



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***

April 19, 2008



The Future Runs on System z

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	z/OS*
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 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

0. Operating System Support

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
11. Will there be a z10 BC?

System z10 EC Operating System Support

Operating System	ESA/390 (31-bit)	z/Architecture (64-bit)
z/OS Version 1 Releases 7 ⁽¹⁾ , 8 and 9	No	Yes
Linux on System z ⁽²⁾ , RHEL 4, 5 & SLES 9, 10	No	Yes
z/VM Version 5 Release 2 ⁽³⁾ and 3 ⁽³⁾	No	Yes
z/VSE Version 3 Release 1 ⁽²⁾⁽⁴⁾	Yes	No
z/VSE Version 4 Release 1 ⁽²⁾⁽⁵⁾	No	Yes
z/TPF Version 1 Release 1	No	Yes
TPF Version 4 Release 1 (ESA mode only)	Yes	No

1. z/OS R1.7 + zIIP Web Deliverable required for z10 EC to enable HiperDispatch
2. Compatibility Support for listed releases. Compatibility support allows OS to IPL and operate on z10 EC
3. Requires Compatibility Support which allows z/VM to IPL and operate on the z10 EC providing System z9 functionality for the base OS and Guests.
4. z/VSE v3. 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities. z/VSE is designed to exploit select features of IBM System z10, System z9, and zSeries hardware.
5. z/VSE V4 is designed to exploit 64-bit real memory addressing, but will not support 64-bit virtual memory addressing

Note: Refer to the z/OS, z/VM, z/VSE subsets of the 2097DEVICE Preventive Planning (PSP) bucket prior to installing a z10 EC

System z10 EC Minimum Operating System Support for New functions – 1

	z/OS(**)	z/VM(**)	Linux on System z(**)	z/VSE(**)	z/TPF TPF(1)(**)
Basic System z10 EC support	1.7 ⁽³⁾	5.2	SLES 9 RHEL 4	3.1	1.1 4.1 ⁽¹⁾
HiperDispatch	1.7 ⁽³⁾	Not Supported	SLES 11 RHEL 6		
STSI for Capacity Provisioning	1.7 ⁽³⁾	5.2	SLES 10 SP2 RHEL 6		
Capacity Provisioning	1.9 ⁽³⁾	Not Supported	IBM work with LDPs ⁽²⁾		
Large Page (1MB)	1.9 ⁽³⁾	Not Supported	IBM work with LDPs ⁽²⁾		
RMF Enhancements for FICON	1.9	Not Supported	Not Supported		
HW Decimal Math Support	1.7 ⁽³⁾	5.3	IBM work with LDPs ⁽²⁾		
z/VM-Mode partitions	1.7 ⁽³⁾	SOD*	IBM work with LDPs ⁽²⁾		
CPACF Enhancements	1.7 ⁽³⁾	5.2 (Guests)	IBM work with LDPs ⁽²⁾	4.1 ⁽³⁾	
Configurable Crypto Express2	1.7 ⁽³⁾	5.2 (Guests)	SLES 9 SP3 RHEL 4.4	3.1	1.1
Dynamically Add Crypto to LPAR	1.7 ⁽³⁾	5.2 (Guests)	SLES 10 SP1 RHEL 5.1		

1. Indicates TPF
2. This function will be provided in a future Linux on System z distribution release/service updates. IBM is working with Linux distribution partners (LDPs) on Kernel space exploitation
3. Additional features, service or Web downloads required

SLES = SUSE Linux Enterprise Server
RHEL = Red Hat Enterprise Linux

** Note: Please refer to the latest PSP bucket for latest PTFs for z10 EC Compatibility and new functions/features support.

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

System z10 EC Minimum Operating System Support for New functions – 2

	z/OS(**)	z/VM(**)	Linux on System z(**)	z/VSE(**)	z/TPF TPF(1)(**)
InfiniBand Coupling Links*	1.7 ⁽³⁾	5.3 (Dynamic I/O)	Not Supported		
STP NTP Client Support	1.7 ⁽³⁾	Not Supported	Not Supported		
OSA-Express3 10 Gbps – CHPID OSD	1.7	5.2	SLES 9 RHEL 4	3.1	4.1 & 1.1
HiperSockets Multi Write Facility (2Q08)	1.9 ⁽³⁾	Not Supported	Not Supported		
HiperSockets Layer 2 Support	Not Supported	5.2 (Guests)	SLES 10 SP2, 11 RHEL 6		
Enhanced FCP caching	Not Supported	5.2			
64-way support	1.9	5.3 (32-way)	SLES 9		1.1
Preserve CTC Logical Path	1.7 ⁽³⁾	5.2	Not Supported		
Fabric Config Support for NPIV	Not Supported	5.2			
1 TB/LPAR (z10 HW supports 1.5 TB)	1.8	Not Supported	SLES 9 (4 TB)		
256 GB	1.8	5.3			

1. Indicates TPF
2. This function will be provided in a future Linux on System z distribution release/service updates. IBM is working with Linux distribution partners (LDPs) on Kernel space exploitation
3. Additional features, service or Web downloads required

SLES = SUSE Linux Enterprise Server
RHEL = Red Hat Enterprise Linux

** Note: Please refer to the latest PSP bucket for latest PTFs for z10 EC Compatibility and new functions/features support.

* 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 z/OS and z/OS.e Support Summary



		z890 (WdfM)	z990*** (WdfM)	z9 EC	z9 BC	z10 EC	End of Service	Coexists with z/OS	Ship Date
z/OS & z/OS.e	1.7	X	X	X	X	X	9/08	1.9	9/05
-"-	1.8	X	X	X	X	X	9/09*	1.10*	9/06
z/OS	1.9	X	X	X	X	X	9/10*	1.11*	9/07

Note: z/OS R1.7 + zIIP Web Deliverable required to use HiperDispatch

z/OS.e - z800, z890 and z9 BC only. Release 1.8 will be the last release of z/OS.e.
Only service-supported releases can coexist in the same sysplex

** Planned. 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.

*** WdfM for EMEA countries. Rest of the world June 30, 2008

z/OS Release 1.9

Improving Usability and Skills

Health Checker improvements & checks, ISPF, DFSMSrmm™, Configuration and Management Usability for Communications Server, HCM, CF management...

Integrating new Applications and Supporting Industry and Open Standards

XML offload to zIIPs & zAAPs, System REXX™, SDSF REXX, Metal C, NFS V4 Server, pthread enhancements, Decimal Floating Point, porting enablement, Binder, PKCS#11...

Extending the Network

Policy-based TCP/IP Routing, Centralized Policy-Based Networking, Expanded use of AT-TLS, FTP Unicode support, new Network Management Interfaces

Scalability and Performance

64-way support, 64-bit GRS, SMF to Logger, TSO/E support for large sequential data sets, Message Flooding Automation, XCF CDS Performance, heap pools, cache alignment, LAN idle, VSCR

Optimize your z/OS environment



Enterprise-Wide Roles

CIM monitoring enhancements, Updated Pegasus server, DFSMSrmm CIM Update, IRMM, ARM 4.1 support

Improving Availability

Improved latch contention detection, CF duplexing and Logger enhancements, SFM improvements, New RRS options ...

Self-Managing Capabilities

WLM support for cross-system routing of zAAP workloads, WLM “Trickle” Support, Promotion of canceled jobs, Start servers in parallel, RMF reporting for CF structures...

Enhancing Security

IPSec offload to zIIP, Additional password phrase support, Kerberos AES & Enhanced CRL support, PKI Services & RACF extensions, z/OS UNIX System Services auditability, Java user and group SAF admin classes, Crypto, NAS AES ...

z/OS 1.9 Summary for z10*

▪ z/OS V1.9

- 64-way support for a single z/OS image
- HiperDispatch
- Up to 4 TB Real Memory (Up to 1.5 TB per server and 1 TB per LPAR on z10 EC Models E56 and E64 servers)
- Hardware Decimal Floating Point**
- Capacity Provisioning
- Large (1 MB) Page support
- Parallel Sysplex support for InfiniBand Coupling links
- System Data Mover (SDM) offload to zIIP
- NTP support for STP
- OSA-Express3 10 Gbps – CHPID OSD
- HiperSockets Multi Write Facility
- Crypto Exploitation
- 4096-bit RSA support

* Additional features, service or Web downloads required

** Level of decimal floating-point exploitation will vary by z/OS release and PTF level.
z/OS V1.7 with PTFs (for High level Assembler support).
z/OS V1.9 with PTFs for full support, for C/C++.

z/OS 1.10 Preview*

Improving Usability and Skills

Configuration Assistant, Health Checker & checks, root migration to zFS, PSP processing, CEEPARM syntax checks, DFSMSrmm and DFSMShsm, ISPF, ...

Integrating new Applications and Supporting Industry and Open Standards

HLASM source-level dbx debugging, submit from z/OS UNIX shell, XML offload & validating parser, CEEROPT for batch, NFS V4 enhancements, FTP from Java, Mixed addressing for METAL C, ...

Extending the Network

Hipersockets Multiple Write, TCP/IP and SNA Performance, multiple VLAN support, Auditability enhancements, Load Balancing Subplex support ...

Scalability and Performance

64-way support, HiperDispatcher, Extended Address Volumes, InfiniBand Coupling, 64-bit Common, Improved XCF locking, Metro Mirror devices in Channel Set 1, ...

Optimize your z/OS environment



Improving Availability

Auto-IPL and SADMP, Dynamic JES2 exits, JES2 NJE connection recovery, Basic HyperSwap support, Expanded ASID reuse, SDM offload to zIIP, RACF database integrity, Consoles 2B, DFSMSrmm recovery ...

Self Managing Capabilities

Policy-based capacity provisioning, Contention Mgt Phase 3, PB delays, DATACLAS JCL overrides, More DFSMSrmm/IRMM integration, zIIP CPU Management ...

Enhancing Security

Password phrase exploitation in z/OS UNIX, rlogin, Language Environment, TSO/E, Kerberos, and LDAP; UTF8 CA support in PKI Services, System SSL and Crypto improvements, Custom Fields ...

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

z/OS 1.10* Summary

- **z/OS V1.10**
 - **z/OS V1.10**
 - **64-way support for single z/OS image**
 - **HiperDispatch support**
 - **PS InfiniBand Coupling support**
 - **Large PAGE support**
 - **Capacity provisioning**
 - **System Data Mover (SDM) offload to zIIP**

* 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 z/VM & z/VSE Support Summary



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

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

*z/VSE V3 can execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities. z/VSE V3 is designed to exploit select features of IBM System z10, System z9, and zSeries hardware.

