



## Mainframe SRDF/A MSC Troubleshooting Guide

EMC Proven™ Professional Knowledge Sharing 2008

Mike Smialek  
Sr. Solutions Architect  
EMC Corporation  
Smialek\_Mike@emc.com

## Table of Contents



.....	1
<b>Introduction</b> .....	<b>1</b>
<b>Purpose and Scope</b> .....	<b>1</b>
<b>Audience and Prerequisites</b> .....	<b>1</b>
<b>Reference Documentation and Resources</b> .....	<b>1</b>
<b>Input Data for Troubleshooting</b> .....	<b>2</b>
<b>Troubleshooting Guide</b> .....	<b>2</b>
<b>1. SCF Log Review</b> .....	<b>2</b>
<b>2. Buffer-to-Buffer Credits Display</b> .....	<b>5</b>
<b>3. SQ SRDFA Display</b> .....	<b>7</b>
<b>4. Service Processor – Event Displays</b> .....	<b>10</b>
<b>5. Recovery Procedures</b> .....	<b>12</b>
<b>6. Performance Considerations</b> .....	<b>13</b>
<b>7. Network Considerations</b> .....	<b>13</b>
<b>8. Balanced SRDF Links</b> .....	<b>14</b>
<b>9. Link Resiliency (Transmit Idle)</b> .....	<b>14</b>
<b>10. Delta Set Extension (5772+ Feature)</b> .....	<b>16</b>
<b>11. STAR Implementation Impact on SRDF/A</b> .....	<b>25</b>

*Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect EMC Corporation's views, processes or methodologies.*

## Introduction

The Mainframe SRDF/A MSC Troubleshooting Guide assumes you have basic experience with SRDF/A and MSC (Multi Session Consistency), and that your z/OS host software and Symmetrix® frames have been installed with SRDF® links available for processing. But something is not working as anticipated. You need to determine the problem and recommend an EMC resolution. SRDF/A and MSC problems vary but typically include:

- SRDF/A and MSC periodically drops processing
- High number of SRDF/A Transmit and Restore Delay messages
- Secondary Delay value higher than desired
- Instability of the links or network equipment
- Network throughput lower than anticipated
- SRDF/A performance issues caused by an unbalanced configuration

## Purpose and Scope

This document describes SRDF/A and MSC monitoring and problem resolution techniques for Symmetrix arrays using Enginuity™ versions 5x71 and later. The techniques provide a general outline of troubleshooting procedures and are not intended to be a complete list of every problem that an SRDF/A Administrator may encounter.

## Audience and Prerequisites

Because troubleshooting procedures are usually manipulating production data, EMC strongly recommends that users of this guide have previous SRDF and TimeFinder experience. [Open a software case with the EMC Software Assistance Center \(SAC\) at 800-782-4362 if you have questions concerning these procedures or problems executing the commands.](#)

## Reference Documentation and Resources

Much of the information in this operations guide is from the following EMC documentation:

- *EMC ResourcePak Base for z/OS v 5.5 Product Guide*
- *Symmetrix SRDF Host Component for z/OS v 5.3 Product Guide*
- *Symmetrix SRDF Host Component for z/OS v 5.3 Message and Code Guide*
- *TimeFinder/Mirror for z/OS v 5.4 Product Guide*
- *TimeFinder/Mirror for z/OS v5.4 Message and Code Guide*
- *TimeFinder/Clone Mainframe SNAP Facility v 5.6 Product Guide*
- *TimeFinder/Clone Mainframe SNAP Facility v5.4 Message and Code Guide*
- *EMC Storage Replication Implementation for Mainframe Service Operations Guide*

## Input Data for Troubleshooting

You can request a great deal of data from the customer that will help with troubleshooting. Focusing only on necessary documentation can help resolve the problem quickly. Here is a list of resources. Some are rather large, like SYSLOG and RMF, so they need to be uploaded to the EMC FTP site.

- SCF and RDF Logs
- z/OS SYSLOG
- WLA or STP data
- Symmetrix Service Processor log
- Network logs or reports
- RMF DASD Device Activity reports
- Open Service Requests
- Primus cases
- Documentation on frame configurations, SCF and RDF INI files

## Troubleshooting Guide

### 1. SCF Log Review

Start SRDF/A and MSC troubleshooting by reviewing messages in the SCF Log. There are informational and error messages that can give you a good idea of what happened. The following SCF log is delivering bad news:

```
00.51.09 STC52877 SCF1236I ASY -- CONTROLLER 00748 RDFGRP(10) SECONDARY DELAY = 294
00.51.09 STC52877 SCF1236I ASY -- CONTROLLER 01248 RDFGRP(08) SECONDARY DELAY = 294
00.51.09 STC52877 SCF1234I ASY -- CONTROLLER 20765 RDFGRP(02) ACT CHANGED ON -> OFF
00.51.12 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.51.12 STC52877 SCF1563I MSC - GROUP=MSCGRP01(8070,10) SER=000187751590 CYCLE SWITCH DELAY - RESTORE
00.51.18 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.51.18 STC52877 SCF1563I MSC - GROUP=MSCGRP01(8070,10) SER=000187751590 CYCLE SWITCH DELAY - RESTORE
00.51.24 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.51.30 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.51.36 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.51.42 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.51.48 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.51.54 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.52.00 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.52.06 STC52877 SCF1562I MSC - GROUP=MSCGRP01(8070,10) SER=000187700748 CYCLE SWITCH DELAY - TRANSMIT
00.52.09 STC52877 SCF1236I ASY -- CONTROLLER 00748 RDFGRP(10) SECONDARY DELAY = 354
00.52.09 STC52877 SCF1236I ASY -- CONTROLLER 01248 RDFGRP(08) SECONDARY DELAY = 354
00.52.10 STC52877 SCF1405E MSC - GROUP=MSCGRP01(8070,10) HOST CLEANUP INVOKED
00.52.12 STC52877 SCF1406E MSC - GROUP=MSCGRP01(8070,10) HOST CLEANUP IS RUNNING
00.52.12 STC52877 SCF1463E MSC - GROUP=MSCGRP01(3093,02) SRDFA IS NOT ACTIVE4
00.52.14 STC52877 SCF1414E MSC - GROUP=MSCGRP01(8070,10) HOST CLEANUP - PHASE2 IS RUNNING
00.52.14 STC52877 SCF1411I MSC - GROUP=MSCGRP01(8070,10) HOST CLEANUP CASE2 RUNNING
00.52.14 STC52877 SCF1409I MSC - GROUP=MSCGRP01(8070,10) PROCESS_FC10-DISCARD INACTIVE CYCLE
00.52.14 STC52877 SCF1409I MSC - GROUP=MSCGRP01(3093,02) PROCESS_FC10-DISCARD INACTIVE CYCLE
00.52.14 STC52877 SCF1409I MSC - GROUP=MSCGRP01(8900,08) PROCESS_FC10-DISCARD INACTIVE CYCLE
00.52.17 STC52877 SCF1413I MSC - GROUP=MSCGRP01(8070,10) HOST CLEANUP IS FINISHED
00.53.09 STC52877 SCF1234I ASY -- CONTROLLER 00748 RDFGRP(10) ACT CHANGED ON -> OFF
00.53.09 STC52877 SCF1234I ASY -- CONTROLLER 01248 RDFGRP(08) ACT CHANGED ON -> OFF
```

