G01

zMainframe Concepts (Part 1)

Brian Hatfield

zSeries Expo | Nov. 1 - 5, 2004

Miami, FL

RETURN TO INDEX

zSeries S/390 Server Hardware Overview Introduction

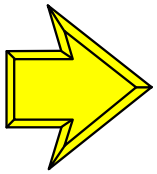

zSeries S/390 Server Hardware Overview

z/OS Operating Systems

Architectural Overview

Common Terms

CMOS - Complimentary Metal Oxide Semiconductor

SE - Support Element

HMC - Hardware Management Console

PU = Processing units

PUs can be assigned as CPs, SAPs, ICFs, IFLs

z900 model 104

Four Engines, 4-Way

Four CPs - Central Processors

IBM @Servers zSeries

zSeries 990 /890 / 900 / 800

or z990 /z890 / z900 / z800

S/390 Parallel Enterprise Server

CPC - Central Processing Complex

CEC - Central Electronic Complex

CPU - Central Processing Unit

Mainframe or Host

Mainframe S/W operating systems

MVS, OS/390, z/OS or → MVS

CHPID - Channel Path ID

PR/SM - Processor Resource Systems Manager

LPAR - Logical Partitioning

(belongs to PR/SM)

EMIF - ESCON Multiple Image Facility

MIF - Multiple Image Facility

(FICON, OSA and CF Links)

CMOS 9672 Hardware Technology

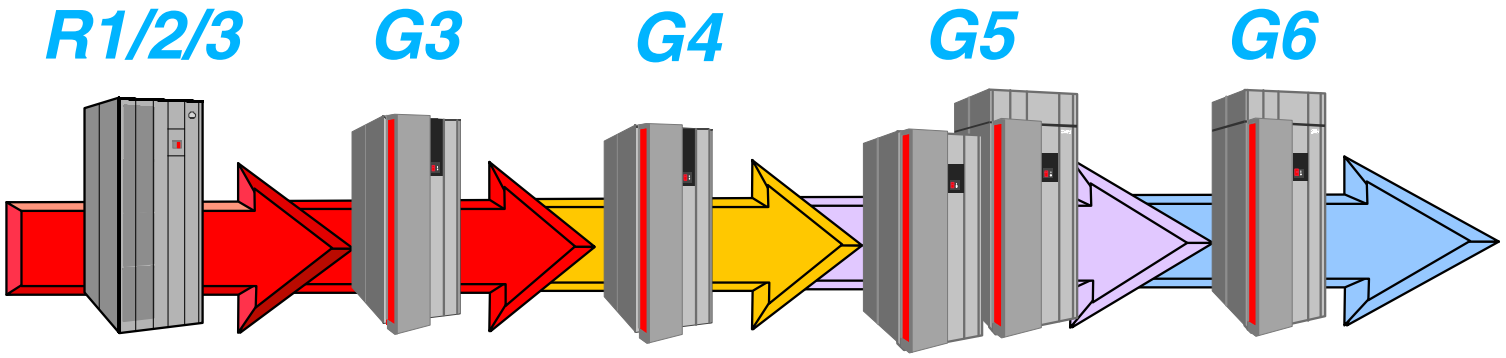

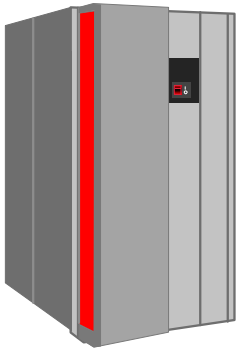

9672

G1 - E0n, P0n, Rn1
G2 - Rn2, Rn3
G3 - Rn4
G4 - Rn5
G5 - Rn6, Yn6
G6 - Xn7, Zn7

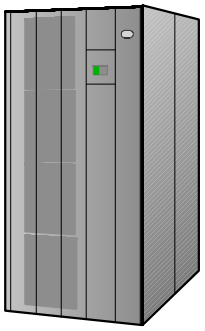

9674

C01
C02
C03
C04
C05
9672 - R06

Multiprise 2003

Multiprise 3000

zSeries Hardware Technology (1 of 2)

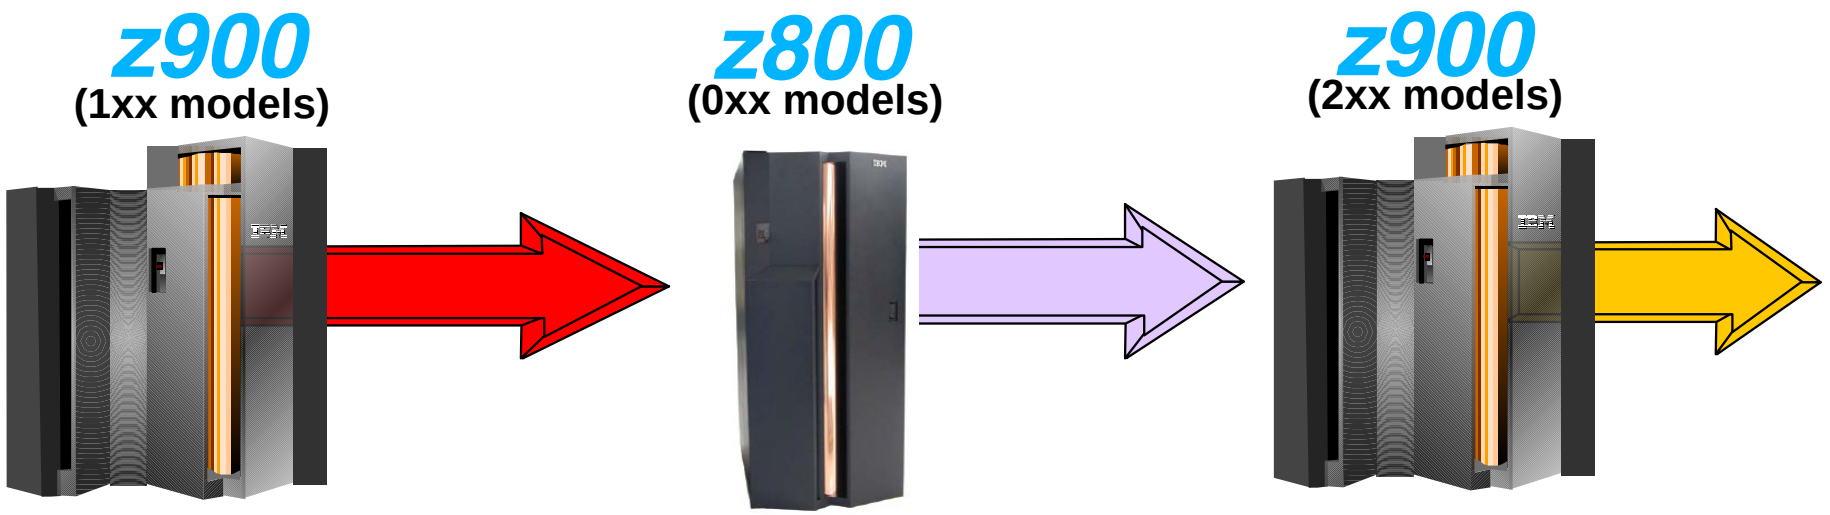

z900 M/T 2064
(October 2000)

- General Purpose
  (101 - 116)
- Capacity models
  (1C1 - 1C9)
- Coupling Facility
  model 100
  (March 2001)

z800 M/T 2066
(February 2002)

- General Purpose
  (0E1 - 004)
- Coupling Facility
  model 0CF
- Dedicated Linux
  model OLF

z900 M/T 2064
(April 2002)

High-performance models

- General Purpose
  (2C1 - 2C9, 210 - 216)

zSeries Hardware Technology (2 of 2)

z990

z890

z990 M/T 2084
(May 2003)

- General Purpose (GA1)
  - A08
  - B16
- Additional models (GA2)
  - C24
  - D32

note - GA2-4th Q 2003
GA3-2nd Q 2004

z890 M/T 2064
(April 2004)

- One General Purpose H/Wmodel
  - A04 (one book - up to 4 CPs)
    Many capacity settings

April 2004 was IBM's 40 year anniversary for the mainframe

IBM @Server zSeries 800 and 900

IBM @Servers zSeries with new z/Architecture

- z/Architecture
  - Based on 64-bit Real and Virtual Storage Addressing
  - Supports trimodal addressing (64-bit,31-bit and 24-bit)
    - ESA/390 supported bimodal addressing (31-bit and 24-bit)
  - Eliminates need of expanded storage
  - Increased register size to support 64-bit instruction/ data addresses
- Intelligent Resource Director
  - LPAR CPU Management
  - Dynamic Channel Path Management
  - Channel Subsystem Priority Queuing
- HiperSockets

- Faster Processor Unit (PU)
  - Up to 20 PUs (z900)
  - Up to 5 PUs (z800)

- z900 Memory
  - Up to 64 GB
- z800 Memory
  - Up to 32 GB

- Channel CHPID Assignment
- Dense Channel Packaging
- New Cabling connectors
- Increased Channel options
  - FICON (z900-96, z800-32)
  - OSA-E (24)
  - PCI-CC (16)
  - PCI-CA (12)
- Increased Total I/O Bandwidth
  - z900 24 GB/sec
  - z800 6 GB/sec
- Increased Parallel Sysplex Connectivity Options
  - Peer mode
  - Compatibility mode
- z900 Upgradable from G5/G6

zSeries Family of Servers

z900

109
108
107
106
105
104
103
102
101

100

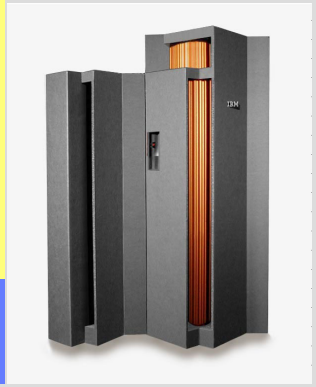

- CMOS 8S with Copper interconnect
- MCUs (Modular Cooling Unit)
- 12 PUs
- Up to nine CPs
- 5 - 32 GB Memory
- 2 Memory cards
- 1.3 ns Cycle time
- CF = Model 100

116
115
114
113
112
111
110
1C9
1C8
1C7
1C6
1C5
1C4
1C3
1C2
1C1

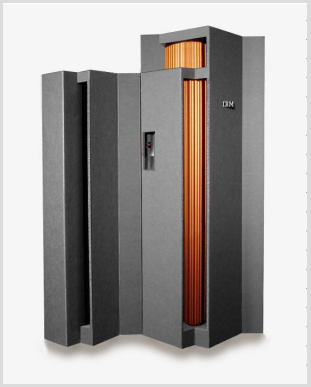

- CMOS 8S with Copper interconnect
- MCUs (Modular Cooling Unit)
- 20 PUs
- Up to 16 CPs
- 10 - 64 GB Memory
- Four Memory cards
- 1.3 ns Cycle time

216
215
214
213
212
211
210
2C9
2C8
2C7
2C6
2C5
2C4
2C3
2C2
2C1

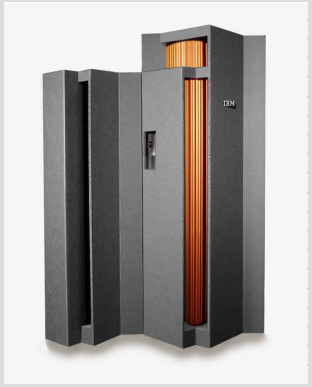

- CMOS 8SE with Copper interconnect
- MCUs (Modular Cooling Unit)
- 20 PUs
- Up to 16 CPs
- 10 - 64 GB Memory
- 4 Memory cards
- 1.09 ns Cycle time

z800

004
003
002
001
0X2
0C1
0B1
0A1
0E1

OCF
OLF

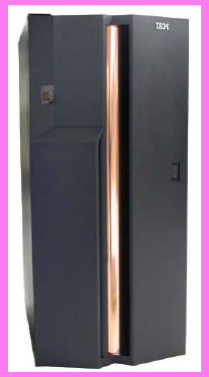

- CMOS 8S with Copper interconnect
- 5PUs
- Up to four CPs
- 8 - 32 GB Memory
- Four Banks of Memory chips
- 1.6 ns Cycle time
- CF = Model OCF
- LINUX model OLF

z800 General Purpose Models Relative Performance

* 001 - 004 is approximately same performance as G6 X17-X47

** 0E1 has one IFL as standard

z900 Models 101 - 216 Relative Performance

- 1nn models offer approx 25% CP and 50% system performance compared to G6 9672 models
- 1Cn models offer a slight performance benefit compared to the 10n models
- 2nn models offer an approx 20% performance benefit over equivalent 1nn models

