

InfoScale™ Operations Manager 9.1 Add-ons User's Guide

Last updated: 2025-12-12

Legal Notice

Copyright ©2025 Arctera US LLC. All rights reserved.

Arctera and the Arctera Logo are trademarks or registered trademarks of Arctera US LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners. This product may contain third-party software for which Arctera is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Arctera product or available at:

<https://www.arctera.io/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and de-compilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Arctera US LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. ARCTERA US LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq." Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Arctera as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Arctera US LLC | www.arctera.io

Technical Support

Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within the company to answer your questions in a timely fashion.

Our support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about our support offerings, you can visit our website at the following URL:

www.arctera.io/support

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available here:

<https://sort.veritas.com/arctera/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

productdocs@arctera.io

You can also see documentation information or ask a question on the Arctera community site:

<https://vox.veritas.com/category/arctera-discussions/discussions/infoscale>

Services and Operations Readiness Tools (SORT)

Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1	Storage Provisioning and Enclosure Migration Add-on 9.1	6
Chapter 1	Provisioning storage	7
	About storage provisioning	7
	About creating a storage template	8
	Creating a storage template using VxFS file systems	9
	Create Storage Template – File system parameters panel options	10
	Create Storage Template – File system advanced panel options	10
	Create Storage Template – Volume parameters panel options	13
	Select LUN Characteristics options	14
	Create Storage Template – Volume advanced panel options	14
	Creating a storage template using NTFS file systems	15
	Creating a storage template using volumes	16
	Updating a storage template	17
	Provisioning storage	18
	Uploading storage templates	19
	Downloading storage templates	19
	Deleting storage templates	20
	Locking storage templates	20
	Unlocking storage templates	20
Chapter 2	Migrating volumes	22
	About volume migration	22
	About the Impact Analysis report for volume migration	23
	Migrating volumes by host	24
	Select LUN Characteristics options	26
	Change layout options	26
	Migrating volumes by enclosure	27
	Migrating volumes by disk group	29
	Pausing or resuming a volume migration	31

	Deleting a scheduled volume migration	31
Section 2	Application Migration Add-on	33
Chapter 3	Introduction to Application Migration Add-on	34
	About Application Migration add-on	34
	Before installing the Application Migration add-on	35
	Installing the Application Migration Add-on	35
	Before uninstalling the Application Migration add-on	36
	Uninstalling the Application Migration add-on	36
Chapter 4	Creating and managing an application migration plan	37
	Supported versions and platforms	38
	User privileges	39
	Prerequisites for creating an application migration plan	40
	Prerequisites for migration to AWS	42
	VVR Replication: Environment variables used in application migration	42
	Creating an application migration plan	44
	Understanding user-defined tasks	52
	Understanding application migration operations	53
	Understanding the Setup Storage operation	53
	Understanding the Rehearse operation	55
	Understanding the Migrate operation	56
	Understanding the cleanup operation	57
	Understanding the tasks executed in each operation	58
	Validations performed before migration plan execution	67
	Executing the application migration plan	68
	Editing an application migration plan	71
	Deleting application migration plan(s)	72
	Exporting application migration plan(s)	72
	Importing application migration plan(s)	73
	Viewing historical runs	74
	Viewing properties of an application migration plan	74
	Application migration logs	75
Index		76

Storage Provisioning and Enclosure Migration Add-on 9.1

- [Chapter 1. Provisioning storage](#)
- [Chapter 2. Migrating volumes](#)

Provisioning storage

This chapter includes the following topics:

- [About storage provisioning](#)
- [About creating a storage template](#)
- [Creating a storage template using VxFS file systems](#)
- [Creating a storage template using NTFS file systems](#)
- [Creating a storage template using volumes](#)
- [Updating a storage template](#)
- [Provisioning storage](#)
- [Uploading storage templates](#)
- [Downloading storage templates](#)
- [Deleting storage templates](#)
- [Locking storage templates](#)
- [Unlocking storage templates](#)

About storage provisioning

Storage provisioning is the process of assigning storage, usually in the form of server disk drive space to optimize the performance of a storage area network (SAN). The storage template is an efficient way to quickly create storage based on the configurations that are mentioned in the template. This mechanism lets you specify storage attributes and conditions, and the template can be used by the storage operators to allocate storage. It ensures that the allocated storage

contains attributes mentioned in the template, and it adheres to the conditions that are specified in the template.

InfoScale Operations Manager storage provisioning lets you create storage templates from existing virtual storage devices such as volumes and file systems. It extracts all the properties of the virtual storage device, and recreates the entire stack. The supported file systems are VxFS and NTFS.

With InfoScale Operations Manager storage provisioning, you determine the virtual device characteristics of the storage object using the file system, volume manager, and operating system utilities and can create a storage template based on this information. A storage template lets you migrate existing virtual devices to new ones with different properties. For example, you can migrate volumes from thick LUNs to thin LUNs.

You can use the Management Server console to create storage templates and provision storage.

You can also use the InfoScale Operations Manager Web services API version 6.1 or later to provision storage using a storage template that you previously created.

For information on using the InfoScale Operations Manager Web services API, refer to the *Arctera Operations Manager Management Server User Guide*.

For information on the InfoScale Operations Manager versions that the add-on is compatible with, refer to the *InfoScale Operations Manager Hardware and Software Compatibility List (HSCL)*.

See [“About creating a storage template”](#) on page 8.

About creating a storage template

You can create a storage template using one of the following methods:

- Modeling it on an existing VxFS file system
See [“Creating a storage template using VxFS file systems”](#) on page 9.
- Modeling it on an existing NTFS file system
See [“Creating a storage template using NTFS file systems”](#) on page 15.
- Modeling it on an existing Storage Foundation volume
See [“Creating a storage template using volumes”](#) on page 16.

Creating a storage template using VxFS file systems

In the Management Server console, you can create a new storage template using VxFS file systems. However, note that the file systems with RAID 5 volumes are not supported.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a storage template using VxFS file systems

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 In the **Actions** pane, click **Create Template**.
- 3 In the **Create Storage Template** wizard, choose **Using a NTFS or VxFS file system**, and click **Next**.
- 4 Select a host, a disk group, and a VxFS file system, and click **Next**.
- 5 In the **Summary** tab, enter a unique storage template name, and optionally an operating system and a description.
- 6 Select the file system object from the left navigation tree and on the **FS Summary** tab, enter a unique name for the file system template, and optionally a description. To view and update the file system template parameters, click the **Parameters** and **Advanced** tabs.

See [“Create Storage Template – File system parameters panel options”](#) on page 10.

See [“Create Storage Template – File system advanced panel options”](#) on page 10.

- 7 Select the volume object from the left navigation tree. Click the **Summary**, **Parameters**, and **Advanced** tabs to view or update the volume name, description and parameters.

See [“Create Storage Template – Volume parameters panel options”](#) on page 13.

See [“Create Storage Template – Volume advanced panel options”](#) on page 14.

- 8 Click **Finish**. The results are displayed. Click **Close**. You can now use the template to provision storage.

See [“Provisioning storage”](#) on page 18.

Create Storage Template – File system parameters panel options

Use this wizard panel to view or update the storage template file system parameters fields when you create or update the storage template.

Table 1-1 Create Template - File system parameters

Field	Description
Size	File system size.
Ask	Lets you change the value while provisioning the storage.
Create file system	Select this check box to create a file system while provisioning the storage.
Mount point prefix	Enter the file system mount point prefix.
Ask	Lets you set the mount point prefix while provisioning the storage.
Block size	File system block size in bytes.
Windows NTFS Parameters	
Volume Label	Enter the volume label for Windows NTFS.
Quick Format	Specify if Quick Format is allowed for the volume.
Compression	Specify if data compression is allowed for the volume.
Allocation Unit Size	Select the required allocation unit size for the volume.

See [“Creating a storage template using VxFS file systems”](#) on page 9.

See [“Creating a storage template using volumes”](#) on page 16.

See [“Updating a storage template”](#) on page 17.

Create Storage Template – File system advanced panel options

Use this wizard panel to view or update the storage template file system advanced fields when you create or update the storage template. These parameters apply to running the `mount` and `vxtunefs` commands.

Note: When you create a storage template based on the VxFS file system, you can set all the parameters. However, some of these parameters are not supported on certain platforms.

Table 1-2 Create Storage Template - Performance Tunables Tab

Field	Description
Intent log size (bytes)	Enter the size of the intent log in the multiples of size byte.
Intent log behavior (mount option)	Select an option to control the timing of flushing the VxFS intent log and other metadata to disk. This affects when operations are guaranteed persistent after a system failure.
I/O feature (mount option)	Select the required mount I/O feature.
Read ahead	Select the required option. This implements an algorithm that allows read ahead to detect more elaborate patterns (such as increasing or decreasing read offsets, or multithreaded file accesses) in addition to simple sequential reads.
Allocation unit size (bytes)	Enter the size of the allocation unit.
Initial extent size (blocks)	Enter a value for the initial extent.
Maximum extent size (bytes)	Enter the maximum size of an extent.
HSM write preallocation	Select the required option. For a file managed by a Hierarchical Storage Management (HSM) application, this option preallocates disk blocks before data is migrated back into the file system.
Maximum disk queue (bytes)	Enter a value for the maximum disk queue that is generated by a single file.
Write throttle	Enter a value to lower the number of dirty pages per file that the file system generates before writing them to disk.

Table 1-3 File Change Log Tunables Tab

Field	Description
FCL keep time	Enter a value for the VxFS File Change Log (FCL) to keep records in the log.
FCL maximum allocated space	Enter the maximum amount of space that can be allocated to the VxFS File Change Log.
FCL open interval	Enter the time interval in seconds within which subsequent opens of a file do not produce an additional FCL record.
FCL write interval	Enter the time, in seconds, that must elapse before the VxFS File Change Log records a data overwrite, data extending write, or data truncate for a file.

Table 1-4 Other Options Tab

Field	Description
Auto mount	Select an option. This lets the file system to be mounted explicitly.
Preferred read request size	Enter the preferred read request size.
Number of parallel read requests	Enter the number of parallel read requests of size <code>read_pref_io</code> that can be outstanding at one time.
Preferred write request size	Enter the preferred write request size.
Number of parallel write requests	Enter the number of parallel write requests of size <code>write_pref_io</code> that can be outstanding at one time.
Maximum inode aging count	Enter the maximum number of inodes to place on an inode aging list.
Minimum inode size	Enter the minimum size to qualify a deleted inode for inode aging.
Large files enabled	Select an option to make the file system is large files capable.

See [“Creating a storage template using VxFS file systems”](#) on page 9.

See [“Creating a storage template using volumes”](#) on page 16.

See [“Updating a storage template”](#) on page 17.

Create Storage Template – Volume parameters panel options

Use this wizard panel to define or update the volume parameters in a storage template when you create or update the storage template.

Table 1-5 Create Storage Template - Volume parameters options

Field	Description
Layout	Volume layout, which can be Concat , Concat Mirror , Mirror Concat , Stripe , Stripe Mirror , Mirror Stripe .
# Columns	The number of columns (for stripe volumes).
# Mirrors	The number of mirrors (for mirrored volumes).
Name prefix	The volume name prefix.
Minimum number of paths	The minimum number of DMP paths for the disks on which the volume is created.
Fast resync	Turn on/off the Storage Foundation fast resync option.

