EMC® Symmetrix® Remote Data Facility (SRDF®)

Product Guide

November, 2014

Rev 05

EMC²

Copyright © 2001 - 2014 EMC Corporation. All rights reserved. Published in the USA.

Published November, 2014

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to the technical documentation and advisories section on the EMC online support website.

CONTENTS

Preface

Chapter 1 SRDF Family of Products

Introduction 18
Specifics of SRDF for VMAX 10K or VMAXe solutions 19
SRDF solutions 20
Automation capabilities 23
Management tools 24
SRDF Host Component for z/OS 24
Solutions Enabler SRDF component 25
Unisphere for VMAX 25
EMC z/OS Storage Manager 25

Chapter 2 Symmetrix Enginuity and SRDF Concepts

Symmetrix logical devices, RAID groups, and mirrors 28
Symmetrix session 30
Virtual Provisioning and thin devices 30
SRDF devices 31
Primary (R1) and secondary (R2) devices 32
Primary devices (R1, R11) 34
Dual-role (cascaded) SRDF devices (R21) 35
Diskless R21 devices (DL R21) 36
Secondary devices (R2, R22) 38
Thin SRDF devices 39
SRDF thick-to-thin support 40
Open systems metadevices as SRDF devices 41
Mainframe RAID 10 SRDF devices 42
EMC Compatible Peer 42
Dynamic SRDF devices 43
SRDF device states 46
Host interface view 46
SRDF view 48
SRDF groups 50
Group types 52
Moving dynamic SRDF devices between SRDF groups 52
System-level SRDF device and group support 56
SRDF links 57
SRDF network protocols and topologies 60
Fibre Channel 60
Gigabit Ethernet (GigE) 62
ESCON 64
SRDF modes of operation 65
Synchronous 65
Asynchronous 65
Semi-synchronous 65
Adaptive copy 66

Chapter 3 SRDF/Synchronous Operations

Enginuity emulations and I/O operations 68
Write operations 69
Remote mirroring 69
Write I/Os to R21 devices 71
Write I/Os and compression 73
Read operations 74
Read I/O if R1 local mirror fails 75
Read I/Os to R2 devices 76
Read I/Os to DL R21 devices 76
Read I/Os to R21 devices 76
Dependent-write consistency and SRDF/CG 77
Dependent-write operations 77
Ensuring data consistency 78
Recovery operations 80
Planned failover to the secondary system 80
Failback to the primary system 82
Recovery for a large number of invalid tracks 83

Chapter 4 SRDF/Asynchronous Operations

Overview 86
SRDF/A benefits 86
Single session mode 86
Multi Session Consistency (MSC) mode 86
Reserve Capacity 87
Write folding 87
Write pacing 88
Tolerance mode 88
Development and availability 89
Single session mode 91
Dependent-write consistency 91
Single session states 95
Single session state transitions 96
Switching to SRDF/A mode 97
Switching from SRDF/A to SRDF/S mode 98
Consistency exempt attribute 99
Recovery scenarios 100
Temporary all links lost 100
Permanent all links lost 100
Session cleanup process 100
Multi Session Consistency mode 103
Entering SRDF/A multi session consistency 103
MSC mode dependent-write consistency 104
Performing an SRDF/A MSC consistent cycle switch 105
MSC mode delta set switching 106
MSC session cleanup process 109
SRDF/A and cache utilization 111
Tunable cache utilization 111
Reserve Capacity 111
Write pacing 111

Chapter 5 Adaptive Copy Operations

Adaptive copy modes 116

Adaptive copy write pending 116
Adaptive copy disk 116
SRDF/Data Mobility 117
Benefits 117
Limitations 118
SRDF/Automated Replication 118
Benefits 120
Limitations 120

Chapter 6 SRDF Multisite Solutions

Concurrent SRDF 122
Setting up a concurrent SRDF relationship 123
Concurrent SRDF/S with SRDF/A 123
Concurrent SRDF/A 123
Migration with concurrent SRDF 124
Concurrent SRDF with independent consistency protection 124
Benefits 125
Requirements 125
Cascaded SRDF 126
Setting up a cascaded SRDF relationship 126
Benefits 127
Requirements 127
Extended Distance Protection 128
Setting up an SRDF/EDP relationship 128
Benefits 129
Requirements 129
SRDF/Star 130
How SRDF/Star works 131
Concurrent SRDF/Star 132
Cascaded SRDF/Star 134
SRDF/Star with SRDF/EDP 136
SRDF/Star benefits 137
SRDF/Star requirements 137
Four-site SRDF solution for open systems host environment 139
Benefits 140
Requirements 140
SRDF/SQAR 141
Requirements 141

Chapter 7 SRDF Interfamily Connectivity

Overview 144
SRDF two-site interfamily connectivity 144
SRDF multi-site interfamily connectivity 146

Chapter 8 SRDF Migration Operations

Overview 150
Full SRDF functionality support and migration only SRDF support 150
Migrating data with concurrent SRDF 153
Replacing R2 devices with new R2 devices 153
Replacing R1 devices with new R1 devices 154
Migrating data with concurrent and cascaded SRDF 156
Replacing R1 and R2 devices with new R1 and R2 devices 156

Chapter 9 SRDF Integration

SRDF and TimeFinder 160
R1 and R2 devices in TimeFinder operations 161
TimeFinder and SRDF adaptive copy mode operations 162
TimeFinder and SRDF/A operations 164
TimeFinder and SRDF/S operations 165
TimeFinder and SRDF in virtual provisioned environments 167
SRDF and open systems clusters 169
SRDF and mainframe automation software 171
SRDF and EMC AutoSwap 171
SRDF and EMC GDDR 172
SRDF and open systems automation software 174
SRDF and VMware environments 175
SRDF and EMC FAST VP 176
Requirements 177
SRDF and EMC RecoverPoint 178
Requirements 178
Limitations 178

FIGURES

	Title	Page
1	Typical SRDF solutions in open systems and mainframe host environments	18
2	Two logical devices, each with a single mirror (RAID 5 and RAID 1)	29
3	Write I/O to thin devices	31
4	SRDF active remote mirroring	32
5	R1 and R2 devices, mirror configuration (Symmetrix view)	33
6	R11 device with two R1 SRDF mirrors, paired with two R2 devices	34
7	R21 device used in cascaded SRDF operations	35
8	DL R21 devices used in SRDF/EDP configuration	36
9	R22 in cascaded and concurrent SRDF/Star solutions	38
10	Striped and concatenated SRDF metadevices	41
11	Host interface and SRDF view	46
12	Configuring SRDF groups	51
13	SRDF groups with devices that mirror between Symmetrix A and Symmetrix B	51
14	Moving SRDF devices between SRDF groups (full move)	53
15	SRDF groups and devices with two SRDF mirrors	54
16	Unidirectional links	57
17	Bidirectional links	58
18	Dual-directional links	58
19	Switched SRDF over Fibre Channel	61
20	Switched and concurrent SRDF over Fibre Channel	62
21	Switched SRDF over GigE	64
22	Enginuity emulations processing I/Os in a simple SRDF environment	68
23	Write I/O flow in a simple SRDF configuration	69
24	Write to R2 if R1 local mirror fails	70
25	Write I/Os to R21 devices	71
26	Three-site configuration with R21 devices	72
27	Read I/O flow from the primary device	74
28	Read I/O if R1 local mirror fails	75
29	Primary and secondary systems	78
30	Failed links between Primary 2 and Secondary 1	78
31	Primary 1, 2, and 3 in a consistency group	79
32	Failed link between Primary 2 and Secondary 2	79
33	Planned failover operation in a two-site configuration (SRDF/S)	81
34	Failover to Symmetrix B, Symmetrix A and production host unavailable	82
35	SRDF/A delta sets and their relationships	92
36	Delta set switching	93
37	SRDF/A single session states	95
38	SRDF/A MSC allowed state transitions	103
39	SRDF/A MSC delta sets and their relationships	105
40	SRDF/A MSC capture delta set collects application write I/O	107
41	SRDF/A MSC delta set switching process	109
42	SRDF/A R2 devices as source devices in TimeFinder operations	114
43	SRDF/DM adaptive copy mode	117
44	SRDF/AR single-hop data flow	118
45	SRDF/AR multi-hop data flow	119
46	Concurrent SRDF topology	122
47	Concurrent SRDF/S with independent consistency protection	124
48	Cascaded SRDF topology	126
49	SRDF/EDP basic topology	128
50	Concurrent SRDF/Star (left) and cascaded SRDF/Star (right)	130

51 Concurrent SRDF/Star solution 132
52 Concurrent R22 SRDF/Star solution 133
53 Cascaded SRDF/Star solution 134
54 Cascaded R22 SRDF/Star solution 135
55 SRDF/EDP in the cascaded SRDF/Star environments with an R2 site 136
56 Four-site SRDF solution 139
57 SRDF/SQAR with AutoSwap environment 141
58 Migrating data and removing the original secondary system (R2) 153
59 Migrating data and replacing the original primary system (R1) 154
60 Migrating data and replacing the original primary (R1) and secondary (R2) systems .. 156
61 TimeFinder/Clone 160
62 SRDF devices in TimeFinder operations 162
63 SRDF/AR single-hop solution 163
64 SRDF/AR multi-hop solution 163
65 Simultaneous TimeFinder/Clone solution prior to Enginuity version 5875.135.91 165
66 Simultaneous TimeFinder/Clone with Enginuity version 5875.135.91 166
67 SRDF/CE two-node, two-cluster configuration 169
68 AutoSwap processing 171
69 Three-site concurrent SRDF/Star with GDDR 172
70 SRDF AutoStart typical configuration 174
71 EMC SRDF and VMware SRM 175

TABLES

	Title	Page
1	Revision History	13
2	Enginuity version required on Symmetrix systems	18
3	SRDF solutions	20
4	SRDF automation software	23
5	RAID schemes used in Symmetrix storage systems	28
6	Device types supported by Virtual Provisioning	30
7	Support for Virtual Provisioning in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX	39
8	Support for Virtual Provisioning in SRDF for VMAX 10K or VMAXe	40
9	Enginuity support for PPRC	42
10	Enginuity support for dynamic SRDF devices	44
11	Primary (R1) device accessibility	49
12	Secondary (R2) device accessibility	49
13	Scale of SRDF support in Symmetrix VMAX 40K, VMAX 20K/VMAX, and DMX systems	56
14	Scale of SRDF support in Symmetrix VMAX 10K and VMAXe systems	56
15	SRDF hardware and software compression support	60
16	Summary of SRDF/A development and availability for Symmetrix VMAX 40K, VMAX 20K/VMAX, and DMX series	89
17	Delta set switching process	93
18	SRDF modes allowed for cascaded SRDF	127
19	SRDF modes allowed for SRDF/EDP	129
20	SRDF two-site interfamily connectivity (VMAX 40K, VMAX 20K/VMAX, DMX)	144
21	SRDF two-site interfamily connectivity for solutions including VMAX 10K or VMAXe	145
22	SRDF multi-site interfamily connectivity	146
23	Limitations of migration only SRDF support	150
24	SRDF/A device-level pacing requirements for TimeFinder operations	164
25	Thin SRDF devices with TimeFinder/Clone or VP Snap	167
26	Thin SRDF devices and TimeFinder/Snap	168

PREFACE

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your EMC representative if a product does not function properly or does not function as described in this document.

Note: This document was accurate at publication time. New versions of this document might be released on the EMC Online Support site. Check the EMC Online Support site to ensure that you are using the latest version of this document.

Purpose

This document introduces SRDF product family, concepts, and theory of operation for the Symmetrix platforms including Symmetrix VMAX 40K, VMAX 20K/VMAX, VMAX 10K, VMAXe, and DMX.

Audience

This document is intended for audiences who require basic knowledge of SRDF.

Related documentation

The following EMC publications provide additional information:

- *EMC Symmetrix VMAX Family Documentation Set* — Contains the hardware platform product guide, TimeFinder product guide for the Symmetrix VMAX Family (10K, 20K, and 40K).
- *EMC Symmetrix System Viewer for Desktop and iPad®* — Illustrates VMAX 10K, VMAX 2Ok, and VMAX 40K system hardware, incrementally scalable system configurations, and available host connectivity offered for Symmetrix systems.
- *EMC SRDF Connectivity Guide*
- *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide*
- *EMC SRDF Host Component for z/OS Product Guide*
- *EMC GDDR for SRDF/SQAR with AutoSwap Product Guide*

Conventions used in this document

EMC uses the following conventions for special notices:

Note: A note presents information that is important, but not hazard-related.

Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities • URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications
Bold	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages Used in procedures for: • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • What the user specifically selects, clicks, presses, or types
Italic	Used in all text (including procedures) for: • Full titles of publications referenced in text • Emphasis, for example, a new term • Variables
`Courier`	Used for: • System output, such as an error message or script • URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
`Courier bold`	Used for specific user input, such as commands
`Courier italic`	Used in procedures for: • Variables on the command line • User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections — the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Revision History

The following table presents the revision history of this document:

Table 1 Revision History

Revision	Description and/or Change	Enginuity Operating Environment
REV 01	• This version combines the previous versions that were split into the following two books: *EMC Symmetrix Remote Data Facility (SRDF) for VMAX 40K, VMAX 20K/VMAX, DMX Series Product Guide A15* *EMC Symmetrix Remote Data Facility (SRDF) for VMAX 10K or VMAXe Series Product Guide A04* • This version refines the information about SRDF interfamily connectivity. "SRDF Interfamily Connectivity" on page 143 provides details. • This version refines the information about SRDF migration solutions. "SRDF Migration Operations" on page 149 provides details.	5876.159.102
REV 02	This version includes the addition of information about the following features: • SRDF/SQAR: "SRDF/SQAR" on page 141 provides details. • Co-existence of SRDF and RecoverPoint CDP: "SRDF and EMC RecoverPoint" on page 178 provides details. • Coordination of FAST VP data movements in multi-site SRDF solutions: "SRDF and EMC FAST VP" on page 176 provides details. This version includes the following updated information: • Conditions under which a dynamic R1/R2 personality swap is not supported: "R1/R2 personality swap" on page 44 provides details. • Requirements for SRDF/Star thick-to-thin support: "SRDF multi-site interfamily connectivity" on page 146 provides details.	5876.229.145
REV 03	Clarification of RAID 10 support in Table 5 on page 28.	
REV 04	Correct typographical error on page 141.	
REV 05	Correct typographical error on page 135.	

Where to get help

EMC support, product, and licensing information can be obtained on the EMC Online Support site as described next.

Note: To open a service request through the EMC Online Support site, you must have a valid support agreement. Contact EMC support personnel for details about obtaining a valid support agreement or to answer any questions about your account.

Product information

For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to the EMC Online Support site (registration required) at:

https://support.EMC.com

Technical support

EMC offers a variety of support options.

Support by Product — EMC offers consolidated, product-specific information on the Web at:

https://support.EMC.com/products

The Support by Product web pages offer quick links to Documentation, White Papers, Advisories (such as frequently used Knowledgebase articles), and Downloads, as well as more dynamic content, such as presentations, discussion, relevant Customer Support Forum entries, and a link to EMC Live Chat.

EMC Live Chat — Open a Chat or instant message session with an EMC Support Engineer.

Symmetrix Customer Procedure Generator — For customized documentation for managing EMC Symmetrix storage systems, go to the EMC Online Support site, and download this stand-alone utility.

The Symmetrix Procedure Generator is updated regularly by EMC and contains a 30-day expiration timer to ensure you always have the latest validated procedures.

Symmetrix System Viewer — Illustrates the system hardware, incrementally scalable system configurations, and available host connectivity offered for Symmetrix VMAX systems, go to the EMC Online Support site, and download this stand-alone utility.

eLicensing support

To activate your entitlements and obtain your Symmetrix license files, visit the Service Center on https://support.EMC.com, as directed on your License Authorization Code (LAC) letter e-mailed to you.

For help with missing or incorrect entitlements after activation (that is, expected functionality remains unavailable because it is not licensed), contact your EMC Account Representative or Authorized Reseller.

For help with any errors applying license files through Solutions Enabler, contact the EMC Customer Support Center.

If you are missing a LAC letter, or require further instructions on activating your licenses through the Online Support site, contact EMC's worldwide Licensing team at licensing@emc.com or call:

- North America, Latin America, APJK, Australia, New Zealand: SVC4EMC (800-782-4362) and follow the voice prompts.
- EMEA: +353 (0) 21 4879862 and follow the voice prompts.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to:

`techpubcomments@emc.com`

CHAPTER 1
SRDF Family of Products

This chapter introduces EMC Symmetrix Remote Data Facility (SRDF) solutions. Topics include:

- Introduction 18
- SRDF solutions 20
- Automation capabilities 23
- Management tools 24

Introduction

The EMC® Symmetrix® Remote Data Facility (SRDF®) family of products offers a range of Symmetrix-based disaster recovery, parallel processing, and data migration solutions for Symmetrix VMAX® Family systems, VMAXe® systems, and DMX™ systems. The VMAX Family includes Symmetrix VMAX® 40K, VMAX® 20K/VMAX®, and VMAX® 10K.

The SRDF solutions require that the Enginuity™ Operating Environment for Symmetrix runs on every Symmetrix system that participates in an SRDF solution. Different Enginuity versions offer different SRDF features. Table 2 on page 18 lists the Enginuity versions required on the Symmetrix systems.

Table 2 Enginuity version required on Symmetrix systems

Symmetrix system	Enginuity version
VMAX 40K	5876.82.57 and higher
VMAX 20K	5876.82.57 and higher
VMAX	5874 and higher
VMAX 10K	5876.82.57 and higher
VMAXe	5875.231.172 and higher
DMX	5773 and lower

SRDF solutions require at least two Symmetrix systems. These systems are also known as the primary and the secondary system. Both sites can be located in the same room, in different buildings within the same campus, or hundreds to thousands of kilometers apart.

Figure 1 on page 18 shows a typical SRDF solution in an open systems host environment (left) and in a mainframe host environment (right).



Figure 1 Typical SRDF solutions in open systems and mainframe host environments

In the open systems host environment, the production host connects to Symmetrix A. Primary devices (R1) in Symmetrix A contain production data. The production data is remotely mirrored to the secondary devices (R2) in Symmetrix B while the production host is in operation, issuing I/O to primary devices in Symmetrix A. Symmetrix A and Symmetrix B are connected to each other through the SRDF links.

In the mainframe host environment, the production host and the remote host are connected to both Symmetrix A and Symmetrix B. By configuring active, recovery, and alternate host paths, the mainframe host can leverage the operating system capabilities to use the SRDF configuration not only as a disaster recovery mechanism, but also as a parallel processing solution.

SRDF disaster recovery solutions are based on active remote mirroring and dependent-write consistent copies of data maintained at one or more remote locations. A dependent-write is a write operation that cannot be issued by an application until a prior, related write I/O operation completes. Dependent-write consistency is required to ensure transactional consistency when the applications are restarted at the remote location.

SRDF products address different Recovery Point Objectives (RPOs) within the required Recovery Time Objective (RTO) where:

- RTO denotes the time allowed for recovery to a specified point of consistency
- RPO denotes the point of consistency to which applications need to recover.

SRDF solutions can operate in the following modes of operation:

- Synchronous mode
- Semi-synchronous mode
- Asynchronous mode
- Adaptive copy mode

These modes of operation address different service level requirements and determine:

- How R1 devices are remotely mirrored to R2 devices across the SRDF links
- How I/Os are processed
- When the acknowledgement is returned to the production host that issued a write I/O command.

"SRDF modes of operation" on page 65 provides additional details about each mode.

Specifics of SRDF for VMAX 10K or VMAXe solutions

SRDF for VMAX 10K or VMAXe solutions are supported with Enginuity versions 5875.231.172 and higher. Two-site SRDF for VMAX 10K or VMAXe solutions are available with Enginuity 5875.231.172 or higher, while three-site SRDF for VMAX 10K or VMAXe solutions are supported with Enginuity 5876.82.57 or higher.

SRDF for VMAX 10K or VMAXe solutions do not support replicating mainframe data between Symmetrix arrays. Other features such as Domino and support for hardware compression, which are available in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX solutions, are not supported in SRDF for VMAX 10K or VMAXe solutions. Notes are used in this document to highlight the features that do not apply to SRDF for VMAX 10K or VMAXe solutions.

SRDF solutions

SRDF solutions offer a range of topologies and modes of operations to meet different service level requirements. Table 3 on page 20 summarizes SRDF solutions.

Table 3 SRDF solutions (page 1 of 3)

Solution highlights	Description	Site topology
SRDF/S • Synchronous mode • No data exposure • SRDF/Consistency Group (SRDF/CG) is offered • Limited distance (up to 125 miles (200 km))	Maintains a real-time (synchronous) mirrored copy of production data (R1 devices) at a physically separated Symmetrix system (R2 devices). Chapter 3, "SRDF/Synchronous Operations," provides additional details.	
SRDF/A • Asynchronous mode • RPO seconds before the point of failure • Unlimited distance	Mirrors data from the R1 devices while maintaining a dependent-write consistent copy of the data on the R2 devices at all times. The copy of the data at the secondary site is typically only seconds behind the primary site. Chapter 4, "SRDF/Asynchronous Operations," provides additional details.	
SRDF/Data Mobility (SRDF/DM) • Adaptive copy mode • Designed for migration • Unlimited distance	Enables fast data transfer from R1 to R2 devices over extended distances. Chapter 5, "Adaptive Copy Operations,"provides additional details.	
SRDF/Automated Replication (SRDF/AR) • Combines SRDF/DM (adaptive copy mode) and SRDF/S (synchronous mode) • Unlimited distance • Disaster restart solution • Requires EMC TimeFinder®	Optimizes bandwidth requirements and provides data protection with dependent-write consistency across long distances. SRDF/AR requires TimeFinder on the primary system. Chapter 5, "Adaptive Copy Operations,"and Chapter 9, "SRDF Integration,"provide additional details.	

Table 3 SRDF solutions (page 2 of 3)

Solution highlights	Description	Site topology
Concurrent SRDF • Combines SRDF/S, SRDF/A, or adaptive copy • Three-site disaster recovery solution using R11 devices • Provides support for device migration	Provides advanced multisite business continuity protection. It supports concurrent primary devices (R11) that can be mirrored concurrently to two R2 devices in one or two remote Symmetrix systems. "Primary devices (R1, R11)" on page 34 and "Concurrent SRDF" on page 122 provide additional details.	
Cascaded SRDF • Combines SRDF/S and SRDF/A • Three-site disaster recovery solution using R21 devices	Provides advanced multisite business continuity protection. Data on the primary site is synchronously mirrored to a secondary (R21) site, and then asynchronously mirrored from the secondary (R21) site to a tertiary (R2) site. "Dual-role (cascaded) SRDF devices (R21)" on page 35 and "Cascaded SRDF" on page 126 provide additional details.	
SRDF/Star • SRDF/CG with SRDF/S and SRDF/A • Three-site disaster recovery solution providing a zero data loss recovery • Cascaded SRDF/Star (top) • Concurrent SRDF/Star (bottom)	Provides advanced multisite business continuity protection and a disaster-restart solution. It offers the ability to differentially establish and protect data amongst surviving sites in a multisite disaster recovery implementation. "SRDF/Star" on page 130 provides additional details.	

Table 3 SRDF solutions (page 3 of 3)

Solution highlights	Description	Site topology
SRDF/SQAR • Four-site implementation of SRDF/S and SRDF/A • For mainframe environments only • EMC GDDR and AutoSwap™ are required	SRDF/SQAR enables differential resynchronization between sites along the perimeter of a 'square' multi-site SRDF topology. "SRDF/SQAR" on page 141 provides additional details.	

For additional software options in your region, contact EMC support personnel.

Automation capabilities

SRDF solutions can be integrated with automation software listed in Table 4 on page 23.

Table 4 SRDF automation software (page 1 of 2)

Description	SRDF site configuration
SRDF/CE • Complete solution for restarting operations in cluster environments such as MSCS with Microsoft Failover Clusters • Can be implemented with SRDF/S and SRDF/A • Additional details in "SRDF and open systems clusters" on page 169	
GDDR (SRDF for VMAX 40K, VMAX 20K, DMX solutions only) • Complete automation solution for restarting mainframe operations • Can be implemented with SRDF/S, SRDF/A, EMC AutoSwap™, TimeFinder and other EMC products • Additional details in "SRDF and EMC GDDR" on page 172	
EMC AutoStart™ • Complete monitoring and automation solution for restarting operations on an alternate local or remote server • Can be implemented with SRDF/S, SRDF/A, TimeFinder and other EMC products • Additional details in "SRDF and open systems automation software" on page 174	

Table 4 SRDF automation software (page 2 of 2)

Description	SRDF site configuration
EMC SRDF Storage Replication Adapter (SRA) for VMware Site Recovery Manager (SRDF for VMAX 40K, VMAX 20K, DMX solutions only) • Complete automation solution for restarting VMware environments • Can be implemented with SRDF/S, SRDF/A, SRDF/Star, and TimeFinder • Additional details in "SRDF and VMware environments" on page 175	

Management tools

EMC support personnel installs and initially configures SRDF at each SRDF site. After SRDF is enabled, you can monitor and control SRDF with one of the available EMC host-based SRDF software suites:

- SRDF Host Component for z/OS — A z/OS subsystem for controlling SRDF processes and monitoring SRDF status by using commands executed from a host. It is not supported in SRDF for VMAX 10K or VMAXe solutions.
- EMC Solutions Enabler — A library of commands that are entered from a command line or from a script.
- Unisphere for VMAX — A browser-based user interface used to manage and monitor SRDF and other Symmetrix operations.
- EMC z/OS Storage Manager (EzSM) — A mainframe services management tool that monitors and reports on Symmetrix mainframe storage from the z/OS operating system perspective. It is not supported in SRDF for VMAX 10K or VMAXe solutions.

SRDF Host Component for z/OS

SRDF Host Component for z/OS is delivered with members of the Mainframe Enablers product family.

EMC SRDF Host Component is a z/OS subsystem for controlling SRDF processes and monitoring SRDF status. You can issue SRDF Host Component commands to both local and remote Symmetrix systems. Commands destined for remote Symmetrix systems are transmitted through local Symmetrix systems through SRDF links. Configuration and status information can be viewed for each device on every Symmetrix system containing SRDF devices.

User interfaces to SRDF Host Component are provided through both the Time Sharing Option (TSO), Interactive System Productivity Facility (ISPF), and batch commands, as well as through the system console. An optional interface is provided for TimeFinder and SRDF commands to centralize commands for both replication products.

The *EMC SRDF Host Component for z/OS Product Guide* located on the EMC Online Support site contains additional information.

The entire Mainframe Enablers suite includes the following components:

- SRDF Host Component for z/OS
- EMC ResourcePak® Base for z/OS
- Consistency Group for z/OS
- TimeFinder/Clone Mainframe Snap Facility
- TimeFinder/Mirror for z/OS
- TimeFinder Utility

Solutions Enabler SRDF component

The Solutions Enabler is a specialized library that consists of commands that can be invoked from a host command line or within scripts. These commands perform control operations on devices and data objects within a Symmetrix storage environment.

Solutions Enabler SRDF extends the basic command set to include SRDF commands for performing control operations on SRDF devices. This component runs on many open systems servers and on mainframe servers that have Linux-based operating system partitions installed.

The *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide* contains additional information.

Unisphere for VMAX

Unisphere for VMAX provides a user interface for the configuration and management of Symmetrix systems and can be used to operate and monitor SRDF remote mirroring functions.

Additionally, Unisphere for VMAX provides health indicators for SRDF/A cache usage, cycle time, and throughput at user-configurable polling intervals.

EMC z/OS Storage Manager

Intended for mainframe users managing complex storage environments, EMC z/OS Storage Manager (EzSM) enables simplified resource management. EzSM can:

- Perform SRDF, TimeFinder/Clone, TimeFinder/Snap, and TimeFinder/Mirror operations
- Perform Group Name Services (GNS) operations
- Monitor and report on Symmetrix mainframe storage.

CHAPTER 2
Symmetrix Enginuity and SRDF Concepts

This chapter explains SRDF concepts. Topics include:

- Symmetrix logical devices, RAID groups, and mirrors 28
- Virtual Provisioning and thin devices 30
- SRDF devices 31
- SRDF device states 46
- SRDF groups 50
- System-level SRDF device and group support 56
- SRDF links 57
- SRDF network protocols and topologies 60
- SRDF modes of operation 65

Symmetrix logical devices, RAID groups, and mirrors

Enginuity presents storage capacity to the host as a series of host-addressable devices, also known as Symmetrix logical devices. When interacting with physical devices (drives), Enginuity views storage associated with a logical device as *RAID groups* and *mirrors.* Enginuity versions 5874 and higher leverage the virtual RAID infrastructure. This infrastructure simplifies configuration changes such as changing the RAID type and benefits SRDF by allowing more flexible SRDF solutions. Enginuity operates at the track level, where a track represents the basic data unit handled in Symmetrix systems.

The RAID group attribute determines how Enginuity lays out data across multiple physical drives that constitute a RAID group. Drives in the same RAID group work in unison to provide data protection, fault detection, error correction, and I/O performance characteristics.

The mirror attribute of a logical device provides information about the physical location and the status of each track associated with that mirror. Enginuity supports up to four mirrors per logical device. If a logical device is configured with multiple mirrors, Enginuity maintains synchronization of the mirrors. If one mirror receives updates and another mirror does not, Enginuity automatically invokes a background process that copies tracks from the mirror with the latest data to another mirror that still holds the old data.

Table 5 on page 28 lists the RAID schemes used in Symmetrix systems. Consult EMC support personnel for configuration rules. Not all RAID schemes are available for all Symmetrix systems, Enginuity versions, and logical device types.

