

What I learned about VSAM RLS SMF measurements



z/OS Performance
Education, Software, and
Managed Service Providers



Creators of Pivotor®

Nancy DiFilippo

Email: nancy.difilippo@EPStrategies.com

Enterprise Performance Strategies, Inc.

3457-53rd Avenue North, #145

Bradenton, FL 34210

<http://www.epstrategies.com>

<http://www.pivotor.com>



Contact, Copyright, and Trademark Notices

Questions?

Send email to Peter at Peter.Enrico@EPStrategies.com, or visit our website at <http://www.epstrategies.com> or <http://www.pivotor.com>.

Copyright Notice:

© Enterprise Performance Strategies, Inc. All rights reserved. No part of this material may be reproduced, distributed, stored in a retrieval system, transmitted, displayed, published or broadcast in any form or by any means, electronic, mechanical, photocopy, recording, or otherwise, without the prior written permission of Enterprise Performance Strategies. To obtain written permission please contact Enterprise Performance Strategies, Inc. Contact information can be obtained by visiting <http://www.epstrategies.com>.

Trademarks:

Enterprise Performance Strategies, Inc. presentation materials contain trademarks and registered trademarks of several companies.

The following are trademarks of Enterprise Performance Strategies, Inc.: **Health Check**[®], **Reductions**[®], **Pivotor**[®]

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM[®], z/OS[®], zSeries[®], WebSphere[®], CICS[®], DB2[®], S390[®], WebSphere Application Server[®], and many others.

Other trademarks and registered trademarks may exist in this presentation

Abstract

What I learned about VSAM RLS SMF measurements

There is a wealth of performance data buried inside SMF Type 42 records relating to VSAM RLS. This presentation is about making sense of that data and sharing what stood out along the way

EPS: We do z/OS performance...

- **Pivotor** – z/OS performance reporting and analysis software and services
 - Not just SMF reporting, but analysis-based reporting based on expertise
 - www.pivotor.com
- **Education and instruction**
 - We teach our z/OS performance workshops all over the world
 - Want a workshop in your area? Just contact me.
- **z/OS Performance War Rooms**
 - Intense, concentrated, and highly productive on-site performance group discussions, analysis and education
 - Amazing feedback from dozens of past clients
- **MSU Reduction Exercises**
 - The goal is to reduce the MSU consumption of your applications and environment
- **Information**
 - We present around the world and participate in online forums
 - <https://www.pivotor.com/content.html>
<https://www.pivotor.com/webinar.html>



z/OS Performance workshops available

During these workshops you will be analyzing your own data!

- Essential z/OS Performance Tuning
 - March 30 – April 3, 2026 (4 days, excl Wednesday the 1st)
- Parallel Sysplex and z/OS Performance Tuning
 - May 12-13, 2026
- WLM Performance and Re-evaluating Goals
 - July 27– 31, 2026 (4 days, excl Wednesday the 29th)
- Also... please make sure you are signed up for our free z/OS educational webinars! (email contact@epstrategies.com)

Like what you see?

- The z/OS Performance Graphs you see here come from Pivotor™
- If you don't see them in your performance reporting tool, or you just want a free cursory performance review of your environment, let us know!
 - We're always happy to process a day's worth of data and show you the results
 - See also: <http://pivotor.com/cursoryReview.html>

- We also have a **free** Pivotor offering available as well

- 1 System, SMF 70-72 only, 7 Day retention
- That still encompasses over 100 reports!

All Charts (132 reports, 258 charts)

All charts in this reportset.

Charts Warranting Investigation Due to Exception Counts (2 reports, 6 charts, [more details](#))

Charts containing more than the threshold number of exceptions

All Charts with Exceptions (2 reports, 8 charts, [more details](#))

Charts containing any number of exceptions

Evaluating WLM Velocity Goals (4 reports, 35 charts, [more details](#))

This playlist walks through several reports that will be useful in while conducting a WLM velocity goal an.

Like what you see?

- Free z/OS Performance Educational 2026 Winter/Spring webinars!
 - The titles for our Winter/Spring 2026 webinars are as follows:
 - ✓ *New Year's Resolutions for z/OS Performance and Capacity People*
 - ✓ *How WLM Makes Decisions*
 - *What I Learned about VSAM RLS SMF Data*
 - *z/OS Performance Spotlight: Some Top Things You May Not Know*
 - *Building a Strong Foundation When You Are New To z/OS Performance*
 - *z15 to z16 to z17 – What has changed? Foundation When You're New to z/OS Performance*
 - *Wait...Do We Need to Re-evaluate our WLM Goals*
 - *Evaluating in the Mainframe Environment*
 - *Managing Workload Manager: Multiple Sysplexes and Asymmetric Sysplexes*
 - *Introduction to z Processor Measurements*
 - *(more to be announced)*

What will we talk about ?

- Overview of VSAM RLS
 - *High level briefing of VSAM RLS*
- Lessons Learned Processing SMF 42
 - Storage Class
 - Data Class
 - CF Cache
 - RLS CF Lock
 - Local Buffer Management (LBM)

VSAM RLS Overview

- What is VSAM RLS?
 - Record Level Sharing for VSAM datasets
 - Record level serialization
 - 64-bit buffering
 - Global caching
 - Result in better performance for VSAM datasets
 - Enables multiple address spaces on multiple systems, and batch jobs to access recoverable VSAM data sets at the same time.
 - Multiple CICS systems can directly access a shared VSAM data set, eliminating the need to ship functions between the application-owning regions (AOR) and file-owning regions (FOR).

VSAM RLS Overview

- Why VSAM RLS?

Problem:

- One data set, many users, many systems
- Serialization of concurrent access is complex
- Data can get lost
- Data sets can be broken

Previous solutions:

- Use GRS and Homegrown methods
- CICS AOR / FOR

VSAM RLS Overview

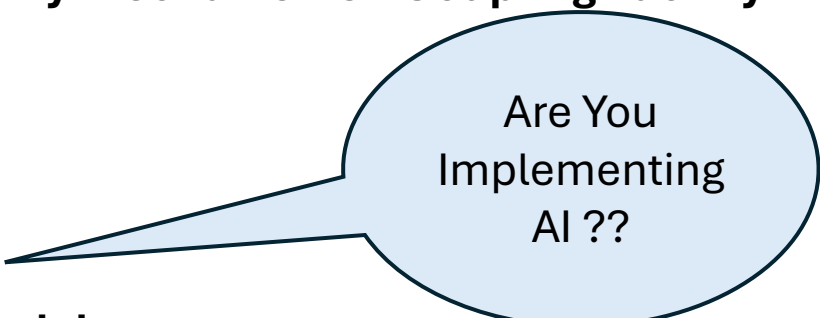
- Why VSAM RLS?

RLS solution:

- VSAM: Record Level Sharing
- All access goes through the SMSVSAM address space
- Record locks maintained in the CF allowing Sysplex-wide serialization
- VSAM control block structure kept in SMSVSAM data spaces or in SMSVSAM address space private area (above the bar).
- Some of the control block structure is copied to the coupling facility.
- **Each application does not have to provide locking and integrity mechanisms / Coupling Facility functionality does this**

