



Racf and the Sysplex

Dan Mapes

Mail Station P301
522 South Road
Poughkeepsie, N.Y. 12601

dmapes@us.ibm.com



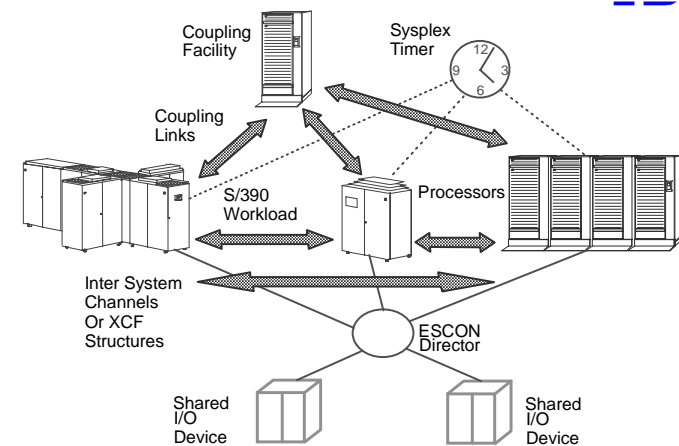
- What is a Parallel Sysplex?
- Why do we need a Sysplex?
- How is the RACF Parallel Sysplex support implemented?
 - ▶ Sysplex communication mode
 - ▶ Datasharing
 - ▶ RACF operational modes
- What are the Racf considerations in a Sysplex environment?
- Security Requirements
 - ▶ Utilities and Structures

What is a Parallel Sysplex?



- Combination of hardware and software
 - ▶ Multiple OS/390 systems linked together into a single image of computing resources.
 - ▶ Infrastructure for creating highly scalable and continuously available applications.
- The heart of the Parallel Sysplex is the Coupling Facility (CF) which provides:
 - ▶ Locking services (IRLM for DB2 and IMS resources)
 - ▶ Caching services (RACF database information)
 - ▶ List or queue services (JES2 checkpoint information)
- A coupling facility may be either a dedicated box, a standalone S/390, or an S/390 logical partition (LPAR)

Parallel Sysplex Elements



Parallel Sysplex Advantages



- High availability
 - ▶ Individual system maintenance won't affect the whole sysplex
- Scalability
 - ▶ Processing power can be added incrementally
- Dynamic workload balancing (WLM)
- Automatic Recovery (ARM)
- Total cost of computing
 - ▶ Parallel Sysplex licensing charge can significantly reduce the cost of IBM software
 - ▶ S/390 CMOS technology offers lower per MIP and environmental costs

© Copyright IBM Corporation, 1999

RACF's Parallel Sysplex Support



- Consists of two Modes of Operation
 - ▶ Sysplex communication: Use of the Cross-System Communication Facility (XCF) to communicate specific SETROPTS and RVARV commands to other systems
 - ▶ Sysplex data sharing: Use of the coupling facility to share data (cache) and serialize resources within a Parallel Sysplex
- Sysplex communication may be used without using sysplex data sharing
- Sysplex data sharing requires the use of sysplex communication

© Copyright IBM Corporation, 1999

RACF's Parallel Sysplex Support



- What are the benefits?
 - ▶ Simplified system administration through the automatic propagation of RVARV and SETROPTS commands
 - ▶ Use of the coupling facility as a cache for RACF data
 - ▶ Ability to share the RACF data base with a larger number of systems
 - ▶ Single system image for security relevant information
 - ▶ ... along with all of the other benefits of a Parallel Sysplex

© Copyright IBM Corporation, 1999

RACF Sysplex Communication



RACF Sysplex Communication

© Copyright IBM Corporation, 1999

RACF Sysplex Communication



- What is RACF Sysplex Communication?
 - ▶ Use of the Cross-System Communication Facility (XCF) to communicate specific SETROPTS and RVARY commands to other systems
 - ▶ Communication enabled through the XCF group IRRXCF00
 - There is one IRRXCF00 data sharing group per sysplex
 - ▶ Propagated commands appear to run longer
 - Commands are running on all systems in the sysplex not just the one the command was entered on.
 - ▶ Systems assume roles for command propagation
 - Coordinator (where the command was first entered) drives the other systems
 - Peers repond to coordinators

