



**SGI NAS Auto-Sync
User Guide v3.1.x**

Copyright © 2013 SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

SGI reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use.

Refer to the latest product announcement or contact your local SGI representative for information on feature and product availability.

This document includes the latest information available at the time of publication.

TRADEMARKS AND ATTRIBUTIONS

SGI, Silicon Graphics, Supportfolio and the SGI logo are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and other countries. Solaris and OpenSolaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. Microsoft and Windows are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks mentioned herein are the property of their respective owners.

Document Number: 007-6305-001

Contents

Preface	v
1 Introduction	1
Introducing Auto-Sync	1
About Auto-Sync Replication Algorithm	2
About Incremental Snapshots	2
About Transport Protocols	3
About the Directions of the Replication Stream	3
About Replicating Content	4
About Changes at Destination	5
2 Managing Auto-Sync	7
About Auto-Sync Installation	7
Viewing Summary Information about an Auto-Sync Service	8
Creating an Auto-Sync Service Instance	8
Modifying the Auto-Sync Properties	13
Enabling an Auto-Sync Service Instance	13
Disabling an Auto-Sync Instance	14
Destroying an Auto-Sync Service Instance	14
Deleting Snapshots Created by Auto-Sync	15
Executing the Auto-Sync Service by Administrative Action	16
Running the Auto-Sync Services in Series	16
Running the Auto-Sync Services in Parallel	17
Stopping the Auto-Sync Service	17
Recovering a Folder	18
Recovering Multiple Folders	18
Reversing the Direction of the Auto-Sync Service	18
Mounting an Auto-Sync Destination Folder	19
Viewing the Auto-Sync Log File	20
Replicating from a Snapshot	21
Viewing the Auto-Sync Statistics	21

3	Advanced Configuration	23
	Creating an Auto-Sync Service Instance with Advanced Options	23
	Changing Paths to Source or Destination Dataset	25
	Backing Up the Data to Other ZFS Storage Systems	26
	Verifying the Versions of Volumes	26
	Backing up to Oracle Solaris 11	26
	About Auto-Sync In-Reverse	27
	About the roles of Auto-Sync Service	28
	Migration Use Case	29
	Disaster Recovery Use Case	31
4	Advanced Settings	33
	Glossary	37

Preface

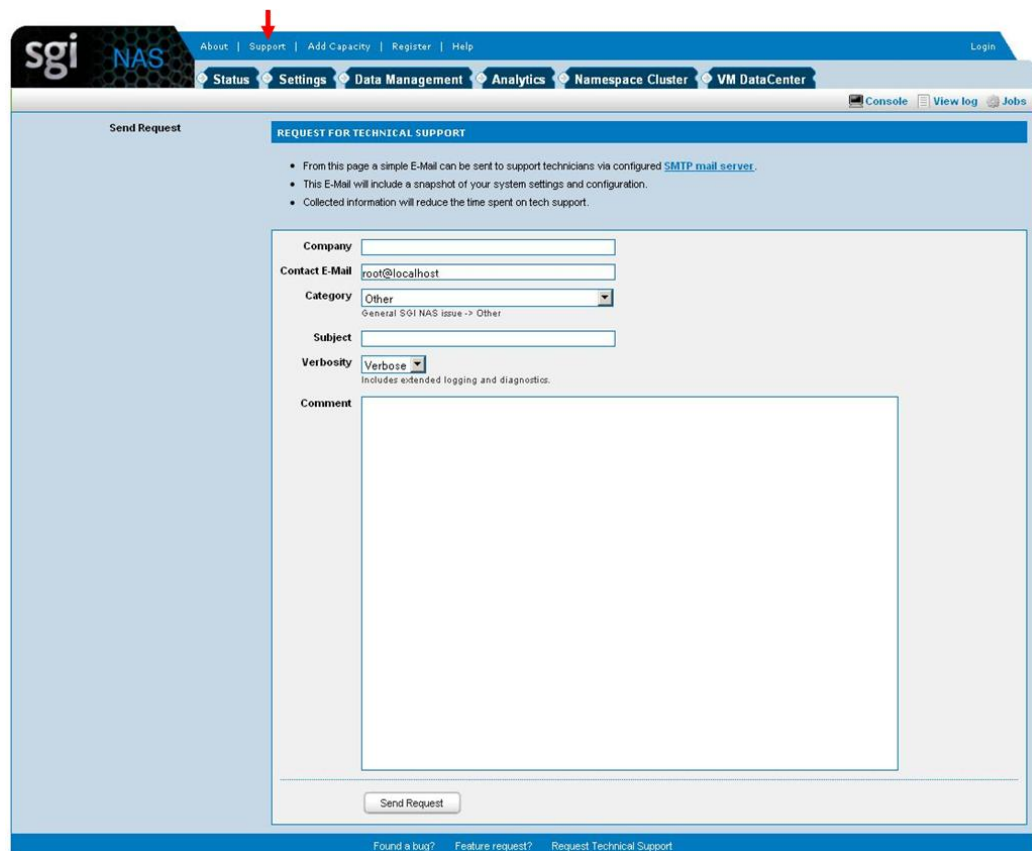
This documentation presents information specific to SGI NAS products. The information is for reference purposes and is subject to change.

Intended Audience

This documentation is intended for Network Storage Administrators and assumes that you have experience with data storage concepts, such as NAS, SAN, NFS, and ZFS.

Contacting Support

To contact support at SGI, click the Support link in NMV as marked with a red arrow on the screen below:



The screenshot shows the SGI NAS web interface. At the top, there is a navigation bar with the SGI logo and the text 'NAS'. To the right of the logo, there are links: 'About', 'Support', 'Add Capacity', 'Register', and 'Help'. A red arrow points to the 'Support' link. Below the navigation bar, there is a secondary menu with links: 'Status', 'Settings', 'Data Management', 'Analytics', 'Namespace Cluster', and 'VM DataCenter'. On the right side of this menu, there are icons for 'Console', 'View log', and 'Jobs'. The main content area is titled 'Send Request' and contains a 'REQUEST FOR TECHNICAL SUPPORT' form. The form includes the following fields and options:

- Company:** A text input field.
- Contact E-Mail:** A text input field containing 'root@localhost'.
- Category:** A dropdown menu with 'Other' selected. Below it, the text 'General SGI NAS issue -> Other' is visible.
- Subject:** A text input field.
- Verbosity:** A dropdown menu with 'Verbose' selected. Below it, the text 'Includes extended logging and diagnostics.' is visible.
- Comment:** A large text area for entering the request details.

At the bottom of the form is a 'Send Request' button. Below the form, there are three links: 'Found a bug?', 'Feature request?', and 'Request Technical Support'.

or type the following NMC command:

```
nmc:/$ support
```

which will then prompt for a subject and message.

Other resources

For licensing questions, please contact your SGI sales or support representative.

Product Support

SGI provides a comprehensive product support and maintenance program for its products.

For a full description of this program, do one of the following:

- See <http://www.sgi.com/support/>.
- If you are in North America, contact the Technical Assistance Center at 1 (800) 800 4SGI or contact your authorized service provider.
- If you are outside North America, see the following website for the appropriate Customer Service phone number: <http://www.sgi.com/support/supportcenters.html>.

