zEnterprise.
A New Dimension in Computing

IBM zEnterprise BladeCenter Extension (zBX)

Trademarks

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

IBM*	Redbooks*	z9*
IBM logo*	Resource Link	z10
ibm.com	REXX	z10 BC
BladeCenter*	System x*	z10 EC
DB2*	System z*	z/OS*
eServer	System z9*	zEnterprise
InfiniBand*	System z10	zSeries*

* Registered trademarks of IBM Corporation

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.

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 System z10 zBX Topics – July 2010

1	What is an Ensemble
2	IBM zEnterprise™ BladeCenter® Extension (zBX) Announcement
3	Positioning IBM Smart Analytic Optimizer for DB2® for z/OS®, V1.1
4	Smart Analytics Optimizer System Design Overview
5	IBM zEnterprise BladeCenter Extension Hardware
6	zBX Hardware Components
7	zBX Ordering and Upgrades
8	zBX Firmware
9	Software for zBX
10	Storage for zBX
11	Systems Management and SE/HMC Updates for zBX
12	zBX Configuration and Planning
13	zBX RAS
14	Summary
15	End of Presentation
16	Back Up Charts
17	xBlades Chassis Overview
18	HS22 xBlade (used for zBX)
19	xBlade – Power and Cooling

New Glossary

Blade	Hardware that provides application-specific services and components. The consistent size and shape (or form factor) of each blade allows it to fit in a BladeCenter chassis.
BladeCenter chassis	A modular chassis that can contain multiple blades, allowing the individual blades to share resources such as the management, switch, power, and blower modules
IBM BladeCenter PS701 Express	The supported POWER7 blade that can be installed in a zBX Model 002
Data Warehouse Blades = IBM Smart Analytics Optimizer	IBM Smart Analytics Optimizer
AMM	Advance Management Module
Node	A single z196 and any optionally attached zBX. A node can be a member of only one ensemble
Ensemble	A collection of one or more zEnterprise nodes (including any optionally attached zBX) that are managed as a single logical virtualized system by the Unified Resource Manager using a Hardware Management Console (HMC)
BPH	Bulk Power Hub
BCN	BladeCenter H
HMC	Hardware Management Console (primary or alternate)
GPFS	General Parallel File System
IEDN	intraensemble data network
INMN	intranode management network
NVM	Network Virtual Manager – component on HMC responsible for MAC address and VLAN assignments for all servers running in the ensemble
OSM	OSA Channel path identifier (CHPID) for INMN
OSX	OSA Channel path identifier (CHPID) for IEDN
PDN	Private Data Network
Unified Resource Manager (zManager)	IBM zEnterprise Unified Resource Manager – The Unified Resource Manager provides energy monitoring and management, goal-oriented policy management, increased security, virtual networking, and data management for the physical and logical resources of a given ensemble
RAIM	Redundant array of independent memory (RAIM). A new technology introduced with z196 designed to provide protection at the direct random access memory (DRAM), dual inline memory module (DIMM), and memory channel level
PSN	Private Support Network
PSCN	Power System Control Network.
SOD	Statement of Direction
zBX	IBM zEnterprise BladeCenter Extension

What is an Ensemble?

What is a zEnterprise Ensemble?

- **A zEnterprise ensemble is a collection of 1 to 8 z196 CPCs with/without zBX managed collectively by the Unified Resource Manager as a single logical virtualized system using the HMC**
- **A zEnterprise node is a z196 CPC with 0 to 4 racks up to 2 BladeCenters per rack**
  - zEnterprise nodes are deployed within a single site
  - A zEnterprise node can be a member of only one ensemble
- **Blade based fit-for-purpose Solutions**
- **Integrated Advanced Virtualization Management**
- **Implements well-defined external interface to Data Center Service Management functions**
- **Virtual Resource Management and Automation**
- **z10 can access the Optimizers, but can't be part of the managed ensemble**

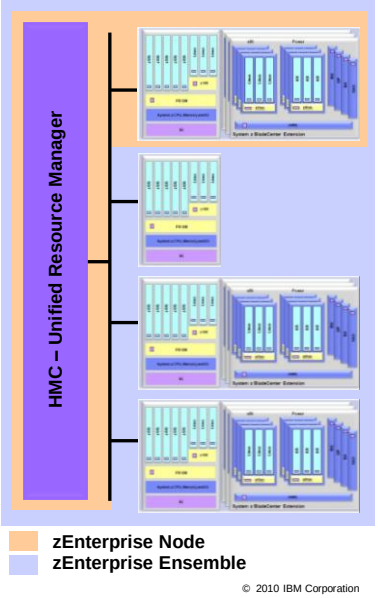

zEnterprise System - z196 + zBX + Unified Resource Manager

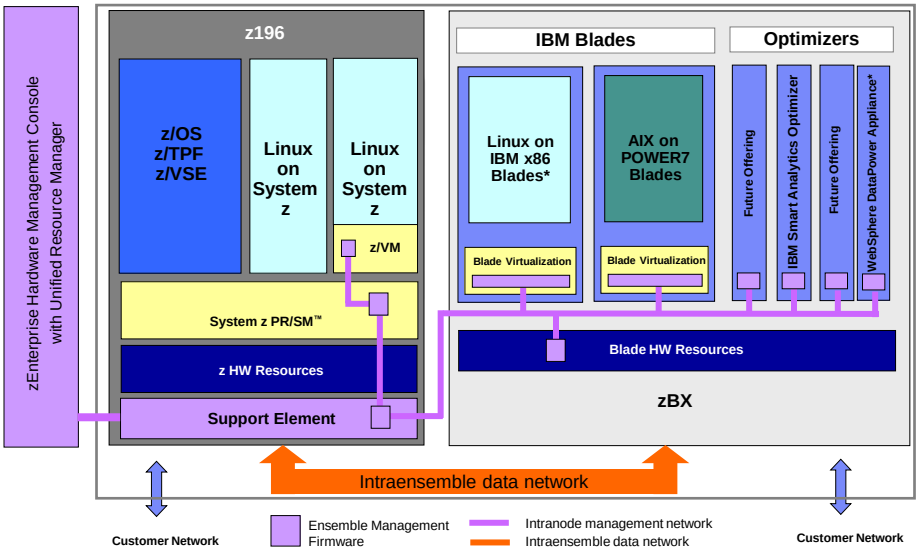

*All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

IBM zEnterprise BladeCenter Extension (zBX)

IBM zEnterprise System

IBM zEnterprise 196 (z196) | **IBM zEnterprise BladeCenter Extension (zBX)**

IBM zEnterprise Unified Resource Manager (zManager)

z196 and zBX Hardware Components Overview

z196

zBX Infrastructure

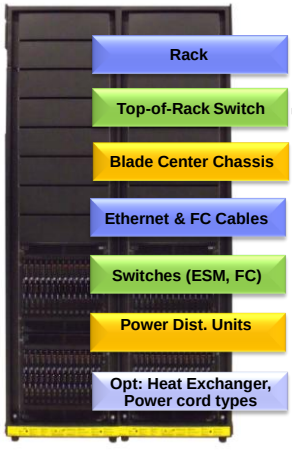

- **New System z Machine Types:**
  - CPC = 2817
  - zBX = 2458 Model 002
  - Customer supplied IBM x86* and POWER7 Blades
- **Key points:**
  1. CPC Availability and Integrity
  2. Independent Lifecycles
     - Asynchronously Upgradeable
     - Enhancements not tied to CPC HW
  3. Data availability not tied to single CPC
     - Shareable between CPCs
  4. Growth
  5. System z service

*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.

zBX Details

- **Machine Type/Model 2458-002**
  - One Model with 5 configurations for IBM Smart Analytics Optimizer
- **Racks – Up to 4 (B, C, D and E)**
  - 42U Enterprise, (36u height reduction option)
  - 4 maximum, 2 chassis/rack
  - 2-4 power line cords/rack
  - Non-acoustic doors as standard
  - Optional Acoustic Doors
  - Optional Rear Door Heat Exchanger (conditioned water required)
- **Chassis – Up to 2 per rack**
  - 9U BladeCenter
  - Redundant Power, cooling and management modules
  - Network Modules
  - I/O Modules
- **Blades (Maximum 112 in 4 racks)**
  - IBM Smart Analytics Optimizer Blades (up to 7 to 56)
    - Can not mix other Blades in the same Chassis
  - Customer supplied POWER7 Blades (up to 112)
  - Customer supplied IBM x86 Blades* (up to 112)
  - WebSphere DataPower Appliances* (up to 28)
  - Non-IBM Smart Analytics Optimizer Blades can be mixed in the same chassis
- **Management Firmware**
  - Unified Resource Manager
- **Top of Rack (TOR) Switches - 4**
  - 1000BASE-T intranode management network (INMN)
  - 10 GbE intraensemble data network (IEDN)
- **Network and I/O Modules**
  - 1000BASE-T and 10 GbE modules
  - 8 Gb Fibre Channel (FC) connected to customer supplied disks
  - IBM Smart Analytics Optimizer uses DS5020 disks
    - DS5020s not shared with Customer supplied Blades

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

Redbooks IBM

zBX Hardware Components

zBX Hardware Components Overview

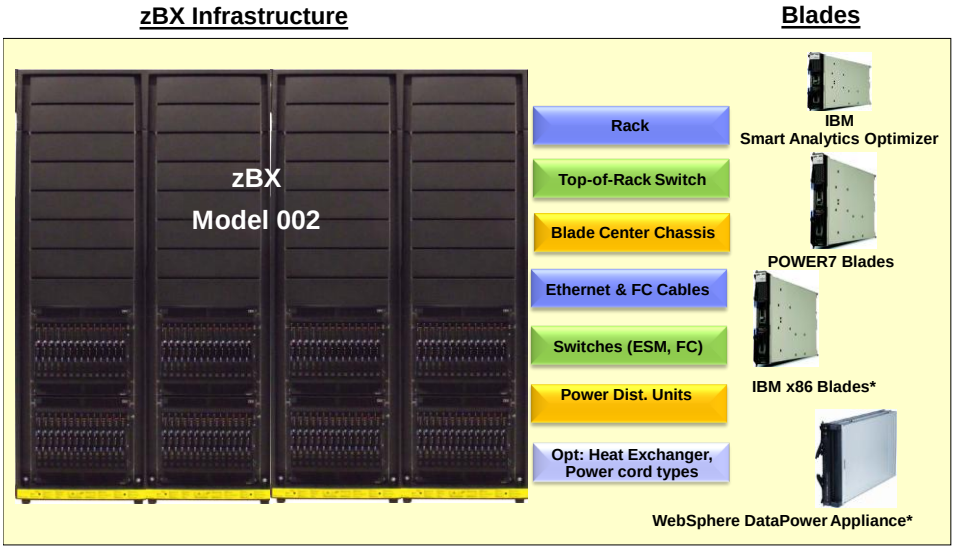

* All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

System Design Overview - continued

- **zBX Hardware Integration**
  - The zBX component hardware consists of blades within a BladeCenter chassis
  - Each BladeCenter can have up to 14 blades.
  - Maximum of 8 BladeCenters
    - One to four 19 inch 42U IBM Server Racks which should be located adjacent to to the System z frames
  - VLAN capable switches in the first rack, which are used for network connectivity between the SEs and the Management Modules/Blades.
  - The BladeCenter power and cooling is not integrated into z196 power.
    - Power comes off wall power.
  - BladeCenter hardware is attached with redundant OSA-Express3 10Gb Ethernet to z196
    - Two OSA-Express3 10 GbE required with the controlling/owning z196
- **Key Design criteria:**
  - CPC Availability and Integrity
  - Independent Lifecycles
    - Asynchronously Upgradeable
    - Enhancements not tied to CPC HW
  - Data availability not tied to single CPC
    - Shareable between CPCs
  - Growth
  - System z service

zBX Rack

- **42U of rack space**
- **Fits easily through standard 2.03m (80") high doorways**
  - External dimensions (HxWxD): 202x65x110cm (79.5"x25.6"x43.3")
  - Supports the optional IBM Rear Door Heat Exchanger and IBM Acoustic Door

BladeCenter

- **Form factor/height rack-mount chassis/9U Blade bays**
- **Power supply module**
  - Up to 4 hot-swap and redundant 2900W AC with load-balancing and failover capabilities. Operating at 200-240V
- **Cooling modules**
  - Two hot-swap and redundant blowers standard, additional fan packs on power supplies
- **Systems Management for hardware**
  - Advanced Management Module standard; additional Management Module for redundancy required

