

NetBackup™ Access Appliance 8.6 Solutions Guide for NetBackup

Access Appliance Solutions Guide for NetBackup

Last updated: 2026-04-13

Legal Notice

Copyright © 2026 Cohesity All rights reserved.

© 2026 Cohesity All Rights Reserved. Veritas, the Veritas Logo and other Veritas Marks are trademarks of Cohesity in the US and/or internationally. The information supplied herein is the confidential and proprietary information of Veritas and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Veritas software and services. Find the terms of Veritas licenses at www.cohesity.com/agreements.

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

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to [Cohesity Support](#), to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the [Cohesity Support Portal](#) to create a new case.
- Click the (?) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the [Cohesity Knowledge Base](#).
- Log in to the [Cohesity Support Portal](#) to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Contents

Chapter 1	Access Appliance integration with NetBackup	
		7
	About Access Appliance as a NetBackup client	7
	About Access Appliance as backup storage for NetBackup	8
	Use cases for long-term data retention	8
Chapter 2	System requirements	11
	Supported configurations and versions for NetBackup with Veritas Data Deduplication	11
	Supported configurations and versions for NetBackup with MSDP-C	12
Chapter 3	Configuring Veritas Data Deduplication with Access Appliance	13
	About Veritas Data Deduplication	14
	Benefits of using Veritas Data Deduplication with Access Appliance	16
	Supported features with different versions of Veritas Data Deduplication	16
	Configuring Veritas Data Deduplication without WORM	17
	Configuring Veritas Data Deduplication using the GUI	18
	Configuring Veritas Data Deduplication using the Access CLISH	23
	Configuring Veritas Data Deduplication with WORM	27
	Configuring Veritas Data Deduplication with WORM using GUI	31
	Configuring Veritas Data Deduplication with WORM using the Access CLISH	34
	Managing Veritas Data Deduplication using the Access CLISH	39
	Managing Veritas Data Deduplication using GUI	48
	Viewing information about Veritas Data Deduplication	48
	Starting or stopping the Veritas Data Deduplication service	49
	Increasing storage for Veritas Data Deduplication	49
	Unconfiguring Veritas Data Deduplication	49

	Adding secondary IP for Veritas Data Deduplication	50
	Deleting secondary IP for Veritas Data Deduplication	51
	Changing the primary IP for Veritas Data Deduplication	52
	Upgrade the deduplication version for Veritas Data Deduplication	53
	Setting affinity node for Veritas Data Deduplication	54
	Accessing Access Appliance storage shell for management tasks	55
	About the Access Appliance storage server shell	55
	Support for multiple domains across networks for Veritas Data Deduplication	74
	Veritas Data Deduplication storage layout	79
	Configuring a Veritas Data Deduplication storage unit on NetBackup	79
	Configuring global deduplication using the Veritas Data Deduplication storage server across the domain	89
	Enabling Veritas Data Deduplication encryption	89
	Support for bandwidth throttling during duplication	91
	Setting up secure communication between Veritas Data Deduplication on Access Appliance and the NetBackup primary server	91
	About MSDP encryption using NetBackup KMS service	92
	Support for NetBackup Auto Image Replication	94
	Separating the duplication and replication network	110
	Setting up Auto Image Replication (AIR) between two Access Appliances which have secure communication enabled	113
Chapter 4	NetBackup Dedupe Direct for Oracle	115
	Support for NetBackup Dedupe Direct for Oracle	115
Chapter 5	Configuring MSDP-C with Access Appliance	117
	Starting the Access S3 server and provisioning the S3 bucket	117
	Configuring Access S3 bucket with NetBackup MSDP-C	122
Chapter 6	Migrating the NetBackup images from existing storage to Veritas Access storage	136
	Migrating NetBackup images from CloudCatalyst/S3/MSDP-C to Veritas Data Deduplication storage	136
	Migrating the NetBackup images from the Veritas Data Deduplication STU to the MSDP-C STU	140

Chapter 7	Configuring Access Appliance with the NetBackup client	141
	Configuring Access Appliance for NetBackup client	142
	Installing the NetBackup client add-on packages	143
	Prerequisites for configuring the NetBackup client	145
	Configuring the NetBackup client	145
	Adding an optional media server	146
	Adding multiple primary servers	147
	Adding file systems to the backup configuration	148
	Removing file systems from backup configuration	148
	Changing the virtual IP address used by the NetBackup client	149
	Upgrade the NetBackup client	149
	Unconfiguring the NetBackup client	150
	Enabling Veritas Data Deduplication catalog backup with NetBackup client	150
	Disabling Veritas Data Deduplication catalog backup from NetBackup client	154
	Displaying the status of NetBackup services	155
	Configuring backup operations using NetBackup or other third-party backup applications	156
	Restoring filesystems backed up with NetBackup client	156
Chapter 8	Configuring isolated recovery environment (IRE)	158
	Requirements	158
	Configuring the network isolation	159
	Configuring an isolated recovery environment using the command line	161
	Configuring an isolated recovery environment on a storage server	162
	Managing an isolated recovery environment on a storage server	165
	Configuring data transmission between a production environment and an IRE storage server	168
Chapter 9	Troubleshooting	172
	Additional resources	172
	Generating Access Appliance S3 server keys using the helper script	172
Index		174

Access Appliance integration with NetBackup

This chapter includes the following topics:

- [About Access Appliance as a NetBackup client](#)
- [About Access Appliance as backup storage for NetBackup](#)
- [Use cases for long-term data retention](#)

About Access Appliance as a NetBackup client


Access Appliance is integrated with Veritas NetBackup so that a NetBackup administrator can back up your Access Appliance file systems to NetBackup primary or media servers and retain the data as per your company policy. Once data is backed up, a storage administrator can delete unwanted data from Access Appliance. The NetBackup primary and media servers that run on separate computers from Access Appliance are licensed separately from Access Appliance.

You can configure NetBackup domain information using the command-line interface (CLI)

The Veritas Access command-line interface has a dedicated `Backup>` menu. From the `Backup>` menu, register the NetBackup client with the NetBackup domain. Information is saved in the `bp.conf` file.

See [“Configuring Access Appliance for NetBackup client”](#) on page 142.

Consolidating storage reduces the administrative overhead of backing up and restoring many separate file systems. Critical file data can be backed up and restored through the NetBackup client on Access Appliance.

Figure 1-1 Backing up Access Appliance using NetBackup

About Access Appliance as backup storage for NetBackup

This document describes how Access Appliance fulfills the needs of NetBackup customers looking for a cost-effective solution for moving away from tape backups, yet retain the backed-up data for the long term.

NetBackup is an enterprise-class heterogeneous backup and recovery application. It provides cross-platform backup functionality to a large variety of Windows, UNIX, and Linux operating systems.

Access Appliance is based on the rock-solid and industry-proven Veritas CFS stack. It offers an AWS-compatible S3 protocol as object storage for NetBackup. Access Appliance also provides deduplication storage for NetBackup by using the NetBackup MSDP.

Use cases for long-term data retention


The following are the use cases for long-term data retention (LTR) using Access Appliance:

Use Case: Access Appliance with Veritas Data Deduplication

- Veritas Data Deduplication technology is installed on top of Access Appliance and integrates with NetBackup. It catalogs and organizes incoming deduplicated backup data and stores it on Access Appliance storage.
- MSDP deduplicates the primary backup data and stores it on the NetBackup server.
- The same deduplicated data is sent to Access Appliance for long-term retention purpose using SLPs through duplication jobs.

- Veritas Data Deduplication can also act as a source or destination while replicating deduplicated data using Auto Image Replication (AIR).

Figure 1-2 Access Appliance Veritas Data Deduplication with NetBackup



Note: Veritas supports a deduplication pool of size upto 2790822931980000 bytes (approximately 2.478 PiB) if the Veritas Data Deduplication version is 20.4.0.1 or later. If the Veritas Data Deduplication version is 19.0.1, you cannot create a deduplication pool of size greater than 1392125484720000 bytes (approximately 1.236 PiB).

If the Veritas Data Deduplication version is 19.0.1 and you want to use 2.478 PiB of storage, the following configurations are possible:


- Configure two instances of Veritas Data Deduplication with 1.236 PiB (approximately) storage each.
- Configure MSDP-C with four buckets of size 600 TiB each.
- Configure Veritas Data Deduplication with 1.236 PiB (approximately) storage and configure MSDP-C with two buckets of size 600 TiB each.

Use the virtual IP of one node for Veritas Data Deduplication and use the virtual IP of another node for MSDP-C.

Use Case: Access Appliance with MSDP-C

- MSDP deduplicates the primary backup data and stores it on the NetBackup server.
- The same deduplicated data is moved to Access Appliance through SLP using MSDP-C.


See [“Configuring Access S3 bucket with NetBackup MSDP-C”](#) on page 122.

Figure 1-3 NetBackup MSDP-C over Access Appliance Object Access protocol (S3)**Use Case: Access Appliance as an S3 connector**

- Backup data is stored in non-deduplicated format on the NetBackup server as primary backup. The same primary backup data is moved to Access Appliance through SLP over the Access Appliance S3 protocol.

Or

- MSDP duplicates the primary backup data and stores it on the NetBackup server.
- The deduplicated data is rehydrated and then moved to Access Appliance through SLP over the Access Appliance S3 protocol.

Figure 1-4 Backup data stored in non-deduplicated format on the NetBackup server as primary backup and the same backup data duplicates to Access Appliance over Access S3 protocol.

Note: Configurations and use cases beyond storing copies for long term retention periods require a Support Exception. Please contact your Veritas representative for details on the Support Exception approval process.

System requirements

This chapter includes the following topics:

- [Supported configurations and versions for NetBackup with Veritas Data Deduplication](#)
- [Supported configurations and versions for NetBackup with MSDP-C](#)

Supported configurations and versions for NetBackup with Veritas Data Deduplication

Table 2-1 Supported versions of NetBackup for Veritas Data Deduplication

Veritas Data Deduplication	Veritas NetBackup primary servers	Veritas NetBackup media servers
20.5.0.1	10.5.0.1 and later versions	10.5.0.1 and later versions
21.0.0.1	11.0.0.1 and later versions	11.0.0.1 and later versions
21.1.0.2	11.1.0.2 and later versions	11.1.0.2 and later versions

Note: The version of the media server should be the same as or lower than the version of the primary server.

Note: Veritas recommends that you have at least 0.1 GB of RAM for 1 TB of storage if you are using the PS cache mechanism and at least 0.5 GB of RAM for 1 TB of storage if you are using the FP cache mechanism.

Supported configurations and versions for NetBackup with MSDP-C

Table 2-2 Supported versions of NetBackup for MSDP-C without WORM

Access Appliance	Veritas NetBackup servers
8.6	10.1 and later versions

Table 2-3 Supported versions of NetBackup for MSDP-C with WORM

Access Appliance	Veritas NetBackup servers
8.6	10.1 and later versions

Configuring Veritas Data Deduplication with Access Appliance

This chapter includes the following topics:

- [About Veritas Data Deduplication](#)
- [Benefits of using Veritas Data Deduplication with Access Appliance](#)
- [Supported features with different versions of Veritas Data Deduplication](#)
- [Configuring Veritas Data Deduplication without WORM](#)
- [Configuring Veritas Data Deduplication with WORM](#)
- [Managing Veritas Data Deduplication using the Access CLISH](#)
- [Managing Veritas Data Deduplication using GUI](#)
- [Accessing Access Appliance storage shell for management tasks](#)
- [Support for multiple domains across networks for Veritas Data Deduplication](#)
- [Veritas Data Deduplication storage layout](#)
- [Configuring a Veritas Data Deduplication storage unit on NetBackup](#)
- [Configuring global deduplication using the Veritas Data Deduplication storage server across the domain](#)
- [Enabling Veritas Data Deduplication encryption](#)
- [Support for bandwidth throttling during duplication](#)

- [Setting up secure communication between Veritas Data Deduplication on Access Appliance and the NetBackup primary server](#)
- [About MSDP encryption using NetBackup KMS service](#)
- [Support for NetBackup Auto Image Replication](#)
- [Setting up Auto Image Replication \(AIR\) between two Access Appliances which have secure communication enabled](#)

About Veritas Data Deduplication

Access Appliance is integrated with a duplication engine which is based on Media Server Deduplication Pool (MSDP) technology for storing backup data. The storage server component of Veritas Data Deduplication runs on the Access Appliance nodes with high availability in active/passive mode. The deduplication plug-in of the NetBackup media server does segmentation and finger printing of the backup data and sends the deduplicated data to Access Appliance. The Veritas Data Deduplication storage server stores and manages the deduplicated data. The deduplication storage server provides high availability to protect against storage, node, and network failures. It supports client direct as well as media server deduplication configurations.

Considerations before configuring Veritas Data Deduplication:

- All storage that is provisioned for Veritas Data Deduplication is displayed as a single storage pool on NetBackup.
- Access Appliance supports a single deduplication pool of size upto 2.478 PiB (approximately) if the Veritas Data Deduplication version is 20.4.0.1 or later. If the Veritas Data Deduplication version is 19.0.1, you can create two storage pools with up to 2.478 PiB (approximately) capacity wherein each storage pool has a capacity of 1.236 PiB (approximately).
- The import operation may fail if you export the Veritas Data Deduplication configuration (on Access Appliance 8.0 or previous versions) and try to import it after upgrading to Access Appliance 8.6 version.
- Maximum cache size is set based on the storage consumption of the container and the number of containers running on that node. It also considers the affinity of the containers running on that node. The maximum cache size is managed internally by the Access Appliance and the user is not required to modify it. During operations such as `dedupe shrink` and `dedupe grow`, this value is tuned as required.
- 500 GB of system memory is required for every 1 PB of storage when you configure the Veritas Data Deduplication server. If you want to configure 1.24

PiB storage, you require 768 GB of system memory as additional system memory kit can be purchased in 386 GB increments. You can install additional memory and retry the configuration. This also applies to upgrade and `dedupe grow` operation.

- Access Appliance supports multiple versions of Veritas Data Deduplication. You can configure Veritas Data Deduplication version of your choice, which is compatible with the NetBackup primary server version in the domain. You can upgrade from a lower version to a higher version of Veritas Data Deduplication without loss of data. For supported versions, See [“Supported configurations and versions for NetBackup with Veritas Data Deduplication”](#) on page 11.
- If you configure Veritas Data Deduplication version 19.0.1, the Finger Print (FP) cache mechanism is used to allocate the cache size. If you configure Veritas Data Deduplication version 20.4.0.1 or later, the Predictive and Sampling cache mechanism (PS) cache mechanism is used to allocate the cache size.
- If you have upgraded to Access Appliance 8.68.6 from an earlier version, the previously configured Veritas Data Deduplication instances continue to use the FP cache mechanism. In order to leverage the 2.478 PiB storage, the type of cache mechanism of the Veritas Data Deduplication container instance should be converted to PS cache, Veritas strongly recommends that the operation be performed under the supervision of Veritas Technical Support.
- If you have newly configured an instance of Veritas Data Deduplication version 20.4.0.1 on Access Appliance version 8.3.100 and then upgraded the appliance to version 8.6, then the type of cache mechanism of the Veritas Data Deduplication container instance automatically gets converted to PS cache.

To convert the cache mechanism of the Veritas Data Deduplication container instance to FP, perform the following steps:

- After the upgrade is completed, login to the MSDP container shell of version 20.4.0.1 using the following command:

```
podman exec -it <container_name> bash
```

- Execute the following commands to disable PS cache inside the MSDP container.

```
/opt/veritas/vxapp-manage/pdde_cmds tune --write  
--enable-local-predictive-sampling-cache false
```

- Write the sample FP value in terms of bytes as follows:

```
/opt/veritas/vxapp-manage/pdde_cmds tune --write  
--max-cache-size "1GiB"
```

- Restart the Veritas Data Deduplication instance using the following commands:

```
dedupe stop <config_name>
dedupe start <config_name>
```

Benefits of using Veritas Data Deduplication with Access Appliance

- Capability of writing the deduplicated data from the media server to Access Appliance.
- Support global deduplication across media servers within the same domain and across different NetBackup domains.
- The storage server components reside directly on the Access Appliance nodes with high availability in active/passive mode to protect against storage, node, and network failures.
- Includes WORM storage that provides immutable and indelible data protection, ensuring data cannot be changed for a determined length of time to protect data against cyber criminal intrusion and internal threats.
- All software meet stringent security standards and are Federal information Processing Standard (FIPS) 140-2–compliant.

Supported features with different versions of Veritas Data Deduplication

Access Appliance supports multiple versions of Veritas Data Deduplication. The [Table 3-1](#) lists the features that are supported with a particular version of Veritas Data Deduplication.

Table 3-1 Supported features for different Veritas Data Deduplication versions

Feature	Veritas Data Deduplication 20.5.0.1	Veritas Data Deduplication 21.0.0.1	Veritas Data Deduplication 21.1.0.2
WORM configuration	Yes	Yes	Yes

Table 3-1 Supported features for different Veritas Data Deduplication versions (*continued*)

Feature	Veritas Data Deduplication 20.5.0.1	Veritas Data Deduplication 21.0.0.1	Veritas Data Deduplication 21.1.0.2
Non-WORM configuration	Yes	Yes	Yes
Multifactor authentication	Yes	Yes	Yes
Changing the primary IP	Yes	Yes	Yes
Removing user	Yes	Yes	Yes
KMS	Yes	Yes	Yes
Secure communication	Yes	Yes	Yes
Encryption	Yes	Yes	Yes
Encryption crawler	Yes	Yes	Yes
Conversion from non-WORM to WORM	Yes	Yes	Yes
Shrink operation	Yes	Yes	Yes
Affinity	Yes	Yes	Yes
IRE	Yes	Yes	Yes

Configuring Veritas Data Deduplication without WORM

You can configure Veritas Data Deduplication without WORM using the Access Appliance GUI as well as command-line interface. Veritas recommends that you configure Veritas Data Deduplication using the GUI as it ensures that the configuration is performed with the appropriate filesystem size, and volume layout along with catalog filesystem creation.

You can configure Veritas Data Deduplication without WORM when the lockdown mode is Normal, Enterprise or Compliance.

For more information, see the *About lockdown modes* section in the *Access Appliance Administrator's Guide*.


Note: When you provision storage using the GUI or the command-line interface, specify the storage pool size and the size configured for Veritas Data Deduplication in MiB, GiB, TiB, or PiB. You can also specify the input size in decimal values.

Configuring Veritas Data Deduplication using the GUI


This section describes how you can configure Veritas Data Deduplication using the Access Appliance GUI.

To configure Veritas Data Deduplication


- 1 Log on to the **Access Appliance** application.
 - Select the **Quick Actions** icon present on the dashboard page. Navigate to **Provision Storage > Storage for Netbackup > Veritas Data Deduplication**. Click **Next**. You have to use this option if any instance of Veritas Data Deduplication is already present on the cluster.
 - Or, navigate to **Settings > Veritas Data Deduplication Management**. Click **Configure**. This option is available only if no Veritas Data Deduplication instances are present on the cluster.
- 2 A default storage pool is created at the time of initial configuration with all the available disks.



- 3 Under the **Activate Policy** tab, select the Veritas Data Deduplication policy.
Click **Complete prerequisites**.



- 4 Complete all the prerequisites such as, selecting the storage pool to activate the Veritas Data Deduplication policy and click **Activate policy**.



After the policy is activated, click **Finish**.

5 Under the **Provision storage** tab, specify all the required fields.

Select storage policy	From the dropdown menu, choose <code>Veritas Data Deduplication policy</code> .
Storage size	Mention the capacity for Veritas Data Deduplication storage. If the Veritas Data Deduplication version is 19.0.1 or lower, the size should be less than 1.236 PiB (approximately). If the Veritas Data Deduplication version is 20.4.0.1 or later, the size should be less than 2.478 PiB (approximately). Note: Reducing the storage for Veritas Data Deduplication is not supported using GUI. You can reduce storage for Veritas Data Deduplication using CLISH. See “Managing Veritas Data Deduplication using the Access CLISH” on page 39. for more details.
Username	Specify the username to be used
Password/Confirm Password:	Enter the password for the above username
VIP (Virtual IP address)	From the drop-down list, select any virtual IP.
Node	Name of the node on which deduplication service is to be configured (default is <i>None</i> which means any node can be configured)
Configuration name	Name for the deduplication service. Default value is <code>msdp-engine</code>. Note: This is an optional parameter when the first deduplication server is configured. It is mandatory when you configure the second deduplication server.
Dedupe version (NBU version)	Select the required version of Veritas Data Deduplication server (NetBackup version) from the drop-down list. This is a mandatory parameter.
Enable encryption	Select this check box to enable encryption.

Provision storage

Veritas Data Deduplication

Configure storage Activate policy Provision storage

Provision storage

Select storage policy
Veritas Data Deduplication

Storage size *
239.00 GB

Username *
Enter username for NetBackup deduplication engine.

Password *
Enter password.

Confirm password *
Enter password.

Dedupe Version (NBU version) *
17.1 (NB-10.1.1)
19.0.1 (NB-10.3.0.1)

VIP (Virtual IP Address) *

Node

Configuration name
Enter the configuration name

Enable encryption ☒

Provision storage

Click **Provision storage**.

If you configure a second deduplication server, then you have to specify the name of the deduplication service in the `Configuration name` field.

