


# Bringing You Up to Date with IBM System z Hardware



**Mike Augustine**  
**z10 BC Offering Manager**  
 maugust@us.ibm.com

*Thinking Beyond Today*



© 2009 IBM Corporation

## Trademarks

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

Certified Used Equipment	Multiprise*	Tivoli*
DB2*	Parallel Sysplex*	z9
DS8000	PR/SM	z10
Dynamic Infrastructure*	Resource Link	z10 BC
ESCON*	S/390*	z10 EC
FICON*	Scalable Architecture for Financial Reporting	z/Architecture*
GDPS*	Sysplex Timer*	z/OS*
HiperSockets	Systems Director Active Energy Manager	z/VM*
HyperSwap	System Storage	z/VSE
IBM*	System z*	zSeries*
IBM eServer	System z9*	
IBM logo*	System z10	
IMS	System z10 Business Class	

\* 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 there from.

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.


# IBM System z10™ Business Class (z10™ BC)

The Modern Mainframe – Innovative, Cool, Affordable



- Machine Type – 2098 Model E10
  - Single frame, air cooled
  - Non-raised floor option available
- Processor Units (PUs)
  - 12 PUs
  - 2 SAPs, standard
  - Zero spares when all PUs characterized
  - Up to 10 PUs available for characterization
    - Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z10 Application Assist Processors (zAAPs), System z10 Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs)
- Memory
  - System Minimum of 4 GB
  - Up to 256 GB for System, including 8 GB Fixed HSA, standard
    - Up to 248 GB for your use
    - Increments: 4, 8 and 32 GB
- I/O
  - Up to 12 I/O Interconnects per System @ 6 GBps each
  - Fiber Quick Connect for ESCON and FICON LX
  - New OSA-Express3 features

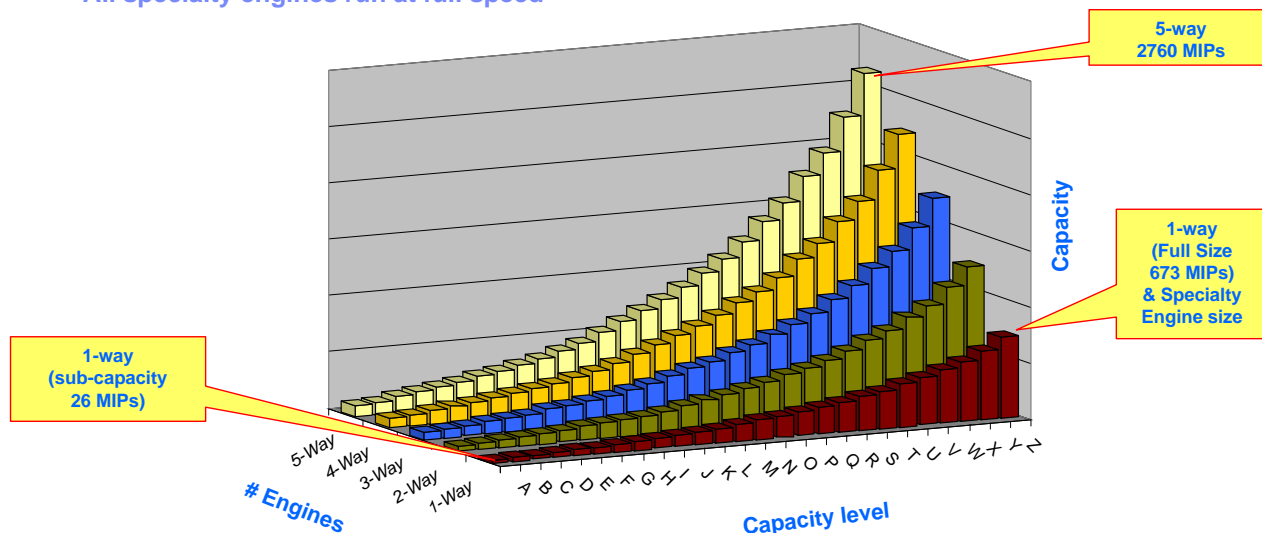
## z10 BC Product Positioning



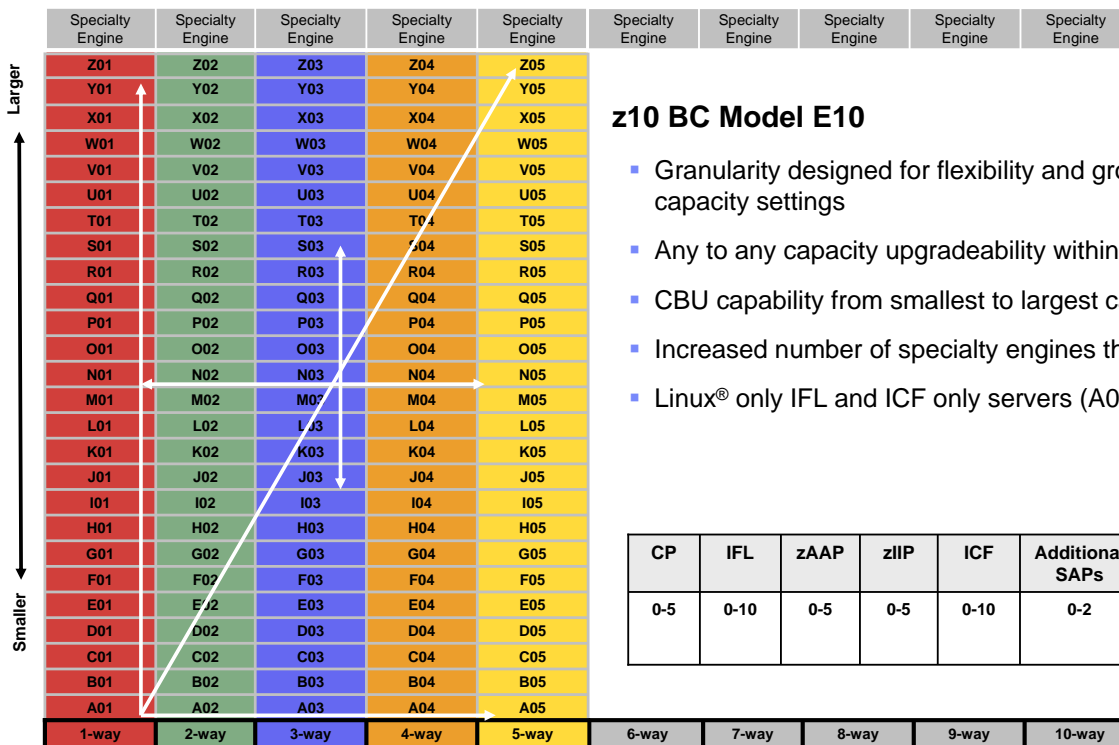
- z10 BC provides increased capacity in a single footprint
  - Faster uniprocessor performance
  - Up to 1.5 times the total system capacity of the z9 BC Model S07
  - Nearly four times the memory of the z9 BC
- Availability improvements via planned outage reduction
- Increased functionality and flexibility in temporary capacity offerings
- Connectivity improvements include bandwidth, throughput, and distance

