B01

zMainframe Concepts (The Big Picture) - Part 1

Brian Hatfield

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

Orlando, FL

RETURN TO INDEX

Session Objectives

In this session we will Discuss

The IBM mainframe evolution - System z Concepts and Overview of the Server design

- What is a Mainframe
- The Mainframe difference
  - engine and channel subsystem usage
- Current System z models and comparisons
- High level Server design and operation
  - Operating System to Channel Subsystem relationship

IBM System z9 and zSeries Mainframe Concepts and Server Overview

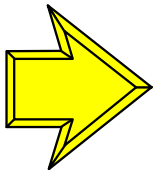

System z9 and zSeries Mainframe Concepts and Server Overview

IBM - Forty Plus Years of Mainframe History

From 1964 to today, the mainframe is alive and kicking

- 1964: The revolutionary IBM System/360 is born.
- By 1970 more than 3,000 different types of businesses and scientific research make use of one of System/360's models.

The IBM System 360 is considered the first General Purpose Business Computer

- April 2006: IBM announces the new IBM System z9 Business class server.
- The System/360's tradition of encompassing every user's every need is continued in the z9 family of servers.

System z9

The Mainframe is Dead or is It?

Haven't you heard the Mainframe is Dead?

1980s - Due to distributed processing, several analysts predict an end to the glass house - The mainframe is dead!

1990s - Several times, early 90s and late 90s, due to many factors, but mainly the rapid growth in PCs and small servers - The mainframe is dead!

"I predict that the last mainframe will be unplugged March 15 1996"
- Stewart Alsop, Infoword 1991

We're baaaack..... Actually we never left

See "Mainframe Mythology" - The Clipper Group Report #TCG2006038R May 2006, identifies 20 mainframe myths and facts

How is the Mainframe Defined

- A very large and expensive computer capable of supporting hundreds, or even thousands, of users simultaneously
  - http://www.webopedia.com

However, the trend in the last several years;

Hardware costs 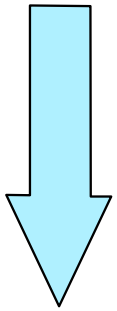

Significantly decreased with each mainframe family - various engine capacities

Software costs 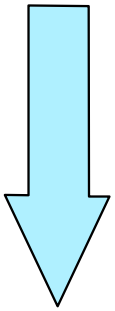

Multiple pricing offerings by IBM and other vendors

RAS H/W & S/W 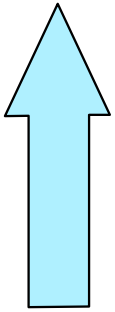

Additional optimization through systems management extensions

Today's mainframe, System z servers use their resources very efficiently, exploiting multiple engine types to process workload and manage input / output operations

System z servers refer to both System z9 family and zSeries family of servers

IBM Mainframe concept (Big Picture)

IBM mainframes have two independent but closely related functions

A logical partition contains the OS which works with application programs. The OS knows device numbers (but not I/O configuration), it works with processors and memory

The CSS can operate independent of OS. Contains I/O configuration data unknown to the OS, has access to memory and it's own processor

The Operating System
(Multiple OSs are possible)

Runs various application programs

Exploits various processor types and usage

Turns over I/O operation to the channel subsystem

The Channel Subsystem
(associated with all I/O adapters)

Works closely with the OS manages I/O operations

Has its own dedicated engine (SAP) and operates independently of the OS engines

Mainframe Differences

- Reliability Availability Serviceability (RAS)
- Reliability Availability Scalability (RAS)
  - Vertically
  - Horizontally
  - on demand
- Security
- Engine Versatility
- Engine Usage
- Resource sharing
  - Resource reallocation
  - Autonomic
- Virtualization
- I/O Bandwidth
  - Extremely large number of available channels
  - Massive amounts of data that can be moved

System z Engines (Processor Units)

Mainframes typically have more engines physically available then are used in most configurations

Mainframe engines (or PUs) are very versatile and can be assigned (characterized) as the following:

- A Central Processor (CP) - ***OS***
- A System Assist Processor (SAP) - ***CSS***
- An Internal Coupling Facility (ICF) - ***CF***
- An Integrated Facility for Linux (IFL) - ***Linux***
- System z Application Assist Processor (zAAP) - ***JAVA***
  - System z9 and zSeries 990/890
- System z9 Integrated Information Processor (zIIP) - ***DB2***
  - System z9 only

The IBM General Purpose Business Computer (The Mainframe)

A key difference between Mainframes and other servers is the amount of data that can be imported/exported to external shared storage devices while the operating system(s) maintain high performance levels processing other workloads

The OS and the CSS work together to efficiently manage data

I/O Bandwidth is the maximum amount of I/O (data measured in Giga Bytes per second) data that a Server can potentially achieve

9672- 8 GB/sec z900- 24 GB/sec z990- 96 GB/sec z9- 172.8 GB/sec

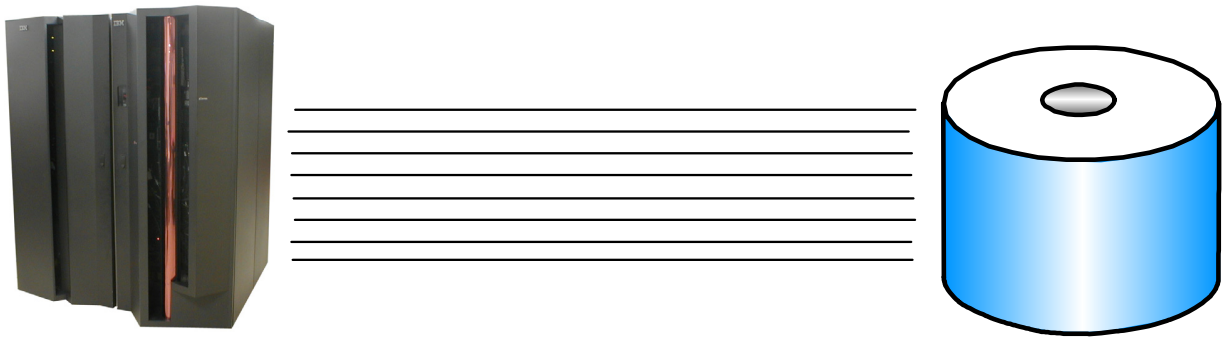

Mainframe's architecture is designed for efficient data movement
Has multipath capability (up to 8 to an LCU) with up to 256 channels per OS

Common Terms

CMOS - Complimentary Metal Oxide Semiconductor

SE - Support Element

HMC - Hardware Management Console

IBM System z Servers

- Includes System z9 and zSeries servers

CPC - Central Processing Complex
CEC - Central Electronic Complex
CPU - Central Processing Unit
Server or Mainframe or Host

Current Mainframe OS systems
z/OS, TPF, VSE, z/VM, Linux

z/OS is often referred to as **MVS**
MVS → OS/390 → z/OS

IBM mainframe architecture

- ESA- Enterprise System Architecture (31-bit addressing)
- z/Architecture (64-bit addressing)

PU = Processing units

PUs can be assigned:

- Central Processor (CP)
- System Assist Processor (SAP)
- Other specialty processors

Channel related terms

SAP(s) is used by the CSS
CSS - Channel Subsystem
LCSS - Logical Channel Subsystem
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

zSeries Hardware Technology (1 of 2)

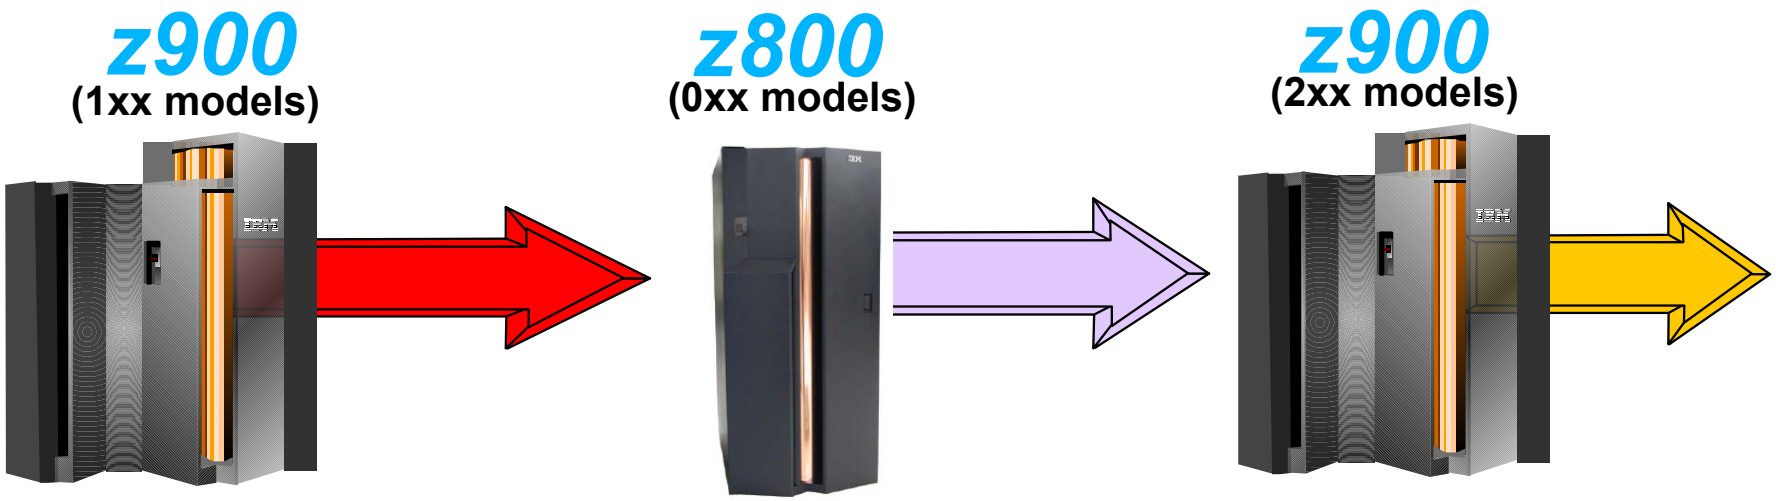

z900 M/T 2064
(October 2000)

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

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)

