



Washington Systems Center

Planning for and Using System Managed CF Structure Duplexing

Riaz Ahmad
IBM Washington Systems Center
Gaithersburg, Maryland

© 2003 IBM Corporation

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*	Magstar*	VSE/ESA
Enterprise Storage Systems	MVS	VTAM*
ESCON*	Netfinity*	WebSphere
FICON	OS/390*	z/Architecture
GDS	Parallel Sysplex*	z/OS
Geographically Dispersed Parallel Sysplex	PR/SM	z/VM
HyperSockets	S/390*	zSeries
IBM*	S/390 Parallel Enterprise Server	

IBM eServer zSeries

* 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.

Topics

- What is CF Duplexing ?
- Hardware requirements
- Software requirements
- Migration
- Summary

Table of contents

- **Overview of Duplexing**
- Hardware Requirements
- Software Requirements
- Migration
- Summary

What is SM CF Duplexing?

- System allocates a secondary instance of a Coupling Facility structure
 - ▶ Replicates the structure contents in the secondary structure to establish duplexing
- System transparently replicates Coupling Facility operations
 - ▶ Update structure contents in both structure instances in a synchronized manner to maintain data consistency between the primary and secondary structure contents
- In the event of a failure affecting one of the instances
 - ▶ System failovers to simplex mode using only the unaffected structure instance with no loss of data

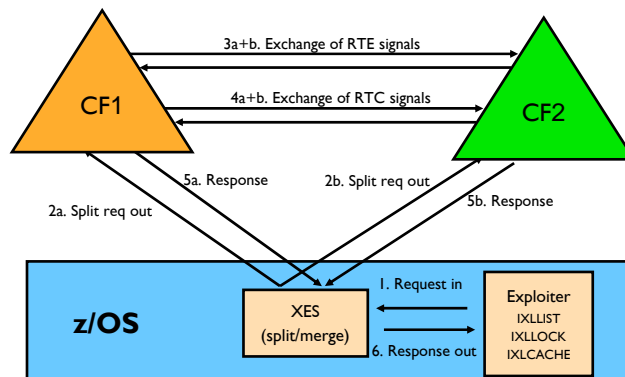
Why Duplex?

- To provide near-continuous availability for CF structure data in case of **hard failure** through rapid and transparent duplexing failover between the primary and secondary structures
 - ▶ CF or structure failure
 - ▶ Loss of CF Connectivity
- To enable configuration simplification and overall cost reduction for Parallel Sysplex®
 - ▶ Enabling the use of non-standalone CFs (e.g. ICFs) for all resource sharing and data sharing environments

Duplexing Candidates

- Structures whose connectors:
 - ▶ Do not support rebuild at all
 - ▶ Support rebuild but
 - Either cannot rebuild the structure contents when the structure and one or more connectors are simultaneously lost as the result of lack of failure-isolation
 - Or reconstruction of the data in failure scenarios is time-consuming or error-prone
 - ▶ Do not support User-Managed Duplexing

CF Duplexing Protocol Flow



Mainline Coupling Facility Requests

- All mainline CF requests are transparently processed appropriately as simplex or duplex
- z/OS internally handles
 - Request splitting and conversion
 - Protocols for driving duplexed requests to the CFs
 - Request response reconciliation
- CFCC internally handles:
 - CF-to-CF signalling to coordinate structure object updates

CFRM Policy

- Determines if and when a structure is to be duplexed
 - DUPLEX(DISABLED) - Never duplexed
 - DUPLEX(ALLOWED) - Manual
 - DUPLEX(ENABLED) - Automatic

Stopping Structure Duplexing

- SETXCF operator command
 - ▶ SETXCF STOP,RB,DUPLEX,STRNAME=xxx,KEEP=OLD
 - ▶ SETXCF STOP,RB,DUPLEX,STRNAME=yyy,KEEP=NEW
- IXLREBLD
- Automatically:
 - ▶ When CFRM policy change disallows duplexing
 - ▶ When duplex command reconciliation detects failure
 - ▶ When failure affects one of the structure instances

Remember !

