


IBM zEnterprise 196 (z196) Hardware Overview and Update

Session ID: 9688

Speaker: Harv Emery



IBM zEnterprise - Freedom by Design



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
- **Session: 9738 – Aug 10, 9:30, Romney White**
- **Labs: 9711, 9718 – Aug 11, 8:00 and 9:30, Hiren Shah**

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
- **Session: 9690 – Aug 9, 1:30, Greg Hutchison**
- **Technology Exchange opens today at 4:00**

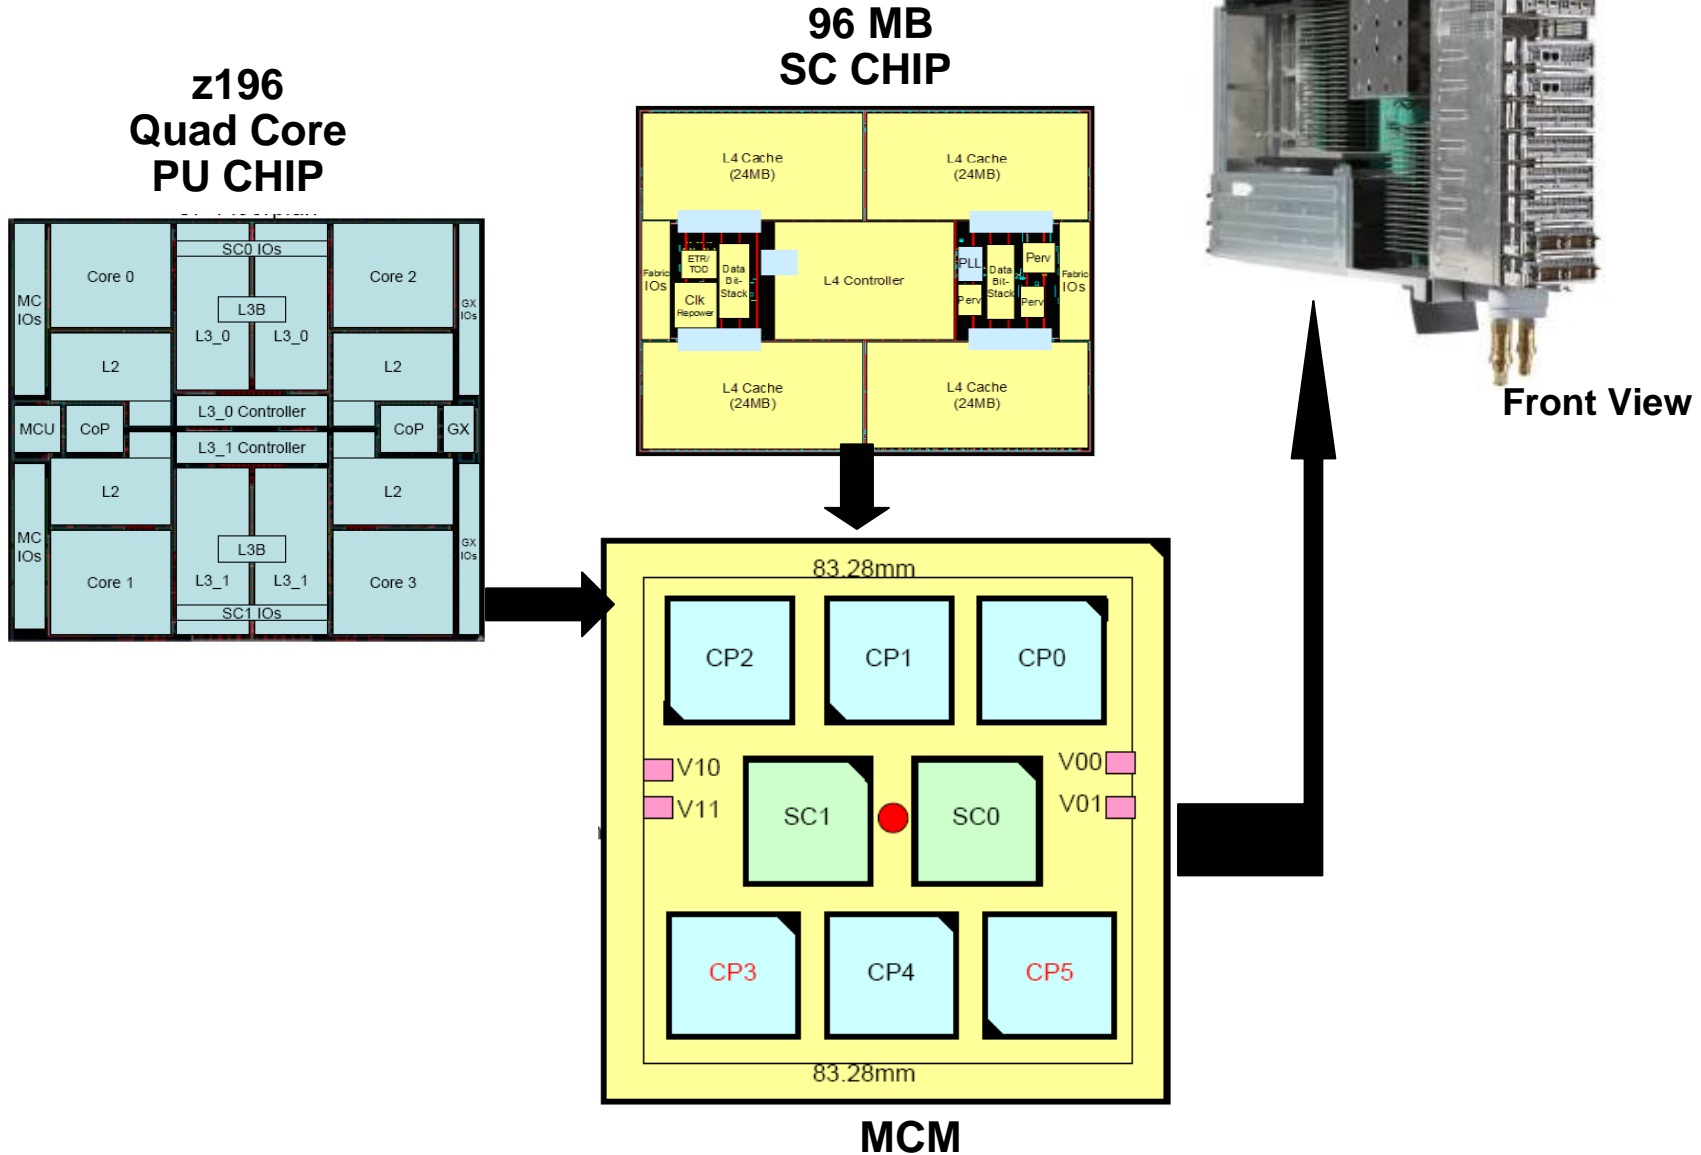
¹ All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