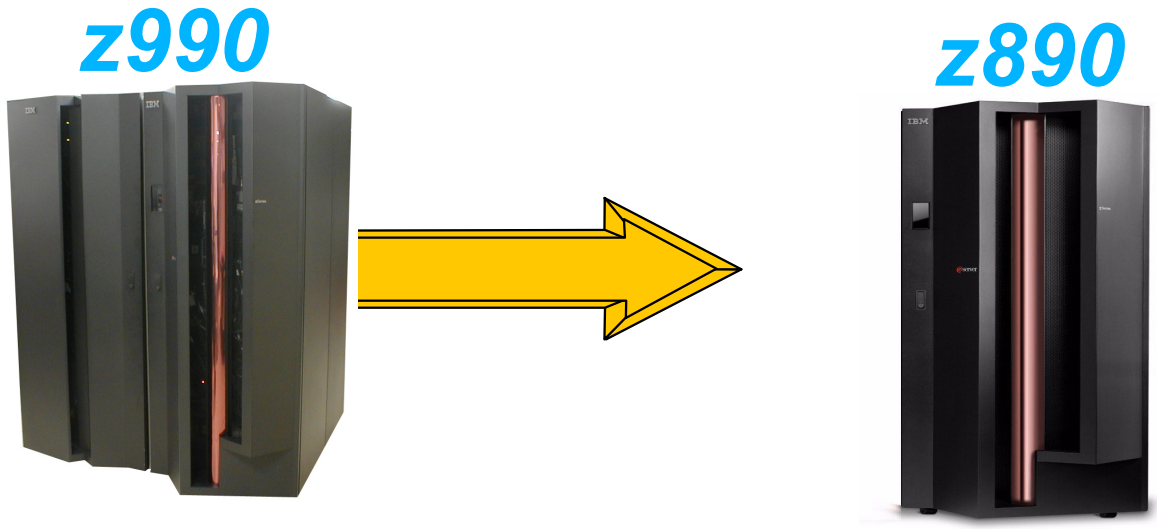

z990 M/T 2084
(May 2003)

- General Purpose hardware models
  - A08 (one book - up to 8 CPs)
  - B16 (two books - up to 16 CPs)
  - C24 (three books - up to 24 CPs)
  - D32 (four books - up to 32 CPs)

z890 M/T 2086
(April 2004)

- One General Purpose hardware model
  - A04 (one book - up to 4 CPs)

  Many capacity settings

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

IBM System z9 Hardware Technology

z9 Enterprise Class
(formally z9-109)

M/T 2094
(July 2005)

- S08 (one book - up to 8 CPs)
- S18 (two books - up to 18 CPs)
- S28 (three book - up to 28 CPs)
- S38 (four books - up to 38 CPs)
- S54 (four books - up to 54 CPs)

The IBM System z9 109 (2005) introduced a new family of servers

The IBM System z9 family
May also be referred to as
➡ System z9

z9 Business Class

M/T 2096
(April 2006)

- R07 (one book - up to 3 CPs)
- S07 (one books - up to 4 CPs)

- General Purpose hardware models

IBM z900 (z800) Introduced z/Architecture

IBM eservers 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 Up-gradable from G5/G6

zSeries 900/800 Family of Servers

z900

109
108
107
106
105
104
103
102
101

100

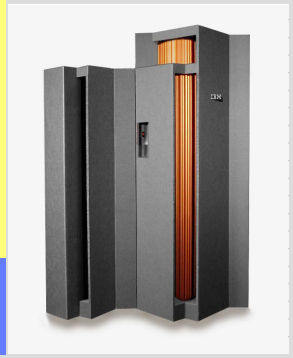

- CMOS 8S with Copper interconnect
- MCUs (Modular Cooling Unit)
- 12 PUs
- Up to nine CPs
- 5 - 32 GB Memory
- 2 Memory cards
- 1.3 ns Cycle time
- CF = Model 100

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

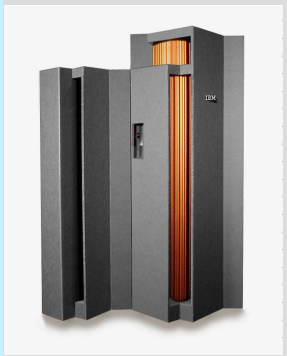

- CMOS 8S with Copper interconnect
- MCUs (Modular Cooling Unit)
- 20 PUs
- Up to 16 CPs
- 10 - 64 GB Memory
- Four Memory cards
- 1.3 ns Cycle time

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

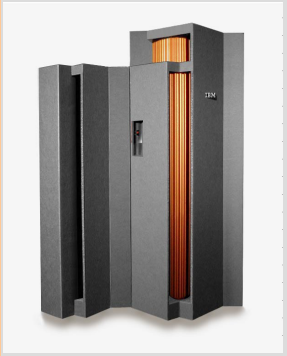

- CMOS 8SE with Copper interconnect
- MCUs (Modular Cooling Unit)
- 20 PUs
- Up to 16 CPs
- 10 - 64 GB Memory
- 4 Memory cards
- 1.09 ns Cycle time

z800

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

OCF
OLF

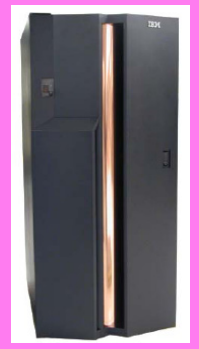

- CMOS 8S with Copper interconnect
- 5PUs
- Up to four CPs
- 8 - 32 GB Memory
- Four Banks of Memory chips
- 1.6 ns Cycle time
- CF = Model OCF
- LINUX model OLF

zSeries 900 Design

Modular Cooling Unit (MCU)

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

N +1 design

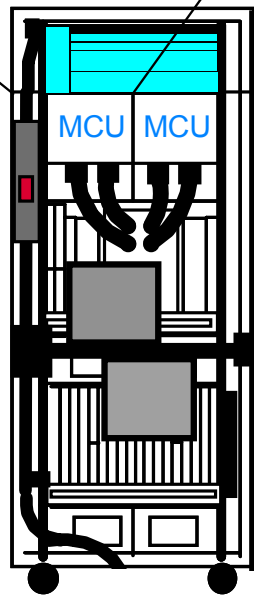

CEC Cage

Two S/Es

I/O Cage

z900 A Frame

May have an additional Z and/or B frame

CEC Cage - Front View

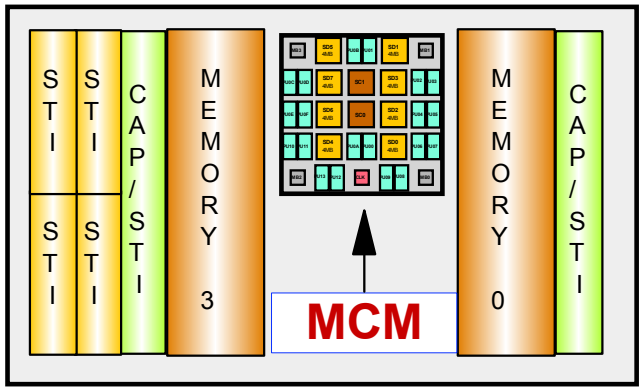

zSeries 900 CPC MCM will contain
- 12 PUs
- 20 PUs

depending on z900 model

CEC Cage - Rear View

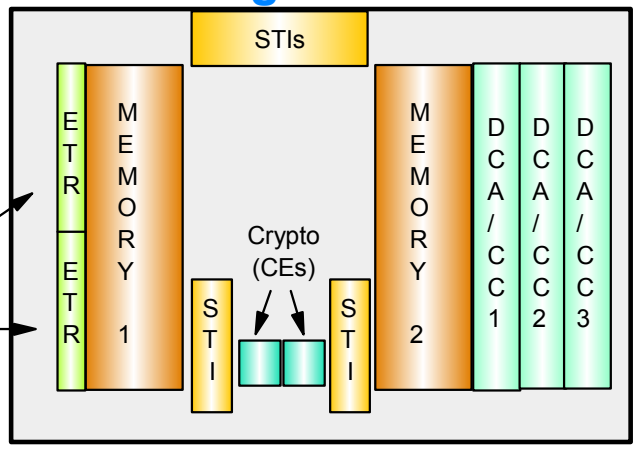

Two ETR cards N+1

The z800 design is similar

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

Recent IBM Mainframe Servers

The IBM System z9 and zSeries 990/890
Designed for on demand business computing

The on demand operating Environment

- **Integrated** - **Open** - **Virtualized** - **Autonomic**

-

IBM zSeries family

2003

zSeries 990
1 to 32 CPs

2004

zSeries 890
sub-uni to 4 CP

IBM System z9 family

2005

z9 Enterprise Class
1 to 54 CPs
Sub-uni capability
– (1 to 8 CPs)

2006

z9 Business Class
sub-uni to 4 CP
0 to 7 specialty engines

System z9 and zSeries 990 / 890 Terminology

- Book
  - A book contains an MCM (processors), memory and STI connections for the MBAs
  - System z9 EC and z990s can have multiple books
  - System z9 BC and z890s have one book
- System z Application Assist Processor (zAAP)
- System z9 Integrated Information Processor (zIIP)
- I/O Subsystem
  - All System z servers have one I/O Subsystem which utilizes one IOCDS and a single HSA
- Logical Channel Subsystem (LCSS)
  - System z9 and zSeries 990/890 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 server (CHPID is unique to a LCSS)

Book Design Concept

Book contents - MCM, Memory cards and MBA card

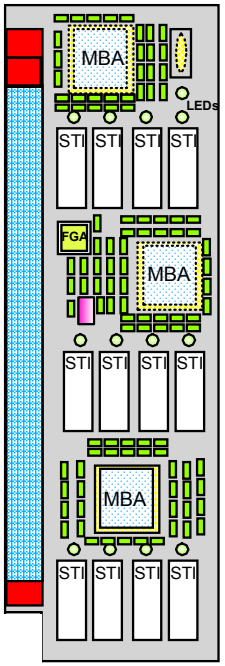

MBA card
- z990 3 MBA chips
- z890 2 MBA chips
- 4 STIs per MBA

z990 Book Example

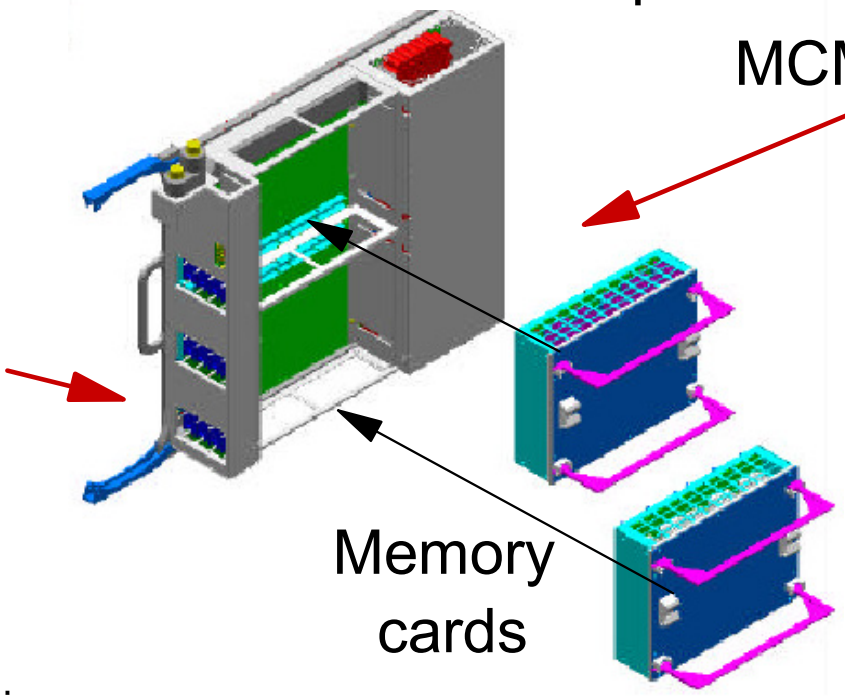

MCM

Memory cards

z990 - Two memory cards
z890 - One memory card

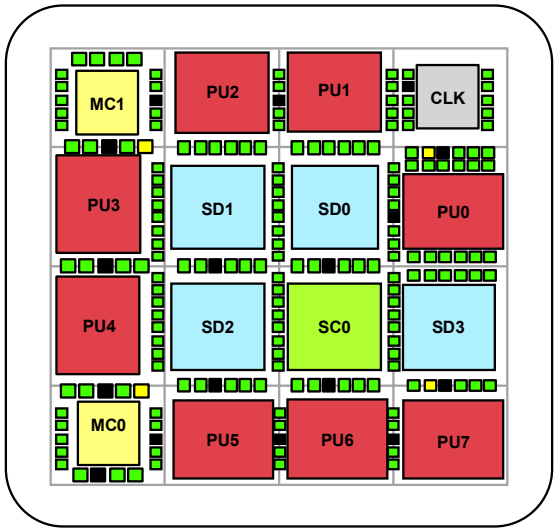

MCM PU usage
- z990- 8 characterizable PUs
- z990- 2 SAPs, 2 Spares
- z890- 4 characterizable PUs
- z890- 1 SAP

