



z/OS® Operating System

VSAM RLS diagnostic commands: DIAG and QUIESCE

@business on demand software

© 2008 IBM Corporation

This presentation covers the VSAM RLS diagnostic commands DIAG and QUIESCE.

What this presentation will cover

- This presentation will cover two new RLS Diagnostic Console Commands and any APARs necessary to run them.
- 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.

The Diag command will show all tasks waiting for latching resources and show which resources are currently held. Then you can manually terminate the SMSVSAM address space or RLS client space (on the correct system) to break up deadlocks.

The Quiesce command will display the status of RLS quiesce automation events in the sysplex.

D SMS,SMSVSAM,DIAG(Contention)

- This command will display latch contention on the system where it is issued.
- Latches are 8 bytes of storage used as a logical means to serialize resources. They are either held, or they are waited on; there are 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.

In VSAM RLS, latches provide system level serialization for shared control structures, ensuring that resources are either currently assigned to a process or are available for use. Because of their exclusive nature, latches tend to go into contention during VSAM RLS hangs. With the DIAG command you can ferret out latch contention and its possible source through a console message that displays any latch contention on that particular system.

Contention displays show latches that have a holder and a waiter with the elapsed waiting time. Although DIAG is not a sysplex-wide command, it can be routed across multiple systems. The DIAG command issues a semaphore-like response to contention: either there is contention or there is not. There is no shared response.

DIAG(Contention) example

- RESOURCE
 - ▶ TYPE -- Type of resource in contention. Almost always a Latch.
 - ▶ ID -- Address location of the 8 byte latch.
- WAITER
 - ▶ JOB NAME -- Job name of Waiter.
 - ▶ ASID -- Address space ID of Waiter.
 - ▶ TASK -- TCB of waiter.
- HOLDER
 - ▶ ASID -- Address space ID of Latch Holder.
 - ▶ TASK -- TCB of Latch Holder.
- ELAPSED TIME
 - ▶ The amount of time that the latch was in Contention.

The term “latch” refers specifically to internal SSF (System Services Facility) latches, unless otherwise explicitly stated. SSF latches have a scope of a single system, and are not shared across a sysplex. The command will operate on a single system basis. It will be accepted in any of the following forms:

```
D SMS,SMSVSAM,DIAG(CONTENTION)
D SMS,SMSVSAM,DIAG(C)
```

The command will result in one of the following outputs:

1. A console message indicating that there is no latch contention on this system.
2. A console message displaying all latch contention active on the system.
3. The VSAM RLS server is not ready.
4. Command rejected because it was issued too soon after the previous command.

DIAG(Contention) example #1

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

```

SMSVSAM shows no external system when experiencing contention. The SMS,DIAG(C) command will return with an IGW343I message and a list of latches in contention. There is one display per waiter. If the latch is in a list of commonly held resources then there will be a DESCRIPTION field detailing the latch. Although only the waiter's job name is displayed (and not the holder's), usually the waiter is the victim. The Diag command is also helpful in detecting deadlock. In case of deadlock, either SMSVSAM or another address space needs to be recycled.

This display reads as follows: (1) TCB 008DA250 is waiting on a latch that TCB 008D7218 holds, (2) TCB 008D7218 is waiting on a latch that TCB 008DC1C8 holds, (3) TCB 008DC1C8 is waiting on a latch that TCB 007F3000 holds, (4) TCBs (008DC1C8, 008D5A48, 008D6938) are all waiting on TCB 007F3000 for the latch at location x'7BAD43B8'.

Since ELAPSED TIME greater than one minute generally indicates a problem, there is a problem with TCB 008D64F8, which has been waiting to obtain Latch 07F1B1D0 for almost two hours (01:47:20) but not with TCB 008DA250, which has been waiting to obtain Latch 7F158C70 for only 6 seconds (00:00:06).

DIAG(Contention) Example #2

