


zSeries Parallel Sysplex Update September, 2004



Angelo F. Corridori
Parallel Sysplex Product Development Team
angeloc@us.ibm.com
www.ibm.com/servers/eserver/zseries/psa

Trademarks

The following are trademarks of International Business Machines Corporation.

ACF/VTAM	Enterprise System/4381	Open Blueprint
AD/Cycle	Enterprise System/9000	OpenEdition*
ADSM	Enterprise Systems Connection Architecture	OSA
Advanced Function Printing	ES/3090	OSA 1
AFP	ES/4381	OSA 2
AIX*	ES/9000	OS/2*
AIX/ESA	ESA/370	OS/390
AOEXPERT/MVS	ESA/390	OS/400*
Automated Operations Expert/MVS	ESCON	Parallel Sysplex
CICS/ESA	FASTService*	Power Prestige
DataHub	FlowMark	PR/SM
DATABASE 2	Hardware Configuration Definition	PS/2*
DataTrade	Hiperbatch	Processor Resource/Systems Manager
DB2*	Hipersorting*	RISC System/6000
DFDSM	Hiperspace	S/360
DFSMS	IBM*	S/370
DFSMS/MVS	IBM S/390 Parallel Enterprise Server	S/390
DFSMDfp	IBM S/390 Parallel Enterprise Server - Generation 3	SAAS
DFSMSdss	IMS/ESA	SAP R3
DFSMSshm	LANRES	Sysplex Timer
DFSMSrmm	Micro Channel*	System/370
Distributed Relational Database Architecture	MQ Series	System/390
DRDA	MVS/DFP	Systems Application Architecture*
Enterprise Systems Architecture/370	MVS/ESA	SystemView
Enterprise Systems Architecture/390	NetView*	VM/ESA
Enterprise System/3090	NQS/MVS	VSE/ESA
	OPC	VTAM
		3090

Note: 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.

Actual performance and environmental costs will vary depending on individual customer configurations and conditions.

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

Agenda

- Hardware Update
- Software Update
- GDPS Update
- Summary

Agenda

- Hardware Update
 - ▶ **zSeries External CFs**
 - ▶ **CF Duplexing and CF Configuration Options**
 - ▶ **Hardware Directions**
- Software Update
- GDPS Update
- Summary

zSeries Coupling Technology

- 64 Bit Architecture
- Dedicated or Shared CPs
- Up to 15 LPs
- Up to 32 GB of storage
- zSeries Peer Channels
 - InterSystem Channels-3 (ISC3)
 - Integrated Cluster Bus-3 (ICB3)
 - Integrated Cluster Bus-4 (ICB4)
- zSeries Compatibility Channels
 - zSeries to 9672/9674
 - InterSystem Channels (ISC)
 - Integrated Cluster Bus (ICB)
- Dynamic CF Dispatch
- Dynamic ICF Expansion

z990 Model 300



- Up to 32 ICFs
- Up to 16 ICB-3
- Up to 32 ISC-3
- Up to 16 ICB-4
- Up to 64 GB/book
- Upgrade to z990
- Cannot upgrade directly from z100

z900 Model 100




- Up to 9 ICFs
- Up to 16 ICB-3
- Up to 32 ISC-3
- Upgrade to z900
- Upgrade from R06 z890...

z800 Model OCF



- Up to 4 ICFs
- Up to 6 ICB-3
- Up to 24 ISC-3
- Upgrade to z800
- Up to 4 ICFs
- Up to 16/8 ICB-3/4
- Up to 48 ISC-3
- Upgrade to z990


zSeries Peer Mode Links



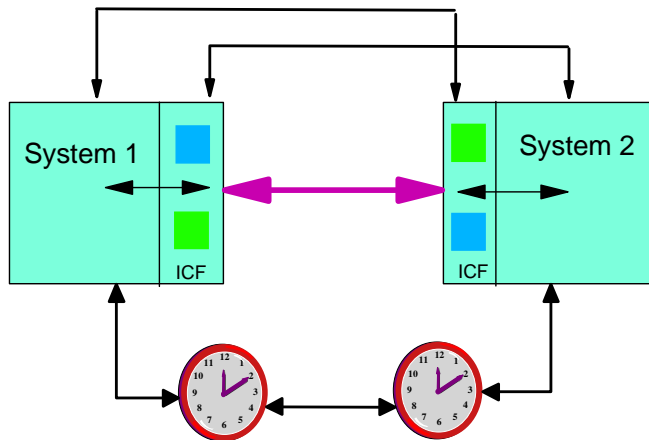
- Reduced number of links
- Dynamic reconfiguration to peer mode
 - ISC, ICB, IC
 - sender or receiver

"Traditional" CF Configuration Recommendations

Data Sharing ...



