

IBM Spectrum Accelerate Family Storage
Replication Adapter
Version 3.1.0

User Guide



Note

Before using this document and the product it supports, read the information in [“Notices” on page 49](#).

Edition notice

Publication number: GA32-1067-10. This edition applies to version 3.1.0 of the IBM Spectrum Accelerate Family Storage Replication Adapter and to all subsequent releases and modifications until otherwise indicated in a newer publication.

© **Copyright International Business Machines Corporation 2009, 2020.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures.....	v
Tables.....	vii
About this guide.....	ix
Who should use this guide.....	ix
Conventions used in this guide.....	ix
Related information and publications.....	ix
Getting information, help, and service.....	ix
Ordering publications.....	x
Feedback.....	x
Chapter 1. Introduction.....	1
Concept diagram.....	2
Volume mapping principles.....	2
IBM Spectrum Accelerate Family SRA in the IBM HyperSwap solution.....	3
Compatibility and requirements.....	4
Before you proceed.....	5
Chapter 2. Preparation.....	7
Verifying the mirroring configuration.....	7
Verifying the VMware vCenter SRM installation.....	8
Setting up site-to-site mirroring from scratch.....	9
Stage 1: Establish mirroring connection.....	10
Stage 2: Create a storage pool for the protected volumes.....	13
Stage 3: Create a storage pool for the recovery volumes.....	13
Stage 4: Create protected volumes.....	14
Stage 5: Create recovery volumes.....	14
Stage 6: Map protected volumes to protected ESXi hosts.....	15
Stage 7: Define recovery ESXi hosts.....	17
Stage 8: Define mirroring for volumes.....	19
Stage 9: Define consistency groups (optional).....	21
Chapter 3. Installation.....	25
Installing the SRA.....	25
Verifying the SRA installation.....	27
Removing the SRA software.....	28
Chapter 4. Usage.....	31
Adding a storage system as an Array Manager.....	31
Initiating SRA operations.....	35
Snapshot creation principles.....	37
Chapter 5. Best practices.....	39
Storage pool snapshot size.....	39
Naming convention.....	39
Volumes in a consistency group.....	39
Immediate log collection.....	39
Reversing replication roles.....	39

SRM operation in the IBM HyperSwap solution.....	40
Planned failover with SRA in an IBM HyperSwap solution.....	40
Unplanned failover with SRA in an IBM HyperSwap solution.....	40
Recovery from rolling disaster with SRA in IBM HyperSwap solution.....	41
Chapter 6. Troubleshooting.....	43
Checking the log file.....	43
Handling warning and error messages.....	44
Notices.....	49
Trademarks.....	50

Figures

1. IBM storage systems in a typical protection and recovery deployment.....	2
2. Best-effort volume mapping principles.....	3
3. IBM Spectrum Accelerate Family SRA with FlashSystem A9000R systems in IBM HyperSwap solution.....	4
4. Targets form for FlashSystem A9000/9000R – Remote System name is "XIV gen4d-64a".....	7
5. Properties information for FlashSystem A9000/9000R – System name is "XIV gen4d-64a".....	8
6. VMware vSphere Client – Site Recovery Manager is installed.....	9
7. System Connectivity form.....	11
8. Connectivity Patch Panel.....	12
9. Create Pool form.....	13
10. Create Volume form.....	14
11. Volume mini-dashboard.....	15
12. Adding new volume mapping to host.....	16
13. Adding new volume mapping for a host.....	16
14. New Host option.....	17
15. Add Host form.....	18
16. Volume mini-dashboard.....	19
17. Adding volume mirroring.....	20
18. Mirroring form.....	20
19. New Consistency Group option.....	21
20. Create Consistency Group form.....	22
21. Create Mirror form – for a CG.....	23
22. Move Mirrored Volumes to a CG option.....	23

23. Installing SAF SRA.....	26
24. Uploading SRA container image.....	26
25. Displaying upload status.....	27
26. Displaying success message.....	27
27. Verifying SAF SRA installation.....	28
28. Deleting SRA from SRM.....	29
29. Confirming to delete SRA.....	29
30. SRA removal success message.....	29
31. Adding array manager.....	31
32. Selecting storage replication adapter.....	32
33. Filling local array manager fields.....	33
34. Filling remote array manager fields.....	33
35. Selecting array pairs to enable.....	34
36. Reviewing settings.....	34
37. Discovering devices.....	35
38. Array Based Replication – Storage Replication Adapters tab.....	37
39. Array Based Replication – Array Pairs tab.....	37
40. Recovery Plans – History tab.....	38
41. Downloading SRM logs.....	43
42. Exporting SRM logs.....	44
43. Downloading generated logs.....	44
44. Saving log file locally.....	44

Tables

1. IBM Spectrum Accelerate Family SRA operation types..... 35

2. SRM operations and snapshot creation..... 38

3. IBM Spectrum Accelerate Family SRA warning messages..... 45

4. IBM Spectrum Accelerate Family SRA error messages..... 45

About this guide

This guide describes how to prepare for, install, configure, and use the IBM Spectrum Accelerate Family Storage Replication Adapter.

Who should use this guide

This guide is intended for system administrators who are familiar with the VMware vCenter, Site Recovery Manager, and vSphere platforms, and with the relevant IBM storage systems.

Conventions used in this guide

These notices are used in this guide to highlight key information.

Note: These notices provide important tips, guidance, or advice.

Important: These notices provide information or advice that might help you avoid inconvenient or difficult situations.



Attention: These notices indicate possible damage to programs, devices, or data. An attention notice appears before the instruction or situation in which damage can occur.

Related information and publications

You can find additional information and publications related to the IBM Spectrum Accelerate Family Storage Replication Adapter on the following information sources.

- [IBM Knowledge Center](#)
- [IBM XIV® Storage System on IBM Knowledge Center](#)
- [IBM Spectrum Accelerate on IBM Knowledge Center](#)
- [IBM FlashSystem® A9000 on IBM Knowledge Center](#)
- [IBM FlashSystem A9000R on IBM Knowledge Center](#)
- [VMware vCenter Site Recovery Manager Documentation](#) (www.vmware.com/support/pubs)
- [VMware Technical Resources](#) (www.vmware.com/technical-resources)
- [VMware knowledgebase](#) (kb.vmware.com)
- For SRA deployment with IBM HyperSwap® solution, also refer to the following VMware documents:
 - [Stretched Clusters & VMware Site Recovery Manager](#)
 - [Site Recovery Manager and Stretched Storage: Tech Preview of a New Approach® to Active-Active Data Centers](#)

Getting information, help, and service

If you need help, service, technical assistance, or want more information about IBM products, you can find various sources to assist you. You can view the following websites to get information about IBM products and services and to find the latest technical information and support.

- [IBM website](#) (ibm.com)
- [IBM Support Portal website](#) (www.ibm.com/storage/support)
- [IBM Directory of Worldwide Contacts website](#) (www.ibm.com/planetwide)

Ordering publications

The IBM Publications Center is a worldwide central repository for IBM product publications and marketing material.