Veritas Access Appliance 8.2 - palak

Provision storage

Veritas Data Deduplication

Configure storage Activate policy Provision storage

Provision storage

Select storage policy
Veritas Data Deduplication

Storage size *
1 GB

Username *
dedupe1

Password *
.....

Confirm password *
.....

Dedupe Version (NBU version) *
17.1 (NB-10.1.1)

VIP (Virtual IP Address) *

Node

Configuration name *
Enter the configuration name

Enable encryption ☒

View the **Recent Activity** panel for the status of the task.

The storage provisioning is now complete.

An alert appears in the GUI when the memory usage of the Veritas Data Deduplication server is more than 80%. The alert displays the amount of memory used and the total RAM usage. If the situation persists for more than 24 hours, an AutoSupport alert is raised and an email is sent.

If you receive this alert, go to **Settings > Veritas Data Deduplication management** and restart the service once you reach this threshold. For information about the issues that are caused when the Veritas Data Deduplication service consumes large amount of memory, see

https://www.veritas.com/support/en_US/article.100045745.

Configuring Veritas Data Deduplication using the Access CLISH

This section describes how you can configure Veritas Data Deduplication without WORM using the Access Appliance command-line interface.

You can use the `dedupe> config` command to configure the deduplication storage server using the given set of filesystems and IP address.

Prerequisites:

- The filesystem should be striped and should be already created with size not exceeding 100 TiB and the IP should be online before you start configuring Veritas Data Deduplication.
- If the total deduplication pool size required is larger than 100 TiB, then create multiple data filesystems.
- If the Veritas Data Deduplication version is 19.0.1, the total data filesystem size should not exceed 1.236 PiB (approximately). If the Veritas Data Deduplication version is 20.4.0.1 or later, the total data filesystem size should not exceed 2.478 PiB (approximately).
- The data filesystems are used to store the deduplicated data and its metadata.
- You are required to create a separate catalog file system to store the Veritas Data Deduplication catalog. The minimum size of the catalog file system should be 5 TiB.
- The deduplication storage server uses the specified IP address to receive the backup data from NetBackup.

To configure the deduplication storage server

- ◆ To configure the first deduplication storage server, enter the following:

```
dedupe> config filesystem1[,filesystem2,...] IP username
[catfs_name] vdd_version [encryption] [worm=no] [minret] [maxret]
[node_name] [config_name]
```

Where:

<i>filesystem1,</i> <i>filesystem2,...</i>	Specifies the data filesystem names.
<i>IP</i>	Specifies the interface IP. You can check for the IP using the <code>network IP address show</code> command.
<i>username</i>	Specifies the username.
<i>catfs_name</i>	Specifies the catalog filesystem of the deduplication server.
<i>vdd_version</i>	Specifies the version of Veritas Data Deduplication sever. This is a mandatory parameter. For supported version, See “Supported configurations and versions for NetBackup with Veritas Data Deduplication” on page 11.
<i>encryption</i>	Enables encryption. The possible values are <code>enable_encryption=yes/enable_encryption=no</code> . If enabled, all data sent to the deduplication server is encrypted but without KMS.
<i>worm</i>	Enables WORM. The possible values are <code>enable_worm=yes/enable_worm=no</code> . This is an optional parameter.
<i>minret</i>	Specifies the minimum image retention period. This is an optional parameter. If you want to specify the <i>config_name</i> parameter, then you should pass the value, <i>vxdefault</i> as the minimum retention value.
<i>maxret</i>	Specifies the maximum image retention period. This is an optional parameter. If you want to specify the <i>config_name</i> parameter, then you should pass the value, <i>vxdefault</i> as the maximum retention value.

<i>node_name</i>	Name of the node on which deduplication service is to be configured. If <i>node_name</i> is provided, the affinity of the server is set to that node. If <i>node_name</i> is not provided, and if there is only one container in the cluster, the affinity is not set to any node. After the second container comes into existence, irrespective of whether affinity (<i>node_name</i>) is provided or not, the affinity is set to the node other than the node where the first container is running, as it does not have any running container with storage consumed. If you want to specify the <i>config_name</i> parameter, then you should pass the value, <i>vxdefault</i> as the node name.
<i>config_name</i>	Name for the deduplication service. Default value is <i>msdp-engine</i>. Note: This is an optional parameter when the first deduplication server is configured. It is mandatory when you configure the second deduplication server.

You are prompted to enter your password. After you specify your password, Veritas Data Deduplication is configured.

If you configure a second deduplication server, then you have to specify the *config_name* parameter and specify the maximum and minimum retention values as *vxdefault*. For example:

```
dedupe> config filesystem1[,filesystem2,...] IP username
[catfs_name] vdd_version [encryption] [worm=no] [minret=vxdefault]
[maxret=vxdefault] [node_name=vxdefault] [config_name]
```

Note: Make a note of the password that you use. You have to use the same password when you configure the storage server on NetBackup and also when you reconfigure Veritas Data Deduplication.

Note: Configuring Veritas Data Deduplication requires an FQDN entry which resolves to a specific IP. You can either have a DNS entry or you can add an entry in the `/etc/hosts` file of both the NetBackup media and primary servers on all the nodes to resolve the FQDN to a specific IP.

Note: If you have to reconfigure deduplication, you have to pass all the filesystems in the same order that is displayed in the output of the `dedupe show` command before deduplication was unconfigured.

If a new IP is used for the deduplication server during reconfiguration, the new IP should map to the old FQDN. You can either have a DNS entry or add an entry in the `/etc/hosts` file on all the nodes to resolve the FQDN to a specific IP.

Configuring Veritas Data Deduplication with WORM

You can configure Veritas Data Deduplication with WORM using both the Access Appliance GUI and command-line interface. Veritas recommends that you configure Veritas Data Deduplication using the GUI as it ensures that the configuration is performed with the appropriate filesystem size, and volume layout along with catalog filesystem creation.

You can configure Veritas Data Deduplication with WORM when the lockdown mode is Enterprise or Compliance.

For more information, see the *About lockdown modes* section in the *Access Appliance Administrator's Guide*.

Note: When you provision storage using the GUI or the command-line interface, specify the storage pool size and the size configured for Veritas Data Deduplication in MiB, GiB, TiB, or PiB. You can also specify the input size in decimal values.

Support for multiple Veritas Data Deduplication pools

Two Veritas Data Deduplication pools are supported on Access Appliance. The pools can be any combination of WORM and non-WORM instances. Veritas supports converting an existing Veritas Data Deduplication pool (without WORM) to a Veritas Data Deduplication pool with WORM.

Enabling Veritas Data Deduplication with WORM

You can enable Veritas Data Deduplication with WORM using the following command:

```
# dedupe enable-worm <minret> <maxret> <config_name>
```

minret

Specifies the minimum image retention period. It can contain: [1-9](s|S).

<i>maxret</i>	Specifies the maximum image retention period. It can contain: [1-9](s S).
<i>config_name</i>	Name for the deduplication service. Default value is <i>msdp-engine</i> . Note: This is an optional parameter when the first deduplication server is configured. It is mandatory when you configure the second deduplication server.

Migrating non-WORM Veritas Data Deduplication to a WORM capable Veritas Data Deduplication

You can change the lockdown mode from Normal to either Enterprise or Compliance mode. Then, you can change the deduplication pool to a WORM capable pool using the `dedupe enable worm` command. This command converts the existing non-WORM pool to a WORM capable pool, without any loss of data. You must update the storage unit in the NetBackup primary server to enable WORM on it so that the new images getting duplicated to this STU are WORM-enabled. Alternatively, you can continue to use the existing STU for storing non-WORM backups and create a new STU which is WORM-enabled, for storing WORM-enabled backups.

The old images stored in the pool are not WORM-enabled. To make the older images WORM-enabled:


Method 1

1. Duplicate the non-WORM images from the Access Appliance storage pool to another Veritas Data Deduplication or MSDP-C storage pool.
2. Expire these images from the Access Appliance storage pool.
3. Duplicate these images back onto the Access Appliance storage pool using the WORM-enabled storage unit.

Method 2

You can convert non-WORM images to WORM-enabled images using the NetBackup GUI.

1. Update the disk pool in the NetBackup UI to refresh the WORM status of the pool.
2. Create a new storage unit using the same disk pool, and select the **Use WORM** option for the new storage unit.




The image shows a 'New Storage Unit' dialog box with the following fields and options:



- Storage unit name:** worm-ltr-stu
- Storage unit type:** Disk (dropdown menu)
- Disk type:** PureDisk (dropdown menu)
- On demand only:** ☒
- Use WORM:** ☒
- Properties and Server Selection:**
 - Select disk pool:** vddinst2pool (dropdown menu) with a **View Properties** button.
 - Media server:**
 - ☒ Use any available media server
 - ☐ Only use the following media servers
 - Media Servers:** A list box containing one entry, 'Media Servers', which is currently selected.
- Maximum concurrent jobs:** 1 (spin box)
- Maximum fragment size:** 51200 Megabytes

At the bottom are **OK**, **Cancel**, and **Help** buttons.

- Using NetBackup UI, list down all the images in the catalog.



4. Select the **duplicate** option for each image, and select the WORM enabled STU.



Method 3

You can convert non-WORM images to WORM-enabled images using the NetBackup command line.

1. Run the `bpimagelist` command to list down the images you want to duplicate.
2. Run the `bpduplicate` command to duplicate the images to the WORM STU.

Configuring Veritas Data Deduplication with WORM using GUI

This section describes how you can configure Veritas Data Deduplication with WORM using the Access Appliance GUI.

To configure Veritas Data Deduplication


- 1 Log on to the Access Appliance application.

Select the **Quick Actions** icon present on the dashboard page, and navigate to **Provision Storage > Storage for NetBackup > Veritas Data Deduplication**. Click **Next**. You have to use this option if any instance of Veritas Data Deduplication is already present on the cluster.

Or, navigate to **Settings > Veritas Data Deduplication Management**. Click **Configure**. This option is available only if no Veritas Data Deduplication instances are present on the cluster.

If you want to configure lockdown mode, refer to the *About lockdown modes* section in the *Access Appliance Administrator's Guide*.

- 2 Under the **Activate Policy** tab, select **Veritas Data Deduplication (WORM)** policy and click **Complete prerequisites**. Select the required storage pools.



- 3 Enter the minimum and maximum retention time details and click **Activate policy**.

Veritas Access Appliance 8.2 -

Provision storage

Veritas Data Deduplication

Configure storage Activate policy Provision storage

Complete following prerequisite to activate Veritas Data Deduplication (WORM)

Select storage pool (Selected storage pools: 1, total size: B

Set WORM

Configure the retention time range that Appliance Administrators will be allowed to select from when setting WORM retention time. The minimum retention time you can set is 1 hour and the maximum is 60 years.

Minimum retention time 1.44 Maximum retention time 2.51 Years

Activate policy

Once the Veritas Data Deduplication (WORM) policy is activated, click **Finish**.

- 4** In the **Provision storage** page, select the Veritas Data Deduplication (WORM) policy and specify the storage requirements. Enter the required information.

Select storage policy	From the dropdown menu, choose the relevant Veritas Data Deduplication policy (WORM)
Storage size	Mention the capacity for Veritas Data Deduplication storage. If the Veritas Data Deduplication version is 19.0.1 or lower, the size should be less than 1.236 PiB (approximately). If the Veritas Data Deduplication version is 20.4.0.1 or later, the size should be less than 2.478 PiB (approximately).
Username	Specify the username to be used
Password/Confirm Password:	Enter the password for the above username
VIP (Virtual IP address)	From the drop-down list, select any virtual IP.
Node	Name of the node on which deduplication service is to be configured (default is <i>None</i> which means any node can be configured)
Configuration name	Name for the deduplication service. Default value is <i>msdp-engine</i> . Note: This is an optional parameter when the first deduplication server is configured. It is mandatory when you configure the second deduplication server.
Dedupe version (NBU version)	Select the required version of Veritas Data Deduplication server (NetBackup version) from the drop-down list. This is a mandatory parameter.
Enable encryption	Select this check box to enable encryption.

Click **Provision storage**.

Veritas Access Appliance 8.2 - palak

Provision storage

Veritas Data Deduplication

Configure storage Activate policy Provision storage

Provision storage

Select storage policy
Veritas Data Deduplication (WORM)

Storage size *
234.65 GB

Username *
Enter username for NetBackup deduplication engine.

Password *
Enter password.

Confirm password *
Enter password.

Deduplicate Version (NBU version) *
16.0 (NB-10.0)
17.1 (NB-10.1.1)
19.0.1 (NB-10.3.0.1)

VIP (Virtual IP Address) *
Node

Configuration name *
Enter the configuration name

Enable encryption

Provision storage

View the **Recent Activity** panel for the status of the task.

If you configure a second deduplication server, then you have to specify the name of the deduplication service in the `Configuration name` field.

Provision storage

Veritas Data Deduplication

Configure storage Activate policy Provision storage

Provision storage

Select storage policy
Veritas Data Deduplication (WORM)

Storage size *
1 GB

Username *
dedupe1

Password *
.....

Confirm password *
.....

Deduplicate Version (NBU version) *
16.0 (NB-10.0)
17.1 (NB-10.1.1)
19.0.1 (NB-10.3.0.1)

VIP (Virtual IP Address) *
10.221.39.75

Node

Configuration name
config1

Enable encryption

Provision storage

Configuring Veritas Data Deduplication with WORM using the Access CLISH

This section describes how you can configure Veritas Data Deduplication with WORM using the Access Appliance command-line interface.

You can use the `dedupe> config` command to configure the deduplication storage server using the given set of file systems and IP address.

Prerequisites:

- The filesystem should be striped and should be already created with size not exceeding 100 TiB and the IP should be online before you start configuring Veritas Data Deduplication.
- If the total deduplication pool size required is larger than 100 TiB, then create multiple data filesystems.
- If the Veritas Data Deduplication version is 19.0.1, the total data filesystem size should not exceed 1.236 PiB (approximately). If the Veritas Data Deduplication version is 20.4.0.1 or later, the total data filesystem size should not exceed 2.478 PiB (approximately).
- The data filesystems are used to store the deduplicated data and its metadata.
- You are required to create a separate catalog file system to store the Veritas Data Deduplication catalog. The minimum size of the catalog file system should be 5 TiB.
- The deduplication storage server uses the specified IP address to receive the backup data from NetBackup.
- The lockdown mode should be enterprise or compliance.

To configure the deduplication storage server with WORM

- ◆ To configure the deduplication storage server, enter the following:

```
dedupe> config filesystem1[,filesystem2,...] IP username
[catfs_name] vdd_version [encryption] [worm=yes] [minret] [maxret]
[node_name] [config_name]
```

Where:

<i>filesystem1, filesystem2,...</i>	Specifies the data filesystem names.
<i>IP</i>	Specifies the interface IP.
<i>username</i>	Specifies the user name.
<i>catfs_name</i>	Specifies the catalog file system of the deduplication server.
<i>vdd_version</i>	Specifies the version of Veritas Data Deduplication sever. This is a mandatory parameter. For supported version, See “Supported configurations and versions for NetBackup with Veritas Data Deduplication” on page 11.
<i>encryption</i>	Enables encryption. The possible values are <code>enable_encryption=yes/enable_encryption=no</code> . If enabled, all data sent to the deduplication server is encrypted but without KMS.
<i>worm</i>	Enables WORM. The possible values are <code>enable_worm=yes/enable_worm=no</code> .
<i>minret</i>	Specifies the minimum image retention period in seconds. It can contain: [1-9](s S). The minimum value is 3600s.
<i>maxret</i>	Specifies the maximum image retention period in seconds. It can contain: [1-9](s S). The maximum value is 1892160000s.

<i>node_name</i>	Name of the node on which deduplication service is to be configured. If <i>node_name</i> is provided, the affinity of the server is set to that node. If <i>node_name</i> is not provided, and if there is only one container in the cluster, the affinity is not set to any node. After the second container comes into existence, irrespective of whether affinity (<i>node_name</i>) is provided or not, the affinity is set to the node other than the node where the first container is running, as it does not have any running container with storage consumed.
<i>config_name</i>	Name for the deduplication service. Default value is <i>msdp-engine</i>. Note: This is an optional parameter when the first deduplication server is configured. It is mandatory when you configure the second deduplication server.

You are prompted to enter your password. After you specify your password, Veritas Data Deduplication is configured.

Note: You have to use the same password when you configure the storage server on NetBackup.

Note: Configuring Veritas Data Deduplication requires a FQDN entry which resolves to a specific IP. You can either have a DNS entry or you can add an entry in the `/etc/hosts` file on all the nodes to resolve the FQDN to a specific IP.

If the FQDN entry is not found, `<catfs>.vdd.com` is used as the FQDN entry during configuration.

Note: If you have to reconfigure deduplication, you have to pass all the file systems in the same order that is displayed in the output of the `dedupe show` command.

If a new IP is used for the deduplication server during reconfiguration, the new IP should map to the old FQDN. You can either have a DNS entry or add an entry in the `/etc/hosts` file on all the nodes to resolve the FQDN to a specific IP.

The default size of the stripe unit of the deduplication file system created is 64K. As a result of the creation of catalog file system and VxFS and VDD optimizations, performance improvements can be seen in optimized duplication and restore.

Managing Veritas Data Deduplication using the Access CLISH

This section describes the various operations that you can perform on the deduplication server using the Access CLISH.

To change the user password

- ◆ You can change the password of the deduplication server user.

To change the password, enter the following:

```
dedupe> passwd user_name config_name
```

You are prompted to enter your current password and the new password. After you confirm the new password, the password gets updated.

Note:

- If the user is not present on NetBackup, you have to add the user on the NetBackup server.

```
# /usr/opensv/pdde/pdcr/bin/spauser -a -u user_name
```

- You also must update the password on NetBackup if you have updated the password for the Veritas Data Deduplication server from Access Appliance.

```
# /usr/opensv/volmgr/bin/tpconfig -update -stype PureDisk  
-storage_server IP address -sts_user_id user_name
```

To start the deduplication storage server

- ◆ The storage server is started on the node on which affinity is set. If only one storage server is configured and affinity is not set, the server starts on the node on which the virtual IP, which you specified in the `dedupe config` command is online.

To start the deduplication storage server, enter the following:

```
Dedupe> start config_name
```

The deduplication storage server should be already configured before you start the deduplication storage server.

If `config_name` is not specified, all the deduplication servers that are in stopped state are started.

To display information the deduplication storage server

- ◆ To display information about the deduplication storage server, enter the following:

```
dedupe> show config_name
```

If *config_name* is not specified, the details of all the configured deduplication servers are displayed.

The information includes details such as the file system(s) being used, the IP, and the cluster node on which the server is running, affinity, encryption, shrink status and errors.

To display the deduplication storage server statistics

- ◆ To display the deduplication storage server statistics, enter the following:

```
dedupe> stats config_name
```

The information includes details like total storage size, free storage, and deduplication ratio. The deduplication storage server should be online for this operation.

If the deduplication server is offline, the values of the attributes are displayed as "-". If *config_name* is not specified, the statistics of all the configured deduplication servers are displayed.

To show the status of the deduplication storage server

- ◆ To show the status of the deduplication storage server, enter the following:

```
dedupe> status config_name
```

The information includes details such as whether the server is running and the cluster node on which the server is running.

If *config_name* is not specified, the status of all the configured deduplication servers is displayed.

To stop the deduplication storage server

- ◆ To stop the deduplication storage server, enter the following:

```
dedupe> stop config_name
```

The deduplication storage server must be configured and running. If any backup job is running when the server is stopped, the job aborts with an error.

If *config_name* is not specified, all the deduplication servers are stopped.

To increase the storage capacity of the deduplication storage server

- ◆ To increase the storage capacity of the deduplication storage server to the given size, enter the following:

```
dedupe> grow size config_name
```

The deduplication storage server should be online for this operation.

The *size* parameter specifies the size to which the deduplication storage size is grown.

The *config_name* is the name of the deduplication server whose storage capacity is to be grown. It must be specified if more than one deduplication server is running on the cluster.

To reduce the storage capacity of the deduplication storage server

- ◆ To reduce the storage capacity of the deduplication storage server to the given size, enter the following:

```
dedupe> shrink size config_name data_correct={yes|no}
```

The *size* parameter specifies the size to which the deduplication storage size is shrunk.

The *config_name* parameter must be specified if more than one deduplication server is running on the cluster.

The *data_correct* parameter is set to *no*, by default, for shrink operation. If *data_correct* is set to *yes*, this command only corrects data errors that were caused by earlier shrink operation failures. If *data_correct* operation needs to be performed, the size parameter must be set to *vxdefault*.

The deduplication storage server is brought offline for this operation. The storage server remains offline after the shrink operation if it was offline before the operation. It is brought online after the shrink operation if it was online earlier. Running the `dedupe shrink` command starts a background process to perform the operation. You can view the status using the `dedupe show` command. It is recommended to perform the shrink operation with size equal to the size of the file system.

