



IBM @server zSeries 990

ITSO - z990 *Technical* Workshop

Introduction - Processor Complex

Technical Operations - Processor Complex

z990 GA1 06/2003, z990 GA2 10/2003

2084-z990

ITSO Poughkeepsie

Ken Trowell

Trademarks

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

CICS*	IBM TotalStorage*	Sysplex Timer*
DB2*	IMS	VM/ESA*
e-business logo*	Multiprise*	VSE/ESA
Enterprise Storage Server*	OS/2*	WebSphere*
ESCON*	OS/390*	z/Architecture
FICON	Parallel Sysplex*	z/OS
FICON Express	PR/SM	z/VM*
HiperSockets	Resource Link	zSeries*
IBM*	S/390*	
IBM @server	S/390 Parallel Enterprise Server	

* 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

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.

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.



- z990 - Processor Introduction subjects
 - Layouts
 - Processor
 - Memory
 - Channel Subsystem Operation
 - I/O support
 - Support
- z990 Processor Technical Operations
 - PU
 - Book to Book communication
 - Book to channel communication
 - PU to Memory Data access
 - Ring Structure
 - Memory addressing
 - I/O operation flow



- IBM eServer - zSeries model 990
 - IBM z990
 - z990
- IBM Machine Type
 - 2084
- 4 Machine type hardware models:
 - 2084 model A08 - 1 Book (GA1)
 - 2084 model B16 - 2 Books (GA1)
 - 2084 model C24 - 3 Books (GA2)
 - 2084 model D32 - 4 Books (GA2)





- Machine Type
 - 2084
- 4 Hardware Models:
 - 2084 model A08 - 1 Book (GA1)
 - 2084 model B16 - 2 Books (GA1)
 - 2084 model C24 - 3 Books (GA2)
 - 2084 model D32 - 4 Books (GA2)
- One or more Books make up the CEC (processor)
- Each Book has 12 PUs:
 - 8 PUs available per book for characterization as:
 - CPs, IFLs, ICFs or additional SAPs
 - 2 PUs standard as SAPs
 - 2 PUs standard as Spares
- Memory:
 - Up to 64 GB per book
 - Up to 128 GB with 2 Books, 192 GB with 3 books and 256 GB with 4 Books
 - 8 GB increments
 - System minimum of 16GB GA2 (GA1 8GB)
- I/O:
 - Each Book has 12 eSTIs at 2 GB/s
 - eSTIs support Book to Channel connectivity
 - Up to 512 channels per CPC are supported
 - Dependent on channel type mix





zSeries 990

Z- Frame

A- Frame



External View - Front

Z- Frame

A- Frame




Internal View - Front

The z990 always has two frames for the CPC




zSeries 990






zSeries 990





z990 System Overview



New function for z990

Updates for [GA2 on 10/31/03](#)

● Processor


- 4 flexible models (A08 - B16 - C24 - D32)
- Unified SMP structure of up to 4 books
 - Concurrent model (book) upgrade
- 64-bit z/Architecture™
- Up to 48 PUs (12 per book), up to 32 characterizable PUs
- CMOS9S-SOI Technology SuperScalar
- Crypto assist for DES and SHA in every PU
- Capacity Upgrade on Demand
- On/Off Capacity on Demand (12/15/03)
 - For CPs (9/15/03)
 - For IFLs (10/7/03)
- Capacity Backup (CBU)
- Customer Initiated Upgrade (CIU) (9/15/03)
- Hybrid Cooling (Air/Liquid)
- LPs, up to 30 defined and 15 active
- LPs, up to 30 defined and 30 active
- Optional ETR attachment

● Memory

- Maximum system memory 256 GB (D32)
 - Minimum system memory 16 GB (10/1/03)
- Card sizes 8, 16, 32 GB (2 cards per book)
- Bi-directional redundant ring memory interconnect between books



z990 System Overview



The z990 is always a two Frame System

New function for z990

Updates for [GA2 on 10/31/03](#)

● I/O (LCSSs and Channels)

- 64-bit Architecture (42/48-bit I/O addressing in hardware)
- From 12 and up to 48 x 2 GB eSelf-Timed Interconnects (eSTIs)
- I/O cage with enhanced power (1 to 3)
- Up to 2 Logical Channel SubSystems (LCSS)
 - Up to 256 channels per LCSS, 512 total
 - Internal channels spanned between LCSSes
 - Dynamic I/O support for LCSS 0
 - Dynamic I/O support for both LCSS 0 and 1
- Up to 48 OSA-Express network connectors
 - QDIO portname relief, VLAN (all zSeries)
- Up to 120 FICON™ Express Channels
- SCSI over Fibre Channel (FCP for Linux)
 - IPL from FCP disk for Linux and SA dump
 - SAN Management - CHBA API
- Up to 16 HiperSockets

● Crypto function


- CPACF in every PU
- PCICA
- PCIXCC (9/19/03)
- TKE 4.0 (9/30/03)
- No CHPID numbers required

● Parallel Sysplex®

- ICB-4 (2 GB/s) ICB-3, ICB-2, ISC-3, IC
- CF Duplexing
- 100 km distance between sites (RPQ)




CEC Cage




zSeries 990 Processor Books


Book 3




Book 0



Book 1




Book 2



- 1, 2, 3 or 4 Books can be installed
 - Book 0 is first - 2084 hardware Model A08
 - Book 0 and 1 - 2084 hardware Model B16
 - Book 0, 1 and 2 - 2084 hardware Model C24
 - Book 0, 1, 2 and 3 - 2084 hardware Model D32





zSeries 990 Processor Book




- Each Book contains
 - 12 Processor Units (PUs), PUs reside on the microprocessor chips located on the Multi-Chip-Module (MCM)
 - 16GB to 64 GB of physical memory, always two memory cards per book, each (in pairs) containing 8, 16, or 32 GB
 - Three Memory Bus Adapters (MBAs), supporting 12 eSelf Timed Interfaces (eSTIs), eSTIs provide the connection to the I/O cages and/or ICB-4 channels (ICB-4 is a z990 to z990 Coupling Facility channel)



z990 Processor Book Riser Card Memory Bus Adapters (MBAs) and STI Ports Z990 Memory, and z990 MCM

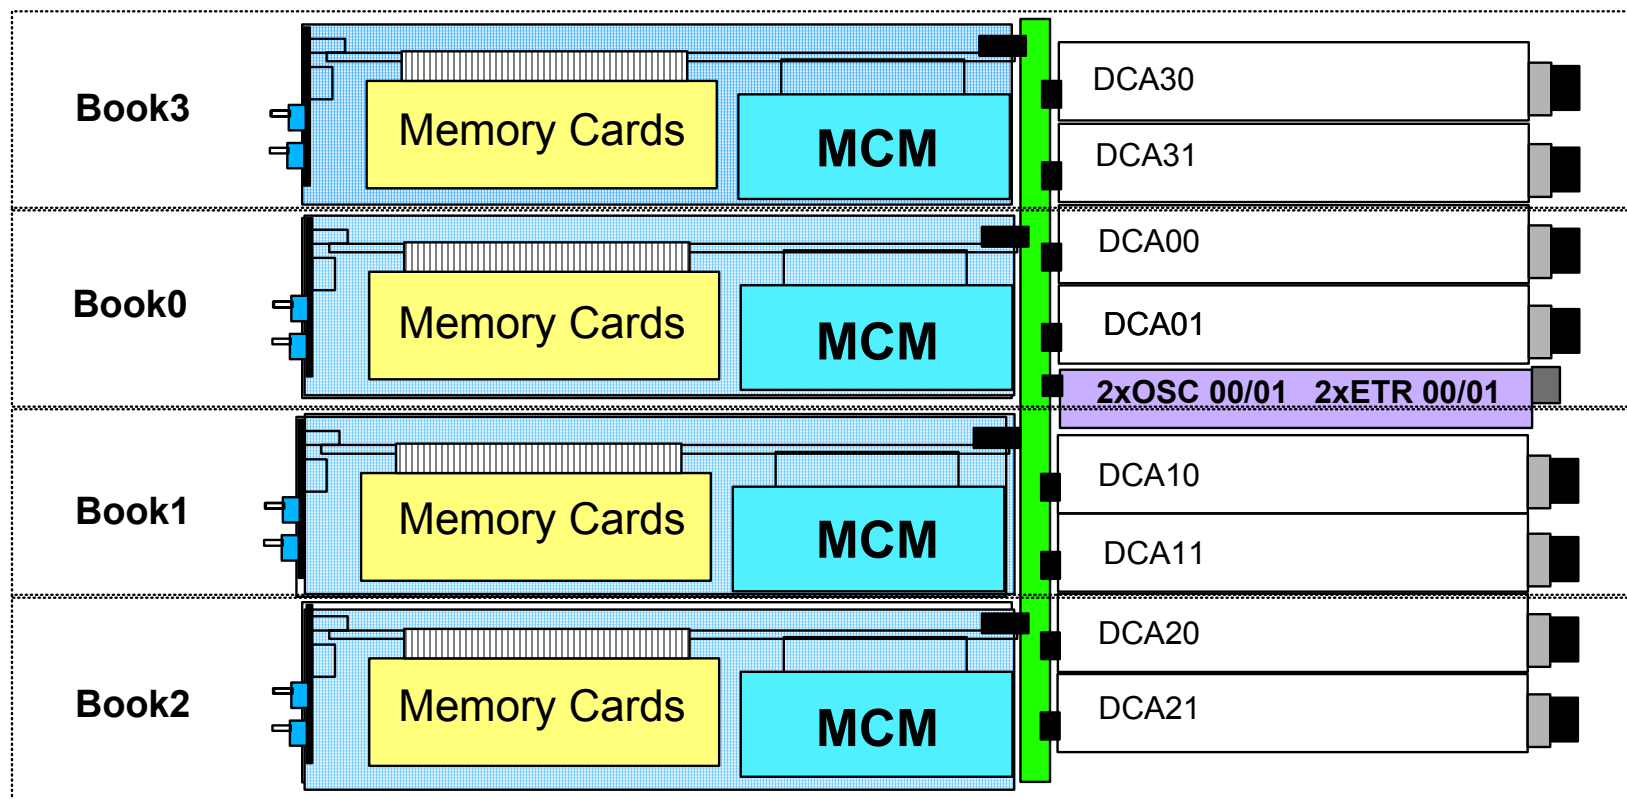




z990 'Book' board and MCM site



z990 CEC cage - Top down view



- Each Book is powered from two Distributed Converter Assemblies (DCA) that reside on the opposite side of the CEC board
- The DCAs can be concurrently maintained, which means that replacement of one DCA can be performed without taking the Book down (the Book is still operational while the DCA is being replaced)
- The ETR connections are via the ETR ports 00 and 01 located at the rear of the CEC cage
 - ETR fiber cable connections are MTR-J

MRU to Book Coolant circuit





- The hardware model number Nxx indicates:
 - N - the number of Books (A, B, C, or D denotes 1, 2, 3, or 4 books respectively)
 - xx - the maximum of PUs available for characterization (8, 16, 24, or 32 PUs)
 - PU characterization is identified by the type of features ordered (CP, IFL, ICF, SAPs)
 - 2 SAPs are standard on each MCM
 - 2 Spares are standard on each MCM
- Two Memory cards are standard with each book
 - Minimum memory at the system level (CEC) is
 - 8 GB - at GA 1 ordering
 - 16GB - at GA 2 ordering

Models	# of MCMs	Max PUs	Maximum Characterizable CP/IFL/ICF*	Standard SAPs	Standard Spares	Memory Cards	Max Memory	Max** Channels
A08	1	12	8	2	2	2	64 GB	512
B16	2	24	16	4	4	4	128 GB	512
C24	3	36	24	6	6	6	192 GB	512
D32	4	48	32	8	8	8	256 GB	512

* Must have a minimum of one CP or IFL or ICF ordered/installed

** Quantity of supported channels vary depending on channel types

ESCON (512), FICON (120), OSA-Express (48), ISC, ICB, IC, IQD

- The z990 has in addition to hardware model numbers, also software model numbers
 - The software numbers are 300, and 301 to 332
 - 3xx, where xx = the number of PU *characterized as CPs* in the CEC/CPC



# of PU Characterized as CPs	2084-A08	2084-B16	2084-C24	2084-D32
0	300	300	300	300
1	301	301	301	301
2	302	302	302	302
3	303	303	303	303
4	304	304	304	304
5	305	305	305	305
6	306	306	306	306
7	307	307	307	307
8	308	308	308	308
9		309	309	309
10		310	310	310
11		311	311	311
12		312	312	312
13		313	313	313
14		314	314	314
15		315	315	315
16		316	316	316
17			317	317
18			318	318
19			319	319
20			320	320
21			321	321
22			322	322
23			323	323
24			324	324
25				325
26				326
27				327
28				328
29				329
30				330
31				331
32				332

2084 xxx software model 300 has no characterized CPs, but it must have at least one or more ICFs and/or IFLs



z/OS Command
'D M=CPU' displays
software model numbers

```

D M=CPU
IEE174I 20.04.48 DISPLAY M 575
PROCESSOR STATUS
ID  CPU                      SERIAL
0    +                      116A3A2084
1    +                      116A3A2084
2    -

CPC ND = 002084.305.IBM.02.0000000026A3A
CPC SI = 2084.305.IBM.02.0000000000026A3A
CPC ID = 00
CPC NAME = A11
LP NAME = A11      LP ID = 11
CSS ID  = 1
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE
CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
CPC SI  SYSTEM INFORMATION FROM STSI INSTRUCTION
CPC ID  CENTRAL PROCESSING COMPLEX IDENTIFIER
CPC NAME CENTRAL PROCESSING COMPLEX NAME
LP NAME  LOGICAL PARTITION NAME
LP ID   LOGICAL PARTITION IDENTIFIER
CSS ID  CHANNEL SUBSYSTEM IDENTIFIER
MIF ID  MULTIPLE IMAGE FACILITY IMAGE IDENTIFIER

```

Hardware model number
incorrectly shows the software
model number. IBM will change
to correctly show the hardware
model number

Software bmodel number

CPC Name the LP name is
displayed in error
IBM will change this show
the correct CPC name



SCZP901: Primary Support Element Workplace (Version 1.8.0)

Views Daily

Groups Exceptions

SCZP901

SCZP901 Details

Instance information

CP Status:	Operating	Activation profile:	DEFAULT
CHPID Status:	Exceptions	Last used profile:	not set via Activate
Group:	CPC	Service state:	Disabled
IOCDS identifier:	A08	Maximum CPs:	5
IOCDS name:	IODF22	Maximum ICFPs:	3

Lockout disruptive tasks: ☐ Yes ☒ No

System mode: Logically partitioned Dual AC power maintenance: Fully Redundant

Alternate SE Status: Operating CP Assist for Cryptographic Functions: Installed

Acceptable CP/CHPID status

<input checked="" type="checkbox"/> Operating -	<input type="checkbox"/> Power save -	<input type="checkbox"/> No power -
<input type="checkbox"/> Not Operating -	<input type="checkbox"/> Exceptions -	<input type="checkbox"/> Status check -
<input checked="" type="checkbox"/> Acceptable -	<input type="checkbox"/> Service Required -	<input type="checkbox"/> Degraded -

Product information

Machine type / model:	002084 / A08-305	Manufacturer:	IBM
Machine serial:	02 - 0026A3A	CPC serial:	000020026A3A
Machine sequence:	000000026A3A	CPC location:	A19B
Plant of manufacture:	02	CPC identifier:	00

Save Change Options Defaults Cancel Help

The z990 SE CPC details panel Machine type information, shows the machine model, the: Hardware model number, and a Software model number

For this CPC the:

Hardware model number = A08

Software model = 305

Also 3 PUs are characterized as 3 ICFPs



```

D M=CPU
IEE174I 12.31.47 DISPLAY M 711
PROCESSOR STATUS
ID  CPU                      SERIAL
0    +                      01796A2084
1    +                      01796A2084
2    +                      01796A2084
3    +                      01796A2084
4    +                      01796A2084
5    +                      01796A2084
6    +                      01796A2084
7    +                      01796A2084
8    +                      01796A2084
9    +                      01796A2084
A    +                      01796A2084
B    +                      01796A2084
C    +                      01796A2084
D    +                      01796A2084
CPC ND = 002084.332.IBM.02.00000001796A
CPC SI = 2084.332.IBM.02.000000000001796A
CPC ID = 00
CPC NAME = TC4P01
LP NAME = TC4P01      LP ID = 1
CSS ID = 0
MIF ID = 1
+ ONLINE      - OFFLINE      . DOES NOT EXIST      W WLM-MANAGED
N NOT AVAILABLE
CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
CPC SI  SYSTEM INFORMATION FROM STSI INSTRUCTION
CPC ID  CENTRAL PROCESSING COMPLEX IDENTIFIER
CPC NAME CENTRAL PROCESSING COMPLEX NAME
LP NAME  LOGICAL PARTITION NAME
LP ID    LOGICAL PARTITION IDENTIFIER
CSS ID   CHANNEL SUBSYSTEM IDENTIFIER
MIF ID   MULTIPLE IMAGE FACILITY IMAGE IDENTIFIER

```

z/OS Command
'D M=CPU' displays
software model numbers

Hardware model number
incorrectly shows the software
model number. IBM will change
to correctly show the hardware
model number

Software bmodel number

LP name displayed in error
IBM will change this show
the correct CPC name



Display of HMC CPC details panel (IBM Test System)

Note: Machine type / models where both the hardware and software model values are shown

G18 Details

Instance information

CP Status:	Operating	Activation profile:	G18
CHPID Status:	Exceptions	Last used profile:	not set via Activate
Group:	CPC	Service state:	Disabled
IOCDS identifier:	A2	Maximum CPs:	32
IOCDS name:	TC4P610	Maximum ICFPs:	0

Lockout disruptive tasks: ☐ Yes ☒ No

System mode: Logically partitioned
Alternate SE Status: None

Dual AC power maintenance: Fully Redundant
CP Assist for Cryptographic Functions: Not Installed

Acceptable CP/CHPID status

<input checked="" type="checkbox"/> Operating - ■	<input type="checkbox"/> Power save - ■	<input type="checkbox"/> No power - ■
<input type="checkbox"/> Not Operating - ■	<input type="checkbox"/> Exceptions - ■	<input type="checkbox"/> Status check - ■
<input checked="" type="checkbox"/> Acceptable - ■	<input type="checkbox"/> Service Required - ■	<input type="checkbox"/> Degraded - ■

Product information

Machine type / model:	002084 / D32-332	Manufacturer:	IBM
Machine serial:	02 - 001796A	CPC serial:	00002001796A
Machine sequence:	00000001796A	CPC location:	A19B
Plant of manufacture:	02	CPC identifier:	00

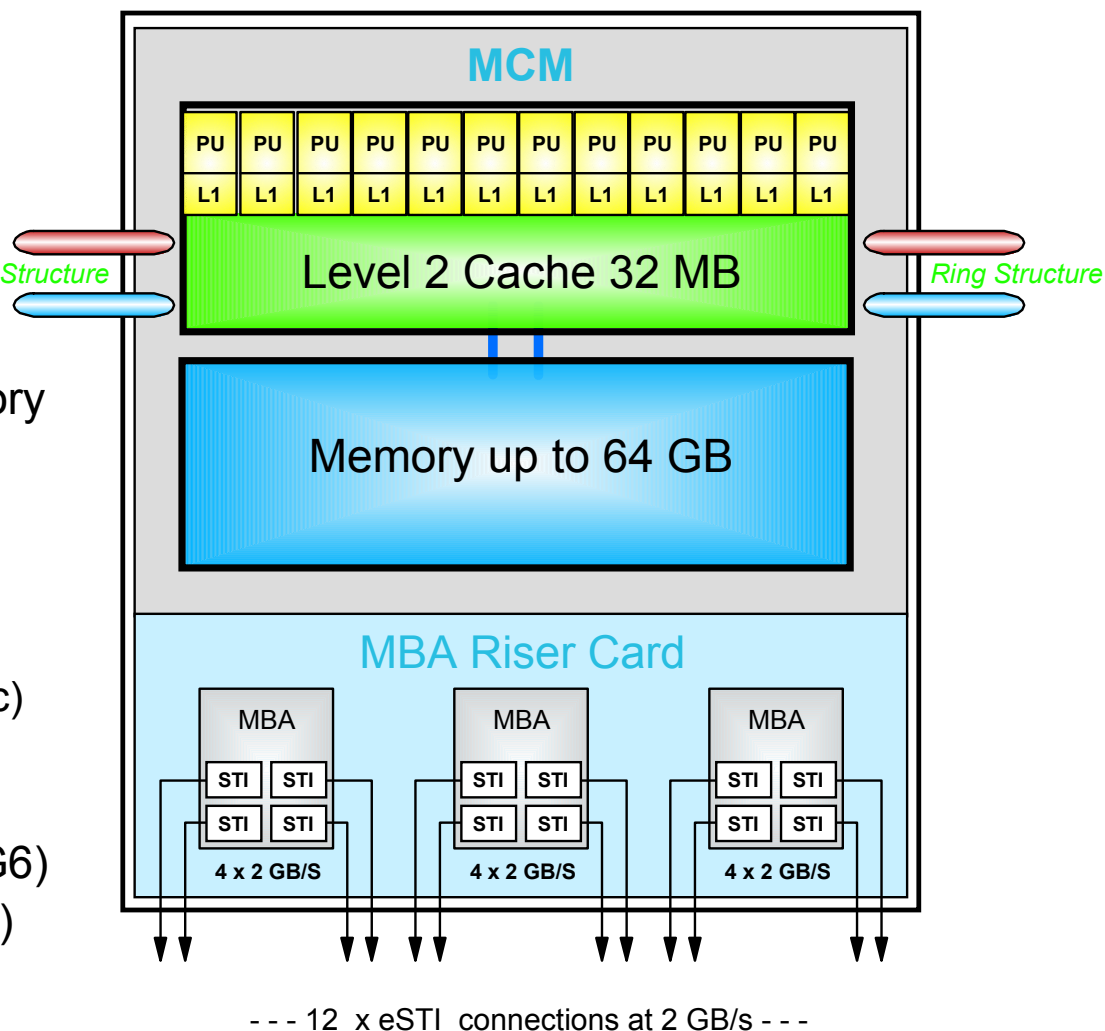
Save Change Options... Upgrade Reasons... Cancel Help



All processors and memory in all books appear as a single Symmetric Multi-Processor (SMP) system

- 12 Processor Units (PU)
 - 8 characterizable PUs
 - 2 SAPs
 - 2 Spare
 - L1 Cache 512 KB (I/D) per PU
- L2 Cache 32 MB
 - Memory subsystem dual ring communication to other Books
- From 16GB (8GB) to 64GB of memory
- 3 MBAs with 12 eSTIs in total eSTIs can support:
 - I/O Cage Domains
 - ▶ 4 I/O slots per I/O domain
 - I/O Cards (ESCON, FICON, etc)
 - Crypto Cards
 - ICB Extender Card
 - ▶ ICB -2 Extender (z990 to G5/G6)
 - ▶ ICB -3 Extender (z990 to z900)
 - ICB-4 connection
 - ▶ To a z990 ICB-4

z990 Processor Book





The type of Processing Units (PUs) that can be ordered (enabled / assigned) on a z990 server include:

- **Central Processors (CPs)**

- ▶ Provides processing capacity exclusively for z/Architecture and ESA/390 instruction sets
- ▶ Runs z/OS, z/VM, VSE/ESA and TPF

- **Integrated Facility for Linux (IFL)**

- ▶ Provides additional processing capacity exclusively for Linux workloads
- ▶ Runs Linux or Linux under z/VM Version 4

- **Internal Coupling Facility (ICF)**

- ▶ Provides additional processing capacity exclusively for the execution of the Coupling Facility Control Code (CFCC) in a CF LPAR


- **System Assisted Processors (SAPs)**

- ▶ SAPs manages the start and ending of I/O operations for all logical partitions and all attached I/O

Note: There are two standard spare PUs in each book. If these have been used and another failure occurs in a characterized PU, then an unassigned PU will temporarily be used to replace the failing PU.



- 12 PUs per book
 - 2 spares
 - 2 standard System Assist Processors (SAPs)
 - 8 PUs for characterization
- Orderable Characterization Features
 - One active CP, ICF or IFL is required
 - CP - FC #0716 (Range: 0 - 32)
 - Unassigned CP - FC #1716 (Replaces downgrade RPQ)
 - Active IFL - FC #0516 (Range: 0 - 32)
 - Unassigned IFL - FC #0517 (Replaces downgrade RPQ)
 - ICF - FC #0518 (Range: 0 - 16)
 - Optional SAP - FC #0519 (Range: 0 - 31)
- Software "3xx" model for software pricing
 - Indicates number of CPs (count of FC #0716)
 - Model 300 for no CPs (ICFs, IFLs only)
 - Model 301 to 332 for 1 to 32 CPs
 - Noted on configuration report
 - Reported by the STSI Instruction






- New Option for z990
 - Feature Code #1716 - "Unassigned CP"
 - Feature Code #0517 - "Unassigned IFL"
- Avoids downgrade RPQ and disruption
- Objective: Purchase capacity with software cost flexibility
- Current S/390 and zSeries Method:
 - Order model
 - Order downgrade RPQ (Disruptive if already installed)
 - Result --> smaller model for S/W cost
 - To restore original model --> Order upgrade RPQ
 - Required special processes, manual effort and tracking
- z990 Enhancement
 - eConfig offers option to order "unassigned" engines
 - Contained within current order process
 - No disruption on new-build machine
 - With proper planning, order nondisruptive upgrade later



- Followed for new build or MES upgrade
- Book order: 0, 1, 2, 3
- Fill each book (8 characterized PUs) before moving to a new book
- Order of assignment
 - CPs
 - IFLs
 - ICFs
 - Optional SAPs
 - Unassigned CPs
 - Unassigned IFLs
- Example configuration: z990 Model B16, with 6 CPs, 1 Unassigned CP, 2 IFLs, and 1 ICF
 - Book 0 - 6 CPs, 2 IFLs (2 SAPs, 2 spare PUs)
 - Book 1 - 1 ICF, 1 Unassigned CP, 6 uncharacterized PUs (2 SAPs, 2 spare PUs)





- CEC with a single MCM (single Book)
 - Two spare PUs per MCM
 - When a 2 Processor chip fails it is fully "spared"
 - Even if only one processor fails
 - When a single Processor chip fails it may be "spared" by a 2 processor chip
 - One replaces the failing chip, the second remains available as a spare
 - 'Call Home' initiated when the first non-reserved PU is used as a spare
 - Try to avoid customer impact
- CEC with Multiple MCMs (Multiple Books)
 - The rules of the single MCM apply
 - The Spares on the MCM with the failing PU are used first
 - Off MCM Sparing is used when there are no spares left on the MCM

of PU failures required to generate an RSF call home

Failure On a:	A08	B16	C24	D32
2 PU Chip	2	3	4	5
1 PU Chip	3*	5*	7*	9*
	2-3	3-5	4-7	5-9


* Running same logical on less physical

* Single dedicated PU (CP, IFL, ICF) logical partitions are not recommended for production logical partitions

Provided there are sufficient functioning PUs available: concurrent Book and PU MES are honored

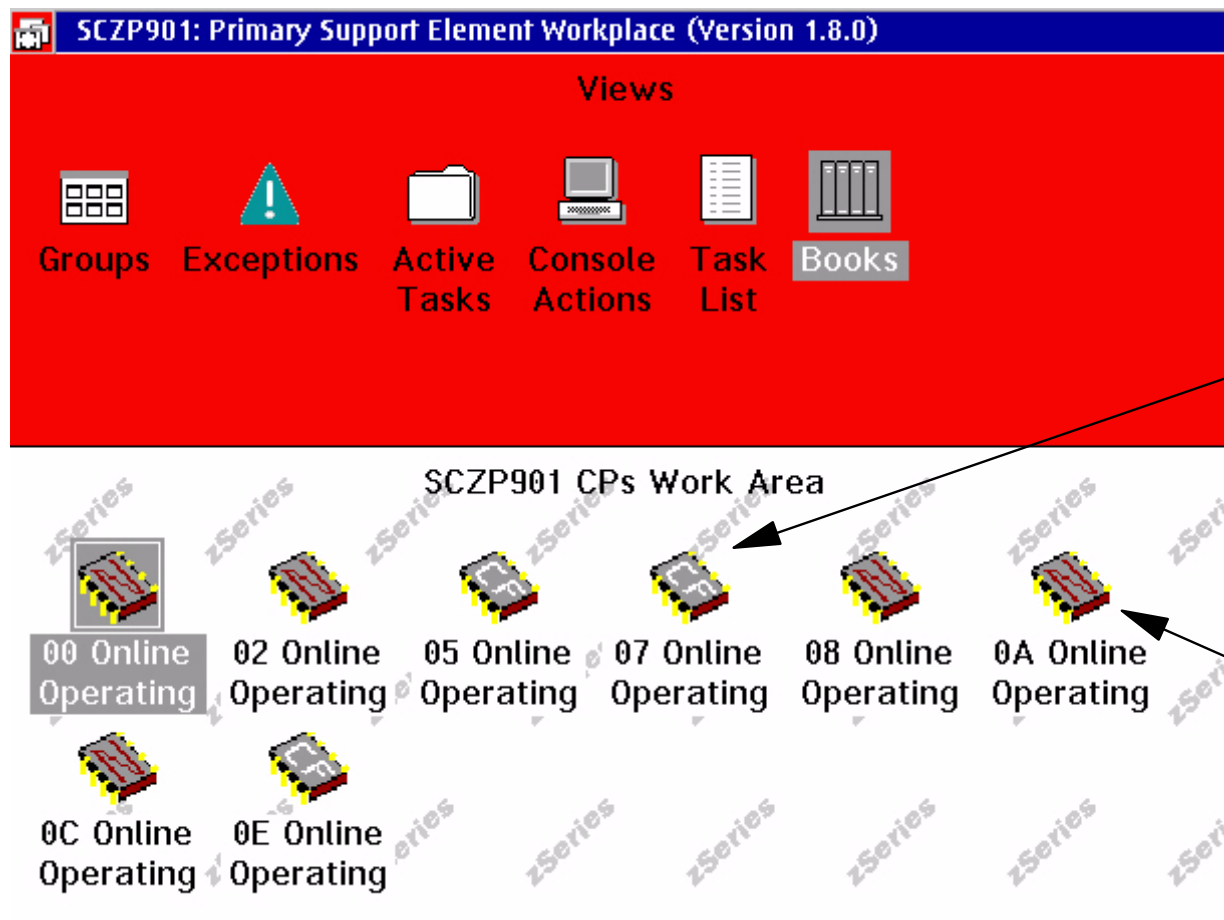


- Advanced 93mm x 93mm MCM
 - 101 Glass Ceramic layers
 - 16 chip sites, 185 capacitors
 - 0.4 km of internal wire
 - 46% smaller than IBM @server zSeries 900 (z900) MCM
 - 23% more I/O connections
 - 133% I/O density improvement
 - 5184 LGA connectors vs 4224 pins for z900



Note: MBAs are not on the MCM


- CMOS 9S-SOI chip Technology
 - PU, SC, SD, and MSC
 - Copper interconnections, 8 copper layers
 - 8 PU chips/MCM (Single and Double PUs)
 - 14.1 mm x 18.9 mm
 - 122 million transistors/PU
 - L1 cache/PU
 - 256 KB I-cache
 - 256 KB D-cache
 - 0.83 ns Cycle Time
 - One Storage Control (SC) chip
 - 17.3 mm x 17.3 mm
 - 98 million transistors
 - 3692 Power Signal I/Os
 - L2 cache crosspoint switch
 - L2 access rings to/from other MCMs
 - L2 access to/from MBAs (off MCM)
 - 4 System Data (SD) cache chips/MCM
 - 17.5 mm x 17.5 mm
 - L2 cache per book
 - 521 million transistors/chip
 - 32 MB
 - Two Memory Subsystem Control (MSC) chips
 - Controls Memory cards (L3) interface to L2
 - One Clock (CLK) chip - CMOS 8S
 - Clock and ETR Receiver



