The IBM zEnterprise® EC12 (zEC12) System: Processor, Memory and System Structure Enhancements

Session ID: 13078
Tuesday, February 5, 2013, 9:30 AM
Grand Ballroom B, San Francisco Hilton

Speaker: Harv Emery

Advanced Technical Skills

Trademarks

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

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

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

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

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

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

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

IBM zEC12 Processor Introduction

IBM zEC12 Functions and Features

Processor, Memory, RAS

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

Parallel Sysplex, Security, I/O

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

Upgradeable from z10 EC and z196

Environmental

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

Ensemble, Platform Management

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

[1]Based on preliminary internal measurements and projections against a z196. Official performance data will be available upon announce and can be obtained online at LSPR (Large Systems Performance Reference) website at: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument . Actual performance results may vary by customer based on individual workload, configuration and software levels.

Introducing the newest members of the zEnterprise System family

The zEnterprise EC12 and zEnterprise BladeCenter Extension Model 003

IBM zEnterprise EC12 (zEC12)

- zEC12 has the industry's fastest superscalar chip with each core at 5.5 GHz
- New innovation to drive availability with IBM zAware and Flash Express
- Optimized for the corporate data serving environment
- Hardware functions boost software performance for Java™, PL/I, DB2®

IBM zEnterprise Unified Resource Manager and zEnterprise BladeCenter® Extension (zBX) Mod 003

- Supports the new zEC12 platform
- Hosts PS701 and HX5 blades
- Provides workload-awareness resource optimization
- Enhancements to System Director support zBX
- System z will continue to expand hybrid computing

Plus more flexibility and function by IBM DB2 Analytics Accelerator for z/OS

- IBM DB2 Analytics Accelerator for z/OS allows deployment of business analytics on the same platform as operational applications
- Analytics and OLTP can be run as the same workload

IBM zEC12 and zBX Model 3 Planned Availability Dates

- **September 19, 2012**
  - New build IBM zEC12 and zBX Model 3 (All announced features except as shown below)
  - Air-cooled z196 upgrade to air-cooled or water-cooled zEC12 (A new zBX Model 3 can be included)
  - Water-cooled z196 upgrade to water-cooled zEC12 (A new zBX Model 3 can be included)
  - z10 EC upgrades to air-cooled or water-cooled zEC12 (A new zBX Model 3 can be included)
  - z196 with zBX Model 2 upgrade to zEC12 with zBX Model 3 (The zBX Model 2 is upgraded to Model 3)
  - TKE 7.2 LIC (FC 0850) on z196 and z114
- **November 7, 2012**
  - Feature changes to an installed build zBX Model 3[1]
  - Add a zBX Model 3 to a previously installed zEC12[1]
  - zBX Model 002 detach, upgrade to zBX Model 003 and attach to an installed zEC12 (zEC12 FC #0031) [1]
  - zEC12 380-415V 3PH line cord sets: 14 foot, FC #8976, and top exit, FC #8977
- **December 7, 2012**
  - Upgrade an installed zEC12 to Model H43, H66, H89 or HA1
- **December 14, 2012**
  - Web download for z/OS V1.13 exploitation of Flash Express memory
- **December 31, 2012**
  - Other feature changes to an installed zEC12[1]
- **1Q2013**
  - zBX Model 3 detach (zEC12 FC #0030) and attach to another installed zEC12
  - DataPower Blade remove or move from one zBX to another zBX
  - z/OS V1.13 exploitation of 2 GB large pages and dynamic reconfiguration of Flash Express memory

Available Now

Note 1. Adding a zBX Model 3 to an installed zEC12 or changing features of an installed zBX Model 3 usually requires feature changes to the zEC12. Those changes are available on November 7, 2012 when made with a zBX add or change.

IBM zEC12 Processor, Memory and System Structure

IBM zEC12 Continues the CMOS Mainframe Heritage Begun in 1994

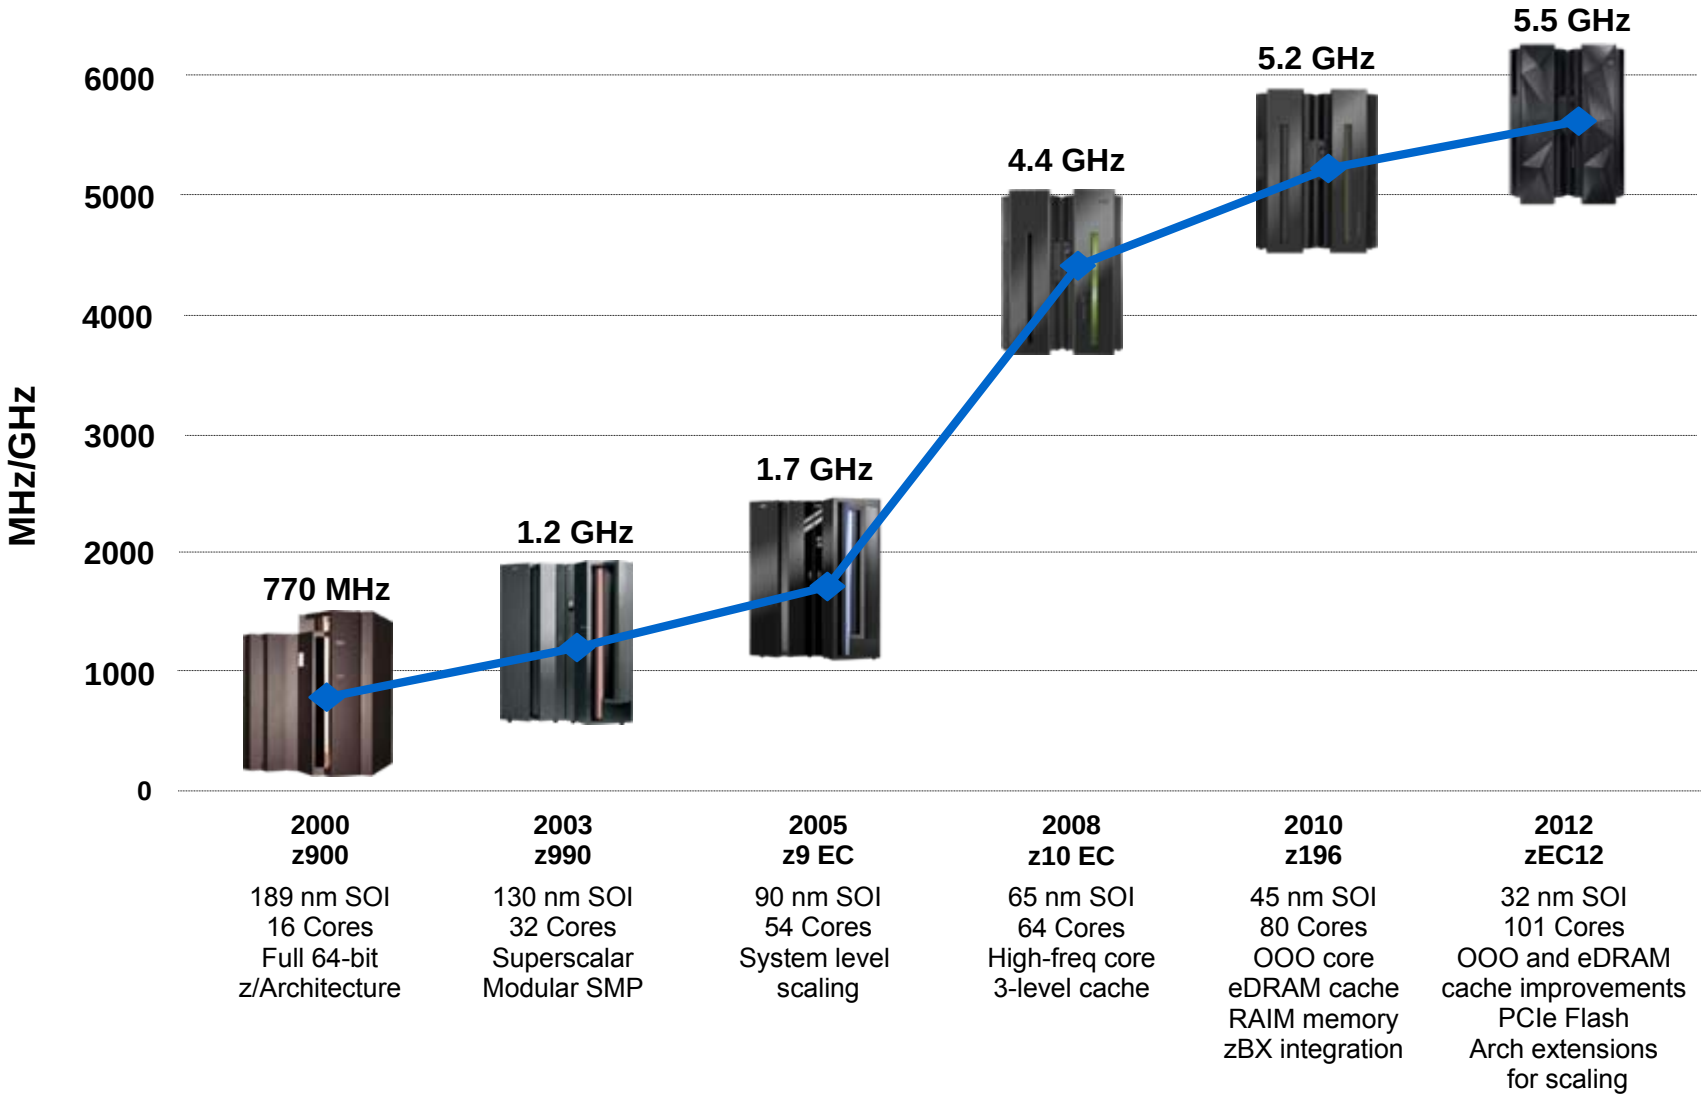

System Offering Overview

z196 2458-002

z10 EC

New Server

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

zEC12 Processor Design

zEC12 Hex Core Processor Chip Detail

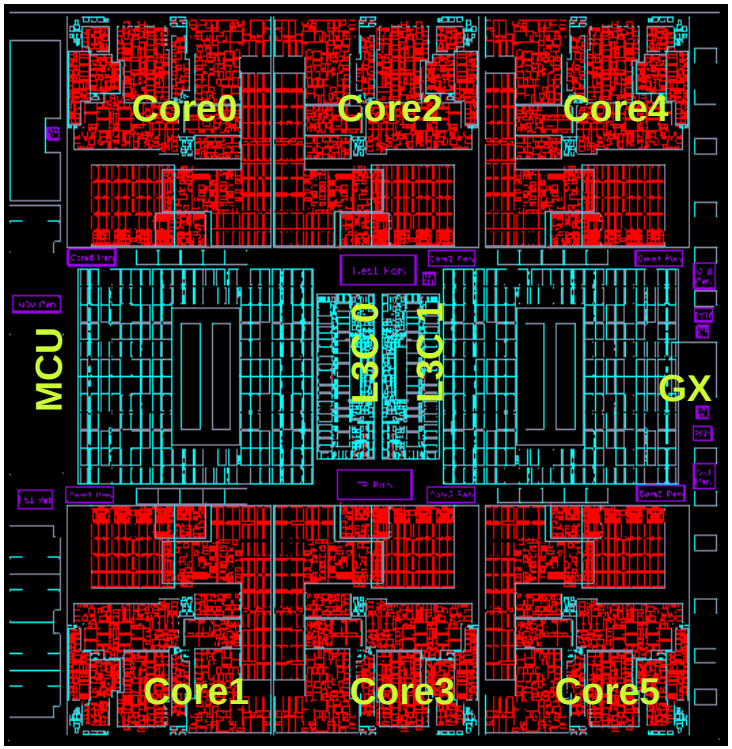

- **13S 32nm SOI Technology**
  - 15 layers of metal
  - 7.68 km wire
- **2.75 Billion Transistors**
- **Chip Area**
  - 597 mm²
  - 23.7mm x 25.2mm
  - 10000+ Power pins
  - 1071 signal I/Os

- **Up to six active cores (PUs) per chip**
  - 5.5 GHz
  - 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**
  - Shared by all six cores
- **Interface to SC chip / L4 cache**
  - 40 GB/sec to each of 2 SCs (5.5 GHz)
- **I/O Bus Controller (GX)**
  - Interface to Host Channel Adapter (HCA)
- **Memory Controller (MC)**
  - Interface to controller on memory DIMMs
  - Supports RAIM design

