

VSAM New Features in z/OS 2.2 *(and revisit those from 2.1)*

Neal Bohling

IBM, VSAM RLS Development

Session 17137



#SHAREorg



SHARE is an independent volunteer-run information technology association
that provides **education, professional networking and industry influence.**

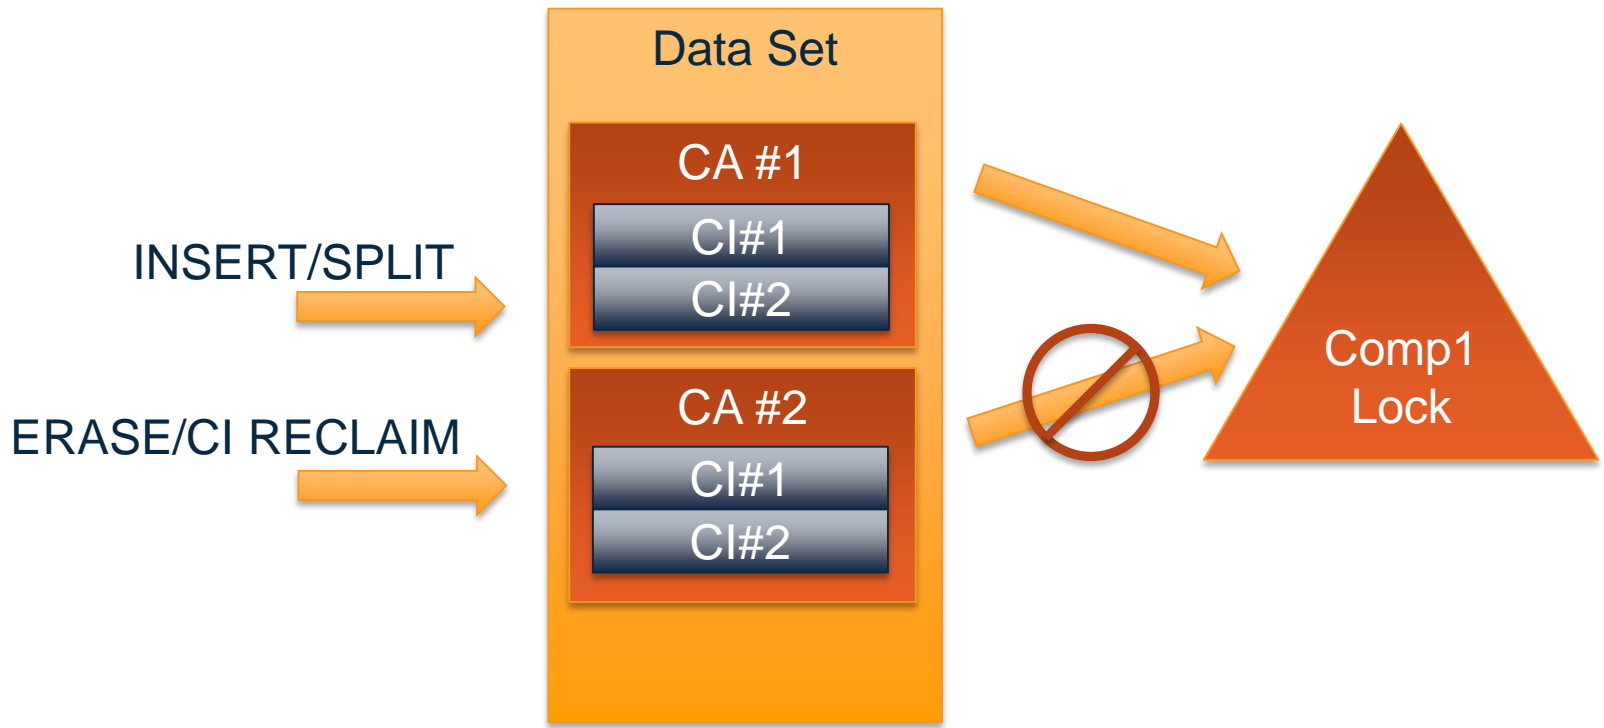


Overview

- z/OS 2.2
 - Index Record Locking (RLS)
 - Primary and Secondary Space Reduction (all)
 - IDCAMS LISTSTAT support (RLS)
 - Chained I/O for Spanned Records (all)
 - LSR Dynamic Buffer Addition (VSAM)
 - Linear DS Constraint Relief (VSAM)
 - Verify Enhancements (VSAM)
 - Other RAS Enhancements (all)