System Managed CF Structure Duplexing



Function


- - Automatic Rebuild for planned reconfiguration
- - Automatic switchover for unplanned outages
- - Automatic duplexing re-establishment
- - Overlapped requests for high performance

Value

- - Consistent Recovery Mechanism
 - ▶ Reduced complexity
- - Faster than structure rebuild
- - Enables a robust "all-ICF" configuration

System-Managed CF Duplexing Exploiters

▶ CICS®	<u>Shared TS, CF data tables, named counter</u>
▶ Comm Server	GR, MNPS (VTAM®), <i>SWSA, Sysplex Ports</i>
▶ DB2®	SCA
▶ DFSMS	RLS lock (VSAM), Common Recall Queue
▶ IMS™	CQS, EMH, VSO structures
▶ IRLM	Lock (DB2 and IMS)
▶ JES2	Checkpoint
▶ MQ	Shared queues
▶ WLM	Shared enclaves, LPAR clusters (IRD)
▶ BatchPipes®	
▶ System Logger	



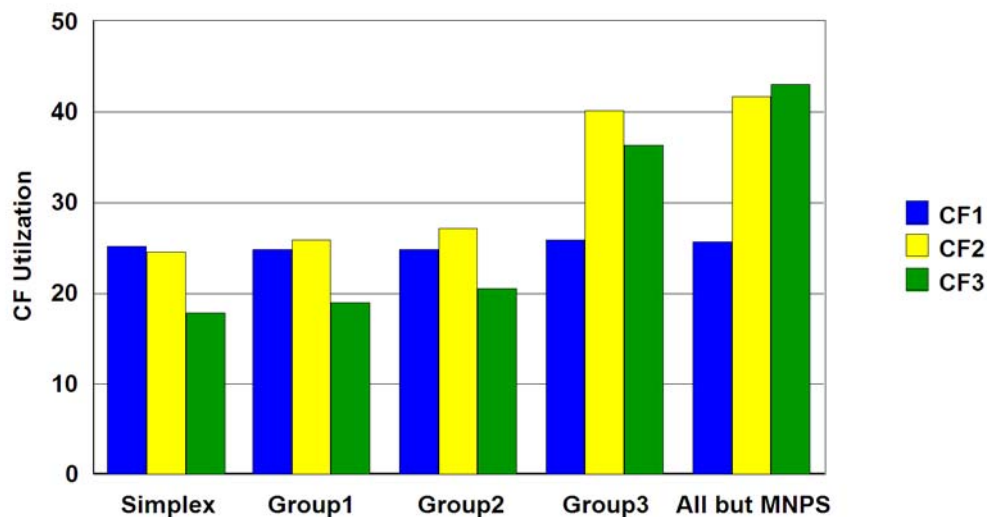
CF Duplexing - Key Requirements

- Hardware
 - ▶ zSeries or G5/G6 Processor
 - ▶ Bidirectional CF to CF connectivity
- Software
 - ▶ z/OS 1.2 or above
 - ▶ Appropriate exploiter release and service
 - ▶ CFDUPLEXING PSP bucket
- CFCC
 - ▶ Level 11 for G5/G6
 - ▶ Level 12 or above for zSeries

CF Duplexing "Readiness Review"

- Readiness Review ran from March, 2003 thru 2/7/2004
 - ▶ Existing readiness review customers supported through September
- CF Duplexing generally available since February
- IBM has assisted about 60 customers with their CF Duplexing implementation
 - ▶ At least 10 customers are using CF Duplexing in production
 - ▶ About 50+ customers continue to test and prepare for CF Duplexing production
 - ▶ All significant problems found in the Readiness Review have been fixed
- Overall Quality
 - ▶ Very stable on existing HW (z900, G5/G6)
 - ▶ Some problems triggered by new HW rollout (z990)
 - ▶ one specific to z990
 - ▶ Customer experience presentations given at Share
- Overall Performance
 - ▶ As expected and described in the CF Duplexing documentation

CF Duplexing Staged Implementation (Internal Test)

