IBM z196 zEnterprise

George Handera

HanderaG@Aetna.com

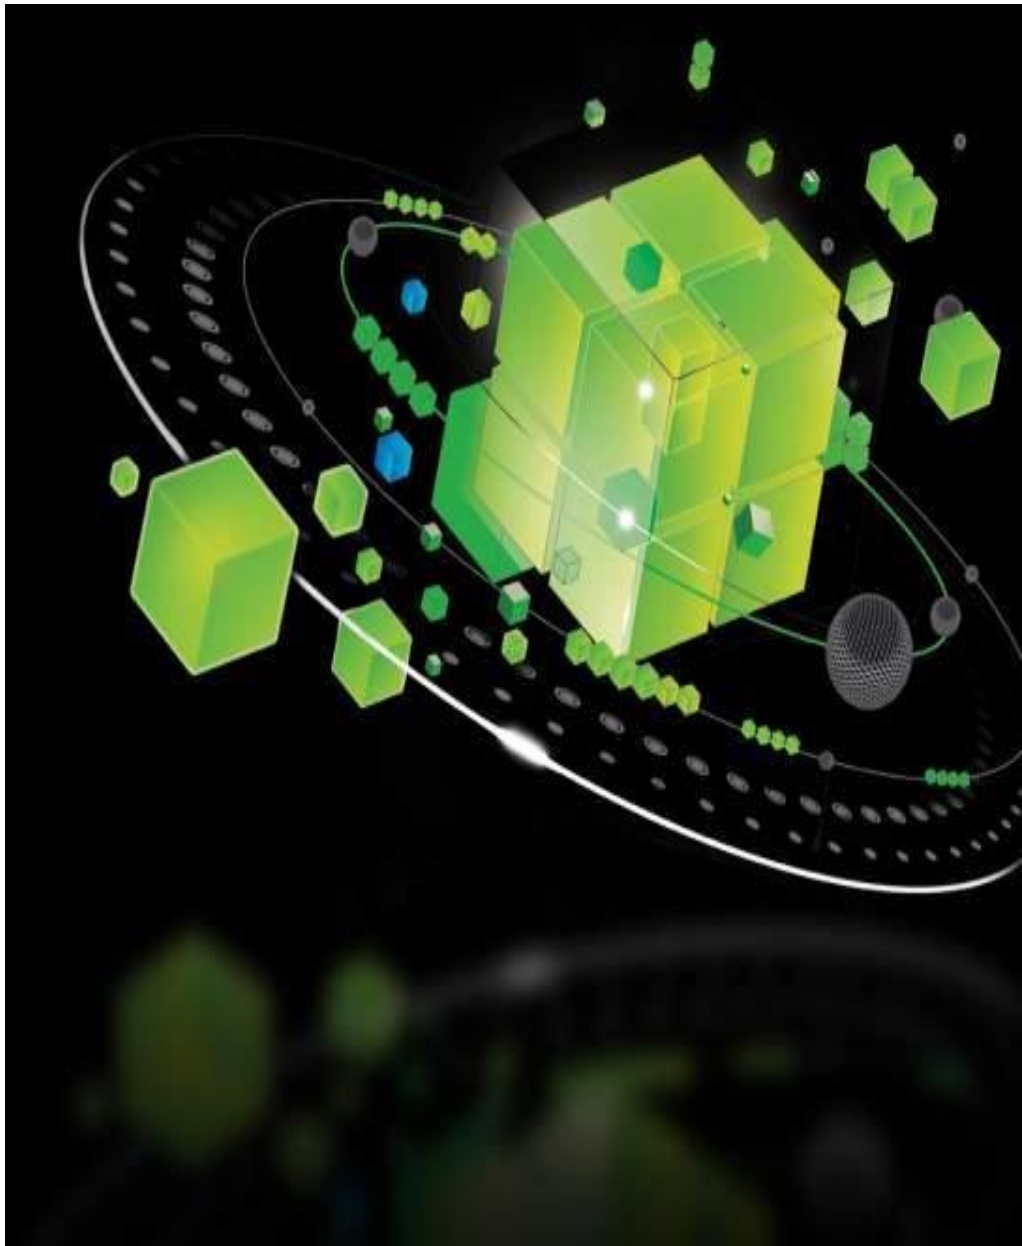

IBM z196 zEnterprise

Trademarks

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

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

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

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

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

Notes:

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

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

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

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

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

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

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

IBM z196 zEnterprise

Agenda

- Introduction/Background
- Implementation considerations
- Hardware feature comparison
- Hiperdispatch
- Experiences and observations

IBM z196 zEnterprise

Introduction/Background

George Handera - SE Capacity

Aetna has 2 datacenters - 25 miles apart

14 - 2097's and 1 - 2817

Each datacenter has a pair of external CF CECs

One datacenter has 1 SYSPLEX over 8 CECs - heavy Datasharing over 16 ICB4 links

Second datacenter 6 SYSPLEXES over 6 CECs plus the ESP box
Datasharing load is light over PSIFB and ICB4 links

IBM z196 zEnterprise

- **What's new with the z196**

- Energy attributes
- New level of cache on the CHIP
- New instruction support for C and Java environments
- Vertical scalability - applies to the LPAR as well as the CEC
- Large Page support evolving - benefits Websphere and DB2 workloads
- ETR timing no longer supported, STP implementation requires
- ICB4 links (high speed Coupling Facility links) no longer supported, PSIFB links best option for ICB4 replacement.

IBM z196 zEnterprise

Implementation Considerations

- Cooling - Water option available, our box is Air cooled
- Power - DC option available, our box runs AC
- STP no longer an option – time to get rid of the timer!
  We initially established a Mixed-CTN to support the 2817
- PSIFB – While Infiniband is not new, the 2097 was the last machine that supports ICB4 links.

  PSIFB becomes the replacement technology on the 2817. Our ESP testing experience focused on PSIFB for the Coupling Facility.

IBM z196 zEnterprise

Feature Comparison Table z9, z10 and z196

Series Model	Number of CPU's	Mips	Max Memory	Chip Cores	Cache Levels	Bus Speed	CPU GHz	CF High Speed Links
Z9 2094	1 - 54	580	512 GB	DUAL	2	2.7 GB	1.7	16-ICB4
Z10 2097	1 - 64	920	1.5 TB	QUAD	3	6.0 GB	4.4	16-ICB4 16/32-PSIFB
Z196 2817	1 - 80	1200	3 TB	QUAD w/ shared cache on chip	4	8.0 GB	5.2	32 PSIFB

IBM z196 zEnterprise

Model Configurations

Number of Books	2094 MODELS	Max Engines	2097 MODELS	Max Engines	2817 MODELS	Max Engines
1	S08	8	E12	12	M15	15
2	S18	18	E26	26	M32	32
3	S28	28	E40	40	M49	49
4	S38	38	E56	56	M66	66
4	S54	54	E64	64	M80	80

IBM z196 zEnterprise

Power

- z10 was introduced as a green machine
  - Power requirements increased substantially for the same book configuration when a z10 was compared to a z9.
  - From a MIP perspective the z9 and z10 were neck to neck on a MIPS per WATT comparison
- **z196** uses less power per book (approx .5 KVA) while delivering significantly more MIPS per book.

IBM z196 zEnterprise

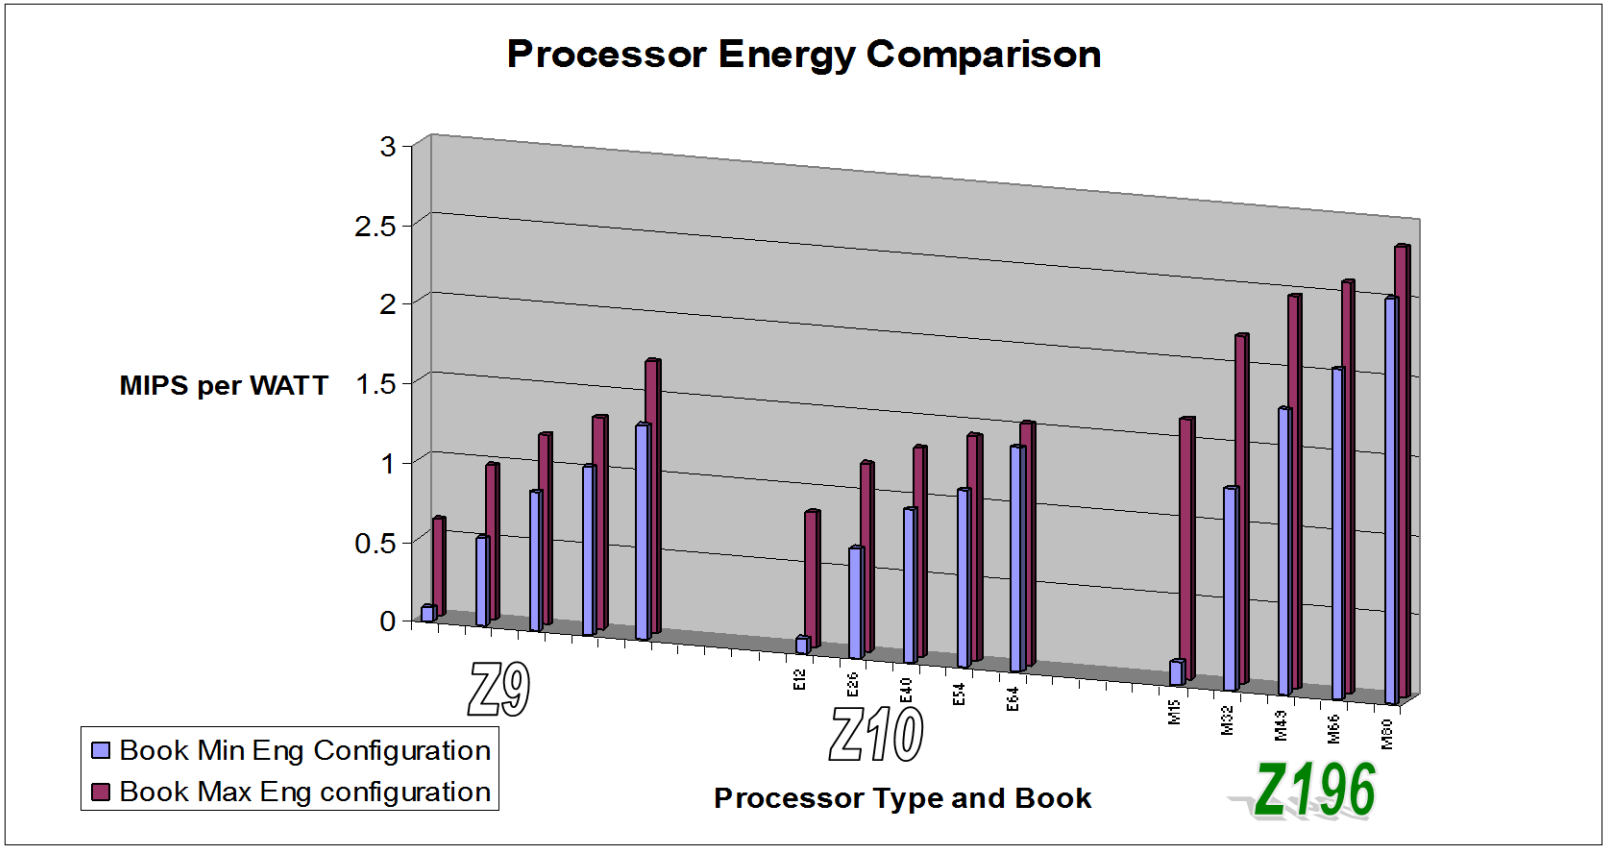

- These numbers were generated by the power estimation tool.
- Power requirements and MIPS per Book were generated using a minimum and maximum processors per book.

IBM z196 zEnterprise

HMC Environmental Display

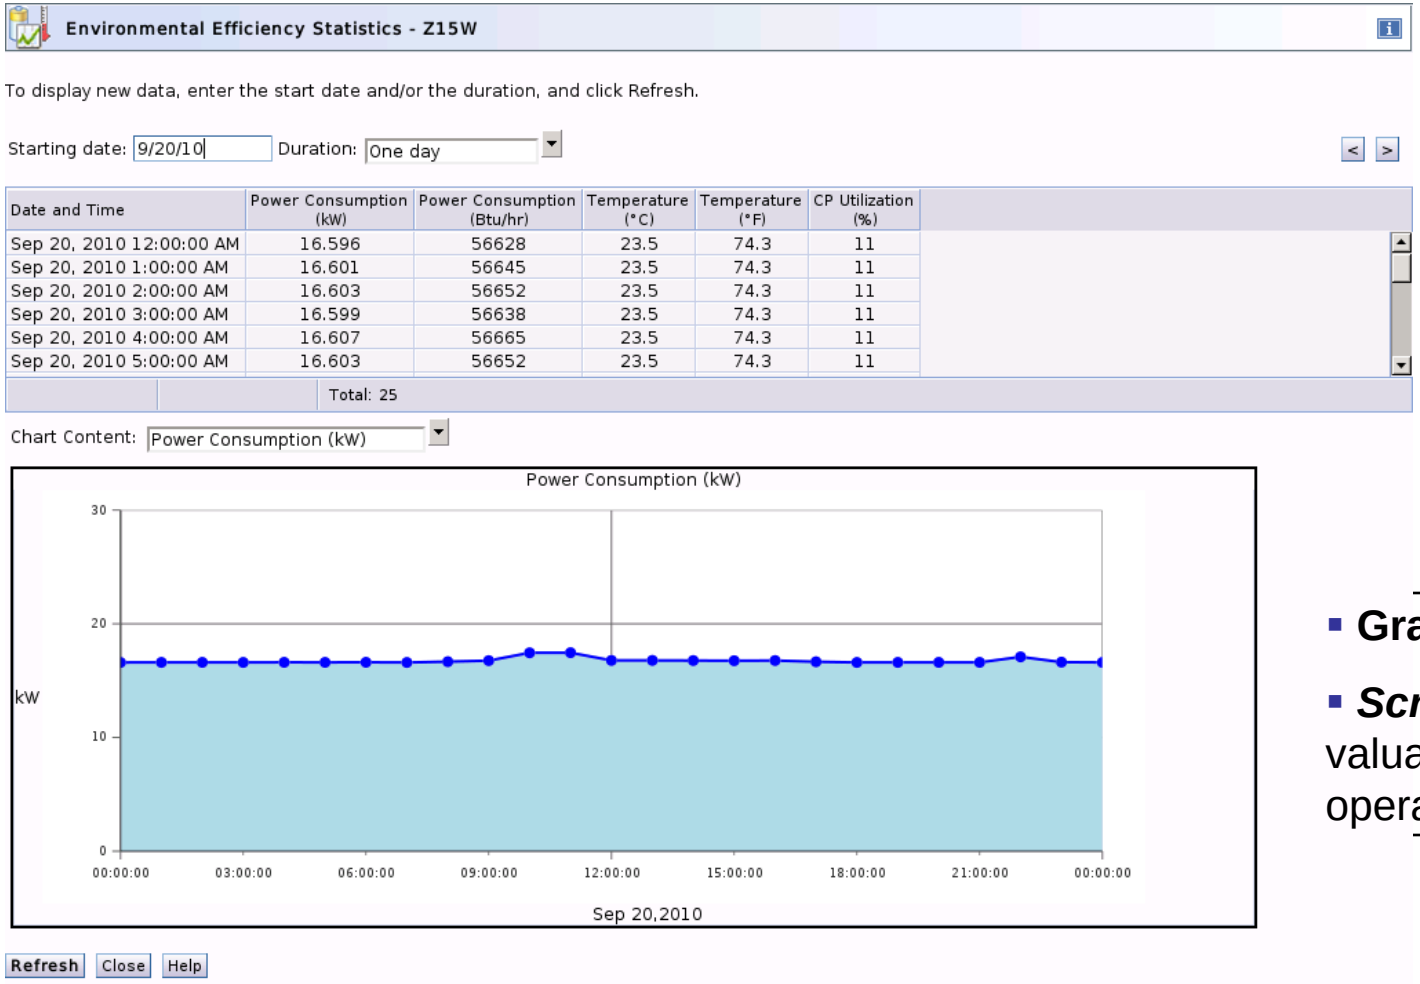

- **Graphic display over time**
- ***Screen capture support*** valuable for providing operations documentation