Introduction

This section includes the following topics:

- [Introducing Auto-Sync](#)
- [About Auto-Sync Replication Algorithm](#)
- [About Incremental Snapshots](#)
- [About Transport Protocols](#)
- [About the Directions of the Replication Stream](#)
- [About Replicating Content](#)

Introducing Auto-Sync

Auto-sync is the SGI NAS data and metadata replication service. Auto-Sync combines the advantages of the built-in zfs send/receive functionality of the underlying OpenSolaris OS with tunable replication schedule.

The Auto-Sync plugin creates snapshots of the selected dataset and saves them in local or remote destination. You can use these snapshots as a backup or archive, as well as in the disaster recovery use case.

Auto-sync provides the following features:

- Replication of the ZFS folders, volumes, snapshots and zvols
- Metadata replication
- Root file system replication
- Replication in reverse-direction
- Tunable replication schedule
- Possibility to store snapshots in a few destinations
- Bandwidth throttling
- Compression and deduplication of the replication stream
- Auto-mounting and Auto-cloning of the snapshots at the destination
- Network trunking support

See Also:

- [Managing Auto-Sync](#)
- [Advanced Settings](#)

About Auto-Sync Replication Algorithm

During the first run Auto-Sync creates an initial snapshot of a source dataset and sends it to destination. The initial snapshot is a full copy of the source dataset. After Auto-Sync creates an initial snapshot, it functions according to the Auto-Sync replication algorithm.

The Auto-Sync replication algorithm includes the following tasks:

1. Creating a snapshot at source

Auto-Sync executes according to schedule or by the administrative action and creates a snapshot of a dataset. The name of the snapshot has the following pattern:

```
@sync-'COUNTER'_'DATE AND TIME'
```

2. Determining the latest identical snapshots

Auto-Sync compares the lists of snapshots at source and at destination and locates the latest pair of identical snapshots.

3. Comparing a snapshot at source with a snapshot at destination

Auto-Sync compares new snapshot at source with the latest snapshot at destination and sends the changes from source to destination.

See Also:

- [About Incremental Snapshots](#)
- [About Transport Protocols](#)
- [About the Directions of the Replication Stream](#)

About Incremental Snapshots

On every run, Auto-Sync verifies the latest identical snapshots at source and at destination. Therefore, every run of the Auto-Sync service must complete successful. If Auto-Sync cannot verify a pair of identical snapshots, then replication may fail or change the status to maintenance.

See Also:

- [About Auto-Sync Replication Algorithm](#)
- [About Transport Protocols](#)
- [About the Directions of the Replication Stream](#)

- [Introducing Auto-Sync](#)

About Transport Protocols

Auto-Sync uses the following transport protocols:

- **Remote Replication (RR)**

A multi-threaded end-to-end replication protocol that is designed from ground up. The protocol uses native ZFS `send/receive` command combined with `rrdaemon`. Auto-Sync uses the Remote Replication protocol for local-to-remote replication.

If you use the RR protocol, both source and destination appliances must be the SGI NAS appliances.

- **ZFS+SSH**

You must use the ZFS+SSH protocol, if you use a non-SGI NAS appliance as a destination for the replication stream.

- **ZFS**

Auto-sync uses built-in ZFS functionality for local to local replication.

About the Directions of the Replication Stream

You can create an Auto-Sync service that runs in one of the following directions:

- **Locally (L2L)**

Both source and replication datasets reside on a local SGI NAS appliance.

- **From Host (L2R)**

Source dataset resides on a local SGI NAS appliance. Destination dataset resides on a remote SGI NAS appliance or any other ZFS system.

You can specify multiple destinations for the local-to-remote Auto-Sync service.

You can reverse the direction of the replication stream after you create the service.

- **To Host (R2L)**

Source dataset is located on a remote SGI NAS appliance or any other ZFS system. Destination dataset resides on a local SGI NAS appliance.

You can reverse the direction of the replication stream after you create the service.

See Also:

- [Reversing the Direction of the Auto-Sync Service](#)

About Replicating Content

When you select a source dataset for the Auto-Sync service, you must specify whether to replicate sub-folders of the dataset or include parent dataset in the replication stream.

If you create an Auto-Sync service instance in NMC and want to replicate the content of the folder you must enable the `dircontent` property. If you want to use the `dircontent` property, you must select the following symbol: `/*`.

The following diagrams compare replication with enabled and disabled replicate content property.

Figure 1-1: Replication with Replicate Content Disabled





Figure 1-2: Replication with Replicate Content Enabled

**See Also:**

- [Advanced Settings](#)
- [Creating an Auto-Sync Service Instance](#)
- [Modifying the Auto-Sync Properties](#)

About Changes at Destination

Auto-Sync does not support any changes in data or metadata in the destination folder. If you mount the destination folder and perform any changes in it, Auto-Sync identifies them as a loss of synchronization between source and destination. Auto-Sync processes this condition as a recoverable error and tries to recover the Auto-Sync service instance. Therefore, it discards all the changes in the destination folder.

If you want to use the destination folder and apply any changes to it, SGI recommends that you disable or delete any Auto-Sync service instances created for this folder.

See Also:

- *SGI NAS User Guide*
- [About Auto-Sync Replication Algorithm](#)
- [About Replicating Content](#)
- [Creating an Auto-Sync Service Instance](#)
- [Advanced Settings](#)

Managing Auto-Sync

This section includes the following topics:

- [Viewing Summary Information about an Auto-Sync Service](#)
- [Creating an Auto-Sync Service Instance](#)
- [Modifying the Auto-Sync Properties](#)
- [Enabling an Auto-Sync Service Instance](#)
- [Disabling an Auto-Sync Instance](#)
- [Destroying an Auto-Sync Service Instance](#)
- [Deleting Snapshots Created by Auto-Sync](#)
- [Executing the Auto-Sync Service by Administrative Action](#)
- [Running the Auto-Sync Services in Series](#)
- [Running the Auto-Sync Services in Parallel](#)
- [Stopping the Auto-Sync Service](#)
- [Recovering a Folder](#)
- [Reversing the Direction of the Auto-Sync Service](#)
- [Viewing the Auto-Sync Log File](#)
- [Replicating from a Snapshot](#)
- [Viewing the Auto-Sync Statistics](#)

About Auto-Sync Installation

The Auto-Sync plugin installs automatically when you install SGI NAS.

❖ *To view plugins, using NMC:*

◆ Type:

```
nmc:/$ show plugin
```

❖ *To view plugins, using NMV:*

1. Click **Settings > Appliance**.

2. In the Administration panel, click **Plugins**.

Viewing Summary Information about an Auto-Sync Service

Summary information page includes the name, replication schedule, status, and other information about Auto-Sync services. You can view summary information about any Auto-Sync service.

❖ *To view the summary information about an Auto-Sync service using NMV:*

1. Click **Data Management > Auto-Services**.
2. In the **Auto-Sync Services** panel, click **Show**.

❖ *To view the summary information about the Auto-Sync service using NMC:*

1. Log in to NMC using an SSH-client.
2. Type the `root` credentials.
3. Type:

```
nmc:/$ show auto-sync
```

This command displays the list of all auto-sync services.

4. Alternatively, type:

```
nmc:/$ show auto-sync <service instance>
```

This command displays detailed information for the selected service instance.

