IBM zEnterprise Hardware News

7th European GSE /IBM TU for z/VSE, z/VM and Linux on System z, Hamburg, 30.9.2013-2.10.2013

Dr. Manfred Gnirss
Senior IT Specialist
IBM Client Center, IBM Germany Lab
gnirss@de.ibm.com

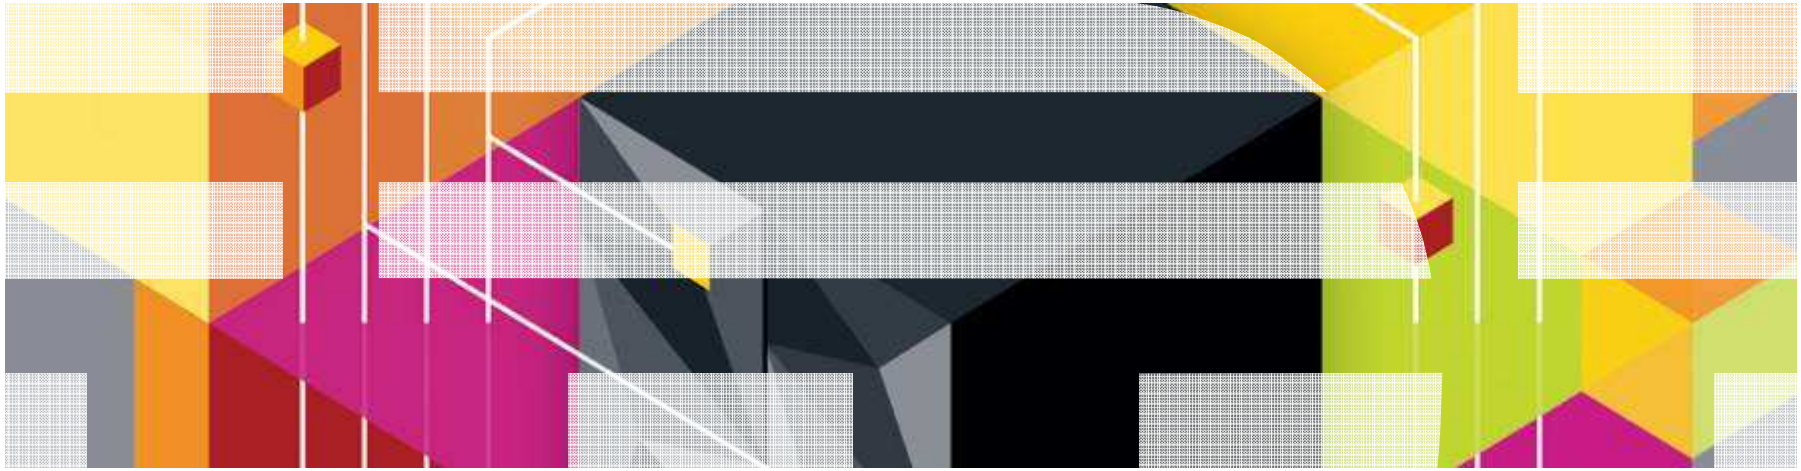

Trademarks

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

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

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

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

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

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads 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). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT").
No other workload processing is authorized for execution on an SE.

IBM offers SEs 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.

The Momentum is Growing around zEnterprise…

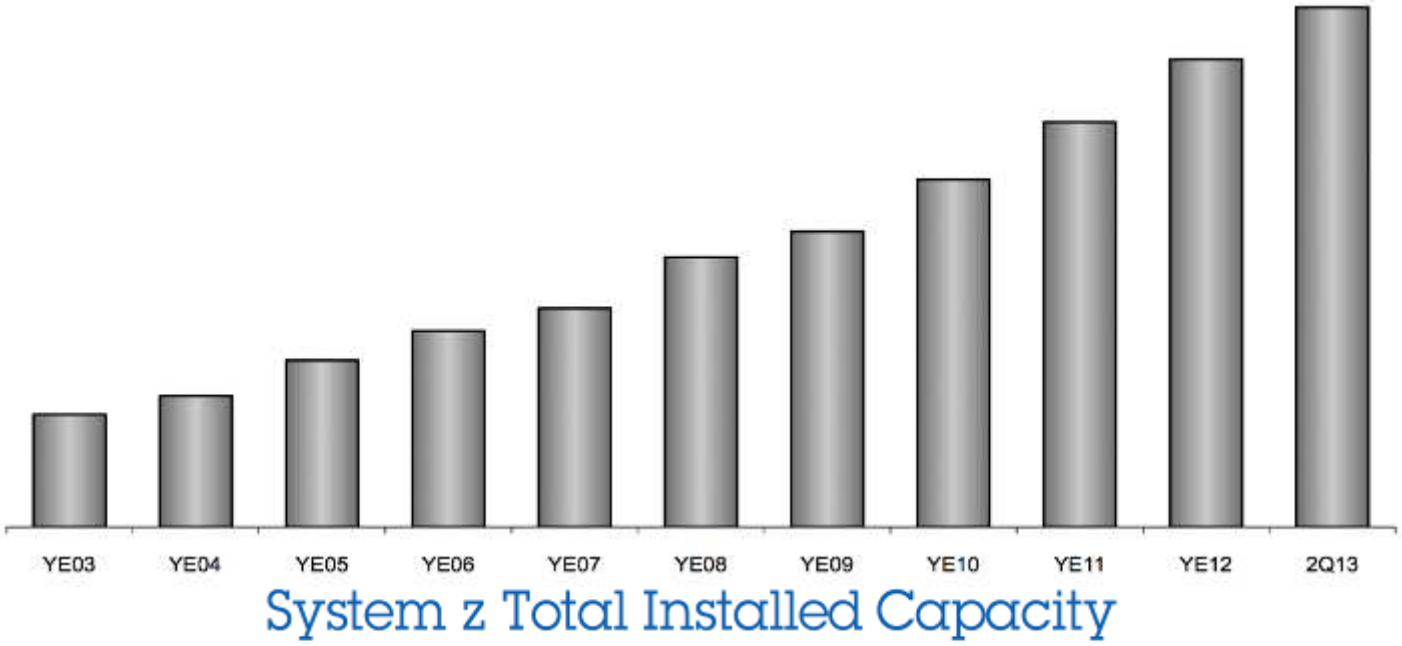

23%

MIPS growth and 10% revenue growth year-to-year

210+

new accounts since 3Q10 zEnterprise launch, with 1/3+ in growth markets

275+

hybrid computing units shipped since 3Q10

3/4+

of Top 100 clients have installed IFLs

1,000+

schools in 67 countries are part of the IBM Academic Initiative for System z

7,400+

ISV apps run on IBM System z; 55 new ISVs added in 1H13

[IFL = Linux-on-z Only Engine]

As of 2Q13

IBM zEnterprise System

An optimized system

Semiconductor Technology

Microprocessor Design

Systems Design

Virtualization & Operating Systems

Compilers & Java Virtual Machine

Optimized Middleware

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

IBM zEnterprise 196 (2817)

- Announced 7/10 – Server with up to 96 PU cores
- 5 models – Up to 80-way
- Granular Offerings for up to 15 CPs
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP, zIIP
- On Demand Capabilities
  - CoD, CIU, CBU, On/Off CoD, CPE
- Memory – up to 3 TB for Server and up to 1 TB per LPAR
  - 16 GB Fixed HSA
- Channels
  - PCIe bus
  - Four LCSSs
  - 3 Subchannel Sets
  - MIDAW facility
  - Up to 240 ESCON channels
  - Up to 288 FICON channels
  - FICON Express8 and 8S
  - zHPF
  - OSA 10 GbE, GbE, 1000BASE-T
  - InfiniBand Coupling Links
- Configurable Crypto Express3
- Parallel Sysplex clustering
- HiperSockets – up to 32
- Up to 60 logical partitions
- Enhanced Availability
- Unified Resource Manager
- Operating Systems
  - z/OS, z/VM, z/VSE, z/TPF, Linux on System z

IBM zEnterprise BladeCenter Extension (2458)

- Announced 7/10
- Model 002 for z196 or z114
- zBX Racks with:
  - BladeCenter Chassis
  - N + 1 components
  - Blades
  - Top of Rack Switches
  - 8 Gb FC Switches
  - Power Units
  - Advance 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 114 (2818)

- Announced 07/11 – Server with up to 18 PU cores
- 2 models – M05 and M10
  - Up to 5 CPs
- High levels of Granularity available
  - 130 Capacity Indicators
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP, zIIP
- On Demand Capabilities
  - CoD, CIU, CBU, On/Off CoD. CPE
- Memory – up to 256 GB for Server
  - 8 GB Fixed HSA
- Channels
  - PCIe bus
  - Two LCSSs
  - 2 Subchannel Sets
  - MIDAW facility
  - Up to 240 ESCON channels
  - Up to 128 FICON channels
  - FICON Express8 and 8S
  - zHPF
  - OSA 10 GbE, GbE, 1000BASE-T
  - InfiniBand Coupling Links
- Configurable Crypto Express3
- Parallel Sysplex clustering
- HiperSockets – up to 32
- Up to 30 logical partitions
- Unified Resource Manager
- Operating Systems
  - z/OS, z/VM, z/VSE, TPF, z/TPF, Linux on System z

New innovations available on zBC12 and zEC12

NEW

Data Compression Acceleration

Reduce CP consumption, free up storage & speed cross platform data exchange

zEDC Express

NEW

High Speed Communication Fabric

Optimize server to server networking with reduced latency and lower CPU overhead

10GbE RoCE Express

ENHANCED

Flash Technology Exploitation

Improve availability and performance during critical workload transitions, now with dynamic reconfiguration; Coupling Facility exploitation (SOD)

Flash Express

ENHANCED

Proactive Systems Health Analytics

Increase availability by detecting unusual application or system behaviors for faster problem resolution before they disrupt business

IBM zAware

NEW

Hybrid Computing Enhancements

x86 blade resource optimization; New alert & notification for blade virtual servers; Latest x86 OS support; Expanding future roadmap

zBX Mod 003; zManager Automate; Ensemble Availability Manager

NEW **zEnterprise compilers (COBOL, PL/I, C/C++)** provide an optimized application infrastructure for increased software performance

zEC12 Overview

- Machine Type
  - 2827
- 5 Models
  - H20, H43, H66, H89 and HA1
- Processor Units (PUs)
  - 27 (30 for HA1) PU cores per book
  - Up to 16 SAPs per system, standard
  - 2 spares designated per system
  - Dependant on the H/W model - up to 20, 43, 66,89, 101 PU cores available for characterization
    - Central Processors (CPs), Internal Coupling Facility (ICFs), Integrated Facility for Linux (IFLs), System z Application Assist Processors (zAAPs), System z Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs) and Integrated Firmware Processor (IFP)
  - Sub-capacity available for up to 20 CPs
    - 3 sub-capacity points
- Memory
  - RAIM Memory design
  - System Minimum of 32 GB
  - Up to 768 GB per book
  - Up to 3 TB for System and up to 1 TB per LPAR
    - 32 GB Fixed HSA, standard
    - 32/64/96/112/128/240/256 GB increments
  - Flash Express
- I/O
  - 6 GBps I/O Interconnects – carry forward only
  - Up to 48 PCIe Gen2 interconnects per System @ 8 GBps each
  - Up to 4 Logical Channel Subsystems (LCSSs)
    - Up to 3 Sub-channel sets per LCSS
- STP - optional (No ETR)

zEC12 New Build Radiator-based Air cooled – Under the covers (Model H89 and HA1) Front view

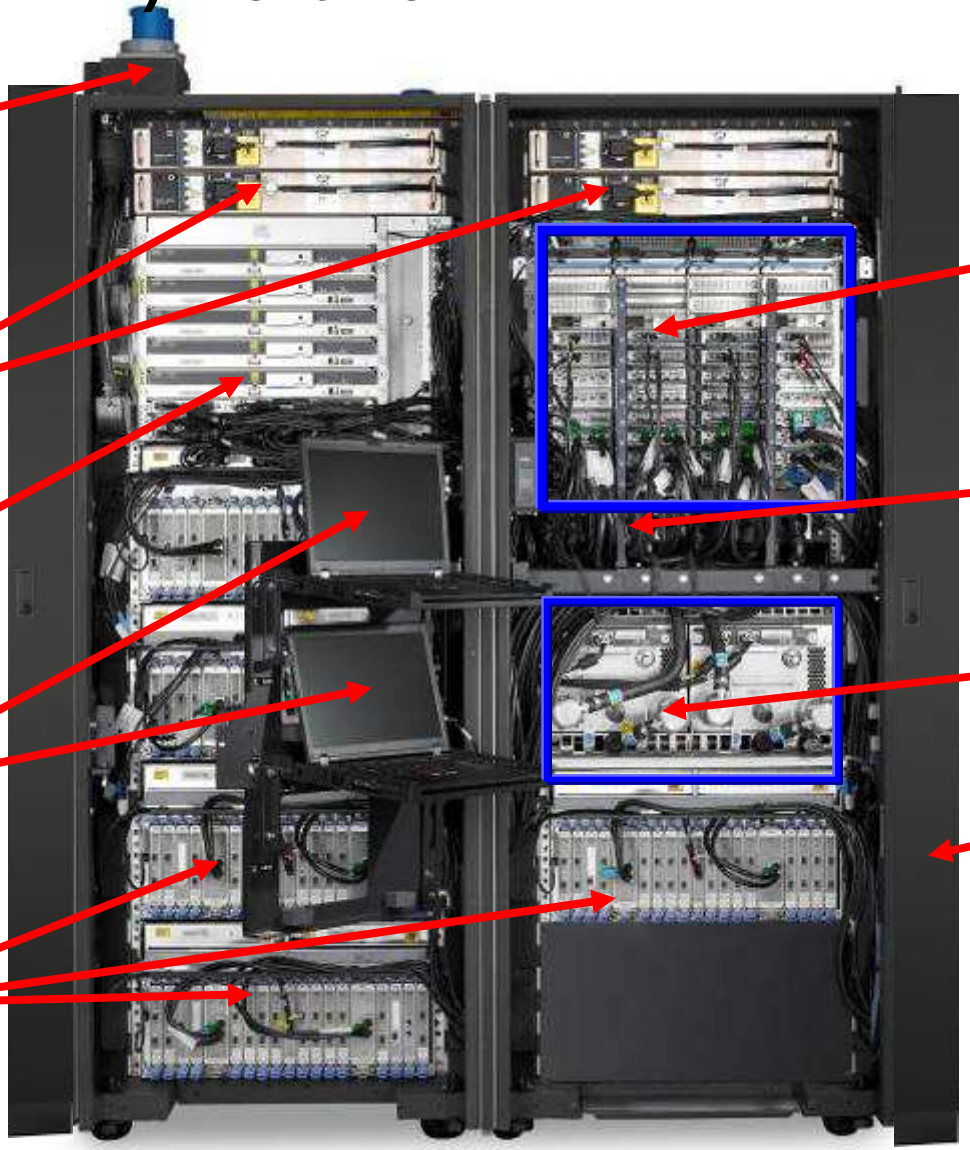

Overhead Power Cables (option)

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 FSP cage controller cards

Radiator with N+1 pumps, blowers and motors

Overhead I/O feature is a co-req for overhead power option

Optional FICON LX Fiber Quick Connect (FQC) not shown

zEC12 Continues the CMOS Mainframe Heritage Begun in 1994

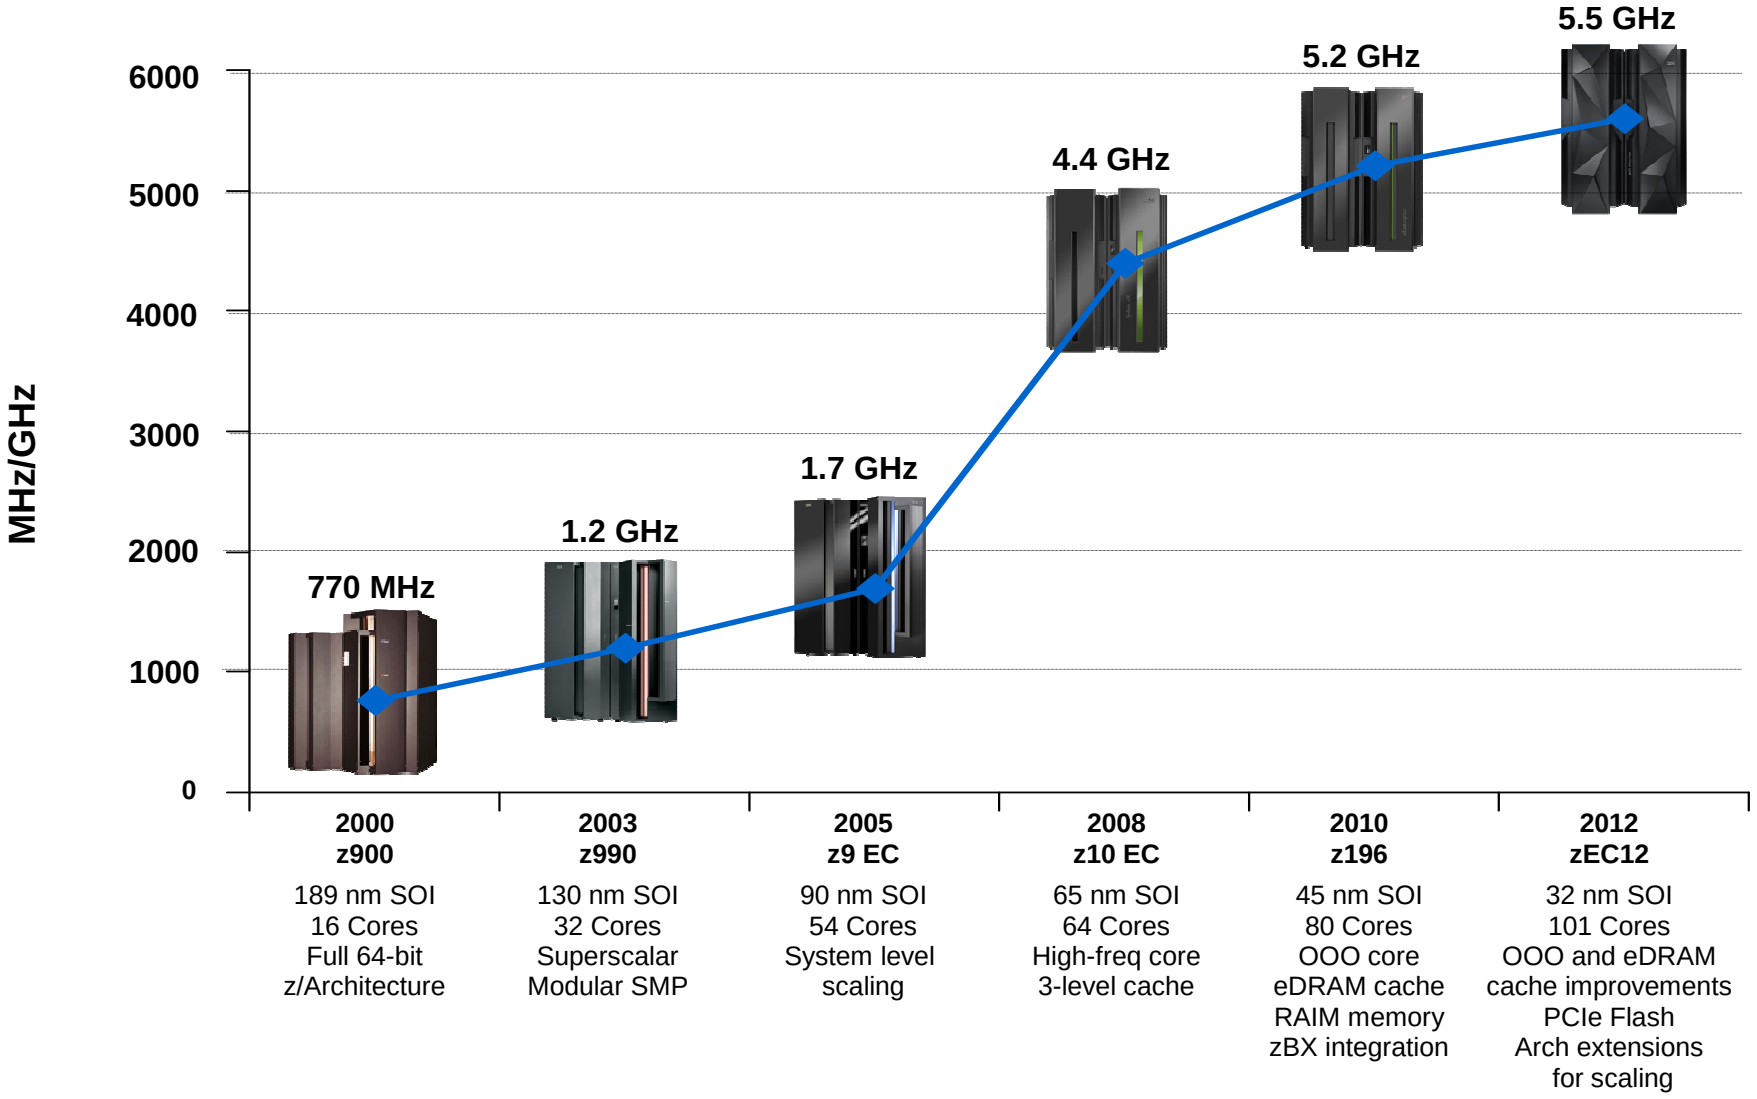

Processor chip optimized for software performance

Exploited by Java, PL/I, compilers, DB2, more

- **Our leadership in microprocessor design supports a boost in performance for all workloads**
  - Second generation out of order execution design
  - Multi-level branch prediction supports complex workloads
- **Larger caches to optimize data serving environments**
  - Almost 2x on chip and 2x additional on book
- **New hardware functions optimized for software performance**
  - ***Transactional Execution Facility*** for parallelism and scale
  - ***Runtime Instrumentation Facility*** is intended to help reduce Java overhead
  - ***2 GB page frames*** are intended to offer performance Improvements for DB2 buffer pools and Java heaps
  - New IBM Enterprise PL/I compiler is planned to exploit and get a performance boost from ***decimal format conversions facility***
  - Up to ***30% improvement in IMS™ throughput*** due to faster CPU, cache and compliers[1]
  - Workloads leveraging ***Flash Express*** with ***Pageable Large Pages*** can see up to a ***8% price performance improvement***[3] over the z196

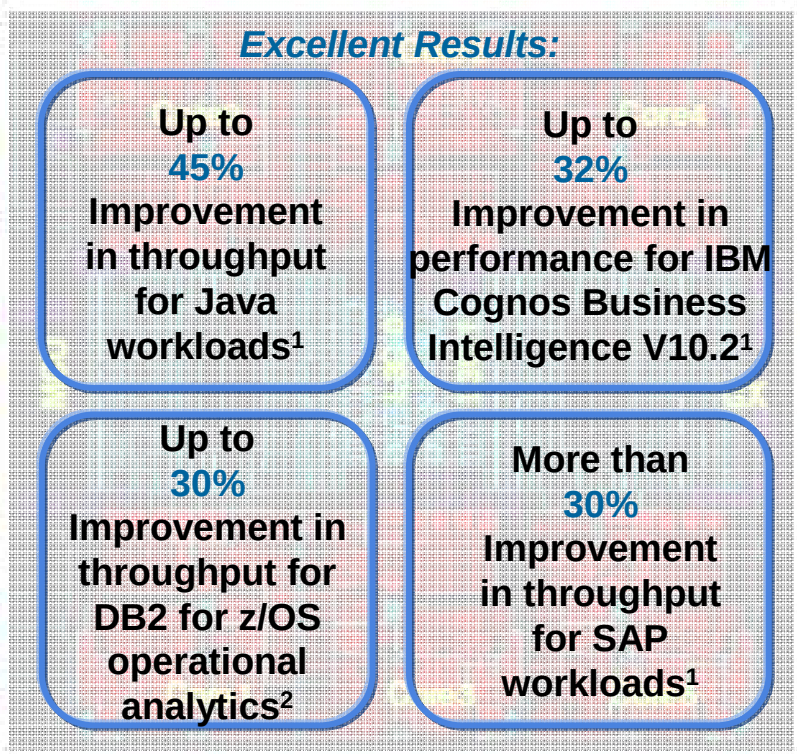

[1] Based on preliminary internal measurements and projections
[2] As measured by the IBM 9700 Solution Integration Center. The measured operational BI workload consists of 56 concurrent users executing a fixed set of 160,860 Cognos reports . Compared DB2 v10 workload running on IBM's z196 w/10 processors to an zEC12 w/10 processors
[3] based on average 5% discount for zEC12 workloads under the AWLC pricing plus up to 3% more performance per MSU with Flash Express.