# z10 BC Sub-capacity Processor Granularity

- The z10 BC has 26 CP capacity levels (26 x 5 = 130)
  - Up to 5 CPs at any capacity level
  - All CPs must be the same capacity level
- The one for one entitlement to purchase one zAAP and/or one zIIP for each CP purchased is the same for CPs of any speed.
  - All specialty engines run at full speed



## The right size for existing and future applications *Innovative, affordable and flexible*



### z10 BC Model E10

- Granularity designed for flexibility and growth with 130 capacity settings
- Any to any capacity upgradeability within the Model
- CBU capability from smallest to largest capacities
- Increased number of specialty engines than z9 BC
- Linux® only IFL and ICF only servers (A00)

CP	IFL	zAAP	zIIP	ICF	Additional SAPs	Spares
0-5	0-10	0-5	0-5	0-10	0-2	0 when fully configured

1-way Capacity Indicator / MIPS / MSU	2-way Capacity Indicator / MIPS / MSU	3-way Capacity Indicator / MIPS / MSU	4-way Capacity Indicator / MIPS / MSU	5-way Capacity Indicator / MIPS / MSU
A01 / 26 / 3	A02 / 48 / 6	A03 / 69 / 9	A04 / 88 / 11	A05 / 107 / 13
B01 / 30 / 4	B02 / 56 / 7	B03 / 80 / 10	B04 / 102 / 13	B05 / 123 / 15
C01 / 38 / 5	C02 / 71 / 9	C03 / 101 / 12	C04 / 129 / 16	C05 / 156 / 19
D01 / 46 / 6	D02 / 85 / 11	D03 / 122 / 15	D04 / 156 / 19	D05 / 189 / 23
E01 / 53 / 7	E02 / 99 / 12	E03 / 141 / 17	E04 / 180 / 22	E05 / 218 / 27
F01 / 59 / 7	F02 / 110 / 14	F03 / 156 / 19	F04 / 200 / 25	F05 / 242 / 30
G01 / 70 / 9	G02 / 130 / 16	G03 / 186 / 23	G04 / 237 / 29	G05 / 287 / 36
H01 / 78 / 10	H02 / 145 / 18	H03 / 207 / 26	H04 / 264 / 33	H05 / 320 / 40
I01 / 88 / 11	I02 / 163 / 20	I03 / 233 / 29	I04 / 298 / 37	I05 / 361 / 45
J01 / 98 / 12	J02 / 182 / 23	J03 / 260 / 32	J04 / 332 / 41	J05 / 402 / 50
K01 / 110 / 14	K02 / 204 / 25	K03 / 292 / 36	K04 / 372 / 46	K05 / 451 / 56
L01 / 130 / 16	L02 / 241 / 30	L03 / 354 / 43	L04 / 440 / 55	L05 / 533 / 66
M01 / 150 / 19	M02 / 279 / 35	M03 / 398 / 49	M04 / 508 / 63	M05 / 615 / 76
N01 / 172 / 21	N02 / 319 / 40	N03 / 456 / 57	N04 / 582 / 72	N05 / 705 / 87
O01 / 193 / 24	O02 / 358 / 44	O03 / 511 / 63	O04 / 653 / 81	O05 / 791 / 98
P01 / 216 / 27	P02 / 401 / 50	P03 / 572 / 71	P04 / 731 / 91	P05 / 885 / 110
Q01 / 242 / 30	Q02 / 449 / 56	Q03 / 641 / 80	Q04 / 819 / 102	Q05 / 991 / 123
R01 / 270 / 33	R02 / 501 / 62	R03 / 715 / 89	R04 / 913 / 113	R05 / 1106 / 137
S01 / 303 / 38	S02 / 563 / 70	S03 / 803 / 100	S04 / 1025 / 127	S05 / 1241 / 154
T01 / 340 / 42	T02 / 631 / 78	T03 / 901 / 112	T04 / 1151 / 143	T05 / 1393 / 173
U01 / 381 / 47	U02 / 707 / 88	U03 / 1010 / 125	U04 / 1289 / 160	U05 / 1561 / 194
V01 / 429 / 53	V02 / 797 / 99	V03 / 1137 / 141	V04 / 1452 / 180	V05 / 1758 / 218
W01 / 480 / 60	W02 / 891 / 111	W03 / 1272 / 158	W04 / 1625 / 202	W05 / 1968 / 245
X01 / 538 / 67	X02 / 999 / 124	X03 / 1426 / 177	X04 / 1822 / 226	X05 / 2206 / 274
Y01 / 615 / 76	Y02 / 1142 / 142	Y03 / 1631 / 202	Y04 / 2083 / 258	Y05 / 2523 / 313
Z01 / 673 / 83 (full speed)	Z02 / 1250 / 155	Z03 / 1784 / 221	Z04 / 2279 / 283	Z05 / 2760 / 342

## z10 Capacity Planning in a nutshell

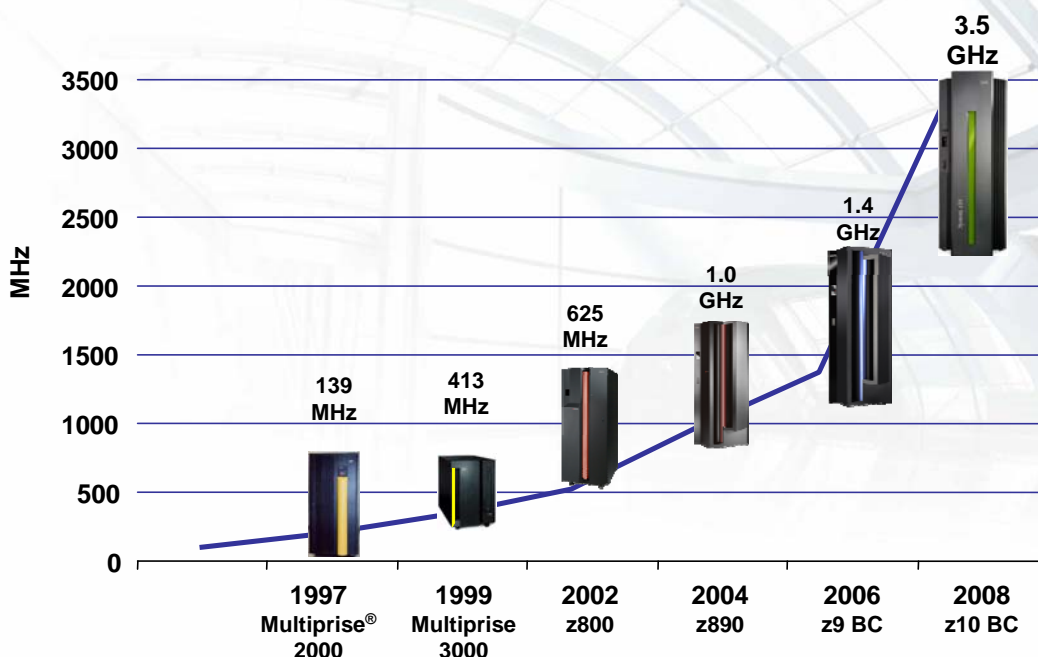


**Don't use "one number" capacity comparisons!**  
**Work with IBM technical support for capacity planning!**  
**Customers can now use zPCR**

The IBM Processor Capacity Reference (zPCR) is a free tool available for download that can be used to size your System z processors.

