



zEnterprise Hardware Overview



Agenda

- zEnterprise Hardware
- Network Connectivity
 - OSA2 to OSA3
- NPIV on 'z' and zFCP

zEnterprise System Hardware Overview Introduction



IBM zEnterprise 196 (z196)



IBM zEnterprise BladeCenter® Extension (zBX™) Model 002



IBM System z10® EC or BC



IBM zEnterprise BladeCenter® Extension (zBX™) Model 001

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)



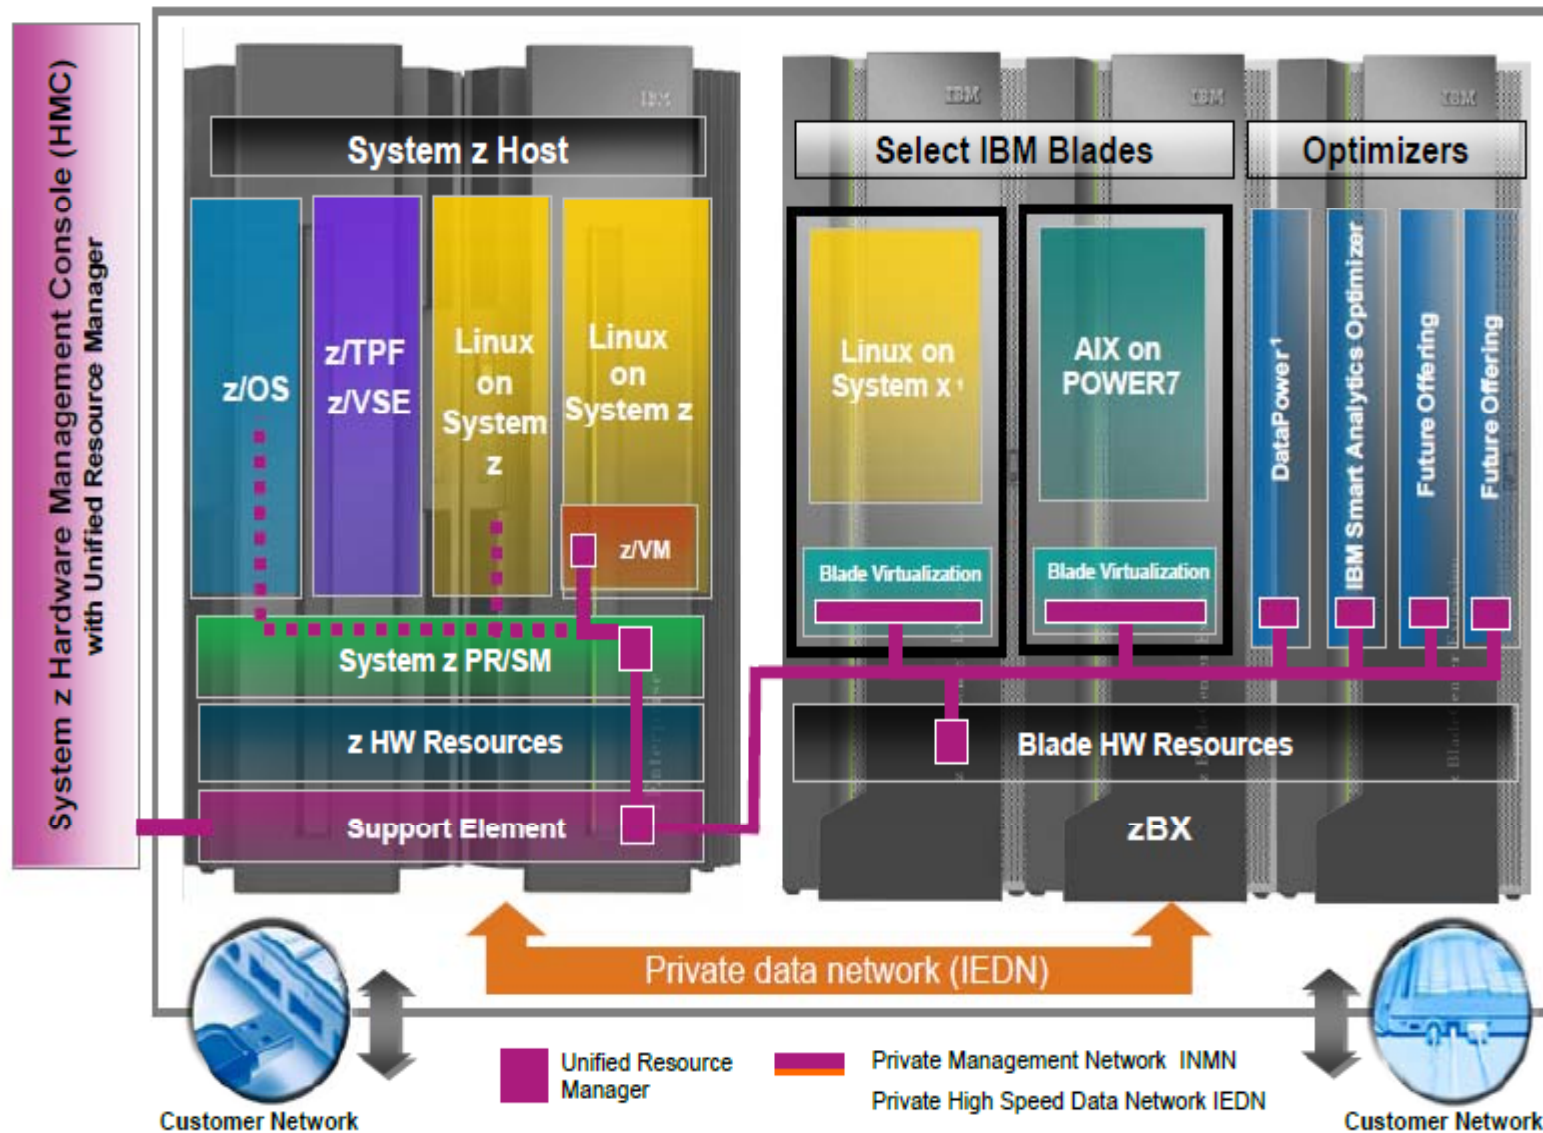
Scale out to a trillion instructions per second:
IBM zEnterprise BladeCenter® Extension (zBX)

- Selected IBM POWER7® blades and IBM System x® 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.