zEnterprise 196 Processors and Memory

IBM zEnterprise - Freedom by Design



z196 PU chip, SC chip and 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

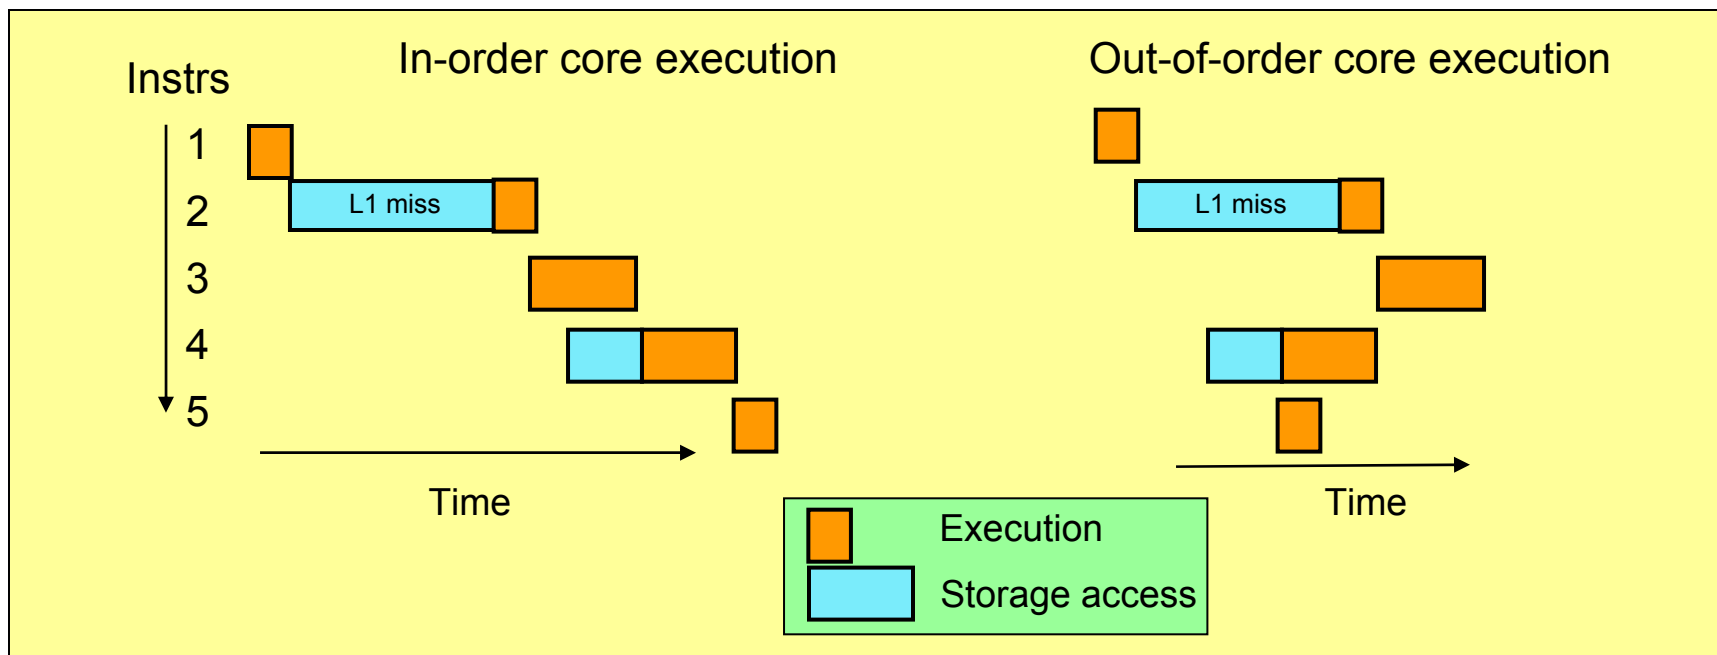
z196 Book Level Cache Hierarchy



Cache Comparison	
z196	z10
L1 – 64K/128K	L1 – 64K/128K
L2 – 1.5 M	L 1.5 – 3 M
L3 – 24 M	N/A
L4 – 192 M	L2 – 48 M

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

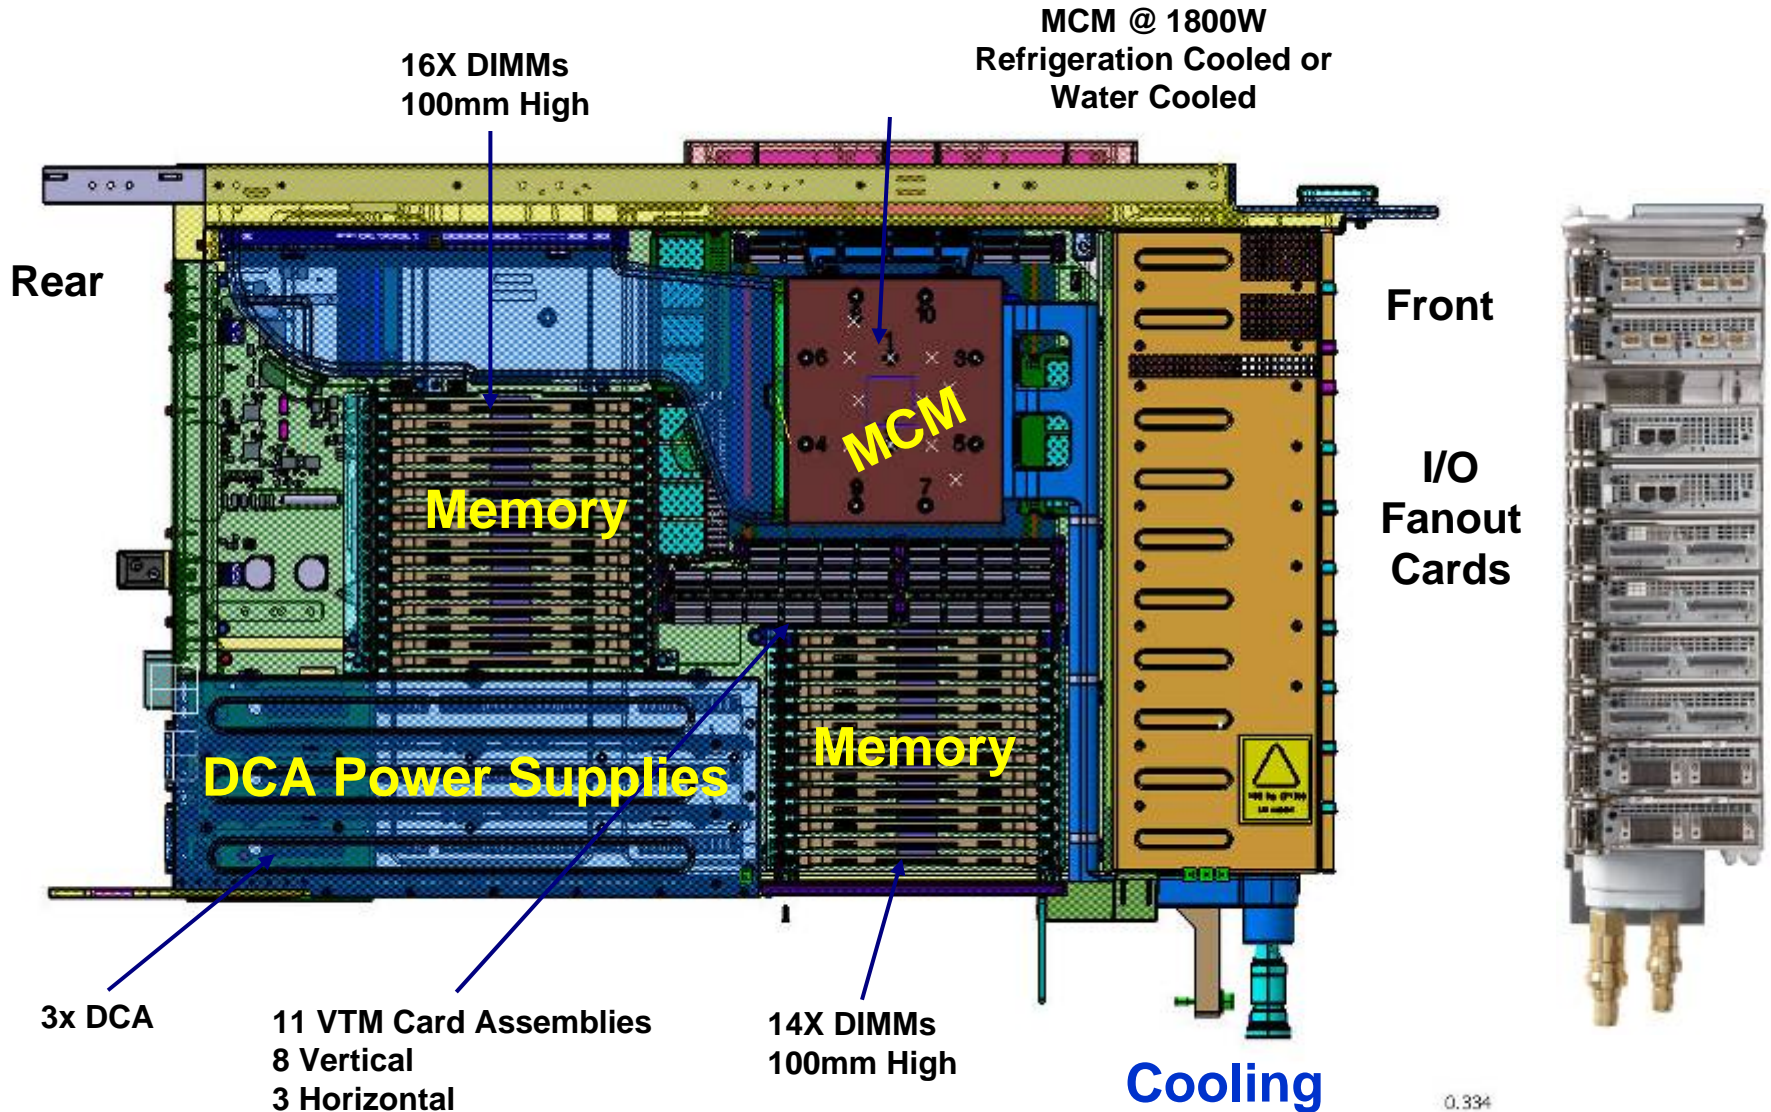
I/O Interconnects: (Same as z10 EC)

