



## IBM System z10 Enterprise Class

### z10 is a TEN / z top 10 List

**David F. Anderson PE**  
**Green Architect**  
**IBM Energy Efficiency Initiative**  
**dfa@us.ibm.com**



**The Future Runs on System z**

**March 3, 2008**

© 2008 IBM Corporation

IBM Systems



## Trademarks

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

DB2	HyperSwap	System z9*
Cool Blue	IBM*	Tivoli®
DRDA*	IBM logo*	WebSphere*
DS8000	OMEGAMON*	z9
ESCON*	Parallel Sysplex*	zArchitecture*
eServer	ResourceLink	zOS*
FICON*	System p	z/VM*
FlashCopy*	System Storage	z/VSE
GDPS*	System x	zSeries*
HiperSockets	System z	

\* 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 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

1. Green Design
2. Faster Engines
3. New Advanced Function
4. Flexible Capacity
5. Improved RAS
6. Support for New Security Standards
7. Expanding the ability to run more diverse workloads
8. Upgradeability
9. Virtualization
10. Energy Efficiency


## Green Design

- **Modular**
  - Processor Book Package
  - I/O cards
- **High Efficiency Power Supplies**
- **Variable Speed Energy Efficient Blowers**
  
- **Share Everything Architecture**
- **Deep Integration**
  
- **A masterpiece of HW, SW, microcode working in unison**

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



## z10 EC Book Layout




## Comprehensive Software Leveraging the Strengths of the z10 EC



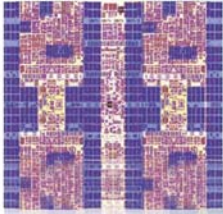
## Faster Engines and more overall Performance

- **Faster Engines**
  - Should need less of them
  - Shorter paths are more efficient
- **Quad Core Technology**
- **State of the art packaging**
- **Share Everything Architecture**
- **Coprocessor and assists**



## Making high performance a reality

- **New Enterprise Quad Core z10 EC processor chip**
  - 4.4 GHz - additional throughput means improved price/performance
  - Cache rich environment optimized for data serving
  - 50+ instructions added to improve compiled code efficiency
  - Support for 1MB page frames
  
- **Hardware accelerators on the chip**
  - Hardware data compression
  - Cryptographic functions
  - Hardware Decimal Floating point
  
- **CPU intensive workloads get performance improvements from new core pipeline design**




**Enterprise Quad Core z10 EC processor chip**

10 IBM Systems

## System z10 EC Multi-Chip Module (MCM)

- **96mm x 96mm MCM**
  - 103 Glass Ceramic layers
  - 7 chip sites
- CMOS 11s chip Technology (65 nm)



### 5 Enterprise Quad Core Processor (PU) chips

- 21.97 mm x 21.17 mm
- 1.0 billion transistors/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



### 2 Storage Control (SC) chips

- 21.11 mm x 21.71 mm
- 1.6 billion transistors/chip
- L2 cache 24 MB/SC chip (48 MB/MCM)
- L2 access to/from other MCM/Books

### 4 SEEPRROM (S) chips


- Product data for MCM, chips and other engineering information

Clock Functions distributed across PU and SC chips

Dense packaging enabling Energy Efficient **Green Computing**

## Improved server performance and scalability with faster and more processors and improved dispatching synergy


- The z10 EC delivers on average 50% more performance in a n-way configuration than a IBM System z9™ Enterprise Class (z9 EC) n-way
  - The uniprocessor is expected to deliver 62% more performance than z9 EC uniprocessor \*
- The z10 EC 64-way offers 70% more server capacity than the largest z9 EC\*\*
- Introducing HiperDispatch for improved synergy with z/OS® operating system to deliver scalability and performance



\* LSPR mixed workload average running z/OS 1.8 - z10 EC 701 versus z9 EC 701


\*\* This is a comparison of the z10 EC 64-way and the z9 EC S54 and is based on LSPR mixed workload average running z/OS 1.8

## Evolution of System z Specialty Engines



Cell Processor\*

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 z10 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.

## System z and Cell Broadband Engine – The Vision A 'Marriage' of Two Technologies that Perfectly Complement Each Other



z today



Cell Blade  
QS20, QS21, QS2x



Integrated and / or  
Networked Cell (NG)

z tomorrow

Preserves the same programming model between Network and Integrated



Aerospace and Defense

Financial Services Sector

Chemicals and Petroleum

Digital Video Surveillance

Digital Media

Information Based Medicine

Electronic Design Automation

## Focused performance boost Hardware Decimal Floating Point

*Up to 10X improvement  
in decimal floating point  
instructions*

- Decimal arithmetic widely used in commercial and financial applications
- Computations often handled in software
- First delivered in millicode on the System z9 - brought improved precision and function
  - Avoids rounding and other problems with binary/decimal conversions
- On z10 EC integrated on every core giving a performance boost to execution of decimal arithmetic
- Growing industry support for hardware decimal floating point standardization
  - Java BigDecimal, C#, XML, C/C++, GCC, DB2 V9, Enterprise PL/1, Assembler
  - Endorsed by key software vendors including Microsoft® and SAP
  - Open standard definition led by IBM



*Bringing high performance computing benefits to  
commercial workloads*



## z10 EC Hardware Decimal Floating Point Accelerator

- Meets requirements of business and human-centric applications
  - ▶ Performance, Precision, Function
  - ▶ Avoids rounding and other problems with binary/decimal conversions
  - ▶ Improved numeric functionality over legacy Binary Coded Decimal (BCD) operations
  - ▶ Much of commercial computing is dominated by decimal data and decimal operations



Single PU Core

- IBM z10 EC Hardware Decimal Floating Point Unit co-developed (HDFU) with POWER6
  - ▶ Common architecture operations and semantics
  - ▶ Common dataflow elements
  - ▶ Mainframe legacy Binary Coded Decimal (BCD) operations mapped onto HDFU in z10 EC
- Growing industry support for DFP standardization
  - ▶ Java BigDecimal, C#, XML, XL C/C++, GCC, DB2 9, Enterprise PL/1, Assembler
  - ▶ Endorsed by key software vendors including Microsoft® and SAP
  - ▶ Open standard definition led by IBM



## New Advanced Function

- HiperSocket Improvements
- Large Pages
- Infiniband links
- Hiper Dispatch
- ...