IBM System z:
Design Comparison for High End Systems

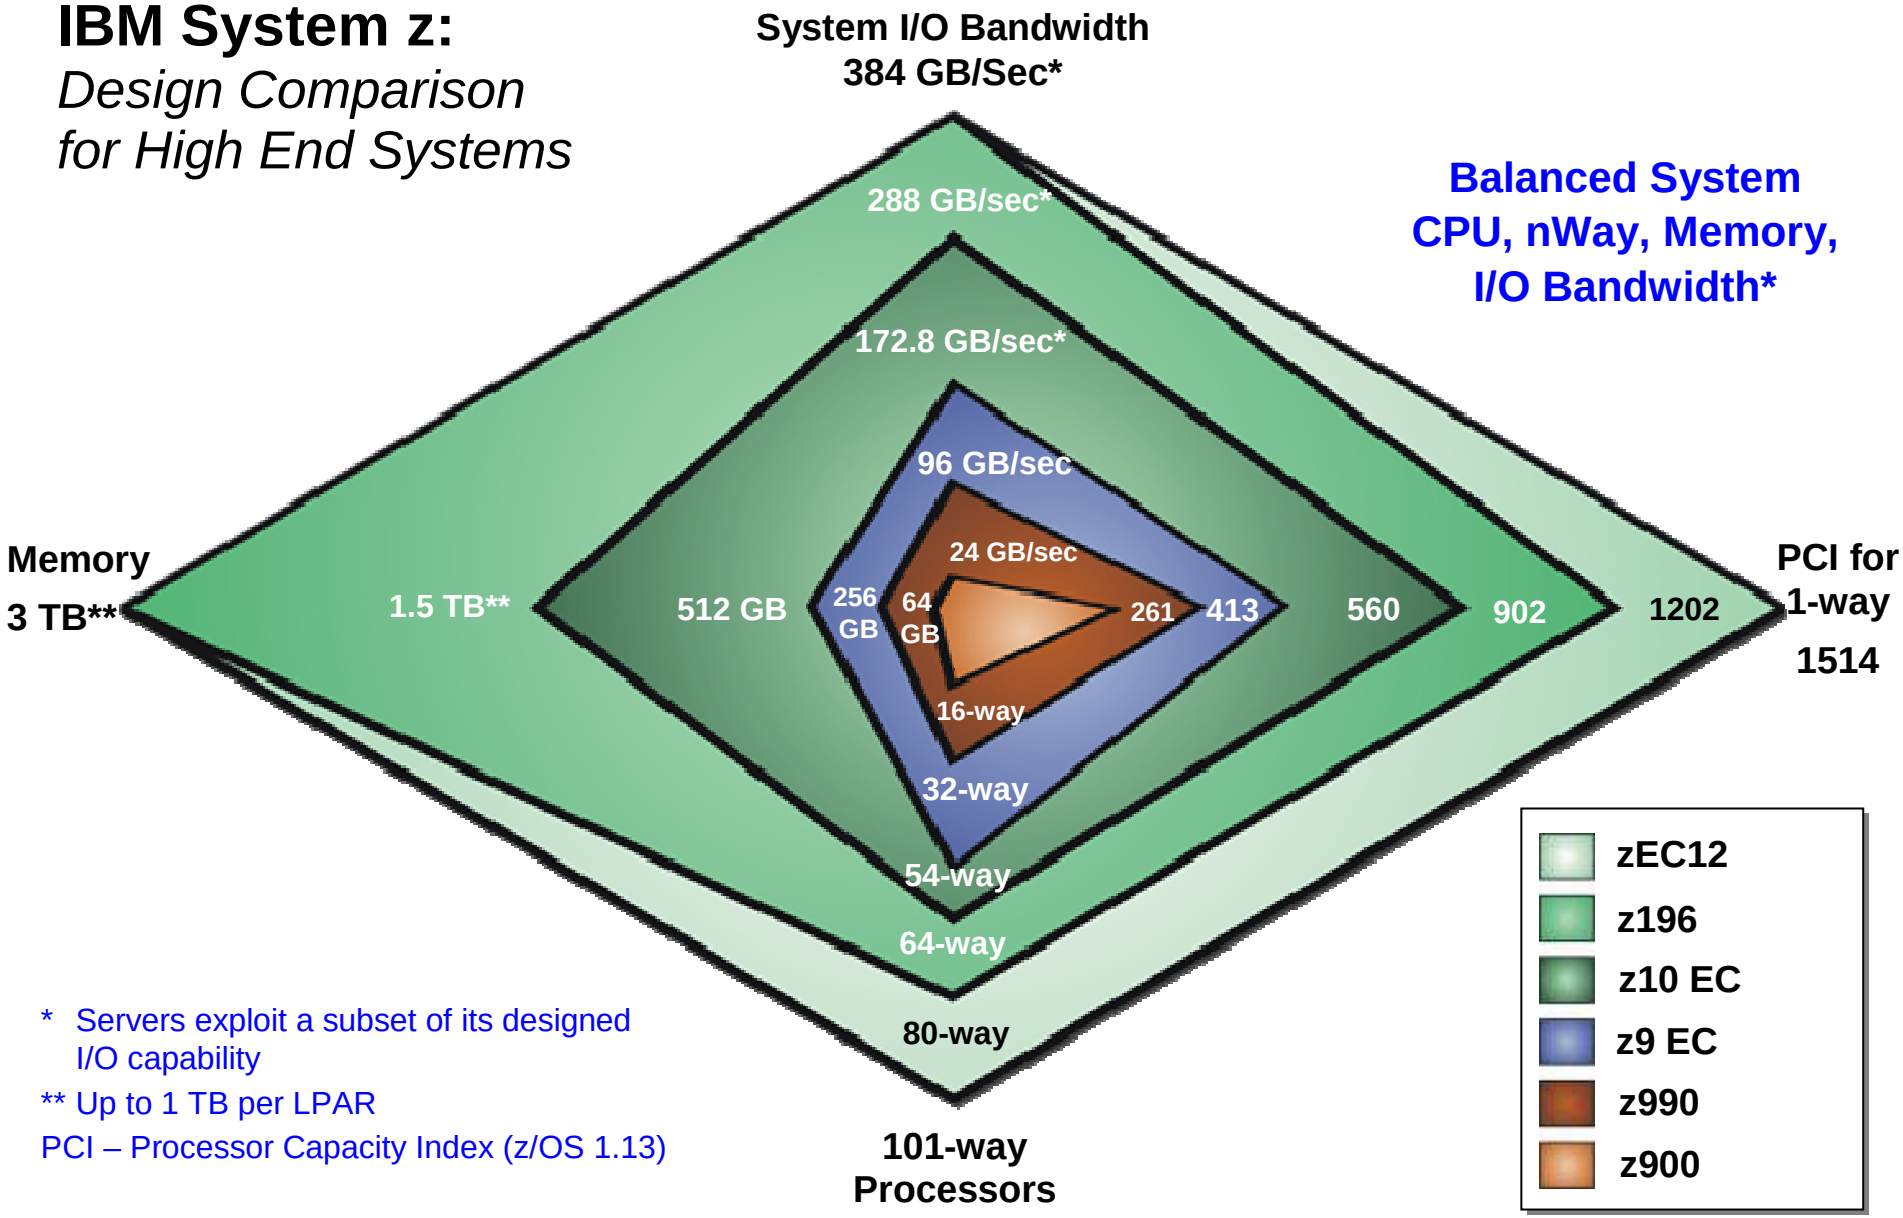

* Servers exploit a subset of its designed I/O capability

** Up to 1 TB per LPAR

PCI – Processor Capacity Index (z/OS 1.13)

zEC12 Hexa Core PU Chip Details

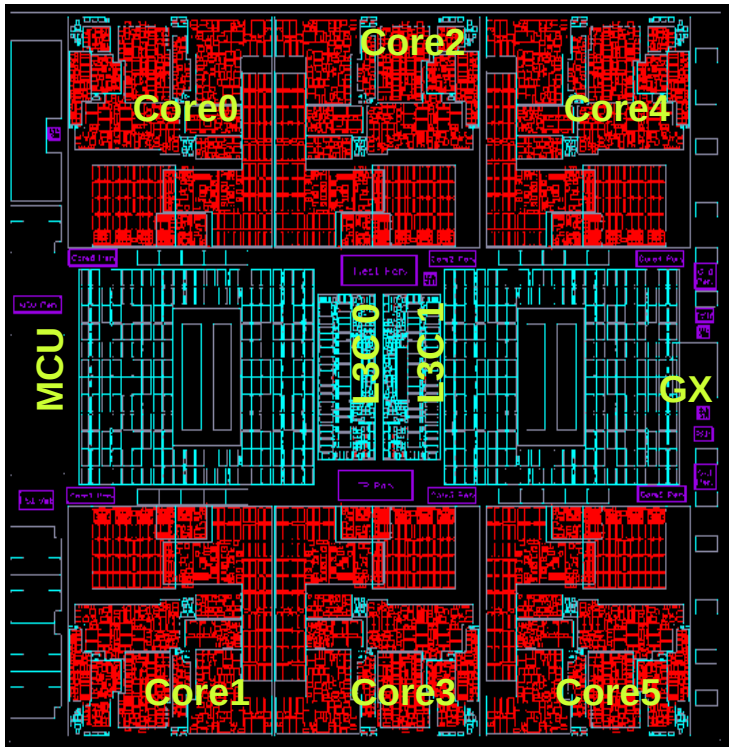

- 13S 32nm SOI Technology
  - 15 layers of metal
  - 7.68 km wire
- 2.75 Billion Transistors

- Chip Area
  - 597 mm^2
  - 23.7mm x 25.2mm
  - 10000+ Power pins
  - 1071 signal I/Os

- Up to Six active cores 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-processors (COP) per core
  - Crypto & compression accelerators
  - Includes 16KB cache
- On chip 48 MB eDRAM L3 Cache
  - Shared by all six cores
- Interface to SC chip / L4 cache
  - 44 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

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

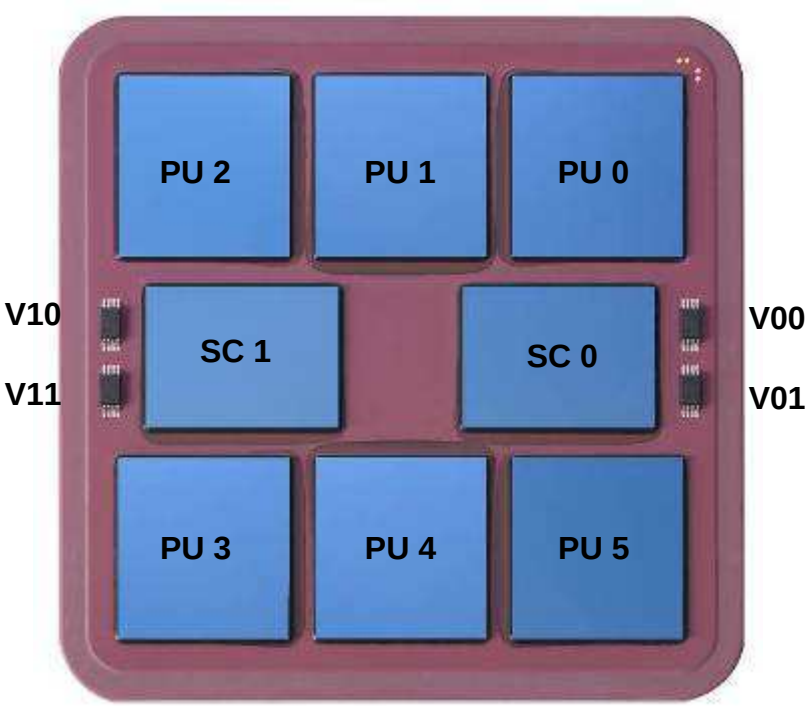

- CMOS 13s chip Technology
  - PU, SC, S chips, 32nm
  - 6 PU chips/MCM – Each up to 6 active cores
    - 23.7 mm x 25.2 mm
    - 2.75 billion transistors/PU chip
    - L1 cache/PU core
      - 64 KB I-cache
      - 96 KB D-cache
    - L2 cache/PU core
      - 1 MB I-cache
      - 1 MB D-cache
    - L3 cache shared by 6 PUs per chip
      - 48 MB
    - 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 other MCMs
  - 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

zEC12 Processor Unit allocation/usage

Model	Books/ PUs	CPs	IFLs uIFLs	zAAPs	zIIPs	ICFs	Std SAPs	Optional SAPs	Std. Spares	IFP
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	0-66*	0-66*	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	0-66*	0-66*	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
  - Each MCM uses PU chips with a combination of 4, 5 and 6 active cores
- New zIIP/zAAP to CP ratio – changed from 1:1 to 2:1
- The maximum number of logical ICFs or logical CPs supported in a CF LPAR is 16
- IFP – integrated firmware processor. The IFP is standard and not defined by the customer. It is used for infrastructure management of 10GbE RoCE Express, zEDC Express
- Concurrent Book Add is available to upgrade from model H20 to model H89

* Based on the 2:1 ratio, the maximum number of zIIPs or zAAPs possible with the largest number of CPs for this Model

Notes: 1. At least one CP, IFL, or ICF must be purchased in every machine
2. zIIP/zAAP to CP ratio is 2:1. Applies to 'banked' CP capacity also.
3. "uIFL" stands for Unassigned IFL

zBC12 Overview

- **Machine Type**
  - 2828
- **2 Models**
  - H06 and H13
  - Single frame, air cooled
  - Non-raised floor option available
  - Overhead Cabling and DC Power Options
- **Processor Units (PUs)**
  - 9 PU cores (using 4 and 5 core PU SCMs) per processor drawer (One for H06 and two for H13)
  - 2 SAPs per system, standard
  - 2 spares designated for Model H13
  - 1 Integrated firmware processor
  - Dependant on the H/W model - up to 6 or 13 PU cores available for characterization
    - Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z Application Assist Processors (zAAPs), System z Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs), Integrated firmware processor (IFP)
  - 156 capacity settings
- **Memory**
  - Up to 512 GB for System including HSA
    - System minimum = 8 GB (Model H06), 16 GB (Model H13)
    - 16 GB fixed HSA separately managed
    - RAIM standard
    - Maximum for customer use 496 GB (Model H13)
    - Increments of 8 or 32 GB
    - Flash Express
- **I/O**
  - Support for non-PCIe Channel features
  - PCIe Gen2 channel subsystem
    - Up to 64 PCIe Channel features
  - Up to 2 Logical Channel Subsystems (LCSSs)
- **STP - optional** (No ETR)

zBC12 Model H13 – Under the covers

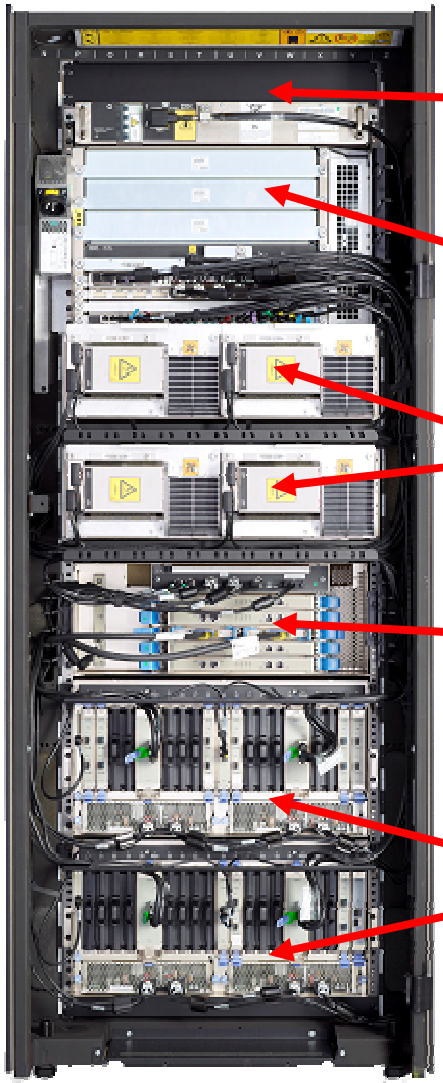

Rear View

Internal Batteries (optional)

