IBM TRAINING

Session G02

IBM System z9 Business Class: Processor, Memory and System Structure

Harv Emery, IBM Americas System z Advanced Technical Support

IBM
SYSTEM z9 AND zSERIES EXPO
October 9 - 13, 2006

Orlando, FL

Trademarks

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

APPN*
CICS*
DB2*
DB2 Connect
DirMaint
e-business logo*
ECKD
Enterprise Storage Server*
ESCON*
FICON*
GDPS*
Geographically Dispersed Parallel Sysplex
HiperSockets
HyperSwap
IBM*
IBM eServer
IBM e(logo)server*
IBM logo*
IMS
Language Environment*
MQSeries*
Multiprise*
NetView*
On demand business logo
OS/390*
Parallel Sysplex*
PR/SM
Processor Resource/Systems Manager
RACF*
Resource Link
RMF
S/390*
Sysplex Timer*
System z9
TotalStorage*
Virtualization Engine
VM/ESA*
VSE/ESA
VTAM*
WebSphere*
z/Architecture
z/OS*
z/VM*
z/VSE
zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

- **IBM System z™ Introduction**
- **System Structure**
- **Model Configurations and Memory**
- **Enhanced Driver Maintenance**
- **Processor Unit Specifics**
- **“On Demand”**
- **Cryptography and Security**

IBM System z9 and zSeries Introduction

IBM System z family

IBM eServer zSeries 990 – z990 (2084)

- Announced 5/03 – first zSeries Superscalar Server with up to 48 PUs
- 4 models – Up to 32-way
- Specialty Engines
  - **CP, IFL, ICF, zAAP**
- On Demand Capabilities
  - **CUoD, CIU, CBU, On/Off CoD**
- Memory – up to 256 GB
- Channels
  - **Four LCSSs**
  - **Up to 1024 ESCON channels**
  - **Up to 240 FICON Express2 channels**
  - **Token-Ring, GbE, 1000BASE-T Ethernet**
  - **Coupling Links**
- Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 30 logical partitions
- Operating Systems
  - **z/OS, z/VM, VSE/ESA, z/VSE, TPF, z/TPF, Linux on zSeries**

IBM eServer zSeries 890 – z890 (2086)

- Announced 4/04 – zSeries Superscalar Server with 5 PUs
- 1 model – Up to 4-way
  - **28 capacity settings**
- Specialty Engines
  - **CP, IFL, ICF, zAAP**
- On Demand Capabilities
  - **CUoD, CIU, CBU, On/Off CoD**
- Memory – up to 32 GB
- Channel
  - **Two LCSSs**
  - **Up to 420 ESCON channels**
  - **Up to 80 FICON Express2 channels**
  - **Networking Adapters (OSA)**
  - **Coupling Links**
- Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 30 logical partitions (Setting 110 – 15 logical partitions)
- Operating Systems
  - **z/OS, z/VM, VSE/ESA, z/VSE, TPF, z/TPF, Linux on zSeries**

IBM System z9 EC (z9 EC) (2094)

- Announced 7/05 as z9-109 - Superscalar Server with up to 64 PUs
- 5 models – Up to 54-way
- Granular Offerings for up to 8 CPs
- Specialty Engines
  - **CP, 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**
  - **Enhanced FICON Express2 and 4**
  - **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, VSE/ESA, z/VSE, TPF, z/TPF, Linux on System z9**

IBM System z9 BC (z9 BC) (2096)

- Announced 4/06 - Superscalar Server with 8 PUs
- 2 models – Up to 4-way
- High levels of Granularity available
  - **73 Capacity Indicators**
- Specialty Engines
  - **CP, 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**
  - **Enhanced FICON Express2 4 Gbps**
  - **10 GbE, GbE, 1000BASE-T**
  - **Coupling Links**
- Configurable Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 30 logical partitions (Model R07 – 15)
- Enhanced Availability
- Operating Systems
  - **z/OS, z/VM, VSE/ESA, z/VSE, TPF, z/TPF, Linux on System z9**

z9 BC Overview

- **Based on System z9 Enterprise Class (z9 EC) technology**
- **High levels of granularity in 2 new models**
- **More engines for more workloads**
  - System z9 Integrated Information Processor (zIIP), System z Application Assist Processor (zAAP), Integrated Facility for Linux (IFL), Internal Coupling Facility (ICF)
- **On demand upgrade capability**
  - Upgradeability
  - On/Off Capacity on Demand (CoD) functions available
- **Enhanced networking and connectivity options**
  - FICON Express4
  - Lower FICON® entry cost now available with new 2-port FICON card
- **Built with System z9's cryptographic and encryption functions**
- **EWLC and Tiered EWLC Software Pricing Structure**
- **Operating system support – similar to z9 EC**
  - z/OS.e will continue to be supported

z9 BC Functions and Features (May 2006)

- Two hardware models with high levels of granularity available – 73 CIs
- Faster Uni Processor
- Up to 7 customer PUs
- Up to 64 GB memory
- Up to 30 LPARs
- Separate PU pool management
- PU Conversions for zAAPs and zIIPs
- CBU for IFL, ICF, zAAP and zIIP
- CBU and On/Off CoD Enhancements
- Redundant I/O interconnect
- Dynamic oscillator switchover
- 54 additional hardware Instructions
- Enhanced Driver Maintenance

- Preview*
  Server Time Protocol

- Hot pluggable/ maintainable MBA/STI fanout cards
- Up to 16 2.7 GB STIs per system
- MIDAW facility
- Multiple Subchannel Sets per LCSS
- 63.75K Subchannels for Set-0
- Increased Number of FICON Express2 and 4 Features
- N_Port ID Virtualization
- IPv6 Support for HiperSockets
- OSA-Express2 1000BASE-T
- FICON and OSA Enhancements
- Crypto Enhancements
- Configurable Crypto Express2
- Availability in EU after implementation of RoHS

*This statement represents IBM's current intentions. IBM development plans are subject to change or withdrawal without further notice.

Note: Please refer to the latest PSP bucket for latest PTFs for new functions/features

z9 BC Model R07

- **Machine Type**
  - 2096
- **Model**
  - R07
- **Processor Units (PUs)**
  - 8 PUs per System
  - 1 SAP per book, standard
  - No dedicated spares
  - 7 PUs available for characterization
    - **1 to 3 Central Processors (CPs)**
    - Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z Application Assist Processors (zAAPs), System z9 Integrated Information Processors (zIIPs), optional System Assist Processors (SAPs)
  - **Up to 15 LPARs**
- **Memory**
  - Minimum of 8 GB
  - Up to 64 GB per System
    - 8 GB increments
- **I/O**
  - Up to 16 STIs per System @ 2.7 GB/s each
  - Total system I/O bandwidth capability of 43.2 GB
  - Up to 2 Logical Channel Subsystem (LCSS)
  - **Up to a maximum of 4 I/O Domains**
    - **Up to 240 channels** – dependent on Channel types

z9 BC Model S07

- **Machine Type**
  - 2097
- **Model**
  - S07
- **Processor Units (PUs)**
  - 8 PUs per System
  - 1 SAP per book, standard
  - No dedicated spares
  - 7 PUs available for characterization
    - **0 to 4 Central Processors (CPs)**
    - Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z Application Assist Processors (zAAPs), System z9 Integrated Information Processor (zIIP), optional System Assist Processors (SAPs)
    - Can have an IFL-only System
  - Up to **30 LPARs**
- **Memory**
  - Minimum of 8 GB
  - Up to 64 GB per System
    - 8 GB increments
- **I/O**
  - Up to 16 STIs per System @ 2.7 GB/s each
  - Total system I/O bandwidth capability of 43.2 GB
  - Up to 2 Logical Channel Subsystems (LCSSs)
  - **Up to a maximum of 7 I/O Domains**
    - **Up to 420 channels** – dependent on Channel types

Upgrade Paths

- **Full upgrades within the z9 (R07 to S07 to z9 EC)**
- **Any to any upgrade from the z890**
- **Upgrade from the z800 model 004**
- **No charge MES upgrades on IFLs and zAAPs**
- **Capability of the System z9 servers to nondisruptively increase computing resources within the server**
  - 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

z9 BC – Under the covers

A Frame

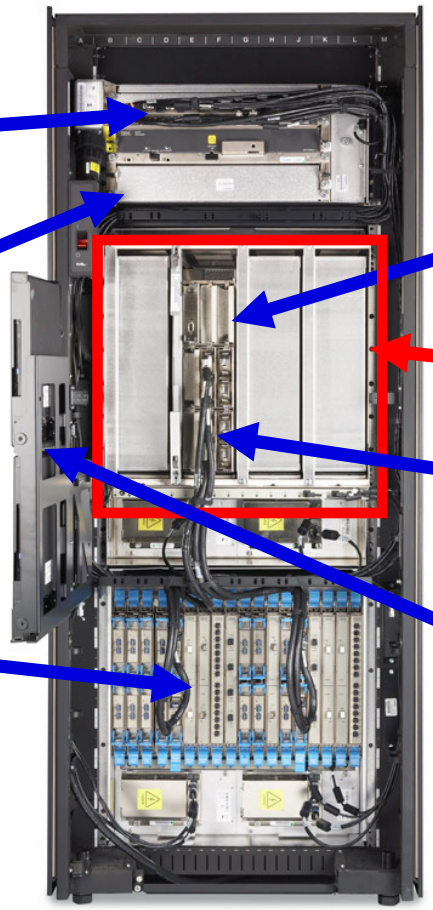

Internal Batteries (optional)

Power Supplies

I/O Cage

Single Processor Book and Memory

CEC Cage

STI Connectors

Support Elements (gate with Laptops swung open)

Front View

Fiber Quick Connect Feature (optional)

