



Global Distributed Storage Engineering
Citi Architecture & Technology Engineering

SRDF/Star for Open Systems

GDSE-06050-SG

**Solutions Guide
Version 2.0**

VERSION CONTROL

Version and Date of Issue	Author	Comment
V0.5 11th December 2006	Chris Clayton	First draft
V0.6 20 th December 2006	Chris Clayton	Incorporated changes requested by Ludwig Derodel under Peer Review
V1.0 24 th January 2007	Chris Clayton	Included Peer review changes requested by Ludwig Derodel. Also added RFC comments from Kevin Leah (EMEA) & Javed Dad (EMEA).
V2.0 25th April 2008	Chris Clayton Reviewed by Pat Fogarty	Updated for support with new 5772 microcode features. Numerous other changes to Best Practice

CONTENTS

Legal Notice	7
Contact	7
Latest Versions	7
1 INTRODUCTION	8
1.1 Intended audience	8
1.2 Out of scope	8
1.3 Important Related Documentation	8
1.4 Changes in V2.0	9
2 INTRODUCTION TO SRDF/STAR	11
2.1 What is SRDF/Star?	11
2.2 SRDF/Star vs SRDF/A	13
2.2.1 Similarities between SRDF/A and SRDF/Star	13
2.2.2 Differences between 2-site SRDF/A and 3-site SRDF/Star	13
2.3 Applications for which SRDF/Star is of benefit	16
2.3.1 Applications which need long-distance zero data loss COB when production site lost	16
2.3.2 Applications which must have on-going COB after loss of production site	16
2.4 Applications for which SRDF/Star is <u>not</u> suitable	16
2.4.1 Applications with multiple data servers	16
2.4.2 Applications with changing storage requirements	17
2.4.3 Cost-sensitive applications	17
2.4.4 Applications that are not already in production	17
3 PREPARATION FOR AN SRDF/STAR IMPLEMENTATION	18
3.1 Complexity and scalability warnings	18
3.1.1 Complexity	18
3.1.2 Scalability limitations	18
3.2 Major Pre-requisites	19
3.3 Overview of concurrent SRDF setup	19
3.4 Planning and Preparation for an SRDF/A implementation – SRDF/Star updates	22
3.4.1 Collection of Basic Application requirements for Star	22
3.4.2 RDF group naming/allocation conventions for Star	23

3.4.3	EMC Solution Qualification (SQ) for Star	25
3.4.4	WAN requirements for Star	25
3.4.5	Network hardware for Star	26
3.4.6	Management server requirements for Star	26
3.4.7	Storage Hardware Requirements for Star	28
3.4.8	DMX bin-file configuration for Star	29
3.5	SRDF/A configuration – SRDF/Star updates	31
3.5.1	Management Server Configuration	31
3.5.2	Network Equipment Configuration	32
3.5.3	Dynamic RDF configuration	32
3.5.4	Point-in-Time (PiT) copies for SRDF resynchronization	33
3.5.5	Device Group definition	36
3.5.6	Enabling SRDF/A replication	36
3.5.7	SRDF/A Replication Monitoring	36
3.6	Concurrent SRDF configuration – SRDF/Star updates	37
3.6.1	Licensing for SRDF/Star	37
3.6.2	Device Group definitions in SRDF/Star environments	37
3.6.3	3-Site SRDF/Star RDF group requirements	38
4	SETTING UP SRDF/STAR	39
4.1	Example 1: Setting up SRDF/Star with one triangle per session	39
4.1.1	Step 1 – Create Concurrent SRDF environment	40
4.1.2	Step 2 – Creating SRDF/Star Composite Groups	42
4.1.3	Step 3 – Creating the SRDF/Star options file	42
4.1.4	Step 4 – Performing the symstar setup operation	42
4.1.5	Step 5 – Creating R2 Composite groups	43
4.1.6	Step 6 - Adding BCV devices to the SRDF/Star configurations	43
4.1.7	Step 7 – Bringing up the SRDF/Star sites for normal operation	44
4.2	Example 2: Configuration SRDF/Star with multiple triangles per session	44
4.3	High-availability of the SRDF/Star environment	45
4.3.1	Star Control server high-availability	45
4.3.2	Storrdfd daemon high-availability	45
5	SITE SWITCHING (FAILOVER) OPERATIONAL PROCEDURES	46
5.1	Introduction	46
5.1.1	Nomenclature	46
5.1.2	Other documentation	46
5.1.3	Site switching multiple Star sessions	46
5.1.4	Notes on impact of Transmit Idle in SRDF/Star configurations	47
5.2	Site B as bunker only – site switching scenarios	47
5.2.1	Server and application level failover/switch scenario	48
5.2.2	Site A failure scenario	48
5.2.3	Regional disaster scenario	49
5.3	Controlling SRDF/Star	50

5.3.1	symstar	50
5.3.2	symstar arguments	50
5.3.3	SRDF/Star States	51
5.4	Using BCVs with Star operations	51
5.4.1	General Best Practice	51
5.4.2	Why make a point-in-time copy?	51
5.5	Executing a planned switch	52
5.5.1	Planned switch workload to Async target site (standard process)	52
5.5.2	Planned switch back from Async target to workload site	54
5.5.3	Planned switch workload to Sync target site (non standard process)	55
5.5.4	Planned switch back from Sync target to workload site	56
5.6	Responding to disaster faults – “unplanned” switches	57
5.6.1	Determining from which site to keep data	57
5.6.2	Scheduling multiple switch operations	58
5.6.3	Switch workload to Async target site (Site C), keeping Sync (Site B) data	58
5.6.4	Switch workload to Async target site (Site C), keeping Async (Site C) data	60
5.7	Switching back to Site A after a temporary disaster	61
5.8	Ongoing SRDF operations after a disaster	61
6	OTHER OPERATIONAL PROCEDURES	63
6.1	List symstar sessions on Symmetrix	63
6.2	Checking status of a specific Star session	63
6.3	Isolating an SRDF/Star site	64
6.3.1	Isolating a site	64
6.3.2	Rejoining an isolated site	64
6.4	Transient remote fault recovery	65
6.5	Rebuilding Star configuration when adding/removing disks	66
6.5.1	Example of adding devices to an SRDF/Star configuration	66
6.6	Monitoring SRDF/Star	69
6.7	Dealing with loss of all RDF daemons	70
6.7.1	How to disable Star consistency so that the DMX can control cycle switching.	71
7	APPENDIX A – SET UP OF SRDF/STAR WITH MULTIPLE TRIANGLES PER SRDF/STAR SESSION	72
7.1	Example showing configuration of Star with two triangles	72
7.1.1	Step 1a – Verify Symmetrix Settings	73
7.1.2	Step 1b - Create Concurrent SRDF environment	73
7.1.3	Step 2 – Creating SRDF/Star Composite Groups	75

7.1.4	Step 3 – Creating the SRDF/Star options file	75
7.1.5	Step 4 – Performing the symstar setup operation	76
7.1.6	Step 5 – Creating R2 Composite groups	76
7.1.7	Step 6 - Adding BCV devices to the SRDF/Star configurations	76
7.1.8	Step 7 – Bringing up the SRDF/Star sites for normal operation	77
7.2	Example showing removal of a triangle from a Star session	77
8	APPENDIX B – LISTING SUMMARY STATUS OF STAR SESSIONS	79
9	APPENDIX C – OUTPUT FROM SYMSTAR SHOW –DETAIL	81
10	APPENDIX D – OUTPUT FROM SYMSTAR QUERY	82

Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citigroup.

CITI, CITIBANK, CITICORP and CITIGROUP are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citigroup and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

Contact

Please address any comments or questions regarding the information contained in this document via email to the “*GT Global Distributed Storage Engineering” distribution list or alternatively to gdse@citigroup.com.

Latest Versions

It is recommended that the latest version of the document be used from the GDSE website

<http://engineering.citigroup.net/gdse/products/>

Newer versions may include expanded scope, or comments on issues discovered after deployment.

1 Introduction

1.1 Intended audience

This document is intended for:

- Storage Integration staff responsible for planning and implementing SRDF/Star environments.
- Storage management staff responsible for the ongoing operation of SRDF/Star environments

1.2 Out of scope

The following topics are neither covered nor referenced in this document:

- Converting an existing 2-site SRDF/S implementation to SRDF/Star.
- Building a 3-Site Storage Wide Area Network (SWAN), and related networking hardware that will be needed.
- Template operational scripts. These are still in preparation and will be referenced in a future update to this document.

1.3 Important Related Documentation

The present document does not repeat material that already exists in other GDSE documentation. It must be read in conjunction with the documentation that is references, in particular those marked below as **Essential Reading**.

Document	Location
Key documents	
<i>Essential Reading.</i> EMC Solutions Enabler Symmetrix SRDF Family CLI Product* guide	Available from EMC. Chapter 4 and Appendix B refer to SRDF/Star.
<i>Essential Reading.</i> SRDF/A for Long Distance Replication Solutions Guide	http://engineering.ti.citigroup.net/gdse/products/platforms/symmetrix/stdsdocs/05002-sg.pdf
Using SYMCLI to Implement SRDF/Star	Available on Powerlink

SRDF/Star Release Notes (5772 & SE6.4)	http://engineering.citigroup.net/gdse/products/platforms/symmetrix/releasenotes/06050-rn.pdf
EMC Symmetrix DMX Best Practice	http://engineering.ti.citigroup.net/gdse/products/platforms/symmetrix/stdsdocs/03038-sg.pdf
Other documentation	
SRDF/Star Scalability Testing – Key findings	http://engineering.citigroup.net/gdse/products/platforms/symmetrix/techinfo/08012-tr.pdf
SRDF/A for Long Distance Replication Release Notes	5x71 microcode http://engineering.ti.citigroup.net/gdse/products/platforms/symmetrix/releasenotes/05002-rn.pdf 5772 microcode http://engineering.citigroup.net/gdse/products/platforms/symmetrix/proposed/07039-rn.pdf
SRDF/A Evaluation	http://engineering.ti.citigroup.net/gdse/products/platforms/symmetrix/stdsdocs/05002-e.pdf
EMC Solutions Enabler 6.x Solutions Guide	SE6.3 http://engineering.ti.citigroup.net/gdse/products/platforms/symmetrix/stdsdocs/05016-sg.pdf SE6.4.2 http://engineering.ti.citigroup.net/gdse/products/platforms/symmetrix/stdsdocs/07060-sg.pdf

1.4 Changes in V2.0

- Updated throughout for RAID5/RAID6 support
- Updated throughout for new 5772 SRDF features
 - Transmit Idle
 - Delta Set Expansion (DSE)
- Various typographical errors corrected in examples
- Section 1.3. Added referece to new SRDF/Star Scalability Key Findings document.

-
- Section 3.1.2. New section on scalability limitations following joint testing with EMC in their lab in Hopkinton.
 - Section 3.4.1. Updated to cover multiple applications sharing a Star session.
 - Section 3.4.6. Updated Best Practice on Star control servers and storrdfd daemon servers.
 - Section 3.4.7.1. Updated to accommodate DSE (5772 or later)
 - Section 3.4.7.2. Updated to accommodate higher limits on RDF groups per RA (5772 or later)
 - Section 3.5.3. Updated to accommodate DSE (5772 or later)
 - Section 3.5.4. Updated with revised point-in-time copy Best Practice
 - Section 5.1.3. New Best Practice on failover of multiple Star sessions
 - Section 4.1. Updated to include correct setup of DSE & Transmit Idle for Async legs
 - Section 4.2. Updated with scalability considerations
 - Section 5.6 Updated with new procedure needed for unplanned failover when Transmit Idle is enabled on the Async leg
 - Section 6.4. Updated with a warning about Site A failures during an Site A → Site C resync.
 - Section 6.7 Dealing with loss of all RDF daemons.
 - Appendix A. Updated to include correct setup of DSE & Transmit Idle for Async legs

2 Introduction to SRDF/Star

2.1 What is SRDF/Star?

SRDF/Star is a three-site DR solution from EMC for Symmetrix storage arrays that uses concurrent data replication technology to replicate data from a primary production site to both a nearby remote site and a distant remote site. Data is transferred in SRDF/Synchronous (SRDF/S) mode to the nearby remote and in SRDF/Asynchronous (SRDF/A) mode to the distant remote site.




Figure 1 - STAR/Star configuration during normal operations

In the event of a production site failure, there would be no data loss because of the synchronous replication to the regional site. SRDF/Star allows you to *incrementally* establish an asynchronous session between the two remote sites, allowing, for example, data at the synchronous target site to quickly update the distance asynchronous target site with missing data to ensure that there is zero data loss at the remote site.


In the event of a regional disruption that knocked out both the workload site and the synchronous target site, SRDF/Star's concurrent RDF setup ensures that there would be only a minimal data loss because of the asynchronous replication to the more distant async target site.

SRDF/Star also allows you to determine which remote target site (sync or async) has the most current data in the event of a rolling disaster that affects the workload site. With a rolling disaster, there is no guarantee that the sync site will be more current than the async site. The capability to

display where the most current data is located helps determine which site's data should be used for failure recovery.

It is important to understand that SRDF/Star functionality is dependent on an underlying concurrent SRDF implementation. Concurrent SRDF (simultaneous SRDF/S and SRDF/A) provides 3-site replication, but SRDF/Star supplies an advanced failover and recover mechanism on top of concurrent SRDF.

The following diagram shows a simplified view of the relationship between different SRDF components. SRDF/Star simply brings together a number of separate SRDF technologies that EMC have been developing over the last few years, into a single advanced DR solution.



All of the components shown above that SRDF/Star depends on are already GDSE certified components (note though that SRDF consistency groups are only certified for use with SRDF/A at the time of writing).

2.2 SRDF/Star vs SRDF/A

It is important to understand some of the key similarities and differences between 2-site SRDF/A and SRDF/Star when deciding which is the most appropriate technology to use. SRDF/Star has both benefits and functional limitations compared with 2-site SRDF/A.

2.2.1 Similarities between SRDF/A and SRDF/Star

- Both offer zero data loss after a server or application level failure
- Both require application level failover when an individual data server fails
- Both can offer application level failover
- Neither guarantees zero data loss in a Rolling Disaster scenario
- Neither guarantees zero data loss in a Regional disaster scenario
- Both protect against total data loss in the event of a regional disaster
- Long distance network costs could be similar if engineered carefully

2.2.2 Differences between 2-site SRDF/A and 3-site SRDF/Star

- Only STAR provides zero data loss (up to the start of a rolling disaster) in the event of loss of the production site
- Only STAR offers ongoing data replication after the loss of Site A
- 2-site SRDF/A is a higher I/O performance solution than SRDF/Star as it does not have a “slow” synchronous replication component.
- STAR requires a third bunker site.
- STAR has approximately 50% extra storage costs. In addition there will in due course be an additional charge for 3-site replication to cover other related costs.
- STAR may have higher networking cost compared with 2-site SRDF/A, depending on how the network component is engineered.
- Failovers with STAR is a much more complex and a slower process than with 2-site SRDF/A. Failover of a Star session will take a minimum of 6-8 minutes. An 8-minute SRDF/Star failover execution time (just an example) for an individual session may sound acceptable until one considers that one would be doing multiple failovers, usually NOT in parallel on a given Symmetrix. **Hence, when multiple sessions (e.g. 12) are being failed over the end-to-end process for all sessions might be nearly 2 hours with SRDF/Star.** With SRDF/A, the entire process would be measured in minutes.

-
- STAR is much less scalable and flexible than 2-site SRDF/A. At the time of writing, EMC only support ~10 Star sessions per Symmetrix frame. **Hence, Star is not suitable for large scale deployment where application level failover is also required.** If many applications require Star, it will be necessary to put multiple applications into the same Star session, and fail them all over together. Application level failover can only be provided where only a small number of applications are resident on a Symmetrix.
 - SRDF/Star setup is complex and hence allocation of additional disk to an application on a regular basis once under Star control cannot usually be supported by operations. Disk allocations would need to be made before Star is set up to accommodate anticipated growth.
 - Long-distance replication data flow under SRDF/Star is controlled by daemons running on servers and not by the more robust DMX itself.

