

Session Title:
IBM zEnterprise 196 (z196) Hardware Overview

Session ID: 8670



Speaker: Harv Emery



Permission is granted to SHARE to publish this presentation in the SHARE Proceedings.
IBM retains its right to distribute copies of this presentation to whomever it chooses.

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*BladeCenter®, DB2®, e business(logo)®, DataPower®, ESCON, eServer, FICON, IBM®, IBM (logo)®, MVS, OS/390®, POWER6®, POWER6+, POWER7, Power Architecture®, S/390®, System p, System p5, System x, System z, System z9®, System z10®, WebSphere®, X-Architecture®, zEnterprise, z9®, z10, z/Architecture®, z/OS®, z/VM®, z/VSE, zSeries®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

Notes:

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

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

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

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

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

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

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

IBM zEnterprise System – Best in Class Systems and Software Technologies

A system of systems that unifies IT for predictable service delivery



Unified management for a smarter system: **zEnterprise Unified Resource Manager**

- Unifies management of resources, extending IBM System z® qualities of service end-to-end across workloads
- Provides platform, hardware and workload management

The world's fastest and most scalable system:
IBM zEnterprise™ 196 (z196)

- Ideal for large scale data and transaction serving and mission critical applications
- Most efficient platform for Large-scale Linux® consolidation
- Leveraging a large portfolio of z/OS® and Linux on System z applications
- Capable of massive scale up, over 50 Billion Instructions per Second (BIPS)

Scale out to a trillion instructions per second:
IBM zEnterprise BladeCenter® Extension (zBX)

- Selected IBM POWER7® blades and IBM x86 Blades¹ for tens of thousands of AIX® and Linux applications
- High performance optimizers and appliances to accelerate time to insight and reduce cost
- Dedicated high performance private network



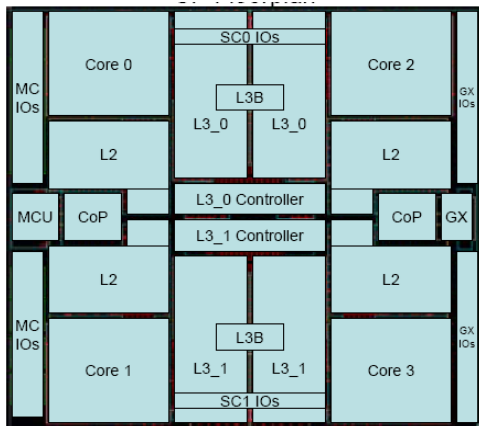
¹ All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

zEnterprise System z196 Processors and Memory

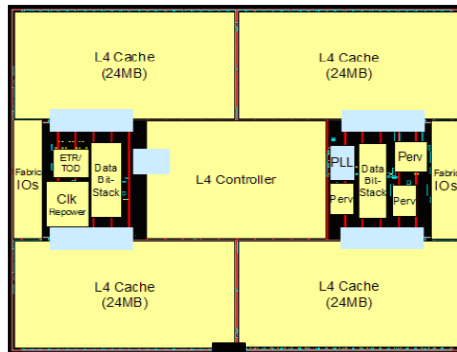


z196 PU chip, SC chip and MCM


z196 Quad Core PU CHIP



96 MB SC CHIP



Front View



MCM

z10 EC MCM vs z196 MCM Comparison

z10 EC MCM

▪ MCM

–96mm x 96mm in size

–5 PU chips per MCM

- Quad core chips with 3 or 4 active cores
- PU Chip size 21.97 mm x 21.17 mm
- 4.4 GHz
- Superscalar, In order execution
- L1: 64K I / 128K D private/core
- L1.5: 3M I+D private/core

–2 SC chips per MCM

- L2: 2 x 24 M = 48 M L2 per book
- SC Chip size 21.11 mm x 21.71 mm

–1800 Watts

z196 MCM

▪ MCM

–96mm x 96mm in size

–6 PU chips per MCM

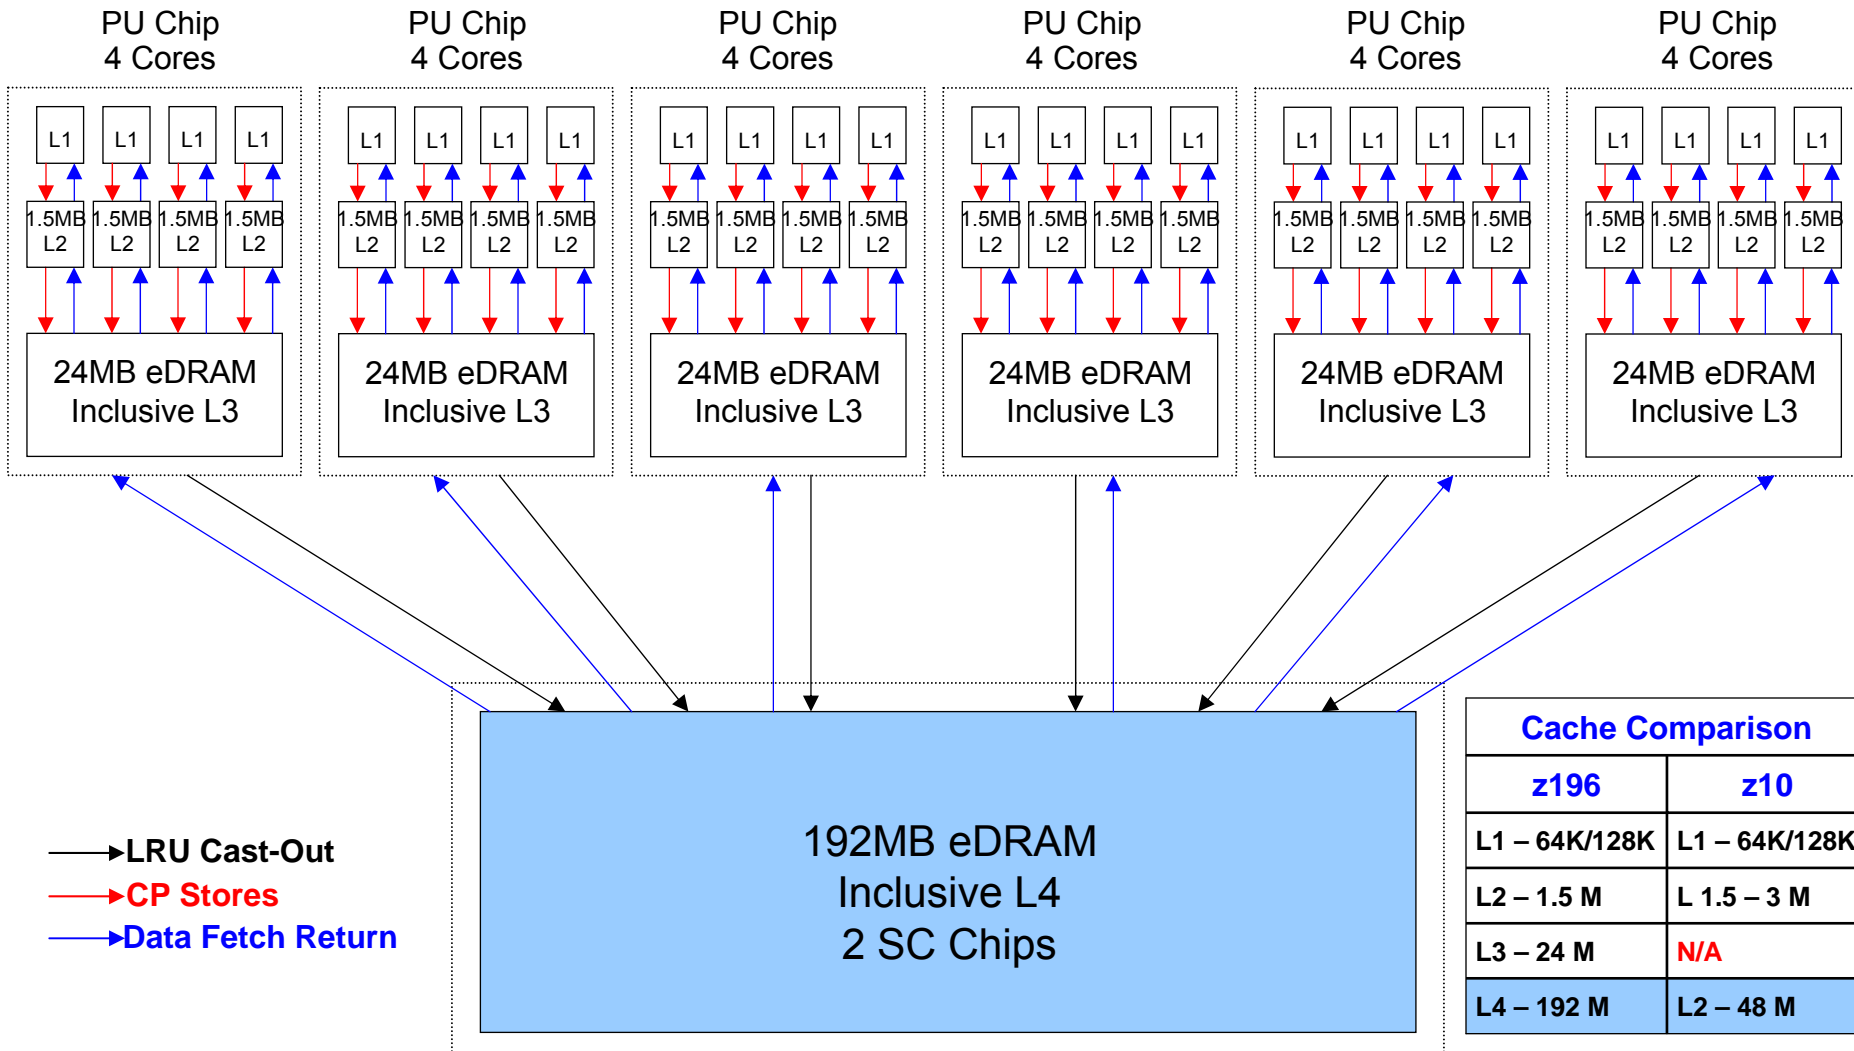
- Quad core chips with 3 or 4 active cores
- PU Chip size 23.7 mm x 21.5 mm
- 5.2 GHz
- Superscalar, OOO execution
- L1: 64K I / 128K D private/core
- L2: 1.5M I+D private/core
- L3: 24MB/chip - shared