Table 5 RAID schemes used in Symmetrix storage systems (page 1 of 2)

RAID scheme	Description
RAID 0	**Note:** RAID 0 is only supported in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX solutions. A RAID 0 group consists of one or multiple drives (RAID members). If the latter, data is striped across all RAID members to improve the speed of write I/Os. The RAID 0 scheme alone does not protect data (unprotected group) in the event of one drive failure.
RAID 1	A RAID 1 group consists of two mirrored RAID members that maintain identical copies of data. RAID 1 protects data in the event of one drive failure. The RAID 1 scheme provides the highest level of performance and availability for business-critical applications.
RAID 5	A RAID 5 group consists of four or eight RAID members. Data and horizontal parity blocks are distributed across all RAID members. RAID 5 provides data protection because data kept by any single RAID member can be recovered as long as other members of the RAID 5 group remain intact. RAID 5 protects data in the event of one drive failure.
RAID 6	A RAID 6 group consists of 8 or 16 RAID members. Data and horizontal and diagonal parity blocks are distributed across all RAID members. RAID 6 protects data in the event of one or two drive failures.

Table 5 RAID schemes used in Symmetrix storage systems (page 2 of 2)

RAID scheme	Description
Symmetrix RAID 10	**Note:** RAID 10 is only supported in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX solutions. RAID 10 is a nested RAID scheme used in mainframe host environments. A RAID 10 group consists of 4, 8, or 16 RAID members. Data is striped (RAID 0) across two, four, or eight drives which are then RAID 1 mirrored. In Symmetrix, a RAID 10 mainframe device is a concatenated device that consists of four Symmetrix logical devices. RAID 10 scheme always protects data in the event of one drive failure and can sustain multiple drive failures as long as at least one member of every mirrored pair in the RAID 10 group remains intact. **Note:** You can create a RAID 10 equivalent in open systems environments by using striped metadevices with each meta member configured as a RAID 1 group. Like in RAID 10 groups, data is striped across mirrored pairs. **Note:** The VMAX 10K arrays do not support RAID 10. Designed for 100% virtually provisioned storage environments, the VMAX 10K array features virtually provisioned (thin) volumes widely striped across RAID 1 disk pairs to provide both I/O concurrency and the RAID 1 protection level. The benefits are equal or superior to those provided by RAID 10 (Mainframe) or striped meta volumes (Open Systems) in the VMAX 20K and VMAX 40K arrays in non-virtually provisioned storage environments.

Figure 2 on page 29 illustrates logical device Symm_01 with a single mirror configured as RAID 5 (3+1) group and logical device Symm_02 with a single mirror configured as RAID 1 group.

Mirrors and their RAID groups shown in Figure 2 on page 29 reside within the same Symmetrix system. These mirrors are also referred to as *local mirrors*.



Figure 2 Two logical devices, each with a single mirror (RAID 5 and RAID 1)

Symmetrix session

Enginuity uses the Symmetrix session mechanism to deliver Enginuity features without interrupting production applications. SRDF/A and SRDF/Star all use the session mechanism.

Symmetrix sessions have the following common characteristics:

- Users configure, initiate, poll, and terminate Symmetrix sessions via Symmetrix host-based control software that typically runs on a dedicated controlling host.
- Enginuity provides all session-specific functionality without requiring any resources from the application hosts.
- Enginuity limits the number of concurrent sessions per Symmetrix device to 16. Symmetrix host-based control software may impose additional restrictions.

Virtual Provisioning and thin devices

EMC Virtual Provisioning™ (VP) is an Enginuity feature available with Enginuity versions 5773 and higher. Virtual Provisioning allows for storage to be allocated and accessed on demand from a pool of storage that services one or many applications. This type of storage has multiple benefits:

- Enables LUNs to be “grown” into over time with no impact to the host or application as space is added to the thin pool.
- Delivers space from the thin pool on demand.
- Provides wide striping for a thin pool.
- Relieves the storage administrator’s efforts of physical device/LUN configuration.

The Virtual Provisioning feature introduces three new concepts: thin devices, data devices, and thin pools. Thin devices can be created with an inflated capacity, because the actual storage space for the data written to the thin devices is on the data devices. In this way, when additional storage is needed, more data devices can be created in the thin pool.

Table 6 on page 30 lists the thin device types that are supported with different Enginuity versions.

Table 6 Device types supported by Virtual Provisioning

Enginuity version	Thin device type supported in VMAX 40K, VMAX 20K/VMAX, DMX	Thin device type supported in VMAX 10K/VMAXe
Prior to 5876.82.57	FBA	FBA
5876.82.57	FBA IBM i 512-byte D910 CKD 3390	FBA
5876.229.145	FBA IBM i 512-byte D910 CKD 3390	FBA IBM i 512-byte D910

Figure 3 on page 31 shows an example of a write I/O to thin devices.



Figure 3 Write I/O to thin devices

With Virtual Provisioning, you can gradually add physical storage as they grow your storage environment without having to reconfigure their host-addressable logical devices. Once a thin device is configured and mapped to the host, the allocated storage space at any point in time is only the space, to which the host has already written, regardless of the storage capacity specified at configuration time. For example, if you initially configure a 100 GB thin device and the host has written 10 KB of data to this thin device, only 10 KB of storage space is allocated at that point in time. That is why a thin device has no pre-allocated storage capacity when initially configured. As the host write I/Os for a thin device arrive, Enginuity allocates only as much storage space as required to service the host write I/O.

"Thin SRDF devices" on page 39 provides details about thin devices in SRDF environments.

SRDF devices

An SRDF device is a Symmetrix logical device paired with another Symmetrix logical device that resides in a remote Symmetrix system. The Symmetrix systems on both sites are connected to each other through SRDF links. An SRDF pair consists of the ***primary*** SRDF device (R1) residing on one and the ***secondary*** SRDF device (R2) residing on another Symmetrix system.

Note: The mainframe Page Packs devices and Stand Alone Dump devices are not recommended to be configured as SRDF devices.

Primary (R1) and secondary (R2) devices

The primary (R1) device contains production data that is *actively mirrored* across the SRDF links to the other member of the SRDF pair, the secondary (R2) device. Active mirroring indicates that the secondary device (R2) maintains a restartable, mirrored copy of the production data at the remote site while the production host is running I/Os to the R1 device.

Figure 4 on page 32 illustrates a typical open systems host (left) and a typical mainframe host SRDF environment (right).



Figure 4 SRDF active remote mirroring

During ordinary SRDF operations in a typical open systems host environment, the production host connected to Symmetrix A has Read/Write access to the R1 device and Read Only (Write Disabled) access to the R2 device. If another host is attached to Symmetrix B, that host also has Read Only (Write Disabled) access to the R2 device.

During ordinary SRDF operations in a typical mainframe host environment, the production host has Read/Write access to the R1 device and no access to the R2 device. The production and the remote host both see the R2 device as Read Only or Not Ready.

The R2 devices resemble local mirrors in the way they receive data from the R1 devices at the production site. Enginuity automatically propagates updates from the R1 devices to the R2 devices, similar to the way it synchronizes multiple local mirrors.

The R1 and the R2 devices are Symmetrix logical devices with a special mirror configuration. The R1 and the R2 devices have local mirrors and at least one *SRDF mirror,* also known as a remote mirror. An SRDF mirror indicates that data is mirrored across the SRDF links at another Symmetrix site.

Figure 5 on page 33 shows the mirror view of the R1 and the R2 devices. In this example, all local mirrors are configured as RAID 1 groups.



Figure 5 R1 and R2 devices, mirror configuration (Symmetrix view)

SRDF mirrors are unique to SRDF devices. An R1 SRDF mirror indicates that data stored on the R1 device (R1 local mirrors) is also remotely mirrored to the R2 device (R2 local mirrors). Likewise, an R2 SRDF mirror indicates that data stored on the R2 device (R2 local mirrors) is remotely mirrored to the R1 device (R1 local mirrors).

When a write I/O arrives to the R1 device in Symmetrix A, Enginuity propagates updates to the local mirror in Symmetrix A and also sends the updates across the SRDF links to update the R2 device. The host perceives this operation as a write I/O to the R1 device.

The important distinction between R1 and R2 devices is their accessibility to the hosts connected to either Symmetrix A or Symmetrix B as illustrated in Figure 4 on page 32 and Figure 5 on page 33. The production host has Read/Write permissions to the R1 device. At the same time, any host connected to Symmetrix B can have either Read Only (Write Disabled) or no permissions (Not Ready) to the R2 device as long as the SRDF relationship between the R1 and R2 devices exists.

Primary devices (R1, R11)

SRDF differentiates between ordinary primary devices as described in "SRDF devices" on page 31 and concurrent primary devices (R11) typically used in three-site SRDF solutions. concurrent R11 devices are configured with local mirrors and two unique R1 SRDF mirrors.

R11 devices maintain mirrored copies of data across the SRDF links at two remote sites. Figure 6 on page 34 illustrates a concurrent SRDF configuration with Symm_14 as the R11 device paired with two R2 devices, Symm_05 in Symmetrix B and Symm_04 in Symmetrix C. As in previous examples, the local mirror is configured as a RAID 1 group.

"Concurrent SRDF" on page 122 provides more details about concurrent SRDF.

Note: The SRDF mode for SRDF devices with multiple SRDF mirrors is set at the SRDF mirror level. For example, one R1 SRDF mirror can operate in synchronous and another in asynchronous mode. Each SRDF mirror must be configured to a separate SRDF group. "SRDF groups" on page 50 provides more details.



Figure 6 R11 device with two R1 SRDF mirrors, paired with two R2 devices

Dual-role (cascaded) SRDF devices (R21)

R21 devices are dual-role SRDF devices typically used in three-site solutions as shown in Figure 7 on page 35. An R21 device can simultaneously assume the roles of an R2 and an R1 device. R21 devices have an R1 SRDF and an R2 SRDF mirror.

Note: The SRDF mode for SRDF devices with multiple SRDF mirrors is set at the SRDF mirror level. For example, the R2 SRDF mirror can operate in synchronous mode and the R1 SRDF mirror in adaptive copy disk mode. Each SRDF mirror must be configured to a separate SRDF group. "SRDF groups" on page 50 provides more details.

The R2 SRDF mirror of the R21 device in Symmetrix B receives data from the R1 SRDF mirror of the R1 device in Symmetrix A. The R2 SRDF mirror of the R2 device in Symmetrix C receives data from the R1 SRDF mirror of the R21 device in Symmetrix B. When data from Symmetrix A reaches Symmetrix B, both SRDF mirrors of the R21 device receive updates, but their behavior differs:

- After the R2 SRDF mirror of the R21 device receives updates, data are written to drives in Symmetrix B. If SRDF operates in synchronous mode, the SRDF emulation sends the acknowledgement to Symmetrix A.
- After the R1 SRDF mirror of the R21 device receives updates, the SRDF emulation sends data across the SRDF links to the R2 SRDF mirror of the R2 device in Symmetrix C as shown in Figure 7 on page 35.

The R2 SRDF mirror of the R21 device can operate in all SRDF modes. However, the R1 SRDF mirror of the R21 device can operate in only adaptive copy disk or asynchronous mode.

The R21 devices appear like the R2 devices to the hosts connected to Symmetrix systems in a cascaded SRDF topology (Figure 7 on page 35). As long as the R1->R21->R2 SRDF relationship is established, no host can have write access to the R21 device. "Cascaded SRDF" on page 126 provides more details about cascaded SRDF.



Figure 7 R21 device used in cascaded SRDF operations

Diskless R21 devices (DL R21)

Diskless R21 devices (DL R21) operate like R21 devices except that DL R21 devices do not store full copies of data and have no local mirrors (RAID groups) configured to them. DL R21 devices store only the differences that would be owed to Symmetrix C if Symmetrix A failed.

Unlike R21 devices that are Symmetrix logical devices, the DL R21 devices are special cache devices that cannot be assigned to the host. At any point in time, data resides in Symmetrix B cache memory and on local mirrors in Symmetrix A and Symmetrix C. Like R21 devices, DL R21 devices have two SRDF mirrors, an R1 SRDF and an R2 SRDF mirror. However, during the creation/deletion of DL R21 devices or during automated failover/failback operations, DL R21 devices may temporarily have only one SRDF mirror. Figure 8 on page 36 illustrates a three-site configuration that uses DL R21 devices in Symmetrix B.



Figure 8 DL R21 devices used in SRDF/EDP configuration

This configuration is also known as the SRDF/EDP solution and looks very similar to that shown in Figure 7 on page 35 except that the R1 SRDF mirror of the DL R21 device operates in adaptive copy write pending mode or in asynchronous mode. "Extended Distance Protection" on page 128 provides more details about SRDF/EDP.

Note: The R1 and R2 devices connected to DL R21 devices must be either both thick or both thin.

DL R21 devices have to be preconfigured within the Symmetrix system before they can be placed in the SRDF relationship. The configuration process involves the following steps:

1. Create DLR1 -> R2 pair between Symmetrix B and Symmetrix C
2. Create R1 -> DLR2 pair between Symmetrix A and Symmetrix B.

DLR1 and DLR2 are intermediate device types that exist only during the configuration process. When step 2 completes, the DL R21 device is created and the resulting configuration is R1 --> DL R21 --> R2

The following considerations apply to DL R21 devices:

- Cannot be assigned to the host.
- Cannot be thin.
- Cannot reside on the Symmetrix VMAX 10K or VMAXe systems.
- R21 and DL R21 devices cannot be used in the same SRDF/A session.
- DL R21 devices and R21 devices can be configured to the same SRDF group if the SRDF/A mode is not used.
- DL R21 devices cannot participate in TimeFinder operations if Enginuity version 5876.82.57 or lower runs on the DL R21 devices.

Secondary devices (R2, R22)

Secondary devices can also be ordinary R2 devices described in "SRDF devices" on page 31 or concurrent R22 devices configured with two R2 SRDF mirrors that receive data from two different R1 devices. R22 devices are supported with Enginuity version 5773 or higher.

Only one of the R2 SRDF mirrors that belong to the R22 device can accept reads or writes from its R1 device at a time. The other SRDF mirror is marked as *blocked* to its R1 device. R22 devices are designed to benefit SRDF/Star solutions that operate by using only the active SRDF links, although a recovery path is in place to ensure that you can fail over to another site in case of a disaster at the production site. "SRDF/Star" on page 130 provides more details about SRDF/Star.

Figure 9 on page 38 illustrates an R22 device used in a cascaded SRDF/Star and in a concurrent SRDF/Star solution. In both cases, the R22 device is configured with two SRDF mirrors, but only the SRDF mirror that accepts I/Os from active SRDF links is enabled at a time.



Figure 9 R22 in cascaded and concurrent SRDF/Star solutions

Both SRDF/Star solutions shown in Figure 9 on page 38 can also be implemented with R2 devices. Having one of the two SRDF mirrors always blocked, the R22 devices function as ordinary R2 devices.

However, the important advantage of R22 devices is that they do not require any reconfiguration steps in the event of an SRDF/Star failover or failback operation. You can quickly perform recovery operations without having to remove old SRDF pairs and create new ones. "Recovery operations" on page 80 provide more details.

Note: R22 devices are required in the SRDF/Star solutions including VMAX 10K or VMAXe systems.

Thin SRDF devices

With Enginuity versions 5773 and higher, SRDF supports Virtual Provisioning and thin devices. Thin devices, also known as Virtual Provisioning devices, are Symmetrix devices that do not need to have physical storage completely allocated when devices are created. In contrast to thin devices, thick devices, also known as standard devices, are Symmetrix devices whose storage capacity is fully allocated when devices are created.

Table 7 on page 39 provides the chronology of support for Virtual Provisioning features starting with Enginuity version 5773 in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX solutions. The new features for each Enginuity version are listed and described.

Table 7 Support for Virtual Provisioning in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX

Enginuity version	SRDF support for Virtual Provisioning (FBA devices)	SRDF support for Virtual Provisioning (CKD devices)
5773	• Supports thin R1 and R2 devices • Supports thin-to-thin connectivity in two-site SRDF solutions	No support
5874	• Supports thin R11, R22, and R21 devices • Supports thin-to-thin connectivity in concurrent SRDF, cascaded SRDF, and SRDF/Star solutions.	No support
5875.135.91	• Supports thin R1 and thin R2 devices in SRDF/EDP solutions. • Supports thick-to-thin connectivity in two-site and concurrent SRDF solutions with migration only SRDF support. "SRDF Interfamily Connectivity" on page 143 provides details.	No support
5875.198.148	• Supports thick-to-thin connectivity in two-site and concurrent SRDF solutions with full SRDF functionality. "SRDF Interfamily Connectivity" on page 143 provides details.	No support
5876.82.57	No new features	• Supports thin R1, R11, R21, R2, R22 devices • Supports thin-to-thin connectivity in two-site SRDF, concurrent SRDF, cascaded SRDF, and SRDF/Star solutions. "SRDF Interfamily Connectivity" on page 143 provides details.
5876.159.102 and higher	• Supports thick-to-thin connectivity in cascaded SRDF and SRDF/Star solutions with full SRDF functionality. "SRDF Interfamily Connectivity" on page 143 provides details.	• Supports thick-to-thin connectivity in two-site SRDF, concurrent SRDF, cascaded SRDF, and SRDF/Star solutions with full SRDF functionality. "SRDF Interfamily Connectivity" on page 143 provides details.

Table 8 on page 40 provides the chronology of support for Virtual Provisioning features in SRDF for VMAX 10K or VMAXe solutions. The new features for each Enginuity version are listed and described.

Table 8 Support for Virtual Provisioning in SRDF for VMAX 10K or VMAXe

Enginuity version	SRDF support for Virtual Provisioning	Description
5875.231.172	• Supports thin R1 and R2 devices • Supports thin-to-thin connectivity in two-site SRDF solutions • Supports thick-to-thin connectivity in two-site SRDF solutions	For thick-to-thin connectivity, "SRDF Interfamily Connectivity" on page 143 provides details.
5876.82.57	• Supports thin R11, R21, and R22 devices • Supports thin-to-thin connectivity in concurrent SRDF, cascaded SRDF, and SRDF/Star solutions • Supports thick-to-thin connectivity in concurrent SRDF solutions	For thick-to-thin connectivity, "SRDF Interfamily Connectivity" on page 143 provides details.
5876.159.102 and higher	• Supports thick-to-thin connectivity in cascaded SRDF and SRDF/Star solutions • Supports thin R1 and R2 devices in SRDF/EDP	For thick-to-thin connectivity, "SRDF Interfamily Connectivity" on page 143 provides details.

SRDF thick-to-thin support

With Enginuity version 5875.135.91 or higher, SRDF supports FBA device pairs which include thick devices on one side and thin devices on another side of the SRDF links. SRDF supports CKD device pairs with mixed thick and thin devices with Enginuity version 5876.159.102 or higher.

To support SRDF thick-to-thin connectivity between different Symmetrix hardware models running different Enginuity versions, specific Enginuity versions are required on the systems containing thick and thin devices. In most cases, SRDF thick-to-thin is supported with full SRDF functionality.

However, for the following pairs running in a two-site SRDF solution or on any leg of a concurrent SRDF solution, SRDF thick-to-thin connectivity is only supported with migration only SRDF:

- Enginuity version 5671 – 5875.135.91 and higher
- Enginuity version 5773 – 5875.135.91

"SRDF Migration Operations" on page 149 provides details about migration only SRDF support.

"SRDF Interfamily Connectivity" on page 143 provides greater details about SRDF thick-to-thin support in each SRDF solutions.

Zero space reclamation

With Enginuity version 5875.135.91 or higher, SRDF supports zero space reclamation for FBA devices. Zero space reclamation is an Enginuity feature that allows you to remotely mirror a thick SRDF device to a thin SRDF device while avoiding mirroring pre-allocated zero data chunks that may be associated with a thick SRDF device. The SRDF advanced zero data detection capability guarantees that thin SRDF devices at Enginuity version

5875.135.91 or higher provide you with storage capacity savings even if the SRDF partner of the thin device is a thick SRDF device. SRDF zero space reclamation is a system-wide setting enabled by default.

With Enginuity version 5876.159.102 or higher, zero space reclamation can be issued from the thick R1 devices to the thin R2 devices across the SRDF links without breaking the SRDF device pair. Zero space reclamation is also supported for CKD devices with Enginuity 5876.159.102 or higher.

Note: Zero space reclamation is available in both full SRDF functionality support and migration only SRDF support.

Open systems metadevices as SRDF devices

Open systems metadevices are comprised of multiple Symmetrix devices, the meta head device, a number of meta members, and the meta tail device. The application host addresses a metadevice through the meta head device and perceives the Symmetrix devices grouped in the metadevice as a single addressable storage device.

Metadevices are used to create storage devices larger than the maximum Symmetrix logical device size. Open systems metadevices can be used in SRDF solutions, if certain requirements are met.

If SRDF devices are open systems metadevices, these R1 and R2 devices must be configured as metadevices with equal numbers of meta members. Figure 10 on page 41 shows striped R1 and R2 metadevices (top) and concatenated R1 and R2 metadevices (bottom). If concatenated SRDF metadevices are used, the R1 and the R2 devices must be of the same size.



Figure 10 Striped and concatenated SRDF metadevices

Note: Figure 10 on page 41 shows logical devices. Physical layout (RAID groups) depends on the RAID group attributes associated with each logical device that belongs to a metadevice. Contact EMC support personnel for more details about additional considerations that apply to remote mirroring of metadevices.

Mainframe RAID 10 SRDF devices

Note: Support for mainframe RAID 10 SRDF devices is only available in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX solutions.

With Enginuity version 5568 or higher, R1 devices configured as RAID 10 groups can be remotely mirrored to R2 devices configured as non-RAID 10 groups as long as the R1 mainframe device size and the R2 mainframe device size are equal. With Enginuity version 5569 or higher, the R2 mainframe devices can be larger than the R1 mainframe devices.

R1 devices only, R2 devices only, or both can be configured as RAID 10 groups.

EMC Compatible Peer

Note: EMC Compatible Peer is only available in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX solutions.

Enginuity versions 5568 and higher support native IBM Peer-to-Peer Remote Copy (PPRC) commands through a Symmetrix feature called EMC Compatible Peer. An SRDF device enters PPRC mode when it receives a **PPRC CESTPAIR** command from the host. Once the device is in PPRC mode, only the PPRC commands from the host can control this device.

The EMC Compatible Peer feature is available in mainframe environments for all CKD device types supported by the IBM PPRC feature. Table 9 on page 42 details Enginuity support for PPRC.

Table 9 Enginuity support for PPRC

Enginuity version	EMC Compatible Peer support
5568	PPRC version 1, architectural level 2
5671 and higher	PPRC version 1, architectural levels 3 and 4 and hyperswap support including failover and failback functionality

Converting to/from PPRC mode devices

With Enginuity versions 5874 or higher running on Symmetrix systems on both sides of the SRDF links, you can convert from the PPRC mode to SRDF devices and from SRDF to PPRC mode devices with minimal downtime between the R1 and R2 devices.

This feature is supported only for CKD 3380 and 3390 devices and requires an online configuration change at the primary site performed by EMC support personnel. The SRDF links remain active during this operation, but the relationship between the R1 and R2 devices is suspended before the conversion process.

After the PPRC mode devices are converted to SRDF devices, either the EMC support personnel or you must resume the SRDF relationship between the primary and the secondary devices.

After the SRDF devices are converted to PPRC mode, you must run the **CESTPATH** and **CESTPAIR** commands.

The following SRDF devices cannot be converted to PPRC mode:

- Static SRDF devices

  Note: EMC Compatible Peer mode requires dynamic SRDF devices. "Dynamic SRDF devices" on page 43 provides more information.

- SRDF devices with two SRDF mirrors (R11, R21, DLR21, R22)
- SRDF devices in an SRDF/Star relationship
- SRDF devices in an active SRDF/A group
- SRDF R1 devices configured to an active SRDF/Consistency Group
- SRDF devices configured to different SRDF groups. Only one SRDF group at a time can be converted.

Both conversion processes support static and dynamic SRDF groups. EMC recommends that static SRDF groups be used with the EMC Compatible Peer feature.

Dynamic SRDF devices

Dynamic SRDF devices are SRDF devices that allow flexible control over the SRDF solution. Dynamic SRDF device attributes are stored dynamically in the mirrored and protected Symmetrix cache memory. Dynamic SRDF devices are only enabled in the Symmetrix configuration file. You can configure and control SRDF devices by using EMC host-based SRDF software.

EMC host-based SRDF software can initiate the following actions to modify dynamic SRDF device attributes:

- Create a new R1/R2 pair relationship from non-SRDF devices.
- Terminate and establish an SRDF relationship with a new R2 device.
- Swap personalities between R1 and R2 devices.
- Move R1/R2 pairs between different SRDF groups.

Note: Creating, terminating, and swapping personality functions are not allowed in Enterprise Systems Connection (ESCON) FarPoint topologies.

Dynamic SRDF devices are supported in the following topologies:

- Fibre Channel SRDF, point-to-point connection
- Fibre Channel SRDF, switched Fibre Channel
- GigE SRDF
- ESCON, point-to-point connection (SRDF for VMAX 40K, VMAX 20K/VMAX, DMX solutions only).

Table 10 on page 44 details Enginuity support for dynamic SRDF devices.

Table 10 **Enginuity support for dynamic SRDF devices**

Enginuity version	Dynamic SRDF feature
5567 and higher	R1/R2 personality swap
5568 and higher	• Create SRDF devices • Create SRDF pairs • Delete SRDF pairs
5x71 and higher	Support for dynamic SRDF/A devices
5773 and higher	Move SRDF pairs

R1/R2 personality swap

Dynamic SRDF devices allow an R1/R2 personality swap. After a personality swap, R1 becomes an R2 device and R2 becomes an R1 device. The personality swap allows the R2 side to take control over operations while being remotely mirrored at the primary site.

This feature is especially useful in application failover operations. If an application fails at the production site, R1/R2 personalities are swapped, the application restarted at the remote site, and production resumed at the secondary site.

The following steps are required to perform an R1/R2 personality swap and to resume SRDF operations:

1. Suspend the SRDF remote mirroring.
2. Perform a personality swap by converting the R1 to R2 and the R2 to R1 devices.
3. Use EMC host-based software to determine the synchronization direction and synchronize the R1 and the R2 devices.
4. Resume remote mirroring.

At this point, the SRDF pair is available for operations, accepting host I/Os at the secondary site where the new R1 device resides and is remotely mirrored to the new R2 device at the primary site.

A dynamic R1/R2 personality swap is not supported under the following conditions:

- If the R2 device is larger than the R1 device.
- If the device to be swapped is participating in an active SRDF/A session.
- In SRDF/EDP topologies where attempted reconfiguration operations would result in an illegal device type such as a DL R2 or DL R1.

- If the device to be swapped is the target device of any TimeFinder or EMC Compatible Flash operations.

The *SRDF Host Component for z/OS Product Guide* and the *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide* provide more information about controlling the R1/R2 swap operations and specific host platform requirements.

SRDF device states

The SRDF devices undergo logical device state transitions like ordinary Symmetrix logical devices and SRDF mirror state transitions that are unique to SRDF devices. Because of the SRDF pairing between the R1 and the R2 devices, the overall state of the SRDF operations is adequately described only by examining both the R1 and R2 devices and their SRDF mirrors.

Note: With Enginuity version 5874 or higher, the state change timestamp provides you with information about the most recent SRDF mirror transitions on the primary and the secondary sites. The timestamp data includes the time, type, and the cause of the last SRDF mirror state change to help you to determine the state of the SRDF device pair.

SRDF device state is therefore determined by a combination of two substates as shown in Figure 11 on page 46:

- Host interface view (reflects the SRDF device state as seen by the host)
- SRDF view (reflects the SRDF mirror states and internal SRDF device states).



Figure 11 Host interface and SRDF view

Host interface view

The host interface view reflects the SRDF device states as seen by the host that connected to the device.

R1 device states

An R1 device can have one of the following states to the host connected to the primary system:

- Read/Write (Write Enabled) — In this state, the R1 device is available for read/write operations. This is the default R1 device state.
- Read Only (Write Disabled) — In this state, the R1 device responds with Write Protected to all write operations to that device.
- Not Ready — In this state, the R1 device responds Not Ready to the host for read and write operations to that device.

Domino modes and data integrity

Note: Domino modes are only supported in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX solutions.

Under certain conditions, the SRDF devices can be forced into the Not Ready state to the host. For example, if the host I/Os cannot be delivered across the SRDF links, Domino attributes are used to stop all subsequent write operations to both R1 and R2 devices to avoid data corruption. While such a shutdown temporarily halts production processing, domino modes can protect data integrity in case of a *rolling disaster*. Domino modes can be set at the device level (domino mode) or at the SRDF group level (link domino mode).

If domino mode is set and the R1 device cannot successfully mirror data to the R2 device, upon the next host write to the R1 device, this device will become Not Ready to the host connected to the primary system.

Link domino mode is set at the SRDF group level at either side of the SRDF links. If this attribute is set and the last available link in the SRDF group fails upon the next host write to any R1 device in the SRDF group, all R1 devices in the SRDF group become Not Ready to their hosts.

Once the Not Ready to the host condition is set, you must re-enable the R1 devices using EMC host-based SRDF software to make them Ready to the host. Both domino modes cause host applications to stop. This is an extreme measure of protecting data integrity. SRDF/Consistency Group (SRDF/CG) implements a more moderate mechanism. "Dependent-write consistency and SRDF/CG" on page 77 contains additional information.

R2 device states

An R2 device can have one of the following states to the host connected to the secondary system:

- Read Only (Write Disabled) — In this state, the secondary (R2) device responds Write Protected to the host for all write operations to that device.
- Read/Write (Write Enabled) — In this state, the secondary (R2) device is available for read/write operations. This state is possible in recovery or parallel processing operations.
- Not Ready — In this state, the R2 device responds Not Ready (Intervention Required) to the host for read and write operations to that device.