- ▶ 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



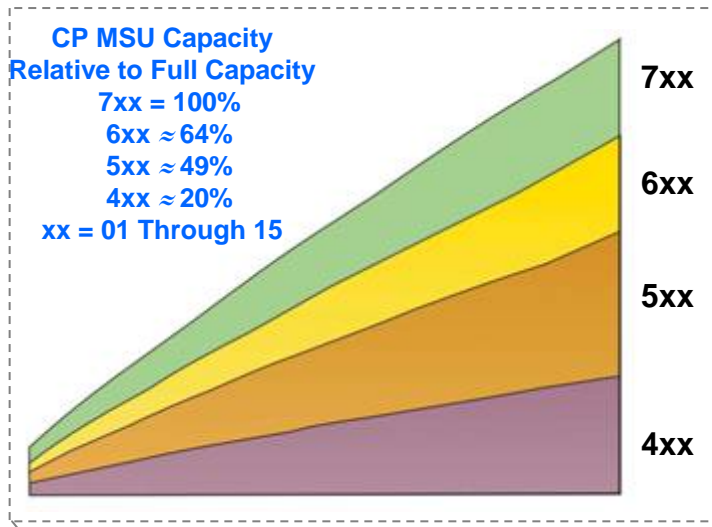
z196 Processor Features

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



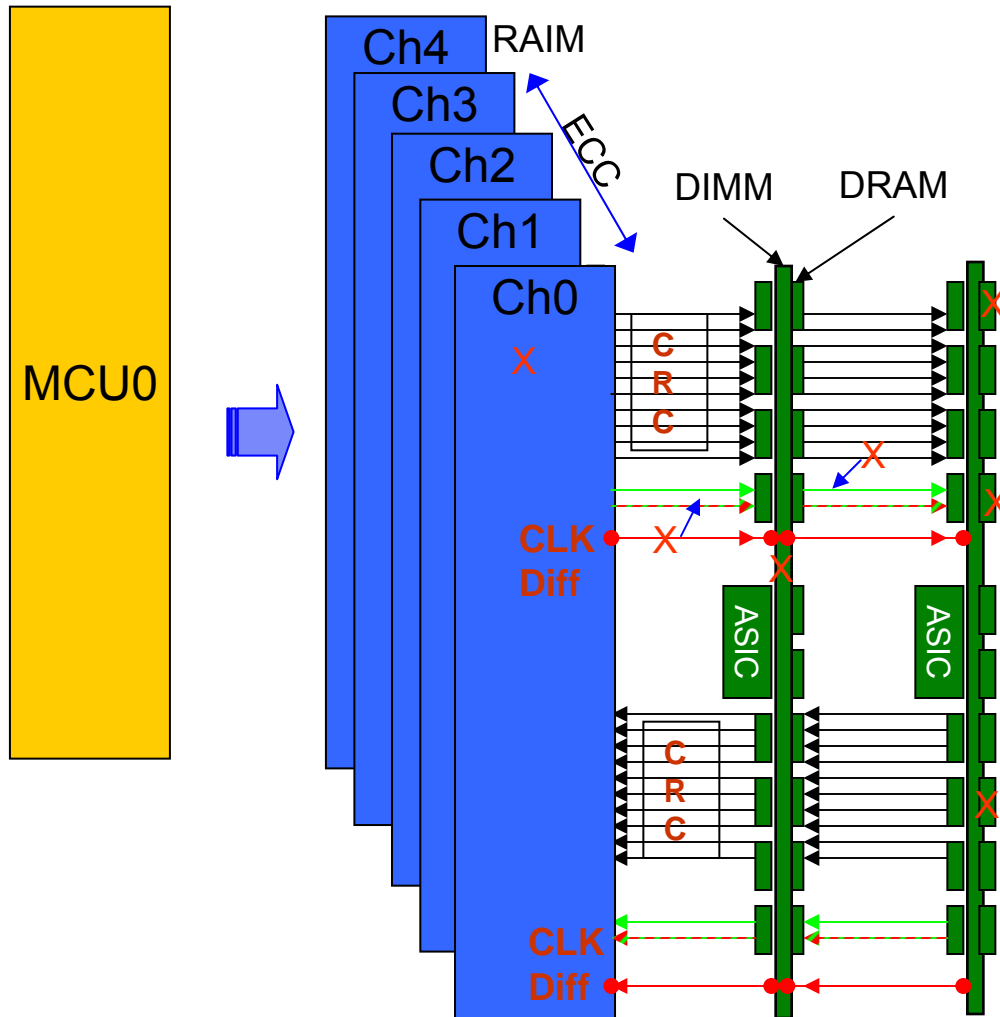
z196 Redundant Array of Independent Memory (RAIM)