For example:

If the file system layout is as follows:

```
Fs1 100TB
Fs2 100TB
Fs3 80TB
Fs4 100TB
```

If you want to reduce the deduplication storage capacity from 380 TB to 200 TB, the `dedupe shrink` operation can be done in two steps.

```
Dedupe> shrink 300T msdp-engine data_correct=no
```

```
Dedupe> shrink 200T msdp-engine data_correct=no
```

Caution: The time taken to perform the `dedupe shrink` operation varies based on the size parameter passed to the `dedupe shrink` command. In order to minimize the downtime required during the `dedupe shrink` operation, Veritas strongly recommends that the operation be performed under the supervision of Veritas Technical Support.

Note:

- Do not perform any file system offline operation or any kind of maintenance activity when the `dedupe shrink` command is in progress.
- Backup jobs should be stopped before performing the `dedupe shrink` operation. They can be resumed once the `dedupe shrink` operation is complete.
- If multiple filesystems are configured with deduplication, then NetBackup client backup should be performed after the `dedupe shrink` operation. This is because the `dedupe shrink` operation may destroy the file system(s) and any backup of deduplication filesystems done before the `dedupe shrink` operation may contain stale configuration entries of the deleted filesystem(s) due to which the deduplication service fails to come online after the `dedupe shrink` operation is completed.

To add a new user to the deduplication storage server

- ◆ To add a new user to the deduplication storage server, enter the following:

```
dedupe> adduser username config_name
```

The user name that is mentioned here is the NetBackup media server user name which you use later to configure the storage server on the NetBackup server.

The `config_name` is the name of the deduplication server to which the user is being added. It must be specified if more than one deduplication server is running on the cluster.

To remove an existing user from the deduplication storage server

- ◆ To remove an existing user from the deduplication storage server, enter the following:

```
dedupe> removeuser username config_name
```

The command prompts the user for the password for the given user name.

The *config_name* is the name of the deduplication server to which the user was added. It must be specified if more than one deduplication server is running on the cluster.

Note: You cannot delete a user if the user has backup images present on the deduplication pool or if the user is the primary user (if the username was passed during deduplication service configuration).

To show the list of users of the deduplication storage server

- ◆ To show the list of users of the deduplication storage server, enter the following:

```
dedupe> listuser config_name
```

The *config_name* parameter must be specified if more than one deduplication server is running on the cluster. If *config_name* is not specified, all the users of the configured deduplication servers are listed.

To unconfigure the deduplication storage server

Ensure that you make a note of the following configuration details before you unconfigure the Veritas Data Deduplication server. These details are required when you reconfigure the service.

- The virtual IP that you used to configure the Veritas Data Deduplication server.
- Credentials to access the Veritas Data Deduplication server.
- ◆ To unconfigure the deduplication storage server, enter the following:

```
Dedupe> unconfig destroy_fs=yes/no config_name
```

The *config_name* is the name of the deduplication server which is being unconfigured. It must be specified if more than one deduplication server is configured on the cluster.

The *destroy_fs* is an optional parameter with default value set to *no*.

If *destroy_fs=no*, the file systems that are used in the deduplication configuration (both data and catalog file systems) are not destroyed. You can later reconfigure deduplication using the same file systems.

If *destroy_fs=yes*:

- If WORM was configured, a check is made to see if there are any backup images present in the system. If backup images are present, you cannot destroy the file system and the `dedupe unconfig` command fails. Before you unconfigure Veritas Data Deduplication and destroy the file systems, ensure that you have removed the storage unit, disk pool and storage server from the NetBackup primary server.
- If WORM is not configured, the file systems are destroyed after unconfiguring deduplication.

To add an IP to the deduplication cluster

- ◆ To add an IP to the deduplication cluster, enter the following:

```
dedupe> addip IP config_name
```

The *config_name* is the name of the deduplication server to which IP is being added. It must be specified if more than one deduplication server is configured on the cluster.

Note: IPs of the same subnet cannot be used as primary and secondary IPs. Ensure that the secondary IP which is being added to the deduplication configuration should not belong to the same subnet as the primary or any of the already existing secondary IPs.

To remove an IP from the deduplication cluster

- ◆ To remove an IP from the deduplication cluster, enter the following:

```
dedupe> removeip IP config_name
```

The *config_name* is the name of the deduplication server from which IP is being removed. It must be specified if more than one deduplication server is running on the cluster.

To change the primary IP of the deduplication server

You can change the primary IP of the deduplication server using the `dedupe changeip` command. The specified new IP should be an existing, free virtual IP in the cluster. You must map the new IP with the FQDN of the old IP in the `/etc/hosts` file before you run the `dedupe changeip` command. This mapping can be done using the `network host add <new_ip> <old_fqdn>` command.

- To change the primary IP of the deduplication server, enter the following:

```
dedupe> changeip IP config_name
```

The *config_name* is the name of the deduplication server to which IP is being added. It must be specified if more than one deduplication server is configured on the cluster.

After changing the IP from Access Appliance, you also have to update the IP from NetBackup. Run the following commands on NetBackup:

```
/usr/openv/netbackup/bin/admincmd/nbemmcmd -deletehost
-machinename <old_ip> -machinetype ndmp
/usr/openv/netbackup/bin/admincmd/nbemmcmd -renamehost -machinename
<old_ip> -machinetype ndmp -newmachinename <new_ip>
```

Note: You are required restart the NetBackup service after these commands get executed.

To add an FQDN to the deduplication server

The FQDNs of the NetBackup master and media servers should be added to the deduplication server. This configuration leads to faster backups on the storage server and should be done prior to adding the storage server in the NetBackup administration console or Web UI. Adding the FQDNs of the NetBackup servers allows for faster name resolution within the docker container.

Do not add the IP and FQDN of the deduplication service to the deduplication server as this can create issues when the MSDP services are started.

- To add an FQDN to the deduplication cluster, enter the following:

```
dedupe> fqdn add IP fqdn config_name
```

The *config_name* parameter must be specified if more than one deduplication server is running on the cluster.

To remove an FQDN from the deduplication cluster

- ◆ To remove an FQDN from the deduplication cluster, enter the following:

```
dedupe> fqdn remove IP config_name
```

The *config_name* parameter must be specified if more than one deduplication server is running on the cluster.

To set predictive cache size of the deduplication server

The *setpcachesize* command sets the predictive cache size of the deduplication server. The deduplication server should be already configured before you set its predictive cache size. During the execution of this command, the deduplication

server is restarted and this may affect any ongoing backups. This operation is supported only via CLISH.

- To set the predictive cache size of the deduplication server to a specific value, enter the following:

```
dedupe> setpcachesize value config_name
```

The *config_name* parameter must be specified if more than one deduplication server is running on the cluster.

To set sampling cache size of the deduplication server

The *setscachesize* command sets the sampling cache size of the deduplication server. The deduplication server should be already configured before you set its sampling cache size. During the execution of this command, the deduplication server is restarted and this may affect any ongoing backups. This operation is supported only via CLISH.

- To set the sampling cache size of the deduplication server to a specific value, enter the following:

```
dedupe> setscachesize value config_name
```

The *config_name* parameter must be specified if more than one deduplication server is running on the cluster.

To grow the catalog file system of the deduplication server

- To grow the catalog file system of the deduplication server to a specific size, enter the following:

```
dedupe> catgrow size config_name
```

The deduplication server should be online for this operation. This operation is not supported from the Access Appliance GUI.

The *config_name* parameter must be specified if more than one deduplication server is running on the cluster.

You can use the `dedupe catgrow <size>` command to grow the catalog file system to the specified size. The default `dedupe grow` command does not modify the size of the catalog file system. If there are multiple data file systems, specify the file system names in the same order in which you specified them in the `dedupe show` command previously.

To set the affinity of a deduplication server to a node

- ◆ To set the affinity of a deduplication server to a node, enter the following:

```
dedupe> affinity set node_name
        container_name
```

The container runs on the node on which affinity is set. This may take some time.

In case there is more than one container, it is mandatory to specify the *node_name*.

If there are multiple instances of Veritas Data Deduplication, note the following:

Use case 1: Two instances of Veritas Data Deduplication are created and the affinity is set to the same node in both instances

- The `MaxCache` size of the container is in proportion to the container's deduplication storage size (0.5GB/1TB of storage size)
- If the combined `MaxCache` size required by both the containers is within 60% of the system memory, then the maximum cache size of the deduplication instance is set according to the calculations (0.5GB/1TB).
- If the total combined `MaxCache` size of the containers exceed 60% of the system memory, then the `MaxCache` size is set in proportion to its storage size with the upper threshold set to 50% of the available memory. The container is also restarted if there is a change in `MaxCache` size to reflect the values inside the containers..

Use case 2: One instance of Veritas Data Deduplication fails over to the other node, either in the presence or absence of a Veritas Data Deduplication instance on the second node

- In case there is only one container in the cluster, then after failover, the migrating container is allocated a `MaxCache` size as was previously set, which is in proportion to its storage size (0.5GB/1TB).
- In case there is already a container running on the first node, then the container's `MaxCache` size is calculated and allocated if sufficient memory is available. Else, for the first container (which is running on the affinity node), the `MaxCache` size threshold is set to 50% of the available memory. For the second container, the `MaxCache` size is capped to 20% of system memory.

To upgrade the deduplication server to a higher version

- ◆ To upgrade the deduplication server to a higher version, enter the following:

```
dedupe> upgrade version config_name
```

The *config_name* parameter must be specified irrespective of the number of deduplication servers configured.

This command upgrades the Veritas Data Deduplication version to the specified version. As part of this command, the deduplication server is restarted and the backup services get disrupted.

For supported versions, See [“Supported configurations and versions for NetBackup with Veritas Data Deduplication”](#) on page 11.

See the `dedupe` man page for more information.

Managing Veritas Data Deduplication using GUI

This section describes the various operations that you can perform on the deduplication server using the GUI.

Viewing information about Veritas Data Deduplication

From the **Settings > Veritas Data Deduplication management** page, you can view the following:

- Configuration names
- Primary IP address for each instance
- Service status for each instance
- Deduplication version of each instance
- Node name for each instance
- Affinity for each instance

You can modify the service status, upgrade the deduplication version (NetBackup version), set affinity and unconfigure Veritas Data Deduplication for any of the instances after clicking on their respective kebab menu.

See [“About Veritas Data Deduplication”](#) on page 14.

See [“Starting or stopping the Veritas Data Deduplication service”](#) on page 49.

See [“Increasing storage for Veritas Data Deduplication”](#) on page 49.

See [“Unconfiguring Veritas Data Deduplication”](#) on page 49.

See [“Setting affinity node for Veritas Data Deduplication”](#) on page 54.

See [“Upgrade the deduplication version for Veritas Data Deduplication”](#) on page 53.

Starting or stopping the Veritas Data Deduplication service

You can start or stop the Veritas Data Deduplication service.

To start or stop the Veritas Data Deduplication service

- 1 Go to **Settings > Veritas Data Deduplication management**. Click on the kebab menu of the instance and select the **Enable/Disable** option to start or stop the service for that instance.
- 2 Confirm that you want to start or stop the Veritas Data Deduplication service. Click **Save**.
- 3 View the **Recent Activity** panel for the status of the task.

Increasing storage for Veritas Data Deduplication

You can increase storage for Veritas Data Deduplication.

To increase storage for Veritas Data Deduplication

- 1 Click on **Veritas Data Deduplication** page after expanding **Provisioned storage** option present on the left navigation bar.
- 2 Click on the kebab menu for the instance on which you want to perform the `grow` operation and click **Grow**.

The **Grow Veritas Data Deduplication Storage** dialog box is displayed.

- 3 Specify the storage options for Veritas Data Deduplication and confirm the settings.

When you grow a Veritas Data Deduplication storage pool, you need to use five disks or volumes. You add disks or volumes in multiples of five.

- 4 View the **Recent Activity** panel for the status of the task.

Note: Reducing the storage for Veritas Data Deduplication is not supported using GUI. You can reduce storage for Veritas Data Deduplication using CLISH.

See [“Managing Veritas Data Deduplication using the Access CLISH”](#) on page 39. for more details.

Unconfiguring Veritas Data Deduplication

You can unconfigure Veritas Data Deduplication.

Ensure that you make a note of the following configuration details before you unconfigure the Veritas Data Deduplication server with **Retain data** option. These details are required when you reconfigure the service.

- The virtual IP that you used to configure the Veritas Data Deduplication server.
- Credentials to access the Veritas Data Deduplication server.
- File system names (including data and catalog file systems).

Note: You can reconfigure Veritas Data Deduplication from the Access Appliance CLISH only.

To unconfigure Veritas Data Deduplication

- 1 Go to **Settings > Veritas Data Deduplication management**. Click on the kebab menu of the instance.
- 2 Click **Unconfigure**.
- 3 The **Unconfigure Veritas Data Deduplication** dialog box is displayed.
Select one of the options - **Retain data** (if you want to keep the Veritas Data Deduplication utilized file systems) or **Delete data** (if you want to delete the file systems) and click **Save**.
- 4 View the **Recent Activity** panel for the status of the task.


Adding secondary IP for Veritas Data Deduplication

You can add a secondary IP for any instance of Veritas Data Deduplication.

To add secondary IP for any instance

- 1 Go to **Settings > Veritas Data Deduplication management**. Click on the configuration name for which you want to add a secondary IP.
- 2 Click **Add IP**.

- 3 The **Add Deduplication Secondary IP** dialog box is displayed. Enter the IP. Click **Next**.



Note: The restrictions which are applicable to the virtual IP field during configuring of Veritas Data Deduplication are also applicable here.


- 4 Click **Finish**.
- 5 View the **Recent Activity** panel for the status of the task.

Deleting secondary IP for Veritas Data Deduplication

You can delete the secondary IP for any instance of Veritas Data Deduplication.

To remove secondary IP for any instance

- 1 Go to **Settings > Veritas Data Deduplication management**. Click on the configuration name for which you want to delete the secondary IP.
- 2 Click **Remove** for the IP that you want to remove. The **Remove Deduplication Secondary IP** dialog box is displayed. Click **Next**.



Remove Deduplication Secondary IP

Remove Secondary IP Address

Cancel Next

- 3 Click **Finish**.
- 4 View the **Recent Activity** panel for the status of the task.

Changing the primary IP for Veritas Data Deduplication

You can change the deduplication primary IP for any instance of Veritas Data Deduplication.

To change the deduplication primary IP for any instance

- 1 Go to **Settings > Veritas Data Deduplication**.
- 2 Click on the instance name for which you want to change the primary IP.
- 3 From the primary IP section, click on the **Edit** icon (pencil symbol).

- 4 A dialog box is displayed. From the drop-down list, select the virtual IP that you want to use as the new primary IP for that instance.

Click **Save**.



- 5 You can click **View all activities** to see the progress of the task.

Note: You must map the new IP with the FQDN of the old IP in the `/etc/hosts` file. You can do this using the `network host add <new_ip> <old_ip_fqdn>` command from CLISH.

Upgrade the deduplication version for Veritas Data Deduplication


You can upgrade the version of deduplication (NetBackup version) for any instance of Veritas Data Deduplication

To upgrade the deduplication version for any instance

- 1 Go to **Settings > Veritas Data Deduplication**.
- 2 Click on kebab menu of the instance for which you want to upgrade the deduplication version (NetBackup version). Select **Upgrade**.

- 3 A dialog box is displayed. From the drop-down list, select the required version of deduplication.

Click **Save**.



- 4 You can click **View all activities** to see the progress of the task.


Setting affinity node for Veritas Data Deduplication

You can set affinity node for any instance of Veritas Data Deduplication.

To set affinity node for any instance

- 1 Go to **Settings > Veritas Data Deduplication**.
- 2 Click on kebab menu of the instance for which you want to set affinity. Select **Set affinity**.

- 3 A dialog box is displayed. From the drop-down list, select the required node.
 Click **Save**.



- 4 You can click **View all activities** to see the progress of the task.

Accessing Access Appliance storage shell for management tasks

To perform some management tasks on any storage server instance, you must open an SSH session to the instance and log in to the storage server shell. Any user with appliance administrator role can login and perform management tasks.

From the storage shell, you can run commands to manage the instance and your storage. See [“About the Access Appliance storage server shell”](#) on page 55.

About the Access Appliance storage server shell

Use the Access Appliance storage server shell to configure the immutable and indelible data from the appliance. The interface provides tab-completed command options.

These are main categories of commands:

- `dedupe`
 This command lets you manage the deduplication service.
 See [“About the dedupe command”](#) on page 56.
- `retention`
 This command lets you manage image retention.
 See [“About the retention command”](#) on page 58.

- **setting**
 This command lets you manage the deduplication and system configuration settings.
 See [“About the setting command”](#) on page 58.
- **support**
 This command lets you access and upload the relevant logs and configuration files for troubleshooting.
 See [“About the support command”](#) on page 62.

About the dedupe command

The `dedupe` commands lets you manage the deduplication services.

The following table describes the options and arguments for the `dedupe` command.

Table 3-2 The options and arguments for the `dedupe` command.

Option and its description	Argument	Description
CRC Check and manage the deduplication CRC service.	<code>state</code>	Display the state of the CRC checking services.
	<code>enable</code>	Enable the CRC checking services.
	<code>enable-fixmode</code>	Enable the CRC checking services with a fix mode.
	<code>disable-fixmode</code>	Disable the fix mode of the CRC checking services.
	<code>disable</code>	Disable the CRC checking services.
	<code>fast</code>	Restart the CRC checking services in a mode.
	<code>fixmode-state</code>	Display the state of the fix mode of the CRC checking services.
CRQP Manage the transaction log (Tlog) process service.	<code>status</code>	Display the status of the CRQP services.
	<code>start</code>	Start the CRQP services.
	<code>info</code>	Show the CRQP services information.

Table 3-2 The options and arguments for the `dedupe` command. (*continued*)

Option and its description	Argument	Description
<code>online-check</code> Manage online checking services.	<code>enable</code>	Enable the online checking services.
	<code>status</code>	Display the status of the online checking services.
	<code>disable</code>	Disable the online checking services.
<code>compaction</code> Manage the MSDP compaction service.	<code>start</code>	Start the deduplication compaction services.
	<code>state</code>	Display the status of the deduplication compaction services.
	<code>disable</code>	Disable the deduplication compaction services.
	<code>enable</code>	Enable the deduplication compaction services.
<code>MSDP</code> Manage the MSDP services.	<code>start</code>	Start the deduplication (MSDP) services.
	<code>status</code>	Display the status of the deduplication (MSDP) services.
	<code>stop</code>	Stop the deduplication (MSDP) services.

Stopping and starting the MSDP services

To stop the MSDP services:

1. Check the health monitor status using the following command:

```
setting health status
```
2. If the health monitor is enabled, stop the monitor using the following command:

```
setting health disable
```
3. After disabling the health monitor, use the following command to stop the MSDP services:

```
dedupe MSDP stop
```

To start the MSDP services:

1. Start the MSDP services using the following command:

```
dedupe MSDP start
```

2. Start the health monitor using the following command:

```
setting health enable
```

About the retention command

The `retention` command lets you manage image retention.

The following table describes the options and arguments for the `retention` command.

Table 3-3 The options and arguments for the `retention` command.

Option and its description	Argument	Description
<code>policy</code> The backup policy to use for the image retention with a retention lock.	<code>list</code>	List the images that have a retention lock.
	<code>audit</code>	Display the audit information for the images with a retention lock.
	<code>disable</code>	Disable the image retention using the backup ID of the image. For example: <pre>retention policy disable backup_ID=<test_123> copynumber=<number></pre> You can find the backup ID and the copy number in the output of the <code>retention policy list</code> command.

About the setting command

The `setting` command lets you manage the deduplication and system configuration settings.

The following table describes the options and arguments for the `setting` command.

Table 3-4 The options and arguments for the `setting` command.

Option and its description	Argument	Description
<code>FIPS</code> Manage FIPS settings.	<code>status</code>	Display the status of FIPS in MSDP.

Table 3-4 The options and arguments for the `setting` command.
(continued)

Option and its description	Argument	Description
network Configure the network settings.	ping	Ping the network for a connection status. Use the <code>ip</code> parameter to enter the IP address of the network to check the connection status.
	ifconfig	Display the IP address of the network.
	route	Display the network route information.
WORM Configure the settings for immutable and indelible storage.	status	Display the status of WORM storage.
	set-max	Specify the maximum duration to keep the storage immutable and indelible. Use the <code>worm_max</code> parameter to specify the maximum duration in seconds to keep the storage immutable and indelible (WORM).
	set-min	Specify the minimum duration to keep the storage immutable and indelible. Use the <code>worm_min</code> parameter to specify the minimum duration in seconds to keep the storage immutable and indelible (WORM).
	show-mode	Display the WORM mode.

Table 3-4 The options and arguments for the `setting` command.
(continued)

Option and its description	Argument	Description
<code>set-MSDP-param</code> Set the parameters in the MSDP configuration files.	<code>write-thread-num</code>	Get and set the <code><WriteThreadNum></code> parameter. Use the <code>value</code> parameter to enter the value that is used by the MSDP parameters.
	<code>spoold-logging</code>	Get and set the <code><SpoolLogging></code> parameter.
	<code>allocation-unit-size</code>	Get and set the <code><AllocationUnitSize></code> parameter.
	<code>data-check-days</code>	Get and set the <code><DataCheckDays></code> parameter.
	<code>max-fp-cache-size</code>	Get and set the <code><MaxFPCacheSize></code> parameter.
	<code>max-retry-count</code>	Get and set the <code><MaxRetryCount></code> parameter.
	<code>spad-logging</code>	Get and set the <code><SpadLogging></code> parameter.
	<code>log-retention</code>	Get and set the <code><LogRetention></code> parameter.