Power Supplies

2 x CPC Drawers, Memory & HCAs

I/O Drawer

PCIe I/O drawers

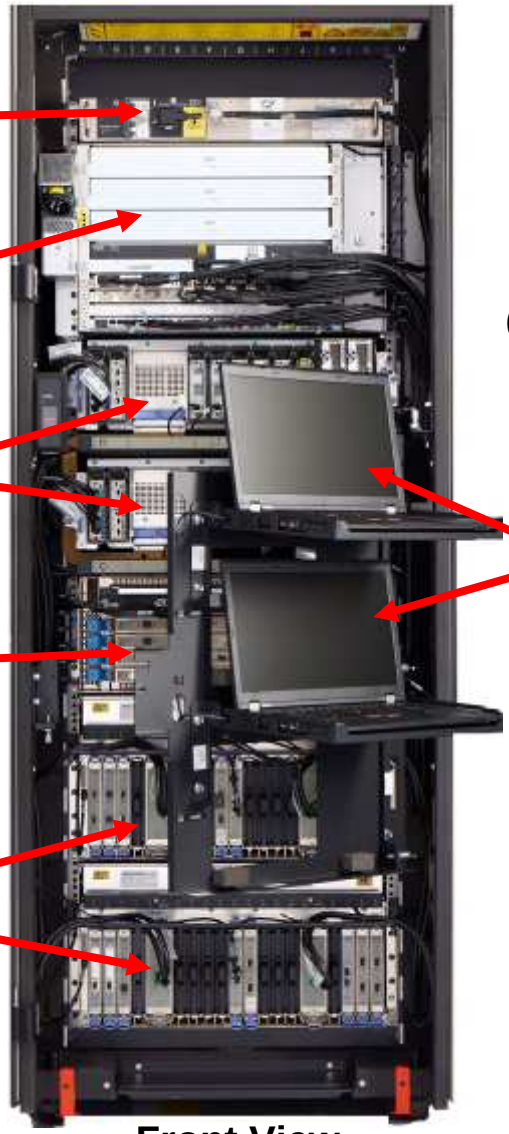

Front View

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

2 x Support Elements

FQC for FICON LX only

zBC12 continues the CMOS Mainframe heritage

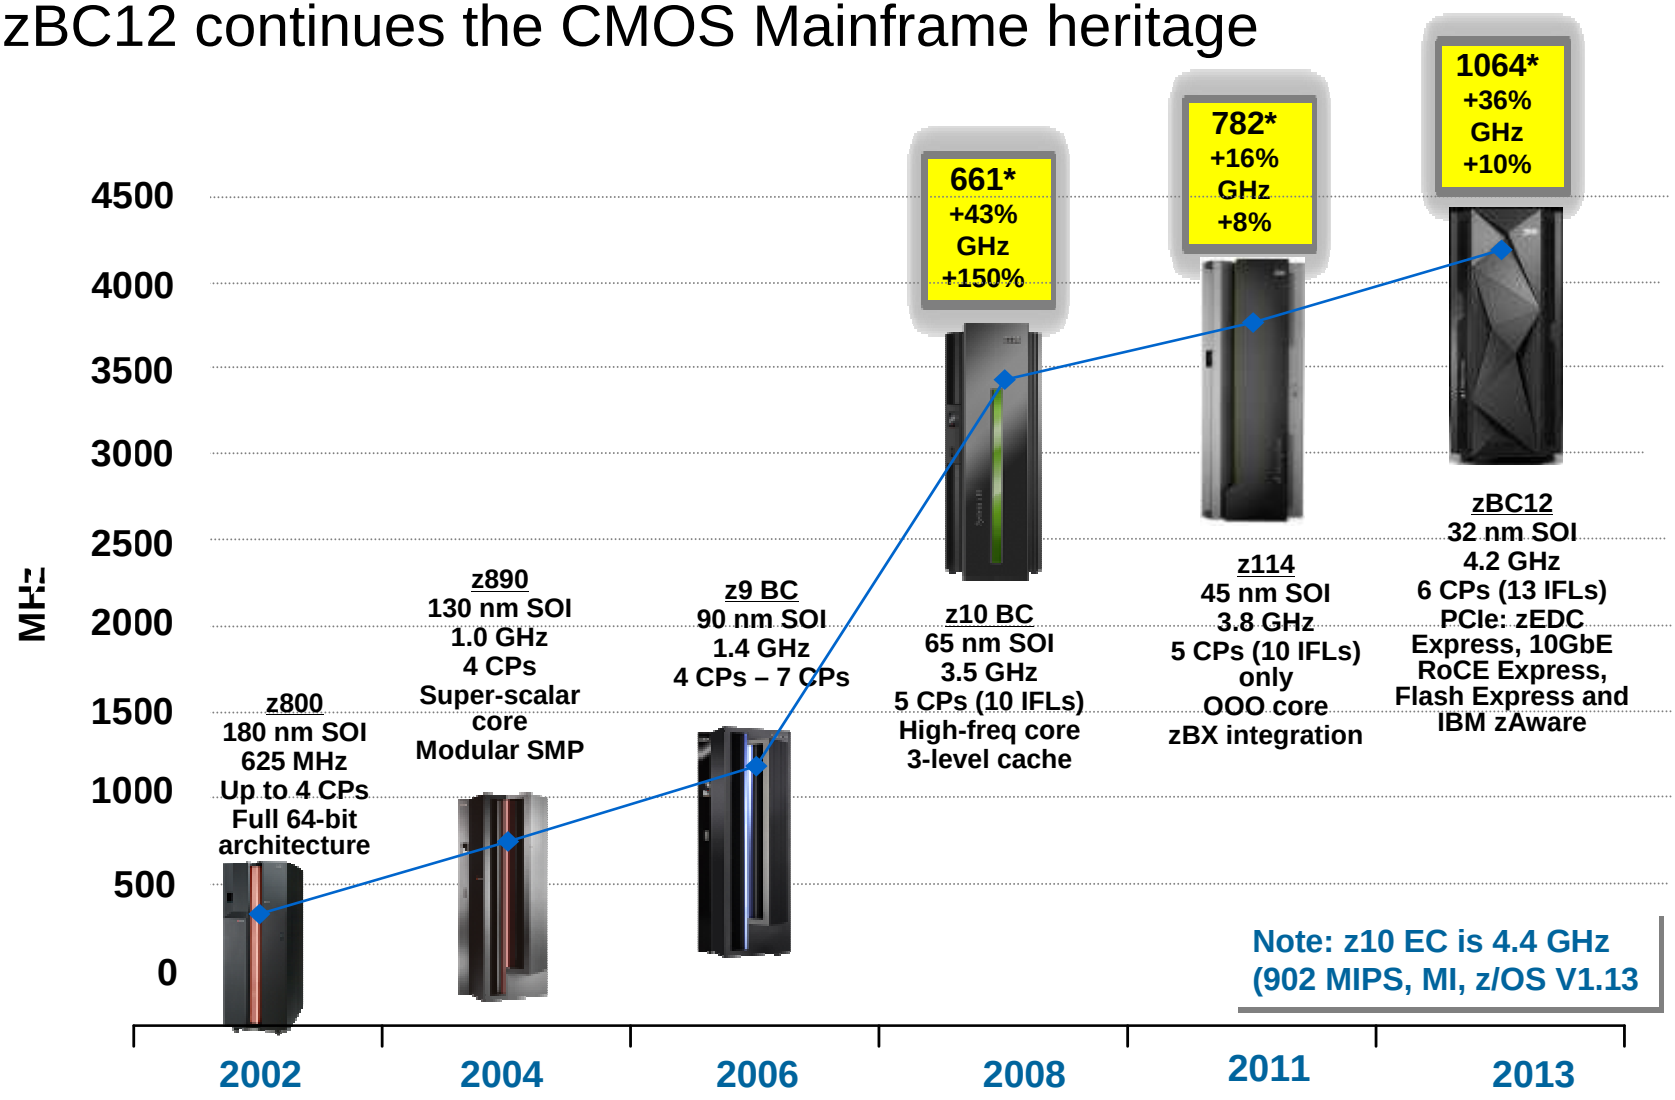

*** MIPS Tables are NOT adequate for making comparisons of System z processors in proposals**

Balanced System Design

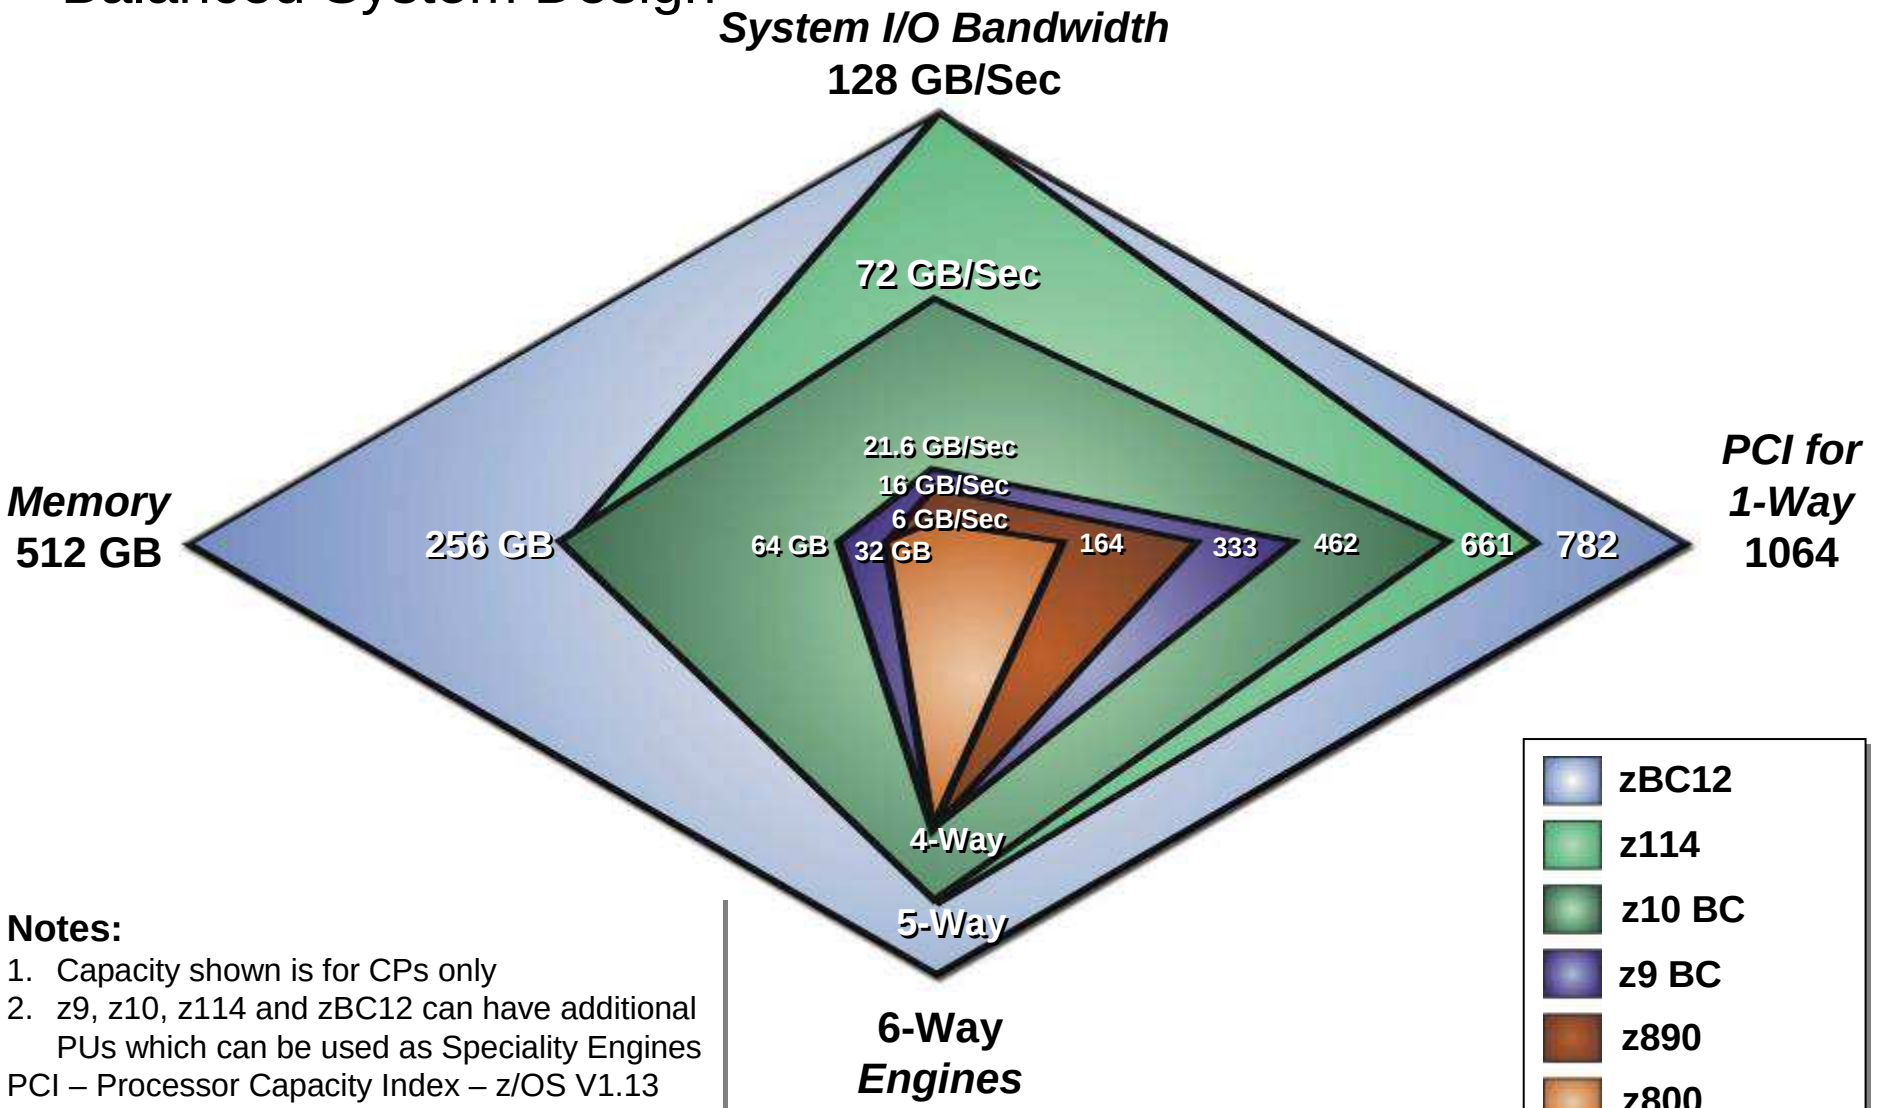

Notes:

1. Capacity shown is for CPs only
2. z9, z10, z114 and zBC12 can have additional PUs which can be used as Speciality Engines

PCI – Processor Capacity Index – z/OS V1.13 (refer to LSPR)

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 ratio of zIIPs and zAAPs for each CP purchased is the same for CPs of any speed.
  - New zIIP/zAAP ratio to CP – changed from 1:1 to 2:1
  - All specialty engines run at full speed
  - Processor Value Unit (PVU) for IFL = 100

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

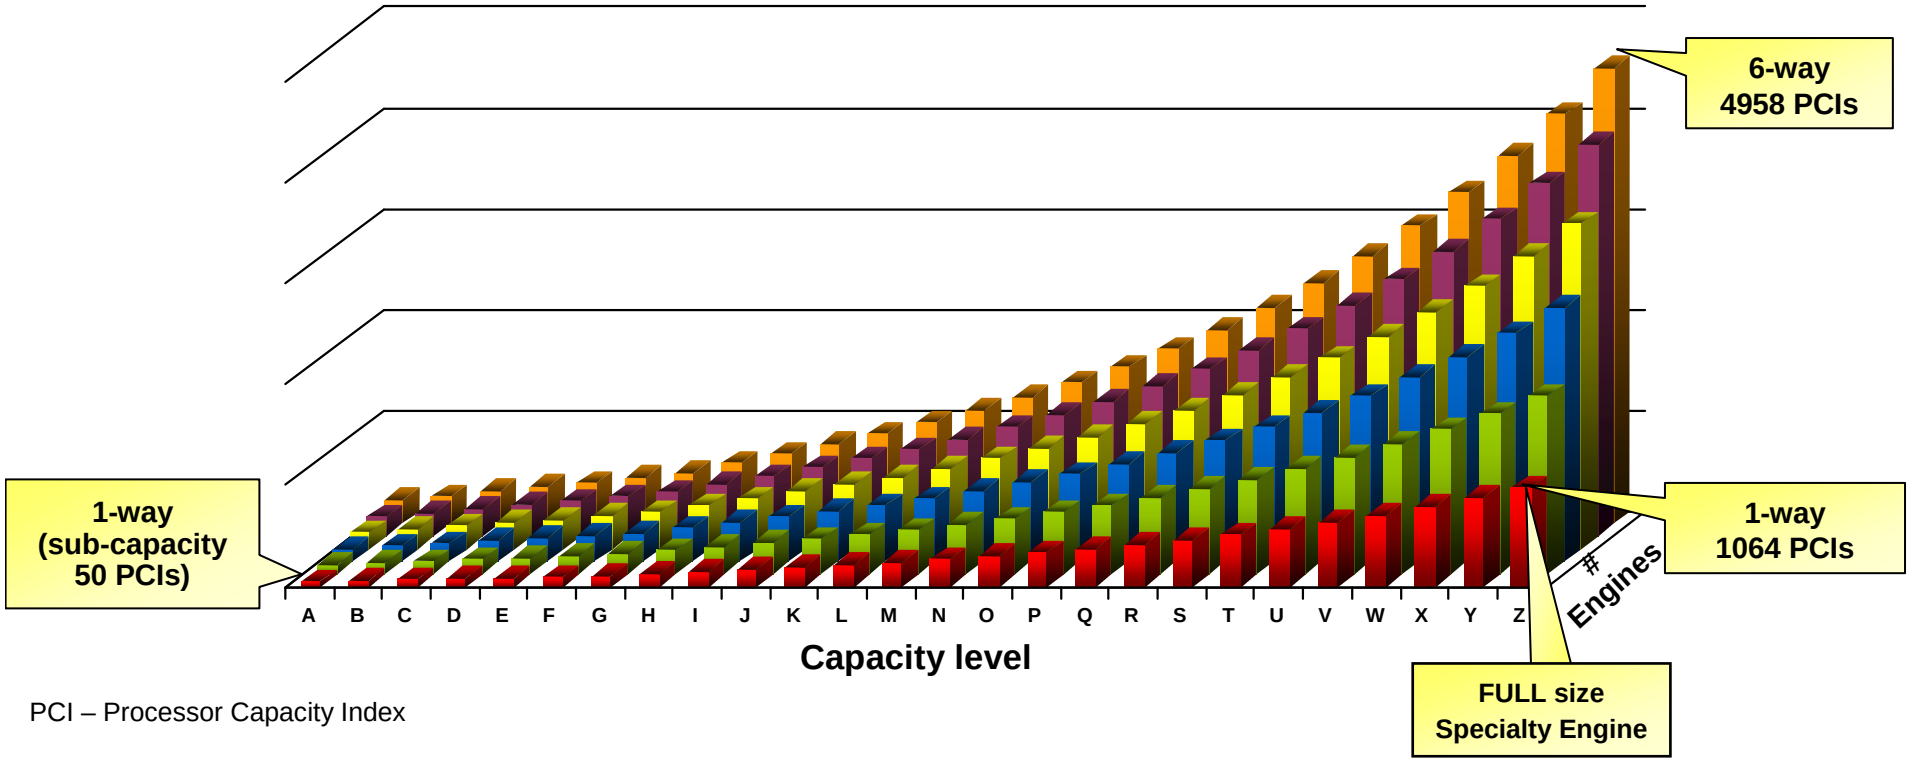

PCI – Processor Capacity Index

zBC12 Hexa Core PU Chip Details

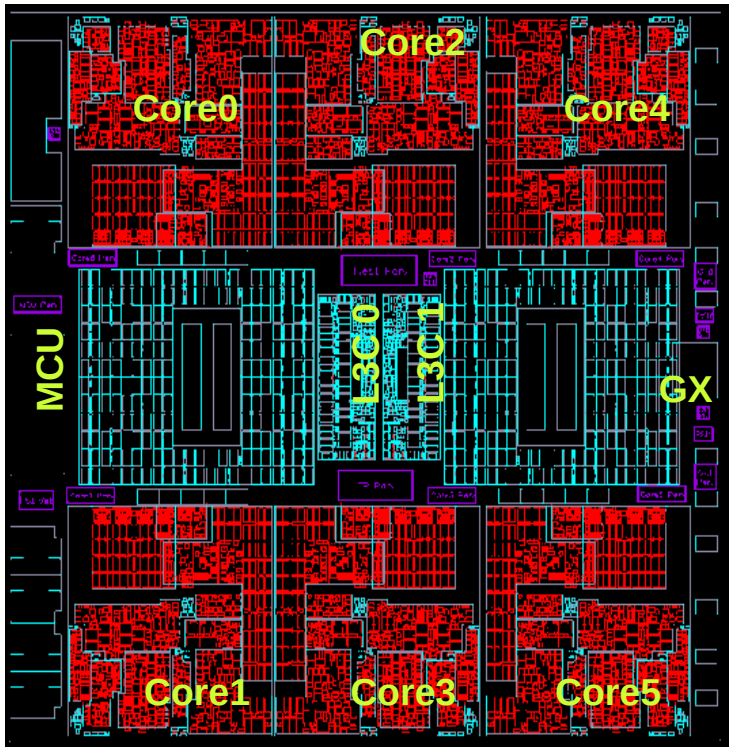

- 13S 32nm SOI Technology
  - 15 layers of metal
  - 7.68 km wire
- 2.75 Billion Transistors

- Chip Area
  - 597 mm^2
  - 23.7mm x 25.2mm
  - 10000+ Power pins
  - 1071 signal I/Os

- Up to Six active cores per chip
  - 4.2 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-processors per core
  - Crypto & compression accelerators
  - Includes 16KB cache
- On chip 24 MB eDRAM L3 Cache
  - Shared by all six cores
- Interface to SC chip / L4 cache
  - 16.8GB/sec to the 2 SCs (4.8 GHz)
- I/O Bus Controller (GX)
  - Interface to Host Channel Adapter (HCA)
- Memory Controller (MC)
  - Interface to controller on memory DIMMs
  - Supports RAIM design

zBC12 PU Details

- Improved Instruction Fetching Unit
  - faster branch prediction
  - 2nd-level branch prediction
  - improved sequential instruction stream delivery
- Improved Out-of-order efficiency
  - better group formation, including regrouping
  - “uncracked” common instructions
  - on-the-fly Culprit/Victim detection for store-load hazards
- Increased Execution/Completion Throughput
  - bigger GCT, speculative completion
  - virtual branch unit
- Innovative Local Data-Cache design
  - store banking in Data-L1
  - unique Data-L2 cache (1M-byte) design
- Dedicated Instruction-L2 cache (1M-byte)
- Optimized Floating-Point Performance
  - Increased physical pool for FPRs
  - fixed-point divide in DFU
- Dedicated Co-Processor
  - Designed for improved start-up latency
  - UTF8<->UTF16 conversion support

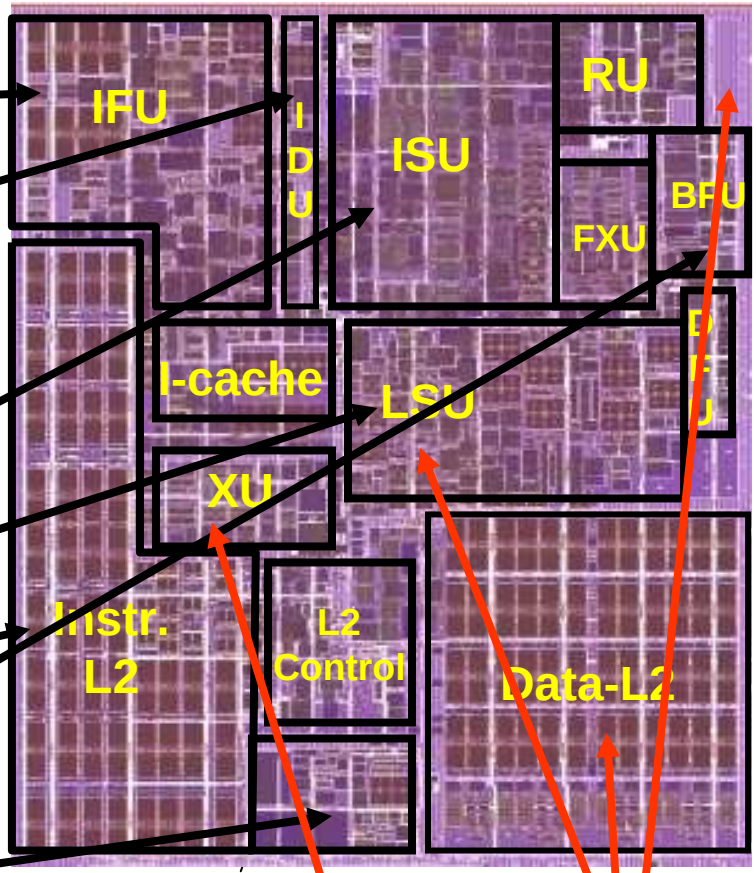

- Main Architectural Extension Support
  - Transactional Memory support
  - Run-time instrumentation support
  - EDAT-2 support

zBC12 SCM vs zEC12 MCM Comparison

zBC12 Single Chip Module (SCM)

- **PU SCM**
  - 50mm x 50mm in size – fully assembled
  - Hexa core chip with 4 or 5 active cores
  - 2 PU SCMs for H06 and 4 PU SCMs for H13
  - PU Chip size 23.7 mm x 25.2 mm
- **SC SCM**
  - 50mm x 50mm in size – fully assembled
  - 1 SC SCM for H06, 2 SC SCMs for H13
  - 192 MB L4 cache per chip
  - SC Chip size 26.72 mm x 19.67 mm