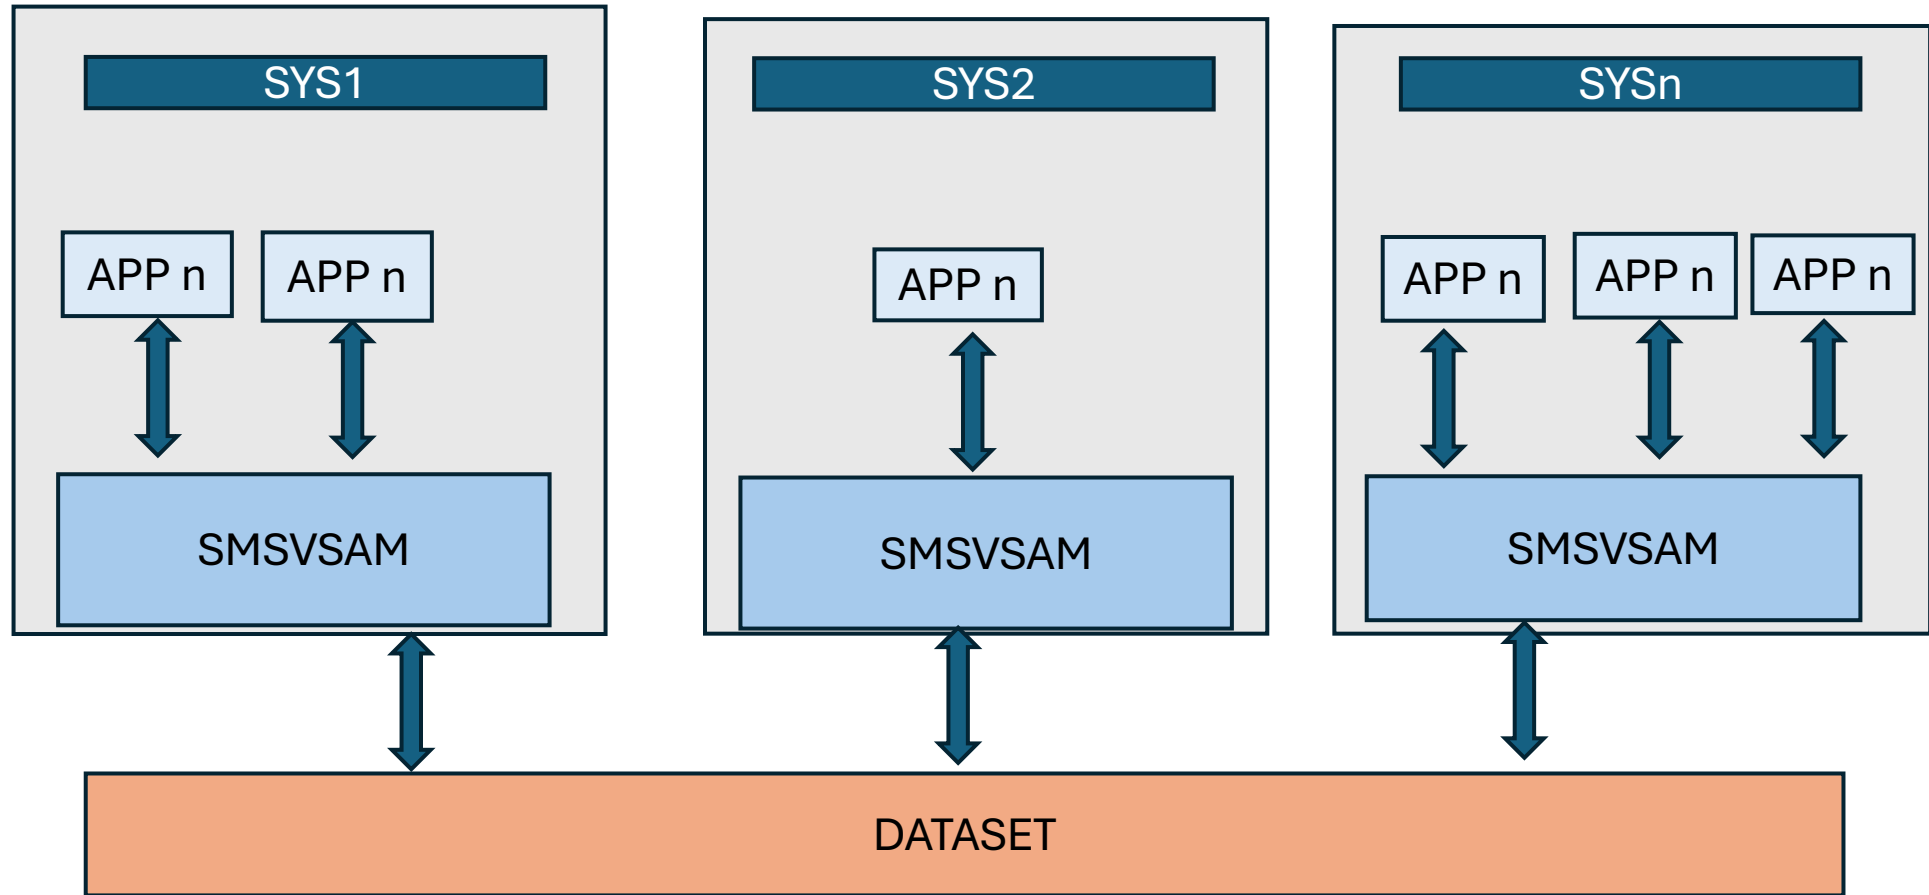
AI FUNCTIONALITY

- VSAM RLS is required for implementation of **EzNoSQL**
- **EzNoSQL – is the datastore to hold data for training the AI Model**



Are You
Implementing
AI ??

VSAM RLS Overview



VSAM RLS Overview

Transactional VSAM (TVS)

Enhance VSAM Record Level Sharing (RLS)

- Built on top of VSAM RLS
- provide data recovery capabilities for any application exploiting VSAM RLS.

RLS allows batch sharing of recoverable data sets for read

- RLS provides locking
- CICS provides logging and two-phase commit protocols

VSAM RLS Overview

IBM Products Exploiting RLS:

- CICS - User application files
- HSM - HSM Control Data sets (MCDS, BCDS, etc.)
- INFOMAN
- SCLM
- IMS (RLS and TVS) – IMS Recon Data Sets
- Catalog – User catalog

VSAM RLS Overview

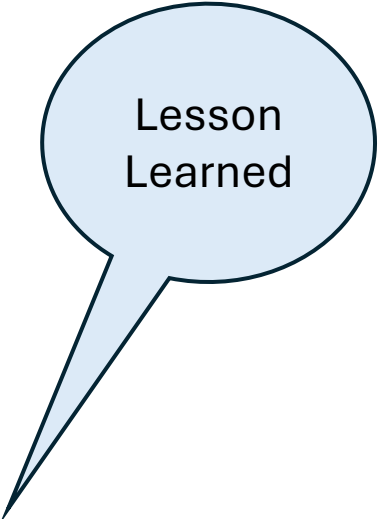
Setup

Coupling Facility :

- Define the Lock Structure (IGWLOCK00)
- Define Cache structures and Cache set

Setup SMSVSAM Address Space

- Must be active on each system in the sysplex that will be using RLS ← one per LPAR
- Define Buffer Pool Size
 - Below 2G : RLS_MAX_POOL_SIZE ← one value of all lpars
 - Above 2G: RLSAboveTheBarMaxPoolSize ← may be different on each LPAR
 - NO TUNING: RLS internally manages the buffers . The user is only responsible for supplying the size.



Lesson Learned

Define an SMS Storage Class and Data Class with RLS Support

- Associate the Storage Class with Cache Set
- RLS mode is determined at data set open time
- Data sets must set RLSAboveTheBar(YES) in the DATA CLASS to use the above the bar pool.

Lessons Learned

Lesson Learned : So Many Metrics

- SMF 42 Subtypes 15-19

- **Subtype 15** - RLS statistics by Storage Class

- **Subtype 16** - RLS statistics by Data Class

- **Subtype 17** - RLS locking statistics

- **Subtype 18** - RLS caching statistics

- **Subtype 19** - BMF statistics

- **Note: Only one system in the sysplex collects the SMF 42 records.**

- “For **Subtypes 15, 16, 17, 18 and 19**, specify the time interval in the IGDSMSxx parmlib member and activate recording on one system in the sysplex”

- <https://www.ibm.com/docs/en/zos/2.5.0?topic=records-record-type-42-x2a-dfsms-statistics-configuration>

- Be aware of SMF 62 and 64

- SMF 62 – Created by RLS OPEN for each ACB.

