



SHARE
Technology • Connections • Results

Boston, Massachusetts



Parallel Sysplex InfiniBand & Aetna's Implementation Experience

Session 7503

Riaz Ahmad – IBM
George Handera - Aetna





Trademarks

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

APPN*	IBM System z	pSeries*
CICS*	IBM System z9	RACF*
DB2*	IBM logo*	Resource Link
DB2 Connect	IMS	RMF
e-business logo*	Infoprint*	S/390*
Enterprise Storage Server*	Language Environment*	S/390 Parallel Enterprise Server
ESCON*	MQSeries*	Sysplex Timer*
FICON	Multiprise*	TotalStorage*
FICON Express	NetView*	VM/ESA*
GDPS*	On demand business logo	VSE/ESA
Geographically Dispersed Parallel Sysplex	OS/2*	VTAM*
HiperSockets	OS/390*	WebSphere*
HyperSwap	Parallel Sysplex*	z/Architecture
IBM	POWER	z/OS*
IBM eServer	PR/SM	z/VM*
IBM @server	Processor Resource/Systems Manager	zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

INFINIBAND, InfiniBand Trade Association and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

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

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

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

Notes:

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

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

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

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice.

Consult your local IBM business contact for information on the product or services available in your area.

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

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

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

Agenda

- Parallel Sysplex InfiniBand (PSIFB) technology
- Aetna's PSIFB Implementation Experience

IMPORTANT

The System z10 will be the last server to support ICB-4 links.

IBM intends not to offer Integrated Cluster Bus-4 (ICB-4) links on future servers as originally stated in US Hardware Announcement 108-154, dated February 26, 2008.

The IBM zEnterprise System



zEnterprise.

A New Dimension in Computing


Parallel Sysplex InfiniBand (PSIFB) Coupling

ready for even the most demanding data sharing workloads

▪ Simplify Parallel Sysplex connectivity

Do more with less

- Can **share physical links** by defining multiple logical links (CHPIDs)
- Can **consolidate** multiple legacy links (ISC and/or ICB)
- Easily address link constraints (e.g. define another CHPID to increase available subchannels instead of having to add physical links)



▪ More flexible placement of systems in a data center







- InfiniBand coupling links (FC 0163 and 0167) take advantage of optical cables **up to 150m long**. No longer restricted to only 7m between System z CECs when using these high performance links.
- InfiniBand coupling link Long Reach (LR FC 0168) features use the same 9 micron fiber cables as ISC-3 and FICON/FCP for **unrepeated distances of up to 10km**, and metropolitan distances with qualified DWDM solutions.

InfiniBand Glossary

Term	Description
Gbps	Gigabits per second
GBps	GigaBytes per second
1x	One “lane”, one pair of fibers
12x	12 “lanes”, 12 pairs of fiber
SDR	Single Data Rate – 2.5 Gbps per “lane” (0.25 GBps)
DDR	Double Data Rate – 5 Gbps per “lane” (0.5 GBps)
12x IB-SDR	12 “lanes” (pairs) for a total link data rate of 3 GBps, 150 meters point-to-point Used with OM3, 2000 MHz-k 50 micron multimode fiber optic cabling with MPO connectors
12x IB-DDR	12 “lanes” (pairs) for a total link data rate of 6 GBps, 150 meters point-to-point Used with OM3, 2000 MHz-k 50 micron multimode fiber optic cabling with MPO connectors
1x IB-SDR LR	One “lane” (one pair), 2.5 Gbps link data rate, unrepeated distance of 10 km Used with 9 micron single mode fiber optic cabling with LC Duplex connectors
1x IB-DDR LR	One “lane” (one pair), 5 Gbps link data rate, unrepeated distance of 10 km Used with 9 micron single mode fiber optic cabling with LC Duplex connectors

Note: The InfiniBand link data rate of 6 GBps or 3 GBps does 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. With InfiniBand coupling links, while the link data rate may be higher than that of ICB, the service times of coupling operations are greater, and the actual throughput may be less than with ICB links.

InfiniBand Architecture

- <http://www.infinibandta.org/home>
- InfiniBand is a communications link primarily used in high-performance computing. Its features include quality of service and failover, and it is designed to be scalable. The InfiniBand architecture specification defines a connection between processor nodes and high performance I/O nodes such as storage devices.
- Key requirements: high-bandwidth and low latency
- InfiniBand Trade Association (IBTA) founded in 1999
- Steering committee members
 -      
- z10 Implementation uses the IBTA defined Link and Physical layers only for:
 - InfiniBand (IFB) for host bus (CEC/CPC to I/O domain)
 - Parallel Sysplex over InfiniBand (PSIFB)



InfiniBand® is a registered trademark of the InfiniBand Trade Association (IBTA)

Overview

System z InfiniBand Implementation

- In z10, copper links used to connect I/O cage to book, optical links used to connect to other z10 or z9
- In z9, optical links used to connect to z10
 - **Cannot** use PSIFB to connect two z9s
- With 12x PSIFB links, all 12 lanes are used, even if only one CHPID is assigned to that link

Overview

PSIFB is a point-to-point architecture which

- Supports both optical (fiber) and electrical (copper) – Supports varying number of "physical lanes" per connecting media - 1, 4, 8, or 12 lanes per link.
 - Cable has 2 wires per lane, send and receive
- Supports multiple "virtual lanes" (up to 16 CHPIDs in System z terminology) over the same physical link.
- Supports varying bandwidths
 - Single data rate (250 MB/sec per physical lane)
 - Used between z9 and z10
 - Double data rate (500 MB/sec per physical lane)
 - Used between a pair of z10s, or within a z10
 - Quadruple data rate (1000MB/sec per physical lane)

z10 Coupling link options

Type	Description	Use	Link data rate	Distance	z10 BC/EC Maximum	z10 Maximum
PSIFB	1x IB-DDR LR	z10 to z10	5 Gbps	10 km unrepeated (6.2 miles) 100 km repeated	12*/32*	64 CHPIDs
PSIFB	12x IB-DDR	z10 to z10 z10 to z9	6 GBps 3 GBps**	150 meters (492 feet)***	12*/32*	
IC	Internal Coupling Channel	Internal communication	Internal speeds	NA	32/32	
ICB-4	Copper connection between OS and CF	z10, z9, z990, z890	2 GBps	10 meters *** (33 feet)	12/16	
ISC-3	Fiber connection between OS and CF	z10, z9, z990, z890	2 Gbps	10 km unrepeated (6.2 miles) 100 km repeated	48/48	

- The maximum number of Coupling Links combined cannot exceed 64 per server (PSIFB, ICB-4, ISC-3). There is a maximum of 64 Coupling CHPIDs (CIB, ICP, CBP, CFP) per server.
- For each MBA fanout installed for ICB-4s, the number of possible customer HCA fanouts is reduced by one
 - * Each link supports definition of multiple CIB CHPIDs, up to 16 per fanout
 - ** z10 negotiates to 3 GBps (12x IB-SDR) when connected to a System z9
 - *** 3 meters (10 feet) reserved for internal routing and strain relief

Note: The InfiniBand link data rates of 6 GBps, 3 GBps, 2.5 Gbps, or 5 Gbps 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. With InfiniBand coupling links, while the link data rate may be higher than that of ICB (12x IB-SDR or 12x IB-DDR) or ISC-3 (1x IB-SDR or 1x IB-DDR), the service times of coupling operations are greater, and the actual throughput may be less than with ICB links or ISC-3 links.

z196 coupling link options

Type	Description	Use	Link data rate	Distance	z196 Maximum	z196 Max links	z196 Max CHPIDs
IC (ICP)	Internal Coupling Channel	Internal communication	Internal speeds	NA	32	NA	128 CHPIDs
InfiniBand (CIB)	12x InfiniBand	z196 & z10 z196 & z10 to z9	6 GBps 3 GBps*	150 meters (492 feet)	32	32	
InfiniBand (CIB)	1x InfiniBand	z196 & z10	5 Gbps or 2.5 Gbps**	10 km unrepeated (6.2 miles) 100 km repeated	32		
ISC-3 (CFP)	InterSystem Channel-3	z196, z10, z9	2 Gbps	10 km unrepeated (6.2 miles) 100 km repeated	48	48	

- **Maximum of 16 InfiniBand fanouts are allowed, 2 links per fanout**
- **Maximum of 128 coupling CHPIDs (ICP, CIB, CFP) per server**
Each InfiniBand link supports definition of multiple CIB CHPIDs, up to 16 per fanout

* z196 & z10 negotiate to 3 GBps when connected to a z9


** May negotiate to 2.5 Gbps when connected to a DWDM

Note: ICB-4 is not supported on z196

Note: The InfiniBand link data rates of 6 GBps, 3 GBps, 2.5 Gbps, or 5 Gbps 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.

Coupling Link Choices - Overview

- **ISC (Inter-System Channel)**
 - ▶ Fiber optics
 - ▶ I/O Adapter card
 - ▶ 10km and longer distances with qualified WDM solutions
- **PSIFB (1x IB)**
 - ▶ Fibre optics – uses same cabling as ISC
 - ▶ 10km and longer distances with qualified WDM solutions
- **PSIFB (12x IB)**
 - ▶ 150 meter max distance optical cabling
 - ▶ Supports multiple CHPIDs per physical link
 - ▶ Multiple CF partitions can share physical link
- **ICB (Integrated Cluster Bus)**
 - ▶ Copper cable plugs close to memory bus
 - ▶ 10 meter max length
- **IC (Internal Coupling Channel)**
 - ▶ Microcode - no external connection
 - ▶ Only between partitions on same processor



Relative Performance
Based on avg data xfer size

PSIFB Configurations Supported

- InfiniBand coupling link attachments on System z9 are no longer limited to standalone coupling facilities.

z/OS CF	z890 / z990	System z9	System z10
z890 / z990	N/A	N/A	N/A
System z9 Standalone CF only	N/A	No	Yes
System z9 Mixed ICF & z/OS	N/A	No	SOD
System z10	N/A	SOD	Yes

Yes

Note: Support of IFB links on System z9 for attachment to System z10 are for migration purposes only. IBM does not intend to support 12x InfiniBand coupling links to connect two System z9 servers.