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381>

## IBM z10 BC continues the CMOS Mainframe heritage



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


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

- z10 BC – Architectural extensions
- Higher frequency CPU

## Making high performance a reality

*Designed for the next evolution of Enterprise applications*

- **Enterprise Quad Core z10 processor chip**
  - 3.5 GHz – additional throughput means improved price/performance
  - Cache rich environment optimized for data serving
  - 50+ instructions added to improve compiled code efficiency
  - Support for 1 MB 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 processor chip

# New levels of application performance

## Designed for an expanded set of workloads

- The z10™ BC can deliver up to 50% more performance for general purpose workloads than an IBM System z9® Business Class (z9® BC)\*
- The uniprocessor can deliver up to 40% more performance than z9 BC uniprocessor\*\*
- Up to 1.9x performance improvements for CPU intensive jobs or tasks
- Up to 10X improvement in decimal floating point instructions
- Up to 10 IFLs for large scale consolidation



### More capacity and engines for traditional growth and consolidation

All performance information was determined in a controlled environment.

\* LSPR mixed workload average running z/OS® 1.9 - z10 BC z05 versus z9 BC z04 \*\* LSPR mixed workload average running z/OS 1.9 - z10 BC z01 versus z9 BC z01

# Reducing costs with System z10 BC - delivering continued price / performance for new workloads

Generation to generation price / performance improvements:	z10 BC
Reduction in software charging units, MSUs, <sup>1</sup> versus z9 BC († Millions of Service Units)	10%
Reduction in software charging units, MSUs, versus z890 or z800 / z900	19% or 27%
Maintenance price per MIPS reduction for equivalent capacity <sup>1</sup>	5%
Maintenance price per MIPS reduction with capacity growth <sup>1</sup>	Up to 10%
Performance improvement for Linux (IFLs), Java (zAAPs) and Integrated Information Processors (zIIPs)	Up to 40%
Typical charge for MES upgrades for IFLs, zAAPs, and zIIPs	0
Technology-driven value	z10 BC
50% price reduction on Specialty engines for System z10 BC <sup>2, 4</sup>	\$47.5 K USD
62% price reduction on System z10 Memory Prices for new workloads when purchased together with Specialty engines <sup>2, 3, 4</sup>	\$2,250 USD




#### Plus

- 100 percent of IBM mainframes are delivered virtualization ready
- On/Off Capacity on Demand (On/Off CoD) enhancements to better manage volatile business requirements


1 – Comparisons shown are z9 BC to z10 BC; 2 - Prices in USD, may vary by country; 3 – Limited to 16GB per engine; 4 – Does not include Internal Coupling Facilities (ICFs)

## z10 BC Upgrade Paths



- Temporary, nondisruptive addition of CP processors, IFLs, ICFs, zAAPs or zIIPs
- New options for reconfiguring specialty engines if the business demands it
- New options for changing On/Off CoD configurations
- Subcapacity CBU engines



## z10 BC – Under the covers Front View





# z10 BC CPC and Memory Drawer Layout



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

# z10 BC PU/SC SCM Components





# z10 BC I/O Drawer

## FRONT VIEW



I/O cards




I/O cards

## REAR VIEW

- Up to 4 I/O drawers supported
  - Up to 8 I/O cards in each drawer – 4 in front and 4 in rear = total of 32
- I/O cards are horizontal; Very important that cables are routed to the side or else concurrent replacement of I/O cards may not be possible
- Concurrent add of drawer
- Concurrent replacement, repair for Systems with >1 I/O drawer
- Drawer can be removed without affecting system input power or power to any other unit

# I/O Subsystem – internal host bus interconnect speeds



STI: Self-Timed Interconnect

# Physical Configuration / Planning



## CONFIGURATION

- ▶ *Single CEC drawer*
  - ▶ **Disruptive memory add**
  - ▶ **Disruptive SCM replacement**
- ▶ *I/O Subsystem:*
  - ▶ **Concurrent add of drawers**
  - ▶ **4 drawers (32 card total)**

## SPACE

- ▶ *Single Frame*
  - ▶ **79" height**  
+3 inches vs. z9 BC  
Height reduction to 72"
  - ▶ **50" frame depth**  
+4 inches vs. z9 BC
  - ▶ **71" depth with covers**  
+9 inches vs. z9 BC

## Energy Management

- ▶ **Same Power / Plugs and service as z890 and z9 BC**
- ▶ **Power Monitoring through IBM Director Power Executive**
- ▶ **Power Executive hosting on Linux for System z**

# z10 BC Configuration Comparisons

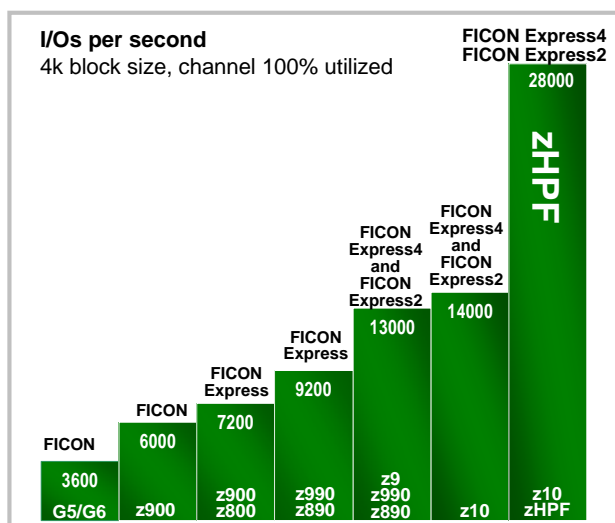
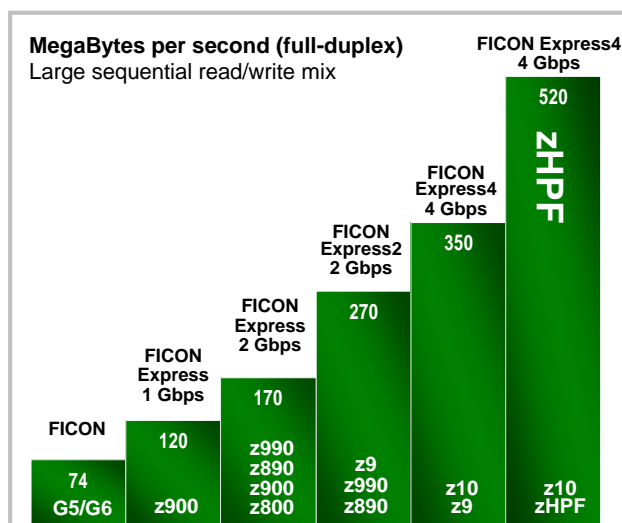
	z9 BC R07	z9 BC S07	z10 BC E10	z10 EC E12
Uniprocessor Performance	470 MIPS		673 MIPS	920 MIPS
System Capacity	26-172 MIPS	193-1748 MIPS	26-2760 MIPS	218-8225 MIPS
Total System Memory	64	64	Up to 256 GB	Up to 384 GB
Configurable Engines	7	7	10	12
Configurable CPs	1-3	0-4	0-5	0-12
LPARS/LCSS	15/1	30/2	30/2	60/4
HiperSockets	16		16	16
I/O Cages/Drawers	1	1	Up to 4	Up to 3
I/O slots per Cage/Drawers	28	28	8	28
FICON Channels	64	112	128	336
OSA Ports (10GbE/1GbE)	16/32	24/48	48/96	48/96
ESCON Channels	240	420	480	1024
STI (z9), IFB (z10) Bandwidth	2.7 GB/sec		6.0 GB/sec	6.0 GB/sec
ICB-4/ISC-3/PSIFB	16/48/00		12/48/12	16/48/32
zIIP/zAAP Maximum Qty	3	3	5	6
Capacity Settings	20	53	130	48
Upgradeable	Upgrade to S07	Upgrade to z9 EC & z10 BC	Upgrade to z10 EC Model E12	Upgrade to z10 EC Models E26, E40, E56 and E64

