

***Unclog Your Systems with z/OS 2.1 –
Something New and Exciting for
Catalog
Part 1 and 2 (Sessions 12977/12978)***

**Terri Menendez
IBM Corporation
STSM RLS/VSAM/Catalog Development
San Jose, CA USA
terriam@us.ibm.com
Jan 2013**

Notices and Disclaimers



- NOTICES AND DISCLAIMERS
- Copyright © 2013 by International Business Machines Corporation.
- No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.
- Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) described herein at any time without notice.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Consult your local IBM representative or IBM Business Partner for information about the product and services available in your area.
- Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.
- THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

Notices and Disclaimers



- This information is provided on an "AS IS" basis without warranty of any kind, express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Some jurisdictions do not allow disclaimers of express or implied warranties in certain transactions; therefore, this statement may not apply to you.
- This information is provided for information purposes only as a high level overview of possible future products. PRODUCT SPECIFICATIONS, ANNOUNCE DATES, AND OTHER INFORMATION CONTAINED HEREIN ARE SUBJECT TO CHANGE AND WITHDRAWAL WITHOUT NOTICE.
- USE OF THIS DOCUMENT IS LIMITED TO SELECT IBM PERSONNEL AND TO BUSINESS PARTNERS WHO HAVE A CURRENT SIGNED NONDISCLOSURE AGREEMENT ON FILE WITH IBM. THIS INFORMATION CAN ALSO BE SHARED WITH CUSTOMERS WHO HAVE A CURRENT SIGNED NONDISCLOSURE AGREEMENT ON FILE WITH IBM, BUT THIS DOCUMENT SHOULD NOT BE GIVEN TO A CUSTOMER EITHER IN HARDCOPY OR ELECTRONIC FORMAT.

Important notes:

- IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.
- IBM makes no warranties, express or implied, regarding non-IBM products and services, including but not limited to Year 2000 readiness and any implied warranties of merchantability and fitness for a particular purpose. IBM makes no representations or warranties with respect to non-IBM products. Warranty, service and support for non-IBM products is provided directly to you by the third party, not IBM.
- All part numbers referenced in this publication are product part numbers and not service part numbers. Other part numbers in addition to those listed in this document may be required to support a specific device or function.
- MHz / GHz only measures microprocessor internal clock speed; many factors may affect application performance. When referring to storage capacity, GB stands for one billion bytes; accessible capacity may be less. Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and the population of all hard disk drive bays with the largest currently supported drives available from IBM.

IBM Information and Trademarks

- The following terms are trademarks or registered trademarks of the IBM Corporation in the United States or other countries or both: the e-business logo, IBM, xSeries, pSeries, zSeries, iSeries.
- Intel, Pentium 4 and Xeon are trademarks or registered trademarks of Intel Corporation. Microsoft Windows is a trademark or registered trademark of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. Other company, product, and service names may be trademarks or service marks of others.
- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not necessarily tested those products in connection with this publication and cannot confirm the accuracy of 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.
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:
 - IBM Director of Licensing
 - IBM Corporation
 - North Castle Drive
 - Armonk, NY 10504-1785
 - U.S.A.



Trademarks



Trademarks

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

BookManager*	Enterprise Storage Server*	IP PrintWay	RMF
CICS*	ES/9000*	Language Environment*	S/370
DB2*	FlashCopy*	Lotus*	S/390*
DB2 Universal Database	GDPS*	Multiprise*	Tivoli*
developerWorks*	HiperSockets	MVS	TotalStorage*
DFSMSdftp	IBM*	Notes*	WebSphere*
DFSMSdss	IBM eServer	OS/390*	z/Architecture
DFSMSHsm	IBM e(logo)server*	Parallel Sysplex*	z/OS*
DFSMSrmm	IBM logo*	RACF*	zSeries*
DFSORT	IMS	RAMAC*	
Domino	InfoPrint*		

Intel is a trademark of the Intel Corporation in the United States and other countries.

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group in the United States and other countries.

* 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.

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.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

Agenda

- ❑ Unclog Your Systems with z/OS 2.1 – Part 1
 - Overview of Enhancements
 - Catalog Issues
 - Current Catalog Design Overview
 - New Catalog Design Overview
 - New Externals:
 - DEFINE/DELETE/ALTER USERCATALOG parameters
 - Catalog MODIFY commands
 - Catalog Vendor Interface
 - Serialization changes
 - System Recovery
 - DSS Parameters
 - Example recover scenario
 - RLS for AMS
 - Toleration maintenance / Recommendations

Agenda

- Unclog Your Systems with z/OS 2.1 – Part 2
 - Performance Options
 - Performance Reports
 - Performance Benchmarks

Unclog Your Systems with z/OS 2.1 – Part 1

Overview of Enhancements

Overview of Enhancements

□ *z/OS 2.1 provides new catalog functionality to:*

- Improve performance by accessing individual user catalogs using VSAM Record Level Sharing (RLS) within a parallel sysplex. Referred to as “RLS for Catalog”.
- Improve integrity and availability with new sysplex wide commands to control access to individual user catalogs within a parallel sysplex.

Catalog Issues

Catalog Issues

❑ *Performance*

- Contention on SYSIGGV2 bcsname when updating catalogs.
- Limited catalog buffering and buffer invalidation.
- Limited VSAM buffers/strings/storage.

❑ *Availability*

- Catalogs need to be split for contention issues.
- Catalogs unavailable for splitting, recovering, and other maintenance activities.