The IBM Publications Center website (www.ibm.com/shop/publications/order/) offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download at no charge. You can also order publications. The publications center displays prices in your local currency.

Sending comments

Your feedback is important in helping to provide the most accurate and highest quality information.

Procedure

To submit any comments about this publication or any other IBM storage product documentation:

- Send your comments by email to ibmkc@us.ibm.com. Be sure to include the following information:
 - Exact publication title and version
 - Publication form number (for example, GA32-1234-00)
 - Page, table, or illustration numbers that you are commenting on
 - A detailed description of any information that should be changed

Chapter 1. Introduction

The IBM Spectrum Accelerate Family Storage Replication Adapter (SAF SRA, or SRA) is a software add-on that integrates with the VMware vCenter Site Recovery Manager (SRM) platform and enables site-to-site failovers together with supported IBM storage systems.

The IBM Spectrum Accelerate Family SRA extends the VMware vCenter Site Recovery Manager capabilities by allowing it to employ the replication and mirroring features for continuous storage availability at both the protected site and recovery site.

Using the IBM Spectrum Accelerate Family SRA, VMware administrators can automate the failover of a supported storage system at the protected (primary) SRM site to a storage system at a recovery (secondary) SRM site. Immediately upon a failover, the ESX and ESXi servers at the recovery SRM site start using the replicated datastores on the mirrored volumes of the secondary storage system.

When the primary site is back online, failback from the recovery site to the primary site can be performed manually or automatically.

IBM Spectrum Accelerate Family SRA can be integrated with IBM HyperSwap solution, as introduced by the FlashSystem A9000 and A9000R storage systems (version 12.1). See [“IBM Spectrum Accelerate Family SRA in the IBM HyperSwap solution”](#) on page 3 for details.

Full description of the IBM HyperSwap solution appears in the Product Overview of the IBM FlashSystem A9000 and A9000R storage systems.

In addition, the SRA is compatible with Multi-site HA/DR solution, as introduced by the IBM FlashSystem A9000 and A9000R storage systems (version 12.3). While not actively supporting this function, the SRA is agnostic to the underlying HA/DR configuration and does not interfere with its operation.

Concept diagram

The following figure illustrates how two supported IBM storage systems are integrated in a typical VMware SRM disaster recovery solution. The IBM Spectrum Accelerate Family SRA is used on the VMware vCenter Site Recovery Manager server at both the protected and recovery sites.




Figure 1. IBM storage systems in a typical protection and recovery deployment

Volume mapping principles

At both the protected and recovery sites, the IBM Spectrum Accelerate Family SRA aims to best achieve the following configuration goals on the supported IBM storage system, if possible:

- Keep the volume LUN numbers upon failover, and reuse the same ones upon failback.
- Map volumes with the same LUN IDs across all the specific ESXi hosts that are requested for the failover.
- In more scenarios, map volumes to a cluster that contains all the specific ESXi hosts that are requested for the failover, rather than mapping to individual hosts.

Note: The above configuration goals are attained on a best-effort basis, depending on the specific configuration scenario at the protected and recovery sites.

The following figure illustrates the volume mapping principles (attained on a best-effort basis):




Figure 2. Best-effort volume mapping principles

For in-depth understanding of how the different IBM Spectrum Accelerate Family SRA operations are carried out, see [“Initiating SRA operations”](#) on page 35.

IBM Spectrum Accelerate Family SRA in the IBM HyperSwap solution

SRM support for active-active storage scenarios can be leveraged with FlashSystem A9000 and A9000R HyperSwap.

HyperSwap high availability is based on active-active pairing of storage systems per volume or per consistency group. Each volume or consistency group pair uses synchronous replication to keep both systems updated at all times. When certain conditions apply, an automatic and completely transparent failover is performed, so that the applications experience no downtime. As soon as the actual failure is recovered, the pair is automatically resynchronized.

As in other high availability solutions, HyperSwap requires a quorum witness component, to avoid split-brain situations. The IBM Spectrum Accelerate Family HyperSwap Quorum Witness is constantly monitoring the status of the related storage systems, and, if necessary, acts as a tiebreaker for conflict resolution.

The HyperSwap solution relies on Asymmetrical Logical Unit Access (ALUA) support to inform the host about the optimized paths to the storage system, and minimize I/O latency. HyperSwap does not require an external solution that provides an active-active solution, and no separate licensing.

With HyperSwap support, SRM stretch clustering capabilities can be leveraged with FlashSystem A9000 and A9000R to:

- Provide cross-site non-disruptive mobility of workloads between active-active datacenters.
- Enable multi-site load balancing
- Increase the availability of workloads through partial or complete site and subsystem failures

Full description of the IBM HyperSwap solution appears in the Product Overview of the IBM FlashSystem A9000 and A9000R storage systems.




Figure 3. IBM Spectrum Accelerate Family SRA with FlashSystem A9000R systems in IBM HyperSwap solution

See “SRM operation in the IBM HyperSwap solution” on page 40 for details on different SRM and HyperSwap deployment scenarios.

Compatibility and requirements

For the complete and up-to-date information about the compatibility and requirements of the IBM Spectrum Accelerate Family Storage Replication Adapter, refer to the latest release notes.

You can find the latest release notes on the [IBM Knowledge Center](#).

Note: Refer to the relevant VMware documentation for information about how to install the compatible versions of vCenter Server and vSphere Client. You should also refer to the latest installation and configuration instructions for ESXi servers.

Important: To learn about the requirements for using VMware vCenter Site Recovery Manager (SRM) at your sites, refer to the Administration Guide that was issued for the SRM version that you are using. The

guide is available at the [VMware vCenter Site Recovery Manager Documentation website](http://www.vmware.com/support/pubs/srm_pubs.html) (www.vmware.com/support/pubs/srm_pubs.html).

Before you proceed

Before you proceed to the Preparation chapter, you must have access to a workstation on which your preferred storage management software is locally installed or accessible via a remote connection.

You can manage the connected storage system using the Hyper-Scale Manager user interface (UI) or CLI (command-line interface).

Chapter 2. Preparation

Preparation is required depending on your specific site configuration.

Prior to installing and using the IBM Spectrum Accelerate Family Storage Replication Adapter (SRA), the following verification tasks are required:

- [Verifying the mirroring configuration](#)
- [Verifying the VMware vCenter SRM installation](#)

To learn about how to prepare your sites from scratch, refer to [“Setting up site-to-site mirroring from scratch”](#) on page 9.

Verifying the mirroring configuration

All storage systems, volumes, and ESXi hosts at both the protected (primary) and recovery (secondary) sites must be properly connected to their remote counterparts and configured for site mirroring.