## z10 EC New Functions and Features (February 2008)

Five hardware models		<b>6.0 GBps InfiniBand HCA to I/O interconnect</b>	
Faster Uni Processor		FCP Performance Improvement	
Up to 64 customer PUs		SCSI IPL included in Base LIC	
36 CP Subcapacity Settings		OSA-Express3 10 GbE (2Q08)*	
Star Book Interconnect		<b>HiperSockets Multi Write Facility enhancements</b>	
Up to 1.5 TB memory		<b>InfiniBand Coupling Links (2Q08)*</b>	
Fixed HSA as standard		STP using InfiniBand (2Q08)*	
<b>Large Page Support (1 MB)</b>		Capacity Provisioning Support	
<b>HiperDispatch</b>		Scheduled Outage Reduction	
Enhanced CPACF SHA 512, AES 192 and 256-bit keys		Improved RAS	
<b>Hardware Decimal Floating Point</b>		FICON LX Fiber Quick Connect	
New Capacity on Demand architecture and enhancements		Power Monitoring support	
<div style="background-color: red; color: white; padding: 5px; display: inline-block;"> <b>SOD: PSIFB for z9 EC &amp; BC for non-dedicated CF Models*</b> </div>			

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.

\* 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 HiperSockets Performance Enhancements

- **HiperSockets Multiple Write Facility**
  - ▶ Performance improvements
    - For the streaming of bulk data over a HiperSockets link between LPARs
    - Allows receiving LPARs to process a much larger amount of data per I/O interrupt
      - z/OS V1.10\*
        - ◆ Transparent to software in receiving LPARs
- **HiperSockets Layer 2 support**
  - ▶ **Hosting of new workloads**
  - ▶ Host non-IP protocols (IPX, NetBIOS, SNA)
  - ▶ **Bridge from and into distributed switched fabrics**
  - ▶ Supports broadcast, unicast, or multicast
  - ▶ VLANs: In Layer 2 the same rules apply as for Layer 3 VLAN handling
    - Linux on System z
  - ▶ Layer 3 applications cannot communicate with Layer 2 applications
  - ▶ z/VM 5.2 or higher – Guest support

*High speed connectivity between LPARs “Network in a box”*

\* 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 Performance Considerations

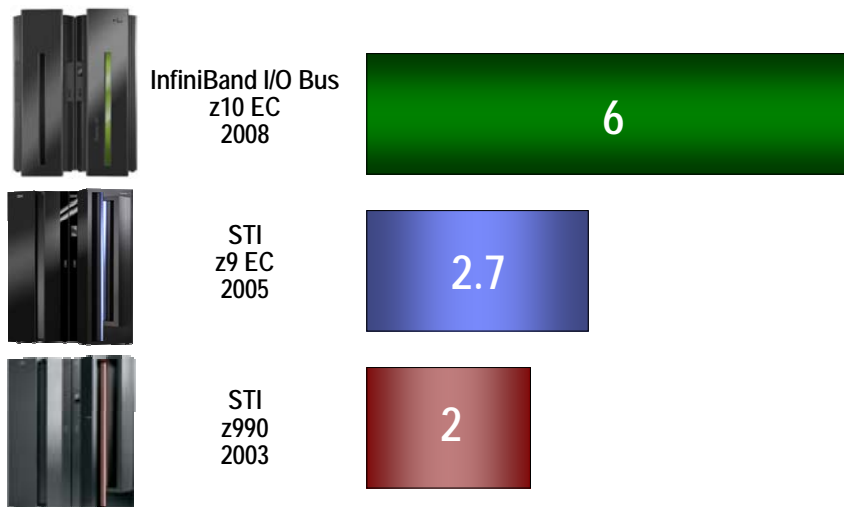
- **Large Page is a special purpose performance improvement feature. It is not recommended for general use. Large page usage provides performance value to a select set of applications**
  - ▶ **These are primarily long running memory access intensive applications**
- **Not all applications benefit from using large pages. Some applications can be severely degraded by the use of large pages**
  - ▶ Short lived processes with small working sets are usually not good candidates for large pages
- **Factors to consider when trying to either estimate the potential benefit or understand measured performance differences of using larger pages instead of 4K pages include:**
  - ▶ Memory Usage
  - ▶ A workload's page translation overhead
- **Large Page Exploiters**
  - ▶ **A future\* release of DB2 will support Large Pages for bufferpools. Default is 4K pages**
  - ▶ **Java 6.0 SR1 for z/OS is planned\* to support Large Pages. Large pages can be used to back the object heap. Default is 4K pages**

\* 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 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**

## I/O Subsystem host bus interconnect speeds in GBps




STI: Self-Timed Interconnect  
 InfiniBand refers to: 12x IB-DDR for z10 EC and 12x IB-SDR for System z9

## Innovation that *Matters!* Grow with Ease Using Online Expansion

### IBM System Storage DS8000

#### Dynamic Volume Expansion


- ▶ Simplifies management by enabling easier, online, volume expansion to support application data growth.
- ▶ Yields more highly available, simplified volume expansion
  - No longer need to backup data, bring volume offline, delete and recreate volume to expand it
  - Can be done on-line with a single command or GUI screen
  - Copy services relationships must be removed before volume expansion
  - Some operating systems-specific actions may be needed on server side before applications can see larger volumes
  - Can be used to non-disruptively migrate from smaller to larger z/OS volumes



## Extended Address Volumes: Taking z/OS storage volumes to the extreme with the DS8000

- **EAV helps relieve constraints to address large capacity needs**
- **Extended Address Volumes (EAV) enables volumes of more than 65,280 cylinders**
  - ▶ 223 GB volumes initially supported on z/OS V1.10\* and IBM System Storage DS8000\*
  - ▶ First exploiter is VSAM –applications that uses VSAM data sets (including DB2 and CICS) can benefit from EAV.
  - ▶ In the future, IBM intends to expand support for EAV with larger volume sizes and support for additional data set types and access methods.\*
- **In the future, EAV can help simplify storage management.**
  - ▶ Manage fewer, large volumes as opposed to many small volumes.
- **DS8000 HyperPAV function complements EAV by allowing the scaling the I/O rates against a single, larger volume**
- **DS8000 Dynamic Volume Expansion can allow non-disruptive migration to larger volume sizes**