Table 3-4 The options and arguments for the `setting` command.
(continued)

Option and its description	Argument	Description
<code>get-MSDP-param</code> Get the parameters from the MSDP configuration files.	<code>write-thread-num</code>	Get and set the <code><WriteThreadNum></code> parameter.
	<code>spoold-logging</code>	Get and set the <code><SpoolLogging></code> parameter.
	<code>allocation-unit-size</code>	Get and set the <code><AllocationUnitSize></code> parameter.
	<code>data-check-days</code>	Get and set the <code><DataCheckDays></code> parameter.
	<code>max-fp-cache-size</code>	Get and set the <code><MaxFPCacheSize></code> parameter.
	<code>max-retry-count</code>	Get and set the <code><MaxRetryCount></code> parameter.
	<code>spad-logging</code>	Get and set the <code><SpadLogging></code> parameter.
	<code>log-retention</code>	Get and set the <code><LogRetention></code> parameter.
<code>MSDP-user</code> Configure the settings for the MSDP user.	<code>random-password</code>	Generates a random password.
	<code>list</code>	List all the MSDP users.
	<code>verify-user</code>	Verify user password and user role.
	<code>reset-password</code>	Reset the password for the MSDP user.
<code>encryption</code> Configure the settings for MSDP encryption.	<code>status</code>	Display the current encryption status.
	<code>enable</code>	Enable encryption for MSDP.
<code>kernel</code> Get the information about the kernel.	<code>search-param</code>	Search for a keyword in the kernel parameters. Use the <code>keyword</code> to specify a keyword to search in the keyword parameters.

About the support command

The `support` command lets you access and upload the relevant logs and configuration files for troubleshooting.

The following table describes the options and arguments for the `support` command.

Table 3-5 The options and arguments for the `support` command.

Option and its description	Argument	Description
MSDP-history Access the MSDP history files.	tail	Append the last 10 lines of each file to standard output. This command requires the following parameters: ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail</code>: <code>-f</code>, <code>-n</code>.
	collect	Collect files for transferring to the target host. This command requires the following parameters: ■ <code>mmin</code> Find files that are modified <code>n</code> minutes ago. Same as <code>mmin</code> option in <code>find</code> command. For example, 15, +15, -15. ■ <code>pattern</code> Specify the naming pattern to find and select matching files and folders. For example, <code>spoold</code>, or <code>spad*</code>. ■ <code>mtime</code> Finds files that are modified <code>n*24</code> hours ago. Same as <code>mtime</code> option in <code>find</code> command. For example, 2, +2, -2.
	ls	List information about the FILES. Use the <code>dir</code> parameter to enter the directory name.
	scp	

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
		Securely transfer selected files to the target host. This command requires the following parameters: ■ <code>scp_port</code> Specifies the port to connect to on the remote host. ■ <code>scp_target</code> Specifies the target host and path for transferring files. Use the <code>user@host:/directory path/</code> format. This is a required parameter.
	<code>cat</code>	Concatenate files and print on the standard output. Use the <code>file</code> parameter to enter the file name.
	<code>grep</code>	Print lines that match patterns. This command requires the following parameters: ■ <code>pattern</code> Specify the naming pattern to find and select matching files and folders. For example, <code>spoold</code>, or <code>spad*</code>. This is a required parameter. ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail: -f</code>, <code>-n</code>.

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
<code>process</code> Displays information about the MSDP processes.	<code>MSDP-process</code>	Display the MSDP processes.
	<code>htop</code>	Display the CPU and memory information.
	<code>atop</code>	Display information about the operating system.
	<code>pidstat</code>	Display the PID statistics.
	<code>memory-usage</code>	Displays the free and used memory.
<code>software</code> Displays information about the software.	<code>show-MSDP-version</code>	Display the MSDP version.
	<code>show-OS-version</code>	Display the operating system information.

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
MSDP-log Access the MSDP log files.	tail	Append the last 10 lines of each file to standard output. This command requires the following parameters: ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail</code>: <code>-f</code>, <code>-n</code>.
	collect	Collect files for transferring to the target host. This command requires the following parameters: ■ <code>mmin</code> Find files that are modified <code>n</code> minutes ago. Same as <code>mmin</code> option in <code>find</code> command. For example, <code>15</code>, <code>+15</code>, <code>-15</code>. ■ <code>pattern</code> Specify the naming pattern to find and select matching files and folders. For example, <code>spoold</code>, or <code>spad*</code>. ■ <code>mtime</code> Finds files that are modified <code>n*24</code> hours ago. Same as <code>mtime</code> option in <code>find</code> command. For example, <code>2</code>, <code>+2</code>, <code>-2</code>.
	ls	List information about the files. Use the <code>dir</code> parameter to enter the directory name.
	scp	

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
		Securely transfer selected files to the target host. This command requires the following parameters: ■ <code>scp_port</code> Specifies the port to connect to on the remote host. ■ <code>scp_target</code> Specifies the target host and path for transferring files. Use the <code>user@host:/directory path/</code> format. This is a required parameter.
	<code>cat</code>	Concatenate files and print on the standard output. Use the <code>file</code> parameter to enter the file name.
	<code>grep</code>	Print lines that match patterns. This command requires the following parameters: ■ <code>pattern</code> Specify the naming pattern to find and select matching files and folders. For example, <code>spoold</code>, or <code>spad*</code>. This is a required parameter. ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail: -f</code>, <code>-n</code>.
<code>hardware</code> Displays information about the hardware.	<code>cpumem</code>	Display the CPU and memory information.

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
MSDP-config Access the MSDP configuration files.	tail	Append the last 10 lines of each file to standard output. This command requires the following parameters: ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail</code>: <code>-f</code>, <code>-n</code>.
	collect	Collect files for transferring to the target host. This command requires the following parameters: ■ <code>mmin</code> Find files that are modified <code>n</code> minutes ago. Same as <code>mmin</code> option in <code>find</code> command. For example, <code>15</code>, <code>+15</code>, <code>-15</code>. ■ <code>pattern</code> Specify the naming pattern to find and select matching files and folders. For example, <code>spoold</code>, or <code>spad*</code>. ■ <code>mtime</code> Finds files that are modified <code>n*24</code> hours ago. Same as <code>mtime</code> option in <code>find</code> command. For example, <code>2</code>, <code>+2</code>, <code>-2</code>.
	ls	List information about the FILES. Use the <code>dir</code> parameter to enter the directory name.
	scp	

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
		Securely transfer selected files to the target host. This command requires the following parameters: ■ <code>scp_port</code> Specifies the port to connect to on the remote host. ■ <code>scp_target</code> Specifies the target host and path for transferring files. Use the <code>user@host:/directory path/</code> format. This is a required parameter.
	<code>cat</code>	Concatenate files and print on the standard output. Use the <code>file</code> parameter to enter the file name.
	<code>grep</code>	Print lines that match patterns. This command requires the following parameters: ■ <code>pattern</code> Specify the naming pattern to find and select matching files and folders. For example, <code>spoold</code>, or <code>spad*</code>. This is a required parameter. ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail: -f</code>, <code>-n</code>.

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
diskio Displays information about the disk I/O.	iostat	Displays the information about the disk I/O.
	vmstat	Displays the information on the wait on the disk I/O.
	nmon	Display the information about the monitor system.
	disk-volume	Display information about the disk volume.

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
syslogs Access the system logs.	tail	Append the last 10 lines of each file to standard output. This command requires the following parameters: ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail</code>: <code>-f</code>, <code>-n</code>.
	collect	Collect files for transferring to the target host. This command requires the following parameters: ■ <code>mmin</code> Find files that are modified <code>n</code> minutes ago. Same as <code>mmin</code> option in <code>find</code> command. For example, 15, +15, -15. ■ <code>pattern</code> Specify the naming pattern to find and select matching files and folders. For example, <code>spoold</code>, or <code>spad*</code>. ■ <code>mtime</code> Finds files that are modified <code>n*24</code> hours ago. Same as <code>mtime</code> option in <code>find</code> command. For example, 2, +2, -2.
	ls	List information about the FILES. Use the <code>dir</code> parameter to enter the directory name.
	scp	

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
		Securely transfer selected files to the target host. This command requires the following parameters: ■ <code>scp_port</code> Specifies the port to connect to on the remote host. ■ <code>scp_target</code> Specifies the target host and path for transferring files. Use the <code>user@host:/directory path/</code> format. This is a required parameter.
	<code>cat</code>	Concatenate files and print on the standard output. Use the <code>file</code> parameter to enter the file name.
	<code>grep</code>	Print lines that match patterns. This command requires the following parameters: ■ <code>pattern</code> Specify the naming pattern to find and select matching files and folders. For example, <code>spoold</code>, or <code>spad*</code>. This is a required parameter. ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail: -f</code>, <code>-n</code>.

Table 3-5 The options and arguments for the `support` command.
(continued)

Option and its description	Argument	Description
<code>proc</code> Access process information about the pseudo-filesystem.	<code>tail</code>	Append the last 10 lines of each file to standard output. This command requires the following parameters: ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail</code>: <code>-f</code>, <code>-n</code>.
	<code>ls</code>	List information about the FILES. Use the <code>dir</code> parameter to enter the directory name.
	<code>cat</code>	Concatenate files and print on the standard output. Use the <code>file</code> parameter to enter the file name.
	<code>grep</code>	Print lines that match patterns. This command requires the following parameters: ■ <code>pattern</code> Specify the naming pattern to find and select matching files and folders. For example, <code>spoold</code>, or <code>spad*</code>. This is a required parameter. ■ <code>file</code> Enter the filename. This is a required parameter. ■ <code>options</code> Enter the supported options for <code>grep</code>: <code>-i</code>, <code>-w</code>, <code>-E</code>, or <code>tail</code>: <code>-f</code>, <code>-n</code>.

Here are some examples of the `support` command usage:

```
support MSDP-log ls dir=spoolld
support MSDP-log ls dir=spad/soflvm08.tec.com/spoolld/spad/072320.log
support MSDP-log cat file=pdde-config.log
support MSDP-log cat file=spad/soflvm08.tec.com/spoolld/spad/072320.log
support MSDP-log tail file=pdde-config.log
support MSDP-log tail file=pdde-config.log options="-f"
support MSDP-log tail file=pdde-config.log options="-n 5"
support MSDP-log tail file=spad/sofia11vm08/spoolld/spad/072320.log
options="-f -n 5"
support MSDP-log grep file=spad* pattern=sessionStartAgent
options="-r"
support MSDP-log grep file=pdde-setup.log pattern="pid 455"
support MSDP-log grep file=pdde-setup.log pattern="Pid" options="-w
-i"
support MSDP-log grep file=spoolld* pattern="Pid" options="-w -i -r"
support MSDP-log grep file=pdde-setup.log pattern="MSDP-X|PureDisk"
options="-w -i -r -E"
```

Examples to uploading logs, history, or config files:

Uploading logs is a two-step process:

- Run the `collect` command for the desired category to collect files of interest. For example:

```
support MSDP-log collect pattern=spoolld* mmin="+2"
support MSDP-history collect
support MSDP-config collect
support syslogs collect pattern=crash mmin="-2"
```

- Transfer files using `scp` command

Run the `scp` command from any category to create a tarball of all previously collected files (from all categories) and transfer the tarball to the target host using the `scp` protocol.

```
support MSDP-config scp scp_target=user@example.com:/tmp
```

Support for multiple domains across networks for Veritas Data Deduplication

You can use a single deduplication pool to backup data from multiple NetBackup domains. The Veritas Data Deduplication server deduplicates backups originating from different subnets. Use the `network ip addr show` command to list the multiple subnets that have been created.

```
Network> ip addr show
```

IP	Netmask/Prefix	Device	Node	Type	Status
192.168.10.5	255.255.255.0	eth5	vaupg_01	Physical	
192.168.10.13	255.255.255.0	eth5	vaupg_02	Physical	
192.168.30.2	255.255.255.0	eth4	vaupg_01	Physical	
192.168.30.10	255.255.255.0	eth4	vaupg_02	Physical	
10.209.192.206	255.255.252.0	eth1	vaupg_01	Physical	
10.209.192.208	255.255.252.0	eth1	vaupg_02	Physical	
10.209.194.38	255.255.252.0	eth1	vaupg_01	Virtual	ONLINE (Con IP)
192.168.10.22	255.255.255.0	eth5	vaupg_01	Virtual	ONLINE
192.168.10.25	255.255.255.0	eth5	vaupg_02	Virtual	ONLINE
192.168.30.15	255.255.255.0	eth4	vaupg_02	Virtual	ONLINE
192.168.30.20	255.255.255.0	eth4	vaupg_01	Virtual	ONLINE

You can use the following command to configure the deduplication server with a virtual IP.

```
Dedupe> config filesystem1[,filesystem2,...] IP username
[catfs_name] vdd_version [encryption] [worm=yes] [minret] [maxret]
[node_name] [config_name]
```

For example:

```
Dedupe> config D3_3533793 192.168.10.22 dedupeuser CAT_3088152 19.0.1
enable_encryption=no enable_worm=yes 3600s 1892160000s cluster1_01 msdp-worm
Access Appliance dedupe INFO V-493-10-3732 Deduplication server configured su
```

You can verify that the IP has been added successfully using the `dedupe show` command.

```
Dedupe> show
```

Parameter	msdp-worm
Catalog_FileSystem	CAT_3088152
Virtual_IP	192.168.10.22
Secondary_IP	-
Filesystem	D3_3533793
Deduplication_Storage_Size	500.0TB
WORM_enabled	True
Hostname	cluster1_01
Server_Status	ONLINE
Minimum_Retention_Value	3600
Maximum_Retention_Value	1892160000
Cache_Type	FP Cache

```
Max_Cache_Size          250GiB
Predictive_Cache_Size   -
Sampling_Cache_Size     -
Percentage_Memory_Usage 5.59%
Affinity                 cluster1_01
Encryption_enabled      False
KMS_enabled             False
KMS_server              -
KMS_keygroup            -
Shrink_status           -
Shrink_errors           -
Dedupe_FQDN             accessvip01.vxindia.veritas.com
Dedupe_Version          19.0.1
```

You can add another IP from a different subnet.

```
Dedupe> addip <IPaddress>
```

For example:

```
Dedupe> addip 192.168.10.25
ACCESS dedupe INFO V-493-10-0 IP has been successfully added to the
deduplication server.
Dedupe> addip 192.168.30.15
ACCESS dedupe INFO V-493-10-0 IP has been successfully added to the
deduplication server.
Dedupe> addip 192.168.30.20
ACCESS dedupe INFO V-493-10-0 IP has been successfully added to the
deduplication server.
```

You can verify that the virtual IP is added successfully. The `dedupe show` command is used to display information about the deduplication server that has been configured. The output of this command includes the file system(s) being used, deduplication server status, storage capacity, secondary IPs used, percentage of memory usage, maximum cache size, and the IP and the cluster node on which the server is running.

```
Dedupe> show
Parameter                      msdp-worm
=====
Catalog_FileSystem             CAT_3088152
Virtual_IP                     192.168.10.22
Secondary_IP                   192.168.10.25,192.168.30.15,192.168.30.20
Filesystem                     D3_3533793
```

```

Deduplication_Storage_Size  500.0TB
WORM_enabled                True
Hostname                    cluster1_01
Server_Status               ONLINE
Minimum_Retention_Value     3600
Maximum_Retention_Value     1892160000
Cache_Type                  FP Cache
Max_Cache_Size              250GiB
Predictive_Cache_Size       -
Sampling_Cache_Size         -
Percentage_Memory_Usage     5.59%
Affinity                    cluster1_01
Encryption_enabled          False
KMS_enabled                 False
KMS_server                  -
KMS_keygroup                -
Shrink_status               -
Shrink_errors               -
Dedupe_FQDN                 accessvip01.vxindia.veritas.com
Dedupe_Version              19.0.1

```

Use the `network ip addr show` command to list the multiple subnets that have been created.

```

Network> ip addr show

```

IP	Netmask/Prefix	Device	Node	Type	Status
192.168.10.5	255.255.255.0	eth5	vaupg_01	Physical	
192.168.10.13	255.255.255.0	eth5	vaupg_02	Physical	
192.168.30.2	255.255.255.0	eth4	vaupg_01	Physical	
192.168.30.10	255.255.255.0	eth4	vaupg_02	Physical	
10.209.192.206	255.255.252.0	eth1	vaupg_01	Physical	
10.209.192.208	255.255.252.0	eth1	vaupg_02	Physical	
10.209.194.38	255.255.252.0	eth1	vaupg_01	Virtual	ONLINE (Con IP)
192.168.10.22	255.255.255.0	eth5	vaupg_01	Virtual	ONLINE (Dedupe IP)
192.168.10.25	255.255.255.0	eth5	vaupg_01	Virtual	ONLINE (Dedupe IP)
192.168.30.15	255.255.255.0	eth4	vaupg_01	Virtual	ONLINE (Dedupe IP)
192.168.30.20	255.255.255.0	eth4	vaupg_01	Virtual	ONLINE (Dedupe IP)

You can also add a second user. The second user is used to add the second NetBackup domain, if required. You should use a different user for every new NetBackup domain.

```
Dedupe> adduser user1
Enter Password:
Confirm Password:
ACCESS dedupe INFO V-493-10-0 Adding new user user2 to deduplication server.
ACCESS dedupe INFO V-493-10-0 Added new user user2 to deduplication server
successfully.
```

You can perform backup and restore operations from the NetBackup domains from the two subnets.

Note: You can remove the virtual IP that you added later after the initial configuration. But you cannot remove the virtual IP specified during the initial configuration.

```
Dedupe> removeip IPaddress
```

For example:

```
Dedupe> removeip 192.168.30.15
ACCESS dedupe INFO V-493-10-0 IP has been successfully removed from the
deduplication server.
```

You can verify that the second virtual IP has been removed from the configuration using the `dedupe show` command.

```
Dedupe> show
Parameter                                msdp-worm
=====
Catalog_FileSystem                       CAT_3088152
Virtual_IP                               192.168.10.22
Secondary_IP                             192.168.10.25,192.168.30.20
Filesystem                               D3_3533793
Deduplication_Storage_Size                500.0TB
WORM_enabled                             True
Hostname                                 cluster1_01
Server_Status                             ONLINE
Minimum_Retention_Value                   3600
Maximum_Retention_Value                   1892160000
Cache_Type                               FP Cache
Max_Cache_Size                           250GiB
Predictive_Cache_Size                     -
Sampling_Cache_Size                       -
Percentage_Memory_Usage                   5.59%
Affinity                                  cluster1_01
```

Encryption_enabled	False
KMS_enabled	False
KMS_server	-
KMS_keygroup	-
Shrink_status	-
Shrink_errors	-
Dedupe_FQDN	accessvip01.vxindia.veritas.com
Dedupe_Version	19.0.1

You can use the `dedupe unconfig` command to free up all the virtual IPs added to the deduplication server. This command also unconfigures Veritas Data Deduplication from the cluster.

```
Dedupe> unconfig destroy_fs=yes config_name
ACCESS dedupe INFO V-493-10-0 Removing deduplication server configuration...
ACCESS dedupe INFO V-493-10-0 Deduplication server unconfigured successfully.
```

Veritas Data Deduplication storage layout

The layout of the Veritas Data Deduplication storage is such that when you configure Veritas Data Deduplication, multiple file systems of 100TiB and a separate catalog file system of size 5TiB are created.

The `dedupe grow` command creates a new file system for every 100TiB. If you perform upgrade and you had previously configured Veritas Data Deduplication, then the old layout of Veritas Data Deduplication is preserved. If a separate catalog file system existed before upgrade, the same catalog file system is present after the upgrade also. If the size of data filesystems is more than 100 TiB, then those filesystems remain as it is but new data filesystems are created for every 100 TiB that is grown using the `dedupe grow` command.

Note: Veritas Data Deduplication configuration requires 5 TiB for creating the catalog file system. This is over and above the size specified for the deduplication pool in the GUI.

Configuring a Veritas Data Deduplication storage unit on NetBackup


Veritas recommends limiting the total number of I/O streams to 80 per storage server. This limit on I/O streams is applicable to the combined total of duplication and replication jobs. You can configure Access Appliance such that multiple

NetBackup domains can write concurrently to the same disk pool. Veritas recommends that you adjust the I/O streams for each NetBackup primary's storage server, and/or set the appropriate SLP windows so that the total I/O streams from all sources do not exceed 80 concurrent streams to and from a single Access storage server.

Note: Since Access can be configured as both a duplication target and a replication target, allowing up to 80 I/O streams can be excessive when many concurrent jobs write to the storage server. If the Access disk pool/Veritas Data Deduplication instance becomes overloaded, the disk pool may become unstable. Veritas recommends that you reduce this limit to a lower value until the disk pool operates reliably.

To create a storage server on NetBackup

- 1 Log on to the NetBackup primary server from the NetBackup web UI.
- 2 Select **Disk storage**.
- 3 The **Add Storage** pop-up window appears. Select **OpenStorage**. Click **Start**.



- 4 The **Add OpenStorage storage server** form appears. Enter the storage server details and your credentials. Specify the **Storage server type** as **PureDisk**. Click **Next**.

