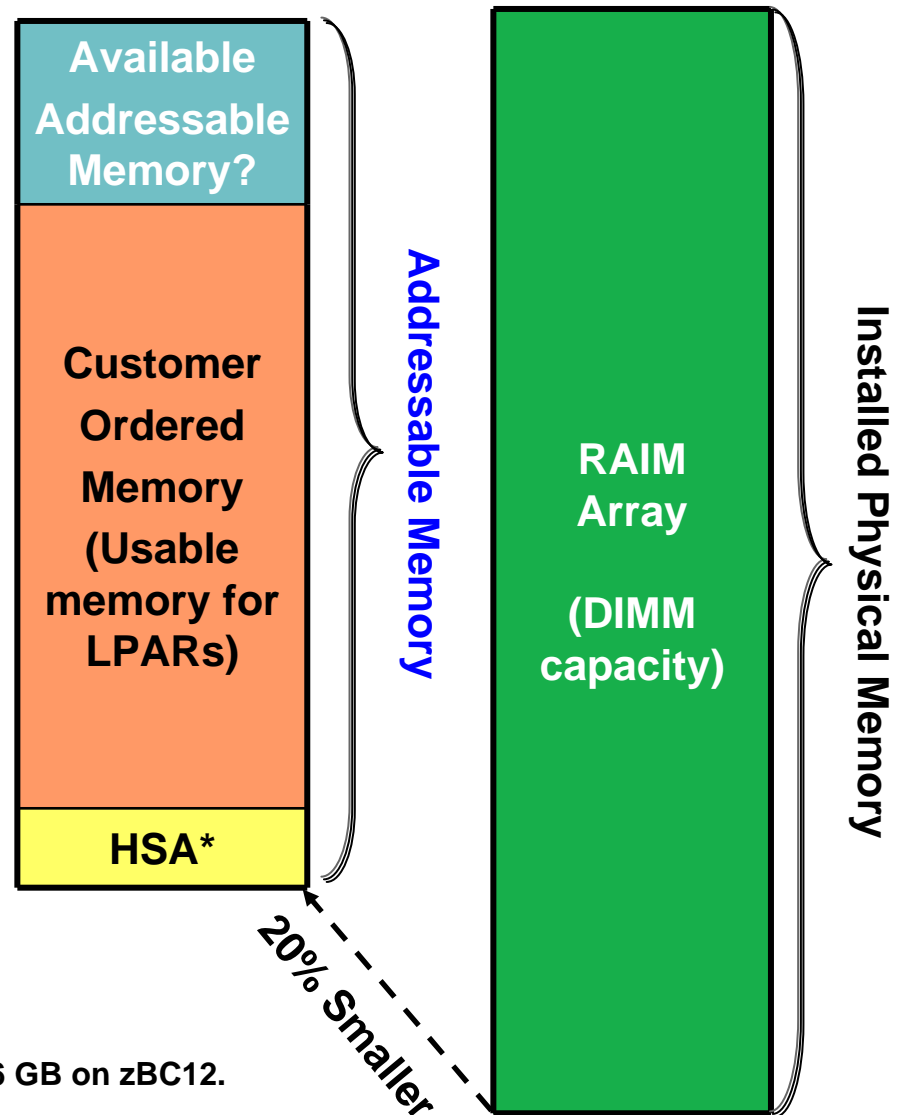
**zBC12: Each channel supports only one DIMM**

## zEC12 and zBC12 Memory Usage and Allocation

- Installed Physical Memory** (DIMM capacity) in configuration reports is **RAIM Array** size. **Addressable Memory** for customer LPARs and HSA is 20 percent smaller.
- Servers are configured with the most efficient configuration of memory DIMMs that can support **Addressable Memory required for Customer Ordered Memory plus HSA**. In some cases, there will be **Available Addressable Memory** that **might** support one or more concurrent LIC CC Customer Memory upgrades with no DIMM changes.
 

Note: DIMM changes require a disruptive POR on zBC12 and on zEC12 Model H20. They are always done without a POR on zEC12 models with multiple books using Enhanced Book Availability (EBA). On those models, some or all LPARs can continue to run during with one book out of service to have DIMMs changed or added. Probably all LPARs, if **Flexible Memory** is selected.
- To determine the size of the largest LIC CC Customer Memory upgrade possible, examine the configurator default **"Memory Plan Ahead Capacity"** field. If the customer requires a LIC CC upgrade larger than that, configure Plan Ahead Memory by selecting a larger "Memory Plan Ahead Capacity" target value.

\*HSA size is 32 GB on zEC12, 16 GB on zBC12.



# zBC12 and zEC12 Memory DIMMs and Plugging



## ▪ zBC12 Memory Plugging

- Ten DIMM slots per drawer supported by two memory controllers with five slots each
- All slots in a drawer must be populated
- Within a drawer all ten slots must be populated with the same size DIMM
- On the Model H13, the two processor drawers can be populated with different DIMM sizes

## ▪ zEC12 Memory Plugging

- Thirty DIMM slots per book supported by three memory controllers with ten slots each
- Within a book all populated slots must have the same size DIMM
- At least ten DIMM slots in each book must be populated
- After the first ten, DIMM slots are populated five at a time

## ▪ Maximum Client Memory Available

- Remember **RAIM** – 20% of DIMM memory is used only for error recovery
- Subtract HSA (zBC12 16 GB, zEC12 32GB) from addressable memory size and round **down** to an offered memory size

DIMM Size	zBC12 Feature (5 DIMMs) RAIM and Addressable Size	zEC12 Feature (5 DIMMs) RAIM and Addressable Size
4 GB	#1600 = 20 GB RAIM , 16 GB Addressable Memory	#1614 = 20 GB RAIM , 16 GB Addressable Memory
8 GB	#1601 = 40 GB RAIM , 32 GB Addressable Memory	Not supported, no feature
16 GB	#1603 = 80 GB RAIM, 64 GB Addressable Memory	#1615 = 80 GB RAIM, 64 GB Addressable Memory
32 GB	#1609 = 160 GB RAIM, 128 GB Addressable Memory	#1618 = 160 GB RAIM, 128 GB Addressable Memory

## zBC12 Model H06 and Model H13 Memory (New Build)



- zBC12 has the same RAIM Memory infrastructure as z114
- Minimum client memory is 8 GB for both models with 8 GB or 32 GB purchase increments
- HSA is 16GB on both models

H06 Physical Memory RAIM GB Addressable Memory GB	Client GB	Increment GB
40 RAIM (10 x 4GB)= 32 for HSA + Client	8 16	8
80 RAIM (10 x 8GB)= 64 for HSA + Client	24 32 40 48	8
160 RAIM (10 x 16 GB)= 128 for HSA + Client	56 64 72 80 88 96 104 112	8
320 RAIM (10 x 32 GB)= 256 for HSA + Client	144 176 208 240	32

H13 Physical Memory RAIM GB Addressable Memory GB	Client GB	Increment GB
80 RAIM (20 x 4GB)= 64 for HSA + Client	16 24 32 40 48	8
120 RAIM (10 x 4GB + 10 x 8 GB)= 96 for HSA + Client	56 64 72 80	8
160 RAIM (20 x 8 GB)= 128 for HSA + Client	88 96 104 112	8
240 RAIM (10 x 8 GB + 10 x 16 GB)= 192 for HSA + Client	144 176	32
320 RAIM (20 x 16 GB) = 256 for HSA + Client	208 240	32
480 RAIM (10 x 16 GB + 10 x 32 GB)= 384 for HSA + Client	272 304 336 368	32
640 RAIM (20 x 32 GB)= 512 for HSA + Client	400 432 464 496	32

Memory upgrades that require DIMM changes are **disruptive**. Plan Ahead Memory can be added to eliminate disruption. On both models, memory upgrades within each row (same color) are **concurrent** without adding Plan Ahead Memory.

## zEC12 Purchased Memory Offerings

Model	Standard Memory GB	Flexible Memory GB
H20	32 - 704	NA
H43	32 - 1392	32 - 704
H66	32 - 2272	32 - 1392
H89	32 - 3040	32 - 2272
HA1	32 - 3040	32 - 2272

- **Purchased Memory** - Memory available for assignment to LPARs
- **Hardware System Area** – Standard **32 GB** outside customer memory for system use
- **Standard Memory** - Provides minimum physical memory required to hold base purchase memory plus **32 GB** HSA
- **Flexible Memory** - Provides additional physical memory needed to support activation base customer memory and HSA on a multiple book zEC12 with one book out of service.  
(Not available on zBC12)
- **Plan Ahead Memory** – Provides additional physical memory needed for a concurrent upgrade (LIC CC change only) to a preplanned target customer memory



# zEC12 Standard and Flexible Purchase Memory Offerings

Increment	GB, Notes	Growth %	Increment	GB, Notes	Growth %	Increment	GB, Notes	Growth %
32 GB	32	100%	96 GB	608	16%	240 GB	1760	15%
	64	50%		704 <sup>1</sup>	13%			
	96	33%		800	12%	256 GB	2016	13%
	128	25%		896	12%		2272 <sup>3</sup>	11%
	160	20%					2528	10%
	192	17%	112 GB	1008	13%		2784	9%
	224	14%					3040 <sup>4</sup>	NA
	256	25%				<b>Notes – Memory Maximums:</b> 1. H20 Standard, H43 Flexible = 704 2. H43 Standard, H66 Flexible = 1392 3. H66 Standard, H89 and HA1 Flexible = 2272 4. HA1 Standard = 3040 (16 GB less than z196 above 1520 GB)		
64 GB	320	20%	128 GB	1136	11%			
	384	17%		1264	10%			
	448	14%		1392 <sup>2</sup>	9%			
	512	19%		1520	16%			

# LPAR Absolute Physical Capacity Setting



# LPAR Absolute Physical Capacity Limit (Image Profile)

**Customize Image Profiles: SSYS : SOSP0B : Processor**

- SSYS
  - SOSP0B
    - General
    - Processor
    - Security
    - Storage
    - Options
    - Load
    - Crypto

