



IBM System z10 Enterprise Class
The Evolution of Mainframe Technology



The Future Runs on System z

Fulvio Capogrosso
Distinguished Engineer
STG IBM Italy

Introducing the IBM System z10™ Enterprise Class

A Marriage of Evolution and Innovation

Evolution

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

Innovation

- 4.4 GHz quad-core chip
- New processor pipeline design
- HyperDispatch
- ‘Just in time’ deployment of capacity resources
- Vision to expand System z capabilities with Cell Broadband Engine™ technology



IBM z10 EC Chip Highlights

- **New high-frequency z/Architecture microprocessor core**

- 4.4 GHz operation
- IBM 65nm SOI technology
- Increased overall pipeline depth
- 50+ instructions added
- Support for 1MB page frames

- **4 cores per die**

- Each with 3MB private 2nd-level cache

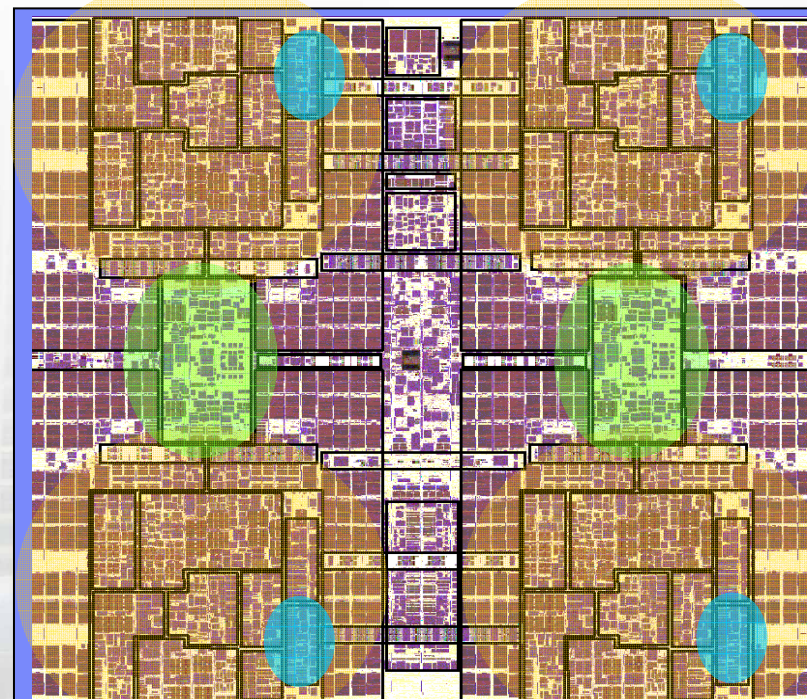
- **On chip accelerator engines**

- Data compression
- Cryptographic functions
- Decimal floating point

- **Cores connected to SMP hub chip for integrated SMP communications**

- **System interfaces**

- 2 x 48 GB/s SMP Hub
- 4 x 13 GB/s Memory
- 2 x 17 GB/s I/O




- 991M Transistors
- 138 Mb SRAM
- 6 km wire
- 21.7 X 20.0 mm die
- 1188 signal / 8765 total chip I/Os


z10 EC Chip Relationship to POWER6™ Chip

- **Siblings, not identical twins**
- **Share lots of DNA**
 - IBM 65nm Silicon-On-Insulator (SOI) technology
 - Design building blocks:
 - Latches, SRAMs, regfiles, dataflow elements
 - Large portions of Fixed Point Unit (FXU), Binary Floating-point Unit. (BFU), Hardware Decimal Floating-point Unit (HDFU), Memory Controller (MC), I/O Bus Controller (GX)
 - Core pipeline design style
 - High-frequency, low-latency, mostly-in-order
 - Many System z and System p designers and engineers working together
- **Different personalities**
 - Very different Instruction Set Architectures (ISAs)
 - Very different cores, cache hierarchy and coherency model
 - SMP topology and protocol
 - Chip organization
 - IBM z10 EC Chip optimized for Enterprise Data Serving Hub

Enterprise Quad Core z10
processor chip




Dual Core POWER6
processor chip



IBM z10 EC SMP Hub Chip



- **Connects multiple IBM z10 processor chips**
 - 48 GB/Sec bandwidth per processor
- **Shared 3rd-Level cache**
 - 24MB SRAM Cache per chip
- **Low-latency SMP coherence fabric**
 - Robust SMP scaling
 - Strongly-ordered architecture
- **Multiple hub chip pairs allow further SMP scaling**