In the **Rules for selecting LUNs** area, you can view, add, and edit rules for selecting the LUNs. If you have specified mirrored layout options, tabs are displayed corresponding to the number of mirrors that you have specified.

For each mirror, you can define the rules that InfoScale Operations Manager can use to select LUNs for hosting the mirror.

- To add a rule, click the plus (+) icon at the bottom of the rules list.
- To remove a rule, click the minus (-) icon.
- To edit an existing rule, double-click it. The **Select LUN Characteristics** dialog box is displayed.

See [“Select LUN Characteristics options”](#) on page 14.

To create mirror plex on multiple enclosures, select the **Mirror across enclosures** check box.

To create stripe volumes on multiple enclosures, select the **Stripe across enclosures** checkbox.

See [“Creating a storage template using VxFS file systems”](#) on page 9.

See [“Creating a storage template using volumes”](#) on page 16.

See [“Updating a storage template”](#) on page 17.

Select LUN Characteristics options

Use this dialog box to view or edit required LUN characteristics for a storage template.

Table 1-6 Select LUN Characteristics options

Field	Description
Enclosure VDIID	Displays the vdiid of the enclosure.
Vendor	Select an enclosure vendor.
Product	Select an enclosure product name.
LUN Type	Select the type required, for example, Thin , Thick , or Any (the default).
Tier	Select a predefined LUN classification. LUN classifications are defined by the storage administrator on the Storage perspective.
Media	Select the type of the media, for example, SSD (solid-state drive) or HDD (hard disk drive), or Any (the default).
Replicated	Select whether to use replicated LUNs.

See [“Creating a storage template using VxFS file systems”](#) on page 9.

See [“Creating a storage template using volumes”](#) on page 16.

See [“Updating a storage template”](#) on page 17.

Create Storage Template – Volume advanced panel options

Use this wizard panel to view or update the storage template volume advanced fields upon creation or update of the storage template.

Table 1-7 Create Storage Template - Volume advanced options

Field	Description
Mode	Enter a value to set the read or write access on the volume.

Table 1-7 Create Storage Template - Volume advanced options (*continued*)

Field	Description
Exclusive	Select On to specify that only one node in the cluster can open an existing volume at a time. Select Off to specify that more than one node in a cluster can open a volume simultaneously

See [“Creating a storage template using VxFS file systems”](#) on page 9.

See [“Creating a storage template using volumes”](#) on page 16.

See [“Updating a storage template”](#) on page 17.

Creating a storage template using NTFS file systems

In the Management Server console, you can create a new storage template using NTFS file systems.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a storage template using NTFS file systems

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 In the **Actions** pane, click **Create Template**.
- 3 In the **Create Storage Template** wizard, choose **Using a NTFS or VxFS file system**, and click **Next**.
- 4 Select a host, a disk group, and a NTFS file system, and click **Next**.
- 5 In the **Summary** tab, enter a unique storage template name, and optionally an operating system and a description..
- 6 Select the file system object from the left navigation tree. Click the **Summary** and **Parameters** tabs to view or update the file system name, description, and parameters.

See [“Create Storage Template – File system parameters panel options”](#) on page 10.

- 7 Select the volume object from the left navigation tree. Click the **Summary** and **Parameters** tabs to view or update the volume name, description, and parameters.
See [“Create Storage Template – Volume parameters panel options”](#) on page 13.
- 8 Click **Finish**. The results are displayed. Click **Close**. You can now use the template to provision storage.
See [“Provisioning storage”](#) on page 18.

Creating a storage template using volumes

In the Management Server console, you can create a storage template using volumes. Storage templates let you store the volume configuration information in the form of a template so that you can use them on other volumes.

Storage Provisioning and Enclosure Migration Add-on supports only Storage Foundation volumes. RAID5 volumes are not supported.

Note: Only hosts with VxFS volumes are displayed on the host table. Depending on which host is selected, the disk group choices change. Only disk groups with Storage Foundation volumes are displayed on the disk group table. Depending on which host and disk group are selected, the volume choices change.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To create a storage template using volumes

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 In the **Actions** pane, click **Create Template**.
- 3 In the **Create Storage Template** wizard, choose **Using a volume**. Click **Next**.
- 4 Select a host, a disk group, and a volume, and click **Next**.
- 5 In the **Summary** tab, enter a unique storage template name and optionally an operating system and a description. .

- 6 Select the file system object from the left navigation tree. Click the **Summary**, **Parameters**, and **Advanced** tabs to view or update the file system name, description and parameters.

See [“Create Storage Template – File system parameters panel options”](#) on page 10.

See [“Create Storage Template – File system advanced panel options”](#) on page 10.
- 7 Select the volume object from the left navigation tree. Click the **Summary**, **Parameters**, and **Advanced** tabs to view or update the volume name, description and parameters.

See [“Create Storage Template – Volume parameters panel options”](#) on page 13.

See [“Create Storage Template – Volume advanced panel options”](#) on page 14.
- 8 Click **Finish**. The results are displayed. Click **Close**.

Updating a storage template

In the Management Server console, you can update or view a storage template.

To perform an update task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To update a storage template

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 In the table, right-click the storage template and select **Update Template**.
- 3 In **Update Template** wizard, optionally modify the template information in the **Summary** tab. The storage template name is required and it must be unique.
- 4 Expand the storage template in the left navigation to view its associated volume and file system.
- 5 Select the file system from the left navigation tree. You can view or update the file system-related information in the following tabs:
 - **FS Summary** tab
You can update the name and description.
 - **Parameters** tab
See [“Create Storage Template – File system parameters panel options”](#) on page 10.

- **Advanced** tab: This tab is displayed only if you selected the **Create File System** check box in the **Parameters** tab.
See [“Create Storage Template – File system advanced panel options”](#) on page 10.
- 6 Select the volume from the left navigation tree. You can view or update the volume-related information in the following tabs:
 - **Volume Summary** tab
 - **Parameters** tab
See [“Create Storage Template – Volume parameters panel options”](#) on page 13.
 - **Advanced** tab
See [“Create Storage Template – Volume advanced panel options”](#) on page 14.
 - 7 To save the updated storage template information, click **Finish**. The results are displayed. Click **Close**.

Provisioning storage

In the Management Server console, you can provision storage using a storage template.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To provision storage

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 In the **Actions** pane, click **Provision Storage**.
- 3 Select a host and a disk group that the storage should be provisioned on. (When you select a host, the associated disk groups are displayed for your selection.) Click **Next**.
- 4 Select a storage template.

If **Ask** is selected in the storage template for the size or for the mount point, you can edit the corresponding fields.
- 5 Click **Finish** to provision storage. The task information is displayed so that you can track the task in the Tasks pane. Click **Close**.

Uploading storage templates

In the Management Server console, you can upload storage templates from your computer to the InfoScale Operations Manager repository. You can upload one template at a time.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To upload a storage template

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 In the **Actions** pane, click **Upload Template**.
- 3 In **Upload Templates** panel, click **Load template file** to navigate to the template that you want to upload and click **Open**.
- 4 Click **Upload** to upload the template to the InfoScale Operations Manager repository.

See [“Downloading storage templates”](#) on page 19.

Downloading storage templates

In the Management Server console, you can download storage templates. You can select multiple storage templates and download them simultaneously. When you download the storage templates, they are saved on your computer in a single file. The downloaded file has a `.json` (JavaScript Object Notation) extension. For example, if you download two storage templates, the application creates a combined downloaded file (`2tpls.json`) that you need to save on your computer.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To download one or more storage templates

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 In the **Actions** pane, click **Download Template**.
- 3 In the **Download Templates** panel, select the templates to download and click **Download**.

See [“Uploading storage templates”](#) on page 19.

Deleting storage templates

In the Management Server console, you can delete storage templates that you created.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To delete one or more storage templates

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 Right-click the storage template and select **Delete Template**.
- 3 In the **Delete Storage Template** panel, confirm the operation. You can delete only those storage templates that are editable, and that you have created.

Locking storage templates

In the Management Server console, you can lock storage templates that you have created.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To lock one or more storage templates

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 Right-click the storage template and select **Lock Storage Template**.
- 3 In the **Lock Storage Template** panel, click **Yes** to lock the storage template. You can lock only those storage templates that you have created.

See [“Unlocking storage templates”](#) on page 20.

Unlocking storage templates

In the Management Server console, you can unlock storage templates that you have created.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To unlock one or more storage templates

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Provisioning**.
- 2 Right-click the storage template and select **Unlock Storage Template**.
- 3 In the **Unlock Storage Template** panel, click **Yes** to unlock the storage template. You can unlock only those storage templates that you have created.

See [“Locking storage templates ”](#) on page 20.

Migrating volumes

This chapter includes the following topics:

- [About volume migration](#)
- [About the Impact Analysis report for volume migration](#)
- [Migrating volumes by host](#)
- [Select LUN Characteristics options](#)
- [Change layout options](#)
- [Migrating volumes by enclosure](#)
- [Migrating volumes by disk group](#)
- [Pausing or resuming a volume migration](#)
- [Deleting a scheduled volume migration](#)

About volume migration

Enterprises frequently need to move (migrate) storage volumes to new disk enclosures. Volume migration can be time-consuming and complex. Storage Provisioning and Enclosure Migration Add-on simplifies the complexities so that you can easily migrate volumes between disk enclosures.

Using the Management Server console, you can migrate volumes in the following ways.

See [“Migrating volumes by host”](#) on page 24.

See [“Migrating volumes by enclosure”](#) on page 27.

See [“Migrating volumes by disk group”](#) on page 29.

Note: Volume migration is supported for managed hosts that run InfoScale Operations Manager (VRTSsfmh package) 5.x and later.

Before migrating volumes from one enclosure to the target disks in another enclosure, consider the following factors:

- The volumes to be migrated
Storage Provisioning and Enclosure Migration Add-on migrates only the Storage Foundation-managed volumes.
A storage array might contain volumes that are managed by a different volume manager than Storage Foundation. It might also contain disks directly used by the hosts. The add-on cannot migrate such volumes and will ignore them.
- The target enclosure
Ensure that the target enclosure is connected to the host to which the source volumes are connected. Both the disks that contain the volumes to be migrated and the target disks must be accessible from the same host.
- The target disks where you want to migrate the volumes
The size and layout of selected volumes determine the number of target disks you need. Ensure that the size and number of the target disks are adequate to migrate the selected volumes. Ensure that you add the disks to the disk group. If you use the disk group migration wizard, you can add the disks using the migration wizard.

Volume migration stops if the host shuts down while migrating the selected volumes to the specified target disks. However, it does not result in loss of data. You can start a new volume migration case to migrate the remaining volumes.

InfoScale Operations Manager lets the storage administrator define LUN classifications. LUN classifications can specify the LUN characteristics, such as RAID levels and replication, that match the SLAs for the applications that are using the volumes. If LUN classifications are defined, you can use them to determine the target LUNs when you set up storage migration.