IBM z196 zEnterprise

Enhanced Activity Display

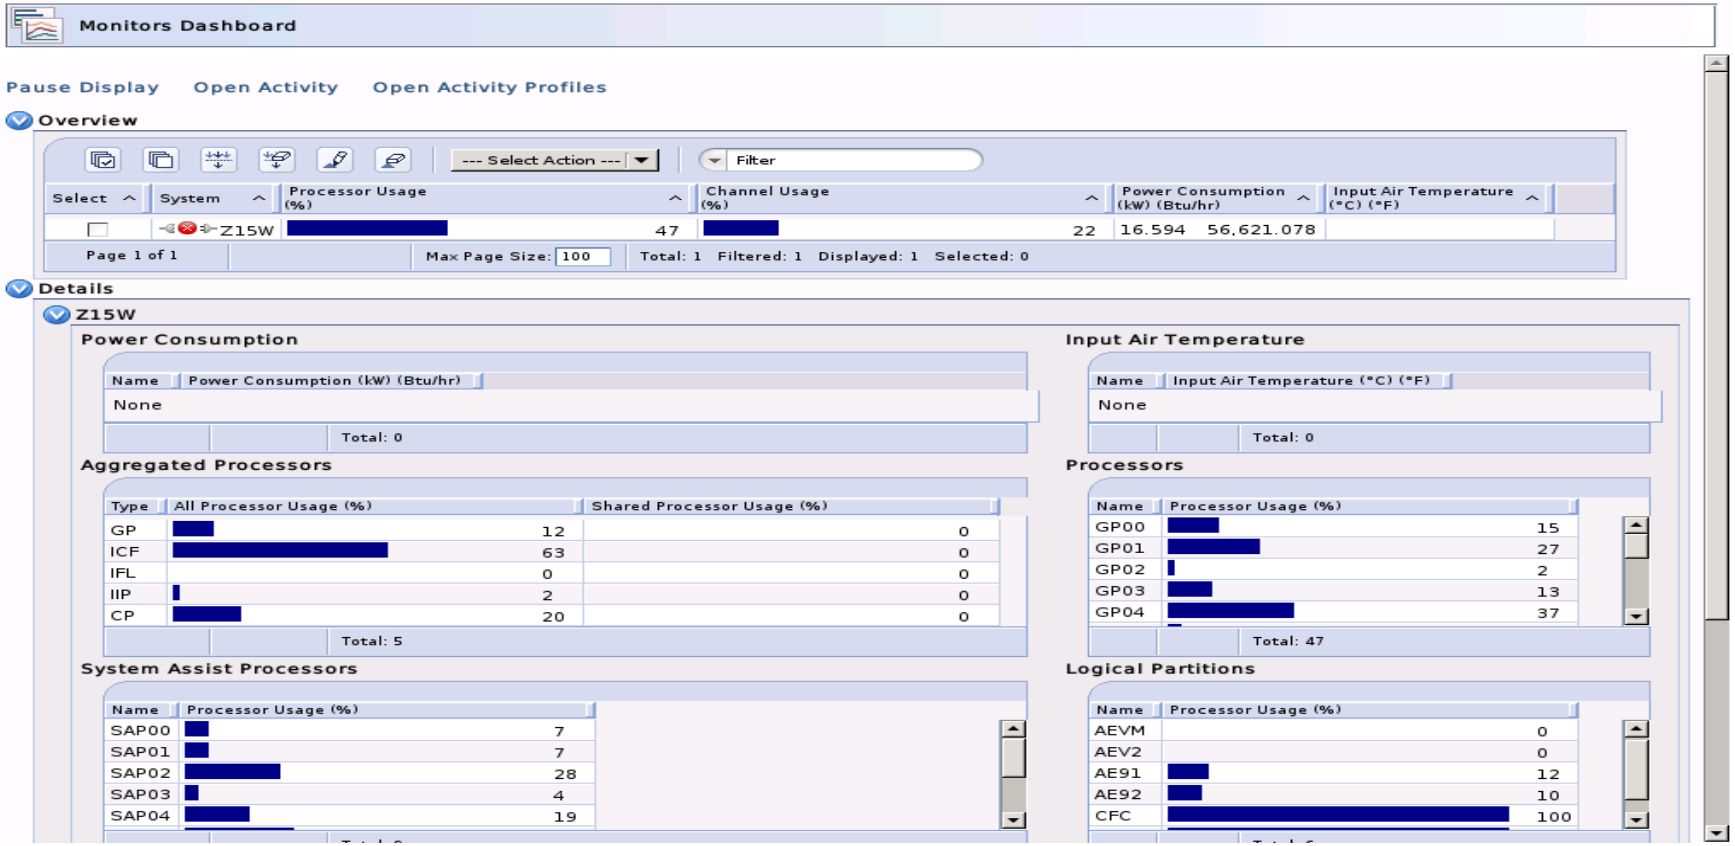

IBM z196 zEnterprise

MCM

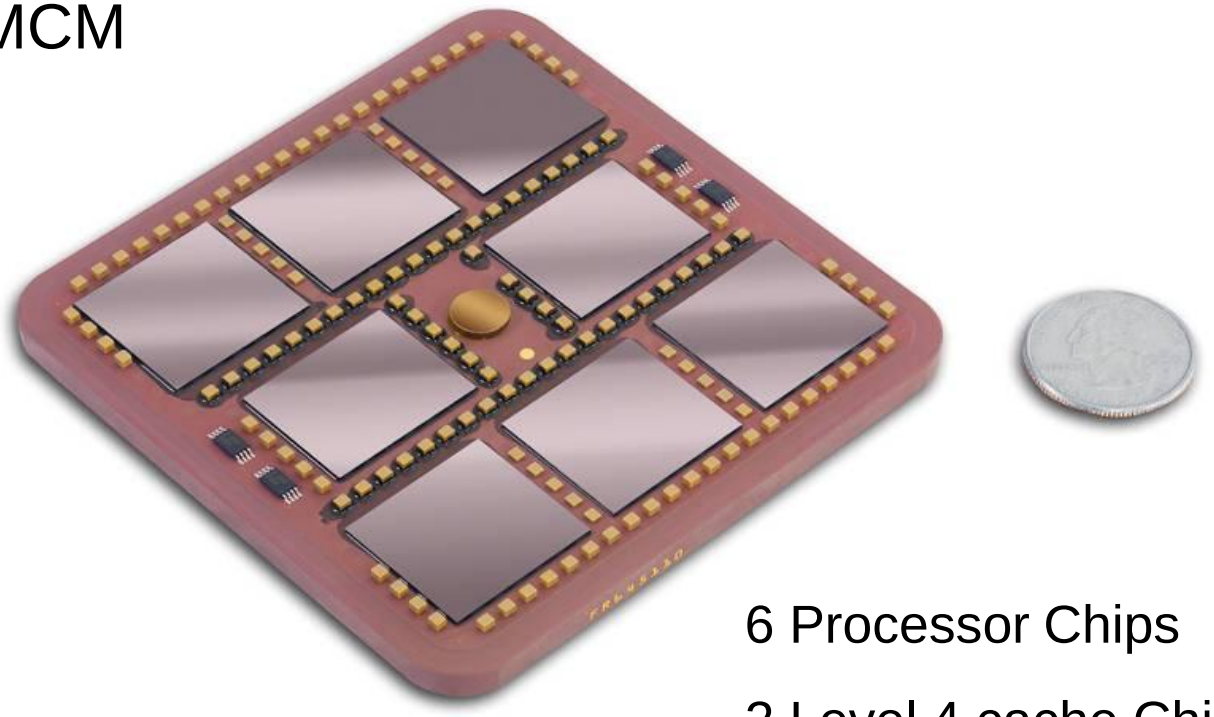

6 Processor Chips

2 Level 4 cache Chips

IBM z196 zEnterprise

2817 Model Characteristics by Book

	1st Book			2nd Book			3rd Book			4th Book			
Model	**Avail CPs**	**SAPs**	**Spares**	**Avail CPs**	**SAPs**	**Spares**	**Avail CPs**	**SAPs**	**Spares**	**Avail CPs**	**SAPs**	**Spares**	**Max Memory for Model**
M15	**15**	**3**	**2**										**752 GB**
M32	**16**	**3**	**1**	**16**	**3**	**1**							**1520 GB**
M49	**16**	**3**	**1**	**16**	**3**	**1**	**17**	**3**	**0**				**2288GB**
M66	**16**	**3**	**1**	**16**	**3**	**1**	**17**	**3**	**0**	**17**	**3**	**0**	**3056GB**
M80	**20**	**3**	**1**	**20**	**3**	**1**	**20**	**4**	**0**	**20**	**4**	**0**	**3056GB**

IBM z196 zEnterprise

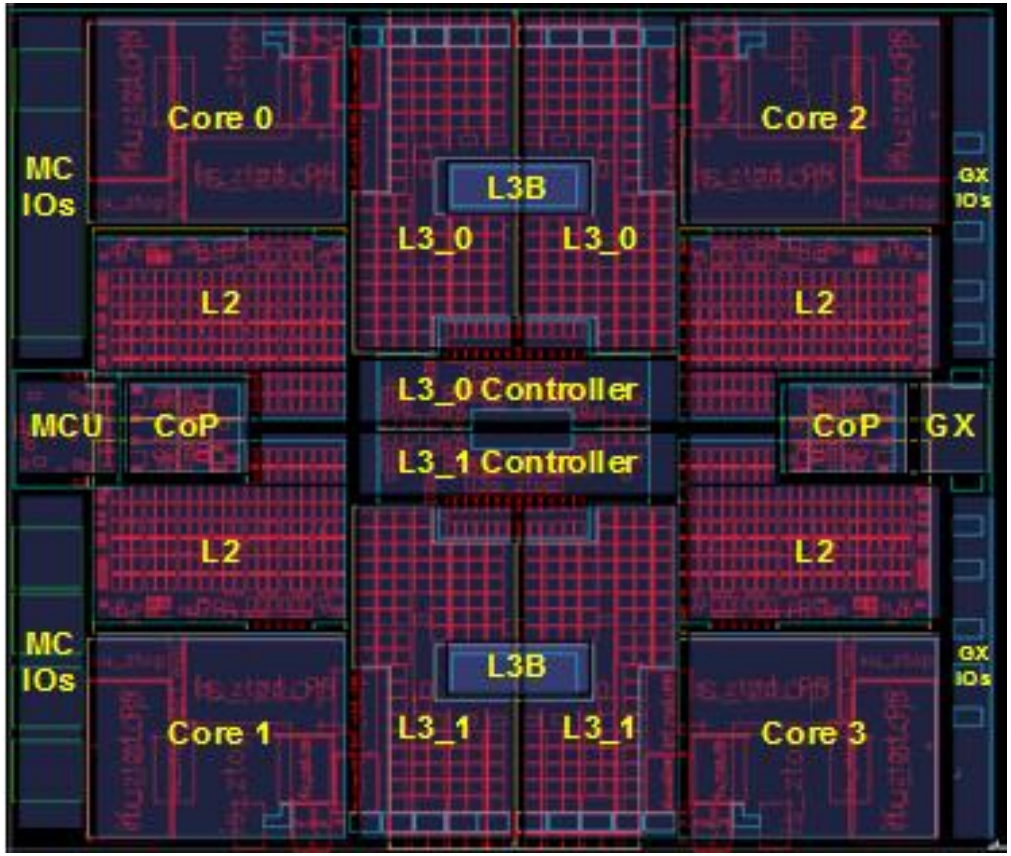

Processor Chip

- 4 Cores 3 cache levels
- Level 1 cache - 128/64 KB
- Level 2 cache - 1.5 MB
  formally known as level 1.5 cache
- Level 3 cache - 24 MB shared by all cores on a chip

Level 4 cache is on 2 separate

96 MB Chips for 192 MB total

IBM z196 zEnterprise

Cache Comparison

z10	Cache Level 1	Cache Level 1.5	Cache Level 2	
	128k Data 64k Instr	3 MB	**book cache**	
z196	Cache Level 1	Cache Level 2	Cache Level 3	Cache Level 4
	128k Data 64k Instr	1.5 MB	**Cores on chip share**	**book cache**

IBM z196 zEnterprise

Hiperdispatch

Hiperdispatch characteristics are an evolving science

- The Hipervisor continues to align physical processors vertically to a subset of logical processors (as introduced with the z10)
- WLM's management of the Affinity nodes has evolved.
  - Node affinity scope has changed from same book to same core (sharing level 3 cache influences this change)
  - WLM now places work on a weighted basis on the affinity nodes. 4 core chips/nodes are assigned more work that a 3 core chip.