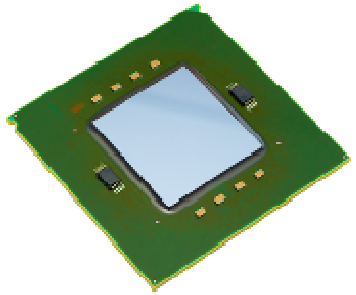

Single PU Chip

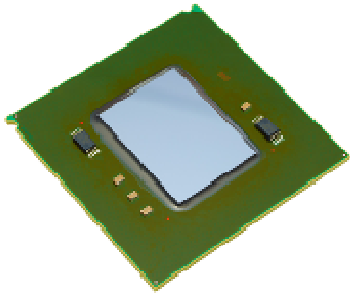

Single SC Chip

zEC12 Multi Chip Module (MCM)

- **MCM**
  - 96mm x 96mm in size
  - 6 PU chips per MCM
    - 6 core chips with 4, 5 or 6 active cores
    - PU Chip size 23.7 mm x 25.2 mm
  - 2 SC chips per MCM
    - 384 MB L4 cache per MCM (in each book)
    - SC Chip size 26.72 mm x 19.67 mm
  - Up to 4 MCMs per System

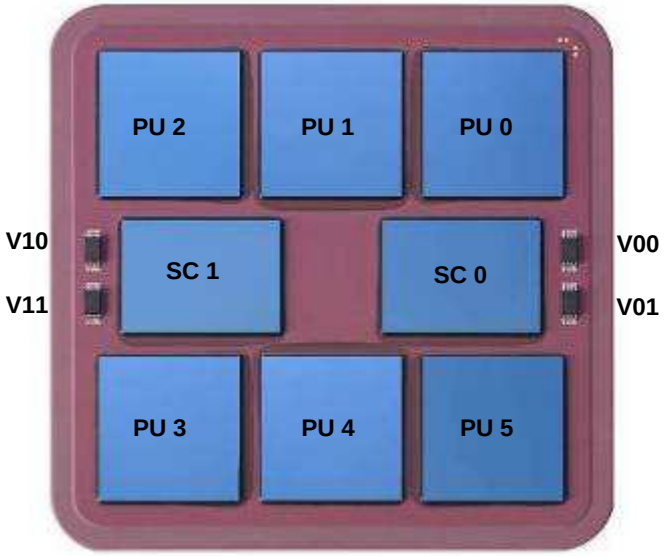

z12BC CPC Drawer (Top View)

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

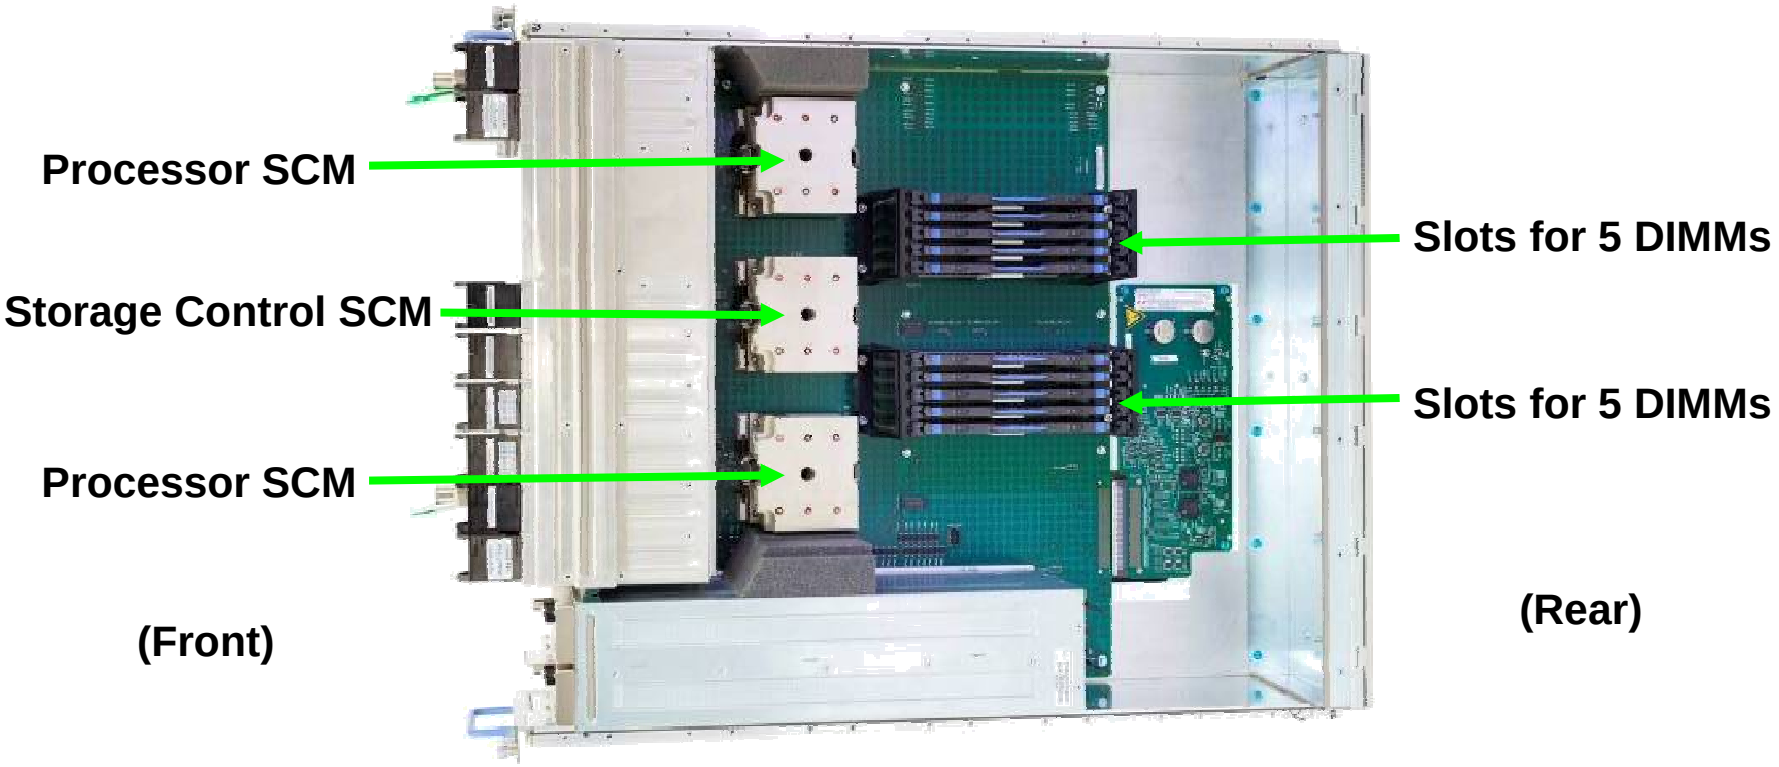

zBC12 Central processor complex (CPC) Drawers (Model H06 and H13)

One CPC Drawer (Model H06)

Two CPC Drawers (Model H13)

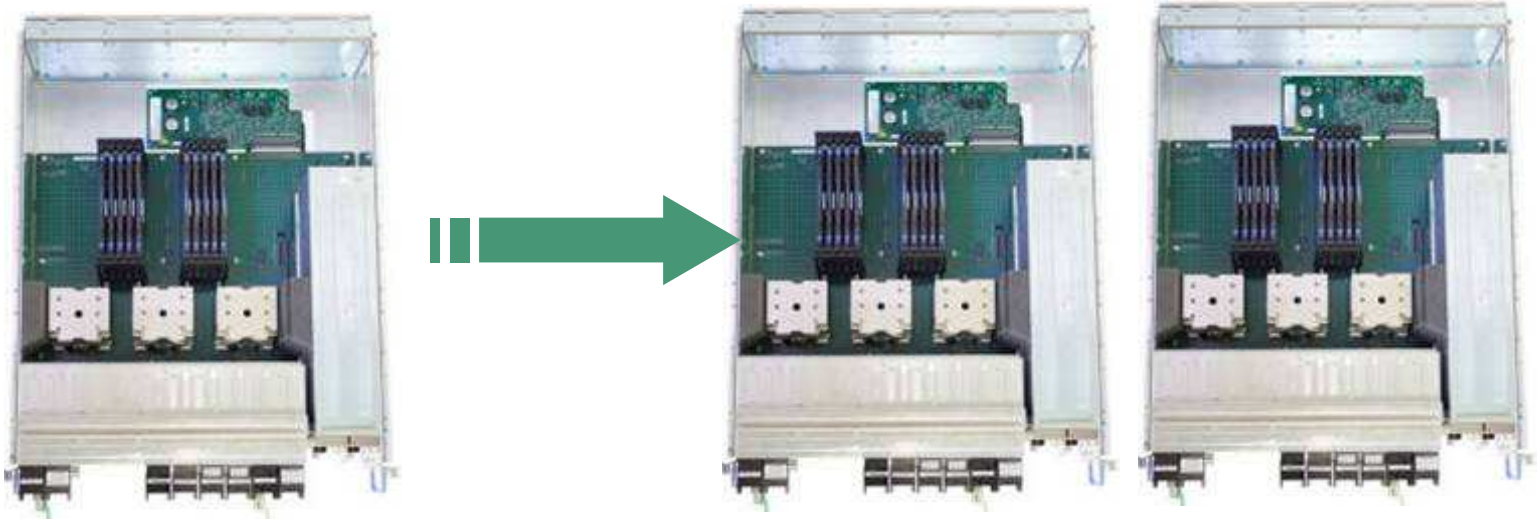

- **System resources split between 2 CPC drawers (Model H13)**
- **Second CPC drawer (Model H13) for:**
  - Increased specialty engine capability up to 13 total Pus for customer use
  - Inreased memory capability (Up to double Model H06 or z114 M10)
  - Increased I/O capability
    - More coupling links than z10 BC (Same number as z114 M10)
    - More I/O features than z10 BC (Same number as z114 M10)

- **Planning Note: Unlike the zEC12 Books, add/remove/repair of a zBC12 CPC drawer is disruptive**

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 CPC drawer
  - 9 per system
    - 2 SAPs standard
    - Up to 6 CPs
    - Up to 6 IFLs/ICFs
    - Up to 4 zIIPs/zAAPs (with Max of 2 CPs)
    - 0 spares when fully configured
    - 1 IFP

- **M/T 2828 – Model H13**
  - Air cooled
  - Single Frame
  - Non-raised floor option available
  - 30 LPARs
- **Processor Units (PUs)**
  - Two CPC drawers
  - 18 per system
    - 2 SAPs standard
    - Up to 6 CPs
    - Up to 13 IFLs/ICFs
    - Up to 8 zIIPs/zAAPs (with Max of 4 CPs)
    - 2 dedicated spares
    - 1 IFP

- **When Model H06 (requires the 2nd processor drawer)?**
  - > 6 Customer PUs
  - > 240 GB memory
  - > 4 Fanouts for additional I/O connectivity – especially PSIFB links
    - Depends - numbers vary for I/O and PCIe I/O drawers, I/O features and PSIFB links

zBC12 Model Structure

Model	Total PUs	CPs	IFLs	zAAPs	zIIPs	ICFs	Std. SAPs	Add'l SAPs	Spares	IFP
H06	9	0-6	0-6	0-4*	0-4*	0-6	2	0-2	0	1
H13	18	0-6	0-13	0-8**	0-8**	0-13	2	0-2	2	1

- Model structure based on number of drawers
  - Model H13 provides specialty engine scale out capabilities
  - Linux only and ICF only servers available
- A minimum of one CP, IFL or ICF must be purchased for each model.
- The zIIP/zAAP to CP ratio is 2:1 i.e for every CP, customer can optionally purchase either two zIIPs or two zAAPs or two zIIPs and two zAAPs in increments of 1
- H06 sparing based on prior Business Class (BC) offerings – no dedicated spares
- H13 sparing based on Enterprise Class (EC) offerings – dedicated spares
  - SAP and PU Allocation/Sparing in the H13
    - Default assignment is one SAP per drawer; one Spare per drawer. Spill and fill CP low to high; spill and fill specialty engines high to low
    - Two defective PUs may cause the default assignment to spill and fill into the second processor drawer. LPAR has the capability to request PU of a specified type to be grouped together in a book/drawer (i.e. LPAR may change the default assignment)
- Disruptive upgrade from H06 to H13 -- No model downgrades
- Upgrades from z10 BC and z114 into either model H06 or H13
- Only the H13 can be upgraded to zEC12 Model H20 (Radiator-based air cooled only)
- IFP – integrated firmware processor. The IFP is standard and not defined by the customer. It is used for infrastructure management of 10GbE RoCE Express, zEDC Express

* With maximum of 2 CPs
** With maximum of 4/5 CPs

System z Cache Topology – z196 vs. zEC12 Comparison

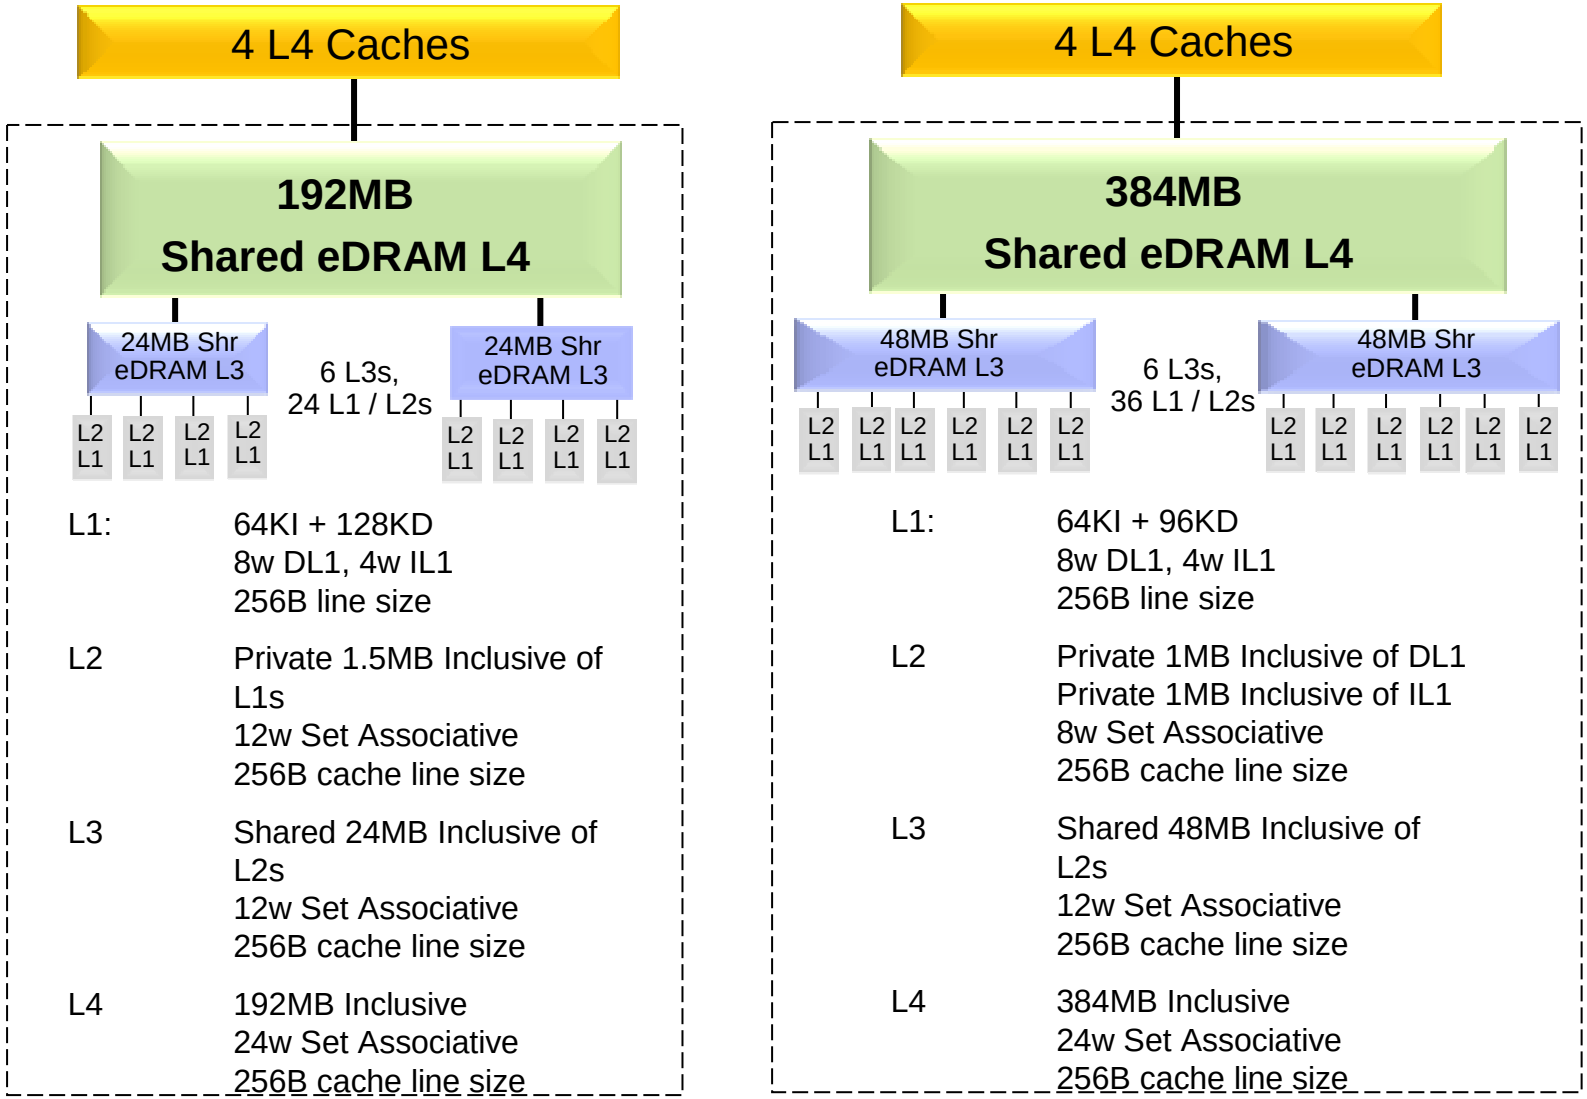

System z Cache Topology – z114 vs. zBC12 Comparison

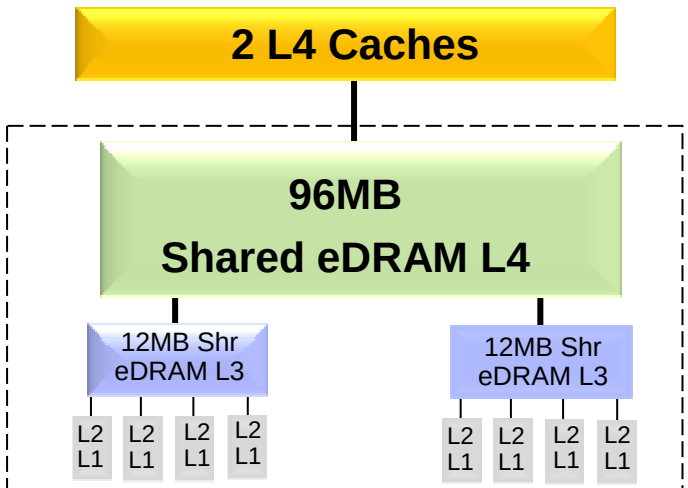

L1:	64KI + 128KD 8w DL1, 4w IL1 256B line size
L2	Private 1.5MB Inclusive of L1s 12w Set Associative 256B cache line size
L3	Shared 12MB Inclusive of L2s 12w Set Associative 256B cache line size
L4	96MB Inclusive 24w Set Associative 256B cache line size

z114

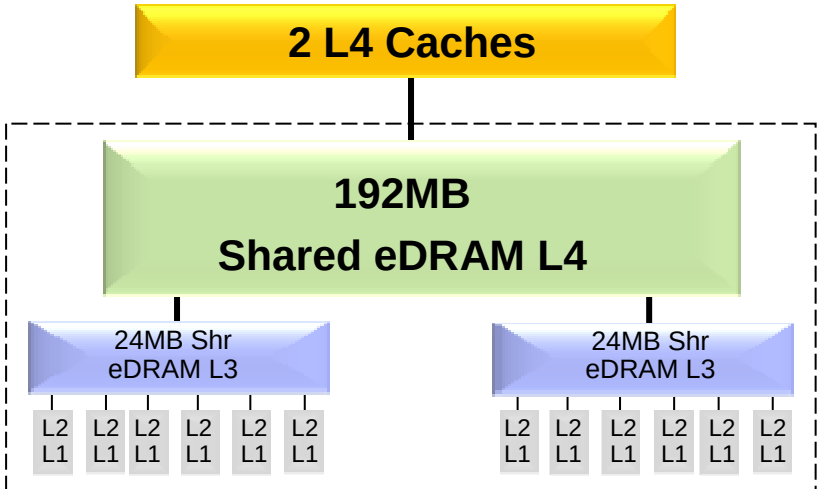

L1:	64KI + 96KD 6w DL1, 4w IL1 256B line size
L2	Private 1MB Inclusive of DL1 Private 1MB Inclusive of IL1 12w Set Associative 256B cache line size
L3	Shared 24MB Inclusive of L2s 12w Set Associative 256B cache line size
L4	192MB Inclusive 24w Set Associative 256B cache line size

zBC12

System z Cache Topology – zBC12 vs. zEC12 Comparison

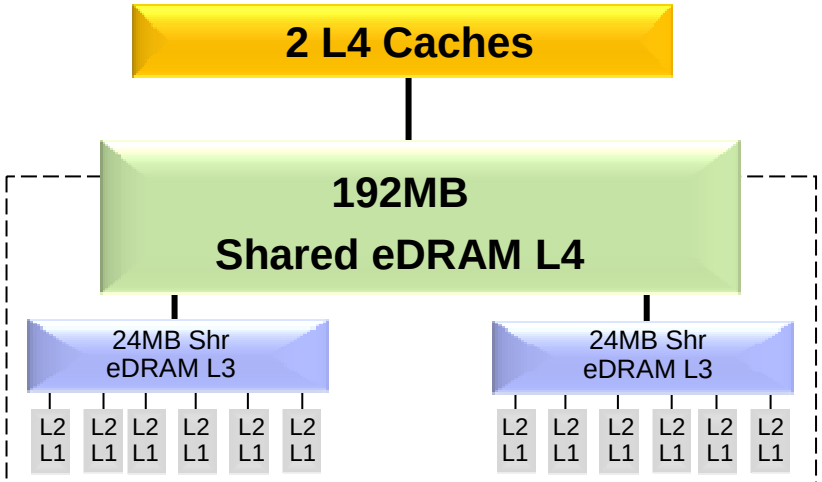

L1: 64KI + 96KD
6w DL1, 4w IL1
256B line size

L2 Private 1MB Inclusive of DL1
Private 1MB Inclusive of IL1
12w Set Associative
256B cache line size

