



# IBM System z10

## Introduction and Hardware Overview



Dr. Stephan Kammerer  
IBM Germany Research & Development

August 2010

[stephan.kammerer@de.ibm.com](mailto:stephan.kammerer@de.ibm.com)



# Trademarks

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

AlphaBlox*	GDPS*	Processor Resource/Systems Manager	System z10
APPN*	HiperSockets	RACF*	System zEnterprise z196
CICS*	HyperSwap	Redbooks*	System/30
Cool Blue	IBM*	Resource Link	Tivoli*
DB2*	IBM eServer	RETAIN*	Tivoli Storage Manager
DFSMS	IBM logo*	REXX	TotalStorage*
DFSMSshsm	IMS	RMF	VSE/ESA
DFSMSrmm	Language Environment*	S/370	VTAM*
DFSORT*	Lotus*	S/390*	WebSphere*
DirMaint	Multiprise*	Scalable Architecture for Financial Reporting	xSeries*
DRDA*	MVS	Sysplex Timer*	z9*
DS6000	OMEGAMON*	Systems Director Active Energy Manager	z10
DS8000	Parallel Sysplex*	System p*	z10 BC
ECKD	Performance Toolkit for VM	System Storage	z10 EC
ESCON*	POWER6	System x	z/Architecture*
FICON*	POWER7	System z	z/OS*
FlashCopy*	PowerPC*	System z9*	z/VM*
	PR/SM		z/VSE
			zSeries*

\* Registered trademarks of IBM Corporation

## The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

### Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.



# Agenda

## **System z10 EC**

**Overview**

**Architecture**

**MCM**

**PU Chip**

**Accelerator**

**SC Chip**

**Book Layout**

**Connectivity**

**Inter Book Connectivity**

**IO Connectivity**

**Ethernet Connectivity**

**Crypto**

**Capacity on Demand**

**Green IT**

**High Availability**

**Operating Systems**

## **System z10 BC**

**Overview**

**Architecture**

## **zEnterprise System z196**

**Overview**

**Architecture**

**Hardware Development**

**Firmware Development**



# Agenda

## **System z10 EC**

### **Overview**

### **Architecture**

**MCM**

**PU Chip**

**Accelerator**

**SC Chip**

**Book Layout**

### **Connectivity**

**Inter Book Connectivity**

**IO Connectivity**

**Ethernet Connectivity**

**Crypto**

**Capacity on Demand**

**Green IT**

**High Availability**

**Operating Systems**

## **System z10 BC**

### **Overview**

### **Architecture**

## **zEnterprise System z196**

### **Overview**

### **Architecture**

**Hardware Development**

**Firmware Development**






## Von 'System/360' (S/360) zu ESA/390 und z/-Architektur

- **1964 S/360**
  - CISC, 24bit Adressierung, 'Real Storage', Uniprozessoren
  - Amdahl, G.M., Blaauw, G.A., and Brooks, F.P.: *Architecture of the IBM System/360*
- **1971 S/370**
  - 'Virtual Storage', Multiprozessor-Unterstützung, ...
- **1981 S/370 XA (Extended Architecture)**
  - 31bit Adressierung (2GB), 'Expanded Storage' (>2GB), 'Channel Subsystem'
  - 'Interpretive Execution': Basis für Logische Partitionierung ('LPAR')
- **1988 ESA/370**
  - ESA = Enterprise Systems Architecture, Logische Partitionierung
  - Ausbau der Speicher-Zugriffsmethoden: Mehr als ein 'address space'
- **1990 ESA/390**
  - 'ESCON' (Enterprise Systems Connection Architecture) Glasfasertechnologie ...
  - Datenkompression, Kryptographie, LPAR Erweiterungen
- **1994 Parallel Sysplex, Übergang von Bipolar zu CMOS Technologie**
  - 'Coupling Facility', Cluster von bis zu 32 x 16-way MultiProzessoren
  - 'FICON' (Fiber Channel Connectivity), Ausbau der Glasfasertechnologie
- **2000 z/-Architektur (64-bit), z900, z800, z990(2003), z890(2004), z9 EC(2005), z9 BC(2006), z10 (2008), zEnterprise (2010)**



# IBM z10 EC Continues the CMOS Mainframe Heritage



- **G4** – 1<sup>st</sup> full-custom CMOS S/390®
- **G5** – IEEE-standard BFP; branch target prediction
- **G6** – Copper Technology (Cu BEOL)

- **z900** – Full 64-bit z/Architecture®
- **z990** – Superscalar CISC pipeline
- **z9 EC** – System level scaling

- **z10 EC** – Architectural extensions



# Do GHz matter?

- **GHz does matter**

- It is the "rising tide that lifts all boats"
- It is especially important for CPU-intensive applications

- **GHz is not the only dimension that matters**

- System z focus is on balanced system design across many factors
  - Frequency, pipeline efficiency, energy efficiency, cache / memory design, I/O design

- **System performance is not linear with frequency**

- Need to use LSPR+ System z capacity planning tools for real client / workload sizing

- **System z has been on consistent path while others have oscillated between extremes**

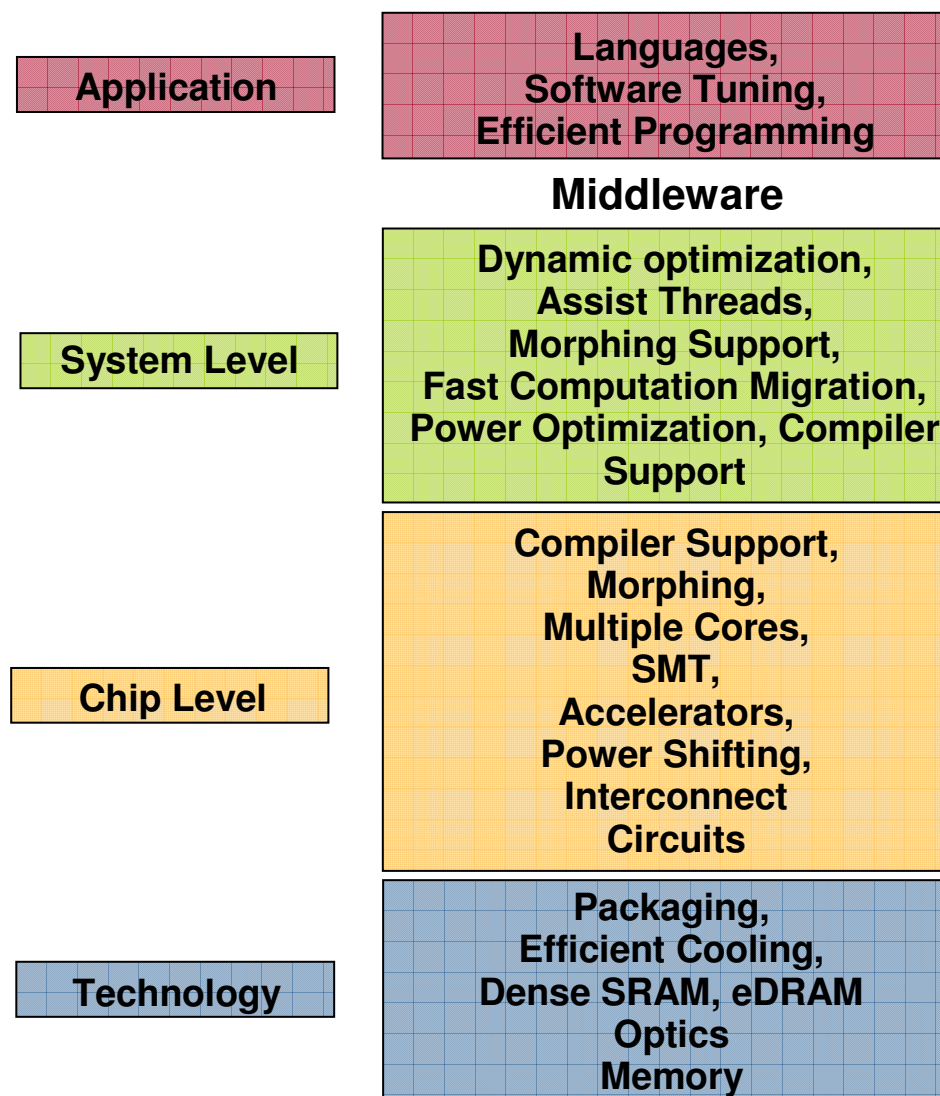
- Growing frequency steadily, with occasional jumps/step functions (G4 in 1997, z10 in 2008)

- **z10 leverages technology to get the most out of high-frequency design**

- Low-latency pipeline
- Dense packaging (MCM) allows MRU cooling which yields more power-efficient operation
- Virtualization technology (etc.) allows consistent performance at high utilization, which makes CPU power-efficiency a much smaller part of the system/data-center power consumption picture



## Industry's Approach to Integrated Systems Performance



- Microprocessor frequency will no longer be the dominant driver of system level performance
- Scale-out and small SMPs will continue to out pace scale-up growth
- Systems will increasingly rely on modular components for continued performance leadership
- Systems will be designed with the ability to dynamically manage and optimize power
- Integration over the entire stack, from semiconductor technology to end-user applications, will replace scaling as the major driver of increased system performance