- **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%			

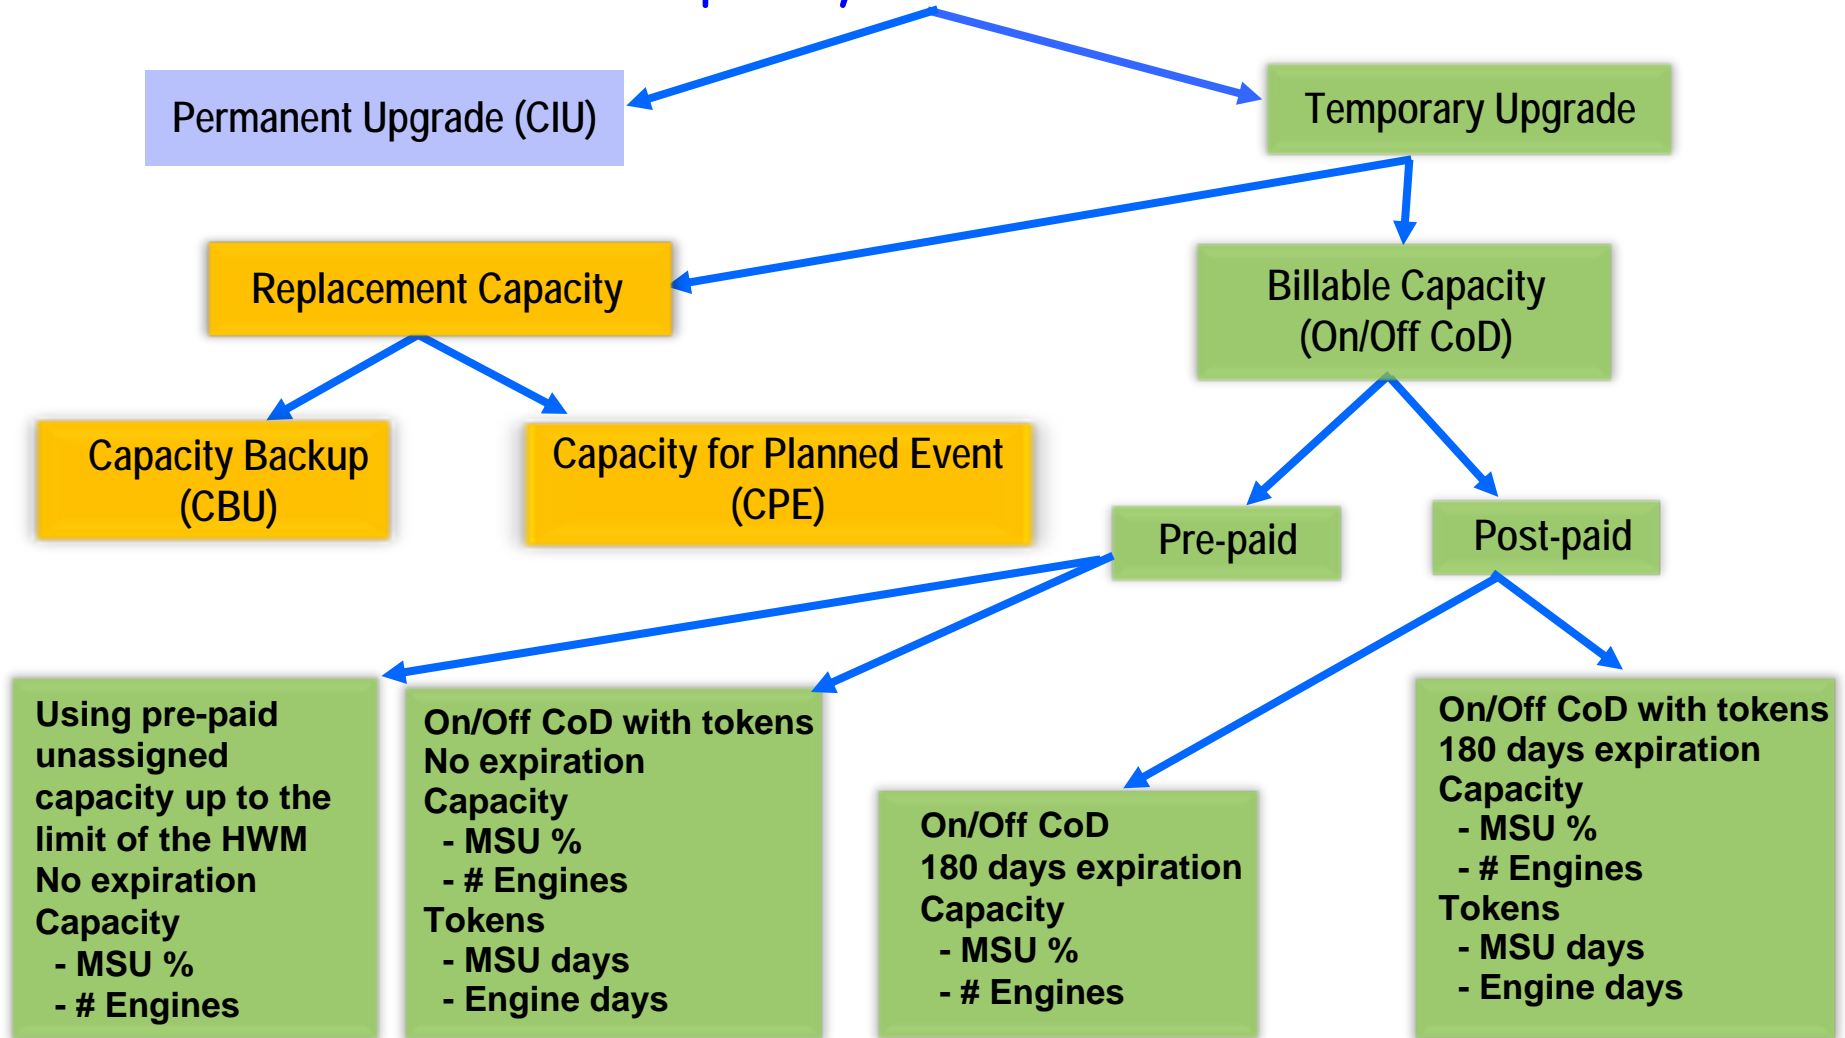
zEnterprise 196 On Demand Capabilities

IBM zEnterprise - Freedom by Design



zEnterprise Basics of Capacity on Demand

Capacity on Demand



zEnterprise 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 196 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.



Exclusive zEnterprise cryptographic capabilities (July 22, 2010)

- **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.


Summary of Cryptographic Enhancements (July 12, 2011)



- **Crypto Express3-1P (z114 only):** A cost reduced one coprocessor version of Crypto-Express3
- **Expanded key support for the AES algorithm:** Adds support for **AES Key Encrypting Keys (AES KEKs) and AES typed keys**. These AES wrapping keys have adequate strength to protect other AES keys for transport or storage. The new AES key types use the variable-length key token. The supported key types are EXPORTER, IMPORTER, and for use in the encryption and decryption services, CIPHER.
- **Enhanced ANSI TR-31 interoperable secure key exchange.** **ANSI TR-31** key wrapping defines a method of cryptographically protecting Triple Data Encryption Standard (TDES) cryptographic keys and their associated usage attributes.
- **Elliptic Curve Cryptography (EC-DH key agreement protocol):** **EC-DH key agreement protocol** allows two parties, each having an elliptic curve public-private key pair, to establish a shared secret over an insecure channel. This shared secret may be used directly as a key, or to derive another key which can then be used to encrypt subsequent communications using a symmetric key cipher such as AES.
- **PIN block decimalization table protection:** To help avoid a decimalization table attack to learn a personal identification number (PIN), a solution is now available in the CCA API to thwart this attack by protecting the decimalization table from manipulation.
- **PKA RSA OAEP with SHA-256:** **Optimal Asymmetric Encryption Padding (RSA OAEP)** is a public-key encryption scheme or method of encoding messages and data in combination with the RSA algorithm and a hash algorithm.

zEnterprise 196 I/O Structure

IBM zEnterprise - Freedom by Design



The zEnterprise 114 and zEnterprise 196 GA2 I/O Infrastructure

Session 9797: Aug 9, 4:30 PM, Southern Hemisphere I - II



New PCIe-based I/O infrastructure

- **New PCIe-based 8 GBps interconnects**
- **New PCIe I/O drawer**
 - Improved port purchase granularity (fewer ports per I/O card)
 - Increased port density compared to the previous I/O drawer or z196 I/O cage
 - Designed for improved power and bandwidth compared to previous I/O cage or z196 I/O drawer