Invalid tracks attribute

If used with a special device attribute, the invalid tracks attribute, the Not Ready state of the R2 device provides information about the state of the SRDF device pair. The invalid tracks attribute can only be set for secondary devices. When the invalid tracks attribute is enabled, SRDF makes the secondary device Not Ready to the host if the secondary device is not synchronized with its R1 device.

The purpose of the invalid tracks attribute is to inform you that data on the R2 devices may not be suitable for use in disaster recovery situations. You then decide whether to proceed using data from R2 devices that were made Not Ready by the invalid tracks attribute. You can reset the Not Ready condition set by the invalid track attribute via EMC host-based SRDF software.

Note: In Symmetrix systems, invalid tracks indicate a special mirror condition. In SRDF solutions, tracks marked as invalid for an SRDF mirror are considered *owed* to the SRDF device across the SRDF links.

SRDF view

The SRDF view reflects the SRDF mirror states and internal SRDF device states, which indicate whether the device is available to send data across the SRDF links, and whether the device is able to receive host-based software commands.

R1 device states

A primary (R1) device can have the following states for SRDF operations:

- Ready — The R1 device is ready for SRDF operations.

  This condition occurs when the R1 device is able to send data across the SRDF links. This condition is true even if local mirror(s) of the R1 device are Not Ready for I/O operations.

  If the R1 local mirror(s) fail and the primary system is still able to send data across the SRDF links to the R2 device, the host connected to the primary system sees the R1 device Write Enabled (Read/Write) since all read or write operations continue uninterrupted to the secondary (R2) device. At the same time, the host interface state for the R2 device to any host connected to the secondary system remains Read Only or Not Ready. This situation demonstrates why SRDF operations are known as remote mirroring—the host is unaware of any failures on the primary site as long as the primary system can send data to the secondary site. The R2 device, write-disabled to any host connected to the secondary system, receives data from the SRDF links as if it were a local mirror.

- Not Ready (SRDF mirror Not Ready) — The R1 device is Not Ready for SRDF operations. In mainframe environments, this R1 SRDF mirror Not Ready state is also referred to as Target Not Ready (TNR). It is important to understand that this is an R1 SRDF mirror state, not the R1 or R2 device state.

  This condition occurs when the remote mirroring is halted and write updates do not propagate to the R2 device. Invalid tracks accumulate on the R1 SRDF mirror, indicating that updates are *owed* to the R2 devices.

Note: When the R2 device is placed into a Read/Write state to the host, the corresponding R1 device is automatically placed into the SRDF mirror Not Ready state.

R2 device states

A secondary (R2) device can have one of the following states for SRDF operations:

- Ready — In this state, the R2 device receives the updates propagated across the SRDF links and can accept SRDF host-based software commands.
- Not Ready — The R2 device cannot accept SRDF host-based software commands, but can still receive updates propagated from the primary system.
- Link blocked (LnkBlk) — This state is unique to R2 SRDF mirrors that belong to R22 devices. In this state, one of the R2 SRDF mirrors cannot receive writes from its associated primary (R1) device. In normal operations, one of the R2 SRDF mirrors of the R22 device is in this state.

The accessibility of a particular SRDF device to the host depends on both the host and the Symmetrix view of the SRDF device state. Table 11 on page 49 and Table 12 on page 49 list host accessibility levels of R1 and R2 devices.

Table 11 Primary (R1) device accessibility

Host interface state	SRDF state	Accessibility
Read/Write	Ready	Read/Write
	Not Ready	Depends on secondary (R2) device availability
Read Only	Ready	Read Only
	Not Ready	Depends on secondary (R2) device availability
Not Ready	Any	Unavailable

Table 12 Secondary (R2) device accessibility

Host interface state	SRDF R2 state	Accessibility
Write Enabled (Read/ Write)	Ready	Read/Write
	Not Ready	Read/Write
Write Disabled (Read Only)	Not Ready	Read Only
	Ready	Read Only
Not Ready	Any	Unavailable

The *EMC Solutions Enabler Symmetrix SRDF CLI Product Guide* or the *Symmetrix SRDF Host Component for z/OS Product Guide* provide more information on how device states are presented in open systems and mainframe operating environments, respectively.

SRDF groups

Every SRDF device that operates in an SRDF configuration must have its SRDF partner device configured in the Symmetrix system across the SRDF links. Every SRDF device must also be added to as many SRDF groups as the number of SRDF mirrors configured to a single SRDF device.

"SRDF devices" on page 31 provides details about SRDF devices and their SRDF mirror configuration.

The SRDF group defines the logical relationship between SRDF devices and SRDF directors on both sides of the SRDF links. It is comprised of a range of SRDF devices and SRDF directors that reside on a given Symmetrix system. The SRDF mirrors that belong to these SRDF devices point to the SRDF partner devices that reside on another Symmetrix system and are configured to the partner SRDF group.

Each SRDF group communicates with the partner SRDF group that is defined in another Symmetrix system located across the SRDF links. Each SRDF group contains in its configuration parameters the partner Symmetrix system identification and the set of SRDF directors that belong to the partner SRDF group.

Note: Multiple SRDF groups can be configured on a single Symmetrix system.

As soon as an empty SRDF group is created on one Symmetrix system, a partner SRDF group must be created on another Symmetrix system. This is how the SRDF directors and ports on both sides of the SRDF links are mapped. The SRDF directors assigned to each group share CPU processing power, SRDF ports, and serve all SRDF devices that are added to a particular SRDF group. SRDF directors on each side of the SRDF links cooperate to support regular SRDF I/O operations. Apart from sharing workload among SRDF directors that belong to the same group, the SRDF groups provide resiliency as they can withstand an isolated incident such as a single SRDF port failure.

After an empty SRDF group is created on the primary system and its partner SRDF group on the secondary system, SRDF device pairs can be added to the SRDF groups on both sides of the SRDF links. SRDF device pairs are created when they are added to the SRDF groups. The SRDF mode of operation is also defined at that time.

Prior to Enginuity version 5876.82.57, EMC recommends that you do not configure an SRDF group in SRDF/S and an SRDF group in SRDF/A on the same directors. With Enginuity version 5876.82.57 or higher running on the primary system, you can configure an SRDF/S group and an SRDF/A group on the same SRDF directors on the primary site. In the mixed mode, the decision on how to allocate CPU resources is based on a group selection policy that can be adjusted. You can use the EMC host-based SRDF software to enable or disable the group selection policy. You can also set and view the SRDF director CPU resource distribution assigned to each type of workload on a specific director.

EMC host-based SRDF software provides commands that allow you to simultaneously create SRDF pairs over a range of SRDF devices and add these SRDF pairs to the corresponding SRDF groups on each Symmetrix system.

The following example summarizes the steps required to configure SRDF devices to SRDF groups in an SRDF configuration that consists of Symmetrix A and Symmetrix B:

1. On Symmetrix A, create a new SRDF group (01), assign SRDF directors to it (for example, SRDF_dir_1, SRDF_dir_2, and SRDF_dir_3) and simultaneously create a partner SRDF group on Symmetrix B. For example, create a partner SRDF group (02) with the following SRDF directors assigned to it: SRDF_dir_4, SRDF_dir_5, and SRDF_dir_6.

   Figure 12 on page 51 illustrates this example.



Figure 12 Configuring SRDF groups

2. Suppose that Symmetrix logical devices 10, 11, and 12 reside on Symmetrix A and Symmetrix logical devices 21, 22, and 23 reside on Symmetrix B. Using the EMC host-based SRDF control software, issue the appropriate command to create the following dynamic SRDF device pairs, add the devices 10, 11, and 12 to the SRDF group 01 and specify the corresponding SRDF mode.

R1	R2
10	21
11	22
12	23

   The SRDF partner devices will be automatically added to the SRDF group 02. Figure 13 on page 51 illustrates the result of this operation.



Figure 13 SRDF groups with devices that mirror between Symmetrix A and Symmetrix B

SRDF directors configured to the SRDF group 01 now work together servicing the R1 devices and sending data across the SRDF links to SRDF directors configured to SRDF group 02 that also work together in Symmetrix B, servicing the R2 devices.

The *SRDF Host Component for z/OS Product Guide* and the *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide* provide more information about controlling SRDF groups.

Group types

Like SRDF devices, SRDF groups can be defined as:

- Static SRDF groups — Created and saved in the Symmetrix configuration file. With Enginuity version 5875.135.91, you can add multiple SRDF groups within a single online configuration change procedure.

  Note: Static SRDF groups are only supported in SRDF for VMAX 40K, VMAX 20K/VMAX, DMX solutions.

  Note: EMC recommends that static SRDF groups be used with EMC Compatible Peer.

- Dynamic SRDF groups — Created and controlled through EMC host-based SRDF software, stored in Symmetrix cache.

  Symmetrix logical devices must be dynamic-capable devices to allow EMC host-based SRDF control software to configure SRDF pairs. When you create SRDF pairs through the EMC host-based SRDF control software, they convert dynamic-capable Symmetrix devices into SRDF devices.

SRDF groups cannot have both GigE and Fibre Channel directors defined to the same SRDF group. If Fibre Channel is used, Switched Fibre must be enabled to support multiple SRDF groups per SRDF director. This setting is enabled by default with Enginuity code levels 5773 or higher.

Moving dynamic SRDF devices between SRDF groups

With Enginuity versions 5773 or higher, you can move dynamic SRDF devices between SRDF groups without having to process a full synchronization. This feature is helpful when you:

- Transition to a different SRDF topology and require minimal exposure on device move operations.
- Add new SRDF devices to an existing SRDF/A group and require fast synchronization with the existing SRDF/A devices in the group. For example, if you want to add more SRDF devices to an existing SRDF/A group, proceed as follows:

  1. Create a new SRDF group and its partner SRDF group across the SRDF links. Set the mode of operation to adaptive copy disk mode.
  2. Add new SRDF device pairs to the newly-created SRDF groups and allow the devices to initially synchronize in adaptive copy disk mode.
  3. Move the SRDF pairs to the existing SRDF/A groups. The newly added devices will synchronize with the rest of the SRDF/A group within a few SRDF/A cycle times.

Figure 14 on page 53 shows an example.



Figure 14 Moving SRDF devices between SRDF groups (full move)

A Dynamic SRDF move operation can be either a full move as shown in Figure 14 on page 53 or a half move; the former moves the SRDF device pair, the latter moves a single device in the SRDF device pair.

If the R2 device is moved in the half move operation, the operation requires that the SRDF R1 mirror be set Not Ready. The R1 device is suspended from the SRDF links.

The following considerations apply to moving SRDF devices between groups:

- In SRDF/Star solutions, SRDF devices cannot be moved between groups when SRDF Consistency Group (SRDF/CG) is used and enabled. The consistency groups must be disabled before the move.
- SRDF devices can be moved between SRDF groups in SRDF/S, SRDF/A and SRDF/A MSC solutions. The SRDF/A solutions require the use of the consistency exempt attribute to allow moving SRDF devices between groups when the SRDF/A session is active.

  "Consistency exempt attribute" on page 99 provides details about this attribute.

SRDF devices with multiple SRDF mirrors

The following rules apply to configuring SRDF groups in multi-site topologies where SRDF devices with two SRDF mirrors are used:

- R11 devices must be configured to two SRDF groups at the primary site. These groups are then paired with two remote SRDF groups, for example, one at the secondary and one at the tertiary site. One R1 SRDF mirror points to an R2 device at the secondary and another R1 SRDF mirror points to an R2 device at the tertiary site.

  Note: The R2 devices that remotely mirror the R11 device can reside on any Symmetrix system that participates in an SRDF solution. If the SRDF solution consists of two Symmetrix systems, both R2 devices may reside on the secondary system or one R2 device on the primary and another R2 device on the secondary system.

- R21 devices and DL R21 devices must be configured to two SRDF groups at the site where the R21 devices and DL R21 devices reside. One SRDF group is paired with an SRDF group at the primary site that contains R1 or R11 devices and another SRDF group is paired with an SRDF group at the tertiary site that contains R2 or R22 devices.
- R22 devices at the tertiary site in SRDF/Star solutions must be configured to two different SRDF groups at the tertiary site as shown in Figure 15 on page 54.



Figure 15 SRDF groups and devices with two SRDF mirrors

In both cases, the R22 devices are configured to SRDF Group 4 and to SRDF Group 7. These two groups are paired with SRDF Group 3 on the primary and SRDF Group 5 on the secondary site.

The difference between cascaded and concurrent SRDF/Star solutions is in the state of the R2 SRDF mirror:

- In cascaded SRDF/Star, the R2 SRDF mirrors that point to R1 devices in SRDF Group 3 are in a Link Blocked state. The active R2 SRDF mirror points to R21 devices in SRDF Group 5. This group is paired with SRDF Group 7 at the tertiary site that contains R22 devices.
- In concurrent SRDF/Star, the R2 SRDF mirrors that point to Group 5 are in a Link Blocked state. The active R2 SRDF mirror points to R1 devices in SRDF Group 3. This group is paired with SRDF Group 4 at the tertiary site that also contains R22 devices.

System-level SRDF device and group support

Table 13 on page 56 and Table 14 on page 56 summarize the scale of SRDF devices and SRDF groups support in Symmetrix VMAX Family, VMAXe, and DMX systems.

Table 13 Scale of SRDF support in Symmetrix VMAX 40K, VMAX 20K/VMAX, and DMX systems

Hardware platform	Symmetrix DMX			Symmetrix VMAX 40K, VMAX 20K/VMAX
Enginuity version	5771	5772	5773	5874 and higher
Device limit for an system[1]	64K	64K	64K	64K
Device limit for an SRDF director	32K	32K	64K	64K
SRDF group limit for an system	64	128	128	250
SRDF group limit for an SRDF director	8	32	32	64

1. 64K devices require a minimum of two SRDF groups.

Table 14 Scale of SRDF support in Symmetrix VMAX 10K and VMAXe systems

Enginuity version	5875.231.172 and higher
SRDF device limit per system	8K
SRDF device limit per SRDF director	8K
SRDF device limit per group	8K
SRDF group limit per system	32
SRDF group limit per SRDF director	32
SRDF port limit per system	8

SRDF links

SRDF links are logical connections between Symmetrix SRDF groups and their ports, which are physically connected by cables, routers, extenders, switches and other network devices.

The EMC SRDF Connectivity Guide provides details about supported cables and network devices.

Note: You should configure two or more SRDF links per SRDF group for redundancy and fault tolerance.

The links between a given pair of Symmetrix systems in an SRDF solution can use one of the following methods to transmit data:

- Unidirectional
- Bidirectional
- Dual-directional

The location of the R1 and R2 devices and the SRDF group configuration at both ends of the SRDF links determines the direction of the SRDF links.

Unidirectional links

If all primary devices (R1) reside in one Symmetrix system and all secondary devices (R2) reside in another Symmetrix system, write operations move in one direction, from primary to secondary devices. These are a *unidirectional* links as data moves in the same direction over all links in the SRDF group.

Figure 16 on page 57 shows a unidirectional SRDF link configuration.



Figure 16 Unidirectional links

Bidirectional links

If an SRDF group contains both primary and secondary devices, write operations move data in *both* directions over the SRDF links for that group. These links are *bidirectional*.

Figure 17 on page 58 illustrates a bidirectional link configuration over ESCON.



Figure 17 Bidirectional links

Dual-directional links

If multiple SRDF groups are used, some sending data in one direction, while others send data in the opposite direction, then the links are considered *dual-directional*. Dual-directional is often used in extended distance SRDF topologies that require data to move in two directions.

Figure 18 on page 58 shows a dual-directional SRDF link configuration.



Figure 18 Dual-directional links

Link states

SRDF links can be either online or offline.

The SRDF links are online when the following occurs:

- The Symmetrix SRDF ports are operational and enabled on both sides of the SRDF configuration.
- The external link infrastructure components are operational.

SRDF links are offline if one or more of the following occurs:

- The SRDF hardware (ports, cables, extenders, routers, and switches) is damaged or disconnected.
- The SRDF directors are offline.

- The Symmetrix systems or devices participating in the SRDF relationship are not configured properly.
- Any combination of the above.

Automatic recovery after all links fail

If all SRDF links fail, the Symmetrix system remembers the SRDF states of the affected SRDF devices and can automatically restore the devices to these states once the SRDF links become operational.

An SRDF device that is in a Ready state just before all links fail is restored to a Ready state when the SRDF links become operational. Likewise, an SRDF device that is in a Not Ready state before the failure is restored to a Not Ready state once the SRDF links become operational. This behavior is also known as pure Ready state behavior.

Automatic recovery can be toggled in the configuration file. If the Prevent Automatic Recovery After All Links Fail setting is on in the configuration file, all devices in the corresponding SRDF group remain Not Ready after all links fail. This setting allows you to perform sanity checks, for example, to inspect data on the affected devices after the SRDF links are restored and before the devices are made Ready to the host.

Prior to Enginuity version 5669, this configuration setting was applied at the Symmetrix system level. At Enginuity version 5669 and higher, this setting is available at the SRDF group level.

Note: The device-level pure Ready state behavior does not apply to the SRDF/A mode of operation.

Link states after Symmetrix power cycle

If the Force RAs Offline After Power Up configuration setting is on, the SRDF links remain offline following a Symmetrix power cycle. When set, this forces the SRDF ports offline after a power outage event and protects data in a rolling disaster scenario that can follow a power outage. The Force RAs Offline After Power Up setting is a Symmetrix system-level configuration setting, affecting all SRDF groups configured on a given Symmetrix system.

SRDF network protocols and topologies

SRDF uses Fibre Channel, Gigabit Ethernet (GigE), or ESCON network protocols to move data across the SRDF links.

Note: ESCON is only supported in SRDF for DMX solutions.

The *EMC SRDF Connectivity Guide* provide details about supported transport layers and performance considerations.

Fibre Channel

Consider the following when deploying SRDF over Fibre Channel:

- Latency — Round trip time between the primary and the secondary systems. Increased latency can be caused by host bus adapters (HBAs), switches, optics (multimode, single-mode), distance extension devices, and Fibre Channel cables.
- Bandwidth — Allows for the necessary amount of bandwidth to handle both present and future requirements.
- Flow control — Hardware-based flow control mechanisms affect distances between links. SRDF flow control allocates the number of I/Os that can be sent from the SRDF layer.
- Fast write/write acceleration — Products that can increase SRDF throughput over extended distances. These products decrease response time for SRDF/S and improve performance over long distances for adaptive copy and SRDF/AR and some SRDF/A solutions.
- Compression — May increase data throughput over the SRDF links. SRDF hardware compression and software compression are available over Fibre Channel and GigE. Table 15 on page 60 summarizes SRDF hardware and software compression support on the different platforms.

Table 15 SRDF hardware and software compression support

	DMX	VMAX 20K/VMAX	VMAX 40K	VMAX 10K/VMAXe
Hardware compression	Yes Over GigE[1] only Enginuity version 5670 or higher	Yes Over GigE[1] only Enginuity version 5874 or higher	Yes Over GigE[1] and Fibre Channel Enginuity version 5876.82.57 or higher	No
Software compression	N/A	Yes Over GigE[1] and Fibre Channel Enginuity version 5874 or higher	Yes Over GigE[1] and Fibre Channel Enginuity version 5876.82.57 or higher	Yes Over GigE[1] and Fibre Channel

1. "Gigabit Ethernet (GigE)" on page 62 provides details about SRDF hardware and software compression over GigE.

Hardware compression for SRDF traffic over Fibre Channel is available through the compression chip on the IO module carrier board. You can control hardware compression through an SRDF group-level attribute by using EMC host-based SRDF software. SRDF hardware compression also requires that the *Supports Hardware Compression* IMPL setting is set to `YES`.

You can enable/disable software compression by using EMC host-based SRDF software.

- Fibre Channel front-end port (FA port)/Fibre Channel SRDF port (RF port) conversion — Supports converting a Fibre Channel front-end port to a Fibre Channel SRDF port and vice versa. You can use EMC host-based software to change the type of the emulation running on a port from Fibre Channel to Fibre Channel SRDF or from Fibre Channel SRDF to Fibre Channel.

Note: FA/RF port conversion is not supported in SRDF/DM solutions.

Switched Fibre Channel

Enginuity versions 5x67 and higher support fully switched open SRDF (Fibre Channel) connections. Switched SRDF Fibre Channel topologies support multiple SRDF groups configured to a single SRDF port and allow greater connectivity and configuration flexibility with a given number of Symmetrix Fibre Channel ports, and are often used in concurrent SRDF and SRDF/Star solutions. Switched SRDF topologies require non-blocking switching devices that interconnect two or more Fibre Channel ports.

Figure 19 on page 61 illustrates Symmetrix A simultaneously communicating with Symmetrix B and Symmetrix C through a Fibre Channel switch and using point-to-point and point-to-multipoint Fibre Channel connectivity with Symmetrix B and Symmetrix C, respectively. Since the primary (R1) devices reside in all Symmetrix systems, the SRDF links are bidirectional or dual-directional.



Figure 19 Switched SRDF over Fibre Channel

Figure 20 on page 62 illustrates Symmetrix A communicating with Symmetrix B and Symmetrix C simultaneously through a Fibre Channel switch using point-to-point and point-to-multipoint connectivity. Symmetrix A also uses concurrent SRDF to mirror its primary devices (R11s) to a second set of secondary devices (R2s) in Symmetrix D. To support concurrent SRDF, each SRDF Fibre Channel port in Symmetrix A is configured to support two SRDF groups. A separate GigE SRDF group in Symmetrix A is configured for concurrent SRDF to Symmetrix D.



Figure 20 Switched and concurrent SRDF over Fibre Channel

Gigabit Ethernet (GigE)

1 GbE SRDF over GigE connections are available with Enginuity version 5773 or higher, and 10 GbE SRDF over GigE connections are available with Enginuity version 5875.198.148 or higher. If implementing SRDF over GigE, consider the following:

- Latency — Round trip time between the primary and the secondary systems. Latency should be as low as possible at given distances separating the sites.
- Network quality — Affects packet loss. Packets should not be sent or delivered out of order to avoid performance degradation and SRDF link instabilities.
- Link speed — Affects the amount of traffic that an SRDF director sends from each port. Speed limits are set in the Symmetrix configuration file.
- Flow control — Hardware-based flow control mechanisms affect distances between links.
- MTU size — Maximum transmission unit (MTU) should not be exceeded or packet fragmentation will occur.

- Compression — May increase data throughput over the SRDF links. Software and hardware compression are supported for SRDF traffic over GigE.

  Hardware compression for SRDF traffic over GigE is available through a compression chip:

  - On the Multi Protocol Channel Director (MPCD) and GigE director boards (Symmetrix DMX systems)
  - On the compression-capable Front End I/O Module Carriers (Symmetrix VMAX 40K, VMAX 20K/VMAX systems).

  Enginuity versions 5874 or lower require that hardware compression be enabled in the Symmetrix configuration file.

  Beginning with Enginuity 5875.135.91, you can control hardware compression through an additional SRDF group-level attribute by using EMC host-based SRDF software. SRDF hardware compression also requires that the *Supports Hardware Compression* IMPL setting is set to `YES`.

  You can enable/disable software compression by using EMC host-based SRDF software

- Encryption — To ensure the privacy of data transmitted between two points, encryption devices are available. Contact EMC support personnel for the list of qualified encryption devices.
- Security — SRDF GigE implementations support IPv4 and IPv6. IPv6 is available on the GigE/IPsec implementations beginning with Enginuity version 5773. IPsec is available with Enginuity versions 5773 and higher.

  Note: IPSec is available only with 1 Gb/s GigE SRDF ports.

SRDF over native GigE

Symmetrix systems feature native IP support. These systems can be directly attached to an IP network through their GigE SRDF ports. SRDF solutions over GigE can use connectivity provided by an existing Ethernet infrastructure and directly access high-speed data transmission conduits using IP.

Figure 21 on page 64 illustrates a switched GigE SRDF solution. Each GigE port in a primary site establishes a logical TCP connection to each GigE port in the remote Symmetrix system.



Figure 21 Switched SRDF over GigE

ESCON

Note: ESCON is an SRDF network protocol for DMX only.

Consider the following when deploying SRDF over ESCON:

- ESCON distances can be increased by using ESCON extender/repeater units or the ESCON director XDF feature.
- FarPoint is an SRDF feature used only with ESCON extended distance solutions and certain ESCON campus solutions to optimize the performance of the SRDF links. For SRDF distances beyond 15 km, or connections through channel extenders, use FarPoint to eliminate distance limitations and minimize performance penalties.
- If ESCON directors are used in SRDF solutions, the ESCON director ports must be configured as static and dedicated.

If you need to move data in both directions across the ESCON SRDF links, you should consider the following suggestions:

- If the additional latency is not as important as the number of Symmetrix ports used for SRDF operations, an ESCON bidirectional configuration is a good choice.
- If the number of ports used for SRDF operation is less important than the additional latency incurred by bidirectional operations, a dual-directional configuration is a better choice.

SRDF modes of operation

The mode of operation determines:

- How R1 devices are remotely mirrored to R2 devices across the SRDF links
- How I/Os are processed
- When the acknowledgement is returned to the production host that issued a write I/O command.

SRDF modes can be categorized as primary and secondary to indicate how they operate together. Once set, the primary mode is the default mode of operation for a given SRDF device, range of SRDF devices, or an SRDF group. Synchronous and semi-synchronous modes are the primary modes of operation while adaptive copy modes are the secondary modes of operation. Asynchronous mode is a separate category.

Synchronous

The synchronous mode maintains a real-time mirror image of data between the R1 and R2 devices. Data must be successfully stored in Symmetrix cache at both the primary and the secondary site before an acknowledgement is sent to the production host at the primary site.

Chapter 3, “SRDF/Synchronous Operations” provides more information.

Asynchronous

The asynchronous mode mirrors R1 devices by maintaining a dependent-write consistent copy of the data on the secondary (R2) site at all times. SRDF/A session data is transferred from the primary to the secondary site in cycles using delta sets. The point-in-time copy of the data at the secondary site is only slightly behind that on the primary site.

SRDF/A has little or no impact on performance at the primary site as long as the SRDF links contain sufficient bandwidth, and the secondary system is capable of accepting the data as quickly as it is being sent.

Chapter 4, “SRDF/Asynchronous Operations” provides more information.

Semi-synchronous

Semi-synchronous mode is supported with Enginuity versions prior to 5773. Semi-synchronous mode allows the R1 and R2 devices to be out of synchronization by one write I/O operation. Data must be successfully written to the primary device before an acknowledgement is sent to the production host. Semi-synchronous mode does not allow another write I/O to the R1 device until the primary system receives an acknowledgement from the secondary system that the first write operation was received in Symmetrix cache at the secondary site.

Note: In the mainframe host environments, semi-synchronous mode is not supported if the Symmetrix system uses Fibre Connection (FICON) host connectivity. In the mainframe Parallel Access Volumes/Multiple Allegiance (PAV/MA) environments, semi-synchronous mode is not recommended because it provides no performance benefit. If you want to use the increased parallelism offered by PAVs, you should run SRDF in synchronous mode.

Adaptive copy

SRDF supports adaptive copy write pending and adaptive copy disk modes as secondary modes of operation. Adaptive copy modes allow the R1 and R2 devices to be more than one I/O out of synchronization. Unlike the asynchronous mode, adaptive copy modes do not guarantee a dependent-write consistent copy of data on R2 devices.

The number of tracks out of synchronization between the R1 and the R2 devices at any given time is determined by the maximum skew value. You can set the maximum skew value by using EMC host-based SRDF software.

As secondary modes of operation, adaptive copy modes revert to the synchronous mode of operation when certain conditions are met.

Chapter 5, “Adaptive Copy Operations” provides more information.

CHAPTER 3
SRDF/Synchronous Operations

This chapter describes SRDF/S operations. Topics include:

- Enginuity emulations and I/O operations 68
- Write operations 69
- Read operations 74
- Dependent-write consistency and SRDF/CG 77
- Recovery operations 80

Enginuity emulations and I/O operations

In Symmetrix systems, the Enginuity operating system provides executables (also known as Enginuity emulations) that manage I/O execution, data protection, integrity, and data availability. Every processor in every Symmetrix director is loaded with specific Enginuity emulation code. Enginuity coordinates independent processors to act as one integrated cached disk system.

In SRDF configurations, Enginuity sends host updates across the SRDF links, for example, from R1 to R2 devices. SRDF I/O flow depends on the SRDF mode of operation.

Figure 22 on page 68 shows how Enginuity emulations participate in the SRDF I/O flow:

- *Host emulations* — Manage I/Os between the host and the Symmetrix system, moving data between the host interface cards and cache in the primary system.
- *SRDF emulations* — Manage data movement across the SRDF links between cache in the primary and cache in the secondary system.
- *Drive emulations* — Manage data movement between cache and drives within a single Symmetrix system.



Figure 22 Enginuity emulations processing I/Os in a simple SRDF environment

Enginuity emulations operate the same way in both simple SRDF configurations and in complex, multi-site SRDF configurations. Enginuity emulations use a messaging or a queue mechanism to communicate and coordinate their activities.

This chapter describes I/O operations in synchronous mode (SRDF/S). Chapter 4, "SRDF/Asynchronous Operations" describes asynchronous operations (SRDF/A).

When operating in synchronous mode, the R1 devices are mirrored to R2 devices in real-time. Real-time mirroring indicates that the Symmetrix system connected to the host continuously sends host updates across the SRDF links as part of the regular I/O flow. Synchronous mode of operation is set at the device level and requires that both R1 and R2 devices operate in synchronous mode.