IBM z196 zEnterprise

RMFIII

```
Samples: 60       System: AE92  Date: 06/30/10  Time: 18.01.00  Range: 60

Partition:   AE92        2817 Model 734
CPC Capacity:    3013   Weight % of Max: ****   4h Avg:     5   Group:    N/A
Image Capacity:  1595   WLM Capping %:    0.0   4h Max:   223   Limit:    N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num   Effect   Total   LPAR  Effect  Total

*CP                      12.0                        0.1     7.5    7.6
AE92           0   227  NO   12.0      21.2    21.4   0.1     7.5    7.5
PHYSICAL                                             0.1            0.1
```

Omegamon

```
 LPAR
+       Model = 2817-734  Serial = 0D7425     CEC MSUs capacity = 3134
+       Number of Physical processors = 47    Dispatch interval = DYNAMIC
+       Number of CPs = 34 Special CPs = 13   Elapsed interval = 00.00.01.005
+
```

IBM z196 zEnterprise

CICS TRANSACTION KEY

Tran Name	Threadsafe	Description
GHDL	**No**	**MQ 100 - .25 MS WAITS + 16 Million instructions per POST**
GHC1	**YES**	**1100 fetches per tran out of cache**
GHC2	**NO**	**1100 fetches per tran out of cache**
GHI0	**NO**	**1000-1100 fetches per tran prefetch I/O**
GHI1	**YES**	**1000-1100 fetches per tran prefetch I/O**

Transaction Driver - completion based initiation

Reads MQ Queue for transaction specifications for:

- Duration
- Transaction name
- Transaction concurrency

IBM z196 zEnterprise

Hiperdispatch

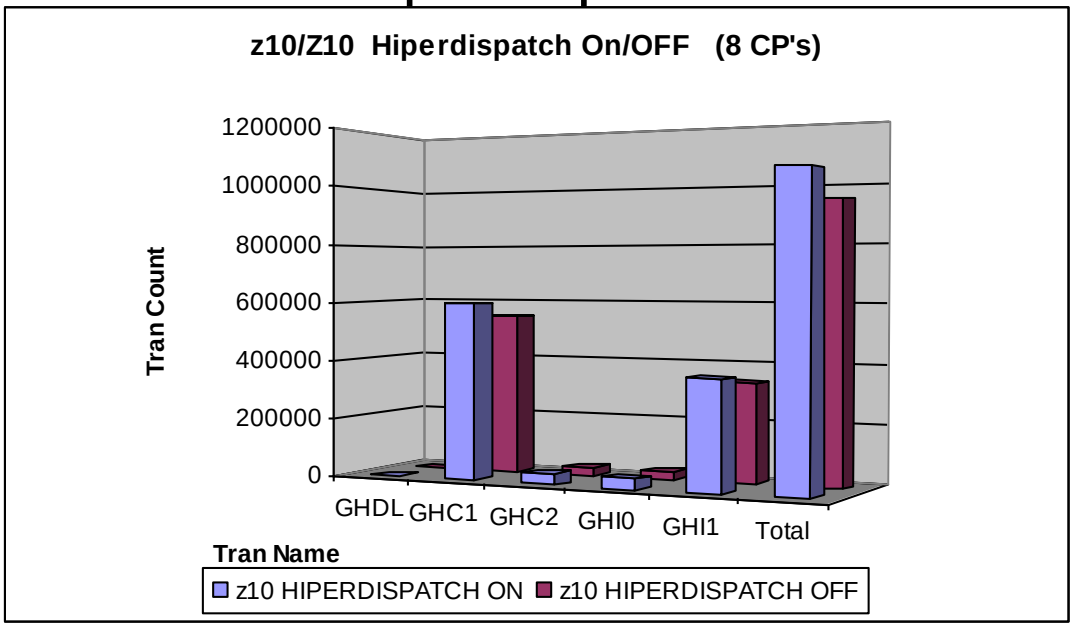

Tran	GHDL	GHC1	GHC2	GHI0	GHI1	total tran
z10 Hiperdispatch ON	**160**	**598667**	**36133**	**35100**	**367649**	**1037709**
z10 Hiperdispatch OFF	**144**	**551831**	**27873**	**27667**	**338888**	**946403**
Tran delta	**16**	**46836**	**8260**	**7433**	**28761**	**91306**
% delta	**11.11%**	**8.49%**	**29.63%**	**26.87%**	**8.49%**	**9.65%**

IBM z196 zEnterprise

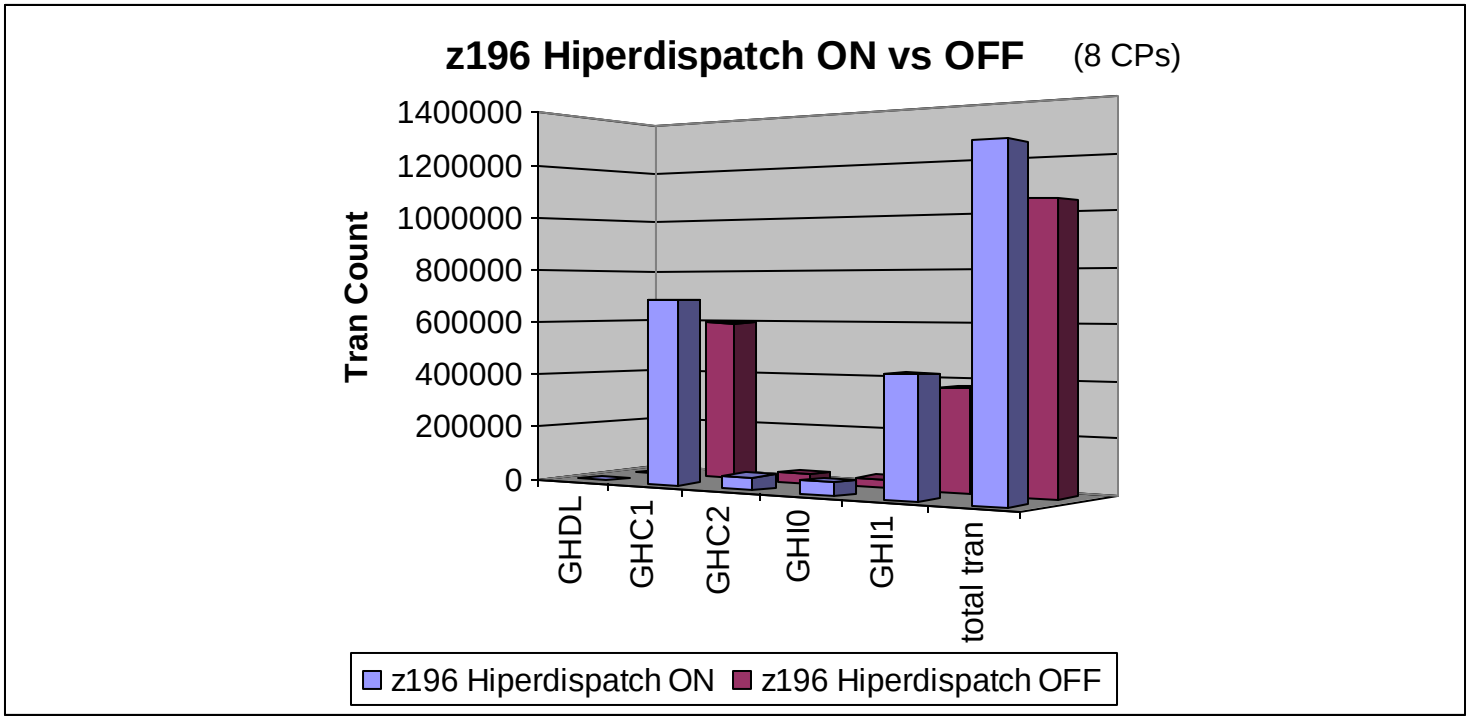

Tran	GHDL	GHC1	GHC2	GHI0	GHI1	total tran
z196 Hiperdispatch ON	180	688194	48813	47881	433296	1218364
z196 Hiperdispatch OFF	168	596079	34169	34001	373456	1037873
Tran delta	12	92115	14644	13880	59840	180491
% delta	7.14%	15.45%	42.86%	40.82%	16.02%	17.39%

IBM z196 zEnterprise

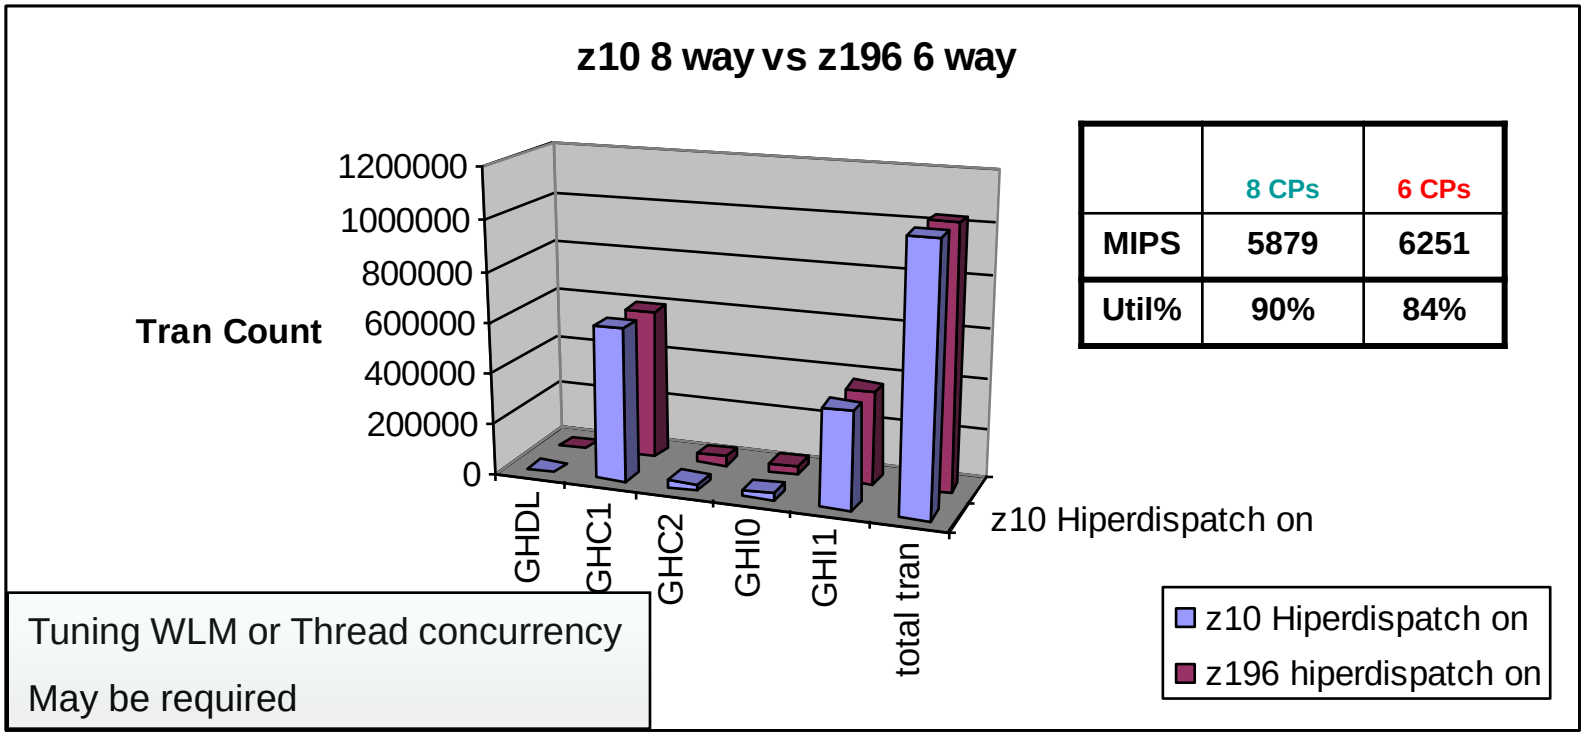

	8 CPs	6 CPs
MIPS	5879	6251
Util%	90%	84%

Tuning WLM or Thread concurrency

May be required

	GHDL	GHC1	GHC2	GHI0	GHI1	total tran
z196 Hiperdispatch ON	172	604553	26864	26853	383130	1041572
z10 Hiperdispatch ON	160	598667	36133	35100	367649	1037709
Tran delta	12	5886	-9269	-8247	15481	3863
% delta	7.50%	0.98%	-25.65%	-23.50%	4.21%	0.37%

IBM z196 zEnterprise

Hiperdispatch

- The z196 processor rewards your ability to hold a dispatch!
- MIPS are nice but point of dispatch must also be considered
- Chatty workloads (CICS QR for example) vs Batch or CICS Threadsafe lose ground to workloads that hold that their dispatch.
- When the point of dispatch is reduced and the MIPS per engine is increased the capability of the QR throughput increases from an engine perspective. BUT the competition for that engine also increases
- Machine upgrades often result in more MIPS being delivered by a reduced number of engines.
- Tuning WLM goals or managing concurrency may be required!

IBM z196 zEnterprise

SYSID	Mon	Day	SH	Hour	CPI	Prb State	Est Instr Cmplx CPI	Est Finite CPI	Est SCPL1M	L1MP	L2P	L3P	L4LP	L4RP	MEMP	Rel Nest Intensity	LPARCPU	Eff GHz	CICS
PAR1	SEP	1	P	16	7.0	31.8	2.8	4.2	50	8.5	69.3	19.0	7.6	3.4	0.7	0.46	1459.4	5.2	QR
PAR2	SEP	1	P	16	4.7	23.4	3.0	1.7	24	7.0	89.2	5.6	4.5	0.1	0.7	0.20	1546.2	5.2	Threadsafe

- Benchmark Description
  - Comprises of CICS transactions and some Batch…
    - All Batch is heavy Update and running on both LPARs
    - The CICS transactions are cloned pairs. One group is left to run in QR mode and the other is marked threadsafe in the CICS PPT definition. This test Focused all the Quasi-Reentrant transactions in one LPAR and all the Threadsafe transactions in the other LPAR. Transaction concurrency was establish in order to drive the LPARs to 90%+ utilization levels.
- Threadsafe Vs QR Results
  - CICS 110s
    - Increase of 52% of transactions
    - Decrease of 42% in CPU per Transaction
    - Decrease of average response time by 67% (3.0x)
  - RMF 72s – CICS Storage Class
    - Ended Transactions up 2.4x
    - Response Time down 3.6x
  - SMF 113s – LPAR
    - CPI down 1.48x from 7.0 to 4.7
      L1MP down 1.5% from 8.5% to 7.0%
      L2P up 19.9% from 69.3% to 89.2%

CICS Threadsafe is an option that may help you reduce CPU cost for applicable transactions by reducing switches between different TCB types

CPU MF example to supplement CICS and RMF performance metrics

As a secondary data source to understand why performance may have changed

These numbers come from a synthetic Benchmark and do not represent a production workload

IBM z196 zEnterprise

SYSID	Mon	Day	SH	Hour	CPI	Prb State	Est Instr Cmplx CPI	Est Finite CPI	Est SCPL1M	L1MP	L15P	L2LP	L2RP	MEMP	Rel Nest Intensity	LPARCPU	Eff GHz	Ded/Shr
AE83	JUN	1	P	TOTAL	6.87	15.4	3.0	3.9	88	4.4	69.5	23.9	0.1	6.6	0.73	464.7	4.40	Ded
AE83	JUN	2	P	TOTAL	7.27	14.9	3.2	4.1	93	4.4	67.2	25.8	0.1	6.9	0.78	612.8	4.40	Ded
AE83	JUN	3	P	TOTAL	6.43	12.5	2.9	3.6	93	3.8	67.3	25.5	0.1	7.1	0.79	627.6	4.40	Ded
AE83	JUN	4	P	TOTAL	7.14	14.8	3.2	4.0	91	4.4	67.5	25.6	0.1	6.8	0.77	553.4	4.40	Ded
	Average				6.93	14.4	3.1	3.9	91	4.3	67.9	25.2	0.1	6.9	0.77	564.6		Ded
AE83	MAY	25	P	TOTAL	7.66	17.6	3.2	4.5	87	5.2	68.6	25.4	0.1	6.0	0.70	641.6	4.40	Shr
AE83	MAY	26	P	TOTAL	7.36	16.1	3.1	4.3	87	4.9	69.0	24.7	0.1	6.3	0.72	621.1	4.40	Shr
AE83	MAY	27	P	TOTAL	7.31	14.5	3.1	4.2	88	4.8	68.5	25.0	0.1	6.4	0.73	395.6	4.40	Shr
	Average				7.44	16.0	3.1	4.3	87	5.0	68.7	25.0	0.1	6.2	0.72	552.8		Shr
Dedicated / Shared					0.93	0.90	0.98	0.90	1.04	0.86	0.99	1.01	0.98	1.11	1.07	1.02		
					1.07	Relative ITR Capacity Ratio Of Dedicated Vs Shared												

