



# **IBM System z10 Business Class for GSE Leipzig October 27, 2008**

by  
**Klaus-Dieter Mueller**  
IBM Executive Briefing Center Europe  
[kmuller@de.ibm.com](mailto:kmuller@de.ibm.com)

October 21, 2008 Announcement



**Smart. Cool. Affordable.**

# Trademarks

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

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

\* 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.

# Introducing the IBM System z10 Business Class (z10 BC)

## *Simplified*



*Change the operational model to reduce cost*

- Operational superiority through extreme virtualization and leading management
- Robustness through leadership security and availability
- Save more by adding further virtualized applications


## *Shared*



*Today's ultimate shared resource pool*

- Processors, memory and channels can be shared across all the applications
- Hundreds of applications can share resources on IBM System z<sup>®</sup> simultaneously

## *Dynamic*



*Automating service delivery to meet new business needs*

- Automated provisioning and management optimize for high value services
- Services can be added or deleted on the fly
- Applications can be integrated with centralized enterprise-wide real time data



*An ideal platform for the Cornerstone of the New Enterprise Data Center*

# Introducing the IBM System z10 Business Class

***Smart, Cool, Affordable***



**z Can Do IT**

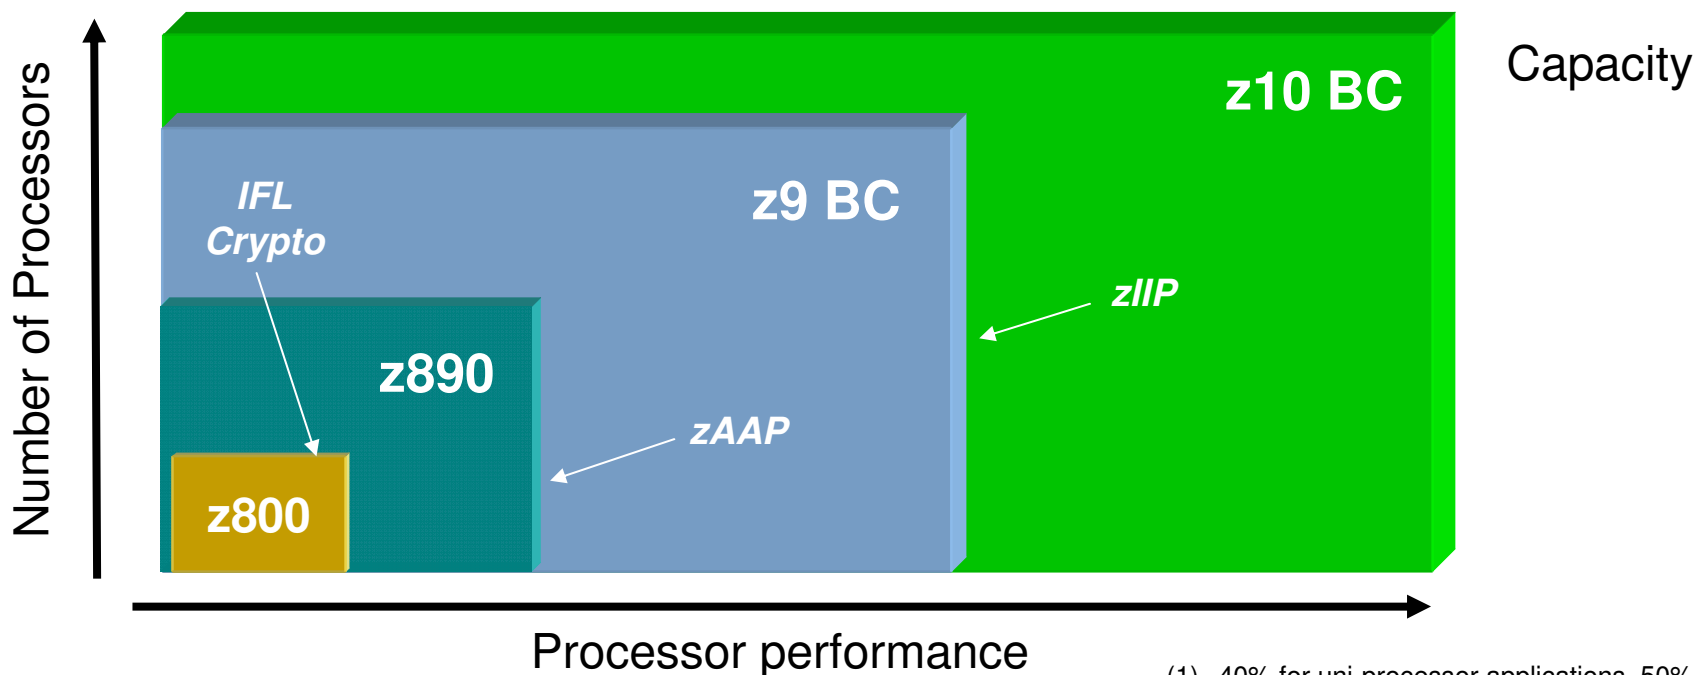


- **Smart** technologies enable you to quickly respond to changing requirements
- **Cool** Dramatically reduce energy costs compared to distributed server alternatives
- **Affordable** A state-of-the-art platform that can help you save through upgrades & consolidation
  - Up to 50% more performance at half the price for incremental Linux workloads <sup>(1)</sup>

(1) Specialty engines and associated memory on z10 BC are priced 50% and 60% less, respectively, compared to the z9 BC. “incremental” refers to the addition or expansion of IFLs on new or upgraded z10 BC. 50% performance gain based on multi-core performance improvement as measured by IBM for 5 IFL cores. 40% for single cores. ICFs excluded.

# z10 BC delivers new levels of performance for today's demanding applications

- Up to ten 3.5 GHz z10 Quad-core processors
  - Up to 5 processors for z/OS and/or z/VSE
  - Up to 10 IFLs for large scale Linux consolidation
  - 2.5x performance boost CPU Intensive workloads and up to 10x improvement in Decimal Floating Point instructions
  
- 40-50% more performance than the highly acclaimed IBM System z9<sup>®</sup> Business Class (z9 BC) <sup>(1)</sup>



(1) 40% for uni-processor applications, 50% for 5-processors

# The Smart and Affordable platform for new workloads

- *Why put new workloads on z?*
  - Efficiencies through co-location of applications and data
  - Huge economies of scale
  - Legendary Quality of Service
  
- *Announcing lower prices for new workloads:*
  - Integrated Facility for Linux (IFL), System z Integrated Information Processor (zIIP) and System z Application Assist Processor (zAAP) reduced by 50% <sup>(1)</sup> z10 BC Only
  - Memory prices reduced by over 60% <sup>(2)</sup>
    - For memory purchased with a new Specialty Engine
    - Valid on IBM System z10 Enterprise Class (z10 EC™) and z10 BC
  - Special pricing promotions from Novell and Red Hat for Linux support

(1) Based on US list price, May vary in other countries. Internal Coupling Engines excluded,

(2) Up to 16GB DRAM per specialty engine



# The Smart upgrade to the z10 BC

- *Unprecedented investment protection*
  - Upgradeable from System z9 BC *and* z890
- *Increased financial value*
  - Specialty engines upgrade typically at no extra cost
    - Nearly 50 - 100% more capacity than z9 BC and IBM eServer™ zSeries® 890 (z890) respectively
  - IBM software and maintenance technology dividend helps lower costs
    - 10% fewer MSUs from z9 BC (19% on z890)
    - A potential average savings of 5% on MLC software on z/OS® workloads (10% average on z890)
    - A potential 5% savings on HW maintenance on a dollar per MIP basis for zero capacity growth, and up to 10% for higher growth
  - Talk to your ISVs about Usage Based Software Pricing to help lower your cost of growing with System z
- *And all the benefits from the latest innovations in the z10 system design*
  - Infiniband®, Instant On-Demand, ....