Storage

- **New PCIe-based FICON Express8S** features

Networking


- **New PCIe-based OSA-Express4S** features

Coupling

- **New** 12x InfiniBand and 1x InfiniBand features (**HCA3-O fanouts**)
 - *12x InfiniBand - decreased service times when using 12x IFB3 protocol*
 - *1x InfiniBand – increased port count*

Note: The z114 and z196 at GA2 will ship with a new LIC Driver, Driver 93g

z196 Connectivity for I/O and Coupling





- **Up to 8 fanout cards per z196 book**
 - M15 (1 book) – up to 8
 - M32 (2 books) – up to 16
 - M49 (3 books) – up to 20
 - M66 and M80 (four books) – up to 24

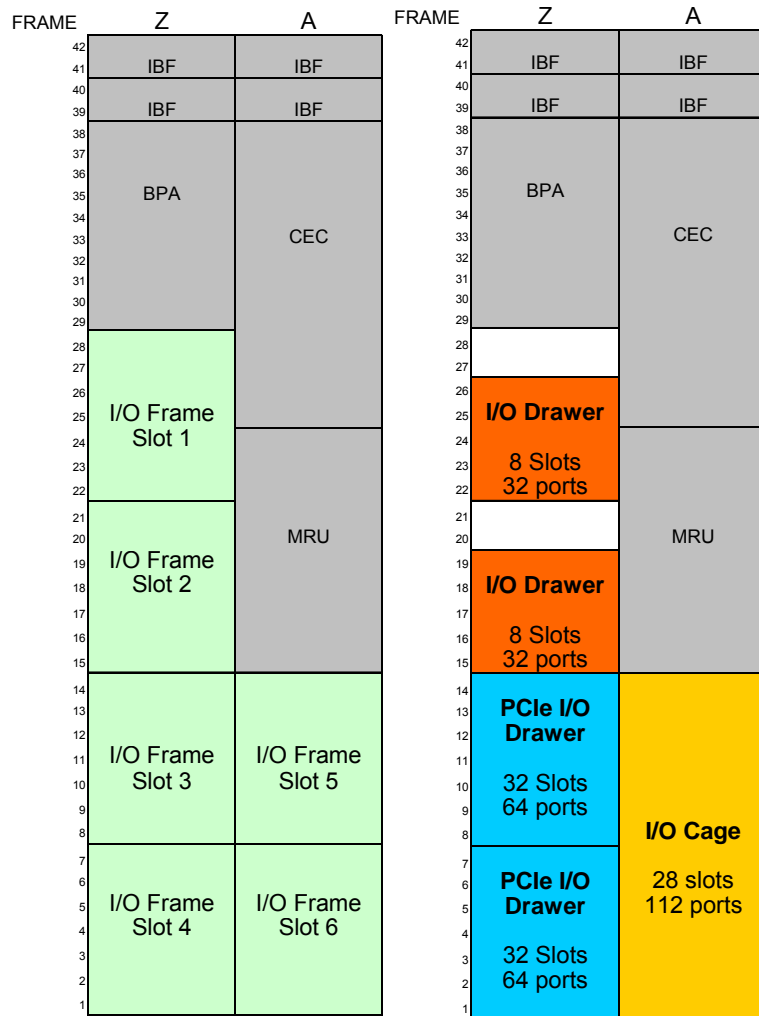
- **I/O fanouts compete for fanout slots with the the InfiniBand HCA fanouts that support coupling:**
 - HCA2-O 12x two InfiniBand DDR links
 - HCA2-O LR two 1x InfiniBand DDR links
 - **HCA3-O two 12x InfiniBand DDR links**
 - **HCA3-O LR four 1x InfiniBand DDR links**

- **PCIe fanout – PCIe I/O Interconnect links**
 Supports two copper cable PCIe 8 GBps interconnects to two 8-card PCIe I/O domain multiplexers.
Always plugged in pairs for redundancy.

- **HCA2-C fanout – InfiniBand I/O Interconnect**
 Supports two copper cable 12x InfiniBand DDR 6 GBps interconnects to two 4-card I/O domain multiplexers.
Always plugged in pairs for redundancy.



z196 Frame Layout for I/O – Air Cooled*



- An **I/O frame slot** is a physical location in the A or Z frame for an I/O cage, I/O drawer or PCIe I/O drawer to be inserted = 7u
- **I/O cage** uses 2 I/O frame slots = 14u
 - 28 four port I/O slots = 112 ports
 - 2 cages maximum (3 with RPQ)
- **PCIe I/O drawer** uses 1 I/O frame slot = 7u
 - 32 two port I/O slots = 64 ports
 - 5 drawers maximum
- **Old I/O drawer** uses 0.7 frame slot = 5u
 - 8 four port I/O slots = 32 ports
 - Requires 2u of free space for future upgrade to the PCIe I/O drawer
 - 6 drawers maximum
 - GA2: Up to 2 on new build
 - GA1: Up to 4 on new build

* Locations differ if water cooled; but the number of I/O frame slots is identical.

z196 GA2 I/O Features supported


Supported features

■ Features – PCIe I/O drawer

- **FICON Express8S**
 - SX and LX
- **OSA-Express4S**
 - 10 GbE LR and SR
 - GbE SX and LX




PCIe I/O drawer




32 I/O slots

■ Features – I/O cage and I/O drawer

- **Crypto Express3**
- **ESCON (240 or fewer)**
- FICON Express8 (*Carry forward or RPQ 8P2534 to fill empty slots*)
- FICON Express4 (*Carry forward only*)
- **ISC-3**
- **OSA-Express3 1000BASE-T**
- OSA-Express3 (*Carry forward or RPQ 8P2534 to fill empty slots*)
 - 10 GbE, GbE
- OSA-Express2 (*Carry forward only*)
 - GbE, 1000BASE-T
- **PSC (Carry forward or new build, no MES add)**



28 slot I/O cage



8 slot I/O drawer

z196 GA2 I/O Connectivity

Features	Offered As	Maximum # of features	Maximum channels	Increments per feature	Purchase increments
ESCON	NB	16	240 channels	1 - 15 active	4 channels
FICON					
FICON Express8S	NB	160	320 channels	2 channels	2 channels
FICON Express8	CF*	72	288 channels	4 channels	4 channels
FICON Express4	CF	72	288 channels	4 channels	4 channels
ISC-3	NB	12	48 links	4 links	1 link
OSA-Express					
OSA-Express4S	NB	48	96 ports	1 (10 GbE) / 2 ports	1 feature
OSA-Express3 1000BASE-T	NB	24	96 ports	4 ports	1 feature
OSA-Express3 10 GbE, GbE	CF*	24	96 ports	2 (10 GbE) / 4 ports	1 feature
OSA-Express2**	CF*	24	48 ports	2 ports	1 feature
Crypto Express3***	NB	8	16 PCIe adapters	2 PCIe adapters	1 feature ***

* Can be carried forward or ordered on MES with RPQ 8P2534 to fill empty slots

** OSA-Express2 10 GbE LR is not supported as a carry forward

*** Two features initially, one thereafter

NB = New Build

CF = Carry Forward

* 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.



Optica PRIZM FICON Converter

<http://www.opticatech.com/>

Supports the elimination of ESCON channels on the host while maintaining ESCON and Bus/Tag-based devices and applications


What is PRIZM?

