IBM System z10™ Enterprise Class

Hardware Overview

CMG – April 15, 2008

Tom Russell – IBM Canada

Thanks to Harv Emery, John Hughes, WSC

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
- Fixed HSA as 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 & 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
- 72 New Instructions
- Improved RAS

No support for Japanese Compatibility Mode (JCM)
No support for MVS Assist instructions

* 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.

z10 EC System Upgrades

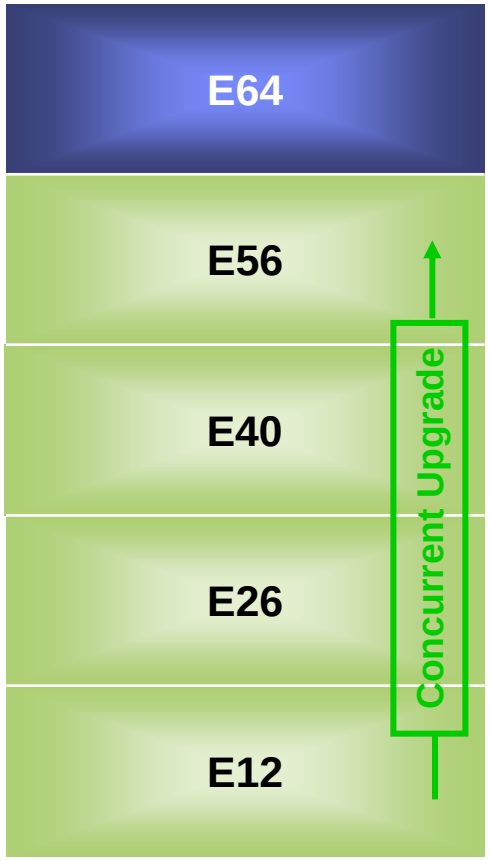

- **z10 EC to higher z10 EC model**
  - Concurrent upgrade of z10 EC Models E26, E40 and E56. Upgrade to E64 is disruptive
  - When upgrading to z10 EC E64, unlike the z9 EC, the first Book is retained
- **Any z9 EC to any z10 EC**
- **Any z990 to any z10 EC**

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.**

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

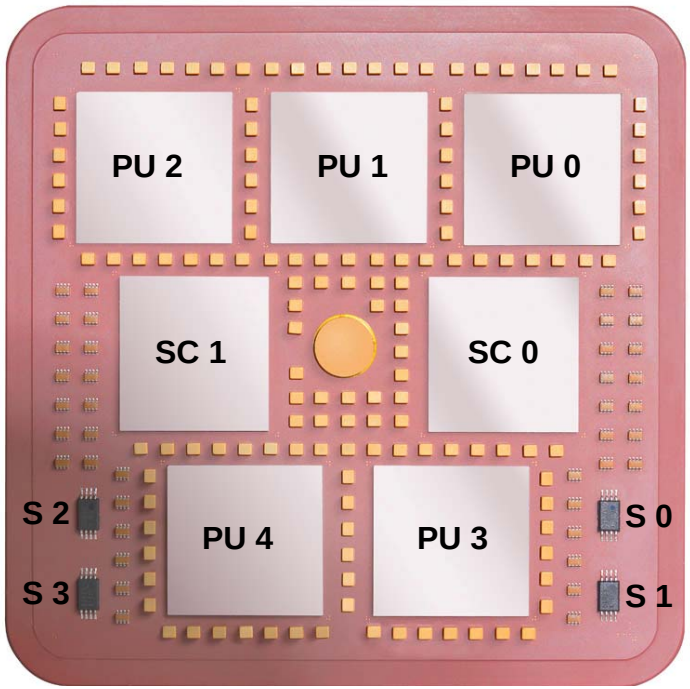

- **CMOS 11s chip Technology**
  - PU, SC, S chips, 65 nm
  - 5 PU chips/MCM – Each up to 4 cores
    - One memory control (MC) per PU chip
    - 21.97 mm x 21.17 mm
    - 994 million transistors/chip
    - L1 cache/PU
      - 64 KB I-cache
      - 128 KB D-cache
    - L1.5 cache/PU
      - 3 MB
    - 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

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

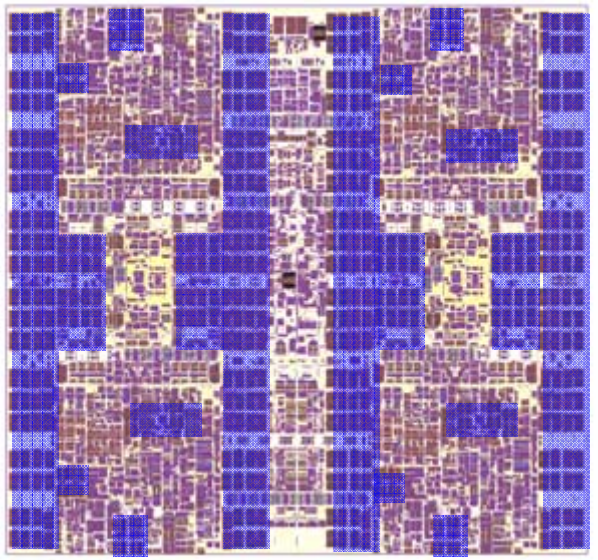

POWER6 Dual Core Chip

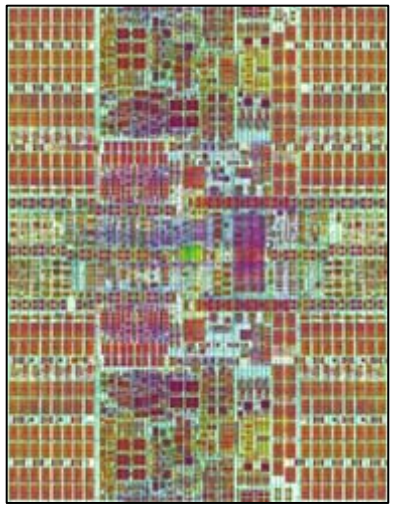

Orderable Processor Features

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.

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**
- All supported z/OS releases (z/OS 1.7 requires the zIIP web deliverable)