Group Name <Not Assigned>

---

*Logical Processor Assignments*

Dedicated processors

Select	Processor Type	Initial	Reserved
<input checked="" type="checkbox"/>	Central processors (CPs)	1	0
<input type="checkbox"/>	System z application assist processors (zAAPs)	0	0
<input type="checkbox"/>	System z integrated information processors (zIIPs)	0	0

---

*Not Dedicated Processor Details*

Initial processing weight  1 to 999  Initial capping

Enable workload manager

Minimum processing weight

Maximum processing weight

Absolute Capping

None

Number of processors (0.01 to 255.0)

- Physical capacity limited for shared logical processors with a granularity of in .01 processor, 50% of a CP in this case
- Operating System Support
  - z/OS V1.12 or later with PTFs
  - z/VM V6.3 with PTFs

Cancel
Save
Copy Profile
Paste Profile
Help

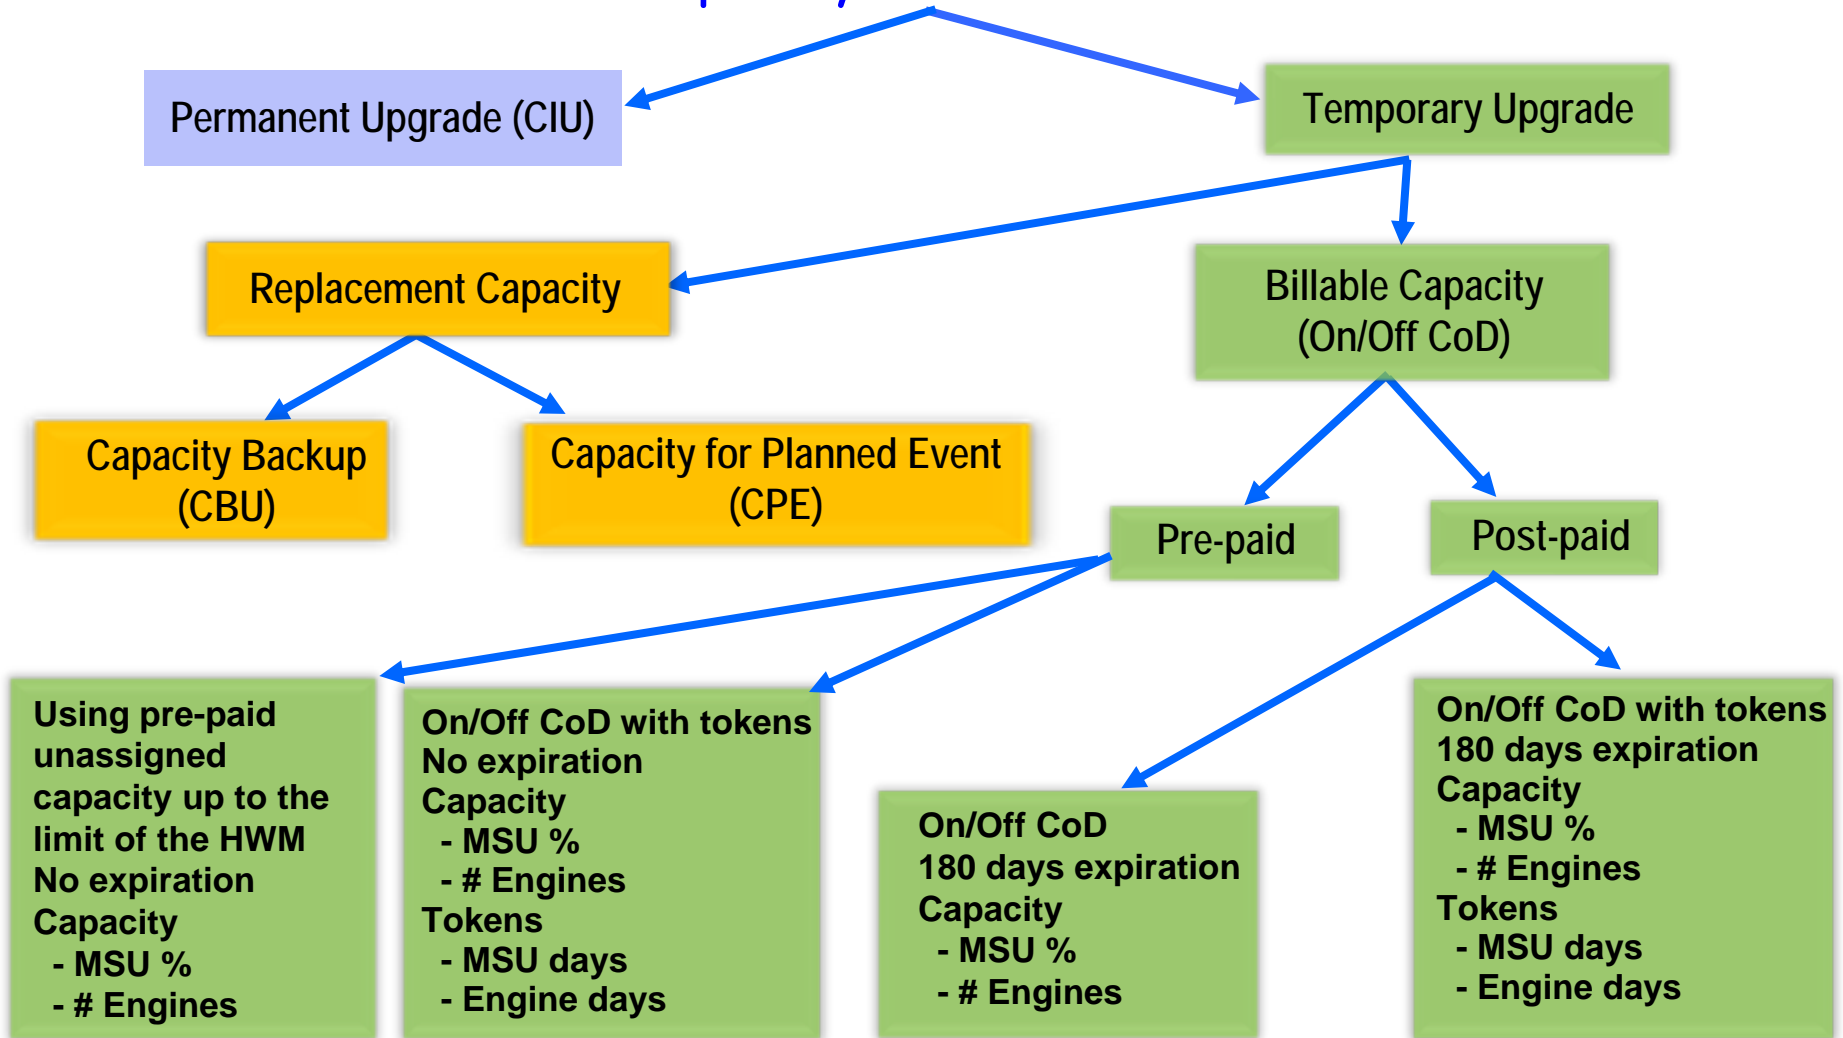
**Note: Screen capture is not GA level. Details may change.**

# On Demand



# zEC12 Basics of Capacity on Demand (Unchanged for zEC12 GA)

## Capacity on Demand



# Installation Planning



## zEC12 Physical Planning

- **Extend / Maintain z196 Datacenter Characteristics**

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

- **Better control of energy usage and improved efficiency in your data center**

- **Improved N+1 air cooled option with radiator**
- **New optional non-raised floor installation**
  - Target “on slab” low cost datacenters
- **Improved water cooled option**
  - **Backup air cooling**
  - **Supports building chilled water up to 20° C**
- Same number of power cords (2 or 4) as “equivalent” z196 configuration
- Maintain 27.5 kW box max input power (same as z10, z196)
- Maintain DC input power capability, overhead I/O cabling option, add overhead power option



## zBC12 Physical Planning

### ▪ **Extend / Maintain Datacenter Characteristics**

- zBC12 one frame system (air cooled, no radiator)
- No significant increase in weight
- Maintain floor tile cutouts for raised floor system (same as z10 BC or z114)

### ▪ **Better control of energy usage and improved efficiency in your data center**

- Same number of power cords (2 or 4) as an “equivalent” z114 configuration
- Maintain box max input power
- All power and I/O cables the same as z114 except:
  - 400V AC is an orderable feature (Was a z114 RPQ)
  - New 380-415v AC bottom and top exit
  - Different rule for non-raised floor than z114