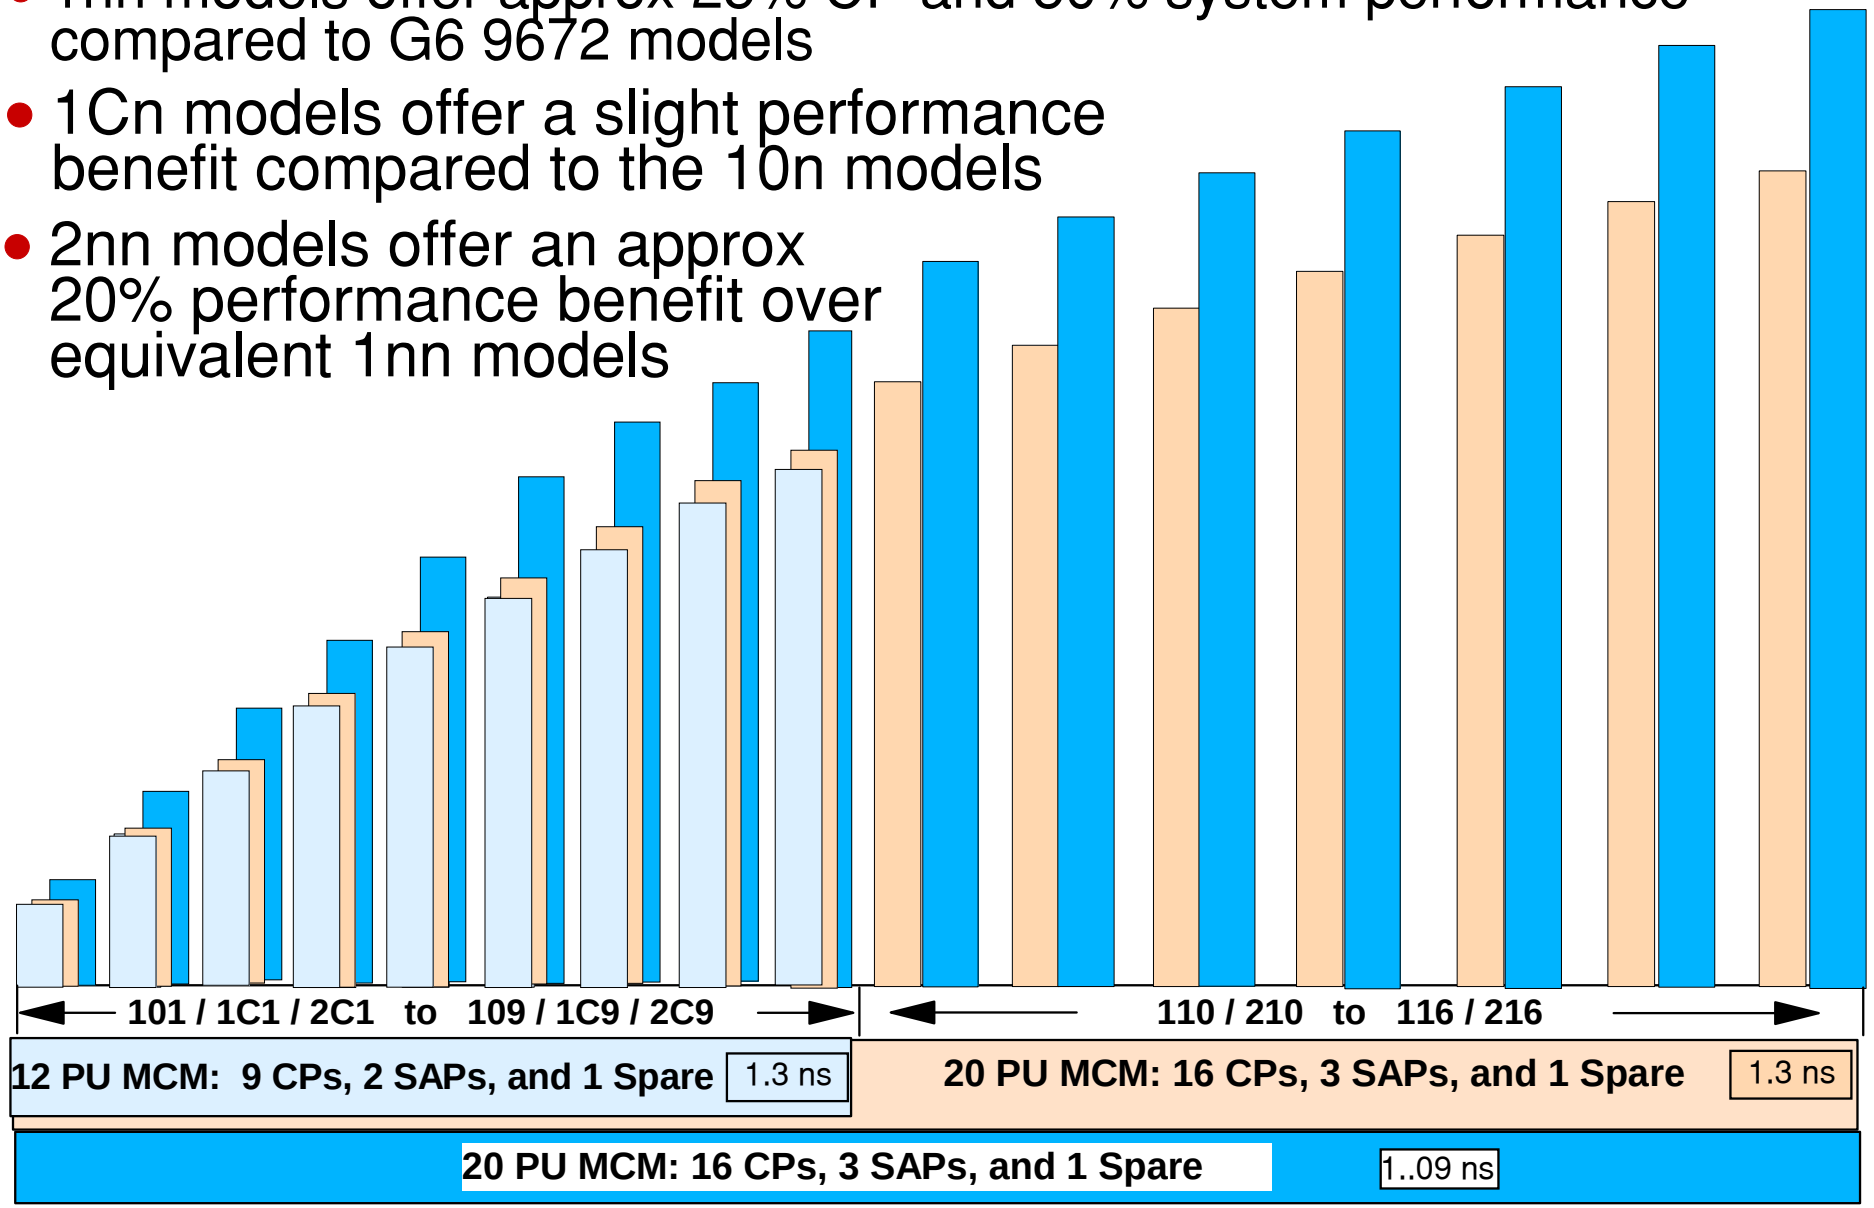

Graph is for illustration purposes only

zSeries 800 CPC Design

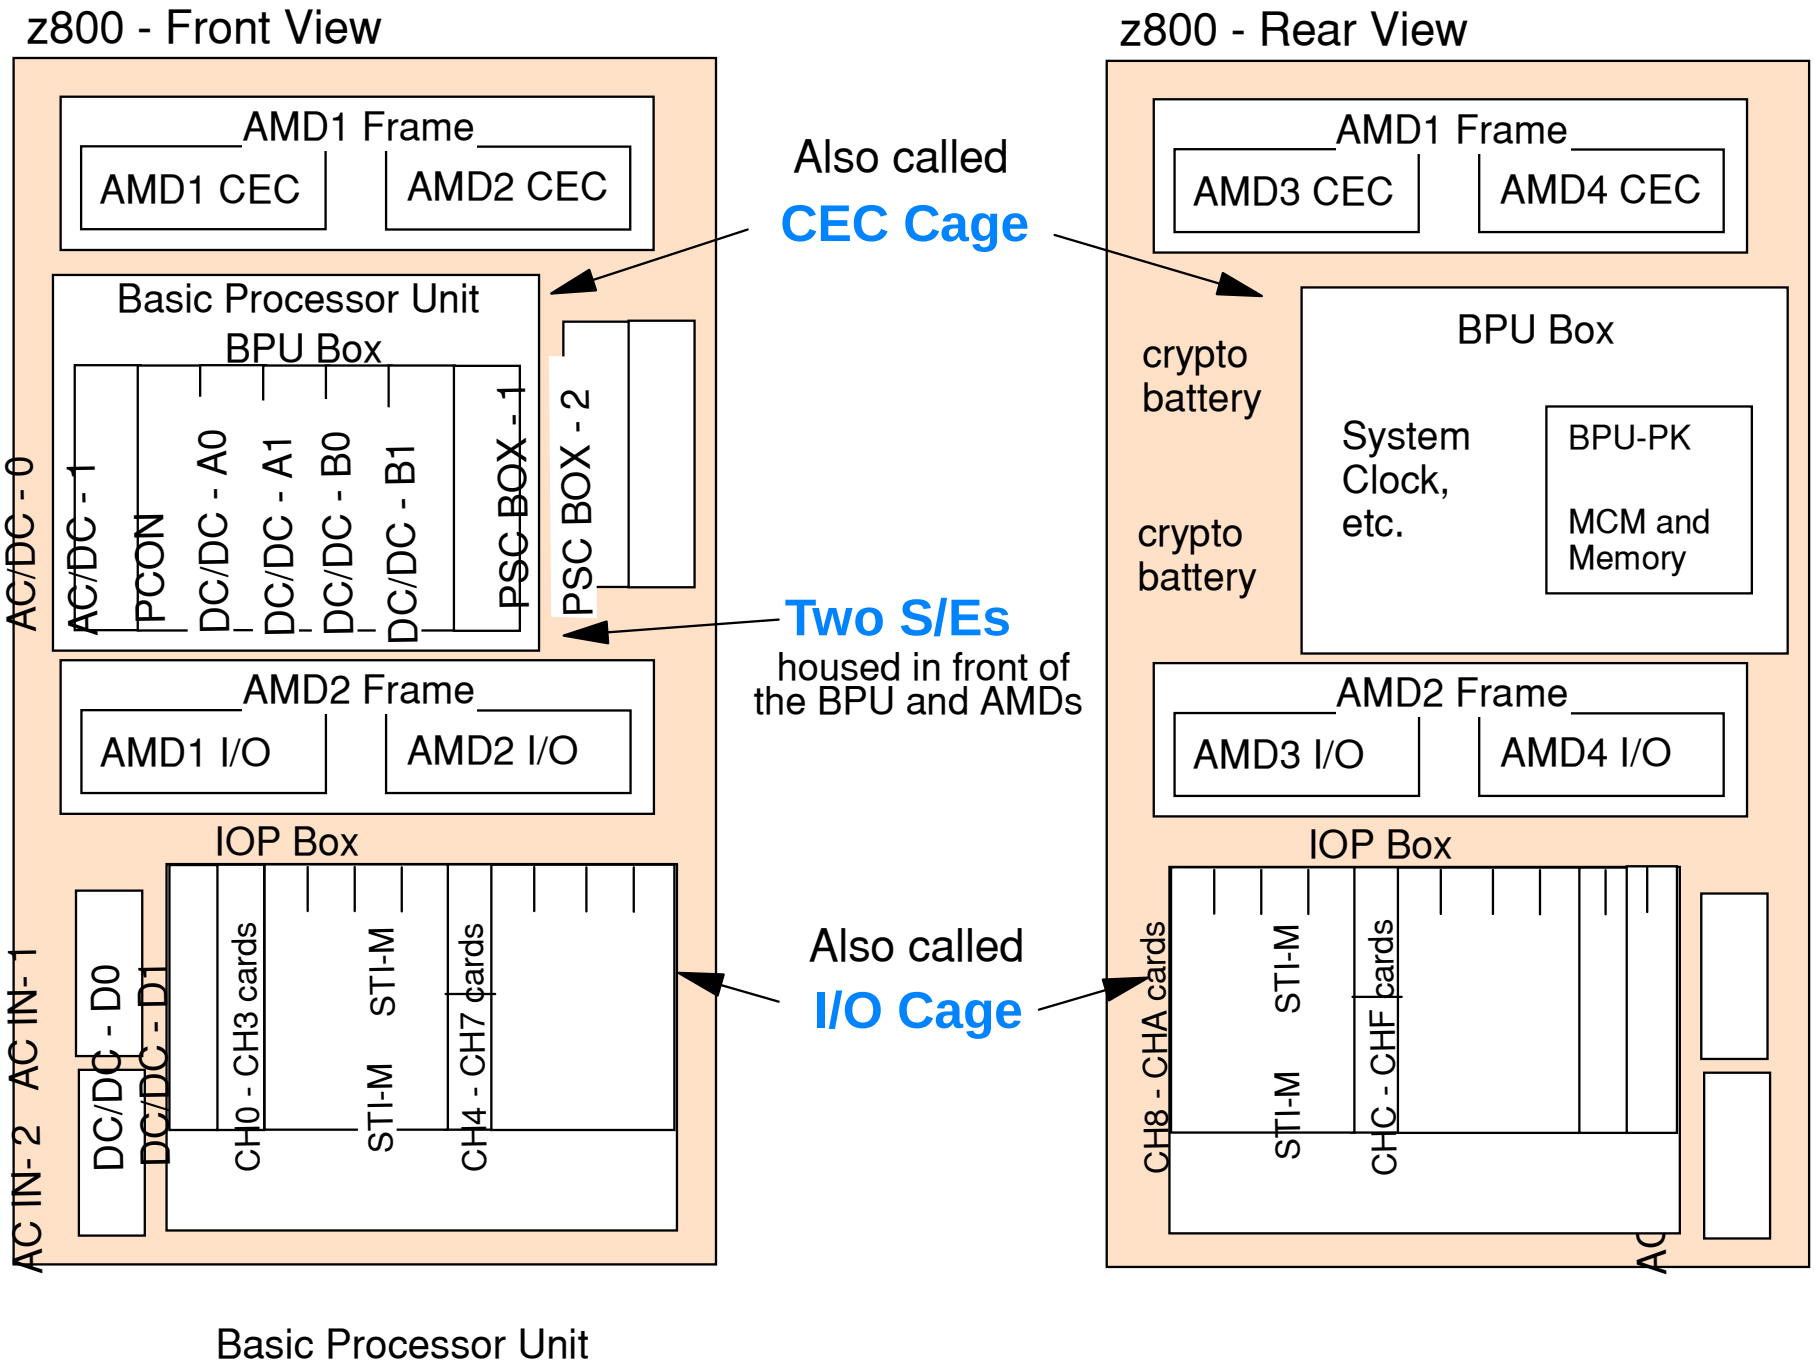

zSeries 900 CPC Design

Modular Cooling Unit (MCU)

- Modular Refrigeration Unit (MRU)
- Motor Scroll Assembly (MSA)
- Motor Drive Assembly (MDA)

N +1 design

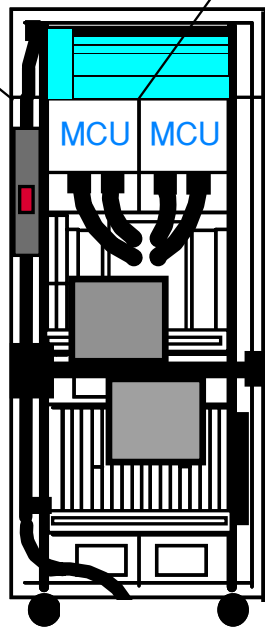

CEC Cage

Two S/Es

I/O Cage

z900 A Frame

May have an additional Z and/or B frame

CEC Cage - Front View

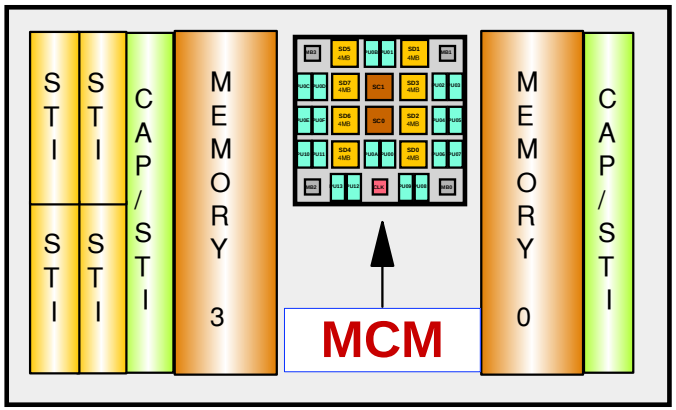

zSeries 900 CPC MCM will contain
- 12 PUs
- 20 PUs

depending on z900 model

CEC Cage - Rear View

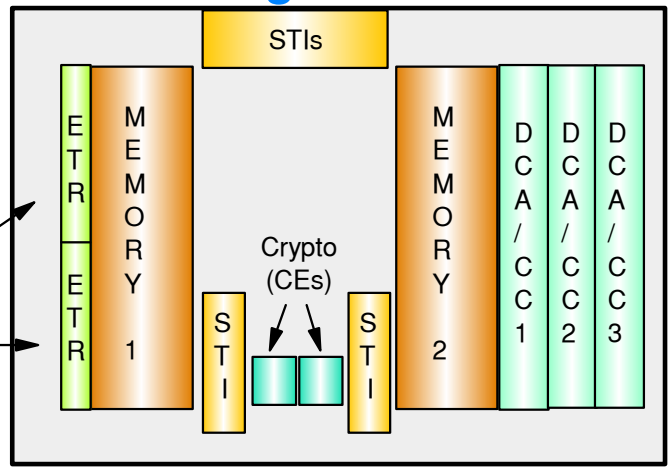

Two ETR cards N+1

The z800 CEC cage holds the Base Processor Unit (BPU) which contain equivalent functionality as the z900
- MCM contains 5 PUs
- Has memory DIMMS

Introducing the IBM eserver zSeries
The z990 and z890 Servers

Designed for on demand business computing

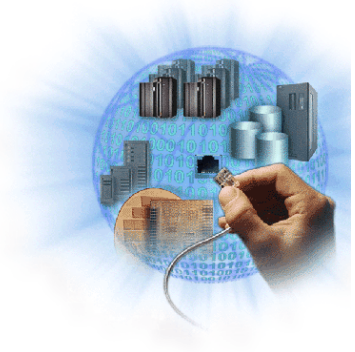

The on demand operating Environment

- Integrated
- Open
- Virtualized
- Autonomic

z890

z990

Even more Virtualization, functionallity flexibility and capacity then previous zSeries

The on demand operating environment enables the on demand business

zSeries 990 / 890 Terminology

Some new Terminology used with z990 / z890 Servers

- Book
  - A book contains an MCM (processors), memory and STI connections for the MBAs
  - A z990 can have multiple books
  - A z890 has one book
- eServer zSeries Application Assist Processor (zAAP)
  - A PU used by the JAVA virtual machine to run JAVA code
- I/O Subsystem
  - All zSeries CPCs has one I/O Subsystem which utilizies one IOCDS and a single HSA
- Logical Channel Subsystem (LCSS)
  - z990 / z890 can have multiple LCSSs
- Physical Channel ID (PCHID)
  - Physical location that can be mapped to a channel path ID (CHPID)
  - A PCHID is unique to the z990 / z890 (CHPID is unique to a LCSS)

IBM @server zSeries 990 Models

- z990 model A08
  - One book with 12 PUs, maximum of 8 PUs can be assigned as CPs
- z990 model B16
  - Two books with 24 PUs, maximum of 16 PUs can be assigned as CPs
- z990 model C24
  - Three books with 36 PUs, maximum of 24 PUs can be assigned as CPs
- z990 model D32
  - Four books with 48 PUs, maximum of 32 PUs can be assigned as CPs

- CMOS 9S-SOI with copper interconnect
- 12 Processor Units (PUs) per MCM, up to 8 as CPs
- 8 - 256 GB Memory
- .8 ns cycle time
- Modular Refrigeration Units (MRU)

z990 models C24 and D32 available at GA2

Note that the system model number no longer reflects the number of CPs

zSeries 990 Software Models

Each z990 model has an additional software model number association. The software model can be used for licensing and MSU purposes. z990 *MSUs range from 70 (301) to 1365 (332)

z990 A08		z990 B16		z990 C24		z990 D32	
S/W model	CPs	S/W model	CPs	S/W model	CPs	S/W model	CPs
301	1	309	9	317	17	325	25
302	2	310	10	318	18	326	26
303	3	311	11	319	19	327	27
304	4	312	12	320	20	328	28
305	5	313	13	321	21	329	29
306	6	314	14	322	22	330	30
307	7	315	15	323	23	331	31
308	8	316	16	324	24	332	32

This chart represents the maximum S/W model that can be assigned to a given H/W model. Different H/W and S/W model combinations are possible depending on factors such as storage and other PU assignments.

***See www-1.ibm.com/servers/eserver/zseries/library/swpriceinfo/hardware.html for current MSU ratings**

Relative Performance of z990 Models

Relative performance scale of new z990 Processors

Number of z990 CPs	Base	Ratio
1 CP	z900-2C1	1.54 - 1.61
8 CPs	z900-2C8	1.52 - 1.56
16 CPs	z900-216	1.51 - 1.55
32 CPS	z900-216	2.46 - 2.98

IBM @server zSeries 890 model A04

- z890 model A04
  - One book with 5 PUs, maximum of 4 PUs can be assigned as CPs
  - Various capacity settings available across ordered CP(s)
  - One standard SAP
  - Eight STIs for I/O connectivity
  - Up to 32 GB of storage

z890

M/T 2086

- CMOS 9S-SOI with copper interconnect
- Five Processor Units (PUs) per MCM, up to 4 as CPs
- 8 - 32 GB Memory
- 1.0 ns cycle time
- Air cooled

Note that the system model number no longer reflects the number of CPs

zSeries 890 Capacity Settings and MSUs

A z890 model A04 has additional capacity settings available. Capacity settings can be used for licensing and MSU purposes.

z890 Capacity Settings and MSU ratings

1-WAY	MSUs	2-WAY	MSUs	3-WAY	MSUs	4-WAY	MSUs
110	4	210	8	310	11	410	15
120	7	220	13	320	20	420	26
130	13	230	26	330	38	430	49
140	17	240	32	340	47	440	62
150	26	250	50	350	74	450	97
160	32	260	62	360	91	460	119
170 - Full 1-way	56	270 - Full 2-way	107	370 - Full 3-way	158	470 - Full 4-way	208

*Capacity settings are reported by certain software instructions as a machine model number

A capacity setting of 070 indicates a z890 with no CPs assigned, This could be a z890 with only IFLs or ICFs or some combination of IFLs and ICFs

See www-1.ibm.com/servers/eserver/zseries/library/swpriceinfo/hardware.html for current MSU ratings

IBM @server zSeries 990/z890 Features

z990 1 to 48 PU engines (32 CPs)
z890 1 to 5 PU engines (4 CPs)
- superScalar design

Greater functionallity Up to 30 Logical Partitions

Spare Engines

zSeries Application Assist Processor (zAAP)

Internal Battery Feature

Security: Increased SSL throughput
- CP Crypto Assist Function (CPACF) Cryptographic-assist instructions
- PCI Cryptographic Accelerators (PCICA)
- PCIX Cryptographic Coprocessors (PCIXCC)

Processor Storage
z990 Up to 256 GB
z890 Up to 64 GB

- Up to 48 (z890 40) OSA-E ports
- Up to 120 (z890 40) FICON channels
- Up to 16 HiperSockets

Multiple Logical Channel Subsystems, 256 Channels per LCSS

Up to 48 (z890 8) 2GB/sec STIs

Coupling Links:
- Internal Coupling Channels (IC)
- InterSystem Channels (ISC)
- Integrated Cluster Bus (ICB)

Spanned Channels
- Channels that may access more than one LCSS

On/Off Capacity on Demand

z990/z890 Processor CEC Cage Design

MRU components

- Modular Refrigeration Unit (MRU)
  - Motor Scroll Assembly (MSA)
  - Motor Drive Assembly (MDA)