Creating an Auto-Sync Service Instance

You can schedule an Auto-Sync service for any dataset, such as volume, folder, snapshot, or zvol.

❖ *To create Auto-Sync service instance using NMV:*

1. Click **Data Management > Auto-Services**.
2. In the **Auto-Sync Services** panel, click **Create**.
3. In the **Create Auto-Sync Service** page, specify the following parameters:

Table 2-1: Auto-Sync Properties

Parameter in NMV	Description
Name	Optional parameter. Specifies a custom name of a service instance.
Direction	Auto-Sync can execute a service in the following directions: <ul style="list-style-type: none"> • Locally • To Remote Host • From Remote Host
Replicate Content	Replicate the content of a folder. If you leave this property unselected, Auto-Sync replicates the parent folder as well.
Periodic Interval	Periodic Interval consists of two parameters: Period and Frequency. Frequency depends on period. For example, if you set period to hourly and frequency to 4, then Auto-Sync runs service instance every 4 hours. If you set period to daily and frequency to 5, then Auto-Sync runs the service instance every 5 days, and so on.
Keep on Source	The number of snapshots that you want to keep at source dataset. Alternatively, specify a period in YDHMS format. Auto-Sync keeps three additional snapshots after retention period expires. Auto-Sync deletes these snapshots on the next scheduled run.
Keep on Destination	The number of snapshots that you want to keep at destination dataset. Alternatively, specify a period in YDHMS format. Auto-Sync keeps three additional snapshots after retention period expires. Auto-Sync deletes these snapshots on the next scheduled run.
If you select Locally	
Local Source Folder/ Zvol	A dataset on the SGI NAS appliance that you want to replicate.
Local Destination Folder	Destination for Auto-Sync service instance.
If you select To Remote Host	
Local Source Folder/ Zvol	A dataset on this SGI NAS appliance that you want to replicate.

Table 2-1: Auto-Sync Properties (Continued)

Parameter in NMV	Description
Multi-destination	<p>You need to specify this parameter, if you select to replicate to remote destination.</p> <p>Multi-destination enables you to set few destinations for the Auto-Sync service instance.</p> <p>Multi-destination includes the following parameters:</p> <ul style="list-style-type: none"> • Host The IP address or host name of the remote destination appliance. • Folder A dataset on the remote SGI NAS appliance that you want to use as a replication destination. • Keep on destination The number of snapshots that you want to keep at destination. Alternatively, specify the time in YDHMS format. You must set this parameter for every destination.
If you select From Remote Host	
Local Destination Folder/Zvol	A dataset on the SGI NAS appliance that you want to use as a destination for replication.
Remote Source Host	The IP address or host name of remote SGI NAS appliance on which the dataset that you want to replicate resides.
Remote Source Folder/Zvol	The dataset that you want to replicate.

4. Optionally, click More Options.

This option enables you to configure advanced settings.

For more information, see [Advanced Settings](#).

5. Click **Create Service**.

❖ *To create Auto-Sync service instance using NMC:*

1. Log in to NMC using an SSH-client.

2. Type the root credentials.

3. Type:

```
nmc:/$ setup auto-sync create
```

4. Select a replication interval.

The options are:

- second
- minute
- hourly

- daily
 - weekly
 - monthly
5. Specify a replication period.
The replication period depends on the replication interval.
 6. Select a source host.
 - a. If source dataset resides on local appliance, press **Enter** and accept the default value `localhost`.
 - b. If source dataset is located on the remote appliance, type the IP address or host name of the source host.

**Note:**

If source dataset is located on the remote SGI NAS appliance,

7. Select a source dataset.
You can select volume, folder, or zvol.

**Note:**

Select `./*`, if you want to replicate the content of the selected dataset, but not the dataset itself.

For more information, see [About Replicating Content](#)

8. Select a destination folder.
 - a. If source dataset resides on local appliance, select one of the following:
 - 1) Press **Enter** to accept the default value `localhost`; then select a local dataset.
 - 2) Type the IP address or host name of the remote SGI NAS appliance; then select the remote dataset.

**Note:**

If destination dataset resides on the remote SGI NAS appliance, you may be prompted to enter the super-user password for the remote SGI NAS appliance to establish SSH-binding.

- c. If source dataset resides on the remote appliance, select a local dataset that you want to use as a replication destination.
9. Specify the number of snapshots to keep at source.
Alternatively, specify a period in YDHMS format.
10. Specify the number of snapshots to keep at destination.
Alternatively, specify a period in YDHMS format.

- 11.** If you selected to replicate from local to remote SGI NAS appliance, you can specify more than one destination. Would you like to specify additional destination? `y`
- a.** Select additional destination host.
 - b.** Select additional destination folder.
 - c.** Specify the number of snapshots to keep at destination.
Alternatively, specify a period in YDHMS format. You must set this parameter for every destination.
- SGI NAS verifies configuration of the service.



Before you configure advanced settings, read [Advanced Settings](#). SGI recommends that you do not modify the advanced settings.

- 12.** Optionally, modify advanced settings.
- d.** In the **Copy zfs**, select `yes` or `no`.
 - e.** In the **Rate limit**, type the maximum allowed speed in KB.
 - f.** Optionally, specify a custom name for the Auto-Sync service instance.
 - g.** Optionally, type a commentary.
- For more information, see [Advanced Settings](#).

- 13.** Confirm the creation of the Auto-Sync service. Type `y`.

Example:

```

1      show auto-sync ':datapool-000' state and properties
      'show auto-sync :datapool-000 -v'

2      show auto-sync ':datapool-000' log
      'show auto-sync :datapool-000 log'

3      show volume 'datapool' I/O statistics
      'show volume datapool iostat'
```

See Also:

- [Modifying the Auto-Sync Properties](#)
- [Advanced Settings](#)
- [Enabling an Auto-Sync Service Instance](#)
- [Executing the Auto-Sync Service by Administrative Action](#)

Modifying the Auto-Sync Properties

After you create the Auto-Sync service instance, you can modify some of the service properties.

❖ *To modify the Auto-Sync properties using NMV:*

1. Click **Data Management > Auto-Services**.
2. In the **Auto-Sync Services** panel, click **Show**.
3. Click on the service that you want to modify.
4. Click Save.

❖ *To modify the Auto-Sync properties using NMC:*

1. Type:

```
nmc:/$ setup auto-sync <service instance> property
```
2. Select the property that you want to modify.
3. Save the changes. Type *y*.
4. Alternatively, to discard the changes, type *n*.

See Also:

- [Auto-Sync Properties](#)
- [Advanced Settings](#)

Enabling an Auto-Sync Service Instance

You can enable an Auto-Sync service instance, in case it was previously disabled by NMS or by Administrator.

❖ *To enable an Auto-Sync service instance using NMV:*

1. Click **Data Management > Auto-Services**.
2. In the **Auto-Sync Services** panel, click **Show**.
3. Select the checkbox near the service that you want to enable.
4. Click **Enable**.

❖ *To enable an Auto-Sync service instance using NMC:*

1. Type:

```
nmc:/$ setup auto-sync <service instance> enable.
```

See Also:

- [Disabling an Auto-Sync Instance](#)
- [Destroying an Auto-Sync Service Instance](#)
- [Viewing the Auto-Sync Log File](#)

Disabling an Auto-Sync Instance

You can disable an Auto-Sync service instance, if you want to temporarily suspend synchronization between source and destination folders.