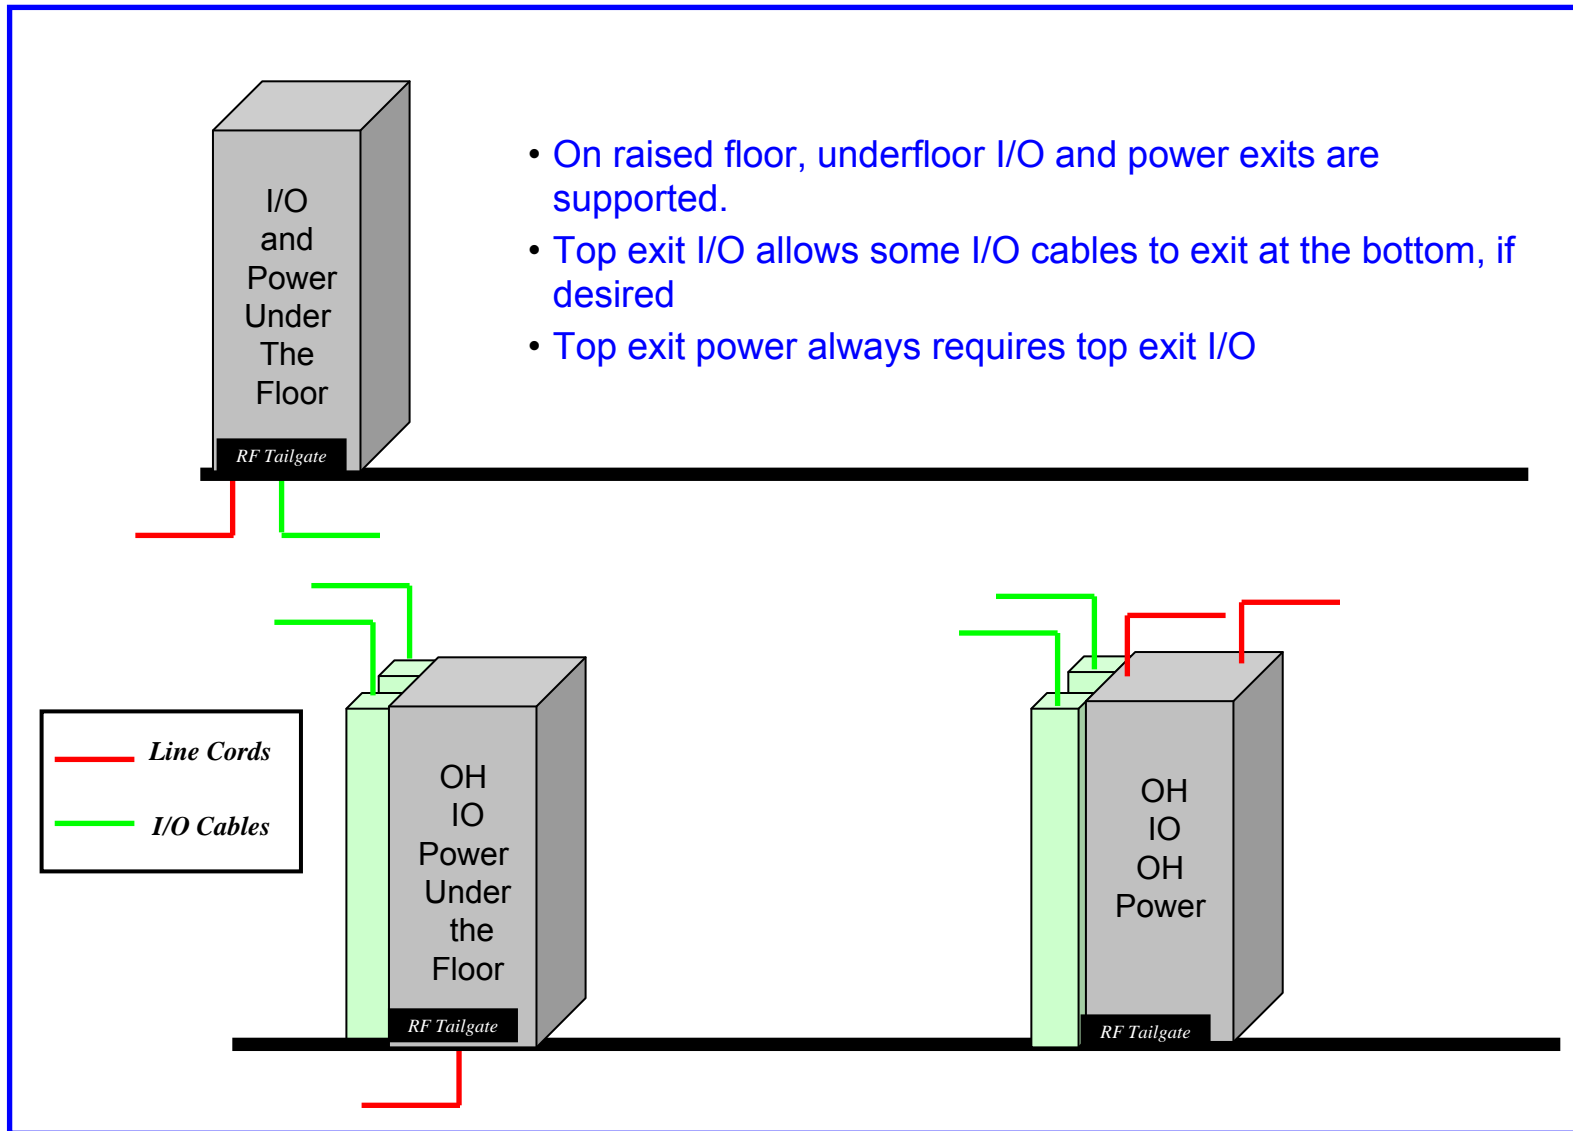


**Always Refer to the Installation Manual for Physical Planning for details:**

**M/T 2828 – GC28-6923**  
**M/T 2458 – GC27-2619-01 (Model 003)**

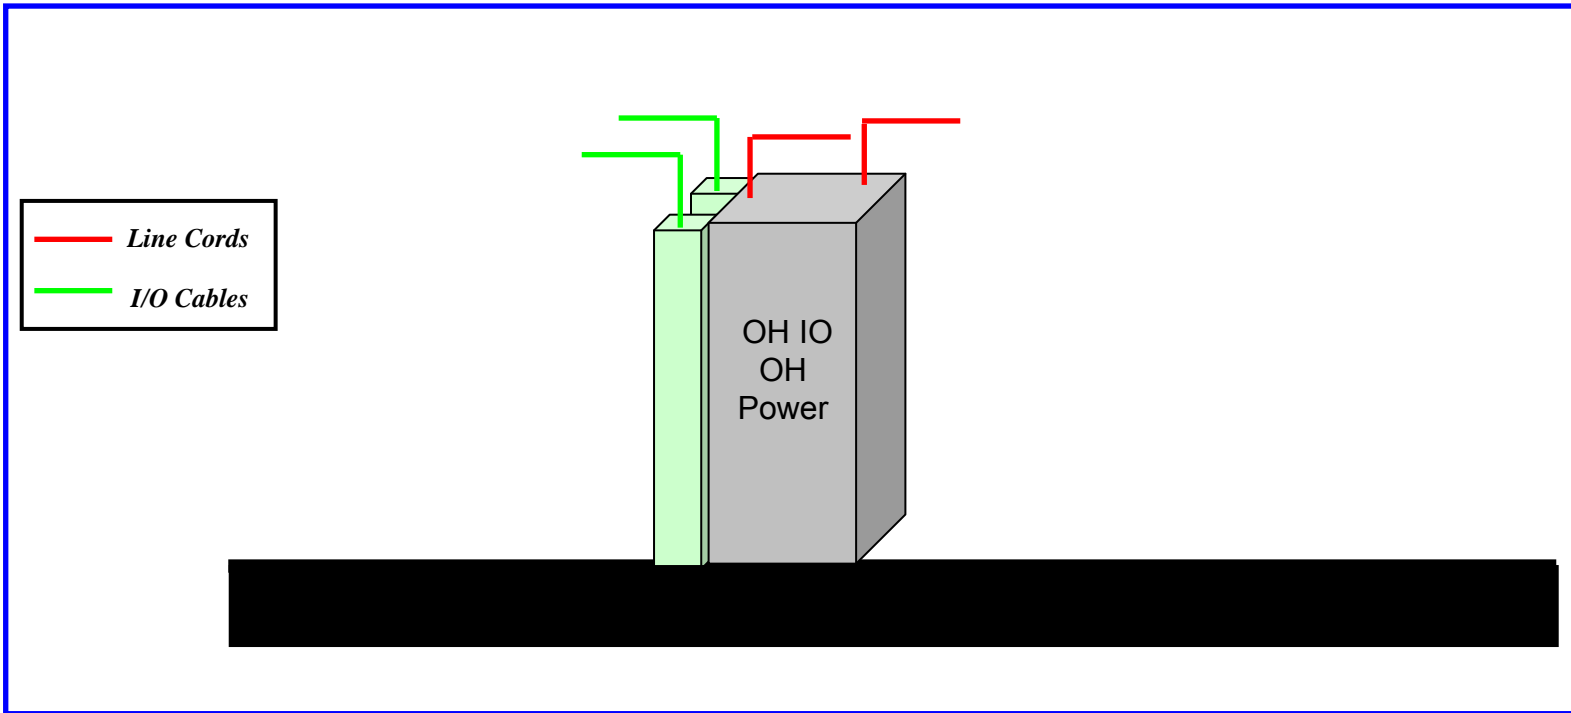


## zBC12 Installation Options on Raised Floor



- On raised floor, underfloor I/O and power exits are supported.
- Top exit I/O allows some I/O cables to exit at the bottom, if desired
- Top exit power always requires top exit I/O

# zBC12 Installation Option Non-Raised Floor



If zBC12 is NOT installed on a raised floor, overhead I/O and overhead power are REQUIRED. No cables may exit at the bottom because there is no tailgate.

Note: z114 allowed bottom I/O and bottom power on NRF so this is a change!

## zBC12 and zEC12 GA2 Hardware Management Console

- **HMC System Unit and LIC Support**

- Required: HMC FC 0092 or 0091 with 16 GB (FC 0091 HMCs shipped prior to zEC12 may have 8 GB) FC 0092 can be ordered new build. FC 0091 can carry forward on upgrade from z10 BC or z114
- A no-charge ECA orderable by IBM service is available to upgrade HMC FC 0091 features of an earlier System z server with more memory and to add HMC Driver 15 LIC for zBC12 or zEC12 at GA2
- Older HMCs (e.g FC 0090) can not be upgraded to control zEC12 or zBC12

- **HMC Display Support**

- 22 inch flat panel FC 6096 (Carry forward or new build)

- **HMC 1000BASE-T LAN Switch – No longer offered**

- FC 0070 10/100/1000BASE-T switch(es) – (Carry Forward Only)
- Recommended Alternative: Compatible customer provided 1000BASE-T switch(es)

- **HMC application LIC for zEC12 and zBC12 does NOT support dial modem use**

(Fulfills the Statement of Direction in Letter 111-167, dated October 12, 2011)

- Use of Broadband (Ethernet) access to RSF is required
- Optional connection to an NTP time source requires use of Ethernet, not dial (If a Pulse-per-Second time source is used, Ethernet is also required)
- If a modem is present on an HMC FC 0091 with the HMC Application in Driver 12 (zEC12 GA) or 15 (zBC12 or zEC12 GA2) installed, it can not be used

- **See the zBC12 or zEC12 Library on Resource Link for the latest publications:**