Write operations

SRDF write operations require data transfer across the SRDF links. Data must be successfully written to cache at the secondary site before a positive command completion status is returned to the host that issued the write command.



Figure 23 Write I/O flow in a simple SRDF configuration

The following occurs during an SRDF write operation illustrated in Figure 23 on page 69:

1. The local host sends a write I/O to Symmetrix A.
2. The host emulations write data to cache and create a write request:
   - For the drive emulations to write (destage) data from cache to the R1 local mirror(s) in Symmetrix A
   - For the SRDF emulations to send data across the SRDF links to cache in Symmetrix B.
3. Having received a request from the host emulations, the SRDF emulations access updated data in cache, frame it according to the SRDF protocol, and transmit it across the SRDF links to Symmetrix B.
4. The SRDF emulations in Symmetrix B receive data from the SRDF links, write it to cache and return an acknowledgement to SRDF emulations in Symmetrix A. The SRDF emulations in Symmetrix B also create a request for drive emulations to write data to the R2 local mirror(s) in Symmetrix B.
5. The SRDF emulations in Symmetrix A forward the acknowledgement to host emulations, which complete the I/O flow by sending a successful command completion status to the host.

Remote mirroring

After data reaches cache in Symmetrix B, drive emulations at the secondary site destage data to the R2 local mirror(s), processing the request created from the SRDF emulations in Symmetrix B which have moved data from the links to cache in Symmetrix B.

By moving data across the SRDF links, the SRDF emulations process write I/Os to R2 devices, servicing a command to the R1 device issued by the host connected to Symmetrix A. The R2 device therefore receives data as if it were another mirror of the R1 device. That is why SRDF is known as a remote mirroring solution.

Note: R2 devices are also host-addressable logical devices while mirrors are not.

"Symmetrix logical devices, RAID groups, and mirrors" on page 28 provides more details about these concepts.

SRDF remote mirroring provides additional protection and maintains data available to the host in case of multiple drive failures at the primary site. If all R1 local mirrors fail during the host write operation, the write I/O proceeds to the R2 device by way of the SRDF links. Figure 24 on page 70 illustrates an example[1]. After the host updates arrive in cache in Symmetrix B, they are written to R2 local mirrors while the corresponding R2 SRDF mirror tracks are marked invalid, indicating to the Symmetrix system that the R1 device in Symmetrix A contains old data which cannot be used by the host.

If no new host write I/Os to the R1 device have arrived since the failure, SRDF automatically synchronizes R1 local mirrors with data from the R2 local mirrors after the R1 local mirrors become available. This is because the R2 SRDF tracks were marked invalid.



Figure 24 Write to R2 if R1 local mirror fails

1. The local mirror is presented as a RAID 1 group but can be any RAID group supported in Symmetrix systems.

Write I/Os to R21 devices

Write I/Os to dual-role R21 devices in cascaded SRDF configurations consist of the following:

- Synchronous I/O flow between Symmetrix A and Symmetrix B as shown in Figure 25 on page 71
- Asynchronous or adaptive copy I/O flow between the synchronous remote site and the tertiary site. Chapter 4, "SRDF/Asynchronous Operations" and Chapter 5, "Adaptive Copy Operations" provide more details.

The synchronous I/O flow is the same as that described in "Write operations" on page 69 except that the R1 device and the R2 SRDF mirror of the R21 device operate in synchronous mode while the R1 SRDF mirror of the R21 device and the R2 device in Symmetrix C operate in either asynchronous or adaptive copy mode.

When a write I/O arrives to cache in Symmetrix B, the SRDF emulation in Symmetrix B sends a positive status across the SRDF links to Symmetrix A (synchronous operations) and also creates a request for SRDF emulations at Symmetrix B to send data across the SRDF links to Symmetrix C. Figure 25 on page 71 illustrates the synchronous I/O flow in a cascaded SRDF topology.



Figure 25 Write I/Os to R21 devices

In cascaded configurations that use DL R21 devices in Symmetrix B, host I/Os can sustain multiple hardware failures in Symmetrix A as shown in Figure 26 on page 72.

If the R1 local mirror in Symmetrix A fails during a host write operation, DL R21 devices in Symmetrix B sends data to R2 devices in Symmetrix C. The failure in Symmetrix A is reflected as follows:

- The R1 device's R1 SRDF is valid, indicating that the R2 device in Symmetrix C contains good data.
- The DL R21 device's R2 SRDF mirror is invalid and its R1 SRDF mirror is valid, indicating that R2 local mirrors in Symmetrix C contain current data that can be used by the host.
- The R2 device in Symmetrix C has its R2 SRDF mirror invalid, indicating that the R1 device in Symmetrix A contains bad data.

After the R1 local mirrors become available, they can receive data from R2 devices in Symmetrix C, provided that no new host I/Os have been issued to the R1 device since the drive failure in Symmetrix A.



Figure 26 Three-site configuration with R21 devices

The following are important differences between synchronous I/O flow to R21 and DL R21 devices:

- DL R21 devices store only the differences between the R1 devices in Symmetrix A and the R2 devices in Symmetrix C.
- The R1 SRDF mirror of the R21 device operates in adaptive copy disk mode while the R1 SRDF mirror of the DL R21 device operates in adaptive copy write pending mode between Symmetrix B and Symmetrix C.
- R21 devices are host addressable; DL R21 devices are not.

Write I/Os and compression

With Enginuity versions 5874 and higher, Enginuity software compression can be applied to SRDF traffic over Fibre Channel and GigE SRDF links. If software compression is enabled, Enginuity compresses data before sending it across the SRDF links.

Symmetrix systems at both sides of the SRDF links must support software compression and must have the software compression feature enabled in the configuration file. You can set software compression by using EMC host-based SRDF software.

Hardware compression is available for GigE SRDF traffic on VMAX Family systems running Enginuity version 5874 or higher and on DMX systems running Enginuity version 5669 or higher.

Hardware compression is available for Fibre Channel SRDF traffic on Symmetrix VMAX 40K systems running Enginuity version 5876.82.57 or higher.

Both Enginuity software and hardware compression can be activated simultaneously for SRDF traffic over GigE and Fibre Channel. Data is first compressed by software and then further compressed by Symmetrix hardware.

Read operations

In ordinary situations, a read operation from the R1 device does not require any involvement of the SRDF emulations.

If the production host issues a read command to the R1 device and data is in cache in Symmetrix A, the host emulation reads data from cache and sends it to the host. If data is not in cache (read miss), the drive emulation first reads the requested data from local drives to cache. In both cases, the read I/O operation does not require the involvement of the SRDF emulation.

In the example shown in Figure 27 on page 74, both hosts are attached to Symmetrix A and Symmetrix B and can read from the R2 device, provided that the R2 device is in Read Only (Write Disabled)[2] state. Such a read operation is still an ordinary read I/O and does not require any action from the SRDF emulations. However, you should proceed with caution when issuing read commands to the R2 device in the course of regular SRDF/S operations.

"Read I/Os to R2 devices" on page 76 provides more information about these precautions.



Figure 27 Read I/O flow from the primary device

2. "SRDF devices" on page 31 provides more details about R1 and R2 device states.

Read I/O if R1 local mirror fails

SRDF devices can process read I/Os that cannot be performed by regular Symmetrix logical devices. If the R1 local mirror fails, the R1 device can still service the request as long as its SRDF state is Ready and the R2 device has good data.

In this case, the SRDF emulations help service the host read request. Since the R1 local mirror is not available, the SRDF emulations bring data from the R2 device to Symmetrix A. The host perceives this as an ordinary read from the R1 device, although the data was read from the R2 device acting as if it was a local mirror. This is because the host writes to the R1 device and the Symmetrix system automatically propagates data to the R2 device. The host sees the R2 device either as Read Only or Not Ready and therefore cannot issue a direct write command to the R2 device as long as the existing SRDF R1-R2 pair relation exists.

SRDF emulations may also assist a read I/O to improve overall host I/O performance. For example, the host may issue a read command to the R1 device while the R1 local mirror undergoes a RAID rebuild operation. If both Symmetrix systems are close enough for the SRDF links to maintain a low latency, the SRDF emulations need less time to bring data from Symmetrix B than the drive emulations need to rebuild the RAID group that belongs to the R1 local mirror. Figure 28 on page 75 shows an example of a read request being serviced from the R2 device while a RAID rebuild operation takes place on the R1 device.



Figure 28 Read I/O if R1 local mirror fails

Read I/Os to R2 devices

The SRDF configurations in Figure 27 on page 74 shows the production and the remote hosts attached to both Symmetrix systems. The R2 devices can be Read Only to any host attached to Symmetrix B.

However, allowing the remote host to read data from the R2 devices is not recommended, because:

- The SRDF/S mode relies upon the application's ability to determine if the data image is the most current. Symmetrix B may not yet know that data currently in transmission on the SRDF links has been sent.
- If the remote host reads data from the R2 device while a write I/O is in transmission on the SRDF links, the host will not be reading the most current data.

EMC strongly recommends that you allow the remote host to read data from the R2 devices while in Read Only mode only when:

- Related applications on the production host are stopped.
- The SRDF writes to the R2 devices are blocked due to a temporary suspension/split of the SRDF relationship.

Note: Some host operating systems do not allow reading data from a device that is in Read Only mode.

Read I/Os to DL R21 devices

DL R21 devices are not visible to the host. Therefore, no direct host read operations can be issued to DL R21 devices.

DL R21 devices can be made read/write capable to the host through EMC host-based SRDF control software, even though they are not host addressable. As soon as R21 or DL R21 devices are made read/write capable, the primary R1 devices automatically transition into the SRDF Not Ready/R1 SRDF mirror Not Ready state.

Read I/Os to R21 devices

Ordinary dual-role R21 devices are host addressable. Despite being dual-role devices, they resemble R2 devices. During regular SRDF operations, the R21 devices are in either Read Only (Write Disabled) or Not Ready state and behave as described in "Read I/Os to R2 devices" on page 76.

Dependent-write consistency and SRDF/CG

Note: With Enginuity versions 5874 or higher, the SRDF/CG option does not require a separate license, it is offered with the SRDF/S license.

In a remotely mirrored environment, data consistency cannot be ensured if one of the writes is remotely mirrored, but its predecessor was not. This could occur if a communication loss affects only a subset of the devices involved in the remote mirroring function.

SRDF/Consistency Group (SRDF/CG) is an SRDF feature designed to ensure the dependent-write consistency of the data distributed across multiple R1 devices. The R1 and R2 devices can be distributed across multiple primary and secondary systems, respectively.

The basic building block of the SRDF/CG product are consistency groups. A consistency group is a set of SRDF devices that may reside on multiple Symmetrix systems and are enabled for database consistency. SRDF devices that belong to the same consistency group act in unison to preserve dependent-write consistency of a database distributed across multiple devices within the consistency group. The consistency group ensures that remote mirroring is suspended for all SRDF devices in a consistency group as soon as one SRDF device in the group fails to send data across the SRDF links.

Dependent-write operations

SRDF/CG is based on dependent-write operations. A dependent-write is a write operation that cannot be issued by an application until a prior, related write I/O operation is completed. An example of a dependent-write is a database update:

1. The DBMS writes to the transaction log.
2. The DBMS writes the data to the actual database.
3. The DBMS writes again to the transaction log.

SRDF/CG prevents a rolling disaster from affecting data integrity at the secondary site. When SRDF/CG detects any write I/O to a device that cannot communicate with its R2 (secondary) device, SRDF/CG suspends the remote mirroring for all devices defined to the consistency group before completing the intercepted I/O and returning control to the application. In this way, SRDF/CG prevents a dependent-write I/O from reaching the secondary site if the previous I/O only gets as far as the primary site.

If SRDF/CG is enabled, the I/Os to the primary devices in the consistency group can still occur even when the R1 devices have their SRDF mirrors Not Ready. Such updates are not immediately sent to the secondary site. However, they are propagated after the affected links are again operational, and data transfer from the R1 devices to the R2 devices resumes.

Ensuring data consistency

Assume an SRDF configuration consists of three Symmetrix systems with R1 devices, and two additional Symmetrix systems with R2 devices. The systems with R1 devices send data to the systems with R2 devices as shown in Figure 29 on page 78.



Figure 29 Primary and secondary systems

Next, assume that the links between primary system 2 and secondary system 1 fail. The primary systems 1 and 3 continue to write data to the secondary site systems 1 and 2 while the primary system 2 does not, as shown in Figure 30 on page 78.

As a result, the copy of the data on the secondary systems 1 and 2 becomes inconsistent.



Figure 30 Failed links between Primary 2 and Secondary 1

Figure 31 on page 79 illustrates an example where the links between primary system 2 and secondary system 1 fail. The consistency group automatically stops primary systems 1 and 3 from sending data to secondary systems 1 and 2, as shown in Figure 32 on page 79. Thus, the dependent-write consistency of the data spanning secondary systems 1 and 2 remains intact.



Figure 31 Primary 1, 2, and 3 in a consistency group



Figure 32 Failed link between Primary 2 and Secondary 2

The SRDF/CG product requires that Symmetrix devices in consistency groups act in unison to maintain data integrity. EMC host-based SRDF software provides the tools used to configure and manage these devices.

The following documents provide more information about consistency groups and consistency group utilities:

- *EMC Consistency Group for z/OS Product Guide*
- EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide.

Recovery operations

This section summarizes recovery operations in two-site SRDF configurations operating in synchronous mode.

Although only a few specific examples are discussed in this section, the flexibility of the SRDF product allows you to meet their specific needs. Many possibilities exist to tailor a remote protection/disaster recovery solution to best fit a particular business objective. For details on how to tailor the SRDF environment to suit specific business needs, contact an EMC Business Continuance Applications consultant.

SRDF recovery operations fall into these categories:

- Planned failover (for testing, maintenance)

  A planned failover is a controlled failover operation to test the robustness of the disaster restart solution. You should always allow for the time and resources required to test their disaster restart solution in a controlled environment to be able to protect their business when a disaster actually happens. Production is temporarily moved to the secondary site during a planned failover.

- Unplanned Failover

  Production processing is moved from the primary to the secondary site due to an unexpected failure of the production host at the primary site, the primary system, or both. The secondary site temporarily becomes the primary/production site.

- Failback

  Once the primary site and/or the production host are restored or repaired, the production processing is resumed at the primary site.

The following sections summarize basic recovery operations in a simple, two-site SRDF solution.

The *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide and the EMC SRDF Host Component for z/OS Product Guide contain additional information about the failover and failback operations.*

Planned failover to the secondary system

A planned failover due to maintenance, upgrade or testing the recovery solution involves moving production applications from Symmetrix A to Symmetrix B, as shown in Figure 33 on page 81. A planned failover involves the following steps:

1. Shut down all applications on the production host.
2. Take all SRDF links between Symmetrix A and Symmetrix B offline to suspend remote mirroring.
3. If SRDF/CG is used, disable consistency groups between Symmetrix A and Symmetrix B.
4. Swap personalities between R1 and R2 devices.
5. At this point, the new R1 devices reside in Symmetrix B and the new R2 devices reside in Symmetrix A. Since the SRDF pairs operated in synchronous mode before the failover process, they are ready to resume production operations in Symmetrix B.

6. Enable consistency groups between Symmetrix B and Symmetrix A.
7. Bring all SRDF links between Symmetrix B and Symmetrix A online to resume remote mirroring.
8. Start production applications from the failover host attached to Symmetrix B.

Figure 33 on page 81 shows an example of a the before and after states of a planned failover using R1/R2 personality swap.



Figure 33 Planned failover operation in a two-site configuration (SRDF/S)

Once the maintenance, upgrades or testing procedures that required a planned failover to Symmetrix B complete, you can repeat the same procedure to return production to Symmetrix A. By moving production between sites using R1/R2 personality swap rather than initiating a failback operation, you can avoid downtime associated with the ordinary failback process.

"Failback to the primary system" on page 82 provides more information about the failback procedures.

Unplanned failover

An unplanned failover involves moving production applications from Symmetrix A to Symmetrix B after an unanticipated outage at the primary site. In this case, the primary site is not available. You have to initiate the failover process that involves changing the R2 device states to Read/Write to the host attached to Symmetrix B.

The following steps summarize an unplanned failover in a two-site configuration operating in synchronous mode:

1. The production host and Symmetrix A become unavailable.
2. Take all SRDF links between Symmetrix A and Symmetrix B offline to suspend remote mirroring.

3. Change the R2 device states to Read/Write to the failover host connected to Symmetrix B.
4. Start applications on the failover host and resume production to write-enabled R2 devices in Symmetrix B.

Once the failover host begins writing to the R2 devices, updates made to the R2 devices accumulate on the R2 SRDF mirrors as invalid tracks owed to the R1 devices. Processing continues on the failover host until the production host and Symmetrix A are operational again. Figure 33 on page 81 shows an example of a failover to the secondary site after the primary site fails.



Figure 34 Failover to Symmetrix B, Symmetrix A and production host unavailable.

Failover to the secondary site in a simple configuration shown in Figure 34 on page 82 can be performed in minutes. You can resume production processing as soon as the applications are restarted on the failover host connected to Symmetrix B.

Unlike the planned failover operation described in "Planned failover to the secondary system" on page 80 which resumes remote mirroring as soon as the SRDF links are online, the failover process shown in Figure 33 on page 81 resumes production at the secondary site, but without remote mirroring until Symmetrix A and the production host become operational and ready for a failback operation.

Failback to the primary system

After the primary host and Symmetrix system containing the primary (R1) devices are again operational, production processing can resume on the primary host. The following steps are required to transfer processing from Symmetrix B to Symmetrix A.

1. Stop I/Os on the failover host connected to Symmetrix B.
2. Make all R2 devices in Symmetrix B Not Ready or Read Only (Write Disabled) to the host.

3. If Symmetrix A was powered off, ensure that SRDF links between Symmetrix A and Symmetrix B are disabled before powering on Symmetrix A.
4. If Symmetrix A running Enginuity 5874 or earlier was powered off and you do not want to discard its changed data, disconnect or disable the SRDF links between Symmetrix A and Symmetrix B before powering on Symmetrix A. In this way, you will prevent changed data from Symmetrix B (secondary) from moving automatically to Symmetrix A (primary).
5. Power on Symmetrix A and make R1 devices Read/Write enabled to the production host.
6. Enable the SRDF links between Symmetrix A and Symmetrix B.
7. Bring the SRDF links online and restart the local host. The R1 devices automatically receive data from the R2 devices which accumulated invalid tracks on their R2 SRDF mirrors during production processing.
8. Once all SRDF pairs are synchronized, enable consistency groups on the SRDF links between Symmetrix A and Symmetrix B.
9. Restart the production host and applications.

Recovery for a large number of invalid tracks

If the R2 devices have handled production processing for a long period of time, there might be a large number of invalid tracks (TBs in size) *owed* to the R1 devices. You can use EMC host-based SRDF control software to resynchronize the primary and secondary devices while the secondary host continues production processing. Once there is a relatively small number of invalid tracks owed to the R1 devices (for example, 50 GB), the failback process can be initiated.

CHAPTER 4
SRDF/Asynchronous Operations

This chapter describes SRDF/A operations. Topics include:

- Overview 86
- Single session mode 91
- Recovery scenarios 100
- Multi Session Consistency mode 103
- SRDF/A and cache utilization 111

Overview

SRDF/A provides a long distance disaster restart solution with fast application response time. This solution is intended for users who must preserve dependent-write consistency within and across the database and application environments at an extended-distance secondary (R2) site and with minimal host application impact.

SRDF/A accumulates host write I/Os in delta sets that reside in the primary and the secondary system's cache, manages dependent-write consistency of each delta set, and transfers delta sets across the SRDF links in cycles.

SRDF/A benefits

SRDF/A provides the following benefits:

- Supports extended data mirroring with database and application consistency
- Features efficient link utilization that results in lower link bandwidth requirements
- Provides a host write response performance close to that observed with non-SRDF devices.
- Maintains a dependent-write consistent point-in-time image on the secondary (R2) devices.
- Supports all current SRDF topologies (point-to-point and switched fabric Fibre Channel, GigE, ESCON).
- Allows failover and failback operations between the primary and the secondary sites.

Single session mode

SRDF/A single session mode refers to implementations where individual SRDF/A sessions, whether in the same or multiple Symmetrix systems, have their cycle switch controlled independently by Enginuity rather than by EMC host-based SRDF control software.

"Single session mode" on page 91 provides additional details.

Multi Session Consistency (MSC) mode

SRDF/A Multi Session Consistency (MSC) mode refers to implementations where multiple primary SRDF groups and multiple secondary SRDF groups require consistency, and are thereby included in the same SRDF/A MSC session. In this environment, the SRDF/A cycle switch is driven by EMC host-based SRDF software to maintain consistency between multiple SRDF/A sessions operating in the SRDF/A single session mode.

"Multi Session Consistency mode" on page 103 provides additional details.

SRDF/A can be controlled in open systems and mainframe host environments through the EMC host-based SRDF software. The *EMC Solution Enabler Symmetrix SRDF Family CLI Product Guide* and the *EMC SRDF Host Component for z/OS Product Guide* provide more details.

The following sections describe the SRDF/A features that guarantee such high levels of data consistency, disaster-restart protection, and low host I/O response time.

Reserve Capacity

Enginuity versions 5772 and higher support the SRDF/A Reserve Capacity feature that keeps the SRDF/A session operational in the event of temporary network resource shortfalls that would have previously suspended the SRDF/A operations. The two functions that implement SRDF/A Reserve Capacity are Transmit Idle and DSE.

Transmit Idle

Transmit Idle keeps the SRDF/A session active after all the SRDF network links have failed. It allows SRDF/A to remain fully active during network outages that cause an *All Links Lost* condition. When the links return, SRDF/A resumes data transfer. SRDF/A Transmit Idle requires the following:

- Both Symmetrix systems must be running Enginuity version 5671 or higher.
- This feature must be enabled on both Symmetrix systems. It is enabled by default with Enginuity version 5772 or higher.

The *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide* and the *EMC SRDF Host Component for z/OS Product Guide* provide additional information on how to use the Transmit Idle feature in SRDF/A sessions.

Note: SRDF/A Transmit Idle is not supported in SRDF ESCON topologies.

Delta Set Extension

SRDF/A DSE keeps the SRDF/A session active if the Symmetrix system approaches the SRDF/A maximum cache utilization limits. DSE pages out (offloads) some or all of the delta set data into preconfigured storage pools known as DSE pools that are associated with SRDF/A groups.

You have to configure a separate DSE pool for each device emulation type (FBA, IBM i, CKD3380 or CKD3390) and can create multiple DSE pools for different SRDF/A groups.

EMC recommends that you configure DSE pools and enable DSE on the primary and on the secondary system. When TimeFinder/Snap sessions are used to replicate either R1 or R2 devices, you must create two separate preconfigured storage pools: DSE and Snap pools.

Consult your EMC representative for details about planning and monitoring preconfigured storage pools.

Write folding

Write folding improves the efficiency of the SRDF network links. If multiple updates to the same location arrive in the same cycle, the SRDF emulations send the most current data across the SRDF links. This feature makes SRDF/A superior to competitive asynchronous mirroring solutions as it decreases network bandwidth consumption and the number of I/Os processed by the SRDF emulations.

Write pacing

Write pacing is an SRDF/A feature that balances cache utilization by extending the host write I/O response time to prevent SRDF/A operational interruptions. The write pacing feature provides the following options:

- The group-level pacing option. This option requires Enginuity version 5874 or higher on both Symmetrix systems. It is enabled for the entire SRDF/A group when slowdowns in host I/O rates, transmit cycle rates, or apply cycle rates occur.
- The device-level pacing option. This option is for SRDF/A solutions in which the SRDF/A R2 devices participate in TimeFinder copy sessions. This option requires Enginuity version 5875.135.91 or higher on both Symmetrix systems.

You can enable or disable each write pacing option. All write pacing options are compatible with each other and with other SRDF/A features including tunable cache utilization, Reserve Capacity, and MSC. With Enginuity version 5876.82.57 or higher and EMC host-based SRDF software, you can view the global write pacing statistics report.

The key benefit of the SRDF/A write pacing is its dynamic, self-paced mechanism. Once enabled, write pacing is employed only when required and applies the appropriate delay to the host write I/O response time to keep the SRDF/A session active. When write pacing is enabled, but pacing is not required to keep the SRDF/A session running, host write I/Os are not paced. The delay applied to the overall host write I/O response time is known as the write pacing delay.

With Enginuity version 5875.135.91 or higher, the write pacing delay can by default extend the host write I/O response time up to 50 milliseconds. You can overwrite the default value with a user-specified maximum write pacing delay that can extend the host I/O response time up to one second. The maximum write pacing delay applies to all write pacing options.

"SRDF/A and cache utilization" on page 111 provides more information about the write pacing feature and how to use it with other SRDF/A features.

Tolerance mode

Tolerance mode is an SRDF/A feature that allows you to balance the performance and data consistency requirements. Setting Tolerance mode ON allows one or more devices to be Not Ready on the link at the R1 side and the SRDF/A session remain active. R2 consistency is not maintained, but this may be acceptable in certain controlled service activities. If all links are lost, Tolerance mode does not keep SRDF/A active. You can use the Tolerance mode option to complete hardware service procedures like device replacement without disabling or dropping SRDF/A.

The EMC host-based SRDF control software implementations of tolerance mode are different for mainframe and open systems environments. The mainframe software defines Tolerance mode as `Tolerance mode` by directly exporting the Enginuity tolerance feature, while the open systems software implements it through `Consistency` enabling/disabling options. The *EMC Symmetrix SRDF Host Component for z/OS Product Guide* and the *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide* provide more details.

By default, EMC host-based SRDF control software does not allow an SRDF/A session participating in MSC operations to have tolerance mode set.

Tolerance mode allows certain conditions that would otherwise interrupt SRDF/A operations, for example:

- If you make some of the R2 devices in the SRDF/A session Read/Write enabled.
- If some of the R1 SRDF mirrors in the SRDF/A session become Not Ready.

Note: When tolerance mode is set, dependent-write consistency is *not* guaranteed.

Development and availability

Table 16 on page 89 lists the highlights of development and availability for SRDF/A for Symmetrix VMAX 40K, VMAX 20K/VMAX, and DMX series.

Most of these SRDF/A features are available in SRDF for VMAX 10K or VMAXe solutions when Enginuity version 5875.231.172 was released, except for MSC and some write pacing enhancements. MSC for Symmetrix VMAX 10K or VMAXe is supported with Enginuity version 5876.82.57 or higher.

Table 16 Summary of SRDF/A development and availability for Symmetrix VMAX 40K, VMAX 20K/VMAX, and DMX series (page 1 of 2)

Enginuity version	SRDF/A feature or enhancement	Feature/enhancement description
5670	SRDF/A	Allowed a single static SRDF group to operate in SRDF/A mode within a single Symmetrix system. This SRDF/A group could not participate in concurrent SRDF operations.
5670.50	Multi Session Consistency (MSC)	Introduced SRDF/A MSC for mainframe host environments. Support was limited to a single static SRDF/A group per Symmetrix system, but multiple Symmetrix systems could now participate in an SRDF/A MSC solution.
5671	Multiple SRDF/A SRDF groups per Symmetrix	Multiple SRDF/A groups (up to 64, depending on the configuration) allowed per Symmetrix system. This feature allowed a primary and secondary SRDF/A group within the same Symmetrix system and therefore the ability to use SRDF/A with bidirectional operations[1].
	SRDF/A MSC for open systems and mainframe	SRDF/A MSC for mainframe and open systems supported multiple Symmetrix systems with no restrictions to the number of SRDF/A groups per Symmetrix system operating in SRDF/A MSC mode.
	Concurrent SRDF/S with SRDF/A	"Concurrent SRDF/S with SRDF/A" on page 123 provides details.
	Dynamic SRDF support	This feature enabled SRDF/A on dynamic SRDF groups.
	Tunable cache utilization	"Tunable cache utilization" on page 111 provides details.
5772	Reserve Capacity	"Reserve Capacity" on page 87 provides details.
5773	Transfer Log	"Transfer Log" on page 101 provides details.

Table 16 Summary of SRDF/A development and availability for Symmetrix VMAX 40K, VMAX 20K/VMAX, and DMX series (page 2 of 2)

Enginuity version	SRDF/A feature or enhancement	Feature/enhancement description
5874	Write pacing - group-level pacing	Group-level pacing is supported in the following SRDF topologies: • Two-site SRDF topologies • Concurrent SRDF topologies (including concurrent SRDF/Star) • The first hop (R1 -> R21) of cascaded SRDF topologies (including cascaded SRDF/Star) "Write pacing" on page 111 provides details.
5875.135.91	Write pacing - device-level pacing	Device-level pacing is supported in the following SRDF topologies: • Two-site SRDF topologies • Concurrent SRDF topologies (including concurrent SRDF/Star) • The first hop (R1 -> R21) of cascaded SRDF topologies (including cascaded SRDF/Star) "Write pacing" on page 111 provides details.
	Concurrent SRDF/A	"Concurrent SRDF/A" on page 123 provides details.
5876.82.57	Write pacing - group-level pacing enhancements	Group-level pacing option is improved to apply write pacing delays on individual SRDF/A R1 devices or a given SRDF/A group to match slowdowns in R2 apply cycle rates.
5876.159.102	Write pacing enhancements	Write pacing is improved to support group-level and device-level pacing options on the second hop (R21 -> R2) of a cascaded SRDF topology. This feature is also supported in cascaded SRDF/Star and SRDF/EDP. Enginuity version 5876.159.102 is required on the Symmetrix system containing the R21(or DL R21) devices.

1. Bidirectional operation within a single SRDF/A group is not supported. All R1 to R2 device operations within a single SRDF/A group still remained unidirectional. "SRDF links" on page 57 contains additional information.

