IBM System z10™ Enterprise Class Overview

Session 2832, February 27, 2008

Harv Emery, IBM, Washington Systems Center
Team Leader, Americas System z™ Hardware ATS

SHARE 110 in Orlando

IBM Systems

Trademarks

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

AIX*
APPN*
Cell Broadband Engine
DB2*
DB2 Connect
DirMaint
DRDA*
Distributed Relational Database Architecture
e-business logo*
ECKD
Enterprise Storage Server*
ESCON*
FICON*
GDPS*
Geographically Dispersed Parallel Sysplex
HiperSockets
HyperSwap
IBM*
eServer
IBM logo*
IMS
InfoPrint*
Language Environment*
MQSeries*
Multiprise*
NetView*
On demand business logo
OS/390*
Parallel Sysplex*
POWER6
PR/SM
Processor Resource/Systems Manager
RACF*
Resource Link
RMF
S/390*
Sysplex Timer*
System z
System z9
System z10
TotalStorage*
Virtualization Engine
VSE/ESA
VTAM*
WebSphere*
z/Architecture
z/OS*
z/VM*
z/VSE
zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

InfiniBand® is a registered trademark of the InfiniBand Trade Association (IBTA).
Intel is a trademark of Intel Corporation in the United States, other countries, or both.
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.
SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

* 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.

System z10 EC New Functions and Features

Five hardware models
Faster Processor Unit (PU)
Up to 64 customer PUs
36 CP Subcapacity Settings
Star Book Interconnect
Up to 1,520 GB memory
Separate fixed HSA standard
Large Page (1 MB)
HiperDispatch
Enhanced CPACF SHA 512, AES 192 and 256-bit keys
Hardware Decimal Floating Point
New Capacity on Demand architecture and enhancements
Capacity Provisioning

SOD: PSIFB for z9 EC & z9 BC for non-dedicated CF Models*

6.0 GBps InfiniBand HCA to I/O interconnect
FICON Enhancements
SCSI IPL included in Base LIC
OSA-Express3 10 GbE **(2Q08)***
HiperSockets enhancements
InfiniBand Coupling Links **(2Q08)***
STP using InfiniBand **(2Q08)***
Standard ETR Attachment
FICON LX Fiber Quick Connect
Power Monitoring support
Scheduled Outage Reduction
Improved RAS

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

IBM System z10 EC Key Dates

- **IBM System z10 Announce – February 26, 2008**
  - First Day Orders
  - Resource Link™ support available
  - Capacity Planning Tools (zPCR, zTPM, zCP3000)
  - SAPR Guide (SA06-016-00) and SA Confirmation Checklist available
- **Availability – February 26, 2008**
  - z10 EC all Models
  - Upgrades from z990, z9 EC to z10 EC
- **Availability – May 26, 2008**
  - Model upgrades within z10 EC
  - Feature Upgrades within the z10 EC – May 26, 2008
- **Planned Availability* – 2Q 2008**
  - OSA Express3 10 GbE LR – the first of a new OSA generation
  - InfiniBand Coupling Links for any z10 EC and ICF-only z9 EC and BC machines
- **New ITSO Redbooks (Draft versions)**
  - z10 EC Technical Introduction, SG24-7515 - February 26, 2008
  - z10 EC Technical Guide, SG24-7516 - February 26, 2008
  - z10 EC Capacity on Demand, SG24-7504 - March, 2008
  - Getting Started with InfiniBand on z10 EC and System z9, SG24-7539 – May, 2008

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

System z10 EC Sessions at SHARE in Orlando

Session	Time and Location	Title	Speaker
2832	Wed, 8:00, Coronado J	z10 EC Overview	Harv Emery
2836	Wed, 9:30, Coronado J	z10 EC I/O Subsystem	Connie Beuselinck
2833	Wed, 11:00, Coronado J	z10 EC Capacity On Demand Overview	Greg Hutchison
8190 8191	Wed, 1:30, Fiesta 6 Wed, 3:00, Fiesta 6	Assembler University: Additions to z/Architecture, Parts 1 & 2	Dan Greiner
2852	Thr, 9:30, Coronado J	z/OS System Programmer's Guide to Migration to z10 EC	Greg Daynes
2860	Thr, 11:00, Coronado J	z10 EC Installation Planning and Migration Considerations	John Hughes
2532	Thr, 3:00, Durango 2	To MIPS or Not to MIPS, That Is the Question! (z10 EC examples)	Gary King

z10 EC Overview

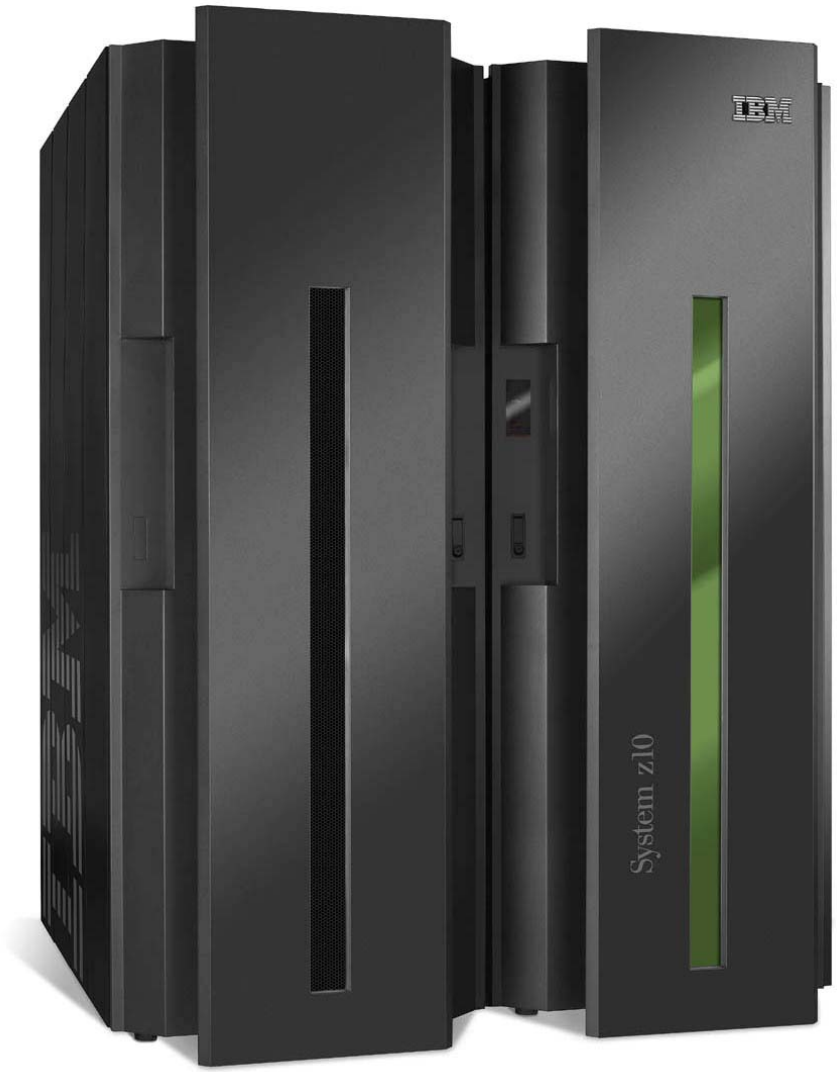

- **Machine Type**
  - 2097
- **5 Models**
  - E12, E26, E40, E56 and E64
- **Processor Units (PUs)**
  - 17 (17 and 20 for Model E64) PUs per book
  - Up to 11 SAPs per system, standard
  - 2 spares designated per system
  - Depending on the H/W model - up to 12, 26, 40, 56 or 64 PUs available for characterization
    - Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z Application Assist Processors (zAAPs), System z9 Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs)
- **Memory**
  - System Minimum of 16 GB
  - Up to 384 GB per book
  - Up to 1.5 TB GB for System
    - Fixed HSA, standard
    - 16/32/48/64 GB increments
- **I/O**
  - Up to 48 I/O Interconnects per System @ 6 GBps each
  - Up to 4 Logical Channel Subsystems (LCSSs)
- **ETR attachment standard**

