

Dell EMC SRDF

Introduction

Rev 01

September 2019

Copyright © 2019 Dell Inc. or its subsidiaries. All rights reserved.

Dell 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.” DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Figures		7
Tables		9
	PREFACE	11
Chapter 1	Introduction	13
	What is SRDF?.....	14
	Disaster recovery.....	14
	High availability.....	16
	Data migration.....	17
	SRDF concepts.....	18
	SRDF device pairs.....	18
	SRDF device states.....	22
	Invalid tracks in SRDF pairs.....	27
	SRDF groups.....	28
	Dynamic devices.....	29
	SRDF modes of operation.....	30
	SRDF consistency.....	32
	Director boards, links, and ports.....	33
Chapter 2	Disaster recovery	35
	2-site configurations.....	36
	3-site configurations.....	39
	Concurrent SRDF solutions.....	41
	Cascaded SRDF solutions.....	42
	SRDF/Star solutions.....	43
	4-site configurations.....	51
	4-site FBA configurations.....	51
	4-site mainframe configurations (SRDF/SQAR).....	52
	SRDF recovery scenarios.....	54
	Planned failover (SRDF/S).....	54
	Unplanned failover	55
	Failback to the primary array.....	56
	Recovery for a large number of invalid tracks.....	56
	Temporary link loss.....	56
	Permanent link loss (SRDF/A).....	56
	SRDF/A session cleanup.....	56
	Failback from R2 devices.....	57
Chapter 3	High availability	59
	SRDF/Metro	60
	SRDF/Metro life cycle.....	62
	SRDF/Metro resilience.....	63
	Device Bias.....	63
	Array Witness.....	64
	Virtual Witness (vWitness).....	65

	Witness negotiation and selection.....	66
	Witness failure scenarios.....	67
	Mobility ID with ALUA.....	68
	Disaster recovery facilities.....	69
	Deactivate SRDF/Metro.....	70
	SRDF/Metro restrictions.....	71
Chapter 4	Data migration	73
	Introduction to data migration using SRDF.....	74
	Non-disruptive migration.....	74
	Migration from VMAX array.....	74
	Migration from VMAX All Flash or VMAX3.....	75
	Migrating data with concurrent SRDF.....	77
	Replacing R1 devices with new R1 devices.....	77
	Replacing R2 devices with new R2 devices.....	79
	Replacing R1 and R2 devices with new R1 and R2 devices.....	80
	Migration-only SRDF.....	81
	Device Migration operations requirements	82
Chapter 5	SRDF I/O operations	83
	SRDF write operations.....	84
	Write operations in synchronous mode.....	84
	Write operations in asynchronous mode.....	84
	Cycle switching in asynchronous mode.....	86
	Write operations in cascaded SRDF.....	89
	SRDF read operations.....	90
	Read operations if R1 local copy fails.....	90
	Read operations from R2 devices.....	90
	SRDF/A resilience and performance features.....	91
	Tunable cache.....	91
	SRDF/A cache data offloading	91
	Transmit Idle.....	92
	Write folding.....	92
	Write pacing	92
Chapter 6	Interfamily compatibility	95
	Overview.....	96
	SRDF supported features.....	96
Chapter 7	Management tools	99
	Solutions Enabler.....	100
	Unisphere.....	100
	SRDF/TimeFinder Manager for IBM i.....	101
	Mainframe management tools.....	102
	Mainframe Enablers.....	102
	Geographically Dispersed Disaster Restart (GDDR).....	103
Chapter 8	More information	105
	Solutions Enabler CLI.....	106
	Unisphere.....	106
	Mainframe Enablers.....	106
	GDDR.....	106

SRDF/TimeFinder Manager for IBM i..... 106
SRDF/Metro vWitness..... 106
SRDF Interfamily Compatibility..... 106
Storage arrays..... 107

FIGURES

1	Data replicated to one additional array.....	14
2	Data replicated to two arrays simultaneously.....	15
3	SRDF/Metro example.....	16
4	Migration example.....	17
5	R1 and R2 devices	18
6	R11 device in concurrent SRDF.....	19
7	R21 device in cascaded SRDF.....	20
8	R22 devices in cascaded and concurrent SRDF/Star.....	21
9	Host interface view and SRDF view of states.....	22
10	Concurrent SRDF topology.....	41
11	Cascaded SRDF topology.....	42
12	Concurrent SRDF/Star.....	44
13	Concurrent SRDF/Star with R22 devices.....	45
14	Cascaded SRDF/Star.....	46
15	R22 devices in cascaded SRDF/Star.....	47
16	Concurrent SRDF/Star with GDDR.....	48
17	Cascaded SRDF/Star with GDDR.....	49
18	SRDF 4-site FBA configuration.....	51
19	SRDF/SQAR with Autoswap environment.....	52
20	Planned failover: before personality swap.....	54
21	Planned failover: after personality swap.....	55
22	Failover to Site B, Site A and production host unavailable.....	55
23	SRDF/Metro.....	60
24	SRDF/Metro life cycle.....	62
25	SRDF/Metro Array Witness and groups.....	64
26	SRDF/Metro vWitness vApp and connections.....	65
27	SRDF/Metro Witness single failure scenarios.....	67
28	SRDF/Metro Witness multiple failure scenarios.....	68
29	Disaster recovery for SRDF/Metro.....	69
30	Migrating data and replacing the original primary array (R1).....	78
31	Migrating data and removing the original secondary array (R2).....	79
32	Migrating data and replacing the original primary (R1) and secondary (R2) arrays.....	80
33	Write I/O flow: simple synchronous SRDF.....	84
34	SRDF/A SSC cycle switching – multi-cycle mode.....	86
35	SRDF/A SSC cycle switching – legacy mode.....	87
36	SRDF/A MSC cycle switching – multi-cycle mode.....	88
37	Write commands to R21 devices.....	89

TABLES

1	SRDF pair states.....	23
2	R1 device accessibility.....	25
3	R2 device accessibility.....	25
4	Possible SRDF device and link state combinations	26
5	SRDF 2-site configurations.....	36
6	SRDF multi-site solutions	39
7	Limitations of the migration-only mode	81
8	SRDF features by hardware platform/operating environment.....	96
9	Unisphere tasks.....	100

PREFACE

As part of an effort to improve its product lines, Dell 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 Dell EMC technical support professional if a product does not function properly or does not function as described in this document.

Note: This document was accurate at publication time. Go to Dell EMC Online Support (<https://support.emc.com>) to ensure that you are using the latest version of this document.

Purpose

This document provides an introduction to the Symmetrix Remote Data Facility (SRDF) and its uses in disaster recovery, high availability, and data migration applications.

Audience

This document is intended for Dell EMC customers who want an overview of SRDF and its applications.

Related documentation

Information on the storage arrays that SRDF runs on is in the following publications:

- Dell EMC PowerMax Family Product Guide
- Dell EMC VMAX All Flash Product Guide for VMAX 250F, 450F, 850F, 950F with HYPERMAX OS
- EMC VMAX3 Family Product Guide for VMAX 100K, VMAX 200K, VMAX 400K with HYPERMAX OS

Where to get help

Dell EMC support, product, and licensing information can be obtained as follows:

Product information

Dell EMC technical support, documentation, release notes, software updates, or information about Dell EMC products can be obtained at <https://support.emc.com> (registration required) or <https://www.dell EMC.com/en-us/documentation/vmax-all-flash-family.htm>.

Technical support

To open a service request through the Dell EMC Online Support (<https://support.emc.com>) site, you must have a valid support agreement. Contact your Dell EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

Technical support

Dell EMC offers a variety of support options.

- Support by Product — Dell EMC offers consolidated, product-specific information on the Web through the Dell EMC Online Support site. The Support by Product web pages (<https://support.EMC.com/products>) 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 Dell EMC Live Chat.

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

Your comments

Your suggestions help us improve the accuracy, organization, and overall quality of the documentation. Send your comments and feedback to: VMAXContentFeedback@emc.com

CHAPTER 1

Introduction

This chapter introduces SRDF, lists its uses, and defines SRDF's concepts.

- [What is SRDF?](#)14
- [SRDF concepts](#)..... 18

What is SRDF?

The Symmetrix Remote Data Facility (SRDF) maintains real-time (or near real-time) copies of data on a production storage array at one or more remote storage arrays. SRDF has three primary applications:

- Disaster recovery
- High availability
- Data migration


This is an introduction to SRDF, its uses, configurations, and terminology. The rest of this section provides a summary of the applications for SRDF ([SRDF device pairs](#) on page 18 explains the device naming conventions used in the diagrams).

Disaster recovery

In disaster recovery, SRDF maintains a real-time copy of the data of one or more devices on a storage array at one or more additional arrays. This provides the means to restart host applications should the main array become unavailable for any reason. Typically, each array is on a separate site from all the others in the SRDF configuration.


For example, this diagram shows a configuration where data is replicated to one additional array:

Figure 1 Data replicated to one additional array



The next example shows data being replicated to two additional arrays simultaneously. This improves redundancy and data security:

Figure 2 Data replicated to two arrays simultaneously



[Disaster recovery](#) on page 35 describes SRDF's disaster recovery facilities, and system configurations in more detail.

High availability


In other SRDF configurations, devices on the primary array are Read/Write accessible to the application host while devices on the additional arrays are Read Only/Write Disabled. However, in a SRDF high availability configuration:

- Devices on the additional array are Read/Write accessible to the application host.
- The application host can write to both sides of the device pair.
- The devices on the additional array assume the same external identity (such as geometry and device identifier) as the devices on the primary array.
- The arrays are up to 200 km (125 miles) apart and are typically in separate fault domains for extra resilience.

This shared identity means that the devices appear as to the application host as a single, virtual device across the two arrays. Using two devices improves the availability of the data they contain. One device can become unavailable without impacting on the host application as the second device continues to operate.

Such a configuration is known as SRDF/Metro which can be deployed in a single, multi-pathed host environment, or in a clustered environment as this diagram shows:

Figure 3 SRDF/Metro example



[High availability](#) on page 59 describes SRDF's high availability capabilities and system configurations in more detail.

Open systems (FBA) only

SRDF/Metro is available in open systems (FBA) and IBM i D910¹ environments only. The mainframe environment has its own high availability called AutoSwap. The publications listed in the Mainframe and GDDR sections of [More information](#) on page 105 contain details of AutoSwap and its capabilities.

1. IBM i D910 requires PowerMaxOS 5978.444.444 or later


Data migration

The data migration capabilities of SRDF enable devices on either side of a 2-array configuration to be replaced with new devices without interrupting disaster recovery operations. To do this, the configuration is enhanced with a third array that contains the new devices and data is replicated to that in addition to the normal operation. Once the replication is complete the devices being replaced can be taken out of the configuration leaving one of the original arrays and the new one.

For example, this diagram shows the replacement of R2 devices with new devices using SRDF migration facilities:

- The initial 2-array topology
- The interim 3-array topology
- The final 2-array topology

Figure 4 Migration example



[Data migration](#) on page 73 describes SRDF's migration capabilities and system configurations in more detail.

SRDF concepts

SRDF device pairs

An SRDF device is a logical device paired with another logical device that resides in a second array. The arrays are connected by SRDF links.

Encapsulated Data Domain devices used for ProtectPoint cannot be part of an SRDF device pair.

Note: ProtectPoint has been renamed to Storage Direct and it is included in the PowerProtect, Data Protection Suite for Apps, or Data Protection Suite Enterprise Edition software.

R1 and R2 devices


A R1 device is the member of the device pair at the source (production) site. R1 devices are generally Read/Write accessible to the application host.

A R2 device is the member of the device pair at the target (remote) site. During normal operations, host I/O writes to the R1 device are mirrored over the SRDF links to the R2 device. In general, data on R2 devices is not available to the application host while the SRDF relationship is active. In SRDF synchronous mode, however, an R2 device can be in Read Only mode that allows a host to read from the R2.

In a typical environment:

- The application production host has Read/Write access to the R1 device.
- An application host connected to the R2 device has Read Only (Write Disabled) access to the R2 device.


Figure 5 R1 and R2 devices



R11 devices

R11 devices operate as the R1 device for two R2 devices. Links to both R2 devices are active. R11 devices are typically used in 3-site concurrent configurations where data on the R11 site is mirrored to two secondary (R2) arrays:

Figure 6 R11 device in concurrent SRDF




R21 devices

R21 devices have a dual role and are used in cascaded 3-site configurations where:

- Data on the R1 site is synchronously mirrored to a secondary (R21) site, and then
- Asynchronously mirrored from the secondary (R21) site to a tertiary (R2) site:

Figure 7 R21 device in cascaded SRDF



The R21 device acts as a R2 device that receives updates from the R1 device, and as a R1 device that sends updates to the R2 device.

When the R1->R21->R2 SRDF relationship is established, no host has write access to the R21 device.

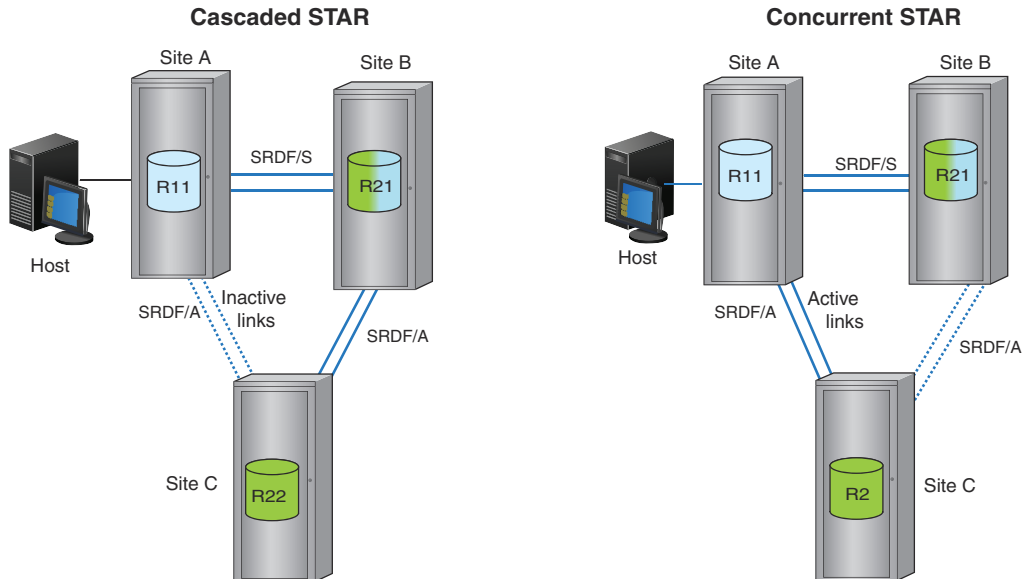
In arrays that run Enginuity, the R21 device can be diskless. That is, it consists solely of cache memory and does not have any associated storage device. It acts purely to relay changes in the R1 device to the R2 device. This capability requires the use of thick devices. Systems that run PowerMaxOS or HYPERMAX OS contain thin devices only, so setting up a diskless R21 device is not possible on arrays running those environments.

R22 devices

R22 devices:

- Have two R1 devices, only one of which is active at a time.
- Are typically used in cascaded SRDF/Star and concurrent SRDF/Star configurations to decrease the complexity and time required to complete failover and failback operations.
- Let you recover without removing old SRDF pairs and creating new ones.

Figure 8 R22 devices in cascaded and concurrent SRDF/Star




SRDF device states

An SRDF device's state is determined by a combination of two views; host interface view and SRDF view, as shown in this diagram.

Figure 9 Host interface view and SRDF view of states

Host interface view
(Read/Write, Read Only (Write Disabled), Not Ready)

Open systems host environment



SRDF view
(Ready, Not Ready, Link Blocked)

Host interface view

The host interface view is the SRDF device state as seen by the application host.

R1 device states

An R1 device presents one of the following states to a host connected to it:

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

R2 device states

An R2 device presents one of the following states to a host connected to it:

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

SRDF view

The SRDF view is composed of the SRDF state and internal SRDF device state. These states indicate whether the device is available to send data across the SRDF links, and able to receive software commands.

R1 device states

An R1 device can have the following states for SRDF operations:

- **Ready**—The R1 device is ready for SRDF operations. The R1 device is able to send data across the SRDF links. True even if local mirror(s) of the R1 device are Not Ready for I/O operations.
- **Not Ready (SRDF mirror Not Ready)**—The R1 device is Not Ready for SRDF operations.

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

An R2 device can have the following states for SRDF operations:

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

Device pair states

Device pairs that are part of an SRDF operation need to be in the correct state. This table lists the states that a device pair can be in.

Table 1 SRDF pair states

Pair State	Description
SyncInProg	Synchronization is currently in progress between the R1 and the R2 devices. There are existing invalid tracks between the two pairs, and the logical links between both sides of an SRDF pair are up.
Synchronized	The R1 and the R2 are currently in a synchronized state. The same content exists on the R2 as the R1, and there are no invalid tracks between the two pairs.
Split	The R1 and the R2 are currently ready to their hosts. However, the links are not ready or, are write disabled.
Failed Over	The R1 is not ready or write disabled. Operations have been failed over to R2.
R1 Updated	The R1 is not ready or write disabled to the host.

