SHARE / Boston
VSAM RLS Best Practices
Session 8062

Douglas Lehr / IBM Corporation
dllehr@us.ibm.com

Agenda

- **A Background into VSAM RLS Processing**
  - VSAM RLS Components associated with I/O Requests
  - A Day in the Life of a GET Request
- **Tuning the VSAM RLS Structures**
  - RLS Performance Areas
  - Local Buffer Pool sizes
  - Cache Structure sizes
  - Lock Structure sizes
- **Parameters related to VSAM RLS performance and tuning**
  - Data Set Level
  - Request Level

Agenda

- **Measurements related to performance and tuning**
  - SMF 64 Records
  - SMF 42 Subtypes 15-19
- **Example RMF Reports**
  - RLSSC – VSAM RLS activity by storage class.
  - RLSDS – VSAM RLS activity by data set.
  - RLSLRU - VSAM LRU overview
- **SMSVSAM Diagnostics**
  - SMSVSAM Display Commands
  - SMSVSAM Dump Collection.

VSAM RLS "I/O" Path - Components

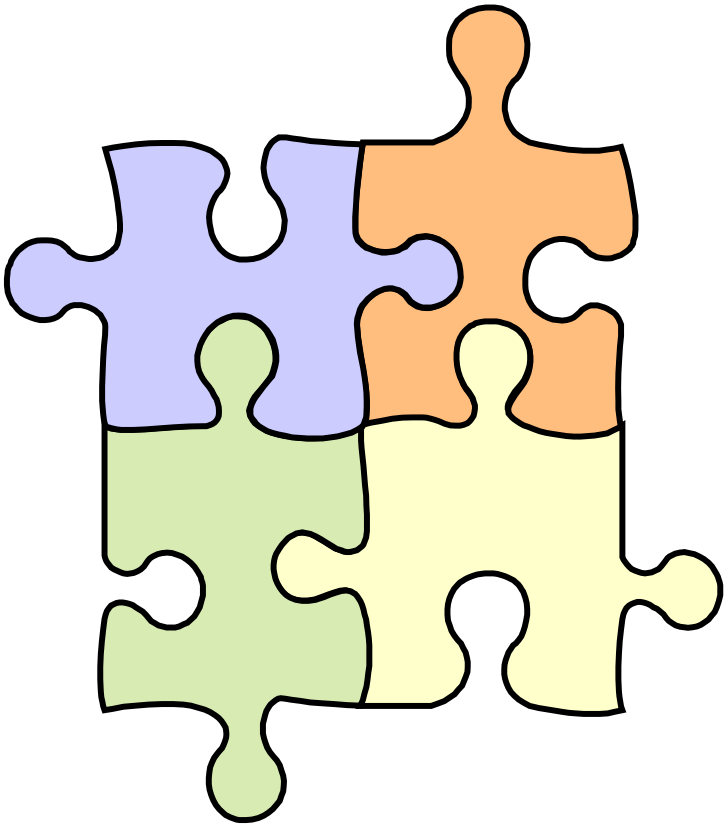

- **VSAM Record Management (VRM)**
  - Provides the VSAM interfaces: GET, PUT, POINT, ERASE, etc.
  - Parameters passed to VRM are through the RPL control block.
- **Storage Management Locking Services (SMLS)**
  - Interfaces with VRM and XCF locking services to obtain, release, and alter locks in the coupling facility lock structure (IGWLOCK00).
- **Ses Cache Manager (SCM)**
  - Interfaces with BMF and XCF caching services to obtain directory elements and read/write data elements to the coupling facility cache structures.
- **Buffer Manager Facility (BMF)**
  - Interfaces with VRM and SCM to locate/add buffers to the local buffer pool.

VSAM RLS "I/O" Path - Components

- **Performance Goal: Spend the least amount of time in the I/O path!**

GET/PUT ↔ VRM ↔ SMLS ↔ XCF Locking services

↔ BMF ↔ SCM ↔ XCF Caching Services

↔ Media Mgr Services (to DASD)

Example of a Get Path

RLS Client AddressSpace

```
OPEN ACB MACRF=RLS,
         RLSREAD=CR
GET Dir,Asy  Key1
  (VRM)
    PC
```

Coupling Facility

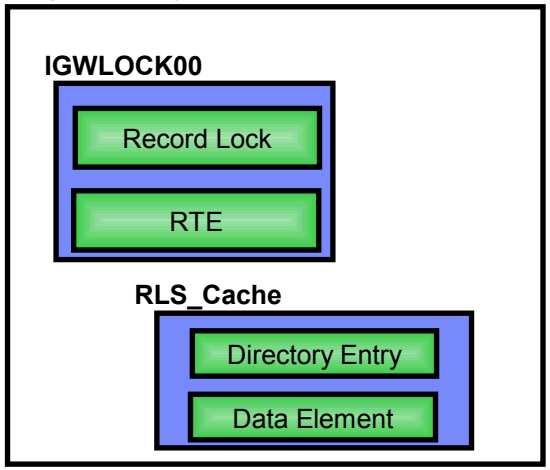

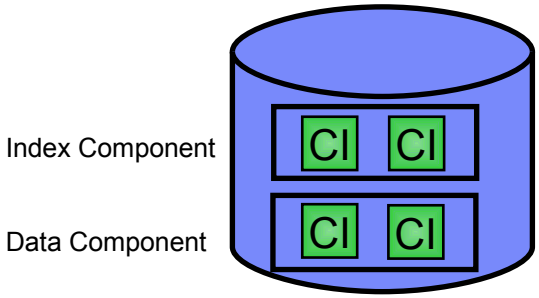

SMSVSAM Address Space

VSAM/RLS Record Management

1. **Index_search**:
   (Call **BMF** to locate Index CIs, if no_buffer Call **SCM** to read from CF or DASD)