Prior to installing and using the IBM Spectrum Accelerate Family SRA, make sure that:

- Your local storage system at the protected site has mirroring connectivity with the target storage system at the recovery site.
- The name of each storage system is unique in both the protected and recovery sites.
- The **Target Name** of any remote storage system is identical to the predefined System Name of that same remote storage system, as illustrated below.




Figure 4. Targets form for FlashSystem A9000/9000R – Remote System name is "XIV gen4d-64a"

Units: GB ▼

↑

Name *

XIV gen4d-64a

Version

12.0.1

System Allocation (GB)

1033

System Effective Capacity (GB)

9819

System S/N (ID)

6013794

Machine Model / Machine Type

415/9835

IP / Hostname

Gen4D-64d

SDS System

No

Figure 5. Properties information for FlashSystem A9000/9000R – System name is "XIV gen4d-64a"

- The storage pools that contain the replicated volumes at both the protected and recovery sites have sufficient size for creating the snapshots of all replicated volumes concurrently. For IBM XIV and Spectrum Accelerate storage systems, allocate a minimum of 17 GB for each volume, with additional 17 GB as spare space.
- Remote mirroring is defined for all volumes that you intend to protect as part of your Disaster Recovery Plan (DRP).
- Your protected (primary) volumes are mapped to the protected ESXi hosts.
- The recovery (secondary) volumes remain unmapped.
- Your recovery (secondary) ESXi hosts are defined as storage system hosts at the recovery site.
- Your recovery ESXi host ports are FC-zoned with the storage system at the recovery site, and are visible by that storage system.



Attention: Confirm with your storage administrator that all the requirements above are met.

Note: To learn how to set up mirroring in a "start from scratch" scenario, refer to [“Setting up site-to-site mirroring from scratch”](#) on page 9.

Verifying the VMware vCenter SRM installation

Before installing the IBM Spectrum Accelerate Family Storage Replication Adapter, make sure that VMware vCenter Site Recovery Manager is already installed and accessible at the protected (primary) site, as well as at the recovery (secondary) site.

On the vSphere client application, go to the Home page and check that the **Site Recovery** icon is displayed under **Solutions and Applications**.




Figure 6. VMware vSphere Client – Site Recovery Manager is installed

Setting up site-to-site mirroring from scratch

This section covers the specific scenario when no previous setup for mirroring has been performed at your sites.

The specific scenario of setting up mirroring from scratch includes the following configuration stages:

- “[Stage 1: Establish mirroring connection](#)” on page 10
- “[Stage 2: Create a storage pool for the protected volumes](#)” on page 13
- “[Stage 3: Create a storage pool for the recovery volumes](#)” on page 13
- “[Stage 4: Create protected volumes](#)” on page 14
- “[Stage 5: Create recovery volumes](#)” on page 14
- “[Stage 6: Map protected volumes to protected ESXi hosts](#)” on page 15
- “[Stage 7: Define recovery ESXi hosts](#)” on page 17
- “[Stage 8: Define mirroring for volumes](#)” on page 19
- “[Stage 9: Define consistency groups \(optional\)](#)” on page 21

Note:

- If your sites are already partially configured for mirroring, perform only the stages or steps that are relevant to your specific case.
- The following sections provide setup examples from version 5.0 of the Hyper-Scale Manager UI. If you are using an older or a newer management package version, refer to its online help or documentation.

Stage 1: Establish mirroring connection

Your first task is to determine which of your storage systems should be used as the protected (primary) system, and which should be used as the recovery (secondary) system.

About this task

Complete the following procedure to define a mirroring connection with the secondary storage system, using the IBM Hyper-Scale Manager UI.

Important: The target storage system at the remote site must already be up and running, as well as added to the IBM Hyper-Scale Manager UI before you can set the mirroring connection to it. In addition, all iSCSI or FC connections to the remote storage system must be operational as well. Confirm with your storage administrator that the remote storage system is ready for establishing the mirroring connection.

Procedure

1. Select **System and Domains Views > System Connectivity**.
2. From the Hub, click on **Targets**.
The list of targets is displayed in the **System Connectivity** form. If there are no targets connected to this system, the display will be empty.
3. From the top of the **System Connectivity** form, click on the plus (+) sign to add a target.
The System Connectivity form is displayed.

The screenshot shows the 'System Connectivity' configuration window. At the top, there's a 'TARGETS' section with a '+' icon. Below it, a status bar indicates 'Status: Creating New Connectivity' with a 'Define Connectivity' button. The main configuration area is divided into two columns. The left column is for the 'Source System', which is 'XIV gen4d-64a'. It has a 'Source Domains' section with a '+' icon and a 'Source Domain' field containing '/Global Space/'. The right column is for the 'Remote System', which is currently empty and marked with a red asterisk. It has a 'Remote Domains' section with a '+' icon. Below these columns is the 'TARGET PARAMETERS' section. It contains two rows of radio buttons: 'Migration' and 'Mirroring / Mobility' (selected), and 'FC' and 'iSCSI' (selected). At the bottom, there are three input fields: 'Max Sync Job Rate (MB/s)' with a value of 300, 'Max Resync Rate (MB/s)' with a value of 300, and 'Max Initialization Rate (MB/s)' with a value of 100.

Figure 7. System Connectivity form

4. Add a remote system that is either monitored by management server or unmonitored and define all of its parameters, including **Mirroring/Mobility**. Click **Apply**.
5. Click **Define Connectivity** to view the Connectivity Patch Panel displaying the defined ports on the source and target systems. A dotted line between the source and target systems shows that the connectivity is not yet defined.
6. Click **Auto Connect**.
The system sees all of the ports on the target and connects the ports. You can switch between the source and target views on the Connectivity Patch Panel to see the ports and the port status defined on both.
7. To see a graphical image of the Connectivity Patch Panel and the connections between the ports of the source and target systems, click on **Connectivity Details**.




Figure 8. Connectivity Patch Panel

The auto-detected physical connections (iSCSI or FC) between the two storage systems are displayed graphically as green arrow lines between the interface connections of both storage systems.

- Make sure that a bi-directional connection is established with the target storage system.
- Fibre Channel (FC) connections can be auto-detected only through proper FC zoning. Confirm with your storage administrator that FC zoning has been properly set in advance.

8. Click **Approve** to use the graphically displayed connections.

Stage 2: Create a storage pool for the protected volumes

In this preparation stage you create a storage pool for the protected volumes.

About this task

After the mirroring between the storage systems is set, create a storage pool for the protected volumes on the primary storage system.

Procedure

1. Select **Pools and Volumes Views > Pools**.
The list of pools is displayed. If there are no pools in this system, the display will be empty.
2. From the Hub, click on the plus (+) sign to add a new pool.
The **Create Pool** form is displayed.




Figure 9. Create Pool form

3. Define the pool name, its size, primary storage system and domain.
4. Click **Create**.
The storage pool is created.

Stage 3: Create a storage pool for the recovery volumes

In this preparation stage you create a storage pool for the recovery volumes.