Table 1 SRDF pair states (continued)

Pair State	Description
	There are no local invalid tracks on the R1 side, and the links are ready or write disabled.
R1 UpdInProg	The R1 is not ready or write disabled to the host. There are invalid local (R1) tracks on the source side, so data is being copied from the R2 to the R1 device, and the links are ready.
ActiveActive	The R1 and the R2 are currently in the default SRDF/Metro configuration which uses an Array Witness or Virtual Witness: <ul style="list-style-type: none"> • There are no invalid tracks between the two pairs. • The R1 and the R2 are Ready (RW) to the hosts.
ActiveBias	The R1 and the R2 are currently in an SRDF/Metro configuration using bias: <ul style="list-style-type: none"> • The user has specified “use bias” during the establish/restore action or the desired Witness is not available. • There are no invalid tracks between the two pairs. • The R1 and the R2 are Ready (RW) to the hosts.
Suspended	The SRDF links have been suspended and are not ready or write disabled. If the R1 is ready while the links are suspended, any I/O accumulates as invalid tracks owed to the R2.
Partitioned	The SYMAPI is currently unable to communicate through the corresponding SRDF path to the remote array. The Partitioned state may apply to devices within an RA group. For example, if SYMAPI is unable to communicate to a remote array from an RA group, devices in that RA group will be marked as being in the Partitioned state. A half pair and a duplicate pair are also reported as Partitioned.
Mixed	A composite SYMAPI device group SRDF pair state. There are different SRDF pair states within a device group.
Invalid	This is the default state when no other SRDF state applies. <ul style="list-style-type: none"> • The combination of the R1 device, the R2 device, and the SRDF link states do not match any other pair state. • This state may occur if there is a problem at the disk director level.
Consistent	The R2 SRDF/A capable devices are in a consistent state. The consistent state signifies the normal state of operation for device pairs operating in asynchronous mode.

Table 1 SRDF pair states (continued)

Pair State	Description
Transmit Idle	The SRDF/A session cannot send data in the transmit cycle over the link because the link is unavailable.

R1/R2 device accessibility

Accessibility of a SRDF device to the application host depends on both the host and the array view of the SRDF device state.

[Table 2](#) on page 25 and [Table 3](#) on page 25 list application host accessibility for R1 and R2 devices.

Table 2 R1 device accessibility

Host interface state	SRDF R1 state	Accessibility
Read/Write	Ready	Read/Write
	Not Ready	Depends on R2 device availability
Read Only	Ready	Read Only
	Not Ready	Depends on R2 device availability
Not Ready	Any	Unavailable

Table 3 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)	Ready	Read Only
	Not Ready	Read Only
Not Ready	Any	Unavailable

SRDF device and link state combinations

Control actions on an SRDF pair may change the SRDF pair state.

Additionally, the state of a device can change if its front-end or back-end directors change in the SRDF links.

The following table lists:

- SRDF pair states resulting from the combination of the states of the source and target devices and the SRDF links.
- The possible R1 or R2 invalid tracks for each SRDF pair state.

Table 4 Possible SRDF device and link state combinations

SRDF pair state	Source (R1) SRDF state	SRDF link state	Target (R2) SRDF state	R1 or R2 invalid tracks
Synchronized	Ready (RW)	Ready (RW)	Not Ready or WD	0
Failed Over	Not Ready or WD	Not Ready	Ready (RW)	—
R1 Updated	Not Ready or WD	Ready (RW) or WD	Ready (RW)	0 ^a
R1 UpdInProg	Not Ready or WD	Ready (RW) or WD	Ready (RW)	>0 ^a
ActiveActive	Ready (RW)	Ready (RW)	Ready (RW)	0
ActiveBias	Ready (RW)	Ready (RW)	Ready (RW)	0
Split	Ready (RW)	Not Ready or WD	Ready (RW)	—
SyncInProg	Ready (RW)	Ready (RW)	Not Ready or WD	>0
Suspended	Any status ^b	Not Ready or WD	Not Ready or WD	—
Partitioned ^c	Any status	Not Ready	Not Available	—
Partitioned ^d	Not Available	Not Ready	Any status	—
Mixed	e	e	e	—
Invalid ^e	Any status ^f	Any status	Any status	—
Consistent	Ready (RW) ^f	Ready (RW)	Not Ready or WD	0 or >0 ^a
Transmit Idle	Ready (RW) ^f	Ready (RW)	Not Ready or WD	—

a. Refers to invalid local (R1) tracks on source.

b. Any status value is possible (Ready, Not Ready, Write Disabled, or Not Available).

c. Viewed from the host locally connected to the source (R1) device

d. Viewed from the host locally connected to the target (R2) device.

e. When no other SRDF states apply, the state defaults to Invalid.

f. The combination of source SRDF, SRDF links, and target SRDF statuses does not match any other SRDF state; therefore, the SRDF state is considered Invalid.

Invalid tracks in SRDF pairs

Invalid tracks are tracks that are not synchronized between the two devices in a SRDF pair. They occur when either member of the pair cannot communicate with its partner; for example due to a failure of the SRDF link between the storage arrays. On both sides of the configuration, the storage arrays record the number of tracks that are owed to the other side.

Once the two devices can communicate once again, the invalid tracks need resolving between the pair. There are two ways to resolve the tracks:

- Copy the modified R1 tracks to the R2 side.
Any tracks that were modified on the R2 side are overwritten with the data for the corresponding tracks on the R1 side.
- Copy the modified R2 tracks to the R1 side.
Any tracks that were modified on the R1 side are overwritten with the data for the corresponding tracks on the R2 side.

Example: Unavailable SRDF link or R2 device

Here, the SRDF link is unavailable for some reason, the R2 device is unavailable, or both the link and the R2 device are unavailable. The R1 device, however, remains write accessible to the application host. While this situation exists, the R1 device receives I/O from the application host, and invalid tracks accumulate on the R1 array.

Once the SRDF link and the R2 device are available again, the array containing the R1 array sends the invalid tracks to the R2 device so that the two devices are synchronized once more.

Example: R1 unavailable

Here, the R1 device has become unavailable for some reason. To maintain service to the application host, processing is moved to the R2 device. That is, the R2 device is made write accessible to the application host, and it receives I/O from that host. While this situation exists, invalid tracks accumulate at the R2 array.

Once the R1 device is available again, the array containing the R2 device sends the invalid tracks to the R1 device. Once the two devices are fully synchronized, processing returns to the R1 device and the R2 device is made write protected to the application host.

SRDF groups

An SRDF group defines the logical relationship between SRDF devices and directors on both sides of a SRDF pair.

Group properties

The properties of an SRDF group are:

- Label (name)
- Set of ports on the local array used to communicate over the SRDF links
- Set of ports on the remote array used to communicate over the SRDF links
- Local group number
- Remote group number
- One or more pairs of devices

The devices in the group share the ports and associated CPU resources of the port's directors.

Advanced properties of an SRDF group include:

- Link Limbo mode – The amount of time that the array's operating environment waits after the SRDF link goes down before updating the link's status.
- Link Domino mode – Specifies whether to force SRDF devices into the Not Ready state to the application host if, for example, host I/Os cannot be delivered across the SRDF link.
- Autolink recovery – Specifies whether SRDF automatically restores the SRDF links when they become available again after an earlier failure.
- Compression – Specifies whether to use compression when sending data over the SRDF links. Both hardware and software compression are available and can be used independently or together.

Types of group

There are two types of SRDF group:

- Static
- Dynamic

Static groups are defined in the local array's configuration file. Dynamic groups are defined using SRDF management tools and their properties stored in the array's cache memory.

On arrays running PowerMaxOS or HYPERMAX OS all SRDF groups are dynamic.

Group membership

An SRDF device is a member of as many SRDF groups as there are mirrors of that device. So, in a simple, 2-site configuration (see [Figure 5](#) on page 18) that consists of R1 and R2 devices, each device is a member of one group. In a concurrent SRDF configuration (see [Figure 6](#) on page 19), the R11 device is a member of two groups, one for each R2 mirror. The R2 devices are each in a single group.

Dynamic devices

Dynamic SRDF devices are SRDF devices that allow flexible control over the SRDF solution. You can configure and control dynamic SRDF devices using the SRDF management tools. Dynamic device attributes are stored in the mirrored and protected array cache memory.

SRDF management tools can modify the attributes of dynamic SRDF devices in these ways:

- 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 SRDF groups.

R1/R2 personality swap

SRDF devices can dynamically swap “personality” between R1 and R2. After a personality swap:

- The R1 in the device pair becomes the R2 device, and
- The R2 becomes the R1 device.

Swapping R1/R2 personalities allows the application to be restarted at the remote site if an application fails at the production site. After a swap, the R2 side (now R1) can control operations while being remotely mirrored at the primary (now R2) site.

An R1/R2 personality swap is not possible:

- 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 diskless R11 or R22 devices are not valid end states.
- If the device to be swapped is the target device of any TimeFinder or Dell EMC Compatible flash operations.

SRDF modes of operation

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

- How R1 devices are remotely mirrored across the SRDF links.
- How I/Os are processed.
- When the host receives acknowledgment of a write operation relative to when the write is replicated.
- When writes “owed” between partner devices are sent across the SRDF links.

The mode of operation may change in response to control operations or failures:

- The primary mode (synchronous or asynchronous) is the configured mode of operation for a given SRDF device, range of SRDF devices, or an SRDF group.
- The secondary mode is adaptive copy. Adaptive copy mode moves large amounts of data quickly with minimal host impact. Adaptive copy mode does not provide restartable data images at the secondary site until no new writes are sent to the R1 device and all data has finished copying to the R2.

Use adaptive copy mode to synchronize new SRDF device pairs or to migrate data to another array. When the synchronization or migration is complete, you can revert to the configured primary mode of operation.

Synchronous mode

SRDF/Synchronous (SRDF/S) maintains a real-time mirror image of data between the R1 and R2 devices. The recommended distance between the devices is 200 km (125 miles) or less because application latency may rise to unacceptable levels at longer distances.

Host writes are written simultaneously to both arrays in real time before the application I/O completes. Acknowledgments are not sent to the host until the data is stored in cache on both arrays.

[Write operations in synchronous mode](#) on page 84 and [SRDF read operations](#) on page 90 have more information about I/O operations in synchronous mode.

Asynchronous mode

SRDF/Asynchronous (SRDF/A) maintains a dependent-write consistent copy between the R1 and R2 devices across any distance with no impact to the application.

Host writes are collected for a configurable interval into “delta sets”. Delta sets are transferred to the remote array in timed cycles.

SRDF/A operations vary depending on whether the SRDF session mode is single or multi-session with Multi Session Consistency (MSC) enabled:

- For single SRDF/A sessions, cycle switching is controlled by the array's operating environment. Each session is controlled independently, whether it is in the same or multiple arrays.
- For multiple SRDF/A sessions in MSC mode, multiple SRDF groups are in the same SRDF/A MSC session. Cycle switching is controlled by SRDF host software to maintain consistency.

[SRDF/A MSC cycle switching](#) on page 87 has more information on I/O operations in asynchronous mode.

Adaptive copy modes

Adaptive copy modes:

- Transfer large amounts of data without impact on the host.
- Transfer data during data center migrations and consolidations, and in data mobility environments.
- Allow the R1 and R2 devices to be out of synchronization by up to a user-configured maximum *skew value*. If the maximum skew value is exceeded, SRDF starts the synchronization process to transfer updates from the R1 to the R2 devices.
- Are secondary modes of operation for SRDF/S. The R1 devices revert to SRDF/S when the maximum skew value is reached and remain in SRDF/S until the number of tracks out of synchronization is lower than the maximum skew.

There are two types of adaptive copy mode:

- [Adaptive copy disk](#) on page 31
- [Adaptive copy write pending](#) on page 31

Adaptive copy disk


In adaptive copy disk mode, write requests accumulate on the R1 device (not in cache). A background process sends the outstanding write requests to the corresponding R2 device. 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, or
- The write requests exceed 50 percent of the primary or secondary array write pending space.

Adaptive copy write pending

In adaptive copy write pending mode, write requests accumulate in cache on the primary array. A background process sends the outstanding write requests to the corresponding R2 device.

Adaptive copy write-pending mode reverts to the primary mode if the device, cache partition, or system write pending limit is near, regardless of whether the maximum skew value specified for each device is reached.

 **Note:** Adaptive copy write pending mode is not available when the R1 side of an SRDF device pair is on an array running PowerMaxOS or HYPERMAX OS.

Domino modes

Under typical conditions, when one side of a device pair becomes unavailable, new data written to the device is marked for later transfer. When the device or link is restored, the two sides synchronize.

Domino modes force SRDF devices into the Not Ready state to the host if one side of the device pair becomes unavailable.

Domino mode can be enabled/disabled for any:

- Device (domino mode) – If the R1 device cannot successfully mirror data to the R2 device, the next host write to the R1 device causes the device to become Not Ready to the host connected to the primary array.
- SRDF group (link domino mode) – If the last available link in the SRDF group fails, the next host write to any R1 device in the SRDF group causes all R1 devices in the SRDF group become Not Ready to their hosts.

Link domino mode is set at the SRDF group level and only impacts devices where the R1 is on the side where it is set.

Geometry Compatibility Mode

In Enginuity 5876, the track size of an FBA device is 64 KB, while in PowerMaxOS 5978 and HYPERMAX OS 5977 the track size is 128 KB. So an array running PowerMaxOS or HYPERMAX OS cannot create a device that is the same size as a device that has an odd of cylinders on an array running Enginuity in a mixed SRDF configuration. However, SRDF requires that the devices in a device pair are the same size.

PowerMaxOS and HYPERMAX OS manage the difference in size automatically using the device attribute Geometry Compatibility Mode (GCM). A device with GCM set is presented as being half a cylinder smaller than its configured size. This enables full functionality in a mixed configuration for SRDF, TimeFinder, SnapVX, and TimeFinder emulations (TimeFinder Clone, TimeFinder VP Snap, and TimeFinder/Mirror) and ORS.

The GCM attribute can be set in two ways:

- Automatically on a target device when it is on an array running PowerMaxOS or HYPERMAX OS and the source device is on an array running Enginuity 5876
- Manually using the Solutions Enabler CLI, Mainframe Enablers SRDF Host Component, or Unisphere

Notes:

- Do not set GCM on devices that are mounted and under the control of a Local Volume Manager (LVM).
- Clear the GCM flag before mapping the device to a host. Otherwise, to clear the attribute, the device must be unmapped from the host which results in a data outage.
- The GCM setting for a device cannot be changed when the target of the data device is already part of another replication session.

SRDF consistency

Many applications (in particular, DBMS), use dependent write logic to ensure data integrity in the event of a failure. A dependent write is a write that is not issued by the application unless some prior I/O has completed. If the writes are out of order, and an event such as a failure, or a creation of a point in time copy happens at that exact time, unrecoverable data loss may occur.

An SRDF consistency group (SRDF/CG) is comprised of SRDF devices with consistency enabled.

SRDF consistency groups preserve the dependent-write consistency of devices within a group by monitoring data propagation from source devices to their corresponding target devices. If consistency is enabled, and SRDF detects any write I/O to a R1 device that cannot communicate with its R2 device, SRDF suspends the remote mirroring for all devices in 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.

SRDF consistency allows you to quickly recover from certain types of failure or physical disasters by retaining a consistent, DBMS-restartable copy of your database.

SRDF consistency group protection is available for both SRDF/S and SRDF/A.

Director boards, links, and ports

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

 **Note:** Two or more SRDF links per SRDF group are required for redundancy and fault tolerance.

The relationship between the resources on a director (CPU cores and ports) varies depending on the operating environment.

PowerMaxOS and HYPERMAX OS

On arrays running PowerMaxOS or HYPERMAX OS :

- The relationship between the SRDF emulation and resources on a director is configurable:
 - One director/multiple CPU cores/multiple ports
 - Connectivity (ports in the SRDF group) is independent of compute power (number of CPU cores). You can change the amount of connectivity without changing compute power.
- Each director has up to 16 front end ports, any or all of which can be used by SRDF. Both the SRDF Gigabit Ethernet and SRDF Fibre Channel emulations can use any port.
- The data path for devices in an SRDF group is not fixed to a single port. Instead, the path for data is shared across all ports in the group.

Mixed configurations: PowerMaxOS or HYPERMAX OS and Enginuity 5876

For configurations where one array is running Enginuity 5876, and the other array is running PowerMaxOS or HYPERMAX OS, the following rules apply:

- On the 5876 side, an SRDF group can have the full complement of directors, but no more than 16 ports on the PowerMaxOS or HYPERMAX OS side.
- You can connect to 16 directors using one port each, 2 directors using 8 ports each or any other combination that does not exceed 16 per SRDF group.

CHAPTER 2

Disaster recovery

This chapter provides more detail on the disaster recovery configurations of SRDF.

- [2-site configurations](#)..... 36
- [3-site configurations](#)..... 39
- [4-site configurations](#)..... 51
- [SRDF recovery scenarios](#)..... 54

2-site configurations

This table shows the 2-site configurations for SRDF.