- SMF 64 – Created by RLS EOVS and CLOSE for each ACB. Stats are on an ACB level

VSAM RLS – Storage Class Measurements

SMF 42 Subtype 15

Purpose: SYSPLEX wide statistics to help tune RLS response time/cache performance for VSAM RLS Storage Class

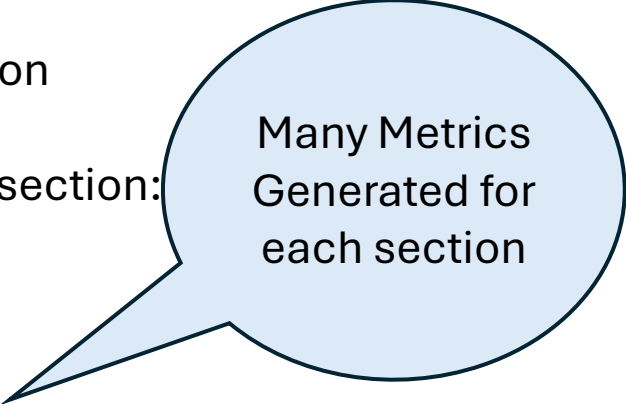
- Response Time is measured in milliseconds
 - Should be less than one millisecond

VSAM RLS – Storage Class Measurements

SMF 42 Subtype 15

- **Data Above 2g and Below 2g for Direct Access statistics and Sequential Access Statistics**

- Sysplex-wide direct access summary section:
- Sysplex-wide direct access record merge (RE-DO) summary section
- Sysplex-wide sequential access summary section:
- Sysplex-wide sequential access record merge (RE-DO) summary section:
- Sysplex-wide sequential access read ahead summary section:
- Sysplex-wide sequential access pre-format summary section:
- SC CF SYS summary section
 - Direct access section:
 - Direct access record merge (RE-DO) summary section:
 - Sequential access section:
 - Sequential access record merge (RE-DO) summary section:
 - Sequential access read ahead summary section:
 - Sequential access pre-format summary section:



Many Metrics
Generated for
each section

VSAM RLS – Storage Class Measurements

SMF 42 Subtype 15

LESSON LEARNED:

- In this and other subtypes with Storage Above and Below 2G, SMF nomenclature is to use same field name but replace “42” with “2A”

Eg. SMF42FCX vs SMF2AFCX

VSAM RLS – Storage Class Measurements

SMF 42 Subtype 15

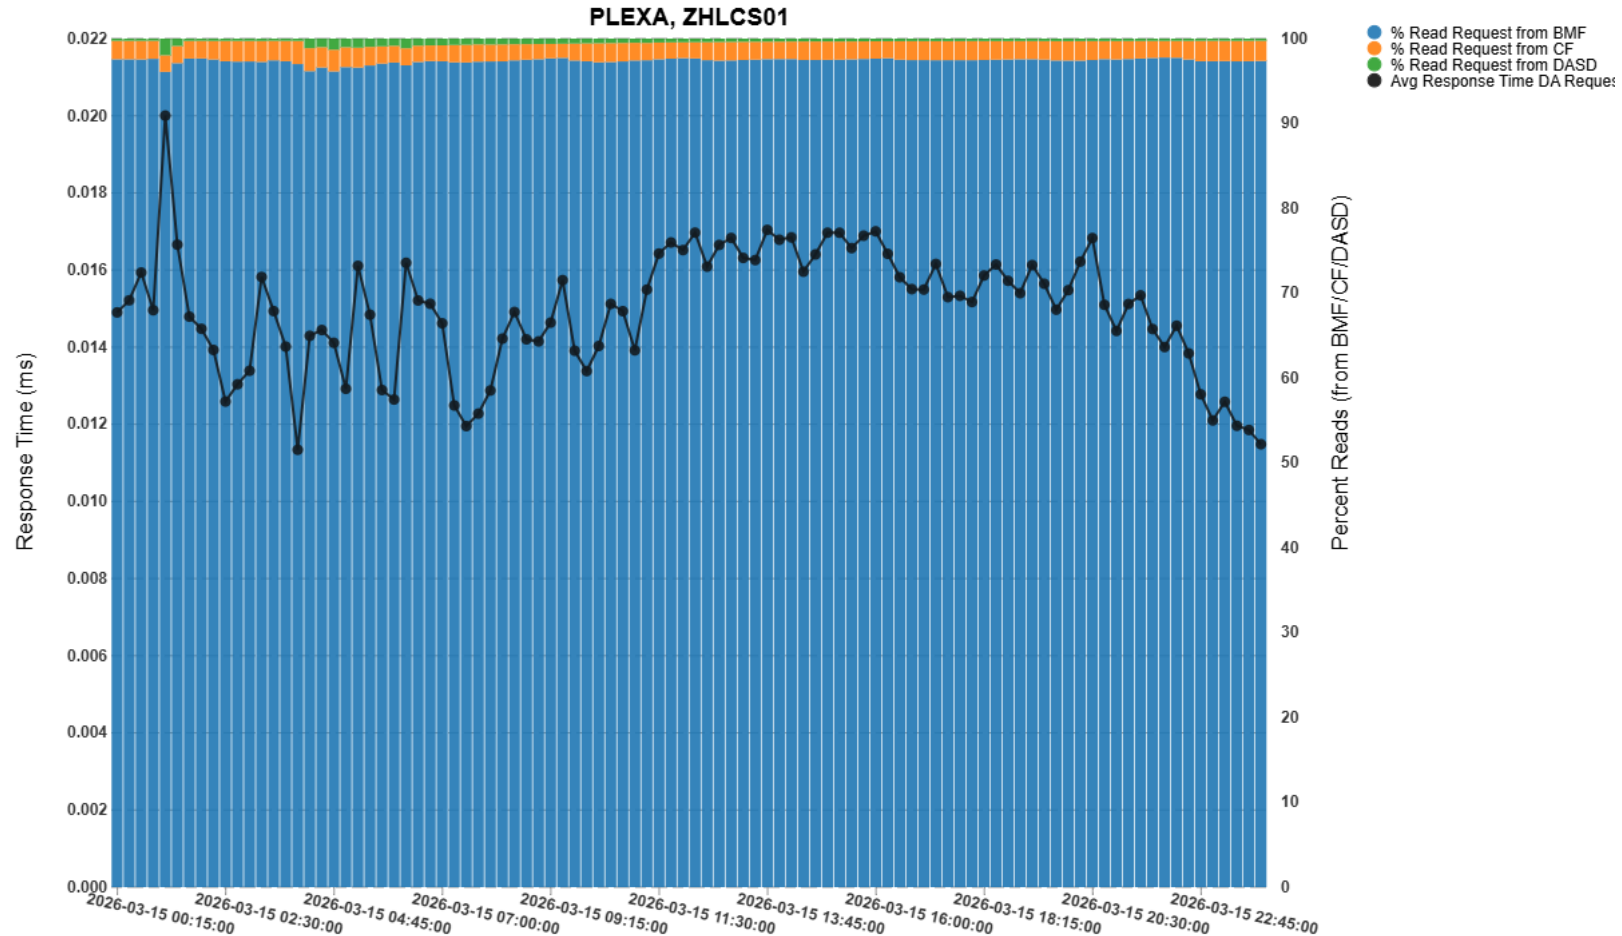
What data should I focus on?

Good measures of health

- Response Time
- Location of the data when it is read

VSAM RLS Storage Class Measurements

Above 2G Avg RT / Read %'s - Cache Set Name By SYSPLEX
Direct Access Requests for Storage Class SCRLS01



VSAM RLS – Storage Class Measurements

SMF 42 Subtype 15

LESSON LEARNED:

- Average and Normalized Response Times in this record are always zero
 - fields SMF2AFCX and SMF2AFCY for 2G storage
 - fields SMF42FCX and SMF42FCY for 31 bit storage

IBM answer: because response time is reported in milliseconds, a response time under a millisecond will be reported as zero. Plan to correct this at some time.

- Since RT is expected to be < 1 MS - IBM provided values will still be helpful
- Can use other fields in the record to calculate
 - Total Time in Milliseconds
 - field smf2afcw for 2G storage/field smf42fcw for 31 bit storage
 - Total BMF Requests
 - field smf2afcb for 2G Storage/field smf42fcb for 31 bit storage
- To calculate the Response Time in milliseconds with a non zero result

smf2afcw/smf2afcb
smf42fcw/smf42fcb

VSAM RLS – Storage Class Measurements

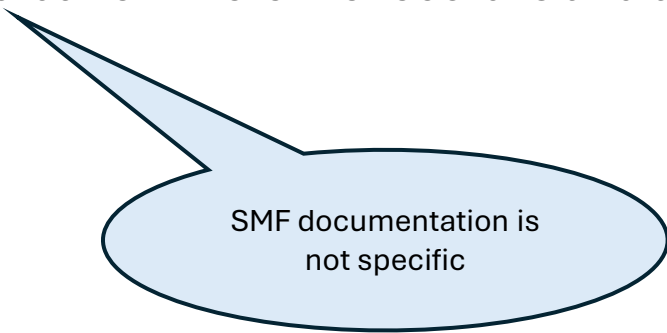
SMF 42 Subtype 15

LESSON LEARNED:

- Total Number of Request is < the Total BMF Request for both Direct Access and Sequential request.
(smf2afcb or smf42fcb) less than (SMF2Afcf or SMF42fcf)

IBM answer: total BMF requests is always > than the total Requests.