Procedure

- Complete the pool creation procedure (see [“Stage 2: Create a storage pool for the protected volumes” on page 13](#)) at the recovery (secondary) site as well. When you do so, you may want to give an identical name to the storage pool at the recovery site.

Stage 4: Create protected volumes

In this preparation stage you create protected volumes.

About this task

After the storage pool is created, create the protected volumes on it.

Procedure

1. Select **Pools and Volumes Views > Volumes**.

The list of volumes is displayed. If there are no volumes in this system, the display will be empty.

2. From the Hub, click on the plus (+) sign to add a new volume.

The **Create Volume** form is displayed.




Figure 10. Create Volume form

3. Define the volume name, the number of volumes, volume size, primary storage system, its domain and the pool you created in the previous stage.

4. Click **Create**.

The volume is created.

Stage 5: Create recovery volumes

In this preparation stage you create recovery volumes.

Procedure

- Complete the volume creation procedure (see: [“Stage 4: Create protected volumes”](#) on page 14) at the recovery (secondary) site as well. When you do so, keep in mind:
 - You might want to give identical names to the volumes at the recovery site.
 - The volumes at the recovery site must have the same size as the volumes at the protected site.

Stage 6: Map protected volumes to protected ESXi hosts

In this preparation stage you map protected volumes to protected ESXi hosts.

About this task

After you have created protected volumes, map the volumes to the primary ESXi hosts that are to be protected by SRM, and then rescan these hosts.

Procedure

1. Select **Pools and Volumes Views > Volumes**.
The list of volumes is displayed.
2. Select one or more protected volumes to which you want to map a host.
The **Volume** mini-dashboard is displayed.



Figure 11. Volume mini-dashboard

3. From the **Volume** mini-dashboard, click the **Mapping** relation icon to display the empty **Volume Mapping** form with the **ADD** icon.
4. In the **Volume Mapping** form, click **ADD** icon to map the volume to an existing host.




Figure 12. Adding new volume mapping to host

The **Host** and **LUN** drop-down lists are displayed in **Volume Mapping** form.

The screenshot shows the 'Volume Mapping' window with a dark blue background. At the top left is a back arrow icon. The title 'Volume Mapping' is at the top center. On the top right, there are two tabs: 'Host' (active, highlighted in teal) and 'Cluster'. Below the tabs is a section titled 'HOSTS' with a teal plus icon to its right. Below 'HOSTS' is a form with two drop-down lists: 'Host' and 'LUN'. Both labels have a red asterisk indicating they are required. Each drop-down list has a teal arrow icon to its right. At the bottom right are 'Cancel' and 'Apply' buttons.

Figure 13. Adding new volume mapping for a host

5. In the **Volume Mapping** form, select a host and a LUN to be mapped, using the respective drop-down lists.
6. Click the plus sign in the right-hand corner if you want to continue the host mapping.
7. Click **Apply** to finish the mapping.

Stage 7: Define recovery ESXi hosts

In this preparation stage you define recovery ESXi hosts.

About this task

After mapping the protected volumes to the protected ESXi hosts, you can start defining your recovery (secondary) ESXi hosts at the recovery site.

Important: The recovery ESXi hosts should be connected over iSCSI or FC to the recovery storage system. For FC, proper zoning must be predefined.

Procedure

1. Click the **New** icon on the toolbar at the top and then select **Host**.




Figure 14. New Host option

The **Add Host** form is displayed.

The screenshot shows the 'Add Host' form with the following fields and sections:

- Name ***: A text input field.
- Type**: A dropdown menu with 'Default' selected.
- PORTS**: A section with a '+' icon. It contains:
 - FC** and **iSCSI** radio buttons.
 - Port Address ***: A dropdown menu.
- SYSTEMS / CLUSTERS / DOMAINS**: A section with a '+' icon. It contains:
 - System ***: A dropdown menu.
 - Cluster**: A dropdown menu.
 - DOMAINS**: A list with a '+' icon.
 - CLUSTER'S DOMAINS**: A list.
- Buttons**: 'Cancel' and 'Apply' buttons at the bottom right.

Figure 15. Add Host form

2. Enter the following host details:

- **Name** – Define name of the host.
- **Type** – Select the host type from the drop down list.
- **PORTS** – Select **FC** or **iSCSI** to define the port type. Depending on the selected port type, set the port address as follows:
 - For an FC port – Select the World Wide Port Name (WWPN) of the port from the drop down list.
 - For an iSCSI port – Enter the iSCSI Qualified Name (IQN) of the port.
- **System** – Select the secondary storage system to which you want to attach the host.
- **Cluster** – If the host is a member of a cluster, select the cluster name from the drop down list. Otherwise, leave this option blank.
- **Domain** – Select the storage system domain or domains in which this host should be defined.
- **Cluster's Domain** – Select the cluster domain or domains in which this host should be defined.

3. Click **Apply**. The host definition is added to the storage system.

4. Repeat steps 1–3 for each recovery host at the recovery site.

Note: The procedure above is for a single host. If some hosts are part of a cluster, go to **Hosts and Clusters Views > Hosts**, hold down the SHIFT key and select these hosts. Then, from the Hub and click **Actions > Clusters > Add/Remove Cluster** to add the hosts to the cluster.

Stage 8: Define mirroring for volumes

In this preparation stage you define mirroring for volumes on IBM storage systems.

About this task

After the recovery hosts are defined as storage system hosts at the recovery site, you can start defining the volume mirroring. The mirroring direction is from the protected site to the recovery site.

Procedure

1. Select **Pools and Volumes Views > Volumes**.
The list of volumes is displayed.
2. Select one or more volumes to be mirrored.
The **Volume** mini-dashboard is displayed.



Figure 16. Volume mini-dashboard

3. From the **Volume** mini-dashboard, click the **Mirror** relation icon to display the empty **Volume Mirroring** form with the **ADD** icon.
4. In the **Mirroring** form, click **ADD** icon to create mirroring.




Figure 17. Adding volume mirroring

The full **Mirroring** form is displayed.

A dark-themed user interface window titled 'Mirroring'. At the top left is a small circular icon with an upward arrow. Below it, the text 'Mirroring status: Creating' is followed by an information icon. The form contains several sections: 'Destination System (Not monitored)' with a dropdown menu showing 'XIV gen4d-64b'; 'Destination Domain' with a dropdown menu showing '/Global Space/'; 'Destination Pool' with a dropdown menu showing 'pool1'; 'Auto Create Volume' with two radio buttons, 'Auto Create Volume' (selected) and 'Choose Existing'; 'Destination Volume' with a text input field containing 'brkdVOLUME'; and 'Sync' with two radio buttons, 'Sync' (selected) and 'Async'. At the bottom right, there are two buttons: 'Cancel' and 'Apply'.

Figure 18. Mirroring form

5. In the **Mirroring** form, select the destination system, domain, pool, and other relevant mirroring details.
6. Click **Apply** to finish the mirroring configuration.