1. The SCF1236I Secondary Delay messages are interesting. The Secondary Delays are issued by RDFGroup number and inform you how many seconds the target box is out of synchronization with the source box. We can see the Secondary Delay is increasing from 294 seconds to 354 seconds.
  - a. How much Secondary Delay the solution can tolerate is a factor of cache, write activity, and available bandwidth. For an active mainframe solution, as much as 600 seconds or more of Secondary Delay may be needed to handle peak write activity. Use SymmMerge to determine how much cache is needed to provide a reasonable amount of Secondary Delay for your implementation.
  - b. Must update SCF parameters to get the Secondary Delay messages:  
**SCF.ASY.MONITOR=ENABLE**  
**SCF.ASY.SECONDARY\_DELAY=100** (Display if 100 seconds or >)  
(Always perform a SC GLOBAL,PARM\_REFRESH after making SCF parameter changes)
  - c. You may want to track Secondary Delay on new or changed implementations to determine the trend and peak write periods.
2. The SCF1562I Transmit Delay and SCF1563I Restore Delay are critical because they can indicate performance issues that might cause SRDF/A to drop, usually with a CACA10. First, you will get the Transmit and Restore Delay messages only when running with MSC. A standalone DMX-to-DMX SRDF/A solution will not produce the delay messages in the SCF log. You may see these messages sporadically in the SCF log during peak processing, but you should not see hundreds of repeated messages.
  - a. SCF1562I Transmit Delay means SRDF/A has a source cycle to Transmit to the target but can not obtain enough network bandwidth to transmit the cycle. As a result, MSC cycle switching is delayed. Some of the typical causes of Transmit Delays include:
    - Network link dropped or halted
    - Not enough network bandwidth for write activity
    - Network links not balanced
    - Buffer-to-Buffer credits need to be increased
    - Many Restore messages causes subsequent Transmit cycles to grow larger
    - Application growth, going to FICON, new higher MIPs processors

b. SCF1563I Restore Delay means MSC is waiting to do the cycle switch but is delayed by the target box because it is still processing the last Transmit cycle. Some of the typical causes of Restore Delays include:

- Target box RAID 5 (7+1) is not as fast as the source box RAID 1 or RAID 10
- Because of BCVs and Clones, target box doesn't have enough cache hit Write Pending limit
- BCVs are established to R2s; try splitting BCVs during peak periods
- Fewer 146 gb physical drives on the target compared to double 73 gb drives on the source
- Standards and BCV volumes on the same physical drives
- Consolidating multiple source boxes into one target

When running SRDF/A under MSC, the solution is only as fast as the slowest box in the configuration! Many customers will track the daily Transmit and Delay messages to determine when they occur, how frequently, and which frames are getting the messages. The text of the SCF log can be loaded into a spreadsheet and sorted by the SCF1562I and SCF1563I to provide answers to these questions. After several days, you can see what is "normal" for your implementation and determine if a SRDF/A CACA10 drop was caused by too many Transmit and/or Restore Delays. The spreadsheet may be simple:

		<u>Boxes Pair 1</u>	<u>Boxes Pair 2</u>	<u>Boxes Pair 3</u>	<u>Total</u>	<u>Max 2ndy Delay</u>
Wed-Thur 9/28/06		<u>S1 - T1</u>	<u>S2 - T2</u>	<u>S3 - T3</u>	<u>Total</u>	
2200-0700	TRANSMIT	232	74	2	308	
9 Hours	RESTORE	88	8	0	96	311 Seconds
Thur-Fri 9/29/06		<u>S1 - T1</u>	<u>S2 - T2</u>	<u>S3 - T3</u>	<u>Total</u>	
2200-0700	TRANSMIT	384	146	6	536	
9 Hours	RESTORE	2	22	0	24	404 Seconds
Fri-Sat 9/30/06		<u>S1 - T1</u>	<u>S2 - T2</u>	<u>S3 - T3</u>	<u>Total</u>	
2200-0051 Drop	TRANSMIT	696	288	16	1000	TRANSMIT Issue
3 Hours 16037655	RESTORE	17	6	0	23	414 Seconds
Sun-Mon 10/01/06		<u>S1 - T1</u>	<u>S2 - T2</u>	<u>S3 - T3</u>	<u>Total</u>	
2200-0700	TRANSMIT	134	32	24	190	
9 Hours	RESTORE	29	3	0	32	248 Seconds

There is an Excel macro called SCF-SYSLOG-Parser-MJS-10-26-07V1.4.xls that can scan any amount of hourly, Shift, Daily, Weekly customer SCF or SYSLOG text log looking specifically for the following SRDF/A messages:

- 1) SCF1233I – Tolerance Changed OFF -> ON
- 2) SCF1234I – Controller (xxxx) ACT Changed ON -> OFF
- 3) SCF1236I – Secondary Delay
- 4) SCF1562I – Transmit Delays by box
- 5) SCF1563I – Restore Delays by box
- 6) SCF1524I – Global Inconsistency
- 7) SCF1586I – Transmit Idle Invoked

The macro will support up to 12 source Boxes and 12 target Boxes providing Transmit and Delay counts by individual box. After entering the last four digits of the source and target Box serial numbers once, the macro remembers the serial numbers so they don't have to be entered again on subsequent executions for the same customer. Send email to [smialek\\_mike@emc.com](mailto:smialek_mike@emc.com) for macro.

3. Other critical messages in the sample SCF log are SCF1405E Host Cleanup Invoked and SCF1406E Host Cleanup Running which means SRDF/A has dropped for some reason and the software is trying to determine the Case 1, 2, or 3 for cleanup actions, (see Recovery Procedures.) These messages are good FIND text editor fields because they take you right to the drop point in the log. The other message, SCF1234I Controller 0748 RDFGRP(10) ACT CHANGED ON -> OFF, indicates box 0748 RDF Group 10 has dropped SRDF/A.

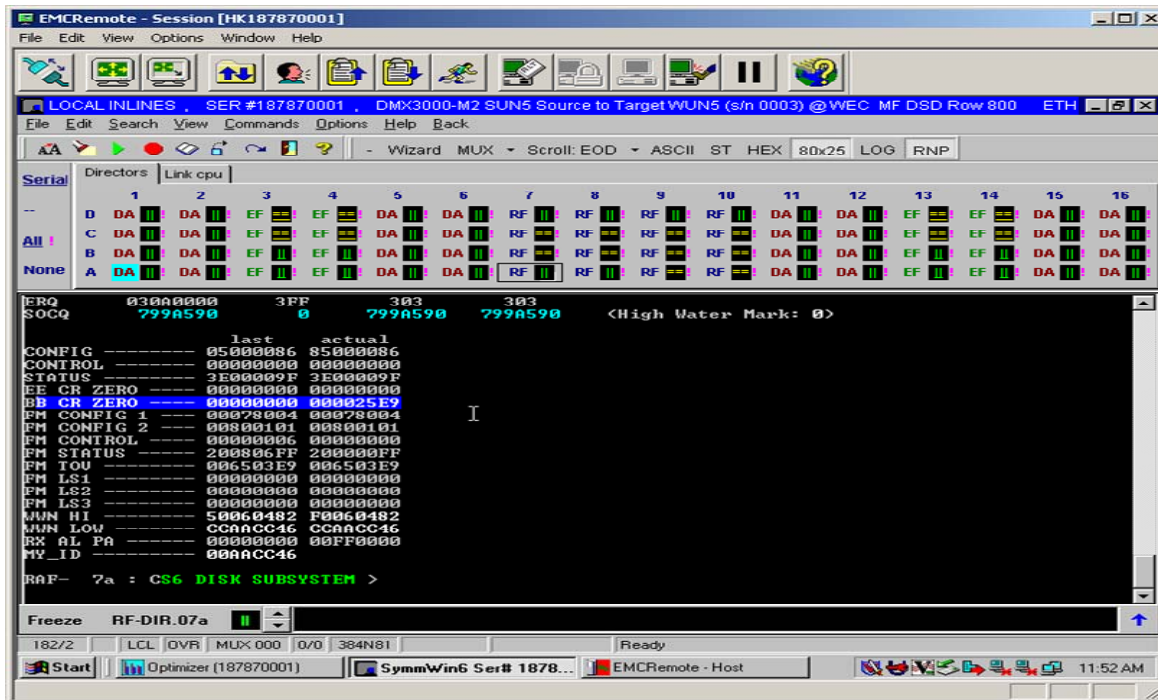
## **2. Buffer-to-Buffer Credits Display**

If you use InLines commands or can get an RTS or CE to execute InLines commands, there are two important network values you can verify which dramatically impact SRDF/A performance. Is the Symmetrix accessing the expected number of switch BB\_Credits and are you periodically maxing out the number available? Sometimes customers complain about throughput because they have a 100 MB/s pipe but the problem is not enough B2B\_Credits to fill the pipe. Typically, switches have between 16-60 B2B\_Credits; verify the value with network personnel.