# It's Cool

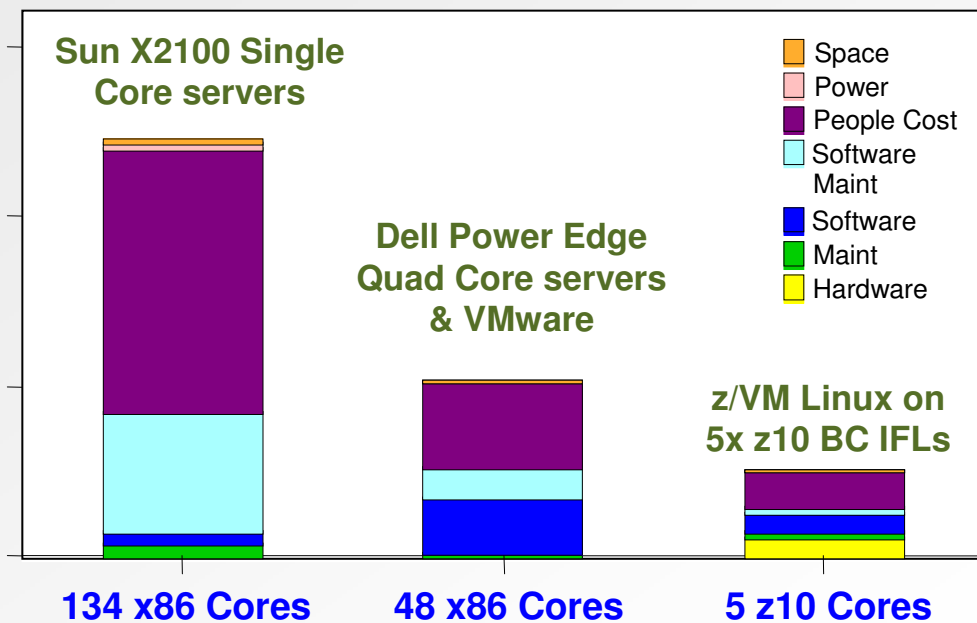
Add one Domino IFL for up to 7,500 users

- 1,000 users for the energy of a 100 watt bulb
- Hardware at less than \$10 per user
- No extra floor space required

## Consolidating 134 Linux servers to 5 IFLs

### Save up to 50% over x86 w/ VMware

Oracle DB Workload, 3-Year Total IT Cost



Your IT Cost may vary

All performance information was determined in a controlled environment. Actual results may vary.

- **Domino and Linux on z**
- **Consolidate: Save 50% over x86 virtualization**





# System z Ecosystem - The growing zCommunity

*Dramatic growth responding to market demand*



- ★ Expanding application portfolio on System z driven by ISV growth:
  - ★ YTD: 130 New ISVs, 600+ New applications/tools
  - ★ 5,000+ applications available
  - ★ 2,450+ applications for Linux on System z
  - ★ Over 1,400 ISVs building applications
  
- ★ IBM Academic Initiative driving Skills growth on System z
  - ★ 50,000 Students attended mainframe education
  - ★ 481 Schools registered offering 29 separate courses and more to come
  - ★ Student Mainframe Contest:
    - ★ 10 contests, more than 1,000 schools, more than 8,000 students

# The modern mainframe for small and medium enterprises

*The mainframe made over – Smart, Cool, Affordable*

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

**Machine Type: 2098**

**Single Model: E10**

**Single Frame**

**Non-raised floor option**



### **Processor Cores:**

- Faster and larger Processor core
  - **z10 Enterprise Quad Core at 3.5 GHz**
  - **2.5x-z9 BC, 3.5x-z890, 5.6x-z800**
- More configurable Processor Cores

### **Memory:**

- Maximum memory
  - **120 GB – Oct08**
  - **248 GB – Jun09**
- 8 GB HSA standard
- Lower 4 GB entry point

### **I/O:**

- New I/O drawer (RAS)
- 6 Gbps InfiniBand® host buses for I/O
- New OSA-Express3
- New High Performance FICON® for System z

# Tracking energy consumption within the infrastructure

## *System z10 – Smart, cool, manageable*


- **Resource Link provides tools to estimate server energy requirements before you purchase a new system or an upgrade**
- **Has energy efficiency monitoring tool**
  - Introduced on IBM System z9 platform in April 2007
  - Power and thermal information displayed via the System Activity Display (SAD)
- **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



# 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.\*

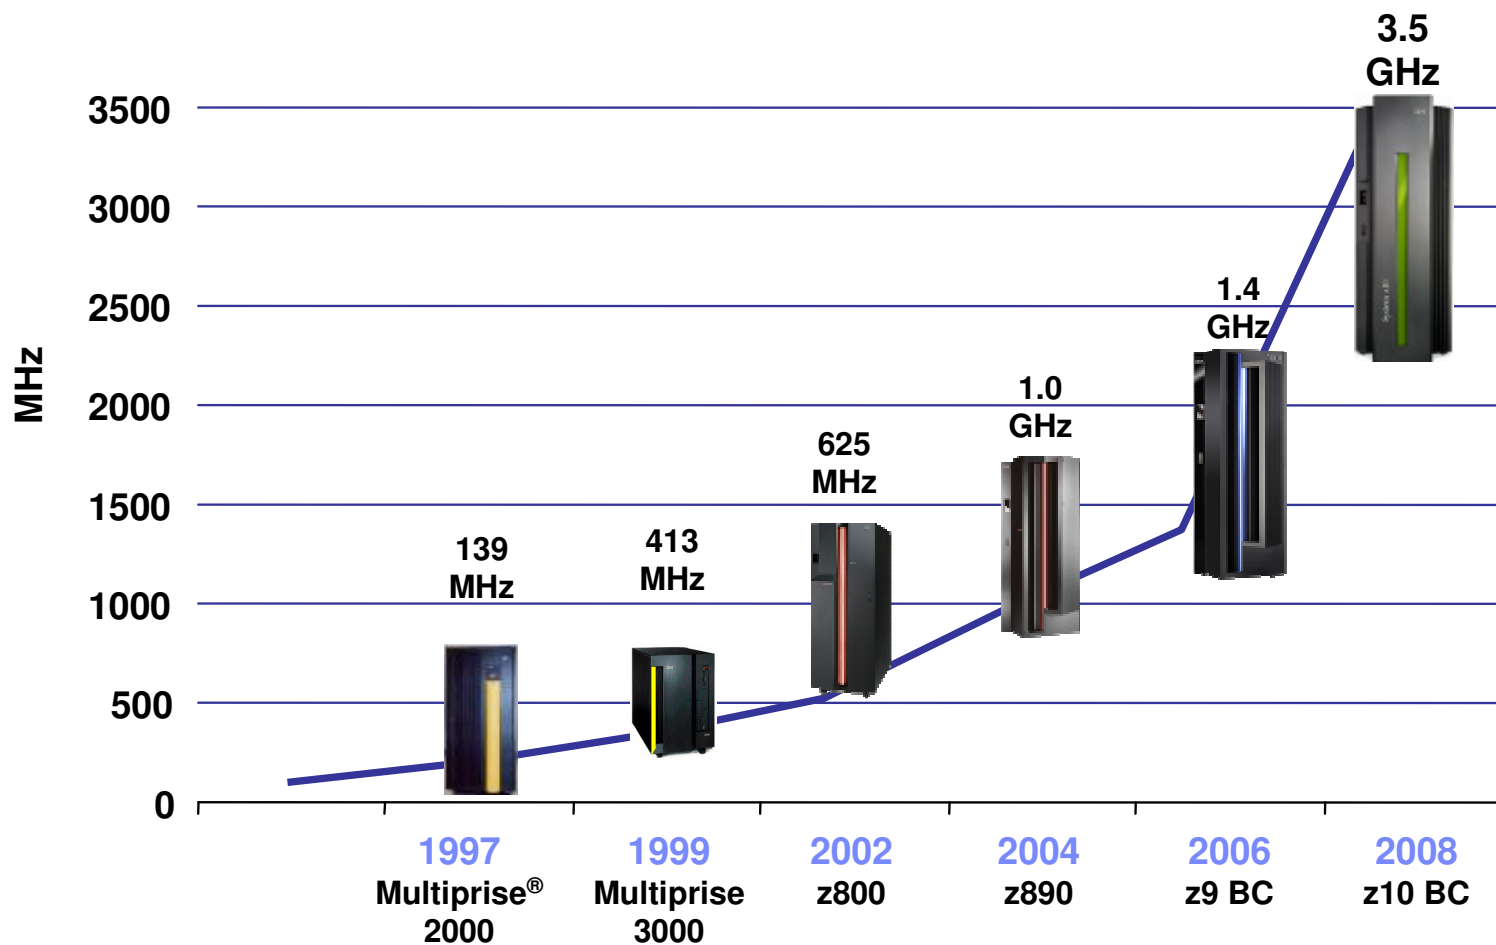
### ***Value***

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 z10 BC continues the CMOS Mainframe heritage



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

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

- z10 BC - Architectural extensions
- Higher frequency CPU

# z10 BC Overview



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

# z10 BC – Under the covers Front View

Power Supplies

Internal Battery (optional)