One application driven request will drive multiple calls to the RLS buffer manager to satisfy the request. The buffer manager has to retrieve the index buffers to find the data buffer where the record is and all of those buffer manager requests are recorded in the BMF requests.



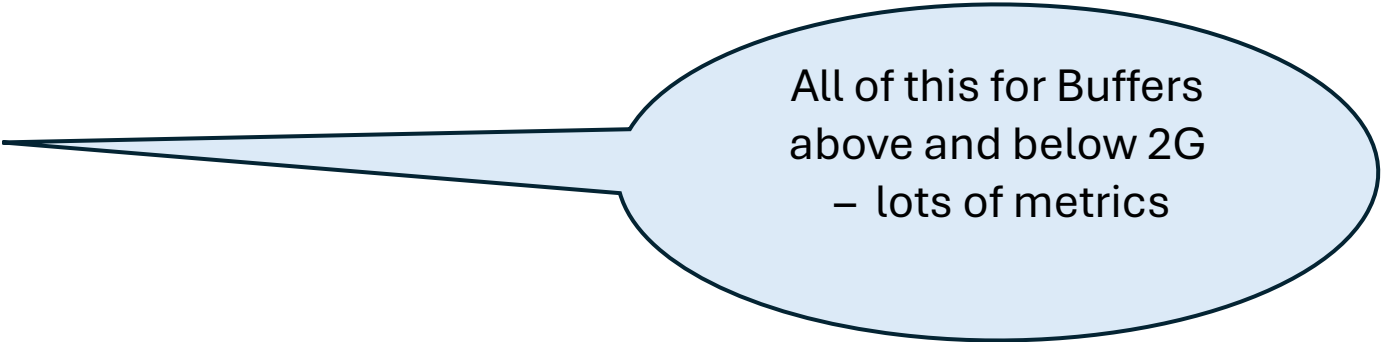
SMF documentation is not specific

VSAM RLS Data Class Measurements

SMF 42 Subtype 16

Purpose: Provide insight into activity for datasets opened under RLS control

- I/O Activity
- RLS Lock Statistics
- Cache Structure Statistics
- Buffer Pool Statistics
- Recovery Log Statistics