Note: z/VSE V4 is designed to exploit 64-bit real memory addressing, but will not support 64-bit virtual memory addressing

** Planned. 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.

*** WdfM for EMEA countries. Rest of the world June 30, 2008

Linux and z/VM on System z

Security capabilities:

- ✓ Privacy
- ✓ Regulatory requirements
- ✓ Identify management
- ✓ zSeries qualities of service
- ✓ Common Criteria Certification
- ✓ Ethical hacking
- ✓ HiperSockets

Operational simplification capabilities:

- ✓ Virtualization
- ✓ Simulation
- ✓ Single point of control
- ✓ Large single system image
- ✓ z/OS similarities and synergies
- ✓ Highly granular resource sharing

Consolidation capabilities:

- ✓ Server, network, storage
- ✓ Staff and skills
- ✓ Applications and utilities



Business resiliency capabilities:

- ✓ High availability
- ✓ Disaster recovery
- ✓ Serviceability
- ✓ Reliability
- ✓ HyperSwap
- ✓ XRC and PPRC

On demand infrastructure:

- ✓ Scale up and scale out
- ✓ Rapid server (de)commissioning
- ✓ Idle servers don't consume system resources

Proximity to data:

- ✓ Can increase transaction throughput
- ✓ Shared data access
- ✓ Integrated storage management
- ✓ HiperSockets

z/VM Version 5 Release 3 New Function Highlights

- **Processor and I/O support**
 - Single-image support for up to 32 CPUs
 - Guest support for specialty processors
 - Enhanced SCSI support
 - HyperPAV support for IBM DS8000
 - FlashCopy® I/O support enhancements
- **Virtualization support**
 - Additional z/VM support for large real memory configurations (256 GB)
 - Collaborative Memory Management Assist
 - OSA-Express2 and OSA-Express3 link aggregation
 - Virtual Switch SNMP agent support
 - Enhanced usability for z/VM Virtual Switch and Guest LAN support
 - Guest program-directed IPL support
 - Guest MIDAW support
 - Guest ASCII console support
- **Networking**
 - z/VM TCP/IP support enhancements
 - Enhanced IP failover
 - Currency support for SSL server
 - Dynamic SSL/TLS support
- **Systems management**
 - z/VM system management API enhancements
 - z/VM integrated systems management with the System z HMC
 - User Directory COMMAND support
 - Asynchronous CP command API
 - RACF support enhancements
 - Password phrase support
 - LDAP server and client support
 - Installation and service enhancements
 - Performance Toolkit support



z/VM Release Support for z10 EC

- **z/VM 5.2 and 5.3 with the PTFs for APARs VM64180 and VM64242 allow guests to exploit the System z10 EC at the System z9 functionality level**
- **Exploitation of selected functions of the System z10 EC include the capability for:**
 - **Dynamic I/O to define, modify, and delete a Coupling using InfiniBand link type CIB when z/VM 5.3 is the controlling LPAR for I/O**
 - **Processors to be dynamically added to or removed from a z/VM LPAR with the PTF for APAR VM64249 and VM64323**
 - **TCP/IP and VSWITCH gaining the performance benefit of OSA-Express3 using only one port per CHPID**
- **Improved memory management algorithms to help benefit paging workloads with large memory environments with the PTF for APAR VM64349**

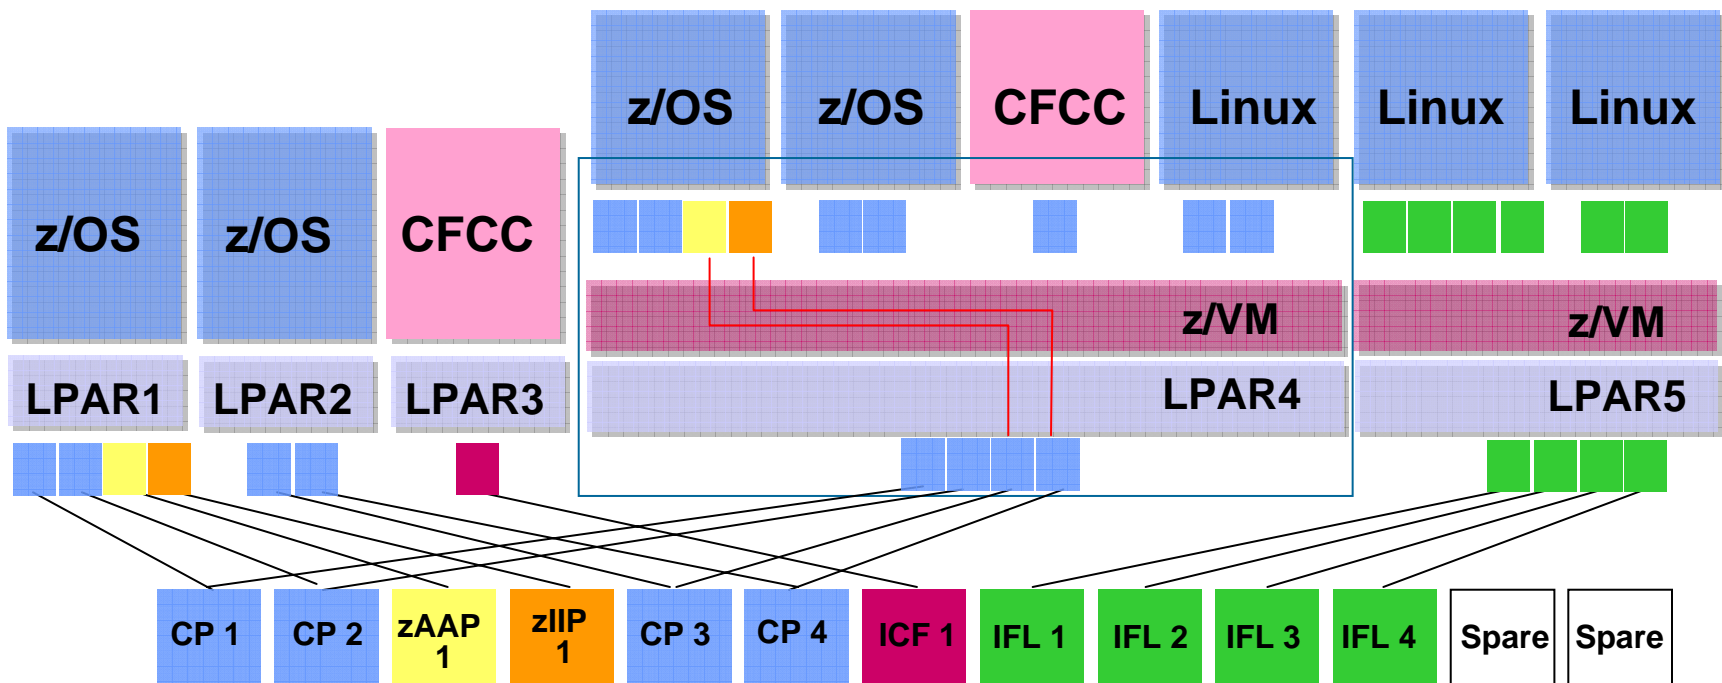
z/VM Support for Specialty Processors

- **z/VM 5.3 introduced support for zAAP and zIIP specialty processors**
 - **System z Application Assist Processors (zAAPs)** – provide an economical Java execution environment for z/OS and z/OS.e
 - **System z9 and System z10 Integrated Information Processors (zIIPs)** – designed to help improve resource optimization and lower the cost for eligible z/OS and z/OS.e workloads by offloading software system overhead from standard Central Processors (CPs); this includes certain DB2 processing
- **z/VM support is provided for z/OS guest exploitation**
 - **Offers additional hardware support for z/OS-on-z/VM development and test support**
- **Two levels of z/VM support:**
 - **Simulation support**
 - **z/VM dispatches virtual zAAPs and zIIPs on real CP engines**
 - **Only possible if the underlying hardware is capable of supporting the real engine type**
 - **Does not require activation of real specialty engine(s) within the mainframe server**
 - **Virtualization support**
 - **z/VM dispatches virtual zAAPs and zIIPs on corresponding real specialty engines**

Specialty Processor Support Example

Simulating Specialty Engines in Virtual Machines

- Allows users to discover the operational aspects of using zAAPs and zIIPs in a z/OS environment without having to purchase real specialty processors
- May help users assess specialty-processor eligible workloads in a z/OS environment
- Provides a function test environment for z/OS workloads that use specialty processors
- Consumes CP processor capacity to host virtual zAAP and zIIP processor cycles*

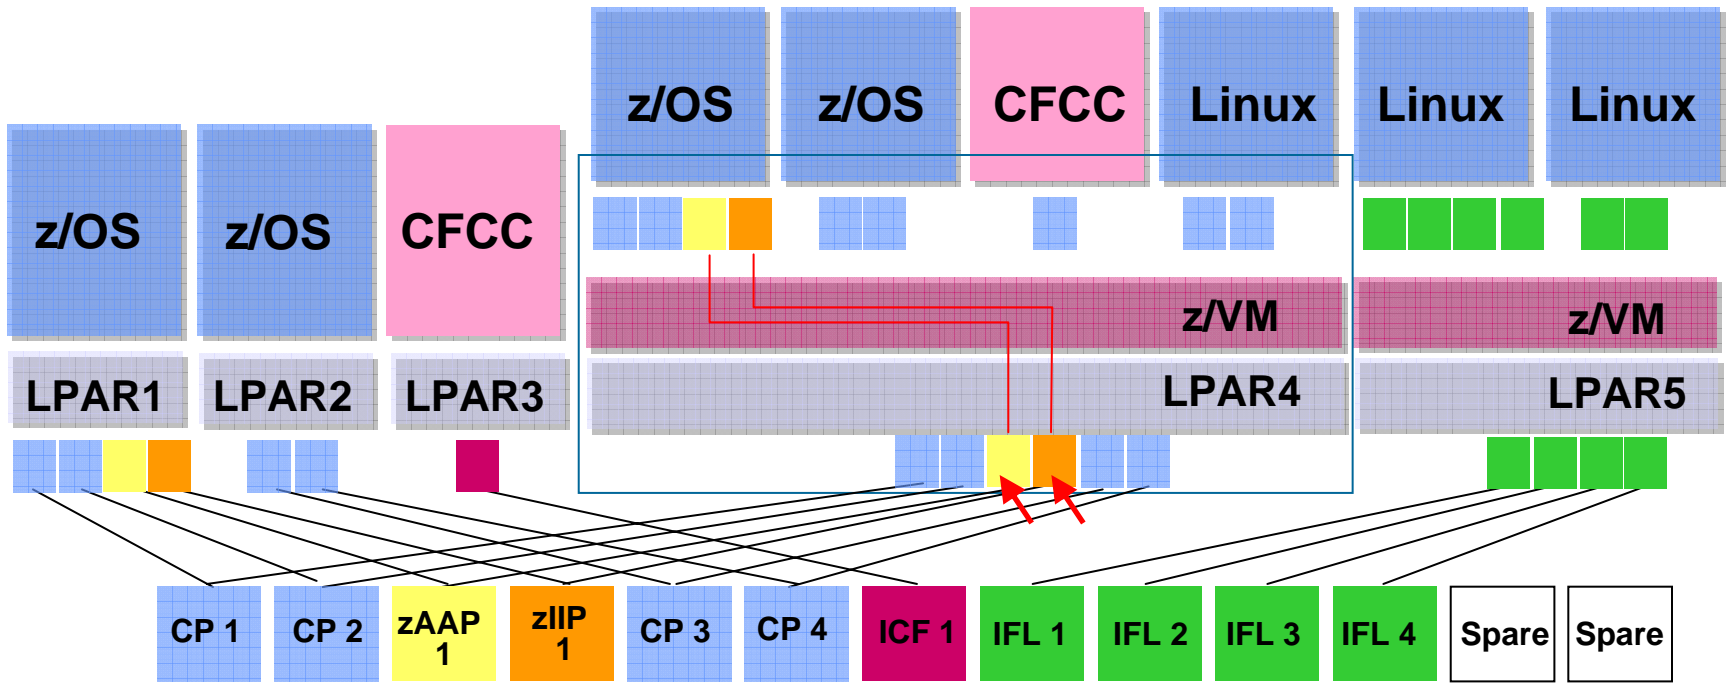


* Specialty engine simulation requires underlying hardware to support the real specialty engine type