CPC (SCMs, Memory, MBA, HCA and FSP) Drawer


2 x Support Elements

I/O Drawer #3


I/O Drawer #2

I/O Drawer #1

I/O Drawer #4



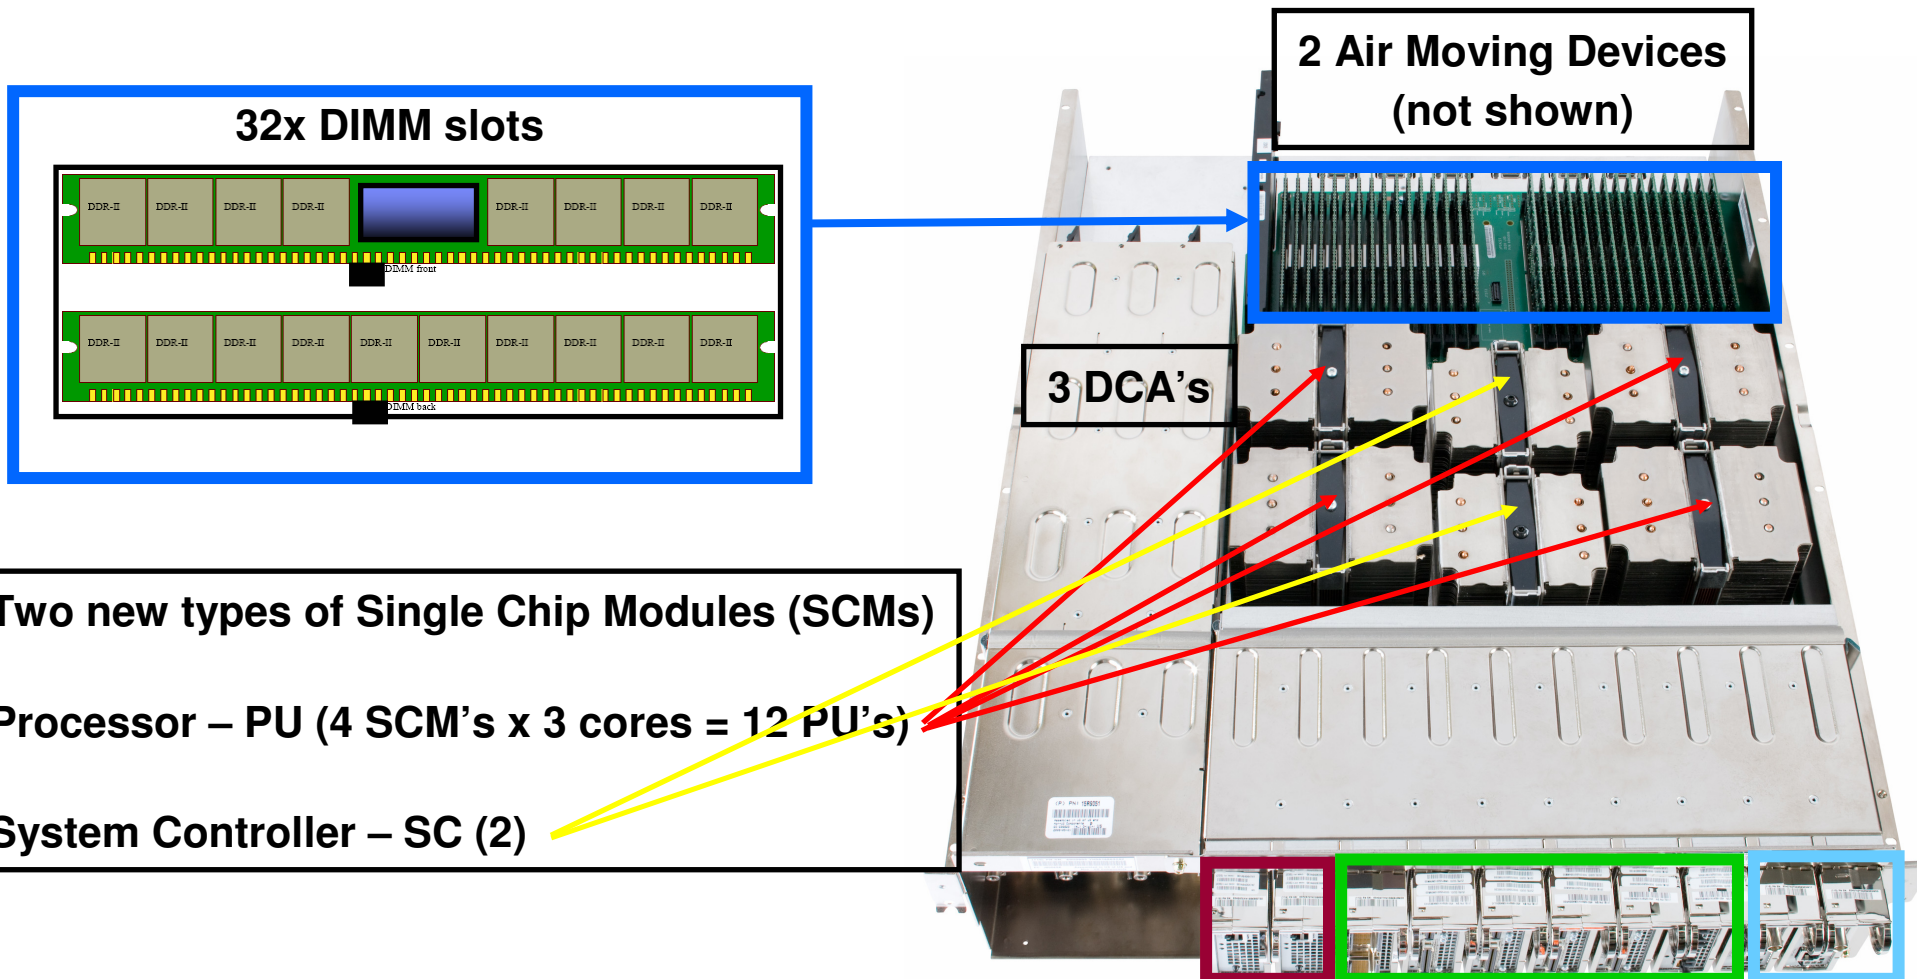
4 x I/O Drawers



Fiber Quick Connect (FQC) Feature (optional – not shown)



# z10 BC CPC and Memory Drawer Layout



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



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

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




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



# z10 BC PU/SC SCM Components



## SCM Components



# z10 BC SCM Vs z10 EC MCM Comparison

## z10 BC SCMs

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



## z10 EC MCM

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

Single PU Chip  
without heatsink



Single SC Chip  
without heatsink





# z10 BC CPC Drawer Components


Up to 32 DIMMS



CPC Drawer



4 x PU and  
2 x SC pluggable  
SCMs



PU chip, SC chips,  
Land Grid Array  
(LGA) socket,  
Indium foil



2 x OSC/ETR Cards


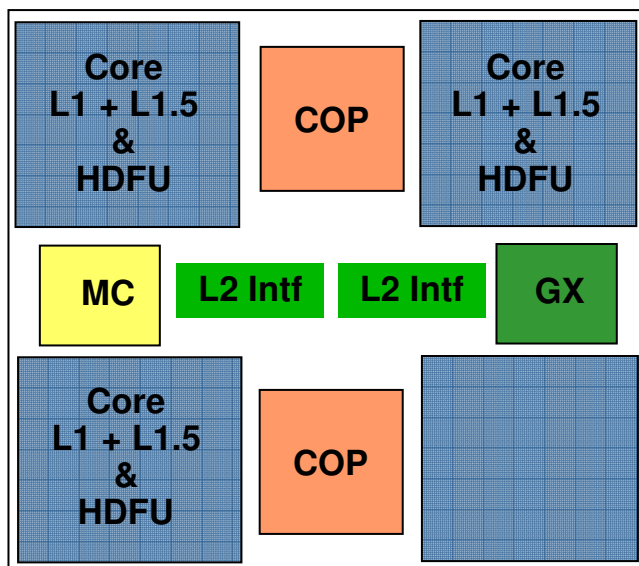


DCA Power



I/O Hub for fanout slots

# z10 BC – Enterprise Quad Core z10 PU Chip



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

# z10 BC Concurrent PU Conversions

- **Must order (characterize one PU core as) a CP, an ICF or an IFL**
- **Concurrent processor upgrade is supported if PUs are available**
  - Add CP, IFL, unassigned IFL, ICF, zAAP, zIIP or optional SAP
- **PU Conversions**
  - Standard SAP cannot be converted to other PU types

To From.....	CP	IFL	Unassigned IFL	ICF	zAAP	zIIP	Optional SAP
CP	x	Yes	Yes	Yes	Yes	Yes	Yes
IFL	Yes	x	Yes	Yes	Yes	Yes	Yes
Unassigned IFL	Yes	Yes	x	Yes	Yes	Yes	Yes
ICF	Yes	Yes	Yes	x	Yes	Yes	Yes
zAAP	Yes	Yes	Yes	Yes	x	Yes	Yes
zIIP	Yes	Yes	Yes	Yes	Yes	x	Yes
Optional SAP	Yes	Yes	Yes	Yes	Yes	Yes	x

Exceptions: Disruptive if ALL current PUs are converted to different types may require individual LPAR disruption if dedicated PUs are converted.