Standard BladeCenter Chassis

BladeCenter Detail

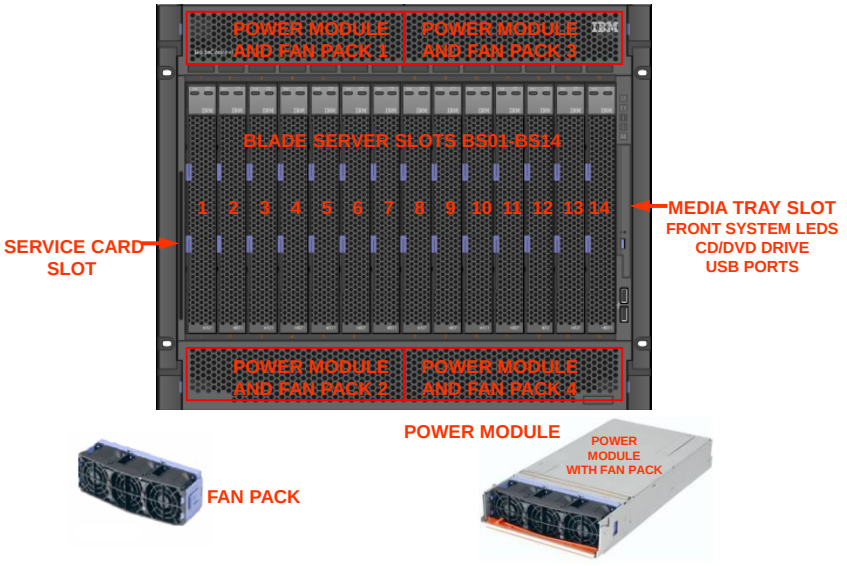

BladeCenter Configuration details

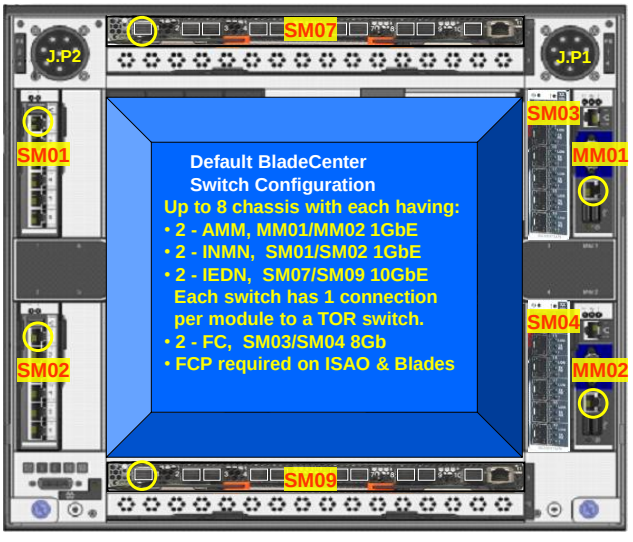

zBX – Interfaces to BladeCenter

- **High Speed Switches Modules**
  - The HSS with Short Range optics is used to connect to System z OSA-Express3 10 GbE adaptors
- **Fibre Channel Switch Modules**
  - QLOGIC 20 port 8 Gbps FC Switch
  - Fibre Channel attachment for customer supplied disk storage
- **Top of Rack (TOR) Switches for INMN**
  - The INMN provides connectivity to the Management Module in the BladeCenter Chassis and the BladeCenter Chassis Ethernet switch network. This allows the SE/HMC to communicate to the management module for service and have the ability to directly interface to the blade

High speed Switch Module

Fibre Channel Switch Module

Top of Rack Switch

Bits and Pieces

- Internal connections pre built and wired in IBM manufacturing
- External connections performed by IBM during installation
  - –Cables and labels provided by the customer
  - –Disk provided by an alternate means
  - –IBM POWER7 Blades provided by an alternate means
- Redundant network components and paths
- Redundant disk connections and paths
- Redundant Top of Rack (TOR) switches
  - –INMN
  - –IEDN

IBM Smart Analytics Optimizer

IBM Smart Analytics Optimizer

- **A special purpose, workload optimized system**
  - Accelerate typical DW queries from traditional database server to the optimizer
  - Based on IBM research prototype
  - Network attached Blades
- **No changes to the applications**
- **Designed to improve performance of typical DW queries 5-10x**
  - Vector processing - evaluation of all predicates in parallel
  - Consistent Query Response Times
  - Achieving linear scaling with the number of CPUs
- **Significant price/performance and TCO improvement**
  - Order of magnitude performance improvement for offloaded queries
  - Significant reduction in DBA effort for tuning offloaded queries (MQTs, indexes, etc.)
  - Managing the data consistently, coherently, and securely on zOS

Total solution remains centrally managed by System z…

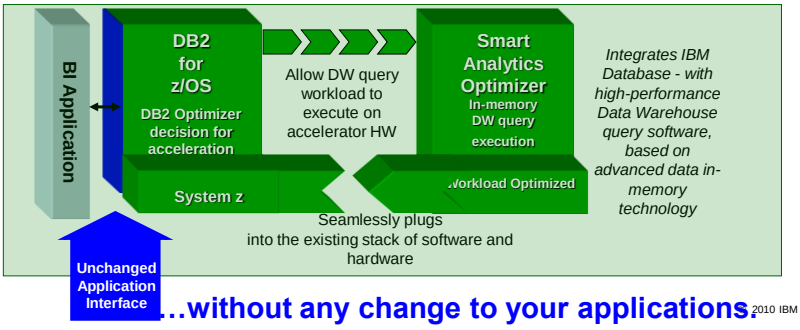

…without any change to your applications.

Smart Analytics Optimizer – A New DB2 Resource Manager

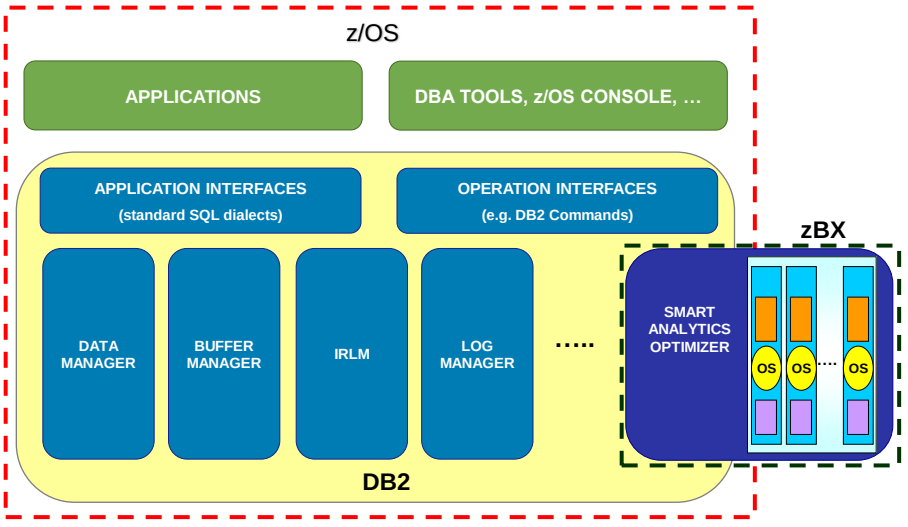

Smart Analytics Optimizer product layout

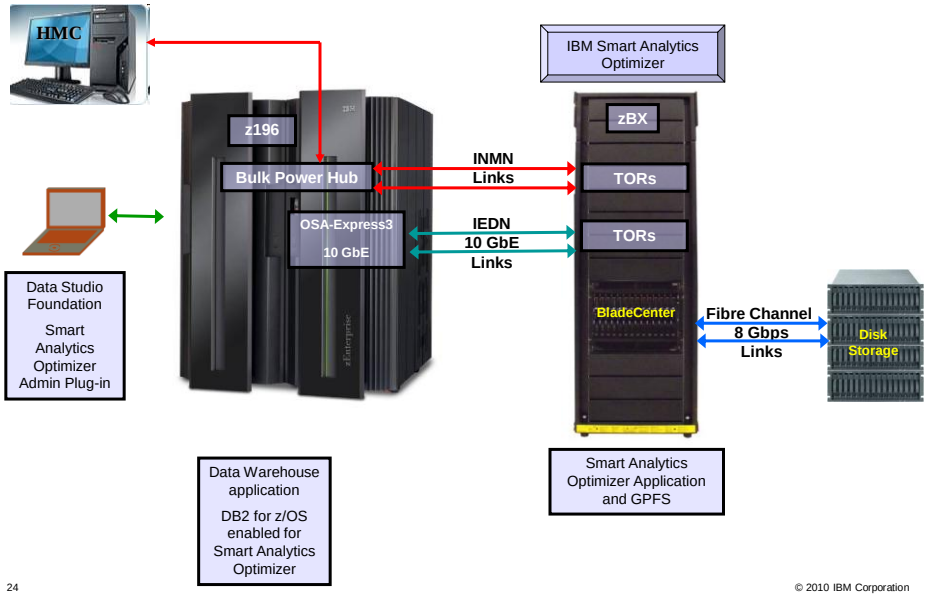

IBM Smart Analytics Optimizer Blade

- **Blade**
  - 2.93 GHz
  - 2 x QUAD Core
  - 12 x 4GB Dimms = 48GB Total
  - 31.4GB SSD
  - 10 GbE interface
  - 8 Gbps FC

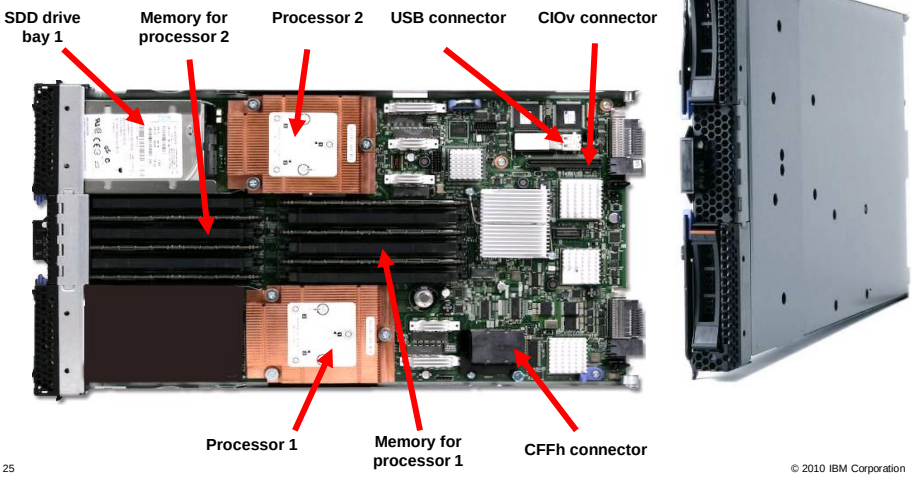

ISAO – Coordinator and Worker node roles

- **The blades can play two different roles. The number of Blades allocated for each are role determined by the Smart Analytics Optimizer**
- **Coordinator node**
  - Coordinates which worker nodes will work in parallel on a specific task/workload
  - ISAO has one or two coordinators in the TOTAL configuration. Others are used/available for 'sparing'
  - If all coordinator nodes fail, queries will be processed by DB2 as if there is no ISAO
  - At any one time, the ISAO processes only one query. The coordinator nodes builds and manages the queues for the query
- **Worker node**
  - Blades work together to solve a large problem

Websites for Additional Information for IBM Smart Analytics Optimizer

IBM Smart Analytics System Optimizer:
External page: http://www-01.ibm.com/software/data/infosphere/smart-analytics-optimizer-z/

Data Warehousing and Analytics
External page: http://www-01.ibm.com/software/data/infosphere/data-warehousing/

Data Warehousing and Business Intelligence on System z
External page: http://www-01.ibm.com/software/data/businessintelligence/systemz/

IBM POWER7 Blades

IBM POWER7 PS701Express Blade

- **General-purpose computing platform**
  - Housed in standard BladeCenter H chassis inside IBM® zEnterprise™ BladeCenter® Extension enclosure
  - Up to 112 blades
    - 14 blades per BladeCenter
    - 2 BladeCenters per rack
    - 4 racks per zBX Model 2
- **Managed by the IBM zEnterprise Unified Resource Manager**
- **Virtualized with firmware-supplied hypervisor**
- **Entitled through System z firmware**
- **Performance and Energy Efficiency**
  - Single-wide 8-core with three configurations
  - POWER7 processor-based blades automatically optimize performance
  - Ideal for highly virtualized environments with demanding commercial workload performance
  - Virtualization performance and scalability
  - Take advantage of the power of IBM's industry-leading UNIX operating system, AIX

✓ 8 cores
✓ Single Wide
✓ 3.0GHz POWER7
✓ Up to 128GB of Memory

POWER7 BUILT ON Power

IBM BladeCenter PS701 (8406-71Y) Configurations for zBX

- IBM BladeCenter PS701 (8406-71Y)
  - POWER7 8 Core Processor
  - 8 Processor Cores activated
  - 1 Processor socket
  - Single wide Blade only
  - 3.0GHz
  - 16 dimm slots (4 or 8 GB DIMMs)
  - 300GB HDD Internal Disk
- 3 Configurations shown are supported
- POWER7 Blades may be acquired by the customer through existing channels or through IBM
- PowerVM Enterprise Edition licence and Software Maintenance Agreement is required for all 8 Cores, and must be maintained for the duration of use.
- PowerVM Enterprise Edition is controlled as zEnterprise Licensed Internal Code (LIC)
  - pHyp 2.1, VIOS 2.1.3
  - Extensions for configuration and systems management: Hardware setup, FFDC, Heartbeat, PPM daemon

