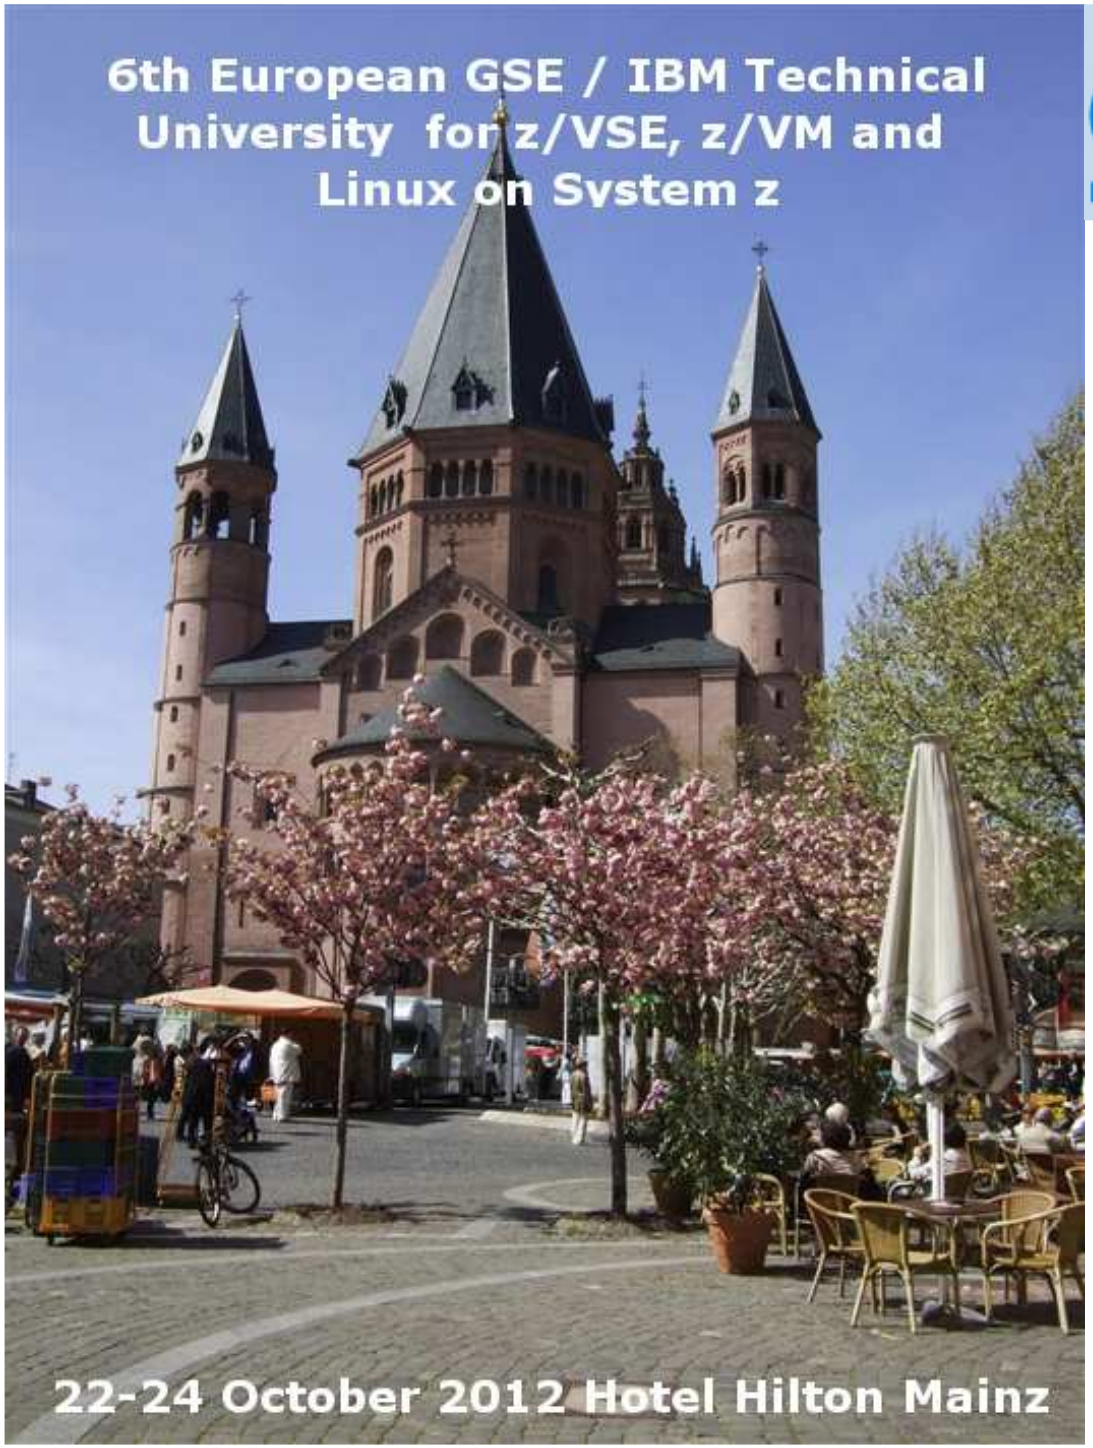

GSE GUIDE SHARE EUROPE Deutsche Region

IBM

IBM zEnterprise - zEC12 News

October 23, 2012

Dr. Manfred Gnirss
gnirss@de.ibm.com
Global Client Center, IBM Germany R&D

© 2012 IBM Corporation

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
- 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
- System z will continue to expand hybrid computing

Plus more flexibility and function by connecting to IDAA

- IBM DB2 Analytics Accelerator (IDAA) allows deployment of business analytics on the same platform as operational applications
- Analytics and OLTP can be run as the same workload

© 2012 IBM Corporation

IBM System z:
Design Comparison for High End Systems

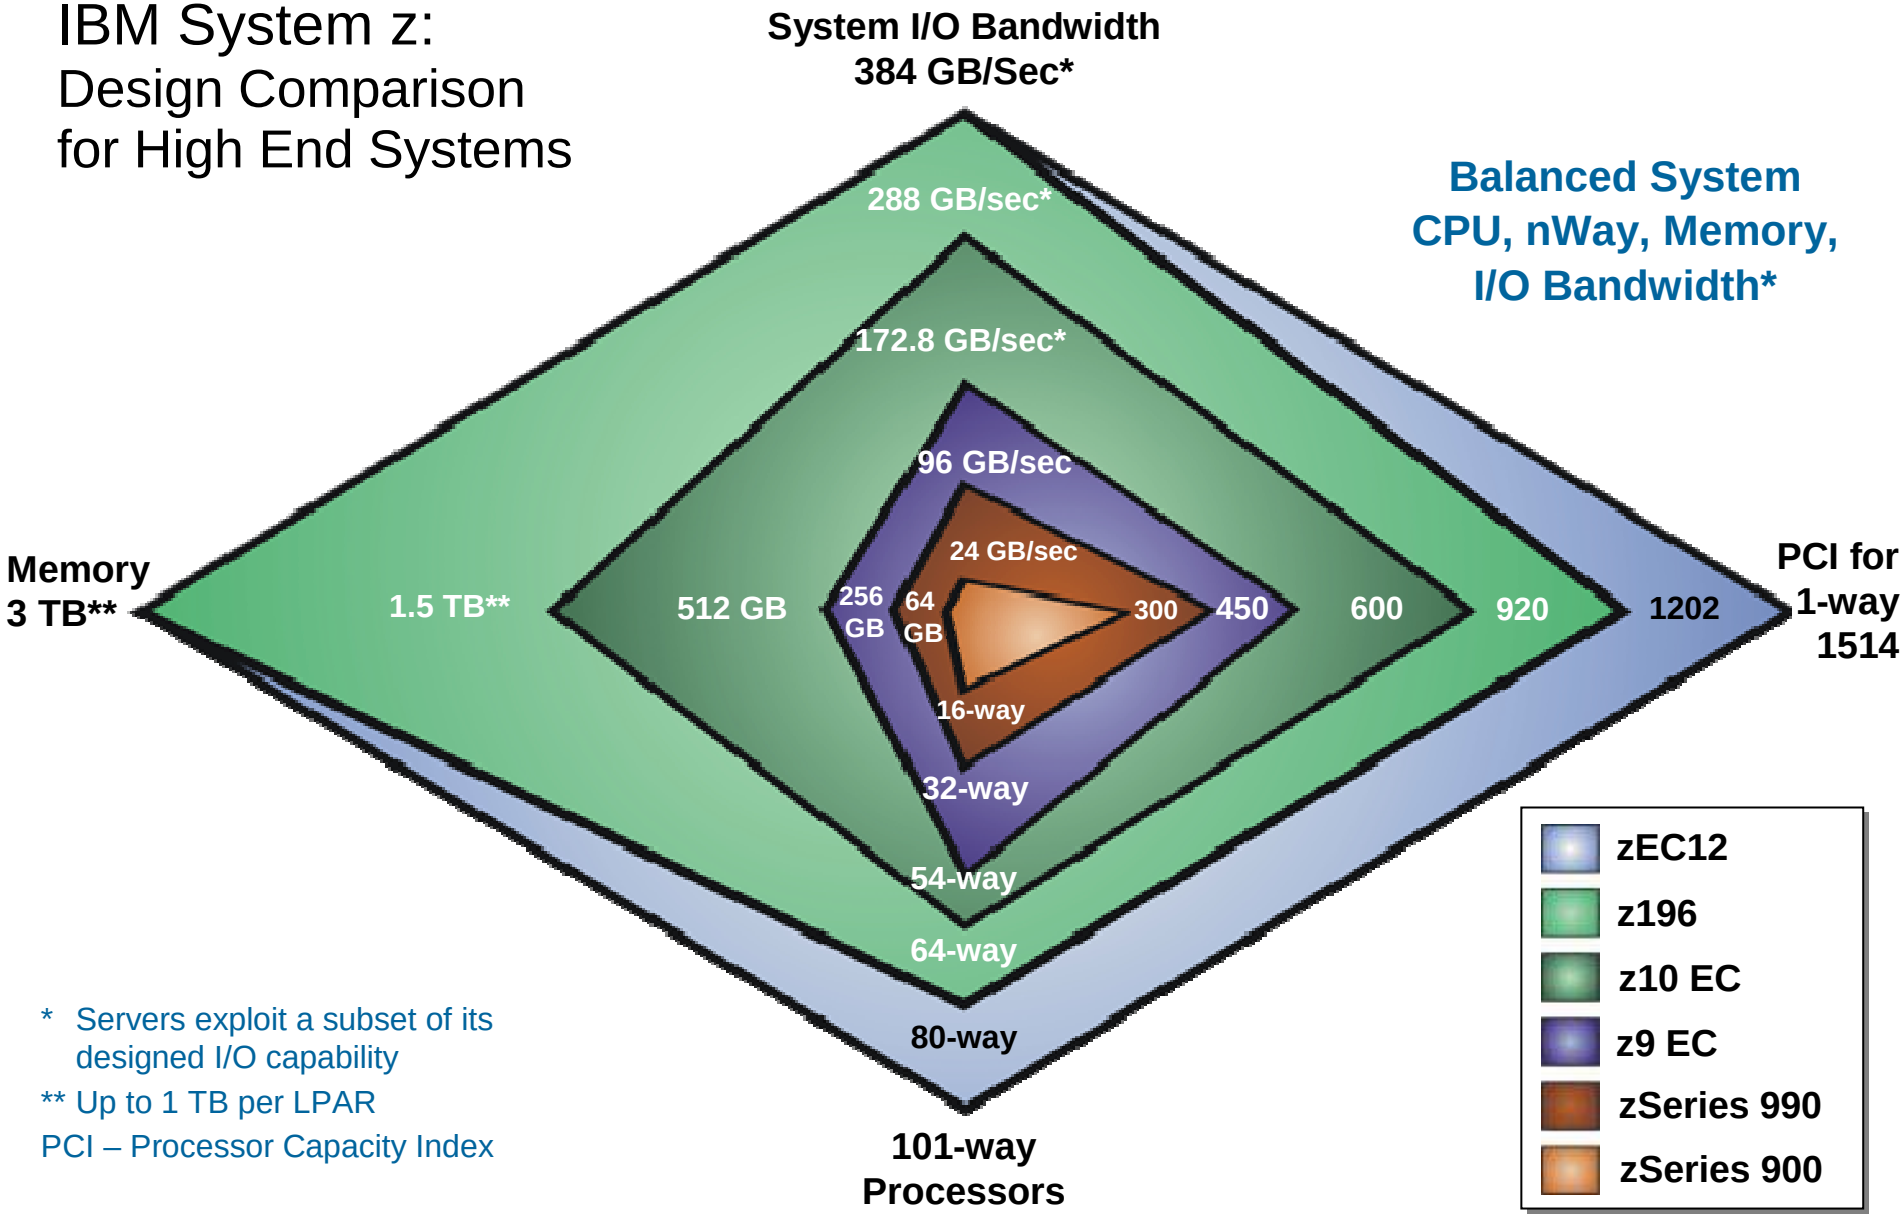

* Servers exploit a subset of its designed I/O capability

** Up to 1 TB per LPAR

PCI – Processor Capacity Index

© 2012 IBM Corporation

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)
  - 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 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)

© 2012 IBM Corporation

zEnterprise EC12 is the core of next generation System z

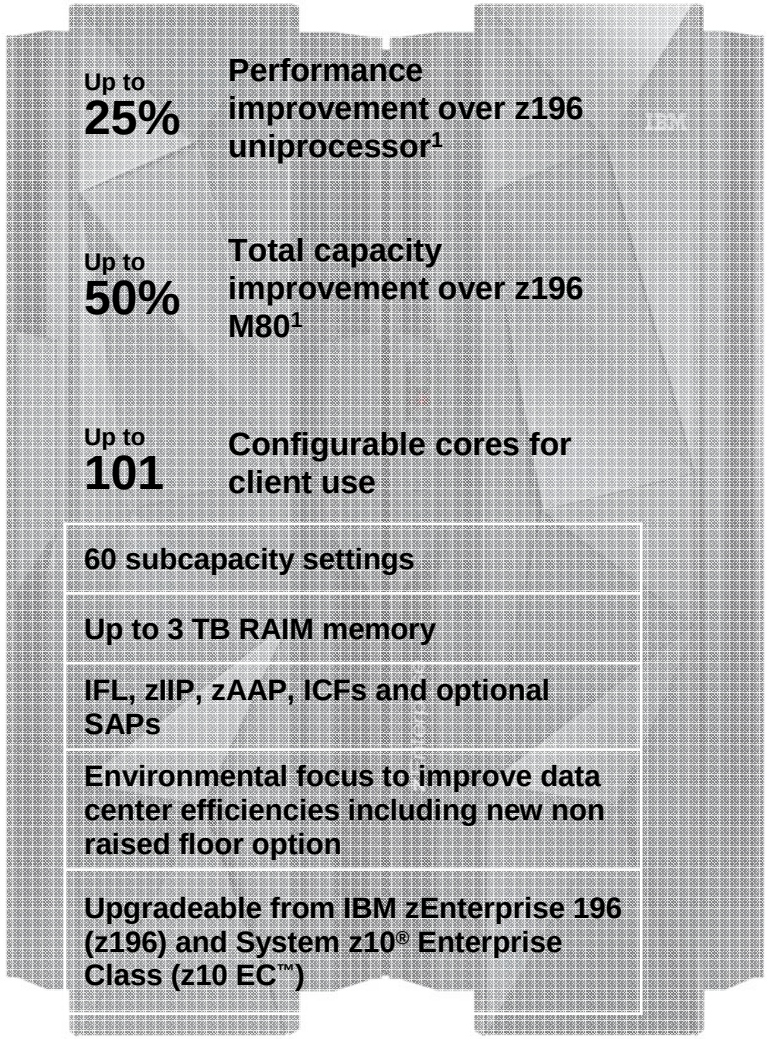

zEC12

Machine Type: 2827

Models: H20, H43, H66, H89, HA1

- Advanced Technology 5.5 GHz processor chip for performance boost for all workloads
  - Over ***78,000 MIPS*** for large scale consolidation
  - ***Larger cache*** for data serving
- Processor chip optimized for software performance
  - Advanced performance functions exploited by ***Java, PL/I, compilers, DB2*** and more
- Innovation to drive availability to superior levels
  - ***IBM zAware*** with out-of-band analytics provide point in time snap-shot of the current state of your business and can help you improve availability
  - ***FLASH Express and pageable large pages*** to drive availability and performance for critical workloads
- Security and reliability are in our DNA
  - High speed ***cryptography integrated as part of the chip***
  - Enhanced support for applications requiring data encryption, cryptographic keys and digital signing with new ***Crypto Express4S***
  - PR/SM designed for ***EAL5+ certification***

[1] Based on preliminary internal measurements and projections. 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.