## z10 BC Channel Type and Crypto Overview

- FICON/FCP up to 32 cards, 128 CHPIDs
  - FICON Express4
  - FICON Express2 (carry forward on upgrade)
  - FICON Express (carry forward on upgrade for FCV)
- Networking
  - OSA-Express3 up to 24 cards, 48 CHPIDs, 96 ports
    - 10 Gigabit Ethernet LR and SR
    - Gigabit Ethernet LX and SX
    - 1000BASE-T Ethernet
  - OSA-Express2 up to 24 cards, 48 CHPIDs & ports
    - 10 Gigabit Ethernet LR (carry forward on upgrade)
    - Gigabit Ethernet LX and SX (limited availability or carry forward on upgrade)
    - 1000BASE-T Ethernet (limited availability or carry forward on upgrade)
  - HiperSockets (internal network)
- ESCON up to 32 cards, 480 CHPIDs
  - April 28, 2009 SOD
    - It is IBM's intent for ESCON channels to be phased out. System z10 EC and System z10 BC will be the last servers to support greater than 240 ESCON channels.
- Coupling Links
  - InfiniBand Coupling Links
    - 12x IB-DDR
    - 1x IB-DDR (z10 EC GA2)
  - ISC-3 (Peer mode only)
  - ICB-4 (except z10 EC E64 and z10 BC E10 with non-raised floor feature)
- Crypto up to 8 cards, 16 engines
  - Crypto Express2
  - 1 or 2 engine options
- Channel types not supported:
  - FICON (pre-FICON Express)
  - OSA-Express
  - ICB-2
  - ICB-3
  - ISC-3 Links in Compatibility Mode
  - PCIXCC and PCICA
  - Parallel (use ESCON Converter)

## z10 High Performance FICON for System z (zHFPF)


- **Simplification of storage area network (SAN) traffic with zHFPF can improve performance**
  - For small data transfers of OLTP and other workloads that exploit the zHFPF protocol, the maximum number of I/Os per second is increased by up to 100%\*
- **Only available on System z10**
  - Supported on FICON Express2 and FICON Express4
- **Requires control unit exploitation and z/OS 1.8 and higher with PTF**
  - IBM System Storage™ DS8000™ Release 4.1



\* Some complex channel programs can not be converted to zHFPF protocol

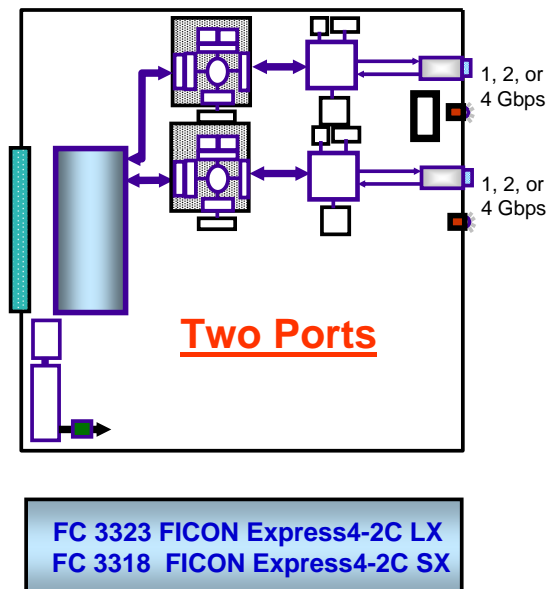
# z10 BC (and EC) FICON Express4

- **FICON Express4 – 4 channels**
- **4 Gbps with Auto-negotiate capability (1, 2, or 4 Gbps)**
- **Up to 128 channels**
- **LX 10 KM, LX 4 KM and SX features**
- **Concurrent repair of optics**
- **Personalize as:**
  - **FC:** Native FICON
    - Support on all operating systems
  - **CTC:** Channel-To-Channel
    - Support on all operating systems
  - **FCP:** Fibre Channel Protocol
    - Support of SCSI devices on/VM, z/VSE, Linux on System z



# z10 BC (only) FICON Express4-2C

- **FICON Express4 - 2 Channels** supports:
  - 4 Gbps with Auto-negotiate capability (1, 2, or 4 Gbps)
  - CHPID types FC and FCP
- **Ordering**
  - One feature = two port increment
  - 2 port card can not be 'upgraded' to 4 ports
- z10 BC FC 3318 and FC 3323 (2 ports) is NOT available on z10 EC and cannot be carried forward from z10 BC to z10 EC





## z10 OSA-Express3

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

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


## z10 OSA-Express3 – 10 GbE

- **10 Gigabit Ethernet LR (Long Reach) and SR (Short Reach)**
  - 1 port per PCI-E adaptor and 2 ports per card
  - Small form factor connector (LC Duplex)
    - LR = Single Mode 9 micron fiber
    - SR = Multimode 50 or 62.5 micron fiber
  - Two CHPIDs, one port each
    - Type OSD (QDIO TCPIP and Layer 2)
- **New Microprocessor and hw data router**
  - Large send packet construction, inspection and routing preformed in hardware instead of firmware
  - Large send for IPv4 traffic
  - Checksum offload
  - Concurrent LIC update
  - Designed to improve performance for standard (1492 byte) and jumbo frames (8992 byte)




# z10 OSA-Express3 GbE – 2 and 4 ports features

- Gigabit Ethernet LX and SX
  - 2 ports (SX only) and 4 ports per feature options
  - Two ports per PCI-E adaptor/CHPID
  - CHPIDs support
    - OSD\* (QDIO TCPIP and Layer 2)
    - OSN (OSA-Express for NCP)
  - Small form factor connector (LC Duplex)