- 1.6 Billion Transistors
- 242 Mb SRAM
- 3 km wire
- 20.8 X 21.4 mm die
- 2419 signal / 7984 total chip I/Os

z10 System Packaging

Chips




MCM



Book



System




	<i>z9</i>	<i>z10</i>
System Config	4 Books, 64 PUs	4 Books, 77 PUs
Max # of CPU.s	54	64
X-Book Topology	Dual Ring	Star Connected
I/O Connections	16 x 2.7 x 4 G/sec	12 x 6 x 4 GB/sec
Max I/O rate	172.8 GB/sec	288 GB/sec
Max Memory	512 GB	1.5 TB + HSA
Cache levels	2	3

Improved Server Performance and Scalability

More faster engines

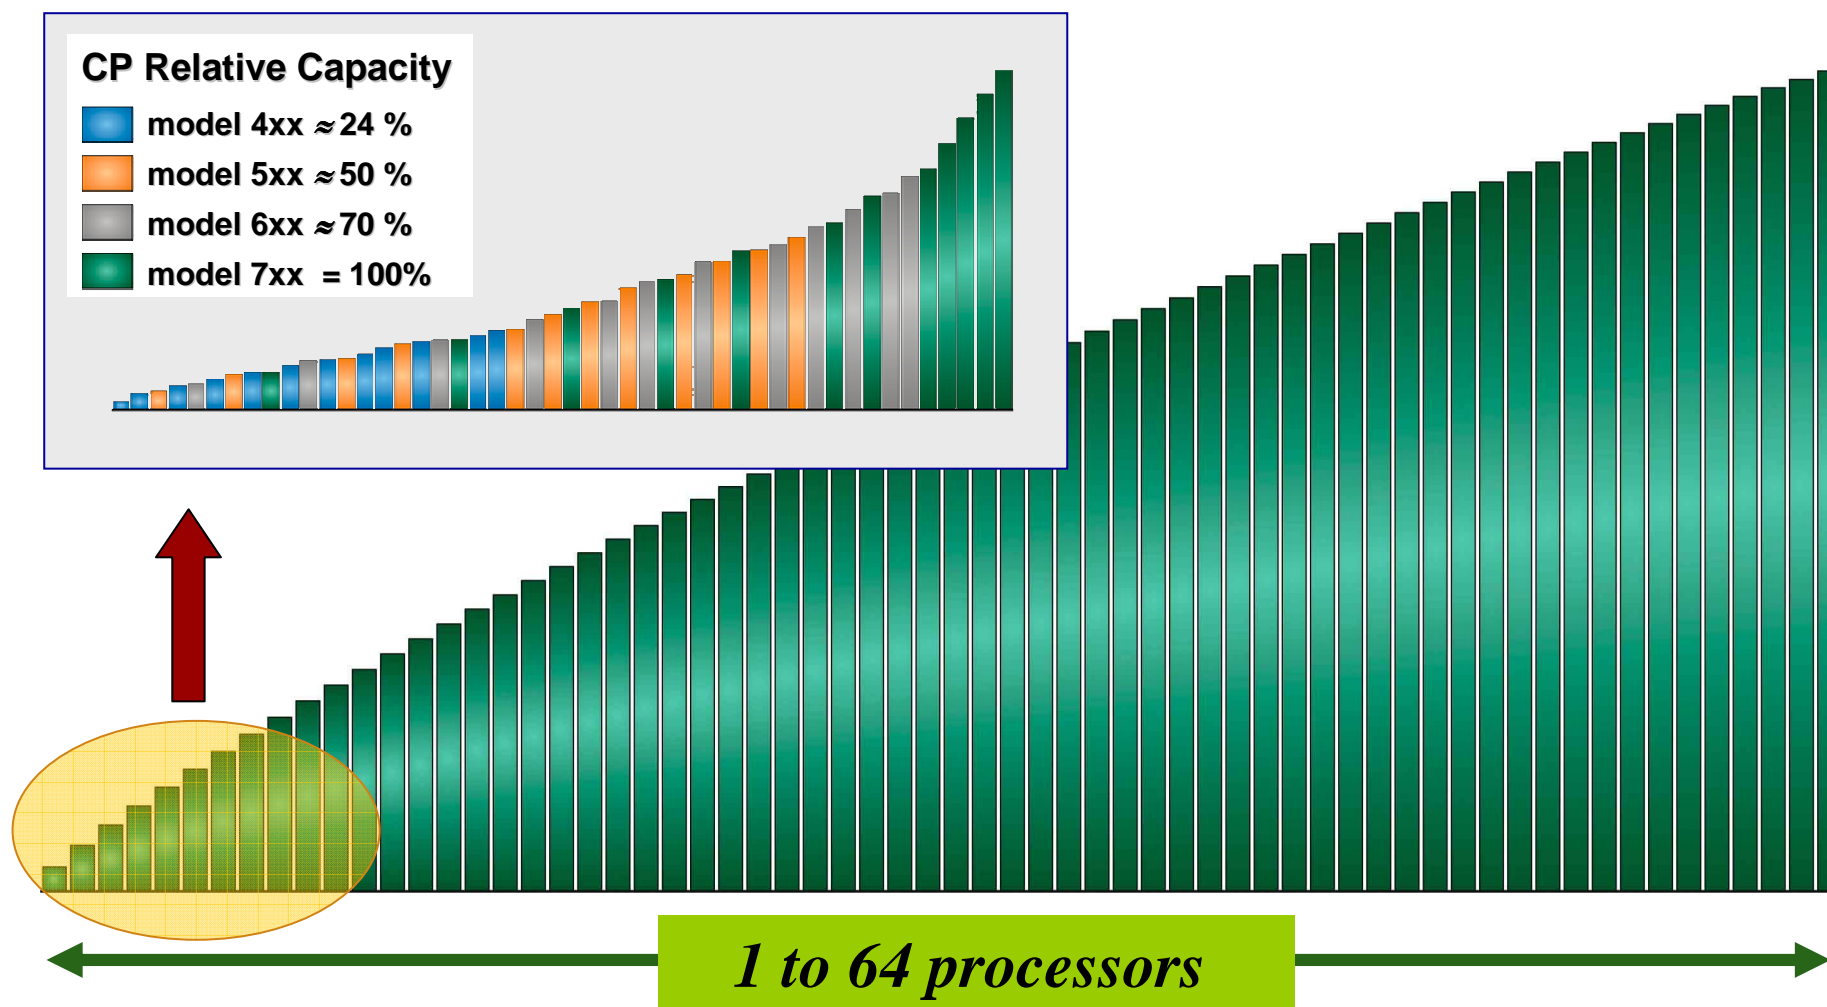
- The z10 EC can deliver, on average up to **50% more performance** in a n-way configuration than a IBM System z9™ Enterprise Class (z9 EC) n-way
- The uniprocessor can deliver **up to 62% more performance** than z9 EC uniprocessor *
- The z10 EC 64-way can deliver **up to 70% more server capacity**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