# Agenda

## **System z10 EC**

Overview

**Architecture**

**MCM**

**PU Chip**

**Accelerator**

**SC Chip**

**Book Layout**

**Connectivity**

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

Green IT

High Availability

Operating Systems

## **System z10 BC**

Overview

Architecture

## **zEnterprise System z196**

Overview

Architecture

Hardware Development

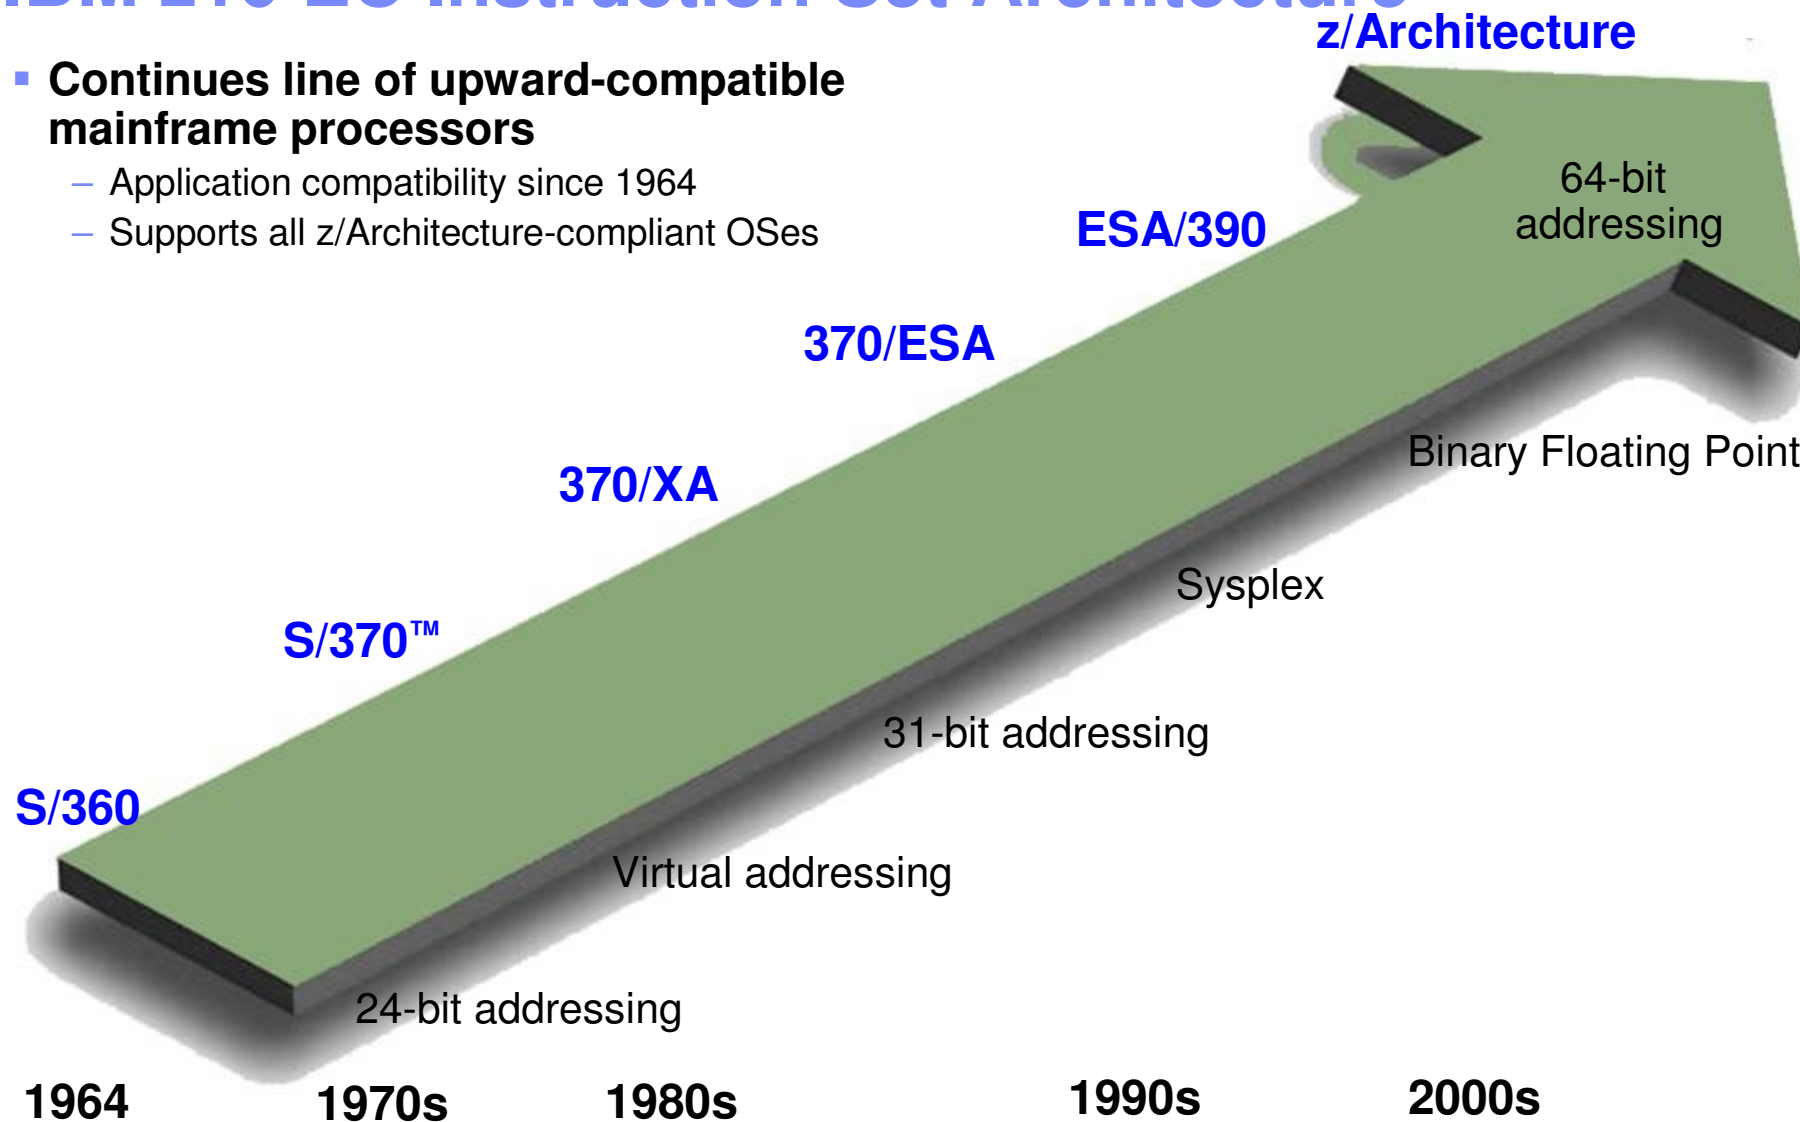
Firmware Development



# IBM z10 EC Instruction Set Architecture

- **Continues line of upward-compatible mainframe processors**

- Application compatibility since 1964
- Supports all z/Architecture-compliant OSes






## 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 1000's 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 1 MB page frames
  - Full hardware support for Hardware Decimal Floating-point Unit (HDFU)


# z10 EC Chip Relationship to POWER6™

- **New Enterprise Quad Core z10 EC processor chip**
- **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 System z and System p® designers and engineers working together
- **Different personalities**
  - Very different Instruction Set Architectures (ISAs)
    - very different cores
  - Cache hierarchy and coherency model
  - SMP topology and protocol
  - Chip organization
  - IBM z10 EC Chip optimized for Enterprise Data Serving Hub