L3 Shared 24MB Inclusive of L2s
12w Set Associative
256B cache line size

L4 192MB Inclusive
24w Set Associative
256B cache line size

zBC12

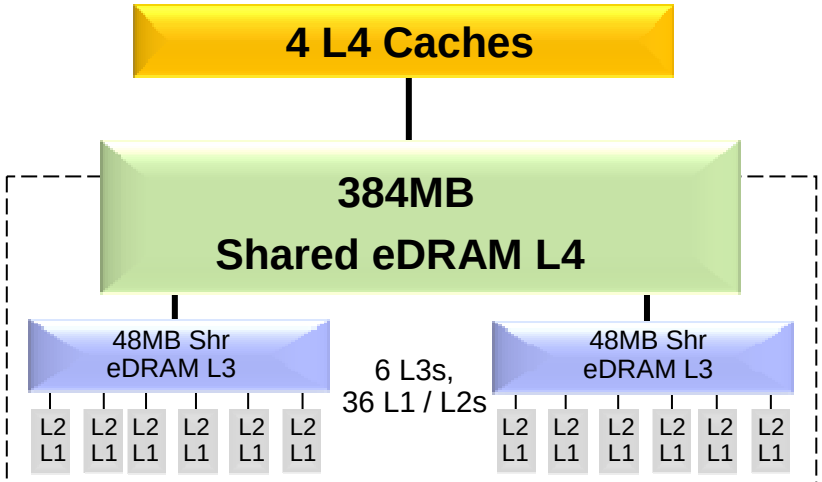

L1: 64KI + 96KD
8w DL1, 4w IL1
256B line size

L2 Private 1MB Inclusive of DL1
Private 1MB Inclusive of IL1
8w Set Associative
256B cache line size

L3 Shared 48MB Inclusive of L2s
12w Set Associative
256B cache line size

L4 384MB Inclusive
24w Set Associative
256B cache line size

zEC12

zEC12 and zBC12 Comparison

	zEC12	zBC12
Nodal Topology	Fully Interconnected	Bi-Nodal
Number of books / drawers	4 books	2 drawers
Max # of PU's for System	120	18
Number of chips per Book/Drawer	20 DIMMs 6 PUs, 2 SCs	2 PUs, 1 SC
1st level cache (PU core)	64KB I, 96KB D	64KB I, 96KB D
2nd level cache (PU core)	2MB private per core (L2)	2MB private per core (L2)
3rd level cache (PU nest)	48MB Shared per 6 cores (L3)	24MB Shared per 6 cores (L3)
4th level cache (SC)	384MB Shared per node (SC L4)	192MB Shared per node (SC L4)
# of PU L3 pipes	2 Main, 2 Store	1 Main, 1 Store
# of SC cache pipes	1 Main (per SC)	1 Main
Memory DIMMs	SuperNova x81 DIMMs	SuperNova x81 DIMMs
Max System Memory Size including HSA	3 TBs on 120 DIMMs	512 GBs on

zBC12 and z114 BC Comparison

	zBC12	z114
Nodal Topology	Bi-Nodal	Bi-Nodal
Number of drawers	1-2 drawers	1-2 drawers
Max # of PU's	9 PU Cores / Drawer	7 PU Cores / Drawer
Number of chips per Drawer	2 PUs, 1 SCs	2 PUs, 1 SC
1st level cache (PU core)	64KB I, 96KB D	64KB I, 128KB D
2nd level cache (PU core)	2MB private per core (L2)	1.5MB private per core (L2)
3rd level cache (PU nest)	24MB Shared per 6 cores (L3)	12MB Shared per 4 cores (L3)
4th level cache (SC)	192MB Shared per node (SC L4)	96MB Shared per node (SC L4)
# of CP L3 pipes	1 Main, 1 Store	1 Main, 1 Store
# of SC cache pipes	1 Main	1 Main
Memory DIMMs	SuperNova x81 DIMMs 10 DIMMs/Drawer	SuperNova x81 DIMMs 10 DIMMs/Drawer
Max System Memory Size (excluding HSA)	240 GB: 1 CPC Drawer 496 GB : 2 CPC Drawer	120 GB: 1 CPC Drawer 248 GB : 2 CPC Drawer

System z I/O Subsystem Internal Bus Interconnect Speeds (GBps)


PCIe
zEC12/zBC12
z196/z114

8 GBps


InfiniBand
z10/z196/z114
zEC12 and
zBC12 for carry forward

6 GBps


STI
z9

2.7 GBps


STI
z990/z890

2 GBps


STI
z900/z800

1 GBps

STI: Self-Timed Interconnect

Out of Order Detail

- Out of order yields significant performance benefit through
  - Re-ordering instruction execution
    - Instructions stall in a pipeline because they are waiting for results from a previous instruction or the execution resource they require is busy
    - In an in-order core, this stalled instruction stalls all later instructions in the code stream
    - In an out-of-order core, later instructions are allowed to execute ahead of the stalled instruction
  - Re-ordering storage accesses
    - Instructions which access storage can stall because they are waiting on results needed to compute storage address
    - In an in-order core, later instructions are stalled
    - In an out-of-order core, later storage-accessing instructions which can compute their storage address are allowed to execute
  - Hiding storage access latency
    - Many instructions access data from storage
    - Storage accesses can miss the L1 and require 10 to 500 additional cycles to retrieve the storage data
    - In an in-order core, later instructions in the code stream are stalled
    - In an out-of-order core, later instructions which are not dependent on this storage data are allowed to execute

Out of Order Execution – z196/z114 vs. zEC12/zBC12

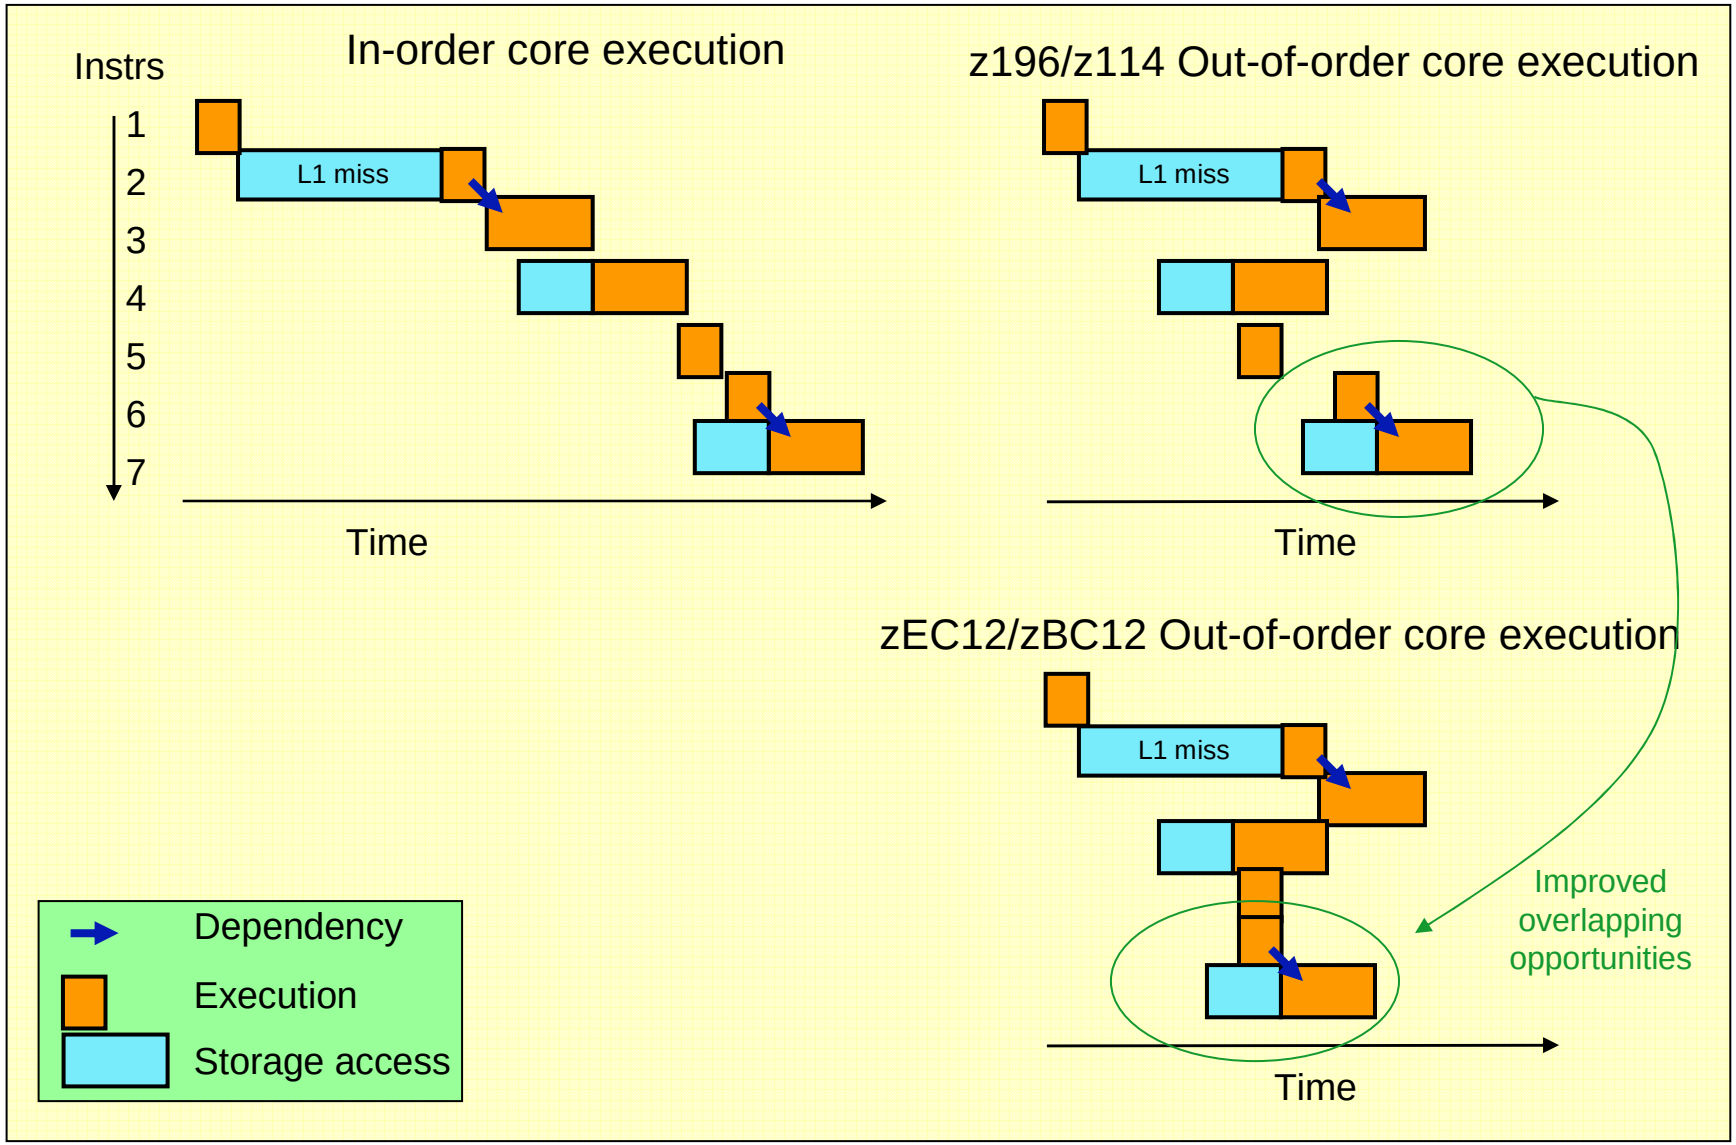

Out of Order - Improved instruction delivery and execution

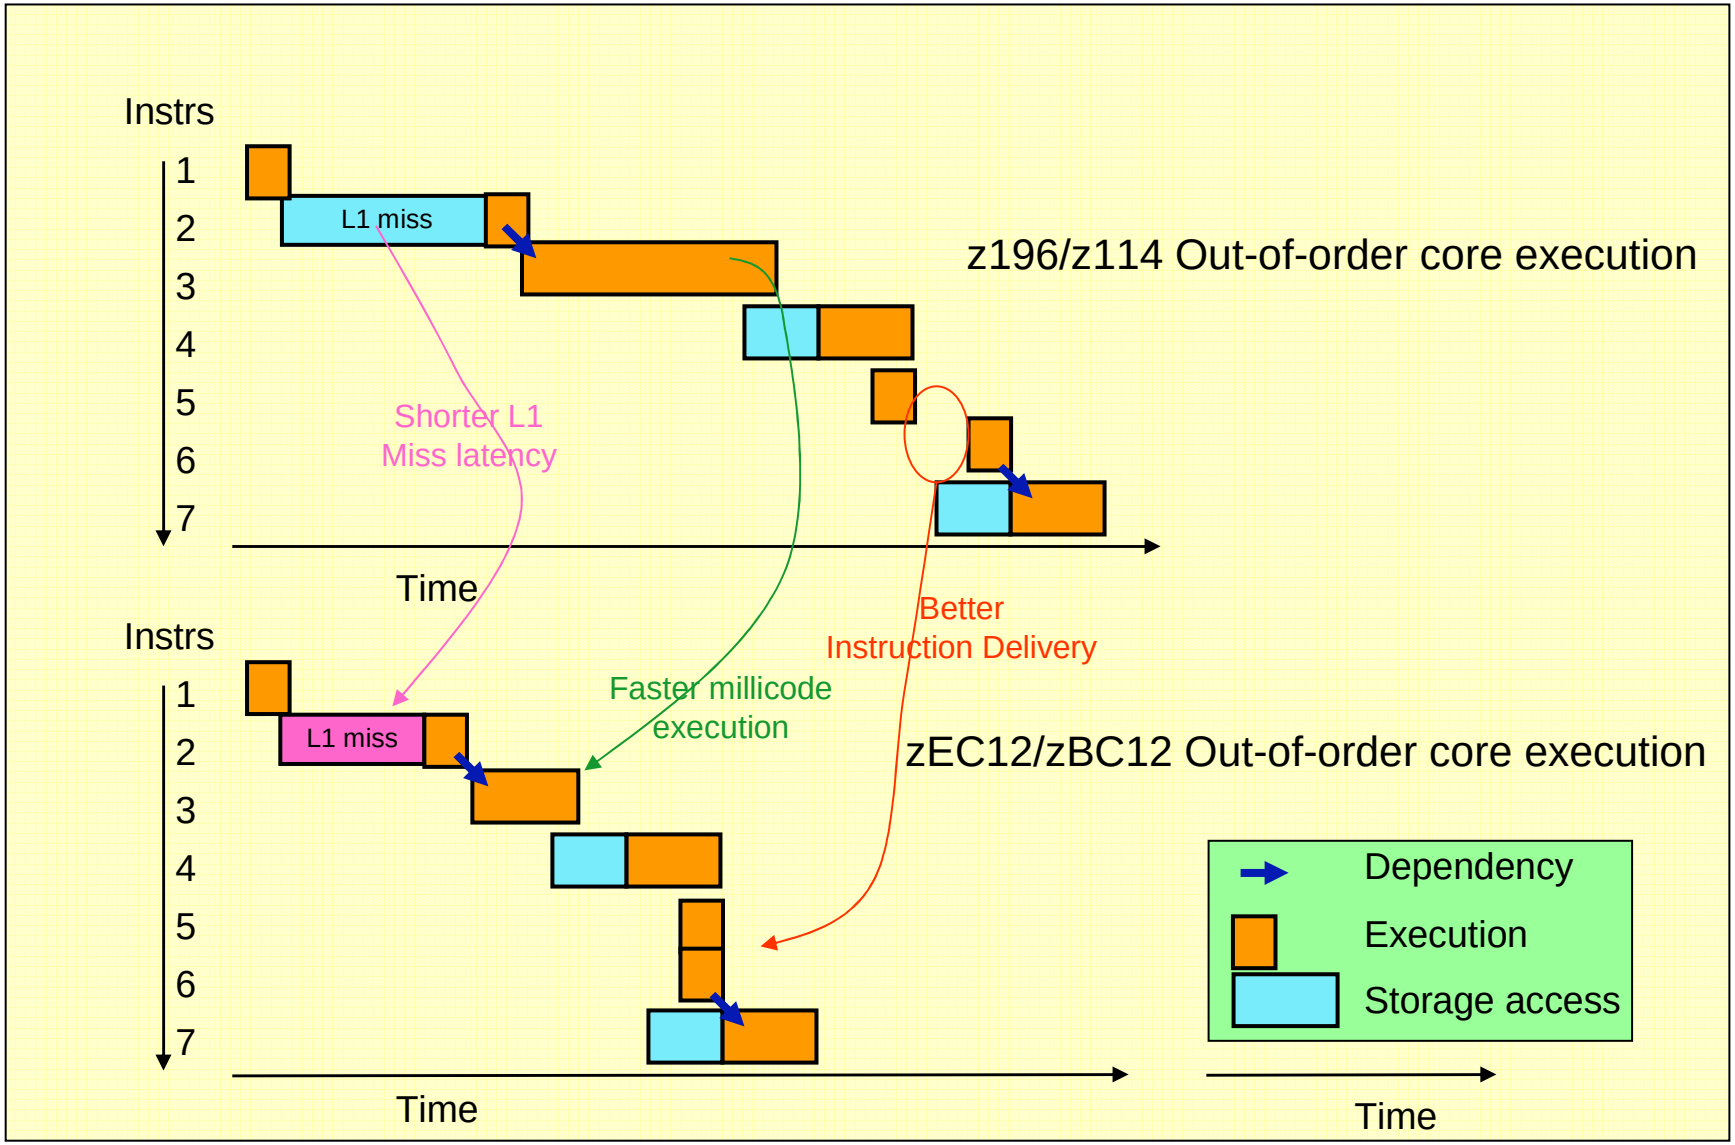

Transactional Execution - (TX) Facility

Transaction:

- Atomic memory accesses for set of instructions
- “All or nothing”
- Collision detection and rollback instead of collision avoidance

Increases **scalability** and **parallelism** for higher transaction throughput

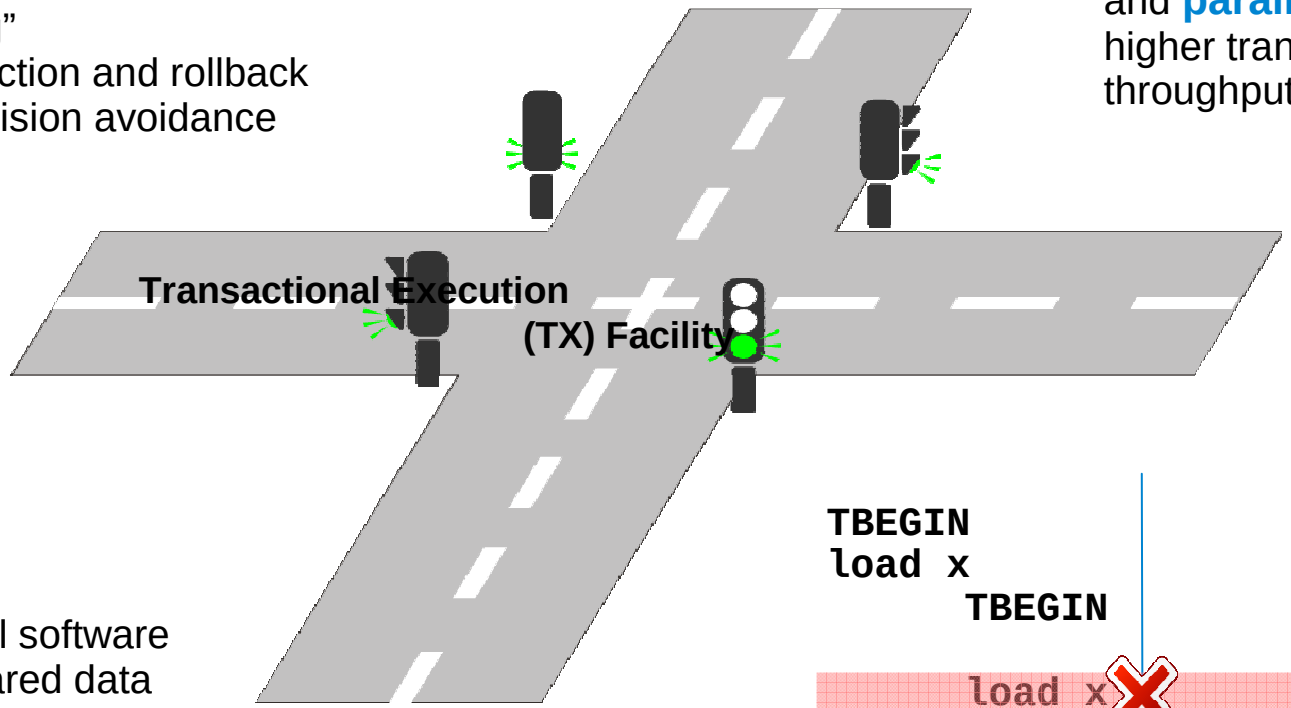

Applications:

- Java
- Highly parallel software relying on shared data structures

```
TBEGIN
load x
      TBEGIN
      load x
x:= x+1
      x:= x+1
store x
      store x
      TEND
TEND
```

What Is Flash Express?