Table 5 SRDF 2-site configurations

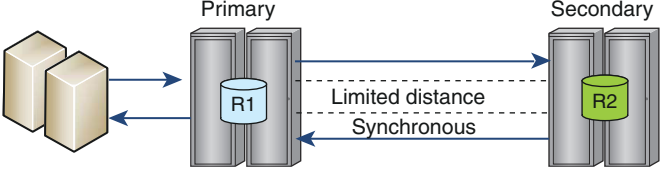
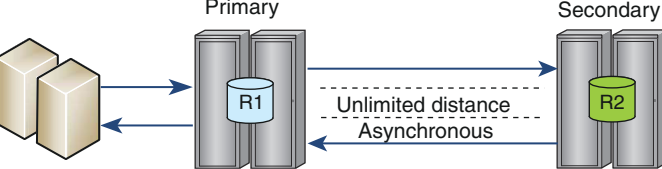
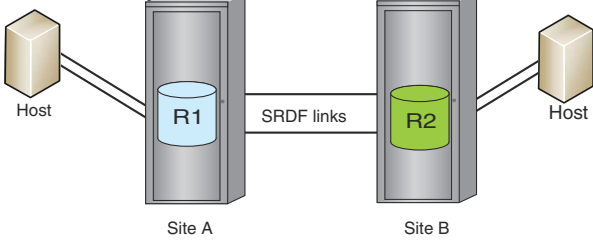
Solution highlights	Site topology
<p>SRDF/Synchronous (SRDF/S) Maintains a real-time copy of production data at a physically separated array.</p> <ul style="list-style-type: none"> • No data exposure • Ensured consistency protection with SRDF/Consistency Group • Recommended maximum distance of 200 km (125 miles) between arrays as application latency may rise to unacceptable levels at longer distances 	 <p>The diagram shows two sites: Primary and Secondary. The Primary site contains a host and a storage array labeled R1. The Secondary site contains a storage array labeled R2. A dashed line between the two sites is labeled 'Limited distance'. A solid double-headed arrow between R1 and R2 is labeled 'Synchronous'.</p>
<p>SRDF/Asynchronous (SRDF/A) Maintains a dependent-write consistent copy of the data on a remote secondary site. The copy of the data at the secondary site is seconds behind the primary site.</p> <ul style="list-style-type: none"> • RPO seconds before the point of failure • Unlimited distance 	 <p>The diagram shows two sites: Primary and Secondary. The Primary site contains a host and a storage array labeled R1. The Secondary site contains a storage array labeled R2. A dashed line between the two sites is labeled 'Unlimited distance'. A solid double-headed arrow between R1 and R2 is labeled 'Asynchronous'.</p>
<p>SRDF/Data Mobility (SRDF/DM) This example shows an SRDF/DM topology and the I/O flow in adaptive copy mode.</p> <ul style="list-style-type: none"> • The host write I/O is received in cache in Site A • The host emulation returns a positive acknowledgment to the host • The SRDF emulation transmits the I/O across the SRDF links to Site B • Once data is written to cache in Site B, the SRDF emulation in Site B returns a positive acknowledgment to Site A <p>Operating Notes:</p> <ul style="list-style-type: none"> • The maximum skew value set at the device level in SRDF/DM solutions must be equal or greater than 100 tracks • SRDF/DM is only for data replication or migration, not for disaster restart solutions 	 <p>The diagram shows two sites, Site A and Site B. Site A has a Host connected to a storage array R1. Site B has a Host connected to a storage array R2. The two storage arrays are connected by a line labeled 'SRDF links'.</p> <p>Note: Data may be read from the drives to cache before it is transmitted across the SRDF links, resulting in propagation delays.</p>

Table 5 SRDF 2-site configurations (continued)

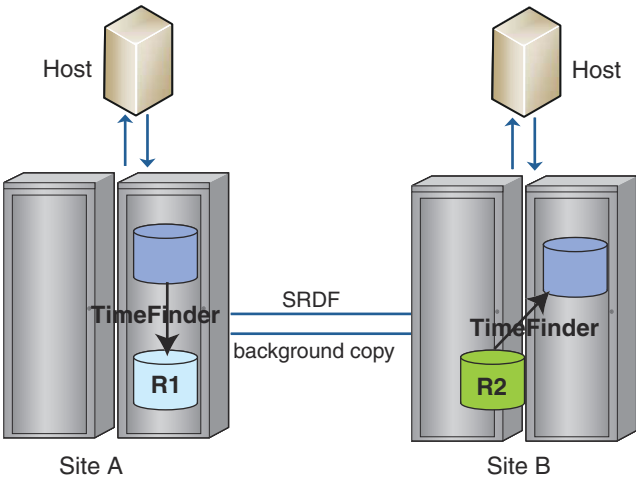
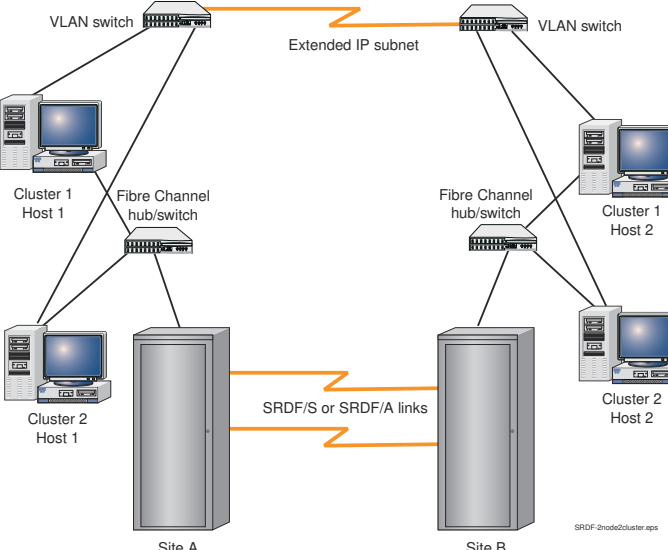
Solution highlights	Site topology
<p>SRDF/Automated Replication (SRDF/AR)</p> <ul style="list-style-type: none"> Combines SRDF and TimeFinder to optimize bandwidth requirements and provide a long-distance disaster restart solution. Operates in 2-site solutions that use SRDF/DM in combination with TimeFinder. 	 <p>The diagram illustrates a two-site SRDF/AR configuration. At Site A, a Host is connected to a TimeFinder component, which is connected to a primary storage R1. At Site B, a Host is connected to a TimeFinder component, which is connected to a secondary storage R2. SRDF links connect the TimeFinder components, and a background copy link connects R1 to R2.</p>
<p>SRDF/Cluster Enabler (CE)</p> <ul style="list-style-type: none"> Integrates SRDF/S or SRDF/A with Microsoft Failover Clusters (MSCS) to automate or semi-automate site failover. Complete solution for restarting operations in cluster environments (MSCS with Microsoft Failover Clusters) Expands the range of cluster storage and management capabilities while ensuring full protection of the SRDF remote replication. 	 <p>The diagram illustrates a two-site SRDF/Cluster Enabler (CE) configuration. Site A has Cluster 1 Host 1 and Cluster 2 Host 1 connected to a Fibre Channel hub/switch, which is connected to a VLAN switch. Site B has Cluster 1 Host 2 and Cluster 2 Host 2 connected to a Fibre Channel hub/switch, which is connected to a VLAN switch. SRDF/S or SRDF/A links connect the Fibre Channel hub/switches. An Extended IP subnet connects the VLAN switches.</p>

Table 5 SRDF 2-site configurations (continued)

Solution highlights	Site topology
<p>SRDF and VMware Site Recovery Manager Completely automates storage-based disaster restart operations for VMware environments in SRDF topologies.</p> <ul style="list-style-type: none"> • The Dell EMC SRDF Adapter enables VMware Site Recovery Manager to automate storage-based disaster restart operations in SRDF solutions. • Can address configurations in which data are spread across multiple storage arrays or SRDF groups. • Requires that the adapter is installed on each array to facilitate the discovery of arrays and to initiate failover operations. • Implemented with: <ul style="list-style-type: none"> ▪ SRDF/S ▪ SRDF/A ▪ SRDF/Star ▪ TimeFinder 	<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p><u>Protection side</u></p> <ul style="list-style-type: none"> • vCenter and SRM Server • Solutions Enabler software </div> <div style="width: 45%;"> <p><u>Recovery side</u></p> <ul style="list-style-type: none"> • vCenter and SRM Server • Solutions Enabler software </div> </div> <p>The diagram illustrates a two-site SRDF configuration. At the top, two ESX Servers are shown, one for the Protection side and one for the Recovery side. Both are connected to their respective IP Networks and have Solutions Enabler software configured as a SYMAPI server. Below the ESX Servers, there are two SAN Fabrics on each side, connected to storage arrays. The storage arrays at Site A (primary) and Site B (secondary) are connected via SRDF mirroring. The ESX Servers are also connected to their respective SAN Fabrics.</p>

3-site configurations

This table shows the 3-site configurations for SRDF.

Table 6 SRDF multi-site solutions

Solution highlights	Site topology
<p>SRDF/Automated Replication (SRDF/AR)</p> <ul style="list-style-type: none"> Combines SRDF and TimeFinder to optimize bandwidth requirements and provide a long-distance disaster restart solution. Operates in 3-site solutions that use a combination of SRDF/S, SRDF/DM, and TimeFinder. 	
<p>Concurrent SRDF</p> <p>3-site disaster recovery and advanced multi-site business continuity protection.</p> <ul style="list-style-type: none"> Data on the primary site is concurrently replicated to 2 secondary sites. Replication to remote site can use SRDF/S, SRDF/A, or adaptive copy. 	
<p>Cascaded SRDF</p> <p>3-site disaster recovery and advanced multi-site business continuity protection.</p> <ul style="list-style-type: none"> 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. First “hop” is SRDF/S. Second hop is SRDF/A. 	

Table 6 SRDF multi-site solutions (continued)

Solution highlights	Site topology
<p>SRDF/Star 3-site data protection and disaster recovery with zero data loss recovery, business continuity protection and disaster-restart.</p> <ul style="list-style-type: none"> • Available in 2 configurations: <ul style="list-style-type: none"> ▪ Cascaded SRDF/Star ▪ Concurrent SRDF/Star • Differential synchronization allows rapid reestablishment of mirroring among surviving sites in a multi-site disaster recovery implementation. • Implemented using SRDF consistency groups (CG) with SRDF/S and SRDF/A. 	<p>The diagram illustrates two SRDF/Star configurations across three sites: Site A, Site B, and Site C.</p> <p>Cascaded SRDF/Star: Site A (R11) is connected to Site B (R21) via SRDF/S. Site B (R21) is connected to Site C (R2/R22) via SRDF/A. A dashed arrow indicates SRDF/A (recovery) from Site B to Site C.</p> <p>Concurrent SRDF/Star: Site A (R11) is connected to Site B (R21) via SRDF/S. Site A (R11) is also connected to Site C (R2/R22) via SRDF/A. A dashed arrow indicates SRDF/A (recovery) from Site B to Site C.</p>

Concurrent SRDF solutions

Concurrent SRDF is a 3-site disaster recovery solution using R11 devices that replicate to two R2 devices. The two R2 devices operate independently but concurrently using any combination of SRDF modes:

- Concurrent SRDF/S to both R2 devices if the R11 site is within synchronous distance of the two R2 sites.
- Concurrent SRDF/A to sites located at extended distances from the workload site.

You can restore the R11 device from either of the R2 devices. You can restore both the R11 and one R2 device from the second R2 device.

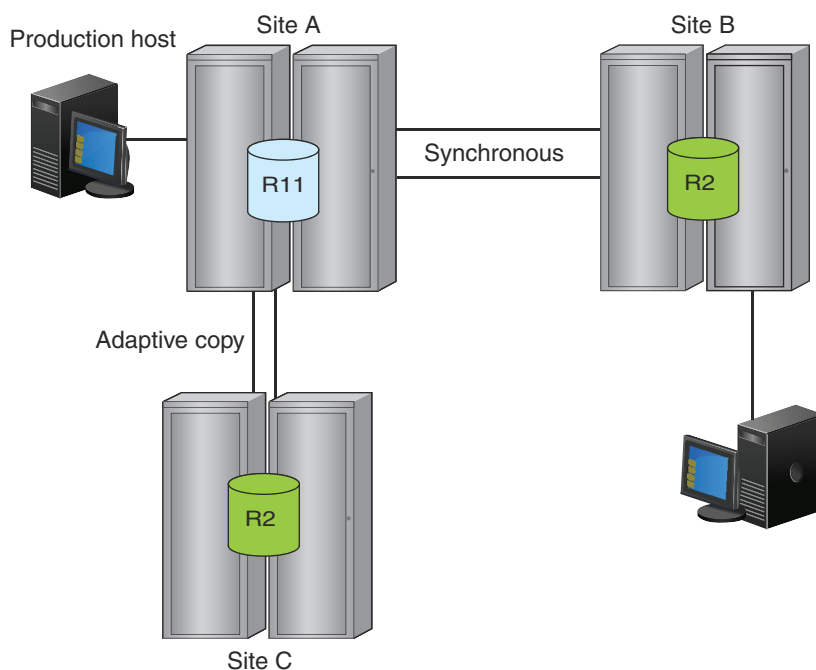
Use concurrent SRDF to replace an existing R11 or R2 device with a new device. To replace an R11 or R2, migrate data from the existing device to a new device using adaptive copy disk mode, and then replace the existing device with the newly populated device.

Concurrent SRDF topologies use Fibre Channel and Gigabit Ethernet.

This example shows:

- The R11 -> R2 in Site B in synchronous mode.
- The R11 -> R2 in Site C in adaptive copy mode.

Figure 10 Concurrent SRDF topology



Concurrent SRDF/S with Engineuity Consistency Assist

If both legs of a concurrent SRDF configuration use SRDF/S, you can leverage the independent consistency protection feature. This feature is based on Engineuity Consistency Assist (ECA) and enables you to manage consistency on each concurrent SRDF leg independently.

If consistency protection on one leg is suspended, consistency protection on the other leg can remain active and continue protecting the primary site.

Cascaded SRDF solutions

Cascaded SRDF provides a zero data loss solution at long distances in the event that the primary site is lost.

In cascaded SRDF configurations, 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.


Cascaded SRDF provides:

- Fast recovery times at the tertiary site
- Tight integration with TimeFinder product family
- Geographically dispersed secondary and tertiary sites

If the primary site fails, cascaded SRDF can continue mirroring, with minimal user intervention, from the secondary site to the tertiary site. This enables a faster recovery at the tertiary site.


Both the secondary and the tertiary site can be failover sites. Open systems solutions typically fail over to the tertiary site.

Figure 11 Cascaded SRDF topology



SRDF/Star solutions

SRDF/Star is a disaster recovery solution that consists of three sites: primary (production), secondary, and tertiary. The secondary site synchronously mirrors the data from the primary site, and the tertiary site asynchronously mirrors the production data.

 **Note:** In mainframe environments, GDDR is required to implement SRDF/Star. For more information, see [SRDF/Star for mainframe systems](#) on page 48 and the appropriate GDDR product guide listed in [More information](#) on page 105.

If an outage occurs at the primary site, SRDF/Star enables you to quickly move operations and re-establish remote mirroring between the remaining sites. When 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 that address different recovery and availability objectives:

- **Concurrent SRDF/Star**—Data is mirrored from the primary site concurrently to two R2 devices. Both the secondary and tertiary sites are potential recovery sites should the primary site fail. Differential resynchronization is used between the secondary and the tertiary sites.
- **Cascaded SRDF/Star**—Data is mirrored first from the primary site to a secondary site, and then from the secondary to a tertiary site. Both the secondary and tertiary sites are potential recovery sites. Differential resynchronization is used between the primary and the tertiary site.

Differential synchronization between two remote sites:

- Allows SRDF/Star to rapidly reestablish cross-site mirroring should the primary site fail.
- Greatly reduces the time required to remotely mirror the selected production site.

If a rolling disaster affects the primary site, SRDF/Star helps you determine which remote site has the most current data. You can select which site to operate from and which site's data to use when recovering from the primary site failure.

If the primary site fails, SRDF/Star enables you to resume asynchronous protection between the secondary and tertiary sites, with minimal data movement.


Concurrent SRDF/Star

In concurrent SRDF/Star solutions, production data on R11 devices replicates to two R2 devices in two remote arrays.

In this example:

- Site B is a secondary site using SRDF/S links from Site A.
- Site C is a tertiary site using SRDF/A links from Site A.
- The (normally inactive) recovery links are SRDF/A between Site C and Site B.

Figure 12 Concurrent SRDF/Star




Concurrent SRDF/Star with R22 devices

SRDF supports concurrent SRDF/Star topologies using R22 devices. R22 devices have two SRDF mirrors, only one of which is active on the SRDF links at a given time. R22 devices improve the resiliency of the SRDF/Star application, and reduce the number of steps for failover procedures.

This example shows R22 devices at Site C.

Figure 13 Concurrent SRDF/Star with R22 devices




Cascaded SRDF/Star

In cascaded SRDF/Star solutions, the synchronous secondary site is always more current than the asynchronous tertiary site. If the synchronous secondary site fails, the cascaded SRDF/Star solution can incrementally establish an SRDF/A session between primary site and the asynchronous tertiary site.

Cascaded SRDF/Star can 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 Site B and Site C to fully synchronize them.

This example shows a basic cascaded SRDF/Star solution.

Figure 14 Cascaded SRDF/Star

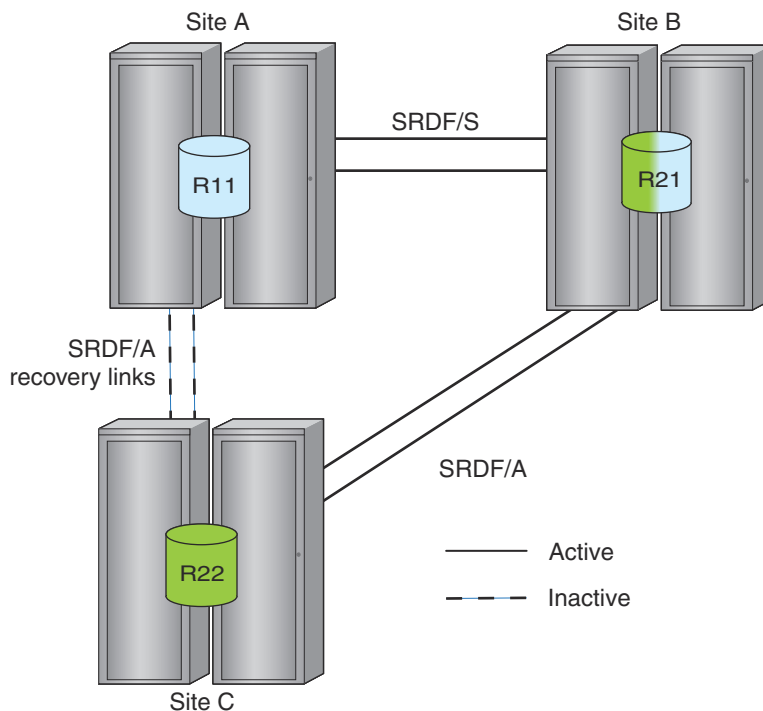


Cascaded SRDF/Star with R22 devices

You can use R22 devices to pre-configure the SRDF pairs required to incrementally establish an SRDF/A session between Site A and Site C in case Site B fails.

This example shows cascaded R22 devices in a cascaded SRDF solution.

Figure 15 R22 devices in cascaded SRDF/Star



In cascaded SRDF/Star configurations with R22 devices:

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

Requirements/restrictions

Cascaded and Concurrent SRDF/Star configurations (with and without R22 devices) require the following:

- 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.
- 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.