**Enterprise Quad Core z10 processor chip**



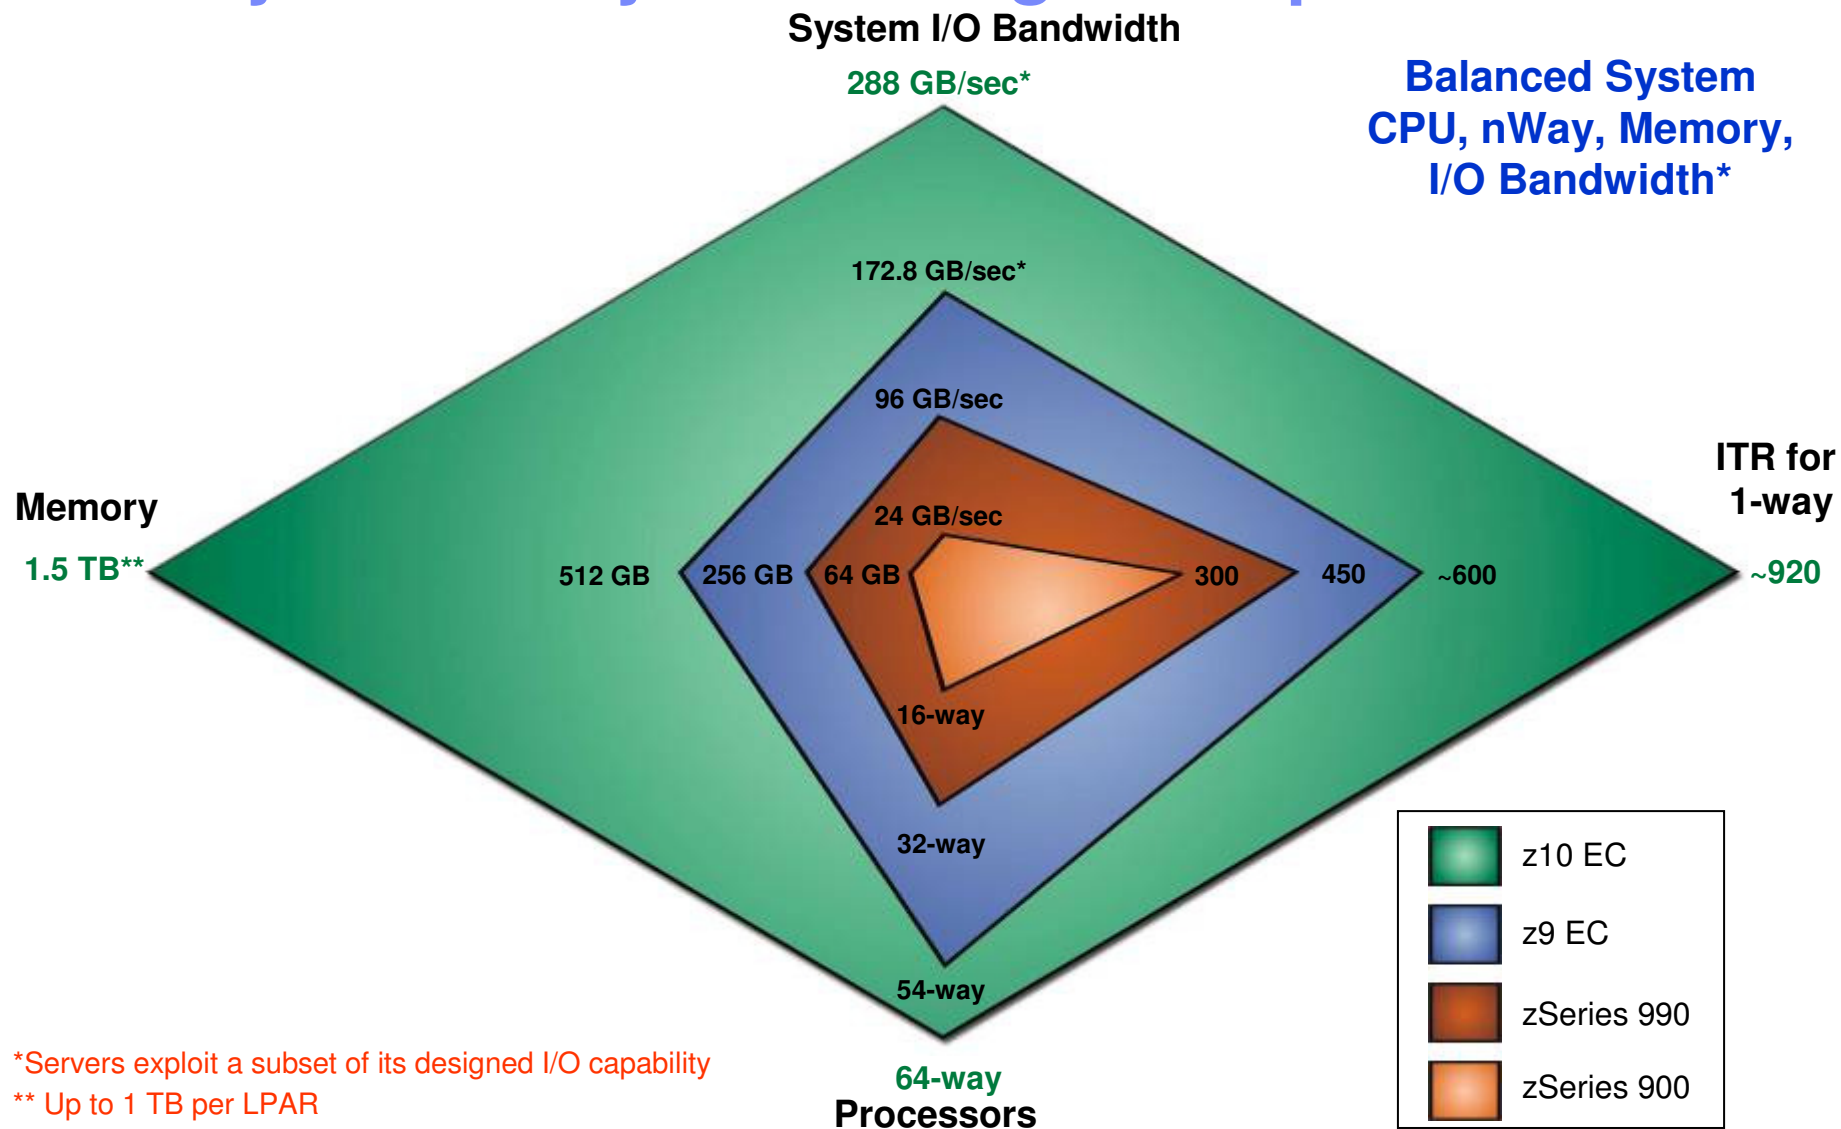
**Dual Core POWER6 processor chip**







# IBM System z: System Design Comparison

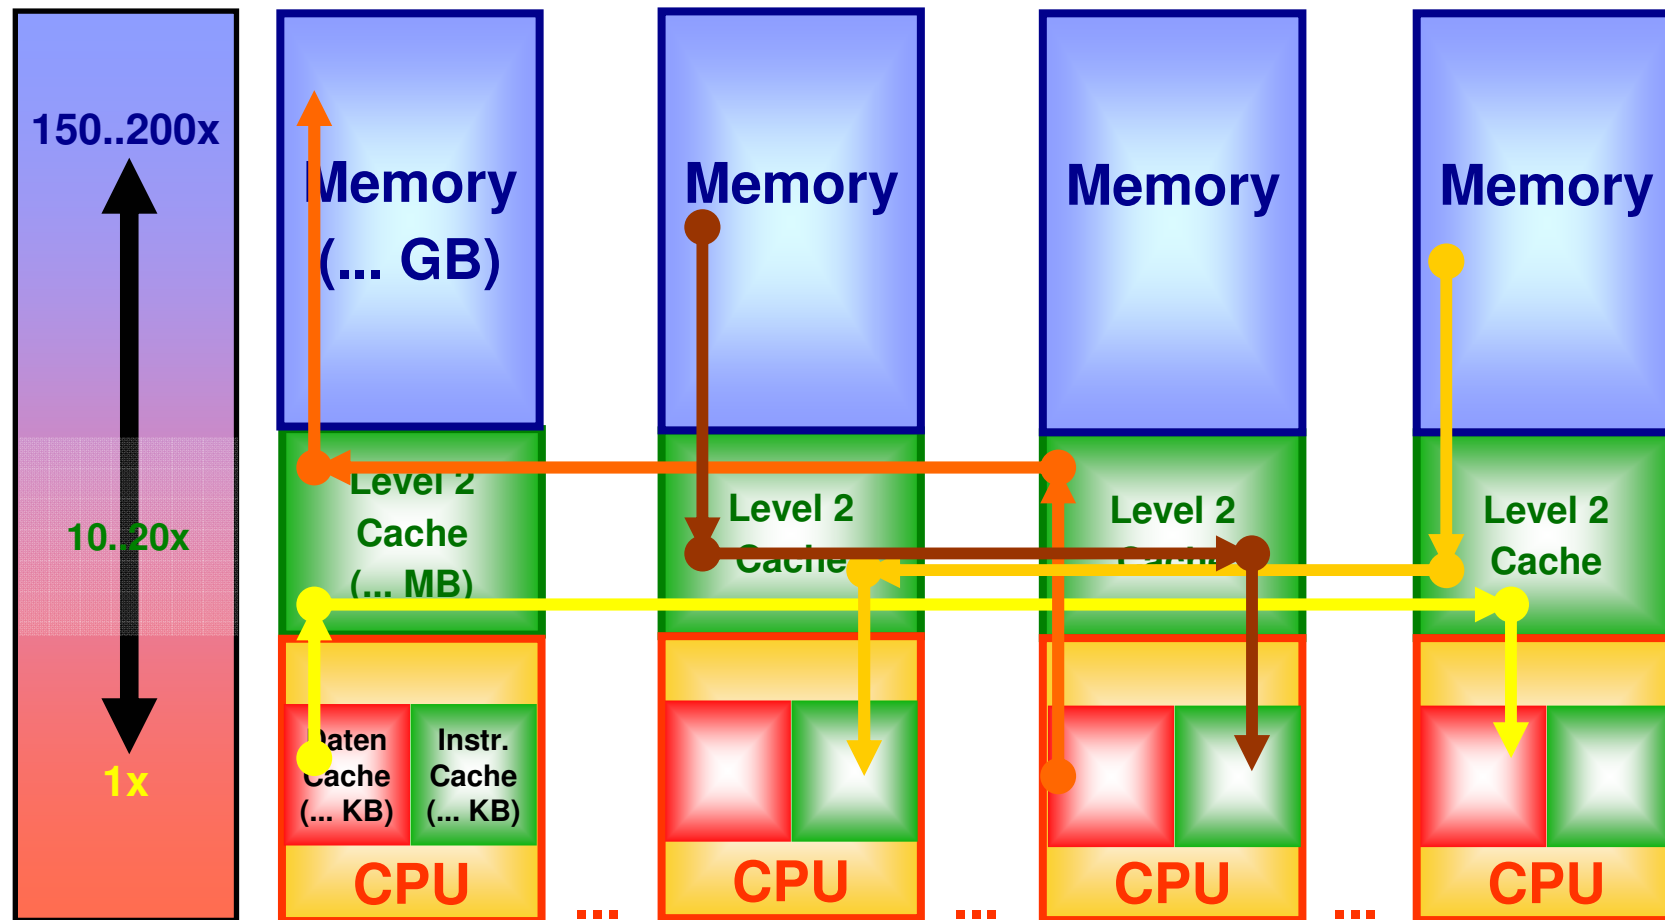


\*Servers exploit a subset of its designed I/O capability

\*\* Up to 1 TB per LPAR

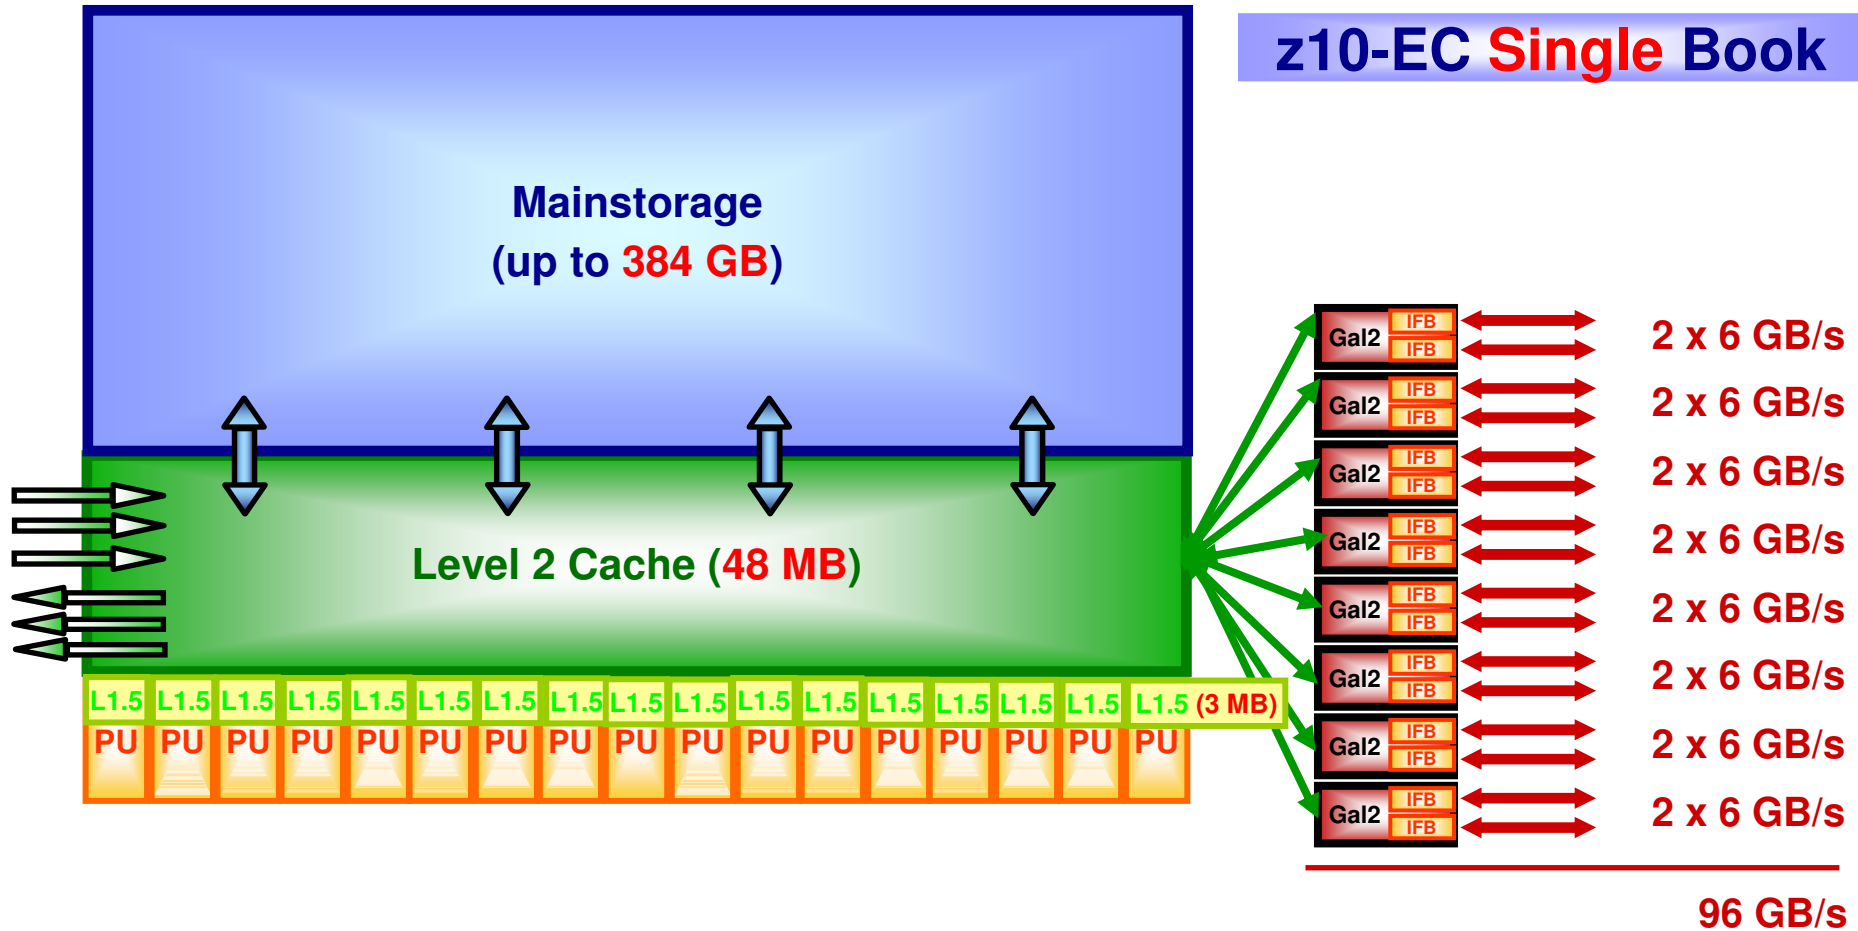
## Scalability: System-Structures optimized for data

The key problem of current microprocessor-systems:  
**Memory access does not scale with CPU-cycletime !**



# z10-EC Systemstructure:

## z10-EC Single Book




## z10 EC Overview



- **Machine Type**
  - 2097
- **5 Models**
  - E12, E26, E40, E56 and E64
- **Processor Units (PUs)**
  - 17 (17 and 20 for Model E64) PU cores per book
  - Up to 11 SAPs per system, standard
  - 2 spares designated per system
  - Dependant on the H/W model - up to 12, 26, 40, 56 or 64 PU cores available for characterization
    - Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z10 Application Assist Processors (zAAPs), System z10 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 for System and up to 1 TB per LPAR
    - 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 Feature, standard**

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






# 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/PU chip
  - L1 cache/PU core
    - 64 KB I-cache
    - 128 KB D-cache
  - L1.5 cache/PU core
    - 3 MB
  - 4.4 GHz
  - 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 – Enterprise Quad Core z10 PU Chip



- **Up to Four cores per PU**
  - 4..4 GHz
  - L1 cache/PU core
    - 64 KB I-cache
    - 128 KB D-cache
  - 3 MB L1.5 cache/PU core
  - Each core with its own Hardware Decimal Floating Point Unit (HDFU)
- **Two Co-processors (COP)**
  - Accelerator engines
    - Data compression
    - Cryptographic functions
  - Includes 16 KB cache
  - Shared by two cores
- **L2 Cache interface**
  - Shared by all four cores
  - Even/odd line (256B) split
- **I/O Bus Controller (GX)**
  - Interface to fanout
  - Compatible with System z9 MBA
- **Memory Controller (MC)**
  - Interface to controller on memory DIMMs

# Evolution of System z Specialty Engines

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



**Internal Coupling Facility (ICF) 1997**



**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 (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.



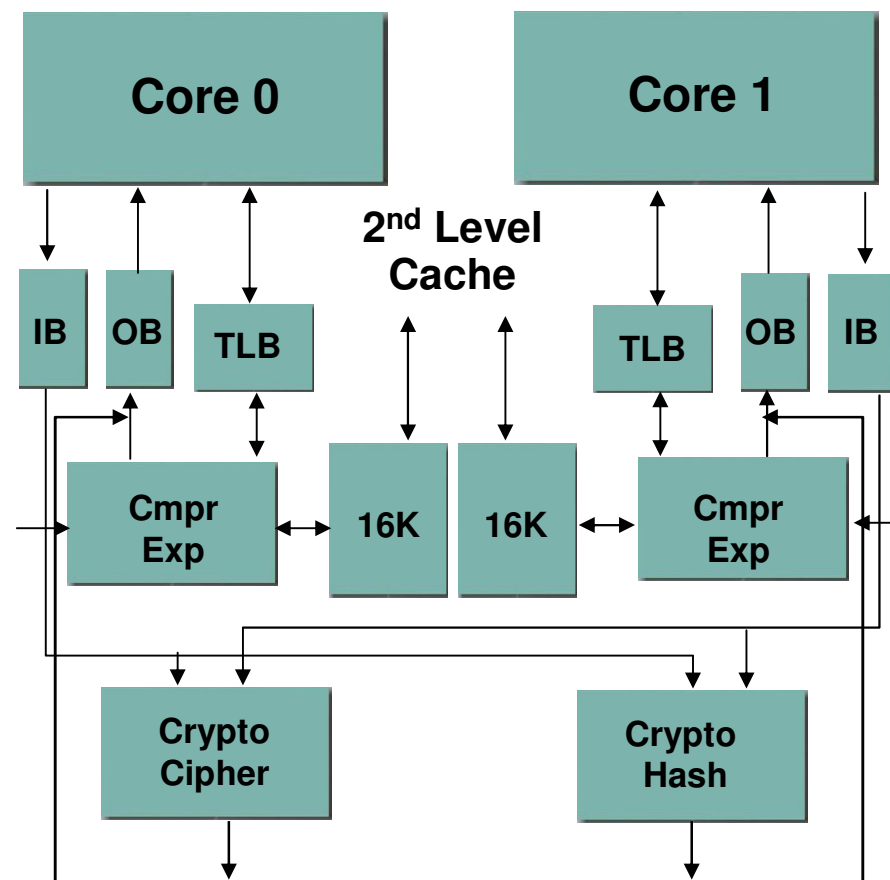


# Processing Units

- **Modern processor offer many specialist processor types. Most of the additional types are not counted as full processors for software charging**
  - Central Processor (CP) - A full z/OS processor
  - System Assistance Processor (SAP) - Used for the I/O subsystem – each machine has at least one
  - Integrated Facility for Linux (IFL) - Special processor for Linux - optional
  - zAAP – Used for JAVA code and XML – optional
  - zIIP – Used for DB2 processing, XML, IPsec & Global Mirror– optional
  - Integrated Coupling Facility (ICF) - For coupling facilities
  - Spares
- **SAPs, IFLs, ICFs, zIIPs, and zAAPs are not counted in the model number and or against software costs**

# z10 Compression and Cryptography Accelerator

- **Data compression engine**
  - Static dictionary compression and expansion
  - Dictionary size up to 64 KB (8K entries)
    - Local 16 KB caches for dictionary data
- **CP Assist for Cryptographic Function (CPACF)**
  - DES (DEA, TDEA2, TDEA3)
  - SHA-1 (160 bit)
  - SHA-2 (224, 256, 384, 512 bit)
  - AES (128, 192, 256 bit)
  - PRNG
- **Accelerator unit shared by two cores**
  - Independent compression engines
  - Shared cryptography engines