# z10 BC – Granularity and scalability

Z01	Z02	Z03	Z04	Z05
Y01	Y02	Y03	Y04	Y05
X01	X02	X03	X04	X05
W01	W02	W03	W04	W05
V01	V02	V03	V04	V05
U01	U02	U03	U04	U05
T01	T02	T03	T04	T05
S01	S02	S03	S04	S05
R01	R02	R03	R04	R05
Q01	Q02	Q03	Q04	Q05
P01	P02	P03	P04	P05
O01	O02	O03	O04	O05
N01	N02	N03	N04	N05
M01	M02	M03	M04	M05
L01	L02	L03	L04	L05
K01	K02	K03	K04	K05
J01	J02	J03	J04	J05
I01	I02	I03	I04	I05
H01	H02	H03	H04	H05
G01	G02	G03	G04	G05
F01	F02	F03	F04	F05
E01	E02	E03	E04	E05
D01	D02	D03	D04	D05
C01	C02	C03	C04	C05
B01	B02	B03	B04	B05
A01	A02	A03	A04	A05
1-way	2-way	3-way	4-way	5-way

## z10 BC Model E10

- Granularity designed for flexibility and growth
  - Up to 130 Capacity Settings
- Any to any capacity upgradeability within the Hardware Model
- CBU capability from smallest to largest capacities
  - For CPs
  - All Specialty Engines and SAPs (always full capacity)
- Increased number of Specialty Engines
- Linux and CF only servers
- Capacity Indicator A00 is for systems with IFL(s) or ICF(s) only
  - Additional options for CBU i.e. CBU from A01 to Z05

Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine
------------------	------------------	------------------	------------------	------------------	------------------	------------------	------------------	------------------	------------------

# z10 BC Model Structure

- **Single Model E10, Machine Type – 2098**
- **Model number indicates PUs available for characterization**
  - Single serial number
  - PU core characterization is identified by number of features ordered
- **Two System Assist Processors (SAPs) per System**
- **z10 BC Capacity Indicators**
  - nxx, where n = subcapacity engine size and xx = number of CPs
  - Total 130 Capacity Indicators
  - A00 for systems with IFL(s) or ICF(s) only. For z9 BC it was Z00.
- **Memory DIMM sizes: 2 GB, 4 GB and 8 GB**
  - Maximum physical memory: 128 GB per System (256 GB, June 30, 2006)
  - Minimum physical installed = 16 GB of which 8 GB is for Fixed HSA
  - For 8 to 32, 4GB increments, from 32 to 120, 8 GB increment, from 120-248, 32 GB increments

Models	PU SCMs	PUs for Customer	Max Available Subcapacity CPs	Standard SAPs	Standard Spares	CP/IFL/ICF/zAAP/zIIP**	Max Customer Memory	Max Channels
E10	4	10	5	2	0	5/10/10/5/5	248 GB***	480 *

Notes:

\* Max is for ESCON channels.

\*\* For each zAAP and/or zIIP installed there must be a corresponding CP. The CP may satisfy the requirement for both the zAAP and/or zIIP. The combined number of zAAPs and/or zIIPs can not be more than 2x the number of general purpose processors (CPs).

\*\*\* Initially 120 GB. 248 GB available June 30, 2009

# z10 BC Memory Configuration, 4 - 120 GB

Memory Offering (GB) *Customer purchased memory*	Memory with 8 GB HSA	# of cards plugged DIMM size (GB)		Physically Plugged
		2 GB	4GB	
<b>4 GB Increments</b>				
4	12	8	0	16
8	16	8	0	16
12	20	12	0	24
16	24	12	0	24
20	28	16	0	32
24	32	16	0	32
28	36	20	0	40
32	40	20	0	40
<b>8 GB Increments</b>				
40	48	24	0	48
48	56	28	0	56
56	64	32	0	64
64	72	0	20	80
72	80	0	20	80
80	88	0	24	96
88	96	0	24	96
96	104	0	28	112
104	112	0	28	112
112	120	0	32	128
120	128	0	32	128

Non-disruptive with LICCC



# z10 BC Memory Configuration, 120 - 248 GB

Memory Offering (GB) *Customer purchased memory*	Memory with 8 GB HSA	# of cards plugged DIMM size (GB) 8 GB	Physically Plugged
<b>32 GB Increments</b>			
<b>152</b>	<b>160</b>	<b>20</b>	<b>160</b>
<b>184</b>	<b>192</b>	<b>24</b>	<b>192</b>
<b>216</b>	<b>224</b>	<b>28</b>	<b>224</b>
<b>248</b>	<b>256</b>	<b>32</b>	<b>256</b>