Single session mode

SRDF/A single session mode refers to implementations where individual SRDF/A sessions, whether in the same or multiple Symmetrix VMAX systems, have their cycle switch controlled independently by Enginuity rather than by EMC host-based SRDF control software.

Unlike traditional ordered write asynchronous approaches, Symmetrix systems implement asynchronous mode host writes from the primary system to the secondary system by using dependent-write consistent delta sets:

- Enginuity groups host write I/Os into delta sets and manages dependent-write consistency for each delta set.
- Enginuity transfers these delta sets by using cycles of operation, one cycle at a time, between the primary system and the secondary system.

With Enginuity version 5x71 and higher, multiple SRDF groups can be enabled for SRDF/A operations within the same Symmetrix system. However, if multiple groups within the same Symmetrix system are enabled for SRDF/A, these sessions operate without knowledge of each other, or of SRDF/A sessions in other Symmetrix systems.

To manage multiple SRDF/A sessions that need to be consistent, a multi session mode is required and EMC host-based SRDF software to performs this multi session management. "Multi Session Consistency mode" on page 103 provides additional information.

Dependent-write consistency

Dependent-write consistency is achieved through the processing of ordered SRDF/A delta sets (cycles) between the primary system and the secondary system:

- The active cycle on the primary system contains the current host writes or N data updates in the capture delta set.
- The inactive cycle contains the N-1 data set that is transferred from the primary system to the secondary system.
- The primary inactive cycle is the transmit delta set and the secondary system inactive cycle is the receive delta set.
- The active cycle on the secondary system contains the N-2 data update in the apply delta set. This is the guaranteed dependent-write consistent image in the event of a disaster or failure.

Figure 35 on page 92 illustrates the delta sets and their relationships.



Figure 35 SRDF/A delta sets and their relationships

The Symmetrix SRDF emulation[3] ensures dependent-write consistency within SRDF/A by obtaining the active cycle number from a single location in cache and assigning that number to each I/O. The SRDF emulation retains that cycle number even if a cycle switch occurs during the life of that I/O. This results in an atomic cycle switch process for dependent-write sequences, even though it is not physically an atomic event across a range of devices. As a result, two I/Os with a dependent relationship between them can be in the same cycle, or the dependent I/O can be in a subsequent cycle.

Delta set switching

Delta set switching is at the core of the SRDF/A active session operation. The following conditions must be met before the cycle switch can take place:

- On the primary system, write I/Os are collected in the capture delta set (N copy of the data), and the minimum cycle time is reached (Step 1 in Figure 36 on page 93).

  The minimum cycle time is a user-specified period, a minimum time that has to elapse between consecutive SRDF/A cycle switches. The actual time between consecutive SRDF/A cycle switches may be longer if all conditions for a cycle switch are not met within the minimum cycle time.

- The previous cycle's transmit delta set (N-1 copy of the data) must have completed data transfer to the receive delta set on the secondary system (Step 2 in Figure 36 on page 93).

- On the secondary system, the previous apply delta set (N-2 copy of the data) is written to cache, and data is marked write pending for the R2 devices.

3. An Enginuity executable responsible for SRDF operations. "Enginuity emulations and I/O operations" on page 68 provides more information.

Table 17 on page 93 and Figure 36 on page 93 present what happens during the delta set switching process.

Table 17 Delta set switching process

On primary system	On secondary system
1. The primary system sends a commit message to the secondary system to begin the secondary system cycle switch (Step 2a in Figure 36 on page 93). 2. Enginuity halts the SRDF transfer (Step 2b in Figure 36 on page 93). 3. Enginuity automatically switches between the capture and the transmit delta set on the primary system (Step 2c in Figure 36 on page 93). 4. The new capture delta set on the primary system is now available to collect new host I/Os (Step 2d Figure 36 on page 93). 5. Once the primary system receives the acknowledgement of the secondary system cycle switch (Step 4 in Figure 36 on page 93), the primary system initiates the SRDF transfer of the transmit delta set (Step 4a in Figure 36 on page 93).	1. The secondary system receives the commit message from the primary system (Step 3 in Figure 36 on page 93), and verifies the apply delta set has been written (Step 3a in Figure 36 on page 93). 2. The secondary system performs a delta set cycle switch between the receive and apply delta sets (Step 3b in Figure 36 on page 93). This preserves the dependent-write consistent copy at the secondary system prior to receiving the next dependent-write consistent copy (Step 3c in Figure 36 on page 93). 3. The secondary system sends an acknowledgement to the primary system (Steps 3d in Figure 36 on page 93). 4. The restore of the apply delta set begins (Step 3e in Figure 36 on page 93). Data in the apply delta set is marked write pending for the R2 devices.



Figure 36 Delta set switching

Delta set switching in concurrent SRDF/A

Cycle switching in concurrent SRDF/A environments has unique operational characteristics because each R1 SRDF mirror of the R11 device belongs to a different SRDF/A group and therefore different SRDF/A session with independent cycle switching. In single session mode, cycle switching on both legs of the concurrent SRDF topology typically occurs at different times. Consequently, each R1 SRDF mirror may hold a different copy of data for the same track at different times.

Host I/Os to an R11 device therefore require special I/O handling because the incoming I/Os are accepted by two R1 SRDF mirrors that operate as if they were separate Symmetrix logical devices.

The following procedure summarizes the I/O flow to the R11 device with both R1 SRDF mirrors operating in SRDF/A mode:

1. A host write command arrives for a track on the R11 device. Enginuity intercepts the command and associates a cache slot to its data. Then, Enginuity adds the newly allocated cache slot to the current cycle of each SRDF/A session.

   Note: Each SRDF/A session counts the cache slot as shared, but part of different SRDF/A session.

2. Suppose a cycle switch happens in one SRDF/A session, for one R1 SRDF mirror, and a new write command to the same track arrives. Enginuity intercepts the write command for the session that switched cycles, duplicates the cache slot, inserts it into the current cycle for that session as an exclusive slot.

   Note: The original shared cache slot still exists and has moved to a different cycle for one and remained in the same cycle for another SRDF/A session. In addition, there is an exclusive slot that belongs to the current cycle of a single R1 SRDF mirror. At this point, two cache slots are associated with the same track.

3. Suppose a cycle switch occurs in the second SRDF/A session and the host again writes to the same slot. Enginuity intercepts this command and duplicates the cache slot the same way as described in step 2 , just for the second SRDF/A session. The slot that was duplicated moves to another cycle for the session that switched cycles and remains in the same cycle for the session that did not switch cycles.

   At this point, three cache slots exist for the same track, two exclusive slots, one for each SRDF/A session and a shared standard slot.

The above example illustrates an extreme case when three cache slots exist for the same track at the time. This can happen when the host I/O profile consists of many sequential write commands to the R11 device. Consider adequate system cache capacity when configuring R11 devices.

Single session states

SRDF/A single session mode can have the following states: Not Ready, Inactive, and Active. Figure 37 on page 95 shows SRDF/A session state transitions.



Figure 37 SRDF/A single session states

Not Ready state

When the SRDF environment is configured, and the SRDF links become active and enabled, by default all SRDF devices are in the Not Ready[4] state. This means that the R1 devices have the R1 SRDF mirror Not Ready.

"SRDF device states" on page 46 provides more details about SRDF states.

EMC host-based SRDF software can be used to make SRDF devices Ready. Once the SRDF devices are in the Ready state, the session transitions to the Inactive state.

Inactive state

The SRDF devices are Ready while the SRDF/A session is Inactive. Host commands are required to transition SRDF/A into an Active state.

Enginuity can also place SRDF single sessions into the Inactive state, if certain conditions are met.

Active state

The Active state is the normal SRDF/A operating mode. The secondary system is either consistent or inconsistent.

The secondary system is consistent once all of the previously owed tracks from the R1 devices have been transferred to the R2 devices. Specifically, when the last cycle containing this data is fully copied to cache and in the N-2 cycle (apply delta set) on the secondary system and has successfully been applied. The consistent active state always represents a dependent-write consistent image of the data. The inconsistent active state means that previously owed tracks from the R1 devices are yet to be transferred to the R2 devices.

Host-based SRDF control commands maintain a dependent-write consistent copy on the secondary system, with the exception of the *deactivate* command.

4. This is the SRDF view of the device state.

Single session state transitions

While the state transition sequence from Not Ready or Inactive to an Active state is common when setting up SRDF/A, transitions from an Active state to a Not Ready or an Inactive state can happen in response to certain conditions or host software commands. The following sections provide some examples of these state transitions.

Note: If the session is in an Active state with an inconsistent secondary system and still transferring accumulated updates, it is impossible to create a dependent-write consistent image on the secondary system by transitioning out of SRDF/A mode.

Drop

When an SRDF/A session is dropped, the R1 devices participating in the SRDF/A session are placed in the Not Ready state. The SRDF/A session transitions from the Active to the Not Ready state.

This state transition can be triggered by:

- Enginuity
  - The primary system write pending limit is reached or the SRDF/A maximum cache value is exceeded
  - If any R1 device has R1 SRDF mirror Not Ready
  - All links are lost and the Transmit Idle function is not enabled
  - If resource constraint conditions exceeding user-defined operating parameters occur
- EMC host-based SRDF software commands (`symrdf suspend -immediate`).

Data not yet sent from the capture and transfer delta sets is marked as owed to the secondary system. At Enginuity versions lower than 5773, data so far received in a partially completed receive delta is marked as owed to the primary system. At Enginuity version 5773 or higher, the Transfer Log function changes this behavior. "Transfer Log" on page 101 contains further information.

This allows for a consistent image of data to be re-synchronized from either system, in either direction, according to user requirements.

EMC recommends capturing the resulting dependent-write consistent data with a set of TimeFinder local replicas prior to any resynchronization. During the resynchronization activity, the dependent-write consistent image at the secondary system is compromised. TimeFinder copies (gold copies) of write-dependent data preserve dependent-write consistency on the secondary system.

Note: Tolerance mode may change the operating behaviors described above. "Tolerance mode" on page 88 provides additional information.

Pend-drop

Another way of dropping an active SRDF/A session and moving it into a Not Ready state is to place the SRDF devices in a Not Ready state only at the end of the current cycle. This transition can be triggered by Enginuity or the EMC host-based SRDF software (`symrdf suspend`). On the primary system, write pending tracks in the active cycle are converted to tracks owed to the secondary system.

By dropping SRDF/A on the cycle boundary by using the pend-drop mechanism, you do not have to resolve owed tracks from the secondary to the primary system when the SRDF/A session resumes. "Transfer Log" on page 101 contains additional information.

Note: In multi session environments, a single session drop often results in the corresponding MSC session drop.

Deactivate

SRDF/A allows an immediate session transition from Active to Inactive while leaving the SRDF devices Ready. Because the devices are left Ready, data continues to flow and the dependent-write consistency of the data at the secondary system is compromised. The capture and transmit delta set's data are marked as owed tracks to the secondary system similar to a resynchronization operation. These owed tracks are not dependent-write consistent.

Switching to SRDF/A mode

You can initiate the session mode switching from any SRDF mode to SRDF/A mode through EMC host-based SRDF software.

SRDF/S to SRDF/A

Before switching from SRDF/S to SRDF/A mode, SRDF verifies that all SRDF devices are Ready, then moves SRDF groups on both sides of the SRDF links to the SRDF/A Active state. As a result, the delta sets are established on both Symmetrix systems and the SRDF/A mechanism is enabled.

If all devices are synchronized when switching from SRDF/S to SRDF/A mode, the secondary system shows a consistent state and the data is dependent-write consistent. If some data is still owed to the secondary site, there is no dependent-write consistency at the secondary system until the last owed track is in the N-2 cycle (apply delta set).

Note: SRDF/A is an SRDF group-level feature; all devices assigned to an SRDF group set to operate in asynchronous mode will operate in asynchronous mode when the SRDF/A state becomes Active.

Adaptive copy write -pending to SRDF/A

When the mode is set to SRDF/A from adaptive copy write-pending mode, all devices are moved into the adaptive Copy Pending Off mode. An SRDF/A session is activated and the drive emulations scan devices for write-pending tracks and add them to the active cycle (capture delta set). Once all write-pending data to the R2 device has been written to cache in the secondary system, the R1 device can transition out of the adaptive Copy Pending Off

mode. When all devices transition out of the adaptive Copy Pending Off mode, two cycle switches are required for the secondary system to report a consistent state and to have its data dependent-write consistent.

Adaptive copy disk mode to SRDF/A

When a host requests a mode switch from adaptive copy disk mode to SRDF/A mode, the transition is immediate. Tracks owed to the secondary system as a result of adaptive copy disk skew are scheduled as resynchronization operations. These are copy I/Os scheduled by the drive emulation to be serviced by SRDF/A.

Each SRDF/A session is allowed to process an equal number of copy tracks per cycle, up to the total of 30,000 copy tracks for all sessions. This mechanism avoids using all cache in the primary system for resynchronization purposes. For example, if five active SRDF/A sessions are active, each SRDF/A session is allowed a maximum of 6,000 copy tracks per cycle (new capture delta set). DSE further limits the number of copy tracks to 10 percent of the 30,000 copy tracks (3,000) that are allowed and have to be equally distributed among all active SRDF/A sessions. For example, if 10 SRDF/A sessions are active while cycle data is paged out to the DSE pool, the maximum number of copy tracks per session is 300.

This is because it is assumed that Symmetrix cache resources are under stress if DSE is paging out. In this case, accomplishing re-synchronization (copy operations) is not as important as keeping SRDF/A sessions from dropping due to reaching the maximum cache usage limits.

Host I/Os continue to be serviced in the current SRDF/A cycles (capture delta set). The length of time to send the tracks owed with asynchronous mode depends on the number of outstanding tracks owed prior to switching to asynchronous mode.

For example, if there is one active SRDF/A session in the Symmetrix system with 90,000 tracks owed to the secondary system, it will take a minimum of three SRDF/A cycle switches to transmit data.

Two additional cycle switches are required to ensure that the data is in the apply delta set cycle (N-2 copy of data). SRDF/A produces a consistent state on the secondary system and a dependent-write consistent copy of data after all resynchronization operations are complete and the two additional cycle switches have occurred.

Switching from SRDF/A to SRDF/S mode

With Enginuity version 5x71 and higher, it is possible to transition to the SRDF/S mode from SRDF/A without losing dependent-write consistency. This is only allowed for SRDF/A single session mode. If you are running SRDF/A in MSC mode, you must first disable MSC, which will leave all SRDF/A sessions in the MSC group running in SRDF/A single session mode.

The following caveats apply:

- The transition is not immediate. Once a transition is requested, it may take some time for SRDF/A mode to transition to SRDF/S mode
- Some performance degradation occurs with synchronous mode while the transition takes place
- The transition requires both Symmetrix systems to be running Enginuity version 5x71 and higher.

Note: Prior to switching session modes, you should capture dependent-write consistent copies to local TimeFinder replicas on both the primary and the secondary system.

Consistency exempt attribute

The consistency exempt attribute is provided for an SRDF/A environment to indicate that the device should be considered exempt from the consistency requirements for the session. This allows dynamic expansion of devices participating in an SRDF/A session without taking the SRDF/A group offline or declaring the entire SRDF/A group inconsistent.

In normal operation, SRDF/A provides a consistent copy of the data on the R2 devices. However, if there are tracks that need to be copied as part of an initial synchronization operation, the data is not consistent until those tracks have been copied to the remote Symmetrix system. Also, when new SRDF device pairs are created, a full synchronization is required. As a result, adding new SRDF devices to an existing SRDF/A group causes the group to appear inconsistent.

To solve this problem, the consistency exempt option is available to allow a new SRDF device pair to be excluded from the consistency check until the tracks have been copied to the secondary system. When the new SRDF pairs are added to the SRDF/A group in consistency exempt mode, the consistency exempt attribute has to be cleared before those devices can be used or considered part of the consistent image.

Recovery scenarios

This section briefly discusses the different recovery scenarios associated with SRDF/A single session mode.

Temporary all links lost

If SRDF/A suffers a temporary loss (‹ 10 seconds by default) of all SRDF links, the SRDF/A state remains active and data continues to accumulate in global memory. This may result in an elongated cycle, but the secondary system dependent-write consistency is not compromised and the primary and secondary system device relationships are not suspended. If you want to keep SRDF/A in an active state during all links lost conditions, SRDF/A Transmit Idle should be used. "Transmit Idle" on page 87 provides more information.

Note: Switching to SRDF/S mode with the link limbo parameter configured for more than 10 seconds could result in an application, database, or host failure if SRDF is restarted in synchronous or semi-synchronous mode.

Permanent all links lost

If SRDF/A experiences a permanent all SRDF links lost condition, all of the devices configured to the SRDF group where the devices were configured are set to a Not Ready state. All data in capture and transmit delta sets is changed from write pending for the R1 SRDF mirror to invalid for the R1 SRDF mirror and is therefore owed to the R2 device. Any new write I/Os to the R1 device are also marked invalid for the R1 SRDF mirror.

All of these tracks are owed to the secondary system once the links are restored.

With Enginuity versions prior to 5773, data in the receive delta set in the secondary system is also marked owed to the R1 devices. These tracks are owed to the primary system. The active cycle (apply delta set) data completes its commit to the R2 devices.

Note: Enginuity versions 5773 or higher with the Transfer Log feature change this behavior. "Transfer Log" on page 101 contains additional information.

When the links are restored, normal SRDF recovery procedures are followed. The metadata representing the data owed is compared and merged based on normal host recovery procedures used by EMC host-based SRDF software. Data is then resynchronized by sending the owed tracks as part of the SRDF/A cycles.

Data on non-consistency exempt devices on the secondary system is always dependent-write consistent in SRDF/A active/consistent state, even when all SRDF links fail. However, the act of starting a resynchronization process compromises the dependent-write consistency until the resynchronization is fully complete and two cycle switches have occurred. For this reason, it is important to use TimeFinder to create a gold copy of the dependent-write consistent image on the secondary system.

Session cleanup process

Once the SRDF/A single session mode is dropped, Enginuity automatically starts a cleanup process:

1. The primary system marks new incoming writes as being owed to the secondary system.
2. The capture and transmit delta sets are discarded, but the data is marked as being owed to the secondary system. All of these owed tracks are sent to the secondary system once SRDF is resumed, as long as the copy direction remains primary to secondary.
3. The secondary system marks and discards the receive delta set only. Data is marked as tracks owed to the primary system. With Enginuity versions prior to 5773, these tracks are scheduled to be sent from the primary system once the SRDF links are restored, as long as the copy direction has not changed.

   Note: The SRDF/A Transfer Log improves this cleanup operation. The following sections contain information related to the SRDF/A Transfer Log.

4. The secondary system makes sure the apply (N-2) delta set is safely applied to disk; this is the dependent-write consistent image.

Note: It is very important to capture a gold copy of the dependent-write consistent data on the secondary system R2 devices prior to any resynchronization. Any resync process compromises the dependent-write consistent image. The gold copy can be stored on a remote set of BCVs or Clones.

Transfer Log

With Enginuity level 5773 and higher, the SRDF/A Transfer Log is provided to manage any needed re-synchronization after an SRDF/A session drop. It allows SRDF/A to know which tracks in the transmit delta set have been successfully transferred to the receive delta set. The overall effect of this function is a significant increase in the speed of the SRDF/A recovery operations. It reduces the time required to recover from an unplanned SRDF/A session drop.

This function requires Enginuity version 5773 or higher at both the primary and secondary sites. The Transfer Log removes the need to manage invalid track information on the secondary site after an unplanned SRDF/A session drop. To perform the Transfer Log recovery, proceed as follows:

1. Make the R1 devices Ready.
2. Activate SRDF/A.

Failback from secondary devices

If a disaster occurs on the primary system, data on the R2 devices represents a dependent-write consistent image and can be used to restart the applications.

After the primary system has been repaired, you can return production operations to the primary system by following procedures described in "Recovery operations" on page 80.

If the failover to the secondary site requires an extended event, the SRDF/A solution can be reversed by using either dynamic SRDF features or a configuration change. SRDF/A can continue operations until a planned reversal of direction can be performed to restore the original SRDF/A primary and secondary relationship.

After the workload has been transferred back to the primary system hosts, SRDF/A can be activated to resume normal asynchronous mode protection.

Multi Session Consistency mode

Mainframe software and Enginuity version 5670.50 and higher support SRDF/A control for multiple Symmetrix systems with a single SRDF group per Symmetrix system. Beginning with Enginuity version 5x71 for mainframe and open systems, SRDF/A is supported in topologies with multiple primary systems and multiple SRDF groups on both the primary and the secondary system.

This feature is referred to as SRDF/A MSC. SRDF/A MSC solutions can also support mixed open systems and mainframe data controlled within the same SRDF/A MSC session.

Achieving data consistency across multiple SRDF/A groups requires that the cycle switch process described in "Delta set switching" on page 92 be coordinated among the participating Symmetrix system SRDF groups or systems, and that the switch occurs during a very brief time period when no host writes are being serviced by any participating Symmetrix system.

"MSC mode delta set switching" on page 106 provides more details.

Entering SRDF/A multi session consistency

SRDF/A MSC requires a single coordination point to drive the cycle switch process in all participating Symmetrix systems. This function is provided by the EMC host-based SRDF control software.

Figure 38 on page 103 shows the SRDF/A MSC state diagram.



Figure 38 SRDF/A MSC allowed state transitions

For the host to control the cycle switch process, the Symmetrix systems must be aware that they are running in MSC mode. EMC host-based SRDF control software performs the following:

- Coordinates the cycle switching for all SRDF/A sessions comprising the SRDF/A MSC solution
- Monitors for any failure to propagate data to the secondary system devices and drops all SRDF/A sessions together to maintain dependent-write consistency
- Performs MSC cleanup operations (if possible).

Note: Simply activating SRDF/A does not place a session in MSC mode. Conversely, exiting MSC mode does not drop or deactivate SRDF/A, it merely places SRDF/A in single session mode. However, if SRDF/A is dropped or deactivated, then MSC mode (along with all SRDF/A single session modes comprising the MSC session) is necessarily terminated and would need to be reentered once SRDF/A was made active again.

As part of the process to enter MSC mode, and with each cycle switch issued thereafter, Enginuity assigns a cycle tag to each new capture cycle. That cycle tag is retained throughout that cycle's life. This cycle tag is a value that is common across all participating SRDF/A sessions and eliminates the need to synchronize the cycle numbers across them. The cycle tag is the mechanism by which dependent-write consistency is assured.

MSC mode dependent-write consistency

I/Os are processed exactly the same way in SRDF/A MSC mode as they are in single session mode:

- The active cycle on the primary system contains the current host writes or N data version in the capture delta set.
- The inactive cycle contains the N-1 data version that is transferred using SRDF/A from the primary system to the secondary system. The primary inactive delta set is the transmit delta set and the secondary system's inactive delta set is the receive delta set.
- The active cycle on the secondary system contains the N-2 data version on the apply delta set. This is the guaranteed dependent-write consistent image in the event of a disaster or failure.

Figure 39 on page 105 illustrates the delta sets and their relationships.



Figure 39 SRDF/A MSC delta sets and their relationships

Performing an SRDF/A MSC consistent cycle switch

SRDF/A MSC mode performs a coordinated cycle switch during a very short window of time when there are no host writes being completed. This time period is referred to as an SRDF/A window. SRDF/A window is an SRDF/A group attribute. It is checked at the start of each I/O. This process involves the following:

1. When the EMC host-based SRDF control software discovers that all the SRDF groups and Symmetrix systems are ready for a cycle switch, it issues a single command to each SRDF group that performs a cycle switch and opens the SRDF/A window.
2. The host emulation obtains the cycle number at the start of each write as it does in single session mode. In multi session mode, the host emulation also checks the SRDF/A window setting. If the setting is on (an open window), the host emulation disconnects upon receiving host write I/O and begins polling to determine when the EMC host-based SRDF control software has closed the window. While the SRDF/A window is open no write I/Os can be issued to R1 devices that participate in an SRDF/A MSC session.
3. The SRDF/A window remains open on each SRDF group and Symmetrix system until the last SRDF group and Symmetrix system in the multi session group acknowledges to the host software that the **open** and **switch** command has been processed. At this point the host software issues a **close** command for each SRDF/A group under MSC control. As a result, dependent-write consistency across the SRDF/A MSC session is ensured.

Note: Enginuity does provide a fail-safe mechanism to ensure that the SRDF/A MSC window does not remain open permanently due to a host software failure. Enginuity closes the SRDF/A MSC window itself if the host software has not closed it within 15 seconds.

MSC mode delta set switching

This section describes how the delta set switching process works for SRDF/A MSC mode. Figure 40 on page 107 and Figure 41 on page 109 show three SRDF/A single sessions combined together to create a single SRDF/A MSC group. There are two primary systems, one with a single SRDF/A group and the other with two SRDF/A groups. The secondary systems have the same configuration as the primary systems have. This is a balanced configuration. Figure 40 on page 107 and Figure 41 on page 109 assume SRDF/A MSC has been activated and two cycle switches have previously occurred.

Before a primary system cycle switch can occur, the following must be true:

- The primary system's transmit delta set (DS) must be empty
- The secondary system's apply delta set must have completed. The N-2 data must be marked write pending for the R2 devices.

Figure 40 on page 107 shows how the current host I/O being collected by the capture delta set on the primary system. The primary system transmit delta set is continuing to send the data to the secondary system receive delta set. The apply delta set is continuing to restore or mark the data write pending to the R2 devices in the secondary systems.



Figure 40 SRDF/A MSC capture delta set collects application write I/O

Figure 41 on page 109 shows that SRDF transfer between the primary system transmit delta sets and the secondary system receive delta sets is complete (Step 2 in Figure 41 on page 109). The primary Symmetrix sends the secondary Symmetrix a transmit complete message (Step 2a in Figure 41 on page 109).

The process continues as follows:

1. The primary systems wait for an acknowledgement from the secondary Symmetrix (Step 2b in Figure 41 on page 109) systems. The primary systems halt the SRDF transfer (Step 2c Figure 41 on page 109).
2. The secondary system apply delta sets complete the restore process by marking the data write pending to the R2 devices (Step 3 in Figure 41 on page 109). When completed, the secondary systems send a restore complete message to the primary systems (Step 3a in Figure 41 on page 109).

3. Once the primary systems receive the restore complete message from the secondary systems, the primary systems respond to polls from the SRDF/A MSC host software with a ready to switch condition (Step 4 Figure 41 on page 109). Again this is because the primary system transmit delta sets are empty and the secondary system apply delta sets have completed the restore process.
4. The SRDF/A MSC host software initiates the primary system cycle switch once all of the participating SRDF groups in the SRDF/A MSC solution report a ready to switch state (Step 5 in Figure 41 on page 109).
5. The write I/Os are deferred long enough for the host software to coordinate the cycle switch across all SRDF groups and the primary systems (Step 5a in Figure 41 on page 109).
6. The primary systems allow the deferred I/O, and a new capture delta sets accept the host I/O (Step 5b in Figure 41 on page 109). The transmit delta sets contain the N-1 copy of dependent-write consistent data.
7. The primary systems send a commit message to the secondary systems (Step 6 in Figure 41 on page 109) once the primary systems cycle switch occurs (Step 5a in Figure 41 on page 109).
8. After receiving the commit message, the secondary systems perform a cycle switch between the receive and apply delta sets (Step 6a in Figure 41 on page 109).
9. The secondary systems now have a new receive delta set available (Step 6b in Figure 41 on page 109).
10. The SRDF transfer process begins from the primary systems to the secondary systems (Step 6c in Figure 41 on page 109). The secondary systems also begin the apply delta set restore process (Step 6d in Figure 41 on page 109).

The delta set switching process completes and the new cycle begins as shown in Step 1 in Figure 40 on page 107.



Figure 41 SRDF/A MSC delta set switching process

MSC session cleanup process

When SRDF/A is deactivated or dropped while in MSC mode, each primary system starts the same cleanup process as in single session mode - it discards all I/O from both the transmit and capture delta sets and marks the corresponding tracks owed to the secondary system as described in "Session cleanup process" on page 100.

The host software does not need to perform any special recovery on the primary system. The process can be summarized as follows:

1. Enginuity at the secondary (R2) Symmetrix system completes the restore of its apply delta set automatically.
2. For each SRDF group, Enginuity discards any receive delta sets that are not complete. If the receive delta set is a complete delta set for each SRDF group, Enginuity marks it *needing cleanup* in cache, awaiting a decision from the host software.

The SRDF/A MSC uses the host cycle tags described in "Entering SRDF/A multi session consistency" on page 103. The host software must use its cycle tags during recovery of the receive delta sets on the secondary system.

The following scenarios must be considered when SRDF/A has been terminated with respect to the receive delta sets on the secondary system:

- All receive delta sets on all secondary systems and SRDF groups have the same tag and are marked needing cleanup. This scenario is the result of the secondary system receiving and acknowledging the transmit complete message in the SRDF/A MSC cycle switch process. The default behavior of the host software commits all of the receive delta sets to deliver the most current dependent-write consistent data to the secondary systems.
- All receive delta sets on all Symmetrix systems have the same tag number, but at least one Symmetrix systems or SRDF/A group does not have a receive delta set marked needing cleanup. This scenario means that Enginuity discarded an incomplete receive delta set.

  The host software *must* discard all receive delta sets for this tag number. The most current consistent image of data is already on the secondary systems in the apply delta sets. The data that was in the discarded receive delta sets are marked as tracks owed to the primary systems.
- Different cycle tags exist within the apply and receive delta sets. In this case, the secondary system SRDF/A sessions can be divided into two groups. The first group has Symmetrix systems with apply delta set cycle tags that match the receive delta set cycle tags of the Symmetrix systems from the second group. In other words, Symmetrix systems from the first group have received the commit message for a certain host cycle, while the Symmetrix systems from the second group have not.