2084 model A08 / 305 (5 x CP) plus 3 x ICF

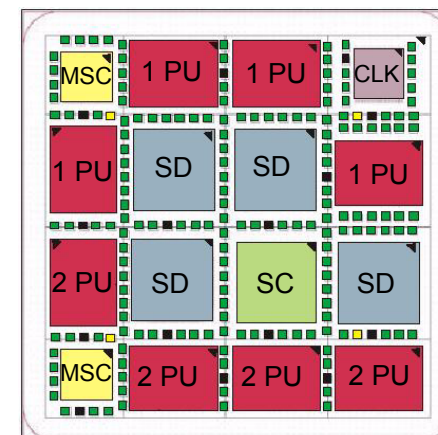
This panel only shows optional characterized PUs
Standard SAPs, Spare PUs, and uncharacterized PUs are not shown

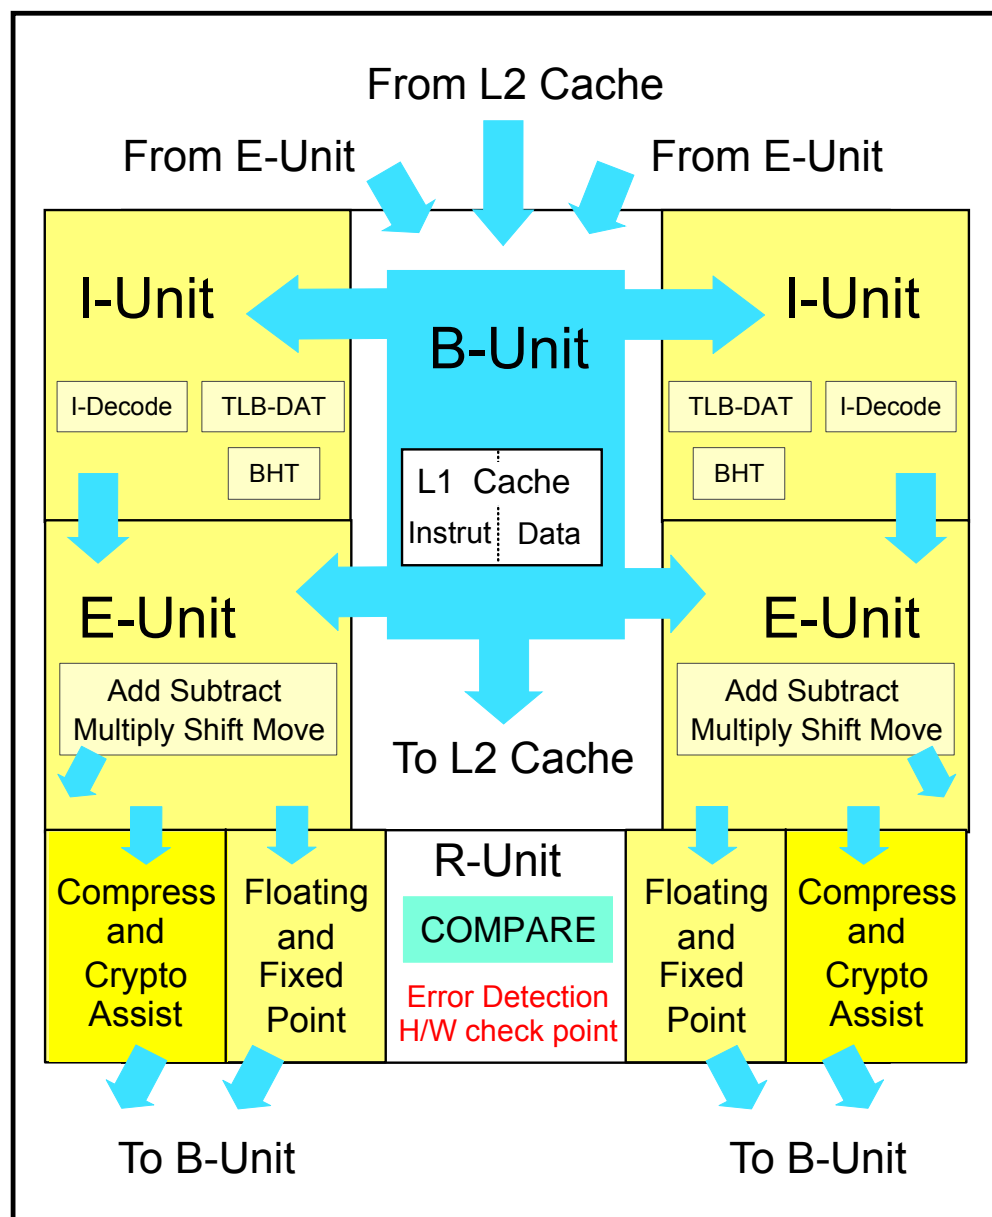
- Interleaved memory
- Dual memory subsystem controllers
- Integrated storage controller
 - ▶ Multiple paths
- Densest L2 Chips
 - ▶ 2.2 times more transistors than z900
- Support for denser PUs





- Multiple instructions per cycle
 - Decode up to two instructions per cycle
 - Execute up to three instructions per cycle
 - In-order instruction issue (execution)
 - Out-of-order instruction and operand fetch
 - Up to two operand fetches or one fetch and one store per cycle
 - Up to 4 L1 cache misses can be processed while servicing L1 cache hits
- Additional PU design improvements
 - New hardware unit in PU for compression, crypto assist, and translation
 - Support for more than 15 LPAR zones for logical partitions
 - Asymmetric mirroring for error detection
 - Allows faster execution with full execution result checking one cycle later
 - Second-level Translation Lookaside Buffer (TLB2)
 - Second-level dynamic address translation (DAT) cache for L1 cache TLBs
 - TLB2 is two-level: **512** entries, 4K entries (First-level TLBs have 512 entries)
 - New 20-bit signed displacement (+/- 512 KB) instructions (RXY, RSY, SIY formats)
 - Improved floating point support
 - Optimized IEEE arithmetic (Java exploitation)
 - New hexadecimal floating-point multiply-and-add type instructions





- Processing Unit (PU)
 - Asymmetric mirroring for error detection
 - I-Unit
 - E-Unit
 - General instruction processing
 - Compress function
 - Crypto assist
 - Floating Point function
 - Complete error detection mechanism
 - Data flow - parity checked
 - Address paths - parity checked
 - L1 Cache - parity checked
 - L1 store through (to L2 store-in)
 - Processor logic (I - E - C/C -F)
 - Duplicated, then compared output
 - Error detection for mis-compare
 - Re-execute from h/w check point
 - Spare for high intermittent solid failure



- Instruction fetching and instruction decode
 - The superscalar design of the z990 microprocessor allows for the decoding of up to two instructions per cycle and the execution of three instructions per cycle. Storage accesses for instruction and operand fetching may occur out-of-sequence, but instruction execution takes place in order.
- Instruction fetching
 - Instruction fetch in non z990 models tries to get as far ahead of instruction decode and execution as possible because of the relatively large instruction buffers available. In the z990 microprocessor, smaller instruction buffers are used. The operation code is fetched from the I-cache and put in instruction buffers that hold pre-fetched instructions awaiting decode.
- Instruction decoding
 - The processor can decode one or two instruction per cycle. The result of the decoding process is queued and subsequently used to form a group.
- Instruction grouping
 - From the instruction queue, one simple branch instruction and up to two general instructions can be issued every cycle. The instructions are taken from the instruction queue and grouped together. The instructions are assembled according to the z990 instruction grouping rules.
 - Compilers should select instructions that best fit with the z990 superscalar microprocessor and abide by the grouping rules to create code that best exploit the superscalar implementation.



- Additional features of the PU cores are:

- **Second Level Translator Lookaside Buffer (TLB)**
 - This feature provides a secondary cache for the DAT, both for the instruction and data caches.
 - The secondary buffers provide an additional 4K of dynamic address mapping over the current 512 entries now available.
- **Programmable Dynamic Address Translation (DAT) translator**
 - This feature allows for updates to the address translation algorithm (mapping virtual addresses to physical memory locations) without any longer requiring silicon changes to implement.
 - The various steps of the translation are stored in a RAM and can be changed on the design process or even in a customer's office if necessary.
 - This feature allows for modifications to the current algorithm to invoke performance changes as well as making the translator extendable to future architectural changes.
- **Floating-point performance optimized for IEEE arithmetic (JAVA exploitation)**
 - Floating point pipeline has now been fully pipelined at 5-stages for both IEEE and legacy Floating Point arithmetic. Each cycle is now utilized as compared to every other in the past.
- **On-chip Cryptographic Assist Functions for eBusiness**
 - This area hardware implemented encryption/decryption used to speed up SSL transactions, VPN data transfer and data repository applications.
- **Asynchronous processor <-> SCE interface**
 - This asynchronous interface was modeled after the MSC <-> SMI interface.
 - It's purpose is to decouple the processor and SCE cycle times allowing each to run at their optimum speed and not forced into fixed multiples of one another.
 - This interface serves to buffer the exchange of data between the two disparate cycle times.
- **Asymmetric Mirroring for Error Detection**
 - This feature provides a mechanism by which the mirrored execution used for error detection occurs 1 cycle delayed from actual operation.
 - This allows the operation execution logic and circuitry to be optimized for performance
 - This comparison execution logic and circuitry can now be implemented at the periphery of the chip where it does not sub-optimize the actual operation implementation.



- Compression Unit on a chip

- Each z990 PU has a Compression Unit on the PU chip, providing better hardware compression performance. The Compression Unit is integrated with the CP Assist for Cryptographic Function benefiting from combining the use of buffers and interfaces.

- CP Assist for Cryptographic Function

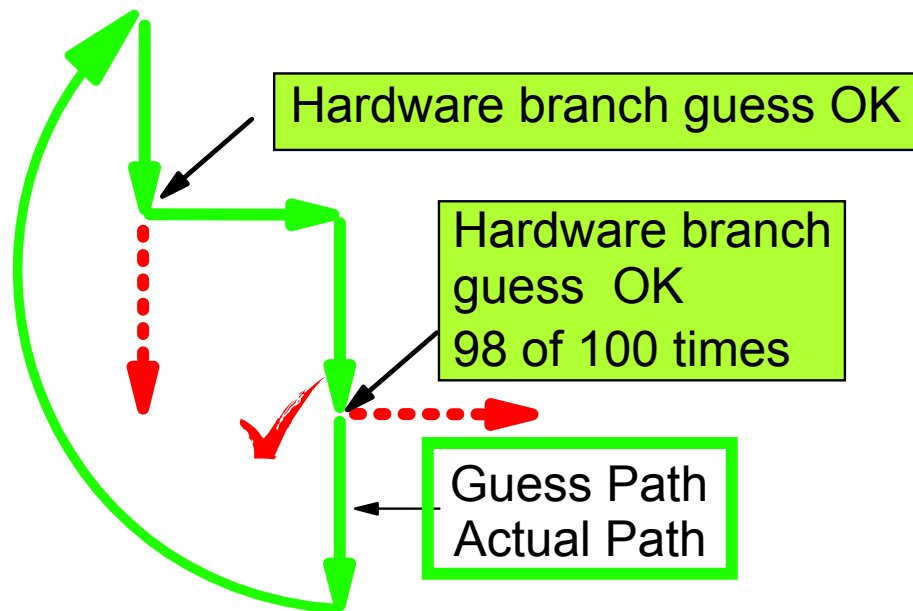

- Each z990 PU has a CP Assist for Cryptographic Function on the chip. The assist provides high performance hardware encryption and decryption support for clear key operations. Five new instructions are introduced with the cryptographic assist function.
- The CP Assist for Cryptographic Function offers a set of symmetric cryptographic functions that enhance the encryption and decryption performance of clear key operations for SSL, VPN and data storing applications that do not require FIPS 140-2 level 4 security. The cryptographic architecture includes DES, T-DES data encryption and decryption, MAC message authorization and SHA-1 hashing.
- The CP Assist for Cryptographic Function complements public key (RSA) functions and the secure cryptographic operations is provided by the PCIXCC cryptographic coprocessor card.



Branch History Table Operation

Without BHT:

With BHT:





● Branch History Table Operation

- The Branch History Table (BHT) implementation on processors has a key performance improvement effect.
 - ▶ The BHT was originally introduced on the IBM ES/9000 9021 in 1990 and has been improved ever since.
- The z990 server BHT offers significant branch performance benefits. The BHT allows each CP to take instruction branches based on a stored BHT, which improves processing times for calculation routines.
- Using a 100-iteration calculation routine as an example (see previous foil)
 - ▶ Hardware preprocesses the branch incorrectly 99 times without the BHT
 - ▶ With a BHT, it preprocesses branch correctly 98 times
- Without BHT, the processor:
 - ▶ At the second branch point (in the example), the CP makes an incorrect branch guess the first time through the loop
 - ▶ Preprocesses instructions for the wrongly guessed branch path (these eventually will not be used)
 - ▶ When the CP discovers the branch not equal to guess, it needs to fetch and then preprocess a new path
 - ▶ Repeats this 98 more times until the last time when the guess matches the actual branch taken.
- With BHT, the processor:
 - ▶ Makes an incorrect branch guess the first time through the loop, at the second branch point in the example.
 - ▶ Preprocesses instructions for the wrongly guessed branch path (these eventually will not be used)
 - ▶ When it discovers the branch not equal to guess, it needs to fetch and then preprocess a new path
 - ▶ Updates the BHT to indicate the last branch action taken at this address.
 - ▶ The next 98 times at this branch point the branch taken will come from the BHT
 - ▶ The last time the guess is wrong.



z990 single book structure

● Interconnection Structure:

- PU

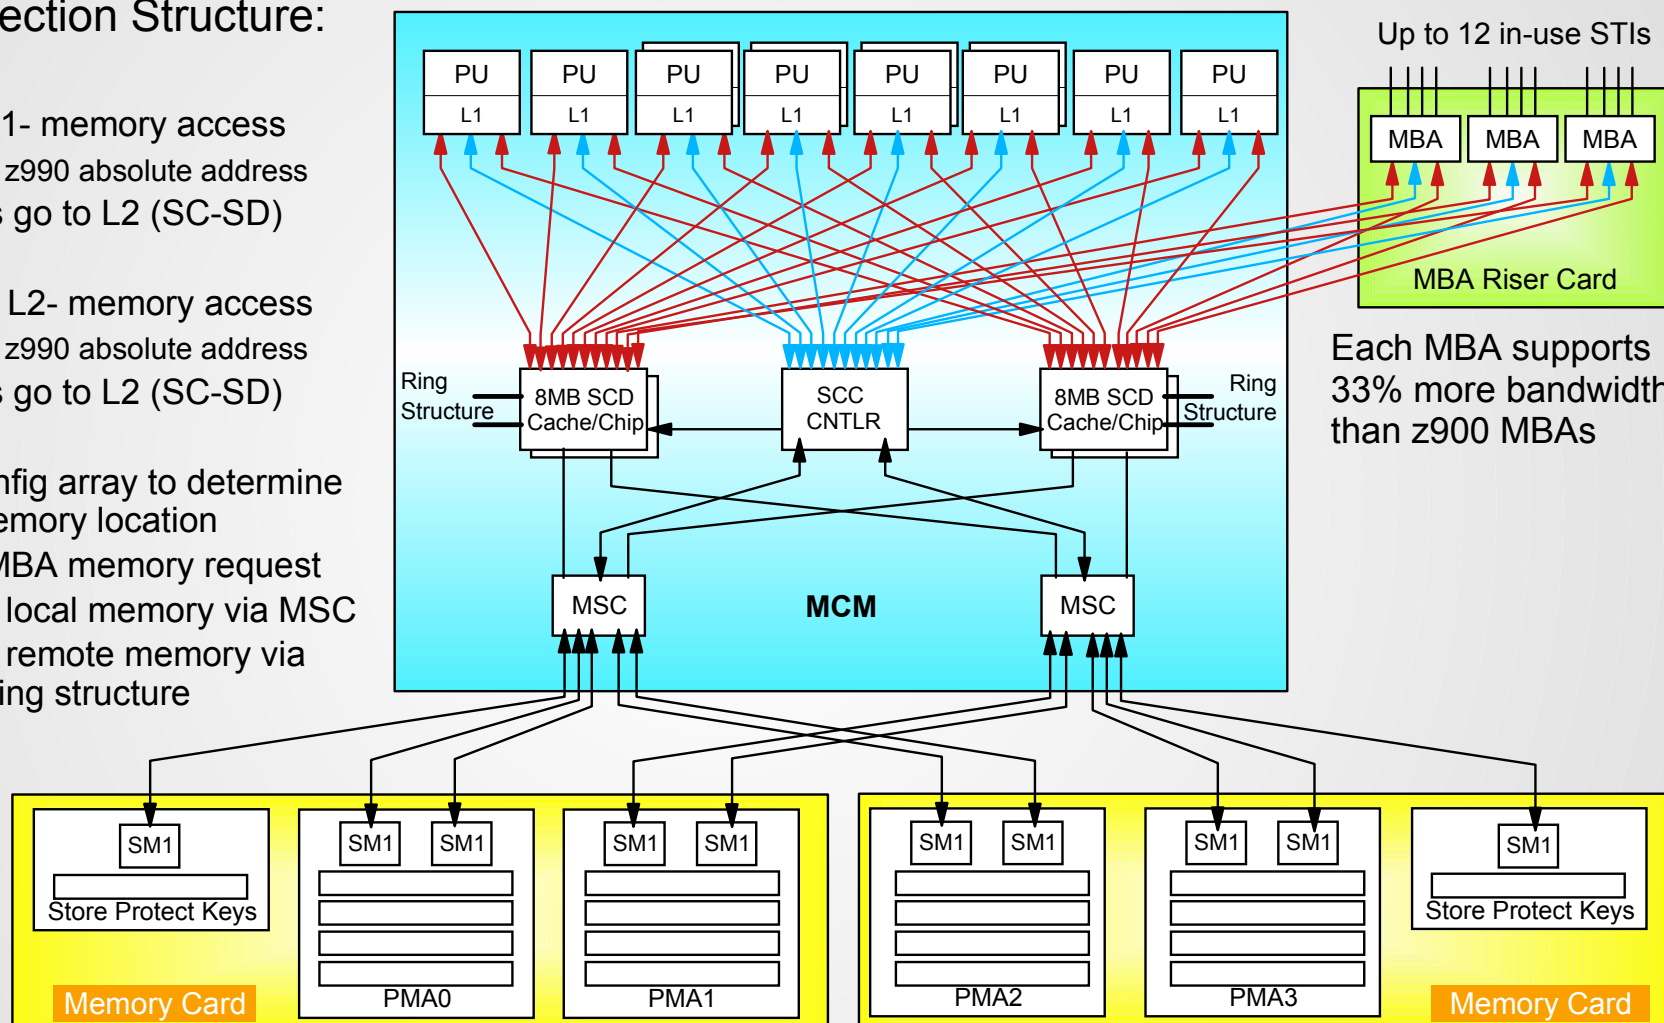
- ▶ PU to L1- memory access
 - Using z990 absolute address
- ▶ L1 miss go to L2 (SC-SD)

- MBA/STI

- ▶ MBA to L2- memory access
 - Using z990 absolute address
- ▶ L1 miss go to L2 (SC-SD)

- L2 miss

- ▶ Use config array to determine data memory location
- ▶ PU or MBA memory request
- ▶ Access local memory via MSC
- ▶ Access remote memory via Book Ring structure

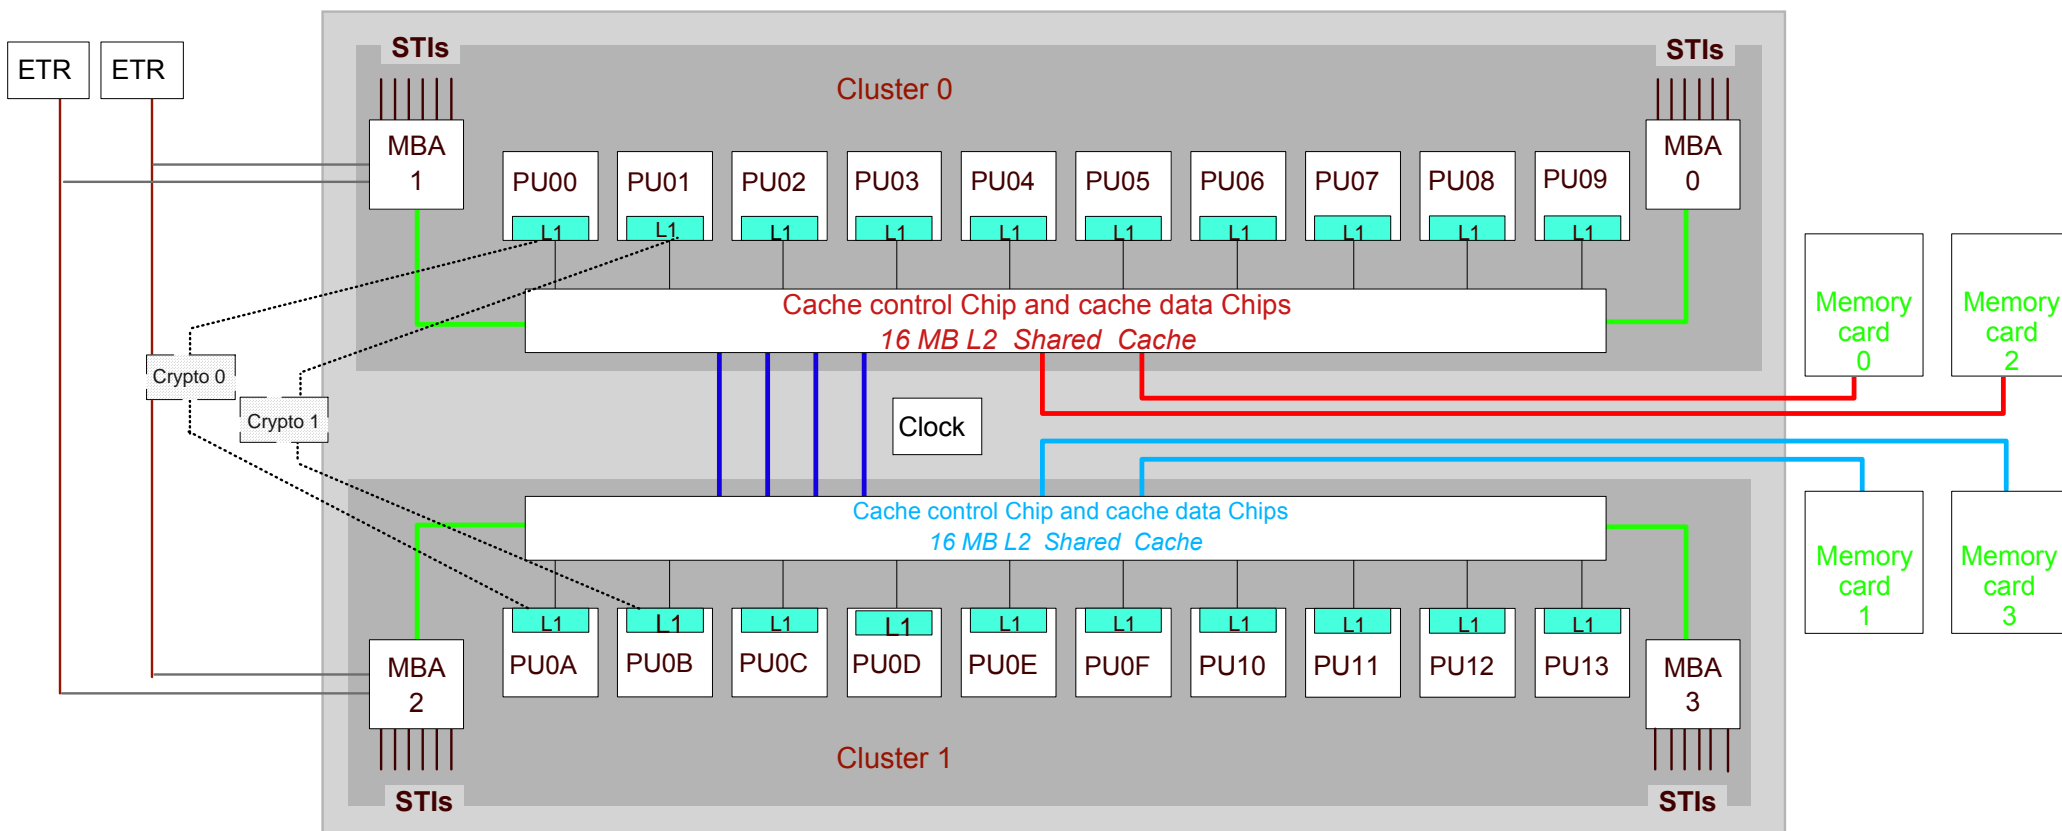




- z990 CEC structure verus the z900 CEC structure
 - z990 supports from 1 to 4 books in a CEC (1 MCM in each book)
 - z990 supports 12 PUs in one Book with the 1 MCM in the Book
 - 12 PUs, upto 64GB Memory and 3 MBAs / 12STIs are all in one book
 - z900 supports 2 clusters of 10 PUs per cluster all in 1 MCM
 - Memory and MBAs/STIs are split between the books



zSeries 900 - Clusters



IBM has built on the approach of multiple z900 clusters into multiple z990 books, while maintaining coherency of all the processor, memory and I/O operations



zSeries 990 - Book to Book, and Book to Channel Connectivity


Book 3

Book 0

CEC Cage

Book 1

Book 2





- z990 Processor Book connectivity
 - Inter-Book connectivity
 - Inter-Book ring communication
 - Supported by the CEC cage board (and as required jumper Books)
 - Used for cross memory communication
 - SCE) to remote memory system controller (via SCE)
 - Used PU to PU communication
 - Book CP to Book CP
 - Book CP to Book SAP, and Book SAP to Book CP
 - Book SAP to Book SAP
 - PU to memory communication
 - Book CP to another Book L2 or MSC
 - Book to Channel connectivity - SAP to STI
 - Support by 12 STI connections per installed Book
 - STI connections are used for to/from communication
 - Book SAP to I/O domain
 - Book STI to ICB-2 extender card
 - Book STI to ICB-3 extender card
 - z990 ICB-4 to z990 ICB-4 communication



- CP data access
 - CP to L1- memory access
 - Using z990 absolute addresses
 - L1 hit - provides data (instruction or operand data) to the CP
 - No L1 hit (L1 miss) go to L2 (SC-SD)
 - Local L2 cache hit
 - Provide data to L1
 - No L2 hit (L2 miss) go to SCE
 - SCE to L2 cross interrogate (for required data)
 - Data in another L2
 - Provide to requesting L2
 - Data not in any L2
 - Get data from Main Storage (via config array addressing lookup)
 - Use SCE config array to determine the data's main storage memory location
 - Access config array with z990 absolute address
 - Access local memory via local MSC
 - Access remote memory via Book Ring structure