❑ *Integrity*

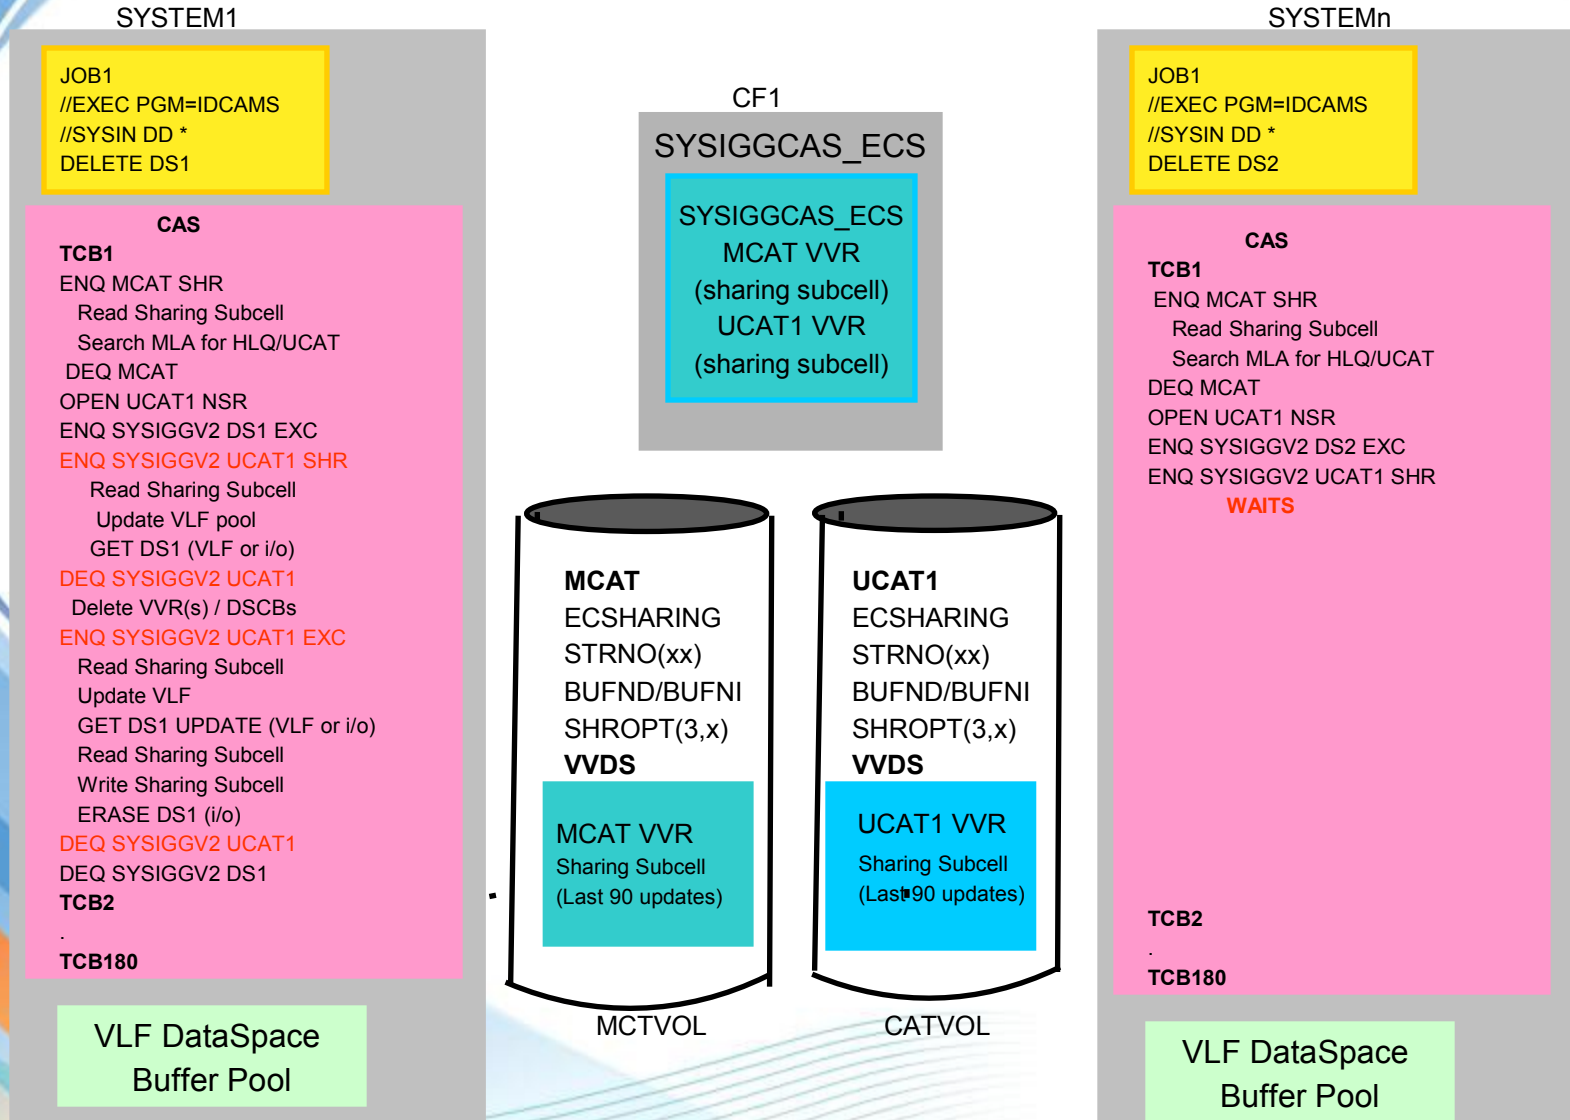
- Catalogs damaged by utilities updating a catalog while catalog is opened by CAS.
- Lack of sysplex control of closing and serializing catalogs.

❑ *Recovery*

- Long/error prone forward recovery procedures.

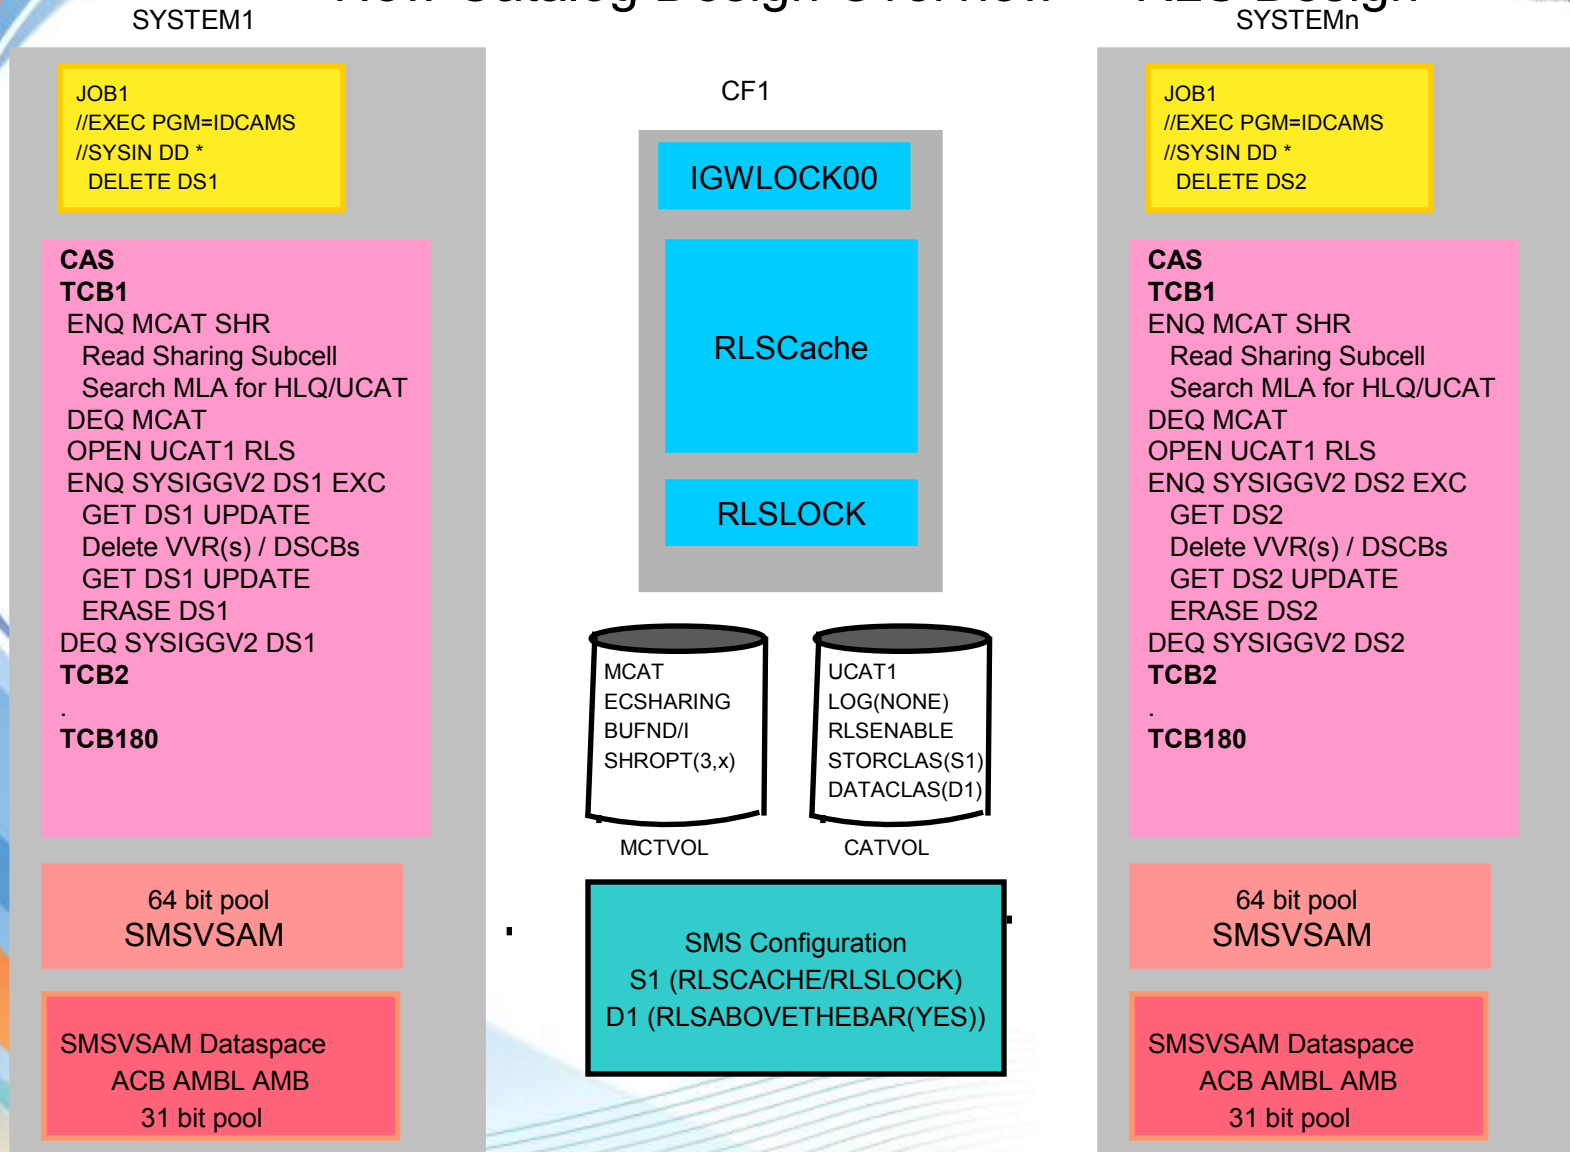
Current Catalog Design

Current Catalog Design Overview – “NonRLS Design”