System z – Maximum Coupling Links and CHPIDs

Server	1x InfiniBand	12x InfiniBand	IC	ICB-4	ICB-3	ISC-3	Max**** External Links	Max***** Coupling CHPIDs
z196	32* M15 - 16	32* M15 - 16	32	N/A	N/A	48	80	128
z10 EC	32** E12 - 16	32* E12 - 16	32	16** (32/RPQ)	N/A	48	64	64
z10 BC	12***	12***	32	12***	N/A	48	64	64
z9 EC	N/A	16 S08 - 12	32	16	16	48	64	64
z9 BC	N/A	12	32	16	16	48	64	64

* Maximum of 32 InfiniBand links of all types on System z196.

** Maximum of 32 InfiniBand links of all types + ICB4 links on System z10 EC. ICB-4 not supported on Model E64

*** Maximum of 12 InfiniBand links of all types + ICB4 links on System z10 BC.

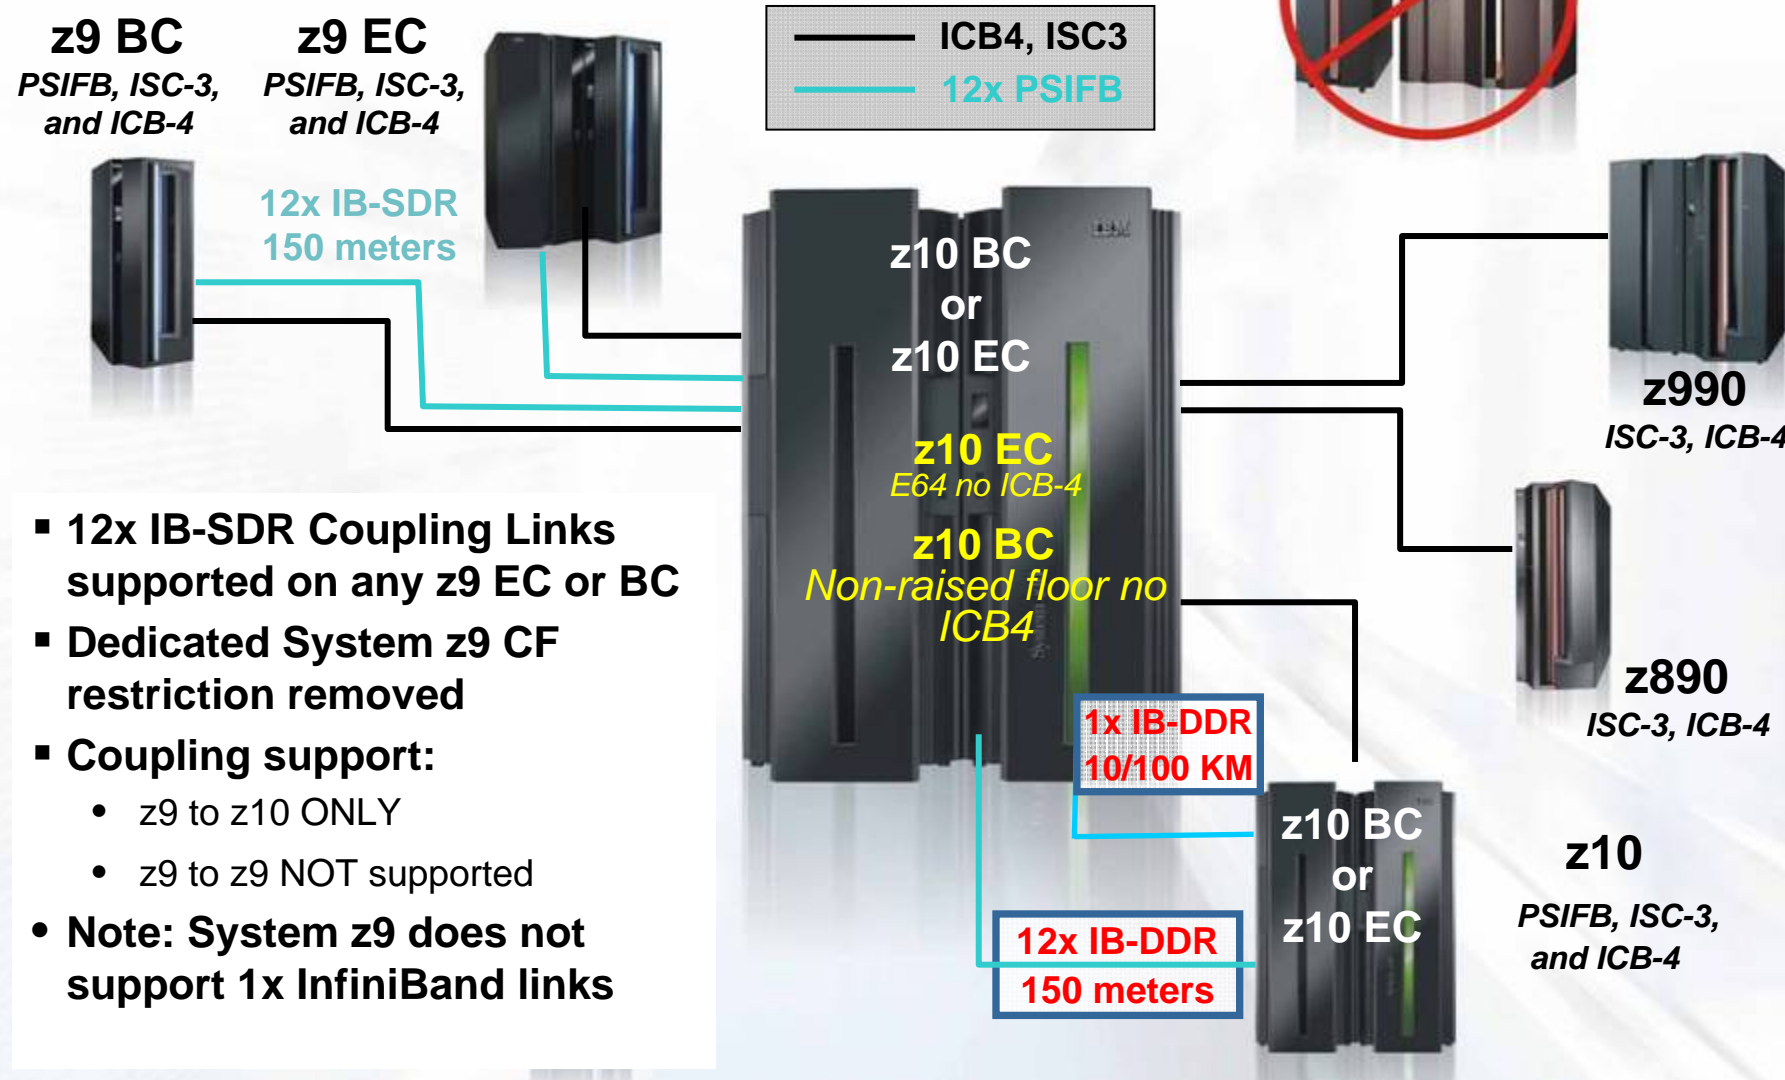
**** Maximum external links is the maximum total number of physical link ports (Does not include IC)