**New! Preview February 2008  
Target 2H08**



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

## z10 EC Coupling Link Options

- **PSIFB\* - 12x IB-DDR for high speed communication at medium distance**
  - ▶ New CHPID – CIB (Coupling using InfiniBand)
  - ▶ New 50 micron OM3 (2000 MHz-km) multimode fiber with MPO connectors
  - ▶ **Up to 150m at 6 GBps**
- **ICB-4 for short distances over copper cabling**
  - ▶ New ICB-4 cables are required
    - z10 EC to z10 EC and z10 EC to System z9/z990/z890
  - ▶ 10 meter distance remains
- **ISC-3 for extended distance over fiber optic cabling**
  - ▶ No change to current cabling
- **Internal Coupling channels (IC)**

\* 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 functional comparison to z9 EC

<b>Processor / Memory</b>	<ul style="list-style-type: none"> <li>▪ Uniprocessor Performance</li> <li>▪ n-way Performance</li> <li>▪ System Capacity</li> <li>▪ Processor Design</li> <li>▪ Models</li> <li>▪ Processing Units (PUs)</li> <li>▪ Granular Capacity</li> <li>▪ Memory</li> <li>▪ Fixed HSA</li> </ul>	<ul style="list-style-type: none"> <li>▪ 62% performance improvement over z9 EC uniprocessor *</li> <li>▪ On average 50% more performance than z9 EC in a n-way configuration</li> <li>▪ 70% system capacity performance improvement over z9 EC 54-way **</li> <li>▪ New 4.4GHz processor chip</li> <li>▪ z10 EC has 5 and z9 EC has 5 models, both with up to 4 books</li> <li>▪ z10 EC has up to 64 PUs to configure, up to 54 on z9 EC</li> <li>▪ z10 EC has up to 100 Capacity settings versus 78 on the z9 EC</li> <li>▪ z10 EC has up to 1.5 TB vs. up to 512 GB on z9 EC</li> <li>▪ z10 EC has fixed 16 GB HSA, z9 EC had HSA carved from purchased memory</li> </ul>
<b>Virtualization</b>	<ul style="list-style-type: none"> <li>▪ LPARs</li> <li>▪ HiperDispatch</li> </ul>	<ul style="list-style-type: none"> <li>▪ z10 EC has up to 64 logical processors in an LPAR versus 54 on z9 EC</li> <li>▪ z10 EC has HiperDispatch for improved synergy with z/OS Operating System to deliver scalability and performance</li> </ul>
<b>Connectivity</b>	<ul style="list-style-type: none"> <li>▪ HiperSockets</li> <li>▪ FICON for SANs</li> <li>▪ Total channels</li> <li>▪ Internal I/O Bandwidth</li> <li>▪ Enhanced I/O structure</li> <li>▪ Coupling</li> <li>▪ Cryptography</li> <li>▪ LAN Connectivity</li> </ul>	<ul style="list-style-type: none"> <li>▪ z10 EC New HiperSockets Layer 2 and Multiple Write Facility</li> <li>▪ Up to 336 FICON channels on z10 EC and z9 EC</li> <li>▪ Same - Up to 1024 channels</li> <li>▪ z10 EC has industry standard 6 GBps InfiniBand supports high speed connectivity and high bandwidth versus z9 EC using 2.7 GBps Self Time Interconnects (STIs)</li> <li>▪ Star L2 Cache Book Interconnect versus Ring Topology interconnect on z9 EC</li> <li>▪ Coupling with InfiniBand<sup>1</sup> – improved distance and potential cost savings</li> <li>▪ Improved AES 192 and 256 and stronger hash algorithm with Secure Hash Algorithm (SHA-512)</li> <li>▪ New OSA-Express3<sup>1</sup> for 10 Gigabit Ethernet connectivity</li> </ul>
<b>On Demand / RAS</b>	<ul style="list-style-type: none"> <li>▪ Capacity Provisioning Mgr</li> <li>▪ RAS Focus</li> <li>▪ Just in Time deployment of Capacity</li> </ul>	<ul style="list-style-type: none"> <li>▪ z10 EC &amp; z/OS (1.9) for policy based advice and automation</li> <li>▪ z10 EC can help eliminate preplanning required to avoid scheduled outages</li> <li>▪ Capacity on Demand offerings CBU and On/Off CoD plus new Capacity for Planned Events are resident on z10 EC</li> </ul>
<b>Environmentals</b>	<ul style="list-style-type: none"> <li>▪ Monitoring</li> </ul>	<ul style="list-style-type: none"> <li>▪ z10 EC displays energy efficiency on SAD screens</li> <li>▪ Utilizes IBM Systems Director Active Energy Manager for Linux on System z for trend calculations and management of other servers that participate</li> </ul>

LSPR mixed workload average running z/OS 1.8 - z10 EC 701 versus z9 EC 701

\*\* This is a comparison of the z10 EC 64-way and the z9 EC S54 and is based on LSPR mixed workload average running z/OS 1.8

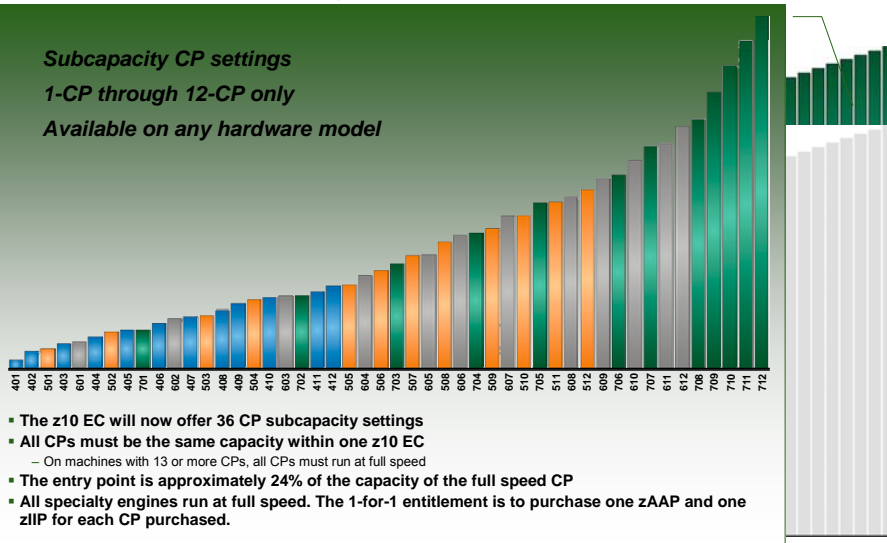
<sup>1</sup> Planned availability 2Q08

## Flexible Capacity

- 100 Capacity Settings
- More ways to get Capacity how and when needed
- CoD
- Capacity Planned Event

## 100 capacity settings to meet your needs

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



- The z10 EC will now offer 36 CP subcapacity settings
- 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 24% of the capacity of the full speed CP
- All specialty engines run at full speed. The 1-for-1 entitlement is to purchase one zAAP and one zIIP for each CP purchased.