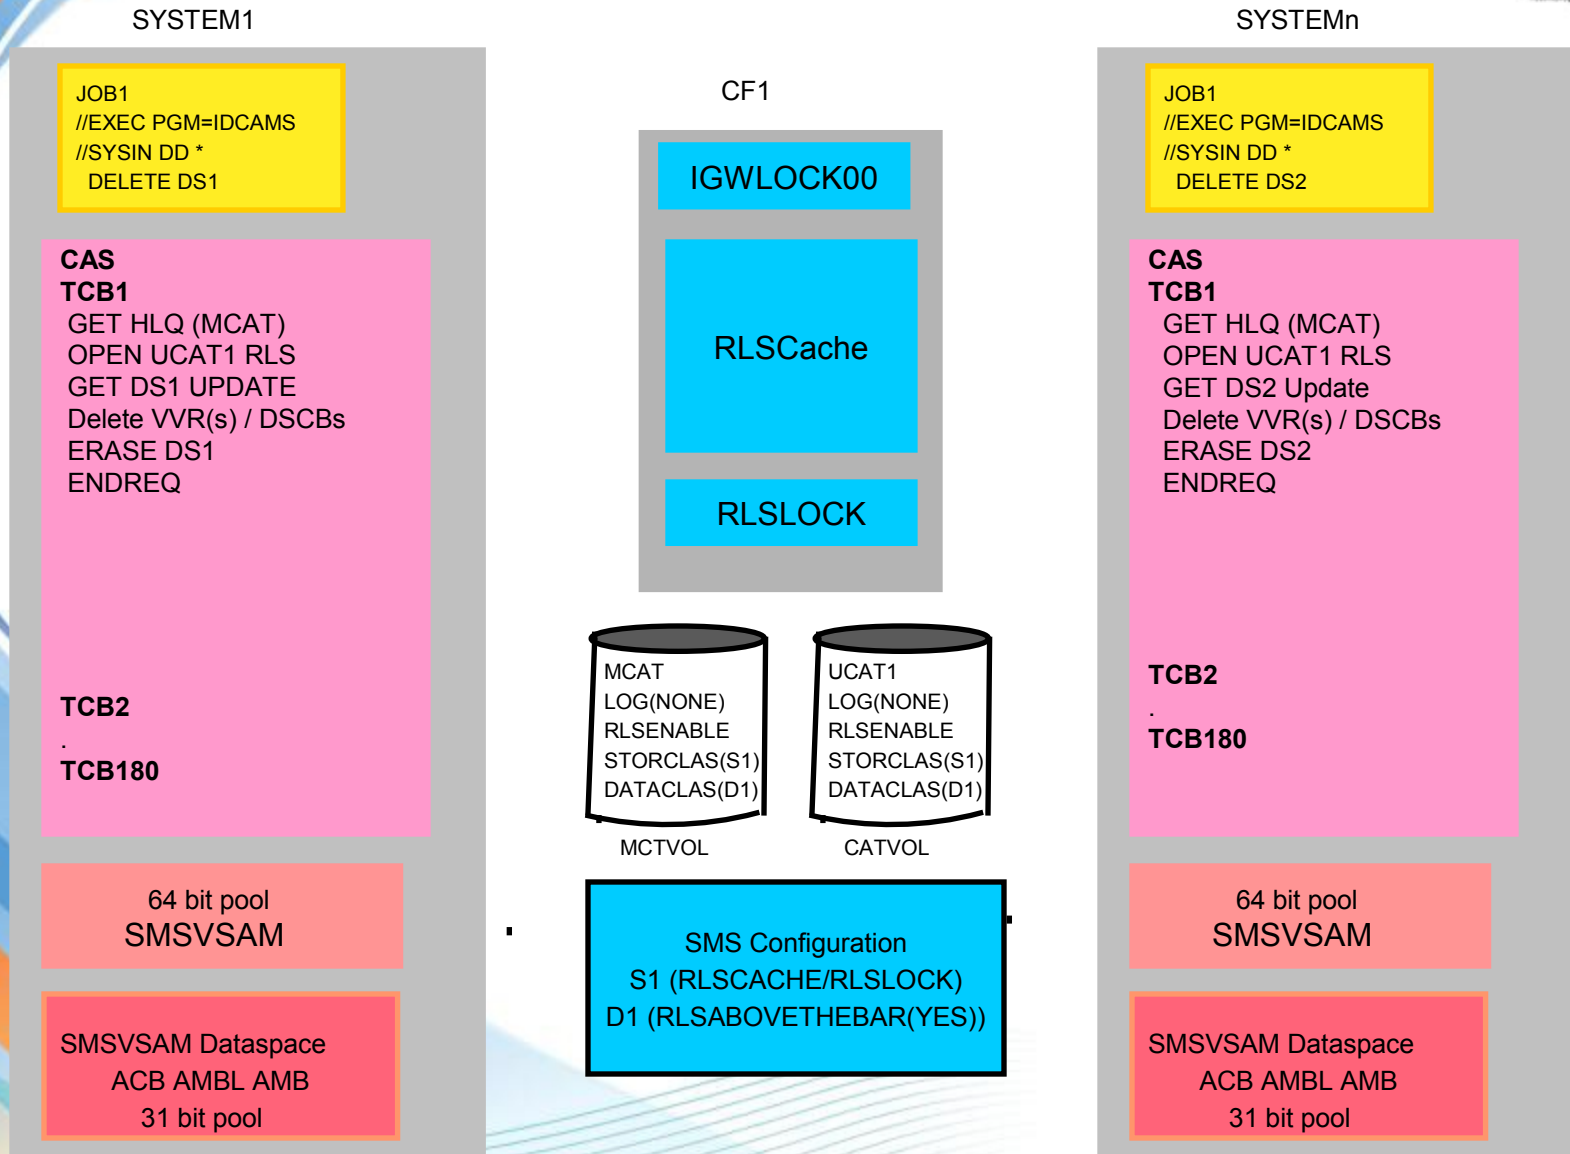


New Catalog Design

New Catalog Design Overview – “RLS Design”



New Catalog Design2 Overview – “RLS Design2”



New Externals

z/OS RLS for Catalog - New Externals

❑ *New DEFINE USERCATALOG parameters:*

DEFINE USERCATALOG(NAME(bcsname) -
LOG(NONE) -
RLSQUIESCE/RLSENABLE -
SUSPEND/RESUME -
RECONNECT
LOCK/UNLOCK
STORCLAS(storclasname))

LOG parameter defines the catalog as non-recoverable and is eligible for RLS access.

RLSQUIESCE indicates the catalog is to be accessed by nonRLS on first access following define.

RLSENABLE indicates the catalog is to be accessed by RLS on first access following define.

SUSPEND indicates all unauthorized catalog requests will be suspended following the define*.

RESUME indicates all catalog requests will be resumed following the define*.

RECONNECT reuses the existing UCONN record (alias information).

LOCK indicates all unauthorized requests will be failed following the define*.