  In this case, the receive cycles of the Symmetrix systems from the second group are *necessarily* complete and the host software *must* force the restore of the apply delta sets. At the same time, host software must discard the receive delta sets of the Symmetrix systems from the first group regardless of their completeness.

SRDF/A and cache utilization

The SRDF/A solutions tend to experience high cache utilization levels. They require:

- Large amount of cache resources on the primary and on the secondary system to handle the SRDF/A cycles and to absorb peaks in host I/O activity.
- A balanced SRDF/A configuration as the primary system, the SRDF/A links, and the secondary system have to work in unison to handle workloads at different host I/O activity and SRDF link throughput levels.

You have several tools, including tunable cache utilization, Reserve Capacity, and write pacing, to manage cache utilization levels and avoid Symmetrix system cache exhaustion. The following paragraphs describe these features and how they interact with one another.

Tunable cache utilization

This feature allows you to set a SRDF/A maximum cache utilization threshold to a certain percentage of the system write pending limit for the following SRDF/A sesstions:

- An individual SRDF/A session in single session mode
- Multiple SRDF/A sessions in single or MSC mode.

You can assign priorities to sessions, keeping SRDF/A active for as long as cache resources allow.

When the SRDF/A maximum cache utilization threshold or the system write pending limit is exceeded, the Symmetrix system exhausts its cache.

By default, the SRDF/A session drops if Symmetrix system cache is exhausted. You can modify this behavior to keep the SRDF/A session running for a user-defined period of time once Symmetrix cache is exhausted. If this condition does not resolve itself at the expiration of the user-defined period, the SRDF/A session drops.

Reserve Capacity

The SRDF/A Reserve Capacity functionality includes SRDF/A DSE and Transmit Idle. DSE pages out (offloads) some or all of the delta set data into preconfigured storage pools known as DSE pools when the system approaches the maximum SRDF/A cache utilization threshold. The Transmit Idle feature keeps the SRDF/A session running when all SRDF/A links are lost, for example, due to the temporary network outages. The SRDF/A session can therefore remain active during short-term network interruptions.

"Reserve Capacity" on page 87 provides details about the Enginuity versions that support the SRDF/A Reserve Capacity.

Write pacing

The SRDF/A write pacing is a dynamic technique that may extend host write I/O response time to prevent cache exhaustion in the primary or in the secondary system. The delay applied to overall host write I/O response time is also known as the write pacing delay. The write pacing feature provides the following options.

- Group-level pacing option — Activated for the entire SRDF/A group when increase in host I/O rates or transmit and/or apply cycle times occur.

- Device-level pacing option — Activated for R1 devices in an SRDF/A group whose R2 partners on the secondary system participates in TimeFinder operations (without -precopy option).

All options are compatible with each other and with other SRDF/A features including tunable cache utilization, Reserve Capacity, and MSC. EMC host-based SRDF software allows you to enable/disable each write pacing option.

Group-level pacing

The SRDF/A group-level pacing option detects imbalances that can lead to the maximum SRDF/A cache utilization threshold either in the primary or in the secondary system. When spikes in host I/O rates occur, or slowdowns make transmit or apply cycle times longer, the group-level pacing option takes a corrective action by extending the host write I/O response time to match slower SRDF/A service rates.

The group-level pacing option may extend the host write I/O response time for a given SRDF/A group to balance the incoming host I/O rates with the SRDF link bandwidth and throughput capabilities when:

- The host I/O rate exceeds the SRDF link throughput.
- Some SRDF links that belong to the SRDF/A group are lost.
- SRDF links experience reduced overall throughput.

Prior to Enginuity version 5876.82.57, group-level pacing was only responsible for managing the host I/O rates so that they would not overrun the data transmittal between R1 and R2 on the SRDF links. Enginuity version 5876.82.57 extends this capability to respond to the following conditions where asynchronous replication is slowed down on the secondary site:

- The write-pending level on an R2 device in an active SRDF/A session reaches the device write-pending limit.
- The apply cycle time on the R2 side is longer than 30 seconds and the R1 capture cycle time (or in MSC, the capture cycle target).

As a result, the enhanced group-level pacing feature can effectively pace host write I/Os in the following operational scenarios:

- Slower apply cycle times on specific R2 devices managed by slower-speed physical drives.
- FAST operations that lead to an imbalance in SRDF/A operations between sites.
- Sparing operations that lead to back-end directors on the R2 side becoming slower in overall restore operations.
- Production I/Os to the R2 side that leads to back-end directors and/or SRDF directors becoming slower in restore operations.
- Restore delays during the pre-copy phase of TimeFinder/Clone sessions before activation.

Note: The enhanced group-level pacing functionality requires the primary and secondary systems run Enginuity version 5876.82.57 or higher. Otherwise, group-level pacing is available without the enhanced functionality.

If group-level pacing is activated, the host write I/Os are paced as follows:

- If you do not specify the maximum write pacing delay, the group-level pacing feature extends the host write I/O response time to match the speed of either the SRDF links or the apply operation on the R2 side, whichever is slower. By default, the group-level pacing feature cannot add more than 50 milliseconds to the overall host response time. The group-level pacing feature always applies the lowest possible write pacing delay required to keep the SRDF/A session active. If cache utilization conditions require a write pacing delay greater than 50 milliseconds, SRDF/A session might drop if SRDF/A Max Cache Usage value is exceed.
- If you specify the maximum write pacing delay, the group-level pacing feature extends the host write I/O response time by using the appropriate write pacing delay to keep the SRDF/A session running, but no greater than the user-specified maximum write pacing delay.

  For example, if you set the write pacing delay at 10 milliseconds while cache utilization conditions can be resolved by using a write pacing delay of 5 milliseconds, the group-level pacing feature extends the overall host write response time by 5 milliseconds. However, if cache utilization conditions require a write pacing delay greater than 10 milliseconds, the SRDF/A session might drop if SRDF/A Max Cache Usage value is exceed.

Device-level (or TimeFinder) pacing

The device-level pacing option, also known as TimeFinder pacing, applies a write pacing delay for individual SRDF/A R1 devices whose R2 counterparts participate in TimeFinder copy sessions. The write pacing delay is used to avoid high SRDF/A cache utilization levels when the R2 devices servicing both the SRDF/A and TimeFinder copy requests experience slowdowns in service rates.

Like the group-level pacing option, the device-level pacing option applies the appropriate write pacing delay when the SRDF/A cache utilization levels require that the host write I/Os be paced. Unlike the group-level pacing option, the device-level pacing option applies the write pacing delay for individual R1 devices, but only if the R2 partners of these R1 devices participate in TimeFinder operations. This way, the device-level pacing option indirectly paces the SRDF/A cycle requests issued to the appropriate R2 devices.

If device-level pacing is activated, the host write I/Os are paced as follows:

- If you do not specify the maximum write pacing delay, the device-level pacing feature always applies the appropriate write pacing delay to the overall host write response time to keep the SRDF/A session active. By default, the maximum write pacing delay is 50 milliseconds.
- If you set the maximum write pacing delay, the device-level pacing feature applies the appropriate write pacing delay to keep the SRDF/A session active up to the user-defined maximum write pacing delay.

If the device-level pacing option does not help resolve the high levels of cache utilization, the SRDF/A session might drop and activating the group-level pacing option is necessary to keep the session up.

Figure 42 on page 114 shows how the device-level pacing affects individual devices.



Figure 42 SRDF/A R2 devices as source devices in TimeFinder operations

"TimeFinder and SRDF/A operations" on page 164 provides additional details about support for TimeFinder operations in the SRDF/A sessions.

Write pacing and Transmit Idle

When write pacing is activated and Transmit Idle is enabled, if cache usage conditions require write pacing, the behavior depends on when all SRDF links are lost:

- If cache usage conditions require write pacing when all SRDF links are lost and, consequently, Transmit Idle is in effect, the host write I/Os continue to be paced.
- If all SRDF/A links are lost, Transmit Idle is in effect, and cache usage conditions require that the group-level pacing option be activated, the host write I/Os are not paced until at least one SRDF link recovers. The SRDF/A device-level pacing feature might not take effect if all SRDF/A links are lost.

Write pacing and DSE

If write pacing and DSE are activated, the write pacing threshold is linked to the ratio of delta set tracks paged out to the DSE pools to the total number of delta set tracks in use by the SRDF/A in a given Symmetrix system. The default threshold value is 95 percent. When threshold values are reached, the write pacing feature takes effect if the host I/O rate exceeds the DSE page out rate.

Write pacing in cascaded SRDF

Prior to Enginuity version 5876.159.102, device-level pacing options are only supported on the first hop (R1 -> R21) of a cascaded SRDF topology. With Enginuity version 5876.159.102 or higher, device-level pacing can also be activated on the second hop (R21 -> R2) of a cascaded SRDF topology. This feature also applies to cascaded SRDF/Star and SRDF/EDP topologies. The following requirements applies:

- The Symmetrix system containing the R21 devices is running Enginuity version 5876.159.102 or higher.
- The corresponding R1 -> R21 device pair is not RW on the SRDF link and is in synchronous mode. Otherwise, the R21 devices cannot be paced.

CHAPTER 5
Adaptive Copy Operations

This chapter describes adaptive copy operations, SRDF/Data Mobility (SRDF/DM) and SRDF/Automatic Replication (SRDF/AR). Topics include:

- Adaptive copy modes 116
- SRDF/Data Mobility 117
- SRDF/Automated Replication 118

Adaptive copy modes

The SRDF adaptive copy modes are adaptive copy write pending and adaptive copy disk mode. These modes can effectively move large amounts of data with minimal host impact, for example, to synchronize the new SRDF device pairs or to migrate data to a newer Symmetrix hardware model using SRDF.

The SRDF adaptive copy modes have a fast application response time because they return a positive acknowledgment to the host as soon as data is written to cache in the primary system.

Note: The SRDF adaptive copy modes do not provide restartable data images at the secondary site and are not designed for disaster restart solutions.

The SRDF adaptive copy modes allow the R1 and R2 devices to be more than one write I/O and up to the maximum skew value out of synchronization. The maximum skew value is the maximum number of the updated R1 device tracks that have not ye been transferred to the R2 device. By default, the maximum skew value is set to 65535 tracks. Once the maximum skew value is reached, the SRDF emulations must start the synchronization process to transfer updates across the SRDF links from the R1 to the R2 devices.

You can set the maximum skew value at the device level by using the EMC host-based SRDF software. Adaptive copy is supported as the secondary mode of operation to synchronous or semi-synchronous mode. The R1 devices revert to the primary mode of operation when the maximum skew value is reached and remain in primary mode until the number of tracks out of synchronization is lower than the maximum skew value.

Adaptive copy write pending

In *SRDF adaptive copy write pending* mode, write requests accumulate in cache *in the primary system.* A background process writes to the corresponding R2 device on the other side of the SRDF links.

Note: SRDF adaptive copy write-pending mode reverts to the specified primary mode if 75% of the write pending limit for the Symmetrix system is reached, regardless of whether the maximum skew value specified for each device is reached.

Adaptive copy disk

In *SRDF adaptive copy disk* mode, the write requests accumulate on the *R1 device* rather than in Symmetrix cache. A background process sends the outstanding write requests to the corresponding R2 device.

With Enginuity version 5773 or higher, the background copy process scheduled to send I/Os from the R1 to the R2 devices can be deferred if:

- The write requests exceed the maximum R2 write pending limits.
- The write requests exceed 50 percent of the secondary system write pending space.

SRDF/Data Mobility

SRDF/Data Mobility (SRDF/DM) is a licensed two-site SRDF data migration and replication solution that only operates in *adaptive copy* modes. The maximum skew value set at the device level in SRDF/DM solutions must be equal or greater than 100 tracks.

Figure 43 on page 117 illustrates an SRDF/DM topology and the I/O flow in adaptive copy mode:

1. The host write I/O is received in cache in Symmetrix A.
2. The host emulation returns a positive acknowledgment to the host.

   In adaptive copy write pending mode, data stays in cache until the SRDF emulation moves it across the SRDF links.

   In adaptive copy disk mode, the drive emulation writes data from cache to disk and schedules a request for the SRDF emulation to move data across the SRDF links. When data is scheduled to be sent across the SRDF links, the drive emulation, if needed, writes data back to cache.

   "Enginuity emulations and I/O operations" on page 68 provides more details about Enginuity emulations.
3. The SRDF emulation transmits the I/O across the SRDF links to Symmetrix B.
4. Once data is written to cache in Symmetrix B, the SRDF emulation in Symmetrix B returns a positive acknowledgement to Symmetrix A.



Figure 43 SRDF/DM adaptive copy mode

Benefits

The benefits of SRDF/DM are:

- It supports all Symmetrix systems and all Enginuity versions that support SRDF.
- Compared to SRDF/S, it has a faster response time.

Limitations

The limitations of the SRDF/DM are:

- It is used for data replication or migration only, not for disaster restart solutions.
- In adaptive copy write pending mode, cache is temporarily consumed until data is transferred across the SRDF links. In adaptive copy disk mode, propagation delays may occur because data may be read from the drives to cache before it is transmitted across the SRDF links.

SRDF/Automated Replication

The SRDF/Automated Replication (SRDF/AR) solution provides a long-distance disaster restart solution. SRDF/AR can operate as follows:

- In two-site topologies that use SRDF/DM in combination with TimeFinder.
- In three-site topologies that use a combination of SRDF/S, SRDF/DM, and TimeFinder. These solutions operate in synchronous mode between Symmetrix A and Symmetrix B and in adaptive copy mode between Symmetrix B and Symmetrix C.

Note: Multi-hop SRDF/AR requires Enginuity version 5876.159.102 or higher on the VMAX 10K or VMAXe systems that participate in the SRDF relationships.

"SRDF and TimeFinder" on page 160 provides more details about SRDF integration with TimeFinder.

Figure 44 on page 118 shows a single-hop SRDF/AR topology. The R1 device is also the TimeFinder target device which is replicated across the SRDF links. TimeFinder thus separates the remote replication and production processing at the production site in Symmetrix A. The R2 device holds an asynchronous restartable copy of the R1 data. The R2 device can be configured as a TimeFinder source device to enable independent asynchronous processing in Symmetrix B.



Figure 44 SRDF/AR single-hop data flow

Figure 45 on page 119 shows a multi-hop SRDF/AR topology. This disaster restart solution operates in synchronous mode between Symmetrix A and Symmetrix B and in adaptive copy mode between Symmetrix B and Symmetrix C. The three-site solution provides a remote mirror copy in Symmetrix B to guarantee disaster restart with zero data loss in the event of the primary site failure and an asynchronous restartable copy in Symmetrix C. TimeFinder is required in Symmetrix B to separate synchronous remote mirroring from remote replication between Symmetrix B and Symmetrix C.



Figure 45 SRDF/AR multi-hop data flow

SRDF/AR is a native Symmetrix solution that does not require a Symmetrix controlling host attached to Symmetrix B because its design fully supports automation from the controlling host connected to Symmetrix A. The controlling host attached to Symmetrix A can issue SRDF and TimeFinder commands to R1 and R2 devices. In addition, the SRDF/CG option and the TimeFinder/CG option guarantee consistent images of data across multiple devices and Symmetrix systems. "Dependent-write consistency and SRDF/CG" on page 77 and "SRDF and TimeFinder" on page 160 provide more information.

Note: With Enginuity versions 5874 or higher, the SRDF/CG option does not require a separate license. The SRDF/CG functionality is available with the SRDF/S license. Similarly, the TimeFinder/CG option is available with TimeFinder/Clone and TimeFinder/Snap license.

Benefits

The benefits of SRDF/AR are:

- The multi-hop solution provides a disaster restart copy in Symmetrix B and a remote copy in Symmetrix C.
- You can achieve dependent-write consistency in SRDF/AR topologies by using SRDF/CG and TimeFinder/CGs.
- SRDF/AR provides a lower cost solution for remote data protection by utilizing lower-cost communication links.
- With Enginuity 5875.135.91, SRDF /AR supports thin devices.

Limitations

In a three-site SRDF/AR multi-hop solution as shown in Figure 45 on page 119, SRDF/S host I/O to Symmetrix A is not acknowledged until Symmetrix B has acknowledged it. This can cause a delay in host response time.

CHAPTER 6
SRDF Multisite Solutions

This chapter describes the SRDF multisite solutions. Topics include:

- Concurrent SRDF 122
- Cascaded SRDF 126
- Extended Distance Protection 128
- SRDF/Star 130
- Four-site SRDF solution for open systems host environment 139
- SRDF/SQAR 141

Concurrent SRDF

Concurrent SRDF is a three-site disaster recovery solution using concurrent primary devices (R11) that are configured with two R1 SRDF mirrors. The two R1 mirrors allow the R11 devices to be mirrored concurrently to two R2 devices that can reside in one or two remote Symmetrix systems. Each R1 SRDF mirror is configured to a different SRDF group.

Note: The R11 devices have to be assigned to two different SRDF groups even if the two partner R2 devices reside on the same remote Symmetrix system.

"Primary devices (R1, R11)" on page 34 provides more information about R11 devices.

Concurrent SRDF topologies are supported on Fibre Channel, Gigabit Ethernet (GigE), and ESCON SRDF topologies. Basic concurrent SRDF solutions require that each R1 SRDF mirror operate in the same primary mode, either both synchronous or both semi-synchronous, but allows either or both SRDF mirrors to be placed into one of the adaptive copy modes. Advanced concurrent SRDF solutions support one or both R1 SRDF mirrors of the R11 device operating in asynchronous mode.

Figure 46 on page 122 shows a concurrent SRDF topology in which the R11 device communicates with the R2 device in Symmetrix B in synchronous mode. Concurrently, the same R11 device communicates with the R2 device in Symmetrix C in one of the adaptive copy modes.



Figure 46 Concurrent SRDF topology

Setting up a concurrent SRDF relationship

Setting up a concurrent SRDF relationship is a two-step process:

1. Create the initial R1 -> R2 pair between Symmetrix A and Symmetrix B.
2. Create the R11 -> R2 pair between Symmetrix A and Symmetrix C.

I/O operational rules as described in "SRDF/Synchronous Operations" on page 67 and "Adaptive Copy Operations" on page 115 also apply to concurrent SRDF topologies:

- When operating in synchronous mode, ending status for an I/O is not presented to the host until the remote Symmetrix system acknowledges receipt of the I/O to the primary system.
- If both SRDF mirrors are operating in synchronous mode, ending status is not presented to the host until both remote Symmetrix systems acknowledge receipt of the I/O.
- If one SRDF mirror operates in synchronous mode and one SRDF mirror operates in adaptive copy or asynchronous mode, ending status is presented to the host when the Symmetrix system across the synchronous SRDF links acknowledges receipt of the I/O.

Note: Concurrent SRDF can also be implemented with SRDF/Star. "Concurrent SRDF/Star" on page 132 contains additional information.

Concurrent SRDF/S with SRDF/A

Concurrent SRDF/S with SRDF/A is an advanced concurrent SRDF solution that allows one R1 mirror of an R11 device to operate in synchronous mode and the other R1mirror to operate in asynchronous mode. A group of R11 devices is remotely mirrored to one secondary site by using SRDF/S, and to another extended distance secondary site by using SRDF/A.

Each R1 mirror of an R11 device is configured to either a synchronous or an asynchronous SRDF group. Prior to Enginuity version 5876.82.57, to avoid latency on the SRDF links, it is recommended that you configure these groups on different SRDF directors to separate the SRDF/A and the SRDF/S traffic on the SRDF links. With Enginuity version 5876.82.57 or higher running on the primary site, you can configure both SRDF/S and SRDF/A groups on the same SRDF director on the primary site.

Concurrent SRDF/A

Enginuity versions 5875.135.91 and higher support another advanced concurrent SRDF solution, concurrent SRDF/A. This solution allows both R1 SRDF mirrors of the R11 device to operate in asynchronous mode. At any time, two SRDF/A sessions exist for the R11 device, one for each R1 SRDF mirror. You can manage these SRDF/A sessions using MSC.

"Multi Session Consistency (MSC) mode" on page 86 provides more details.

Migration with concurrent SRDF

Concurrent SRDF provides the ability to replace an existing R1 or R2 device with a new device in an SRDF pair. During migration, a concurrent SRDF relationship is established to transfer data from an existing device to a new device using adaptive copy disk mode. After this data transfer, the existing device is replaced with the newly populated device in the SRDF pair.

"Migrating data with concurrent SRDF" on page 153 provides more information about concurrent SRDF topologies that support thick-to-thin migration.

Concurrent SRDF with independent consistency protection

You can leverage the independent consistency protection feature in concurrent SRDF solutions that use SRDF/S on both legs of the concurrent SRDF topology (SRDF/S/S). This feature is based on Enginuity Consistency Assist (ECA) and enables you to independently manage consistency on each concurrent SRDF leg. If consistency protection on one leg is suspended, consistency protection on the other leg can remain active and continue protecting the primary site as shown in Figure 47 on page 124.

Concurrent SRDF with independent consistency assist feature is available with Enginuity versions 5874 and higher. The following options are available for managing consistency groups in a concurrent SRDF/S/S configuration:

- Independent consistency group protection
- Joint consistency group protection.



Figure 47 Concurrent SRDF/S with independent consistency protection

Benefits

The benefits of concurrent SRDF are:

- Extended distance solution with SRDF/A
- Support for device migration
- Integration with ECA.

Requirements

Requirements for concurrent SRDF interfamily connectivity apply. "SRDF multi-site interfamily connectivity" on page 146 provides details.

Cascaded SRDF

Cascaded SRDF is a three-site disaster recovery solution where data from a primary (R1) site is synchronously mirrored to a secondary (R21) site, and then asynchronously mirrored from the secondary (R21) site to a tertiary (R2) site. The core benefit behind a cascaded topology is its inherent capability to continue mirroring, with minimal user intervention, from the secondary site to the tertiary site in the event that the primary site fails. This enables a faster recovery at the tertiary site, assuming that the tertiary site is where you want to restart production operations. Figure 48 on page 126 illustrates a cascaded SRDF topology.

Both, the secondary and the tertiary site can be used as failover sites. Open systems SRDF solutions typically fail over to the tertiary site while mainframe solutions using AutoSwap typically fail over to the secondary site.



Figure 48 Cascaded SRDF topology

Cascaded SRDF uses dual-role SRDF devices (R21 devices) on the secondary site which acts as both an R2 to the primary site and an R1 to the tertiary site.

"Dual-role (cascaded) SRDF devices (R21)" on page 35 provides more information about R21 devices.

Setting up a cascaded SRDF relationship

Setting up a cascaded SRDF relationship is a two-step process:

1. Create the initial R1 –› R21 pair between Symmetrix A and Symmetrix B for the first hop. SRDF/S, SRDF/A, adaptive copy disk mode, or adaptive copy write-pending mode is allowed over the first hop.
2. Set up the R21 –› R2 pair between Symmetrix B and Symmetrix C for the second hop. SRDF/S, SRDF/A or adaptive copy disk mode is allowed over the second hop.

   The most common implementation is SRDF/S mode for the first hop and SRDF/A mode for the second hop.

Note: Only one hop (R1 –› R21 or R21 –› R2) can be asynchronous at a time. If R1 -› R21 is in asynchronous mode, R21 -› R2 must be in adaptive copy disk mode.

Table 18 on page 127 lists the SRDF modes allowed for cascaded SRDF.

Table 18 SRDF modes allowed for cascaded SRDF

R1 - R21	R21 - R2
Synchronous	Asynchronous
Synchronous	Adaptive copy disk
Asynchronous	Adaptive copy disk
Adaptive copy disk[1]	Asynchronous
Adaptive copy disk[1]	Adaptive copy disk
Adaptive copy write pending[1]	Asynchronous
Adaptive copy write pending[1]	Adaptive copy disk

1. The use of adaptive copy mode on the first leg does not provide full time consistency of the R21 or R2 devices.

Note: Cascaded SRDF can also be implemented with SRDF/Star. "Cascaded SRDF/Star" on page 134 contains additional information.

Benefits

The benefits of cascaded SRDF are:

- Faster recovery times at the tertiary site enabled by the continuous mirroring from the secondary site to the tertiary site in the event of the primary site failure.

  Note: In single session mode, SRDF continues mirroring data from the secondary to the tertiary site if the primary site fails. In MSC mode, user intervention is required.

- Tight-integration with TimeFinder product family.
- Management capability by using the current storage management portfolio of software products.
- Geographically dispersed secondary and tertiary sites.

Requirements

Requirements for cascaded SRDF interfamily connectivity apply. "SRDF multi-site interfamily connectivity" on page 146 provides details.

Extended Distance Protection

The SRDF/Extended Distance Protection (EDP) functionality is a licensed SRDF feature that offers a long distance disaster recovery (DR) solution. Figure 49 on page 128 shows an SRDF/EDP topology.

SRDF/EDP is achieved through a cascaded SRDF setup, where a VMAX 40K or VMAX 20K/VMAX system at a secondary site uses DL R21 devices to capture only the differential data that would be owed to the tertiary site in the event of a primary site failure. It is the data on the diskless R21 devices that helps these topologies achieve a zero Recovery Point Objective (RPO).

"Diskless R21 devices (DL R21)" on page 36 provides more information about DL R21 devices. The DL R21 devices must be preconfigured within the Symmetrix system prior to being placed in an SRDF relationship.

Note: The R1 and R2 devices at each end must both be thick devices or both be thin devices. "SRDF multi-site interfamily connectivity" on page 146 provides details for thick-to-thin support in SRDF/EDP.



Figure 49 SRDF/EDP basic topology

Setting up an SRDF/EDP relationship

Setting up a SRDF/EDP relationship is a two-step process:

1. Create the DLR1 --> R2 pair between Symmetrix B and Symmetrix C.
2. Create the R1 --> DLR2 pair between Symmetrix A and Symmetrix B. Now, the configuration is R1 --> DLR21 --> R2.

Note: The first hop R1 -> R21 cannot operate in asynchronous mode.

Table 19 on page 129 lists the SRDF modes allowed for SRDF/EDP.

Table 19 SRDF modes allowed for SRDF/EDP

R1 - DLR21	DLR21 - R2
Synchronous	Asynchronous
Synchronous	Adaptive copy write pending
Adaptive copy disk[1]	Asynchronous
Adaptive copy disk[1]	Adaptive copy write pending
Adaptive copy write pending[1]	Asynchronous
Adaptive copy write pending[1]	Adaptive copy write pending

1. The use of adaptive copy mode on the first leg does not provide full time consistency of the R21 or R2 devices.

Benefits

The benefits of SRDF/EDP are:

- A long distance mirroring and disaster restart solution with the ability to achieve zero RPO at the tertiary site.
- SRDF/EDP can also be implemented with SRDF/Star. SRDF/Star differential relationship support between Symmetrix A and Symmetrix C is intended for failover operations.

  "Cascaded SRDF/Star" on page 134 contains additional information.

Requirements

The following requirements apply to SRDF/EDP:

- Requirements for SRDF/EDP interfamily connectivity apply. "SRDF multi-site interfamily connectivity" on page 146 provides details.
- CKD devices are only supported in all thick SRDF/EDP environments.

SRDF/Star

SRDF/Star is a three-site disaster recovery solution that consists of the primary (production), secondary, and the tertiary site. The secondary site synchronously mirrors the production site, and the tertiary site asynchronously mirrors the production data. In the event of the primary site outage, EMC's SRDF/Star solution allows you to quickly move operations and re-establish remote mirroring between the remaining sites. Once conditions permit, you can quickly rejoin the primary site to the solution, resuming the SRDF/Star operations.

SRDF/Star operates in concurrent and cascaded environments. Figure 50 on page 130 shows two basic topologies that address different recovery and availability objectives:

- Concurrent SRDF/Star positions the secondary site or the tertiary site as potential recovery sites, and provides differential resynchronization between the secondary and the tertiary site. To achieve this positioning, some level of reconfiguration intervention is required to access point-of-disaster data.

  "Concurrent SRDF/Star" on page 132 contains additional information.

- Cascaded SRDF/Star positions the secondary or the tertiary site as potential recovery site with minimal intervention required to access point-of-disaster data. This solution provides differential synchronization between the primary and the tertiary site. "Cascaded SRDF/Star" on page 134 contains additional information.



Figure 50 Concurrent SRDF/Star (left) and cascaded SRDF/Star (right)

It is differential synchronization between two remote sites that allows SRDF/Star to rapidly reestablish cross-site mirroring in the event of the primary site failure. Differential synchronization between the remote sites dramatically reduces the time required to remotely mirror the new production site. SRDF/Star also provides a mechanism to determine which remote site has the most current data in the event of a rolling disaster that affects the primary site.

In all cases, you maintain the ability to select which site to operate from and which site's data to use when recovering from the primary site failure.

How SRDF/Star works

Suppose a concurrent SRDF/Star disaster-restart solution where Symmetrix A resides at the production site, Symmetrix B resides at the synchronous secondary site, and Symmetrix C resides at the asynchronous tertiary site as shown in Figure 50 on page 130. This concurrent SRDF/Star solution operates as follows:

- Normal operations —This refers to SRDF/Star remote mirroring to two remote sites, session management, and differential B-to-C or C-to-B resynchronization in SRDF/Star solutions. In normal operations, SRDF devices in Symmetrix A are in SRDF/S relationship with SRDF devices in Symmetrix B and in concurrent SRDF/A relationship with SRDF devices in Symmetrix C. SRDF/CG is required between Symmetrix A and Symmetrix B and an active host is required at Symmetrix A.
- Planned failovers — This refers to host-based automation processed at the synchronous remote site to reconfigure SRDF/Star in support of planned site failovers between the synchronous SRDF sites.
- Unplanned failovers — This refers to host-based automation processed at the tertiary site to perform SRDF/Star reconfiguration.