CPI – Cycles per Instruction

Prb State - % Problem State

Est Instr Cmplx CPI – Estimated Instruction Complexity CPI (infinite L1)

Est Finite CPI – Estimated CPI from Finite cache/memory

Est SCPL1M – Estimated Sourcing Cycles per Level 1 Miss

L1MP – Level 1 Miss %

L15P – % sourced from Level 2 cache

L2LP – % sourced from Level 2 Local cache (on same book)

L2RP – % sourced from Level 2 Remote cache (on different book)

MEMP - % sourced from Memory

Rel Nest Intensity – Reflects distribution and latency of sourcing from shared caches and memory

LPARCPU - APPL% (GCPs, zAAPs, zIIPs) captured and uncaptured

Eff GHz – Effective gigahertz for GCPs, cycles per nanosecond

IBM z196 zEnterprise

DB2 V10 testing - HIS data

Field defs on next page

	CPI	PRBSTATE	L1MP	L2P	L3P	L4LP	L4RP	LPARBUS	MEMP	MIPSEXE(	ESTICCPI	ESTFINCF	ESTSCP1I	RNI	EFFGHZ	TLB1MISS	TLB1CYCL	PTEPCTM	SYSTEM
05JAN2011:06:59:	5.3	25.7	8.5	85.1	8.3	4.3	1.5	1633.5	3.9	15971.4	2.7	2.6	31.3	0.7	5.2	8.5	41.4	15.9	AE91
05JAN2011:07:59:	5.4	25.6	8.4	85.0	8.4	4.3	1.5	1654.7	3.9	16092.4	2.7	2.6	31.3	0.6	5.2	8.5	41.8	16.0	AE91
05JAN2011:06:59	5.1	25.6	8.4	85.2	8.2	5.0	0.8	1614.3	4.5	16506.5	2.6	2.5	29.7	0.7	5.2	8.6	42.0	16.9	AE92
05JAN2011:07:59	5.1	25.3	8.4	85.0	8.3	5.0	0.8	1636.0	4.5	16637.0	2.6	2.5	29.9	0.7	5.2	8.7	42.4	17.2	AE92
06JAN2011:06:59	5.3	25.6	8.4	84.9	8.4	4.3	1.5	1637.1	3.9	16006.6	2.7	2.6	31.1	0.6	5.2	8.5	42.9	18.0	AE91
06JAN2011:07:59	5.3	26.4	8.3	84.9	8.4	4.3	1.5	1664.2	3.9	16479.4	2.7	2.6	31.1	0.6	5.2	8.5	43.2	18.1	AE91
06JAN2011:06:59	5.1	25.3	8.4	85.0	8.3	5.0	0.8	1622.1	4.5	16543.5	2.6	2.5	29.8	0.7	5.2	9.1	44.1	18.8	AE92
06JAN2011:07:59	5.1	25.1	8.3	85.1	8.3	5.0	0.8	1641.7	4.5	16836.4	2.6	2.5	29.8	0.7	5.2	9.1	44.2	18.9	AE92

- AE91 had 7.5GB (20%) large page allocated
- AE92 had **NO** Large Page allocated
- BP's were Pagefix=yes on Jan05 and Pagefix=no on Jan 6

IBM z196 zEnterprise

HIS Field Definitions

CPI	NUM	8	6.1	CYCLES*PER*INSTRUCTION
EFFGHZ	NUM	8	6.1	EFFECTIVE*GIGAHERTZ*CYCLES*PER NANO
ESTFINCP	NUM	8	6.1	ESTIMATED*CPI FROM*FINITE*CACHE/MEM
ESTICCPI	NUM	8	6.1	ESTIMATED*INSTRUCTION*COMPLEXITY*CPI
ESTSCP1M	NUM	8	6.1	ESTIMATED*SOURCING*CYCLES*PER L1 MISS
LPARBUSY	NUM	8	6.1	LPARCPU*PERCENT*CAPTURED AND*UNCAPTURED
L1MP	NUM	8	6.1	LEVEL*1*MISS*PERCENT
L15P	NUM	8	6.1	PERCENT*SOURCED*FROM*L1.5*CACHE
L2LP	NUM	8	6.1	PERCENT*SOURCED*FROM*L2*SAME BOOK
L2P	NUM	8	6.1	PERCENT*SOURCED*FROM*L2*CACHE
L2RP	NUM	8	6.1	PERCENT*SOURCED*FROM*L2*DIFFEERNT*BOOK
L3P	NUM	8	6.1	PERCENT*SOURCED*FROM*L3*SAME CHIP CACHE
L4LP	NUM	8	6.1	PERCENT*SOURCED*FROM*L4*SAME BOOK
L4RP	NUM	8	6.1	PERCENT*SOURCED*FROM*L4*DIFFERENT*BOOK
MEMP	NUM	8	6.1	PERCENT*SOURCED*FROM*MEMORY
MIPSEXEC	NUM	8		EXECUTED*MIPS
PRBSTATE	NUM	8	6.1	PERCENT*PROBLEM*STATE
PTEPCTMI	NUM	8	6.1	PAGETABLE*ENTRY*PCT OF TLB*MISSES
RNI	NUM	8	6.1	RELATIVE*NEST*INTENSITY
TLB1CYCL	NUM	8	6.1	CYCLES*PER*TLB*MISS
TLB1MISS	NUM	8	6.1	TLB*CPU MISS*PERCENT OF*TOTAL CPU

IBM z196 zEnterprise

DB2 V10 testing - RMF Large Page

System Summary

-- Memory Objects --			--------- Frames ----------				--- Area Used % ----		
Common	Shared	Large	Common	Fixed	Shared	1 MB	Common	Shared	1 MB
21	8	204	3485	718	48209	204	0.0	0.0	5.0

		Service		---- Memory Objects ---				Frames	----- Bytes -----		
Jobname	C	Class	ASID	Total	Comm	Shr	Large	1 MB	Total	Comm	Shr
BBNS001S	S	STCHI	0280	290	0	1	0	0	13.6G	0	50.0M
BBNS001	S	STCHI	0043	261	0	1	0	0	11.6G	0	50.0M
DBPADBM1	S	SYSSTC	0250	201	0	2	195	195	1185G	0	160G
BBN7ACRS	S	SYSSTC	0027	91	0	0	0	0	4043M	0	0
DBMADBM1	S	SYSSTC	0248	63	0	1	0	0	146G	0	128G
MQX9MSTR	S	SYSSTC	0266	26	0	0	0	0	347M	0	0
SMSPDSE1	S	SYSTEM	0009	24	0	0	0	0	88.0M	0	0
MQX6MSTR	S	SYSSTC	0219	20	0	0	0	0	206M	0	0
TRACE	S	SYSTEM	0004	19	0	0	0	0	19.0M	0	0
DBXADBM1	S	SYSSTC	0249	15	0	2	9	9	1185G	0	160G

IBM z196 zEnterprise

DB2 V10 testing - HIS Observations

Comparing AE91 to AE92

- TLB1MISS improved 1.8% with Large Page
- TLB1CYCL Improved 1.9% with Large Page
- PTEPCTMI showed a 6.9% improvement with Large Page

- CPI increased 4.7% with Large Page
- PRBSTATE showed an improvement of 3.5% with Large Page
- ESTICCPI increased 3.7%
- MIPSEXEC showed an improvement of 3.4% with Large Page

IBM z196 zEnterprise

DB2 V10 testing - HIS Conclusions

Large Page processing by DB2 showed an increase in the CPI (Cycles Per Instruction) which must be weighed by the fact that the PRBSTATE mix of instructions increased. The instruction complexity increased 3.7% and indicated by the ESTICCPI.

This indicates that the productive processing or application machine path (business logic vs service support) consumed a greater portion of our processor capacity.

Overall we seen to have gained about 2% in productivity which becomes meaningful in an installation with 124 CPs and 21 zAAP engines.

Another factor to consider is the fact that the AE91 LPAR sourced a L4 cache from the local book 16% less than AE92 (L4LP) and increase the access to L4RP (remote book) by 46%.

Based on these findings Large Page will be implemented at Aetna for DB2 V10.

IBM z196 zEnterprise

zIIP consumption in DB2 V10 was a pleasant finding as they were not observed in DB2 V9. The DBM1 address space shows zIIP consumption attributed to an Enclave that appears to be classified under the MSTR address space.

Samples: 1799 System: AE92 Date: 01/05/11 Time: 08.03.00 Range: 1800 Sec

		Service	--- Time on CP % ---			----- EAppl % -----		
Jobname	CX	Class	Total	AAP	IIP	CP	AAP	IIP
T8D1EASM	BO	ONLISO1	230.7	0.0	0.0	230.7		0.0
T8D1EASP	BO	ONLISO1	229.8	0.0	0.0	229.8		0.0
T8D1EASO	BO	ONLISO1	228.7	0.0	0.0	228.7		0.0
T8D1EASN	BO	ONLISO1	228.4	0.0	0.0	228.4		0.0
T8D1EASQ	BO	ONLISO1	227.7	0.0	0.0	227.7		0.0
T8D1EASR	BO	ONLISO1	227.6	0.0	0.0	227.6		0.0
DBPBDBM1	S	SYSSTC	18.2	0.0	0.0	18.3		25.9
DBUBP921	B	BATISO03	8.9	0.0	0.0	8.9		0.0
DBUBP925	B	BATISO03	8.9	0.0	0.0	8.9		0.0