Stage 9: Define consistency groups (optional)

If your volumes require replication in consistency with each other (for example, when a datastore consists of more than one volume), create a consistency group for these volumes.

About this task

Define a mirrored consistency group on a storage system, which will contain previously created mirrored volumes.

Procedure

1. Click the **New** icon on the toolbar at the top and then select **Consistency Group**.





Figure 19. New Consistency Group option

The **Create Consistency Group** form is displayed.



The image shows a 'Create Consistency Group' form in a dark-themed interface. At the top, the title 'Create Consistency Group' is displayed. Below the title is a diagram showing a central circle labeled 'CG' with a 'NEW' tag, connected by dashed lines to four surrounding circles containing icons for different storage systems: a database, a folder, a server rack, and a network switch. Below the diagram is a form with four fields: 'Name' (text input), 'System' (dropdown menu), 'Domain' (dropdown menu with an information icon), and 'Pool' (dropdown menu). Each field has a red asterisk indicating it is required. At the bottom right of the form are 'Cancel' and 'Create' buttons.

Figure 20. Create Consistency Group form

2. Enter the consistency group's name, storage system, domain and select the storage pool on which it should be created. Then, click **Create**.
3. Define mirroring for the consistency group (CG) you have created. The mirroring should be of the same type and characteristics of the mirrored volumes that you intend to add to the group. Then, click **Apply**.

Mirroring status: Creating ⓘ

Destination System *

NO_CODE-form.mirroring.destination.label.cg_name

☐ Sync
☒ Async

NO_CODE-form.mirror.label.... ⓘ *
☐ Automatically set Schedule ⓘ
☐ Manually set Schedule

☒ Activation State: Active ⓘ
☐ Activation State: Inactive

Cancel Apply

Figure 21. Create Mirror form – for a CG

4. Select **Pools and Volumes Views > Volumes**.
The list of volumes is displayed.
5. Select one or more mirrored volumes which you want to add to the mirrored consistency group.
6. From the Hub, select **Actions > Consistency Group > Move to a Group**.




Figure 22. Move Mirrored Volumes to a CG option

The mirrored volumes are now part of the mirrored consistency group.

Chapter 3. Installation

After the required preparation has been performed, you can start the IBM Spectrum Accelerate Family SRA installation.

This chapter describes:

- [“Installing the SRA” on page 25](#)
- [“Verifying the SRA installation” on page 27](#)
- [“Removing the SRA software” on page 28](#)

Installing the SRA

Use this procedure to install the IBM Spectrum Accelerate Family Storage Replication Adapter.

Before you begin

Do not upgrade Spectrum Accelerate Family SRA while the VMware SRM service is running and there are pending operations related to the Spectrum Accelerate Family SRA to be updated. Always check VMware SRM service status, and stop the service if service is active.

About this task

Complete the following procedure on the VMware vCenter SRM server to install the IBM Spectrum Accelerate Family SRA.

Note: The Storage Replication Adapter 3.1.0 is installed on the Site Recovery Manager (SRM) Virtual Appliance.

Procedure

1. From the SRM Appliance Management interface, navigate to **Storage Replication Adapters** tab and click **NEW ADAPTER**.




Figure 23. Installing SAF SRA

2. On the next window, click **UPLOAD**.




Figure 24. Uploading SRA container image

3. Select and upload locally downloaded SRA container image `ibm-xiva9000-sra.3.1.0.tar.gz`.
The next window displays the upload status. After the SRA container image is successfully uploaded, a new window displays the success message.




Figure 25. Displaying upload status




Figure 26. Displaying success message

What to do next

Verify that the IBM Spectrum Accelerate Family SRA is properly installed, as explained in [“Verifying the SRA installation”](#) on page 27.

Verifying the SRA installation

Verify the IBM Spectrum Accelerate Family SRA installation.

Procedure

1. From the vSphere client application, on the **Site Pair** tab, navigate to **Configure > Array Based Replication > Storage Replication Adapters** page.
2. Click **RESCAN ADAPTERS**.

vmw

Site Recovery

vmlab-vcsa-1.ssd.hursley.ibm.com - vmlab-vcsa-2.ssd.hursley.ibm.com
P-Site - DR-Site

Site Pair

Protection Groups

Recovery Plans

Summary

Issues

Configure

Array Based Replication

Storage Replication Adapters

Array Pairs

Network Mappings

Folder Mappings

Resource Mappings

Storage Policy Mappings

Placeholder Datastores

Advanced Settings

Permissions

Recovery Plans History

Storage Replication Adapters

vmlab-vcsa-1.ssd.hursley.ibm.com
vmlab-vcsa-2.ssd.hursley.ibm.com

RESCAN ADAPTERS

IBM Spectrum Accelerate Family SRA

Status	✓ OK
Version	3.1.0
Vendor	IBM Corp.
Install Location	ibm-xiva9000-sra:3.1.0
Vendor URL	https://www.ibm.com/support/knowledgecenter/SSEQ8L/landing
Supported Array Models	IBM Corp., IBM XIV IBM Corp., IBM Spectrum Accelerate IBM Corp., IBM FlashSystem A9000/A9000R
Supported Software	IBM XIV Remote Mirroring 10.2 IBM XIV Remote Mirroring 11 IBM FlashSystem A9000/A9000R Remote Mirroring 12
Stretched Storage	Supported

Figure 27. Verifying SAF SRA installation

3. Verify that correct version details are displayed for the IBM Spectrum Accelerate Family SRA.

Removing the SRA software

Remove the IBM Spectrum Accelerate Family SRA from the SRM.

Procedure

1. From the SRM Appliance Management interface, navigate to the **Storage Replication Adapters** page.
2. Click **Overflow** menu icon and click **Delete**.




Figure 28. Deleting SRA from SRM

- On the next window, select the two check boxes and click **DELETE**.
The success message is displayed at the lower right area of the window.




Figure 29. Confirming to delete SRA



Figure 30. SRA removal success message

Chapter 4. Usage

This chapter describes the usage of the IBM Spectrum Accelerate Family SRA together with VMware SRM for performing dependable Disaster Recovery Planning (DRP) and protecting VMware sites that utilize supported IBM storage resources.

This includes:

- “Adding a storage system as an Array Manager” on page 31
- “Initiating SRA operations” on page 35
- “Snapshot creation principles” on page 37

Adding a storage system as an Array Manager

After the IBM Spectrum Accelerate Family SRA is installed, you can start adding IBM storage systems (referred to as *array managers*) to your protected and recovery sites, as described in the following procedure.

Procedure

1. From the vSphere client application, on the **Site Pair** tab, navigate to **Configure > Array Based Replication > Array Pairs**.




Figure 31. Adding array manager

2. On **Array Pairs** tab, click **+ ADD** to launch **Add Array Pair** page.
3. On **Storage replication adapter** tab of **Add Array Pair** page, select a storage replication adapter and click **NEXT**.