- A purpose-built appliance designed exclusively for IBM System z; enables ESCON devices to be connected to FICON channels or fabrics
- Allows ESCON devices to connect to FICON channels and FICON fabrics/networks
 - Prizm also supports attachment of parallel (bus/tag) devices to FICON channels via ESBT module
- Converts 1 or 2 FICON channels (CHPID type FC) into 4, 8 or 12 ESCON channels
 - Replace aging ESCON Directors with PRIZM (maintenance savings)
 - Achieve streamlined infrastructure and reduced Total Cost of Ownership
- Qualified by the IBM Vendor Solutions Lab in POK for all ESCON devices; qualified for connectivity to Brocade and Cisco FICON switching solutions
 - Refer to: <http://www-03.ibm.com/systems/z/hardware/connectivity/index.html>
 - Products -- > FICON / FCP Connectivity -- > Other supported devices
- PRIZM is available via IBM Global Technology Services: ESCON to FICON Migration offering (#6948-97D)



zEnterprise 196 Storage Connectivity

IBM zEnterprise - Freedom by Design



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



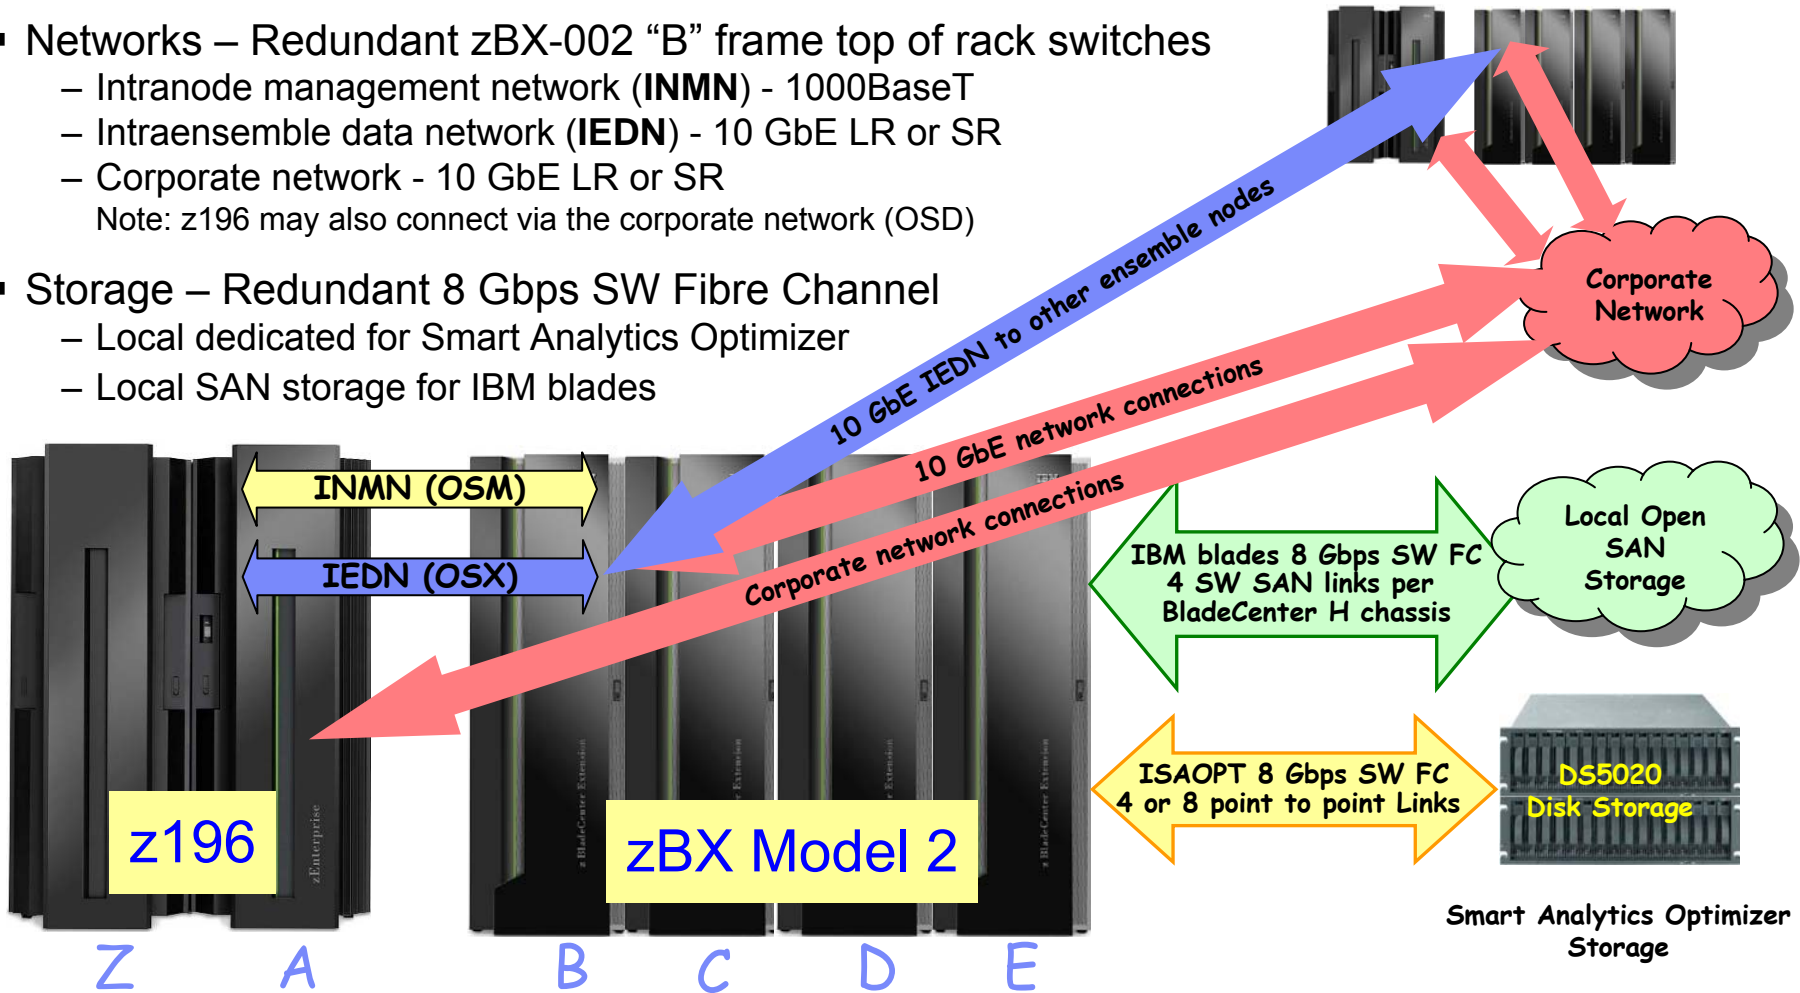
zEnterprise 196 Network Connectivity

IBM zEnterprise - Freedom by Design




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




zEnterprise 196 Parallel Sysplex

IBM zEnterprise - Freedom by Design





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



z114 and z196 GA2 InfiniBand Coupling Fanouts

Description	F/C	Ports	Comments
HCA3-O LR 1x IB DDR 	0170	4	PSIFB coupling (10 km unrepeated, 100 km with DWDM) Double port density. More subchannels per CHPID.
HCA3-O 12x IB DDR 	0171	2	PSIFB coupling (150 m) Improved responsiveness (HCA3-O to HCA3-O)
HCA2-O 12x IB-DDR	0163	2	Coupling (150 meters) Also available on z10 EC, z10 BC. Required for 12x connection to System z9 HCA1-O.
HCA2-O LR 1x IB-DDR Carry Forward only	0168	2	Coupling (10 km unrepeated, 100 km with DWDM) Also available on z10 EC, z10 BC

Note: Coupling fanouts compete for slots with the HCA2-C and PCIe fanouts for I/O drawers and cages.