***** Maximum coupling CHPIDs defined in IOCDs includes IC and multiple CHPIDs defined on InfiniBand physical links.

z10 Parallel Sysplex coexistence and coupling connectivity

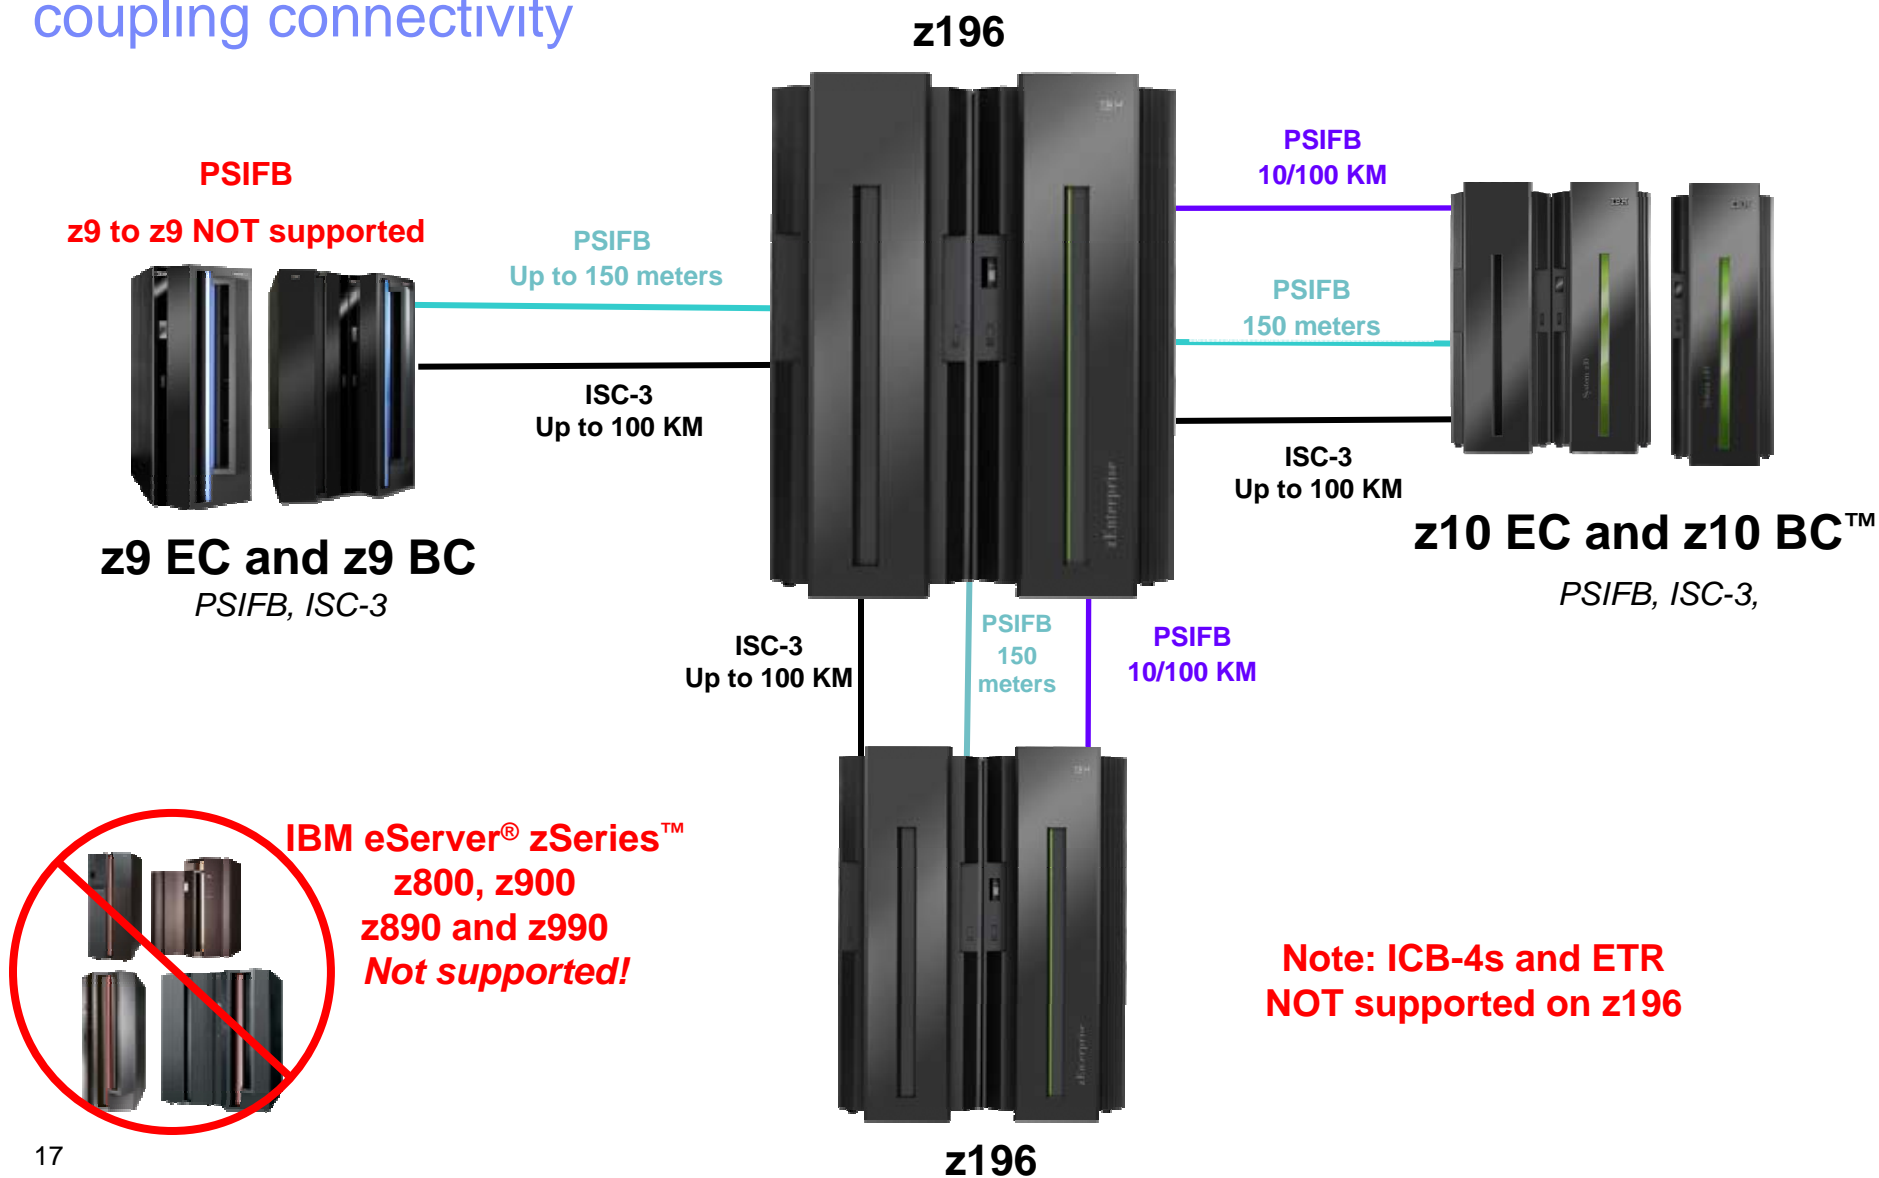


z800, z900
Not supported!



- 12x IB-SDR Coupling Links supported on any z9 EC or BC
- Dedicated System z9 CF restriction removed
- Coupling support:
 - z9 to z10 ONLY
 - z9 to z9 NOT supported
- **Note: System z9 does not support 1x InfiniBand links**

z196 Parallel Sysplex coexistence of Servers/CFs and coupling connectivity



System z9 12x IB-SDR Coupling Links

12x IB-SDR
Minimum – 0
Maximum – z9 BC 12 links Maximum – z9 EC 16 links (Model S08 12 links)
Order increment – 2 ports
Distance – 150 meters
OM3 fiber optic cables

Check System z9 Driver 67 Exception Letter on Resource Link for any restrictions.

System z9 does NOT support 1x IB-DDR InfiniBand Coupling Links

System z9 to System z9 InfiniBand Coupling Link connectivity is not supported

Note: The InfiniBand link data rates of 6 GBps, 3 GBps, 2.5 Gbps, or 5 Gbps 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. With InfiniBand coupling links, while the link data rate may be higher than that of ICB (12x IB-SDR or 12x IB-DDR) or ISC-3 (1x IB-SDR or 1x IB-DDR), the service times of coupling operations are greater, and the actual throughput may be less than with ICB links or ISC-3 links.

Up to 16 CHPIDs – across 2 ports




- Point-to-point up to 150 m (492 ft)
- 2 ports per HCA1-O fanout
- 12x IB-SDR (3 GBps – 12 lanes 2.5 Gbps)
 - z10 to any System z9
 - **Dedicated System z9 CF restriction removed**
- OS Support for non-dedicated CFs
 - z/OS 1.8 – definition and exploitation (z/OS 1.7 with Lifecycle Extension, 5637-A01)
 - z/VM V5.3 – definition only

Dynamic I/O configuration to define, modify and delete a CHPID when z/VM 5.3 is the controlling LPAR for dynamic I/O
- **Disruptive Outage Required for First PSIFB MES Addition for General Purpose System z9 Servers**

System z10 12x IB-DDR Coupling Link

12x IB-DDR
Minimum – 0
Maximum – z10 BC 12 links Maximum – z10 EC 32 links (Model E12 16 links)
Order increment – 2 ports
Distance – 150 meters
OM3 fiber optic cables

Up to 16 CHPIDs – across 2 ports



- Point-to-point up to 150 m (492 ft)
 - 2 ports per HCA2-O fanout
 - 12x IB-DDR (6 GBps – 12 lanes at 5 Gbps)
 - z10 to z10
 - 12x IB-SDR (3 GBps – 12 lanes at 2.5 Gbps)
 - **z10 to any System z9**
 - Autonegotiated to z9
 - OS Support for z10 CFs
 - z/OS 1.8 – definition and exploitation (z/OS 1.7 with Lifecycle Extension, 5637-A01)
 - z/VM V5.3 – definition only
- Dynamic I/O configuration to define, modify and delete a CHPID when z/VM 5.3 is the controlling LPAR for dynamic I/O


Check System z10 Driver 76 Exception Letter on Resource Link for any restrictions.

Note: The InfiniBand link data rates of 6 GBps, 3 GBps, 2.5 Gbps, or 5 Gbps 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. With InfiniBand coupling links, while the link data rate may be higher than that of ICB (12x IB-SDR or 12x IB-DDR) or ISC-3 (1x IB-SDR or 1x IB-DDR), the service times of coupling operations are greater, and the actual throughput may be less than with ICB links or ISC-3 links.


z196 InfiniBand coupling links

Type	Speed	Distance	Fanout	Cabling
12x InfiniBand	6 or 3 GBps	150 meters	HCA2-O	50μ MM (OM3) fiber
1x InfiniBand	5 or 2.5 Gbps	10 km	HCA2-O LR	9μ SM fiber

Up to 16 CHPIDs – across 2 ports



Up to 16 CHPIDs – across 2 ports

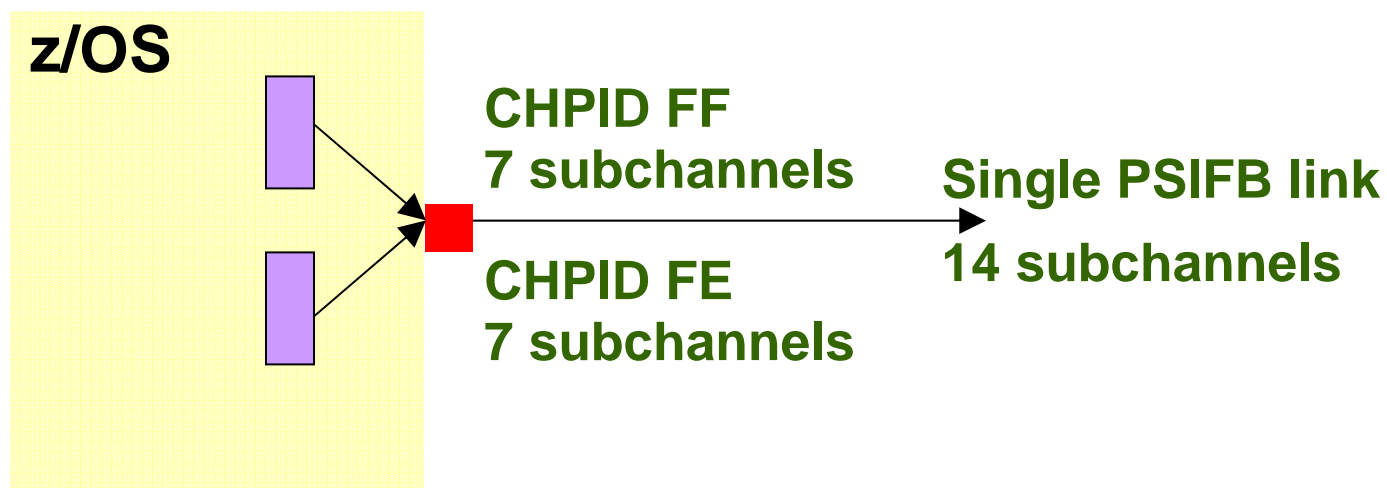


- **Ports exit from the front of a book**
Does not use I/O card slots
- **12x InfiniBand – z196, z10, z9**
 - DDR at 6 GBps
 - z196 and z10
 - SDR at 3 GBps
 - z196 & z10 to z9
 - First addition to z9 is disruptive
 - z9 to z9 connection not supported
- **1x InfiniBand – z196 and z10 (not z9)**
 - DDR at 5 Gbps
 - SDR at 2.5 Gbps (if DWDM requires)