Out of Order Execution – z196 Vs zEC12

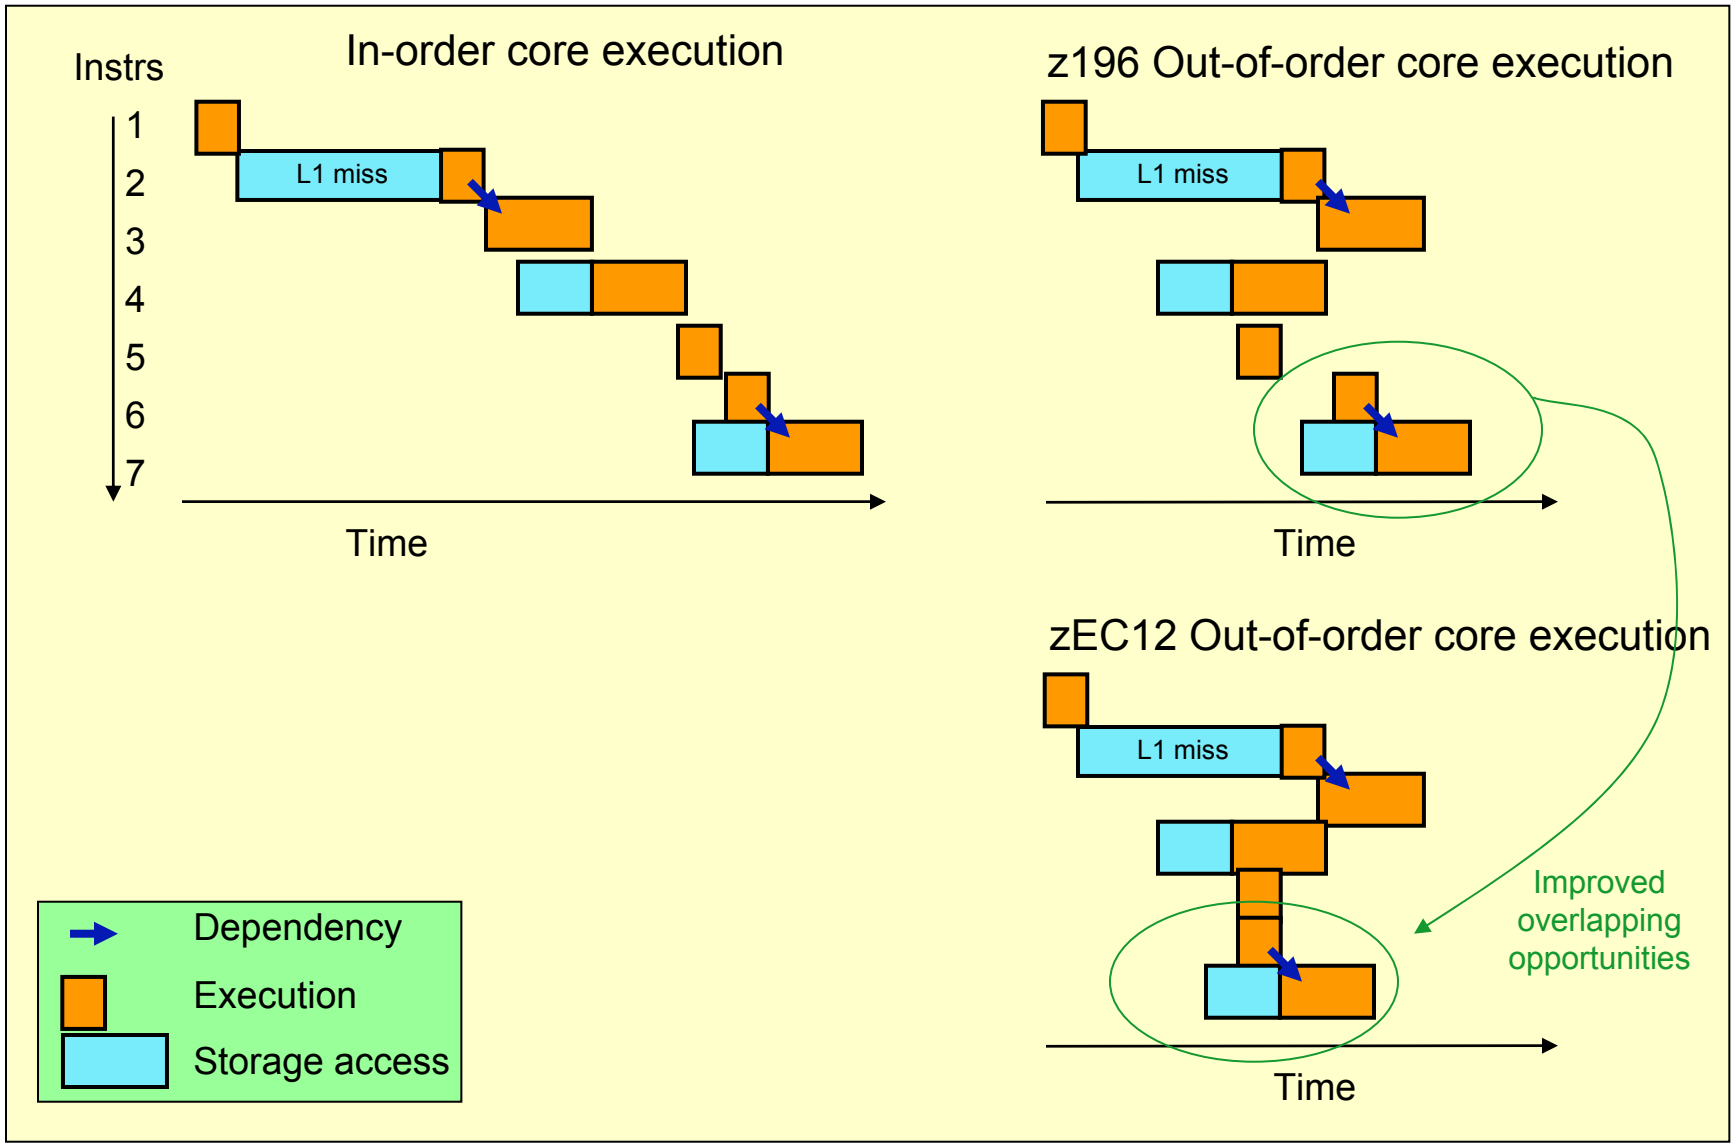

zEC12 OOO - Improved instruction delivery and execution

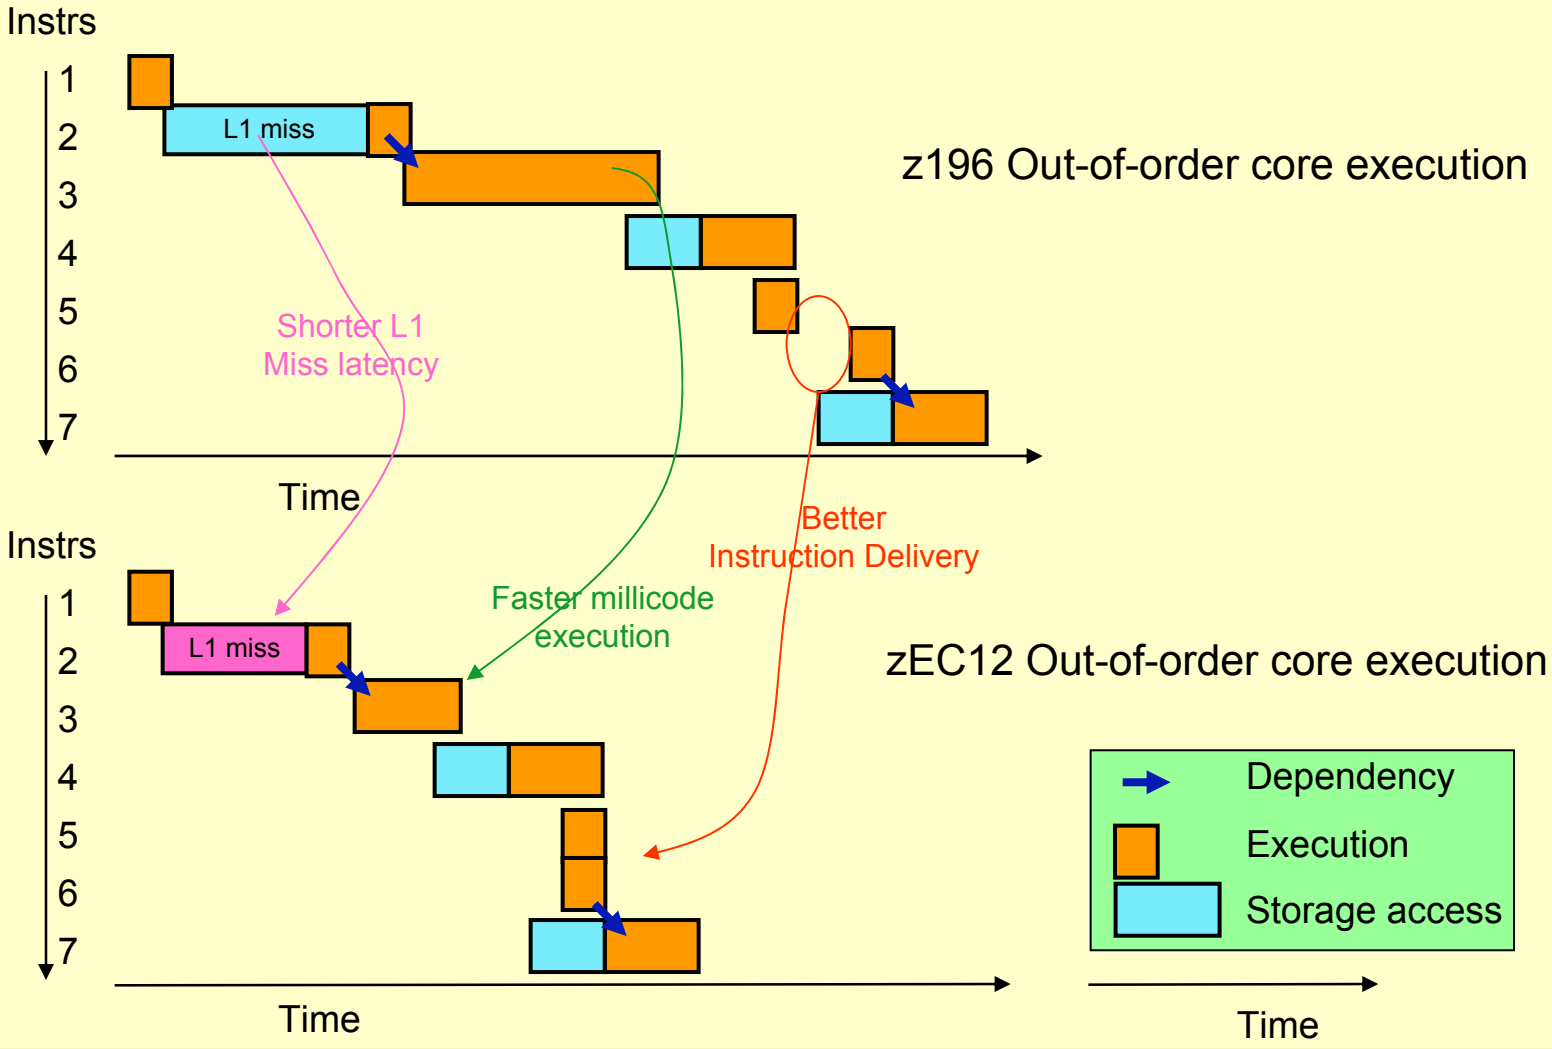

zEC12 Compression and Cryptographic Coprocessor

- **Coprocessor dedicated to each core (Was shared by two cores on z196)**
  - Independent compression engine
  - Independent cryptographic engine
  - Available to any processor type
  - Owning processor is busy when its coprocessor is busy
- **Data compression/expansion engine**
  - Static dictionary compression and expansion
- **CP Assist for Cryptographic Function**
  - 290-960 MB/sec bulk encryption rate
  - 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