The following shows the output of the Symm 6 E1 InLines command. The two register fields you want to review are BB CR ZERO and FM CONFIG 1. Each tick in the BB CR ZERO hexadecimal field is 10us elapsed time with NO B2B credits available. In the display BB CR ZERO = 0000, but if it had a value of 0010 it would mean we had no B2B credits 16 times.

To determine the B2B credits available from the switch, look at the FM CONFIG 1 register's byte 01 (assuming you count 0, 1, 2, 3, etc...). On the display, FM CONFIG 1 = 00078004. Since byte 01 = 07 hex, this box is talking to a switch with 7 B2B\_Credits available. If it had a value of 00108004 the hex 10 = decimal 16. This is the maximum number of B2B\_Credits available from the switch, NOT the number being used. Said another way, I would not expect to see the 07 change in our example if we did multiple E1 displays whereas the BB CR ZERO could change between displays. Typically, the customer network personnel can obtain similar B2B\_Credit information from the switch.

## BB CR ZERO incrementing



E1, on Symm 6    EF,E1 on Symm 7

Each tick is 10us elapsed with no B2B credits available



### 3. SQ SRDFA Display

The SQ SRDFA is one display that can help you understand what is happening. This should be executed about every five minutes to determine the SRDFA activity during the peak processing cycles, especially with new implementations. The display output will go to the RDF Log and z/OS SYSLOG. In the following example, things are not good. Pay attention to Secondary Consistent (?) which means SRDFA doesn't know if the data is consistent. Add the Capture Cycle Size 171,180 + the Transmit Cycle Size 0 x 56K per cache slot for a DMX2000 = 10 GB of cache just to hold the cycles. By the way, this DMX2000 had 45 GB of usable cache. Secondary Delay of 323 seconds and Cleanup Running (Y) tells you SRDFA has dropped on this box!

```
01.48.22 STC22138 EMCMN00I SRDF-HC : (5177) #SQ SRDFA,836B
01.48.22 STC22138 EMCQR00I SRDF-HC DISPLAY FOR (5177) #SQ SRDFA,836B 410
410 MY SERIAL # MY MICROCODE
410 -----
410 000187700748 5671-54
410
410 MY GRP ONL PC OS GRP OS SERIAL OS MICROCODE SYNCHDIR FEATURE
410 -----
410 LABEL TYPE AUTO-LINKS-RECOVERY LINKS_DOMINO MSC_GROUP
410 -----
410 10 Y F 10 000187751590 5671-54 G(R1>R2) SRDFA I MSC
410 Box-S1 STATIC AUTO-LINKS-RECOVERY LINKS-DOMINO:NO (MSCGRP01)
410 -----
410 PRIMARY SIDE: CYCLE NUMBER 7,811 MIN CYCLE TIME 30
410 SECONDARY CONSISTENT ( ? ) TOLERANCE ( N )
410 CAPTURE CYCLE SIZE 171,180 TRANSMIT CYCLE SIZE 0
410 AVERAGE CYCLE TIME 57 AVERAGE CYCLE SIZE 60,490
410 TIME SINCE LAST CYCLE SWITCH 209 DURATION OF LAST CYCLE 114
410 MAX THROTTLE TIME 0 MAX CACHE PERCENTAGE 94
410 HA WRITES 1,522,704,296 RPTD HA WRITES 720,119,145
410 HA DUP. SLOTS 13,513,070 SECONDARY DELAY 323
410 LAST CYCLE SIZE 148,915 DROP PRIORITY 33
410 CLEANUP RUNNING ( Y ) MSC WINDOW IS OPEN ( N )
410 MSC ACTIVE ( Y ) ACTIVE SINCE 12/03/2006 02:11:41
410 CAPTURE TAG C0000000 00003F26 TRANSMIT TAG C0000000 00003F25
410 GLOBAL CONSISTENCY ( Y ) STAR RECOVERY AVAILABLE ( N )
410 -----
410 END OF DISPLAY
```

- **Min Cycle Time**
  - The target minimum cycle time in seconds. SRDF/A will try to execute the cycles in this time interval but the cycles can be longer than the specified value.
  - A cycle can be active until SRDF/A reaches the MAX CACHE PERCENTAGE that it is allowed to use. If the limit of cache is reached, the SRDF/A session is terminated and a bitmap session is activated.
  - When MSC is active, the Min Cycle Time is not used. Rather, the value specified in the parameter MSC\_CYCLE\_TARGET of SRDF Host Component is used.
- **Secondary Consistent**
  - Secondary Consistent is a Y/N flag that indicates whether the secondary side is consistent during SRDF/A operations.
  - Y      SRDF/A is consistent
  - N      SRDF/A is not consistent
  - ?      SRDF/A is not active and the data on the secondary side may or may not be consistent. When SRDF/A is not active, SRDF/A cannot determine the consistency of the secondary side.

**Note:** After reaching a point of consistency, SRDF/A will continue to preserve a consistent copy in the secondary side.

- **Capture Cycle Size**
  - Number of cache slots currently in the active cycle.
- **Transmit Cycle Size**
  - Number of cache slots left in the cycle being transmitted to the secondary side.
- **Average Cycle Size**
  - Average number of cache slots in the past sixteen cycles.
- **Time since last cycle switch**
  - Number of seconds since the last time SRDF/A has cycle switched.
- **Duration of last cycle**
  - Number of seconds the last cycle lasted.
- **Average Cycle Time**
  - Average number of seconds in the past sixteen cycles.

- **Max Throttle Time**
  - Maximum Throttle Time indicates how long SRDF/Asynchronous will slow the host adapters once cache limits are reached. If the value is 0, then once Cache limits are reached, SRDF/Asynchronous is dropped.
  - If the value is 65535, then the host adapters will work at write pending limits speed indefinitely.
  - If any other value, then that is the number of seconds the host adapters will work at write pending limits speed before SRDF/Asynchronous will be dropped.
- **Max Cache Percentage**
  - Maximum Cache Percentage is the percentage of cache that SRDF/Asynchronous is allowed to use. In the initial release of SRDF/A, the Maximum Cache Percentage is 100%.
- **HA Writes**
  - Number of tracks written by the host adapters.
- **RPTD HA Writes**
  - Total number of tracks written multiple times in a cycle by the host adapters.
- **HA DUP SLOTS**
  - Host Adapter Duplicated Slots is the number of times a slot had to be duplicated because it was written to in multiple cycles.
- **Secondary Delay**
  - Secondary Delay is the approximate time the data on the secondary side is behind the primary side.
- **Last Cycle Size**
  - The Last Cycle Size is the size of the complete previous cycle.
- **Cleanup Running**
  - Cleanup Running is a (Y/N) flag:
    - **Y** The secondary side will reject non-SRDF/Asynchronous for a small window of time (approximately 30 seconds). Cleanup only runs immediately after SRDF/Asynchronous goes from Active to Inactive state. Cleanup prevents RDF-RSUM, REFRESH RFR-RSUM, or VALIDATE INVALIDATE from being run on the SRDF/Asynchronous devices.
    - After the cleanup is finished, the RDF-RSUM, REFRESH RFR-RSUM, or, VALIDATE INVALIDATE commands may be run.
    - **N** Cleanup is not running.
- **MSC window is open**
  - The Host Managed Consistency Window is a small time frame that the cycle switch must be run in when running in MSC. When the MSC window is open, all write I/O's to SRDF/Asynchronous primary devices are disconnected. Read I/O's continue to run.

- **MSC Active**
  - Host Active is a (Y/N) flag:
  - **Y** The SRDF/Asynchronous session is part of a multiple SRDF/Asynchronous session group. The multiple SRDF/Asynchronous session group is a group of multiple SRDF/Asynchronous sessions that are having the cycle switch coordinated by the Host.
  - When SRDF/Asynchronous is not active and MSC “Y,” SRDF/Asynchronous is deactivated when MSC was active.
  - **N** SRDF/Asynchronous is not running in MSC mode.
- **Active since**
  - Active since is the date and time that the SRDF/Asynchronous session joined MSC.
- **Capture tag**
  - The Capture Tag is the tag for the data in the capture cycle. The Capture Tag verifies the multiple SRDF/Asynchronous sessions in the MSC group are coordinated.
  - Capture Tag functions in the same manner as cycle number when SRDF/Asynchronous is active and MSC is not active.
- **Transmit tag**
  - The Transmit Tag is the tag for the data in the transmit cycle. The Transmit Tag verifies that the multiple SRDF/Asynchronous sessions in the MSC group are coordinated. When MSC is active, the Transmit Tag functions in the same manner as the cycle number when SRDF/Asynchronous is active and MSC is not active.