z9 BC Improved granularity and scalability

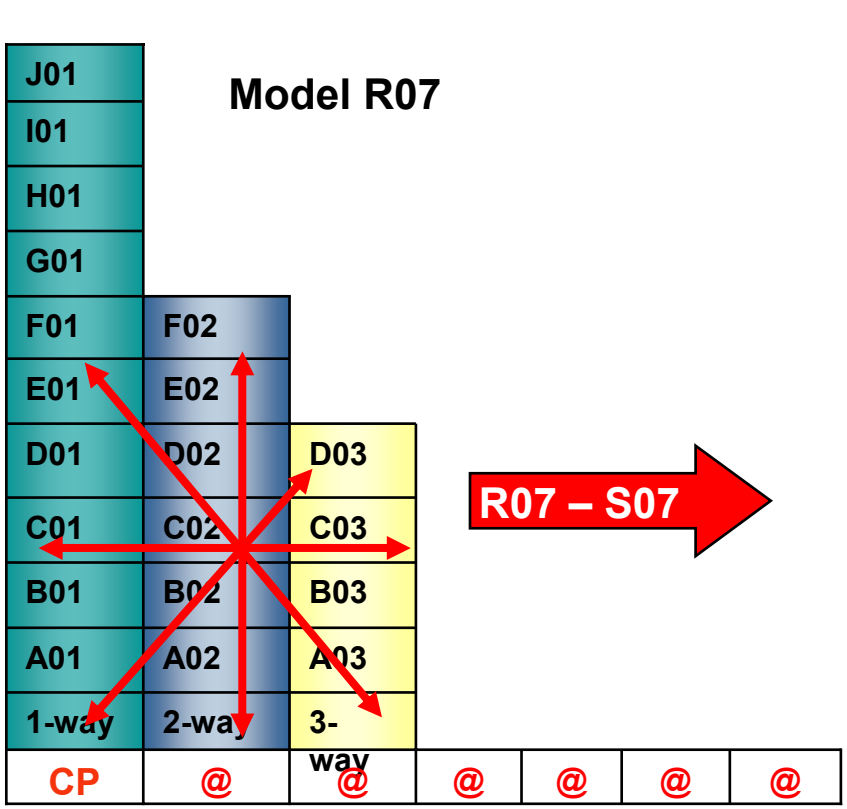

R07 – S07

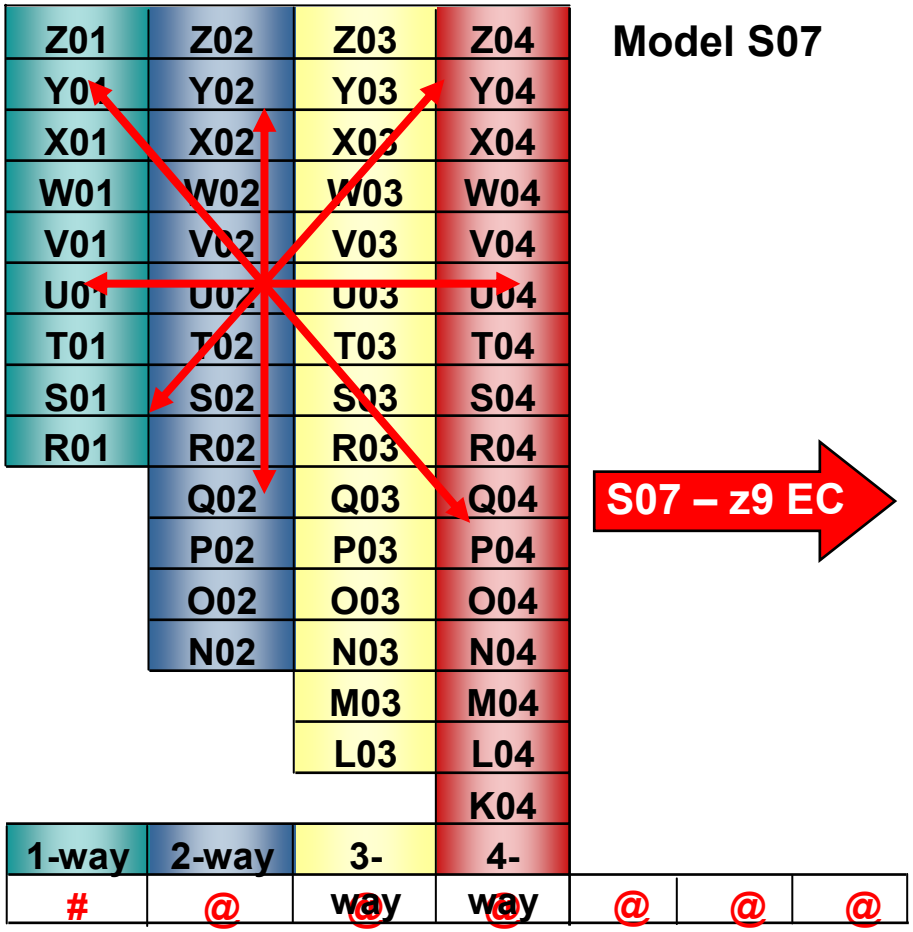

S07 – z9 EC

Full on demand upgradeability in the family

- Model R07 must have minimum 1 CP engine
- Model S07 may be a full IFL or ICF system
- Model R07 upgradeable to model S07
- Model S07 upgradeable to z9 EC Model S08

 # = CP or IFL or ICF

@ = Any Specialty Engines. zAAPs and zIIPs have T & Cs

Processor Granularity z9 BC R07

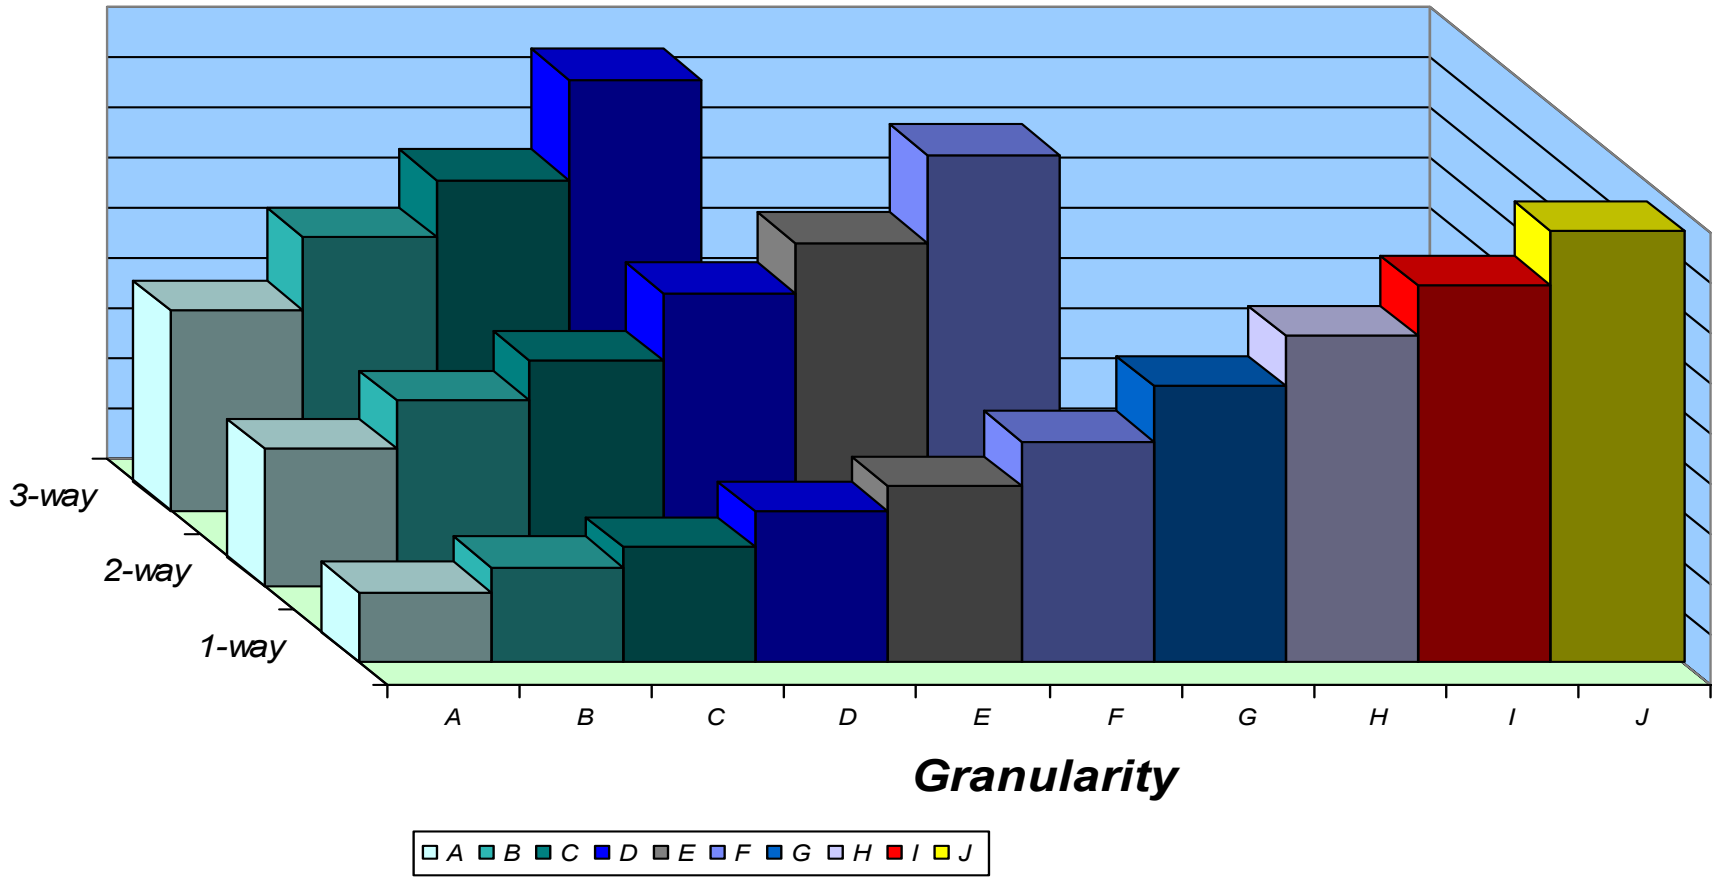

Processor Granularity z9 BC S07

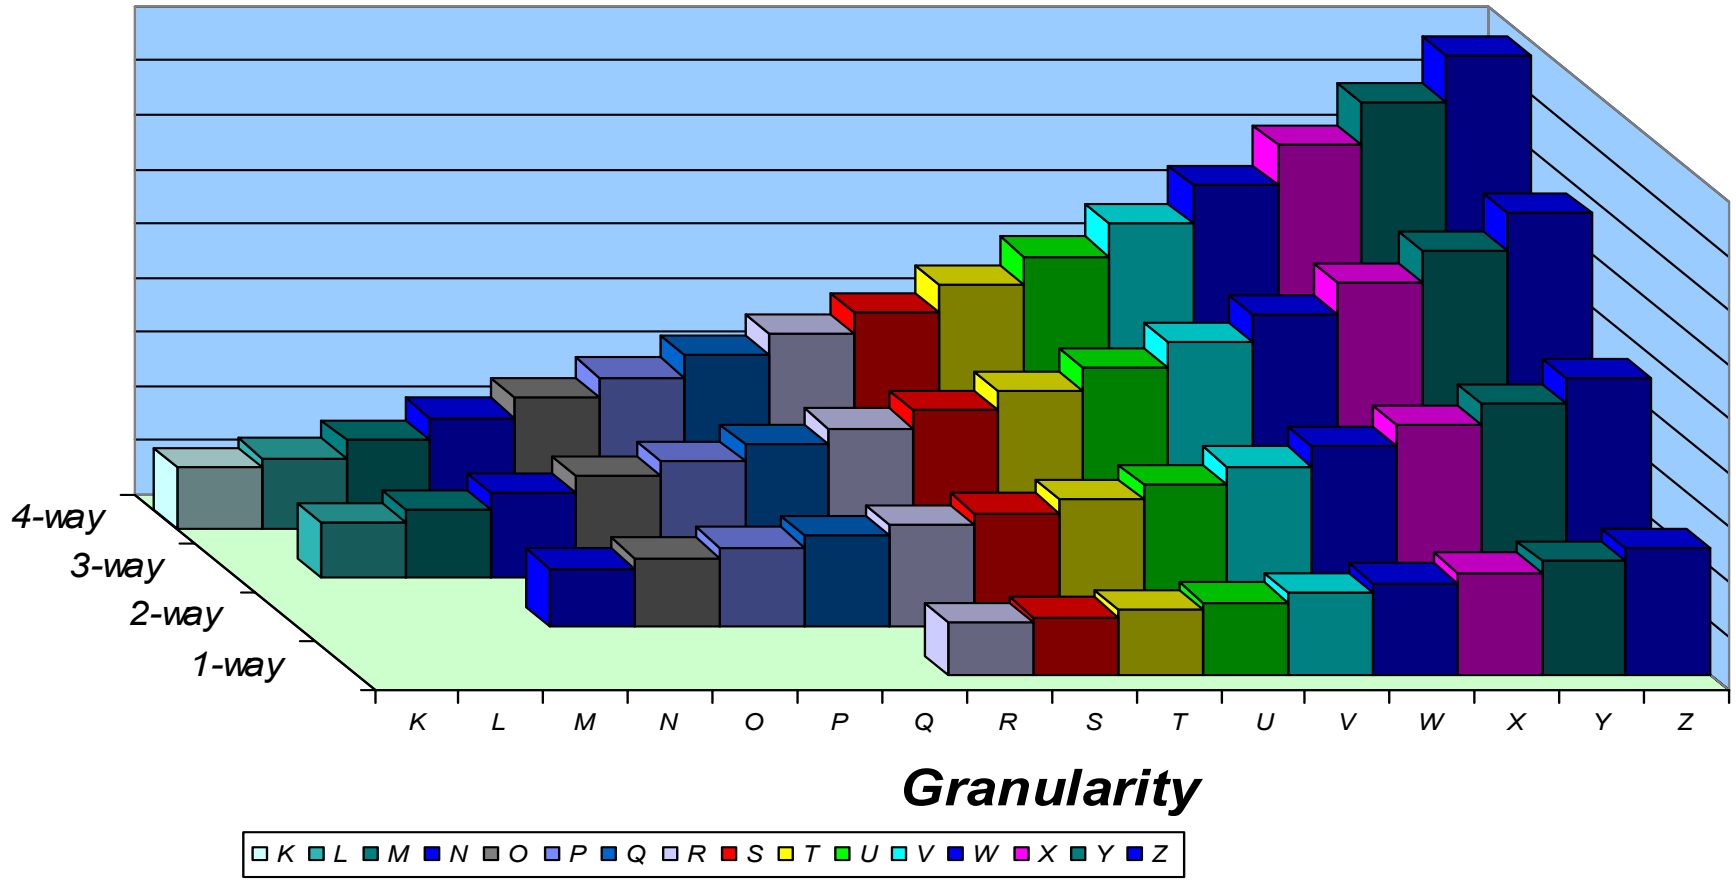

z9 BC – Delivering increased capacity and performance

- **Greater granularity and scalability**
  - Two models with one machine type (2096)
    - 1 to 4-way high performance server standard engines
    - Entry model with 1 to 3-way standard engines
    - Up to a 7-way with specialty engines
  - 73 capacity settings for a 260% increase in flexibility over z890
  - Delivers up to 34% more capacity with the same low entry point as the z890
  - Up to 34% hardware performance improvement for Linux (IFLs), Java (zAAPs) and Internal Coupling Facilities (ICFs)
  - New zIIP for data serving workloads
  - Double the memory – up to 64 GB per server
- **Improved I/O Performance**
  - 40% more FICON channels – up to 112
  - Up to 170% more bandwidth than z890
  - Can improve FICON performance with Modified Indirect Data Address Word (MIDAW) facility
  - Double the FICON concurrent I/O operations from 32 to 64 on FICON channel
  - Multiple Subchannel Sets (MSS) for an increased number of logical volumes