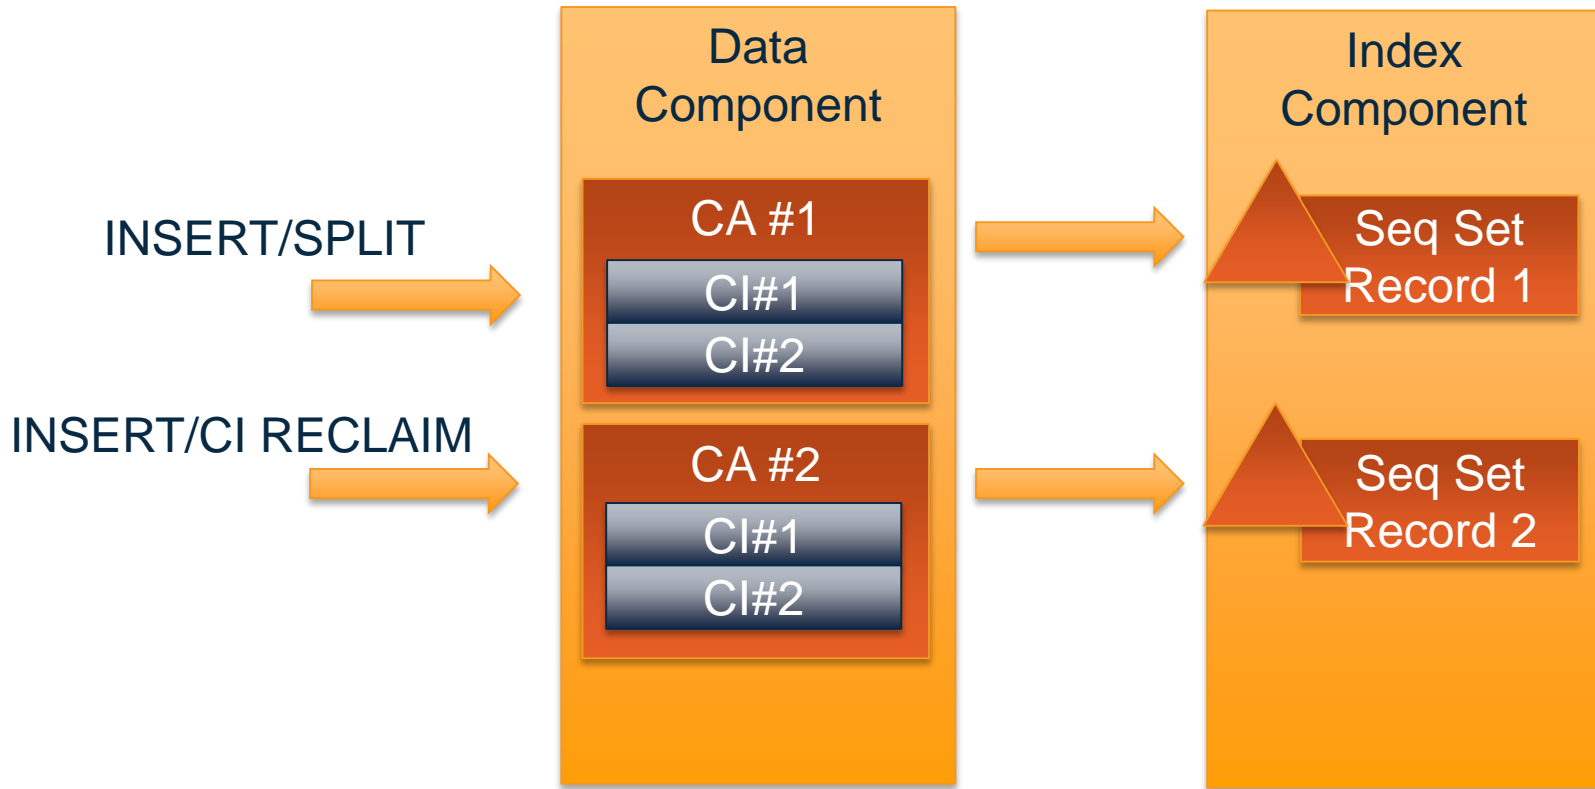
Index Record Locking

Current Split Logic



Only one split per data set can occur simultaneously

With Index Record Locking



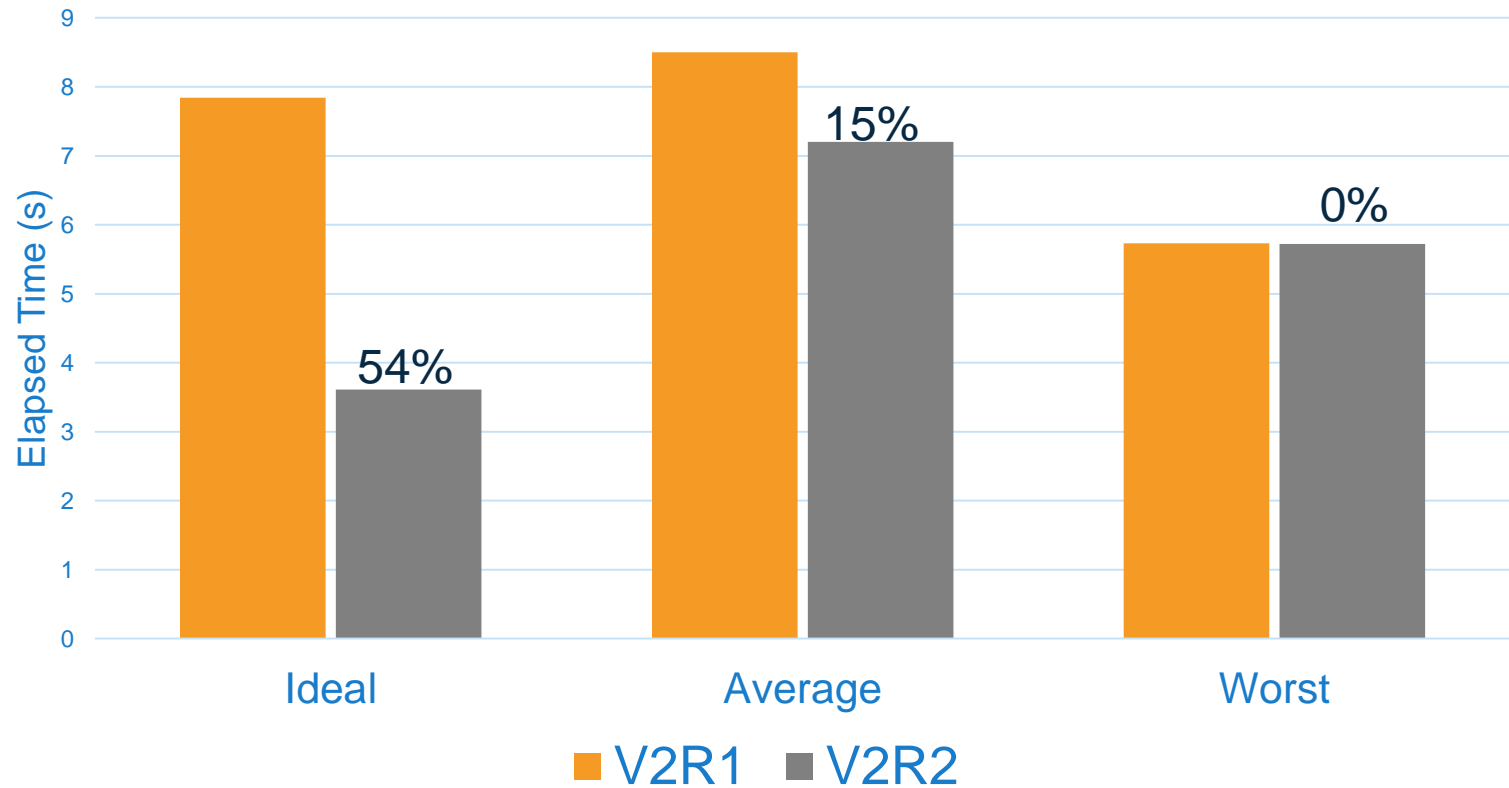
Locks the sequence set Index Record corresponding to the CA

Performance Measurements

- Three workloads on zEC12 / 8 logical CPs, CF 1 CP
- **Ideal:**
 - 30 regions spread across 3 systems
 - Records insert into different CAs throughout the DS
- **Average:**
 - 30 regions across 3 systems
 - Random inserts that may include some CA splits
- **Worst:**
 - Single task on a single system
 - Inserting randomly through the data set

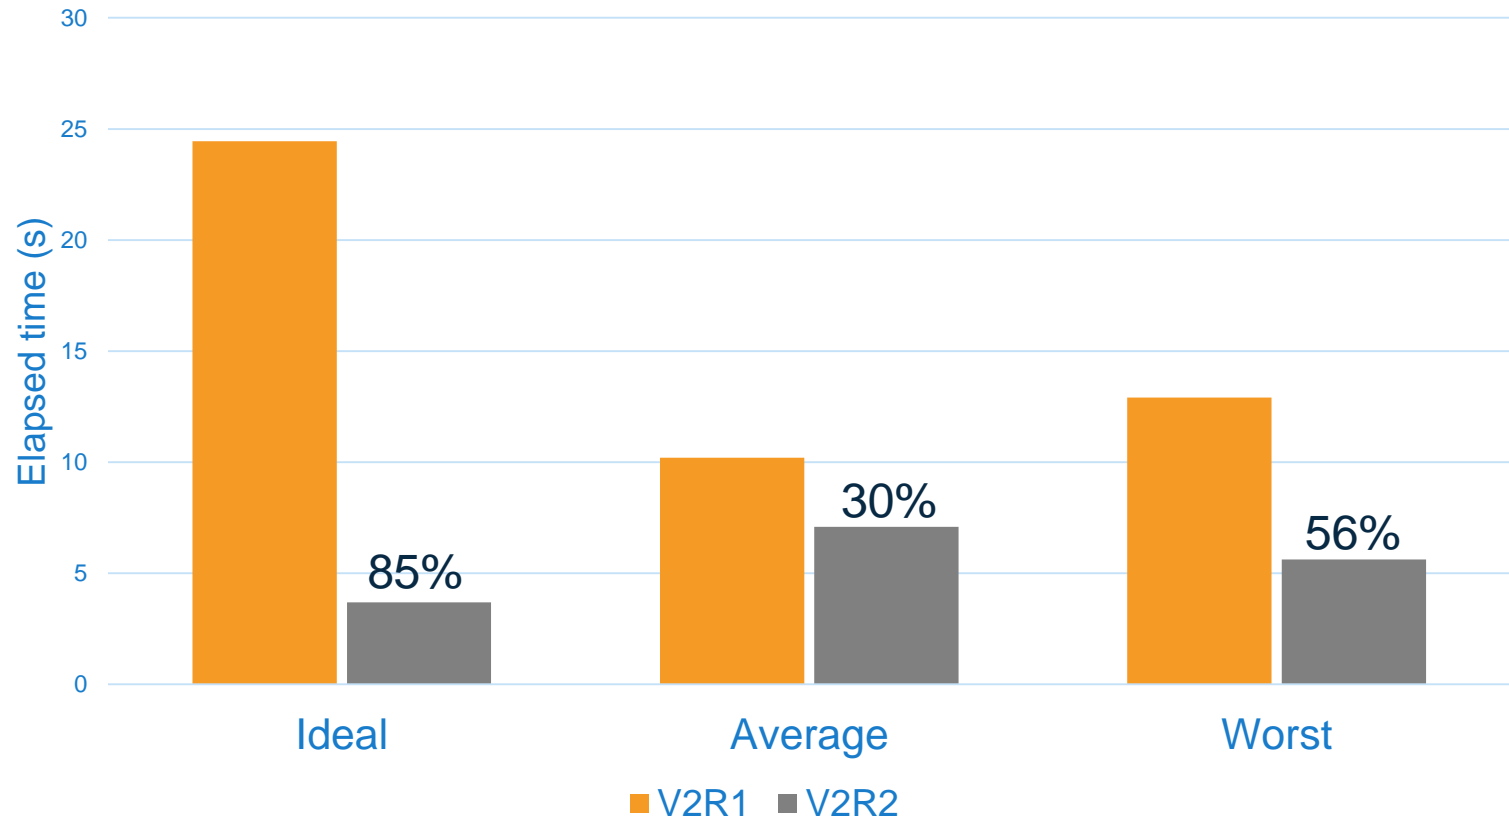
Performance Improvement

Batch RLS Non-BWO Workload Runtime




Performance Improvement

Batch RLS BWO Workload Runtime




Performance Improvement

SMSVSAM CPU (Best Case)