The following table attempts to summarize some of the similarities and differences between the SRDF/S, 2-site SRDF/A and the 3-site SRDF/Star solutions. The questions are phrased such that an affirmative answer is a positive characteristic.

	2-site SRDF/S	2-site SRDF/A	3-site SRDF/Star
Certified by Global Engineering?	Yes (Strategic solution)	Yes (Strategic solution)	Yes (Strategic, <i>but for selected applications only</i>)
Typical separation of production and COB sites	Tens of km	Thousands of km	Thousands of km
Protects against total loss of all data availability after regional disaster?	No	Yes	Yes
Typical amount of data loss after server failure	None	None	None
Can support individual data server level failover (where > 1 data server)?	Yes	<i>In practice, No</i>	No
Can offer application level failover	Yes	Yes	<i>Yes, with scalability limitations</i>
Typical amount of data loss after production site failure	None	<i>30-60 sec</i>	None
Guarantees no data loss in a rolling disaster scenario?	No	No	No
Typical amount of data loss at remote site after wide-scale disaster at production location	N/A	30-60 sec	30-60 sec
Number of sites requiring storage	2	2	3
Ongoing replication to a remote site after loss of original production site	No	No	Yes
Application I/O performance not limited by synchronous replication	No	Yes	No
Cob copy of data available for use under normal operations	No	No	No

A comparison between three Citigroup-certified SRDF models

2.3 Applications for which SRDF/Star is of benefit

2.3.1 Applications which need long-distance zero data loss COB when production site lost

SRDF/Star allows you to *incrementally* establish an asynchronous session between the two remote sites, allowing, for example, data at the synchronous target site to quickly update the distance asynchronous target site with missing data to ensure that there is zero data loss at the remote site.

It must be stressed that even SRDF/Star cannot guarantee zero data loss in a rolling disaster scenario. A rolling disaster describes disaster marked by different start and end points i.e. non-instantaneous events. Depending on the order in which the IT infrastructure fails, one cannot guarantee zero data loss in this case. For example, if the synchronous replication links go down before the production storage or servers, Star will not be able to provide zero data loss at the long distance site.

Note that the rolling disaster issue already exists today with the Synchronous SRDF replication solution used today. One will get zero data loss up until the start of the rolling disaster, but applications may continue to run for a while during the finite period that the site becomes disabled.

This issue is being mentioned, as Star is a high-cost solution being put in explicitly over and above SRDF/A to provide zero data loss.

In Summary: No replication solution, EMC or otherwise, can guarantee zero-data loss in the event of a rolling disaster.

2.3.2 Applications which must have on-going COB after loss of production site

Being a 3-site solution, SRDF/Star is able to provide SRDF replication between two remaining sites should one site be rendered inoperable. Note, however, that current Global Data Centre Strategy standard state that Citigroup's 3-site replication solution will not have servers at Site B. Site B will purely be used as a data bunker. Site C will be the DR site protecting Site A. Site B as bunker-only configuration, does not provide true redundant protection (ongoing COB protection) after a site or regional disaster as there are no servers at Site B.

2.4 Applications for which SRDF/Star is not suitable

2.4.1 Applications with multiple data servers

Star can only be used to failover at the RDF group level. Citigroup intends to map applications to RDF groups on a 1:1 basis to ensure data consistency between the different components of the application. Hence, if a data server is one of several that support a given application, it will not be possible to failover just that one data server as can be done today with SRDF/S. With both SRDF/Star and SRDF/A failover must be at the RDF group level and hence ALL data servers that comprise the application would need to failover together, even though only one of them may have failed.

In Summary: SRDF/Star only supports failover of entire RDF groups. Entire applications will need to be failed over, each with their own Star session.

2.4.2 Applications with changing storage requirements

The process for modifying the storage used by an SRDF/Star session is very complex, and includes dismantling and rebuilding the Star session for the application from scratch whenever LUNs are added or removed from the application. Hence, there is a very significant management overhead for applications with changing storage requirements.

2.4.3 Cost-sensitive applications

SRDF/Star requires storage at three sites and additional network infrastructure. As such, it is the most expensive SRDF solution offered today, and only suitable for applications that can justify the additional expense and which really need the benefits offered by SRDF/Star.

2.4.4 Applications that are not already in production

At the time of writing, EMC needs to size an SRDF/A (and hence an SRDF/Star) environment based on the actual I/O profile of the applications to be protected. Hence, it is difficult at present to size an SRDF/Star environment for new applications that cannot have their I/O profile measured up front.

Hence, Citigroup is primarily certifying SRDF/Star for use with existing application with I/O profiles that can be measured, as EMC cannot size a generic system. Engineering recommends against the installation of SRDF/Star infrastructure without clearly defined target applications and hence identifiable sizing requirements. If an SRDF/Star-enabled environment must be built, contact the author to discuss this requirement further.

3 Preparation for an SRDF/Star implementation

3.1 Complexity and scalability warnings

3.1.1 Complexity

SRDF/Star is considerably more complex than 2-site SRDF/S or 2-site SRDF/A to both configure, but more importantly to operate and manage. It is essential that Storage Management teams make themselves very familiar with Star operational tasks by reading the documents referenced and possibly going on EMC SRDF/Star training, if available.

It is *strongly* recommended that Storage Management teams configure a couple of LUNs into a permanent test Star configuration to allow them to test their processes and procedures, and for training new staff.

3.1.2 Scalability limitations

3-Site SRDF/Star is not as scalable as 2-Site SRDF/A for a number of reasons, but primarily because of:

- Star session failover times result in longer overall RTO if many sessions are to be failed over, as failover of sessions is sequential, at present (SE6.4).
- Lower EMC support limits for Star sessions per array

At the time of writing, the following scalability limits are supported by EMC & GDSE.

Description	EMC support limit
# Star sessions per source DMX	32 ¹
# Star sessions per primary Star control server	32 ²
# Star sessions per <u>active pair</u> of storrdfd daemons (one per SE host)	32 ³

¹ If TimeFinder is in use at Site B then this will be reduced to 30.

² A single server could control either 32 Star sessions on a single DMX array or, for example, 16 on each of two DMX arrays.

³ Cycle switching will not occur reliably with less than 2 RDF daemons running when 32 Star sessions are active.

For example, a pair of Star Control servers (primary and backup), both running storrdfd daemons can support a maximum of 32 Star sessions. When scaling to the 32 session limit, a third storrdfd daemon should also be running on a third host at the production site (Site A) as a minimum, to allow for the failure of a single Star management server.

If the number of applications to be placed on an array or to be managed by a pair of management servers exceeds the above limits, then it will be necessary to place multiple applications in the same Star Session (see [Section 4.2](#)). The consequence of this is that all the applications sharing the Star session will need to be failed over together, as a single unit.

3.2 Major Pre-requisites

The following significant pre-requisites must be met before SRDF/Star can be implemented.

1) A concurrent SRDF environment.

SRDF/Star			
Concurrent SRDF		Dynamic RDF	SRDF Consistency Groups
SRDF/S	SRDF/A		

As the above figure shows, SRDF/Star requires a pre-existing concurrent SRDF environment before it can be configured. This document contains references to other documents that described how to build a concurrent SRDF environment, in particular, the GDSE SRDF/A Solutions Guide. The present document only provides information that is specific to SRDF/Star.

SRDF/Star has certain additional Symmetrix hardware requirements over and above those for Concurrent SRDF. These additional Star-specific requirements are highlighted in the present document.

2) SRDF/Star-specific networking requirements.

Concurrent SRDF requires connectivity between Production site (Site A) and both SRDF target site Symmetrix (Sites B & C). For SRDF/Star, IP connectivity is also required between the Symmetrix environments at Site B and Site C, with dedicated SRDF/A bandwidth. This is over and above network requirements of standard concurrent SRDF. **It is essential that the regional network engineering team be engaged as soon as a Star deployment is proposed, as a triangular Storage Wide Area Network (SWAN) will be required to support SRDF/Star** (See [Section 3.4.4](#)).

3.3 Overview of concurrent SRDF setup

Design, planning and configuration of 2-Site SRDF/A and 3-Site Concurrent SRDF is described in detail in the document “**SRDF/A for Long Distance Replication**”. That detailed information is not repeated in the present document. The main steps are summarized in the tables below, with references to the relevant sections in the SRDF/A Solutions Guide, with additional further references to updated Best Practice the present document, as appropriate.

The creation of a concurrent SRDF environment is a very significant piece of work. The above SRDF/A Solutions Guide must be studied in detail, and the installation planned carefully. Planning and configuration of the SRDF/A component is the major part of setting up a concurrent SRDF environment and is the focus of much of the discussion.

There are some notable differences between the requirements for a Star and a non-Star concurrent SRDF environment. Hence, where there are differences in Best Practice or standards between the SRDF/A Solution guide and the present document, the present document is the definitive guide for Star implementations.

The structure of the rest of this section on preparing for an SRDF/Star implementation may seem a little odd as it is dictated by the layout of the existing SRDF/A Solutions Guide, which it heavily references. This section is in effect just a set of additional or updated STAR-specific Best Practice recommendations. The SRDF/A Solutions Guide is still the primarily reference. Here the discussion is broken up into three main sections:

- Planning and Preparation for an SRDF/A implementation
- SRDF/A configuration
- SRDF/Star-specific Concurrent SRDF configuration

Planning and Preparation for an SRDF/A implementation	SRDF/A Solutions Guide section reference	Section reference for additional or updated Best Practice for SRDF/Star in the present document
Collection of Basic Application Requirements	4.2	3.4.1
RDF group naming/allocation conventions	N/A	3.4.2
EMC Solution Qualification (SQ)	4.3	3.4.3
WAN requirements	N/A	3.4.4
Network Hardware	4.5	3.4.5
Management server requirements	4.6	3.4.6
Storage Hardware requirements	4.7	3.4.7
DMX bin file configuration	4.8	3.4.8

SRDF/A configuration	SRDF/A Solutions Guide section reference	Section reference for additional or updated Best Practice for SRDF/Star in the present document
Management Server Configuration	5.1	3.5.1
Network Equipment Configuration	5.2	3.5.2
Dynamic RDF configuration	5.3	3.5.3
Target site Point-in-Time image for SRDF/A resynchronization	5.4	3.5.4
Device Group Definition	5.5	3.5.5
Enabling SRDF/A replication	5.6	3.5.6
SRDF/A Replication Monitoring	5.8	3.5.7

The following sections in the SRDF/A Solutions Guide that refer to Concurrent SRDF should be dis-regarded when deploying SRDF/Star: The superceding Best Practice is located in the present document, as indicated below.

SRDF/Star-specific Concurrent SRDF configuration	SRDF/A Solutions Guide section reference	<i>Superceding Best Practice for SRDF/Star in this document</i>
Licensing for SRDF/Star	7.2	3.6.1
Device Group definitions in SRDF/Star environments	7.3	3.6.2
3-Site SRDF/Star RDF group requirements	7.4	3.6.3

Finally, throughout this document, the following Site naming convention is used:

Site Role	Site Role Short name
Production site / Source site	Site A
Synchronous SRDF target site / bunker site	Site B
Asynchronous SRDF target site / COB site	Site C

3.4 Planning and Preparation for an SRDF/A implementation – SRDF/Star updates

The following information supplements rather than replaces (unless explicitly stated) the referenced Best Practice in the SRDF/A Solutions Guide.

Planning and Preparation for an SRDF/A implementation	SRDF/A Solutions Guide section reference	Section reference for additional Best Practice for SRDF/Star in the present document
Collection of Basic Application Requirements	4.2	3.4.1
RDF group naming/allocation conventions	N/A	3.4.2
EMC Solution Qualification (SQ)	4.3	3.4.3
WAN requirements	N/A	3.4.4
Network Hardware	4.5	3.4.5
Management server requirements	4.6	3.4.6
Storage Hardware requirements	4.7	3.4.7
DMX bin file configuration	4.8	3.4.8

3.4.1 Collection of Basic Application requirements for Star

Collection of detailed application requirements is key to a successful SRDF/Star implementation. For SRDF/Star the critical information that needs to be collected for EACH application is:

1. Name of application. The application must already exist in production, as it will be necessary to measure its I/O profile.
2. Confirm if SRDF/Star is a suitable protection scheme for the application. See Sections [2.2](#), [2.3](#) and [2.4](#) for information on applicable environments.
3. Amount of usable production storage required.
4. RDF groups to be used by application (see section 3.4.2 below)
5. If TimeFinder/Mirror is used with the application at production site. This is not supported with SRDF/Star due to mirror position limitations, and some changes to the application setup would be needed. For example, the application may need to start using TimeFinder/Mirror Clone Emulation mode, rather than “Classic” TimeFinder/Mirror.
6. Determine if the application has to have storage-enforced I/O consistency at the DR site with any other applications.
 - a. If not, then a “single triangle” implementation of SRDF/Star will be needed for this application. **This will be the normal configuration for Star.**
 - b. If Storage-enforced I/O consistency is needed with other applications, then a “multiple triangle” implementation of SRDF/Star would be needed for this and the other applications with which consistency is required. Each application will have its own triangle of RDF groups, but RDF replication will be done in a consistent manner across the groups. Note that all applications sharing a single Star session

will be controlled as a single entity for DR purposes i.e. all those applications must be failed over together.

7. If multiple applications need to share a Star Session due to scalability limitations, then a “multiple triangle” implementation of SRDF/Star should be used (see [Section 4.2](#)). Each application will have its own triangle of RDF groups. Only place those applications in the same Star session that is acceptable to the business to failover over together. For example, the business may agree that all applications on the same physical server can be grouped together.

Note that most applications will rely on application-enforced consistency at the DR site, and hence a “single triangle” implementation will be sufficient.

3.4.2 RDF group naming/allocation conventions for Star

SRDF/Star uses RDF group naming. This is a nickname that can be used to refer to a RDF link (or group of links) to a remote site. Star operations are then executed using this nickname, avoiding having to refer to RDF group number in commands, and hence removing a significant source of human error.

It is recommended that you use physical location names as the RDF group nicknames. For example, if your production site were in New York, your bunker site in New Jersey and the remote DR site in London, then the following naming convention would be suitable.

Star Site	Nickname
Site A	NewYork
Site B	NewJersey
Site C	London

In addition, it is necessary to allocate a total three RDF groups that each application will use for replication. These three RDF groups define the “triangle” over which SRDF/Star operates.

Any sensible naming convention can be used, although it is strongly recommended to use the same RDF group number at each end of a link. This will consume more RDF group numbers but makes troubleshooting issues and management of DSE much simpler. An example of an RDF group usage convention is shown in the table below. The key objective is to have a simple convention, and not to use “random” RDF group numbers.

Application	Dynamic RDF group between Site A and Site B	Dynamic RDF group between Site A and Site C	Dynamic RDF group between Site B and Site C
App A	4	5	6
App B	7	8	9
App C	10	11	12
App D	13	14	15
App E	16	17	18
App F	19	20	21
App G	22	23	24
App H	25	26	27
App I	28	29	30
App J	31	32	33

The RDF groups 1 through 3 are reserved for underlying static RDF group upon which the dynamic RDF group used by Star are built (see [Section 3.6.3](#))

Note that this naming convention differs from that recommended for Concurrent SRDF in the SRDF/A Solutions Guide as it has to accommodate the RDF recovery links between Sites B and C.




Figure 2 - Example naming convention for SRDF/Star RDF groups showing two applications

In the above example, App A is using RDF groups 4, 5 & 6 and App B is using 7, 8, & 9 for the synchronous leg, asynchronous leg and recovery links respectively.

The entries in the dotted boxes refer to the RA names over which the RDF groups run.

3.4.3 EMC Solution Qualification (SQ) for Star

Every SRDF/Star implementation requires qualification by EMC. The local EMC account team should be contacted as soon as an SRDF/Star implementation is proposed, and instructed to request a Solutions Qualifier (SQ) for STAR Open Systems. It is predominately the SRDF/A component that requires the level of detailed planning.

The SQ process for Star is likely to vary with time EMC refine and update their processes. Hence, a detailed description of the current process is not given here.

