



IBM Washington Systems Center

z/OS 1.6 WSC Early Support Experience



Riaz Ahmad
IBM Washington Systems Center
Gaithersburg, Maryland

© 2004 IBM Corporation

IBM Washington Systems Center



Trademarks

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

APPN*	IBM logo*	Virtual Image Facility
DB2*	IMS	VM/ESA*
e-Business logo*	Mugstar*	VSE/ESA
Enterprise Storage Systems	IMS	VTAM*
ESCON*	Netfinity*	WebSphere
FICON	OS/390*	z/Architecture
GDPS	Parallel Sysplex*	z/OS
Geographically Dispersed Parallel Sysplex	PR/SM	z/VM
HiperSockets	S/390*	zSeries
IBM*	S/390 Parallel Enterprise Server	



* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation

LINUX is a registered trademark of Linus Torvalds

Penguin (Tux) complements of Larry Ewing

Tivoli is a trademark of Tivoli Systems Inc.

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

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.

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.

IBM considers a product "Year 2000 ready" if the product, when used in accordance with its associated documentation, is capable of correctly processing, providing and/or receiving date data within and between the 20th and 21st centuries, provided that all products (for example, hardware, software and firmware) used with the product properly exchange accurate date data with it. Any statements concerning the Year 2000 readiness of any IBM products contained in this presentation are Year 2000 Readiness Disclosures, subject to the Year 2000 Information and Readiness Disclosure Act of 1998.

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.

© 2004 IBM Corporation

Agenda

- z/OS 1.6 Install
- What we tested
- Overall experience



Who we are and what we do ?

- The IBM Washington Systems Center
 - Advanced Technical Support, IBM Americas
- Provide Technical Support
 - Early Support Programs
 - World-wide Q&A support
 - Critical situation support
 - Product introductions
 - White papers, flashes
 - etc...

zSeries Architecture Level Set

- z/OS Release 1.6 and future releases will no longer support the ESA/390 Architecture
- The z/OS 1.6 and higher will support only the following servers or equivalent
 - IBM @server zSeries servers
 - z800, z900, z890 and z990
 - Otherwise Waitstate 088-14
- The following servers do not support z/Architecture™ and will not be able to run the z/OS 1.6, planned GA for September 2004:
 - S/390 Parallel Enterprise Server™ - Generation 5 (G5) and Generation 6 (G6)
 - All models of the Multiprise® 3000 Enterprise Server or comparable server and other previous processors

z/OS Support



		G3-G4	G5/G6 Multiprise 3000	z800	z890	z900	z990	End of Service	Coexists with z/OS...	Planned Ship Date
OS/390	2.10	X	X	X	X ^c	X	X ^c	09/04	1.4	
z/OS	1.1		X	X		X		3/04	1.4	
	1.2		X	X	X ^c	X	X ^c	10/04	1.5	
	1.3		X	X	X ^c	X	X ^c	3/05	1.6	
	1.4		X	X	X	X	X	3/07	1.7	
	1.5		X	X	X	X	X	3/07	1.8	
	1.6			X	X	X	X	9/07	1.8	9/04
	1.7			X	X	X	X	9/08	1.9	9/05

x^c - Compatibility support only

IBM Bimodal Accomodation Offering is available for z/OS 1.2, 1.3, and 1.4. It will not be provided for z/OS 1.5

> 1.1 was withdrawn from service in March 2004

> 1.4 was orderable through September 9, 2004

> 1.6 became orderable September 10, 2004

1.5 was also orderable through September 9, 2004

z/OS.e V1R3/R4/R5/R6/R7 will only run on a z800 and z890 server

Positioning for z/OS 1.6

- Read Documentation and PSP buckets
 - z/OS Migration, z/OS and z/OS.e: Planning for Installation
 - Software PSP buckets: ZOSV1R6: ZOSGEN, SERVERPAC, ..
 - Hardware PSP buckets: 2084DEVICE, 2086DEVICE, 2064DEVICE, 2066DEVICE
- DASD Storage Requirements

	z/OS 1.4	z/OS 1.5	z/OS 1.6
Target	4,646	5,160	5,277
DLIBs	6,295	7,212	7,338
HFS	2,200	2,200	2,800