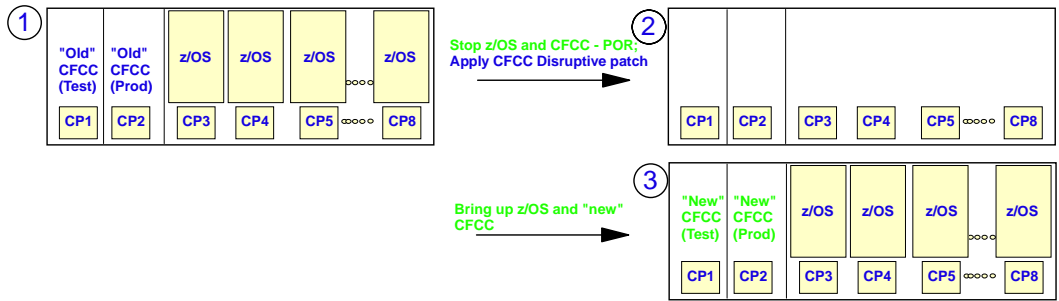


Advantages of a zSeries CF for CF Duplexing

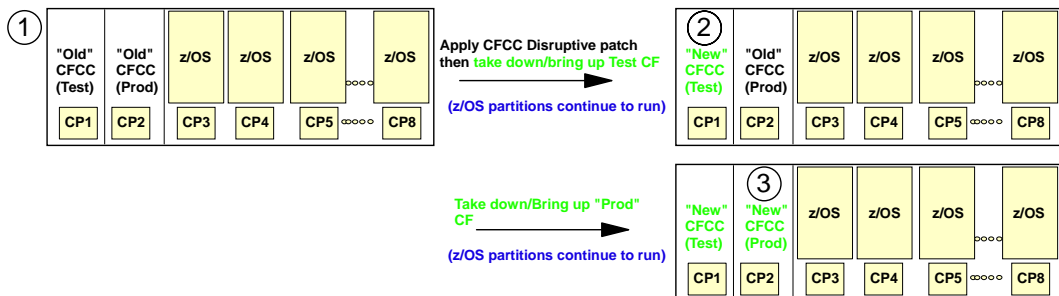
- 48 tasks (up from 16)
 - ▶ Increased CF command parallelism
- 7 subchannels (up from 2) per peer link
 - ▶ Increased CF command parallelism
- Reduced connectivity requirements using peer links
 - ▶ z/OS to CF as well as CF to CF
- Path busy redrive optimization on zSeries
 - ▶ Improved path busy management
- Faster processor speed

z890 / z990 Enhanced CFCC Patch Apply

BEFORE: Disruptive to entire CEC




AFTER: Disrupts CF Partitions only



"New" CF Configuration Recommendations

Data Sharing ...



Which CF Configuration is Right for Me?

- **It Depends !**
 - ▶ On the factors that are most important to your business
 - cost
 - availability
 - system management
 - ▶ Much less on the technical factors associated with your Parallel Sysplex implementation
 - link technology
 - max. size of sysplex
 - etc.
 - ▶ See next page for summary.....

External vs. Internal CF Summary Comparison

	External	Internal
General Characteristics		
Acquisition Cost	Good	Better
CF Upgrade/Migration	Independent of server upgrade	Automatically with server upgrade
CFCC microcode	same	same
Processor HW	same	same
Affects SW License charges	No	No (unless running on a general purpose CP)
Links		
ISC	Yes	Yes
ICB	Yes	Yes
IC	Yes (if partitioned)	Yes
Max.Connectivity(Differences only)		
9672 G5 ICB-2	24	6-18 (model dependent)
z900 ISC-3	42 (RPQ)	32
z800 ICB-3	6	5
Systems Management		
Server Upgrade (e.g. new driver)	No Effect	CF out of service during upgrade
CF Upgrade (e.g. disruptive CFCC patch)	CF out of service during upgrade	CF and z/OS images on the server are out of service during upgrade (z990/890 avoids z/OS outage)
Capacity Planning (storage)	CF owns all storage	Allocated among z/OS(s) and CF(s)
Can be partitioned	Yes	Yes
Availability	All External	All Internal
Resource Sharing	Yes	Yes, with careful structure placement
Data Sharing	Yes	Yes, with use of SM CF Duplexing