System z9 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

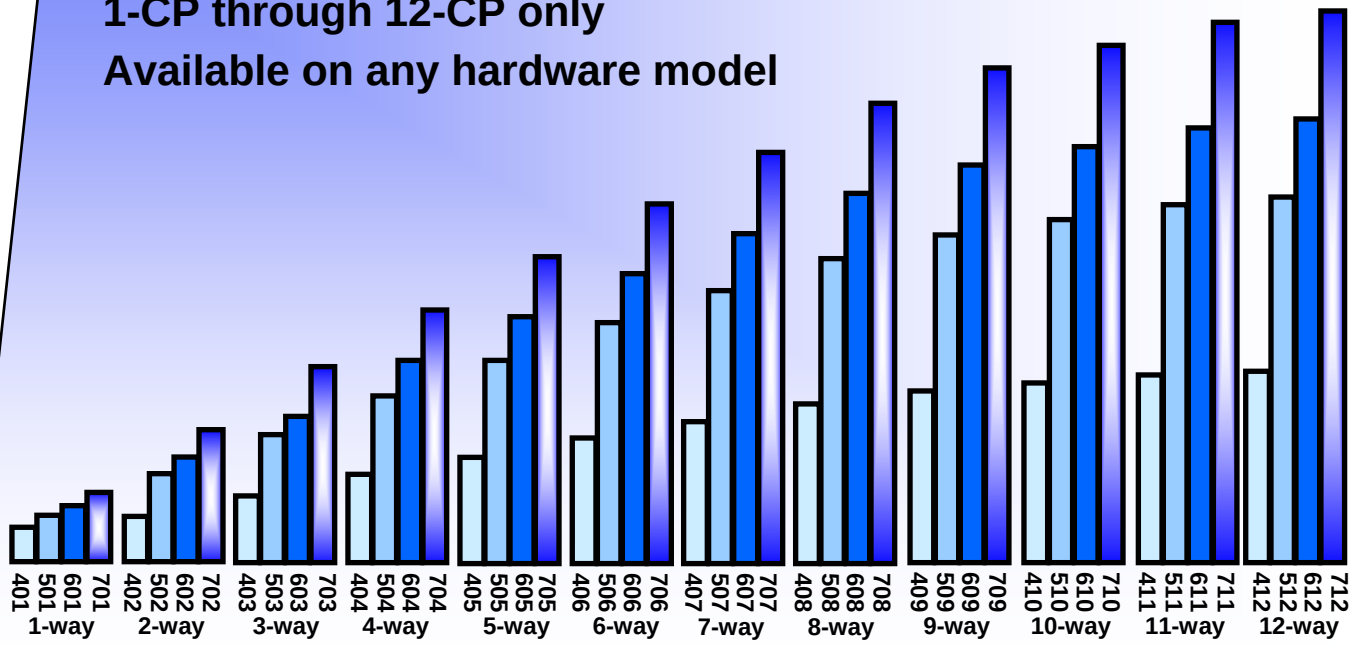

- The System z10 EC will offer 36 CP subcapacity settings with the first twelve or fewer CPs (general purpose) engines.
  - All CPs must be the same capacity within one z10 EC
  - On machines with 13 or more CPs, all CPs must run at full speed
- The entry point is approximately xx% of the capacity of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

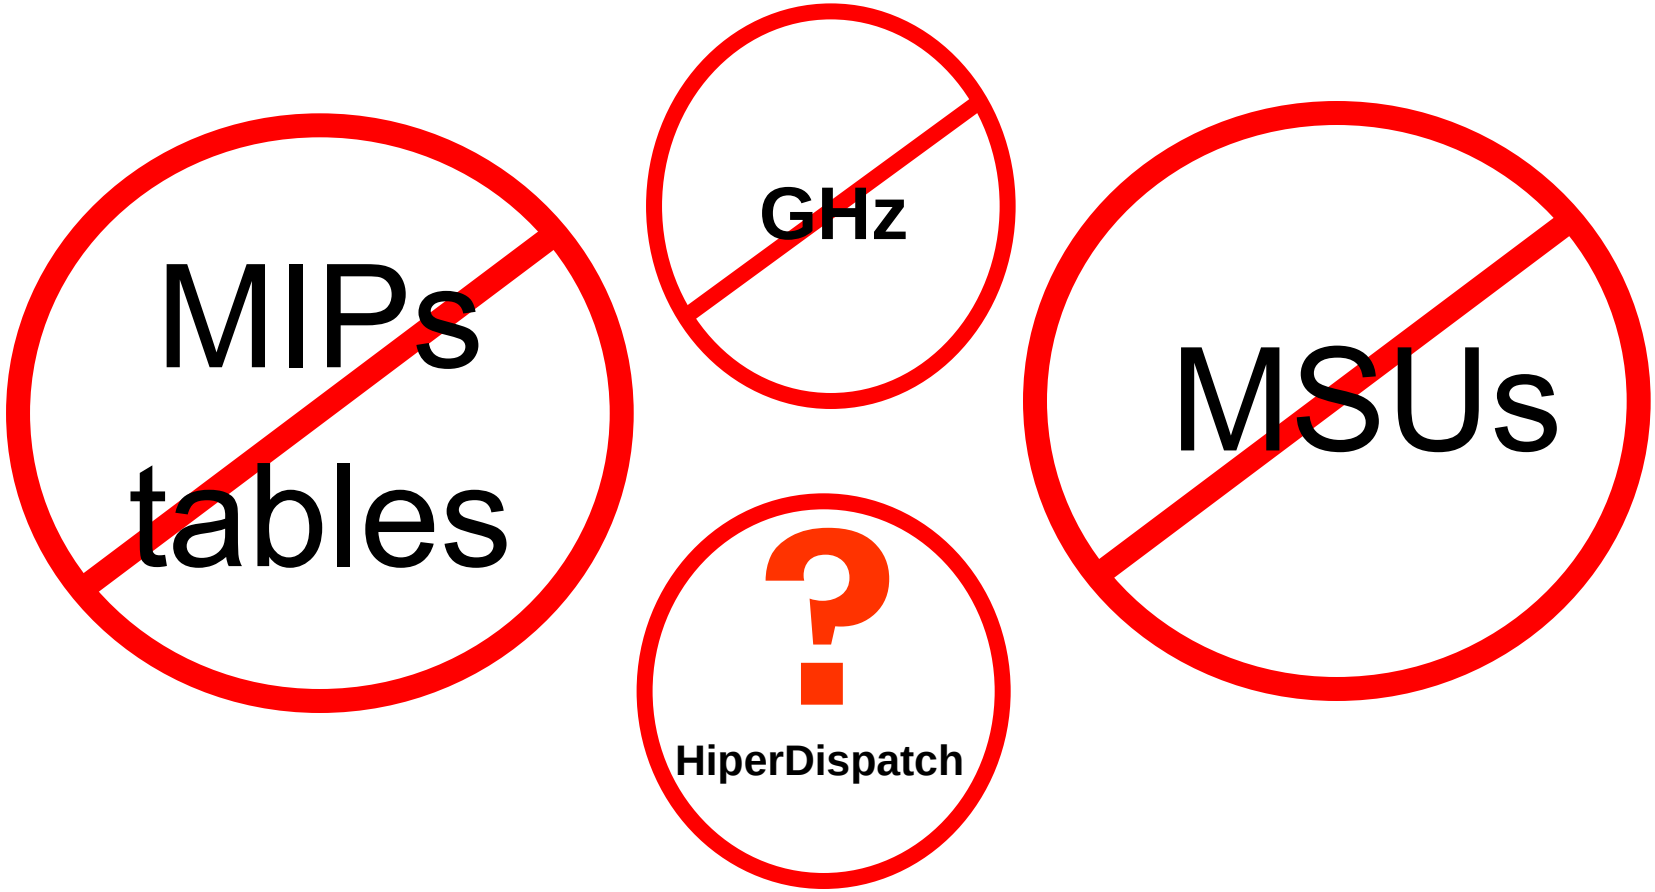

Don't use "one number" capacity comparisons!
Work with IBM technical support for capacity planning!
Customers can now use zPCR