© Copyright IBM Corporation, 1999

RACF Sysplex Communication




- What should you know about?
 - ▶ Some SETROPTS commands will run holding RACF resources
 - SYSZRACF.DSDTPREP...
 - Prevents concurrent SETROPTS commands
 - SYSZRACF.DSDTDSDT...
 - Prevents mode changes (i.e. datashare/nodatashare)
 - SYSZRACF.SETROPTS
 - Recovery Trigger
 - ▶ Leave jobs owning these resources alone
 - They should not affect normal system throughput
 - No need to cancel users owning these for system throughput concerns
 - Don't let operations know what they are?

© Copyright IBM Corporation, 1999

Command Propagation



RVARY	ACTIVE INACTIVE SWITCH NODATASHARE DATASHARE
SETROPTS	RACLIST NORACLIST RACLIST REFRESH GLOBAL REFRESH GLOBAL REFRESH GENERIC REFRESH WHEN(PROGRAM) WHEN(PROGRAM) REFRESH



Command propagation occurs when RACF is configured for sysplex communication

© Copyright IBM Corporation, 1999

Sysplex Communication Requirements



- Serialization
 - ▶ Two Qnames are important
 - SYSZRAC2
 - Prevents profile updates during the SETROPTS process
 - SYSZRACF
 - Prevents concurrent execution of SETROPTS
 - ▶ SYSZRACF and SYSZRAC2 resource names must not be in the exclusion Resource Name List (RNL)
 - ▶ Systems using sysplex communication only may share the RACF database with systems outside of the sysplex
 - However, manual propagation of SETROPTS and RVARY commands is required for those systems which are outside of the sysplex.

© Copyright IBM Corporation, 1999

Sysplex Communication Requirements



- General requirements required:
 - ▶ RACF Version 2, or OS/390 Security Server
 - ▶ MVS 4.2
 - ▶ XCF signaling must be active
 - ▶ RACF database must be on shared DASD

Benefits of Sysplex Communication




- What are the benefits of Sysplex Communication?
 - ▶ Eliminate errors in manually propagating SETROPTS and RVARV commands
 - You don't have to remember the system that required the refresh, it happens on all systems.
 - Change those scheduled refreshes such that they only run on one system!
 - ▶ Ease migration to data sharing mode
 - Rvary datashare can be issued (remember every system must share the same database!)
 - ▶ Resident data blocks are enabled for the backup RACF database (Datasharing mode is not required).

RACF Data Sharing



RACF Data Sharing


Sharing a RACF Database - Before



Sharing a RACF Database - Before



- Synchronization with RESERVE/RELEASE
- Index and Data Blocks cached locally in Resident Data Blocks
- After update by another system, local cache must be invalidated.



© Copyright IBM Corporation, 1999

RACF Data Sharing



- What is RACF Data Sharing?
 - ▶ Use of the coupling facility as a large data buffer for the RACF database
 - ▶ Elimination of some serialization when reading the RACF database
 - We use read serialization whenever we can (such as ALTERI requests)
 - First logon uses ALTERI serialization
 - ▶ Use of global ENQs instead of RESERVEs to serialize the use of the RACF database
 - ▶ Can be used in single-system sysplexes ("monoplex"), or multi-system sysplexes. Cannot be used in XCFLOCAL mode.

© Copyright IBM Corporation, 1999

RACF Data Sharing



- Migration Considerations
 - ▶ GRS RNLS
 - Installations using RNLS to exclude SYSZRACF.database_name (they like reserves)
 - Subsequently convert to datasharing (but did not read the datasharing "some assembly required" section of the documentation)
 - Results in no database serialization and other problems with propagated commands
 - ▶ MIM
 - Process = select means you have to spell out the resource queue names to MIM.
 - Look out for exemptions

© Copyright IBM Corporation, 1999


Sysplex Data Sharing Requirements



- What are the database sharing rules?
 - ▶ All systems must:
 - Have compatible data set name tables (DSNT)
 - Note: RACF stores only the DSN in the DSNT but OW35420 enabled volser checking for IPLs
 - Have identical database range tables (RRNG)
 - *Recommendation: Share the DSNT and RRNG*