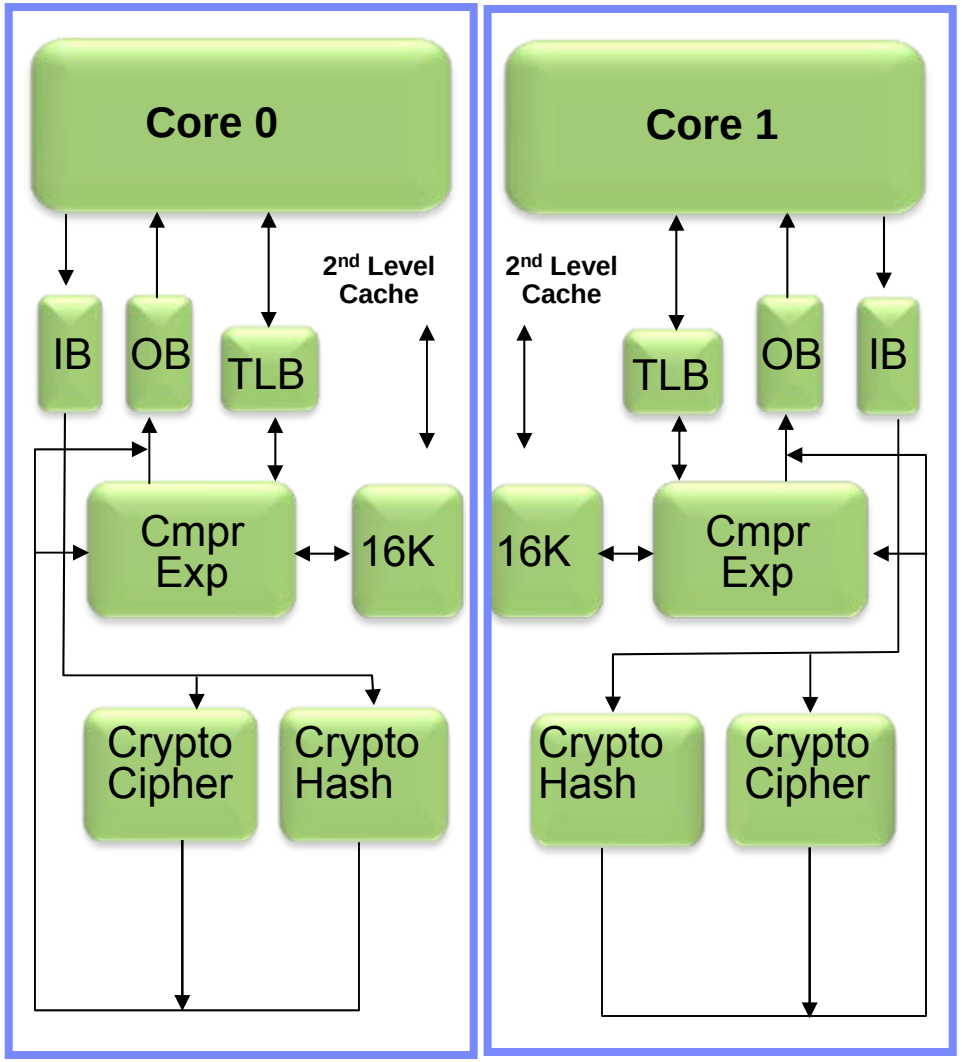

zEC12 Architecture Extensions

- **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
  - Designed for near term exploitation by Java longer-term opportunity for DB2, z/OS, 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 new Enterprise PL/I compiler**

zEC12 Storage Control (SC) Chip Detail

- **CMOS 13S 32nm SOI Technology**
  - 15 layers of metal
- **Chip Area –**
  - 526 mm²
  - 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
  - 384 MB per Book
- **6 CP chip interfaces**
- **3 Book fabric interfaces**
- **1 Clock domain**
- **4 Unique chip voltage supplies**

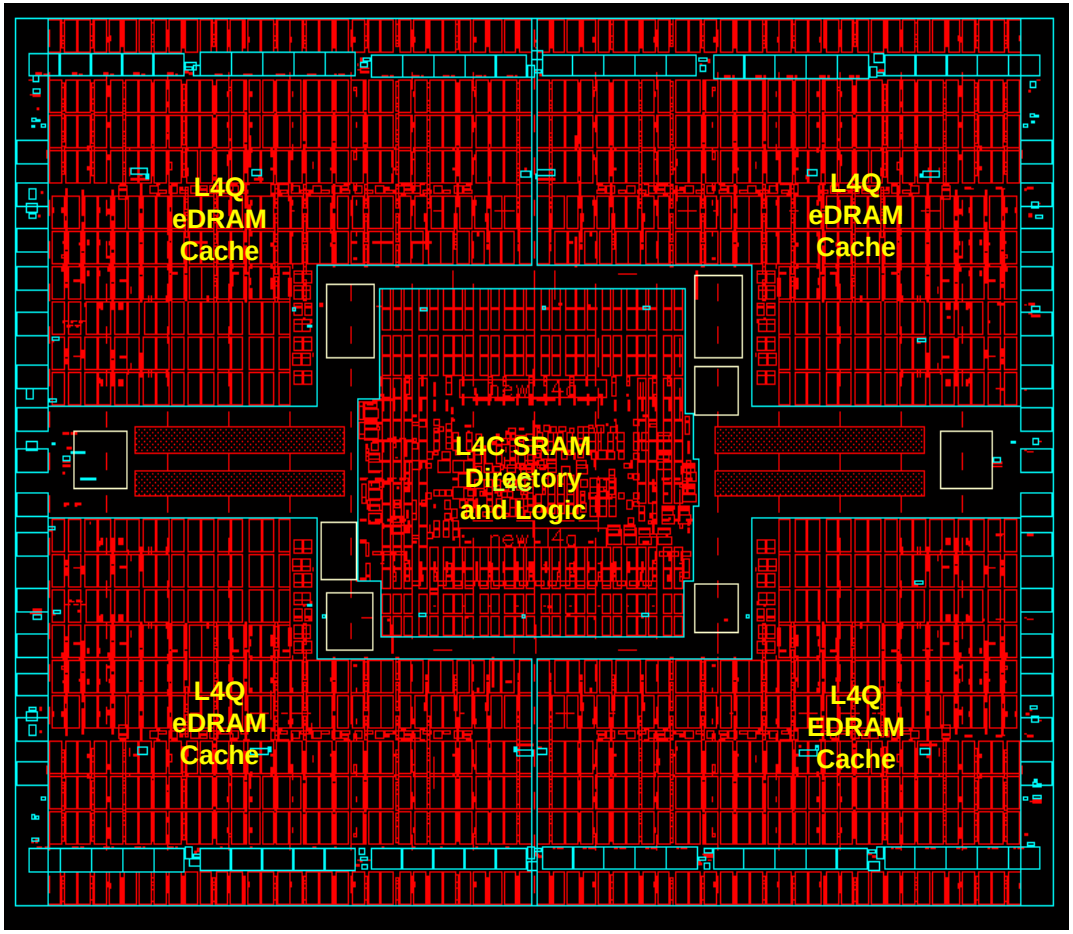

zEC12 PU chip, SC chip and MCM

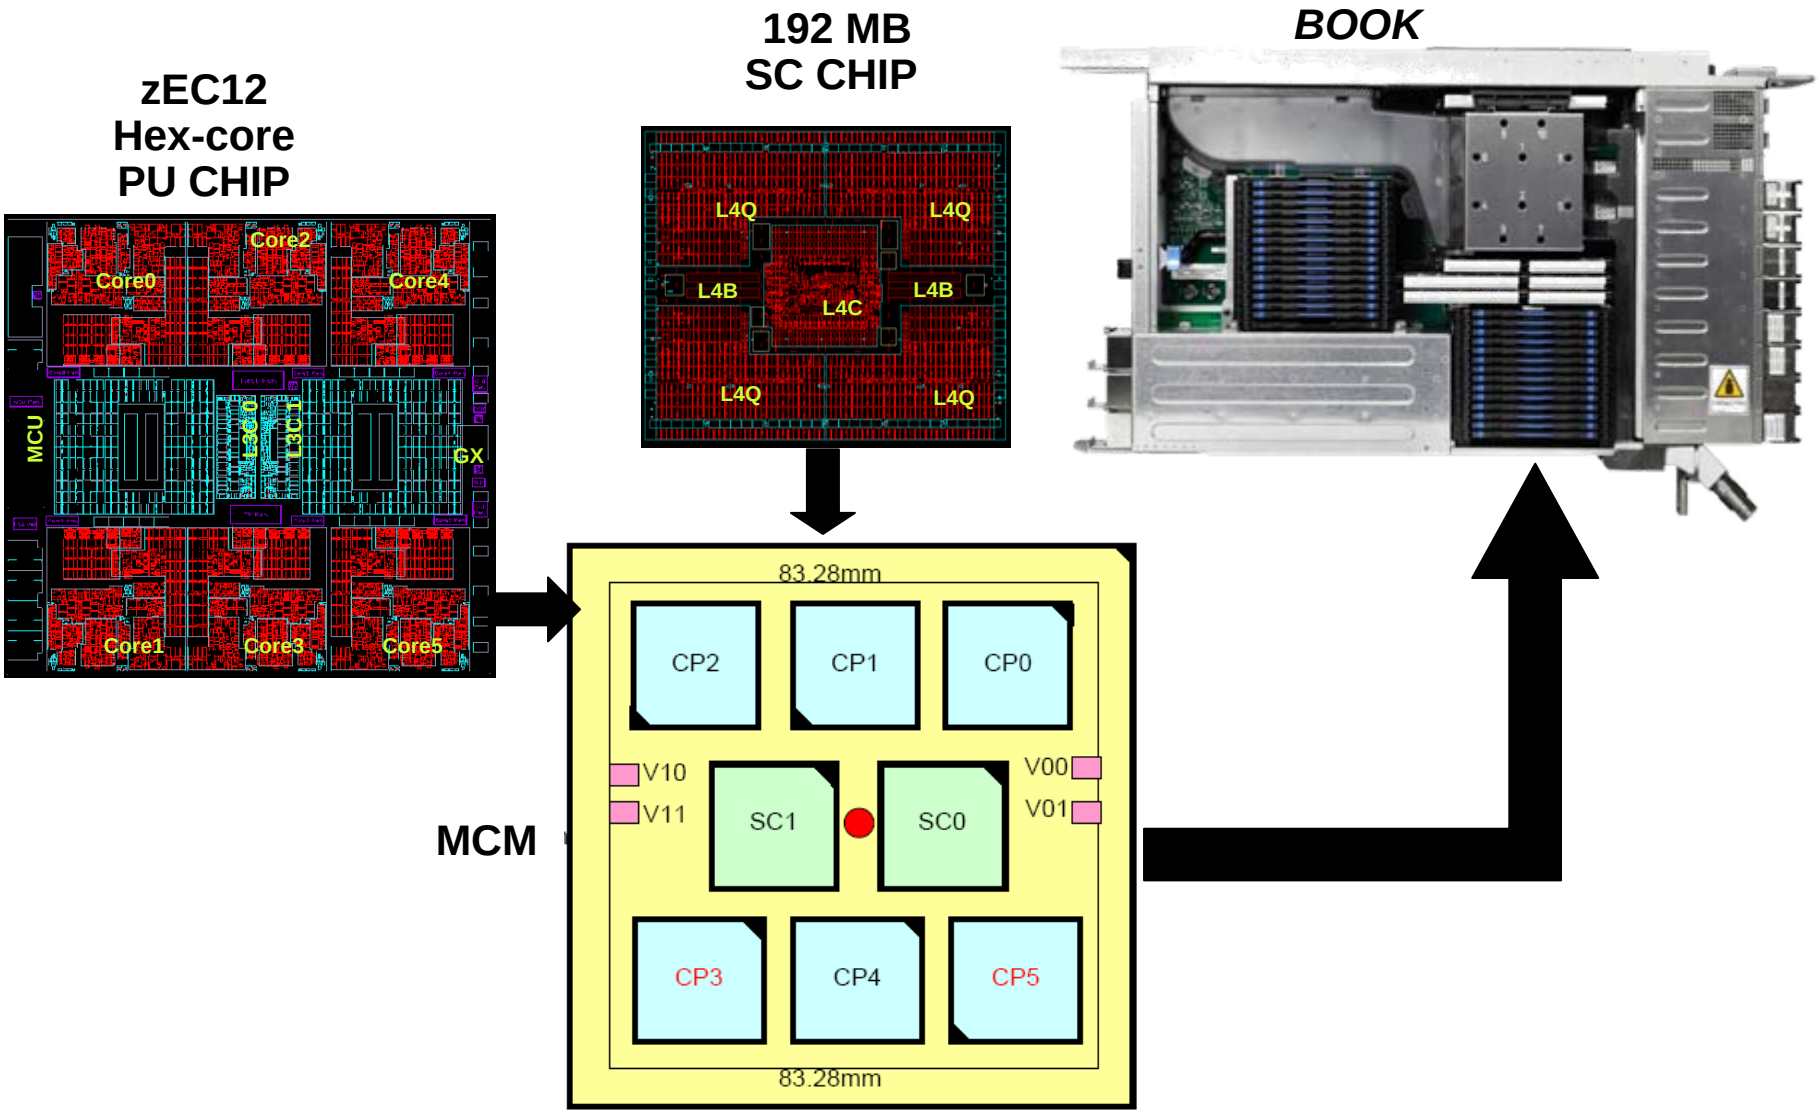

zEC12 Multi-Chip Module (MCM) Packaging

- 96mm x 96mm MCM
  - 102 Glass Ceramic layers
  - 8 chip sites
- 7356 LGA connections
  - 27 and 30 way MCMs
  - Maximum power used by MCM is 1800W