Evolution of System z Specialty Engines

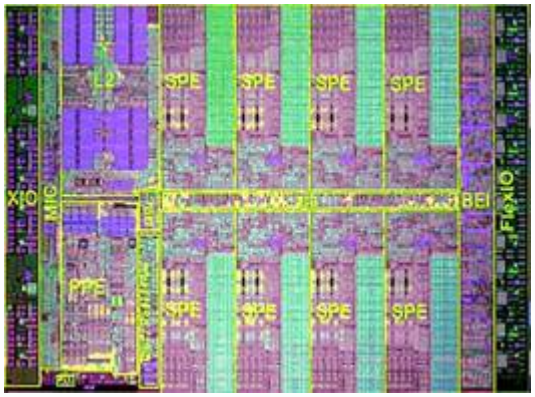

Cell Broadband Engine™

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

Integrated Facility for Linux (IFL) 2000

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***

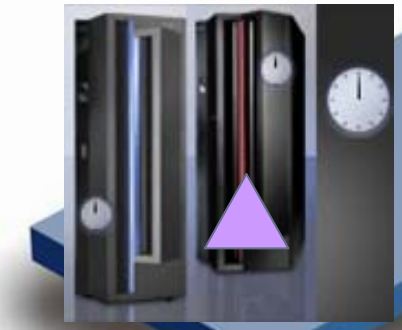

Internal Coupling Facility (ICF) 1997

***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.

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

New Instructions

- **A large variety of facilities added in the System z10 EC**
  - **General-Instructions Extension Facility (72 new)**
  - **Execute-Extension Facility (1 new)**
  - **Parsing-Enhancement Facility (2 new)**
  - **Compare-and-Swap-and-Store Facility 2 (new function)**
  - **Message-Security-Assist Extensions (new functions)**
  - **Enhanced-DAT Facility (1 new, 3 changed)**
  - **Configuration-Topology Facility (1 new, 1 changed)**
- **Potential for:**
  - **Significant performance improvement**
  - **Enhanced capabilities**
  - **Simpler code**

Examples – New Instructions

- **Compare and Branch**
  - **Replaces Compare, followed by Branch on Condition**
- **Compare and Trap**
  - **Program Check (Data Exception) if Compare is True**
- **New Immediate instructions**
  - **Saves having storage references**

- **Primary motivation: Performance**
  - **ARCH(8) TUNE(8) options in compilers**
  - **I expect Java to get the most benefit**

IBM System z10 EC Capacity on Demand (CoD)

Removing Road Blocks

- **A permanent upgrade cannot occur while CBU or On/Off CoD is active.**
- **Only one solution can be active at a time**
- **Limited to the permanent capacity**
  - **After a permanent capacity upgrade, the old temporary contract may become useless.**
- **Cannot add temporary capacity while a Concurrent Book Add is in progress.**
- **No CBU-like replacement capacity offering where a disaster is not involved.**
- **When On/Off CoD or CBU records are activated/deactivated, all processors defined in those records must be activated/deactivated.**
- **The HMC requires connectivity to the IBM Support System to obtain temporary records or verify passwords at the time of activation.**
  - **HMC connectivity or response time is a potential inhibitor.**
  - **The process to activate capacity can take too long.**
- **No way to determine which capacity is billable versus replacement**
- **Drastic system slow down occurs if CBU or CBU test expires**
- **Automation provides only limited control**

The Basics – Temporary Upgrades

- **Capacity Backup (CBU)**
  - **Predefined capacity for disasters on a other "lost" server(s)**
  - **Concurrently add CPs, IFLs, ICFs, zAAPs, zIIPs, SAPs**
  - **Pre-paid**
- **Capacity for Planned Events (CPE)**
  - **CBU-like offering, when a disaster is not declared**
  - **Example: System migration (push/pull) or relocation (data center move)**
  - **Predefined capacity for a fixed period of time (3 days)**
  - **Pre-paid**
- **On/Off Capacity on Demand (On/Off CoD)**
  - **Satisfy periods of peek demand for computing resources**
  - **Concurrent 24 hour rental of CPs, IFLs, ICFs, zAAPs, zIIPs, SAPs**
  - **Supported through a new software offering – Capacity Provisioning Manager (CPM)**
  - **Post-paid**

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 Memory

z10 EC – HSA considerations

- **HSA of 16GB provided as standard outside of purchased memory**
- **The HSA has been designed to eliminate planning for HSA. Preplanning for HSA expansion for configurations is eliminated because HCD/IOCP will, via the IOCDS process, always reserves HSA space for:**
  - 4 CSSs
  - 15 LPs in each CSS (total of 60 LPs)
  - Subchannel set-0 with 63.75k devices in each CSS
  - Subchannel set-1 with 64k devices in each CSS
  - All the above are designed to be activated and used with dynamic I/O changes

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

- **Customer 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**
- **LIC CC controls purchased memory**
- **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.**

IBM System z10 EC Cryptography

z10 EC CP Assist for Cryptographic Functions (CPACF)

Integrated Cryptographic Service Facility (ICSF)

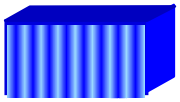

Crypto Express2

PU PU PU PU PU PU PU PU

CP Assist for Cryptographic Function

z9

z10 EC

High performance clear key symmetric encryption/decryption

z10 EC Cryptographic Support

- **CP Assist for Cryptographic Function (CPACF)**
  - Standard on every CP and IFL
  - Supports DES, TDES, AES and SHA
  - Pseudo Random Number Generation (PRNG)
  - New to z10 EC
    - Advanced Encryption Standard (AES) – 192 and 256
    - Secure Hash Algorithm (SHA) – 384 and 512
- **Crypto Express2**
  - Two configuration modes
    - Coprocessor (default)
      - Federal Information Processing Standard (FIPS) 140-2 Level 4 certified
    - Accelerator (configured from the HMC)
  - Three configuration options
    - Default set to Coprocessor
  - Concurrent Patch
- **Dynamic Add Crypto to LPAR**
  - No recycling of LPAR
  - No POR required
- **TKE 5.0 workstation with TKE 5.2 LIC**
  - **Diskette drive support – read only**