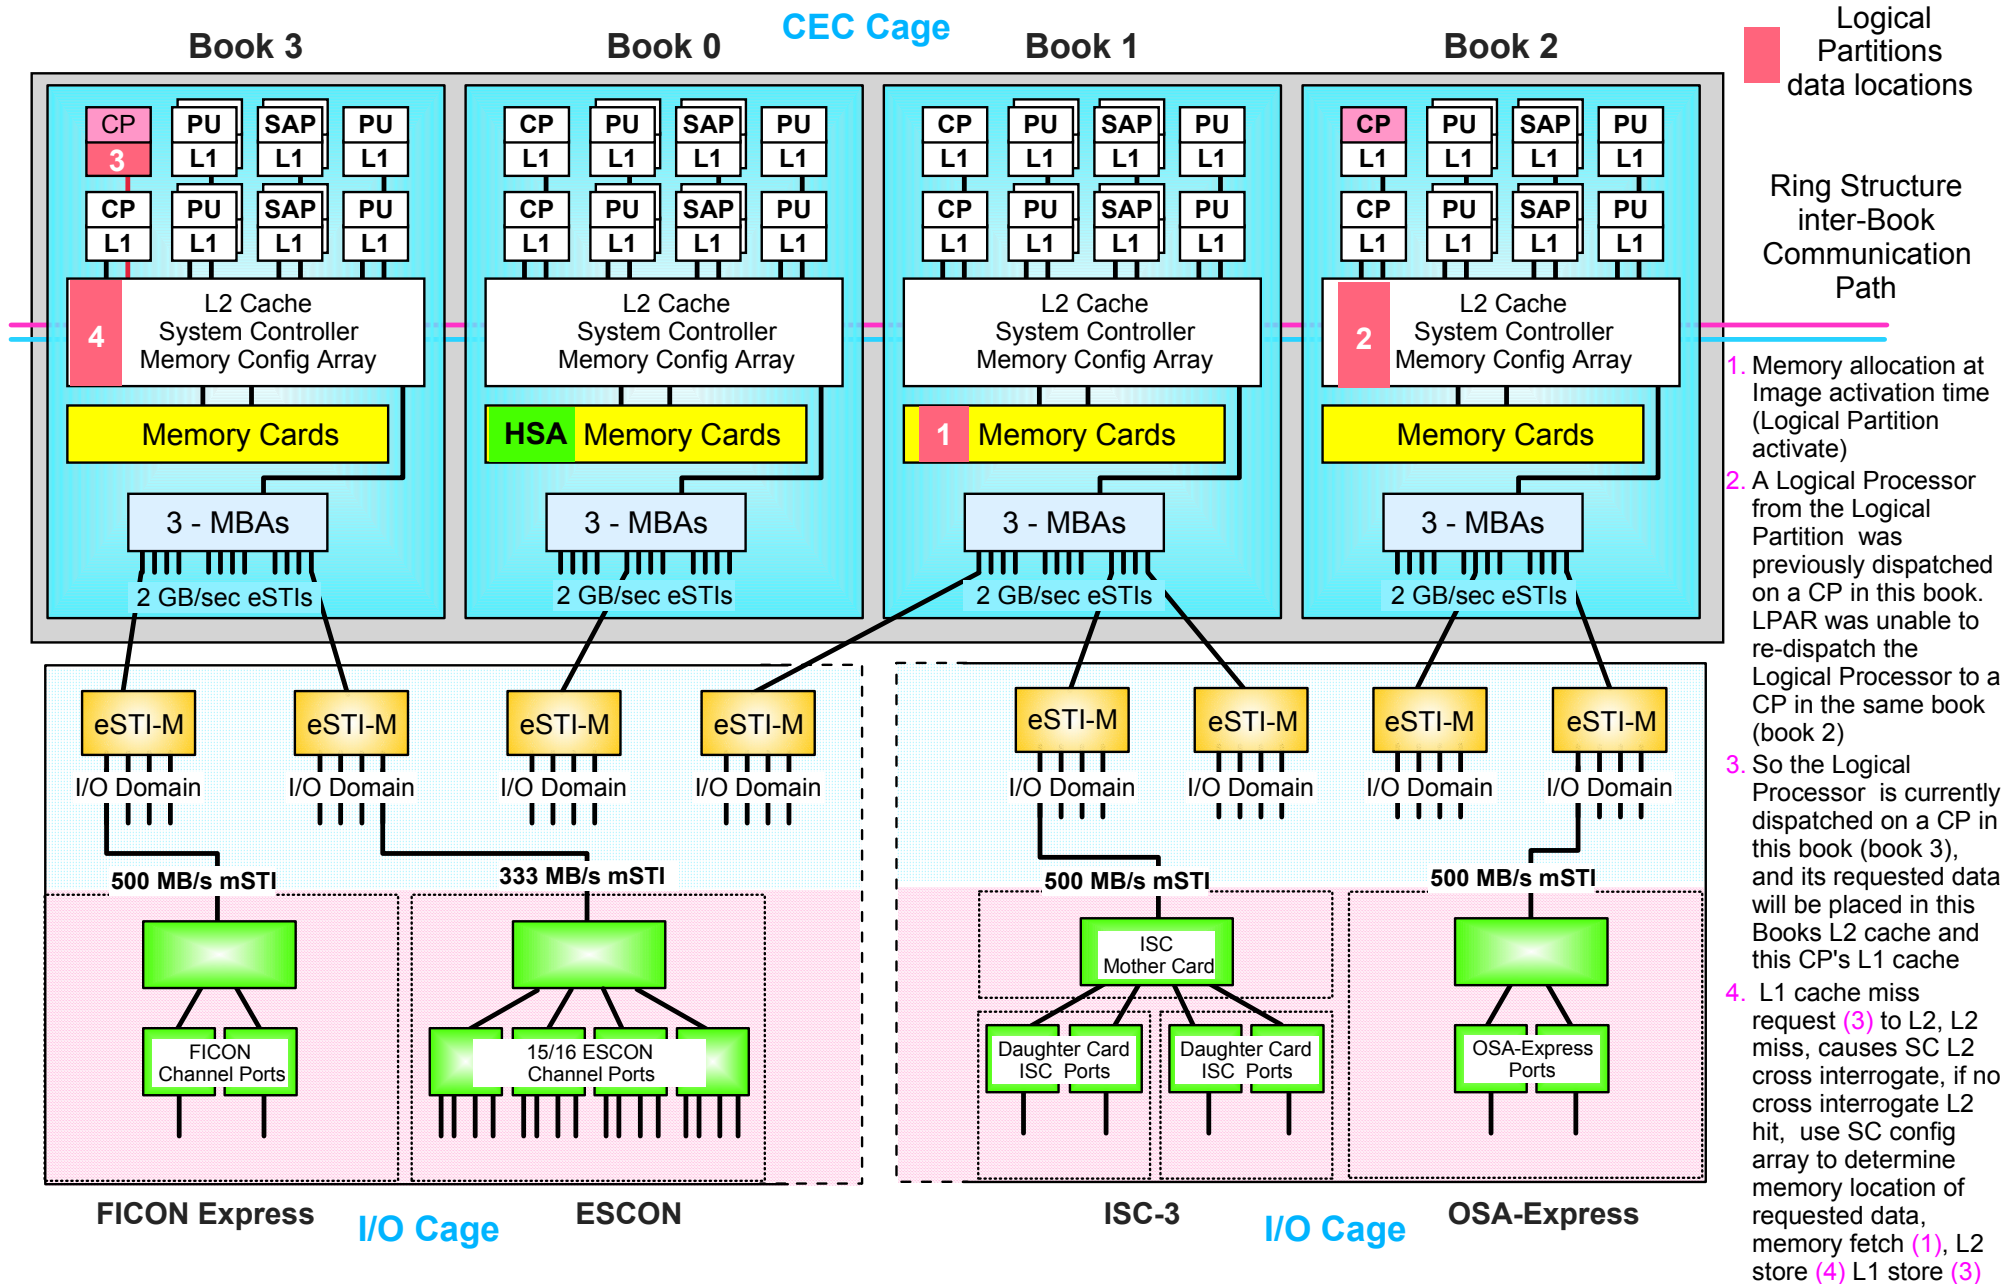


- Logical Partition Data

- Data for a logical partition can reside in a number of different storage areas of the z990 processor (CPC)
 - ▶ L1 Cache
 - The L1 cache acts as an immediate high speed data depository for a logical processor (a dispatched logical processor) running in a CP
 - Data for other logical processors previous dispatched on this PU L1 may age
 - ▶ L2 cache
 - The L2 cache acts as a processor book immediate high speed data depository for all the PUs *in the same book*
 - Data for a Logical Processor, if not in the physical processor L1 cache (L1 cache miss) may be in any L2 cache or main memory in the z990
 - ▶ Main Memory (Central Storage)
 - A logical partitions data (Instructions and Data) that currently resides in Real Address memory addressing will always be in Main Storage. However latest copy may be in any L2 cache in the z990.
 - The SCE in the same book as the dispatched logical processor determines which Main Memory to access (if requested data is in none of the L2 caches) by the use of the SCE config array (each SCE has a Main Storage config array)
 - Memory is assigned to a logical partition when the logical partition is activated or reserve storage is brought online



zSeries 990 CPC Logical Structure - Logical Processor - Data Access





1. Memory allocation at Image activation time (Logical Partition activate)
 - LPAR tries to allocate all in the same book
2. A Logical Processor from the Logical Partition was previously dispatched on a CP in this book. LPAR was unable to re-dispatch the LP to a CP in the same book
 - LPAR tries to dispatch a logical processor in the same book where it was last running
3. For this example the Logical Processor is currently dispatched on a CP in book, which is not the book where it last was dispatched, and also not in the book where the logical partitions memory resides. Data requests for this CP will eventually be stored in this CPs L1 cache
4. If an L1 cache miss occurs, then:
 - A request (3) is made to L2
 - If a L2 miss occurs, the SC will perform an L2 cross interrogate
 - If no L2 hit,
 - Use SC config array to determine memory location
 - Memory fetch (1), L2 store (4) L1 store (3)



The z990, inter-book ring structure, PU design and operation, PU L1 cache management (I/D), PU TLB, L2 cache management (processor cache), Image activation memory assignment, SC configuration address management (LP memory address to storage location mapping), and LPARs logical processor dispatching algorithm, all add up to a highly efficient single structured coherent SMP, and along with its I/O connectivity structure makes the z990 adaptable to many different processing environments.



- There are 2 levels of cache (L1 and L2)
- Level 1 Cache
 - Each processor chip now contains 1 or 2 PUs (Processing Units)
 - Each PU has its own private L1 Cache (512KB I/D)
 - The L1 cache is physically located on the processor chip, and is subdivided into I-Cache (for instructions) and D-Cache (for data)
 - L1 is addressed using the z990 absolute address
 - The L1 cache uses a Store-through policy, meaning all updates (changes) are immediately sent to the L2 cache
 - This serves 2 purposes:
 - Availability in event of a hardware defect in the L1 cache: a copy of the data is always available in the L2 cache (or main storage) for recovery.
 - Allows for dynamic CP chip sparing, by insuring that no storage updates are lost in the L1 if a physical processor chip (and its associated L1) is removed from the configuration (spared)



- Level 2 Cache

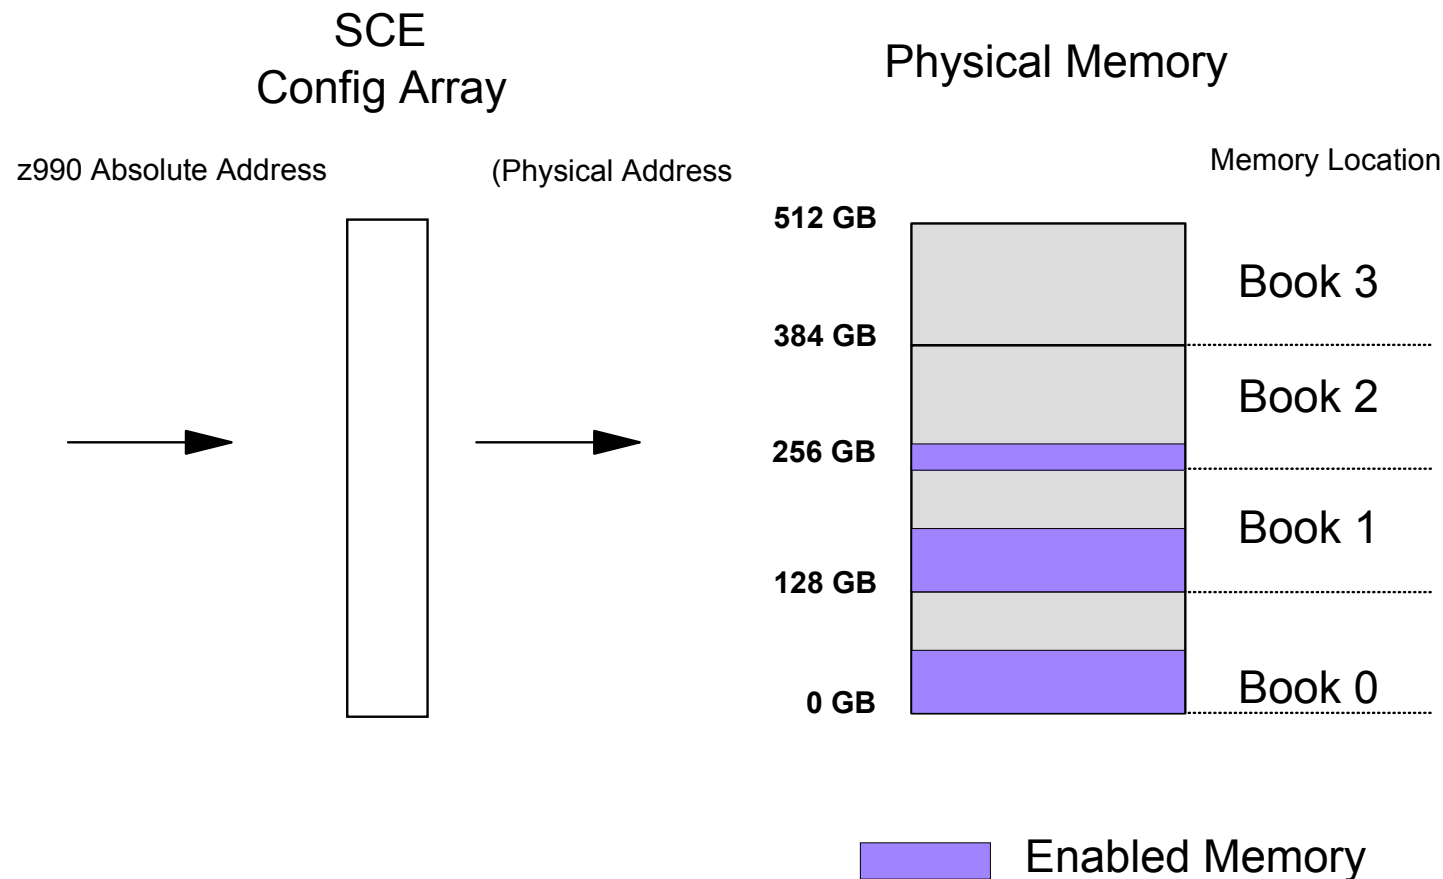
- The SCE provides a shared L2 Cache for use by all the processors in the same book, and (to a lesser extent) by I/O
 - ▶ The L2 cache is fundamentally a processor cache, not a memory cache
 - ▶ This means that Data is normally resident in the L2 cache on the book where it is being used by a processor PU - CP), not in the book where the associated memory address resides
 - ▶ The L2 cache makes no distinction between instructions and data; unlike the L1 cache, everything is stored in one cache and is treated in exactly the same way



- The SCE provides 4 basic functions in the z990:
 - It interconnects 4 books of 12 processing units (PUs) by way of a ring fabric
 - It provides a 2nd level shared cache between the private L1 cache in each processor unit and the memory cards
 - It acts as a cache for certain I/O data accesses by the MBA (STI and channel)
 - It provides data management functions for data sharing between processors and I/O, with the responsibility to maintain data coherency so that each processor is always working with the most recent copy of data




- On the z990 256GB of maximum physical address space is sliced into 4 contiguous blocks where each block is assigned to the memory in the processor book)

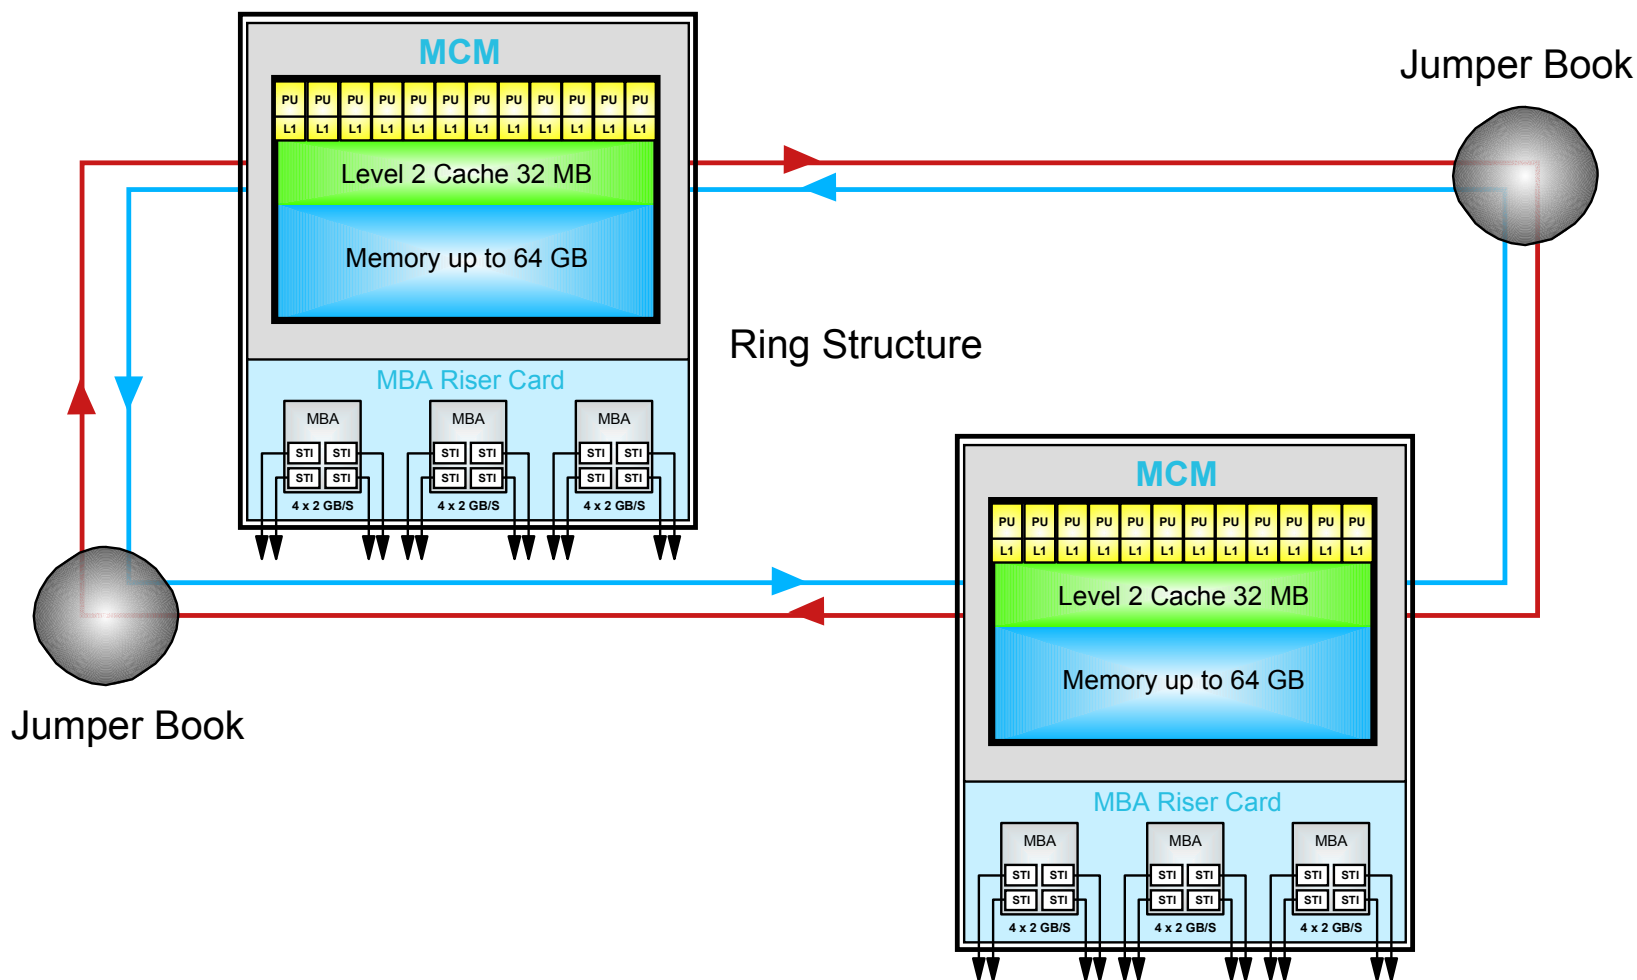




- Access from one SCE to another SCE is via the inter-book ring structure communication path




- Concentric communication loops or rings are constructed such that in a 4-book system each book is only connected to two others
- A memory-coherent director optimizes ring traffic and filters out cache traffic



- The ring structure consists of two rings (one running clockwise, the other counterclockwise)
- In a two or three Book configuration, jumper Book(s) (installed in the CEC cage) complete the ring
 - Jumper Books are not needed for a single-Book configuration
- Additional Books can be inserted into the ring non-disruptively
 - The jumper book is removed and replaced by a Book





2084 z990 model A08 (1 book)




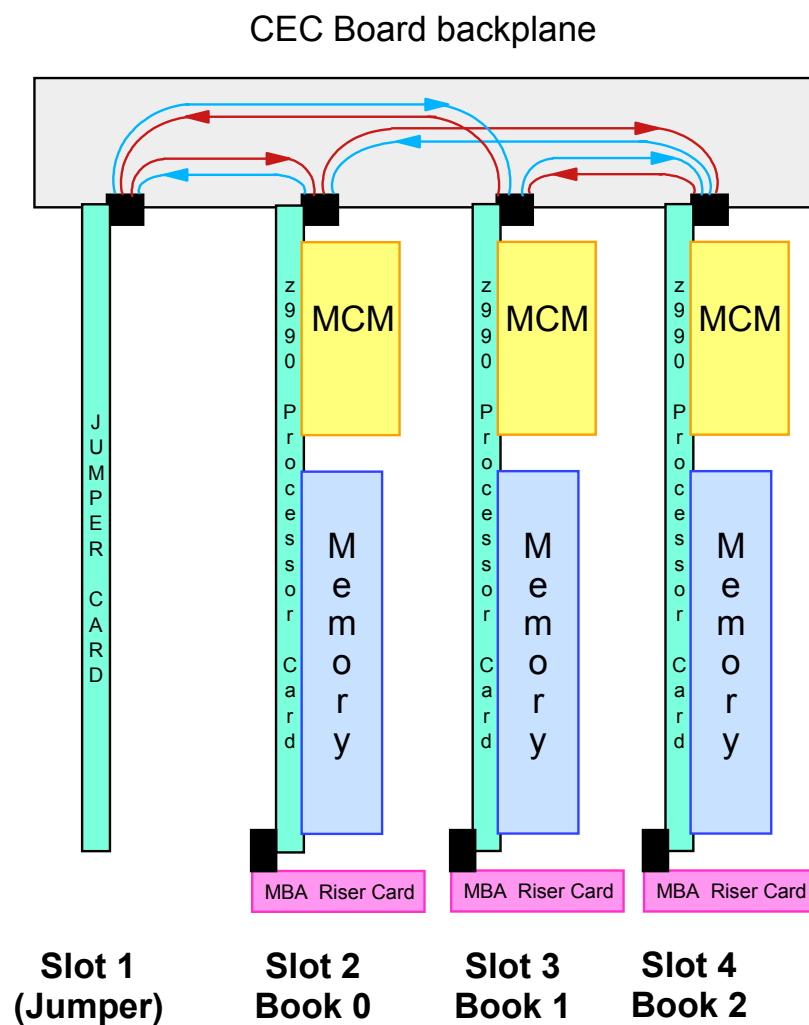


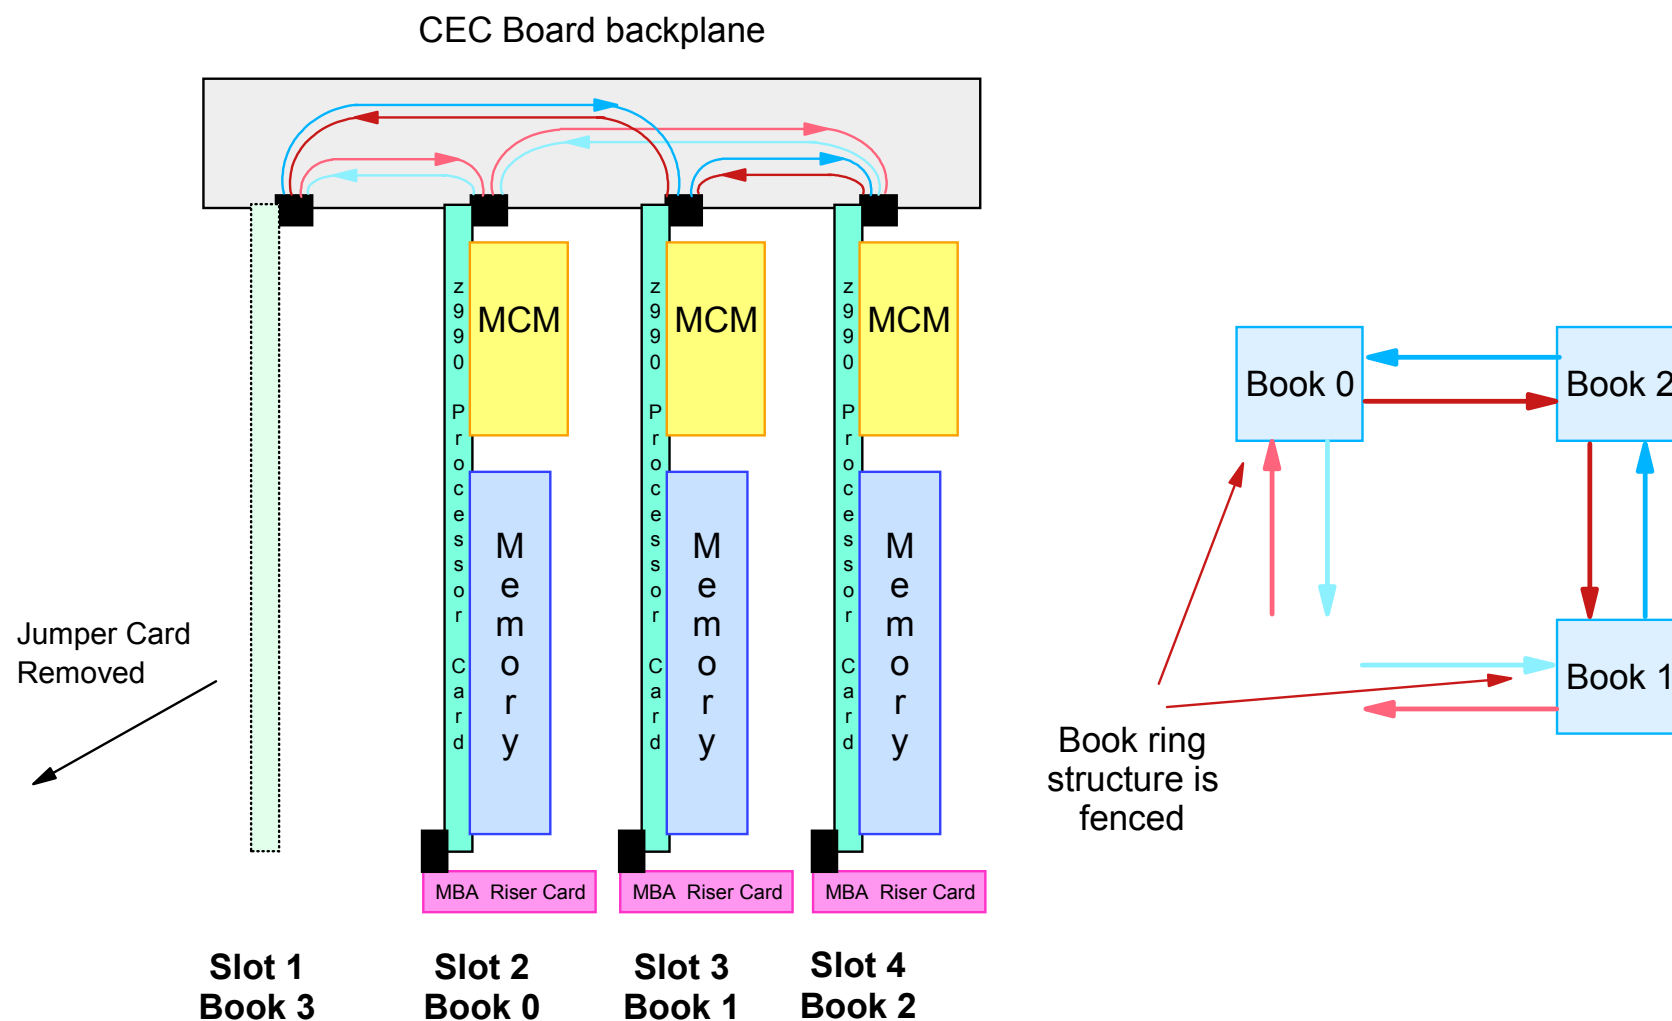
2084 z990 model B16 (2 books)





2084 z990 model C24 (3 books)




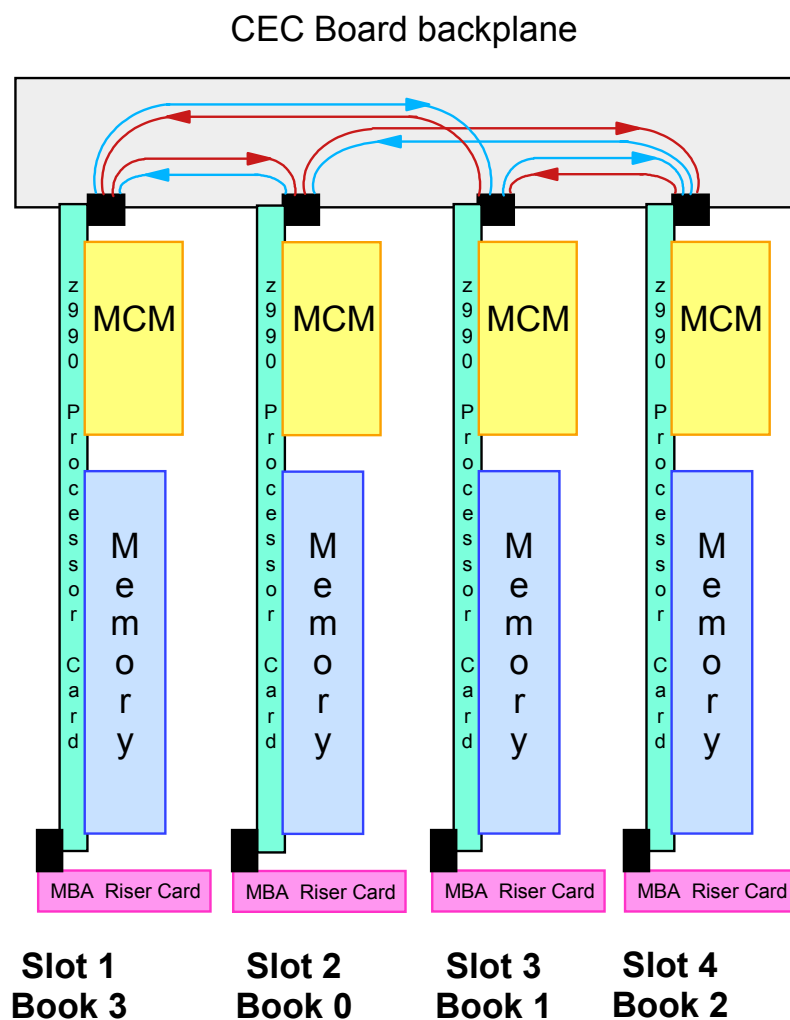


• Open Ring Support

- Needed when performing concurrent upgrade activity (e.g Book)
- Removing a Book Jumper card breaks the links between 2 Books in the system (book 0 and 1 in this example) by activating a logical ring fence with the SCE, the 2 end Books will only have one incoming ring buss and one outgoing ring buss (to book 2)



2084 z990 model D32 (4 books)






Memory Cards	Maximum Memory Size	Memory Cards Per MCM	Memory Granularity (GB)
8 GB	16 GB	2	8 and 16
16 GB	32 GB	2	24 and 32
32 GB	64 GB	2	40, 48, 56 and 64

- Two cards with the same amount of enabled memory are always present in each book
- Total z990 system memory appears as a single coherent memory
- All memory is shared within the system (but dedicated to LPs)
- Cache structure ensures uniform memory access
- LICCC specifications determine how much memory is available for use

Remember: PR/SM™ (LPAR) manages all CPs and memory as a single large SMP. The book location of resources is a secondary issue. A single logical partition (logical processor) can use resources (CPs, memory, STI paths) located in different books.



- Memory is purchased in 8 GB increments
- Memory card sizes - 8, 16, or 32 GB/card = 16, 32 or 64 GB/book
- Offered memory sizes - all multiples of 8 GB:
 - Model A08 - 8GB (GA1) 16GB (GA2) to 64 GB (FC#2601, FC#2608)
 - Model B16 - 8GB (GA1) 16GB (GA2) to 128 GB (FC#2601, FC#2616)
 - Model C24 - 8GB (GA1) 16GB (GA2) to 192 GB (FC#2601, FC#2624) - **GA2**
 - Model D32 - 8GB (GA1) 16GB (GA2) to 256 GB (FC#2601, FC#2632) - **GA2**
- New build memory configuration
 - Use the smallest possible memory cards
 - LIC enable memory in 8 GB increments to balance purchased memory across all installed books
 - Example: B16 with 16 GB
 - ▶ Two 8 GB cards (16 GB) in each book
 - ▶ 8 GB LIC enabled in each book
- MES memory add
 - Concurrent
 - ▶ LIC enable increments on installed cards
 - ▶ LIC enable increments on cards added with a new book - **GA2**
 - Disruptive - Memory card change





- Prerequisites for concurrent add
 - Must have spare memory capacity on currently installed cards (Rows and **boxes** in the table)
 - Or add memory with a new book - GA2
- Concurrent Add to partition
 - Must predefine additional memory to partition as "Reserved Storage"
 - Add to z/OS™ partition using Dynamic Storage Reconfiguration (DSR/2)
 - z/OS comand 'CF STOR(E=1),ON'
- **Notes:**
 1. All combinations NOT shown
 2. Cards shown are for new build only. Cards added by MES may differ.