External vs. Internal CF Summary Comparison

	External	Internal
General Characteristics		
Acquisition Cost	Good	Better
CF Upgrade/Migration	Independent of server upgrade	Automatically with server upgrade
CFCC microcode	same	same
Processor HW	same	same
Affects SW License charges	No	No (unless running on a general purpose CP)
Links		
ISC	Yes	Yes
ICB	Yes	Yes
IC	Yes (if partitioned)	Yes
Max.Connectivity(Differences only)		
9672 G5 ICB-2	24	6-18 (model dependent)
z900 ISC-3	42 (RPQ)	32
z800 ICB-3	6	5
Systems Management		
Server Upgrade (e.g. new driver)	No Effect	CF out of service during upgrade
CF Upgrade (e.g. disruptive CFCC patch)	CF out of service during upgrade	CF and z/OS images on the server are out of service during upgrade (z990/890 avoids z/OS outage)
Capacity Planning (storage)	CF owns all storage	Allocated among z/OS(s) and CF(s)
Can be partitioned	Yes	Yes
Availability	All External	All Internal
Resource Sharing	Yes	Yes, with careful structure placement
Data Sharing	Yes	Yes, with use of SM CF Duplexing

z990 GA3/z890 GA1 Coupling related enhancements

- 48 ISC-3 links maximum on CF (up from 32 Standard)
- CFCC Patch Apply Enhancement for Disruptive CFCC patches
 - ▶ New type of patch for CFCC which is "concurrent" in that it does not require a CEC POR, and is disruptive only to the CF partition on which it is being activated
 - ▶ NO more CEC PORs to install CFCC maintenance - either fully concurrent, or deactivate/reactivate CF partition only!
- Read Castout Class cache command performance enhancements
 - ▶ Performance benefit for very large DB2 GBP cache structures
- CFCC level 13 (also available on z800/900)


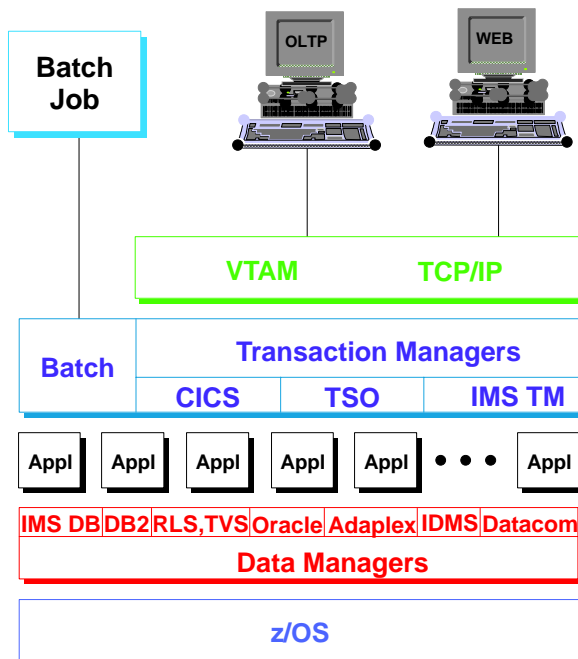
z990 GA4 / z890 GA2 Coupling Related Enhancements

- CFCC Level 14
 - ▶ CFCC Dispatcher Restructured
 - Provide a "true ready queue" for CF tasks
 - Maintain all work queues in CSN order for FIFO command execution
 - Favor completion of work in progress over starting new requests
 - Latch manager fairness (latches granted in CSN order)
 - ▶ Expected benefits
 - most significant benefits expected for targeted environments (e.g. large number of CF CPs)
 - most "typical" Parallel Sysplex configurations will not see much difference

Agenda

- Hardware Update
- Software Update
 - ▶ **Subsystems**
 - ▶ Base Control Program
- GDPS Update
- Summary