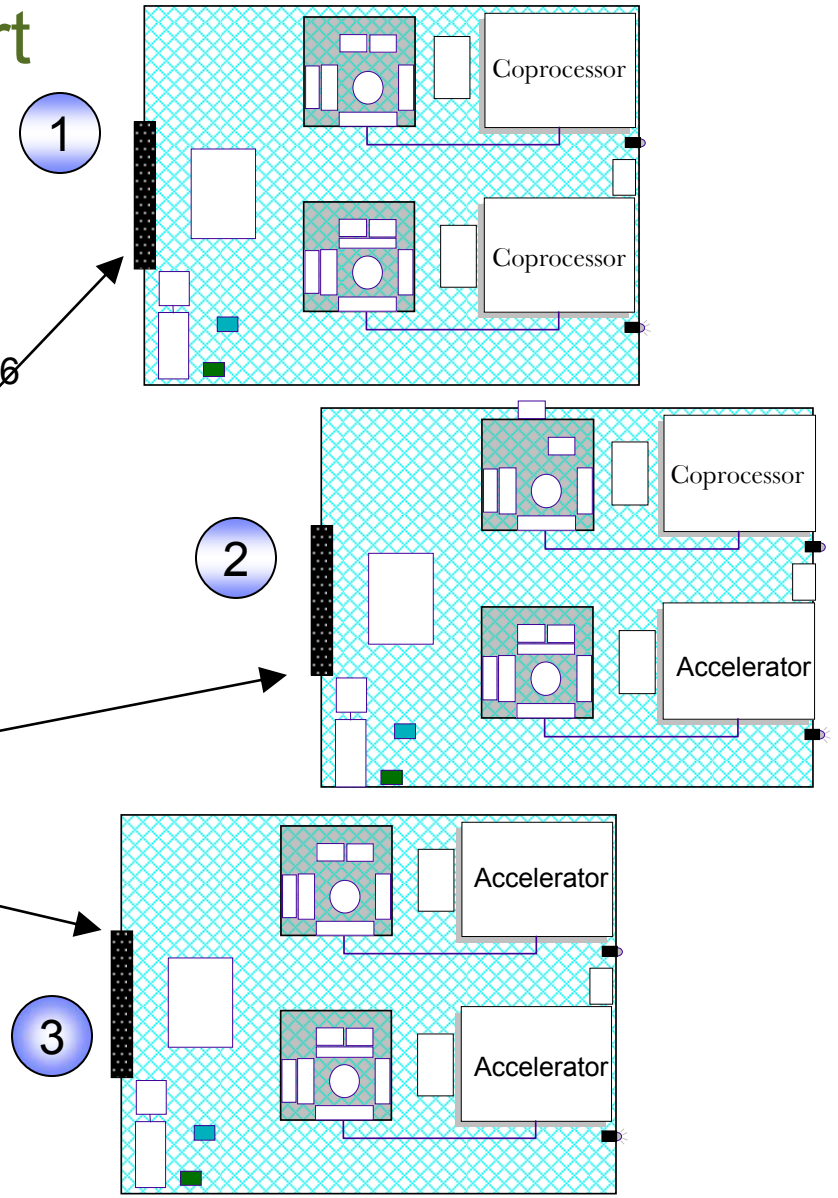

IBM System z10 EC Availability

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-

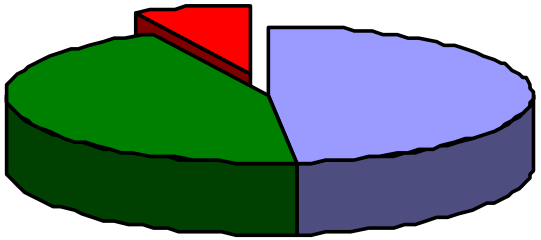

- Scheduled (CIE+Disruptive Patches + ECs)
- Planned - (MES + Driver Upgrades)
- Unscheduled (UIRA)

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**

- **CoD – Flexible Acitvation/Deactivation**
- **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

IBM System z10 EC I/O Structure

z10 EC Book Layout – Under the covers

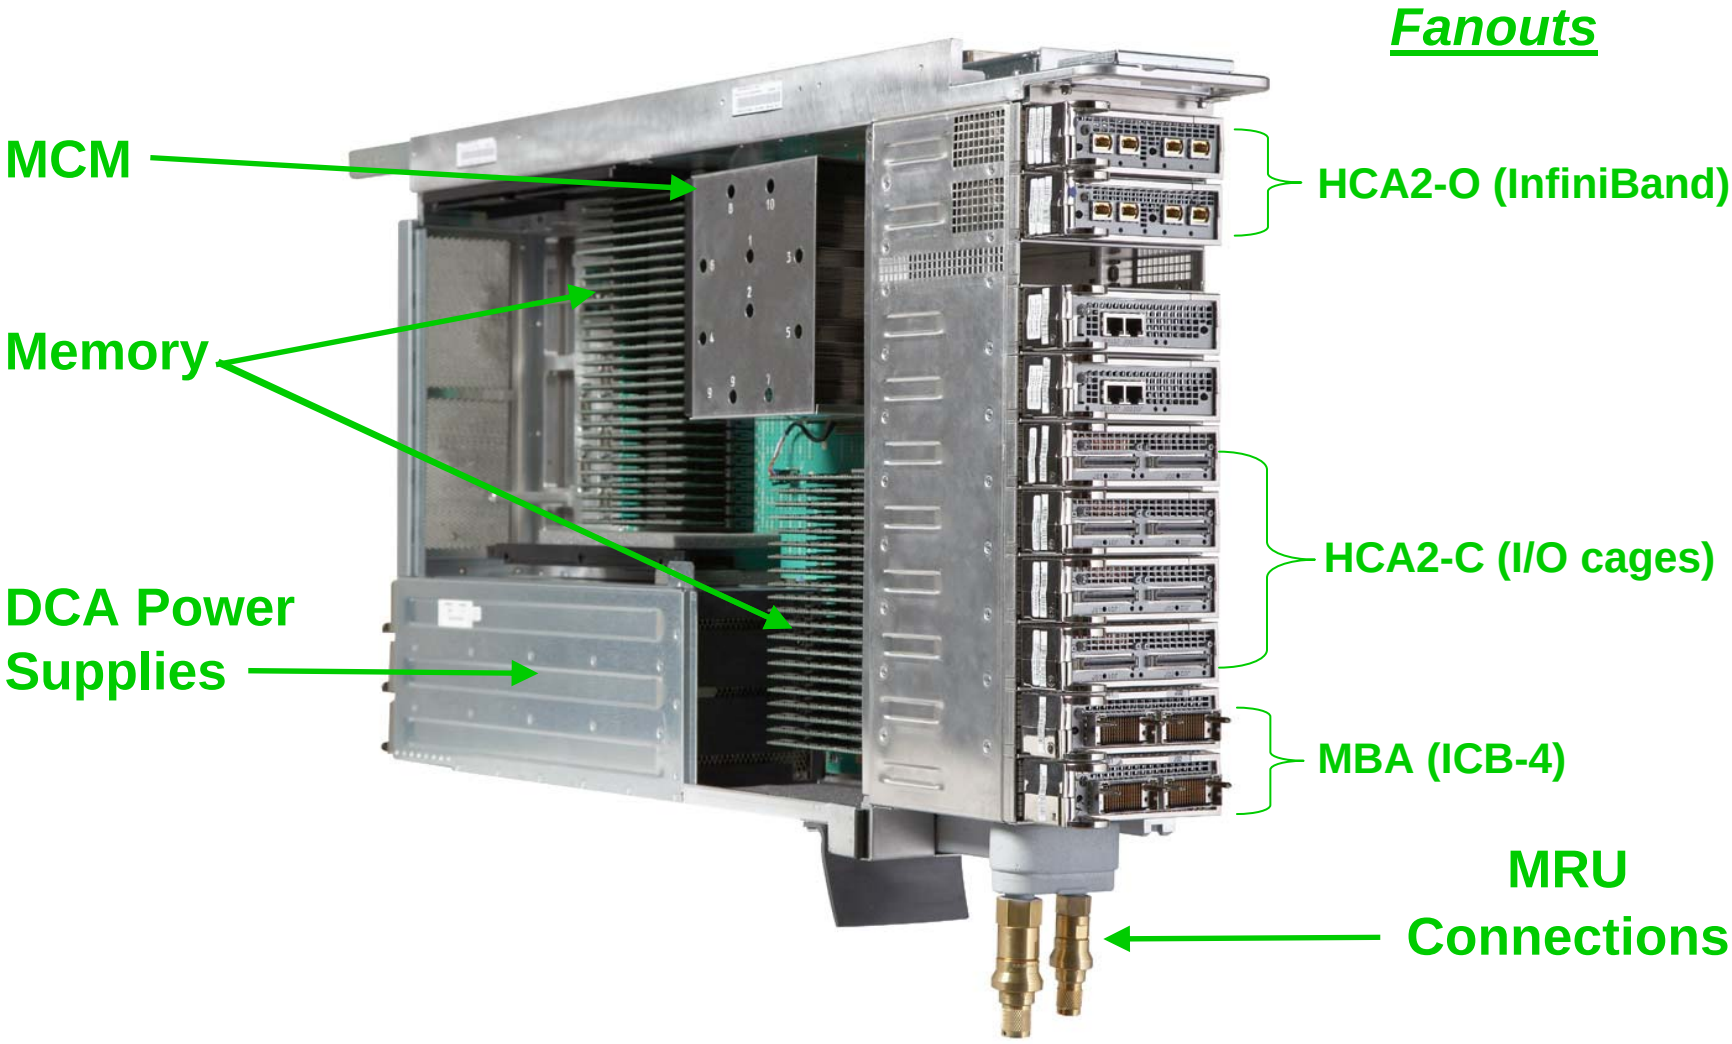

z10 EC – Under the covers (Model E56 or E64)