2. **Lock_Record;**
   (Call **SMLS** to obtain record lock)

3. **Get_Data_CI:**
   (Call **BMF** to locate Data CI, If no_buffer Call **SCM** to read from CF or DASD

4. **UnLock_Record:**
   (Call **SMLS** to release record lock)

RLSAboveTheBarPool

2,000,000M

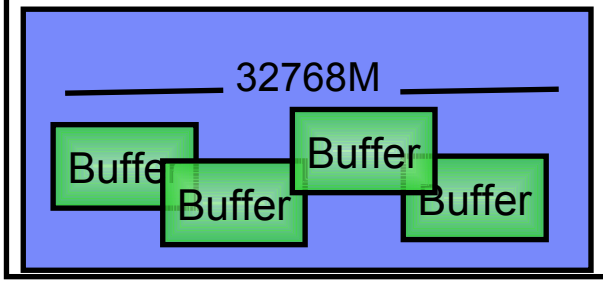

Dataspace

2,000M

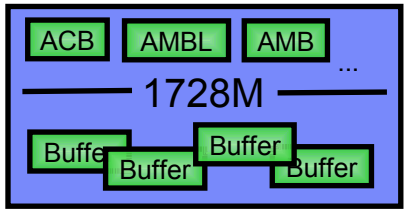

VSAM RLS Performance Areas

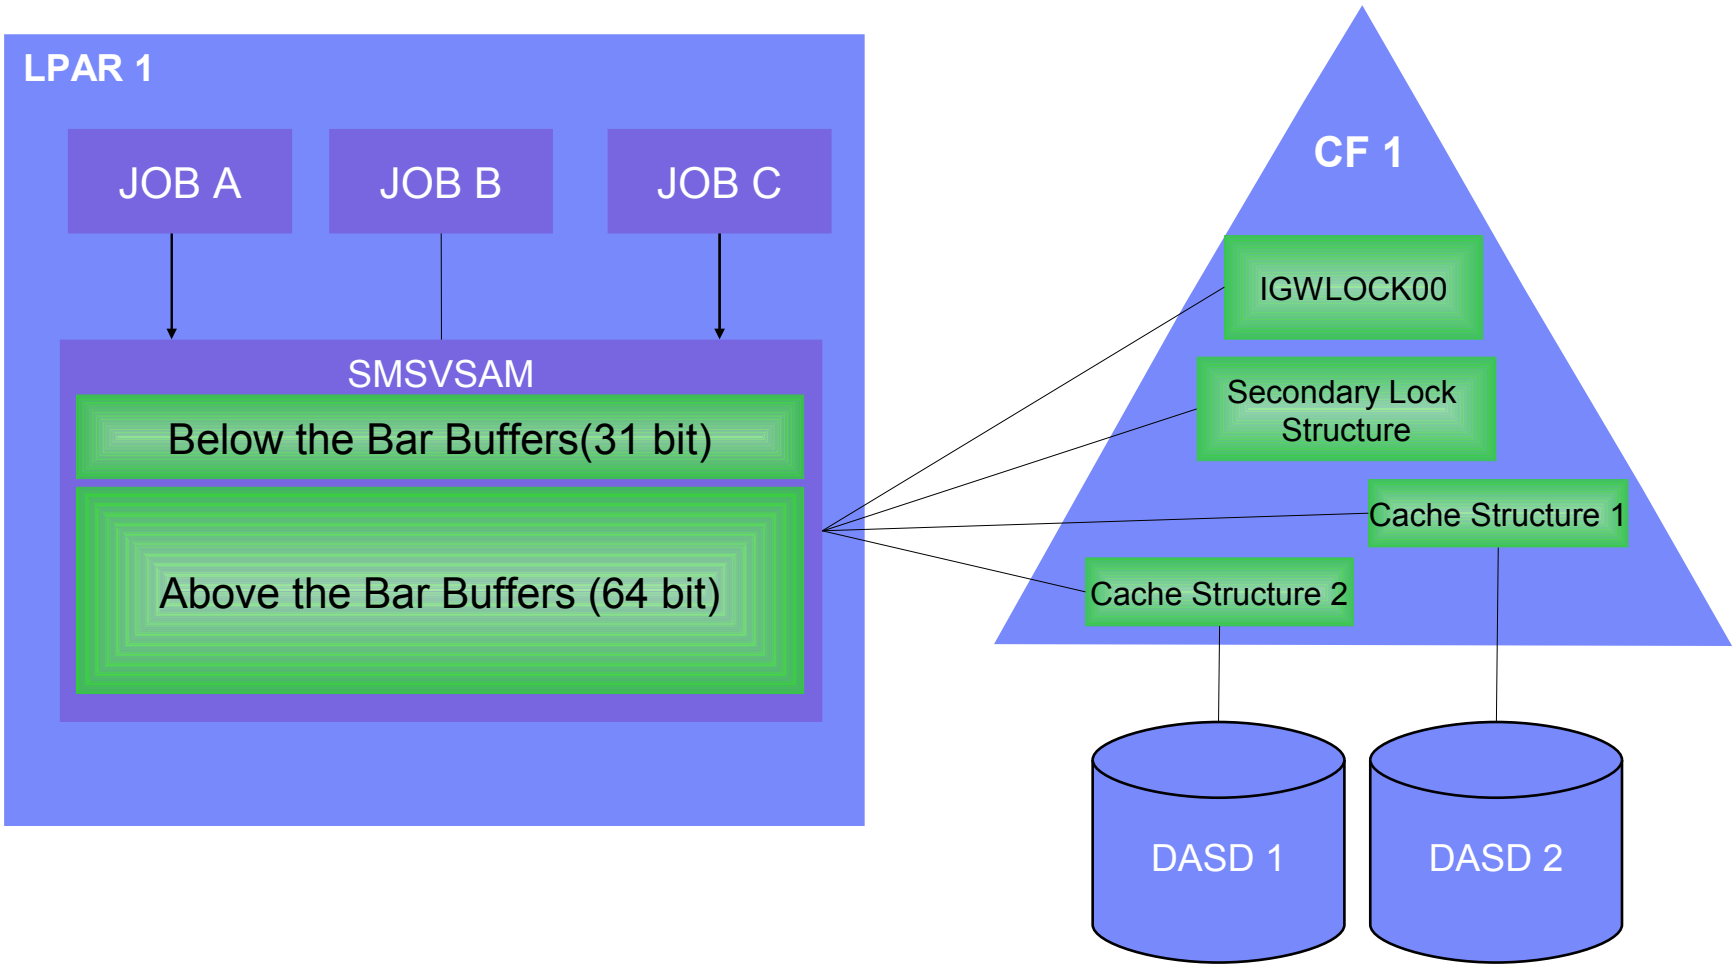

Local Buffer Pools

- 31 Bit Buffer Pool sizes - 10M to 1728M.
  - Defined using the RLS_MAX_POOL_SIZE(nnnn) Parameter
  - Same on all systems.
- 64 Bit Buffer Pool sizes - 0, or a value between 500M and 2,000,000M
  - Defined using the RLSAboveTheBarMaxPoolSize(sysname1,nnnn) Parameter
  - Can be specified per system.
- Pool Size values are a goal for which the LRU tries to maintain. If more buffers are required at any given time, the pool may temporarily exceed the values set.
- Total size of buffer pools should not exceed amount of real storage. A paged out buffer is immediately freed by the LRU.

VSAM RLS Buffer LRU Below the Bar(31 bit)

- **LRU Comprised of 4 different levels**
  - **Normal Mode (0-80%)**
    - Buffer_Limit = 240
    - Buffer_Count = Buffer_Count + 1
    - Cycle stays at 15 seconds