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 Unified Resource Manager Exploitation

- z/OS V1.10¹ and higher plus PTFs
 - z/VM 6.1¹ plus PTFs
 - Linux¹ on System z
 - Novell SUSE SLES 10 and SLES 11
 - Red Hat RHEL 5
 - Note: Distributors determine future release support
 - AIX on POWER7 Blades
 - AIX 5.3 (Technology Level 12) and later in Power 6 and 6+ compatibility mode
 - AIX 6.1 (Technology Level 5) and later
 - Linux on System x (Statement of Direction*)
 - Applications – Designed to support all applications supported on the above operating systems
1. Older releases of z/OS, older versions of z/VM, and other operating systems supported on z196 can run on a z196 in an ensemble but cannot be managed by or benefit from Unified Resource Manager function.



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

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

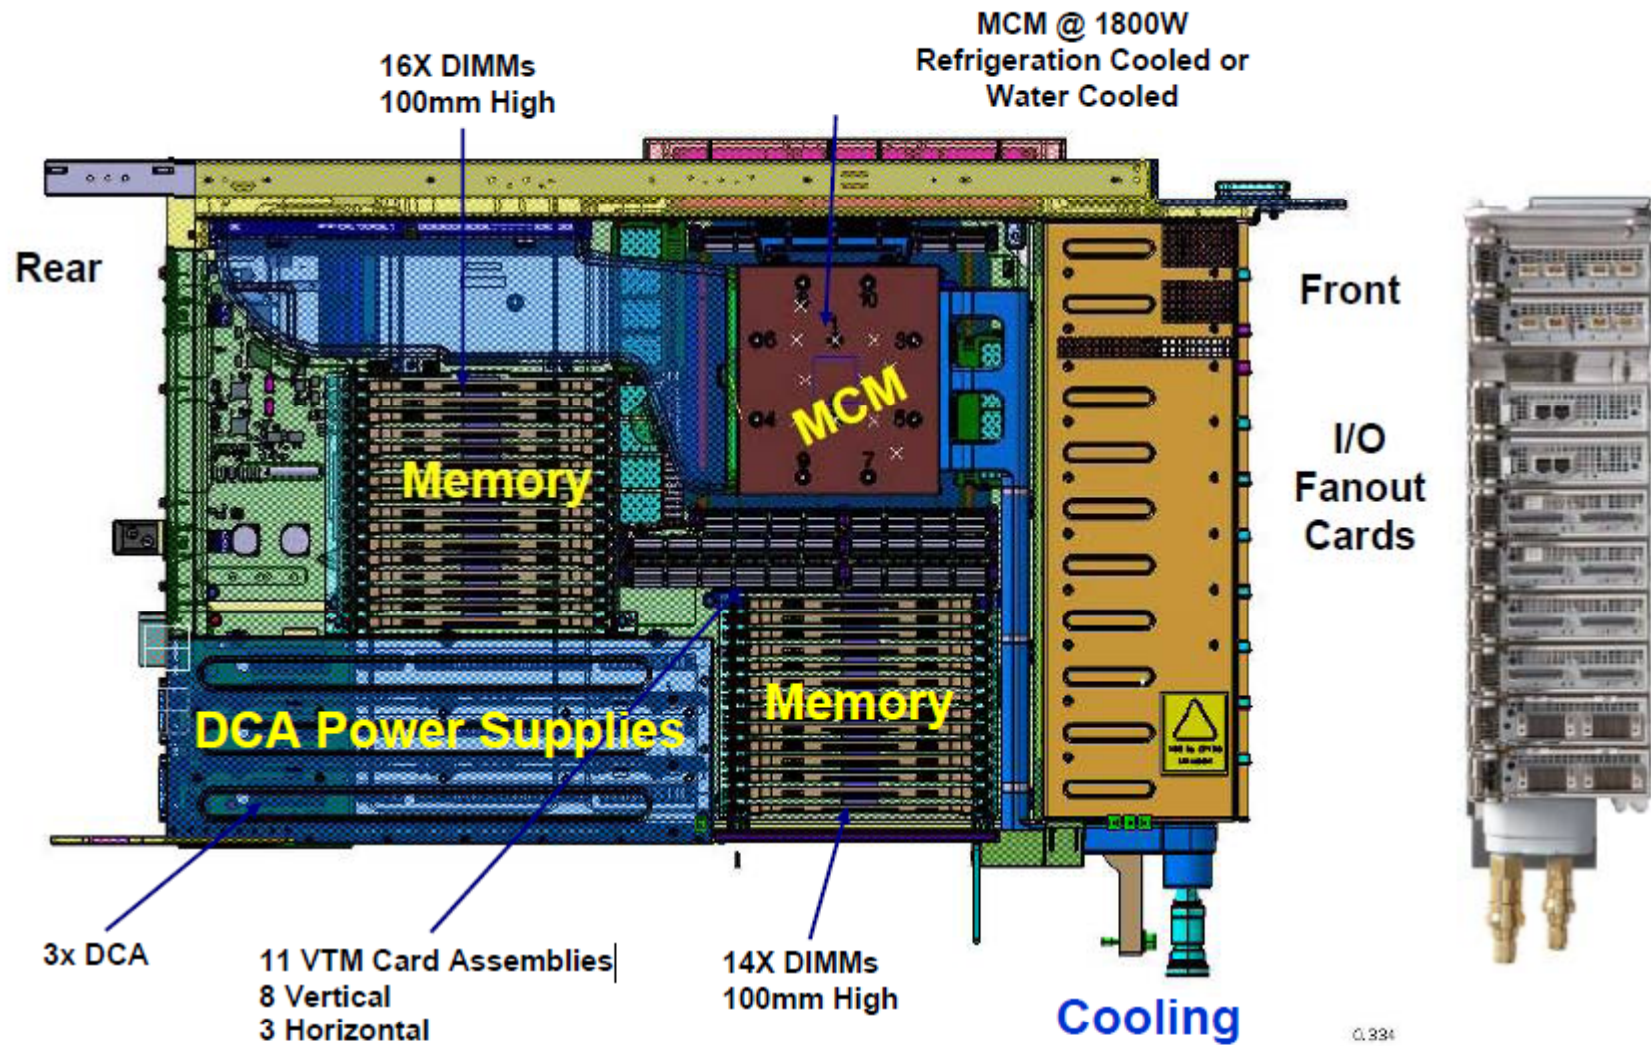
– 1800 Watts

z196 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**

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

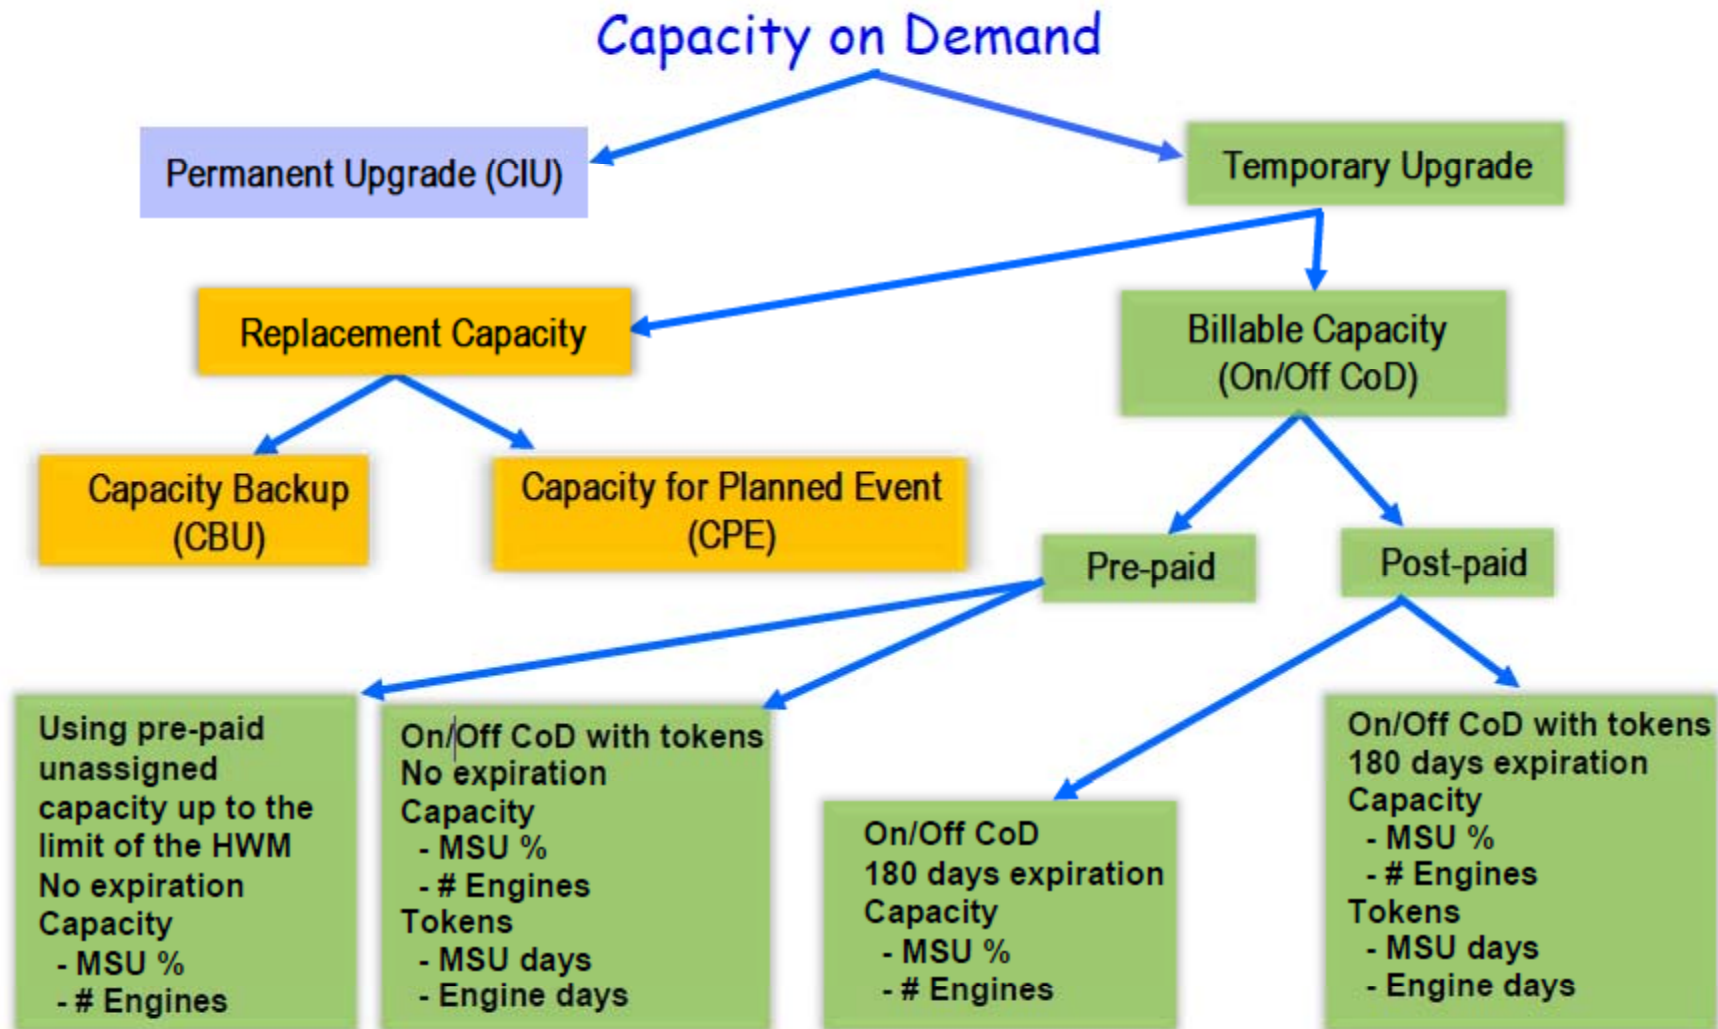
Notes: 1. At least one CP, IFL, or ICF must be purchased in every machine
 2. One zAAP **and** one zIIP may be purchased for each CP purchased even if CP capacity is "banked".
 3. "uFL" stands for Unassigned IFL

z196 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 Basics of Capacity on Demand