- Ensuring System requirements are satisfied
 - Hardware & Software for Driving and Target systems
 - Coexistence requirements
- Migration Actions

Driving System Requirements - ServerPac

- Minimum Any one of the following:
 - z/OS V1.3 with PTFs
 - Customized Offerings Driver (subset of z/OS V1.4 with SMP/E V3.2)
- Minimum requirement to service the new target system after the initial ServerPac install
 - SMP/E 3.3
 - z/OS 1.6 Binder
 - z/OS 1.6 HLASM

Target System Requirement

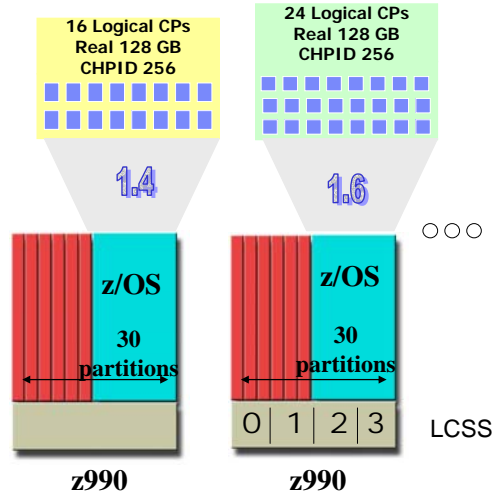
- Hardware Requirements
 - Processor Requirements:
 - z/OS V1.6 runs on an IBM zSeries Server only
 - Coupling Facility: see ibm.com/eserver/zseries/psocftable.html
- Software Requirements
 - Coexistence Software (on other Target Systems)
 - Correct Levels of IBM Non-z/OS and Non-IBM Products
 - Functional Requirements
 - Only on z800 or z890: cannot IPL z/OS under an LPAR named ZOSExxxx
- z/OS Bimodal Migration Accommodation offering
 - **NOT available for z/OS 1.5 and later releases**

z/OS 1.6 Install Enhancements

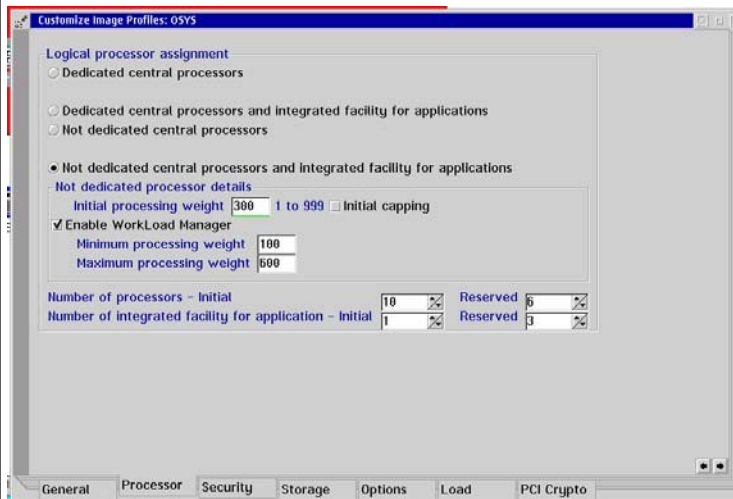
- ServerPac:
 - ServerPac can be delivered electronically!!!
 - Fewer jobs
 - Only mount RIM tape once
 - Automated retrieval of DASD volume information
 - Automatic block size optimization
 - Support for zFS file systems
- Documentation:
 - z/OS Migration contains all supported migration paths to z/OS 1.6
 - z/OS R3 to z/OS R6 (the whole book), z/OS R4 to z/OS R6, and z/OS R5 to z/OS R6
 - z/OS Migration can be used just for your particular path!
 - Only seeing actions associated with your migration