System z9 book similar, contains multiple memory and MBA cards, up to 16 STI connections, 8, 10 or 14 characterizable PUs per book depending on server model. Two spares PU available per z9 EC server model

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)

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.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 eserver 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.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 functionality
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 Crypto Accelerators (PCICA)
- PCIX Crypto Coprocessors (PCIXCC)
- PCI Crypto Express2

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

- Up to 48 (z890 40) OSA-E ports (OSA Express2)
- Up to 120/240 (z890 40/80) FICON Express and Express2 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

IBM System z9 Enterprise Class models

- z9 model S08
  - One book with 12 PUs, maximum of 8 PUs can be assigned as CPs
    - A book contains memory and STI connections
- z9 model S18
  - Two books with 24 PUs, maximum of 18 PUs can be assigned as CPs
- z9 model S28
  - Three books with 36 PUs, maximum of 28 PUs can be assigned as CPs
- z9 model S38
  - Four books with 48 PUs, maximum of 38 PUs can be assigned as CPs
- z9 model S54
  - Four books with 64 PUs, maximum of 54 PUs can be assigned as CPs

z9 EC

M/T 2094

- CMOS 10K with copper interconnect
- Up to twelve or sixteen Processor Units (PUs) per MCM.
- 16 - 512 GB Memory
- .6 ns cycle time
- Modular Refrigeration Units (MRU)

Note - system model number does not reflect the number of ordered CPs

System z9 EC Full Capacity Models and MSU

Each z9 EC model has an additional software capacity number association. The capacity number can be used for licensing and MSU purposes. z9 EC full capacity *MSUs range from 81 (701) to 2409 (754)

z9 S08		z9 S18		z9 S28		z9 S38		z9 S54	
capacity number	CPs	capacity number	CPs	capacity number	CPs	capacity number	CPs	capacity number	CPs
701	1	709	9	719	19	729	29	739	39
702	2	710	10	720	20	730	30	740	40
:	:	:	:	:	:	:	:	:	:
:	:	:	:	:	:	:	:	:	:
707	7	717	17	727	27	737	37	753	53
708	8	718	18	728	28	738	38	754	54

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

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

System z9 EC Sub Capacity Models and MSUs

Z9 EC Sub Capacity Settings and *MSU ratings

The capacity number can be used for licensing and MSU purposes.

	Model capacity identifier	MSUs	Model capacity identifier	MSUs	Model capacity identifier	MSUs
1-WAY	401	28	501	53	601	66
2-WAY	402	54	502	104	602	127
3-WAY	403	78	503	152	603	184
4-WAY	404	102	504	197	604	240
5-WAY	405	124	505	240	605	292
6-WAY	406	144	506	279	606	339
7-WAY	407	164	507	317	607	385
8-WAY	408	182	508	352	608	428

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

Relative Performance of z9 EC Models

Relative performance scale of z9 EC Servers

Number of z9109 CPs	Base	Ratio
1 CP	z990-301	1.35
8 CPs	z990-308	1.34
18 CPs	z990-318	1.34
28 CPs	z990-328	1.32
38 CPs	z990-332	1.51
54 CPs	z990-332	1.95

Note: For MSU values, refer to: www.ibm.com/servers/eserver/zseries/library/swpriceinfo/
For ITRs refer to: www.ibm.com/servers/eserver/zseries/lspr/zSerieszOS.html

* CI = Capacity Indicator and refers to number of installed CPs. Reported by STSI instruction. Model 700 does not have any CPs.

IBM System z9 Business Class Models

z9 Business Class

- z9 model R07
  - One book with 8 PUs, one standard SAP
  - Up to 7 PUs can be characterized
    - 1 to 3 as CPs
    - Up to 6 as specialty processors
  - Up to 15 LPs and 240 channels

M/T 2096

- z9 model S07
  - One book with 8 PUs, one standard SAP
  - Up to 7 PUs can be characterized as:
    - 0 to 4 as CPs
    - Up to 7 as specialty processors
  - Up to 30 LPs and 420 channels

- CMOS 10K
- Up to eight Processor Units (PUs) per MCM.
- 8 - 64 GB Memory
- .7 ns cycle time
- Air cooled

Models R07 and S07 have many different capacity settings

Note - A book contains memory and STI connections

IBM System z9 BC Capacity Settings

A z9 BC model has many capacity settings available.
Capacity settings can be used for licensing and MSU purposes.

z9 R07 model

Capacity Indicator	1 CP MSU	2 CP MSU	3 CP MSU
A	4	7	10
B	5	10	15
C	6	12	18
D	8	16	23
E	10	19	N/A
= = =	= = =	= = =	= = =
J	24	N/A	N/A

z9 S07 model

Capacity Indicator	1 CP MSU	2 CP MSU	3 CP MSU	4 CP MSU
K	N/A	N/A	N/A	30
L	N/A	N/A	28	36
= = =	= = =	= = =	= = =	= = =
W	47	92	134	174
X	53	103	150	195
Y	59	115	166	216
Z	67	130	189	246

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

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

IBM System z9 New Functions and Features

System z9 is the next step in the evolution of the mainframe family

IBM System z9 offers zSeries features and functions introduced by the zSeries 990 and much more

- New faster Uni Processors
- Up to 60 LPARs
- CBU for IFL, ICF, zAAP and zIIP
- Spare Engines
- Enhanced CPACF with AES, PRNG and SHA-256
- Configurable Crypto Express2
- Separate PU pool management
- Redundant I/O interconnect
- Hot Pluggable /maintainable MBA/STI fanout cards
- Enhanced Driver maintenance
- Enhanced Book availability
- MIDAW facility
- FCP N-port virtualization
- OSA-Express2 1000BASE-T
- OSA-Express2 (OSA for NCP)
- Dynamic oscillator switchover
- 54 additional hardware instructions

z9 EC Five H/W models
- 1 to 54 assignable PUs

z9 BC Two H/W models
- 1 to 7 assignable PUs

Up to 512 GB Processor Memory

- Up to 4 Multiple Logical Channel Subsystems with Multiple Subchannel Sets (MSS)
  - 63.75K Subchannels for Set-0
  - 64 K subchannels for Set-1
  - more than doubles the amount of subchannels previously available
- 256 Channels per LCSS
- Up to 1024 Channels
- Up to 1024 ESCON ports
- Up to 48 OSA-E ports
- Up to 336 FICON Express2 ports
- Up to 16 HiperSockets
  - with IPV6 support
- Up to 16 2.7GB/sec STIs per book

Frame and Cage Design Overview

System z9 EC and zSeries 990
- 1 to 3 I/O cages

System z9 BC and zSeries 890
- 1 I/O cage

Each I/O cage up to 20 FICON adapters

MRU components

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

Hybrid cooling design

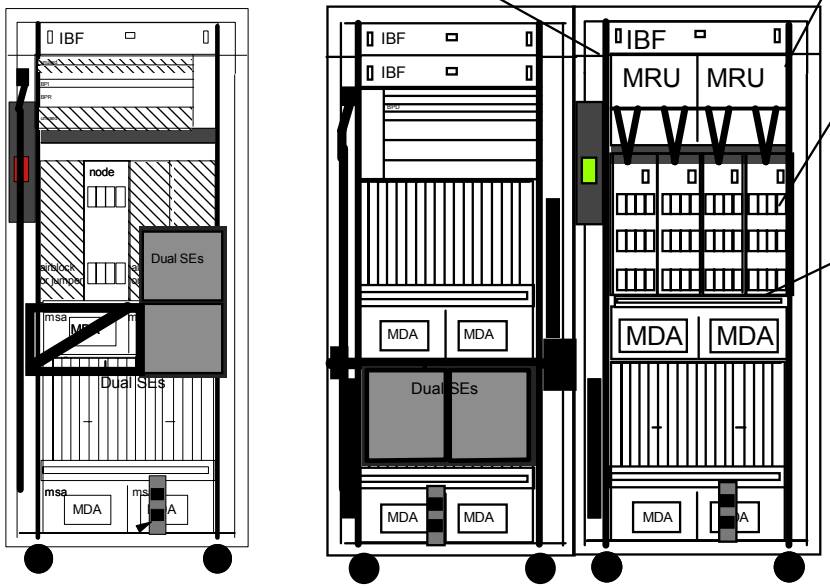

A frame | Z frame | A frame

z9 BC, z890 | z9 EC, z990

CEC Cage - Front View

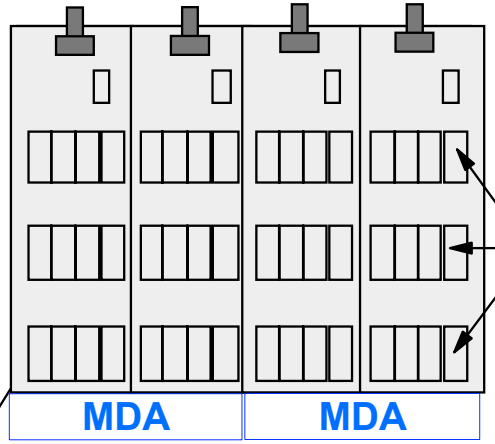

System z9 EC and z990
One to four books

System z9 - Up to 16 STIs per book
z990 - 12 STIs per book

System z9 BC and z890
One book
System z9- Up to 16 STIs
z890- 8 STIs

CEC Cage - Rear View

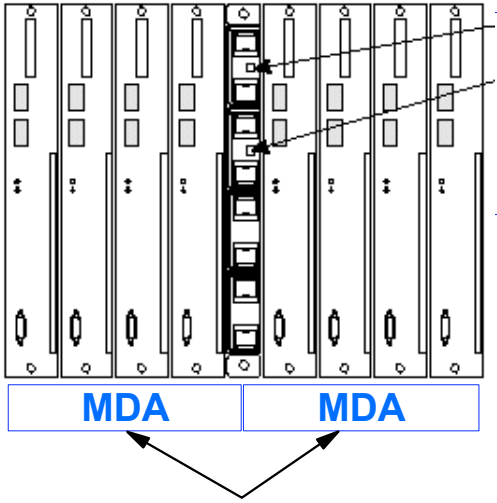

Two ETR and multiple DCA cards located in rear of CEC cage (N+1)

