



Auto-Sync

User Guide

3.1.6

Copyright © 2014 Nexenta Systems, ALL RIGHTS RESERVED

Notice: No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose, without the express written permission of Nexenta(hereinafter referred to as "Nexenta").

Nexenta reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. Nexenta products and services only can be ordered under the terms and conditions of Nexenta' applicable agreements. All of the features described in this document may not be available currently. Refer to the latest product announcement or contact your local Nexenta sales office for information on feature and product availability. This document includes the latest information available at the time of publication.

Nexenta is a registered trademark of Nexenta in the United States and other countries. All other trademarks, service marks, and company names in this document are properties of their respective owners.

Product	Versions supported
NexentaStor™	3.1.6
Auto-Sync	3.1.6

Contents

Preface	v
1 Introduction	1
About Auto-Sync	1
About Auto-Sync Replication Algorithm	2
About Incremental Snapshots	2
About Transport Protocols	3
About the Directions of Replication	3
About Replicating Content	4
About Changes at Destination	5
2 Managing Auto-Sync	7
Before You Start to Use Auto-Sync	8
Viewing Summary Information About an Auto-Sync Service	8
Creating an Auto-Sync Service Instance	8
Establishing the SSH-binding for an Auto-Sync Service	14
Modifying the Auto-Sync Properties	15
Enabling an Auto-Sync Service Instance	15
Disabling an Auto-Sync Service Instance	16
Deleting an Auto-Sync Service Instance	16
Deleting Snapshots Created by Auto-Sync	17
Resuming an Auto-Sync Service	18
Executing an Auto-Sync Service by Administrative Action	18
Running Auto-Sync Services in Series	19
Running Auto-Sync Services in Parallel	19
Stopping an Auto-Sync Service	19
Restarting an Auto-Sync Service	20
Enabling the Reversing of an Auto-Sync Service	20
Recovering a Folder	21
Recovering Multiple Folders	22
Reversing an Auto-Sync Service	22

Mounting an Auto-Sync Destination Folder	23
Viewing the Auto-Sync Log File	24
Replicating from a Snapshot	25
Viewing the Auto-Sync Statistics	25
3 Advanced Configuration	27
Creating an Auto-Sync Service Instance with Advanced Options	27
Changing Paths to Source or Destination Dataset	29
About the Auto-Sync Reverse-Service	30
About Roles of the Auto-Sync Service	31
Migration Use Case	31
Disaster Recovery Use Case	34
4 Advanced Settings	37
Glossary	41

Preface

This documentation presents information specific to Nexenta 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.

Documentation History

The following table lists the released revisions of this documentation.

Table 1: Documentation Revision History

Revision	Date	Description
3000-at_syn_v3.1.6-000046-A	June, 2014	GA

Contacting Support

Choose a method for contacting support:

- Visit the Nexenta customer portal <http://nexenta.force.com/customerportal> or partner portal <http://nexenta.force.com/partnerportal>. Log in and browse a knowledge base.
- Using the NexentaStor user interface, NMV (Nexenta Management View):
 - a. Click **Support**.
 - b. Complete the request form.
 - c. Click **Send Request**.
- Using the NexentaStor command line, NMC (Nexenta Management Console):
 - a. At the command line, type `support`.
 - b. Complete the support wizard.

Comments

Your comments and suggestions to improve this documentation are greatly appreciated. Send any feedback to doc.comments@nexenta.com and include the documentation title, number, and revision. Refer to specific pages, sections, and paragraphs whenever possible.

Introduction

This section includes the following topics:

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

About Auto-Sync

Auto-sync is a NexentaStor data and metadata replication service that combines 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 a local or remote destination. You can use these snapshots as a backup or an archive, as well as in the disaster recovery use case.

Auto-sync provides the following features:

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

See Also:

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

About Auto-Sync Replication Algorithm

When the Auto-Sync service instance runs for the first time, 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. On every run, it creates a snapshot of the selected dataset. You can also run an Auto-Sync service manually. Auto-Sync uses the following naming pattern for snapshots:

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

2. Determining the latest identical snapshots at source and destination

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

3. Sending delta from source to destination

Auto-Sync sends changes from source snapshot to destination snapshot.

See Also:

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

About Incremental Snapshots

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

See Also:

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

- [About 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 by Nexenta. The protocol uses native ZFS `send/receive` combined with the `rrdaemon` service. Auto-Sync uses Remote Replication protocol for local-to-remote replication.

If you use the RR protocol, both source and destination, and remote to local appliances must be the NexentaStor appliances.

- **ZFS + SSH**

Another replication option for local-to-remote and remote-to-local replication. The ZFS+SSH protocol is more secure comparing to the RR protocol. However, the replication speed is slower.

- **ZFS**

Auto-sync uses built-in ZFS functionality for replication from one local folder to another.

About the Directions of Replication

You can create an Auto-Sync service that replicates data to one of the following directions:

- **Locally (L2L)**

Both source and replication datasets reside on local NexentaStor appliance.

- **From Host (L2R)**

The source dataset resides on local NexentaStor appliance. The destination dataset resides on a remote NexentaStor appliance.

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

You can change the direction of replication after you create an Auto-Sync service instance.

- **To Host (R2L)**

Source dataset is stored on a remote NexentaStor appliance. Destination dataset resides on local NexentaStor appliance.

You can change the direction of replication after you create an Auto-Sync service instance.

See Also:

- [Reversing an 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 the enabled and disabled replicate content (`dircontent`) property.