z/OS 1.6 – Scaling to new Heights

- Scale up
- Up to 24 processors in a single z/OS image
 - Total of 24 general purpose and zAAP processors
 - Prior to z/OS 1.6, max processors per z/OS image is 16
 - IBM plans to support 32 engines in a single LPAR in 2005
- Scale out
 - Up to 30 partitions in a single zSeries server
 - Support for 4 Logical Channel SubSystems (LCSS)
- Scale out with higher availability
 - Up to 32 images in a zSeries Parallel Sysplex®



z990 – 24 CPs in an LPAR - Image Profile



z990 allows up to 32 processors total. That is, the sum of Initial and Reserved processors of all types (e.g. CPs plus zAAPs) up to 32.

Up to 32 is valid even on an A08 or B16 because of concurrent book add support.

z/OS 1.6 supports up to 24 including reserve.

Dynamic LPAR Name (Reserved LPAR) z890 and z990


```

Session B - [24 x 80]
File Edit View Communication Actions Window ZipPrint Help
-----
Partition List
-----
C CBDPPAF0 Row 1 of 8
S Command ==> Scroll ==> PAGE
P Select one or more partitions, then press Enter. To add, use F11.
Processor ID . . . . : P20846A3
Configuration mode . : LPAR
Channel Subsystem ID : 2
-----
/ Partition Name Number Usage + Description
- CF21 A CF/OS Coupling facility
- TC4T21 1 OS Production A
* - TC4T22 2 OS Production B
- TC4T23 3 OS Development
- TC4T24 4 OS Test
- * 5 OS Reserved partition 1
- * 6 OS Reserved partition 2
- * 7 OS Reserved partition 3
***** Bottom of data *****
-----
F
b 20/006
Connected to remote server/host boehcd3 using port 2
    
```

- A reserved partition is defined with partition name placeholder ' * '.
- It has a MIF ID (partition number) and a usage type assigned. It may contain a description.
- Reserved partitions can not be assigned to access or candidate lists of channel paths or devices.
- z/OS 1.6 dynamic I/O can name a reserved LPAR or change a named LPAR to ' * '.

OSA-Integrated Console Controller (OSA-ICC)

- Console Controller for z990 and z890
 - System Console support for multiple LPARs
- Exclusive to one or both ports of an OSA-Express 1000BASE-T Ethernet
- Minimum software:
 - z/OS V1.3, z/OS.e 1.3 with APAR OA05738
 - HCD with PTF for APAR OA03689
- Supports Ethernet-attached TN3270E emulated sessions
- Can coexist in configurations using IBM 2074 and older 3174 control units



No coaxial cable support or Token-Ring support

IBM @server zSeries (z890 and z990) Application Assist Processor* (zAAP)



- A new zSeries Application Assist Processor (zAAP)
 - ▶ zAAPs are designed for Java code execution
 - ▶ z/OS JVMs assists with execution of code from standard processors to zAAPs
 - ▶ JVM executes Java code on zAAPs
 - ▶ Calls (JNIs) to DB subsystems (e.g., DB2) switch back to standard processors
 - ▶ DB returns (JNIs) to Java switch back to zAAPs
- Designed to provide a Single Tier integrated application and database serving environment
 - ▶ Advanced performance and QoS over typical 2-tier front-end application server TCP/IP connected back-end data server platforms
- When configured with general purpose engines in logical partitions running z/OS, zAAPs may help increase general purpose processors productivity

*For customers planning to run WebSphere Application Server V5 under z/OS on an IBM eServer zSeries 890 (z890) or 990 (z990) server, you may be eligible to take advantage of some of the benefits that the zSeries Application Assist Processor (zAAP) provides for a limited time on your z890 or z990 prior to migrating to z/OS 1.6. Contact your local IBM or authorized business partner hardware sales specialist for more information.