  - **Maintenance Mode (80-120%)**
    - Buffer_Limit = Buffer_Limit - 1
    - Buffer_Count = Buffer_Count + 1
    - Cycle stays at 15 seconds
  - **Accelerated Mode (120-200%)**
    - Buffer_Limit = Buffer_Limit - 4
    - Buffer_Count = Buffer_Count + 1
    - Cycle stays at 15 seconds
    - Buffer stealing will be attempted.
  - **Panic Mode (> 200%)**
    - Buffer_Limit = Buffer_Limit - 8
    - Buffer_Count = Buffer_Count + 1
    - Cycle reduces to run every 5 seconds
    - Buffer stealing will be attempted.

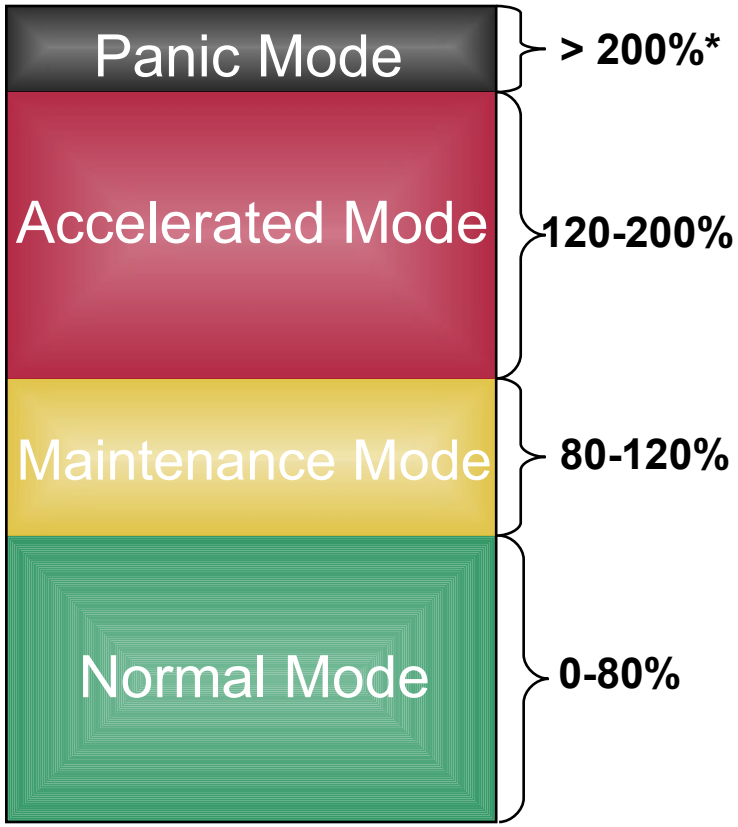

Percentage of RLS_MAX_POOL_SIZE in use

* *200% of RLS_MAX_POOL_SIZE OR > 1728MB Limit*

VSAM RLS Buffer LRU Above the Bar(64 bit)

- **LRU Comprised of 4 different modes**
  - **Normal Mode (0-80%)**
    - Buffers older than 60 minutes will be tossed.
  - **Maintenance Mode(80-90%)**
    - Buffers older than 60 minutes will be tossed.
  - **Accelerated Mode(90-100%)**
    - Buffers older than 30 minutes will be tossed.
    - Buffer stealing will be attempted. If no available buffers a new one will be generated.
  - **Panic Mode(>100%)**
    - Buffers older than 5 minutes will be tossed.
    - Buffer stealing will be attempted. If no available buffers a new one will be generated.