Figure 1-1: Replication with Replicate Content Disabled

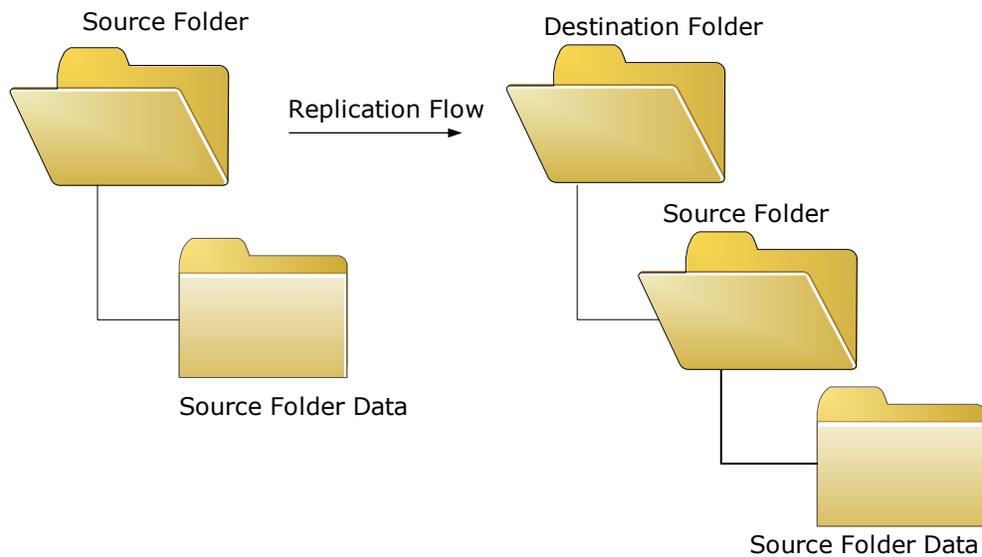
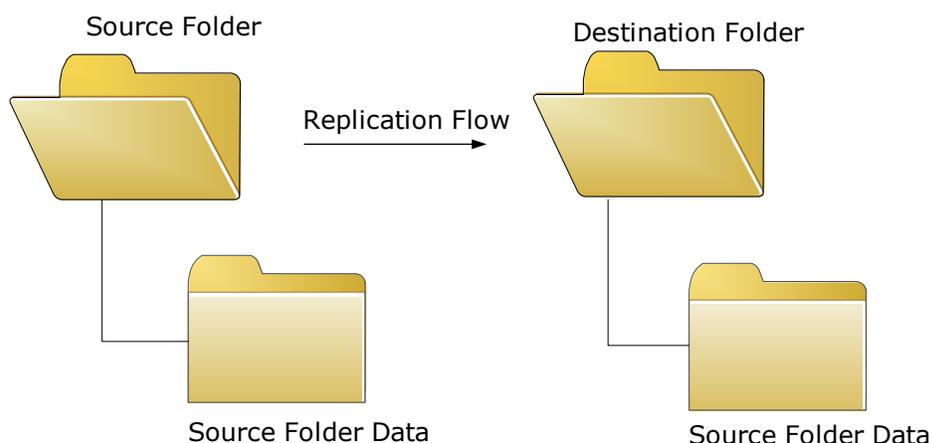


Figure 1-2: Replication with Replicate Content Enabled**See Also:**

- [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 at destination. If you mount the destination folder and change any data, Auto-Sync identifies the changes as a loss of synchronization between source and destination. Auto-Sync processes this condition as a recoverable error and attempts to fix 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 changes to it, disable or delete any Auto-Sync service instances that you have earlier created for this folder.

See Also:

- [About Auto-Sync Replication Algorithm](#)
- [About Replicating Content](#)
- [Creating an Auto-Sync Service Instance](#)
- [Advanced Settings](#)

This page intentionally left blank

Managing Auto-Sync

This section includes the following topics:

- [Before You Start to Use Auto-Sync](#)
- [Viewing Summary Information About an Auto-Sync Service](#)
- [Creating an Auto-Sync Service Instance](#)
- [Establishing the SSH-binding for an Auto-Sync Service](#)
- [Modifying the Auto-Sync Properties](#)
- [Enabling an Auto-Sync Service Instance](#)
- [Disabling an Auto-Sync Service Instance](#)
- [Deleting an Auto-Sync Service Instance](#)
- [Deleting Snapshots Created by Auto-Sync](#)
- [Resuming an Auto-Sync Service](#)
- [Executing an Auto-Sync Service by Administrative Action](#)
- [Running Auto-Sync Services in Series](#)
- [Running Auto-Sync Services in Parallel](#)
- [Stopping an Auto-Sync Service](#)
- [Restarting an Auto-Sync Service](#)
- [Recovering a Folder](#)
- [Reversing an Auto-Sync Service](#)
- [Viewing the Auto-Sync Log File](#)
- [Replicating from a Snapshot](#)
- [Viewing the Auto-Sync Statistics](#)

Before You Start to Use Auto-Sync

Auto-Sync is included with the NexentaStor Enterprise Edition and NexentaStor Trial Version. Before you start to use Auto-Sync, verify that plugin is installed. Go to **Settings > Appliance > Plugin** page and find `autosync` in the list of installed plugins. If the plugin is not in that list, install the plugin from the repository.

See Also:

- *NexentaStor Installation Guide*, Section *Accessing the Plugins*
- *NexentaStor Installation Guide*, Section *Installing Plugins*

Viewing Summary Information About an Auto-Sync Service

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

❖ *To view 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 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 specified 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.