Specialty Processor Support Example

Using Real Specialty Engines in Virtual Machines

- Allows users to test and verify z/OS specialty processor support on the real hardware
- Users can maximize real specialty processor utilization by sharing processors among production and test LPARs
- Consumes specialty processor capacity to host virtual zAAP and zIIP processor cycles

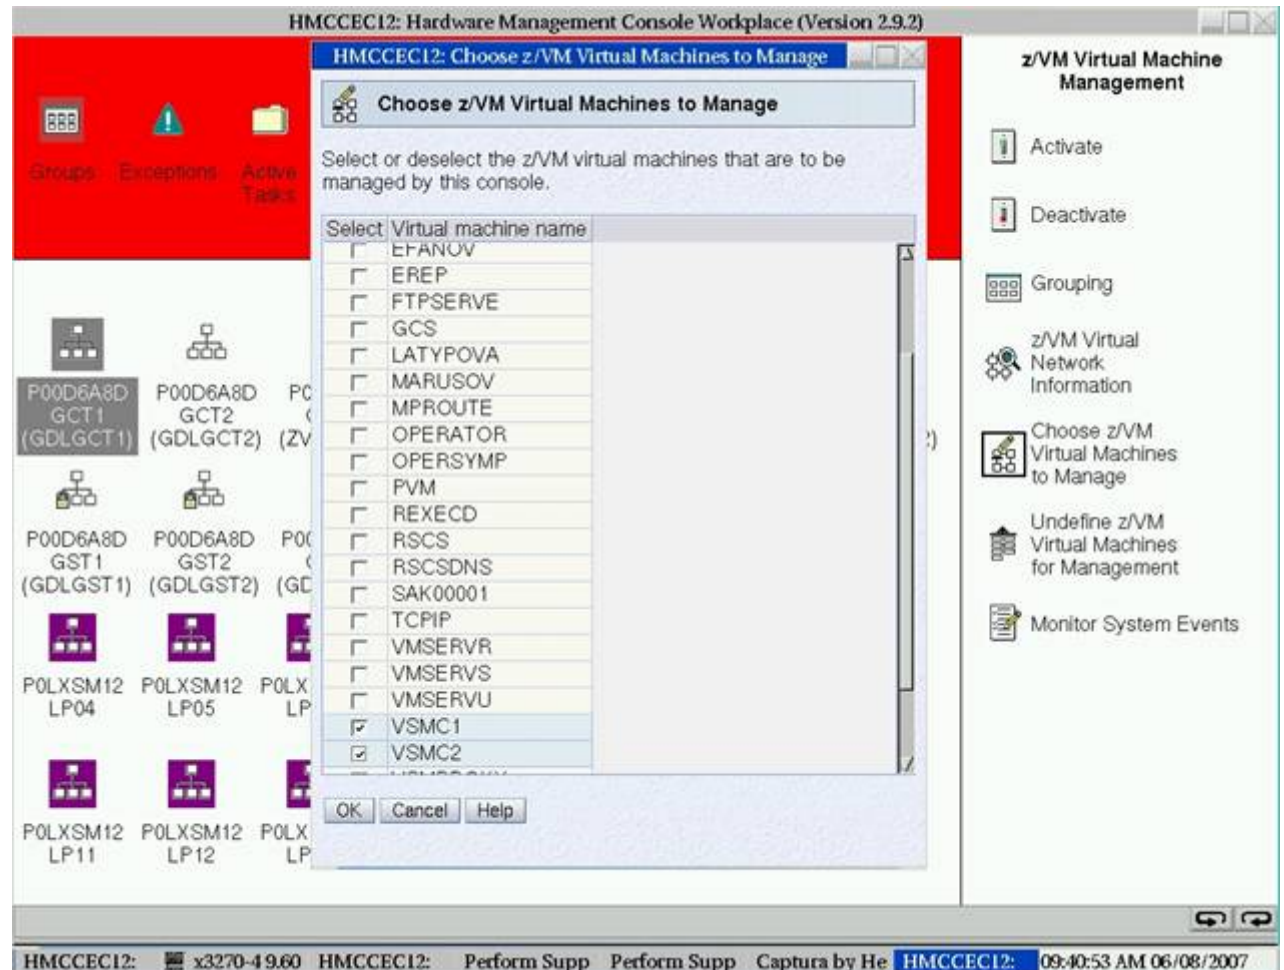


Additional z/VM Guest Support

- **ASCII console support**
 - Allows the real System z ASCII console to be dedicated to a guest system
 - Can facilitate recovery of a Linux guest system during an emergency situation
 - Supports a VT220 data stream
- **Modified Indirect Data Address Words (MIDAW) support**
 - Allows guest use of MIDAWs when z/VM is running on MIDAW-capable servers
 - z/Architecture MIDAW facility offers an alternative to using CCW data chaining in channel programs
 - May reduce channel, director, and control unit overhead by reducing number of CCWs and frames that have to be processed
 - May improve I/O throughput, especially on faster FICON channels
 - Allows z/OS guests to exercise their MIDAW support in a z/VM test environment
- **Program-directed IPL support**
 - Enables a virtual machine to programmatically re-IPL using CCW-type or List-Directed (SCSI) IPL architectures
 - Linux is the exploiter of this function
 - Includes support to allow the setting and storing of IPL parameters

z/VM Integrated Systems Management using the HMC

- **Included in z/VM V5.3**
 - Allows basic z/VM functions to be performed from HMC
 - Network connection not required
 - Uses SCLP hardware interface to access z/VM systems management APIs
 - Requires PTFs for APARs VM64233 and VM64234
- **Supported operations:**
 - View z/VM guests
 - Activate z/VM guests
 - Deactivate z/VM guests
 - Display guest configuration and status
- **Supported systems:**
 - z10 EC, z9 EC and BC
 - z800, z900, z890, z990



z/VSE Comparison

- **z/VSE V3.1***
 - **ESA/390 (31-bit) mode only**
 - up to 2GB real processor storage
 - **System z9 EC and z9 BC**
 - **zSeries 990, 890, 900, 800**
 - **Multiprise[®] 3000 & S/390 G5/G6**
 - **GMLC, GOLC, zELC, TWLC, etc.**
 - **HiperSockets**
 - **CPACF**
 - **Crypto Express2 (configurable)**
 - **FCP/SCSI disks & NPIV**
 - **DS8000, DS6000[™], ESS**
 - **FICON Express2 & 4**
 - **OSA-Express2**
 - **31-bit buffers for ACF/VTAM[®] (via PTF)**
 - **TS1120 encrypting tape**
- **z/VSE V4.1**
 - **z/Architecture (64-bit) mode only**
 - up to 8 GB real processor storage
 - **System z9 EC and z9 BC**
 - **zSeries 990, 890, 900, 800**
 - **MWLC Pricing Metric (System z9 only)**
 - **Full-capacity and sub-capacity mode**
 - **HiperSockets**
 - **CPACF + enhancements**
 - **Crypto Express2 (configurable)**
 - **FPC/SCSI disk & NPIV + point-to-point**
 - **DS8000, DS6000, ESS**
 - **FICON Express2 & 4**
 - **OSA-Express2**
 - **31-bit buffers for ACF/VTAM**
 - **TS1120 encrypting tape**

*z/VSE V3 can operate in 31-bit mode only. It does not implement z/Architecture and specifically does not implement 64-bit mode capabilities. z/VSE V3 is designed to support selected features of IBM System z hardware

z/VSE Version 4 Release 1 Highlights

- **Hardware Support**
 - z/Architecture mode only⁽¹⁾
 - Up to 8 GB real storage
 - IBM System z9 EC / z9 BC servers
 - IBM eServer zSeries 990, 890, 900, 800
 - Open Systems Adapter (OSA)-Express2 and FICON Express4 adapters
 - IBM System Storage DS8000 and DS6000 (both as ECKD™ and FCP-attached SCSI)
 - N_Port ID Virtualization (NPIV)
- **Encryption enhancements**
 - CPACF enhancements (AES-128)⁽²⁾
 - Configurable Crypto Express2 (add accelerator option)
 - 2048-bit RSA keys with Crypto Express2
 - System managed encryption with an IBM System Storage TS1120⁽²⁾
 - Secure FTP⁽²⁾
- **Security**
 - Basic Security Manager enhancements
 - Enhanced logging and reporting
- **New Midrange Workload License Charge (MWLC) pricing metrics**
 - Based on Capacity Measurement Tool (CMT)
 - Attractive full-capacity MWLC price points
 - Sub-capacity MWLC option for added price/performance
- **Interoperability improvements**
 - SOA and interoperability improvements
 - VTAPE enhancements
 - VSAM Redirector enhancements
 - Backup using Tivoli Storage Manager™ (TSM)
 - TSM running on supported non-VSE server
- **Networking**
 - ACF/VTAM for VSE/ESA V4.2 enhancements
 - 31-bit I/O buffer support
 - TCP/IP for VSE/ESA enhancements
 - OSA-Express2 OSN (Open System Adapter for NCP) support
- **Fast Service Upgrade from z/VSE 3.1**
- **Requires z/VM 5.2 (or later) if running under VM**

1. z/VSE V4 is designed to exploit 64-bit real memory addressing, but will not support 64-bit virtual memory addressing

2. Provided via PTF after GA

System z TPF and z/TPF Support



		z890 (WdfM)	z990* (WdfM)	z9 EC	z9 BC	z10 EC	End of Service	Ship Date
TPF	4.1	x ^c	x ^c	x ^c	x ^c	x	TBD	2/01
z/TPF	1.1	x	x	x	x	x	TBD	9/05

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

x^c – Supports up to 30 LPARs with PJ29309

* WdfM for EMEA countries. Rest of the world June 30, 2008

TPF Products

▪ TPF4.1

- 31-bit mode
- z/OS development environment
- 16 engines/processor
- 16 processors/cluster
- Sysplex Timer (9037)
- PSLC
- 4000 modules
- Near end of life
- EOM/EOS comments

▪ z/TPF 1.1

- 64-bit mode
- Linux development environment
 - GCC
 - HLASM for Linux
- 32 processors/cluster
- Up to 84* engines/processor
- 40,000 modules
- STP planned
- No zIIP/zAAP support
- WLC

z/TPF

- **A 64-bit operating system that allows you to move legacy applications into an open development environment, leveraging large scale memory spaces for increased speed, diagnostics and functionality**
- **The open development environment allows access to commodity skills and enhanced access to open code libraries, both of which can be used to lower development costs**
- **Large memory spaces can be used to increase both system and application efficiency as I/Os for memory management can be eliminated**

z10 EC z/TPF and TPF Support

- **Compatibility for:**
 - TPF 4.1
 - z/TPF 1.1
- **z/TPF 1.1 support for:**
 - Up to 64 CPs
 - InfiniBand Coupling
 - OSA
 - H/W Crypto (Crypto Express2 – Accelerator option)
- **Exploitation items for GA1 (Delivered via PUT tapes:)**
 - OSA frame invalidation - Required for OSA-Express3 (TPF 4.1 & z/TPF)
 - Multi-port OSA adapter support (z/TPF)