Veritas NetBackup™

Add OpenStorage storage server

1 Basic properties 2 Media server 3 Review

pub25-q45.amazonaws.com

Storage server details

Storage server type *
PureDisk

Storage server name *
pub25-q45.amazonaws.com

Username *
Enter user name

Password *
Enter password

Cancel Previous Next

- 5 In the **Select media server** form, select the media server from the list of media servers that is available. Click **Select**.

Veritas NetBackup™

Select media server

Search...

Name
(Primary server)

Showing 1-1 of 1 (1 selected) Rows per page: 100

Cancel Select


- 6 You have the option to add additional media servers. If you want to add additional media servers, click **Add**.
- Click **Next** to continue.

The screenshot shows the 'Add OpenStorage storage server' form in the Veritas NetBackup interface. The left sidebar contains a navigation menu with options: Workloads, Storage, Disk storage, Media servers, Snapshot Manager, Storage lifecycle policies, Storage units, Tape storage, Catalog, and Detection and reporting. The main content area has a progress bar at the top with three steps: 'Basic properties' (completed with a green checkmark), 'Media server' (current step, indicated by a blue circle with the number 2), and 'Review' (indicated by a blue circle with the number 3). Below the progress bar, the section is titled 'Select additional media servers'. It contains a large empty box with a folder icon and the text 'You have not added any additional media servers.' Below this box is a blue 'Add' button. At the bottom right of the form are three buttons: 'Cancel', 'Previous', and 'Next'.


- 7 In the **Add OpenStorage storage server** form, click **Save**.

The screenshot shows the 'Add OpenStorage storage server' form in the Veritas NetBackup interface, now at the 'Review' step. The left sidebar is the same as in the previous screenshot. The progress bar at the top now shows 'Basic properties' (completed with a green checkmark), 'Media server' (completed with a green checkmark), and 'Review' (current step, indicated by a blue circle with the number 3). The main content area shows the 'Storage server type: OpenStorage Technology (OST) (PureDisk)' and a 'Media server' section with the text 'You have not added any media servers.' Below this section are three buttons: 'Cancel', 'Previous', and 'Save'.

- 8 Verify that the storage server is created successfully. Click **Add disk pool**.



- 9 In the **Select a storage server** form, select the storage servers. Click **Select**.



10 In the **Add MSDP disk pool** form, specify the disk pool name. Click **Next**.

The screenshot shows the 'Add MSDP disk pool' form in the Veritas NetBackup interface. The left sidebar contains navigation links: Dashboard, Activity monitor, Recovery, Protection, Workloads, Storage, Disk storage (selected), Media servers, Snapshot Manager, Storage lifecycle policies, Storage units, Tape storage, Catalog, Detection and reporting, Credential management, Hosts, and Bare Metal Restore. The main form area is titled 'Add MSDP disk pool' and has a progress bar with four steps: 1. Disk pool options (active), 2. Volumes, 3. Replication, and 4. Review. The form fields include: 'Storage server name' with a 'Change' button; 'Features' with the text 'Accelerator, A.I.R., Instant access, WORM capable'; 'Disk pool name' with the value 'lfr-pool' and a '0' icon; 'Description' with a placeholder 'Enter description'; a checkbox for 'Limit I/O streams' with the text 'Concurrent read and write jobs affect disk performance. Limit I/O streams to prevent disk overload.'; 'High water mark' with the value '98 %'; and 'Low water mark' with the value '80 %'. At the bottom right are 'Cancel', 'Previous', and 'Next' buttons.

11 In the **Add MSDP disk pool** form, click **Next**.

The screenshot shows the 'Add MSDP disk pool' form in the Veritas NetBackup interface, now at Step 2: Volumes. The progress bar shows 'Disk pool options' as completed (green checkmark) and 'Volumes' as the current step (blue circle). The main form area displays a search bar and a table of available volumes. The table has columns: Name, Available space, Total size, and Encryption. One volume, 'PureDiskVolume', is selected (radio button). Below the table, it says 'Showing 1-1 of 1 (1 selected)'. At the bottom right are 'Cancel', 'Previous', and 'Next' buttons.

Name	Available space	Total size	Encryption
PureDiskVolume	8.41 GB	8.75 GB	Yes

12 If you want to add any replication targets, click **Add**. Click **Next** to continue.

Veritas NetBackup™

Add MSDP disk pool

✓ Disk pool options ✓ Volumes 3 Replication 4 Review

You have not added any replication targets

Add

Cancel Previous Next

13 Verify all the details in the **Add MSDP disk pool** form, click **Finish**.

Veritas NetBackup™

Add MSDP disk pool

✓ Disk pool options ✓ Volumes ✓ Replication 4 Review

lfr-pool
Storage category: MSDP

Storage server

Storage server name: [redacted] Features: Accelerator, A.I.R., Instant access, WORM capable

Disk pool

Volumes


Name	Available space	Total size	Encryption	Replication	WORM capable	WORM mode
PureDiskVolume	8.49 GB	8.74 GB	No	None		

1 Records


Low water mark: 80% High water mark: 98%

Cancel Previous Finish

- 14** The disk pool gets created successfully. Select **Add storage unit** to add a storage unit.



- 15** In the **Add MSDP storage unit** form, enter the storage unit name and specify the maximum concurrent jobs that need to be run from the drop-down list. Click **Next**.



16 In the **Add MSDP storage unit** form, select a disk pool and click **Next**.

The screenshot shows the 'Add MSDP storage unit' configuration window in Veritas NetBackup. The left sidebar contains navigation links: Dashboard, Activity monitor, Recovery, Protection, Workloads, Storage, Disk storage, Media servers, Snapshot Manager, Storage lifecycle policies, Storage units (highlighted), Tape storage, Catalog, Detection and reporting, Credential management, Hosts, and Bare Metal Restore. The main area has a progress bar with four steps: Basic properties (checked), Disk pool (active), Media server, and Review. Below the progress bar, the 'Select a disk pool' section features a search bar and a table with the following data:

Name	Used space	Volumes	Storage type	Storage server	Replication	WORM capable
lt-pool	263.28 MB of 8.74 GB used	PureDiskVolume	PureDisk	10.221.40.128	None	

Below the table, there is a checkbox for 'On demand only' which is currently unchecked. At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Next' (which is highlighted in blue).

- 17** In the **Add MSDP storage unit** form, under **Select media server**, choose the **Manually select** option. Click **Next**.

The screenshot shows the 'Add MSDP storage unit' form in the Veritas NetBackup interface. The left sidebar contains navigation links: Dashboard, Activity monitor, Recovery, Protection, Workloads, Storage (selected), Disk storage, Media servers, Snapshot Manager, Storage lifecycle policies, Storage units (highlighted), Tape storage, Catalog, Detection and reporting, Credential management, Hosts, and Bare Metal Restore. The main panel is titled 'Add MSDP storage unit' and has three tabs: Basic properties (checked), Disk pool, and Media server (active). Under the 'Media server' tab, the 'Select media server' section has two radio buttons: 'Allow NetBackup to automatically select' (unchecked) and 'Manually select' (checked). Below this is a table with columns 'Name', 'NetBackup version', and 'OS platform'. One row is visible with a value of '10.3' under 'NetBackup version' and 'Linux' under 'OS platform'. At the bottom right are 'Cancel', 'Previous', and 'Next' buttons.

- 18** Verify that the storage unit is added successfully. Click **Save**.

The screenshot shows the 'Add MSDP storage unit' form in the Veritas NetBackup interface, now at the 'Review' step. The left sidebar is the same as in the previous screenshot. The main panel has three tabs: Basic properties (checked), Disk pool (checked), and Media server (checked). The 'Basic properties' tab is active, showing the 'lfr-pool-stu' storage unit. Below this, the 'Disk pool' section shows a table with columns 'Name', 'Available space', 'Raw size', 'Replication', and 'Volumes'. One row is visible with values: 'lfr-pool', '8.48 GB', '8.74 GB', 'None', and 'PureDiskVolume'. The 'Media server' section shows a table with columns 'Name', 'NetBackup version', and 'OS platform'. One row is visible with values: 'dfr-pool-stu-10.3', '10.3', and 'Linux'. At the bottom right are 'Cancel', 'Previous', and 'Save' buttons.

Configuring global deduplication using the Veritas Data Deduplication storage server across the domain

Access Appliance supports the configuration of multiple Logical Storage Units (LSU) with Veritas Data Deduplication. Using the global deduplication feature, multiple NetBackup domains can access the Veritas Data Deduplication storage server in Access Appliance. By default, only one NetBackup domain is preconfigured with the Veritas Data Deduplication storage server. To configure multiple domains, you need to create a new user for each domain. Different NetBackup domains require a different Veritas Data Deduplication user. A specific user can access data only from the domain to which the user has been added.

To add a new user to the deduplication storage server, enter the following:

```
dedupe> adduser username
```

Where *username* is the user name of the second domain user.

The command should be run on the Access Appliance cluster.

After a new user name is added to the storage server in Access Appliance, you can use the credentials of the new user to add a STU on NetBackup in the new domain. Any backups taken from the new domain are visible only to the new user. The NetBackup domain can access the catalog and data of the user. When you configure Auto Image Replication and add the storage target, use the Veritas Data Deduplication user name used by the target NetBackup domain.

To show the list of users of the deduplication storage server, enter the following:

```
dedupe> listuser
```

The command should be run on the Access Appliance cluster.

Note: You must use the deduplication IP address of Veritas Data Deduplication host that was previously configured with the first domain in the **Add Storage Server** form when you provide the details of the storage server.

Enabling Veritas Data Deduplication encryption

Access Appliance supports Veritas Data Deduplication encryption for all backup and deduplication jobs. This section describes how to enable encryption for the backup jobs.

You can do it in two different ways.

Method 1:

Note: In NetBackup long-term data retention (LTR) use case, the server side refers to Access Appliance and the client refers to the media server.

For details on the MSDP compression and encryption settings, see the *MSDP compression and encryption settings matrix* section in the *Configuring deduplication* chapter in the *Veritas NetBackup Deduplication Guide* on SORT.

For details on configuring encryption for MSDP backups, see the *Configuring encryption for MSDP backups* section in the *Configuring deduplication* chapter in the *Veritas NetBackup Deduplication Guide* on SORT.

Method 2:

You can use the Access Appliance command line to enable the encryption feature on Veritas Data Deduplication servers. If encryption is enabled in the Veritas Data Deduplication server, then any data coming from any media server is encrypted.

To enable encryption on Veritas Data Deduplication server

- ◆ To enable encryption on the deduplication storage server, enter the following:

```
dedupe> encryption enable config_name
```

This command should be run on the Veritas Data Deduplication server. To enable encryption, the Veritas Data Deduplication server should be in start state. Once encryption is enabled, it encrypts the forthcoming data segments of data containers.

The `config_name` parameter is an optional parameter for the `dedupe encryption enable` command when only one deduplication server is present in the cluster. It is mandatory if more than one deduplication server is present.

Enabling encryption crawler on Veritas Data Deduplication servers

The encryption crawler searches all MSDP pools to check for unencrypted data. It traverses all the existing data containers and if a data segment is found to be not encrypted, that segment is encrypted by the encryption crawler.

To enable the encryption crawler on Veritas Data Deduplication server

- ◆ To enable encryption crawler on the deduplication storage server, enter the following:

```
dedupe> encryption crawler enable config_name
```

This command should be run on the Veritas Data Deduplication server. To enable encryption crawler, encryption should be enabled on the Veritas Data Deduplication server.

The `config_name` parameter is an optional parameter for the `dedupe encryption crawler enable` command when only one deduplication server is present in the cluster. It is mandatory if more than one deduplication server is present.

To display the status of the encryption crawler of all deduplication servers

- ◆ To display the status of the encryption crawler of all deduplication servers, enter the following:

```
dedupe> encryption crawler status
```

To display the status of the encryption crawler of a specific deduplication server

- ◆ To display the status of the encryption crawler of a specific deduplication server, enter the following:

```
dedupe> encryption crawler status <dedup_server>
```

Support for bandwidth throttling during duplication

Veritas Data Deduplication supports bandwidth throttling during duplication.

For more details, see https://www.veritas.com/support/en_US/article.100037976

Setting up secure communication between Veritas Data Deduplication on Access Appliance and the NetBackup primary server

You can setup secure communication between Veritas Data Deduplication on Access Appliance and the NetBackup primary server.

Prerequisites:

- The `/usr/opensv/netbackup/bp.conf` file on the NetBackup primary server should have `MSDP_SERVER` entry for the MSDP server's FQDN.
- After adding the entry in the `bp.conf`, run the `/usr/opensv/netbackup/bin/admincmd/bprdreag -rereadconfig` command.

For clusters where DNS is not configured, the container's host name is created internally using the catalog file system name. For example, if the catalog file system name is `catfs1`, the container's host name is set to `catfs1.vdd.com`. For such setups, the `MSDP_SERVER` entry on the primary server's `bp.conf` file should point to the container's host name. To query the container's host name, use the `support software show-os-version` command in the MSDP restricted shell. Or if you have access to the host running the container, you can log on to the host and run the `docker exec <container_name> hostname` command.

You can enable secure communication between the deduplication server and NetBackup primary server using the `securecomm nbca enable` command. The `securecomm nbca enable` command enables Host ID certificate based secure communication between the deduplication server and the NetBackup primary, which is also the Certificate Authority. The authorization token should be generated on the NetBackup primary. A reissue token is required if the certificate is to be regenerated. This command only supports secure communication using a NetBackup Certificate Authority.

```
dedupe> securecomm nbca enable primary_server token config_name
```

primary_server FQDN of the NetBackup primary server which is also the KMS server.

token Token which is to be generated on the NetBackup primary.

config_name Name for the deduplication service.

Note: This is an optional parameter when only one deduplication server is present in the cluster. It is mandatory if more than one deduplication servers are present.

About MSDP encryption using NetBackup KMS service

Access Appliance incorporates Key Management Server (KMS) with Media Server Deduplication Pool. MSDP encryption carries out segment-level encryption and assigns a unique encryption key for every data segment. A customer key is retrieved

from NetBackup KMS to encrypt the segment key. Key creation and activation actions must be done manually (or using scripts) by the user.

The version of the NetBackup primary server which is to be used as the KMS server should be 10.1.1 for KMS encryption.

Pre-requisite:

Set up secure communication between Access Appliance and the primary server in its domain.

See [“Setting up secure communication between Veritas Data Deduplication on Access Appliance and the NetBackup primary server”](#) on page 91.

Managing MSDP KMS encryption using CLISH

You can use the `kms enable` command to enable encryption in the deduplication server using KMS. The key group specified should have an active key for encryption to work. Ensure that secure communication has been set up between the deduplication server and the NetBackup primary server before enabling KMS. The deduplication server is restarted as a part of the `kms enable` command.

```
dedupe> kms enable primary_server keygroup config_name
```

<i>primary_server</i>	FQDN of the NetBackup primary server which is also the KMS server.
<i>keygroup</i>	Name of the key group whose active key is used to encrypt data.
<i>config_name</i>	Name for the deduplication service in which KMS is to be enabled. Note: This is an optional parameter when only one deduplication server is present in the cluster. It is mandatory if more than one deduplication servers are present.

The `kms enable` command can be used to enable external KMS in Access Appliance MSDP. Ensure that the NetBackup primary server is already configured to work with the external KMS.

The encryption and KMS status of the deduplication servers can be verified with the `dedupe show` command. After you enable the KMS server, the existing unencrypted data can be encrypted with KMS encryption by running the encryption crawler.

Multi domain scenario

If the MSDP storage server is being used in multiple NetBackup domains, backup and restores from other domains fail if secure communication has been set up with the NetBackup primary server or the NetBackup KMS server. To avoid the backup

and restore failure from other domains, you must set up trust between the NetBackup primary server with which secure communication has been set up, and the rest of the NetBackup primary servers. You can set up trust using the same steps that you use to set up AIR between two appliances.

See [“Setting up Auto Image Replication \(AIR\) between two Access Appliances which have secure communication enabled”](#) on page 113.


For detailed steps, refer to the *NetBackup™ Deduplication Guide* on [SORT](#).

Support for NetBackup Auto Image Replication


You can use the NetBackup Auto Image Replication (AIR) feature to replicate from one media to another media with Veritas Data Deduplication storage: You have to create a trust relationship between two NetBackup primary/media servers to configure NetBackup AIR.

To set up trust between NetBackup primary servers

- 1 Log on to the NetBackup primary server (*NBU1*) from the web UI.
- 2 Navigate to **Hosts > Host properties**. Select the **Edit primary server** tab.



- 3 Click on **Servers > Trusted primary servers**. Click **Add**. Provide the target NetBackup primary server (*NBU2*) FQDN. Click on the **Validate Certificate Authority** option.



Add trusted primary server

When you establish trust between two primary servers, it allows the servers to securely communicate and perform replication operations.

Note: A trusted primary server is not required to be added for Isolated Recovery Environment.

Trusted primary server *

punsm6049p-12v01p2.vxindia.veritas.com

The primary server must be at NetBackup 8.1 or later.

Validate Certificate Authority

Cancel Create trust

- 4 After verifying the CA fingerprint, click **Next**.

- 5 Generate token on the target NetBackup primary server (*NBU2*) and use it for authentication. Or, you can specify the target NetBackup primary server (*NBU2*) credentials for authentication. Click on **Create trust**. The NetBackup primary server (*NBU2*) is listed in the **Trusted primary servers** tab.

Add trusted primary server
✕

punsm6049p-12v01p2.vxindia.veritas.com

The primary server must be at NetBackup 8.1 or later.

Validate Certificate Authority

☐ Specify authentication token of the trusted primary server

☒ Specify credentials of the trusted primary server

User name *

root

For a Windows host, specify the domain name along with the username. For example, domain name\username.

Password *

.....

Enter the password. If multi-factor authentication is configured for the user on the target host, append the one-time password.

Cancel

Create trust

- 6 Repeat the same steps on the target NetBackup primary server (*NBU2*) to set up trust with the NetBackup primary server (*NBU1*).

To configure storage servers for NetBackup Auto Image replication (AIR)

- 1 Configure the deduplication servers on the Access Appliances (both source and target) for AIR configuration.
- 2 In case of AIR with KMS Encryption support, See [“Setting up Auto Image Replication \(AIR\) between two Access Appliances which have secure communication enabled”](#) on page 113.

- 3** Add storage servers on the target NetBackup primary server (*NBU2*) with the target Access Appliance MSDP server's FQDN and credentials.
- 4** Log on to the target NetBackup primary server (*NBU2*) with the web UI. Navigate to **Storage > Storage configuration > Add Storage server**. Configure the storage server, disk pool and STU to complete the storage server configuration on the target.
- 5** Log on to the source NetBackup primary server (*NBU1*) with the web UI. Navigate to **Storage > Storage configuration > Add Storage server**. Add the storage server.

- 6 While adding the MSDP disk pool, add the replication target. Select the trusted primary server (*NBU2*) and the target storage server.

Add replication targets

Search...

Trusted primary server

☒ punsm6049p-12v01p2.vxindia.ve...

1 Records (1 selected)

Select target storage server

These settings apply only to A.I.R. between NetBackup domains.

Search...

Target storage server	Target volume	Target storage server type
☒ dl365g10p-014v49p3.vxindia.veri...	PureDiskVolume	PureDisk

1 Records (1 selected)

Login credentials for the replication target storage server:

Cancel Add

Scroll down and provide login credentials for the replication target storage servers. Click **Add**.

Add replication targets

×

Select target storage server

These settings apply only to A.I.R. between NetBackup domains.

Search...

Q

Target storage server	Target volume	Target storage server type
☒ dl365g10p-014v49p3.vxindia.veri...	PureDiskVolume	PureDisk

1 Records (1 selected)

Login credentials for the replication target storage server:

Username *

dedupeuser

Password *

.....

Cancel

Add

7 Review the MSDP pool configuration and click **Finish**.

Disk pool configuration

×

Replication configuration operations continue in the background when this window is closed. You can review the status of these operations on the storage review page.

✓

Configuring the disk pool.

✓

✓

Validating the credentials and configuring replication for storage options


1 of 1 successful

✓


Target storage server: dl365g10p-014v49p3.vxindia.veritas.com

Close

- 8 Create an STU to complete the storage server configuration for the source.



- 9 Rlogin to the target NetBackup primary (*NBU2*) and navigate to the earlier configured disk pool. Click on the disk pool name and navigate to replication target option. Click **Add**.



- 10 Select a trusted primary server (*NBU1*) and storage server. Enter the password for the deduplication service on the storage server.

Add replication targets

Note: A replication target is not required to be added for Isolated Recovery Environment in this window.

Select trusted primary server

Search...

Trusted primary server

☒ punsm6049p-12v01p4.vxindia.ve...

1 Records (1 selected)

Select target storage server

These settings apply only to A.I.R. between NetBackup domains.

Search...

Target storage server	Target volume	Target storage server type
☒ dl365g10p-014v54p1.vxindia.veri...	PureDiskVolume	PureDisk

1 Records (1 selected)

Cancel

Add

Add replication targets

Select target storage server

These settings apply only to A.I.R. between NetBackup domains.

Search...

Target storage server

Target volume

Target storage server type

dl365g10p-014v54p1.vxindia.veri...