**Note: The above memory options will be available June 30, 2009**

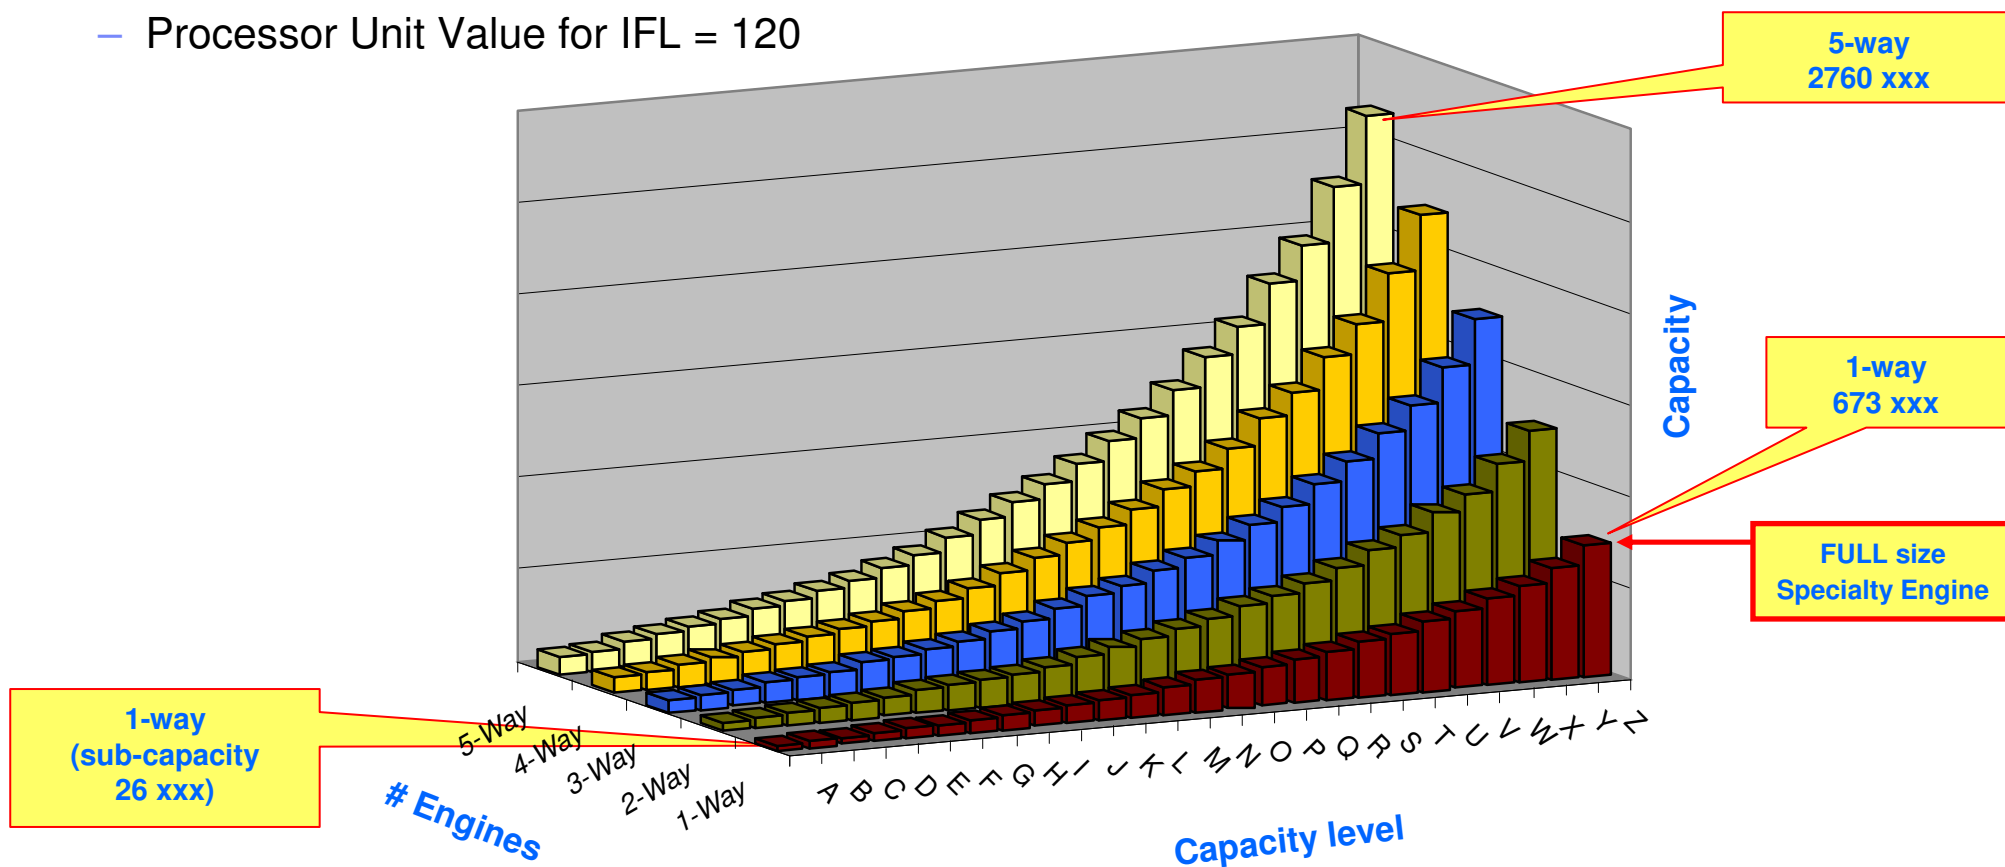
# z10 BC Plan Ahead Memory

- **Provides the ability to plan for non-disruptive memory upgrades**
  - Memory cards are pre-installed based on planned target capacity
- **Pre-installed memory is activated by installing a new LICCC**
  - Orderable via Resource Link by the customer (CIU upgrade)
  - Orderable as an ordinary MES by IBM
  - Memory upgrade orders use the pre-installed memory first
- **Pre-planned memory install**
  - FC 1991 - Charged when physical memory is installed used to track the quantity of physical increments of plan ahead memory capacity
    - Cost part pre-paid
  - Increment size of 4 GB (based on minimum memory purchase increment)
- **Pre-planned memory activation**
  - FC 1992 - Charged when Plan Ahead Memory is enabled based on the amount of Plan Ahead memory that is being activated
    - Remaining cost paid at time of activation
  - Subsequent memory upgrade orders will use up the Plan Ahead memory first
- **Plan Ahead Memory is NOT temporary CoD or CBU memory (Removing memory is disruptive)**

# 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
  - Processor Unit Value for IFL = 120

Number of z10 BC CPs	Base Ratio	Ratio z9BC to z10BC
1 CP	z9 BC Z01	1.40
2 CPs	z9 BC Z02	1.36
3 CPs	z9 BC Z03	1.30
4 CPs	z9 BC Z04	1.28
5 CPs	Z9 BC Z04	1.54



