



Global Leader in Software-Defined Storage.

Auto-Sync

User Guide 4.0.5

Date: December, 2016

Subject: Auto-Sync

Software: NexentaStor

Software Version: 4.0.5

Part Number: 3000-at_sync-4.0.5-000058-A

Copyright © 2016 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 Systems (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 Systems’ 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 Systems sales office for information on feature and product availability. This document includes the latest information available at the time of publication.

Nexenta, NexentaStor, NexentaEdge, and NexentaConnect are registered trademarks of Nexenta Systems 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™	4.0.5

Contents

Preface	v
1 Introduction	1
Definitions	1
About Auto-Sync Usage Scenarios	1
Auto-Sync Features	2
Auto-Sync Deployment Topologies	2
Best Practices	4
About Auto-Snap	4
About Auto-Sync Replication Algorithm	5
About Incremental Snapshots	5
About Transport Protocols	6
About the Directions of Replication	6
About Replicating Content	7
About Changes at Destination	8
2 Managing Auto-Sync	9
Before You Start to Use Auto-Sync	10
Viewing Summary Information About an Auto-Sync Service	10
Creating an Auto-Sync Service Instance	10
Establishing the SSH-binding for an Auto-Sync Service	17
Binding a NexentaStor Appliance to a Virtual IP Address	17
View the bindings	18
Modifying the Auto-Sync Properties	18
Enabling an Auto-Sync Service Instance	19
Disabling an Auto-Sync Service Instance	19
Deleting an Auto-Sync Service Instance	20
Cloning Auto-Sync Snapshots	21
Deleting Auto-Sync Snapshots	22
Resuming an Auto-Sync Service	23
Unlocking a Secondary Dataset	23

Repairing an Auto-Sync Service that Failed During the Initial Replication	24
Executing an Auto-Sync Service by Administrative Action	25
Running Auto-Sync Services in Series	26
Running Auto-Sync Services in Parallel	26
Stopping an Auto-Sync Service	26
Restarting an Auto-Sync Service	27
Flipping the Direction of Replication	27
Viewing the Auto-Sync Log File	28
Saving Auto-Sync Log Files	29
Replicating from a Snapshot	30
Viewing the Auto-Sync Statistics	30
3 Advanced Configuration	32
Creating an Auto-Sync Service Instance with Advance Options	32
About Roles of the Auto-Sync Service	34
Creating Multiple Services	35
Changing Paths to Primary or Secondary Dataset	36
Disaster Recovery Use Case	37
Migration Use Case	39
Managing Snapshots from Secondary Appliance	40
About Auto-Sync Replication in the HA Cluster Environment	42
Configuration Settings for Replication Between non-HA and HA Cluster	43
Configuration Settings for Replication Between HA Cluster and HA Cluster	43
Configuration settings for Replication from HA Cluster to Single Node	45
4 Advanced Settings	47
Glossary	51

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.

Product Versions Applicable to this Documentation:

Revision	Date	Description
3000-at_sync-4.0.5-000058-A	December, 2016	GA

Contacting Support

Methods for contacting support:

- Send your support questions and requests to support@nexenta.com.
- Using the NexentaStor user interface, NMV (Nexenta Management View):
 - a. Click **Support**.
 - b. Select an action:
 - **Send by email**
Send the support request to the Nexenta support email.
 - **Save to disk**
Saves the support information to the `/var/tmp` directory on the NexentaStor appliance.
 - c. Complete the request form.
 - d. Click **Make 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:

- [Definitions](#)
- [About Auto-Sync Usage Scenarios](#)
- [Auto-Sync Features](#)
- [Auto-Sync Deployment Topologies](#)
- [Best Practices](#)
- [About Auto-Sync Replication Algorithm](#)
- [About Incremental Snapshots](#)
- [About Transport Protocols](#)
- [About the Directions of Replication](#)
- [About Replicating Content](#)
- [About Changes at Destination](#)

Definitions

This document uses the following terms.

Terms	Definition
Local	Local represents the node on which the Auto-Sync service is run.
Remote	Remote represents the node where data is synchronized from or synchronized to.
The above two terms are generally used to represent the flow of replication.	
Primary	Primary represents the site where the primary dataset exists. In other words it represents the location on which the network clients access the data.
Secondary	Secondary represents the Disaster Recovery site.
Local Volume	NexentaStor enables you to aggregate the disks in the system to Local volume otherwise known as Pools.
Folders	Folders represent File systems that are allocated to the local volume/pools.
zVol	zVol represents block devices that resides on the pool.

About Auto-Sync Usage Scenarios

Auto-Sync is a NexentaStor data and metadata replication service that combine advantages of the built-in ZFS send/receive functionality of the underlying distribution based on illumos™ with tunable replication schedule. This section describes how Auto-Sync can be used in two primary use cases.

Auto-Sync generates snapshots at the Primary appliance on a set schedule. You can create snapshots of a selected dataset and save them in one or more Secondary appliances. The destination can be either local pools on the same NexentaStor appliance or on a remote NexentaStor appliance. You can create multiple service from the same Primary appliance dataset to different Secondary appliance datasets.

Use these snapshots as a backup or an archive, as well as in the disaster recovery use case to handle network failures or sudden power outages.

Auto-Sync Features

Auto-Sync is NexentaStor's periodic asynchronous replication facility supporting block and file datasets with the following features:

- Replication of ZFS folders, snapshots and zvols
- Flip direction replication to restore data in case of a disaster
- Fully configurable tunable replication schedule
- Bandwidth throttling
- Compression and deduplication of the replication stream
- Network trunking support
- Delivers Recovery Point Objectives (RPO) as low as 10 mins

Auto-Sync Deployment Topologies

This section provides information to help you configure Auto-Sync based on your requirements. You may run Auto-Sync service on a single NexentaStor appliance for local replication or on multiple NexentaStor appliances for distributed replication. You can use different topologies to perform data replication.

This list contains the most supported topologies. For all the topologies mentioned below, run the Auto-Sync service from the Secondary site. You can schedule an Auto-Sync service for any dataset, such as local volume, list of folders/zvols, snapshot, or zvol on local volume. Use any one of these topologies to migrate a dataset or datasets between two pools in the same NexentaStor appliance or on different NexentaStor appliances. When replicating datasets, if the service is interrupted the replication can be continued during the next run.

- **Data migration of a dataset, or datasets between two pools on the same NexentaStor appliance**

All features and functionalities are supported in this replication service. When you run Auto-Sync service from Pool A to Pool B, the Secondary dataset in pool B is locked for certain operations. You need to unlock datasets and/or transfer ownership to perform the following:

- create snapshots from the secondary dataset
- delete, rename, and clone the existing snapshots

See [Unlocking a Secondary Dataset](#)

- **Single node appliance to Single remote node appliance**

To create an Auto-Sync service from the local host appliance to a remote appliance, you must first establish the ssh-binding to the remote host by using IP address of the remote host. To facilitate the Auto-Sync service, you may use any supported transport protocol. For more information on the different protocols, see [About Transport Protocols](#). When you run Auto-Sync service from one host to another, and if one of them fails you can recover and restore the datasets.

- **HA cluster to Single node appliance**

To create an Auto-Sync service in this HA Cluster environment, setup SSH bindings between both nodes at Primary site and the single node at Secondary site. Setup Auto-Sync between Primary cluster and Secondary appliance dataset using Virtual IP(VIP) on the Primary appliance HA cluster.

- **HA cluster to HA cluster**

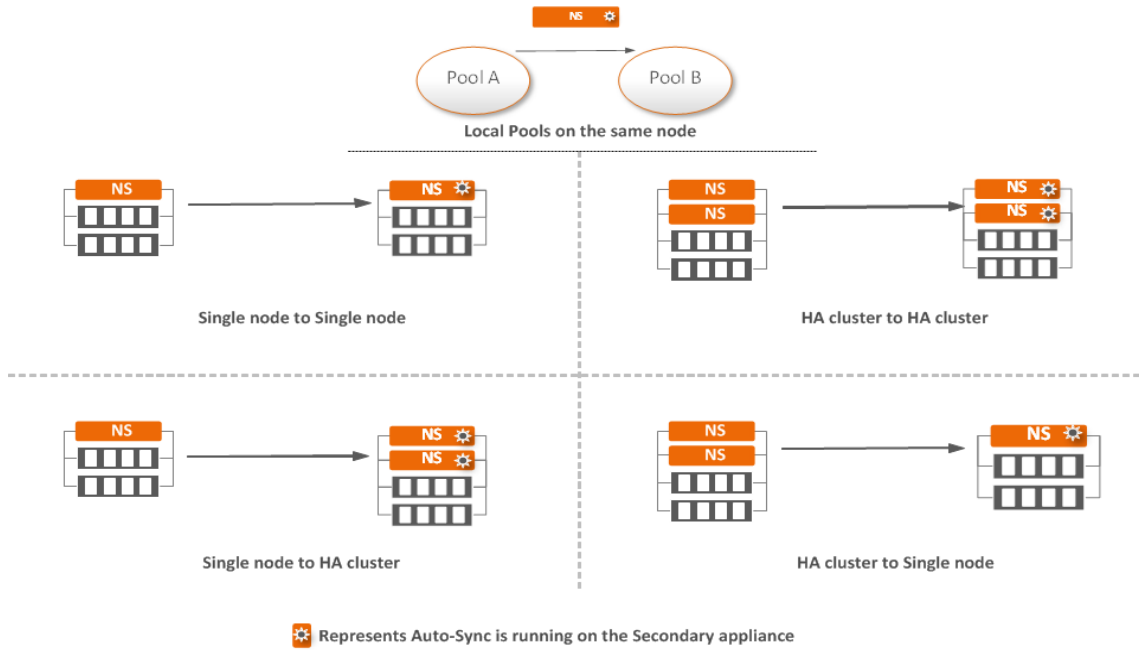
To create an Auto-Sync service in this HA Cluster environment, set up SSH bindings between all nodes. Each Primary cluster node should be SSH accessible to each Secondary cluster node of the replicated cluster group by using the physical IP addresses, not VIPs. However, establish the Auto-Sync datapath between VIPs of each HA Cluster.