© 2012 IBM Corporation

zEC12 Vs z196 Vs z10 EC Vs z9 EC capacity comparison

- z9 EC up to 54-way: 193 – 18,505 PCIs
- z10 EC up to 64-way: 214 – 31,826 PCIs
- z196 up to 80-way: 240 – 52,286 PCIs
- zEC12 up to 101-way: 240 – over 78,426 PCIs

Note: For zEC12, z/OS V1.13 LSPR numbers quoted. For the rest z/OS 1.11

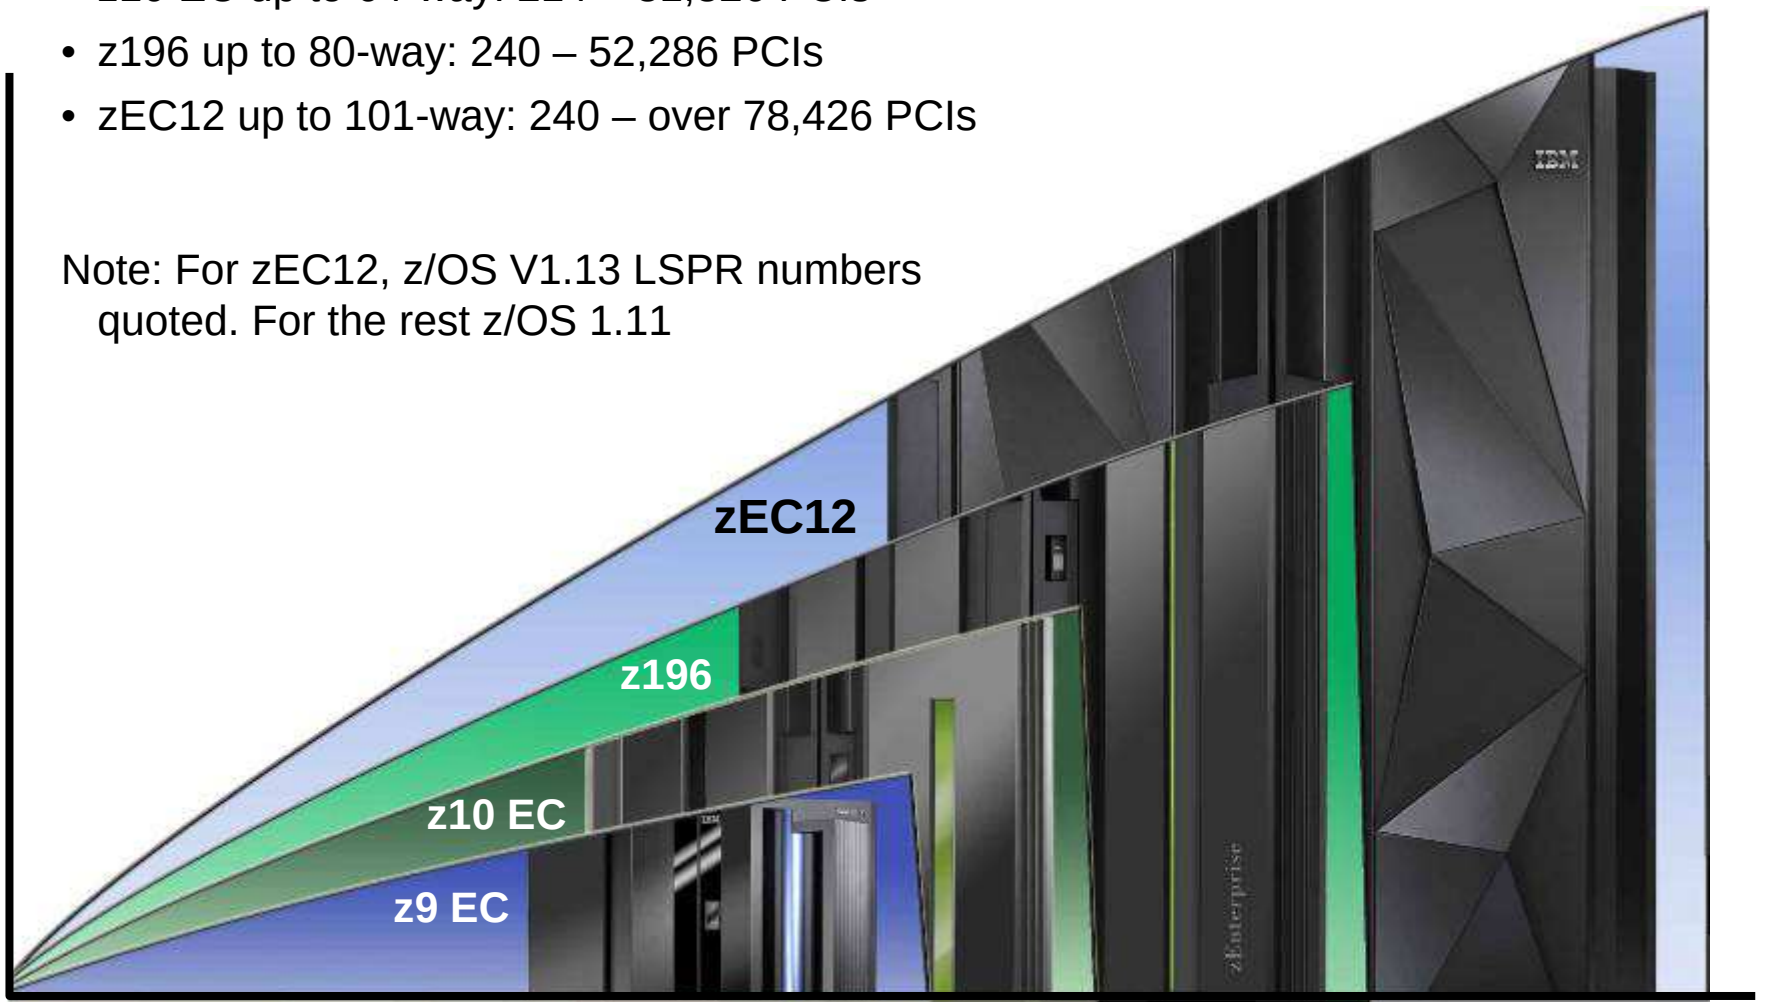

© 2012 IBM Corporation

zEC12 Continues the CMOS Mainframe Heritage Begun in 1994

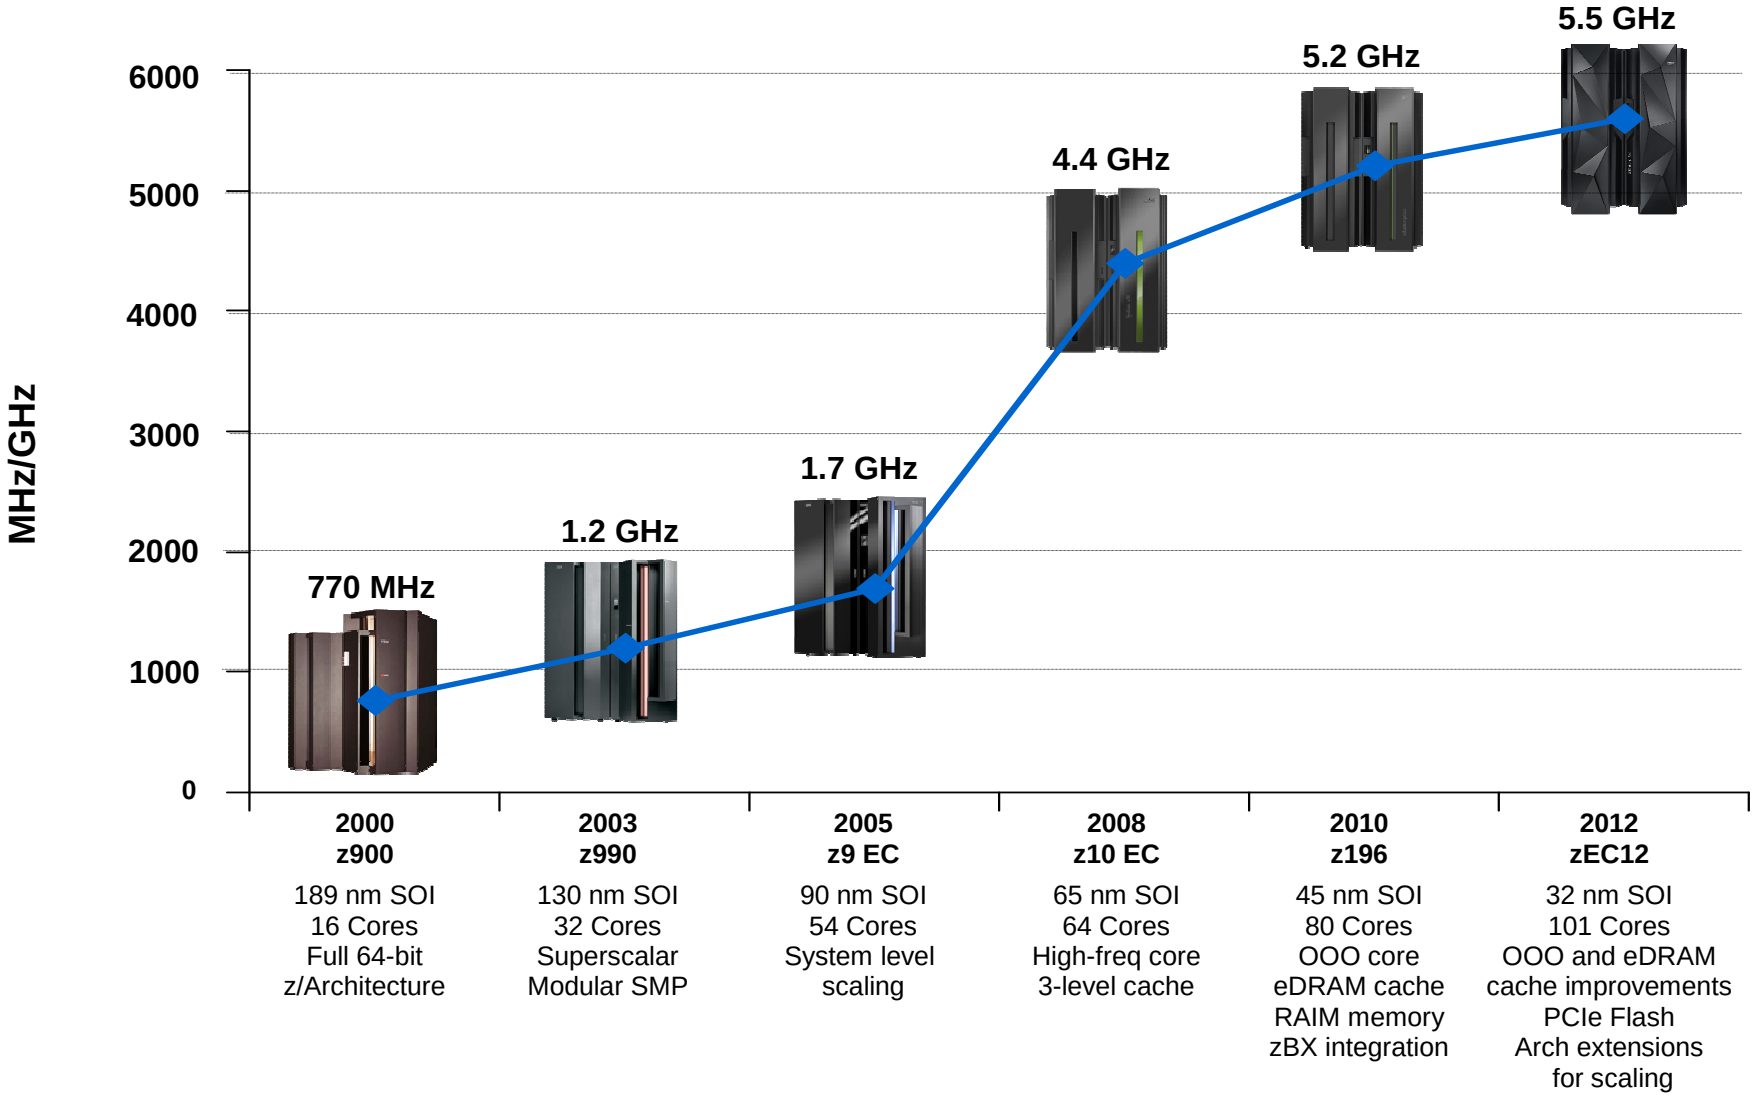

© 2012 IBM Corporation

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 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 scalability
  - ***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
  - Up to ***30% improvement in IMS throughput*** due to faster CPU and cache, compliers, and more[1]
  - New IBM Enterprise PL/I compiler is planned to exploit and get a performance boost from ***decimal format conversions facility***

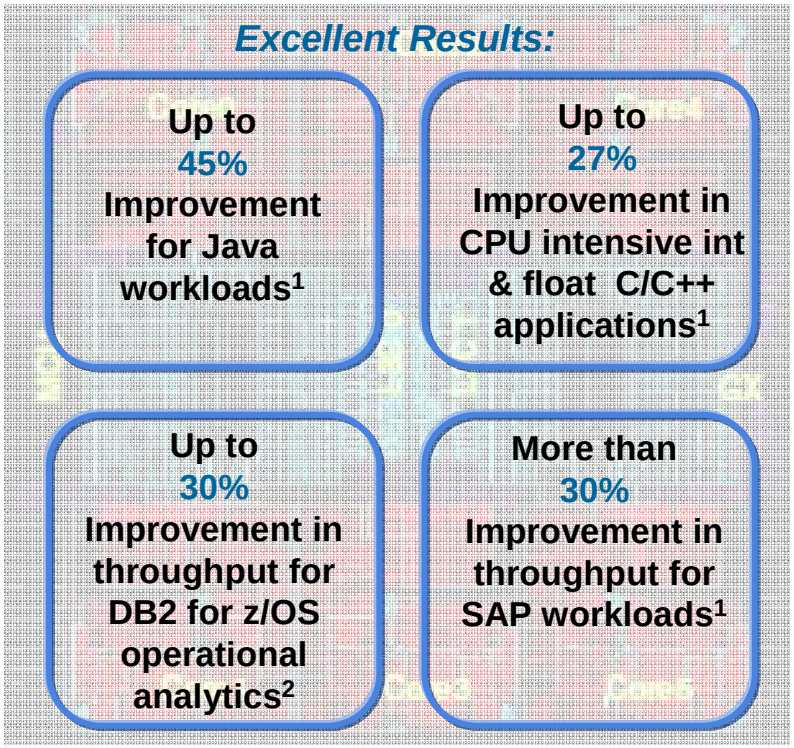

[1] Based on preliminary internal measurements and projections
[2] Aa measured by the IBM 9700 Solution Integration Center. The measured operational BI workload consists of 56 concurrent users executing a fixed set of 160,860 Cognos reports . Compared DB2 v10 workload running on IBM's z196 w/10 processors to an zEC12 w/10 processors

© 2012 IBM Corporation

zEC12 – Overall Attributes Highlights (compared to z196)

- 50% more cores in a CP chip
  - Up to 5.7% faster core running frequency
  - Up to 25% capacity improvement over z196 uni-processor
- Bigger caches and shorter latency
  - Total L2 per core is 33% bigger
  - Total on-chip shared L3 is 100% bigger
  - Unique private L2 designed to reduce L1 miss latency by up to 45%
- 3rd Generation High Frequency, 2nd Generation Out of Order Design
  - Numerous pipeline improvements based on z10 and z196 designs
  - # of instructions in flight is increased by 25%
- New 2nd level Branch Prediction Table for enterprise scale program footprint
  - 3.5x more branches
- Dedicated Co-processor per core with improved performance and additional capability
  - New hardware support for Unicode UTF8<>UTF16 bulk conversions
- Multiple innovative architectural extensions for software exploitation

© 2012 IBM Corporation

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

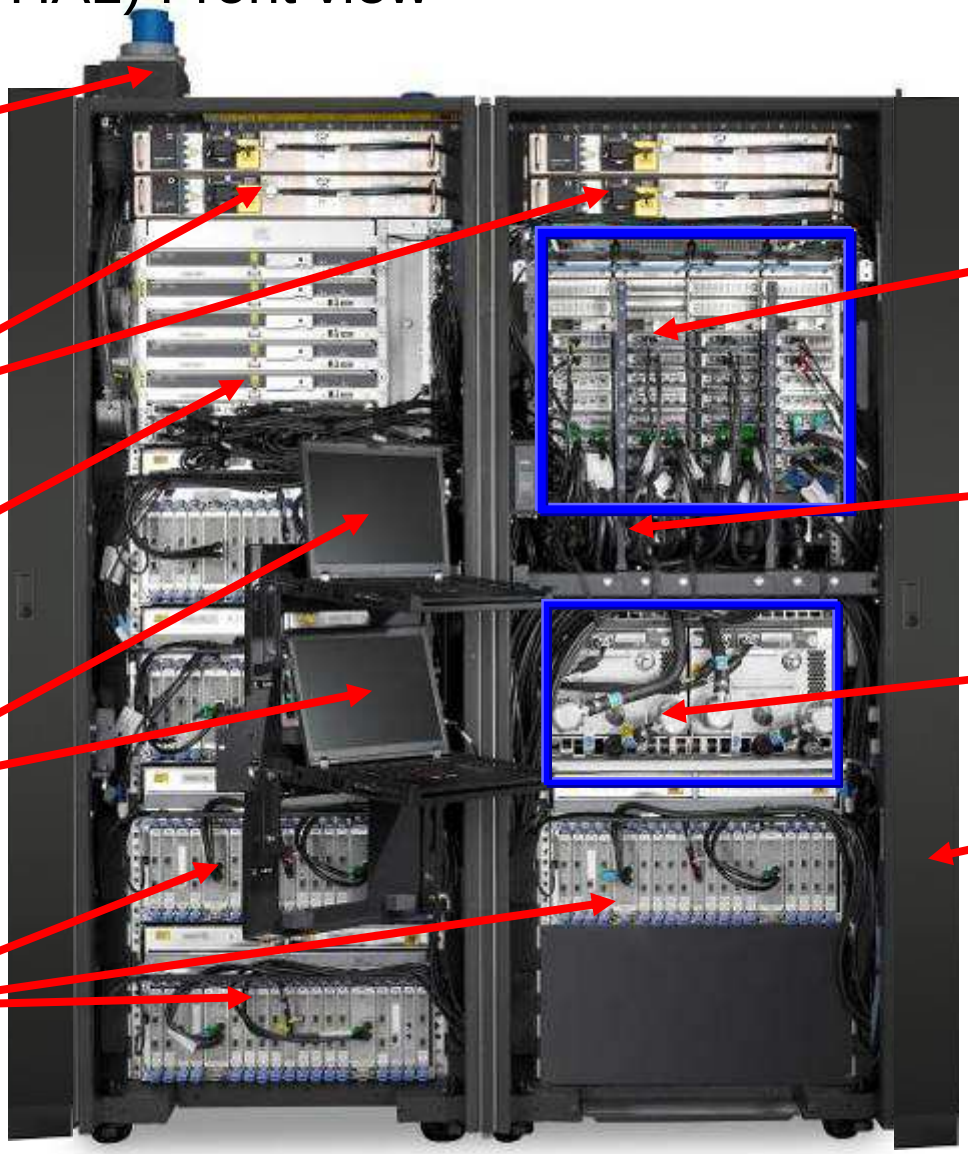

Overhead Power Cables (option)

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

© 2012 IBM Corporation

zEC12 Water cooled (Upgraded from a z196) - Under the covers (Model H89 or HA1) Front view