# z10 BC Upgrade Paths

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



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

# I/O Subsystem – internal host bus interconnect speeds



**InfiniBand I/O Bus**  
z10  
2008



**STI**  
z9  
2005



**2.7 GBps**



**STI**  
z990/z890  
2003



**2 GBps**



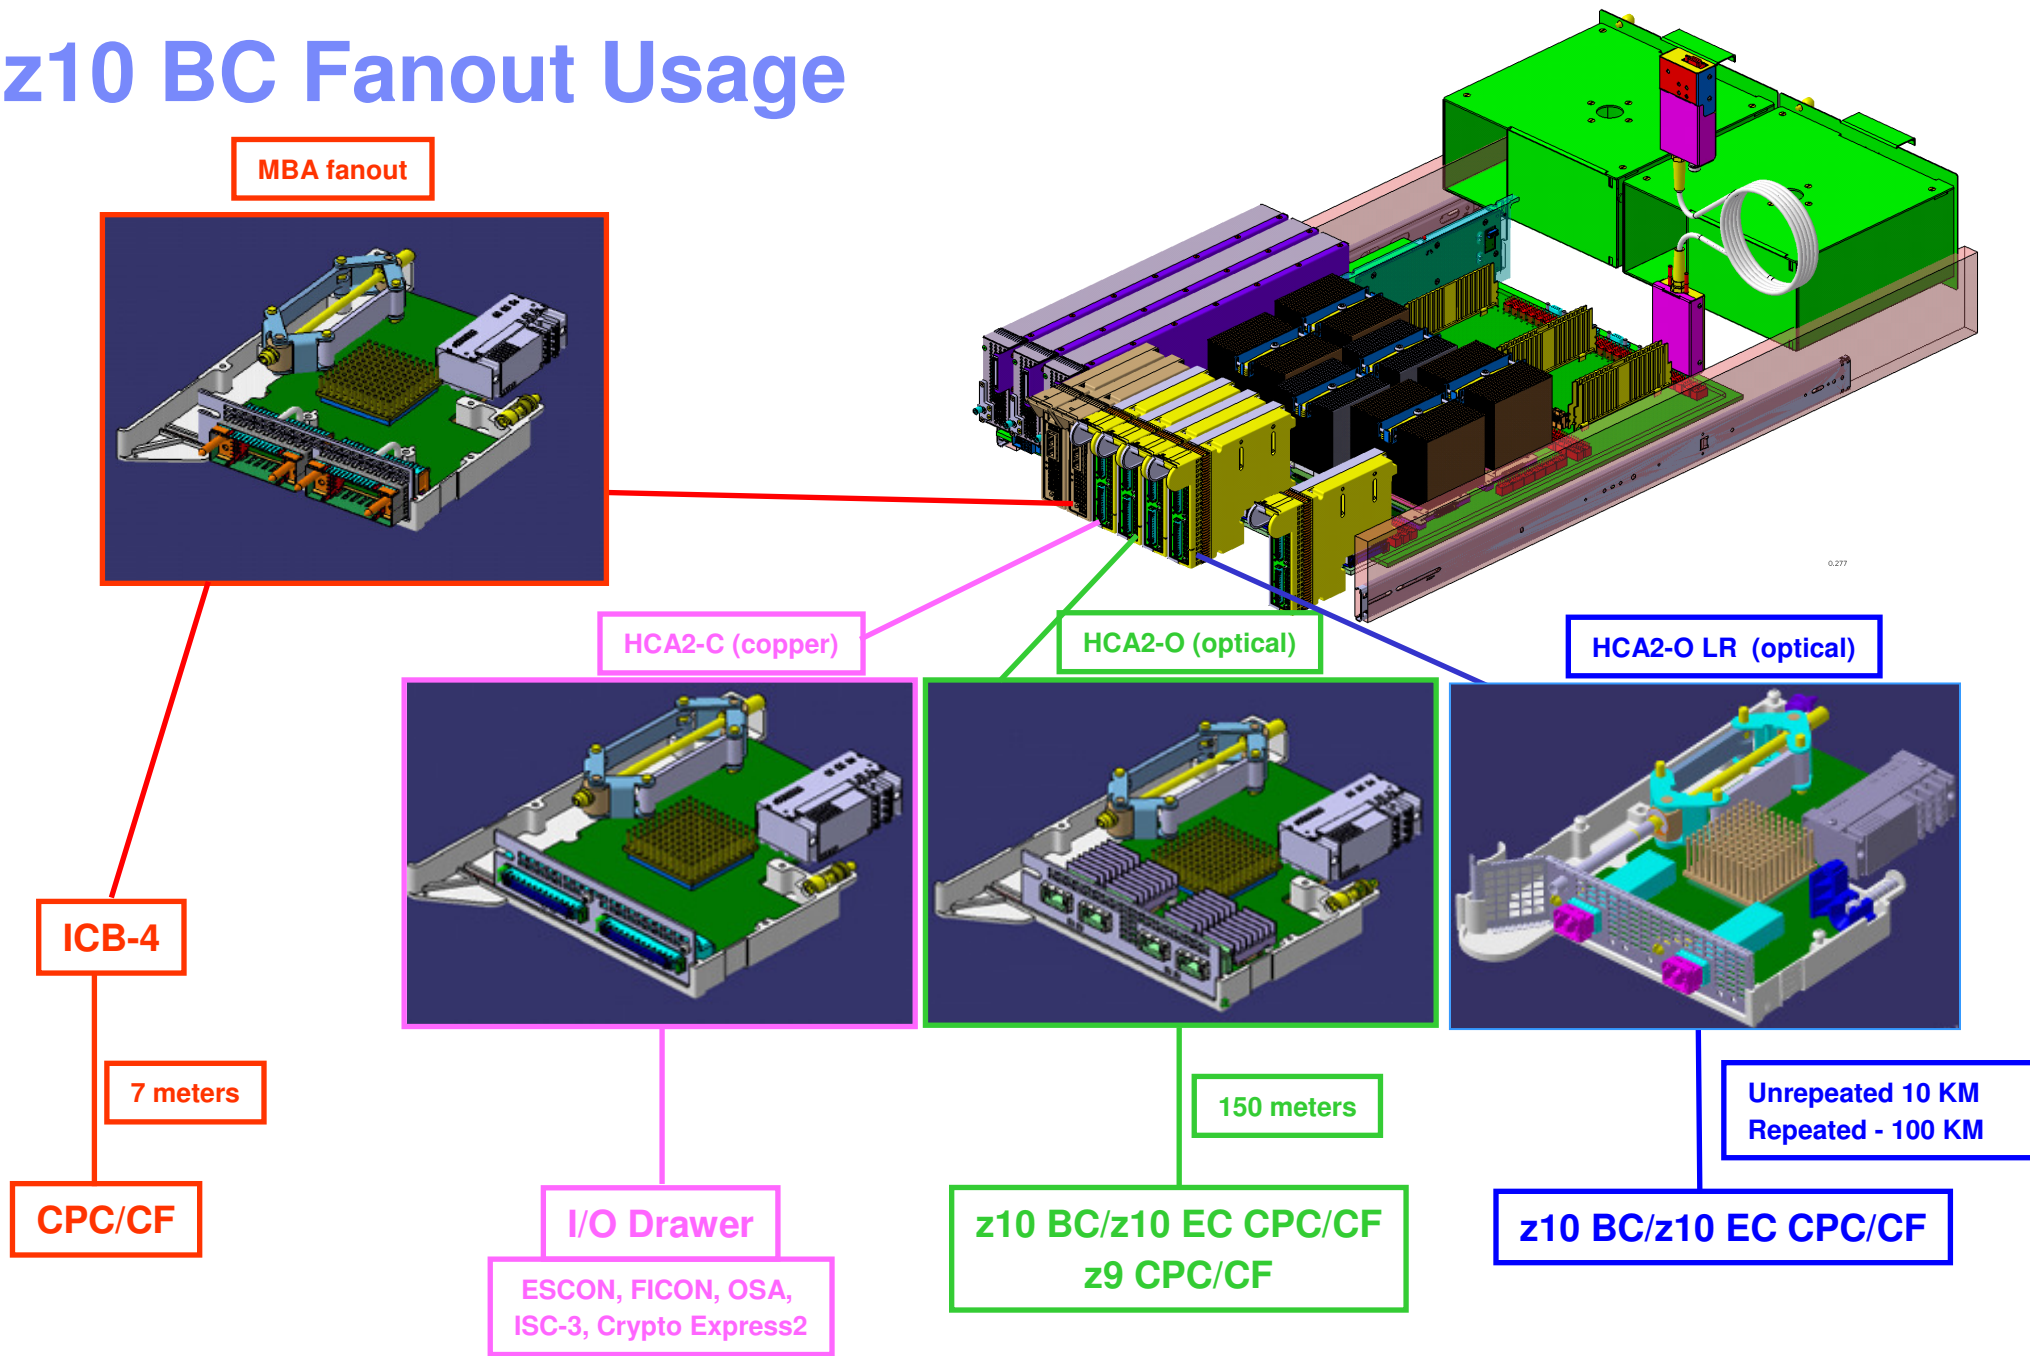
**STI**  
z900/z800  
200x



**1 GBps**



STI: Self-Timed Interconnect

# z10 BC Fanout Usage



# z10 BC I/O Drawer

## FRONT VIEW



## REAR VIEW

- The new I/O drawer together with logic board are field replaceable units (FRUs)
- Drawer can be removed without affecting system input power or power to any other unit
- Up to 4 I/O drawers supported
  - Up to 8 I/O cards in each drawer
  - 4 in front and 4 in rear
- 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
- Concurrent replacement, repair for Systems with >1 I/O drawer
- No support for I/O cage from prior Servers

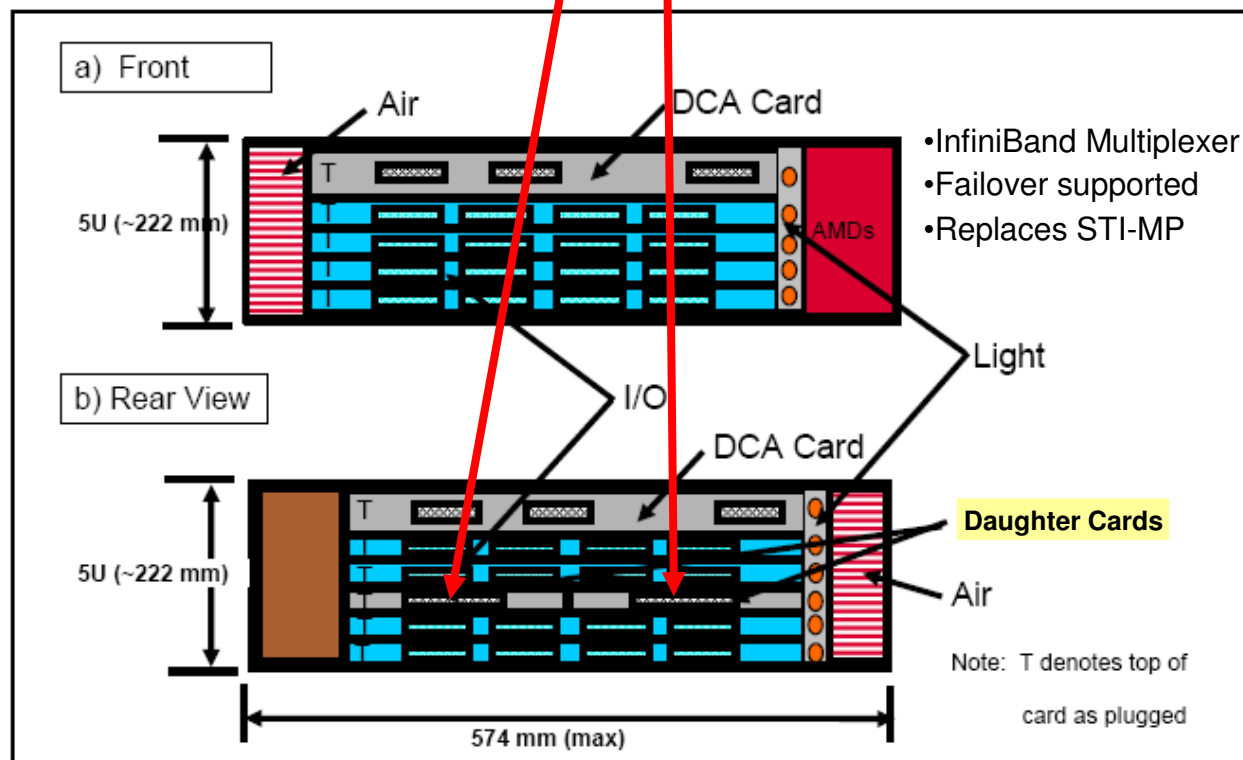
# z10 BC I/O Drawer Structure

- I/O expansion via HCA2-C 12x IB-DDR Copper 2 port
- The number of drawers is determined by the total I/O card count
- The number of drawers could lower the total number of PSIFB Links and/or ICB-4's





HCA2-C (copper)

- 0 I/O cards = 0 drawers
  - PSIFB Links/ICB-4
- 01-08 I/O cards = 1 drawer
- 09-16 I/O cards = 2 drawers
- 17-24 I/O cards = 3 drawers
- 24-32 I/O cards = 4 drawers