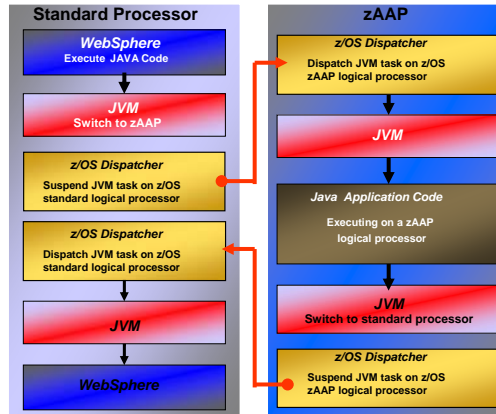
zAAP Characteristics

- Orderable by feature code (FC 6520 for z890 and FC 0520 for z990), up to one for each CP or unassigned CP (for z990) configured on the processor
- The zAAP only executes Java code
- zAAPs are designed so that users can manage the use of CPs such that Java code runs only on a CP, only on a zAAP, or on both
- Subsystems that exploit zAAPs
 - ▶ WAS 5.1
 - ▶ CICS/TS 2.3
 - ▶ DB2 V7 and V8
 - ▶ IMS V8, V9
 - ▶ WebSphere WBI for z/OS
- Other Software
 - ▶ z/OS 1.6*
 - ▶ SDK 1.4.1 with PTF UQ88783 (JVM)

* Prerequisites

zAAP Architecture and Workflow: Executing Java under the IBM JVM Control

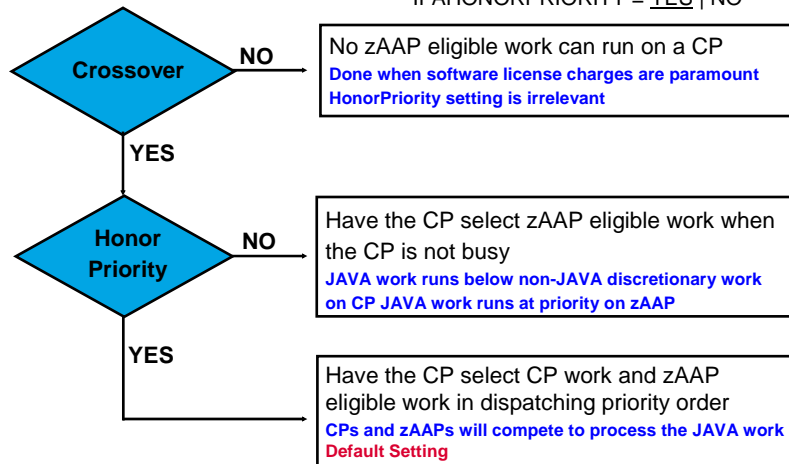
- IBM JVM, parts of LE runtime, and z/OS Supervisor needed to support JVM execution on zAAPs
- IBM JVM communicates with z/OS dispatcher when Java code is to be executed
- When Java is to be executed, the work unit is "eligible" to be dispatched on a zAAP
- z/OS dispatcher attempts to dispatch zAAP eligible work on a zAAP (when present)
- If there is insufficient zAAP capacity available, or standard processors are idle, the dispatcher may dispatch zAAP eligible work on a standard processor



SYS1.PARMLIB Options - IEAOPTxx

IFACROSSOVER = YES | NO

IFAHONORPRIORITY = YES | NO



Display Matrix Command

```

D M=CPU
IEE174I 11.10.01 DISPLAY M 608
PROCESSOR STATUS
ID  CPU                SERIAL
00  +                  023A6A2084
01  +                  023A6A2084
02  +A                 023A6A2084
03  +                  023A6A2084
04  +                  023A6A2084
CPC ND = 002084.B16.IBM.02.000000023A6A
CPC SI = 2084.315.IBM.02.000000000023A6A
CPC ID = 00
CPC NAME = OSYS
LP NAME = OOSP2      LP ID = 2
CSS ID = 0
MIF ID = 2