This process will involve data collection for the target applications and subsequent analysis of the data. Hence, all applications in scope will need to be confirmed before this work can take place. This analysis must be completed before any orders are placed, as it will be used to determine network bandwidth requirements, DMX hardware configuration and other hardware requirements.

The whole process can take up to 2 months, although with focussed effort, the duration could be reduced to 3-4 weeks. The local account team will be responsible for the completion of the SQ, but will need significant input from the local Citigroup staff. The SQ form is very detailed (hundreds of questions) and requires the collection of current performance data.

Hence, Citigroup is primarily certifying SRDF/Star for use with existing application with I/O profiles that can be measured, as EMC cannot size a generic system. Engineering recommends against the installation of SRDF/Star infrastructure without clearly defined target applications and hence identifiable sizing requirements. If an SRDF/Star-enabled environment must be built, contact the author to discuss this requirement further.

3.4.4 WAN requirements for Star

As described in the SRDF/A Solutions Guide, Concurrent SRDF requires:

- Redundant DWDM connectivity between Site A and Site B for SRDF/S replication
- Redundant, dedicated IP bandwidth between Site A and Site C for SRDF/A replication

In addition, SRDF/Star requires IP bandwidth between Site B and Site C for the “recovery links” (see previous Figure 2) used for:

- the transfer of data from Site B to Site C during an unplanned switch to Site C
- ongoing SRDF/A replication from Site C to Site B

Global standards do not exist for designing such “triangular” Storage Wide Area Networks (SWANs) that can be referenced. Japan already has a triangular SWAN in place, which is used for the existing Concurrent SRDF environment in that location. At the time of writing, EMEA is in the process of determining how best to provide similar network connectivity.

At the time of writing it appears that region-specific solutions will need to be created to accommodate different regional constraints. **It is essential that the regional network engineering team be engaged as soon as a Star deployment is proposed.**

One possible design for a SWAN network is shown in Figure 3, for illustrative purposes only.




Figure 3 - Example SWAN design for Tokyo (courtesy of Geoff Taylor)

Note that the design in Figure 3 gives full redundancy but with minimal additional long-distance network costs compared with a redundant 2-site long distance SRDF/A solution.

3.4.5 Network hardware for Star

The exact detail of the networking hardware needed for an SRDF/Star deployment depends on the nature of solution the triangular SWAN. For example, some possible triangular SWAN deployments may route traffic from Site A to Site C via Site B (for one redundant leg), as shown in Figure 3. Similarly, some solutions may do all routing at the IP layer where as others will do some routing between Sites A and B at the Fibre Channel level.

A future revision of this document may be able to give more specific advice, or references to more detailed documentation as regional triangular SWAN designs, and associated storage networking hardware are finalized.

3.4.6 Management server requirements for Star

Redundant management servers will be needed at Sites A and C as a minimum. Star, like 2-Site SRDF/MSC, depends on RDF daemons running on the management servers and hence these systems must be highly available (see [Section 4.3](#)).

If at least one active storrdfd daemon is not available at the production site, SRDF/A cycle switching will stop. Hence, the Star management servers at the production site should be taken down for the briefest periods and never both at the same time. When more than ~16 Star sessions are planned, you should aim to always have at least 2 RDF daemons running at all times.

Star does not strictly need dedicated Star management servers. Existing Solutions Enabler (SE) servers could be used. However, Star is critically dependent on the storrdfd daemon that runs on SE servers. A minimum of two daemons are needed for redundancy, three if supporting more than ~20 Star sessions.

Hence, if a very small number of Star sessions are to be supported, one could use an existing pair of SE servers at each site, each of which must run the storrdfd, and one of which must always be up (with the daemon running).

For a more scalable solution (up to 32 Star sessions), it is recommended that a pair of servers at the Production site (Site A) be dedicated for Star Control, and that storrdfd daemons also be run on the non-Star SE servers for resilience and load-balancing.

At the Async target site (Site C) it is possible to save money by using the non-Star SE management servers, as they will not be controlling RDF replication, except when operations have switched to Site C. If Site C is also to be used at some point as a production site for Star protected applications, then also install two Star-dedicated servers in anticipation (hence providing a total of 4 hosts running the storrdfd daemon).

Servers are not needed at Sync target site (Site B), unless switching production to Site B is planned. At the time of writing, this is not Citi's strategic direction.

Site	Size	Management Servers required	Comments
Production (Site A)	Large installation – up to 32 Star Sessions	2 dedicated Star servers <u>in addition</u> to any non-Star SE servers <u>recommended</u> . Run storrdfd on all servers for redundancy and load balancing.	At least <u>two</u> storrdfd daemons, and hence two hosts must be up <u>at all times</u> to regulate SRDF/A cycle switching.
	Small installation – up to 16 Star Sessions	2 dedicated Star servers in addition to any non-Star SE servers <u>optional</u> . Run storrdfd on all servers for redundancy and load balancing.	At least <u>one</u> storrdfd daemons, and hence one host must be up <u>at all times</u> to regulate SRDF/A cycle switching.
Sync Target (Site B)		<u>Optional</u> : Use existing SAN management servers for Star control and storrdfd daemons.	Servers only needed if switching production to Site B

Async target (Site C)		2 dedicated Star servers in addition to any non-Star SE servers <u>optional</u> . Run storrdfd on all servers for redundancy and load balancing.	If Site C is also to be used at some point as a production site for Star protected applications, then install two Star-dedicated servers in anticipation.
--------------------------	--	--	---

Configuration and management of SRDF/Star via Solutions Enabler is only supported on Solaris. See SRDF/Star Release Notes for the version to be used, and any hot-fixes that must be applied.

Server CPU loading by the storrdfd daemon and symstar commands is modest and hence, there are no specific minimum hardware requirements for the management servers.

3.4.7 Storage Hardware Requirements for Star

It is strongly recommended that a dedicated set of Symmetrix be deployed for applications requiring SRDF/Star protection.

The following requirements are over-and-above those defined in the SRDF/A Solution Guide Section 4.7:

3.4.7.1 Cache

The Symmetrix at Site B may become the target of asynchronous replication after loss of Site A and when the workload has been moved to Site C. Hence, it needs additional cache to support SRDF/A. EMC must define the amount of cache needed via the Solutions Qualification (SQ) process, as it is dependent on workload and available bandwidth to the asynchronous target sites.

With 5772 or later, DSE and Transmit Idle should be used between Sites A and C to avoid SRDF/A dropping due to temporarily network outages or workload peaks.

3.4.7.2 Additional RA requirements

5x71 microcode

All three Symmetrix will need sufficient RAs to be capable of running concurrent SRDF. RAs should not be used for SRDF/A and SRDF/S simultaneously. EMC recommends that no more than 8 RDF groups (including the underlying static RDF group i.e. 7 applications) are configured per RA port.

The requirements are best illustrated with an example:

Site A Symmetrix

Site A → Site B: SRDF/S RA ports needed = # Star-protected applications / 7 * 2

Site A → Site C: SRDF/A RA ports needed = # Star-protected applications / 7 * 2

Site B Symmetrix

Site B → Site A: SRDF/S RA ports needed = # Star-protected applications / 7 * 2

Site B → Site C: Recovery link RA ports needed = # Star-protected applications / 7 * 2

Site C Symmetrix

Site C → Site A: SRDF/A RA ports needed = # Star-protected applications / 7 * 2

Site C → Site B: Recovery link RA ports needed = # Star-protected applications / 7 * 2

For example, if you wanted to support a maximum of 12 applications with SRDF/Star you would need 8 RA ports on each Symmetrix (based on 5x71 EMC limits).

The above algorithm assumes that the determining factor of the number of RA ports is the limit of 8 RDF sessions per port. Additional RAs may be required if heavy I/O applications require that < 7 applications are configured per pair of RA ports.

5772 microcode or later

All three Symmetrix will need sufficient RAs to be capable of running concurrent SRDF. RAs should not be used for SRDF/A and SRDF/S simultaneously. With 5772 or later, the number of RDF groups per RA is no longer the limiting factor; the total number of Star sessions supported per array will be *less* than the number of RDF groups per RA supported. However, more than one pair of RAs may be needed to provide sufficient bandwidth for the target applications. Analysis of the workload will determine the number of RAs needed (see [Section 3.4.3](#)).

3.4.7.3 Balanced Configuration

For an SRDF/Star installation, the Symmetrix at all three sites require:

- Identical symmetrix models
- Identical microcode levels
- Identical disk type for RDF devices
- Identical RAID type for RDF devices
- Additional BCV storage at Asynchronous Target site (see [Section 3.5.4](#))

Any deviations from the above would need to be approved by EMC on a case-by-case basis.

3.4.8 DMX bin-file configuration for Star

The following requirements are **over-and-above, or replace** (as indicated) those defined in the SRDF/A Solution Guide (Section 4.8 of that document):

3.4.8.1 Microcode levels

The latest GDSE certified microcode versions should be used. The minimum versions that SRDF/Star will work with are 5671.5459 (DMX2), 5771.8795 (DMX3) & 5772.83.75 (DMX4).

3.4.8.2 Dynamic RDF bin file requirements

A static RDF group (RDFG 3) will also need to be created between all ports on DMX units at Sites B and C that are to be used for the RDF recovery links. This is in addition to the static RDF groups

referred to in the SRDF/A Solutions Guide Section 4.8.2 between Sites A and B, and Sites A and C.

3.4.8.3 BCV Requirements at target sites

This is a complex issue with SRDF/Star and is discussed in detail in [Section 3.5.4](#).

3.4.8.4 BCV usage constraints at source site

It is not possible to use a TimeFinder/Mirror BCV with a mirrored volume that is the source of concurrent SRDF replication. If BCV-type functionality is required, it will be necessary to use either TimeFinder/Mirror Clone emulation or TimeFinder/Clone technology to create PIT copies at the source site. Clones have already been certified for use in this circumstance by GDSE.

3.4.8.5 Additional Gatekeeper requirements

See Section 3.5.1.6 of this document for recommendations on Gatekeeper requirements.

3.4.8.6 Hot Spares and Symptomiser

Although covered in the SRDF/A Solutions Guide, it must be stressed that with Concurrent SRDF it is not possible to use Hot Spares on the source DMX, as all four mirror positions are used up for local mirroring (RAID1 & RAID5) and two SRDF relationships.

With SRDF/Star, the situation is more extreme. The target DMX units should also not use Hot Sparing as a spare mirror position is needed by Star during certain switching operations. Such switching operations will fail if Hot Spare activities are taking place and consuming the final mirror position.

Similarly, as Symptomiser needs a spare mirror position, it is not compatible with SRDF/Star enabled DMX.

3.5 SRDF/A configuration – SRDF/Star updates

The following information supplements rather than replaces (unless explicitly stated) the referenced Best Practice in the SRDF/A Solutions Guide.

SRDF/A configuration	SRDF/A Solutions Guide section reference	Section reference for additional Best Practice in the present document for SRDF/Star
Management Server Configuration	5.1	3.5.1
Network Equipment Configuration	5.2	3.5.2
Dynamic RDF configuration	5.3	3.5.3
Target site Point-in-Time image for SRDF/A resynchronization	5.4	3.5.4
Device Group Definition	5.5	3.5.5
Enabling SRDF/A replication	5.6	3.5.6
SRDF/A Replication Monitoring	5.8	3.5.7

3.5.1 Management Server Configuration

The following requirements are **over-and-above, or replace** those defined in the SRDF/A Solution Guide Section 5.1:

3.5.1.1 *Solutions Enabler*

Configuration and management of SRDF/Star via Solutions Enabler is only supported on Solaris. See SRDF/Star Release Notes for the version to be used, and any hot-fixes that must be applied.

3.5.1.2 *Software Licensing*

Software licensing for Star is discussed in this document in [Section 3.6.1](#)

3.5.1.3 */etc/system*

Configure as per SRDF/A Solutions Guide Section 5.1

3.5.1.4 */var/symapi/config/options*

Configure as per SRDF/A Solutions Guide Section 5.1

3.5.1.5 */var/symapi/config/daemon options*

GNS remote mirror replication does not work for composite groups containing devices using concurrent SRDF, and therefore cannot be used in an SRDF/Star environment. GNS can, however, be used to ensure changes to composite definitions are automatically visible on other local SAN management servers.

3.5.1.6 Gatekeeper requirements

When undertaking parallel operations, there must be sufficient gatekeepers to satisfy simultaneous requests, without causing either gatekeeper starvation, or worse failure.

The following settings in the daemon_options file were found to support SE6.4 Star without issues:

```
storrdfd:rdfd_num_dedicated_gks=20*  
storrdfd:rdfd_gk_pool_size=40
```

* Note that if parameter is set to >21, it defaults to 2!

This results in 20 GKs dedicated for RDF daemon usage, and 20 in the pool for the base daemon to use.

Hence, 40 gatekeepers (20 dedicated plus 20 pool) will support the largest Star environment that EMC support today (maximum 32 sessions per array), and should be the default deployment. Note that reducing these numbers may cause failures (e.g. 20 dedicated plus 12 pool gatekeepers were found to cause issues with 32 parallel operations).

3.5.1.7 Symacl configuration requirements

Configure as per SRDF/A Solutions Guide Section 5.1

3.5.2 Network Equipment Configuration

As mentioned in [Section 3.4.5](#), the exact detail of the network equipment configuration needed for an SRDF/Star deployment depends on the nature of solution the triangular SWAN.

A future revision of this document may be able to give more specific instructions when triangular SWAN designs are clarified. FC-IP conversion units will be needed for both the Site A $\leftarrow \rightarrow$ Site C leg, and also the new SRDF/Star-specific Site B $\leftarrow \rightarrow$ Site C recovery links.

3.5.3 Dynamic RDF configuration

Dynamic RDF should be configured as per the procedures in SRDF/A Solutions Guide Section 5.3, with the following variations:

- Dynamic RDF groups need to be configured for each application for each of the following DMX pairs:
 - Site A $\leftarrow \rightarrow$ Site B
 - Site A $\leftarrow \rightarrow$ Site C
 - Site B $\leftarrow \rightarrow$ Site C
- RDF groups to be used in each case as defined by RDF group allocation convention (see [Section 3.4.2](#) for an example convention).

- 5772 or later: DSE and Transmit Idle should be set up as a minimum between Site A and Site C. DSE setup between Sites B and C is only of use once Site A is lost and SRDF is running C → B. Hence, it is recommended that DSE be set up between B and C once Site A is lost. For further information on setting up DSE, see the SRDF/A Solutions Guide, Appendix G

3.5.4 Point-in-Time (PiT) copies for SRDF resynchronization

3.5.4.1 SRDF resync Point-in-time copy issues

As described in the SRDF/A Solutions Guide, PiT copy storage is needed at sites that require R2 protection during SRDF/A resynchronizations.

With SRDF/Star there are some additional complications that need to be considered:

- There can be issues with mirror position exhaustion when TimeFinder/Mirror is used for providing the Point-in-Time copies with RAID1+0 standards. Great care must be taken when undertaking symstar failover and recovery operations when BCVs are in use. These operations generally involve createpair operations, which will need a spare mirror position at each end of the new link to execute successfully. If BCVs are established, for example, for PiT resync protection at a target site, then symstar failover/recovery operations can fail due to lack of mirror positions at Site C. Best Practice on dealing with this issue is detailed below.
- As in principle SRDF/A replication may be taking place to either Sites A, B or C in an SRDF/Star configuration, there is an increased requirement for PiT copies:
 - Site A can be an SRDF/A target site after a planned switch to Site C
 - Site B can be an SRDF/A target site following a disaster switch to Site C when Site A is no longer available
 - Site C is the SRDF/A target site under normal operating conditions (Production running out of Site A)


Best Practice for providing PiT copy protection as a reasonable cost is detailed below.

3.5.4.2 PiT Copy requirements at each site

The following is recommended Best Practice for the PiT copy storage requirements in the SRDF/Star configuration where Site B is a bunker site only.