❖ *To enable an Auto-Sync service instance using NMV:*

1. Click **Data Management > Auto-Services**.
2. In the **Auto-Sync Services** panel, click **Show**.
3. Select the checkbox near the service that you want to enable.
4. Click **Disable**.

❖ *To enable an Auto-Sync service instance using NMC:*

1. Type:

```
nmc:/$ setup auto-sync <service instance> enable.
```

See Also:

- [Enabling an Auto-Sync Service Instance](#)
- [Destroying an Auto-Sync Service Instance](#)
- [Viewing the Auto-Sync Log File](#)

Destroying an Auto-Sync Service Instance

You can destroy an Auto-Sync service instance to stop synchronizing source and destination folders. If you destroy an Auto-Sync service instance, you cannot resume it.



Note:

If you destroy an Auto-Sync service instance, you do not delete the snapshots that it has created.

❖ *To destroy an Auto-Sync service instance using NMV:*

1. Click **Data Management > Auto-Services**.
2. In the **Auto-Sync Services** panel, click **Show**.
3. Select the checkbox near the service that you want to enable.
4. Click **Delete**.

❖ *To destroy an Auto-Sync service instance using NMC:*

1. Type:

```
nmc:/$ setup auto-sync <service instance> destroy
```

2. Confirm the operation, type *y*.

See Also:

- [Enabling an Auto-Sync Service Instance](#)
- [Destroying an Auto-Sync Service Instance](#)
- [Viewing the Auto-Sync Log File](#)

Deleting Snapshots Created by Auto-Sync

You may want to delete the snapshots that Auto-Sync service created to free space on disk or to discard changes.

❖ *To delete snapshots created by Auto-Sync using NMV:*

1. Click **Data Management > Auto-Services**.
2. In the **Auto-Sync Services** panel, click **Show**.
3. In the **Snapshots** column, click on the snapshot count.
4. Optionally, filter the results.
5. Select the checkboxes near the snapshots that you want to delete.
6. Alternatively, select the checkbox near the **Snapshot** column to mark all snapshots on this page for deletion.
7. Click **Delete**.

❖ *To delete snapshots created by Auto-Sync using NMC:*

1. Type:

```
nmc:/$ setup folder <folder> destroy_snapshots -n
<template>
```

Example:

```
nmc:/$ setup folder datapool/accounts destroy_snapshots
-n AutoSync*
```

System response:

The following list of snapshots going to be deleted:

```
datapool/accounts@AutoSync-0_2012-11-29-16-01-00
datapool/accounts@AutoSync-1_2012-12-03-12-00-07
datapool/accounts@AutoSync-1_2012-12-03-15-00-07
datapool/accounts@AutoSync-0_2012-12-03-16-00-25
```

Ready to delete? (y/n)

2. Confirm the operation, type **y**.

Executing the Auto-Sync Service by Administrative Action

You can run the existing Auto-Sync service instance any time, not waiting for the scheduled time. You may need this functionality to immediately synchronize the source and destination datasets of any service.

- ❖ *To execute the Auto-Sync service by administrative action using NMV:*
 1. Click **Data Management > Auto-Services**.
 2. In the **Auto-Sync Services** panel, click **Show**.
 3. Select the checkbox near the service that you want to enable.
 4. Click **Start**.
 5. Alternatively, you can execute an Auto-Sync service on the Auto-Sync properties page.
 - a. Click on the Auto-Sync service instance that you want to run.
 - b. In the **Update Auto-Sync Service: <service instance>** pane, click **Run Now**.
- ❖ *To execute the Auto-Sync service by administrative action using NMC:*
 1. Type:

```
nmc:/$ setup auto-sync <service instance> run
```
 2. Press any key to quit.

See Also:

- [Enabling an Auto-Sync Service Instance](#)
- [Disabling an Auto-Sync Instance](#)
- [Destroying an Auto-Sync Service Instance](#)
- [Modifying the Auto-Sync Properties](#)
- [Viewing the Auto-Sync Log File](#)

Running the Auto-Sync Services in Series

Execution of the Auto-Sync service may impact performance of the SGI NAS appliance. If you have multiple Auto-Sync services that run at the same time, you may want to configure them to execute one after another.

This functionality is only available in NMC.

- ❖ *To run the Auto-Sync services in series using NMC:*
 1. Type:

```
nmc:/$ setup auto-sync serialize-all
```

System response:

```
<service instance> has been serialized
<service instance> has been serialized
```

See Also:

- [Running the Auto-Sync Services in Parallel](#)

Running the Auto-Sync Services in Parallel

You can disable execution of the Auto-Sync services in series and configure them to run in parallel. It may affect performance of the SGI NAS appliance.

This functionality is only available in NMC.

❖ *To run Auto-Sync services in parallel using NMC:*

1. Type:

```
nmc:/$ setup auto-sync unserialize-all
```

System response:

```
<service instance> has been serialized
<service instance> has been serialized
```

See Also:

- [Running the Auto-Sync Services in Parallel](#)

Stopping the Auto-Sync Service

You can forcefully stop an Auto-Sync service instance in case of emergency.

This functionality is only available in NMV.

❖ *To stop Auto-Sync service using NMV:*

- 1.** Click **Data Management > Auto-Services**.
- 2.** In the **Auto-Sync Services** panel, click **Show**.
- 3.** Select one or more checkboxes near service that you want to stop.
- 4.** Click **Abort**.
- 5.** Alternatively, you can stop an Auto-Sync service on the Auto-Sync properties page.
 - a.** Click on the Auto-Sync service instance that you want to stop.
 - b.** In the **Update Auto-Sync Service: <service instance>** pane, click **Abort**.

See Also:

- [Disabling an Auto-Sync Instance](#)
- [Destroying an Auto-Sync Service Instance](#)
- [Enabling an Auto-Sync Service Instance](#)

Recovering a Folder

Auto-Sync can restore the synchronization between source and destination. You may want to use this functionality to restore an Auto-Sync service, if you have applied changes to destination folder.

This functionality is only available in NMC.

❖ *To recover a folder using NMC:*

1. Type:

```
nmc:/$ setup auto-sync recover-folder -n <folder_name>
```

System response:

```
Recovery completed successfully for folder: <folder_name>
```

Recovering Multiple Folders

You can recover multiple folders using a single command.

This functionality is only available in NMC.

❖ *To recover a folder using NMC:*

1. Type:

```
nmc:/$ setup auto-sync recover-folder -a
```

Reversing the Direction of the Auto-Sync Service

Auto-Sync can reverse the direction of the replication stream. Therefore, source and destination change the roles.

When you change the direction of a service instance, Auto-Sync performs the following tasks:

1. Sets the access permissions of source dataset to read-only.
2. Enables write access at destination folder.
3. Creates a new Auto-Sync service instance with the same parameters at destination SGI NAS appliance.



Make sure that source folder is available. Otherwise, you may lose the data from last replication.