DDR = double data rate, SDR = single data rate

Multiple Channel Paths

- Up to 16 CHPIDs across the two ports of single InfiniBand coupling HCA
 - More subchannels per physical link
 - **NOT** more subchannels per CHPID
 - Can connect to multiple CF LPARs

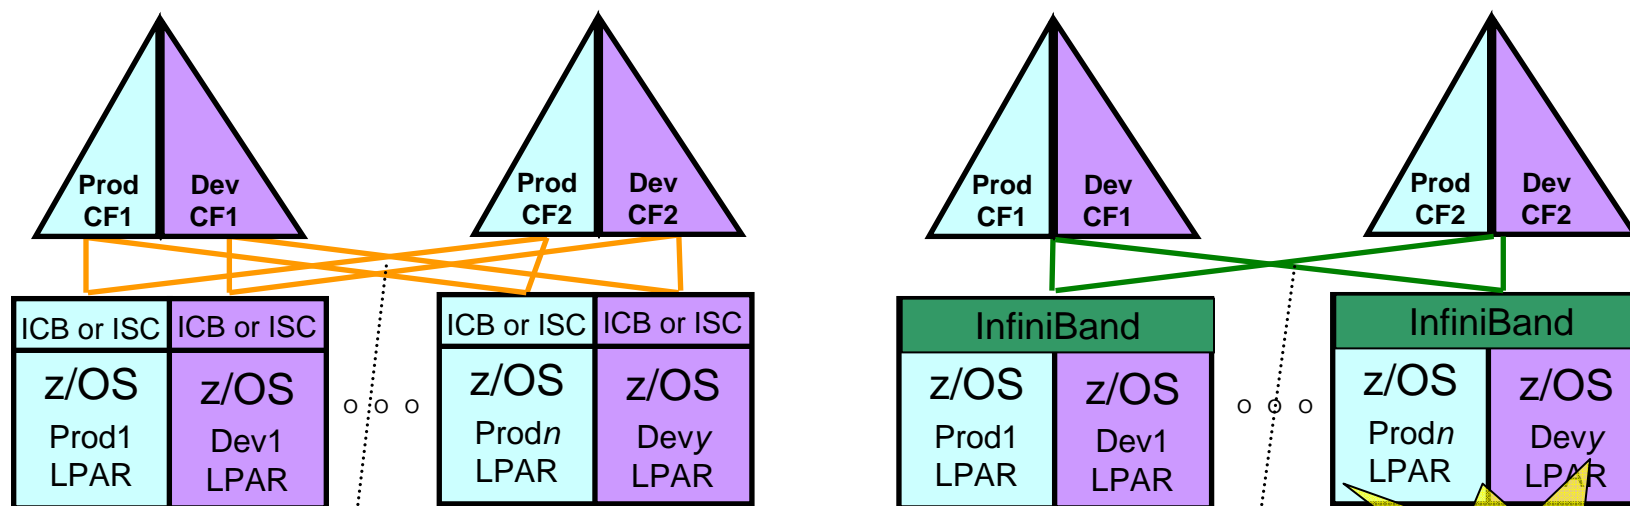


- MIF uses same address, 7 subchannels per CHPID

System z – Parallel Sysplex InfiniBand coupling (PSIFB)

Lower Cost Coupling Infrastructure – consolidating coupling links

- Can consolidate multiple ISC -or- ICB links with InfiniBand (PSIFB)
- Level of consolidation possible depends upon system configuration and workload
 - Do not define more than 8 CHPID per InfiniBand feature (although up to 16 CHPIDs are supported) without a careful review by IBM Advanced Technical Support or Development.
 - Rarely would more than 8 CHPIDs help a configuration -- and fully configured implementations (all 16 CHPIDs defined) may even perform worse than just using 8 CHPIDs.



Systems can share the IFB link

Planning considerations

▪ PSIFB links

- Requires z/OS 1.7 or later with PTFs

IBM Life Cycle Extension for z/OS 1.7 (**Expires on 09/30/2010**)

- Still have limit of 64 CF link CHPIDs per CPC
- Still have limit of 7 subchannels per CF link CHPID

But you can overcome this by defining multiple CHPIDs over the same link

Valid to have one port on a HCA2 card talking to a z9 (SDR), and the other port talking to z10 (DDR)

Planning considerations

▪ **Distance support:**

- 12x PSIFB supports max of 150 meters
- 1x DDR on z10 (ONLY) supports up to 10km unrepeated
- 1x DDR or SDR on DWDM supports up to 100km
 - DWDM support determines whether DDR or SDR is used

- Installation of first HCA on a z9 is disruptive - requires a POR to enable new IOP code

Planning considerations

Adapter ID (AID)

- New to z10 EC and System z9
 - The AID is used in defining CIB CHPIDs in HCD/IOCP

- When installed each HCA is assigned an Adapter ID (AID)
 - The AID has a number range of 00-1F

 - On z10 the AID is permanently assigned to an HCA, based on the HCA serial number, for as long as it's installed in the same CEC

Planning considerations

- **Same validation rules as for ISC and ICB channel paths and STP links**
 - CIB channel path can only be connected to another CIB channel path.
 - When a production IODF is built, all CIB channel paths have to be connected
 - A spanned CIB channel path must have defined the same Adapter ID (AID) and port for all channel subsystems where it is defined

Definition of InfiniBand Coupling Link

▪ **The Definition**

- Defining new channel path type CIB (Coupling using IB)
- Requires specification of Host Channel Adapter (HCA) ID and port number

▪ **Characteristics:**

- CIB channel path can be DED, REC, SHR or SPAN
- Up to 16 CHPIDs per HCA2-O
 - Maximum of 16 CHPIDs per AID (Adapter ID)
 - 16 CHPIDs can be shared across the two ports of the HCA2-O
- No PCHID value
- Point-to-point connections for HCA2-O/HCA2-O LR and via patch panel for HCA2-O LR only
- Target server identified by CSYSTEM on CHIPID statement
- Local server identified by LSYSTEM on ID statement