If you create an Auto-Sync service from this appliance to a remote appliance, you must first establish the ssh-binding.

For more information, see [Establishing the SSH-binding for an Auto-Sync Service](#).

- ❖ To create an 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 an Auto-Sync service instance. If you do not specify a custom name, Auto-Sync assigns the name automatically.
Direction	Auto-Sync can execute a service in the following directions: <ul style="list-style-type: none"> • Locally (L2L) • To Remote Host (L2R) • From Remote Host (R2L)
Periodic Interval	Periodic Interval consists of two parameters: Period and Frequency. Frequency depends on Period. For example, if you set Period to <i>hourly</i> and Frequency to <i>4</i> , then Auto-Sync runs the service instance every 4 hours. If you set Period to <i>daily</i> and frequency to <i>5</i> , then Auto-Sync runs the service instance every 5 days, and so on.
Replicate Content	Replicate content of a folder. If you leave this property unselected, Auto-Sync replicates parent folder as well. For more information, see: About Replicating Content .
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 the retention period expires. The next time the service starts, Auto-Sync deletes these snapshots.
Keep on Destination	The number of snapshots that you want to keep on the destination dataset. Alternatively, specify a period in YDHMS format. Auto-Sync keeps three additional snapshots after the retention period expires. The next time the service starts, Auto-Sync deletes these snapshots.
Local Source Folder/Zvol	A dataset on the NexentaStor 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 selected the local-to-remote replication.</p> <p>Multi-destination enables you to set a few destinations for the Auto-Sync service instance. 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 NexentaStor appliance that you want to use as a replication destination. ● Keep on destination A 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.
Local destination folder/zvol	A dataset on the NexentaStor appliance that you want to use as a destination for replication.
Schedule	<p>Schedule consists of two parameters: Period and Frequency. Frequency depends on Period.</p> <p>For example, if you set Period to <i>hourly</i> and Frequency to <i>4</i>, then Auto-Sync runs the service instance every 4 hours. If you set Period to <i>daily</i> and frequency to <i>5</i>, then Auto-Sync runs the service instance every 5 days, and so on.</p>
Transport Protocol	<p>Replication protocol.</p> <p>Auto-Sync provides the following protocols:</p> <ul style="list-style-type: none"> ● zfs+rr Auto-Sync uses Remote Replication protocol for local-to-local and local-to-remote replication by default. ● zfs+ssh You can use the ZFS over ssh protocol for local-to-local and local-to-remote replication. The zfs+ssh protocol is more secure comparing to RR. However, the replication performance is slower comparing to RR protocol. ● zfs Auto-Sync uses the ZFS built-in capabilities to sync two datasets located on the same NexentaStor appliance. For more information, see About Transport Protocols.

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

Parameter in NMV	Description
Remote Destination Host	Only for local to remote replication. IP address or host name of the remote host to which you plan to replicate the data. The remote host must be bind by SSH.
Destination folder	Local or remote destination folder. For L2R replication, you can assign multiple destination folders. Multi-destination includes the following parameters: <ul style="list-style-type: none"> ● Host The IP address or host name of the remote destination appliance. ● Folder A dataset on the remote NexentaStor appliance that you want to use as a replication destination. ● Keep on destination A 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.
Compression	Specifies the inline compression algorithm for the replication stream.
	Level
	Specifies compression level. The options are: <ul style="list-style-type: none"> ● 0 to 9 level 9 - Best compression

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 an Auto-Sync service instance, using NMC:*

1. Log in to NMC using SSH-client.

2. Type the root credentials.

3. Type:

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

4. Select a replication interval.

The options are:

- minute
- hourly

- daily
- weekly
- monthly



Set the replication interval to a value greater than 10 minutes.

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



If a source dataset is located on a remote NexentaStor appliance, type the IP address.

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



Select `./*`, if you want to replicate the content of a selected dataset, but not the dataset itself.
For more information, see [About Replicating Content](#)

8. Select a destination folder.
 - a. If the source dataset resides on this NexentaStor 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 NexentaStor appliance; then select a remote dataset.
-



If a destination dataset resides on a remote NexentaStor appliance, you may be prompted to type the super-user password for the remote NexentaStor appliance to establish the SSH binding.

- c. If a source dataset resides on a remote appliance, select a local dataset that you want to use as replication destination.
9. Specify a number of snapshots to keep at source.

Alternatively, specify a period in YDHMS format.

- 10.** Specify a number of snapshots to keep at destination.

Alternatively specify a period in YDHMS format.

- 11.** If you selected to replicate from the local to remote NexentaStor appliance, you can specify more than one destination.

Would you like to specify additional destination? *y*

- a.** Select a destination host.
- b.** Select a destination folder.
- c.** Specify a number of snapshots to keep at destination.

Alternatively, specify a period in YDHMS format. You must set this parameter for every destination.

NexentaStor verifies configuration of the service.



Before you configure advanced settings, read [Advanced Settings](#). Nexenta recommends that you keep the default settings.

- 12.** Optionally, modify advanced settings.

- a.** In Copy zfs, type yes or no.
- b.** In Rate limit, type the maximum allowed speed in KB.
- c.** Optionally, specify a custom name for the Auto-Sync service instance.
- d.** Optionally, type a commentary.