PureDiskVolume

PureDisk

1 Records (1 selected)

Login credentials for the replication target storage server:

Username *

dedupeuser

Password *

.....

Cancel

Add

11 Verify that the replication target is added.

Veritas NetBackup™

Dashboard

Activity monitor

Recovery

Protection

Workloads

Storage

Disk storage

Media servers

Snapshot Manager

Storage lifecycle policies

Storage units

Tape storage

Catalog

Detection and reporting

tgair_pool

Low water mark

High water mark

80%

95%

Volume options

Volumes

Name

Available space

Total size

Encryption

Replication

PureDiskVolume

28.06 GB

28.32 GB

No

Target

1 Records

Replication targets

Search...

Target storage server

Target storage server type

Source volume

Target volume

dl365g10p-014v54p1.vxindia.veri...

PureDisk

PureDiskVolume

PureDiskVolume

1


NetBackup AIR process:

1. Backup - This is performed on the originating primary server (source domain, *NBU1*). Clients are backed up according to a backup policy that is indicated by the SLP specified in the Policy storage selection. The SLP must include at least one replication operation to a similar storage in the target domain.

2. Duplication - Images are duplicated to the source domain secondary storage server.
3. Replication - Images are replicated from the source domain to the target domain.
4. Device notifies NetBackup - The storage server in the target domain recognizes that a replication event has occurred. It notifies the NetBackup primary server in the target domain.
5. Import - NetBackup imports the image immediately, based on an SLP that contains an import operation. NetBackup can import the image quickly because the metadata is replicated as part of the image.

To create Storage LifeCycle Policy (SLP) for AIR

- 1 Create SLP Policy (import) on the target: cluster.
 - On the target NetBackup primary (*NBU2*), navigate to **Storage > Storage lifecycle policies**. Click **Add**.



- Provide the SLP name. Click **Add** to add the import operation. Select **Import** from the drop-down. Select the destination storage STU. Click **Create**.

New operation

Properties Window

Source storage: Not applicable

Operation
Import

Destination storage attributes
Destination storage
tgt_all_sfu

Retention
Retention type
Target retention

Import
☐ Override default priority
Operation Priority
0
A higher number is greater priority.

Advanced

Cancel Create

- Verify the import operation, storage and SLP name. Click **Create**.

Storage Lifecycle Policy

Storage lifecycle policy Validation report

Storage lifecycle policy name
import_SLP

Data classification
No data classification


Priority for secondary operations
0
A higher number is greater priority.

+ Add

Operation	Window	Target primary	Storage	Storage type	Volume pool
☐ Import	Default_24x7_Window		tgt_all_sfu	PureDisk	


Cancel Create

- Verify that the Import SLP has been added.



2 Create SLP Policy (backup, duplication and replication) on the source NetBackup primary (NBU1):

- Log on to the source NetBackup primary (*NBU1*). Navigate to **Storage > Storage lifecycle policies**. Click **Add**.



- Provide the SLP name. Click **Add**.

Storage Lifecycle Policy

Storage lifecycle policy Validation report

Storage lifecycle policy name: Backup_SLP Data classification: No data classification Priority for secondary operations: 0

You haven't added any operations yet.

State of secondary operation processing: ☒ Active ☐ Postponed

To find impact on policies associated with this SLP due to change in configuration click here. Check for conflicts with backup policies

Cancel Create

- Select **Backup** from the drop-down. Select the destination storage pool as the local MSDP pool which stores the primary data backup. Provide retention type and retention period. Click **Create**.

New operation

Properties

Source storage: Not applicable

Operation: Backup

Destination storage attributes

Destination storage: local_msdp_slt

Volume pool: NetBackup

Media owner: Any


Retention

Retention type: Fixed


Retention period: 2 weeks (Retention Level 1)

Cancel Create


- Select the created backup operation on the SLP window, click on **Add child** to add the duplication operation as a child for backup operation.




- Set the destination storage attributes.




- Select the created duplication operation on the SLP window. Click **Add child** to add the replication operation as a child for duplication operation.




- Set the destination storage attributes. Choose the target primary server and target import SLP policy.



- Verify that the backup, duplication and replication operations are added correctly. Click **Create**.



- 3 On the source NetBackup primary server (*NBU1*), create a policy and select the policy storage as the SLP you created. Configure the policy **Attributes**, **Schedules**, **Clients** and **Backup selections**.



- Run the policy manually. The backup job is responsible for the data backup. The duplication job is responsible for duplicating backed up data on source storage server. The replication job is responsible for the replication of data to the target NetBackup primary server.

Job ID	Type	Client or display name	Job state	Status code	Policy name	Schedule	Schedule type	Elapsed time	State	Active elapsed	Accel
19	Duplication	pernm04lhp12-0719k-vc...	Done	0	SLP_Backup_SLP	Default_24hr7...		00:02:04	Done	00:02:04	
18	Backup	pernm04lhp12-0719k-vc...	Done	0	AIR_Policy1	Full backup	Full backup	00:00:10	Done	00:00:10	
17	Backup	pernm04lhp12-0719k-vc...	Done	0	AIR_Policy1	Full backup	Full backup	00:00:12	Done	00:00:12	
16	Backup	pernm04lhp12-0719k-vc...	Done	0	SLP_Backup_SLP	Default_24hr7...		00:12:22	Done	00:12:22	
15	Duplication	pernm04lhp12-0719k-vc...	Done	0	SLP_Backup_SLP	Default_24hr7...		00:02:04	Done	00:02:04	
14	Backup	pernm04lhp12-0719k-vc...	Done	0	AIR_Policy1	Full backup	Full backup	00:00:10	Done	00:00:10	
13	Backup	pernm04lhp12-0719k-vc...	Done	0	AIR_Policy1	Full backup	Full backup	00:00:12	Done	00:00:12	
12	Duplication	pernm04lhp12-0719k-vc...	Done	0	SLP_Backup_SLP	Default_24hr7...		00:02:04	Done	00:02:04	
11	Backup	pernm04lhp12-0719k-vc...	Done	0	AIR_Policy1	Full backup	Full backup	00:00:10	Done	00:00:10	
10	Backup	pernm04lhp12-0719k-vc...	Done	0	AIR_Policy1	Full backup	Full backup	00:00:12	Done	00:00:12	
9	Duplication	pernm04lhp12-0719k-vc...	Done	0	SLP_Backup_SLP	Default_24hr7...		00:02:07	Done	00:02:07	
8	Backup	pernm04lhp12-0719k-vc...	Done	0	AIR_Policy1	Full backup	Full backup	00:00:10	Done	00:00:10	
7	Backup	pernm04lhp12-0719k-vc...	Done	0	AIR_Policy1	Full backup	Full backup	00:00:10	Done	00:00:10	
6	Backup	pernm04lhp12-0719k-vc...	Done	0	AIR_Policy1	Full backup	Full backup	00:00:12	Done	00:00:11	
5	Image Cleanup		Partial success	1				00:00:02	Done	00:00:02	
4	Image Cleanup		Partial success	1				00:00:01	Done	00:00:01	
3	Image Cleanup		Partial success	1				00:00:01	Done	00:00:01	
2	Image Cleanup		Partial success	1				00:00:01	Done	00:00:01	
1	Image Cleanup		Partial success	1				00:00:01	Done	00:00:01	

- On the target NetBackup primary server (**NBU2**), verify that the **Activity** monitor lists the import job.

Job ID	Type	Client or display name	Job state	Status code	Policy name	Schedule	Schedule type	Elapsed time	State
7	Import from replica		Done	0	SLP_import_SLP	Default_24hr7...		00:00:02	Done
6	Import from replica		Done	0	SLP_import_SLP	Default_24hr7...		00:00:03	Done
5	Image Cleanup		Partial success	1				00:00:01	Done
4	Image Cleanup		Partial success	1				00:00:01	Done
3	Image Cleanup		Partial success	1				00:00:01	Done
2	Image Cleanup		Partial success	1				00:00:01	Done
1	Image Cleanup		Partial success	1				00:00:01	Done

Separating the duplication and replication network

You can separate the duplication and replication network by performing an AIR from the source Access Appliance MSDP server to the target Access Appliance MSDP server. With the default configuration, the AIR traffic follows the same network route as the duplication traffic (eth4). You can separate both the traffic by running the following steps on the Access Appliance cluster.

Consider the following scenario:

- The deduplication server is configured on source Access Appliance cluster (*NBU1*) with primary IP as *10.10.10.11* and FQDN as *primary1.access.com*.
- The deduplication server is configured on the target cluster (*NBU2*) with primary IP as *10.10.10.12* and FQDN as *primary2.access.com*.
- The primary IPs are online on device *eth4*.
- The storage server is configured with *primary1.access.com* on (*NBU1*) and *primary2.access.com* on (*NBU1*).
- *10.10.10.13* and *10.10.10.14* are unused IPs with FQDNs as *vip1.access.com* and *vip2.access.com* respectively.

To separate the duplication and replication network

- 1 Log on to source Access Appliance cluster with management console IP.
- 2 On the source cluster, add a virtual IP. Select a device other than the one already used for deduplication of the primary IP.

```
network> ip addr add <virtual_ip> <netmask> [device] [nodename] [fqdns]
```

Example:

```
network> ip addr add 10.10.10.13 255.255.240.0 eth5
```

- 3 Repeat step 2 with the other IP on the target Access Appliance cluster.

Example:

```
network> ip addr add 10.10.10.14 255.255.255.0 eth5
```

- 4 On the source cluster, add a route to the new target virtual IP.

```
network ip route add all <target_vip> 255.255.255.255 via
<source_gateway> dev <device_name> scope=global
```

Where

target_vip	Newly added virtual IP on target Access Appliance cluster.
source_gateway	Gateway IP of the source cluster
device_name	Device on which the newly added virtual IP is present on the source cluster.

Example:

```
network> ip route add all 10.10.10.14 255.255.255.255 via 10.10.10.1
dev eth5
```

- 5 Add the virtual IP to Access Appliance MSDP server on both the source and target cluster with the respectively added virtual IPs.

```
dedupe> addip <virtual_ip> [config_name]
```

Example:

On the source cluster:

```
dedupe> addip 10.10.10.13 source-engine
```

On target: cluster:

```
dedupe> addip 10.10.10.14 target-engine
```

- 6 Add the mapping of the target virtual IP to the target deduplication primary FQDN on the source deduplication instance.

On the source cluster:

```
dedupe> fqdn add <target_vip> <target_primary_fqdn> [config_name]
```

Where

target_vip	Newly added virtual IP on the target Access Appliance cluster.
target_primary_fqdn	FQDN with which the storage server has been added to the target NetBackup primary server.
config_name	Source deduplication instance's configuration name.

Example:

```
dedupe> fqdn add 10.10.10.14 primary2.access.com source-engine
```

- 7 If you want to replicate from the target cluster to the source cluster, you can perform the same steps on the target cluster with the source cluster IP.

On the target cluster:

```
dedupe> fqdn add 10.10.10.13 primary1.access.com target-engine
```

Setting up Auto Image Replication (AIR) between two Access Appliances which have secure communication enabled

You can set up AIR between two Access appliances which have secure communication enabled.

Perform the following steps on the AIR source and target Access Appliance deduplication servers when secure communication is enabled:

- Generate certificates using the remote NetBackup primary as CA. Use the `setting certificate get-CA-certificate` command in the MSDP restricted shell to get the CA certificate of the remote NetBackup primary server.
- Create a host ID-based certificate using the `setting certificate get-certificate` command in the MSDP restricted shell.

- This needs to be done on both the Access clusters participating in the AIR workflow.

You can now set up trust between the NetBackup domains and set up AIR.

See [“Support for NetBackup Auto Image Replication”](#) on page 94.

For detailed steps on setting up trust, see the *NetBackup™ Deduplication Guide* on [SORT](#).

Note: Access Appliance 8.6 release supports only NetBackup CA for MSDP secure communication. External CA in MSDP is not supported.

NetBackup Dedupe Direct for Oracle

This chapter includes the following topics:

- [Support for NetBackup Dedupe Direct for Oracle](#)

Support for NetBackup Dedupe Direct for Oracle

Access Appliance supports NetBackup Dedupe Direct for Oracle, which is a lightweight plug-in for performing RMAN backups to MSDP storage (Veritas Data Deduplication pool). This enables client-side deduplication to minimize network traffic and improves overall backup speed. The integration is enabled with RMAN and empowers the DBA to perform all backup and restore operations.

This solution helps Oracle DBAs to control the entire protection and lifecycle stages without a NetBackup client. Oracle DBAs can deploy this plug-in with owned system permission (non-root) and perform backup or restore jobs. NetBackup Dedupe Direct for Oracle is a separate package and can be installed on the Oracle host.

Prerequisites for installing NetBackup Dedupe Direct for Oracle:

- Veritas Data Deduplication should be configured.
- The *spa* user with *app* role should be configured from Access Appliance CLISH.

```
# dedupe adduser username [role] [config_name]
```

username	Username of the deduplication server
----------	--------------------------------------

role	Enter one of the supported roles (app admin) [admin]
------	--

config_name	Name of the deduplication service. Default value is vxdefault.
-------------	--

You can use the `dedupe listuser` command to show the user and its role.

```
# dedupe listuser
```

You can install the NetBackup Dedupe Direct for Oracle on the Oracle database server using the Oracle user.

- Change the ownership permission of the plugin tar file.

```
chown oracle:oinstall netbackup-dedupe-direct-for-oracle.linuxR_x86.tar
```

- Login with the Oracle user and extract the tar file.
- Change the directory to the extracted folder and install the NetBackup Dedupe Direct for Oracle using the following command:

```
./sbt_config -install
```

Supported versions:

- Oracle database version from 19C, 21C
- OS version: RHEL 7.x or RHEL 8.x

For more information on installing the NetBackup Dedupe Direct for Oracle plug-in on the Oracle database server, see the *NetBackup Dedupe Direct for Oracle* chapter in the *NetBackup™ for Oracle Administrator's Guide* on [SORT](#).

Configuring MSDP-C with Access Appliance

This chapter includes the following topics:

- [Starting the Access S3 server and provisioning the S3 bucket](#)
- [Configuring Access S3 bucket with NetBackup MSDP-C](#)

Starting the Access S3 server and provisioning the S3 bucket


You can use the following procedure to start the S3 server and provision the S3 bucket.

Note:

- The size of the bucket for MSDP-C should not be more than 1 PiB.
- For NetBackup MSDP-C, WORM-enabled bucket is only supported with Access Appliance Compliance mode.

To start the S3 server and provision the S3 bucket using the Access GUI


- 1 In **S3 buckets for NetBackup** area, click **Provision storage**.



Or

Navigate to **Quick Actions > Provision Storage > Storage for NetBackup > S3 without cloud**. Click **Next**.

- 2 Click **Configure S3 server**.



- 3 Select the storage pool and click **Start S3 server**. If you want to enable SSL for S3 server. then turn on the **Enable SSL** toggle button.

Start S3 server

This step configures S3 server and make it online. This configuration enables the Veritas Access server to create S3 bucket(s) on-premises. **This step will also start the server and bring it online after the default settings are configured.**

Select storage pool :

	Storage pool	Total size	Available size
☒	pool_default	250.00 GB	201.50 GB

Items per page: 2 1 - 1 of 1 < >

i Access keys and secret keys should be regenerated if this setting is changed in future.

☐ Enable SSL

Start S3 server

- 4 Click **Activate LTR Policy**.

S3 Setup

Configure storage

✓

Creates a storage pool required for storage provisioning

Configure S3 Server

✓

Configures S3 Server and make it online

Activate LTR Policy

✓

Activates On-Premises LTR policy

Generate S3 Keys

☁

Generates Access and Secret keys required for storage provisioning

Provision Storage

🗄

Creates on-premises bucket to store NetBackup images

Back

- 5 If you want to configure NetBackup MSDP-C on a non-WORM bucket, select **NetBackup LTR S3 Bucket policy** and click **Complete prerequisites**.

If you want to configure NetBackup MSDP-C on a WORM bucket, select **NetBackup LTR S3 Bucket (WORM) policy** and click **Complete prerequisites**.

Select and activate the policy for storage provisioning.

Complete prerequisites

Policy name	Description	Capabilities	Prerequisite
☐ NetBackup LTR S3 bucket	This policy provisions an S3 bucket to stor...	Fault Tolerance	✓ Storage pool created
☒ NetBackup LTR S3 Bucket (WORM)	This policy provisions an S3 bucket to stor...	Fault Tolerance,WORM	✓ Storage pool created ✓ Compliance mode configured

Items per page: 5 1 - 2 of 2 < >

Activate policy

- 6** For non-WORM use case, select the storage pool. Click **Activate policy**.
 For WORM use case, select the storage pool.

Complete following prerequisite to activate NetBackup LTR S3 Bucket (WORM)

Select storage pool :Selected storage pools:1, total size:201.50 GB

☒	Storage pool ↑	Total size	Available size
☒	pool_default	250.00 GB	201.50 GB

Items per page: 5 1 ~ 1 of 1 < > >>

- 7** Perform this step only for WORM use case. During activation, under **Set WORM**, the minimum and maximum retention period should be left as blank for MSDP-C buckets. This will be applied from NetBackup during configuration. Click **Activate policy**.

Set WORM

Configure the retention time range that Appliance Administrators will be allowed to select from when setting WORM retention time. The minimum retention time you can set is 1 hour and the maximum is 60 years.

Minimum retention time Maximum retention time

0 0 Hours ▾

- 8** Generate S3 keys. You can use any local/AD/LDAP user credentials. Under **Name**, enter the username for which you want to create keys and under **Password**, specify the user password.
 Click on **Generate S3 Keys**.

S3 Setup

Generate S3 access key and secret key.

i You need an access and secret key to access the S3 buckets.

Enter user credentials to generate the keys.

Name *

admin

Password *

- 9 Note down the access key and secret key.

S3 Setup

Configure storage

Configure S3 Server

Activate LTR Policy

Generate S3 Keys

Provision Storage

Generate S3 access key and secret key.

Generating S3 keys

View details

Save the access key and secret key in a safe location for future use. These keys are required for storage provisioning.

Access key


NTM0MDMyMzFIMDZjZml

Secret key

YmRjYmRiZGYxZGNmMTRkMzM4YWE2OTZIMDZjYTU

Finish

10 Click Provision Storage.



11 Select the storage policy according to the activation, set the bucket size as per the requirement, and specify the access and the secret keys.

Provision storage

i This step provides storage to store NetBackup images. You have to enter the access and secret keys created in the previous step.

On-premises storage options

Select storage policy
 NetBackup LTR S3 Bucket (WORM) ▼

Bucket size
 100 GB ▼

Access key

Secret key

Configure storage

Configuring Access S3 bucket with NetBackup MSDP-C

This section describes how to configure an Access S3 bucket with NetBackup MSDP-C.

Note: The steps and screen shots provided in this section are for MSDP-C configuration with NetBackup version 10.1. For any other version of NetBackup, refer to the NetBackup™ Deduplication Guide on [SORT](#).

Pre-requisite to configuring MSDP-C with Access S3 server configured in SSL mode

If Access S3 is started with SSL, perform the following steps to add the Access S3 certificate to the OS trust store of NetBackup media servers:

1. Copy the certificate obtained after running the `objectaccess certificate show` CLISH command to the `s3cert.pem` file. Append the following file with the obtained certificate on the NetBackup servers:

For NetBackup 9.x: `/usr/openv/var/global/wmc/cloud/cacert.pem`

For NetBackup 10.x: `/usr/openv/var/global/cloud/cacert.pem`

2. Run the following commands on the NetBackup primary server:

```
cp s3cert.pem /usr/share/pki/ca-trust-source/anchors/
```

```
update-ca-trust force-enable
```

```
update-ca-trust extract
```

```
update-ca-trust
```


Pre-requisite to configuring MSDP-C with WORM

If you want to configure NetBackup MSDP-C WORM with Access S3 WORM bucket, then you must add the volume to the bucket using the `msdpclutil` command from the NetBackup server's bash.

To add volume to the bucket

- 1 Log in to NetBackup primary and set the following environment variables:


```
export MSDPC_ACCESS_KEY=<access-key>
export MSDPC_SECRET_KEY=<secret-key>
export MSDPC_REGION=us-east-1
export MSDPC_PROVIDER=vtas-access
export MSDPC_ENDPOINT=<endpoint name> (find it from GUI Dashboard.
Go to Settings > S3 management. For example: http://<cluster>:8143)
```



- 2 Run the following command to create a volume. Use the bucket name from the Access Appliance GUI (**Provisioned storage > Buckets**).

```
/usr/openv/pdde/pdcr/bin/msdpclutil create --bucket <bucketname>
--volume <volumename> --mode COMPLIANCE --max <maxdays>
--min <mindays> --live <date>
```


Refer to the `msdpclutil` command's help for more details.




Log in to the NetBackup GUI and perform the following steps to configure MSDP-C with Access S3 bucket, create disk pool and create STU.

To configure MSDP-C with Access S3 bucket

- 1 Go to **Storage > Storage Configuration > Storage servers** and add the MSDP storage server on the NetBackup media server using which you want to configure MSDP-C.
- 2 Go to **Storage > Storage Configuration > Disk pools** and add a new pool.



- 3 Select a storage server and enter the name of pool and click **Next**.