Hybrid cooling design

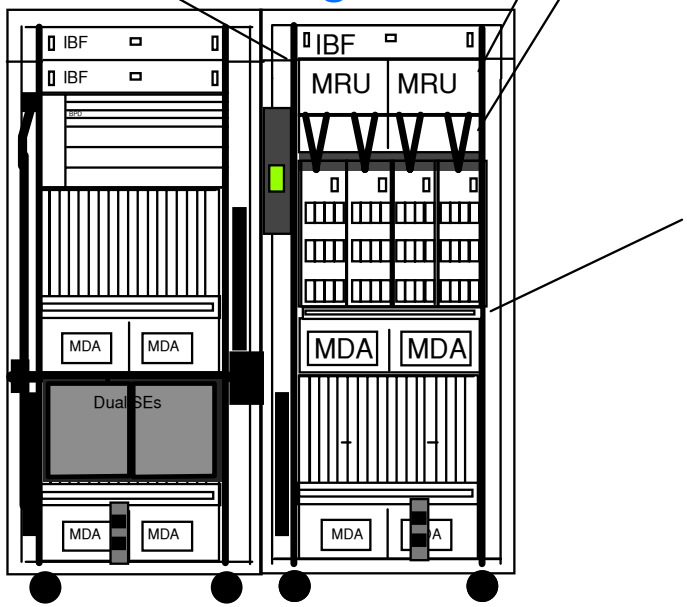

Z frame A frame

CEC Cage - Front View

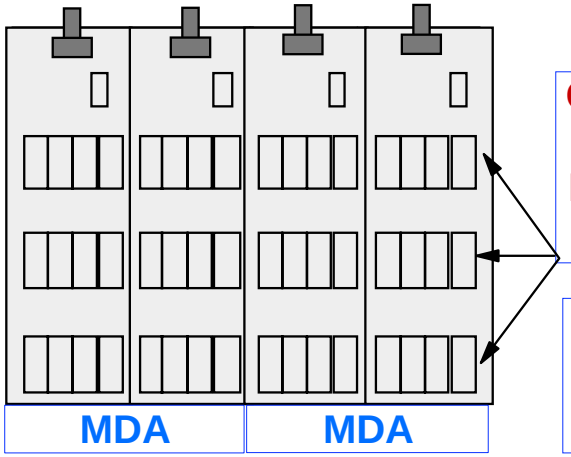

One to four books

Each book has 12 STI connections

z890
One book with 8 STI connections

CEC Cage - Rear View

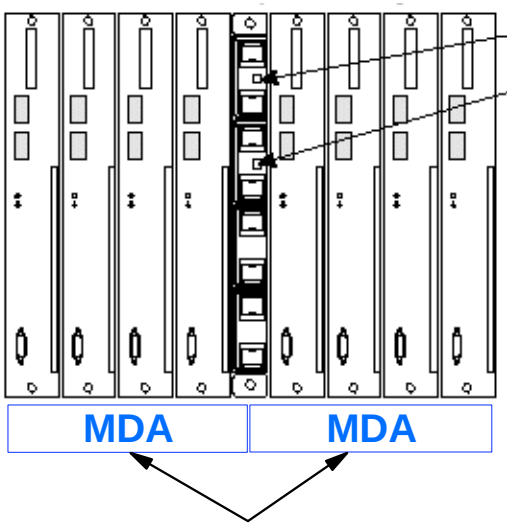

Two ETR and multiple DCA cards located in rear of CEC cage (N+1)

These MDAs are system activated in case of MRU failure

* z890 - A frame only, one book, MDA used for cooling (N+ 1)

Processor Units Assignments

- A PU can be assigned (characterized) as the following:
  - A Central Processor (CP)
  - A System Assist Processor (SAP)
  - An Internal Coupling Facility (ICF)
  - An Integrated Facility for Linux (IFL)
  - zSeries Application Assist Processor (zAAP) **z990/z890 only**
- Unassigned PUs are considered to be spare PUs
  - Any spare PU can be used for CP, SAP, ICF or IFL sparing
  - Spare PUs can also be used for dynamic upgrades (On/Off CoD, CUoD, CIU, or CBU)
  - Number of spare PUs is dependent on PU configuration and model
  - z900 comes with at least one spare PU
  - z990 comes with at least two spare PUs

Large Systems Storage Hierarchy

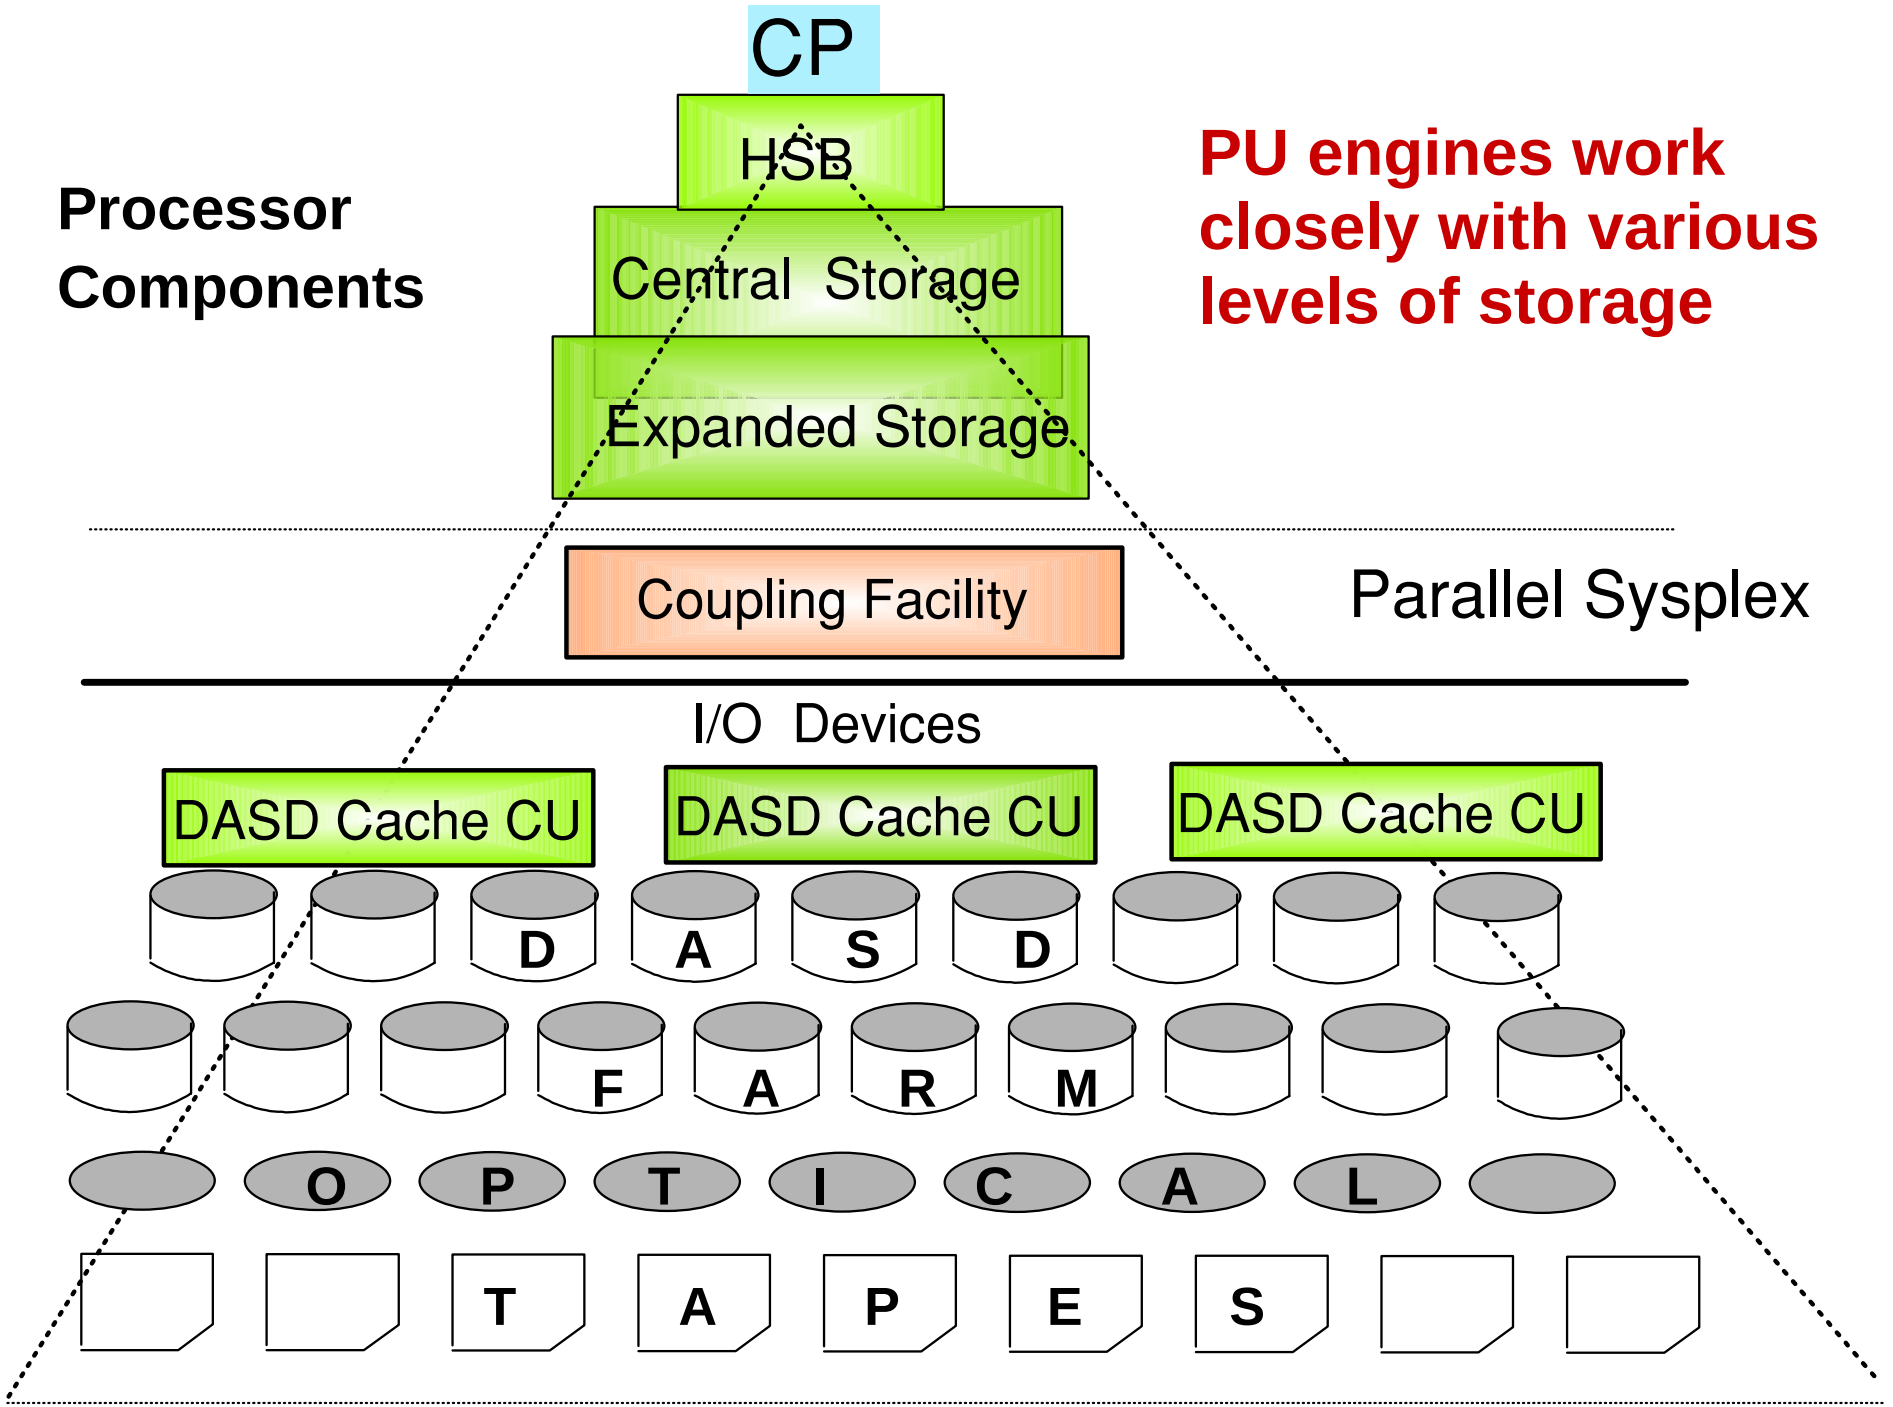

z800 Five-PU MCM System Structure

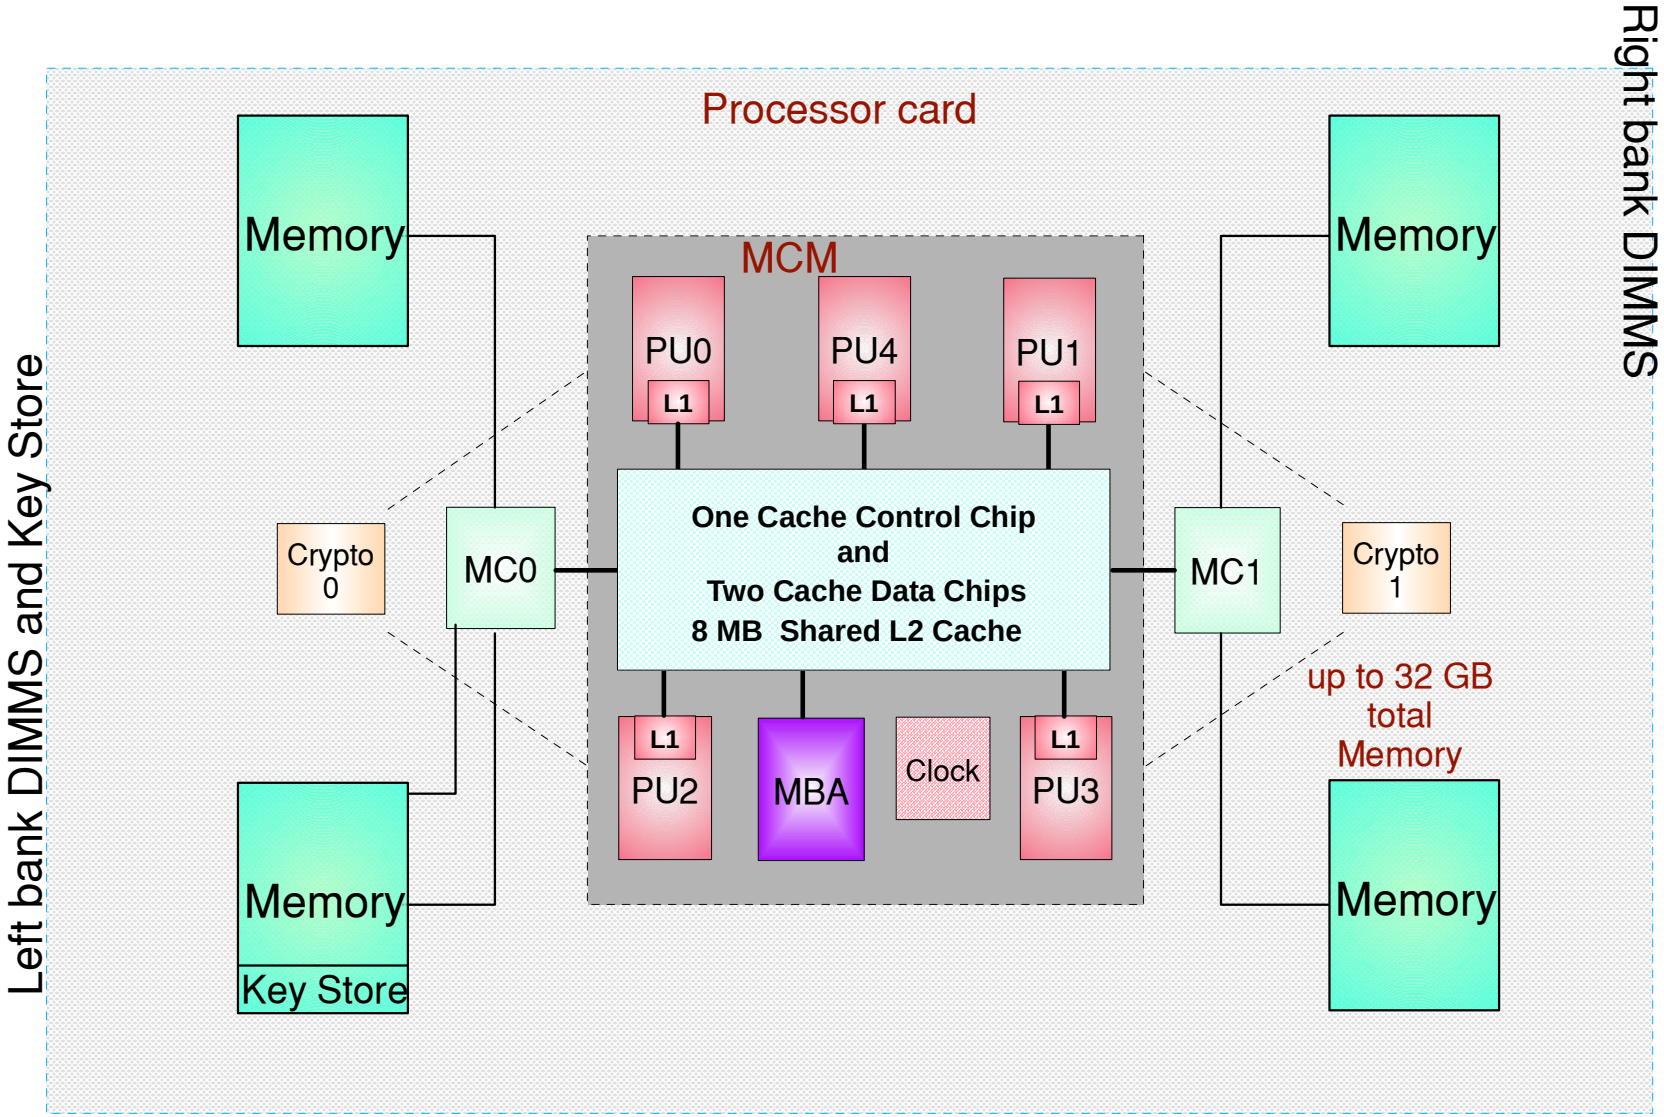

z900 20 PU MCM System Structure

Number of PUs (12 or 20), Memory cards (2 or 4)