# z10 Hardware Decimal Floating Point Unit



- **Decimal arithmetic widely used in commercial and financial applications**
  - Computations often handled in software
  - Avoids rounding and other problems with binary/decimal conversions
- **On IBM System z9 delivered in millicode** – brought improved precision and function
- **On IBM System z10 integrated on every core** – giving a performance boost to execution of decimal arithmetic
- **Growing industry support for hardware decimal floating point standardization**
  - Open standard definition led by IBM, endorsed by key ISVs including Microsoft and SAP
  - Java BigDecimal, C#, XML, C/C++, GCC, DB2 V9, Enterprise PL/1, Assembler
- **z/OS V1.9 Hardware Decimal Floating Point support requires:**
  - High Level Assembler (z/OS V1.8)
  - Enterprise PL/1
  - XL C/C++ with PTF
  - Debug tool (in support of C/C++, PL/1, and HLASM)
  - dbx (in support of C/ C++)
  - DB2 9 for z/OS (allows you to define DFP data in DB2)

*Bringing high performance computing benefits to commercial workloads*




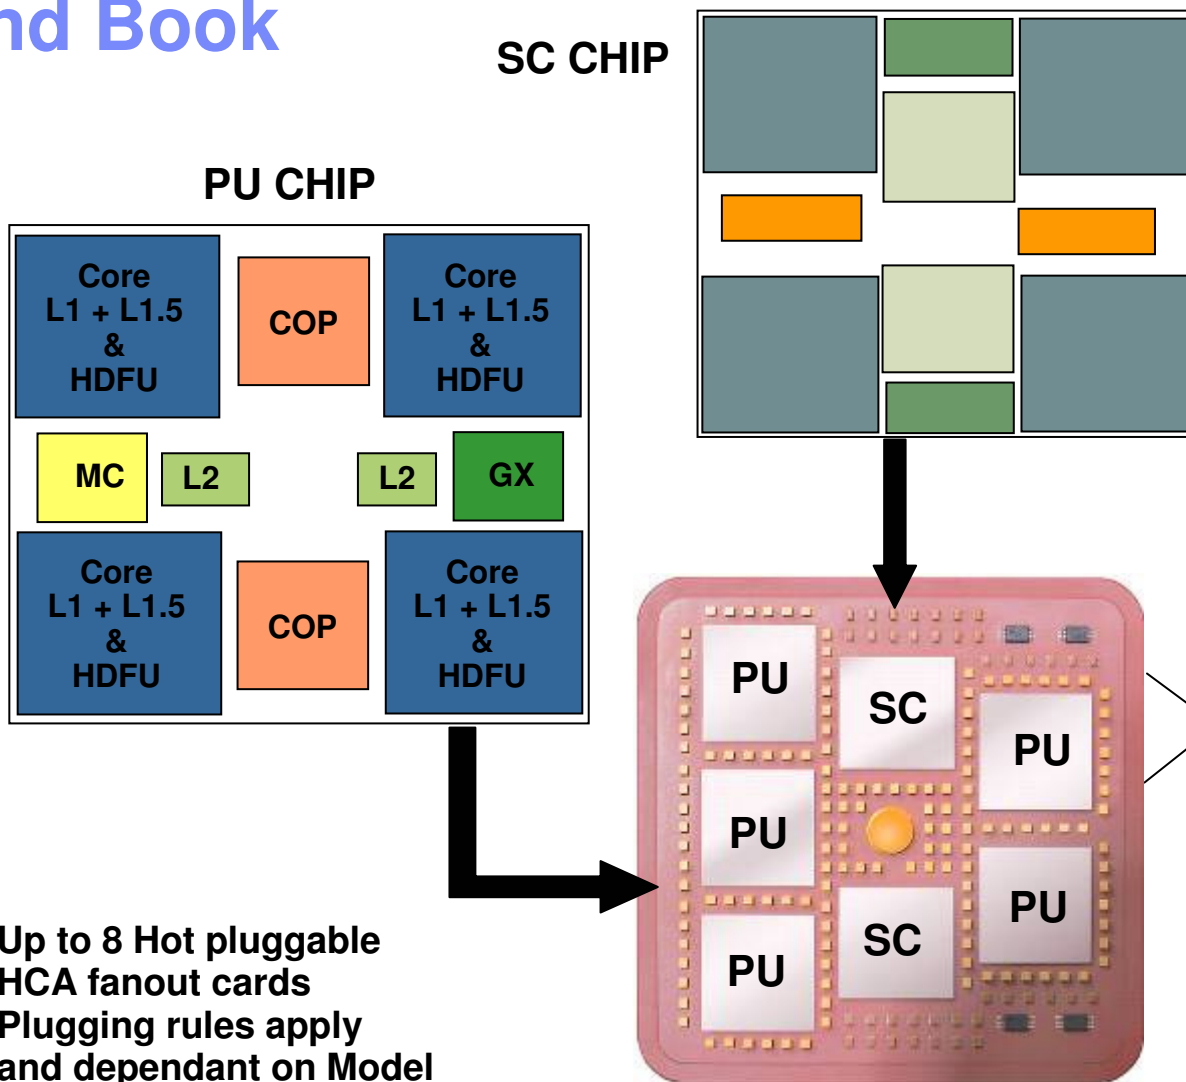
Single PU core

# z10 EC SC Hub Chip



- **Connects multiple z10 PU chips**
  - 48 GB/Sec bandwidth per chip
- **Shared Level 2 cache**
  - 24 MB SRAM Cache
  - Extended directory
    - Partial-inclusive discipline
  - Hub chips can be paired
    - 48 MB shared cache
- **Low-latency SMP coherence fabric**
  - Robust SMP scaling
  - Strongly-ordered architecture
- **Multiple hub chips/pairs allow further SMP scaling**

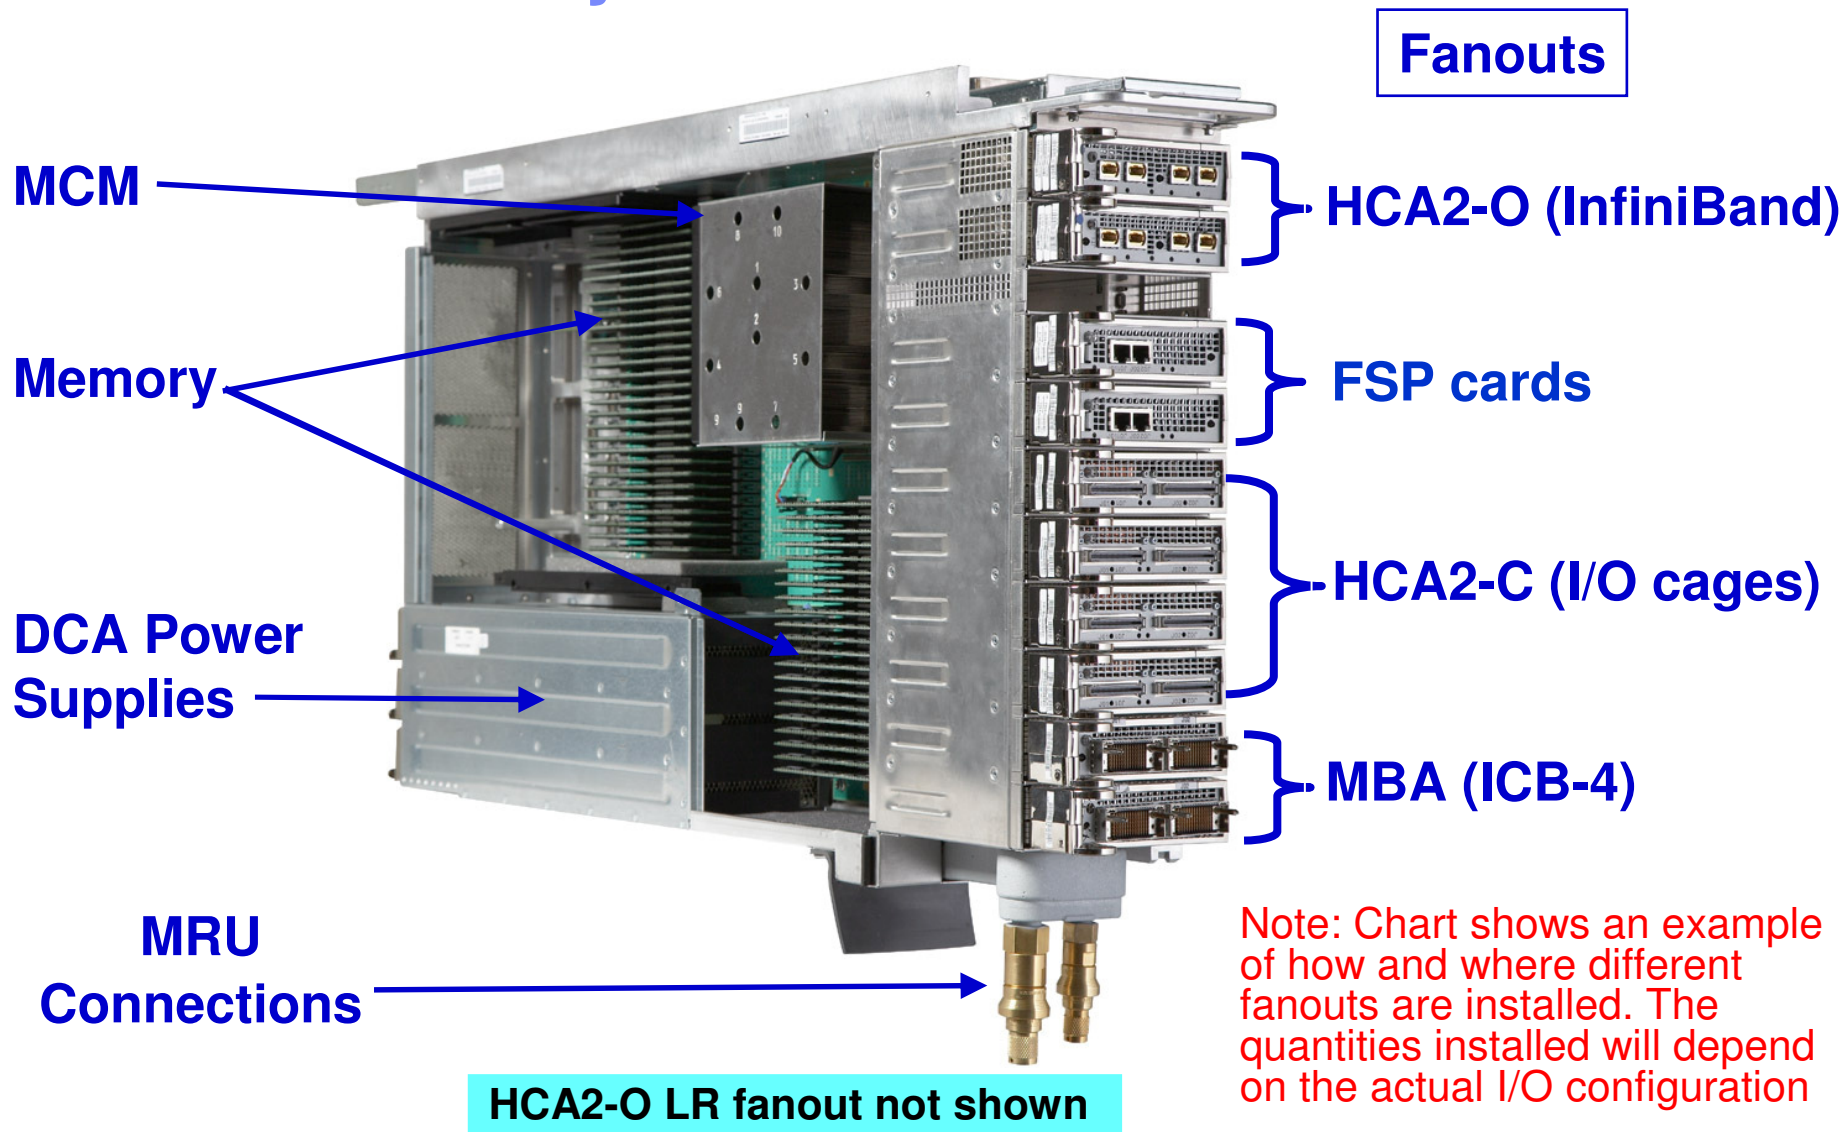
# z10 EC Processor/Memory/HCA and Book



- Up to 8 Hot pluggable HCA fanout cards
- Plugging rules apply and dependant on Model