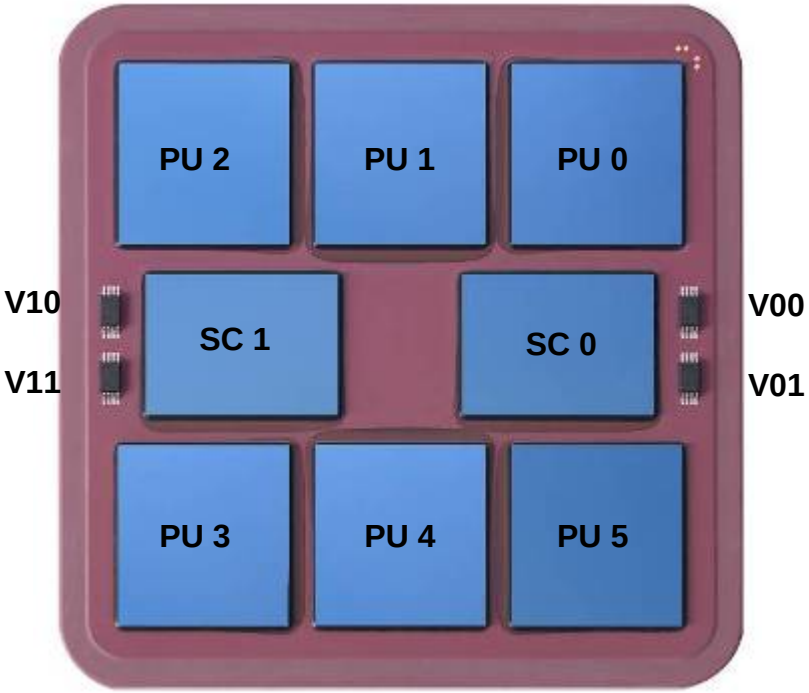

- CMOS 13s chip Technology
  - PU, SC, S chips, 32nm
  - 6 PU chips/MCM – Each up to 6 active cores
    - 23.7 mm x 25.2 mm
    - 2.75 billion transistors/PU chip
    - L1 cache/PU core
      - 64 KB I-cache
      - 96 KB D-cache
    - L2 cache/PU core
      - 1 MB I-cache
      - 1 MB D-cache
    - L3 cache shared by 6 PUs per chip
      - 48 MB I+D cache
    - 5.5 GHz
  - 2 Storage Control (SC) chip
    - 26.72 mm x 19.67 mm
    - 3.3 billion transistors/SC chip
    - L4 Cache 192 MB per SC chip (384 MB/book)
    - L4 access to/from MCMs in other books
  - 4 SEEPROM (S) chips – 1024k each
    - 2 x active and 2 x redundant
    - Product data for MCM, chips and other engineering information
  - Clock Functions – distributed across PU and SC chips
    - Master Time-of-Day (TOD) function is on the SC

z196 EC MCM vs zEC12 MCM Comparison

z196 MCM

- **MCM**
  - **–96mm x 96mm in size**
  - **–6 PU chips per MCM**
    - **Quad core chips with 3 or 4 active cores**
    - **PU Chip size 23.7 mm x 21.5 mm**
    - **5.2 GHz**
    - **Superscalar, OoO execution**
    - **L1: 64 KB I / 128 KB D private/core**
    - **L2: 1.5 MB I+D private/core**
    - **L3: 24 MB/chip – shared**
  - **–2 SC chips per MCM**
    - **L4: 2 x 96 MB = 192 MB L4 per book**
    - **SC Chip size 24.5 mm x 20.5 mm**
  - **–1800 Watts**

zEC12 MCM

- **MCM**
  - **–96mm x 96mm in size**
  - **–6 PU chips per MCM**
    - **Hex-core chips with 4 to 6 active cores**
    - **PU Chip size 23.7 mm x 25.2 mm**
    - **5.5 GHz**
    - **Improved superscalar and OoO execution**
    - **L1: 64 KB I / 96 KB D private/core**
    - **L2: 1 MB I / 1 MB D private/core**
    - **L3: 48 MB/chip - shared**
  - **–2 SC chips per MCM**
    - **L4: 2 x 192 MB = 384 MB L4 per book**
    - **SC Chip size 26.72 mm x 19.67 mm**
  - **–1800 Watts**

zEC12 Book Level Cache Hierarchy

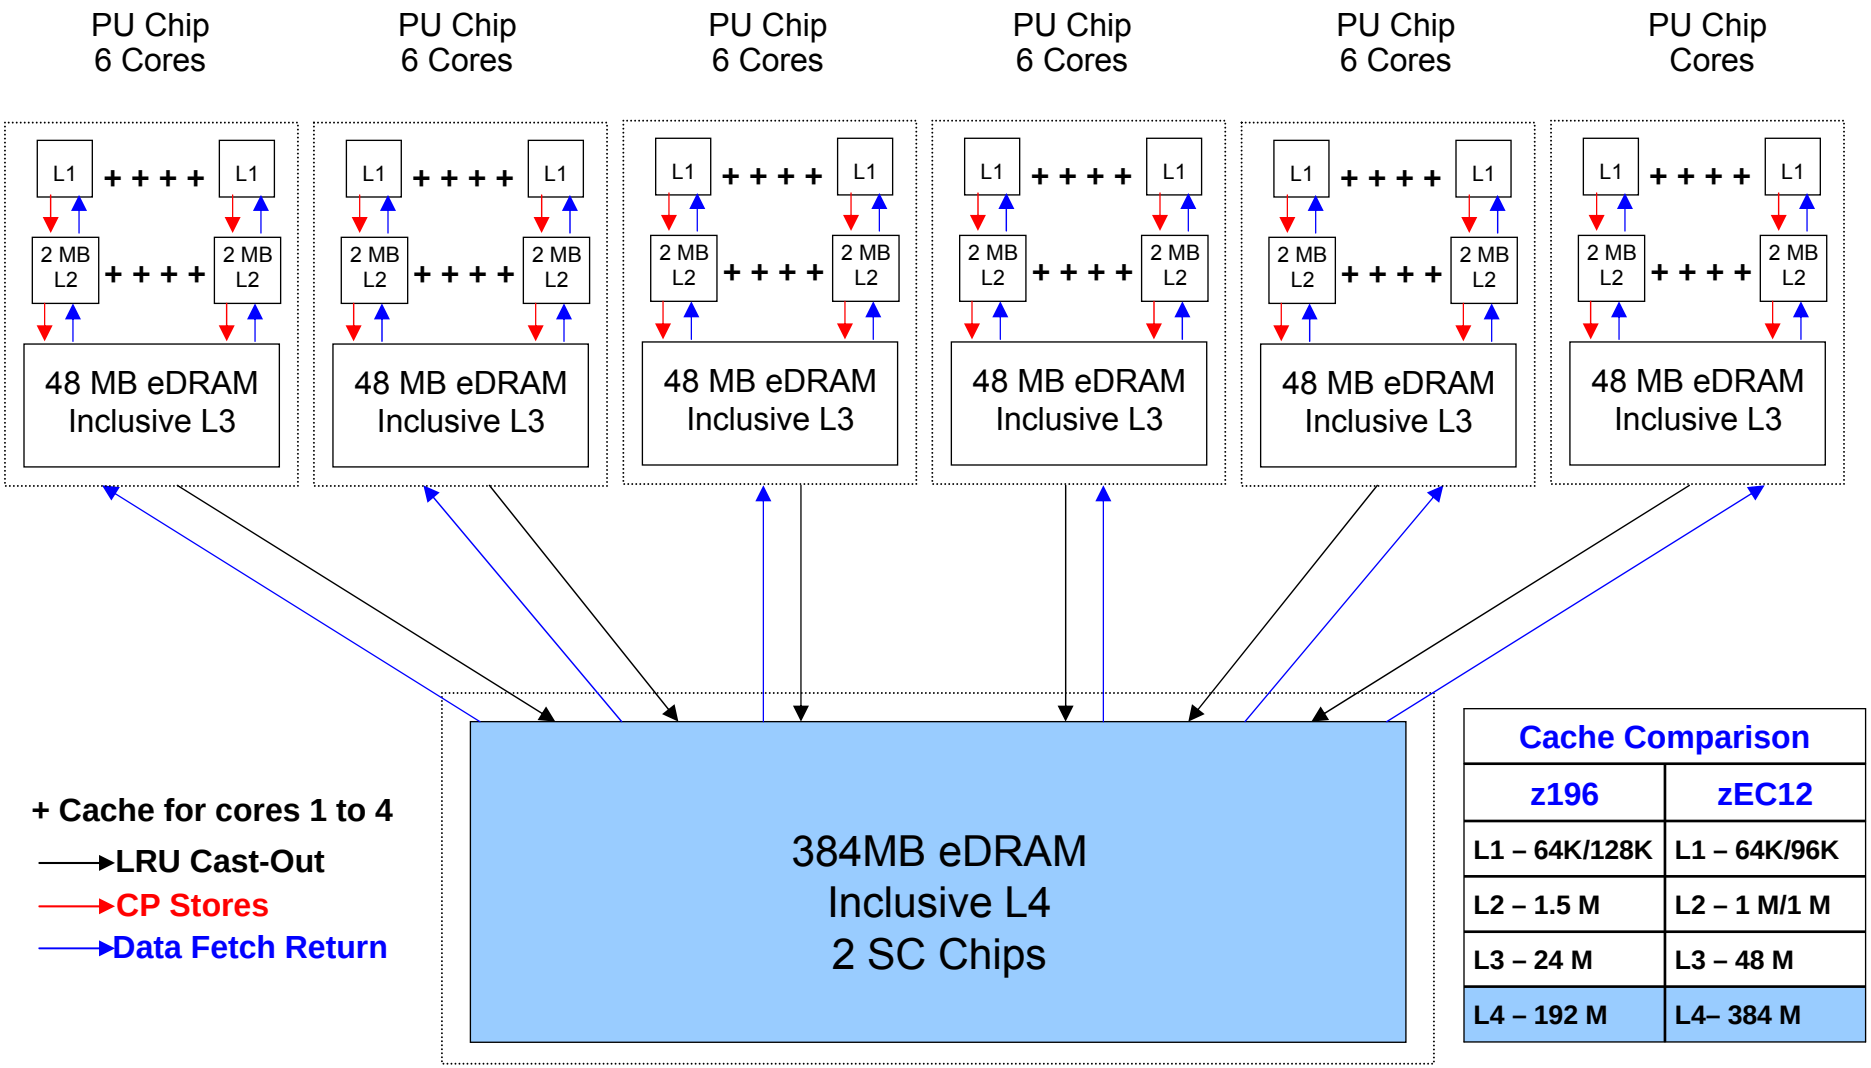

Cache Comparison	
z196	zEC12
L1 – 64K/128K	L1 – 64K/96K
L2 – 1.5 M	L2 – 1 M/1 M
L3 – 24 M	L3 – 48 M
L4 – 192 M	L4– 384 M