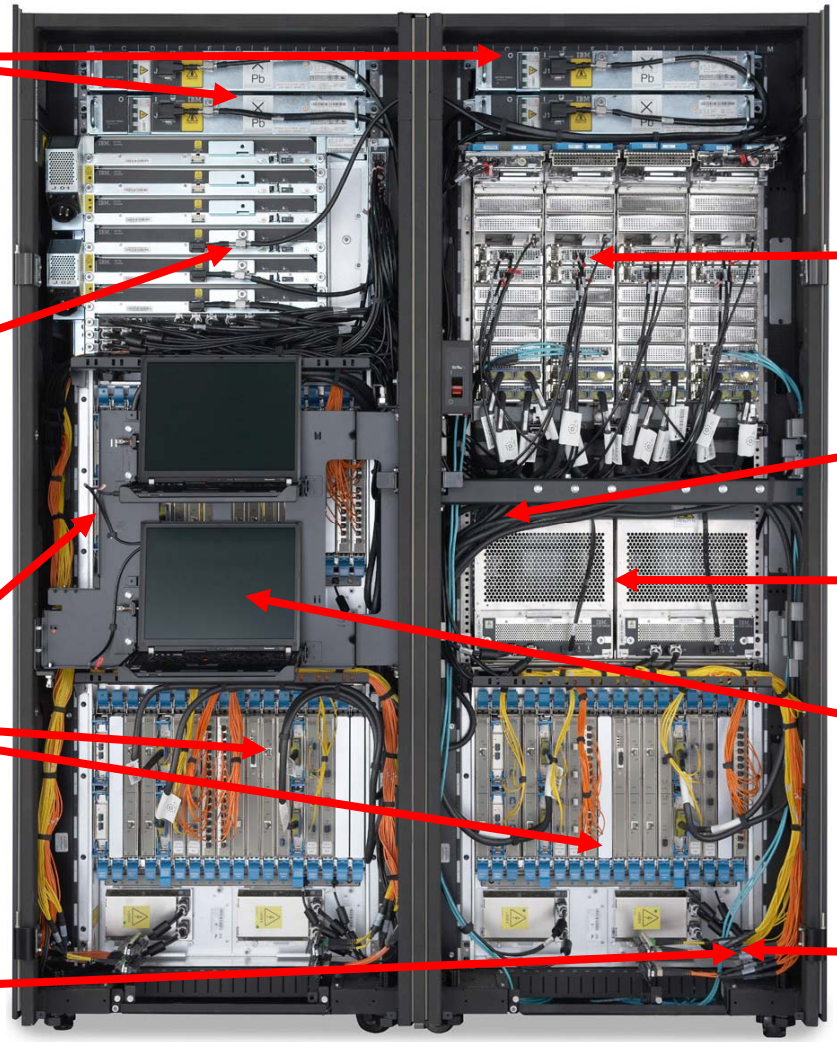

Internal Batteries (optional)

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

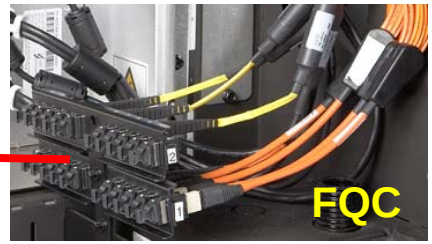

z10 EC Channel Type and Crypto Overview

- **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)
    - Layer 2 support
- **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.

Connectivity for Coupling and I/O

- **Up to 8 fanout cards per book**
  - Up to 16 ports per book
    - 48 Port System Maximum
- **Fanout cards - InfiniBand pairs dedicated to function**
  - **HCA2-C fanout – I/O Interconnect**
    - Supports all I/O, ISC-3 and Crypto Express2 cards in I/O cage domains

Up to 16 CHPIDs – across 2 ports

  - **HCA2-O fanout – InfiniBand Coupling***
    - New CHPID type – CIB for Coupling
      - Fiber optic external coupling link

2 CHPIDs – 1 per port

  - **MBA fanout *(Not available on Model E64)***
    - ICB-4
    - New connector and cables

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

SAPs, I/O Buses, Links, and I/O Connectivity

Model	Books/ PUs	Std SAPS	Opt SAPs	Maximum HCAs/ Buses	Maximum PSIFB+ICB4 Links + I/O cards	Maximum I/O Cards + PSIFB/ICB4 Links	Max FICON/ ESCON CHPIDs
E12	1/17	3	0-3	8/16	16 + 0 Cards	64 + 0 PSIFB + ICB4	256/ 960
E26	2/34	6	0-7	16/32	32 + 0 Cards ICB-4 limit 16	84 + 8 PSIFB + ICB4	336/ 1024
E40	3/51	9	0-11	20/40	32 + 32 Cards ICB-4 limit 16	84 + 16 PSIFB + ICB4	336/ 1024
E56	4/68	10	0-18	24/48	32 + 64 Cards ICB-4 limit 16	84 + 24 PSIFB + ICB4	336/ 1024
E64	4/77	11	0-21	24/48	32 + 64 Cards No ICB-4	84 + 24 PSIFB	336/ 1024

Note: Only TPF may need Opt SAPs for normal workload
Note: PSIFB and ICB4 do not reside in I/O cages
Note: Plan Ahead for up to 2 additional I/O cages
This assumes no PSC24V power sequence cards
a. 0 to 24 I/O cards – 1 cage
b. 25 to 48 I/O cards – 2 cages
c. 49 to 84 I/O cards – 3 cages

Note: Include Crypto Express2 cards in I/O card count
Limits:
a. 4 LCSSs maximum
b. 15 partitions maximum per LCSS, 60 maximum
c. 256 CHPIDs maximum per LCSS, 1024 maximum

InfiniBand glossary

Term	Description
Gbps	Gigabits per second
GBps	GigaBytes per second
1x	One “lane”, one pair of fibers
12x	12 “lanes”, 12 pairs of fiber
SDR	Single Data Rate – 2.5 Gbps per “lane” (0.3 GBps)
DDR	Double Data Rate – 5 Gbps per “lane” (0.5 GBps)
12x IB-SDR	12 “lanes” (pairs) for a total link data rate of 3 GBps, 150 meters point-to-point Used with OM3, 2000 MHz-k 50 micron multimode fiber optic cabling with MPO connectors
12x IB-DDR	12 “lanes” (pairs) for a total link data rate of 6 GBps, 150 meters point-to-point Used with OM3, 2000 MHz-k 50 micron multimode fiber optic cabling with MPO connectors
1x IB-DDR LR*	One “lane” (one pair), 5 Gbps link data rate, unrepeated distance of 10 km Used with 9 micron single mode fiber optic cabling with LC Duplex connectors