Windsor Computer Center
CPU Configuration
Pre PSIFB

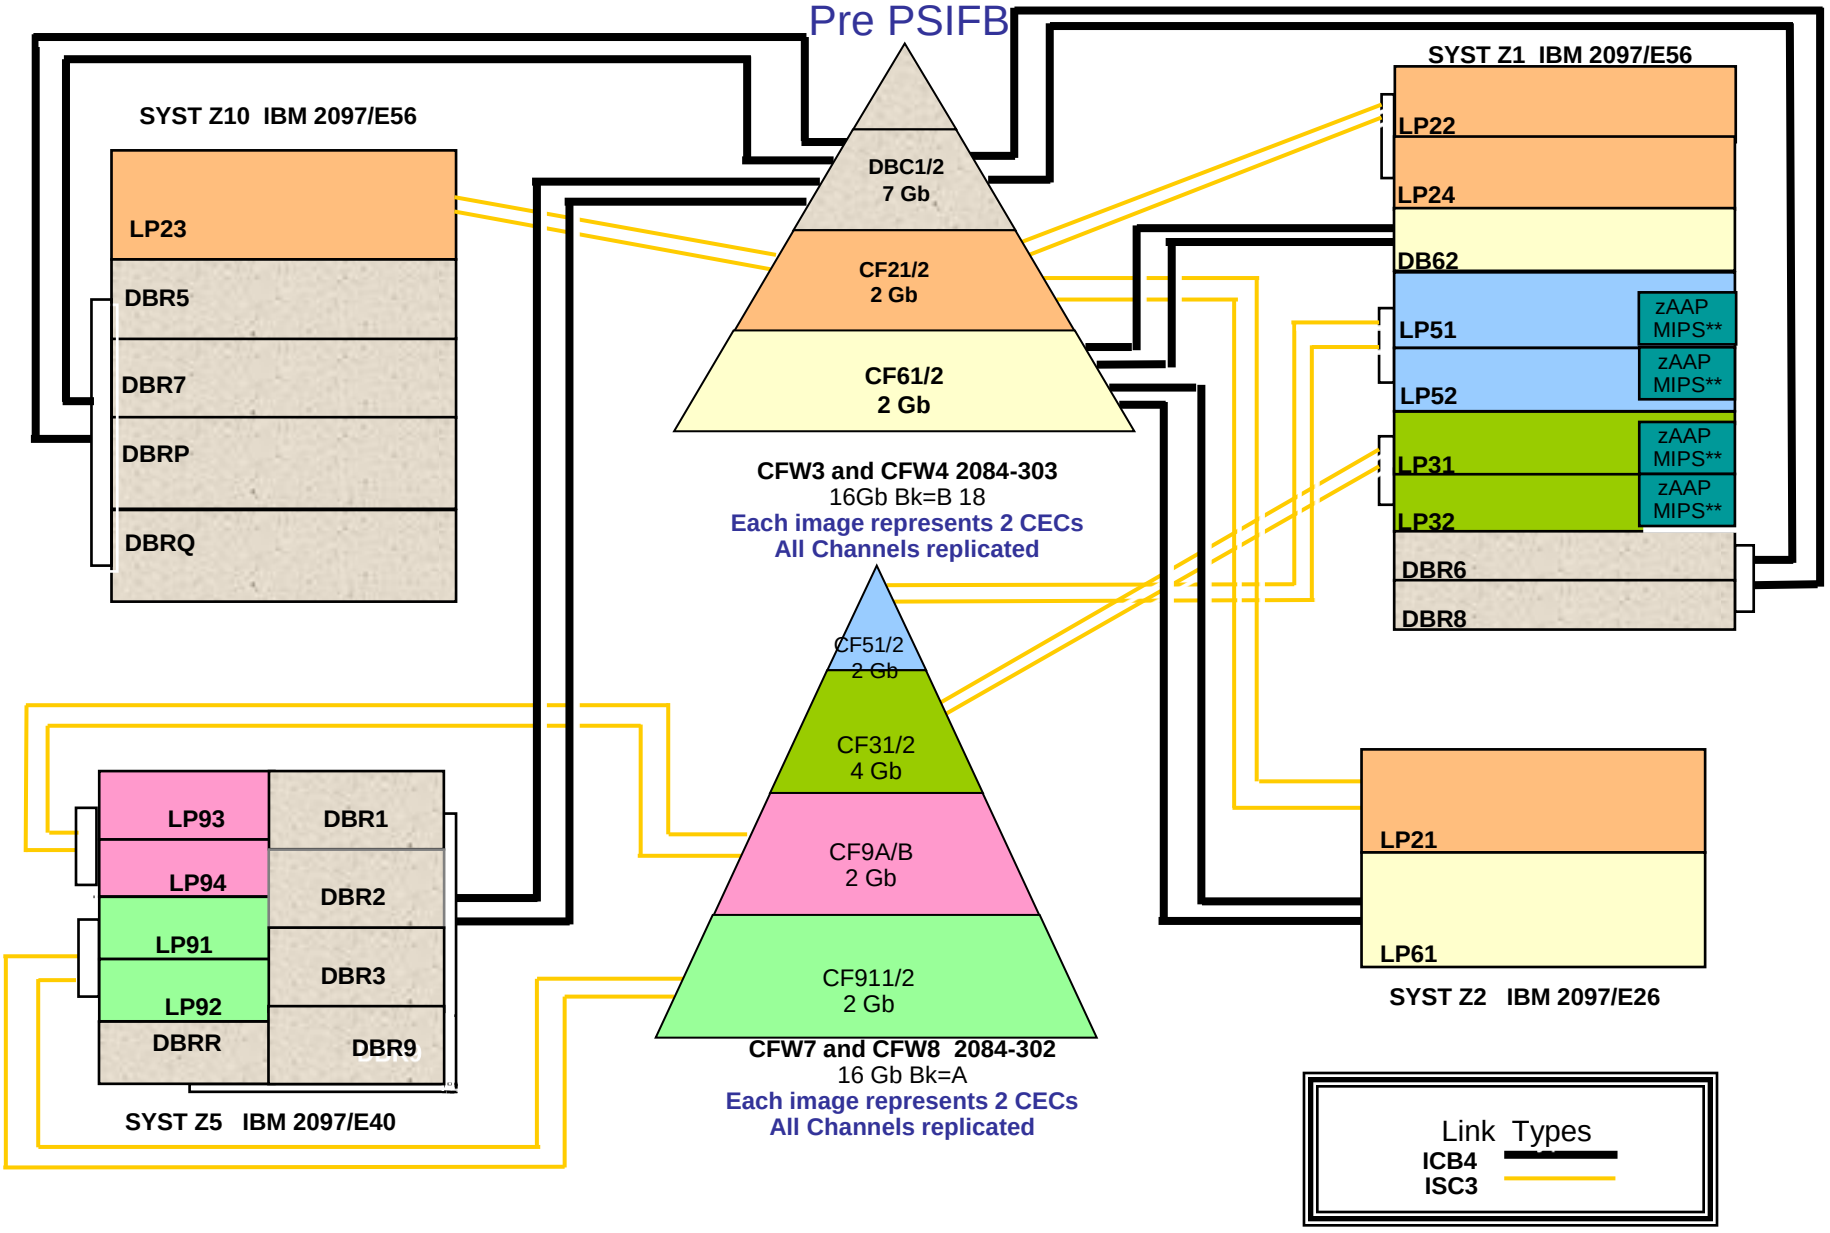

Windsor Computer Center
CPU Configuration
Pre PSIFB

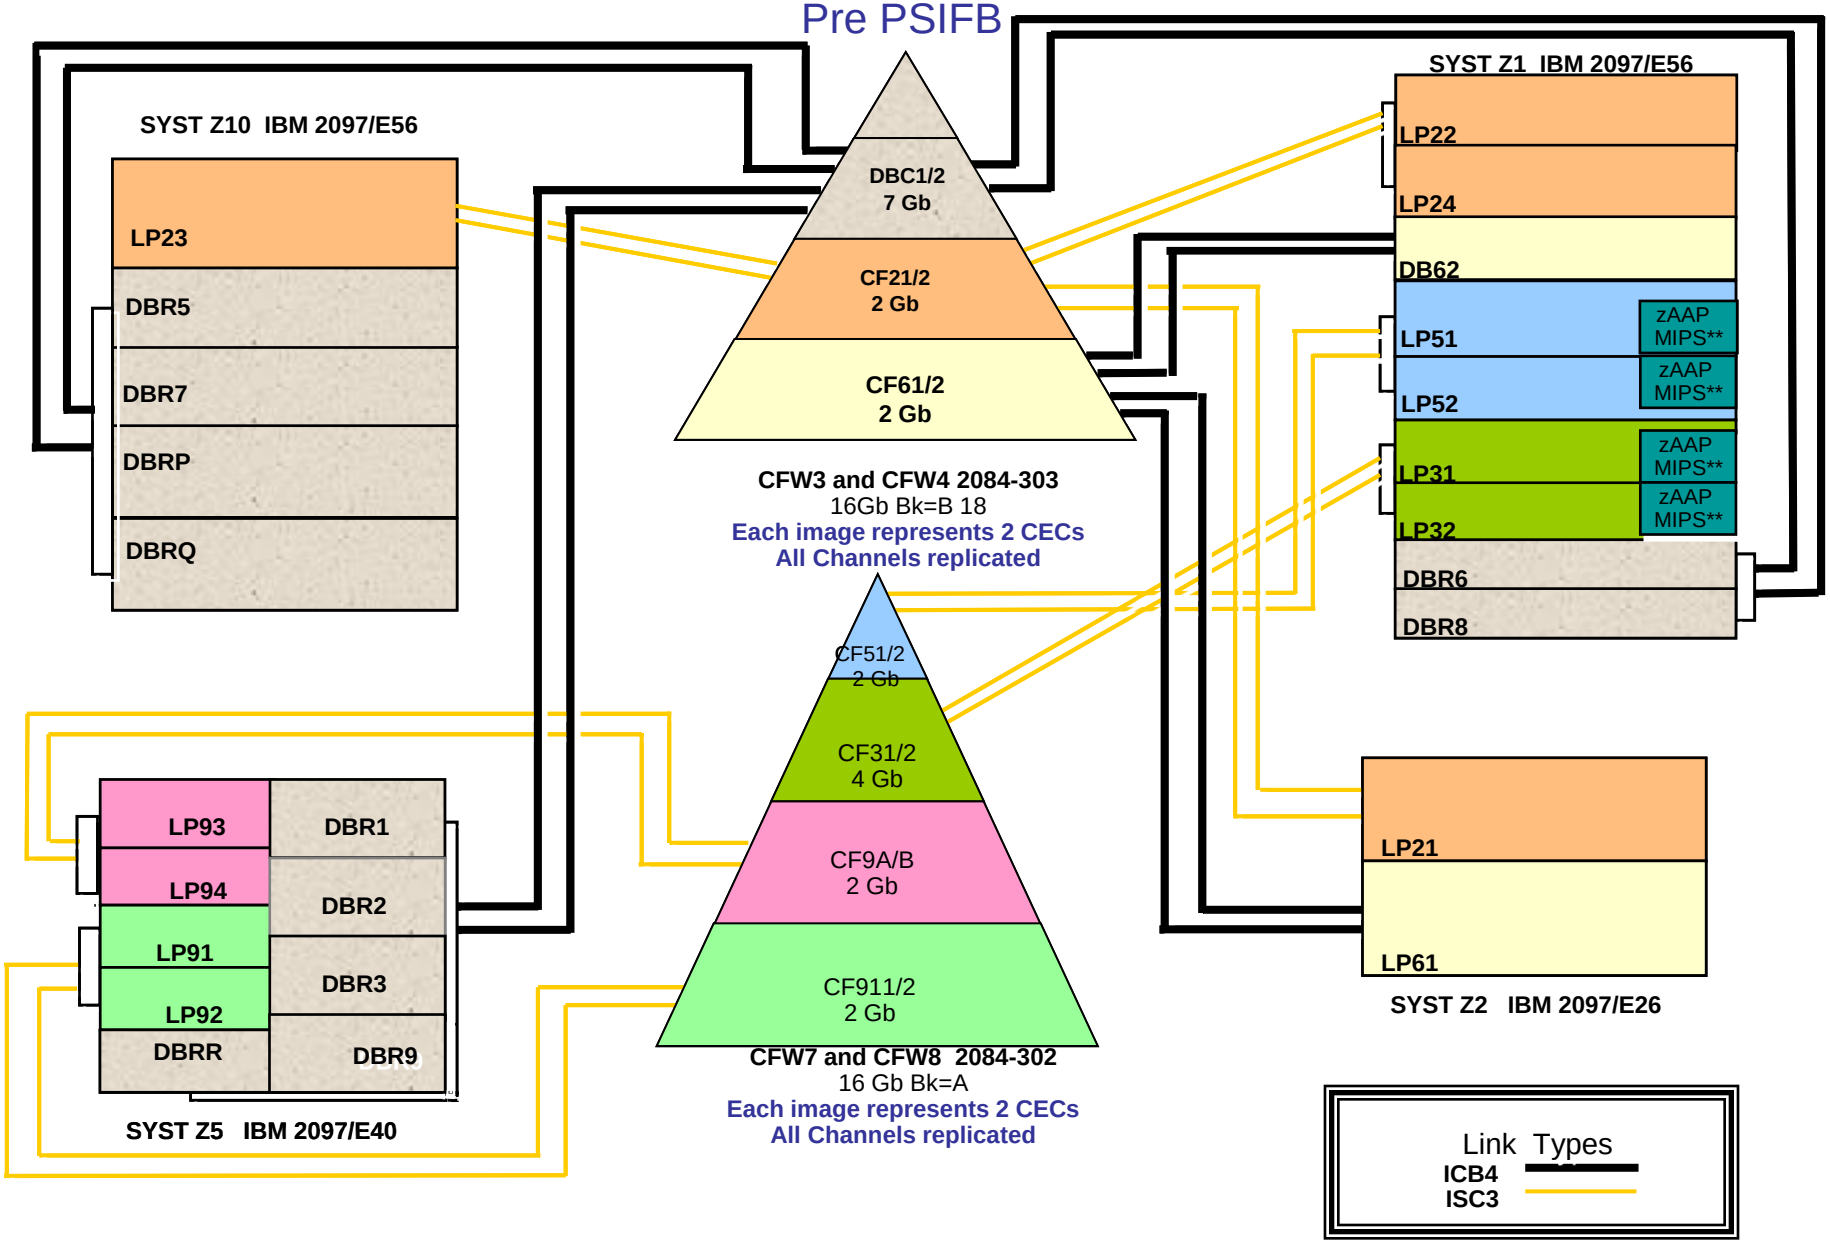

Windsor Computer Center
CPU Configuration
Post PSIFB

ZFW3 and ZFW4 2097-706

Each image represents 2 CECs
All ICB4 Channels replicated
(PSIFB unchanged to support both CF's)