- Also referred to as Storage Class Memory (SCM)
- Flash Express is internal storage implemented via NAND Flash SSDs (Solid State Drives) mounted in PCIe Flash Express feature
  - Plugs into PCIe I/O drawers in pairs
  - Data security provided on the feature
  - A pair provides 1.4 TB of useable storage (1 TB = 2^{30} bytes)
  - A maximum of 4 pairs (4 x 1.4 = 5.6 TB) are supported in a system
- Internal Flash Express 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 Express in zEC12/zBC12 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
- Another application is DUMP
- zEC12 GA2 and zBC12 New Function
  - CF exploitation of Flash Express memory (SoD)*

Relative Access Times for different technologies

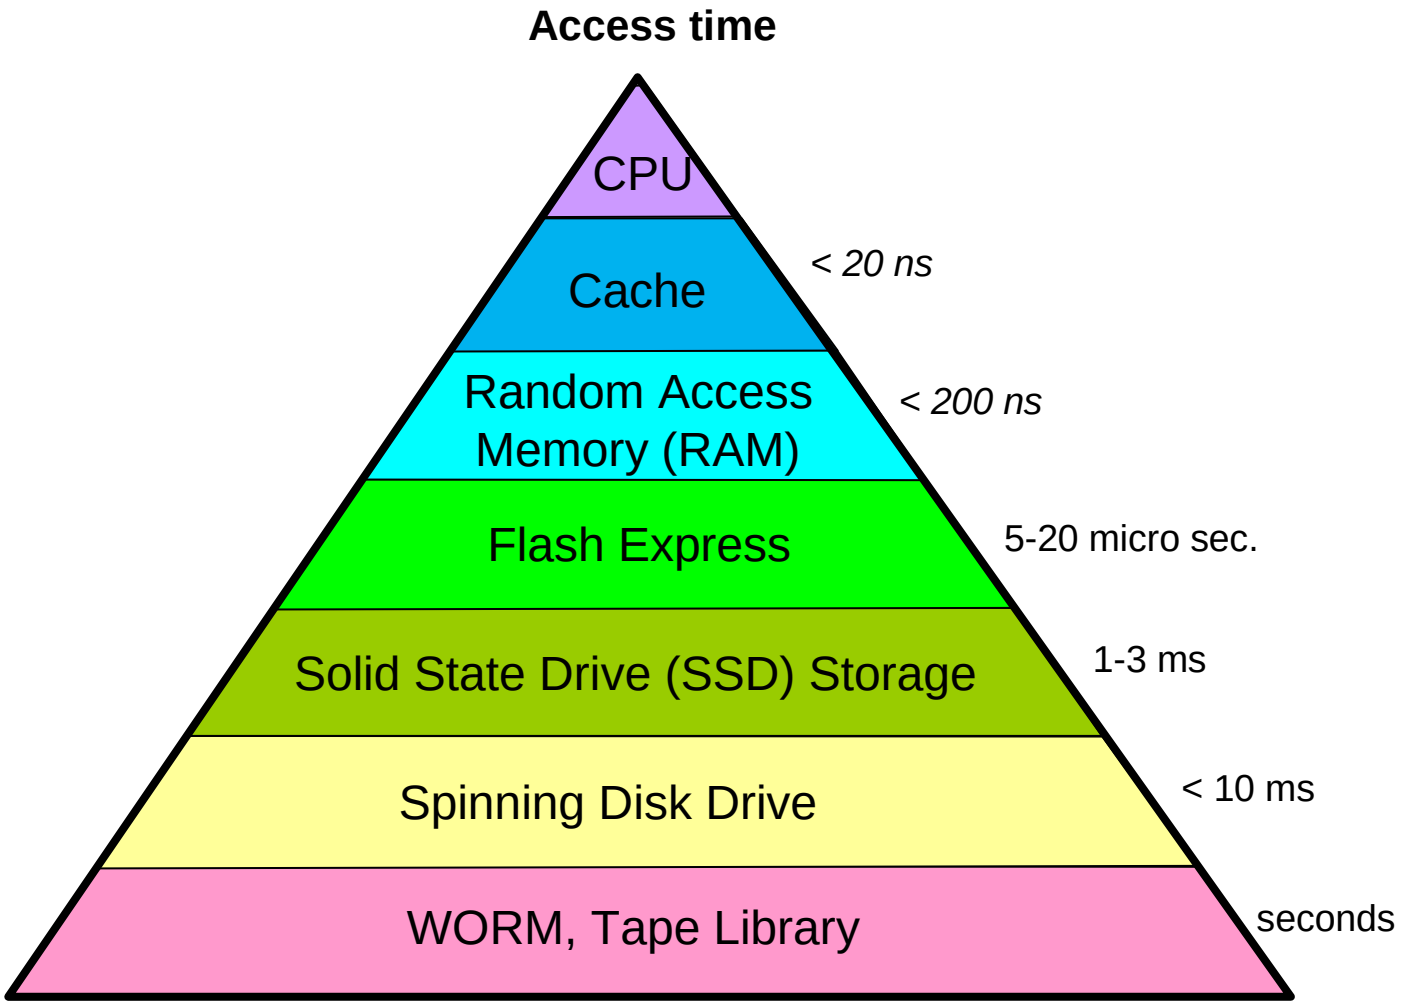

Flash Express Exploitation on zEC12/zBC12

- Flash Express exploitation:
  - z/OS
    - z/OS V2.1
    - z/OS V1.13 with PTFs and the z/OS V1.13 RSM Enablement Offering Web deliverable installed
      - Pagable Large Pages (1 MB)
    - z/OS V1.13 enabling PTFs for RSM enhancements
      - Flash Dynamic Reconfiguration
      - Optional PLPA and COMMON Page data sets z/OS V1.13 Language Environment
  - IMS 12 Common Queue Server exploitation with APAR# PM66866
  - IMS 13 Common Queue Server storage; IMS 13 DB private storage pools*
  - DB2 V10* planned
  - Java SDK7 SR3 and by extension exploiters such as
    - CICS Transaction Server 5.1
    - WAS Liberty Profile v8.5
    - IMS 12 available October 2013
    - DB2 V11
    - Traditional WAS 8.0.0x and Traditional WAS 8.5.5 (future) **
  - **Note:** Even device driver is inlcuded in Linux for System z distributions (RHEL 6.4 and SUSE SLESS 11 SP3), based on the support available today, IBM recommends for customers running virtualized and non-virtualized Linux environments on System z to use IBM FlashSystem. (z/VM does not use Flash Express and does not provide guest access to Flash Express.)

Flash Express

**Traditional WAS support is planned for a future date
*DB2 10 support for V10 with APARs is planned. IMS 13 support is planned

zBC12 I/O Features supported

Supported features

PCIe I/O drawer

32-slot PCIe I/O drawer

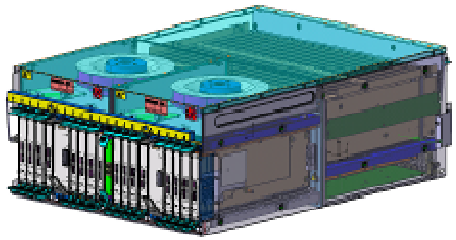

- The drawer itself does NOT carry forward but the features do

- **Features that Carry Forward**
  - OSA-Express4S GbE LX and SX, OSA-Express4S 10 GbE LR and SR
  - FICON Express8S 10Km LX and FICON Express8S SX
- **Features - New Build**
  - FICON Express8S 10Km LX and FICON Express8S SX
  - OSA-Express5S GbE LX and SX, OSA-Express5S 10 GbE LR and SR, OSA-Express5S 1000BASE-T
  - Crypto Express4S
  - Flash express
  - 10GbE RoCE Express
  - zEDC Express

I/O drawer

8-slot I/O drawer

- The drawer itself can carry forward. It cannot be ordered on new build or Migration Offering
- One I/O drawer can be carried forward, two I/O drawers requires an RPQ 8P2733

- **Features that Carry Forward**
  - OSA-Express3 GbE LX and SX, OSA-Express3 10 GbE LR and SR, OSA-Express3 1000BASE-T
  - OSA-Express3-2P 1000BASE-T, OSA-Express3-2P GbE SX
  - FICON Express4-2C SX, FICON Express4 10KM LX and FICON Express4 SX
  - FICON Express8 10KM LX, FICON Express8 SX
  - Crypto Express3, Crypto Express3-1P
  - ISC-3
  - **Not Supported: ESCON, older FICON, FICON Express4 4 KM LX (2 and 4 ports) , OSA-Express2, PSC**

FICON performance on System z

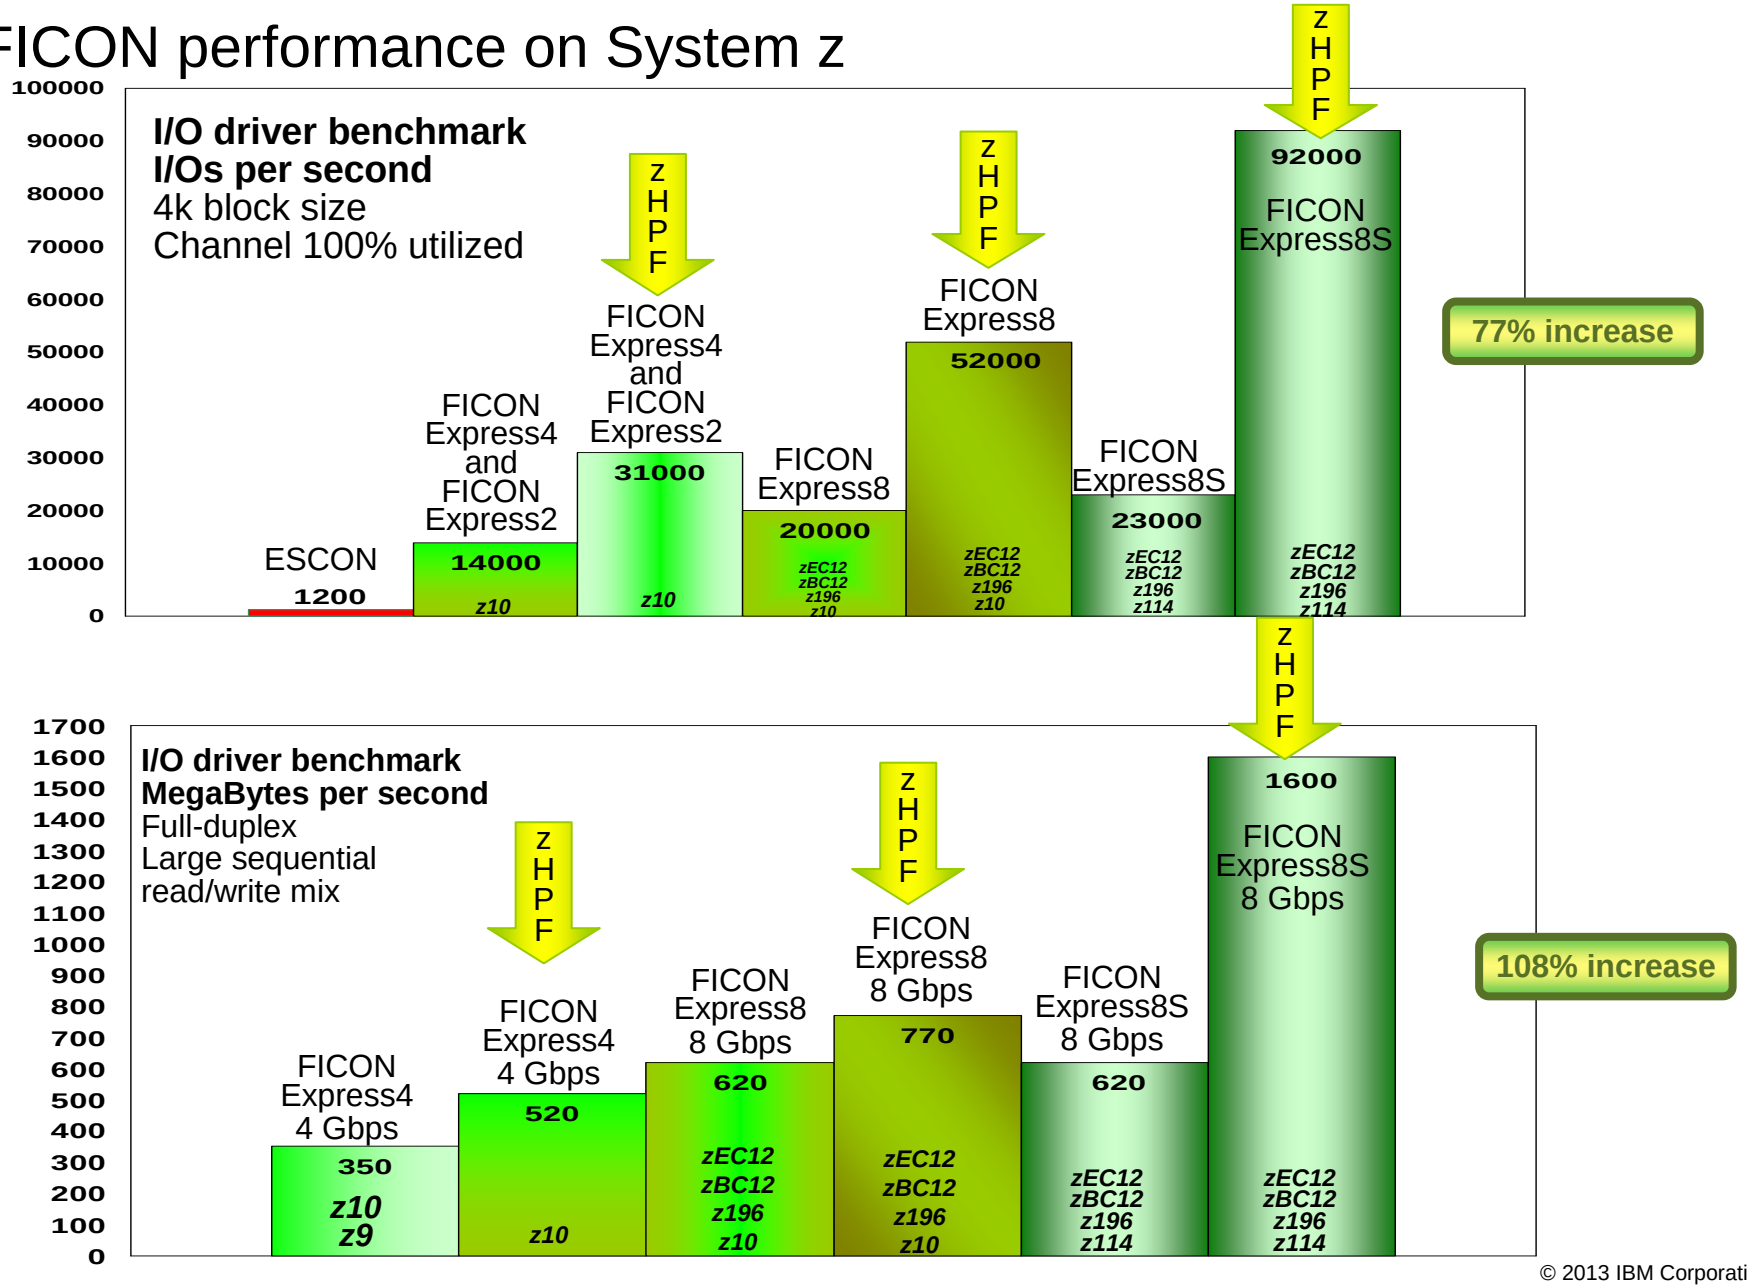

FCP performance on System z

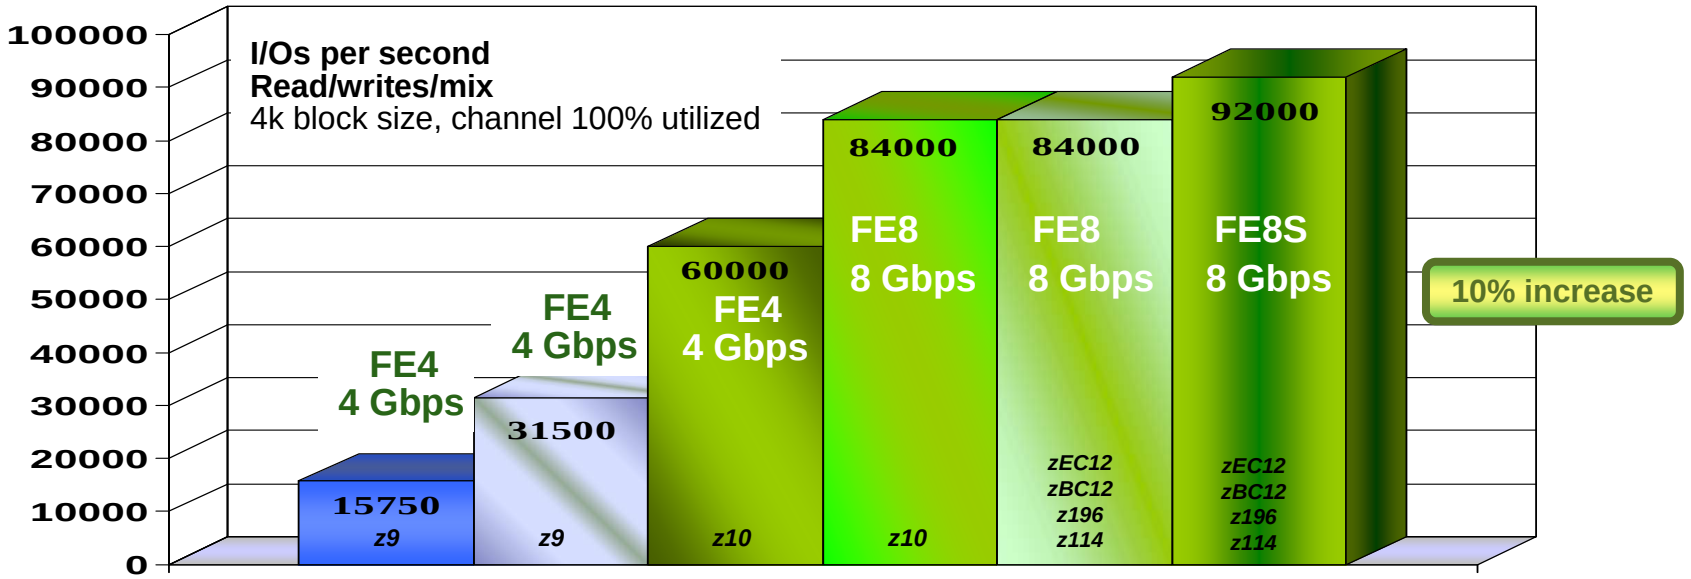

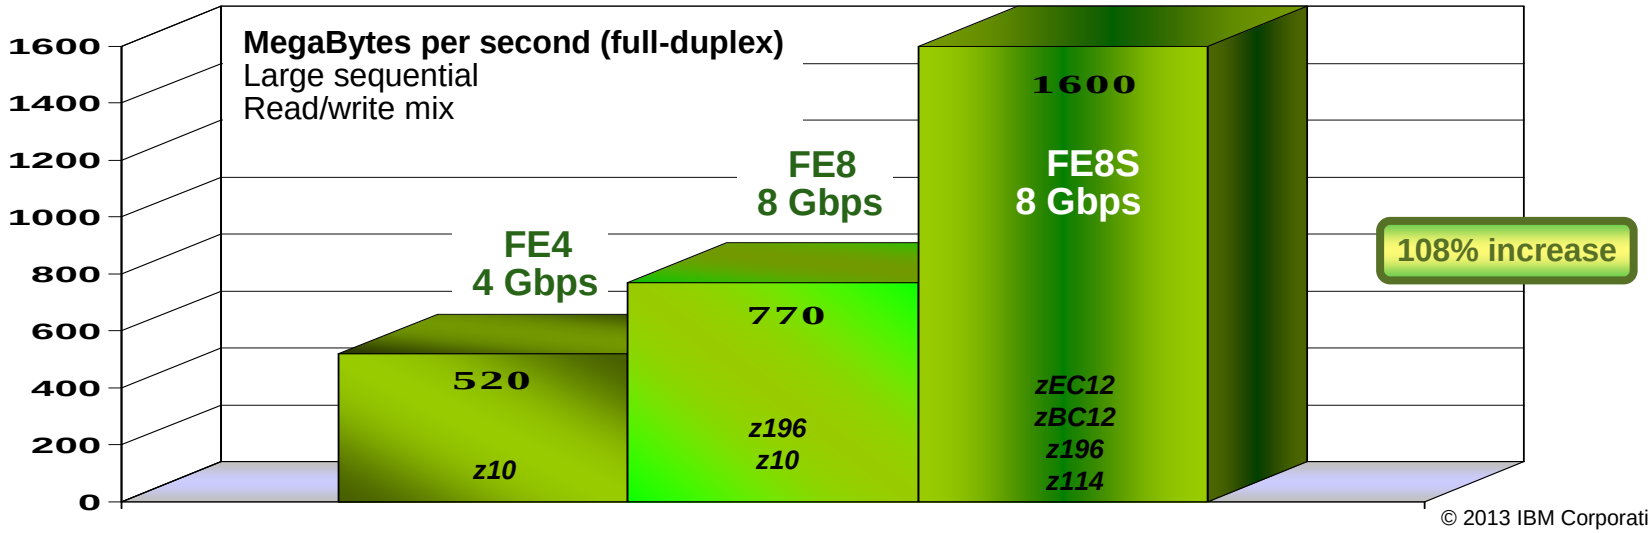

OSA-Express4S and 5S 10 GbE performance (laboratory)