System z Linux Support



	z890 (WdfM)	z990** (WdfM)	z9 EC	z9 BC	z10 EC	Ship Date
SLES 9	x	x	x	x	x	08/04
SLES 10	x	x	x	x	x	07/06
RHEL 4	x	x	x	x	x	02/05
RHEL 5	x	x	x	x	x	03/07

	General support	Extended support	Self support
SLES 9 support*	07/30/2009	07/30/2011	07/30/2014
SLES 10 support*	07/31/2011	07/31/2013	07/31/2016
	Full support	Deployment support	Maintenance support
RHEL 4 support*	08/31/2007	02/29/2008	02/29/2012
RHEL 5 support*	03/31/2010	09/30/2010	03/31/2014

For latest information and details contact your Linux distributor

Recommendation: use SLES 10 or RHEL 5 for new z10 EC projects

SLES = SUSE Linux Enterprise Server
RHEL = Red Hat Enterprise Linux

* Support dates may be changed by Linux distributors

** WdfM for EMEA countries. Rest of the world June 30, 2008

Linux on System z – Compatibility Plan for z10 EC

Compatibility

- z10 EC tolerated by existing Linux on System z distributions* (most recent service levels):
 - **Novell SUSE SLES9**
 - **Novell SUSE SLES10**
 - **Red Hat RHEL4**
 - **Red Hat RHEL5**

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

Linux on System z – Exploitation Plan for z10 EC

Exploitation

- z10 EC exploitation will be supported by the following Linux on System z distributions*:
 - **Novell SUSE SLES10 SP2**
 - **Novell SUSE SLES11**
 - **Red Hat RHEL6**

***For latest information and details contact your Linux distributor. IBM is working with its Linux distribution partners that selected exploitation for System z10 EC functionality will be provided in future Linux on System z distribution releases or versions**


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**

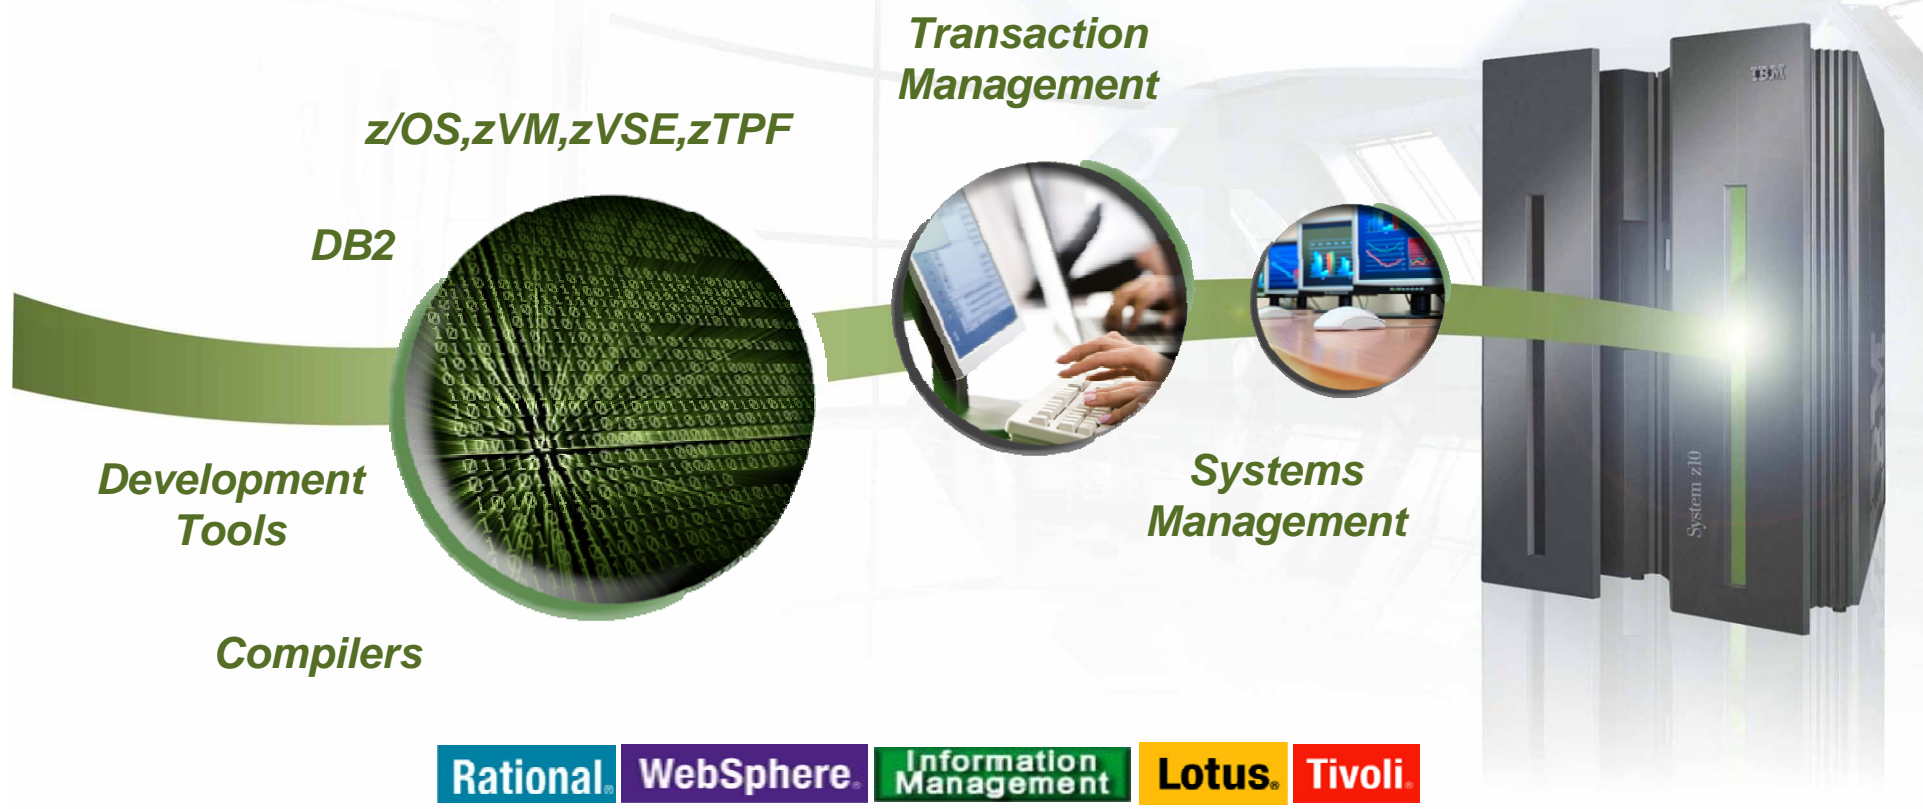
z10 EC – Under the covers (Model E56 or E64)



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**

IBM z10 EC Continues the CMOS Mainframe Heritage




- G4 - 1st full-custom CMOS S/390®
- G5 - IEEE-standard BFP; branch target prediction
- G6 - Cu BEOL

- IBM eServer zSeries 900 (z900) - Full 64-bit z/Architecture®
- IBM e Server zSeries 990 (z990) - Superscalar CISC pipeline
- z9 EC - System level scaling

- z10 EC - Architectural extensions

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**

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 SEEPRAM (S) chips

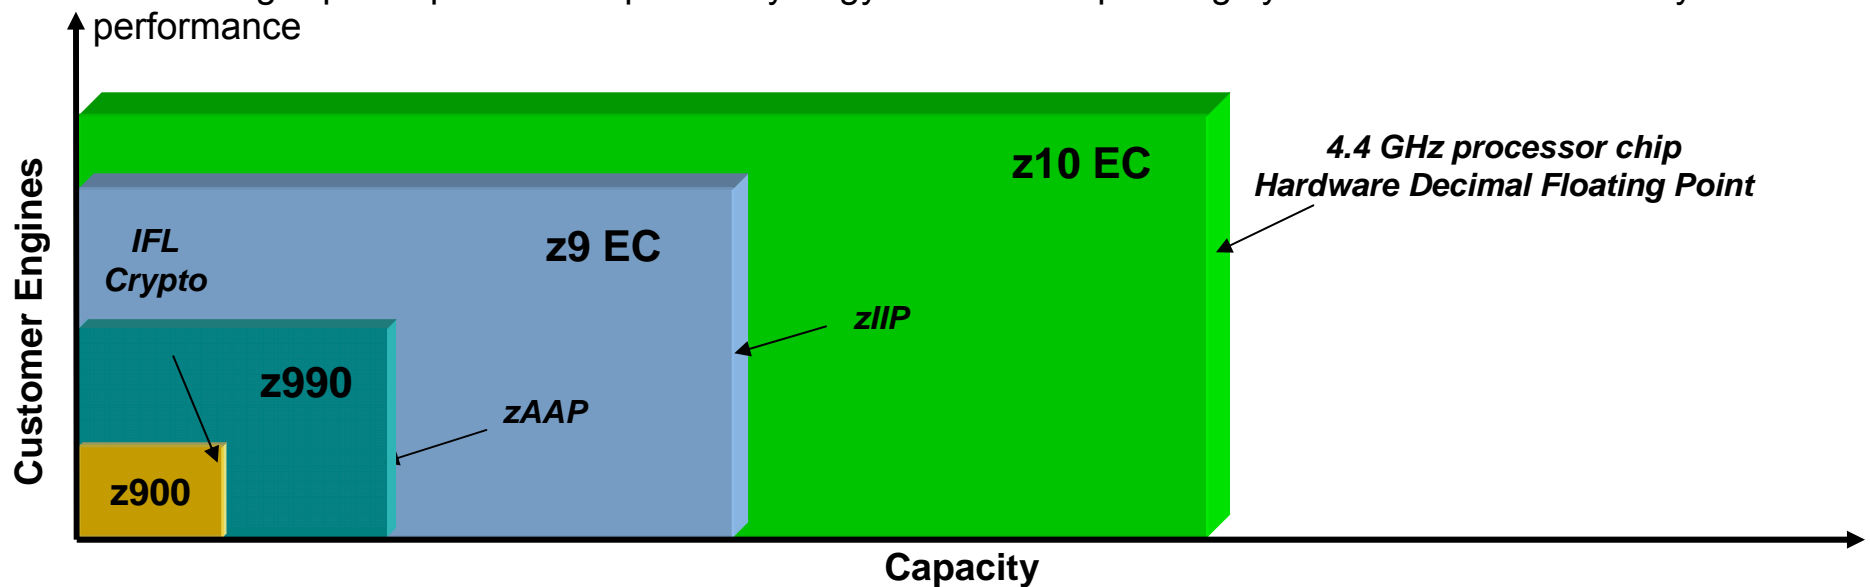
- Product data for MCM, chips and other engineering information

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



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

Evolution of System z Specialty Engines

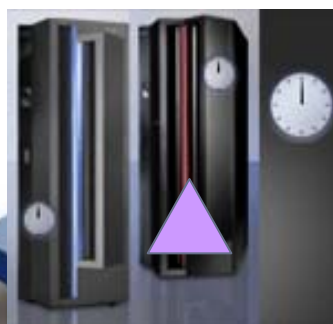


Cell Processor*

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



Integrated Facility for Linux (IFL) 2000



Internal Coupling Facility (ICF) 1997



System z Application Assist Processor (zAAP) 2004

Eligible for zAAP:

- Java execution environment
- z/OS XML



IBM System 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



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
Faster Uni Processor
Up to 64 customer PUs
36 CP Subcapacity Settings
Star Book Interconnect
Up to 1.5 TB memory
Fixed HSA as standard
Large Page Support (1 MB)
HiperDispatch
Enhanced CPACF SHA 512, AES 192 and 256-bit keys
Hardware Decimal Floating Point
New Capacity on Demand architecture and enhancements



6.0 GBps InfiniBand HCA to I/O interconnect
FCP Performance Improvement
SCSI IPL included in Base LIC
OSA-Express3 10 GbE (2Q08)*
HiperSockets Multi Write Facility enhancements
InfiniBand Coupling Links (2Q08)*
STP using InfiniBand (2Q08)*
Capacity Provisioning Support
Scheduled Outage Reduction
Improved RAS
FICON LX Fiber Quick Connect
Power Monitoring support

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



InfiniBand I/O Bus
z10 EC
2008



STI
z9 EC
2005



STI
z990
2003



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

Processor / Memory

- Uniprocessor Performance
- n-way Performance
- System Capacity
- Processor Design
- Models
- Processing Units (PUs)
- Granular Capacity
- Memory
- Fixed HSA
- 62% performance improvement over z9 EC uniprocessor *
- On average 50% more performance than z9 EC in a n-way configuration
- 70% system capacity performance improvement over z9 EC 54-way **
- New 4.4GHz processor chip
- z10 EC has 5 and z9 EC has 5 models, both with up to 4 books
- z10 EC has up to 64 PUs to configure, up to 54 on z9 EC
- z10 EC has up to 100 Capacity settings versus 78 on the z9 EC
- z10 EC has up to 1.5 TB vs. up to 512 GB on z9 EC
- z10 EC has fixed 16 GB HSA, z9 EC had HSA carved from purchased memory

Virtualization

- LPARs
- HiperDispatch
- z10 EC has up to 64 logical processors in an LPAR versus 54 on z9 EC
- z10 EC has HiperDispatch for improved synergy with z/OS Operating System to deliver scalability and performance

Connectivity

- HiperSockets
- FICON for SANs
- Total channels
- Internal I/O Bandwidth
- Enhanced I/O structure
- Coupling
- Cryptography
- LAN Connectivity
- z10 EC New HiperSockets Layer 2 and Multiple Write Facility
- Up to 336 FICON channels on z10 EC and z9 EC
- Same - Up to 1024 channels
- 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)
- Star L2 Cache Book Interconnect versus Ring Topology interconnect on z9 EC
- Coupling with InfiniBand ¹ – improved distance and potential cost savings
- Improved AES 192 and 256 and stronger hash algorithm with Secure Hash Algorithm (SHA-512)
- New OSA-Express³¹ for 10 Gigabit Ethernet connectivity

On Demand / RAS

- Capacity Provisioning Mgr
- RAS Focus
- Just in Time deployment of Capacity
- z10 EC & z/OS (1.9) for policy based advice and automation
- z10 EC can help eliminate preplanning required to avoid scheduled outages
- Capacity on Demand offerings CBU and On/Off CoD plus new Capacity for Planned Events are resident on z10 EC

Environmentals

- Monitoring
- z10 EC displays energy efficiency on SAD screens
- Utilizes IBM Systems Director Active Energy Manager for Linux on System z for trend calculations and management of other servers that participate