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

z196 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.
 - **ESCON channels to be phased out. 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.**

z196 I/O Drawer

Front view



I/O cards



Rear view

I/O cards

STI-A mother card

Air exhaust

- **Introduced with z10 BC**
 - Up to 8 I/O cards in each drawer
4 in front and 4 in rear
- **Concurrent add, repair and replacement for systems with more than one I/O drawer**
- **Drawer can be removed without affecting system input power or power to any other unit**
- **Drawers are favored on z196**
- **New Build Examples**
 - Up to 32 I/O cards use 1 to 4 drawers
 - 33 to 72 I/O cards use 1 or 2 z10 I/O cages plus up to 2 drawers
- **I/O cards are horizontal**
- **IBM Service will route cables to the side so as not to block concurrent replacement of I/O cards or drawers**

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**
 - **RPQ – 8P2507 (Please don't.)**
 - **More than 72 I/O feature cards**
 - **RPQ – 8P2506 (Please don't. REALLY!)**
 - **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

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


z196 Overhead Cabling Option (FC 7942)




Overhead cabling is designed to provide increased flexibility and increase air flow in raised-floor environments.

- **Design is side cable “chimneys” or raceways**
 - Fiber I/O and Fiber Trunking
 - 1000BASE-T Ethernet (copper)
 - Supports FQC / FTS options (side mounted brackets)
 - Additional floor cutouts NOT required

- **Physical Planning**
 - Weight add 95 lbs per frame
 - Width add 12 inches
 - Height add 5.5 inches



z196 Front View Overhead Cabling Option (FC 7942)





zEnterprise z196 Network Connectivity




z196


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**
 - Gigabit Ethernet – OSD TCP/IP or and OSN for the communication controller for Linux
- **Port usage in 2-port 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



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
<b style="color: blue;">Unified Resource Manager	<b style="color: blue;">OSM	<b style="color: blue;">Connectivity to intranode management network (INMN)
OSA for NCP (LP-to-LP)	OSN	NCPs running under IBM Communication Controller for Linux (CDLC)

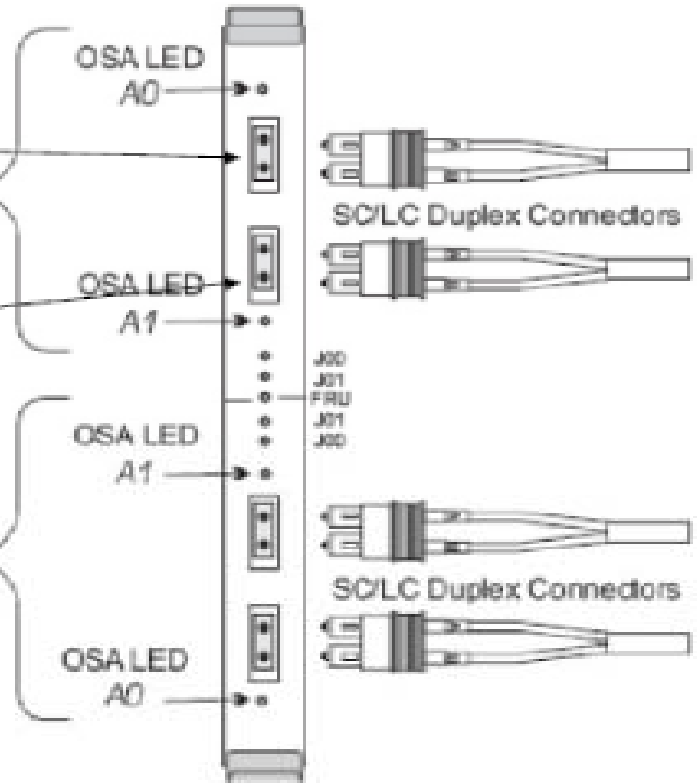
OSA-Express3 - Two Ports per CHPID Definitions