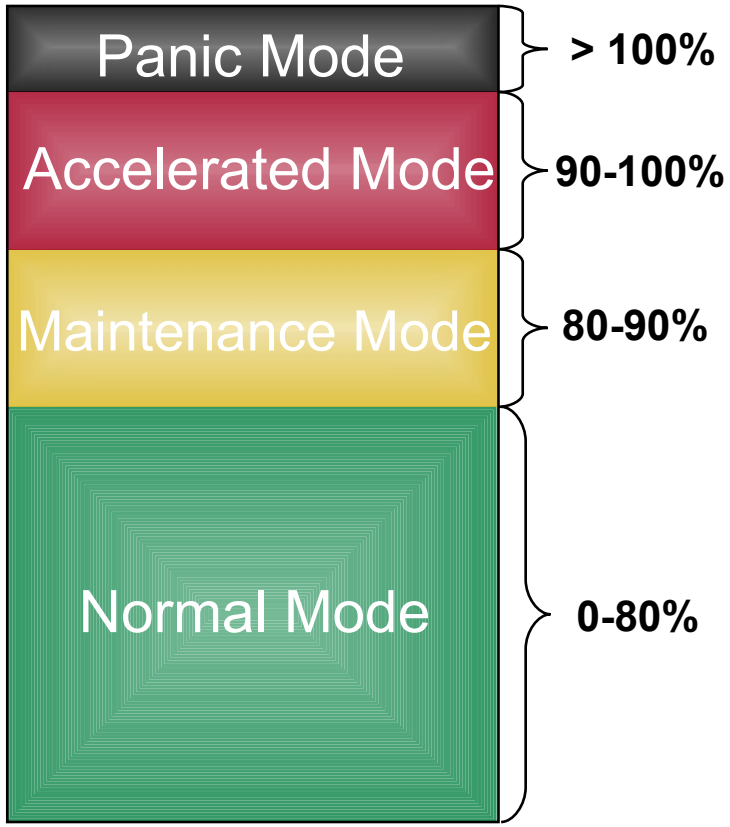

Percentage of RLSAboveTheBarMaxPoolSize in use

Recommendations for the Local Buffer Pool Sizes:

- 31 Bit Buffer Pool (<850):
  - Allows for 680MB (80%) of buffers to reside in the 31 bit pool for one hour. Allows for a doubling of the pool to 1700M before panic mode sets in by exceeding the 1728M limit.
  - Must have adequate cache structure sizes.
- 64 Bit Buffer Pool (<32768):
  - Must have matching amount of real storage.
  - Must have adequate cache structure sizes !!!!
  - Recommended when LRU for the 31 bit pool is frequently in accelerated/panic mode.
  - Great for applications that need to traverse large amounts of data repeatedly.

Sizing the RLS Cache Structures

- The “ideal” cache structure size:
  - Total_Cache_Structure_sizes = Sum of Buffer Goals
    - 31 Bit buffers = ((RLS_Max_Pool_Size) * Number of Systems)
    - 64 Bit buffers = (RLSAboveTheBarMaxPoolSize(system1) + … +RLSAboveTheBarMaxPoolSize(systemn)
  - Assumes the following:
    - RLS_MaxCFFeaturelevel(A) - caching all data
    - No sharing of data across the sysplex.
    - If more than one cache structure to be allocated, Data sets are “evenly” distributed (size, number, amount of data accessed) between the individual cache structures.

Sizing the RLS Cache Structures

- Example:
  - RLS_Max_Pool_Size(850)
  - Number of Systems = 2
  - RLSAboveTheBarMaxPoolSize(System1,2048)
  - RLSAboveTheBarMaxPoolSize(System2,4096)
    - Cache_Structure_Size = (850*2) + 2048 + 4096 = 7844M
  - Cache structure sizes less than the ideal amount should be closely monitored for directory reclaims.

Optimizing Cache Usage:

- RLS_MaxFeatureLevel(Z):
  - RLS will cache CIs less than 4096 only. Saves space in the RLS CF cache structures by not caching large CI sizes.
  - Advantage if data is read only and remains valid in the local buffer pool.
- RLS_MaxFeatureLevel(A):
  - RLS will cache CIs up to 32K.
  - Requires more space in the RLS CF cache structures.
  - Advantage when shared data is updated across the sysplex.
  - Allows use RLS CF Cache feature in Data Class

False Invalids

- When a cache structure is too small for current buffer configuration
  - Cache structure will invalidate a buffer entry to make room for new requests.
  - False Invalids cause additional requests to I/O!
  - Tracked in
    - RMF panels RLSSC, and RLSDS
    - SMF 42 subtypes 15-16

Final Thoughts on Cache vs. Buffer Sizing.