- ❖ *To reverse the direction of the Auto-Sync service using NMV:*
 1. Click **Data Management > Auto-Services**.
 2. In the **Auto-Sync Services** panel, click **Show**.
 3. Click on the Auto-Sync service instance that you want to reverse.
 4. In the **Update Auto-Sync Service: <service instance>** pane, click **Reverse**.
 5. In the confirmation dialog box, click **Ok**.
- ❖ *To reverse the direction of the Auto-Sync service using MMC:*
 1. Type:

```
nmc:/$ setup auto-sync <folder_name> reverse-direction
```

 - If reverse-direction is disabled you may get the following:
System response:
Reverse-direction feature is currently disabled
Do you want enable 'reverse-direction' feature?
 - a. Type `y`.
System response:
'Reverse-direction' feature successfully enabled
 - b. Type the reverse-direction command once again.
 2. Confirm the operation, type `y`.
System response:
Reversed direction for 1 auto-sync instance

See Also:

- [Recovering a Folder](#)
- [Enabling an Auto-Sync Service Instance](#)
- [Disabling an Auto-Sync Instance](#)
- [Advanced Configuration](#)
- [About Auto-Sync In-Reverse](#)
- [Migration Use Case](#)
- [Disaster Recovery Use Case](#)

Mounting an Auto-Sync Destination Folder

You can mount the Auto-Sync destination dataset and use as a regular folder. You can change data or metadata in the mounted folder. However, when Auto-Sync executes the next time, these changes are handled as a

lack of synchronization between source and destination. Auto-Sync overwrites the changes that you made in the destination folder with data from source.



Warning:

SGI recommends that you do not mount the destination folder.

However, you can delete an Auto-Sync service instance. After that Auto-Sync mounts the destination folder automatically.

This functionality is only available in NMC.

❖ *To mount an Auto-Sync destination folder:*

1. Type:

```
nmc:/$ setup auto-sync :source-initial-000 mount-destination
```

System response:

```
Mounted destination for auto-sync instance
```

See Also:

- [About Changes at Destination](#)
- [About Incremental Snapshots](#)
- [About Auto-Sync Replication Algorithm](#)
- [Advanced Settings](#)

Viewing the Auto-Sync Log File

The Auto-Sync log file provides detailed information about activity of the Auto-Sync service. You can monitor the Auto-Sync log file to verify that Auto-Sync jobs complete successfully. The level of detail in the Auto-Sync log file depends on the Auto-Sync property `trace_level`.

❖ *To view the Auto-Sync log file using NMV:*

1. Click **Data Management > Auto-Services**.
2. In the **Auto-Sync Services** panel, click **Show**.
3. Click on the Auto-Sync service instance.
4. In the **Update Auto-Sync Service: <service instance>** pane, click **View Log**.
Scroll down to see the required event.
5. Optionally, in the **Lines to show** field, adjust the number of lines that you want to view.

❖ *To reverse the direction of the Auto-Sync service using NMC:*

1. Type:

```
nmc:/$ show auto-sync <folder_name> log
```

This command displays all messages in the Auto-Sync log file.

2. Alternatively, type:

```
nmc:/$ show auto-sync <folder_name> logtail
```

This command displays last 30 lines of the Auto-Sync log file.

See Also:

- [Advanced Settings](#)

Replicating from a Snapshot

You can specify a snapshot that you want to use as an initial snapshot for the Auto-Sync replication.

This functionality is only available in NMC.

❖ *To replicate from a snapshot using NMC:*

1. Type:

```
nmc:/$ setup auto-sync create -S <snapshot>
```

2. Follow the steps described in [Creating an Auto-Sync Service Instance](#).

You must select a parent dataset of the selected snapshot as a source dataset.

See Also:

- [Advanced Settings](#)

Viewing the Auto-Sync Statistics

Auto-Sync provides the real-time statistics. You can monitor the Auto-Sync work and perform the actions that are required to correct any errors.



You can monitor the Auto-Sync statistics only when an Auto-Sync service executes.

❖ *To view the Auto-Sync statistics using NMV:*

1. Click **Data Management > Auto-Services**.

2. In the **Auto-Sync Services** panel, click **Show**.

View the progress bar.

❖ To view the Auto-Sync statistics using NMC:

1. Type:

```
nmc:/$ setup auto-sync <service instance> stats
```

See Also:

- [Advanced Settings](#)
- [Creating an Auto-Sync Service Instance](#)
- [Modifying the Auto-Sync Properties](#)
- [Stopping the Auto-Sync Service](#)
- [Viewing the Auto-Sync Log File](#)

Advanced Configuration

This section includes the following topics:

- [Creating an Auto-Sync Service Instance with Advanced Options](#)
- [Changing Paths to Source or Destination Dataset](#)
- [Backing Up the Data to Other ZFS Storage Systems](#)
 - [Verifying the Versions of Volumes](#)
 - [Backing up to Oracle Solaris 11](#)
- [About Auto-Sync In-Reverse](#)
 - [Migration Use Case](#)
 - [Disaster Recovery Use Case](#)



This chapter describes advanced configuration for Auto-Sync service instances. SGI recommends that you use this functionality, if you are an experienced SGI NAS user. SGI does not recommend to use any of these settings if you configure the Auto-Sync service for the first time.

Creating an Auto-Sync Service Instance with Advanced Options

You can create an Auto-Sync service instance with advanced options. This functionality is only available in NMC.

❖ *To create an Auto-Sync service with advanced options using NMC:*

1. Type:

```
nmc:/$ setup auto-sync create <option>
```

2. Add an option from [Table 3-1](#) to the command in [Step 1](#).

Table 3-1: Advanced Options for the Create Auto-Sync Command

Option	Description
-s	Source dataset. You can specify volume, folder, zvol, or snapshost. Example: <code>nmc:/\$ setup auto-sync create -S <snapshot></code>
-d	Destination dataset. You can specify volume or folder.
-R	Resume an Auto-Sync service. This option enforces Auto-Sync to skip validation of source and destination. You may want to use this functionality to re-create the service that already contains the Auto-Sync marker snapshots in both source and destination folders.
-e	Exclude from the replication stream the folders that match a pattern or a list of patterns. Example: <code>nmc:/\$ setup auto-sync create -e *.txt</code> Use the standard Unix patterning syntax for name filtering. By default, all files and folders are included in the replication stream.
-x	This option enables you to specify a custom suffix for the latest snapshot. Default value is <code>latest</code> .
-1	This option makes an Auto-Sync service instance to run only once at scheduled date and time. The Auto-Sync service instance changes the status to maintenance after execution.
-T <time of day>	Time of a day.
-D <day of the month>	Date of a month.
-i <interval>	Time interval. The options are: <ul style="list-style-type: none"> • Second • Minute • Hourly • Daily • Weekly • Monthly
-p <period>	The period property depends on the interval. For example, if <code>interval</code> is set to <code>hourly</code> and <code>period</code> is set to 4, then the Auto-Sync service instance executes every 4 hours.
-c	Automatically create a clone of the replicated service at destination.