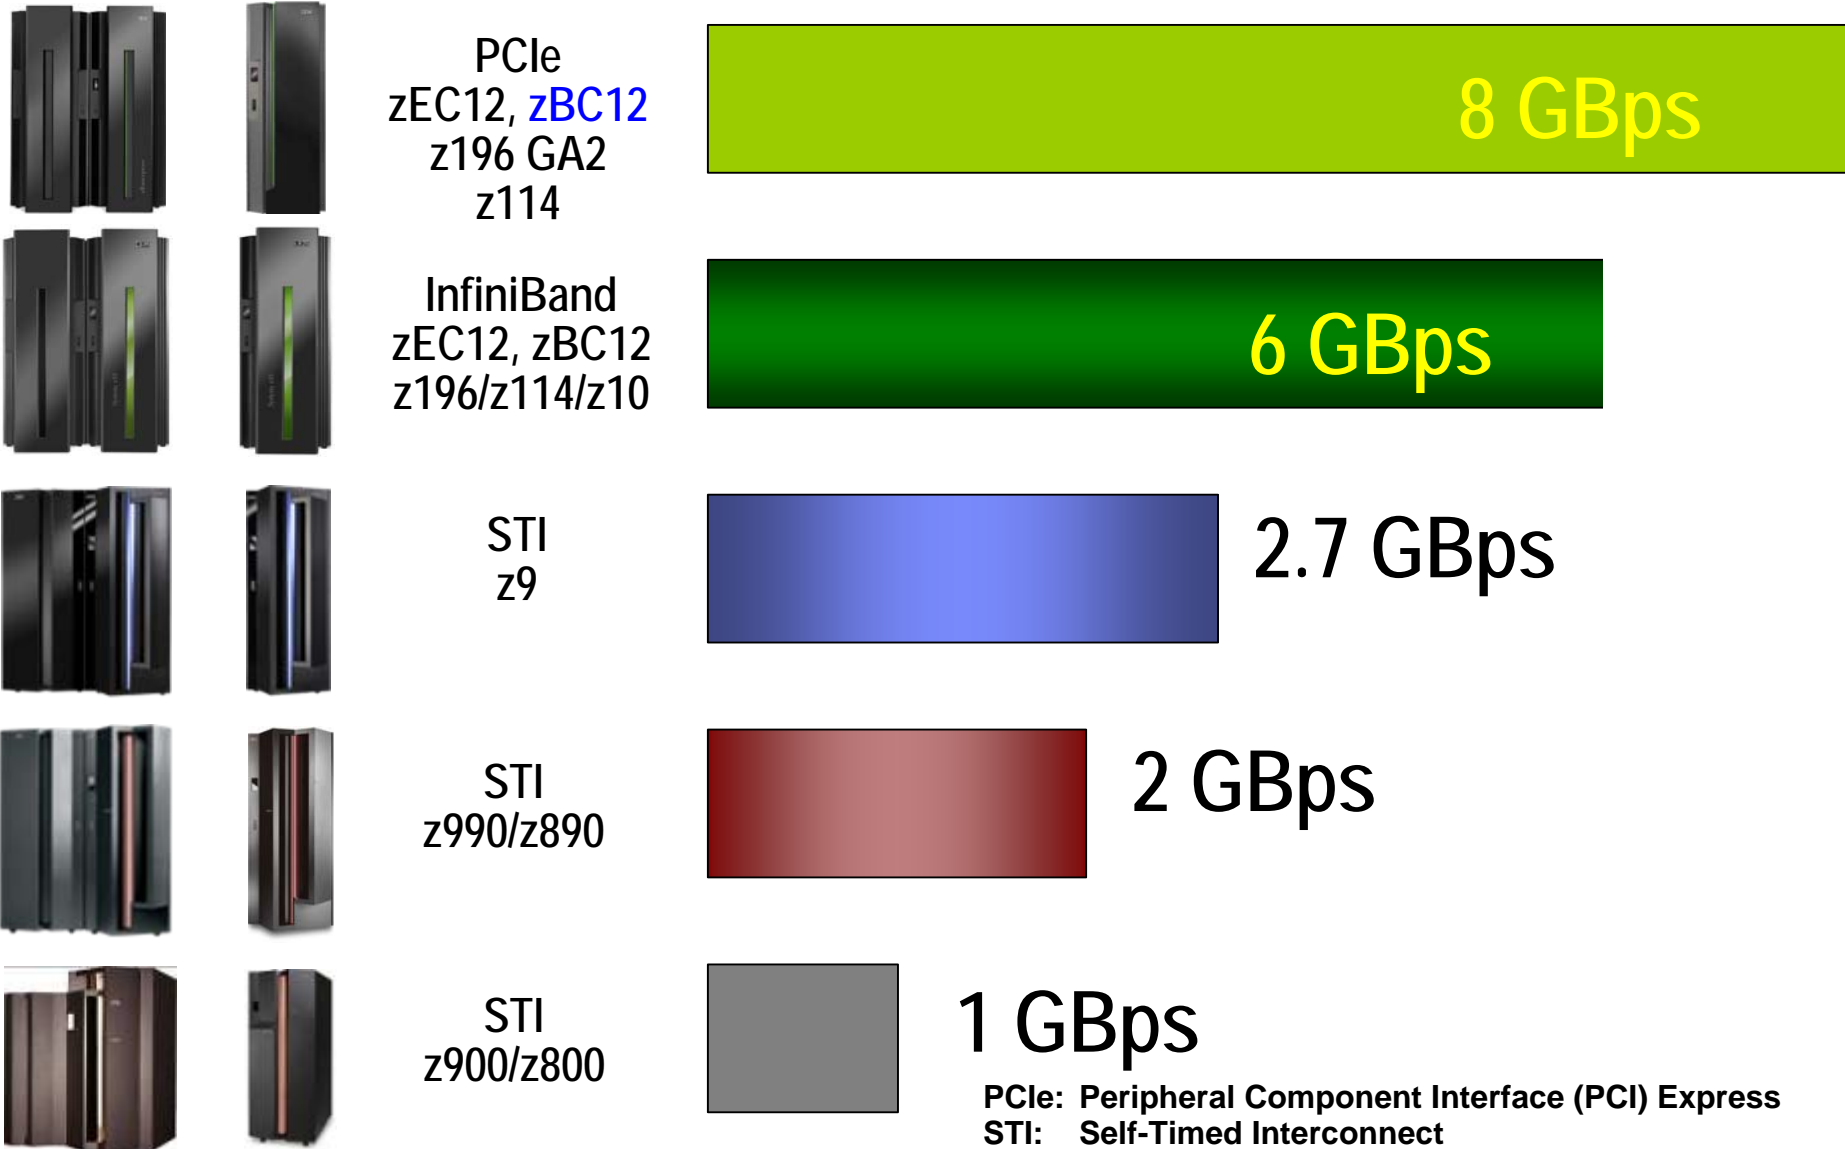
- “Installation Manual for Physical Planning” for HMC FC 0092 feature physical characteristics
- “Integrating the HMC Broadband RSF into your Enterprise”
- “Hardware Management Console Operations Guide” and “Support Element Operations Guide”

# zBC12 and zEC12 GA2 I/O Subsystem Introduction

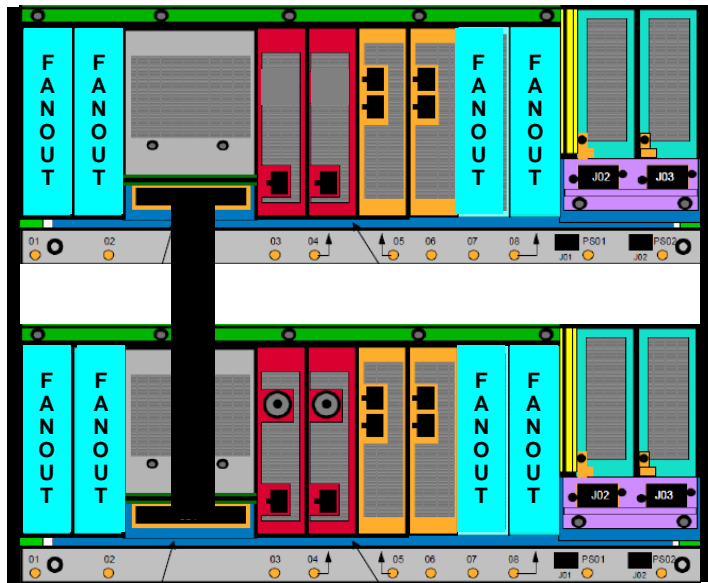
For more detail on I/O, plan to attend **Session 13706**,  
**9:30 AM Wednesday in Hynes Room 312**



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



# zBC12 Connectivity for I/O and Coupling

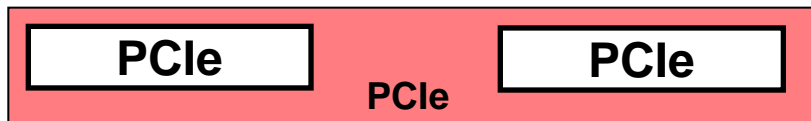


Drawer 2  
H13 only

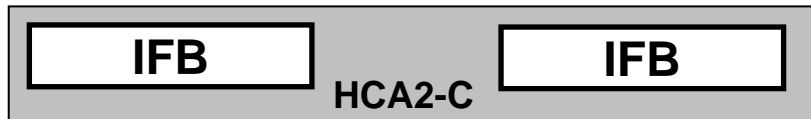
Drawer 1  
H06 and H13

- **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**

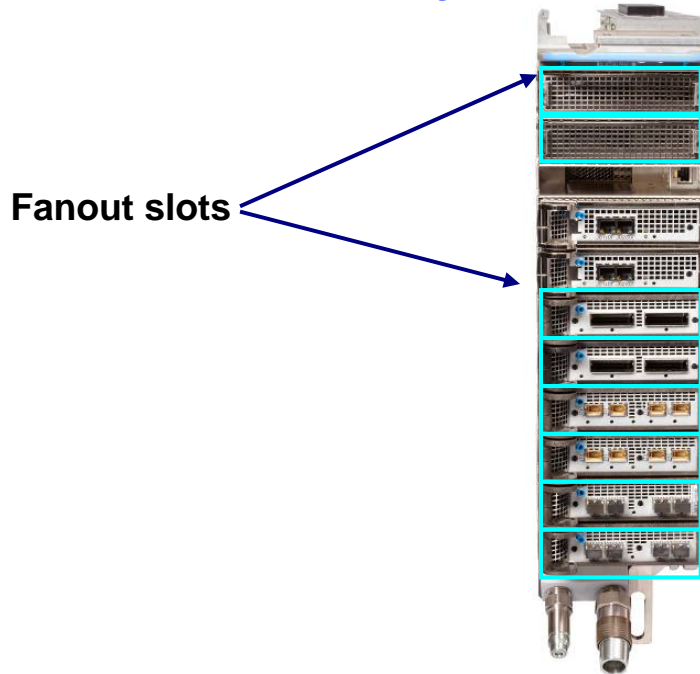


- **PCIe fanout – PCIe I/O Interconnect links**  
Supports two PCIe 8 GBps interconnects on copper cables to two 8-card PCIe I/O domain switches.  
**Always plugged in pairs for redundancy.**