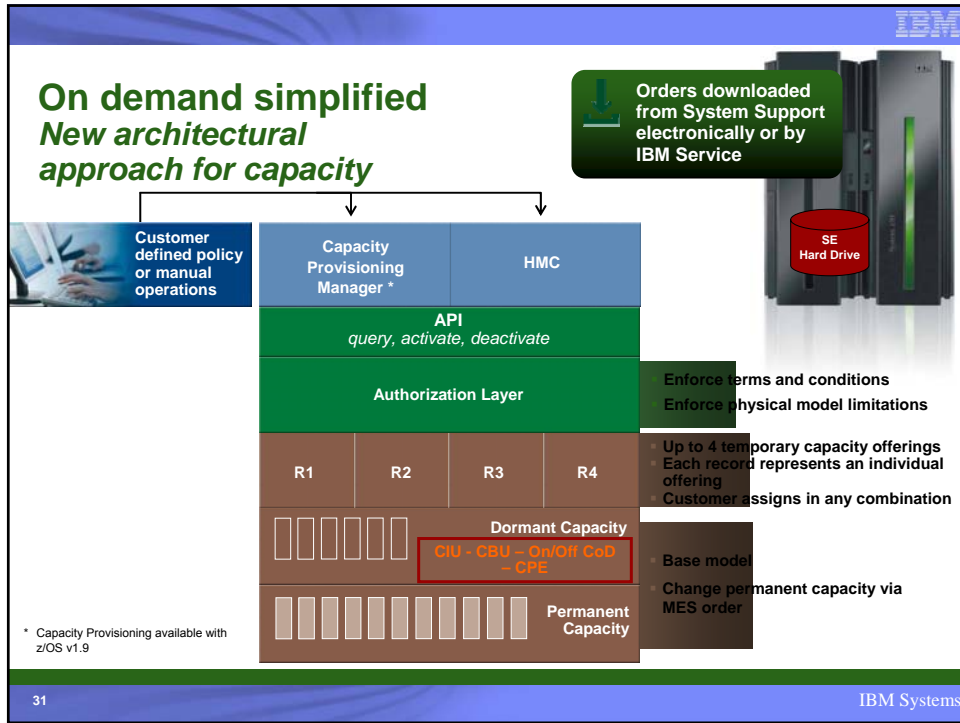
## Just in time capacity gives you control

- **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



## CoD – New architectural approach for z10 EC

- Resources can be activated in any amount up to defined limit
  - **Customer can customize activation real-time, based on circumstances**
  - **Eliminates unique record to be managed for all possible permutations**
  - **Dynamic changes in activation level without reloading records**
- As records expire or are consumed, the resources will be deactivated
  - System will no longer 'throttle' but deactivate capacity when records expire
  - Will not deactivate if removing dedicated engines or last of that engine type
- Various record limits can be dynamically updated / replenished
  - Changes possible even if record is currently active
- Ability to perform permanent upgrades while temporary capacity is active
  - Configuration dependent
  - Allows quick conversion of temporary capacity to permanent
- Separation of Temporary LICCC record download/add and activation/deactivation of resources
- API enhancements to support use by Capacity Provisioning Manager
  - Capacity Provisioning Manager provides policy based automation



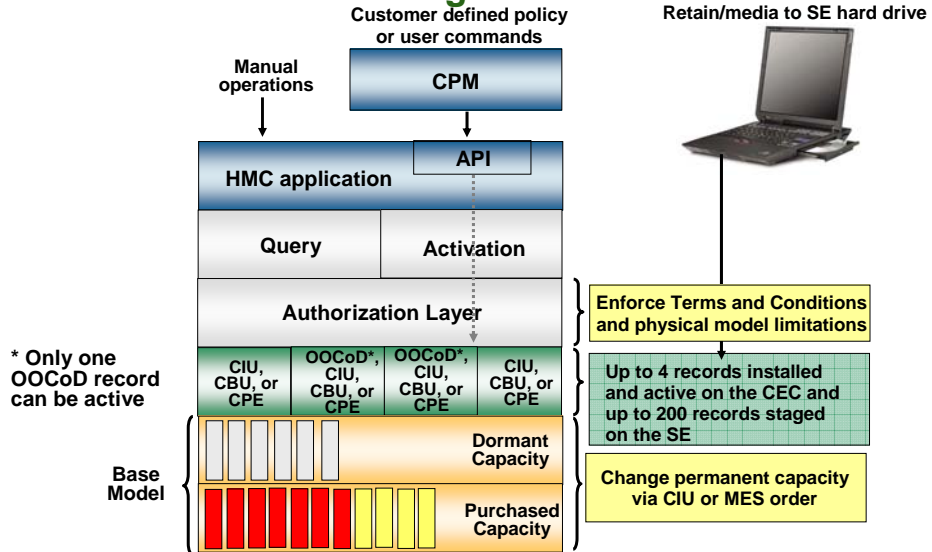
- 
- ## z10 EC CoD Offerings - Details
- **CIU**
    - Customer Initiated Upgrade for permanent increase in capacity
  - **CBU**
    - For disaster recovery
  - **Capacity for Planned Event (CPE)**
    - To replace capacity lost within the enterprise due to a planned event such as a facility upgrade or system relocation
  - **OOCoD**
    - Production Capacity
    - Post-pay with unlimited capacity usage
- 32 IBM Systems