* 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

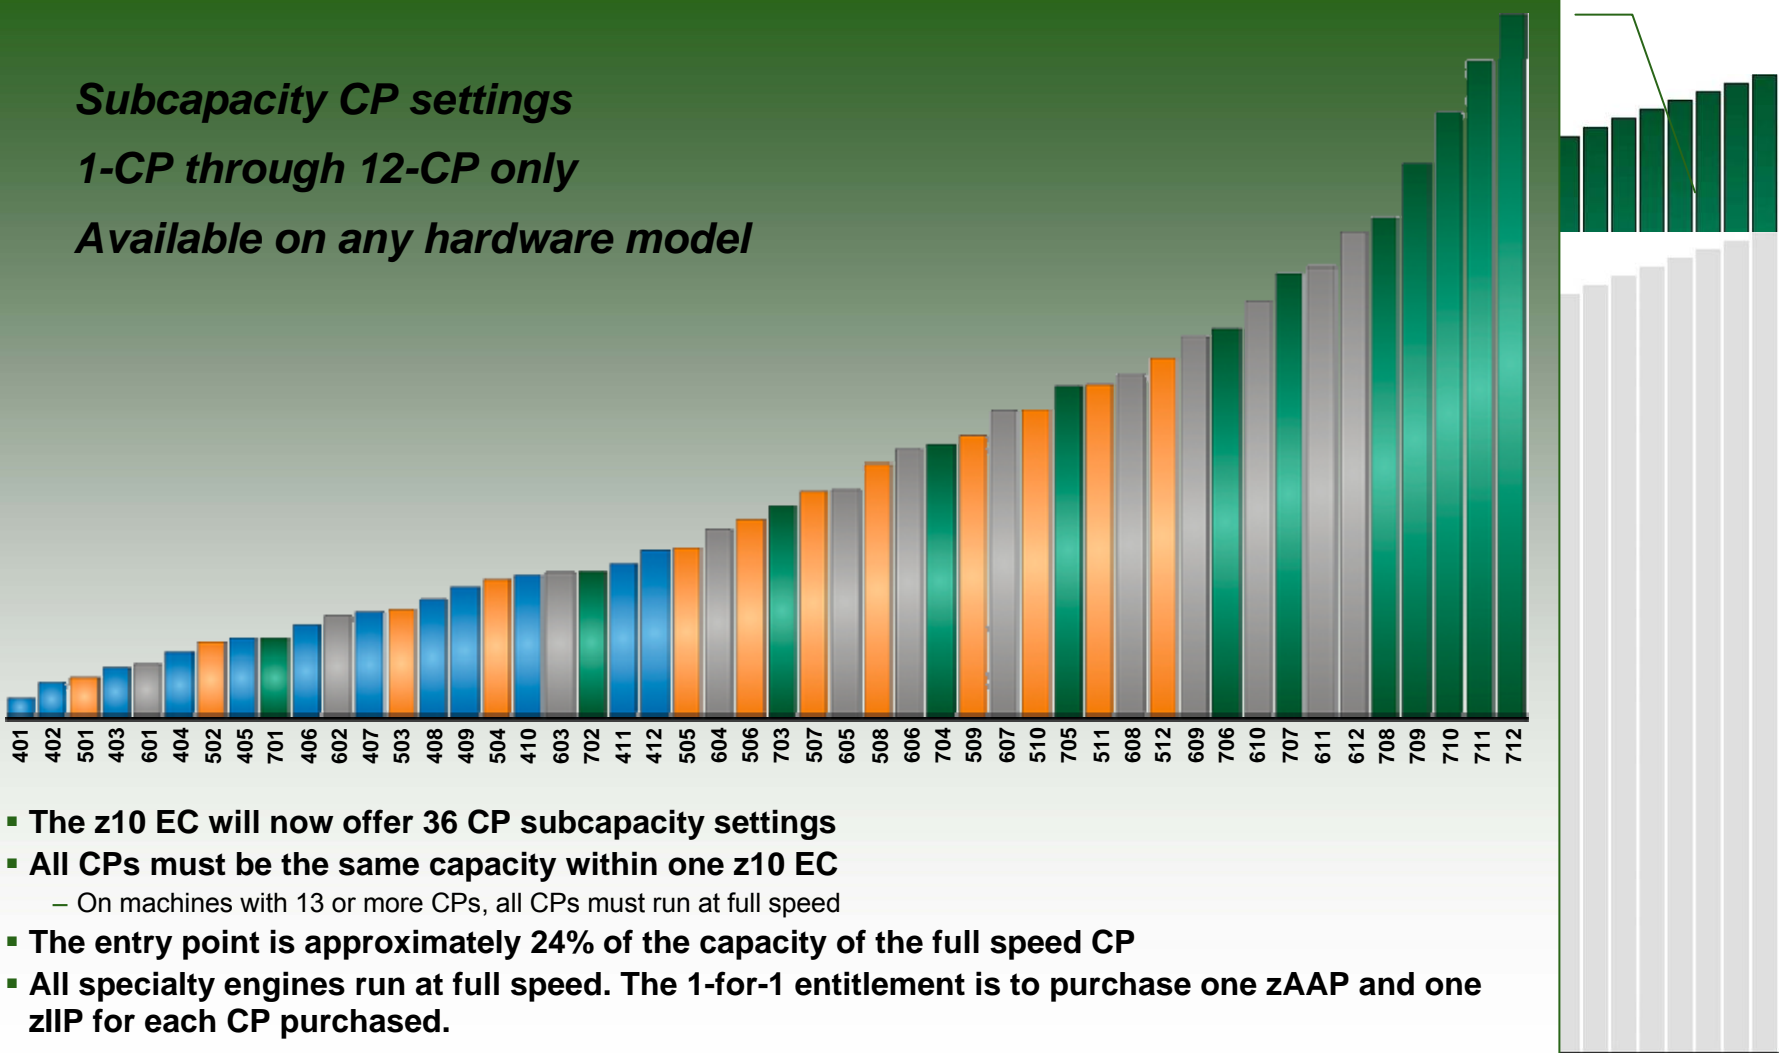
¹ Planned availability 2Q08

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

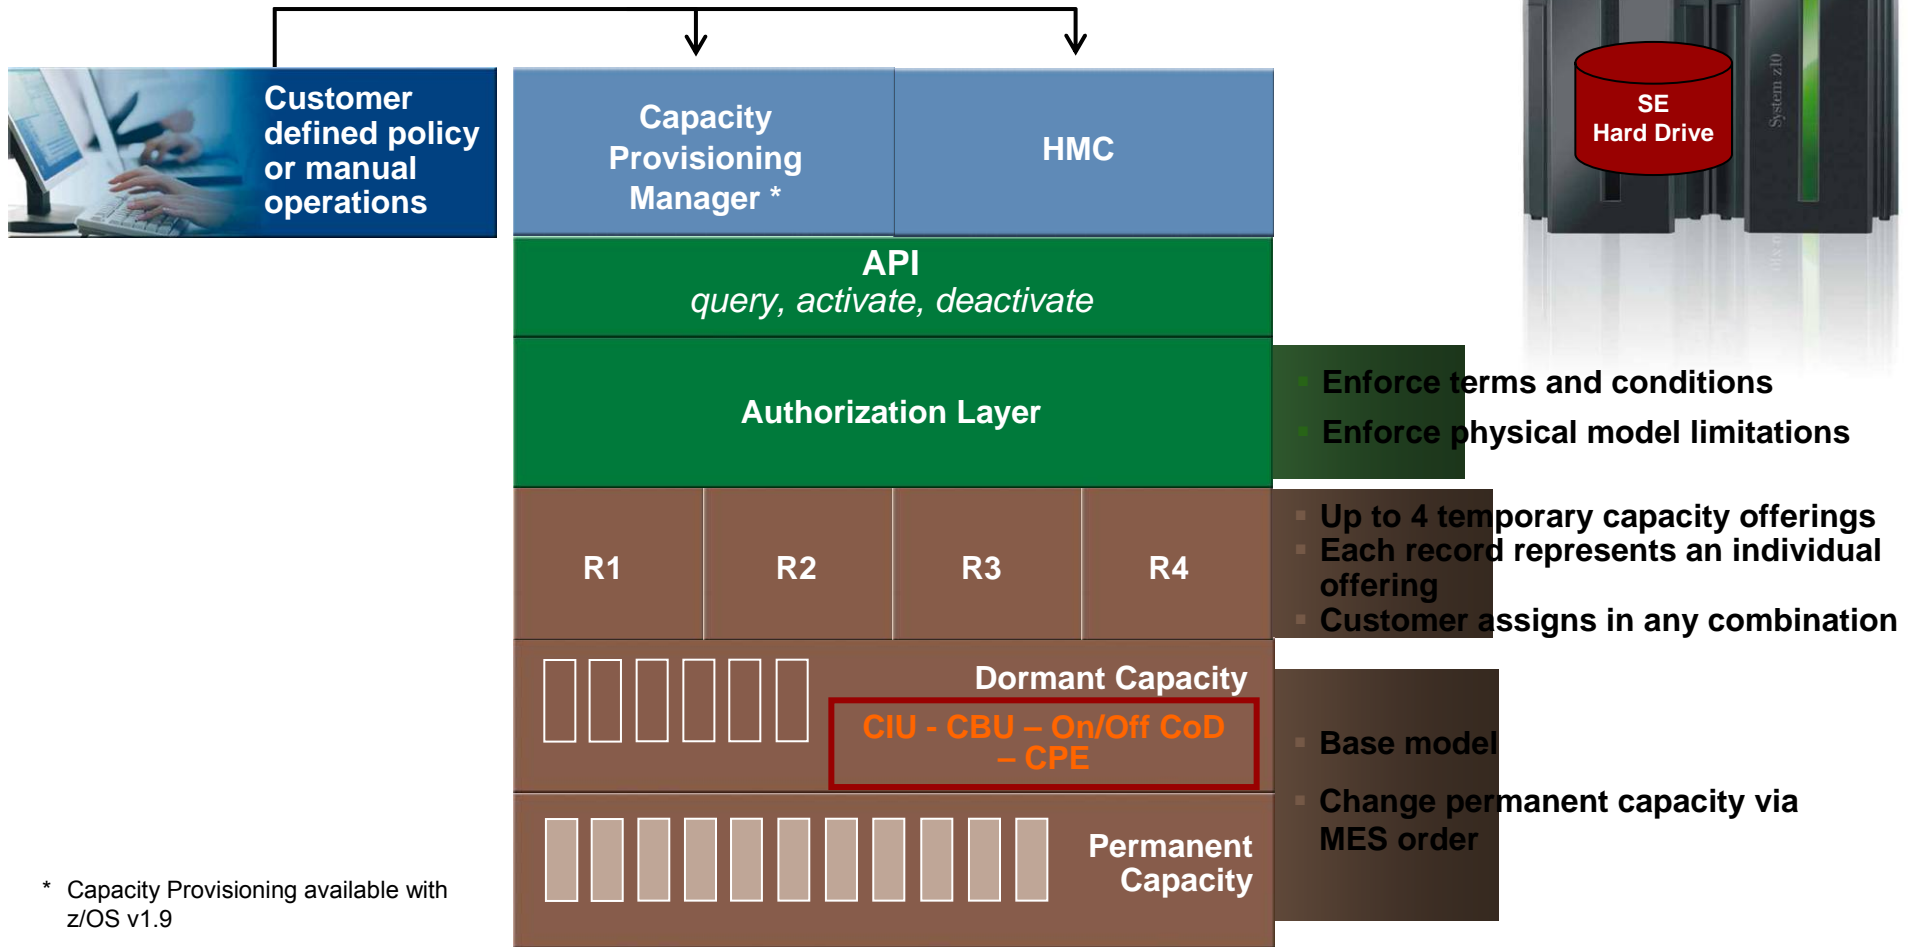


CoD – New architectural approach for z10 EC

- 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

On demand simplified New architectural approach for capacity

↓
Orders downloaded from System Support electronically or by IBM Service



* Capacity Provisioning available with z/OS v1.9

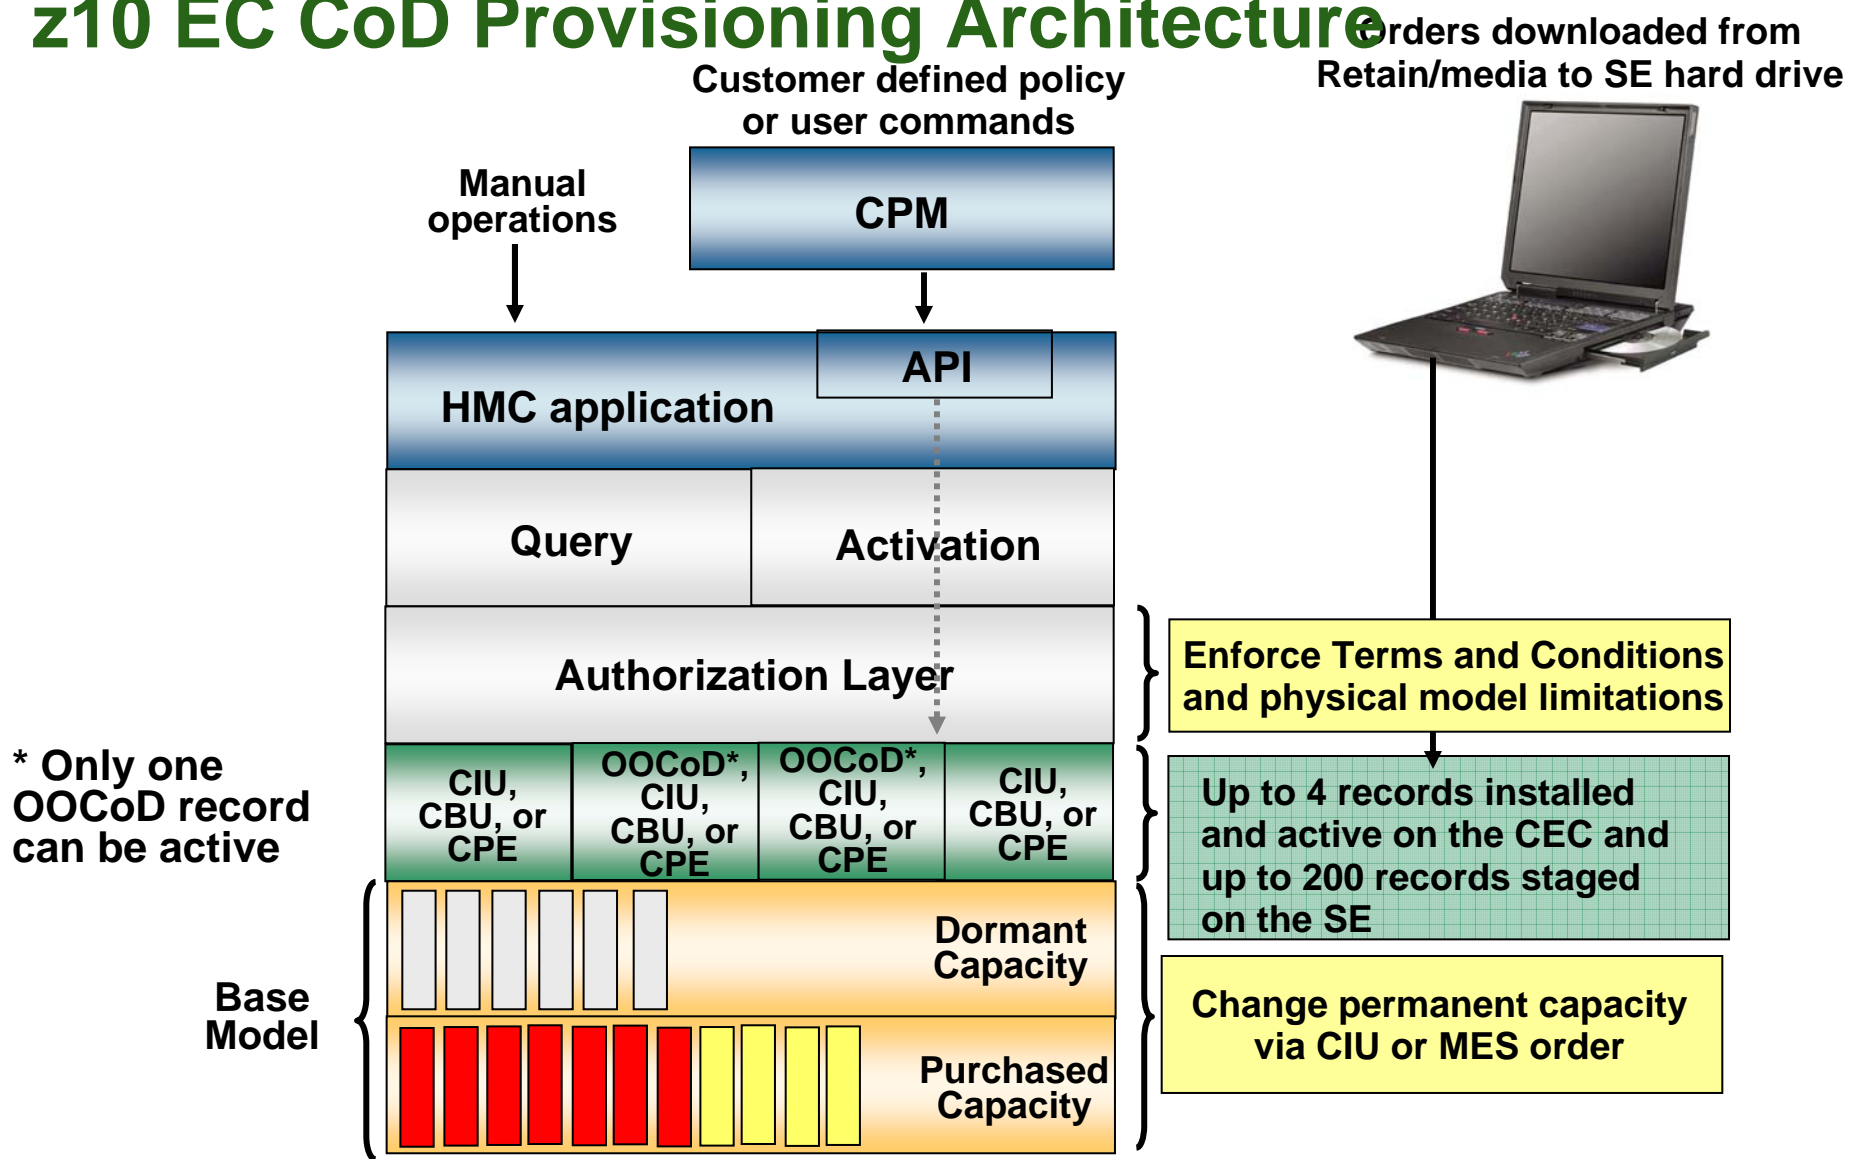
z10 EC CoD Offerings - Details

- **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**

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



CBU Comparison – z9 EC versus z10 EC

	z9	z10
Granularity	All on / All off	Granular
Customer exceeds terms	Reduce machine capacity	Removed automatically, if possible
End of term	CBU record does not expire	CBU record expires
Number of CBU orders	Buy one, apply one	Buy many, apply many
Terms	Usually 5 years	Variable, 1-5 years