IBM System z10 EC Upgrades

z10 EC System Upgrades

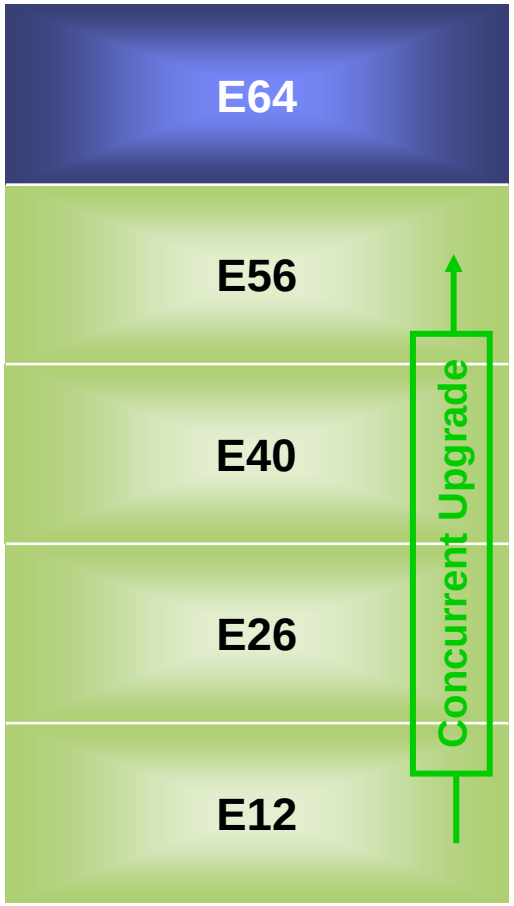

- **z10 EC to higher z10 EC model**
  - Concurrent upgrade of z10 EC Models E26, E40 and E56. Upgrade to E64 is disruptive
  - When upgrading to z10 EC E64, unlike the z9 EC, the first Book is retained
- **Any z9 EC to any z10 EC**
- **Any z990 to any z10 EC**

z10 EC – Under the covers (Model E56 or E64)

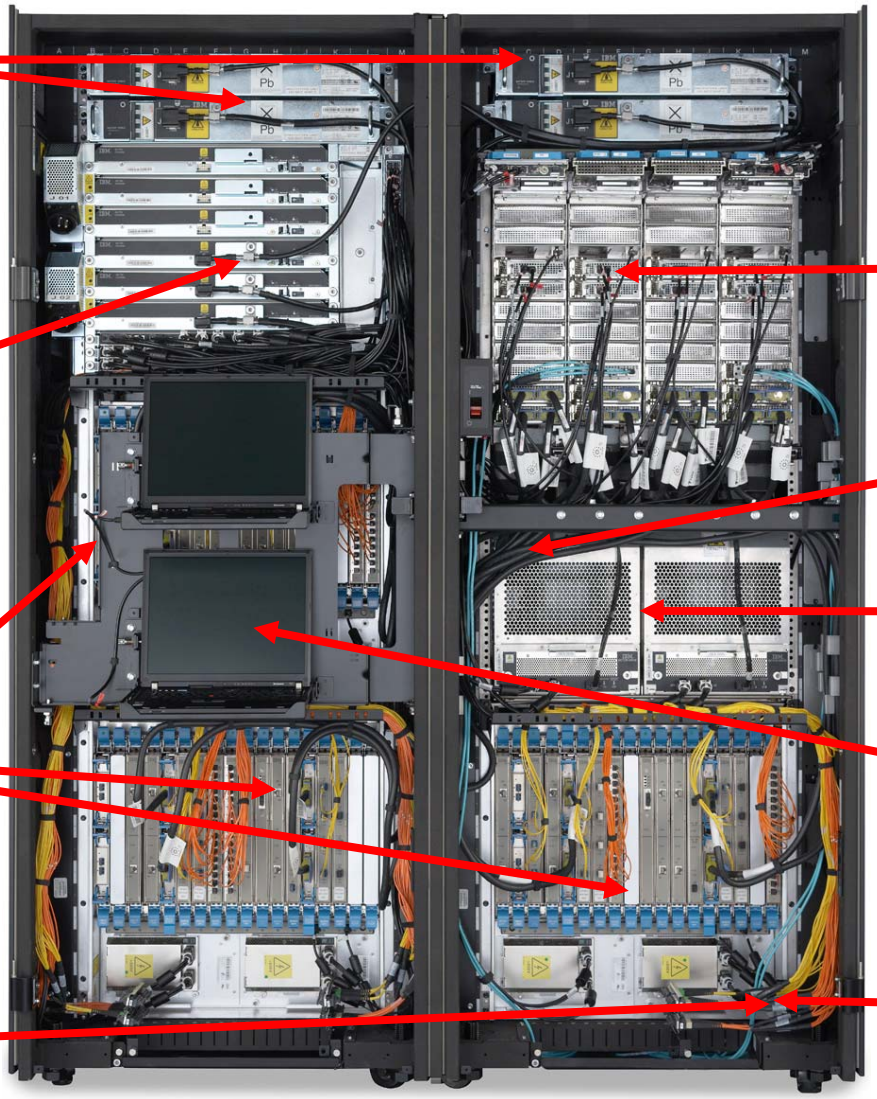

Internal Batteries (optional)

Power Supplies

3x I/O cages

Fiber Quick Connect (FQC) Feature (optional)

Processor Books, Memory, MBA and HCA cards

InfiniBand I/O Interconnects

Cooling Units

2 x Support Elements

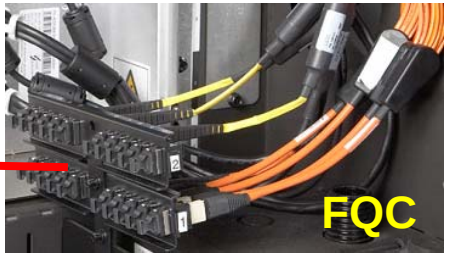

IBM System z10 EC Processor Architecture

z10 EC Chip Relationship to POWER6™

- **Siblings, not identical twins**
- **Share lots of DNA**
  - IBM 65nm Silicon-On-Insulator **(**SOI) technology
  - Design building blocks:
    - Latches, SRAMs, regfiles, dataflow elements
  - Large portions of Fixed Point Unit (FXU), Binary Floating-point Unit. (BFU), Hardware Decimal Floating-point Unit (HDFU), Memory Controller (MC), I/O Bus Controller (GX)
  - Core pipeline design style
    - High-frequency, low-latency, mostly-in-order
  - Many designers and engineers
- **Different personalities**
  - Very different Instruction Set Architectures (ISAs)
    - very different cores
  - Cache hierarchy and coherency model
  - SMP topology and protocol
  - Chip organization
  - IBM z Chip optimized for Enterprise Data Serving Hub