Locking Performance Improvement

Component 1_1 Locking and Contention Across BWO and Non-BWO Batch Workloads



Index Record Locking Other Notes

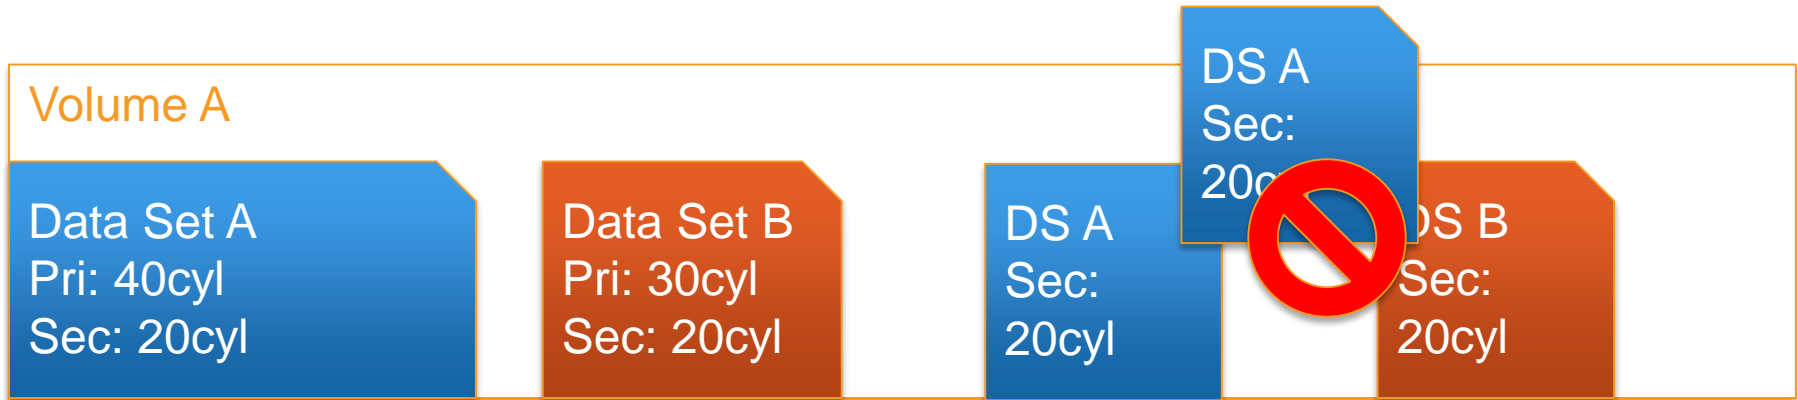
- **Greatly reduces SPLIT pain point**
- Locks at the CA (sequence set index record) level during:
 - CI SPLIT
 - CI RECLAIM
 - Spanned record activity
- CA split / reclaim activity will still use data-set level lock

- New lock, new SMF fields (Component 1 class 4)
 - Obtain (SMF42FPH)
 - True Contention (SMF42FPI) – note, these are cross-plex contention
 - False Contention (SMF42FPJ)
 - Release (SMF42FPK)

- Lower releases will require toleration OA42676

Primary and Secondary Space Reduction Improvement

Current Space Allocation Process



If last volume, extend fails...

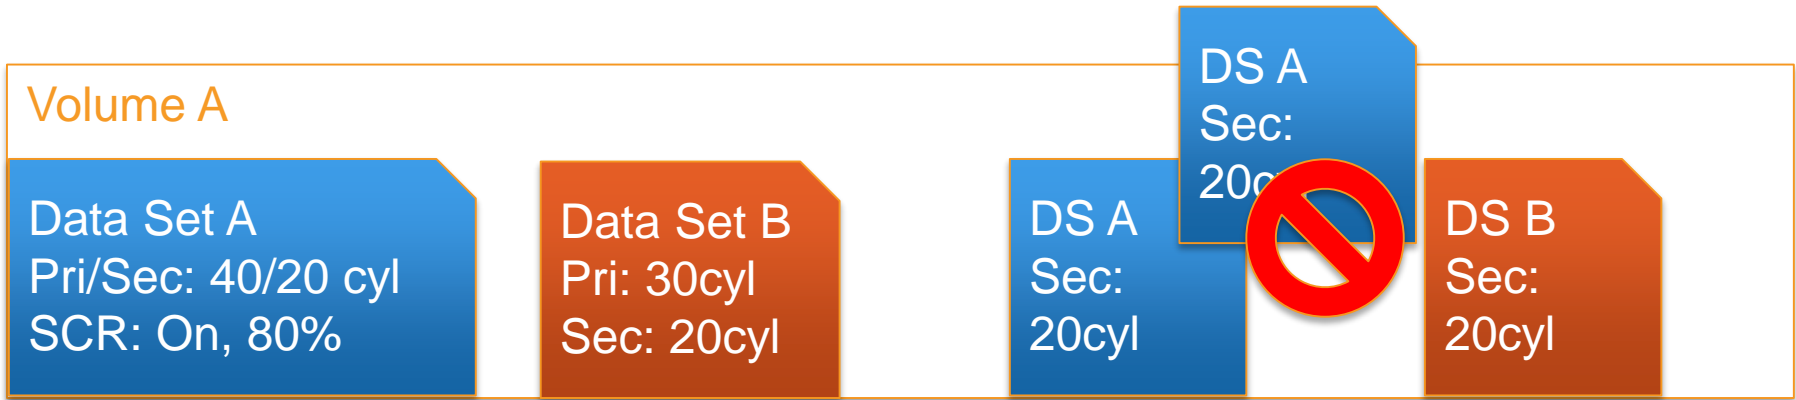
Space allocation processing

- **Old Method:**
 - Space is obtained by data-set defined amounts.
 - For primary, retries once with “best fit”, or fails
 - For secondary, fails request or moves to next volume
- **New Method:**
 - Space is obtained by data-set defined amounts
 - DADSM will return the largest extent that meets SCR requirements
 - ONE CALL – values are passed to DADSM

Fields

- **Data Class fields:**
 - Space Constraint Relief: Y
 - Reduce Space Up to (%): 0-99%
- **Note: Reduce Space Up To is a subtraction**
- Ex: 100cyl primary, 50cyl secondary, SCR set to 80%
 - For primary: $100\text{cyl} * (1-.80) = 20\text{cyl}$
 - For secondary: $50\text{cyl} * (1-.8) = 10\text{cyl}$
- Will return largest available extent that fits SCR threshold
 - For primary, could be 20, 30, 50, or 70 cyl

Updated Space Reduction



Updated Space Reduction

Volume A

Data Set A
Pri/Sec: 40/20 cyl
SCR: On, 80%

Data Set B
Pri: 30cyl
Sec: 20cyl

DS A
Sec:
20cyl

DS A
Sec:
15cyl

DS B
Sec:
20cyl

Space amount is reduced by up to 80% -
In this case, it was reduced by 25% to find 15cyl of space..