Note: The InfiniBand link data rates do not represent the performance of the link. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

z114 and z196 GA2 Parallel Sysplex Coupling Connectivity

z9 EC and z9 BC S07

IFB 12x SDR, ISC-3

z9 to z9 IFB is NOT supported



HCA1


ISC-3, 2 Gbps
10/100 km

12x IFB, 3 GBps
Up to 150 m

ISC-3, 2Gbps, 10/100 km

12x IFB, 3 Bps, 150 m

z196



HCA3-O LR

OR

HCA2-O LR*

HCA3-O

OR

HCA2-O

HCA2-O

OR

HCA3-O

12x IFB3
or IFB
6 GBps
150 m

HCA3-O

OR

HCA2-O


HCA2-O

z114

HCA2-O

z10 EC and z10 BC

IFB 12x and 1x, ISC-3,



HCA2-O LR

HCA2-O

HCA2-O LR

1x IFB, 5 Gbps
10/100 km

ISC-3, 2 Gbps
10/100 km

12x IFB, 6 GBps
150 m

HCA2-O LR*

OR

HCA3-O LR

1x IFB, 5 Gbps
10/100 km

HCA3-O LR

OR

HCA2-O LR*

1x IFB, 5 Gbps, 10/100 km

ISC-3, 2 Gbps, 10/100 km

Note: ICB-4 and ETR are NOT supported on z196 or z114

*HCA2-O LR carry forward only on z196 and z114

Note: The InfiniBand link data rates do not represent the performance of the link. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.



**z800, z900
z890 and z990**
Not supported!

z196 coupling link GA2 connectivity summary and comparison

Features	Minimum # of features	Maximum # of features	Maximum connections	Increments per feature	Purchase increments
HCA3-O LR (1x)	0	12	48 links (M15: 32*)	4 links	4 links
HCA3-O (12x)	0	16	32 links (M15: 16*)	2 links	2 links
ISC-3	0	12	48 links	4 links	1 link
HCA2-O LR (1x) – CF only	0	16 (M15-8*)	32 links (M15:16*)	2 links	2 links
HCA2-O (12x)	0	16 (M15-8*)	32 links (M15:16*)	2 links	2 links


Server	1x IFB (HCA3-O LR)	12x IFB & 12x IFB3 (HCA3-O)	1x IFB (HCA2-O LR)	12x IFB (HCA2-O)	IC	ICB-4	ICB-3	ISC-3	Max External Links	Max Coupling CHPIDs
z196	48 M15 – 32*	32 M15 – 16* M32 – 32*	32 M15 – 16* M32 – 32*	32 M15 – 16* M32 – 32*	32	N/A	N/A	48	104 (1)	128
z10 EC	N/A	N/A	32 E12 – 16*	32 E12 – 16*	32	16 (32/RPQ)	N/A	48	64	64
z9 EC	N/A	N/A	N/A	HCA1-O 16 S08 - 12	32	16	16	48	64	64

* Uses all available fanout slots. Allows no other I/O or coupling

** A z196 M49, M66 or M80 supports a maximum 96 extended distance links (48 1x IFB and 48 ISC-3) plus 8 12x IFB links.
 A z196 M32 supports a maximum 96 extended distance links (48 1x IFB and 48 ISC-3) plus 4 12x IFB links*.
 A z196 M15 supports a maximum 72 extended distance links (24 1x IFB and 48 ISC-3) with no 12x IFB links*.

zEnterprise 196 Physical Planning

IBM zEnterprise - Freedom by Design




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.

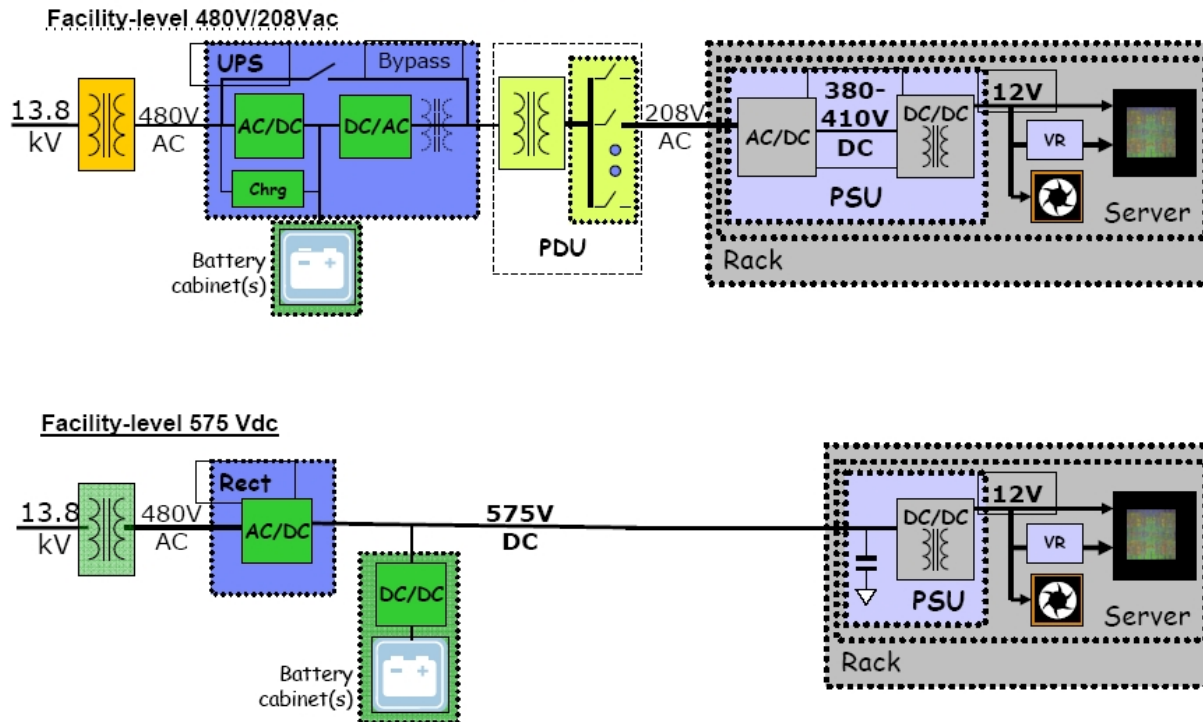
z196 optional high voltage DC power

- 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?

IBM zEnterprise - Freedom by Design



zEnterprise z196 Backup

IBM zEnterprise - Freedom by Design




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, and zIIP are 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

Current FICON Express4 and OSA-Express2 Statements of Direction* July 12, 2011 Announcements

- **The IBM zEnterprise 196 and the IBM zEnterprise 114 will be the last System z servers to support FICON Express4 features:** IBM plans not to offer FICON Express4 features as an orderable feature on future System z servers. In addition, FICON Express4 features **cannot be carried forward** on an upgrade to such follow-on servers. Enterprises should begin migrating from FICON Express4 to FICON Express8S features.
 - For z196, this new Statement of Direction restates the SOD in Announcement letter 110-170 of July 22, 2010.
- **The IBM zEnterprise 196 and the IBM zEnterprise 114 will be the last System z servers to support OSA-Express2 features:** IBM plans not to offer OSA-Express2 features as an orderable feature on future System z servers. In addition, OSA-Express2 features **cannot be carried forward** on an upgrade to such follow-on servers. Enterprises should begin migrating from OSA-Express2 features to OSA-Express4S 10 GbE and GbE features and OSA-Express3 1000BASET features.
 - For z196, this new Statement of Direction restates the SOD in Announcement letter 110-170 of July 22, 2010.