```

WSC Environment and Test

- z/OS 1.6 part of ESP
- One zAAP on B16 (2084-315)
- Benchmarks
 - Executed JSP that calculated value of Pi to a user defined decimal position
 - Executed CICS 2.3 JAVA transactions which called DB2 V7 data bases
 - Executed USS JAVA workload which called DB2
- Methodology
 - Executed workloads with SDK 1.3 on z/OS 1.5 to get baseline and obtain trace data needed for the estimation tool
 - Executed workload on z/OS 1.6 with varying values for IFACROSSOVER and IFAHONORPRIORITY
- Measurements:
 - Response Time, CPU per tran, Transaction count, various RMF APPL% values and zAAP busy

WSC Early Results

- Importance of switch rate on performance and throughput
 - CICS workload had an extremely high switching rate due to the method of calling DB2 and using the result set
 - JNI call backs to set variables in the JVM environment
 - Highlighted the need to review the switch rate when evaluating applications
 - Introduction of a new metric - Microseconds per switch
 - Amount of Java execution time per switch
 - Updated the JNI call back functions to identify additional switch points which were not productive
 - May reduce the need for microseconds per switch
 - High switch rates can influence the accuracy of the estimation tool
 - On each switch 2 TIMEUSED macros are used to get the CPU time
 - One of the macro invocations is included in the zAAP offloadable time

WSC Early Results

- Need to evaluate and update any performance management reports which report utilization
- Update capacity planning methods for zAAPs
- Update capture ratio calculations to include zAAP times
- Update IEFACTRT exit to add information on zAAP time – *If you Care !*
 - IEF374I message contains non-IFA and normalized IFA time

z/OS 1.6 RAS Items

- z/OS 1.6 includes several RAS (Reliability, Availability, Serviceability) enhancements:
- Activate an additional GTF trace specification when others have already been active
- The size of the SMF buffer is user controlled
 - Controls via PARMLIB specification - SMFPRMxx
 - Up to 1GB can be used to cache SMF records prior to writing to SMF data sets
- Execute Channel Program (EXCP) and EXCPVR VSCR
 - Many of the control blocks have moved above the line
 - Help reduce outages caused by storage shortages when large number of requests are active concurrently.
- New restartable PDSE address space
 - Help avoid unscheduled IPLs to recover from failure in PDSE processing

Reduction in IPL Time – OA07335

- IOS will parallelize operations to DASD during the MSI (Master Scheduler Initialization)
- 60-70% reduction in time during IOS Dynamic Pathing validation
- Early tests at one customer has the following results
 - 9500 DASD volume configuration
 - Time reduction from 17 minutes to just over one minute
- This support will be in base z/OS 1.6
 - APAR is closed, support rolled back to z/OS 1.4
 - UA12740 (z/OS V1R4), UA12741, (z/OS V1R5)
- You can use IPCS to view your IOS Dynamic Pathing Time by issuing **IPLDATA STATUS** command with default **SOURCE** set to **ACTIVE**

```

*** IEEMB860 Statistics ***
ILRTMRLG  00:00:00.546  ASM
IECVIOSI  00:00:09.847  IOS dynamic pathing
ATBINSYS  00:00:00.023  APPC
  
```

z/OS 1.6

- GRS SYNCHRES defaults to YES
 - If using previous default, be prepared for this change
- JOBCAT/STEP CAT are disabled by default
 - If need this support, you must issue F CATALOG command
 - F CATALOG,ENABLE(JOBSTEP CAT)
- DFSMSHsm now uses RACF for HSM command authority
 - This is not an option, you must define profiles for Hxxxxxx commands
- Language Environment LE RTLS support has been removed
 - If you use this support, there will be failures
- Temporary File System (TFS) supports larger sizes
 - TFS is an in-storage physical file system that supports in-storage mountable file systems
 - in z/OS 1.5 and later TFS is 64-bit resident by default
 - Don't forget to add MEMLIMIT keyword in SMFPRMxx
 - MEMLIMIT defaults to zero

z/OS 1.6

- SMF Type 30 records are updated to support zAAP
 - zAAP Statistics
- RMF support for zAAP
- RACF Templates loaded from a load module at initialization time
 - Both in z/OS 1.5 and z/OS 1.6
 - Templates from RACF data base are not used
 - An error message will be issued if you have not updated templates in the data base
- ICSF FMID HCR770A is shipped with z/OS 1.6
 - z990 and z890 enhancements to Cryptographic Support (FMID HCR770B) is available from the Web