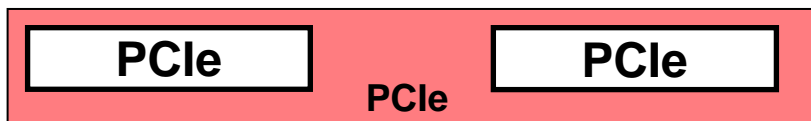


- **HCA2-C fanout – InfiniBand I/O Interconnect (CF only)**  
Supports two 12x InfiniBand DDR 6 GBps interconnects on copper cables to two 4-card I/O domain multiplexers.  
**Always plugged in pairs for redundancy.**

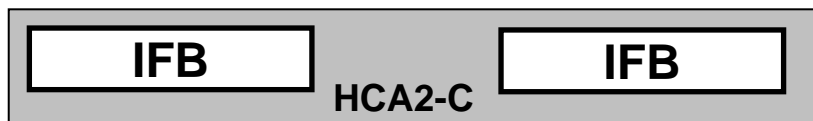
## zEC12 Connectivity for I/O and Coupling



- **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**
  -

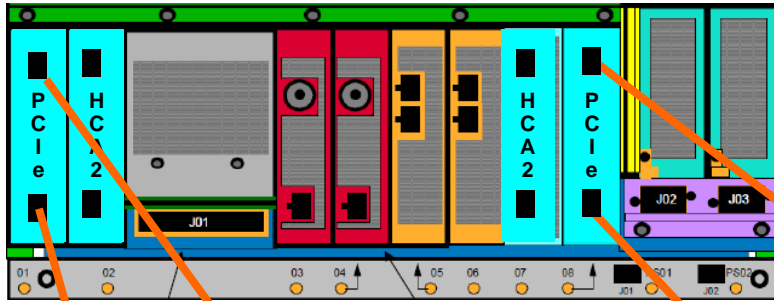


- **PCIe fanout – PCIe I/O Interconnect links**  
Supports two copper cable PCIe 8 GBps interconnects to two 8-card PCIe I/O domain multiplexers.  
**Always plugged in pairs for redundancy.**



- **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.**

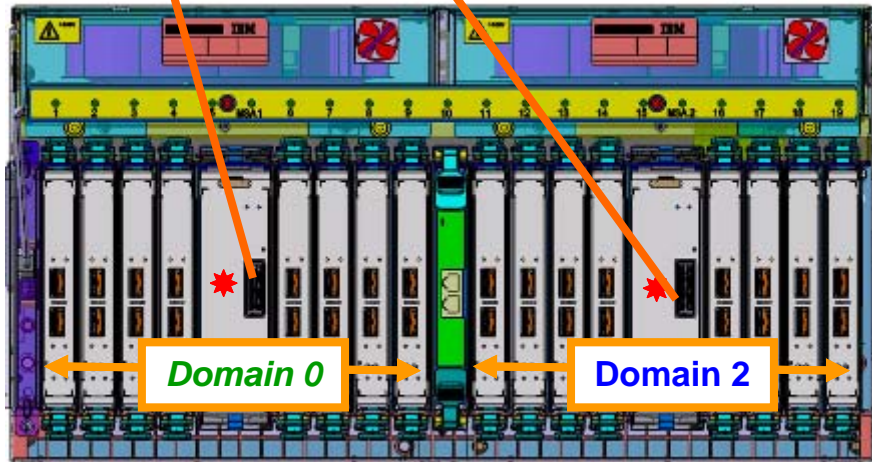
# zBC12 Redundant I/O Interconnect Example – One PCIe I/O drawer



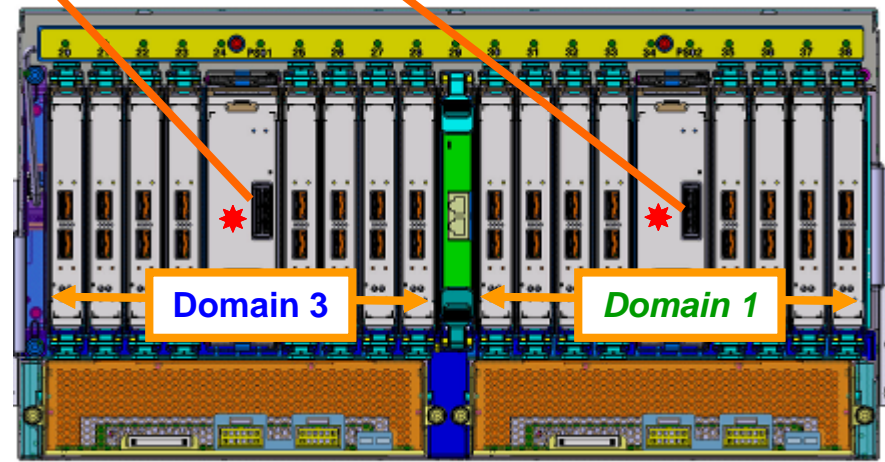
- Different PCIe Fanouts Support Domain Pairs:
  - 0 and 1
  - 2 and 3
- Normal operation each PCIe interconnect in a pair supports the eight I/O cards in its domain.
- Backup operation: One PCIe interconnect supports all 16 I/O cards in the domain pair.

PCIe switch cards (\*)

Front



Rear





# Supported I/O Features



## 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.”

# 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



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

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

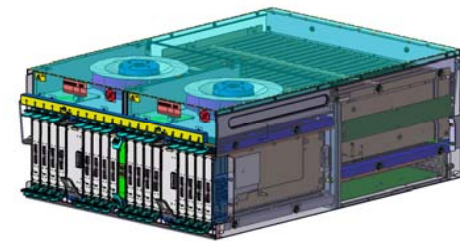


## New Build Features

### ■ 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

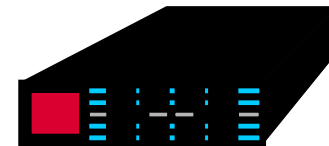
PCIe I/O drawer



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.

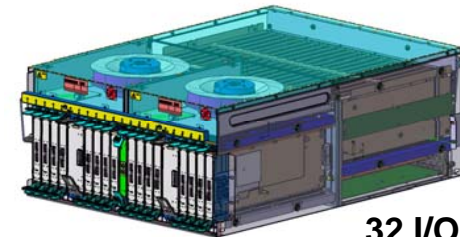


## Carry Forward Features

### ■ Features – PCIe I/O drawer

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

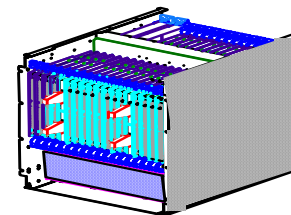
PCIe I/O drawer



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

# THANK YOU



## Session 13707

Please complete session evaluations online at [SHARE.org/BostonEval](http://SHARE.org/BostonEval)



# Technical Backup Charts



# zEC12 and zBX Model 3 New GA2 Functions and Features



## Processor, Memory, RAS

LPAR absolute physical capacity setting

## Ensemble, Platform Management

Unified Resource Manager support for ensembles with zEC12, zBC12, zBX Model 3, z196, z114, and zBX Model 2

Unified Resource Manager enhancements including KVM hypervisor “c groups” exploitation for virtual servers on System x and Ensemble Availability Manager (EAM) monitoring, error analysis, and CPU share policy goal management

Refreshed LIC for zBX TOR switches, BladeCenter modules, and Blades

Upgraded POWER7 and System x blade Hypervisor Levels

Continued incremental improvements



**Upgradeable from z10 EC and z196**

## Parallel Sysplex, Security, I/O

IBM zEnterprise Data Compression (zEDC): zEDC Express feature

10 GbE Converged Ethernet (CE) with Remote Direct Memory Access (RDMA) support: 10 GbE RoCE Express feature

Integrated firmware processor (IFP) for zEDC Express and 10GbE RoCE Express Support Functions

FICON – 24k devices per channel

OSA-Express5S 10 GbE, GbE, and 1000BASE-T with SFP transceivers

Security: Crypto PKCS 11 (EP11) mode stage 2 and CCA 4.4 enhancements

Trusted Key Entry (TKE) 7.3 firmware

CFCC Level 19 – Thin Interrupt Support to enable production use of shared logical processors in some cases

CFCC Level 19 – MQ Shared Queue overflow Flash Express Exploitation

STP – Improved SE time accuracy



# zBC12 Offering Plan Content Summary

<b>Machine Type 2828</b>
2 Models – H06, H13 Single frame, air cooled Non-raised floor option available Overhead Cabling and DC Power Options

<b>Base Content</b>
<b>Processors</b> 13 available PUs Up to 6 CPs at 26 capacity points, 50 to 4958 MIPS zIIP/zAAP/IFL/ICF/IFP (integrated firmware processor) Up to the equivalent of 8733 MIPS with 13 IFL's Uni- Processor Capacity 1064 MIPS 9 PU cores (using 4 and 5 core PU SCMs) per drawer (One for H06 and two for H13) 2 SAPs per system, standard 2 spares designated for Model H13
<b>Memory</b> System minimum = 16 GB 16 GB HSA separately managed RAIM standard Maximum 496 GB (model H13) Increments of 8 or 32 GB
<b>I/O</b> Up to 8 Legacy Channel Cards (16 via RPQ) Up to 64 PCIe Channel Cards Concurrent Add, Remove & Replace of IO drawers
<b>Coupling</b> Up to 16 12x PSIFB ports (model H13) Up to 32 1x PSIFB ports (model H13)