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

** Comparison of the z10 EC 64-way and the z9 EC S54 based on LSPR mixed workload average running z/OS 1.8

Total 100 capacity settings to meet your needs



Helping to Get You Connected to Your World

- **Improved performance and flexibility for connectivity**
 - Up to 12 I/O interconnections (6GB/sec each) per book (48 per System)
- **Broad set of options to meet your needs**
- **Excellent investment protection when you upgrade to the z10 EC**

Within the server


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

To the Data


- FICON/FCP
 - FICON® Express4 (4Gbps)
 - **FICON Express2**
 - **FICON Express**
- ESCON®



To the Network

- OSA-Express3¹ 
 - 10 Gigabit Ethernet
- OSA-Express2
 - 1000BASE-T Ethernet
 - Gigabit Ethernet LX and SX
 - 10 Gigabit Ethernet LR

For Clustering

- InfiniBand Coupling Links¹ 
- ICB-4
- ISC-3 (peer mode only)
- STP - NTP Client Support
- Support for n-2 and above servers

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

¹ Planned availability 2Q08

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**

3 Sub-capacity levels available up to 12 CPs

Total 100 capacity levels



Processor Units (PUs):

- One to four book modular design
- Up to 20 PUs per book
 - New core sparing technology
 - Configurable PUs to meet your needs with specialty engines (e.g. CPs, IFLs, ZIIPs, ZAAPs, SAPs)
 - More SAPs per system
- Specialty engines always run at full speed