PSIFB – Deployment at Aetna

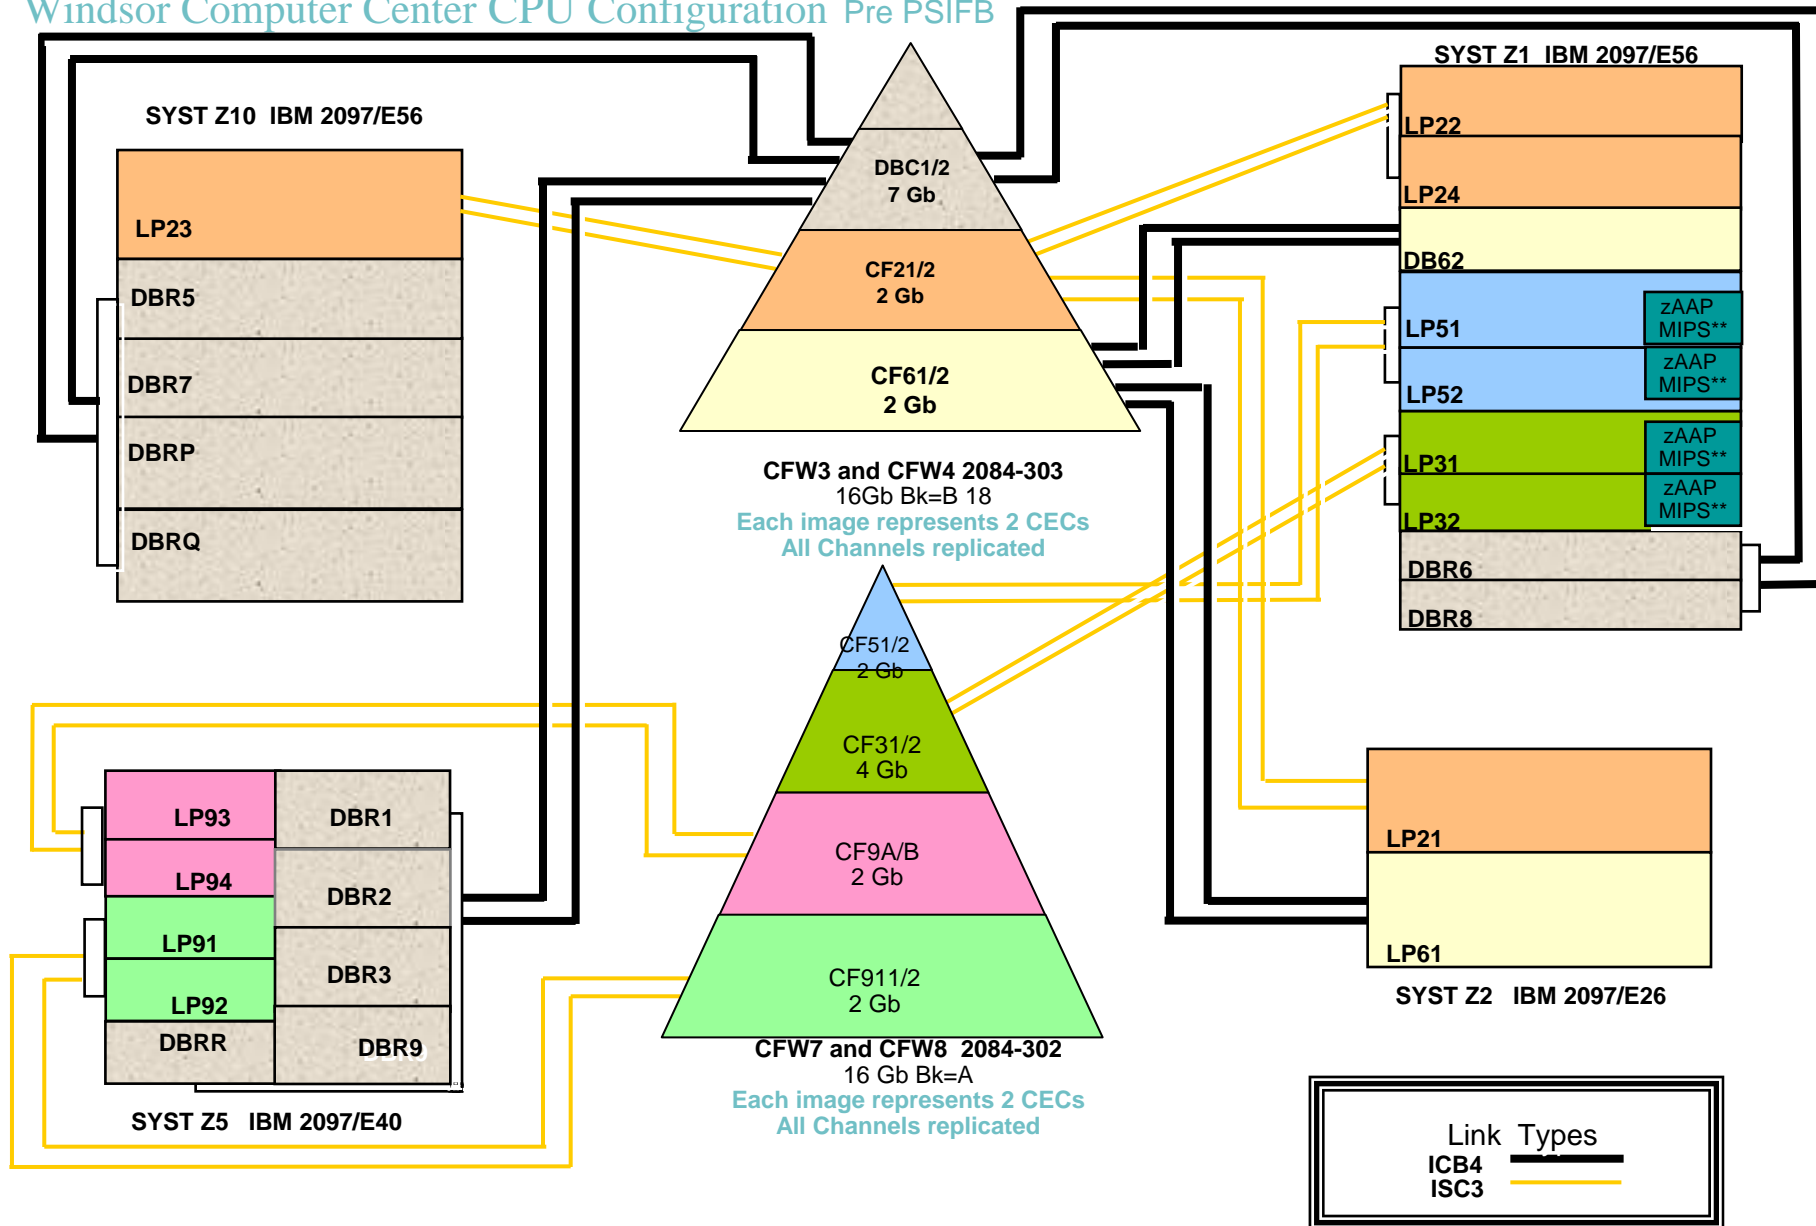
- Datacenter overview
 - 14 z10 2097's
 - 4 External CF 2097's
 - 2 Datacenters (8 CECs/6 CEC's)