Site	SRDF resync. PiT copy storage required?	Comments
Site A – Workload	Optional. Only of use if you plan to regularly run workload	Required when an application has been failed over to Site C. Used to protect Site A Async target R2 volumes

	from Site C Needed when an application is running from Site C after <u>planned</u> switch	during a resync after SRDF/A link outage.
Site B – Sync target (bunker only)	Not recommended due to cost and low likelihood of usage Only need when an application is running from Site C after <u>loss of Site A</u>	Site B would only ever be the target of asynchronous replication <u>after the loss of Site A</u> when operations have been failed over to Site C.
Site C – Async target	Recommended	Required during normal operations to protect Site C Async target R2 volumes during a resync after SRDF/A link outage. <u>This is to ensure that there is ALWAYS a consistent copy at Site C.</u> Also of use during an UNPLANNED switch operation to Site C async target, when keeping Site B sync target data as async disks will be inconsistent during the restore step of the switch. Nominally also of value during an UNPLANNED switch operation to sync target site before the connect commands to async target is issued. However, the Citigroup model only use Site B as a bunker site so this should never happen.



The examples in this document assume use of BCVs at Site C only. They make no assumptions about what TimeFinder/Mirror mode is being used.

3.5.4.3 TimeFinder modes

Use of TimeFinder/Mirror Clone emulation mode is preferred in this configuration, as switch operations require a free mirror position, and will fail if a non-clone emulation BCV is established at the R2 end.

“Classic” TimeFinder/Mirror can be used for SRDF resync. PiT copies, but any BCV copies **MUST** be split off before a switch can take place. As this rule applies to both SRDF/A resync BCVs AND user BCVs, use of “Classic” Timefinder/Mirror by users should not be permitted in an SRDF/Star environment. Note that with RAID5 and RAID6, TimeFinder/Mirror Clone Emulation is automatically configured.

The examples in [Section 5](#) will work with either TimeFinder/Mirror mode as they always suggest breaking off the BCV before a split operation.

3.5.4.4 Mapping PiT device to Standards

There is no SRDF/Star specific advice.

3.5.4.5 Adding BCVs to Composite Groups

The process for adding BCVs to Composite Groups in an SRDF/Star configuration is shown in the examples in [Section 4](#).

3.5.5 Device Group definition

The instructions in both the SRDF/A Solutions Guide Sections 5.5 and 7.3 on defining device groups in a concurrent SRDF environment **should be disregarded** for the following reasons:

- Composite Groups are used for SRDF/Star control, and not device groups
- Subset failover is not supported with SRDF/Star

Follow the Best Practice in the [Section 4](#) of the present document that relates to use of Composite Groups for Star control.

3.5.6 Enabling SRDF/A replication

This is not necessary, as the SRDF/Star setup process will enable SRDF replication on both legs.

3.5.7 SRDF/A Replication Monitoring

At the time of writing, the advice in the SRDF/A Solutions Guide Section 5.8 should be followed. See also [Section 6.6](#) in this document.

3.6 Concurrent SRDF configuration – SRDF/Star updates

Concurrent SRDF configuration	SRDF/A Solutions Guide section reference	Section references for <u>superceding</u> Best Practice in this document for SRDF/Star
Licensing for SRDF/Star	7.2	3.6.1
Device Group definitions in SRDF/Star environments	7.3	3.6.2
3-Site SRDF/Star RDF group requirements	7.4	3.6.3

3.6.1 Licensing for SRDF/Star

Use of SRDF/Star requires additional licenses over and above the EMC “High-Availability” bundle. The following assumes that Star sessions do not span frames (which is not supported by GDSE).

License	Included in H/A-SRDF Bundle?	Required at Site A?	Required at Site B?	Required at Site C?
SRDF/Star	No	Yes	Yes	Yes
First SRDF license	Yes	Yes	Yes	Yes
Additional SRDF license*	No	Yes	Yes	No*
TimeFinder/Clone	Yes	Optional**	No	Yes
TimeFinder/CG	No	Optional**/**	No	Optional***
SRDF/CG	No	Yes	No	Optional****

* If the first SRDF license is being used for SRDF/A, then an SRDF/DM license may also be needed to allow SRDF mode to be changed to Adaptive Copy.

** If BCVs have been configured at Site A.

*** The TimeFinder/Consistency Groups license is needed to create an IO consistency Point-in-time copy of the R2 prior to resynchronisation. It is only needed for the SRDF/A target DMX and only if copies are made when SRDF is still running. The Best Practice for SRDF/A resyncs is to make the copy when SRDF is down, and hence is optional. If, however, a PiT copy is ever needed for COB testing and the application is running when the copy is made, then this license would be needed to guarantee consistency of the copy.

**** Only required at Site C if you intend running multiple-triangle Star workloads from Site C. Very expensive license.

3.6.2 Device Group definitions in SRDF/Star environments

The instructions in both the SRDF/A Solutions Guide Sections 5.5 and 7.3 on defining device groups in a concurrent SRDF environment **should be disregarded** for the following reasons:

-
- Composite Groups are used for SRDF/Star control, and not device groups
 - Subset failover is not supported

Follow the Best Practice in the [Section 4](#) of the present document that relates to use of Composite Groups for Star control.

3.6.3 3-Site SRDF/Star RDF group requirements

As documented in the SRDF/A Solutions Guide, Concurrent SRDF requires SRDF/S between site A and Site B, and SRDF/A between Site A and Site C.

SRDF/Star additionally requires the capability to run SRDF between Sites B and C.

The RA requirements for all three legs have already been covered in [Section 3.4.7.2](#). Here we define the RDF group requirements:

- All ports to be used for SRDF between Site A and Site B should be configured by EMC for use with static RDF group 1 between Sites A and B.
- All ports to be used for SRDF between Site A and Site C should be configured by EMC for use with static RDF group 2 between Sites A and C
- All ports to be used for SRDF between Site B and Site C should be configured by EMC for use with static RDF group 3 between Sites B and C.

The dynamic RDF groups to be used for individual Star sessions (as described in [Section 3.4.2](#)) will be created by Citigroup SAN management teams on top of these static RDF groups. A worked example of the creation of these dynamic RDF groups is given in the next section on setting up SRDF/Star.

4 Setting up SRDF/Star

A concurrent SRDF environment is a pre-requisite for setting up SRDF/Star and must be built first, as per the instructions in the previous section, and references therein. In particular, all physical connectivity must be in place between sites first, and underlying static RDF groups between corresponding ports must also be in place. The examples shown below include the commands needed to set up the specific dynamic RDF groups for individual Star sessions.

The process for setting up Star is described in the SRDF/Star section of the **EMC Solutions Enabler Symmetrix SRDF CLI Product Guide**, and that document should be carefully studied. The process is also covered here in brief as the EMC documentation does have some rather significant omissions relating to set up of the dynamic RDF groups, and configuration of multiple-triangle Star sessions. Here worked examples are provided for:

- Example 1 Configuration SRDF/Star with one triangle per SRDF/Star session
- Example 2: Configuration SRDF/Star with multiple triangles per SRDF/Star session

4.1 Example 1: Setting up SRDF/Star with one triangle per session

The configuration is defined by the following diagram (Figure 5) and table.




Figure 5 – Single RDF triangle in an SRDF/Star session

The Star session in this example will have the following properties:

Site A location & Symmetrix ID	New York – Symm 0822
Site B location & Symmetrix ID	New Jersey – Symm 1497
Site C location & Symmetrix ID	London – Symm 1135
Composite Group Name	flexcg
Synchronous Leg dynamic RDF group	4
Synchronous Leg RDF group name	NewJersey
Asynchronous Leg dynamic RDF group	5
Asynchronous Leg RDF group name	London
Recovery Leg dynamic RDF group	6
BCVs at Site A?	No
BCVs at Site B?	No
BCVs at Site C?	Yes
DSE threshold (Site A ↔ Site C)	40%

4.1.1 Step 1 – Create Concurrent SRDF environment

4.1.1.1 Step 1a – Verify Symmetrix Settings

For each symmetrix to be used, confirm required Symmetrix settings

```
SiteA> symcfg list -v -sid <SID>
```

- Concurrent RDF Configuration State is enabled
- Dynamic RDF Configuration State is enabled
- Concurrent Dynamic RDF Configuration is enabled
- RDF Data Mobility Configuration State is disabled

```
SiteA> symcfg list -rdfg <static RDFG>
```

- Prevent Auto Link Recovery is enabled
- Prevent RAs Online is enabled

4.1.1.2 Step 1b - Create Concurrent SRDF environment

For a detailed description of the steps involved in setting up dynamic RDF groups and device pairs, see the SRDF/A Solutions Guide, Section 5.3.

1) Set up RDF group on Synchronous leg from Site A (0822) to Site B (1497)

```
SiteA> symrdf addgrp -label DYN4 -rdfg 4 -sid 0822 -dir 13C -remote_rdfg 4 -remote_sid 1497 -remote_dir 9C -nop
```

```
SiteA> symrdf modifygrp -add -rdfg 4 -sid 0822 -dir 14C -remote_dir 10C -nop
```

2) Confirm RDF group set up correctly

```
SiteA> symcfg list -sid 0822 -switched -ra all
```

3) Set up RDF groups on Async leg from Site A (0822) to Site C (1135)

```
SiteA> symrdf addgrp -label DYN5 -rdfig 5 -sid 0822 -dir 7D -link_limbo 60 -remote_rdfg 5 -remote_sid 1135 -remote_dir 7D -rem_link_limbo 60 -nop
```

```
SiteA> symrdf modifygrp -add -rdfig 5 -sid 0822 -dir 8D -remote_dir 8D -nop
```

If using Transmit Idle (5772 or later), then link_limbo & rem_link_limbo be set to 10 (the default if neither is specified)

If DSE is to be used for this RDF group (5772 or later), then it need to be configured on both the source and target symmetrix. The following assumes the DSE pool has already been configured at both Site A and Site C. (see SRDF/A Solutions Guide Appendix G)

```
SiteA> symconfigure -f <command-file> -sid 0822 commit
SiteA> symconfigure -f <command-file> -sid 1135 commit
```

where command-file contains:

```
set rdf group 5, rdfa_dse_pool=DSE emulation=fba;
set rdf group 5, rdfa_dse_threshold=40;
set rdf group 5, rdfa_dse_autostart=ENABLE;
```

4) Confirm RDF group set up correctly

```
SiteA> symcfg list -sid 0822 -switched -ra all
```

5) Set up RDF group on recovery leg between Site B (1497) and Site C (1135)

```
SiteA> symrdf addgrp -label DYN6 -rdfig 6 -sid 1497 -dir 7C -link_limbo 60 -remote_rdfg 16 -remote_sid 1135 -remote_dir 7C -rem_link_limbo 60 -nop
```

If using Transmit Idle (5772 or later), then link_limbo & rem_link_limbo be set to 10 (the default if neither is specified)

```
SiteA> symrdf modifygrp -add -rdfig 6 -sid 1497 -dir 7D -remote_dir 8C -nop
```

```
SiteA> symcfg list -sid 1135 -switched -ra all
```

6) Confirm RDF group set up correctly

```
SiteA> symcfg list -sid 1135 -switched -ra all
SiteA> symcfg list -sid 1497 -switched -ra all
```

7) Next create dynamic RDF device pairs

7a) Create the Sync leg first

```
SiteA> symrdf -file ./pairsfile_0822_DYN4 -sid 0822 -rdfig 4 createpair -type rdf1 -establish -nop
```

7b) Now do the async leg

```
SiteA> symrdf -file ./pairsfile_0822_DYN4 -sid 0822 -rdfg 5 createpair -type rdf1
-establish -nop
```

8) Set modes to adaptive copy disk mode for all legs

```
SiteA> symrdf -g DG4 set mode acp_disk -rdfg 4 -sid 0822 -nop
SiteA> symrdf -g DG4 set mode acp_disk -rdfg 5 -sid 0822 -nop
```

4.1.2 Step 2 – Creating SRDF/Star Composite Groups

1) Create composite group

```
SiteA> symcg create flexcg -type RDF1 -rdf_consistency
```

2) Add concurrent SRDF devices to composite group

```
SiteA> symcg -cg flexcg -sid 0822 -rdfg 4 addall dev
```

3) Now define RDF group names and recovery RDF group

```
SiteA> symcg -cg flexcg -rdfg 0822:4 set -name NewJersey
SiteA> symcg -cg flexcg -rdfg 0822:5 set -name London
SiteA> symcg -cg flexcg set -rdfg 0822:4 -recovery_rdfg 6
SiteA> symcg -cg flexcg set -rdfg 0822:5 -recovery_rdfg 6
```

4.1.3 Step 3 – Creating the SRDF/Star options file

Create the SRDF/Star options of the following form:

```
SiteA> cat flexcg_opt

SYMCLI_STAR_WORKLOAD_SITE_NAME=NewYork
SYMCLI_STAR_SYNCTARGET_SITE_NAME=NewJersey
SYMCLI_STAR_ASYNC_TARGET_SITE_NAME=London
SYMCLI_STAR_USE_ADAPTIVE_COPY=both
SYMCLI_STAR_ADAPTIVE_COPY_TRACKS=30000
SYMCLI_STAR_ACTION_TIMEOUT=1800
```

The options file needs to be stored in a standard location as it will be needed again to:

1. dismantle a Star configuration to add/remove disk
2. rebuild a Star configuration

A good location would be wherever you currently store SRDF/A pairs files used for dynamic RDF configuration.

The same options file may be used for many SRDF/Star configurations, if the same parameters are to be used.

4.1.4 Step 4 – Performing the symstar setup operation

Setup Star using this options file

```
SiteA> symstar -cg flexcg -options flexcg_opt setup
```

Now sftp or scp the Star configuration definition file

```
/usr/emc/API/symapi/config/STAR/def/flexcg
```

for this Star session to an identical location on all other Star controller servers at both the local site, and at the synchronous (if in use) and asynchronous target sites.

4.1.5 Step 5 – Creating R2 Composite groups

These commands are issues at the SRDF target sites.

```
SiteC> symstar -cg flexcg buildcg -site London -nop
```

If Star control servers are installed at Site B:

```
SiteB> symstar -cg flexcg buildcg -site NewJersey -nop
```

4.1.6 Step 6 - Adding BCV devices to the SRDF/Star configurations

1) Add BCVs at Site C by associating the BCVs with the composite groups at both Site A and Site C.

Note that the symbcv command does not support the syntax `-rdf -rdfg name:<RDF group Name>`

```
Site A> symbcv -cg flexcg -sid 0822 associateall dev -range <start:end> -rdf -rdfg 5
```

```
Site C> symbcv -cg flexcg -sid 1135 associateall dev -range <start:end>
```

2) Synchronize the BCVs

Before attempting the full establish, ensure that Clone emulation is enabled if not using RAID-5 BCVs. Clone emulation is used if the clone emulation environment variable is set to ENABLED.

```
Site A> SYMCLI_CLONE_EMULATION=ENABLED  
Site A> export SYMCLI_CLONE_EMULATION
```

One can confirm if this (and indeed other) SYMCLI environment variables have been set with the following command:

```
Site A> symcli -def
```

```
Symmetrix Command Line Interface (SYMCLI) Version V6.3 (Edit Level: 623)  
built with SYMAPI Version V6.3 (Edit Level: 623)
```

```
Current settings of the SYMCLI environmental variables:
```

```
SYMCLI_CLONE_EMULATION: ENABLED
```

Note that the establish command must be executed with the `-star` option

```
SiteA> symmir -cg flexcg establish -star -full -rdf -opt_rag -rdfg name:London
```

`-opt_rag` flag is used to create BCV/Standard pairings that do not give DA and/or spindle contention.

4.1.7 Step 7 – Bringing up the SRDF/Star sites for normal operation

```
SiteA> symstar -cg flexcg connect -site NewJersey -nop
SiteA> symstar -cg flexcg connect -site London -nop

SiteA> symstar -cg flexcg protect -site NewJersey -nop
SiteA> symstar -cg flexcg protect -site London -nop

SiteA> symstar -cg flexcg enable -nop
```

4.2 Example 2: Configuration SRDF/Star with multiple triangles per session

It is possible to configure SRDF/Star with multiple triangles per Star session i.e. have multiple RDF groups per leg is single Star session. There are a number of reasons why one might wish to do this:

1. You wish to have one set of RDF groups per application, but have multiple applications bound together from a consistency perspective, in the same way that SRDF/MSR binds multiple 2-site SRDF/A sessions together.
2. One wishes to maintain the capability to remove from Star control an application that is bound from a consistency perspective with other applications
3. One is forced to place multiple applications into the same Star session due to scalability limitations. In this case, applications which the business agrees can be failed over together (e.g. on the same physical server) should be placed in the same Star session.