Returns LARGEST available extent on volume that matches
SCR setting.

Space Reduction Enhancement

- Works for both Primary and Secondary amounts
- For VSAM, extents must be a multiple of CA size
- VSAM/RLS/PDSE/BAM/SAM all supported
 - Except striped data sets
- Data set must be SMS-managed
- To enable, set Data Class fields:
 - Space Constraint Relief: Y
 - Reduce Space Up to (%): anything > 0
- New SMF fields:
 - SMF64SSR – if secondary space reduction was used
 - SMF64NTA – size of extent returned in Tracks

LISTSTAT – Statistics while VSAM is OPEN

LISTSTAT Command

- New command

IDCAMS SHCDS LISTSTAT('cluster')

- Provides point-in-time plex-wide statistics for open VSAM data sets using RLS
- Subset of LISTCAT and SMF64 information, but does not require CLOSE
- Available via OA42435 for 1.13 and 2.1

LISTSTAT Output

SHCDS LISTSTAT(NB.RLS.TEST2)

LIST STATISTICS (LISTSTAT):

CLUSTER-----NB.RLS.TEST2

DATA-----NB.RLS.TEST2.DATA

TOTAL RECORDS-----	101
RECORDS DELETED-----	0
RECORDS INSERTED-----	1
RECORDS UPDATED-----	0
RECORDS RETRIEVED-----	0
HI-A-RBA-----	829440

INDEX-----NB.RLS.TEST2.INDEX

TOTAL RECORDS-----	1
CA RECLAIMS-----	0
RECLAIMED-CA REUSES---	0
RECORDS UPDATED-----	0
RECORDS RETRIEVED-----	0
HI-A-RBA-----	33792
HI-LEVEL-RBA-----	0

CI SPLITS-----	0
CA SPLITS-----	0
EXCPS-----	207
EXTENTS-----	1
FREE SPACE-----	774144
HI-U-RBA-----	829440

CI SPLITS-----	0
CA SPLITS-----	0
EXCPS-----	209
EXTENTS-----	1
FREE SPACE-----	32768
HI-U-RBA-----	1024
INDEX LEVELS-----	1

Chained I/O for Spanned Records

Chained I/O Simile



Morbi rutrum auctor magna sed sollicitudin. Donec vel est metus. Integer luctus pharetra dignissim. Nulla et nulla mi. Nulla consequat magna



a urna gravida maximus. Ut a tincidunt justo, eu scelerisque lectus. Phasellus obortis urna diam, nec tincidunt lorem faucibus et. Etiam



Vs.





Morbi rutrum auctor magna sed sollicitudin. Donec vel est metus. Integer luctus pharetra dignissim. Nulla et nulla mi. Nulla consequat magna
urna gravida maximus. Ut a tincidunt justo, eu scelerisque lectus.
Phasellus obortis urna diam, nec tincidunt lorem faucibus et. Etiam





Chained I/O

- Old Way – Segment Spanned Records into parts and loop



Chained I/O

- New Way – Combine segments into one I/O

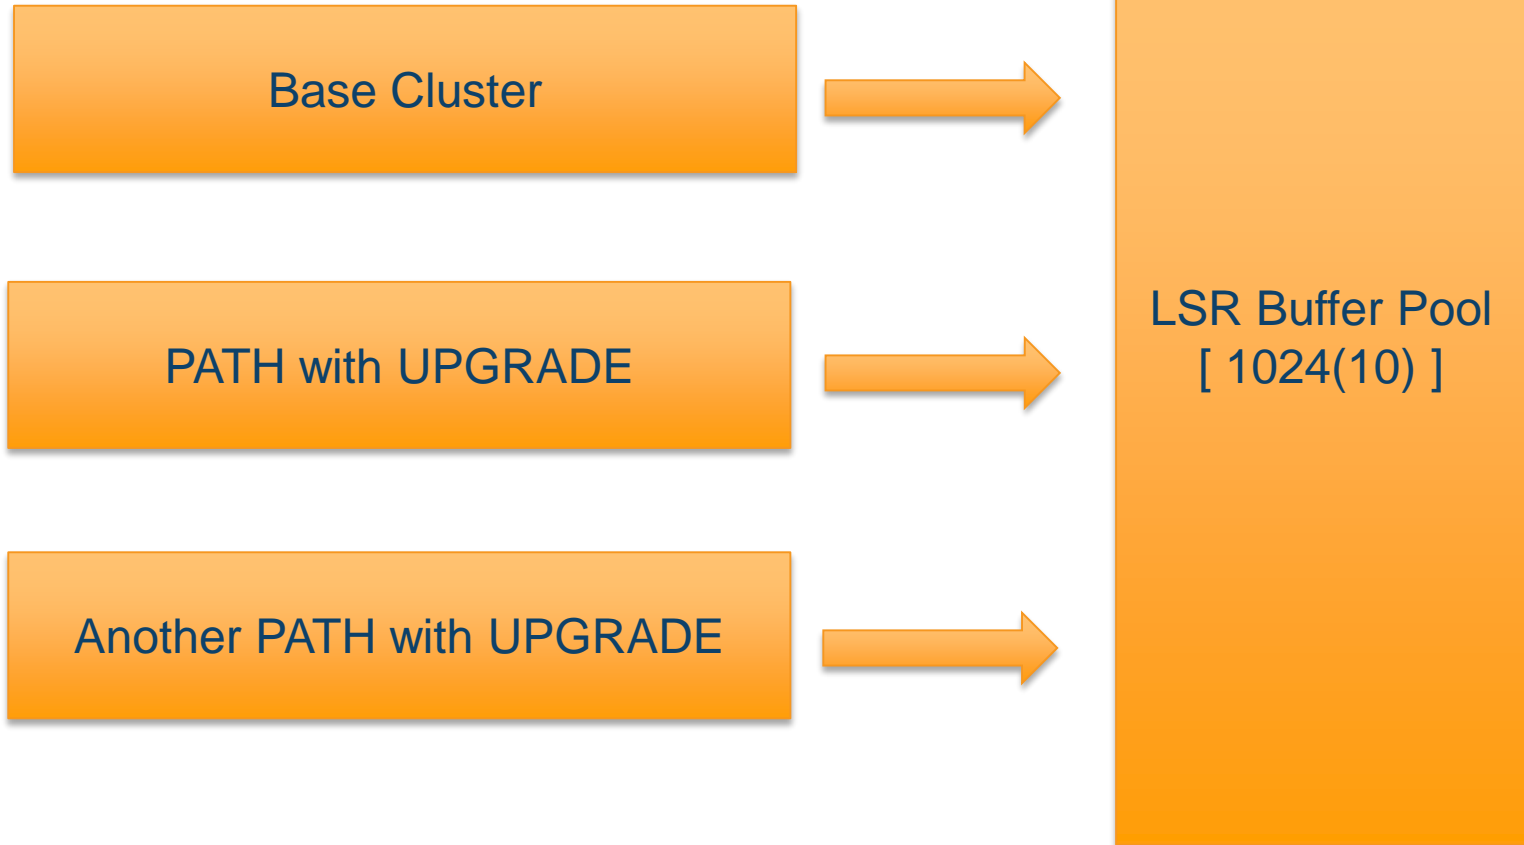


Chained I/O

- VSAM uses CIO for READ, PUT, and ERASE
 - Supports NSR
 - LSR and GSR not currently supported
- RLS uses CIO for PUT, ERASE (not read)
 - Stages updates in buffer/cache
- No co-existence maintenance needed for lower releases
- Benefits:
 - Reduces I/O overhead when using spanned records
 - Avoids the x'8C' that can occur if I/O or system fails, ABEND, or cancel mid-record

Dynamic Buffer Addition

LSR Dynamic Buffer Addition



LSR Dynamic Buffer Addition

- **GOAL:** Avoid failing a request due to buffer shortages
- **SOLUTION:** Dynamically add buffers to LSR pool as needed
- Invoked when LSR processing receives “no buffers available”
- Expands current pool with same settings
- New message: IDA9990I – indicates addition occurred
- NSR/GSR not supported
- Cross-memory mode, SRB mode, and TCB Key 9 not supported