–2 SC chips per MCM

- L4: 2 x 96 M = 192 M L4 per book
- SC Chip size 24.5 mm x 20.5 mm

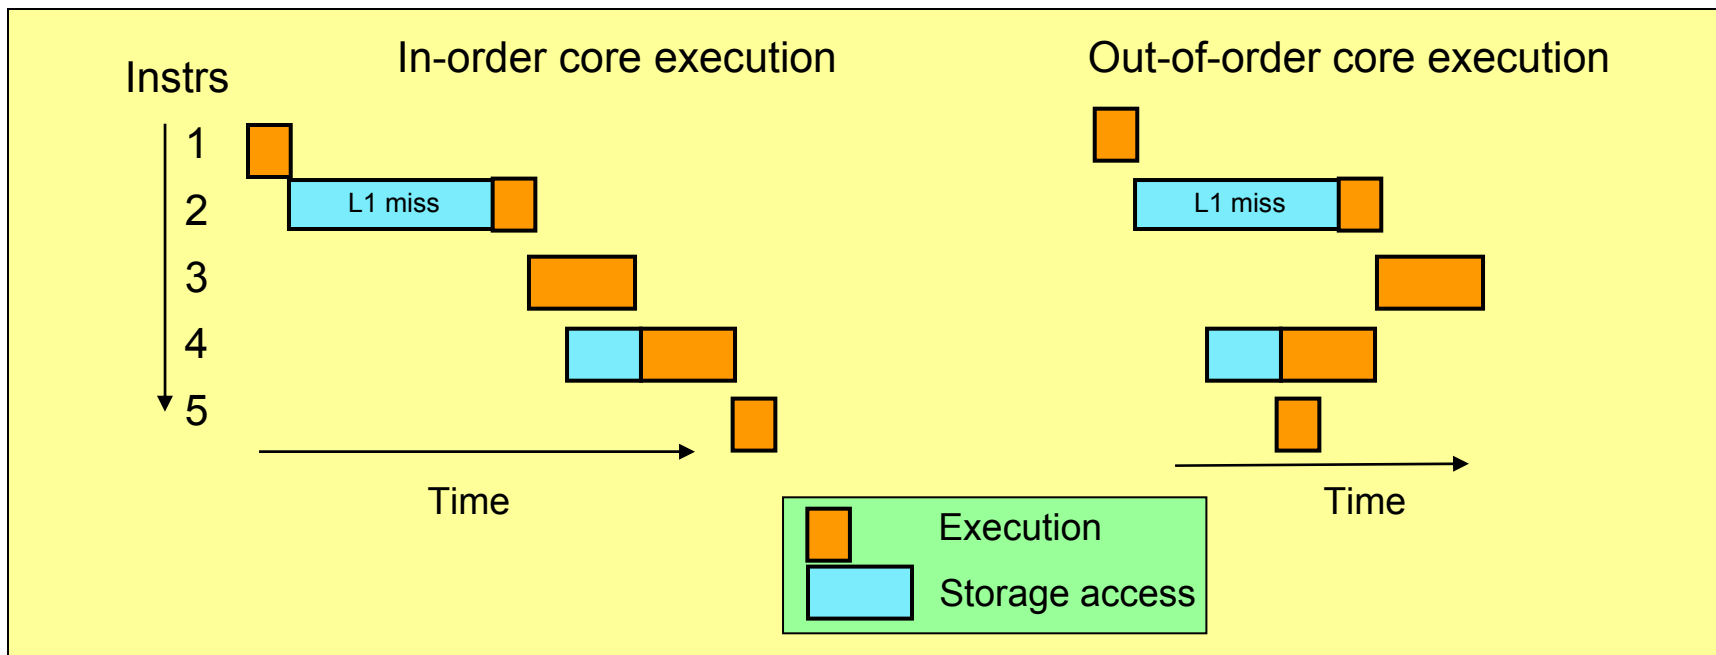
–1800 Watts

z196 Book Level Cache Hierarchy



z196 Out-of-Order (OOO) Value



- **z196 has the first System z CMOS out-of-order core**
- **z196 has the first System z out-of-order core since 1991**
- **OOO yields significant performance benefit for applications through**
 - Re-ordering instruction execution
 - Later (younger) instructions can execute ahead of an older stalled instruction
 - Re-ordering storage accesses and parallel storage accesses
- **OOO maintains good performance growth for traditional apps**



z196 New instructions and instruction enhancements Designed to provide new function and improve performance

- **High-Word Facility** (30 new instructions)
 - Independent addressing to high word of 64 bit General Purpose Registers
 - Effectively provides software with 16 additional registers for arithmetic
- **Interlocked-Access Facility** (12 new instructions)
 - Interlocked (atomic) load, value update and store operation in a single instruction
- **Load/Store-on-Condition Facility** (6 new instructions)
 - Load or store conditionally executed based on condition code
 - Dramatic improvement in certain codes with highly unpredictable branches
- **Distinct-Operands Facility** (22 new instructions)
 - Independent specification of result register (different than either source register)
 - Reduces register value copying
- **Population-Count Facility** (1 new instruction)
 - Hardware implementation of bit counting ~5x faster than prior software implementations
- **Floating-Point-Extension Facility** (21 new instructions, 34 instruction enhancements)
- **Message-Security Assist Extensions 3 and 4** – (5 new instructions, 6 instruction enhancements)
- **And more**

System Offering Overview



z196

z196 machine type: 2817

Processors

- ▶ 20 or 24 available cores per book
- ▶ Sub-capacity available up to 15 CPs
 - ▶ 3 sub-capacity points
- ▶ 2 spares designated per system

Memory for customer purchase

- ▶ System minimum = 32 GB
 - ▶ 16 GB separate HSA
- ▶ Maximum: 3TB / 768 GB per book
- ▶ Increments: 32 to 256 GB

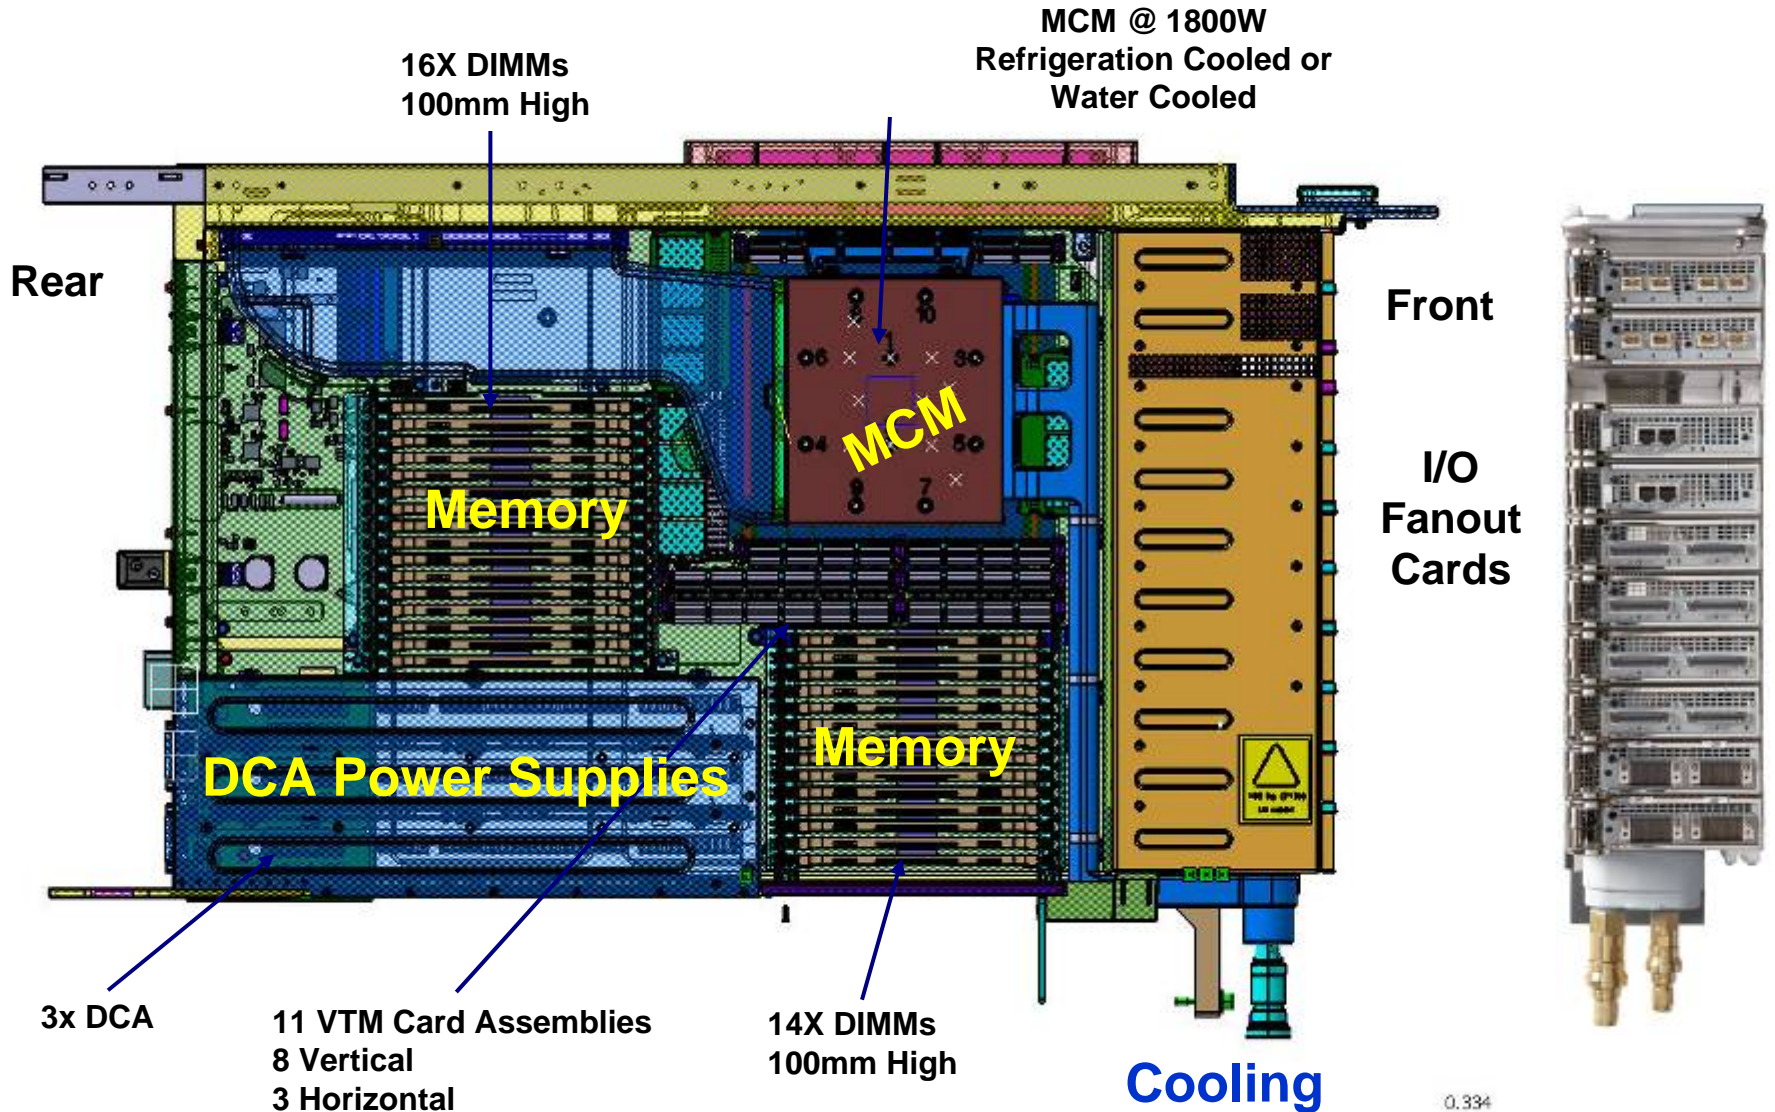
I/O Interconnects: (Same as z10 EC)