Enterprise Quad Core z10 Processor Chip

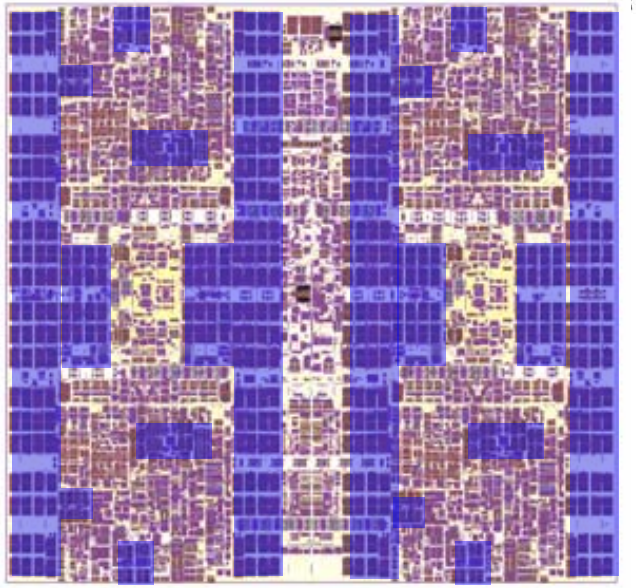

POWER6 Dual Core Chip

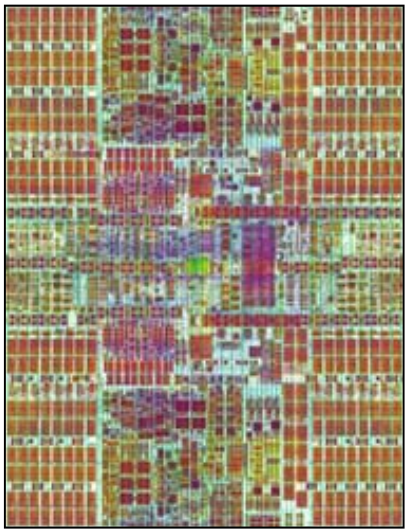

z10 EC Architecture

- **Continues line of upward-compatible mainframe processors**
- **Rich CISC Instruction Set Architecture (ISA)**
  - 894 instructions (668 implemented entirely in hardware)
  - 24, 31, and 64-bit addressing modes
  - Multiple address spaces robust inter-process security
  - Multiple arithmetic formats
  - Industry-leading virtualization support
    - High-performance logical partitioning via PR/SM
    - Fine-grained virtualization via z/VM scales to 1000s of images
  - Precise, model-independent definition of hardware / software interface
- **Architectural extensions for IBM z10 EC**
  - 50+ instructions added to improve compiled code efficiency
  - Enablement for software/hardware cache optimization
  - Support for 1MB page frames
  - Full hardware support for Hardware Decimal Floating-point Unit (HDFU)

z10 EC Multi-Chip Module (MCM)

- **96mm x 96mm MCM**
  - 103 Glass Ceramic layers
  - 7 chip sites
  - 7356 LGA connections
  - 17 and 20 way MCMs

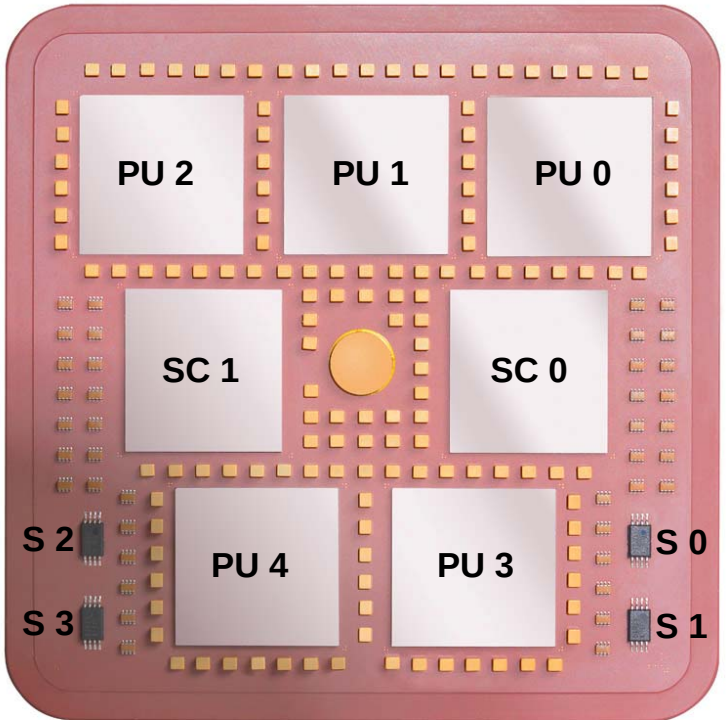

- **CMOS 11s chip Technology**
  - PU, SC, S chips, 65 nm
  - 5 PU chips/MCM – Each up to 4 PU cores
    - One memory control (MC) per PU chip
    - 21.97 mm x 21.17 mm
    - 994 million transistors/PU chip
    - L1 cache/PU core
      - 64 KB I-cache
      - 128 KB D-cache
    - L1.5 cache/PU core
      - 3 MB unified I + D
    - 4.4 GHz
    - Approx 0.23 ns Cycle Time
    - 6 Km of wire
  - 2 Storage Control (SC) chip
    - 21.11 mm x 21.71 mm
    - 1.6 billion transistors/chip
    - L2 Cache 24 MB per SC chip (48 MB/Book)
    - L2 access to/from other MCMs
    - 3 Km of wire
  - 4 SEEPROM (S) chips
    - 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) and 9037 (ETR) functions are on the SC

Orderable Processor Features

Model	Books/PUs	CPs	IFLs uIFLs	zAAPs zIIPs	ICFs	Opt Saps	Std Saps	Std Spares
E12	1/17	0 - 12	0 - 12 0 - 11	0 - 6 0 - 6	0-12	0-3	3	2
E26	2/34	0 - 26	0 - 26 0 - 25	0 - 13 0 - 13	0-16	0-7	6	2
E40	3/51	0 - 40	0 - 40 0 - 39	0 - 20 0 - 20	0-16	0-11	9	2
E56	4/68	0 - 56	0 - 56 0 - 55	0 - 28 0 - 28	0-16	0-18	10	2
E64	4/77	0 - 64	0 - 64 0 - 63	0 - 32 0 - 32	0-16	0-21	11	2

Note: A minimum of one CP, IFL, or ICF must be purchased on every model.

Note: One zAAP and one zIIP may be purchased for each CP purchased.

Note: System z10 EC is designed not to require Optional SAPs for production workloads except sometimes for TPF or z/TPF workloads.

System z10 EC CP Subcapacity (12 or Fewer CPs)

CP Capacity Relative to Full Speed
7nn = 100%
6nn ≈ aa%
5nn ≈ bb%
4nn ≈ cc%
nn = 01 Through 12