- **Single node appliance to HA cluster**

To create an Auto-Sync service in this HA Cluster environment, set up SSH bindings between the non-HA NexentaStor appliance and each of the HA Cluster nodes. Local hosts should be SSH-accessible to each node of the replicated cluster group by using the primary address and to the whole HA-cluster group by using shared IP addresses (VIPs).

In any of the supported deployment topologies, you may configure different retention policies at the Primary appliance and at the Secondary appliances

Figure 1-1: Most supported Auto-Sync deployment topologies.

**See Also:**

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

Best Practices

- Nexenta highly recommends to run the Auto-Sync service on the Secondary site. Use longer term retention on the snapshots of the secondary appliance to recover data for a longer period.
- Even though replication might be possible between a NexentaStor appliance and a generic ZFS system, Nexenta does not recommend it and Nexenta Support Service will only engage on Auto-Sync configuration between NexentaSto appliances.
- Auto-Sync service can be configured between Primary site dataset and Secondary site dataset but Secondary site dataset should not be used or shared. Auto-Snap is not supported on Secondary site dataset.
- Auto-Snap and Auto-Sync can be run concurrently on the same Primary site dataset but they should be managed from one side of a setup.
- To have a better control of Auto-Sync after a primary site failure, the Auto-Sync service should always be run from the Secondary site.
- A dataset can be the target of only one Auto-sync service and a service can have only one schedule per target. However, same Primary site dataset can be shared by multiple Auto-Sync services and each can have different targets.

- Use Flip Direction feature to restore datasets in the case of disaster; reverse service feature is removed in this release.

Warning:

- To recover datasets locally in addition to remote recovery you should combine Auto-snap and Auto-sync. Local refers to where Auto-Sync is running.
 - Auto-Snap should not be run on the Auto-Sync Secondary appliance dataset.
-

About Auto-Snap

Auto-Snap provides the following advantages:

- Automatically creates a point-in-time representation of a dataset
- Performs snapshots on a regular basis (daily, monthly, hourly)
- Creates snapshots of the root folder and its descendants
- Assists in creating a retention policy

About Auto-Sync Replication Algorithm

When the Auto-Sync service instance runs for the first time, Auto-Sync creates an initial snapshot of a Primary site dataset and sends it to Secondary appliance. The initial snapshot is a full copy of the Primary appliance 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 Primary appliance

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 Primary appliance and Secondary appliance

Auto-Sync compares the lists of snapshots at Primary appliance and at Secondary appliance and locates a pair of latest identical snapshots.

3. Sending incremental stream from the common snapshot (the incremental Primary site) to the current snapshot (the incremental Secondary site)

By comparing the set of snapshots on the Primary and the Secondary site, Auto-Sync identifies if it needs to send a snap from Primary to Secondary. Instead of copying the new snapshot from the Primary site, Auto-Sync generates and transfers only the changes from Primary site snapshot to Secondary site snapshot.

See Also:

- [About Incremental Snapshots](#)

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

About Incremental Snapshots

For the Auto-Sync service to complete successfully, on every run, Auto-Sync verifies the latest identical snapshots at Primary site and Secondary site. 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 Usage Scenarios](#)

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 the ground up by Nexenta. The protocol uses libzfs combined with the `rrdaemon` service. Auto-Sync uses Remote Replication protocol for local-to-remote and remote-to-local replication. RRP supports multiple TCP connections for a given replication stream and allows TCP connections of any given session to use multiple paths (IP routes) between the Primary site and the Secondary site hosts, to maximize resource usage, increase redundancy, and ultimately improve performance of the replication.

If you use the RR protocol, both Primary and Secondary sites 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 compared 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 Primary datasets and replication datasets reside on local NexentaStor appliance.

- **From Local to Remote Host (L2R)**

The Primary site dataset resides on local NexentaStor appliance. The Secondary site dataset resides on a remote NexentaStor appliance.

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

- **From Remote to Local Host (R2L)**

Primary site dataset is stored on a remote NexentaStor appliance. Secondary site dataset resides on local NexentaStor appliance.

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

Note: Local represents the node on which the Auto-Sync service is run.

About Replicating Content

When you select a Primary appliance 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 file system, you must enable the `dircontent` property. If you want to use the `dircontent` property, you must select the following symbol: `/*`. In NMV select the **Replicate Content** option.

Note: Do not disable the **Replicate Content** option when replicating between pools and in a clustered environment.

The following diagrams compare replication with the enabled and disabled replicate content (`dircontent`) property.

Figure 1-2: Replication with replicate content disabled.

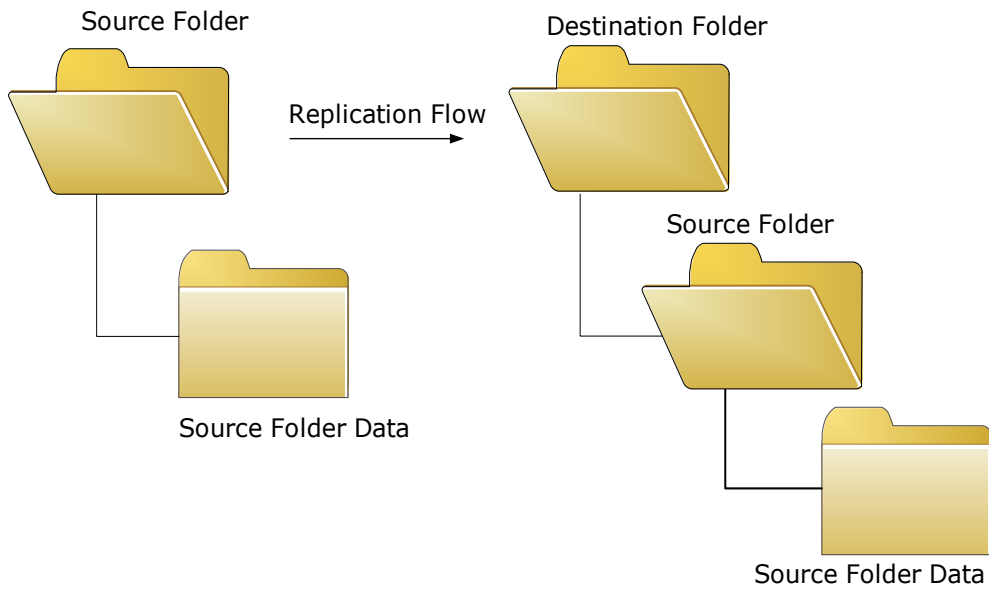



Figure 1-3: Replication with replicate content enabled.

**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 change any data in the destination folder, Auto-Sync identifies the changes as a loss of synchronization between Primary appliance and Secondary appliance. 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.