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

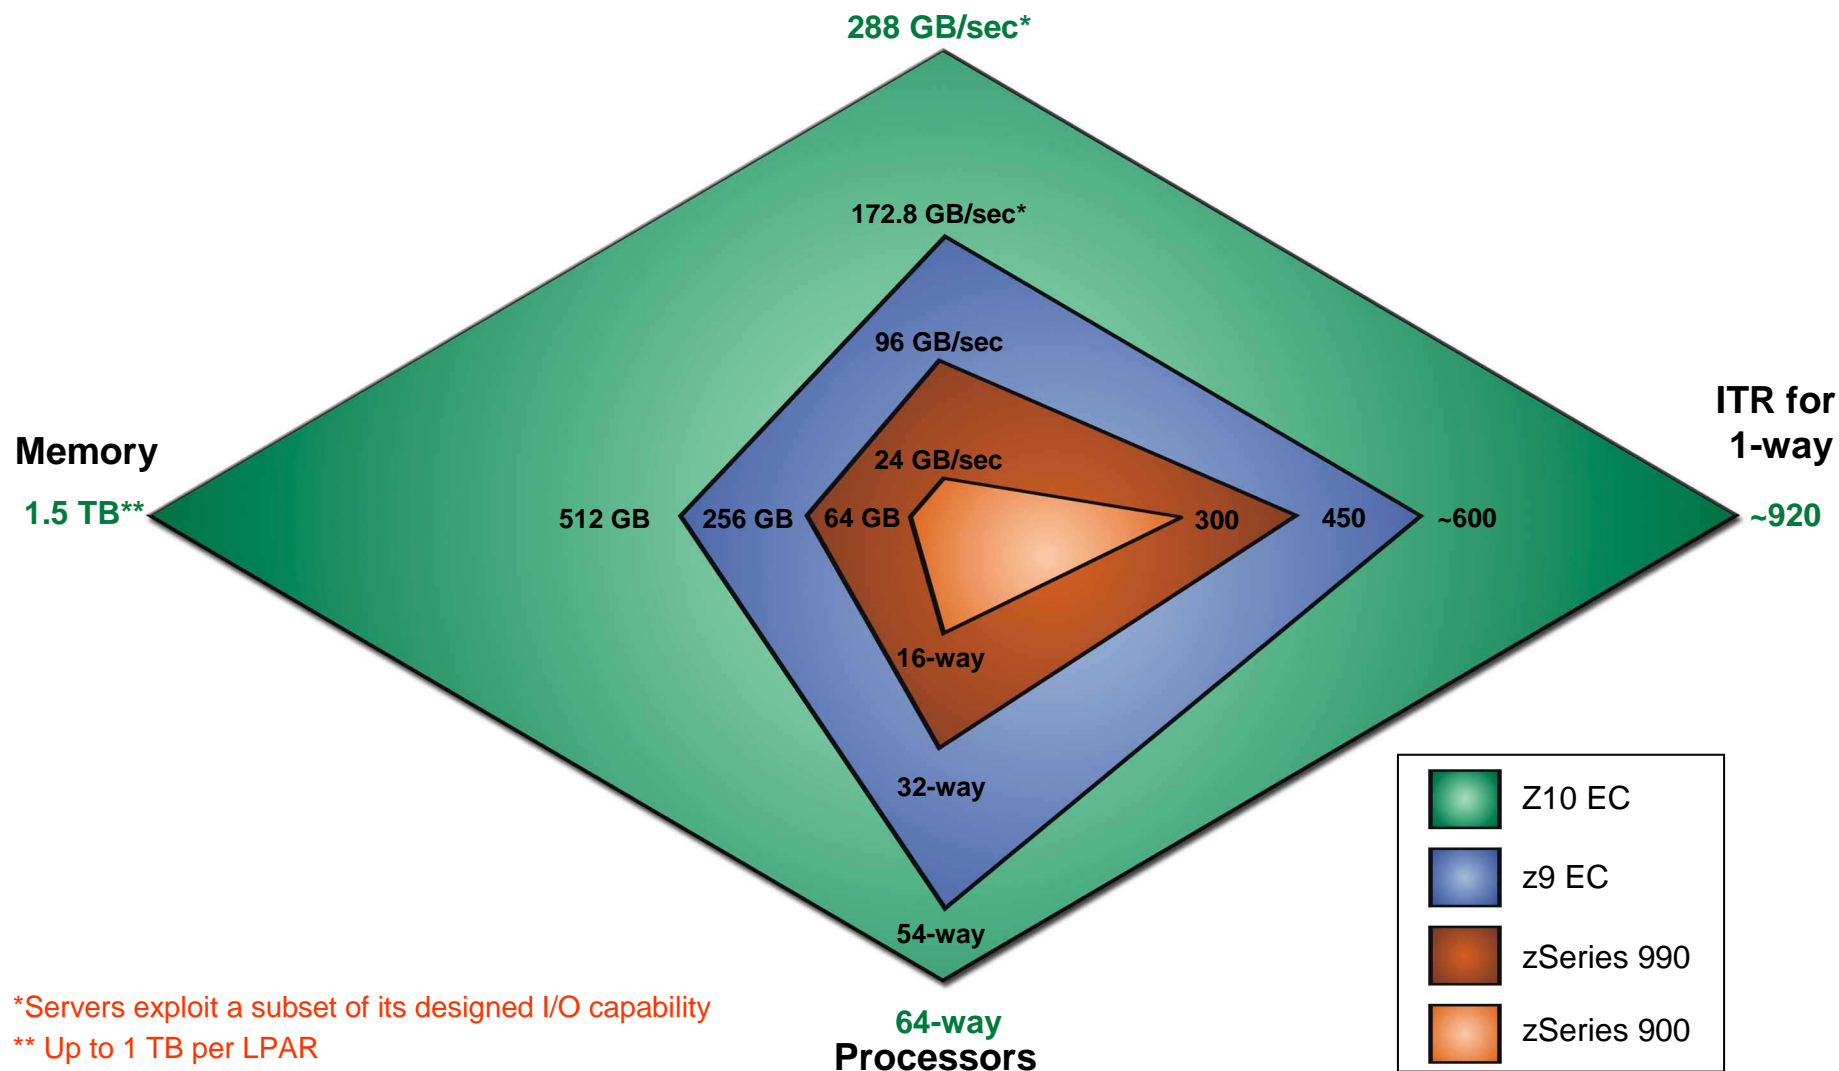
¹ Planned availability 2Q08

System z10 Upgrade Options



IBM System z: Balanced System Growth

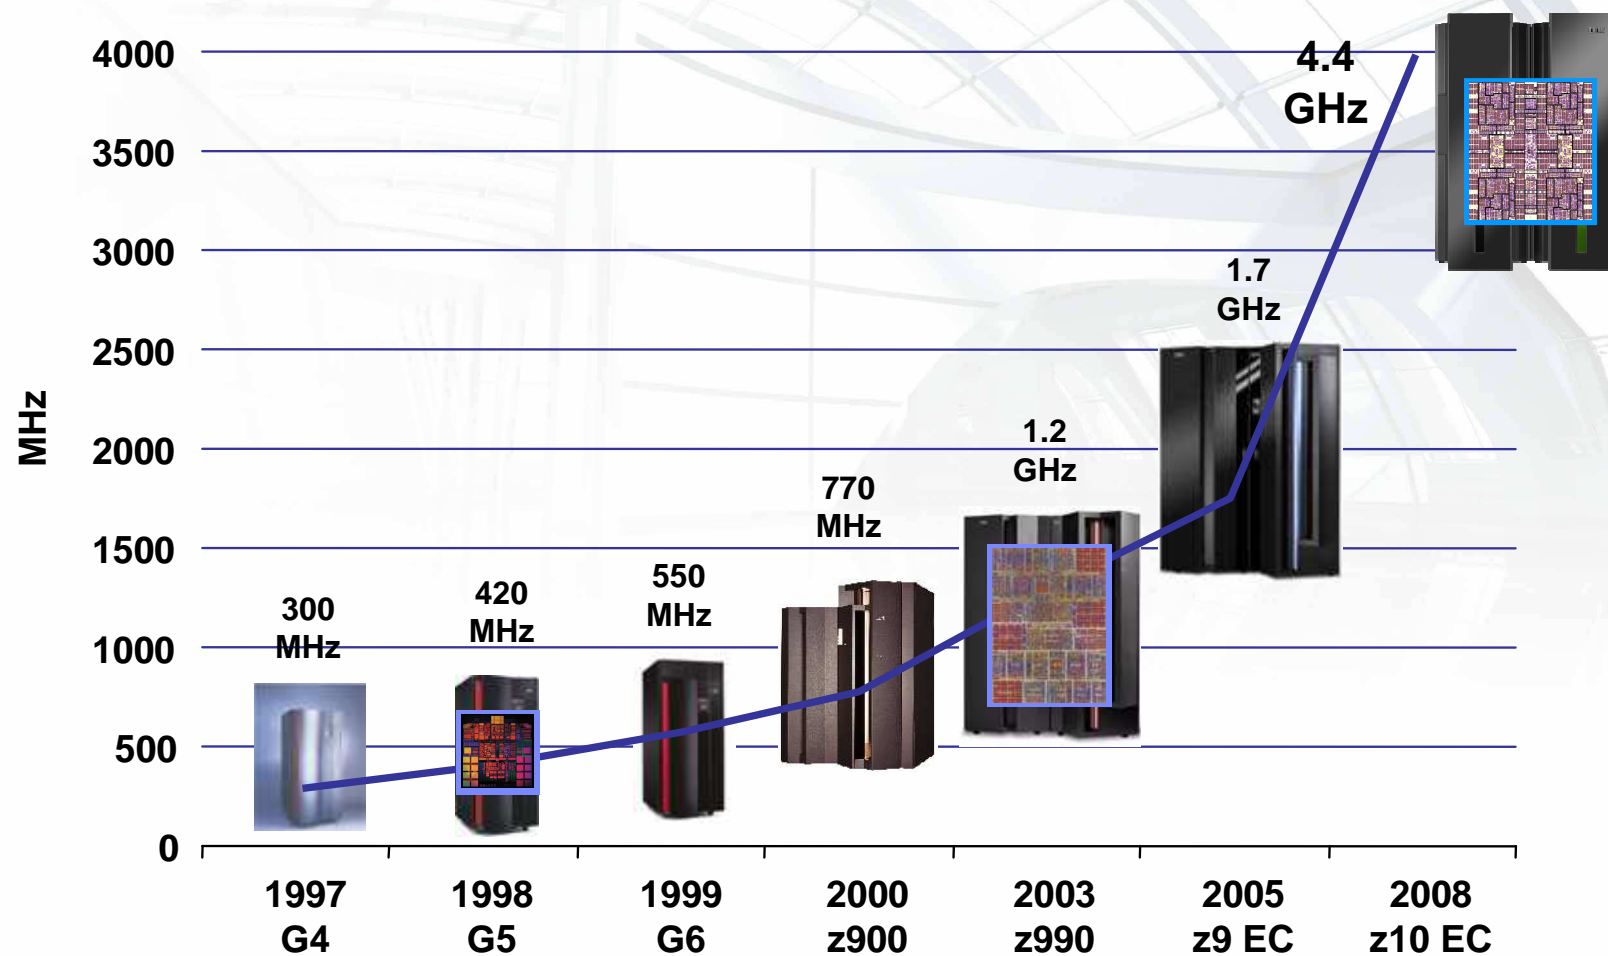
System I/O Bandwidth



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

** Up to 1 TB per LPAR

IBM z10 EC Continues the CMOS Mainframe Heritage



HiperDispatch: a z10 Unique Performance Function

- **In modular systems with NUMA (Non Uniform Memory Access) memory design CPUs have different distance-to-memory attributes**
 - Memory accesses can take a variable number of cycles depending upon cache structure / local or remote array accessed
- **Hardware cache optimization occurs when a given unit of work is consistently dispatched on the same physical CPU**