Memory Card Size and Number of Cards

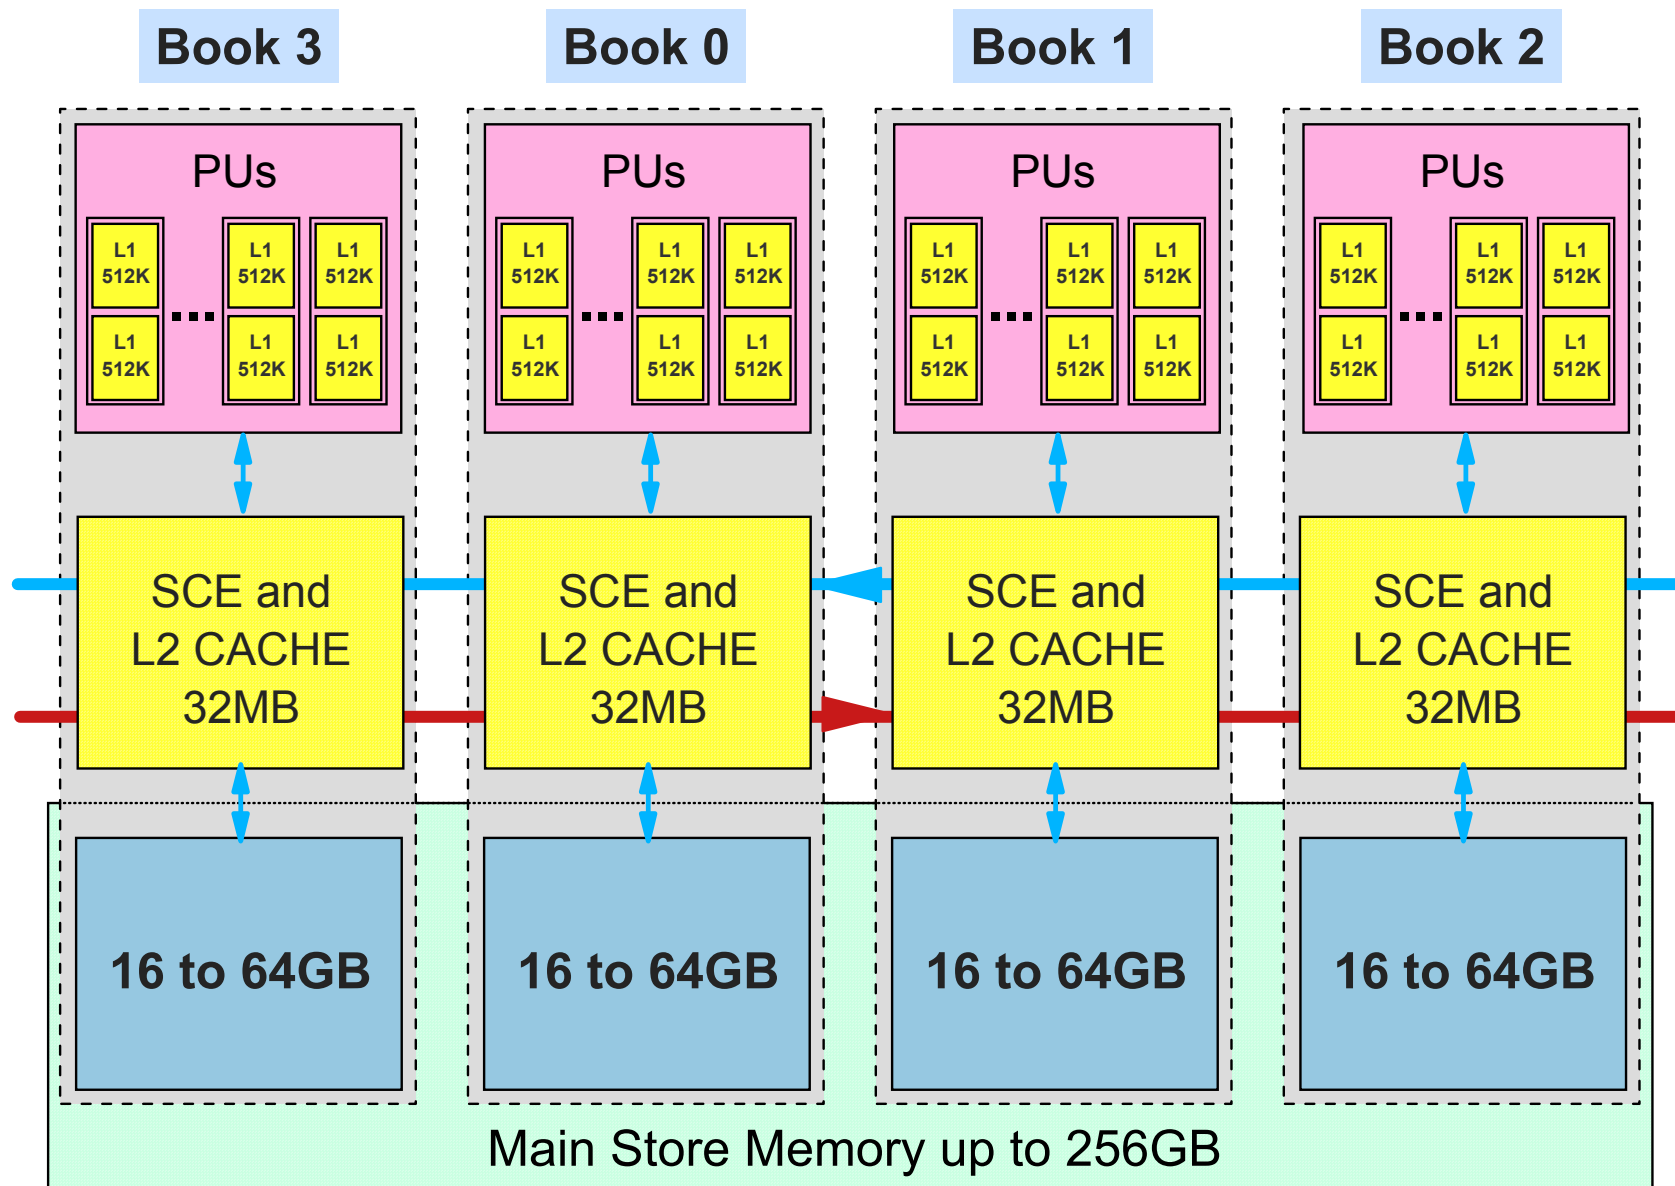
Purchased Memory (GB)	Model A08	Model B16	Model C24 (GA2)	Model D32 (GA2)
8 16	8 GB x 2	8 GB x 4	8 GB x 6	8 GB x 8
24 32	16 GB x 2	8 GB x 4	8 GB x 6	8 GB x 8
40	32 GB x 2	16 GB x 2 8 GB x 2	8 GB x 6	8 GB x 8
48	32 GB x 2	16 GB x 2 8 GB x 2	8 GB x 6	8 GB x 8
56 64	32 GB x 2	16 GB x 4	16 GB x 2 8 GB x 4	8 GB x 8
---	---	---	---	---
128	N/A	32 GB x 4	32 GB x 2 16 GB x 4	16 GB x 8
---	---	---	---	---
256 (GA2)	N/A	N/A	N/A	32 GB x 8

not every orderable memory size is shown in the above chart



- Example configuration: z990 Model B16, with 40 GB of memory, 6 CPs, 1 Unassigned CP, 2 IFLs, and 1 ICF
- Memory Assignment
 - Book 0 - Two 16 GB cards, 24 GB LIC enabled (8 GB available)
 - Book 1 - Two 8 GB cards, 16 GB LIC enabled (0 GB available)
- PU Characterization
 - Book 0: 6 CPs, 2 IFLs, plus 2 SAPs, 2 spare PUs
 - Book 1: 1 ICF, 1 Unassigned CP, 6 uncharacterized PUs, plus 2 SAPs, 2 spare PUs
- But remember:

PR/SM[™] manages all of this as a single large SMP. The book location of resources is a secondary issue. A single logical partition can/will use resources located in multiple books.





● Main Storage

- Main Storage/Memory (sometimes referred to as L3 or storage) is the system's architected storage, and holds data for all installed memory addresses (but maybe not the most recent copy if changed data is still held in L2 cache).
 - ▶ There is only one copy of each address in storage, stored on one of the system's PMA cards (PMA = Physical Main Array).
 - ▶ On the z990 the 256GB of maximum physical address space is sliced up into 4 contiguous blocks where each block is assigned to a Book.
 - ▶ This means on partially populated Book(s) there will be memory holes. Physical address bits 25:26 (highest ordered bits) of the storage increment in the Configuration Table targets the book where the physical memory exists for the system address. Storage within a book is partitioned among the installed memory card(s) by system address bits 53 and 55.
 - ▶ For the 2 PMA cards installed in a Book, address bit 53 selects from the 2 PMA cards of the book and address bit 55 selects from the 2 logical PMA's within a PMA card.
 - ▶ Additional levels of storage include Expanded Storage (sometimes referred to as L4), and external DASD and tape storage devices.



- Memory Coherent Directory (RA Bit)
 - Used to determine whether local memory has been accessed by any remote book
 - Intent is to avoid unnecessary ring traffic
 - ▶ Only launch new request on ring if data may exist in remote book's cache
 - RA (Remote Access) bit is kept in Config Array
 - ▶ Even though part of Config Array, not part of regular Config Array entry
 - ▶ Granularity of RA bit is 16M of memory
 - ▶ Set whenever remote CP fetch which targets local memory is processed on book
 - ▶ Reset when Config Array entry is invalidated
- Multiple Coherence Points
 - IM - Intervention Master
 - ▶ Most recent book to cache in new data or requested book ownership status change is tagged as "IM" for the storage line address
 - ▶ By definition there can only be one book with "IM" status
 - ▶ Sources data (XI Castout) immediately with an intermediate response in reverse direction of incoming request
 - ▶ Can reject all subsequent conflicting requests from other books
 - ▶ Overrides any kind of responses from Memory Master
 - MM - Memory Master
 - ▶ This is the book containing the Memory location of the requested storage line address
 - ▶ Only one book can have the "MM" status
 - ▶ Returns data from memory only when neither responses from both rings have a hit status
 - ▶ Can reject all subsequent conflicting requests from other books
 - ▶ Cannot override any kind of responses from Intervention Master



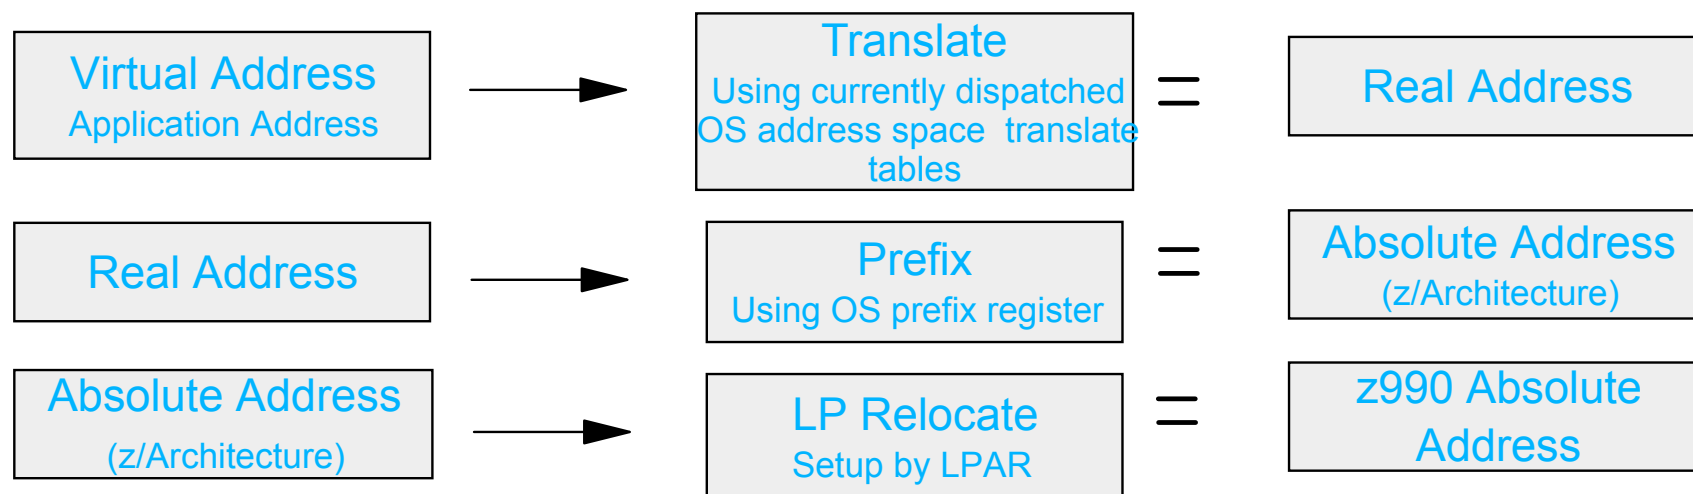
- Remote Storage Controller

- The Remote Storage Controller (RSC) is responsible for all ring communications in the CEC. If any CP or MBA on one book needs to reference or alter data residing in another book, the RSC logic is employed to perform the necessary actions. In addition to data transfers, the RSC logic also maintains complete cache coherency by ensuring data is never owned by processors on more than one book, and that the memory coherent directory reflects the possible shared state of data.
- The major functional units of the RSC are:
 - ▶ RSC Fetch Controller (RFAR)
 - ▶ RSC Store Controller (RSAR)
 - ▶ RSC Millicode Controller (RMAR)
 - ▶ RSC Pipe Priority
 - ▶ RSC Ring Response Priority
 - ▶ RSC Address Interlock




- z990 Absolute Logical Address Generation

- ▶ Assume application is using z/Architecture Virtual Addressing and not Real addressing
- Virtual address (CP) to a Real Address
 - ▶ Translate to a Real Address
 - The operating system sets up the translate controls (CRs) and translate tables
- Real Address (CP) to a z/Architecture Absolute Address
 - ▶ Prefix to an z/Architecture Absolute Address
 - If the Real Address is equal to 0KB - 8 KB (PSA) - forward prefix
 - If the Real Address is equal to the Prefix Register - reverse prefix
 - If the Real Address is not equal to the PSA address or not equal to the prefix register address - then no prefix and the Real Address is an Absolute Address
- z/Architecture absolute address to z990 Absolute Address
 - ▶ Add the 'Relocate' offset (CP PR/SM) to a z990 Absolute Address
 - LPAR sets up the Logical Partition relocate offset at Image Activation time
- z990 Absolute Address to z990 Book / memory address
 - ▶ Use to z990 absolute address to address the z990 Config Array to get z990 Book/memory physical address



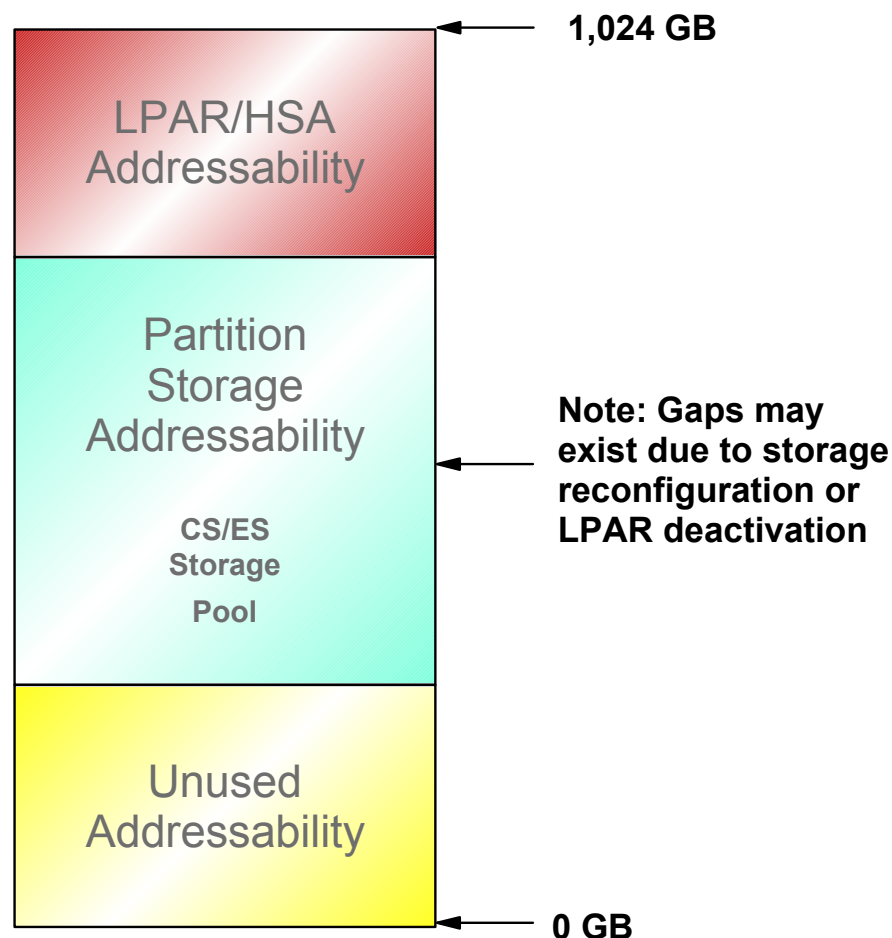
L1 cache hit - if no L1 cache hit go to L2
 Local L2 cache hit - if no L2 cache go to local SCE
 SCE L2 cross interrogate - and if the data not in any L2
 SCE uses z990 absolute address to address its config array to obtain the main storage physical address



Data is held L2 requesting Books L2 cache (data is still in Main Storage)
 If data is later changed - current copy is in both L1 (store thru) and L2 cache (store in)
 Store L2 changed data in Main Storage when it ages
 (Under certain LP dispatch conditions it can be moved to another L2 cache)



- Storage Addressability for LPAR/HSA is allocated top down beginning at (1,024 GB = 1 TB)
- Partition Storage Addressability for OS, applications, and I/O is assigned below HSA at LPAR activation
 - ▶ **OS, Application, I/O operations**
 - ▶ **Origin address is assigned top/down by default but a specific Origin can be requested**
 - ▶ **All Initial and Reserved CS and ES takes addressability at LPAR activation**
- In Book/Memory Physical Storage
 - ▶ **LPAR/HSA starts at book/memory physical address 0 (in book 0)**
 - ▶ **LPAR/HSA size can exceed 2 GB**
 - ▶ **Physical storage assigned to LPARs is above this.**






- z990 2084 Config Array

- The Config Array is a mapping of the z990 Absolute Addresses to a z990 Book/Memory Physical Addresses
- Note: L1 and L2 Caches are effectively z990 Absolute Address based
 - ▶ To access data in one of these Caches (L1 and L2) the effective z990 Absolute Address is used
 - ▶ L3 Storage (main memory) is z990 Book/Memory Physical Address based
 - ▶ Therefore when requested data is not in the CPs L1 cache or any L2 cache it needs to be fetched from L3, there is a need to use the z990 Absolute Address to address the SC Config Array to obtain a Book/Memory Physical Address




● z990 2084 Config Array - Memory Address to Book Mapping






- 64-bit z990 Absolute Address with Logical Partition relocate applied



- Bits 0 --- 23 must be 0 or invalid, maximum is 1 TB = 2^{40}
- Bits 24 --- 37 index the Configuration Array
 - 14 bits = 16 K entries
- Bits 38 --- 63 low order bits
 - 26 bits = 64 MByte increment size
 - 64 MByte times 16 K = 1024 GByte = 1 TByte

- 64-bit Hardware Physical Address



- Bits 0 --- 24 must be 0 or invalid, maximum is 512 GB = 2^{39} (128 GB/Book)
- Bits 25 - 26 = Book 00, 01, 10, 11 (0, 1, 2, 3)
- Bit 27 = 0 (Maximum supported memory per Book = 64 GB, not 128 GB)
- Bits 28 - 37 = Hardware Physical Increment Address from Array
 - 10 bits = 1,024 increments (64 MB each) per Book
- Bits 38 - 63 copied from low order absolute address



- The Config Array has 6 major purposes:
 1. Translates Absolute Physical Addresses to Book/Memory Physical Addresses.
 2. Starting with the z990, identifies which book contains the physical storage that is associated with an Absolute Physical Address. The first two bits of a Physical Address are the book id.
 3. Enables segments of Absolute Physical Addresses to be reassigned (remapped) to different segments of Physical Storage.
 4. Enables Absolute Physical Addresses that are greater than the maximum physical storage size by remapping the Absolute Physical Address to an installed Physical Address. To support LPAR the Config Array allows an Absolute Physical Address range that is at least twice the size of maximum physical storage installed on a machine. i.e. maximum Absolute Physical Address may be at least twice the maximum Physical Address.
 5. Detects invalid addresses. Absolute Physical Addresses that do not have physical storage associated with them are invalid. At least half the Absolute Physical Addresses in the Config Array are always invalid because the Config Array allows Absolute Physical Addresses that are at least twice the maximum Physical Storage that can be installed on a machine. When the maximum Physical Storage is not installed on a machine many Absolute Physical Addresses will be invalid. This implies that if all the storage on a book is not installed then there will be gaps in total physical storage range seen by CP Millicode / LPAR.
 6. Starting with the z990, the Config Array keeps track of which storage segments have been accessed by remote books and thus may have storage resident in remote L2 caches.



- The Config Array has 3 major outputs:
 1. A main storage Physical Address that is generated from a z990 Absolute Address and from an entry in the Config Array. It is used to obtain the book Id and to access L3 storage.
 2. Invalid Address (IVA) Reject signal. It is used to block L2 cache, L3 Storage accesses and send back IVA responses to the CP or MBA.
 3. Remote Access (RA) bits to reduce unnecessary time consuming remote book interrogation the Config Array implements sectorized Remote Access bits.



● Config Array: Facts

1. Config Array supports absolute addresses up to 1 Tera byte. (Absolute Address 24:63.) Therefore maximum Physical Address is 512 Gig. (Physical Address 25:63.)
2. A fixed 64 Meg segment size is associated with each entry in Config Array.
3. Each entry in the Config Array consists of a "Valid" Bit plus 13 physical address bits. Always Physical Address bits 25:37. First two Physical Address bits 25:26 are always the book Id. Physical Address 27:63 allows up to 128 Gig of storage per book. There will be gaps in physical address range because only up to 64 Gig of storage is supported on a book.
 - ▶ Note: Physical Address bit 27 is always 0 due to the 64 Gig limit in supported storage.
 - ▶ Note: Physical Address bits 38:63 are copied from Absolute Address bits 38:63.
4. Entries in the Config Array are addressed/indexed by z990 Absolute Address bits 24:37.
5. The Config Array logically contains 16K of "On-Line" entries and 16K of "Off-Line" entries.
 - ▶ 1 Tera divided by 64 Meg is 16K.
6. For timing and performance reasons a complete copy of the dual (On-Line / Off-Line) Config Arrays is implemented on each of the two SCC pipes per SCC chip. SCC hardware takes care of keeping Config Arrays on both pipes in sync.
7. On a multi book machine there is a complete copy of the dual (On-Line / Off-Line) Config Arrays on each book.
 - ▶ 1 SCC chip per book.
 - ▶ SCC hardware takes care of broadcasting Config Array updates to any set of books based on book masks in Config Array Write commands sent by CP Millicode.
8. The Config Array is protected by ECC with UE/CE reported to CP Millicode via UBus polling.
 - ▶ CEs are corrected on the fly
 - ▶ UEs cause System Checkstops



IBM @server zSeries 990

ITSO - z990 Technical Workshop (06/2003)

Introduction - Processor Complex

LPAR Support

2084-z990

ITSO Poughkeepsie



- LPAR gets optimal performance from the z990 by:
 - Ensuring best storage allocation and processor allocation at image activation time
 - Attempts to allocate all storage for a logical partition from the same Book
 - For dedicated processors (CPs), try to allocate CPs on same book as memory
 - Best dispatching of a shared logical CP to the best physical CP
 - Any CP on last dispatched Book
 - Any CP on the Book where the LP memory resides
 - Any CP




- z990 LPAR Enhancements
 - More than 16 physical engines can be managed - from 1 to 32
 - (SOD - More than 16 logical engines per partition)
 - Multiple logical channel subsystems
 - Many internal changes to PR/SM
 - Limited external change for the customer
 - New extended I/O measurement facility
 - Store Channel Path Status (STCPS) now a "no-op"
 - See also Z990 Profiles and z/OS Displays
 - Awareness of book structure
 - Awareness of CP, L2 cache, and Memory locations
 - Efficient resource allocation to logical partitions at activation
 - Efficient dispatching decisions
 - More than 15 logical partitions - from 1 to 30 may be activated
 - "Logical Partition Identifier" now two hex digits
 - Change to Store CPU ID (STIDP) instruction
 - Large storage / large concurrent memory upgrade
 - Partition virtualization of storage increment size for memory allocation and reconfiguration
 - Very large increase in available central memory addressability



- "Partition Identifier" and "Partition Number">
 - "Partition Identifier" of two hex digits, 1 (or 01) to 3B) instead of one 1 to F
 - ▶ Still specified on the "General" page of the image profile
 - ▶ Still returned by the Store CPUID (STIDP) instruction
 - ▶ Now an 8 bit field, instead of 4 bits
 - ▶ Logical CPU ID no longer returned by STIDP (No room left!)
 - ▶ Logical CPU ID still returned by Store CPU Address (STAP)
 - ▶ Users of old STIDP format require change for compatibility
 - XES for sysplex communication
 - Path Group Identifier (SPID) for I/O (setting up a path group)
 - ICKDSF serialization
 - Independent Software Vendors (ISVs) may have issues (modifications)
 - HCD "Partition Number" (aka IOCP "Resource Number - MIF ID")
 - ▶ Is not unique across all the LPARs between the LCSSs
 - ▶ Is unique for Logical Partitions within the same LCSS
 - ▶ "Partition Number (MIF ID)" range is still x'1 - F' for LCSS 0, and x'1 - F' LCSS 1
 - ▶ Still used to support host MIF channel sharing (e.g. ESCON, FICON etc)
 - Architected limit for MIF channel SHARING is 15 partitions
 - Architected limit for LCSS.MIF channel SPANNING 240 partitions
 - Used to identify the host (zSeries) path source to the target CU.I/O
 - ▶ Problem: Sysplex policy used "Partition Number" to identify a CF partition
 - ▶ Solution: "Partition Identifier" used instead for identification



z990 new STIDP Result Format



Programs that use this must change

- ▶ PGID (OS/390®, VM, VSE, LINUX/390, ICKDSF)
- ▶ XES, RMF™
- ▶ ISV software (should use STSI for the LP ID)

STSI = Store System Information



```
D M=CPU
IEE174I 02.05.56 DISPLAY M 359
PROCESSOR STATUS
ID  CPU
```

```
SERIAL
0C6A3A2084
0C6A3A2084
```

```
CPC ND = 002084.305.IBM.02.0000000026A3A
CPC SI = 2084.305.IBM.02.00000000000026A3A
CPC ID = 00
CPC NAME = A0C
LP NAME = A0C          LP ID =  C
CSS ID  = 0
MIF ID  = C
```

```
+ ONLINE      - OFFLINE      . DOES NOT EXIST      W WLM-MANAGED
N NOT AVAILABLE
```

z/OS 1.4 running on processor 2084 - 26A3A
All logical processors (CPUs) in the same logical partition show the same CPU serial number.
CPU serial number format:

Logical Partition ID (1 byte)
CEC Serial # (2 bytes)
Machine type (2 bytes)

CPC name = SCZP901 (not shown)
LP name= A0C
LCSS = 0
Logical Partition ID = 0C
Image ID (MIF ID) = C

```
D M=CPU
IEE174I 02.03.50 DISPLAY M 887
PROCESSOR STATUS
ID  CPU
```

```
SERIAL
116A3A2084
116A3A2084
```

```
CPC ND = 002084.305.IBM.02.0000000026A3A
CPC SI = 2084.305.IBM.02.00000000000026A3A
CPC ID = 00
CPC NAME = A11
LP NAME = A11          LP ID = 11
CSS ID  = 1
MIF ID  = 1
```

```
+ ONLINE      - OFFLINE      . DOES NOT EXIST      W WLM-MANAGED
N NOT AVAILABLE
```

z/OS 1.4 running on processor 2084 - 26A3A
All logical processors (CPUs) in the same logical partition show the same CPU serial number.
CPU serial number format:

Logical Partition ID (1 byte)
CEC Serial # (2 bytes)
Machine type (2 bytes)

CPC name = SCZP901 (not shown)
LP name= A11
LCSS = 1
Logical Partition ID = 0C
Image ID (MIF ID) = 1



```
D M=CPU
IEE174I 13.40.01 DISPLAY M 219
PROCESSOR STATUS
ID  CPU
0    +
1    +
2    -
```

SERIAL
156A3A2084
156A3A2084

```
CPC ND = 002084.305.IBM.02.0000000026A3A
CPC SI = 2084.305.IBM.02.00000000000026A3A
CPC ID = 00
CPC NAME = A13
LP NAME = A13
CSS ID = 1
MIF ID = 3
```

LP ID = 15

```
+ ONLINE      - OFFLINE      . DOES NOT EXIST      W WLM-MANAGED
N NOT AVAILABLE
```

z/OS 1.4 running on processor 2084 - 26A3A
All logical processors (CPUs) in the same logical partition show the same CPU serial number.
CPU serial number format:

Logical Partition ID (1 byte)
CEC Serial # (2 bytes)
Machine type (2 bytes)

CPC name = SCZP901 (not shown)
LP name= A0C
LCSS = 0
Logical Partition ID = 0C
Image ID (MIF ID) = C

```
D M=CPU
IEE174I 22.09.00 DISPLAY M 071
PROCESSOR STATUS
ID  CPU
0    +
1    +
```

SERIAL
020ECB2064
120ECB2064

```
CPC ND = 002064.1C7.IBM.02.0000000010ECB
CPC SI = 2064.1C7.IBM.02.00000000000010ECB
CPC ID = 00
```

```
+ ONLINE      - OFFLINE      . DOES NOT EXIST      W
```

```
CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIP
```

z/OS 1.2 running on processor 2064 - 10ECB
All logical processors (CPUs) in the same logical partition show a different CPU serial number.
CPU serial number format:

Logical Processor address (4 bits)
Logical Partition ID (4 bits) (from Image Profile)
CEC Serial # (2 bytes)
Machine type (2 bytes)



Customize Activation Profiles : SCZP901

Description A11 image profile

Partition identifier 11

Mode
 ESA/390
 ESA/390 TPF
 Coupling facility
 LINUX Only

Clock type assignment
☒ Standard time of day
☐ Logical partition sysplex timer offset

General Processor Security Storage Options Load PCI Crypto

SCZP901:A03
 SCZP901:A04
 SCZP901:A05
 SCZP901:A06
 SCZP901:A07
 SCZP901:A08
 SCZP901:A09
 SCZP901:A0A
 SCZP901:A0B
 SCZP901:A0C
 SCZP901:A11
 SCZP901:A12
 SCZP901:A13
 SCZP901:A14
 SCZP901:A15

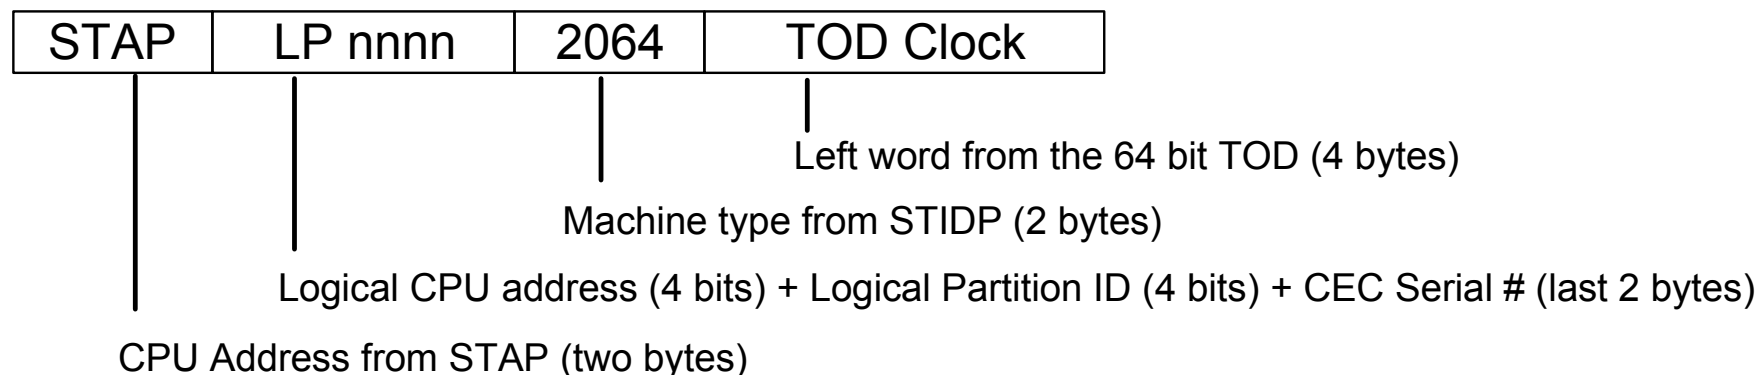
The Logical partition 'Partition Identifier' is a 1 or 2 digit unique hexadecimal value from 0 to 3F
 Some logical partition software operating system / functional users may not support any Logical Partition ID from the full range of partition Identifiers