- ▶ 6 GB/sec
- ▶ Up to 16 per book (8 fanouts)
- ▶ Up to 48 per CEC (24 fanouts)

▶ Capacity compared to z10 EC

- ▶ z196 M80 compared to z10 EC E64 60% more capacity
- ▶ Equal “n-way” – 1.3 to 1.5 ITR ratio depending on workload
- ▶ Some workloads could gain up to 30% additional improvement if optimized to new z196 instructions and architecture

z196 Book Layout



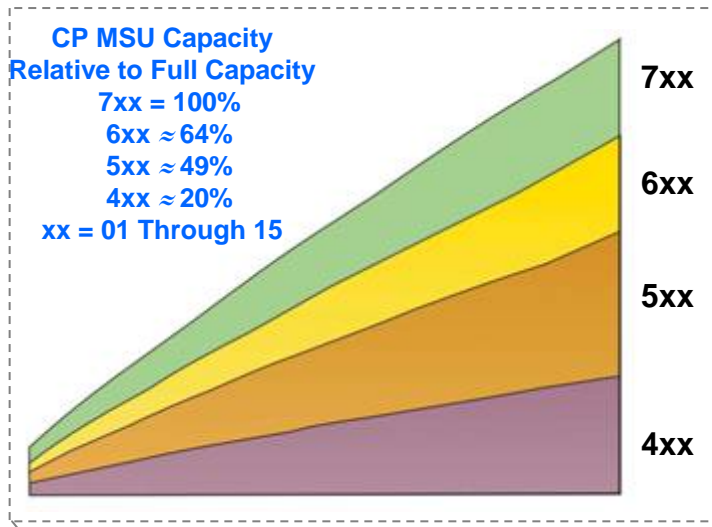
z196 Processor Features

Model	Books/ PUs	CPs	IFLs uIFLs	zAAPs	zIIPs	ICFs	SAPs Std	Optional SAPs	Std. Spares
M15	1/20	0-15	0-15 0-14	0-7	0-7	0-15	3	0-4	2
M32	2/40	0-32	0-32 0-31	0-16	0-16	0-16	6	0-10	2
M49	3/60	0-49	0-49 0-48	0-24	0-24	0-16	9	0-15	2
M66	4/80	0-66	0-66 0-65	0-33	0-33	0-16	12	0-20	2
M80	4/96	0-80	0-80 0-79	0-40	0-40	0-16	14	0-18	2

- ▶ z196 Models M15 to M66 use books each with a 20 core MCM (two 4-core and four 3-core PU chips)
- ▶ Concurrent Book Add is available to upgrade from model to model (except to the M80)
- ▶ z196 Model M80 has four books each with a 24 core MCM (six 4-core PU chips)
- ▶ **Disruptive** upgrade to z196 Model M80 is done by book replacement


- Notes:
1. At least one CP, IFL, or ICF must be purchased in every machine
 2. One zAAP **and** one zIIP may be purchased for each CP purchased even if CP capacity is “banked”.
 3. “uFL” stands for Unassigned IFL

z196 Full and Sub-Capacity CP Offerings



- **Subcapacity CPs, up to 15, may be ordered on ANY z196 model. If 16 or more CPs are ordered all must be full 7xx capacity**
- **All CPs on a z196 CEC must be the same capacity**
- **All specialty engines run at full capacity. The one for one entitlement to purchase one zAAP and one zIIP for each CP purchased is the same for CPs of any capacity.**
- **Only 15 CPs can have granular capacity but other PU cores may be characterized as full capacity specialty engines**

MSU Sub Capacity




z196 Redundant Array of Independent Memory (RAIM)

- **System z10 EC memory design:**
 - Four Memory Controllers (MCUs) organized in two pairs, each MCU with *four* channels
 - DIMM technology is Nova x4, 16 to 48 DIMMs per book, plugged in groups of 8
 - 8 DIMMs (4 or 8 GB) per feature – 32 or 64 GB physical memory per feature
Equals 32 or 64 GB for HSA and customer purchase per feature
 - 64 to 384 GB physical memory per book = *64 to 384 GB for use (HSA and customer)*

- **z196 memory design:**
 - Three MCUs, each with *five* channels. The fifth channel in each z196 MCU is required to implement memory as a *Redundant Array of Independent Memory (RAIM)*. *This technology adds significant error detection and correction capabilities. Bit, lane, DRAM, DIMM, socket, and complete memory channel failures can be detected and corrected, including many types of multiple failures.*
 - DIMM technology is SuperNova x81, 10 to 30 DIMMs per book, plugged in groups of 5
5 DIMMs (4, 16 or 32 GB) per feature – 20, 80 or 160 GB physical RAIM per feature
Equals 16, 64 or 128 GB for use per feature. *RAIM takes 20%. (There is no non-RAIM option.)*
 - 40 to 960 GB RAIM memory per book = *32 to 768 GB of memory for use*
(Minimum RAIM for the M15 is 60 GB = 48 GB = 16 GB HSA plus 32 GB customer memory)

- **For both z196 and z10**
 - *The Hardware System Area (HSA) is 16 GB fixed, outside customer memory*
 - *In some cases, offering granularity can prevent purchase of all available memory in a book*

z196 RAIM Memory Controller Overview



2- Deep Cascade
Using Quad High DIMMs

Layers of Memory Recovery

ECC

- Powerful 90B/64B Reed Solomon code

DRAM Failure

- Marking technology; no half sparing needed
- 2 DRAM can be marked
- Call for replacement on third DRAM

Lane Failure

- CRC with Retry
- Data – lane sparing
- CLK – RAIM with lane sparing

DIMM Failure (discrete components, VTT Reg.)

- CRC with Retry
- Data – lane sparing
- CLK – RAIM with lane sparing

DIMM Controller ASIC Failure

- RAIM Recovery

Channel Failure

- RAIM Recovery

z196 Purchase Memory Offerings

Model	Standard Memory GB	Flexible Memory GB
M15	32 - 704	NA
M32	32 - 1520	32 - 704
M49	32 - 2288	32 - 1520
M66	32 - 3056	32 - 2288
M80	32 - 3056	32 - 2288

- **Purchase Memory** - Memory available for assignment to LPARs
- **Hardware System Area** – Standard 16 GB outside customer memory for system use
- **Standard Memory** - Provides minimum physical memory required to hold base purchase memory plus 16 GB HSA
- **Flexible Memory** - Provides additional physical memory needed to support activation base customer memory and HSA on a multiple book z196 with one book out of service.
- **Plan Ahead Memory** – Provides additional physical memory needed for a concurrent upgrade (LIC CC change only) to a preplanned target customer memory