- A balance between Buffer and Cache Sizes must be made for optimum performance
- Buffer = LARGE and Cache = small
  - Bottleneck in cache will occur between buffers and DASD
  - Cache will invalidate buffer entries to make room for new requests, creating thrashing.
  - False Invalids occur when the cache structure needs more space than it has to satisfy CI requests.
- Buffer = small and Cache = LARGE
  - Cache will try to keep as much data in it as possible
  - Buffer LRU's will be the constraining factor, easier to address

Sizing the RLS Lock Structures

- Lock_Structure_Size = 10M * number_of_Systems * Lock_entry_Size
  - Lock_entry_Size (depends on the CFRM MAXSYSTEM value):
    - MAXSYSTEM <= 7 Lock_entry_size = 2
    - MAXSYSTEM >= 8 & <24 Lock_entry_size = 4
    - MAXSYSTEM >=24 & <=32 Lock_entry_size = 8
- Example: MAXSYSTEM = 16 and 8 systems in sysplex
    - IGWLOCK00 = 10M * 4 * 8 = 320M
- Small lock structures result in increased false contention rates. Contention (true or false) result in asynchronous lock requests.
- Refer to RMF CF Activity Report for IGWLOCK00 or D SMS,CFLS command for contention rates.
- Recommended false contention rate is <.5%

DataSet Level Parameters:

- DATACLAS:
  - RLSCFCACHE (ALL/UDATES/NONE)
    - ALL - (default) cache data and index CIs
    - Updates – Cache CIs for write requests only.
    - None - Cache index CIs only.
  - RLSAboveTheBar(YES/NO)
    - Must also specify a non-zero RLSAboveTheBarMaxPoolSize before 64 bit buffering will occur.

Which One Do I Choose?

- RLSAboveTheBar(NO)
  - Recommended for heavy insert and update datasets.
- RLSAboveTheBar(YES)
  - Recommended for heavy read data sets, where data is accessed multiple times within a one hour timeframe.
  - The current design of 64 bit buffering uses 10-20% more CPU for equivalent 31 bit requests, however, the large 64 bit pool size allows for increased buffer hits over the 31 bit pool. The CPU increase will be fixed in a future release.

DataSet Level Parameters:

- Request Level Parameters:

  ACB:

  – RLSREAD (NRI/CR/CRE)

  - NRI - (default) No Read Integrity (will not get record lock)
  - CR - Consistent Read (will get/release record lock)
  - CRE – Consistent Read Extended (will get record lock, lock released at commit (recoverable data sets only).

Request Level Parameters

- Request Level Parameters:

  RPL:

  – OPTCD:

    - ASY/SYN - Asynchronous/Synchronous (SRB vs TCB)
    - DIR/SEQ/SKP – Direct/Sequential/Skip Sequential
    - NRI/CR/CRE - No Read Integrity/Consistent Read/Consistent Read Extended.

Performance Measurements

- SMF 64 Records
  - Cut by EOV and CLOSE on a ACB basis, fields since open:
    - SMF64DLR - number of logical records
    - SMF64DDE - number of delete requests
    - SMF64DIN - number of insert requests
    - SMF64DUP - number of update requests
    - SMF64DRE - number of retrieve requests
    - SMF64BMH - number of BMF hits in the local buffer pool.
    - SMF64CFH - number of CF hits in the RLS cache structure.
    - SMF64RIO - number of requests read from DASD
    - SMF64DEP - total number of requests.
    - SMF64NLR - number of logical records at open

Performance Measurements

- SMF 42 Subtypes 15, 16, 17, 18, 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.
  - 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,ALL operator command.

RMF Monitor III Displays

RMF Monitor III

- All statistics gathered from interface call with SMSVSAM address spaces
- Collected on RMF defined intervals
- Any SMSVSAM data in RMF monitor reports also found in SMF 42 records

RMF Monitor III - Sysplex Reports

RMF Sysplex Report Selection Menu

Selection ===>

Enter selection number or command for desired report.

Sysplex Reports

1	SYSSUM	Sysplex performance summary	(SUM)
2	SYSRTD	Response time distribution	(RTD)
3	SYSWKM	Work Manager delays	(WKM)
4	SYSENQ	Sysplex-wide Enqueue delays	(ES)
.			
10	**RLSSC**	VSAM RLS activity by storage class	(RLS)
11	**RLSDS**	VSAM RLS activity by data set	(RLD)
12	**RLSLRU**	VSAM LRU overview	(RLL)

RLSSC - VSAM RLS Activity by Storage Class

- Provides VSAM RLS activity data by storage class regarding direct/sequential requests accessing the local buffers, the CF cache structures and DASD.
- Assists in deciding the size of CF cache structures, buffer pool, and the overall performance of the application.

RLSSC - Sysplex Totals View

```
                  RMF V1R8   VSAM RLS Activity  - SYSPLEX        Line 1 of 12
Command ===>                                                  Scroll == => HALF
Samples: 59      Systems: 2    Date: 10/31/06  Time: 13.16.00  Range: 60    sec
               < 2GB / > 2GB
LRU Status    :  Good / Accel
Contention % :    0.0 /  0.0
False Cont % :    0.0 /  0.0
Stor Class  Access  Resp   -------- Read ----------  ------ BMF -------  Write
                    Time    Rate  BMF%   CF%  DASD%  Valid%  False Inv%  Rate
RLS_SC1
 Below 2GB    DIR    0.004  665.6  88.2   0.5   11.3    100       0.01      0.00
              SEQ    0.000   0.00   0.0   0.0    0.0    0.0       0.00      0.00
 Above 2GB    DIR    0.004  665.6  88.2   0.5   11.3    100       0.01      0.00
              SEQ    0.000   0.00   0.0   0.0    0.0    0.0       0.00      0.00
RLS_SC2
 Below 2GB    DIR    0.005  200.0  90.5   0.0    9.5    100       0.00      0.00
              SEQ    0.000   0.00   0.0   0.0    0.0    0.0       0.00      0.00
```