SRDF/Star for open systems

Solutions Enabler controls, manages, and automates SRDF/Star in open systems environments. Session management is required at the production site.

Host-based automation is provided for normal, transient fault, and planned or unplanned failover operations.

Dell EMC Solutions Enabler Symmetrix SRDF CLI Guide provides detailed descriptions and implementation guidelines.

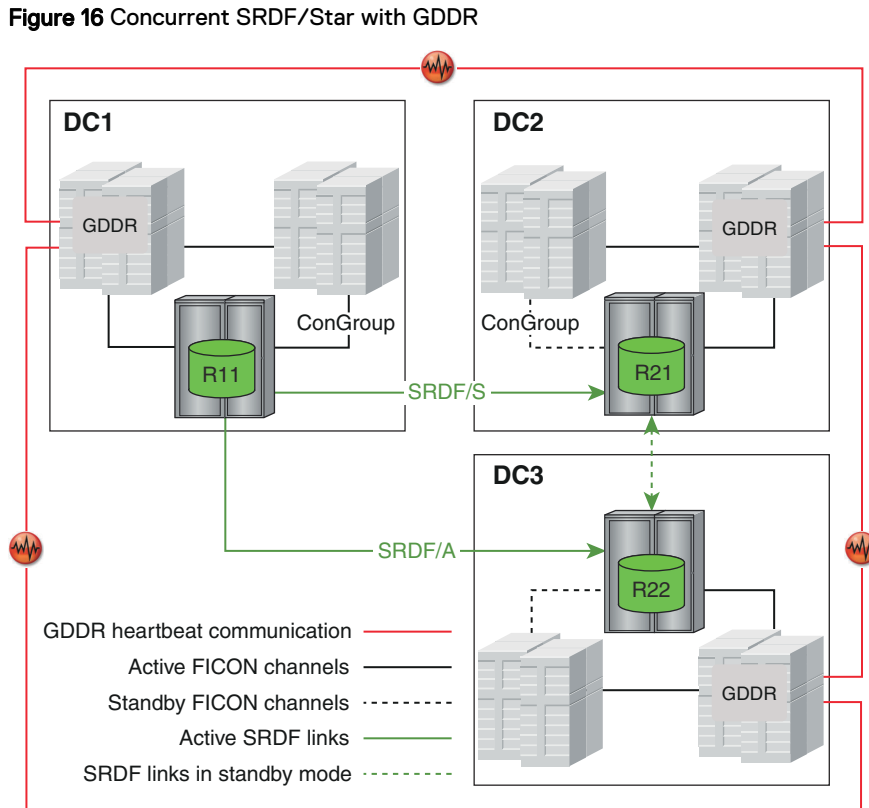
In cascaded and concurrent configurations, a restart from the asynchronous site may require a wait for any remaining data to arrive from the synchronous site. Restarts from the synchronous site requires no wait unless the asynchronous site is more recent (the latest updates need to be brought to the synchronous site).

SRDF/Star for mainframe systems

The SRDF Host Component for z/OS creates and manages the standard 2-site SRDF configurations along with Cascaded and Concurrent 3-site configurations. Geographically Dispersed Disaster Restart (GDDR) creates and manages SRDF/Star configurations in a mainframe environment.

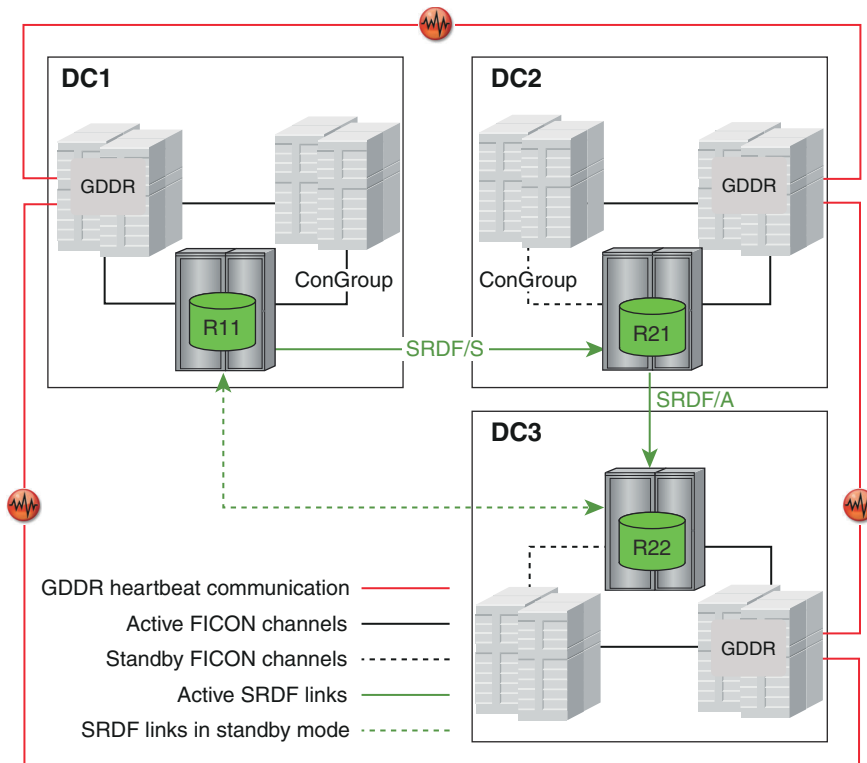
The structure of SRDF/Star configurations differs when GDDR is involved as each site has a GDDR control system. Each control system monitors the overall SRDF/Star configuration, detects failure conditions, and acts on those conditions (for example, by restarting operations at an alternate site).

This is an example of a Concurrent SRDF/Star configuration with GDDR:



This is an example of a Cascaded SRDF/Star configuration with GDDR:

Figure 17 Cascaded SRDF/Star with GDDR



Mainframe-specific variations of SRDF/Star

GDDR provides some variations of SRDF/Star configurations that are unique to the mainframe environment.

2-site SRDF/Star

In the 2-site SRDF/Star configuration, there are three storage arrays as in any other SRDF/Star configuration, but there are GDDR controls systems at DC1 and DC3 only. This means that if there is a failure at DC1 (the primary site), operations can be restarted at DC3 only.

Asynchronous SRDF/Star (Star-A)

Asynchronous SRDF/Star configurations are similar to other SRDF/Star configurations but all SRDF links use SRDF/A. This enables both DC2 in addition to DC3 to be remote from DC1 and so provide better protection from a site failure at DC1.

SRDF/Star with Autoswap

Autoswap is a facility to move (swap) workloads from volumes in one set of storage arrays to volumes in another set of arrays without interrupting host processing. The combination of Autoswap with SRDF/Star provides near-continuous availability through device failover between sites DC1 and DC2, while also providing disaster restart capabilities at site DC3.

SRDF/Star restrictions

- GNS Remote Mirroring is NOT supported with STAR configurations .
- Devices that are part of an RP configuration, cannot at the same time, be part of an SRDF/Star configuration.

- The RDF groups that are part of a STAR CG cannot contain any devices that are not part of the Star CG.
- Devices that are part of a STAR CG should not be controlled outside of `symstar` commands.
- Devices that are part of an SRDF/Metro configuration cannot at the same time be part of an SRDF/Star configuration.
- If any array in a SRDF/Star configuration is running HYPERMAX OS, Solutions Enabler 8.1 or higher is required in order to manage that configuration.
- If any array in a SRDF/Star configuration is running PowerMaxOS, Solutions Enabler 9.0 or later is required in order to manage that configuration.
- Each SRDF/Star control host must be connected to only one site in the SRDF/Star triangle. A Star control host is where the `symstar` commands are issued.
- A minimum of one SRDF daemon must be running on at least one host attached locally to each site. This host must be connected to only one site in the SRDF/Star triangle. The host could be the same as the Star control host but is not required unless using `symstar modifycg`. Dell EMC strongly recommends running redundant SRDF daemons on multiple hosts to ensure that at least one SRDF daemon is available to perform time-critical, consistency monitoring operations. Redundant SRDF daemons avoid service interruptions caused by performance bottlenecks local to a host.
- SRDF/A recovery links are required.
- SRDF groups cannot be shared between separate SRDF/Star configurations.
- R22 devices are required in SRDF/Star environments that include VMAX 10K or VMAXe arrays.
- CKD striped metadevices are not supported.
- R2 devices cannot be larger than their R1 devices.
- Composite groups consisting of device groups are not supported.
- Devices enabled as part of consistency groups cannot at the same time be part of an SRDF/Star configuration.
- Devices cannot be BCV devices.
- Every device must be dynamic SRDF (R1 and R2 capable).
- BCV device management must be configured separately.
 - ① **Note:**
Dell EMC strongly recommends that you have BCV device management available at both the synchronous and asynchronous target sites.
- With Engenuity 5876.159.102 and higher, a mixture of thin and (non-diskless) thick devices is supported.
 - ① **Note:**
If the thick device is on a DMX array running Engenuity 5876 and higher, thick-to-thin migration is supported if the array is running Engenuity 5876.163.105 and higher.

4-site configurations

4-site configurations provide extra data protection. There are configurations for both FBA and mainframe environments.

4-site FBA configurations

The four-site SRDF solution for open systems host environments replicates FBA data by using both concurrent and cascaded SRDF topologies.


Four-site SRDF is a multi-region disaster recovery solution with higher availability, improved protection, and less downtime than concurrent or cascaded SRDF solutions.

Four-site SRDF solution offers multi-region high availability by combining the benefits of concurrent and cascaded SRDF solutions.

If two sites fail because of a regional disaster, a copy of the data is available, and you have protection between the remaining two sites. You can create a four-site SRDF topology from an existing 2-site or 3-site SRDF topology. Four-site SRDF can also be used for data migration.

This is an example of the four-site SRDF solution.

Figure 18 SRDF 4-site FBA configuration

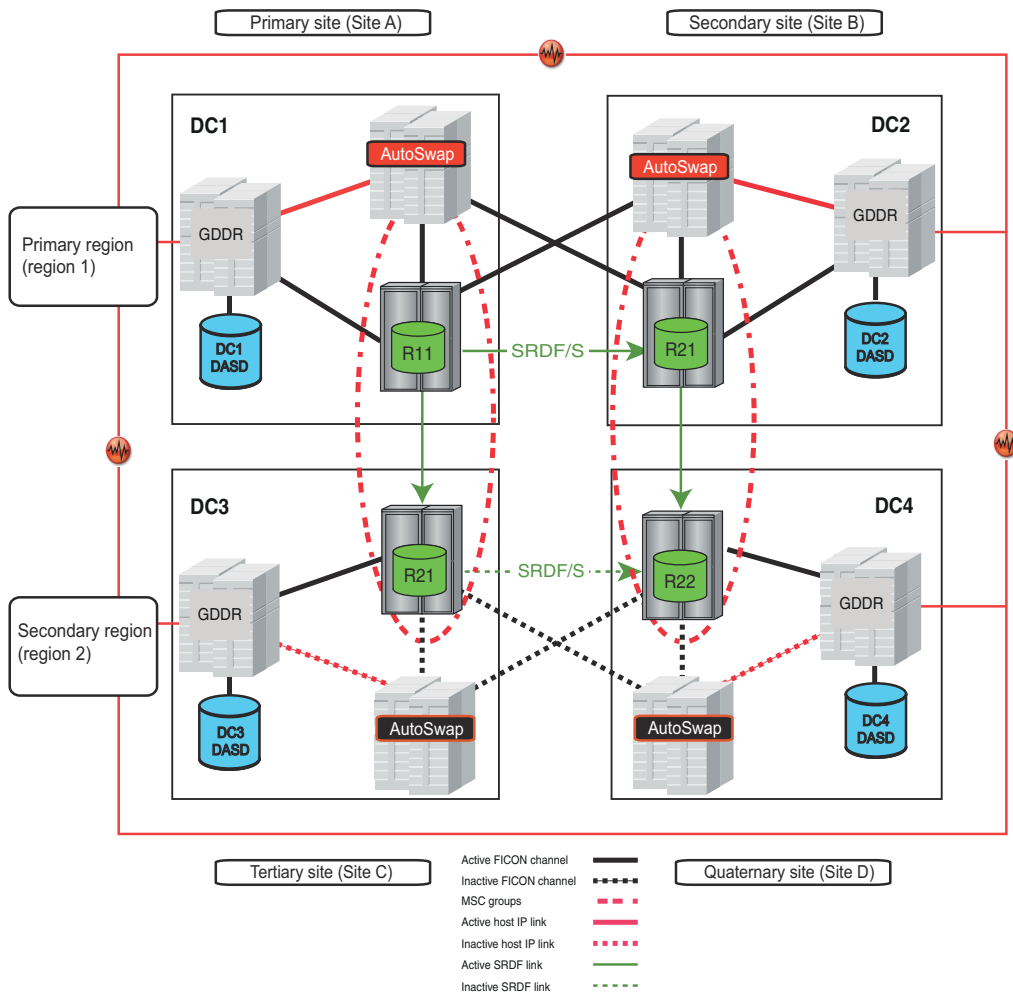


4-site mainframe configurations (SRDF/SQAR)

SRDF/SQAR (Symmetrix Quadrilateral Asynchronous Replication) is a 4-site implementation of SRDF/S and SRDF/A that enables differential resynchronization between sites along the perimeter of a 'square' SRDF topology. Dell EMC GDDR is required to be able to implement SRDF/SQAR.

Dell EMC support for the SRDF/SQAR configuration provides the ability to recover from a single or dual unplanned site outage in one region with SRDF/S protection established differentially between the recover sites in another region. This enables rapid resumption of a workload with SRDF/S and Autoswap protection in another region. In certain failure scenarios, it also provides zero data loss recovery across regions.

Figure 19 SRDF/SQAR with Autoswap environment



The diagram shows four Dell EMC GDDR control systems with their independent heartbeat communication paths, separate from the production disk and computer facilities. Each of the managed z/OS systems has Dell EMC Autoswap and Dell EMC Consistency Groups (ConGroup) installed.

Each GDDR SRDF/SQAR environment manages two consistency groups (one active, one defined) and two Multi-Session Consistency MSC groups (both active). A consistency group is a named group of source (R1) volumes managed by the ConGroup application as a unit. An MSC group is a named group consisting of multiple SRDF groups operating in SRDF/A mode, managed by the Dell EMC MSC control software feature as a single unit. The relationship between Site A (DC1) and Site B (DC2) is maintained through SRDF/S replication of primary disk images at DC1 to DC2, while SRDF/A replication maintains out of region, mirrored data at Site C (DC3) and Site 4 (DC4).

Requirements and restrictions

- SRDF/SQAR is required to be configured with the MSC High Availability feature with a second SCF instance and MSC configured using weight factor =2.
- SRDF Host Component actions that change the devices defined to the SQAR MSC groups require the MSC tasks to be down at the time of the change.
- Connectivity is provided only along the perimeter of the SQAR topology. Cross site connectivity (for example, Site A to Site D) is not supported. Therefore, traditional three-site SRDF/Star as a recovery configuration is not available.
- In the case of a single site failure it is important to know which SRDF/A site is more current. The existing SRDF/A secondary time-of-day value is used to determine which site is ahead.

SRDF recovery scenarios

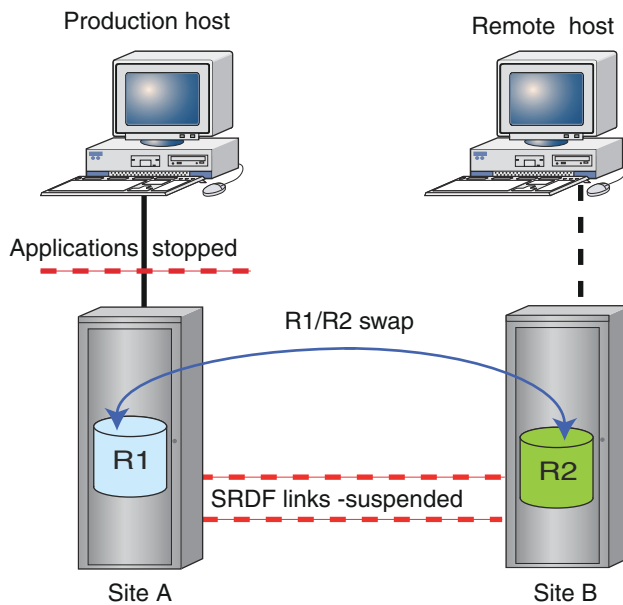
This section describes recovery scenarios in 2-site SRDF configurations.

Planned failover (SRDF/S)

A planned failover moves production applications from the primary site to the secondary site in order to test the recovery solution, upgrade or perform maintenance at the primary site.

This diagram shows a 2-site SRDF configuration before the R1 <-> R2 personality swap:

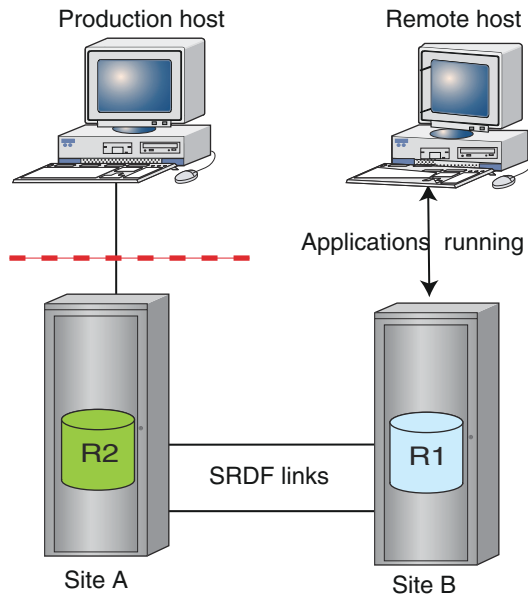
Figure 20 Planned failover: before personality swap



- Applications on the production host are stopped.
- SRDF links between Site A and Site B are suspended.
- If SRDF/CG is used, consistency is disabled.

The next diagram shows a 2-site SDRF configuration after the R1 <-> R2 personality swap.