LSR Dynamic Buffer Addition

- New Message:
 - IDA9990I VSAM DBA ADDED xxxx DATA|INDEX BUFFERS of yyyy BYTES EACH TO SHRPOOL zzz BECAUSE THERE WERE INSUFFICIENT BUFFERS TO PROCESS THE REQUEST.
 - RECOMMENDATION: FOR PERFORMANCE, REBUILD THE SHARED POOL WITH AN INCREASE IN SIZE.
- LSR hash built with original pool size, so added buffers may affect performance due to hash conflicts
- Final note:
 - Prevents errors from minor space miscalculations
 - **Not meant to replace well-sized LSR pools**

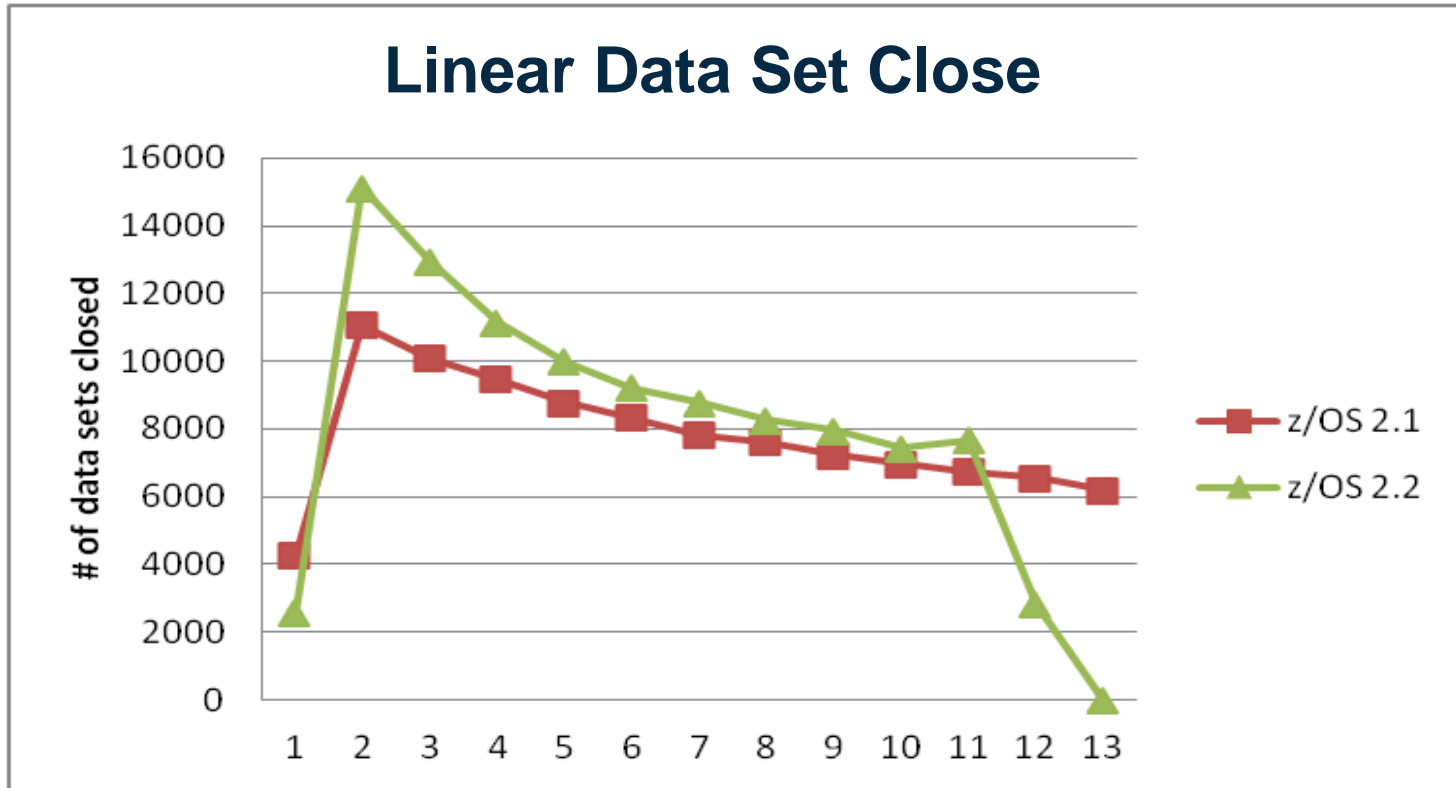
Linear Data Set Constraint Relief

Constraint Relief for Linear VSAM

- **Constraint #1**
 - Below the bar control blocks
 - Limits number of open data sets
- **Solution:**
 - Move some control blocks above the bar
 - Allows for many more open data sets

- **Constraint #2**
 - Large numbers of data sets create large numbers of AMBL blocks
 - CLOSE processing takes time to find the right one
- **Solution:**
 - Use a tree instead of a chain
 - Significant improvements to CLOSE processing time with large numbers of open data sets

Close Speed Improvement



- 13% improvement for 100k data sets (11:20 vs 13m)
- Note – the more data sets, the better the close performance

Verify Enhancements

Verify Enhancements

- IDCAMS VERIFY only fixes a small number of problems
- IDCAMS EXAMINE can find many more
- z/OS 2.2 adds functionality to EXAMINE / VERIFY
- Examine can now pass error information to VERIFY
- VERIFY can now identify errors and try to repair
- Only a few added in 2.2, but framework there to add more.

Verify Enhancements

- **IDCAMS VERIFY**
 - Corrects end of file information (HURBA / VVR / catalog)
 - Repairs behind the scenes if previous close failed
- **IDCAMS VERIFY RECOVER**
 - Completes or backs out any interrupted CA reclaim activity
- **IDCAMS EXAMINE / VERIFY RECOVER (New)**
 - EXAMINE stores information about any problems
 - VERIFY can then read that information and attempt to repair
 - z/OS 2.2 will repair:
 - IDC11718I DATA COMPONENT HIGH-USED RBA IS NOT EQUAL TO CA SIZE
 - IDC11728I DATA FOUND IN EMPTY CI
 - IDC11724I DATA COMPONENT CA NOT KNOWN TO SEQUENCE SET

Other RAS Enhancements

- Additional feedback in message IDA9999I
 - VSAM auto dump may sometime fail, and when it does it issues IDA9999I without helpful information
 - IDA9999I updated to output RPL feedback and JOBNAME:
IDA9999I VSAM AUTO DUMP FAILED TO TAKE A
DUMP FOR RPL FEEDBACK CODE rpl_feedback_code
DUE TO SDUMPX RSN/RC sdump_reason/return_codes
FOR JOB jobname
- Non-SMS EOV Failures – added additional cleanup
- Catalog Statistics update – Permanently records stats in VVR for catalogs