Note: If planned failovers are performed, local replicas created with TimeFinder are required at all three sites to retain any consistency images created either during a planned or an unplanned event. Host automation is required for failovers.

SRDF/Star control for mainframe

The host-based MSC task at the primary (R1) site controls normal SRDF/Star operation. MSC performs session management at Symmetrix B and when necessary at Symmetrix C. The MSC session management task maintains the information needed to perform differential synchronization between Symmetrix B and Symmetrix C.

SRDF Host Component, Consistency Group, support utilities, and documented procedures are also available to users to resynchronize sites and manage reconfigurations.

Automation for SRDF/Star mainframe host environments is provided by EMC GDDR. "SRDF and EMC GDDR" on page 172 contains additional information.

SRDF/Star control for open systems

You can use Solutions Enabler software to control, manage, and automate the SRDF/Star processes at the production site through the **symstar** command. Session management is required at Symmetrix A. The host software and Enginuity maintain the information needed to perform the differential synchronization between Symmetrix B and C.

Host-based automation is provided for normal, transient fault, unplanned failover (switch), and planned failover (switch) operations. This automation is delivered in Solutions Enabler with the SRDF/Star license. Detailed descriptions and implementation guidelines can be found in the *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide*.

Concurrent SRDF/Star

In concurrent SRDF/Star operations, Symmetrix B serves as the secondary site and the target of the SRDF/S links from Symmetrix A. Symmetrix C serves as the tertiary site and the secondary site of the SRDF/A links from Symmetrix A. The recovery links are the SRDF/A links between Symmetrix C and Symmetrix B.

Figure 51 on page 132 shows a concurrent SRDF/Star solution.



Figure 51 Concurrent SRDF/Star solution

Setting up a concurrent SRDF/Star relationship

Proceed as follows to set up a concurrent SRDF/Star relationship:

1. Create the initial R1 --> R2 pairs between Symmetrix A and Symmetrix B.
2. Create the R1 --> R2 pairs between Symmetrix A and Symmetrix C.

If Symmetrix A fails, you can fail over to Symmetrix B and resume remote mirroring after creating new SRDF pairs between Symmetrix B and Symmetrix C. Reconfiguration steps at Symmetrix B and Symmetrix C are necessary before SRDF mirroring can be resumed. To avoid these configurations steps, you can configure R22 (instead of R2) devices in Symmetrix C.

"Concurrent SRDF/Star with R22 devices" on page 133 provides more details.

Concurrent SRDF/Star with R22 devices

SRDF supports concurrent SRDF/Star using concurrent R22 devices. Figure 52 on page 133 shows R22 devices in Symmetrix C. R22 devices have two SRDF mirrors, only one of which is allowed to be active on the SRDF links at a given time.

R22 devices simplify SRDF/Star failover situations, improve the resiliency of the SRDF/Star application, and reduce the number of steps involved in failover procedures.

"Secondary devices (R2, R22)" on page 38 provides more information about the R22 device type.

Note: R22 devices are required in the SRDF/Star solutions including VMAX 10K or VMAXe systems.



Figure 52 Concurrent R22 SRDF/Star solution

Setting up a concurrent SRDF/Star relationship using R22 devices

Proceed as follows to set up a concurrent SRDF/Star relationship using R22 devices:

1. Create the initial R1 --> R2 pairs between Symmetrix A and Symmetrix B.
2. Create the R1 --> R2 pairs between Symmetrix A and Symmetrix C.
3. Create the R1 --> R2 pairs (using the inactive R2 SRDF mirror of the R22 device) between Symmetrix B and Symmetrix C to avoid configuration steps in the event of the primary site failure.

Cascaded SRDF/Star

Cascaded SRDF/Star enhances the basic cascaded SRDF functionality by using the SRDF/Star differential resynchronization. Figure 53 on page 134 shows a basic cascaded SRDF/Star solution. Symmetrix B serves as the secondary site and the target of the SRDF/S links from Symmetrix A. Symmetrix C serves as the tertiary site and the secondary of the SRDF/A links from Symmetrix B. The recovery SRDF/A links are between Symmetrix C and Symmetrix A.

Cascaded SRDF/Star uses a mechanism to determine when the current active R1 cycle (capture) contents reach the active R2 cycle (apply) over the long-distance SRDF/A links. This minimizes the amount of data that must be moved between Symmetrix B and Symmetrix C to fully synchronize them.



Figure 53 Cascaded SRDF/Star solution

In cascaded SRDF/Star solutions, the synchronous secondary site (Symmetrix B) is always more current than the asynchronous secondary site (Symmetrix C). If Symmetrix B fails, the cascaded SRDF/Star solution offers the ability to incrementally establish an SRDF/A session between Symmetrix A and Symmetrix C.

Setting up a cascaded SRDF/Star relationship

Proceed as follows to set up a cascaded SRDF/Star relationship:

1. Create the R1 --> R2 pairs between Symmetrix B and Symmetrix C.
2. Create the R1 --> R2 pairs between Symmetrix A and Symmetrix B.

The basic cascaded SRDF/Star solution requires that you configure new SRDF pairs between Symmetrix A and Symmetrix C before they can incrementally establish an SRDF/A session between these two systems. If R22 devices are used in Symmetrix C, these configuration steps are not needed.

"Cascaded SRDF/Star with R22 devices" on page 135 provides more details.

Cascaded SRDF/Star with R22 devices

Figure 54 on page 135 shows cascaded SRDF R22 mode of operation. This mode has the same topology as the basic cascaded SRDF/Star except that it requires R11 devices in Symmetrix A and R22 devices in Symmetrix C. The SRDF links between Symmetrix A and Symmetrix C are passive and serve as the recovery path. By using R22 devices in Symmetrix C, you can preconfigure the SRDF pairs required to incrementally establish an SRDF/A session between Symmetrix A and Symmetrix C in case Symmetrix B fails.

Note: R22 devices are required in the SRDF/Star solutions including VMAX 10K or VMAXe systems.



Figure 54 Cascaded R22 SRDF/Star solution

Setting up a cascaded SRDF/Star relationship using R22 devices

Proceed as follows to set up a cascaded SRDF/Star relationship with R22 devices:

1. Create the R1 --> R2 pairs between Symmetrix B and Symmetrix C.
2. Create the R1 --> R2 pairs between Symmetrix A and Symmetrix B.
3. Create the R1 --> R2 pairs (using the inactive SRDF R2 mirror) between Symmetrix A and Symmetrix C (to avoid configuration steps in case Symmetrix B fails).

Note: "Secondary devices (R2, R22)" on page 38 provides more information about the R22 device type.

Consider the following when creating a cascaded SRDF/Star solution with R22 devices:

- All devices at the production site in Symmetrix A must be configured as concurrent (R11) devices paired with R21 devices in Symmetrix B and R22 devices in Symmetrix C.
- All devices at the synchronous site in Symmetrix B must be configured as R21 devices.
- All devices at the asynchronous site in Symmetrix C must be configured as R22 devices.

SRDF/Star with SRDF/EDP

Enginuity versions 5874 and higher support SRDF/EDP topologies in SRDF/Star environments. These topologies require DL R21 devices in Symmetrix B (synchronous site) and R2 or R22 devices in Symmetrix C (asynchronous site). Figure 55 on page 136 shows an example with R2 devices configured in Symmetrix C.



Figure 55 SRDF/EDP in the cascaded SRDF/Star environments with an R2 site

The DL R21 devices operate in cascaded mode, mirroring host updates from Symmetrix A and simultaneously sending data to Symmetrix C. Since no full data copies are available in Symmetrix B, only Symmetrix C can be used as a disaster-restart site if Symmetrix A fails.

"Diskless R21 devices (DL R21)" on page 36 provides more information about DL R21 devices.

An SRDF/EDP solution configured with thin R1 and thin R2 devices can coexist with another SRDF/EDP solution configured with thick R1 and R2 devices in the same SRDF/Star implementation.

Setting up an SRDF/Star relationship using SRDF/EDP

Proceed as follows to set up an SRDF/Star with SRDF/EDP:

1. Create the DLR1 --> R2 pairs between Symmetrix B and Symmetrix C.
2. Create the R1 --> DLR2 pairs between Symmetrix A and Symmetrix B.

The resulting configuration is R11 --> DL R21 --> R2.

Like cascaded SRDF/Star, SRDF/Star with SRDF/EDP also supports R22 devices at the recovery site in Symmetrix C to avoid configuration steps if Symmetrix B fails.

Setting up an SRDF/Star relationship using SRDF/EDP and R22 devices

Proceed as follows to set up an SRDF/Star with SRDF/EDP:

1. Create the DLR1 --> R2 pairs between Symmetrix B and Symmetrix C.
2. Create the R1 --> DLR2 pairs between Symmetrix A and Symmetrix B.
3. Create the R1 --> R2 pairs between Symmetrix A and Symmetrix C using the inactive R2 SRDF mirror of the R22 device.

SRDF/Star benefits

SRDF/Star provides the following features and benefits:

- The ability to maintain protection and business continuance despite the loss of any site in a three-site topology (primary, secondary, tertiary).
- The ability to resume asynchronous protection between the secondary and tertiary sites, with minimal data movement, in the event of a primary site failure.
- EMC host-based SRDF software support for control functions.

Note: The *EMC Solutions Enabler Symmetrix SRDF CLI Product Guide* provides additional information.

SRDF/Star requirements

It is strongly recommended that all SRDF devices be locally protected and that each SRDF device is configured with TimeFinder to provide local replicas at each site.

The following requirements apply to SRDF/Star:

- Requirements for SRDF/Star interfamily connectivity apply. "SRDF multi-site interfamily connectivity" on page 146 provides details.

- Requirements for SRDF thick-to-thin support apply. "SRDF thick-to-thin support" on page 40 provides details.
- R22 devices are required in SRDF/Star solutions with VMAX 10K or VMAXe system involved.
- All SRDF/Star device pairs must be of the same geometry and size.
- All SRDF groups including inactive ones must be defined and operational prior to entering SRDF/Star mode.
- If multiple SRDF/Star solutions participate in the same SRDF/Star session, a mix of thin and thick SRDF/Star solutions is allowed.
- If SRDF/EDP is used in an SRDF/Star solution, the requirements for SRDF/Star and SRDF/EDP need to be all met. "Extended Distance Protection" on page 128 provides details about SRDF/EDP.

Four-site SRDF solution for open systems host environment

The four-site SRDF solution for open systems host environment replicates FBA data by using both concurrent and cascaded SRDF topologies. It is a multi-region disaster recovery solution with higher availability, improved protection, and limited downtime than the individual concurrent or cascaded SRDF solution.

The four-site SRDF can also be used for data migration. "Migrating data with concurrent and cascaded SRDF" on page 156 provides details.

Figure 56 on page 139 illustrates an example of the four-site SRDF solution. If two sites fail because of a regional disaster, the copy of the data is available and you can continue to have protection between the remaining two sites. You can configure a four-site SRDF topology from an existing two-site or three-site SRDF topology.



Figure 56 Four-site SRDF solution

Consider the following rules when you configure a four-site SRDF topology:

- Use adaptive copy disk mode while doing large data movements.
- Use SRDF/S when distance requirements allow, after copies are complete.
- Use SRDF/A when distance requirements demand it, after copies are complete
- On the cascaded legs, only one hop (R11 -> R21 or R21 -> R2) can operate in SRDF/A at a time.

Benefits

The four-site SRDF solution offers multi-region high availability by combining the benefits of concurrent and cascaded SRDF solutions. "Concurrent SRDF" on page 122 and "Cascaded SRDF" on page 126 provide details.

Requirements

The following requirements apply to the four-site SRDF:

- Requirements for four-site SRDF interfamily connectivity apply. "SRDF multi-site interfamily connectivity" on page 146 provides details.
- Requirements for SRDF thick-to-thin support apply. Refer to requirements for both concurrent SRDF and cascaded SRDF in "SRDF thick-to-thin support" on page 40 for details.
- CKD devices are not supported in four-site SRDF solutions.
- Only one hop of the cascaded legs can be asynchronous at a time.

SRDF/SQAR

SRDF/SQAR (Symmetrix Quadrilateral Asynchronous Replication) is a four-site implementation of SRDF/S and SRDF/A for mainframe host environments. SRDF/SQAR enables differential resynchronization between sites along the perimeter of a 'square' multi-site SRDF topology. EMC GDDR is required to implement SRDF/SQAR.

EMC GDDR support for the SRDF/SQAR configuration provides the ability to recover from a single or dual unplanned site outage in one region, with local SRDF/S protection established differentially between the recovery sites in another region. This enables you to quickly resume a workload with SRDF/S and Autoswap protection in another region.

Figure 57 on page 141 illustrates an example of SRDF/SQAR with AutoSwap environment.



Figure 57 SRDF/SQAR with AutoSwap environment

The *EMC SRDF Host Component for z/OS Product Guide* and the *EMC GDDR for SRDF/SQAR with AutoSwap Product Guide* provide more details about SRDF/SQAR.

Requirements

The following requirements apply to SRDF/SQAR:

- Requirements for SRDF/SQAR interfamily connectivity apply. "SRDF multi-site interfamily connectivity" on page 146 provides details.

- Mainframe Enabler version 7.5 or higher
- Autoswap version 4.1 or higher

CHAPTER 7
SRDF Interfamily Connectivity

This chapter describes SRDF two-site interfamily connectivity and multi-site interfamily connectivity. Topics include:

- Overview 144
- SRDF two-site interfamily connectivity 144
- SRDF multi-site interfamily connectivity 146

Overview

SRDF two-site and multi-site solutions allow you to configure different Symmetrix hardware models running different Enginuity versions within the same SRDF solution. This chapter describes the basic rules for SRDF two-site and multi-site interfamily connectivity.

SRDF two-site interfamily connectivity

Table 20 on page 144 describe the basic rules of interfamily connectivity for the two-site SRDF solutions for VMAX 40K, VMAX 20K/VMAX, and DMX.

Table 20 SRDF two-site interfamily connectivity (VMAX 40K, VMAX 20K/VMAX, DMX)

Enginuity version	Supported interfamily connectivity	SRDF thick-to-thin support for FBA devices	SRDF thick-to-thin support for CKD devices

5671	Connect to DMX:
	• 5671 – 5671 and higher
	Connect to VMAX 40K, VMAX 20K/VMAX:
	• 5671 – 5875.135.91 and higher
		• The Symmetrix system containing the thick devices needs to run Enginuity version 5671, 5773, or 5875.135.91 or higher.
		• The Symmetrix system containing the thin devices needs to run Enginuity version 5875.135.91 or higher.
		• The connectivity with the following pairs only supports migration only SRDF:
		  • 5671 – 5875.135.91 or higher
		  • 5773 – 5875.135.91
		"Full SRDF functionality support and migration only SRDF support" on page 150 provides details about migration only SRDF support.
		• Full SRDF functionality support for the other cases.
			• The Symmetrix system containing the thick devices needs to run Enginuity version 5671, 5773, 5875.286.218, or 5876.159.102 or higher.
			• The Symmetrix system containing the thin devices needs to run Enginuity version 5876.159.102 or higher.
			• The connectivity with the following pairs only supports migration only SRDF:
			  • 5671 – 5876.159.102 or higher
			"Full SRDF functionality support and migration only SRDF support" on page 150 provides details about migration only SRDF support.
			• Full SRDF functionality support for the other cases.

5773	Connect to DMX: • 5773 – 5671, 5772, 5773 Connect to VMAX 40K, VMAX 20K/VMAX: • 5773 – 5875.135.91 and higher		
5875	Connect to DMX: • 5875 – 5671, 5773 (The minimum version of 5875 is 5875.135.91) Connect to VMAX 40K, VMAX 20K/VMAX: • 5875 – 5875 (The minimum version of 5875 is 5875.135.91) • 5875 – 5876.82.57 (The minimum version of 5875 is 5875.135.91) • 5875 – 5876.159.102 and higher (The minimum version of 5875 is 5875.286.218)		
5876	Connect to DMX: • 5876 – 5671, 5773 (The minimum version of 5876 is 5876.82.57) Connect to VMAX 40K, VMAX 20K/VMAX: • 5876.82.57 – 5875.135.91 and higher • 5876.159.102 and higher – 5875.286.218 and higher		

For connectivity below Enginuity 5671 and other details for SRDF two-site interfamily connectivity, refer to EMC Online Support site (https://support.EMC.com). On the home page of the EMC Online Support site, search with the keyword `SRDF Two-site Interfamily Connectivity`.

Table 21 on page 145 describe the basic rules of interfamily connectivity for the two-site SRDF solutions including VMAX 10K or VMAXe systems.

Table 21 SRDF two-site interfamily connectivity for solutions including VMAX 10K or VMAXe

Enginuity version (VMAX 10K or VMAXe)	Supported interfamily connectivity	SRDF thick-to-thin support
5875	Connect to VMAX 10K or VMAXe: 5875 – 5875 (The minimum version of 5875 is 5875.231.172) 5875 – 5876.82.57 (The minimum version of 5875 is 5875.231.172) 5875 – 5876.159.102 and higher (The minimum version of 5875 is 5875.286.218) Connect to DMX: 5875 – 5773 (The minimum version of 5875 is 5875.231.172)	The system containing the thick devices must be DMX running 5773.

5876	Connect to VMAX 10K or VMAXe:
	5876.82.57 – 5875.231.172 and higher
	5876.159.102 – 5875.286.218 and higher
	Connect to DMX:
	5876.82.57 and higher – 5773
	Connect to VMAX 40K, VMAX 20K/VMAX:
	5876.82.57 (VMAX 10K or VMAXe) – 5876.82.57 and higher (VMAX 40K, VMAX 20K/VMAX, thin only)
	5876.159.102 and higher (VMAX 10K or VMAXe) – 5876.82.57 (VMAX 40K, VMAX 20K/VMAX, thin only)
	5876.159.102 and higher (VMAX 10K or VMAXe) – 5876.159.102 and higher (VMAX 40K, VMAX 20K/VMAX)
		• If VMAX 10K or VMAXe running 5876.82.57 is involved in the connectivity, the system containing the thick devices must be DMX running 5773.
		• If VMAX 10K or VMAXe running 5876.159.102 is involved in the connectivity, The system containing the thick devices can be the followings:
		  • DMX running 5773
		  • VMAX 40K or VMAX 20K/VMAX running 5876.159.102 or higher

The EMC Online Support site (https://support.EMC.com) provides greater details about SRDF two-site interfamily connectivity options. On the home page of the EMC Online Support site, search with the keyword `SRDF Two-site Interfamily Connectivity`.

SRDF multi-site interfamily connectivity

Table 22 on page 146 describes the basic rules of interfamily connectivity for each SRDF multi-site solution as well as the SRDF thick-to-thin support.

Table 22 SRDF multi-site interfamily connectivity (page 1 of 2)

SRDF solution	Basic rules of interfamily connectivity	SRDF thick-to-thin support (FBA and CKD; with or without VMAX 10K or VMAXe)

Concurrent SRDF	<u>Without VMAX 10K or VMAXe systems</u>
	• The minimum Enginuity version required on R11 site is Enginuity 5567.
	• The R2 sites can run any Enginuity versions that are supported in two-site connectivity between R11 and R2 sites.
	<u>With VMAX 10K or VMAXe systems</u>
	• If VMAX 10K or VMAXe systems are on R11 site:
	  • Enginuity 5876.82.57 or higher is required on the VMAX 10K or VMAXe systems.
	  • DMX systems participating in the solution needs to run Enginuity 5773.
	  • The corresponding SRDF two-site interfamily connectivity restrictions apply.
	• If VMAX 10K or VMAXe systems are not on R11 site:
	  • The corresponding SRDF two-site interfamily connectivity restrictions apply
		• For thick-to-thin connectivity on each concurrent leg, refer to SRDF thick-to-thin support in two-site SRDF solutions. Table 20 on page 144 and Table 21 on page 145 provide details.

Cascaded SRDF	<u>Without VMAX 10K or VMAXe systems</u>
	• The minimum Enginuity version required on R21 site is Enginuity 5773.
	• The R1 or R2 site can run other Enginuity versions compatible with Enginuity version of the R21 site.
	<u>With VMAX 10K or VMAXe systems</u>
	• If VMAX 10K or VMAXe systems are on R21 site:
	  • Enginuity 5876.82.57 or higher is required on all VMAX Family systems in the solution.
	  • DMX systems running Enginuity 5773 are supported in the solution if Enginuity 5876.159.102 or higher is running on the VMAX 10K or VMAXe systems on R21 site.
	• If VMAX 10K or VMAXe systems are not on R21 site:
	  • Enginuity 5876.82.57 or higher is required on the VMAX 10K or VMAXe systems in the solution.
	  • The Symmetrix systems on R21 site can be either DMX systems running Enginuity 5773 or VMAX 40K, VMAX 20K/VMAX systems running Enginuity 5876.82.57 or higher.
	  • The other Symmetrix system in the solution can run Enginuity versions compatible with the Enginuity version of systems on R21 site.
		• The Symmetrix system containing the thick devices needs to run Enginuity version 5773, or 5876.159.102 or higher.
		• The Symmetrix VMAX Family system containing the thin devices needs to run Enginuity version 5876.159.102 or higher.

Table 22 SRDF multi-site interfamily connectivity (page 2 of 2)

SRDF/EDP	Without VMAX 10K or VMAXe systems • The minimum Enginuity version required on the DL R21 site is Enginuity 5875.135.91. • The R1 or R2 site can run Enginuity version 5773 (DMX), or 5875.135.91 or higher (VMAX 40K, VMAX 20K/VMAX). With VMAX 10K or VMAXe systems • The DL R21 site must be VMAX 40K, VMAX 20K/VMAX, and Enginuity 5876.159.102 or higher is required on all three sites. • The DMX systems are not allowed in a SRDF/EDP solution with VMAX 10K or VMAXe systems involved.	Without VMAX 10K or VMAXe systems • Thick-to-thin connectivity in SRDF/EDP is only supported for FBA devices. To support thin R1 and thin R2 devices in SRDF/EDP, Enginuity 5875.135.91 or higher is required on all three sites. • CKD devices are supported in all thick SRDF/EDP environments.
Four-site SRDF	• Only FBA devices are supported. • The DMX systems in the solution must run Enginuity version 5773. • The VMAX systems in the solution must run Enginuity version 5875.135.91 or higher. • The VMAX 20K or VMAX 40K systems in the solution must run Enginuity version 5876.82.57 or higher. • The VMAX 10K or VMAXe systems in the solution must run Enginuity version 5876.159.102 or higher. • The corresponding SRDF interfamily connectivity restrictions apply.	• For thick-to-thin support on the concurrent legs, refer to details about SRDF thick-to-thin support in concurrent SRDF. • For thick-to-thin support on the cascaded legs, refer to details about SRDF thick-to-thin support in cascaded SRDF.

SRDF/Star	• For concurrent SRDF/Star, refer to details about concurrent SRDF interfamily connectivity.
	• For cascaded SRDF/Star, refer to details about cascaded SRDF interfamily connectivity.
	• If R22 devices are used, the R22 site can be DMX running Enginuity 5773 or VMAX family systems running Enginuity 5875.135.91 or higher.
	• For SRDF/Star with SRDF/EDP, the requirements for SRDF/EDP also apply.
		• Prior to Enginuity version 5876.163.105, Enginuity 5876.159.102 is required on all three sites to support thick-to-thin connectivity. DMX running Enginuity 5773 is not supported in thick-to-thin connectivity in SRDF/Star. DMX running Enginuity 5773 is allowed in all thick SRDF/Star solutions.
		• With Enginuity version 5876.163.105 or higher, DMX is supported in thick-to-thin connectivity in SRDF/Star with the following limitations:
		  • The Symmetrix system containing the thick devices needs to run Enginuity version 5773, or 5876.163.105 or higher.
		  • The Symmetrix VMAX Family system containing the thin devices needs to run Enginuity version 5876.163.105 or higher.
		  • With Enginuity 5876.163.105, only one DMX is allowed in the SRDF/Star solution with mixed thick and thin devices. This limitation is removed in Enginuity 5876.229.145 and higher.

SRDF/SQAR	• Mainframe environments only. VMAX 10K or VMAXe is not supported. • Enginuity version 5876.159.102 or higher is required on all Symmetrix systems in the environment.	• SRDF thick-to-thin support is available with full SRDF functionality.

For corresponding two-site interfamily connectivity restrictions that apply to multi-site interfamily connectivity, the EMC Online Support site (https://support.EMC.com) provides greater details about SRDF two-site interfamily connectivity options. On the home page of the EMC Online Support site, search with the keyword `SRDF Two-site Interfamily Connectivity`.

CHAPTER 8
SRDF Migration Operations

This chapter describes the SRDF migration operations. Topics include:

- Overview 150
- Migrating data with concurrent SRDF 153
- Migrating data with concurrent and cascaded SRDF 156

Overview

SRDF migration is typically used in technology refresh operations to migrate production data from an older Symmetrix system running an older Enginuity version to the latest Symmetrix hardware model running the latest Enginuity version. You can migrate your data from thick to thick environments, or from thin to thin environments. You can also move your data from a thick to a thin environment with SRDF thick-to-thin support. Once the data migration process is complete, the production environment is typically moved to the Symmetrix system to which the data was migrated.

SRDF migration solutions by using two-site, three-site, or four-site SRDF are available to address the different requirements. You are able to effectively move large amounts of data across the SRDF links, while keeping remote replication for protection by taking advantage of adaptive copy modes, SRDF/S, and SRDF/A.

In open systems host environments, you can use the `symrdf migrate` functionality of Solutions Enabler to reduce migration resynchronization times while replacing either the R1 or R2 devices in an SRDF two site topology. The *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide* provides more details about the `symrdf migrate` functionality.

In both mainframe and open systems host environments, multiple types of migration requirements can be readily accomplished by using multi-site SRDF solutions and procedures. The *EMC Symmetrix SRDF Host Component for z/OS Product Guide* and *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide* provide more information about multi-site SRDF operations in corresponding host environments.

EMC support personnel is available to assist with the planning and execution of your individual migration project requirements.

Full SRDF functionality support and migration only SRDF support

In most of the cases, you can migrate your data with full SRDF functionality support, including disaster recovery and other advanced SRDF features.

In the other cases where the full SRDF functionality support is not available, you can move your data across the SRDF links by using the migration only SRDF support. Compared with *full SRDF functionality support*, the *migration only SRDF support* provides only limited SRDF functionality that enables you to move data across the SRDF links. Table 23 on page 150 lists the limitations on SRDF operations and features in migration only SRDF support.

Table 23 Limitations of migration only SRDF support (page 1 of 2)

SRDF operations or features	Limitations during migration
Unprotected mirror	No support
R2 to R1 copy	Support with restrictions[1]
R2 to R1 copy with write	Support with restrictions[1]
Failover	No support
Failback	No support
SRDF/Star	No support
Domino	No support

Table 23 Limitations of migration only SRDF support (page 2 of 2)

SRDF/A features: • DSE • Consistency Group • ECA • MSC	No support
Dynamic SRDF: • Create SRDF pairs • Delete SRDF pairs • R1/R2 personality swap • Move SRDF pairs	No support
TimeFinder operations: TimeFinder/Mirror [2] TimeFinder/Clone TimeFinder/Snap[2] TimeFinder VP Snap	Support only on R1
Virtual LUN Virtual Pool (VP) mobility	No support
Online configuration change	Support with restrictions[3]
Static to dynamic configuration change	No support
Out-of-family Non-Disruptive Upgrade (NDU)	No support

1. Support only when devices rebuild from un-rebuildable RAID group failures
2. For SRDF for Symmetrix VMAX 40K, VMAX 20K/VMAX, DMX solutions only
3. Suspend migration prior to online upgrade and online configuration changes if the changes affect the group or devices being migrated. If the changes do not affect the migration group, they are allowed without suspending migration

Note: The unsupported or restricted operations listed in Table 23 on page 150 are limitations on SRDF groups involved in migration process, but full support is still provided for SRDF groups that meet full support requirements and are not involved in the migration.

For more details about the unsupported features in different host environments, contact EMC support personnel.

The following SRDF interfamily connectivity cases are solutions for DMX only, and are supported with migration only SRDF:

- Enginuity version 5267 – 5670, 5671, 5771, 5772
- Enginuity version 5567 – 5670, 5671, 5771, 5772, 5773
- Enginuity version 5568 – 5773

The following SRDF interfamily connectivity cases are supported with migration only SRDF when you migrate data between thick and thin devices. Full SRDF functionality support is available in all thick or all thin environments:

- Enginuity version 5671 – 5875.135.91 and higher
- Enginuity version 5773 – 5875.135.91

Zero space reclamation, introduced with Enginuity version 5875.135.91, is available with both full SRDF functionality support and migration only SRDF support.

Migrating data with concurrent SRDF

The concurrent SRDF allows you to non-disruptively migrate data between Symmetrix systems along one concurrent SRDF leg while keeping remote mirroring for protection along the other concurrent SRDF leg. Once the migration process completes, the concurrent SRDF topology is removed, thereby resulting in a two-site SRDF topology.

Note: Specific hardware models and Enginuity versions are required for migrating data between different platforms. Refer to "SRDF multi-site interfamily connectivity" on page 146 for details.

Replacing R2 devices with new R2 devices

Figure 58 on page 153 illustrates an example of replacing the original R2 devices with new R2 devices. It shows the initial two-site topology, the migration process, and the final SRDF topology.



Figure 58 Migrating data and removing the original secondary system (R2)

In open systems host environments, you can complete the migration process by following steps below with the assistance of `symrdf migrate` functionality:

1. Configuring concurrent R11 devices on the primary system.
2. Adding a Symmetrix system (new R2) to one leg of the concurrent SRDF topology.
3. Migrating data from R11 to new R2 devices.
4. Removing one leg of the concurrent SRDF along with the original secondary (R2) Symmetrix system.