- New microprocessor and hw data router
  - Large send packet construction, inspection and routing preformed in hardware instead of firmware
  - Large send for IPv4 traffic
  - Checksum offload
  - Concurrent LIC update




**\* NOTE:** To use 2-Ports per OSD CHPID, the following is required – z/VSE V4.1+, z/VM V5.2+, z/OS.e V1.8, z/OS V1.8+, zTPF 1.1 PUT 4 with APARs. If this support isn't installed, only port zero on each CHPID (PCI-E adaptor) is 'visible' to the Operating System

# z10 OSA-Express3 – 1000BASE-T, 2 & 4 port features

- 1000BASE-T Ethernet
- Two and Four ports per feature options
  - RJ45, Cat 5 UTP
  - Two ports per PCI-E adaptor/CHPID
  - CHPIDs Supported
    - OSC (OSA-Integrated Console)
    - OSD\* (TCPIP and Layer 2)
    - OSE (Non-QDIO TCPIP and SNA/APPN)
    - OSN (OSA-Express for NCP)
- New microprocessor and hw data router
  - Large send for IPv4 traffic
  - Checksum offload
  - Concurrent LIC update
  - Auto-negotiation 10/100/1000
  - Large send packet construction, inspection and routing performed in hardware



- Notes:
- For CHPID type OSC, Port 0 is only used. Port 1 not 'visible' to OS
  - For CHPID type OSD, OSE, both ports on each PCI-E adaptor are used
  - Each PCI-E adaptor can be defined as the same or different CHPID type



**\* NOTE:** To use 2-Ports per OSD CHPID, the following is required – z/VSE V4.1+, z/VM V5.2+, z/OS.e V1.8, z/OS V1.8+, zTPF 1.1 PUT 4 with APARs. If this support isn't installed, only port zero on each CHPID (PCI-E adaptor) is 'visible' to the Operating System


# Reduce risk with IBM's world-class Business Resiliency solutions

- New I/O drawer with concurrent add/replace for drawers 2-4
- Preplanning capabilities to avoid future planned outages, e.g. dynamic LPAR allocation without a system outage and plan ahead memory
- Basic HyperSwap™ improves storage availability
- Updated CP Assist for Cryptographic Function (CPACF) support
- Integrated cryptographic accelerator
- Tamper-resistant Crypto Express2 feature Audit logging on new Trusted Key Entry (TKE) 5.3 with optional Smart Card reader
- System z10 BC – [EAL5 certified on May 4, 2009](#)




## z10 Cryptographic Support

- **CP Assist for Cryptographic Function (CPACF)**
  - Standard on every CP and IFL
  - **Supports the following algorithms:**
    - DES, TDES, AES-128, [AES-192](#), [AES-256](#)
    - SHA-1, SHA-224, SHA-256, [SHA 384 & SHA 512](#)
    - Pseudo random Number Generation (PRNG)
    - SHA-1, SHA-256, and SHA-512 are shipped enabled
  - [Up to 4096-bit RSA keys](#)
  - Random Number Generation Long (8 bytes to 8096 bits)



**Crypto Express2-1P**  
z10 BC only

- **Crypto Express2**
  - Two features – 1 or 2 Coprocessor options
    - minimum of 2 features required
  - Two configuration modes
    - Coprocessor (default), designed to provide both “secure key” and “Public key” function
    - Accelerator (configured from the HMC), designed to provide only “Public key” function with enhanced performance
  - Three configuration options (Default set to Coprocessor)
  - Concurrent Patch
- **Dynamic add crypto to LPAR**
  - No recycling of LPAR
  - No POR required



**Crypto Express2**  
z10 BC and z10 EC

## New HMC Functions

- **HMC Instant Messenger**
  - 2-way conversation different HMC/SE or same HMC/SE
  - Broadcast to all logged on users
- **HMC Screen Capture to removeable media (USB memory stick)**
- **Scheduled Operations for Change LPAR Weights**
- **IPV6 Support**
- **IBM Systems Director Active Energy Manager (power monitor of IBM Systems)**
- **z/VM Tower Enhancements**
  - z/VM directory manager if present
  - Support adding, deleting, altering z/VM guests
  - Support viewing and altering the Virtual Machine Resource Manager (VMRM) configuration
  - Support basic z/VM TCP/IP configuration
  - Allow the z/VM "MAP server" to run on the HMC, thus allowing for CIM management for multiple z/VM images to be accomplished via the HMC



## 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* – On/Off Capacity on Demand (On/Off CoD), Capacity Backup Upgrade (CBU) and the new Capacity for Planned Event (CPE)
- **No interaction with IBM at time of activation**
  - Broader 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
  - Extend CBU entitlements through the purchase of additional CBU years
- **Offerings can be reconfigured or replenished dynamically**
  - Modification possible even if offering is currently active
  - Some permanent upgrades permitted while temporary offerings are active





## System z10 BC Operating System Support

Operating System	ESA/390 (31-bit)	z/Architecture (64-bit)
z/VSE Version 4 Release 1 <sup>(2)(5)</sup> and 2 <sup>(5)</sup>	No	Yes
z/VSE Version 3 Release 1 <sup>(2)(4)</sup>	Yes	No
Linux on System z <sup>(2)</sup> , Red Hat RHEL 5, & Novell SUSE SLES 10 and 11	No	Yes
Linux on System z <sup>(2)</sup> , Red Hat RHEL 4, & Novell SUSE SLES 9	Yes	Yes
z/VM Version 5 Release 3 <sup>(3)</sup> and 4	No*	Yes
z/OS Version 1 Releases 8, 9 and 10	No	Yes
z/OS Version 1 Releases 7 <sup>(1)(2)</sup>	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 System z10 to enable HiperDispatch on System z10 (does not require a zIIP). z/OS V1.7 support was withdrawn September 30, 2008. The Lifecycle Extension for z/OS V1.7 (5637-A01) makes fee-based corrective service for z/OS V1.7 available through September 2009. With this Lifecycle Extension, z/OS V1.7 supports the z10 BC server. Certain functions and features of the z10 BC server require later releases of z/OS. For a complete list of software support, see the PSP buckets and the Software Requirements section of the System z10 BC announcement letter, dated October 21, 2008.
2. Compatibility Support for listed releases. Compatibility support allows OS to IPL and operate on a z10 BC.
3. Requires Compatibility Support which allows z/VM to IPL and operate on the System z10 providing System z9 functionality for the base OS and Guests. \*z/VM supports 31-bit and 64-bit 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. **z/VSE V3.1 support will be withdrawn effective July 31, 2009.**
5. z/VSE V4 is designed to exploit 64-bit real memory addressing, but will not support 64-bit virtual memory addressing.
6. **z/VM 5.2 support was withdrawn effective April 30, 2009.**

**Note: Refer to the z/OS, z/VM, z/VSE subsets of the 2098DEVICE Preventive Planning (PSP) bucket prior to installing a z10 BC**