Subcapacity CP settings
1-CP through 12-CP only
Available on any hardware model

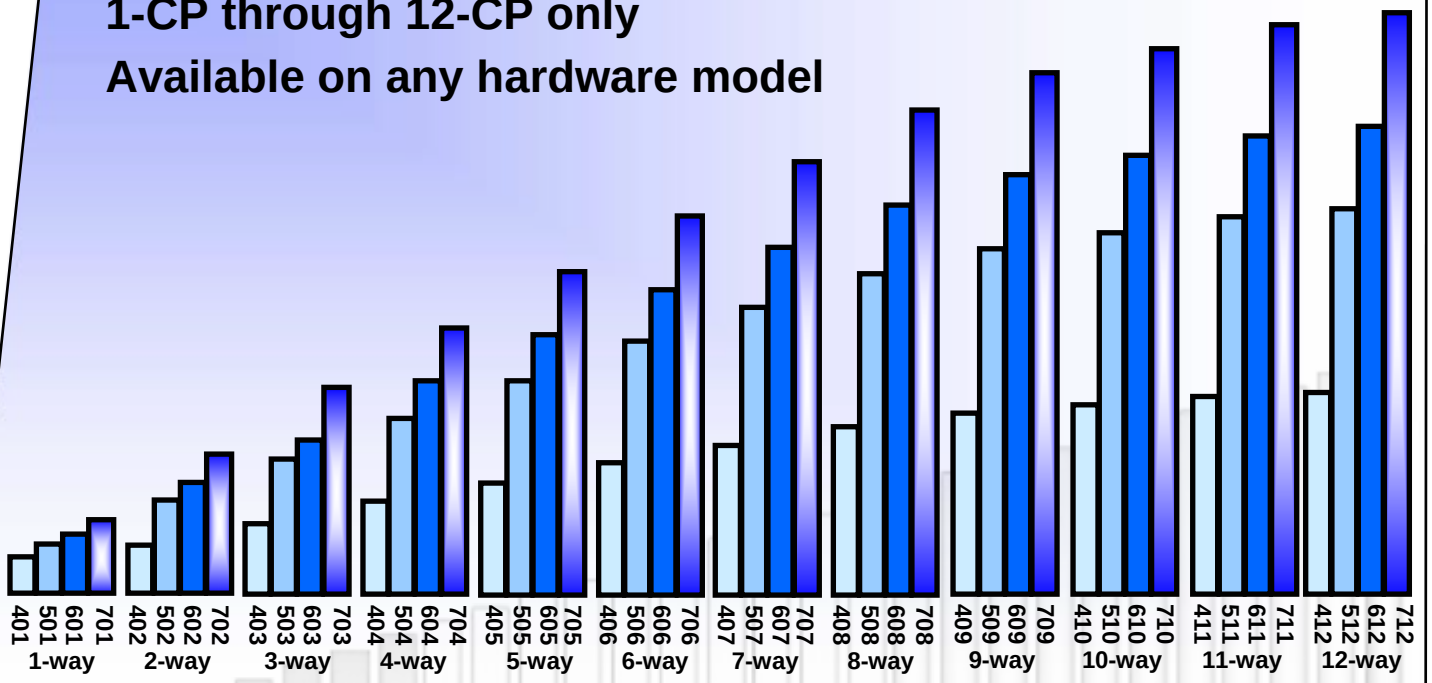

- The System z10 EC will offer 36 CP subcapacity settings with the first twelve or fewer CPs (general purpose) engines.
  - All CPs must be the same capacity within one z10 EC
  - On machines with 13 or more CPs, all CPs must run at full speed
- The entry point is approximately xx% of the capacity of the full speed CP
- **All specialty engines run at full speed**. The one for one entitlement to purchase one zAAP and one zIIP for each CP purchased is the same for CPs of any speed.

LSPR Ratios and MSU Values for System z10 EC

	z10 EC to z9 EC Ratios	z10 EC MSU Values*
LSPR mixed workload average, multi-image for z/OS 1.8 with HiperDispatch active on z10 EC!		
Uni-processor	1.62	115 for 701
16-way z10 EC to 16-way z9 EC	1.49	1,264 for 716
32-way z10 EC to 32-way z9 EC	1.49	2,200 for 732
56-way z10 EC to 54-way z9 EC	1.54	3,395 for 756
64-way z10 EC to 54-way z9 EC	1.70	3,739 for 764

* ***Reflects Mainframe Charter Technology Dividend.***

z10 EC Capacity Planning in a nutshell

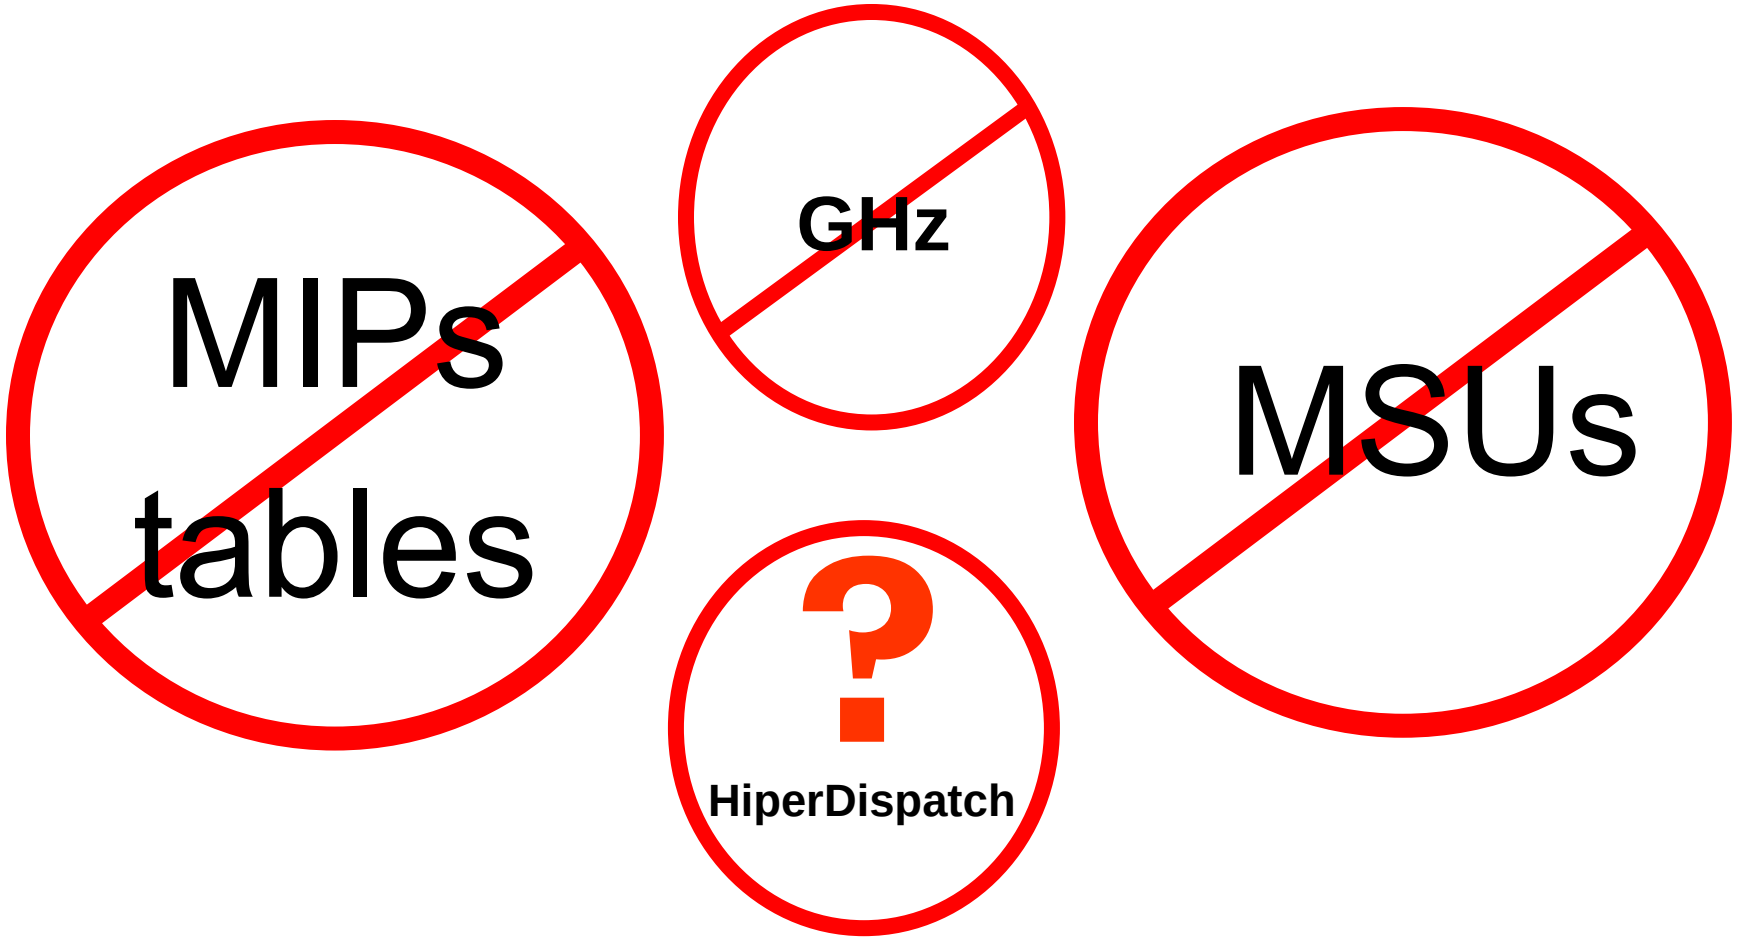

Don't use "one number" capacity comparisons!
Work with IBM technical support for capacity planning!
Customers can now use zPCR

Evolution of System z Specialty Engines

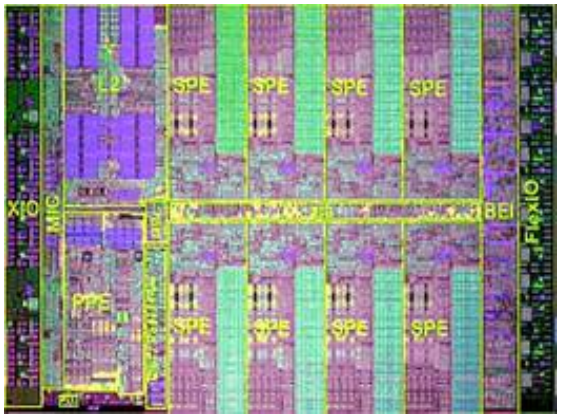

Cell Broadband Engine™

Building on a strong track record of technology innovation with specialty engines – DB Compression, SORT, Encryption, Vector Facility

Integrated Facility for Linux (IFL) 2000

Internal Coupling Facility (ICF) 1997

System z Application Assist Processor (zAAP) 2004

Eligible for zAAP:

- **Java™ execution environment**
- **z/OS XML**

IBM System z9 Integrated Information Processor (IBM zIIP) 2006

Eligible for zIIP:

- **DB2 remote access and BI/DW**
- **ISVs**
- **New! IPSec encryption**
- **z/OS XML**
- **z/OS Global Mirror***

***SOD: IBM plans to enhance z/VM in a future release to support the new System z10 EC capability to allow any combination of CP, zIIP, zAAP, IFL, and ICF processor-types to reside in the same z/VM LPAR**

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

IBM System z10 EC Capacity on Demand (CoD)

System z10 EC Capacity on Demand Reinvented!