#### **4. Service Processor – Event Displays**

The following is a Service Processor Event Display for RA Director 13d. Buried in the display are event codes that can shed light on the SRDF/A processing. These codes are also displayed in the z/OS SYSLOG as part of the IEA480E Service Alert message in the SER=REFCODE=046D field. For example, on line 07 we see 046D that means all RDF links are not operational. Not surprisingly, the next line indicates a CACA20 which usually means a network problem. Some other interesting event codes include:

- E475 – a remote mirrored drive is in a “not ready” state
- E47D – an SRDF adapter link problem (One message per LPAR, per CHPID)
- 0513 – Link Timeout errors
- 053D – Link Resets
- 147E – An RDF link is now operational
- E46E – All RDF links are now operational

D- 13d @[205ECF0,0] 8000;02;ff;0603;10000000 00000000; 0 ; EB, bcm=6

UTILITY EB -- Display EVENTS symptom codes : TIME: DEC/06/05 10:31:58

-----  
\*\* RAF- 13d \*\* LOCAL SHORT TRACE

Legend: PC: Program cntr, C: Cyl, H: Head, x: Repetition cnt  
P: Pipe & Addr, Ed: Err dependent, Id: Dir id

```
00 5A07E3E04 x01 DEC/05/05 20:31:31:769:082 (-14:00:27) PC:00311074
   C:---- H:- P:- 0.00000000 Ed:00000003 00000001 Id:1C IF:-
01 4A6272120 x02 DEC/05/05 20:31:31:107:715 (-14:00:27) PC:002E4A0C
   C: 0 H:0 P:- 0.00000000 Ed:00000000 00000000 Id:1C IF:- DV: 0
02 5A07D3E04 x01 DEC/05/05 20:31:30:959:958 (-14:00:28) PC:00311328
   C:---- H:- P:- 0.00000000 Ed:00000001 00000000 Id:1C IF:-
03 3A6272F21 x01 DEC/05/05 20:31:30:957:677 (-14:00:28) PC:002EADEC
   C: C76 H:8 P:- 0.00000000 Ed:343D0000 00000000 Id:1C IF:- DV:FFF0
04 528046E04 x01 DEC/05/05 20:31:30:485:357 (-14:00:28) PC:00312248
   C:---- H:- P:- 0.00050015 Ed:00000001 00000000 Id:1C
05 528047E04 x01 DEC/05/05 20:31:30:483:253 (-14:00:28) PC:00312248
   C:---- H:- P:- 0.00050015 Ed:00000001 00000000 Id:1C
06 4A63E2120 x01 DEC/05/05 20:31:27:863:163 (-14:00:31) PC:002E4A0C
   C: 0 H:0 P:- 0.00000000 Ed:343D0000 00000000 Id:1C IF:- DV: 20
07 328046D04 x01 DEC/05/05 20:31:27:855:496 (-14:00:31) PC:00312248
   C:---- H:- P:- 0.00050003 Ed:00000001 00000000 Id:1C
08 2A5CACA20 x01 DEC/05/05 20:31:27:630:402 (-14:00:40) PC:00062070
   C:---- H:- P:- 0.00000000 Ed:00000004 00000000 Id:1C
09 5A07D3E04 x01 DEC/05/05 20:31:27:622:255 (-14:00:40) PC:00311328
   C:---- H:- P:- 0.00000000 Ed:00000001 00000000 Id:1C IF:-
0A 3A0FE3F00 x01 DEC/05/05 20:31:27:619:990 (-14:00:40) PC:003224E8
   C: 98A H:A P:0 0.4B0BAB60 Ed:00000000 080000CE Id:1C IF:A DV: 30
0B 4A0053D03 x01 DEC/05/05 20:31:25:744:721 (-14:00:42) PC:002EF9A8
   C: 98B H:3 P:- 0.00000000 Ed:00000000 00000000 Id:1C IF:- DV: 30
0C 3A0051302 x01 DEC/05/05 20:31:25:742:748 (-14:00:42) PC:002EECE4
   C: 98B H:3 P:- 0.00000000 Ed:00000000 00000000 Id:1C IF:- DV: 30
0D 4A0053D03 x01 DEC/05/05 20:31:25:740:165 (-14:00:42) PC:002EF9A8
   C: 98B H:2 P:- 0.00000000 Ed:00000000 00000000 Id:1C IF:- DV: 30
0E 3A0051302 x01 DEC/05/05 20:31:25:748:270 (-14:00:42) PC:002EECE4
   C: 98B H:2 P:- 0.00000000 Ed:00000000 00000000 Id:1C IF:- DV: 30
0F 4A0053D03 x01 DEC/05/05 20:31:25:745:824 (-14:00:42) PC:002EF9A8
```

## 5. Recovery Procedures

Recovery from a dropped SRDF/A MSC session requires execution of detailed procedures.

- Immediately run the MSC CleanUp job (SCFRDFME) at the source, assuming the links are available to the R2s. If the links are down, the MSC CleanUp job must be run at the target site before proceeding. Use an MSC Group name comprised of 8 characters, otherwise you have to embed spaces in CleanUp parameters.

Example:

PGM=SCFRDFME,PARM='Y,0300,SHORT ' with three trailing blanks. Important: *The 0300 UCB address should be an SRDF/A device in any of the MSC Group frames. Do not use the MSC GateKeeper UCB!*

- R2 target side must have a “recovery” or “starter” system to run the MSC CleanUp job, ready R2 devices, make R2s read/write, and bring R2s online in the event of a real disaster.
- Three MSC recovery scenarios:

Case 1 - All R2 receive cycles have same tag and are completed – MSC action Commit Receive Cycle

Case 2 – (the most likely) - All R2 receive cycles have same tag but one or more but not all are complete – MSC Action Discard Receive Cycle

Case 3 - Apply cycle tags of some R2 Symms match receive cycle tags of one or more other R2 Symms (not all R2 receive cycles were committed) – MSC Action Commit and Discard Receive Cycle

### 4. SRDF/A Symmetrix error codes:

- CACA10 – Source system max write pending limit reached. Could be a cache full or R2 target destaging issue. Review the Transmit and Restore Delay messages.
- CACA11 – SRDF/A dropped after throttling the host longer than the throttle time limit
- CACA20 – SRDF/A target device has become TNR on the link; probably link issue.
- CACA30 – Target SRDF/A device made TNR on the link at the R1 side
- CACA40 – All links to target frame were lost for a period of time greater than set in the BIN link limbo parameter. Default is ten seconds which can be increased up to 60 seconds if you receive 7D3 and 7E3 messages. Don't increase from the ten seconds default if using SRDF/S.
- CACA41 – SRDF/A dropped after exceeding Transmit Idle timer
- CACA50 – SRDF/A MSC mode cycle switch window did not close within 5 seconds. Check to make sure SCF is started with SUB=MSTR and dispatch priority =SYSSTC
- CACA60 – SRDF/A timed out waiting on HA's in the state machine (HA's were writing into the inactive cycle and exceeded the time limit)
- CACA62 – R1 issued a commit to R2 side, but found the R2 side's SRDF/A is inactive
- CACA64 – The R1 side found R2 side's SRDF/A inactive or the R2 side found the R1 side's SRDF/A inactive.
- CACA70 – Power down caused SRDF/A to drop
- CACA71 – No online RDF Directors

- CACA8x – When running DCP with 5772 or >, exhausting cache will not produce a CACA10 message. Instead a CACA8x is produced with x = partition Id.