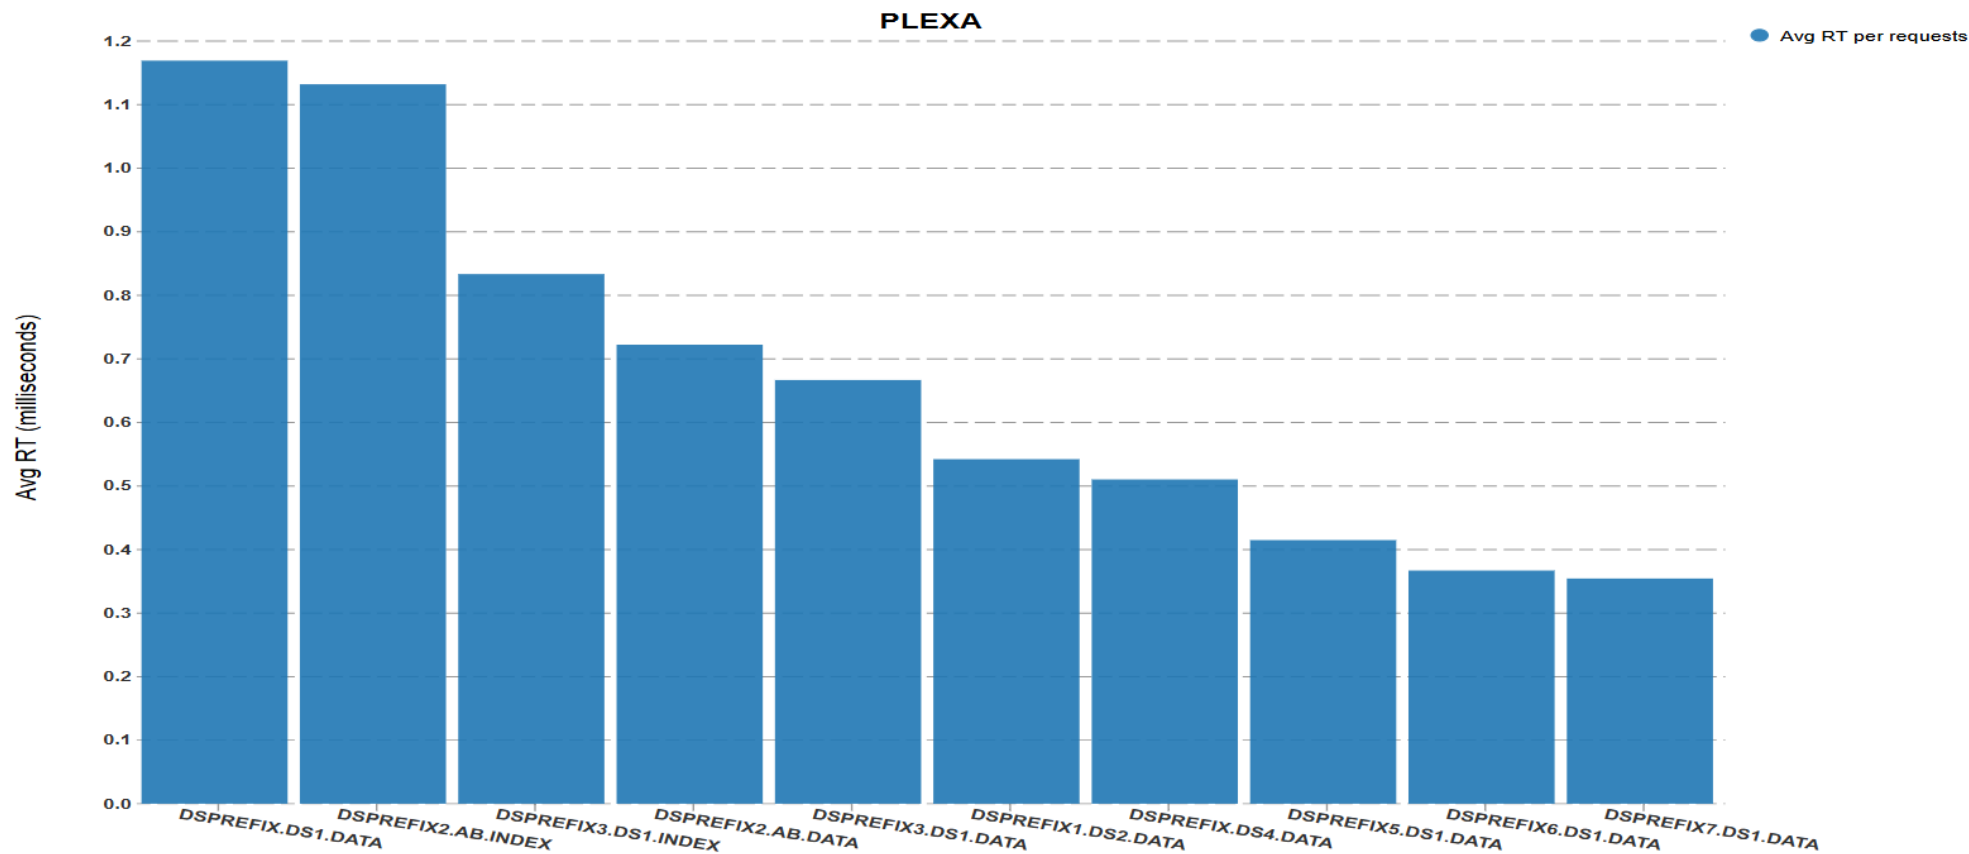


All of this for Buffers
above and below 2G
– lots of metrics

VSAM RLS Data Class Measurements

SMF 42 Subtype 16

Top RT for Direct Access RT Requests By Sysplex
for Storage Above 2G-



VSAM RLS CF Lock Measurements

SMF 42 Subtype 17

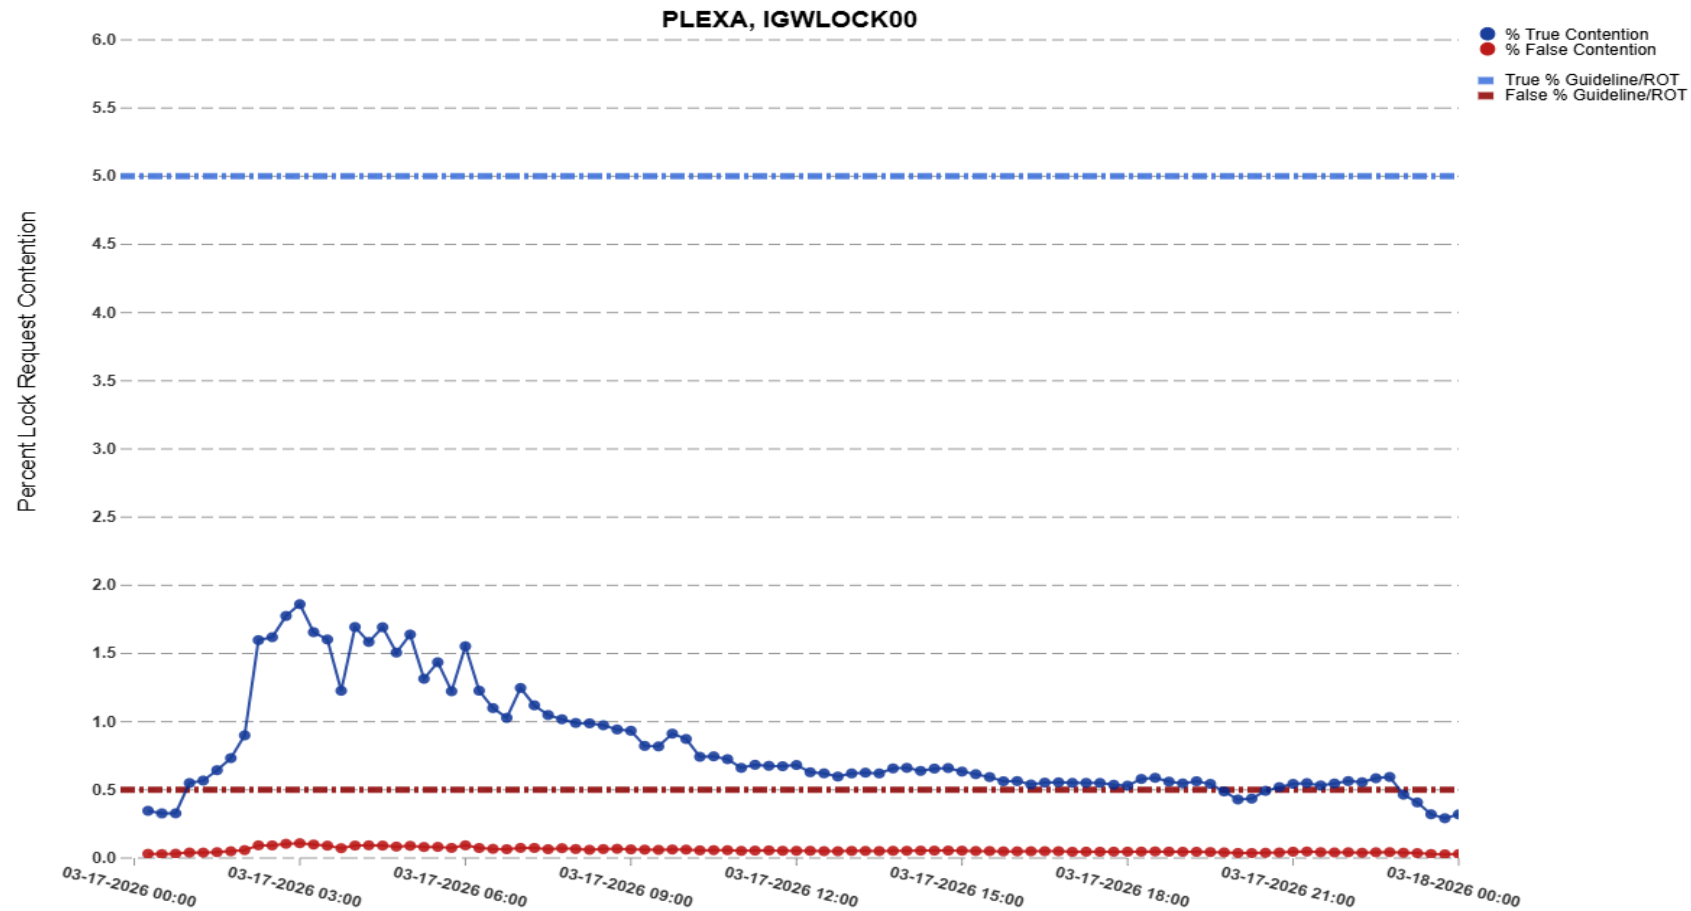
Purpose: Report on Coupling Facility Lock Structure Usage by the SMSVSAM address space

- Lock name is IGWLOCK00
 - IGWLOCK00 is required for SMSVSAM Server Setup
 - Up to 256 secondary locks can be defined - used to isolate workloads
 - Define Lock Set name in the SMS storage class - link to a specific CF lock structure
- Use to
- Identify Contentions
- Identify % Lock Structure Full

VSAM RLS CF Lock Measurements

SMF 42 Subtype 17

RLS CF Lock Requests Contention % by Sysplex/Lock

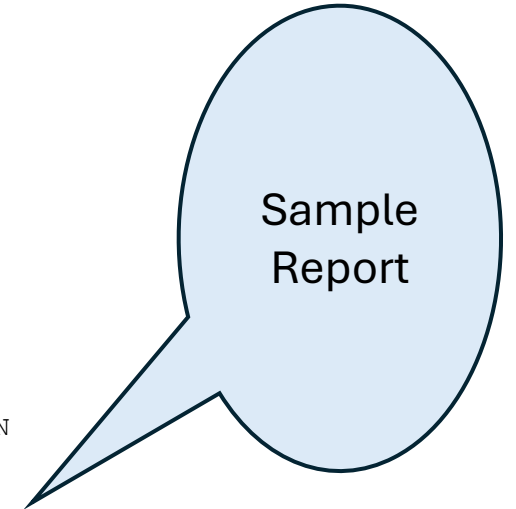


VSAM RLS CF Lock Measurements

Sample SMF 42 Reports

Lock contention:

START TIMESTAMP OF THE INTERVAL	STORAGE CLASS NAME	REC REQUESTS OB/ALT/PRO	LOCK REQUESTS CAUSED CONTENTION	REQUESTS TRUE DIWA LOCKS	COMP_1 CLASS_1 DIWA TRUE	COMP_1 CLASS_1 DIWA CONTENTIONS	COMP_1 CLASS_ DIWA FALSE	COMP_1 CLASS_ DIWA CONTENTION
25AUG2008:03:00:00.01	SCDVRLS1	307	0	0		0	0	0
25AUG2008:03:00:00.01	SCPVRLS1	533887	22	1537		0	0	0
25AUG2008:03:00:00.01	SCSVRLSD	0	0	0		0	0	0
25AUG2008:03:00:00.01	SCSVRLSP	3259867	1007	13390		2552	0	0
25AUG2008:03:00:00.01	SCSVRLSS	3112	0	10		0	0	0



Fields used in above report: SMF42FAC SMF42FOA SMF42FOB SMF42FOH SMF42FOI SMF42FOJ

VSAM RLS CF Lock Measurements

Lesson learned

- SMF 74 Subtype 4 contains coupling facility metrics for ALL locks including IGWLOCK00
- Are SMF 74 CF Lock Counts for IGWLOCK00 the same as those in the SMF 42 Subtype 17 ??
 - **Surprisingly no**
 - SMF 74 Subtype 4 field r744strc > the SMF42 subtype 17 field smf42hca

IBM Answer:

SMF 42 reports the number of "record" lock count.

SMF 74 is reporting on all locks obtained

Explanation: each data set that is opened in RLS mode has 6 special locks associated with them and SMF 74 is reporting on all of those.

VSAM RLS CF Lock Measurements

Lesson learned : Message IGW346I (from D SMS command) will identify specific jobs that are waiting for and holding locks.