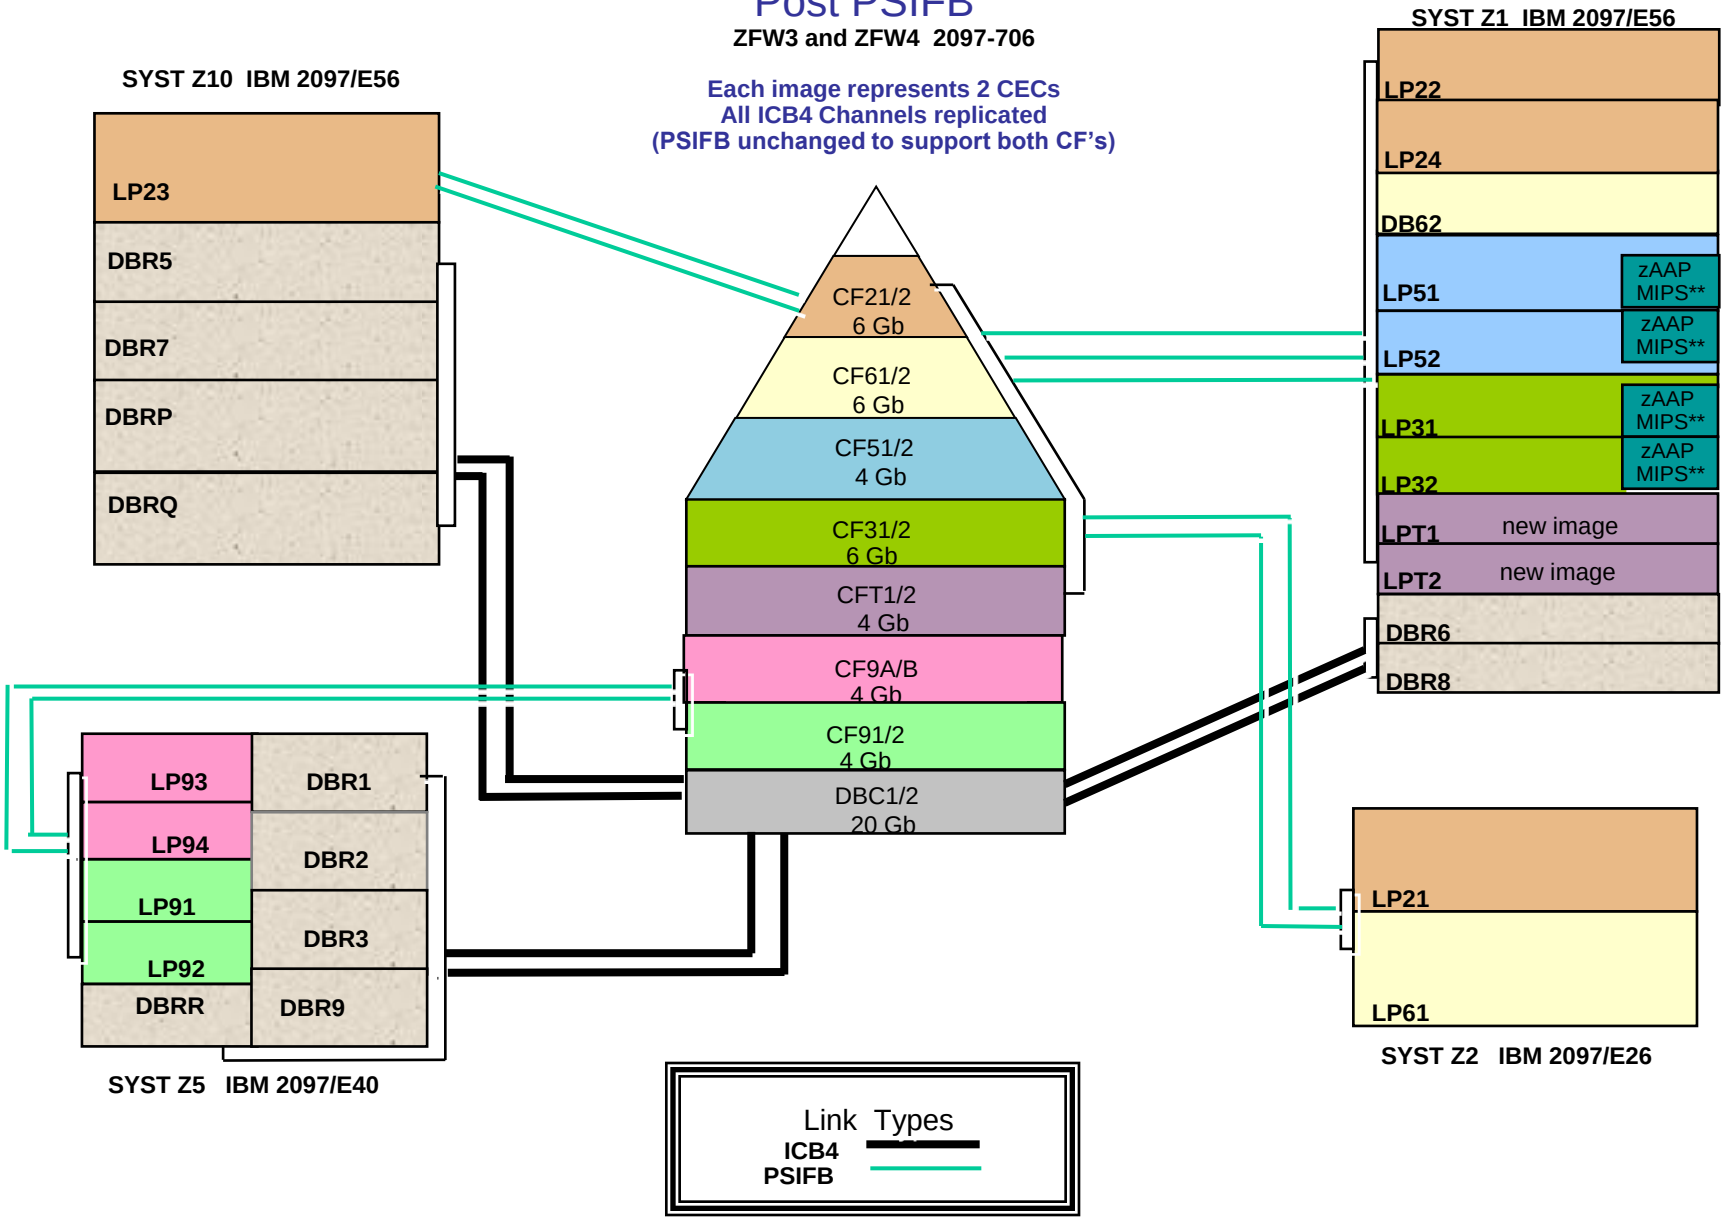

Windsor Computer Center

CPU Configuration

Post PSIFB

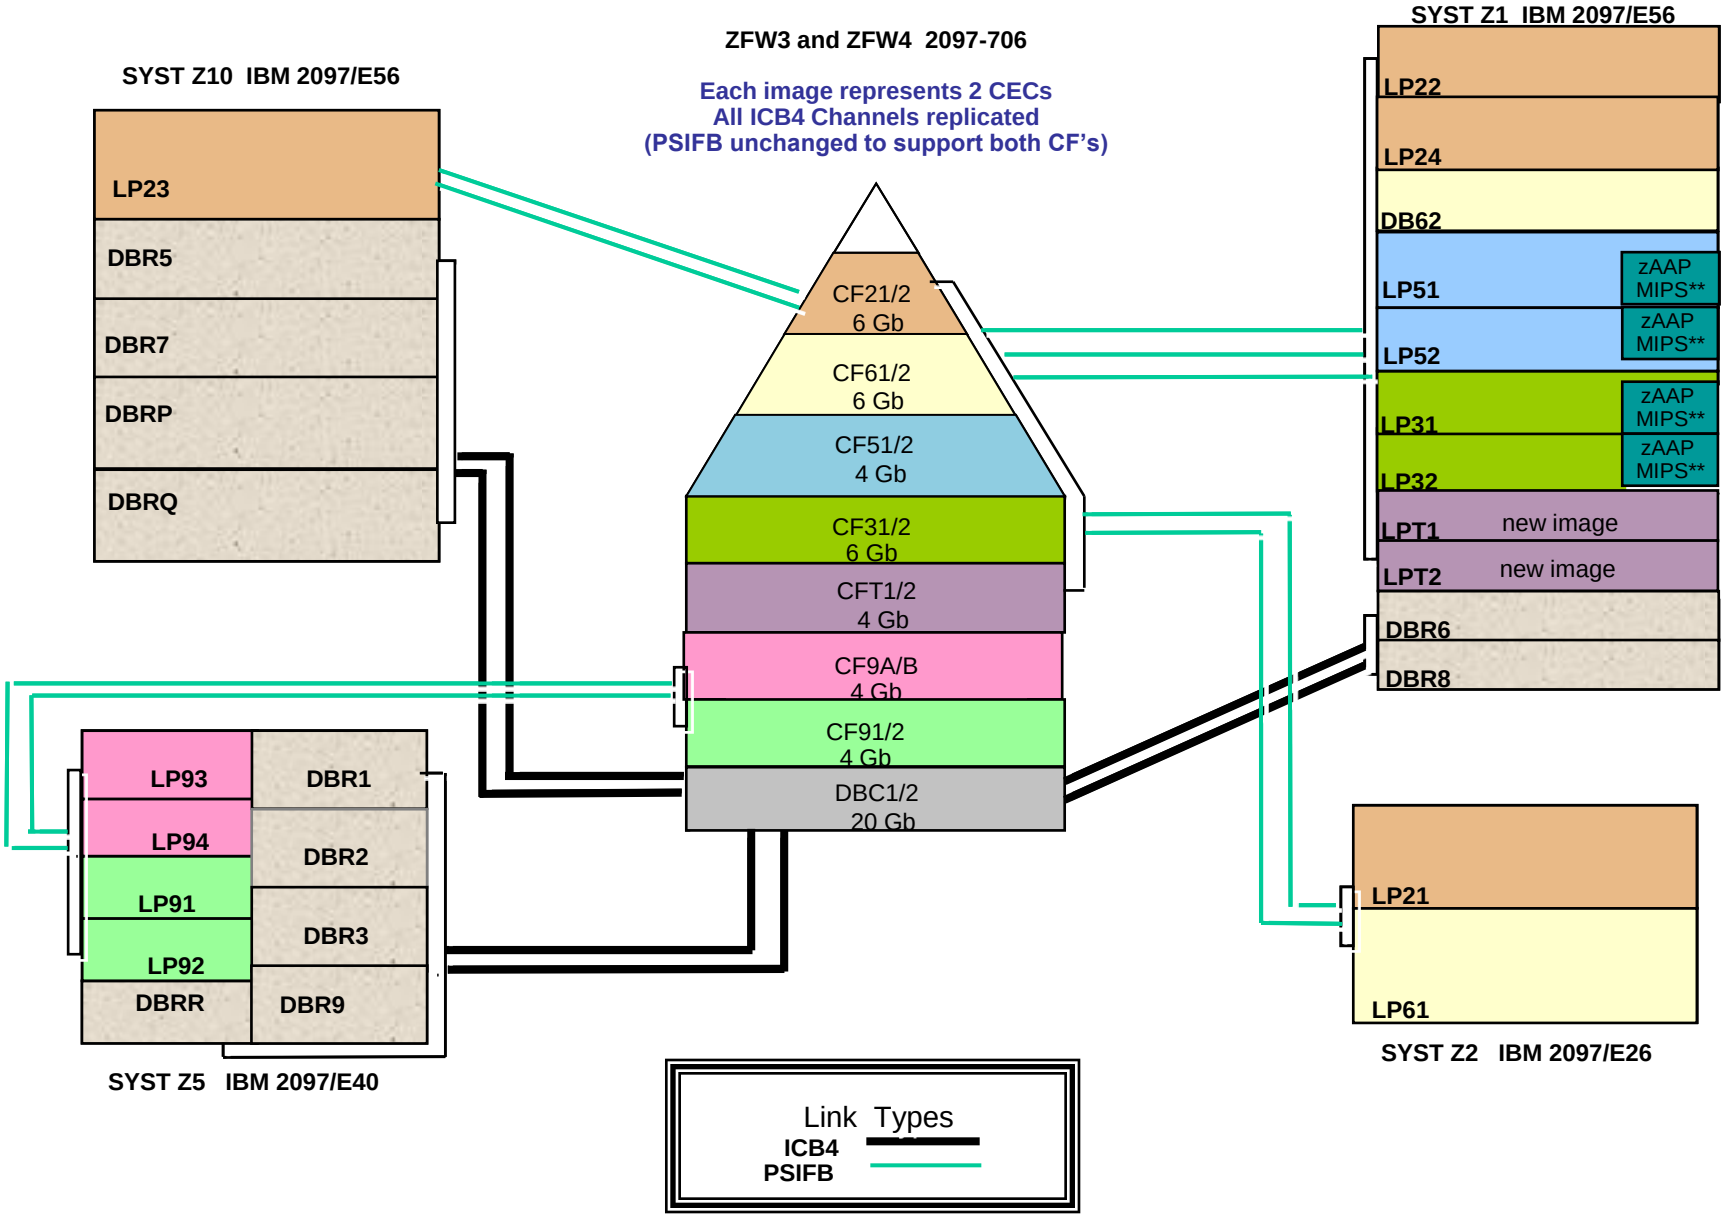

IBM z196 zEnterprise

PSIFB Connectivity

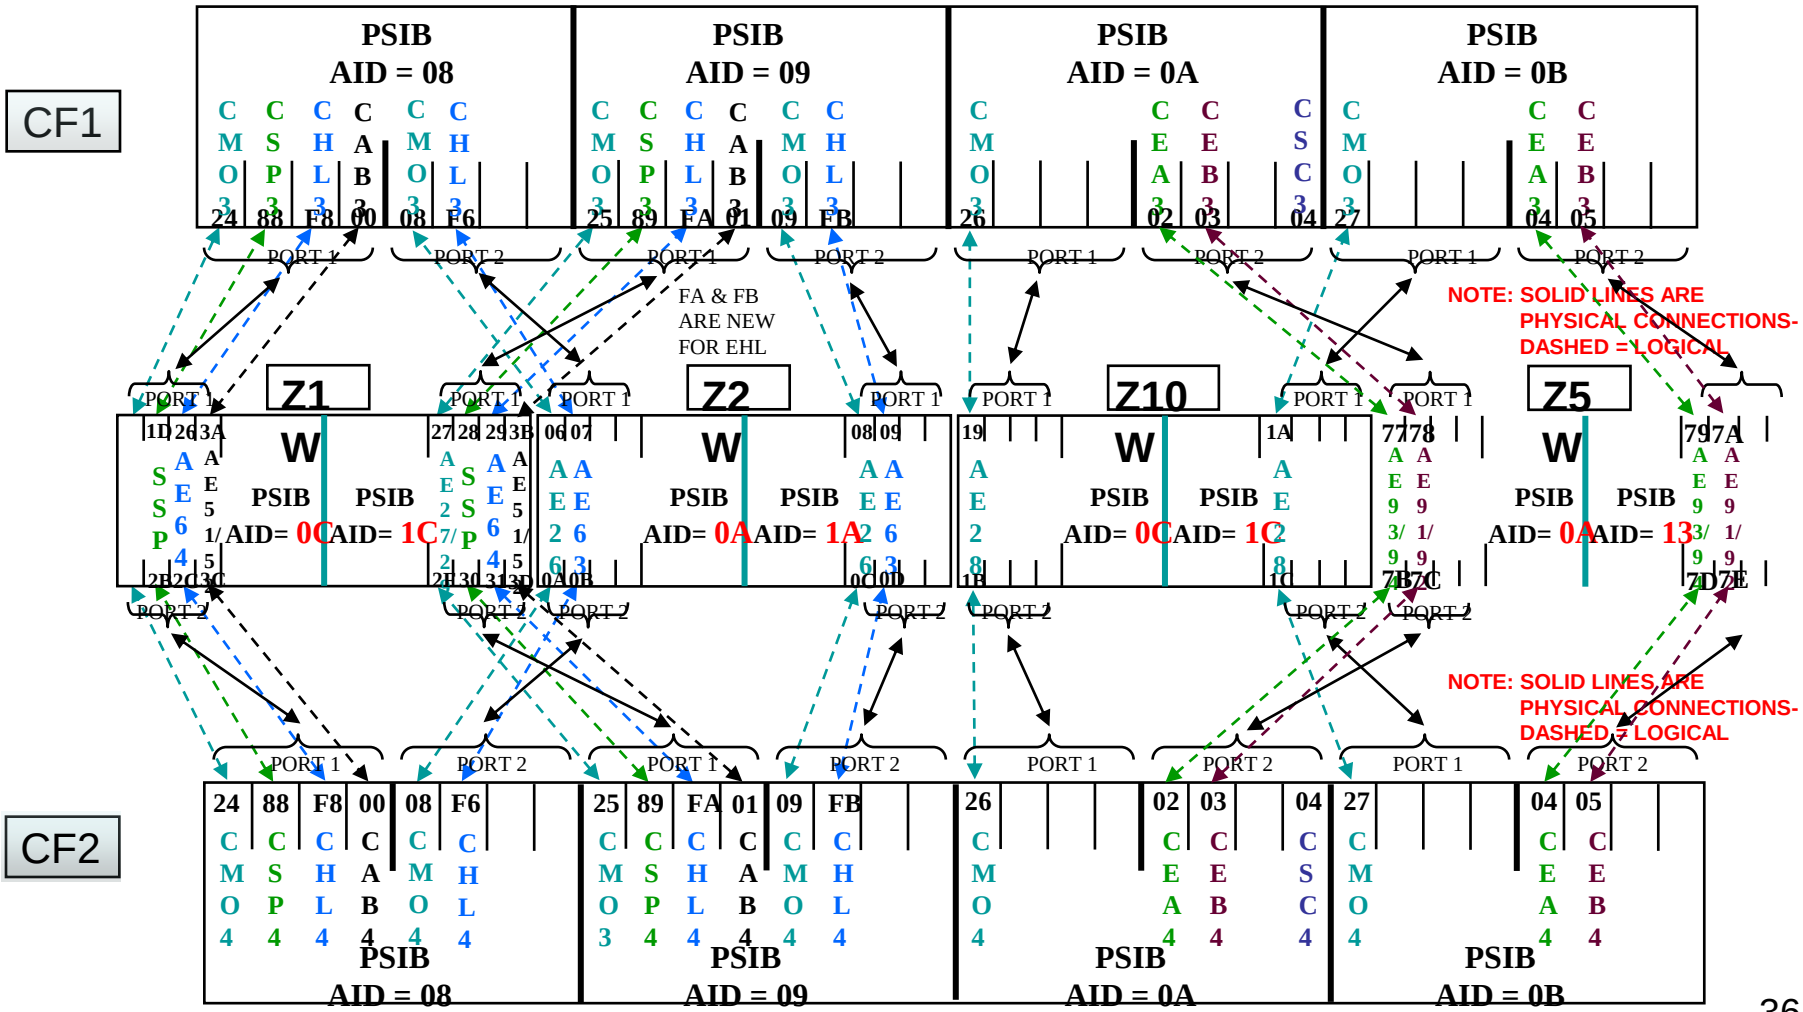

IBM z196 zEnterprise

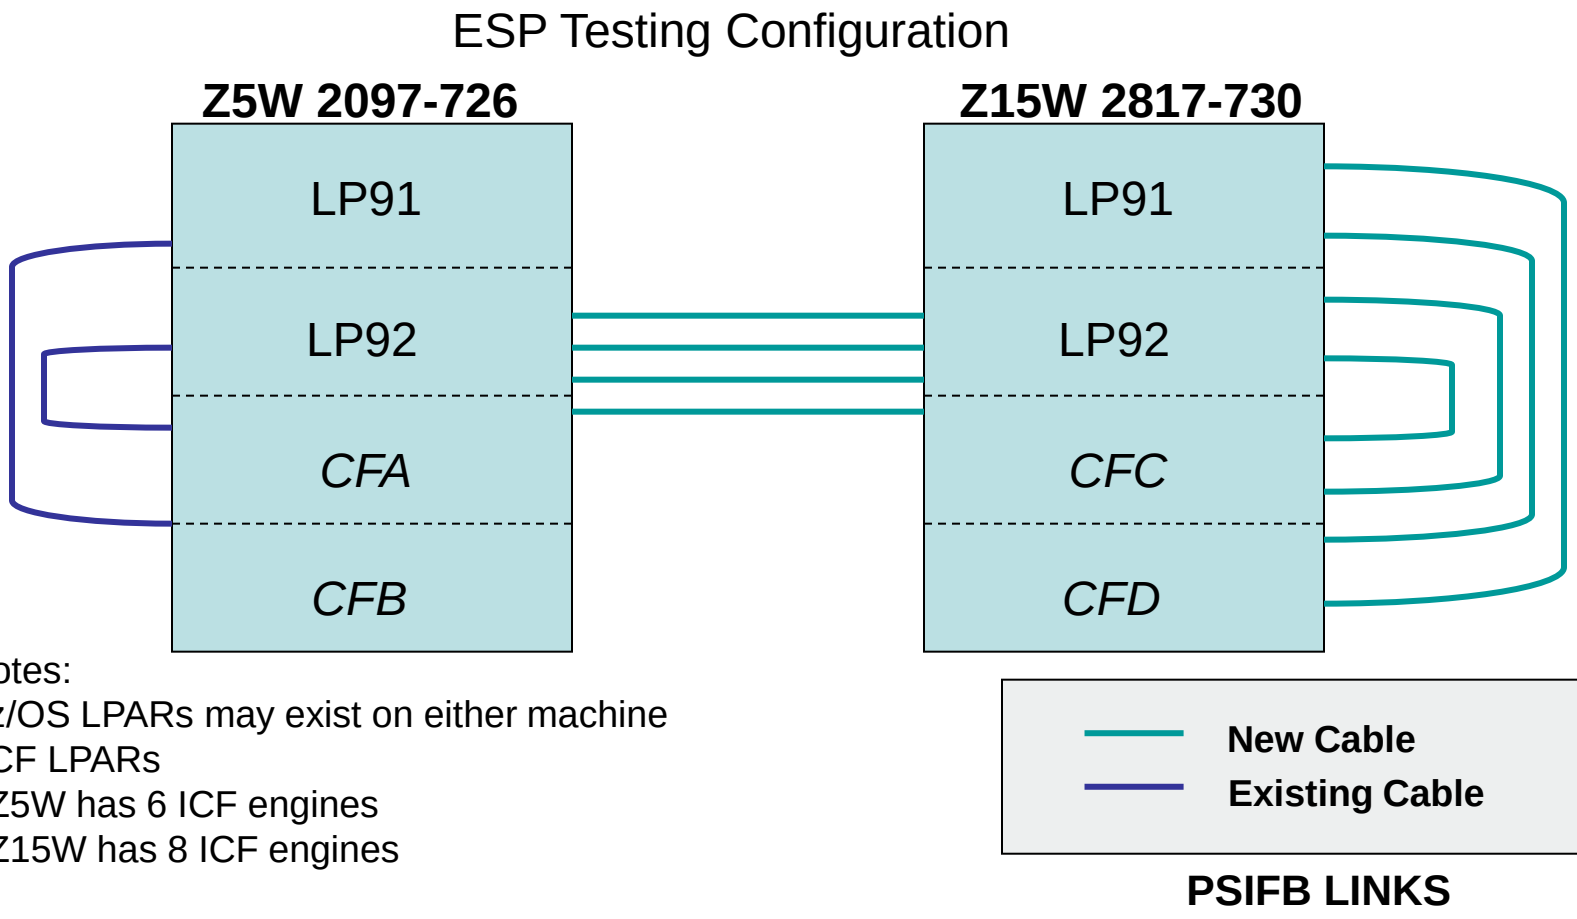

Notes:
- z/OS LPARs may exist on either machine
- CF LPARs
- Z5W has 6 ICF engines
- Z15W has 8 ICF engines

IBM z196 zEnterprise

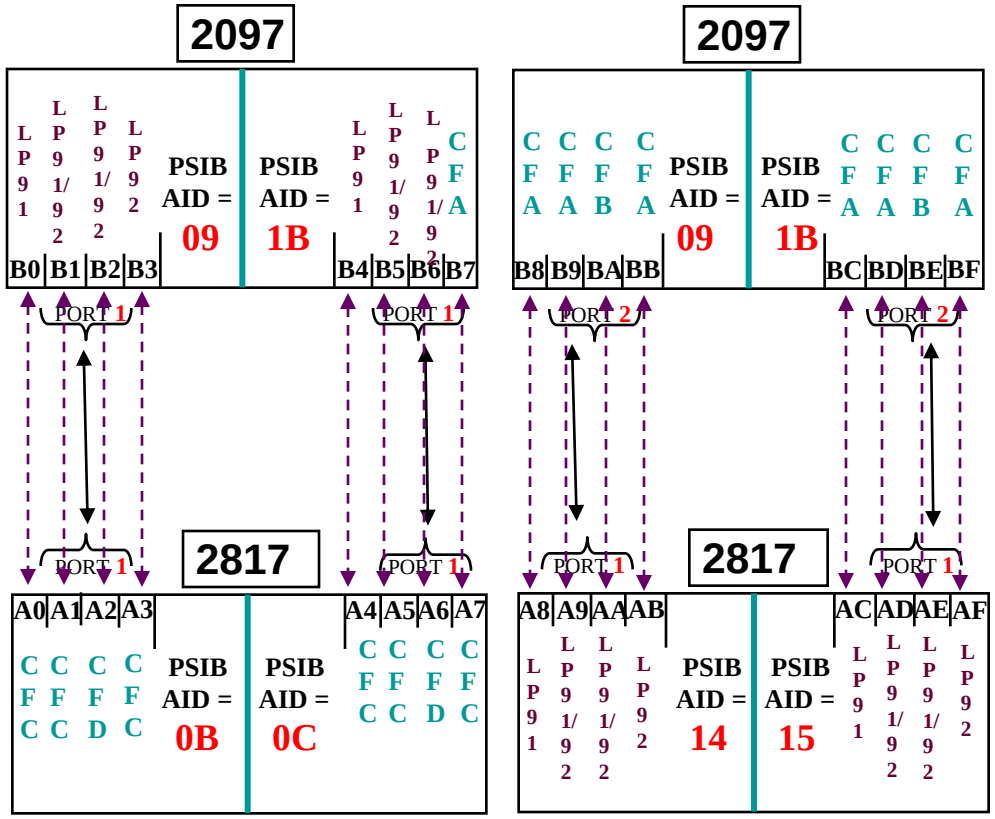

ESP - **Z5W TO Z15W** INFINIBAND CONNECTS

IBM z196 zEnterprise

Z15W ESP INFINIBAND CONNECTS

CPU = **Z-NEXT** (CHP/MODE/LPAR)

AID = 0B 20 Meter	**AID = 1B** 20 Meter	**AID = 0C** 10 Meter	**AID = 1B** 10 Meter	**AID = 14** 20 Meter	**AID = 1C** 20 Meter	**AID = 15** 10 Meter	**AID = 1C** 10 Meter
PORT 2	**PORT 1**	**PORT 2**	**PORT 2**	**PORT 2**	**PORT 1**	**PORT 2**	**PORT 2**
50 S 91	60 D CFC	54 S 91	64 D CFC	58 S 92	68 D CFC	5C S 92	6C D CFC
51 S 91/92	61 D CFC	55 S 91/92	65 D CFC	59 S 91/92	69 S CFC	5D S 91/92	6D D CFC
52 S 91	62 D CFD	56 S 91	66 D CFD	5A S 92	6A D CFD	5E S 91/92	6E S CFD
53 S 91/92	63 D CFD	57 S 91/92	67 D CFD	5B S 91/92	6B D CFD	**5F D CFC**	**6F D CFD**
2	**1**	**2**	**2**	**2**	**1**	**2**	**2**

NOTE: CF LPAR's 'CFC' & 'CFD' ARE IN LCSS 1 (91/92 ARE IN LCSS 0)