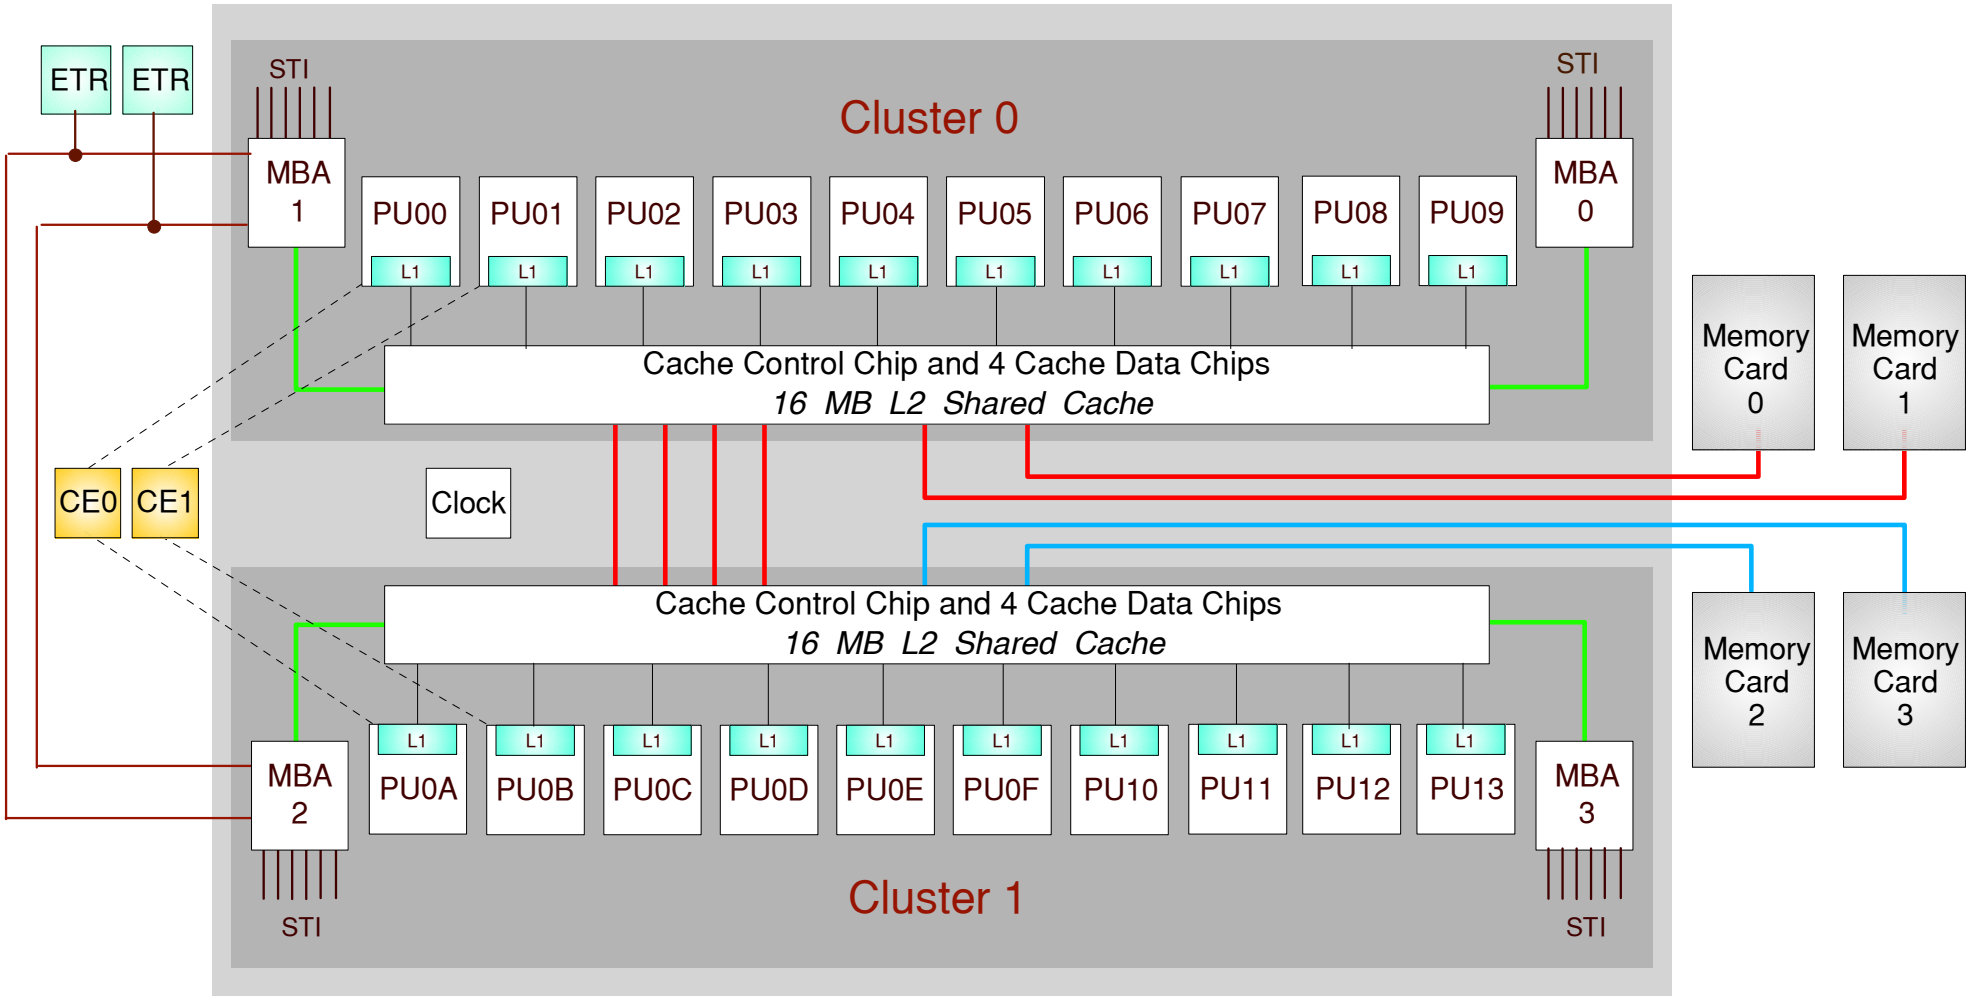

Self Timed Interfaces (STIs) provide all connectivity to I/O via various channel cards and cable connections

- STIs on the z900 and z800 operate at 1GB/sec

z990/z890 Book and MCM System Structure

z990 Book configuration

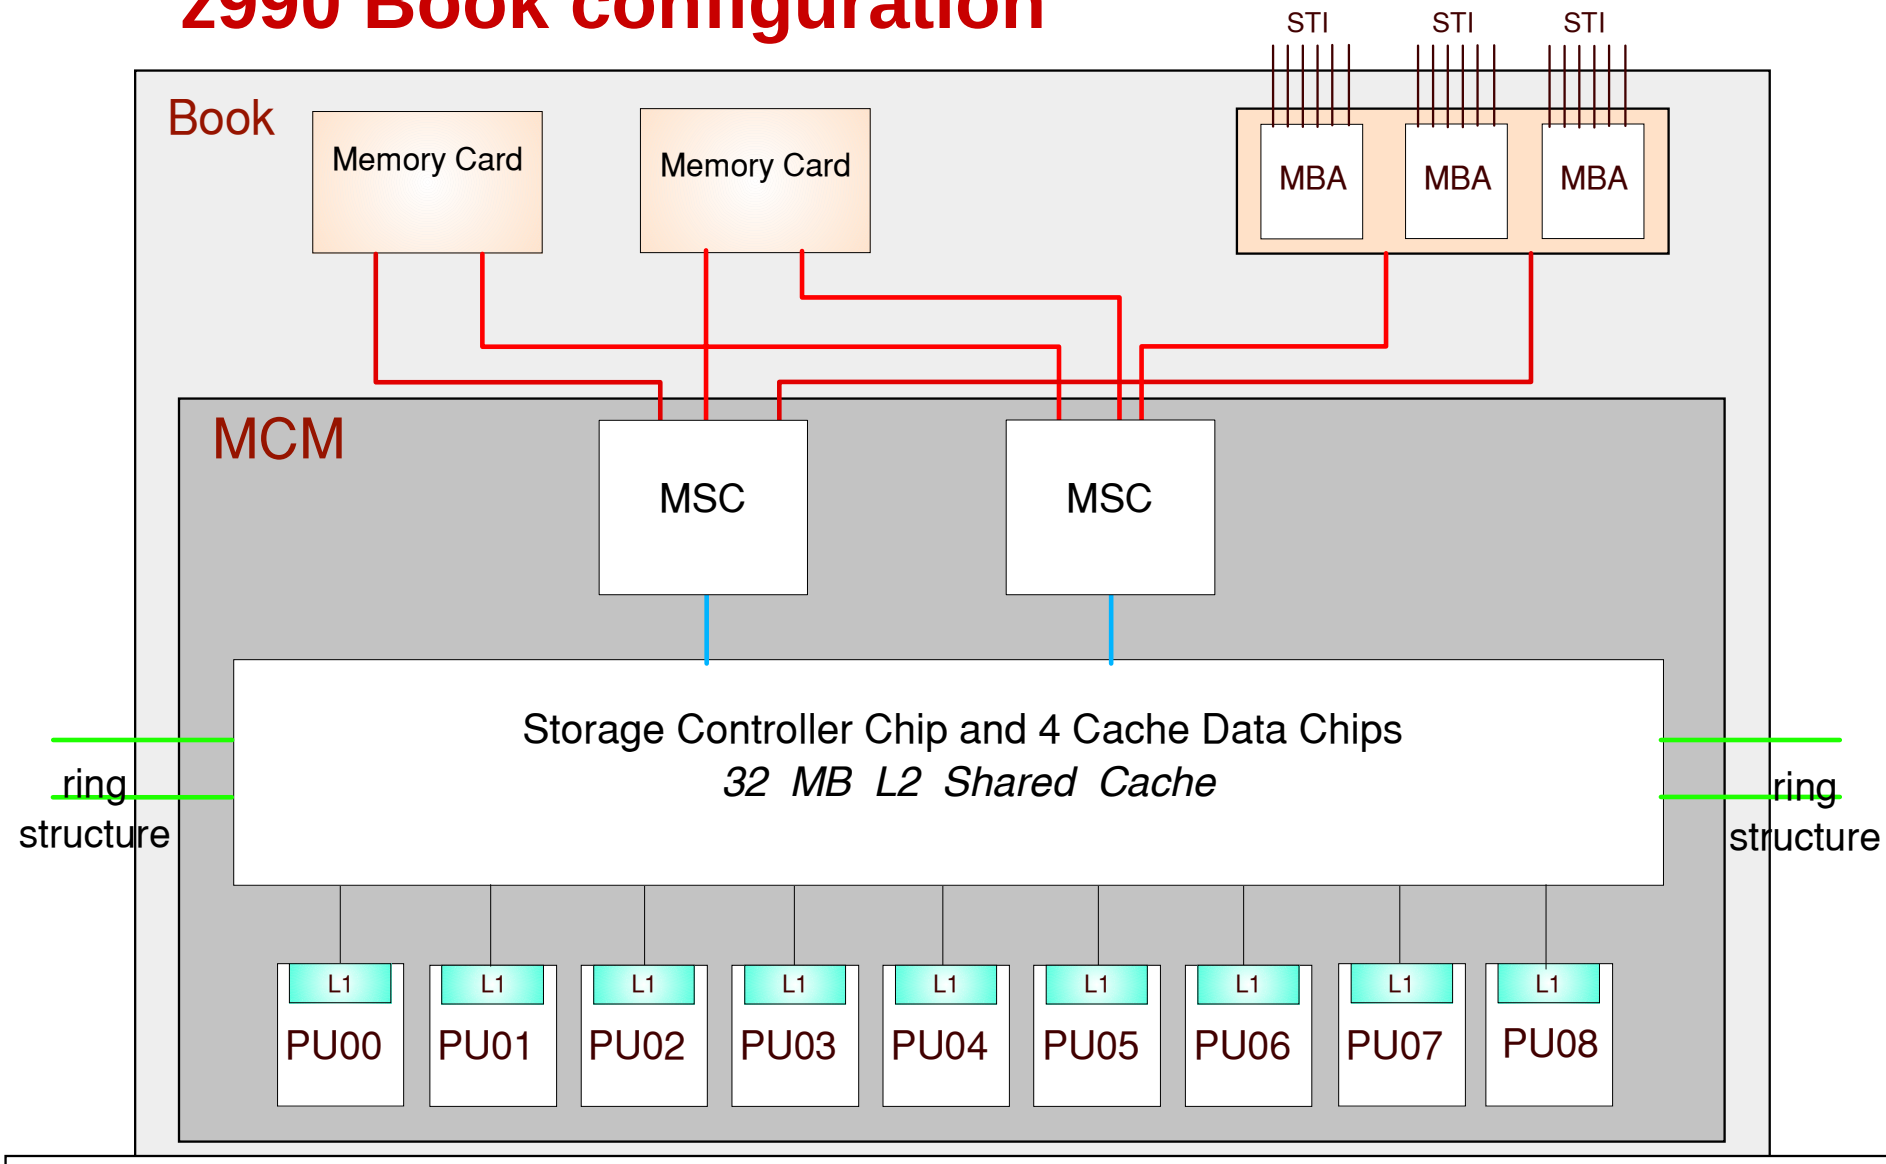

z890 Book configuration similar, with the following exceptions

One Memory card, Two MBAs, 5 PUs

z990 Memory Ring and Book Structure

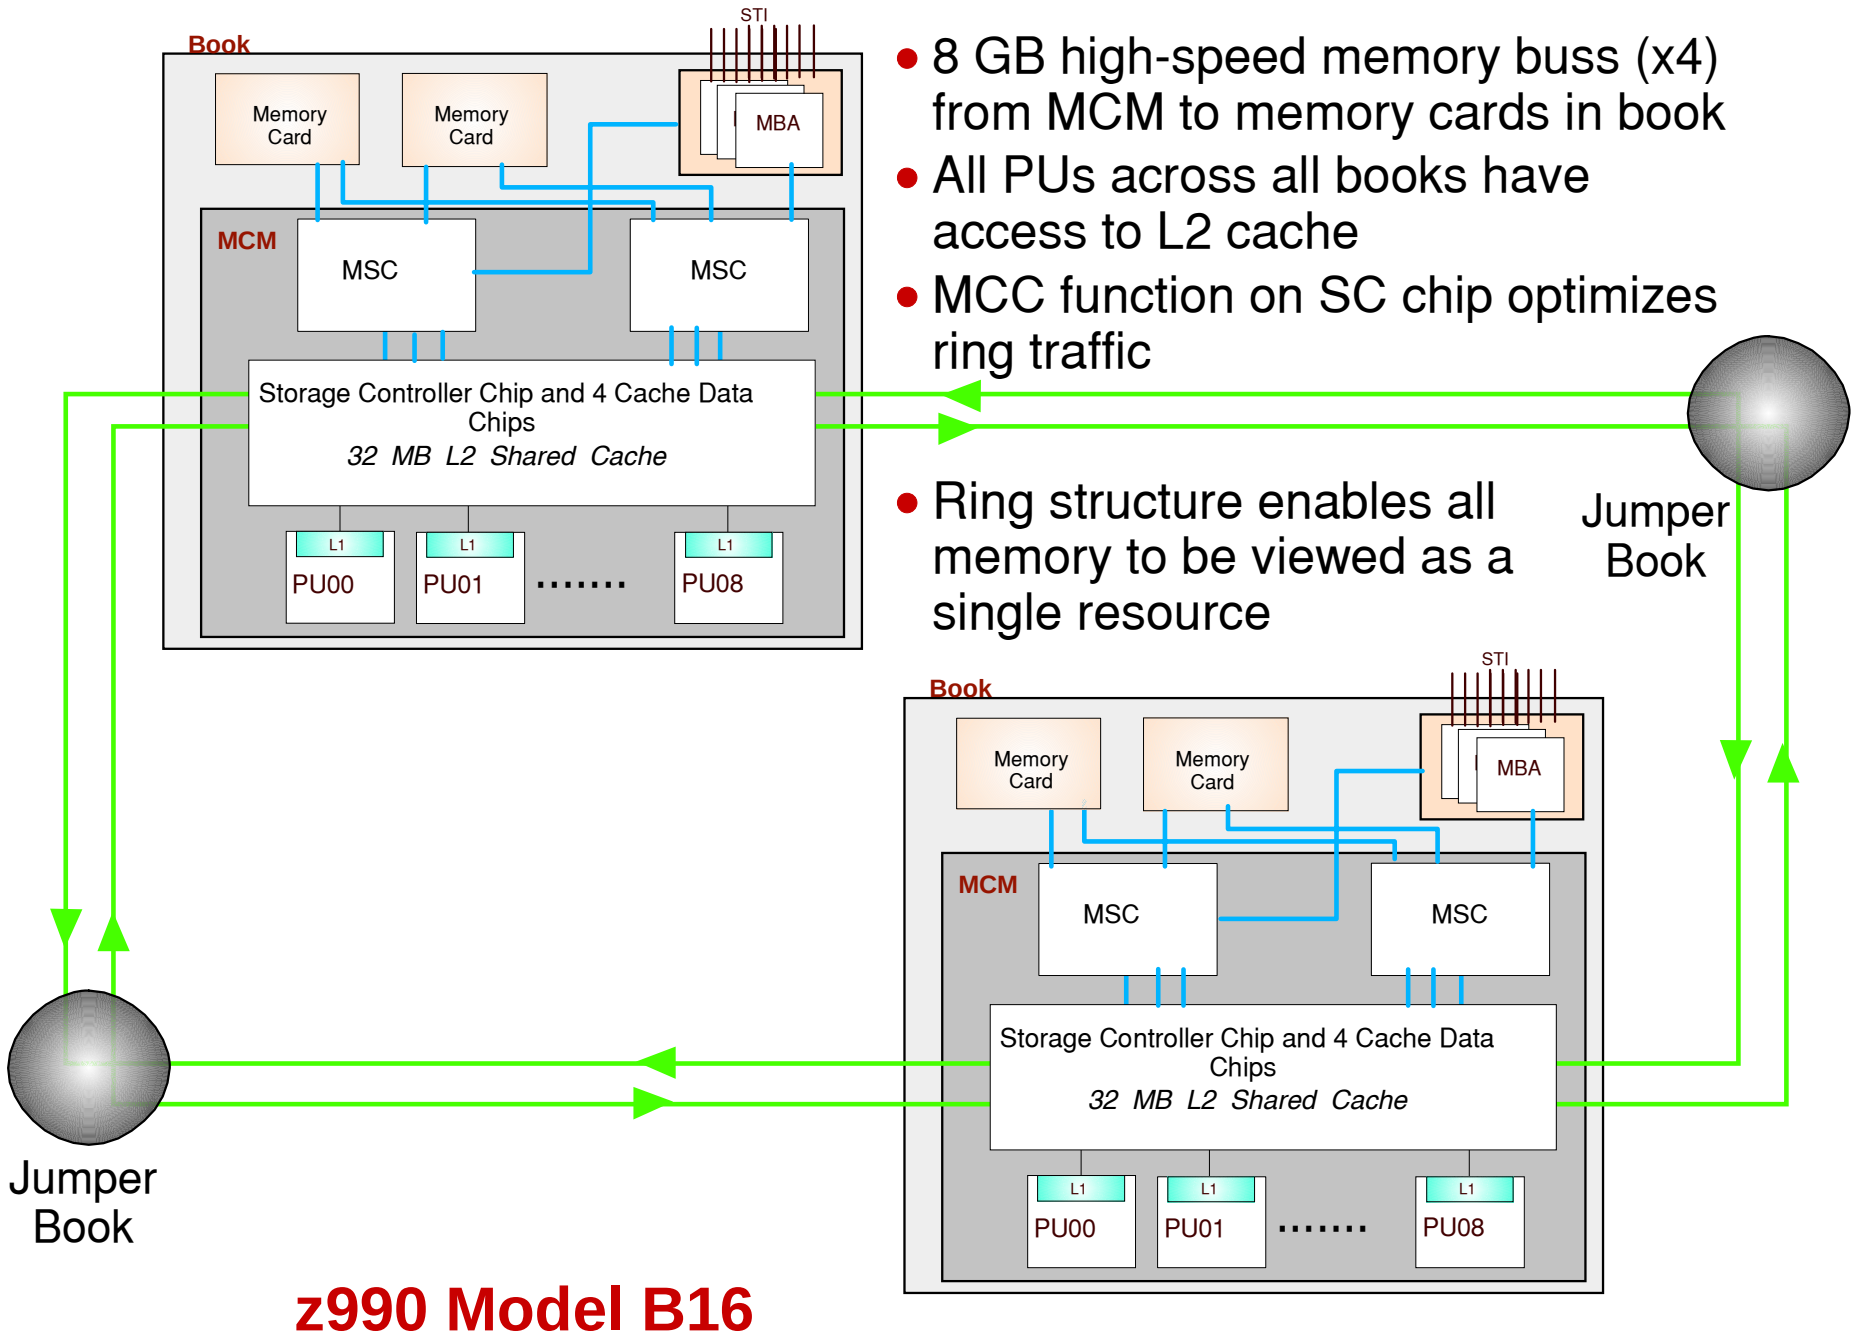

- 8 GB high-speed memory buss (x4) from MCM to memory cards in book
- All PUs across all books have access to L2 cache
- MCC function on SC chip optimizes ring traffic
- Ring structure enables all memory to be viewed as a single resource