Overview

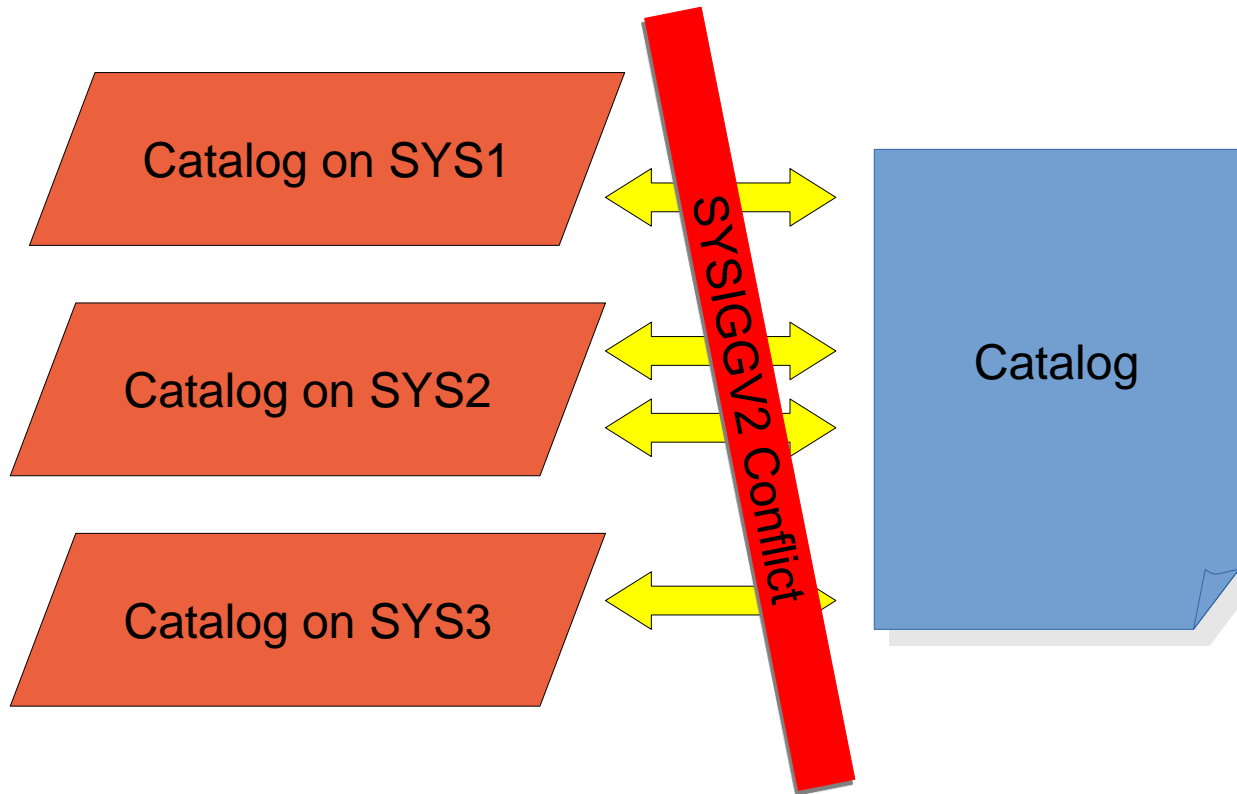
- z/OS 2.1
 - RLS for Catalogs
 - Dynamic Volume Count for RLS
 - Directory Only Caching (RLS)
 - OMEGAMON XE Support (RLS)
 - ACCBIAS in the Data Class (VSAM)
 - Other enhancements

RLS User Catalogs

Current Catalog Limitations

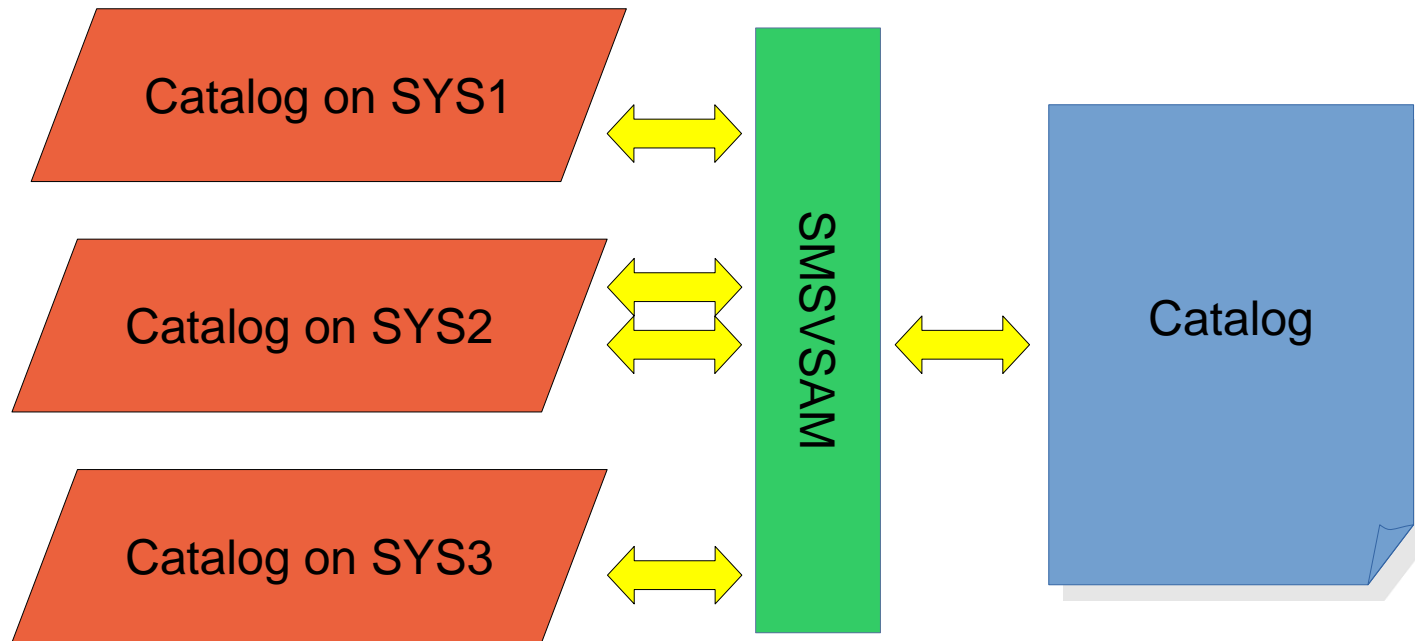
- **Performance**
 - Contention on SYSIGGV2 when updating
 - Limited catalog buffering and buffer invalidation
 - Limited VSAM buffers/strings/storage
- **Availability**
 - Catalogs need to be split to resolve contention
 - Catalogs unavailable during split / recovery
- **Integrity**
 - Catalogs can be damaged by utilities updating while OPEN
 - Lack of SYSPLEX control and serialization
- **Recovery**
 - Long / error prone forward recovery process

Regular Catalog Access



***Potential contention on
SYSIGGV2 'ucat' during updates***

RLS VSAM Access



***SMSVSAM is responsible for serialization.
Serialization is at the RECORD level instead of DS
No more SYSIGGV2 'ucat' ENQ contention***

Improvements RLS Offers

- **Reduced contention**
 - Eliminates SYSIGGV2 'UCAT' ENQ contention
 - Plans to remove the SYSIGGV2 'sphere' ENQ
 - No need to split catalogs to lower contention
- **Higher throughput**
 - Significant improvement in elapsed time & CPU
 - Much shorter wait times
- **Improved control**
 - Suspend / resume ALL catalogs, plex-wide
 - Prevents un-serialized updates

RLS for Catalogs

- **To enable:**
 - IDCAMS ALTER ucat LOG(NONE)
 - F CATALOG,RLSENABLE(ucat)
 - IEC352I MODIFY CATALOG cat.name TO STATE RLSENABLE SUCCESSFUL
 - F CATALOG,RLSQUIESCE(ucat)
- **Notes:**
 - Only available on z/OS 2.1 and up
 - < 1.13 need toleration maintenance
 - IDCAMS tools are updated to use RLS mode (RLSSOURCE)

Performance Benchmark Test

Test	Elapsed Time (min)		CPU* (sec)		Deltas	
	Non-RLS	RLS	Non-RLS	RLS	Elapsed	CPU*
DELETE	80.42	8.42	1269.3	298.7	89.51%	76.46%
DEFINE	48.84	21.42	685.6	130.8	56.13%	80.91%
SEQ READ	7.40	5.03	65.1	75.2	32.08%	-15.52%
DIR READ (first sys)	26.77	20.33	94.0	109.6	24.1%	-14.3%
DIR READ (second sys)	26.86	20.29	95	109.9	24.5%	-13.5%

***CPU in GRS, CATALOG may see a small increase – best to compare per request**

Test environment: Z10 2097 E12, 3 LPARs, 7 CPUs, 1 CF, z/OS 2.1

Catalog parms: TASKMAX=180, CISIZE(32768) and CISIZE(4096), STRNO(255)