IBM z196 zEnterprise

My First Command

```
16:59:57.66           D XCF,C
16:59:57.67           IXC357I  16.59.57  DISPLAY XCF 333
                      SYSTEM AE92 DATA
                         INTERVAL    OPNOTIFY     MAXMSG    CLEANUP       RETRY    CLASSLEN
                              165         168       4096         15          10         956

                         SSUM ACTION  SSUM INTERVAL  SSUM LIMIT     WEIGHT  MEMSTALLTIME
                              PROMPT            168         N/A          1            NO

                         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/ASYNC THRESHOLD:                26
                         DUPLEX SYNC/ASYNC THRESHOLD:                 26
                         SIMPLEX LOCK SYNC/ASYNC THRESHOLD:           26
                         DUPLEX LOCK SYNC/ASYNC THRESHOLD:            28
```

IBM z196 zEnterprise

Pre Z196 - Z10 All ICB4 (1 SYSPLEX)

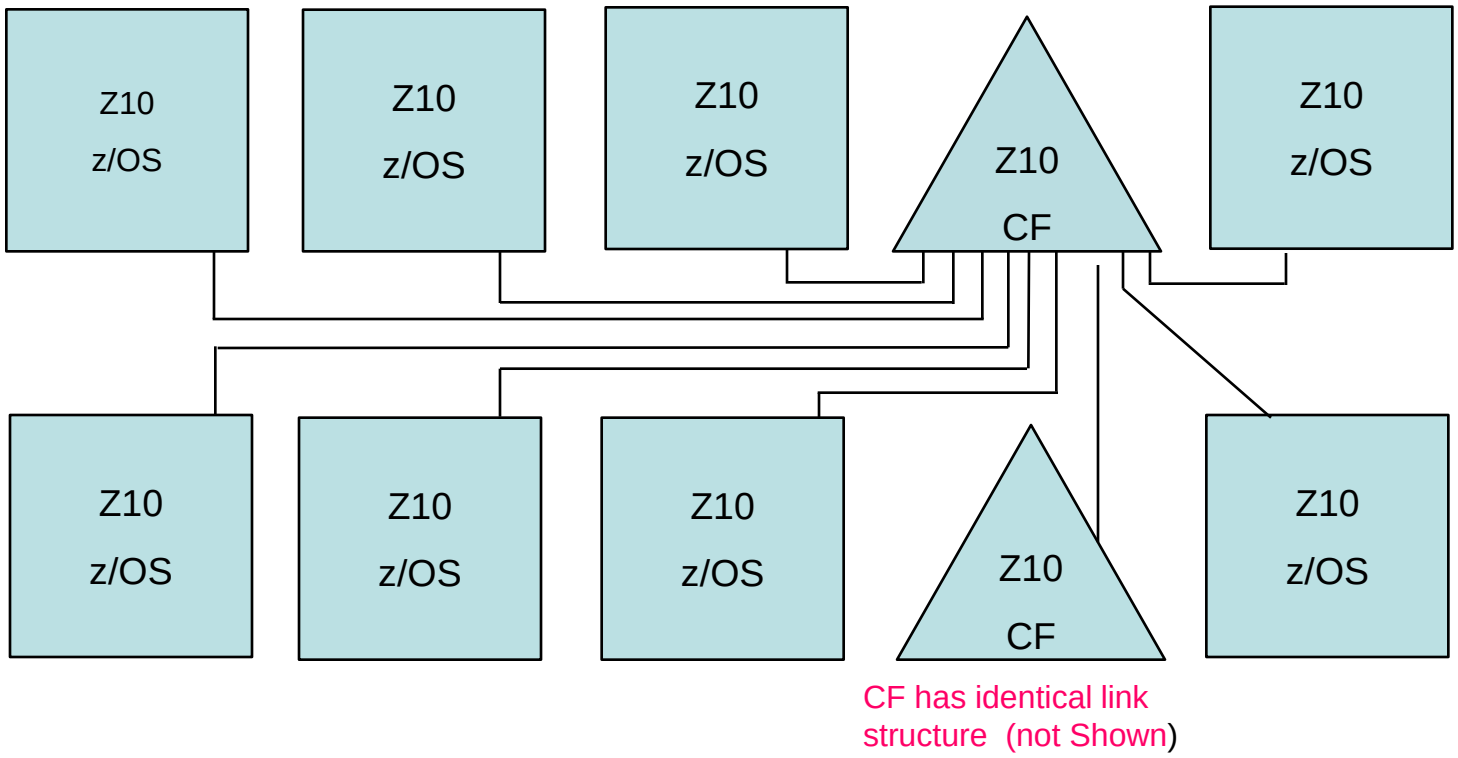

To prepare for roll of the floor the Z10 z/OS CPCs received HCA2-O PSIFB cards

IBM z196 zEnterprise

Z196 and PSIFB connectivity introduced to the SYSPLEX

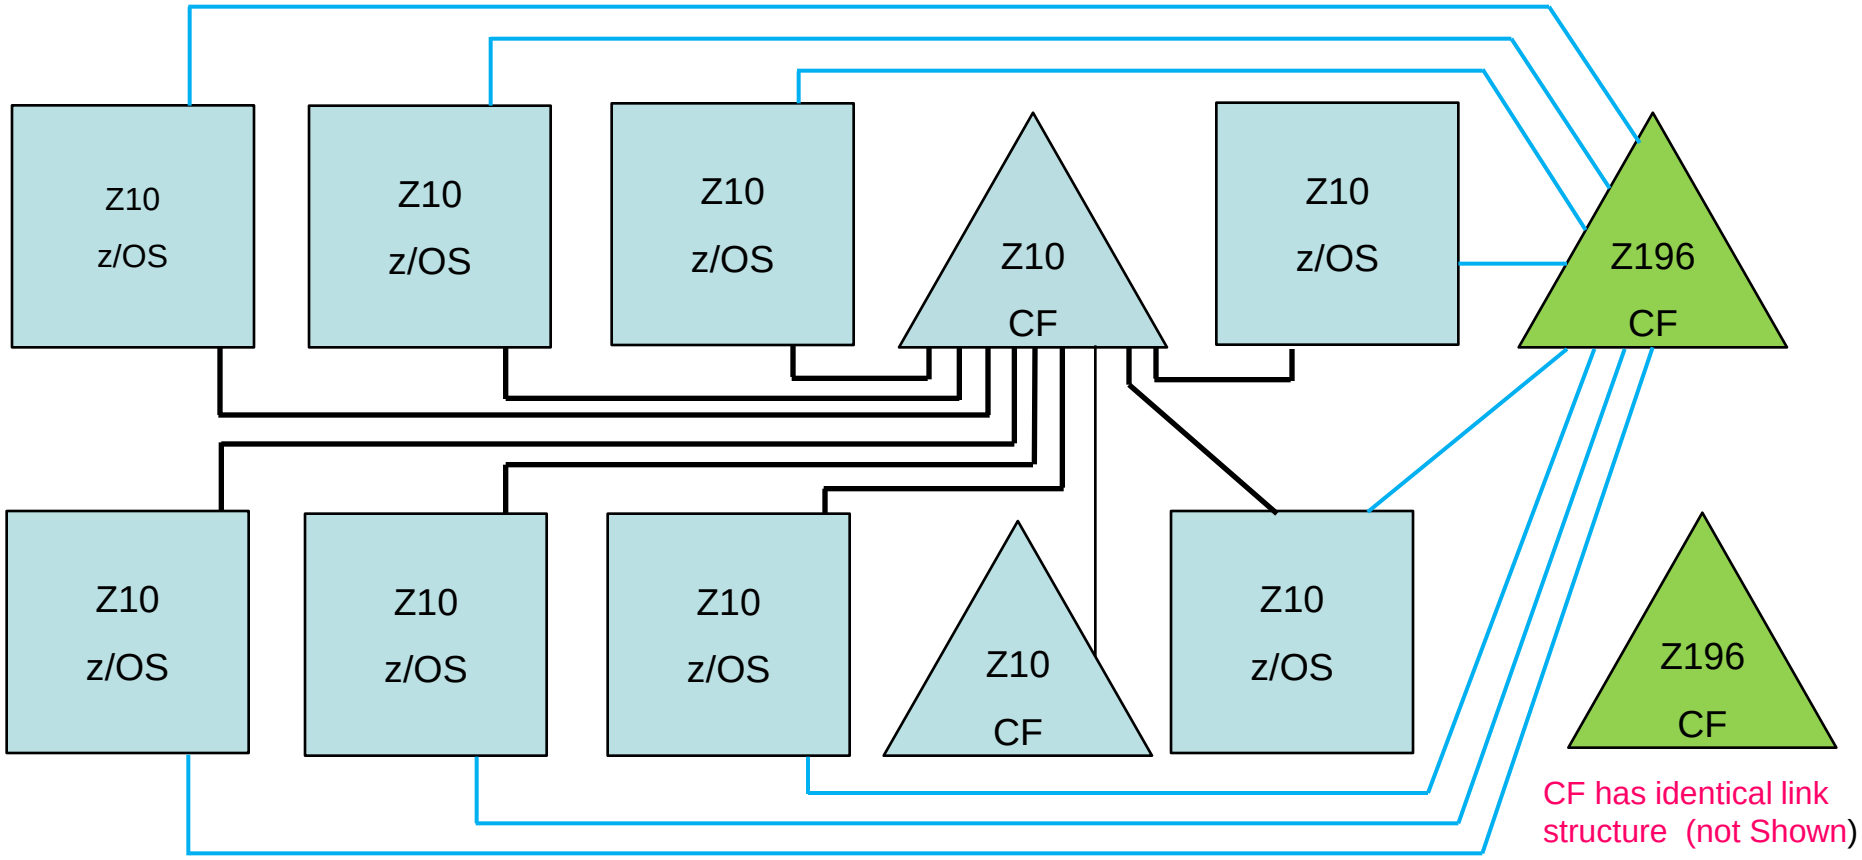

New Z196 CF CPCs installed and PSIFB connectivity was implemented in addition to the Z10 ICB4 links.