```

09.55.23 SYSTEM1          d sms,smsvsam,diag(c)
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 7BAD43B8    SMSVSAM  003A  008D5A48      003A  007F3000      00:22:09
LATCH 07F1B1D0    SMSVSAM  003A  007F3000      003A  008D5A48      00:22:09
LATCH 07F1B1D0    SMSVSAM  003A  008D64F8      003A  008D5A48      00:22:24
LATCH 07F1B1D0    SMSVSAM  003A  008D6CF0      003A  008D5A48      00:23:30

```

This shows the DIAG output for a system that is experiencing latch contention. In the DIAG output, RESOURCE ID is the address of the 8-byte area of virtual storage that represents a latch. For each held latch, the DIAG command also displays information about the job waiting for a resource (WAITER) and the process holding the resource (HOLDER), such as address space ID and TCB address. Note that ELAPSED TIME is not the entire time held, but rather the time, in minutes, that it was in contention.

The problem here is that two processes that are each holding a resource needed by the other. Task1 (in blue) is holding a latch that Task2 (in red) needs. Conversely, Task2 is holding a latch for which Task1 is waiting. In multiprocessing terms, this is called a deadlock. To resolve it, you must cancel one of the two tasks involved.

DIAG – What does it all mean?

- The DIAG command will help you determine which systems are actually experiencing a hang, as opposed to which are just victims.
- Any latch contention with an elapsed time of more than a few seconds is most likely stuck.
- Sometimes the holding ASID is not 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 GRS,C shows the ENQ (enqueue) contention. It lists all the TCBs (Task Control Blocks) which are waiting for an ENQ. It is possible that a TCB can be holding a latch and waiting for an ENQ, thereby becoming a victim as well. Therefore, only both commands together paint a complete contention picture.

The DIAG command provides latch contention on demand, determining the sources of any hangs or deadlocks caused by them. With this functionality, you will more easily see the causes of hangs and deadlocks, thus leading to the correct action. These actions could range from canceling batch jobs and CICS regions (client requests) to terminating an instance of the SMSVSAM server. Threads can hold one or more latches, and contend for other resources (such as, ENQ, Record Lock, Special Lock, and so on). Threads can also be in contention for a latch, potentially waiting indefinitely.

Whether it is a real or merely perceived hang, latches play a role. By displaying the status of VSAM RLS internal latches, even cross memory threads, the Diag command keeps you from unnecessarily restarting the SMSVSAM server, a disruptive procedure that affects all clients sysplex-wide, when all you need to do is cancel the thread holding the key latch.

D SMS,SMSVSAM,QUIESCE

- The quiesce display will show any outstanding quiesce events.
- A Quiesce event is a decision to allow unhindered 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.

One key feature in VSAM RLS is the ability to prevent RLS from accessing a specific data set. This action, called *quiescing* the data set, allows non-RLS applications unhindered access to the files without interference from RLS applications. To resume RLS access to the data set, you *unquiesce* it.

You might, for example, temporarily quiesce a data set from CICS online transactions to allow it to be opened by nightly batch jobs, or other non-RLS users of the dataset, and later unquiesce it to resume VSAM RLS operations. The DISPLAY,QUIESCE command, which is a system-wide command, has been around for several releases, but in z/OS V1R9, it was enhanced to provide more accurate data.

QUIESCE example

- **SPHERE NAME:** Name of dataset being quiesced/unquiesced.
- **SYSTEM NAME:** System where command was issued.
- **SUB-SYSTEM NAME:** Name of CICS region involved in Quiesce event.
 - ▶ **SCHEDULED:** Time when the quiesce event was issued.
 - ▶ **COMPLETED:** Time when CICS region responded to event.
 - ▶ **ELAPSE:** Time between when the quiesce event was issued and it was finally responded to by CICS.

Note that the ELAPSE field is not always zero. But if the COMPLETED field is blank, it means that it has not responded yet.

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:  C11AORP1          .27.50      00.00.00      57.02.55
```

```
  SUB-SYSTEM NAME:  C11AORP2          .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.
```

This shows the DISPLAY,QUIESCE output for a system with VSAM RLS data sets. DISPLAY,QUIESCE shows which data sets are being quiesced or unquiesced, along with the system on which the request was entered, the time of the request, and how long it has been since the quiesce or unquiesce event began. DISPLAY,QUIESCE also shows the names and status of any subsystems that are involved in the quiesce or unquiesce event.