RLSABOVETHEBAR(NO) RLSCFCACHE(ALL) RLSMAXPOOLSIZE(100M) CF Cache size 1G

Catalog RLS vs Catalog VLF at z/OS 2.1

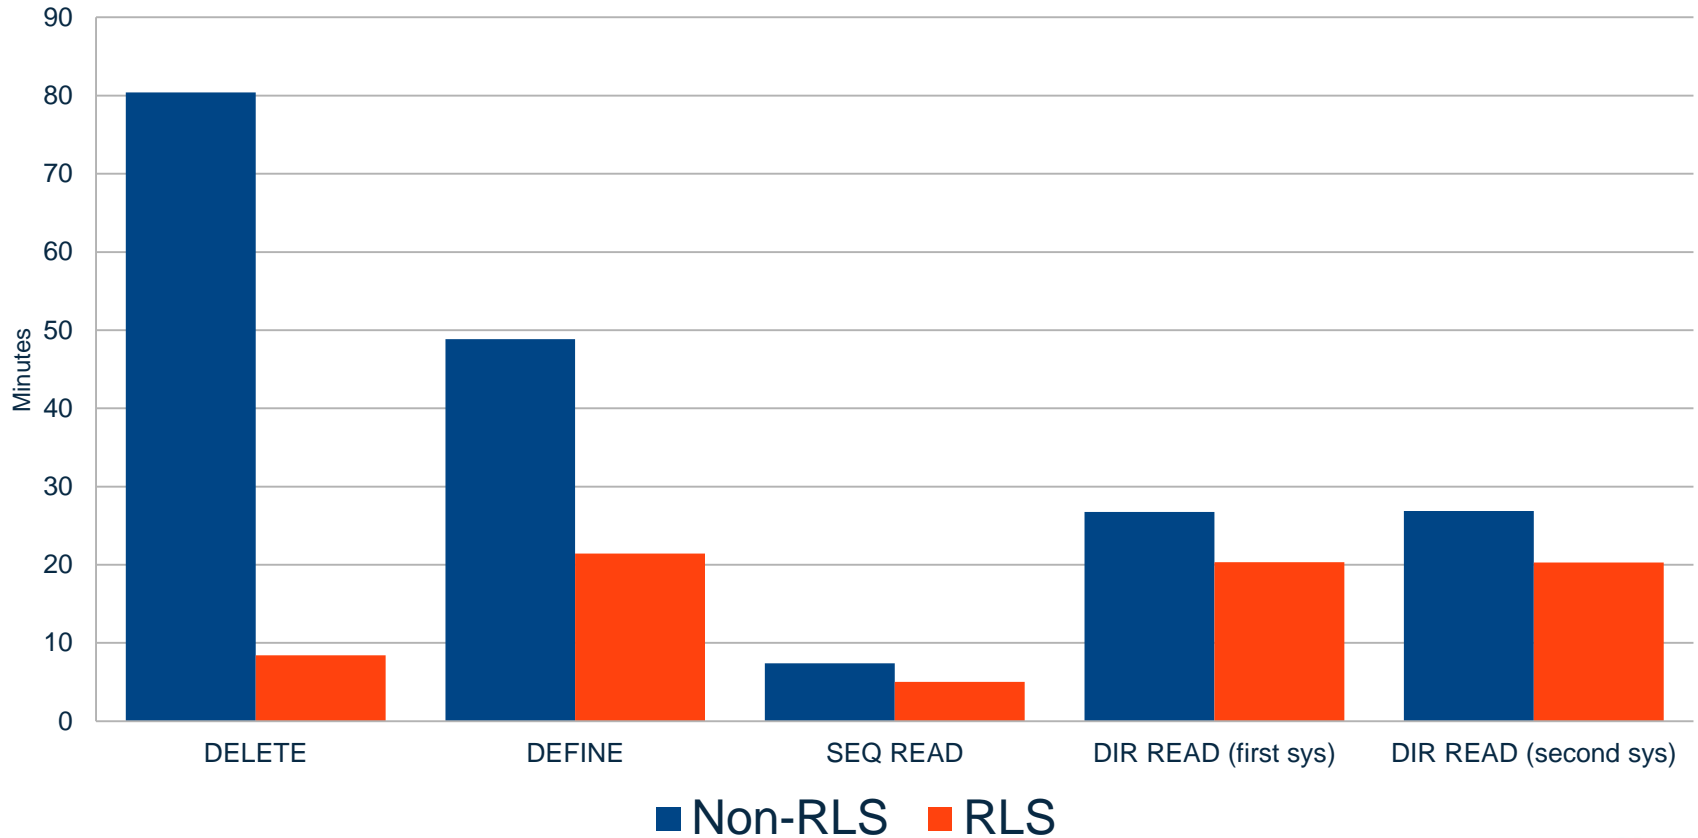
Tests: 300,000 data sets, 100 jobs using 1000 data sets on each LPAR

Source: "Unclog your Systems with z/OS 2.1 – Something New and Exciting for Catalog" by Terri Menendez, IBM

Spring 2013 Session #12977, 12978

Performance Benchmark Test

Elapsed Time RLS vs Non-RLS User Catalog Access



RLS for Catalogs Summary

- **Eliminates bulk of SYSIGGV2 contention**
- **Allows plex-wide serialization at the record level**
- z/OS 2.1 + only
- User catalogs only

- **AMS (IDCAMS) tools support RLS**
 - REPRO, PRINT, IMPORT, EXPORT

Dynamic Volume Count for RLS

Dynamic Volume Count for RLS

- Dynamic Volume Count added to RLS EOV
- During EOV, if out of candidates, but $DVC > VolCNT$, RLS will call ALTER ADDVOL to add volume to catalog
- DVC must be set in Data Class

Space Constraint Relief . . .	N	(Y or N)
Reduce Space Up To (%) . . .	30	(0 to 99 or blank)
Dynamic Volume Count . . .	10	(1 to 59 or blank)

Directory Only Caching

RLS Caching Modes

ALL

Data CIs and Index CIs stored
READ or WRITE will add CI to cache

NONE

Cache on index CIs
READ or WRITE will add CI to cache

**UPDATES
ONLY**

Data CIs and Index CIs stored
Only WRITES will update cache

DIRONLY

No CI data is stored
READ or WRITE will update interest

Inside a Cache

Directory Entry

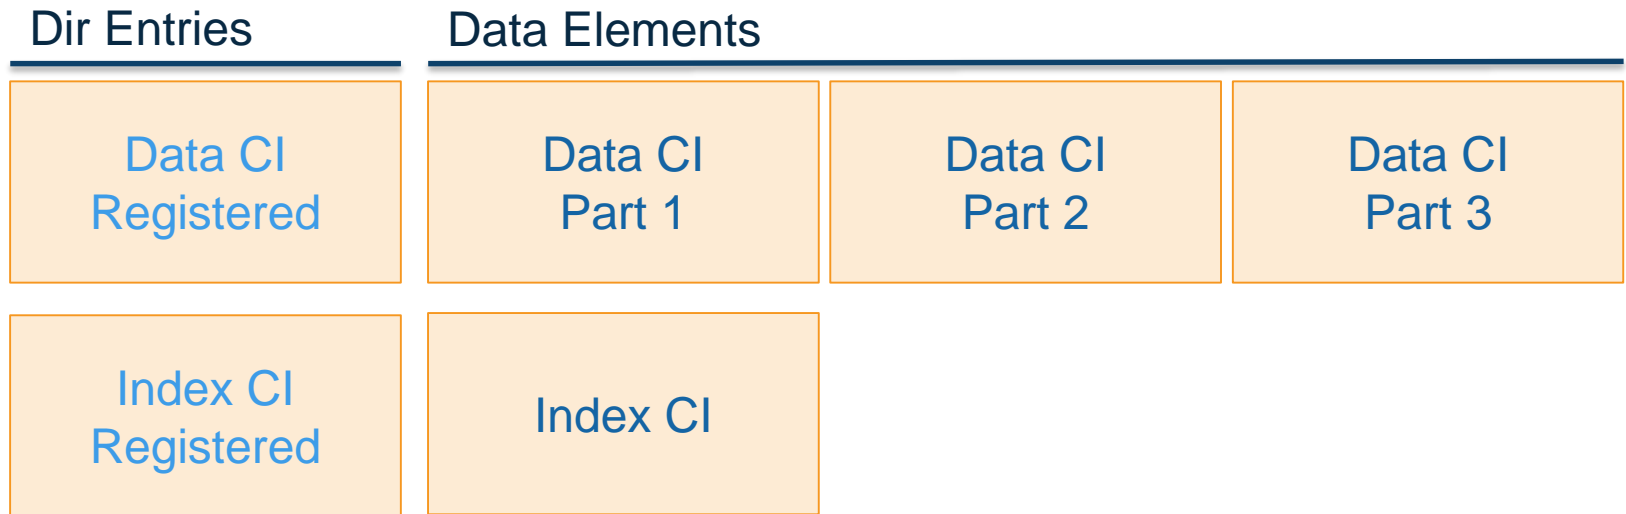
- Holds control information
- One per CI
- Relatively Small