## z10 EC CoD – Key Enhancements

- All offering records are resident on machine
  - No connection or passwords required at time of activation
  - Records are changed only when customer places order for new / updated offering
- Multiple records can be simultaneously active
  - Each has independent controls and policy
  - Each can be activated / deactivated in any sequence
- Individual record can be used to temporarily reach multiple configurations
  - Customer determines level of resources activation real time based on circumstances (i.e. multiple use for a single On/Off CoD record, even during a permanent upgrade)
  - All movement between configurations is concurrent
- More flexibility to configure offering limits
- Ability to perform upgrades while temporary resources are active
  - Modification of record entitlement performed dynamically and concurrently
- “Capacity Provisioning Manager” provides policy based advice and automation

## z10 EC CoD Provisioning Architecture




## Improved RAS

- Reduction in Parts
- Additional Fault Tolerance
- Elimination of many pre preplanning requirements

## Keeping your system available is key to our design Continuing our RAS focus helps avoid outages

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			✓

Increased Focus over time

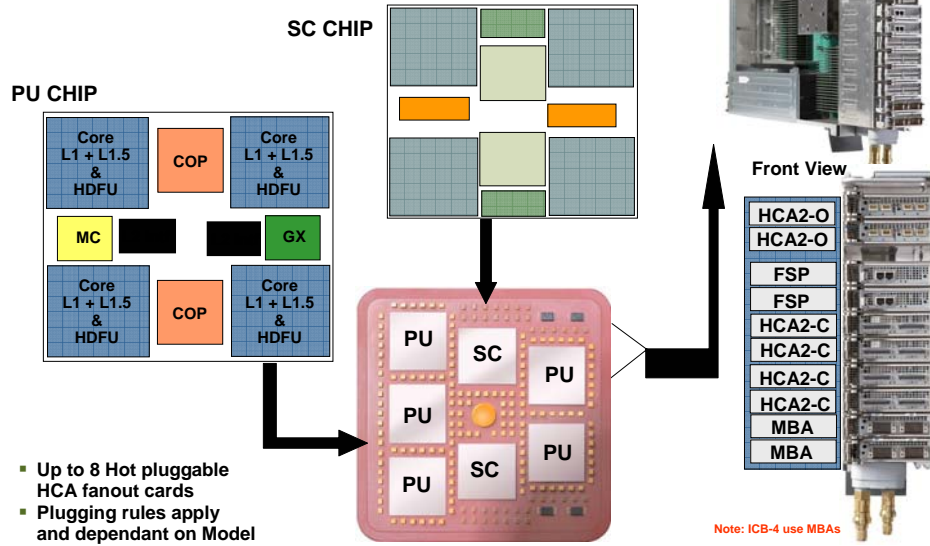
## z10 EC Enhancements designed to avoid Outages

- 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

## z10 EC Processor/Memory/HCA and Book



## 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



## Protecting with IBM's world-class Business Resiliency solutions

- **Preplanning capabilities to avoid future planned outages, e.g. dynamic LPAR allocation without a system outage**
- **100 available capacity settings – 30% more than z9 EC**
- **Integrated enterprise level resiliency for heterogeneous data center disaster recovery management**
- **Policy driven flexibility to add capacity and backup processors**
- **Basic HyperSwap™ improves storage availability \***
- **Integrated cryptographic accelerator**
  - Advanced Encryption Standard (AES) 192 and 256 and Stronger hash algorithm with Secure Hash Algorithm (SHA-512)
- **Tamper-resistant Crypto Express2 feature**
  - Supports high levels of security for demanding applications
  - Fully programmable and configurable
  - High scale performance for SSL transactions
- **Trusted Key Entry (TKE) 5.2 with optional Smart Card reader**
- **System z – the only platform that is EAL5 certified1**



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

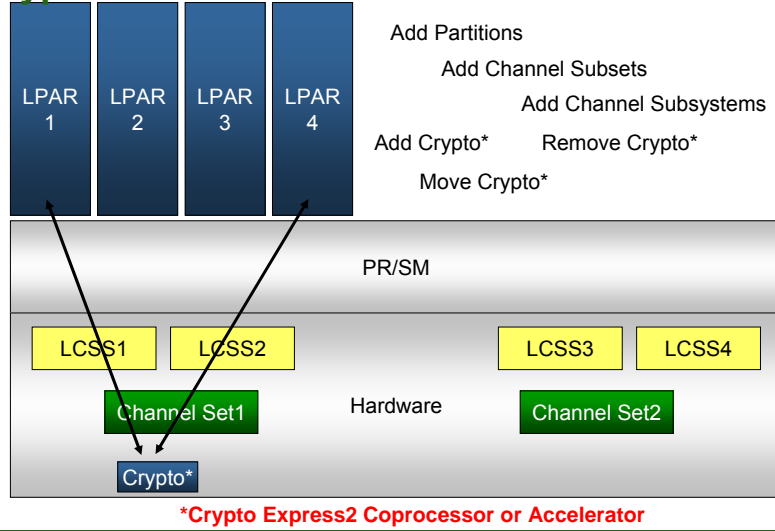
## **z10 EC Service Enhancements designed to avoid Scheduled Outages**

- **Concurrent firmware fixes**
  - MCLs
- **Concurrent Driver Upgrades**
  - Enhanced Driver Maintenance
- **Concurrent parts replacement**
  - N + 1, Fanouts, Power supplies, MRU etc.
- **Concurrent Hardware Upgrades**
  - Book adds, microcode upgrades, I/O cards etc.

## **z10 EC Preplanning Improvements designed to avoid Planned Outages**

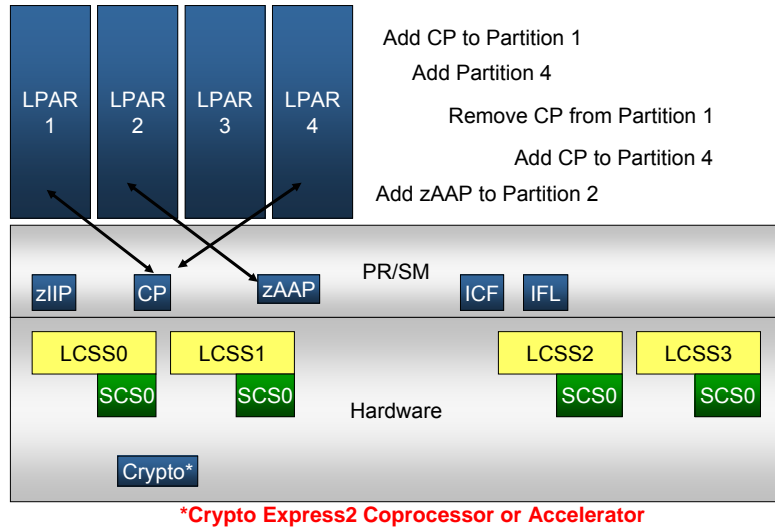
- **Capacity on Demand (CoD) – Flexible Customer initiated Upgrades**
- **Designed for Elimination of unnecessary CBU passwords**
- **Enhanced Driver Maintenance**
  - Multiple 'from' sync points supports for Enhanced Driver Maintenance
- **Designed to eliminate a logical partition deactivate/activate/IPL**
  - Dynamic Change to the Logical Processor Definition and Logical Crypto – z/VM 5.3
  - Dynamic Change to the Logical Cryptographic Coprocessor Definition – z/OS ICSF