ESA/390 mode can be used for OS/390 z/OS - z/VM with the required software support for the z990

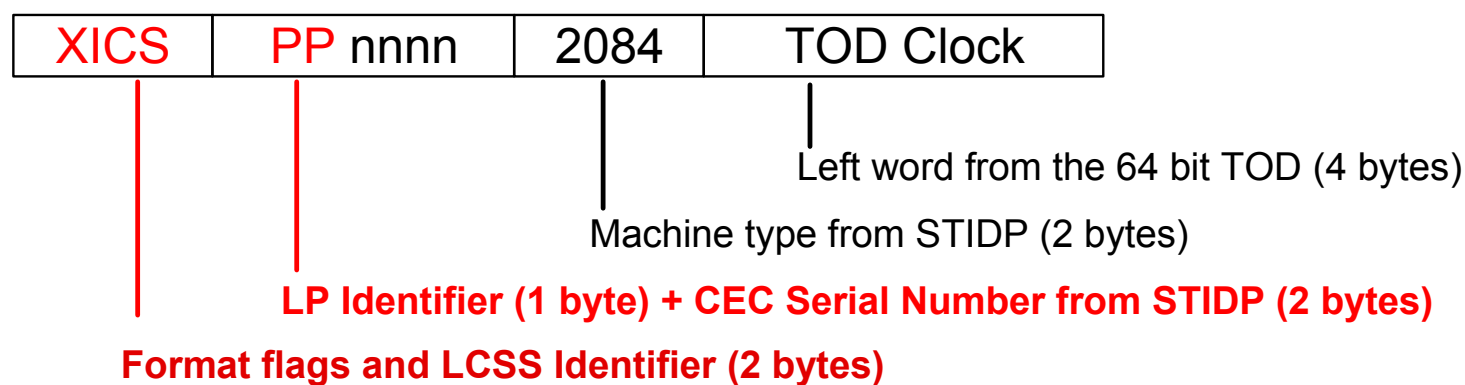


z990 new Path Group Identifier (PGID)

The z900 format and the z990 new format



STAP = Store CPU Address (Store Address Processor)


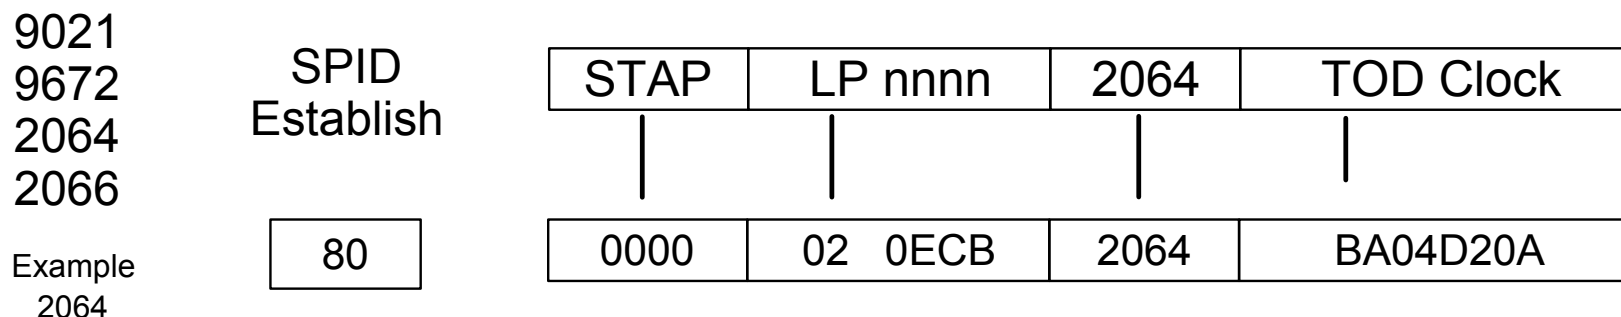


XICS = z/OS created field



z990 new Path Group Identifier (PGID)

The z900 format and the z990 new format



SPID value must always be unique for:

Each IPL on the same machine/image (TOD)

Each IPL for different images on the same machine -same 2064 CSS (LP) or same 2084 LCSS (PP)

Each IPL for different images in different machines (serial number)

Each IPL for different images on the same machine but different 2084 LCSS (PP)

Each IPL for different images on the same machine -but different 9021 sides (logical CP #)

Each IPL for different 9021 basic mode sides on the same machine-(Physical CP #)

The new format field is required by z/OS MIH recovery when a device is reserved by another system

Needs to know the SPID format to support the following message:

IOS431I DEVICE devn RESERVED TO CPU=serialmodn,LPAR ID=gg, SYSTEM=...



- "Partition Identifier" and "Partition Number"
 - HCD "Partition Number" (aka IOCP "Resource Number - MIF ID")
 - ▶ Is not unique across all the LPARs between the LCSSs
 - ▶ Is unique for Logical Partitions within the same LCSS
 - ▶ "Partition Number (MIF ID)" range is still x'1 - F' for LCSS 0, and x'1 - F' LCSS 1
 - ▶ Still used to support host MIF channel sharing (e.g. ESCON, FICON etc)
 - Architected limit for MIF channel **SHARING** is 15 partitions
 - Architected limit for LCSS.MIF channel **SPANNING** 240 partitions
 - **Used to identify the host (zSeries) path source to the target CU.I/O**
 - ▶ **Problem: Sysplex policy used "Partition Number" to identify a CF partition**
 - ▶ **Solution: "Partition Identifier" used instead for identification**



z990 IOCP Statements

```

ID      MSG1='IODF22 checking Resource statement',          *
        MSG2='SYS6.IODF22 - 2003-10-09 11:01',SYSTEM=(2084,1), *
        TOK=('SCZP901',00800000A6A3A2084110126700103282F000000000, *
        00000000,'03-10-09','11:01:26','SYS6','IODF22')
RESOURCE PARTITION=((CSS(0),(A0A,A),(A0B,B),(A0C,C),(A0D,
0E,E),(A0F,F),(A01,1),(A02,2),(A03,3),(A04,4),(A05,5),(A*
06,6),(A07,7),(A08,8),(A09,9)),(CSS(1),(A1A,A),(A1B,B),( *
A1C,C),(A1D,D),(A1E,E),(A1F,F),(A11,1),(A12,2),(A13,3),( *
A14,4),(A15,5),(A16,6),(A17,7),(A18,8),(A19,9))),          *
MAXDEV=((CSS(0),64512),(CSS(1),64512))
CHPID  PATH=(CSS(0),00),SHARED,                              *
        PARTITION=((A0A,A0B,A0C,A01,A02,A03,A04,A05,A06,A07,A08, *
A09),(=)),PCHID=120,TYPE=OSD

```

z990 MIF ID specify



- S/390 9672 , z900, and z800 memory granularity
 - Increment for memory assignment and reconfiguration is global
 - Increment size depends on total physically installed system memory
 - Increment size change is disruptive - limits concurrent memory add
- z990 memory granularity
 - Physical increment size fixed at 64 MB
 - Expanded memory granularity always 64 MB
 - Central memory granularity is virtualized for each LP
 - LP central memory increment is determined according to the size of the larger of the two central memory elements defined in the activation profile
 - Initial central memory or Reserved central memory
 - **Review MVS™ RSU parameter.** Large z990 increment size may result in too much memory being reserved for reconfiguration after migration unless the new RSU options introduced in OS/390 2.10 are used.

	Large Element Size	Granularity
	64 MB to 32 GB	64 MB
	>32 GB to 64 GB	128 MB
	>64 GB to 128 GB	256 MB
	>128 GB to 256 GB	512 MB

GA1 Limit = 128 GB

➔

GA2 Limit = 256 GB

➔



z990 LPAR Memory Granularity

Customize Activation Profiles : SCZP901

Central storage
Amount (in megabytes)

Initial
Reserved

Storage origin

☒ Determined by the system
☐ Determined by the user

Origin

Expanded storage
Amount (in megabytes)

Initial
Reserved

Storage origin

☒ Determined by the system
☐ Determined by the user

Origin

Central storage amount
You may specify up to 128 GB of central storage.
Determine what your OS will support

Storage origin (Central storage)
It is recommended that you use the
'Determined by the system' option

Expanded storage option
Some OSs do not support Expanded storage
One example is z/OS (64-bit) running on a z990
Before you use this option, determine for your OS
use if it supports expanded storage

Storage origin (Expanded storage)
If required, it is recommended that you use the
'Determined by the system' option

SCZP901:A03
SCZP901:A04
SCZP901:A05
SCZP901:A06
SCZP901:A07
SCZP901:A08
SCZP901:A09
SCZP901:A0A
SCZP901:A0B
SCZP901:A0C
SCZP901:A11
SCZP901:A12
SCZP901:A13
SCZP901:A14
SCZP901:A15



Storage Information

z990LPAR Memory Granularity

Input/output configuration data set (IOCDs): A3 IODF20

Available storage: 22976 MB

Central Storage Allocation

Name	Origin	Initial	Current	Maximum	Gap	Expanded Storage Element
A0D	1045120	1024	1024	1024	0	
A0E	1044096	1024	1024	1024	0	
A0F	1043072	1024	1024	1024	0	
A1E	1041024	2048	2048	2048	0	
A1F	1040000	1024	1024	1024	0	

Expanded Storage Allocation

Name	Origin	Initial	Current	Maximum	Gap
A0B	1046528	2048	2048	2048	0
A14		0	0	0	
A13		0	0	0	
A12		0	0	0	
A11		0	0	0	

This storage mapping information was from a 1 Book system

There was 64GB of memory installed.

System origin assigned storage (image profile option) is assigned to a logical partition (image) in POR image activation order (or later when an image is manually activated) starting downwards from the highest z990 implemented storage address (1TB) (that is available and contiguous)



IBM @server zSeries 990

ITSO - z990 Technical Workshop (06/2003)

Introduction - Processor Complex

I/O Support

2084-z990

ITSO Poughkeepsie



zSeries 990



Front View



- Channel cards are installed in I/O cages
- There are 28 Channel card I/O slots per I/O cage
- Multiple channel ports per channel card depending on the channel type
- Channel cards are installed in I/O slots in I/O cages, up to 4 channels card in a group of I/O slots are part of the same I/O domains, and supported by the same eSTI connection and eSTI-M card
- A channel card (with all the channels on the channel card) in an I/O domain is supported (owned) by just one SAP
 - Other SAPs will not communicate to channels on this channel card
- Other SAPs can support the other channel cards in the same I/O domain
 - In his case up to 4 SAPs can use the same STI to communicate to their own channel cards



zSeries 990 CPC Logical Structure - CP and Channel Access

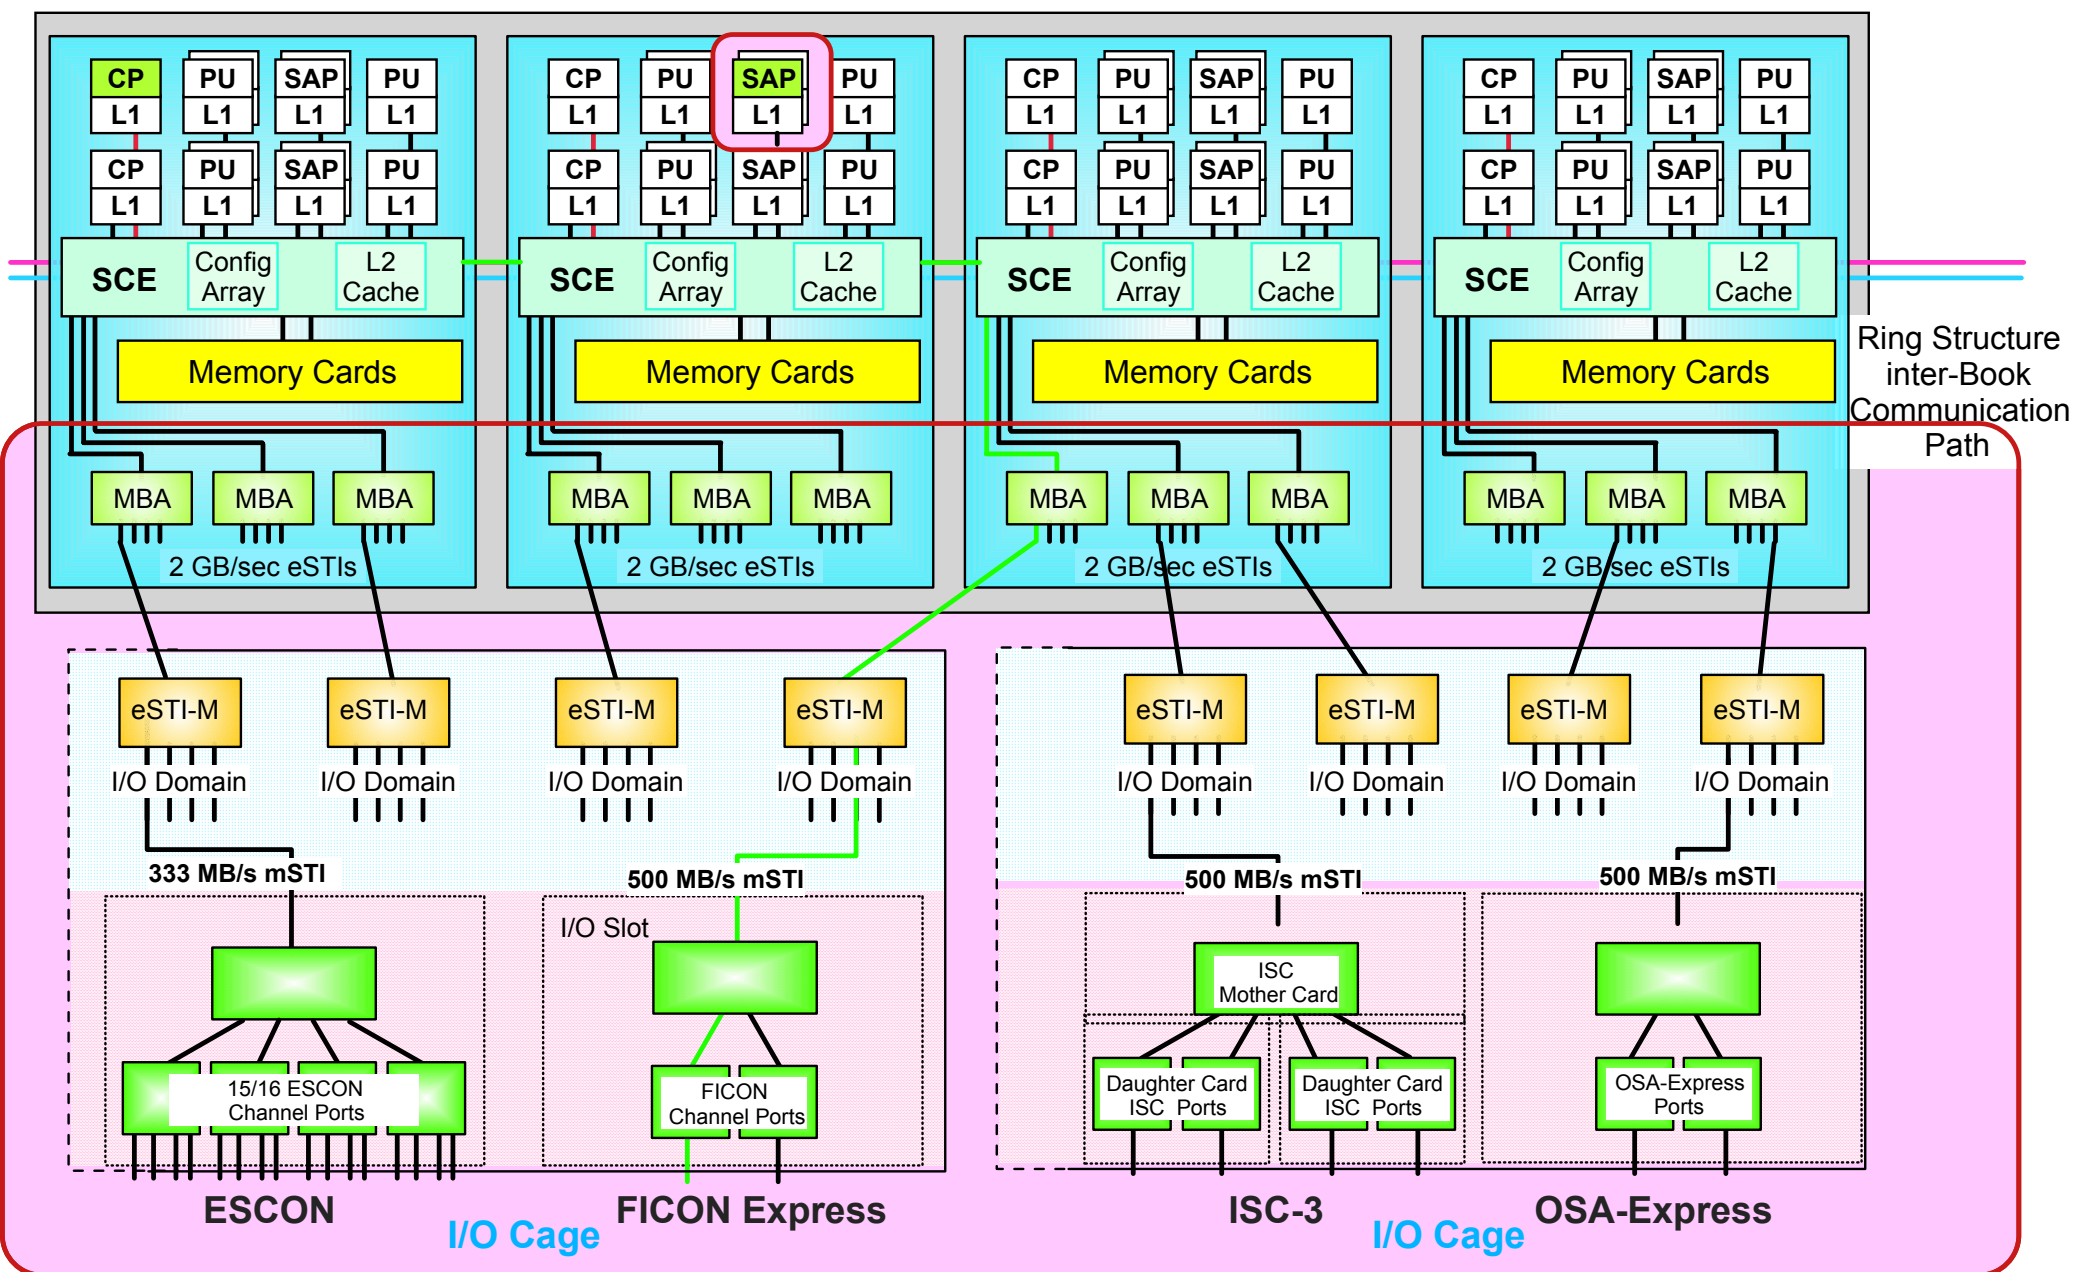
Book 3

Book 0

CEC Cage

Book 1

Book 2





- A 2084 model A08 has 12 eSTI connections and can therefore support up to 12 I/O domains
- Channel cards that can be installed in and supported by an I/O domain are:
 - ESCON - up to 15 usable channels per ESCON channel card
 - FICON-Express - 2 channels per FICON channel card
 - OSA-Express - 2 channels (ports) per OSA-Express card
 - ISC-3 - up to 4 ISC-3 channels per mother (1) / daughter (2) per I/O slot
 - PCICA - 2 Crypto Coprocessors per PCICA card (no CHPID assigned)
 - PCIXCC - 1 Crypto Coprocessors per PCICA card (no CHPID assigned)
- Other channel types that can be in I/O domain card slots but not support by the I/O *domain* eSTI connection are:
 - ICB-2 extender - uses one eSTI input connection directly to the extender card
 - Up to 2 ICB-2 channels per extender card
 - ICB-3 extender - uses one eSTI input connection directly to the extender card
 - Up to 2 ICB-3 channels per extender card
- Channel type that is not installed in an I/O cage
 - ICB-4 - direct eSTI connect from one z990 CEC cage to another z990 CEC cage
- Defined channel types that use no physical location are:
 - IC - Internal coupling channel (each ICP define channel requires a CHPID number)
 - Hipersocket (each IQD define channel requires a CHPID number)




Model A08

12 STIs

24 GBytes

Up to 48 channel
feature cards




24 GBytes for I/O



24 Channel
feature cards



24 Channel
feature cards



6 Domains used



in each of 2 I/O cage. Some configs
may need 3 cages Supported by
the 12 STIs available for I/O
Only 35 ESCON channel cards are
required to support 512 ESCON
channels

Model B16




24 STIs

48 GBytes

Up to 84 channel
feature cards



42 GBytes for I/O
if desired








All 21 Domains can be used
7 in each I/O cage
Supported by the 24 STIs
available for I/O

Model C24

36 STIs

72 GBytes

Up to 84 channel
feature cards







Potential available resources for ESCON®, FICON™ Express,
ISC-3, OSA-Express as well as ICB, PCICA, and PCIXCC.

Model D32

48 STIs

96 GBytes


Up to 84 channel
feature cards



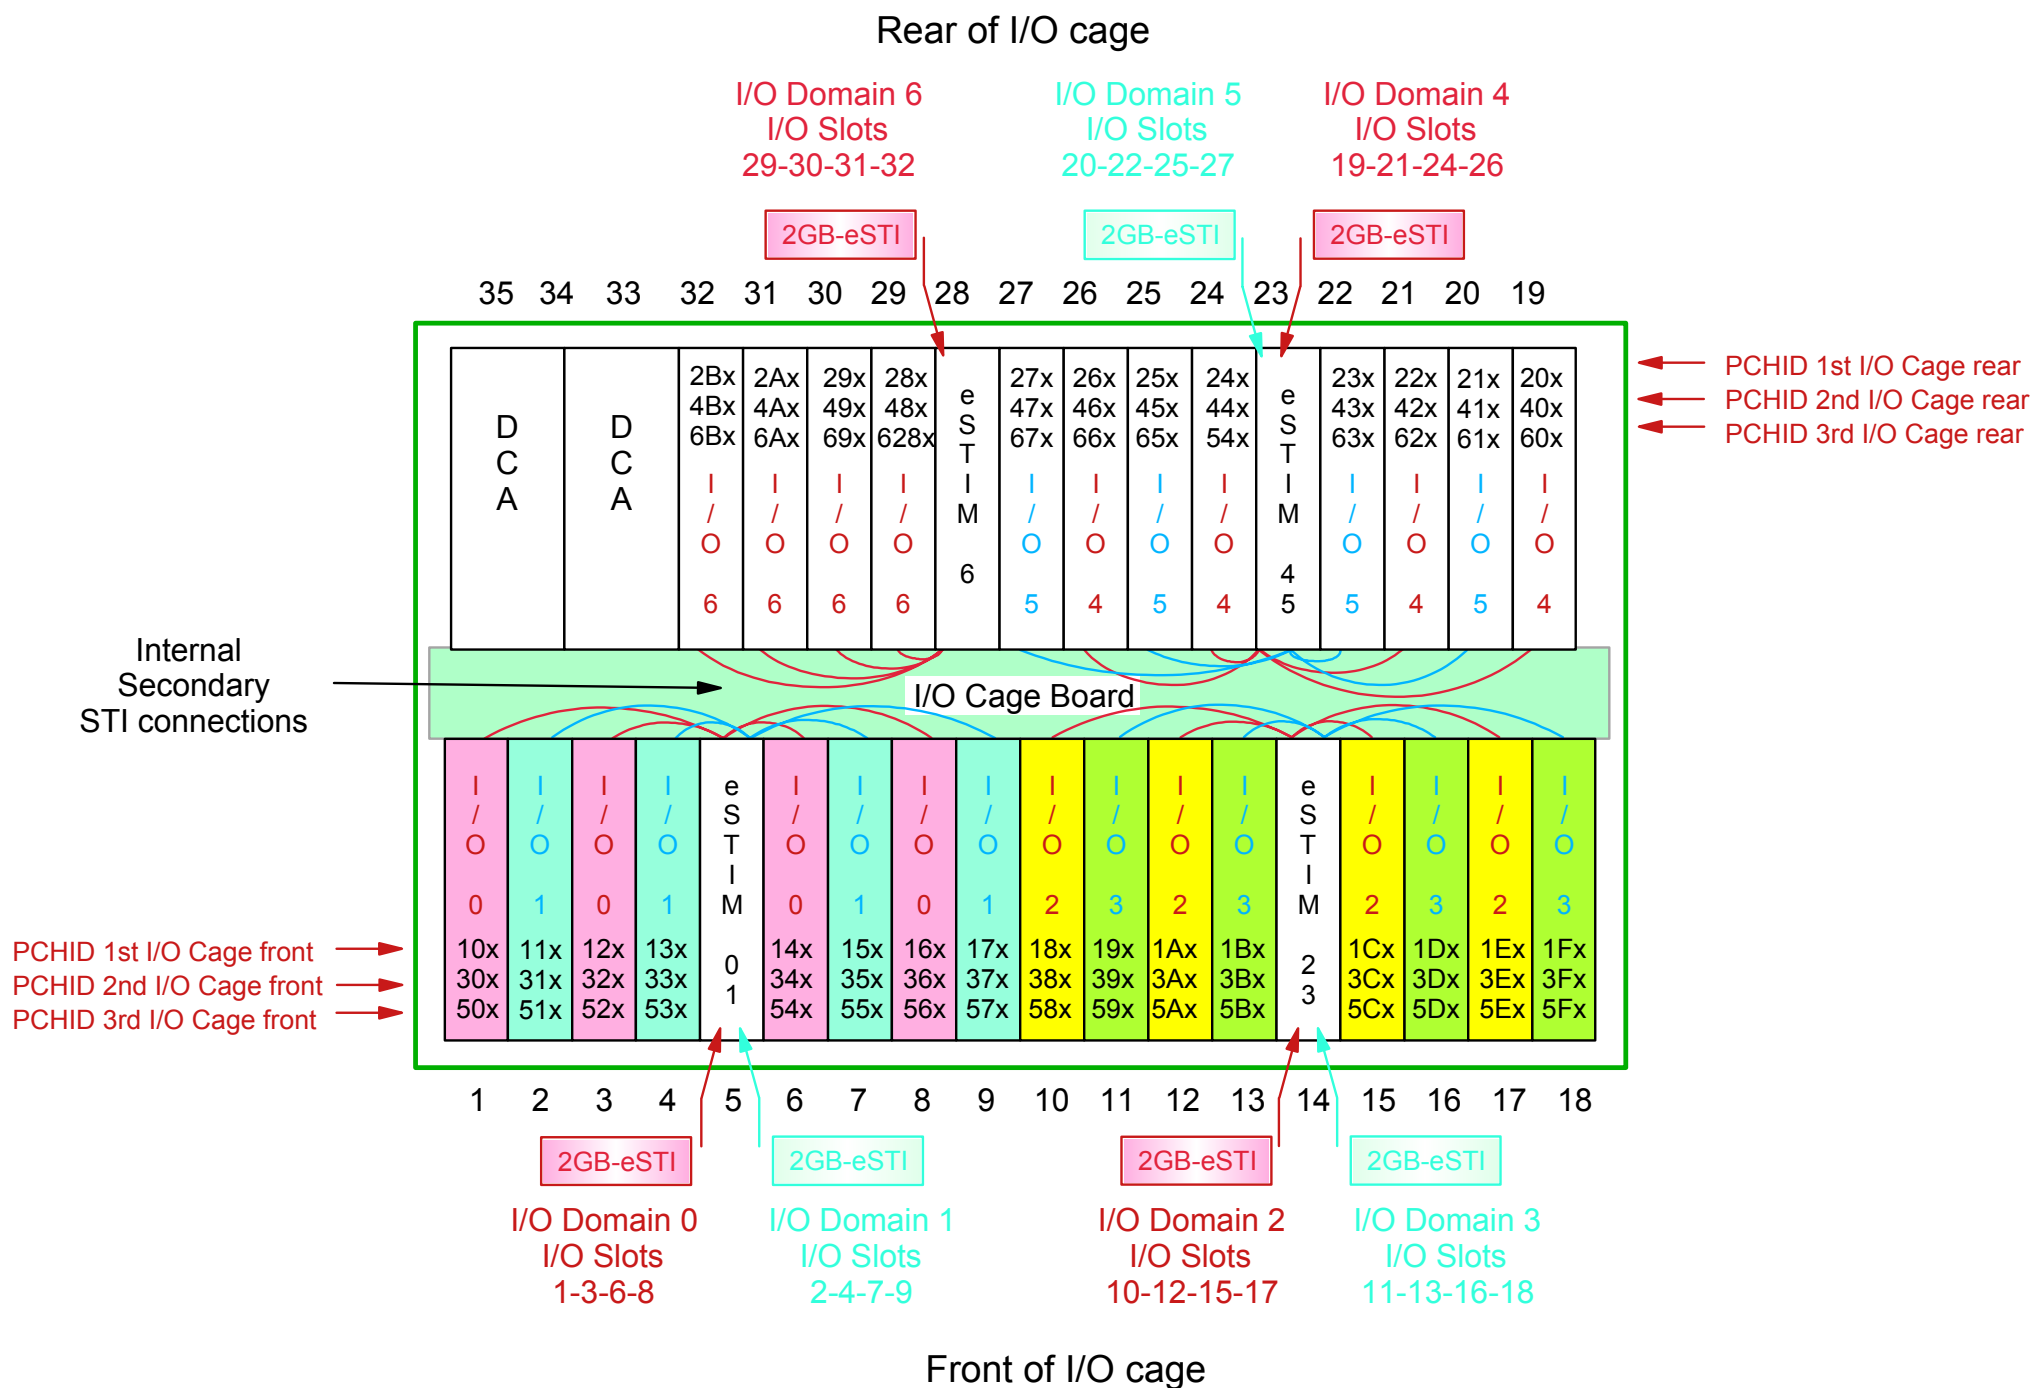
28 Channel
feature cards



28 Channel
feature cards



28 Channel
feature cards



FCH023 Slots 1 - 18



Initial Hardware

Product	Description	Qty
2084-B16	zSeries 990	1
0112	Cargo Cage Airflow Cd	32
0152	Cosmos CEC	1
0225	STI AFB-C 1/2 Airflow	10
0228	ICB Cable (T-REX to T-REX)	4
0322	eSTI-M Card	6
0716	CP	4
1008	8GB Memory Cards (For record purposes only)	4
1216	16-Way Processor	1
1364	OSA-Express GbE LX	2
2319	FICON Express LX	20
2323	16 Port ESCON Card (15 useable per card)	2
2324	ESCON Channel Port (4 chpid groups enabled)	5
2604	32GB LICCC Enabled Memory	1
3062	SE with Token Ring and Ethernet	2
3393	ICB-4	4
4404	4 CPs (For record purposes only)	1
6154	ETR 1 Port	2
7722	B16, 2 I/O Cage, Frame	1
8993	250V 60A Line Cord (US/Can/Japan)	1
9964	Site Tool Kit	1