RLSSC - System/CF Structure View

```
                  RMF V1R8 VSAM RLS Activity   - SYSPLEX        Line 1 of 23
Command ===>                                                  Scroll == => HALF
Samples: 120     Systems: 2    Date: 10/31/06  Time: 13.25.00  Range: 120   Sec
               < 2GB / > 2GB
LRU Status    :  Good / Accel  Storage Class : RLS_SC1
Contention % :    0.0 /  0.0   Cache Set     : PUBLIC1
False Cont % :    0.0 /  0.0
System/CF   Access  Resp   -------- Read ----------  ------ BMF -------  Write
                    Time     Rate  BMF%   CF%  DASD%  Valid%  False Inv%  Rate
*ALL
  Below 2GB  DIR    0.000  14.98  83.0   0.0   17.0    100       0.00      0.00
             SEQ    0.000   0.00   0.0   0.0    0.0    0.0       0.00      0.00
  Above 2GB  DIR    0.000  14.98  83.0   0.0   17.0    100       0.00      0.00
             SEQ    0.000   0.00   0.0   0.0    0.0    0.0       0.00      0.00
SYS1
 CACHE01
  Below 2GB  DIR    0.000   7.49  83.0   0.0   17.0    100       0.00      0.00
             SEQ    0.000   0.00   0.0   0.0    0.0    0.0       0.00      0.00
  Above 2GB  DIR    0.000   7.49  83.0   0.0   17.0    100       0.00      0.00
             SEQ    0.000   0.00   0.0   0.0    0.0    0.0       0.00      0.00
```

RLSDS - VSAM RLS Activity by Data Set

- Provides VSAM RLS activity view by VSAM data sets.
- The information is grouped by VSAM spheres. For each data set, sysplex wide totals will be displayed for direct and sequential access.
- By VSAM components for which data collection has been requested are presented via the operator command V SMS,MONDS.

RLSDS - Sysplex Totals View

```
                  RMF V1R8 VSAM RLS Activity   - SYSPLEX        Line 1 of 20
Command ===>                                                  Scroll == => HALF
Samples: 120     Systems: 2    Date: 10/31/06  Time: 13.25.00  Range: 120  Sec
               < 2GB / > 2GB
LRU Status    :  Good / Accel
Contention % :    0.0 /  0.0
False Cont % :    0.0 /  0.0
Sphere/DS    Access  Resp    -------- Read ----------  ------ BMF -------  Write
                     Time     Rate  BMF%   CF%  DASD%  Valid%  False Inv%  Rate
BMAI.VSAMIN.MEGA
 BMAI.VSAMIN.MEGA.AIX.DATA
  Below 2GB  DIR     0.003    0.01   0.0   0.0    100     0.0       0.00      0.00
             SEQ     0.000    0.00   0.0   0.0    0.0     0.0       0.00      0.00
  Above 2GB  DIR     0.003    0.01   0.0   0.0    100     0.0       0.00      0.00
             SEQ     0.000    0.00   0.0   0.0    0.0     0.0       0.00      0.00
 BMAI.VSAMIN.MEGA.AIX.INDEX
  Below 2GB  DIR     0.003    0.03  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 2GB  DIR     0.003    0.03  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
```

RLSDS - System View

```
                RMF V1R8 VSAM RLS Activity   - SYSPLEX         Line 1 of 9
Command ===>                                                  Scroll == => HALF
Samples: 120     Systems: 2    Date: 10/31/06  Time: 13.25.00  Range: 120  Sec
              < 2GB / > 2GB
LRU Status    : Good  / Accel                      Storage Class : RLS_SC1
Contention % :   0.0 /  0.0                 Cache Set     : RLSCSET
False Cont % :   0.0 /  0.0            Data Set      : RLSADSW.VF01D.INVENTOR.IND
System/CF   Access  Resp   -------- Read ----------  ------ BMF -------  Write
                    Time    Rate  BMF%   CF%  DASD%  Valid%  False Inv%  Rate
*ALL
 Below 2GB   DIR    0.000  24.27  97.7   2.2   0.1    100       0.00      0.00
             SEQ    0.000   0.00   0.0   0.0   0.0    0.0       0.00      0.00
 Above 2GB   DIR    0.000  24.27  97.7   2.2   0.1    100       0.00      0.00
             SEQ    0.000   0.00   0.0   0.0   0.0    0.0       0.00      0.00
SYS1
 RLSCACHE01
 Below 2GB   DIR    0.000   0.20   100   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
SYS2
 RLSCACHE01
 Above 2GB   DIR    0.000  24.07  97.7   2.2   0.1    100       0.00      0.00
             SEQ    0.000   0.00   0.0   0.0   0.0    0.0       0.00      0.00
```

RLSLRU - Local Buffer Manager LRU Statistics

- Provides Local Buffer Manager LRU statistics for each system.
- Helps in adjusting the LRU goal and the local cache size.