Parallel Sysplex Software Structure



Single System Image & High Availability Connection

Dynamic Workload Balancing

Applications Unchanged

Data Sharing

Base Services
Hardware Interfaces

Communications Server Support

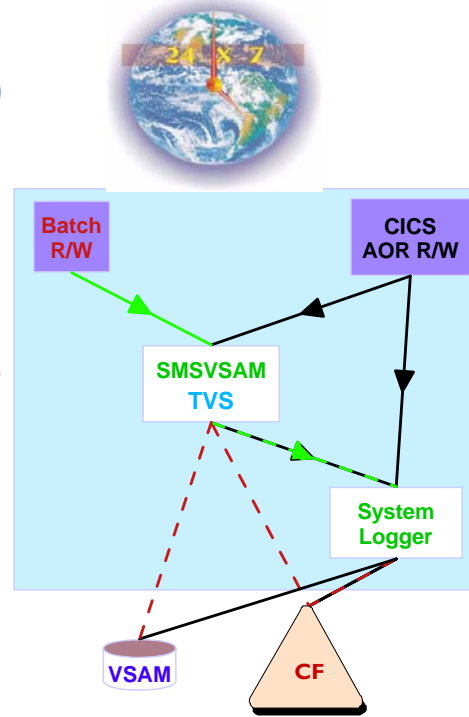
- z/OS 1.4
 - ▶ Sysplex wide security associations
 - ▶ Sysplex wide source VIPA
- z/OS 1.5
 - ▶ Increased VIPADISTRIBUTE ports (4 -> 64)
 - ▶ Increased maximum DVIPAS per stack (256 -> 1024)
 - ▶ Sysplex Distributor enhancements
 - timer based "stickyness" for IP address, DVIPA, and ports
 - round robin distribution
 - ▶ Enhanced DVIPA activation in a sysplex

Communications Server Support (cont.)

- z/OS 1.6
 - ▶ Extended configuration control for outbound connections
 - outbound connections use same IP address as inbound
 - job specific source IP addressing
 - ▶ Self monitoring for improved availability of TCP/IP stacks
 - ▶ Can use sysplex distributor without z/OS stack being a general intermediate routing node
 - ▶ All sysplex functions support IPv6
 - ▶ z/OS Load Balancing Advisor (LBA) - 4Q04
 - for outboard load balancers
 - supports all IP based application workloads
 - sysplex scope
 - z/OS 1.4 and above
- Directions
 - ▶ Improved automomics to reduce/eliminate operator actions
 - ▶ Improved usability for DVIPAs
 - ▶ Improved intra-sysplex flexibility for sysplex distributor and DVIPA
 - ▶ Improved sysplex distributor load balancing decisions and end user response time

Transactional VSAM (DFSMStvs)

- Allows batch update sharing concurrent with CICS on-line access to recoverable data
 - ▶ Addresses the batch window for CICS and batch sharing of VSAM recoverable data sets
 - ▶ Allows multiple batch update programs to run concurrently with CICS against the same files
- "Transactionalizes" batch access to VSAM data sets
 - ▶ Provides two-phase commit capability
 - ▶ VSAM becomes a data base like IMS DL/1 and DB2
- Enables 24 x 7 availability
- Optional priced feature on z/OS 1.4 or later
 - ▶ GA 5/2003
- Exploitation
 - ▶ May require changes to applications to fully exploit



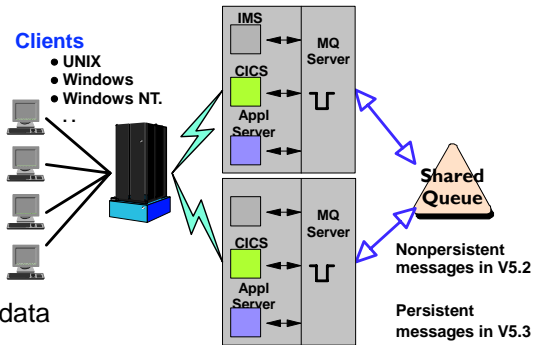
TLLBPARz370

Key Components for Transactional Support

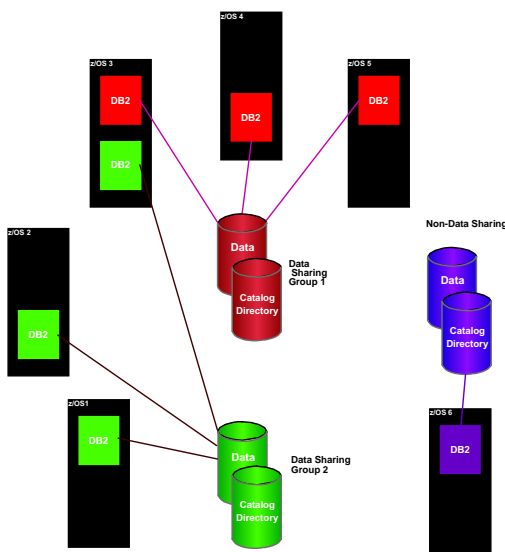
- **Application program**
 - ▶ User defines the data set as Recoverable unless already defined by CICS
 - ▶ Uses RRS Callable Services macros to commit (SRRCMIT) & backout (SRRBACK)
 - ▶ Uses VSAM MACROS (Open, Close, Get, Put, etc.)
- **TVS**
 - ▶ Interacts with MVS Logger (forward recovery & backout logs)
 - ▶ Built on top of and uses VSAM RLS for locking & caching in CF and for I/O
 - ▶ Participates with RRS on commit/rollback
- **MVS Recoverable Resource Services (RRS)**
 - ▶ Provides 2-phase commit, backout
- **CICSVR** required for forward recovery using forward recovery logs