<b>Extended Content</b>
<b>From zEC12 GA1:</b> Crypto Express4S IBM zAware Flash Express
<b>IO enhancements</b> zEDC Express 10GbE RoCE Express OSA-Express5S
<b>zManager Enhancements</b> CPU management for xBlades support Ensemble Availability Manager
<b>Security Enhancements</b> Crypto EP11 enhancements
<b>Coupling Facility Enhancements</b> CFCC Flash Exploitation Coupling Thin Interrupts
<b>Migration</b>
Disruptive upgrade from H06 to H13 Upgrades from z10 BC and z114 Upgrades to zEC12 H20, radiator-based air cooled only, from zBC12 H13

# 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 <sup>(1)</sup>	Up to 3 <sup>(1)</sup>
I/O slots per I/O drawers/ PCIe I/O drawers	8	8/32		8/32 <sup>(2)</sup>	
FICON® Channels	128	128		128 <sup>(3)</sup>	
OSA Ports	96	96		96	
ESCON® Channels	480	240		0 <sup>(4)</sup>	
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 <sup>(8)</sup> /PSIFB	12/48/12	0/48/8 -16	0/48/16 - 32	0 <sup>(5)</sup> /32/8 -16 <sup>(6)</sup>	0 <sup>(5)</sup> /32/16 - 32 <sup>(7)</sup>
zIIP/zAAP Maximum Qty	5	2	5	Depends on CPs	Depends on CPs
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

\* 2<sup>nd</sup> 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

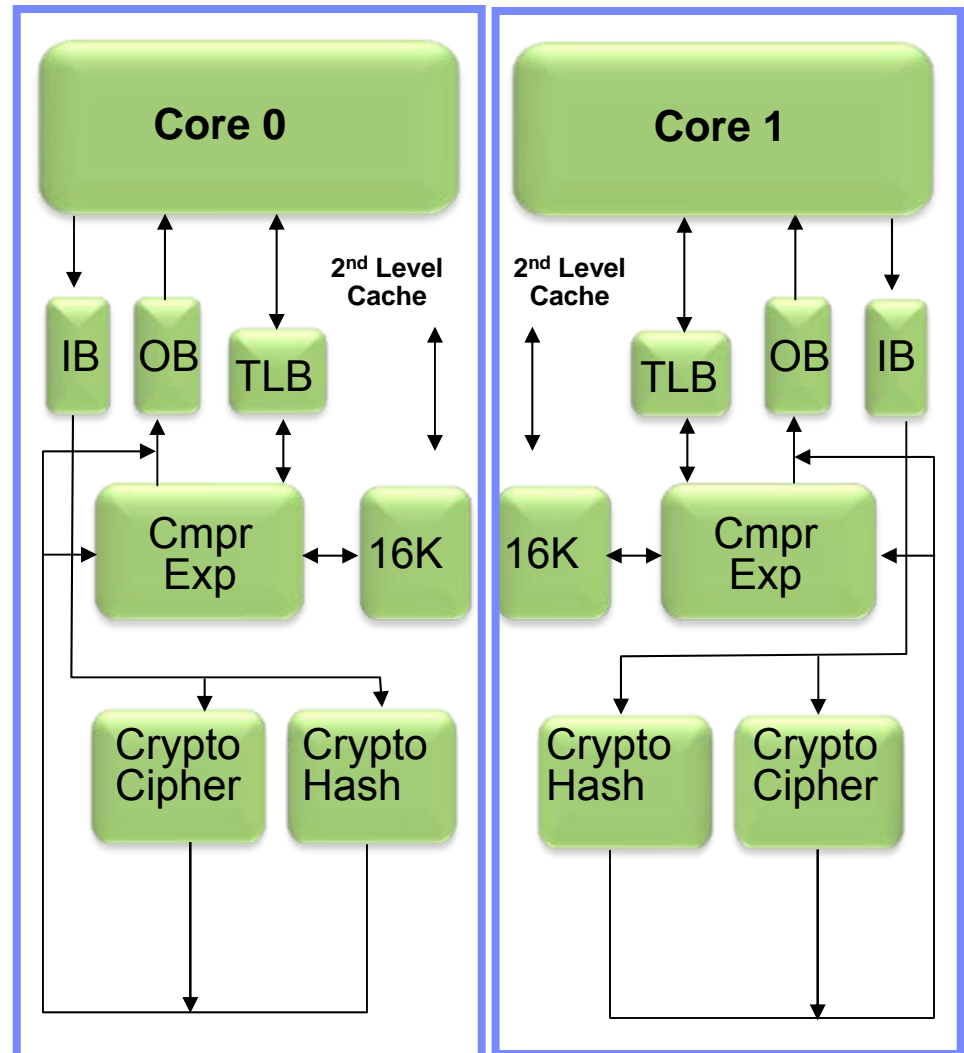
# IBM System z Config Comparisons, zBC12 vs. zEC12 Model H20

	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 <sup>(3)</sup>		256
OSA Ports (10GbE/1GbE/1000BASE-T)	48/96/96		48/96/96
ESCON® Channels	0 <sup>(4)</sup>		0 <sup>(4)</sup>
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 <sup>(8)</sup> /PSIFB	0 <sup>(5)</sup> /32/8 - 16 <sup>(6)</sup>	0 <sup>(5)</sup> /32/16 - 32 <sup>(7)</sup>	0 <sup>(5)</sup> /48/16 - 32
zIIP/zAAP Maximum Qty	Depends on CPs	Depends on CPs	Depends on CPs
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

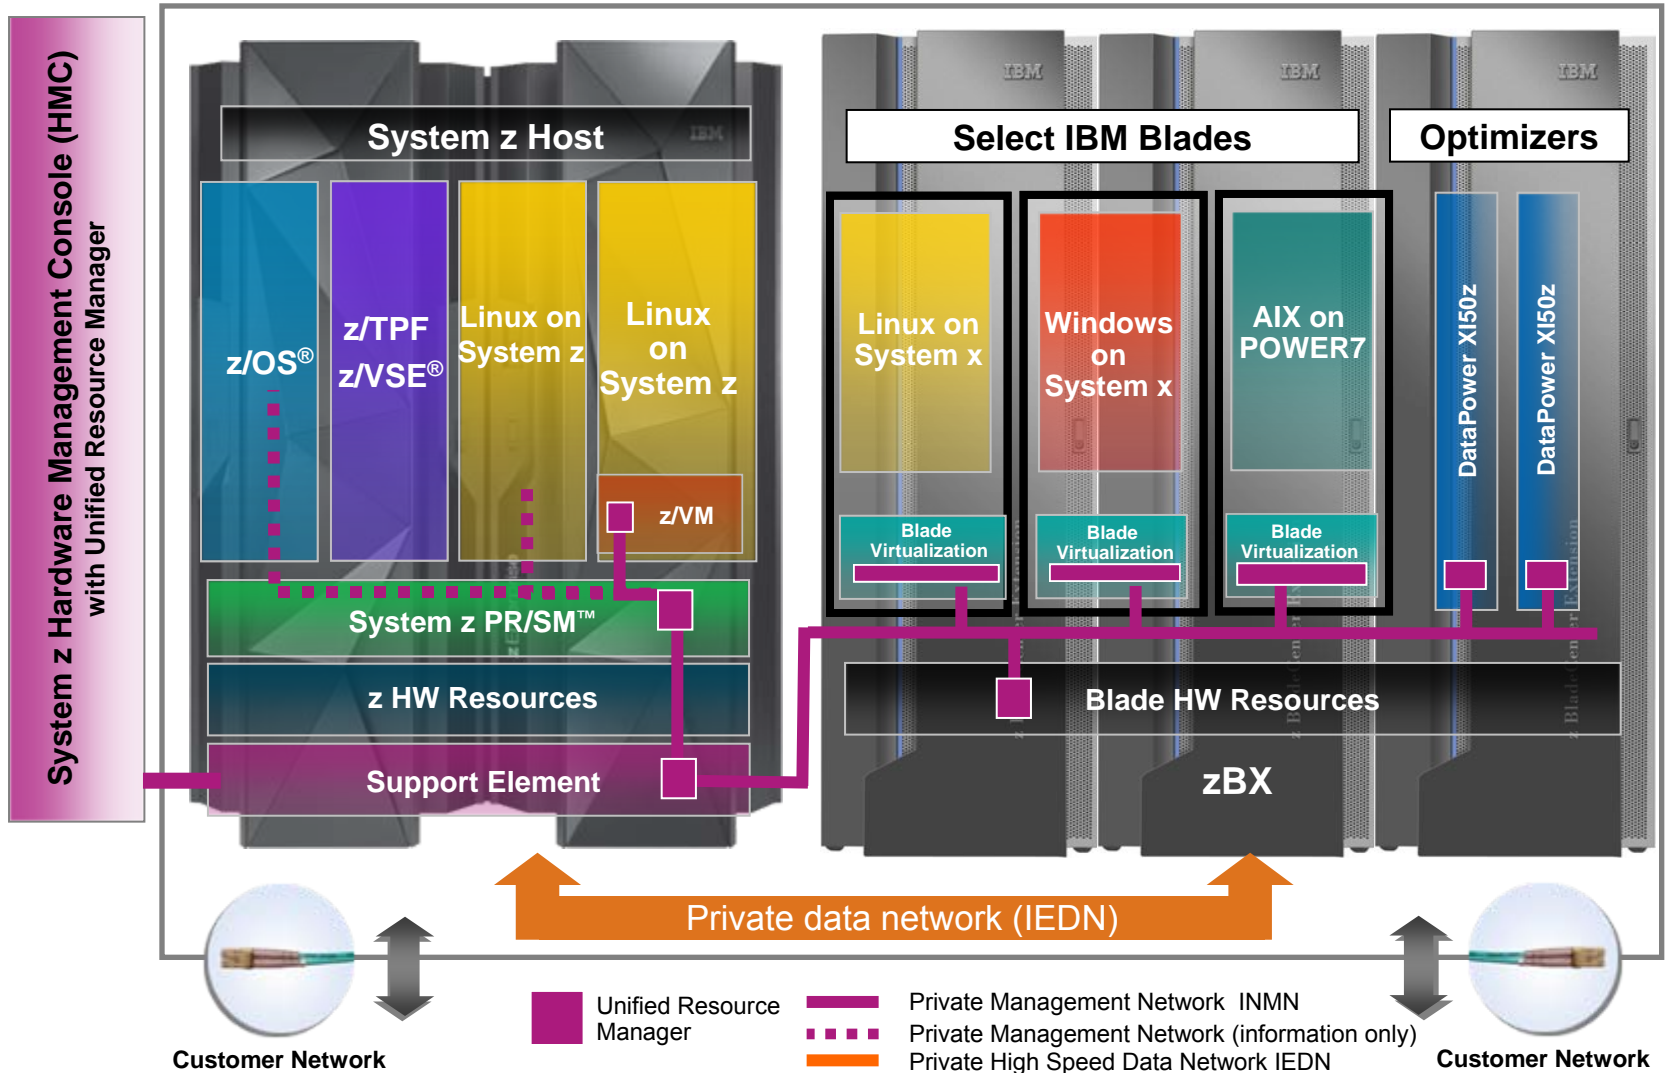
# zEC12 & zBC12 Compression and Cryptographic Coprocessor