About the Impact Analysis report for volume migration

When you migrate volumes using Storage Provisioning and Enclosure Migration Add-on, the migration wizard can generate an Impact Analysis report, in addition to the migration summary displayed in the wizard. The impact analysis shows additional information about the objects that are associated with the volumes being migrated. You can analyze the potential outcome of the volume migration

and if necessary change the options you have selected before executing the volume migration.

You can view information such as the following.

- The details about the databases that may be affected by the volume migration.
- The details about Business Applications that may be affected by the volume migration.

See [“About volume migration”](#) on page 22.

Migrating volumes by host

You can migrate volumes that reside on an enclosure to one or more other enclosures. This topic covers migration by host. You can also migrate by enclosure or disk group. Disk group migration is available only on the UNIX/Linux hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Before you migrate volumes, review the planning information in the following:

See [“About volume migration”](#) on page 22.

In addition, verify the following:

- If volumes have been set up for creating snapshots, unprepare those volumes.
- Ensure that disks are provisioned on the new enclosures and are added to the disk group.
- Ensure that the enclosures are added to InfoScale Operations Manager.

To migrate volumes by host

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
- 2 In the **Actions** pane, click **Migrate Volumes By Host**.
- 3 In the **Migrate Volumes By Host** wizard, specify a name and description for the migration task so that you can track the task status. Click **Next**.
- 4 Select the host and then select the source enclosure. Click **Next**.

- 5 Select one or more enclosures as the targets to which you want to migrate the volumes.

Under **Select LUN Characteristics**, specify the LUN requirements.

See [“Select LUN Characteristics options”](#) on page 26.

Click **Next**.

- 6 Select the volumes to migrate. Verify the source and target layout. If you want to change the target layout, double-click the volume and select the new layout on the **Change Layout** window.

See [“Change layout options”](#) on page 26.

You can also choose from the following options:

- Keep the same number of paths to the volume (the default) or set a minimum number of paths.
- Optionally, select the check box to delete the original volumes after successfully migrating the volumes to the target enclosure.

Click **Next**.

- 7 Choose one of the following:

- Schedule the time of the migration.
- To run immediately, select **Run now**.

Warning: Review the information under **Migration Summary** before starting the migration. For more information, click **Impact Analysis**. A separate window opens to display impact analysis information regarding affected objects.

- 8 When done reviewing the **Migration Summary** and **Impact Analysis**, click **Previous** if you want to go back and change anything. Otherwise, click **Finish** and close the confirmation window.

Scheduled migrations are listed on the **Migration Status (By Schedule)** tab. You can check on the status of a specific volume on the **Migration Status (By Volume)** tab.

See [“Migrating volumes by enclosure”](#) on page 27.

See [“Migrating volumes by disk group”](#) on page 29.

Select LUN Characteristics options

When migrating volumes, you can specify required LUN characteristics for the target enclosure.

Table 2-1 Select LUN characteristics options

Field	Description
Enclosure VDIID	Displays the vdiid of the enclosure.
Vendor	Select an enclosure vendor. The default is the selected enclosure vendor.
Product	Select an enclosure product name. The default is the selected enclosure product name.
LUN Type	Select the type required, for example, Thin , Thick , or Any (the default).
Classification	Optionally, select a predefined LUN classification. LUN classifications are defined by the storage administrator on the Storage perspective.
Media	Select the type of the media, for example, SSD (solid-state drive) or HDD (hard disk drive), or Any (the default).
Raid Type	Select the RAID type, for example, RAID0.
Replicated	Select whether to use replicated LUNs.

See [“Migrating volumes by host”](#) on page 24.

See [“Migrating volumes by enclosure”](#) on page 27.

Change layout options

Use this panel to change the volume layout of the target volume when selecting volumes for migration.

Table 2-2 Change volume layout options

Field	Description
Target Layout	Select one of the following: ■ concat - A layout style that is characterized by the subdisks that are arranged sequentially and contiguously. ■ concat-mirror - A combination of concat and mirror, where mirror represents extra copies of the data for data protection. ■ mirror -concat - A combination of mirror and concat. ■ striped - A layout technique that spreads data across several physical disks using stripes. The data is allocated alternately to the stripes within the subdisks of each plex. ■ stripe-mirror - A combination of stripe and mirror.
No. of mirrors	Select the number of mirrors, for mirrored layouts. The default is 2.
Stripe width	For striped layouts, select the stripe width. The default is 128.
No. of columns	For striped layouts, select the number of columns. The default is 2.

See [“Migrating volumes by host”](#) on page 24.

See [“Migrating volumes by enclosure”](#) on page 27.

See [“Migrating volumes by disk group”](#) on page 29.

Migrating volumes by enclosure

You can migrate volumes that reside on an enclosure to one or more other enclosures. This topic covers migration by enclosure. You can also migrate volumes by host and by disk group. Disk group migration is available only on the UNIX/Linux hosts.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Before you migrate volumes, review the planning information in the following:

See [“About volume migration”](#) on page 22.

In addition, verify the following:

- If volumes have been set up for creating snapshots, unprepare those volumes.
- Ensure that disks are provisioned on the new enclosures.
- Ensure that the enclosures are added to InfoScale Operations Manager.

To migrate volumes by enclosure

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
- 2 In the **Actions** pane, click **Migrate Volumes By Enclosure**.
- 3 In the **Migrate Volumes By Enclosure** wizard, specify a name and description for the migration task so that you can track the task status. Click **Next**.
- 4 Select the source enclosure and click **Next**.
- 5 Select one or more enclosures as the targets to which you want to migrate the volumes.

Under **Select LUN Characteristics**, specify the LUN requirements and click **Next**.

See [“Select LUN Characteristics options”](#) on page 26.

- 6 Select a host and click **Next**.
- 7 Select the volumes to migrate. Verify the source and target layout. If you want to change the target layout, double-click the volume and select the new layout.

See [“Change layout options”](#) on page 26.

You can also choose from the following options:

- Keep the same number of paths to the volume (the default) or set a minimum number of paths.
- Optionally, select the check box to delete the original volumes after successfully migrating the volumes to the target enclosure.

Click **Next**.

- 8 Choose one of the following:
 - Schedule the time of the migration.
 - To run immediately, select **Run now**.

Warning: Review the information under **Migration Summary** before starting the migration. For more information, click **Impact Analysis**. A separate window opens to display impact analysis information regarding affected objects.

- 9 When done reviewing the **Migration Summary** and **Impact Analysis**, click **Previous** if you want to go back and change anything. Otherwise, click **Finish** and close the confirmation window.

Scheduled migrations are listed on the **Migration Status (By Schedule)** tab. You can check on the status of a specific volume on the **Migration Status (By Volume)** tab.

See [“Migrating volumes by host”](#) on page 24.

See [“Migrating volumes by disk group”](#) on page 29.

Migrating volumes by disk group

You can migrate a disk group that resides on an enclosure to one or more other enclosures. While setting up the migration you can optionally add free, uninitialized disks from the target enclosures to the disk group. This disk group operation is available only on the UNIX/Linux hosts and requires Storage Provisioning and Enclosure Migration Add-on.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

Before you migrate volumes by disk group, review the planning information in the following:

See [“About volume migration”](#) on page 22.

In addition, verify the following:

- If volumes have been set up for creating snapshots, unprepare those volumes.
- Ensure that disks are provisioned on the new enclosures. You can use the migration wizard to add disks to the disk group.
- Ensure that the enclosures are added to InfoScale Operations Manager.

To migrate volumes by disk group

- 1 In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
- 2 Expand the Organization or **Uncategorized Hosts** to locate the host.

- 3 Expand the host and expand **Disk Groups** to locate the disk group.
- 4 Right-click the disk group and click **Migrate Storage**.
- 5 In the **Migrate Volumes by Disk Group** wizard, specify a name and description for the migration task so that you can track the task status. Click **Next**.
- 6 Select the source and target enclosure. Optionally, to add disks to the disk group, select the check box **Add free/uninitialized disks to the disk group**. Click **Next**.
- 7 If you selected the option to add disks, select one or more disks to add to the disk group from the list of available free (initialized) or uninitialized disks. Click **Next**.
- 8 Select the volumes to migrate. Verify the source and target layout. If you want to change the target layout, double-click the volume and select the new layout on the **Change Layout** window.

See ["Change layout options"](#) on page 26.

You can also choose from the following options:

- Keep the same number of paths to disks for selected volumes (the default) or set a minimum number of paths.
- Optionally, select the check box to delete the original volumes after successfully moving the volumes to the target enclosure.

Click **Next**.

- 9 Choose one of the following:
 - Schedule the time of the migration.
 - To run immediately, select **Run now**.

Warning: Review the information under **Migration Summary** before starting the migration. For more information, click **Impact Analysis**. A separate window opens to display impact analysis information regarding affected objects.

Click **Next**.

- 10 When done reviewing the **Migration Summary** and **Impact Analysis**, click **Previous** if you want to go back and change anything. Otherwise, click **Finish** and close the confirmation window.

Scheduled migrations are listed in the **Solutions** area for the **Storage Migration** solution on the **Migration Status (By Schedule)** tab. You can check on the status of a specific volume on the **Migration Status (By Volume)** tab.

See [“Migrating volumes by host”](#) on page 24.

See [“Migrating volumes by enclosure”](#) on page 27.

Pausing or resuming a volume migration

If a volume migration task is not complete, you can pause the migration and later resume it. The Pause/Resume operation is only available on a UNIX/Linux managed host with InfoScale Operations Manager version 6.0 or later.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To pause a volume migration

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
- 2 In the **Migration Status (by Schedule)** tab, right-click the scheduled migration task and click **Pause**.
- 3 In the **Pause Migration** wizard, select the volumes to pause and click **Finish**.

To resume a volume migration

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
- 2 In the **Migration Status (by Schedule)** tab, right-click the paused migration task and click **Resume**.
- 3 In the **Resume Migration** wizard, select the volumes to resume and click **Finish**.

Deleting a scheduled volume migration

If a scheduled volume migration is not in progress, you can delete it.

To perform this task, your user group must be assigned the Admin role on the host or the Server perspective. The permission on the host may be explicitly assigned or inherited from a parent Organization.

To delete a scheduled volume migration

- 1 In the Management Server console, go to the **Server** perspective, expand **Solutions** in the left pane, and click **Storage Migration**.
- 2 In the **Migration Status (by Schedule)** tab, right-click the scheduled migration task and click **Delete Schedule**.

Application Migration Add-on

- [Chapter 3. Introduction to Application Migration Add-on](#)
- [Chapter 4. Creating and managing an application migration plan](#)

Introduction to Application Migration Add-on

This chapter includes the following topics:

- [About Application Migration add-on](#)
- [Before installing the Application Migration add-on](#)
- [Installing the Application Migration Add-on](#)
- [Before uninstalling the Application Migration add-on](#)
- [Uninstalling the Application Migration add-on](#)

About Application Migration add-on

The Application Migration add-on allows you to migrate applications that are under Cluster Server (VCS) management from one cluster to another. The application migration operation is less complex and can be accomplished with minimal manual intervention. Application migration can be across operating systems, architectures, or virtualization technologies. You can migrate data using:

- VxVM Mirroring
- VVR Replication

You can migrate an application between different:

- Platforms—AIX, Linux, and Solaris
- Environments—Physical-to-physical, physical-to-virtual, virtual-to-virtual, and virtual-to-physical, on-premise to AWS cloud
- InfoScale versions

From Application Migration add-on version 7.3.1 onwards, you can migrate your applications from an on-premise cluster to a cluster on AWS EC2 instance. You need to select data-migration type as VVR Replication in order to migrate your VCS workload to AWS. For more information about VVR and to understand the basic concepts, refer to the *Arctera InfoScale Replication Administrator's Guide*.