Websphere MQ for z/OS

- Availability
 - ▶ Workload Balancing
 - ▶ Planned maintenance easier
- Administration
 - ▶ Simple, scalable administration
 - ▶ Single name space to describe resources
 - ▶ Fewer resources to define
 - Shared queue definitions and control data
 - ▶ Single system to control and administer
- ARM Support
 - ▶ Restart failed queue manager or mover in place
 - ▶ Move MQ and related subsystems to an available LPAR
- System-Managed CF Structure Duplexing Support
- Requirements for Shared Queue
 - ▶ CFCC Level 9, OS/390 V2.9 with (RRS) activated, and DB2 V5 datasharing



DB2 Data Sharing Directions



- V7
 - ▶ Restart "Light"
 - ▶ N, N-2 Migration/Fallback Support
 - ▶ Structure Rebuild Enhancements
 - ▶ Group attach enhancements
 - ▶ Persistent CF structure sizes
 - ▶ System Managed Duplexing of Lock and SCA Structures
- V8
 - ▶ Batched GBP writes and cast-outs
 - improved performance especially for batch updates
 - ▶ Reduced global contention for tablespace L-locks
 - improved performance especially for OLTP
 - ▶ More efficient index split processing
 - ▶ Less Disruptive LPL Recovery
 - ▶ Restart "Light" Recovery Enhancements
 - ▶ Immediate Write PH1 becomes system default
 - ▶ Enhanced DISPLAY GBPOOL output
 - ▶ DPSIs for more efficient member affinity routing
 - ▶ DDF sub-groups

Parallel Sysplex Training Environment



- IBM Learning Services
- Obtain, maintain, and improve Parallel Sysplex skills
 - ▶ Operators and System Programmers
 - ▶ Practice Sysplex activities
 - Manipulating structures
 - Recovering from errors
 - ▶ Helps customers install Parallel Sysplex in minimal time
- For More Information:
 - ▶ www.ibm.com/services/learning/spotlight/zseries/parallelsysplex.pdf

Agenda

- Hardware Update
- Software Update
 - ▶ Subsystems
 - ▶ **Base Control Program**
- GDPS Update
- Summary

Console Requirements

- 1. Improve Message delivery
 - ▶ Storage and XCF utilization
 - ▶ Varying speeds and feeds
- 2. Improve Console State Information
 - ▶ Data replicated on every system in the sysplex
 - ▶ State changes, system join/leave cause flurries of activity
- 3. Eliminate Limit of 99 MCS/SMCS/Subsystem consoles in a sysplex
 - ▶ Installation constraint

Console Enhancements Feature of z/OS 1.4

- Targets WTO buffer shortage problems
 - ▶ 40-45% of console related multi-system outages
 - ▶ Mixed levels of z/OS supported
 - compatibility PTFs required
 - ▶ Can re-IPL a single image to fall out of stage 1
- Improved Availability in sysplex and non-sysplex environments
 - ▶ Robust management of WTO message storage improves system availability
 - ▶ Multiple-line messages not queued until completed
 - ▶ Emphasis placed on maintaining overall health of the system and sysplex in stressed environments
- More efficient message broadcast and deletion
 - ▶ Use of XCF ordered delivery to help ensure that messages and message deletion requests arrive in the order that they were sent
 - ▶ Simplified message routing decision making
 - Decision made at target, not source system




***This is an optional feature for z/OS 1.4
This function is part of the base in z/OS 1.5***

Console Enhancements Staging

- Stage 1a - solves problem 1
 - ▶ Delivered in z/OS 1.4.2 feature and z/OS 1.5 base
- Stage 1b - improves RAS
 - ▶ future z/OS release
- Stage 2 - solves problems 2 & 3
 - ▶ future z/OS release

zFS, z/OS V1R6 and earlier Sysplex



- zFS and HFS are
 - not "sysplex aware" for R/W file systems
 - USS determines which system is to own a R/W file system and forwards all requests to that system
 - If a system fails, USS moves R/W file systems owned by that system to another system
 - "sysplex aware" for RO file systems