NO SMF 42 Subtype 17 data that provides this information.

```
IGW346I VSAM RLS LOCK STATUS FOR STRUCTURE: IGWLOCK00 ,  
CURRENT NUMBER OF HOLDERS/WAITERS FOR LOCKS: 51,  
DISPLAY OF 50 HIGH ELAPSED TIME FOR HOLDERS/WAITERS FOLLOW:
```

```
-----  
LOCK_CL   ID      JOB NAME ASID TASK      EXC/SHR STATUS ELAPSED  
-----
```

```
CLUSTER NAME: SYSPLEX.KSDS0001
```

```
RECORD  92F92302 OA51165F 0032 007FF1C0 SHR HOLD 00:06:34  
RECORD  96E4F243 OA51165F 0032 007FF1C0 SHR HOLD 00:06:34  
RECORD  9AD0C184 OA51165F 0032 007FF1C0 SHR HOLD 00:06:34  
RECORD  C8A48E7B OA51165C 003B 007FF1C0 EXC HOLD 00:02:41  
RECORD  C8A48E7B OA51165G 0043 007FF1C0 SHR WAIT 00:01:53  
RECORD  C8A48E7B OA51165H 0044 007F8238 SHR WAIT 00:01:18  
RECORD  76E3170C OA51165A 0033 007FF1C0 EXC HOLD 00:00:34  
RECORD  7E766DD0 OA51165A 0033 007FF1C0 EXC HOLD 00:00:34  
RECORD  8609C495 OA51165A 0033 007FF1C0 EXC HOLD 00:00:34  
RECORD  C8A48E7B OA51165J 0045 007F8238 EXC WAIT 00:00:33
```

VSAM RLS CF Cache Measurements

SMF 42 Subtype 18

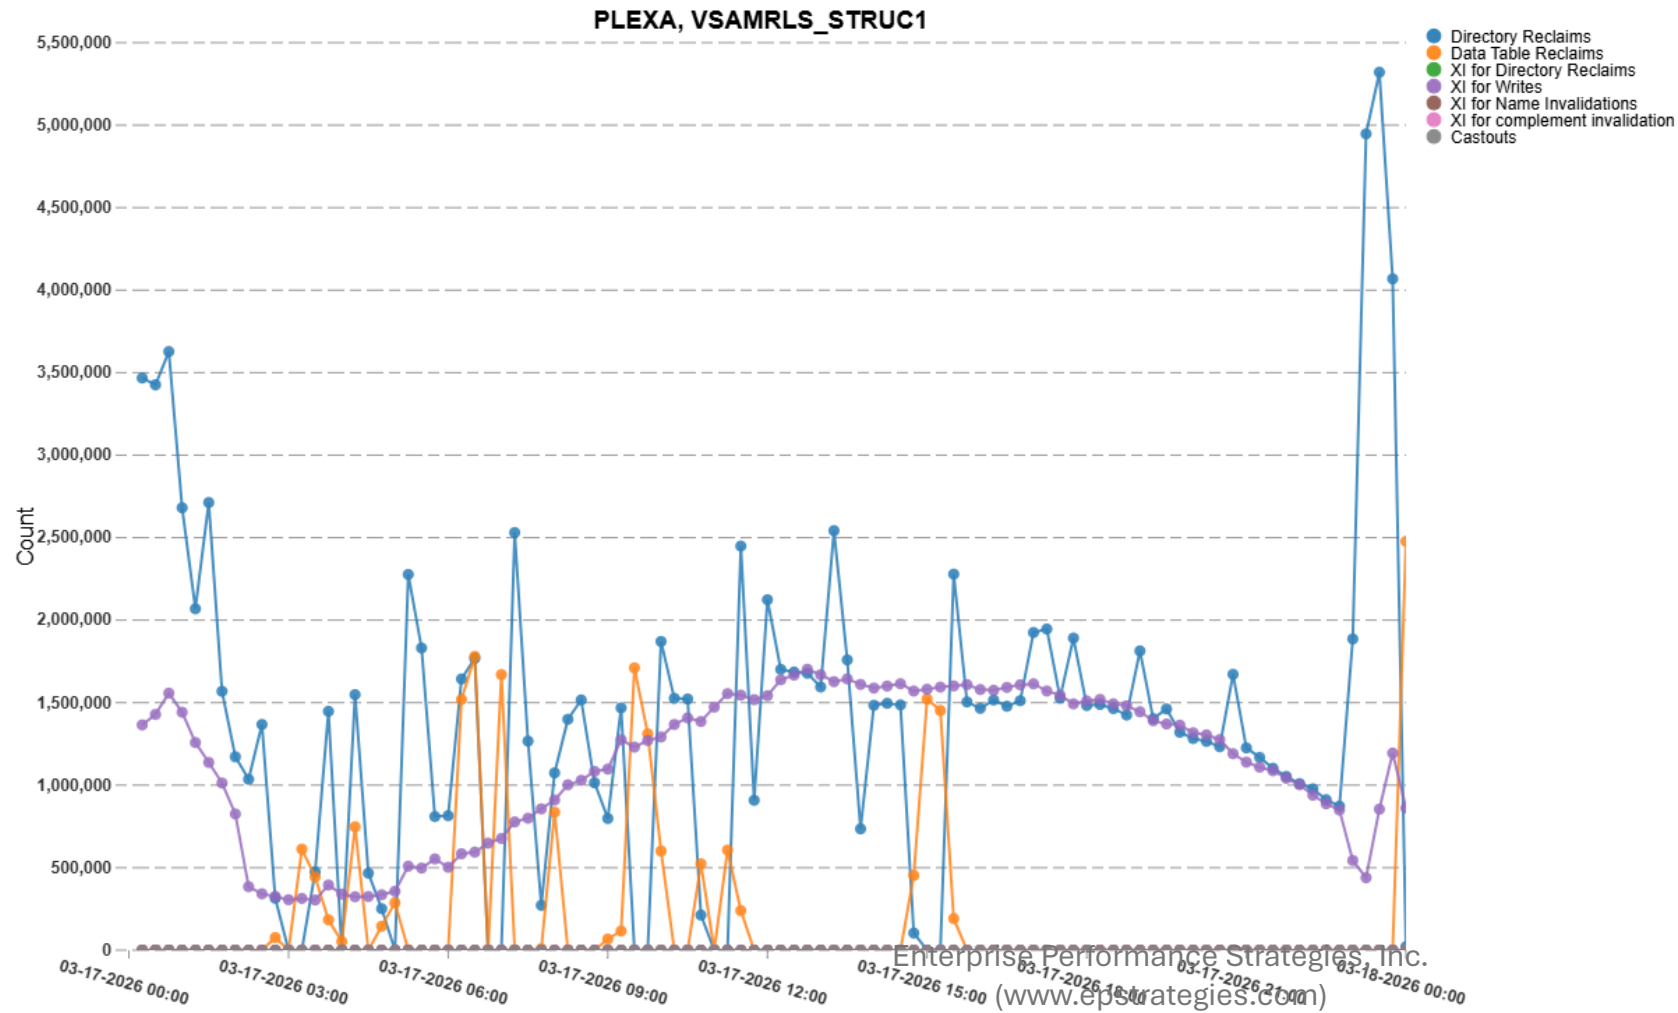
Purpose : to monitor how well VSAM RLS is utilizing the Coupling Facility Cache

- Cache Structure Usage
- Read/Write Activity
- Cache Hits/Misses
- Cross-Invalidation Stats

VSAM RLS CF Cache Measurements

SMF 42 Subtype 18

RLS CF Cache Structure Summary



VSAM RLS CF Cache Measurements

SMF 42 Subtype 18

Lesson Learned

- Summary section : reporting on cache usage at the overall level
- Total Section: provides cache usage at a more granular level more details such as cache structure name

Summary Section: contains multiple entries for same time period but there are no differences in the Read Hit counts

Total Section: contains only one entry for the same time period but the Read Hit Count > a single entry in the Summary

Values in the Summary section entries for a given interval add to the value in the Totals section.

VSAM RLS CF Cache Measurements

SMF 42 Subtype 18

Lesson Learned (continued)

Example

For TOTALS SECTION :

field **SMF42IAF** reported a value of **1342644** ← read count

For SUMMARY SECTION:

two identical entries

field **SMF42ICF** reports a value of **671322** ← read count

Doubled does match the value in the Totals Section

IBM Answer: There are both internal and external cache control blocks per cache structure which seem to include half the data. IBM was investigating to determine if this is the intended behaviour.

The good news is that it doesn't matter !