See [“Supported versions and platforms”](#) on page 38.

To migrate an application, you must create an application migration plan using the **Create Migration Plan** wizard. After you create a plan, you must execute the migration plan.

See [“Creating an application migration plan”](#) on page 44.

See [“Executing the application migration plan”](#) on page 68.

The add-on also allows you to:

- Pause and resume the operation for manual verification and correction, if required.
- Integrate custom scripts in the operation as per application requirements.
- Migrate application dependencies.
- Understand source cluster configuration and create target cluster configuration.
- Perform endian changes to the data as per architecture requirements.
- Rehearse the steps before the actual migration operation.

Before installing the Application Migration add-on

To install the Application Migration add-on, you must be familiar with Arctera InfoScale Operations Manager concepts and documentation.

Install Arctera InfoScale Operations Manager Management Server, if it is not already installed in your environment.

For more information, see the *Arctera InfoScale Operations Manager Installation and Configuration Guide*.

Installing the Application Migration Add-on

You can either install only the Application Migration add-on or you can perform a full installation (install Arctera InfoScale Operations Manager Management Server and add-ons).

You can use the Management Server to install the Application Migration add-on. Download the add-on from the Services Operations Readiness Tools (SORT)

website. In the Management Server console, upload the add-on to the repository using the **Upload Solutions** option, and then install.

For the detailed instructions on downloading, uploading, and installing the Arctera InfoScale Operations Manager add-ons, refer to the *Arctera InfoScale Operations Manager Installation and Configuration Guide*.

Note: After you install the add-on, you must restart the Arctera InfoScale Operations Manager Web Server.

Before uninstalling the Application Migration add-on

Before you uninstall the Application Migration add-on, ensure that you have deleted all the existing application migration plans. If an application migration plan exists, the add-on uninstallation fails.

Before deleting the application migration plan, you can take a backup of the plan by exporting it. If you reinstall the add-on, you can import this plan and use it.

- See [“Exporting application migration plan\(s\)”](#) on page 72.
- See [“Importing application migration plan\(s\)”](#) on page 73.

Uninstalling the Application Migration add-on

You can uninstall the Application Migration add-on from the Management Server. To perform this task, your user group must be assigned the Admin role on the Management Server perspective.

For the detailed instructions on uninstalling Arctera InfoScale Operations Manager add-ons, refer to the *Arctera InfoScale Operations Manager Management Server Installation and Configuration Guide*.

Creating and managing an application migration plan

This chapter includes the following topics:

- Supported versions and platforms
- User privileges
- Prerequisites for creating an application migration plan
- Prerequisites for migration to AWS
- VVR Replication: Environment variables used in application migration
- Creating an application migration plan
- Understanding user-defined tasks
- Understanding application migration operations
- Understanding the cleanup operation
- Understanding the tasks executed in each operation
- Validations performed before migration plan execution
- Executing the application migration plan
- Editing an application migration plan
- Deleting application migration plan(s)
- Exporting application migration plan(s)
- Importing application migration plan(s)

- [Viewing historical runs](#)
- [Viewing properties of an application migration plan](#)
- [Application migration logs](#)

Supported versions and platforms

The following platforms are supported:

- AIX
- Linux
- Solaris

[Table 4-1](#) lists the supported product or component and versions:

Table 4-1 Supported version(s)

Product or component	Supported version(s)
InfoScale Operations Manager Management Server	8.0.2 and later
InfoScale Operations Manager managed host	8.0 and later
SFHA SFCFSHA	8.0 SP1 RP4 to 8.0.2
InfoScale Availability InfoScale Storage InfoScale Enterprise	8.0 and later
Cluster Server version is from 8.0 and later	Managed host version of the cluster must be 8.0.2 or later It is recommended to have the same version as of the Management Server. Note: You must uninstall the older VRTSsfmh package and install the newer VRTSsfmh package. If you upgrade the package, the required functionality may not be available.

Table 4-1 Supported version(s) (*continued*)

Product or component	Supported version(s)
Cluster Server version is from 8.0 and later	Managed host version of the cluster can be the same version as of the cluster. It is recommended to upgrade to the same version as of the Management Server.

Licenses required for the managed hosts

Mirroring

- If your product version is lower than 7.0, you must have one of the following licenses:
Storage Foundation Enterprise Edition with Cluster Server
Or
Storage Foundation Enterprise Edition and Cluster Server
- If your product version is 7.0 or higher, you must have one of the following licenses:
InfoScale Enterprise
Or
InfoScale Storage and InfoScale Availability

Replication

- If your product version is lower than 7.0, you must have one of the following licenses:
Storage Foundation Enterprise Edition with Cluster Server and VVR
Or
Storage Foundation Enterprise Edition with VVR + Cluster Server
- If your product version is 7.0 or higher, you must have one of the following licenses:
InfoScale Enterprise
Or
InfoScale Storage and InfoScale Availability

User privileges

Table 4-2 lists the user privilege required in the Availability and Server perspectives for performing the various tasks.

Table 4-2 Tasks and user privileges

Task	Availability perspective: Admin	Availability perspective: Admin	Availability perspective: Guest	Availability perspective: Guest
	Server perspective: Admin	Server perspective: Guest	Server perspective: Admin	Server perspective: Guest
Setup Storage	Yes	No	No	No
Rehearse	Yes	No	No	No
Migrate	Yes	No	No	No
Create	Yes	No	No	No
Edit	Yes	No	No	No
Delete	Yes	Yes	No	No
Import	Yes	Yes	No	No
Export	Yes	Yes	No	No
Historical Runs	Yes	Yes	Yes	Yes
Properties	Yes	Yes	Yes	Yes

Prerequisites for creating an application migration plan

In this document, source cluster refers to the cluster from which the service group managing the application must be migrated and target cluster is the cluster to which the service group managing the application must be migrated to.

Ensure that the following prerequisites are met before creating an application migration plan:

- Add all the nodes of the source and target cluster to Arctera InfoScale Operations Manager (IOM) as managed hosts.
- Application data must reside on a Volume Manager (VxVM) disk group.

The add-on migrates application data only on VxVM volumes (DiskGroup or CVMVolDG resources). The add-on does not migrate application data on native disks or other volume managers.

- Source cluster must be in running state.
- VCS and VxVM versions in the target cluster must be same or higher than the version in the source cluster.
- All DiskGroup or CVMVolDG type resources in the service groups being migrated must be online.
- Install and configure VCS cluster on the target managed hosts. The target cluster must also be in running state.
- If Cluster Volume Manager (CVM) is configured in the source cluster, install and configure CVM in the target cluster also.
- In the target cluster configuration file, include type definition files for the resource types that are being migrated.
- On the source cluster, all service groups must be online.
- On the managed hosts of the source and target clusters, the XPRTLD daemon must be always running.
- On the target cluster, install and configure the application that is being migrated according to the application and operating system (OS) requirements.
- Ensure that the application has proper service group dependencies because application migration from source to target cluster happens by migrating the service groups in which the application is configured. Dependencies for the selected service group are automatically included as a part of the migration plan.
- Execute the `vxlist (/opt/VRTSsfmh/bin/vxlist)` utility in the source and target cluster node. The output must not contain any errors.
- On the target cluster, identify the IPs that the application will use.
- VVR Replication—On the source and target cluster nodes, identify IPs to create replication links. Ensure that these IPs are not used.
- Depending on the Data Migration type you select, the migration plan detects any pre-existing linked break-off snapshot or VVR replication configured for each disk group between the nodes in the source and target clusters. If no pre-existing configuration is detected, the migration plan creates it and you must identify the list of disks to be used to create either the mirror disk group or remote site disk group. The disks must be under VxVM control and must have enough free space. The total size of the disks you provide must be more than the used size of the disk group in the source cluster. The number of disks

must be at least same or more considering the volume layout in the source cluster.

- VxVM Mirroring—Disks must be shared across all the nodes in the source and target cluster.
- VVR Replication—Disks must be shared between all nodes in the target cluster only.

Prerequisites for migration to AWS

To migrate an application from on-premise to AWS using Application Migration add-on:

- Bi-directional connectivity should be present between AWS EC2 instance and on-premise cluster. This can be achieved by establishing a tunnel between on-premise network and AWS VPC in which the EC2 instances are present.
- Necessary ports should be opened for adding AWS EC2 instances to InfoScale Operations Manager residing on-premise and for establishing VVR between on-premise cluster nodes and AWS EC2 instances.

To view the ports required to be opened for VVR , refer to the section *Port numbers used by VVR in InfoScale Replication Administrator's Guide*.

To view the ports required to be opened for Managed Host communication to Management Server, refer to the section *Network and firewall requirements in InfoScale Operations Manager Installation and Configuration Guide*.

- Once communication is verified, add the AWS EC2 instance as managed host to InfoScale Operations Manager using Private IP of the instances.

VVR Replication: Environment variables used in application migration

In the source and target cluster nodes, configure every disk group that must be migrated with a Storage Replicator Log (SRL) volume. During Rehearse operation, create cache volume in every disk group on the secondary site so that space-optimized snapshots can be created. The application migration plan creates the required SRL volume and cache volume. For creating these volumes, 10% of the used size of the disk group is taken as the default size.

To override the default size of both SRL volume and cache volume, mention the size required for a disk group using specific variables in the `appmig.conf` file. The file is not present by default. On the node where the disk group is online, create

the file in the `/opt/VRTSsfmh/etc` directory. In the file, define the variable for each disk group.

For example, for a set of service groups to be migrated, if one disk group is online on one node and another disk group is online on another node, create the file for each node and define the variable in the file where the corresponding disk groups are online.

Arctera recommends keeping one copy of the environment file with all necessary entries and copy the same file across all nodes in the source and target clusters.

Note: Mention the size in megabytes.

Define the following environment variables in the file:

- `SRLVOL_SIZE_DiskGroup`
- `CACHEVOL_SIZE_DiskGroup`

Example

To specify 1 GB size for SRL volume and 2 GB size for cache volume for disk group dg1, define the variables as follows:

- `SRLVOL_SIZE_dg1 = 1024M`
- `CACHEVOL_SIZE_dg1 = 2048M`

Scenarios in which the environment file (if created) is read

- During plan creation from the source cluster node the file checks whether:
 - Sufficient space is available in the source disk group for creating SRL volume.
 - Sum of the size of disks selected for creating disk group in the target cluster is greater than the sum of the used space in the source disk group, plus SRL volume and cache volume size calculated by the plan.
- During the Setup Storage operation, the SRL volume size is retrieved from the file in the source cluster node and the size is used to create the SRL volume for the disk group in the source and target cluster nodes.
- During the Rehearse operation, the cache volume size is retrieved from the file in the target cluster node and the size is used to create the cache volume for the disk group in the target cluster node.