© Copyright IBM Corporation, 1999

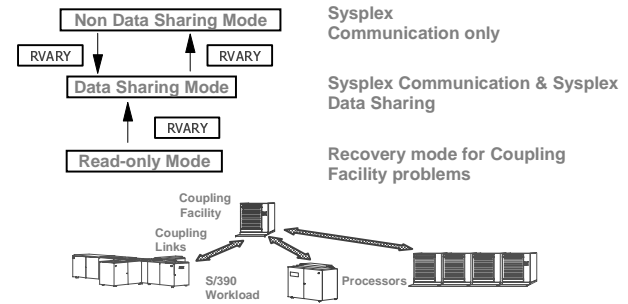
RACF Operational Modes



- ★ RACF "Classic" Mode: No change from previous RACF releases
- ★ Non Data Sharing Mode: Sysplex communication only
- ★ Data Sharing Mode: Exploits Coupling Facility
- ★ Read-only Mode: Recovery mode for Coupling Facility problems

© Copyright IBM Corporation, 1999

RACF Operational Modes ...



Sage Advice:

Sysplex Communication will allow the RVARY DATASHARE command to be issued...Ensure that it isn't attempted unless you are ready for Data Sharing Mode

© Copyright IBM Corporation, 1999

RACF Read Only Mode



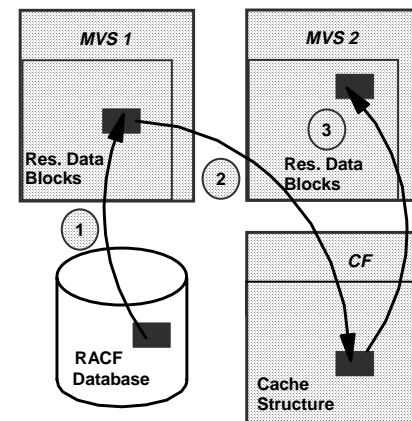
- Entered when the system is not properly connected to the CF.
 - ▶ IPL but no CF available and datasharing is requested, or RVARY datashare was used without CF connected.
 - ▶ Use RVARY to leave datasharing mode
- A system loses connectivity to the CF and there are no other CF's available.
 - ▶ Someone pulls the wrong plug
 - ▶ Use Rvary No datashare if this system is a key player in the plex.
- Redundant CF's allow rebuild based on SFM policy and REBUILDPERCENT in the CFRM.
 - ▶ See OW30814 for OS390R3&+ with no SFM policy

© Copyright IBM Corporation, 1999

Racf Caching: Reading RACF Data



1. A block that is not in cache is read from the database into the local cache
2. The block is copied into the cache structure and associated (register interest)
3. If MVS 2 references the block, it is copied from the cache structure into the Resident Data Blocks and associated (register interest)

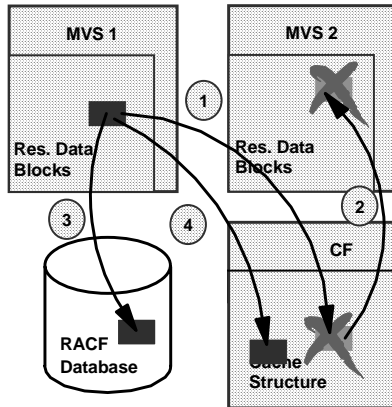


© Copyright IBM Corporation, 1999

RACF Caching: Updating a Block



1. Invalidate the copy of the block in the cache structure
2. Cross-invalidation lets the copies of the block in all other sharing systems become invalid
3. Update the block in the RACF database
4. The block is copied to the cache structure and associated (register interest)



© Copyright IBM Corporation, 1999

Data Sharing Requirements



- What is required to use data sharing?
 - ▶ RACF Version 2, or OS/390 Security Server
 - ▶ MVS 5.1 (or higher)
 - ▶ XCF signaling must be active
 - ▶ RACF database must be on shared DASD
 - ▶ SYSZRACF and SYSZRAC2 resource names must not be in the exclusion Resource Name List (RNL) and MIM users have to manage these Qnames
 - ▶ System must be connected to a coupling facility
 - ▶ RACF coupling facility structures must be defined in the coupling facility resource management (CFRM) policy