Use of multiple triangle sessions is hoped not to be a likely occurrence, but this more complex configuration has been fully tested and is documented here. The following examples can be found in Appendix A.

- Configuration of SRDF/Star with multiple triangles per SRDF/Star sessions
- Removing a triangle from a multi-triangle Star setup.

4.3 High-availability of the SRDF/Star environment

4.3.1 Star Control server high-availability

Here a Star Control server is defined as a server used to issue Star control commands i.e. symstar commands.

- Configure minimum of two Solutions Enabler (SE) servers are needed for Star control (for redundancy). Additional dedicated Star control servers should be used at Site A for larger installations (see [Section 3.4.6](#) for Best Practice on use of dedicated Star Control servers). Each server used for Star control (i.e. run symstar commands) must have an identical copy of the /var/symapi/config/STAR/def directory to the original SE server on which Star was configured. This file must be manually copied (sftp or scp) from the original SE server.
- Each SE server used for Star control must have a Composite Group (CG) defined for each Star session. If GNS is in use, additional local SE servers will see the CGs defined when Star sessions were initially setup on the primary local SE server.

If the above are done then the additional servers can be used to control cycle switching and issue symstar commands if the original Star control server is not available.

4.3.2 Storrdfd daemon high-availability

The storrdfd daemon should be run on any available SE servers (Star Control servers and any other non-Star SAN management servers available) for redundancy and load-balancing.

- Each SE server (Star control server or non-Star SE server also available) should have a running storrdfd. This is needed to enable async leg cycle switching in an SRDF/Star configuration and also to allow symstar and certain symrdf commands to be run. SRDF/A Cycle switching will not occur if a storrdfd daemon is not running on a local server at the production site. Also the Symstar command will not run properly if the daemon is not running on the Star control server issuing the commands.
- Storrdfd can be made highly available by following the instructions in Section 6.3.2 in the Citigroup **SRDF/A for Long Distance Replication Solutions Guide**. This is utterly essential.

If the above are done then the additional servers can be used to control cycle switching if one SE server (or storrdfd on that server) is not available.

Note that even when no daemons are running on the workload site, Star will trip successfully if Site A is lost.

5 Site Switching (Failover) Operational Procedures

5.1 Introduction

5.1.1 Nomenclature

Throughout this document, the following Site naming convention is used:

Site Role	Site Role Short name
Production site / Source site	Site A
Synchronous SRDF target site / bunker site	Site B
Asynchronous SRDF target site / COB site	Site C

In addition, the reader should be aware that EMC refer to the failover operations as **switching** operations. This language is used throughout this document for consistency with the EMC documentation.

The site that is currently running production, is also known as the **workload site**. The workload site can move between Sites A, B and C. Star operations also refer to the nickname assigned to the sites. These nicknames, which reflect physical sites, do not change regardless of where the production is currently running.

5.1.2 Other documentation

There are two key EMC documents on SRDF/Star that any operator of SRDF/Star must become very familiar with. Both can be found in the **EMC Solutions Enabler Symmetrix SRDF Family CLI Product guide**:

- **Chapter 4 – SRDF/Star.** Discusses use of the symstar command for setup, management and failover data operations in the event of a site or link failure.
- **Appendix B – SRDF/Star State Reference.** Describes the allowable SRDF/Star configuration states for using the symstar command arguments. It is critical that operators understand these different states and how to move between them.

The present document covers key operational tasks, but the EMC documentation gives important pre-requisite detailed background and additional information, in particular reference material on the complex issue of SRDF/Star states, and an overview of all symstar control operations.

5.1.3 Site switching multiple Star sessions

In the event of a Site or Array failure, it will be necessary to switch multiple Star sessions to Site C.

Detailed testing of site switching multiple Star sessions has been conducted for SE6.4 (see “SRDF/Star Scalability Testing – Key findings” document referenced in Section 1.3). It is recommended that switching operations be done sequentially with SE6.4 as attempts to run

multiple switches in parallel on the same Star control server can lead to locking problems (e.g. the SYMAPI DB, amongst others), resulting in very extended runtimes for the switch operations.

To speed site switching of many application in a disaster, it is recommended that only the disconnect switch operations be executed initially for each Star session. The subsequent protect and connect steps to provide ongoing COB protection can be executed after all Star sessions have been switched and applications have access the Site C disks.

The order in which Star sessions will be switched in the event of a disaster should be agreed with the business in advance.

Template scripts to help automate switching of multiple Star sessions will be referenced in a future release of this document after testing as part of in the EMEA Star deployment.

5.1.3.1 EMC initiated failover

At the time of writing, EMC do not support zero-data loss failover of Star sessions for an entire array. Star switching must be executed by the customer.

5.1.4 Notes on impact of Transmit Idle in SRDF/Star configurations

The following apply when Transmit Idle is enabled on the Async leg (5772 or later).

- The following controls are blocked on the async leg at the workload site when in the TransIdle state: Protect, Halt, Enable, Disconnect, Reset, Unprotect, and Isolate. Hence, for example, it will not be possible to do a planned switch as the halt operation cannot be run. An unplanned switch would need to be executed instead, if really necessary.
- When a link is in Transmit Idle state, the Star state for that link will be “protected” as SRDF/A has not actually dropped.
- A switch command cannot be executed at the Async target site when the async leg is in Transmit Idle state (as seen from the Async target site). A workaround is given in the unplanned switch sections below.

5.2 Site B as bunker only – site switching scenarios

The Global Data Center Strategy team has determined that Citigroup's 3-site replication solution will not have servers at Site B. Site B will purely be used as a data bunker. Site C will be the DR site protecting Site A.

In this section the different failure and subsequent failover scenarios are discussed in some detail for the Site B as bunker-only configuration. Note that this configuration does not provide true redundant protection (ongoing COB protection) after a site or regional disaster as there are no servers at Site B.

5.2.1 Server and application level failover/switch scenario

Star failover is at the Composite Group (CG) level (device groups are not used). Individual devices within a subset cannot be failed over, or “switched” as it is known for SRDF/Star. Hence, data server level failover can only be supported in a Star configuration where each data server has its own CG and Star session.

It is likely that individual applications would have their own composite groups. If a data server does fail and it is necessary to switch to Site C, an entire CG must be switched i.e. all data servers in the CG (although of course often an application may only have a single data server). Use of high-availability clustering would ensure that switch to Site C is not needed in the event of just a single data server failure.

If Application at Site A fails, production will resume at **Site C** with zero data loss, as Site A storage is still available. **Failover will be via a planned switch to Site C.**




Figure 6 - Application level failover/switch (with ongoing COB)

Asynchronous Replication from COB/Recovery site (Site C) to the Site A could then continue i.e. Site A is now providing COB for Site C. **When failed data server(s) are available again, production should be reverted back to Site A via a planned switch as soon as possible to regain zero data loss capability again.**

5.2.2 Site A failure scenario

If Site A fails cleanly, production will resume at **Site C** with potentially zero data loss as Site C, as Star can differentially transfer data that is missing from Site C from the up-to-date Synchronous target Site B. **This is a unique SRDF/Star function.** If there is a rolling disaster, there may be some data loss if SRDF/S links between Site A and Site B fail before either the application servers or the Site A Symmetrix fail.

Asynchronous Replication from COB/Recovery site (Site C) to the Site B could then start i.e. Site B is now providing COB for Site C from a storage perspective only, as there are no servers there. There would arguably be limited value in replicating to Site B in this manner, but the infrastructure

will be in place, and will offer a degree of extra protection, of the data at least, but not a true COB facility at Site B.

If Site A storage becomes available again, production should be reverted back to Site A as soon as possible to ensure that the application is Star-protected again.




Figure 7 – Site A failure Scenario (limited ongoing COB)

5.2.3 Regional disaster scenario

If Site A and Site B both fail, production will resume at Site C, with potentially some data loss (see below). If failed sites become available again (failure may simply be a wide-spread power outage), production can revert back to Site A.




Figure 8 – Regional Disaster Scenario (no ongoing COB)

5.3 Controlling SRDF/Star

5.3.1 symstar

When SRDF is under Star control, it must be controlled using the **symstar** command. The **symrdf** command is not used, and indeed non-query **symrdf** commands are blocked with Star is configured. The **-star** qualifier can be used to issue a **symrdf** command to an SRDF session under Star control, but under normal circumstance, this should never be used. It would only be used to fix up a Star session should there be a **symstar** failure that leaves SRDF in an abnormal state that cannot be fixed with **symstar** commands, or to deal with Transmit Idle issues (see [Section 5.6](#)).

Symstar commands are complex compound commands that run many **symrdf** commands under the covers. Hence, they can take several minutes to run. You should never abort a **symstar** session part way through, as it may leave the Star session in an abnormal state that may be difficult to recover from.

Symstar commands are intended to be re-runable should they fail part way through. Hence, if a command does fail, or is accidentally aborted, then one should first attempt to re-run the command and see if it is able to complete. If not, then re-run again with the **-force** flag. This flag should be used with care as it forces an operation when RDF is not in a normal pre-requisite state for the operation. This will attempt the requested action even if the STAR system is not in the correct state (symstar carefully checks system states are as expected before attempting any operations).

As SRDF/Star blocks most symrdf commands, most of the processes and scripts referenced in the SRDF/A Solutions Guide do not apply to, or will not work in an SRDF/Star environment and should not be used. Alternate operational practices are given in this document.

Template scripts to help automate certain **symstar** command invocations for multiple Star sessions will be referenced in a future release of this document after testing as part of in the EMEA Star deployment. The rest of this section covers the procedures for individual Star sessions.

5.3.2 symstar arguments

The **symstar** command covers virtually all aspects of SRDF/Star operations, and has many arguments. These are described in the EMC SRDF/Star product documentation referenced earlier, and is not repeated here. However, it is important to note that:

- Most commands are executed at the current workload site (i.e. the site where production is currently located).
- Exception: **symstar switch** command are run from the site to which the workload is being switched to i.e. the target site.
- Exception: **symstar cleanup** commands are run at the target site.
- Exception: Query and show commands can be run from any site.

5.3.3 SRDF/Star States

As indicated earlier, to perform a symstar command, the SRDF/Star configuration needs to be in an allowable system state. Otherwise, a message is returned, stating that SRDF/Star is not in a state that permits the particular operation that you are attempting to perform.

The allowable states for each action are very clearly detailed in the **EMC Solutions Enabler Symmetrix SRDF Family CLI Product guide, Appendix B**. Anyone involved with managing an SRDF/Star environment must read this above reference and become very familiar with the different SRDF/Star states and which operations can be used to move Star between these different states.

5.4 Using BCVs with Star operations

As mentioned earlier in [Section 3.5.4](#), BCVs should be configured into an SRDF/Star environment for protection during SRDF resyncs. This section describes Best Practice for when to use these BCVs.

5.4.1 General Best Practice

1. BCVs should usually be established with their R2 standards under normal operations
2. BCVs **MUST** be split off before attempting a switch operation when using RAID1+0 and “Classic” TimeFinder/Mirror. This is necessary to free up a mirror position, which will be used during the switch operation as new RDF relationships are created on-the-fly. The switch operation will not be permitted if there are insufficient mirror positions available to complete the switch. For further details of this issue, see the document **Open Systems SRDF/Star (SE6.3) Evaluation and Test Results** (Section 5.1).

As a consequence, the Best Practice for BCVs in SRDF/Star is to use TimeFinder/Mirror Clone Emulation. “Classic” TimeFinder/Mirror can be used, but care must be taken before switch operations to split off the BCVs to avoid the above Mirror position exhaustion problem.

5.4.2 Why make a point-in-time copy?

The BCV point-in-time copies are needed to ensure that there is a consistent copy of the data at all times at the target site during an SRDF resynchronization exercise (e.g. following a network outage or a cessation of replication). If this is not done in a 2-site SRDF configuration, then there is a danger that if the source site is lost during the resync, then there will be no consistent copies of the data. Hence for 2-site SRDF/A the recommendation is to always split off a BCV copy of the R2 prior to a resync operation.

In a 3-site SRDF/Star solution, there are more copies of the data available. Hence, the following Best Practice is recommended:

In a **planned switch**, there will be 3 identical copies of the data (one at each site) following the halt operations. Hence, as no data will be flowing, it is not necessary to take a BCV copy before the switch is executed from a data protection perspective. If one site is lost, there will still be two

good copies at the other two sites. However, the BCV will need to be split off to create sufficient mirror positions for the split operation, if using "Classic" Timefinder/Mirror.

During an **unplanned switch to Site C**, following the loss of Site A, a BCV copy should be split off from the R2 at Site C prior to the switch taking place:

- To create the spare mirror position (required if using RAID1)
- To create a consistent point-in-time copy during a differential resync with Site B (i.e. when switching to Site C but keeping the data from Site B)

Summary: Always ensure point-in-time protection BCVs are split off prior to a switch operation

5.5 Executing a planned switch

This section describes how to execute planned switch operations between sites. Planned switch operations are used when the current workload site is available. Planned switches are used when the storage infrastructure is still intact. Hence, planned switches are used:

- when moving production from one site to another following a data server or application failure.
- when moving production back to its original site, following a previous switch operation.
- For certain types of COB testing.

5.5.1 Planned switch workload to Async target site (standard process)

This operation would take place for the most common scenario where there has been a data server or application failure and it is necessary to switch operations to the remote COB/DR site. In this example, the switch will be from Site A to Site C, as Site C is the designated COB/DR site in Citigroup implementations of Star. A planned switch will take at least 5 minutes to execute.

As the storage infrastructure is still operational, there will be no data loss at Site C.

Note that this switch provides on-going COB (Site C to Site A replication), but not zero-data loss protection as there is no local bunker site near Site C in the Star configuration.

This procedure assumes that BCVs are only installed at Site C. Note that when running out of Site C, Star cannot be enabled and one is not running in a zero-data loss mode. Production should be moved back to Site A as soon as possible so that Star can be enabled again, providing the zero data loss capability.

1) Shutdown application at Site A

It is essential that the application is not accessing the disks when the halt operation takes place. Volume manager disks groups ideally should be exported.

2) Halt Star operations at Site A (current workload site).

Halt operation write-disables devices at the workload site. Hence, it is vital that the application is down before a symstar halt command is issued. After a halt operation, all sites have identical data. During testing, a halt operation typically took 2-3 minutes to complete.

```
SiteA> symstar -cg <cgname> halt -nop
```

3) Split off point-in-time copy at Site C

```
SiteA> symmir -cg <cgname> split -star -rdf -rdfg name:<SiteC> -nop
```

or

```
SiteC> symmir -cg <cgname> split -star -nop
```

4) Switch operations to Site C

All subsequent steps take place at the site at which operations are being brought up (Site C). During testing, this switch step would typically take about 4 minutes to complete.

```
SiteC> symstar -cg <cgname> switch -site <SiteC> -nop
```

5) Start SRDF/A replication back to Site A (the original workload site)

```
SiteC> symstar -cg <cgname> connect -site <SiteA> -nop
```

```
SiteC> symstar -cg <cgname> protect -site <SiteA> -nop
```

6) Connect to Site B

This starts adaptive copy to Site B. This is required to keep Site B from being too far out of date compared with Sites A and C (although it will be offering no form of data protection as the data at Site B will be inconsistent). It is also necessary to have Site B connected to allow a later halt operation prior to a switch back to Site A.

```
SiteC> symstar -cg <cgname> connect -site <SiteB> -nop
```

5.5.2 Planned switch back from Async target to workload site

This operation would take place once the application is ready to switch production operations back to Site A.

1) Shutdown application at Site C

It is essential that the application is not accessing the disks when the halt operation takes place. Volume manager disk groups ideally should be exported.