## 6. Performance Considerations

You choices in an SRDF/A environment will impact performance. For example:

- Should you use Adaptive Copy Write or Adaptive Copy Disk? Because Adaptive Copy Write is competing for the same cache slots as SRDF/A, the general consensus is Adaptive Copy Disk is better. Also, don't mix Adaptive Copy Write and Adaptive Copy Disk in the same box.
- With Adaptive Copy, you probably want to use a QOS setting to ensure the initial copy does not impact production. But what is the right value? On a DMX2, a QOS=2 does about 220 IOP/s per RA while a QOS=4 does only 60 IOP/s per RA.
- Be careful when proposing a fan-in solution that has multiple source frames going to one target. Scrutinize the target cache, disk, and RAID protection to ensure the target can handle all source frames peaking at the same time.
- Though supported, avoid putting SRDF/A and SRDF/S workload on the same RAs. The SRDF/S workload has a higher priority and could cause SRDF/A to backup and drop.
- With SRDF/A, how should you handle Page Packs, SYSDA, SORTWORK, TEMP packs. etc.? You still do not want to replicate Page Packs because the SRDF/A window could cause delays in paging I/O that might impact system performance. Since you need to initialize the VOLSERS for these volumes at the R2 recovery site, most shops put these volumes in their own RAGroup and only put them into Adaptive Copy on the weekend or when they do an IPL. Since Page Packs are replicated only weekly, make sure the customer does not put application datasets on the Page Packs!
- Couple dataset volumes have special requirements. The System Logger CDS (LOGR) used by CICS journaling must be replicated for recovery. All other couple datasets (SYSPLEX, ARM, SFM, WLM, OMVS, and CFRM) are required but have their contents re-created at IPL. Therefore, they can be handled like Page Packs using Adaptive Copy only.
- If you have multiple RF links and they are not balanced, make sure Flow Control Resource Shortage monitoring is disabled on each RF port. Invoke the E7, LINK, FLOW and A2, F3 RDF statistics InLine commands

## 7. Network Considerations

- The Link Limbo default defined in the BIN file is ten seconds. If you receive 7D3 and 7E3 error messages, you should turn on (Link Resiliency) Transmit Idle. *Warning:* don't increase Link Limbo from the ten seconds default if using SRDF/S.
- You must have a switch between the Symmetrix and DWDM equipment when using GigE, otherwise you will get numerous link, reset, TNR volumes, and phone home errors. You will also want to rate-limit GigE so you don't overrun the bandwidth.
- Gather available network charts and statistics to correlate Symmetrix information with network activity.

- Events which might cause bandwidth issues:
  - Not enough link bandwidth or other competing activity when using shared bandwidth
  - Bursts of application I/Os; remember, a lot of spikes can hide in a 15 minute interval
  - Long peak periods (hours between 10 PM and 4 AM are typical)
  - Unexpected growth of workload (new application processing, accounts, subsidiaries)
  - New CPUs with more processing MIPS
  - Going from ESCON to FICON
  - Not allocating enough RAs
  - Adding more disk capacity in the source frames
  - Changing the network hardware or configuration
  - Implementing network features such as data encryption

### **8. *Balanced SRDF Links***

Unbalanced SRDF links can jeopardize SRDF/A performance. This is not always obvious so you must review STP or WLA data to see if multiple SRDF links in a box are sending approximately the same amount of MB per second. If, for example, you see two links sending 40 MB/s and the other two sending 12 MB/s, you have unbalanced links that are reducing the SRDF/A throughput. There are three Primus cases you want to check which may alleviate the imbalance.

1. Make sure Echo is OFF. Primus EMC133253 – Fix 31144 in 5671.58.64 and 5771.92.99 resolves this issue. On each RF, issue the FC, HRBT, 'OFF' to disable echo frames.
2. Make sure Flow Control=Enabled Primus. EMC84712 – Fix 23733 in 5671.45.52 and 5771.81.90 resolves this issue. On each RF issue the D0,TUNE,'JFC','ON' to enable Flow Control.
3. Make sure Resource Shortage=Disabled Primus. EMC134213 – Fix 31214 in 5671.58.64 resolves this issue for DMX or DMX2. On each RF issue the D0,TUNE,'JFCR','OFF' to disable Resource Shortage.

### **9. *Link Resiliency (Transmit Idle)***

A consistent point-in-time copy at the Disaster Recovery site is necessary to ensure that SRDF/A is running continuously. But it's not unusual for new SRDF/A implementations to experience some network instability during initial testing or during adverse weather conditions. Normally, if RDF experiences a "Links Lost" 046D Symmetrix Event Code, SRDF/A will drop immediately with a CACA20 RDF mirror Not Ready or CACA40 All RDF links are down.



Link Resiliency allows the SRDF/A RDF group to enter transmit idle state following the expiration of the Link Limbo timer. In transmit idle state, the remote SRDF mirror remains ready on the link even though the link is physically down long as there is sufficient cache resources for SRDF/A to continue operation. Data is not sent to the remote devices, but SRDF/A remains active. Therefore, RPO (Recovery Point Objective) increases while in transmit idle state.

#### **Scenario of events with Transmit Idle activated:**

1. Transmit Idle must be Online for both the source and target frames.
2. All Customer SRDF links drop and are unavailable.
3. The link limbo timer begins. If BIN default was changed, must make Link Limbo = 10 seconds.
4. When the Link Limbo timer expires, Transmit Idle is activated for SRDF/A RDF group(s) SCF1586I.
5. All primary devices in SRDF/A RDG group(s) remain ready on the remote link.
6. During Transmit Idle time, the R1 active cycle (Capture delta set) continues to capture data for new host writes for as long as cache is available.
7. Data remains in the R1 inactive cycle (Transmit delta set) since it can't be propagated to the R2.
8. To stop SRDF/A MSC processing, issue the #SC SRDFA,LCL(8800,08),DROP\_SIDE command to one of the source frames. MSC will detect the Transmit Idle state and drop other frames in the MSC Group. F EMCSCF,MSC,PENDDROP and #SC SRDFA,LCL(8800,08),PEND\_DROP will be rejected because there is no cycle switching while Transmit Idle is active.
9. If the links do not recover, SRDF/A eventually drops due to cache thresholds being reached.
10. If the links recover, SRDF/A continues to operate provided SRDF/A is active on both source and target sites.

#### **SRDF Host Component Minimums V5.3(57), V5.4(04), V5.5 and commands:**

1. Turn on Transmit Idle local: #SC SRDFA,LCL(8800,08),TRANSMIT\_IDLE,ONLINE
2. Turn off Transmit Idle local: #SC SRDFA,LCL(8800,08),TRANSMIT\_IDLE,OFFLINE
3. Drop SRDF/A if Transmit Idle is active: #SC SRDFA,LCL(8800,08),DROP\_SIDE (Issue from R1 side)

Once Transmit Idle is ONLINE, it stays ONLINE even between an MSC,PENDDROP or SRDF/A drop.

## 10. Delta Set Extension (5772+ Feature)

SRDF/A and remote replication stops require manual restart procedures when the amount of cache needed by SRDF/A exceeds the amount of available cache. This can be caused by write spikes in the workload, new SRDF/A devices added to link, or network issues. The Link Resiliency (Transmit Idle) feature was provided allowing SRDF/A to remain active during an “All RDF links lost” event as long as cache limitations are not exceeded. The second line of defense is Delta Set Extension (DSE) which prevents SRDF/A from dropping by allowing MSC to offload (spillover) data that would normally stay in cache to preconfigured DSE volumes.

DSE is not designed to compensate for inadequate network bandwidth or unbalanced SRDF/A configurations (frames, cache, number of disks, capacity of disks, or RAID protection). DSE must be defined and activated on both sides of the RDF relationship. If DSE is enabled, the Link Resiliency feature is automatically enabled. Each RA Group can be assigned its own cache operational threshold which, when exceeded, will trigger activation of the DSE process. Therefore, RA Groups with higher thresholds will keep their data in cache longer at the expense of RA Groups with lower thresholds. RPO (Recovery Point Objective) increases while in DSE state but host writes should not be impacted.