UNLOCK indicates all requests will be allowed following the define*.

STORCLAS specifies a RLS capable storage class (CACHESET/LOCKSET)

* Authorized users are those ids permitted to RACF facility class IGG.CATLOCK with read access. These commands apply to all catalogs (RLS and nonRLS).

z/OS RLS for Catalog – New Externals

❑ *New DELETE/ALTER USERCATALOG parameters:*

DELETE USERCATALOG(NAME(bcsname) -
NODISCONNECT)

NODISCONNECT does not delete the UCONN record (alias information).

Catalogs accessed by RLS at the time of the delete will be closed sysplex wide prior to the delete.

ALTER 'bcsname' –

LOG(NONE)

RLSQUIESCE/RLSENABLE -

SUSPEND/RESUME –

LOCK/UNLOCK –

NULLIFY(LOG))

LOG parameter defines the catalog as non-recoverable and is eligible for RLS access.

RLSQUIESCE invokes a sysplex wide close of the catalog for RLS, and enables the catalog for nonRLS access on the next reference.

RLSENABLE invokes a sysplex wide close of catalog for nonRLS, and enables the catalog for RLS access on the next reference.

SUSPEND invokes a sysplex wide close of the catalog and suspends unauthorized requests in the client space.

RESUME invokes a sysplex wide resume of suspended requests.

LOCK invokes a sysplex wide close of the catalog and fails unauthorized requests.

UNLOCK invokes a sysplex wide unlock of the catalog.

NULLIFY removes the LOG attribute and resets RLS related flags. Catalog is not longer eligible to RLS access. Catalog must be closed at the time of the alter.

SHCDS CFRESETDS(bcsname) removes LOG parameter and resets RLS flags for closed catalogs.

z/OS RLS for Catalog – New Externals

❑ *New Catalog MODIFY Commands:*

Switching between RLS and nonRLS access:

F CATALOG,RLSENABLE(ucatname*)

Sysplex wide switch to RLS access: closes catalog(s) for nonRLS and enables catalog(s) for RLS access. The catalog(s) will be opened for RLS access on the next reference.

F CATALOG,RLSQUIESCE(ucatname*)

Sysplex wide switch to nonRLS access: closes catalog(s) for RLS access and enables catalog(s) for nonRLS access. The catalog(s) will be opened for nonRLS on next reference.

F CATALOG,RLSQUIESCE,SYSTEM

Sysplex wide switch to nonRLS access: closes catalogs enabled and referenced for RLS by this system and enables the catalogs for nonRLS access. The catalog(s) will be opened for nonRLS on next reference.

F CATALOG,RLSENABLE,SYSTEM

Sysplex wide switch to RLS access: closes catalogs enabled and referenced for nonRLS by this system and enables the catalogs for RLS access. The catalog(s) will be opened for RLS on the next reference.

z/OS Catalog – Externals

❑ *New Catalog MODIFY Commands (cont):*

Catalog recovery commands:

F CATALOG,RECOVER,LOCK(ucaname*)

This command will quiesce and close the catalog(s) across the sysplex, LOCK the catalog(s), then fail new requests.

F CATALOG,RECOVER,UNLOCK(ucaname*)

This command will unlock the catalog(s).

F CATALOG,RECOVER,SUSPEND(ucaname*)

This command will quiesce and close the catalog(s) requests across the sysplex. Requests from permitted users will be allowed to complete.

F CATALOG,RECOVER,RESUME(ucaname*)

This command will resume waiting requests across the sysplex.

Note: The above commands apply to all catalogs (RLS and nonRLS).

z/OS RLS for Catalog – Externals

❑ Updated Catalog MODIFY Commands:

Display commands:

```
F CATALOG,ALLOCATED
```

```
IEC348I Allocated Catalogs
```

```
CAS*****  
* FLAGS -VOLSER-USER-CATALOG NAME % *  
* YS-RU- XP0202 0001 USERCAT1 1 *  
*****  
* Y/N-ALLOCATED TO CAS, S-SMS, V-VLF, I-ISC, C-CLOSED, D-DELETED, *  
* R-SHARED, A-ATL, E-ECS, K-LOCKED,U-RLS SHARED,W-SUSPEND *  
*CAS*****
```

Above example shows shared usercatalog USERCAT1 as allocated and open for RLS access.

z/OS RLS for Catalog – Externals

❑ *Catalog Vendor Interface:*

IGGQUIES – Available for purchase in the Advanced Customization Guide

NAME IGGQUIES FUNCTION=function,
SPHERE=catalogname,
RETCODE=retcode,
RSNCODE=rsncode,
ERROR_DATA=address,
MF=I

IGGQUIES FUNCTION=QUICOPY. This interface will quiesce all update requests and wait any new update requests across the sysplex for the specified RLS catalog.

IGGQUIES FUNCTION=QUICEND. This interface will resume update requests for the specified RLS catalog.

IGGQUIES FUNCTION=RLSENABLE. This interface will close a catalog across the sysplex and enable the catalog for RLS access on the next reference.

IGGQUIES FUNCTION=RLSQUIESCE. This interface will close a catalog across the sysplex and enable the catalog for nonRLS access on the next reference.

IGGQUIES FUNCTION=SUSPEND/RESUME. This interface will close a catalog across the sysplex and wait/resume unauthorized requests in the client address space.

IGGQUIES FUNCTION=LOCK/UNLOCK. This interface will close a catalog across the sysplex and fail/allow unauthorized requests.

z/OS RLS for Catalog – Externals

❑ *DFDSS Parameters:*

DUMP –

DATASET(INCLUDE(bcsname)) -
OUTDDNAME(DASD2))

DFDSS will invoke a **Quiesce for Copy (QUICOPY/QUICEND)** for catalogs opened for RLS. The QUICOPY will suspend update requests only in order to obtain a sharp copy of the catalog. The QUICEND will resume update requests. For nonRLS catalogs, DSS will use existing serialization (SYSIGGV2).

RESTORE -

INDDNAME(INDD1) OUTDDNAME(OUTDD1) -
DATASET(INCLUDE(bcsname)) **BCSRECOVER(LOCK/SUSPEND)**

LOCK will invoke a sysplex wide close of the catalog and lock the catalog (failing new unauthorized requests), if the catalog is not already locked or suspended. If DSS locks the catalog, it will unlock it at the end of the restore.

SUSPEND will invoke a sysplex wide close of the catalog and suspend new unauthorized requests in the client space if the catalog is not already locked or suspended. If DSS suspends requests, it will resume the Requests at the end of the restore.

Note: The catalog must be locked or suspended either through the BCSRECOVER option, or from prior commands, otherwise the restore will be failed.

z/OS RLS for Catalog – Externals

❑ *Serialization Changes:*

- Replace SYSIGGV2 bcsname resource with VSAM RLS record level locking (Catalog opened for RLS access).
- SMSVSAM will hold SYSIGGV2 bcsname SHARE while a catalog is opened for RLS access (ensures catalog data integrity from programs relying on SYSIGGV2 to serialize the catalog):

S=SYSTEMS SYSIGGV2 UCAT1

SYSNAME	JOBNAME	ASID	TCBADDR	EXC/SHR	STATUS
SYSTEM1	SMSVSAM	000A5	008FF270	SHARE	OWN

z/OS RLS for Catalog – Externals

❑ System Recovery (SMSVSAM Termination Scenario):

```
F CATALOG,ALLOCATED
IEC351I CATALOG ADDRESS SPACE MODIFY COMMAND ACTIVE
IEC348I ALLOCATED CATALOGS 266
*CAS*****
* FLAGS -VOLSER-USER-CATALOG NAME                % *
* YSU-R- XP0202 0001 RLSCAT                        1 *
*****
* Y/N-ALLOCATED TO CAS, S-SMS, V-VLF, I-ISC, C-CLOSED, D-DELETED , *
* R-SHARED, A-ATL, E-ECS SHARED, K-LOCKED, U-RLS, W-SUSPENDED *
*CAS*****
V SMS,SMSVSAM,TERMINATESERVER
```

← RLSCAT opened for RLS access

← Should have used the FORCE SMSVSAM,ARM if workloads are not quiesced.