Figure 21 Planned failover: after personality swap



When the maintenance, upgrades or testing procedures are complete, use a similar procedure to return production to Site A.

Unplanned failover


An unplanned failover moves production applications from the primary site to the secondary site after an unanticipated outage at the primary site.

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

Unlike the planned failover operation, an unplanned failover resumes production at the secondary site, but without remote mirroring until Site A becomes operational and ready for a failback operation.

This diagram shows failover to the secondary site after the primary site fails.

Figure 22 Failover to Site B, Site A and production host unavailable.



Failback to the primary array

When the primary host and array containing the primary (R1) devices are operational once more, an SRDF failback allows production processing to resume on the primary host.

Recovery for a large number of invalid tracks

If the R2 devices have handled production processing for a long period of time, there may be a large number of invalid tracks owed to the R1 devices. SRDF control software can resynchronize the R1 and R2 devices while the secondary host continues production processing. Once there is a relatively small number of invalid tracks owed to the R1 devices, the failback process can take place.

Temporary link loss

In SRDF/A configurations, if a temporary loss (10 seconds or less) of all SRDF/A links occurs, the SRDF/A state remains active, and data continues to accumulate in global memory. This may result in an elongated cycle, but the secondary array dependent-write consistency is not compromised and the primary and secondary array device relationships are not suspended.

[Transmit Idle](#) on page 92 can keep SRDF/A in an active state during all links lost conditions.

In SRDF/S configurations, if a temporary link loss occurs, writes are stalled (but not accumulated) in hopes that the SRDF link comes back up, at which point writes continue.

Reads are not affected.

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.

Permanent link loss (SRDF/A)

If all SRDF links are lost for more than link limbo or Transmit Idle can manage:

- All of the devices in the SRDF group 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. These tracks are owed to the secondary array once the links are restored.

When the links are restored, normal SRDF recovery procedures are followed:

- Metadata representing the data owed is compared and merged based on normal host recovery procedures.
- Data is resynchronized by sending the owed tracks as part of the SRDF/A cycles.

Data on non-consistency exempt devices on the secondary array is always dependent-write consistent in SRDF/A active/consistent state, even when all SRDF links fail. 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 array.

SRDF/A session cleanup

When an SRDF/A single session mode is dropped, SRDF:

- Marks new incoming writes at the primary array as being owed to the secondary array.

- Discards the capture and transmit delta sets, and marks the data as being owed to the secondary array. These tracks are sent to the secondary array once SRDF is resumed, as long as the copy direction remains primary-to-secondary.
- Marks and discards only the receive delta set at the secondary array, and marks the data as tracks owed to the primary array.

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

Failback from R2 devices

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

After the primary array has been repaired, production operations can return to the primary array as described in [SRDF recovery scenarios](#) on page 54.

If the failover to the secondary site is an extended event, the SRDF/A solution can be reversed by issuing a personality swap. 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 array hosts, SRDF/A can be activated to resume normal asynchronous mode protection.

CHAPTER 3

High availability

This chapter provides more detail on the high availability configurations that SRDF/Metro provides for open systems (FBA) and IBM i D910 application hosts.

- [SRDF/Metro](#) 60
- [SRDF/Metro life cycle](#)..... 62
- [SRDF/Metro resilience](#)..... 63
- [Mobility ID with ALUA](#)..... 68
- [Disaster recovery facilities](#)..... 69
- [Deactivate SRDF/Metro](#).....70
- [SRDF/Metro restrictions](#)..... 71

SRDF/Metro

In traditional SRDF, R1 devices are Read/Write accessible. R2 devices are Read Only/Write Disabled.


In SRDF/Metro configurations:

- R2 devices are Read/Write accessible to application hosts.
- Application hosts can write to both the R1 and R2 side of the device pair.
- R2 devices assume the same external device identity (geometry, device WWN) as the R1 devices.

This shared identity means that the R1 and R2 devices appear to application hosts as a single, virtual device across the two arrays.

SRDF/Metro can be deployed in either a single, multipathed host or in a clustered host environment.

Figure 23 SRDF/Metro



Hosts can read and write to both the R1 and R2 devices:

- In a single host configuration, a single host issues I/O operations. Multipathing software directs parallel reads and writes to each array.
- In a clustered host configuration, multiple hosts issue I/O operations to both sides of the SRDF device pair. Each cluster node has dedicated access to an individual storage array.
- In both single host and clustered configurations, writes to the R1 or R2 devices are synchronously copied to the paired device. SRDF/Metro software resolves write conflicts to maintain consistent images on the SRDF device pairs. The R1 device and its paired R2 device appear to the host as a single virtualized device.

Other characteristics of SRDF/Metro are:

- SRDF/Metro is managed using either Solutions Enabler or Unisphere.
- SRDF/Metro requires a license on both arrays.
- Storage arrays can simultaneously contain SRDF groups operations and SRDF groups traditional SRDF operations.
- The arrays can be up to 200 km (125 miles) apart.

Key differences in SRDF/Metro compared to traditional SRDF

In SRDF/Metro configurations:

- R2 device is Read/Write accessible to the host.
- Hosts can write to both R1 and R2 devices.
- Both sides of the SRDF device pair appear to the hosts as the same device.
- The R2 device assumes the personality of the primary R1 device (such as geometry and device WWN).
- Two extra SRDF pair states:
 - ActiveActive for configurations using the Witness options (Array and Virtual)
 - ActiveBias for configurations using bias

i **Note:** R1 and R2 devices should not be presented to the cluster until they reach one of these two states and present the same WWN.

Device management

All device pairs in an SRDF/Metro group are managed together for all supported operations, with the following exceptions:

- Create pair operations can add devices to the group.
- Delete pair operations can delete a subset of the SRDF devices in the SRDF group.

Failure recovery

If the link fails or other failures occur, SRDF/Metro must decide which side of a device pair remains accessible to the host. The available options are:

- **Device Bias option:** Device pairs for SRDF/Metro have a *bias* attribute. By default, the create pair operation sets the bias to the R1 side of the pair. That is, if the device pair becomes Not Ready (NR) on the SRDF link, the R1 (bias side) remains accessible to the hosts. The R2 (nonbias side) is inaccessible to the hosts. However, if there is a failure on the R1 side, the host loses all connectivity to the device pair. The Bias option cannot make the R2 device available to the host.
- **Witness option:** A Witness is a third party that mediates between the two sides to help decide which remains available to the host if there is a failure. The Witness method allows for intelligently choosing on which side to continue operations when the bias-only method may not result in continued host availability to a surviving nonbiased array. The Witness option is the default.

SRDF/Metro provides two types of Witnesses, Array and Virtual:

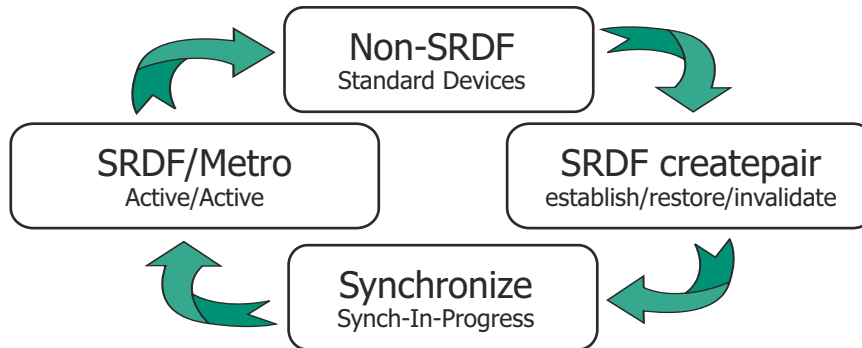
- **Array Witness :** The operating environment on a third array acts as the mediator to decide the side of the device pair that remains R/W accessible to the host. It gives priority to the bias side, but should that side be unavailable the nonbias side remains available.
- **Virtual Witness (vWitness) option:** vWitness provides the same functionality as the Array Witness option, only it is packaged to run in a virtual appliance, not on the array.

[SRDF/Metro resilience](#) on page 63 has more information about these failure-recovery mechanisms.

SRDF/Metro life cycle

The life cycle of an SRDF/Metro configuration begins and ends with an empty SRDF group and a set of non-SRDF devices, as shown in this diagram:

Figure 24 SRDF/Metro life cycle



The life cycle of an SRDF/Metro configuration includes the following steps and states:

- **Create device pairs** in an empty SRDF group.
Create pairs indicating that they are to operate in an SRDF/Metro configuration.
If all the SRDF device pairs are Not Ready (NR) on the link, the create pair operation can be used to add more devices into the SRDF group.
- **Make the device pairs Read/Write (RW)** on the SRDF link.
Use the establish or restore options to make the devices Read/Write (RW) on the SRDF link.
Alternatively, use the invalidate option to create the devices without making them Read/Write (RW) on the SRDF link.
- **Synchronize** the device pairs.
When the devices in the SRDF group are Read/Write (RW) on the SRDF link, invalid tracks begin synchronizing between the R1 and R2 devices.
The direction of synchronization is controlled by the establish (R1 to R2) or restore (R2 to R1) operation.
- **Activate SRDF/Metro**
Device pairs transition to the ActiveActive pair state when:
 - Device federated personality and other information is copied from the R1 side to the R2 side.
 - Using the information copied from the R1 side, the R2 side sets its identity as an SRDF/Metro R2 when queried by host I/O drivers.
 - R2 devices become accessible to the hosts.

When all SRDF device pairs in the group transition to the ActiveActive state, hosts can discover the R2 devices with federated personality of R1 devices. SRDF/Metro manages the SRDF device pairs in the SRDF group. A write to either side of the SRDF device pair completes to the host only after it is transmitted to the other side of the SRDF device pair, and the other side has acknowledged its receipt.
- **Add/remove devices** to/from an SRDF/Metro group.
In arrays that run HYPERMAX OS, the group must be in either Suspended or Partitioned state to add or remove devices. Either the create pair or move pair operations add device to a group. In HYPERMAX OS arrays, the create pair operation initializes the devices before adding them

to the SRDF/Metro session. The move pair operation adds the devices to the session without initializing them.

In arrays that run PowerMaxOS both the create pair and move pair add devices to a SRDF/Metro session without initializing them.

Use the delete pair or move pair operations to delete all or a subset of device pairs from the SRDF group. Removed devices return to the non-SRDF state. If all device pairs are removed from the group, the group is no longer controlled by SRDF/Metro. The group can be reused either as a SRDF/Metro or non-Metro group.

- **Deactivate SRDF/Metro**

If all devices in an SRDF/Metro group are deleted, that group is no longer part of an SRDF/Metro configuration.

You can use the create pair operation to repopulate the SRDF group, either for SRDF/Metro or for non-Metro.

SRDF/Metro resilience

If a SRDF/Metro device pair becomes Not Ready (NR) on the SRDF link, SRDF/Metro decides which side of the device pair to remain accessible to hosts (the winning side).

The following sections explain the methods that SRDF/Metro provides for making that decision.

Device Bias

In an SRDF/Metro configuration, PowerMaxOS and HYPERMAX OS use the link between the two sides of each device pair to ensure consistency of the data on each sides. If the device pair becomes Not Ready (NR) on the RDF link, PowerMaxOS and HYPERMAX OS choose the bias side of the device pair to remain accessible to the hosts, and makes the non-bias side of the device pair, and of all other device pairs in the SRDF/Metro group, inaccessible. This prevents data inconsistencies between the two sides of the RDF device pair.

When adding device pairs to an SRDF/Metro group (using the create pair operation), PowerMaxOS and HYPERMAX OS configure the R1 side of the pair as the bias side. You use an option in the SRDF management software to indicate the use of Device Bias.

If the Witness options are not used, the establish and restore operations also require an indication that the Device Bias method is in use.

When the SRDF/Metro devices pairs are configured to use bias, their pair state is ActiveBias.

Bias can be changed when all device pairs in the SRDF/Metro group have reached the ActiveActive or ActiveBias pair state.

Device bias is the fallback method used when either of the witness methods is unable to determine the winning side or when there is a witness failure.

Array Witness

When using the Array Witness method, SRDF/Metro uses a third "witness" array to determine the winning side. The witness array runs one of these operating environments:

- PowerMaxOS 5978.144.144 or later
- HYPERMAX OS 5977.945.890 or later
- HYPERMAX OS 5977.810.784 with ePack containing fixes to support SRDF N-x connectivity
- Engenuity 5876 with ePack containing fixes to support SRDF N-x connectivity


In the event of a failure, the witness decides which side of the Metro group remains accessible to hosts, giving preference to the bias side. The Array Witness method allows for choosing which side operations continue when the Device Bias method may not result in continued host availability to a surviving non-biased array.

The Array Witness must have SRDF connectivity to both the R1-side array and R2-side array. SRDF remote adapters (RA's) are required on the witness array with applicable network connectivity to both the R1 and R2 arrays.

For redundancy, there can be multiple witness arrays but only one witness array is used by an individual Metro group; the two sides of the Metro group agree on the witness array to use when the Metro group is activated. If the auto configuration process fails and no other applicable witness arrays are available, SRDF/Metro uses the Device Bias method.

The Array Witness method requires two SRDF groups; one between the R1 array and the witness array, and a second between the R2 array and the witness array. Neither group contains any devices.

Figure 25 SRDF/Metro Array Witness and groups



SRDF/Metro management software checks that the Witness groups exist and are online when carrying out establish or restore operations. SRDF/Metro determines which witness array an SRDF/Metro group is using, so there is no need to specify the Witness. Indeed, there is no means of specifying the Witness.

When the witness array is connected to both the SRDF/Metro paired arrays, the configuration enters Witness Protected state.

When the Array Witness method is in operation, the state of the device pairs is ActiveActive.

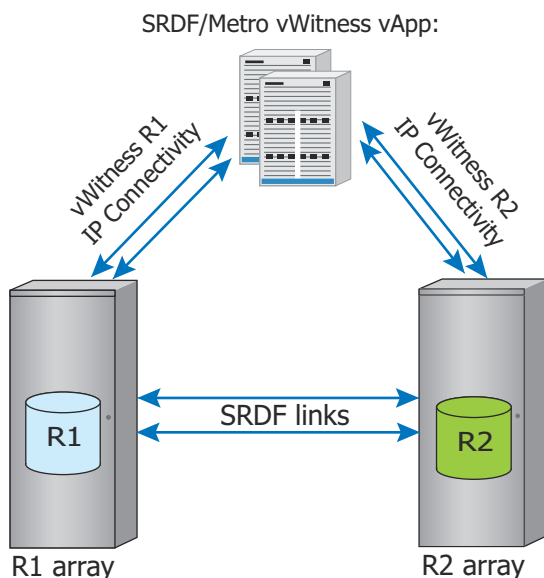
If the witness array becomes inaccessible from both the R1 and R2 arrays, PowerMaxOS and HYPERMAX OS set the R1 side as the bias side, the R2 side as the non-bias side, and the state of the device pairs becomes ActiveBias.

Virtual Witness (vWitness)

vWitness is a third resiliency option. It has similar capabilities to the Array Witness method, except that it is packaged to run in a virtual appliance (vApp) on a VMware ESX server, not on an array.

The vWitness and Array Witness options are treated the same in the operating environment, and can be deployed independently or simultaneously. When deployed simultaneously, SRDF/Metro favors the Array Witness option over the vWitness option, as the Array Witness option has better availability. For redundancy, you can configure up to 32 vWitnesses.

Figure 26 SRDF/Metro vWitness vApp and connections



The management guests on the R1 and R2 arrays maintain multiple IP connections to redundant vWitness virtual appliances. These connections use TLS/SSL to ensure secure connectivity.

Once you have established IP connectivity to the arrays, you can use SRDF management software to:

- Add a new vWitness to the configuration. This does not affect any existing vWitnesses. Once the vWitness is added, it is enabled for participation in the vWitness infrastructure.
- Query the state of a vWitness configuration.
- Suspend a vWitness. If the vWitness is currently servicing an SRDF/Metro session, this operation requires a force flag. This puts the SRDF/Metro session in an unprotected state until it renegotiates with another witness, if available.
- Remove a vWitness from the configuration. Once removed, SRDF/Metro breaks the connection with vWitness. You can only remove vWitnesses that are not currently servicing active SRDF/Metro sessions.

Witness negotiation and selection

At start up, SRDF/Metro needs to decide which witness to use for each SRDF/Metro group. Each side of the SRDF/Metro configuration maintains a list of witnesses, that is set up by the administrator. To begin the negotiation process, the non-bias side sends its list of witnesses to the bias side. On receiving the list, the bias side compares it with its own list of witnesses. The first matching witness definition is selected as the witness and the bias side sends its identification back to the non-bias side. The two sides then establish communication with the selected witness. The two sides repeat this process for each SRDF/Metro group. Should the selected witness become unavailable at any time, the two sides repeat this selection algorithm to choose an alternative.

Intelligent witness management

When both sides run PowerMaxOS, the negotiation process is enhanced to include a decision on the winning side in the event of a failure. The selection of the winning side is based on (in priority order):

1. The side that has connectivity to the application host (requires PowerMaxOS 5978.444.444 or later)
2. The side that has a SRDF/A DR leg
3. Whether the SRDF/A DR leg is synchronized
4. The side that has more than 50% of the RA or FA directors that are available
5. The side that is currently the bias side

The two sides regularly repeat this selection process for each SRDF/Metro group to ensure that the winning side remains the one that is most preferable. This means that the winning side may change during the course of the SRDF/Metro session.


Witness failure scenarios

These diagrams show how SRDF/Metro reacts to various failure scenarios when either Witness option is in use.