Creating an application migration plan

Table 4-3 describes the fields in the Create Migration Plan wizard.

Table 4-3 Create Migration Plan wizard options

Field	Description
Components Selection	
Name	Enter a name for the migration plan. The maximum length of the name should be 128 characters. The following characters are allowed: ■ A to Z ■ a to z ■ 0 to 9 ■ -
Description	Enter a description or any additional information for the migration plan.
Select data migration type	Select one of the following: ■ VxVM Mirroring ■ VVR Replication
Select a service group to migrate using this plan	Click Browse and select the service group that you plan to migrate using this plan. Arctera recommends that the service group you select for migration must be online in the source cluster
Select a target cluster to migrate to using this plan	Click Browse and select the target cluster to which you plan to migrate to using this plan.
Cluster Information & Dependency	
Cluster Information	Displays the source and target cluster details. The following details are displayed: ■ OS version ■ VCS version ■ VxVM version Note: VCS and VxVM versions in the target cluster must not be lower than the version in the source cluster.

Table 4-3 Create Migration Plan wizard options (*continued*)

Field	Description
Use existing service group(s) in the target cluster	Option is enabled and selected if the target cluster has service groups of the same name that are part of the entire dependency tree for migration from source.
Create new service group(s) in the target cluster	Option is enabled and selected if there are no service groups or resource names in the target cluster that matches the name of the service group or the resources present in the entire service group dependency tree selected for migration. In this scenario, the application migration plan creates all the service groups selected as a part of the migration. For Cluster Volume Manager (CVM) migrations, this option is enabled even if the CVM group is present in the target cluster as well as in the service group dependency tree which is to be migrated. Arctera recommends that you must not deselect the CVM service group during migration. If you opt to deselect the CVM service group, dependency with CVM service group is not created in the target cluster and may affect while bringing the service group online with correct dependency.
Dependent Service Groups for Migration	You can deselect any dependent service group(s) that need not be migrated. However, Arctera does not recommend deselecting any service group that is in the dependency tree. This may affect the application when it is brought online in the target cluster.
Target Disk Information—Panel appears if you select VxVM Mirroring	
Service Group to be Migrated	Displays the name of the service group which has a CVMVolDg or DiskGroup type resource configured in it.
Disk Group	Displays the name of the disk group. Only disk groups configured with CVMVolDg or DiskGroup resources are displayed. If a disk group is not configured under any resource, that disk group is not displayed.

Table 4-3 Create Migration Plan wizard options (*continued*)

Field	Description
Action	■ Use Existing—Option is enabled only if there are existing linked break-off snapshots for any source disk groups part of the migration. If snapshot is found for any disk group, the corresponding disk group is shown as Mirrored. Select this option if you want to use the mirror disk group pre-configured for migration. ■ Create New—Option allows you to create a new mirror disk group for a disk group. Option is enabled by default if there are no pre-configured mirror disk groups for a disk group. In this scenario, the corresponding disk group is marked as Not Mirrored. ■ Ignore—You can ignore a particular disk group resource from being migrated. Arctera does not recommend to ignore any disk group as the data present in these disk groups will not be migrated.
Sub-Action	Select Disks—If you opt to create a new disk group, click Select Disks to select the set of disks for creating the mirror disk group. The Select Disk(s) window displays only the set of disks in the source cluster that are shared between all the target cluster managed hosts. You can select multiple disks so that the combined space of the selected disks is greater than the source disk group space; this is validated when you click Next.
RVG Information—Panel appears if you select VVR Replication	
Service Group to be migrated	Displays the name of the service group which has a CVMVolDg or DiskGroup type resource configured in it.
Disk Group	Displays the name of the disk group. Only disk groups configured with CVMVolDg or DiskGroup resources are displayed. If a disk group is not configured under any resource, that disk group is not displayed.

Table 4-3 Create Migration Plan wizard options (*continued*)

Field	Description
Action	■ Use Existing—Option is auto-selected if the disk group has existing RVGs and for all RVGs secondary sites are added. The secondary site for all RVGs should be one of the nodes of the selected target cluster. If a disk group is marked as Use Existing, the workflow does not take any action on these disk groups and you must establish replication for these disk groups to ensure data sync. This is validated during Rehearse and Migrate operations. ■ Create New—If a disk group is marked as Create New, the workflow creates the necessary entities on the source and target cluster nodes for replication to be established for that disk group. Option is auto-selected for the disk group when one of the conditions is met: ■ Has no RVG created. ■ Has RVGs created but no secondary added. ■ Has RVGs created and few or all RVGs has secondary sites added. None of the secondary site of any RVG should be a node which is part of the target cluster. ■ Ignore—If a disk group is marked as Ignore, the disk group is not included for migration as part of the workflow. Option is auto-selected for the disk group when one of the conditions is met: ■ If RVG does not exist in the disk group and the disk group does not have sufficient space to create SRL volume. ■ Has RVGs created and few or all RVGs have a secondary site added, out of which few RVGs have secondary site as one of the nodes of the selected target cluster and few RVGs have secondary site as a node which is not part of the selected target cluster. Configuration issues mentioned above which caused the disk group to be marked as Ignore can be resolved as follows: ■ For the first scenario, either increase the size of the source disk group or define the required SRL volume size in the environment file according to the free space available in the disk group. ■ For the second scenario, maintain consistent configuration across all RVGs in a disk group. After resolving the configuration issues, click Previous to go back to the previous page and then click Next to revalidate. Note: You can also ignore a disk group as part of the migration process, but Arctera does not recommend this.

Table 4-3 Create Migration Plan wizard options (*continued*)

Field	Description
Sub-Action	Select Disks—If a disk group is marked as Create New, disk information must be provided. Click Select Disks for each disk group to provide the target disks on which the disk group for replication needs to be created. Note: Combined space of disks selected should be greater than the sum of space used in the source disk group, calculated SRL volume size, and cache volume size. You must not select the same disk as part of two disk groups.
RLink Details	RLink Details—If a disk group is marked as Create New, RLink details must be provided. Click RLink Details to provide RLink information for each disk group: ■ NIC for each node of source and target cluster where the IP for RLink needs to be plumbed. ■ IP which is to be used for the disk group on source cluster and target cluster for setting up RLinks for replication. Note: The NIC selected for a disk group for all nodes of source and target cluster should have same version of pre-existing IP plumbed on it (either IPv4 or IPv6). IP which is provided for a source cluster and target cluster should be unique for disk groups across service groups. Same IP can be given for disk groups which are part of the same service group.
Network Configuration	
Service Group Name	Displays the name of the service group under which network resources are configured.
Resource Name	Displays the name of IP-related resources.
Network Parameters	The following are the network parameters: ■ Device ■ IP ■ Netmask
Source Details	Displays the network details of the source.

Table 4-3 Create Migration Plan wizard options (*continued*)

Field	Description
Target Details	Displays the network details of the target. You can edit the target network details. The target details must be in the same format as the source. As an example, for IPMultiNIC resource, if NIC field displays <code>net1, 1, net2, 2</code> in source, the target information that you provide must be in the format <code>eth0,1,eth1,2</code> .
Task Customization	
Tasks	Displays the list of all tasks that will be executed as a part of the following application migration operations: ■ Setup Storage ■ Rehearse ■ Migrate
Add Pause	Allows you to add a pause after any predefined task.
Add Script	Allows you to add a task for custom script execution on a specific cluster node.
Delete	Allows you to delete a pause task or custom script execution task that you added. You cannot delete a predefined task.
Summary	Displays a summary of the application migration plan.

While creating the application migration plan, a query is executed on the source and target cluster nodes and this could cause a delay when moving from the Cluster Information & Dependency panel to the Target Disk Information panel. The query fetches the following details:

- VxVM Mirroring—Existing mirroring configuration
- VVR Replication:
 - Existing RVG configuration
 - Space validation on source disk group
 - Space required for target disk group
 - Network details like NIC, IP, and netmask of all the source and target cluster nodes

To create an application migration plan

- 1 In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane and click **Application Migration Plan**.

- 2 In the Migration Plan pane, right-click and click **Create**.

The Create Migration Plan wizard opens.

- 3 In the Components Selection panel, specify a name and description, select the data migration type, and select a service group and target cluster.

Click **Next**.

- 4 In the Cluster Information & Dependency panel, you can view the cluster details and modify the dependent service groups, if required.

Use existing service group(s) in the target cluster and **Create new service group(s) in the target cluster** options are disabled and you will not be able to proceed if there are few service groups or resources in the target cluster which are present in the service group dependency tree that is being migrated. In this scenario, you have to cancel the operation, do one of the following, and then create the plan again:

- Delete the conflicting service groups/resources on the target—This enables the **Create new service group(s) in the target cluster** option.

Or

- Create the cluster configuration in the target cluster manually—This enables the **Use existing service group(s) in the target cluster** option.

Click **Next**.

If you selected VxVM Mirroring as the Data Migration Type, go to [5](#).

If you selected VVR Replication as the Data Migration Type, go to [6](#).

- 5 In the Target Disk Information panel, you can view the disk groups and select disks for creating mirror disk groups, if required.

If any disk group which is mirrored with the source disk group is configured under DiskGroup or CVMVolDg resource, this disk group is also displayed and is marked **Not Mirrored**. However, the source disk group is marked as **Mirrored**. In this scenario, you can ignore the mirror disk group and choose **Use Existing** or **Create New** for the source disk group.

If the disk group is mirrored and the disks used to create the mirror disk group are not present in all nodes of the target cluster, the disk group is marked as **Mirrored**. However, the **Use Existing** option will not be enabled.

If the disk group resources are not online in the source cluster or the XPRTLD daemon is not running in the source cluster nodes, the wizard will not be able to detect the state of the disk groups. In this scenario, the disk groups will not be marked as Mirrored or Not Mirrored.

If there are no disk groups for migration or if you select **Ignore** for all the disk groups listed, an error is displayed and you will not be able to proceed.

Click **Next**. Go to [7](#).

- 6 In the RVG Information Panel, you can view the disk groups, select disks, and enter RLINK details for creating replication configuration.

Disk groups are marked as **Replicated** if RVGs are configured in it; it does not depend on whether a secondary site is added or not to these RVGs. Disk groups are marked as **Not Replicated** if the disk groups does not have any RVGs configured.

If there are no disk groups for migration or if you select **Ignore** for all the disk groups listed, an error is displayed and you will not be able to proceed.

Click **Next**.

- 7 The Network Configuration panel displays details of network resources configured in any service groups being migrated. You can edit the target network details of these resources.

Note: In the Rehearse operation, IP-related resources are converted to FileOnOff type resources as part of cluster configuration in the target so that resource dependencies are maintained and to ensure that there will be no network connectivity to the applications which are brought online in the target, but NIC related resources are configured.

In Migrate operation, both IP and NIC related resources are configured in the target cluster.

For VVR Replication, IP resources configured to monitor IPs used for any existing replication configuration are not displayed.

Click **Next**.

- 8 In the Task Customization panel, you can view the list of all tasks that will be executed as a part of each operation. You cannot modify or delete a predefined task. After any predefined task in an operation, you can add a pause task or custom script execution task. These tasks are executed in the order specified during the course of the operation. You can also see whether a task is marked as critical or not. If any critical task fails, the entire operation fails.