2) Halt Star operations at Site C (current workload site)

Halt operation write-disables devices at the current workload site. Hence, it is vital that the application is down before a symstar halt command is issued. After a halt operation, all sites have identical data. During testing, a halt operation typically took 2-3 minutes to complete.

```
SiteC> symstar -cg <cgname> halt -nop
```

3) Switch operations to Site A

All subsequent steps take place at the site at which operations are being brought up (Site A). During testing, this particular switch step would typically take about 5-6 minutes to complete.

```
SiteA> symstar -cg <cgname> switch -site <SiteA> -nop
```

4) Bring composite group under Star protection

```
SiteA> symstar -cg <cgname> connect -site <SiteB> -nop
```

```
SiteA> symstar -cg <cgname> connect -site <SiteC> -nop
```

```
SiteA> symstar -cg <cgname> protect -site <SiteB> -nop
```

```
SiteA> symstar -cg <cgname> protect -site <SiteC> -nop
```

```
SiteA> symstar -cg <cgname> enable -nop
```

5) Re-establish point-in-time copy at Site C

Re-establish point-in-time copy at Site C that was split off during switch to Site C.

```
SiteA> symmir -cg <cgname> establish -star -rdf -rdfg name:<SiteC> -nop
```

or

```
SiteC> symmir -cg <cgname> establish -star -nop
```

5.5.3 Planned switch workload to Sync target site (non standard process)

This is **not** a process that is anticipated to be used within Citigroup, as the Data Centre Strategy is for Site B to be a serverless site and hence there is little point in switching to Sync target site. However, the process is documented here should there arise a circumstance where this is required. Star control servers would need to have been configured in Site B to execute a switch to Site B.

A switch to Site B (sync target) allows Star to be re-enabled, and hence the process differs from switch to Site C (async target).

This procedure assumes that BCVs are only installed at Site C.

1) Shutdown application at Site A

It is essential that the application is not accessing the disks when the halt operation takes place. Volume manager disk groups ideally should be exported.

2) Halt Star operations at Site A (current workload site)

Halt operation write-disables devices at the workload site. Hence, it is vital that the application is down before a symstar halt command is issued. After a halt operation, all sites have identical data.

```
SiteA> symstar -cg <cgname> halt -nop
```

3) Split off point-in-time copy at Site C

```
SiteA> symmir -cg <cgname> split -star -rdf -rdfg name:<SiteC> -nop
```

4) Switch operations to Site B

All subsequent steps take place at Site B.

```
SiteB> symstar -cg <cgname> switch -site <SiteB> -nop
```

5) Bring composite group under Star protection

```
SiteB> symstar -cg <cgname> connect -site <SiteA> -nop  
SiteB> symstar -cg <cgname> protect -site <SiteA> -nop
```

```
SiteB> symstar -cg <cgname> connect -site <SiteC> -nop  
SiteB> symstar -cg <cgname> protect -site <SiteC> -nop  
SiteB> symstar -cg <cgname> enable -nop
```

5.5.4 Planned switch back from Sync target to workload site

This operation would take place once the application is ready to switch production operations back to Site A.

1) Shutdown application at Site B

It is essential that the application is not accessing the disks when the halt operation takes place. Volume manager disks groups ideally should be exported.

2) Halt Star operations at Site B (current workload site)

```
SiteB> symstar -cg <cgname> halt -nop
```

3) Switch operations to Site A

```
SiteA> symstar -cg <cgname> switch -site <SiteA> -nop
```

4) Bring composite group under Star protection

```
SiteA> symstar -cg <cgname> connect -site <SiteB> -nop  
SiteA> symstar -cg <cgname> connect -site <SiteC> -nop
```

```
SiteA> symstar -cg <cgname> protect -site <SiteB> -nop  
SiteA> symstar -cg <cgname> protect -site <SiteC> -nop
```

```
SiteA> symstar -cg <cgname> enable -nop
```

5) Re-establish point-in-time copy at Site C

```
SiteA> symmir -cg <cgname> establish -star -rdf -rdfg name:<SiteC> -nop
```

5.6 Responding to disaster faults – “unplanned” switches

This section describes how to execute an “unplanned” switch. An unplanned switch is one where the switch is forced by failure of the current production Symmetrix or site (Site A). It is not used when a server or application fails. In that circumstance, a planned switch should be used.

SRDF/Star is very flexible in these circumstances. There are four different unplanned switch scenarios available:

Operation	Citigroup Supported	Notes
Switch workload to Async target site (Site C), keeping Synchronous Target data from Site B	Yes	Most common usage in Citigroup. Offers zero data loss in non-rolling disaster scenario.
Switch workload to Async target site (Site C), keeping Asynchronous Target data at Site C	Yes	Would only be used in a rolling disaster where the async site has more recent data than the sync site.
Switch workload to Sync target site (Site B), keeping Synchronous data at Site B	No	If switching to Sync site, this would be the model used. Note, however, that Site B is a bunker site in the Citigroup model and contains no servers.
Switch workload to Sync target site (Site B), keeping Asynchronous data from Site C	No	Only applies in a rolling disaster where the async site has more recent data than the sync site.

However, only switches to the Async target site (Site C) are supported in Citigroup, as the Sync target site (Site B) are intended to be serverless. Hence, unplanned switches to the Sync target site (Site B) are not documented here. Sync site switches have been thoroughly tested and hence this document could easily be updated to support these should Data Centre strategy change in the future.

5.6.1 Determining from which site to keep data

A rolling disaster describes a disaster marked by different start and end points i.e. non-instantaneous events. Depending on the order in which the IT infrastructure fails, one cannot guarantee zero data loss in a rolling disaster. For example, if the synchronous replication links go down before the production storage or servers, Star will not be able to provide zero data loss at the long distance site. One can choose if the data from the Site B or Site C is used to start production at Site C.

In a non-rolling disaster scenario, the Sync target site will be more up-to-date than the Async target and the switch to Site C should be executed with the –keep <Site B> qualifier.

In a rolling disaster, it is possible that the Async site may end up with more recent data. If that is the case, the switch to Site C should be executed with the –keep <Site C> qualifier.

The most up-to-date site can be determined using the symstar query command:

```
Site C> symstar -cg <cgname> query | grep "most Current"
```

```
Asynchronous Target Site Data most Current      : No
```

If the answer is no, switch to Site C keeping Sync (Site B) data – [Section 5.6.3](#)

If the answer is yes, switch to Site C keeping Async (Site C) data – [Section 5.6.4](#)

5.6.2 Scheduling multiple switch operations

In the case of a disaster, there may be multiple Star sessions per Symmetrix to switch to the DR site. However, it is recommended by both EMC and Citigroup Storage Engineering that these be done sequentially with SE6.4 to avoid resource contention (e.g. gatekeepers) and limit the impact to a single application should there be issues during the switching operation. See [Section 5.1.3](#) or the separate document "SRDF/Star Scalability Testing – Key findings" (see Section 1.3 for reference) for an expanded discussion.

Sessions on separate Symmetrix and separate management servers could be switched in parallel.

Template scripts to help automate switching of multiple Star sessions will be referenced in a future release of this document after testing as part of in the EMEA Star deployment.

5.6.3 Switch workload to Async target site (Site C), keeping Sync (Site B) data

This procedure should be used in a disaster scenario when Site A Symmetrix is unavailable and output from `symstar -cg <cgname> query` shows:

```
Asynchronous Target Site Data most Current      : No
```

1) Check state of Star session

```
SiteC> symstar -cg <cgname> query
```

One would expect to see that the following (or very similar):

```
System State:
{
  1st_Target_Site      : Failed
  2nd_Target_Site      : PathFail;CleanReq
  STAR                 : Tripped
}
```

Also all the RDF pair states should be "Partitioned" or "Transmit Idle" state.

2) Split off point-in-time copy at Site C

This should be done, regardless of RAID type, to ensure a consistent point-in-time copy at Site C at all times, as Site C will inconsistent whilst “missing” data is copying from Site B to Site C>

```
SiteC> symmir -cg <cgname> split -star -nop
```

3a) Force async leg out of Transmit Idle state

If Transmit Idle is enabled on the asynchronous leg (5772 or later), then as seen from Site C, the SiteA ↔ Site RDF link will be in Transmit Idle state (rather than partitioned) and the 2nd_Target_site Star state will still be marked as “protected”.

In this state it is not possible to execute a site switch to Site C. Instead, it is necessary (at the time of writing) to first manually suspend the async leg with the following command:

```
SiteC> symrdf -cg <cgname> suspend -immediate -rdfg name:<SiteA> -star -force -nop
```

before attempting the the switch operation.

3b) Clean Async target site

This may be needed to clean up the Async Target site when SRDF/A is not halted in a clean manner (e.g. during a disaster at Site A).

```
SiteC> symstar -cg <cgname> clean -site <SiteC> -nop
```

4) Switch operations to Site C

```
SiteC> symstar -cg <cgname> switch -site <SiteC> -keep <SiteB> -nop
```

When this operation has completed, the application can be restarted at Site C.

5) Start SRDF/A replication to Site B to provide on-going data protection

2-Site SRDF/A replication can now be started from Site C to Site B.

Note that if you are switching multiple separate star sessions, you may wish to first complete the switches for all Star sessions before executing any protect or connect commands in order to bring all applications on-line as quickly possible.

```
SiteC> symstar -cg <cgname> protect -site <SiteB> -nop
```

Note that it is not necessary to execute a

```
SiteC> symstar -cg <cgname> connect -site <SiteB> -nop
```

command in this case as the “switch to Site C” command causes a connect command to be executed for Site B, when keeping remote data as Site B has to be connected to allow the differential update of Site C from Site B to take place.

5.6.4 Switch workload to Async target site (Site C), keeping Async (Site C) data

This procedure should be used in a disaster scenario when Site A Symmetrix is unavailable and output from `symstar -cg <cgname> query` shows:

```
Asynchronous Target Site Data most Current    :   Yes
```

1) Check state of Star session

```
SiteC> symstar -cg <cgname> query
```

One would expect to see that the following (or very similar):

```
System State:
{
  1st_Target_Site           : Failed
  2nd_Target_Site           : PathFail;CleanReq
  STAR                      : Tripped
}
```

Also all the RDF pair states should be “Partitioned” or “Transmit Idle” state.

2) Split off point-in-time copy at Site C

```
SiteC> symmir -cg <cgname> split -star -nop
```

3a) Force async leg out of Transmit Idle state

If Transmit Idle is enabled on the asynchronous leg (5772 or later), then as seen from Site C, the SiteA ↔ Site RDF link will be in Transmit Idle state (rather than partitioned) and the 2nd_Target_site Star state will still be marked as “protected”.

In this state it is not possible to execute a site switch to Site C. Instead, it is necessary (at the time of writing) to first manually suspend the async leg with the following command:

```
SiteC> symrdf -cg <cgname> suspend -immediate -rdfg name:<Site A> -star -force -nop
```

before attempting the the switch operation.

3b) Clean Async target site

This may be needed to clean up the Async Target site when SRDF/A is not halted in a clean manner (e.g. during a disaster at Site A).

```
SiteC> symstar -cg <cgname> clean -site <SiteC> -nop
```

4) Switch operations to Site C keeping Site C data

```
SiteC> symstar -cg <cgname> switch -site <SiteC> -keep <SiteC> -nop
```

When this operation has completed, the application can be restarted at Site C.

5) Start SRDF/A replication to Site B to provide on-going data protection

2-Site SRDF/A replication can now be started from Site C to Site B.

Note that if you are doing multiple star switches, you may wish to first complete the switches for all Star sessions before executing any protect or connect commands in order to bring all applications on-line as quickly possible.

```
SiteC> symstar -cg <cgname> connect -site <SiteB> -nop  
SiteC> symstar -cg <cgname> protect -site <SiteB> -nop
```

5.7 Switching back to Site A after a temporary disaster

If the unavailability of Site A is temporary, in due course one will want to switch operations back to Site A. The process is as follows:

1) Connect Site A

```
SiteC> symstar -cg <cgname> connect -site <SiteA> -nop
```

2) Execute a planned switch from Site C to Site A

This process is documented in [Section 5.5.2](#)

5.8 Ongoing SRDF operations after a disaster

It is not possible to give detailed instruction on day-to-day operations after a disaster has struck as it depends heavily on the nature of the disaster. For example, is the disaster temporary or permanent? If temporary, its duration will heavily influence what specific post-disaster actions are taken. Hence, any advice here can only talk in terms of generalizations:

- If Site A is permanently out of action, and one is reduced to running SRDF/A from Site B to Site C, it would be worth deconfiguring Star and setting up as a conventional 2-Site SRDF/A configuration to permit a conventional SRDF failover from Site C to Site B should this then become necessary. Testing of further SRDF/Star switches when only two sites remain is not documented by EMC and has not been tested by GDSE, as Star is only targeted at present for use in Citigroup for zero-date loss, not on-going COB.
- If Site A is temporarily unavailable, then the Star configuration should remain in place in anticipation of Site A rejoining the configuration. Should Site C subsequently fail, native symrdf failover commands can be forced to execute when Star is configured using the –

`star` qualifier. SRDF/Star can also be deconfigured using the `symstar setup -clear` command (see [Section 6.5](#) for an example).

6 Other Operational Procedures

6.1 List symstar sessions on Symmetrix

There is no EMC command to do this. However, on a Star-enabled Symmetrix there will typically only be a composite group defined where it is needed for Star control. Hence, the list of composite groups will usually correspond to the list of Star sessions.

An alternative check is to look at the Star definition files in `/var/symapi/config/STAR/def`. There will be a definition file for each Star session.

On a system with many Star sessions, the script similar to that in Appendix B can be used to give an overview of the status of all the Star sessions visible on a Star control server. Follow up individual symstar query commands will then be needed for composite groups where the Star status is not “protected”, as this indicates that Star is not fully enabled.

6.2 Checking status of a specific Star session

The following commands can be run from any site.

```
# symstar -cg <cgname> show
```

will display the contents of the internal SRDF/Star definition file that was created when symstar setup command was run. When run with the `-detail` option, it also shows all the devices in the SRDF/Star session, as shown in Appendix C.

```
# symstar -cg <cgname> query
```

will display local and remote Symmetrix array information and the status of SRDF pairs in the composite group. An example is shown in Appendix D.

This is a very useful command for checking the status of a Star session. In particular, the author uses this command to check

- Current workload site
- Site nicknames
- Checking Star system state
- Last symstar operation performed, with time and status of that operation
- If Async target site data is most current
- RDF modes
- RDF pair states

6.3 Isolating an SRDF/Star site

This process is described in the **EMC Solutions Enabler Symmetrix SRDF CLI Product Guide**, but is included here for completeness. This process would be used for example, to avoid Star tripping, if one of the remote sites were having a powerdown, or perhaps for a DR test at the remote site.

6.3.1 Isolating a site

Isolating a remote target site from the Star environment:

- Disables RDF Consistency to the Remote Target site
- Perform an SRDF 'split' operation to make the remote Target's R2 devices RW to the host.

```
SiteA> symstar -cg <cgname> disable -nop  
SiteA> symstar -cg <cgname> isolate -site <site> -nop
```

Where <site> is the nickname for the site to be isolated.

6.3.2 Rejoining an isolated site

1) Split off point-in-time copy at Site C, if rejoining site C

```
SiteA> symmir -cg <cgname> split -star -rdf -rdfig name:<SiteC> -nop
```

2) Rejoin isolated sites

```
SiteA> symstar -cg <cgname> disconnect -site <site> -nop  
SiteA> symstar -cg <cgname> connect -site <site> -nop  
SiteA> symstar -cg <cgname> protect -site <site> -nop  
SiteA> symstar -cg <cgname> enable -nop
```

Where <site> is the nickname for the site that was isolated and is being brought back under Star protection.

3) Re-establish point-in-time copy at Site C, if previously split off

```
SiteA> symmir -cg <cgname> establish -star -rdf -rdfig name:<SiteC> -nop
```

6.4 Transient remote fault recovery

A transient fault is defined as a fault with the network or the remote storage array that does not disrupt the production site. Data transmission to the remote site is of course impacted. Once the network or remote storage array is available again, it may be necessary to reset the system to provide Star protection again.