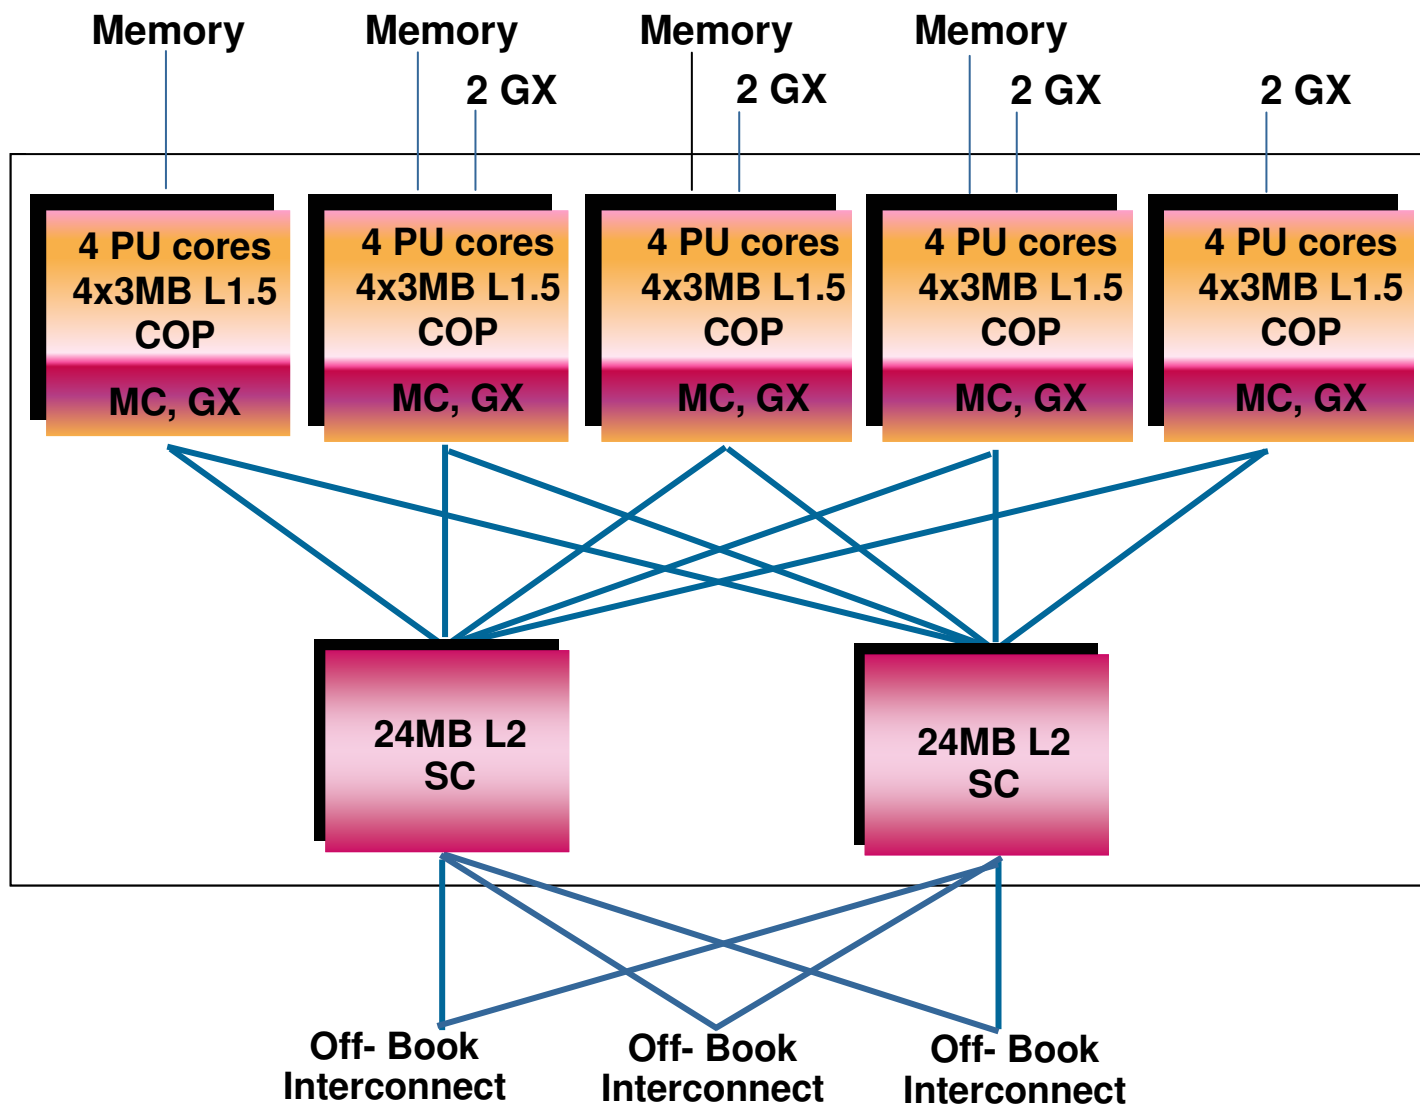


# z10 EC Book Layout – Under the covers





# 20 PU MCM Structure





# Agenda

## **System z10 EC**

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

**Connectivity**

**Inter Book Connectivity**

**IO Connectivity**

**Ethernet Connectivity**

**Crypto**

Capacity on Demand

Green IT

High Availability

Operating Systems

## **System z10 BC**

Overview

Architecture

## **zEnterprise System z196**

Overview

Architecture

Hardware Development

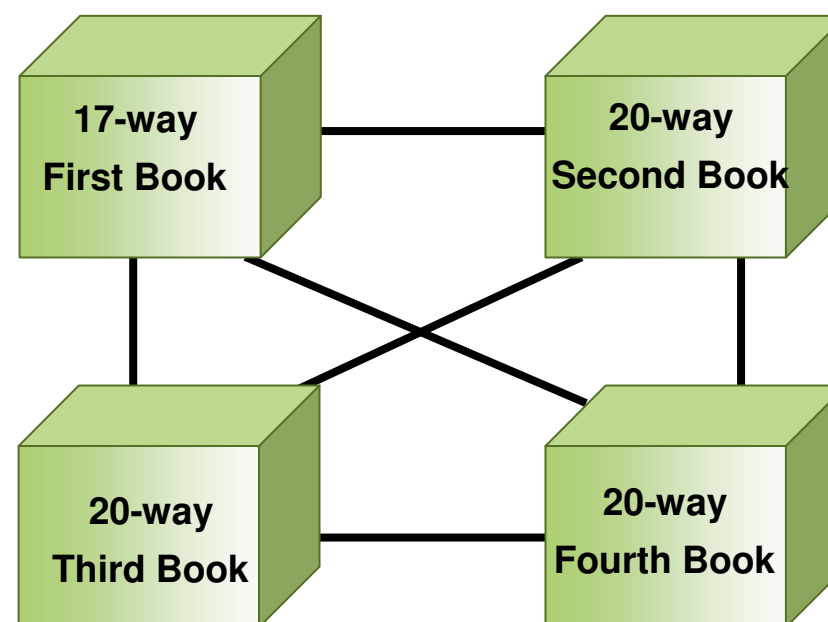
Firmware Development



## z10 EC – Inter Book Communications – Model E64

- The z10 EC Books are fully interconnected in a point to point topology as shown in the diagram
- Data transfers are direct between Books via the Level 2 Cache chip in each MCM.
- Level 2 Cache is shared by all PU chips on the MCM

### 77-way CEC





# Connected to your world

- Improved performance and flexibility for connectivity
- Broad set of options to meet your needs
- Excellent investment protection when you upgrade to the z10 EC

## *Within the server*

- HiperSockets™
  - Multi Write Facility
  - Layer2 support
- Integrated console controller
- Integrated communications controller support

## *To the Data*

- FICON/FCP
  - FICON® Express4
  - **FICON Express2**
  - **FICON Express**
  - ESCON®



## *To the Network*

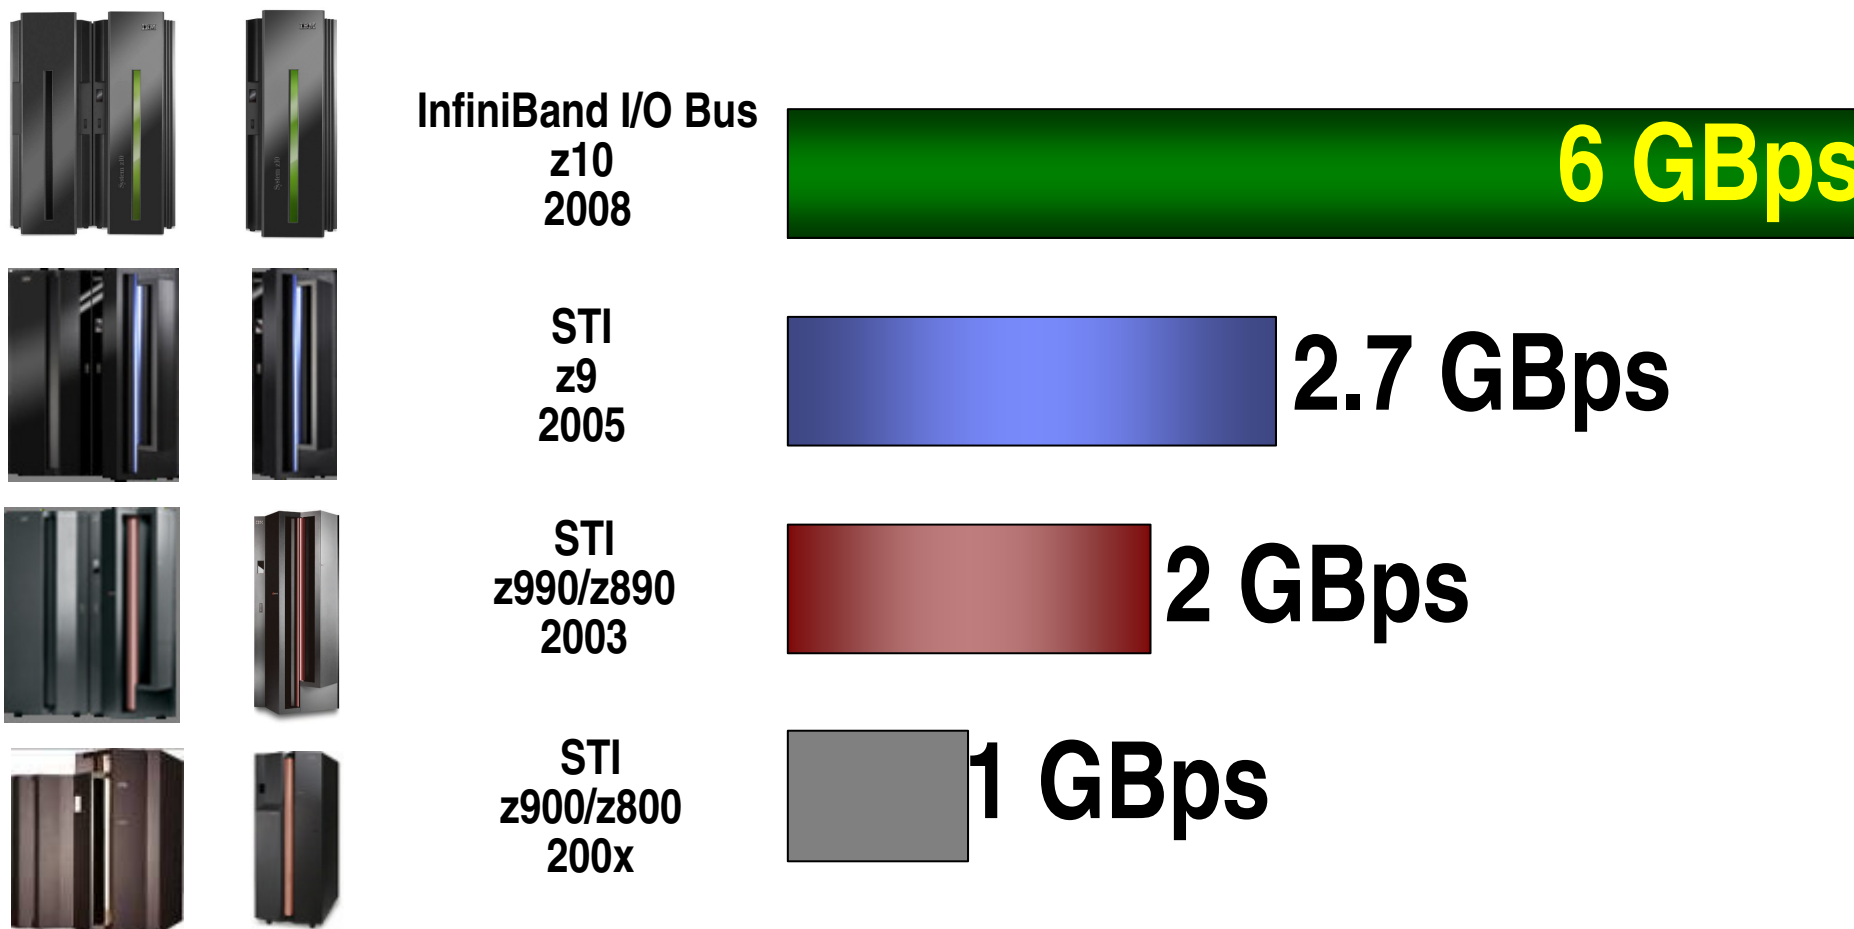
- OSA-Express3
  - 10 Gigabit Ethernet
- OSA-Express2
  - 1000BASE-T Ethernet
  - Gigabit Ethernet LX and SX
  - 10 Gigabit Ethernet LR

## *For Clustering*

- InfiniBand Coupling Links
- ICB-4
- ISC-3 (peer mode only)
- IC (define only)
- STP - NTP Client Support
- Support for n-2 and above servers

\* Note: Red items carry forward on a Machine MES only, not available for new system orders