These MDAs are system activated in case of MRU failure

System z9 and z890 - A frame only, one book, MDA used for cooling (N+ 1)

Large Systems Storage Hierarchy

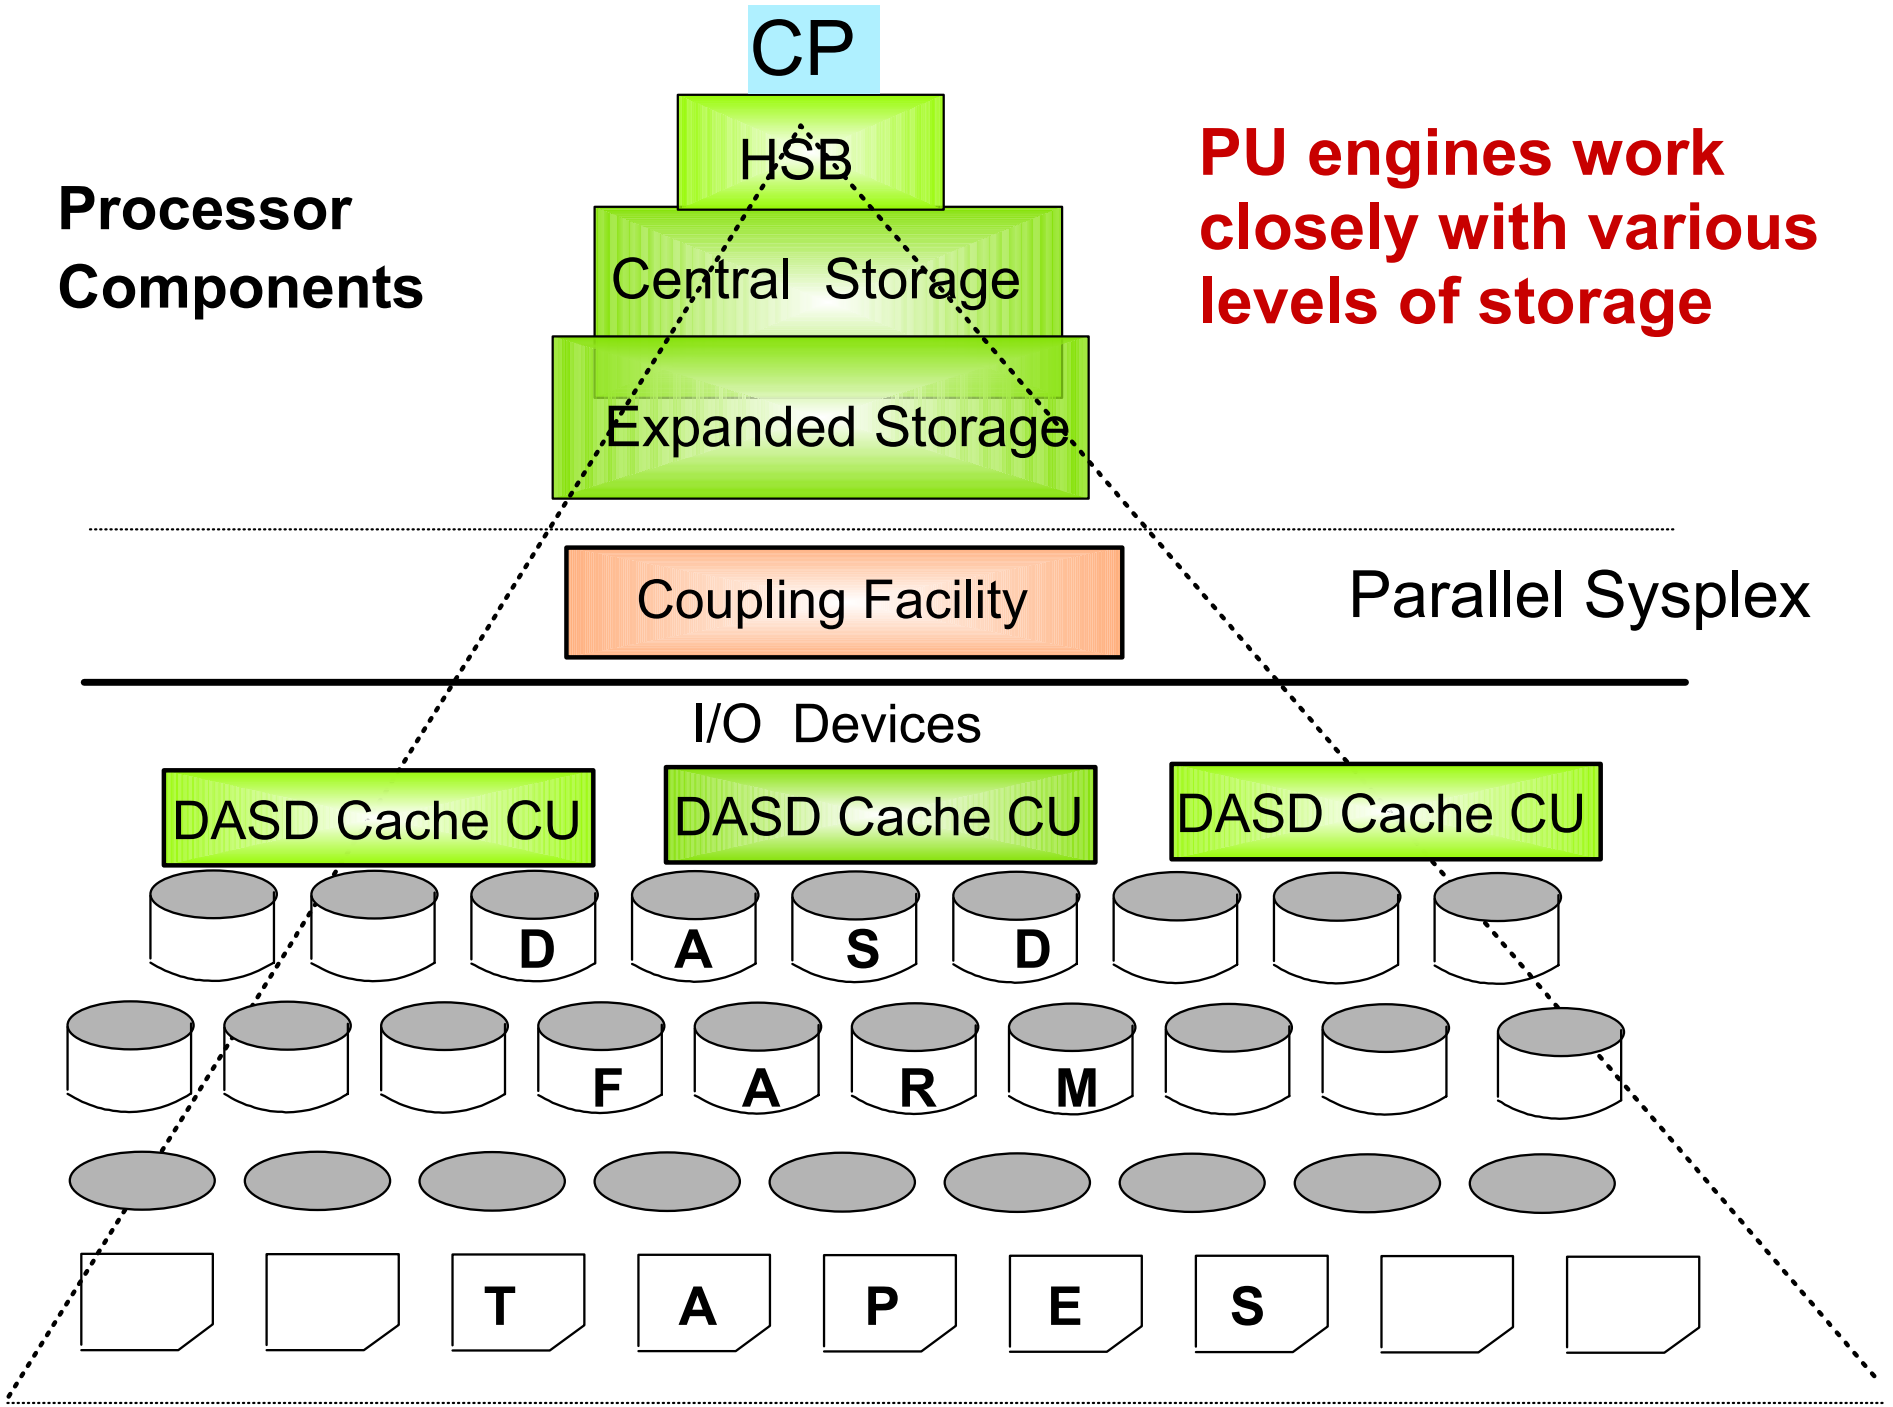

z990/z890 Book and MCM System Structure

z990 Book configuration

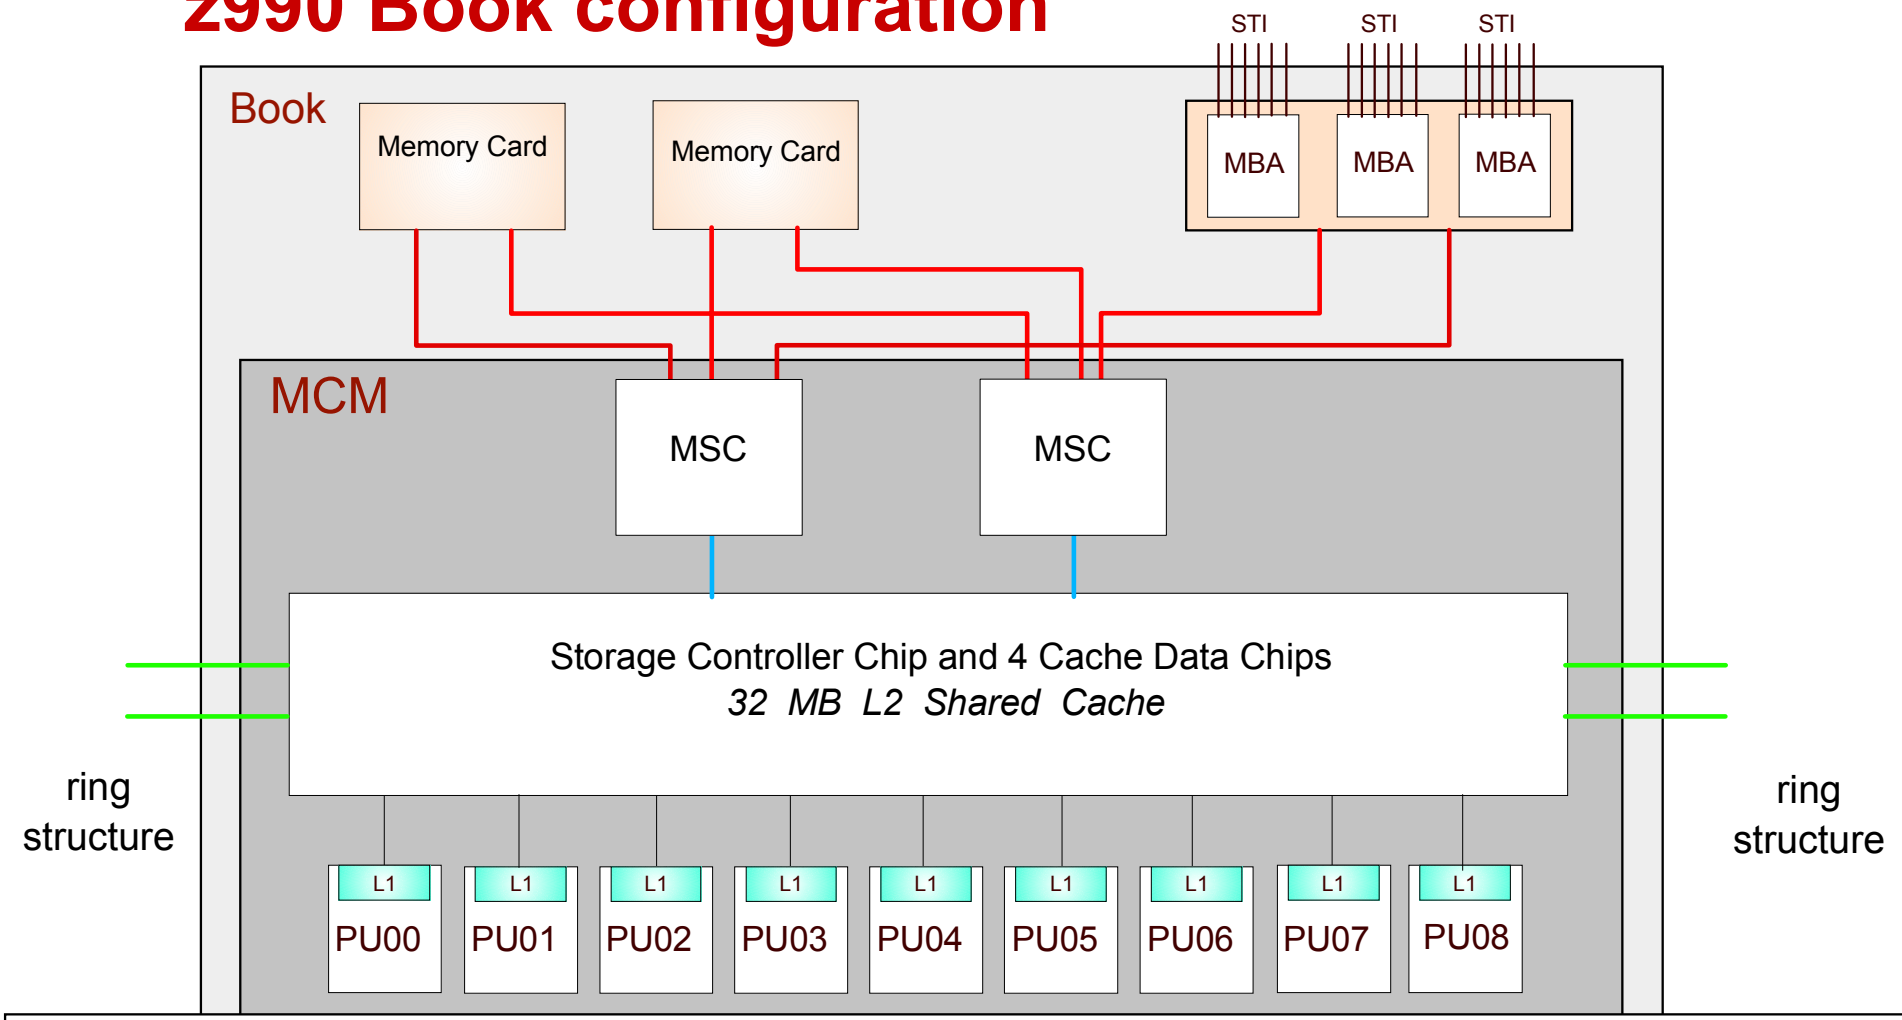

- **z890 Book configuration similar, with the following exceptions One Memory card, Two MBAs, 5 PUs**
- **System z9 book concept is similar, however number of PUs, memory, MBA and STIs configurations are slightly different**