An example of a RLSMAXPOOLSIZE = 500 and RLSABOVETHEBARMAXPOOLSIZE = 2000 specification:

```
                  RMF V1R8 VSAM LRU Overview   - SYSPLEX          Line 1 of 2
Command ===>                                                  Scroll == => HALF
Samples: 120     Systems: 2    Date: 10/31/06  Time: 13.25.00  Range: 120   Sec
MVS          Avg CPU  - Buffer Size -  Accel  Reclaim  ------ Read -----
System        Time      Goal   High      %       %     BMF%   CF%  DASD%
SYS1
 Below 2GB   0.023       500     352    0.0     0.0    82.0   8.0   10.0
 Above 2GB   3.543       2000    1552   0.0     0.0    95.0   2.0    3.0
SYS2
 Below 2GB   4.457       500     612    100.0   0.0    31.2   10.5   58.3
```

RLSLRU - Buffer Counts by Pool View

```
                RMF VSAM LRU Overview - Buffer Counts by Pool
The following details are available for MVS System: SYS2
Press Enter to return to the Report panel.
Fixed Storage  Below 2 GB :      0    Above 2 GB :      0
Real Storage % Below 2 GB :      0    Above 2 GB :      0
----------------- Below 2 GB ----  ---- Above 2 GB ----
                 Low    High     Avg     Low    High     Avg
Fixed Pages        0       0       0       0       0       0
Buffer Counts by Pool:
                                                      More: +
  2K             163     318     226       0       0       0
  4K             713    1537    1299       0       0       0
  6K               0       0       0       0       0       0
  8K               0       0       0       0       0       0
 10K               0       0       0       0       0       0
 12K               0       0       0       0       0       0
 14K               0    5000       0       0       0       0
 16K             460     678     656       0       0       0
 . . .
 30K               0    9000       0       0       0       0
 32K               0    7949       0       0       0       0
```

VSAM RLS Display Commands

Two New Display Commands

- This section will cover three RLS Diagnostic Console Commands
- D SMS,CFLS(lock_structure)
  - Will display lock structure statistics to the console
- D SMS,SMSVSAM,DIAG(Contention)
  - Will display any latch contention on SMSVSAM resources.
- D SMS,SMSVSAM,QUIESCE
  - Will display any outstanding Quiesce activity for the SMSVSAM address space.

Display SMS,CFLS(lock_structure_name)

- Displays Contention values for IGWLOCK00(default) or specified lock structure
  - Monitors locking statistics on four levels, minute, hour, 8 hour, and day.
  - Keeps track of the following data
    - LockRate = Number of lock requests that have come in per second
    - CONTRATE = The percentage of lock requests that encountered contention
    - FCONTRATE = The Percentage of lock requests that encountered false contention
    - WaitQLen = Average number of requests waiting behind locks.
- Recommended false contention rate is <.5%

Display SMS,CFLS

- Sample Results from the CFLS display

```
IGW320I 00:01:48 Display SMS,CFLS(IGWLOCK00 )
PRIMARY STRUCTURE:IGWLOCK00 VERSION:C57C859902B9264E SIZE:95232K
RECORD TABLE ENTRIES:34438 USED:747
SECONDARY STRUCTURE:IGWLOCK00 VERSION:C58548B866A7576E SIZE:95232K
RECORD TABLE ENTRIES:34438 USED:747
LOCK STRUCTURE MODE: DUPLEXED  STATUS: ENABLE
System    Interval    LockRate    ContRate  FContRate    WaitQLen
J90       1 Minute        15.1       0.000       0.331        0.00
J90        1 Hour         11.9       0.000       0.047        0.01
J90        8 Hour          8.2       0.001       0.039        0.00
J90         1 Day         17.3       0.001       0.042        0.01
  (09)    1 Minute          7.5       0.000       0.161        0.00
  (09)     1 Hour           7.0       0.000       0.020        0.00
  (09)     8 Hour           6.6       0.001       0.033        0.00
  (09)      1 Day          20.9       0.001       0.029        0.00

***************** LEGEND ******************
  LockRate = number of lock requests per second
  CONTRATE = % of lock requests globally managed
  FCONTRATE = % of lock requests falsely globally managed
  WaitQLen = Average number of requests waiting for locks
```

D SMS,SMSVSAM,DIAG(Contention)

- This command will display Latch contention on the system where it's issued.

- Latches are 8 bytes of storage used as a logical means to serialize resources. They are either held, or they are waited on… no shared requests.

- The Contention display will show the Latch address, the holder and any waiters if the latch is in contention.

- The display will show Elapsed time, for how long a latch was in contention.

DIAG(Contention) Example

- Sample Results from the Diag display

```
09.55.29 SYSTEM1           IGW343I VSAM RLS DIAG STATUS (V.01)
|---RESOURCE----| |------ WAITER ------| |--HOLDER---| ELAPSED
TYPE         ID    JOB NAME ASID   TASK   ASID   TASK     TIME
-------- -------- -------- ---- -------- ---- -------- --------
LATCH    7F158C70 SMSVSAM  003A 008DA250 003A 008D7218 00:00:06
    DESCRIPTION: IGWLYSPH - SHM OBJECT POOL
LATCH    7F151E78 SMSVSAM  003A 008D7218 003A 008DC1C8 00:00:21
    DESCRIPTION: IGWLYDTS - SHM OBJECT POOL
LATCH    7BAD43B8 SMSVSAM  003A 008DC1C8 002D 007F3000 00:19:09
LATCH    7BAD43B8 SMSVSAM  003A 008D5A48 002D 007F3000 00:22:09
LATCH    7BAD43B8 SMSVSAM  003A 008D6938 002D 007F3000 00:33:23
LATCH    07F1B1D0 SMSVSAM  003A 008D64F8 003A 008D6CF0 01:47:20
LATCH    07F1D3B8 SMSVSAM  003A 008D6CF0 0000 00000000 11:23:30
```

DIAG – What does it all mean?