z990 Model B16

zSeries Channel Connections - STIs

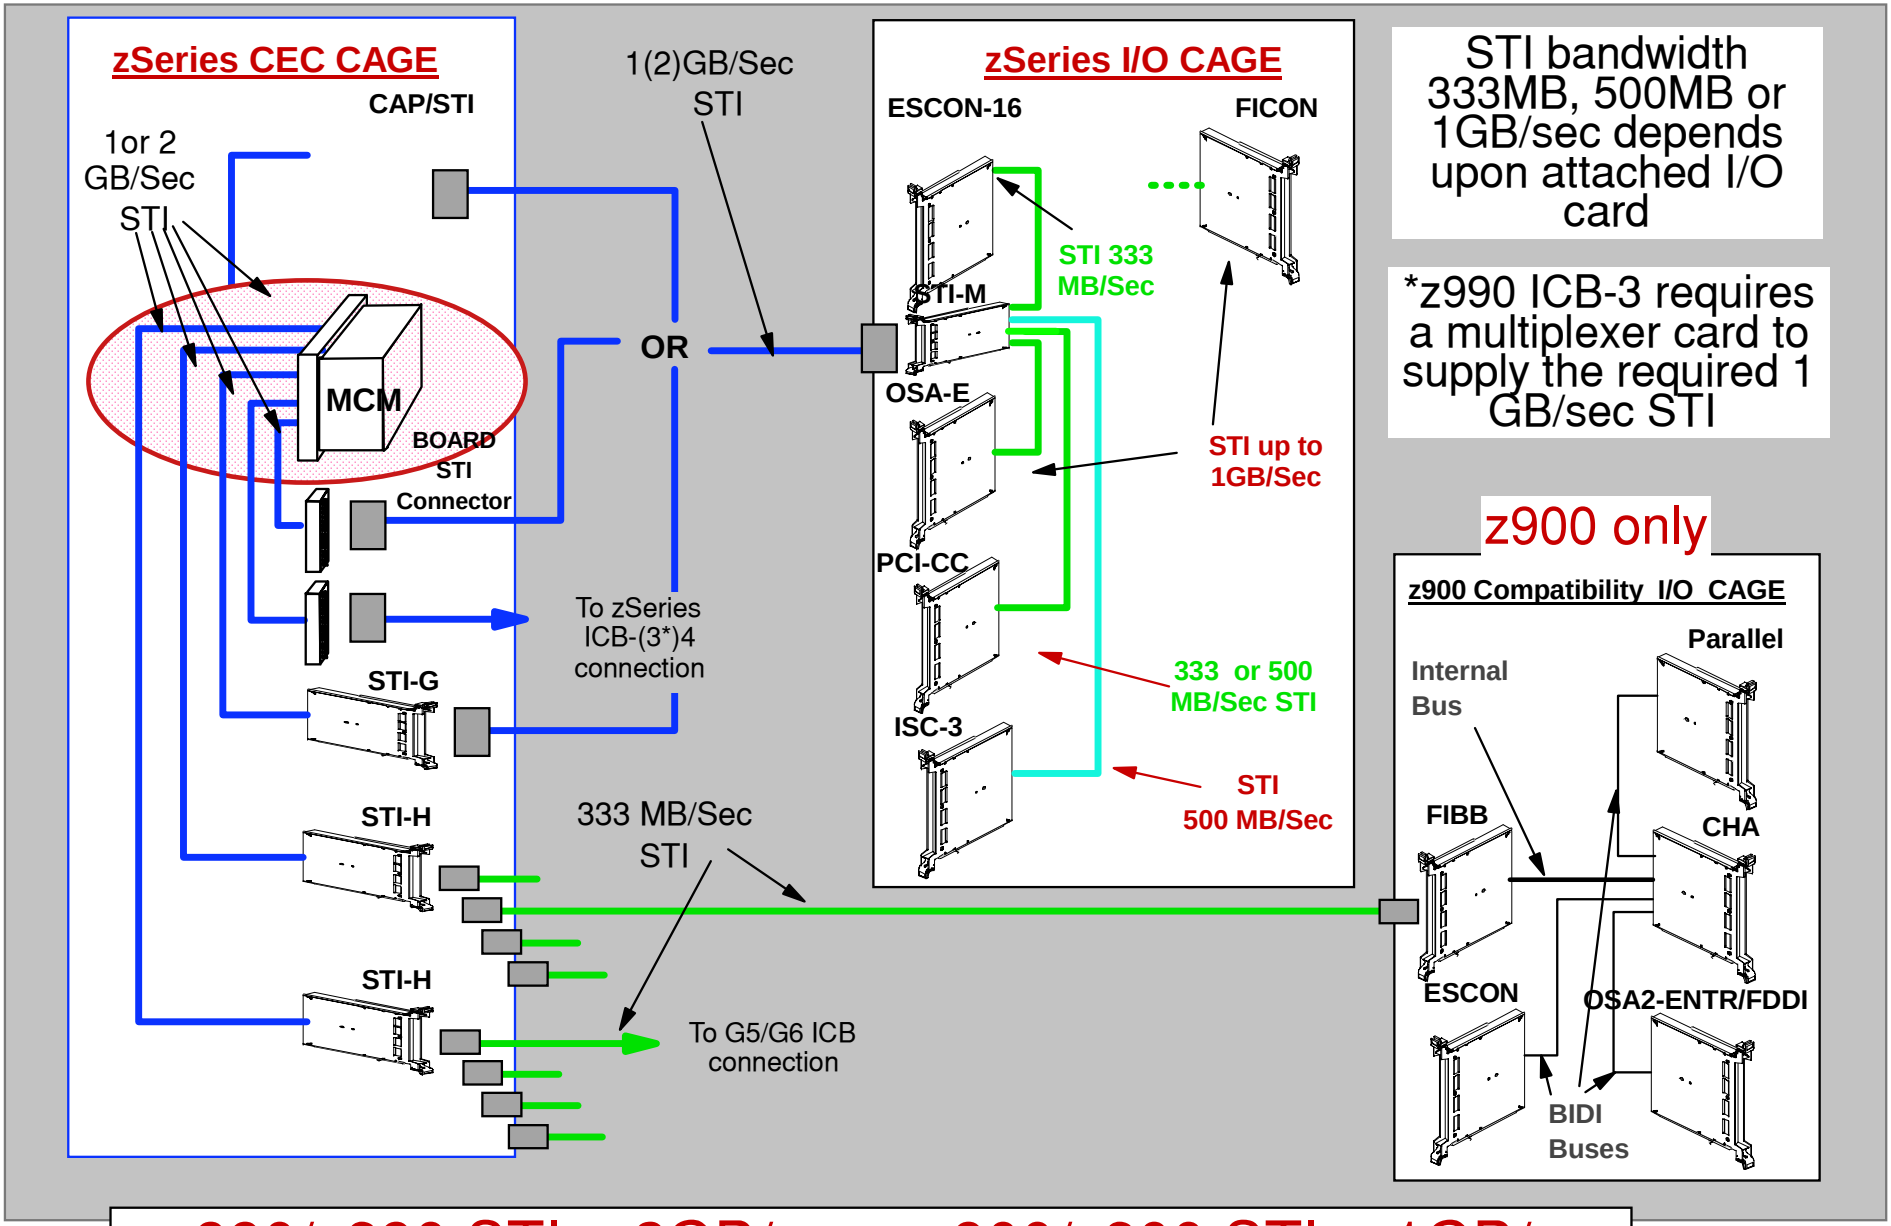

- z990/z890 STI = 2GB/sec, z900/z800 STI = 1GB/sec
- Only z900 supports I/O capatability cage

CPs, SAPs, STIs - Putting the Pieces Together

CPs, SAPs, cache, and I/O adapters via the STIs all work together to process instructions and I/O requests.

- CPs process instructions
- SAPS work with the CSS and processes I/O requests via the I/O adapters

CPC architecture in conjunction with the operating system provide

- Hardware registers and formatted areas of storage
  - The Hardware System Area (HSA) stores information that the CSS needs to process requests
  - The Program Status Word (PSW) register (one for each CP) contains information required for the execution of the currently active program.
    - Status, interrupts, instruction sequencing
    - Status of the CP can be changed by loading a new PSW or a PSW swap

PSW Swaps and Interrupts (Concept)

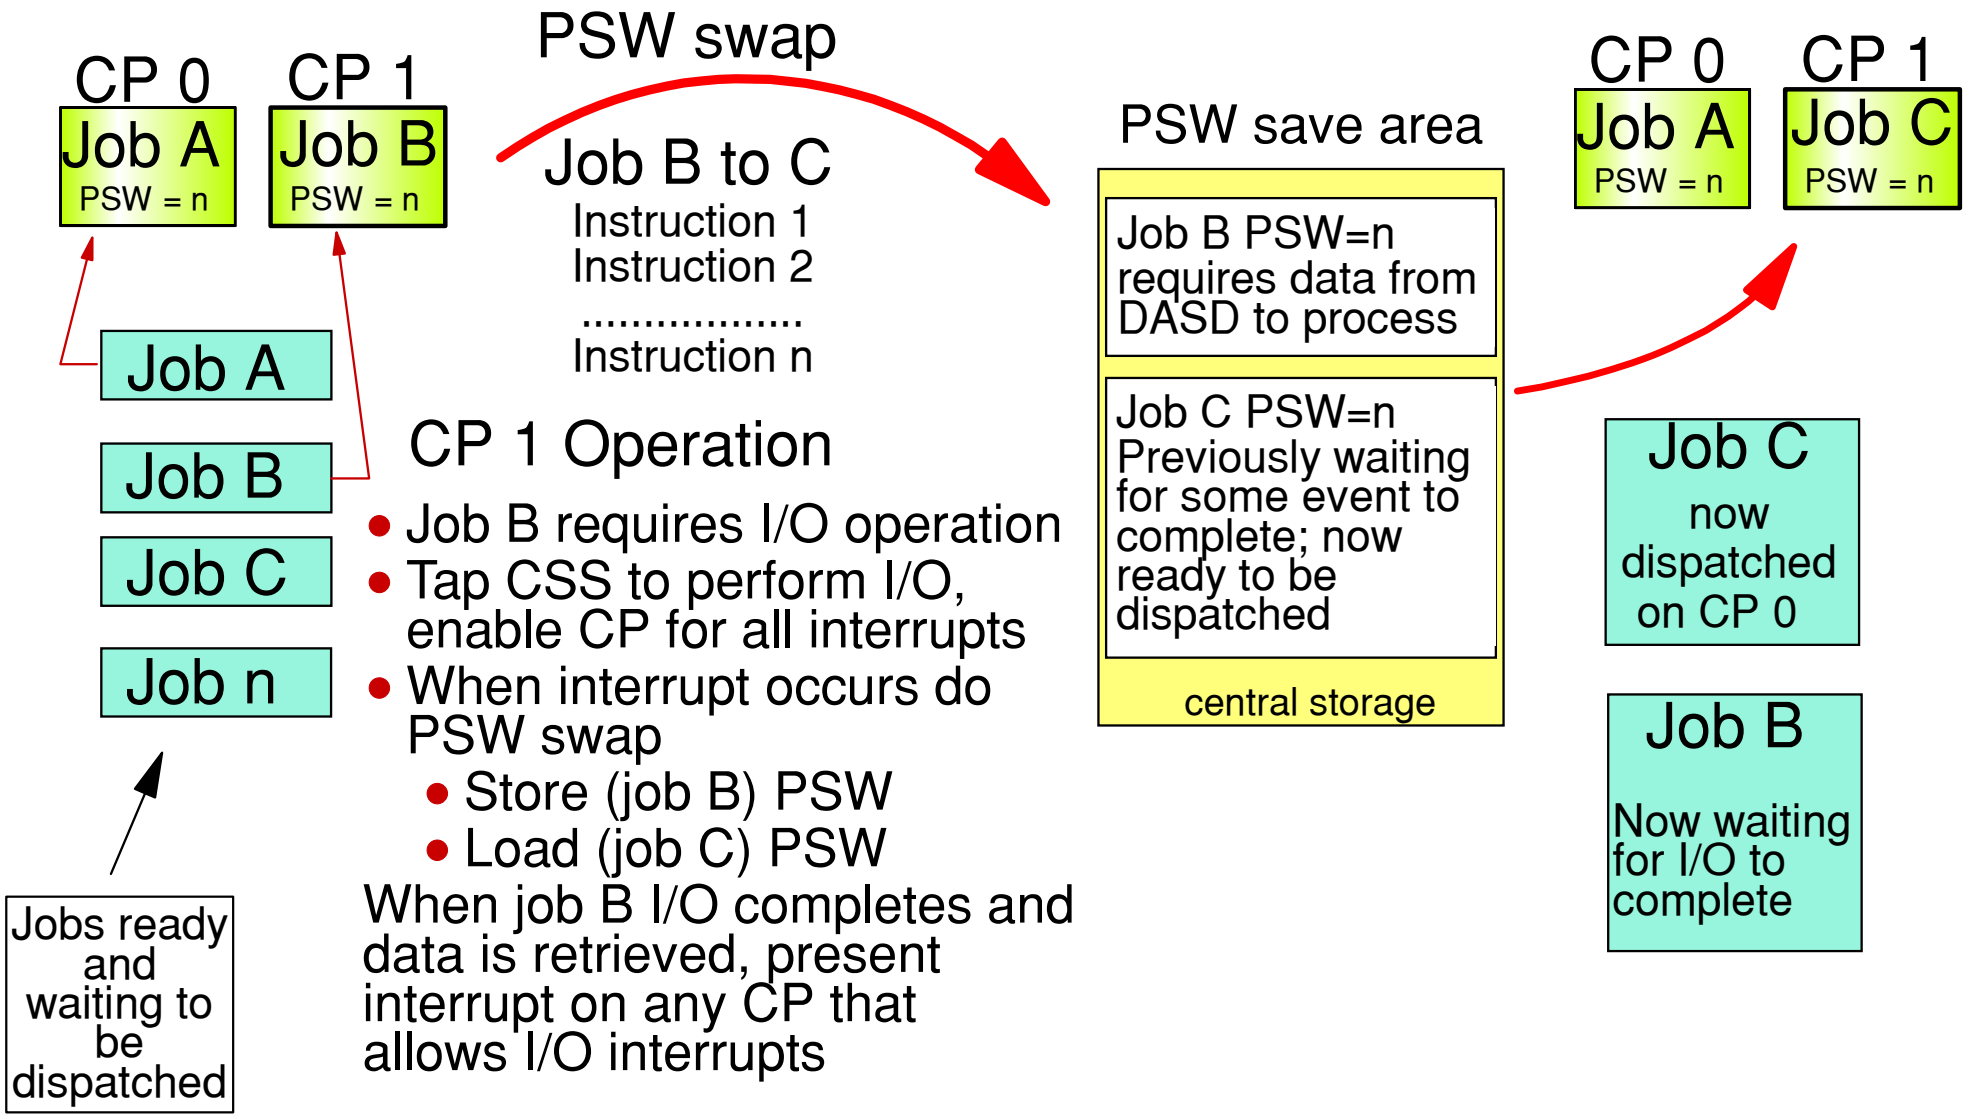

Program Status Word (PSW) - One for each CP, contains status information and next instruction address to be processed.

CSS I/O Operation Overview

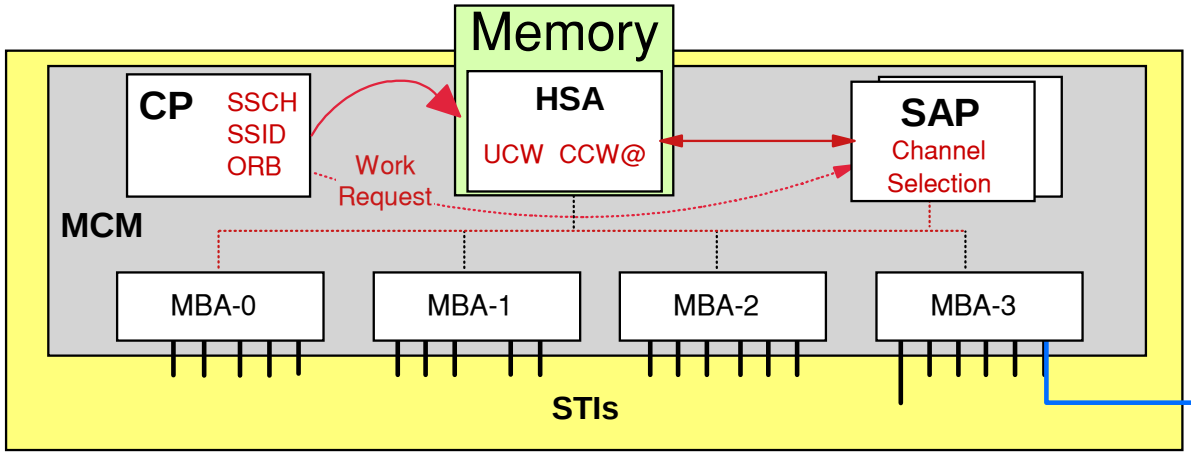

- The program working with the access method and IOS provides the channel program and other information in various control blocks. Start subchannel is issued to the CSS to start the I/O operation
- The CSS utilizes a SAP to perform all I/O functions. The CSS works with the subchannel stored in HSA.
- HSA contains reserved storage that is used for specialized functions. Subchannels used for channel operations contains status, channel paths, and other necessary information for I/O operations to a given device. There is one subchannel for every I/O device.
- CCWs and data is passed to the MBA, and exits the MCM through backboard wiring or external cables to the selected channel card. Connected to the channel cards are external fiber or copper cables.