- Duplexing is **NOT** a "no brainer" solution that can be turned on for all structures without any planning
- Through the CFRM policy, the installation has **full control** over whether to duplex any particular structure
 - ▶ Test/stage duplex implementation
 - ▶ Exploitation availability
 - ▶ Performance cost/benefit analysis for a given structure, given actual workload characteristics

Table of Contents

Overview of Duplexing

••• **Hardware Requirements**

Software Requirements

Migration

Summary

Hardware Requirements


- IBM @server zSeries™ or G5/G6 Servers
- Two or more Coupling Facilities
 - ▶ IBM @server zSeries z990 or z890
 - CFCC Level 12 Driver 52
 - ▶ IBM @server zSeries z800 or z900
 - CFCC Level 12 Driver 3G
 - ▶ IBM 9672 G5/G6
 - CFCC Level 11 Driver 26
 - ▶ Latest microcode maintenance level
 - ▶ Standalone CF must be R06, z900-100, z800-0CF, z890, z990
- CF-to-CF links

CF-to-CF Links and Capacity

Bidirectional communication required for duplexed command execution

- ▶ Provided via standard Coupling Facility links of any type
 - ▶ Peer-mode links – zSeries
 - ▶ CF support for sender channels
 - Requires APAR OW45976 for HCD
- For availability and capacity:
 - ▶ Configure redundant CF-to-CF links
- Loss of CF-to-CF connectivity during signal exchanges can cause structure to fall out of duplexing
- It may be necessary to provide additional CPU, link
 - ▶ Duplex operations consume more resources than simplex operations
 - ▶ May be offset by existing excess capacity
 - ▶ Depends on which structures are duplexed

Hardware Configuration



Provides complete failure-isolation from all structure connectors

Hardware Configuration

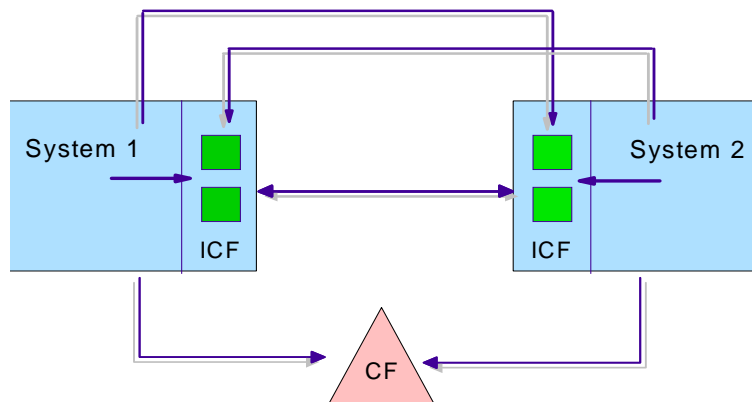


Table of Contents

Overview of Duplexing

Hardware Requirements

➤ Software Requirements

Migration

Summary

Software Requirements

- Operating System
 - ▶ z/OS 1.2 or higher
 - ▶ APAR OW41617
- CFRM Couple Data Set
 - ▶ New format
 - ▶ APAR OW41617 enables format utility to process SMDUPLEX keyword
- CFRM Policy Updates

Software Requirements

- Exploiting subsystems
 - ▶ May require installation of new software levels or PTFs to enable duplexing support
 - ▶ Document their migration procedures
 - ▶ All exploiter instances must be at a duplexing capable level before relevant structure is duplexed

Exploiters

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

CFRM Policy Updates

▪ Update structure statements

- ▶ DUPLEX (ALLOWED|ENABLED)
- ▶ Preference List must contain at least two duplex capable Coupling Facilities
 - Allocation algorithm changed
 - Most preferred CF may not be

▶ SIZE/INITSIZE

```

STRUCTURE NAME(DSNZPLEX_LOCK1)
          SIZE(150000)
          INITSIZE(132000)
          PREFLIST(CF1,CF2,CF103)
          REBUILDPERCENT(1)
          DUPLEX( ALLOWED )
  