CHPIDSTART

30214401

PCHID REPORT

Oct 29, 2003

Machine: 2084-B16 NEW1

Book/Jack/MBA	Cage	Slot	F/C	PCHID/Ports	Comment
0/J.07/1	A19B	06	3393	017/J07	
0/J.11/2	A19B	06	3393	01B/J11	
1/J.07/1	A19B	10	3393	027/J07	
1/J.11/2	A19B	10	3393	02B/J11	
0/J.00/0	A01B	01	1364	100/J00 101/J01	
1/J.00/0	A01B	02	2319	110/J00 111/J01	
0/J.00/0	A01B	03	2319	120/J00 121/J01	
1/J.00/0	A01B	04	2319	130/J00 131/J01	
0/J.00/0	A01B	06	2319	140/J00 141/J01	
1/J.00/0	A01B	07	2319	150/J00 151/J01	
0/J.00/0	A01B	08	2319	160/J00 161/J01	
1/J.00/0	A01B	09	2319	170/J00 171/J01	
0/J.04/1	A01B	10	2319	180/J00 181/J01	
0/J.04/1	A01B	12	2319	1A0/J00 1A1/J01	
0/J.04/1	A01B	15	2319	1C0/J00 1C1/J01	
0/J.04/1	A01B	17	2323	1E0/J00 1E1/J01 1E2/J02 1E3/J03 1E4/J04 1E5/J05 1E6/J06 1E7/J07 1E8/J08 1E9/J09	
0/J.08/2	Z01B	09	2319	370/J00 371/J01	
1/J.04/1	Z01B	10	2319	380/J00 381/J01	
1/J.04/1	Z01B	12	2319	3A0/J00 3A1/J01	
1/J.04/1	Z01B	15	2319	3C0/J00 3C1/J01	
1/J.04/1	Z01B	17	2323	3E0/J00 3E1/J01 3E2/J02 3E3/J03 3E4/J04 3E5/J05 3E6/J06 3E7/J07 3E8/J08 3E9/J09	

Legend:

A19B Top of A frame
A01B Bottom of A frame
Z01B Bottom of Z frame
1364 OSA Express GbE LX
2319 FICON Express LX
2323 ESCON Channel 16 Ports
3393 ICB 4 Function Enablement



- z990 CP to SAP to MBA/STI to Channel Card / Channel to Channel Interface

1. CP

- ▶ SSCH instruction
 - GPR 1 - Subchannel number (for the required device)
 - ORB (Operations Request Block) has the required information for the SSCH
 - Transfer contents of the ORB to the subchannel (in HSA)
 - Subchannel also has 'next-SAP-to-use' indicator
 - Inform the SAP that there is I/O work to do

2. SAP

- ▶ Fetch subchannel
- ▶ Determine which channel of defined and available channels to use
- ▶ Inform channel via MBA and STI there is work to do (Example: FICON channel)

3. Channel (Example: FICON channel)

- ▶ Priorities work
- ▶ Fetch Subchannel
- ▶ Fetch 1st CCW and data (for an I/O write operation)
- ▶ Fetch 2nd CCW and data (for an I/O write operation)
- ▶ Initiate I/O operation to the CU/device



zSeries 990 CPC Logical Structure - CP and Channel Access

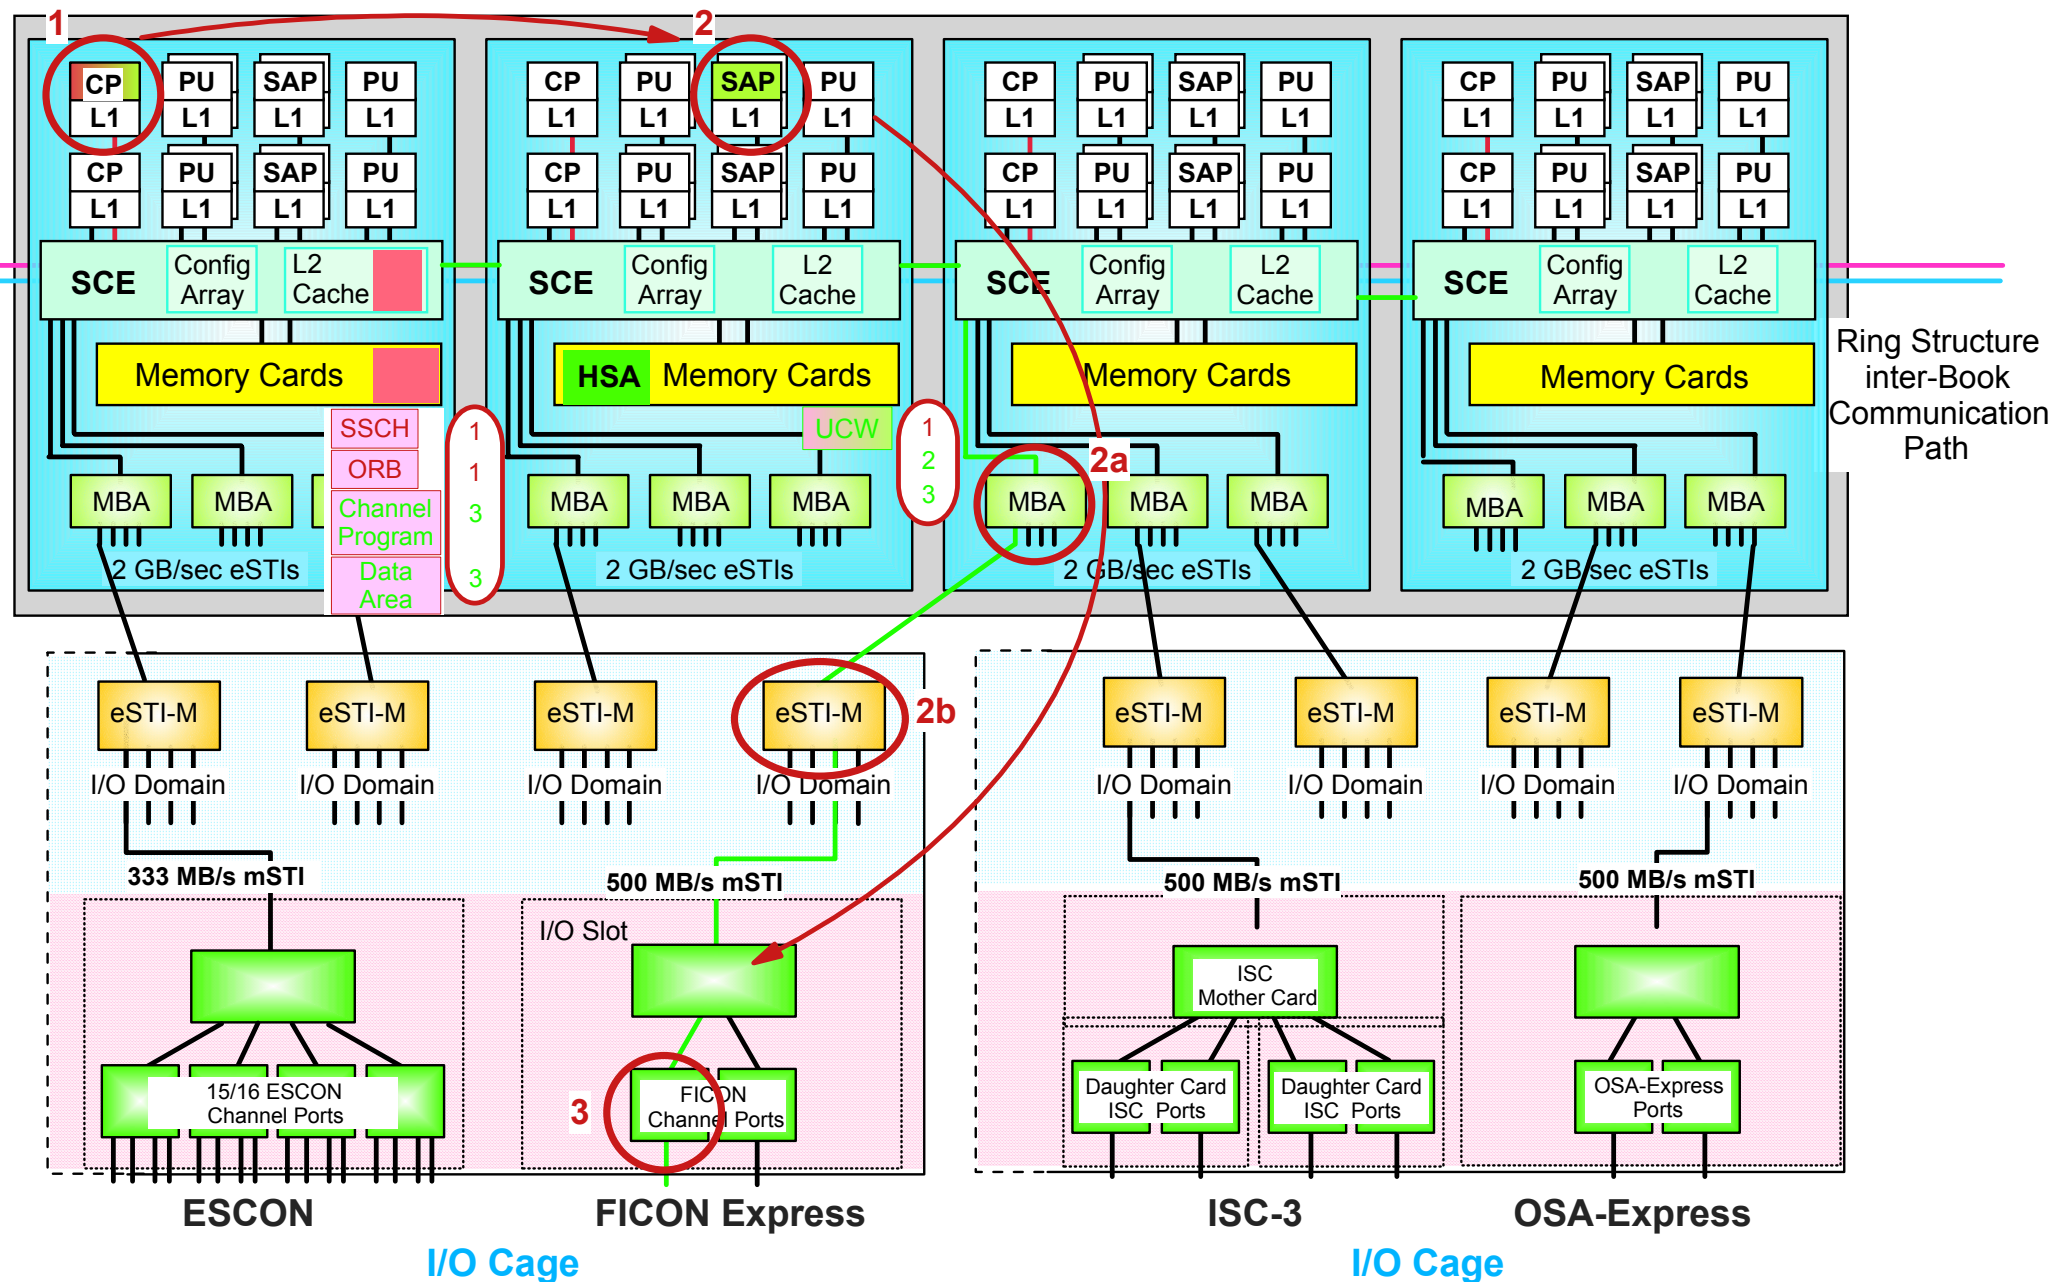
Book 3

Book 0

CEC Cage

Book 1

Book 2






- z990 Processor Book connectivity

- Inter-Book connectivity


- Inter-Book ring communication
 - Supported by the CEC cage board (and as required jumper Books)
 - Used for cross memory communication
 - Memory system controller to memory system controller
 - Used PU to PU communication
 - CP to CP
 - CP to SAP, and SAP to CP
 - SAP to SAP

- Book to Channel connectivity

- Support by 12 STI connections per installed Book
 - STI connections are used for
 - Book (SAP) to I/O domain communication
 - Book to ICB-2 and ICB-3 communication
 - z990 ICB-4 to z990 ICB-4 communication



- STI connections to I/O cages and ICB links are driven from the Memory Bus Adapters (MBAs) that are located on a separate MBS riser card in the book
- Each book has three MBAs each driving four STIs, resulting in 12 STI per book



MBA Riser card



- Processor features

CEC cage, Books, PUs, Memory, STIs



Processor features	Quantity	2084 Hardware Models
Books	1 (Book 0) 2 (Book 0, 1) 3 (Book 0, 1, 2) 4 (Book 0, 1, 2, 3)	A08 B16 C24 D32
Characterized PUs (1 to 32) Up to 8 per book Not included are 2 SAPs - 2 Spares / Book	1 to 8 1 to 16 1 to 24 1 to 32	A08 B16 C24 D32
Memory	8 GB to 64 GB 8 GB to 128 GB 8 GB to 192 GB 8 GB to 256 GB	A08 B16 C24 D32
Book STIs 12 to 48 12 per book (CEC to I/O connection)	12 24 36 48	A08 B16 C24 D32



- Power, heat, cooling requirement depends on installed features



I/O Feature Type	Feature Cards	Maximum Channels Ports Connections by type
ESCON	28 cards (I/O slots)	420 channels
FICON Express	20 cards (I/O slots)	40 channels
OSA-Express	20 cards (I/O slots)	40 ports (channels)
ISC-3	8 ISC-M I/O slots 16 ISC-D cards	32 ports (channels)
ICB-2 (STI-2)	4 cards	8 connections
ICB-3 (STI-3)	8 cards	16 connections
PCIXCC	4 cards (I/O slots)	4 coprocessor (PCHID)
PCICA	2 cards (I/O slots)	4 accelerators (PCHID)

Note:

Number of I/O slots in one I/O cage is 28
 Maximum combined FICON-Express, OSA-Express,
 PCICA/PCIXCC features is 20 for a single I/O cage configuration.
 Only 2 PCICA cards are supported per I/O cage



- Power, heat, cooling requirement depends on installed features



I/O Feature Type	Feature Cards	Maximum Channels Ports Connections by type
ESCON	35 cards (I/O slots)	512 channels
FICON Express	40 cards (I/O slots)	80 channels
OSA-Express	24 cards (I/O slots)	48 ports (channels)
ISC-3	8 ISC-M I/O slots 16 ISC-D cards	32 ports (channels)
ICB-2 (STI-2)	4 cards	8 connections
ICB-3	8 cards	16 connections
PCIXCC	4 cards (I/O slots)	4 coprocessor (PCHID)
PCICA	4 cards (I/O slots)	8 accelerators (PCHID)

Note:

Number of I/O slots in two I/O cages is 56

Maximum combined FICON-Express, OSA-Express, PCICA/PCIXCC features is 40 for a 2 I/O cage configuration.

Only 2 PCICA cards are supported per I/O cage



- Power, heat, cooling requirement depends on installed features




I/O Feature Type	Feature Cards	Maximum Channels Ports Connections by type
ESCON	35 cards (I/O slots)	512 channels
FICON Express	60 cards (I/O slots)	120 channels
OSA-Express	24 cards (I/O slots)	48 ports (channels)
ISC-3	8 ISC-M I/O slots 16 ISC-D cards	32 ports (channels)
ICB-2 (STI-2)	4 cards	8 connections (channels)
ICB-3 (STI-3)	8 cards	16 connections (channels)
PCIXCC	4 cards (I/O slots)	4 coprocessor (PCHID)
PCICA	6 cards (I/O slots)	12 accelerators (PCHID)

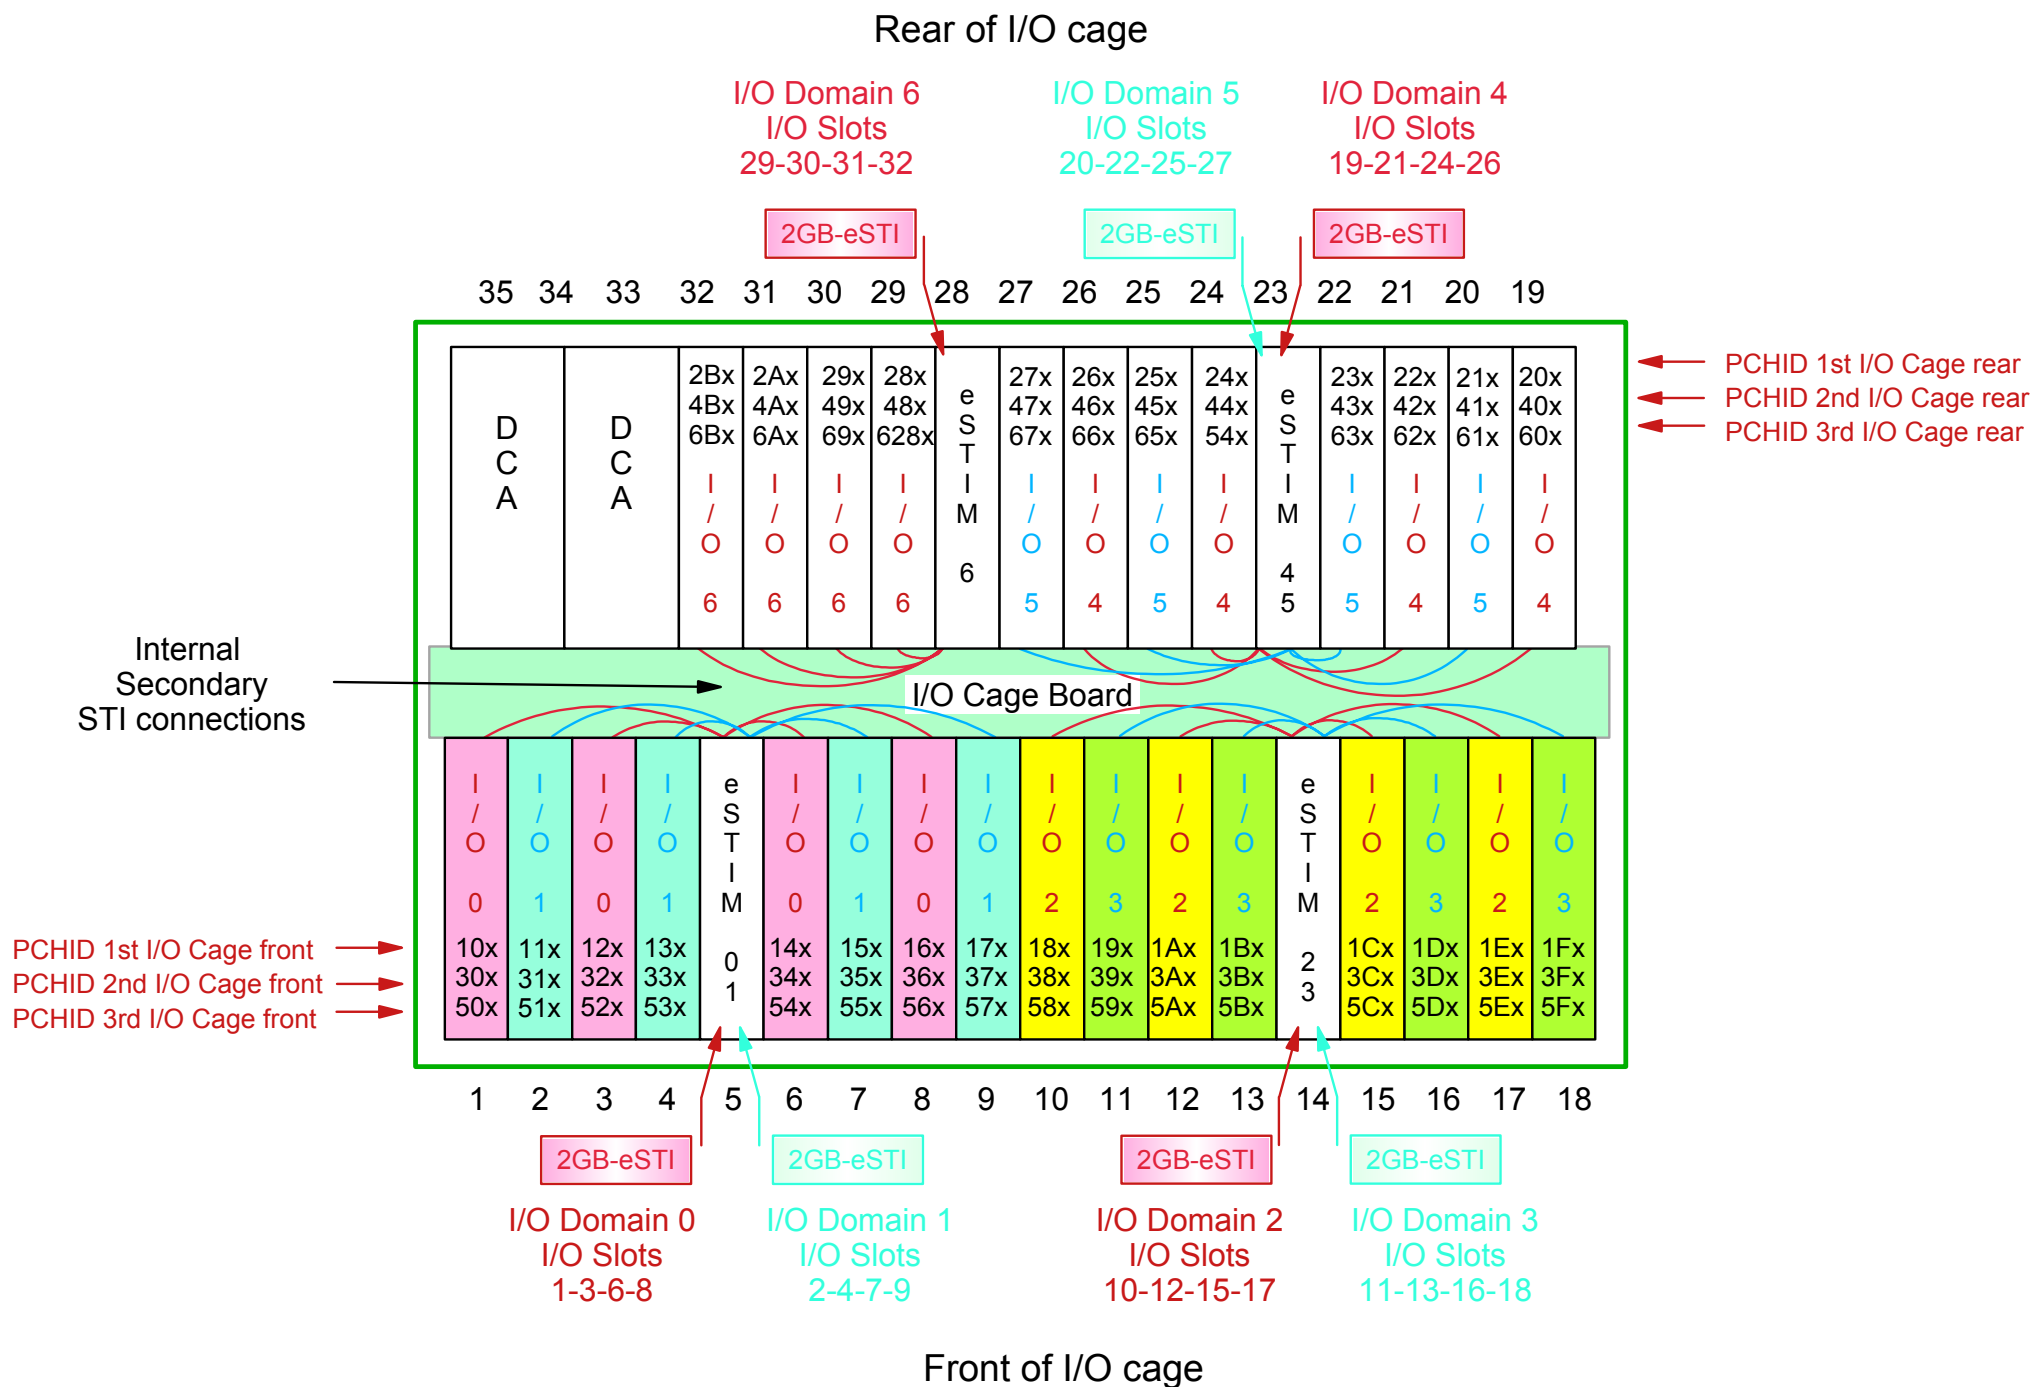
Note:

Number of I/O slots in three I/O cages is 84
 Maximum combined FICON-Express, OSA-Express,
 PCICA/PCIXCC features is 60 for a 3 I/O cage configuration
 Only 2 PCICA cards are supported per I/O cage



- PCHID number assignments to I/O cage

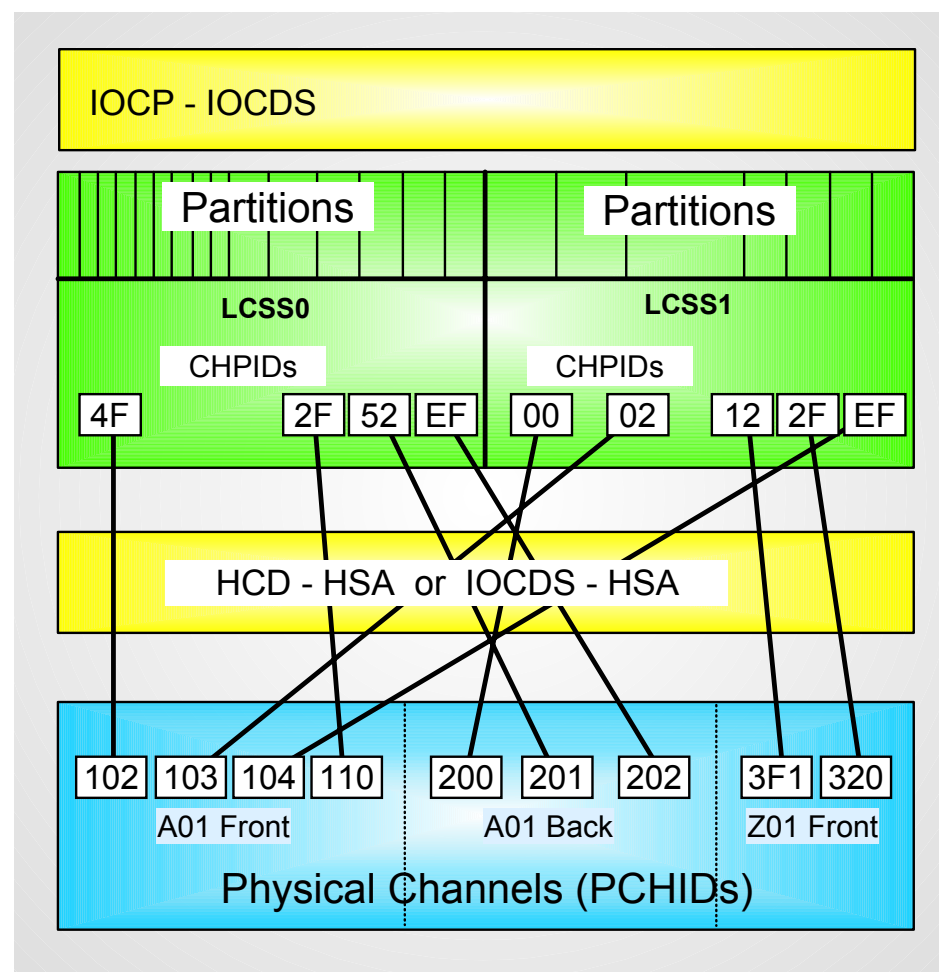


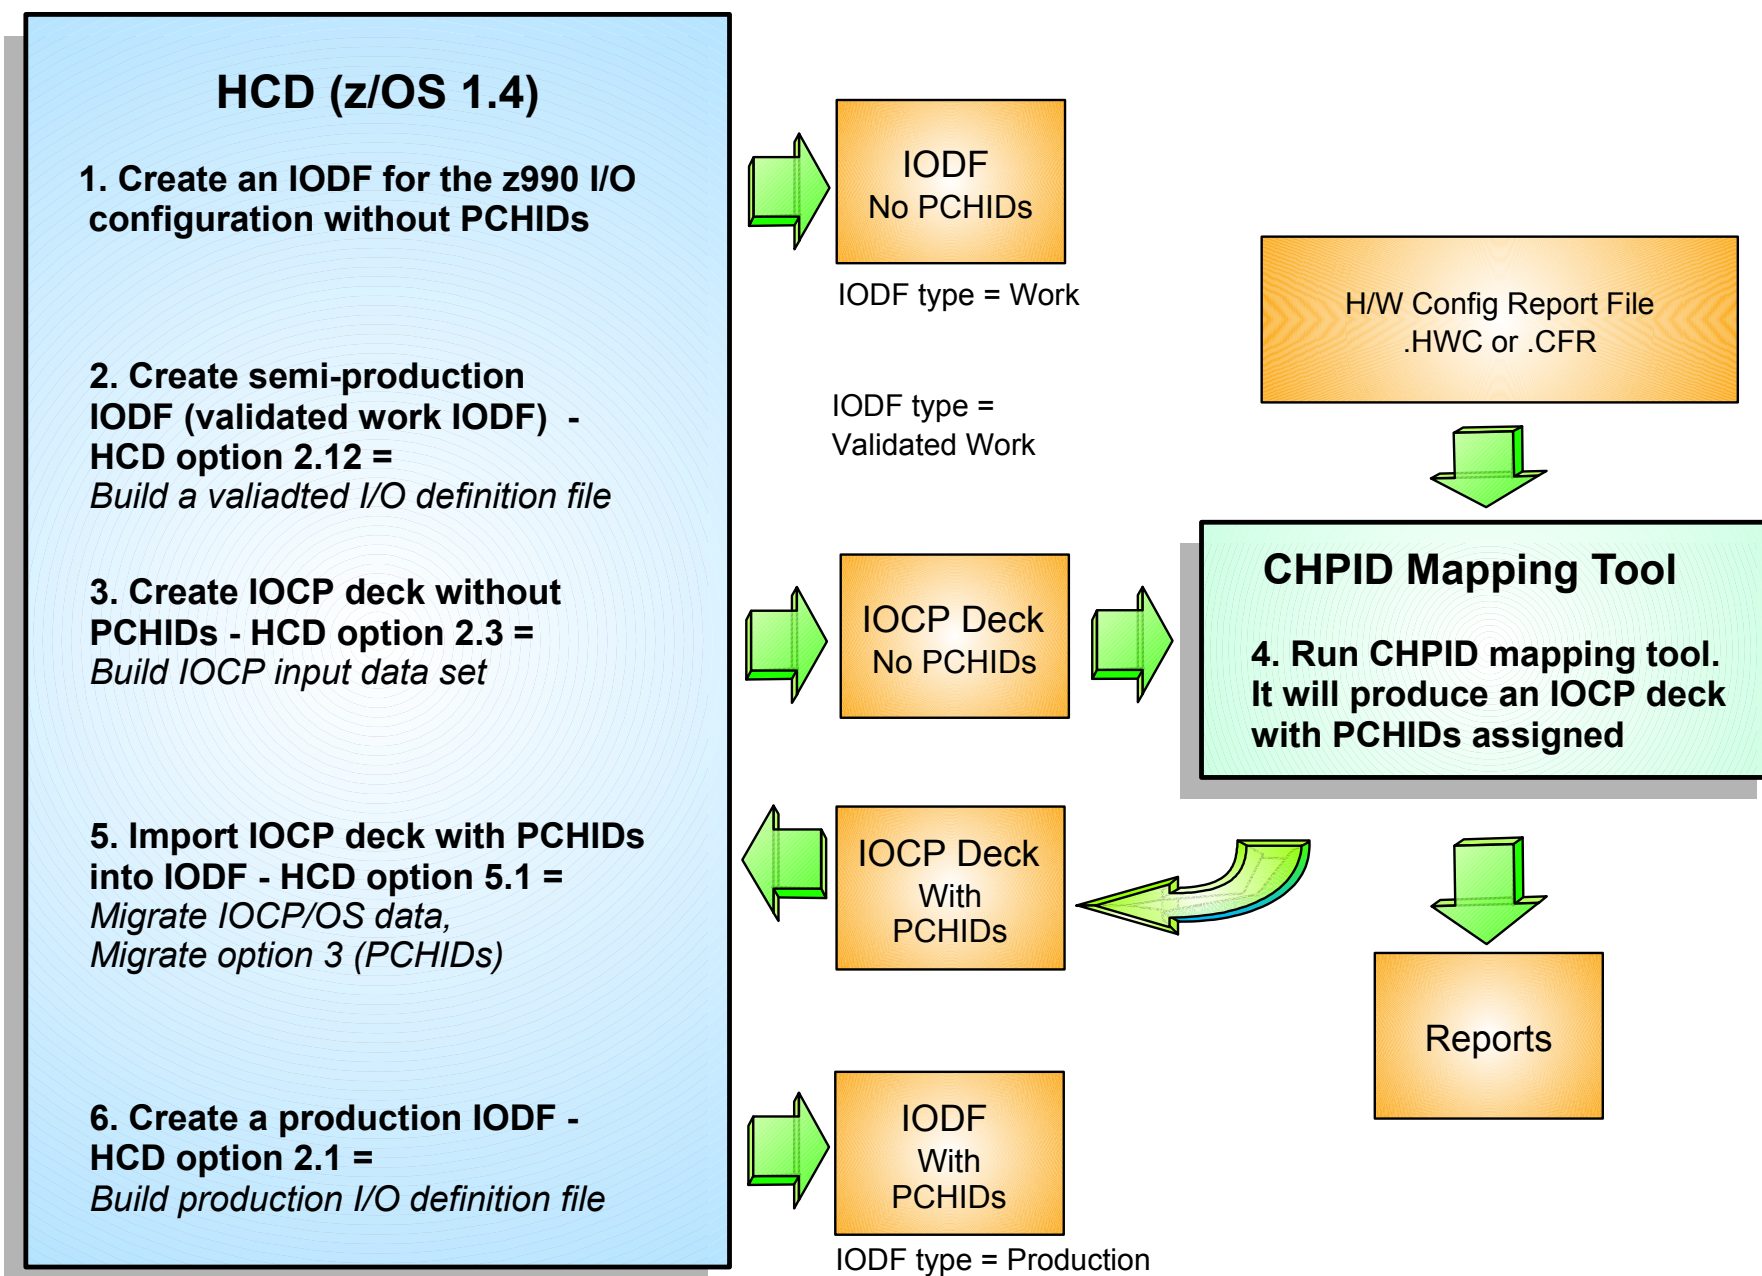




- CHPID assignments to PCHPID (physical channel)