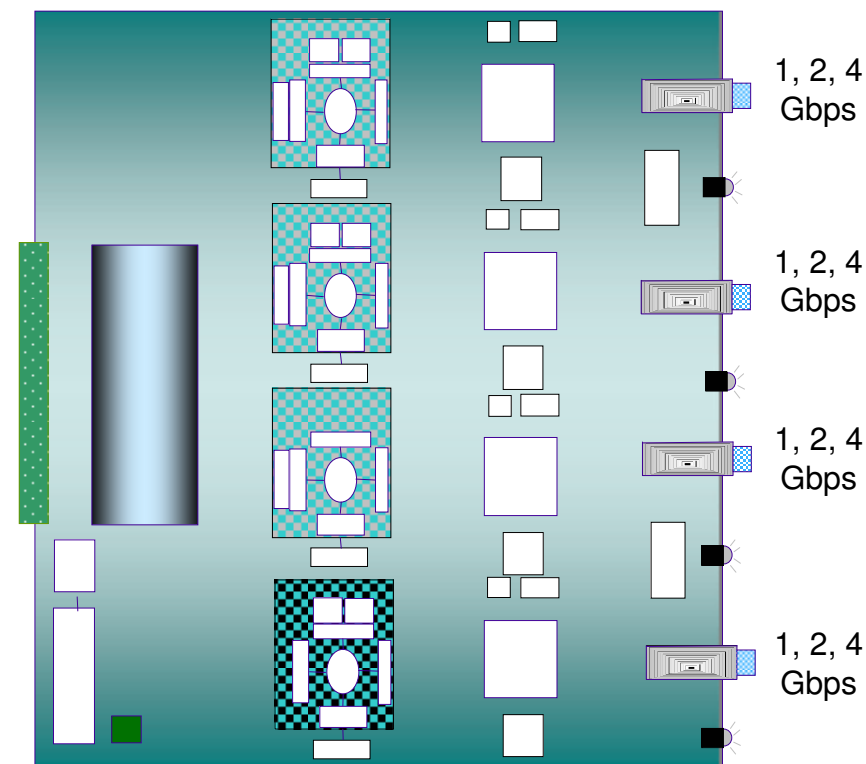
# I/O Subsystem host bus interconnect speeds in GBps



STI: Self-Timed Interconnect

# z10 EC FICON Express4

- **FICON Enhancements**
  - High Performance FICON for System z (zHPF) for FICON Express4
  - Improved performance at extended distance for FICON Express4 (and FICON Express2) features
- **1, 2, 4 Gbps auto-negotiated**
- **Up to 336 ports**
- **LX 10 KM, LX 4 KM and SX features**
  - A 10 KM LX transceiver is designed to interoperate with a 4 KM LX transceiver
- **Concurrent repair of optics**
- **Personalize as:**
  - FC (native)
    - Channel-To-Channel (CTC)
      - z/OS, z/VM, z/VSE, z/TPF, TPF, Linux on System z
    - FCP (Fibre Channel Protocol)
      - Support of SCSI devices
        - z/VM, z/VSE, Linux on System z



**FC 3321 FICON Express4 10 KM LX**  
**FC 3324 FICON Express4 4 KM LX**  
**FC 3322 FICON Express4 SX**




## z10 OSA-Express3

- **Double density of ports compared to OSA-Express2**
  - Reduced CHPIDs to manage
  - Reduced I/O slots
  - Reduced I/O cages or I/O drawers
  - Up to 96 LAN ports versus 48
- **Designed to reduce the minimum round-trip networking time between z10 BC & z10 EC systems (reduced latency)**
  - Designed to improve round trip at the TCP/IP application layer
    - OSA-Express3 10 GbE
      - **45%** improvement compared to the OSA-Express2 10 GbE
    - OSA-Express3 GbE
      - **45%** improvement compared to the OSA-Express2 GbE
  - Designed to improve throughput (mixed inbound/outbound)
    - OSA-Express3 10 GbE
      - 1.0 GBytes/s @ 1492 MTU
      - 1.1 GBytes/s @ 8992 MTU
      - 3-4 times the throughput of OSA-Express2 10 GbE
      - 0.90 of Ethernet line speed sending outbound 1506-byte frames
      - 1.25 of Ethernet line speed sending outbound 4048-byte frames

The above statements are based on OSA-Express3 performance measurements performed in a test environment on a System z10 EC and do not represent actual field measurements. Results may vary.

# z10 Cryptographic Support

- **CP Assist for Cryptographic Function (CPACF)**
  - Standard on every CP and IFL
  - **Supports the following algorithms:**
    - DES, TDES, AES-128, AES-192, AES-256
    - SHA-1, SHA-224, SHA-256, SHA 384 & SHA 512
    - Pseudo random Number Generation (PRNG)
    - SHA-1, SHA-256, and SHA-512 are shipped enabled
  - UP to 4096-bit RSA keys
  - Random Number Generation Long (8 bytes to 8096 bits)
- **Crypto Express2**
  - Two features – 1 (z10 BC only) and 2 Coprocessor option (minimum of 2 features)
  - Two configuration modes
    - Coprocessor (default)
      - Federal Information Processing Standard (FIPS) 140-2 Level 4 certified
      - Accelerator (configured from the HMC)
  - Three configuration options (Two for 1 Coprocessor option)
    - Default set to Coprocessor
  - Concurrent Patch
  - Secure Key AES (128, 192 and 256 bit) support
  - Support for 13 through 19 Personal Account Numbers
- **Dynamic Add Crypto to LPAR**
  - No recycling of LPAR
  - No POR required



**Crypto Express2**



# Agenda

## **System z10 EC**

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

**Capacity on Demand**

Green IT

High Availability

Operating Systems

## **System z10 BC**

Overview

Architecture

## **zEnterprise System z196**

Overview

Architecture

Hardware Development

Firmware Development

# Just in time capacity gives you control

- **Permanent and temporary offerings – with you in charge**
  - Permanent offerings – Capacity Upgrade on Demand (CUoD), Customer Initiated Upgrade (CIU)
  - Temporary offerings include On/Off Capacity on Demand (On/Off CoD), 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



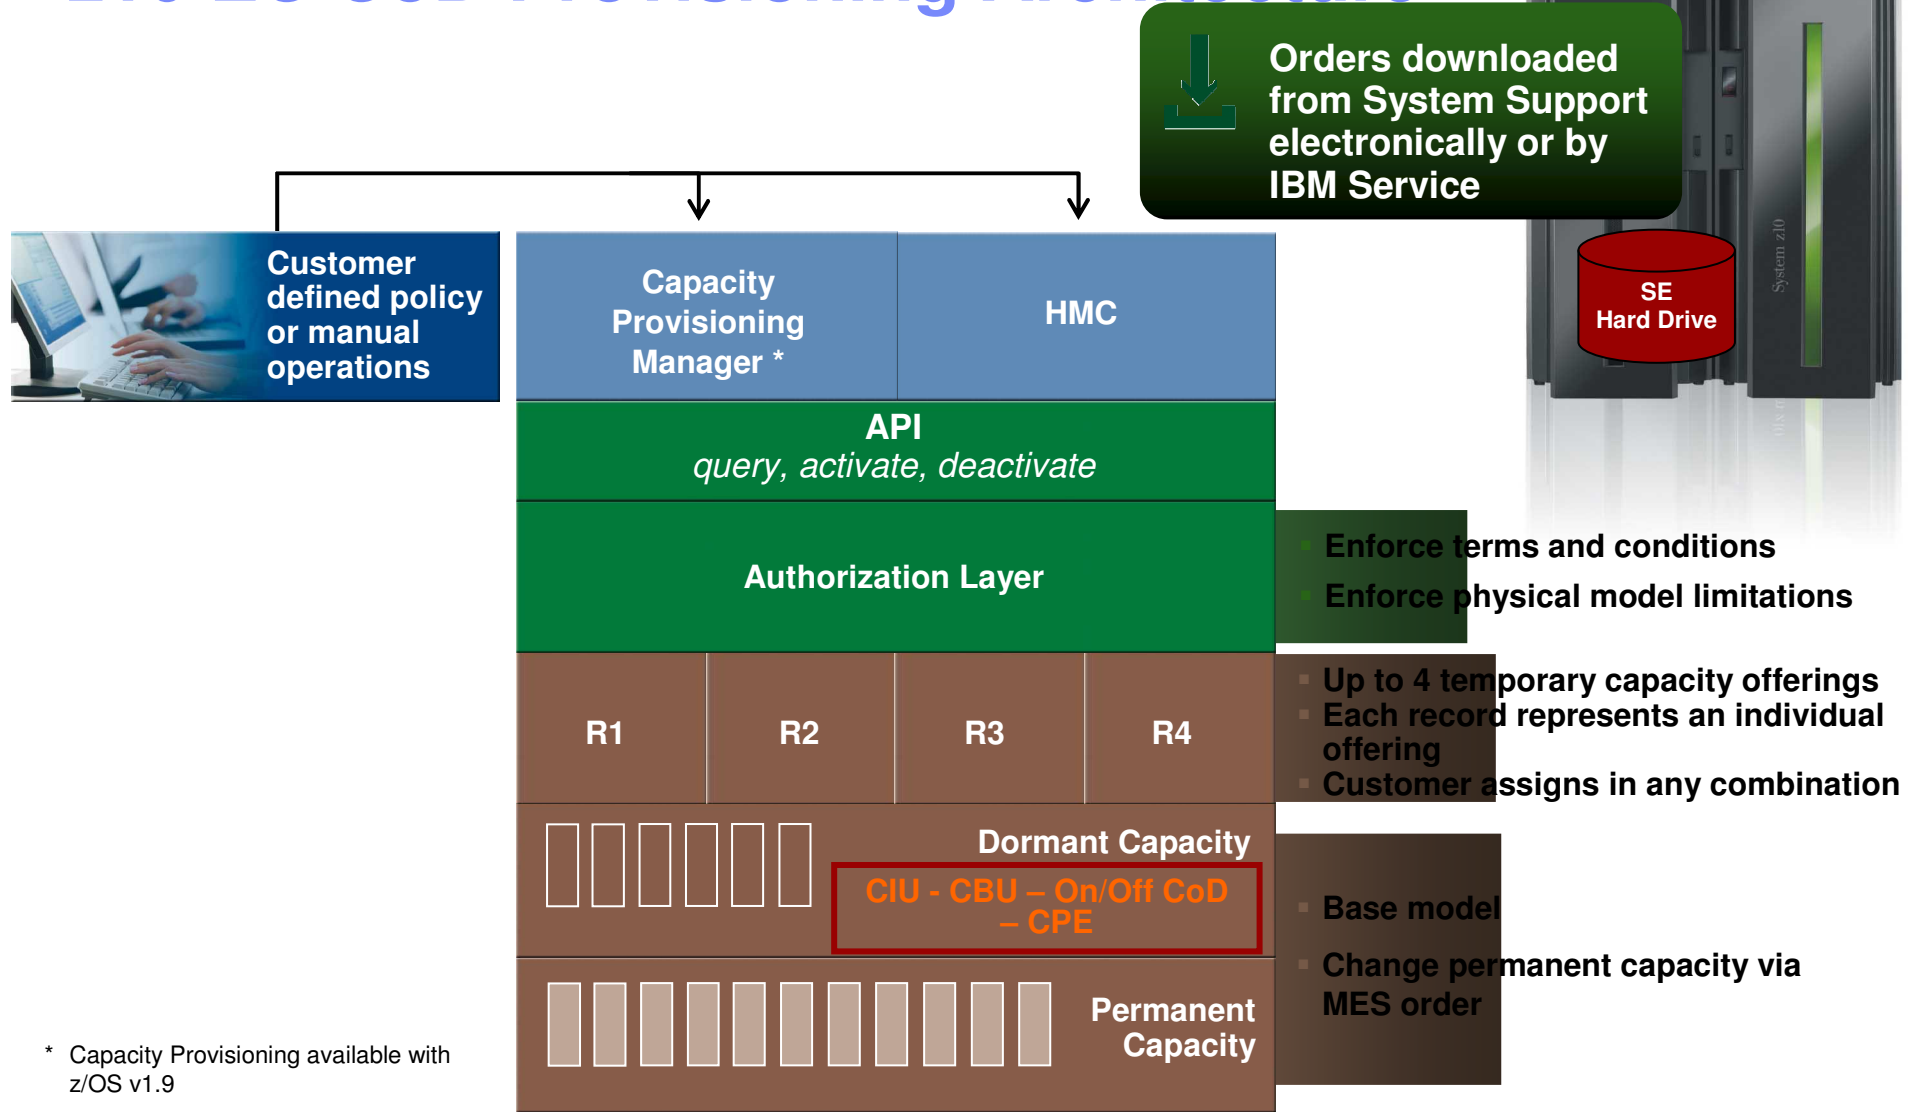




# z10 CoD Offerings

- **On-line Permanent Upgrade**
  - Permanent upgrade performed by customer (previously referred to Customer Initiated Upgrade - CIU)
- **Capacity Backup (CBU)**
  - For disaster recovery
  - Concurrently add CPs, IFLs, ICFs, zAAPs, zIIPs, SAPs
  - Pre-paid
- **Capacity for Planned Event (CPE)**
  - To replace capacity for short term lost within the enterprise due to a planned event such as a facility upgrade or system relocation
  - Predefined capacity for a fixed period of time (three days)
  - Pre-paid
- **On/Off Capacity on Demand (On/Off CoD)**
  - Production Capacity
  - Supported through software offering – Capacity Provisioning Manager (CPM)
  - Payment:
    - Post-paid or Pre-paid by purchase of capacity tokens
    - Post-paid with unlimited capacity usage
    - On/Off CoD records and capacity tokens configured on Resource Link
- **Customer Initiated Upgrade (CIU)**
  - Process/tool for ordering temporary and permanent upgrades via Resource Link
  - Permanent upgrade support:
    - Un-assignment of currently active capacity
    - Reactivation of unassigned capacity
    - Purchase of all PU types physically available but not already characterized
    - Purchase of installed but not owned memory