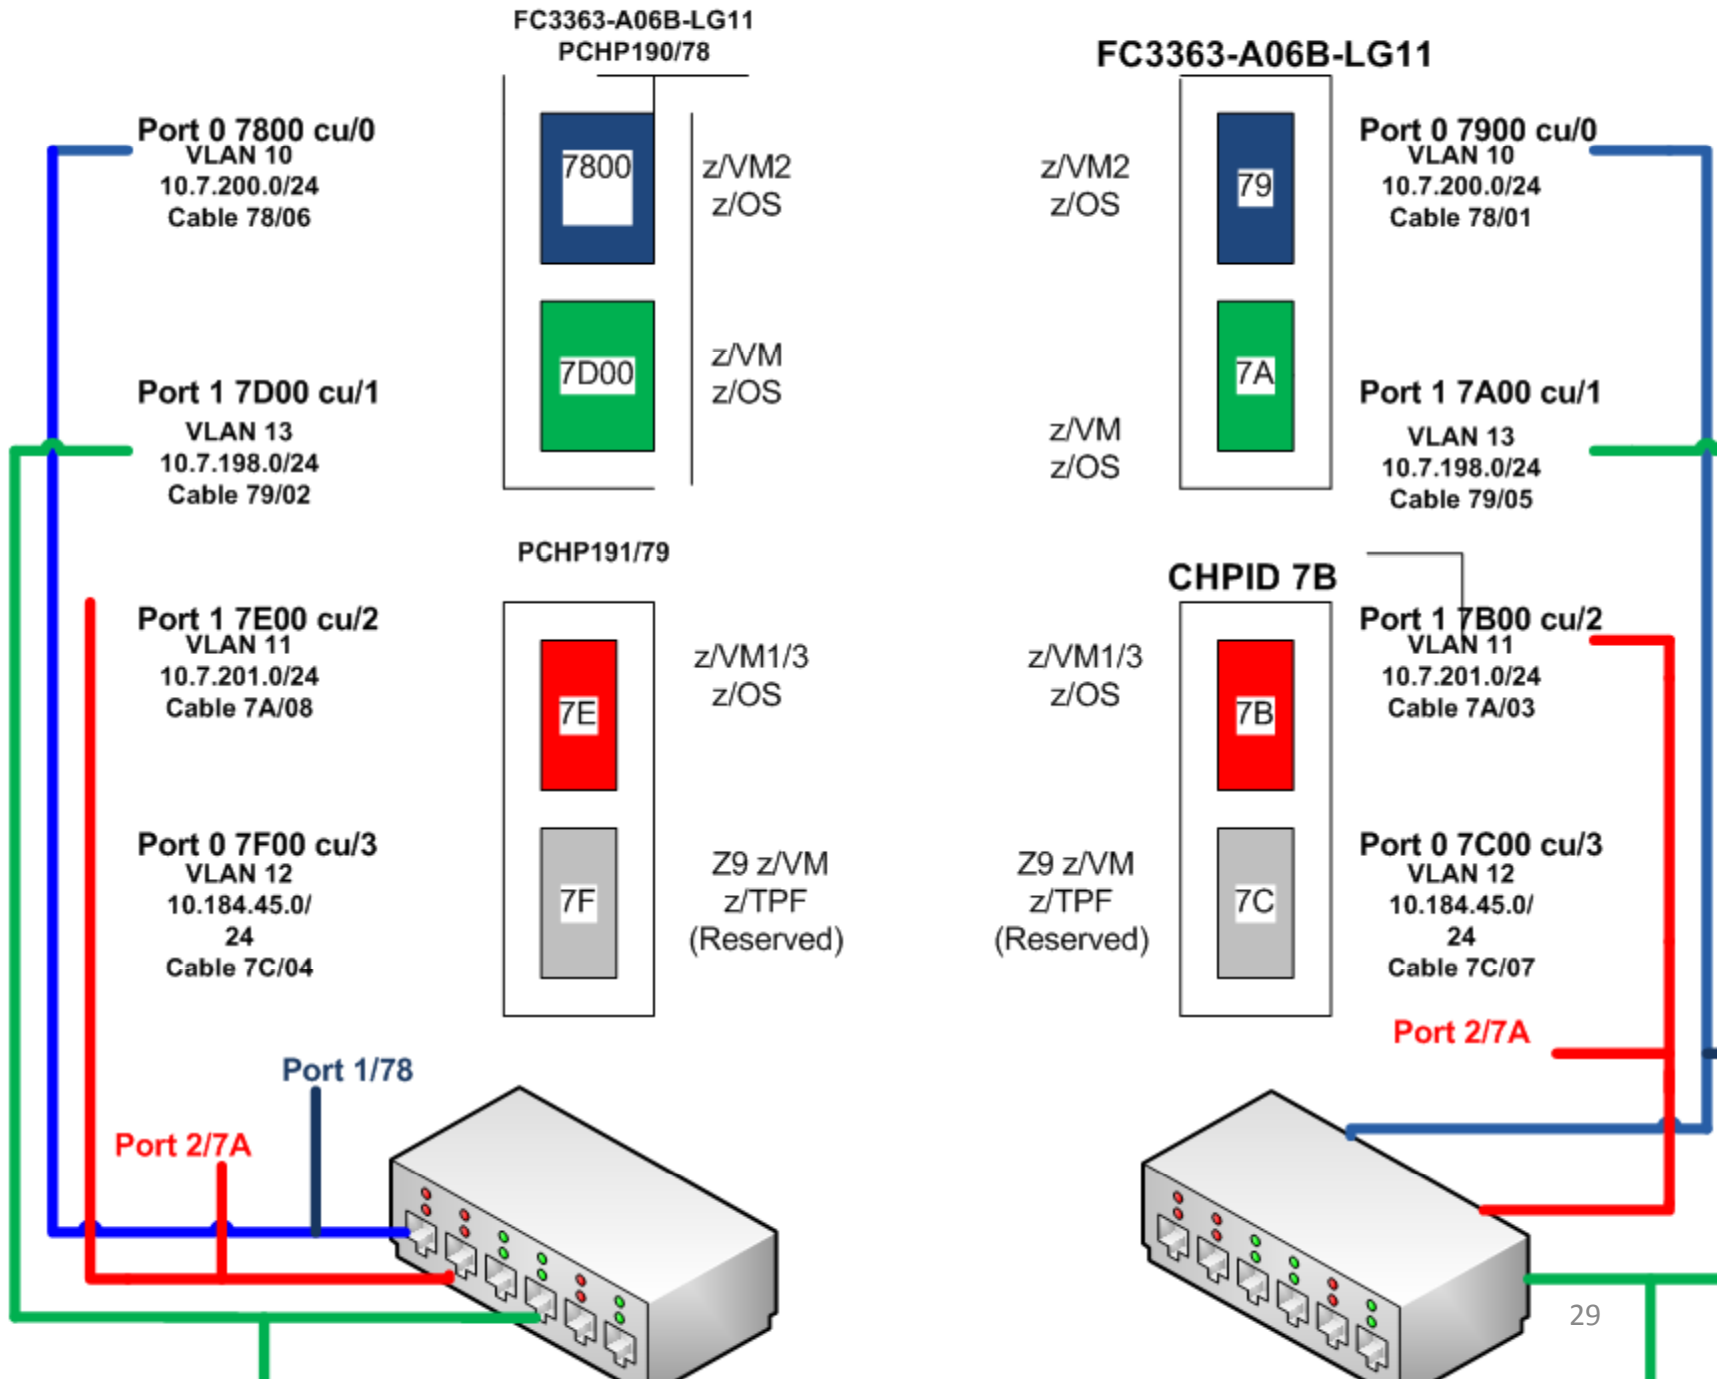
```
OSA2980 VBUILD TYPE=TRL
OSA2980P TRLE LNCTL=MPC,
READ=2980,
WRITE=2981,
DATAPATH=2982,
PORTNAME=(OSA2980),
MCPLEVEL=QD80
```

```
OSA2986 VBUILD TYPE=TRL
OSA2986P TRLE LNCTL=MPC,
READ=2986,
WRITE=2987,
DATAPATH=2988,
PORTNAME=(OSA2986),
PORTNUM=1
MCPLEVEL=QD80
```