See [“Understanding user-defined tasks”](#) on page 52.

Click **Next**.

- 9 The Summary panel provides a summary of the application migration plan. Review the plan and click **Finish** to create the application migration plan.

Click **Close** to close the wizard.

See [“Executing the application migration plan”](#) on page 68.

See [“Understanding application migration operations”](#) on page 53.

See [“Understanding the cleanup operation”](#) on page 57.

Understanding user-defined tasks

Pause task

During the application migration process, the operation can be paused for user intervention. When you create the plan, you can insert a pause task after any predefined task in an operation. When the plan is executed, the operation pauses

for an indefinite period at the point where the pause task was inserted. Click **Resume** to resume the operation.

Predefined pause tasks are added after the tasks listed below so that you can validate that task:

- Rehearse tasks—Cluster configuration, service group online, and service group offline.
- Migrate task—Cluster configuration.

Custom script execution task

You can execute custom scripts on cluster nodes during the application migration process. Scripts can be used to perform or validate application-specific tasks or any other tasks which you may want to perform during the operation. When you create the plan, you can add a custom script execution task after any predefined task in an operation. The script that will be executed should not accept any arguments. Any output from the script is logged in the operation execution log on the host where the script is executed.

You can mark a custom script execution task as critical or not. By default, the task is marked as non-critical. If the task is marked as critical, the return code obtained from the script after execution is used to determine the status of the task. If the return code is:

- Zero: Task succeeded
- Non-zero: Task failed

Understanding application migration operations

The following are the three main operations executed during application migration:

- Setup storage—See [“Understanding the Setup Storage operation”](#) on page 53.
- Rehearse—See [“Understanding the Rehearse operation”](#) on page 55.
- Migrate—See [“Understanding the Migrate operation”](#) on page 56.

Understanding the Setup Storage operation

In this operation, the storage tasks are taken care of.

Mirroring

If you chose to create mirror disk groups, the mirror disk groups are created and data sync is ensured.

The operation aborts if all volumes of all the disk groups being migrated is mirrored with the corresponding volumes in the mirror disk group. In this scenario, 100% sync between all volumes is not verified. This check is performed as part of the Rehearse or Migrate operation. If only some volumes are mirrored, the system is cleaned up so that the Setup Storage operation can be performed again. Only for the disk groups where the mirror was created by the application migration plan are cleaned up; mirrors created earlier are not cleaned up.

Replication

In this operation, Replication configuration is created and replication is established between the source cluster node and target cluster for disk groups marked as **Create New**. Initially, a basic verification is done to ensure that all disk groups marked **Use Existing** has RVGs configured and secondaries added and disk groups marked with **Create New** should have RVG configuration only if it was pre-created before creating the plan. All verifications are done in par with the configuration discovered during plan creation. Once verification is complete, DCM log is added to all volumes in the disk group after which SRL Volume is created as per the size calculations. Primary RVG is then created on the disk groups marked as **Create New**.

Migration plan creates a single RVG for each disk group and will have all volumes of the disk group made part of it. If RVGs are already pre-created on a disk group, these RVGs are used. Disk groups, volumes, and volume sets (if present) are then created on the target cluster node according to its size calculated from the source cluster node. After this is done, the target cluster node is added as a secondary site to the Primary RVG after which the replication is started.

Setup Storage operation completes only when Replication Sync is 100% complete. If any application writes are happening to the volumes on the source, it may take more time for the sync to complete. If you notice the sync is not progressing for a long time, check whether SRL has not overflowed in the logs and if so, it is recommended to reduce or stop application writes so that sync gets completed. Replication configuration created by the plan will be with default values and finetuning is also not performed.

On the source cluster, IP resources are added to the service group where the disk group resource is configured. These resources monitors the IPs used for replication. If the disk group is shared (CVM), IP resource and RVGLogowner resources are created in a separate failover service group with the name `amVvR_sg_failover_plan_name` and is linked to the service group which has the CVMVolDG resource configured.

On the target cluster, a service group is created by the name `amVvR_sg_plan_name`. IP and disk group resources for the entities created by the plan is added to this service group. In case of a shared disk group, a parallel service group by the name

`amVvR_sg_parallel_plan_name` is created in which CVMVolDG resource is added and another failover service group is created `amVvR_sg_failover_plan_name` which contains the IP and RVGLogowner resources. The operation aborts if all RVGs in all disk groups which are part of the migration have secondaries added to it.

Understanding the Rehearse operation

In this operation, you can bring application online on the target cluster and test the application before performing the actual migration.

Mirroring

In this operation, the sync status of all volumes are checked after which cluster configuration of the selected service group is discovered in the source and translated to the target. The mirror disk group is then detached from the source cluster nodes and endian changes are performed on all volumes of the mirror disk group in the target cluster nodes.

After the endian changes are done, the service groups are brought online in the target cluster. After ensuring the application is running fine in the target, the service groups are taken offline and the target cluster configuration is removed.

Before removing the cluster configuration, a backup of the configuration is taken on the first node of the target cluster in the `/etc/VRTSvcs/conf/config` directory and the name will be in the following format:

```
main.cf_plan_name.date.time
```

The volumes in the mirror disk groups are then reattached to the corresponding volumes in the source disk group.

The operation aborts if all volumes between the source and mirror disk groups are not completely synced.

Replication

Initially, all RVGs are checked to ensure 100% sync. At times, if the application is writing to the volumes, data sync might be in progress to the secondary site and hence the Rehearse operation will not proceed. In such scenario, you can reduce or stop application writes so that data remains in sync.

Pre-requisites are done on target cluster nodes in order to create space optimized snapshot of volumes such as preparing volumes, cache volume, and cache object creation. IP, DiskGroup/CVMVolDg, RVGLogowner resources are then removed from the target cluster for the disk groups which are being migrated as part of the plan. This is to ensure no duplicate resources appear for an entity during cluster translation.

Cluster configuration of the selected service group and its dependencies are then discovered on the source and translated to the target. The file systems on all mounted volumes of the disk groups which are part of the replication is then frozen on the source for a moment and space optimized snapshots of these volumes are taken on the target with the help of VVR In-Band Control Messaging (vxibc). After the snapshots are taken, endian changes are performed on these snapshots and the service groups are brought online on the target. When service groups are brought online, the snapshot volumes gets mounted on the target cluster. Applications started on the target can write on these snapshot volumes until the cache volume fills up. After ensuring the application is running fine on the target, the service groups are taken offline and the target cluster configuration is removed.

Before removing the cluster configuration, the configuration is backed up on the first node of the target cluster in the `/etc/VRTSvcs/conf/config` directory and the name is in the following format:

```
main.cf_plan_name.date.time
```

The disk groups are then imported on the target cluster node and the snapshots are destroyed. The IP, DiskGroup/CVMVolDg, and RVGLowner resources are re-created on the target, as required, and then brought online.

Understanding the Migrate operation

In this operation, the actual migration happens. Service groups on the source are taken offline, so there will be a downtime when the operation is in progress. After the operation is complete, you can use the target cluster to run the application.

Mirroring

In this operation, the service groups in the source are taken offline before detaching the mirror disk groups from the source cluster nodes. This is to ensure data integrity before the final migration. After the mirror disk groups are detached from the source cluster nodes, endian changes are performed on these volumes in the target cluster nodes and the service groups are brought online in the target cluster.

The operation aborts if all volumes between the source and mirror disk groups are not completely synced.

Replication

In this operation, verification is done to ensure 100% sync between the source and target clusters. Cluster configuration is completely translated and service groups are taken offline from the source. Replication is stopped for all RVGs of

all disk groups part of the migration. Endian changes are performed on volumes part of replication on the target after which service groups are brought online. Secondary site is removed from the primary RVG and the primary RVG is also removed for all the disk groups marked as **Create New**.

Understanding the cleanup operation

Cleanup operation is an internal operation which takes place whenever a critical task fails in an operation. This operation takes care of cleaning up the managed hosts to ensure the system is in a clean and known state so that the operation can be performed again after taking corrective measures.

The status or the progress of this operation is not displayed in the user interface. To view the status of this operation, check the logs available in the managed hosts.

Mirroring: Steps performed as a part of cleanup:

- Setup Storage—If you had selected the **Create New** option in the wizard, mirror is removed for all volumes between source and target disk groups. Mirror disk groups are deleted.
- Rehearse and Migrate—All volumes of the mirror disk groups are reattached with the corresponding volumes in the source disk group. If the plan created the target cluster configuration, the configuration is backed up and then removed.

Replication: Steps performed as a part of cleanup:

- Setup Storage—Cleanup is performed only for disk groups marked as **Create New**. For these disk groups, replication is stopped (if established), secondary site is deleted and primary RVG is removed if the RVG was created by the plan. Disk groups created on the target cluster by the plan is also removed.
- Rehearse—Snapshots created for volumes are removed from the secondary. Cache object and cache volume are deleted. VCS resources for disk groups, IPs, etc. are restored. If the plan created the target cluster configuration, the configuration is backed up and then removed.
- Migrate—Replication, if stopped, is started again. If the plan created the target cluster configuration, the configuration is backed up and then removed.

Understanding the tasks executed in each operation

Endian changes

The application operation uses the `fscdsconv` utility to perform endian changes on the target cluster node. This utility creates a recovery file for the purpose of conversion. The default path for the creation of the recovery file is:

`/var/opt/VRTSsfmh/tmp`. The file system `/var` on the target cluster node must have more than 5 GB space.

If the file system does not have sufficient space for creating the file

- 1 On all nodes of the target cluster, create the `/opt/VRTSsfmh/etc/appmig.conf` file.
- 2 In the `appmig.conf` file, define the `APPMIG_RECOVERY_PATH=path` variable. Here, *path* is the full directory path where the recovery file needs to be created.

To check how much space is required in a VxFS file system for the recovery file as used by the `fscdsconv` utility, see: <http://www.veritas.com/docs/000012748>

Cluster configuration translation at the service group level

The following points describe how cluster configuration is translated at the service group level:

- All the nodes in the target cluster would be part of the `SystemList` for the service group in the target.
- Service groups on the source cluster that is being migrated are translated to the target cluster along with dependencies.
- While migrating service groups between CVM clusters, the CVM service group is not translated, but its dependencies are translated. To ensure that the dependencies are translated, Arctera recommends that you should not deselect the CVM service group during plan creation.
- Any setting with the following service group level attributes will be ignored in the target cluster:
 - `SystemZones`
 - `VCSi3Info`
 - `TriggerPath`
 - `ClusterList`
 - `Frozen`

- TriggersEnabled
- TypeDependencies
- AdministratorGroups
- OperatorGroups
- Administrators
- Operators
- Guests
- Triggers configured in the source cluster will not be translated to the target cluster. If you require triggers, trigger files must be manually copied and trigger attributes must be manually configured in the target cluster. You can use the pause task after the **Configure Cluster on Target** task or you can create a custom script to automate this.
- If the source and target operating systems are different, the ContainerInfo attribute at the group level is skipped in the target cluster.
- If a service group attribute is localized in the source cluster and the number of nodes in the target cluster is more than the number of nodes in the source cluster SystemList, the localized attribute is converted to a global attribute.
- If the source cluster has offline group dependencies and the target cluster is a single node cluster, offline group dependencies are deleted from the target cluster. Remote dependencies are converted to local dependencies.
- All site dependencies are deleted from the target cluster.