System z9 Book and Memory Ring Structure

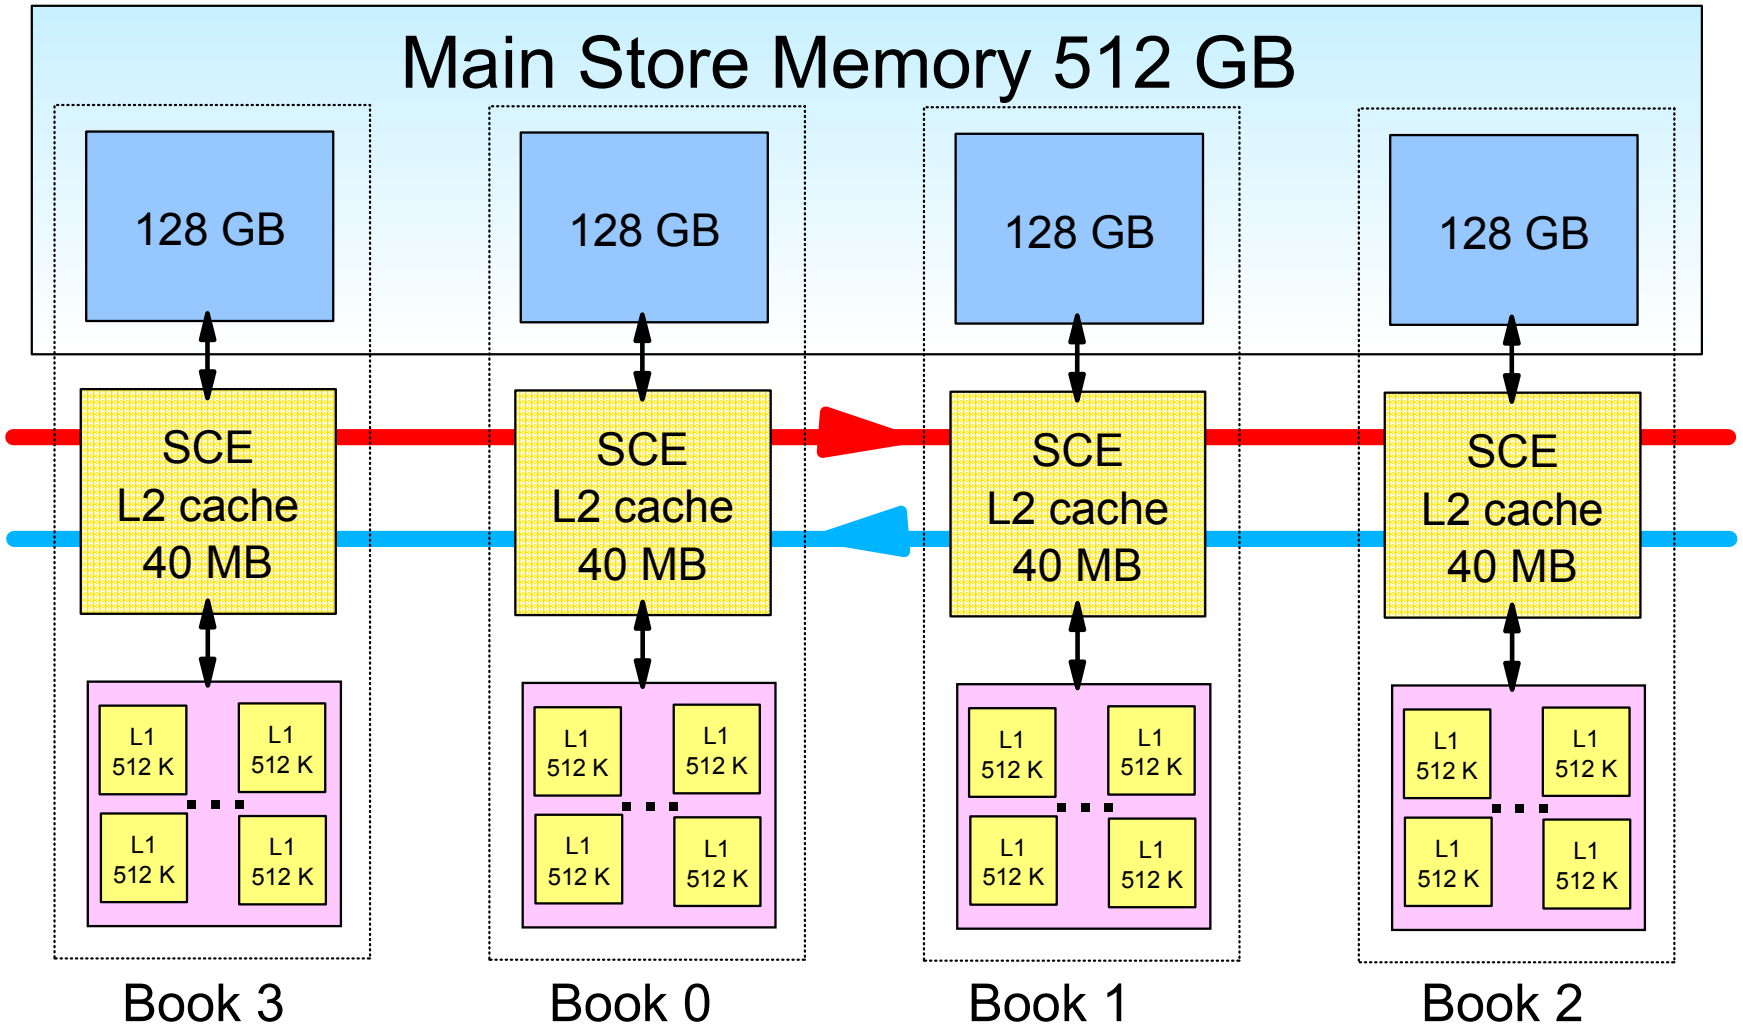

z9 EC - 1 to 4 books, each book has 12 or 16 PUs, memory cards and MBA fan out cards with up to 16 STIs
z9 BC - 1 book, 8 PUs, memory cards and MBA fan out cards with up to 16 STIs

CEC to I/O Cage Physical Connectivity

STI cables provide the physical connectivity between the CEC and I/O cages - All external I/O operations use STIs