- 13.** Complete the wizard by typing *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'
```



When you export a volume, NexentaStor destroys all Auto-Sync services created for this volume.

See Also:

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

- [Executing an Auto-Sync Service by Administrative Action](#)

Establishing the SSH-binding for an Auto-Sync Service

SSH-binding, aka SSH Public Key Based Authentication, is an easy way to make two NexentaStor appliances communicate through secure connection. When you create an ssh-binding the public key from the remote NexentaStor appliance is written to the authorized key file on the local appliance.

You need to create the ssh-binding for:

- Local-to-remote replication
- Remote-to-local Auto-Sync replication
- Replication between NexentaStor appliance and the HA Cluster group

If you create the SSH-binding to replicate data from or to an HA Cluster group, see [Creating an Auto-Sync Service in the HA Cluster Environment](#).

❖ *To establish the SSH-binding for Auto-Sync service, using NMV:*

1. Click **Settings > Network**.
2. In the **Network** panel, click **SSH-Bind**.
3. In the **Remote Server** field, type the IP address of the remote NexentaStor appliance.
4. In the **Remote User** field, type the name of the user with root credentials.
5. Type the remote super user password.

The new SSH-binding displays in the **Summary Network Settings** window, in the **SSH Bound Hosts** list.

❖ *To establish an SSH-binding for Auto-Sync service, using NMC:*

1. Type:

```
nmc:/$ setup network ssh-bind
```
2. Type the IP address or host name of a remote NexentaStor appliance.
3. Type the remote appliance Super User password.
4. Repeat [Step 1](#) - [Step 3](#) on the remote node.



If ssh-binding fails, you can manually configure the `/etc/hosts/` file, which contains the Internet host table.

Type `setup appliance hosts` to access the file.)

Modifying the Auto-Sync Properties

After you create an 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. Modify properties.
5. 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 changes. Type `y`.
4. Alternatively, discard changes by typing `n`.

See Also:

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

Enabling an Auto-Sync Service Instance

You can enable an Auto-Sync service instance if it was previously disabled by system 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 a checkbox near the service that you want to enable.
4. Click **Enable**.

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

- ◆ Type:

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

See Also:

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

Disabling an Auto-Sync Service Instance

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

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

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

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

1. Type:

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

See Also:

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

Deleting an Auto-Sync Service Instance

You can delete an Auto-Sync service instance to stop synchronizing source and destination datasets.

If you destroy an Auto-Sync service instance, you can later create a new service with the same source and destination. Auto-Sync will resume the service using the last snapshot created before the deletion.

For more information, see [Resuming an Auto-Sync Service](#).



If you destroy an Auto-Sync service instance, Auto-Sync does not delete the snapshots that it has created. However, you can specify cleanup options in the Auto-Sync destroy dialog.

❖ *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. Click on the Auto-Sync service that you want to delete.
4. Click **Delete**.
5. Alternatively, you can delete one or multiple Auto-Sync services from the **Data Management > Auto-Services** page by selecting the services and clicking **Delete selected**.

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

1. Type:

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

See Also:

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

Deleting Snapshots Created by Auto-Sync

You may want to delete the snapshots created by Auto-Sync 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 checkboxes near the snapshots that you want to delete.
6. Alternatively, select a 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/zvol <folder>/<zvol> destroy-snapshots
-n <template>
```

Example:

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

```
nmc:/$ setup zvol datapool/info 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 by typing **y**.

Resuming an Auto-Sync Service

If you delete an Auto-Sync service, you can later create a new service with the same source and destination. Auto-Sync resumes the replication using the latest snapshot created by the deleted service.

❖ *To resume an Auto-Sync service, using NMC:*

1. Type:

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

2. Create a new Auto-Sync service specifying the same source and destination folders as in the deleted service. Follow the steps described in [Creating an Auto-Sync Service Instance](#).

Executing an Auto-Sync Service by Administrative Action

You can run an existing Auto-Sync service instance any time without waiting for the scheduled time. You may need this functionality to immediately synchronize source and destination datasets of an Auto-Sync service.

❖ *To execute an 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 check box near the service that you want to start.

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

Running Auto-Sync Services in Series

Multiple Auto-Sync services that run at the same time, may impact performance of your NexentaStor appliance. If you have multiple Auto-Sync services scheduled to run for the same time, you may want to configure them to execute one after another.

This functionality is only available in NMC.

❖ *To run 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 Auto-Sync Services in Parallel](#)

Running Auto-Sync Services in Parallel

You can disable execution of Auto-Sync services running in series and configure them to run in parallel. It may decrease performance of your NexentaStor appliance.

By default, Auto-Sync services run in parallel. You can stop an Auto-Sync service instance in case of emergency.

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 unserialized  
<service instance> has been unserialized
```

See Also:

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

Stopping an Auto-Sync Service

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

This functionality is only available in NMC.

- ❖ *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 the 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 Service Instance](#)
- [Deleting an Auto-Sync Service Instance](#)
- [Enabling an Auto-Sync Service Instance](#)

Restarting an Auto-Sync Service

If an Auto-Sync destination becomes unavailable during replication, Auto-Sync locks the source folder. Therefore, when the destination folder becomes available, the Auto-Sync services that use this folder as destination may not recover automatically. You can recover these Auto-Sync services using the `restart` command.

- ❖ *To restart an Auto-Sync service, using NMC:*

1. Type:

```
nmc:/$ setup auto-sync <name> restart
```

System response:

```
Warning: restart causes abnormal service termination. Do you want to continue? (y/n)
```

2. Type `y`.

System response:

```
Yes  
Restarted 1 auto-sync service instance
```

Enabling the Reversing of an Auto-Sync Service

To use the `reverse-direction` command, you must enable the `reverse-direction` property for the selected service. When you enable the `reverse-direction` property, Auto-Sync creates a copy of the selected Auto-Sync service instance on the destination NexentaStor appliance.