zEC12 Book Layout

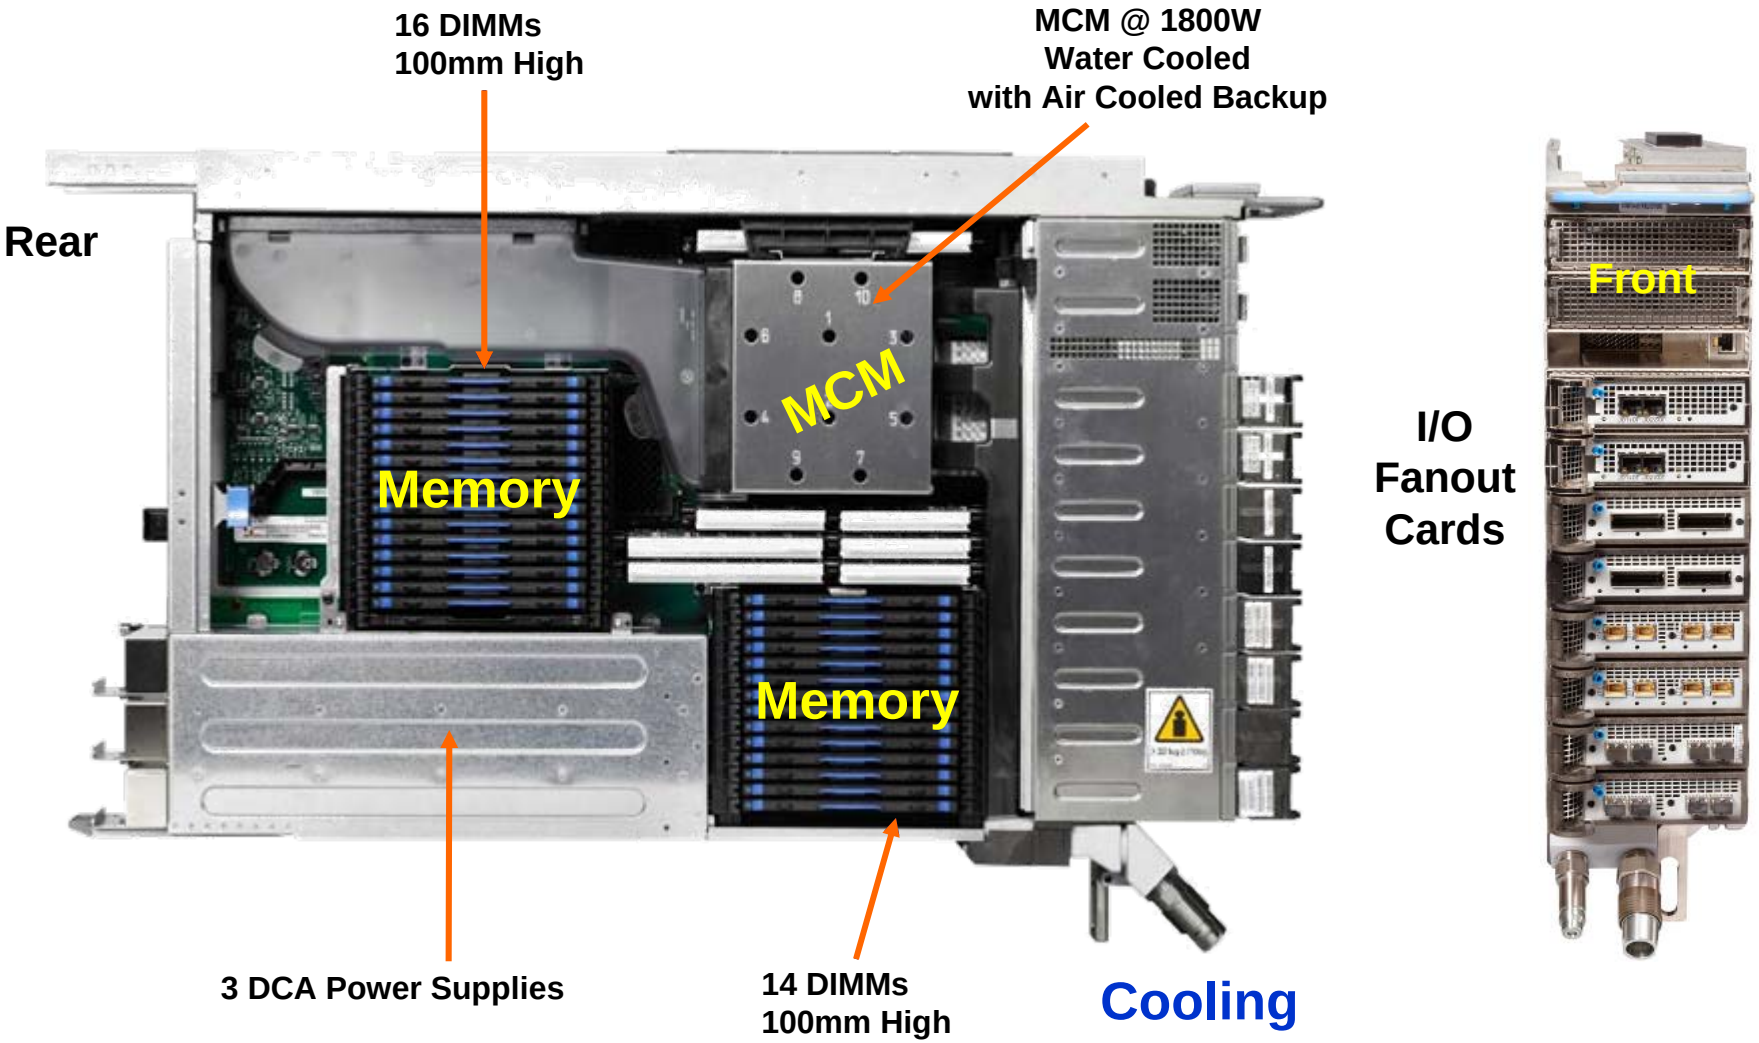

Note: Unlike the z196, zEC12 Books are the same for the Air and Water cooled Systems

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

(Front View)

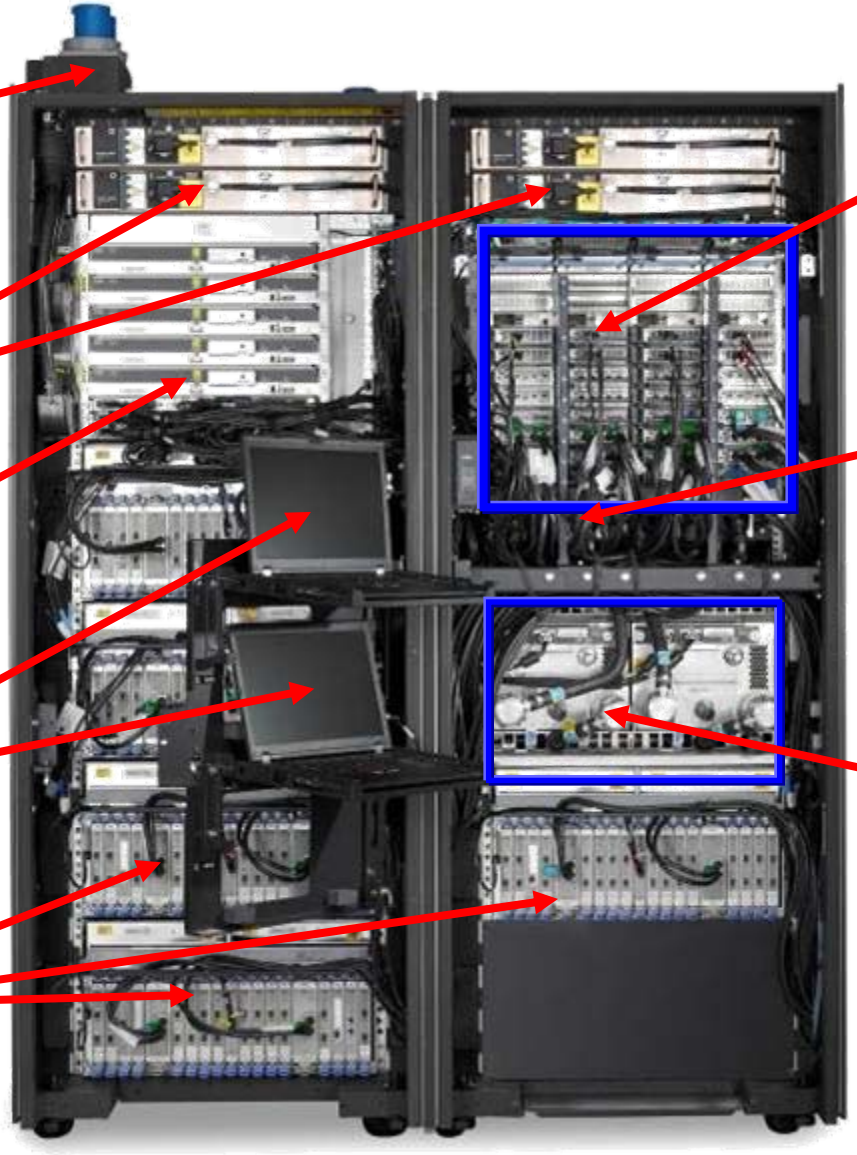

Overhead Power Cables (option)

Internal Batteries (option)

Power Supplies

2 x Support Elements

PCIe I/O drawers (Maximum 5 for zEC12)

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

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

N+1 Radiator-based Air Cooling Unit

Optional FICON LX Fiber Quick Connect (FQC) not shown

zEC12 Processor Features

Model	Books/ PUs	CPs	IFLs uIFLs	zAAPs	zIIPs	ICFs	Std SAPs	Optional SAPs	Std. Spares	Rsvd. PUs
H20	**1/27**	**0-20**	**0-20 0-19**	**0-10**	**0-10**	**0-20**	**4**	**0-4**	**2**	**1**
H43	**2/54**	**0-43**	**0-43 0-42**	**0-21**	**0-21**	**0-43**	**8**	**0-8**	**2**	**1**
H66	**3/81**	**0-66**	**0-66 0-65**	**0-33**	**0-33**	**0-66**	**12**	**0-12**	**2**	**1**
H89	**4/108**	**0-89**	**0-89 0-88**	**0-44**	**0-44**	**0-89**	**16**	**0-16**	**2**	**1**
HA1	**4/120**	**0-101**	**0-101 0-100**	**0-50**	**0-50**	**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 Reserved PU** is not available for customer purchase
- Concurrent Book Add is available to upgrade from model H20 to model H89

Notes: 1. At least one CP, IFL, or ICF must be purchased in every machine
2. One zAAP **and** one zIIP may be purchased for each CP purchased even if CP capacity is “banked”.
3. “uIFL” stands for Unassigned IFL

zEC12 Processor Book Assignment and Purchased Processor Placement

zEC12 Processor Book Assignment														
Model	Total PUs (cores)	1st Book - LG06				2nd Book - LG15			3rd Book - LG10			4th Book - LG01		
		Avail PUs	Std SAPs	Spare PUs	Rsvd PU	Avail PUs	Std SAPs	Spare PUs	Avail PUs	Std SAPs	Spare PUs	Avail PUs	Std SAPs	Spare PUs
H20	27	20	4	2	1	-	-	-	-	-	-	-	-	-
H43	54	21	4	1	1	22	4	1	-	-	-	-	-	-
H66	81	22	4	0	1	22	4	1	22	4	1	-	-	-
H89	108	22	4	0	1	22	4	1	22	4	1	23	4	0
HA1	120	25	4	0	1	25	4	1	25	4	1	26	4	0

- **Books in zEC12 Models H20 to H89 have 27 PU MCMs. Books in the Model HA1 have 30 PU MCMs.**
- **MCMs are built with a combination of processor chips with 4, 5 or 6 active PUs (cores)**
- **Concurrent Book Add provides concurrent model upgrade except to zEC12 Model HA1**
  - Upgrade to zEC12 Model HA1 is done by disruptive replacement of all books
- **Purchased Processors do NOT have fixed chip or book placement. (That was last done on z990)**
  - Purchased processor placement in books and on PU chip cores happens at Power on Reset (POR) and during concurrent add or remove of purchased processors by MES, OnDemand upgrade or downgrade, Concurrent Book Add, or Enhanced Book Availability book remove or add.
  - ICFs and IFLs are assigned to books and PU chips beginning with the highest working downward – Fill and Spill
  - CPs, zAAPs and zIIPs are assigned to books and PU chips beginning with the lowest working upward – Fill and Spill
  - PR/SM Dynamic Processor Reassignment in cooperation with z/O HiperDispatch can reassign CPs, zAAPs and zIIPs

IBM zIIP and zAAP Simplification Statement of Direction

- **IBM System z Integrated Information Processor (zIIP) and IBM System z Application Assist Processor (zAAP) Simplification[1]**
  - IBM zEC12 is planned to be the last high-end System z server 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 plans to provide a PTF for APAR OA38829 on z/OS V1.12 and V1.13 in September 2012 to remove the restriction that prevents zAAP-eligible workloads from running on zIIP processors when a zAAP is installed on the server. This is intended only to help facilitate migration and testing of zAAP workloads on zIIP processors. This works on any System z server that supports zIIPs and zAAPs, not just on zEC12.

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.