Figure 32. Selecting storage replication adapter

4. On **Local array manager** tab, enter a name for the array manager and fill the following **Storage system connectivity parameters** fields. Click **NEXT** after filling all the fields.

- **First management IP address / hostname**
- **Second management IP address / hostname**
- **Third management IP address / hostname**
- **Username**
- **Password**

Note:

- **First management IP address / hostname** is the primary IP address or hostname for the storage system.
- **Second management IP address / hostname** and **Third management IP address / hostname** are the two alternative IP addresses.
- **Username** and **Password** fields require the administrator user name and password for accessing the storage system.

Important:

- You must fill all the three fields for adding IP addresses. Fill the primary IP address in the First field. If there is only one alternative IP address, add that in both the Second and Third alternative address fields. If there is no alternative IP address, use the primary IP address in all the three fields.

Figure 33. Filling local array manager fields

5. On **Remote array manager** tab, enter a name for the array manager and fill the **Storage system connectivity parameters** fields, like in previous step. Click **NEXT** after filling all the fields.

Important: DO NOT select **Do not create a remote array manager now.** check box if you want to create a remote array manager in this step. If you select this check box, you will be required to create a remote array manager from the remote SRM.

Figure 34. Filling remote array manager fields

6. On **Array pairs** tab, select the applicable check boxes to enable the array pairs and click **NEXT**.




Figure 35. Selecting array pairs to enable

- On **Ready to complete** tab, review the selected settings for array manager and click **FINISH**.
The next window displays the array pairs.




Figure 36. Reviewing settings

- Select an array pair to view its devices. If newly replicated devices are not displayed, click **DISCOVER DEVICES**.
All the replicated devices for the selected array pair are now displayed on the screen.

Array Pairs

[Learn more](#)

[+ ADD](#) [ARRAY PAIR](#) [ARRAY MANAGER PAIR](#)

	Array Pair	Array Manager Pair	Last Array Manager Ping
	SME587 ↔ SME891	R-Site ↔ P-Site	Success, 12/6/20, 2:21:26 AM -0800
	SME587 ↔ SME901	R-Site ↔ DR-HS	Success, 12/6/20, 2:21:26 AM -0800
	SME891 ↔ SME587	PSite ↔ RSite	Success, 12/7/20, 2:00:41 AM -0800
	SME891 ↔ SME901	PSite ↔ DR-HS	Success, 12/7/20, 2:00:41 AM -0800

4 items

DISCOVER DEVICES

Device (192.168.150.73)	Datstore	Status	Device (192.168.150.73)
Altran-SRA-HS-TC2_001		Failover Complete (Stretche...	Altran-SRA+
Altran-SRA-HS-TC2_002		Failover Complete (Stretche...	Altran-SRA+
Altran-SRA-HS-TC2_003		Failover Complete (Stretche...	Altran-SRA+

Figure 37. Discovering devices

Initiating SRA operations

The IBM Spectrum Accelerate Family SRA supports different VMware vCenter SRM operations, as detailed in the following table.

Table 1. IBM Spectrum Accelerate Family SRA operation types	
Operation type	Description or reference
Add supported IBM storage systems	See “Adding a storage system as an Array Manager” on page 31.
Refresh volume information (display updated information)	See Figure 39 on page 37.
Create protection groups and recovery plans for storage volumes and datastores that utilize these volumes.	See Figure 38 on page 37.
Test (failover test)	Test the storage replication configuration by creating snapshots at the recovery site and then mapping these snapshots, in the following manner: - In case of synchronous mirroring, the snapshots are created at the recovery site and are mapped to the ESXi host or hosts. - In case of asynchronous mirroring, the last replicated snapshot is duplicated at the recovery site and then mapped to the ESXi host or hosts. - At the end of the failover test workflow, these temporary snapshots are unmapped and deleted.
Cleanup	Delete existing snapshots after the Test operation. Only the snapshots that were created by the IBM Spectrum Accelerate Family SRA are deleted.

Table 1. IBM Spectrum Accelerate Family SRA operation types (continued)

Operation type	Description or reference
Recovery (failover)	Switch the operation to the recovery site in case of a planned migration or an unplanned disaster (when the protected site is unavailable), in the following manner: 1. At the recovery site, the Spectrum Accelerate Family SRA creates snapshots of the volumes as backup. If the recovery is part of a planned migration (as opposed to unplanned disaster), the protected site volumes are first locked to read-only state. 2. The Spectrum Accelerate Family SRA promotes the recovery site volumes to Master volumes (role). 3. The Spectrum Accelerate Family SRA maps these volumes to the ESXi host or hosts.
Reprotect	Reverse the replication direction and mirroring from failed-over devices to the primary protected devices, in the following manner: 1. The Spectrum Accelerate Family SRA creates a fail-safe snapshot of the previously protected volumes. 2. The Spectrum Accelerate Family SRA unmaps the volumes at the original protected site and sets these volumes as the mirror targets. 3. The Spectrum Accelerate Family SRA changes the role of the previously protected volumes to Slave volumes (role). 4. The Spectrum Accelerate Family SRA activates the mirroring and restores the replication between the two sites.

Note: Depending on the mirroring type:

- If the mirroring is asynchronous, the IBM Spectrum Accelerate Family SRA starts the synchronization job before the initiation of the Test, Recovery, and Reprotect operations.
- If the mirror is synchronous, the IBM Spectrum Accelerate Family SRA only verifies that the mirroring state is indeed synchronized before initiating the Test, Recovery, and Reprotect operations.
- For IBM HyperSwap solution, if there is no target connectivity between the peers and the recovery storage system is in the Slave role, the IBM Spectrum Accelerate Family SRA cannot perform failover from the protected to recovery site.

Tip:

- To better visualize the Test, Recovery, and Reprotect operations, see [“Concept diagram”](#) on page 2.
- For more information about the snapshot aspects of Test, Recovery, and Reprotect, see [“Snapshot creation principles”](#) on page 37.

The following figures show the typical contents of the following tabs, which are available when a specific storage system is selected on the Array Based Replication panel:

- **Summary** – Displays general details regarding the selected storage system and its interoperability with the IBM Spectrum Accelerate Family SRA
- **Manage** – Displays the details of storage systems that are paired with the selected storage system. You can enable or disable the pairing as needed. In addition, it displays the details of mirrored volumes that are defined on the selected storage system, including information regarding the mirroring direction, remote volume, related datastore, protection group and consistency group.




Figure 38. Array Based Replication – Storage Replication Adapters tab




Figure 39. Array Based Replication – Array Pairs tab

For more detailed information about how to operate Site Recovery Manager and perform complete DRP for your VMware server sites, refer to the relevant VMware vCenter Site Recovery Manager documentation.

Snapshot creation principles


Snapshots are created when the following SRM operations are initiated: Test, Recovery, and Reprotect.

Table 2 on page 38 summarizes the snapshot creation principles per SRM operation type.