Figure 27 SRDF/Metro Witness single failure scenarios

S1	R1 side of device pair
S2	R2 side of device pair
W	Witness Array/vWitness
↔	SRDF links
↔↔	SRDF links/IP connectivity*
X	Failure/outage


* Depending on witness type




S1 and S2 remain accessible to host
S2 wins future failures
S1 calls home




S1 and S2 remain accessible to host
Move to bias mode
S1 and S2 call home




S1 failed
S2 remains accessible to host



S1 remains accessible to host
S2 suspends

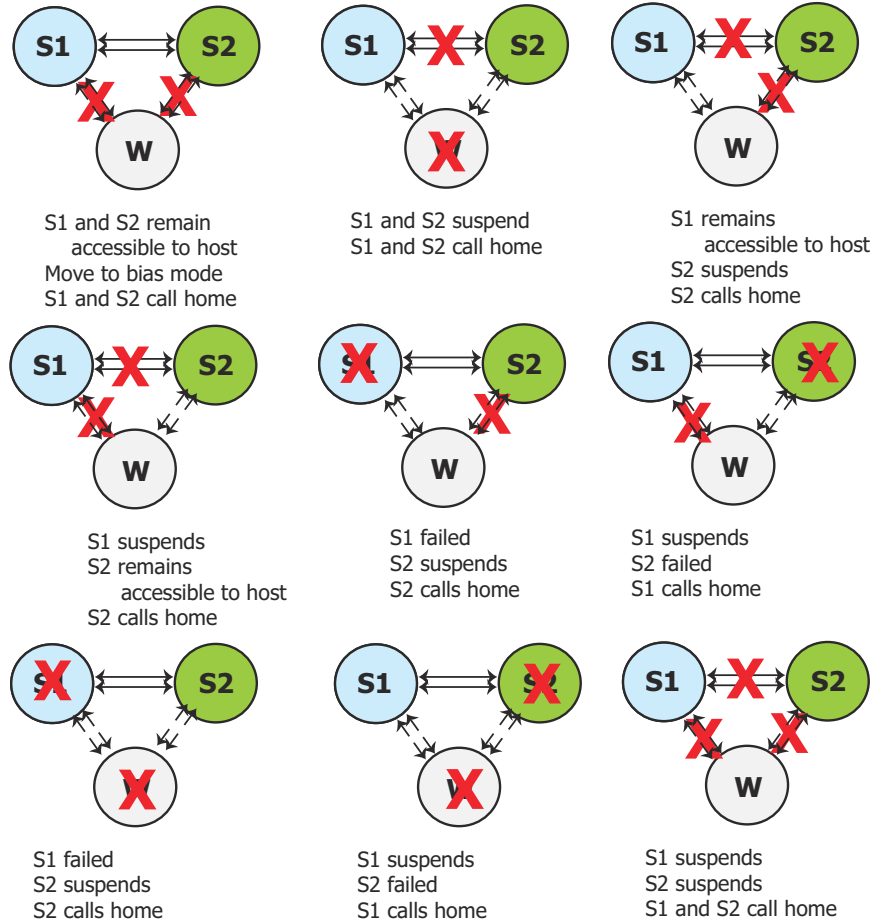


S1 and S2 remain accessible to host
S1 wins future failures
S2 calls home



S2 failed
S1 remains accessible to host

Figure 28 SRDF/Metro Witness multiple failure scenarios



Mobility ID with ALUA


Mobility ID with Asymmetric Logical Unit Access (ALUA) assigns a unique identifier to a device in a system. This identifier enables the device to be moved between arrays without the need for any reconfiguration on the host. PowerMaxOS brings Mobility ID with ALUA capabilities to SRDF/Metro. So, when both sides run PowerMaxOS you can specify the Mobility ID in the createpair operation in place of the regular device identifier.

Disaster recovery facilities

Devices in SRDF/Metro groups can simultaneously be part of device groups that replicate data to a third, disaster-recovery site.

Either or both sides of the Metro region can be replicated. You can choose which ever configuration that suits your business needs. The following diagram shows the possible configurations:

Figure 29 Disaster recovery for SRDF/Metro



Note that the device names differ from a standard SRDF/Metro configuration. This reflects the change in the devices' function when disaster recovery facilities are in place. For instance, when the R2 side is replicated to a disaster recovery site, its name changes to R21 because it is both the:

- R2 device in the SRDF/Metro configuration
- R1 device in the disaster-recovery configuration

Replication modes

As the diagram shows, the links to the disaster-recovery site use either SRDF/Asynchronous (SRDF/A) or Adaptive Copy Disk. In a double-sided configuration, each of the SRDF/Metro arrays can use either replication mode.

Operating environment

In a HYPERMAX OS environment, both SRDF/Metro arrays must run HYPERMAX OS 5977.945.890 or later. The disaster-recovery arrays can run Enginuity 5876 and later or HYPERMAX OS 5977.691.684 and later.

In a PowerMaxOS environment, both SRDF/Metro arrays must run PowerMaxOS 5978.144.144 or later. The disaster recovery arrays can run PowerMaxOS 5978.144.144 and later, HYPERMAX OS 5977.952.892 and later, or Enginuity 5876.288.195 and later.

Deactivate SRDF/Metro

To terminate a SRDF/Metro configuration, remove all the device pairs from the group by either deleting them or moving them to another SRDF group.

Note: In HYPERMAX OS, the devices must be in Suspended state in order to delete them.

When all the devices in the SRDF/Metro group have been deleted, the group is no longer part of an SRDF/Metro configuration. You can delete or move a subset of device pairs from the group. The SRDF/Metro configuration terminates only when the last pair is removed.

Delete one side of a SRDF/Metro configuration

You can remove devices from only one side of a SRDF/Metro configuration, using SRDF management software in an operation called half delete. You can perform this operation on all devices or a subset of them on one side of a SRDF/Metro group. The devices must be in the Partitioned SRDF state to perform the half delete operation.

After the half delete operation:

- The devices on the side where the operation was performed are no longer SRDF devices.
- The devices at the other side of the SRDF group retain their configuration as SRDF/Metro devices.

If all devices are deleted from one side of the SRDF group, that side of the SRDF group is no longer part of the SRDF/Metro configuration.

Restore native personality to a federated device

Devices in SRDF/Metro configurations have federated personalities. When a device is removed from an SRDF/Metro configuration, the device personality can be restored to its original, native personality.

Some restrictions apply to restoring the native personality of a device which has federated personality as a result of a participating in a SRDF/Metro configuration:

- Requires HYPERMAX OS 5977.691.684 or later or PowerMaxOS 5978.
- The device must be unmapped and unmasked.
- The device must have a federated WWN.
- The device must not be an SRDF device.

- The device must not be a ProtectPoint device.

SRDF/Metro restrictions

Some restrictions and dependencies apply to SRDF/Metro configurations:

- Both the R1 and R2 side must be running HYPERMAX OS 5977.691.684 or later or PowerMaxOS 5978.
- Only non-SRDF devices can become part of an SRDF/Metro configuration.
- The R1 and R2 must be identical in size.
- In an SRDF/Metro group, all the R1 devices must be on one side of the SRDF link and all the R2 devices on the other side.
- Devices cannot have Geometry Compatibility Mode (GCM) or User Geometry set.
- Online device expansion is not supported.
- The create pair, with the establish option, establish, restore, and suspend operations apply to all devices in the SRDF group.
- In HYPERMAX OS, control of devices in an SRDF group which contains a mixture of R1s and R2s is not supported.
- An SRDF/Metro configuration contains FBA or IBM i D910² devices only. It cannot contain CKD (mainframe) devices.

Interaction restrictions

Some restrictions apply to SRDF device pairs in an SRDF/Metro configuration with TimeFinder and Open Replicator (ORS):

- Open Replicator is not supported.
- Devices cannot be BCVs.
- Devices cannot be used as the target of the data copy when the SRDF devices are RW on the SRDF link with either a SyncInProg or ActiveActive SRDF pair state.
- A snapshot does not support restores or re-links to itself.

2. IBM i D910 requires PowerMaxOS 5978.444.444 or later.

High availability

CHAPTER 4

Data migration

This chapter has more detail on the data migration facilities of SRDF.

- [Introduction to data migration using SRDF](#) 74
- [Non-disruptive migration](#) 74
- [Migrating data with concurrent SRDF](#) 77
- [Migration-only SRDF](#) 81
- [Device Migration operations requirements](#) 82

Introduction to data migration using SRDF

Data migration is a one-time movement of data from one array (the source) to another array (the target). Typical examples are data center refreshes where data is moved from an old array after which the array is retired or re-purposed. Data migration is *not* data movement due to replication (where the source data is accessible after the target is created) or data mobility (where the target is continually updated).

After a data migration operation, applications that access the data reference it at the new location.

This chapter introduces the migration capabilities that are based on SRDF. These capabilities are available for open host (FBA) systems only. Mainframe systems have data migration capabilities but they are based on other technologies.

Non-disruptive migration

Non-disruptive migration (NDM) is a method for migrating data from one array to another with application downtime. The migration typically takes place within a data center.

This table shows the migration facilities that NDM provides:

Source	Target
VMAX running Engenuity 5876	PowerMax running PowerMaxOS 5978
VMAX running Engenuity 5876	VMAX All Flash or VMAX3 running HYPERMAX OS 5977
VMAX All Flash or VMAX3 running HYPERMAX OS 5977	PowerMax or VMAX All Flash running PowerMaxOS 5978

Migration from VMAX array

Migrating from a VMAX array uses SRDF in Pass-through mode. In this mode, the application host can access data on both source and target devices while the migration is in progress. PowerMaxOS or HYPERMAX OS on the target ensures that the source processes all I/O operations sent to the target.

Process

The steps in the migration process are:

1. Set up the environment – configure the infrastructure of the source and target array, in preparation for data migration.
2. On the source array, select a storage group to migrate.
3. If using NDM Updates, shut down the application associated with the storage group.
4. Create the migration session – copy the content of the storage group to the target array using SRDF.
When creating the session, optionally specify whether to move the identity of the LUNs in the storage group to the target array.
5. When the data copy is complete:
 - a. If the migration session did not move the identity of the LUNs, reconfigure the application to access the new LUNs on the target array.

- b. Cutover the storage group to the PowerMax, VMAX All Flash, or VMAX3 array array.
 - c. Commit the migration session – remove resources from the source array and those used in the migration itself. The application now uses the target array only.
6. If using NDM Updates, restart the application.
 7. To migrate further storage groups, repeat steps 2 on page 74 to 6 on page 75.
 8. After migrating all the required storage groups, remove the migration environment.

Other features

Other features of migrating from VMAX to PowerMax, VMAX All Flash, or VMAX3 are:

- Data can be compressed during migration to the PowerMax, VMAX All Flash, or VMAX3 array
- Allows for non-disruptive revert to the source array
- There can be up to 50 migration sessions in progress simultaneously
- NDM does not require an additional license as it is part of PowerMaxOS or HYPERMAX OS
- The connections between the application host and the arrays use FC; the SRDF connection between the arrays uses FC or GigE

Devices and components that cannot be part of an NDM process are:

- CKD devices
- eNAS data
- ProtectPoint and FAST.X relationships along with their associated data

Migration from VMAX All Flash or VMAX3

Migrating from a VMAX All Flash or VMAX3 array running HYPERMAX OS 5977 to an array running PowerMaxOS 5978 uses a modified form of SRDF/Metro. So both the source and target arrays are visible and accessible to the application host while the migration takes place.

Process

Normal flow

The steps in the migration process that is normally followed are:

1. Set up the migration environment – configure the infrastructure of the source and target array, in preparation for data migration.
2. On the source array, select a storage group to migrate.
3. If using NDM Updates, shut down the application associated with the storage group.
4. Create the migration session optionally specifying whether to move the identity of the LUNs in the storage group to the target array – copy the content of the storage group to the target array using SRDF/Metro.
During this time the source and target arrays are both accessible to the application host.
5. When the data copy is complete:
 - a. if the migration session did not move the identity of the LUNs, reconfigure the application to access the new LUNs on the target array.
 - b. Commit the migration session – remove resources from the source array and those used in the migration itself.
6. If using NDM Updates, restart the application.
7. To migrate further storage groups, repeat steps 2 to 6 on page 75.

8. After migrating all the required storage groups, remove the migration environment.

Alternate flow

There is an alternative process that pre-copies the data to the target array before making it available to the application host. The steps in this process are:

1. Set up the migration environment – configure the infrastructure of the source and target array, in preparation for data migration.
2. On the source array, select a storage group to migrate.
3. Use the precopy facility of NDM to copy the selected data to the target array.
Optionally, specify whether to move the identity off the LUNS in the storage group to the target array.
While the data copy takes place, the source array is available to the application host, but the target array is unavailable.
4. When the copying of the data is complete: use the Ready Target facility in NDM to make the target array available to the application host also.
 - a. If the migration session did not move the identity of the LUNs, reconfigure the application to access the new LUNs on the target array.
 - b. If using NDM Updates, restart the application.
 - c. Commit the migration session – remove resources from the source array and those used in the migration itself. The application now uses the target array only.
5. To migrate further storage groups, repeat steps 2 on page 76 to 4 on page 76.
6. After migrating all the required storage groups, remove the migration environment.

Other functions

Other NDM facilities that are available for exceptional circumstances are:

- Cancel – to cancel a migration that has not yet been committed.
- Sync – to stop or start the synchronization of writes to the target array back to source array. When stopped, the application runs on the target array only. Used for testing.
- Recover – to recover a migration process following an error.

Other features

Other features of migrating from VMAX3, VMAX All Flash or PowerMax to PowerMax are:

- Data can be compressed during migration to the PowerMax array
- Allows for non-disruptive revert to the source array
- There can be up to 50 migration sessions in progress simultaneously
- Does not require an additional license as NDM is part of PowerMaxOS
- The connections between the application host and the arrays use FC; the SRDF connection between the arrays uses FC or GigE

Devices and components that cannot be part of an NDM process are:

- CKD devices
- eNAS data
- ProtectPoint and FAST.X relationships along with their associated data

Migrating data with concurrent SRDF

In concurrent SRDF topologies, you can non-disruptively migrate data between arrays along one SRDF leg while remote mirroring for protection along the other leg.

Once the migration process completes, the concurrent SRDF topology is removed, resulting in a 2-site SRDF topology.


Replacing R1 devices with new R1 devices

This diagram show the use of migration to replace the R1 devices in a 2-site configuration. The diagram shows the:

- Initial 2-site topology
- Interim 3-site migration topology
- Final 2-site topology

After migration, the new primary array is mirrored to the original secondary array.

Figure 30 Migrating data and replacing the original primary array (R1)




Replacing R2 devices with new R2 devices

This diagram shows the use of migration to replace the R2 devices in a 2-site configuration. The diagram shows the:

- Initial 2-site topology
- Interim 3-site migration topology
- Final 2-site topology

After migration, the original primary array is mirrored to a new secondary array.

Figure 31 Migrating data and removing the original secondary array (R2)




Replacing R1 and R2 devices with new R1 and R2 devices

This diagram shows the use of migration to replace both the R1 and R2 arrays in a 2-site configuration. The diagram shows the :

- Initial 2-site topology
- Migration topology using a combination of concurrent and cascaded SRDF
- Final 2-site topology

Figure 32 Migrating data and replacing the original primary (R1) and secondary (R2) arrays



Migration-only SRDF

In some cases, you can migrate data with full SRDF functionality, including disaster recovery and other advanced SRDF features.

In cases where full SRDF functionality is not available, you can move the data using migration-only SRDF.

This table lists SRDF common operations and features and whether they are available in SRDF groups during SRDF migration-only environments.

Table 7 Limitations of the migration-only mode

SRDF operations or features	Whether supported during migration
R2 to R1 copy	Only for device rebuild from un-rebuildable RAID group failures.
Failover, failback, domino	Not available
SRDF/Star	Not available
SRDF/A features (DSE, Consistency Group, ECA, MSC)	Not available
Dynamic SRDF operations (create, delete, and move SRDF pairs; R1/R2 personality swap)	Not available
TimeFinder operations	Only on R1
Online configuration change or upgrade	<ul style="list-style-type: none"> If the changes affect the group or devices being migrated, migration must be suspended before and while the upgrade or configuration changes take place. If the changes do not affect the group or devices being migrated, migration can continue while the upgrade or configuration changes take place.
Out-of-family Non-Disruptive Upgrade (NDU)	Not available

Device Migration operations requirements

- Each array must have a unique ID (sid).
- The existing SRDF device and the new devices must be dynamic R1 or R2 capable.

PowerMaxOS and HYPERMAX OS

- Devices that are part of an SRDF/Metro configuration cannot be migrated.
- Adaptive copy write pending mode is not supported when the R1 side of the RDF pair is on an array running PowerMaxOS or HYPERMAX OS.
For configurations where the R1 side is on an array running PowerMaxOS or HYPERMAX OS, and the R2 side is running Enginuity 5876, the mode of the new device pair is set to the RDF mode of the R1 device being replaced.
- The Geometry Compatibility Mode (GCM) attribute allows devices on arrays running PowerMaxOS or HYPERMAX OS to be paired with devices on arrays running Enginuity 5876 that have an odd number of cylinders. When GCM is set, migration operations are subject to the following restrictions:
 - If the new device is on an array running PowerMaxOS or HYPERMAX OS:
 - If the R1 device is being replaced:
If the existing R2 device is on an array running Enginuity 5876 with an odd number of cylinders, the migration is allowed if the new device can be made the same size using the GCM attribute.

If the existing R2 device is on an array running PowerMaxOS or HYPERMAX OS with GCM set, the migration is allowed if the new device can be made the same size by setting the GCM attribute.
 - If the R2 is being replaced:
If the existing R1 device is on an array running Enginuity 5876 with an odd number of cylinders, then the migration is allowed if the new device can be made the same size by setting the GCM attribute.

If the existing R1 device is on an array running PowerMaxOS or HYPERMAX OS with GCM set, the migration is allowed if the new device can be made the same size by setting the GCM attribute.
 - If the new device is on an array running Enginuity 5876 and has an odd number of cylinders:
 - If the R1 is being replaced:
If the existing R2 device is on an array running Enginuity 5876, the new device must be the same configured size

If the existing R2 device is on an array running PowerMaxOS or HYPERMAX OS with GCM set, the migration is allowed if the new device has the same GCM size as the R2 device.
 - If the R2 is being replaced:
If the existing R1 device is on an array running Enginuity 5876, the new device must be the same configured size.