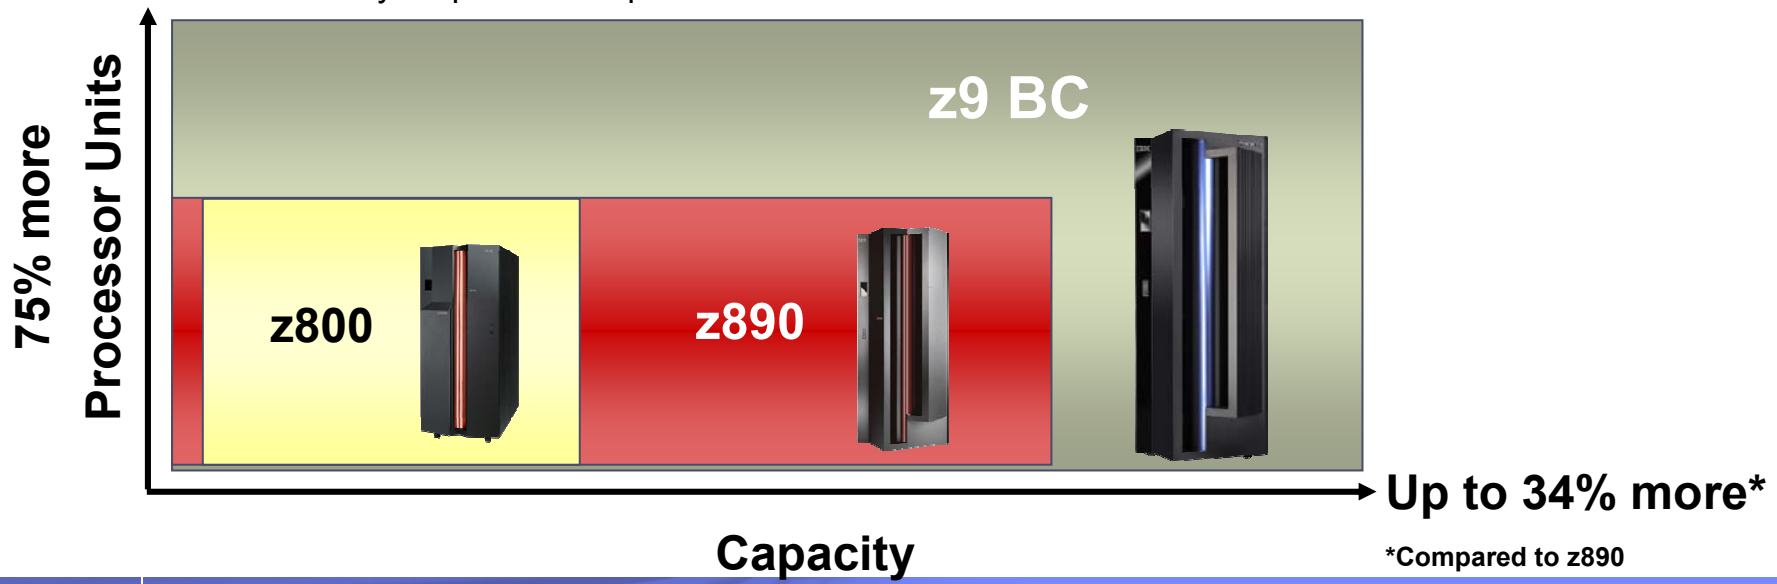

Overlapping Capacity

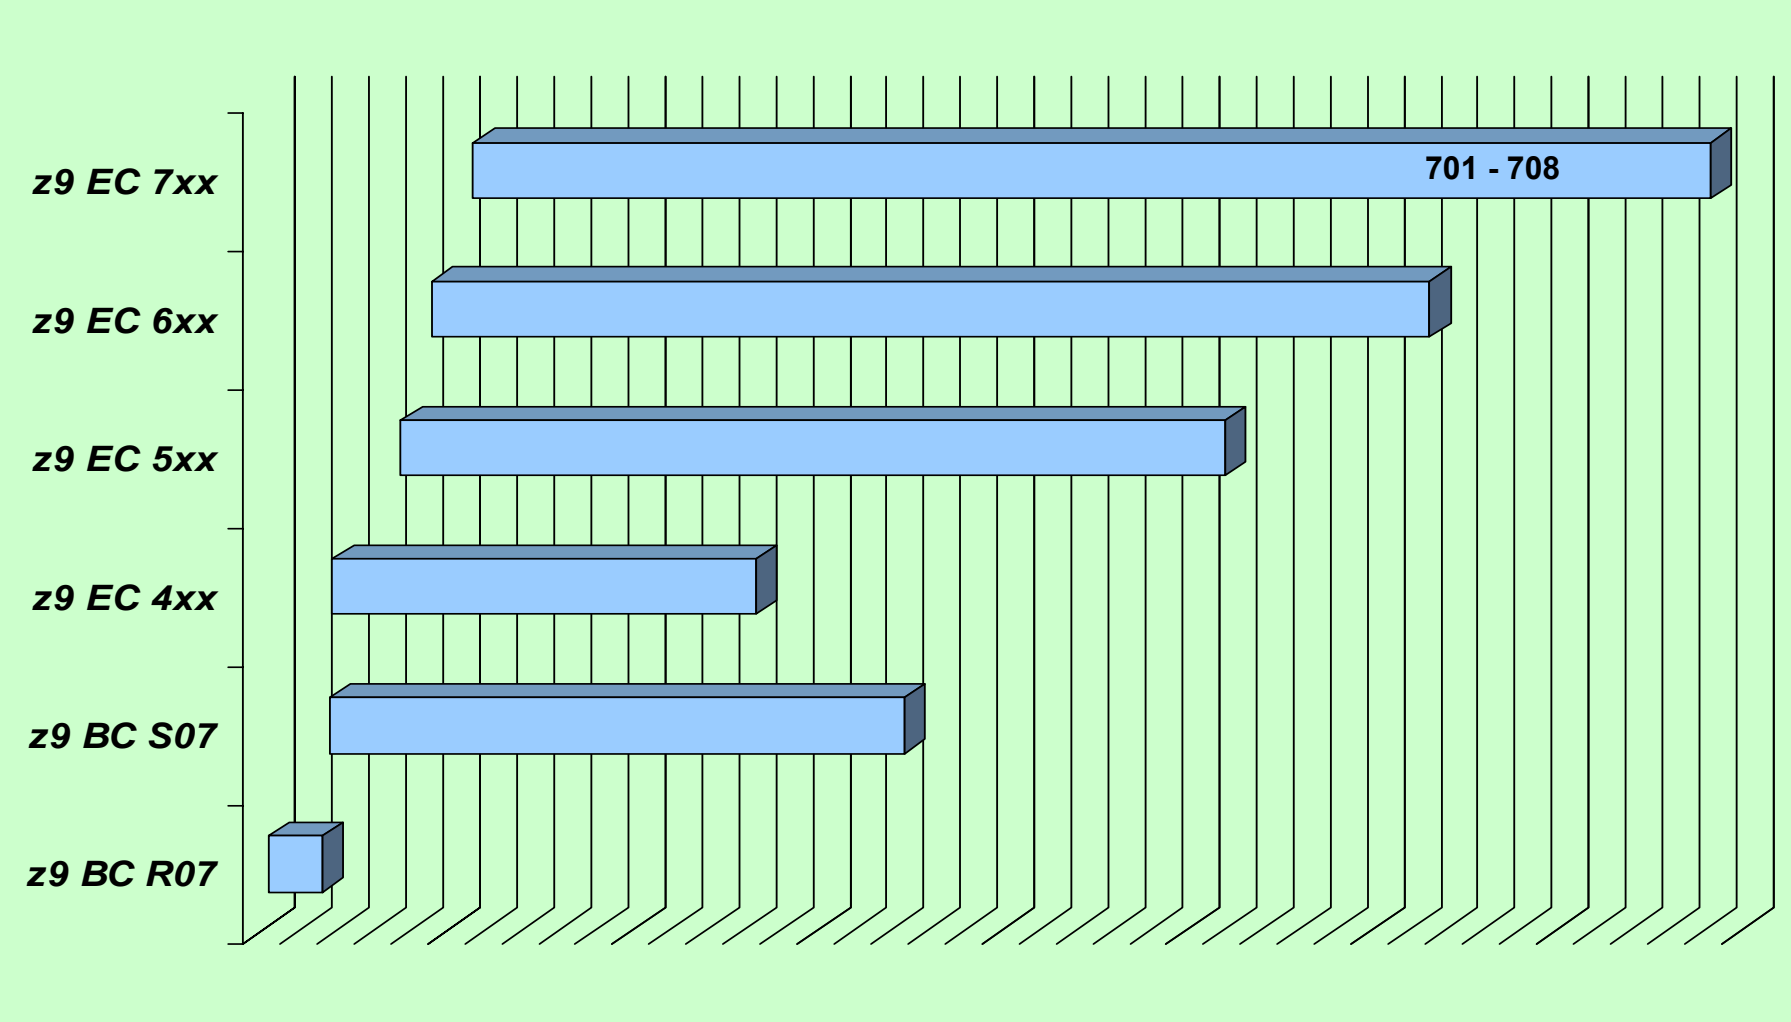

System z9 Capacity Planning in a nutshell

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

System z9 has a 10% MSU "technology dividend" compared to z990/z890
System z9 has a 19% MSU "technology dividend" compared to z900/z800