If a critical batch run needs to access this data set, and a CICS region has yet to respond to the event, the batch jobs will not be able to open the data set. Likewise, if a data set is to be unquiesced, but one or more CICS regions have not responded to the event, the data set will remain inaccessible to all CICS regions.

In this example, the request was issued at 12:30 am (00.27.50) and has been waiting 57 hours (57.02.55) to complete. While region C11AORP2 has responded to the event as soon as it arrived (the COMPLETED and SCHEDULED values are the same), region C11AORP1 has yet to respond (COMPLETED=00.00.00). Since the COMPLETED value is 00.00.00, that CICS region has not responded to the quiesce. As far as RLS is concerned, quiesce events never time out, and so a request can wait forever.

This display is only for quiesce events affecting CICS regions on this system only. As a result, the command needs to be issued per system. Once all non-responsive CICS regions are cancelled, SMSVSAM will no longer require a response from them, and the quiesce can finish. If there are no quiesce events, the IGW540I message will reflect this.

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, instead of SMSVSAM, can be terminated to allow the quiesce event to finish. Thus saving a SMSVSAM wide outage.

In the previous slide, it is better to cancel one CICS region (C1AORP1) so that quiesce can finish, than to recycle SMSVSAM, which will clear all RLS address spaces. Quiesce events do not timeout. While SMSVSAM will wait forever for a Quiesce event to finish, CICS can issue an Unquiesce to cancel the Quiesce event. Unquiescing a dataset will notify every single system or CICS region in the sysplex and wait for every single system or region to respond.

There are several ways to quiesce a VSAM RLS data set. Two common methods are: (1) An explicit CICS request, (2) DFSMSdss™, when migrating or backing up a VSAM RLS data set. In either case, CICS notifies all regions in the sysplex that use the data set. CICS then either relinquishes control of the data set immediately, or completes processing before giving up control of the data set. After the batch processing or the data set migration is completed, the data sets are then unquiesced. Here, CICS is notified that the data sets are again ready to be accessed. This time around, however, it is not known which CICS regions currently need the data set or will need it in the future. To cover all bases, all CICS regions in the sysplex that are registered with a SMSVSAM address space are notified that a response is requested.

Even a lowly CICS test region must respond to an unquiesce event. If any CICS regions fail to respond to the notification, whether because a CICS region still needs the data set, or that lowly test region does not have enough dispatch priority, the entire unquiesce event will hang. If a system is running with a large number of CICS regions, it can be difficult to tell who responded to which events. The DISPLAY,QUIESCE command, which displays current quiesce and unquiesce events in progress for the system, helps pinpoint the root cause. If there is no activity, the DISPLAY output reflects this.

DIAG and Quiesce APARs

- APARs needed in order to 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.
- APARs needed in order to use 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.

While DIAG and DISPLAY, QUIESCE both provide detailed information about hang scenarios, neither will take corrective action. It is up to you to decide the best course of action to resolve the hang. These commands can help you shorten the list of possible culprits and narrow the problem down to a single system or region in the sysplex.

Also, you should dump SMSVSAM across the sysplex, along with any additional address spaces that show up in either display. Send the dump to IBM Customer Support before taking any action to resolve the hang.

Don't let a wait in VSAM RLS processing leave you hanging. By using DIAG and DISPLAY, QUIESCE to help detect the location of hangs related to SMSVSAM, you might avoid the need for a sysplex-wide outage.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_zOSV1R0_diag_quiesce.ppt

This module is also available in PDF format at: [../zOSV1R0_diag_quiesce.pdf](#)

You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

CICS IBM z/OS

A current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. 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. 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.

Information is provided "AS IS" without warranty of any kind. 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 NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not 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.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

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.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance 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 or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2008. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.