- **Coprocessor dedicated to each core (Was shared by two cores on z196, z114)**
  - Independent compression engine
  - Independent cryptographic engine
  - Available to any processor type
  - Owing processor is busy when its coprocessor is busy
- **Data compression/expansion engine**
  - Static dictionary compression and expansion
- **CP Assist for Cryptographic Function**
  - DES (DEA, TDEA2, TDEA3)
  - SHA-1 (160 bit)
  - SHA-2 (244, 256, 384, 512 bit)
  - AES (128, 192, 256 bit)
  - CPACF FC #3863 (No Charge) is required to enable some functions and is also required to support Crypto Express4S or Crypto Express3 features



# Putting zEnterprise System to the Task

*Use the smarter solution to improve your application design*



# Delivering to Smarter Computing with zEnterprise

## Hybrid computing on System z keeps getting smarter

### Vision:

- zEnterprise is a hybrid system and always will be
- zEnterprise will continue to invest in improving the virtualization and management capabilities
- zEnterprise will seamlessly integrate with emerging Cloud-based IT management paradigms
- zEnterprise will more tightly integrate with STG Portfolio over time
- zEnterprise and STG will continue to leverage the Tivoli® portfolio to deliver enterprise wide management capabilities across all STG systems including PureSystems

### Current Focus

- Monitor and management of heterogeneous resources.
- DataPower XI50z appliance can help simplify, govern, and enhance the security of XML, Web and IT services
- Support for AIX on POWER7, and Linux and Windows on System x.
- Extending mgt functions of Unified Resource Manager with APIs
- Support of zBX on zEC12
- Enhanced Tivoli monitoring, discovery and availability/automation for zBX resources
- Image mgmt support for x86 and power blades though ISD

### Delivery of new hybrid function

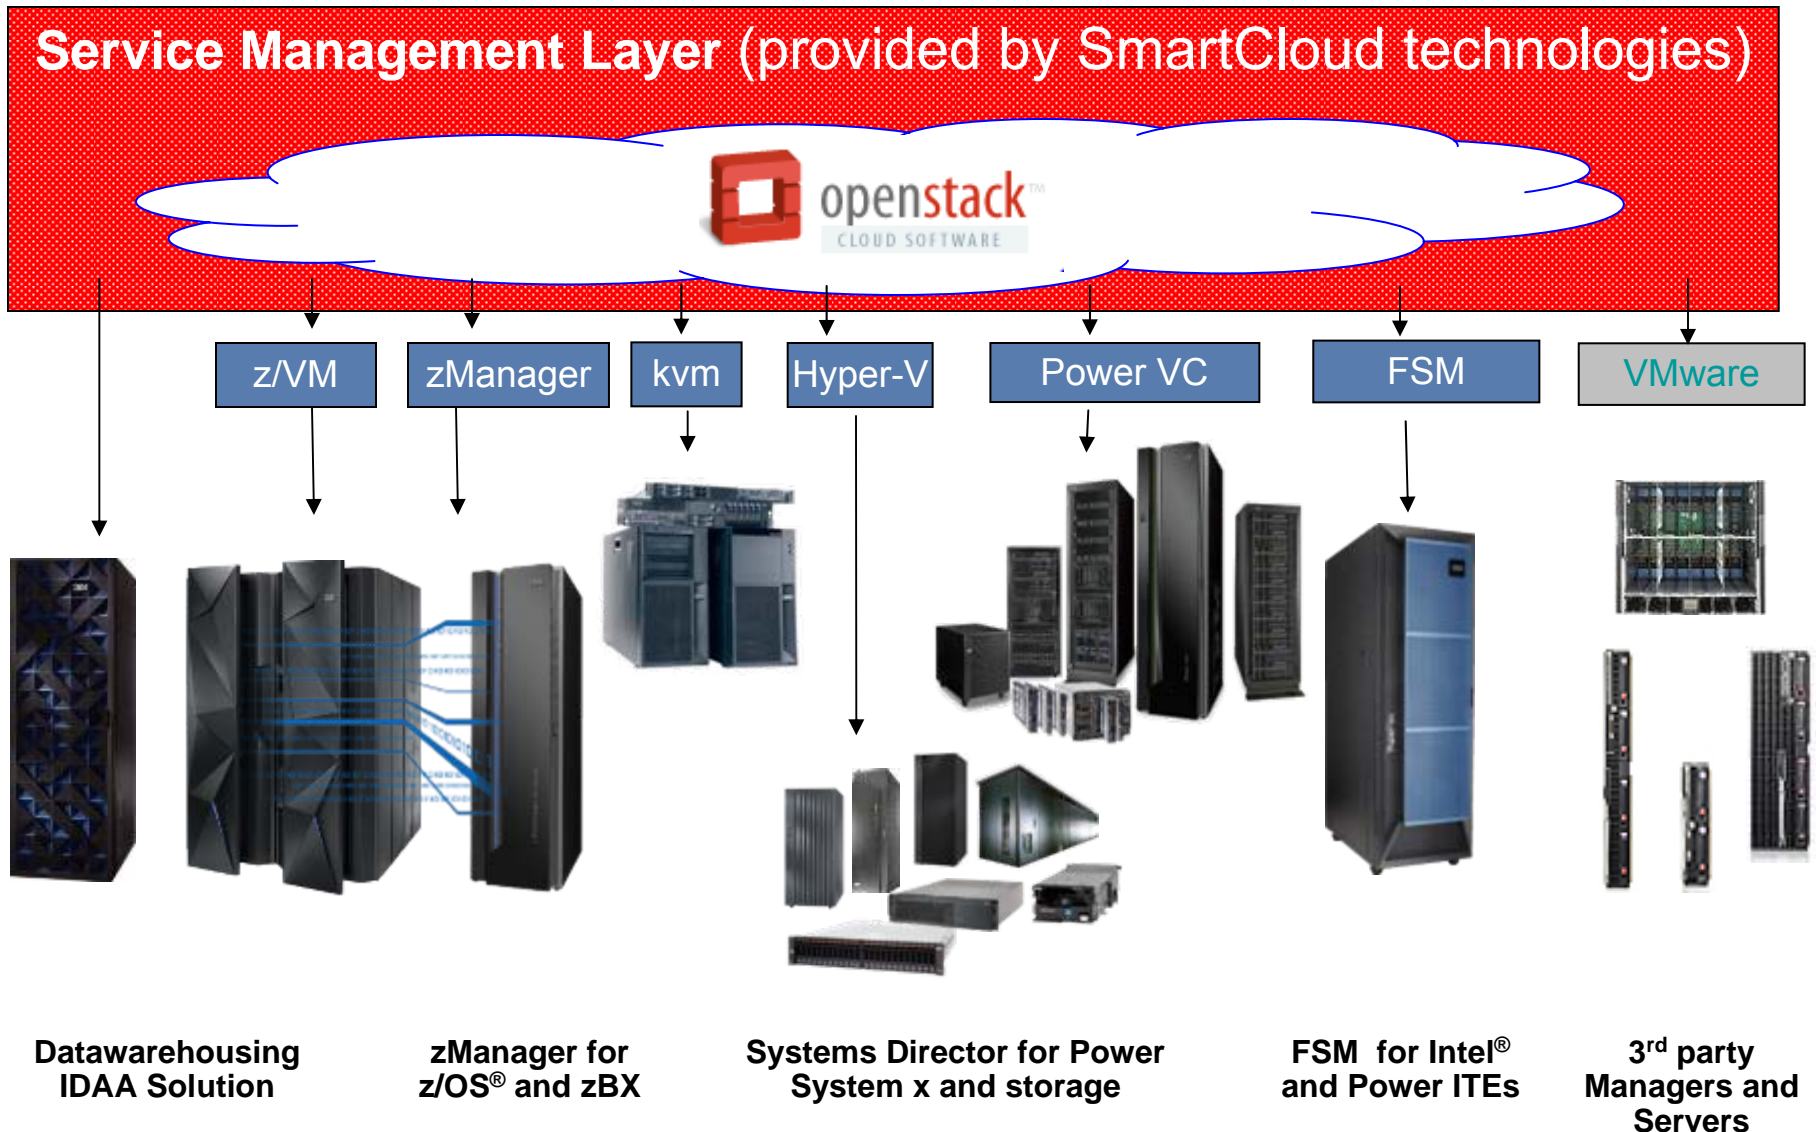
- CPU Management for x86 blades (Automate Bundle)
- GDPS automated site recovery for zBX
- Support of next generation hardware technologies in the zBX
- Ensemble Availability Manager – monitoring and reporting
- Support for additional versions of Windows Server and Linux
- zBX firmware currency
- Stand-alone zBX node
- Support for next generation DataPower Appliance