VSAM RLS CF Cache Measurements

SMF 42 Subtype 18

Lesson Learned

Cache Partition Summary field **SMF42IDG** records the cache partition type : DIR / SEQ /COMBINED

- SMF 42 Subtype 18 entries for all 3 types.
- DATA is only RECORDED for TYPE COMBINED

IBM Answer: Even though the records contain all the other partitions the values are always 0s

VSAM RLS LBM LRU Measurements

SMF 42 Subtype 19

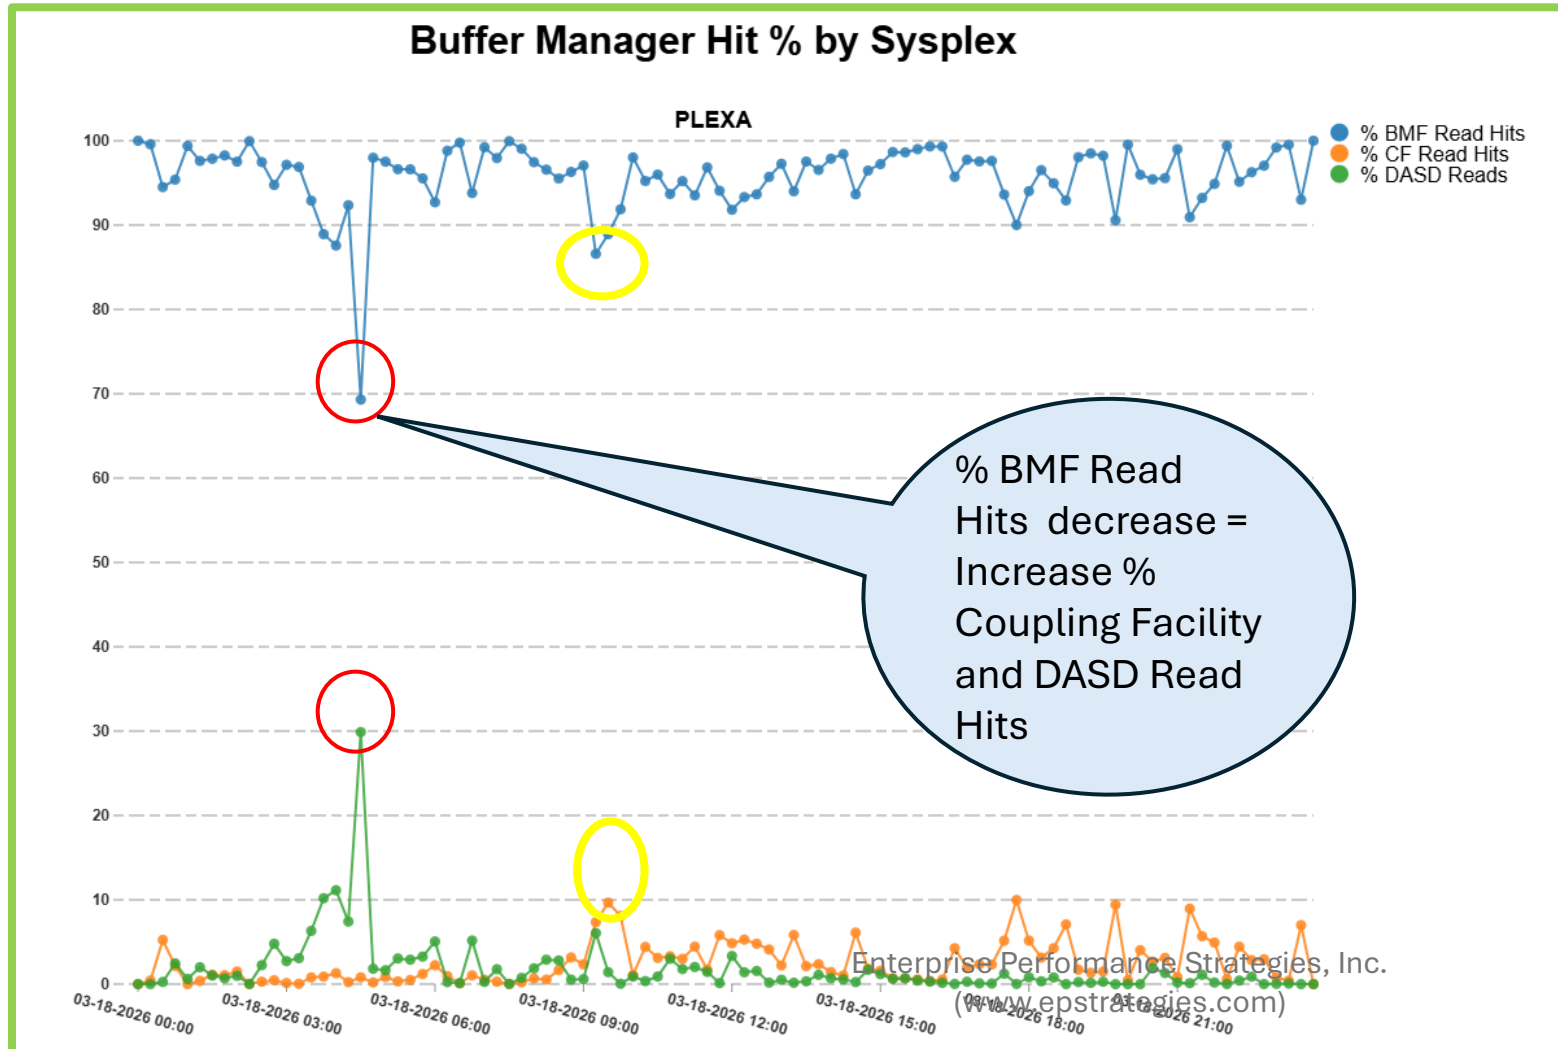
Purpose : Helps monitor local buffer management measurements.

- Storage Above and Below measurements
- Metrics on CPU used for Buffer Pool Management

- LRU → Least Recently Used

VSAM RLS LBM LRU Measurements

SMF 42 Subtype 19



VSAM RLS LBM LRU Measurements

SMF 42 Subtype 19

Least Recently Used Algorithms

Modes:

- Normal
 - Maintenance
 - Accelerated
 - Panic
-
- Mode is determined from % the buffer used compared to RLS max pool size
$$\text{RLSAboveTheBarMaxPoolSize} / \text{RLS_Max_Pool_Size}$$
 - Actions depend upon the mode
 - Actions are different depending upon buffer above or below 2G

VSAM RLS LBM LRU Measurements

SMF 42 Subtype 19

Least Recently Used Algorithms for Storage Above 2G

- Normal mode: < 80% of RLSAboveTheBarMaxPoolSize.
 - Buffers stay indefinitely.
- Maintenance Mode: > 80% but < 90% of RLSAboveTheBarMaxPoolSize .
 - Buffers 60 minutes or older will be released
- Accelerated Mode: > 90% but < 100% RLSAboveTheBarMaxPoolSize.
 - Buffers 30 minutes are older will be released
- Panic Mode: > 100% RLSAboveTheBarMaxPoolSize.
 - Buffers 5 minutes are older will be released.
 - Requests for new buffers will first be stolen
 - **If there are no buffers to steal, the request will sleep until LRU runs.**