- **CHPID numbers are no longer pre-assigned**
 - i.e. there are **NO** default CHPID numbers
- **Customer MUST assign CHPIDs to channels**
 - **HCD/IOCP Process, and/or**
 - **z990 CHPID Mapping Tool**
- CHPID assignment
 - Define the channel to an LCSS(s)
 - Associate the CHPID number to a physical channel port location (PCHID)
 - CHPID numbers are still 00 - FF and must be unique within an LCSS
- Physical channel location, known as the PCHID, is assigned by manufacturing and reported by eConfig in the PCHID report
- Except for ESCON® sparing, a PCHID relates directly to a jack location on a channel card in a specific I/O slot, in a specific I/O cage
 - **Other exception**
 - ▶ IC and Hipersocket - no PCHID
 - ▶ ICB4 - assigned to CPC cage







CHPIDSTART

30214401

PCHID REPORT

Oct 29, 2003

Machine: 2084-B16 NEW1

Book/Jack/MBA	Cage	Slot	F/C	PCHID/Ports	Comment
0/J.07/1	A19B	06	3393	017/J07	
0/J.11/2	A19B	06	3393	01B/J11	
1/J.07/1	A19B	10	3393	027/J07	
1/J.11/2	A19B	10	3393	02B/J11	
0/J.00/0	A01B	01	1364	100/J00 101/J01	
1/J.00/0	A01B	02	2319	110/J00 111/J01	
0/J.00/0	A01B	03	2319	120/J00 121/J01	
1/J.00/0	A01B	04	2319	130/J00 131/J01	
0/J.00/0	A01B	06	2319	140/J00 141/J01	
1/J.00/0	A01B	07	2319	150/J00 151/J01	
0/J.00/0	A01B	08	2319	160/J00 161/J01	
1/J.00/0	A01B	09	2319	170/J00 171/J01	
0/J.04/1	A01B	10	2319	180/J00 181/J01	
0/J.04/1	A01B	12	2319	1A0/J00 1A1/J01	
0/J.04/1	A01B	15	2319	1C0/J00 1C1/J01	
0/J.04/1	A01B	17	2323	1E0/J00 1E1/J01 1E2/J02 1E3/J03 1E4/J04 1E5/J05 1E6/J06 1E7/J07 1E8/J08 1E9/J09	
0/J.08/2	Z01B	09	2319	370/J00 371/J01	
1/J.04/1	Z01B	10	2319	380/J00 381/J01	
1/J.04/1	Z01B	12	2319	3A0/J00 3A1/J01	
1/J.04/1	Z01B	15	2319	3C0/J00 3C1/J01	
1/J.04/1	Z01B	17	2323	3E0/J00 3E1/J01 3E2/J02 3E3/J03 3E4/J04 3E5/J05 3E6/J06 3E7/J07 3E8/J08 3E9/J09	

Legend:

A19B Top of A frame
A01B Bottom of A frame
Z01B Bottom of Z frame
1364 OSA Express GbE LX
2319 FICON Express LX
2323 ESCON Channel 16 Ports
3393 ICB 4 Function Enablement



```

CH          ID      MSG1='testfr',                                     *
                MSG2='KMT1.IODF03.WORK - 2003-06-27 07:30',          *
                SYSTEM=(2084,1),                                     *
                TOK=('FR2084',000000090ECB2064073026880103178F00000000,0*
                0000000,'03-06-27','07:30:26',' ',' ',' ')
RESOURCE PARTITION=((CSS(0),(LP1,1),(LP2,2),(LP3,3),(LP4,4)),(*
                CSS(1),(LPA,1),(LPB,2),(LPC,3),(LPD,4))),          *
                MAXDEV=((CSS(0),24000),(CSS(1),36000))
CHPID PATH=(CSS(0),80),SHARED,                                     *
                PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=61,TYPE=FC
CHPID PATH=(CSS(0),81),SHARED,                                     *
                PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=61,TYPE=FC
CHPID PATH=(CSS(0),82),SHARED,                                     *
                PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=61,TYPE=FC
CHPID PATH=(CSS(0),83),SHARED,                                     *
                PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=61,TYPE=FC
CHPID PATH=(CSS(0),90),SHARED,                                     *
                PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=62,TYPE=FC
CHPID PATH=(CSS(0),91),SHARED,                                     *
                PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=62,TYPE=FC
CHPID PATH=(CSS(0),92),SHARED,                                     *
                PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=62,TYPE=FC
CHPID PATH=(CSS(0),93),SHARED,                                     *
                PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=62,TYPE=FC
                -----
CNTLUNIT CUNUMBR=3300,                                           *
                PATH=((CSS(0),80,81,82,83,90,91,92,93),(CSS(1),80,81,82,*
                83,90,91,92,93)),UNITADD=((00,256)),              *
                LINK=((CSS(0),6160,6161,6162,6163,6260,6261,6262,6263),(*
                CSS(1),6160,6161,6162,6163,6260,6261,6262,6263)),  *
                CUADD=3,UNIT=2105
CNTLUNIT CUNUMBR=4000,                                           *
                PATH=((CSS(0),84,85,86,87,94,95,96,97),(CSS(1),84,85,86,*
                87,94,95,96,97)),UNITADD=((00,256)),              *
                LINK=((CSS(0),6164,6165,6166,6167,6264,6265,6266,6267),(*
                CSS(1),6164,6165,6166,6167,6264,6265,6266,6267)),  *
                CUADD=0,UNIT=2105

```

z990 CMT IOCP Input (with PCHID)



```
CH      ID      MSG1='testfr',                                     *
MSG2='KMT1.IODF03.WORK - 2003-06-27 07:30',                       *
SYSTEM=(2084,1),                                                    *
TOK=('FR2084',0000000090ECB2064073026880103178F000000000,0*
0000000,'03-06-27','07:30:26',' ',' ')
RESOURCE PARTITION=((CSS(0),(LP1,1),(LP2,2),(LP3,3),(LP4,4)),(*
CSS(1),(LPA,1),(LPB,2),(LPC,3),(LPD,4))),                          *
MAXDEV=((CSS(0),24000),(CSS(1),36000))
CHPID PATH=(CSS(0),80),SHARED,                                     *
PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=61,TYPE=FC,              *
PCHID=1E0
CHPID PATH=(CSS(0),81),SHARED,                                     *
PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=61,TYPE=FC,              *
PCHID=1F0
CHPID PATH=(CSS(0),82),SHARED,                                     *
PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=61,TYPE=FC,              *
PCHID=261
CHPID PATH=(CSS(0),83),SHARED,                                     *
PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=61,TYPE=FC,              *
PCHID=441
CHPID PATH=(CSS(0),90),SHARED,                                     *
PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=62,TYPE=FC,              *
PCHID=151
CHPID PATH=(CSS(0),91),SHARED,                                     *
PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=62,TYPE=FC,              *
PCHID=351
CHPID PATH=(CSS(0),92),SHARED,                                     *
PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=62,TYPE=FC,              *
PCHID=341
CHPID PATH=(CSS(0),93),SHARED,                                     *
PARTITION=((LP1,LP2,LP3,LP4),(=)),SWITCH=62,TYPE=FC,              *
PCHID=3F0
-----
CNTLUNIT CUNUMBR=3300,                                           *
PATH=((CSS(0),80,81,82,83,90,91,92,93),(CSS(1),80,81,82,*
83,90,91,92,93)),UNITADD=((00,256)),                               *
LINK=((CSS(0),6160,6161,6162,6163,6260,6261,6262,6263),(*
CSS(1),6160,6161,6162,6163,6260,6261,6262,6263)),               *
CUADD=3,UNIT=2105
CNTLUNIT CUNUMBR=4000,                                           *
PATH=((CSS(0),84,85,86,87,94,95,96,97),(CSS(1),84,85,86,*
87,94,95,96,97)),UNITADD=((00,256)),                               *
LINK=((CSS(0),6164,6165,6166,6167,6264,6265,6266,6267),(*
CSS(1),6164,6165,6166,6167,6264,6265,6266,6267)),               *
CUADD=0,UNIT=2105
```



```

                                Channel Path List      Row 21 of 40 More: <
>
Command ===> _____ Scroll ===>
CSR

Select one or more channel paths, then press Enter. To add, use F11.

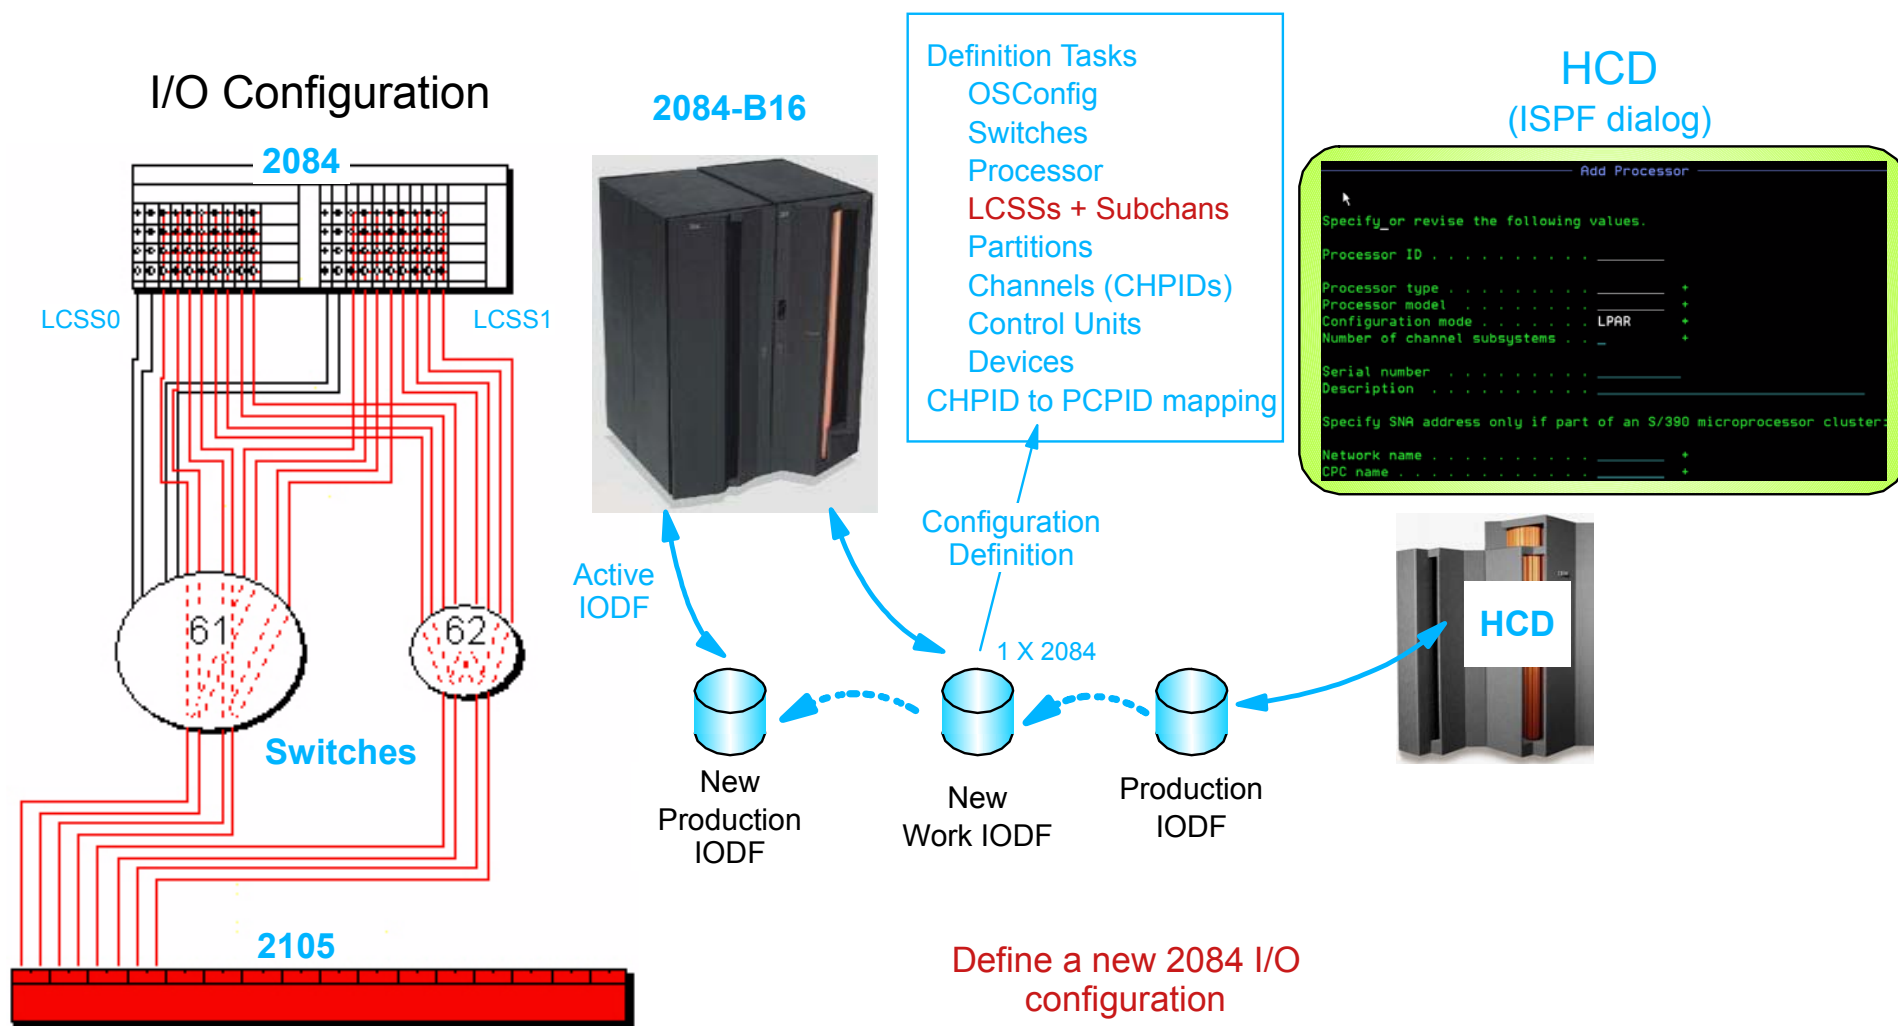
Channel Subsystem ID : 0
1=LP1      2=LP2      3=LP3      4=LP4      5=
6=          7=          8=          9=          A=
B=          C=          D=          E=          F=

                                I/O Cluster  ----- Partitions 0x -----
/ CHPID Type+ Mode+ Mngd  Name +      1 2 3 4 5 6 7 8 9 A B C D E F
PCHID
_ 80      FC      SHR      No      _____ a a a a _ _ _ _ _ _ _ _ _ _ 1E0
_ 81      FC      SHR      No      _____ a a a a _ _ _ _ _ _ _ _ _ _ 1F0
_ 82      FC      SHR      No      _____ a a a a _ _ _ _ _ _ _ _ _ _ 261
_ 83      FC      SHR      No      _____ a a a a _ _ _ _ _ _ _ _ _ _ 441
_ 84      FC      SHR      No      _____ a a a a _ _ _ _ _ _ _ _ _ _ 340
_ 85      FC      SHR      No      _____ a a a a _ _ _ _ _ _ _ _ _ _ 1E1
_ 86      FC      SHR      No      _____ a a a a _ _ _ _ _ _ _ _ _ _ 150
_ 87      FC      SHR      No      _____ a a a a _ _ _ _ _ _ _ _ _ _ 350
F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset
F7=Backward

F8=Forward      F9=Swap      F10=Actions      F11=Add      F12=Cancel
F13=Instruct
F19=Left      F20=Right      F22=Command

```

Use Channel Path List and an PF20 to display the PCHIDs





2064	2084	Comments
OS/390 HCD release 9 Create work IODF	z/OS 1.4 HCD Create work IODF	
Operating Systems EDTs Esoteric Consoles	Operating Systems - option 1.1 EDTs Esoteric Consoles	
Switches (ESCON and FICON) Ports	Switches - option 1.2 (ESCON and FICON) Ports	
Processors Requires 2064 PIT Type, Mode	Processors - option 1.3 Requires 2084 PIT Type, Mode , #CSSs	
--- N/A ---	(L)CSSs - option 1.3.s CSS.ID Max # Devices for each CSS	No HSA expansion support on the 2084 SE. 2084 always supports the maximum PCUs, CUHs, PCHIDs # of subchannels specified in HCD
Partitions (for the processor) Name unique to a processor	Partitions (for a CSS) - option 1.3.s.p LP names are unique across all CSSs	IOCP changes from using Partition numbers to MIF_ID for a partition
Channels (for the processor) Type, Mode	Channels (for a CSS) - option 1.3.s.s Type, Mode - SPANNED, PCHID	2084 channels do not have default CHPID numbers (require PCHIDs) 2084 crypto function does not use CHPID numbers.
CU Type, Paths, Link Addresses	CU - option 1.4 Type, then define paths to each CSS CSS.Paths , Link Addresses	Define ALL processors and CSSs (with partitions and channels) before CU definitions
I/O Device	I/O Device - option 1.5	



z990 Processor Complex


Define - processor *MYCEC*, 2084 model A/B/C/D, mode LPAR

LCSSs + IDs + MaxDevs, Partitions: Names, MIF.IDs, (Partition IDs - profile)

Channels - + Dedicated - Reconfigurable - Shared - Spanned

MYCEC - 2084 model xxx - LPAR

2084 model A/B/C/D



2084 CPC (Books)
PUs
Memory
Channels

PCHID 010B	PCHID 010C	PCHID 010D	PCHID 0160
---------------	---------------	---------------	---------------



PCHID 020A

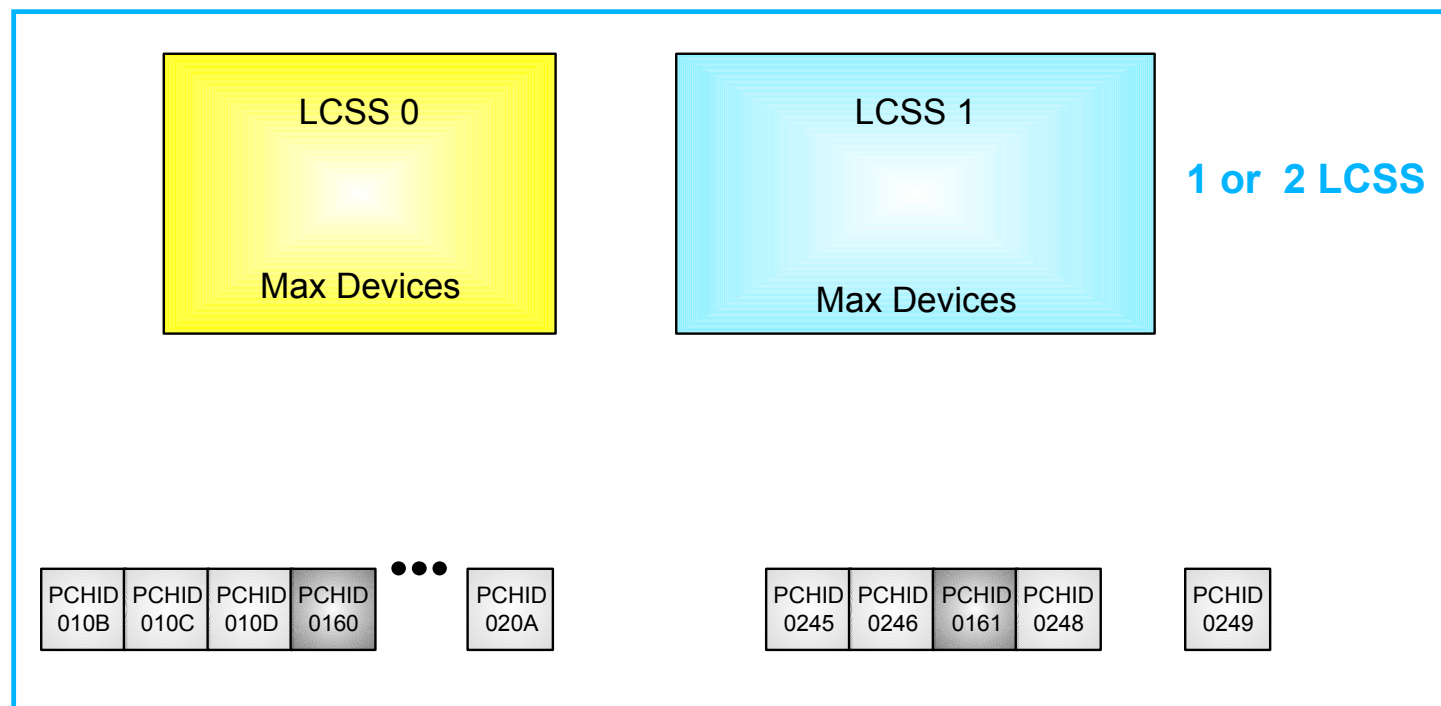
PCHID 0245	PCHID 0246	PCHID 0161	PCHID 0248
---------------	---------------	---------------	---------------

PCHID 0249



Define - processor *MYCEC*, 2084 model A/B/C/D, mode LPAR
LCSSs + IDs + MaxDev, Partitions: Names, MIF.IDs, (Partition IDs - profile)
Channels - + Dedicated - Reconfigurable - Shared - Spanned

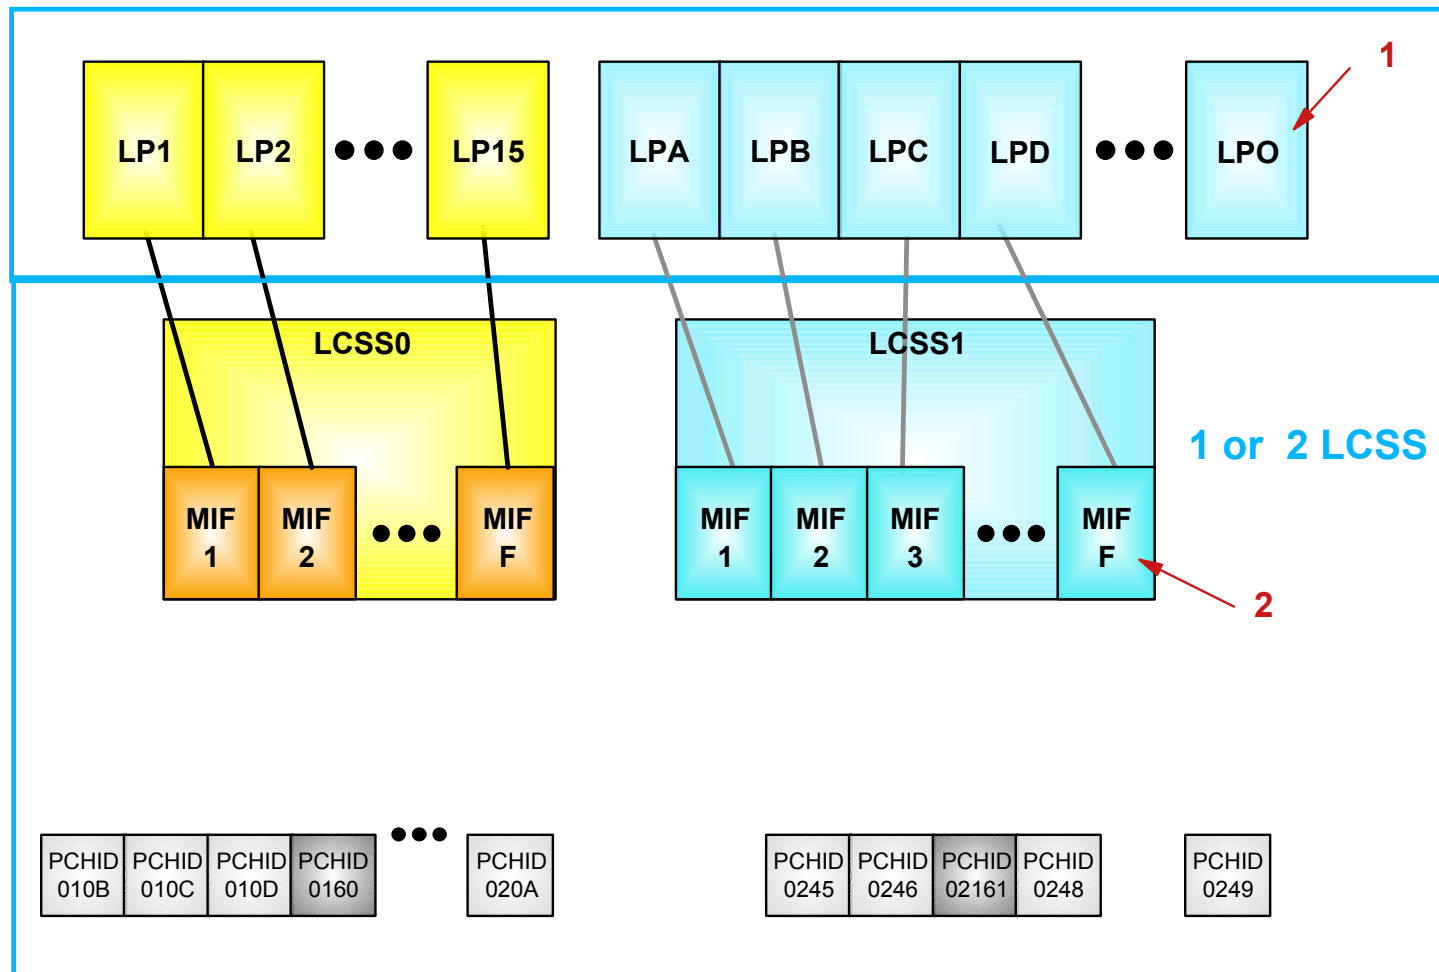
MYCEC - 2084 model xxx - LPAR





Define - processor *MYCEC*, 2084 model A/B/C/D, mode LPAR
LCSSs + IDs + MaxDev, **Partitions: Names, MIF.IDs, (Partition IDs - profile)**
Channels - + Dedicated - Reconfigurable - Shared - Spanned

MYCEC - 2084 model xxx - LPAR

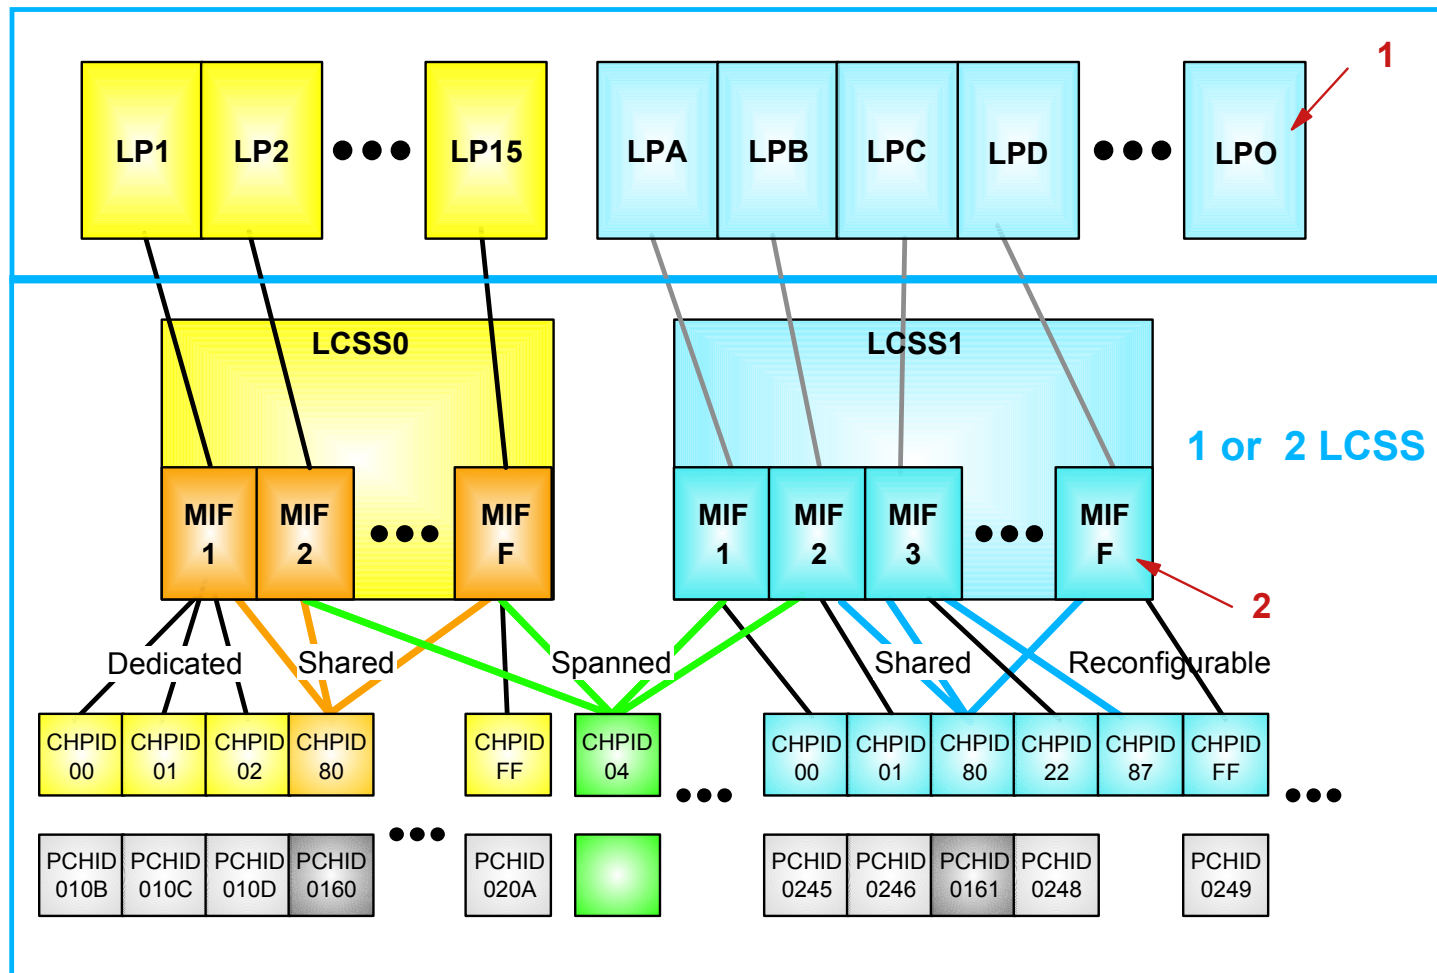


1. Logical partition names are specified in the I/O definition process (HCD or IOCP) must be a unique name for each of the logical partitions across all the LCSSs in the z990
2. Logical partition MIF.IDs are specified in the I/O definition process (HCD or IOCP) and must be a unique value x'1' to x'F' for each of the logical partitions in an LCSSs in the z990