See Also:

- [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:

- [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](#)
- [Cloning Auto-Sync Snapshots](#)
- [Deleting Auto-Sync Snapshots](#)
- [Resuming an Auto-Sync Service](#)
- [Unlocking a Secondary Dataset](#)
- [Repairing an Auto-Sync Service that Failed During the Initial Replication](#)
- [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](#)
- [Flipping the Direction of Replication](#)
- [Viewing the Auto-Sync Log File](#)
- [Saving Auto-Sync Log Files](#)
- [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 Community Edition. 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

The summary information page includes name, service owner, 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_name>
```

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 local volume, list of folders/zvols, snapshot, or zvol on local volume.

If you create an Auto-Sync service from the local host Primary appliance to a remote Secondary appliance, you must first establish the ssh-binding to the remote host by using local IP addresses of the remote host.

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	Parameter in NMC	Description
Name	instance	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. Can only contain alphanumeric characters (a-z, 0-9) in addition to the following three special characters: underscore (_), hyphen (-)
	comment	Description of Auto-Sync service.
Direction	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)
Local or Remote destination folder	to-url	Destination for replication
Replicate Content	dircontent	Replicate content of selected Primary datasets or replicate entire Primary. For more information, see: About Replicating Content .
Period Multiplier	Period Multiplier	Set this parameter while creating an auto-sync service or when modifying the service property using NMV and NMC. Use this parameter for more frequent Primary appliance snapshots and less frequent replication. For example, if you set the period_multiplier to N and greater than 1, the service will run as scheduled but will replicate only every Nth time. All the intermediate (N - 1) scheduled times the service will simply generate its next snapshot at the Primary appliance, and then exit. To set the period multiplier using NMC: <pre>nmc:/\$ setup auto-sync :service_name property period_multiplier</pre> Set a new value for the period multiplier.

Table 2-1: Auto-Sync properties.

Parameter in NMV	Parameter in NMC	Description
Schedule	frequency -p <period> -H <hour> -N <minute> -D <day> -T <time>	Schedule consists of two parameters: Period and Frequency. Frequency depends on Period. For example, if you set Period to <i>hourly</i> and Frequency to 4, then Auto-Sync runs the service instance every 4 hours. If you set Period to <i>daily</i> and frequency to 5, then Auto-Sync runs the service instance every 5 days, and so on.
Transport Protocol	proto	Replication protocol. Auto-Sync provides the following protocols: <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. RRP is a multi-connection session-based layer 5 (L5) transport over TCP. It supports multiple TCP connections for a given replication stream and allows TCP connections of any given session to use multiple paths (IP routes) between the Primary appliance and the Secondary hosts, to maximize resource usage, increase redundancy, and ultimately improve performance of the replication. • 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 compared to RR. However, the replication performance is slower compared 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.
Remote Source Host	from-url to-url	Primary host for replication.
Local or Remote Source Folder/Zvol	from-url	Primary dataset for replication.
Remote Destination Host	to-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 bound by SSH.
Destination folder	to-fs	Local or remote Secondary dataset.

Table 2-1: Auto-Sync properties.

Parameter in NMV	Parameter in NMC	Description
Separate Interfaces For Data	<code>trunk</code>	Specifies additional physical paths that a TCP connection uses for replication. Multipathing ensures that in case one path fails, replication traffic is routed through another path. Multipathing also aids in replication workload balancing. To use multipathing, configure additional network interfaces on the Primary appliance and Secondary NexentaStor appliances before creating an Auto-Sync service. You can use multipathing only with the <code>zfs+rr</code> protocol.
Keep on Source (Retention Policy)	<code>keep_src</code>	The number of snapshots that you want to keep at Primary 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 (Retention Policy)	<code>keep_dst</code>	The number of snapshots that you want to keep on the Secondary 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.
Replicate recursively	<code>create auto-sync -r</code>	If you enable this option, all nested folders and subfolders are included in replication.
Compression	<code>zip_level</code>	Calibrates the type and the speed of compression. "Best speed" indicates the best performance and the fastest compression (less compression). "Good performance" is biased towards high compression at expense of speed/performance. "Best compression" indicates the best compression and potentially slower performance. The compression is disabled by default.
Type	<code>Compression type</code>	Type of the compression algorithm. The options are: <ul style="list-style-type: none"> • None • lz4 • zlib_def • zlib_huff • zlib_rle • zlib_filtered • zlib_fixed You cannot specify compression for local to local services.

Table 2-1: Auto-Sync properties.

Parameter in NMV	Parameter in NMC	Description
Level	Compression level	Specifies compression policy. The options are: <ul style="list-style-type: none"> • Best speed • Good performance (except lz4) • Best compression
Flip direction flags		
Run replication one last time	last_sync	Defines whether to run Auto-Sync replication before executing <code>flip-direction</code> . Default value is enabled.
Ignore failures	ignore_failures	Ignores the faults during the last Auto-Sync replication that runs before <code>flip-direction</code> . Default value is disabled.

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. Type:

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

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

2. Select a replication interval.

The options are:

- minute
- hourly
- daily
- weekly
- monthly

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



Tip:

3. Specify a replication period.

The replication period depends on the replication interval.

4. Select a preferred direction of replication.

Use the arrow keys to navigate around the options or use hjkl keys.

5. Select a preferred transport protocol.
6. Select a retention policy for the source/Primary dataset.
7. Select a Source host.
 1. If the Primary dataset resides on the local NexentaStor appliance, press **Enter** and accept the default value `localhost`.
 2. If the Primary dataset is located on the remote appliance, type the IP address or host name of the Primary appliance host.

Note: If a Primary dataset is located on a remote NexentaStor appliance, you may view the hosts that are already ssh-bound. An error message displays if you specify a host that do not have ssh-binding.

8. Select a Source dataset.
You can select a folder, or a zvol.
9. Select a retention policy for the Secondary dataset.

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

10. Select a multi-connection option if you need to use different data paths for this replication service.
11. Select a destination folder on the Secondary appliance.
 1. If the Secondary 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) Select an IP address or host name of the remote NexentaStor appliance; then select a remote dataset.

Note: If a Secondary 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.

3. If a Primary dataset resides on a remote appliance, select a local dataset that you want to use as replication destination.
12. Specify the number of snapshots to keep at Source.
Alternatively, specify a period in YDHMS format.
13. In the **Replicate recursive** field, type `y` or `n`.
If you type `y`, the Auto-Sync service replicates the dataset and all nested folders. If you type `n`, the service replicates the dataset leaving out all nested folders.

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

14. Specify the number of connections.

The default number of TCP connections is 2.

System response:

```
Source URL: zfs+rr://localhost/<dataset>
Destination URL: zfs+rr://<remote-host>/<folder>
Validating parameters...OK
Check compatibility...OK
Check available space on destination...OK
```

15. Specify settings for deduplication by typing `y` or `n`.
16. Specify whether to copy ZFS properties by typing `y` or `n`.
17. Type **Rate limit** in KB/sec.
18. Select **Compression type**:
 - None
 - lz4
 - zlib_def
 - zlib_huff
 - zlib_rle
 - zlib_filtered
 - zlib_fixed
19. Select a **Compression level** from the following options:
 - Best speed
 - Good performance
 - Best compression
20. Type custom name prefix for the snapshot.
21. Optionally, specify a custom name for the Auto-Sync service instance.
22. Optionally, type a descriptive commentary.

System response:

```
About to create a new auto-sync service. (y/n)
```

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

23. 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'
```

Note:

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, also known as 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 [Configuration Settings for Replication Between non-HA and HA Cluster](#).

❖ *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 **Binding type** field, select **Regular**.
4. In the **Remote Server** field, type the IP address of the remote NexentaStor appliance.
5. In the **Remote User** field, type the name of the user with root credentials.
6. Type the remote super user password.
7. Optionally, select the **Bidirectionally** checkbox.
8. Click **Bind**.

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.

Binding a NexentaStor Appliance to a Virtual IP Address

Before you complete the steps described in this section, read and complete the steps in [About Auto-Sync Replication in the HA Cluster Environment](#) and [Configuration Settings for Replication Between non-HA and HA Cluster](#).

❖ *To bind a NexentaStor appliance to a virtual IP address:*

1. Click **Settings > Network**.
2. In the **Network** panel, click **SSH-Bind**.
3. In the **Binding type** field, select **VIP**.
4. Select a virtual IP address.

Note: You must first bind this NexentaStor appliance to each node of the HA Cluster as described in [About Auto-Sync Replication in the HA Cluster Environment](#).

5. Click **Bind**.

❖ *To bind a NexentaStor appliance to a virtual IP address, 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. Type the VIP address.

View the bindings

❖ *To view the bindings, using NMC:*

1. Type:

```
nmc:/$ show network ssh-bindings
```

System response:

HOST	PINGABLE	SSH-ACCESSIBLE	IS-APPLIANCE
root@10.3.70.85	Yes	Yes	Yes
root@10.3.70.93	Yes	Yes	Yes
root@10.3.35.182	Yes	Yes	Yes
root@10.3.70.86	Yes	Yes	Yes

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_name> 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-user 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:*

1. Type:

```
nmc:/$ setup auto-sync <service instance_name> 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 Primary and Secondary 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_name> 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 Primary and Secondary datasets.

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

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

Note: 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. To destroy the service without customizing any cleanup options:
 1. Click on the Auto-Sync service that you want to delete by selecting the checkbox.
 2. Click **Delete**.

If you need to customize the cleanup options on the Primary appliance and on the Secondary appliance, choose the next deleting option.

4. Alternatively, to destroy the service with customized cleanup options follow the steps listed here:
 1. Select the autosync service that you want to delete by clicking directly on the service name.
 2. In the Update Auto-Sync service window, click **Delete**.
Delete Auto-Sync service window opens
 3. Optionally, force the `_delete_` operation by selecting **Force stop**.
You may need to force the operation if the Auto-Sync service is running. However, using the force option is not recommended. Forcing the service to stop running may cause partial and incomplete results.
 4. Optionally, select the cleanup options on the Primary appliance and on the Secondary appliance:
You may choose to destroy the service snapshots at the Primary appliance or on the Secondary and destroy the replicated datasets on the Secondary.
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**.
You cannot specify the force or cleanup options on this page.

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

1. Type:

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

System response:

```
Cleanup(s) : (Use SPACEBAR for multiple selection)
Destination Properties Destination snapshots Destination datasets Source
properties Source snapshots
```

2. Optionally, select cleanup options.

To selectively destroy the services for a specific interval choose the `-i` option.

To force destroy the service, even if the service is currently running choose the `-f` option. Forcing the service to stop may cause incomplete results depending on the service instance and timing.

3. Confirm the operation by typing `y`.

See Also:

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

About Auto-Sync Snapshots

You can take ownership of an AutoSync snapshot and consequently, the snapshot will not be affected by the retention policy. The owner of the snapshot will have the following privileges:

- Create writable clones of the snapshot to use them as regular datasets
Before cloning an Auto-Sync snapshot, remove the user owned label by using the `take-ownership` command
- Mount clones and manage as needed
- Delete the associated clones
- Revert ownership of the snapshot to be managed by AutoSync

Cloning Auto-Sync Snapshots

❖ *To clone an Auto-Sync snapshot, using NMC:*

1. Remove user owned label from the snapshot that you want to clone by typing:

```
nmc:/$ setup snapshot <snapshot_name> take-ownership
```

2. Clone the Auto-Sync snapshot by typing:

```
nmc:/$ setup snapshot <snapshot_name> clone
```

3. Type a new path to the snapshot in `<volume>/<folder>` format.

System response:

NAME	USED	AVAIL	REFER	MOUNTED	QUOTA	DEDUP	COMPRESS
pool1/test	1K	9.78G	106K	yes	none	off	on

4. Optionally, restore the Auto-Sync service ownership for the snapshot:

```
nmc:/$ setup snapshot <snapshot-name> restore-ownership
```

Warning:

If the snapshot is currently owned by a user or if it is used by another Auto-Sync service, you can still clone this snapshot using Take Ownership property but you can not delete or rename this snapshot.

Take Ownership property is unique to the Primary and Secondary datasets. When replicating a snapshot, the ownership does not get transferred to the snapshot on the Secondary site.

Deleting Auto-Sync Snapshots

You may want to delete the snapshots created by Auto-Sync to free space on disk or to discard changes. In order to delete a snapshot, you must unlock the datasets first. Refer to [Unlocking a Secondary Dataset](#).

- ❖ *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 Primary appliance and Secondary appliance. 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 Primary appliance and folders on the Secondary appliance as in the deleted service. Follow the steps described in [Creating an Auto-Sync Service Instance](#).

Unlocking a Secondary Dataset

When you create an Auto-Sync service, Auto-Sync sets a special marker on the Secondary dataset. This marker protects the dataset and its nested datasets from being modified or deleted.

Typically, when you successfully delete an Auto-Sync service, the marker is deleted as well. However, if an Auto-Sync service is deleted incorrectly, or if a Primary appliance becomes unavailable, the Secondary dataset may become locked and unavailable for any modifications.

To remove the Auto-Sync marker and unlock the Secondary dataset, use the `unlock-dataset` feature. The `unlock-dataset` feature verifies whether the dataset is used by Auto-Sync. If the dataset is not used, the command removes Auto-Sync recorded properties and Auto-Sync mark for the specified dataset and its children. This feature makes the Secondary dataset available for any standard operations.

Note: While unlocking, ensure that the selected volume and its children are not destinations for existing active auto-sync services.
This feature can be applied for both the Primary and the Secondary datasets.

❖ *To unlock a Secondary dataset, using NMV:*

1. Click **Data Management > Auto Services**.
2. In the **Auto-Sync Services** pane, click **Unlock Dataset**.
A dialog window with a list of Secondary dataset displays.
3. Select a Secondary dataset from the drop-down list.
4. Optionally, select the **Recursive checkbox**.
5. Click **Unlock**.

❖ *To unlock a Secondary dataset, using NMC:*

1. Type:

```
nmc:/$ setup auto-sync unlock-dataset -n pool1 -r -w
```

Ensure that the selected volume and its children are not destinations for existing active Auto-Sync services.

2. Alternatively, you can type:

```
nmc:/$ setup <folder | zvol> <dataset_name> unlock-dataset
```

3. Specify an option:

<code>-n <folder></code>	Specify a folder from which you want to remove the Auto-Sync marker.
<code>-h</code>	use <code>-h</code> for more information.
<code>-r</code>	use <code>-r</code> for recursive unlock of selected dataset and child datasets/snapshots
<code>-w</code>	use <code>-w</code> option to unlock selected dataset and wait for process to complete

Repairing an Auto-Sync Service that Failed During the Initial Replication

During the initial replication, Auto-Sync creates a snapshot of a full dataset and then transfers it to the Secondary dataset. On the next scheduled run of the Auto-Sync service, Auto-Sync creates another snapshot, compares this snapshot with the snapshot at Secondary site, and then sends the changes from Primary appliance to the Secondary site.

If Auto-Sync cannot detect identical snapshots at Primary appliance and Secondary, the Auto-Sync service instance will fail during the initial replication. Therefore, during the next run of the Auto-Sync service, it will try to recover missing datasets that failed during the first replication. You can use the `reinitialize` command in the following scenarios:

- if the automated recover fails,
 - if the Secondary should be cleaned totally,
 - to repeat the transfer of the initial snapshot.
- ❖ *To repair an Auto-Sync service that failed during the initial replication, using NMC:*

1. Type:

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

System response:

```
This action will try to remove the <service> service's corresponding
datasets and snapshots at the destination side and reinitialize the
service
Would you like to continue? (y/n)
```

2. Type y.

System response:

```
The <service> service has been successfully reinitialized
Reinitialized 1 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 Primary and Secondary 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 checkbox 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.
 1. Click on the Auto-Sync service instance that you want to run.
 2. 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_name> 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

While 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 forcibly 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 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.
 1. Click on the Auto-Sync service instance that you want to stop.
 2. 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 Secondary site becomes unavailable during replication, Auto-Sync locks the Primary appliance folder. Therefore, when the Secondary 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`.

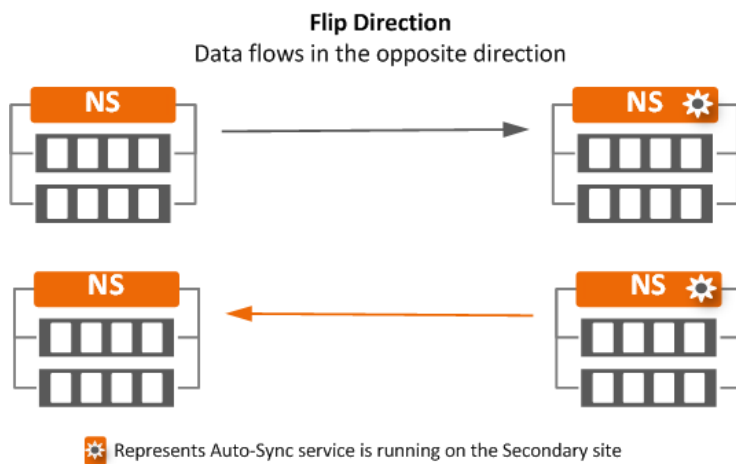
System response:

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

Flipping the Direction of Replication

The `flip-direction` command changes the direction of the Auto-Sync service replication. You may use this feature to restore data back to Primary site as part of a dataset recovery operation. `Flip-direction` does not create a copy of the Auto-Sync service at Secondary site.

Figure 2-1: Example of a Flip-Direction.



❖ *To flip the direction of replication, using NMV:*

1. Click **Data Management > Auto-Services**.
2. Under **Auto-Sync Services**, click **Show**.
3. Select an Auto-Sync service instance.
4. Click **Flip Direction**.

❖ *To flip the direction of replication, using NMC:*

1. Type:

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

You can specify the following properties for `flip-direction` command: `sync_first` and `ignore_sync_fail`. See [Auto-Sync properties](#).

Note:

Do not execute Flip Directions in the case of two or more datasets on the same Secondary appliance. You can change the direction of replication only if you configure local-to-remote or remote-to-local Auto-Sync services.

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 NMC:*

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](#)

Saving Auto-Sync Log Files

You can collect and save Auto-Sync log files to the `/tmp` directory. After you execute the `collect-logs` command, Auto-Sync creates a folder named after the Auto-Sync service and includes the following log files into this folder:

- Log file of the Auto-Sync service
- Configuration file of the Auto-Sync service
- If the Auto-Sync service uses the RR protocol, Auto-Sync collects log files for the Remote Replication Manager (`rrmgr`) from the local appliance and for the Remote Replication Daemon (`rrdaemon`) from the remote appliance.

The Auto-Sync log files are automatically included into the support request when you create it. If you do not save the Auto-Sync log files, they are not generated and not included in the support request.

For more information, see [Contacting Support](#).

This functionality is only available in NMC.