z196 Standard and Flexible Purchase Memory Offerings

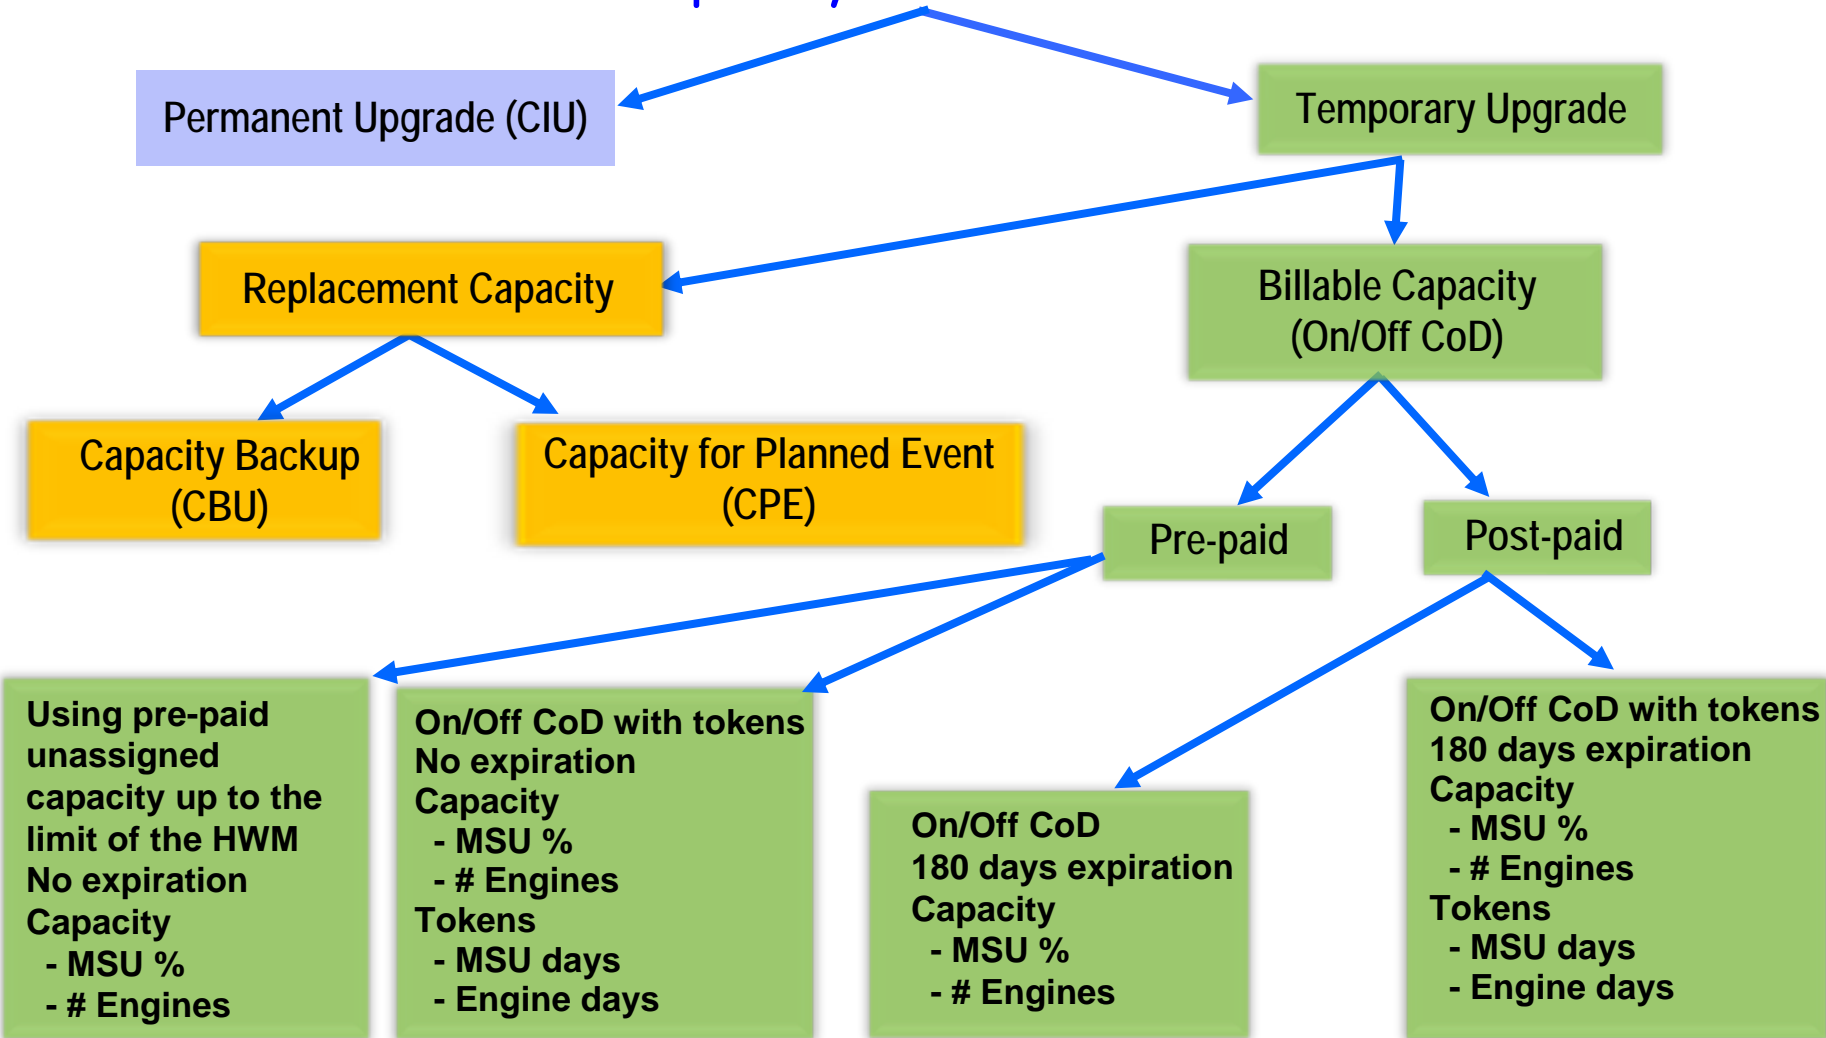
Increment	GB, Notes	Growth %	Increment	GB, Notes	Growth %	Increment	GB, Notes	Growth %
32 GB	32	100%	96 GB	608	16%	256 GB	1776	17%
	64	50%		704 ¹	13%		2032	13%
	96	33%		800	12%		2288 ³	11%
	128	25%		896	12%		2544	10%
	160	20%					2800	9%
	192	17%	112 GB	1008	13%		3056 ⁴	NA
	224	14%						
	256	25%						
						<p>Notes – Memory Maximums:</p> <p>1. M15 Standard, M32 Flexible = 704</p> <p>2. M32 Standard, M49 Flexible, (z10 EC Standard)= 1520</p> <p>3. M49 Standard, M66 and M80 Flexible = 2288</p> <p>4. M66 and M80 Standard = 3056</p>		
64 GB	320	20%	128 GB	1136	11%			
	384	17%		1264	10%			
	448	14%		1392	9%			
	512	19%		1520 ²	17%			

zEnterprise z196 On Demand Capabilities



z196 Basics of Capacity on Demand

Capacity on Demand



z196 Capacity on Demand Exclusives


- **On/Off Capacity on Demand (On/Off CoD) Improvements**
 - **Auto Replenishment of On/Off CoD records**
 - Automatic extension of the expiration date of installed On/Off CoD records
 - **On/Off CoD Administrative Test**
 - On/Off CoD “zero capacity” records for testing and training
- **CIU Improvement**
 - **Purchase of permanent unassigned engines**
 - On permanent engine purchased, Resource Link will provide the customer with the options to adjust the active capacity mark
- **Pre Installed Capacity Backup and Capacity for Planned Event Records**
 - Systems manufactured with records installed instead of staged to the Support Element
 - Limitation: Only up to 4 records can be installed
(Note: This limitation will effect very few orders.)
- **Miscellaneous Panel Enhancements**
 - Add Capacity Level Increase (CLI) to panels aka. “Speed Steps”
 - Warning for “Last Real Activation” for CBU
 - Peak Usage Marks for processor & MSU tokens (“consumption rate

zEnterprise z196 Cryptographic Capabilities



z196 Crypto Express3 2-P (Introduced z10 EC GA3)

- Earlier cryptographic features not supported
- Supported: 0, 2, 3 – 8 features = 0, 4, 6 – 16 cryptographic engines.
Each can be individually configured as Coprocessor or Accelerator.




z196 New and exclusive cryptographic capabilities

- **Elliptic Curve Cryptography Digital Signature Algorithm**, an emerging public key algorithm expected eventually to replace RSA cryptography in many applications. ECC is capable of providing digital signature functions and key agreement functions. The new CCA functions provide ECC key generation and key management and provide digital signature generation and verification functions compliance with the ECDSA method described in **ANSI X9.62 "Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Algorithm (ECDSA) "**. ECC uses keys that are shorter than RSA keys for equivalent strength-per-key-bit; RSA is impractical at key lengths with strength-per-key-bit equivalent to AES-192 and AES-256. So the strength-per-key-bit is substantially greater in an algorithm that uses elliptic curves.
- **ANSI X9.8 PIN security** which facilitates compliance with the processing requirements defined in the new version of the **ANSI X9.8 and ISO 9564 PIN Security Standards** and provides added security for transactions that require Personal Identification Numbers (PIN).
- **Enhanced Common Cryptographic Architecture (CCA)**, a Common Cryptographic Architecture (CCA) key token wrapping method using Cipher Block Chaining (CBC) mode in combination with other techniques to satisfy the key bundle compliance requirements in standards including **ANSI X9.24-1** and the recently published **Payment Card Industry Hardware Security Module (PCI HSM) standard**.
- **Secure Keyed-Hash Message Authentication Code (HMAC)**, a method for computing a message authentication code using a secret key and a secure hash function. It is defined in the standard **FIPS 198, "The Keyed-Hash Message Authentication Code "**. The new CCA functions support HMAC using SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 hash algorithms. The HMAC keys are variable-length and are securely encrypted so that their values are protected.
- **Modulus Exponent (ME) and Chinese Remainder Theorem (CRT)**, RSA encryption and decryption with key lengths greater than 2048-bits and up to 4096-bits.