 - Requires zOS Dispatcher, PR/SM dispatcher, hardware support integration
- **On z10 EC the hardware, the firmware (PR/SM) and the operating system (z/OS) now tightly collaborate to obtain the hardware's full potential**
- **Two components**
 - Dispatcher Affinity (DA) - New z/OS Dispatcher
 - Vertical CPU Management (VCM) - New PR/SM Support

Helping to drive down the cost of IT

More workloads can benefit from zIIP and ZAAP

- **zIIP designed to help integrate data across the 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
- **zAAP designed to help implement new application technologies on System z**
 - z/OS XML System Services (introduced with z/OS V1.9 and rolled back to V1.8 and V1.7) helps make hosting XML data and transactions on System z more attractive. DB2 9 and Enterprise Cobol V4.1 are the first exploiters
- **zIIPs and zAAPs offer economic value**
 - price for z10 EC is same as for z9 EC and no charge MES upgrades available when moving to new technology



new!



new!



Available
9/07



*IBM System z Application Assist Processor and (zAAP)
IBM System z Integrated Information Processor (zIIP)*

System z and Cell Broadband Engine – The Vision

A ‘Marriage’ of Two Technologies that Perfectly Complement Each Other



← *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




***Bringing high performance computing benefits to
commercial workloads***



Keeping your system available is key to our design

Continuing our RAS focus to help 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