- PSIFB – when and where

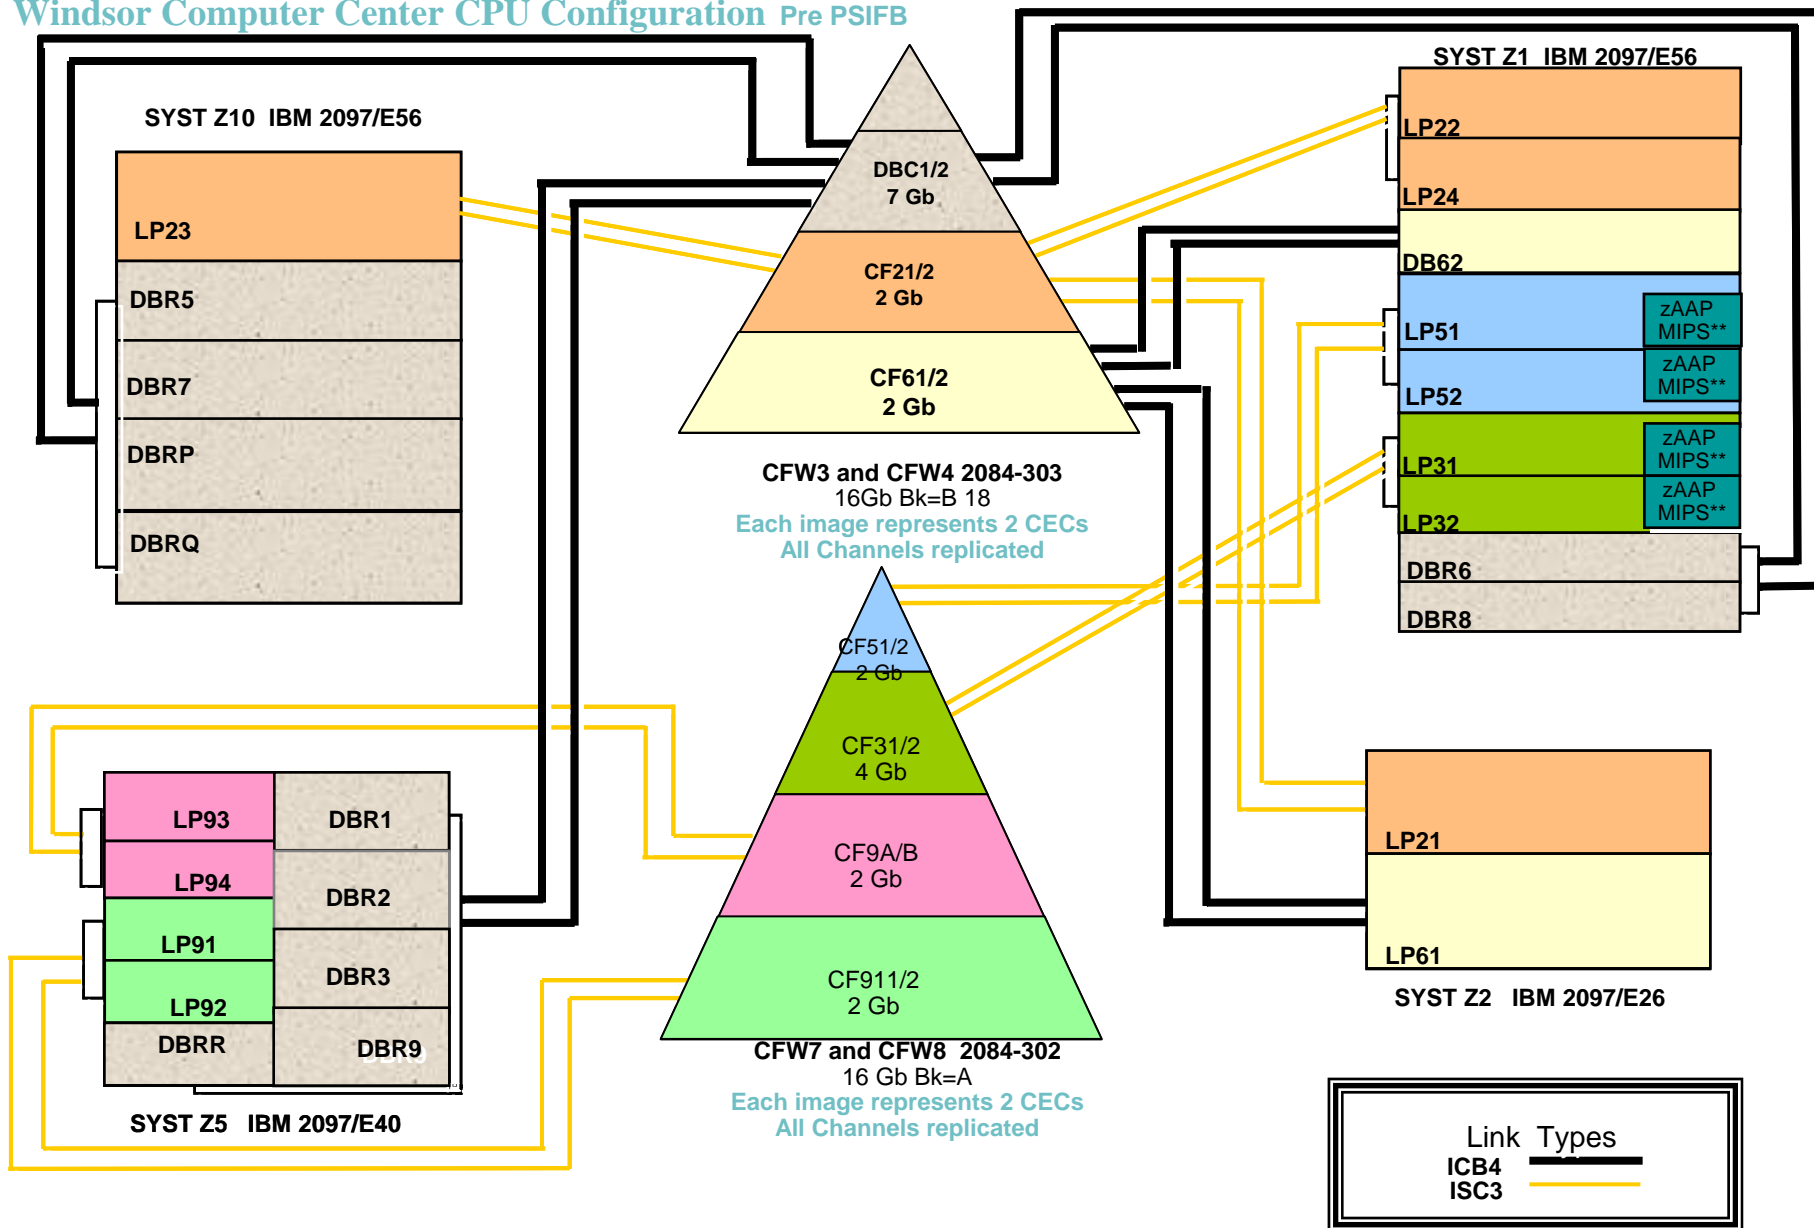
- Performance

handerag@aetna.com

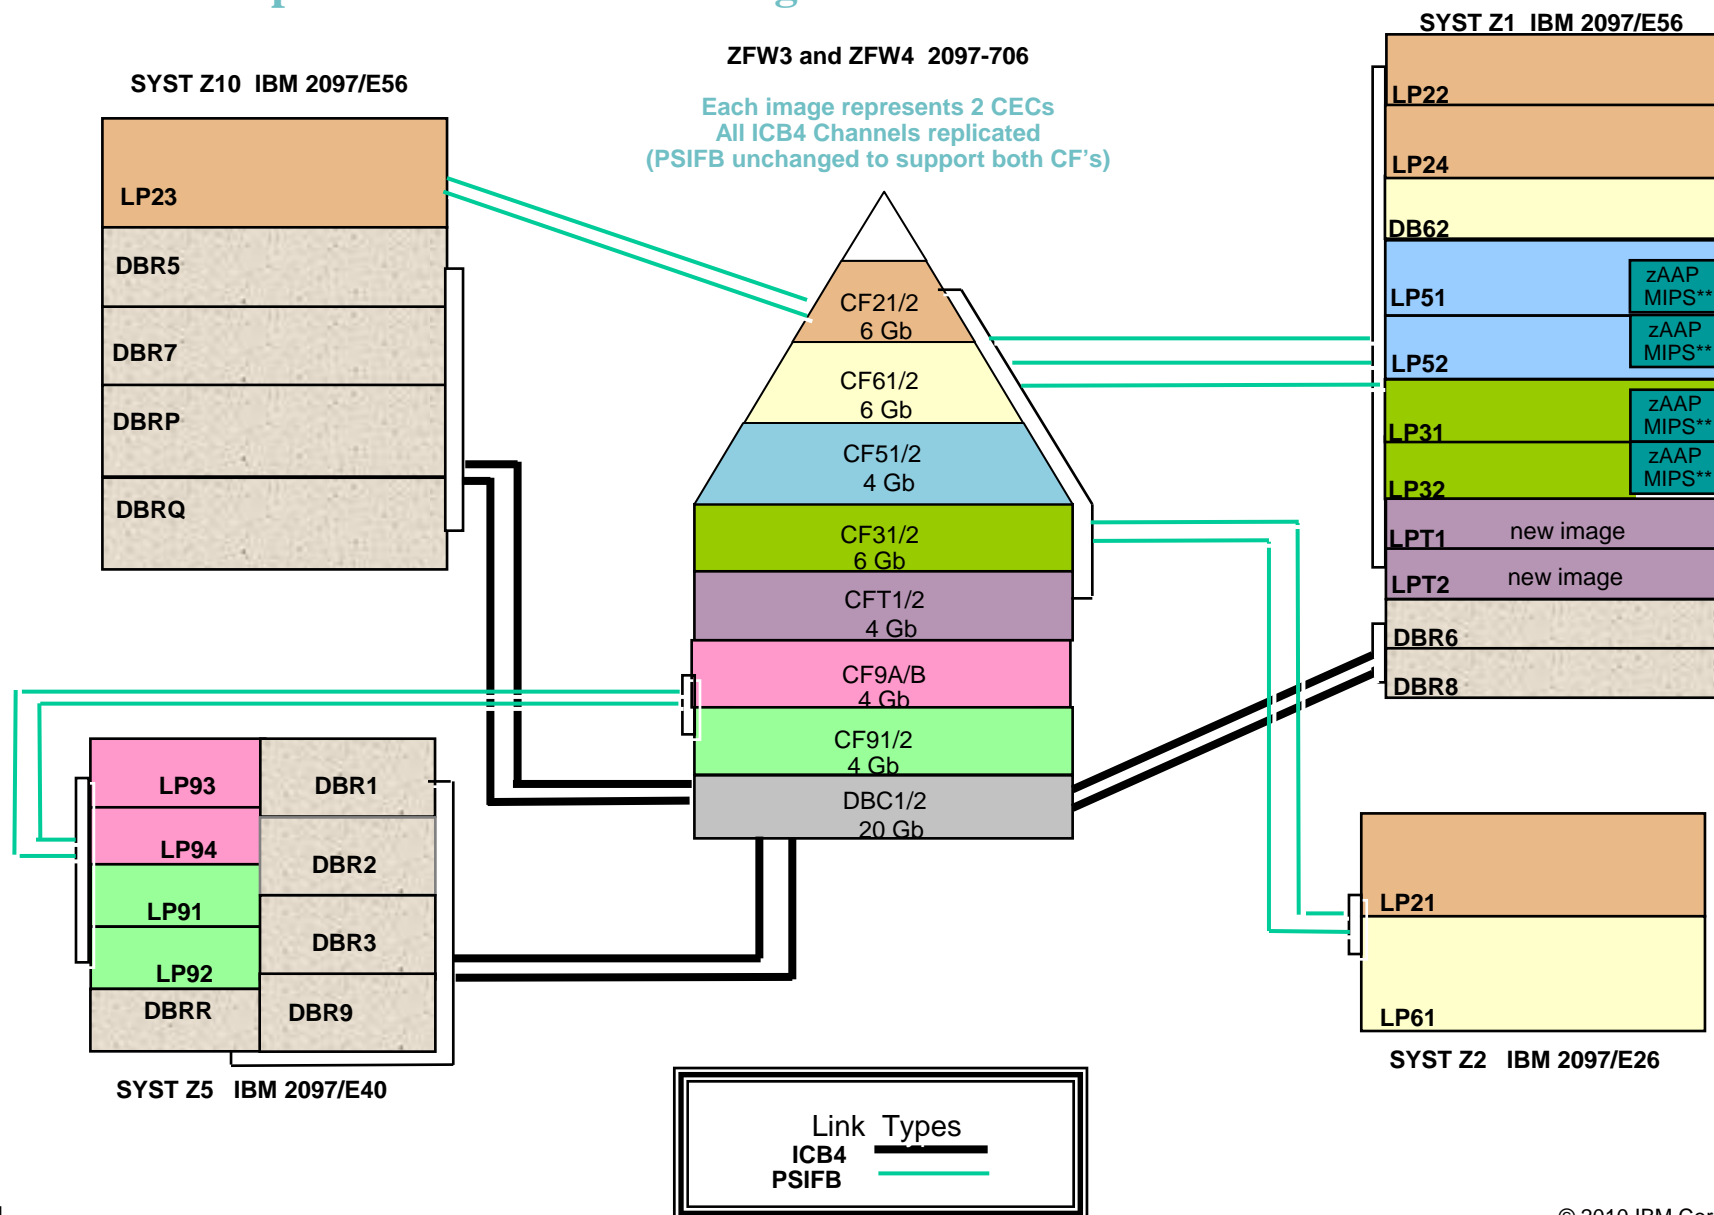
Windsor Computer Center CPU Configuration Pre PSIFB