***IGW408I SMSVSAM SUCCESSFULLY TERMINATED AT END OF MEMORY**

(request issued to RLSCAT)

```
IEC251I 016-0609,CATALOG,ALLOCATE,SYS00006,,,RLSCAT
IEF196I IEC251I 016-0609,CATALOG,ALLOCATE,SYS00006,,,RLSCAT
```

← Catalog closed for RLS, SMSVSAM not avail.

(5 minutes later)

***64 IEC365D SMSVSAM ADDRESS SPACE STILL NOT AVAILABLE FOR CATALOG ADDRESS SPACE. REPLY RLSQUIESCE OR CANCEL**

```
V SMS,SMSVSAM,ACTIVE
```

```
IGW414I SMSVSAM SERVER ADDRESS SPACE IS NOW ACTIVE. 338
```

```
IEE400I THESE MESSAGES CANCELLED - 64.
```

← SMSVSAM issues ENF and WTOR is cancelled

(request issued to RLSCAT)

```
* YSU-R- XP0202 0001 RLSCAT                1 *
```

← RLSCAT reopened for RLS access

z/OS RLS for Catalog - Externals

❑ *Example Catalog Forward Recovery Scenario:*

1) Backup catalog:

- A) DFDSS or AMS will invoke the QUICOPY/QUICEND interface for RLS eligible catalogs. For nonRLS catalogs, DSS will use existing logic (SYSIGGV2).
- B) OEM backup utilities must detect RLS eligible catalogs, and invoke the IGGQUIES interface to QUICOPY/QUICEND update requests. The CSI or FAMS interfaces can be used to retrieve the LOGPARM field to test if the catalog is eligible for RLS.

2) To forward recover the catalog to a different volume:

- A) Suspend catalog requests against the current target catalog: F CATALOG,RECOVER,SUSPEND(ucat1) or ?IGGQUIES FUNCTION=SUSPEND. This step suspends all new catalog requests for ucat1 and closes the catalog across the sysplex.
- B) DELETE ucat1 NODISCONNECT. This step deletes the DSCBs and VVRs and saves the alias information.
- C) DEFINE ucat1 RECONNECT VOL(volser2) SUSPEND and optionally LOG(NONE) RLSQUIESE or RLSENABLE. This step defines the catalog, using existing UCONN record and updates the volser to volser2, and ensures the SUSPEND option is still set.
- D) IMPORT CONNECT VOL(volser2) ALIAS on sharing systems with separate MCATs. This step updates the UCONN with volser2
- E) Restore backup copy:
 - a) Using DSS Restore to backup copy. DSS will detect the catalog is suspended and is a empty pre-existing target, and will proceed to copy the data from the backup file using its own i/o method. With the suspended catalog there is no need for any further serialization (i.e. sysiggv2).
 - b) OEM restore utilities should follow DSS's example.
- F) Apply SMF record updates to bring the new target up to date since last backup (if backup was done prior to suspending this catalog). Use ICFRU or other equivalent catalog utility to apply the SMF records.
- G) Resume catalog requests: F CATALOG,RECOVER,RESUME(ucat1) or IGGQUIES FUNCTION=RESUME. Catalog requests will now be allowed to open the catalog across the sysplex with either RLS or nonRLS.

RLS for AMS

- Enhance IDCAMS commands PRINT, REPRO, IMPORT and EXPORT to be able to open VSAM data sets using RLS.
- A new optional keyword **RLSSOURCE** ({NO|YES| QUIESCE}) and/or **RLSTARGET** ({NO|YES|QUIESCE}) will be implemented for the PRINT, REPRO, IMPORT and EXPORT commands:
 - PRINT: **RLSSOURCE** ({NO|YES|QUIESCE}) for the input data set
 - REPRO: **RLSSOURCE** ({NO|YES|QUIESCE}) and **RLSTARGET** ({NO|YES|QUIESCE})
 - IMPORT: **RLSTARGET** ({NO|YES|QUIESCE}) for the target data sets
 - EXPORT: **RLSSOURCE** ({NO|YES|QUIESCE}) for the source data sets

Toleration Maintenance / Recommendations

- The SMSVSAM address space must be active on down level systems when opening a catalog previously accessed by RLS.
- The Following APARs are required on down level systems when accessing catalogs which may have been accessed for RLS on z/OS 2.1:
 - OA36403
 - OA36409
 - OA36916
 - OA36492
 - OA36422
 - OA36414

Unclog Your Systems with z/OS 2.1 – Part 2

Agenda

- ❑ Unclog Your Systems with z/OS 2.1 – Part 2
 - Performance Options
 - Performance Reports
 - Performance Benchmarks

Performance Options

Performance Options

☐ Buffering/Sharing Options:

○ NonRLS

- NSR Buffering
- In Storage Cache (ISC)
- Virtual LookAside Facility (VLF)
- Enhanced Catalog Sharing (ECS)

○ RLS

- Buffering
- Caching
- Locking
- Data Set
- Example of RLS Sysplex Configuration

Performance Options – Buffering/Sharing

❑ Buffering/Sharing (nonRLS):

○ NSR Strings/Buffers:

➤ DEFINE USERCATALOG:

❑ STRNO(2/n) – specifies the number of VSAM strings (i.e. number of concurrent requests per system). Max=255.

❑ BUFND(n) – specifies the number of data buffers (ACBBUFND).

❑ BUFNI(n) – specifies the number of index buffers (ACBBUFNI).

➤ Catalog will adjust BUFND and BUFNI based on the following algorithm:

$BUFND = \text{MIN}(32767, \text{MAX}(\text{ACBBUFND}, ((32400/\text{CISIZE}) + 1 + \text{STRNO})))$

$BUFNI = \text{MIN}(32767, \text{MAX}(\text{ACBBUFNI}, (\#\text{IndexLevels} + 1 + \text{STRNO})))$

➤ Too low of a STRNO will result in contention on the SYSZRPLW resource.

Performance Options

❑ Buffering (nonRLS):

○ In Storage Cache (ISC)

- Default buffering option.
- Intended for non-shared catalogs or small read only catalogs.
- Each usercat is allowed a fixed amount of space (3041 BCS records).
- Space is reclaimed via LRU algorithm.
- MCAT is allowed unlimited space.
- For shared usercats, any update on a sharing system will invalidate the entire ISC.

Performance Options

❑ Buffering (nonRLS):

○ Virtual Lookaside Facility (VLF)

- Intended for shared catalogs.
- Total cache size for all catalogs determined by the MAXVIRT parameter in the COFVLFxx parmlib member.
- Range of MAXVIRT is 256-524288 (2 gig) 4K blocks.
- Space is reclaimed (trimmed) via LRU algorithm.
- Catalog tracks the last 90 updates to the catalog via the sharing subcell in the catalog's data VVR). The subcell tracks catalog updates from sharing systems and is used to keep the VLF pool coherent. More than 90 updates from sharing systems between references results in the purge of the VLF cache.

Performance Options

❑ Sharing (nonRLS):

○ Enhanced Catalog Sharing (ECS)

- Intended for shared catalogs.
- The ECS function extends the catalog sharing design by placing the catalog's data VVR in a CF cache structure (SYSIGGCAS_ECS).
- By placing the data VVR in the CF, the i/o to retrieve it can be avoided (at least one i/o per catalog request).

Performance Options

☐ RLS

○ Buffering:

- System managed buffering via LRU algorithm.
- Choice of 2 buffer pools (default 31 bit, optional 64 bit)
- Catalog CIs placed in buffer pool and registered with a CF cache structure.

○ Caching:

- CF Cache assigned via catalog's storclas.
- CF cache acts as a global buffer pool.
- Catalog's CIs are registered with cache for XI purposes.
- Catalog CIs (by default) reside in the cache structure.

Performance Options

❑ RLS (cont.)

○ Locking:

- CF lock structure (IGWLOCK00) assigned by default.
- Optional lock structure assigned by storclas.
- Contains data set level locks used to serializing splits/reclaims, spanned updates, and to communicate
- Contains record level locks for serializing access to individual catalog records.
- Catalog design uses No Read Integrity (NRI) when reading catalog records (i.e. no shared locks).