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

z9 EC and z10 EC System Structure for I/O

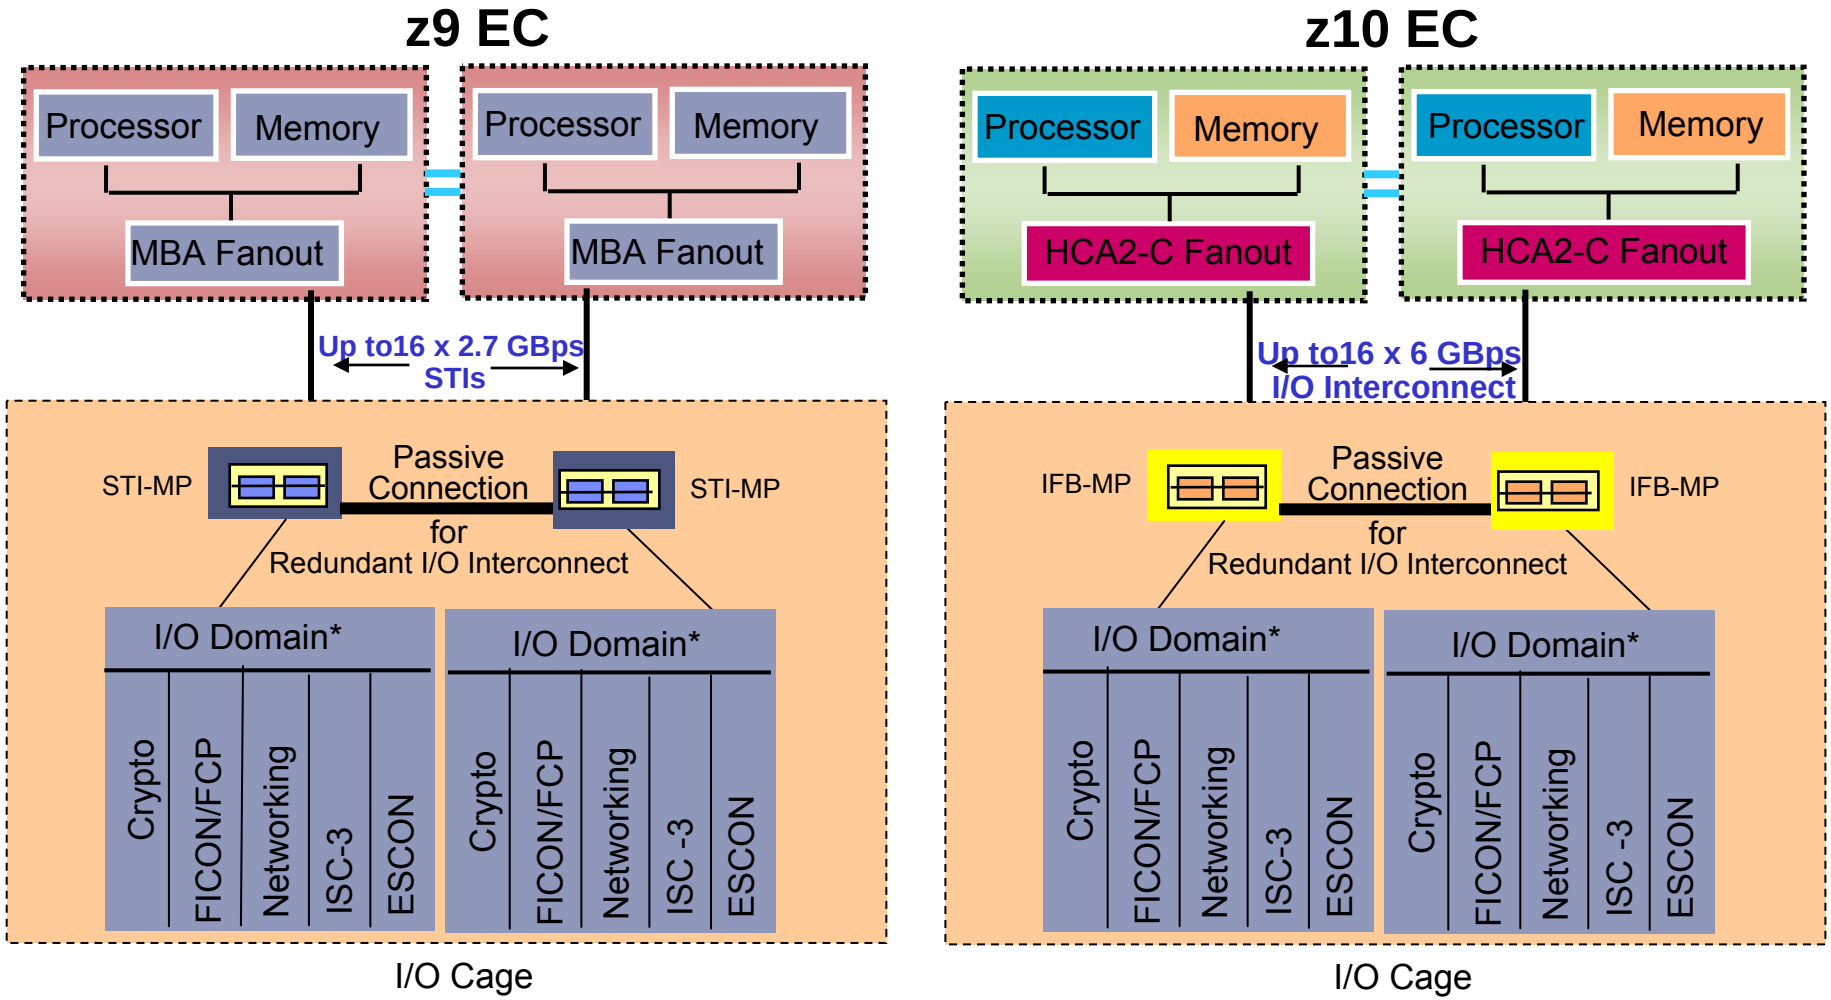

***Note: Each I/O domain supports up to 4 features only**

System z10 EC FICON Enhancements

- **Extended Distance FICON (CHPID type FC) performance enhancements**
  - Enhancement to the industry standard FICON architecture (FC-SB-3)
    - Implements a new protocol for 'persistent' Information Unit (IU) pacing that can help to optimize link utilization
    - Requires supporting Control Unit(s) (e.g. DS8000 at new level)
  - Designed to improve performance at extended distance
    - May benefit z/OS Global Mirror (previously called XRC)
    - May simplify requirements for channel extension equipment
  - Transparent to operating systems
  - Applies to FICON Express4 and Express2 channels
- **Enhancements for Fibre Channel Protocol (FCP) performance**
  - Designed to support up to 80% more I/O operations per second compared to System z9 for small block (4 kB) I/O operations on a FICON Express4 channel
  - Transparent to operating systems
  - Applies to FICON Express4 and Express2 channels (CHPID type FCP) communicating to SCSI devices. (Improvement on FICON Express2 is expected to be less than on FICON Express4)

OSA-Express3 – 10 GbE (2Q2008)

- **Double the port density compared to 10 GbE OSA-Express3**
- **Designed to Improve Performance for standard and jumbo frames**
- **10 Gigabit Ethernet LR (Long Reach)**
  - Two ports per feature
  - Small form factor connector (LC Duplex) single mode
  - CHPID type OSD (QDIO)

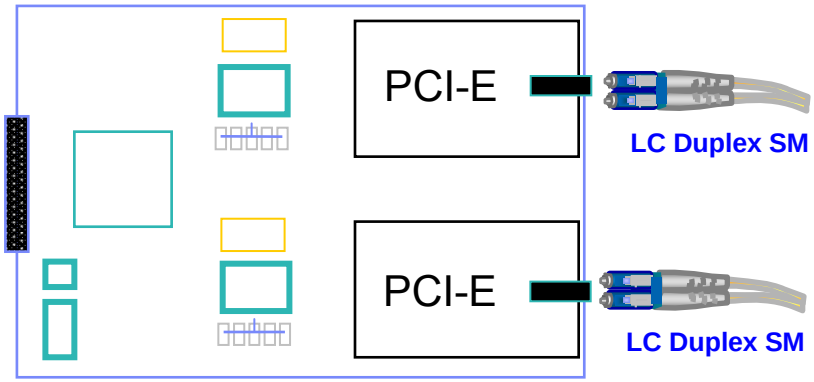

z10 EC InfiniBand PSIFB* Coupling Connectivity (2Q2008)