Table 3-1: Advanced Options for the Create Auto-Sync Command

Option	Description
-l <number>	Rate limit, Kb/s. Specifies the maximum speed limit for the replication stream.
-u <custom name>	A custom name of the Auto-Sync service instance. By default, Auto-Sync creates a service with the following naming pattern: dataset-index_number. Example: datapool/docs-000 where datapool/docs is the name of the dataset and 000 is an index number of the Auto-Sync service.
-l <comment>	Optional commentary.
-r	Replicate a dataset recursively. Auto-Sync ignores this option, if source dataset is zvol.
-N	Replicate a dataset non-recursively.
-g	Create an Auto-Sync service instance with enabled reverse-direction command.
-Z	If you specify this option, then Auto-Sync automatically unmaps a zvol during the reverse-direction operation.
-E	Run an Auto-Sync service instance in daemon mode.
-S	Create an Auto-Sync service instance using a snapshot as source dataset.
-t	Enables you to specify an IP addresses or a host names for TCP trunking. Requires more than one network interface located on the source appliance. This property enables Auto-Sync to split the network traffic across specified NICs for higher bandwidth throughput.
-F	Generate fault notifications about the actions of the Auto-Sync service instance. This option enables you to identify, track and resolve any issues that affect Auto-Sync work. Auto-Sync also forwards a fault report by e-mail, if mailer is properly configured.

Changing Paths to Source or Destination Dataset

You can change the paths to source or destination dataset, if you previously changed location of the datasets, modified an IP address of the source or destination host, or performed any other changes.

This functionality is only available in NMC.

❖ *To change the paths to source or destination datasets using NMC:*

1. Type:

```
nmc:/$ setup auto-sync <service instance> edit-paths
```

2. In the `From host` field, type the name of the new source host.
3. In the `From fs`, type the path to new location of the source dataset.
4. In the `To host` field, type the name of the new destination host.
5. In the `To fs`, type the path to new location of the destination dataset.
6. If you have multiple destinations for the Auto-Sync service or want to add a new destination, type `y`. Otherwise, type `n`.

Auto-Sync prints the detailed information about the Auto-Sync service.

Backing Up the Data to Other ZFS Storage Systems

Auto-Sync can replicate data from an SGI NAS appliance to other ZFS storage systems.

Supported ZFS storage systems:

- Oracle Solaris 10 x32-bit
- Oracle Solaris 11 x32-bit



SGI recommends that you perform this operation if you are an advanced ZFS user.

Verifying the Versions of Volumes

Before you create an Auto-Sync service instance, you must verify that SGI NAS and the destination ZFS storage system use identical volume versions.

❖ *To verify volume version:*

1. On SGI NAS:
 - a. Log in to NMC using an ssh client.
 - b. Type:


```
nmc:/$ show volume version
```
2. On the destination ZFS storage system:
 - a. Log in to console.
 - b. Type:


```
# zpool get version <volume>
```

Backing up to Oracle Solaris 11

After you verify versions of volumes, you can start to back up data to the ZFS storage system.

- ❖ *To back up to Oracle Solaris 11:*
 1. On the Oracle Solaris 11:
 - a. Log in to Oracle Solaris 11 command line interface using an ssh client.
 - b. Enable the root login. Type:


```
# rolemod -K type=normal root
```
 - c. Modify /etc/ssh/sshd_config:


```
PermitRootLogin yes
```
 - d. Restart the ssh service. Type:


```
# svcadm restart ssh
```
 2. On the SGI NAS appliance:
 - a. Log in to NMC using an ssh client.
 - b. Bind the Oracle Solaris 11 host using ssh. Type:


```
nmc:/$ setup network ssh-bind
```
 - c. Create an Auto-Sync service instance. Type:


```
nmc:/$ create auto-sync -s zfs+ssh://<localhost>/<dataset>
```
 - d. Follow the on-screen instructions.
For more information, see [Creating an Auto-Sync Service Instance](#).
 - e. In the **Destination host** field, type an IP address or host name of the remote ZFS Storage System.

See Also:

- [Creating an Auto-Sync Service Instance](#)
- [Modifying the Auto-Sync Properties](#)
- [Advanced Settings](#)

About Auto-Sync In-Reverse

Auto-Sync provides the capability to reverse the direction of the replication stream. Then you can switch the roles of source and destination without re-creating the service.

When you create an Auto-Sync service instance, you must enable the `reverse-direction` property that controls the possibility to run an Auto-Sync service in-reverse. You can only modify this property in NMC.

If you enable the `reverse-direction`, Auto-Sync creates a service instance on source SGI NAS appliance and a copy of the auto-sync service instance on the destination appliance.

You can reverse the direction of the replication, if you configure local-to-remote or remote-to-local Auto-Sync service.



You must bind local and remote appliances using the Fully Qualified Domain Names (FQDN).

You can use Auto-Sync in-reverse in the following use cases:

- **Migration**
A planned change of the replication direction.
- **Disaster Recovery**
Changing the roles of source and destination, if source dataset is unavailable. You must enable the `reverse-direction` property in advance.

See Also:

- [Migration Use Case](#)
- [Disaster Recovery Use Case](#)

About the roles of Auto-Sync Service

Auto-Sync service may perform the following functions:

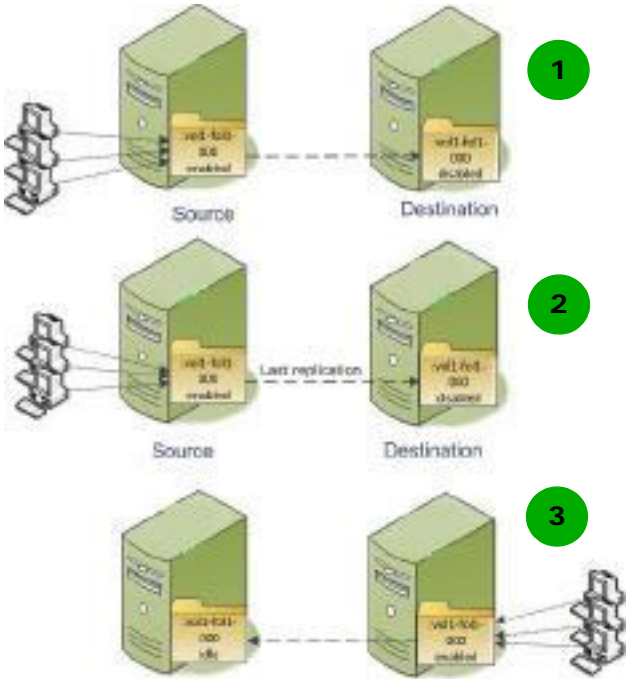
- **Master**
The Auto-Sync service that performs replication from source to destination SGI NAS appliance. By default, any Auto-Sync service that you create is a master service. The name of a master Auto-Sync service appears in blue in the list of the Auto-Sync services.
- **Non-master**
When you execute the `reverse-direction` command, Auto-Sync creates a non-master Auto-Sync service instance on the destination SGI NAS appliance. The name of a non-master Auto-Sync service appears in grey in the list of the Auto-Sync services.

Migration Use Case

In this use case, both source and destination SGI NAS appliances are accessible. SGI NAS Administrator performs a planned migration of an Auto-Sync service instance from the source SGI NAS appliance to destination appliance. You may want to use this functionality during the scheduled maintenance period.

The following diagram describes the migration use case.

Figure 3-1: Migration Use Case



#	Description
1	Normal Auto-Sync work. The Auto-Sync service <code>:vol1-f01-000</code> replicates data from source to destination. The <code>reverse-direction</code> property is enabled. The Auto-Sync service instance <code>:vol1-f01-000</code> is enabled at source and disabled at destination. The Auto-Sync service instance at source performs a <i>master</i> role. The Auto-Sync service instance at destination performs a <i>non-master</i> role.