CHPID 02
Ports 0 and 1

CHPID X
Ports 0 and 1





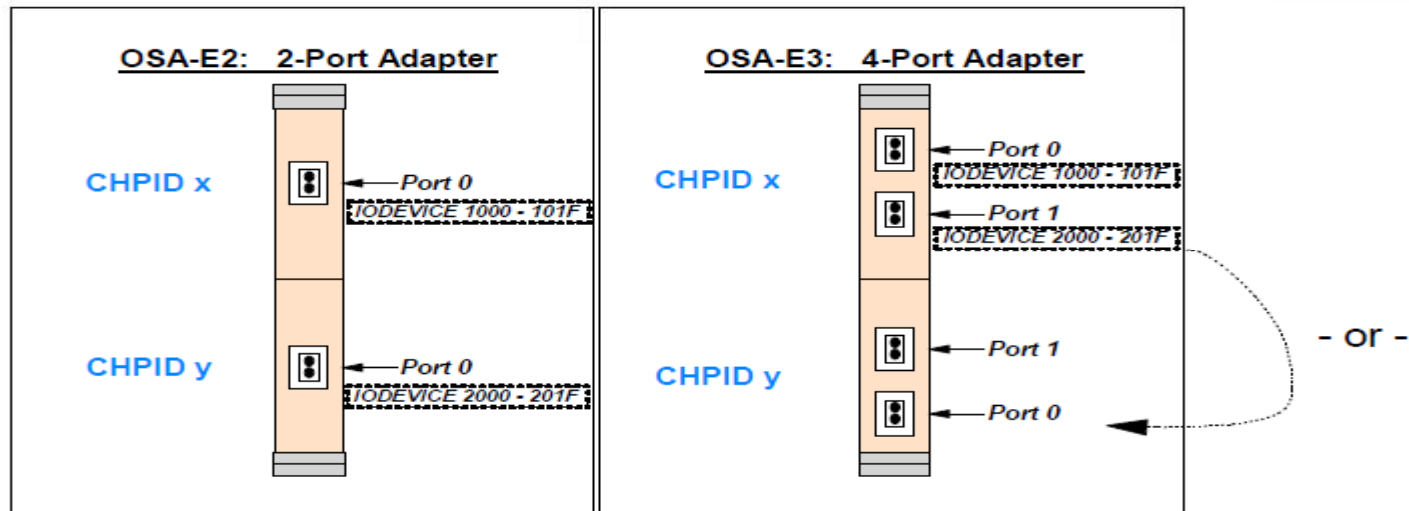
Feature Code	PHYSICAL IO	PCHP PORT/TCPIP PORT	TCPIP IODEVICE	Cable id #	VLAN ID/IP Address	Workload
FC3363-OSA Express3 GbE SX	Z01B LG02					
	310 0.79(S)	J00/PORT0	7900-790F, 7A00-7A0F	01	VLAN 10 -10.7.200.0/24	zVM2/zOS
	310 0.79(S)	J01/PORT1	7900-790F, 7A00-7A0F	05	VLAN 13-10.7.198.0/24	zVM2/zOS
	311 0.7B(S)	J02/PORT1	7B00-7B0F, 7C00-7C0F	03	VLAN 11-10.7.201.0/24	z/VM1/3-zOS
	311 0.7B(S)	J03/PORT0	7B00-7B0F, 7C00-7C0F	07	VLAN 12-10.184.45.0/24	Z9 z/VM-zTPF

CNTLUNIT CUNUMBR=7900,PATH=((CSS(0),79),(CSS(1),79)),UNIT=OSA
 IODEVICE ADDRESS=(7900,016),CUNUMBR=(7900),UNIT=OSA
IODEVICE ADDRESS=(7A00,015),UNITADD=10,CUNUMBR=(7900),UNIT=OSA
 IODEVICE ADDRESS=(7AFE,001),CUNUMBR=(7900),UNIT=OSAD

Assigning Device Addresses (QDIO): Alternative 3 (2 device ranges)



A Smarter Plan



```

CNTLUNIT CUNUMBR=1000,PATH=((CSS(0),10)),UNIT=OSA
IODEVICE ADDRESS=(1000,032),CUNUMBR=(1000),UNIT=OSA (A)
IODEVICE ADDRESS=(10FE,001),CUNUMBR=(1000),UNIT=OSAD
IODEVICE ADDRESS=(2000,032),UNITADD=20,CUNUMBR=(1000),UNIT=OSA
    
```

Multi-Port OSA

- System z10 OSA-Express3 support
 - Multiple ports per adapter on one CHPID
- Allow port numbers to be specified for Virtual Switch real devices
- Virtualization enables any single port number to be used
- Report port number in QUERY VSWITCH, QUERY PORT, QUERY LAN, QUERY NIC

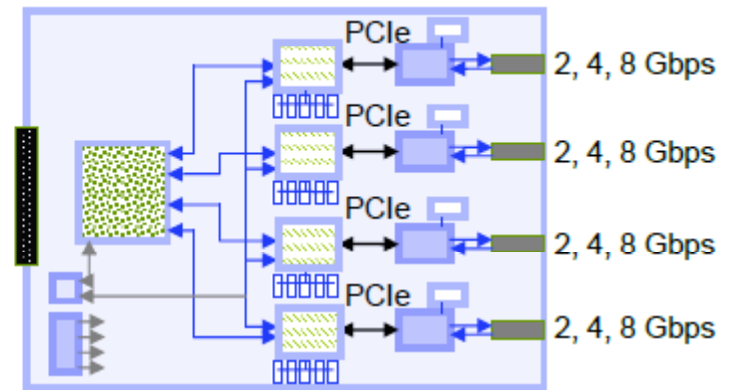
Multi-Port OSA ...

```
▶▶-DEfIne-VSWITCh-name-...-RDEV-nnnn.Pnn-...▶▶
▶▶-Set-VSWITCh-name-...-RDEV-nnnn.Pnn-...▶▶
▶▶-Set-PORT GROUp-groupname-...-JOIn-nnnn-...▶▶
    -LEAve-nnnn.Pnn
▶▶-MODIfy-PORT GROUp-groupname-...-JOIn-nnnn-...▶▶
    -LEAve-nnnn.Pnn
```



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**

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



3325 – 10KM LX, # 3326 – SX



NPIV on 'z' to support FCP devices

The WWPN Tool for System z10=> FCP Channels assists you in handling configuration files which are required or generated by System z machines when FCP Channels are installed. In particular, it helps during installation of new machines and machine upgrades.