In both open systems and mainframe host environments, you can also manually perform the migration process illustrated in Figure 58 on page 153. EMC support personnel is available to assist with the planning and execution of your migration projects.

After migration, the original primary system is mirrored to a new secondary system. If you want to continue replacing the R1 devices with new R1 devices, refer to "Replacing R1 devices with new R1 devices" on page 154 for details.

Replacing R1 devices with new R1 devices

Figure 59 on page 154 illustrates an example of replacing the original R1 devices with new R1 devices. It shows the initial two-site topology, the migration process, and the final SRDF topology.



Figure 59 Migrating data and replacing the original primary system (R1)

In open systems host environments, you can complete the migration process by following steps below with the assistance of `symrdf migrate` functionality:

1. Configuring concurrent R11 devices on the primary system.
2. Adding a Symmetrix system (new R2) to one leg of the concurrent SRDF topology.
3. Migrating data from R11 to new R2 devices.
4. Terminating all SRDF pairings along both legs of the concurrent SRDF topology.
5. Converting new R2 to R1 devices.
6. Configuring the SRDF pairing between new R1 and old R2 devices.

In both open systems and mainframe host environments, you can also manually perform the migration process illustrated in Figure 59 on page 154. EMC support personnel is available to assist with the planning and execution of your migration projects.

After migration, the new primary system is mirrored to the original secondary system. If you want to continue replacing the R2 devices with new R2 devices, refer to “Replacing R2 devices with new R2 devices” on page 153 for details.

Migrating data with concurrent and cascaded SRDF

You can use the combination of concurrent SRDF and cascaded SRDF to replace both R1 and R2 devices at the same time.

Note: Specific hardware models and Enginuity versions are required for migrating data between different platforms. Refer to requirements for concurrent SRDF and cascaded SRDF in "SRDF multi-site interfamily connectivity" on page 146 for detailed requirements.

Replacing R1 and R2 devices with new R1 and R2 devices

Figure 60 on page 156 illustrates an example of replacing both R1 and R2 devices with new R1 and R2 devices at the same time. It shows the initial two-site topology, the migration process, and the final topology.



Figure 60 Migrating data and replacing the original primary (R1) and secondary (R2) systems

In open systems host environments, you can complete the migration process by following steps below with the assistance of `symrdf migrate` functionality:

1. Configuring concurrent R11 devices on the primary system.
2. Adding a cascade of two Symmetrix systems.
3. Migrating data from R11 to new R21 and new R2 devices.
4. Removing the new cascaded legs with new R21 and R2 devices from the concurrent SRDF topology.

In both open systems and mainframe host environments, you can also manually perform the migration process illustrated in Figure 60 on page 156. EMC support personnel is available to assist with the planning and execution of your migration projects.

CHAPTER 9
SRDF Integration

This chapter describes SRDF and its integration with other EMC products. Topics include:

- SRDF and TimeFinder 160
- SRDF and open systems clusters 169
- SRDF and mainframe automation software 171
- SRDF and open systems automation software 174
- SRDF and VMware environments 175
- SRDF and EMC FAST VP 176
- SRDF and EMC RecoverPoint 178

SRDF and TimeFinder

The TimeFinder family of products are Symmetrix local replication solutions designed to create point-in-time copies of critical data. You can configure backup sessions, initiate copies, and terminate TimeFinder operations from mainframe and open systems controlling hosts using EMC Symmetrix host-based TimeFinder control software.

The TimeFinder local replication solutions include TimeFinder/Clone, TimeFinder/Snap, and TimeFinder VP Snap. TimeFinder/Clone creates full-device and extent-level point-in-time copies. TimeFinder/Snap creates pointer-based logical copies that consume less storage space on physical drives. VP Snap provides the efficiency of Snap technology with improved cache utilization and simplified pool management. Each solution guarantees high data availability.

Note: TimeFinder/Snap is only supported in SRDF for Symmetrix VMAX 40K, VMAX 20K/VMAX, DMX solutions.

The production device, also known as the TimeFinder source device, is continuously available to the production application. The backup device, also known as the TimeFinder target device, becomes available to its host as soon as you initiate a TimeFinder session. This allows you to initiate backups and have backup copies immediately available to auxiliary applications. Figure 61 on page 160 illustrates TimeFinder/Clone operations.



Figure 61 TimeFinder/Clone

Note: TimeFinder/Mirror is not supported on Symmetrix VMAX Family systems. The TimeFinder/Clone product provides TimeFinder Clone Emulation mode to support TimeFinder/Mirror commands. This allows customers to run their TimeFinder/Mirror automation software using TimeFinder/Clone infrastructure.

The *EMC Symmetrix TimeFinder for VMAX 40K, VMAX 20K/VMAX Series Product Guide* provides more information about TimeFinder.

TimeFinder is tightly integrated with SRDF solutions. You can use TimeFinder and SRDF products to complement each other in environments that require both local replication and SRDF remote mirroring. For example, TimeFinder/Clone is often used to create local gold copies of the SRDF devices for recovery operations and for testing disaster recovery solutions. SRDF/Star, GDDR (see "SRDF and EMC GDDR" on page 172), and AutoStart (see "SRDF and open systems automation software" on page 174) all use TimeFinder local replication solutions.

The following are the key benefits of TimeFinder integration with SRDF:

- **Remote command support for simplified automation** — EMC host-based control software commands can be transferred across the SRDF links, which allows you to issue a single command to the primary system to initiate TimeFinder operations on both the primary and secondary system.
- **Consistent data images across multiple devices and Symmetrix systems** — The SRDF/Consistency Group (SRDF/CG) feature is used in SRDF/S solutions to guarantee that a dependent-write consistent image of production data on the R1 devices is created across the SRDF links.

  The TimeFinder/Consistency Group (TimeFinder/CG) feature guarantees that a consistent point-in-time image of data written across multiple local devices (TimeFinder source devices) is created on another set of local devices (TimeFinder target devices).

  SRDF/CG and TimeFinder/CG both use the Enginuity Consistency Assist infrastructure. By using TimeFinder/CG in an SRDF configuration, you can create dependent-write consistent local and remote images of production data across multiple devices and Symmetrix systems.

Note: The SRDF/A single session mode solution guarantees dependent-write consistency across the SRDF links and does not require SRDF/CG. MSC mode requires host software to manage consistency among multiple sessions.

R1 and R2 devices in TimeFinder operations

R1 and R2 devices can participate in TimeFinder operations to create local replicas. The following rules apply:

- You can use R1 devices and R2 devices[5] as source devices in TimeFinder operations.
- You can use R1 devices as target devices in TimeFinder/Clone and TimeFinder/Snap, but not in VP Snap.
- You can use R2 devices as target devices in TimeFinder, if SRDF remote mirroring is not active. This restriction applies because the R2 device cannot receive data from its R1 partner device across the SRDF links and from the local TimeFinder source device at the same time. If you want to use R2 devices as TimeFinder target devices, you must first suspend the SRDF remote mirroring.

5. Prior to Enginuity version 5876.159.102, R2 devices of the second hop (R21 -> R2) in a cascaded SRDF topology cannot be used as the source devices in TimeFinder/Clone (without -precopy), TimeFinder/Snap, or TimeFinder VP Snap operations. This restriction is removed in Enginuity version 5876.159.102.

Figure 62 on page 162 shows an example.



Figure 62 SRDF devices in TimeFinder operations

TimeFinder and SRDF adaptive copy mode operations

The most widely used implementation of TimeFinder and SRDF adaptive copy mode operations is SRDF Automated Replication (SRDF/AR).

SRDF/AR described in "SRDF/Automated Replication" on page 118 requires TimeFinder to:

- Separate remote replication (single-hop solution) or remote mirroring (multi-hop solution) from production processing on both Symmetrix systems.
- Create consistent local copies of data in Symmetrix systems that participate in an SRDF/AR solution.

Figure 63 on page 163 shows an SRDF/AR single-hop solution where the production device on the primary system also participates as the source device in TimeFinder operations. The TimeFinder target device is the SRDF R1 device, remotely mirrored to the SRDF R2 device.



Figure 63 SRDF/AR single-hop solution

The R2 device on the secondary system is the TimeFinder source device, replicated to the TimeFinder target device, which can be mapped to a host connected to the secondary system. By replicating R2 devices on the secondary system, you can deploy secondary system's resources without interrupting SRDF operations.

Such implementation separates SRDF operations from production processing on the primary and secondary systems. The multi-hop SRDF/AR solution shown in Figure 64 on page 163 uses the same concept on the second leg of the SRDF configuration.



Figure 64 SRDF/AR multi-hop solution

Since the SRDF links are used to send SRDF and TimeFinder control commands from the local controlling host to the local and remote Symmetrix systems, this solution allows you to automate SRDF/AR operations by issuing control commands only to Symmetrix A.

TimeFinder and SRDF/A operations

TimeFinder can be used in SRDF/A solutions to create local copies of the SRDF devices on the primary system and local copies of their SRDF partner devices on the secondary system. Because SRDF/A and large-scale TimeFinder copy operations heavily utilize system cache, Enginuity version 5875.135.91 provides the SRDF/A device-level pacing feature that allows you to balance cache utilization.

The SRDF/A device-level pacing option is designed for SRDF/A solutions in which the SRDF/A R2 devices also participate in TimeFinder operations as TimeFinder source devices. To support most of the TimeFinder operations in SRDF/A solutions, it is required that SRDF/A device-level pacing is activated on the corresponding R1 devices and supported in the SRDF/A session. Table 24 on page 164 lists the SRDF/A device-level pacing requirements for different TimeFinder operations.

Table 24 SRDF/A device-level pacing requirements for TimeFinder operations

TimeFinder operations	SRDF/A device-level pacing requirements	Enginuity version on R1 and R2 devices
Full-device TimeFinder/Clone (with -precopy)	No	Any versions supporting integration of TimeFinder operations and SRDF/A
Extent level TimeFinder/Clone	No	Any versions supporting integration of TimeFinder operations and SRDF/A
Full-device TimeFinder/Clone (without -precopy)	Active	5875.135.91 or higher[1]
TimeFinder/Snap	Active	5875.135.91 or higher[1]
TimeFinder VP Snap	Active	5876.82.57 or higher[1]

1. To support TimeFinder operations on the second hop (in asynchronous mode) of cascaded SRDF, Enginuity version 5876.159.102 is required on the R21 devices.

Prior to Enginuity version 5876.159.102, you can only use TimeFinder/Clone (with -precopy) to create local copies of the SRDF R2 devices on the second hop of cascaded SRDF. With Enginuity version 5876.159.102 and EMC host-based SRDF software, you can enable SRDF/A device-level pacing on the second hop (R21 -> R2) of a cascaded SRDF topology (including cascaded SRDF/Star and SRDF/EDP), so the R2 or R22 devices of the second hop can be used as the source devices in TimeFinder operations with the following requirements:

- If the SRDF/A session for the R21 -> R2 device pair is active, the following requirements must be met:
  - The state of the R21 -> R2 device pair cannot be TransIdle.
  - SRDF/A device-level pacing must be activated and supported in the SRDF/A session.

- If the SRDF/A session for the R21 -> R2 device pair is not active, SRDF/A device-level pacing must be configured for autostart on the SRDF group that contains the R1 mirror of the R21 device.

"SRDF/A and cache utilization" on page 111 and "Device-level (or TimeFinder) pacing" on page 113 provide more information.

TimeFinder and SRDF/S operations

SRDF/S solutions support any type of TimeFinder copy sessions running on R1 and R2 devices as long as the conditions described in "R1 and R2 devices in TimeFinder operations" on page 161 are met.

The following sections explain simultaneous TimeFinder/Clone solutions that employ SRDF/S for remote mirroring and TimeFinder/Clone for local replication.

Simultaneous TimeFinder/Clone prior to Enginuity 5875.135.91 for mainframe environments only

Simultaneous TimeFinder/Clone operations with Enginuity versions prior to 5875.135.91 involve the following processes:

1. SRDF/S remotely mirrors R1 devices in Symmetrix A to R2 devices in Symmetrix B.
2. You initiate TimeFinder/Clone copy sessions to replicate R1 devices (TimeFinder/Clone source devices) to another set of R1 devices (TimeFinder/Clone target devices) in Symmetrix A.
3. SRDF/S sends all data that arrives for TimeFinder/Clone target devices in Symmetrix A to their R2 partner devices in Symmetrix B.

Figure 66 on page 166 shows an example of simultaneous TimeFinder/Clone solution prior to Enginuity version 5875.135.91.



Figure 65 Simultaneous TimeFinder/Clone solution prior to Enginuity version 5875.135.91

In this scenario, data travelling across the SRDF/S links includes host updates to TimeFinder/Clone target devices and data locally copied from the TimeFinder/Clone source devices.

TimeFinder/Clone target devices in Symmetrix B do not capture the point-in-time copies until they are synchronized with their SRDF partner devices in Symmetrix A.

Simultaneous TimeFinder/Clone and failover scenarios

With Enginuity version 5875.135.91, simultaneous TimeFinder/Clone operations provide failover capabilities in mainframe host environments using full device or dataset level TimeFinder/Clone operations.

With Enginuity version 5875.135.91 running on both local and remote systems, simultaneous TimeFinder/Clone operations do not move data across the SRDF/S links to create remote images on the R2 TimeFinder/Clone target devices in Symmetrix B. You issue a special TimeFinder/Clone copy request to Symmetrix A. This single request to Symmetrix A simultaneously creates TimeFinder/Clone copies on both sides of the SRDF/S relationship.

Figure 66 on page 166 shows an example.



Figure 66 Simultaneous TimeFinder/Clone with Enginuity version 5875.135.91

These simultaneous TimeFinder/Clone operations involve the following processes:

1. The R1 devices (TimeFinder/Clone source devices) in Symmetrix A must be synchronized with their SRDF partners in Symmetrix B.
2. Once the SRDF/S pairs are synchronized, you issue a special TimeFinder/Clone copy request to Symmetrix A. This request is transferred across the SRDF/S links to Symmetrix B to simultaneously initiate TimeFinder/Clone copy sessions from TimeFinder/Clone source to TimeFinder/Clone target devices on both sides of the SRDF/S relationship.

When simultaneous TimeFinder/Clone copy sessions start in Symmetrix A and in Symmetrix B, both systems hold the same data image. The TimeFinder/Clone source devices in Symmetrix B capture current data since they are synchronized with their SRDF partners in Symmetrix A. The TimeFinder/Clone target devices in Symmetrix B capture the same point-in-time copy as their SRDF partners in Symmetrix A because of simultaneous TimeFinder/Clone operations that take place on both sides of the SRDF/S relationship.

By preserving the same data image as Symmetrix A, including in-process clone copies, Symmetrix B can become the failover site and can assure that all data made available to the host after the failover operation is current data.

This feature can be used with concurrent SRDF on the SRDF/S leg of a concurrent SRDF configuration. It is not supported in SRDF/A, cascaded SRDF, or in SRDF/Star configurations and does not support SRDF pairs in which the R2 devices are larger than their R1 partners.

The *EMC Mainframe Enablers TimeFinder/Clone SNAP Facility Product Guide* provides details about controlling simultaneous TimeFinder/Clone operations using mainframe host-based control software.

TimeFinder and SRDF in virtual provisioned environments

TimeFinder operations are supported for thin SRDF devices. The rules described in "R1 and R2 devices in TimeFinder operations" on page 161 also apply to thin SRDF devices.

Note: With Enginuity version 5773, you cannot use TimeFinder/Mirror in virtually provisioned SRDF environments. This restriction applies only to Enginuity version 5773. Enginuity versions lower than 5773 do not support Virtual Provisioning while Enginuity versions higher than 5773 no longer support TimeFinder/Mirror.

Table 25 on page 167 and Table 26 on page 168 list the SRDF and TimeFinder support for Virtual Provisioning.

Table 25 Thin SRDF devices with TimeFinder/Clone or VP Snap

Thin SRDF device type	SRDF mode	TimeFinder/Clone or VP Snap role	Enginuity version		
			5773	5874	5875.135.91 and higher
R1	SRDF/A	source or target	No	Yes	Yes
R2	SRDF/A	source	No	Yes[1]	Yes[2]
R1	SRDF/S or SRDF adaptive copy	source or target	Yes	Yes	Yes
R2	SRDF/S or SRDF adaptive copy	source	Yes	Yes	Yes

1. Only full-device TimeFinder/Clone with the precopy option is supported.
2. Full-device TimeFinder/Clone and extent level TimeFinder/Clone copy sessions that use any copy option are supported as long as the SRDF/A write pacing is enabled. TimeFinder/Clone with the precopy option is allowed without the SRDF/A write pacing feature enabled.

Table 26 Thin SRDF devices and TimeFinder/Snap

Thin SRDF device type	SRDF mode	TimeFinder/Snap role	Enginuity version		
			5773	5874	5875.135.91 and higher
R1	SRDF/A	source or target (virtual)	No	Yes	Yes
R2	SRDF/A	source	No	No	Yes[1]
R1	SRDF/S or SRDF adaptive copy	source or target (virtual)	No	Yes	Yes
R2	SRDF/S or SRDF adaptive copy	source	No	Yes	Yes

1. The SRDF/A device-level pacing must be enabled.

SRDF and open systems clusters

This section provides an overview of SRDF integration with EMC and other cluster software to support automated recovery operations in open systems cluster environments.

SRDF/CE enables automated or semi-automated site failover and uses SRDF/S or SRDF/A with Microsoft Failover Clusters (MSCS). SRDF/CE allows Windows Server 2003 and Windows Server 2008 Enterprise and Datacenter editions running Microsoft Failover Clusters to operate across a single pair of SRDF-connected systems as geographically distributed clusters. Up to 64 MSCS clusters can share the same SRDF pair. Using SRDF links, SRDF/CE expands the range of cluster storage and management capabilities while ensuring full protection of the SRDF remote mirroring. Two Symmetrix systems are connected through SRDF to provide automatic failover of SRDF-mirrored devices during MSCS node failover.

Figure 67 on page 169 shows an example of a two-node, two-cluster SRDF/CE configuration.



Figure 67 SRDF/CE two-node, two-cluster configuration

SRDF/Star with Veritas Cluster Server provides high availability and automated failover through storage-based mirroring and server clustering through SRDF/S and SRDF/Star with Veritas Cluster Server.

Hewlett-Packard MC/ServiceGuard with MetroCluster software and EMC SRDF is an integrated open systems disaster-tolerant solution that combines geographically dispersed HP-UX clusters into a single manageable system to provide automatic application failover.

When a site becomes inoperable due to a system failure or a disaster event, the application switch occurs. This action automatically transfers control of the affected application to another HP 9000 system located at a secondary site, while also automatically read/write-enabling R2 devices associated with the transferred application.

HP MetroCluster for EMC SRDF allows cluster support between sites up to 25 miles (40 km) apart with an intersite link, which is used to provide heartbeat communication between cluster nodes. HP Continental Clusters extend the distance limits while operating in a similar manner.

Sun Solaris Cluster campus topology provides application resumption with automated failover with up to 125 miles (200 km) between Symmetrix sites using SRDF/S.

If the primary site becomes inoperable in the event of a data center failure, fault, or disaster event, the resource groups associated with the Solaris Cluster systems automatically fail over to the systems at the secondary site.

Sun Solaris Cluster Geographic Edition enables independent geographically dispersed Solaris Clusters to provide automatic failure detection with semi-automated failover of applications up to 125 miles (200 km) between SRDF/S sites and unlimited distance between SRDF/A sites.

If the primary site becomes inoperable due to a system failure or disaster event, the secondary cluster detects the failure and notifies the system administrator who can then request a takeover with a single command. This action performs SRDF state validations, then proceeds to bring the applications online in the secondary cluster.

SRDF and mainframe automation software

This section provides an overview of SRDF for VMAX 40K, VMAX 20K/VMAX, and DMX integration with EMC mainframe automation software including EMC AutoSwap and EMC Geographically Dispersed Disaster Restart (GDDR).

SRDF and EMC AutoSwap

EMC AutoSwap is a mainframe solution that integrates with the SRDF/S product offering and provides transparent switching between Symmetrix DASD subsystems. AutoSwap extends Parallel Sysplex redundancy to the disk systems and allows your to totally shutdown the primary system and continue production operations on the secondary system without disturbing applications.

AutoSwap can swap all z/OS servers to the remotely mirrored Symmetrix systems. This enables you to power down the entire DASD center or perform required maintenance.

Figure 68 on page 171 shows an example of EMC AutoSwap processing.



Figure 68 AutoSwap processing

AutoSwap requirements

AutoSwap has the following requirements:

- SRDF/S (also SRDF/CG if Enginuity version lower than 5874 runs on the primary system)
- Sufficient User Control Blocks (UCBs) must be available for all primary and secondary devices
- ResourcePak Base for z/OS.

The *EMC AutoSwap Product Guide* contains additional information.

SRDF and EMC GDDR

EMC Geographically Dispersed Disaster Restart (GDDR) is a mainframe software product that automates business recovery following both planned outages and disaster situations, including the total loss of a data center. EMC GDDR achieves this goal by providing monitoring, automation and quality controls to the functionality of many EMC and third-party hardware and software products required for business restart. As EMC GDDR restarts production systems following disasters, it does not reside on the same servers that it is seeking to protect. EMC GDDR resides on separate logical partitions (LPARs) from the host servers that run your application workloads.

EMC GDDR is installed on a control LPAR at each site. Each EMC GDDR node is aware of the other EMC GDDR nodes through network connections between each site. This awareness allows EMC GDDR to perform the following:

- Detect disasters
- Identify survivors
- Nominate the leader
- Recover business at one of the surviving sites.

Figure 69 on page 172 shows a three-site GDDR configuration, the concurrent SRDF/Star with AutoSwap.



Figure 69 Three-site concurrent SRDF/Star with GDDR

EMC GDDR has no limitations on the number of EMC Symmetrix systems that can be managed. Any limitations are subject to restrictions in EMC hardware and software.

EMC GDDR can be used with the following SRDF configurations:

- SRDF/S with SRDF/CG — The two-site SRDF/S with SRDF/CG configuration provides disaster restart capabilities at the secondary site.
- SRDF/S with AutoSwap — The two-site SRDF/S with AutoSwap configuration provides for near-continuous availability through device failover between sites.
- SRDF/A — The two-site SRDF/A configuration provides disaster restart capabilities at the secondary site.
- SRDF/Star — The three-site SRDF/Star configuration provides disaster restart capabilities at either the secondary or tertiary site. Concurrent and cascaded SRDF support further minimize the RTO at the tertiary site.
- SRDF/Star with AutoSwap — The three-site SRDF/Star with AutoSwap configuration provides for near-continuous availability through device failover between the primary and secondary sites as well as disaster restart capabilities at the tertiary site. Concurrent and cascaded SRDF support further minimize the RTO at the tertiary sites.

EMC GDDR has been designed to be customized to operate in any of these configurations. EMC GDDR functionality is controlled by a parameter library. During EMC GDDR implementation, this parameter library is customized to reflect the following:

- The prerequisite software stack
- The desired data center topology (two-site versus three-site, synchronous or asynchronous).

Note: The *EMC GDDR Product Guide* provides details on GDDR use and operation.

The following software is required to run EMC GDDR:

- Any version of z/OS currently supported by IBM
- IBM hardware Management Console (HMC) API
- CA-OPS/MVS
- EMC SRDF/Host Component for z/OS
- EMC ResourcePak Base for z/OS
- EMC Consistency Group for z/OS
- EMC TimeFinder for z/OS
- EMC AutoSwap for z/OS
- EMC TimeFinder/Clone Mainframe SNAP Facility (optional)

Note: The EMC Online Support site (https://support.EMC.com) contains more details about the software versions required to run AutoSwap.

SRDF and open systems automation software

This section provides an overview of SRDF integration with EMC AutoStart to support automated restart of an application on an open systems server.

EMC AutoStart integrates with SRDF/S and SRDF/A to automate the restart of an application on an alternate local or remote open systems server. AutoStart also automates the failback and failover operations.

Figure 70 on page 174 shows AutoStart nodes in a two-site SRDF topology.

When a host is connected to the primary site, it writes to R1 devices. If the primary site becomes inaccessible, AutoStart initiates a failover and restarts the application at the R2 secondary site.



Figure 70 SRDF AutoStart typical configuration

AutoStart provides add-on application modules for Microsoft SQL Server 2005 and Exchange 2003, 2007, and 2010 as well as Oracle 10g and 11g on Windows, Linux, Solaris, AIX, HP-UX, and VMware ESX Server and guest operating systems, and Microsoft Hyper-V and guest operating systems.

Note: The add-on AutoStart application modules and data source modules require the base AutoStart product. The EMC Online Support site (https://support.EMC.com) contains more details about the operating system requirements and software versions required to run AutoStart.

SRDF and VMware environments

This section provides an overview of SRDF integration with VMware™ Site Recovery Manager to automate storage-based disaster restart operations in SRDF topologies.

Symmetrix storage systems implement the VMware-defined specifications as an application referred to as the storage replication adapter (SRA). The EMC SRDF Adapter is an SRA that enables VMware Site Recovery Manager to interact with Symmetrix systems to automate storage-based disaster restart operations in Symmetrix SRDF solutions.

The EMC SRDF Adapter can also address configurations in which data are spread across multiple storage systems or SRDF groups. The adapter must be installed on each Symmetrix system to facilitate the discovery of systems and replicated LUNs, and to initiate failover operations. It also ensures that the storage replication and the virtual machines are established properly.

Figure 71 on page 175 shows how SRDF integrates with VMware Site Recovery Manager to simplify recovery operations.



Figure 71 EMC SRDF and VMware SRM

Note: The EMC Online Support site (https://support.EMC.com) contains more details about the operating system requirements and software versions required to run VMware Site Recovery Manager.

SRDF and EMC FAST VP

EMC Symmetrix Fully Automated Storage Tiering for Virtual Pools (FAST VP) automates tiered storage strategies in Virtual Provisioning environments. It easily moves workloads between Symmetrix tiers as performance characteristics change over time, where a Symmetrix tier is a specification of a type of storage and a specification of a set of resources (disk groups/virtual pools) from which the storage is selected.

The following drive technologies are supported by Symmetrix VMAX Family systems:

- Flash drives
- Fibre Channel (FC) drives
- SAS drives
- SATA drives
- Federated Tier Storage (FTS) eDisks

FAST VP proactively monitors workloads on thin devices to identify the busy rate of the data and perform the following data movement, based on user-defined Symmetrix tiers and FAST policies:

- Move the most-used data to the fastest storage, such as Flash drives.
- Move the least-used data to the slower storage, such as SATA or FTS eDisks.
- Maintain the remaining data on Fibre Channel or SAS drives.

Data movement executed during these promotion or demotion activities is performed nondisruptively, without affecting business continuity and data availability. You can also use EMC host-based SRDF software to compress the least-used thin data. The EMC Online Support site (https://support.EMC.com) contains more details about FAST VP.

FAST VP is interoperable with SRDF solutions. The FAST VP software can act independently on both the local and remote VMAX Family systems. With Enginuity version 5876.82.57 or higher and EMC host-based SRDF software, you can coordinate data movements on both sides of the SRDF links in a two-site SRDF topology. With Enginuity version 5876.229.145 or higher, this capability is extended to multi-site SRDF topology.

You can enable or disable the coordination option. If you enable the coordination option with Solutions Enabler to send coordination information between the local and the remote systems, the following unbalanced configuration can be avoided:

- In an SRDF failover scenario, the remote system has different performance characteristics than the local production system, from which it failed over.
- In an SRDF/A environment, FAST VP data movements on the production R1 system could result in an unbalanced configuration between R1 and R2, where the performance characteristics of the R2 device are lower than that of the paired R1 device.

Requirements

To reduce the risk of SRDF/A session dropping, the following requirements apply to FAST VP:

- EMC recommends that you enable the SRDF/A device-level pacing option in the FAST VP environments. "Write pacing" on page 88 provides additional details.
- EMC recommends that you enable DSE to protect the SRDF/A session from a longer duration of unbalance. "Reserve Capacity" on page 87 provides additional details about DSE.

SRDF and EMC RecoverPoint

EMC RecoverPoint is a data replication solution designed to protect application data on heterogeneous SAN-attached servers and storage arrays, including Symmetrix arrays. Continuous Data Protection (CDP), Continuous Remote Replication (CRP), and Concurrent Local and Remote Replication (CLR) are RecoverPoint configurations that address different service level requirements.

With Enginuity version 5876.229.145, SRDF and RecoverPoint CDP can co-exist on the same source device in a two-site SRDF solution, which provides the ability to perform remote data protection with SRDF and logical data protection with RecoverPoint at the same time. The SRDF R1 devices are allowed to be tagged for RecoverPoint, and the RecoverPoint tagged devices are allowed to be configured as SRDF R1 devices.

This feature is supported for both static and dynamic SRDF devices, and is available in the following SRDF modes of operations:

- Synchronous
- Asynchronous
- Adaptive copy

The *EMC Solutions Enabler Symmetrix SRDF CLI Product Guide* provides details about controlling operations.

Requirements

Co-existence of SRDF and RecoverPoint CDP requires the following software:

- Enginuity version 5876.229.145 or higher
- Solutions Enabler version 7.6 or higher
- RecoverPoint version 4.0 or higher

Limitations

With Enginuity version 5876.229.145, SRDF and RecoverPoint CDP can co-exist on the same source device with the following limitations:

- The SRDF R1 devices must be configured in a two-site SRDF solution.
- The SRDF R1 BCV devices are not supported.
- If the SRDF R1 device is tagged as a RecoverPoint CDP source device, RecoverPoint consistency needs to be disabled before performing an SRDF failback operation or any operations that may cause data to flow from the R2 to R1 device.