- The DIAG command will help customers and Level 2 reps determine which systems are reporting elongated latch requests.

- Any latch contention with an elapsed time of more than a few seconds is most likely stuck waiting on something.

- Sometimes the holding ASID isn't SMSVSAM, but another address space, such as a CICS region. Canceling that region could avoid an SMSVSAM wide outage.

- DIAG should be used in conjunction with D GRS,C to determine if any of the latch holders are waiting on ENQ's

D SMS,SMSVSAM,QUIESCE

- The quiesce display will show any outstanding quiesce events.
- A Quiesce event is a decision to shutdown/allow CICS access to an RLS dataset.
- Any CICS regions that are part of the quiesce event will show up in the display
- If there is not an event in progress you will receive an IGW540I rejecting the command.

QUIESCE Example

- Sample Results from the Quiesce display

- **With** Quiesce Activity

```
IGW540I 13.30.45 DISPLAY SMS,SMSVSAM,QUIESCE
 MVS1
  SPHERE NAME: DLLEHR.TEST1
  SYSTEM NAME: MVS1     START TIME:    .27.50 TOTAL ELAPSE TIME: 57.02.55
  PARTICIPATING SUB-SYSTEM STATUS:      SCHEDULED:  COMPLETED:  ELAPSE:
   SUB-SYSTEM NAME:  CI1AORP1             .27.50     00.00.00   57.02.55
   SUB-SYSTEM NAME:  CI1AORP2             .27.50     .27.50     00.00.00
```

- **Without** Quiesce Activity

```
IGW540I 07.54.28 DISPLAY SMS,SMSVSAM,QUIESCE
DISPLAY SMSVSAM QUIESCE SPHERE IS REJECTED.
NO QUIESCE EVENTS ARE ACTIVE ON THIS SYSTEM.
```

QUIESCE – What does it all mean?

- The Quiesce command is invaluable to determine what CICS regions are holding up a quiesce request.

- Any CICS regions that have 00.00.00 in the completed section are most likely in trouble and should be dumped along with SMSVSAM before termination

- In most scenarios the CICS region can be terminated instead of SMSVSAM to allow the quiesce event to finish. Thus saving a SMSVSAM wide outage

DIAG and Quiesce APARs

- APARs needed to successfully use the DIAG command.
  - DIAG Introduced in base code for z/OS 1.8
  - OA17556 – PRE z/OS 1.8 retro fit for DIAG command. -- CLOSED
- APARs needed to for the Quiesce display
  - New Quiesce display in base code for z/OS 1.9
  - OA21101 – PRE z/OS 1.9 retro fit for Quiesce command -- CLOSED

Collecting Dumps

- A majority of needed information resides within the SMSVSAM address, and dataspaces
  - SMSVSAM, and MMFSTUFF Dataspaces should always be dumped
- Most functions require communication with other SMSVSAMs in the sysplex.
  - Sysplex wide dumps are a MUST!

- Example of Sysplex wide console dump

```
DUMP COMM=(some meaningful dump title)
        R xx,JOBNAME=(*MASTER*,XCFAS,SMSVSAM),CONT
        R yy,DSPNAME=('SMSVSAM'.*,'XCFAS'.*),CONT
        R nn,SDATA=(PSA,NUC,SQA,LSQA,SUM,RGN,GRSQ,LPA,TRT,CSA,XESDATA),CONT
        R zz,REMOTE=(SYSLIST=(*('SMSVSAM')),DSPNAME,SDATA),END
```

- *Recommended that Dump command kept in IEADMCxx parmlib member to ensure correct doc is collected*

References

- VSAM Demystified (SG24-6105-01)
- z/OS V1R11.0 DFSMS Storage Administration Reference (SC26-7402-13)
- z/OS V1R11.0 DFSMS Using Data Sets z/OS V1R10.0-V1R11.0 (SC26-7410-09)

Acknowledgements

- Special Thanks to the following people for their contributions
  - DFSMS RLS Level2 and Development teams
  - Terri Menendez for her work on the Buffer Performance
  - Kathleen Ting and Pooja Bhargava for their work with RMF III
  - Roity Prieto and Jason Lee for their work on the Display commands

Notices & Disclaimers

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 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 notinfringe IBM's intellectually 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 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 & Disclaimers

The performance data contained herein was obtained in a controlled, isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While IBM has reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere.

The responsibility for use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's or user's ability to evaluate and integrate them into their operating environment. Customers or users attempting to adapt these techniques to their own environments do so at their own risk. IN NO EVENT SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO,
LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY.

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 anyother 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.

Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Trademarks

DFSMSdfp, DFSMSdss, DFSMShsm, DFSMSrmm, IBM, IMS, MVS, MVS/DFP, MVS/ESA, MVS/SP, MVS/XA, OS/390, SANergy, and SP are trademarks of International Business Machines Corporation in the United States, other countries, or both.

AIX, CICS, DB2, DFSMS/MVS, Parallel Sysplex, OS/390, S/390, Seascape, and z/OS are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Domino, Lotus, Lotus Notes, Notes, and SmartSuite are trademarks or registered trademarks of Lotus Development Corporation. Tivoli, TME, Tivoli Enterprise are trademarks of Tivoli Systems Inc. in the United States and/or other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

Other company, product, and service names may be trademarks or service marks of others.