CEC cage

I/O cage

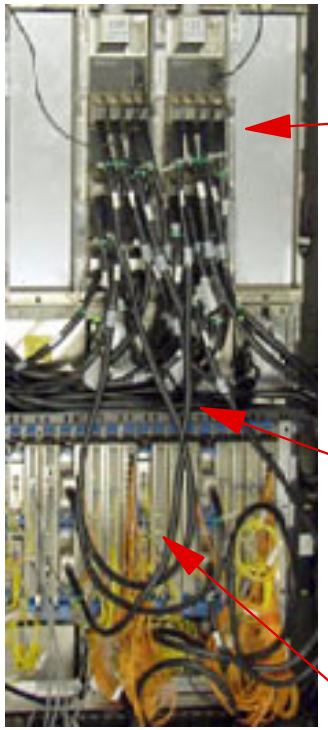

- Book located in the CEC cage
  - Contains Memory Bus Adapters (MBA) that provide a data path into memory and the MCM for all external I/O operations
  - Self Timed Interface (STI) cables provide the connectivity from the CEC cage (MBA) to the I/O cage
  - One STI cable provides bandwidth for multiple channel adapters in one I/O domain

zSeries Channel Connections - STIs

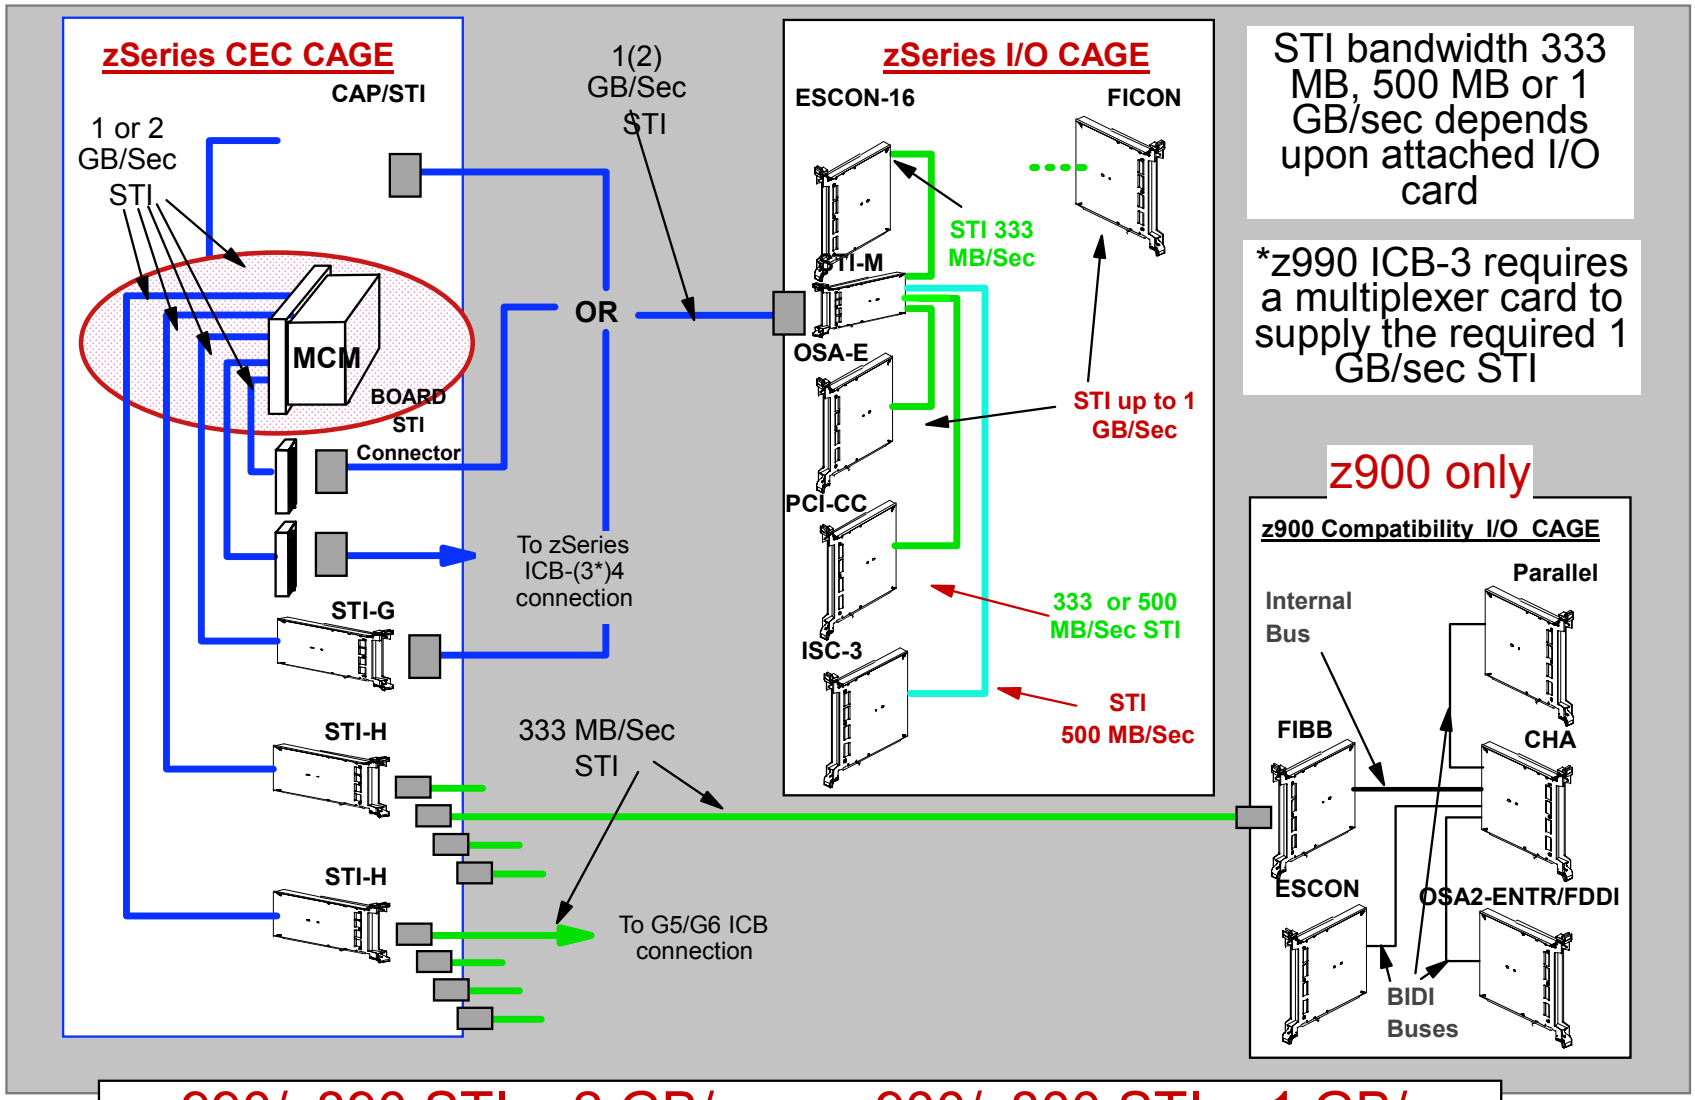

- z990/z890 STI = 2 GB/sec, z900/z800 STI = 1 GB/sec
- Only z900 supports I/O capability cage