Before using the `reverse-direction` functionality, verify that your environment meets the following prerequisites:

- TCP and SSH connectivity in the opposite direction.
- The remote host is connected using SSH by using its Fully Qualified Domain Name (FQDN).
- The NexentaStor appliance on which this Auto-Sync service instance runs is not protected by a firewall.
- All IP addresses used for the Auto-Sync replication have associated hostname(s).



Note:

Nexenta recommends that you use the `reverse-direction` command for advanced configurations.

❖ *To enable the reversing of an Auto-Sync service, using NMC:*

1. Type:

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

System response:

...

```
Reverse-direction capability is currently disabled.
Do you want to activate 'reverse-direction'?
```

2. Type `y`.

Recovering a Folder

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

This functionality is only available in NMC.

❖ *To recover a folder using NMC:*

◆ Type:

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

System response:

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

❖ *To make a destination folder available for read/write operations even if the folder is not managed by auto-sync service:*

◆ Type:

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

Example:

❖ *For one folder:*

◆ Type:

```
nmc:/$ setup auto-sync recover-folder -n backup/to/www
```

Recovering Multiple Folders

You can recover multiple folders using a single command.

This functionality is only available in NMC.

❖ *To recover multiple folders using NMC:*

◆ Type:

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

Reversing an Auto-Sync Service

Auto-Sync can reverse an Auto-Sync service. Therefore, source and destination can change the roles.

When you reverse a service, Auto-Sync performs the following tasks:

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



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



If you use multipathing for an Auto-Sync service, it will be disabled when you reverse the Auto-Sync service.

❖ *To reverse an Auto-Sync service, using NMC:*

1. Type:

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

- If the `reverse-direction` property is disabled, NMC displays the following message:

```
reverse-direction feature is currently disabled  
Do you want to enable 'reverse-direction' feature?
```

a. Type `y`.

System response:

```
'Reverse-direction' feature successfully enabled
```

- b.** Type the reverse-direction command 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 Service Instance](#)
- [Advanced Configuration](#)
- [About the Auto-Sync Reverse-Service](#)
- [Migration Use Case](#)
- [Disaster Recovery Use Case](#)

See Also:

- [Reversing an Auto-Sync Service](#)

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.



Nexenta 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 an 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. Select an Auto-Sync service instance.
4. In the **Update Auto-Sync Service: <service instance>** pane, click **View Log**.
Scroll down to view the required event.
5. Optionally, in the **Lines to show** field, adjust the number of lines that you want to view.

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

1. Type:

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

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

2. Alternatively, type:

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

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

3. To display the output of the Auto-Sync log file in real time, type:

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

See Also:

- [Advanced Settings](#)

Replicating from a Snapshot

You can specify a snapshot that you want to use as the initial snapshot for the Auto-Sync replication. Therefore, you can resume the replication from this snapshot.

This functionality is only available in NMC.

❖ *To replicate from a snapshot using NMC:*

1. Type:

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

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

You must specify 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 promptly correct system errors.



Note:

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

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

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

View the progress bar.

SHOW SYNCING SERVICES										
Instance	Direction	Progress	Speed	Time	Snapshots	Frequency	Last run	Duration	View log	
<input type="checkbox"/> info-lists-000		Task 1 of 1; 288.00 MB of 420.75 MB	49.23 MB/s	00:00:02	1	weekly	18:46:06, Feb04	22s		
<div style="display: flex; justify-content: space-between; gap: 10px;"> Enable Disable Delete Start Abort </div>										

❖ To view the Auto-Sync statistics, using NMC:

1. Type:

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

System response:

TCP CONNECTIONS	SNEXT	RNEXT	TRANSFER
10.3.60.38.40652-10.3.60.56.2001	164900682	283889218	-
10.3.60.38.40652-10.3.60.56.2001	164900682	283889218	0 B
10.3.60.38.40652-10.3.60.56.2001	164900682	283889218	0 B
10.3.60.38.40652-10.3.60.56.2001	164900682	283889218	0 B

See Also:

- [Advanced Settings](#)
- [Creating an Auto-Sync Service Instance](#)
- [Modifying the Auto-Sync Properties](#)
- [Stopping an 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](#)
- [About the Auto-Sync Reverse-Service](#)
 - [Migration Use Case](#)
 - [Disaster Recovery Use Case](#)



Warning:

This chapter describes advanced configuration that you configure for Auto-Sync services. Nexenta recommends that you use this functionality if you are an experienced NexentaStor user. Nexenta does not recommend to use these settings if use Auto-Sync for the first time.

Creating an Auto-Sync Service Instance with Advanced Options

You can create an Auto-Sync service instance using 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 a volume, a folder, a zvol, or a snapshot. Example: <code>nmc:/\$ setup auto-sync create -S <snapshot></code>
-d	Destination dataset. Specify a volume or a folder.