zEnterprise z196 I/O Structure



z196 Book Layout



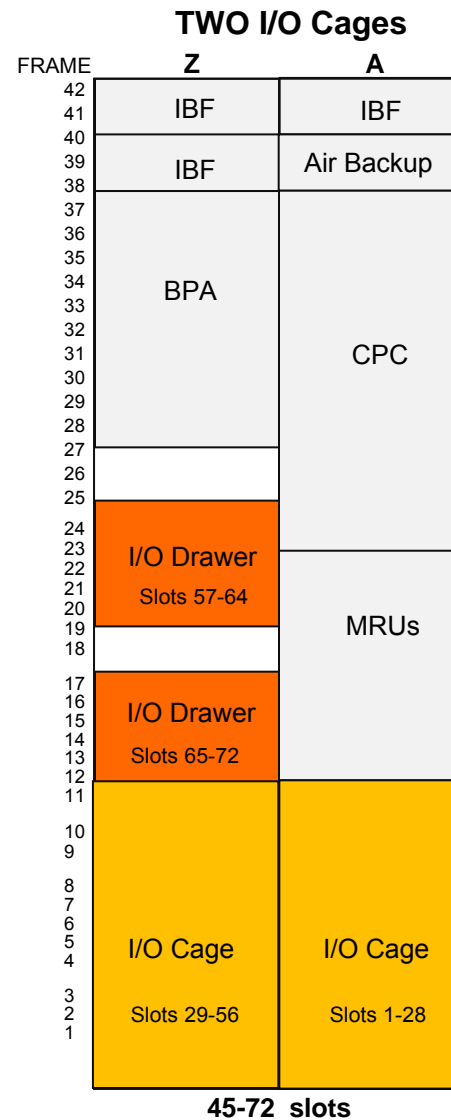
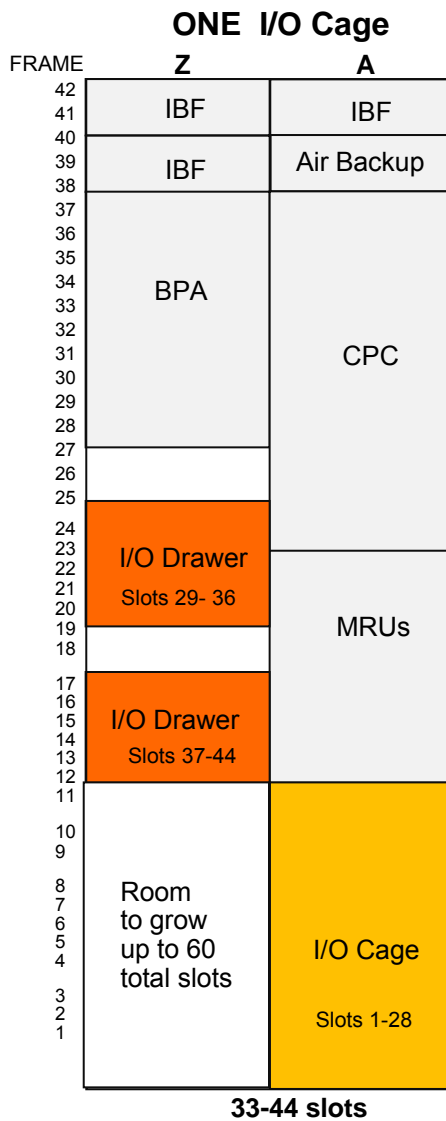
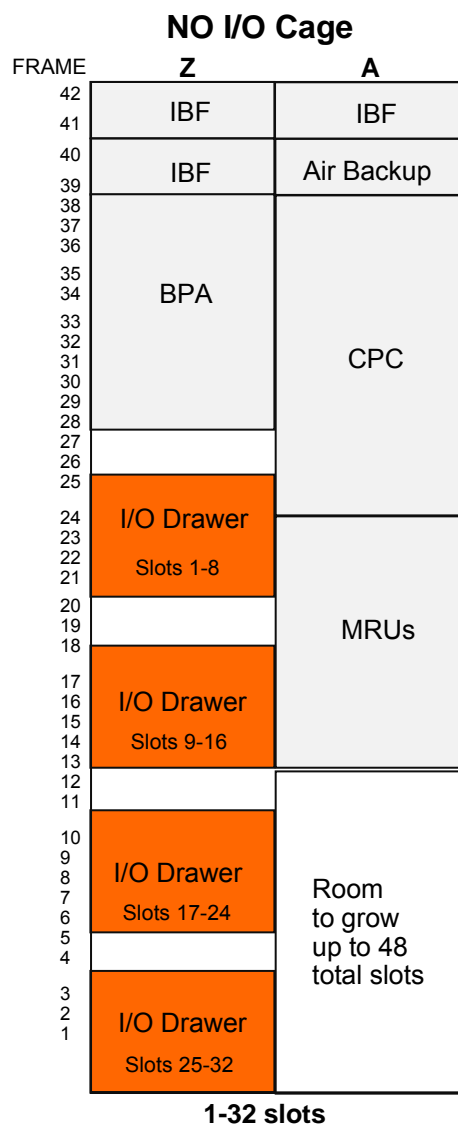
Storage Connectivity Has Gotten Easier and Performance Better

Designed, developed and tested together is key to unlocking value

- Simplified configuration of FICON® disk and tape with z/OS discovery and auto-configuration (zDAC)
- zHPF enhancements allows for increased exploitation transparently to applications and middleware
- Introduction of hot pluggable I/O drawer
- Extending for storage growth with new three subchannel sets per LCSS



z196 I/O Drawer and Cage Configurations (New Build)



Note: Drawings Not to Scale

SHARE 116 in Anaheim, March 1, 2011

z196 Channel Type and Crypto Overview

- **I/O Channels**
 - **FICON Express8**
 - *FICON Express4 (CF only on type upgrade)*
 - **ESCON – (240 or fewer channels)**
- **OSA-Express (Up to 24 features)**
 - **OSA-Express3**
 - **10 Gigabit Ethernet LR and SR**
Intraensemble data network (IEDN) requires two 10 GbE CHPIDs (LR or SR) on two different feature cards. OSX CHPID type.
 - **Gigabit Ethernet LX and SX**
 - **1000BASE-T Ethernet**
Intranode Management Network (INMN) requires two 1000BASE-T CHPIDs on two different feature cards. OSM CHPID type.
 - *OSA-Express2 (CF only on type upgrade)*
 - *1000BASE-T Ethernet*
 - *Gigabit Ethernet LX and SX*
- **HiperSockets (Define only, no additional charge)**
 - **Up to 32 (was 16)**
- **Coupling Links**
 - **Up to 80 external coupling ports (was 64)**
 - **Up to 128 CHPIDs (was 64)**
 - **InfiniBand Coupling Links (Up to 32)**
 - 12x InfiniBand DDR
 - 1x InfiniBand DDR
 - **ISC-3 (Up to 48, Peer mode only)**
 - **IC (Define only, no additional charge)**
- **Crypto**
 - **Crypto Express3 (Up to 8 features)**
 - **New function**
- **Not supported:**
 - **More than 240 ESCON channels**
 - **More than 72 I/O feature cards**
 - **FICON (before FICON Express4)**
 - **FCV – ESCD Model 5 Bridge Card**
 - **OSA-Express2 10 GbE LR**
 - **OSA-Express (pre OSA-Express2)**
 - **ICB-4 and earlier ICB**
 - **Crypto Express2 and earlier**
 - **Sysplex Timer (ETR) Attachment**

Bold – available on new build *CF – carry forward on upgrade of a z9 EC or z10 BC*

New ESCON Statement of Direction* - February 15, 2011

- **The IBM zEnterprise 196 (z196) will be the last high-end server to support ESCON channels:** IBM plans not to offer ESCON channels as an orderable feature on high-end System z servers which follow the z196 (machine type 2817). In addition, ESCON channels **cannot be carried forward** on an upgrade to such a follow-on server. This plan applies to channel path identifier (CHPID) types CNC, CTC, CVC, and CBY and to feature code numbers 2323 and 2324. System z customers should continue migrating from ESCON to FICON. Alternate solutions are available for connectivity to ESCON devices. IBM Global Technology Services offers an ESCON to FICON Migration solution, Offering ID #6948-97D, to help facilitate migration from ESCON to FICON. This offering should help customers to simplify and manage a single physical and operational environment.
- Notes:
 - This new Statement of Direction supersedes the previous ESCON SOD in Announcement letter 110-170 of July 22, 2010. It also confirms the SOD in Announcement letter 109-230 of April 28, 2009 that “ESCON Channels will be phased out.”
 - This SOD does **NOT** say that the z10 BC will be the last midrange server to support ESCON channels or the last to offer ESCON channels as an orderable feature.

*All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these statements of general direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Previous I/O Statements of Direction


- ★ **The z196 is planned to be the last high end System z server to support FICON Express4 and OSA-Express2.** Clients are advised to begin migration to FICON Express8 and OSA-Express3.
 - ★ **The z196 is planned to be the last high end System z server on which ~~ESCON channels~~, ISC-3 links, and Power Sequence Control features can be ordered. Only when an installed server with those features is field upgraded to the next high System z server will they be carried forward.** Clients are advised to begin migration to FICON Express8, InfiniBand links, and alternate means of powering control units on and off.
 - **It is IBM's intent for ESCON channels to be phased out.** System z10 EC and System z10 BC will be the last servers to support more than 240 ESCON channels
 - **The System z10 will be the last server to support connections to the Sysplex Timer (9037).** Servers that require time synchronization, such as to support a base or Parallel Sysplex, will require Server Time Protocol (STP). STP has been available since January 2007 and is offered on the System z10, System z9, and zSeries 990 and 890 servers.
 - **ICB-4 links to be phased out.** IBM intends to not offer Integrated Cluster Bus-4 (ICB-4) links on future servers. IBM intends for System z10 to be the last server to support ICB-4 links as originally stated in Hardware Announcement 108-154, dated February 26, 2008
 - **The System z10 will be the last server to support Dynamic ICF expansion.** This is consistent with the Statement of Direction in Hardware Announcement 107-190, dated April 18, 2007: "IBM intends to remove the Dynamic ICF expansion function from future System z servers."
- ★ All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

zEnterprise z196 Storage Connectivity




z196 FICON Express8

- **Auto-negotiate to 2, 4, or 8 Gbps**
1 Gbps devices not supported point to point
- **Connector - LC Duplex**
- **Four LX ports (FC #3325)**
 - 9 micron single mode fiber
 - Unrepeated distance - 10 km (6.2 miles)
 - Receiving device must also be LX
- **Four SX ports (FC #3326)**
 - 50 or 62.5 micron multimode fiber (50 micron fiber is preferred)
 - Unrepeated distance varies fiber type and link data rate
 - Receiving device must also be SX
- **LX and SX performance is identical**
- **Additional buffer credits supplied by a director or DWDM are required to sustain performance beyond 10 km**