**Resource Link**

Site search

Planning

Education

Library

Fixes

Problem solving

Services

Tools

Product tools

Customer Initiated Upgrade

Feedback

IBM Systems > System z > Resource Link > Tools >

WWPN tool

The new worldwide port name (WWPN) tool provides advance SAN preplanning so you are ready before a new System z10 server arrives.

Downloads

- [WWPN Tool - Linux Version 1.1.2 \(23MB\)](#)
- [WWPN Tool - Windows Version 1.1.2 \(23MB\)](#)

WWPN IOSN (I/O Serial Number)


- [View WWPN IOSN](#)

Before you use the WWPN tool you may want to review the following

- [WWPN Tool FAQ](#)
- [WWPN Documentation](#)

IBM provides an e-mail contact wwpntool@de.ibm.com where users can send problems specific to the WWPN Prediction Tool. Responses to customer e-mails will be provided on an as-available basis

The illustration below depicts the individual steps that need to be followed:



- 1. Export system configuration file
- 2. Create input file

3. Perform WWPN assignments

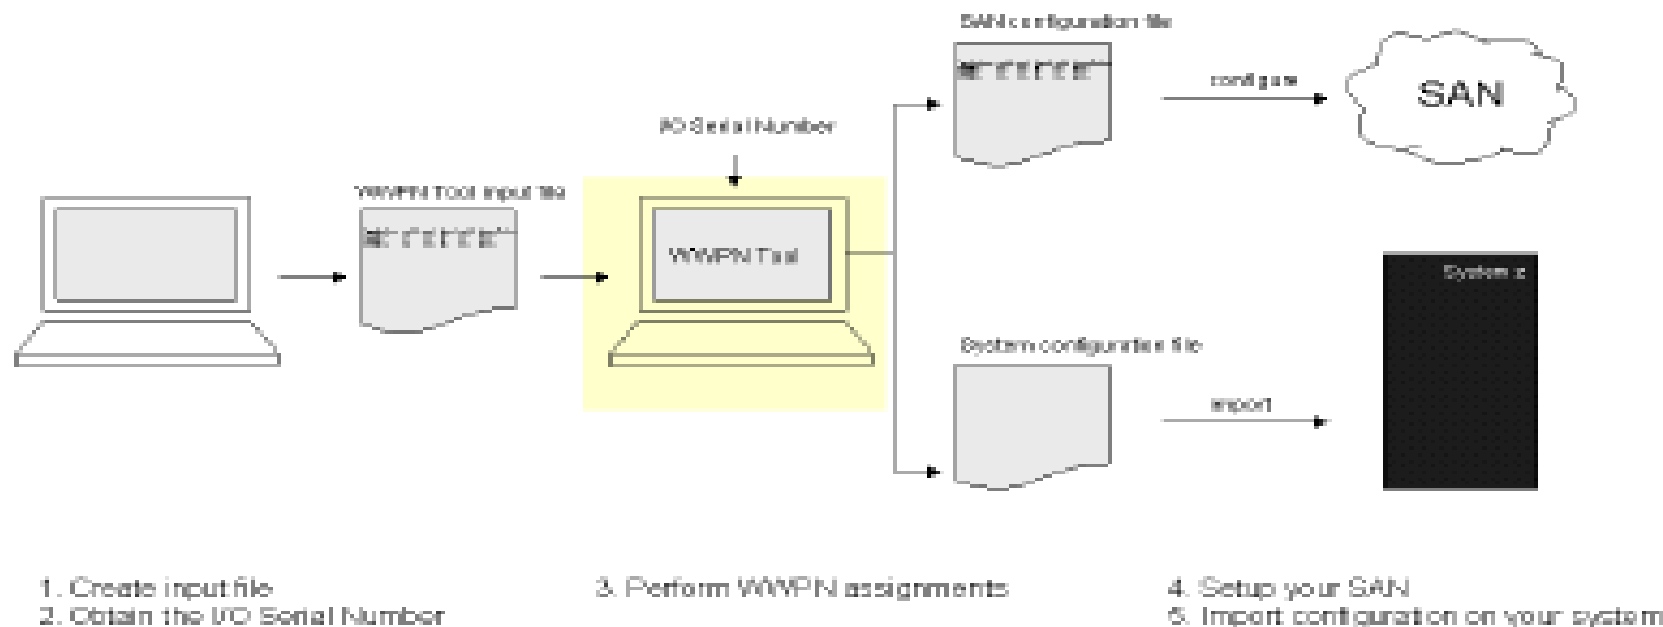
- 4. Setup your SAN
- 5. Import configuration on your system

The *NPIV system configuration file* must be imported on your System z machine when it becomes available. After the next (or initial) power-on reset of that machine, the information from this *NPIV system configuration file* is used to access FCP SANs and devices.

To create these two files, the WWPN Tool requires the following input:

- The I/O serial number of the target System z machine.
- The FCP I/O device definitions of the target machine, provided in form of an NPIV SAN configuration template.