Data Element

- Holds the stored data
- Up to 4k in size
- Many per CI (depending on CISIZE)
- Ex: 16k CI would require 4 DE

Caching Modes

Mode: ALL or UPDATESONLY



Caching Modes

Mode: NONE

Dir Entries

Data Elements

Data CI
Registered

Index CI
Registered

Index CI

Caching Modes

Mode: DIRONLY

Dir Entries

Data CI
Registered

Index CI
Registered

Data Elements

Directory Only Caching

- **Useful for:**

- Data sets used by only one system
- Write-only data sets
- Reduce cache space requirements

- **To Enable:**

- Data Class definition (page 6)

RLS CF Cache Value	A	(A, N, U or D)
RLS Above the 2-GB Bar	. . .	N	(Y or N)

- Requires Rls_MaxCFFeatureLevel(A)
- Toleration: OA36443, OA36415

Directory Only Caching

- New Caching Mode – Directory Only (DIRONLY)
- A true “no caching” option
 - Does not cache any data or index Cis
 - Uses cache to register interest, but stores no data
- Useful for:
 - Data sets that are only used in one system
 - Write-only data sets

OMEGAMON XE Support

OMEGAMON XE Support

- RLS at z/OS 2.1 Provides interfaces for OMEGAMON
- OMEGAMON XE v520 offers RLS Panels
 - 15 new TEP workspaces
 - Many other changes to integrate RLS information
- Monitors all the same info as SMF42

- Requirements:
 - Omegamon XE V520 or higher
 - Maintenance: OA41786, OA42288, OA42798, OA43380, OA43381, OA43376, OA45578, OA44589
- See Share Pittsburg session #15548

Example Display

RLS Summary - dem21lnx.democentral.ibm.com - Vickie Dault

File Edit View Help

Navigator View: Physical

- DFSMShsm Status MVSA
- Tape Management Status
- Record Level Sharing**
- Dataset Attributes System Summary
- Dataset Group Summary
- SMS Configuration
- Storage Toolkit
- Copy Services
- System Automation for z/OS
- Tivoli Decision Support for z/OS
- WebSphere Agent
- z/OS Management Console

Storage Dashboard Physical

RLS Lock Structure CF Details

Lock Name	Sysplex Name	Entries Used Pct	Total Entries	Used Entries
IGWLOCK00	DEMOPLX	0.0	3593	2

RLS Overview

Lock Name	System Sysplex Name	Systems Reporting	Interval	Lock Contention Pct	Lock Contention Min	Lock Contention Max	Lock Contention Avg	Lock Contention Std Dev
IGWLOCK00	DEMOPLX	3	1 day	0.000	0.000	0.001	Sysplex	
IGWLOCK00	DEMOPLX	3	8 hours	0.330	0.000	0.009	Sysplex	
IGWLOCK00	DEMOPLX	3	1 minute	1.650	0.000	0.555	Sysplex	

Lock Structure Summary

Lock Table Name	System Sysplex Name	DIWA Lock Requests	DIWA Lock Requests per Minute	DIWA Lock True Contention Pct	DIWA Lock True Contention Min	ATE Lock False Cont Pct	ATE Lock False Cont Per Minute	Excp Path True Contention Pct	Excp Path False Cont Lock Req Per Minute	Main Path True Contention Pct	Main Path True Cont Lock Req Per Minute	Rec Lock True Contention Pct	Rec Lock Req True Con Minute	Upgrade Locks True Cont Pct	Upp Locks True Cont Per Minute	DIWA Lock False Contention	DIWA Lock False Contention Minute	Main Path Lock R
IGWLOCK00	DEMOPLX	0	0.00	0.0	0.00	0.0	0.00	0.0	0.00	0.0	0.00	0.0	0.00	0.0	0.00	0.0	0.00	

Buffer LRU Summary

Location	System Sysplex Name	BMF Panic Mode Pct	BMF Panic Mode	BMF Accelerated Mode Pct	BMF Accelerated Mode	XCF Castout Lock Cont Retry Pct	XCF Castout Lock Retries	Current BMF Read Hit Pct	Min BMF Read Hit Pct	Cur Lock
Above the bar	DEMOPLX	0.0	0	0.0	0	0.0	0	0.0	0.0	
Below the bar	DEMOPLX	0.0	0	0.0	0	15.6	8	0.0	0.0	

Storage Class Summary

Storage Class	System Sysplex Name	Average Response Time	DIWA Lock Requests	DIWA Lock Requests per Minute	DIWA Lock True Contention Pct	DIWA Lock True Contention Min	DIWA Lock Percent	BMF Requests	BMF Requests per Minute	BMF False Invalid percent	BMF False Invalids	BMF False Invalids per Minute	Lock Requests	Lock Requests per Minute	Lock Contention Percent	False Lock Contention Pct	True Lock Contention Pct	Lock Req True Contention Min	Direct Access Requests Total	Direct Request
RLSC	DEMOPLX	1.47	0	0.00	0.0	0.00	0.0	3706	1,140.30	0.0	2	0.61	891	274.15	0.0	0.0	0.0	0.00	2742	

Hub Time: Fri, 08/01/2014 07:27 PM Server Available RLS Summary - dem21lnx.democentral.ibm.com - Vickie Dault

ACCBIAS in Data Class

ACCBIAS

- ACCBIAS and RMODE31 added to data class
- Previously, these were available on JCL, but not SMS

```

                                DATA CLASS ALTER                                Page 2 of 6
Command ===>

SCDS Name . . . : SYS1.SMS.V2R1.SCDS
Data Class Name : DCRLSNC

To ALTER Data Class, Specify:

Data Set Name Type . . . . . EXT      (EXT, HFS, LIB, PDS, Large or blank)
  If Ext . . . . . R                    (P, R or blank)
  Extended Addressability . . N        (Y or N)
  Record Access Bias . . . . . S      (S, U, DO, DW, SO, SW or blank)
  RMODE31 . . . . . ALL             (ALL, BUFF, CB, NONE or blank)
  Space Constraint Relief . . . Y      (Y or N)
  Reduce Space Up To (%) . . 0        (0 to 99 or blank)
  Dynamic Volume Count . . . 20      (1 to 59 or blank)
  System Managed Buffering . .                (1K to 2048M or blank)

```

Other Changes

- SHOWCB macro updated with new keywords:
 - BUFNOL – # of buffers allocated to data set (LSR or SMB)
 - BUFUSE – # of buffers in use
- LOGREPLICATE – keyword added to IDCAMS
 - Specifies whether VSAM is eligible for replication
- VSAM RAS Enhancements
- RLS 64-bit buffering enhancement
 - Moved some control information above the bar
- RLS Serialization change – moving toward GRS latches

VSAM New Features in z/OS 2.2 *(and revisit those from 2.1)*

Neal Bohling

IBM, VSAM RLS Development

Session 17137



#SHAREorg



SHARE is an independent volunteer-run information technology association
that provides **education, professional networking and industry influence.**