Inbound Streams – 1492 Byte MTUs

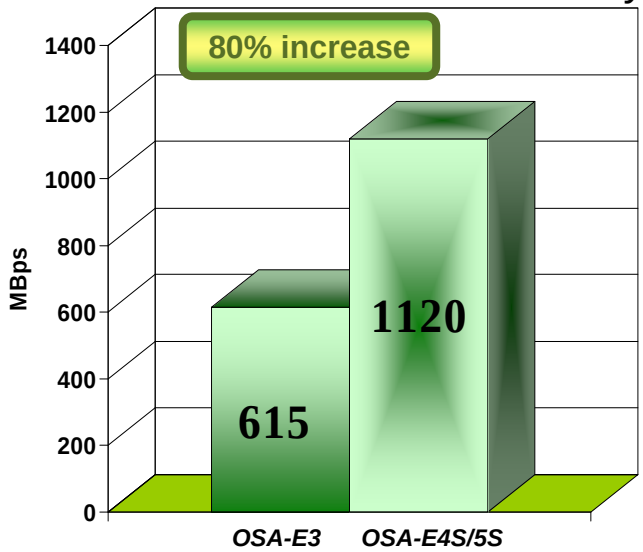

Mixed Streams – 1492 Byte MTUs

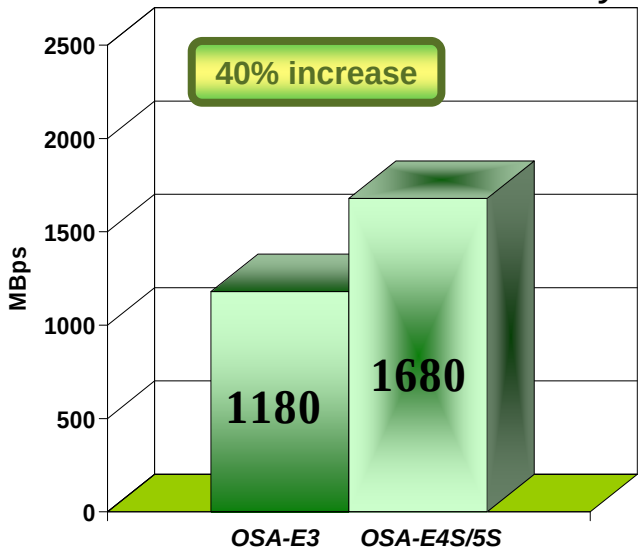

Inbound Streams – 8000 Byte MTUs

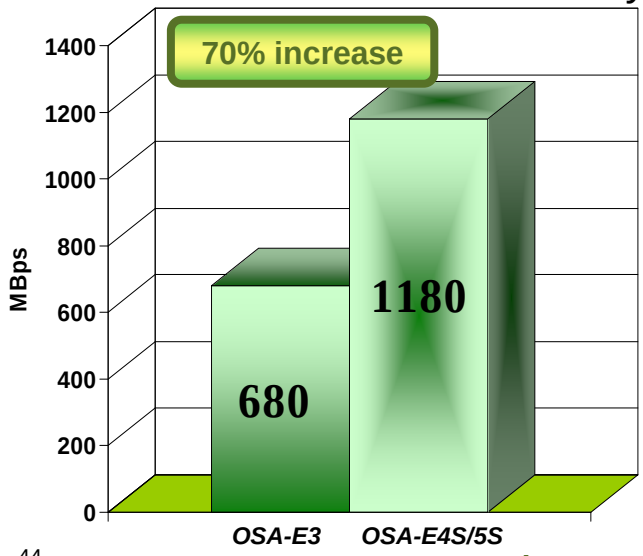

Mixed Streams – 8000 Byte MTUs

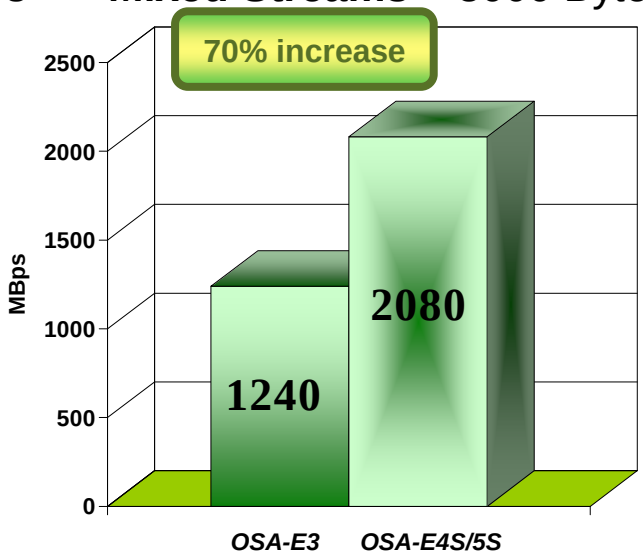

Notes:
- AWM on z/OS
- z/OS is doing checksum
- 1 megabyte per second (MBps) is 1,048,576 bytes per second
- MBps represents payload throughput (does not count packet and frame headers)

Maximum transmission units (MTUs)

System z Crypto History

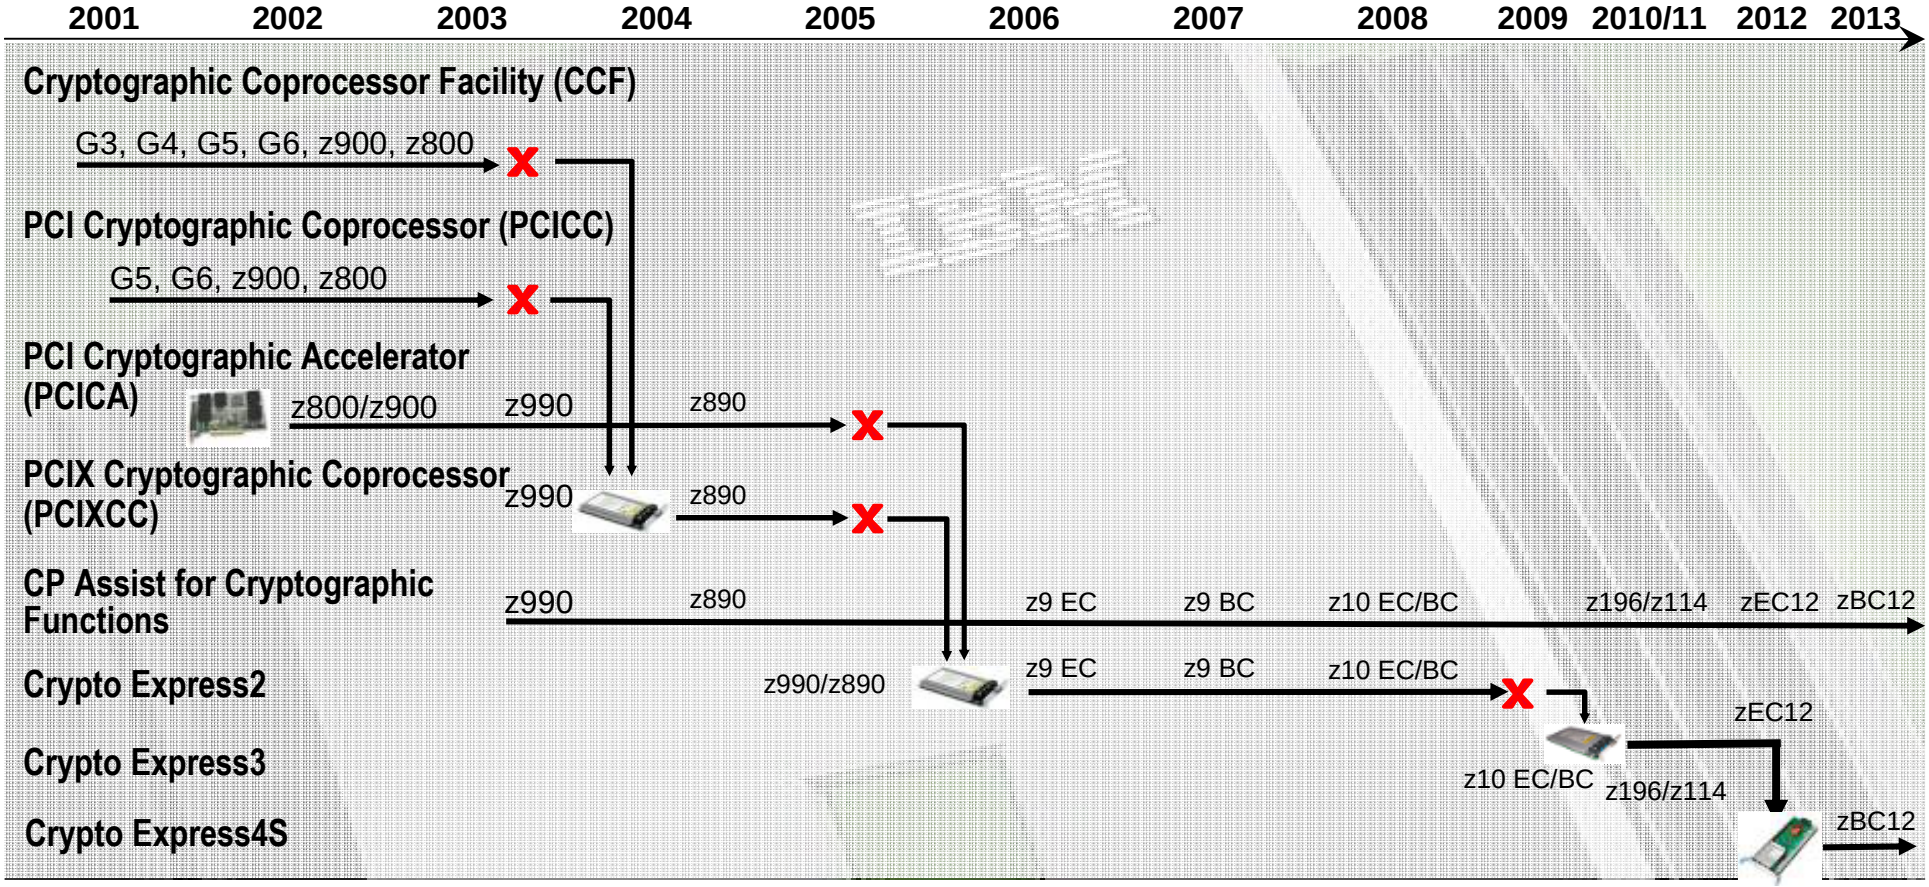

- Cryptographic Coprocessor Facility – Supports “Secure key” cryptographic processing
- PCICC Feature – Supports “Secure key” cryptographic processing
- PCICA Feature – Supports “Clear key” SSL acceleration
- PCIXCC Feature – Supports “Secure key” cryptographic processing
- CP Assist for Cryptographic Function allows limited “Clear key” crypto functions from any CP/IFL
  - NOT equivalent to CCF on older machines in function or Crypto Express2 capability
- Crypto Express2 – Combines function and performance of PCICA and PCICC
- Crypto Express3 – PCIe Interface, additional processing capacity with improved RAS
- Crypto Express4S - IBM Standard PKCS #EP11

CPACF - The CP Assist For Cryptographic Functions

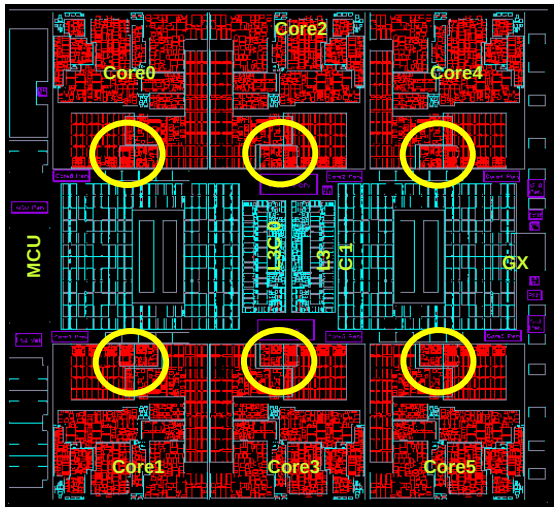

Supported Algorithms	Clear Key	Protected Key
DES, T-DES	Y	Y
AES128	Y	Y
AES192	Y	Y
AES256	Y	Y
SHA-1	Y	N/A
SHA-256	Y	N/A
SHA-384	Y	N/A
SHA-512	Y	N/A
PRNG	Y	N/A

- Provides a set of symmetric cryptographic functions and hashing functions for:
  - Data privacy and confidentiality
  - Data integrity
  - Random Number generation
  - Message Authentication
- Enhances the encryption/decryption performance of clear-key operations for
  - SSL
  - VPN
  - Data storing applications
- Available on every Processor Unit defined as a CP, IFL, zAAP and zIIP
- Supported by z/OS, z/VM, z/VSE, z/TPF and Linux on System z
- Must be explicitly enabled, using a no-charge enablement feature (#3863),
  - SHA algorithms enabled with each server
- Protected key support for additional security of cryptographic keys
  - Crypto Express3 or Crypto Express4S required in CCA mode

Crypto Express4S (July 2012)

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

- **Three configuration options for the PCIe adapter**
  - Only one configuration option can be chosen at any given time
  - Switching between configuration modes will erase all card secrets
    - Exception: Switching from CCA to accelerator or vice versa
- **Accelerator (No change compared to z196)**
  - For SSL acceleration
  - Clear key RSA operations
- **Enhanced: Secure IBM CCA coprocessor (default)**
  - **Optional**: TKE workstation (#0841) with TKE LIC 7.2 (#850) for security-rich, flexible key entry or remote key management
- **New: IBM Enterprise PKCS #11 (EP11) coprocessor**
  - Designed for extended evaluations to meet public sector requirements
    - Both FIPS and Common Criteria certifications
  - **Required*:*** TKE workstation (#0841) with TKE LIC 7.2 (#850) for management of the Crypto Express4S when defined as an EP11 coprocessor

What is an Integrated firmware processor (IFP)?

- Integrated firmware processor (IFP)
  - The IFP (previously known as Reserved PU) is allocated from the pool of PUs available for the whole System
    - Unlike other characterized PUs, the customer doesn't pay for the IFP
    - It's a single PU dedicated solely for the purpose of supporting the native PCIe features (10GbE RoCE Express and zEDC Express) and is initialized at POR if these features are present
    - It has two Resource Groups (RGs) which have firmware for the 10GbE RoCE Express and zEDC Express features
  - The IFP is not a customer useable PU and not 'visible' to the customer
    - Usage is very small and has negligible impact on the MP effect on the System
    - Acts exactly like any other PU during error or failover scenarios i.e. sparing

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.

IBM zEnterprise Data Compression (zEDC)

New data compression offering that can reduce resource usage

What is it?

✓ *A combined software (z/OS V2.1) and hardware (zEDC Express) solution designed to help reduce resource consumption, disk utilization and optimize cross platform exchange of data*

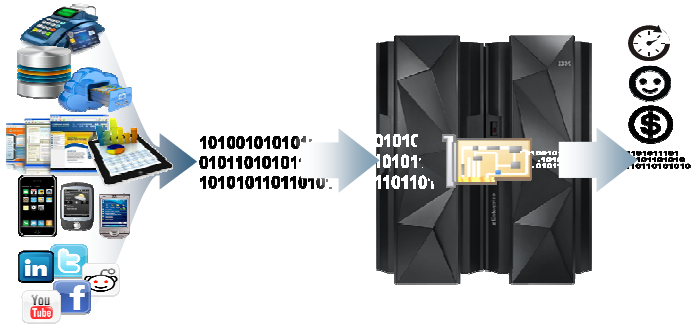

How is it different

- **Performance:** Efficient alternative for larger files. Reduced CPU overhead for SMF diagnostic jobs.
- **Efficient:** Optimized algorithms scan text to locate the re-use of phrases and refers back to earlier references
- **Industry Standard:** Compatible with open zlib based compression – widely used across all platforms
- **Economical:** Reduced DASD space requirements and improved effective bandwidth without significant CPU overhead***

15% reduction in elapsed time for SMF extraction with up to **40%** reduction for CPU time*

Logger overhead reduced by up to **30%** **

* When running an SMF extraction/dump against an SMF logstream with records compressed by zEDC
** The amount of data sent to an SMF logstream can be reduced by up to 75% using zEDC compression – reducing logger overhead
*** SOD for BSAM/QSAM access methods

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

zEDC Feature Details

- Implemented as a native PCIe adapter
- Provides cross platform compliant compression and decompression defined by RFC1951 (DEFLATE)
- Sharable across up to 15 partitions
- Up to 8 devices per CPC
- Support for z/OS

z/OS V2.1 and z/VM V6.3 Exploitation of zEDC Express

- Introductory Use with SMF Log Data: Clients running SMF using logger that are looking to reduce the logger overhead or collect additional data
- Clients such as a clearing house, financial institution or direct marketing agencies that are sending and receiving large files
- SoD: Customers with large volumes of extended format BSAM/QSAM sequential data*
- SoD: Clients that have purchased flash on DS8870 and want to use it more efficiently when storing extended format BSAM/QSAM sequential data*
- SoD: Clients that use Java today where they create a stream of compressed data*
- SoD for z/VM V6.3 guest exploitation support*

*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.

RoCE - RDMA over Converged Ethernet

- RDMA based technology has been available in the industry for many years – primarily based on Infiniband (IB)
  - RDMA technology provides the capability to allow hosts to logically share memory
  - IBM requires a completely unique network eco system (unique hardware such as host adapters, switches, host application software, system management software/firmware, security controls, etc.) – IB is common in the HPC market
- RDMA technology is now available on Ethernet – RDMA over Converged Ethernet (RoCE)
  - RoCE uses existing Ethernet fabric (switches with Global Pause enabled)
- Host software exploitation options fall into two general categories:
  - Native / direct application exploitation (many variations)
  - Transparent application exploitation (e.g. sockets based)
- The opportunity created by RoCE also creates a need for a common RoCE software / application communications strategy that preserves:
  - The performance benefits of RDMA and
  - Does not compromise the existing IP based qualities of service
- ***Focus is on providing a transparent Sockets over RDMA solution***

Shared Memory Communications – Remote Direct Memory Access (SMC-R) Definition

- Shared Memory Communications – Remote Direct Memory Access (SMC-R) is a new communication protocol aimed at providing transparent acceleration for sockets-based TCP/IP applications and middleware
  - Remote Direct Memory Access (RDMA) technology provides low latency, high bandwidth, high throughput, low processor utilization attachment between hosts
  - SMC-R utilizes RDMA over Converged Ethernet (RoCE) as the physical transport layer
- SMC-R is built on the following concepts:
  - RDMA enablement of the communications fabric
  - Partitioning a part of OS host real memory into buffers and using RDMA technology to access this memory
  - Establishing an ‘out of band’ connection over which data is passed to the partner peer using RMDA writes and signaling

RDMA (Remote Direct Memory Access) Technology Overview

- Key attributes of RDMA
  - Enables a host to read or write directly from/to a remote host's memory ***without*** involving the remote host's CPU
    - By registering specific memory for RDMA partner use
    - Interrupts ***still required*** for notification (i.e. CPU cycles are not completely eliminated)
  - Reduced networking stack overhead by using streamlined, low level, RMDA interfaces
    - No requirement for TCP/IP protocols/stack, sockets, etc.
      - Low level APIs such as uDAPL, MPI or RDMA verbs allow optimized exploitation
        - > *For applications/middleware willing to exploit these interfaces*
  - Key requirements:
    - A reliable "lossless" network fabric (LAN for layer 2 data center network distance)
    - An RDMA capable NIC (RNIC) and ethernet fabric – switches (recommended) or point to point

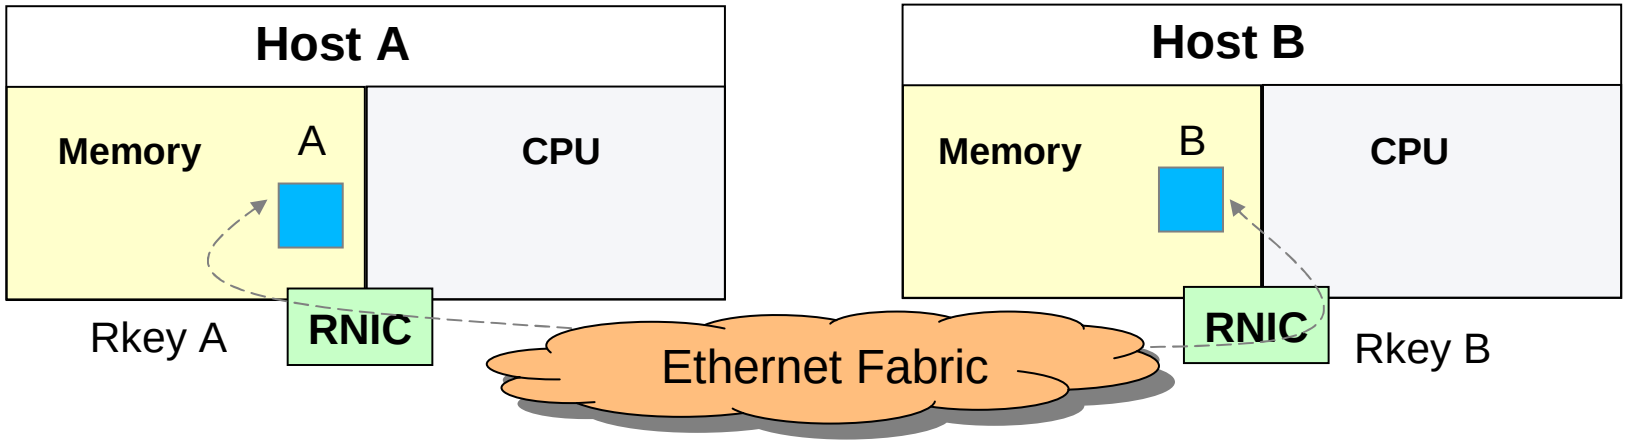

Dynamic Transition from TCP to SMC-R

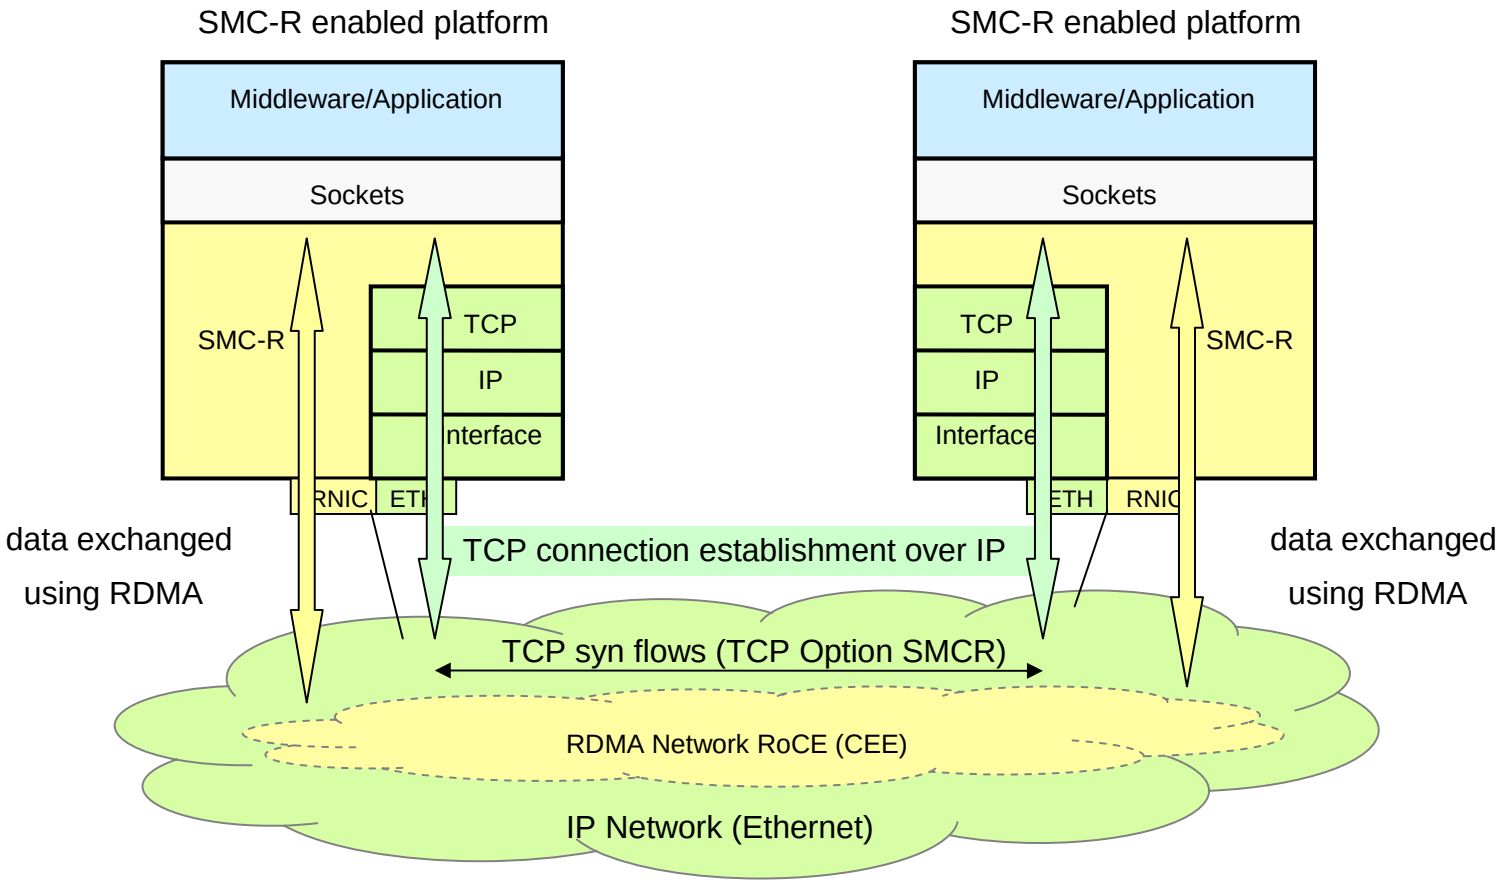

Dynamic (in-line) negotiation for SMC-R is initiated by presence of TCP Option (SMCR)