## z10 EC Outage Elimination examples for I/O and Crypto



## z10 EC Outage Elimination examples for LPARs

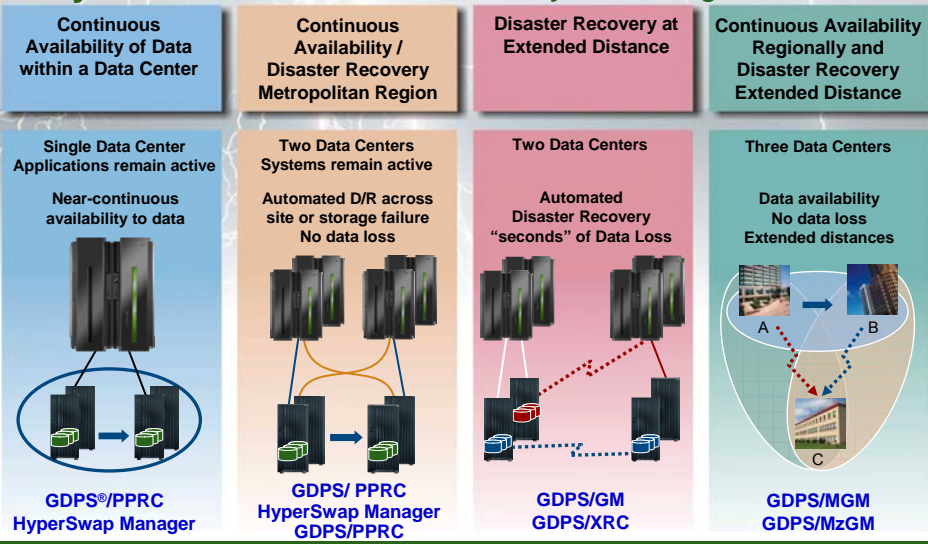
Dynamically Add/Remove CPs, zIIPs, zAAP, ICFs, IFL to Partition w/o Preplanning



## z10 EC RAS Summary

- z10 EC is designed to deliver the industry leading RAS customers expect from System z servers
- z10 EC RAS is designed to reduce all sources of outages by reducing Unscheduled, Scheduled and Planned outages
- Planned outages are further designed to be reduced by eliminating pre-planning requirements
- Designed to reduce need for Power-on-Reset
- Designed to eliminate a logical partition deactivate/activate/IPL

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




## Support for new Security Standards

- **Extension of Open Industry Standards**
- **Additional functionality**

## z10 EC Compression and Cryptography Accelerator

- **Data compression engine**
  - ▶ Static dictionary compression and expansion
  - ▶ Dictionary size up to 64KB (8K entries)
    - Local 16KB 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 2 cores**
  - ▶ **Independent compression engines**
  - ▶ **Shared cryptography engines**





## Encryption Authentication – Digital Certificate Authority on System z10

### Customer Needs

- Deploy secure transactions and network connections
- Securely exchange encryption keys
- Manage the lifecycle of digital certificates

### System z Value Proposition

- Save thousands to millions of dollars by replacing third party hosting of digital certificates
- Full certificate life cycle management
- Highly available and highly secure



### Saving costs

- Relatively low MIPS to drive thousands of certificates using no-cost feature of z/OS

### Industry certification

- Certified by Identrust, a global leader in trusted identity solutions recognized by global financial institutions, government agencies and businesses

### System z robustness

- Leverages resiliency and disaster recovery features
- Highly secure using mainframe access and audit controls and encryption features
- Services available to accelerate

### New ISV support

- Venafi™ solutions for centralized digital certificate management now support z/OS Certificate Authority

## System z as the central repository for sensitive data

Leverage the mainframe security policies and processes that have been developed over many years in your enterprise

- Security-rich holistic design to help protect system from malware, viruses, and insider threats
- Encryption solutions to help secure data from theft or compromise
- Tivoli tools allowing you to address compliance needs with more confidence

The industry's most securable platform!



## The System z and System Storage encryption solution delivers integrated security

### *z/OS Centralized Key Management*

- Can help to protect and manage keys across entire enterprise
  - Highly secure and available key data store
  - Long term key management
  - Disaster recovery capabilities
  - Access control and audit-ability
- Single point of control
- Leverage robustness of z/OS and IBM System z with over a decade of production use



### *System Storage Tape Encryption*

- Designed to provide:
  - ▶ z/OS encryption controlled via Data Policy (SMS) and user Policy (JCL)
  - ▶ Open systems encryption controlled via data source, VolSer or drive
  - ▶ Avoid Host MIPS overhead
  - ▶ Minimize impact to existing processes and applications

## Agenda

1. Green Design
2. Faster Engines
3. New Advanced Function
4. Flexible Capacity
5. Improved RAS
6. Support for New Security Standards
7. Expanding the ability to run more diverse workloads
8. Upgradeability
9. Virtualization
10. Energy Efficiency

## Helping to drive down the cost of IT Now even more workloads can benefit from zIIP

- **zIIP can help to integrate data across the enterprise by optimizing resources and lowering the cost of ownership for eligible data and transaction processing workloads**

- Centralized data serving - First to exploit zIIP were workloads such as BI, ERP, and CRM applications running on distributed servers with remote connectivity to DB2® V8
- Network encryption - zIIP becomes an IPSec encryption engine helpful in creating highly secure connections in an enterprise
- Serving XML data - zIIP is enabled for XML parsing, first to exploit this is inserting and saving DB2 9 XML data over DRDA®
- Remote mirror - zIIP becomes a data mirroring engine with zIIP assisted z/OS Global Mirror function (zGM, formerly XRC) helpful in reducing server utilization at recovery site (with z/OS V1.8 and above)
- Exploiting of zIIPs by ISVs


Available 9/07

new!

new!