Performance Options

❑ RLS (cont.)

○ Data Set /Catalog:

➤ Dataclac controls 64 bit buffering and CF caching options.

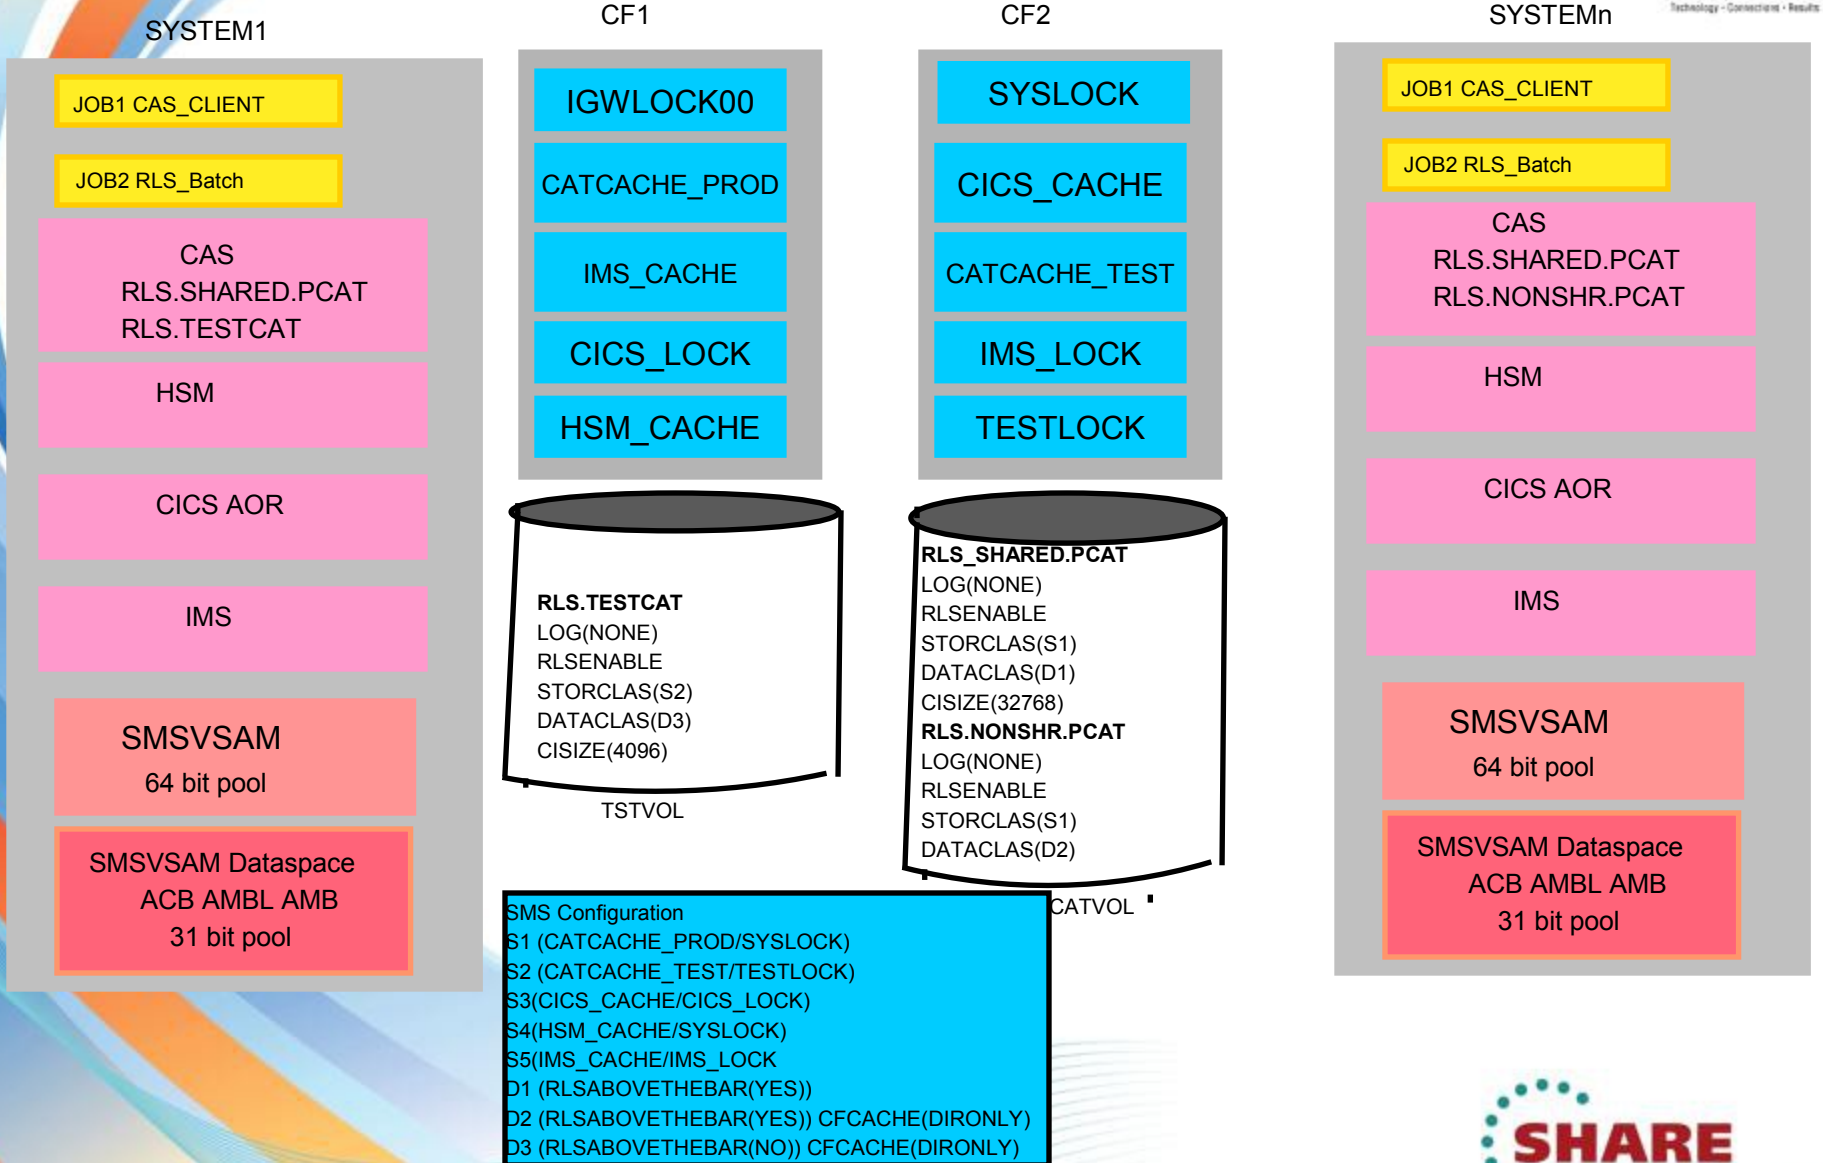
➤ CISIZE becomes important for RLS catalogs:

❑ Removal of SYSIGGV2 catname resource allows for lower level contention issues such as CI split/reclaim, CI, or spanned updates.

❑ Larger CISIZE recommended for high insert/erase catalogs or sequentially accessed.

❑ Smaller CISIZE recommended for high update catalogs.