## IBM Lifecycle Extension for z/OS V1.7

- **For customers that have not completed their migration from z/OS V1.7 to z/OS V1.8 or z/OS V1.9, IBM provides fee-based corrective service for z/OS V1.7 after its end of support date of September 30, 2008.**
  - For z/OS V1.7 only
  - Only provides defect support
  - Up to two years of support starting October 1, 2008
  - Six month minimum duration required
- **This Lifecycle Extension for z/OS V1.7 does not change coexistence, migration, and fallback policies for z/OS.**
  - z/OS V1.7 coexists, migrates, and falls back between z/OS V1.8 and z/OS V1.9.
- **Additional information**
  - Lifecycle Extension for z/OS V1.7 Announcement letter  
<http://www.ibm.com/systems/z/os/zos/>  
(Announcement section)
  - Lifecycle Extension for z/OS V1.7 FAQ  
<ftp://ftp.software.ibm.com/common/ssi/sa/st/n/zsq03023usen/ZSQ03023USEN.PDF>
  - z/OS Installation and Migration  
<http://www.ibm.com/systems/z/os/zos/installation/>
  - Pricing  
<http://ibm.com/zseries/swprice/other/>



## System z strategy

### 1. *Innovate to address the IT infrastructure challenges of today and the future*

- Further simplify, consolidate and reduce the costs of an IT infrastructure
- Integrate, virtualize and coherently manage the multiple and varied elements of business applications
- Scale up and leverage System z strengths in data serving

### 2. *Extend strengths of System z*

- Invest for continued leadership in System z: performance, virtualization, enterprise security, enterprise business continuity
- Extend System z best of breed capabilities to a broader set of workloads
- Deploy optimized technologies for specific applications or components

### 3. *Expand the ecosystem and support core applications that our clients want*

- Recruit new solutions and solution providers and integrators
- Expand skills and capabilities across the globe



## The road ahead for Dynamic Infrastructure with System z

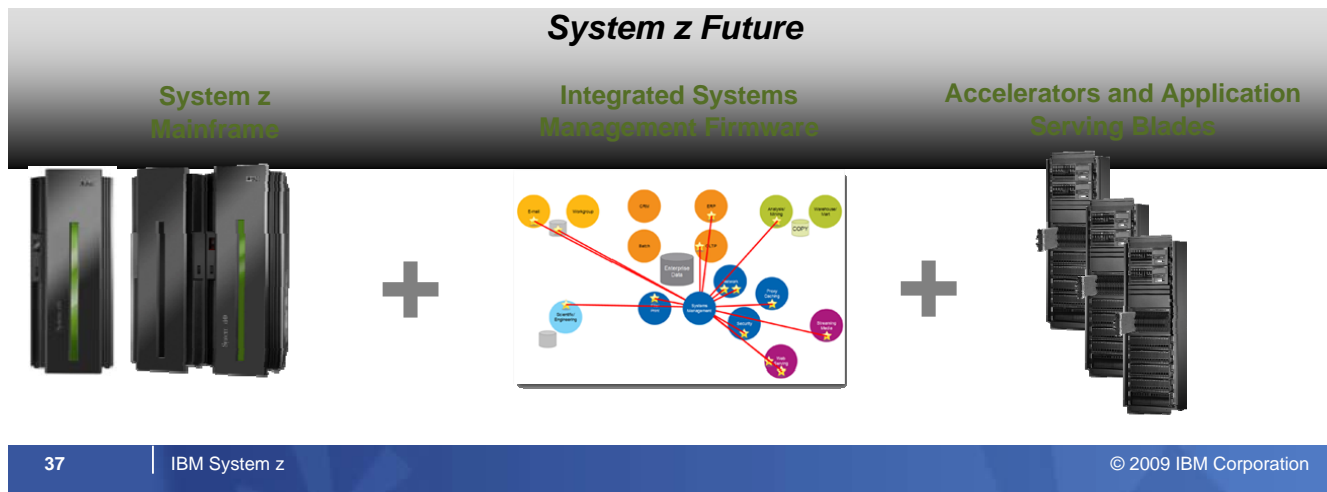
Our goal is to extend mainframe qualities to heterogeneous platforms within a Dynamic Infrastructure to Support Critical Applications



- **End-to-End Systems Management**
- **Policy based Automation Across the Applications Stack**
- **Mainframe Security**
- **Consolidated Disaster Recovery**
- **Improved Economies of Scale and Efficiency**

## Extending System z management and QoS to non-System z technologies

- A single management and policy framework across Web serving, transaction, and database to lower the cost of enterprise computing
- Mainframe QoS characteristics will be extended to acceleration appliances and application servers to manage risk
- The dynamic resource management of the mainframe is extended to all devices within a multi-tier architecture to improve service




## System z: leadership capabilities for a Dynamic Infrastructure Now...



### Role of System z today:


- Secure and resilient enterprise data hub
- Enterprise server for mission critical applications requiring high levels of availability and security (eg, OLTP)
- Highly efficient consolidation platform for exceptional cost savings



### Mainframe Qualities:

- High application-level availability, not just Hardware or OS
- Iron-clad security (EAL-5)
- Extreme scalability
- Integrated capabilities for workload management, provisioning, etc.
- Extensive monitoring and audit capabilities

## System z: leadership capabilities for a Dynamic Infrastructure ... and in the future



### Role of System z today:

- Secure and resilient enterprise data hub
- Enterprise server for mission critical applications requiring high levels of availability and security (eg, OLTP)
- Highly efficient consolidation platform for exceptional cost savings

### Role of System z tomorrow

- An extremely Cost-efficient platform across broader enterprise workloads
- Multi-tier business application host for a wider range of critical applications
- System z QoS (RAS) and management extended to heterogeneous platforms and applications

*Thank you for your  
time and for doing  
business with IBM*

## Additional Reference Charts

## New Statements of Direction\*


- **The System z10 will be the last server to support connections to the Sysplex Timer® (9037).** Servers that require time synchronization, such as to support a base or Parallel Sysplex, will require Server Time Protocol (STP). STP has been available since January 2007 and is offered on the System z10, System z9, and zSeries 990 and 890 servers.
- **The System z10 will be the last server to support Dynamic ICF expansion.** This is consistent with the Statement of Direction in Hardware Announcement 107-190, dated April 18, 2007: "IBM intends to remove the Dynamic ICF expansion function from future System z servers."
- **ICB-4 links to be phased out.** IBM intends to not offer Integrated Cluster Bus-4 (ICB-4) links on future servers. IBM intends for System z10 to be the last server to support ICB-4 links as originally stated in Hardware Announcement 108-154, dated February 26, 2008.
- **ESCON channels to be phased out.** It is IBM's intent for ESCON channels to be phased out. System z10 EC and System z10 BC will be the last server to support greater than 240 ESCON channels.
- **IBM intends to support optional water cooling on future high-end System z servers.** This cooling technology will tap into building chilled water that typically exists within the datacenter for computer room air conditioning systems. External chillers or special water conditioning will typically not be required. Water cooling technology for high-end System z servers will be designed to deliver improved energy efficiencies.
- **IBM intends to support the ability to operate from High Voltage DC power on future System z servers.** This will be in addition to the wide range of AC power already supported. A direct HV DC datacenter power design can improve data center energy efficiency by removing the need for an additional DC to AC inversion step.