z/OS 1.6

- Improved diagnostic data collection for FICON
 - Easier to identify the root causes of FICON errors
- Improved zFS address space availability
 - Diagnostic data collection will proceed in parallel with other processing
 - With z/OS 1.6, zFS parameters can be in multiple PARMLIB members
- z/OS Communication Server
 - Option for separate address spaces for TN3270 servers
 - Alternate route selection for SNA and Enterprise Extender (EE)
 - Dynamic major node definition changes
 - VTAM persistent sessions forced takeover
- RACF CDT (Class Descriptor Table) can be dynamically updated

z/OS 1.6 AUTOMOVE System List Enhancement

- Previously, all system names MUST be defined in the AUTOMOVE INCLUDE list in order to participate in the sysplex that are permitted to takeover the file systems, if the original system leaves
- Now a wildcard character * support is added
 - Systems need not be explicitly listed in the include system list
 - Easy to maintain a large number of systems in a sysplex
- Works on systems that are at z/OS 1.6 in the Sysplex
- BPXPRMxx updates
 - AUTOMOVE(INCLUDE,SYSA,*)
- Console command
 - SETOMVS FILESYS,FILESYSTEM='AHMAD.HFS,AUTOMOVE(I,SYSA,*)

z/OS 1.6 AUTOMOUNT Enhancement


- Previously, temporary change to add new AUTOMOUNT managed directories in an existing automount policy were not allowed
 - Maintained multiple copies of master and map files of the automount policy
- Now, with z/OS 1.6 Automount Enhancement
 - Add new automount managed directory to existing automount policy
 - Use the master and the map file from either PDS, PS or a File
 - This function can be used in a shared HFS environment at 1.6
 - /usr/sbin/automount -a
 - “-a” option appends new policy to existing policy
 - Master file from a member of a PDS
 - /usr/sbin/automount -a '//ahmad.jcl(master)'

HFS and zFS AUTOMOUNT

- Previously, a generic automount policy could not automount both HFS and zFS file systems
 - All file systems must be of the same type
- Now, AUTOMOUNT enhancement supports
 - A generic AUTOMOUNT policy that can manage both HFS and zFS
 - Dynamically determines the File System type and directs the mount to the appropriate Physical File System

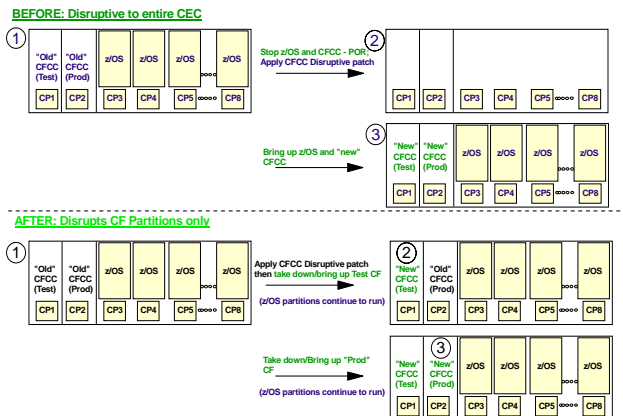
TN3270 Server – A Separate Address Space Option

- Prior to z/OS 1.6, TN3270 Server runs as a subtask of the IBM TCP/IP stack address space
- In z/OS 1.6, you are provided with options:
 - Run the TN3270 server as a separately started address space from TCP/IP
 - Continue to run TN3270 server as a subtask of the TCP/IP address space
- Benefits to running TN3270 as a separate address space
 - Allows for prioritization of TCP/IP address space vs. TN3270 server
 - Less likely for TN3270 server failure to cause a total TCP/IP failure
 - Allow for easier problem diagnosis for both TCP/IP and TN3270
 - Easier controls for starting and stopping the server



z990/z890 Enhanced CFCC Patch Apply

- ▶ CFCC Patch Apply Enhancement designed to help prevent a disruptive CFCC microcode load from impacting LPARs not related to the CFCC work
- ▶ Requires z990 at GA3, z890 at GA