Example RLS Configuration



Performance Reports

Performance Reports

- ❑ Reports common to all catalogs (RLS/NONRLS):
 - F CATALOG,REPORT,PERFORMANCE
 - F CATALOG,REPORT,CACHE
 - F CATALOG,REPORT,CATSTATS
 - F CATALOG,REPORT,CATSTATX (new!)
 - SMF 6x Records
- ❑ Reports for RLS data sets (all releases) and catalogs (z/OS 2.1):
 - RMFMON III:
 - RLS Activity by Storage Class
 - RLS Activity by Data Set
 - RLS VSAM LRU Overview
 - D SMS,CFLS
 - SMF 42 Subtypes 15-19

Performance Reports – All Catalogs

❑ F CATALOG,REPORT,PERFORMANCE

04/14/2011 CAS*****

```

* DR      542 00000090  IEC352I CATALOG ADDRESS SPACE MODIFY COMMAND COMPLETED
LR      542 00000090  SYS1  RESPONSES -----
DR      542 00000090  IEC351I CATALOG ADDRESS SPACE MODIFY COMMAND ACTIVE
DR      542 00000090  IEC359I CATALOG PERFORMANCE REPORT
DR      542 00000090  *CAS*****
DR      542 00000090  * Statistics since 3:22:12.68 on 02/22/2011 *
DR      542 00000090  * ----CATALOG EVENT---- --COUNT-- ---AVERAGE--- *
DR      542 00000090  * Entries to Catalog      469,827K   30.211 MSEC *
DR      542 00000090  * BCS ENQ Shr Sys        923,098K   9.354 MSEC *
DR      542 00000090  * BCS ENQ Excl Sys       29,951K   30.340 MSEC *
DR      542 00000090  * BCS DEQ                1199,324K   0.088 MSEC *
DR      542 00000090  * VVDS RESERVE CI        88,395K   0.338 MSEC *
DR      542 00000090  * VVDS DEQ CI           88,395K   0.103 MSEC *
DR      542 00000090  * VVDS RESERVE Shr     1027,641K   0.384 MSEC *
DR      542 00000090  * VVDS RESERVE Excl    23,859K   0.566 MSEC *
DR      542 00000090  * VVDS DEQ            1051,500K   0.079 MSEC *
DR      542 00000090  * SPHERE ENQ Excl Sys   30,047K   0.365 MSEC *
DR      542 00000090  * SPHERE DEQ           30,047K   0.068 MSEC *
DR      542 00000090  * CAXWA ENQ Shr         97         0.005 MSEC *
DR      542 00000090  * CAXWA DEQ            99         0.006 MSEC *
DR      542 00000090  * VDSPM ENQ           1014,647K   0.202 MSEC *
DR      542 00000090  * VDSPM DEQ           1014,647K   0.005 MSEC *
DR      542 00000090  * RPL ENQ              1,146K   0.448 MSEC *
DR      542 00000090  * RPL DEQ              1,146K   0.039 MSEC *
DR      542 00000090  * BCS Get              5494,167K   0.194 MSEC *
DR      542 00000090  * BCS Put               16,626K   1.790 MSEC *
DR      542 00000090  * BCS Erase            14,175K   1.691 MSEC *
DR      542 00000090  * VVDS I/O            1177,563K   0.779 MSEC *
DR      542 00000090  * VLF Delete Major     22,459     0.054 MSEC *
DR      542 00000090  * VLF Create Minor     44,914K   0.506 MSEC *
DR      542 00000090  * VLF Retrieve Minor   486,973K   0.032 MSEC *
DR      542 00000090  * VLF Delete Minor    31,489K   0.290 MSEC *
DR      542 00000090  * VLF Define Major         1         0.308 MSEC *
DR      542 00000090  * VLF Identify        32,859     0.031 MSEC *

```

Performance Reports – All Catalogs

ENQUEUE ACTIVITY

ENQUEUE ACTIVITY

z/OS V1R12 SYSTEM ID P13 DATE 06/16/2010 INTERVAL 06.07.134
 RPT VERSION V1R12 RMF TIME 18.29.00 CYCLE 1.000 SECONDS

ENQUEUE DETAIL ACTIVITY					GRS MODE: STAR		-%QLEN DISTRIBUTION-				AVG Q -		REQUEST TYPE -		TOTAL			
-NAME-	----- CONTENTION TIME -----				-- JOBS AT MAXIMUM CONTENTION--						LNTH -	EXCL- - SHARE-		EVENT				
MAJOR	MIN	MAX	TOT	AVG	----- OWN -----	----- WAIT ----	1	2	3	4+		MIN	MAX	MIN	MAX			
MINOR					TOT	NAME	TOT	NAME										
					SYSNAME	SYSNAME												
SYSZRPLW					1 CATALOG (E)	P13	1 CATALOG (E)	P13	80.0	20.0	0.0	0.0	1.20	1	2	0	0	4
USERCAT.DB2.MTTR	0.000	0.000	0.000	0.000														
SYSZIGGV2 (SYSTEMS)					1 CATALOG(E)	P13	17 CATALOG(S)	P13	9.4	9.6	11.5	69.5	38.59	0	17	0	16	3472
USERCAT.DB2.MTTR	0.000	5.612	321.56	0.092														
	(seconds)																	

Performance Reports – All Catalogs

❑ F CATALOG,REPORT,CACHE

```

16.36.28 SYSTEM1      IEC359I CATALOG CACHE REPORT
*CAS*****
* HIT% -RECORDS- -SEARCHES --FOUND-- -DELETES- -SHR UPD- --PURGE-- *
*
* USERCAT1                                (VLF) *
* 27% 9,832      16,543      4,566      4,268      0          2      *
* USERCAT2                                (ISC) *
* 88% 22         350         310         36         0          4      *
* SYS1.MVSRES.MASTCAT                    (ISC) *
* 83% 3,688      3,633      3,018         9         0          0      *
*CAS*****
    
```

- ❑ Hit% rate should be >20% to compensate for the overhead of using VLF.
- ❑ A high number of PURGES is an indication of unbalanced access to the catalog. The catalog may not be a good candidate for VLF or the ISC.

Performance Reports – All Catalogs

❑ F CATALOG,REPORT,CATSTATS

```
*CAS*****  
*          ADDS  UPDATES   GETS  GETUPD  DELETES  BUFNI  BUFND  STRNO  *  
*                                                                 *  
* USERCAT2                                                                 *  
*          2          0          2          0          0          5          9          2  *  
*CAS*****
```

- ❑ The STRNO represents the number of VSAM catalog requests (RPLs) that can run in parallel.
- ❑ Too low of a STRNO will result in contention on SYSZRPLW.

Performance Reports – All Catalogs

□ F CATALOG,REPORT,CATSTATX(catname*) (new via APAR OA40654 and z/OS 2.1)

*CAS*****

* CATALOG NAME = catname *

* INSERTS (ADDS) = 12,486 *

* UPDATES = 3,758 *

* RETRIEVES = 21,816 *

* RETRIEVES FOR UP = 9,760 *

* ERASES (DELETES) = 3,456 *

* **CA-RECLAIMS = 16** * ← New stat added in APAR OA40654

* **CA-REUSES = 4** * ← New stat added in APAR OA40654

* BUFNI SETTING = 4 *

* BUFND SETTING = 4 *

* STRNO SETTING = 2 *

* **CPU TIME = 30 USEC** * ← Added in z/OS 2.1

* **ELAPSED TIME = 20 MSEC** * ← Added in z/OS 2.1

*CAS*****

Performance Reports – All Catalogs

- **SMF Type 6x:**

- Type 60** - VVR Updated

- Type 61** - ICF Define

- Type 62** - VSAM OPEN

- Type 64** - VSAM CLOSE (new for catalogs in z/OS 1.13)

- Type 65** - ICF Delete

- Type 66** - ICF Alter

Performance Reports – RLS Only

□ F RMF MONITOR III SYSPLEX RLS Activity by Storage Class (sysplex wide)

ERB3BUF RMF V1R8 VSAM RLS Activity - SYSPLEX Line 1 of 14
 Command ==> Scroll ==> CSR

LRU Status : Good
 Contention % : 0.0
 False Cont % : 0.0

StorClas	Access	Resp	-----Read-----				BMF	----- Write	
		Time	Rate	BMF%	CF%	DASD%	Valid%	False	Inv% rate
STORCLAS1									
Below 2G	DIR	0.004	0.18	50.0	0.0	50.0	100	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
Above 2G	DIR	0.000	0.00	00.0	0.0	00.0	0.0	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00

Performance Reports – RLS Only

▣ F RMF MONITOR III SYSPLEX RLS Activity by Data Set (sysplex wide)