LSPR PCI Comparison Between z196 and zEC12 1-Way[1]

z196 701 1,202 (Measured with z/OS V1.11 and V1.13*)
zEC12 701 1,514 (Measured with z/OS V1.13*)
IFL PVU rating remains at 120 (No change from z10 and z196)
* LSPR ratios and PCI measurements have been revised based on z/OS V1.13)

- Balanced performance growth
  - Across broad range of workloads
  - Scalable from 1 to 101 processors
- LSPR: 1.25 x z196 constant n-way
  - 32nm SOI technology
  - Improved core and cache designs
- 1.5X z196 total capacity
  - 27 or 30 cores per MCM, up from 24
  - 120 max cores, up from 96
- Continued full-stack performance focus
  - New z/Architecture features
  - Compiler optimization for zEC12

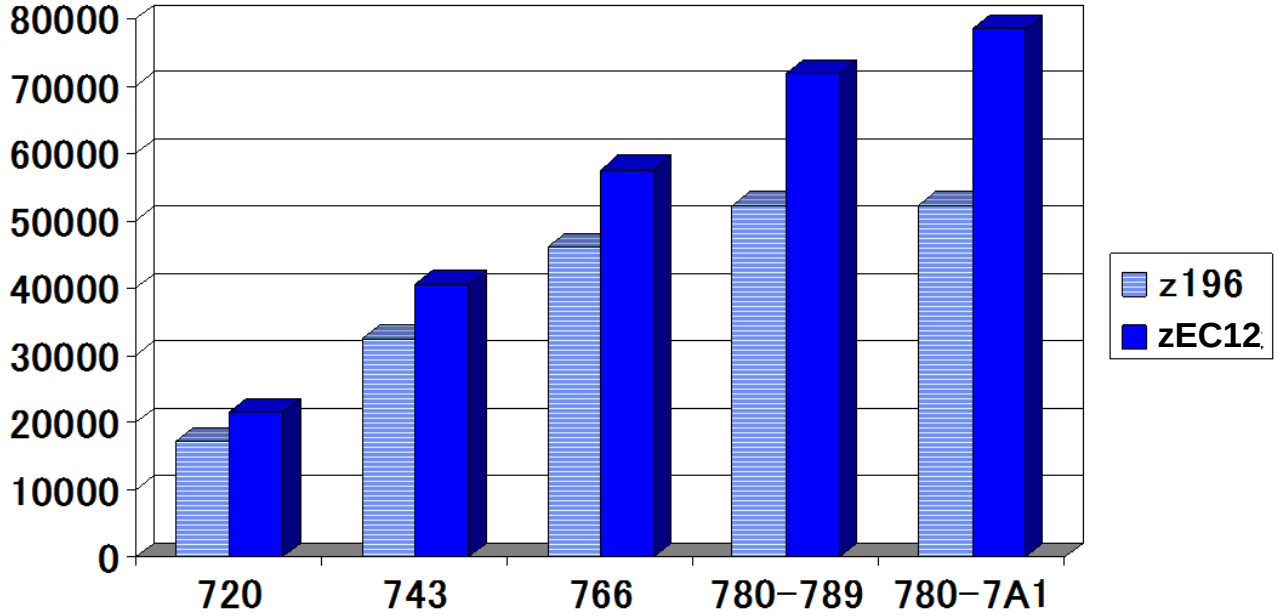

[1]Based on preliminary internal measurements and projections against a z196. Official performance data will be available upon announce and can be obtained online at LSPR (Large Systems Performance Reference) website at: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument . Actual performance results may vary by customer based on individual workload, configuration and software levels.

zEC12 Full and Sub-Capacity CP Offerings

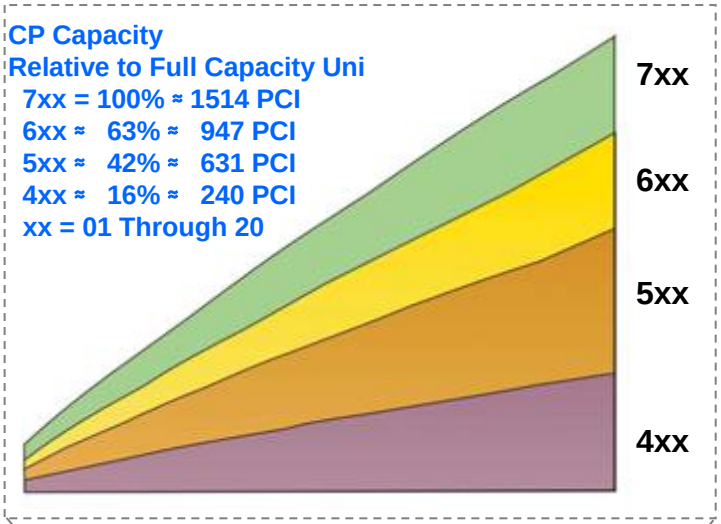

- **Subcapacity CPs, up to 20, may be ordered on ANY zEC12 model. If 21 or more CPs are ordered all must be full 7xx capacity**
- **All CPs on a zEC12 CEC must be the same capacity**
- **All specialty engines run at full capacity. The one for one entitlement to purchase one zAAP and one zIIP for each 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**

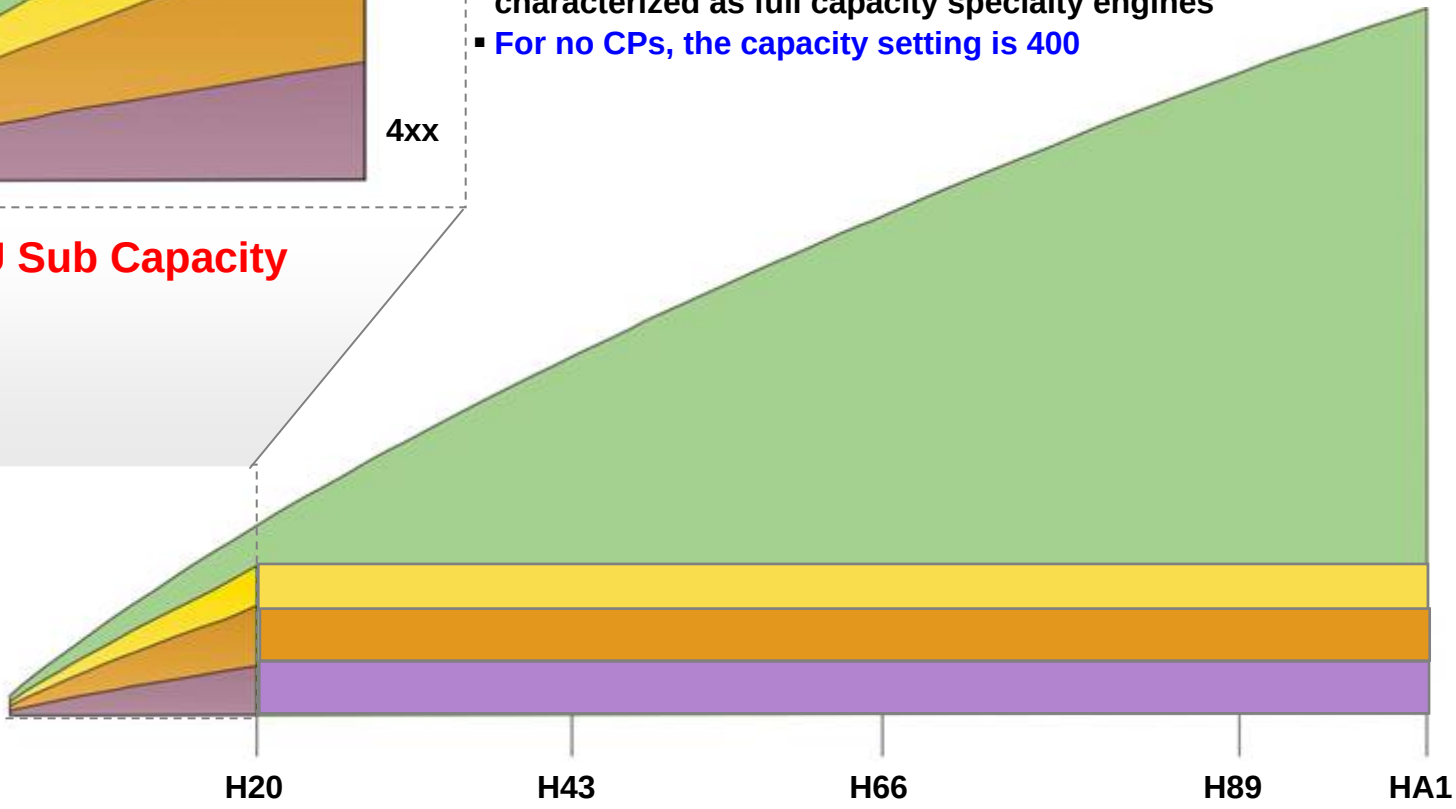

zEC12 Memory Design

zEC12 RAIM Memory Controller Overview

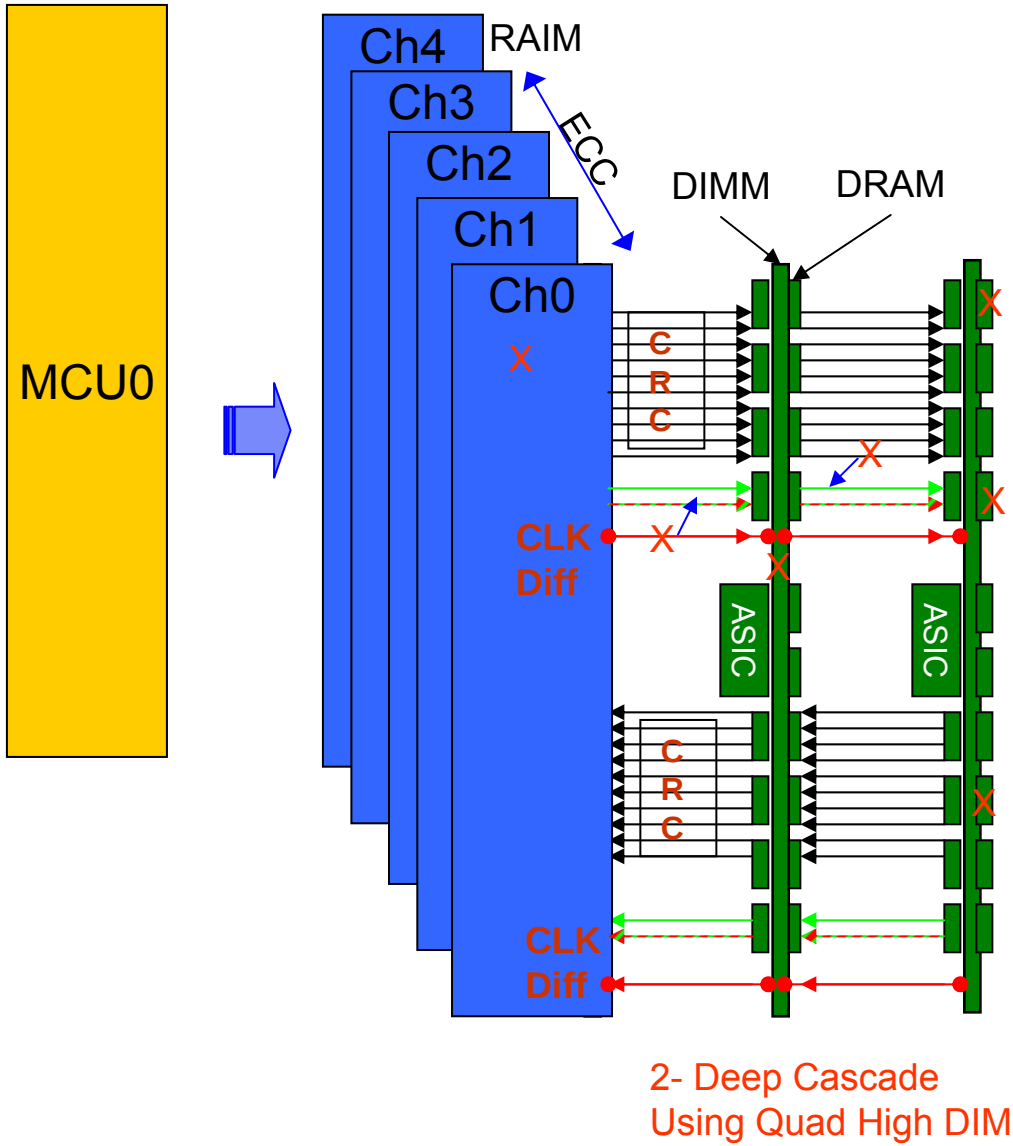

2- Deep Cascade
Using Quad High DIMMs

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 Purchase 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**

- ***Purchase 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.**
- ***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 1	13%			
	96	33%		800	12%	256 GB	2016	13%
	128	25%		896	12%		2272 3	11%
	160	20%					2528	10%
	192	17%	112 GB	1008	13%		2784	9%
	224	14%					3040 4	NA
	256	25%						
64 GB	320	20%	128 GB	1136	11%			
	384	17%		1264	10%			
	448	14%		1392 2	9%			
	512	19%		1520	16%			

Notes – Memory Maximums:

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)

Flash Express

Introducing Flash Express

- zEC12 will offer optional Flash Express memory cards
  - Supported in PCIe I/O drawer with other PCIe I/O cards
  - Pairs of cards for availability
  - No HCD/IOCP definition required
- Assign flash memory to partitions like main memory
  - Assignment is by minimum/maximum memory amount, not by feature
  - Each partition's flash memory is isolated like main memory
  - Dynamically increase the partition maximum amount of flash
  - Dynamically configure flash memory into and out of the partition
- Options to solve many different problems
  - Flash Memory is much faster than spinning disk
  - Flash Memory is much slower than main memory
  - Flash Memory takes less power than either
- The system z Software Stack has a staged plan to exploit flash memory
  - z/OS 1.13 (4 GB LPAR or larger) plus PTFs, JAVA - SDK 7, etc.
  - Future Exploiters: WAS Liberty Profile, DB2, IMS 12, etc
  - Future: Linux on System z (IBM is working with distributors)

Time to Read Data measured in System z Instructions
Real Memory: (256B line) 100 Instructions
Flash Memory (4K page) 100K Instructions
External Disk (4K page) 5,000K Instructions

What Is Flash Express?