Blade	FC#	Config 1	Config 2	Config 3
Processor 3.0GHz@150W		1	1	1
Processor Activations	8412	8	8	8
Memory kits		32 GB	64 GB	128 GB
8 GB (2 x 4 GB)	8208	4	8	0
16 GB (2 x 8 GB)	8209	0	0	8
HDD 300GB	8274	1	1	1
CFFh 10GbE	8275	1	1	1
CIOv 8Gb FC	8242	1	1	1
PowerVM	5228	8	8	8
Required SW	PID			
PowerVM EE SW License PID	5765-PVE	8	8	8
PowerVM EE 1 YR SWMA PID	5771-PVE	Choose Qty 8 of 1 YR or 3 YR		
PowerVM EE 3 YR SWMA PID	5773-PVE	Choose Qty 8 of 1 YR or 3 YR		

Warranty and Maintenance
Separate blade warranty is NOT required if in a zBX under IBM maintenance.
zBX maintenance includes 24x7 on-site support for parts (including blades) and service during the 1 year System z warranty and subsequent post warranty maintenance terms

Reference – ITSO Redpaper REDP-4655
IBM BladeCenter PS700, PS701, and PS702 Technical Overview and Introduction

IBM BladeCenter PS701 (8406-71Y) – Inside POWER7 blade

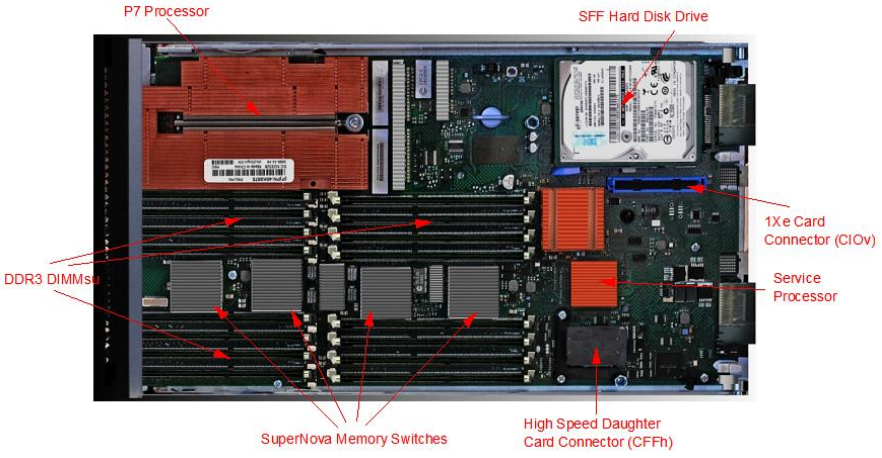

POWER7 Firmware

- **PowerVM Enterprise Edition is the hypervisor**
  - pHyp 2.1
  - VIOS 2.1.3
  - Extensions for configuration systems management
    - Hardware setup
    - FFDC
    - Heartbeat
    - PPM daemon
- **Packaged and managed as firmware on Support Element**
  - Loaded automatically at power on
  - Serviced via System z MCL process
  - Full set of operational controls
- **Managed as a closed system from the Hardware Management Console (HMC)**
  - No AMM access
  - No VIOS access, so no VIOS add-ons supported. (e.g., ITM hypervisor agent, or PCM drivers)
  - Enables creation and management of virtual servers via the Unified Resource Manager

POWER7 Configuration

- **CPU (8-core PS701)**
  - Each CPU has capacity of 1.0 processing units (PUs)
    - VIOS default capacity is 0.1 PUs per CPU (8 shared virtual CPUs on PS701)
    - Each dedicated CPU requires 1.0 PUs
  - Each shared virtual CPU requires at least 0.1 PUs
  - Up to 7.2 PUs across 8 shared CPUs or up to 7 dedicated CPUs (one for VIOS)
  - No over-commitment of PUs
- **Memory**
  - No over-commitment.
  - VIOS requires at least 4GB
- **Network**
  - Two (primary/backup) network adapters on intraensemble data network (IEDN) and customer network
  - Two (primary/backup) network adapters on intranode management network (INMN)
  - All networks IPv6-enabled
- **Storage**
  - Local blade disk dedicated to hypervisor
  - Two FCP adapters for SAN volume access, set up for multi-pathing
  - No hypervisor data on SAN volumes
- **ISO Images**
  - Customer ISO images stored on local blade disk

POWER7 Virtual Servers

- Shared or dedicated virtual CPUs
- Dedicated memory
- Network
  - Virtual network adapters for IEDN and customer network
  - Virtual server assigned 40-bit MAC ID range from ensemble MAC ID pool
  - Virtual network adapter assigned 48bit MAC ID from virtual server's MAC ID range
- Storage
  - External disks are virtual SCSI devices mapped to dedicated FC LUNs
- Virtual Optical Device
  - Customer ISO images (stored on local blade disk) can be associated with virtual server
- Boot Device
  - Virtual disk
  - Virtual optical device
  - Virtual network adapter
- Operating Systems
  - AIX 5.3 (Technology Level 12) and later running in POWER6 and 6+ compatibility mode
  - AIX 6.1 (Technology Level 5) and later
- Platform Performance Management
  - Guest Performance Management Provider (GPMP) can be installed and run in guest OS
- Applications
  - Certified for AIX, POWER7 and PowerVM Enterprise Edition

List of Storage Devices Supported by PS701 IBM BladeCenter® Express in IBM zEnterprise™ System*

- **IBM**
  - DS3400, DS3500, DS3950
  - DS4100, DS4200, DS4700, DS4800
  - DS5020, DS5100, DS5300
  - DS6000™
  - DS8100, DS8300, DS8700
  - SVC 2145
  - XIV®
  - Shark 2107
- **EMC**
  - Symmetrix 3300, 5000 , 8000
  - Symmetrix V-MAX
  - DMX 800, 1000, 1000P, 2000, 2000P, 3000
  - DMX -3, DMX-4
  - Clariion CX3 all models, CX4 all models, CX300, CX400, CX500, CX600, CX700, AX4-5
- **Hitachi**
  - Lightning 9910, 9960, 9970, 9980
  - USP 100, 600, 1100
  - NSC55
  - USP V
- **HP**
  - EVA 3, 4X, 5, 6X, 8X
  - XP 10K, 12K, 48, 128, 512,1024

* - default MPIO Path Control Module support only

http://www-03.ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start_over=yes

WebSphere® DataPower® Appliance*

WebSphere DataPower[1] Appliance in the zBX

Purpose-built hardware for simplified deployment and hardened security

What is it?

The IBM WebSphere DataPower appliance (SOD)[1] integrated in the zEnterprise System, can help simplify, govern, and enhance the security of XML and IT services by providing connectivity, gateway functions, data transformation, protocol bridging, and intelligent load distribution.

How is it different?

- **Security:** VLAN support provides enforced isolation of network traffic with secure private networks. And integration with RACF® security.
- **Improved support:** Monitoring of hardware with "call home" for current/expected problems and support by System z Service Support Representative.
- **System z packaging:** Increased quality with pre-testing of blade and zBX. Upgrade history available to ease growth. Guided placement of blades to optimize.
- **Operational controls:** Monitoring rolled into System z environment from single console. Time synchronization with System z. Consistent change management with Unified Resource Manager.

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

WebSphere DataPower Appliance*

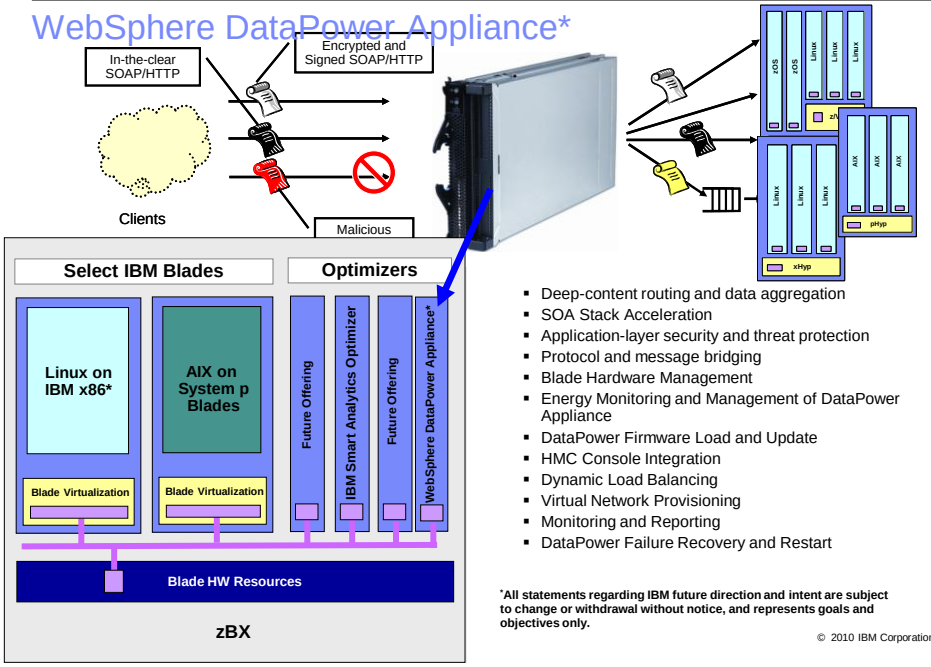

- Deep-content routing and data aggregation
- SOA Stack Acceleration
- Application-layer security and threat protection
- Protocol and message bridging
- Blade Hardware Management
- Energy Monitoring and Management of DataPower Appliance
- DataPower Firmware Load and Update
- HMC Console Integration
- Dynamic Load Balancing
- Virtual Network Provisioning
- Monitoring and Reporting
- DataPower Failure Recovery and Restart

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

zBX Networking and Connectivity

z196 – What are the INMN, IEDN and Customer networks

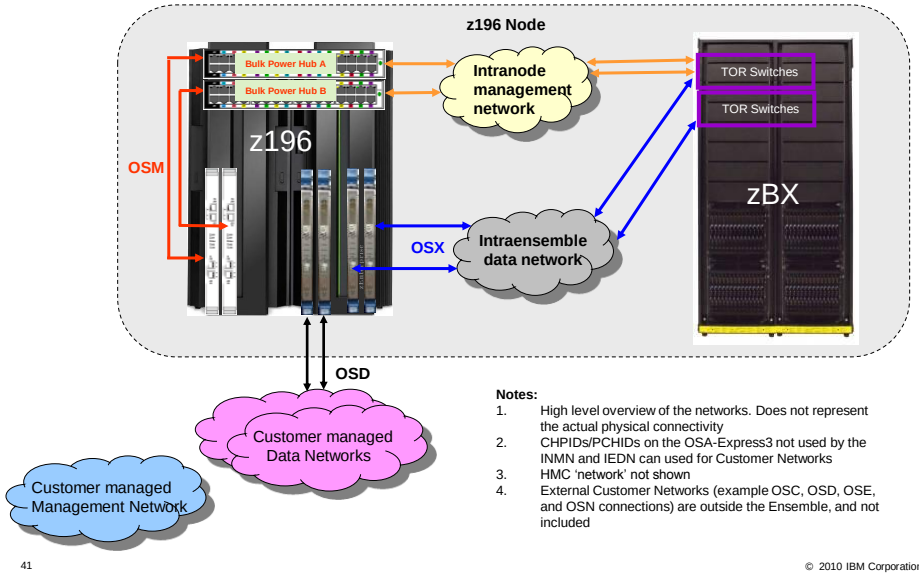

Notes:

1. High level overview of the networks. Does not represent the actual physical connectivity
2. CHPIDs/PCHIDs on the OSA-Express3 not used by the INMN and IEDN can used for Customer Networks
3. HMC 'network' not shown
4. External Customer Networks (example OSC, OSD, OSE, and OSN connections) are outside the Ensemble, and not included

OSA Express3 10GbE 2458-002 Configurations / Connections

- **Each 2817 CPC must have 2 OSA-Express3 10GbE ports connected to the 2458 for the intraensemble data network (IEDN)**
  - The 2 ports provide a redundant configuration for failover purposes in case one link fails
  - For availability, the two ports should be spread across 2 different OSA-Express3 10 GbE cards/features within the same CPC, but in different I/O domains if possible.
  - The IEDN is exclusively for data traffic between the 2817 and 2458. The 2458 can have a maximum of 16 physical IEDN connections (i.e. 8 pairs of OSA-Express3 10 GbE ports)
- **The 2458 should be installed adjacent to the controlling z196 CPC for service requirements**
  - Use of 2 OSA-Express3 10 GbE SR (Short Reach) features (FC3371) or 2 OSA-Express3 10 GbE LR (Long Reach) features (FC3370) cabled directly from the controlling CPC to the 2458 is the **recommended** configuration option to protect the closed physical security of the IEDN
  - Blade Center HSSM supports SR optics from the TOR switch
  - **Blade Center HSSM uses DAC (Direct Attach Cables) from the TOR IEDN switches**

Continued

OSA Express3 10GbE 2458-002 Configurations / Connections

- **The pairs of connections between the CPC's and the 2458 can be either SR/LR (direct connect).**
- **It's the customers responsibility to supply all the cables for the IEDN (z196 OSA-Express3 10 GbE to 2458-002 TOR Switch)**
- **Configuration definitions and usage**
  - OSA Express3 10 GbE ports (CHPID) can be defined in the IOCDS as SHARED or DEDICATED
    - SHARED between all or selected LPARS on the same CPC is allowed and recommended
    - SHARED can be restricted by the PARTITION keyword in the CHPID statement to allow for a subset of LPARs on the CPC to use the same 2458-002
    - The 2 OSA Express3 10 GbE port pairs defined as SHARED can be shared by all the DB2 member LPARs on the same z196 CPC using the 2458-002.
    - DEDICATED, restricts the 2 OSA Express3 10 GbE ports to a single LPAR
    - A combination of SHARED and DEDICATED OSA Express3 10 GbE port pairs can be defined within the maximum of 16 physical connections supported by the 2458-002

Ensemble networking

1. **Intranode management network (OSM)**
   - 2 ports from 2 different OSA Express-3 1000BASE-T Ethernet adapters, for redundancy.
   - To allow the HMC, to talk to the System z hypervisors, within the ensemble.