- ❖ *To save Auto-Sync log files for all Auto-Sync services, using NMC:*

1. Type:

```
nmc:/$ setup auto-sync collect-logs
```

- ❖ *To save Auto-Sync log file for a selected Auto-Sync service, using NMC:*

1. Type:

```
nmc:/$ setup auto-sync <service name> collect-logs
```

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>
```

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

You must specify parent dataset of the selected snapshot as a Primary dataset.

See Also:

- [Advanced Settings](#)

Viewing the Auto-Sync Statistics

Auto-Sync provides the real-time statistics. You can monitor how the Auto-Sync service progresses and promptly correct system errors, if any found.

Note: You can monitor Auto-Sync statistics only when the 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 as shown here.

SHOW SYNCING SERVICES										
<input type="checkbox"/> 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 Advance Options](#)
- [About Roles of the Auto-Sync Service](#)
- [Changing Paths to Primary or Secondary Dataset](#)
- [Disaster Recovery Use Case](#)
- [Migration Use Case](#)
- [Managing Snapshots from Secondary Appliance](#)
- [About Auto-Sync Replication in the HA Cluster Environment](#)

Warning:

This chapter describes the advanced configuration options for Auto-Sync. Nexenta recommends that you use this functionality only if you are an experienced NexentaStor Auto-Sync user.

Creating an Auto-Sync Service Instance with Advance Options

You can create an Auto-Sync service instance using multiple advance options available through NMC.

❖ *To create an Auto-Sync service with advance 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: Advance Options for the Create Auto-Sync command.

Option	Description
-s	Primary dataset. You can specify a folder, a zvol, or a snapshot. Example: <code>nmc:/\$ setup auto-sync create -S <snapshot></code>
-d	Secondary dataset. Specify a folder.
-R	Resume the Auto-Sync service. This option forces Auto-Sync to skip validation of Primary and Secondary datasets. You may want to use this functionality to re-create the service that already contains Auto-Sync marker snapshots in both Primary and Secondary datasets.

Table 3-1: Advance Options for the Create Auto-Sync command. (Continued)

Option	Description
-e	Exclude from replication the folders that match a pattern or a list of patterns. Example: <code>nmcc:/\$ setup auto-sync create -e *-test</code> Use standard Unix pattern 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 the day.
-D <day of the month>	Day of the 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.
-l <number>	Rate limit, Kb/s. Specifies the maximum speed limit for replication.
-u <custom name>	A custom name for the Auto-Sync service instance. Auto-Sync uses the following naming pattern: <code>dataset-index_number</code> . Example: <code>datapool/docs-000</code> Where <code>datapool/docs</code> is the name of the dataset and <code>000</code> 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 Primary dataset is a zvol.
-N	Replicate the dataset non-recursively.
-Z	If you specify this option, then Auto-Sync automatically unmaps the Primary zvol.
-E	Run an Auto-Sync service instance in daemon mode.

Table 3-1: Advance Options for the Create Auto-Sync command. (Continued)

Option	Description
-S	Create an Auto-Sync service instance using a snapshot as Primary 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 Primary. 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.

About Roles of the Auto-Sync Service

Auto-Sync service may act as:

- **Master**

An Auto-Sync service replicates data from Primary to Secondary 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. Auto-Sync replication management logic is stored on the appliance where you create the service.

See [List of Master Auto-Sync services at Primary.](#)

Figure 3-1: List of Master Auto-Sync services at Primary.

SHOW SYNCING SERVICES										
<input type="checkbox"/>	Instance	Direction	Progress	Speed	Time	Snapshots	Frequency	Last run	Duration	View log
<input type="checkbox"/>	performance1-fol1-000		online (attempts made: 11)	-	-	9	hourly	00:00:05,Nov22	0s	
<input type="checkbox"/>	docs-user--guides-000		online	-	-	-	hourly	N/A	0s	
<input type="checkbox"/>	docs-installation-guides-000		online (attempts made: 3)	-	-	-	hourly	00:28:58,Nov22	0s	
<input type="checkbox"/>	docs-reference--guides-000		Task 1 of 1; 0.00 KB of 15.00 KB	0.00 KB/s	-	-	hourly	N/A	0s	
<input type="checkbox"/>	performance1-legal--docs-000		Task 1 of 1; 0.00 KB of 15.00 KB	0.00 KB/s	-	-	hourly	N/A	0s	
<input type="checkbox"/>	test1-archive-000		disabled	-	-	-	hourly	N/A	0s	
<input type="checkbox"/>	test1-data-000		disabled	-	-	-	hourly	N/A	0s	
<input type="checkbox"/>	test1-fol2-000		disabled	-	-	-	hourly	N/A	0s	
<input type="checkbox"/>	test1-list-000		disabled	-	-	-	hourly	N/A	0s	

Figure 3-2: List of Master Auto-Sync services at Secondary.

SHOW SYNCING SERVICES										
Instance	Direction	Progress	Speed	Time	Snapshots	Frequency	Last run	Duration	View log	
performance1-fol1-000		disabled (attempts made: 9)	-	-	-	hourly	23:54:26,Nov21	0s		
test1-fol2-000		online (attempts made: 1)	-	-	-	hourly	00:33:22,Nov22	0s		
test1-archive-000		online	-	-	4	hourly	00:33:24,Nov22	15s		
test1-data-000		Task 1 of 1; 0.00 KB of 15.00 KB	0.00 KB/s	-	2	hourly	00:30:23,Nov22	0s		
test1-list-000		online	-	-	4	hourly	00:33:24,Nov22	14s		
docs-installation-guides-000		online	-	-	-	hourly	N/A	0s		
docs-reference-guides-000		disabled	-	-	-	hourly	N/A	0s		
docs-user-guides-000		disabled	-	-	-	hourly	N/A	0s		
performance1-legal-docs-000		online	-	-	-	hourly	00:33:02,Nov22	0s		

Creating Multiple Services

You can create multiple services from same Primary site to different Secondary sites. Secondary datasets can be either on the local or remote appliance. When you create multiple services, use the Snapshot name prefix option to set unique name for different services. You can use this functionality in various disaster recovery scenarios. If one of the Secondary appliances goes down, Auto-Sync continues to replicate the datasets to the other Secondary appliance. And when the Secondary appliance comes back, Auto-Sync continues to replicate to both the Secondary appliances.

❖ *To set unique name for different services, 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, click **More Options**
4. In the **Snapshot name prefix** field, type unique prefix for each service.

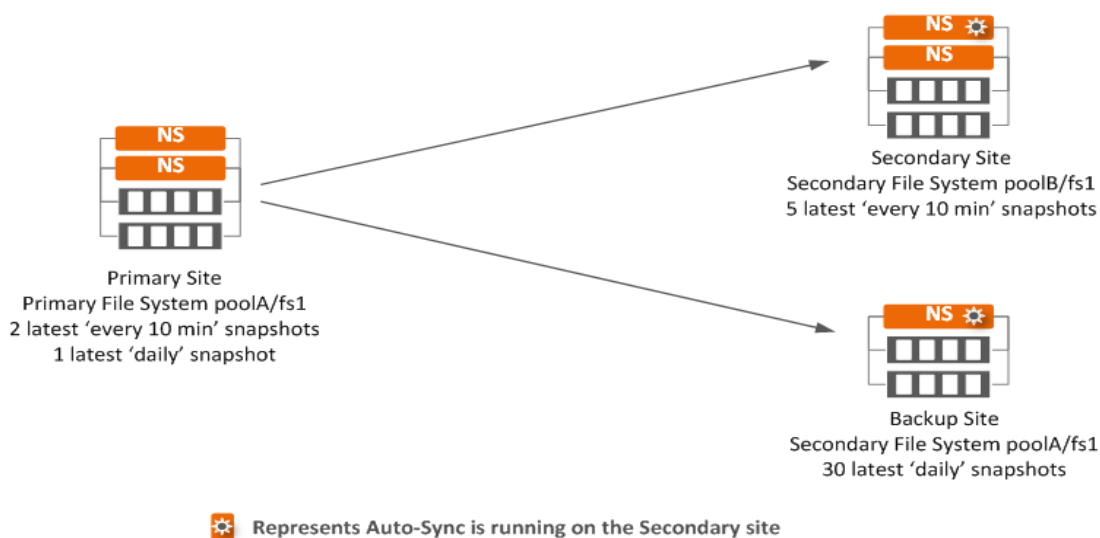
When creating multiple services from the same Primary dataset, you can set unique retention policies for each Secondary site. To learn more about Replicating Intermediate Snapshots, refer Advanced Settings (Chapter 4)

Multiple services replicate its own snapshot and might include the other snapshots too based on the “Replicate Intermediate Snapshot” option. If this option is disabled, the service will generate incremental stream from the first snapshot to the second snapshot. If the option is enabled, the service will include the intermediate snapshots along with the service owned snapshots. On the contrary, when you destroy a service, it will delete only its own snapshots and similarly the retention policy allows the Auto-Sync service to delete only its own snapshot after the period expires.

Note: You can force to destroy all snapshots by choosing the “Force Destroy All Snapshots” option under Force flag. When this option is enabled, it destroys all outdated snapshots on the Secondary, including manually created and the foreign ones that were created by the other services, according to the retention policy of this service. On the Primary site, this option deletes only the snapshots created by that service.

- ❖ *To force destroy all snapshots, using NMC:*