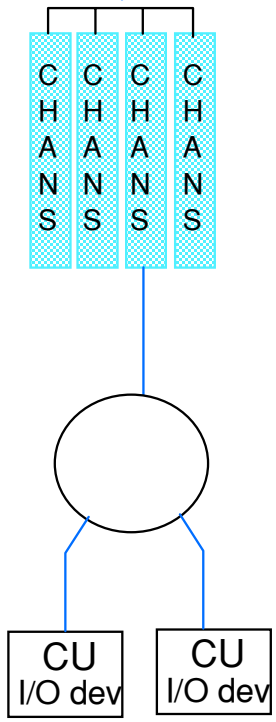

z/OS Operating Systems

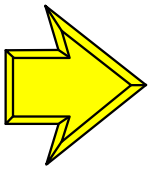

zSeries S/390 Server Hardware Overview

z/OS Operating Systems

Architectural Overview

From MVS to z/OS

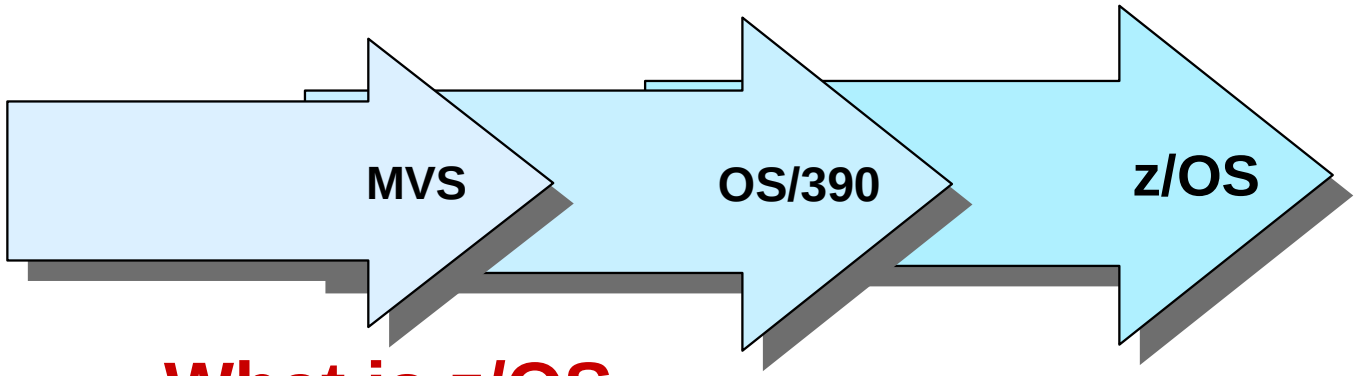

What is z/OS

- The ultimate mainframe operating system
- A packaging of over 70 different functions
  - Base operating system
  - Many industry exclusives:
    - Workload Manager, Parallel Sysplex, Intelligent Resource Director
  - Key e-business Services:
    - Networking, security, storage management, distributed files, and print
  - UNIX built right into the base
  - Exploiting the technologies of the IBM Eserver zSeries 800 and IBM Eserver zSeries 900 servers

z/OS V1 R1 Overview

- **z/OS V1 R1** (Announced Oct 2000) **for the zSeries z900**
  - Available March 2001
  - z/Architecture provides new platform function (zPlatform) that works in combination with z/OS
  - Major functions supported with z/OS V1R1 in addition to 64-bit addressing capabilities
    - ⋆ Intelligent Resource Manager
    - ⋆ Workload License Charges (WLC)
    - ⋆ MSYS for setup
  - Major functions to be supported in other z/OS releases for the **zSeries z900 and z800**
    - zSeries Entry License (zELC), HiperSockets, MSYS for operations, CF duplexing.....

z/OS V1 R2 to V1 R4

- **z/OS V1R2** (Aug 2001) for **zSeries 900**
  - Enhancements including MSYS for operations
  - **Project eLiza** (now Autonomic Computing) next stage of new tools for zSeries
  - **z/OS & *z/OS.e V1R3** (Feb 2002) for **z900 and z800**
    - z/OS.e, sometimes called z/OS lite, offers a reduced price OS for workloads that are consolidated from other non-IBM platforms; traditional workloads are restricted
- **z/OS & *z/OS.e V1R4** (Aug 2002) for **z900 and z800**
  - Enhancements in MSYS, networking IP protocol, PKI support and workload balancing in a sysplex.....
  - With z/OS1.4, z/OS is transitioning from a 6-month release cycle to an annual release cycle

*Note: z/OS.e is only available for zSeries z800 CPCs

z/OS V1 R5 to V1 R6

- **z/OS and *z/OS.e V1R5** (Mar 2004) for **z900/z800 and z990**
  - Multilevel Security
  - Enhancements to
    - Self-optimization of WebSphere applications
    - Backup and recovery of DB2 data
    - Performance for DFSORT
    - Intrusion Detection Services
    - Managing print across the enterprise
    - Last OS release to suport 9672 CPCs
- **z/OS and *z/OS.e V1R6** (Sept 2004) **z900/z800/z990/z890**
  - Support for more than 16 engines in a single image
  - 64-bit application development support for C/C++
  - Requires a zSeries CPC

z/OS.e Overview

- A specially priced z/OS offering available at a fraction of the cost of z/OS
- Runs on z800 or z890:
  - Engine-level pricing
  - Engine-level granularity
- For enterprise and e-business applications only, some limitations apply:
  - Some z/OS base elements and features are disabled
  - Cannot run traditional workloads such as CICS, IMS, COBOL, and FORTRAN.
- Same code base as z/OS:
  - Unique Program Number for Ordering

z890

z800

z/OS.e Workloads

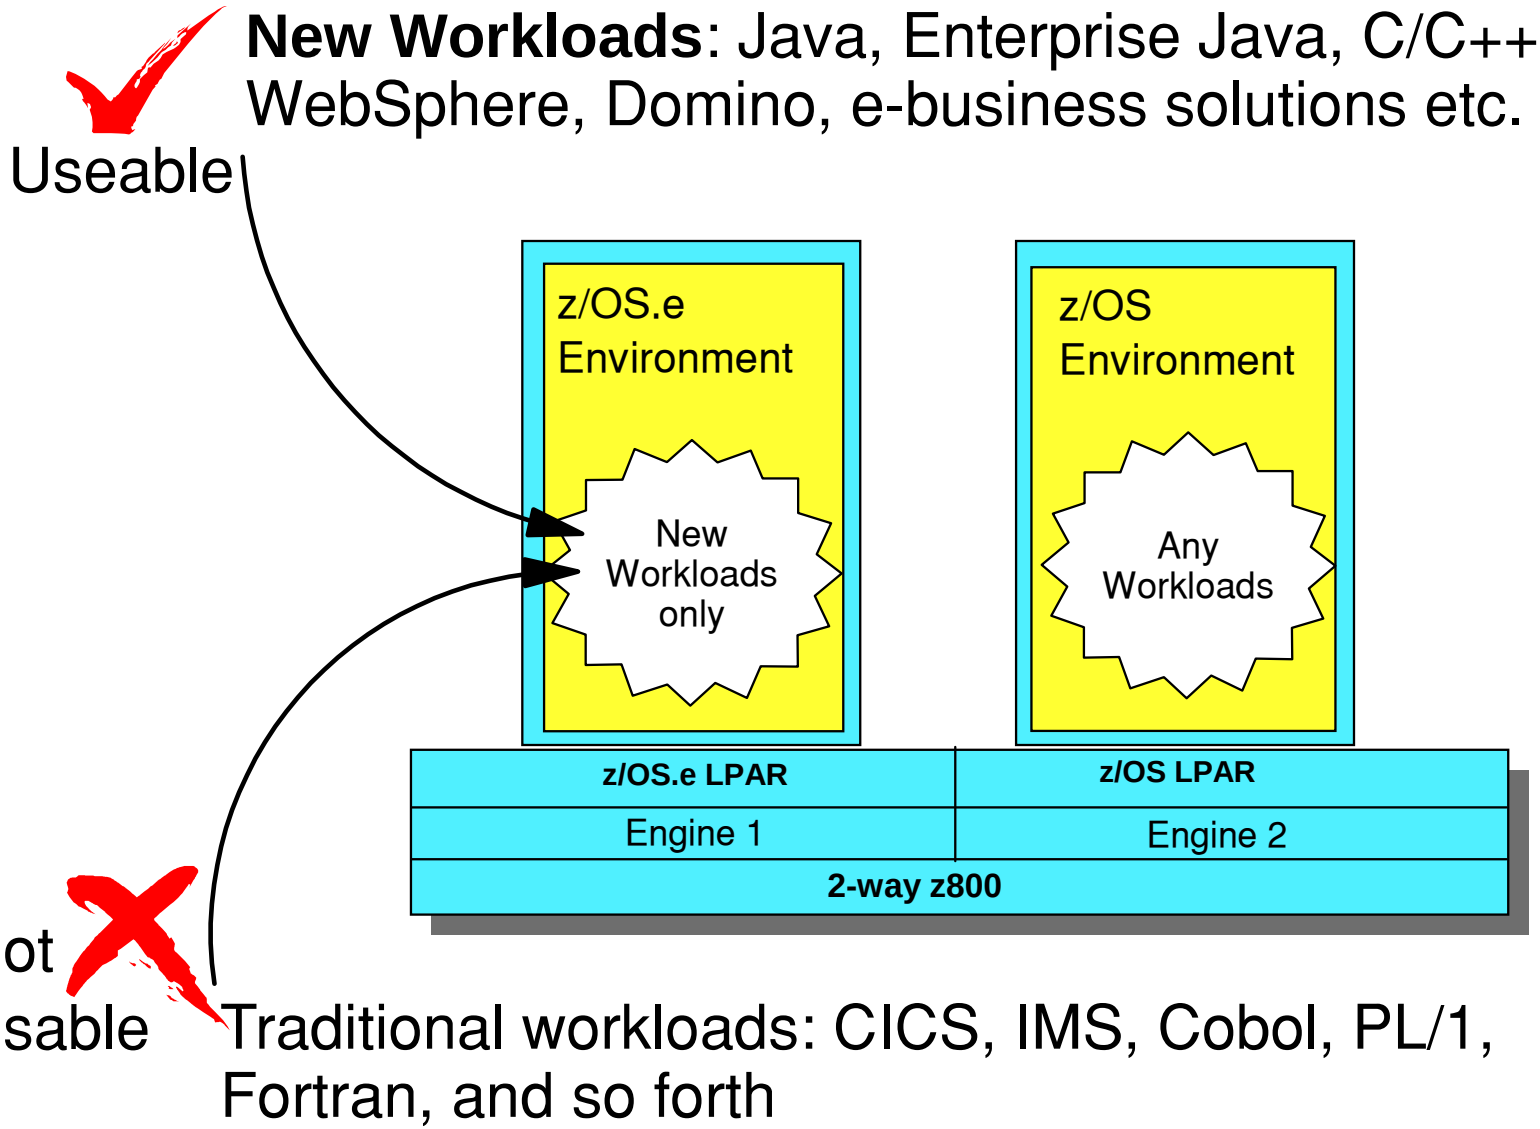

OS/390 and z/OS Product Release Cycle

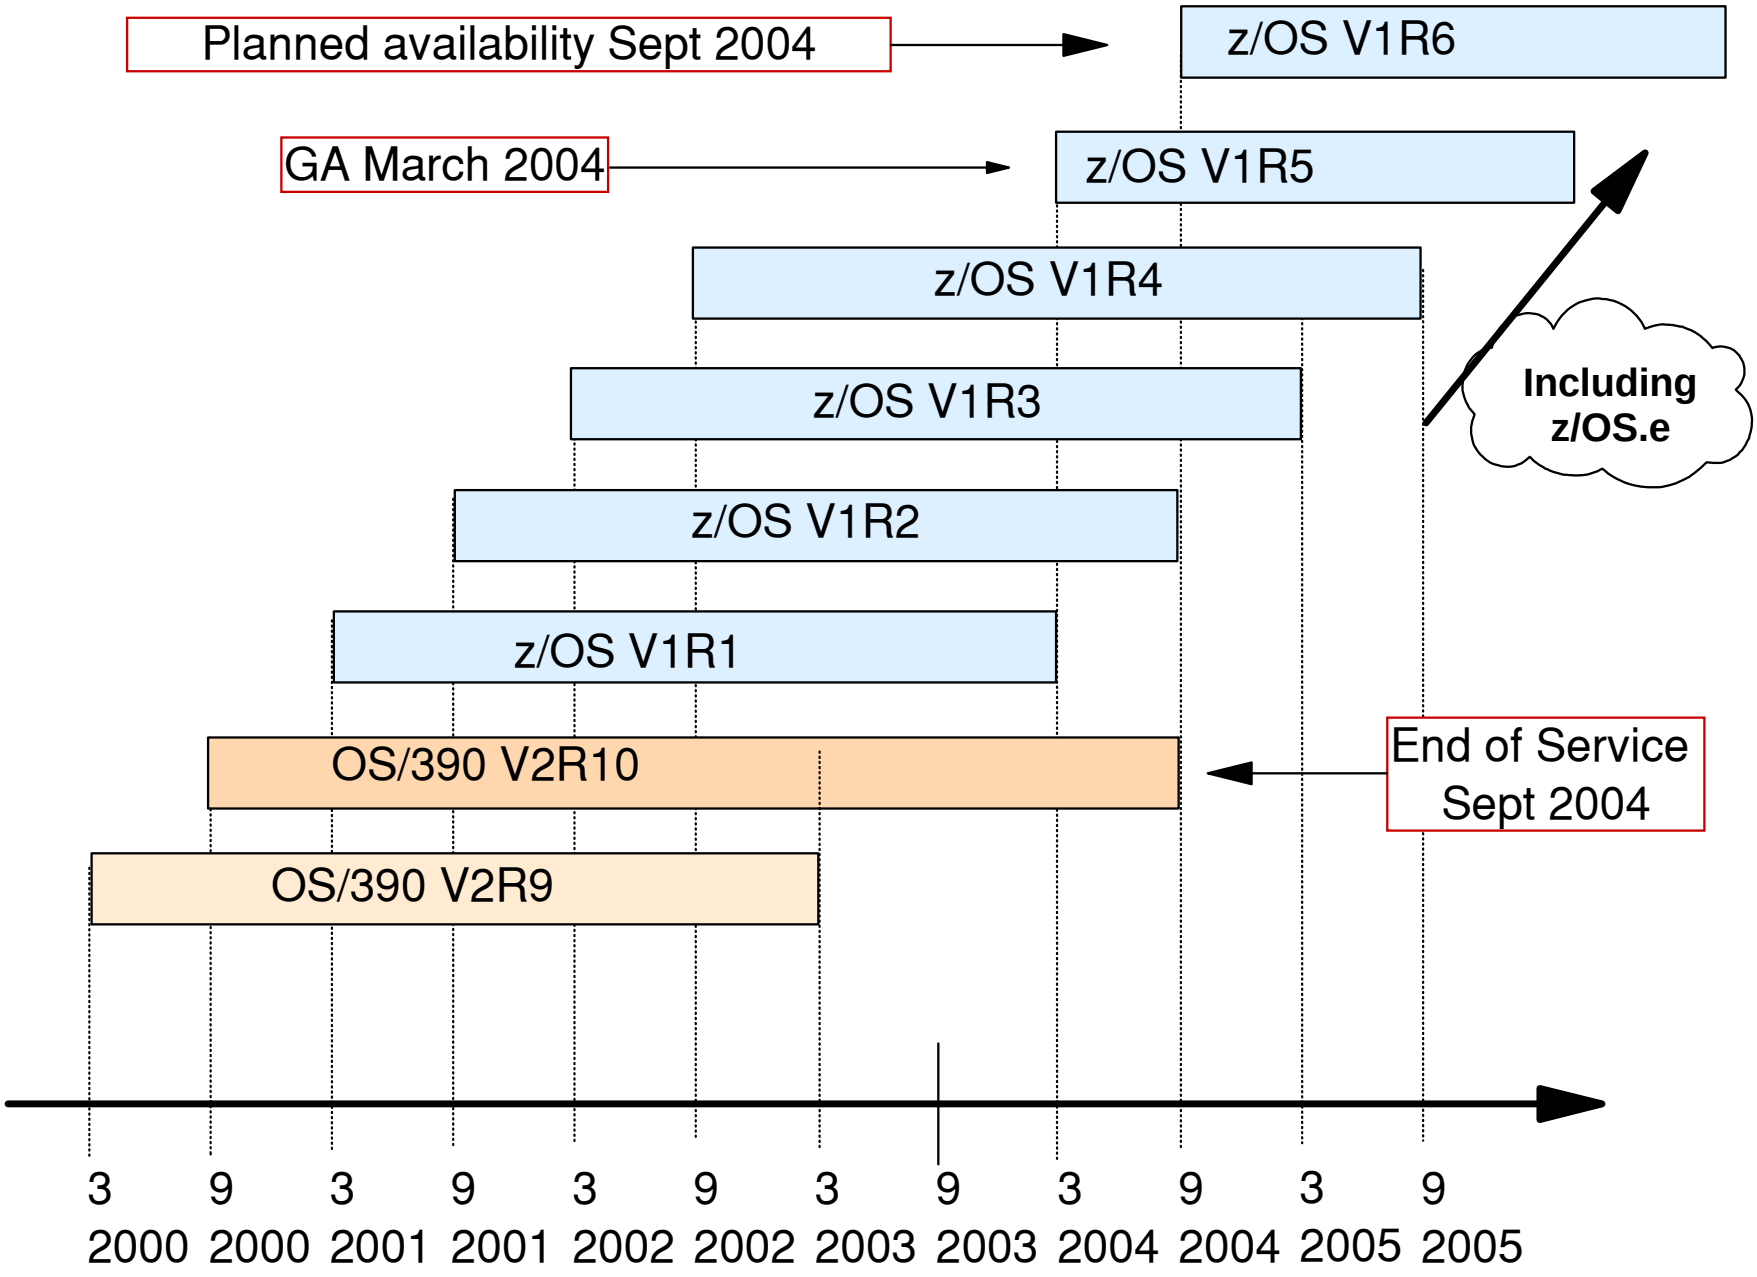

OS/390 and z/OS Hardware support

IBM S/390 Multiprise 3000

IBM S/390 Parallel Enterprise Server G5/G6

S/390 Servers

OS/390 V2 R10 to z/OS V1 R5 only

- ESA/390 mode (31-bit)
- z/OS V1 R6 and above NOT supported

zSeries

zSeries 800 / 900 Servers

OS/390 V2 R10 to z/OS V1 R5 and above

- ESA/390 mode (31-bit)
- z/Architecture mode (64-bit)

zSeries 890 / 990 Servers

OS/390 V2 R10, z/OS V1 R2 and above

- ESA/390 mode (31-bit)
- z/Architecture mode (64-bit)

*z/OS on zSeries intended to run in z/Architecture only

z/OS.e V1R3 through R6 are supported on z800 /z890 servers, in z/Architecture mode only (64-bit mode)

Architectural Overview

zSeries S/390 Server Hardware Overview

z/OS Operating Systems

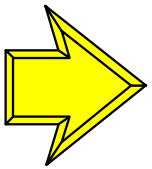

Architectural Overview

Architecture Overview

Today, two architectural modes exist, ESA and z/Architecture OS/390 and z/OS software support these modes. Both OS's use a wide range of H/W registers and control blocks that describe the available resources and addressing available for the supported mode.