2. **Intraensemble data network (OSX)**
   - A pair of OSA Express-3 10 GbE adapters, for redundancy.
   - To allow the System z applications to communicate between OS images to share data.
   - To allow the System z application to communicate to the zBX
   - Ensemble zBX to zBX communications.

3. **Existing customer network**
   - A pair of 10 GbE connections in the zBX TOR Switch
   - For CPC's or switches not in the ensemble

1 intranode management network
2 intraensemble data network
3 existing customer network

OSA-Express3 CHPID Types

- **Two new OSA CHPID types are created to support the new zBX networks.**
- **There are now up to 6 types of NETWORK OSA CHPID's.**

3

- **Existing data networks - defined as OSC, OSD, OSE, and OSN CHPID's**
  - Existing customer provided and managed OSA ports used for access to the current customer external networks. (no changes)

1. **Intranode management network – defined as OSM CHPID's**
   - OSA-Express3 1000base-T
     - Configured as an OSM CHPID port for node management network to be connect to zEnterprise Ensemble CPC via a new ethernet switch A/B J07.

2. **Intraensemble data networks – defined as OSX CHPID's**
   - OSA-Express3 10 GbE (LR or SR) configured as an OSX CHPID, fiber port for IEDN.

1 intranode management network
2 intraensemble data network
3 existing customer network

CHPID Types OSX and OSM

1 OSM (INMN)

FC 3367
OSA-Express3 1000BASE-T Ethernet

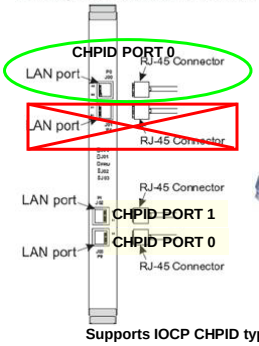

OSA Express3
1GbE
2 CHPIDS
2 PORTS/CHPID

FC3367

CAT 6 ETH CABLE

MUST USE
CHPID PORT 0

Supports IOCP CHPID types:
OSC, OSD, OSE, OSN, and OSM (ONLY 1000BASE-T).

2 OSX (IEDN)

OSA Express3
10 GbE
2 CHPIDS
1 PORT/CHPID

IEDN Distances
- SM (Long Range Optics) 10 km (6.2 miles)
- MM (Short Range Optics)
  50 micron at 2000 MHz-km: 300 meters (984')
  50 micron at 500 MHz-km: 82 meters (269')
  62.5 micron at 200 MHz-km: 33 meters (108')

FC3370 (LR)
Single Mode 9 micron LC duplex

FC3371 (SR)
Multi Mode 50/62.5 micron LC duplex

Supports IOCP CHPID types:
OSD and OSX (ONLY 10 GbE).

OSM IOCDS EXAMPLE:

```
• CHPID PCHID=191,PATH=(CSS(0,1,2,3),23),TYPE=OSM,CHPARM=01,SHARED,   …
• CNTLUNIT CUNUMBR=0910,PATH=((CSS(0),23)),UNIT=OSM
• IODEVICE ADDRESS=(0910,15),CUNUMBR=(0910),UNIT=OSA,UNITADD=00,      *
           MODEL=M,DYNAMIC=YES,LOCANY=YES
```

OSX IOCDS EXAMPLE:

```
•CHPID PCHID=5E1,PATH=(CSS(0,1,2,3),2F),TYPE=OSX,SHARED,   ….
•CNTLUNIT CUNUMBR=09F0,PATH=((CSS(0),2F)),UNIT=OSX
• IODEVICE ADDRESS=(09F0,15),CUNUMBR=(09F0),UNIT=OSA,UNITADD=00,      *
           MODEL=X,DYNAMIC=YES,LOCANY=YES
```


CHPARM=01 indicates that the channel path is managed by Dynamic Channel Path Mgmt (DCM)

2458-002 Intranode management network 1

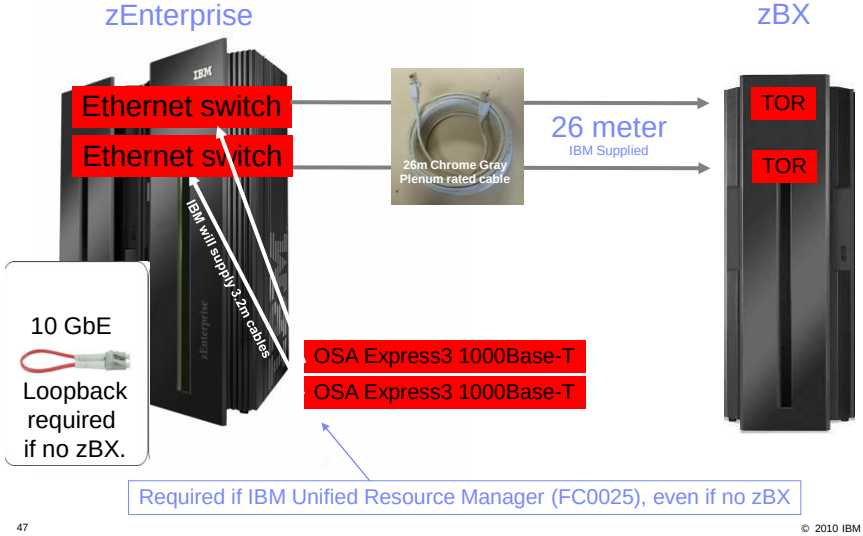

Required if IBM Unified Resource Manager (FC0025), even if no zBX

zBX Top of Rack (TORs) for Ensemble Networking

1 **Intranode management network (INMN)**
- Uses both of the built-in 1GbE interfaces on blade for redundancy
- Each interface connects to separate 1Gb ESM on BladeCenter chasis.
- Link is run from 1Gb ESM to INMN TOR
- Allows the HMC to talk to the hypervisor on the blade. Virtual Servers cannot access the INMN.

2 **Intraensemble data network (IEDN)**
- Uses dual port 10GbE adapter (FC 8275) for redundancy
- Each port connects to 10Gb ESM on BladeCenter chassis
- Link is run from 10Gb ESM to TOR
- Allows virtual servers on each blade to share data with the other applications throughout the ensemble

3 **Existing customer network**
- Several ports on IEDN TOR reserved for external connections
- Isolated from IEDN by VLAN definitions.

2458-002 IEDN Redundant Top of Rack (TOR) Switch Connections

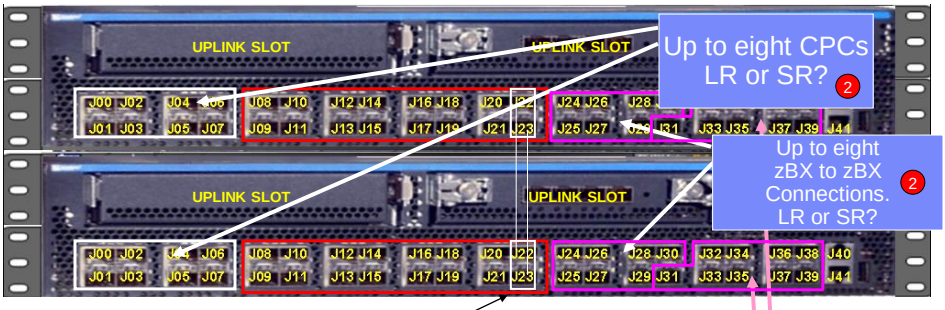

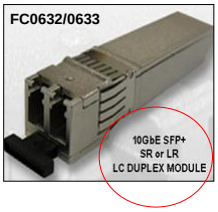

EX4500
- **SFP+ = 10GbE Optical SR or LR**
- **DAC = 10GbE Direct Attach Cables.**

Customer Network LR or SR? 3

- **SWITCH JACK PLUGGING RULES:**
- **J00 - J07 are SFP+ reserved for Host OSX IEDN connections.**
- **J08 - J23 are DAC reserved for BC IEDN, SM07/SM09 connections.**
- **J22 / J23 are 1 Meter DAC for Switch to Switch**
- **J24 - J30 are SFP+ reserved for zBX to zBX IEDN connections.**
- **J31 - J39 are SFP+ reserved for customer (PINK) IEDN connections.**
- **J40 Console Port**
- **J41 IEDN Switch Management Port**

Optic p/n 45W4743 - 10GE sfp+ SR - has a black handle
Optic p/n 45W4744 - 10GE sfp+ LR - has a blue handle

z196 / zBX Ensemble INMN & IEDN connections

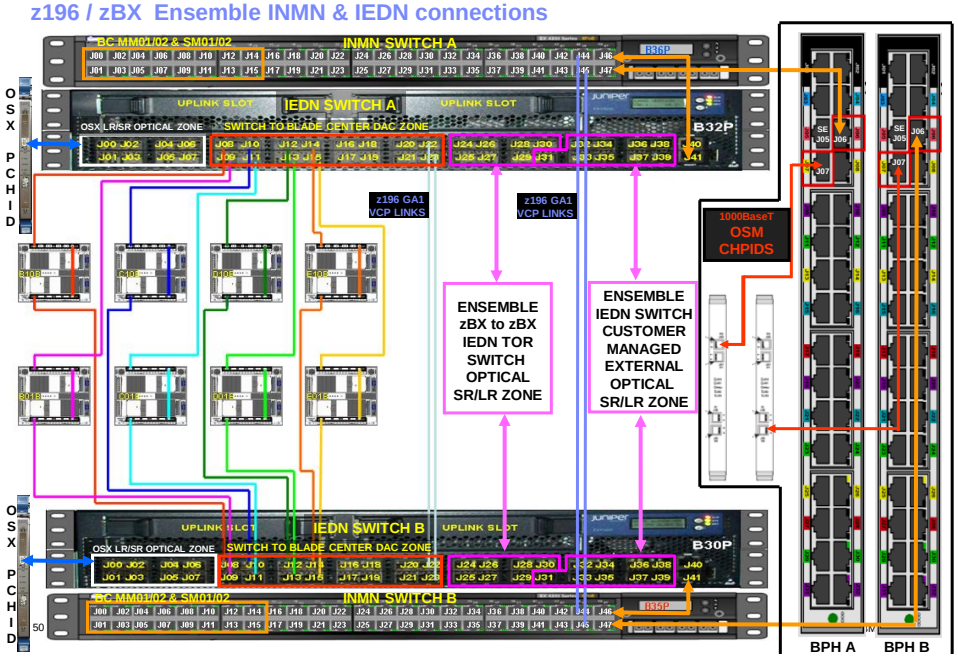

z196/2458-002 MAX CPC/NODE ENSEMBLE

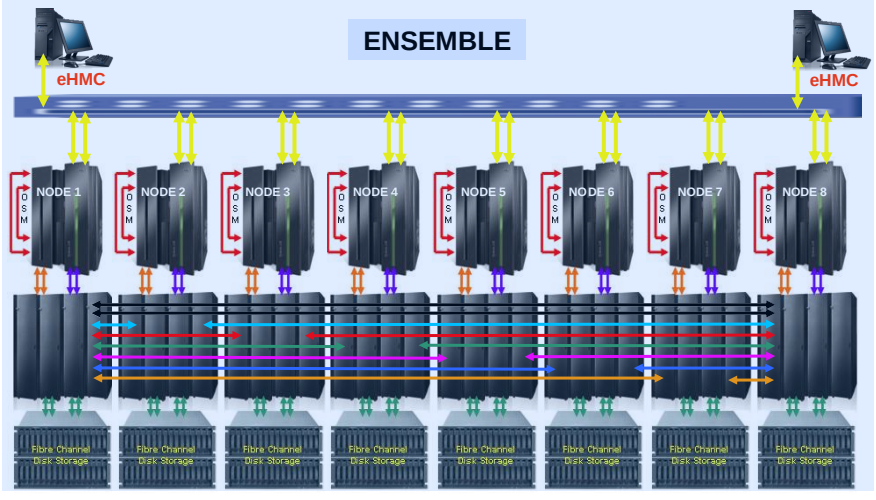

zBX to zBX IEDN Connections 2

1. Intra-Node Management Network
2. Intra-Ensemble Data Network
3. Existing Customer Network

Big Picture 2458-002 Configuration 8 CPC's / 8 PAIR OF SR/LR MIX IEDN CONNECTIONS.