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

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

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

Option	Description
-u <custom name>	A custom name for the Auto-Sync service instance. Auto-Sync uses the following naming pattern: dataset-index_number. Example: datapool/docs-000 Where datapool/docs is the name of the dataset and 000 is the index number of an Auto-Sync service.
-I <comment>	Descriptive commentary.
-r	Replicate the dataset recursively. Auto-Sync ignores this option if the source dataset is a zvol.
-N	Replicate the dataset non-recursively.
-g	Create an Auto-Sync service instance with the enabled reverse-direction property.
-Z	If you specify this option, then Auto-Sync automatically unmaps the source 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 parameters, IP addresses or host names, for TCP trunking. This functionality requires more than one network interface to be present at source. This property enables Auto-Sync to distribute network traffic across specified NICs for higher bandwidth throughput.
-F	Generate fault notifications about actions of the Auto-Sync service instance. This option enables you to identify, track, and resolve any issues that affect the Auto-Sync normal operation. Auto-Sync also sends a fault report by e-mail if mailer is properly configured.
-O	Destroys the source and destination snapshots non-recursively. This option prevents recursive zfs operations. If you enable this option then zfs destroy operation performs directly on each nested dataset.

Changing Paths to Source or Destination Dataset

You can change the paths to source or destination datasets if you changed the location of these datasets, modified the IP address of the source or destination host, or applied 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 a path to new location of the source dataset.

4. In the `To host` field, type the name of a new destination host.

5. In the `To fs` field, type the path to a new location of a destination dataset.

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

About the Auto-Sync Reverse-Service

During normal operation, Auto-Sync replicates data from source to destination folder. For maintenance or migration purposes, you may want to change the direction of replication and switch the roles of the source and destination appliances. This eliminates the need to re-create the affected Auto-Sync service.

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

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

You can change the direction of replication only if you configure local-to-remote or remote-to-local Auto-Sync services.



Note:

To enable this functionality, bind local and remote appliances using 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 for maintenance purposes.

- **Disaster Recovery**

Changing the roles of source and destination when source dataset is unavailable. You must enable the `reverse-direction` property in advance.

See Also:

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

About Roles of the Auto-Sync Service

Auto-Sync service may act as:

- **Master**

An Auto-Sync service that replicates data from source to destination NexentaStor appliance. By default, all Auto-Sync service instances that you create are *master* services. The name of a master Auto-Sync service appears in green in the list of Auto-Sync services on the **Auto-Services > Show** page.

See [List of Master and Non-master Auto-Sync Services at Source](#).

- **Non-master**

When you execute the `reverse-direction` command, Auto-Sync creates a *non-master* Auto-Sync service instance at destination. The name of a non-master Auto-Sync service appears in gray in the list of Auto-Sync services.

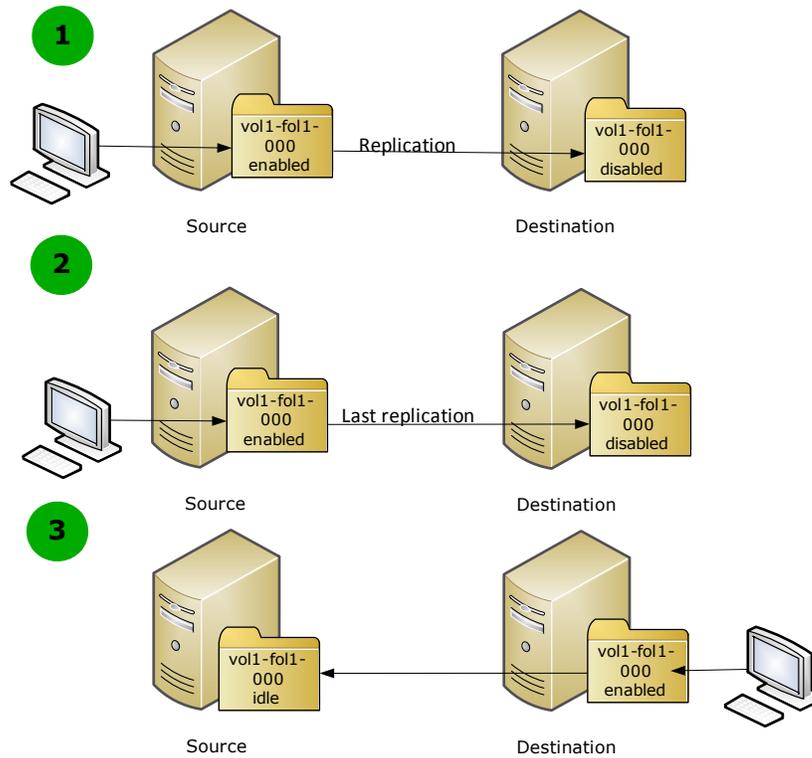
See [List of Master and Non-master Auto-Sync Services at Destination](#).

Migration Use Case

When you integrate new hardware into your IT infrastructure, you may need to migrate your data from one NexentaStor appliance to another. To migrate the data, execute the steps described in this section. Perform the migration during a scheduled maintenance period to minimize the service downtime.

The following diagram describes the migration use case.

Figure 3-1: Migration Use Case



#	Description
1	<p>Normal Auto-Sync operation.</p> <p>The Auto-Sync service <code>:vol1-fo11-000</code> replicates data from source to destination. The <code>reverse-direction</code> property is enabled. The Auto-Sync service instance <code>:vol1-fo11-000</code> is enabled at source and disabled at destination. The Auto-Sync service instance at source acts as <i>master</i>. An Auto-Sync service instance at destination acts as <i>non-master</i>.</p>
2	<p>Migration of the Auto-Sync service.</p> <p>Execute the <code>reverse-direction</code> command using NMC at source. Auto-Sync completes the following tasks:</p> <ol style="list-style-type: none"> 1. Copies 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.

#	Description
3	Starting the service at destination. Auto-Sync completes the following tasks: <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 destination 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>.