- **Permanent and temporary offerings – with you in charge**
  - **Permanent offerings – Capacity Upgrade on Demand (CUoD), Customer Initiated Upgrade (CIU)**
  - **Temporary offerings**
    - **Additional capacity - On/Off Capacity on Demand (On/Off CoD)**
    - **Replacement capacity - Backup Upgrade (CBU) and a new one – Capacity for Planned Event (CPE)**
- **No customer interaction with IBM at time of activation**
  - **Broader customer ability to order temporary capacity**
- **Multiple offerings can be in use simultaneously**
  - **All offerings on Resource Link**
  - **Each offering independently managed and priced**
- **Flexible offerings may be used to solve multiple situations**
  - **Configurations based on real time circumstances**
  - **Ability to dynamically move to any other entitled configuration**
- **Offerings can be reconfigured or replenished dynamically**
  - **Modification possible even if offering is currently active**
  - **Some permanent upgrades permitted while temporary offerings are active**
- **Policy based automation capabilities**
  - **Using Capacity Provisioning Manager with z/OS 1.9**
  - **Using scheduled operations via HMC**

IBM System z10 EC
HiperDispatch and Large Page Support

z9 vs z10 EC CEC Structure

	z9 EC	z10 EC
SMP Configuration	S54 4 books, 64 PUs	E64 4 books, 77 PUs
Topology	Dual Ring One or Two Hops	Fully Connected
Jumper Books	Yes	No
Max Memory	Up to 512GB - HSA?	Up to 1,520 GB + 16 GB HSA
Cache Levels	L1 per PU L2 per Book	L1 and L1.5 per PU L2 per Book
Page Sizes	4 KB	4 KB and 1 MB

z10 EC HiperDispatch

- **HiperDispatch – z10 EC unique function**
  - Dispatcher Affinity (DA) - New z/OS Dispatcher
  - Vertical CPU Management (VCM) - New PR/SM Support
- **Hardware cache optimization occurs when a given unit of work is consistently dispatched on the same physical CPU**
  - Up till now software, hardware, and firmware have had pride in the fact of how independent they were from each other
  - Non-Uniform-Memory-Access has forced a paradigm change
    - CPUs have different distance-to-memory attributes
    - Memory accesses can take a number of cycles depending upon cache level / local or remote repository accessed
- **The entire z10 EC hardware/firmware/OS stack now tightly collaborates to obtain the hardware's full potential**

z10 EC HiperDispatch – PR/SM Functionality

- **New PR/SM Support**
  - Topology information exchanged with z/OS
    - z/OS uses this to construct its dispatching queues
  - Classes of logicals
    - High priority allowed to consume weight
      - Tight tie of logical processor to physical processor
    - Low priority generally run only to consume white space
- **Firmware Support (PR/SM, millicode)**
  - New z/OS invoked instruction to cause PR/SM to enter “Vertical mode” to assign vertical LPs subset and their associated LP to physical CP mapping (based upon LPAR weight)
  - Enables z/OS to concentrate its work on fewer vertical processors
  - Key in PR/SM overcommitted environments to reduce the LP competition for physical CP resources
  - Vertical LPs are assigned High, Medium, and Low attributes
  - Vertical low LPs shouldn’t be used unless there is logical white space within the CEC and demand within LPAR

Large Page Support

- **Issue: Translation Lookaside Buffer (TLB) Coverage shrinking as % of memory size**
  - Over the past few years application memory sizes have dramatically increased due to support for 64-bit addressing in both physical and virtual memory
  - TLB sizes have remained relatively small due to low access time requirements and hardware space limitations
  - TLB coverage today represents a much smaller fraction of an applications working set size leading to a larger number of TLB misses
  - Applications can suffer a significant performance penalty resulting from an increased number of TLB misses as well as the increased cost of each TLB miss
- **Solution: Increase TLB coverage without proportionally enlarging the TLB size by using large pages**
  - Large Pages allow for a single TLB entry to fulfill many more address translations
  - Large Pages will provide exploiters with better TLB coverage
- **Benefit:**
  - Designed for better performance by decreasing the number of TLB misses that an application incurs

IBM System z10 EC Memory

z10 EC Memory Offering and Assignment

Model	Standard Memory GB	Flexible Memory GB
E12	16 - 352	NA
E26	16 - 752	32 - 352
E40	16 - 1136	32 - 752
E56	16 - 1520	32 - 1136
E64	16 - 1520	32 - 1136

- **Memory Granularity for ordering:**
  - **16 GB: Std - 16 to 256; Flex - 32 to 256**
  - **32 GB: Std - 288 to 512; Flex - 288 to 512**
  - **48 GB: Std - 560 to 944; Flex - 560 to 944**
  - **64 GB: Std - 1008 to 1520; Flex - 1008 to 1136**
- **16 GB separate fixed HSA standard**
- **Maximum Physical Memory: 384 GB per book, 1.5 TB per system**
  - **Up to 48 DIMMs per book**
  - **64 GB minimum physical memory in each book**
  - **Physical Memory Increments:**
    - **32 GB – Eight 4GB DIMMs (FC #1604) Preferred if can fulfill purchase memory**
    - **64 GB – Eight 8 GB DIMMs (FC #1608) Used where necessary**
- **For Flexible, if required, 16 GB "Pre-planned Memory" features (FC # 1996) are added to the configuration.**

z10 EC Concurrent Memory Upgrades

- **LIC enable additional memory to the physical limit of the installed cards and memory configuration**
  - Designed to be possible and concurrent in many but not all configurations
- **Add a book with additional memory**
  - Designed to be possible except for Models E56 and E64
- **Exploit Enhanced Book Availability to change memory card configuration in existing books**
  - Exploits the capability to remove, upgrade and return a book concurrently
  - Simplest with flexible memory and PU configurations
  - May be possible with standard memory and PU configurations depending on LPAR configuration
  - Customer pre-planning required may require acquisition of additional hardware resources
  - Not possible on Model E12

Note: Concurrent memory upgrades above are designed not to require CEC activation (POR).
z/OS with “reserved memory” configured in the LPAR profile can add memory to a running partition. Otherwise adding memory to a partition requires deactivation, profile change and activation of the partition, which is designed not to be disruptive to other partitions.