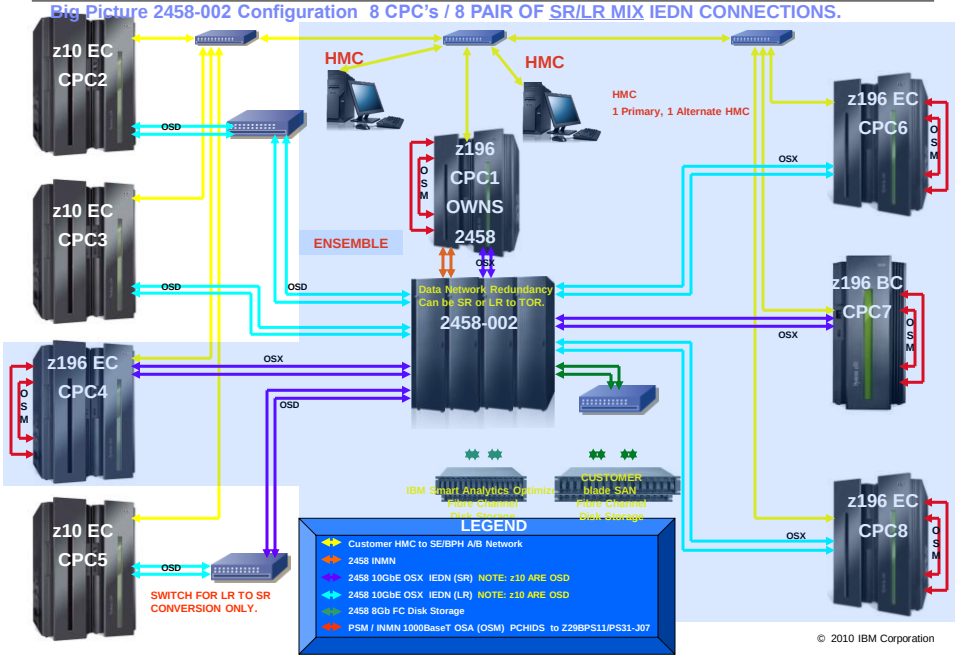

Redbooks IBM

zBX Configuration and Planning

Always Refer to the Installation Manual for Physical Planning for details:

M/T 2817 – GC28-6897
M/T 2458 Model 002 – GC28-2611

zBX – Full Configuration, front view

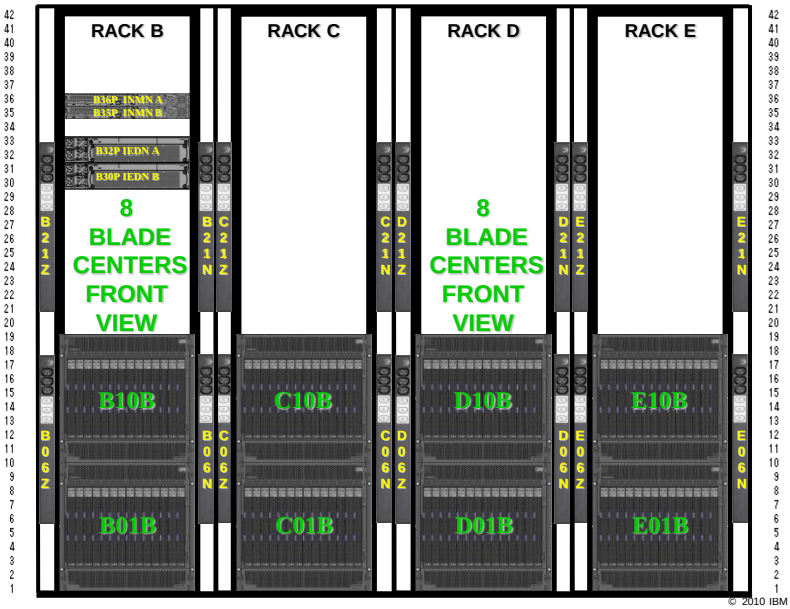

zBX – Full Configuration, rear view

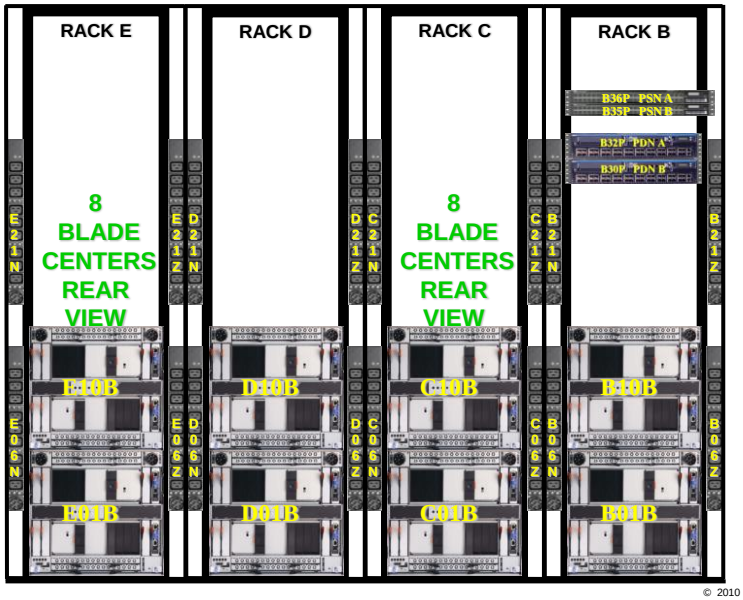

zBX Front Door Options

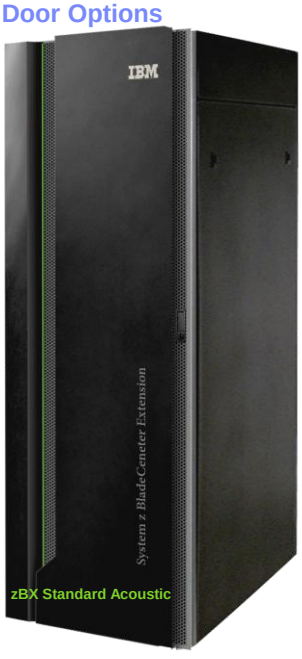

zBX Standard Acoustic

zBX Optional Acoustic

z196 and zBX Blade Physical Planning

- **Blade Configurations**
  - Ordered with z196 Configuration
  - IBM Smart Analytics Optimizer – Up to 56 blades (7, 14, 28, 42, 56)
    - Dedicated IBM DS5020 storage and SAN cables required to attach to the zBX must be ordered separately
  - POWER7 Blades – Up to 112 blades (0, 1 to 112 – Increment 1)
    - Quantity ordered on z196 provides BladeCenter space in the zBX frames
    - Supported Blades must be ordered separately by customer
    - Supported SAN or NAS attached storage and cables to attach to the zBX must be ordered separately
- **zBX BladeCenters – Shipped with zBX to support ordered blade configuration**
  - Installed and cabled by IBM Service
  - Up to 14 blades per BladeCenter
  - IBM Smart Analytics Optimizer blades are isolated in dedicated BladeCenters
  - POWER7 and IBM x86 blades share BladeCenters
  - N+1 configuration for power, SAN, and networks

zBX Rack Physical Planning and Power Planning

- **Standard 42U racks (height reduced shipping orderable)**
  - Installed and cabled by IBM Service
  - Must be within 25 m of the controlling z196
  - Configured to support the required BladeCenter configuration
  - Include redundant top of rack network switches for management, data, and external networks
  - Support two (or one) BladeCenters per rack
  - Include intra-zBX cables for power, data network, management network, external network and SAN required to connect rack to rack and BladeCenter to BladeCenter. Cables provide redundant connectivity to N+1 components for high availability.
  - Include redundant power distribution units, two per BladeCenter in the rack
- **Power**
  - Separate from z196 CPC power, separate Power ON/OFF controlled by z196 HMC/SE
  - One cord required per PDU (included or customer provided depending on power feature ordered)
  - Two or four power services per rack (one per PDU)
    - Typical for USA: IEC 309 60A/208V/3PHASE Delta
    - Power consumed will be MUCH less than 2,900 watt rated power per PDU
- **Cooling**
  - Front to back air flow (cool aisle to hot aisle) just like the z196
  - Optional rear door heat exchanger

zBX Cooling

- **BladeCenter Cooling modules (per chassis)**
  - Two (1+1) hot-swap and redundant blowers standard
  - Additional fan packs on power supplies
- **The blower speeds vary depending on the ambient air temperature at the front of the BladeCenter unit and the temperature of internal BladeCenter components.**
  - If the ambient temperature is 25° C (77° F) or below, the BladeCenter unit blowers will run at their minimum rotational speed, increasing their speed as required to control internal BladeCenter temperature.
  - If the ambient temperature is above 25° C (77° F), the blowers will run faster, increasing their speed as required to control internal BladeCenter unit temperature.
- **If a blower fails, the remaining blower will run full speed and continues to cool the BladeCenter unit and blade servers.**
  - Failed blower must be replaced to restore cooling redundancy

zBX Optional Rear Door Heat eXchanger

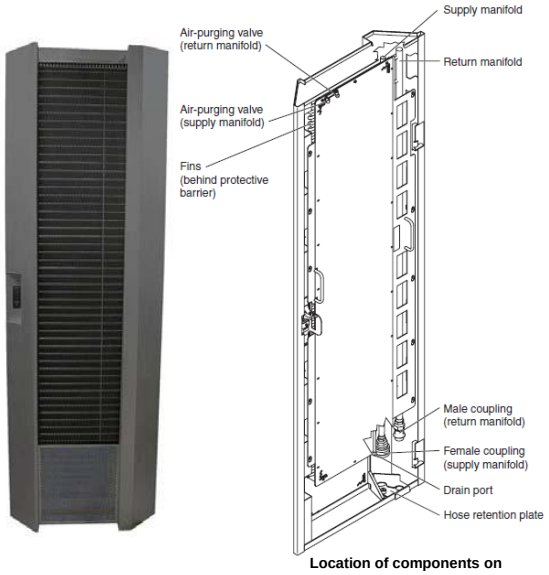

Location of components on the Rear Door Heat eXchanger.

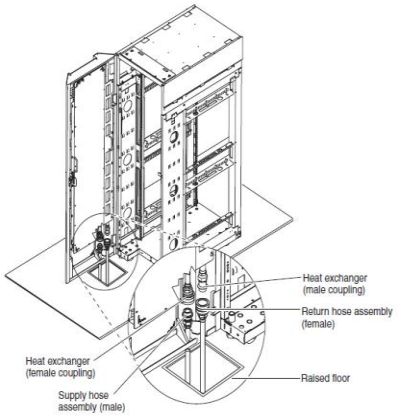

Raised-floor hose management example: Hoses exit through floor tile at the door hinge.

http://www-03.ibm.com/systems/x/hardware/options/cooling.html

z10 zBX Rear Door – options

- **Option 1 – Standard door**
- **Option 2 – Rear Door Heat Exchanger - Feature Code #0540**
  - Requires customer conditioned water
  - Refer to the Heat Exchanger Rear Door Planning Guide to ensure that the heat is properly dissipated. http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphad_p5/iphadexchangeroverview.htm.
  - When at that link, just search on 7014-T42
- **There are two circumstances which can be considered for the Rear Door Heat eXchanger (RDHX).**
  1. Order the RDHX as part of the initial order for the zBX.
  2. If not sure if an RDHX is needed, contact IBM Systems & Technology Group (STG) Lab Services.

  http://www.coolcentric.com/

There is no MES to change to or from a rear door heat exchanger.