# z10 EC CoD Provisioning Architecture



\* Capacity Provisioning available with z/OS v1.9



# Agenda

## **System z10 EC**

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

**Green IT**

High Availability

Operating Systems

## **System z10 BC**

Overview

Architecture

## **zEnterprise System z196**

Overview

Architecture

Hardware Development

Firmware Development

# Tracking energy consumption within the infrastructure

- **Resource Link provides tools to estimate server energy requirements before you purchase a new system or an upgrade**
- **Has energy efficiency monitoring tool**
  - Introduced on IBM System z9 platform in April 2007
  - Power and thermal information displayed via the System Activity Display (SAD)
- **IBM Systems Director Active Energy Manager™ (AEM) for Linux on System z V3.1**
  - Offers a single view of actual energy usage across multiple heterogeneous IBM platforms within the infrastructure
  - AEM V3.1 energy management data can be exploited by Tivoli enterprise solutions such as IBM Tivoli Monitoring, IBM Tivoli Usage and Accounting Manager, and IBM Tivoli OMEGAMON® XE on z/OS
  - AEM V3.1 is a key component of IBM's Cool Blue™ portfolio within Project Big Green



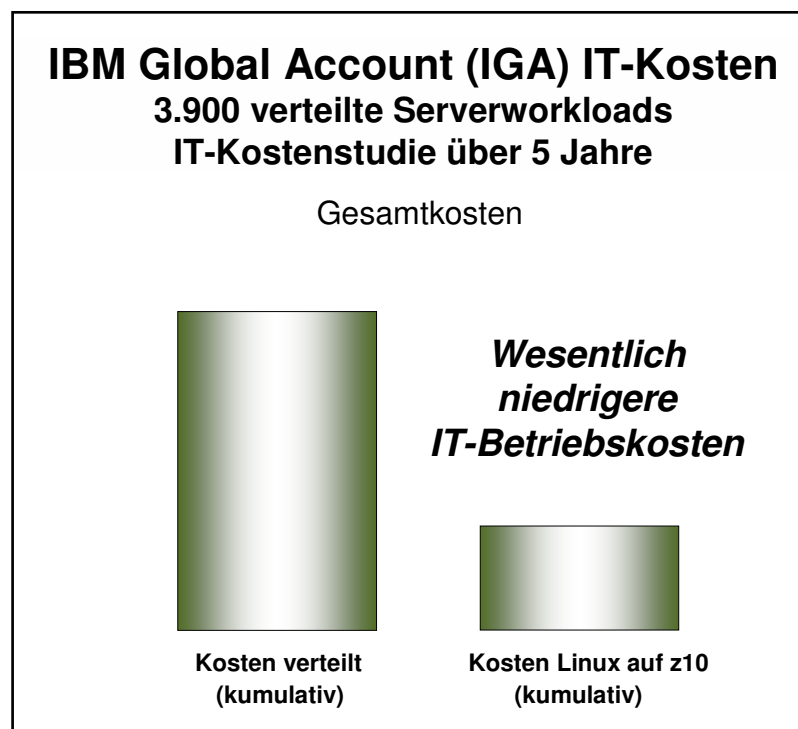
**z10 EC 64-way offers a 15% improvement in performance per kWh over z9 EC 54-way**

# IBM konsolidiert verteilte Server und erzielt Einsparungen

## Erwartetes Ergebnis bei IBM:

- **Entschlackte Umgebung mit deutlich weniger Hardware**
  - 3.900 verteilte Server-Images werden auf 15-20 System z10 konsolidiert
  - Wesentlicher Anstieg der durchschnittlichen Auslastung
- **Weniger Personalkosten durch Virtualisierung**
- **Wesentliche Senkung der Softwareausgaben**
- **85% weniger Stellflächenbedarf im Rechenzentrum durch konsolidierte Server**
  - Ermöglicht weiteres Wachstum
  - Bessere Ausnutzung der Anlagen
- **80% weniger Energieverbrauch**
- **Möglichkeit, mehr Anwendungen auf System z einzusetzen**

## \$250M+ Ersparnisse in 5 Jahren





# Agenda

## System z10 EC

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

Green IT

**High Availability**

Operating Systems

## System z10 BC

Overview

Architecture

## zEnterprise System z196

Overview

Architecture

Hardware Development

Firmware Development



# The right level of business continuity protection for your business .....GDPS family of offerings

**Continuous Availability of Data within a Data Center**


**Continuous Availability / Disaster Recovery Metropolitan Region**

**Disaster Recovery at Extended Distance**

**Continuous Availability Regionally and Disaster Recovery Extended Distance**

**Single Data Center Applications remain active**


**Near-continuous availability to data**



**GDPS®/PPRC  
HyperSwap Manager**

**Two Data Centers Systems remain active**


**Automated D/R across site or storage failure  
No data loss**



**GDPS/ PPRC  
HyperSwap Manager  
GDPS/PPRC**

**Two Data Centers**


**Automated Disaster Recovery  
“seconds” of Data Loss**



**GDPS/GM  
GDPS/XRC**

**Three Data Centers**

**Data availability  
No data loss  
Extended distances**



**GDPS/MGM  
GDPS/MzGM**





# Agenda

## System z10 EC

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

Green IT

High Availability

Operating Systems

## System z10 BC

Overview

Architecture

## zEnterprise System z196

Overview

Architecture

Hardware Development

Firmware Development

# Operating systems

## ***z/OS***

- Providing intelligent dispatching on z10 EC for performance
- Up to 64-way support
- Simplified capacity provisioning on z10 EC
- New high availability disk solution with simplified management
- Enabling extreme storage volume scaling
- Facilitating new zIIP exploitation

## ***z/TPF***

- Support for 64+ processors
- Workload charge pricing
- Exploit encryption technology

## ***z/VSE™***

- Interoperability with Linux on System z
- Exploit encryption technology
- MWLC pricing with sub-capacity option



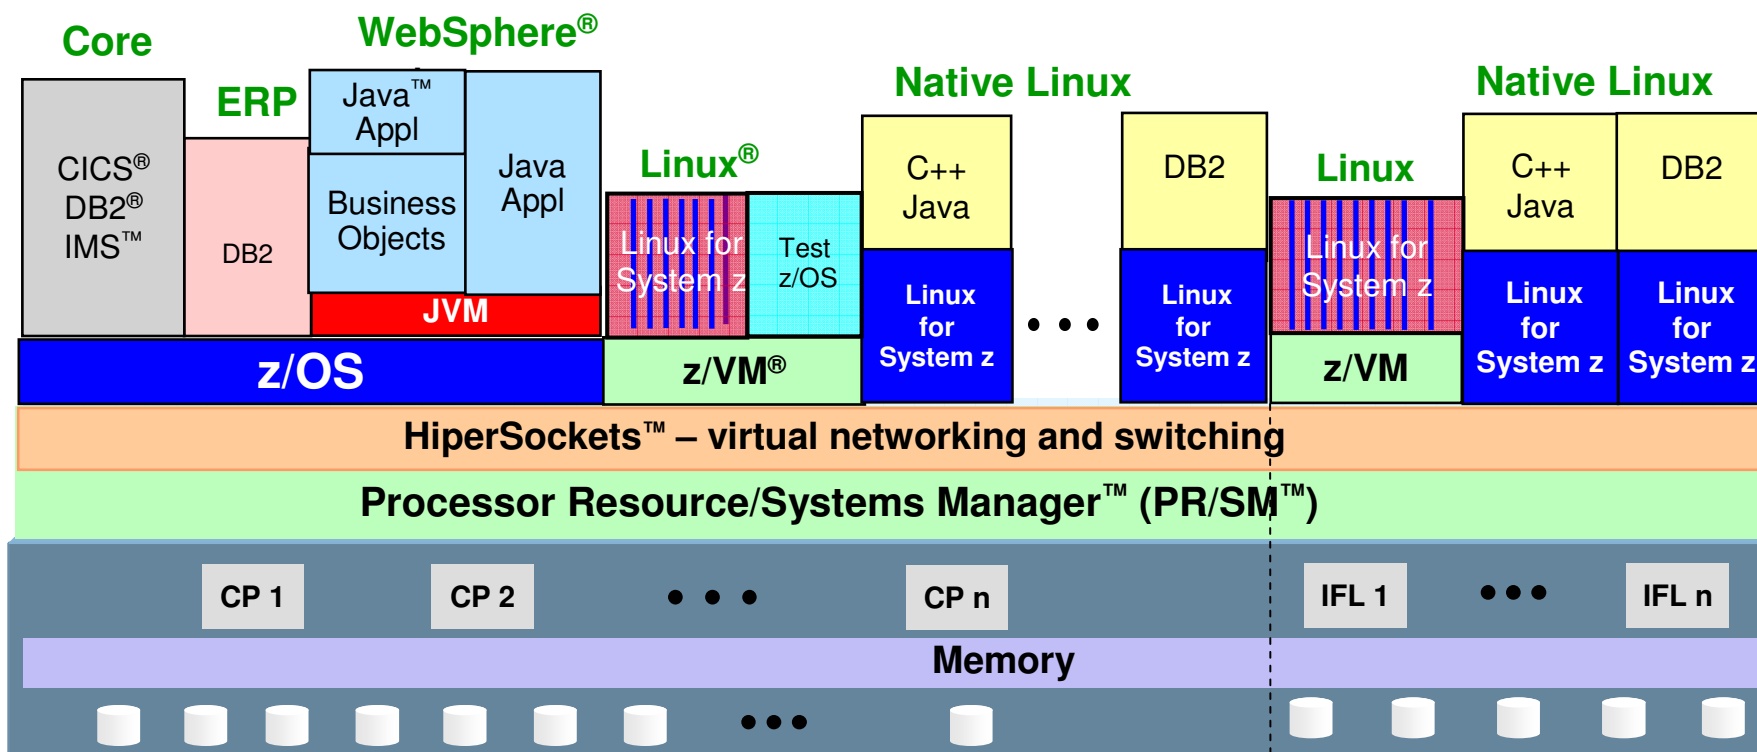
## ***Linux on System z***

- Large Page Support improves performance
- Linux CPU Node Affinity is designed to avoid cache pollution
- Software support for extended CP Assist instructions AES & SHA

## ***z/VM***

- Consolidation of many virtual images in a single LPAR
- Enhanced management functions for virtual images
- Larger workloads with more scalability

# System z – The Ultimate Virtualization Resource



- Massive, robust consolidation platform; virtualization is built in, not added on
- Up to 60 logical partitions on PR/SM; 100's to 1000's of virtual servers on z/VM
- Virtual networking for memory-speed communication, as well as virtual layer 2 and layer 3 networks supported by z/VM
- Most sophisticated and complete hypervisor function available
- Intelligent and autonomic management of diverse workloads and system resources based on business policies and workload performance objectives



# Agenda

## System z10 EC

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

Green IT

High Availability

Operating Systems

## System z10 BC

Overview

Architecture

## zEnterprise System z196

Overview

Architecture

Hardware Development

Firmware Development



# IBM System z10 Business Class

## Hardware Overview





## Businesses are realizing the value of the mainframe within their IT infrastructure – and their business


- A highly utilized, virtualized, scalable, optimized resource for consolidating workloads to help lower overall operating cost and improve energy efficiency
- A highly secure enterprise data server – when your infrastructure is secure, your business is secure
- A backbone for an enterprise SOA hub to enable integration of applications and processes, and add flexibility – when your infrastructure is flexible, your business is flexible
- Integrates new workloads including Linux<sup>®</sup>, Java<sup>™</sup> and Open Standards with demonstrably lower Total Cost of Ownership (TCO)



***Today's Mainframe:  
The Smart choice for an optimized, scalable, secure, resilient infrastructure***



# IBM z10 BC continues the CMOS Mainframe heritage



- Multiprise 2000 – 1<sup>st</sup> full-custom Mid-range CMOS S/390
- Multiprise 3000 – Internal disk, IFL introduced on midrange

- IBM eServer zSeries 800 (z800) - Full 64-bit z/Architecture<sup>®</sup>
- IBM eServer zSeries 890 (z890) - Superscalar CISC pipeline
- z9 BC - System level scaling

- z10 BC - Architectural extensions
- Higher frequency CPU





# Agenda

## System z10 EC

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

Green IT

High Availability

Operating Systems

## System z10 BC

Overview

Architecture

## zEnterprise System z196

Overview

Architecture

Hardware Development

Firmware Development




## z10 BC Overview





- **Machine Type**
  - 2098
- **Single Model – E10**
  - Single frame, air cooled
  - Non-raised floor option available
- **Processor Units (PUs)**
  - 12 PU cores per System
  - 2 SAPs, standard
  - Zero spares when all PUs characterized
  - Up to 10 PUs available for characterization
    - Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z10 Application Assist Processors (zAAPs), System z10 Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs)
- **Memory**
  - System Minimum of 4 GB
  - Up to 128 GB for System, including HSA (up to 256 GB, June 30, 2009)
    - 8 GB Fixed HSA, standard
    - Up to 120 GB for customer use (up to 248 GB, June 30, 2009)
    - 4, 8 and 32 GB increments (32 GB increment, June 30, 2009)
- **I/O**
  - Up to 12 I/O Interconnects per System @ 6 GBps each
  - 2 Logical Channel Subsystems (LCSSs)
  - Fiber Quick Connect for ESCON and FICON LX
  - New OSA-Express3 Features
  - ETR feature, standard