12x IB-DDR
6 GBps
Minimum – 0
Maximum – 32 ports
Order increment – 2 ports
Distance – 150 meters
OM3 fiber optic cables

Up to 16 CHPIDs – across 2 ports

- **Point-to-point up to 150 m (492 ft)**
- **Maximum of 16 HCA2-O fanouts**
  - 2 ports per HCA2-O fanout
  - Up to 16 CHPIDs per HCA1-O fanout
    - Distribute across 2 ports as desired
- **12x IB-DDR (6 GBps)**
  - z10 EC to z10 EC
- **12x IB-SDR (3 GBps)**
  - z10 EC to System z9 Dedicated Coupling Facility
- **OS Support::**
  - **z/OS 1.7 + zIIP Web Deliverable**
  - **z/VM 5.3 – Dynamic I/O Suppoort**

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

System z9 InfiniBand PSIFB* Coupling Connectivity (2Q2008 – Dedicated CF only; SOD – Any z9)

12x IB-SDR
3 GBps
Minimum – 0
Maximum – 16 ports
Order increment – 2 ports
Distance – 150 meters
OM3 fiber optic cables

Up to 16 CHPIDs – across 2 ports

- **Point-to-point up to 150 m (492 ft)**
- **Maximum of 8 HCA1-O fanouts**
  - 2 ports per HCA1-O fanout
  - Up to 16 CHPIDs per HCA1-O fanout
    - Distribute across 2 ports as desired
- **12x IB-SDR (3 GBps)**
  - z10 EC to System z9 Dedicated Coupling Facility
- **OS Support for non-dedicated CFs**
  - Support: z/OS 1.7 + zIIP Web Deliverable
  - Dynamic I/O configuration to define, modify and query a CHPID when z/VM 5.3 is the controlling LPAR for I/O

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

System z – Supported Coupling Links

System	PSIFB** (2Q2008)	ICB-4	ICB-3	ISC-3	IC	Max # Links
z10 EC	32*	16* Except E64	N/A	48	32	64
z9 Dedicated CF	16	16	16	48	32	64
Any z9	SOD**	16	16	48	32	64
z990	N/A	16	16	48	32	64
z890	N/A	8	16	48	32	64

* Maximum of 32 PSIFB + ICB4 links on System z10 EC. ICB-4 not supported on Model E64.

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

z10 Parallel Sysplex Coexistence and Coupling Connectivity

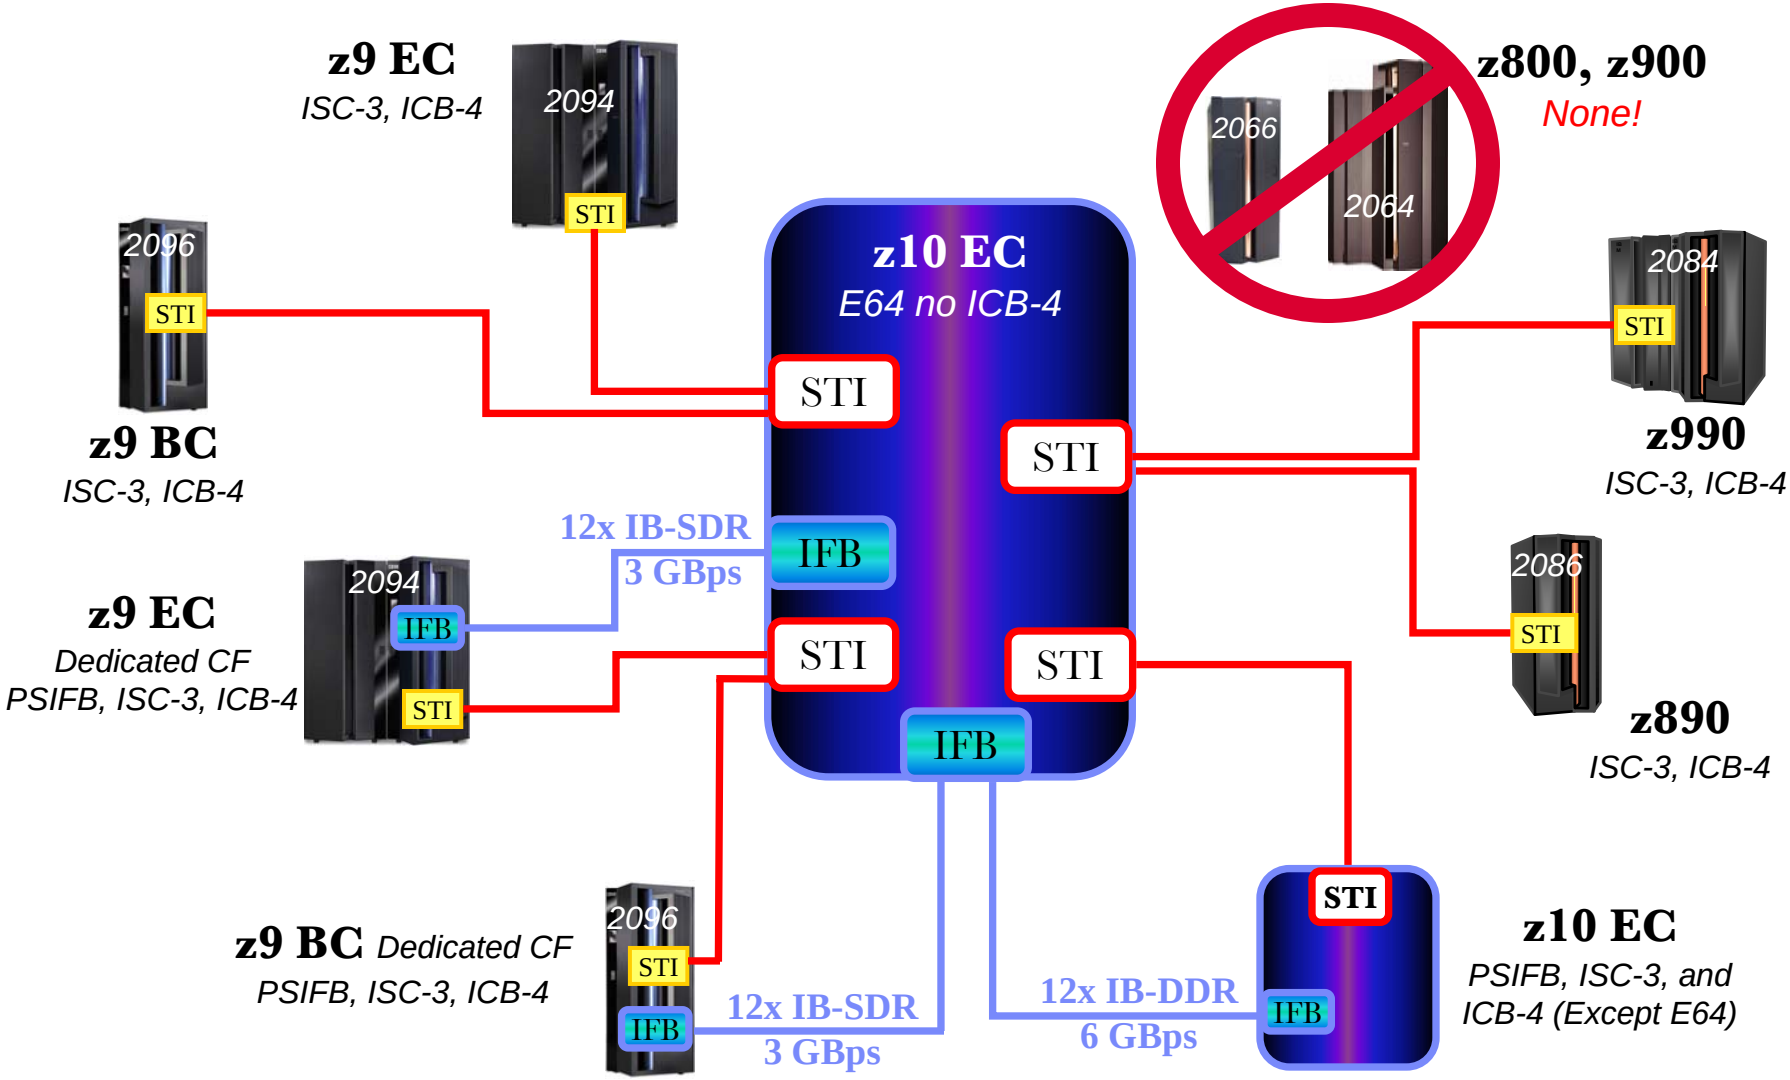

Machine Types and Coupling Technologies Planned to be Supported by z10 EC Servers