There are some configuration considerations with DSE:

1. DSE can potentially elongate the Capture, Transmit, Receive and Apply SRDF/A cycles, All source and target frames in an MSC Group must have defined and activated DSE.
2. Paging tasks run on DAs (16% max per CPU) and RAs (25% max per CPU). Current DAs and RAs should be less than 50% utilized to successfully implement DSE.
3. Get DSE messages if running MSC. A single SRDF/A pair without MSC will not produce DSE messages
4. For simplicity, make all parameter values the same for all frames in the same MSC Group

DSEPOOL devices are configured like TimeFinder Virtual SNAP devices with the following considerations:

1. SRDF/A DSEPOOLS can have up to four unique pool types; FBA, CKD3390, CKD3380, or AS400.
2. DSEPOOLS can be shared among multiple SRDF/A RAGroups
3. A SRDF/A session can be associated with either zero or one DSEPOOL of each type
4. Devices can be added to the DSEPOOLS from the DEFAULT\_POOL and removed from the DSEPOOL once the pool is disabled and drained of SRDF/A tracks

5. DSEPOOL devices are not host accessible but do consume a Symmetrix device number
6. The devices should be spread across all DAs
7. Since DSE performance depends upon the speed of disk writes, DSEPOOL devices should be defined as RAID 1, though RAID 5 can also be used but not RAID 10 or RAID 6.

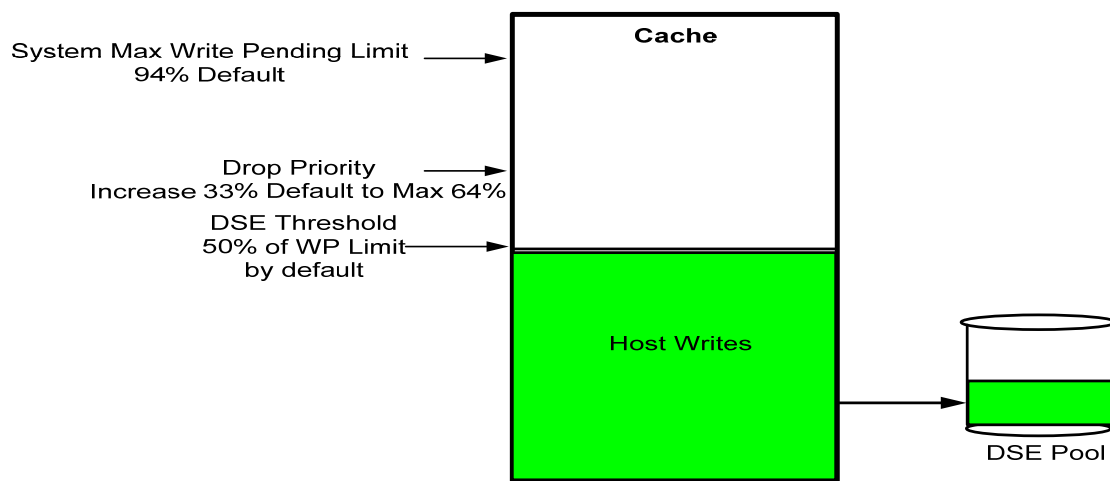
If the DMX-3 running 5772 code is coming close to being cache-constrained while the links are lost, and DSE (Delta Set Extension) is enabled along with Transmit Idle, then DSE de-stages the SRDF/A cache slots to DSE pool devices allowing SRDF/A to remain active.

## Reference Materials

1. "SRDF/A Reserve Capacity Enhancement: Transmit Idle and Delta Set Extension for Mainframe Environments" Technical Notes P/N 300-004-9914 Rev A01 May 16, 2007
2. Various PowerPoint presentations from 2007 EMC World and 5772 training

## What Triggers DSE?

DSE will be triggered when the total number of cache slots used by all SRDF/A sessions in the system, whether the sessions are using DSE or not, expressed as a percentage of the system write pending limit, exceeds the Session Page-out Threshold of the session.



## Determining Number of DSEPOOL Devices

Determining the exact amount of DSEPOOL capacity requires many details such as the number of random and sequential write per second, random and sequential block size, write folding, excess bandwidth, and more. A simpler approach will provide adequate DSEPOOL capacity for processing:

- 1) Reviewing the following #SQ SRDFA 8800 display, determine the largest Capture Cycle Size which is 371,180 slots which caused our DMX3 to drop SRDF/A because Cleanup is running.
- 2) Multiply 371,180 X 56,664 = 21,032,543,520 bytes in the Capture Cycle
- 3) A MOD3 has 2,838,016,440 bytes divided into 21,032,543,520 = 7.4 MOD3s required to hold cycle
- 4) The DSEPOOL requires sufficient capacity to hold both the Capture and Transmit Cycles in the source frame. So, at minimum, multiple 7.4 MOD3s X 2 = 14.8 MOD3s are required.

```

01.48.22 STC22138 EMCMN00I SRDF-HC : (5177) #SQ SRDFA,8800
01.48.22 STC22138 EMCQR00I SRDF-HC DISPLAY FOR (5177) #SQ SRDFA,8800 410
410 MY SERIAL # MY MICROCODE
410 -----
410 000000001248 5772-83
410
410 MY GRP ONL PC OS GRP OS SERIAL OS MICROCODE SYNCHDIR FEATURE
410 -----
410 LABEL TYPE AUTO-LINKS-RECOVERY LINKS_DOMINO MSC_GROUP
410 -----
410 83 Y F 08 000000001248 5772-83 G(R1>R2) SRDFA I MSC
410 Box83 STATIC AUTO-LINKS-RECOVERY LINKS-DOMINO:NO (MSCGRP01)
410 -----
410 PRIMARY SIDE: CYCLE NUMBER 7,811 MIN CYCLE TIME 30
410 SECONDARY CONSISTENT ( ? ) TOLERANCE ( N )
410 CAPTURE CYCLE SIZE 371,180 TRANSMIT CYCLE SIZE 0
410 AVERAGE CYCLE TIME 57 AVERAGE CYCLE SIZE 60,490
410 TIME SINCE LAST CYCLE SWITCH 209 DURATION OF LAST CYCLE 114
410 MAX THROTTLE TIME 0 MAX CACHE PERCENTAGE 94
410 HA WRITES 1,522,704,296 RPTD HA WRITES 720,119,145
410 HA DUP. SLOTS 13,513,070 SECONDARY DELAY 323
410 LAST CYCLE SIZE 148,915 DROP PRIORITY 33
410 CLEANUP RUNNING ( Y ) MSC WINDOW IS OPEN ( N )
410 SRDFA TRANSMIT IDLE ( N ) SRDFA DSE ACTIVE ( N )
410 MSC ACTIVE ( Y ) ACTIVE SINCE 12/03/2006 02:11:41
410 CAPTURE TAG C0000000 00003F26 TRANSMIT TAG C0000000 00003F25
410 GLOBAL CONSISTENCY ( Y ) STAR RECOVERY AVAILABLE ( N )
410 -----

```

The size of the Symmetrix needs to be considered because we want a MOD3 behind every DA CPU to provide maximum DSEPOOL performance. Ideally the DSEPOOL should be RAID 1, as follows:

- 1) DMX1500: 1 DA pair, 8 Drive Loops = 8 usable x 2 rows of MOD3s = 16 devices mirrored
- 2) DMX2500: 2 DA pair, 16 Drive Loops = 16 usable x 1 row of MOD3s = 16 devices mirrored or,  
16 usable x 2 rows of MOD3s = 32 devices mirrored
- 3) DMX3500: 3 DA pair, 24 Drive Loops = 24 usable x 1 row of MOD3s = 24 devices mirrored
- 4) DMX4500: 4 DA pair, 32 Drive Loops = 32 usable x 1 row of MOD3s = 32 devices mirrored

Remember, we need 14.8 DSEPOOL devices. If the frame is a DMX1500, the 16 devices would probably be sufficient. But with a DMX2500, 16 devices are too close, so another row of MOD3s would provide 32 devices to handle probable growth. Total DSEPOOL capacity cannot exceed 512 times the amount of usable cache. If a box has 16 gb of usable cache X 512 = 8TB of spillover maximum.