- **zIIPs offer economics to help you**

- **PLUS** zIIP price is same for z10 EC as z9 EC and we offer no charge MES upgrades when moving to new technology



## System z10 Designed for New Customer Workloads


**New System z10 based solutions for workloads and Industry processes that leverage the full power of System z10**

- Operational Business Intelligence on System z10
- [SAP Business Intelligence Accelerator on System z10](#)
- Consolidation & Virtualization Server Optimization and Integration Services for System z10
- Encryption Authentication - Digital Certificate Authority on System z10
- SOA Core Banking Renovation & Payments Framework for System z10



## Improved server performance and scalability with faster and more processors and improved dispatching synergy

- The z10 EC delivers on average 50% more performance in a n-way configuration than a IBM System z9™ Enterprise Class (z9 EC) n-way
  - The uniprocessor is expected to deliver 62% more performance than z9 EC uniprocessor \*
- The z10 EC 64-way offers 70% more server capacity than the largest z9 EC\*\*
- Introducing HiperDispatch for improved synergy with z/OS® operating system to deliver scalability and performance



**Significant capacity for traditional growth and consolidation**

\* LSPR mixed workload average running z/OS 1.8 - z10 EC 701 versus z9 EC 701

\*\* This is a comparison of the z10 EC 64-way and the z9 EC S54 and is based on LSPR mixed workload average running z/OS 1.8

## Agenda

1. Green Design
2. Faster Engines
3. New Advanced Function
4. Flexible Capacity
5. Improved RAS
6. Support for New Security Standards
7. Expanding the ability to run more diverse workloads
8. Upgradeability
9. Virtualization
10. Energy Efficiency

## Protecting your investment in IBM technology

- **Designed to protect your investment by offering upgrades from z9 EC and z990 to the z10 EC**
- **Full upgradeability within the System z10 family**
  - Upgrade to Model E64 will require a planned outage
- **Temporary or permanent growth when you need it**
  - New provisioning architecture



## Continuing the modular design for flexibility Facilitates upgradeability and availability

**IBM System z10 Enterprise Class (z10 EC)**  
**Machine Type: 2097**

**5 Models:** E12, E26, E40, E56, E64



### Processor Units (PUs):

- One to four book modular design
- Sub-capacity available up to 12 CPs
- Enterprise Quad Core technology – 4.4 GHz
- Enhanced capacity 64-way model
- 17 PUs per book (17 and 20 for Model E64)
  - New core sparing technology
  - More SAPs per system
  - Configurable PUs allow you to design the system to meet your needs (e.g. CPs, specialty engines, SAPs)

### Memory:

- Up to 1.5 TB / 384 GB per book
- 16 GB HSA separately managed and not included in customer purchased memory
- Books connected in star topology via L2 cache

### I/O:

- 6 GBps InfiniBand host buses for I/O
- FICON™/FCP Enhancements
- New OSA-Express3 10 GbE <sup>1</sup>
- InfiniBand Coupling Links <sup>1</sup>

<sup>1</sup> Planned availability 2Q08

## z10 EC Concurrent Memory Upgrades

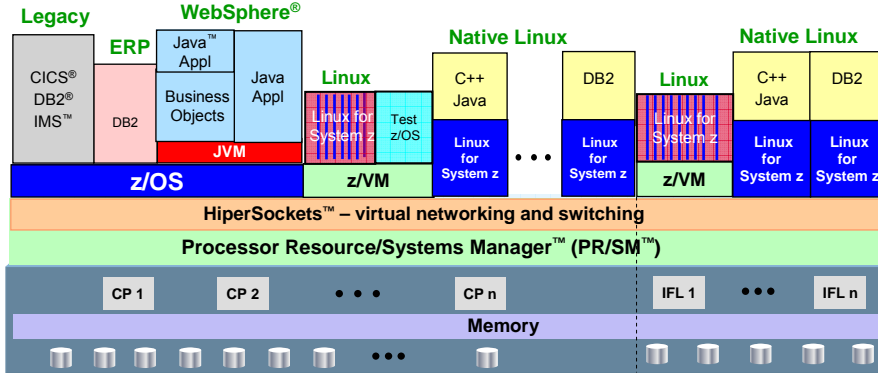
- **LIC enables additional memory to the physical limit of the installed cards and memory configuration**
  - ▶ Designed to be possible and concurrent in many but not all configurations
- **Add a book with additional memory**
  - ▶ Designed to be possible except for Models E56 and E64
- **Exploit Enhanced Book Availability to change memory card configuration in existing books**
  - ▶ Not possible on Model E12
  - ▶ Exploits capability for concurrent book remove, upgrade and return
  - ▶ Designed to be possible 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

Note: Concurrent memory upgrades are designed not to require a Power-on-Reset (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 disruptive to that partition only.

## Agenda

1. Green Design
2. Faster Engines
3. New Advanced Function
4. Flexible Capacity
5. Improved RAS
6. Support for New Security Standards
7. Expanding the ability to run more diverse workloads
8. Upgradeability
9. Virtualization
10. Energy Efficiency

## 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

## Consolidation & Virtualization Server Optimization and Integration Services for System z10

### Customer Needs

- Stop server sprawl
- Reduce energy and related facility costs
- Improve security and reliability
- Increase flexibility for changing business needs
- Get control of workload/systems management

### System z Value Proposition

- Allows customers to rely on GTS consolidation experts skilled on System z10 platform technologies to assess current IT deployment, design and implement changes
- Measure and document improvements and related cost savings
- Gives customers the assistance needed for consolidation projects

... The Power of many



... The Simplicity of ONE

### GTS Services on System z10

- Based on internally developed assessment tools combined with GTS subject matter expertise and System z skills offered to clients on a custom basis tailored to their specific enterprise needs

### Solution Components

- System z10 with z/VM®, Linux
- Services that include:
  - assessment,
  - planning,
  - design,
  - implementation

## Consolidation with Linux gets a “green light”

### System z servers may help customers become more energy efficient:

- Deploy energy efficient technologies – reduce energy consumption and save floor space

### Economics of IFLs and z/VM® help to drive down the cost of IT

- IFLs attractively priced, have no impact on z/OS license fees, and z/VM and Linux software priced at real engine capacity
- ‘No charge’ MES upgrades available when upgrading to new technology




## IT Cost Savings powered by z/VM Virtualization on z10 EC

### Your IT Cost may vary:

- **91% Less Hardware**
  - 304 x86 Processor Cores vs 26 IFLs
- **Potential for dramatic reductions in software expense for processor based licenses**
- **Reductions in power and cooling**
  - 81% Savings in KWatts and Energy Costs in this scenario
- **30% Less Space**
- **93% People savings**
- **Increased processor utilization**
- **Industry leading Security**

### Consolidating 760 Linux servers z/VM versus x86 Virtualization Oracle DB Workload 3-Year Total IT Cost

**\$29.5M Savings versus  
x86 without Virtualization**



Energize your IT savings with z10 EC.