4 Select **Add volume**.

For MSDP-C with WORM use case, enter the volume name which was provided as part of `msdpclldutil` command during S3 configuration.

For MSDP-C with non-WORM use case, you can enter any volume name.

Volume

Add volume

Volume name *

testvol

5 Under **Select cloud storage provider**, search for Veritas Access.

Select cloud storage provider

1 item selected

Cloud storage provider	Description	Storage API type
☐ Microsoft Azure	Microsoft Azure Storage Service	Azure
☐ Microsoft Azure Government	Microsoft Azure Storage Service	Azure
☐ NetBackup Recovery Vault Azure	NetBackup Recovery Vault Azure Storage Service	Azure
☐ NetBackup Recovery Vault Azure Government	NetBackup Recovery Vault Azure Storage Service	Azure
☒ Veritas Access	Veritas Access Scale-out software-defined storage	S3

5 Records (1 selected)

Cancel

Select

- 6 Add **Region name** and **Location Constraint** as *us-east-1*. Service URL should be endpoint name (<http://<cluster>:8143>). Note that the endpoint name refers to only the *<cluster>* part listed in the URL and not the entire URL.

You can find the Access S3 endpoint URL by running the `objectaccess show` command from the Access CLISH. Specify the HTTP and HTTPS port number which is mentioned in S3 endpoint URL. Click **Add** to proceed.

The screenshot shows a 'Add a region' dialog box with the following fields and values:

- Region name ***: us-east-1
- Location constraint ***: us-east-1
- Service URL ***: s3.hwclust3350
- Endpoint access style**: Path style (dropdown menu)
- HTTP port ***: 8143
- HTTPS port ***: (empty)

Buttons at the bottom right: Cancel, Add.

7 Add access key and secret key.

Access details for the account

Access key ID *

NDvjMTlmAGUSNzBhMDQ

Secret access key *

.....

Advanced settings

Security

☐ Use SSL

Proxy

☐ Use proxy server

8 Uncheck the **Use SSL** option, if Access S3 server is configured on HTTP. If Access S3 server is configured with SSL/HTTPS then select the **Use SSL** option with **Authentication and data transfer**.

Security

☒ Use SSL

☐ Authentication only

☒ Authentication and data transfer

☐ Check certificate revocation (IPv6 not supported for this option)

- 9 You can either enter an existing bucket name or select a cloud bucket by clicking **Retrieve list**. Click **Next**.

Cloud buckets

☒ Enter an existing cloud bucket name

Cloud bucket name
 88c143b8-7917-4aeb-86cf-8a008d01a1des3bucket

☐ Select or create a cloud bucket

☒ Select or create a cloud bucket

Complete all required fields to view available cloud buckets.

Retrieve list

- 10 Skip the addition of replication target by clicking **Next**.

Add MSDP disk pool

✓ Disk pool options ✓ Volumes 3 Replication Review


You have not added any replication targets

Add


Activate Windows
 Go to Settings to activate Windows.

Cancel Previous Next


- 11
- Review the configuration and click **Finish**. The disk volume and disk pool will get configured.



- 12
- Go to **Add storage unit** and select the created pool. For MSDP-C with WORM use case, check the **Lock until expiration** option to enable WORM.




- 13 Create storage lifecycle policies. Go to **Storage > Storage lifecycle policies**.



- 14 Add Storage lifecycle policy (SLP) with backup on the primary STU and deduplication on the secondary STU (the STU which is configured with MSDP-C).

For MSDP-C with WORM use case, the retention period selected should be between the maximum and minimum value given while configuring the volume using the `msdpclutil` command.

- 15 After defining the SLP, create a policy to perform backup. Go to **Protection > Policies**.



- 16 Add a policy name. Under **Policy storage**, select the SLP created in the previous step.

Attributes

Schedules

Clients

Backup selections

Policy name

policy1

Policy type

Standard

Destination

Data classification

No data classification

Policy storage

slp1 (Storage lifecycle policy)

Policy volume pool

- 17 Go to **Schedules** to create automatic schedules.

Attributes

Start window

Exclude dates

Name *

full

Type of backup *

Full Backup

☐ Synthetic backup

☐ Accelerator forced rescan

Destination

☐ Multiple copies

☐ Override policy storage selection

slp1 (Storage lifecycle policy)

Schedule type

☐ Calendar

☐ Retries allowed after the run day

☒ Frequency

10

hours

Retention *

Determined by SLP

Cancel

Add

18 Add the client.

Add client - Policy pol1

Server

r6515-002v027.vxindia.veritas.com

Client name

r6515-002v027.vxindia.veritas.com

☒ Detect client operating system

Hardware and operating system

HP-UX-IA64, HP-UX11.31

Cancel

Add

- 19
- In the **Add backup selection** dialog box, select the file or directory that needs to be backed up.

Add backup selection

Server

ssnasdl380-02-v03.vxindia.veritas.com

Construct a list of pathnames (and directives, if applicable) to add to the selection list.

Pathname or directive

▼ | 📁 +

List of pathnames and directives to add to the selection list.

/data

Cancel

Add

20

Click **Create**.

21

Activate the policy.

Dashboard

Activity monitor

Recovery

Protection

Protection plans

Policies

Workloads

Cloud

Policies

+ Add

Search

☐	Name	Type	Data classification	Storage	Volume paths	Media owner	Checkpoint interval (min)	Job Policy	Priority
☐	pol1	Standard		s3	netbackup	any	0	0	1

Edit

Copy policy

Delete

Activate

Interval backup

Migrating the NetBackup images from existing storage to Veritas Access storage

This chapter includes the following topics:

- [Migrating NetBackup images from CloudCatalyst/S3/MSDP-C to Veritas Data Deduplication storage](#)
- [Migrating the NetBackup images from the Veritas Data Deduplication STU to the MSDP-C STU](#)

Migrating NetBackup images from CloudCatalyst/S3/MSDP-C to Veritas Data Deduplication storage

You can migrate the backup images from the CloudCatalyst/S3/MSDP-C Storage Unit (STU) to the Veritas Data Deduplication STU.

Table 6-1

When CloudCatalyst is configured as the secondary/LTR storage in the NetBackup policy

[To migrate the NetBackup images from the CloudCatalyst STU to the Veritas Data Deduplication STU when CloudCatalyst is configured as the secondary storage in the NetBackup SLP](#)

Table 6-1 (continued)

When CloudCatalyst is configured as the primary storage in the NetBackup policy	To migrate the NetBackup images from the CloudCatalyst STU to the Veritas Data Deduplication STU when CloudCatalyst is configured as the primary storage in the NetBackup policy
When S3 is configured as the primary storage in the NetBackup policy	To migrate the NetBackup images from the S3 STU to the Veritas Data Deduplication STU when S3 is configured as the primary storage in the NetBackup policy
If you want to migrate the NetBackup images from the MSDP-C STU to the Veritas Data Deduplication STU	To migrate the NetBackup images from the MSDP-C STU to the Veritas Data Deduplication STU

To migrate the NetBackup images from the CloudCatalyst STU to the Veritas Data Deduplication STU when CloudCatalyst is configured as the secondary storage in the NetBackup SLP

- 1 Configure the Veritas Data Deduplication STU.
 - 2 Change the SLP to point to the Veritas Data Deduplication STU instead of the CloudCatalyst STU so that the new images get duplicated on the Veritas Data Deduplication STU.
 - 3 As optimized duplication from CloudCatalyst to Veritas Data Deduplication is not supported, the administrator is required to add the `RESUME_ORIG_DUP_ON_OPT_DUP_FAIL = TRUE` flag in the `/usr/openv/netbackup/bp.conf` file of both the primary and media server.
 - 4 Go to the **NetBackup Management > Catalog** and select the **Action as Duplicate** from the drop-down menu.
 - 5 In the **NetBackup Management > Catalog**, right-click on the image that is stored on the CloudCatalyst STU and that you want to duplicate on the Veritas Data Deduplication STU. Select **Duplicate**.
 - 6 In the pop-up window that appears, select the Veritas Data Deduplication STU and click **OK**.
 - 7 In the confirmation screen, click **OK** to start duplication.
- You can check the progress of the image duplication in the **NetBackup Activity Monitor**.

- 8 After the completion of duplication, the status of the duplication job appears as **The requested operation was partially successful**. In the job detail status, some error messages appear that are expected.
- 9 After duplication is complete, expire the backup image stored on the CloudCatalyst STU.

To migrate the NetBackup images from the CloudCatalyst STU to the Veritas Data Deduplication STU when CloudCatalyst is configured as the primary storage in the NetBackup policy

- 1 Configure the Veritas Data Deduplication STU.
- 2 Change the SLP to point to the Veritas Data Deduplication STU instead of the CloudCatalyst STU so that the new images get created on the Veritas Data Deduplication STU.
- 3 As optimized duplication from CloudCatalyst to Veritas Data Deduplication is not supported, the administrator is required to add the `RESUME_ORIG_DUP_ON_OPT_DUP_FAIL = TRUE` flag in the `/usr/openv/netbackup/bp.conf` file of both the primary and media server.
- 4 Go to the **NetBackup Management > Catalog** and select the **Action as Duplicate** from the drop-down menu.
- 5 In the **NetBackup Management > Catalog**, right-click on the image that is stored on the CloudCatalyst STU and that you want to duplicate on the Veritas Data Deduplication STU. Select **Duplicate**.
- 6 In the pop-up window that appears, select the Veritas Data Deduplication STU and click **OK**.
- 7 In the confirmation screen, click **OK** to start duplication.

You can check the progress of the image duplication in the **NetBackup Activity Monitor**.
- 8 After the completion of duplication, the status of the duplication job appears as **The requested operation was partially successful**. In the job detail status, some error messages appear that are expected.
- 9 After duplication is complete, expire the backup image stored on the CloudCatalyst STU.

To migrate the NetBackup images from the S3 STU to the Veritas Data Deduplication STU when S3 is configured as the primary storage in the NetBackup policy

- 1 Configure the Veritas Data Deduplication STU.
- 2 Change the SLP to point to the Veritas Data Deduplication STU instead of the S3 STU so that the new images get created on the Veritas Data Deduplication STU.
- 3 Go to the **NetBackup Management > Catalog** and select the **Action** as **Duplicate** from the drop-down menu.
- 4 In the **NetBackup Management > Catalog**, right-click on the image that is stored on the S3 STU and that you want to duplicate on the Veritas Data Deduplication STU. Select **Duplicate**.
- 5 In the pop-up window that appears, select the Veritas Data Deduplication STU and click **OK**.
- 6 In the confirmation screen, click **OK** to start duplication.
You can check the progress of the image duplication in the **NetBackup Activity Monitor**.
- 7 After duplication is complete, expire the backup image stored on the S3 STU.

To migrate the NetBackup images from the MSDP-C STU to the Veritas Data Deduplication STU

- 1 Configure the Veritas Data Deduplication STU.
- 2 Change the existing SLP to point to the Veritas Data Deduplication STU instead of the MSDP-C STU so that the new images will get duplicated to the Veritas Data Deduplication STU.
- 3 Go to **NetBackup Management > Catalog** and select the **Action** as **Duplicate** from the drop-down menu.
- 4 In **NetBackup Management > Catalog**, right-click on the image that is stored on the MSDP- C STU which you want to duplicate on the Veritas Data Deduplication STU. Select **Duplicate**.
- 5 In the pop-up window that appears, select the Veritas Data Deduplication STU and click **OK**.
- 6 In the confirmation screen, click **OK** to start duplication. You can check the progress of the image duplication in the **NetBackup Activity Monitor**.
- 7 After duplication completes, expire the backup image stored on the MSDP-C STU.

Migrating the NetBackup images from the Veritas Data Deduplication STU to the MSDP-C STU

You can migrate the NetBackup images from the Veritas Data Deduplication STU to the MSDP-C STU.

To migrate the NetBackup images from the Veritas Data Deduplication STU to the MSDP-C STU which is configured on bucket of Veritas Access Objectaccess server

- 1 Configure MSDP-C with Veritas Access Objectaccess bucket.
- 2 Change the existing SLP to point to the MSDP-C STU instead of the Veritas Data Deduplication STU so that the new images will get duplicated to the MSDP-C STU.
- 3 Go to **NetBackup Management > Catalog** and select the **Action** as **Duplicate** from the drop-down menu.
- 4 In **NetBackup Management > Catalog**, right-click on the image that is stored on the Veritas Data Deduplication STU which you want to duplicate on the MSDP-C STU. Select **Duplicate**.
- 5 In the pop-up window that appears, select the MSDP-C STU and click **OK**.
- 6 In the confirmation screen, click **OK** to start duplication. You can check the progress of the image duplication in the **NetBackup Activity Monitor**.
- 7 After duplication completes, expire the backup image stored on the Veritas Data Deduplication STU.

Configuring Access Appliance with the NetBackup client

This chapter includes the following topics:

- [Configuring Access Appliance for NetBackup client](#)
- [Installing the NetBackup client add-on packages](#)
- [Prerequisites for configuring the NetBackup client](#)
- [Configuring the NetBackup client](#)
- [Adding an optional media server](#)
- [Adding multiple primary servers](#)
- [Adding file systems to the backup configuration](#)
- [Removing file systems from backup configuration](#)
- [Changing the virtual IP address used by the NetBackup client](#)
- [Upgrade the NetBackup client](#)
- [Unconfiguring the NetBackup client](#)
- [Enabling Veritas Data Deduplication catalog backup with NetBackup client](#)
- [Disabling Veritas Data Deduplication catalog backup from NetBackup client](#)
- [Displaying the status of NetBackup services](#)

- [Configuring backup operations using NetBackup or other third-party backup applications](#)
- [Restoring filesystems backed up with NetBackup client](#)

Configuring Access Appliance for NetBackup client

Access Appliance 8.6 supports NetBackup client versions 10.0, 10.1.1, 10.2, 10.3, 10.4, 10.5, 10.5.0.1, 11.0, 11.0.0.1 and 11.1.0.2.

Certificates are deployed as a part of the NetBackup client configuration to enable secure communication between NetBackup primary and NetBackup client to perform backup and restore related operations successfully.

To back up your data present on filesystems which are present on Access Appliance with NetBackup, you must register the installed and configured NetBackup primary server with Access Appliance.

To configure NetBackup Client for Access Appliance, perform the following tasks in the following order:

Table 7-1 Workflow for configuring the NetBackup client

Task	Details
Make sure that the prerequisites are met.	See “Prerequisites for configuring the NetBackup client” on page 145.
Configure the NetBackup client.	See “Configuring the NetBackup client” on page 145.
Add an optional media server.	See “Adding an optional media server” on page 146.
Add multiple primary servers.	See “Adding multiple primary servers” on page 147.
Add file systems for backup.	See “Adding file systems to the backup configuration” on page 148.

You can also perform the following tasks:

Table 7-2 Other operations that you can perform on the NetBackup client

Task	Details
Unconfigure the NetBackup client.	See “Unconfiguring the NetBackup client” on page 150.

Table 7-2 Other operations that you can perform on the NetBackup client
(continued)

Task	Details
Remove file systems from backup configuration.	See “Removing file systems from backup configuration” on page 148.
Change the virtual IP address used by NetBackup client.	See “Changing the virtual IP address used by the NetBackup client” on page 149.
Upgrade the NetBackup client.	See “Upgrade the NetBackup client” on page 149.
Enable Veritas Data Deduplication catalog backup with NetBackup client.	See “Enabling Veritas Data Deduplication catalog backup with NetBackup client” on page 150.
Disable Veritas Data Deduplication catalog backup with NetBackup client.	See “Disabling Veritas Data Deduplication catalog backup from NetBackup client” on page 154.
Remove file systems from backup configuration.	See “Removing file systems from backup configuration” on page 148.
Display the current status of the NetBackup client.	See “Displaying the status of NetBackup services” on page 155.
Verify that Access Appliance is configured with the NetBackup client.	See “Displaying the status of NetBackup services” on page 155.

If Access Appliance is configured with IPv6 addresses, you have to configure IPv6 support for the NetBackup host as well.

Perform the following steps to configure IPv6 support for the NetBackup host:

- Set the `IP_ADDRESS_FAMILY` option in the NetBackup `bp.conf` file for the host to `AF_UNSPEC`.

```
# bpsetconfig IP_ADDRESS_FAMILY = AF_UNSPEC
```

- You can view the current setting by running the `bpgetconfig` command.

```
# bpgetconfig
IP_ADDRESS_FAMILY = AF_UNSPEC
```

- Restart the services after making this change.

Installing the NetBackup client add-on packages

The supported NetBackup clients can be installed as add-on packages.

Considerations while installing the NetBackup client add-on packages:

- If the nodes are configured in the cluster, the NetBackup client add-on packages are installed on both nodes.
- If the nodes are not configured in the cluster, you have to download and install the NetBackup client add-on packages on both the nodes.
- To create a file system, refer the section *Creating a file system* section in the *Access Appliance Administrator's Guide* on [SORT](#). The file system used for the NetBackup client configuration should be a minimum of 10 GB with a 5-way striped layout. The client configuration uses 1 GB or less of space of the file system, and the remaining space is used for NetBackup Accelerator track logs. The NetBackup file system is used to store NetBackup client configuration and NetBackup Accelerator track logs. You can grow the file system as per your requirement.
- The virtual IP should be pingable and online. You can check the status of the virtual IP using the `network ip addr show` command.

To install the add-on packages

- 1 Download the add-on packages for the NetBackup client from the Download Centre:
 - VRTSaccess-app-ADDON-NBUClient-10.0-1.2.1-1.x86_64.rpm
 - VRTSaccess-app-ADDON-NBUClient-10.1.1-1.2.1-1.x86_64.rpm
 - VRTSaccess-app-ADDON-NBUClient-10.2-1.2.1-1.x86_64.rpm
 - VRTSaccess-app-ADDON-NBUClient-10.3-1.2.1-1.x86_64.rpm
 - VRTSaccess-app-ADDON-NBUClient-10.4-1.2.x86_64.rpm
 - VRTSaccess-app-ADDON-NBUClient-10.5-1.2.1-1.x86_64.rpm
 - VRTSaccess-app-ADDON-NBUClient-10.5.0.1-1.2.1-1.x86_64.rpm
 - VRTSaccess-app-ADDON-NBUClient-11.0-1.2.1-1.x86_64.rpm
 - VRTSaccess-app-ADDON-NBUClient-11.0.0.1-1.2.1-1.x86_64.rpm
 - VRTSaccess-app-ADDON-NBUClient-11.1.0.2-1.2.1-1.x86_64.rpm
- 2 Go to the Appliance CLISH. Enter the `system software share open` command.
- 3 Mount the path specified as a part of the `share open` command on NFS client.

- 4 Copy the add-on packages to the `/inst/patch/appliance/available/` folder using NFS.

For more details, see the *Downloading software updates to an Access appliance node using a client share* section in the *Veritas Access Appliance Upgrade Guide*.

- 5 In the Appliance CLISH, enter the `system software downloaded` command to confirm that the add-on packages are downloaded successfully.
- 6 Enter the `system software install-update update-name=<add-on pkg name>` command to install the package.
- 7 Enter the `system software installed-addons` command to confirm that the package has been installed.
- 8 Enter the `system software share close` command to close the shares.

To remove the add-on packages

- 1 Go to the Appliance CLISH. Enter the `system software rollback-update update-name=<pkg name>` command to uninstall the packages.
- 2 Enter the `system software installed-addons` command to confirm that the package has been removed.

Prerequisites for configuring the NetBackup client

Before configuring the NetBackup client for Access Appliance, you must have completed the following:

- Add the valid licenses on the NetBackup primary server.
- Make sure that host names of the primary server and the NetBackup client installed on Access Appliance can be resolved and the reverse lookup is working.
- The file system used for the NetBackup client configuration should be a minimum of 10 GB with a 5-way striped layout. The client configuration uses 1 GB or less of space of the file system, and the remaining space is used for NetBackup Accelerator track logs. The size of the file system must be decided based on the backup policies that have been created. See https://www.veritas.com/support/en_US/article.100045746 for more details.

Configuring the NetBackup client

The `backup config` CLISH command configures the specified version of the NetBackup client. The NetBackup client version must be 8.1 or higher. Required certificates are generated as part of this command.

To configure the NetBackup client

- ◆ Configure the NetBackup client.

```
Backup> config <virtual_name> <virtual_IP> fs_name  
client_version auth_token primary_server_name [primary_server_ip]
```

virtual_name	Hostname to be used by a NetBackup client configuration on the nodes. Make sure that name can be resolved through DNS and its IP address can be resolved back to name through a DNS reverse lookup. Also make sure that name resolves to the IP address parameter <i>ipaddr</i> that is provided for configuration.
virtual_IP	Virtual IP address to be used by the NetBackup client configuration on the nodes. Make sure that IP address can be resolved back to the name we use for configuration.
fs_name	Name of the file system used to store NetBackup client configuration.
client_version	NetBackup client version to be installed.
auth_token	Authentication token from NetBackup primary server to enable backup/restore in a secure environment.
primary_server_name	FQDN of the primary server
primary_ip	IP of the server to be used as the NetBackup primary server. Specify this parameter if DNS is not configured in the cluster. This is an optional parameter.

See the `backup_config(1)` man page for detailed examples.