Table 2. SRM operations and snapshot creation		
VMware vCenter SRM operation	Results	Snapshot name format
Test (a site-to-site failover test is performed)	Snapshots of the target volumes are created, unlocked for read-write, and then mapped and used at the recovery site. Note: These volume snapshots are unmapped and deleted when the Cleanup operation is initiated after the Test operation is concluded.	sra_test_<volume name>_<time stamp>
Recovery (a site-to-site failover is performed)	Snapshots of the target volumes are created at the recovery site before these volumes are opened for read-write operations.	sra_synced_<volume name>
Reprotect (the previously protected site becomes the recovery site)	Snapshots of the original source volumes are created before these volumes can be used as the mirroring target.	sra_synced_<volume name>

Note: When using consistency groups, the name of each consistency group replaces the volume in the snapshot name. For more information about consistency groups, refer to the relevant IBM storage system documentation.

To view the history log of these SRM operations, navigate to **Recovery Plans > History**.



Operation	Result	Date	Duration	User
Reprotect	Success	Wednesday, December 2, 2020 5:23:21 AM	23 s	VSPHERE.LOCAL\Administrator
Recovery	Success	Wednesday, December 2, 2020 5:17:42 AM	2 m 2 s	VSPHERE.LOCAL\Administrator
Reprotect	Success	Wednesday, December 2, 2020 5:12:28 AM	50 s	VSPHERE.LOCAL\Administrator
Recovery	Success	Wednesday, December 2, 2020 5:10:40 AM	15 s	VSPHERE.LOCAL\Administrator

Figure 40. Recovery Plans – History tab

Chapter 5. Best practices

This chapter summarizes recommended practices when using the IBM Spectrum Accelerate Family SRA. Consider the following recommendations for:

- [Storage pool snapshot size](#)
- [Naming convention](#)
- [Volumes in a consistency group](#)
- [Immediate log collection](#)
- [Reversing replication roles](#)
- [IBM Spectrum Accelerate Family SRA in the IBM HyperSwap solution](#)

Storage pool snapshot size

Consider allocating extra storage pool snapshot space for storage volumes on XIV and Spectrum Accelerate storage systems.

As mentioned in “Verifying the mirroring configuration” on page 7, the storage pools that contain the replicated volumes at both the protected and recovery sites should have sufficient size for creating the volume snapshots.

For XIV and Spectrum Accelerate storage systems, the requirement is to allocate a minimum of 17 GB for each volume, with additional 17 GB as spare space. No such limitation exists for FlashSystem A9000 and A9000R.

As a best practice, if your volumes are working in a high write rate, consider allocating larger storage pool space for snapshots.

Naming convention

Naming storage pools and volumes across different sites might need identical names.

To avoid confusion and to attain a 100% duplicated recovery site, give identical names to the storage pools and volumes at both the protected and recovery sites.

Volumes in a consistency group

If the volumes of a single datastore are not grouped together, datastore consistency issues may arise.

To ensure datastore consistency in Recovery or Test operations, place all volumes of a single datastore in one consistency group. For more information about consistency groups, refer to the relevant IBM storage system documentation.

Immediate log collection

The SRA log information might be overwritten due to fast log filling and rotation.

To ensure that the full information is provided to IBM Support in case of a technical issue, it is recommended to collect the SRA log immediately, as explained in [“Checking the log file”](#) on page 43.

Reversing replication roles

When mirroring is enabled, the master volume or consistency group is designated as 'primary' and the slave volume or consistency group is designated as 'secondary'. If for any reason, the replication roles

have been reversed, the IBM Spectrum Accelerate Family SRA may fail to identify the current replication state properly.

If the administrator manually switches volume or consistency group role in the IBM HyperSwap solution (using the **ha_switch_roles** command), the SRM presents the new state as a failover and considers the VMs as unprotected. In this case, change the highly-available peer designation manually (using the **ha_change_designation** command) to indicate that protection preference has moved to the other side. Alternatively, switch the volume/consistency group role to its original state and allow SRM to protect the VMs on the protected site.

To optimize host I/O performance, by preferring local vs remote connectivity, use the following XCLI commands:

- **mirror_list -x** to check the current replication role designations.
- **mirror_change_destination** to switch the volume or consistency group replication roles.
- **ha_change_designation** to switch designations of the highly-available peers from primary to secondary, and vice versa.

For full description of the relevant XCLI commands, refer to the relevant IBM storage system documentation.

SRM operation in the IBM HyperSwap solution

SRM support for active-active storage scenarios can be leveraged with FlashSystem A9000 and A9000R HyperSwap.

The following sections describe SRM operation with HyperSwap in these scenarios:

- [“Planned failover with SRA in an IBM HyperSwap solution” on page 40.](#)
- [“Unplanned failover with SRA in an IBM HyperSwap solution” on page 40.](#)
- [“Recovery from rolling disaster with SRA in IBM HyperSwap solution” on page 41.](#)

Planned failover with SRA in an IBM HyperSwap solution

SRM support for planned failover is with A9000 and A9000R HyperSwap.

When the planned migration operation is initiated, SRM uses IBM Spectrum Accelerate Family SRA to instruct the storage systems to:

1. Switch HyperSwap volume roles from Primary to Secondary, and vice versa.
2. Switch HyperSwap volume roles designations from Primary to Secondary, and vice versa.

If the VMotion option is enabled, live VM migration between protected and recovery sites is performed without noticeable downtime. However, for successful VMotion operation, the relevant HyperSwap relations must be active and synchronized. Without VMotion, the VMs are shut down at protected site and then switched on at the recovery site. In this case, a short period of downtime is expected.

Unplanned failover with SRA in an IBM HyperSwap solution

If the protected site goes offline, or its storage system becomes unavailable, you can perform an unplanned failover.

Protected site failure

If there is protected site failure (including FlashSystem A9000 or A9000R storage system, SRM server and vSphere ESXi hosts), the volume at the recovery site becomes preferred. After the failover, you can initiate the SRM disaster recovery to activate VMs on the recovery site. After the protected site goes online, the status of its storage system remains as *Master unavailable*, and there is no mapping between the protected host and the recovery storage system. To re-activate site protection, initiate the SRM reprotect procedure to:

1. Define the recovery site as protected, and vice versa.

2. Switch storage system designations from Primary to Secondary, and vice versa.

Once reprotection is completed, the sites are fully protected again.

The result is a complete failover where now the recovery site has become Primary and it is protected by the formerly protected site. In order to fail back to the formerly protected site, use a planned failover with migration to failover, following by a reprotect.

Protected storage system failure

If a storage system fails at the protected site, or it loses connectivity to its target and the Quorum Witness, the volume at the recovery site becomes preferred. Due to an SRM known issue, the protected VMs become unprotected and cannot be failed over, using SRM. However, the VMs continue to run, using storage at the recovery site. When the protected storage system is back online, use the following manual procedure to restore the initial protection state:

1. Verify that the protected storage became operational again.
2. Change the protected storage system role to Slave (**ha_change_role** command).
3. Activate the HyperSwap relation on the recovery (Master) storage system (**ha_activate** command).
4. Wait for the HyperSwap relation to synchronize.
5. Switch the storage system roles (**ha_switch_roles** command).