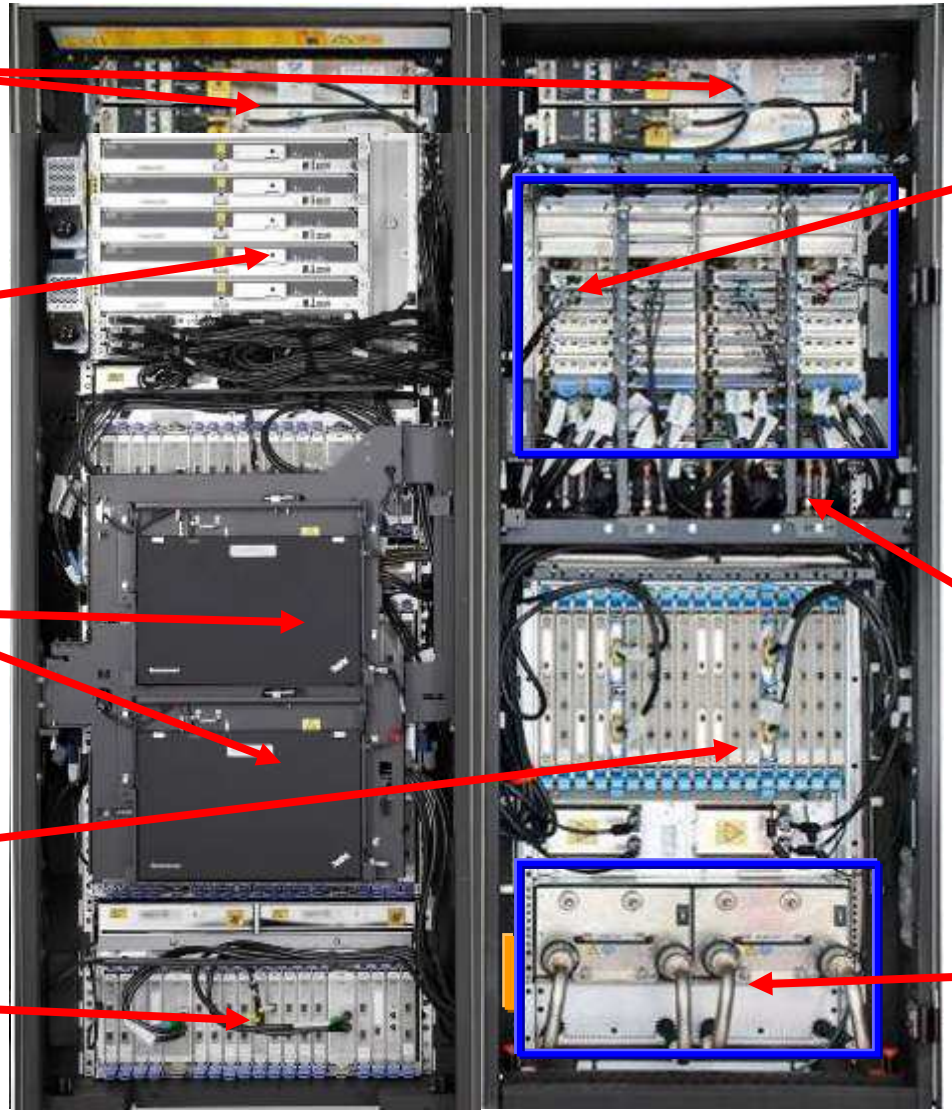

© 2012 IBM Corporation

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

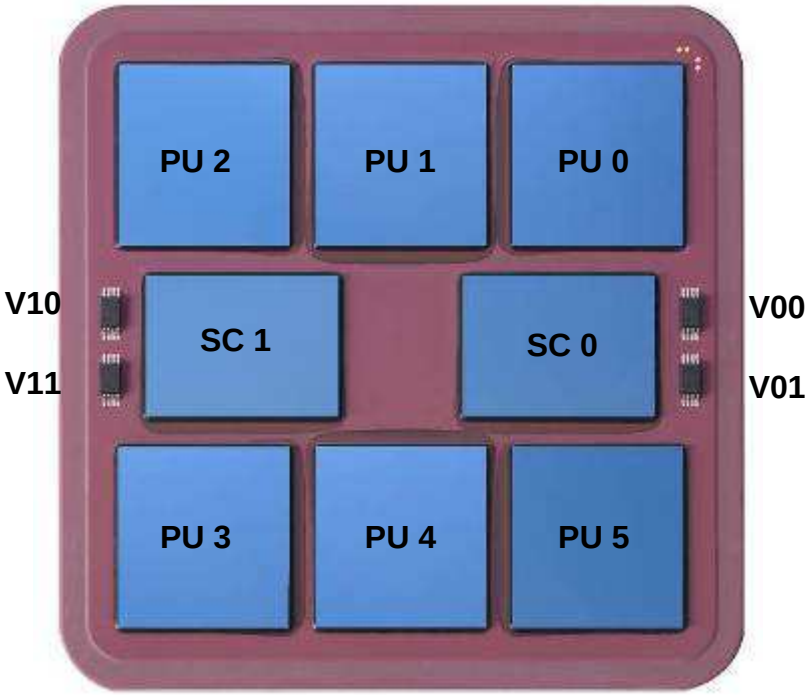

- CMOS 13s chip Technology
  - PU, SC, S chips, 32nm
  - 6 PU chips/MCM – Each up to 6 active cores
    - 23.7 mm x 25.2 mm
    - 2.75 billion transistors/PU chip
    - L1 cache/PU core
      - 64 KB I-cache
      - 96 KB D-cache
    - L2 cache/PU core
      - 1 MB I-cache
      - 1 MB D-cache
    -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

© 2012 IBM Corporation

zEC12 Hexa Core PU Chip Details

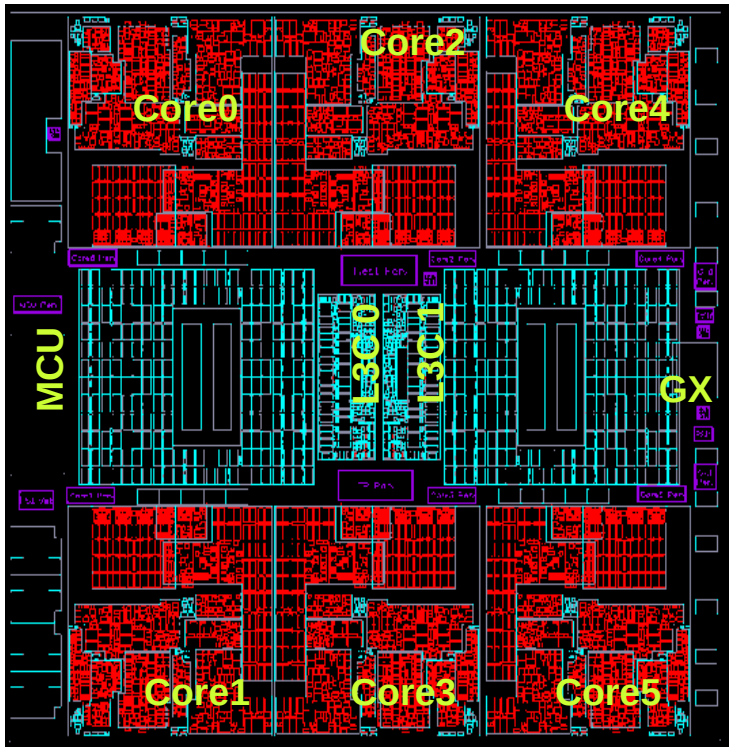

- 13S 32nm SOI Technology
  - 15 layers of metal
  - 7.68 km wire
- 2.75 Billion Transistors

- 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

© 2012 IBM Corporation

zEC12 Book Layout

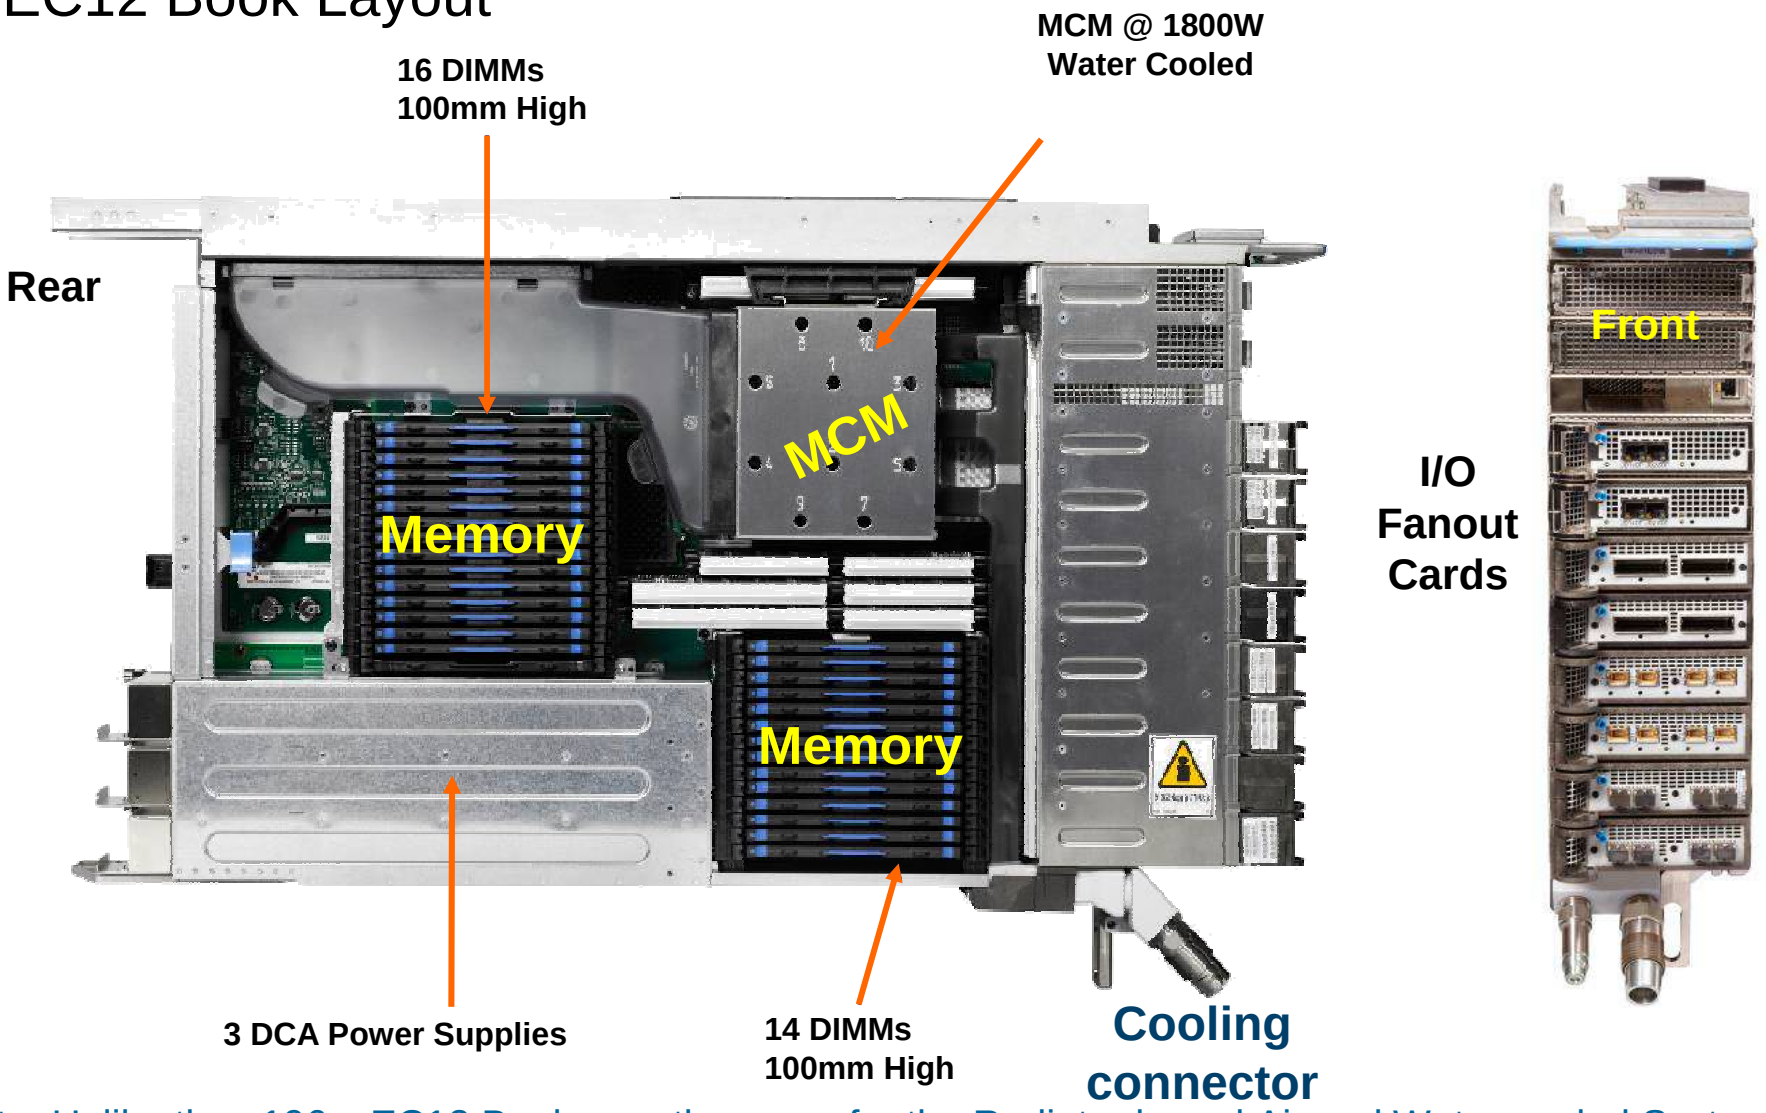

Note: Unlike the z196, zEC12 Books are the same for the Radiator based Air and Water cooled Systems

© 2012 IBM Corporation

zEC12 Processor Unit allocation/usage

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
  - Each MCM uses PU chips with a combination of 4, 5 and 6 active cores
- The maximum number of logical ICFs or logical CPs supported in a CF LPAR 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

© 2012 IBM Corporation

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 exploitation by Java; longer-term opportunity for DB2, z/OS, others
- Runtime instrumentation
  - Real-time information to software on dynamic program characteristics
  - Enables increased optimization in JVM/JIT recompilations
  - Additional exploitation opportunities in the works
- 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
- Decimal format conversions
  - Enable broader exploitation of Decimal Floating Point facility by COBOL programs

© 2012 IBM Corporation

Performance Drivers with zEC12

- Hardware
  - Memory subsystem
    - Focus on keeping data "closer" to the processor unit
    - Larger L2 cache
    - Double size of both chip-level and book-level shared caches
  - Processor
    - Improved Out-Of-Order execution
    - Better branch prediction – added BTB2 (Branch Target Buffer)
    - Add another execution unit pair
    - Improved grouping of instructions (pipeline)
    - Up to 6 processor units per chip
  - Up to 101 configurable processor units
  - 3 sub-capacity settings
- HiperDispatch
  - Exploits new chip configuration
- Variability amongst workloads
  - Moving from z196 to zEC12 expected to be less than last few migrations
    - Fairly balanced improvements to both processor and memory subsystem
  - Moving from z10 and earlier machines expected to be similar to moves to z196