If Transmit Idle is enabled on the asynchronous leg (5772 or later), and SRDF/A does not drop during a transient fault on that leg, then no recovery is needed as Star will not trip. As long as SRDF/A does not drop, then the asynchronous leg will automatically recover when the network fault is resolved, and hence no manual fault recovery is needed.

In the following example, <site> is the nickname of the site that experienced the transient fault. This procedure assumes that BCVs are only installed at Site C.

1) Issue a reset command

Warning: Note that once the following reset command is issued, one cannot execute differential resync between Sites B and C again until the symstar enable in Step 3 has completed. The reset is, however, necessary before the site affected can be brought back under Star protection. In the unlikely event that Site A should fail during the resync, you will not be able to use symstar switch commands to bring production up at Site C after the reset is issued, as a switch requires the Star state to either be PathFailed; Star Tripped or Halted. You would need to use the data on the point in time copy made in Step 2.

```
SiteA> symstar -cg <cgname> reset -site <site> -nop
```

2) Split off point-in-time copy at Site C, if rejoining site C

This is only needed if the fault was on the link from Workload Site to Site C, and SRDF/A dropped as a result. It is required as Site C will become inconsistent after the “symstar connect” command whilst data written at Site A whilst the network was down is copied to Site C.

```
SiteA> symmir -cg <cgname> split -star -rdf -rdfg name:<SiteC> -nop
```

3) Bring site under Star protection again

```
SiteA> symstar -cg <cgname> connect -site <site> -nop
SiteA> symstar -cg <cgname> protect -site <site> -nop
SiteA> symstar -cg <cgname> enable -nop
```

4) Re-establish point-in-time copy at Site C, if previously split off

```
SiteA> symmir -cg <cgname> establish -star -rdf -rdfg name:<SiteC> -nop
```

6.5 Rebuilding Star configuration when adding/removing disks

The process of adding or removing devices from an SRDF/Star configuration is particularly onerous, and the reason that highly dynamic environments are not at all well suited for SRDF/Star. This section gives an overview and examples of how to add and remove devices.

The undesirable aspects of this process include:

- SRDF protection is disabled for much of the duration of the process
- The entire Star configuration must be rebuilt from scratch

It is strongly recommended that applications under SRDF/Star protection be allocated sufficient storage before Star is set up to accommodate future growth, and avoid the following process.

The following steps can be executed in several different ways. In this Solutions Guide, a relatively simple process is documented, in an attempt to minimize human error. It will not be the fastest process, but avoids shortcuts that may lead to errors, both from humans or subsequently from SRDF/Star. This reconfiguration of Star should only be executed during a green zone for the application, as SRDF protection is suspended during the process.

6.5.1 Example of adding devices to an SRDF/Star configuration

1) Take a point-in-time copy at Site C

```
SiteA> symmir -cg <cgname> split -star -rdf -rdfg name:<SiteC> -nop
```

2) Disable Star and SRDF protection

It is necessary to unprotect both sites (which will change replication modes to adaptive copy) before the Star configuration can be cleared

```
SiteA> symstar -cg <cgname> disable -nop  
SiteA> symstar -cg <cgname> unprotect -site <SiteC> -nop  
SiteA> symstar -cg <cgname> unprotect -site <SiteB> -nop
```

Replication mode to both sites will now be adaptive copy, and the remote R2 volumes will be inconsistent if any data is changing at the R1 site.

2) Clear Star setup and remove Composite Group

Execute a `symstar setup` command with the `-clear` flag. Use the same options file originally used to set up Star

```
SiteA> symstar -cg <cgname> -options <optionsfile> -clear setup -nop
```

Note that it is not strictly necessary to remove the Composite Groups (CGs), but this author prefers to build the CGs from scratch rather than add devices to them.

```
SiteA> symcg -cg <cgname> rmall -rdfg <Sync RDF group> -nop
```

```
SiteA> symcg delete <cgname> -force
```

```
SiteB> symcg -cg <cgname> rmall -rdfg <Sync RDF group> -nop  
SiteB> symcg delete <cgname> -force
```

```
SiteC> symcg -cg <cgname> rmall -rdfg <Async RDF group> -nop  
SiteC> symcg delete <cgname> -force
```

3) Add new devices into the Asynchronous RDF group and establish

There are several approaches to doing this, and the following example is just one of several approaches:

```
# symrdf -file <newdev_pairsfile> -sid <sid> -rdfg <Async RDF group>  
createpair -type rdf1 -establish -nop
```

where newdev_pairsfile only contains the new devices being added to the RDF group

If removing devices from a Star configuration, use the deletepair option and a pairsfile that lists the devices to be removed.

4) Add new devices into the Synchronous RDF group and establish

```
# symrdf -file <newdev_pairsfile> -sid <sid> -rdfg <Sync RDF group>  
createpair -type rdf1 -establish -nop
```

If removing devices for a Star configuration, use in the deletepair option and a pairsfile that lists the devices to be removed.

5) Add devices into the Composite Group at Site A

1) Recreate composite group with all devices now the RDF group

```
SiteA> symcg create <cgname> -type RDF1 -rdf_consistency
```

2) Add concurrent SRDF devices to composite group

```
SiteA> symcg -cg <cgame> -sid <sid> -rdfg <Sync RDF group> addall dev
```

3) Now define sync, async and RDF group names again

```
SiteA> symcg -cg <cgname> -rdfg <sid>:<Sync RDF group> set -name <SiteB>  
SiteA> symcg -cg <cgname> -rdfg <sid>:<Async RDF group> set -name  
<SiteC>  
SiteA> symcg -cg <cgname> set -rdfg <sid>:<Sync RDF group> -  
recovery_rdfg <Recovery RDF group>  
SiteA> symcg -cg <cgname> set -rdfg <sid>:<Async RDF group> -  
recovery_rdfg <Recovery RDF group>
```

For example:

```
SiteA> symcg create flexcg -type RDF1 -rdf_consistency
SiteA> symcg -cg flexcg -sid 0822 -rdfg 4 addall dev
SiteA> symcg -cg flexcg -rdfg 0822:4 set -name NewJersey
SiteA> symcg -cg flexcg -rdfg 0822:5 set -name London
SiteA> symcg -cg flexcg set -rdfg 0822:4 -recovery_rdfg 6
SiteA> symcg -cg flexcg set -rdfg 0822:5 -recovery_rdfg 6
```

6) Run symstar setup against new Composite Group

This step should not be executed until the synchronizations of the new devices have completed. Note that for a larger number of devices this could take some time and during that interval, the application is not protected. This is why adding disk to an existing Star configuration is not recommended. Instead, Star-protected systems should have sufficient disk space for anticipated growth allocated in advance, if possible.

```
SiteA> symstar -cg <cgname> -options <optionsfile> setup
```

Now sftp the Star configuration definition file

```
/usr/emc/API/symapi/config/STAR/def/<cgname>
```

for this Star session to an identical location on all other Star controller servers at both the local site, and at the synchronous (if in use) and asynchronous target sites.

7) Create R2 Composite groups

These commands are issued at the SRDF target sites.

```
SiteC> symstar -cg flexcg buildcg -site < SiteC> -nop
```

If Star control servers are installed at Site B:

```
SiteB> symstar -cg flexcg buildcg -site <SiteB> -nop
```

8) Bring SRDF/Star up for normal operations

Note that this step is undertaken as early as possible to minimize the amount of time that the application is not protected by SRDF.

```
SiteA> symstar -cg <cgname> connect -site <SiteB> -nop
SiteA> symstar -cg <cgname> connect -site <SiteC> -nop

SiteA> symstar -cg <cgname> protect -site <SiteB> -nop
SiteA> symstar -cg <cgname> protect -site <SiteC> -nop

SiteA> symstar -cg <cgname> enable -nop
```

9) Add BCV devices the SRDF/Star configuration

1) Add BCVs at Site C by associating the BCVs with the composite groups at both Site A and Site C.

Note that the `symbcv` command does not support the syntax `-rdf -rdrg name:<RDF group Name>`

```
Site A> symbcv -cg <cgname> -sid <SiteA sid> associateall dev -range  
<start:end> -rdf -rdrg <Async RDF group>
```

```
Site C> symbcv -cg <cgname> -sid <SiteC sid> associateall dev -range  
<start:end>
```

2) Synchronize the BCVs

Note that this command must be executed with the `-star` option. It can be executed from with Site A or Site C. Choose one or the other.

```
SiteA> symmir -cg <cgname> establish -star -full -rdf -rdrg name:<RDF  
group Name>
```

OR

```
SiteC> symmir -cg <cgname> establish -star -full -nop
```

[SURELY NOT FULL FOR ALL DEVICES]

6.6 Monitoring SRDF/Star

There is no SRDF/Star-specific monitoring. The strategic method in Citi for monitoring Symmetrix events, including all SRDF modes, is ECC. However, although ECC will indicate, for example, that one or more SRDF/A sessions have dropped on a given Symmetrix, it will not provide detail such as which sessions have dropped.

Once ECC alerts that there is an SRDF problem on an array, it is necessary to then use an element manager (just as SYMCLI) to investigate the exact problem further.

For SE6.2 onwards, an event daemon is available from EMC that can be used to monitor Symmetrix Operations. This can be used in Citi to provide troubleshooting information to supplement ECC alerts. The following events relating to SRDF can be monitored for (inexhaustive list):

- SRDF/A session
 - SRDF/A session inactive
 - SRDF/A session dropped; write pending limit reached. Host throttling is disabled
 - SRDF/A session dropped: write pending limit reached. Host throttling is enabled
 - SRDF/A session dropped: device not ready. Tolerence mode is off
 - SRDF/A session dropped: device not ready through consistency group
 - SRDF/A dropped: no SRDF links operational
 - SRDF/A dropped: timeout in MSC mode
- SRDF CG
 - RDF CG trip event triggered
- SRDF Link
 - No SRDF links in an RDF group are operational
 - Single RDF link in an RDF group is not operational

Note that the alerts include RDF CG trip events, as well as issues with underlying SRDF sessions, which could cause a Star trip.

Within Citi, SE6.4.2 is the first release in which the Event Daemon is supported. The SE 6.4 Solutions Guide provides details on how to configure the Event Daemon. Please configure the `daemon_options` file as per the SE 6.4 Solutions Guide to ensure logging of SRDF events (and indeed all events) to files `/var/symapi/log/symevent.log*`.

In addition, the legacy existing procedures for monitoring SRDF/A documented in Section 5.8.1 of the SRDF/A Solutions Guide can also be used if ECC is not available in your location.

6.7 Dealing with loss of all RDF daemons

If for some reason no functional `storrd` daemons are running at the production site to control SRDF/A cycle switching, then this must be fixed as a matter of urgency otherwise SRDF/A will eventually drop due to lack of buffer space at the production site.

Determine if daemons are running:

- If not, restart. Note, however, that `storrd` daemons should be setup to automatically restart (see Section 6.3 of the SRDF/A Solutions Guide)
- If they are running, stop and restart daemons in case they are hung.

```
SiteA> stordaeon shutdown -immediate storrd
```

If they will not shutdown, kill -9 the processes. Then restart the daemons.

```
SiteA> stordaeon start storrd
```

If daemons are now up but cycle switch is still not working, or daemons cannot be restarted (e.g. all Site A `storrd` servers down), you have to make a choice. Either:

A) Disable consistency so that the DMX can control SRDF/A cycle switching itself. This option would keep Site C 30 – 60 seconds behind SiteA, but “zero data loss” recovery at Site C would no longer be available (as Star is not running)

B Rely on Star to update Site C when a site switch is required. However, once Site A → Site C replication drops:

1. the site switch operation may now be a very lengthy process as Site C will be an increasingly long time behind Sites A/B. This may be unacceptable to the business.
2. if the Star switch (keeping data from Site C) fails, then data loss could be very great.

You should consider during implementation what your policy should be depending on the application needs. If no policy has been agreed, it recommended by the author that the default action be to disable star consistency to ensure Site C does not get increasingly far behind production.

6.7.1 How to disable Star consistency so that the DMX can control cycle switching.

```
SiteA> symstar -cg <CGname> disable -nop  
siteA> symcg -cg <CGname> disable -star -rdfg name:<SiteC> -nop
```

You will not longer be Star protected and “zero data loss” failovers to Site C will no longer be possible, but at least Site C will be being updated. You effectively have a concurrent SRDF setup at this time.

When healthy daemons are available again, then re-enable consistency protection.

```
SiteA> symcg -cg <CGname> enable -star -rdfg name:<SiteC> -nop
```

If necessary, connect and protect SiteC

```
SiteA> symstar -cg <CGname> connect -site <SiteC> -nop  
SiteA> symstar -cg <CGname> protect -site <SiteC> -nop
```

Finally, enable Star protection.

```
SiteA> symstar -cg <CGname> enable -nop
```

7 Appendix A – Set up of SRDF/Star with multiple triangles per SRDF/Star session

7.1 Example showing configuration of Star with two triangles

The following diagram and table define the configuration that is being set up in this example.




Figure 9 – Two triangles in a single SRDF/Star session

The Star session will have the following properties:

Site A location & Symmetrix ID	New York - Symm 0822
Site B location & Symmetrix ID	New Jersey - Symm 1135
Site C location & Symmetrix ID	London - Symm 1497
Composite Group Name	newcg
Synchronous Leg dynamic RDF groups	50 & 51
Synchronous Leg RDF group name	NewJersey
Asynchronous dynamic Leg RDF groups	52 & 53

Asynchronous Leg RDF group name	London
Recovery Leg dynamic RDF groups	54 & 55
BCV at Site A?	No
BCV at Site B?	No
BCV at Site C?	Yes
DSE Threshold (Site A ↔ Site B)	40%

Note that the RDF groups in this example do not conform to the naming convention in [Section 3.4.2](#), as this test was undertaken before the naming convention has been determined.

7.1.1 Step 1a – Verify Symmetrix Settings

For each symmetrix to be used, confirm required Symmetrix settings

```
SiteA> symcfg list -v -sid <SID>
```

- Concurrent RDF Configuration State is enabled
- Dynamic RDF Configuration State is enabled
- Concurrent Dynamic RDF Configuration is enabled
- RDF Data Mobility Configuration State is disabled

```
SiteA> symcfg list -rdfg <RDFG>
```

- Prevent Auto Link Recovery is enabled
- Prevent RAs Online is enabled

7.1.2 Step 1b - Create Concurrent SRDF environment

For a detailed description of the steps involved in setting up dynamic RDF groups and device pairs, see the SRDF/A Solutions Guide, Section 5.3.