- Flash Express is internal storage implemented via NAND Flash SSDs (Solid State Drives) mounted in PCIe Flash Express feature cards
  - Which plug into PCIe I/O drawers in pairs
  - Data security provided on the cards
  - A pair provides 1.4 TB of storage class memory (1 TB = 2^{30})
  - A maximum of 4 pairs are supported in a system
- Internal Flash is accessed using the new System z architected EADM (Extended Asynchronous Data Mover) Facility
  - An extension of the ADM architecture used in the past with expanded storage
  - Access is initiated with a Start Subchannel instruction
  - Subchannels used were previously reserved
  - Definition in IOCDS is not required
- The main application of internal Flash in zEC12 GA-1 is paging store for z/OS
  - Where it provides advantages in resiliency and speed
  - With pageable large pages being introduced in tandem for exceptional performance

More in Session 13086: Flash Express Introduction, Uses and Benefits
Today: 1:30 PM, Imperial Ballroom A

Security of Data on Flash Express

- System z internal flash can be used for paging, dumping and .....
  - It can contain all data, including audited personally identifiable data
- Client data on flash is protected by 128 bit AES Encryption
  - Done using hardware encryption at the device like IBM's Disk and Tape encryption
- Key Management is provided based on a Smart Card in each Support Element
- End of life Audit is based on access to the Smart Card, not access to the Flash Memory
  - Secure Cryptographic Erase well understood

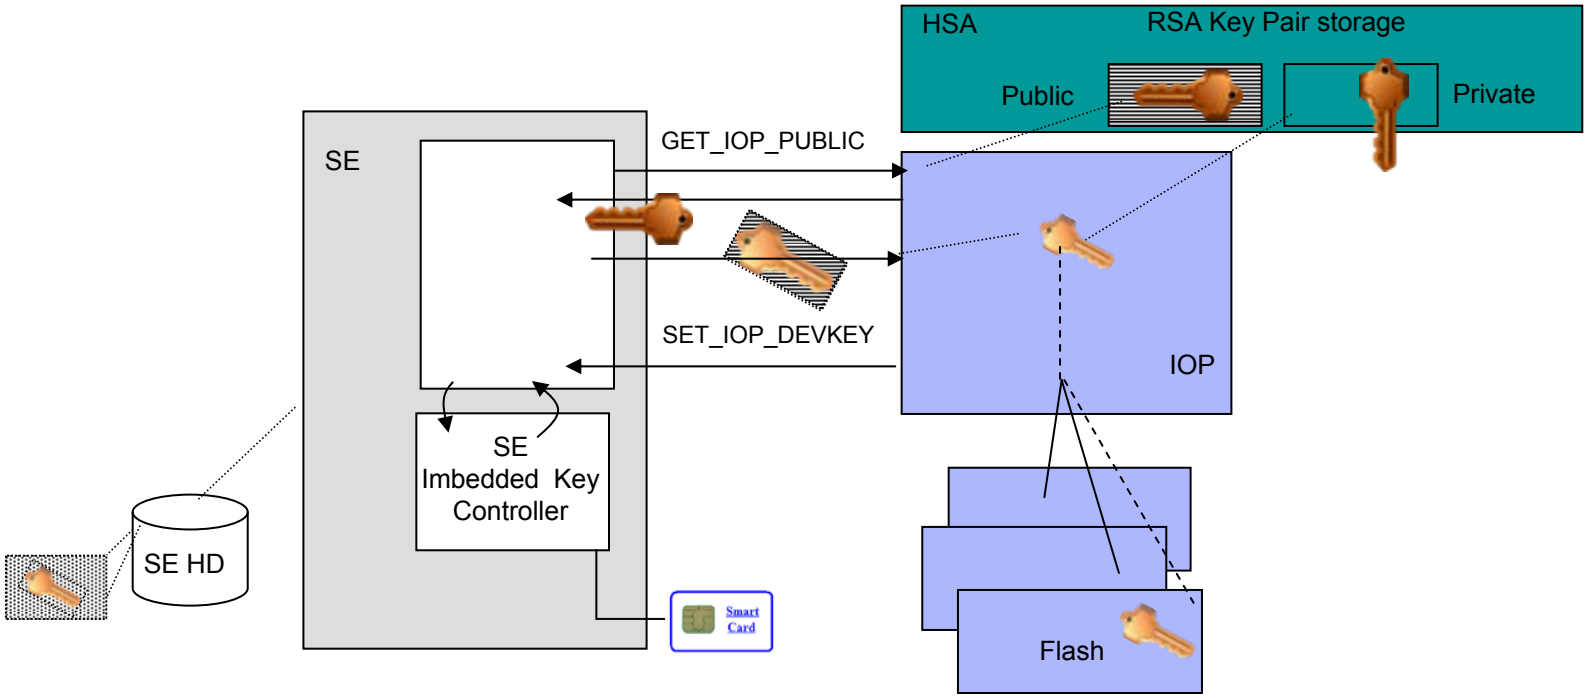

Flash Express PCIe Adapter Card

NEW

Four Solid State Drives (SSDs)

On Demand

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

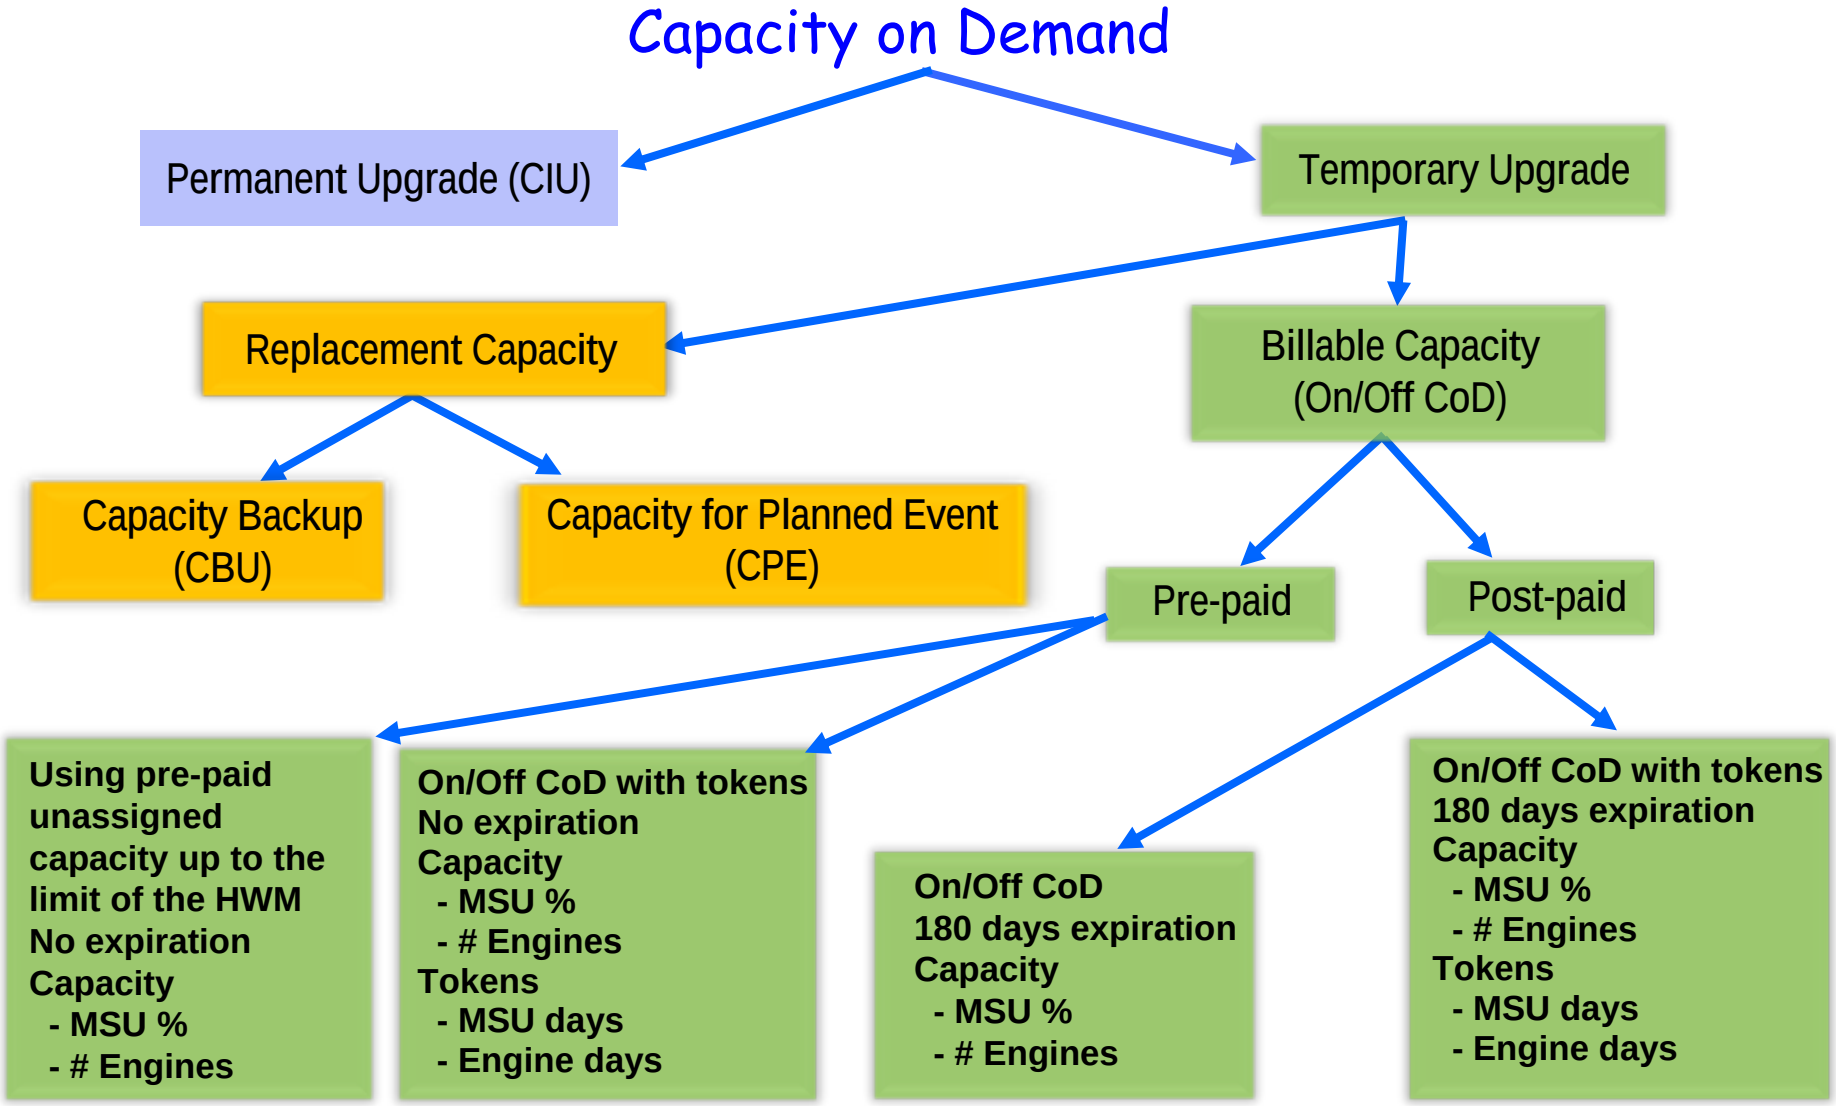

Crypto

Crypto Express4S

NEW

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

Three configuration options for the PCIe adapter

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

- Accelerator
  - For SSL acceleration
  - Clear key RSA operations

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

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

Crypto Enhancements

- **IBM Enterprise Public Key Cryptography Standards #11 (EP11)**
  - Based on PKCS #11 specification v2.20 and more recent amendments
  - Designed to meet Common Criteria EAL 4+ and FIPS 140-2 Level 4
  - Conforms to Qualified Digital Signature (QDS) Technical Standards