The illustration below depicts the required steps:



	A	B	C	D	E	F	G	H	
1	##partitionName	cssId	LPID	chpId	ssId	deviceNumber	Virtual wwpn	npiv mode	current
2	ZVM5C0L3	0	3	40	0	4000	c05076eeaf800000	Yes	
3	ZVM5C0L3	0	3	40	0	4001	c05076eeaf800004	Yes	
4	ZVM5C0L3	0	3	40	0	4002	c05076eeaf800008	Yes	
5	ZVM5C0L3	0	3	40	0	4003	c05076eeaf80000c	Yes	
6	ZVM5C0L3	0	3	40	0	4004	c05076eeaf800010	Yes	
7	ZVM5C0L3	0	3	40	0	4005	c05076eeaf800014	Yes	
8	ZVM5C0L3	0	3	40	0	4006	c05076eeaf800018	Yes	
9	ZVM5C0L3	0	3	40	0	4007	c05076eeaf80001c	Yes	
10	ZVM5C0L3	0	3	40	0	4008	c05076eeaf800020	Yes	
11	ZVM5C0L3	0	3	40	0	4009	c05076eeaf800024	Yes	
12	ZVM5C0L3	0	3	40	0	400a	c05076eeaf800028	Yes	
13	ZVM5C0L3	0	3	40	0	400b	c05076eeaf80002c	Yes	
14	ZVM5C0L3	0	3	40	0	400c	c05076eeaf800030	Yes	
15	ZVM5C0L3	0	3	40	0	400d	c05076eeaf800034	Yes	
16	ZVM5C0L3	0	3	40	0	400e	c05076eeaf800038	Yes	
17	ZVM5C0L3	0	3	40	0	400f	c05076eeaf80003c	Yes	
18	ZVM5C0L3	0	3	41	0	4100	c05076eeaf800040	Yes	
19	ZVM5C0L3	0	3	41	0	4101	c05076eeaf800044	Yes	
20	ZVM5C0L3	0	3	41	0	4102	c05076eeaf800048	Yes	
21	ZVM5C0L3	0	3	41	0	4103	c05076eeaf80004c	Yes	
22	ZVM5C0L3	0	3	41	0	4104	c05076eeaf800050	Yes	
23	ZVM5C0L3	0	3	41	0	4105	c05076eeaf800054	Yes	
24	ZVM5C0L3	0	3	41	0	4106	c05076eeaf800058	Yes	
25	ZVM5C0L3	0	3	41	0	4107	c05076eeaf80005c	Yes	
26	ZVM5C0L3	0	3	41	0	4108	c05076eeaf800060	Yes	
27	ZVM5C0L3	0	3	41	0	4109	c05076eeaf800064	Yes	
28	ZVM5C0L3	0	3	41	0	410a	c05076eeaf800068	Yes	
29	ZVM5C0L3	0	3	41	0	410b	c05076eeaf80006c	Yes	
30	ZVM5C0L3	0	3	41	0	410c	c05076eeaf800070	Yes	
31	ZVM5C0L3	0	3	41	0	410d	c05076eeaf800074	Yes	
32	ZVM5C0L3	0	3	41	0	410e	c05076eeaf800078	Yes	
33	ZVM5C0L3	0	3	41	0	410f	c05076eeaf80007c	Yes	
34	ZVM5C0L3	0	3	41	0	4100	c05076eeaf800080	Yes	

zEnterprise z196 Physical Planning

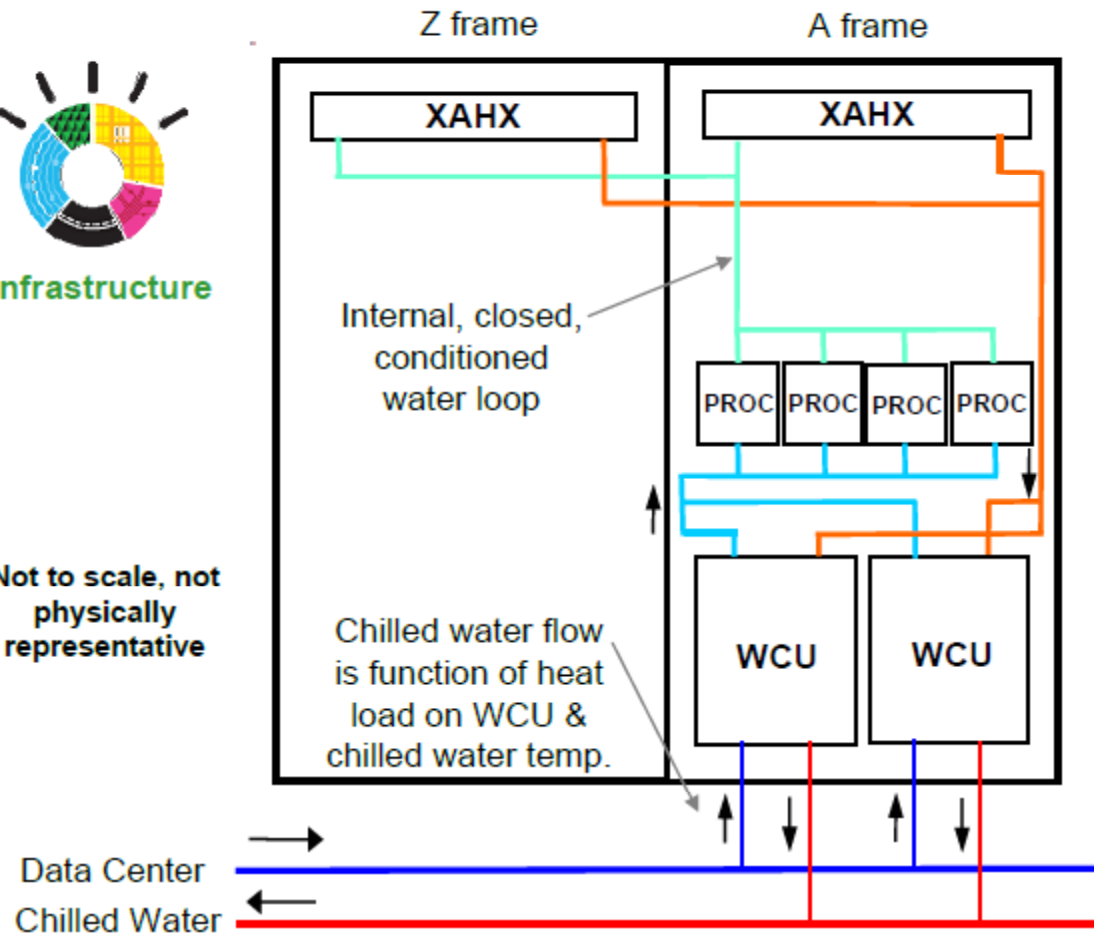


z196 optional water cooling



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.

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



Thank you

W I C O I M Infinity

Questions?

The Future Runs on System z