IBM System z10 EC Cryptography

System z Crypto over time

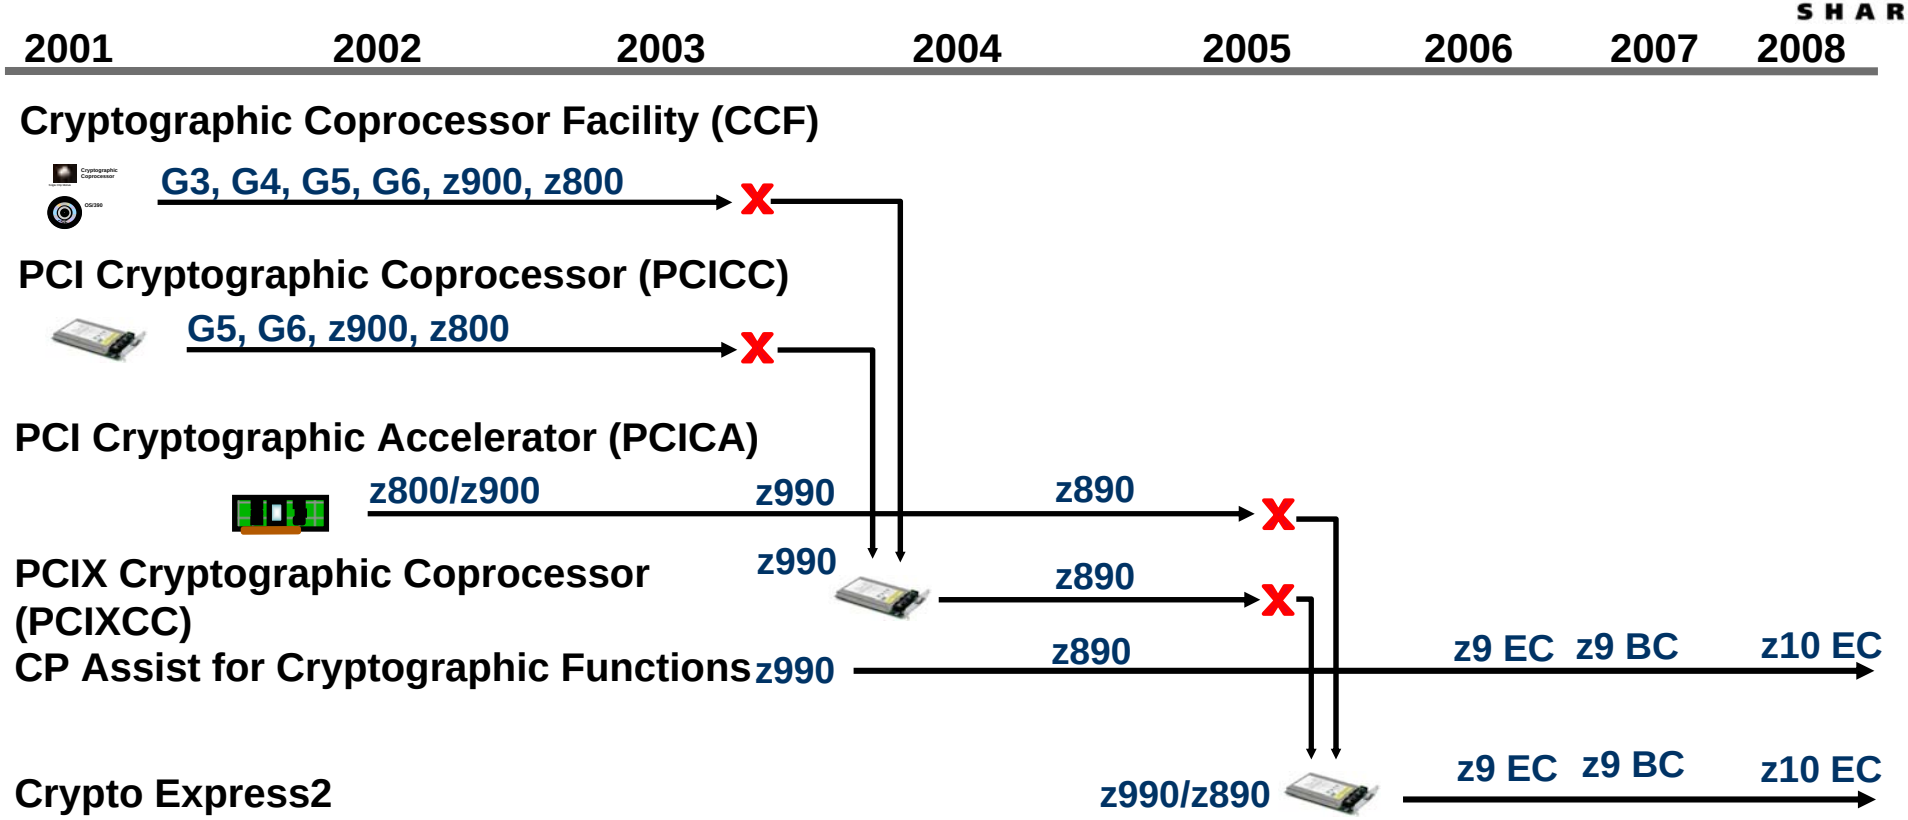

- **Cryptographic Coprocessor Facility – Supports "Secure key" cryptographic processing**
- **PCICC Feature – Supports "Secure key" cryptographic processing**
- **PCICA Feature – Supports "Clear key" SSL acceleration**
- **PCIXCC Feature – Supports "Secure key" cryptographic processing**
- **CP Assist 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**

z10 EC CP Assist for Cryptographic Functions (CPACF)

Integrated Cryptographic Service Facility (ICSF)