 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 destroy
 4. Click **More Options** button
 5. In the **Force Flags** window, select **Force: Destroy All Snapshots**

Figure 3-3: Example of Auto-Sync to multiple destinations.



Changing Paths to Primary or Secondary Dataset

You can change the paths to Primary or Secondary datasets if you moved the datasets to a new location, modified the IP address of the Primary or Secondary host, or applied any other changes.

This functionality is only available in NMC.

- ❖ *To change the paths to Primary or Secondary 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 Primary host.
3. In the `From fs`, type a path to new location of the Primary dataset.
4. In the `To host` field, type the name of a new Secondary host.
5. In the `To fs` field, type the path to a new location of a Secondary dataset.

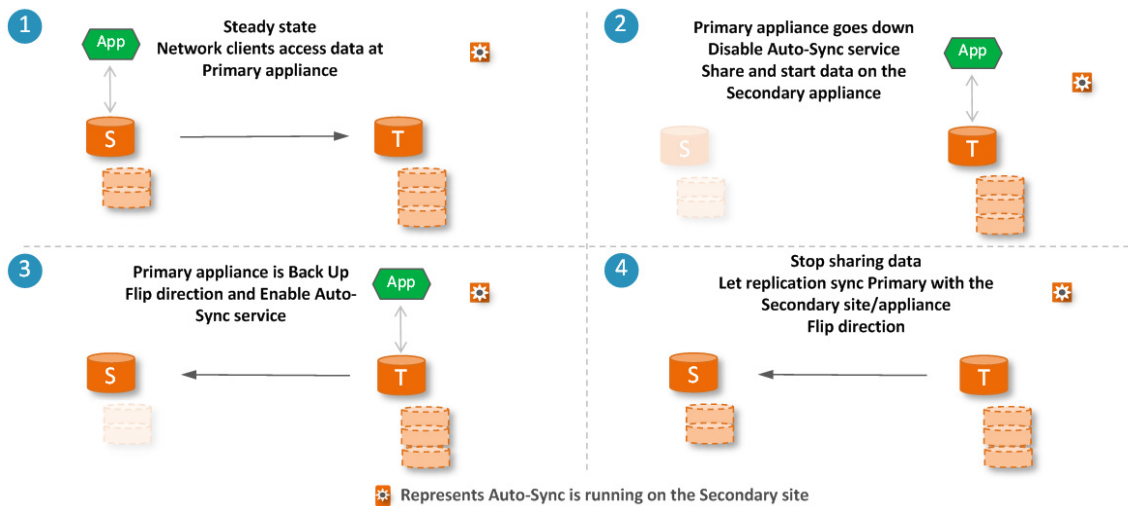
Auto-Sync prints detailed information about the Auto-Sync service.

Disaster Recovery Use Case

In the event of a Primary appliance failure and if it is temporarily unavailable, follow the steps described in this section, so that network clients can access the data at the Secondary appliance. After you recover the Primary appliance, you can switch the service roles back to normal operation.

To restore the snapshots that was taken right before the system failure to the Primary appliance, you must have the Auto-Sync plugin on the Secondary appliance and the Secondary appliance must have the latest replication.

Figure 3-4: Example of a Disaster Recovery use case.






Table 3-2: Steps to restore normal Auto-Sync operation using Flip-Direction feature in case of a disaster.

#	Description
1	<p>Normal Auto-Sync operation.</p> <p>A normal Auto-Sync operation is executed as listed here:</p> <ul style="list-style-type: none"> • Allow the network clients to access the data on the primary appliance. • Run the Auto-Sync service on the Secondary appliance. <p>Auto-Sync service :vol11-fol-000 takes snapshots and replicates them from Primary appliance to Secondary appliance.</p>
2	<p>The Primary NexentaStor appliance becomes unavailable</p> <p>In the event of a disaster and when the Primary appliance becomes unavailable, the network clients can access their data from the Secondary appliance. And if the Primary appliance cannot be recovered, the Secondary appliance can be used as the primary appliance. To set the Secondary as the primary appliance, do the following:</p> <ol style="list-style-type: none"> <p>1. Disable Auto-Sync service</p> <ol style="list-style-type: none"> 1. In NMV > Data Management > Auto Services > Select the Auto-Sync instance you want to disable > Click Disable <p>This is done to ensure that the datasets in the Secondary appliance do not get overwritten in case the Primary appliance comes back.</p> <p>2. Share and start the client application</p> <p>This allows the network clients to modify the datasets on the Secondary appliance.</p>

Table 3-2: Steps to restore normal Auto-Sync operation using Flip-Direction feature in case of a disaster.

#	Description
3 & 4	<p>The Primary NexentaStor appliance is back to operation</p> <p>After you recover the Primary appliance, you can switch the service roles back to normal operation.</p> <p>3. Execute the <code>flip-direction</code> command on the Secondary appliance</p> <ol style="list-style-type: none"> 1. In NMV > Data Management > Auto Services > Click the Auto-Sync instance you want to Flip > Click <code>More Options</code> 2. Turn OFF <code>Run replication one last time</code> <p>This is done to prevent the Auto-Sync service from running the replication one last time before flipping the direction.</p> <ol style="list-style-type: none"> 3. Select <code>Force Receive</code> <p>This lets the zfs receive even if the Secondary snapshot is different from the Primary snapshot.</p> <ol style="list-style-type: none"> 4. Click Save and then click Flip Direction <p>5. Select the Auto-Sync service instance and click Enable</p> <p>6. Stop the network clients from accessing the data on the Secondary appliance</p> <p>7. Run the Auto-Sync Service now to Sync Primary with the Secondary</p> <ol style="list-style-type: none"> 1. In NMV > Data Management > Auto Services > Click the Auto-Sync instance you want to run > Click <code>Run Now</code> <p>Data starts replicating from the Secondary appliance to the Primary appliance. Wait until data is replicated</p> <p>2. Run the <code>flip-direction</code> command at the Secondary</p> <p>Now the data flows from the Primary appliance to the Secondary appliance</p> <ol style="list-style-type: none"> 1. In NMV > Data Management > Auto Services > Click the Auto-Sync instance you want to Flip > Click <code>More Options</code> 2. Turn ON <code>Run replication one last time</code> <p>This is done to sync the datasets between the Primary appliance and the Secondary appliance.</p> <p>3. Remove <code>Force Receive</code></p>
5	<p>Restoring normal Auto-Sync operation.</p> <ol style="list-style-type: none"> 1. In the Primary appliance, reshare the datasets and verify 2. Let the network clients access the data from the Primary appliance

- [Flipping the Direction of Replication](#)

Migration Use Case

When you integrate new hardware into your IT infrastructure, you may need to migrate your data from a Primary NexentaStor appliance to a Secondary appliance. To migrate the data, execute the steps described in this section. Whenever possible run the Auto-Sync service on the Secondary appliance. Perform the migration during a scheduled maintenance period to minimize the service downtime

Table 3-3: Steps to restore the data from Secondary back to Primary.

#	Description
1	<p>Normal Auto-Sync operation</p> <p>A normal Auto-Sync operation is executed as listed here:</p> <ul style="list-style-type: none"> • Allow the network clients to access the data on the Primary appliance. • Run the Auto-Sync service on the Secondary appliance. <p>Auto-Sync service :vol1-f01-000 replicates data from Primary appliance to Secondary appliance.</p>
2	<p>Migration of the datasets (Failover)</p> <p>To migrate the data from the Primary appliance to the Secondary, execute the steps described below.</p> <ol style="list-style-type: none"> 1. Create Auto-Sync service on the Secondary appliance. You may set different data paths for the replication service. More options are available in NMC to configure the service. 2. Start Auto-Sync service at the Secondary appliance. This migrates all the data from the Primary appliances to the Secondary appliance 3. Check the status of the service 4. Check the datasets and snapshot on the Secondary appliance
3	<p>When the Primary NexentaStor appliance is shut down intentionally</p> <p>When the Primary appliance is shut down for maintenance or any other possible reasons, start and share the data on the Secondary appliance. Now you can let the network clients access their data from the Secondary appliance.</p> <p>See Step 2 from Disaster Recovery Use Case</p>
4	<p>Restoring the data from Secondary back to the Primary (Failback)</p> <p>Ideally, you want to restore your data from the Secondary back to the Primary appliance before you failback. This allows users who are currently accessing their data on the Secondary because of failover, to continue accessing their data from the Primary appliance.</p> <p>To restore data from Secondary appliance to Primary appliance: Follow Steps 3, 4 and 5 from Disaster Recovery Use Case</p>

Managing Snapshots from Secondary Appliance