If the existing R1 device is on an array running PowerMaxOS or HYPERMAX OS with GCM set, the migration is allowed if the new device has the same GCM size as the R1.

CHAPTER 5

SRDF I/O operations

This chapter shows how SRDF handles write and read operations. In addition, there is information on the performance and resilience features of SRDF/A.

- [SRDF write operations](#)..... 84
- [SRDF read operations](#)..... 90
- [SRDF/A resilience and performance features](#)..... 91

SRDF write operations

This section describes SRDF write operations:

- Write operations in synchronous mode
- Write operations in asynchronous mode
- Cycle switching in asynchronous mode
- Write operations in cascaded SRDF


Write operations in synchronous mode

In synchronous mode, data must be successfully written to cache at the secondary site before a positive command completion status is returned to the application host that issued the write command.

The following diagram shows the steps in a synchronous write operation:

1. The application host sends a write command to the local array. The host emulations write data to cache and create a write request.
2. SRDF emulations frame updated data in cache according to the SRDF protocol, and transmit it across the SRDF links.
3. The SRDF emulations in the remote array receive data from the SRDF links, write it to cache and return an acknowledgment to SRDF emulations in the local array.
4. The SRDF emulations in the local array forward the acknowledgment to host emulations which pass it on to the application host.

Figure 33 Write I/O flow: simple synchronous SRDF



Write operations in asynchronous mode

When SRDF/A is in use, the primary array collects host write operations into delta sets and transfers them in cycles to the secondary array. The primary array acknowledges the host write operations as soon as they are written to its cache.

SRDF/A sessions behave differently depending on:

- Whether they are managed individually (Single Session Consistency (SSC)) or as a consistency group (Multi Session Consistency (MSC)).
 - With SSC, the SRDF group is managed individually. The primary array's operating environment controls cycle switching. SRDF/A cycles are switched independently of any other SRDF groups on any array in the solution. [Cycle switching in asynchronous mode](#) on page 86 has more details.
 - With MSC, the SRDF group is part of a consistency group spanning all associated SRDF/A sessions. SRDF host software coordinates cycle switching to provide dependent-write

consistency across multiple sessions, which may also span arrays. The host software switches SRDF/A cycles for all SRDF groups in the consistency group simultaneously. [SRDF/A MSC cycle switching](#) on page 87 has more details.

- The number of transmit cycles supported at the R1 side. Enginuity 5876 supports only a single cycle. PowerMaxOS and HYPERMAX OS support multiple cycles queued to be transferred.

Data in a delta set is processed using four cycle types:

- **Capture cycle**—Incoming I/O is buffered in the capture cycle on the R1 side. The host receives immediate acknowledgment.
- **Transmit cycle**—Data collected during the capture cycle is moved to the transmit cycle on the R1 side.
- **Receive cycle**—Data is received on the R2 side.
- **Apply cycle**—Changed blocks in the delta set are marked as invalid tracks and destaging to disk begins.
A new receive cycle is started.

The operating environment running on the R1 side determines when the next capture cycle can begin. It also determines the number of cycles that can be in progress simultaneously.

PowerMaxOS 5978 or HYPERMAX OS 5977 – Multicycle mode

If both arrays in the configuration run PowerMaxOS or HYPERMAX OS, SRDF/A operates in multicycle mode. There can be two or more cycles on the R1 side, but only two cycles on the R2 side:

- On the R1 side:
 - One Capture
 - One or more Transmit
- On the R2 side:
 - One Receive
 - One Apply

Cycle switches are decoupled from committing delta sets to the next cycle.

When the preset Minimum Cycle Time is reached, the R1 data collected during the capture cycle is added to the transmit queue. Then a new R1 capture cycle begins. There is no wait for the commit on the R2 side before starting a new capture cycle.

The transmit queue holds cycles waiting to be transmitted to the R2 side. Data in the transmit queue is committed to the R2 receive cycle when the current transmit cycle and apply cycle are empty.

Queuing enables smaller cycles of data to be buffered on the R1 side and reduces the size of delta sets transferred to the R2 side.

The SRDF/A session can adjust to accommodate changes in the solution. If the SRDF link speed decreases or the apply rate on the R2 side increases, more SRDF/A cycles can be queued on the R1 side.

Multicycle mode increases the robustness of the SRDF/A session and reduces spillover into the DSE storage pool.

Enginuity 5876

If either array in the solution is running Enginuity 5876, SRDF/A operates in legacy mode. There are two cycles on the R1 side, and two cycles on the R2 side:

- On the R1 side:

- One Capture
- One Transmit
- On the R2 side:
 - One Receive
 - One Apply

Each cycle switch moves the delta set to the next cycle in the process.

A new capture cycle cannot start until both the transmit cycle on the R1 side and the apply cycle on the R2 side are empty.

Cycle switching can occur within the preset Minimum Cycle Time. However, it can also take longer since it depends on both:

- The time taken to transfer the data from the R1 transmit cycle to the R2 receive cycle
- The time taken to destage the R2 apply cycle

Cycle switching in asynchronous mode

The number of capture cycles supported at the R1 side varies depending on whether one or both the arrays in the solution are running PowerMaxOS or HYPERMAX OS.

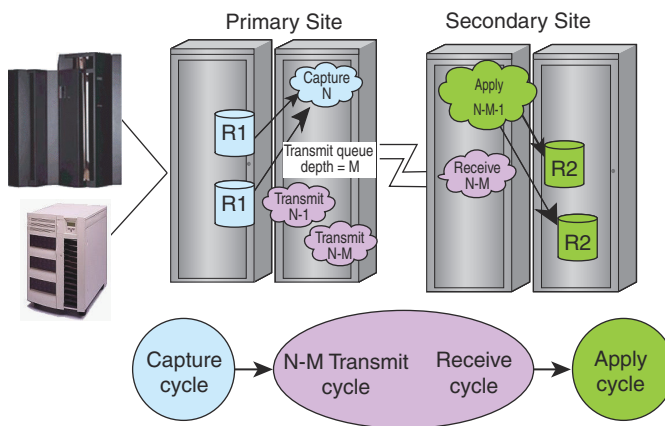
PowerMaxOS or HYPERMAX OS

SRDF/A SSC sessions where both arrays are running PowerMaxOS or HYPERMAX OS have one or more Transmit cycles on the R1 side (multi-cycle mode).

The following diagram shows multi cycle mode:

- Multiple cycles (one capture cycle and multiple transmit cycles) on the R1 side, and
- Two cycles (receive and apply) on the R2 side.

Figure 34 SRDF/A SSC cycle switching – multi-cycle mode



In multi-cycle mode, each cycle switch creates a new capture cycle (N) and the existing capture cycle (N-1) is added to the queue of cycles (N-1 through N-M cycles) to be transmitted to the R2 side by a separate commit action.

Only the data in the last transmit cycle (N-M) is transferred to the R2 side during a single commit.

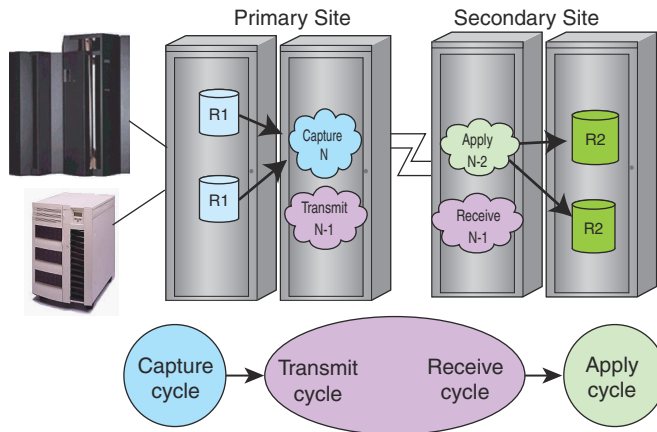
Enginuity 5876

SRDF/A SSC sessions that include an array running Enginuity 5876 have one Capture cycle and one Transmit cycle on the R1 side (legacy mode).

The following diagram shows legacy mode:

- 2 cycles (capture and transmit) on the R1 side, and
- 2 cycles (receive and apply) on the R2 side

Figure 35 SRDF/A SSC cycle switching – legacy mode



In legacy mode, there are conditions must be met before an SSC cycle switch can take place:

- The previous cycle's transmit delta set (N-1 copy of the data) must have completed transfer to the receive delta set on the secondary array.
- On the secondary array, 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.

SSC cycle switching in concurrent SRDF/A

In single session mode, cycle switching on both legs of the concurrent SRDF topology typically occurs at different times.

Data in the Capture and Transmit cycles may differ between the two SRDF/A sessions.

SRDF/A MSC cycle switching

SRDF/A MSC:

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


PowerMaxOS 5978 or HYPERMAX OS 5977

SRDF/A MSC sessions, where both arrays are running PowerMaxOS or HYPERMAX OS, have two or more cycles on the R1 side (multi-cycle mode).

Note: If either the R1 side or R2 side of an SRDF/A session is running PowerMaxOS or HYPERMAX OS, Solutions Enabler 8.x or later (for HYPERMAX OS) or Solutions Enabler 9.0 or later (for PowerMaxOS) is required to monitor and manage MSC groups.

The following diagram shows the cycles on the R1 side (one capture cycle and multiple transmit cycles) and 2 cycles on the R2 side (receive and apply) for an SRDF/A MSC session when both of the arrays in the SRDF/A solution are running PowerMaxOS or HYPERMAX OS.

Figure 36 SRDF/A MSC cycle switching – multi-cycle mode



SRDF cycle switches all SRDF/A sessions in the MSC group at the same time. All sessions in the MSC group have the same:

- Number of cycles outstanding on the R1 side
- Transmit queue depth (M)

In SRDF/A MSC sessions, the array's operating environment performs a coordinated cycle switch during a window of time when no host writes are being completed.

So that it can establish consistency, MSC temporarily suspends write operations from the application host across all SRDF/A sessions. MSC resumes those write operations once there is consistency. There is a timeout associated with the suspension of write operations to protect against failure of MSC. Should the timeout expire, write operations from the application host resume.

Enginuity 5876

SRDF/A MSC sessions that include an array running Enginuity 5876 have only two cycles on the R1 side (legacy mode).

In legacy mode, the following conditions must be met before an MSC cycle switch can take place:

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

To achieve consistency through cycle switching, MSC suspends write operations from the application host in the same way as it does when both arrays run PowerMaxOS 5978 or HYPERMAX OS 5977. It also uses the timeout to protect against the failure of MSC while synchronization is in progress.

Write operations in cascaded SRDF

In cascaded configurations, R21 devices operate as:

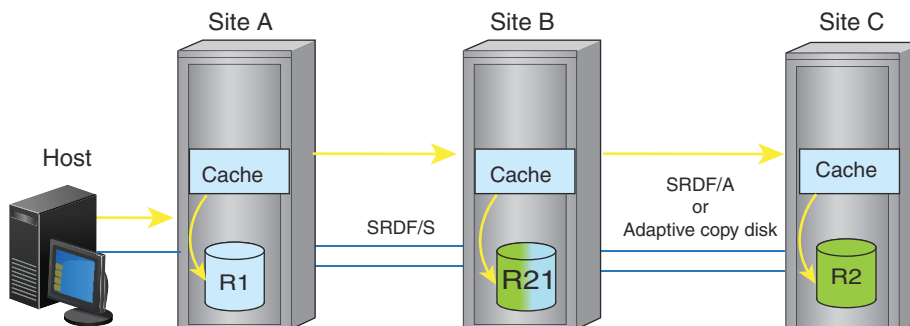
- R2 devices to devices in the R1 array
- R1 devices to devices in the R2 array

I/O to R21 devices includes:

- Synchronous I/O between the production site (R1) and the closest (R21) remote site.
- Asynchronous or adaptive copy I/O between the synchronous remote site (R21) and the tertiary (R2) site.
- You can Write Enable the R21 to a host so that the R21 behaves like an R2 device. This allows the R21 -> R2 connection to operate as R1 -> R2, while the R1 -> R21 connection is automatically suspended. The R21 begins tracking changes against the R1.

This diagram shows the synchronous I/O flow in a cascaded SRDF topology.

Figure 37 Write commands to R21 devices



When a write command arrives in cache at Site B:

- The SRDF emulation at Site B sends a positive status back across the SRDF links to Site A (synchronous operations), and
- Creates a request for SRDF emulations at Site B to send data across the SRDF links to Site C.

SRDF read operations

Read operations from the R1 device do not usually involve the SRDF emulations:

- For read “hits” (the production host issues a read to the R1 device, and the data is in local cache), the host emulation reads data from cache and sends it to the host.
- For read “misses” (the requested data is not in cache), the drive emulation reads the requested data from local drives to cache.

Read operations if R1 local copy fails

In SRDF/S, SRDF/A, and adaptive copy configurations, SRDF devices can process read I/Os that cannot be processed by regular logical devices. If the R1 local copy fails, the R1 device can still service the request as long as its SRDF state is Ready and the R2 device has good data.

SRDF emulations help service the host read requests when the R1 local copy is not available:

- The SRDF emulations bring data from the R2 device to the host site.
- 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 copy.

PowerMaxOS or HYPERMAX OS

Arrays running PowerMaxOS or HYPERMAX OS cannot service SRDF/A read I/Os with Delta Set Extension (DSE). So, spillover is not invoked during a SRDF/A restore operation until that restore operation is complete. [SRDF/A cache data offloading](#) on page 91 contains more information about DSE.

Read operations from R2 devices

Reading data from R2 devices directly from a host connected to the R2 is not recommended, because:

- SRDF/S relies on the application’s ability to determine if the data image is the most current. The array at the R2 side 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 transit on the SRDF links, the host is not reading the most current data.

Dell 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.

In a mainframe environment, however, PPRC does allow secondary devices to be defined as read only. Dell EMC supports this with extra controls in SCF, AutoSwap, and ConGroup software. Such a configuration is known as Host Read Only (HRD).

SRDF/A resilience and performance features

Operational problems that can occur in a SRDF/A configuration include:

- Unbalanced SRDF/A configurations or I/O spikes can cause SRDF/A solutions to use large amounts of cache.
- Transient network outages can interrupt SRDF sessions.
- An application may write to the same record repeatedly.

This section describes the SRDF/A features that address these problems.

Tunable cache

You can set the SRDF/A maximum cache utilization threshold to a percentage of the system write pending limit for an individual SRDF/A session in single session mode and multiple SRDF/A sessions in single or MSC mode.

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

By default, the SRDF/A session drops if array cache is exhausted. You can keep the SRDF/A session running for a user-defined period. You can assign priorities to sessions, keeping SRDF/A active for as long as cache resources allow. If the condition is not resolved at the expiration of the user-defined period, the SRDF/A session drops.

Use the features described below to prevent SRDF/A from exceeding its maximum cache utilization threshold.

SRDF/A cache data offloading

If the system approaches the maximum SRDF/A cache utilization threshold, Delta Set Extension (DSE) offloads some or all of the delta set data. DSE can be configured/enabled/disabled independently on the R1 and R2 sides. However, Dell EMC recommends that both sides use the same configuration of DSE.

DSE works in tandem with group-level write pacing to prevent cache over-utilization during spikes in I/O or network slowdowns.

Resources to support offloading vary depending on the operating environment running on the array.

PowerMaxOS or HYPERMAX OS

PowerMaxOS and HYPERMAX OS offload data into a Storage Resource Pool. One or more Storage Resource Pools are pre-configured before installation and used by a variety of functions. DSE can use a Storage Resource Pool pre-configured specifically for DSE. If no such pool exists, DSE can use the default Storage Resource Pool. All SRDF groups on the array use the same Storage Resource Pool for DSE. DSE requests allocations from the Storage Resource Pool only when DSE is activated.

The Storage Resource Pool used by DSE is sized based on your SRDF/A cache requirements. DSE is automatically enabled.

Enginuity 5876

Enginuity 5876 offloads data to a DSE pool that you configure. There must be a separate DSE pool for each device emulation type (FBA, IBM i, CKD3380 or CKD3390).

- In order to use DSE, each SRDF group must be explicitly associated with a DSE pool.
- By default, DSE is disabled.
- 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.

Mixed configurations: PowerMaxOS or HYPERMAX OS and Enginuity 5876

If the array on one side of an SRDF device pair is running PowerMaxOS or HYPERMAX OS and the other side is running a Enginuity 5876, the SRDF/A session runs in Legacy mode.

- DSE is disabled by default on both arrays.
- Dell EMC recommends that you enable DSE on both sides.

Transmit Idle

During short-term network interruptions, the transmit idle state indicates that SRDF/A is still tracking changes but is unable to transmit data to the remote side.

Write folding

Write folding improves the efficiency of your SRDF links.

When multiple updates to the same location arrive in the same delta set, the SRDF emulations send the only most current data across the SRDF links.

Write folding decreases network bandwidth consumption and the number of I/Os processed by the SRDF emulations.

Write pacing

SRDF/A write pacing reduces the likelihood that an active SRDF/A session drops due to cache exhaustion. Write pacing dynamically paces the host I/O rate so it does not exceed the SRDF/A session's service rate. This prevents cache overflow on both the R1 and R2 sides.

Use write pacing to maintain SRDF/A replication with reduced resources when replication is more important for the application than minimizing write response time.

You can apply write pacing to groups, or devices for individual RDF device pairs that have TimeFinder/Snap or TimeFinder/Clone sessions off the R2 device.

Group pacing

SRDF/A group pacing adjusts the pace of host writes to match the SRDF/A session's link transfer rate. When host I/O rates spike, or slowdowns make transmit or apply cycle times longer, group pacing extends the host write I/O response time to match slower SRDF/A service rates.

When DSE is activated for an SRDF/A session, host-issued write I/Os are paced so their rate does not exceed the rate at which DSE can offload the SRDF/A session's cycle data to the DSE Storage Resource Pool.

Group pacing behavior varies depending on whether the maximum pacing delay is specified:

- If the maximum write pacing delay is not specified, SRDF adds up to 50 ms to 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.
- If the maximum write pacing delay is specified, SRDF adds up to the user-specified maximum write pacing delay to keep the SRDF/A session running.