CFRM Performance Enhancement - z/OS 1.4

- Improved CFRM performance during recovery operations will improve Parallel Sysplex availability
- Recovery actions (CEC failure, CF failure, Structure rebuild) may result in large records requiring large I/O bandwidth
 - ▶ originally: 64 structures with 32 connectors
 - ▶ currently: 512 structures with 255 connectors
 - >60X increase
- up to 30% CFRM CDS I/O performance improvement
- reduces possibility of I/O bottleneck on CFRM CDS
- Improved Parallel Sysplex availability

CF Structure REALLOCATE Command

- Improved ease of use for all Parallel Sysplex customers
- Most beneficial to those customers who
 - ▶ want specific structures to reside in specific CFs
 - e.g. load balancing
 - ▶ have CFs with differing characteristics
 - e.g. capacity, storage, speed
 - ▶ use duplexed structures (user or system managed)
 - ▶ have more than 2 CFs

The XCF REALLOCATE Command - 1/4





- Operator needs to clear out CF1 for maintenance
- Operator issues setxcf stop,duplex,keep=new for structures E' B'
- Operator issues setxcf rebuild,cfname=CF1,loc=other

Active CFRM Policy

```



STRUCTURE NAME(C) PRELIST(CF1,CF2)
STRUCTURE NAME(E) PRELIST(CF1,CF2) DUPLEX(ALLOWED)
STRUCTURE NAME(B) PRELIST(CF2,CF1) DUPLEX(ALLOWED)
STRUCTURE NAME(A) PRELIST(CF1,CF2)
STRUCTURE NAME(G) PRELIST(CF2,CF1)
STRUCTURE NAME(D) PRELIST(CF1,CF2)
    ○
    ○
    
```

The XCF REALLOCATE Command - 2/4



Operator configures CF offline and given to CE

The XCF REALLOCATE Command - 3/4





- Operator needs to repopulate CF1 according to PREFLIST
- Operator issues setxcf REALLOCATE
 - All structures get relocated according to PREFLIST

Active CFRM Policy

```

STRUCTURE NAME(C) PREFLIST(CF1,CF2)
STRUCTURE NAME(E) PREFLIST(CF1,CF2) DUPLEX(ALLOWED)
STRUCTURE NAME(B) PREFLIST(CF2,CF1) DUPLEX(ALLOWED)
STRUCTURE NAME(A) PREFLIST(CF1,CF2)
STRUCTURE NAME(G) PREFLIST(CF2,CF1)
STRUCTURE NAME(D) PREFLIST(CF1,CF2)
o
o
    
```

The XCF REALLOCATE Command - 4/4



- Operator establishes duplexing
 - setxcf start,duplex,strname=b
 - setxcf start duplex,strname=e
 - Etc.

Active CFRM Policy

```

STRUCTURE NAME(C) PREFLIST(CF1,CF2)
STRUCTURE NAME(E) PREFLIST(CF1,CF2) DUPLEX(ALLOWED)
STRUCTURE NAME(B) PREFLIST(CF2,CF1) DUPLEX(ALLOWED)
STRUCTURE NAME(A) PREFLIST(CF1,CF2)
STRUCTURE NAME(G) PREFLIST(CF2,CF1)
STRUCTURE NAME(D) PREFLIST(CF1,CF2)
o
o
    
```

XCF/XES Directions

- >1023 XCF Members per Group
- Healthchecker incorporated into z/OS and restructured
- XES Lock Constraint Relief
- CFRM Performance Enhancements
- XCF Buffer Exhaustion

Agenda


- Hardware Update
- Software Update
- GDPS Update
 - ▶ **Recent Enhancements**
 - **Hyperswap**
 - **Extended Distance**
 - ▶ **Directions**
- Summary

zSeries Continuous Availability

Single System




Parallel Sysplex



1 to 32 Systems

GDPS





Site 1

Site 2


- **Built In Redundancy**
- **Capacity Upgrade on Demand**
- **Capacity Backup**
- **Hot Pluggable I/O**
- **Addresses Planned/Unplanned HW/SW Outages**
- **Flexible, Nondisruptive Growth**
 - ▶ Capacity beyond largest CEC
 - ▶ Scales better than SMPs
- **Dynamic Workload/Resource Management**
- **Addresses Site Failure/Maintenance**
- **Sync/Async Data Mirroring**
 - ▶ Eliminates Tape/Disk SPOF
 - ▶ No/Some Data Loss
- **Application Independent**