- **Option 3 – Noise Reduction Rear Door – Feature Code #0543**

Rear Door Heat eXchanger (option)

zBX Power

- **Power not integrated with z196 power**
- **Required customer supplied outlets**
- **Electrical input**
  - Single-phase
    - Input voltage:
      - Minimum: 200 V ac
      - Maximum: 240 V ac
  - 2 or 4 PDUs per RACK w/ 2-4 line cords
    - Independent line cords 4.3m length
  - Up to 4 hot swap & redundant power supplies (2980w each)
- **For z196:**
  - Individual blade power can be logically tied to system power or independent based on Menu selection at the SE/HMC

zBX Power

- **BladeCenter power is not provided from the System z in z196**
  - Power comes off wall power, and it should generally be always on (Standby Power On state), **but…….**
  - Blade power is associated with the System z power
    - If the System z is powered on, all blades will be powered on
    - If the System z is powered off, those same blades may be powered off
    - Single System z system owns/manages zBX hardware in a shared CPC configuration
    - If more than one System z196 is using the zBX, a menu option is available to allow the controlling z196 to power off individual blades
- **Individual blades have their power controlled (power off) by the SE/HMC of the attached System z primarily for service reasons**

zBX Power Options

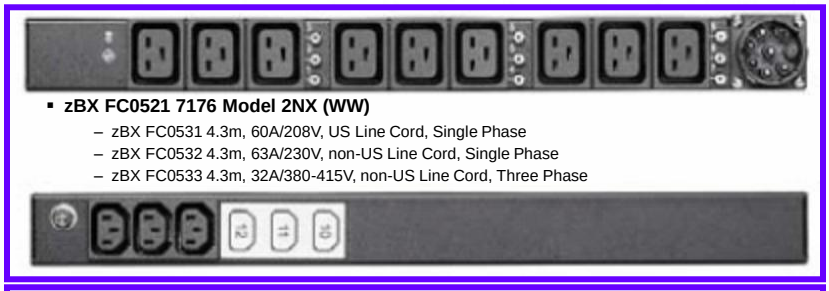

- **zBX FC0521 7176 Model 2NX (WW)**
  - zBX FC0531 4.3m, 60A/208V, US Line Cord, Single Phase
  - zBX FC0532 4.3m, 63A/230V, non-US Line Cord, Single Phase
  - zBX FC0533 4.3m, 32A/380-415V, non-US Line Cord, Three Phase

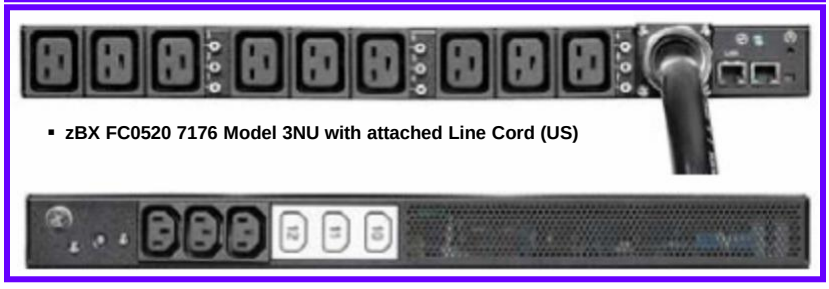

- **zBX FC0520 7176 Model 3NU with attached Line Cord (US)**

zBX Power and Cooling – Summary

- **BladeCenters each operate from two fully-redundant power supplies**
- **These redundant PDUs each have their own line cords, allowing the system to survive the loss of customer power to either line cord**
- **If power is interrupted to one of the PDUs, the other PDU will pick up the entire load and the BladeCenter will continue to operate without interruption**
- **Therefore the line cord(s) for each PDU must be wired to support the entire power load of the system.**

Always refer to zBX Installation Manual for Physical Planning

z196 / zBX Basic Ensemble HMC/SE BPH connections per CPC/Node

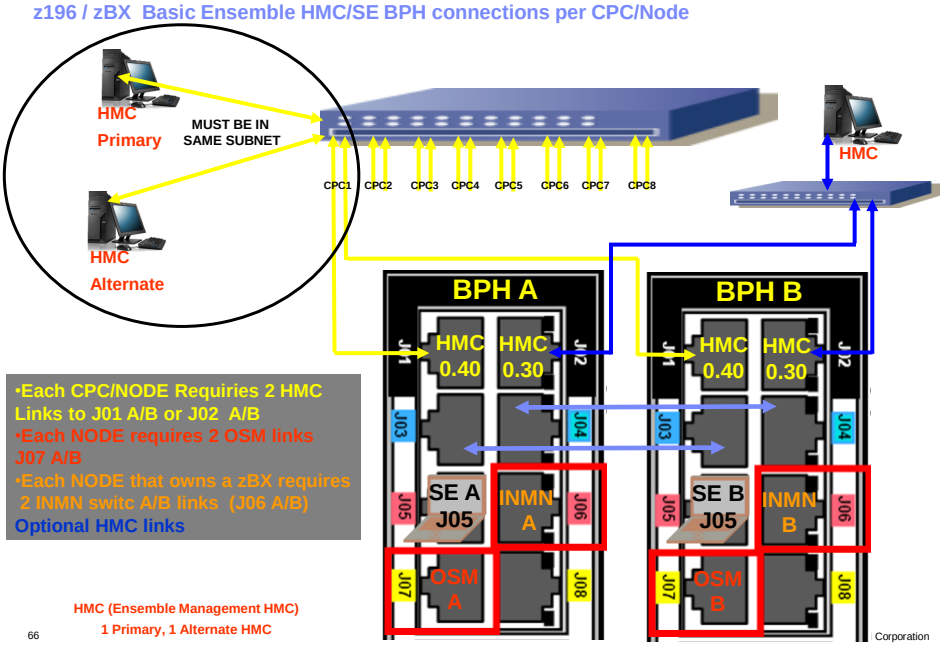

- Each CPC/NODE Requiries 2 HMC Links to J01 A/B or J02 A/B
- Each NODE requires 2 OSM links J07 A/B
- Each NODE that owns a zBX requires 2 INMN switc A/B links (J06 A/B)

Optional HMC links

HMC (Ensemble Management HMC)
1 Primary, 1 Alternate HMC

The zEnterprise – Ensembles and Nodes

Ensemble - a group of CPC/Nodes **(Maximum 8 CPC's in an Ensemble)**
Node - a CPC with/without a zBX

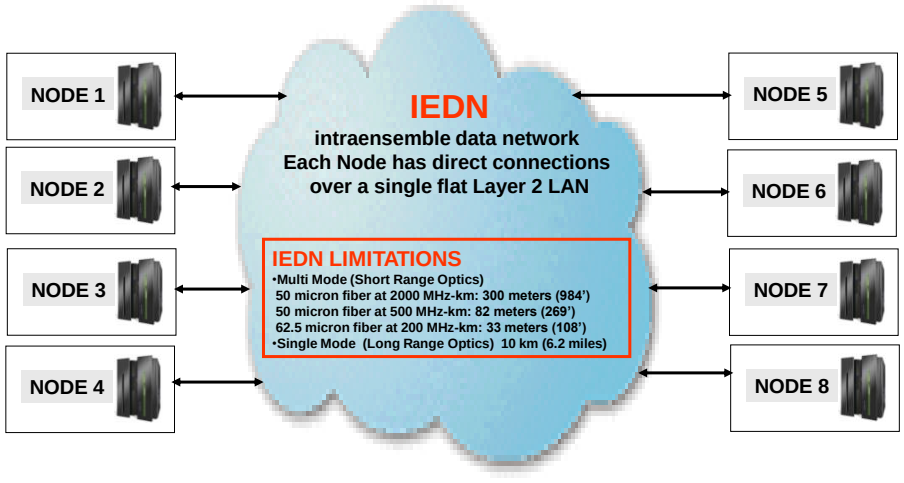

z196/2458-002 SINGLE CPC ENSEMBLE EXAMPLE

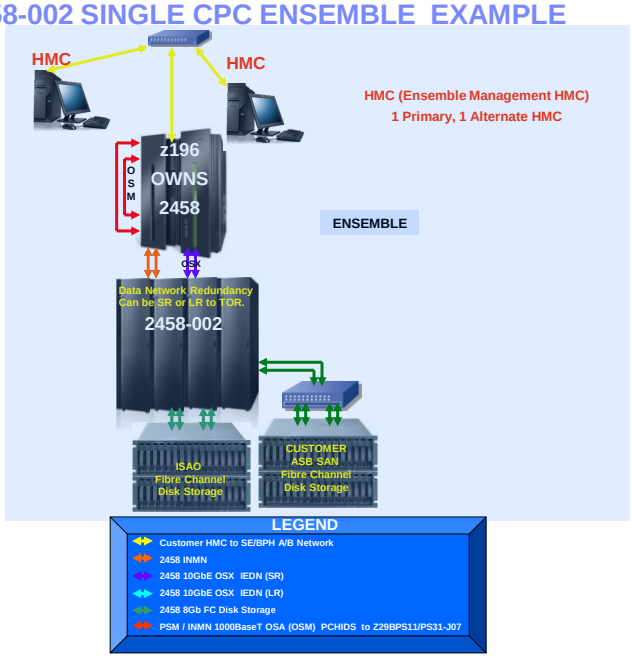

z196/2458-002 DUAL CPC ENSEMBLE EXAMPLE

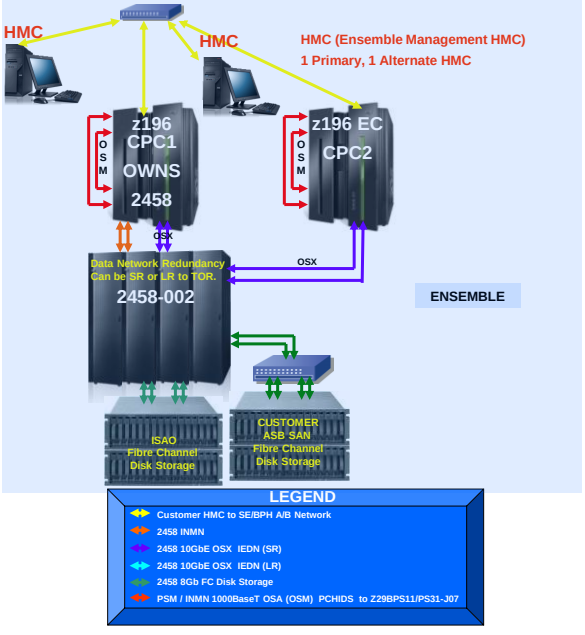

z196/2458-002 CO-LOCATED 2 CPC ENSEMBLE EXAMPLE

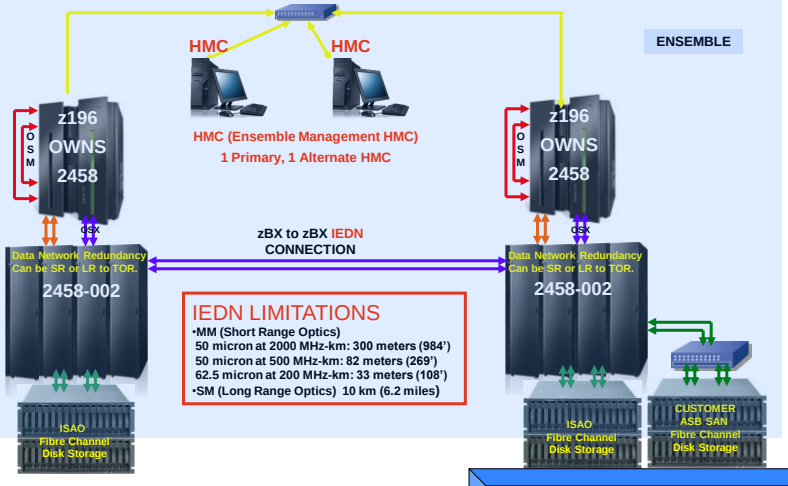

The above illustrates the network connectivity for two co-located zHybrid nodes.
The main point of this figure is to illustrate the fact that the TOR switches within each Node are connected together on the same physical layer 2 LAN.
The connection between the nodes is accomplished with fiber cable (between TOR switches) with no external switches.
This imposes a restriction on the limits of physical proximity of the two nodes.
The two nodes reside on the same physical LAN.
All IP addresses must come from the same IP subnet (or subnets when multiple VLANs are used).

z196/2458-002 2-ENSEMBLE EXAMPLE

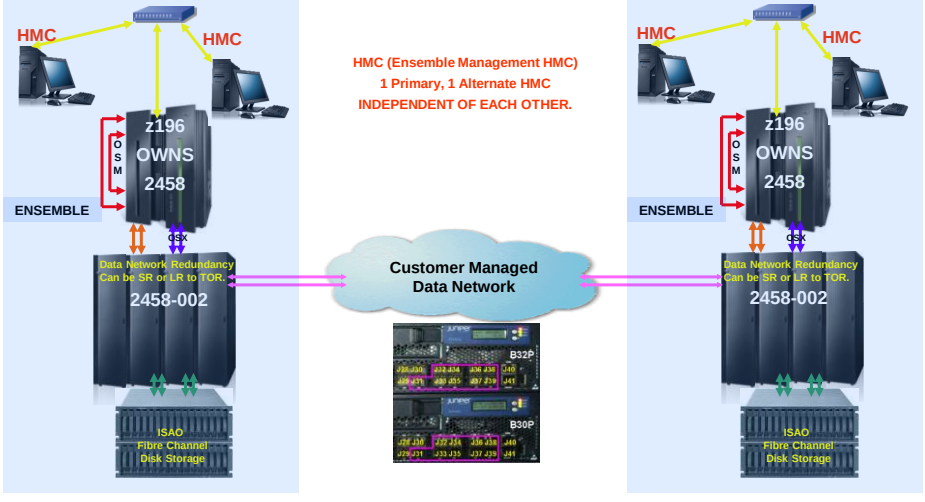

Big Picture Ensemble (EXAMPLE) mix of z196 zBX & z10, Config. SR/LR MIX IEDN CONNECTIONS.