z10 EC Physical Planning

Top of machine must be clear
to allow backup cooling airflow.

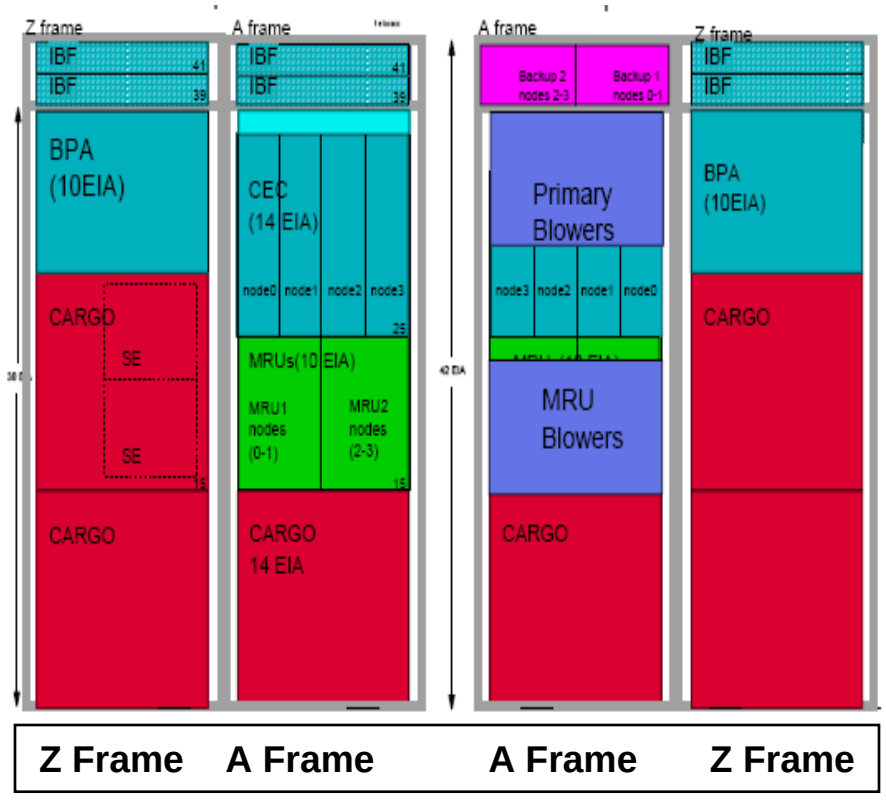

Dimension	z10 EC	z9 EC	Change z9 EC to z10
Height Full/Red	79.3/72.1 in 2015/1832 mm	76.4/72.1 in 1941/1832 mm	+2.9/+0 in +74/+0 mm
Width	61.7 in 1568 mm	61.6 in 1565 mm	+.1 in +3 mm
Depth w/o Covers	50 in 1270 mm	46.1 in 1171 mm	+3.9 in +99 mm
Depth with Covers	71 in 1803 mm	62.1 in 1577 mm	+8.9 in +226 mm

Electrical Service Requirements:

NEW!

	1 I/O cage	2 I/O cage	3 I/O cage
1 book	2x60A	2x60A	2x60A
2 book	2x60A	4x60A	4x60A
3 book	4x60A	4x60A	4x60A
4 book	4x60A	4x60A	4x60A

Same power plugs/service as z990 and z9 EC, but large configurations need 4.

15% better performance/kWh than z9 EC

Check the latest IMPP!

System z10 EC Exploitation: *z/OS Support Summary*

Release	z10 GA1 Support	InfiniBand Coupling Links	65535 MP Factors	Crypto Toleration	Crypto Exploitation	HiperDispatch	Decimal Floating Point **	Large Memory > 128GB (1TB)	Greater than 54 CPs (64)	Large Page Support	Capacity Provisioning	RMF FICON Enhancements
z/OS 1.7	P	P	P	W P	W	N	P	N	N	N	N	N
z/OS 1.7 w/zIIP	P	P	P	W P	W	P	P	N	N	N	N	N
z/OS 1.8	P	P	P	W P	W	P	P	B	N	N	N	N
z/OS 1.9	P	B P	B	W P	W	B P	B P	B	B	B P	B P	B

Legend
B – FMID in Base product (assumes service identified in z9 EC PSP Bucket is installed)
W – FMIDs shipped in a Web Deliverable
P – PTFs required
N – Not Supported

**** Level of decimal floating-point exploitation will vary by z/OS release and PTF level.**

z/VSE & z/VM Support Summary

		z890 (WdfM)	z990	z9 EC	z9 BC	z10	End of Market	End of Service	Ship Date
z/VSE*	3.1	x	x	x	x	x	5/08	TBD	3/05
	4.1	x	x	x	x	x	TBD	TBD	3/07
z/VM	5.2	x	x	x	x	x	6/07	4/09**	12/05
	5.3	x	x	x	x	x	TBD	9/10**	6/07

Note: z/VM requires Compatibility Support which allows z/VM to IPL and operate on the z10 providing z9 functionality for the base OS and Guests

*z/VSE V3 can execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities.
z/VSE V3 is designed to exploit select features of IBM System z9 and zSeries hardware.
Note: z/VSE V4 is designed to exploit 64 bit real memory addressing, but will not support 64-bit virtual memory addressing

Linux on System z – Plans for z10 EC

Program Support

- IBM Systems Director Active Energy Manager (AEM) for Linux on System z

Compatibility

- Existing Linux on System z distributions* (most recent service levels):
  - Novell SUSE SLES9
  - Novell SUSE SLES10
  - Red Hat RHEL4
  - Red Hat RHEL5

Exploitation*

- IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases or versions for:
  - Capacity Provisioning
  - Large Page Support
  - CPACF Enhancements
  - Dynamic Change of Partition Cryptographic Coprocessors
  - HiperSockets Layer 2 Support

***For latest information and details contact your Linux distributor.**

IBM System z10 EC - TPF and z/TPF Support

		z890	z990	z9 EC	z9 BC	z10	End of Service	Ship Date
TPF	4.1	x[c]	x[c]	x[c]	x[c]	x[c]	TBD	2/01
z/TPF	1.1	x	x	x	x	x	TBD	9/05

- **z/TPF Migration Portal**
  - http://www.**ibm.com**/tpf/ztpfmigration
- **A PRPQ for HLASM running on Linux on z is available**
  - z/TPF uses the GNU Cross Compiler (GCC) running under Linux for System z

x[c] – Supports up to 30 LPARs with PJ29309

Any Questions?

Tom_Russell@ca.ibm.com