## Agenda

1. Green Design
2. Faster Engines
3. New Advanced Function
4. Flexible Capacity
5. Improved RAS
6. Support for New Security Standards
7. Expanding the ability to run more diverse workloads
8. Upgradeability
9. Virtualization
10. Energy Efficiency

## Tracking energy consumption within the infrastructure

- ResourceLink™ provides tools to estimate server energy requirements before you purchase a new system or an upgrade
- Offers a 15% improvement in performance per kWh over z9 EC
- 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)
- New 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





- IBM System z
- ## System z Energy Efficiency
- **Mainframe focus on system and data center efficiency**
    - ▶ Consolidation of many servers onto one system
    - ☞ Consistent performance at sustained high utilization
      - Resilience in face of changing workloads
    - ▶ Leverages virtualization capabilities of PR/SM, z/VM, z/OS
  - **System z designs are optimized for scale-up data serving**
    - ▶ SMP Hub design enables robust scaling across wide spectrum of workloads
    - ▶ Centralized SMP fabric minimizes fabric logic per core
      - Extended on z10 EC via 4-core processor chip
    - ▶ MRU cooling allows dense package and reduces leakage power
    - ▶ Extensive hardware support for multi-level virtualization
  - **Chip-level power optimization applied to IBM z10 EC design**
    - ▶ Local clock gating to limit maximum dynamic power
    - ▶ Millicode sleep mode for wait/spare/stop states
- z10TLB\_068 IBM Systems

## What are Energy Efficiency Certificates?

- Also known as White Tags
- A new tradable attribute similar (but different process) to RECs
- Represents the value of energy not used (conserved) in Data Centers or other facilities' efficiency programs
- An emerging method to verify your programs by a 3<sup>rd</sup> party
- Created through the implementation of energy conservation projects on the demand side



69

## z10 EC Electrical Energy / Thermal Information

- **Support from System z9 GA3**
  - ▶ System Activity Display (SAD)
  - ▶ Resource Link's Power Estimator tool
  - ▶ Weekly call-home data
- **Enhancements for z10 EC**
  - ▶ SAD
  - ▶ Resource Link's Power Estimator tool
  - ▶ **Weekly call-home data includes power cord information**
  - ▶ **IBM Systems Director Active Energy Manager for Linux on System z support for z10 EC**

## IBM Systems Director Active Energy Manager Introduction

- **IBM Systems Director Active Energy Manager (AEM) is an energy management solution building block that returns true control of energy costs to the customer**
- **AEM is a cornerstone of the IBM energy management framework and is leading edge in the Industry**
- **In tandem with chip vendors Intel and AMD and consortiums like the Green Grid, AEM supports the IBM initiative to deliver price performance per square foot**
- **AEM runs on Windows, Linux on IBM System x™, Linux on IBM System p™, and Linux on IBM System z. Refer to its documentation for more specific information.**

## IBM System z10 EC – 4 Steps to Maximizing your IT ROI

1



### **Lower the running costs of existing IT**

Reduce the operating costs of existing workloads with the improved price performance and technology driven dividends of a new z10 EC mainframe.

2



### **Manage growth, complexity and risk**

Scalable products and solutions you can trust to more easily and securely manage the complex world of IT.

3



### **Go green and save**

Cut costs and “go green” with leadership energy-efficient hardware, consolidation and virtualization capabilities on System z10 EC

4



### **Realize innovation**

Technology that makes innovation real in your business and sets you apart from the competition.

## The New Enterprise Data Center



- **New economics:** Virtualization breaks the lock between IT users and IT resources
- **Rapid service delivery:** Service management enables visibility, control and automation to deliver quality service at any scale
- **Aligned with business goals:** Real-time integration of transactions, information and analytics - and delivery of IT as a service

IBM

© 2008 IBM Corporation

An evolutionary new model for  
**efficient enterprise  
IT delivery**  
addresses a changing  
landscape



IBM

© 2008 IBM Corporation

## Stages of adoption

### ***Simplified***



### ***Shared***



### ***Dynamic***



IBM

© 2008 IBM Corporation

## Summary

- Companies in every industry are exploiting advanced technology to gain competitive advantage
- Infrastructure complexity and rising energy costs are driving higher operational costs for companies and constraining their growth
- The New Enterprise Data Center offers an evolutionary new model for efficient enterprise IT delivery addressing this changing landscape



IBM

© 2008 IBM Corporation

**IBM.**

## How can you help to Green Data Centers, save Costs and maybe the Earth?

- Reduce waste and inefficiencies
- Power down resources when not in use
- Push Utilization levels to manageable limits
- Include Energy Consumption in Business Case
- Take a Holistic approach to System Evaluations
- Conduct a “Green Assessment”
  - Green Initiative
- Build a roadmap to an Optimized IT Infrastructure
  - Create a plan from which tactical decisions can be made
  - Leverage IBM OIT Offerings

© 2008 IBM Corporation 77

**IBM**

## IBM System z10™ Enterprise Class (z10™ EC) A marriage of evolution and revolution

### Evolution

- Scalability and virtualization to reduce cost and complexity
- Improved efficiency to further reduce energy consumption
- Improved security and resiliency to reduce risk
- New heights in storage scalability and data protection

### Revolution

- 4.4 GHz chip to deliver improved performance for **CPU intensive workloads**
- **‘Just in time’ deployment of capacity resources**
- Vision to expand System z capabilities with Cell Broadband Engine™ technology

78 IBM Systems

Thank  
You