© 2012 IBM Corporation

System z Cache Topology – z196 vs. zEC12 Comparison

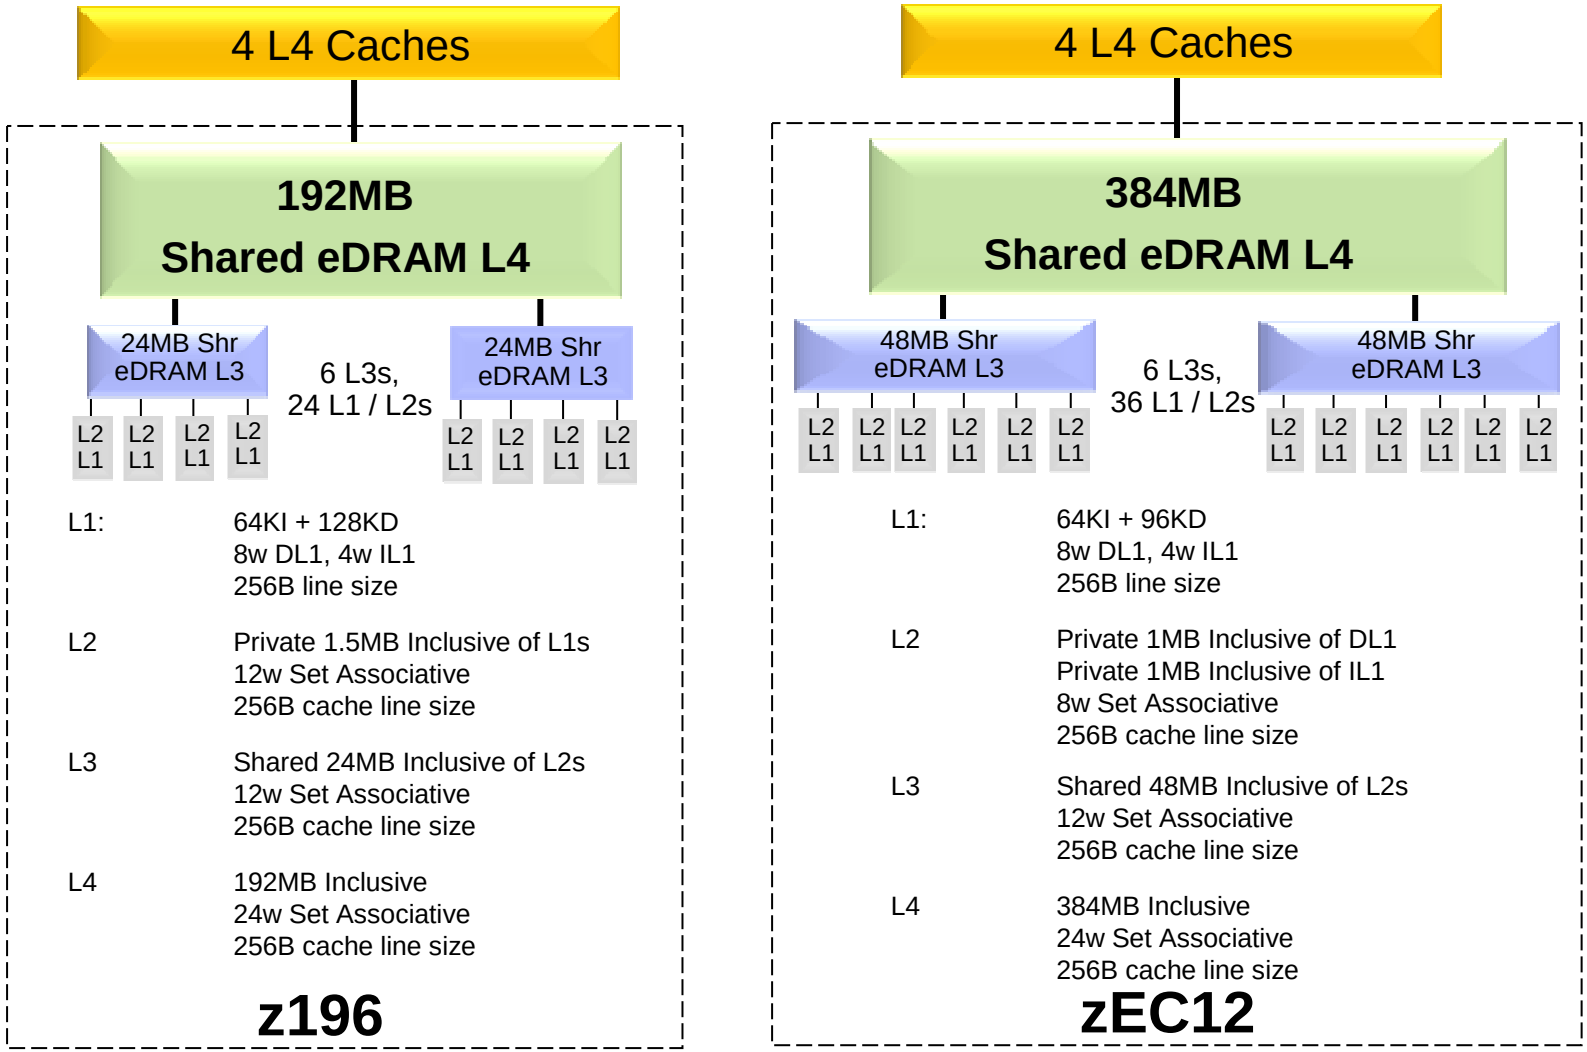

© 2012 IBM Corporation

z196 and zEC12 System Compression and Cryptography Accelerator

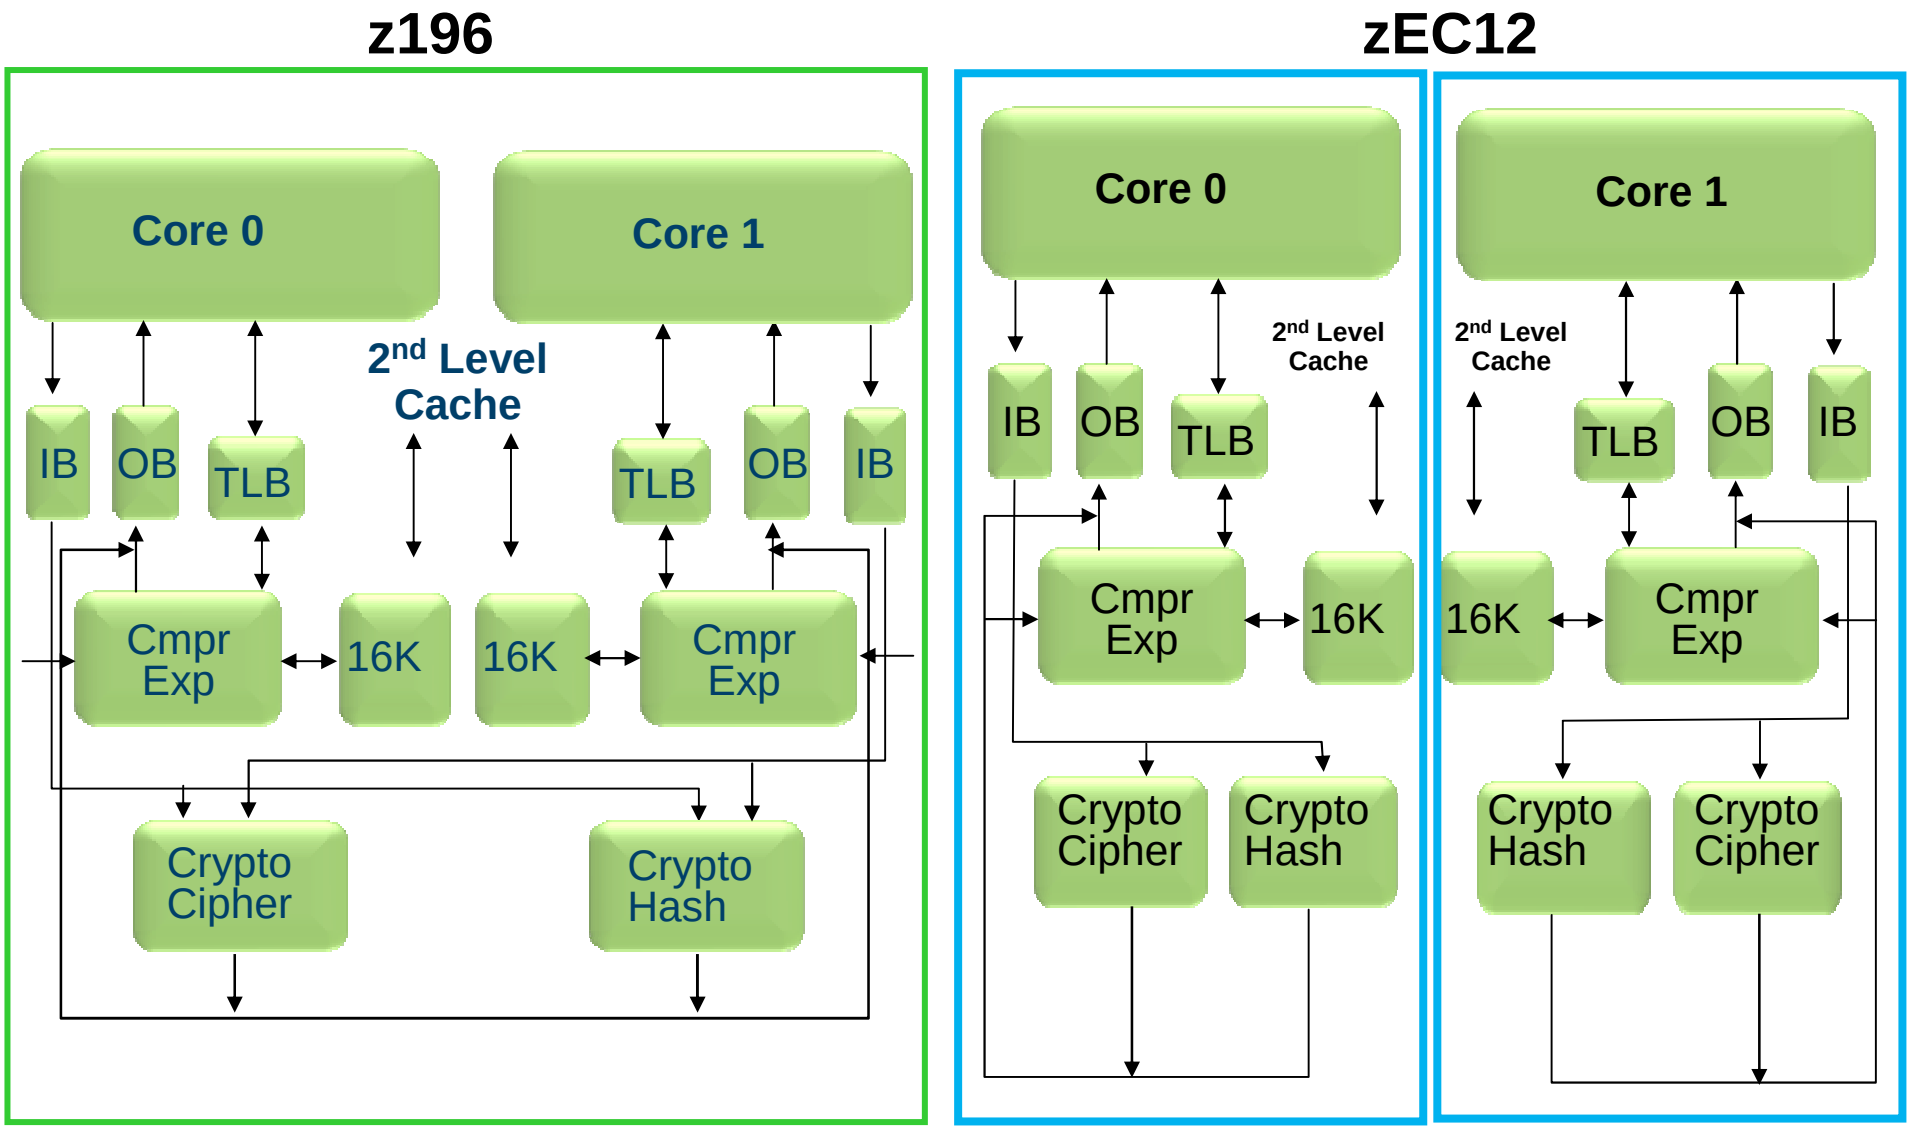

© 2012 IBM Corporation

Out of Order Execution – z196 Vs zEC12

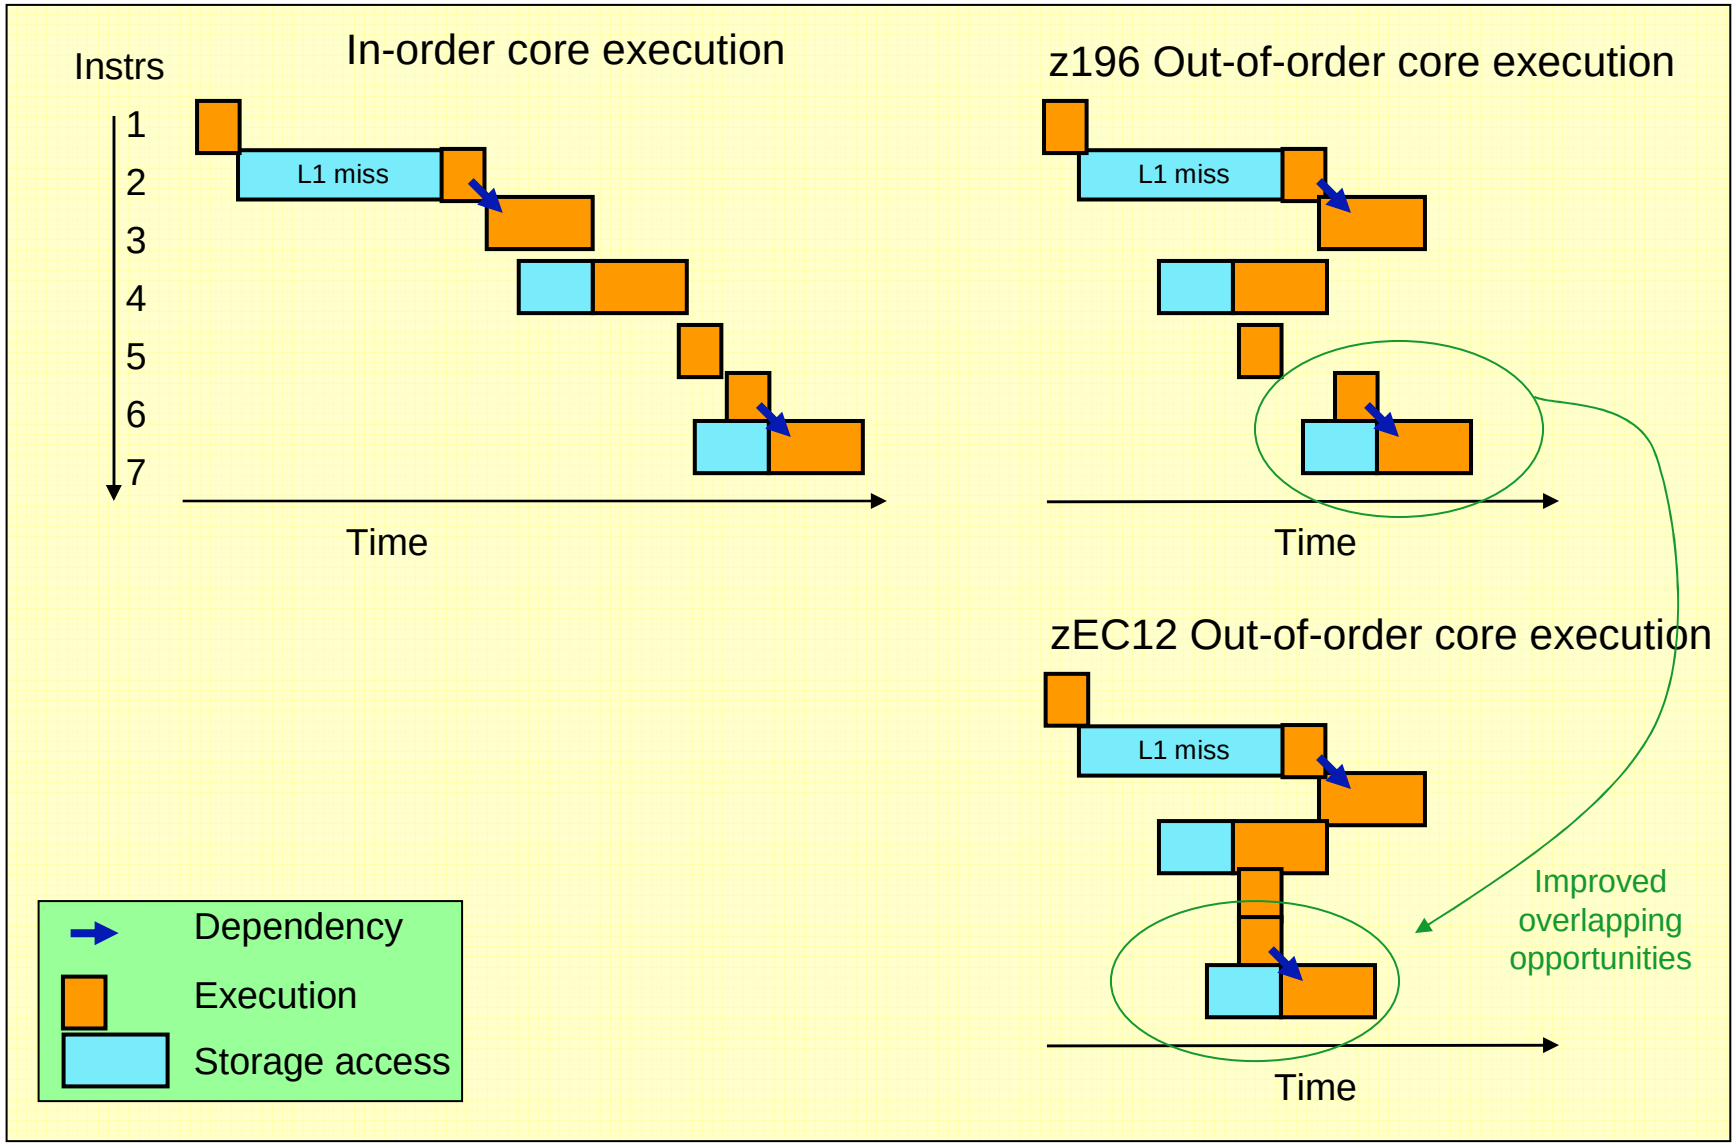

© 2012 IBM Corporation

zEC12 OoO - Improved instruction delivery and execution

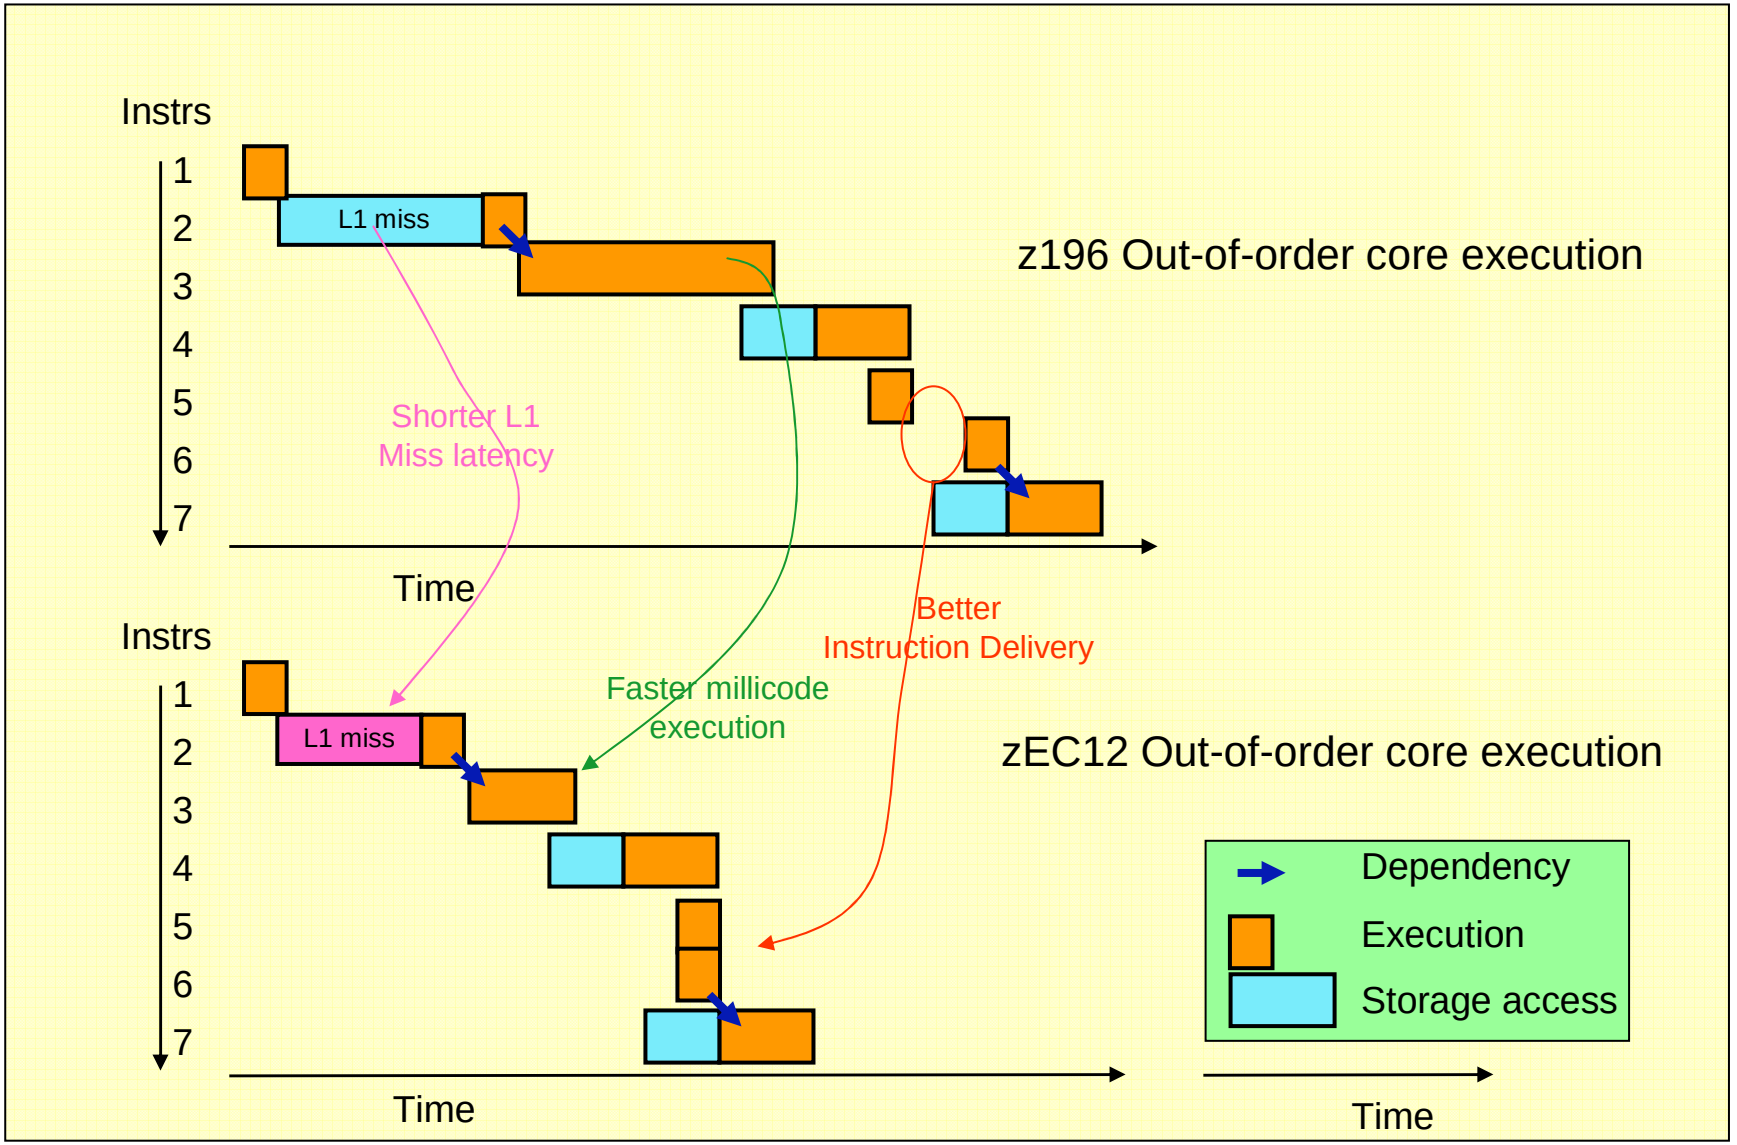

© 2012 IBM Corporation

Introducing Flash Express

- Flash Express is intended to improve System z availability
  - Slash latency delays from paging
    - Flash Memory is much faster than spinning disk
    - Flash Memory is much slower than main memory
    - Flash Memory takes less power than either
  - Make your start of day processing fast
  - Designed to eliminate delays from SVC Dump processing
- zEC12 offers optional Flash Express memory cards
  - Supported in PCIe I/O drawer with other PCIe I/O cards
  - Installed in pairs for availability
  - No HCD/IOCP definitions required
- Assign Flash Memory to partitions like main memory
  - Assignment is by 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