```

Table of Contents

Overview of Duplexing

Hardware Requirements

Software Requirements

➤ Migration

Summary

Migrate to Duplexing

- Determine which structures are to be duplexed
- Upgrade capacity as needed to support duplex structures
 - ▶ Coupling Facility storage
 - ▶ Coupling Facility CPs
 - ▶ Links
- Get Coupling Facility to appropriate CFLEVEL
- Configure CF-to-CF links
- Configure System-to-CF links to avoid path-busy

Migrate to CF Duplexing

- Migrate systems to level of z/OS that supports duplexing
 - Check the CFDUPLEXING PSP Bucket
- Migrate exploiter software to level supporting system-managed CF Duplexing
- Migrate CFRM CDS to level supporting SMDUPLEX
- Define and activate CFRM policy with DUPLEX parameter for desired structures
- Document operational procedures for failure scenarios
- Monitor and Tune

Migration caveats

- CFRM CDS must be formatted with SMDUPLEX record
- Only duplexing-capable z/OS systems can use a CFRM CDS formatted with SMDUPLEX
- Downlevel systems cannot join such a sysplex and use CFRM once the CDS with SMDUPLEX is in use
- Fallback to CDS without SMDUPLEX
 - Sysplex-wide IPL
- You can turn duplexing on/off via policy change
 - **Does NOT require sysplex-wide IPL**
- New connections that do not support duplexing will be unable to connect as long as the structure remains duplexed.
- Avoid starting duplexing for a structure until all connectors have been migrated to a software level that supports duplexing

Operational caveats

- System Managed duplexing is built upon system managed rebuild processing and user managed duplexing
- While duplexed, a structure is considered to be indefinitely in a rebuild phase
 - The Duplex Established phase of a duplexing rebuild process
- Therefore, display commands and messages issued for duplexed structures will use rebuild terminology
 - Rebuild
 - Duplexing Rebuild
 - Old Structure
 - New Structure

Operational caveats

- Policy change remains pending until physical instance(s) of structure allocated under old CFRM policy are deallocated
- Policy changed while structure duplexed, need to stop duplexing, start rebuild, then start duplexing to resolve pending changes
- Policy changed for simplex structure that is then duplexed, the new instance is allocated under the pending policy.
 - Need to stop duplexing (keep the new one), then start duplexing to resolve pending changes
- System managed CF structure duplexing is not started for a structure that has a pending policy change

Table of Contents

Overview of Duplexing

Hardware Requirements

Software Requirements

Migration

➤ **Summary**

Benefits of CF Duplexing

- Near-continuous availability for CF structure data
- Enable use of an all-ICF solution
- Moving structures during planned reconfiguration for exploiters that don't support structure rebuild

Cost of CF Duplexing

- Potential configuration/setup cost
 - CF storage resources for duplexed structure instances
 - CF-to-CF links
 - System-to-CF links
 - System and CF processor resources
- Coupling efficiency cost
 - Duplexed command service time “deltas” resulting from the duplexing protocols
 - z/OS (XES) pathlength deltas resulting from splitting/merging duplexed commands
 - Impact varies with request rate and read/write ratio of workload

Cost of CF Duplexing

- Some of the additional resource requirements may be satisfied by excess capacity already in place for high availability and failover capability
- Recovery for failures in a simplex environment requires excess capacity so workload can continue to run on surviving systems, surviving CFs, etc.
- Recovery for failures in a duplexed environment may actually consume less resource because the structures operate in simplex mode until failed component is restored

Additional Information

- System Managed CF Structure Duplexing (GM13-0103)
- z/OS Planning for Installation (GA22-7504)
- z/OS MVS™ Setting Up a Sysplex (SA22-7625)

▪ ibm.com/servers/eserver/zseries/ps0

▪ ibm.com/servers/eserver/zseries/zos/bkserv