\* 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.

# The Mainframe Charter

## The evolution of System z

community  
innovation  
value



### Community

Support programs designed to foster vitality in the IBM mainframe community, helping to promote a strong application portfolio and world-class support services.\*

### Innovation

Provide leadership in innovation to enhance the use of the IBM mainframe to support increasingly integrated and flexible business processes for the on demand business.\*

### Value

Enhance the value proposition and lower the cost of computing of mainframe solutions in a way that is compelling, clear, and consistent.\*

Explore the  
integral components  
of the IBM community  
ecosystem

\* Excerpted from the Mainframe Charter – August 2003

## IBM System z family

### IBM System z9 EC (2094)



- Announced 7/05 - Superscalar Server with up to 64 cores
- 5 models – Up to 54-way
- Granular Offerings for up to 8 CPs
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP, zIIP
- On Demand Capabilities
  - CUoD, CIU, CBU, On/Off CoD
- Memory – up to 512 GB
- Channels
  - Four LCSSs
  - Multiple Subchannel Sets
  - MIDAW facility
  - 63.75 subchannels
  - Up to 1024 ESCON channels
  - Up to 336 FICON channels
  - FICON Express2 and 4
  - OSA 10 GbE, GbE, 1000BASE-T
  - Coupling Links
- Configurable Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 60 logical partitions
- Enhanced Availability
- Operating Systems
  - z/OS, z/VM, z/VSE, TPF, z/TPF, Linux on System z

### IBM System z9 BC (2096)



- Announced 4/06 - Superscalar Server with 8 cores
- 2 models – Up to 4-way
- High levels of Granularity available
  - 73 Capacity Indicators
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP, zIIP
- On Demand Capabilities
  - CUoD, CIU, CBU, On/Off CoD
- Memory – up to 64 GB
- Channels
  - Two LCSSs
  - Multiple Subchannel Sets
  - MIDAW facility
  - 63.75 subchannels
  - Up to 420 ESCON channels
  - Up to 112 FICON channels
  - FICON Express2 and 4 Gbps
  - OSA 10 GbE, GbE, 1000BASE-T
  - Coupling Links
- Configurable Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 30 logical partitions
- Enhanced Availability
- Operating Systems
  - z/OS, z/OS.e, z/VM, z/VSE, TPF, z/TPF, Linux on System z

### IBM System z10 EC (2097)



- Announced 2/08 - Server with up to 77 cores
- 5 models – Up to 64-way
- Granular Offerings for up to 12 CPs
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP, zIIP
- On Demand Capabilities
  - CoD, CIU, CBU, On/Off CoD, CPE
- Memory – up to 1.5 TB for Server and up to 1 TB per LPAR
- Channels
  - Four LCSSs
  - Multiple Subchannel Sets
  - MIDAW facility
  - 63.75 subchannels
  - Up to 1024 ESCON channels
  - Up to 336 FICON channels
  - FICON Express4 and 2 Gbps
  - OSA 10 GbE, GbE, 1000Base-T
  - InfiniBand Coupling Links
- Configurable Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 60 logical partitions
- Enhanced Availability
- Operating Systems
  - z/OS, z/VM, z/VSE, TPF, z/TPF, Linux on System z

### IBM System z10 BC (2098)



- Announced 10/08 – Server with 12 cores
- Single model – Up to 5-way
- High levels of Granularity available
  - 130 Capacity Indicators
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP, zIIP
- On Demand Capabilities
  - CoD, CIU, CBU, On/Off CoD, CPE
- Memory – up to 248 GB
- Channels
  - Two LCSSs
  - Multiple Subchannel Sets
  - MIDAW facility
  - 63.75 subchannels
  - Up to 480 ESCON channels
  - Up to 128 FICON channels
  - FICON Express4 and 2 Gbps
  - OSA 10 GbE, GbE, 1000BASE-T
  - InfiniBand Coupling Links
- Configurable Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 30 logical partitions
- Enhanced Availability
- Operating Systems
  - z/OS, z/OS.e, z/VM, z/VSE, TPF, z/TPF, Linux on System z

# IBM System z Announcement & Withdrawal Dates

2097 (z10 EC)	2098 (z10 BC)	2094 (z9 EC)	2096 (z9 BC)	2066 (z800)	2064 (z900)	2086 (z890)	2084 (z990)
Available February, 26, 2008	Available October, 28, 2008	Available July 27, 2005	Available April 27, 2006	Withdrawn December 31, 2005	Withdrawn June 30, 2006	Withdrawn September 30, 2008	Withdrawn June 30, 2008

### Announcement Letters

- 2094 - IBM United States Announcement 105-241  
- IBM United States Announcement 106-293
- 2096 - IBM United States Announcement 106-287
- 2097 - IBM United States Announcement 108-154  
- IBM United States Announcement 108-296
- 2098 - IBM United States Announcement 108-754

### Withdrawal Letters

- 2066 - IBM United States Withdrawal Announcement 905-220
- 2064 - IBM United States Withdrawal Announcement 906-060
- 2086 - IBM United States Withdrawal Announcement 907-134  
- IBM United States Hardware Announcement 108-328
- 2084 - IBM United States Withdrawal Announcement 908-057

Note: Even though machines have been withdrawn from marketing, valid upgrade paths from these models to current models are still supported.

## z10 BC – Functions and Features Overview - 1

One Hardware Model
New CPC and I/O Drawers
Faster Uni Processor – 3.5 Ghz
Up to 5 CPs
130 CP Capacity Settings
Replaceable Processor Unit and System Controller - Single Chip Modules
Up to 248 GB Usable Memory
Separate 8 GB of fixed HSA
HiperDispatch
Capacity Provisioning Support
Large Page Support (1 MB)
Hardware Decimal Floating Point



Hot Pluggable I/O Drawer
InfiniBand (IFB) I/O interconnect
InfiniBand Coupling Links
STP over InfiniBand
FICON® Express4 Performance Improvement
New Family of OSA-Express3
HiperSockets Layer 2 Support
Enhanced CPACF with SHA 512, AES 192 and 256-bit keys
CBU & On/Off CoD Increased Functionality and Flexibility
Scheduled Outage Reduction
Improved RAS
Power Monitoring Support

## z10 – Functions and Features Overview – 2



- Plan Ahead memory
- Capacity Provisioning Enhancements  
– Capacity Tokens
- LPAR Dynamic PU reassignment
- Change weights of LPARs
- HMC Instant Messenger
- HMC and z/VM System Management
- CFCC Level 16
- STP Enhancements
- PSIFB replacement for ISC-3
- High Performance FICON for System z
- Crypto TKE 5.3 Enhancements
- z/VM Image Mode
- QDIO Connection Isolation