Group pacing balances 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.
- Reduced throughput on the SRDF links.
- **Enginuity 5876 only:** The write-pending level on an R2 device in an active SRDF/A session reaches the device write-pending limit.
- **Enginuity 5876 only:** The apply cycle time on the R2 side is longer than 30 s and the R1 capture cycle time (or in MSC, the capture cycle target).

Group pacing can be activated by configurations or activities that result in slow R2 operations, such as:

- Slow R2 physical drives resulting in longer apply cycle times.
- Director sparing operations that slow restore operations.
- I/O to the R2 array that slows restore operations.

Note: On arrays running Enginuity 5876, if the space in the DSE pool runs low, DSE drops and group SRDF/A write pacing falls back to pacing host writes to match the SRDF/A session's link transfer rate.

Device (TimeFinder) pacing

PowerMaxOS or HYPERMAX OS

SRDF/A device write pacing is not supported or required for asynchronous R2 devices in TimeFinder or TimeFinder SnapVX sessions when either array in the configuration is running PowerMaxOS or HYPERMAX OS, including:

- R1 PowerMaxOS or HYPERMAX OS - R2 PowerMaxOS or HYPERMAX OS
- R1 PowerMaxOS or HYPERMAX OS - R2 Enginuity 5876
- R1 Enginuity 5876 - R2 PowerMaxOS or HYPERMAX OS

Enginuity 5876

SRDF/A device pacing applies a write pacing delay for individual SRDF/A R1 devices whose R2 counterparts participate in TimeFinder copy sessions.

SRDF/A group pacing avoids high SRDF/A cache utilization levels when the R2 devices servicing both the SRDF/A and TimeFinder copy requests experience slowdowns.

Device pacing avoids high SRDF/A cache utilization when the R2 devices servicing both the SRDF/A and TimeFinder copy requests experience slowdowns.

Device pacing behavior varies depending on whether the maximum pacing delay is specified:

- If the maximum write pacing delay is not specified, SRDF adds up to 50 milliseconds to the overall host write response time to keep the SRDF/A session active.
- If the maximum write pacing delay is specified, SRDF adds up to the user-defined maximum write pacing delay to keep the SRDF/A session active.

Device pacing can be activated on the second hop (R21 -> R2) of a cascaded SRDF and cascaded SRDF/Star, topologies.

Device pacing may not take effect if all SRDF/A links are lost.

Write pacing and Transmit Idle

Host writes continue to be paced when:

- All SRDF links are lost, and

- Cache conditions require write pacing, and
- Transmit Idle is in effect.

Pacing during the outage is the same as the transfer rate prior to the outage.

CHAPTER 6

Interfamily compatibility

This chapter has more detail on the compatibility between different families of Dell EMC storage arrays.

- [Overview](#) 96
- [SRDF supported features](#) 96

Overview

SRDF can operate between different operating environments and arrays. Arrays running PowerMaxOS or HYPERMAX OS can connect to arrays running older operating environments. In mixed configurations where arrays are running different versions, SRDF features of the lowest version are supported.

PowerMax, VMAX All Flash, and VMAX3 arrays can connect to:

- PowerMax arrays running PowerMaxOS
- VMAX 250F, 450F, 850F, and 950F arrays running HYPERMAX OS
- VMAX 100K, 200K, and 400K arrays running HYPERMAX OS
- VMAX 10K, 20K, and 40K arrays running Enginuity 5876 with an Enginuity ePack

This interfamily connectivity allows you to add the latest hardware platform/operating environment to an existing SRDF solution, enabling technology refreshes.

Note: When you connect between arrays running different operating environments, limitations may apply. Information about which SRDF features are supported, and applicable limitations for 2-site and 3-site solutions is available in the *SRDF Interfamily Connectivity Information*.

SRDF supported features

The SRDF features supported on each hardware platform and operating environment are:

Table 8 SRDF features by hardware platform/operating environment

Feature	Enginuity 5876		HYPERMAX OS 5977	PowerMaxOS
	VMAX 40K, VMAX 20K	VMAX 10K	VMAX3, VMAX 250F, 450F, 850F, 950F	PowerMax 2000, PowerMax 8000
Max. SRDF devices/SRDF emulation (either Fibre Channel or GigE)	64K	8K	64K	64K
Max. SRDF groups/array	250	32	250	250
Max. SRDF groups/SRDF emulation instance (either Fibre Channel or GigE)	64	32	250 ^{ab}	250 ^{cd}
Max. remote targets/port	64	64	16K/SRDF emulation (either Fibre Channel or GigE)	16K/SRDF emulation (either Fibre Channel or GigE)
Max. remote targets/SRDF group	N/A	N/A	512	512
Fibre Channel port speed	2/4/8 Gb/s 16 Gb/s on 40K	2/4/8/16 Gb/s	16 Gb/s	16 Gb/s
GbE port speed	1 /10 Gb/s	1 /10 Gb/s	1 /10 Gb/s	1 /10 Gb/s

Table 8 SRDF features by hardware platform/operating environment (continued)

Feature	Enginuity 5876		HYPERMAX OS 5977	PowerMaxOS
	VMAX 40K, VMAX 20K	VMAX 10K	VMAX3, VMAX 250F, 450F, 850F, 950F	PowerMax 2000, PowerMax 8000
Min. SRDF/A Cycle Time	1 sec, 3 secs with MSC	1 sec, 3 secs with MSC	1 sec, 3 secs with MSC	1 sec, 3 secs with MSC
SRDF Delta Set Extension	Supported	Supported	Supported	Supported
Transmit Idle	Enabled	Enabled	Enabled	Enabled
Fibre Channel Single Round Trip (SiRT)	Enabled	Enabled	Enabled	Enabled
GigE SRDF Compression				
Software	Supported <ul style="list-style-type: none"> VMAX 20K VMAX 40K: Enginuity 5876.82.5 7 or higher 	Supported	Supported	Supported
Hardware	Supported <ul style="list-style-type: none"> VMAX 20K VMAX 40K: Enginuity 5876.82.5 7 or higher 	N/A	Supported	Supported
Fibre Channel SRDF Compression				
Software	Supported <ul style="list-style-type: none"> VMAX 20K VMAX 40K: Enginuity 5876.82.5 7 or higher 	Supported	Supported	Supported
Hardware	Supported	N/A	Supported	Supported

Table 8 SRDF features by hardware platform/operating environment (continued)

Feature	Enginuity 5876		HYPERMAX OS 5977	PowerMaxOS
	VMAX 40K, VMAX 20K	VMAX 10K	VMAX3, VMAX 250F, 450F, 850F, 950F	PowerMax 2000, PowerMax 8000
	<ul style="list-style-type: none"> VMAX 20K: N/A VMAX 40K: Enginuity 5876.82.5 7 or higher 			
IPv6 and IPsec				
IPv6 feature on 10 GbE	Supported	Supported	Supported	Supported
IPsec encryption on 1 GbE ports	Supported	Supported	N/A	N/A

- If both arrays are running HYPERMAX OS or PowerMaxOS, up to 250 RDF groups can be defined across all of the ports on a specific RDF director, or up to 250 RDF groups can be defined on 1 port on a specific RDF director.
- A port on the array running HYPERMAX OS or PowerMaxOS connected to an array running Enginuity 5876 supports a maximum of 64 RDF groups. The director on the HYPERMAX OS or PowerMaxOS side associated with that port supports a maximum of 186 (250 – 64) RDF groups.
- If both arrays are running HYPERMAX OS or PowerMaxOS, up to 250 RDF groups can be defined across all of the ports on a specific RDF director, or up to 250 RDF groups can be defined on 1 port on a specific RDF director.
- A port on the array running HYPERMAX OS or PowerMaxOS connected to an array running Enginuity 5876 supports a maximum of 64 RDF groups. The director on the HYPERMAX OS or PowerMaxOS side associated with that port supports a maximum of 186 (250 – 64) RDF groups.

CHAPTER 7

Management tools

This chapter contains an overview of the tools that enable you to manage an SRDF environment.

- [Solutions Enabler](#) 100
- [Unisphere](#) 100
- [SRDF/TimeFinder Manager for IBM i](#) 101
- [Mainframe management tools](#) 102

Solutions Enabler

Solutions Enabler provides a comprehensive command line interface (SYMCLI) to manage your storage environment.

SYMCLI commands are invoked from the host, either interactively on the command line, or using scripts.

SYMCLI is built on functions that use system calls to generate low-level I/O SCSI commands. Configuration and status information is maintained in a host database file, reducing the number of enquiries from the host to the arrays.

Use SYMCLI to:

- Configure array software (For example, TimeFinder, SRDF, Open Replicator)
- Monitor device configuration and status
- Perform control operations on devices and data objects

Solutions Enabler also has a Representational State Transfer (REST) API. Use this API to access performance and configuration information, and provision storage arrays. It can be used in any programming environments that supports standard REST clients, such as web browsers and programming platforms that can issue HTTP requests.

Unisphere

Unisphere is a web-based application that provides provisioning, management, and monitoring of arrays.

With Unisphere you can perform the following tasks:

Table 9 Unisphere tasks

Section	Allows you to:
Home	View and manage functions such as array usage, alert settings, authentication options, system preferences, user authorizations, and link and launch client registrations.
Storage	View and manage storage groups and storage tiers.
Hosts	View and manage initiators, masking views, initiator groups, array host aliases, and port groups.
Data Protection	View and manage local replication, monitor and manage replication pools, create and view device groups, and monitor and manage migration sessions.
Performance	Monitor and manage array dashboards, perform trend analysis for future capacity planning, and analyze data.
Databases	Troubleshoot database and storage issues, and launch Database Storage Analyzer.
System	View and display dashboards, active jobs, alerts, array attributes, and licenses.
Events	View alerts, the job list, and the audit log.
Support	View online help for Unisphere tasks.

Unisphere also has a Representational State Transfer (REST) API. With this API you can access performance and configuration information, and provision storage arrays. You can use the API in any programming environment that supports standard REST clients, such as web browsers and programming platforms that can issue HTTP requests.

SRDF/TimeFinder Manager for IBM i

Dell EMC SRDF/TimeFinder Manager for IBM i is a set of host-based utilities that provides an IBM i interface to SRDF and TimeFinder.

This feature allows you to configure and control SRDF or TimeFinder operations on arrays attached to IBM i hosts, including:

- SRDF: Configure, establish and split SRDF devices, including:
 - SRDF/A
 - SRDF/S
 - Concurrent SRDF/A
 - Concurrent SRDF/S
- TimeFinder:
 - Create point-in-time copies of full volumes or individual data sets.
 - Create point-in-time snapshots of images.

Extended features

SRDF/TimeFinder Manager for IBM i extended features provide support for the IBM independent ASP (IASP) functionality.

IASPs are sets of switchable or private auxiliary disk pools (up to 223) that can be brought online/offline on an IBM i host without affecting the rest of the system.

When combined with SRDF/TimeFinder Manager for IBM i, IASPs let you control SRDF or TimeFinder operations on arrays attached to IBM i hosts, including:

- Display and assign TimeFinder SnapVX devices.
- Execute SRDF or TimeFinder commands to establish and split SRDF or TimeFinder devices.
- Present one or more target devices containing an IASP image to another host for business continuance (BC) processes.

Access to extended features control operations include:

- From the SRDF/TimeFinder Manager menu-driven interface.
- From the command line using SRDF/TimeFinder Manager commands and associated IBM i commands.

Mainframe management tools

There are tools for managing SRDF configurations in a mainframe environment:

- Mainframe Enablers
- GDDR

Mainframe Enablers

Mainframe Enablers (MFE) is a suite of products for managing and monitoring Dell EMC storage systems in a mainframe environment. The entire suite consists of:

- SRDF Host Component for z/OS
- ResourcePak Base for z/OS
- Autoswap for z/OS
- Consistency Groups for z/OS
- TimeFinder SnapVX
- Data Protector for z/Systems(zDP)
- TimeFinder/Clone Mainframe Snap Facility
- TimeFinder/Mirror for z/OS
- TimeFinder Utility

In the context of SRDF, only the SRDF Host Component for z/OS, TimeFinder/Mirror for z/OS, plus these components of the ResourcePak for z/OS are relevant:

- SRDF/A Monitor
- WPA Monitor
- SRDF/AR

SRDF Host Component for z/OS

SRDF Host Component for z/OS is a z/OS subsystem for controlling SRDF processes and monitoring SRDF status using commands issued from a host. With the SRDF Host Component you can manage these SRDF variants:

- SRDF/S
- SRDF/A
- SRDF/DM
- SRDF/AR
- SRDF/CG
- SRDF/Star
- SRDF/SQAR

You can issue SRDF Host Component commands to both local and remote storage systems. Commands destined for remote storage systems are transmitted through local storage systems using SRDF links. Configuration and status information can be viewed for each device on each storage system that contains SRDF devices.

There are user interfaces for the SRDF Host Component for the batch commands and through the system console.

SRDF/A Monitor

SRDF/A Monitor is a facility for managing and monitoring SRDF/A operations. It is a component of the ResourcePak Base for z/OS. SRDF/A Monitor:

- Discovers storage systems that are running SRDF/A and monitors the state of the SRDF/A groups
- Collects and writes System Management Facility (SMF) data about the SRDF/A groups
- Optionally, calls a user exit to perform user-defined actions when it detects a change in the state of a SRDF/A group
- Optionally, invokes SRDF/A automatic recovery procedures to recover a dropped SRDF/A session

WPA Monitor

SRDF/A Write Pacing extends the availability of SRDF/A by enabling you to prevent conditions that can result in cache overflow. The SRDF/A Write Pacing Monitor, a component of the ResourcePak Base for z/OS, gathers information about write pacing activities in a storage system. The data is collected for each:

- SRDF/A group by the storage system
- SRDF device by the SRDF group and the storage system

The data includes:

- Changes in the ARMED state by device
- Total paced delay by device
- Total paced track count by device
- Changes in the ENABLED/SUPPORTED/ARMED/PACED state for the SRDF/A group
- Total paced delay for the SRDF/A group
- Total paced track count for the SRDF/A group

The WPA Monitor writes the collected information as SMF records.

SRDF/AR process management

SRDF/AR automates data copying across SRDF links to provide a logically consistent, restartable image of data at a remote (recovery) site. That image can be used should a disaster occur at the production site.

SRDF/AR automatically propagates the restartable image to the recovery site in a way that is transparent to the host application or database. The result is a series of consecutive data *consistency points* that you use as the basis for restarting host applications at the recover site.

The ResourcePak Base for z/OS and TimeFinder/Mirror for z/OS components of Mainframe Enablers provide commands to configure, manage, monitor, start, pause, restart, and stop SRDF/AR processes.

Geographically Dispersed Disaster Restart (GDDR)

GDDR automates business recovery following both planned outages and disaster situations, including the total loss of a data center. Using the PowerMax, VMAX All Flash, or VMAX architecture and the foundation of SRDF and TimeFinder replication families, GDDR eliminates any single point of failure for disaster restart plans in mainframe environments. GDDR intelligence automatically adjusts disaster restart plans based on triggered events.

GDDR does not provide replication and recovery services itself. Rather GDDR monitors and automates the services that other Dell EMC products and third-party products provide that are

required for continuous operations or business restart. GDDR facilitates business continuity by generating scripts that can be run on demand. For example, scripts to restart business applications following a major data center incident, or resume replication following unplanned link outages.

Scripts are customized when invoked by an expert system that tailors the steps based on the configuration and the event that GDDR is managing. Through automatic event detection and end-to-end automation of managed technologies, GDDR removes human error from the recovery process and allows it to complete in the shortest time possible.

The GDDR expert system is also invoked to automatically generate planned procedures, such as moving compute operations from one data center to another. This is the gold standard for high availability compute operations, to be able to move from scheduled DR test weekend activities to regularly scheduled data center swaps without disrupting application workloads.

CHAPTER 8

More information

This chapter shows where there is further information available on some of the subjects mentioned in other chapters. All documents are available from the Dell EMC support web site (<http://www.support.emc.com>).

- [Solutions Enabler CLI](#)..... 106
- [Unisphere](#)..... 106
- [Mainframe Enablers](#)..... 106
- [GDDR](#)..... 106
- [SRDF/TimeFinder Manager for IBM i](#)..... 106
- [SRDF/Metro vWitness](#)..... 106
- [SRDF Interfamily Compatibility](#)..... 106
- [Storage arrays](#)..... 107

Solutions Enabler CLI

Solutions Enabler SRDF Family CLI User Guide

Unisphere

Unisphere for PowerMax Online Help

Unisphere for VMAX Online help

Unisphere for PowerMax REST API Concepts and Programmer's Guide

Unisphere for VMAX REST API Concepts and Programmer's Guide

Mainframe Enablers

SRDF Host Component for z/OS Product Guide

ResourcePak Base for z/OS Product Guide (contains information about SRDF/A Monitor, WPA Monitor, and SRDF/AR process management)

TimeFinder/Mirror for z/OS Product Guide (contains information about configuring, managing, and monitoring SRDF/AR)

AutoSwap for z/OS Product Guide

Consistency Groups for z/OS Product Guide

GDDR

GDDR for SRDF/Star Product Guide

GDDR for SRDF/Star with AutoSwap Product Guide

GDDR for SRDF/Star-A Product Guide

GDDR for SRDF/SQAR with AutoSwap Product Guide

GDDR for SRDF/A Product Guide

GDDR for SRDF/S with AutoSwap Product Guide

GDDR for SRDF/S with ConGroup Product Guide

SRDF/TimeFinder Manager for IBM i

SRDF/TimeFinder Manager for IBM i Product Guide

SRDF/Metro vWitness

SRDF/Metro vWitness Configuration Guide

SRDF Interfamily Compatibility

SRDF and NDM Interfamily Connectivity Information

Storage arrays

PowerMax Family Product Guide

VMAX All Flash Product Guide

VMAX 3 Product Guide

Symmetrix VMAX Family with Enginuity Product Guide

More information