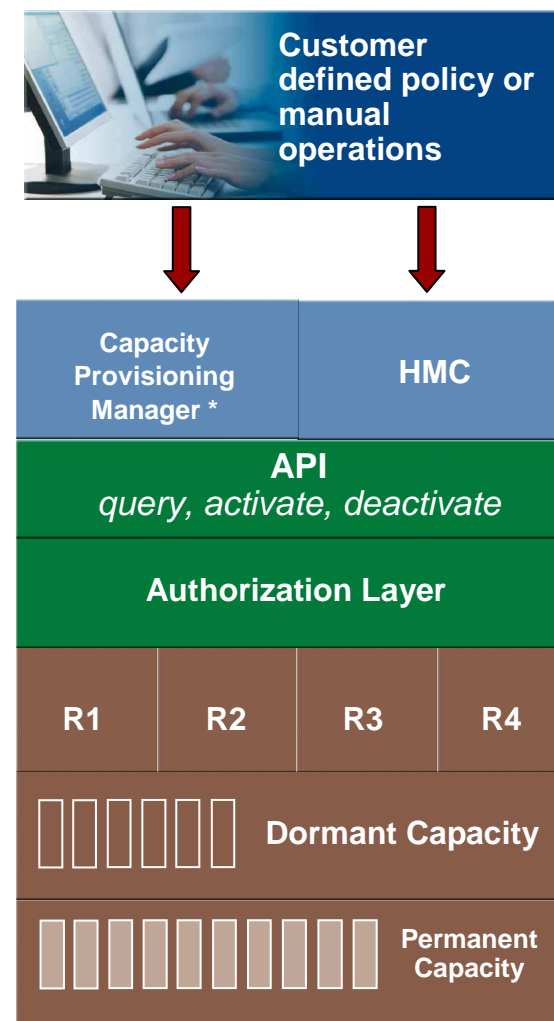
- DIMM FRU indicators
- Single Processor Core Checkstop
- Single Processor Core Sparing
- Point to Point SMP Fabric (not a ring)
- Rebalance PSIFB and I/O Fanouts
- Redundant 100Mb Ethernet service network w/ VLAN

- 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

New Provisioning Architecture


Just in Time Capacity Under Your Control

- **Permanent and temporary offerings**
 - Everything resident in the SE Hard Drive
- **No customer interaction with IBM at time of activation**
- **Multiple offerings can be in use simultaneously**
 - Each offering independently managed and priced
- **Flexible offerings may be used to solve multiple situations**
 - 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



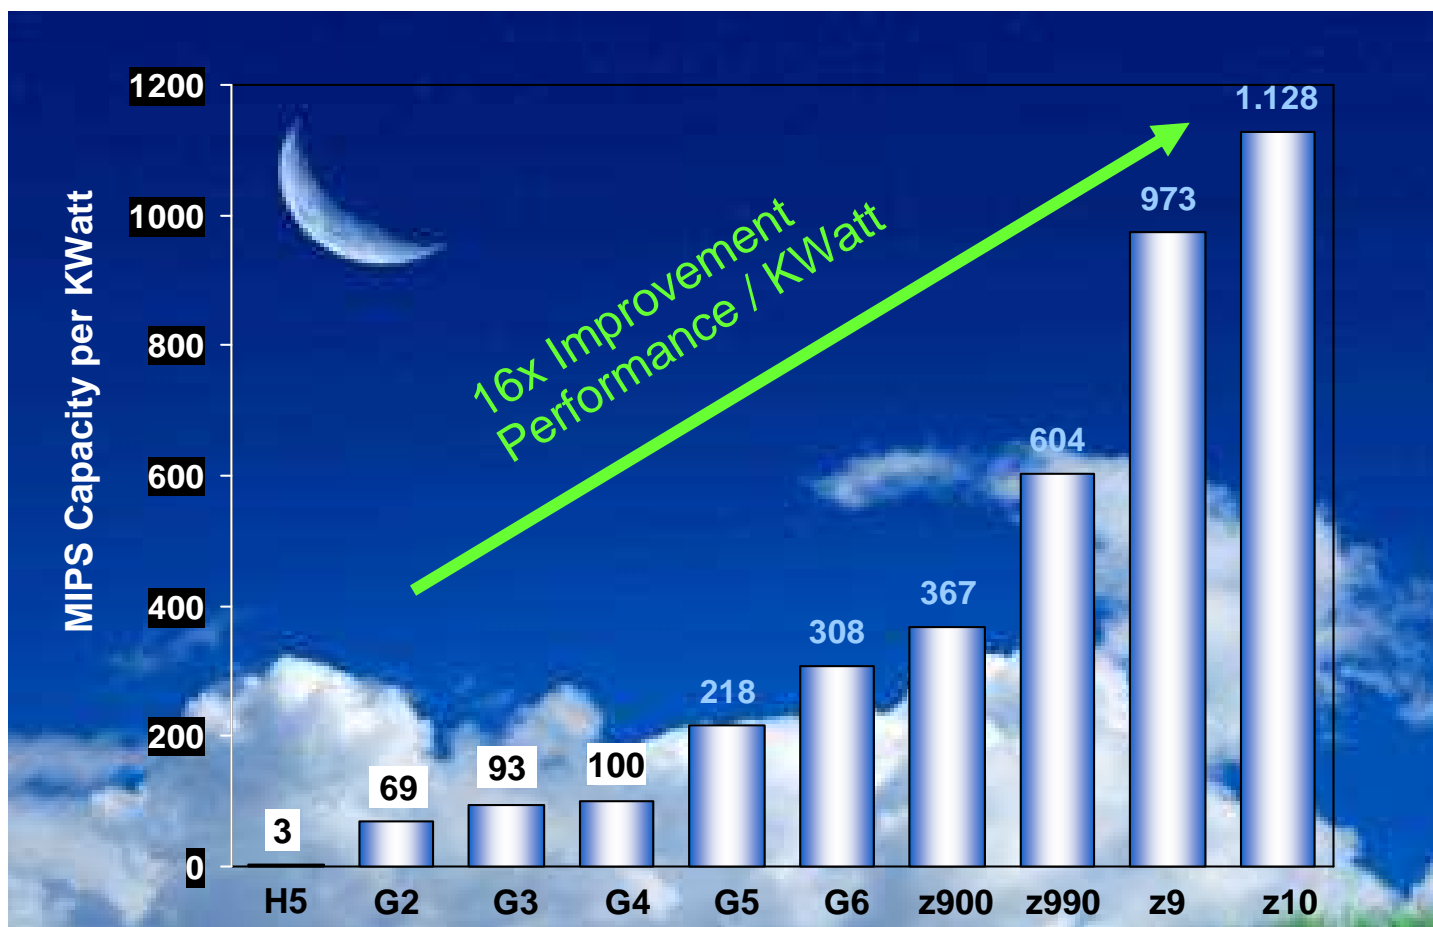
Tracking Energy Consumption within the Infrastructure