Cluster configuration translation at the resource level

The following points describe how cluster configuration is translated at the resource level:

- Resource attributes part of the service groups being migrated are translated from the source to target cluster. For example, if the DiskGroup attribute exists in the service group, the attribute value is translated from the source cluster to the target cluster.
- VxVM Mirroring—For each disk group resource, another disk group resource is created by the application migration operation. The name of this resource will start with `amMir_`. This is for internal housekeeping of snap disk group created for mirroring. This resource is skipped in the target cluster.
- If the source and target operating systems are different, ResContainerInfo and ContainerOpts attributes are skipped in the target cluster.

- If a resource attribute is localized in the source cluster and the number of nodes in the target cluster is more than the number of nodes in the source cluster SystemList, the localized attribute is converted to a global attribute.
- During Rehearse operation, IP resources like IP, IPMultiNIC, and IPMultiNICB are not created in the target cluster. This is to avoid accidental connection of clients to the application during Rehearse operation. These resources are converted to FileOnOff type to maintain resource dependencies.
- If any resource type configured in the source cluster is not available in target cluster, the corresponding resources are removed from the target cluster during cluster configuration.
- If the source and target cluster operating systems are different, IPMultiNICB resources are converted to IP type and MultiNICB type resources are converted to NIC type.
- VVR Replication—In the Rehearse operation, Volume and VolumeSet type resources managing volumes and volume sets being replicated are converted to FileOnOff type resource in the target cluster. This is done as snapshot volumes created in the target cluster cannot be managed using these resources.
- VVR Replication—In the Migrate operation, Volume and VolumeSet type resources managing volumes and volume sets being migrated are converted to FileOnOff type in the source cluster. This is to ensure taking the service group offline does not fail in the source cluster.
- If target cluster node is detected as an AWS EC2 instance, for every IP type resource an AWSIP type resource is created in the target configuration.

List of tasks executed

[Table 4-4](#) lists predefined tasks executed as a part of each operation. Each predefined task in an operation is marked as **Critical** or not. The table also lists the tasks for which the cleanup operation is performed and the state of the managed hosts after the operation.

If a task marked as critical fails, the operation aborts. If a task is not critical, the operation continues even if it fails. For most critical tasks, failure necessitates a cleanup of the system and an internal cleanup operation is performed.

See [“Understanding the cleanup operation”](#) on page 57.

Table 4-4 Mirroring—List of tasks executed

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Setup Storage			
Deploy scripts on the managed host	Yes	No	No change.
Verify mirror configuration	Yes	No	No change.
Prepare disk group for mirroring	Yes	Yes	After cleanup, system reverts to the state before the Setup Storage operation. Perform the Setup Storage operation again.
Create mirror between source disk group	Yes	Yes	After cleanup, system reverts to the state before the Setup Storage operation. Perform the Setup Storage operation again.
Snapshot sync progress	Yes	No	System will be in existing state and mirror will be under sync on the managed host.
Remove scripts from the managed host	No	No	System will be in existing state.
Rehearse			
Deploy scripts on the managed host	Yes	No	No change.
Verify mirror configuration	Yes	No	No change.
Configure target cluster	No	No	Target cluster configuration is created partially. You must validate the configuration and make necessary changes during the Pause task before proceeding.
Pause plan for user validation	Yes	No	No change.
Detach mirror disk group	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, mirror is reattached to the source. Target cluster configuration is removed after taking backup. Perform the Rehearse operation again.

Table 4-4 Mirroring—List of tasks executed (*continued*)

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Online service group in the target cluster	No	No	Service groups will be in partial state. Check the reason in the Pause task.
Pause plan for user validation	Yes	No	No change.
Offline service group in the target cluster	No	No	Service groups will be in partial state. Take all the service groups offline and check whether the mirror disk groups are deported; if not, manually deport the mirror disk groups in the Pause task.
Pause plan for user validation	Yes	No	No change.
Remove service group from the target cluster	Yes	No	Task will not fail.
Endian changes on the mirror disk group	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, mirror is reattached to the source. Perform the Rehearse operation again.
Reattach mirror disk group in the source cluster	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, mirror is reattached to the source. Perform the Rehearse operation again.
Snapshot sync progress	Yes	No	System will be in existing state. Mirror will be under sync on the managed host.
Remove scripts from the managed host	No	No	System will be in existing state.
Migrate			
Deploy scripts on the managed host	Yes	No	No change.
Verify snap mirror configuration	Yes	No	No change.
Configure target cluster	No	No	Target cluster configuration is partially created. You must validate the configuration and make necessary changes before proceeding.

Table 4-4 Mirroring—List of tasks executed (*continued*)

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Pause plan for user validation	Yes	No	No change.
Offline service group in the source cluster	Yes	No	Service groups will be in partial state. Do the following: ■ Check the reason why taking the service group offline failed. ■ Fix the issue. ■ Bring the service groups online. ■ Start the Migrate operation.
Detach snap mirror disk group	Yes	Yes	During cleanup, mirror is reattached to the source and target cluster configuration is removed. Bring the service groups online in the source cluster and perform the Migrate operation again.
Endian changes on the mirror disk group	Yes	Yes	During cleanup, mirror is reattached to the source and target cluster configuration is removed. Bring the service groups online in the source cluster and perform the Migrate operation again.
Online service group in the target cluster	No	No	Service groups will be in partial state in the target cluster.
Remove scripts from the managed host	No	No	System will be in existing state.

Table 4-5 Replication—Lists of tasks executed

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Setup Storage			
Deploy scripts on managed hosts	Yes	No	No change.
Verify replication configuration	Yes	No	No change.

Table 4-5 Replication—Lists of tasks executed (*continued*)

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Configure replication on source and target clusters	Yes	Yes	After cleanup, system reverts to the state before the Setup Storage operation. Check and resolve the issue. Perform the Setup Storage operation again.
Set up replication between source and target clusters	Yes	Yes	After cleanup, system reverts to the state before the Setup Storage operation. Check and resolve the issue. Perform the Setup Storage operation again.
Create RVGLowner type resource for shared disk group(s)	Yes	Yes	Task will not fail.
Replication sync progress	Yes	No	System will be in existing state and replication will be under sync between source and target cluster nodes.
Remove scripts from managed hosts	No	No	System will be in existing state.
Rehearse			
Deploy scripts on managed hosts	Yes	No	No change.
Verify replication configuration	Yes	No	No change.
Delete replication-related VCS resources on the target cluster	Yes	Yes	Task will not fail.
Perform target cluster configuration	No	No	Target cluster configuration is created partially. You must validate the configuration and make necessary changes during the Pause task before proceeding.
Pause plan for user validation	Yes	No	No change.
Set up volumes on target cluster for creating snapshot	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, snapshots are deleted. Target cluster configuration is removed. Replication-related resources are re-created on the target cluster. Check and resolve the issue. Restart the Rehearse operation.

Table 4-5 Replication—Lists of tasks executed (*continued*)

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Create snapshots on target by freezing VxFS mount points and replication	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, snapshots are deleted. Target cluster configuration is removed. Replication-related resources are re-created on the target cluster. Check and resolve the issue. Restart the Rehearse operation.
Perform endian changes on snapshot volumes on the target cluster	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, snapshots are deleted. Target cluster configuration is removed. Replication-related resources are re-created on the target cluster. Check and resolve the issue. Restart the Rehearse operation.
Bring service groups online on the target cluster	No	No	Service groups will be in partial state. Check the reason for the failure and correct it.
Pause plan for user validation	Yes	No	No change.
Take service groups offline on the target cluster	No	No	Service groups will be in partial state. Take all the service groups offline and during the Pause task, check whether all the disk groups are deported.
Pause plan for user validation	Yes	No	No change.
Remove target cluster configuration	Yes	Yes	Task will not fail.
Remove snapshot volumes and re-create replication-related VCS resources on the target cluster	Yes	Yes	After cleanup, system reverts to the state after the Setup Storage operation. During cleanup, snapshot removal is reattempted and replication-related resources are re-created on the target cluster.
Remove scripts from managed hosts	No	No	No change.
Migrate			
Deploy scripts on managed hosts	Yes	No	No change.

Table 4-5 Replication—Lists of tasks executed (*continued*)

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Verify replication configuration	Yes	No	No change.
Delete replication-related VCS resources on the target cluster	Yes	Yes	Task will not fail.
Perform target cluster configuration	No	No	Target cluster configuration is created partially. You must validate the configuration and make necessary changes during the Pause task before proceeding.
Pause plan for user validation	Yes	No	No change.
Reconfigure the source cluster to prevent plan failure	Yes	Yes	During cleanup, cluster configuration performed on the target is removed. Restart the Migrate operation
Take service groups offline on the source cluster	Yes	No	Service groups will be in partial state. Do the following: ■ Check the reason why taking the service group offline failed. ■ Fix the issue. ■ Bring the service groups online. ■ Start the Migrate operation.
Stop replication and remove the secondary site	Yes	Yes	During cleanup, secondary sites which might have been removed will be re-added and replication is started. Service groups will be in offline state on the source cluster. Bring the service groups back online. Check and resolve the issue. Restart the Migrate operation.
Perform endian changes on replicated volumes	Yes	Yes	During cleanup, secondary sites which might have been removed will be re-added and replication is started. Service groups will be in offline state on the source cluster. Bring the service groups back online. Check and resolve the issue. Restart the Migrate operation.
Bring service group online on the target cluster	No	No	Service groups will be in partial state on the target cluster.

Table 4-5 Replication—Lists of tasks executed (*continued*)

Task	Is the task critical?	Will cleanup be performed if task fails?	What happens to the system if the task fails?
Remove replication configuration from source and target cluster	No	No	Replication configurations added on the source and target cluster disk groups will not be fully cleaned up.
Remove scripts from managed hosts	No	No	No change.

Validations performed before migration plan execution

Before a plan is executed, validations are performed in order to check the state of the clusters, service groups, and the resources monitoring the disk groups.

In some scenarios:

- If the validation fails, the plan will be marked as an **Invalid Plan**.
- Even though the plan is marked as Invalid, the plan can be executed after taking necessary corrective steps.
- After taking the corrective steps, you must edit the plan before executing it.

[Table 4-6](#) lists the validations performed before a plan is executed and the state of the plan if the validation fails. It also lists when the plan will be marked as invalid or not and the necessary corrective steps.

Table 4-6 Validations performed

Object	State	Is the plan invalid?	Will the plan remain invalid forever?	Corrective Action
Source cluster	Deleted	Yes	Yes	Edit the plan, select any service group of an existing source cluster that needs to be migrated and execute the plan again.
Source cluster	Faulted	No	No	Resolve the issues in the source cluster. Ensure that the cluster is not in faulted state in InfoScale Operations Manager and execute the plan.