3325 – 10KM LX, # 3326 – SX



Small Form Factor Pluggable (SFP) optics.
Concurrent repair/replace action for each SFP

z196 zHPF supports data transfers larger than 64 k bytes

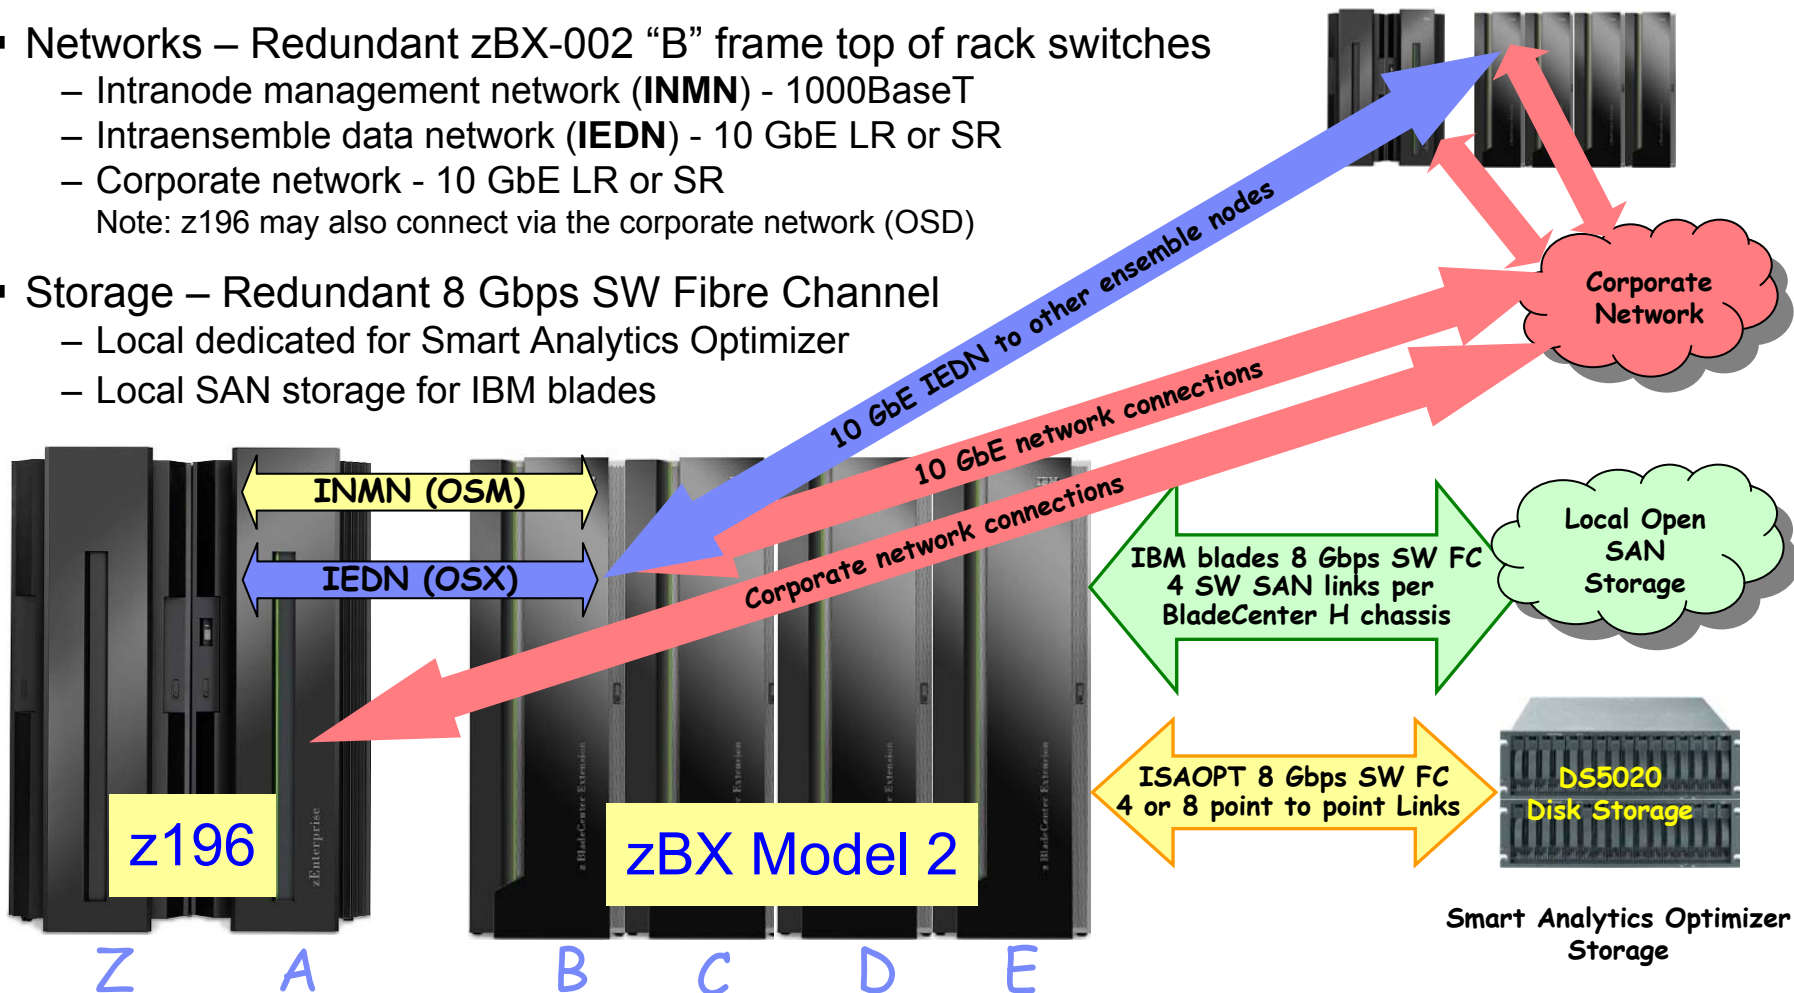
- **zHPF multi-track data transfers are no longer limited to 64 k bytes**
 - Up to 256 tracks can be transferred a single operation
 - Eliminating the 64 k byte limit is designed to allow a FICON Express8 channel to fully exploit its available bandwidth
 - This enhancement is exclusive to z196
- **Designed to help provide**
 - Higher throughput for zHPF multi-track operations
 - With lower response time
- **Requires:**
 - FICON Express8 or FICON Express4 channel
 - CHPID TYPE=FC definition
 - Control unit support for zHPF (no change to previous zHPF support)
- **z/OS operating system support**

zEnterprise z196 Network Connectivity



zEnterprise z196 and zBX Model 2 Ensemble Connectivity


- Networks – Redundant zBX-002 “B” frame top of rack switches
 - Intranode management network (**INMN**) - 1000BaseT
 - Intraensemble data network (**IEDN**) - 10 GbE LR or SR
 - Corporate network - 10 GbE LR or SR
 - Note: z196 may also connect via the corporate network (OSD)
- Storage – Redundant 8 Gbps SW Fibre Channel
 - Local dedicated for Smart Analytics Optimizer
 - Local SAN storage for IBM blades




z196 OSA-Express3 (fiber optic)

- **Double the port density of OSA-Express2**
- **Reduced latency & improved throughput**
 - Ethernet hardware data router
- **Improved throughput – standard & jumbo frames**
 - New microprocessor
 - New PCI adapter
- **CHPID types**
 - 10 Gigabit Ethernet – OSD TCP/IP or **OSX for intraensemble data network**
Ensemble requires two OSX CHPIDs on two separate feature cards of the same type
 - Ensemble requires two OSX
 - Gigabit Ethernet – OSD TCP/IP or and OSN for the communication controller for Linux
- **Port usage in 2-port Gigabit Ethernet CHPIDs**
 - OSD both with operating system support
 - OSN does not use any ports

	OSA-Express2	OSA-Express3
Microprocessor	500 MHz – 10 GbE 448 MHz – 1 GbE	667 MHz
PCI bus	PCI-X	PCIe G1




10 GbE LR #3370,
10 GbE SR #3371



CHPID shared by two ports
GbE LX #3362, GbE SX #3363

z196 OSA-Express3 1000BaseT

- Auto-negotiation to 10, 100, 1000 Mbps
- Double the port density of OSA-Express2
- Reduced latency & improved throughput
 - Ethernet hardware data router
- Improved throughput – standard & jumbo frames
 - New microprocessor
 - New PCI adapter
- Port usage in 2-port CHPIDs
 - OSC, OSD, OSE both
 - **OSM port 0 only**
 - OSN does not use ports
- **Ensemble requires two OSM CHPIDs on two different feature cards**