## z10 BC to z9 BC Functional Comparison

<b>Processor / Memory</b>	<ul style="list-style-type: none"> <li>▪ Uniprocessor 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>▪ Up to 1.40 performance improvement over z9 BC uniprocessor*</li> <li>▪ Up to 1.50 times system capacity performance improvement over z9 BC**</li> <li>▪ 3.5 GHz processor chip for z10 BC vs. 1.4 GHz for z9 BC</li> <li>▪ z10 BC has 1 and z9 BC has 2 models</li> <li>▪ z10 BC has up to 10 cores to configure, up to 7 on z9 BC</li> <li>▪ z10 BC has up to 130 Capacity settings vs. 73 on the z9 BC</li> <li>▪ z10 BC has up to 256 GB vs. up to 64 GB on z9 BC</li> <li>▪ z10 BC has fixed 8 GB HSA, z9 BC had HSA 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 BC has up to 5 logical processors in an LPAR vs. 4 on z9 BC</li> <li>▪ z10 BC 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 BC New HiperSockets Layer 2 and Multiple Write Facility</li> <li>▪ Up to 128 FICON channels on z10 BC vs. 112 on z9 BC</li> <li>▪ z10 BC can support up to 480 vs. 420 for z9 BC</li> <li>▪ z10 BC has industry standard 6 GBps InfiniBand supports high speed connectivity and high bandwidth vs. z9 BC using 2.7 GBps Self Time Interconnects (STIs)</li> <li>▪ New I/O Drawer</li> <li>▪ Coupling with InfiniBand<sup>1</sup> – improved distance and potential cost savings</li> <li>▪ Improved AES 192 and 256, stronger hash algorithm with Secure Hash Algorithm (SHA-512) and support for longer Personal Account Numbers up to 19</li> <li>▪ New family of OSA-Express3 features for z10</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> <li>▪ Enhanced I/O structure</li> </ul>	<ul style="list-style-type: none"> <li>▪ z10 &amp; z/OS (1.9) for policy based advice and automation</li> <li>▪ z10 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</li> <li>▪ z10 BC has 'hot-pluggable' I/O drawers</li> </ul>
<b>Environmentals</b>	<ul style="list-style-type: none"> <li>▪ Monitoring</li> </ul>	<ul style="list-style-type: none"> <li>▪ z10 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 1.9 - z10 BC Z01 versus z9 BC Z01

\*\* This is a comparison of the z10 BC 10-way and the z9 BC 7-way and is based on LSPR mixed workload average running z/OS 1.9



## z9 BC to z890 Functional Comparison

<b>Processor / Memory</b> <ul style="list-style-type: none"> <li>▪ Uniprocessor Perf.</li> <li>▪ System Capacity</li> <li>▪ Processor Design</li> <li>▪ Models</li> <li>▪ Processing Units (PUs)</li> <li>▪ Granular Capacity</li> <li>▪ Memory</li> </ul>	<ul style="list-style-type: none"> <li>▪ Up to 1.35x performance improvement over z890 Capacity Setting 170 uniprocessor *</li> <li>▪ Up to 1.30x system capacity performance improvement over z890 Capacity Setting 470 **</li> <li>▪ z9 BC adds instructions, new technology, larger L2 cache</li> <li>▪ z890 has one and z9 BC has 2 hardware models, both with a book</li> <li>▪ z9 BC has up to 7 PUs to configure Vs 4 on z890</li> <li>▪ z9 BC has up to 73 Capacity settings</li> <li>▪ z9 BC has up to 64 GB vs. up to 32 GB on z890</li> </ul>
<b>Virtualization</b> <ul style="list-style-type: none"> <li>▪ LPARs</li> <li>▪ FCP</li> <li>▪ PR/SM™ Dispatching</li> </ul>	<ul style="list-style-type: none"> <li>▪ z9 BC has up to 30 LPARs</li> <li>▪ z9 BC supports N_Port ID Virtualization</li> <li>▪ z9 BC has separate Processor core pools for CPs, IFLs, ICFs, zAAPs, zIIPs</li> </ul>
<b>Connectivity</b> <ul style="list-style-type: none"> <li>▪ HiperSockets</li> <li>▪ FICON for SANs</li> <li>▪ OSA for LANs</li> <li>▪ Total channels</li> <li>▪ Internal I/O Bandwidth</li> <li>▪ Enhanced I/O structure</li> </ul>	<ul style="list-style-type: none"> <li>▪ Both - 16 HiperSockets, z9 BC adds IPv6 support, z890 IPv4 only</li> <li>▪ Up to 112 FICON channels on z9 BC, 80 on z890</li> <li>▪ Both - 48 OSA ports, z9 BC adds OSA-Express2 1000BASE-T</li> <li>▪ Same - Up to 420 channels</li> <li>▪ z9 BC has up to 16 STIs at 2.7 GB/s, z990 has 8 STIs at 2GB/s</li> <li>▪ z9 BC has RII for availability, z890 Single I/O Bus</li> </ul>
<b>On Demand / RAS</b> <ul style="list-style-type: none"> <li>▪ Capacity BackUp</li> <li>▪ MBA repair</li> <li>▪ STI repair</li> <li>▪ Driver Upgrade</li> </ul>	<ul style="list-style-type: none"> <li>▪ For CPs, IFLs, ICFs, zAAPs, zIIPs on z9 BC, CPs only on z890</li> <li>▪ Concurrent on z9 BC, unscheduled outage on z890</li> <li>▪ Concurrent on z9 BC without loss of I/O (RII). z890, I/O for STI lost</li> <li>▪ Concurrent on z9 BC versus scheduled outage on z990<sup>1</sup></li> </ul>

\* LSPR mixed workload average running z/OS 1.7 – z9 BC Z01 versus z890 capacity setting 170  
 \*\* This is a comparison of the z9 BC 7-way and the z890 4-way and is based on LSPR mixed workload average running z/OS 1.7  
 Note 1: Concurrent driver upgrades possible only at certain MCL levels.

## Consolidation with Linux may help drive down cost

### ***Economics of IFLs<sup>1</sup> and z/VM virtualization 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
- 50% price reduction on IFLs for System z10 BC, => \$47,500<sup>2</sup>
- 62% price reduction on System z10 memory prices for new workloads when purchased with Specialty Engines<sup>2,4</sup>, => \$2250 per GB<sup>2,3</sup>
- Typically MES upgrades are priced at no charge when moving to new technology

### ***System z may help customers to become more energy efficient with a Dynamic Infrastructure***

- Consolidation onto System z may reduce energy consumption and save floor space

Over 2450 applications  
 are supported on  
 Linux on System z -  
 15% growth in 2008



<sup>1</sup> Integrated Facility for Linux - IFL  
<sup>2</sup> Prices are stated in US currency and may vary by country.  
<sup>3</sup> Limited to 16 GB per engine.  
<sup>4</sup> Specialty engines do not include Internal Coupling Facilities (ICF)