Table 4-6 Validations performed (*continued*)

Object	State	Is the plan invalid?	Will the plan remain invalid forever?	Corrective Action
Target cluster	Deleted	Yes	Yes	Edit the plan, select an existing target cluster to which migration needs to be performed and execute the plan.
Target cluster	Faulted	No	No	Resolve the issues in the target cluster. Ensure that the cluster is not in faulted state in InfoScale Operations Manager and execute the plan.
Service group	Deleted	Yes	Yes	Edit the plan and reselect any existing service group in the source cluster that needs to be migrated and execute the plan.
Dependent service groups	Deleted	Yes	No	Add the dependent service group that was deleted and execute the plan.
Disk group resource	Deleted	Yes	No	Add the disk group resource that was deleted and execute the plan.
Disk group resource offline	Offline	No	No	Bring the disk group resource online and execute the plan.
Disk group	Deleted	Yes	Yes	Edit the plan so that the current state of the managed hosts is detected by the wizard. Save the plan and execute it.
Volume	Deleted	Yes	Yes	Edit the plan so that the current state of the managed hosts is detected by the wizard. Save the plan and execute it.

Executing the application migration plan

After creating the migration plan, you must execute the plan.

If you had selected VVR Replication as the data migration type, you must have a valid VVR license and you must start the following services on the source and target cluster nodes before executing the plan:

- vxnm-vxnetd
- vras-vradmind
- vxrsyncd

To execute the application migration plan

- 1 In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane and click **Application Migration Plan**.
- 2 Right-click the plan that you created, click **Setup Storage**, and click **OK** in the Execute Migration Plan window.

On the **Latest Run** area of the screen, you can view the Setup Storage tasks that are being executed and the corresponding status.

The **Status** column of the Migration Plan window displays the overall status of the Setup Storage operation.

To verify whether mirror disk groups have been created in the source cluster in case of mirroring and whether remote disk groups have been created in the target cluster in case of replication:

- In the Management Server console, go to the **Server** perspective and expand **Manage** in the left pane.
- Expand the Organization, or **Uncategorized Hosts** to locate the managed host.
- Expand the managed host of the cluster and then expand **Disk Groups**.

To verify whether RVGs have been created in the disk group in case of replication:

- Expand the managed host in the target cluster.
- Expand RVGs.

If you had chose to pause any task as outlined in step 8 of [Creating an application migration plan](#) section, click **Resume** to resume the task that was paused.

- 3 Go to step 1. Right-click the plan that you created, click **Rehearse**, and click **OK** in the Execute Migration Plan window. You can execute the Rehearse operation multiple times.

On the **Latest Run** area of the screen, you can view the Rehearse tasks that are being executed and the corresponding status.

The **Status** column of the Migration Plan window displays the overall status of the Rehearse operation.

If you had chose to pause any task as outlined in step 8 of [Creating an application migration plan](#) section, click **Resume** to resume the task that was paused.

- 4 Go to step 1. Right-click the plan that you created, click **Migrate**, and click **OK** in the Execute Migration Plan window.

You can execute the Migrate operation only once. After you start the Migrate operation, you cannot stop or revert it. You must perform this operation only after executing all the other operations.

On the **Latest Run** area of the screen, you can view the Migrate tasks that are being executed and the corresponding status.

The **Status** column of the Migration Plan window displays the overall status of the Migrate operation.

If you had chose to pause any task as outlined in step 8 of [Creating an application migration plan](#) section, click **Resume** to resume the task that was paused.

- 5 To verify whether the service groups have been migrated successfully to the target cluster:
 - In the Management Server console, go to the **Availability** perspective and expand **Manage** in the left pane.
 - Expand the Organization, or **Uncategorized Clusters** to locate the cluster.
 - Go to the source cluster and expand **Service Groups**, the service groups must be offline.
 - Go to the target cluster and expand **Service Groups**, the service groups must be online.

Refer to the following sections:

See [“Understanding the tasks executed in each operation”](#) on page 58.

See [“Validations performed before migration plan execution”](#) on page 67.

Editing an application migration plan

After you create a plan, you can edit it if you need to modify any of the entries. When you edit a plan, the service group dependencies and disk group details are not populated from the previously saved plan, but are fetched afresh. This is to ensure that the current state of the service groups and disk groups in the source cluster are populated when you edit the plan.

Before you edit a plan, note the following:

- If you edit the plan before the Setup Storage operation:
 Mirroring—Re-select the disks for any mirror disk group which needs to be created for each source disk group.
 Replication—Re-select the disks and re-enter the RLINK details.
- Mirroring—If you edit a plan after the Setup Storage or Rehearse operation, mirror disk groups are also be displayed in the **Target Disk Information** panel. You must select the **Use Existing** option for source disk groups and select **Ignore** for mirror disk groups, if these were created as part of Setup Storage. The mirror disk groups will have a naming convention as
amMir_source_diskgroup_name.
- If you edit the plan after it is imported, any custom script execution task which was part of the imported plan must be removed and new custom script execution tasks must be added as necessary by uploading the custom scripts copied when the plan was exported.

To edit an application migration plan

- 1 In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
- 2 In the Migration Plan pane, right-click and click **Edit**.
 The Edit Migration Plan wizard opens.
- 3 Edit the fields in the migration plan as necessary.
- 4 Click **Save** to save the changes in the existing plan or click **Save As** to save the changes as a new plan.

Note: If you want to use the Save As option, you must change the plan name.

Deleting application migration plan(s)

You can delete a single plan or multiple plans. To select multiple plans, hold the **Ctrl** key and select the plans.

To delete application migration plan(s)

- 1 In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
- 2 In the Migration Plan pane, select the plan(s), right-click, and select **Delete**.
- 3 In the Delete Migration Plan(s) window, click **OK**.

The selected plan(s) are deleted.

Exporting application migration plan(s)

You can export any application migration plan to your system. It can be a single plan or multiple plans. Exported plans are saved in JavaScript Object Notation (JSON) format. Use a JSON viewer to view the contents of the exported file. The file created upon export can be used for importing the plan, if a plan needs to be imported on same or a different Management Server.

When you export multiple plans simultaneously, only a single JSON file is created. This file contains all the entries of the exported plans. You can edit the JSON file, but you must not modify the JSON syntax structure of the file. If you modify the JSON syntax structure, the import operation fails. If multiple plans are part of a single JSON file, you can remove the entries of a plan from the JSON file without modifying the syntax structure. Arctera recommends to change the value of only the following entries in the JSON file. Based on the value of these entries, the value of the remaining entities are fetched when the plan is being imported:

- `planName`—Name of the plan
- `description`—Description of the plan
- `selectedSGName`—Name of the selected service group
- `sourceClusName`—Name of the source cluster
- `targetClusName`—Name of the target cluster

When you edit the plan after exporting, you need to edit only the names of the entities listed above; the IDs are fetched during the import operation. If the entities mentioned are not available on the Management Server where the plan is being imported, the import operation fails.

Note: Custom scripts part of the plan are not exported. You have to manually copy the custom scripts either from the Management Server (path is available in the exported JSON file) or from the location from which you uploaded the custom scripts and use it when you edit the plan.

To export application migration plan(s)

- 1 In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
- 2 In the Migration Plan pane, select a plan, right-click, and click **Export**.
The Export Plan(s) window opens.
- 3 Click **Yes**.
The selected plan(s) are exported.

See [“Importing application migration plan\(s\)”](#) on page 73.

Importing application migration plan(s)

You can import any application migration plan from your system. It can be a single plan or multiple plans. To import multiple plans, all the plans must be part of the same file.

The following validations are performed during the import operation. If any of these validations fail, the plan will not be imported.

- No plan with same name as the plan being imported must exist.
- Plan name must follow recommendations.
See [“Creating an application migration plan”](#) on page 44.
- The names of the source cluster, target cluster, and the selected service group mentioned in the JSON file used for import must be available on the Management Server in which the plan is being imported.

If the plan is successfully imported, the plan description is prefixed with the [Imported] tag.

After the plan is imported, Arctera recommends to edit the plan to validate all the entries. If any custom script execution tasks were part of the plan, these entries are displayed in the Task Customization panel of the wizard. As the custom scripts may not be available in the Management Server where the plan is imported, delete these pre-existing tasks and add new custom script execution tasks in the same place by uploading custom scripts copied during the export operation.

To import application migration plan(s)

- 1 In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
- 2 In the Migration Plan pane, right-click and click **Import**.
- 3 Click **Browse** to select the plan(s).
- 4 Click **Import**.

The Import Plan(s) wizard opens.

The selected plans are imported.

See [“Exporting application migration plan\(s\)”](#) on page 72.

Viewing historical runs

The Historical Runs option allows you to view the execution history of a plan.

To view historical runs

- 1 In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
- 2 In the Migration Plan pane, select a plan, right-click and click **Historical Runs**.

The Historical Runs window displays the historical runs of the selected plan.

Viewing properties of an application migration plan

To view the properties of an application migration plan

- 1 In the Management Server console, go to the **Availability** perspective, expand **Solutions** in the left pane, and click **Application Migration Plan**.
- 2 In the Migration Plan pane, select a plan, right-click and click **Properties**.

The Properties window displays the following:

- Summary of the plan
- Disk group details
- Custom scripts and its location on the Management Server
- Network details
- RVG details (Displayed only if the Data Migration Type is VVR Replication)

Application migration logs

The following table lists the name and location of the log files. You can use these logs for troubleshooting.

Table 4-7 Log files

Name	Path
Add-on installation and uninstallation related logs	■ Windows MS: C:\ProgramData\Symantec\VRTSsfmcs\logs\vc sappmig_deployment.log ■ Linux MS: /var/opt/VRTSsfmcs/logs/vc sappmig_deployment.log
Application migration wizard logs	■ Windows MS: C:\ProgramData\Symantec\VRTSsfmcs\logs\WebDebugLog.txt C:\ProgramData\Symantec\VRTSsfmcs\logs\tomcat.log ■ Linux MS: /var/opt/VRTSsfmcs/logs/WebDebugLog.txt /var/opt/VRTSsfmcs/logs/tomcat.log
Application migration execution logs	AIX, Linux, and Solaris managed hosts: /var/opt/VRTSsfmh/logs/app_migration_plan_name.log

Index

C

- creating storage templates using NTFS file systems 15
- creating storage templates using volumes 16
- creating storage templates using VxFS file systems 9

D

- deleting storage templates 20
- downloading storage templates 19

I

- Impact Analysis report for volume migration 23

L

- locking storage templates 20

M

- migrating volumes 22
- moving disk groups
 - between enclosures 29
- moving volumes
 - between enclosures 24, 27

O

- overview
 - Storage Provisioning and Enclosure Migration Add-on 7

P

- provisioning storage 18

S

- storage provisioning 18
 - overview 7
- Storage Provisioning and Enclosure Migration Add-on
 - about 7
- storage template
 - creating using NTFS file systems 15

- storage template *(continued)*

- creating using volumes 16
 - creating using VxFS file systems 9
 - deleting 20
 - downloading 19
 - locking 20
 - unlocking 20
 - updating 17
 - uploading 19
 - using to provision storage 18

T

- template. *See* storage template

U

- unlocking storage templates 20
- updating storage templates 17
- uploading storage templates 19

V

- volume migration
 - about 22
 - by disk group 29
 - by enclosure 27
 - by host 24
 - deleting from schedule 31
 - impact analysis 23
 - pausing or resuming 31