Some H/W registers

- PSW - Program Status word
  - Controls instruction sequence and state of CP
- GR or GPR - General Purpose Registers (16)
  - Contains data used by instructions, arithmetic function, accumulator, special purpose
- CR - Control Registers (16)
  - Maintains and manipulate control information outside of PSW
- AR - Access Registers (16)
  - Used for dynamic address translation (DAT) when PSW is in access register mode

Storage and Control Blocks

- The OS uses areas of storage in predetermined formats, control blocks, that applications can exploit.
- The application runs in a virtual address space. The address space is divided into private and common areas which also contain control blocks like the CVT and PSA
  - Address Space - contains the range of virtual storage locations that can be addressed
  - CVT - Communications vector table, contains the address of the next control block to be used when a chain of control blocks are to be constructed
  - PSA - Prefixed Save Area, critical information stored here, such as location of new PSWs to use during interrupts

z/Architecture (New Extensions)

- PSW expanded to 16 bytes

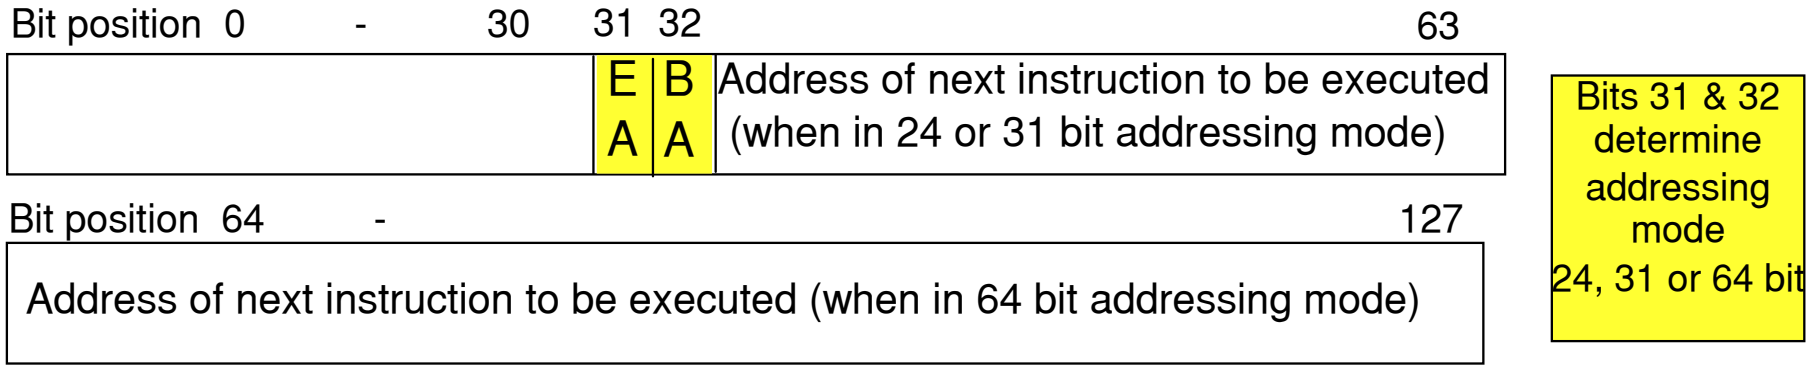

- 64-bit general purpose registers and control registers

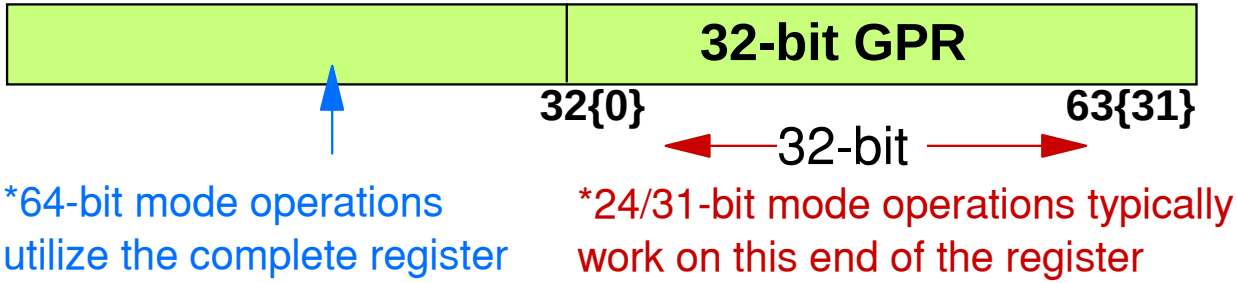

- Up to three additional levels of DAT, called region tables
- 8 KB prefix area for PSWs and register save areas
- A SIGNAL PROCESSOR order to switch processor modes at IPL

Dynamic Address Translation Region Tables

The ASCE describes an address space

- A virtual space is described by translation tables
- With a real space the virtual addresses are translated to the same real address; there are no translation tables