CFCC Level Support

- CFCC Level 12
 - 64-bit addressing (large structure support, no 2GB control store limit)
 - Multitasking/parallelism enhancements (48 tasks)
 - System-Managed CF Structure Duplexing Support (note that G5/G6 CFs support duplexing via CFLEVEL 11)
 - Support for "batched" CF cache structure writes, castouts, and cross-invalidates (with DB2 Version 8 and z/OS 1.4 exploitation)
- CFCC Level 13
 - z990 and z890; DRIVER 55; GA 5/2004
 - Performance enhancement for castout processing for large DB2 GBP cache structures
 - Support rolledback for z900 and z800; DRIVER 3G; GA on 6/2004
 - NO ENHANCED PATCH APPLY Support for z900/z800

CFCC Level Support

- CFCC Level 14
 - Available on z890 GA2 and z990 GA4
- Functions and Potential benefits
 - Improvements in Coupling Facility Dispatcher
 - Serialization and latch management designed to better manage CF work
- Requirements and Support
 - z890/z990 hardware LIC support (GA October 29,2004)
 - z/OS 1.2 and z/OS.e 1.3 or haigher
 - APAR fix for OA08556
 - z/OS 1.4 and higher
 - APAR fix for OA09742 to allow connectors to request structure allocation in CFCC Level 14
 - CF Structure storage sizing
 - CFSizer

CF Structure REALLOCATE Command

- New SETXCF START,REALLOCATE command provides a simple, broad-based CF structure placement "optimization" capability:
 - Clears all CFRM "policy change pending" conditions
 - Moves all simplex structures into their "most preferred" CF location
 - Moves all duplexed structure instances into their two "most preferred" CF locations, in the correct order
 - Serial, one-structure-at-a-time processing to minimize any disruption caused by reallocation actions
 - Issues a message describing the evaluation process for each allocated structure
 - Issues a summary message upon completion of all structures, summarizing actions
 - Can simplify CF structure movement during disruptive CF maintenance scenarios
- Supported now on z/OS 1.4 and above via APAR OA03481
- Washington Systems Center Flash 10285
- OA03993 Serviceability APAR closed, PTFs available – *Lots of goodies*

Greater than 1023 XCF Members Per Group

- Constraint relief item in z/OS 1.6
 - SPE rollback to z/OS 1.4 and above also (OA04034)
- Allows larger numbers of CICS regions in the sysplex to join the single CICS XCF group
- This support allows the Sysplex CDS to be formatted to support up to 2047 XCF members per group
- Downlevel systems without this support will not be able to use a Sysplex CDS formatted for >1023 XCF members per group, they will not be able to join the sysplex
- Don't *overformat* the Sysplex CDS - performance considerations

ESP Summary

- z/OS 1.6 Install went smoothly overall – Good code quality
- More PDSEs in the system has implications of LPA cost
 - SIEALNKE PDSE need to be LNKLISTed
 - Combines several old libraries
 - SEZALOAD PDSE replaces SEZALINK
- Since z/OS 1.5, ServerPac now delivers ROOT HFS and Product HFSES
 - Mountpoints need to be added to BPXPRMxx
- SMPLTS elimination in SMP/E 3.2 (z/OS 1.5)
 - Removes the need for large LTS
 - Requires access to DLIBs
- We encountered problems with the installation WebSphere 5.1 on z/OS 1.6. The problems have since been fixed

ESP Summary . . .

- We do not have ISV products at the Washington Systems Center and did not test any
- Customers in the z/OS 1.6 ESP reported problems:
 - MIM/MIA Release 4.6 Service Pack 1 required for z/OS 1.6
 - EMC Symmetrix 8000 and Symmetrix DMX Storage Processors may experience device time-outs with OA07335
 - TopSecret
- Please contact your ISVs