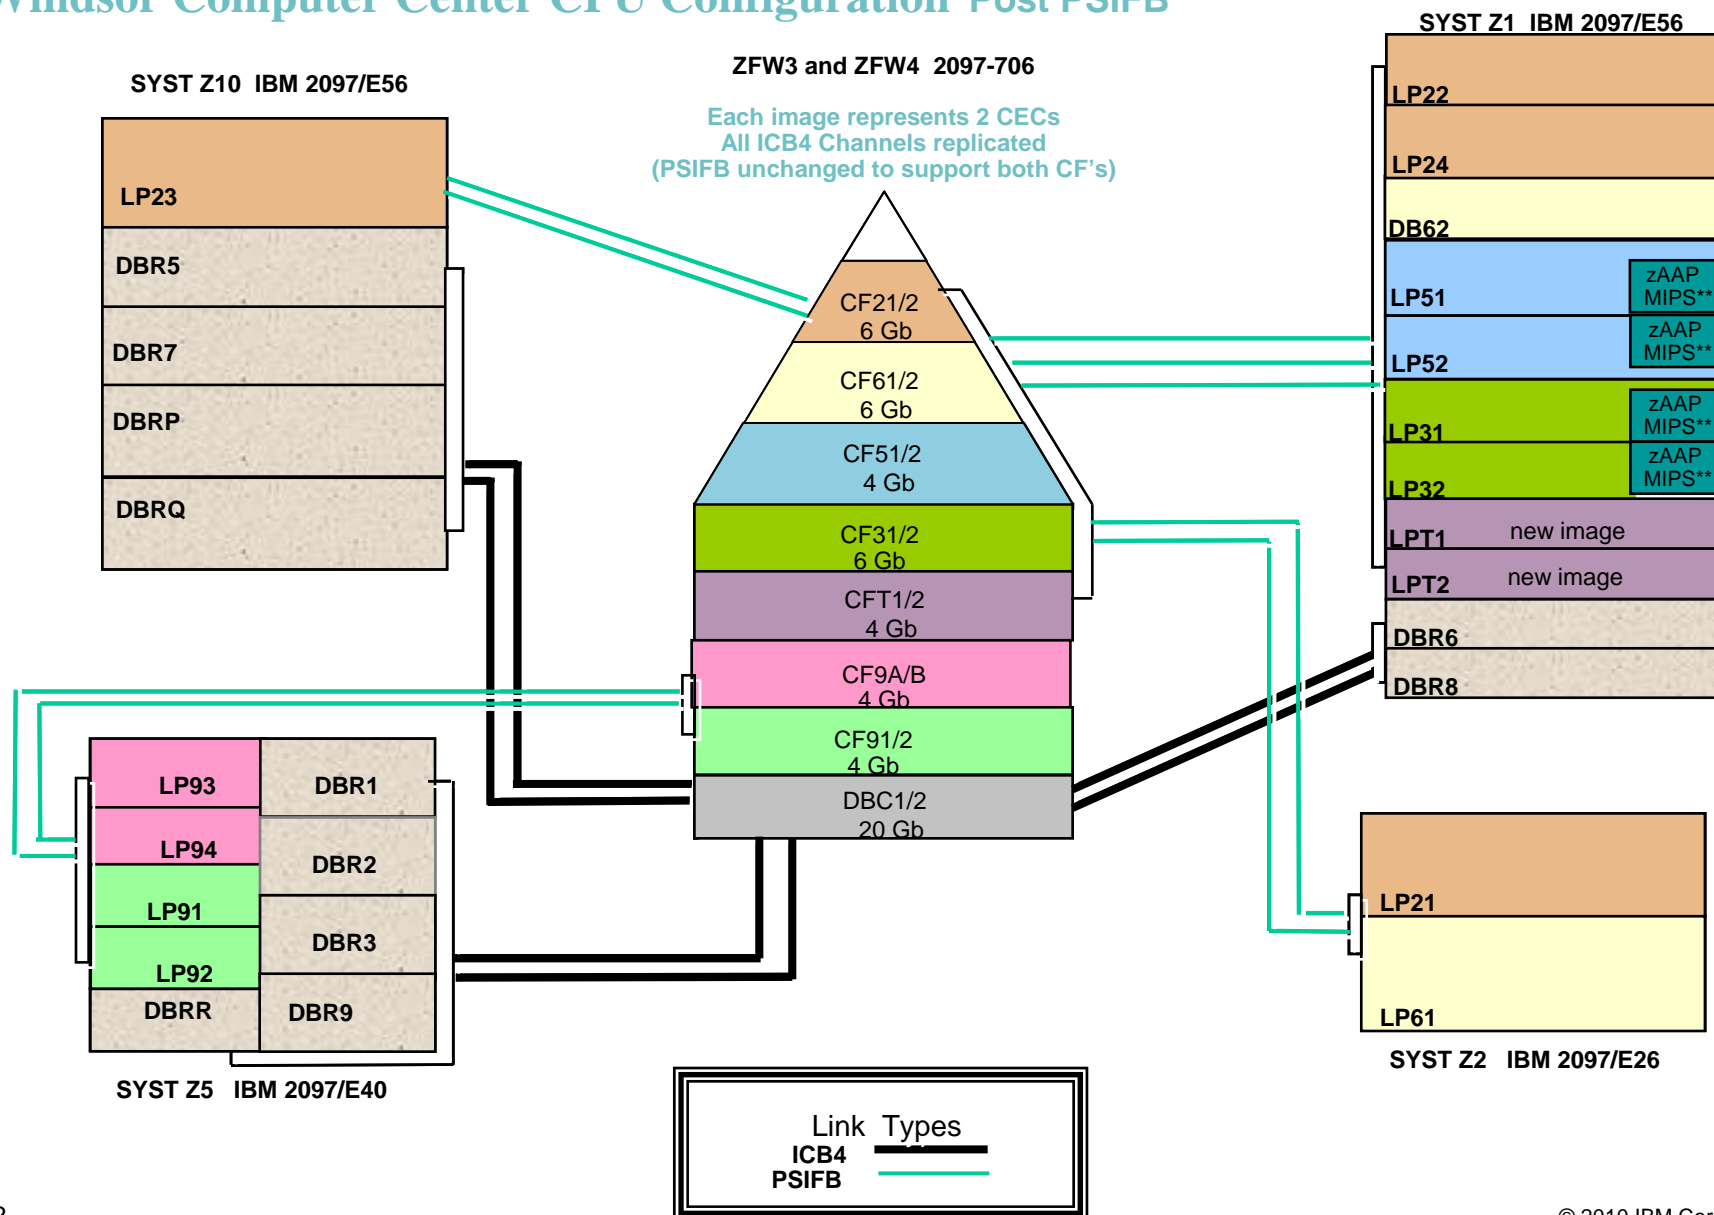
Windsor Computer Center CPU Configuration Pre PSIFB



Windsor Computer Center CPU Configuration Post PSIFB



Windsor Computer Center CPU Configuration Post PSIFB



PSIFB Heuristic Display

`-D XCF,C`

IXC357I 07.34.53 DISPLAY XCF 585

SYSTEM LP91 DATA

INTERVAL	OPNOTIFY	MAXMSG	CLEANUP	RETRY	CLASSLEN
165	168	4096	15	10	956

SSUM ACTION	SSUM INTERVAL	SSUM LIMIT	WEIGHT	MEMSTALLTIME
-------------	---------------	------------	--------	--------------

DEFAULT USER INTERVAL: 165

DERIVED SPIN INTERVAL: 165

DEFAULT USER OPNOTIFY: + 3

MAX SUPPORTED CFLEVEL: 16

MAX SUPPORTED SYSTEM-MANAGED PROCESS LEVEL: 16

SIMPLEX SYNC/ASYNCHRESHOLD: 26

DUPLEX SYNC/ASYNCHRESHOLD: 26

SIMPLEX LOCK SYNC/ASYNCHRESHOLD: 26

DUPLEX LOCK SYNC/ASYNCHRESHOLD: 28

PSIFB CF and RMFIII

```

Samples: 1800   Systems: 2   Date: 07/13/09   Time: 08.00.00   Range: 1800   Sec
----- Coupling Facility -----   ----- Processor -----   Request   - Storage --
Name      Type  Model Lv1  Dyn  Util% Def Shr Wgt Effect  Rate      Size  Avail
CFIA      2097   E40  16   OFF  15.4  1  0    1.0   14615    20G   17G
CFIB      2097   E40  16   OFF  10.5  1  0    1.0   26208    20G   18G

CF: CFIA      Type  ST System  CF  --- Sync ---   ----- Async -----
              Util  Rate  Avg  Rate  Avg  Chng  Del
              %    %    %    %    %    %    %
Structure Name
DSNDBMG_GBP20  CACHE A  *ALL    99.2 11545 26 2970 113 0.0 0.0
              CACHE AE91    5723 26 1474 114 0.0 0.0
              CACHE AE92    5822 26 1496 112 0.0 0.0

CF: CFIB      Type  ST System  CF  --- Sync ---   ----- Async -----
              Util  Rate  Avg  Rate  Avg  Chng  Del
              %    %    %    %    %    %    %
Structure Name
DSNDBMG_LOCK1 LOCK  A  *ALL    89.5 24388 13 134.1 56 0.0 0.0
              LOCK  AE91    12162 13 67.5 57 0.0 0.0
              LOCK  AE92    12226 13 66.6 56 0.0 0.0
  
```

PSIFB CF and RMFIII

```

Samples: 1800   Systems: 2   Date: 09/01/09   Time: 09.00.00   Range: 1800   Sec

```

Coupling Facility					Processor					Request	Storage	
Name	Type	Model	Lvl	Dyn	Util%	Def	Shr	Wgt	Effect	Rate	Size	Avail
CFIA	2097	E40	16	OFF	23.0	1	0		1.0	18757	20G	17G
CFIB	2097	E40	16	OFF	13.3	1	0		1.0	32894	20G	18G

CF:	Type	ST	System	CF	--- Sync ---			----- Async -----			
				Util	Rate	Avg	Rate	Avg	Chng	Del	
Structure Name				%		Serv		Serv	%	%	
DSNDBMG_GBP20	CACHE	A	*ALL	99.5	15467	22	3149	132	0.0	0.0	
	CACHE		AE91		7544	22	1459	133	0.0	0.0	
	CACHE		AE92		7923	22	1689	130	0.0	0.0	
CF:	Type	ST	System	CF	--- Sync ---			----- Async -----			
				Util	Rate	Avg	Rate	Avg	Chng	Del	
Structure Name				%		Serv		Serv	%	%	
DSNDBMG_LOCK1	LOCK	A	*ALL	92.7	31584	13	135.9	53	0.0	0.0	
	LOCK		AE91		14303	13	75.9	53	0.0	0.0	
	LOCK		AE92		17281	13	60.0	54	0.0	0.0	

PSIFB CF and RMFIII

Samples: 1800 Systems: 2 Date: 09/01/09 Time: 12.00.00 Range: 1800 Sec

----- Coupling Facility -----					----- Processor -----					Request	- Storage -	
Name	Type	Model	Lvl	Dyn	Util%	Def	Shr	Wgt	Effect	Rate	Size	Avail
CFIA	2097	E40	16	OFF	36.1	1	0		1.0	51078	20G	17G
CFIB	2097	E40	16	OFF	0.9	1	0		1.0	926.0	20G	18G

CF: **CFIA**

Structure Name	Type	ST	System	CF Util %	--- Sync ---		----- Async -----			
					Rate	Avg Serv	Rate	Avg Serv	Chng %	Del %
DSNDBMG_GBP20	CACHE	A	*ALL	64.2	15306	22	3287	129	0.0	0.0
	CACHE		AE91		7613	23	1702	127	0.0	0.0
	CACHE		AE92		7694	22	1585	131	0.0	0.0
DSNDBMG_LOCK1	LOCK	A	*ALL	35.1	31689	16	457.4	55	0.0	0.0
	LOCK		AE91		15739	16	246.7	55	0.0	0.0
	LOCK		AE92		15951	16	210.7	55	0.0	0.0