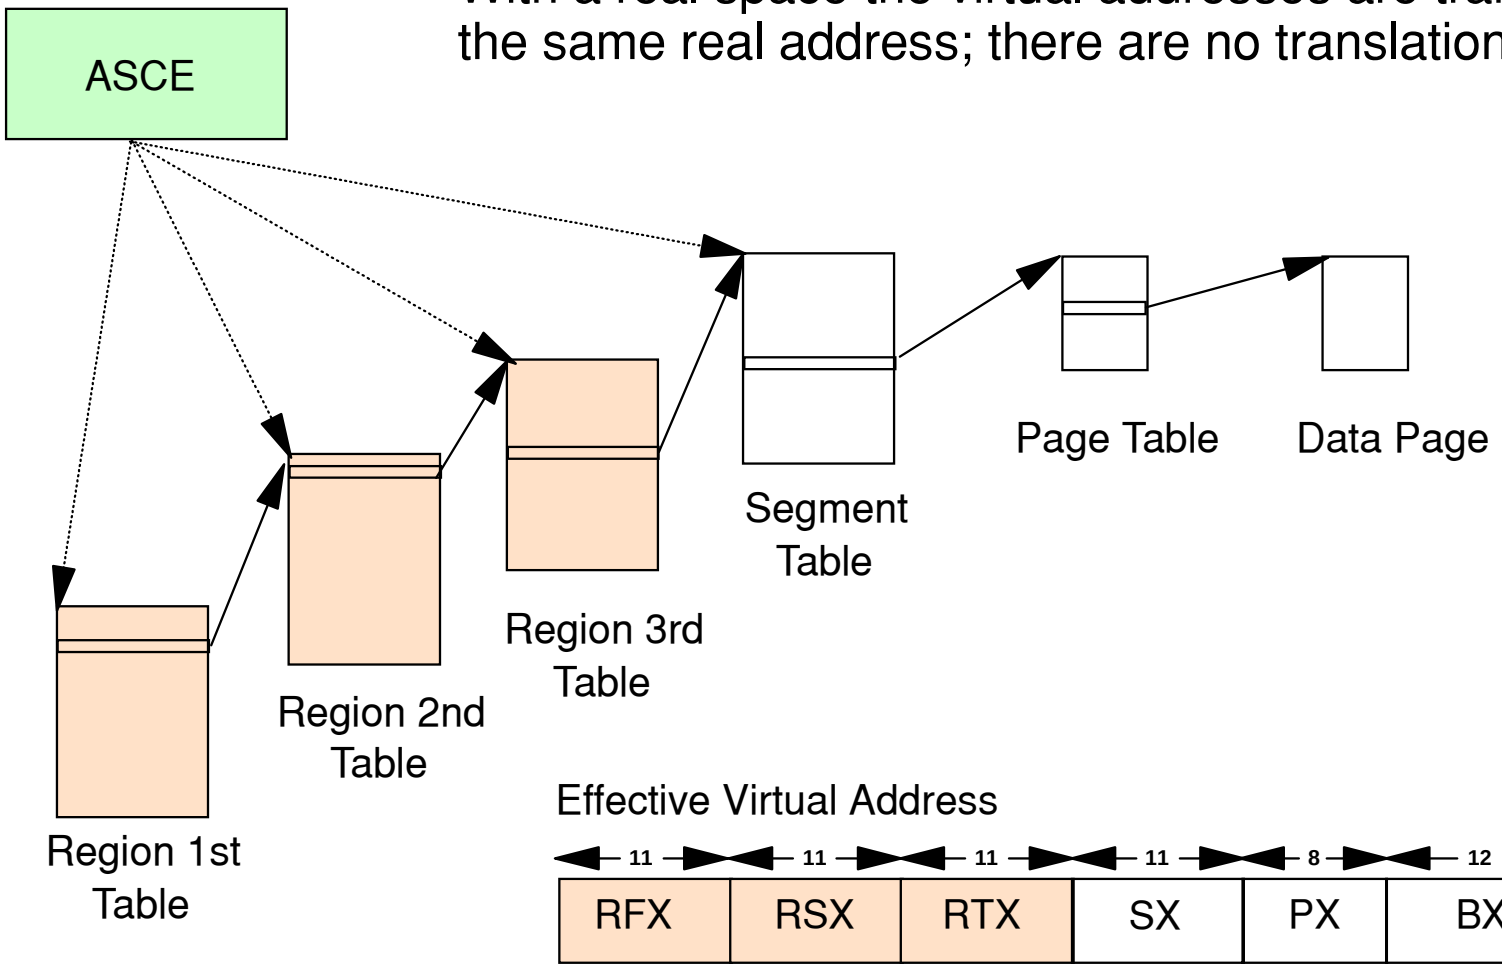

Enterprise Server Storage, Real and Auxiliary

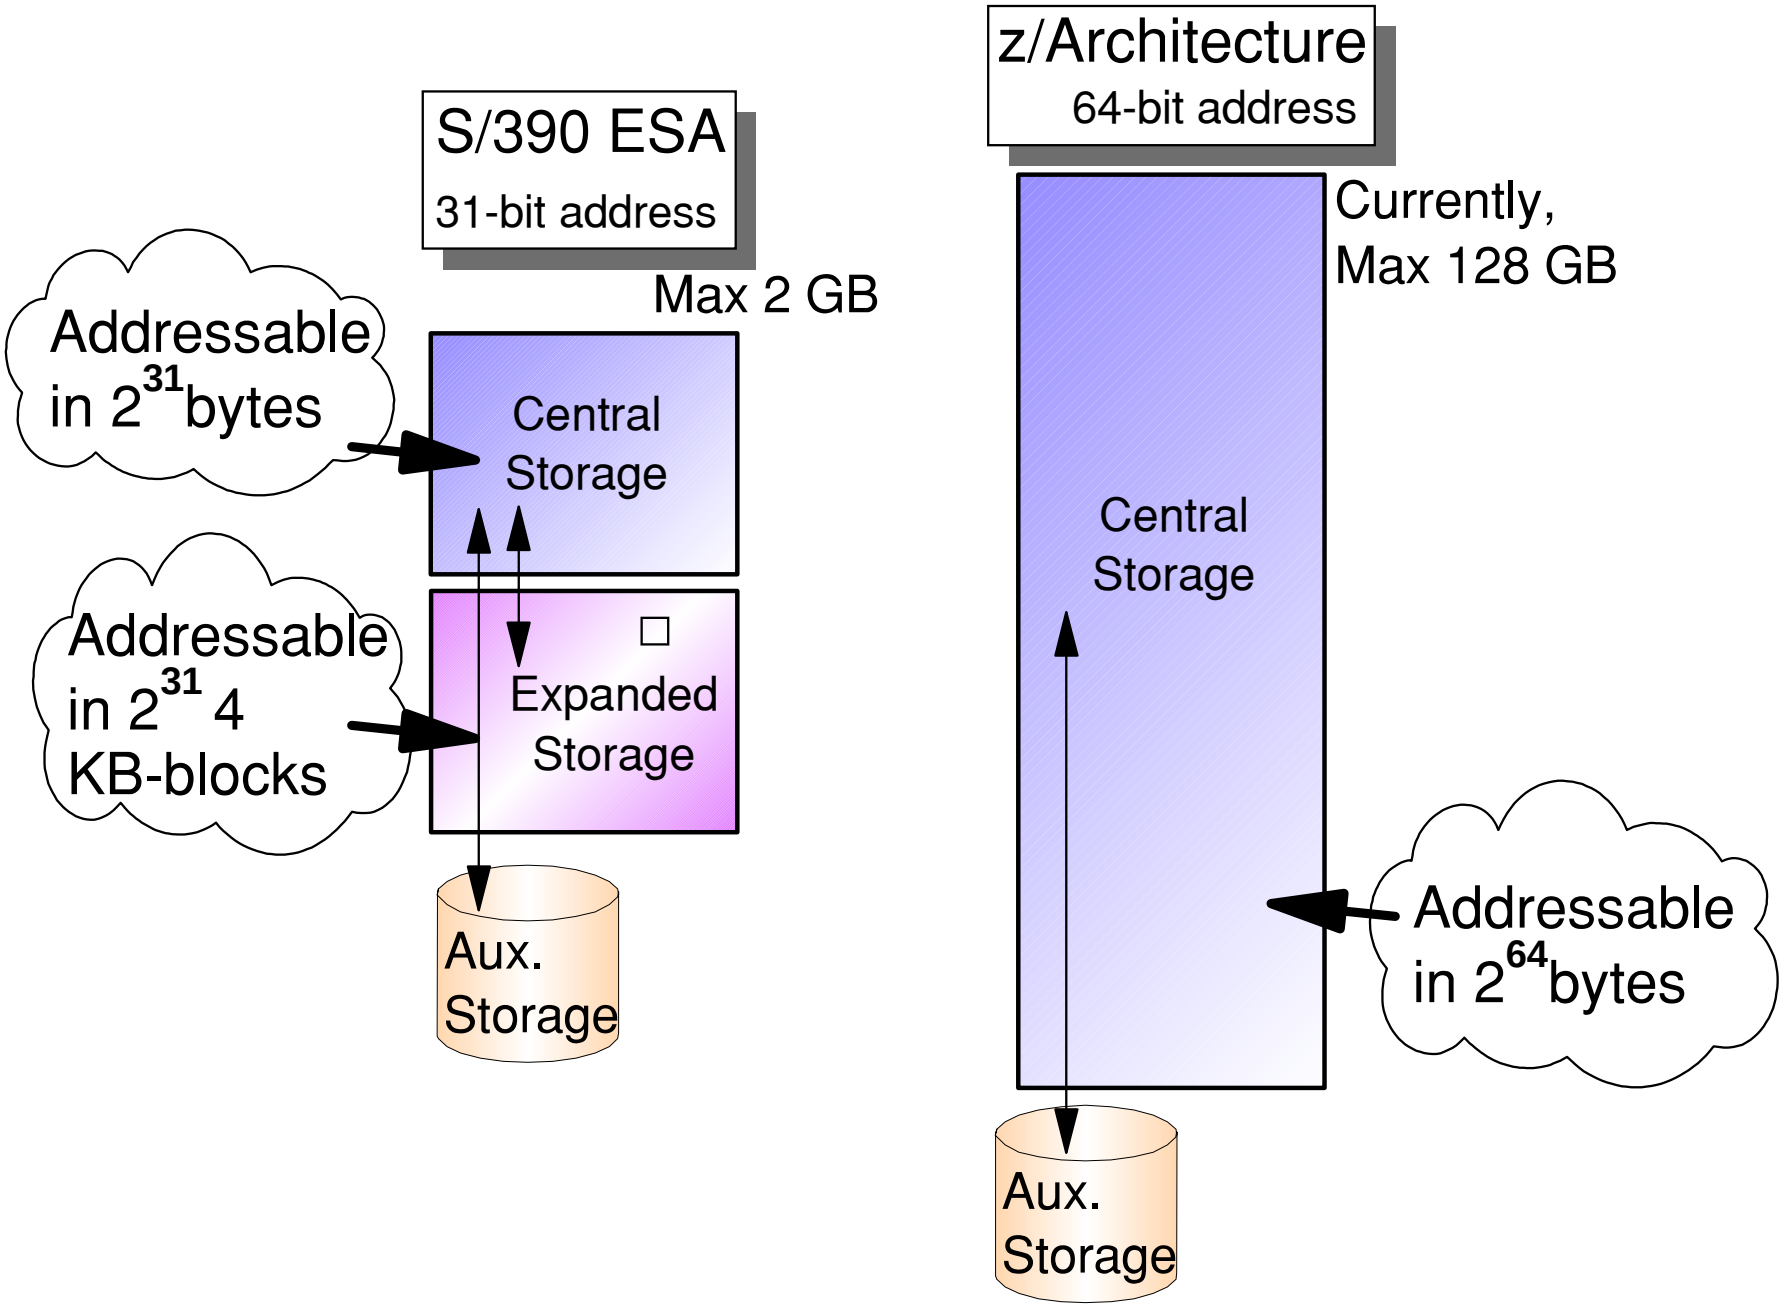

Enterprise Server Storage and Virtual

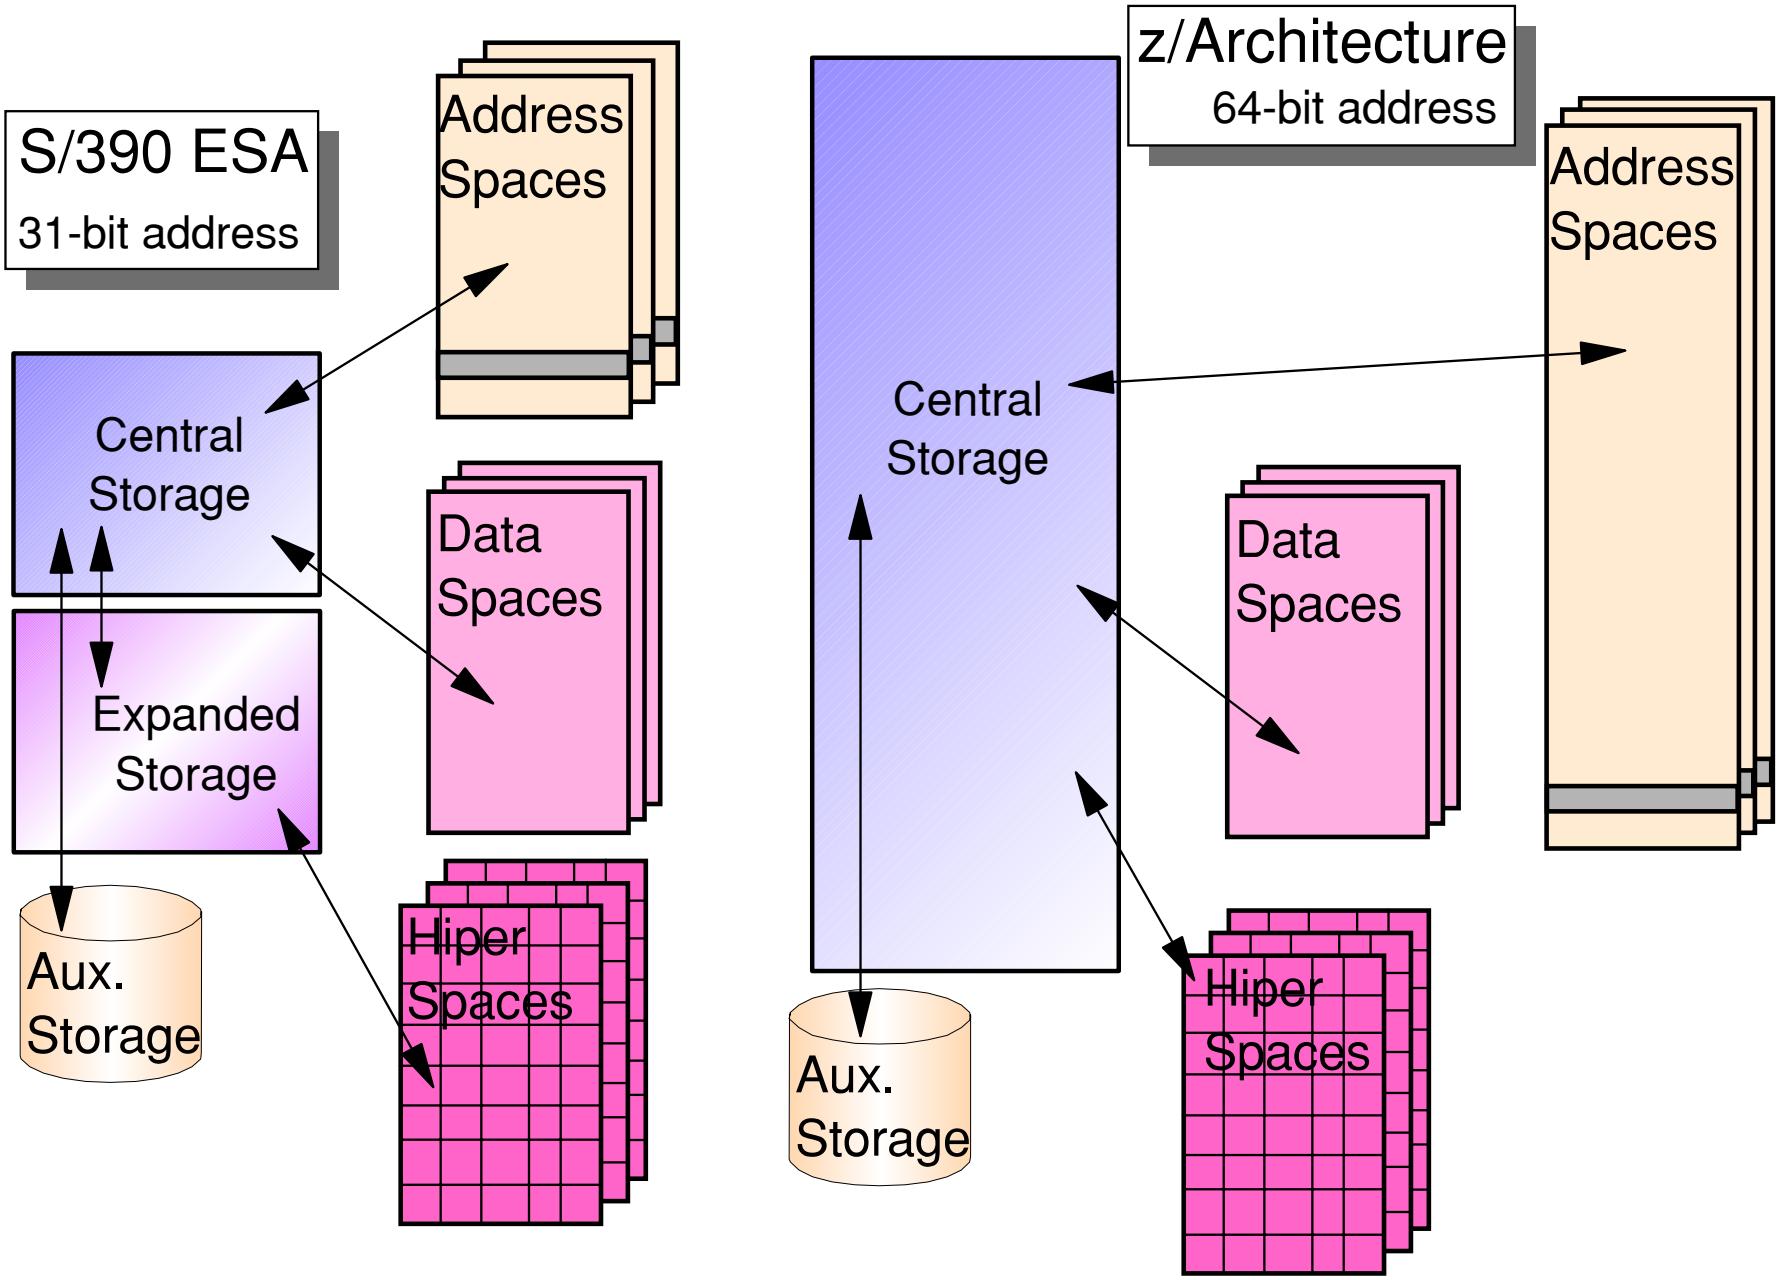

Virtual Storage Usage

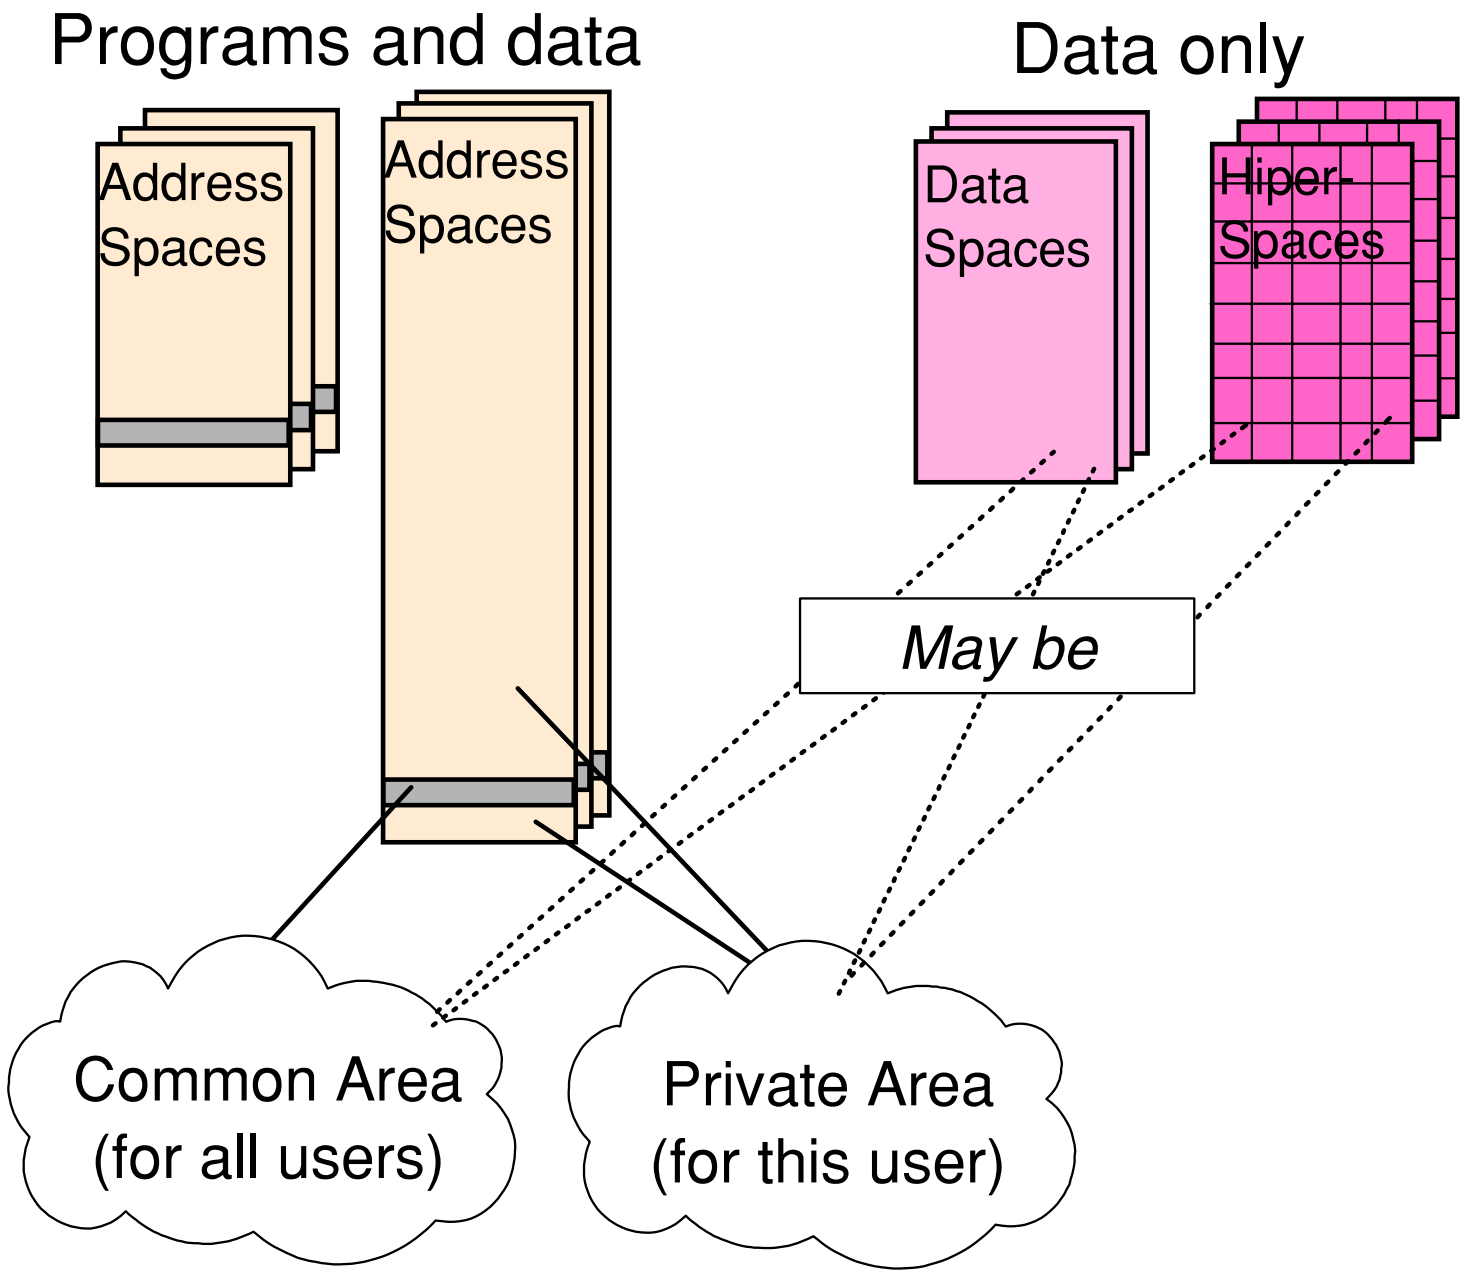

Address Space Layout

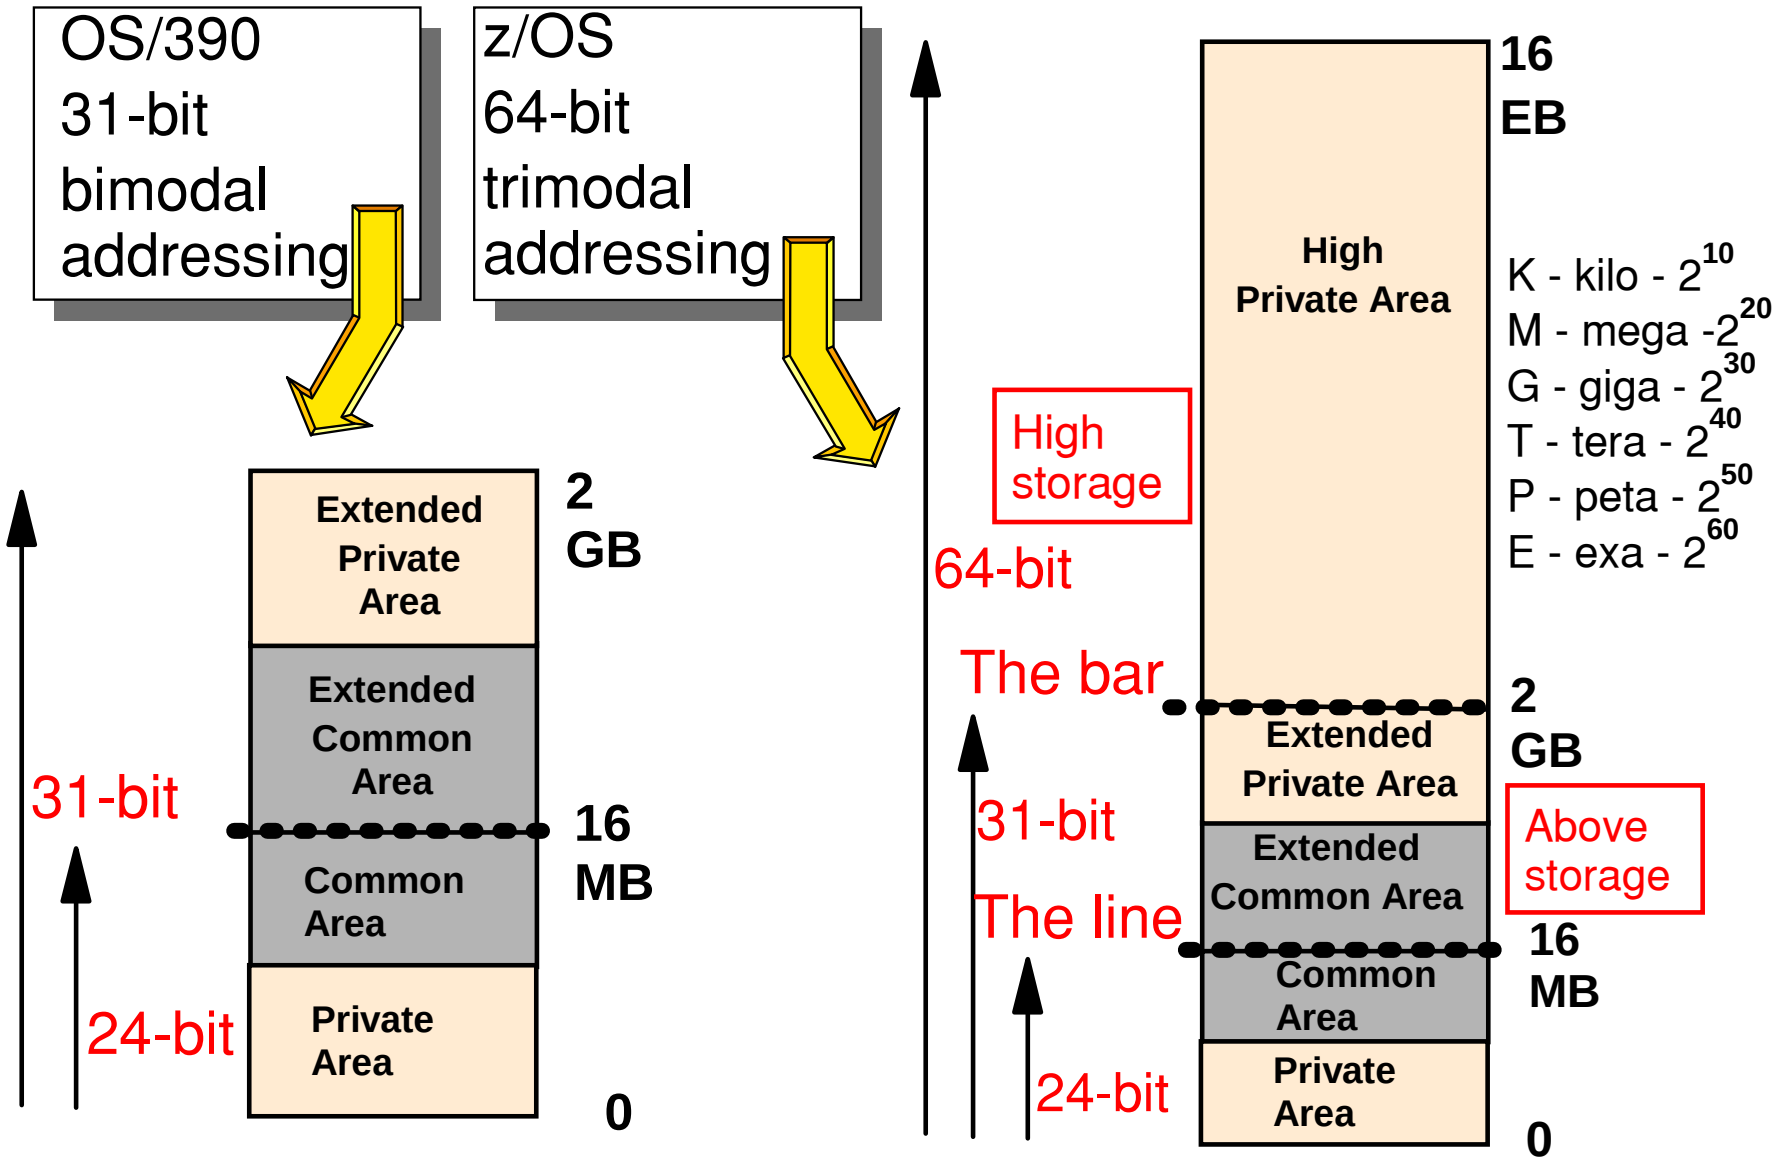

Note: Not to scale. 8+ billion traditional address spaces in one 64-bit address space. (billion = million million)

zSeries Educational Offerings

Sysplex / zSeries course offerings

- H4016 (2 days) HMC Class
- H4041 (3) Plex Ops & Recovery (sysplex only)
- H4057 (5) Plex Ops & Recovery (H4016 & H4041)
- ES902 (5) Advanced Plex Recovery
- ES420 (5) Sysplex Implementation
- ES830 (5) CSAR (Complex Systems Availability & Recovery)
- ES820 (2) zSeries Mainframe Environment (A Technical Overview)
- OZ09 (2) z/Architecture for z900
- OZ05 (2) z990 Technical Update & Configuration Requirements
- ES321 (2) FICON Environment (Native & Brigde)
- ES960 (4) HCD and Dynamic I/O
- ES270 (3) z/OS and OS/390 System Operations

RETURN TO INDEX