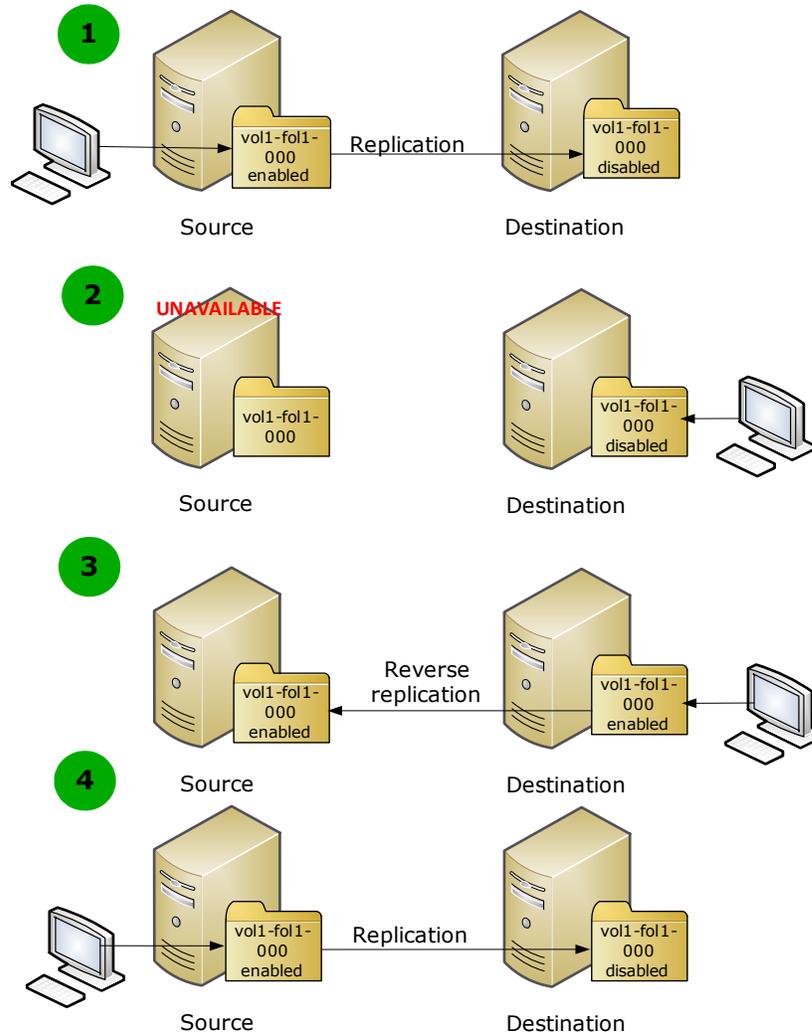
See Also:

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

Disaster Recovery Use Case

In the event of a system failure on the source NexentaStor appliance, follow the steps described in this section, so network clients can access the data at destination. After you recover the source appliance, you can switch the service roles back to normal operation. Ensure that the `reverse-direction` property is enabled in advance, so you can use the reverse capability in case a system failure occurs.

Figure 3-2: Disaster Recovery Use Case



#	Description
1	<p>Normal Auto-Sync operation.</p> <p>The Auto-Sync service <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 at source and disabled at destination. The Auto-Sync service instance at source acts as <i>master</i>. The Auto-Sync service instance at destination acts as <i>non-master</i>.</p>
2	<p>The source NexentaStor appliance becomes unavailable.</p> <ol style="list-style-type: none"> 1. Verify that the source NexentaStor appliance is not pingable from the destination NexentaStor. If the source appliance is available and you execute the <code>reverse-direction</code> command at destination, you may lose the changes that were accumulated after the last successful replication. 2. Run the <code>reverse-direction</code> command at destination. Auto-Sync performs the following: <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 <i>master</i>.
3	<p>The source NexentaStor 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 act as <i>master</i>.</p>
4	<p>Restoring normal Auto-Sync operation.</p> <ol style="list-style-type: none"> 1. Run the <code>reverse-direction</code> command at source. 2. Wait for the Auto-Sync service to run according to schedule or start the service manually at source. 3. After the service starts at source, Auto-Sync completes the following tasks: <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 <i>master</i>.

See Also:

- [About the Auto-Sync Reverse-Service](#)
- [Advanced Settings](#)
- [Reversing an Auto-Sync Service](#)

This page intentionally left blank

Advanced Settings

The following table describes advanced settings for Auto-Sync services.

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

Property name in NMV	Property name in NMC	Description
Misc. ZFS Flags	general-flags	You can set up different flags for the Auto-Sync service.
Do not Copy ZFS properties	do_not_copy_zfs_properties	Excludes ZFS properties from the replication.
	unmap_zvols	Automatically unmaps zvols during the reverse replication.
	serialized	Run services in serial. All auto_sync services run one after another.
Non-recursive replication	non_recursive_replication	Enables you to replicate dataset non_recursively without nested folders.
	source_is_rootfs	Uses root filesystem as source dataset.
	dont_mount_on_destroy	Do not mount destination on service destroy operation.
	daemon	Run the Auto-Sync service in background.
	schedule_once	Run the Auto-Sync service once according to schedule, then change the service state to maintenance.
Use deduplication	Deduplicate Replication Stream	Deduplicate the replication stream. Deduplication consumes additional CPU resources. However, it reduces bandwidth usage. Deduplication works if both the zfs systems (appliances) support this feature.
Use copy properties	copy_zfs	Copy ZFS properties to destination dataset.
Trace level	trace_level	A level of detail of messages in the Auto-Sync log file.
Rate Limit	rate_limit	Limits transfer rate of auto sync services. Specify rate_limit in Kb/s that auto-sync uses for replication. Default value is 0.