### Defining the DSEPOOL Devices

The DSEPOOL devices can be any size volumes, like MOD3s, specifically defined to hold spillover tracks from cache once the DSE Threshold is reached. DSEPOOL devices should be on the fastest drives in the box.

When the BIN is loaded, let's assume we have (32) CKD devices at Symmetrix device numbers 0100-011F which will be in the 'DEFAULT\_POOL'. One pool must be created, if applicable, for each emulation type 3380, 3390, AS400 and FBA. To create the unique DSEPOOL pool for 3390 CKD devices, there are four tasks for both local and remote frames:

- 1) Create a new pool name to use with DSE called MY3390POOL
- 2) Add the (32) devices from the 'DEFAULT\_POOL' to MY3390POOL
- 3) Enable the (32) devices in MY3390POOL
- 4) Display volumes in MY3390POOL

You can complete these steps by executing the ResourcePak Base program ESFGPMBT that can be found as member CNFGPOOL in SCFSAMPLIB. In the following JCL example, 8800 is a Gatekeeper pointer to the box. MY3390POOL can be any 12 characters. Do not use spaces.

```
//jobname      JOB (EMC),pgmmrname,CLASS=A,MSGCLASS=X
//QCOPYRUN    EXEC PGM=ESFGPMBT,REGION=4M
//STEPLIB     DD  DISP=SHR,DSN=ds-prefix.LINKLIB
//SCF$V570    DD  DUMMY
//SYSPRINT    DD  SYSOUT=*
//SYSUDUMP    DD  SYSOUT=*
//SYSOUT      DD  SYSOUT=*
//GPMPRINT    DD  SYSOUT=*
//QCOUTPUT    DD  SYSOUT=*
//EMCQCAPI    DD  SYSOUT=*
//GPMINPUT    DD  *
CONFIGPOOL CREATE (LOCAL(UNIT(8800)) TYPE(DSEPOOL) POOL(MY3390POOL))
CONFIGPOOL ADD (LOCAL(UNIT(8800)) TYPE(DSEPOOL) POOL(MY3390POOL) -
                DEV(0100-011F) MEMBERSTATE(ENABLE))
CONFIGPOOL DISPLAY(LOCAL(UNIT(8800)) TYPE(DSEPOOL) POOL(MY3390POOL))
CONFIGPOOL CREATE (REMOTE(UNIT(8800)) TYPE(DSEPOOL) POOL(MY3390POOL))
CONFIGPOOL ADD (REMOTE(UNIT(8800)) TYPE(DSEPOOL) POOL(MY3390POOL) -
```

DEV(0100-011F) MEMBERSTATE(ENABLE))  
 CONFIGPOOL DISPLAY(REMOTE(UNIT(8800)) TYPE(DSEPOOL) POOL(MY3390POOL))

**Output from Local CONFIGPOOL DISPLAY:**

```
CONFIGPOOL DISPLAY(LOCAL(UNIT(8800)) POOL(MY3390POOL) TYPE(DSEPOOL))
LOGPOOL DEVICE INFORMATION FOR LOGPOOL - MY3390POOL
  -DEVICE-  -STATUS-  TYPE  --USED--  --FREE--  -DRAIN?-
00000100  ACTIVE    CKD    0    50145    NO
00000101  ACTIVE    CKD    0    50145    NO
00000102  ACTIVE    CKD    0    50145    NO
00000103  ACTIVE    CKD    0    50145    NO
Etc. to 011F
```

Congratulations. You now have a usable DSE Pool for 3390 cache spillover.

**Other CONFIGPOOL Commands using the Same Syntax**

- 1) **DRAIN:** copy tracks from a single DSEPOOL device to other devices in the same DSEPOOL, cannot be done with Inlines commands.
- 2) **UNDRAIN:** change a single draining DSEPOOL device back to Active
- 3) **DISABLE:** change DSEPOOL from Active to Inactive
- 4) **ENABLE:** change DSEPOOL from Inactive to Active
- 5) **REMOVE:** remove Inactive devices, must DISABLE first, from DSEPOOL
- 6) **DELETE:** delete DSEPOOL, must REMOVE DSEPOOL first

To DELETE an active DSEPOOL you have to DRAIN, DISABLE, REMOVE, and then DELETE.

**Define SCF Threshold Monitoring Parameters**

There are special ResourcePak Base SCF parameters that allow you to define a cache percentage at which a DSE message or user routine will be invoked. In the following examples, the cache utilization will be checked:

- 1) Every 5 minutes and message issued once cache utilization is between 60-65%
- 2) Every 2 minutes and UserExit invoked multiple times if cache utilization is between 66-90%

```
SCF.SDV.LIST=ENABLE
SCF.SDV.01.LIST=PERCENT=(60,65)
SCF.SDV.01.LIST=DURATION=5
SCF.SDV.01.LIST=ACTION=MESSAGE (User Appended String)
SCF.SDV.01.LIST=FREQ=ONCE
SCF.SDV.01.LIST=ENABLE
SCF.SDV.01.LIST=GATEKEEPER(8800)
SCF.SDV.02.LIST=PERCENT=(66,90)
SCF.SDV.02.LIST=DURATION=2
```

SCF.SDV.02.LIST=ACTION=USEREXIT(Helpdesk Exit)

SCF.SDV.02.LIST=FREQ=REPEAT

SCF.SDV.02.LIST=ENABLE

SCF.SDV.02.LIST=GATEKEEPER(8800)

Setting the Threshold percentage can be difficult. You want a large enough cache percentage to allow normal SRDF/A processing without invoking DSE prematurely. Twenty percent may be appropriate for testing, but you don't want 20% for production because tracks would quickly go to the DSEPOOL. At the same time, you don't want to make the threshold 90% because by the time DSE is invoked, there isn't enough time for DSE to free the necessary cache slots before cache is full. You may want to use a value between 50-60%.

### **Activate DSE for SRDF/A and MSC**

There are several SRDF Host Component commands which activate and control DSE. In the following examples, 8800 is the GateKeeper and 08 is the RDFGroup. Commands can be executed while SRDF/A and MSC are active.

1. Activate DSE: **#SC SRDFA\_DSE,LCL(8800,08),ACT**
2. Change DSE threshold from default 50% to 60%: **#SC SRDFA\_DSE,LCL(8800,08),THRESHOLD,60**
3. Make DSE persistent: **#SC SRDFA\_DSE,LCL(8800,08),AUTO\_ACT,ON**

The default for DSE is OFF. The AUTO\_ACT, ON will automatically activate DSE when SRDF/A starts, even between drop/resume processing. Perform an AUTO\_ACT, OFF to stop the automatic DSE activation.

## Changes to SRDF/A Query Display

The normal SRDF/A query will now include new a DSE field (red, in the display below) that indicates DSE status. Notice that SRDFA Transmit Idle is also available but has not been invoked because the feature code is SRDF A MSC, not SRDF T MSC.

```
08.38.19 S0041609 EMCMN00I SRDF-HC : (574) #SQ SRDFA,LCL(8800,08)
08.38.19 S0041609 EMCQR00I SRDF-HC DISPLAY FOR (574) #SQ SRDFA,LCL(8800,08) 491
491 MY SERIAL # MY MICROCODE
491
491 000000001248 5772-83
491
491 MY GRP ONL PC OS GRP OS SERIAL OS MICROCODE SYNCHDIR FEATURE
491 -----
491 LABEL TYPE AUTO-LINKS-RECOVERY LINKS_DOMINO MSC_GROUP
491 -----
491 08 Y F 08 000000001248 5772-83 (R1>R2) SRDFA A MSC
491 BOX83 DYNAMIC AUTO-LINKS-RECOVERY LINKS-DOMINO:NO (MSCGRP01 )
491 -----
491 PRIMARY SIDE: CYCLE NUMBER 15,451 MIN CYCLE TIME 30
491 SECONDARY CONSISTENT ( Y ) TOLERANCE ( N )
491 CAPTURE CYCLE SIZE 0 TRANSMIT CYCLE SIZE 0
491 AVERAGE CYCLE TIME 22 AVERAGE CYCLE SIZE 0
491 TIME SINCE LAST CYCLE SWITCH 12 DURATION OF LAST CYCLE 15
491 MAX THROTTLE TIME 0 MAX CACHE PERCENTAGE 94
491 HA WRITES 3,132,696 RPTD HA WRITES 4,930
491 HA DUP. SLOTS 397 SECONDARY DELAY 27
491 LAST CYCLE SIZE 0 DROP PRIORITY 33
491 CLEANUP RUNNING ( N ) MSC WINDOW IS OPEN ( N )
491 SRDFA TRANSMIT IDLE( Y ) SRDFA DSE ACTIVE ( Y )
491 MSC ACTIVE ( Y ) ACTIVE SINCE 10/24/2006 08:34:11
491 CAPTURE TAG C0000000 00000011 TRANSMIT TAGC0000000 00000010
491 GLOBAL CONSISTENCY ( Y ) STAR RECOVERY AVAILABLE( N )
491 -----
```