Define - processor *MYCEC*, 2084 model A/B/C/D, mode LPAR
LCSSs + IDs + MaxDev, Partitions: Names, MIF.IDs, (Partition IDs - profile)
Channels: CHPID & type, Dedicated - Reconfigurable - Shared - Spanned

MYCEC - 2084 model xxx - LPAR



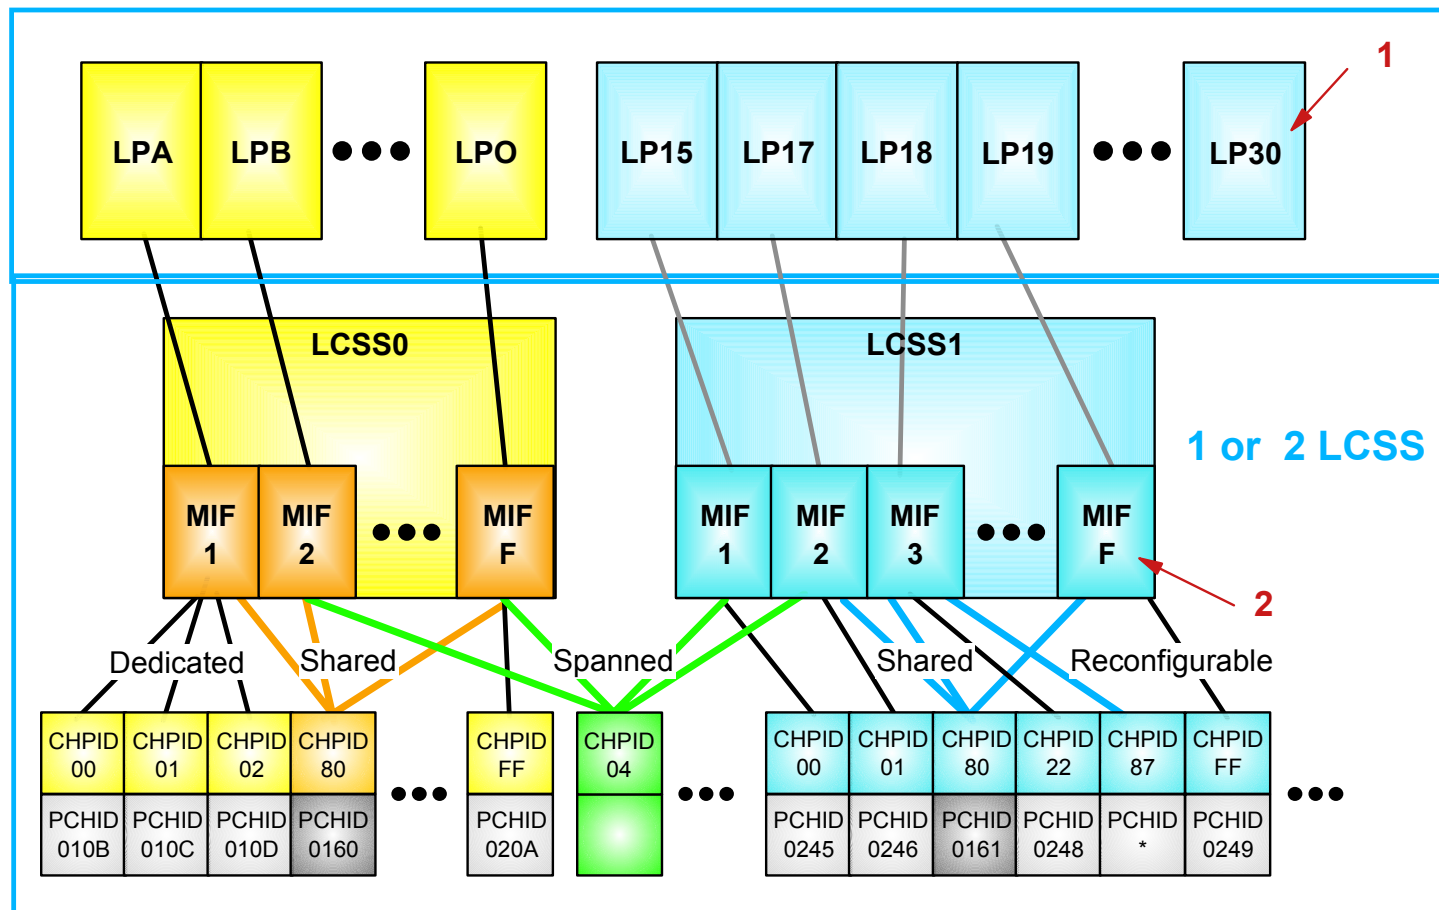
1. Logical partition names are specified in the I/O definition process (HCD or IOCP) and must be a unique name for each of the logical partitions across all the LCSSs in the z990
2. Logical partition MIF.IDs are specified in the I/O definition process (HCD or IOCP) and must be a unique value x'1' to x'F' for each of the logical partitions in an LCSSs in the z990

- **Spanned CHPID support for HyperSockets™ (IQD) and Internal Coupling (IC) channels 10/31/03**
- Allows sharing of IC and IQD channels among LPs in multiple Logical Channel Subsystems (LCSS)



Define - processor *MYCEC*, 2084 model A/B/C/D, mode LPAR
LCSSs + IDs + MaxDev, Partitions: Names, MIF.IDs, (Partition IDs - profile)
Channels: CHPID & type, Dedicated - Reconfigurable - Shared - Spanned
Map the CHPIDs to PCHIDs

MYCEC - 2084 model xxx - LPAR

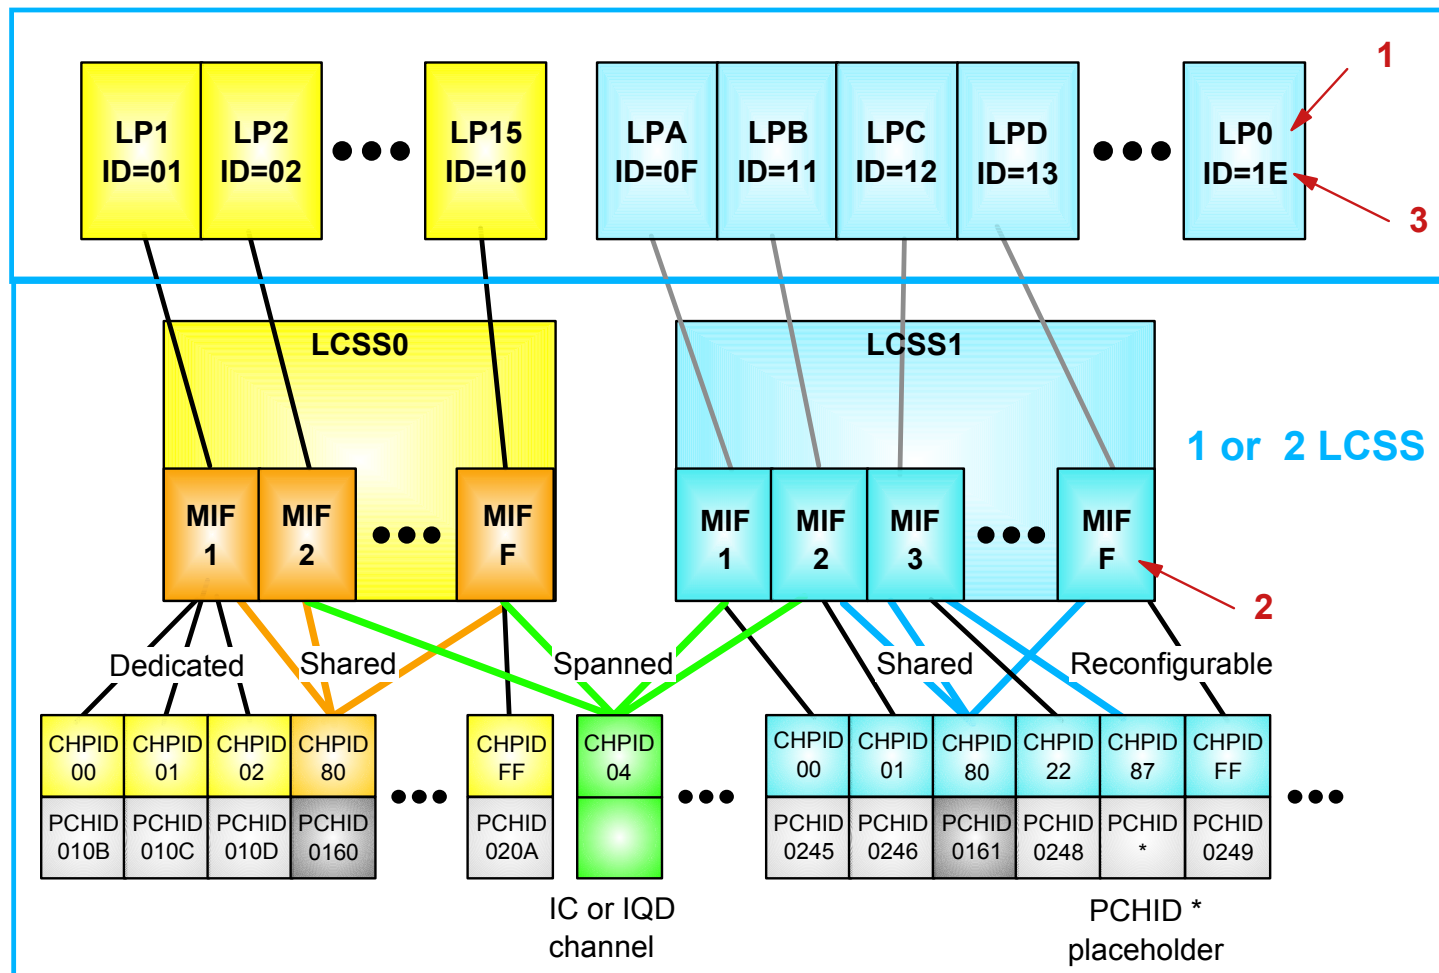


1. Logical partition names are specified in the I/O definition process (HCD or IOCP) and must be a unique name for each of the logical partitions across all the LCSSs in the z990
2. Logical partition MIF.IDs are specified in the I/O definition process (HCD or IOCP) and must be a unique value x'1' to x'F' for each of the logical partitions in an LCSSs in the z990

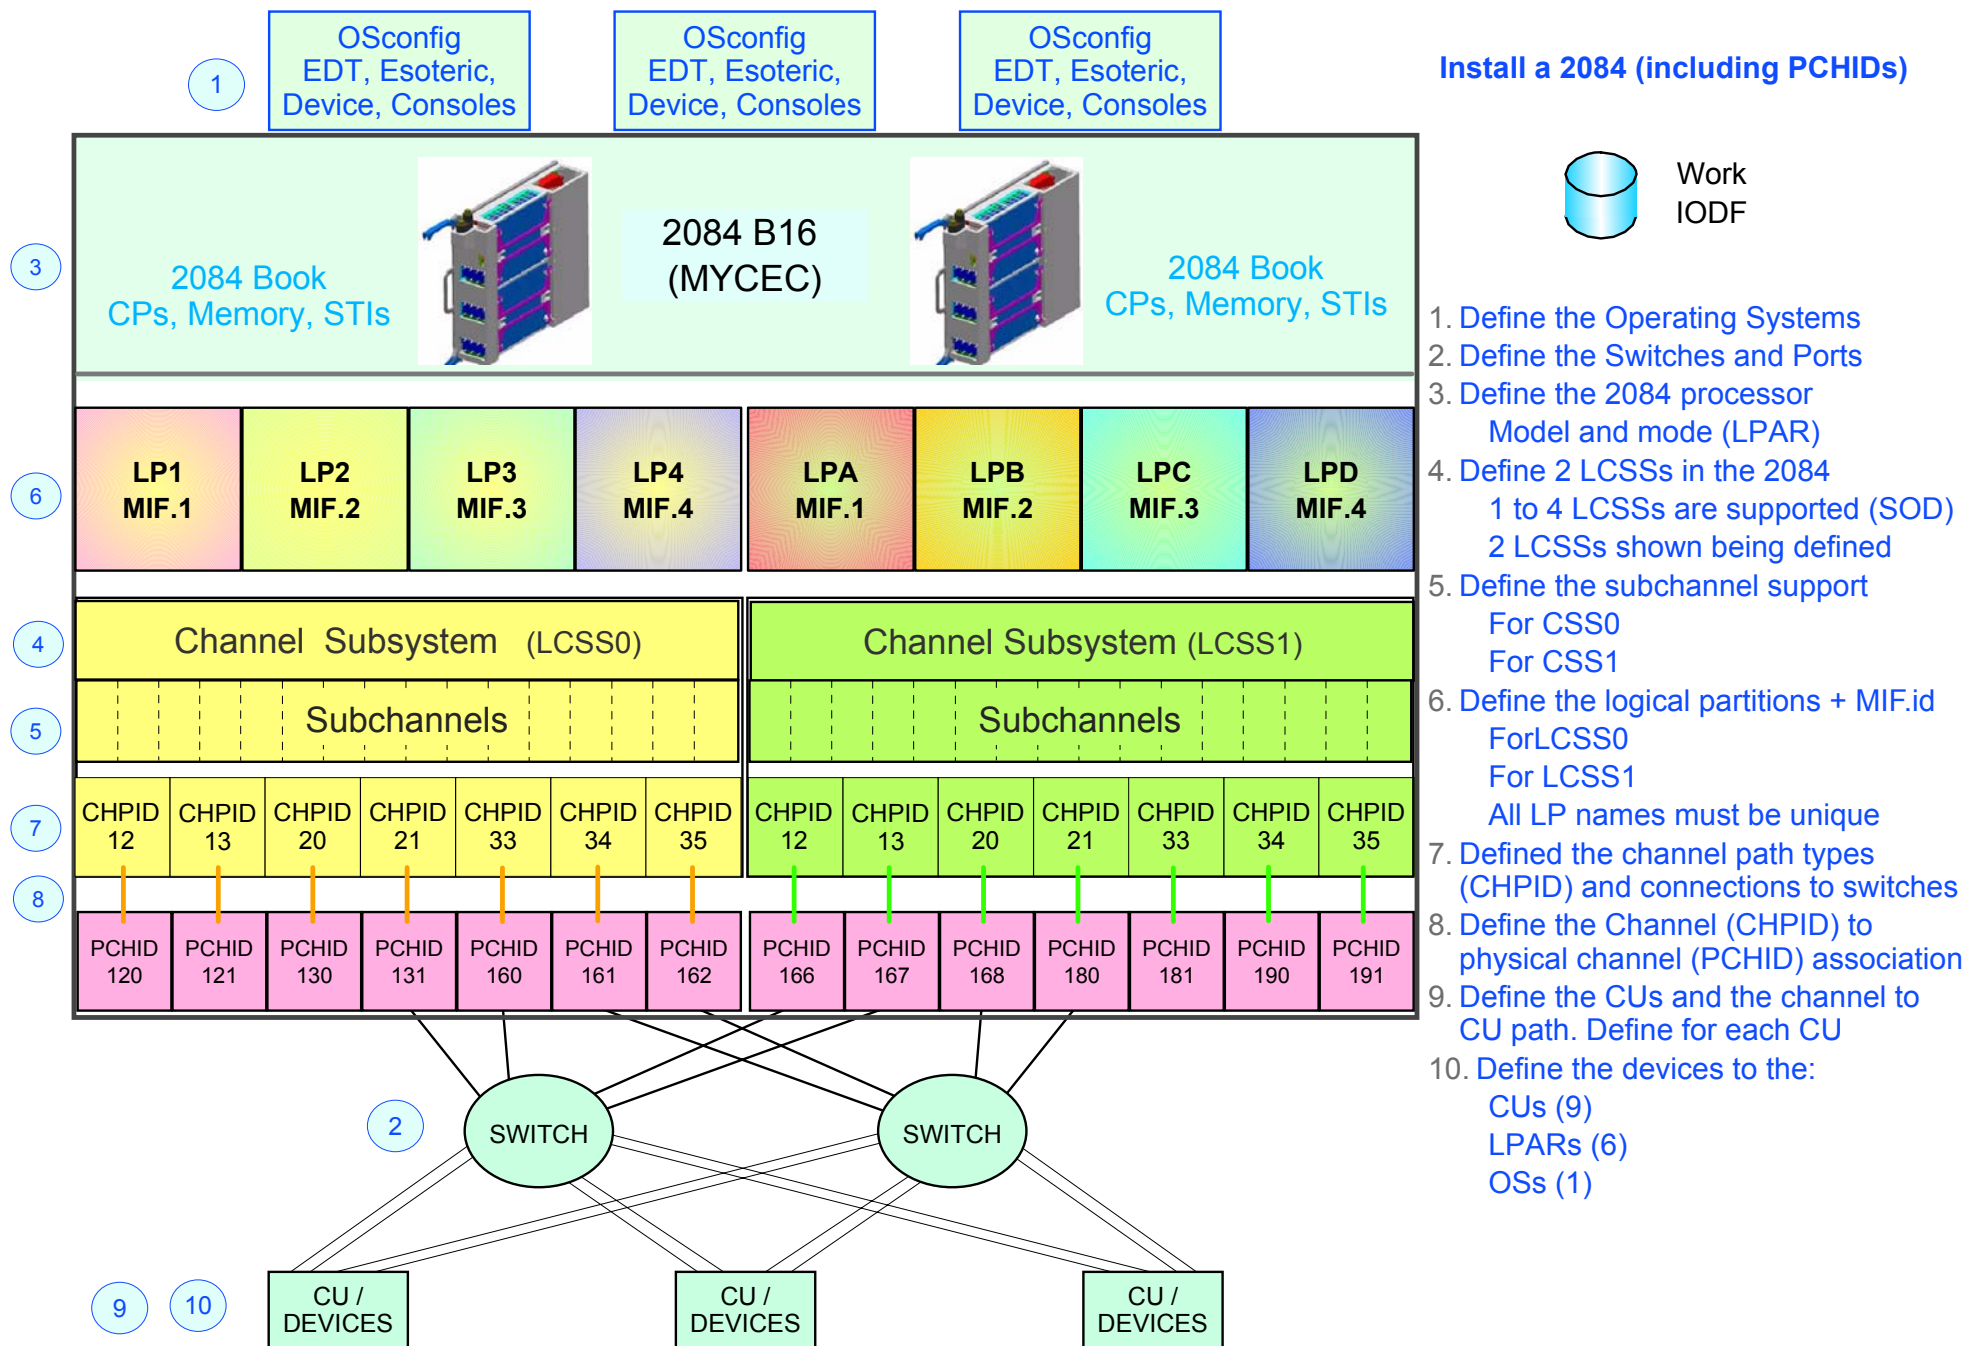
- Specifying PCHID as an * value (HCD CHPID placeholder) requires HCD APAR

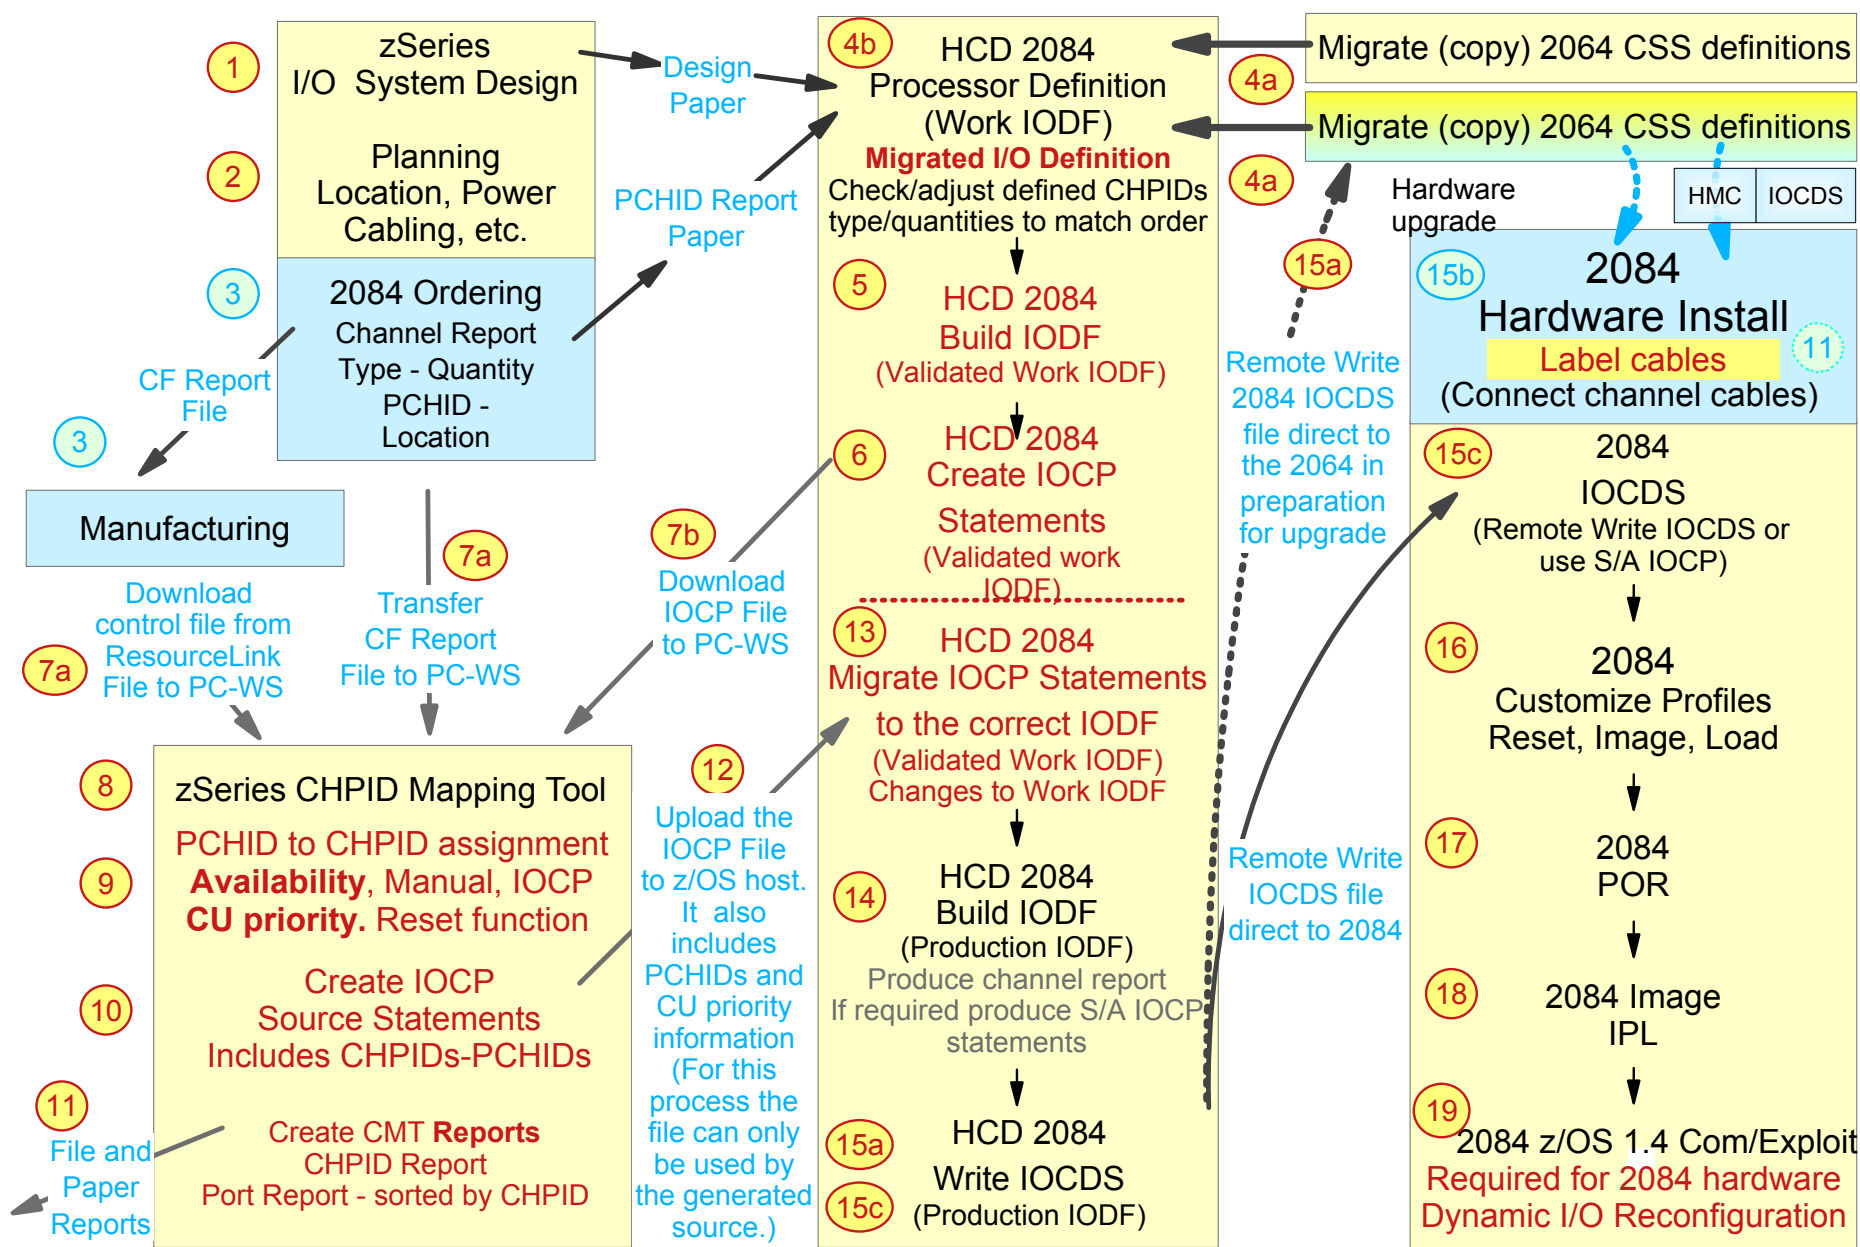


Define - processor *MYCEC*, 2084 model A/B/C/D, mode LPAR
LCSSs + IDs + MaxDev, Partitions: Names, MIF.IDs, (Partition IDs - profile)
Channels: CHPID & type, Dedicated - Reconfigurable - Shared - Spanned
Map the CHPIDs to PCHIDs
Customize the Image Profiles



1. Logical partition names are specified in the I/O definition process (HCD or IOCP) and must be a unique name for each of the logical partitions across all the LCSSs in the z990
2. Logical partition MIF.IDs are specified in the I/O definition process (HCD or IOCP) and must be a unique value x'1' to x'F' for each of the logical partitions in an LCSSs in the z990
3. The Logical Partition ID is specified (by the user) in the z990 Image profile and must be a unique value from x'00' to x'3F' for each of the logical partitions across all the LCSSs in the z990







- Planning items - configuring / ordering
 - No Parallel channels (BL or BY), no OSA-2 channels (OSA)
 - ▶ No 4 Port ESCON channel card, also requires a cable connector change - IBM duplex to MT-RJ
 - ▶ No FICON 1.5 channel card, also requires cable connector change - SC to LC
 - Consider 'Plan-Ahead' for channels for additional I/O cages - prevents disruption
- Planning items - 2084 processor - any change to the following is disruptive
 - Number of LCSSs (1 or 2)
 - Quantity of logical partitions per LCSS and for the 2084 in total
 - Logical partition names - they must be unique across the CSSs
 - Number of Subchannels per LCSS
 - If any of the above is change, HCD will not allow a dynamic change
- Planning items - I/O configuration definition
 - FICON CTC (FCTC) target 2084 CUADD - must be CSS.ID + MIF.ID
 - FCTC support change required for the 9672
 - If CTCs are defined, always produce an HCD CTC connection report for checking
 - CF Connectivity - must be defined by the user
 - IQD connectivity
 - ICP connectivity - allows connections across LCSSs
- Planning items - I/O definition file, and HCD IOCP statements
 - Do not change a validated work IODF after a generation of the IOCP statements
 - The HCD validated work IOCP statement file must only be used by the CMT
 - The HCD validated work IOCP statement file CANNOT be used by any IOCP program
- Planning items - HCD panel changes
 - HCD hierarchical changes - use HCD options 1.3.s.s to display Channel Path List panel
 - Use the HCD channel path list panel (1.3.s.s) and PF20 to display PCHIDs
 - Use the HCD channel path panel and PF20 + PF20 to display the SPANNED channels



IBM @server zSeries 990

ITSO - z990 Technical Workshop (06/2003)

Introduction - Processor Complex
Feature Support and Availability

2084-z990

ITSO Poughkeepsie



Initial Hardware June 16, 2003	MES Changes Sept 15, 2003	Additional Changes Oct 31, 2003	Hardware Statement of Direction
2084 Models A08 and B16		2084 Models C24 and D32	> 16 CPs per LPAR
		Concurrent Model Upgrade (Book Add)	
128 GB Storage	Add memory	256 GB storage	
2 LCSSs (0 and 1)			4 LCSSs
30 LPARs defined			60 LPARs supported
15 LPARs activated		30 LPARs activated	
LCSS 0 dynamic I/O reconfiguration		LCSS 1 dynamic I/O reconfiguration	
512 channels	Add channels		Spanned external channels (except ESCON)
16 Port ESCON			
FICON Express			
OSA-Express			
16 Hipersockets		Spanned Hipersockets	
IC channels		Spanned IC Channels	
ISC-3			
ICB-4, ICB-3, ICB-2			No ICB-2 on future machines
		FCP Base	
		FCP Enhancements	FCP disk IPL support
No Parallel Channels no ATM, or FDDI			



Initial Hardware June 16, 2003	MES Changes Sept 15, 2003	Additional Changes Oct 31, 2003	Hardware Statement of Direction
CBU setup and activate			
CIU only	On/Off Capacity on Demand or CIU activation		
	CUoD MES Upgrade		
			CFCC change reduce outage
CP crypto assist			
PCICA crypto		PCIXCC crypto	
		TKE 4.0 workstation	TKE Smart Card Reader
			Secure Key entry via TKE
			PCIXCC Linux support
			PCIXCC z/VM guest support



IBM @server zSeries 990

z990 Processor Introduction
z990 Processor Technical Information
End of Presentation

2084 - z990
ITSO Poughkeepsie