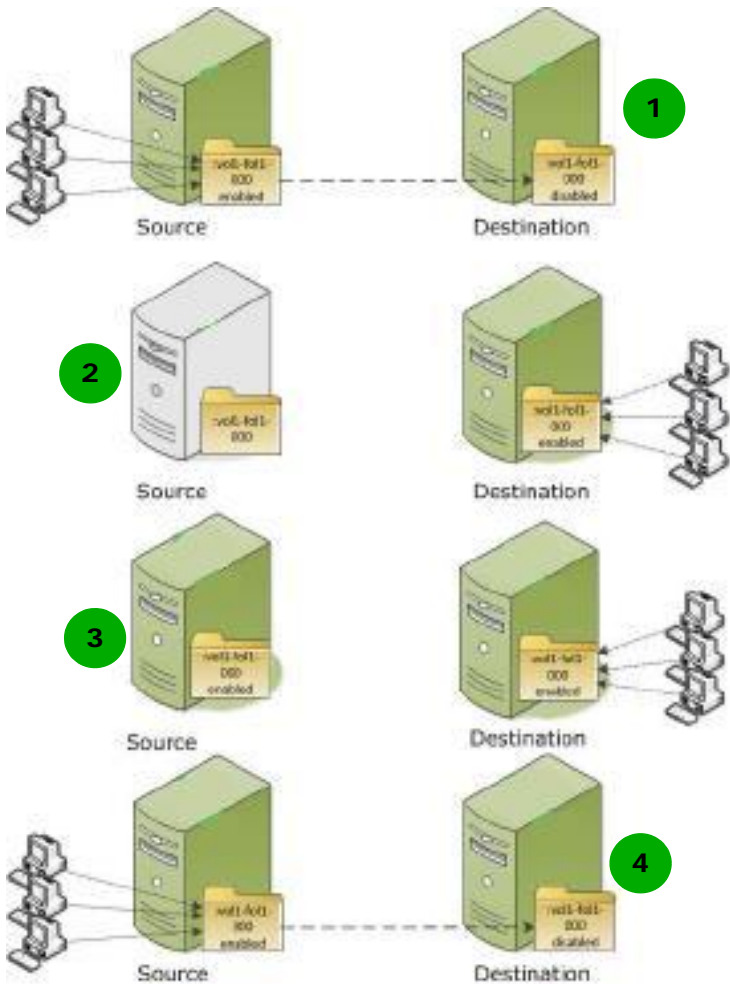
#	Description
2	<p>Migration of the Auto-Sync service.</p> <p>An SGI NAS Administrator executes the <code>reverse-direction</code> command using NMC on source appliance.</p> <p>Auto-Sync performs the following:</p> <ol style="list-style-type: none"> 1. Copies the properties of the Auto-Sync service from source to destination. 2. Changes the <code>read-only</code> property of the source dataset to <code>on</code>. 3. Replicates the latest snapshots.
3	<p>Starting the service at destination.</p> <p>Auto-Sync performs the following:</p> <ol style="list-style-type: none"> 1. Changes the <code>read-only</code> property of the destination dataset to <code>off</code>. 2. Sets the role of the Auto-Sync service instance at source to <code>non-master</code>. 3. Enables the Auto-Sync service instance at destination appliance. 4. Starts the Auto-Sync service instance at the destination appliance according to schedule. <ul style="list-style-type: none"> • Alternatively, you can start the service manually. 5. Sets the role of the Auto-Sync service instance at destination to <code>master</code>.

- [About Auto-Sync In-Reverse](#)
- [Creating an Auto-Sync Service Instance](#)
- [Advanced Settings](#)
- [Reversing the Direction of the Auto-Sync Service](#)
- [Executing the Auto-Sync Service by Administrative Action](#)

Disaster Recovery Use Case

You may want to use the disaster recovery use case if source SGI NAS appliance is not available. You must enable the `reverse-direction` property in advance. Otherwise, you cannot use the disaster recovery capability.

Figure 3-2: Disaster Recovery Use Case



#	Description
1	Normal Auto-Sync work. The Auto-Sync service called <code>:vol1-fol-000</code> replicates data from source to destination. The <code>reverse-direction</code> property is enabled. The service <code>:vol1-fol1-000</code> is enabled on the source appliance and disabled on the destination appliance. The Auto-Sync service instance at source performs a <i>master</i> role. The Auto-Sync service instance at destination performs a <i>non-master</i> role.

#	Description
2	<p>The source SGI NAS appliance is unavailable.</p> <p>Make sure that the source SGI NAS appliance is not pingable from the destination SGI NAS. If the source SGI NAS appliance is pingable, and you execute the <code>reverse-direction</code> command from the destination SGI NAS appliance, you may lose the changes on the source SGI NAS appliance that were accumulated since last successful replication.</p> <p>SGI NAS Administrator executes the <code>reverse-direction</code> command on the destination SGI NAS appliance; then Auto-Sync performs the following:</p> <ol style="list-style-type: none"> 1. Changes the <code>readonly</code> property of the replicated dataset to <code>off</code>. 2. Enables the Auto-Sync service instance at destination. 3. Changes the role of the Auto-Sync service instance at destination to <code>master</code>.
3	<p>The source SGI NAS appliance is back to operation.</p> <p>Both source and destination appliances have <code>readonly</code> property set to <code>off</code>. Both Auto-Sync services are enabled. Both Auto-Sync services perform a role of a <code>master</code>.</p>
4	<p>Restoring the normal Auto-Sync work.</p> <p>SGI NAS Administrator executes the <code>reverse-direction</code> command on the source SGI NAS appliance.</p> <p>You must either wait for the Auto-Sync to execute according to schedule, or start the service manually at source.</p> <p>After the service starts at source, Auto-Sync performs the following:</p> <ol style="list-style-type: none"> 1. Changes the <code>readonly</code> property of the replicated dataset at destination to <code>on</code>. 2. Enables the Auto-Sync service instance at source. 3. Changes the role of the Auto-Sync service instance at source to <code>master</code>.

See Also:

- [About Auto-Sync In-Reverse](#)
- [Advanced Settings](#)
- [Reversing the Direction of the Auto-Sync Service](#)

4

Advanced Settings

The following table provides descriptions of advanced settings that you can use to configure an Auto-Sync service.

Table 4-1: Advanced Settings of the Auto-Sync

Property name in NMV	Property name in NMC	Description
Use copy properties	copy zfs	Copy ZFS properties with replication stream.
Compression level	zip_level	Level of replication stream compression, value range is from 1 to 9. Default value is 0 — no compression.
Rate Limit	rate_limit	Limit of the network traffic in Kb/s that Auto-Sync uses for replication. Default value is 0.
Use deduplication	Deduplicate Replication Stream	Use deduplication on replication stream. Deduplication consumes CPU resources, but lowers the network bandwidth utilization. Deduplication is disabled by default.
	sync-recursive	Synchronize nested folders.
Auto Clone	auto-clone	Automatically clone replicated dataset at destination
Exclude folders	exclude	Exclude folders that match a pattern or alist of patterns fromthe replication stream. Use standard Unix patterning syntax for name filtering. By default, all files and folders are included in the replication stream.
Trace level	trace_level	A level of detail of the messages in the Auto-Sync log file.
Force		Forces replication, even if Auto-Sync identifies change at destination.
Service Retry		If enabled, Auto-Sync automatically restarts the service in case of failure.