- **IBM Common Cryptographic Architecture (CCA)**
  - Secure Cipher Text Translate
  - DUKPT for derivation of MAC and Encryption Keys
  - Wrap weaker keys with stronger keys for security and standards compliance
  - Compliance with new Random Number Generator standards
  - EMV enhancements for applications supporting American Express cards

- **IBM Trusted Key Entry (TKE) 7.2 Licensed Internal Code (LIC)**
  - Support for Crypto Express4S defined as a CCA coprocessor
  - Support for Crypto Express4S as a Enterprise PKCS #11 coprocessor
  - Support for new DES operational keys
  - New AES CIPHER key attribute
  - Allow creation of corresponding keys
  - New smart card part 74Y0551
  - Support for 4 smart card readers
  - Support for stronger key wrapping standards
  - Compatible with current TKE Workstation hardware

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

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

Note 1: All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

I/O Subsystem Introduction

More detail in Session 13077
11:00 Today, Grand Ballroom B

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

Front view

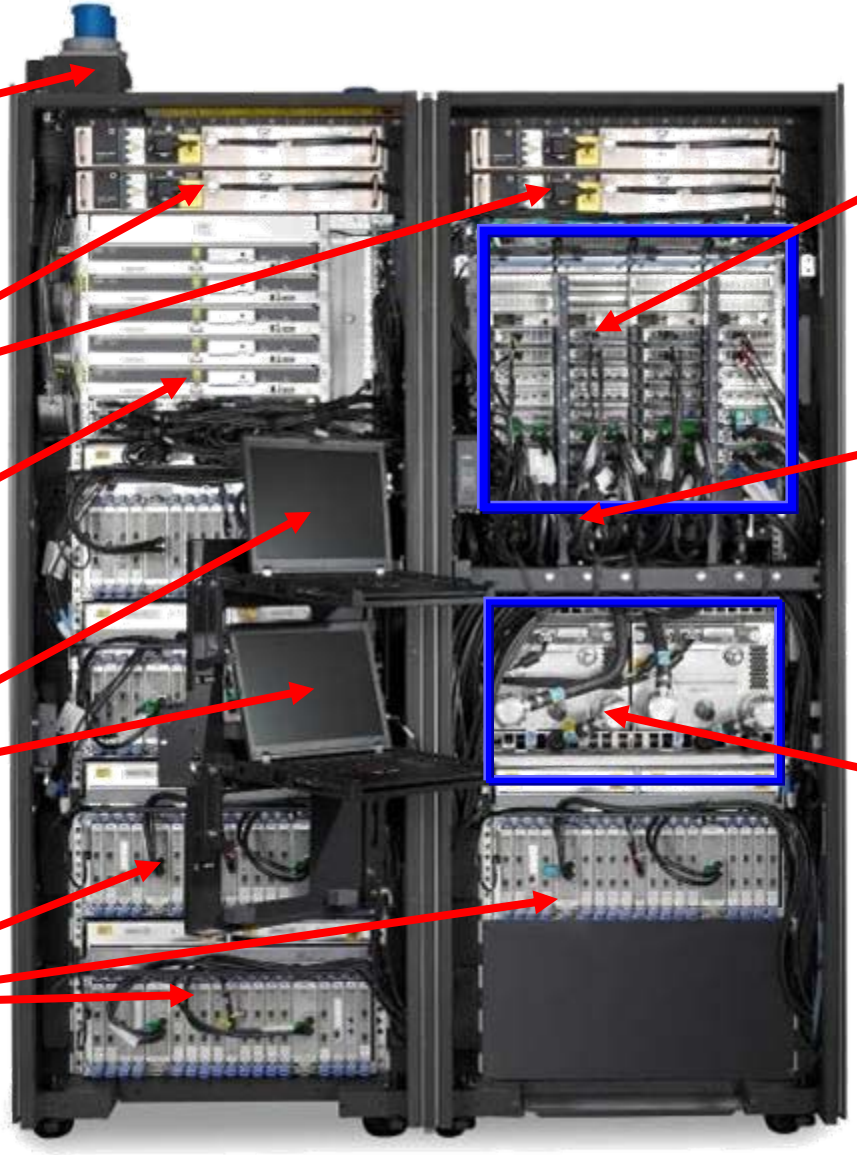

Overhead Power Cables (option)

Internal Batteries (option)

Power Supplies

2 x Support Elements

PCIe I/O drawers (Maximum 5 for zEC12)

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

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

N+1 Radiator-based Air Cooling Unit

Optional FICON LX Fiber Quick Connect (FQC) not shown

zEC12 Connectivity for I/O and Coupling

Note: Optional and disruptive "STI Rebalance" FC 2400 on model upgrade (book add) was eliminated on z196. All fanout types in zEC12 and z196 are concurrently rebalanced on a model upgrade.

- **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)
  - HCA2-O LR two 1x InfiniBand DDR links (CF)
  - **HCA3-O two 12x InfiniBand DDR links**
  - **HCA3-O LR four 1x InfiniBand DDR links**
  -

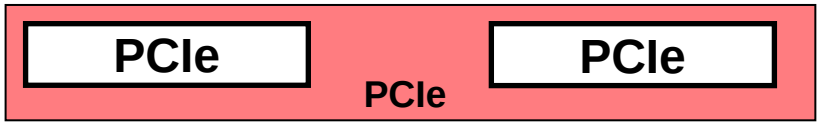

- **PCIe fanout – PCIe I/O Interconnect links**
  Supports two copper cable PCIe 8 GBps interconnects to two 8-card PCIe I/O domain multiplexers.
  Always plugged in pairs for redundancy.

- **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 only)
  Always plugged in pairs for redundancy.

zEC12 I/O Features supported

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

Supported features

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

PCIe I/O drawer

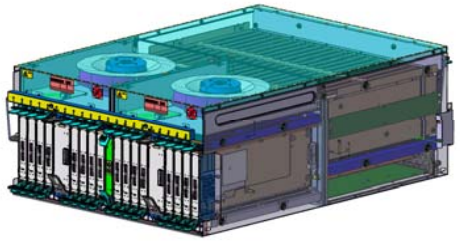

32 I/O slots

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

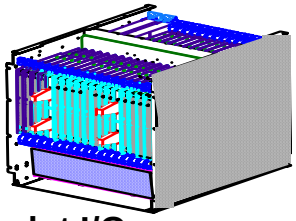

28 slot I/O cage

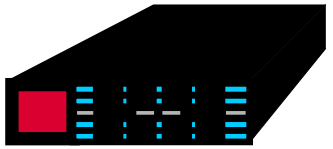

8 slot I/O drawer

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

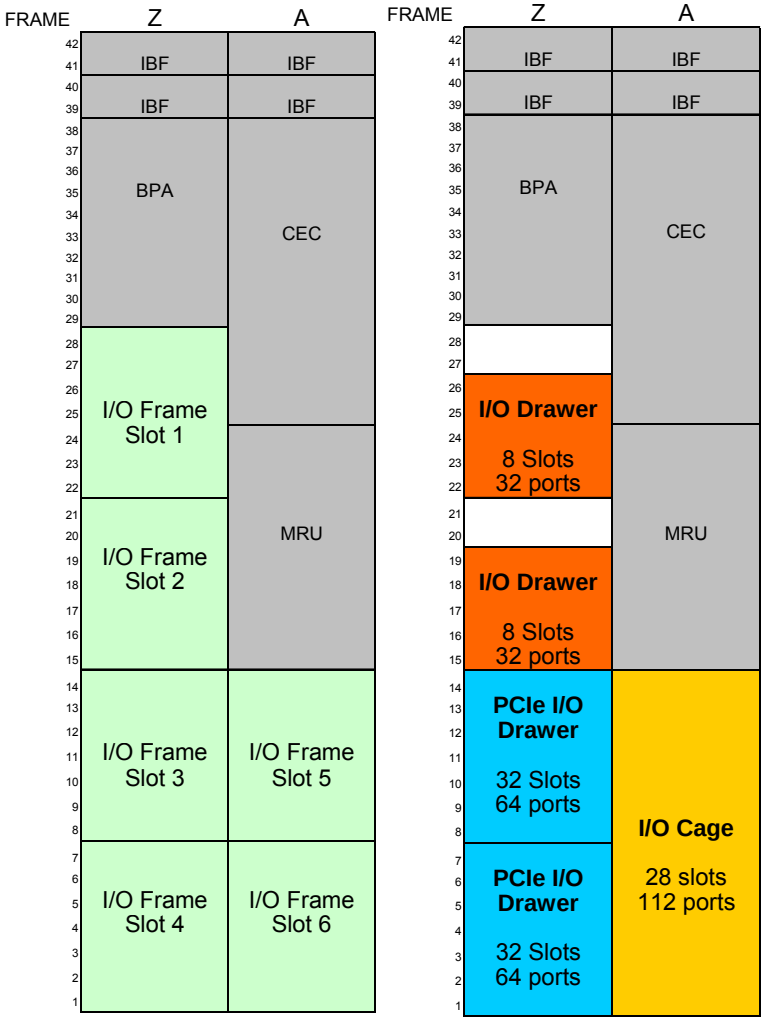

- An **I/O frame slot** is a physical location in the A or Z frame for an I/O cage, I/O drawer or PCIe I/O drawer to be inserted = 7u
- **PCIe I/O drawer** uses 1 I/O frame slot = 7u
  - 32 two port I/O slots = 64 ports
  - 5 drawers maximum
- **I/O cage** uses 2 I/O frame slots = 14u
  - 28 four port 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.**

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 20º 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

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 is 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**)

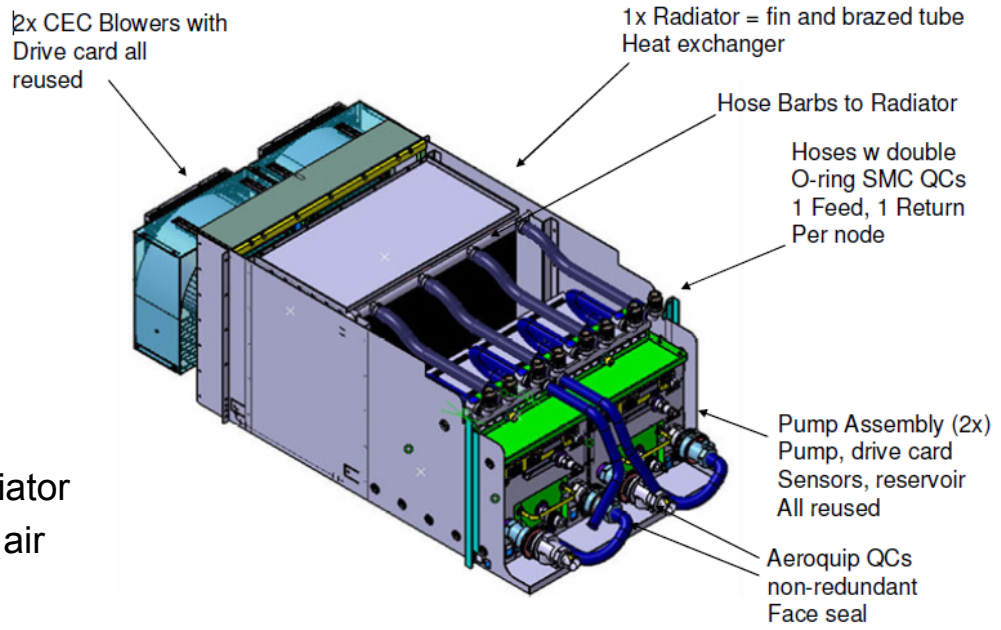

- Normal operation design:
  - Heat removed by water circulating to the radiator
  - Fans exhaust heat from the radiator to room air

- 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

zEC12 Hardware Management Console

- **zEC12 requires HMC FC 0091 with 16 GB memory and Driver 12 LIC**
  - Can be ordered new or carried forward from z10 or z196 and upgraded
  - An earlier System z server can have its FC 0091 HMCs upgraded by ordering ECA 309 to the server
    - Upgrades all FC 0091 HMC features of the server to 16 GB and Driver 12 LIC
  - Older HMCs (eg FC 0090) can not upgraded to control zEC12

- **HMC Display Support**
  - 22 inch flat panel FC 6096 (Carry forward or new build)

- **HMC 1000BASE-T Ethernet LAN Switch – No longer offered**
  - FC 0070 10/100/1000BASE-T switches – (Carry Forward Only)
  - Alternative: Compatible switches provided by the customer

- **HMC application LIC at Driver 12 for zEC12/zBX Model 3 does NOT support dial modem**
  (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)
  - Modems on installed HMC FC 0091 hardware will not work with the HMC application LIC required to support zEC12 and zBX Model 3
  - Also applies to new HMC FC 0091 orders for z196 and z114 shipped after September 19, 2012

THANK YOU

Discussion and questions...