# z10 BC FANOUT Cards – Summary



**CPC Drawer - FRONT**

Description of Fanout Cards	F/C	Ports	Available	Comments
HCA2-C 12x IB-DDR	0162	2	z10	To I/O Drawers
HCA2-O 12x IB-DDR	0163	2	z10	Coupling (150 meters)
MBA	0164	2	z10	ICB-4 Coupling
HCA2-O LR 1x IB-DDR	0168	2	z10	Coupling (10 KM Unrepeated)
HCA1-O 12x IB-SDR	0167	2	z9 only	Required on z9 for PSIFB to z10

# z10 BC Channel Type and Crypto Overview

## ▪ FICON/FCP

- FICON Express4
- FICON Express2 (carry forward on upgrade)
- FICON Express (carry forward on upgrade for FCV)

## ▪ Networking

- OSA-Express3
  - 10 Gigabit Ethernet LR and SR
  - Gigabit Ethernet LX and SX
  - 1000BASE-T Ethernet
- OSA-Express2
  - 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 (Define only)

## ▪ ESCON

## ▪ STP

## ▪ Coupling Links

- InfiniBand Coupling Links
  - 12x IB-SDR, 12x IB-DDR
  - 1x IB-SDR, 1x IB-DDR
- ISC-3 (Peer mode only)
- ICB-4 (except z10 EC E64 and z10 BC E10 with non-raised floor feature)
- IC (Define only)

## ▪ Crypto

- Crypto Express2
  - Configurable Coprocessor or Accelerator


## ▪ 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)

Note: ICB-4 cables are available as features. All other cables are sourced separately

# z10 CP Assist for Cryptographic Functions (CPACF)

Integrated Cryptographic Service Facility (ICSF)



Crypto Express2

PU PU PU PU PU

CP Assist for Cryptographic Function

z9



DES  
TDES  
AES-128  
SHA-1, 256  
PRNG



z10




DES  
TDES  
AES-128, 192, 256  
SHA-1, SHA-2:  
(224, 256, 384, 512)  
PRNG

*High performance clear key symmetric encryption/decryption*


# z10 Cryptographic Support

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



**Crypto Express2-1P**

**z10 BC only**



**Crypto Express2**

**z10 BC and z10 EC**

# System z – RAS Design Focus

## ▪ High Availability (HA)

- The attribute of a system designed to provide service during defined periods, at acceptable or agreed upon levels and masks UNPLANNED OUTAGES from end-users. It employs Fault Tolerance; Automated Failure Detection, Recovery, Bypass Reconfiguration, Testing, Problem and Change Management.

## ▪ Continuous Operations (CO)

- Attribute of a system designed to continuously operate and mask PLANNED OUTAGES from end-users. It employs non-disruptive hardware and software changes, non-disruptive configuration, software coexistence.

## ▪ Continuous Availability (CA)

- Attribute of a system designed to deliver non-disruptive service to the end user 7 days a week, 24 HOURS A DAY (there are no planned or unplanned outages). It includes the ability to recover from a site disaster by switching computing to a second site.



# z10 BC Enhancements designed to avoid Outages

- **Continued Focus on Firmware Quality**
- **Memory Subsystem Improvements**



- **DIMM FRU indicators**
- **Single Processor Core Checkstop**
- **Single Processor Core Sparing**
- **Redundant 100Mb Ethernet service network w/ VLAN**

## *NEW for z10 BC*

- **Concurrent add/replaceable I/O drawer**
- **Plan ahead memory**
- **Serviceability enhancements of SAN for both FICON and FCP**

- **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**
  - 8 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 BC Field Replaceable SCMs



- The SCM is the Field Replaceable Unit (FRU)
- z10 BC has no guaranteed IBM spare which is consistent with z9 BC
- The SCM will be replaced only if a PU core has failed and a spare is not available (failed into customer PUs)
  - exception being one-way systems where the SCM will be replaced when there are no available PUs left.
  - exception being customer with CIU/CPE/OOCoD where the SCM will be replaced on every PU failure.



# z10 RAS Summary

- **z10 is designed to deliver the industry leading RAS customers expect from System z servers**
- **z10 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**