*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.

Current ESCON Statement of Direction*

July 12, 2011 Announcements

- **The IBM zEnterprise 196 and the IBM zEnterprise 114 will be the last System z servers to support ESCON channels:** IBM plans not to offer ESCON channels as an orderable feature on future System z servers. In addition, ESCON channels **cannot be carried forward** on an upgrade to such follow-on servers. This plan applies to channel path identifier (CHPID) types CNC, CTC, CVC, and CBY and to featured 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 simplify and manage an all FICON environment with continued connectivity to ESCON devices if required.
- Notes:
 - For z196, this new Statement of Direction restates the SOD in Announcement letter 111-112 of February 15, 2011. It also confirms the SOD in Announcement letter 109-230 of April 28, 2009 that “ESCON Channels will be phased out.”

*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.

Current Power Sequence Controller Statement of Direction* July 12 , 2011 Announcements

- **The IBM zEnterprise 196 and the zEnterprise z114 are the last System z servers to support the Power Sequence Controller (PSC) feature.** IBM intends to not offer support for the PSC (feature #6501) on future System z servers after the z196 (2817) and z114 (2818). PSC features cannot be ordered and cannot be carried forward on an upgrade to such a follow-on server.

- **Notes:**
 - This is a revision to the PSC statement of general direction published October 20, 2009, IBM System z10 - Delivering Security-Rich Offerings to Protect Your Data, Hardware Announcement 109-678.

 - The PSC optional feature provides the ability to power control units with the required hardware interface on and off from the System z server.

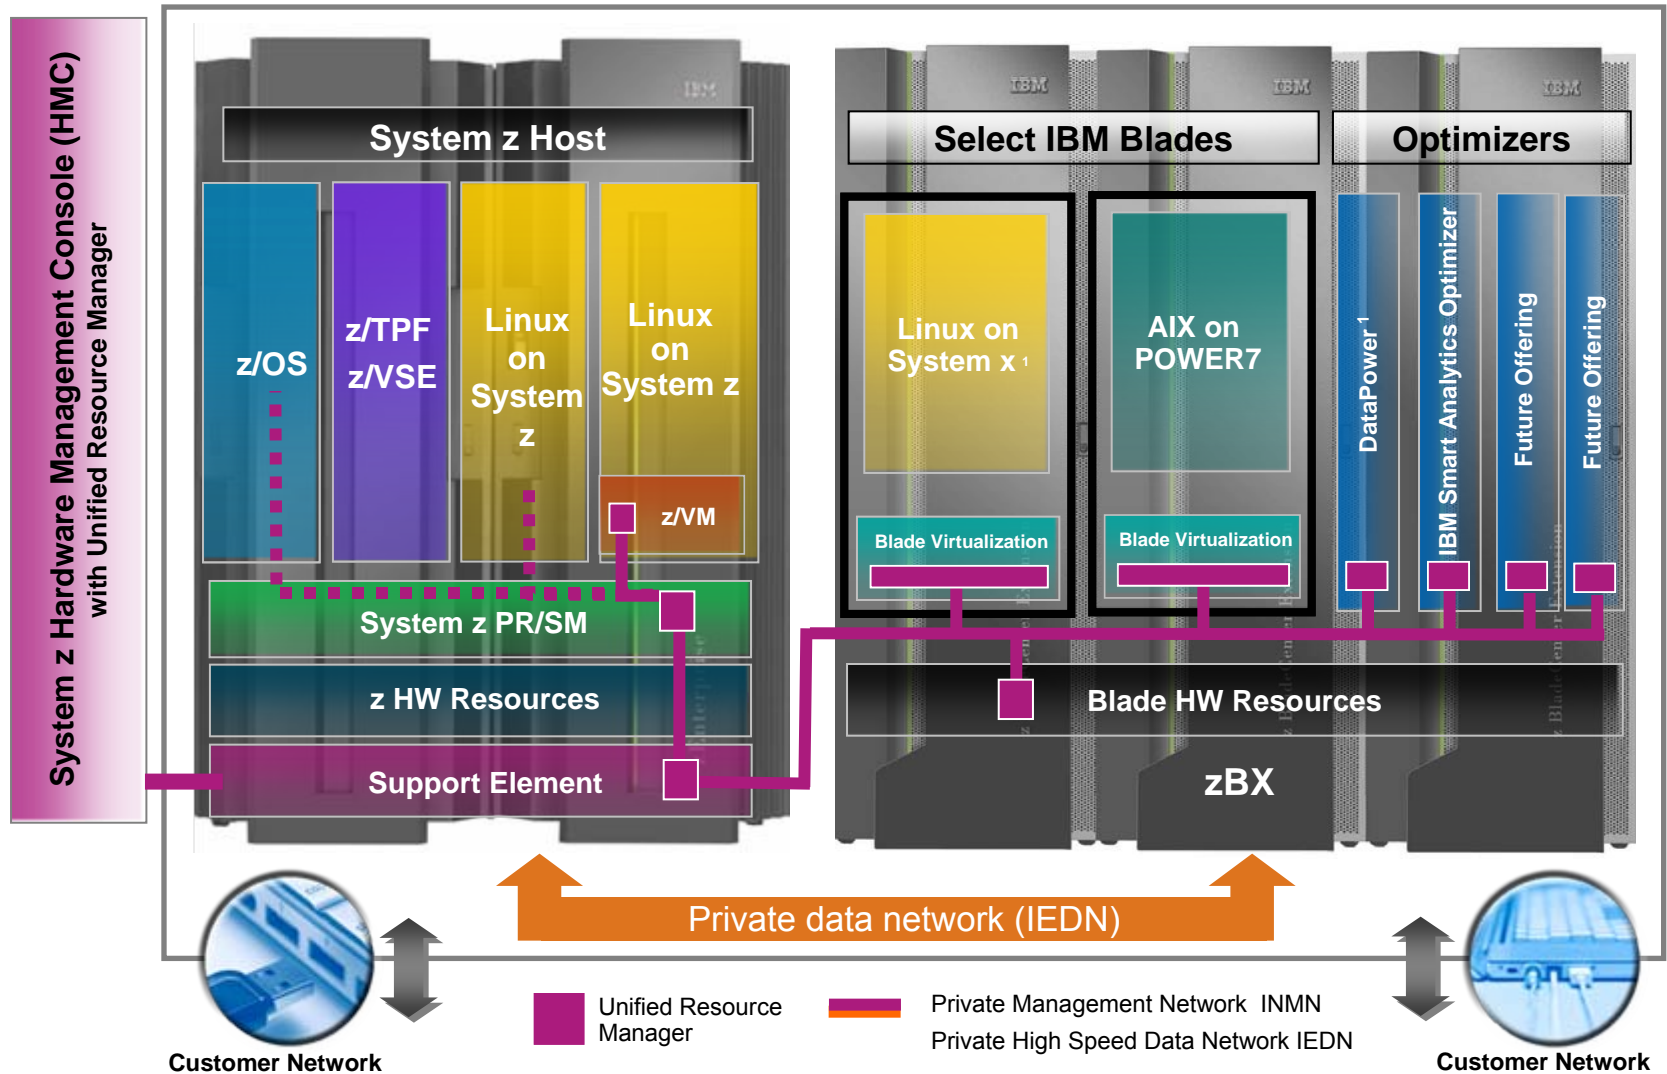
***All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these statements of general direction is at the relying party's sole risk and will not create liability or obligation for IBM.**

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.

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.

zEnterprise 196 Overview and Update Final Slide

IBM zEnterprise - Freedom by Design