IBM System z9 BC System Structure

z9 BC Processor Book Layout

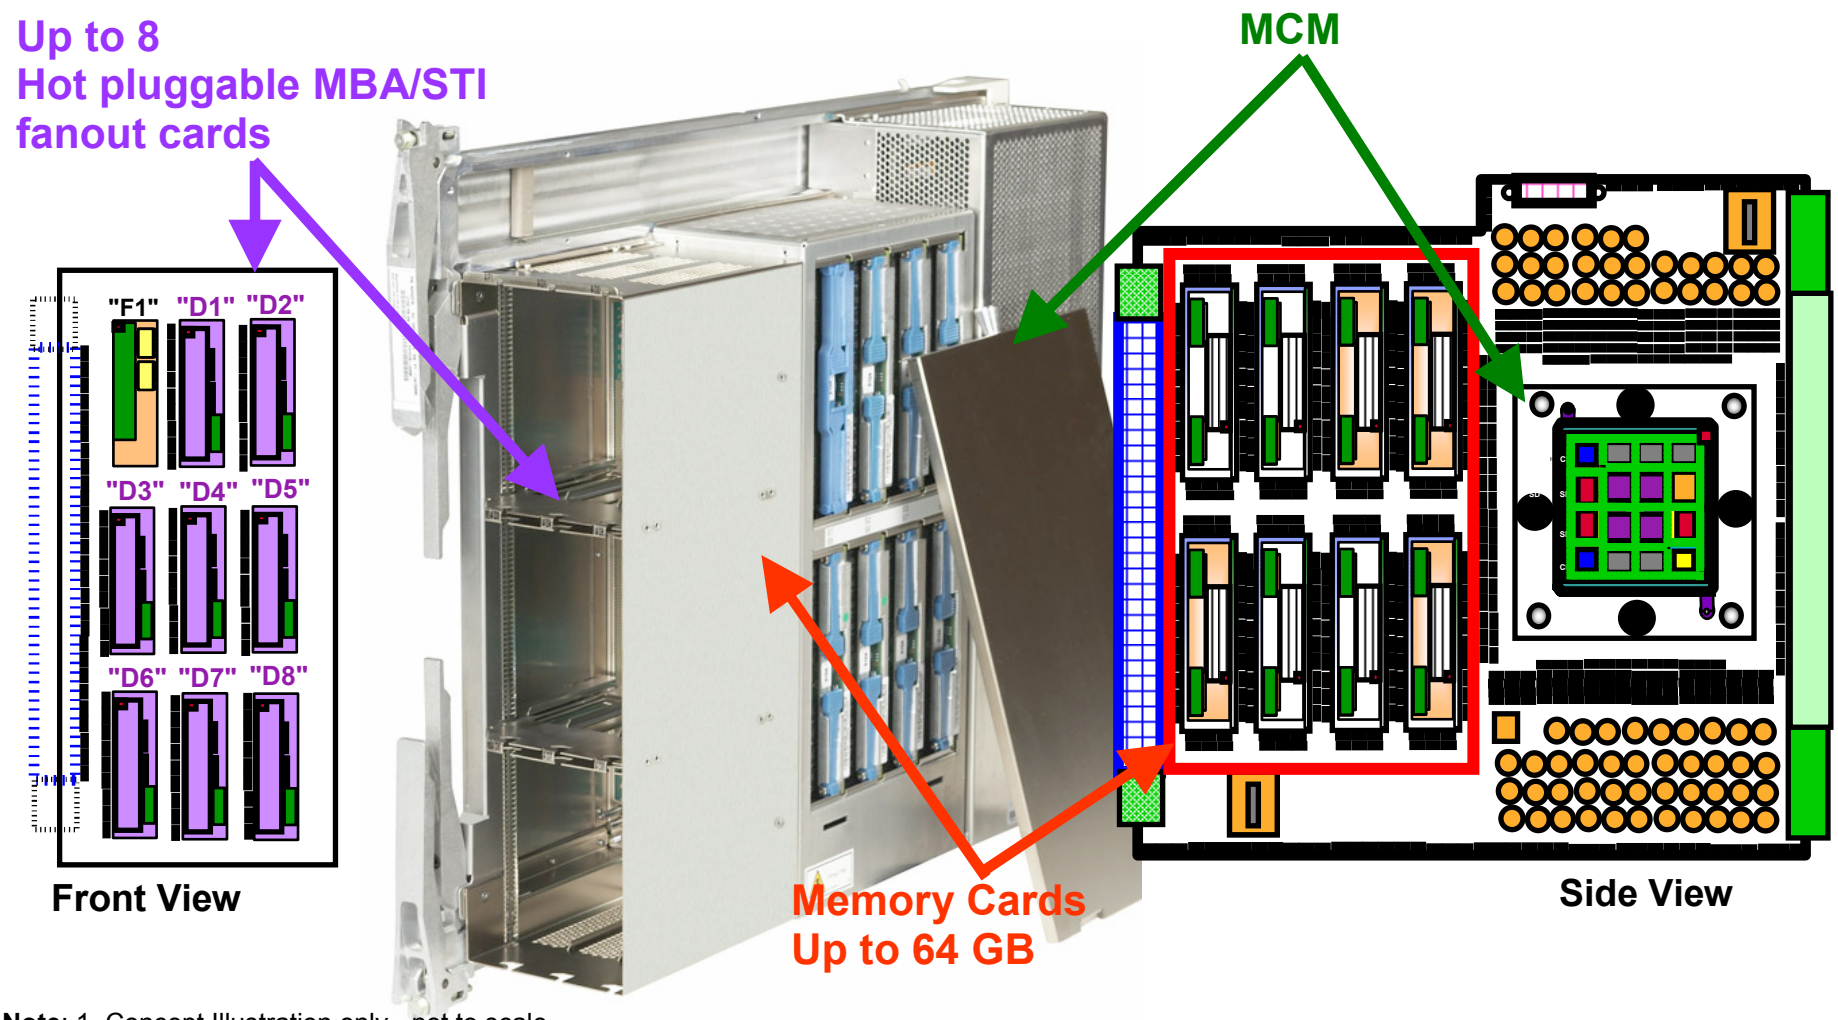

Note: 1. Concept Illustration only - not to scale
2. 4 or 8 pluggable Memory Cards
3. Each MBA fanout card is hot-pluggable and has 2 STIs

z9 BC 8-way MCM

- Advanced 95mm x 95mm MCM
  - 102 Glass Ceramic layers
  - 16 chip sites, 217 capacitors
  - 0.545 km of internal wire

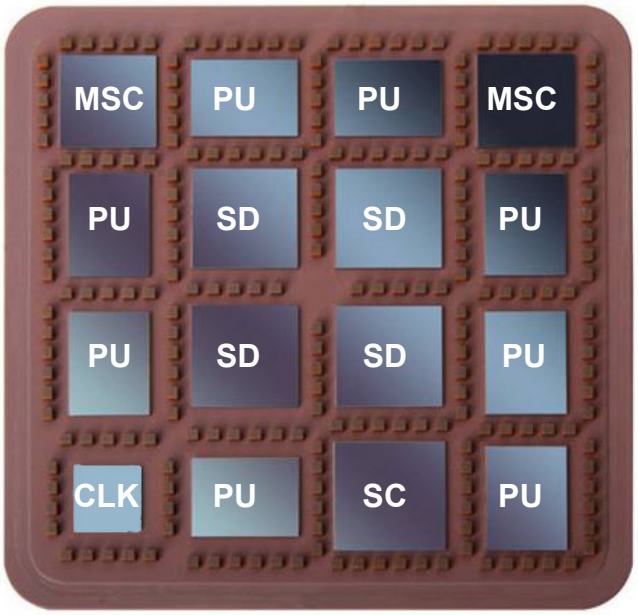

- CMOS 10Ks0 chip Technology
  - PU, SC, SD and MSC chips
  - Copper interconnections, 10 copper layers
  - 8 PU chips/MCM
    - 15.78 mm x 11.84 mm
    - 121 million transistors/chip
    - L1 cache/PU
      - 256 KB I-cache
      - 256 KB D-cache
    - 0.7 ns Cycle Time
  - 4 System Data (SD) cache chips/MCM
    - 15.66 mm x 15.40mm
    - L2 cache per Book
      - 660 million transistors/chip
      - 40 MB
  - One Storage Control (SC) chip
    - 16.41mm x 16.41mm
    - 162 million transistors
    - L2 cache crosspoint switch
    - L2 access rings to/from other MCMs
  - Two Memory Storage Control (MSC) chips
    - 14.31 mm x 14.31 mm
    - 24 million transistors/chip
    - Memory cards (L3) interface to L2
    - L2 access to/from MBAs (off MCM)
  - One Clock (CLK) chip - CMOS 8S
    - Clock and ETR Receiver

z9 BC 8 PU Logical Structure

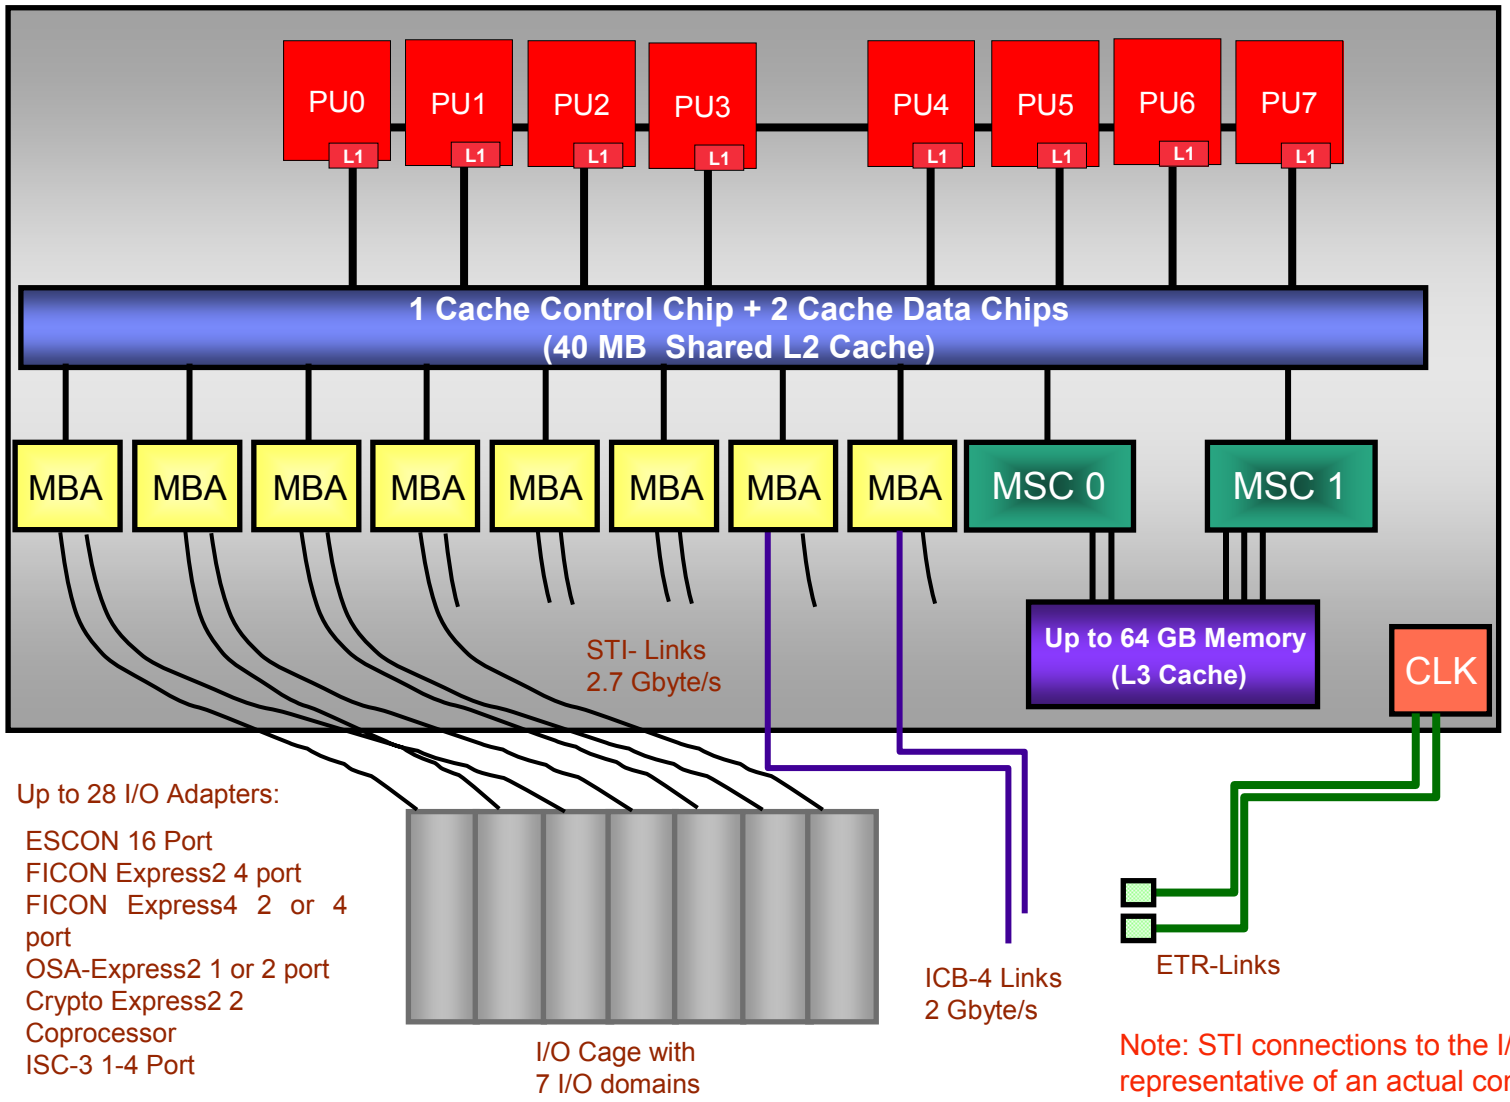

- PU 4 always is SAP, each PU is single core and Crypto Assist
- L2 contains Storage Controller and Storage Data Chips
- Book Package contains PU, Storage Controllers, L2 and L3 Caches, 2.7 GB/sec Self-Timed Interfaces
- Seven of 16 STIs can attach to the I/O cage. Rest can be used for ICB-4s

Note: STI connections to the I/O cage are NOT representative of an actual configuration.

IBM System z9 BC Model Configurations and Memory

z9 BC Model Structure

- One machine type – 2096 – two hardware models, R07 and S07
- Model number indicates PUs available for characterization
  - Single serial number
  - PU characterization is identified by number of features ordered
- One System Assist Processors (SAPs) per System
- z9 BC software models
  - nxx, where n = subcapacity engine size and xx = number of CPs
    - For Model R07 n = A up to J and xx = 1 to 3
    - For Model S07 n = K up to Z and xx = 1 to 4
  - Total 73 Capacity Indicators for software models
    - 20 for Model R07 and 53 for Model S07

Models	MCMs	Available PUs	Max Available Subcapacity CPs	Standard SAPs	Standard Spares	CP/IFL/ ICF/zAAP/zIIP ****	Max Memory	Max Channels
R07*	1	8	3	1	0	3/6/6/3/3	64 GB	240 ***
S07**	1	8	4	1	0	4/7/7/3/3	64 GB	420 ***

Notes:
* Must have a minimum of 1 CP
** Must have a minimum of 1 CP, IFL or ICF
*** 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).