# z10 BC – Under the covers Front View



# z10 BC CPC and Memory Drawer Layout



- Two new types of Single Chip Modules (SCMs)
- Processor – PU (4 SCM's x 3 cores = 12 PU's)
- System Controller – SC (2)


2 Flexible Support Processor (FSP) card slots providing support for the Service Network subsystem (hot swappable)

6 fanout card slots providing support for the I/O subsystem and/or coupling


2 card slots for the oscillator/ETR function (standard) – dynamic switchover support



# z10 BC PU/SC SCM Components



**PU Chip**





**Heatsink**

**PU Chip**


**LGA**

**Assembled SCM  
with heatsink**



**SC Chip**

**CPC Drawer with  
4 PU SCMs, 2 SC SCMs**



# z10 BC SCM Vs z10 EC MCM Comparison


## z10 BC SCMs

- **PU SCM**
  - 50mm x 50mm in size – fully assembled
  - Quad core chip with 3 active cores
  - 4 PU SCMs per System with total of 12 cores
  - PU Chip size 21.97 mm x 21.17 mm
- **SC SCM**
  - 61mm x 61mm in size – fully assembled
  - 2 SC SCMs per System
  - 24 MB L2 cache per chip
  - SC Chip size 21.11 mm x 21.71 mm

**Single PU Chip  
without heatsink**




**Single SC Chip  
without heatsink**




## z10 EC MCM

- **MCM**
  - 96mm x 96mm in size
  - 5 PU chips per MCM
    - Quad core chips with 3 or 4 active cores
    - PU Chip size 21.97 mm x 21.17 mm
  - 2 SC chips per MCM
    - 24 MB L2 cache per chip
    - SC Chip size 21.11 mm x 21.71 mm
  - Up to 4 MCMs for System



# z10 BC Upgrade Paths

- Can enable dynamic and flexible capacity growth for mainframe servers
- Temporary capacity upgrade available through On/Off Capacity on Demand
- Temporary, nondisruptive addition of CP processors, IFLs, ICFs, zAAPs or zIIPs
- New options for reconfiguring specialty engines if the business demands it
- New options for changing On/Off CoD configurations
- Subcapacity CBU engines



- Full upgrades within the z10 BC
- Any to any upgrade from the z9 BC
- Any to any upgrade from z890
- No charge MES upgrades on IFLs, zAAPs and zIIPs





# Agenda

## System z10 EC

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

Green IT

High Availability

Operating Systems

## System z10 BC

Overview

Architecture

## **zEnterprise System z196**

Overview

Architecture

Hardware Development

Firmware Development



# IBM zEnterprise z196

## Hardware Overview





# What is the IBM zEnterprise System?

## IBM zEnterprise Unified Resource Manager

An integrated System z management facility responsible for zEnterprise platform management

Unifies management of resources, extending System z qualities of service across the zEnterprise System

## IBM zEnterprise 196

Ideal for large scale data and transaction serving and mission critical applications

Capable of massive scale up, over 50 Billion Instructions per Second (BIPS)



## IBM zEnterprise BladeCenter Extension (zBX)

Selected IBM POWER7 blades and IBM x86 blades\* for AIX and Linux applications

High performance optimizers and appliances

Dedicated high performance private network


\*Statement of Direction

# The New Processor Generation



## zEnterprise Processor

- 5.2 GHz
- 45 nm SOI Technology
- 1.4 billion transistors
- 13 layer metal
- 3.5 Km wires
- 512 mm<sup>2</sup>
- 8.1K Power C4's
- 1.1K Signal C4's



## P7 Processor

- 4.1 GHz
- 45 nm SOI Technology
- 1.2 billion transistors
- 567 mm<sup>2</sup>



# Agenda

## System z10 EC

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

Green IT

High Availability

Operating Systems

## System z10 BC

Overview

Architecture

## **zEnterprise System z196**


Overview

**Architecture**

Hardware Development

Firmware Development


# zEnterprise Multiple Chip Module (MCM)



- 96mm x 96mm MCM
  - >192 MB L4 Cache
  - 103 layers Glaskeramik
  - 7356 LGA connections
  - Up to 24 aktiv processor cores



# zEnterprise - (z196) Water Cooled System






# Moving to Integrated System with Accelerators and Co-Processors

**OLTP**

**Distributed**  
**Integrated**



## Values

- Reduced latency
- Simple operations
- Reduced points of failure
- Simpler and faster problem determination
- Reduced communication path lengths and costs
- Coordinated management and resource sharing
- Supports both componentisation and simplification





## We know, you hear it all the time: Mainframe Misperceptions

MYTH	REALITY
<p>"Antiquated technology" Only for legacy systems, not modern workloads</p>	<ul style="list-style-type: none"><li>▪ Can provision new virtual servers (images) in minutes rather than days/ weeks , add capacity on the fly</li><li>▪ Runs modern technologies such as Java®, C/C++, XML, Linux, Web 2.0, SOA</li><li>▪ Unrivaled availability for mission critical workloads</li></ul> <p>▪ <b>And now zEnterprise.. So who is the dinosaur?</b></p>
<p>"Too Expensive"</p>	<ul style="list-style-type: none"><li>▪ System z Solution Editions package HW, SW maintenance services at very attractive prices , close or equivalent to distributed, including for SAP</li><li>▪ The most efficient server for large scale consolidation. System z can reduce per-core licensing costs up to 28 to 1 compared to x86 environments</li><li>▪ With unrivaled virtualization and scale, System z can be significantly less cost than distributed systems for running multiple, diverse workloads</li><li>▪ Low cost server specialty engines (zAAP, zIIP, IFL) enable many applications (e.g. Java®, XML, Linux) and large scale consolidation on the platform</li></ul>
<p>"No Apps Support"</p>	<ul style="list-style-type: none"><li>▪ Over 6,500+ applications available for System z</li><li>▪ Over 3,000+ Linux applications are supported on System z; 18% growth in 2008</li><li>▪ Over 1,700+ ISVs building applications for System z</li></ul>
<p>"Skills Shortage"</p>	<ul style="list-style-type: none"><li>▪ Over 60,000 students educated worldwide, adding 5,000 more in China by 2010</li><li>▪ Over 600 schools enrolled globally</li><li>▪ 50,000 Students completed courses to date</li></ul>



# Agenda

## System z10 EC

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

Green IT

High Availability

Operating Systems

## System z10 BC

Overview

Architecture

## zEnterprise System z196





Overview

Architecture

## Hardware Development

Firmware Development

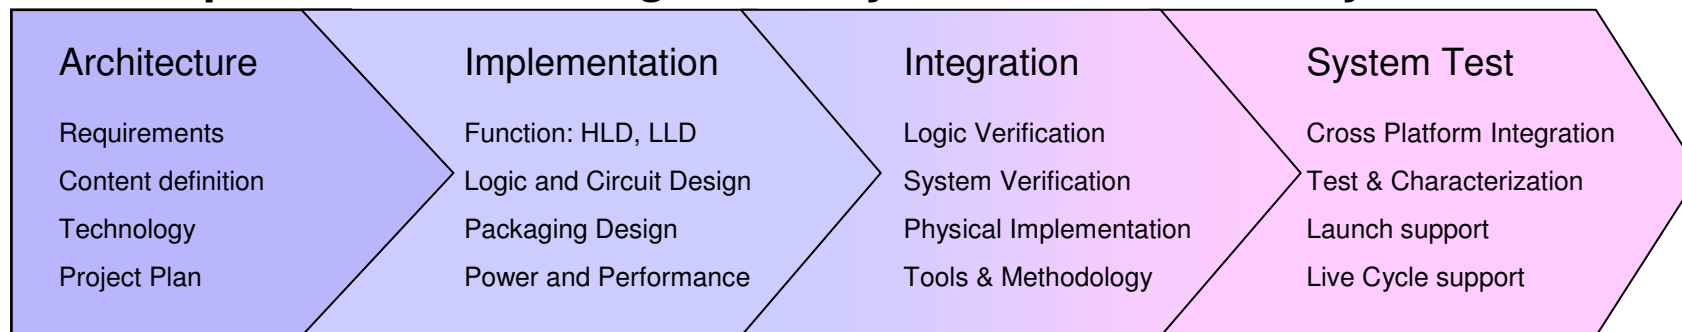
# System Hardware Development



## Concept

## Design & Verify

## Qualify & Deliver





# Agenda

## System z10 EC

Overview

Architecture

MCM

PU Chip

Accelerator

SC Chip

Book Layout

Connectivity

Inter Book Connectivity

IO Connectivity

Ethernet Connectivity

Crypto

Capacity on Demand

Green IT

High Availability

Operating Systems

## System z10 BC

Overview

Architecture

## zEnterprise System z196


Overview

Architecture

Hardware Development

**Firmware Development**

# System z Firmware Development

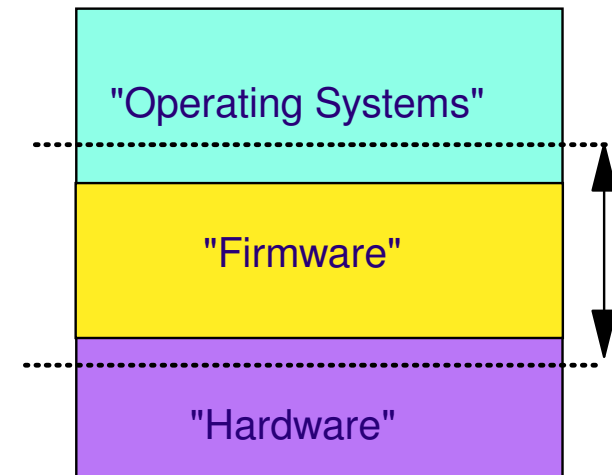


## Base HW Support: Basic Values

- Initialisation
- HW Access und Debug
- RAS Functions
- Customer Operations

## Complex Functions: Platform Values

- LPAR, Sysplex
- Complex RAS: Self Healing, Multi-pathing...
- On Demand Functions ( CP / Memory capacity)
- Security / Encryption
- Systems / Ensemble Management and Monitoring
- Power management / Energy Efficiency
- Concurrent Functions e.g. Driver Upgrade



### A simple definition:

Firmware is all the software or code delivered with or 'firm' to the server platform.

# System z Firmware (LIC) Components

## Support Element


System-Startup  
System Control,  
Serviceability

C, C++, JAVA,  
Perl, Assembler

## HMC

Control of multiple zSeries  
Systems via SE

C, C++, Java



## CEC-LIC

### Millicode

390/ Instructions Assembler  
Hardware Access Functions

### LPAR Hypervisor

Virtualization PL/X, Assembler

### I390

Complex System Functions  
IO-Driver PL.8, C  
Autonomic Functions

## PSCN

Cage / Power Control  
SE-CEC Connection

C, C++



\*\*\* STOP: 0x00000404 (0x00000000,0x12345678,0x69696969,0x97f86a65)  
PAGE\_NOT\_FOUND\*\*\* ADDRESS 09034672 has base at 09037000 - abort

CPUID:GenuineIntel Pentium Irql:1f SYSUER: 0xf0096404

Dll Base	DateStmp	Name	Dll Base	DateStmp	Name
80100000	336546ff	- ntoskrnl.exe	80010000	331054d1	- iexplore.exe
80037386	33104f49	- atapi.sys	80070000	3341a4d9	- outlook.exe
88013faa	33107s98	- epsy.sys	8021bf09	3342ab4f	- disk.sys
80037389	3310354d	- CLASS2.SYS	80102848	334013fb	- NTice.sys
80037390	3310098d	- Siwvid.sys	f95cb000	3340101f	- Ntfs.sys
800ae122	3310h896	- Floppy.sys	f1035800	336735af	- Parport.sys
800bf392	33100f89	- KSecDD.sys	f031fde0	34dc3ccf	- Beep.sys
801cca04	3310h854	- i8042prt.sys	f95cb000	3340101f	- mouclass.sys
80037398	331034fg	- kbdc class.sys	f3e3b000	33104f49	- ctrl2cap.sys
80037399	3310f987	- VIDEOPORT.SYS	f42cb000	3340101f	- msfs.sys
80037400	3310g653	- vga.sys	f63cb640	3340101f	- iexplore.exe
80037401	33106g89	- npfs.sys			

Address	dword	dump	Build [1381]	- Name			
c0948732	3340101f	8021bf09	80037391	12345678	80037401	80a0b00f	- Ntfs.sys
80948732	80070000	80037386	12345678	8021bf09	1c005638	df00eabf	- rnl.exe

Restart and set the recovery options in the system control panel or the /CRASHDEBUG system start option. If this message reappears, Contact your system administrator or technical support group.

## Changing Tires While Driving...



### z10 UIRA Status

System MTBF: 30+ yrs

FW MTBF: 100+ yrs



**Noch Fragen?**

01010101010101010101010101

**Danke !**