- **z10 EC offers a 15% improvement in performance per kWh over z9 EC**
- **ResourceLink™ provides tools to estimates server energy requirements before you purchase a new system or an upgrade**
- **Z10 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 provides highly energy efficient technology

15% improvement in performance per kWh over z9 EC



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



Comprehensive Software Leveraging the Strengths of the z10 EC



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

z/OS V1.10 Preview - Integration with the z10 EC

Supporting System z innovation, raising the IT bar and taking System z to the next level of...

... scalability and performance

- HiperDispatch for intelligent dispatching of work for optimized performance¹
- Up to 1TB of real memory² and 64 processors (zIIPs, zAAPs, and CPs)³ per LPAR
- Extended Address Volume (EAV) capability for large storage volumes, improved storage management^{4,5}
- Large (1 MB) pages expected to reduce memory management overhead for exploiting applications³
- Support for Hardware Decimal Floating Point enables high performance computing for your commercial workloads³
- Support for InfiniBand Coupling Links^{1,6}

... networking and connectivity

- Policy-based networking helps create a network responsive to your application needs¹
- Automatic intrusion defense capabilities⁴

...availability

- Basic HyperSwap – for high availability disk^{3,*}
- Parallel Sysplex and GDPS enhancements

... simplified operations

- Capacity Provisioning Manager can monitor systems and dynamically activate / deactivate capacity³
- New z/OS Management Facility – planned – a single, modern, Web-browser based management console for z/OS, intended to simplify day to day operations and administration of a z/OS system. *

....improved economics

- Additional XML exploitation of specialty engines³
- zIIP assisted z/OS Global Mirror (XRC)³

(1) available with z/OS V1.7 with appropriate maintenance

(2) available with z/OS V1.8 and appropriate maintenance, 1TB memory on z10 E56 and E64 only

(3) available with z/OS V1.9 and appropriate maintenance

(4) planned for z/OS V1.10

(5) with appropriate storage

(6) Planned availability 2Q08

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



*The Future Runs
on System z...*

... And Your Future Begins Today!

Trademarks

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

DB2*	IBM eServer	System z10
Cool Blue	IBM logo*	Tivoli*
DRDA*	OMEGAMON*	WebSphere*
DS8000	Parallel Sysplex*	z9
ESCON*	ResourceLink	z10
FICON*	System p	zArchitecture*
FlashCopy*	System Storage	z/OS*
GDPS*	System x	z/VM*
HiperSockets	System z	z/VSE
HyperSwap	System z9*	zSeries*
IBM*		

* 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.

InfiniBand is a trademark and service mark of the InfiniBand Trade Association.

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.