System z9 BC Memory Cards

- **Memory cards 2, 4, or 8 GB**
- **Configurations of 4 or 8 cards per book**
- **Physical memory may be larger than purchased memory**
- **Purchased memory enabled by LIC-CC**
  - 8 GB Increments
  - 8 to 64 GB per system

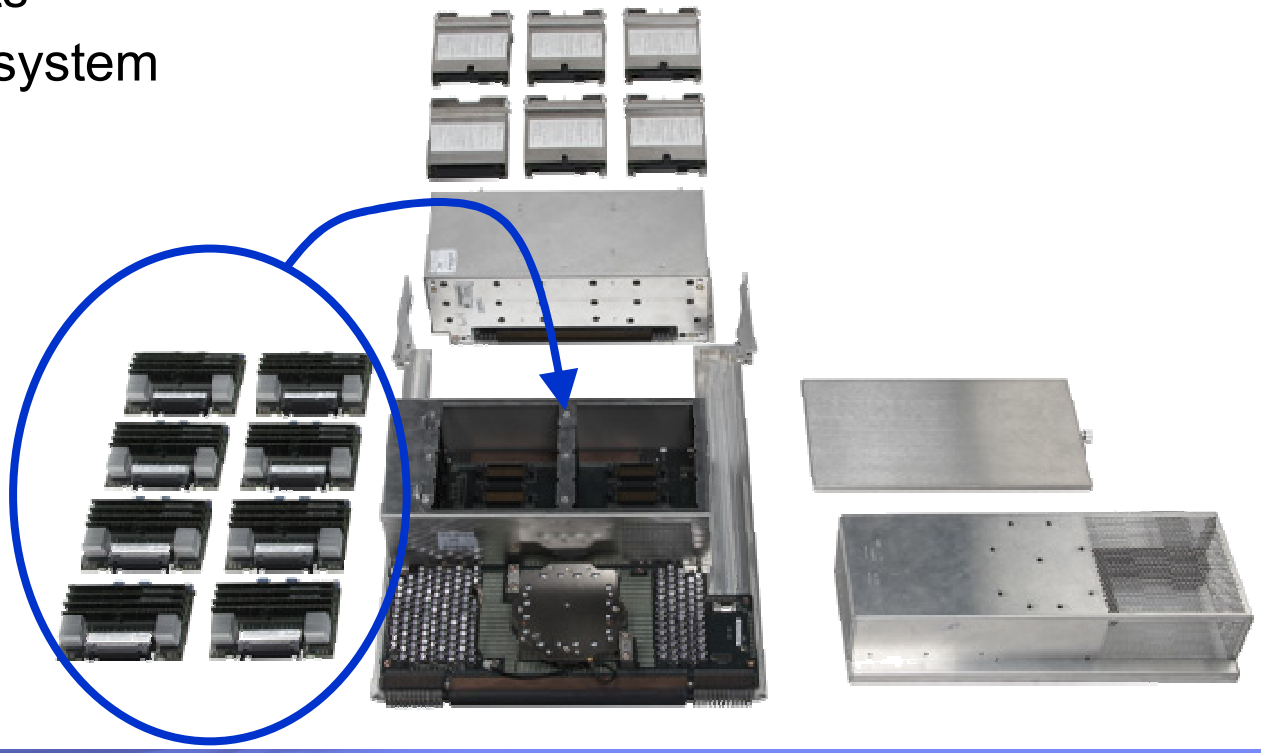

z9 BC Memory Upgrade Options

From	To	To	To	To	To	To	To
8 GB	16 GB	24 GB	32 GB	40 GB	48 GB	56 GB	64 GB
16 GB	-	24 GB	32 GB	40 GB	48 GB	56 GB	64 GB
24 GB	-	-	32 GB	40 GB	48 GB	56 GB	64 GB
32 GB	-	-	-	40 GB	48 GB	56 GB	64 GB
40 GB	-	-	-	-	48 GB	56 GB	64 GB
48 GB	-	-	-	-	-	56 GB	64 GB
56 GB	-	-	-	-	-	-	64 GB
64 GB	-	-	-	-	-	-	-

Red - Disruptive upgrade

Green - Concurrent upgrade

Card Sizes = 2 GB, 4 GB and 8 GB

IBM Enhanced Driver Maintenance

System z9 Enhanced Driver Maintenance

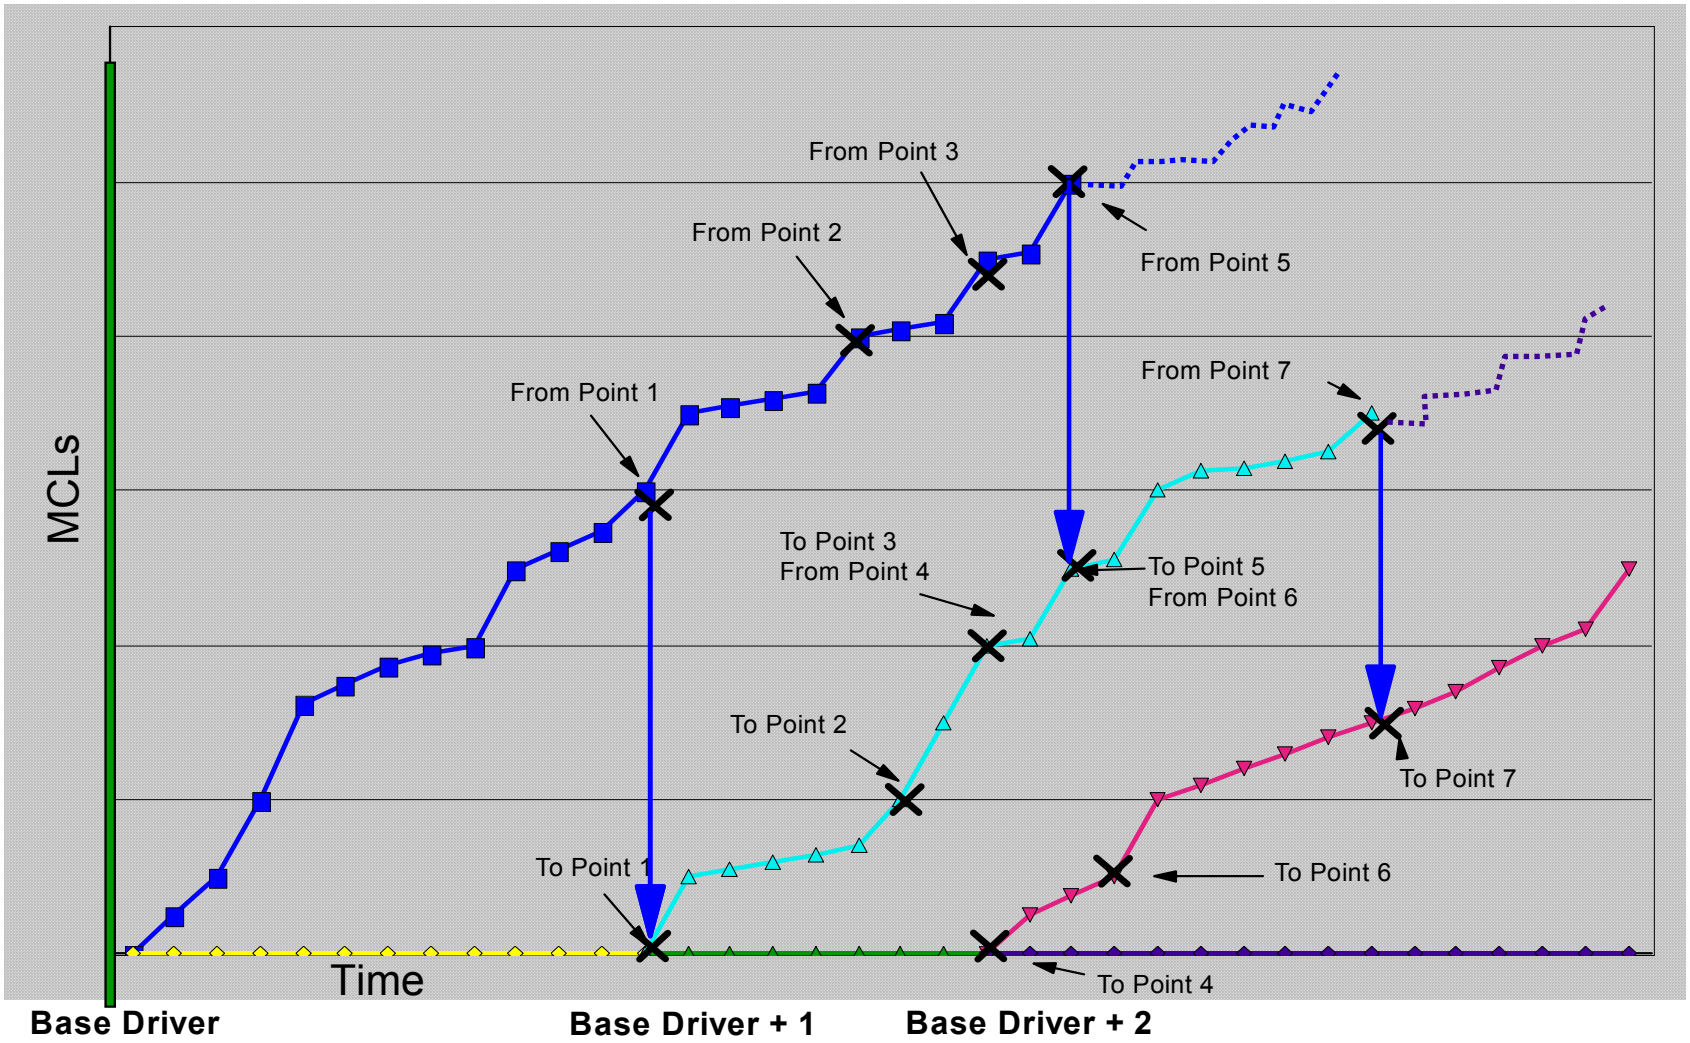

IBM System z9 PU Specifics

Technology evolution with specialty engines

Building on a strong track record of technology innovation with specialty engines, IBM introduces the System z9 Integrated Information Processor

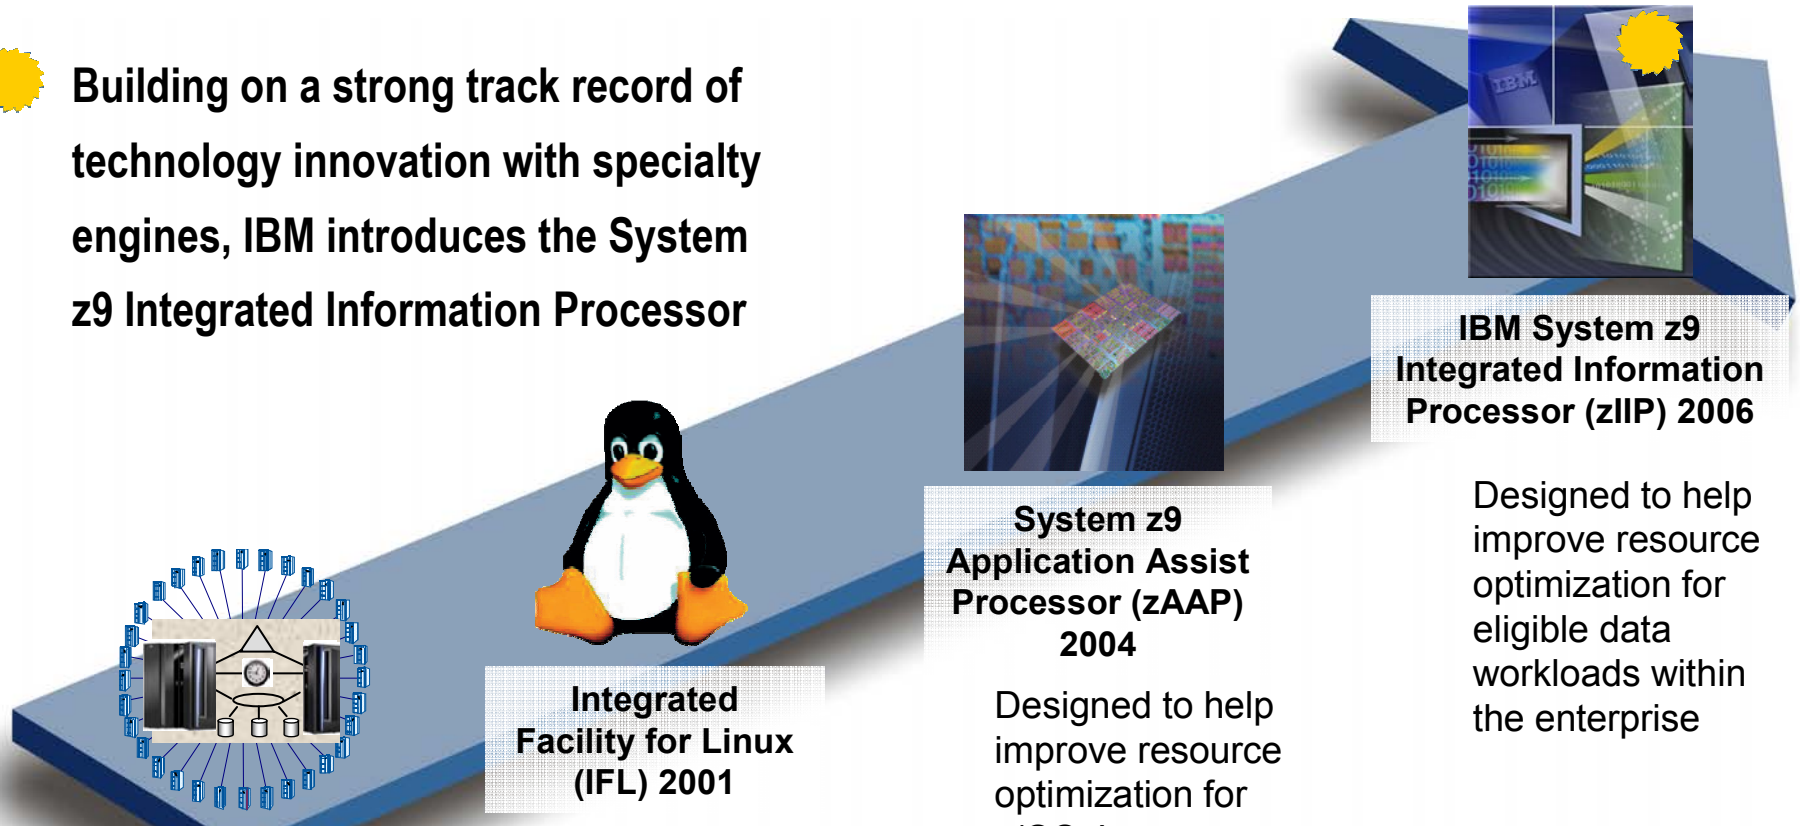

Internal Coupling Facility (ICF) 1997

Centralized data sharing across mainframes

Integrated Facility for Linux (IFL) 2001

Support for new workloads and open standards

System z9 Application Assist Processor (zAAP) 2004

Designed to help improve resource optimization for z/OS Java technology-based workloads