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

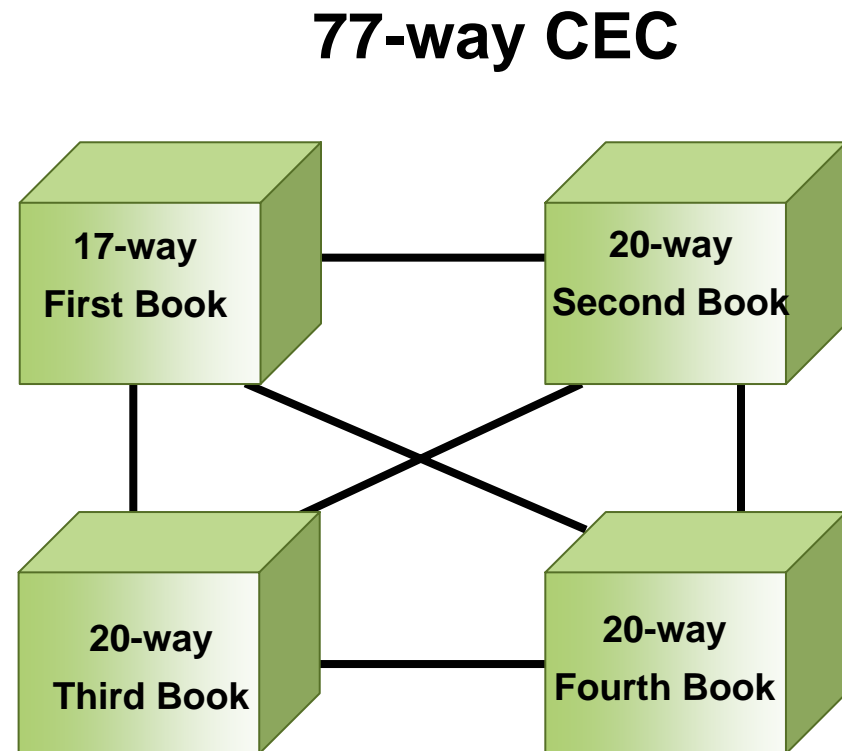
- Continued Focus on Firmware Quality
- Reduced Chip Count on MCM
- Memory Subsystem Improvements

- 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 – 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.

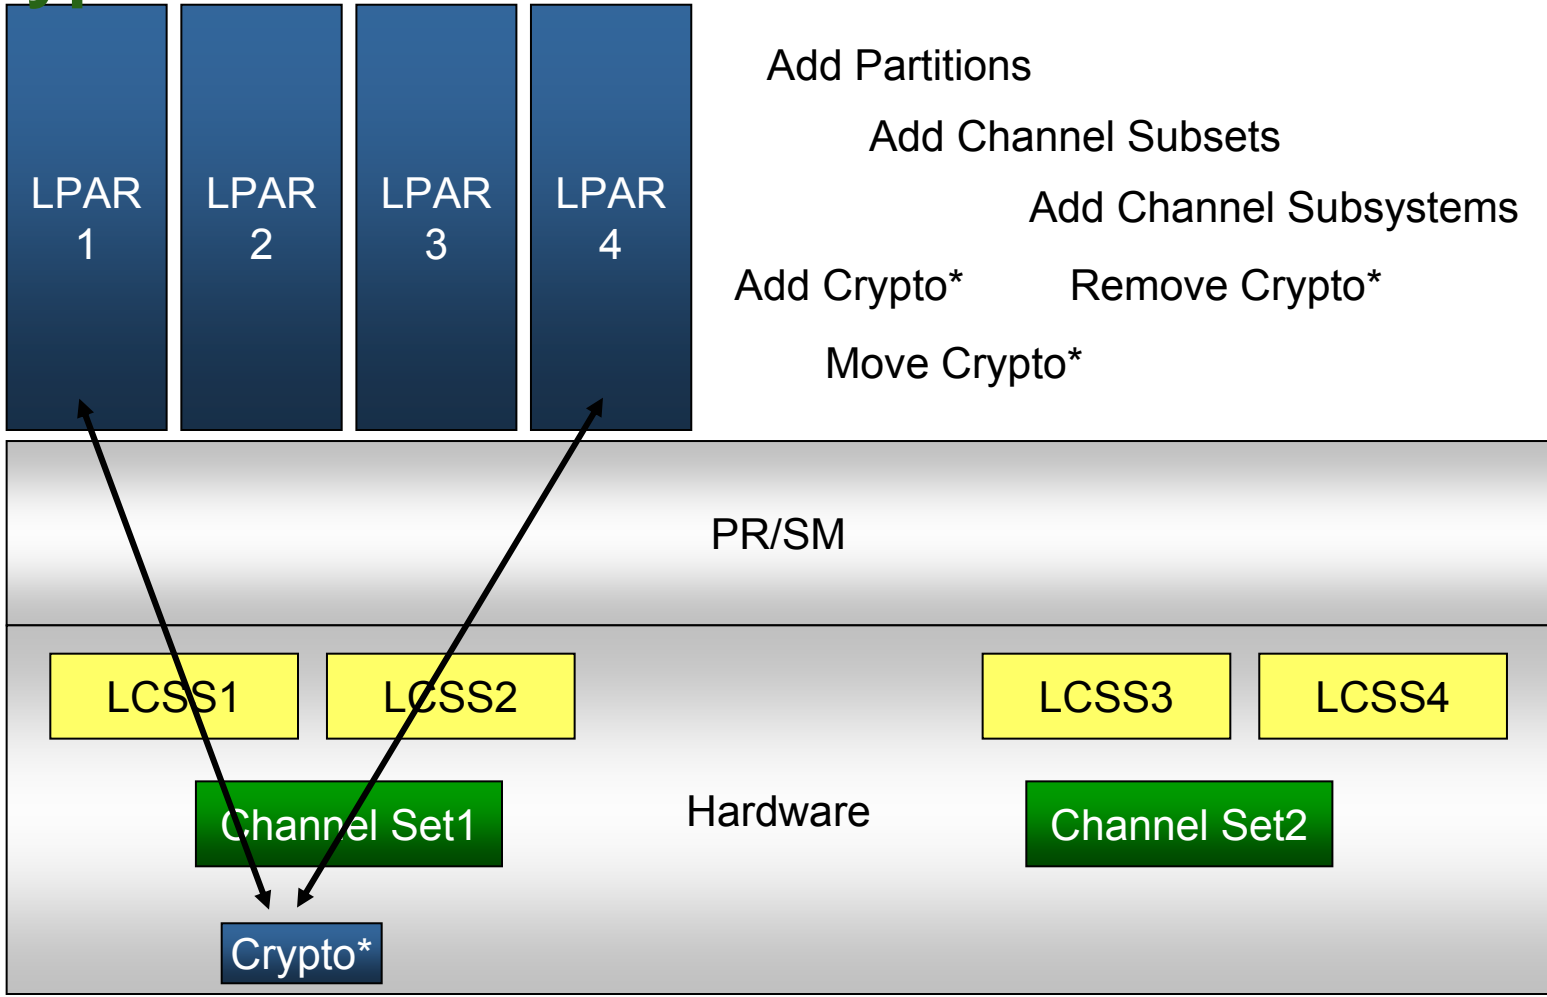
z10 EC Service Enhancements designed to avoid Scheduled Outages

- **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

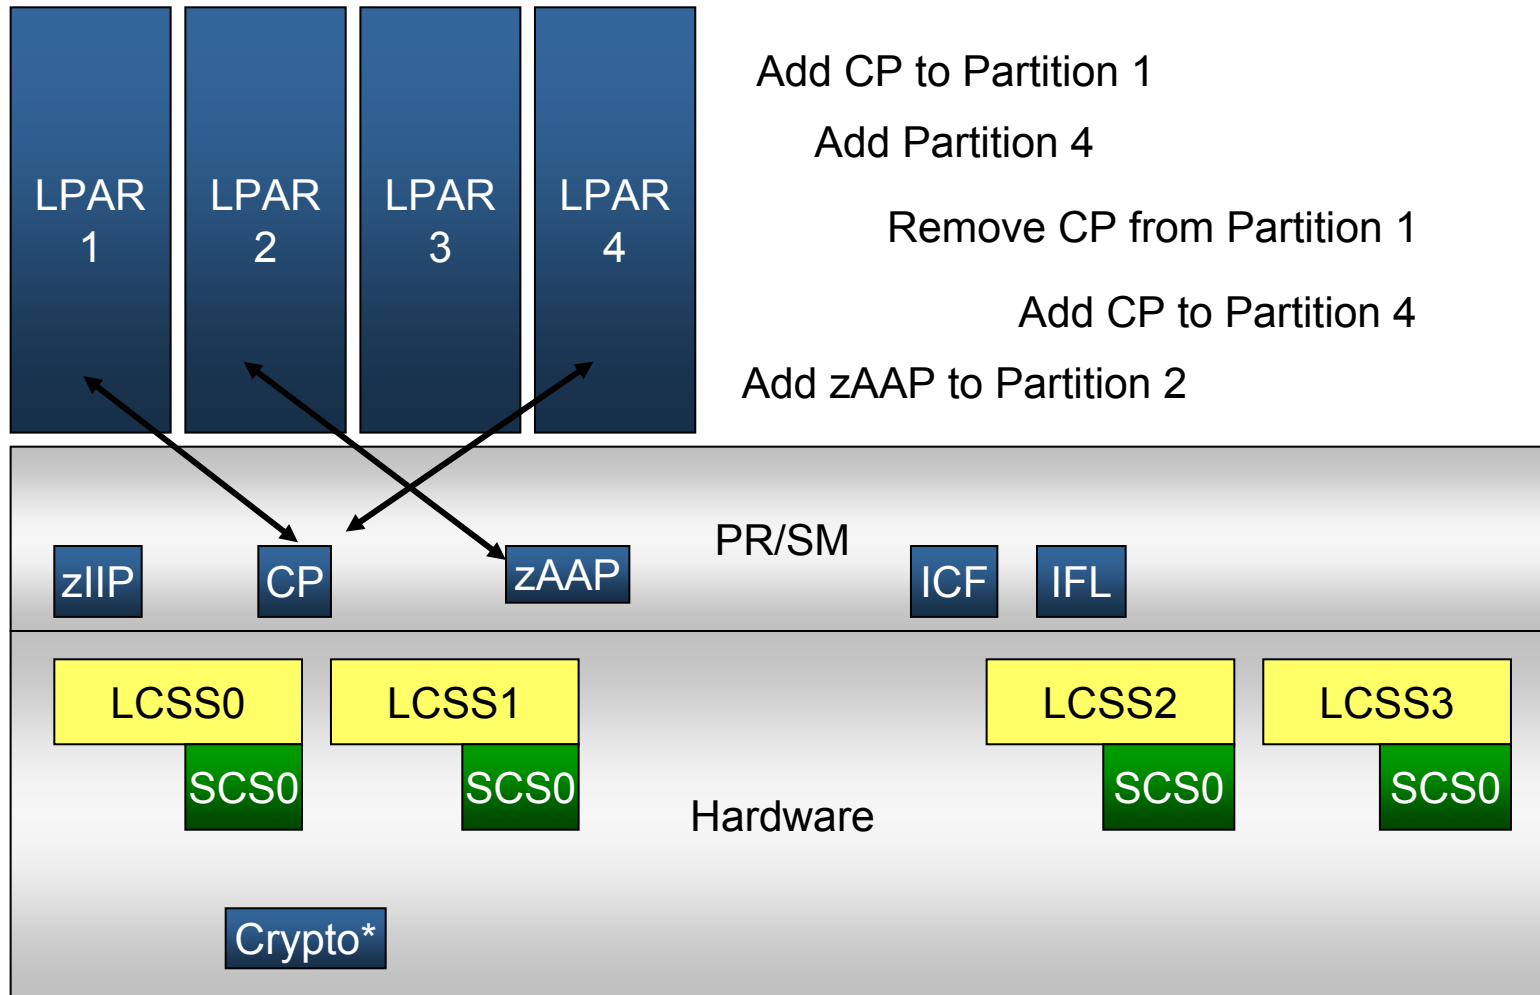


- Add Partitions
- Add Channel Subsets
- Add Channel Subsystems
- Add Crypto*
- Remove Crypto*
- Move Crypto*