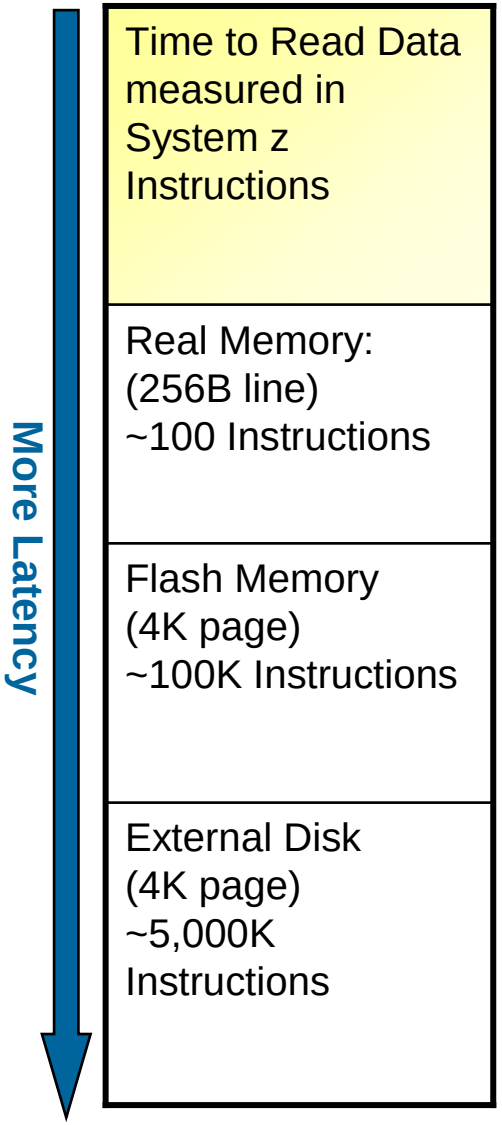

© 2012 IBM Corporation

Relative Access Times for different technologies

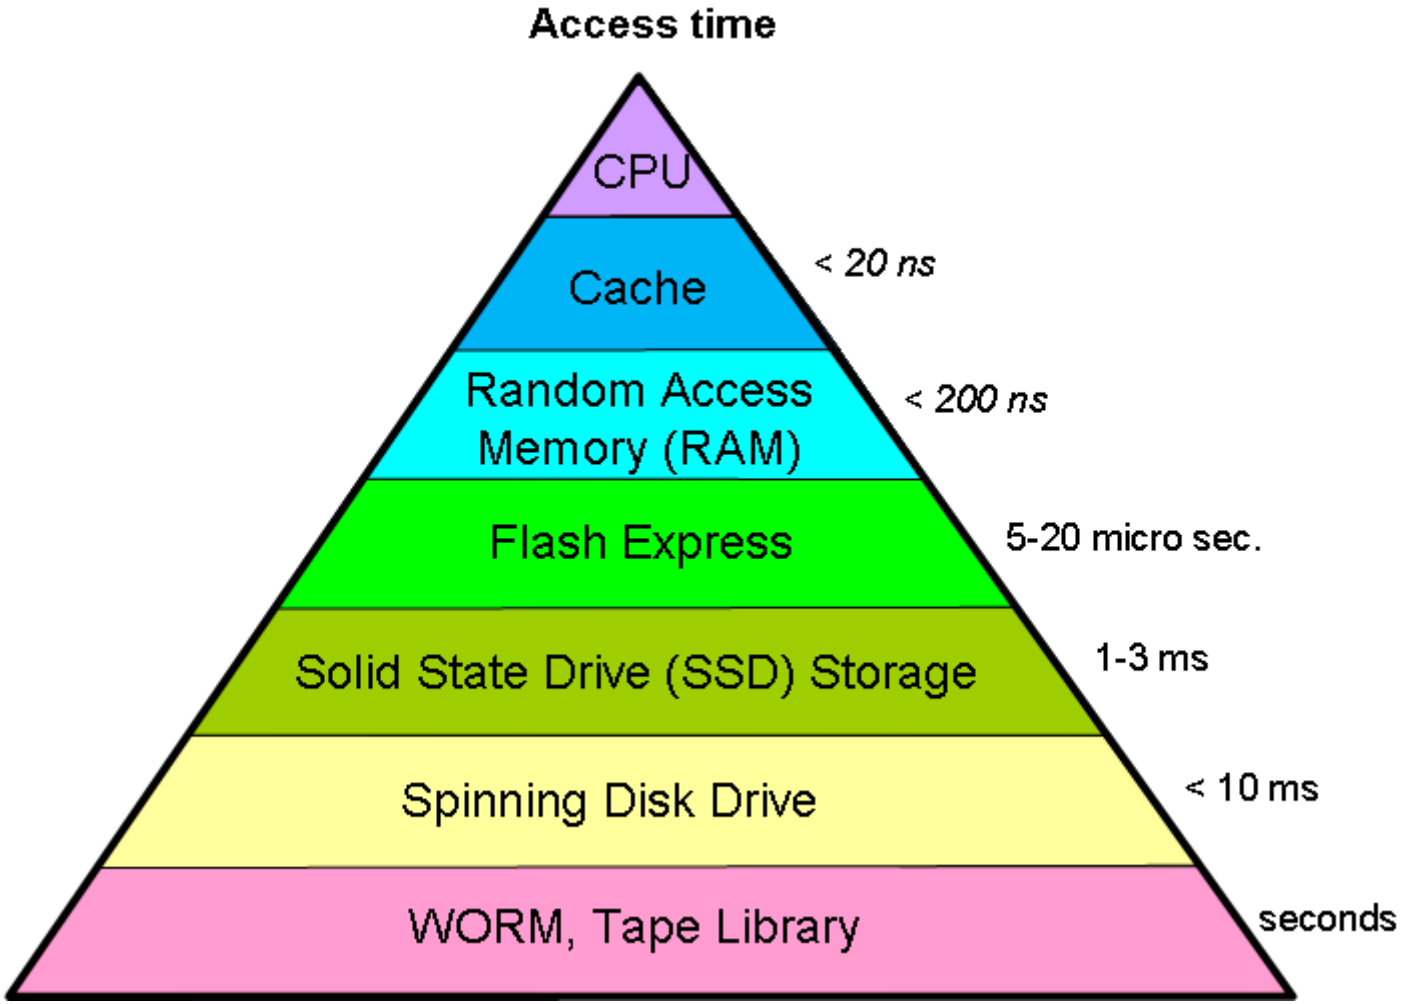

© 2012 IBM Corporation

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 cards
  - Plugs into PCIe I/O drawers in pairs
  - Data security provided on the feature cards
  - A pair provides 1.6 TB of storage
  - A maximum of 4 pairs are supported in a system
- 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 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

© 2012 IBM Corporation

z/OS FLASH Use Cases

- Paging
  - z/OS paging subsystem will work with mix of internal Flash and External Disk
    - Self Tuning based on measured performance
    - Improved Paging Performance, Simplified Configuration
  - Begin Paging 1 MB Large Pages only to Flash
    - Exploit Flash's random I/O read rate to gain CPU performance by enabling additional use of Large Pages. Currently large pages are not pagable.
  - Begin Speculative Page-In of 4K Pages
    - Exploit Flash's random I/O read rate to get Improved Resilience over Disruptions.
    - Market Open, Workload Failover,

© 2012 IBM Corporation

z/OS Flash Use Cases …

- Dumping
  - Minimize SVC Dump duration, System impact
    - Flash performance during SDUMP
    - Flash performance after SDUMP
  - Reduce Stand Alone Dump duration
    - Read time for paged out data

© 2012 IBM Corporation

Typical Customer Configurations for FLASH

- Flash card pair memory size is 1.6TB
  - Min: 1 Card Pair
- Typical customer configuration is 6 to 8 LPARs per CPC and 40 GB - 80GB for paging configuration dataset size
- Even with 10 LPARs per CPC, each LPAR has 160 GB of flash memory available for its paging datasets, more than double the current typical customer configuration
  - All paging data can easily reside on Flash
  - Data will preferably go to flash and only go to disk (if any) when flash is full
  - No intelligent placement of data on internal flash needed

© 2012 IBM Corporation

Flash Express PCIe Adapter Card

Must have 4 x SSD cards. Each 400 GBs

© 2012 IBM Corporation

Flash Express - Twin-Cable to Form a RAID 10 Mirrored Pair

- Flash Express cards are always installed in pairs
  - Maximum 4 pairs in a System
- Installed in a PCIe I/O Drawer in 2 different I/O Domains
  - Maximum of 2 pairs installed in a drawer
  - One Flash Card per Domain only
  - Greater than 2 pairs will require a second PCIe I/O Drawer
  - eConfig will reserve a slot in each Domain in case Flash Express is ordered in the future
  - Cards first installed in the front of the installed drawers (slots 1 and 14) before using the rear slots (25 and 33)
- Flash Express Cards are cabled together to form a RAID 10 Mirror for redundancy
- Data on the Flash Card is protected with a unique key stored on the Support Element (SE) harddisk
  - Only useable on the system with the key that encrypted it
  - Secure Key Store is implemented via a 'Smart Card' that plugs into the SE Smart Card reader
  - Smart Card contains both a unique key personalised for each system and a small Crypto engine that can perform a set of security functions within the card
  - AES encryption

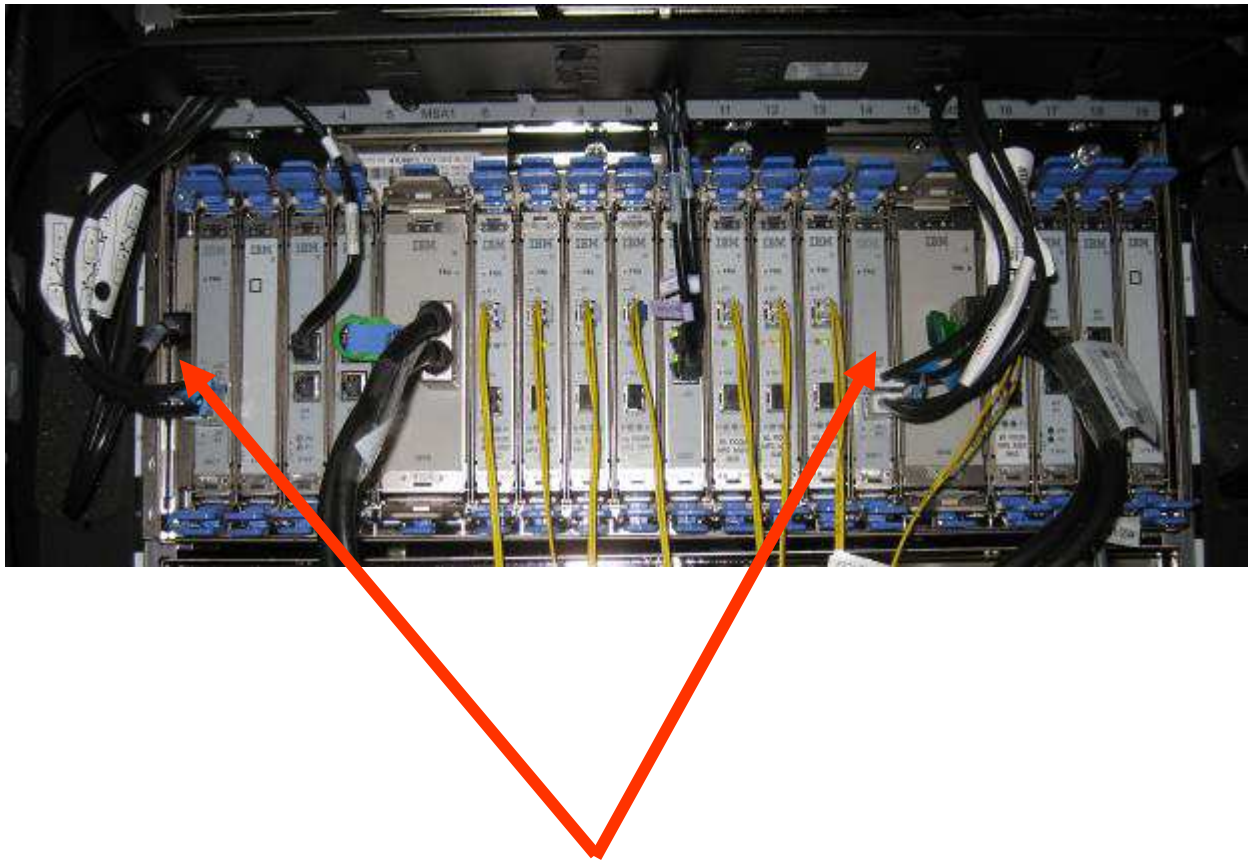

Twin-Cable interconnect between the two Flash Express cards in slots 1 and 14

© 2012 IBM Corporation

Allocating Flash Express

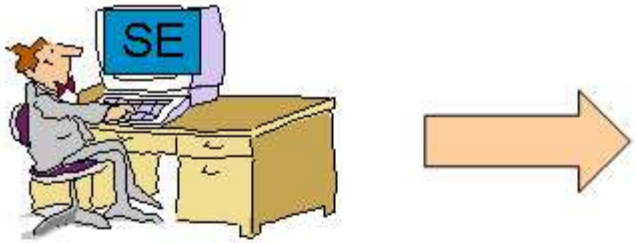

Allocating Flash to a partition

- The initial and maximum amount of Flash Memory available to a particular logical partition is specified at the SE or HMC via a new Flash Memory Allocation panel
- Can dynamically change maximum amount of Flash Memory available to a logical partition
- Additional Flash Memory (up to the maximum allowed) can be configured online to a logical partition dynamically at the SE or HMC
  - For z/OS this can also be done via an operator command
- Can dynamically configure Flash Memory offline to a logical partition at the SE or HMC
  - For z/OS this can also be done via an operator command
- Predefined sub-channels, no IOCDS
  - Sub-channels are allocated from the .25K reserved in sub-channel set 0

P00MNXK4: Manage Flash Memory Allocation

Manage Flash Memory Allocation - P00MNXK4

Summary

Allocated:	32 GB	Storage increment:	16 GB
Available:	0 GB	Rebuild Complete	0 %
Uninitialized:	0 GB		
Unavailable:	0 GB		
Total:	32 GB		

Partitions

--- Select Action ---

Select	Partition Name	Status	IOCDS	Allocated (GB)	Maximum (GB)
◉	LP1	Inactive	A0,A1,A2,A3	32	240
○	LP2			0	0
○	LP3			0	0
○	LP4			0	0
○	LP5			0	0
○	LP6			0	0
○	LP7			0	0
○	LP8			0	0

OK | Apply | Refresh | Cancel | Help

© 2012 IBM Corporation

Flash Express Virtualization

- Full virtualization of physical Flash PCIe cards across partitions, software sees an Abstracted Flash Storage Space…
  - Allows each logical partition to be configured with its own SCM address space
  - Allocate Flash to partitions by amount, not card size
  - Ability to change underlying technology while preserving API
- No Hardware Specifics in Software
  - Error Isolation, Transparent mirroring, Centralized diagnostics, etc.
  - Hardware Logging, FRU Call, Recovery: Independent of software

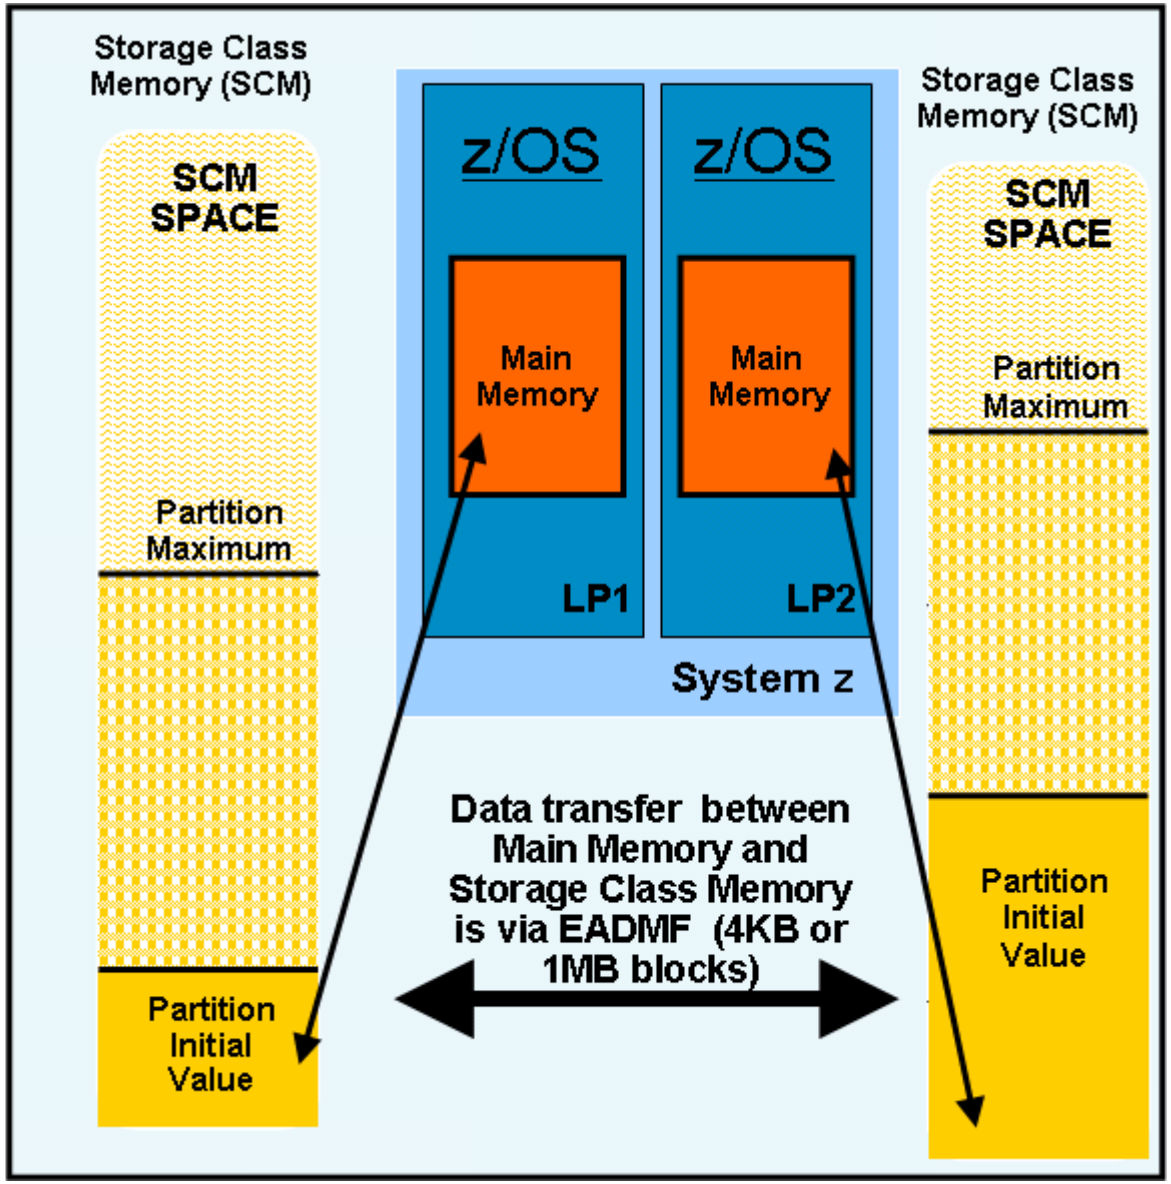

© 2012 IBM Corporation

IBM z/OS Solutions Address Problem Determination

Solutions Available:		Rules based	Analytics / Statistical model	Examines message traffic	Self Learning	Method
z/OS Health Checker	•Checks configurations •Programmatic, applies to IBM and ISV tools •Can escalate notifications	✓				Rules based to screen for conditions
z/OS PFA	•Trending analysis of z/OS system resources, and performance •Can invoke z/OS RTD		✓		✓	Early detection
z/OS RTD	•Real time diagnostics of specific z/OS system issues	✓		✓		After an incident
IBM zAware	•Pattern based message analysis •Self learning •Provides aid in diagnosing complex z/OS problems, including cross sysplex, problems that may or may not bring the system down		✓	✓	✓	Diagnosis Useful before or after an incident

- IBM zAware Uniquely analyzes messages in context to determine unusual behaviors
- IBM zAware Uniquely understands and tunes its baseline to compare against your current activity
- IBM zAware does not depend on other solutions, manual coding of rules, and is always enabled to watch your system

© 2012 IBM Corporation

IBM zAware delivers smarter message monitoring capabilities

- **Out-of-band** high speed analytics application delivered as an integrated firmware stack
- Analyzes system messages to provide a **near real-time view of your system**
- Uses **self learning** to recognize message patterns of your environment
- User-friendly web interface
- **Easy drill down** techniques help identify problematic messages and unusual patterns to help speed up diagnostic time
- Provides information, in **XML format**, that can feed other processes or tools
  - Tivoli Integrated Service Management intends to work with IBM zAware to provide alert and event notifications*

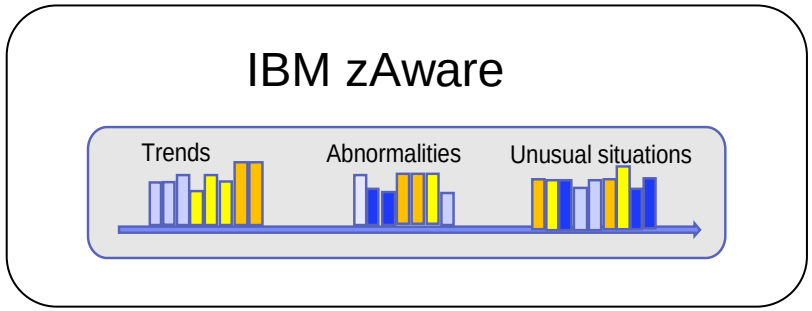

*All statements regarding IBM future direction and intent are subject to change or withdrawal without notice and represents goals and objectives only.