IBM z196 zEnterprise

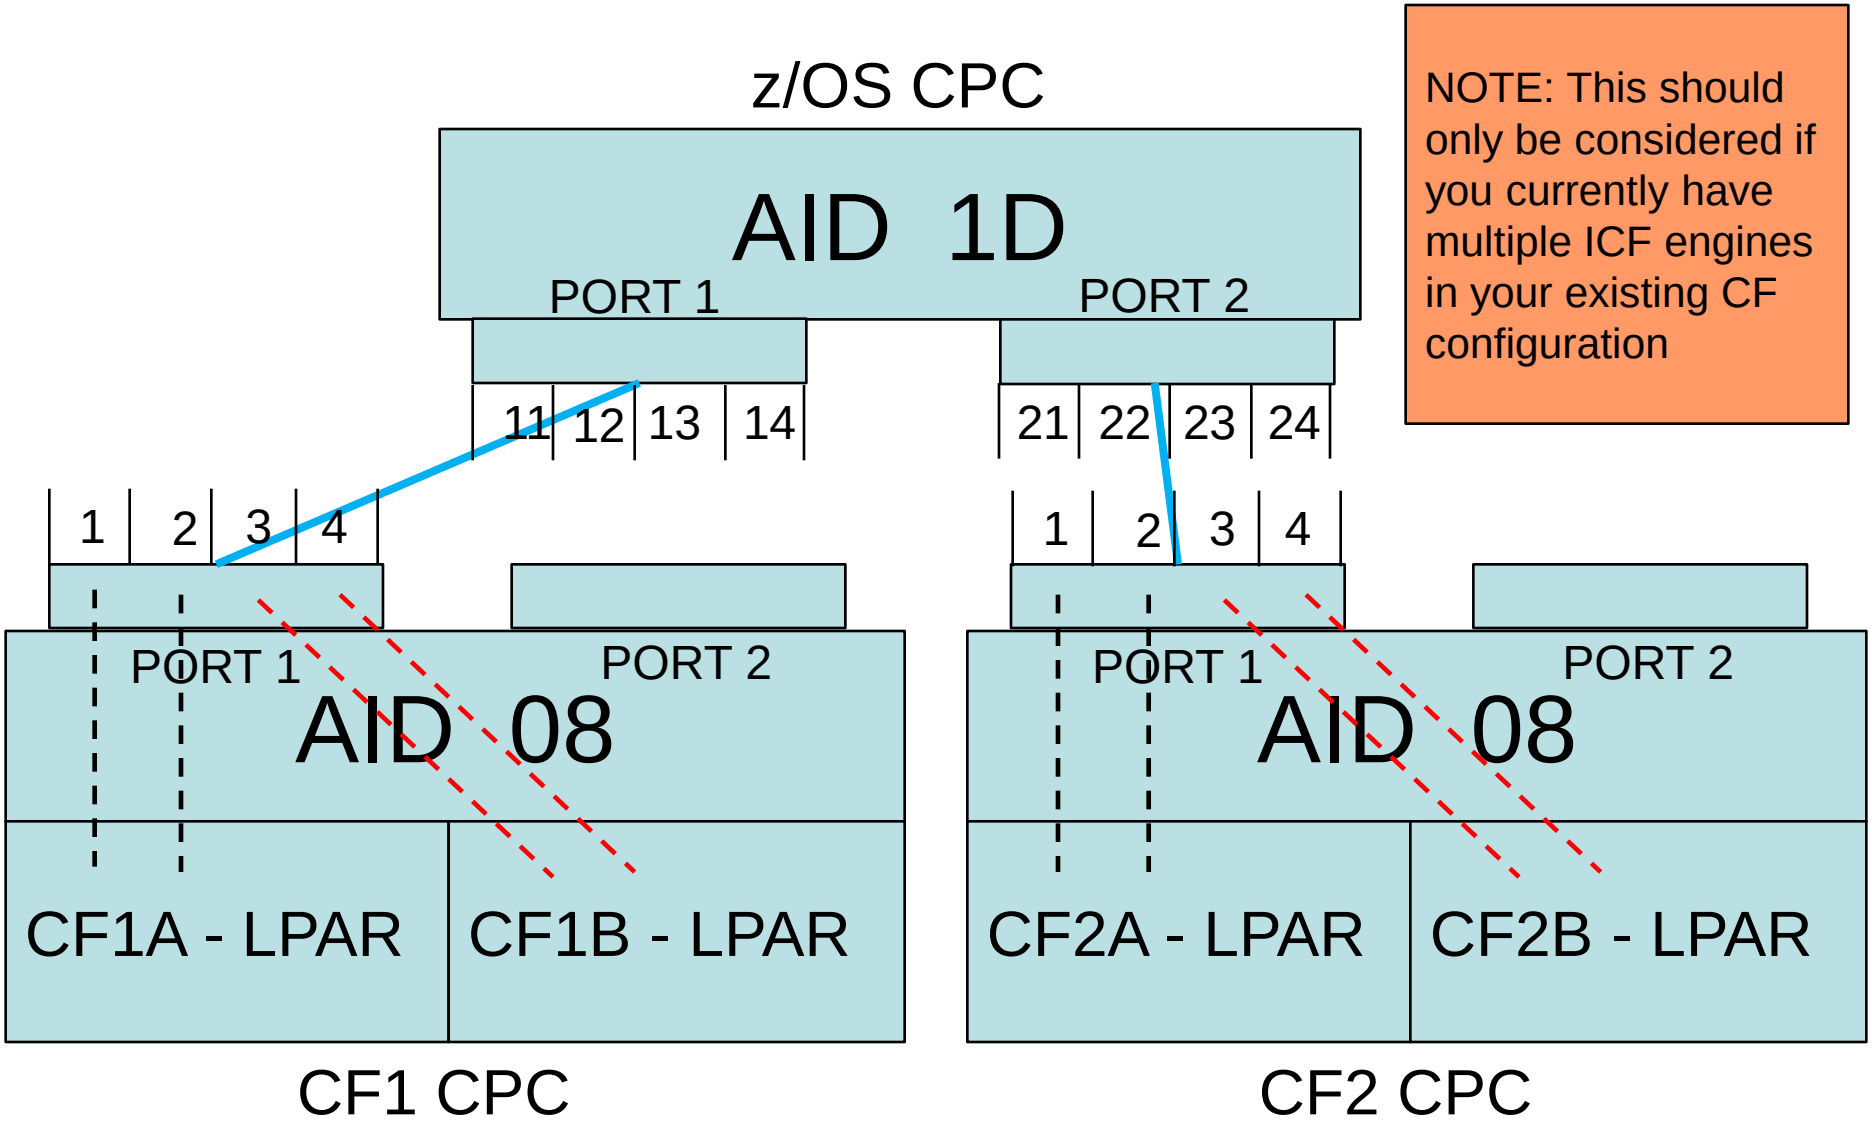

NOTE: This should only be considered if you currently have multiple ICF engines in your existing CF configuration

IBM z196 zEnterprise

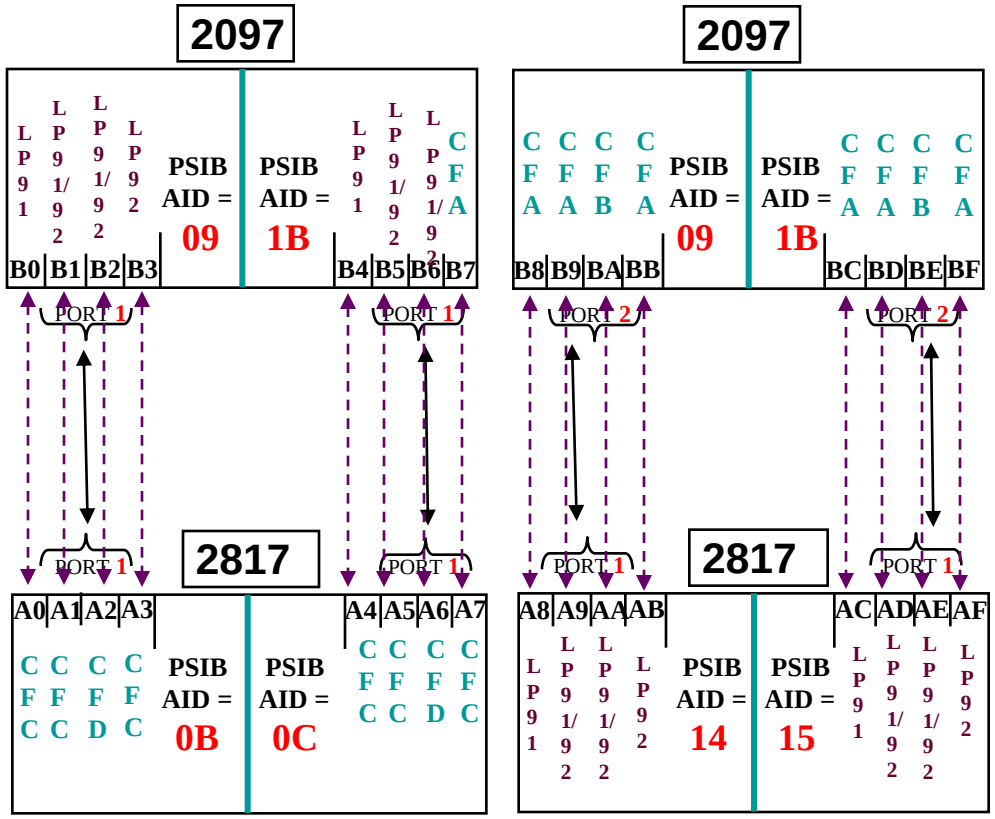

ESP - **Z5W TO Z15W** INFINIBAND CONNECTS

IBM z196 zEnterprise

Switching from ICB4 to PSIFB

- Alter the CFRM Policy to include the new Coupling Facility LPARs
- Place Z10 CF's in Maintmode
  SETXCF START,MAINTMODE,CFNAME=(CF1,CF2)
- Move the structures to the new Z196 CF's
  SETXCF START,REALLOCATE
- Verify structures moved
  D XCF,CF,CFNAME=(CF1,CF2)
- Observe
- To backoff
  SETXCF STOP,MAINTMODE,CFNAME=(CF1,CF2)
  SETXCF START,MAINTMODE,CFNAME=(CF1A,CF1B,CF2A,CF2B)
  SETXCF START,REALLOCATE

IBM z196 zEnterprise

Results/Recommendations

- INFINIBAND performed better than our ICB4 environment / rec: at least 28 Subchannel Buffers for each LPAR. With ICB4 links I had bursts of activity that would overrun my subchannel buffers.
- CF Processor Utilization decreased dramatically due to a reduction in the MP effect and the increased cycle rate on the z196
  We went from 2 CF only CPCs with 5 dedicated engines on 2 LPARS, to a configuration with 2 LPARs on each CPC… Each CPC had a 2 way and 3 way dedicated engine configuration.
  The 3 way is targeted with our “Loved ones” (SYNC)
  The 2 Way gets the ASYNC traffic.

Note: CFs only know they have a request. They do not know if the request is SYNC or ASYNC, but the z/OS LPARs do. Sharing SYNC and ASYNC requests tends to increase the ASYNC service times.

IBM z196 zEnterprise

Samples: 60 Systems: 10 Date: 04/20/11 Time: 16.01.00 Range: 60 Sec

CF Policy: POLICY2 Activated at: 04/19/11 01.02.01

Coupling Facility					Processor				Request	Storage	
Name	Type	Model	Lvl	Dyn	Util%	Def	Shr Wgt	Effect	Rate	Size	Avail
CFM1	2097	E12	16	OFF	0.0	5	0	5.0		30G	30G
CFM2	2097	E12	16	OFF	0.0	5	0	5.0		30G	30G
CF1A	2817	M15	17	OFF	12.3	3	0	2.8	92662	30G	22G
CF1B	2817	M15	17	OFF	8.4	2	0	1.9	18312	22G	21G
CF2A	2817	M15	17	OFF	24.9	3	0	2.8	85490	30G	21G
CF2B	2817	M15	17	OFF	13.9	2	0	1.9	30850	22G	21G

IBM z196 zEnterprise

Lock Structure Comparison - Peak Hour

CF Structure Name=DSNDB3G_LOCK1 Hour of Day=14

Week of Year	Day of Month	Reqs - Synch	Req Time - Synch	Reqs - Asynch	Req Time - Asynch	Defined Processors Utilization	Requests Changed from Synch to Asynch	AVESYTM	Reqs - Asynch	AVEASYTM	Requests Completed - Total
16	14	184834753	2292.5972	38804	3.418064	13.22 %	4469	.000012403	38804	.000088085	200046152
17	20	200053247	1941.6895	601	0.245176	8.64 %	351	.000009706	601	.000407947	212059129

- Total and SYNC request rate increased
- Total SYNC service time decreased
- Average SYNC service time dramatically reduced

IBM z196 zEnterprise

Summary of 12 hour Weekday Activity for SYNC and ASYNC Activity

Sysplex Name=AEPLEX04 Time Zone=1

Obs	Year of Century	Week of Year	Day of Month	_TYPE_	_FREQ_	Reqs - Synch	Req Time - Synch	Reqs - Asynch
1	11	16	13	0	4184	5403693742	79379.144	1646294832
2	11	17	19	0	4207	5885592760	69484.639	1665926063
ZONE						11289286502	148863.78	3312220895

Obs	Req Time - Asynch	Requests Changed from Synch to Asynch	Requests Completed - Total	AVESYTM	AVEASYTM
1	91693.984	2753375	7173079565	.000014690	.000055697
2	79962.040	841730	7791578304	.000011806	.000047999
ZONE	171656.02	3595105	14964657869		

IBM z196 zEnterprise

Questions?