Alternatively, use the following manual procedure to move VMs to the recovery site:

1. Power off the VMs on the protected site.
2. Power off the SRM server or service on the protected site.
3. Once SRM server on the protected site is down, initiate a failover to power on the VMs on the recovery site.
4. After the failed protected storage system becomes operational again, power on the SRM server or service if it was powered off at the protected site.
5. Run SRM failover to migrate the VMs to the recovery site.
6. Initiate SRM reprotect to define the recovery site as protected, and vice versa.

Recovery from rolling disaster with SRA in IBM HyperSwap solution

Rolling disaster is a chain of events leading to multiple failures, which in some cases may result in data inconsistency.

For example, losing connectivity over mirroring link leads to claiming of the HyperSwap volume by the protected storage. A subsequent failure of the protected storage itself will result in a potential data loss if SRM failover is triggered. In order to avoid this, IBM Spectrum Accelerate Family SRA prevents the failover and notifies the user that manual intervention is required. To re-activate site protection, use this procedure:

- Verify the protected storage is unavailable and isolated from all ESXi hosts at the protected and recovery sites. This is a critical step to prevent data inconsistency.
- Establish connection to the recovery site storage system.
- Manually change roles of the recovery site volumes from secondary to primary.
- Initiate the SRM failover procedure at the recovery site.

Chapter 6. Troubleshooting

This chapter can help you solve technical problems that you may encounter when using the IBM Spectrum Accelerate Family SRA.

If you encounter an error, refer to the following troubleshooting sections:

- “Checking the log file” on page 43
- “Handling warning and error messages” on page 44

Note: For up-to-date information about known issues and possible workarounds, refer to the latest IBM Spectrum Accelerate Family Storage Replication Adapter release notes.

Checking the log file

Events are recorded separately at each site in a log file (each site with its own log file).

About this task

The IBM Spectrum Accelerate Family SRA log file records events at the local site. When encountering an issue and you are not sure whether the problem is at the local site or the remote site, you can collect a log file from the local site as well as from the remote site. These two log files may provide helpful information.

Note:

- Always attach the log files when opening a new support request as it helps the support team to have more detailed information about the technical problem you encountered.
- Upon any technical issue that you might encounter, try to collect the SRA log immediately or as fast as possible, as it might be overwritten due to rapid log file filling and rotation.

Perform the following procedure to retrieve the SRA log together with other SRM logs at a given site (local or remote).

Procedure

1. From the SRM Appliance Management interface, navigate to the **Summary** page.



Figure 41. Downloading SRM logs

2. Click **DOWNLOAD SUPPORT BUNDLE** to download the logs.

Note: This operation may take several minutes to complete.




Figure 42. Exporting SRM logs

3. Once the logs are generated, click on **DOWNLOAD** to save the server logs locally.




Figure 43. Downloading generated logs




Figure 44. Saving log file locally

Handling warning and error messages

The following tables summarize the different possible warnings or errors you might encounter when working with the IBM Spectrum Accelerate Family SRA.

[Table 3 on page 45](#) lists the warning messages, and [Table 4 on page 45](#) lists the error messages.

Table 3. IBM Spectrum Accelerate Family SRA warning messages

Warning code number	Warning text	Recommended action
500	The replication process for this device is already in progress.	No further action is required.
2001	No access groups were provided during the Failover Test operation. Snapshots were created but could not be mapped.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
2002	Snapshot space in the storage pool may exceed its limit if snapshots are created for the requested groups or devices.	Extend the snapshot size for the pool, or delete old snapshots.
2003	Could not find hosts on the storage system that matches the provided initiators.	Verify that your ESX/ESXi hosts are defined in the storage arrays.
2004	The mapping operation has failed.	Check the SRA log file, request assistance from your storage administrator, or open a support request.

Table 4. IBM Spectrum Accelerate Family SRA error messages

Error code number	Error text	How to resolve
100	Command is not supported. An internal error has occurred.	Open a support request and receive specific guidance from the support team. In your request, attach the relevant log file or files.
101	Locale is not supported. An internal error has occurred.	Change to English locale or open a support request and receive specific guidance from the support team.
1001	Failed to connect.	Make sure that the IP address or hostname is correct, and verify your user credentials.
1002	Already in Test state.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1003	The snapshot for this storage device was not found.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1004	Removal of mapping for this device has failed.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1005	The secondary volume or consistency group is not in consistency state.	Wait for the initialization or synchronization to complete, and then retry the operation.

Table 4. IBM Spectrum Accelerate Family SRA error messages (continued)

Error code number	Error text	How to resolve
1006	Failover operation could not be performed for this device.	Check the SRA log file, request assistance from your storage administrator, or open a support request
1007	The device could not be verified for a failover operation.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1008	The mapping removal and snapshot deletion operations could not be performed.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1009	Could not find the remote mirroring definition for the storage device.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1010	The restore replication operation has failed.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1011	The reverse replication operation has failed.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1012	Could not establish an immediate synchronization.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1013	There is insufficient space allocated in the storage pool for creating new snapshots for this device.	Extend the storage pool, or delete old snapshots.
1014	The operation has failed.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1015	SRA internal error has occurred.	Open a support request and provide the information detailed in the log file.
1016	The failover preparation operation has failed.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1017	There is no mirror connectivity between the protected site and the recovery site.	Request assistance from your storage administrator.
1018	Failed to create snapshot for the storage device.	Check the SRA log file, contact your storage administrator, or open a support request.

Table 4. IBM Spectrum Accelerate Family SRA error messages (continued)

Error code number	Error text	How to resolve
1019	Already in Failover state.	Check the SRA log file, request assistance from your storage administrator, or open a support request.
1020	There are no target mirroring connections for this storage device.	Make sure that the mirroring connectivity for your recovery storage device is properly set. For assistance, contact your storage administrator.
1021	Unsupported storage system microcode version has been detected.	Make sure that the microcode version of the relevant storage system that you are trying to connect is supported. For more information, refer to the release notes.
1022	Insufficient user privileges.	Make sure that the user account you are using has the required privileges.
1023	Mirroring error: inactive mirroring.	Check the SRA log file or use the Hyper-Scale Manager UI to find the cause of this error.

Notices

These legal notices pertain to the information in this IBM Storage product documentation.

This information was developed for products and services offered in the US. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
USA*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119*

Armonk, NY 10504-1785
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of the International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the [Copyright and trademark information website](http://www.ibm.com/legal/us/en/copytrade.shtml) (www.ibm.com/legal/us/en/copytrade.shtml).

VMware, ESX, ESXi, vSphere, vCenter, and Site Recovery Manager are trademarks or registered trademarks of VMware Corporation in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.



Printed in USA

GA32-1067-10