© 2012 IBM Corporation

How can IBM zAware Improve Problem Determination?

- **Identify messages indicating a possible z/OS incident is happening**
  - *Which image is behaving abnormally?*
    - Examines unique messages
    - High score generated by
      - unusual messages or message patterns
  - *When did this unusual behavior start?*
    - For a selected 10 minute interval either the current 10 minute interval or past intervals
    - **Which message ids** are unusual?
    - **How often** did the message occur?
    - **When** did the message start to occur?
  - *Were similar messages issued in the past?*
    - Similar characteristics, Same pattern?
- **After a change has been made**
  - *Are unusual messages being issued following changes ?*
    - New software levels (operating system, middleware , applications)
    - Updated system settings / system configurations
- **When diagnosing the cause of an intermittent problem**
  - *Are new unusual messages being issued in advance of the problem?*
  - Are more messages issued then expected?
  - Are messages issued out of normal pattern or context?

Vertical bar shows the number of unique messages in a 10 minute interval

Scoring of messages color coded from common (blue) to rare (orange)

Finds Anomalies that would be Hard to Detect

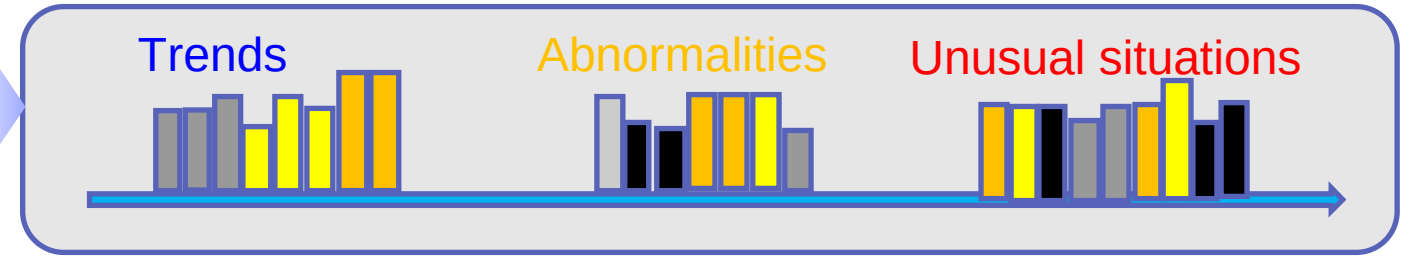

© 2012 IBM Corporation

Snapshots – Drill Down

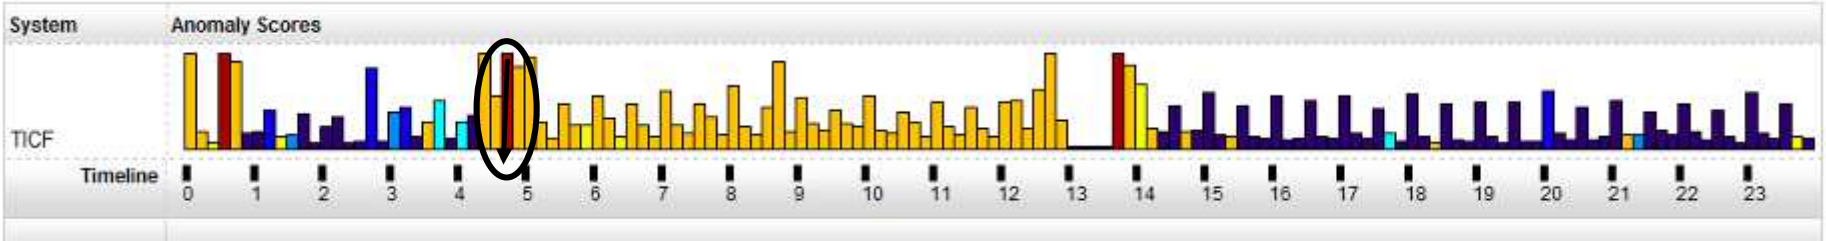

Interval Contribution Score	Appear Count	Cluster ID	Rarity Score ▼	Time Line	Component	Message ID	Message Example
0.894	17	-	101		ILR	ILR032I	PAGE DATA SET HAS BEEN USED BY ANOTHER SYSTEM: 042 DATA SET NAME - SYS1.TIPG00.COMMON VOLUME SERIAL - TIPG00 DEVICE NUMBER - 87FC SYSTEM NAME - TICF [...] DATA SET LAST UPDATED AT
0.043	1	-	101		ILR	ILR030A	PAGE DATA SET MAY BE IN USE: 032 DATA SET NAME - SYS1.TIPG00.PLPA VOLUME SERIAL - TIPG00 DEVICE NUMBER - 87FC SYSTEM NAME - TICF [...] DATA SET LAST UPDATED AT 09:37:31 ON
0.043	1	-	101		ILR	ILR031A	REPLY 'DENY' TO PREVENT ACCESS, 'CONTINUE' TO ALLOW USE OF SYS1.TIPG00.PLPA.
0.043	1	-	101		IEA	IEA380I	THIS SYSTEM IS NOW OPERATING IN STP TIMING MODE.
0.043	1	-	101		AOF	AOF604I	AUTOMATION PAUSED BY OPERATOR REQUEST
0.043	1	-	101		CNZZ	CNZZ033E	SPECIFIC MESSAGE THRESHOLD REACHED FOR IOS291I
0.043	1	-	101		IEAVEH	IEAVEH071E	HiperDispatch is expected to be enabled but it is disabled
0.291	6	-	101		EZZ	EZZ7871I	NO MATCHING INTERFACE STATEMENTS FOR 10.103.247.18 (MPC_TICF_TW PR)

© 2012 IBM Corporation

A Closer Look Inside IBM zAware

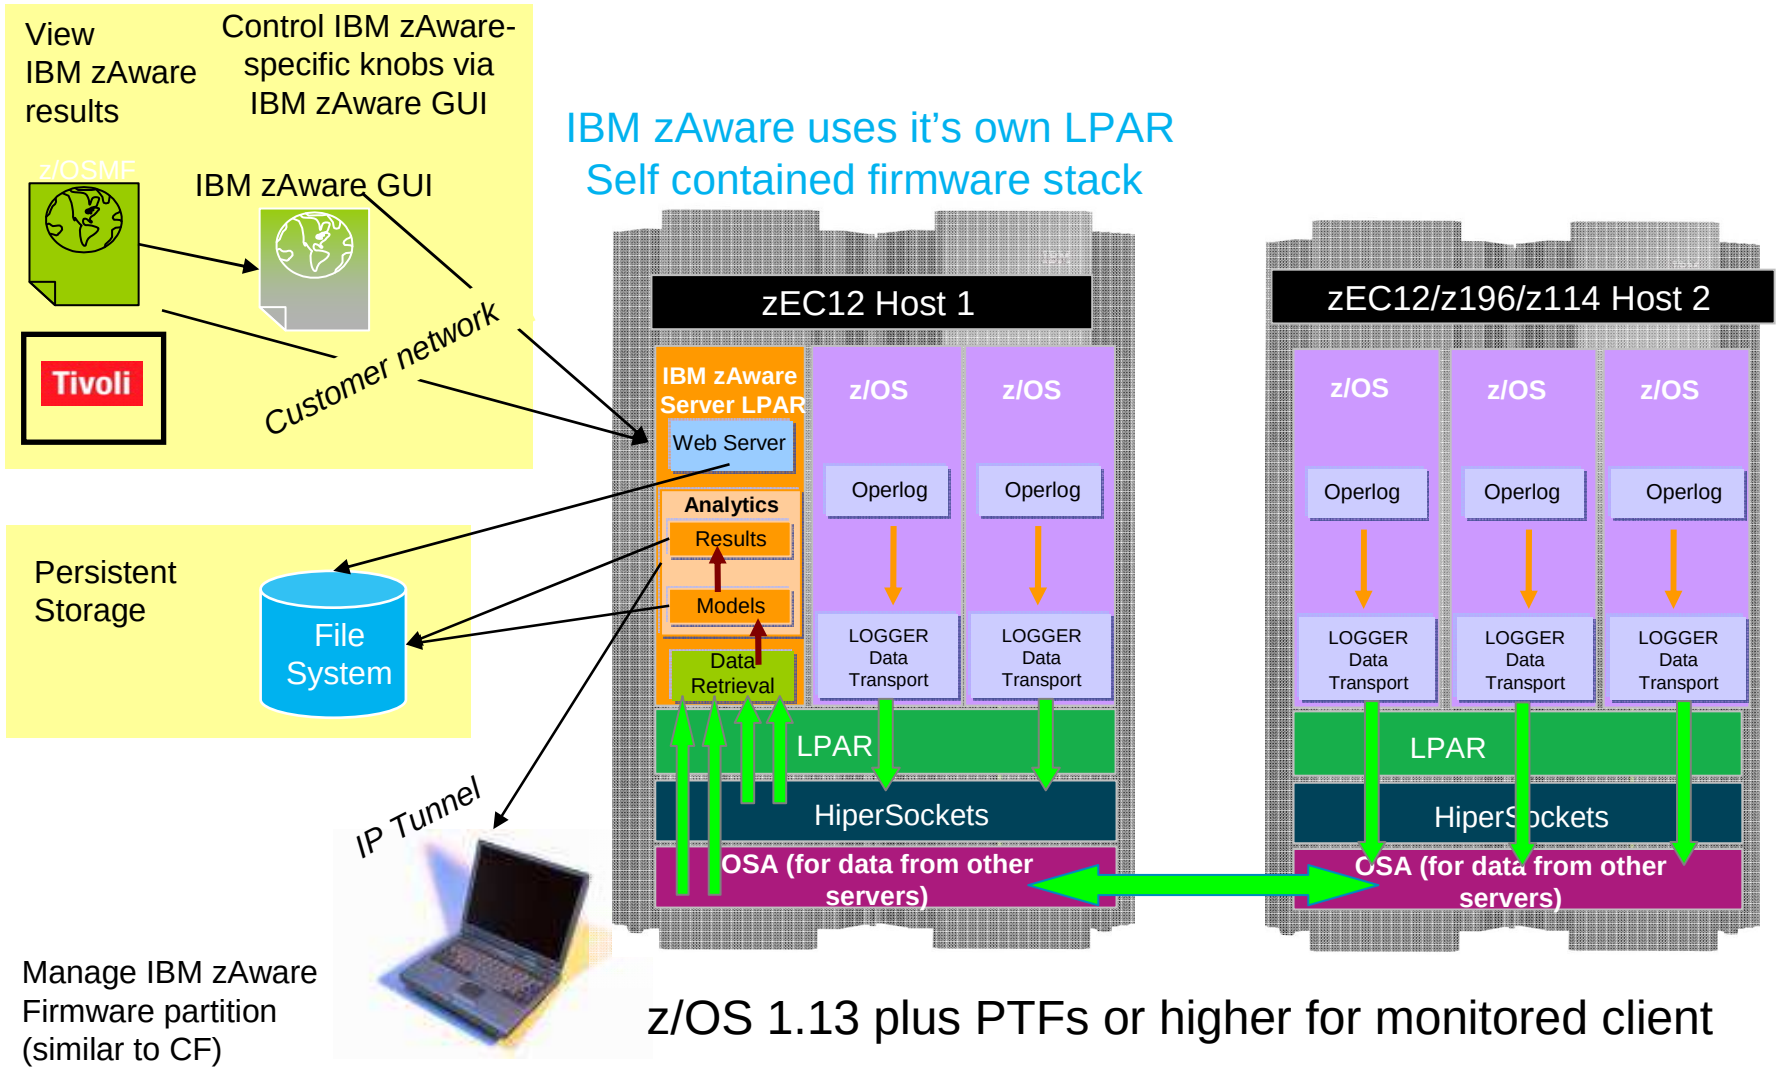

z/OS 1.13 plus PTFs or higher for monitored client

© 2012 IBM Corporation

IBM zAware Operating Requirements

- zEC12 to host IBM zAware Server
  - IBM zAware requires it's own LPAR and runs it's own self-contained firmware stack.
    - This will reduce the number of LPARs available for customer use
  - IBM zAware processor resources can be IFL or General Purpose CP
  - Memory and DASD resources are dependent on the number of monitored clients, amount of message traffic, length of time data retained
    - Memory - Min 6 GB + 256 MB. 6 GB for the first 6 z/OS Clients + 256 MB required for additional z/OS Clients above 6.
    - DASD ~ 500 GB (ECKD)
    - IBM zAware uses Logical Volume Manager (LVM) to aggregate multiple physical devices into a single logical device
  - Network: HiperSockets or OSA ports – for both gathering of instrumentation data, and outbound alerting/communications
    - Need dedicated IP address for partition
- IBM zAware Monitored Clients
  - IBM zAware monitored clients can be on any System z Server running z/OS 1.13 + PTFs
  - IBM zEnterprise 196 (z196), IBM zEnterprise 114 (z114), etc., and can share log files via IP network with IBM zAware server
- 90 days historical syslog or OPERLOG data to initially prime IBM zAware

© 2012 IBM Corporation

IBM zAware Features and 'Connections'

- Definitions:
  - The **IBM zAware host system** is the zEC12 that hosts the IBM zAware partition. In most cases, the host server will also have partitions on it that are being monitored. There may be multiple IBM zAware host partitions on one zEC12, but there will only be one IBM zAware FC 0011 feature (no additional charge for multiple host partitions).
  - An **IBM zAware monitored client** is a z/OS partition that sends OPERLOG files for processing to an IBM zAware partition. The server that has partitions that are IBM zAware monitored clients, must be running z/OS 1.13 with required exploitation software. There may be multiple z/OS partitions (monitored clients) on the server.
  - The **IBM zAware environment** is the collection of the IBM zAware host system and the IBM zAware monitored clients that are sending information to the IBM zAware host system.
  - An **IBM zAware connection** is used by eConfig to represent a set of 10 CPs associated with servers that are either the IBM zAware host system or the IBM zAware monitored clients. The IBM zAware environment may be less than total 10 CPs and in that situation will be rounded up to 10. Caution: do not confuse the term 'connection' in this example with a network connection.
  - A **Disaster Recovery (DR) IBM zAware server** is a zEC12 with no-charge firmware to run IBM zAware in a disaster situation.