CHPID shared by two ports
1000BaseT # 3367

	OSA-Express2	OSA-Express3
Microprocessor	448 MHz	667 MHz
PCI bus	PCI-X	PCIe G1

Mode	TYPE	Description
OSA-ICC	OSC	TN3270E, non-SNA DFT, OS system console operations
QDIO	OSD	TCP/IP traffic when Layer 3, Protocol-independent when Layer 2
Non-QDIO	OSE	TCP/IP and/or SNA/APPN/HPR traffic
Unified Resource Manager	OSM	Connectivity to intranode management network (INMN)
OSA for NCP (LP-to-LP)	OSN	NCPs running under IBM Communication Controller for Linux (CDLC)

zEnterprise z196 Parallel Sysplex




z196 Coupling Links

- **ICB-4 – No longer supported**
- **ETR – No longer supported**
- Fanout, not I/O slot, is used for PSIFB
- All coupling links support STP
- Sysplex Coexistence – z10 EC and BC and z9 EC and BC only




z196



System z CFCC Level 17

- **CFCC Level 17 allows:**
 - Up to 2047 CF structures (CFCC 16 allowed 1024).
Allowing more CF structures to be defined and used in a sysplex permits more discrete data sharing groups to operate concurrently, and can help environments requiring many structures to be defined, such as to support SAP or service providers
 - Improved CFCC diagnostics & Link Diagnostics
- **Structure and CF Storage Sizing with CFCC level 17**
 - May increase storage requirements when moving from CFCC Level 16 (or below) to CF Level 17
 - Using the **CFSizer** Tool is recommended
 - <http://www.ibm.com/systems/z/cfsizer/>
- **Greater than 1024 CF Structures requires a new version of the CFRM CDS**
 - All systems in the sysplex must to be at z/OS V1.12 or have the coexistence/preconditioning PTF installed.
 - Falling back to a previous level (without coexistence PTF installed) is NOT supported without sysplex IPL



zEnterprise z196 Physical Planning




z196 optional water cooling



Infrastructure

Not to scale, not physically representative



- A Smarter IT for a Smarter Planet™
- Each book has a water cooled cold plate for the processor MCM
- Water Cooling Unit (WCU) design is N+1 with independent chilled water connections
 - One WCU can support system without cycle steering
 - Connects to ordinary building chilled water (like AC units and unlike water cooled rear doors)
- Rear Door Exhaust Air Heat Exchanger (XAHX)
 - Removes heat from exhaust air at back of both frames
 - Provides an air cooling back-up mode for robustness
- Designed to reduce the heat load exhausted to air by 60-65%
 - ~10 kW system heat load to air maximum (5 kW per frame)
 - ~2 kW Input energy savings for a maximum power system
 - ~2.5 kW additional power savings to cool the reduced air heat load

The water cooling option must be ordered with a new build or machine type upgrade. It is not available as a z196 MES change after installation.


z196 optional high voltage DC power

- A Smarter IT System for a Smarter Planet
- Using high voltage DC power can save, on average, 1 to 3% of power by eliminating DC to AC and AC to DC conversion losses




Infrastructure

AC vs DC Distribution



System z196 with optional water cooling and Overhead I/O Dimension changes compared to the z10 EC

- Depth: Water Cooled option adds 4 inches to the rear (with reference to floor cutouts)
- Width: Overhead I/O Option adds 11 -12 inches side to side
5.5 – 6 inches to the outside edge of the A and Z frames (with reference to floor cutouts)
- Height: Overhead I/O Option adds 5.5 – 6 inches (Reduced height shipping to 71 inches available)
- Weight: Overhead I/O Option adds ~ 200 pounds, Water Cooled Option adds ~ 100 pounds
- z196 must be installed on a raised floor



Note: All dimensions are approximate

z196 Processor Statements of Direction

- ★ **The IBM zEnterprise 196 is planned to be the last high-end server to support dial-up modems for use with the Remote Support Facility (RSF), and the External Time Source (ETS) option of Server Time Protocol (STP).** Enterprises should begin migrating from dial-up modems to broadband for RSF connections.

- **Notes:**

1. Broadband already provides superior support to dial for these RSF and STP functions.
- 2. Use of broadband RSF significantly improves serviceability for both z196 and z10.**
See on Resource Link: **“System z Broadband RSF”**, Z121-0244-03

- **IBM intends to support optional water cooling on future high-end System z servers.** This cooling technology will tap into building chilled water that typically exists within the datacenter for computer room air conditioning systems. External chillers or special water conditioning will typically not be required. Water cooling technology for high-end System z servers will be designed to deliver improved energy efficiencies
- **IBM intends to support the ability to operate from High Voltage DC power on future System z servers.** This will be in addition to the wide range of AC power already supported. A direct HV DC datacenter power design can improve data center energy efficiency by removing the need for an additional DC to AC inversion step

- ★ **All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.**

Thank you

Questions?

The Future Runs on System z



zEnterprise z196 Backup

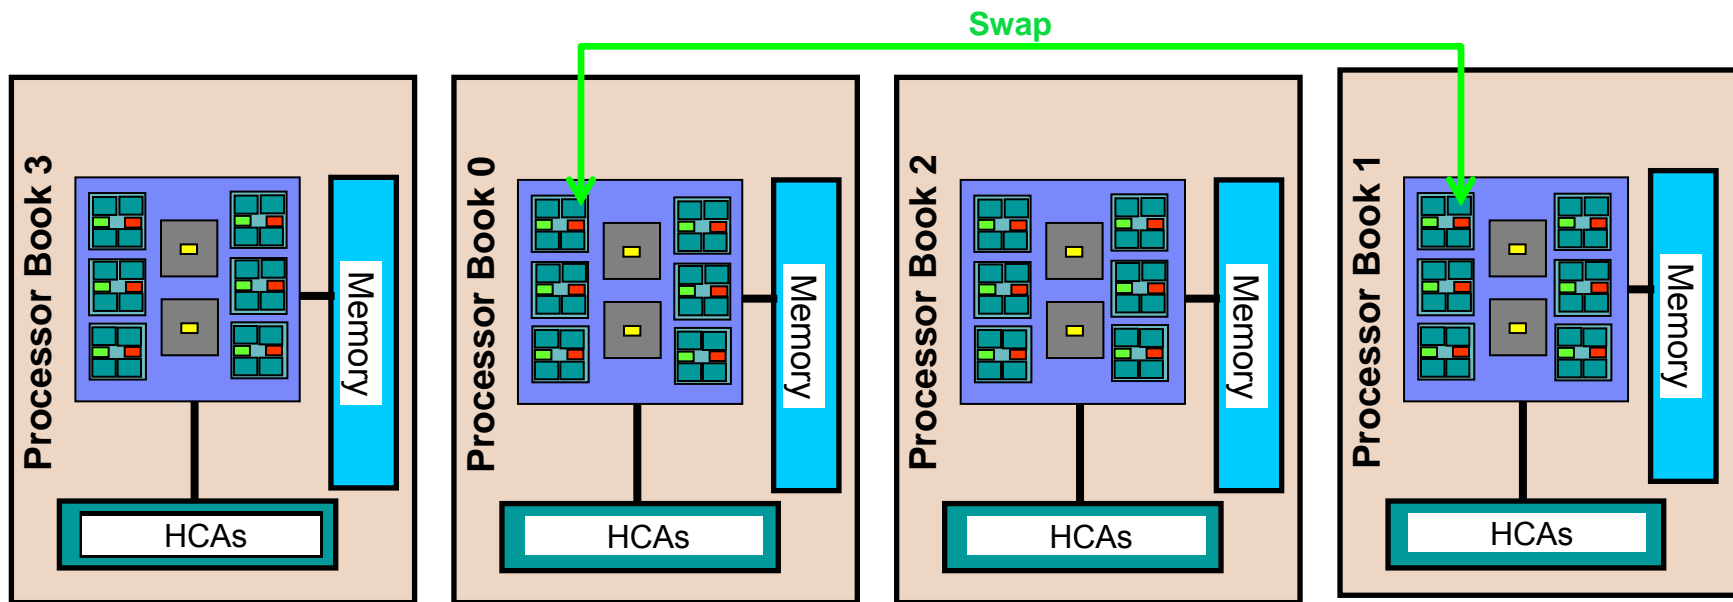


z196 Key System z Structure Fundamentals

- **Robust redundant design for critical system components (e.g. N+1 power components)**
- **Extensive error detection, correction, and recovery**
- **Dynamic processor core sparing to standard spare cores**
- **Concurrent hardware upgrade and service**
 - “Repair and Verify” support for Service Representatives
 - Concurrent Book Add (concurrent model upgrade)
 - Enhanced Book Availability (concurrent book upgrade or service on multiple book models)
 - Concurrent Licensed Internal Code (LIC) update
 - Enhanced Driver Maintenance (concurrent upgrade to a new level of LIC function)
- **Remote Support Facility (“Phone Home” Support)**
 - Problem reporting, log upload, and other service support
 - Download of LIC updates for concurrent upgrade and service
- **Capacity on demand – Permanent Upgrade, Capacity Backup, Planned Event, On/Off Capacity**
 - Concurrent processor enablement and unassignment
 - Capacity Provisioning Manager for on demand temporary capacity
- **Concurrent Logical Partition (LPAR) resource reassignments**
 - Logical processor assignment, weighting and capping
 - Logical cryptographic coprocessor assignment
 - LPAR capacity and group capacity definitions
 - Memory reconfiguration
 - Intelligent Resource Director (IRD) driven by workload management
 - HiperDispatch and LPAR Dynamic Processor Reassignment
- **z/Architecture support (e.g. Trimodal Addressing – For programs with 24, 31, or 64-bit addressing)**

z196 LPAR Dynamic PU Reassignment

- PR/SM dynamic relocation of running processors to different processor cores
- Designed to optimize physical processor location for the current LPAR logical processor configuration
- Swap an active PU with a different active PU in a different book
 - Designed Benefit: Better L3 and L4 cache reuse
 - CP, zAAP, zIIP, IFL and ICF supported
 - Triggers: Partition activation/deactivation, machine upgrades/downgrades, logical processors on/off
- Designed to provide the most benefit for:
 - Multiple book machines
 - Dedicated partitions and wide partitions with HiperDispatch active



z196 Flexible Memory

- **Provides additional physical memory needed to support activation all purchased memory and HSA on a multiple book z196 with one book out of service for**
 - Scheduled concurrent book upgrade (e.g. memory)
 - Scheduled concurrent maintenance
 - Concurrent repair of a book “fenced” during Activation (POR)
 - *Note: All of the above can be done without Flexible Memory; but, all purchased memory will not be available for use in most cases. Some work may have to be shut down or not restarted.*
- **Offered on M32, M49, M66 and M80 in:**
 - 32 GB increments from 32 GB to 256 GB
 - 64 GB increments from 320 GB to 512 GB
 - 96 GB increments from 608 GB to 896 GB (*M32 limit 704 GB*)
 - 112 GB increment to 1008 GB
 - 128 GB increments from 1136 GB to 1520 GB (*M49 limit 1520 GB*)
 - 256 GB increments from 1776 GB to 2288 GB
- **Selected by checking the “Flexible” box when configuring memory**
- **Additional physical memory, if required, is added to the configuration and priced as “Plan Ahead Memory”**

z196 Plan Ahead Memory

- **Provides the capability for concurrent memory upgrades without exploitation of Enhanced Book Availability, Licensed Internal Code upgrades**
 - Memory cards are pre-installed to support target Plan Ahead capacity
 - Available on all System z196 models
 - Can be ordered with standard memory on any z196 model
(*Standard plus Plan Ahead will NOT be Flexible in most cases.*)
 - Can be ordered with Flexible memory on a multiple book z196 model