TCP connection transitions to SMC-R allowing application data to be exchanged using RDMA

10GbE RoCE Express Feature

- **Configuration**
  - One 10 GbE SFP+ port enabled per feature
  - Each feature must be dedicated to one LPAR (Reconfiguration is supported)
  - Recommended minimum configuration two features per LPAR for redundancy
  - A switched connection requires an enterprise-class 10 GbE switch
    - SR Optics, *Global Pause* enabled (as defined by the IEEE 802.3x standard) & *Priority Flow Control (PFC)* disabled
  - Point-to-point connection is supported Not defined as a CHPID and does not consume a CHPID number
  - Up to 16 features supported by zBC12/zEC12
  - Link distance up to 300 meters over OM3 50 micron multimode fiber
- **Exploitation and Compatibility**
  - Exclusive to zEC12 GA2 and zBC12
  - z/OS V2.1 Exploitation with PTFs

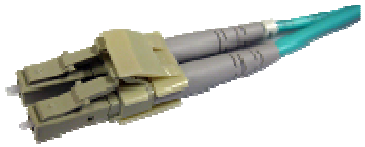

FC 0411 10GbE RoCE Express - OM3 fiber recommended

Shared Memory Communication – RDMA utilizing the 10GbE RoCE Express

IBM zAware - Identifies Unusual System Behavior

IBM zAware contains sophisticated analytics, applies IBM insight, and machine learning to understand your unique system.

Monitoring	Detection	Frequency	Reporting
• Supports IBM and non IBM middleware and applications • Monitors OPERLOG in a sysplex or monoplex • Assigns a message anomaly score to help identify potential issues • Maximum distance client to host IBM zAware is 3500 km	• Detects anomalies other solutions might miss • Can find the rare or infrequent message • Can detect an unusual number of normal messages • Can detect messages issued out of context	• Samples every 2 minutes • 10 minute interval • Uses 90 day rolling baseline; a utility provided to populate baseline; flexibility provided	• Near real time analysis • Intuitive reporting – both high level and drill down • Color coded browser display • XML output can feed ISVs or processes – Tivoli® provides alert and event notifications

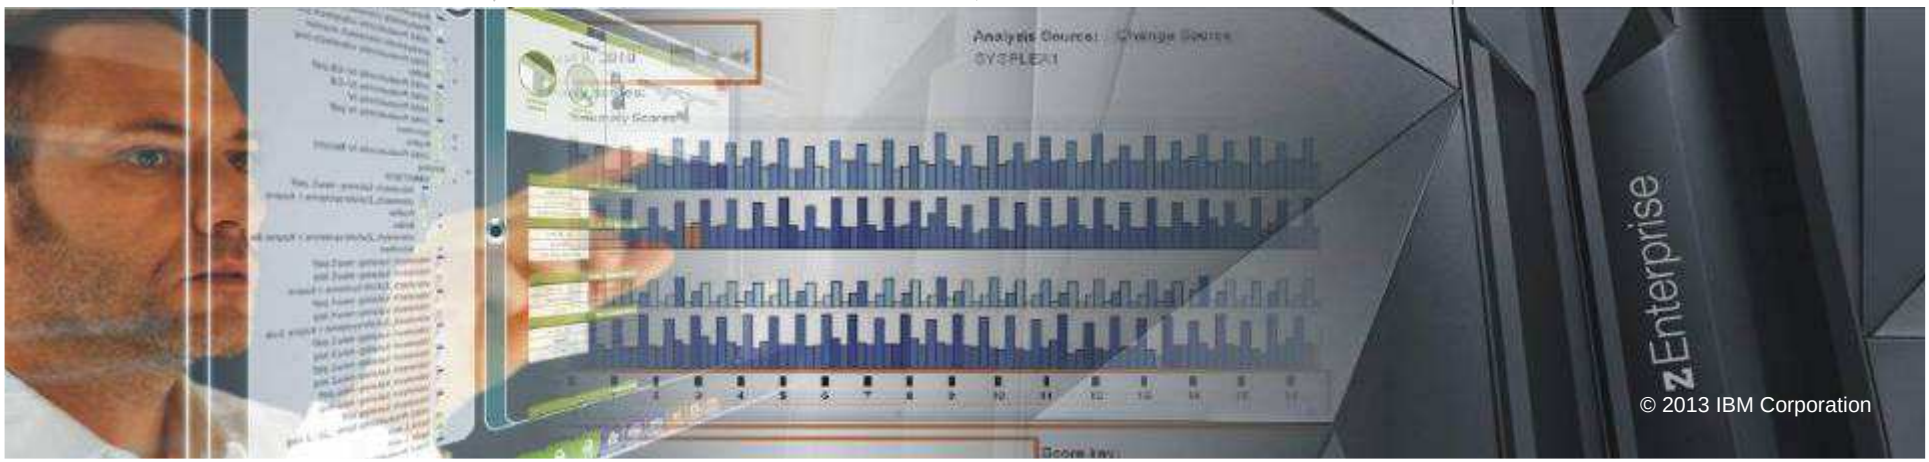

Problem Determination in 2 clicks instead of Hours

Each ten minute reporting interval (2 minutes sampling) is represented by a bar. Height and color coding signal potential anomalies

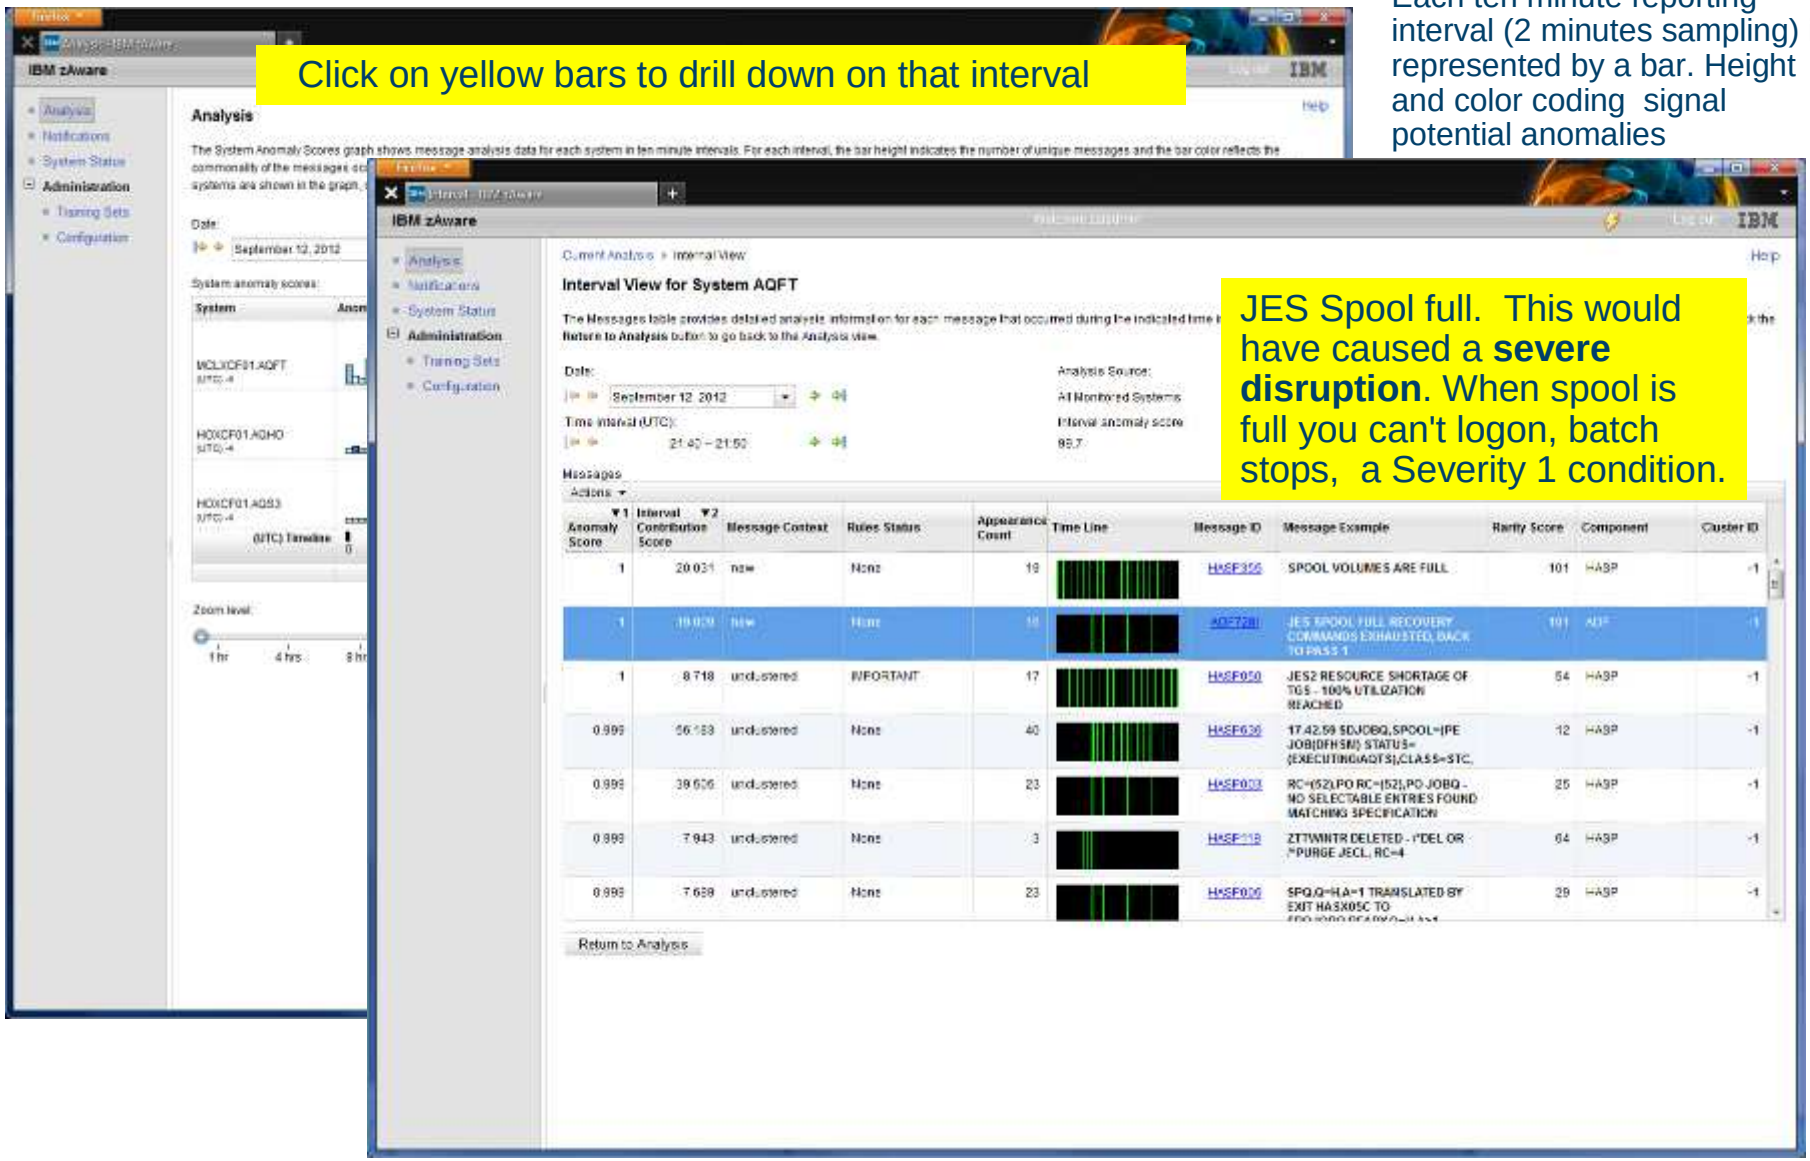

IBM zEnterprise BC12 - 2013

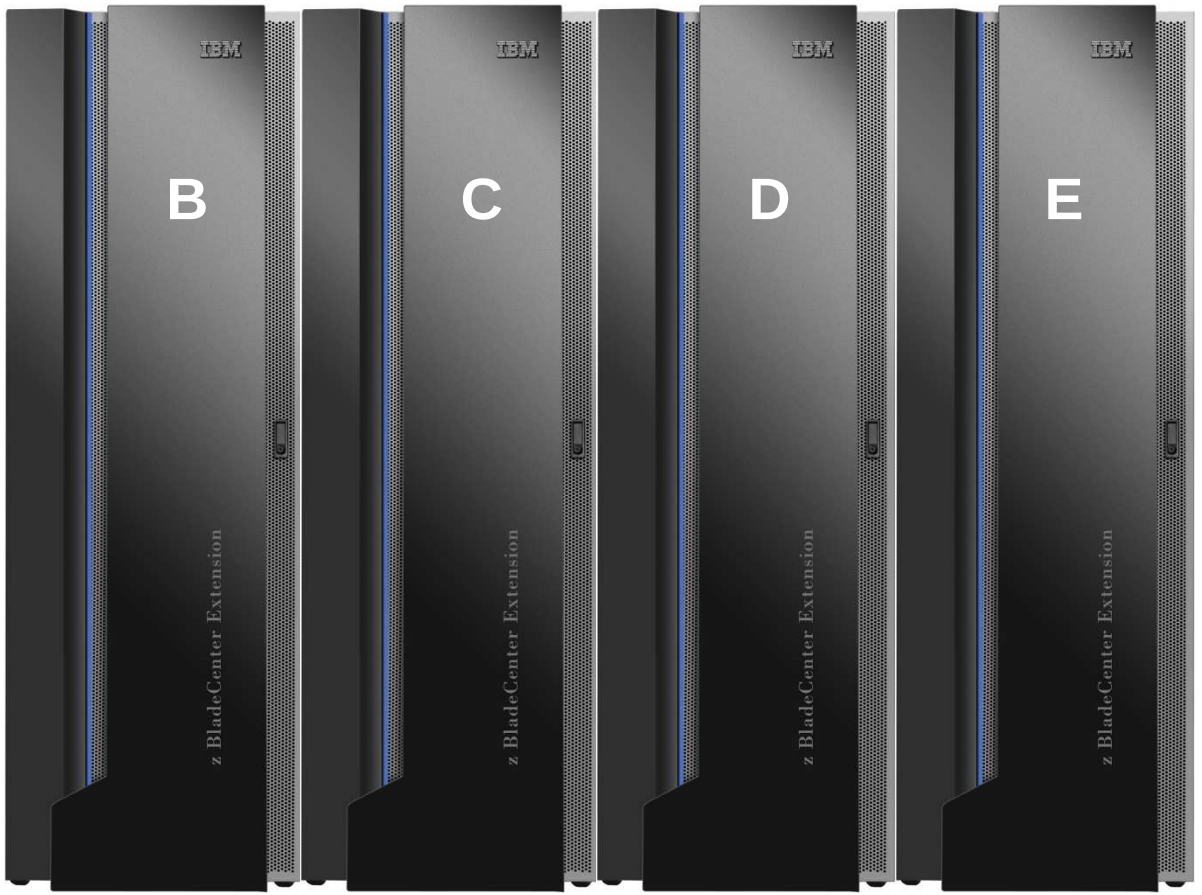

IBM zEnterprise BC12 (zBC12)

IBM zEnterprise BladeCenter Extension (zBX Model 003)

IBM zEnterprise Unified Resource Manager

Note: Covers for zBX. New zBX Model 003s will be shipped with new design covers with a blue stripe. zBX Model 002s upgraded to Model 003s will retain existing design covers

Appendix

IBM System z Generations

N-5

z900
•Announced 10/2000
•770 MHz
•Up to 16 assignable cores
•CP, IFL, ICF
•Up to 64 GB Memory

z800
•Announced 2/2002
•625 MHz
•Up to 4 assignable cores
•CP, IFL, ICF
•Up to 32 GB Memory

N-4

z990
•Announced 5/2003
•1.2 GHz
•Up to 32 assignable cores
•CP, IFL, ICF, zAAP
•Up to 256 GB Memory

z890
•Announced 4/2004
•1.0 GHz
•Up to 4 assignable cores
•CP, IFL, ICF, zAAP
•Up to 32 GB Memory

N-3

z9 Enterprise Class
•Announced 7/2005
•1.7 GHz
•Up to 54 assignable cores
•CP, IFL, ICF, zAAP, zIIP
•Up to 512 GB Memory

z9 Business Class
•Announced 4/2006
•1.4 GHz
•Up to 7 assignable cores
•CP, IFL, ICF, zAAP, zIIP
•Up to 64 GB Memory

N-2

z10 Enterprise Class
•Announced 2/2008
•4.4 GHz
•Up to 64 assignable cores
•CP, IFL, ICF, zAAP, zIIP
•Up to 1.5 TB Memory

z10 Business Class
•Announced 10/2008
•3.5 GHz
•Up to 10 cfg cores (5 CP)
•CP, IFL, ICF, zAAP, zIIP
•Up to 248 GB Memory

N-1

zEnterprise 196
•Announced 7/22/2010
•5.2 GHz
•Up to 80 assignable cores
•CP, IFL, ICF, zAAP, zIIP
•Up to 3 TB Memory

zEnterprise 114
•Announced 7/12/2011
•3.8 GHz
•Up to 10 cfg cores (5 CP)
•CP, IFL, ICF, zAAP, zIIP
•Up to 256 GB Memory

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/ PCIe I/O drawers	Up to 4	Up to 4	Up to 3	Up to 3(1)	Up to 3(1)
I/O slots per I/O drawer/ PCIe I/O drawer	8	8/32		8/32(2)	
FICON® Channels	128	128		128(3)	
OSA Ports	96	96		96	
ESCON® Channels	480	240		0(4)	
IFB host bus Bandwidth PCIe Gen2 Bandwidth	6.0 GB/sec(IFB)	6.0 GB/sec (IFB) 8.0 GB/sec (PCIe)		6.0 GB/sec (IFB) 8.0 GB/sec (PCIe)	
ICB-4/ISC-3(8)/PSIFB	12/48/12	0/48/8 -16	0/48/16 - 32	0(5)/32/8 -16(6)	0(5)/32/16 - 32(7)
zIIP/zAAP Maximum Qty	5	2	5	4 (with Max of 2 CPs)	8 (with Max of 4/5 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)	

Core Design changes compared to z114

- Reduce Finite Cache Penalty
  - New L1/L2 cache structure for L2 latency reduction for both instruction and data misses
    - Bigger L2 (1M-Byte per instruction and Data L2 cache each) with shorter latency
  - D-TLB install buffer to improve d-TLB hit rates, improved hit under miss processing in TLB2
  - Various HW prefetcher and SW prefetch (PFD) handling improvements
- Improve OSC (fetch / store conflict) scheme
  - Additional culprit/victim tracking based on instruction TEXT (not just history based)
- Enhance branch prediction structure and sequential instruction fetching
  - secondary BTB (BTB2) effectively providing 33% more branches
  - Faster prediction throughput in BTB1 by usage of a Fast re-Indexing Table (FIT)
  - Improve sequential instruction stream delivery
- Millicode Handling
  - Selective HW execution of prior millicoded instructions (TR/TRT; STCK*)
  - Streamlined millicode entry to reduce hardware interlocks for commonly used instructions
    - e.g. MVCP
  - MVCL improvements
    - New millicode assist to prefetch data into L4 (in addition to architected addition to prefetch into L3)
    - Overlapping miss handling between
      - L1/2<>L3
      - L3<>L4
      - L4<>Memory
    - Speed up both aligned and unaligned data-moves through caches and through memory

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)