© 2012 IBM Corporation

IBM zAware Ordering

- Features for ordering IBM zAware will be located on the Processor usage panel of eConfig.
  - FC 0011 represents that IBM zAware is installed on the 'host' CPC.
  - FC 0101 represents the quantity of IBM zAware host/client connections
  - FC 0102 represents the quantify of IBM zAware connections on a DR server

  Note that features FC 0101 and FC 0102 are mutually exclusive – if you have one then you can't have the other. Also in most cases the the number of FC 0102 features on DR should match the number of FC 0101 features on the IBM zAware host server.
- IBM zAware connections are orderable based on the quantity of CPs on the host machine plus the quantity of CPs in the monitored client machine(s) up to the High Water Mark of the CPC
- Minimum quantity available would be the quantity of CPs on the host (ordering) machine.
- A Disaster Recovery option is available. It represents that IBM zAware is installed on a Disaster Recovery machine. In this case FC 0102 will be assigned to represent the quantity of connections. FC 0102 is a zero-priced feature.
- In addition to the IBM zAware features, additional hardware may be required for the IBM zAware partition e.g CP or IFL capacity, memory, storage and OSA prots for network connectivity

© 2012 IBM Corporation

All zEC12 site – CPC failure

- Single active IBM zAware
- IBM zAware feature on both zEC12 CPCs only one active
- On CPC failure automation starts alternative IBM zAware

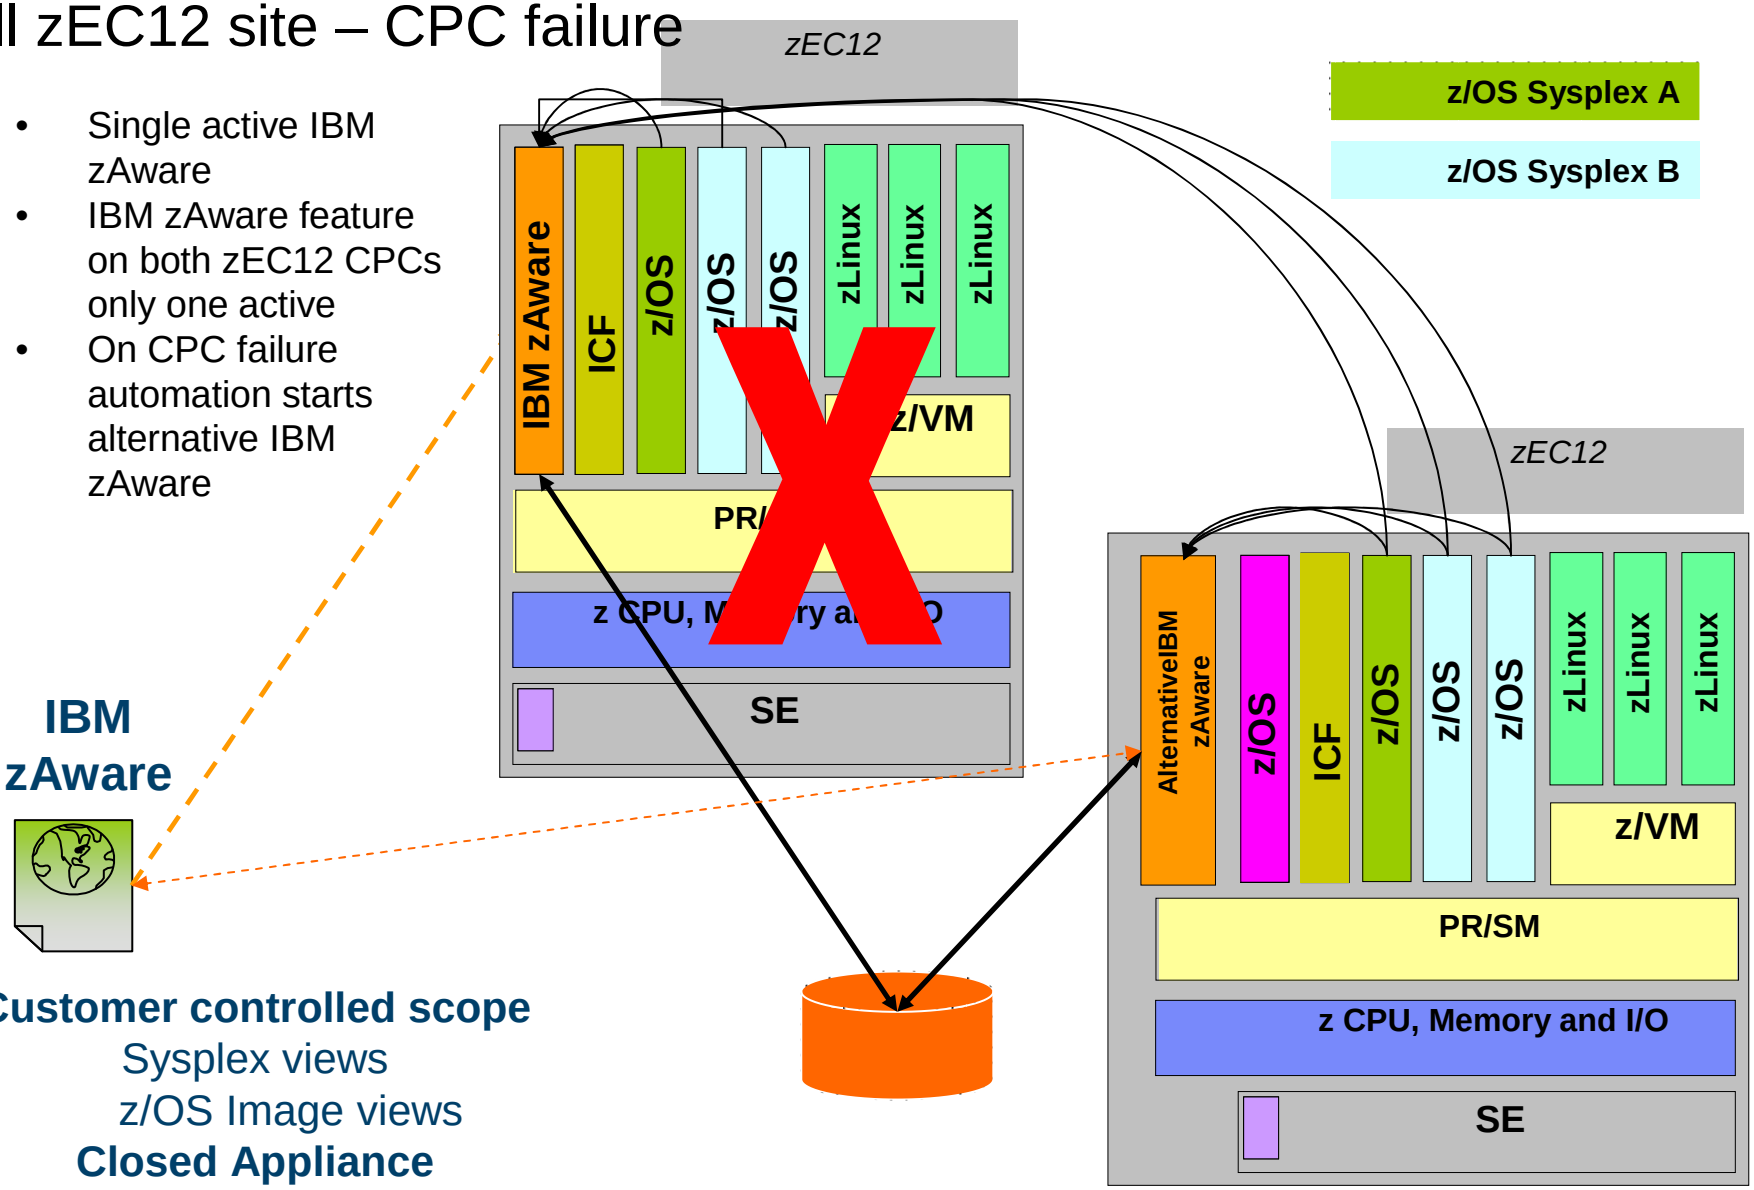

© 2012 IBM Corporation

Mixed site – CU failure

- Single active IBM zAware
- IBM zAware feature on both zEC12 CPCs only one active
- On CPC failure automation starts alternative IBM zAware

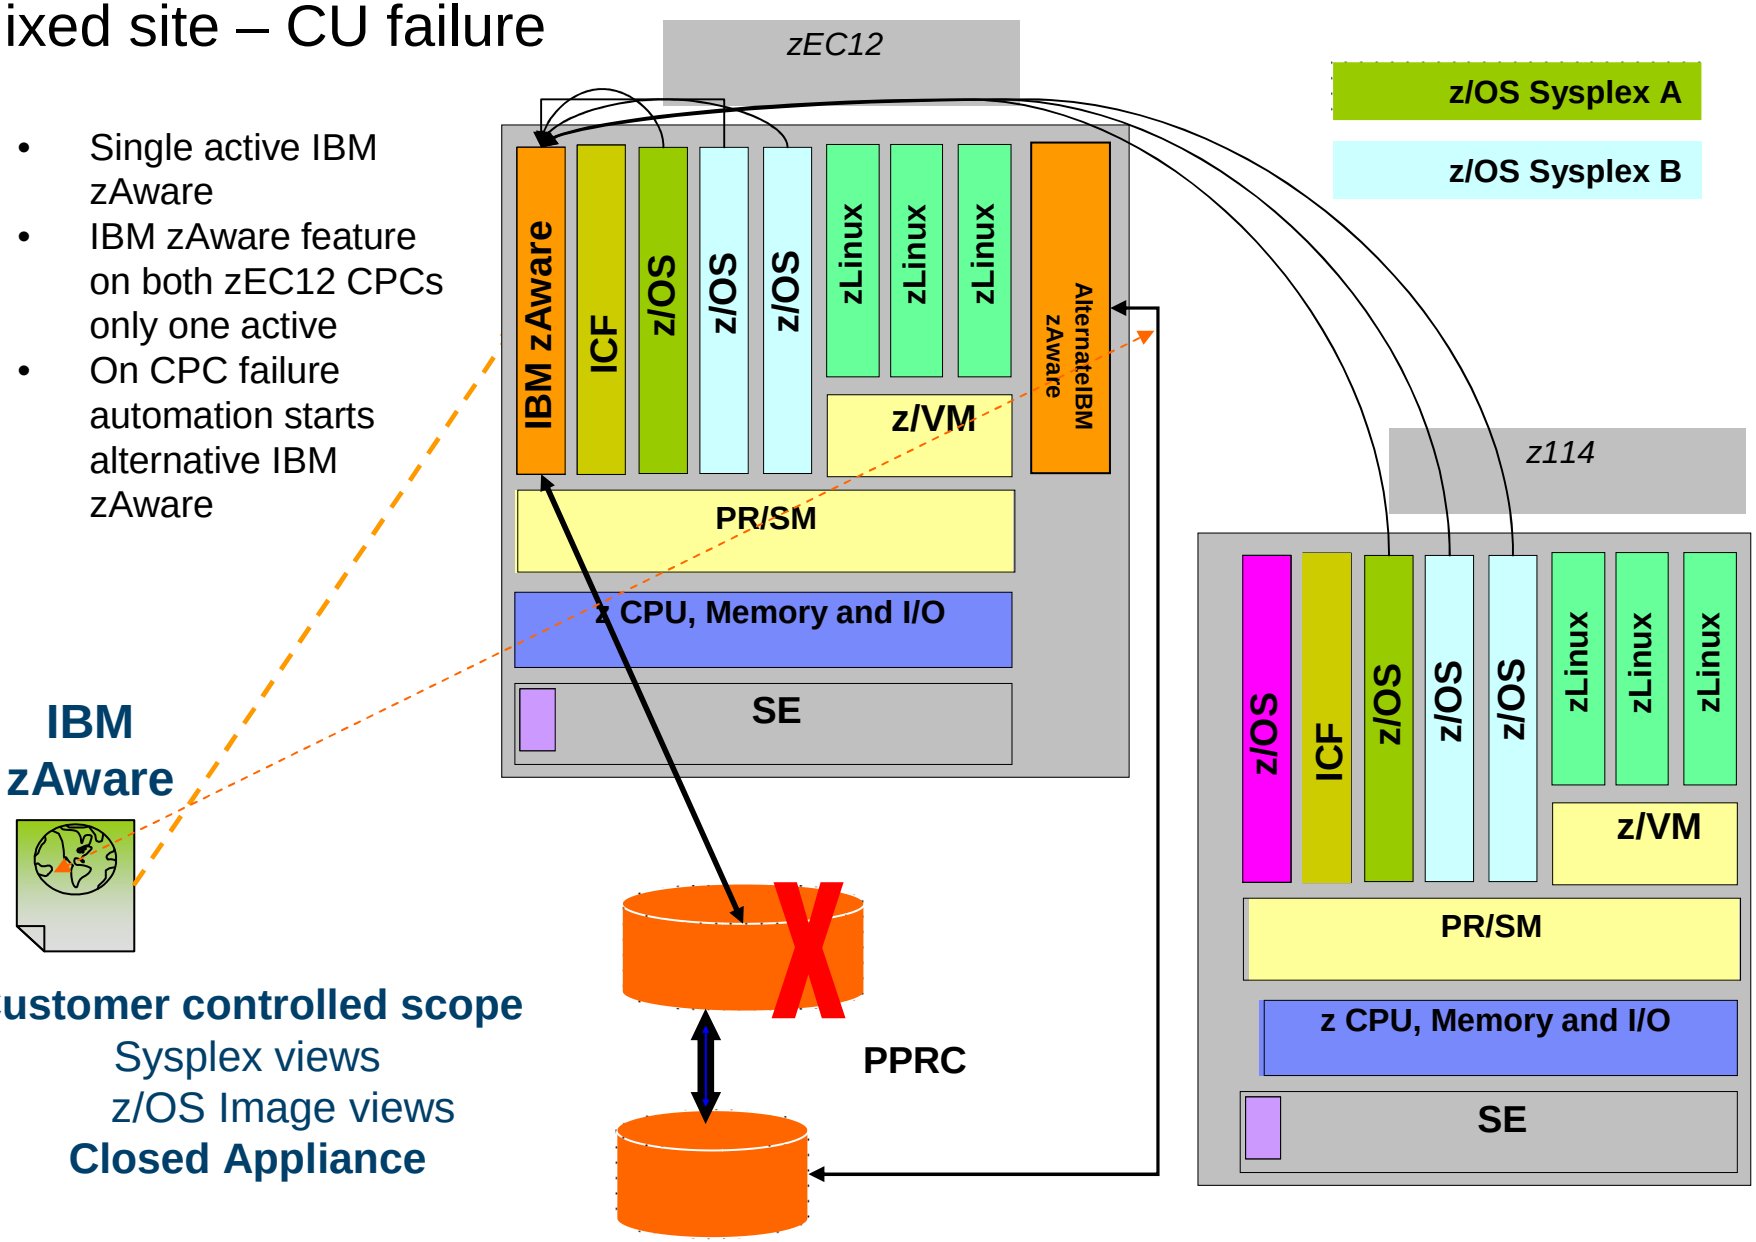

© 2012 IBM Corporation

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


PCIe zEC12/z196/z114

8 GBps


InfiniBand z10/z196/z114 & zEC12 for carry forward

6 GBps


STI z9

2.7 GBps


STI z990/z890

2 GBps


STI z900/z800

1 GBps

STI: Self-Timed Interconnect

© 2012 IBM Corporation

zEC12 I/O infrastructure

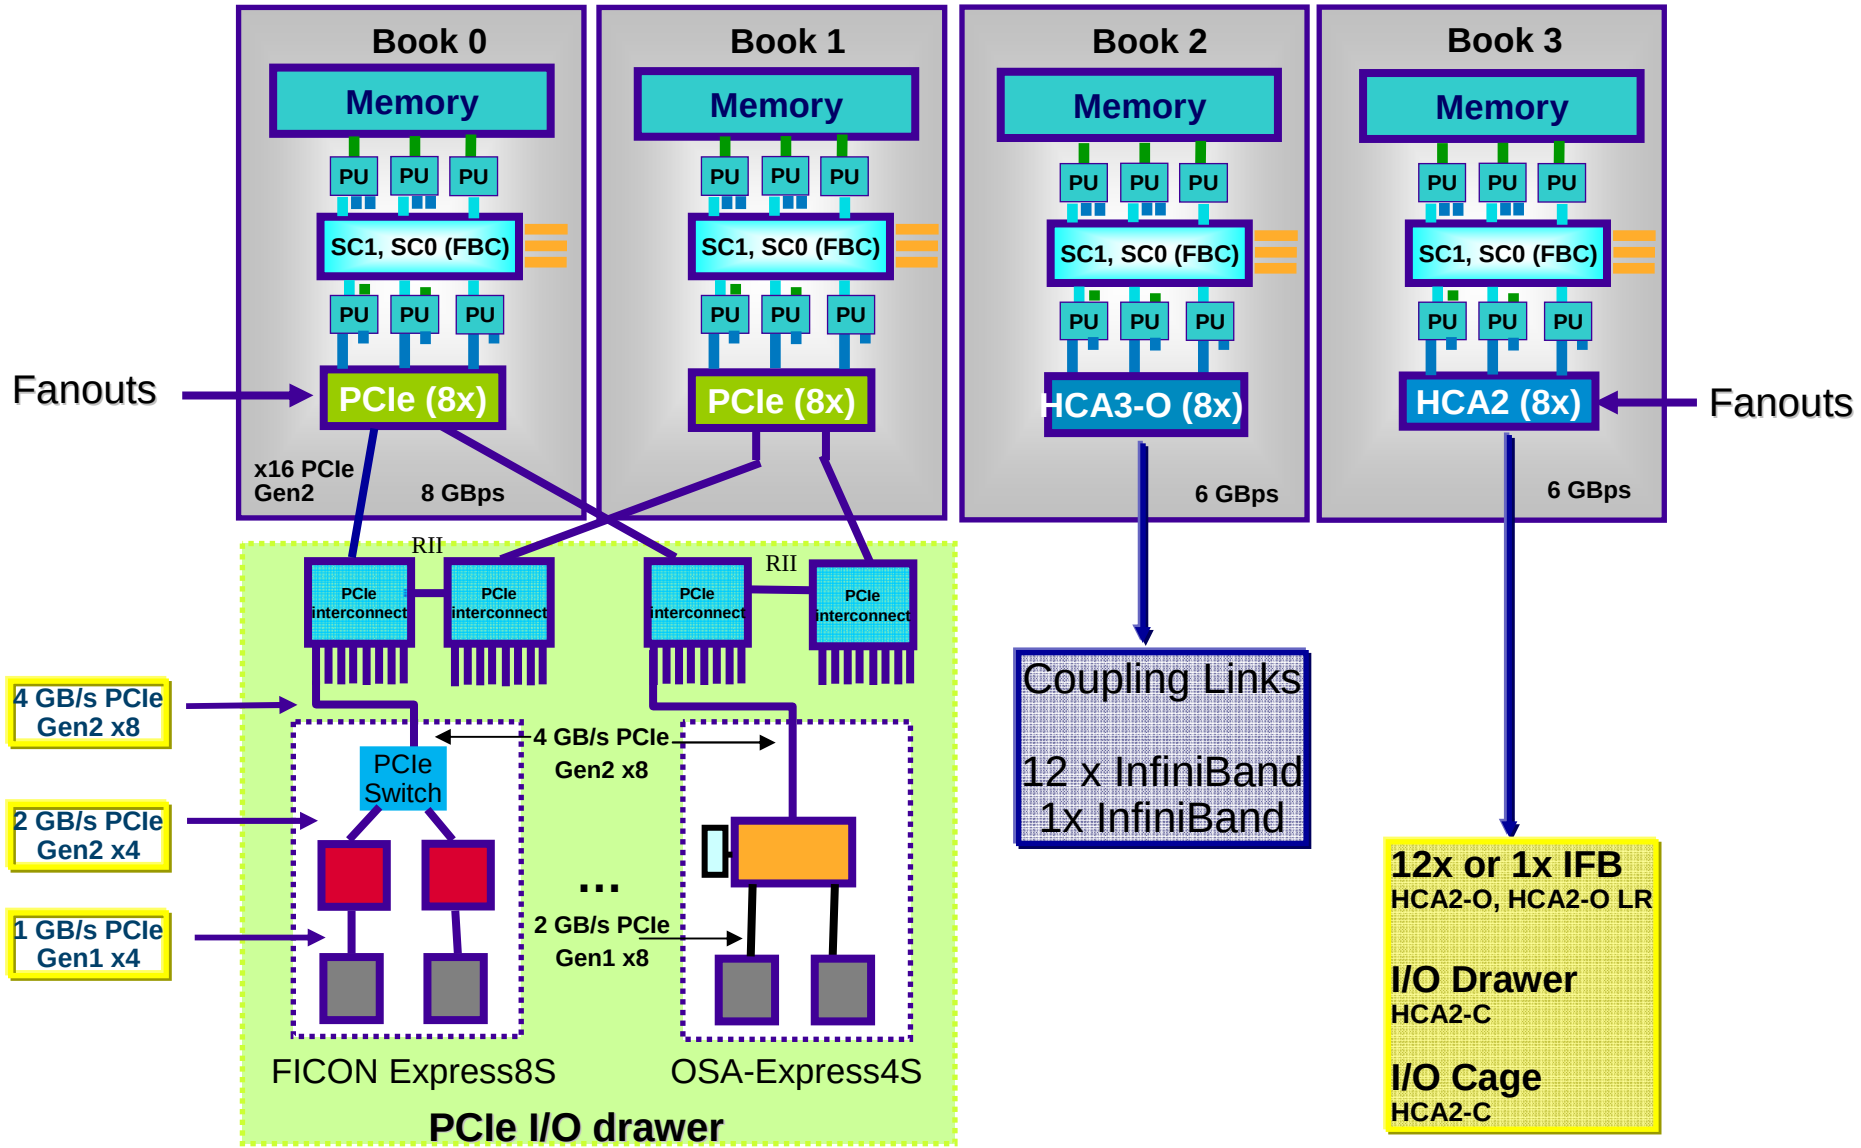

© 2012 IBM Corporation

zEC12 GA1 Features Supported – I/O Cage, I/O Drawer, PCIe I/O Drawer

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
    - 1000BASE-T

PCIe I/O drawer

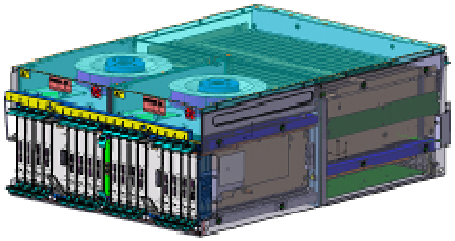

32 I/O slots

- Features – I/O cage and I/O drawer (Carry Forward ONLY on System MES – Maximum of 44 features. No ADDs)
  - **Not Supported: ESCON, older FICON, FICON Express4 LX 4 km, OSA-Express2, PSC**
  - Crypto Express3
  - FICON Express8
  - FICON Express4 10 km LX and SX
  - ISC-3
  - OSA-Express3
    - 10 GbE, GbE, 1000BASE-T