You can separately manage the snapshots copied from the Primary appliance on the Secondary appliance. In order to do so, transfer the ownership of the Auto-Sync snapshots to the user, using the Take Ownership feature at the Secondary appliance. Once the user becomes the owner, snapshots will not be affected by the retention policy set on them but user cannot delete or rename these snapshots. To manage these snapshots, clone them on the Secondary appliance. On these cloned datasets you can read-write, share, rename or destroy as needed.

The following use case is an example to show how to leverage the Take Ownership property to manage the snapshots on the Secondary appliance.

Figure 3-5: Example of managing clones at Secondary site.

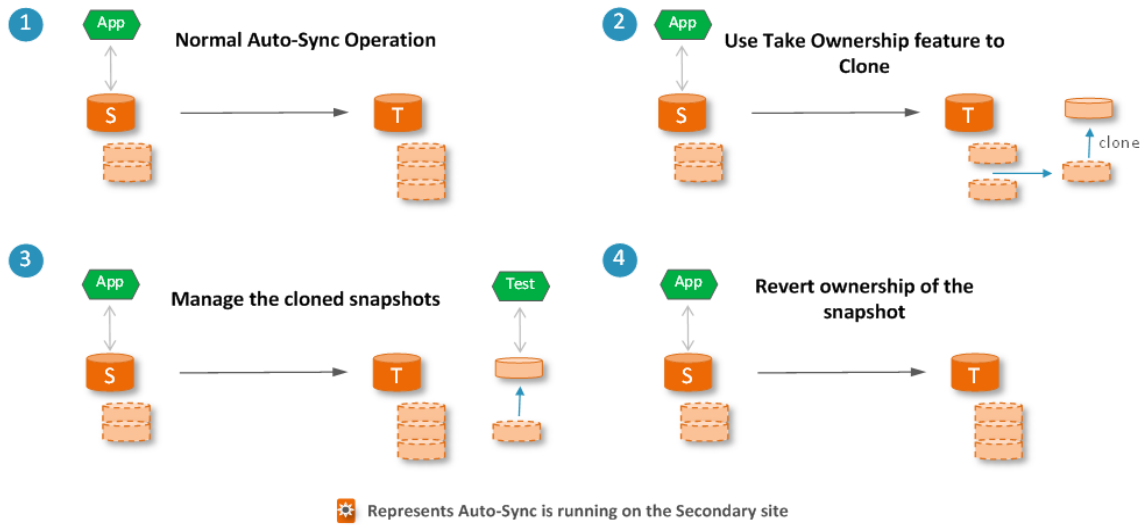


Table 3-4: Steps to manage cloned snapshots at the Secondary site.

#	Description
1	<p>Normal Auto-Sync operation</p> <p>A normal Auto-Sync operation is executed as listed here:</p> <ul style="list-style-type: none"> • Allow the network clients to access the data on the Primary appliance. • Run Auto-Sync service on the Secondary appliance. • Auto-Sync service generates snapshots at the Primary appliance on a set schedule or manually • Auto-Sync service :vol1-f01-000 replicates data from Primary appliance to Secondary appliance. <p>Note: You can set different retention policies at the Primary and Secondary sites</p> <p>Example 3-1:</p> <ul style="list-style-type: none"> • keep 2 “10 min” snapshots at Primary site • Keep 3 “10 min” snapshots at Secondary site • Save daily snapshot at Secondary, and keep 3 copies <p>At this point the snapshots are owned by the Auto-Sync service</p>
2	<p>Use Take Ownership feature to clone</p> <ul style="list-style-type: none"> • Use Take Ownership feature to transfer the ownership of the Auto-Sync snapshots to the user. • Now you can create writable clones of the snapshot <p>To clone the snapshots at the Secondary appliance, see Cloning Auto-Sync Snapshots</p> <p>Once you transfer the ownership of the snapshots to the user, the snapshots and its clone are no longer controlled by Auto-Sync service.</p>
3	<p>Manage the cloned snapshots</p> <ul style="list-style-type: none"> • Now you can replicate the cloned snapshot. <p>Auto-Sync still continues to run, creating and deleting according to the retention policy independent of the snapshot owned by the user.</p>
4	<p>Reverting ownership of the snapshot</p> <ul style="list-style-type: none"> • Now you may return the snapshot to be managed by Auto-Sync by reverting the ownership or delete the snapshot and its associated clones. <p>To revert the ownership, see Cloning Auto-Sync Snapshots</p>

About Auto-Sync Replication in the HA Cluster Environment

NexentaStor supports Auto-Sync replication between an HA Cluster and a non-HA Cluster. Auto-Sync replication is also possible between two NexentaStor HA Clusters.

You can configure Auto-Sync replication for the following:

- From a dataset under the HA Cluster control to a dataset on a standalone NexentaStor appliance
- From a dataset on a standalone NexentaStor appliance to a dataset under the HA Cluster control
- From a dataset under the HA Cluster control to another dataset under an HA Cluster control

Configuration Settings for Replication Between non-HA and HA Cluster

Setting Auto-Sync to replicate data from a non-HA NexentaStor appliance to a shared volume under cluster control requires additional configuration.

Note: During the failover, all running Auto-Sync services must be stopped.

You must configure the following:

- Bind the non-HA NexentaStor appliance to both nodes of the HA Cluster.
- Create ssh-bindings between the non-HA NexentaStor appliance and each of the HA Cluster nodes.
- Bind the non-HA appliance to the virtual IP address (VIP) of the shared volume.

When one of the HA nodes becomes unavailable, the shared volume service fails over to the healthy node and remains available through VIP. Therefore, Auto-Sync does not interrupt replication.

Do not bind the non-HA appliance to the IP address of the HA Cluster node. The shared volume will become unavailable when the node goes down. Therefore, the Auto-Sync replication will be interrupted.

- Select a shared volume or a folder on a shared volume as replication destination.

When you bind an appliance through VIP, you must select a shared volume as destination. Local volumes are not available through VIP.

Note: When you create an Auto-Sync service to establish a replication schedule to and from clustered appliances, management components are stored on the appliance on which you create the service. Therefore, it is recommended that you create the Auto-Sync service on the secondary node as a Remote to Local service. In that case, management components are stored on the secondary node. This will significantly simplify recovery operations in case of a failure on primary node.

Configuration Settings for Replication Between HA Cluster and HA Cluster

Setting Auto-Sync to replicate data from a HA Cluster to a shared volume under cluster control requires additional configuration.

Configure the following:

- Create SSH binding from each Primary node to all the Secondary nodes and vice versa
- The environment shown in the figure below requires a total of four SSH bindings
 - Primary-Node 1 to Secondary-Node 1
 - Primary-Node 1 to Secondary-Node 2
 - Primary-Node 2 to Secondary-Node 1

- Primary-Node 2 to Secondary-Node 2
- SSH bindings must be configured to the physical IP addresses not VIPs. However, Auto-Sync datapath must be established between VIPs of each HA Cluster

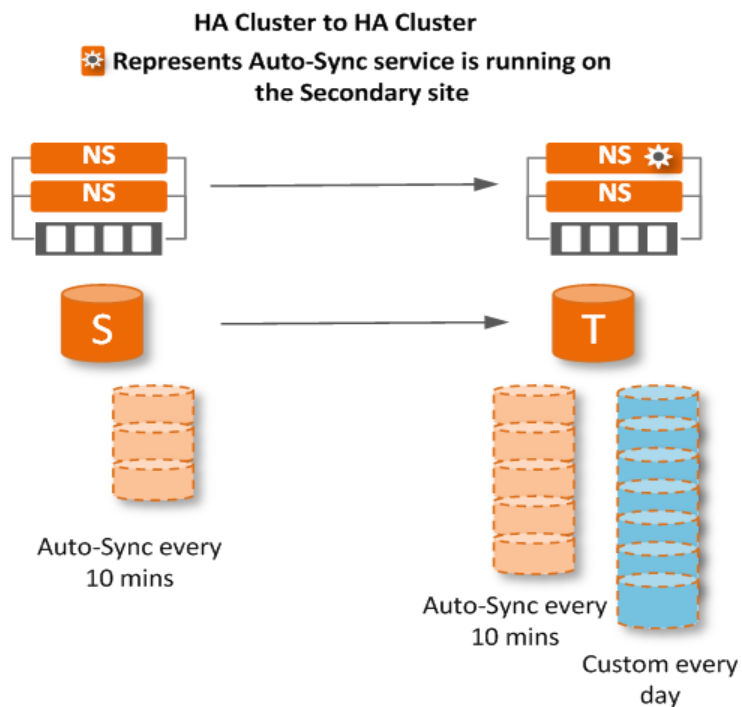
The following example shows how to leverage the Take Ownership feature to manage the daily snapshots extracted from the Auto-Sync snaps.

Example:

- Keep 3 '10 min' snapshots at Primary site
- Keep 5 '10 min' snapshots at the Secondary site
- Replicate every 10 minutes between Primary HA cluster and the Secondary cluster.
- Save daily snapshot at the Secondary site and use longer term retention policy at the Secondary site.

You can use Take Ownership feature on the Auto-Snaps at the Secondary site to set up cron jobs as needed.

Figure 3-6: Example of daily snapshot at the Secondary site.



-
- Note:**
- If the bindings are configured properly, AutoSync service between HA cluster will persist across HA failover at the Primary and HA failover at the Secondary.
 - If an HA failover occurs while a snapshot is being transferred, the transfer is aborted and AutoSync will pick up the next snapshot instance.
- In this scenario RPO will be 2x the snapshot schedule interval.
-

Configuration settings for Replication from HA Cluster to Single Node

- ❖ *To create an Auto-Sync service in this HA Cluster environment:*
 1. Setup SSH bindings between both nodes at primary site and the single node at the Secondary site
 2. Run AutoSync from the Secondary site
 3. Setup AutoSync between Primary and Secondary dataset using VIP on the Primary HA cluster
 4. Setup AutoSync schedules and retention policies.

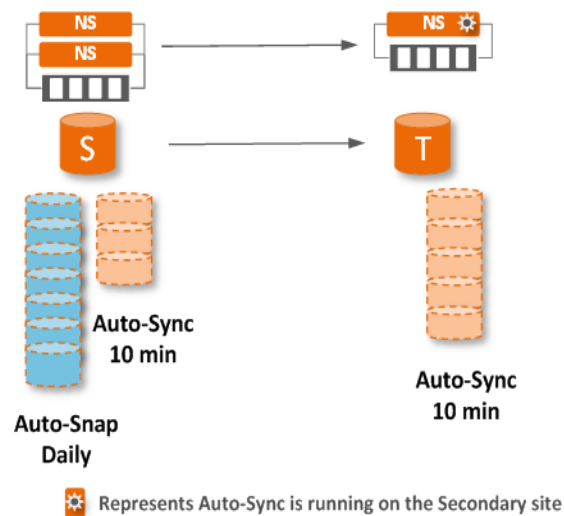
The following example shows how to run Auto-Snap service on the Primary appliance along with the Auto-Sync service in this HA cluster environment.

Example:

- Keep 3 '10 min' snapshots at Primary site
- Keep 5 '10 min' snapshots at the Secondary site
- Replicate every 10 minutes between HA cluster and the single node on the secondary site.
- Setup Auto-Snap on the Primary appliance with the daily schedule to keep 5 snapshots.

You can use Take Ownership feature on the Auto-Snaps at the Primary site or take your own snapshots as needed.

Figure 3-7: Example of Auto-Snap service running at the Primary site.



See Also:

- [Establishing the SSH-binding for an Auto-Sync Service](#)
- [Binding a NexentaStor Appliance to a Virtual IP Address](#)

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

Note:

For disaster recovery to and from HA Cluster configurations, set up remote-to-local (R2L) Auto-Sync service on the NexentaStor appliance(s) that is planned to be used as a disaster recovery site which is the Secondary site. Use the flip-direction command to replicate data from Secondary site (disaster recovery site) to Primary site.

Advanced Settings

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

Table 4-1: Auto-Sync advanced settings.

Property name in NMV	Property name in NMC	Description
Misc. ZFS Flags	<code>general-flags</code>	You can set up different flags for the Auto-Sync service.
Deduplicate Stream	<code>deduplicate_stream</code>	Use deduplication of replication stream. All duplicated data inside one Auto-Sync iteration will be transferred only once. This option will consume CPU power, but will lower utilization of network bandwidth. Will work if both ZFS systems (appliances) support this feature.
Do not Copy ZFS properties	<code>do_not_copy_zfs_properties</code>	Excludes ZFS properties from the replication.
	<code>unmap_zvols</code>	Automatically unmaps zvols.
	<code>skip_readonly_check</code>	Skip verification of the read-only property settings at Secondary.
	<code>daemon</code>	Run the Auto-Sync service in background.
	<code>schedule_once</code>	Run the Auto-Sync service once according to schedule, then change the service state to maintenance.
Use deduplication	<code>Deduplicate Replication Stream</code>	Deduplicate the replication stream. Deduplication consumes additional CPU resources. However, it reduces bandwidth usage. Deduplication is disabled by default.
Use copy properties	<code>copy_zfs</code>	Copy ZFS properties to Secondary dataset.
Trace level	<code>trace_level</code>	Show verbose information for service while running. Default value is 10. Bigger value gives more verbosity. For example 10, 20, 30.
Rate Limit	<code>rate_limit</code>	Limit of network traffic in Kb/s that Auto-Sync uses for replication. Default value is 0.
Snapshot Name Prefix	<code>marker_name</code>	Custom name prefix for service snapshots.
Mark latest as	<code>latest-suffix</code>	A custom suffix for the latest snapshot. Default value is <code>latest</code> .

Property name in NMV	Property name in NMC	Description
Fault Reports	<code>nmc command: show faults - displays results for trigger</code>	<ul style="list-style-type: none"> • handle all RRP errors: parse, convert into NMS faults • allow to track failures on per-service basis • check “space constraints” on Primary and Secondary
Fault Reports	<code>show trigger autosync-svc <instance name> faults</code>	<ul style="list-style-type: none"> • shows details for selected autosync instance
Force Flags	<code>force_flags</code>	You set flags that force some Auto-Sync operations.
Force: Destroy same name snapshots	<code>destroy_same_name_snapshots</code>	Enables you to destroy same-name snapshots at the Secondary that are conflicting with the Primary snapshots.
Force: Destroy missing datasets	<code>destroy_missing_datasets</code>	Destroy Secondary datasets that are not present anymore at the Primary.
Force: Do Rollback	<code>rollback</code>	Enables you to rollback the Secondary dataset to the most recently replicated snapshots.
Force: Receive	<code>receive</code>	Executes all operations at Secondary -F flag. Therefore, before executing the receive command, Auto-Sync rolls back the Secondary folder to the latest snapshot.
Force: Destroy All Snapshots	<code>destroy_all_snapshots</code>	Destroys all outdated snapshots, including manually created and foreign ones, according to the retention policy.
Force: Receive Once	<code>receive_once</code>	Force the <code>zfs recv</code> operation at Secondary. Auto-Sync first rolls back the filesystem to the latest snapshot and then performs the <code>receive</code> operation. This option applies only for the next service run. After that, the option discards the setting. Default value is <code>No</code> .

Property name in NMV	Property name in NMC	Description
	<code>non_recursive_zfs_destroy</code>	<p>Auto-Sync stores snapshots at Primary and Secondary 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 Primary dataset includes a multi-layer hierarchy of nested folders, deletion of the snapshots created for this dataset may significantly affect performance.</p> <p>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.</p>
	<code>schedule_manual</code>	The Auto-Sync service does not use the replication schedule. System Administrator runs the service.
Replicate Intermediate Snapshots		<p>Default value of this property is Enabled. You can disable this property to skip the intermediate snapshots and to send only the changes between the replication snapshots with each update. Following list describes the two options.</p> <ul style="list-style-type: none"> disabled: generates 'zfs send' incremental stream from the first snapshot (the incremental Primary) to the second snapshot (the incremental Secondary). enabled: generates 'zfs send' stream that includes ALL intermediate snapshots (including user- and autosnap-created) between the first and the second snapshot
Exclude folders	<code>exclude</code>	Do not replicate the folders that match a pattern or a list of patterns. Use standard Unix pattern syntax for name filtering. By default, Auto-Sync replicates all files and folders.
Middle Buffer Size	<code>mbuffer_size</code>	<p>The size of middle buffer. Type a value in the range of 16-1024 MB.</p> <p>Middle buffer is memory used by RR daemon to advance the data transfer. Nexenta recommends to use default setting for this property.</p>
Replicate recursively	<code>sync-recursive</code>	Include nested folders and their subfolder into replication.

Property name in NMV	Property name in NMC	Description
TCP Connections	nconn	Specifies the number of TCP connections per data path. You can use one or many network interfaces and assign multiple TCP connections. So for multiple network interfaces you can use up to 16 TCP connections for each network interface. Multiple network interfaces and TCP connections enable you to spread out the replication workload and get better network bandwidth.
RR command line options	options	<p>Command line options for Remote Replication protocol include:</p> <ul style="list-style-type: none"> -n — number of TCP connections. Default value is 4. -P — maximum size of a network packet for replication (PDU). Default value is 1024 kB. <p>Changing the RR command line option is not recommended unless you have sufficient information about the values to provide for these settings.</p>
Comment	comment	Descriptive commentary.
Volume Management (HA configuration)		<p>Force termination of all local running NMS services.</p> <p>Session termination logic for both the Primary and the Secondary side on local and remote auto-sync services.</p> <ul style="list-style-type: none"> Terminate outgoing sessions (rrmgr) Terminate incoming sessions (rrdaemon)
Advanced Properties (available in NMC only)		
	t trunk	<p>Enables you to specify IP addresses or host names for TCP trunking. This functionality requires more than one network interface present on the Primary and Secondary appliance. This property enables Auto-Sync to distribute the network workload 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	Create Auto-Sync service instance using a snapshot as a Primary dataset.

Property name in NMV	Property name in NMC	Description
	marker_name	Marker name is an identifier that declares a type of snapshot. The default marker name is AutoSync. Example: datapool/accounts@AutoSync-1_2012-12-04-09-00-11

Glossary

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.

Service Management Facility (SMF)

A component of the underlying Illumos 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

451 El Camino Real, Suite 201
Santa Clara, CA 95050
USA

3000-at_sync-4.0.5-000058-A