### SRDFA DSE ACTIVE (Y)

(N) = DSE is not active

(Y) = DSE is active

**SRDFA A MSC** = SRDF/A with MSC is running normally or

**SRDFA T IDLE** = SRDF/A is running with a single frame and Transmit Idle is active or

**SRDFA I MSC** = SRDF/A is Inactive with MSC or

**SRDFA T MSC** = SRDF/A is running with MSC and Transmit Idle is active or

**SRDFA T STAR** = SRDF/A is running with STAR and Transmit Idle is active

## DSE Query Display

The DSE query provides specific information about DSE activity. In this example, the DSE was triggered by the Capture Cycle in the source box.



EMCQR00I SRDF - HC DISPLAY FOR (74) ## SQ SRDFA\_DSE,LCL(8800,08) 495

MY SERIAL # MY MICROCODE

-----  
000000001248 5772 -83

MY GRP	ONL	PC	OS GRP	OS SERIAL	OS MICROCODE	SYNCHDIR	FEATURE
				AUTO - LINKS	- RECOVERY	LINKS_DOMINO	MSC_GROUP
08	Y	F	F1	000000001248	5772 -83	-	SRDFA A MSC
BOX83		DYNAMIC		AUTO LINKS RECOVERY		LINKS_DOMINO:NO	(MSCGRP01)

-----  
PRIMARY SIDE: CYCLE NUMBER 7,913 SRDFA DSE ACTIVE ( Y )  
THRESHOLD PERCENTAGE 60 SRDFA DSE AUTO ACTIVATE ( Y )  
CAPTURE CYCLE SIZE 17,876 TRANSMIT CYCLE SIZE 0  
DSE USED TRACKS 5,602 DSE USED TRACKS 0  
DSE MDATA TRACKS 16 DSE MDATA TRACKS 0

-----  
FBA POOL NAME DSE 3390 POOL NAME DSE MY3390POOL  
AS400 POOL NAME DSE 3380 POOL NAME DSE

-----  
END OF DISPLAY

### SRDF Host Component Minimum V5.5+ and commands:

Define four possible DSE Pool types; must be done for each source and target frame.

The remote version RMT of these commands has the same syntax.

1. #SC SRDFA\_DSE,LCL(8800,08),3390\_POOL,P(MY3390POOL) – get message EMCCRA3R
2. #SC SRDFA\_DSE,LCL(8800,08),3380\_POOL,P(MY3380POOL) – get message EMCCRA4R
3. #SC SRDFA\_DSE,LCL(8800,08),A400\_POOL,P(MYA400POOL) – get message EMCCRA5R
4. #SC SRDFA\_DSE,LCL(8800,08),FBA\_POOL,P(MYFBAPOOL) – get message EMCCRA2R

Purge four possible DSE Pool types; must be done for each source and target frame.

1. #SC SRDFA\_DSE,LCL(8800,08),3390\_POOL,P()
2. #SC SRDFA\_DSE,LCL(8800,08),3380\_POOL,P()
3. #SC SRDFA\_DSE,LCL(8800,08),A400\_POOL,P()
4. #SC SRDFA\_DSE,LCL(8800,08),FBA\_POOL,P()

Activate and Deactivate DSE; must be done for each source and target frame.

1. #SC SRDFA\_DSE,LCL(8800,08),ACT - message EMCCRA7R is generated
2. #SC SRDFA\_DSE,LCL(8800,08),DEACT - message EMCCRA8R is generated

Set the DSE cache threshold; must be done for each source and target frame.

1. #SC SRDFA\_DSE,LCL(8800,08),THRESHOLD,70

Note: Percentage can be 20-100% with 50% as the default, message EMCCRA6R is generated

Enable and disable DSE Auto-Activation; must be done for each source and target frame.

1. #SC SRDFA\_DSE,LCL(8800,08),AUTO\_ACT,ON - message EMCCRA9R generated
2. #SC SRDFA\_DSE,LCL(8800,08),AUTO\_ACT,OFF - message EMCCRA9R generated

Note: Default is OFF; when ON, DSE is persistent through drop/resume processing

### **Resource Pak Base Minimum 5.7+ and SCF messages:**

1. EMCU00 through EMCU05 messages for batch pool utility

### **Symmetrix Event Codes for DSE**

1. A5.60CA through A5.6ACA

## **Operational Considerations:**

1. SRDF/A DSE is not supported on ESCON RAs; Fiber RFs and GigE REs.
2. DSE can be used with Dynamic Cache Partitioning (DCP) in 5772.83 and above.
3. If SRDF/A drops with DSE active, an extra cleanup pass is needed, extending cleanup processing.
4. DSE doesn't require MSC to be running.
5. Link Resiliency can run without DSE, but Link Resiliency will be enabled if using DSE.
6. If DSE is invoked, DSE must be active for all MSC SRDF/A sessions, otherwise MSC will fail.
7. MSC cycle switching will not occur until the Transmit and Receive cycles are empty.
8. Must execute the DSEPOOL and activate commands on both source and target sides.
9. If R2 side is too slow, DSE will exaggerate the problem.

## **11. STAR Implementation Impact on SRDF/A**

SRDF/A and MSC processing is exactly the same under STAR (Symmetrix Triangular Asynchronous Replication) but there are several additional considerations which will impact SRDF/A. Specifically:

1. Must run MSC even if only a single frame pair configuration
2. STAR automation (EMCSPE) requires definition of GNS (Global Name Services) RDF Groups. Also make sure Alias definitions used by STAR automation are in Host Component INI file.
3. STAR requires a License Feature Code (LFC) defined in the SCF INI file
4. Will have three levels of consistency; lowest to highest: SRDF/A, MSC, and finally STAR
5. Dynamic RDF devices are required, define as DRX volumes making them R1- and R2-capable
6. Both STAR links from the primary data center cannot be SRDF/A; one link has to be SRDF/S
7. Only one SRDF/STAR configuration per MSC started task
8. Autoswap is not supported with STAR; GDDR is required

## Required STAR Software Levels

1. SRDF Host Component V5.5.0+  
**Note:** In V5.4+ module, EHCMSME (SRDF/A Recovery) and EHCMSM6 (STAR Recovery) in HC.LINKLIB were introduced with an alias to the old SCFRDFME and SCFRDFM6 modules.  
Set ALLOW\_CRPAIR\_NOCOPY=YES
2. ResourcePak Base V5.7+  
Set SCF.MSC.ADCOPY.ONDROP=YES to put all MSC Group devices into Adaptive Copy Disk
3. Consistency Groups V6.3+
4. TimeFinder/Clone Mainframe SNAP V5.7+ and/or
5. TimeFinder/Mirror V5.5+
6. Enginuity 5671, 5771 or 5772+

## Verify BIN File Settings for STAR

1. Switched RDF Configuration State is Enabled
2. Concurrent RDF Configuration State is Enabled
3. Dynamic RDF Configuration State is Enabled
4. Concurrent Dynamic RDF Configuration is Enabled
5. RDF Data Mobility Configuration State is Disabled
6. RDF Directors are either Fibre-Switched or GigE, no ESCON
7. For SRDF Group Level Settings:
  - a. Prevent Auto Link Recovery is Enabled
  - b. Prevent RAs Online Upon Power On is Enabled