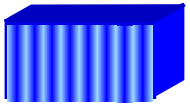

Crypto Express2

PU PU PU PU PU PU PU PU

CP Assist for Cryptographic Function

z9

z10 EC

High performance clear key symmetric encryption/decryption

IBM System z10 EC I/O Structure

z10 EC Book Layout – Under the covers

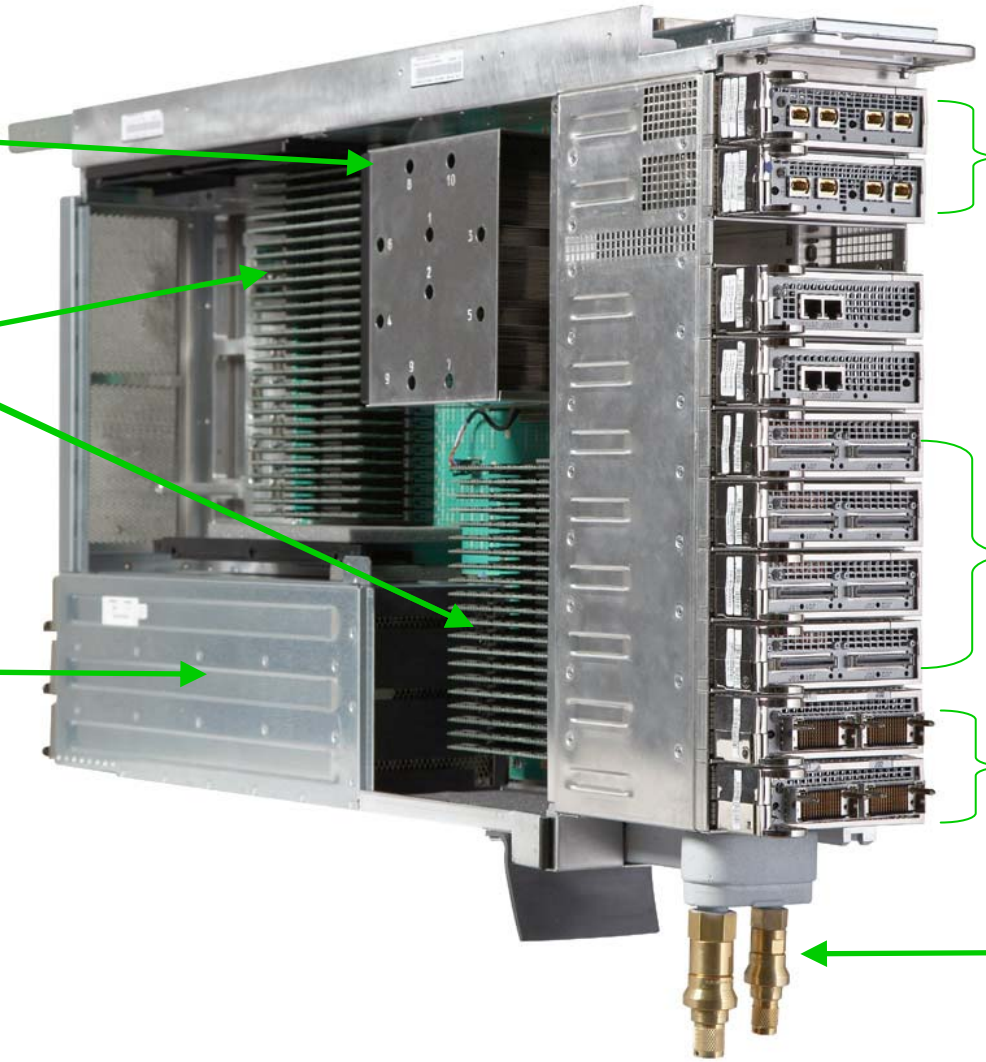

z10 EC Channel Type and Crypto Overview

- **FICON/FCP**
  - FICON Express4
  - FICON Express2 (carry forward only)
  - FICON Express (carry forward only)
- **Networking**
  - OSA-Express3 (2Q2008)
    - 10 Gigabit Ethernet LR
  - OSA-Express2
    - 1000BASE-T Ethernet
    - Gigabit Ethernet LX and SX
    - 10 Gigabit Ethernet LR
  - HiperSockets (Define only)
- **ESCON**
- **Coupling Links**
  - InfiniBand (PSIFB) – 2Q2008
  - ISC-3 (Peer mode only)
  - ICB-4 (Not available on Model E64)
  - IC (Define only)
- **Time Features**
  - **STP - Optional**
  - **ETR Attach – Standard**
- **Crypto**
  - Crypto Express2
    - Configurable Coprocessor or Accelerator
- **Channel types not supported:**
  - OSA-Express
  - ICB-3
  - Features not supported on System z9

Note: ICB-4 cables are available as features.
All other cables are sourced separately

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

IBM System z10 EC Availability

System z – RAS Design Focus

- **High Availability (HA)**
  - The attribute of a system designed to provide service during defined periods, at acceptable or agreed upon levels and masks UNPLANNED OUTAGES from end-users. It employs Fault Tolerance; Automated Failure Detection, Recovery, Bypass Reconfiguration, Testing, Problem and Change Management
- **Continuous Operations (CO)**
  - Attribute of a system designed to continuously operate and mask PLANNED OUTAGES from end-users. It employs non-disruptive hardware and software changes, non-disruptive configuration, software coexistence
- **Continuous Availability (CA)**
  - Attribute of a system designed to deliver non disruptive service to the end user 7 days a week, 24 HOURS A DAY (there are no planned or unplanned outages). It includes the ability to recover from a site disaster by switching computing to a second site

System z10 EC continues to focus on RAS

Keeping your system available is key to our total design

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

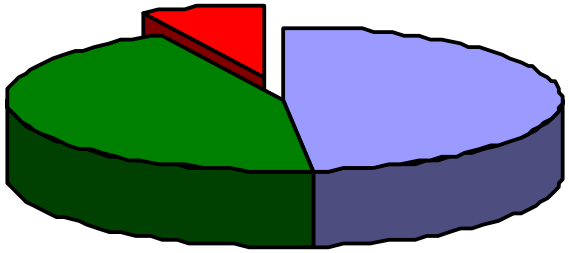

- Scheduled (CIE+Disruptive Patches + ECs)
- Planned - (MES + Driver Upgrades)
- Unscheduled (UIRA)

Impact of Outage

	Prior Servers	z9 EC	z10 EC
Unscheduled Outages	✓	✓	✓
Scheduled Outages	✓	✓	✓
Planned Outages		✓	✓
Preplanning requirements			✓

z10 EC Increased Focus

z10 EC Enhancements designed to avoid Outages

- **Continued Focus on Firmware Quality**
- **Reduced Chip Count on MCM**
- **Memory Subsystem Improvements**

- **DIMM FRU indicators**
- **Single Processor Core Checkstop**
- **Single Processor Core Sparing**
- **Point to Point SMP Fabric (not a ring)**
- **Rebalance PSIFB and I/O Fanouts**
- **Redundant 100Mb Ethernet service network w/ VLAN**

- **Elimination of unnecessary CBU passwords**
- **Enhanced Driver Maintenance (EDM) Upgrades**
  - Multiple “from” sync point support
  - Improved control of channel LIC levels
- **Reduce Pre-planning to Avoid POR**
  - 16 GB for HSA
  - Dynamic I/O Enabled by Default
  - Add Logical Channel Subsystem (LCSS)
  - Change LCSS Subchannel Sets
  - Add/Delete Logical Partitions
- **Reduce Pre-Planning to Avoid LPAR Deactivate**
  - Change Partition Logical Processor Config
  - Change Partition Crypto Coprocessor Config
- **CoD – Flexible Activation/Deactivation**

IBM System z10 EC Backup

System z10: Obsolete Instructions Removed

- **Instructions present on System z9 and removed on System z10**
  - Hexadecimal Operation Codes: E503, E504, E505, E506, E507
  - Use on System z10 will result in Operation Exception – 0C1 ABEND in software

- **Removed Instruction Archeology**
  - Removed instructions were originally introduced as part of the “IBM System/370: Assists for MVS” to provide various assist and lock manipulation functions
    - These five were carried forward to System z9
    - Former “MVS Assist” instruction Op Code E502 was recycled years ago for the z/Architecture 64-bit Store Real Address (STRAG) instruction
    - Other “MVS Assist” instructions have been gone from hardware for many years
  - Removed instructions have not been shipped in any IBM software for many years
  - Removed instructions have never been part of z/Architecture as documented in Reference: **z/Architecture Principles of Operations**, SA22-7832

System z10: New Instructions Added (1 of 4)

Mnemonic	Format	OpCode	Instruction Name
ASI	SIY	EB6A	ADD IMMEDIATE (32←8)
AGSI	SIY	EB7A	ADD IMMEDIATE (64←8)
ALSI	SIY	EB6E	ADD LOGICAL WITH SIGNED IMMEDIATE (32←8)
ALGSI	SIY	EB7E	ADD LOGICAL WITH SIGNED IMMEDIATE (64←8)
CRB	RRS	ECF6	COMPARE AND BRANCH (32)
CGRB	RRS	ECE4	COMPARE AND BRANCH (64)
CRJ	RIE	EC76	COMPARE AND BRANCH RELATIVE (32)
CGRJ	RIE	EC64	COMPARE AND BRANCH RELATIVE (64)
CRT	RRF	B972	COMPARE AND TRAP (32)
CGRT	RRF	B960	COMPARE AND TRAP (64)
CGH	RXY	E334	COMPARE HALFWORD (64←16)
CHHSI	SIL	E554	COMPARE HALFWORD IMMEDIATE (16←16)
CHSI	SIL	E55C	COMPARE HALFWORD IMMEDIATE (32←16)
CGHSI	SIL	E558	COMPARE HALFWORD IMMEDIATE (64←16)
CHRL	RIL	C65	COMPARE HALFWORD RELATIVE LONG (32←8)
CGHRL	RIL	C64	COMPARE HALFWORD RELATIVE LONG (64←16
CIB	RIS	ECFE	COMPARE IMMEDIATE AND BRANCH (32←8)
CGIB	RIS	ECFC	COMPARE IMMEDIATE AND BRANCH (64←8)
CIJ	RIE	EC7E	COMPARE IMMEDIATE AND BRANCH RELATIVE (32←8)
CGIJ	RIE	EC7C	COMPARE IMMEDIATE AND BRANCH RELATIVE (64←8)

System z10: New Instructions Added (2 of 4)