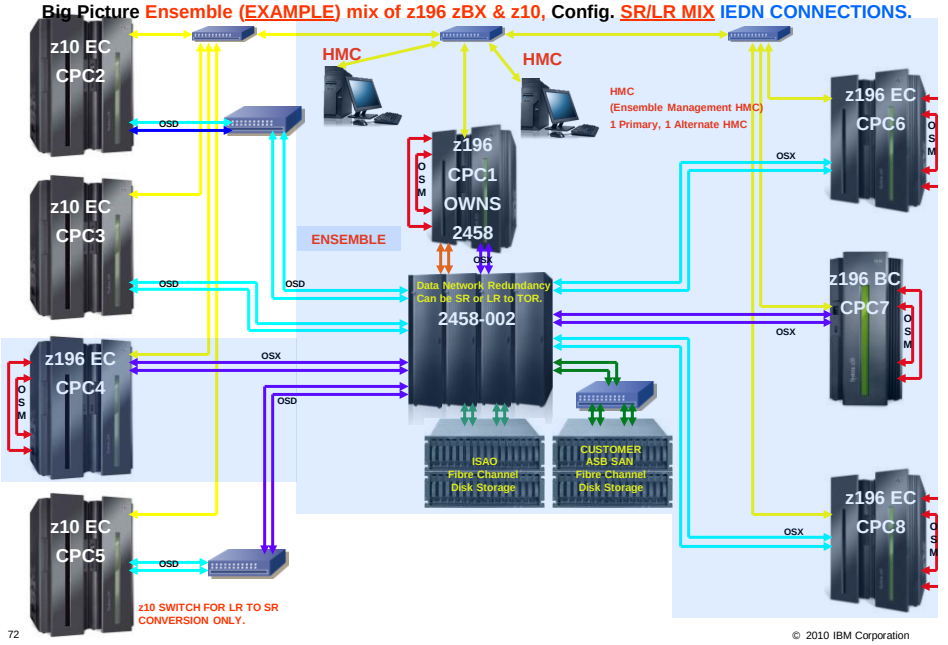

z196 INMN/IEDN Network, Big Picture ENSEMBLE, ZERO 2458-002, 1 CPC's & NO Switches.

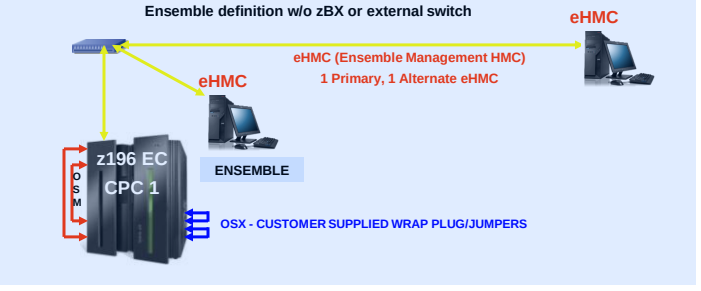

Special Case1 CEC only 0 zBX

zOS/zVM/zLinux Ensemble participation without the need for ZBX.

Dual 1000BaseT to BPH for Management Comm between HMC/SE to O/S

10G OSA (OSX) - Customer provided wrap plug/jumpers.

Big Picture Ensemble (EXAMPLE) mix of zBX, & NO zBX Config. SR/LR MIX IEDN CONNECTIONS.

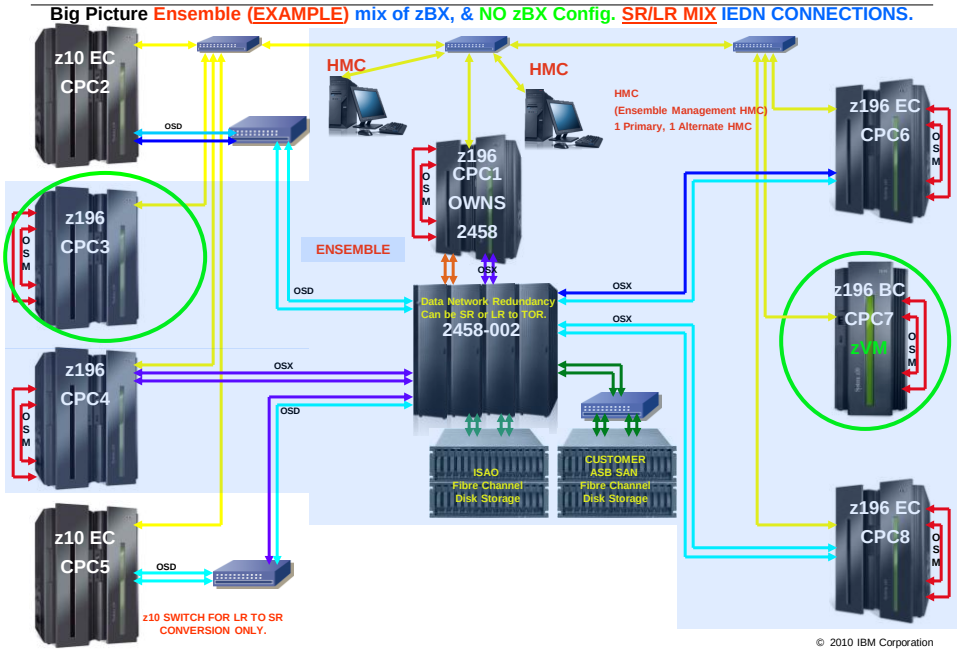

Redbooks IBM

Storage for IBM Smart Analytics Optimizer and POWER7 Blades

IBM Smart Analytics Optimizer Disk Attachment Details

- **Includes two 20 port – 8 Gb FC switches in each BladeCenter to allow connectivity to disk**
  - Must be directly attached
  - Supports 8 Gbps, 4 Gbps, 2 Gbps
    - 1 Gbps is NOT supported
  - Allows for connectivity to:
    - DS5020 with 1 TB HDD

16 Drives	16 Drives	16 Drives	32 Drives	32 Drives
4 ports	4 ports	8 ports	8 ports	8 ports

- **Disk is not part of the integrated Smart Analytics Optimizer offering**
  - Customer is responsible for:
    - supplying disk (separate order)
    - disk cabling
    - disk configuring

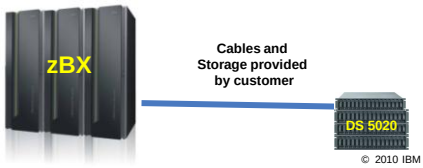

zBX 4 BLADECENTER FC CONNECTIONS

4 BLADECENTER, Smart Analytics Optimizer ONLY FC DISK STORAGE CASCADE CONNECTIONS

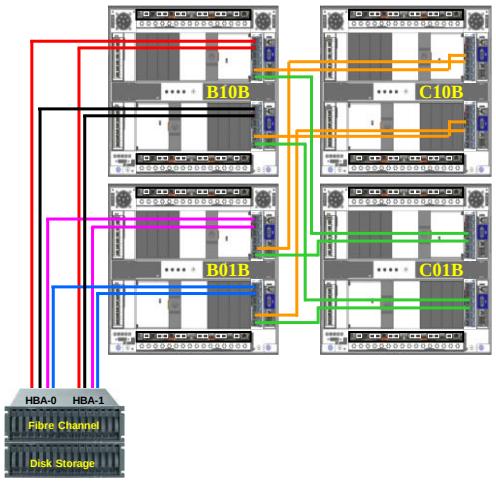

SM03/04 JACK LOCATIONS

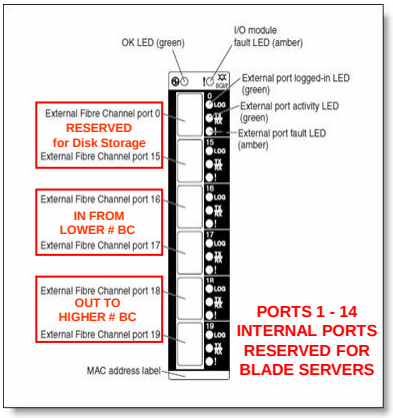

1-2 BladeCenter, IBM Smart Analytics Optimizer ONLY 8Gb FC Disk Storage cross-connect fabric

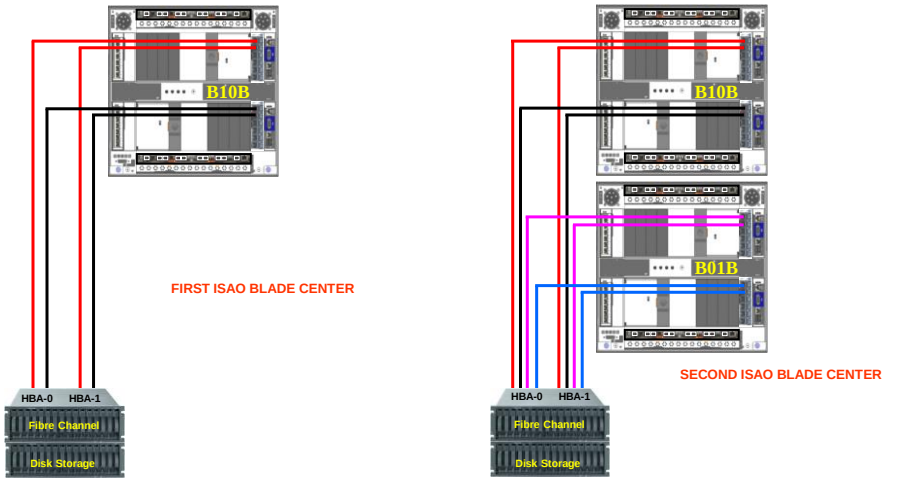

3-4 BladeCenter, IBM Smart Analytics Optimizer ONLY 8Gb FC Disk Storage cross-connect fabric