1) Set up RDF group on Asynchronous leg from 0822 (Site A) to 1497 (Site C)

```
symrdf addgrp -label DYN50 -rdfg 50 -sid 0822 -dir 13C -link_limbo 60 -remote_rdfg 50 -
remote_sid 1497 -remote_dir 9C -rem_link_limbo 60 -nop
symrdf modifygrp -add -rdfg 50 -sid 0822 -dir 14C -remote_dir 10C -nop

symrdf addgrp -label DYN51 -rdfg 51 -sid 0822 -dir 13C -link_limbo 60 -remote_rdfg 51 -
remote_sid 1497 -remote_dir 9C -rem_link_limbo 60 -nop
symrdf modifygrp -add -rdfg 51 -sid 0822 -dir 14C -remote_dir 10C -nop
```

If using Transmit Idle (5772 or later), then link_limbo & rem_link_limbo be set to 10 (the default if neither is specified)

If DSE is to be used for this RDF group (5772 or later), then it need to be configured on both the source and target symmetrix. The following assumes the DSE pool has already been configured at both Site A and Site C. (see SRDF/A Solutions Guide Appendix G)

```
SiteA> symconfigure -f <command-file> -sid 0822 commit
```

```
SiteA> symconfigure -f <command-file> -sid 1497 commit
```

where command-file contains:

```
set rdf group 50, rdfa_dse_pool=DSE emulation=fba;
set rdf group 50, rdfa_dse_threshold=40;
set rdf group 50, rdfa_dse_autostart=ENABLE;
set rdf group 51, rdfa_dse_pool=DSE emulation=fba;
set rdf group 51, rdfa_dse_threshold=40;
set rdf group 51, rdfa_dse_autostart=ENABLE;
```

2) Confirm RDF group set up correctly

```
symcfg list -sid 0822 -switched -ra all
```

3) Set up RDF groups on Synchronous leg from 0822 (Site A) to 1135 (Site B)

```
symrdf addgrp -label DYN52 -rdfg 52 -sid 0822 -dir 7D -remote_rdfg 52 -remote_sid 1135 -
remote_dir 7D -nop
symrdf modifygrp -add -rdfg 52 -sid 0822 -dir 8D -remote_dir 8D -nop

symrdf addgrp -label DYN53 -rdfg 53 -sid 0822 -dir 7D -remote_rdfg 53 -remote_sid 1135 -
remote_dir 7D -nop
symrdf modifygrp -add -rdfg 53 -sid 0822 -dir 8D -remote_dir 8D -nop
```

4) Confirm RDF group set up correctly

```
symcfg list -sid 0822 -switched -ra all
```

5) Set up RDF group on recovery leg between 1497 (Site C) and 1135 (Site B)

```
symrdf addgrp -label DYN54 -rdfg 54 -sid 1497 -dir 7C -remote_rdfg 54 -link_limbo 60 -
remote_sid 1135 -remote_dir 7C -nop
symrdf modifygrp -add -rdfg 54 -sid 1497 -dir 7D -remote_dir 8C -rem_link_limbo 60 -nop

symrdf addgrp -label DYN55 -rdfg 55 -sid 1497 -dir 7C -remote_rdfg 55 -link_limbo 60 -
remote_sid 1135 -remote_dir 7C -nop
symrdf modifygrp -add -rdfg 55 -sid 1497 -dir 7D -remote_dir 8C -rem_link_limbo 60 -nop
```

If using Transmit Idle (5772 or later), then link_limbo & rem_link_limbo be set to 10 (the default if neither is specified)

6) Confirm RDF group set up correctly

```
symcfg list -sid 1135 -switched -ra all
symcfg list -sid 1497 -switched -ra all
```

7) Next create dynamic RDF device pairs

7a) Create the Async legs first, creating the device groups

```
symrdf -file ./pairsfile_0822_DYN50 -sid 0822 -rdfg 50 createpair -type rdf1 -establish
-nop
symrdf -file ./pairsfile_0822_DYN51 -sid 0822 -rdfg 51 createpair -type rdf1 -establish
-nop
```

7b) Now do the Sync legs (note no need to specify the device group)

```
symrdf -file ./pairsfile_0822_DYN50 -sid 0822 -rdfg 52 createpair -type rdf1 -establish
-nop
symrdf -file ./pairsfile_0822_DYN51 -sid 0822 -rdfg 53 createpair -type rdf1 -establish
-nop
```

8) Set modes to adaptive copy disk mode for all legs

```
symrdf -g DG50 set mode acp_disk -rdfg 50 -sid 0822 -nop
symrdf -g DG50 set mode acp_disk -rdfg 52 -sid 0822 -nop
symrdf -g DG51 set mode acp_disk -rdfg 51 -sid 0822 -nop
symrdf -g DG51 set mode acp_disk -rdfg 53 -sid 0822 -nop
```

7.1.3 Step 2 – Creating SRDF/Star Composite Groups

1) Create composite group

```
symcg create newcg -type RDF1 -rdf_consistency
```

2) Add concurrent SRDF devices to composite group

2a) Add concurrent device for Triangle 1

```
symcg -cg newcg -sid 0822 -rdfg 50 addall dev # Triange 1
```

2b) Add concurrent device for Triangle 2

```
symcg -cg newcg -sid 0822 -rdfg 51 addall dev # Triange 2
```

3) Now define RDF group names and recovery leg RDF groups

Triangle 1

```
symcg -cg newcg -rdfg 0822:50 set -name London
symcg -cg newcg -rdfg 0822:52 set -name NewJersey
symcg -cg newcg set -rdfg 0822:50 -recovery_rdfg 54
symcg -cg newcg set -rdfg 0822:52 -recovery_rdfg 54
```

Triangle 2

```
symcg -cg newcg -rdfg 0822:51 set -name London
symcg -cg newcg -rdfg 0822:53 set -name NewJersey
symcg -cg newcg set -rdfg 0822:51 -recovery_rdfg 55
symcg -cg newcg set -rdfg 0822:53 -recovery_rdfg 55
```

7.1.4 Step 3 – Creating the SRDF/Star options file

Create the SRDF/Star options of the following form:

```
SiteA> cat newcg_opt
```

```
SYMCLI_STAR_WORKLOAD_SITE_NAME=NewYork
SYMCLI_STAR_SYNC_TARGET_SITE_NAME=NewJersey
SYMCLI_STAR_ASYNC_TARGET_SITE_NAME=London
SYMCLI_STAR_USE_ADAPTIVE_COPY=both
SYMCLI_STAR_ADAPTIVE_COPY_TRACKS=30000
SYMCLI_STAR_ACTION_TIMEOUT=1800
```

The options file needs to be stored in a standard location as it will be needed again to:

1. dismantle a Star configuration to add/remove disk
2. rebuild a Star configuration

A good location would be wherever you currently store SRDF/A pairs files used for dynamic RDF configuration.

The same options file may be used for many SRDF/Star configurations, if the same parameters are to be used.

7.1.5 Step 4 – Performing the symstar setup operation

Setup Star using this options file

```
SiteA> symstar -cg newcg -options newcg_opt setup
```

Now sftp the Star configuration definition

```
/usr/emc/API/symapi/config/STAR/def/newcg
```

file for this Star session to an identical location on all other Star controller servers at both the local site, and at the synchronous (if in use) and asynchronous target sites.

7.1.6 Step 5 – Creating R2 Composite groups

These commands are issues at the SRDF target sites.

```
SiteC> symstar -cg flexcg buildcg -site London -nop
```

If Star control servers are installed at Site B:

```
SiteB> symstar -cg flexcg buildcg -site NewJersey -nop
```

7.1.7 Step 6 - Adding BCV devices to the SRDF/Star configurations

1) Add BCVs at Site C by associating the BCVs with the composite groups at both Site A and Site C

Note that the symbcv command does not support the syntax `-rdf -rdfg name:<RDF group Name>`

```
Site A> symbcv -cg newg -sid 0822 associateall dev -range <02C1:02C5> -rdf -rdfg 50  
Site A> symbcv -cg newg -sid 0822 associateall dev -range <02C9:02CD> -rdf -rdfg 51
```

```
Site C> symbcv -cg newcg -sid 1497 associateall dev -range <02C1:02C5>  
Site C> symbcv -cg newcg -sid 1497 associateall dev -range <02C9:02CD>
```

2) Synchronize the BCVs

Before attempting the full establish, ensure that Clone emulation is enabled if not using RAID-5 BCVs. Clone emulation is used if the clone emulation environment variable is set to ENABLED.

```
Site A> SYMCLI_CLONE_EMULATION=ENABLED
Site A> export SYMCLI_CLONE_EMULATION
```

One can confirm if this (and indeed other) SYMCLI environment variables have been set with the following command:

```
Site A> symcli -def

Symmetrix Command Line Interface (SYMCLI) Version V6.3 (Edit Level: 623)
built with SYMAPI Version V6.3 (Edit Level: 623)

Current settings of the SYMCLI environmental variables:

SYMCLI_CLONE_EMULATION: ENABLED
```

Note that the establish command must be executed with the `-star` option

```
SiteA> symmir -cg flexcg establish -star -full -rdf -opt_rag -rdfg name:London
```

7.1.8 Step 7 – Bringing up the SRDF/Star sites for normal operation

```
SiteA> symstar -cg newcg connect -site NewJersey -nop
SiteA> symstar -cg newcg connect -site London -nop

SiteA> symstar -cg newcg protect -site NewJersey -nop
SiteA> symstar -cg newcg protect -site London -nop

SiteA> symstar -cg newcg enable -nop
```

7.2 Example showing removal of a triangle from a Star session

The following example demonstrates how a triangle can be removed from a multi-triangle star setup so that it can be manipulated outside of Star. In brief the process involves unconfiguring star, removing the RDF group from the Composite Group and then re-defining star setup using the new composite group.

In this example, we remove the triangle 1 setup in the previous example.

1) Deconfigure star

```
SiteA> symstar -cg newcg disable
SiteA> symstar -cg newcg unprotect -site NewJersey
SiteA> symstar -cg newcg unprotect -site London
SiteA> symstar -cg newcg -options newcg_opt -clear setup
```

2) Remove RDF group concurrent devices from composite groups

```
SiteA> symcg -cg newcg rmall -rdfg 50
SiteB> symcg -cg newcg rmall -rdfg 50
```

```
SiteC> symcg -cg newcg rmall -rdfg 52
```

3) Remove BCVs from Composite Groups

```
Site A> symbcv -cg newcg -sid 0822 disassociate dev 02C1 -rdf  
Site A> symbcv -cg newcg -sid 0822 disassociate dev 02C5 -rdf
```

```
Site C> symbcv -cg newcg -sid 0822 disassociate dev 02C1  
Site C> symbcv -cg newcg -sid 0822 disassociate dev 02C5
```

4) Remove RDF group name

RDF group names are reset by setting to the null string

Recovery RDF groups are reset by setting to zero.

```
SiteA> symcg -cg newcg -rdfg 0822:50 set  
SiteA> symcg -cg newcg -rdfg 0822:52 set  
SiteA> symcg -cg newcg set -rdfg 0822:50 -recovery_rdfg 0  
SiteA> symcg -cg newcg set -rdfg 0822:52 -recovery_rdfg 0
```

5) Set up Star for the new reduced composite group

Follow instruction from Step 4 in previous section for setting up Star.

8 Appendix B – Listing summary status of Star sessions

The following perl script gives an example of how a summary report of Star status can be generated for composite groups visible on a Unix SAN management server.

```
#!/export/opt/perl/5.8.0/bin/perl

#####
#
# script starlist
#
# Invocation: starlist
# Return codes:
#           0      Success
#           1      Error running symcg list command
#           2      Error running symstar query command
#
#####
#
# Overview:
#
# 1) Get a list of composite groups
# 2) Check Star status for each composite group
#
#####
#
# Variables used:
#
# $cg      Composite Group
#
#####
#
# Get list of composite groups visible on this server
#
open(CGLIST,"/usr/symcli/bin/symcg list |") || exit (1);
while(<CGLIST>) {
    if (/RDF/) {
        @cg_list_output = split(" ",$_);
        $cg=$cg_list_output[0];

#
# Determine if Star defined for each CG, and if so star status
#

        open(STARQUERY,"/usr/symcli/bin/symstar -cg $cg query |") || exit (2);

        print "-----\n";
        print "Star status for composite group $cg\n";
        while(<STARQUERY>) {
            next if /STAR Consistency/;
            next if /STAR Information/;
            if (/STAR/ || /Target_Site/) {
                @query_output = split(" ",$_);
                print "  $cg: $query_output[0] $query_output[2]\n";
            }
        }
    }
}
exit (0);
```

Example output from the above script is shown below.

```
Star status for composite group flexcg
  flexcg: 1st_Target_Site Protected
  flexcg: 2nd_Target_Site Disconnected
-----
Star status for composite group forty

STAR definition for Cgname: forty was not found
-----
Star status for composite group newcg
  newcg: 1st_Target_Site Protected
  newcg: 2nd_Target_Site Protected
  newcg: STAR Protected
-----
Star status for composite group test1
  test1: 1st_Target_Site Protected
  test1: 2nd_Target_Site Protected
  test1: STAR Protected
-----
Star status for composite group test2
  test2: 1st_Target_Site Protected
  test2: 2nd_Target_Site Protected
  test2: STAR Protected
```

In this artificial example:

1. Star is not enabled for composite group flexcg as the 2nd Target Site is disconnected
2. There is no Star session defined for composite group forty.
3. Composite Groups newcg, test1 and test2 are fully protected by Star; both Target sites are marked as Protected, and the status is "STAR protected"

9 Appendix C – Output from symstar show –detail

Example of symstat show –detail output for a Star session with two triangles, set up as per configuration detailed in Appendix A.

```
#symstar -cg newcg show -detail

Composite Group Name      : newcg

Synchronous Target Site Information:
{
-----
Source (R1) View          Target (R2) View
-----
Symmetrix      RD          Symmetrix      RD
ID             F           ID             F           Rcvy
               G   Dev     ID             G   Dev     RDFG
-----
000187700822 53  0209 000187721135 53  0209 55
               020D          020D
000187700822 52  0201 000187721135 52  0201 54
               0205          0205
-----
}

Asynchronous Target Site Information:
{
-----
Source (R1) View          Target (R2) View
-----
Symmetrix      RD          Symmetrix      RD
ID             F           ID             F           Rcvy
               G   Dev     ID             G   Dev     RDFG
-----
000187700822 51  0209 000187721497 51  0209 55
               020D          020D
000187700822 50  0201 000187721497 50  0201 54
               0205          0205
-----
}

Option file settings:
{
WorkloadSite      : NewYork
SyncTargetSite    : NewJersey
AsyncTargetSite    : London
Adaptive_Copy_Tracks : 30000
Action_Timeout    : 1800
Term_Sddf         : NO
}
```

10 Appendix D – Output from symstar query

An example of symstar query output for a Star-enabled Composite Group. This is a very useful command for checking the status of a Star session. In particular, the author often uses this command to check

- Current workload site
- Site nicknames
- Checking Star system state
- Last symstar operation performed, with time and status of that operation
- If Async target site data is most current
- RDF modes
- RDF pair states

Each of the above are marked in bold in the output below:

```
# symstar -cg newcg query
```

```
Site Name                               : NewYork

Workload Site                          : NewYork
1st Target Site                       : NewJersey
2nd Target Site                       : London

Composite Group Name                    : newcg
Composite Group Type                    : RDF1

Workload Data Image Consistent          : Yes
System State:
{
  1st Target Site                      : Protected
  2nd Target Site                      : Protected
  STAR                                : Protected
}

Last Action Performed                 : Enable
Last Action Status                   : Successful
Last Action Timestamp               : 10/24/2006_09:50:22
```

```
STAR Information:
{
  STAR Consistency Capable              : Yes
  STAR Consistency Mode                 : STAR
  Synchronous Target Site               : NewJersey
  Asynchronous Target Site              : London
  Differential Resync Available          : Yes
  R2 Recoverable                        : Yes
  Asynchronous Target Site Data most Current : No
}
```

```
1st Target Site Information:
{
  Site Name                             : NewJersey
  RDF Consistency Capability             : SYNC
  RDF Consistency Mode                  : SYNC
  Site Data Image Consistent            : Yes
}
```

Workload Site					Target Site				
-----					-----				
ST	LI	ST	M						
RD A	N Rem	RD A	O						
Symm F T R1 Inv R2 Inv K Symm F T R1 Inv R2 Inv D RDF Pair									
ID G E Tracks Tracks S ID G E Tracks Tracks E STATE									

```

00822 53 RW      0      0 RW 21135 53 WD      0      0 S Synchronized
00822 52 RW      0      0 RW 21135 52 WD      0      0 S Synchronized

Totals: -- ----- -- -- ----- - -----
          RW      0      0 RW      WD      0      0 S Synchronized
}

```

2nd Target Site Information:

```

{
Site Name                      : London
RDF Consistency Capability     : MSC
RDF Consistency Mode          : MSC
Site Data Image Consistent     : Yes

```

Workload Site										Target Site									
ST					LI					ST					M				
RD A					N Rem					RD A					O				
Symm	F	T	R1 Inv	R2 Inv	K	Symm	F	T	R1 Inv	R2 Inv	D	RDF	Pair						
ID	G	E	Tracks	Tracks	S	ID	G	E	Tracks	Tracks	E	STATE							
00822	51	RW	0	0	RW	21497	51	WD	0	0	A Consistent								
00822	50	RW	0	0	RW	21497	50	WD	0	0	A Consistent								
Totals: -- ----- -- -- ----- - -----																			
RW 0 0 RW WD 0 0 A Consistent																			

Legend:

Modes:

Mode of Operation: A=Async, C=Adaptive Copy, S=Sync, O=Other, M=Mixed