IBM System z9 Integrated Information Processor (zIIP) 2006

Designed to help improve resource optimization for eligible data workloads within the enterprise

System z9 PU Characterization

- **The type of Processor Units (PUs) that can be ordered on z9 BC:**
  - **Central Processor (CP)**
    - **Provides processing capacity for z/Architecture™ and ESA/390 instruction sets**
    - **Runs z/OS, z/VM, VSE/ESA, z/VSE, TPF4, z/TPF, Linux for System z9 and zSeries and Linux under z/VM or Coupling Facility**
    - **System z9 has Capacity Marker features NOT Unassigned CP features**
  - **IBM System z9 Application Assist Processor (zAAP)**
    - **Under z/OS, the Java™ Virtual Machine (JVM) assists with Java processing to a zAAP**
  - **IBM System z9 Integrated Information Processor (zIIP)**
    - **Provides processing capacity for selected workloads e.g. DB2 UDB for z/OS V8 workloads executing in SRB mode**
  - **Integrated Facility for Linux (IFL)**
    - **Provides additional processing capacity for Linux workloads**
  - **Internal Coupling Facility (ICF)**
    - **Provides additional processing capacity for the execution of the Coupling Facility Control Code (CFCC) in a CF LPAR**
  - **Optional System Assist Processors (SAP)**
    - **SAP manages the start and ending of I/O operations for all Logical Partitions and all attached I/O**

z9 BC Concurrent PU Conversions

- **Must order (characterize one PU 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**
- **Conversion of unassigned IFL to any other PU type direct is supported on the z9 BC (This is different than the z9 EC). Conversion to unassigned IFL is via a IFL**

From/To->	CP	IFL	Unassigned IFL	ICF	zAAP	zIIP
CP	x	Yes	No	Yes	Yes	Yes
IFL	Yes	x	Yes	Yes	Yes	Yes
Unassigned IFL	Yes	Yes	x	Yes	Yes	Yes
ICF	Yes	Yes	No	x	Yes	Yes
zAAP	Yes	Yes	No	Yes	x	Yes
zIIP	Yes	Yes	No	Yes	Yes	x

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

PR/SM™ Hypervisor™ PU Dispatching "Pools"

- **PU Pool – Physical PUs to dispatch to online logical PUs**
- **z9 BC with 2 CPs, 1 ICF, 2 IFLs, 1 zIIP and 1 zAAP**
  - CP pool contains 2 CP engines
  - **ICF pool** contains 1 ICF
  - **IFL pool** contains 2 IFLs
  - **zAAP pool** contains 1 zAAP
  - **zIIP pool** contains 1 zIIP
  - **z/OS LPAR can have different CP, zAAP and zIIP weights**

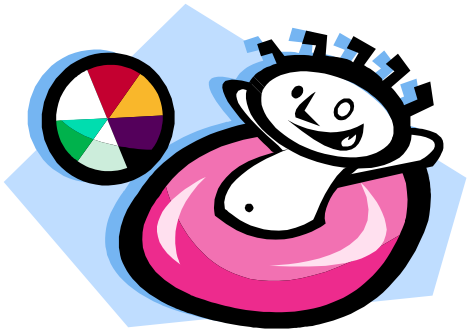

- z890 with 1 CP, 1 ICF, 1 IFL, and 1 zAAP
  - CP pool contains 1 CP engines
  - **Specialty pool** contains 3 engines – ICFs, IFLs, zAAPs
  - z/OS LPAR zAAP weight is set equal to the initial CP weight

IBM System z9 “On Demand”

System z9 Capacity Upgrade on Demand

- **CUoD is designed to support addition of processors and/or memory or concurrent type conversion among CPs, IFLs, and ICFs without disruption to workloads running on the machine - no power-off, power-on. Includes**:
  - Addition of CP, ICF, IFL, zIIP and zAAP
    - Includes turning on (assigning) Unassigned“ IFL features
  - LIC enabling additional memory increments
  - Concurrent z9 EC model upgrade (Concurrent Book Add)
  - Concurrent z9 EC memory upgrade exploiting **Enhanced Book Availability**
- **CUoD capabilities can be exploited by IBM ordered/installed MES upgrade**
- **Some CUoD capabilities can be exploited by customer controlled upgrades:**
  - **Capacity Backup (CBU)** – temporary emergency upgrades
  - **Customer Initiated Upgrade (CIU)** – permanent upgrades
  - **On/Off Capacity on Demand (On/Off CoD)** – temporary on-demand upgrades
- Notes:
  1. CUoD is built on a base of concurrent “hot-plug” maintenance
  2. I/O feature adds and removes are also nondisruptive but not really "CUoD“
  3. Customer planning and operator action are required to take full advantage of CUoD. To avoid a planned outage, it may be necessary to predefine LPAR profiles with “reserved” resource specified. It may also be necessary to use z/OS or z/VM dynamic I/O capabilities. In some cases, disruption of certain LPARs is required following a concurrent hardware change.

System z9 and zSeries "On Demand"

Capacity on Demand	Server
Capacity BackUp	System z9, z990, z890, z900, z800
Capacity Upgrade on Demand	System z9, z990, z890, z900, z800
Customer Initiated Upgrade	System z9, z990, z890, z900, z800
On/Off Capacity on Demand	System z9, z990, z890

	Capacity BackUp	Capacity Upgrade on Demand	Customer Initiated Upgrade	On/Off Capacity on Demand (temp. cap.)
z900	Yes, CP only	Yes (CP, I/O, ICF, IFL, Memory[2])	Yes (CP, IFL, ICF, Memory[2])	No
z800	Yes, CP only	Yes (CP[1], I/O, ICF, IFL)	Yes (CP[1], IFL, ICF)	No
z990	Yes, CP only	Yes (CP, I/O, ICF, IFL, zAAP, Memory[2])	Yes, (CP, IFL, ICF, zAAP, Memory[2])	Yes, CP, IFL, ICF, zAAP
z890	Yes, CP only	Yes (CP[1], I/O, ICF, IFL, zAAP, Memory[2])	Yes, (CP[1], IFL, ICF, zAAP, Memory[2])	Yes, CP[1], IFL, ICF, zAAP
System z9	Yes, CP, IFL, zAAP, zIIP, ICF	Yes (CP[1], I/O, ICF, IFL, zAAP, zIIP, Memory[2])	Yes, (CP[1], IFL, ICF, zAAP, zIIP, Memory[2])	Yes, CP[1], IFL, ICF, zAAP

[1] For z9, z890, z800 – Includes CP speed changes

[2] Limited to upgrades possible with installed memory cards

System z9 BC Capacity Backup Upgrade

- For customers who have a requirement for robust Disaster Recovery
- What Is It?
  - Temporary, non-disruptive addition of one or more CPs, IFLs, ICFs, zAAPs, or zIIPs
    - Memory and channels are not included
  - Must plan ahead for memory and connectivity requirements
  - Contract between IBM and customer
  - Count of CBU Features is the number of IFLs, ICFs, zAAPs or zIIPs to be added. For CPs, if within the same 'speed' it's the number of CPs to be added. If going to a different 'speed' CP, it's the TOTAL number of CPs which will be active.
    - FC: 7870 – 7895 CBU CP features, one for each target CP speed, 10 choices on R07, 16 on S07
    - FC: 7821 for CBU IFL
    - FC: 7822 for CBU ICF
    - FC: 7824 for CBU zAAP
    - FC: 7825 for CBU zIIP
  - Count of active PU features plus CBU features limited to available PUs
- Non-disruptive temporary upgrade or test process
  - Execute CBU from HMC
  - CBU features activate as CPs, IFLs, ICFs, zAAPs, or zIIPs
  - Configure additional logical CPs ON to active partition
    - Predefine as "Reserved" PU(s)
- Non-disruptive downgrade process
  - Required after recovery or test completed
  - Follow procedures to quiesce workload
  - Configure CBU PU(s) OFF or deactivate using Logical Partition
  - Execute downgrade from HMC

Note: Upgrades are non-disruptive only where there is sufficient hardware resource available and provided pre-planning has been done

z9 BC CBU Enhancements

- **CBU for Specialty Engines**
  - CBU is available for CPs, IFLs, ICFs, zAAPS and zIIPs
  - FULL size specialty engines for CBU
  - During CBU can't reduce engine count or convert engine types of the base machine
- **CBU for CPs**
  - No change for FULL size CPs
  - CBU for sub-capacity CPs
    - In CBU mode, CP count must be equal or greater than base, active CPs
    - In CBU mode, CP capacity measured in MSUs must increase if the CP configuration changes

MCI	1-way	2-way	3-way	4-way
Txx	34	66	95	124
Uxx	38	**73**	106	138
Vxx	42	82	119	155
Wxx	47	92	134	174

MCI = Machine Capacity Identifier

Numbers = MSUs

= Valid CBU, Base = U02 (73 MSUs)

= Invalid target for CBU

- **CBU Pricing**
  - Price is per CBU feature
    - Requires one CBU CP feature to change the speed of an existing CP or add an additional CP
  - It may not be cost effective to change CP speed:
    - In the above example, if MCI U02 has CBU target MCI U04, 2 CBU CP features are required; but, CBU target MCI W03 requires 3 CBU features even though it provides a smaller upgrade.

System z9 On/Off Capacity on Demand (On/Off CoD)

- **For customers who have a requirement for short-term additional capacity, temporary, non-disruptive addition of CPs, IFLs, ICFs, zAAPs and/or zIIPs**
  - **Requires a contract between IBM and the customer and features on the machine**
    - **CIU Enablement (#9898) and On/Off CoD Enablement (#9896)**
  - **Add temporary CPs, IFLs and/or ICFs up to the number of base active CPs plus base unassigned CP capacity, base active and unassigned IFLs, and base ICFs, respectively**
  - **Add temporary zAAPs up to the quantity of base zAAPs**
    - **Temporary zAAPs plus base zAAPs may not exceed the quantity of temporary CPs plus base active CPs plus base unassigned CP capacity.**
  - **Add temporary zIIPs up to the quantity of base zIIPs.**
    - **Temporary zIIPs plus base zIIPs may not exceed the quantity of temporary CPs plus base active CPs plus base unassigned CP capacity.**
  - **Billing through one or more of the following features:**
    - **On/Off CoD Active CP-Day for base CP speed, On/Off CoD Active IFL Day, On/Off Active ICF-Day, or On/Off CoD Active zAAP-Day, or On/Off CoD Active zIIP-Day**
    - **HW charges calculated monthly and billed in arrears based on 24-hour usage period**

- **Non-disruptive temporary upgrade and downgrade**
  - **Order On/Off CoD upgrade on Resource Link**
  - **Download and install upgrade On/Off CoD upgrade LIC from RSF using an HMC**
  - **Deactivate On/Off CoD upgrade using an HMC (same action as CBU undo)**

Note: Upgrades are non-disruptive only where there is sufficient hardware resource available and provided pre-planning has been done

System z9 On/Off Capacity on Demand (On/Off CoD)