z9 Channel Connections - STIs to I/O Cage

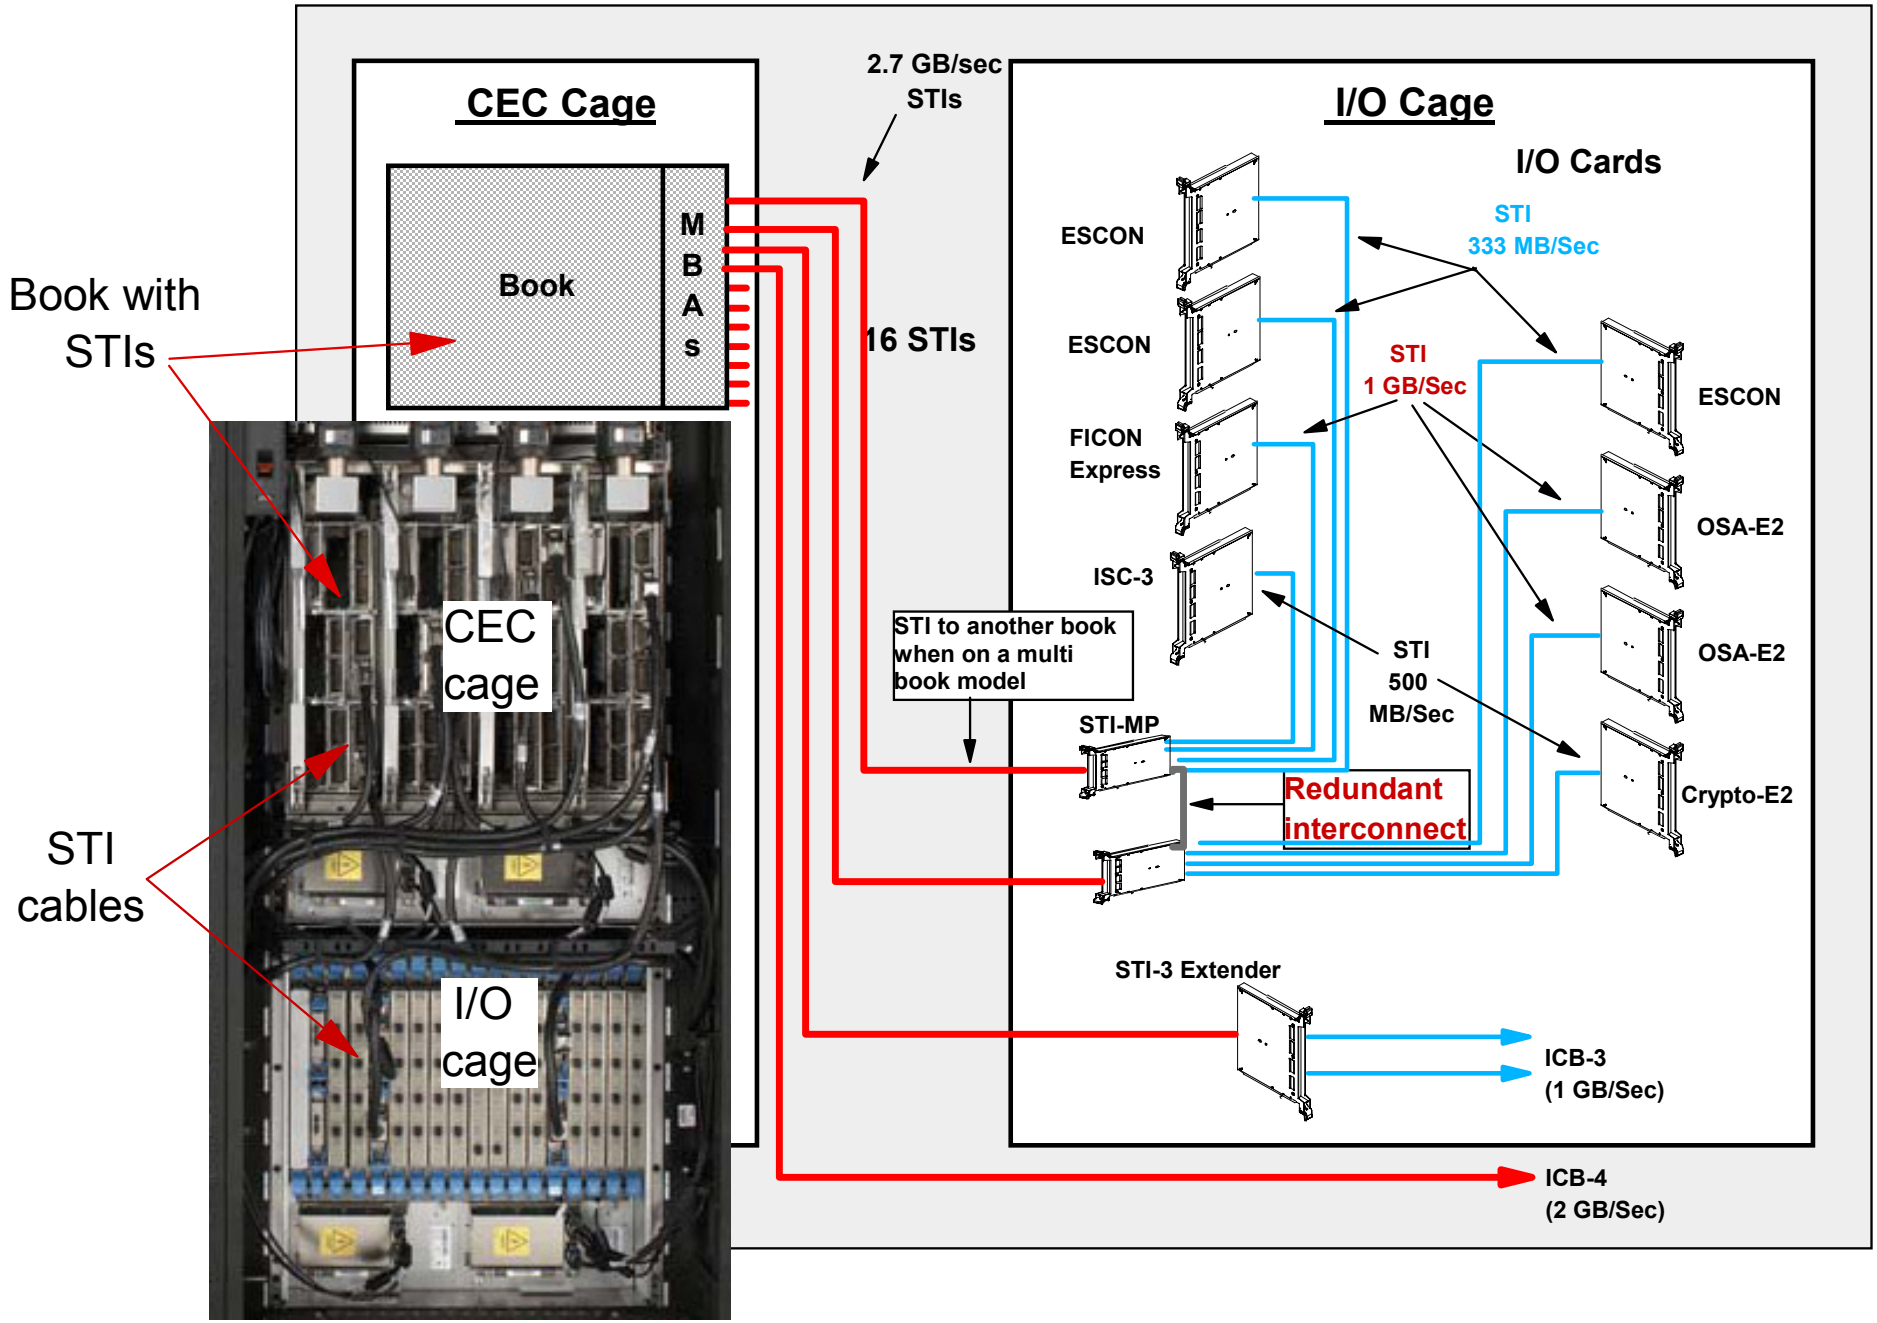

CPs, SAPs, STIs - Putting the Pieces Together

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

System z 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

z/OS (OS) I/O Operation Overview

A closer look at how the mainframe OS works with the CSS

z/OS operates in an LP
z/OS works with many
application programs

One CSS can work with many operating systems

Channel Subsystem

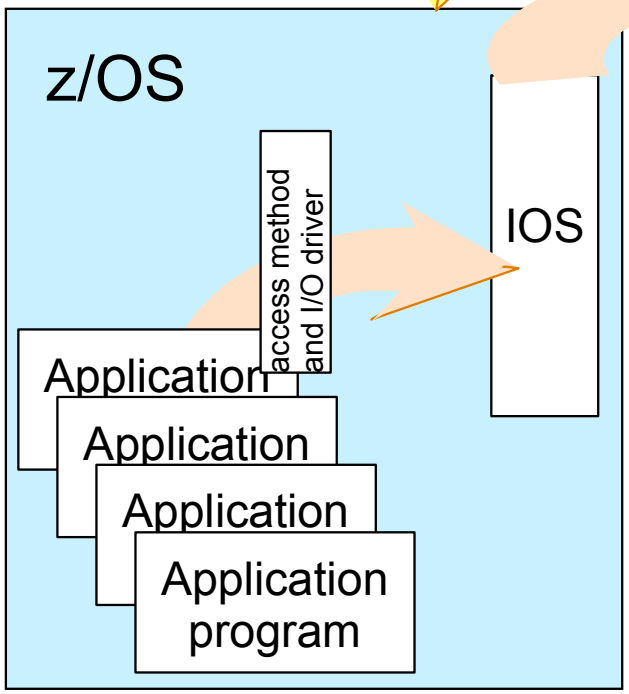

- Input Output Supervisor (IOS) is the z/OS component that interfaces between the OS and the CSS
- IOS issues the Start Subchannel (SSCH) command to the CSS
- SSCH passes the Operation Request Block (ORB) to the CSS
- The ORB contains the Address of the first CCW, other control information and main storage location to transfer data to/from for the I/O operation

Application programs(also called User programs or jobs) use resources like CPs, memory and UCBs

PSW Swaps and Interrupts (Concept)