© Copyright IBM Corporation, 1999

Data Sharing Requirements



- What are the rules?
 - ▶ All of the sysplex communication rules, plus all systems must:
 - ▶ Share the same RACF database
 - ▶ Must be in the SAME sysplex
 - ▶ Database must be on the SAME volser
 - ▶ Operating in a compatible mode (data sharing or read-only)
 - ▶ Connect to the coupling facility that contains the RACF data structures

© Copyright IBM Corporation, 1999

RACF Sysplex Enablement



- ICHRDSNT contains the necessary switches

```

DC  ALL(2)                * number of pairs
DC  CL44'RACFDB.DSP1'    * primary ds
DC  CL44'RACFDB.DSB1'    * back up ds
DC  ALL(255)             * resident blocks
DC  X'CC'                * stats and profiles
DC  CL44'RACFDB.DSP2'    * primary ds
DC  CL44'RACFDB.DSB2'    * back up ds
DC  ALL(255)             * resident blocks
DC  B'11001100'         * stats, profiles, switches
    
```

- ▶ Don't care (at one time this requested buffers)
- ▶ Don't care
- ▶ Data Sharing is requested
- ▶ Enable Sysplex Communication
- ▶ Don't care
- ▶ Don't care
- ▶ All Updates, including statistics
- ▶ All Updates, minimal statistics (recommended)

© Copyright IBM Corporation, 1999

Sysplex Security Aspects



- IXCMIAPU
 - ▶ Administration utility for updating the policy datasets
 - ▶ The FACILITY class allows policy listings or updates
 - MVSADMIN.XCF.CFRM READ (CFRM report)
 - MVSADMIN.XCF.CFRM UPDATE (CFRM changes)
- Connectors need authorization to the structures via the FACILITY class
 - ▶ RACFDS runs TRUSTED (no STARTED entry required)
 - ▶ Datasharing FACILITY class checks
 - IXLSTR.IRRXCF00_P001 ALTER (primary)
 - IXLSTR.IRRXCF00_B001 ALTER (back up)
 - ▶ Others such as DB2 may not be TRUSTED.
 - Check individual program requirements

© Copyright IBM Corporation, 1999

Benefits Of Data Sharing



- What are the benefits of RACF Data Sharing?
 - ▶ Reduced I/O to the RACF database by using the coupling facility as a cache
 - ▶ Less contention for the RACF database
 - Elimination of RESERVE/RELEASE serialization
 - Improved serialization when reading the RACF database
 - ▶ Improved cache management by invalidating only the affected cached data blocks
 - ▶ Potential for faster access to the backup database due to increased caching
 - ▶ ... plus all of the other benefits of the Parallel Sysplex

© Copyright IBM Corporation, 1999

Reading and Reference Material



- OS/390 Security Server (RACF) Planning: Installation and Migration (GG24-4405)
- OS/390 Security Server (RACF) Systems Programmer's Guide (SC28-1913)
- OS/390 Parallel Sysplex Overview (GC28-1860)
- OS/390 Setting Up a Sysplex (GC28-1779)

© Copyright IBM Corporation, 1999

Reminder: Service End Dates



- RACF 1.9.0 and RACF 1.9.2
 - ▶ Ended on 26 April, 1997
- RACF 2.1
 - ▶ Ends on 06/30/99
 - ▶ Announcement letter 997-230 contains the details
- RACF 2.2
 - ▶ End of service not yet announced
- OS/390 Releases:
 - ▶ OS/390 R1, R2, R3, V2R4 service continued until at least January 31, 2001
 - ▶ Announcement letter 298-049

© Copyright IBM Corporation, 1999

RACF Home Page



- <http://www.s390.ibm.com/racf>
 - ▶ Latest release information on RACF
 - ▶ Links to announcement letters
 - ▶ Sample code
 - DBSYNC to compare two RACF data bases
 - RACFDB2 to migrate DB2 access control to RACF
 - RACFICE to create reports
 - OS390ART for a web-based reporting tool
 - RACTRACE tracing facility
 - ▶ Frequently asked questions (FAQ)
 - ▶ RACF user group information
 - ▶ RACF List (RACF-L) internet forum information