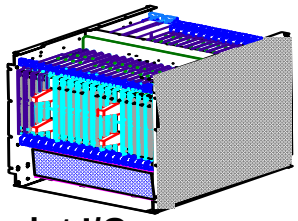

28 slot I/O cage

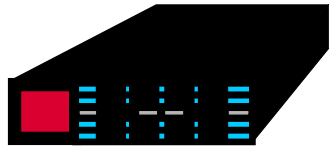

8 slot I/O drawer

© 2012 IBM Corporation

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

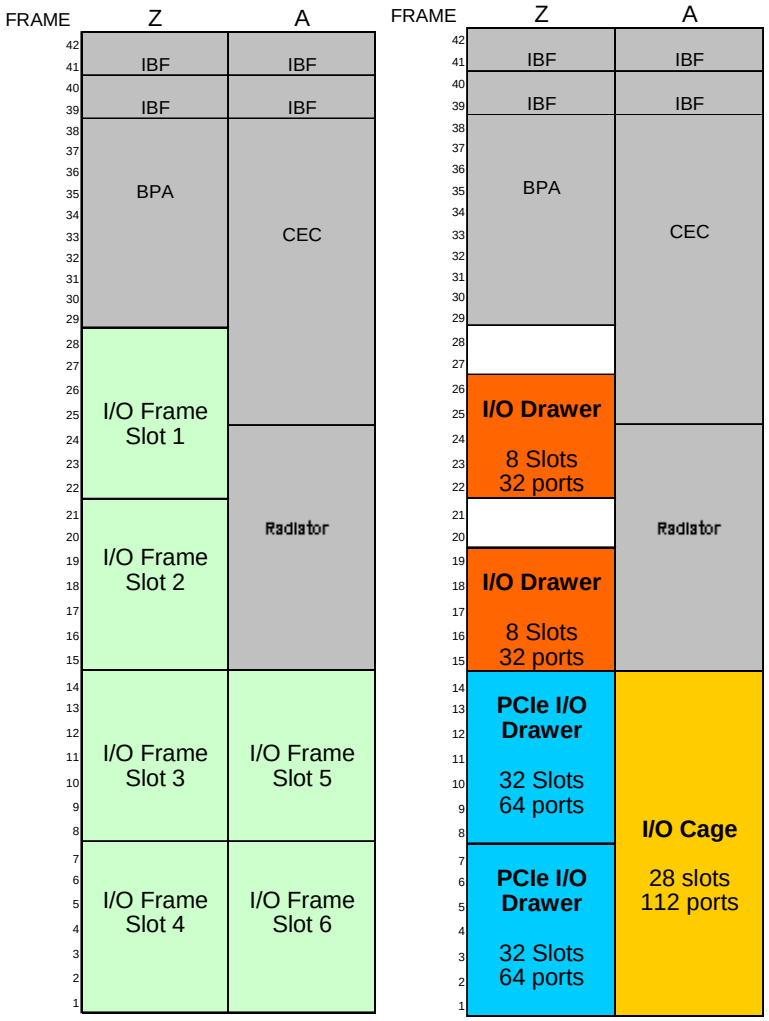

- An I/O frame slot is a physical location in the A or Z frame for an I/O cage, I/O drawer or PCIe I/O drawer to be inserted = 7u
- PCIe I/O drawer uses 1 I/O frame slot = 7u
  - 32 two port I/O slots = 64 ports
  - 5 drawers maximum
- I/O cage uses 2 I/O frame slots = 14u
  - 28 four port I/O slots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

© 2012 IBM Corporation

ESCON Statement of Direction - February 15, 2011

- The IBM zEnterprise 196 (z196) will be the last high-end server to support ESCON channels: IBM plans not to offer ESCON channels as an orderable feature on high-end System z servers which follow the z196 (machine type 2817). In addition, ESCON channels cannot be carried forward on an upgrade to such a follow-on server.
- This plan applies to channel path identifier (CHPID) types CNC, CTC, CVC, and CBY and to feature code numbers 2323 and 2324. System z customers should continue migrating from ESCON to FICON. Alternate solutions are available for connectivity to ESCON devices.
- IBM Global Technology Services offers an ESCON to FICON Migration solution, This offering should help customers to simplify and manage a single physical and operational environment.
- Notes:
  - This Statement of Direction superseded the previous ESCON SOD in Announcement letter 110-170 of July 22, 2010. It also confirms the SOD in Announcement letter 109-230 of April 28, 2009 that "ESCON Channels will be phased out."

  - No exceptions or RPQs. zEC12 does NOT have any code for ESCON

© 2012 IBM Corporation

FICON Convertor Support

IBM

OPTICA

Partnership

Old → IBM 9034 (ESCON to Parallel) → Optica 34600 FXBT (ESCON to Parallel) → Optica PRISM (FICON to ESCON) → **Current**

PRIZM is available from IBM GTS Site & Facilities as part of the EFM Service (ESCON to FICON Migration - offering # 6948-97D)

- http://www.**ibm.com/**services/us/index.wss/itservice/igs/a1026000?cm_re=masthead-_-itservices-_-site
- The order process for PRIZM is the same as it is for IBM cabling systems

© 2012 IBM Corporation

Channel Sub-System Enhancements for zEC12

- The zEC12 channel subsystem has been enhanced with new channel path selection algorithms designed to provide improved throughput and I/O service times when abnormal conditions occur.
- Abnormal conditions include the following:
  - Multi-system work load spikes
  - Multi-system resource contention in the SAN or at the CU ports
  - SAN congestion
  - Destination port congestion
  - Firmware failures in the SAN, channel extenders, DWDMs, control units
  - Hardware failures - link speeds did not initialize correctly
  - Cabling Errors
  - Dynamic changes in fabric routes
- When conditions occur that causes an imbalance in performance (I/O latency/throughput) the channel subsystem will bias the path selection away from poorer performing paths toward the well performing paths.
- This is accomplished by exploiting the in-band I/O instrumentation and metrics of System z FICON and zHPF protocols and new intelligent algorithms in the channel subsystem to exploit this information.

© 2012 IBM Corporation

System z Crypto History

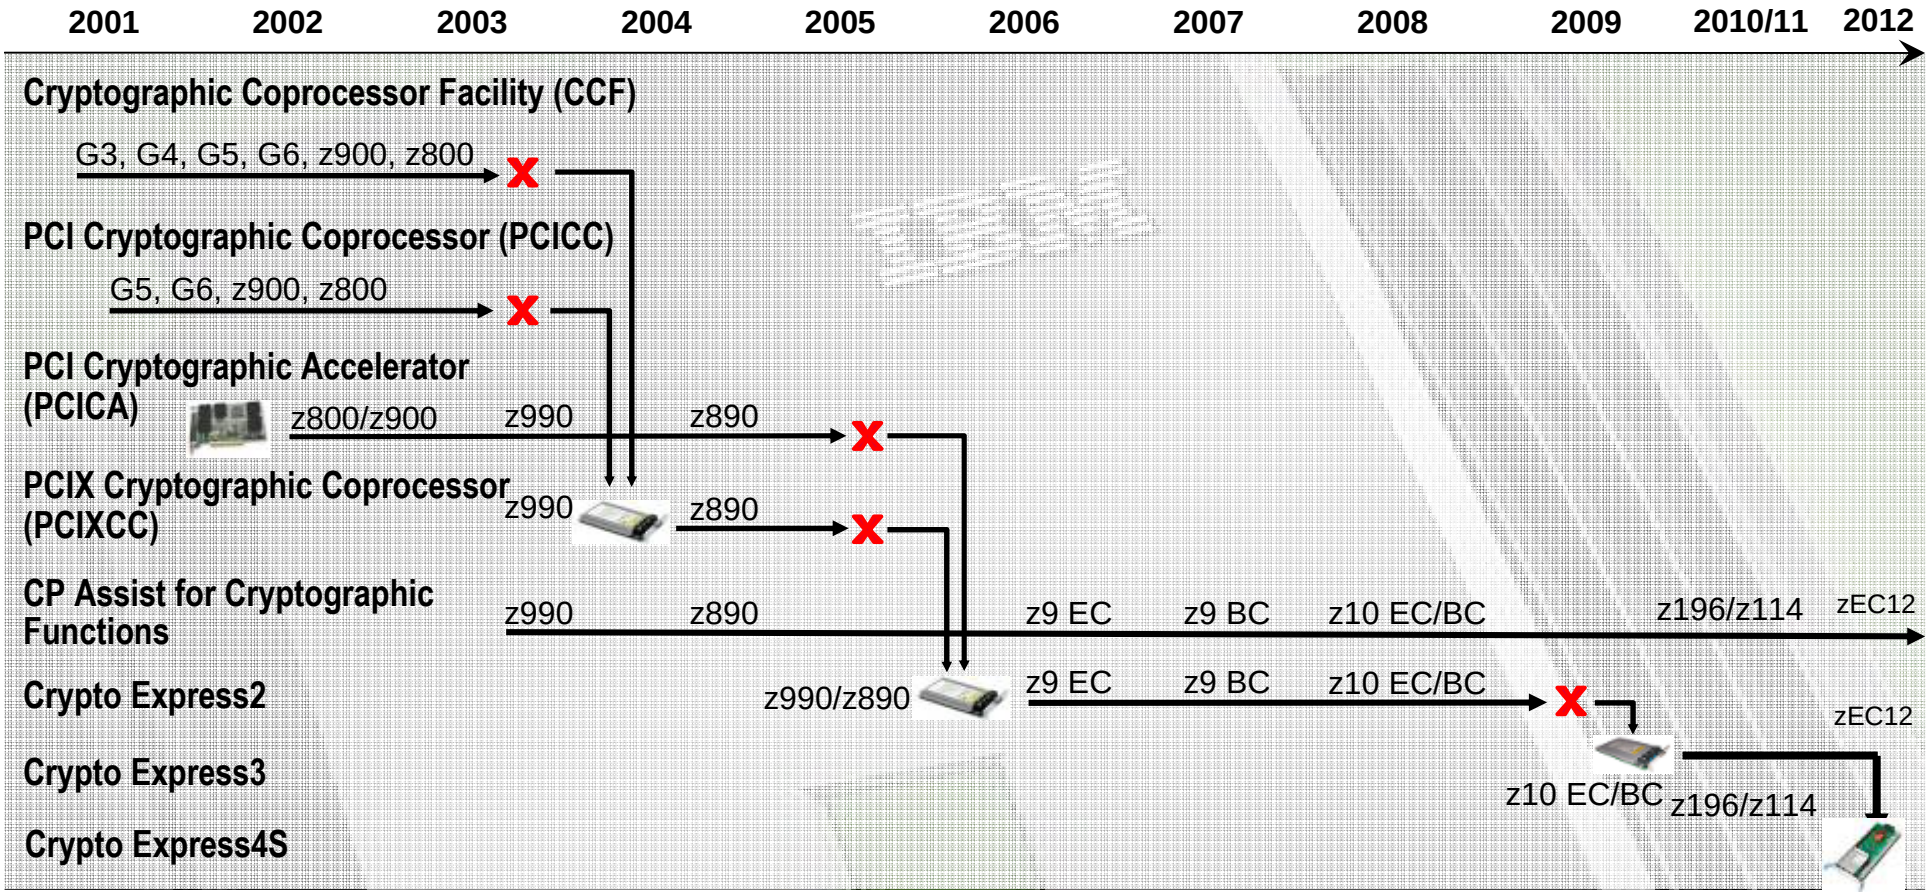

- Cryptographic Coprocessor Facility – Supports “Secure key” cryptographic processing
- PCICC Feature – Supports “Secure key” cryptographic processing
- PCICA Feature – Supports “Clear key” SSL acceleration
- PCIXCC Feature – Supports “Secure key” cryptographic processing
- CP Assist for Cryptographic 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

© 2012 IBM Corporation

zEC12 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 5+ 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

© 2012 IBM Corporation

IBM Enterprise PKCS #11 LIC

- Industry standardized set of services that adhere to the PKCS specifications
  - Based on PKCS #11 specification v2.20 and more recent amendments
- Supports secure PKCS #11 keys
  - Keys that never leave the secure boundary of the coprocessor unencrypted
- Designed to meet Common Criteria (EAL 5+) standards and FIPS 140-2 Level 4 requirements
  - Targets the public sector where industry standard services are required
  - Certifications tailored to meet requirements of this market place
- Conforms to the Qualified Digital Signature (QDS) Technical Standards
  - Becoming a mandate by the European Union
  - High quality electronic signatures
    - Trusted to the same extent as hand written signatures
  - Uses: Smart passports, national id cards, …
- EP11 is intended to broaden the appeal of System z and hardware crypto
- Supported on Crypto Express4S only

© 2012 IBM Corporation

Crypto Express4S

- One PCIe adapter per feature
  - Initial order – two features
- FIPS 140-2 Level 4
- Installed in the PCIe I/O drawer
- Up to 16 features per server
- Prerequisite: CPACF (FC 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 (FC 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 (FC 0841) for management of the Crypto Express4S when defined as an EP11 coprocessor
  - Supported on Crypto Express4S only

© 2012 IBM Corporation

System z overall RAS Strategy
…..Continuing our RAS focus helps avoid outages

Sources of Outages
Pre z9
-Hrs/Year/Syst-

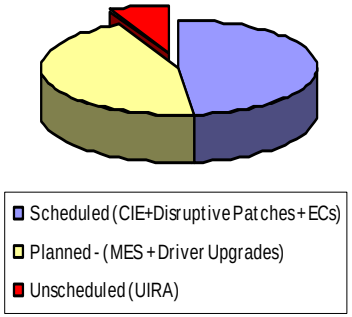

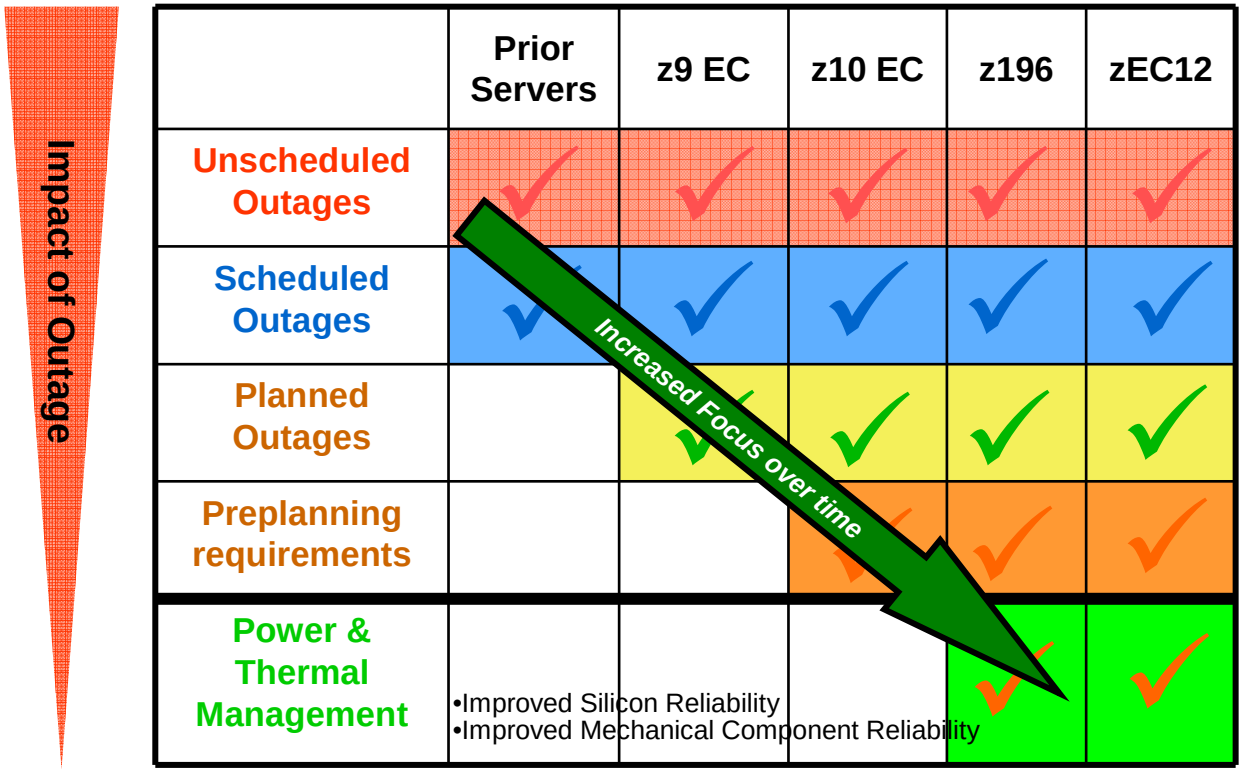

© 2012 IBM Corporation

Simplify and reduce cost with IBM zEnterprise

- An Integrated system of multiple architectures for optimizing the deployment of multi-tier workloads
- Creating a single point of control for management and administration to reduce operational overheads by up to 80%, including:
  - Power and Facilities
  - Labor
  - Software Licenses

Lower cost of acquisition

Reduce cost of ownership

- Based on IBM analysis of a large Financial Services company Datacenter. See details on ibm.com/systems/zenterprise/
- Deployment configurations based on IBM studies and will vary based on workload characteristics. Price calculations based on publicly available US list prices, prices will vary by country.

© 2012 IBM Corporation

zEnterprise
Client Optimized Systems

- Multi-Architecture System for z/OS, AIX, Linux and Windows
- Centrally managed through the Unified Resource Manager
- Best fit when data or applications exist on System z and clients desire z governance

PureSystems
Integrated Expert Systems

- Multi-Architecture system for AIX, i/OS, Linux and Windows
- Centrally managed resources though PureSystems managers – Flex System Manager (FSM)
- Best fit when data and applications run on a combination of POWER and System x architecture

Today: Clients can also attach IBM zEnterprise and IBM PureSystems (via Ethernet) to gain benefits of simplified management and lower IT infrastructure costs for all workloads.
IBM's Tivoli service management platform allows for integration for improving delivery of business services.

In future: Tighter integration of these two systems. Today's investment in either will gain value over time.

© 2012 IBM Corporation

© 2012 IBM Corporation