PSIFB CF and RMFIII

Samples: 1800 Systems: 2 Date: 09/09/09 Time: 07.30.00 Range: 1800 Sec

----- Coupling Facility -----					----- Processor -----				Request	- Storage -		
Name	Type	Model	Lvl	Dyn	Util%	Def	Shr	Wgt	Effect	Rate	Size	Avail
CFIA	2097	E40	16	OFF	21.6	1	0		1.0	32639	20G	17G
CFIB	2097	E40	16	OFF	0.6	1	0		1.0	1033	20G	18G

Samples: 1800 Systems: 2 Date: 09/09/09 Time: 07.30.00 Range: 1800 Sec

CF: CFIA	Type	ST	System	CF	--- Sync ---		----- Async -----			
Structure Name				Util	Rate	Avg	Rate	Avg	Chng	Del
				%		Serv		Serv	%	%
DSNDBMG_GBP20	CACHE	A	*ALL	62.9	9862	21	1746	171	0.0	0.0
	CACHE		AE91		4818	21	890.1	166	0.0	0.0
	CACHE		AE92		5044	20	856.2	177	0.0	0.0
DSNDBMG_LOCK1	LOCK	A	*ALL	36.0	20542	15	218.8	68	0.0	0.0
	LOCK		AE91		9635	15	114.7	67	0.0	0.0
	LOCK		AE92		10907	14	104.2	68	0.0	0.0

PSIFB CF and RMFIII

Samples: 1800 Systems: 2 Date: 09/09/09 Time: 09.00.00 Range: 1800 Sec

----- Coupling Facility -----					----- Processor -----				Request	- Storage -		
Name	Type	Model	Lvl	Dyn	Util%	Def	Shr	Wgt	Effect	Rate	Size	Avail
CFIA	2097	E40	16	OFF	17.9	1	0		1.0	15019	20G	17G
CFIB	2097	E40	16	OFF	12.0	1	0		1.0	32018	20G	18G

Samples: 1800 Systems: 2 Date: 09/09/09 Time: 09.00.00 Range: 1800 Sec

CF: **CFIA**

Structure Name	Type	ST	System	CF	Util	--- Sync ---		----- Async -----			
						Rate	Avg	Rate	Avg	Chng	Del
					%		Serv		Serv	%	%
DSNDBMG_GBP20	CACHE	A	*ALL		97.7	12374	20	2271	210	0.0	0.0
	CACHE		AE91			6166	20	1163	208	0.0	0.0
	CACHE		AE92			6209	20	1108	213	0.0	0.0

CF: **CFIB**

Structure Name	Type	ST	System	CF	Util	--- Sync ---		----- Async -----			
						Rate	Avg	Rate	Avg	Chng	Del
					%		Serv		Serv	%	%
DSNDBMG_LOCK1	LOCK	A	*ALL		92.1	30145	12	72.7	84	0.0	0.0
	LOCK		AE91			14741	13	39.7	84	0.0	0.0
	LOCK		AE92			15404	12	33.0	85	0.0	0.0

ICB4 CF and RMFIII

```

Samples: 1200      Systems: 10      Date: 03/12/10      Time: 10.05.00      Range: 1200 Sec
---- Coupling Facility ----      ---- Processor ----      Request - Storage --
Name      Type      Model Lvl  Dyn  Util% Def Shr Wgt Effect Rate      Size Avail
CFM1      2097      E12   16   OFF  23.8  5   0    4.5   126K   30G   22G
CFM2      2097      E12   16   OFF  16.8  5   0    4.6   71700  30G   21G

Samples: 1200      Systems: 10      Date: 03/12/10      Time: 10.05.00      Range: 1200 Sec
CF: CFM1
Type ST System      CF      --- Sync ---      --- Async ---
Structure Name      Util %      Rate      Avg Serv      Rate      Avg Serv      Chng %      Del %
DSNDB3G_LOCK1      LOCK A *ALL      18.2      49925      10      148.5      86      0.0      0.3
LOCK LP80      4606      10      66.2      87      0.0      0.0
LOCK LP81      9248      9      0.6      61      0.4      8.1
LOCK LP83      4962      11      10.0      49      0.6      0.6
LOCK LP85      2285      10      0.4      155      0.0      8.5
LOCK LP86      4666      11      69.9      90      0.0      0.0
LOCK LP87      9634      10      0.4      122      0.4      28.1
LOCK LP88      3109      10      0.4      55      0.8      15.0
LOCK LP89      11416     10      0.6      46      0.6      15.5
DSNDB3G_GBP10      CACHE A *ALL      6.9      9936      12      514.6      80      0.2      0.2
CACHE LP80      379.6     15      18.0     175      0.0     0.0
CACHE LP81      2377      11      142.0    74      0.0     0.0
CACHE LP83      3295      12      171.5    65      0.4     0.4
CACHE LP85      245.1     12      15.9     95      0.0     0.0
CACHE LP86      357.1     16      13.6    139      0.0     0.0
CACHE LP87      1126      12      49.7     88      0.0     0.0
CACHE LP88      826.5     13      48.4     76      0.2     0.2
CACHE LP89      1330     13      55.5     85      0.1     0.1
DSNDB3G_GBP20      CACHE A *ALL      12.3     8393      12      1033     84      0.2     0.2
CACHE LP80      317.4     15      23.4    120      0.0     0.0
CACHE LP81      1434      11      278.1    84      0.0     0.0
CACHE LP83      3181      12      238.5    70      0.7     0.7
CACHE LP85      161.9     13      23.5     90      0.0     0.0
CACHE LP86      340.9     14      17.7    126      0.0     0.0
CACHE LP87      1297      13      271.8    96      0.0     0.0
CACHE LP88      418.9     12      51.4     66      0.0     0.0
CACHE LP89      1242      13      128.9    83      0.0     0.0

```

PSIFB CF and RMFIII

Samples: 120 Systems: 2 Date: 03/14/10 Time: 18.16.00 Range: 120 Sec

----- Coupling Facility -----					----- Processor -----				Request	- Storage -		
Name	Type	Model	Lvl	Dyn	Util%	Def	Shr	Wgt	Effect	Rate	Size	Avail
EHL3	2097	E12	16	OFF	0.2	1	0		1.0	686.3	6050M	5483M
EHL4	2097	E12	16	OFF	4.0	1	0		1.0	6221	6050M	3773M

Samples: 120 Systems: 2 Date: 03/14/10 Time: 18.16.00 Range: 120 Sec

CF: ALL	Type	ST	System	CF	--- Sync ---		----- Async -----			
Structure Name				Util	Rate	Avg	Rate	Avg	Chng	Del
				%		Serv		Serv	%	%
DSNDB7G_GBP10	CACHE	A	*ALL	N/A	2836	14	425.3	186	23.8	24.6
	CACHE		LP61		216.0	13	1.1	197	0.0	0.0
	CACHE		LP62		2620	15	424.2	186	23.8	24.6
DSNDB7G_GBP20	CACHE	A	*ALL	N/A	2166	16	590.6	161	16.1	16.8
	CACHE		LP61		267.8	13	3.9	203	0.0	0.0
	CACHE		LP62		1898	16	586.8	161	16.2	16.9

PSIFB CF and RMFIII



Samples: 240 Systems: 2 Date: 08/03/10 Time: 09.23.00 Range: 240 Sec

CF: CFIA Type ST System CF --- Sync --- ----- Async -----

Structure Name	Type	ST	System	CF	Util %	Rate	Avg Serv	Rate	Avg Serv	Chng %	Del %
DSNDBMG_GBP20	CACHE	A	*ALL		27.4	17445	19	4014	125	0.0	0.0
	CACHE		AE91			8311	18	1594	135	0.0	0.0
	CACHE		AE92			9134	19	2420	117	0.0	0.0
DSNDBMG_GBP21	CACHE	A	*ALL		7.9	2318	22	1485	114	0.0	0.0
	CACHE		AE91			1157	25	902.7	114	0.0	0.0
	CACHE		AE92			1161	18	582.0	114	0.0	0.0
DSNDBMG_GBP22	CACHE	A	*ALL		3.8	1148	21	756.1	112	0.0	0.0
	CACHE		AE91			642.0	24	484.8	110	0.0	0.0
	CACHE		AE92			506.3	18	271.3	114	0.0	0.0
DSNDBMG_LOCK1	LOCK	A	*ALL		58.8	181K	12	1036	62	0.0	0.0
	LOCK		AE91			62628	12	248.1	70	0.0	0.0
	LOCK		AE92			119K	13	787.7	60	0.0	0.0

Samples: 300 Systems: 2 Date: 08/03/10 Time: 10.00.00 Range: 300 Sec

CF: CFC4 Type ST System CF --- Sync --- ----- Async -----

Structure Name	Type	ST	System	CF	Util %	Rate	Avg Serv	Rate	Avg Serv	Chng %	Del %
DSNDBMG_GBP20	CACHE	A	*ALL		22.0	16026	15	2755	140	0.0	0.0
	CACHE		AE91			8134	15	1149	153	0.0	0.0
	CACHE		AE92			7891	16	1606	130	0.0	0.0
DSNDBMG_GBP21	CACHE	A	*ALL		8.2	4273	16	1311	119	0.0	0.0
	CACHE		AE91			2241	16	939.9	107	0.0	0.0
	CACHE		AE92			2032	15	370.8	148	0.0	0.0
DSNDBMG_GBP22	CACHE	A	*ALL		4.6	2677	16	811.6	114	0.0	0.0
	CACHE		AE91			1402	16	572.5	103	0.0	0.0
	CACHE		AE92			1275	15	239.1	140	0.0	0.0
DSNDBMG_LOCK1	LOCK	A	*ALL		65.2	195K	11	73.3	49	0.0	0.0
	LOCK		AE91			78850	10	0.8	53	0.0	0.0
	LOCK		AE92			117K	11	72.4	49	0.0	0.0

References

- Overview
 - <http://www.redbooks.ibm.com/redpieces/pdfs/sg247539.pdf>
 - Getting Started with InfiniBand on System z10 and System z9
- Considerations
 - <http://www-03.ibm.com/systems/z/advantages/pso/whitepaper.html>
 - Coupling Facility Configurations Guidelines
- Tools
 - <http://w3-03.ibm.com/support/americas/wsc/cps.html>
 - zCP3000

Thank You !