- **For customers who have a requirement for short-term additional capacity, temporary, non-disruptive addition of CPs, IFLs, ICFs, zAAPs and/or zIIPs**
  - Requires a contract between IBM and the customer and features on the machine
    - CIU Enablement (#9898) and On/Off CoD Enablement (#9896)
  - Add temporary CPs, IFLs and/or ICFs up to the number of base active CPs plus base unassigned CP capacity, base active and unassigned IFLs, and base ICFs, respectively
  - Add temporary zAAPs up to the quantity of base zAAPs
    - Temporary zAAPs plus base zAAPs may not exceed the quantity of temporary CPs plus base active CPs plus base unassigned CP capacity.
  - Add temporary zIIPs up to the quantity of base zIIPs.
    - Temporary zIIPs plus base zIIPs may not exceed the quantity of temporary CPs plus base active CPs plus base unassigned CP capacity.
  - Billing through one or more of the following features:
    - On/Off CoD Active CP-Day for base CP speed, On/Off CoD Active IFL Day, On/Off Active ICF-Day, or On/Off CoD Active zAAP-Day, or On/Off CoD Active zIIP-Day
    - HW charges calculated monthly and billed in arrears based on 24-hour usage period

- **Non-disruptive temporary upgrade and downgrade**
  - Order On/Off CoD upgrade on Resource Link
  - Download and install upgrade On/Off CoD upgrade LIC from RSF using an HMC
  - Deactivate On/Off CoD upgrade using an HMC (same action as CBU undo)

Note: Upgrades are non-disruptive only where there is sufficient hardware resource available and provided pre-planning has been done

System z9 On/Off CoD Enhancements – 1

- **Full Function Test**
  - One no-charge test per Server contract.
  - Enables customer to use On/Off CoD function and install/remove additional capacity
  - A maximum duration of 24 hours commencing with the download and activation of an On/Off CoD order.
  - On/Off CoD tests that do exceed 24 hours in duration will be treated in their entirety as billable On/Off CoD upgrades.
- **Introducing special Administrative On/Off CoD Test**
  - Enables customers to order 'zero' quantity features via Resource Link for:
    - Pre-staging On/Off CoD order
    - Activating and deactivating 'zero' quantity On/Off CoD
  - To allow customer staff to order/test/rehearse/document whole On/Off CoD process without incurring any cost. Zero quantity features = zero cost
  - Unlimited number of tests
  - No time period restrictions

System z9 On/Off Capacity on Demand Enhancements – 2

- **Additional flexibility for On/Off CoD function**
  - With On/Off CoD already activated, customer can add additional On/Off CoD capacity without having to restore system to 'purchased' capacity – as today
  - Customer can keep adding or removing capacity using On/Off CoD without have to go back to 'purchased' capacity using On/Off CoD
    - Linked to purchased capacity rule i.e. maximum capacity customer can have with On/Off CoD is 2 x of purchased capacity
    - Limit controlled by Capacity Marker feature.
  - Customer charged for additional capacity on 24 hour basis
  - If customer increases capacity multiple times during a 24 hour period, charge applies to the highest amount of capacity activated
  - Total of 60 minutes grace period allowed within a 24 hour On/Off CoD "day"
- **Enhancement for Automation Code**
  - New API being provided
  - Currently Operations have to select On/Off CoD from the HMC
  - New API is designed to enable customers to modify automation code to use SNMP to send command with order number to HMC to enable activation
  - Can provide improved flexibility of operation of On/Off CoD function.

System z9 and zSeries CIU and On/Off CoD

- **Order CIU and CoD "right to use" features**
  - Qualification, contracting, and pricing
  - Resource Link ID Authorization
- **Customer CIU or On/Off CoD order or On/Off CoD test order (up to 24 hours)**
  - Configure upgrade on Resource Link
  - Secondary Approval (Option)
  - Resource Link communicates with Remote Support Facility (RSF) to stage order and prepare download
  - Customer notified order ready
- **Access Support Element (SE) using Hardware Management Console (HMC)**
  - "Perform Model Upgrade"
  - Code obtained using RSF and installed on target machine

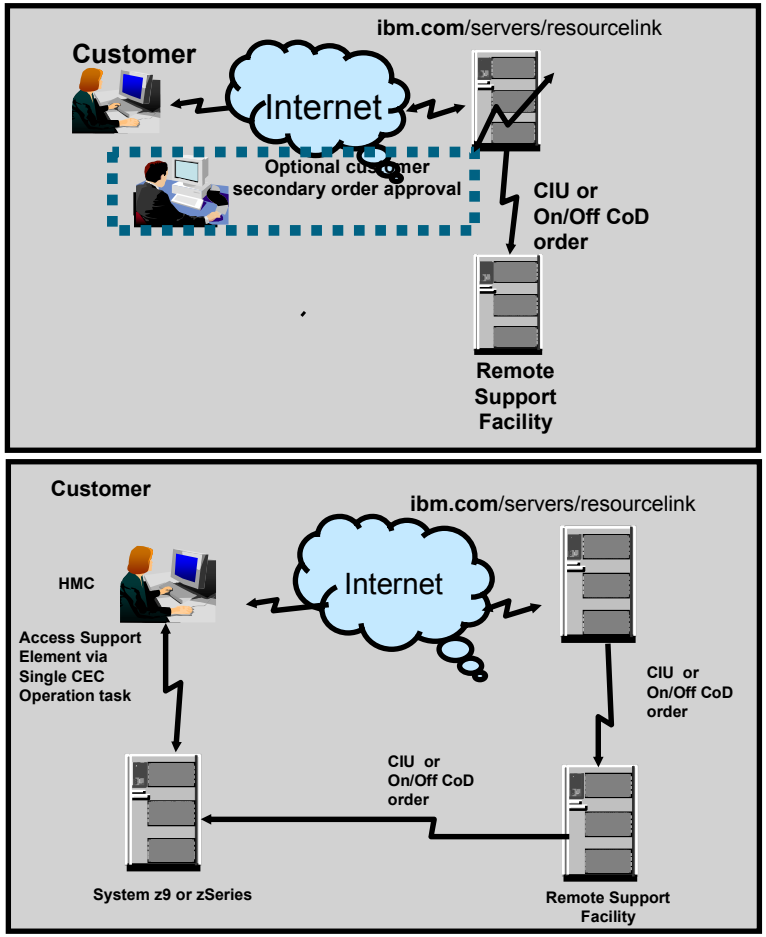

IBM System z9 Cryptography and Security

System z9 and zSeries Crypto Roadmap

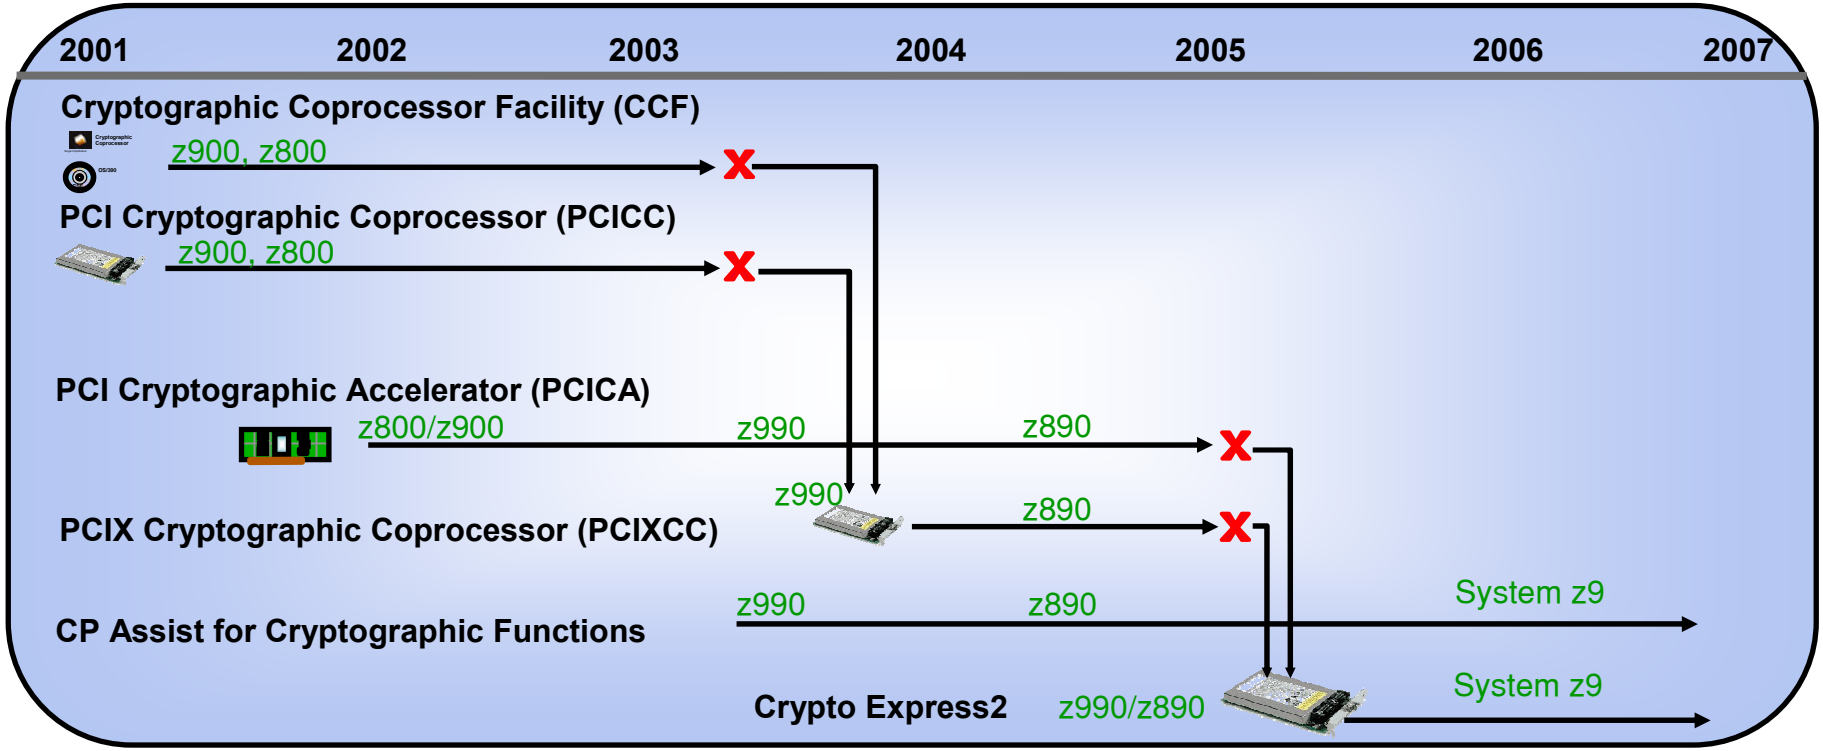

- **System z9, z990, and z890 include NO standard cryptographic coprocessor function**
- **CP Assist for Cryptographic Function (message security assist) Optional Feature #3863**
  - Provides instructions and access to cryptographic functions in every PU
  - Supports limited clear key processing **running on the PU** – Compute intensive!
  - **NOT equivalent to CCF on older machines in function or offload capability**
- **Migration to System z9, z990 or z890 when CCF, PCICC, PCIXCC or PCICA is in use on an older machine almost always requires Crypto Express2.**

System z9 CP Assist for Cryptographic Function (CPACF)

- **High performance cryptographic instructions in every PU but NOT an offload engine**
  - Clear key cryptographic processing, hashing and random number generation
  - Pseudo-optimized for low-latency SSL transactions
- **Five capabilities, three System z9 exclusive:**
  - **Advanced Encryption Standard (AES)**
    - **128 bit keys**
  - **Secure Hashing – 256 (SHA-256)**
  - **Pseudo-random Number Generation (PRNG)**
  - Data Encryption Standard (DES) and Triple DES
    - Up to 2**64 byte message, interruptible execution
  - Secure Hashing (SHA-1)
- **CPACF Enabler Feature FC #3863**
  - No additional charge export control feature
  - Required to enable AES, DES/TES, and PRNG (SHA-1 and SHA-256 are always enabled)
  - Required to order Crypto Express2
  - **Recommended on EVERY system if allowed by law**