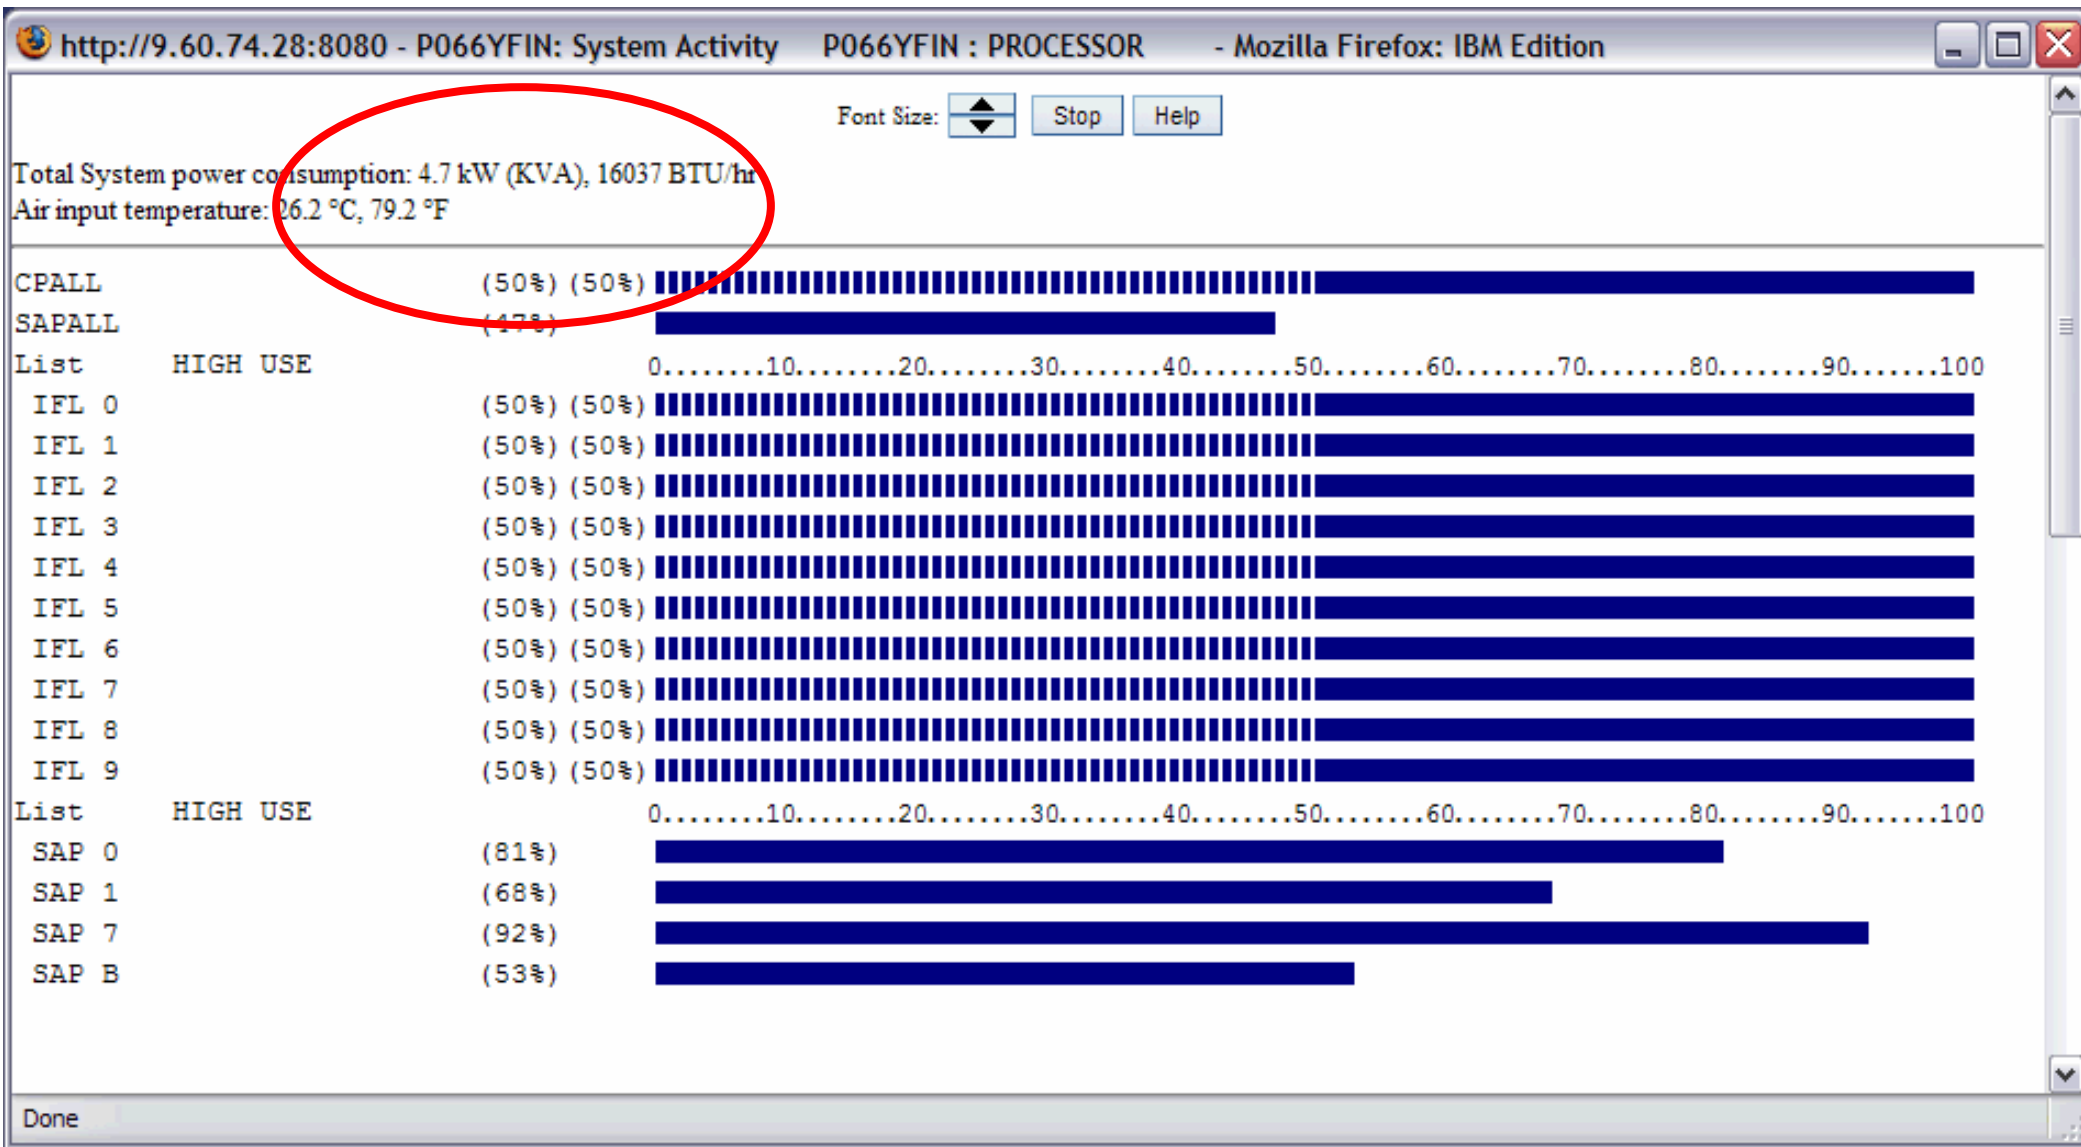
# z10 CoD Offerings

- **On-line Permanent Upgrade**
  - Permanent upgrade performed by customer (previously referred to Customer Initiated Upgrade - CIU)
- **Capacity Backup (CBU)**
  - For disaster recovery
  - Concurrently add CPs, IFLs, ICFs, zAAPs, zIIPs, SAPs
  - Pre-paid
- **Capacity for Planned Event (CPE)**
  - To replace capacity for short term lost within the enterprise due to a planned event such as a facility upgrade or system relocation
  - Predefined capacity for a fixed period of time (3 days)
  - Pre-paid
- **On/Off Capacity on Demand (On/Off CoD)**
  - Production Capacity
  - Supported through software offering – Capacity Provisioning Manager (CPM)
  - Payment:
    - Post-paid or Pre-paid by purchase of capacity tokens
    - Post-paid with unlimited capacity usage
    - On/Off CoD records and capacity tokens configured on Resource Link
- **Customer Initiated Upgrade (CIU)**
  - Process/tool for ordering temporary and permanent upgrades via Resource Link
  - Permanent upgrade support:
    - Un-assignment of currently active capacity
    - Reactivation of unassigned capacity
    - Purchase of all PU types physically available but not already characterized
    - Purchase of installed but not owned memory

# z10 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**

# System Activity Display with Power Monitoring



# z10 BC Highlights and Physical Dimensions

## ▪ Highlights

- Slight increase in size. Height Reduction feature may be required for doors with openings less than 202 cm / 79.5 in.
- Single or Three Phase power – same power cable as z9 BC
  - Option for Zero Halogen power cables
- Non- Raised floor option

	<b>z9 BC</b>	<b>z10 BC</b>
<b># of Frames</b>	1 Frame	1 Frame
<b>Height (with covers)</b>	194.1 cm / 76.4 in (40 EIA)	201.5 cm / 79.3 in (42 EIA)
<b>Width (with covers)</b>	78.5 cm / 30.9 in	77.0 cm / 30.3 in
<b>Depth (with covers)</b>	157.7 cm / 62.1 in	180.6 cm / 71.1 in
<b>Height Reduction</b>	178.5 cm / 70.3 in (EIA)	180.9 cm / 71.2 in (EIA)
<b>Width Reduction</b>	None	None
<b>Machine Area</b>	1.24 Sq. Meters / 13.31 Sq. Feet	1.42 Sq. Meters / 15.22 Sq. Feet
<b>Service Clearance</b>	3.03 Sq. Meters / 32.61 Sq. Feet (IBF Contained within the Frame)	3.50 Sq. Meters / 37.62 Sq. Feet (IBF Contained within the Frame)

Always refer to the z10 BC IMPP (GC28-6875) for detailed planning information

# z10 BC System Power

- **z10 BC maximum configuration calculated AC input power (Statistical Maximum)**
  - All systems should draw less power than this
  - Typical systems will draw less power than this

	1 I/O Drawer	2 I/O Drawers	3 I/O Drawers	4 I/O Drawers
<b>normal room (&lt;28 degC)</b>	<b>3.686 kW</b>	<b>4.542 kW</b>	<b>5.308 kW</b>	<b>6.253 kW</b>
<b>warm room (&gt;=28 degC)</b>	<b>4.339 kW</b>	<b>5.315 kW</b>	<b>6.291 kW</b>	<b>7.266 kW</b>

- **30 Amp plug capacity (208 VAC)**
  - 5.5 kW single phase or unbalanced 3 phase
    - Supports up to 2 I/O drawers
  - 8.9 kW balanced 3 phase
    - Supports all system configurations – have balanced 3 phase feature
    - Plug 2 additional BPR's per side

Always refer to the z10 BC IMPP (GC28-6875) for detailed planning information

# System z10 BC Operating System Support

Operating System	ESA/390 (31-bit)	z/Architecture (64-bit)
z/OS Version 1 Releases 8, 9 and 10	No	Yes
z/OS Version 1 Releases 7 <sup>(1)(2)</sup> with IBM Lifecycle Extension for z/OS V1.7	No	Yes
Linux on System z <sup>(2)</sup> , Red Hat RHEL 4, & Novell SUSE SLES 9	Yes	Yes
Linux on System z <sup>(2)</sup> , Red Hat RHEL 5, & Novell SUSE SLES 10	No	Yes
z/VM Version 5 Release 2 <sup>(3)</sup> and 3 <sup>(3)</sup> and 4	No*	Yes
z/VSE Version 3 Release 1 <sup>(2)(4)</sup>	Yes	No
z/VSE Version 4 Release 1 <sup>(2)(5)</sup> and 2 <sup>(5)</sup>	No	Yes
z/TPF Version 1 Release 1	No	Yes
TPF Version 4 Release 1 (ESA mode only)	Yes	No

- z/OS V1.7 support on the z10 BC requires the Lifecycle Extension for z/OS V1.7, 5637-A01. The Lifecycle Extension for z/OS R1.7 + zIIP Web Deliverable required for z10 to enable HiperDispatch on 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.
- Compatibility Support for listed releases. Compatibility support allows OS to IPL and operate on z10 BC
- Requires Compatibility Support which allows z/VM to IPL and operate on the z10 providing System z9 functionality for the base OS and Guests. \*z/VM supports 31-bit and 64-bit guests
- 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 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 2098DEVICE Preventive Planning (PSP) bucket prior to installing a z10 BC**

Thank

you