GDPS Solution

- **Continuous Availability and Disaster Recovery solution**
 - ▶ GDPS/PPRC (Peer to Peer Remote Copy (PPRC) - Synchronous)
 - ▶ Multisite Sysplex (fiber distance between sites 100 km - max)
 - ▶ No or limited data loss in unplanned failover - user policy
 - ▶ Planned and Unplanned reconfiguration support
- **Disaster Recovery solution**
 - ▶ GDPS/XRC (eXtended Remote Copy (XRC) - Asynchronous)
 - ▶ Supports any distance
 - ▶ Production systems in Site 1
 - ▶ Limited data loss to be expected in unplanned failover
 - ▶ GDPS initiates restart of production systems in Site 2
- **Common functions (GDPS/PPRC and GDPS/XRC)**
 - ▶ GDPS solution extended to tape resident data
 - Peer-to-Peer Virtual Tape Server support (PtP VTS)
 - ▶ Point-in-time copy created (FlashCopy™)
 - Maintains D/R readiness during resynchronization
 - Perform D/R testing while maintaining D/R readiness
 - ▶ Management of zSeries Operating Systems (Linux® for zSeries, z/VM™, VSE/ESA™)




PtP VTS



GDPS/PPRC Hyperswap Function

Brings different technologies together to provide a comprehensive application and data availability solution




- GDPS/PPRC HyperSwap substitutes PPRC secondary for PPRC primary devices
 - No operator interaction - GDPS-managed
 - Can swap large number of devices - fast
 - Changes status in secondary disk subsystem
 - Changes status in primary disk subsystem - if available
 - Nondisruptive - applications keep using same device addresses
- GDPS/PPRC with HyperSwap helps reduce RTO for
 - Planned disk and site reconfiguration
 - Unplanned disk and site reconfiguration

PPRC Peer-to-Peer Remote Copy
UCB Unit Control Block

GDPS/PPRC distance extended up to 100km

- What has changed?
 - Timer links from Sysplex Timer® to CPC extended from 50 km to up to 100 km
 - ISC3 (Coupling Facility) links (Peer mode) extended from 50 km to up to 100 km
 - Mid-span optical amplification may be required (DWDM and fiber dependent)
- What has not changed?
 - Links between Sysplex Timers CANNOT exceed 40 km
 - Latency of signals over fiber (10 us/km)
 - Impact on application performance MUST be evaluated for each customer
- Where can I locate the Sysplex Timers?
 - Intermediate site along fiber route (less than 40 km) from either data center
 - Different building on same campus as main data center
 - Different floor of same building as main data center
- What do I order?
 - RPQ 8P2262 (IBM S/390®)
 - RPQ 8P2263 (IBM eServer™ zSeries®)

GDPS/PPRC 100km example



Redundant ISC3 cross site links from Site 1 to Site 2 also extended to up to 100 km max

Geographically Dispersed Parallel Sysplex

Continuous Availability and Disaster Recovery

Directions

- ◆ Increase GDPS functionality
- ◆ Enhance GDPS recovery capabilities
- ◆ Increase GDPS/PPRC distance capability to 100 Km
- ◆ Extend GDPS solution into the open server/storage environment

Strategy: GDPS is the industry leading Continuous Availability / Disaster Recovery heterogeneous e-business infrastructure platform that will allow customers meet their business objectives

Agenda

- Hardware Update
- Software Update
- GDPS Update
- **Summary**

Parallel Syplex Technology Directions

Technology Element	Current Implementation	Technology Directions
Servers/CFs	z800, z900, z890, z990	Future server support/enhancements
CFCC Microcode Function	CFCC level 14	Performance, Serviceability, Hardware support
Connectivity/Links	ISC-3, ICB-4 peer mode	Link speed/function enhancements
z/OS	z/OS 1.4, 1.5, 1.6	Improved availability, Constraint relief, Enhanced sharing, Complexity reduction
Subsystems	CICS, DB2, IMS, DF/SMS, MQ, Websphere, Comm. Server	Improved availability, Enhanced/new sharing capability, Improved workload balancing
GDPS	Disk/Tape Mirroring, Automation	Enhanced functionality, Increased distance (PPRC), Reduced Recovery Time, Include open server/storage

Parallel Sysplex Value