Mnemonic	Format	OpCode	Instruction Name
CIT	RIE	EC72	COMPARE IMMEDIATE AND TRAP (32←16)
CGIT	RIE	EC70	COMPARE IMMEDIATE AND TRAP (64←16)
CLRB	RRS	ECF7	COMPARE LOGICAL AND BRANCH (32)
CLGRB	RRS	ECE5	COMPARE LOGICAL AND BRANCH (64)
CLRJ	RIE	EC77	COMPARE LOGICAL AND BRANCH RELATIVE (32)
CLGRJ	RIE	EC65	COMPARE LOGICAL AND BRANCH RELATIVE (64)
CLRT	RRF	B973	COMPARE LOGICAL AND TRAP (32)
CLGRT	RRF	B961	COMPARE LOGICAL AND TRAP (64)
CLHHSI	SIL	E555	COMPARE LOGICAL IMMEDIATE (16←16)
CLFHSI	SIL	E55D	COMPARE LOGICAL IMMEDIATE (32←16)
CLGHSI	SIL	E559	COMPARE LOGICAL IMMEDIATE (64←16)
CLIB	RIS	ECFF	COMPARE LOGICAL IMMEDIATE AND BRANCH (32←8)
CLGIB	RIS	ECFD	COMPARE LOGICAL IMMEDIATE AND BRANCH (64←8)
CLIJ	RIE	EC7F	COMPARE LOGICAL IMMEDIATE AND BRANCH RELATIVE (32←8)
CLGIJ	RIE	EC7D	COMPARE LOGICAL IMMEDIATE AND BRANCH RELATIVE (64←8)
CLFIT	RIE	EC73	COMPARE LOGICAL IMMEDIATE AND TRAP (32←16)
CLGIT	RIE	EC71	COMPARE LOGICAL IMMEDICAL AND TRAP (64←16)
CLRL	RIL	C6F	COMPARE LOGICAL RELATIVE LONG (32)
CLHRL	RIL	C67	COMPARE LOGICAL RELATIVE LONG (32←16)
CLGRL	RIL	C6A	COMPARE LOGICAL RELATIVE LONG (64)

System z10: New Instructions Added (3 of 4)

Mnemonic	Format	OpCode	Instruction Name
CLGHRL	RIL	C66	COMPARE LOGICAL RELATIVE LONG (64←16)
CLGFRL	RIL	C6E	COMPARE LOGICAL RELATIVE LONG (64←32)
CRL	RIL	C6D	COMPARE RELATIVE LONG (32)
CGRL	RIL	C68	COMPARE RELATIVE LONG (64)
CGFRL	RIL	C6C	COMPARE RELATIVE LONG (64←32)
ECAG	RSY	EB4C	EXTRACT CACHE ATTRIBUTE
EXRL	RIL	C60	EXECUTE RELATIVE LONG
LAEY	RXY	E375	LOAD ADDRESS EXTENDED
LTGF	RXY	E332	LOAD AND TEST (64←32)
LHRL	RIL	C45	LOAD HALFWORD RELATIVE LONG (32←16)
LGHRL	RIL	C44	LOAD HALFWORD RELATIVE LONG (64←16)
LLHRL	RIL	C42	LOAD LOGICAL HALFWORD RELATIVE LONG (32←16)
LLGHRL	RIL	C46	LOAD LOGICAL HALFWORD RELATIVE LONG (64←16)
LLGFRL	RIL	C4E	LOAD LOGICAL RELATIVE LONG (64←32)
LRL	RIL	C4D	LOAD RELATIVE LONG (32)
LGRL	RIL	C48	LOAD RELATIVE LONG (64)
LGFRL	RIL	C4C	LOAD RELATIVE LONG (64←32)
MVCOS	SSF	C80	MOVE WITH OPTIONAL SPECIFICATIONS
MVHHI	SIL	E544	MOVE (16←16)
MVHI	SIL	E54C	MOVE (32←16)

System z10: New Instructions Added (4 of 4)

Mnemonic	Format	OpCode	Instruction Name
MVGHI	SIL	E548	MOVE (64←16)
MFY	RXY	E35C	MULTIPLY
MHY	RXY	E37C	MULTIPLY HALFWORD
MSFI	RIL	C21	MULTIPLY SINGLE IMMEDIATE
MSGFI	RIL	C20	MULTIPLY SINGLE IMMEDIATE
PFD	RXY	E336	PREFETCH DATA
PFDRL	RIL	C62	PREFETCH DATA RELATIVE LONG
PFMF	RRE	B9AF	PERFORM FRAME MANAGEMENT FUNCTION
PTF	RRE	B9A2	PERFORM TOPOLOGY FUNCTION
RNSBG	RIE	EC54	ROTATE THEN AND SELECTED BITS
RXSBG	RIE	EC57	ROTATE THEN EXCLUSIVE OR SELECTED BITS
RISBG	RIE	EC55	ROTATE THEN INSERT SELECTED BITS
ROSBG	RIE	EC56	ROTATE THEN OR SELECTED BITS
STHRL	RIL	C47	STORE HALFWORD RELATIVE LONG
STRL	RIL	C4F	STORE RELATIVE LONG (32)
STGRL	RIL	C4B	STORE RELATIVE LONG (64)

New book for new instructions: **z/Architecture Principles of Operations**, SA22-7832-**06**

System z10: New and Changed Instruction Formats

RIE Format

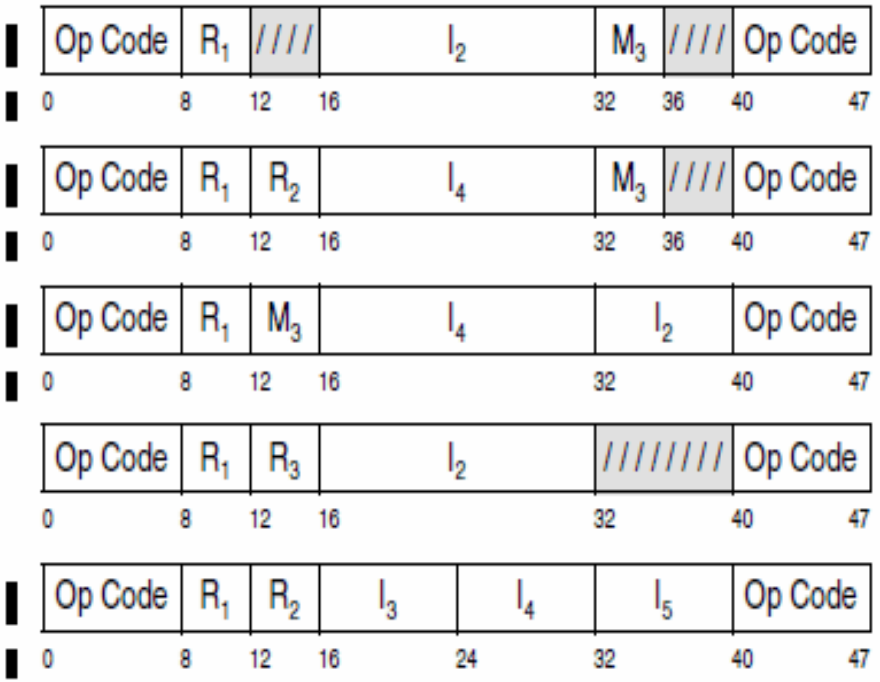

RIS Format

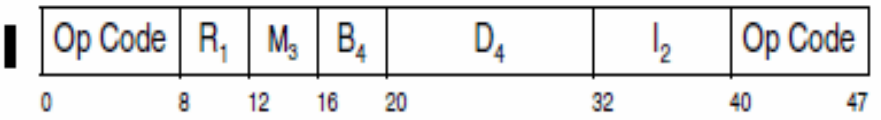

RRS Format

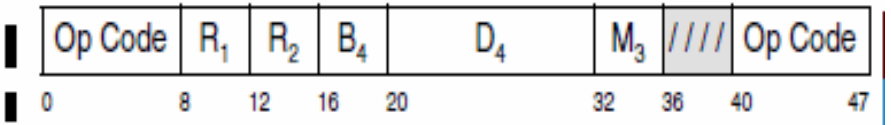

SIL Format

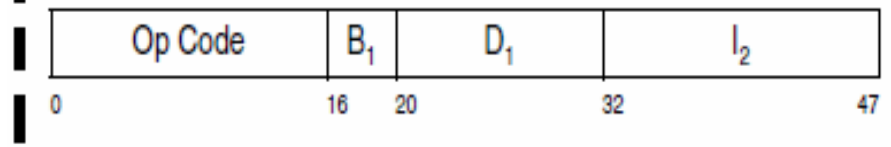

zEnd