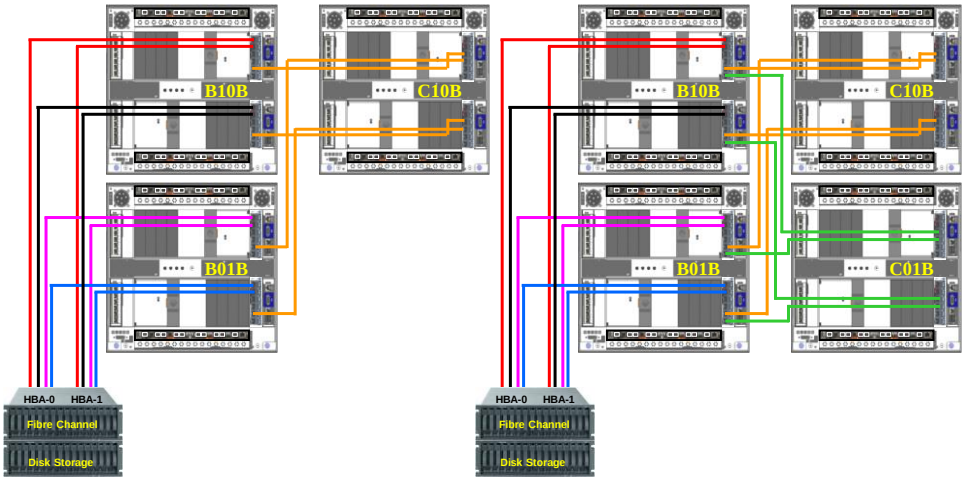

ISAO 4BC FC DISK STORAGE CROSS CONNECT FABRIC & IBM Blades/DP MIX EXAMPLE #1

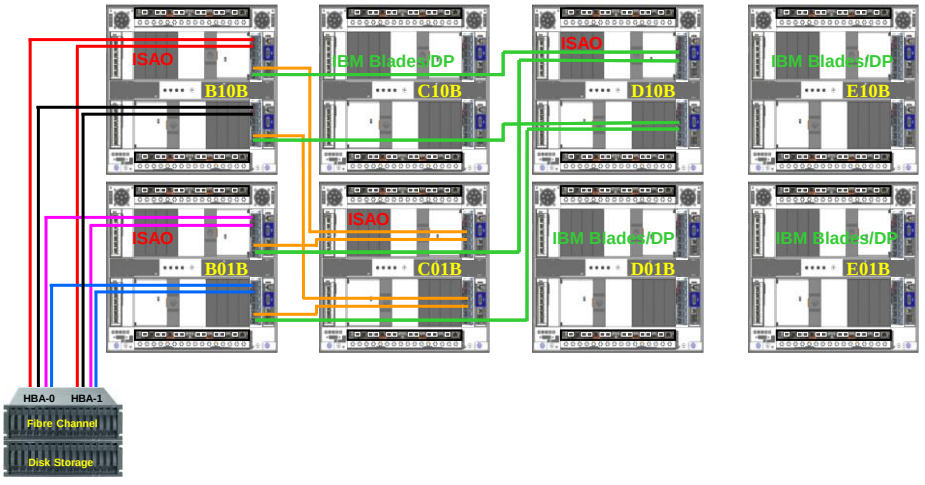

IBM Blades/DP chassis DOES NOT cable to the same Disk Storage as ISAO chassis

ISAO 4BC FC DISK STORAGE CROSS CONNECT FABRIC & IBM Blades/DP MIX EXAMPLE #2

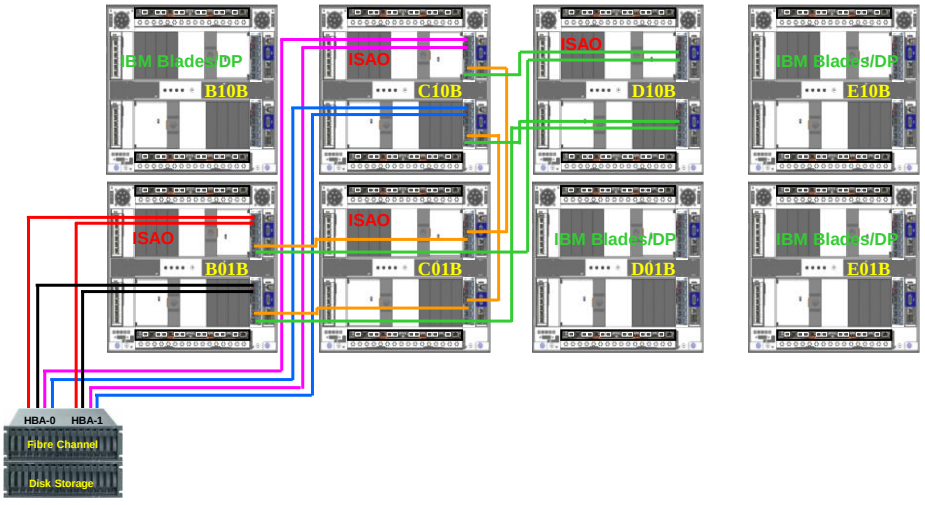

IBM Blades/DP chassis DOES NOT cable to the same Disk Storage as ISAO chassis

ISAO 4BC FC DISK STORAGE CROSS CONNECT FABRIC & IBM Blades/DP MIX EXAMPLE #3

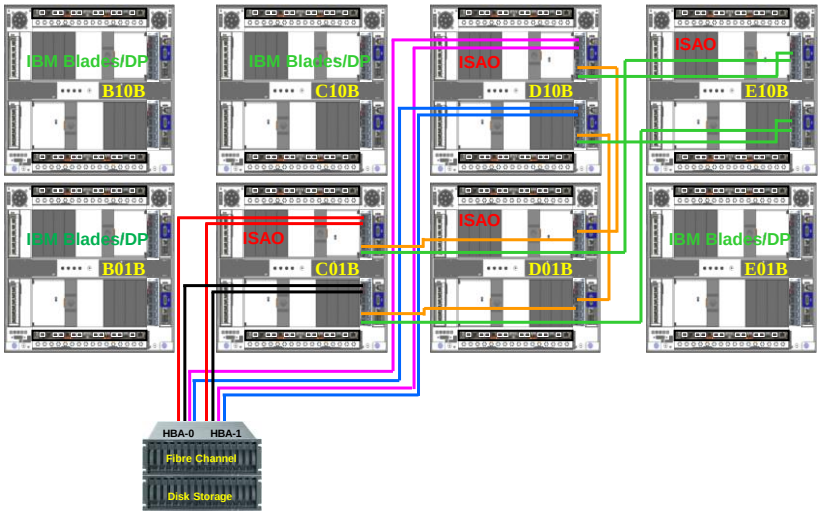

IBM Blades/DP chassis DOES NOT cable to the same Disk Storage as ISAO chassis

ISAO 4BC FC DISK STORAGE CROSS CONNECT FABRIC & IBM Blades/DP MIX EXAMPLE #4

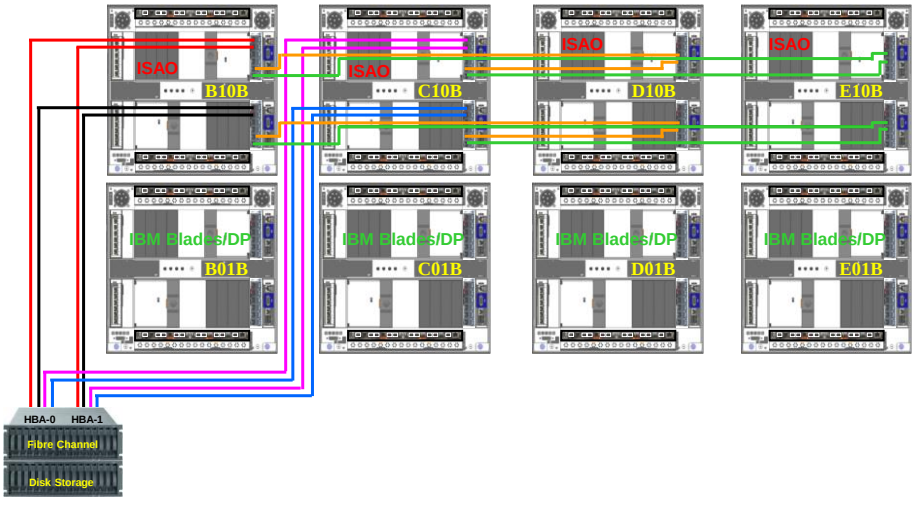

IBM Blades/DP chassis DOES NOT cable to the same Disk Storage as ISAO chassis

ISAO 4BC FC DISK STORAGE CROSS CONNECT FABRIC & IBM Blades/DP MIX EXAMPLE #5

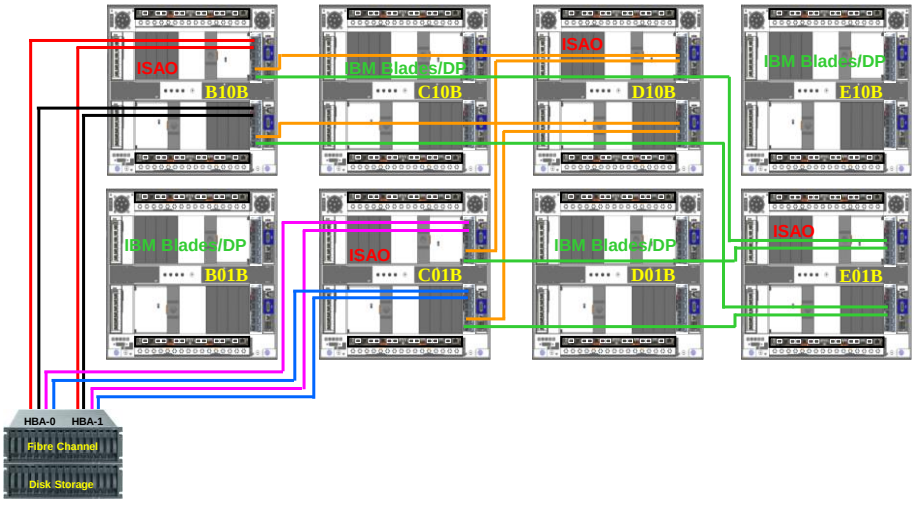

IBM Blades/DP chassis DOES NOT cable to the same Disk Storage as ISAO chassis

FC DISK STORAGE CONNECTIONS FOR ASB/DP ONLY EXAMPLE ONLY

IBM Blades/DP chassis DOES NOT cable to the same Disk Storage as ISAO chassis

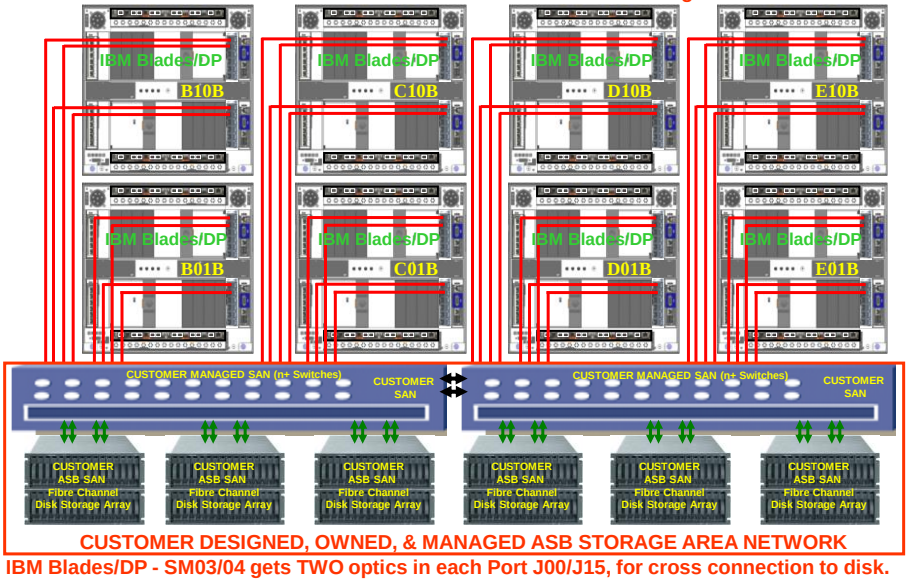

CUSTOMER DESIGNED, OWNED, & MANAGED ASB STORAGE AREA NETWORK

IBM Blades/DP - SM03/04 gets TWO optics in each Port J00/J15, for cross connection to disk.

CUSTOMER provides all ASB SAN H/W & SAN ZONING. No BladeCenter interconnect.

Hardware Management Consoles

zEnterprise and HMC's

New HMC

- New HMC feature Code 0091, New Switch feature code 0070
- Additional HMC's required for Unified Resource Manager and z196 zBX (if installed)
  - Alternate HMC used for Unfied Resource Manager is allocated for backup purposes only, cannot be used for daily HMC activities. Consider the need for additional HMC's (command center, computer room, etc).
- Can have a mix of traditional and Ensemble HMC's

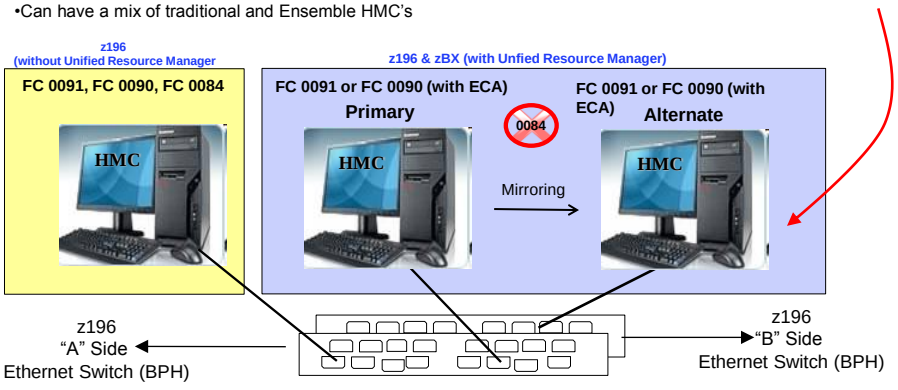

New Switch FC 0070, multiple switches required for redundancy

Primary and Alternate Hardware Management Consoles

- **Any V2.11.0 HMC can become the Primary HMC that controls the ensemble.**
  - The Primary HMC can perform all non-ensemble HMC functions on CPCs that aren't members of the ensemble.
- **The HMC that creates an ensemble (the HMC that performed the "Create Ensemble" wizard) becomes the Primary HMC.**
  - Function will be available in November 2010 when the zBX becomes available.
- **The Alternate HMC is specified when executing the "Create Ensemble" wizard.**
  - Any V2.11.0 HMC is eligible to be an Alternate HMC after running the “Manage Alternate Hardware Management Console task”.
- **The title of Primary Hardware Management Console and Alternate Hardware Management Console will appear on the Login HMC panel and the title line once you are logged in.**
  - The default HMC titles will change to these titles when the ensemble is created.
  - The titles will revert back to the default if the ensemble is deleted.
- **A Primary HMC is the only HMC that can perform ensemble related management tasks (create virtual server, manage virtual networks, create workload ....)**
- **Disaster Recover site**
  - Initially, we do not have the capability for offsite backup of the ensemble-specific data.
  - Future enhancement

Dank u
Dutch

Merci
French

Спасибо
Russian

Gracias
Spanish

شكراً
Arabic

감사합니다
Korean

Tack så mycket
Swedish

धन्यवाद
Hindi

תודה רבה
Hebrew

Obrigado
Brazilian Portuguese

谢谢
Chinese

Dankon
Esperanto

Thank You

ありがとうございます
Japanese

Trugarez
Breton

Danke
German

Tak
Danish

Grazie
Italian

நன்றி
Tamil

děkuji
Czech

ขอบคุณ
Thai

go raibh maith agat
Gaelic