Table 4-1: Advanced Settings of the Auto-Sync (Continued)

Property name in NMV	Property name in NMC	Description
Mark latest as	latest-suffix	A custom suffix for the latest snapshot. Default value is <code>latest</code> .
Force Flags	force	You set flags that enforce some Auto-Sync operations.
	destroy_any_snapshots	Enables you to destroy snapshots manually created at destination.
	destroy_folders	Enables you to destroy folders at destination.
	rollback	Enables you to rollback the destination folder. You may want to enable this property to synchronize source and destination datasets if changes were applied to destination.
	force_receive	Force a rollback of the file system to the most recent snapshot before performing the receive operation.
	destroy_snapshots	Destroy manually created snapshots at destination.
	use_any_snapshots	Enables you to use any snapshots for replication including manually created ones.
Service Retry	retry	If enabled, Auto-Sync automatically restarts the service in case of failure.
	non_recursive_zfs_destroy	Auto-Sync stores snapshots at source and destination for the period defined by retention policy. When the period expires, Auto-Sync deletes all snapshots created for this Auto-Sync service at the same time. If the source dataset includes a multi-layer hierarchy of nested folders, deletion of the snapshots created for this dataset may significantly affect performance. Enabling the <code>non_recursive_zfs_destroy</code> flag enforces Auto-Sync to delete snapshots one after another which decreases the impact on the NexentaStor performance. You may want to use this flag if the replicated dataset contains multiple nested folders.
Exclude folders	exclude	Do not replicate the folders that match a pattern or a list of patterns. Use standard Unix patterning syntax for name filtering. By default, Auto-Sync replicates all files and folders.
Compression level	zip_level	Level of replication stream compression. Value range is from 1 to 9. Default value is 0 — no compression.

Table 4-1: Advanced Settings of the Auto-Sync (Continued)

Property name in NMV	Property name in NMC	Description
Middle Buffer Size	<code>mbuffer_size</code>	The size of middle buffer. Type a value in range of 16-1024 KB. Middle buffer is memory used by RR daemon to fasten the data transfer. Nexenta recommends to use default setting for this property.
Replicate recursively	<code>sync-recursive</code>	Includes nested folders and their subfolder into replication.
Auto Clone	<code>auto-clone</code>	Automatically clone replicated dataset at destination.
RR command line options	<code>options</code>	Command line options for Remote Replication protocol include: <ul style="list-style-type: none"> • <code>-n</code> — number of TCP connections. Default value is 4. • <code>-P</code> — maximum size of a network packet for replication (PDU). Default value is 1024 kB. Changing the RR command line options is not recommended unless you have sufficient information about the values to provide for these settings.
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	<code>comment</code>	Descriptive commentary.
Advanced Properties (available in NMC only)		
	<code>trunk</code>	Enables you to specify IP addresses or host names for TCP trunking. This functionality requires more than one network interface present on the source and destination appliance. This property enables Auto-Sync to distribute the network workload across specified NICs for higher bandwidth throughput. Example: <code><ssh_bound_host1>/<add_ip1>, <add_ip2> <ssh_bound_host2>/<add_ip3>, <add_ip4></code>
	<code>from-snapshot</code>	Create Auto-Sync service instance using a snapshot as a source dataset.

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 redundant copies of data and consequently reduces storage requirements. NexentaStor 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 NexentaStor appliance, on the remote NexentaStor 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.

NexentaStor

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

NMC

The Nexenta Management Console (NMC) is a command line interface that enables you to execute most of the NexentaStor functions.

NMS

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

NMV

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

R—Z**Recursive**

Recursive is a property of the Auto-Sync service that enables copying of a dataset with nested datasets.

Remote Replication (RR)

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

The RR protocol uses the client-server model for transferring data stream and includes the following components:

- **Remote Replication Manager** (`rrmgr`) — the client multi-threaded application that initiates remote connection between source and destination appliances. `rrmgr` can act as sender in local-to-remote replication and as receiver in remote-to-local replication.
- **Remote Replication Daemon** (`rrdaemon`) — a Service Management Facility (SMF) service that starts when the system boots. RR daemon can handle multiple client requests at the same time. `rrdaemon` acts as receiver in local-to-remote replication and as sender in remote-to-local replication.

Reverse Direction

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

Service Management Facility (SMF)

A component of the underlying Solaris Operating System. SMF manages system and network services, such as `ssh`, `snmp`, `nfs`, etc.

For more information, see <http://www.illumos.org/man/5/smf>.

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 NexentaStor 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, replication, 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 filesystem and then streams it to the `zfs receive` command. The `zfs receive` command creates a new filesystem 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

A virtual block device created over ZFS. Zvols are used to expose SCSI targets (iSCSI or FC) to hosts. You can also use replication and snapshotting services on zvols.

This page intentionally left blank

Global Headquarters

455 El Camino Real
Santa Clara, California 95050

Nexenta EMEA Headquarters

Camerastraat 8
1322 BC Almere
Netherlands

Nexenta Systems Italy

Via Vespucci 8B
26900 Lodi
Italy

Nexenta Systems China

Room 806, Hanhai Culture Building,
Chaoyang District,
Beijing, China 100020

Nexenta Systems Korea Chusik Hoesa

3001, 30F World Trade Center
511 YoungDongDa-Ro
GangNam-Gu, 135-729
Seoul, Korea