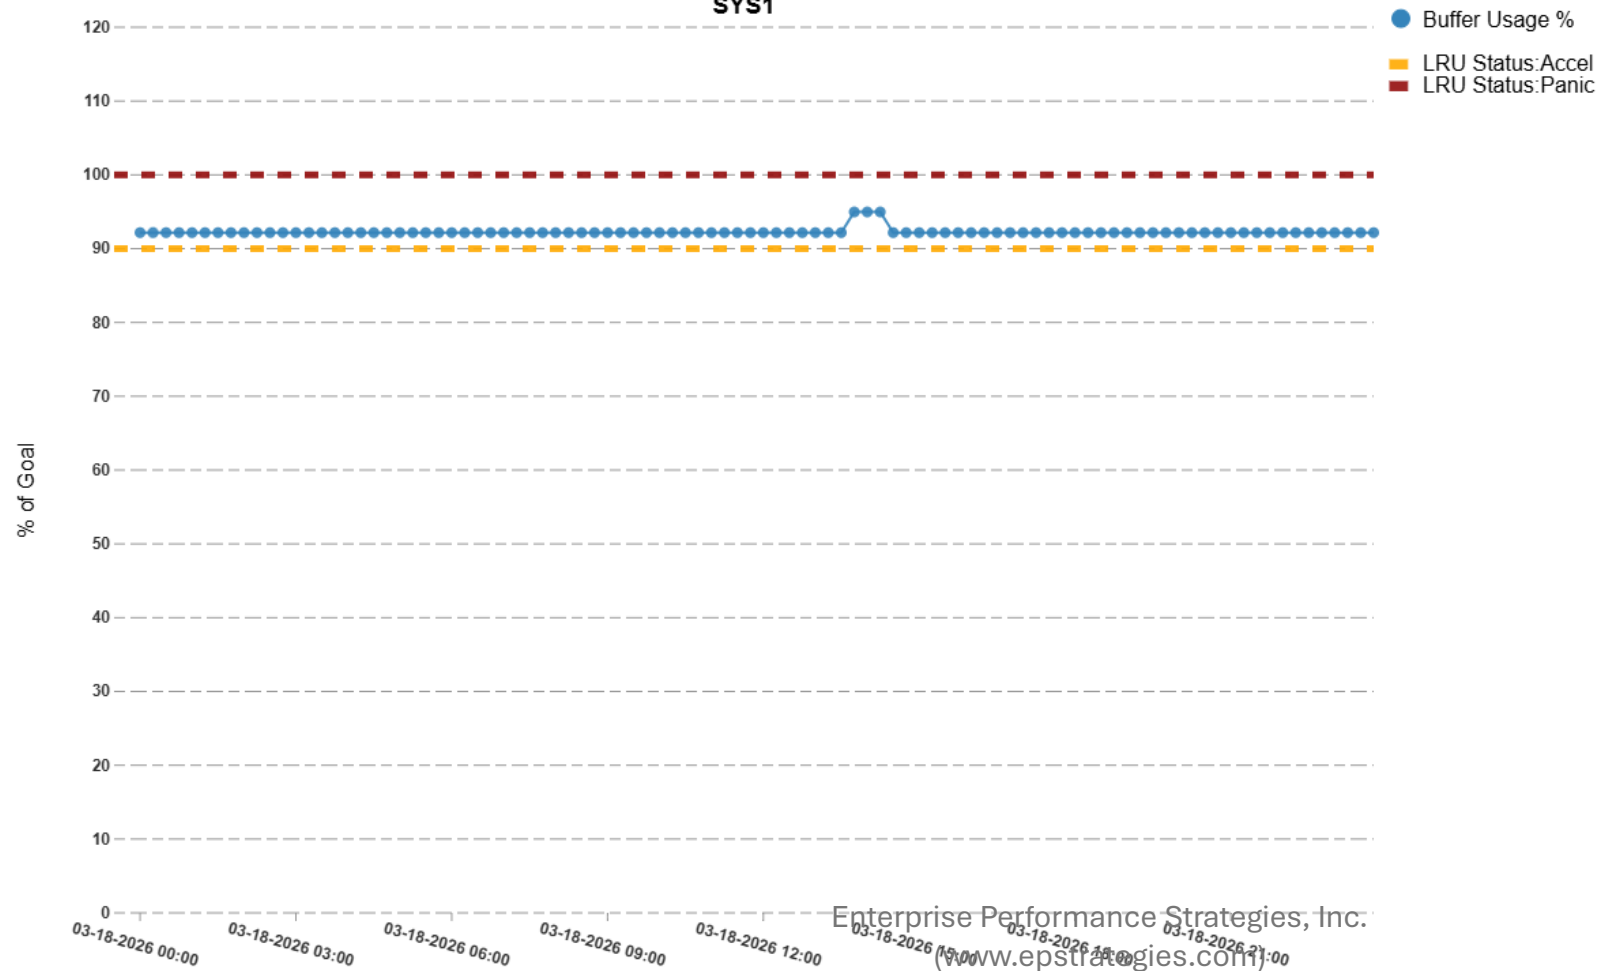
VSAM RLS LBM LRU Measurements

SMF 42 Subtype 19

Local Buffer Usage as % of Goal - By System

Storage Above 2G

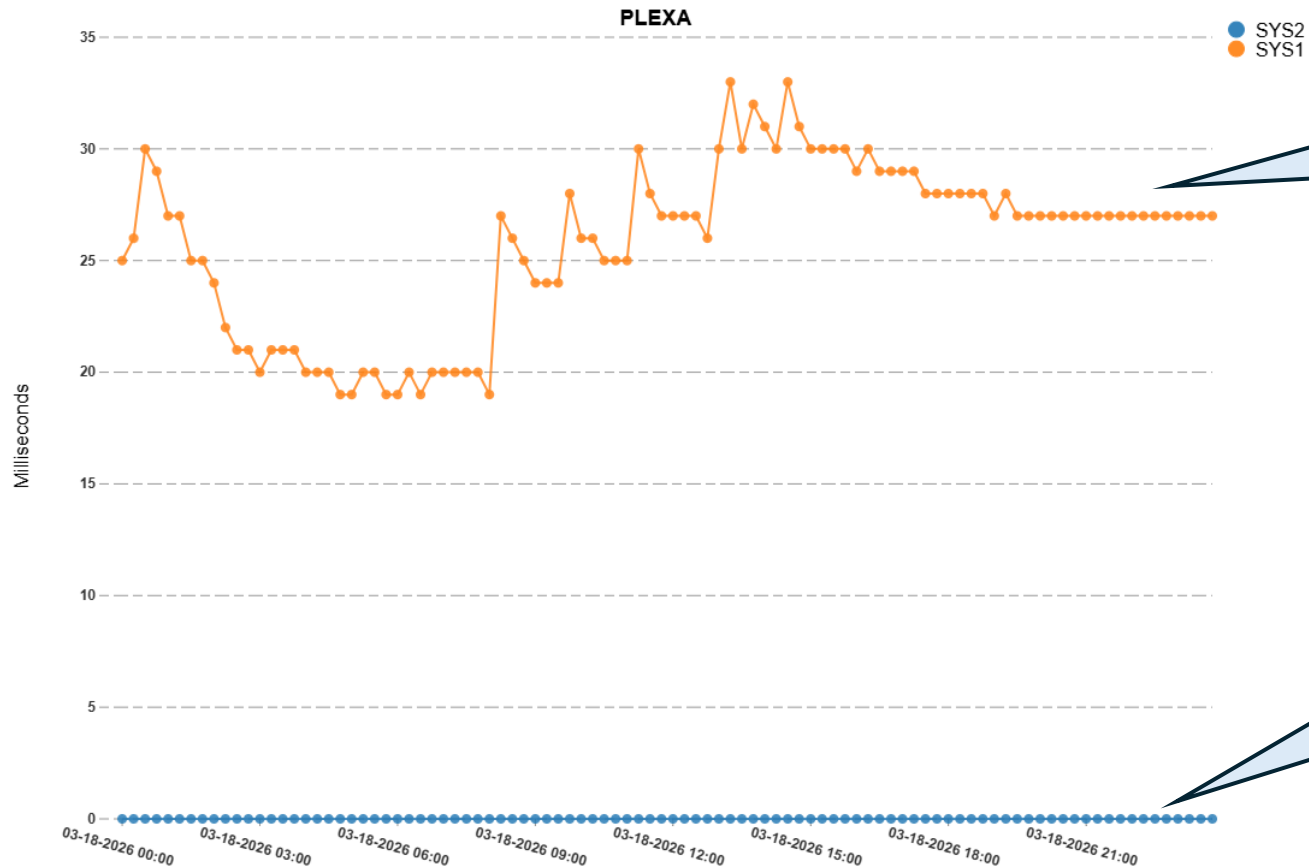
SYS1



VSAM RLS LBM LRU Measurements

SMF 42 Subtype 19

AVG CPU Time for LRU Monitoring by System
Storage Above 2G



SYS1: In Accelerated Mode - CPU used by SMSVSAM for Buffer Monitoring

SYS2: there is nothing to do so CPU consumption is zero

VSAM RLS LBM LRU Measurements

SMF 42 Subtype 19

Least Recently Used Algorithms for Storage Below 2G

- Normal mode: < 80% of RLS Max Pool Size.
 - Buffers stay indefinitely.
- Maintenance Mode: > 80% but < 120% of RLS Max Pool Size.
 - Buffers 60 minutes or older will be released
- Accelerated Mode: > 120% but < 2 times RLS Max Pool Size.
 - Buffers 30 minutes or older will be released
- Panic Mode: > 2 times RLS Max Pool Size or hits 1.7 G.
 - Buffers 5 minutes or older will be released.
 - Requests for new buffers will first be stolen.
 - **If there are no buffers to steal, the request will sleep until LRU runs.**

VSAM RLS SMF Measurements

Summary :

There is a plethora of metrics and naming is sometimes confusing

- 5 Different Subtypes
- Metrics for Direct Access and Sequential Access where applicable
- Metrics for Buffers Above and Below 2G

Recommendation: Initially stick to basic information within each record

If the metric is confusing connect with IBM ! They have been very helpful

VSAM RLS SMF Measurements

Questions ?