ERB3BUF RMF V1R8 VSAM RLS Activity - SYSPLEX Line 1 of 14
 Command ==> Scroll ==> CSR

LRU Status : Good
 Contention % : 0.0
 False Cont % : 0.0

Sphere/DS Access Resp -----Read----- BMF ----- Write
 Time Rate BMF% CF% DASD% Valid% False Inv% rate

USERCAT1

Below 2G	DIR	0.004	0.18	0.0	0.0	100.0	100	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
Above 2G	DIR	0.000	0.00	00.0	0.0	00.0	0.0	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00

USERCAT1.INDEX

Below 2G	DIR	0.004	0.18	100.0	0.0	0.0	100	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
Above 2G	DIR	0.000	0.00	00.0	0.0	00.0	0.0	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00

Performance Reports – RLS Only

□ F RMF MONITOR III SYSPLEX VSAM LRU Overview (sysplex wide)

Samples: 59 Systems: 6 Date: 07/27/06 Time: 12.38.50 Range: 10

MVS System	Avg CPU Time	- Buffer Goal	Size High	- Accel %	Reclaim %	----- Read BMF%	----- CF%	DASD%
------------	--------------	---------------	-----------	-----------	-----------	-----------------	-----------	-------

SYS1

Below 2GB	0.1147	100M	16M	0.0	0.0	50.0	0.0	50.0
Above 2GB	0.112	500M	0M	0.0	0.0	00.0	0.0	00.0

Performance Reports – RLS Only

❑ *D SMS,CFLS (Coupling Facility Lock Structure)*

J80 12152 15:30:56.26 D SMS,CFLS

IGW320I 15:30:56 Display SMS,CFLS(IGWLOCK00)

PRIMARY STRUCTURE:IGWLOCK00 VERSION:C99DC09480021972 SIZE:200704K

RECORD TABLE ENTRIES:434612 USED:984

SECONDARY STRUCTURE:IGWLOCK00 VERSION:C9A3EFBCF6FC3610

SIZE:200704K

RECORD TABLE ENTRIES:434612 USED:984

LOCK STRUCTURE MODE: DUPLEXED STATUS: ENABLE

System	Interval	LockRate	ContRate	FContRate	WaitQLen
J80	1 Minute	1239.1	0.065	0.042	0.33
J80	1 Hour	373.9	0.099	0.038	0.04
J80	8 Hour	-----	-----	-----	-----
J80	1 Day	-----	-----	-----	-----
(13)	1 Minute	344.9	0.029	0.014	0.10
(13)	1 Hour	108.4	0.041	0.015	0.03

Performance Reports – RLS Only

□ SMF 42 Subtypes 15-19:

Subtype 15 - RLS statistics by Storage Class

Subtype 16 - RLS statistics by Data set

Must use V SMS,MONDS(spherename),ON to collect subtype 16 statistics. Must turn on data set collection when using RMF III:

Subtype 17 - RLS locking Statistics for IGWLOCK00

Subtype 18 - RLS caching Statistics

Subtype 19 - BMF statistics

Note: Only one system in the sysplex collects the SMF 42 records. The system collecting the records is displayed in the D SMS,SMSVSAM operator command.

RLS for Catalog Benchmarks

RLS for Catalog - Benchmarks

❑ SYSPLEX Configuration:

- Processor z10 2097-E12
- 3 LPARs in the plex
- 7 CPUs (Online - Dedicated for the test)
- 1 CF (Internal with 1 dedicated CP) with 34GB

❑ Relevant parameters:

- Catalog TASKMAX = 180 (default)
- Catalog CFSIZE(32768) and CFSIZE(4096)
- Catalog STRNO(255)
- RLS Dataclas (RLSABOVETHEBAR(NO)) and RLSCFCACHE(ALL) – (default)
- RLSMAXPOOLSIZE(100M) (default)
- CF cache structure 1Gig

RLS for Catalog - Benchmarks

- ❑ Comparison made:
 - RLS Catalog vs VLF Catalog on z/OS 2.1
- ❑ Five workloads analyzed:
 1. Direct random data set delete
 2. Direct random data set define
 3. Generic (sequential) read
 4. Direct first read
 5. Direct first read sharing system

RLS for Catalog – Benchmarks (cont.)

❑ Workload #1 Data set DELETE:

- Define one usercatalog (CISIZE 32768) and load with 300,000 non-VSAM data sets.
- Each lpar runs 100 jobs (300 total).
- Each job deletes 1000 nonvsam data sets (300,000 total).

❑ Results:

NonRLS Elapse	RLS Elapse	RLS to NonRLS %Delta
80.42 min	8.42 min	89.51

NonRLS CPU	RLS CPU	RLS to NonRLS %Delta
1269.3 sec	298.7 sec	76.46

* CPU measured in the CATALOG, GR5 - SMSVGA

RLS for Catalog – Benchmarks (cont.)

❑ Workload #2 Data set DEFINE:

- Define one usercatalog (CISIZE 32768).
- Each lpar runs 100 jobs (300 total).
- Each job defines 1000 nonvsam data sets (300,000 total).

❑ Results:

NonRLS Elapse	RLS Elapse	RLS to NonRLS %Delta
48.84 min	21.42 min	56.13

NonRLS CPU	RLS CPU	RLS to NonRLS %Delta
685.6 sec	130.8 sec	80.91

* CPU measured in the CATALOG, GRS, SMSVSAM, and XCFAS address spaces.

RLS for Catalog – Benchmarks (cont.)

❑ Workload #3 Sequential (Generic) Read:

- Define one usercatalog (CISIZE 32768) and load with 300,000 non-VSAM data sets.
- Run one job that does a generic search of all 300,000 data sets (i.e. LISTC ENT(HLQ.*) ALL).

❑ Results:

NonRLS Elapse	RLS Elapse	RLS to NonRLS %Delta
7.40 min	5.03 min	32.08

NonRLS CPU	RLS CPU	RLS to NonRLS %Delta
65.1 sec	75.2 sec	-15.52

* CPU measured in the CATALOG, GRS, SMSVSAM, and XCFAS address spaces.

RLS for Catalog – Benchmarks (cont.)

❑ Workload #4 First Direct Read:

- Define one usercatalog (CISIZE 4096) and load with 300,000 non-VSAM data sets.
- Run one job from System1 that does a TSO ALLOCATE for all 300,000 data sets. Buffer pools cleared prior to running the job.

❑ Results:

NonRLS Elapse	RLS Elapse	RLS to NonRLS %Delta
26.77 min	20.33 min	24.1

NonRLS CPU	RLS CPU	RLS to NonRLS %Delta
94.0 sec	109.60 sec	-14.3

* CPU measured in the CATALOG, GRS, SMSVSAM, and XCFAS address spaces.

RLS for Catalog – Benchmarks (cont.)

- ❑ Workload #5 First Direct Read Sharing System:
 - Run one job from System2 that does a TSO ALLOCATE for all 300,000 data sets from same catalog accessed in Workload #4. Buffer pools cleared prior to running the job.
- ❑ Results:

NonRLS Elapse	RLS Elapse	RLS to NonRLS %Delta
26.86 min	20.29 min	24.5

NonRLS CPU	RLS CPU	RLS to NonRLS %Delta
95.0 sec	109.90 sec	-13.5

* CPU measured in the CATALOG, GRS, SMSVSAM, and XCFAS address spaces.

RLS for Catalog – Benchmarks (cont.)

- ❑ Workload #5 First Direct Read Sharing System (Cont.):
 - Results:
 - Cache hits on sharing system reduce elapse time for RLS

SYSTEM	NonRLS Elapse	RLS Elapse	RLS to NonRLS %Delta
SYSTEM1	26.77 min	20.33 min	-24.1
SYSTEM2	26.86 min	20.29 min	-24.5

Summary

- ❑ z/OS 2.1 provides new optional method for improving catalog performance.
- ❑ Catalogs can grow without the need to be split for performance reasons.
- ❑ New sysplex wide control of catalogs to suspend or lock catalogs to improve availability.
- ❑ New vendor interface to quiesce updates or suspend lock/catalogs to improve integrity.
- ❑ New performance measurements at a catalog level.

Questions ???