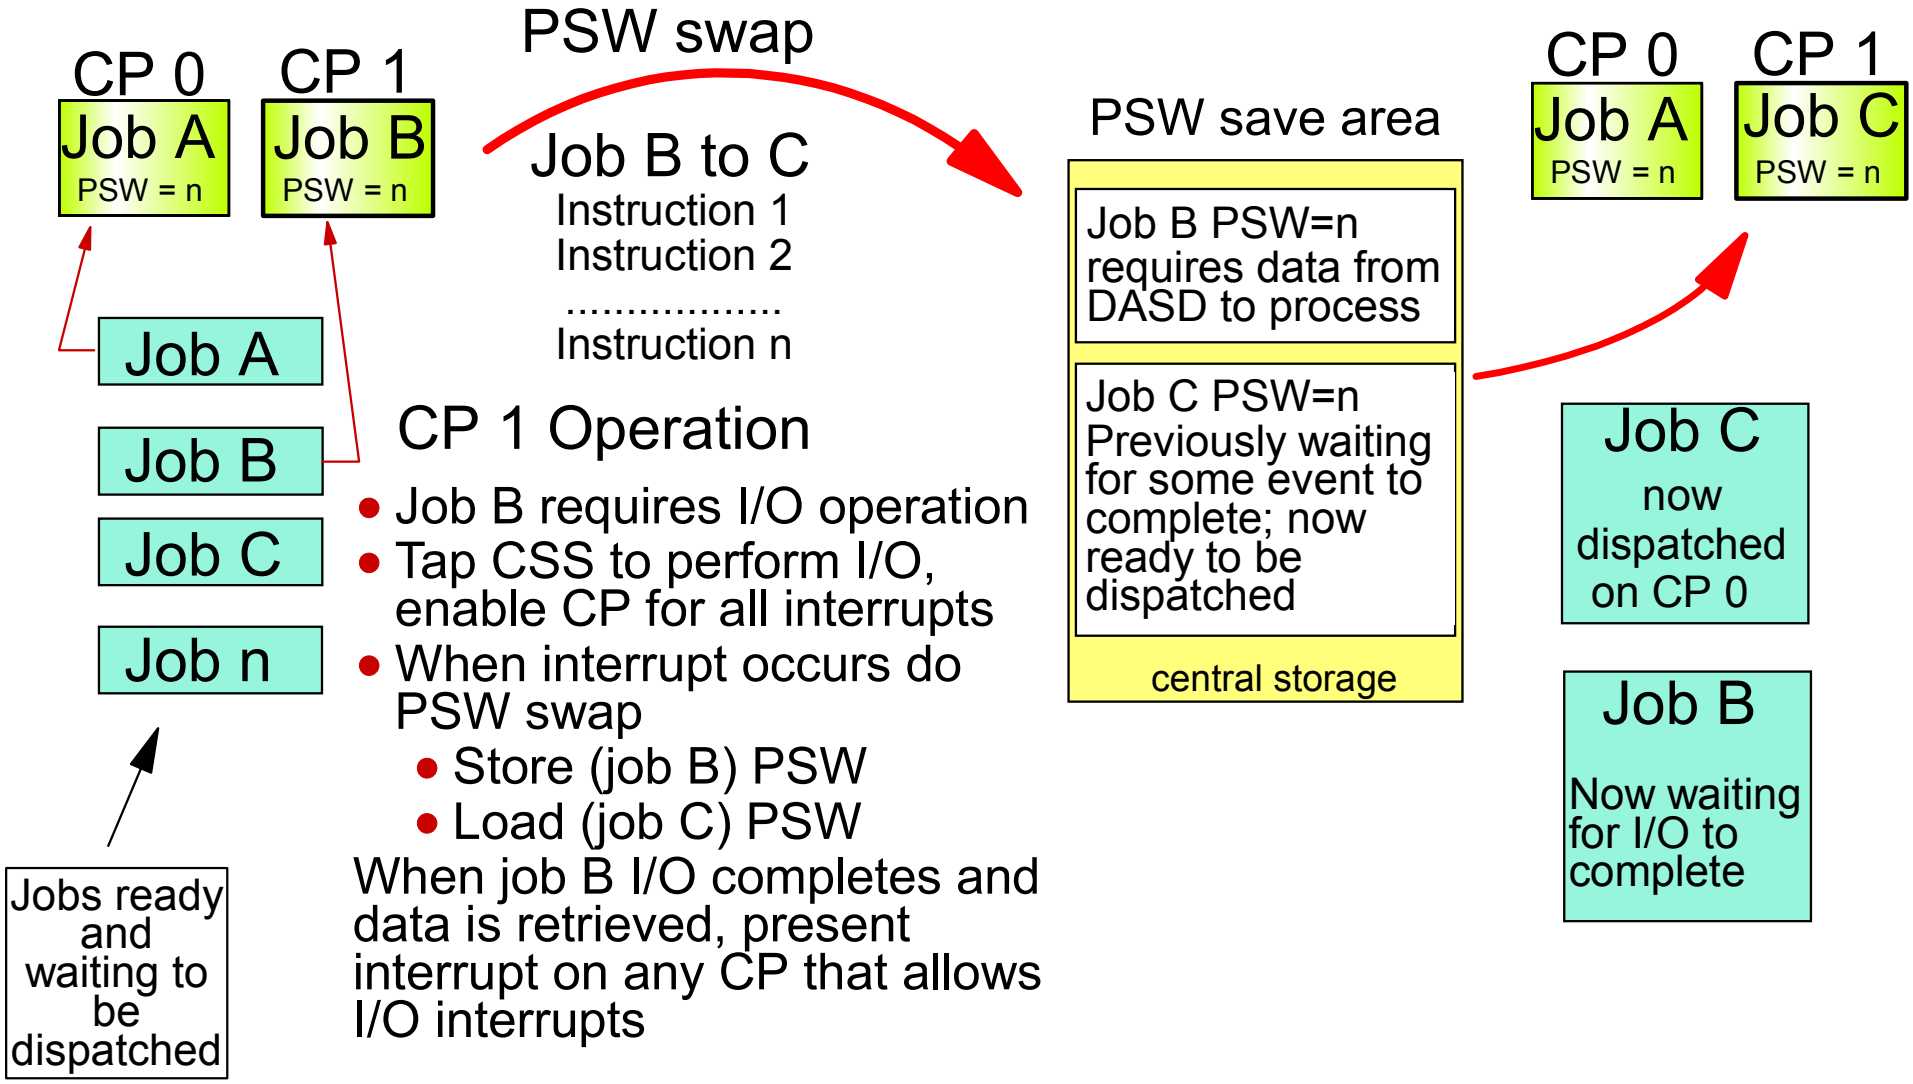

Program Status Word (PSW) - One for each CP, contains status information and next instruction address to be processed.

CSS I/O Operation Overview

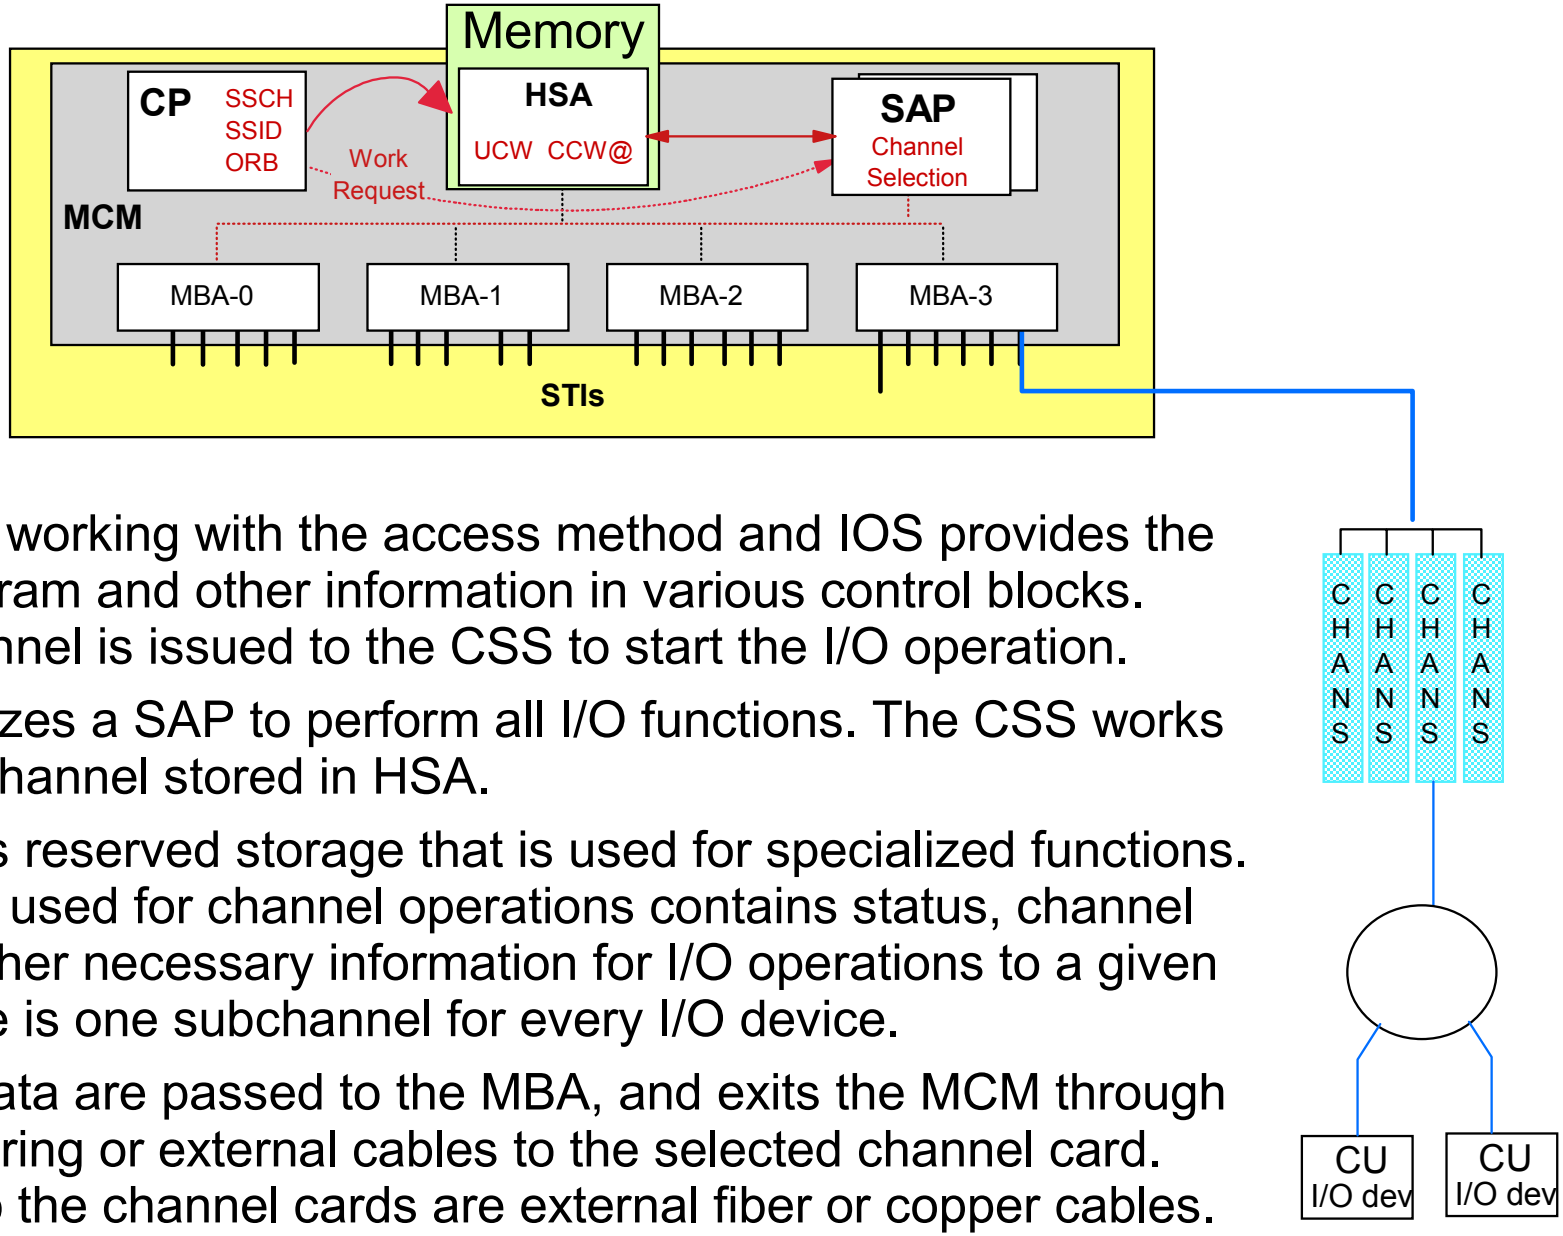

- The program working with the access method and IOS provides the channel program and other information in various control blocks. Start subchannel is issued to the CSS to start the I/O operation.
- The CSS utilizes a SAP to perform all I/O functions. The CSS works with the subchannel stored in HSA.
- 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 are 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.

That's it Folks!

Questions?

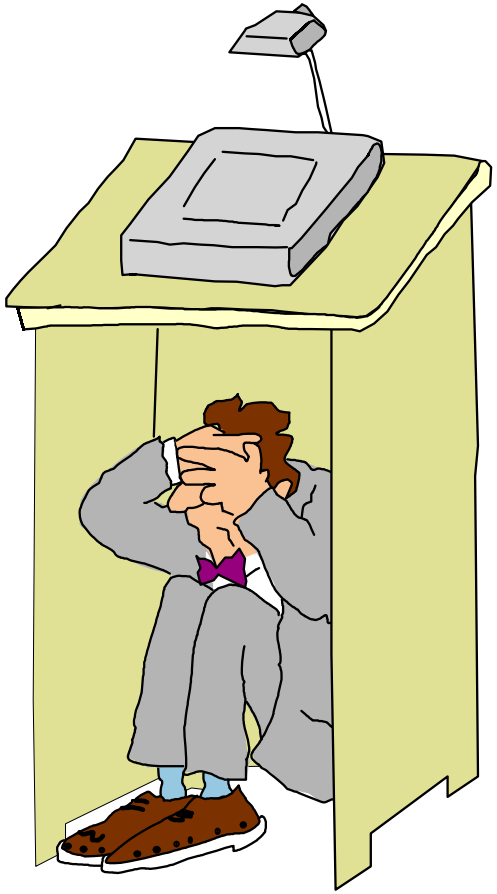

System z Educational Offerings

Sysplex / System z 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) Plex Implementation
- ES830 (5) CSAR (Complex Systems Availability & Recovery)
- ES820 (2) System z Mainframe Environment (A Technical Overview)
- OZ05 (2) System z9 / zSeries 990/890 Technical Update & Configuration
- OZ96 (5) zSeries Channel Architecture, ESCON/FICON Operation and PD
- ES326 (3) FICON(fc,fcv,fcp) Planning, Implementation, Operation and PD
- ES960 (4) HCD and Dynamic I/O
- ES270 (3) z/OS and System Operations

RETURN TO INDEX