Table 4-1: Advanced Settings of the Auto-Sync

Property name in NMV	Property name in NMC	Description
	force: 0 1 2 4 8 16 32	You can set the following force flags in NMC: <ul style="list-style-type: none"> • 0 — Disables all forced operations • 1 — Enables all forced operations, except <code>zfs recv -F</code>. This value is set by default. • 2 — Enables operation of destroy for manually created snapshots at destination • 4— Enables folders destroying at destination • 8 — Enables rollback of the destination folder • 16 — Enables <code>zfs recv -F</code> operation • 32 — Enables using of any snapshots for replication
RR command line options	options	This option provides the possibility to specify a number of TCP connections (-n) and maximum size of package size for replication (-P). Default values are: <ul style="list-style-type: none"> • n=4, which means four TCP connections • P=1024, which means that the maximum PDU size is 1024 kB. SGI recommends to modify these settings only for advanced users.
Mark latest as	latest-suffix	This option enables you to specify a custom suffix for the latest snapshot. Default value is <code>latest</code> .
Auto-Mount Location		Automatically mounts the filesystem that Auto-Sync replicates at destination. Type the absolute UNIX path to the location where you want to mount the destination folder. You cannot modify this property after you create an Auto-Sync service instance. Default value is empty, i.e. auto-mount disabled.
Comment	comment	Optional commentary.
Advanced Properties (available only in NMC)		

Table 4-1: Advanced Settings of the Auto-Sync

Property name in NMV	Property name in NMC	Description
	proto	<p>Replication protocol.</p> <p>Auto-Sync provides the following protocols:</p> <ul style="list-style-type: none"> • zfs+rr <p>Auto-Sync uses native Remote Replication protocol by default. It requires both source and destination to be SGI NAS appliances.</p> <ul style="list-style-type: none"> • zfs+ssh <p>You can use the zfs+ssh protocol to create Auto-Sync service that replicates source dataset located at SGI NAS to destination dataset located at non-SGI NAS appliance.</p> <ul style="list-style-type: none"> • zfs <p>Auto-Sync uses the zfs built-in capabilities to sync two datasets located on the same SGI NAS appliance.</p>
	trunk	<p>Enables you to specify the IP addresses or host names for TCP trunking. This functionality requires more than one network interface located on the source appliance. This property enables Auto-Sync to split the network bandwidth across specified NICs for higher bandwidth throughput.</p> <p>Example:</p> <pre><ssh_bound_host1>/<add_ip1>, <add_ip2> <ssh_bound_host2>/<add_ip3>, <add_ip4></pre>
	from-snapshot	<p>Create Auto-Sync service instance using a snapshot as a source dataset.</p>
	marker_name	<p>Marker name is an identifier that declares the type of snapshot. The default marker name is AutoSync.</p> <p>Example:</p> <pre>datapool/accounts@AutoSync-1_2012-12-04-09-00-11</pre>
	mbuffer_size	<p>The size of the middle buffer. You must enter a value in range of 16-1024.</p>

See Also:

- [Creating an Auto-Sync Service Instance](#)
- [Modifying the Auto-Sync Properties](#)

A—I

Auto-Sync Service Instance

A replication schedule for a selected dataset. By default, the name of the Auto-Sync service instance has the following format:

`dataset_name-N`, where `dataset_name` is the name of the source dataset, and `N` - is the index number of the service.

For example, if you create a first Auto-Sync service instance for folder `pool-docs`, Auto-Sync assigns the following default name: `pool/docs-000`. If you create another replication schedule for this dataset, Auto-Sync assigns the following default name: `pool-docs-001`.

Deduplication

Deduplication eliminates the redundant copies of data and consequently reduces storage requirements. SGI NAS and Auto-Sync uses in-line deduplication. Therefore, you must consider your CPU resources before you enable the deduplication. You must also analyze your data and determine, whether your data gets benefits from deduplication space savings.

Destination Dataset or Destination

A folder or a volume that you use as a destination of the Auto-Sync replication. You can select a dataset on the local SGI NAS appliance, on the remote SGI NAS appliance, or on the remote non-SGI NAS ZFS storage appliance.

Full Backup

Full, complete replica of all of the datasets in the specified source. Provides for a more secure method for backing up data. In case of disk failure, the files are easily restored from a single backup set.

Incremental Backup

Backs up only the changes since the last backup operation. This is less secure than a full backup. In order to restore a file, all of the incremental backups must be present.

J—Q

Marker snapshots

A pair of identical snapshot for the same Auto-Sync service. One marker snapshot is located at source and other at destination. Before every replication iteration, Auto-Sync compares the pair of marker snapshots. Then Auto-Sync sends the difference between marker snapshots.

SGI NAS

SGI NAS is a fully featured NAS/SAN open storage appliance, that leverages the advantages of ZFS.

NMC

The SGI NAS Management Console (NMC) is a command line interface that enables you to execute all SGI NAS functions.

NMS

The SGI NAS Management Server is a service that controls all SGI NAS services and runners. It receives and processes requests from NMC and NMV and returns the output.

NMV

The SGI NAS Management View (NMV) is a web-based graphical User interface that enables you to perform most SGI NAS functions.

R—Z**Recursive**

Recursive is a property of the Auto-Sync service that enables you to run Auto-Sync service in-reverse direction. If you set this property to `on`, then you can execute `reverse-direction` command in the NMC. This functionality is only available in NMC.

Remote Replication (RR)

A native replication protocol for the SGI NAS Auto-Sync. Auto-Sync uses the RR protocol for the replication from local source to remote destination.

Reverse Direction

A function that enables you to swap source and destination of the Auto-Sync service instance.

Snapshot

A read-only copy of the dataset at a particular point in time. You can perform the following actions with a snapshot: clone, rollback a folder to a snapshot, access the data.

Source Dataset or Source

A folder, a volume, a zvol, or a snapshot that you want to use a source of the replication. You can select a dataset on the local or remote SGI NAS appliance.

Trunking

The port trunking or link aggregation enables you to set up multiple network connections for a replication stream. This helps to increase the speed of the Auto-Sync replication.

VMware ESXi

An enterprise class hypervisor that provides a software virtualization environment.

VMware ESXi Cluster

A collection of two or more ESXi hosts. In the NV4V environment, you use the ESXi cluster to load balance and better utilize any resources.

VMware vCenter Server

It is the centralized management tool for the vSphere suite. VMware vCenter Server enables you to manage multiple ESX servers and virtual machines (VMs) from different ESXi servers through a single console application.

VMware vSphere client

A Microsoft Windows desktop application that enables you to access VMware ESXi and VMware vCenter.

ZFS

Zettabyte File System (ZFS) is 128-bit file system that provides features, such as data integrity verification, disk management, snapshots, transactional operations, and so on.

ZFS send/receive

ZFS send/receive is the main working mechanism of the Auto-Sync replication. The `zfs send` command creates a snapshot of a files sytem and then streams it to `zfs receive` command. The `zfs receive` command creates a new file system out of this replication stream. ZFS

ZIL

ZFS Intent Log is a component of a hybrid storage pool that speeds up write operation. Usually, SSD drives are used as ZIL devices.

Zvol

Interface layers available in the ZFS