***Crypto Express2 Coprocessor or Accelerator**

z10 EC Outage Elimination examples for LPARs

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



- Add CP to Partition 1
- Add Partition 4
- Remove CP from Partition 1
- Add CP to Partition 4
- Add zAAP to Partition 2

***Crypto Express2 Coprocessor or Accelerator**

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

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**

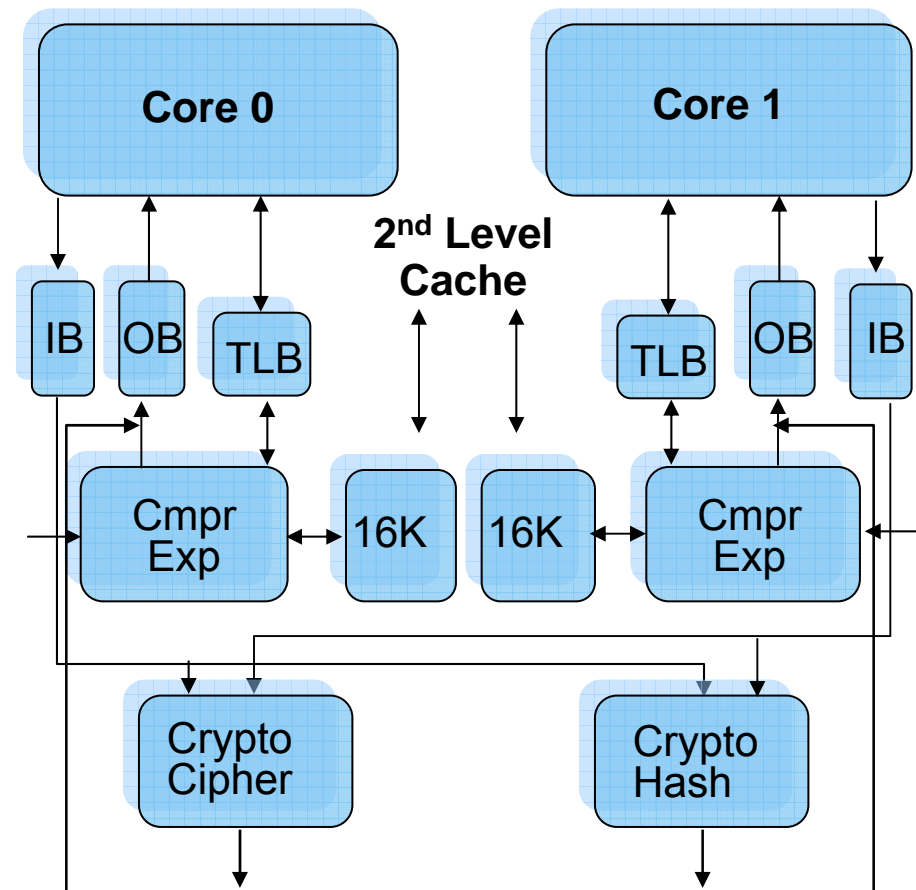


Support for new Security Standards

- **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

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



**IBM System z10 Integrated Information Processor and
IBM System z9 Integrated Information Processor**

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



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 ¹
- InfiniBand Coupling Links ¹

¹ 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

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

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

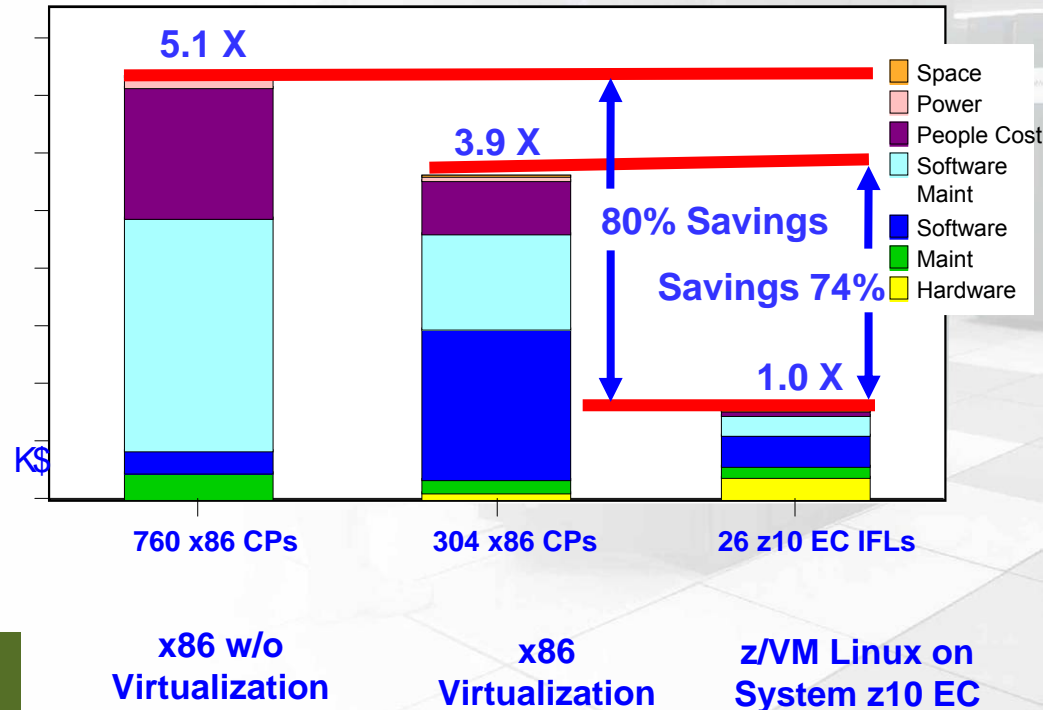
Your IT Cost may vary:

- **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

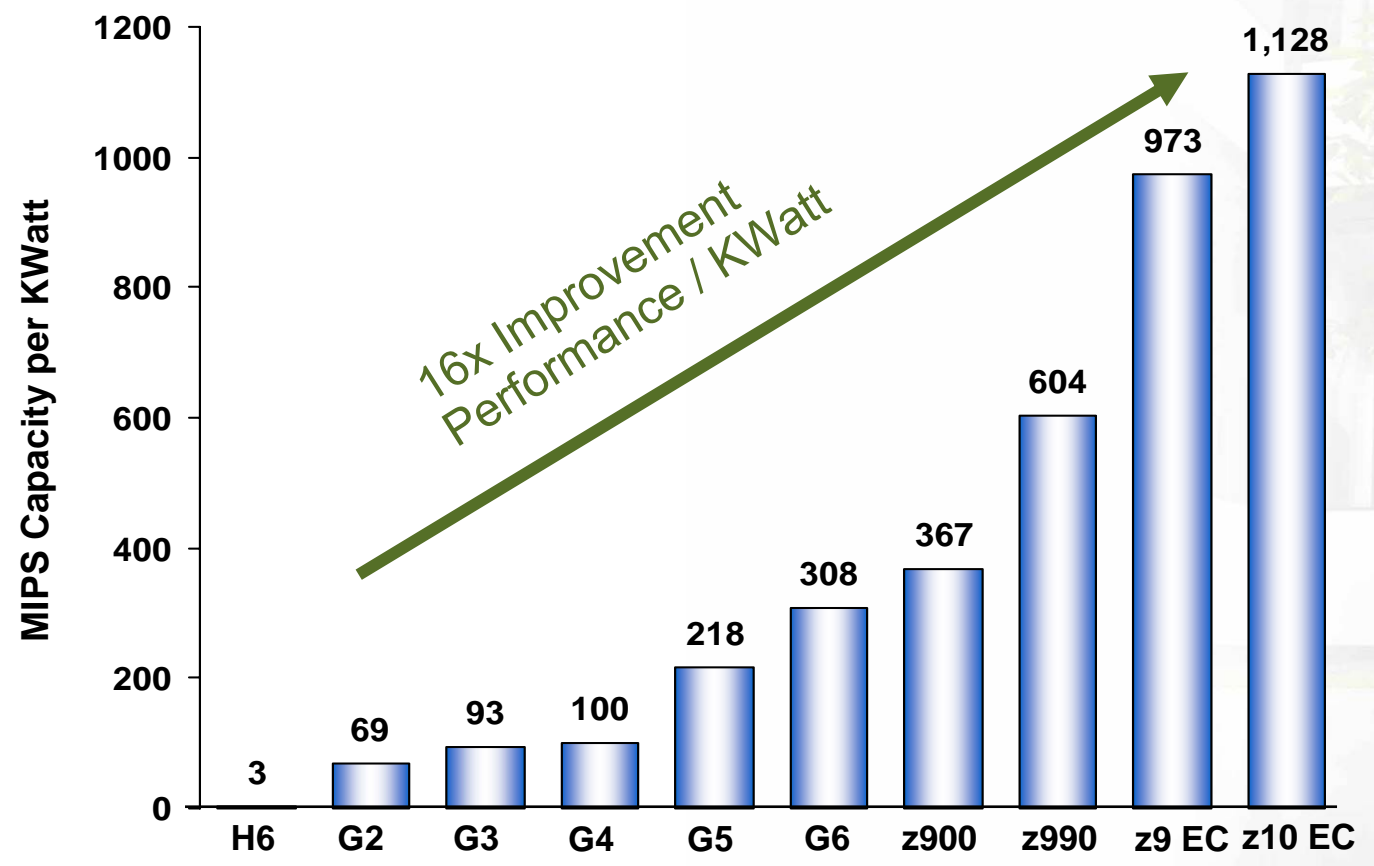
1. Green Design
2. Faster Engines
3. New Advanced Function
4. Flexible Capacity
5. Improved RAS
6. Support for New Security Standards
7. Expanding the ability to run more diverse workloads
8. Upgradeability
9. Virtualization
10. Energy Efficiency

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



System z10 EC is Cost Effective and can help you Go Green by delivering *highly energy efficient technology*



**1/24 the
Energy
Use in 1/5
the Floor
Space***

* Comparison is versus x86 Blade servers without virtualization, reflecting a current-day consolidation.

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

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 3rd party
- Created through the implementation of energy conservation projects on the demand side



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.**

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

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



Stages of adoption


Simplified



Shared



Dynamic



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 and ...



Thank
You