- **Pre-planned memory features are chargeable**
 - Charge for memory hardware needed to enable the selected plan ahead target.
 - FC #1996 – One feature for each 16 GB (20 GB RAIM) of additional hardware needed

- **Pre-planned memory activation is chargeable**
 - Subsequent memory upgrade orders will use Plan Ahead Memory first
 - Charged when Plan Ahead Memory is enabled by concurrent LIC upgrade
 - Add FC #1901, Delete FC #1996 – For each 16 GB of memory activated for use

- **Note: Plan Ahead Memory is NOT temporary, On Demand memory**
Temporary memory is not offered because Memory LIC downgrade is disruptive.

z196 Manufacturing Installation of CBU and CPE Records

- **Ordered CBU and CPE records on z196 will be installed by manufacturing**
 - Saves customer or IBM Service time after machine installation
 - Prepares the records for immediate manual or automated activation
 - Exclusive to z196
 - Limitation: Only up to 4 records can be installed
(Note: This limitation will effect very few orders.)

- **On System z10**
 - Records are shipped staged on the support element
 - Records must be installed on-site prior to use
 - While installation is a simple process, it takes time
 - If not done in a timely manner, record activation for a test or response to an emergency could be delayed

z196 Key System z I/O Fundamentals

- **Robust redundant design for critical I/O components (e.g. Redundant I/O Interconnect)**
- **Concurrent add, remove and service for I/O hardware**
- **Concurrent channel path (CHPID)and device definition to enable added hardware for use**
- **Concurrent Licensed Internal Code (LIC) update for I/O features**
- **Four logical channel subsystems (LCSS) predefined**
 - 15 logical partitions predefined in each (60 total)
 - Up to 256 CHPIDs in each
 - Multiple Subchannel Sets each with 64 K subchannels for I/O operations
- **Multiple Image Facility – CHPID sharing among Logical Partitions (LPARs) in an LCSS**
- **Spanning – LCSS sharing of CHPIDs**
- **I/O operations managed by System Assist Processors (SAPs)**
- **Multiple path support - up to 8 I/O paths per I/O device for availability and performance**
 - Channel Subsystem I/O path selection
 - Extensive support for I/O error retry and recovery
- **HiperSockets memory to memory internal network connections among LPARs**
- **Extensive Fibre Channel Support**
 - FICON, System z High Performance FICON, and Fibre Channel Protocol (FCP)
 - FCP N_Port Identifier Virtualization (NPIV)
- **Parallel Sysplex clustering support**
 - Coupling Facility
 - Coupling links
 - Server time protocol

z196 SAPs, I/O Buses, Links, and I/O Connectivity

Model	Books/PU cores	Standard SAPS	Optional SAPs	Maximum I/O Fanouts/ Buses	Maximum PSIFB Links + I/O cards	Maximum I/O Cards + PSIFB Links	Max FICON/ ESCON CHPIDs
M15	1/20	3	0-4	8/16	16 + 0 Cards	56 + 0 PSIFB	224/ 240
M32	2/40	6	0-10	16/32	32 + 0 Cards	72 + 12 PSIFB	288/ 240
M49	3/60	9	0-15	20/40	32 + 32 Cards	72 + 20 PSIFB	288/ 240
M66	4/80	12	0-20	24/48	32 + 56 Cards	72 + 28 PSIFB	288/ 240
M80	4/96	14	0-18	24/48	32 + 56 Cards	72 + 28 PSIFB	288/ 240

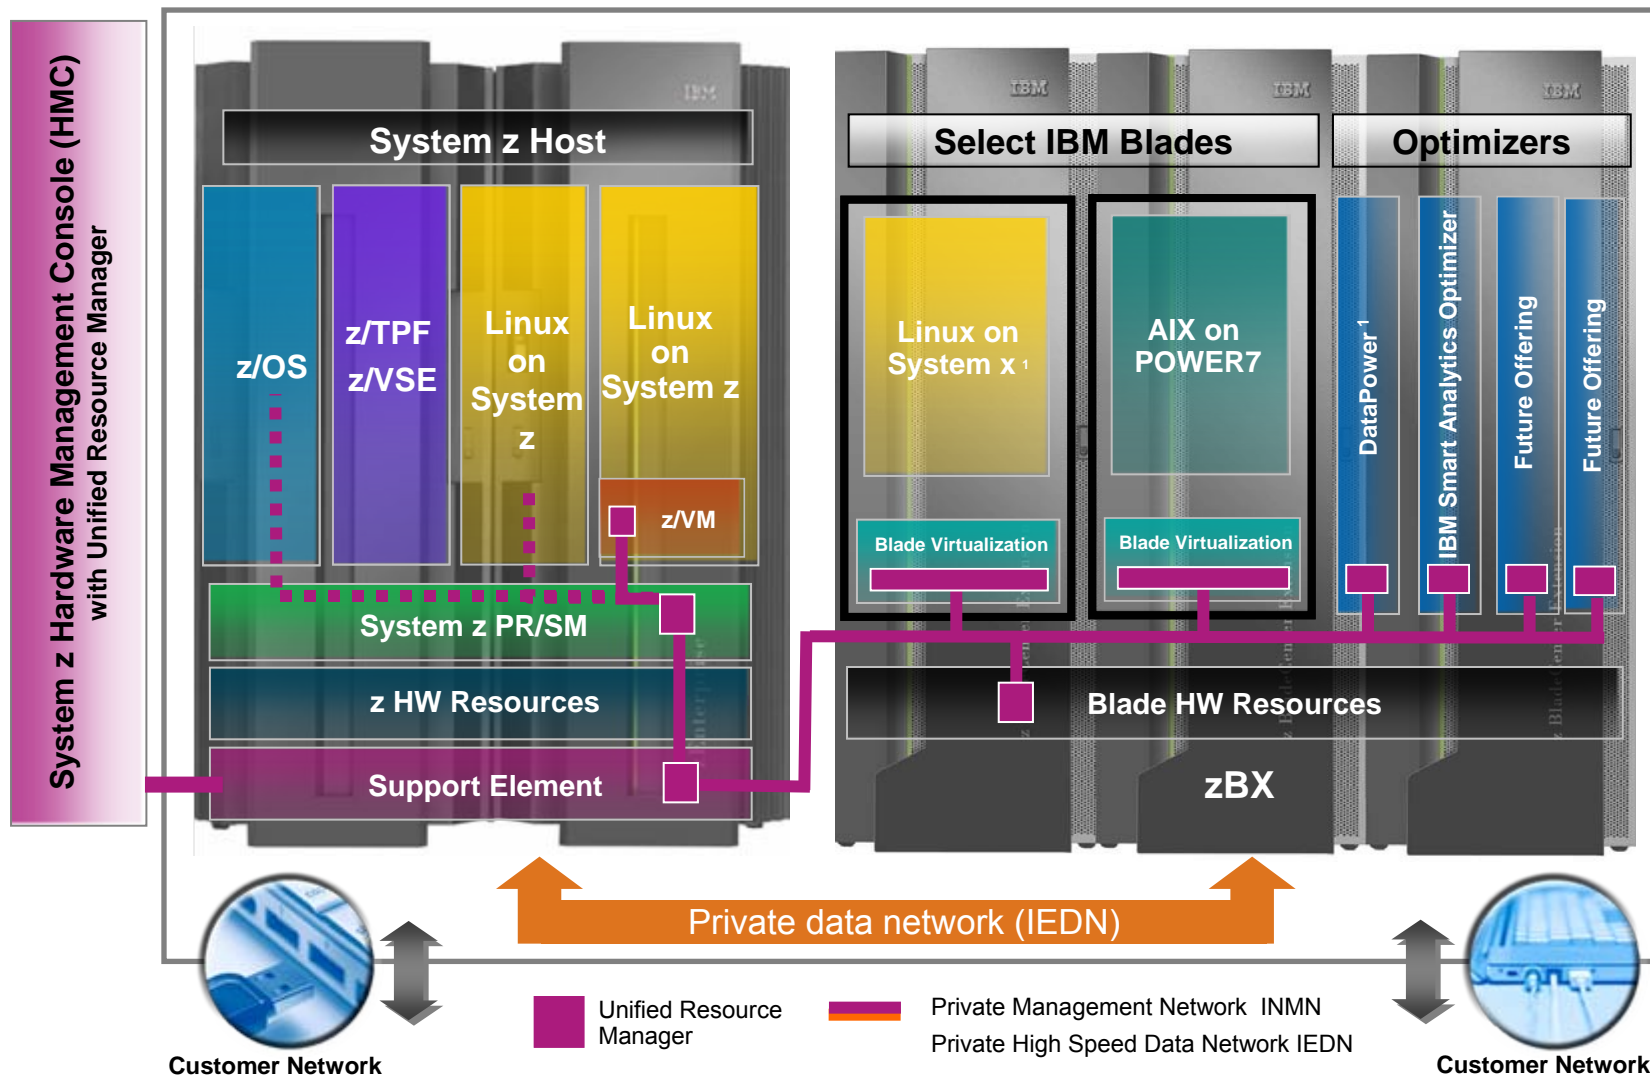
Note: Only z/TPF may need Opt SAPs for normal workload

Note: Include Crypto Express3 cards in I/O card count
Fundamental Limits:

- a. 4 LCSSs maximum
- b. 15 partitions maximum per LCSS, 60 maximum
- c. 256 CHPIDs maximum per LCSS

Putting zEnterprise System to the task

Use the smarter solution to improve your application design



¹ All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.