### Future Vision

- Support of zBX with the next generation server
- Technology configuration extensions in the zBX
- CPC and zBX
  - Continued investment in proving the virtualization and management capabilities for hybrid computing environment
  - Enablement of Infrastructure as a Service (IAAS) for Cloud
- Unified Resource Manager Improvements and extensions
  - Guest Mobility
  - Monitoring Instrumentation
  - Autonomic Management functions
- Integration with STG Portfolio
- zEnterprise and STG will continue to leverage the Tivoli portfolio to deliver enterprise wide mgmt capabilities across all STG systems

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

# Hybrid computing model integrated and enabled for Cloud





# I/O and Coupling Overview



# zEC12 I/O Feature Cards at GA2



Features	Offered As	Maximum # of features	Maximum channels/adapters	Increments per feature	Purchase increments
<b>FICON</b>					
<b>FICON Express8S</b>	<b>NB</b>	160	320 channels	2 ports/channels	2 channels
FICON Express8	CF <sup>1</sup>	44	176 channels	4 ports/channels	<b>CF Only</b>
FICON Express4 10km LX, SX	CF <sup>1</sup>	44	176 channels	4 or 2 ports/channels	<b>CF Only</b>
<b>ISC-3 Coupling</b>	CF <sup>1</sup>	12	48 links	1, 2, 3 or 4 links	<b>CF Only</b>
<b>OSA-Express</b>					
<b>OSA-Express5S<sup>4</sup></b>	<b>NB</b>	48	96 ports	1 (10 GbE) or 2 ports	1 feature/channel
OSA-Express4S	CF <sup>1</sup>	48	96 ports	1 (10 GbE) or 2 ports	<b>CF Only</b>
OSA-Express3	CF <sup>1</sup>	24	96 ports	2 (10 GbE) / 4 ports	<b>CF Only</b>
<b>Crypto</b>					
<b>Crypto Express4S<sup>4</sup></b>	<b>NB</b>	16	16 coprocessors	1 coprocessor	1 feature <sup>2</sup>
Crypto Express3 <sup>***</sup>	CF <sup>1</sup>	8	16 coprocessors	2 coprocessors	<b>CF Only</b>
<b>Special Purpose</b>					
<b>10 GbE RoCE<sup>4</sup></b>	<b>NB</b>	16	16 usable SR ports	1 usable SR port	1 feature <sup>3</sup>
<b>zEDC Express<sup>4</sup></b>	<b>NB</b>	8	8 coprocessors	1 coprocessor	1 feature <sup>3</sup>
<b>Flash Express<sup>4</sup></b>	<b>NB</b>	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
<b>FICON</b>					
<b>FICON Express8S</b>	<b>NB</b>	64	128 channels	2 ports/channels	2 channels
FICON Express8	CF <sup>1</sup>	8, 16 RPQ*	64 channels	4 ports/channels	<b>CF Only</b>
FICON Express4 10km LX, SX	CF <sup>1</sup>	8, 16 RPQ*	64 channels	4/2 ports/channels	<b>CF Only</b>
<b>ISC-3 Coupling</b>	CF <sup>1</sup>	8, 12 RPQ*	48 links	1, 2, 3 or 4 links	<b>CF Only</b>
<b>OSA-Express</b>					
<b>OSA-Express5S<sup>4</sup></b>	<b>NB</b>	48	96 ports	1 (10 GbE) or 2 SFPs	1 feature/channel
OSA-Express4S	CF <sup>1</sup>	48	96 ports	1 (10 GbE) or 2 ports	<b>CF Only</b>
OSA-Express3	CF <sup>1</sup>	8, 16 RPQ*	64 ports	2 (10 GbE) / 4 ports	<b>CF Only</b>
<b>Crypto</b>					
<b>Crypto Express4S<sup>4</sup></b>	<b>NB</b>	16	16 coprocessors	1 coprocessor	1 feature <sup>2</sup>
Crypto Express3	CF <sup>1</sup>	8	16 coprocessors	2/1 coprocessors	<b>CF Only</b>
<b>Special Purpose</b>					
<b>10 GbE RoCE<sup>4</sup></b>	<b>NB</b>	16	16 usable SR ports	1 usable SR port	1 feature <sup>3</sup>
<b>zEDC Express<sup>4</sup></b>	<b>NB</b>	8	8 coprocessors	1 coprocessor	1 feature <sup>3</sup>
<b>Flash Express<sup>4</sup></b>	<b>NB</b>	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
<b>zEC12</b>	64 H20 – 32* H43 – 64*	32 H20 – 16* H43 – 32*	32 <sup>(4)</sup> H20 – 16* H43 – 32*	32 <sup>(4)</sup> H20 – 16* H43 – 32*	32	N/A	48 <sup>(4)</sup>	112 <sup>(1)</sup> H20 – 72* <sup>(2)</sup> H43 – 104* <sup>(1)</sup>	128
<b>zBC12</b>	H13 – 32* H06 – 16*	H13 – 16* H06 – 8*	H13 – 16* H06 – 8*	H13 – 16* H06 – 8*	32	N/A	48 <sup>(4)</sup>	H13 – 72* <sup>(2)</sup> H06 – 56* <sup>(3)</sup>	128
<b>z196</b>	48 M15 – 32*	32 M15 – 16* M32 – 32*	32 M15 – 16* M32 – 32*	32 M15 – 16* M32 – 32*	32	N/A	48	104 <sup>(1)</sup> M15 - 72* <sup>(2)</sup> M32 – 100* <sup>(1)</sup>	128
<b>z114</b>	M10 – 32* M05 – 16*	M10 – 16* M05 – 8*	M10 – 12* M05 – 8*	M10 – 16* M05 – 8*	32	N/A	48	M10 – 72* <sup>(2)</sup> M05 – 56* <sup>(3)</sup>	128
<b>z10 EC</b>	N/A	N/A	32 E12 – 16*	32 E12 – 16*	32	16 (32/RPQ)	48	64	64
<b>z10 BC</b>	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

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

- 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)

## zEC12 Improved N+1 Radiator-based Air Cooling

- Closed loop water cooling N+1 pump system replaces modular refrigeration units (MRUs) used for air cooling in z196 and z10 EC

- No connection to chilled water required
- Fits in same space as z196 MRUs
- Water added to the closed loop system during installation
- **New “Fill and Drain Tool” used by BOTH radiator cooled and water cooled zEC12)**

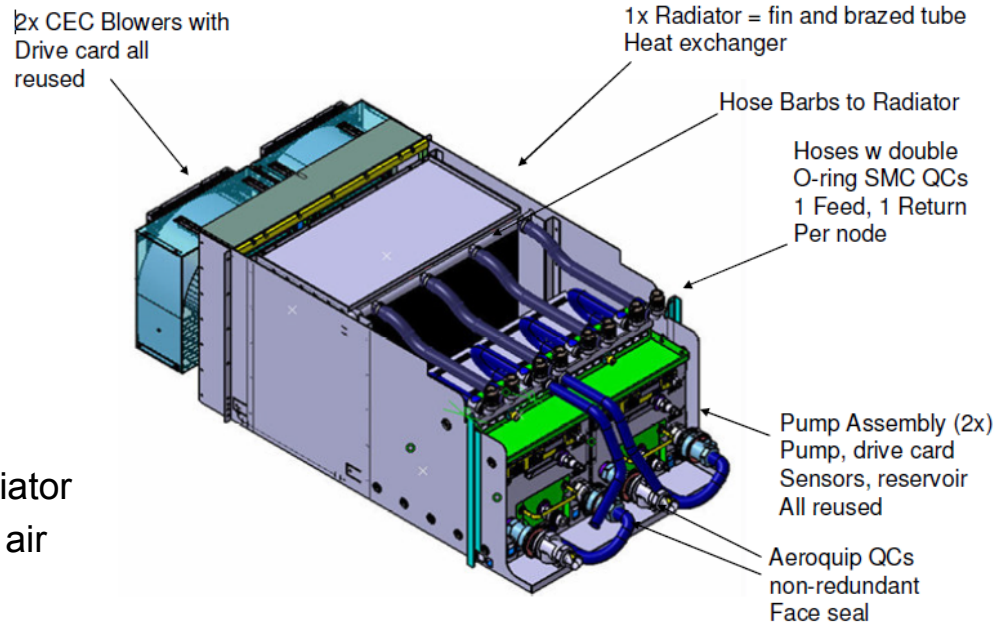
- Normal operation design:
  - Heat removed by water circulating to the radiator
  - Fans exhaust heat from the radiator to room air

- Radiator Cooled backup operation design

- N+1 pump/blower failure: Cooling maintained by closed loop water system without “cycle steering” slow down. Concurrent repair.

- **Improved Water Cooled backup operation design**

- Water cooling system failure: Cooling maintained by backup fans as in the z196 air cooled option with MRUs. “Cycle steering” slow down if needed to maintain operation. Concurrent repair



# zEC12 Fill and Drain Tool (FDT)



## System Fill Procedure

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

Approximate FDT unit dimensions:

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



# zBC12 and zEC12 GA2 Summary



# 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 <sup>(1)</sup>	Up to 3 <sup>(1)</sup>
I/O slots per I/O drawers/ PCIe I/O drawers	8	8/32		8/32 <sup>(2)</sup>	
FICON® Channels	128	128		128 <sup>(3)</sup>	
OSA Ports	96	96		96	
ESCON® Channels	480	240		0 <sup>(4)</sup>	
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 <sup>(8)</sup> /PSIFB	12/48/12	0/48/8 -16	0/48/16 - 32	0 <sup>(5)</sup> /32/8 -16 <sup>(6)</sup>	0 <sup>(5)</sup> /32/16 - 32 <sup>(7)</sup>
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

\* 2<sup>nd</sup> 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 <sup>(3)</sup>		256
OSA Ports (10GbE/1GbE/1000BASE-T)	48/96/96		48/96/96
ESCON® Channels	0 <sup>(4)</sup>		0 <sup>(4)</sup>
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 <sup>(8)</sup> /PSIFB	0 <sup>(5)</sup> /32/8 - 16 <sup>(6)</sup>	0 <sup>(5)</sup> /32/16 - 32 <sup>(7)</sup>	0 <sup>(5)</sup> /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

# zEC12 Processor and Memory Design



# Last Slide