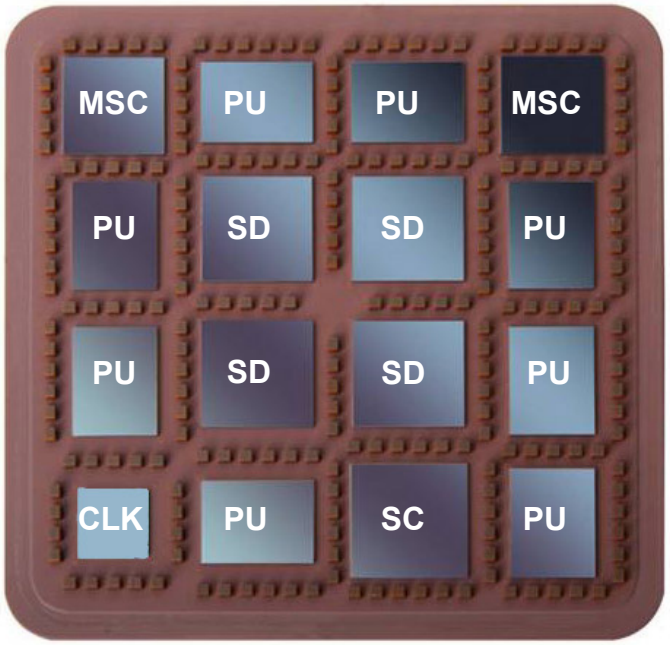

System z9 Crypto Express2 Feature

- **Dual Integrated Cryptographic Coprocessors**
  - Individually configurable as:
    - **Secure Coprocessor** (default), designed to provide both "Secure key" and "Public key" function
    - **Accelerator**, designed to provide only "Public key" function with enhanced performance
  - Current applications expected to run without change
- **Secure Coprocessor mode is fully programmable and supports User Defined Extensions (UDX)**
- **Scalable (no CP affinity) –**
  - Supported Crypto Express2 configurations: 0, 2, 3, 4, 5, 6, 7, or 8 features (but NOT 1 feature)
  - Plugs into an I/O card slot (no external cables)
  - Up to 8 Crypto Express2 features can plug into a single I/O cage
- **Designed for FIPS 140-2 Level 4 Certification**
- **Trusted Key Entry (TKE) support (optional)**
  - If TKE configured, TKE 5 is required on System z9
  - Updated user interface compared to TKE 4.x
  - Secure operational and master key loading
  - Smart Card Reader support

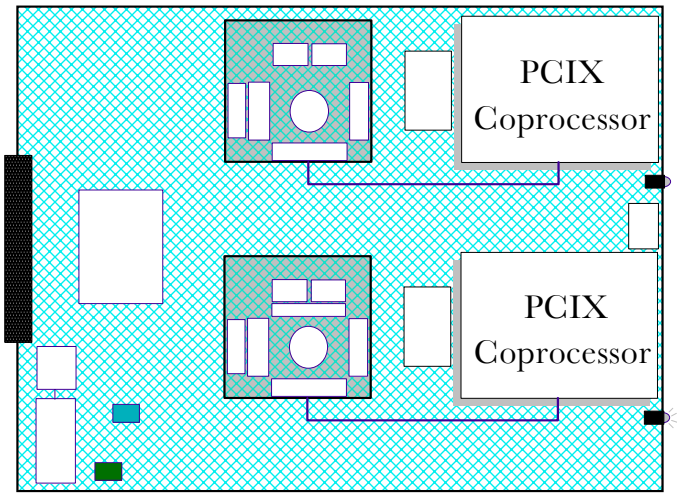

System z9 Crypto Express2 Configuration

- **Secure Coprocessor** (default)
  - Designed to provide both Secure key" and "Public key" function with performance equivalent to Crypto Express2 on z990
  - Designed to provide "Secure key" function with improved performance compared to PCIXCC on z990 (requires multitasking)
  - Designed to provide "Public key" function with performance equivalent to PCICA on z990
  - No configuration action required

- **Accelerator**
  - Designed to provide only "Public key" function with enhanced performance compared to the Secure Coprocessor configuration
  - **Must be configured using the HMC**

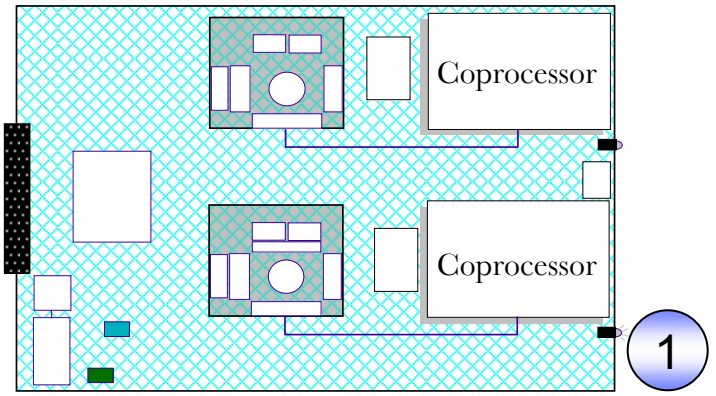

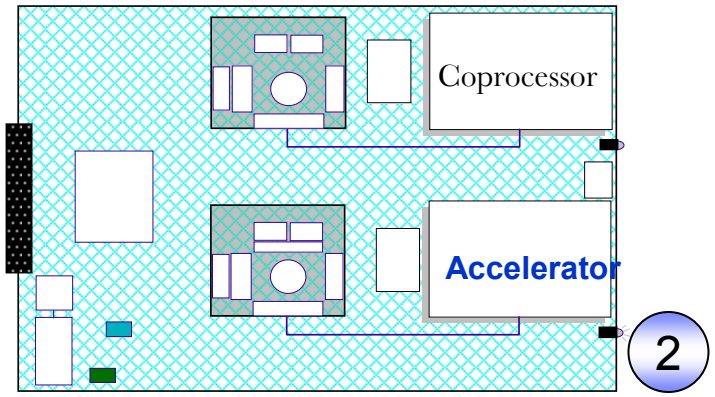

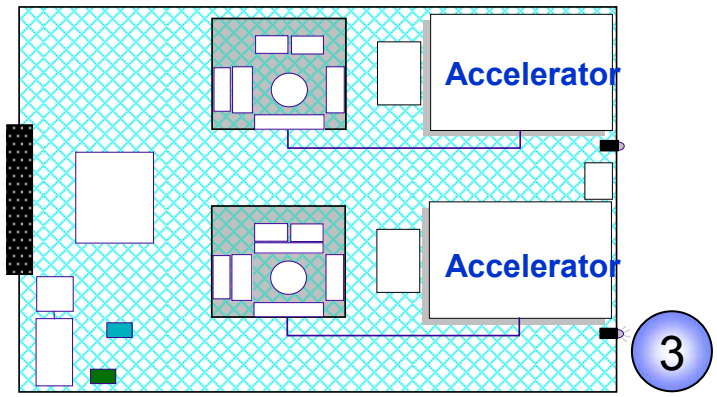

System z9 Cryptographic Enhancements

- **Remote loading of initial Automated Teller Machines (ATM) keys**
  - Typically, a new ATM has none of the financial institutions keys installed.
  - Remote Key Loading refers to the process of loading Data Encryption Standard (DES) keys to Automated Teller Machines (ATMs) from a central administrative site without the need for personnel to visit each machine to manually load DES keys.
  - This has been done by manually loading each of the two cleartext key parts individually and separately into ATMs.
  - Manual entry of keys is one of the most error-prone and labor-intensive activities that occur during an installation, making it expensive for the banks and financial institutions.
- **Remote Key Loading possible Benefits**
  - Provides a mechanism to load initial ATM keys without the need to send technical staff to ATMs.
  - Reduces downtime due to key entry errors.
  - Reduces service call and key management costs.
  - Improves the ability to manage ATM conversions and upgrades.
- **ATMs which can support the IBM System z9 Remote Key Load function:**
  - ATMs supporting the mechanisms, outlined in following standards, should be compatible with IBM System z9 Remote Key Load solution:
    - ISO/IEC 11770-3: Information Technology, Security Techniques, Key Management, Part 3: Mechanisms Using Asymmetric Techniques.
    - ANS X9.24-2 (Draft): Retail Financial Services, Symmetric Key Management, Part 2: Using Asymmetric Techniques for the Distribution of Symmetric Keys.

System z9 Further Cryptographic Enhancements

- **Improved key Exchange with Non-CCA Cryptographic systems:**
  - IBM Common Cryptographic Architecture (CCA) employs Control Vectors to control usage of cryptographic keys.
  - Non-CCA systems use other mechanisms, or may use keys that have no associated control information.
- **Possible benefits of improved key exchange:**
  - Enhances the ability to exchange keys between CCA systems, and systems that do not use Control Vectors.
  - Allows the CCA system owner to define permitted types of key import and export.
    - Can help to prevent uncontrolled key exchange that can open the system to an increased threat of attack.
- **These enhancements are exclusive to System z9 and supported by z/OS 1.6 or higher with a web download**

System z9 Further Cryptographic Enhancements

- **ISO 16609 CBC Mode T-DES MAC enhancement**
- **Supports the requirements for Message Authentication, using symmetric techniques.**
- **The Integrated Cryptographic Service Facility (ICSF) will use the following callable services to access the ISO 16609 CBC Mode T-DES MAC enhancement in the Cryptographic coprocessor:**
  - MAC Generate (CSNBMGN)
  - MAC Verify (CSNVMVR)
  - Digital Signature Verify (CSNDDSV)
- **ISO 16609 CBC mode T-DES MAC is accessible through ICSF function calls made in the Cryptographic Adapter Segment 3 Common Cryptographic Architecture (CCA) code.**
- **This enhancement is exclusive to System z9 and supported by z/OS 1.6 or higher with a web download**

System z9 and zSeries Cryptographic Technology

- Continues to provide flexible Secure Sockets Layer (SSL) acceleration
- Continues to provide competitive symmetric performance in a security-rich environment
- Provides integration of Crypto features via ICSF
- Focuses on required certifications and open standards
- Continues to improve performance
  - Each Crypto Express2 feature on a System z9, with both adapters configured as accelerators is designed to provide up to 6000* SSL handshakes per second

z900/z800 – Dec. 2000/ May 2002
2 SCMs on CEC Board -
CMOS7s+ PCICC/PCICA (10/01)

G6 – June 1999
2 Chips on Processor
MCM - CMOS5x +
PCICC (6/99)

G5 – Sept. 1998
2 Chips on Processor
MCM - CMOS5x +
PCICC (6/99)

G4 – Sept. 1997
SCMs on Planar
Board - CMOS5x

G3 – June, 1997
SCMs on Planar
Board - CMOS5x

System z9 - September 2005, May 2006
Crypto Express2

z990/z890 – January 2005
Crypto Express2

z890 – May 2004
PCIXCC/PCICA

z990 - September 2003
PCIXCC

z990 - June 2003
CPACF/PCICA

z900/z800 – Dec. 2000/ May 2002
2 SCMs on CEC Board -
CMOS7s+ PCICC/PCICA (10/01)

*These measurements are examples of the maximum transactions/second achieved in a lab environment with no other processing occurring and do not represent actual field measurements. Details available upon request.

System z9 Trusted Key Entry (TKE) Workstation 5.0

- **Optional TKE Workstation:**
  - **The only TKE feature that supports System z9**
  - Orderable on System z9, z990, z890, and z900
  - TKE 5.0 LIC: FC 0855
    - **Requires TKE 5.0 hardware**
  - TKE 5.0 hardware: FC 0859 -> FC 0839
    - **Requires TKE 5.0 LIC**
    - xSeries-based system unit, keyboard, flat panel, mouse
    - PCI-X Crypto Coprocessor
    - **Ethernet connectivity only**
  - Optional Smart Card Reader: FC 0887
  - Optional Additional Smart Cards: FC 0888
- **TKE 5.0 Hardware and LIC support to enter secure cryptographic keys for:**
  - System z9 Crypto Express2
  - z990 and z890: PCIXCC and Crypto Express2
  - z900 and z800: CCF and PCICC

NEW!

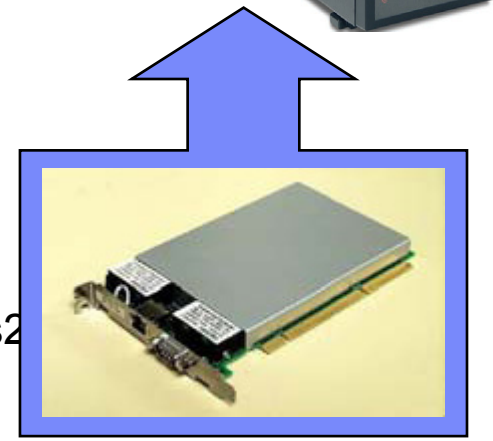

PCI-X Crypto Coprocessor

System z9 Security Summary

- **System z9 EC PR/SM – Common Criteria EAL5 Certification, March 2006**
- **Crypto Express2 feature**
  - Support high levels of security for demanding applications – fully programmable
  - Designed to meet FIPS 140-2 level 4 certification
  - Offers high-scale performance for SSL transactions
  - Integration using ICSF
- **Trusted Key Entry optional smart card reader**
  - Smart Cards – Certified to meet FIPS 140-2 Level 2
- **Common Criteria (ISO/IEC 15408) Evaluation Assurance Levels Reference: http://niap.nist.gov/cc-scheme/**
  - z/OS 1.7 – Controlled Access Protection Profile (CAPP) EAL4+ and Labeled Security Protection Profile (LSPP) EAL4+, February 2006
  - z/VM V5.1 with RACF® for z/VM – Controlled Access Protection Profile (CAPP) EAL3+ and the Labeled Security Protection Profile (LSPP) EAL3+, October 2005
  - Novell SUSE SLES 9 and Red Hat RHEL 4 – Controlled Access Protection Profile (CAPP) EAL4+

zEnd

System z9