Adding an optional media server

You can optionally add a media server to be used by the NetBackup client. If the media server name is not added, the local NetBackup client uses the NetBackup primary server as the NetBackup media server. But it is not mandatory that the NetBackup primary server be the NetBackup media server. In production environments, the NetBackup media server is separate from the NetBackup primary server.

To add an optional NetBackup media server

- ◆ Add an optional NetBackup media server.

```
Backup> netbackup media-server add media_server_name
[media_server_IP]
```

For example:

```
Backup> netbackup media-server add FQDN_of_primary_server
```

Make sure that the server hostname can be resolved through DNS, and its IP address can be resolved back to the NetBackup media server hostname through the DNS reverse lookup. If DNS is not configured in the cluster, specify the *media_server_IP* in the `/etc/hosts` file of all nodes of the cluster. to create a mapping between the server hostname and its IP address.

You can also delete an already configured NetBackup media server.

To delete an already configured NetBackup media server

- ◆ Delete an already configured NetBackup media server.

```
Backup> netbackup media-server delete server
```

See the `backup_netbackup(1)` man page for detailed examples.

Adding multiple primary servers

The NetBackup primary server is registered with Access Appliance during NetBackup client configuration so that it can communicate with the Access Appliance. If necessary, you can add multiple primary servers with the same NetBackup client.

To add an additional NetBackup primary server

- ◆ Add an additional new NetBackup primary server to the existing client configuration. Required certificates of the new primary server are generated as part of this command.

```
Backup> netbackup primary-server add auth_token primary_server
[primary_ip]
```

<code>primary_server</code>	FQDN of the primary server
<code>auth_token</code>	Authorization token to generate th host ID based certificates
<code>primary_ip</code>	IP address of the primary server. This is optional field when the primary server is not resolvable to the FQDN,

To delete a NetBackup primary server

- ◆ Remove a primary server from the existing client configuration.

```
Backup> netbackup primary-server del primary_server
```

where *primary_server* is the FQDN of the primary server.

CA and host ID certificates for the specified primary server is deleted and its associated entry from the configuration files also removed.

See the `backup` man page `fs` module for detailed examples.

Adding file systems to the backup configuration

You can add the file systems required to be backed up with the NetBackup client. Once the file system is configured for backup, it is available for backup and restore operation from NetBackup primary server side.

To add file systems for backup

- ◆ Add the file system to the NetBackup client configuration to enable backup and restore operations from NetBackup primary server.

```
Backup> fs add fs_name
```

See the `backup` man page `fs` module for detailed examples.

Removing file systems from backup configuration

You can remove the file system from the NetBackup client configuration. Once the file system is removed from the backup configuration, it is not available for backup and restore operations from NetBackup primary server side.

To remove the file systems for backup

- ◆ Remove the file system from the NetBackup client configuration.

```
Backup> fs remove fs_name
```

See the `backup` man page `fs` module for detailed examples.

Changing the virtual IP address used by the NetBackup client

You can alter the virtual IP address of the NetBackup client. This address is a highly-available virtual IP address in the cluster.

Note: Configure the virtual IP address using the `backup virtual-ip set` command so that it is different from all of the other virtual IP addresses, including the console server IP address and the physical IP addresses that are used to install Access Appliance. Use the `network ip addr show` command to display the currently assigned virtual IP addresses on Access Appliance.

See the `backup_virtual-ip(1)` man page for detailed examples.

To change the virtual IP address used by NetBackup

- ◆ Configure the virtual IP address of NetBackup on Access Appliance.

```
Backup> virtual-ip set ipaddr [device]
```

This command configures the built-in NetBackup client to use *ipaddr* as its virtual IP address and *device* as the network interface. A network mask can be provided with the IP address as the prefix length for the *ipaddr*.

Ensure that the NetBackup primary server is able to resolve the NetBackup client's FQDN to the new IP.

Upgrade the NetBackup client

You can upgrade the NetBackup client on your appliance by specifying the version. Ensure that the client image version that is specified is pre-installed on the cluster. Also, the version to be installed must be higher than the already configured version. Download the Add-Ons before you begin the upgrade operation.

To upgrade the NetBackup client to a specified version

- ◆ Upgrade the NetBackup client on the appliance by specifying the version.

```
Backup> upgrade <version>
```

Unconfiguring the NetBackup client

The `backup unconfig` command removes the existing NetBackup client configuration.

To unconfigure the NetBackup client

- ◆ Removes the NetBackup client configuration for the client,

```
Backup> unconfig
```

See the `backup_unconfig(1)` man page for detailed examples.

Enabling Veritas Data Deduplication catalog backup with NetBackup client

You can configure the built-in NetBackup client to perform a backup and restore of the deduplication catalog.

To enable the Veritas Data Deduplication catalog backup with NetBackup client

- 1 Enable the NetBackup client to perform a backup and restore of the deduplication catalog.

```
Backup> dedupe-catalog enable config_name
```


`config_name` is an identifier and must be provided for any operation to be performed on the deduplication server if more than one deduplication server is running in the cluster.

- 2 Create a new policy. Navigate to the **Attributes** tab, select *MSDP local STU (Storage unit)* in **Policy storage** drop-down. You have to create a local STU as data is written to the media server.


Note: During policy creation, you must select the **Policy Type** as Standard and also select the **Cross Mount Point** check box for successful backup operation.

- 3 Navigate to the **Schedules** tab and create a new schedule as per your requirement.

- 4 Navigate to the **Clients** tab. Click **New** and add the client name (FQDN) which is mentioned during backup configuration (Virtual Name). Select the check box **Detect client operating system** and click **OK**.




- 5 Navigate to the **Backup Selection** tab. Click **Browse**.



Double click on the `vx` folder. The `dedupe_catalog` folder appears. Double-click on the `dedupe_catalog` folder, and then select the `msdp-engine` folder. It lists all the files and directories which are to be backed up. Click **OK**.

Note: When you run the `dedupe-catalog enable` command from CLISH, all the catalog files and folders that are required for back up are automatically added to the `config_name` folder under the `dedupe_catalog` folder. You need not select any files manually. Here, `config_name` is the name of the deduplication server.



- 6 If there are multiple Veritas Data Deduplication engines, select the engine for which you want to perform the catalog backup and click **OK**. Select the `dedupe_catalog` folder if you want catalog backup of both Veritas Data Deduplication engines in a single policy.
- 7 Run the policy manually or as per schedule.

Disabling Veritas Data Deduplication catalog backup from NetBackup client

You can disable the NetBackup client from performing a backup and restore of the deduplication catalog.

To disable the Veritas Data Deduplication catalog backup from the NetBackup client

- ◆ Disable Veritas Data Deduplication catalog backup from the NetBackup client.

```
Backup> dedupe-catalog disable config_name
```

`config_name` is an identifier and must be provided for any operation to be performed on the deduplication server if more than one deduplication server is running in the cluster.

Displaying the status of NetBackup services

Use the backup commands to display the status of the NetBackup services.

See the following man pages for detailed examples:

- `backup_show(1)`
- `backup_status(1)`
- `backup_start(1)`
- `backup_stop(1)`

To display NetBackup services

- ◆ Display the current NetBackup services.

```
Backup> show
```

If the settings were configured while the backup and the restore services were online, run the `backup stop` command followed by the `backup start` command to make the changes come into effect.

To display the status of backup services

- ◆ Display the status of backup services.

```
Backup> status
```

If the NetBackup server is started and online, then `Backup> status` displays any on-going backup or restore jobs.

To start or stop backup services

- 1 Start the backup services.

```
Backup> start [nodename]
```

You can also change the status of a virtual IP address to online after it has been configured using the `Backup> virtual-ip` command. This command applies to any currently active node in the cluster that handles backup and restore jobs.

- 2 Stop the backup services.

```
Backup> stop
```

You can also change the status of a virtual IP address to offline after it has been configured using the `Backup> virtual-ip` command.

Configuring backup operations using NetBackup or other third-party backup applications

You can backup Access Appliance using NetBackup client capability, or backup applications from other third-party companies that use the standard NFS mount to backup over the network.

For information on NetBackup, refer to the *Veritas NetBackup* product documentation set.

The Backup commands configure the local NetBackup installation of Access Appliance to use an external NetBackup primary server or media server. When NetBackup is installed on Access Appliance, it acts as a NetBackup client to perform IP-based backups of Access Appliance file systems.

Restoring filesystems backed up with NetBackup client

This section describes how you can restore filesystems (which are backed up using the NetBackup client) that get corrupted in the Access Appliance environment.

You can backup a filesystem with the NetBackup client using the `backup fs add` command. The backup of the filesystem is created on the NetBackup media server. If the filesystem goes corrupt and is destroyed, you can restore the filesystem.

To restore a corrupted filesystem

- 1 Create a new filesystem. You can use the same name as that of the existing filesystem if a filesystem with that name is no longer present, else use a different name.
- 2 Add the filesystem back to the NetBackup client's backup filesystem configuration using the `backup fs add fs_name` command.
- 3 Restore the data backed up from the filesystem to the respective path on the new filesystem.

To restore the MSDP catalog filesystem

- ◆ Restore the MSDP catalog filesystem.

If the NetBackup client is enabled with deduplication catalog backup and the catalog filesystem goes corrupt, you will need to format the catalog filesystem.

- Disable the catalog backup.

```
backup> dedupe-catalog disable
```

- Format the catalog filesystem or re-create it with the same name.
- Add the catalog filesystem as a backup filesystem.

```
backup> fs add fs_name
```

- From the deduplication catalog backup, restore the content that was backed up to `/vx/<catalog_fs>/dedupe/databases` by selecting **Restore to a different location** on the NetBackup primary.
- Enable the deduplication catalog to continue backing up the deduplication catalog.

```
backup> dedupe-catalog enable
```

Note: In case the MSDP catalog filesystem is destroyed, ensure that a new filesystem is created with same name as the old MSDP catalog filesystem. You can then follow the steps to add, restore and enable the catalog filesystem.

Configuring isolated recovery environment (IRE)

This chapter includes the following topics:

- [Requirements](#)
- [Configuring the network isolation](#)
- [Configuring an isolated recovery environment using the command line](#)
- [Configuring an isolated recovery environment on a storage server](#)
- [Managing an isolated recovery environment on a storage server](#)
- [Configuring data transmission between a production environment and an IRE storage server](#)

Requirements

The requirements to configure isolated recovery environment (IRE) in a Pull model are as follows:

- Access Appliance: 8.2 or later
- Storage server: 19.0.1 or later
- NetBackup: 10.2 or later

[Table 8-1](#) lists the supported configuration for MSDP source and targets for isolated recovery environment.


Table 8-1 Supported configuration for MSDP source and targets

Source	Target	Deployment models	NetBackup version	Replication - OptDup model
Veritas Data Deduplication 20.4.0.1	Veritas Data Deduplication 20.5.0.1	Access Appliance 8.6	10.5.0.1	Pull model
Veritas Data Deduplication 20.5.0.1	Veritas Data Deduplication 21.0.0.1		11.0.0.1	
Veritas Data Deduplication 21.0.0.1	Veritas Data Deduplication 21.1.0.2		11.1.0.2	
Veritas Data Deduplication 21.1.0.2	Veritas Data Deduplication 21.1.0.2			

Configuring the network isolation

To support the AIR for IRE, network communication from the IRE MSDP storage server to the production MSDP storage server is required. IRE MSDP storage server initiates the network connection. IRE with AIR works even when the production MSDP server does not have network access to the IRE MSDP server.

Figure 8-1 Network isolation example



Configure network isolation in the firewall

Configure the firewall at the IRE domain to deny all the inbound and the outbound connections. It helps to protect all the hosts in IRE domain from the cyber attacks. You must allow the IRE MSDP server outbound connection. For IRE replication, the IRE MSDP server must have network access to the production MSDP server

through the ports 10082 and 10102. The IRE MSDP server also must have network access to the production primary server using the port 1556.

If you cannot allow unidirectional network access (allow only outbound connection) on the firewall, you can allow bidirectional network for the IRE MSDP server. IRE air gap in the IRE MSDP server still denies all the inbound connections.

Configure IRE air gap on IRE MSDP server

Pre-requisites:

- Set up secure communication between Access Appliance and the primary server in its respective domain.

See [“Setting up secure communication between Veritas Data Deduplication on Access Appliance and the NetBackup primary server”](#) on page 91.

This establishes secure communication between the production Access Appliance Veritas Data Deduplication and the production primary. This also establishes secure communication between IRE Access Appliance Veritas Data Deduplication and IRE primary.

Note: Ensure that you adhere closely to the steps provided in this section to setup secure communication with primary server of its respective domain.

- Set up Auto Image Replication between two Access Appliances with secure communication enabled.

See [“Setting up Auto Image Replication \(AIR\) between two Access Appliances which have secure communication enabled”](#) on page 113.

Note: Cloud tiering is not supported on the target cluster when IRE is configured.

Air gap in IRE MSDP server does the following:

- Allows the network connections from servers in IRE domain. Connectivity between MSDP server and NetBackup primary or media servers is required to make the MSDP server functional.

Add the subnets or IP addresses of the IRE domain to the allowed subnet list. The IP addresses in the subnet list have direct network access to the MSDP server.

For example:

```
setting ire-network-control allow-subnets  
subnets=<subnet1>,<subnet2>,<ip address>
```

Note the following:

- The list must have at least the primary server, the media servers, and the DNS server in IRE domain.
- If you have a different subnet for your management network, ensure that this subnet is also included in the list, else you will lose access to appliance node and the management console.
- Do not add subnets or IP addresses from the domains outside the IRE domain.
- After adding the allowed subnets and schedule, you may have to wait for couple of minutes for the rules to take effect.
- Be cautious while adding the IPs and subnets to avoid cluster wide inaccessibility.
- All the existing connections that are not in allowed subnets are disconnected.
- Enables a unidirectional network access (allow outbound connection from IRE MSDP server to the other domains) in IRE air gap window. By default, the window is 24 hours per day.

All the inbound connections that are not in the allowed subnet list are denied.

Configuring an isolated recovery environment using the command line

See the following topics to configure an isolated recovery environment using the command line:

Table 8-2 IRE configuration using the command line

Task	Description
Configure an isolated recovery environment on a storage server.	See “Configuring an isolated recovery environment on a storage server” on page 162.
Manage an isolated recovery environment on a storage server.	See “Managing an isolated recovery environment on a storage server” on page 165.
Configure the data transmission between a production environment and an IRE storage server.	See “Configuring data transmission between a production environment and an IRE storage server” on page 168.

Configuring an isolated recovery environment on a storage server

You can configure an isolated recovery environment (IRE) on a storage server to create an air gap between your production environment and a copy of the protected data. The air gap restricts network access to the data except during the timeframe when data replication occurs. This feature helps to protect against ransomware and malware.

To configure the IRE, you need a production Access Appliance environment and a target storage server on a supported Veritas appliance. Check the appliance documentation for compatibility.

The production environment does not require any additional steps for this feature. Use the following procedure to configure an IRE on the target storage server from the MSDP restricted shell. You can login to the MSDP restricted shell by logging in to the target MSDP server IP (deduplication IP) with credentials of a cluster local user having administrator role.

Note: All the MSDP restricted shell commands should be executed on the IRE/target storage server.

To configure an IRE

- 1 If A.I.R. is not configured on the production domain, continue to the next step.

If A.I.R. is already configured on the production domain, log in as a local user with administrator role (same user that was used for Access Appliance CLISH login). Run the following command to show the SLP windows for replication from the primary server to the server.

```
setting ire-network-control show-slp-windows
production_primary_server=<production domain>
production_primary_server_username=<production username>
ire_primary_server=<IRE domain> ire_primary_server_username=<IRE
username>
```

Where:

- *<production domain>* is the fully qualified domain name (FQDN) of the primary server in your production environment.
- *<production username>* is the username of the NetBackup primary user with permission to list SLPs and SLP windows in the production environment. For Windows users, enter the username in the format **<domain name>\<username>**. For other users, enter the username only.

- *<IRE domain>* is the FQDN of the primary server in the IRE. Use the same hostname that you used for the target primary server when you configured the SLPs in the production environment.
- *<IRE username>* is the username of a IRE NetBackup primary user with permission to list SLPs and storage units in the IRE. For Windows users, enter the username in the format **<domain name>\<username>**. For other users, enter the username only.

For example:

```
production_primary_server=examplePrimary.domain.com
production_primary_server_username=example_user
ire_primary_server=exampleIREPrimary.domain.com
ire_primary_server_username=example_user1
```

The following is an example output of the command:

```
EveryDayAtNoon:
SLPs: SLP1
Sunday start: 12:00:00 duration: 00:59:59
Monday start: 12:00:00 duration: 00:59:59
Tuesday start: 12:00:00 duration: 00:59:59
Wednesday start: 12:00:00 duration: 00:59:59
Thursday start: 12:00:00 duration: 00:59:59
Friday start: 12:00:00 duration: 00:59:59
Saturday start: 12:00:00 duration: 00:59:59

WeeklyWindow:
SLPs: SLP2
Sunday start: 10:00:00 duration: 01:59:59
Monday NONE
Tuesday NONE
Wednesday NONE
Thursday NONE
Friday NONE
Saturday start: 10:00:00 duration: 01:59:59
```

This example shows two SLP windows:

- A daily window for one hour starting at noon.
- A weekly window for two hours starting at 10:00 A.M.

- 2 Based on the requirements for your environment, determine a schedule and take note of it. For an existing A.I.R. environment, the schedule must accommodate the SLP windows that you viewed in the previous step.

You can set a daily schedule that is open at the same time each day, or you can set a different schedule for each day of the week.

In the previous example, you can accommodate both SLP windows with either of the following:

- A daily schedule from 10:00 A.M. to 1:00 P.M.
- A schedule from 12:00 P.M. to 1:00 P.M. on Monday through Friday and a schedule from 10:00 A.M. to 1:00 P.M. on Saturday and Sunday

Note: If the production environment and the IRE are in different time zones, the schedule must begin only once per day in both time zones. For example, if one environment is in the Asia/Kolkata time zone and the other is in the America/New_York time zone, the following schedule in Kolkata is not supported: Tuesday start time 22:00:00 and Wednesday start time 03:00:00. When these start times get converted to the New York time zone, they become Tuesday start time 12:30:00 and Tuesday start time 17:30:00, which is not supported.

- 3 Run the following command to configure which subnets and IP addresses are allowed to access the storage server:

```
setting ire-network-control allow-subnets subnets=<CIDR subnets  
or IP addresses>
```

Where *<CIDR subnets or IP addresses>* is a comma-separated list of the allowed IP addresses and subnets, in CIDR notation.

For example:

```
setting ire-network-control allow-subnets  
subnets=10.80.120.208,10.84.48.0/20
```

Note: The IRE primary server, the IRE media servers, and the DNS server for the IRE must be included in the allowed list. If all of these servers are in the same subnet, only the subnet is required to be in the allowed list.

- 4 Run the following command to set the daily air gap schedule:

```
setting ire-network-control set-schedule start_time=<time>
duration=<duration> [weekday=<0-6>]
```

Where [weekday=<0-6>] is an optional parameter to indicate the day if you need to set different schedules for different days. 0 is Sunday, 1 is Monday, etc.

For example:

```
setting ire-network-control set-schedule start_time=10:00:00
duration=03:00:00 weekday=0
```

- 5 Before you can send data between the production domain and the IRE storage server, you must add MSDP reverse connections and add the replication operation.

See [“Configuring data transmission between a production environment and an IRE storage server”](#) on page 168.

Note: When multiple Veritas Data Deduplication instances are configured in an IRE domain cluster, and subnets and schedules are set using the MSDP restricted shell for all the Veritas Data Deduplication instances; if any of the instances is stopped and the cluster is restarted, the outbound connection rules corresponding to that instance is lost after the restart operation. This is by design and the outbound connection rules, if any, are applied back only if the stopped instance(s) are started back.

Managing an isolated recovery environment on a storage server

Once you have configured an isolated recovery environment (IRE) on a storage server, you can manage it from the MSDP restricted shell with a local user with administrator role (same user that was used for Access Appliance CLISH login). Use the following commands.

- To view the SLP windows from the primary server to the server:

```
setting ire-network-control show-slp-windows
production_primary_server=<production domain>
production_primary_server_username=<production username>
ire_primary_server=<IRE domain> ire_primary_server_username=<IRE
username>
```

Where:

- *<production domain>* is the fully qualified domain name (FQDN) of the primary server in your production environment.
- *<production username>* is the username of a NetBackup primary user with permission to list SLPs and SLP windows in the production environment. For Windows users, enter the username in the format **<domain name>\<username>**. For other users, enter the username only.
- *<IRE domain>* is the FQDN of the primary server in the IRE. Use the same hostname that you used for the target primary server when you configured the SLPs in the production environment.
- *<IRE username>* is the username of a IRE NetBackup primary user with permission to list SLPs and storage units in the IRE. For Windows users, enter the username in the format **<domain name>\<username>**. For other users, enter the username only.

For example:

```
production_primary_server=examplePrimary.domain.com
production_primary_server_username=example_user
ire_primary_server=exampleIREPrimary.domain.com
ire_primary_server_username=example_user1
```

Note: The SLP replication window on the production domain must be configured to be open at the same time as the IRE schedule.

- To list the MSDP reverse connections:

```
setting ire-network-control list-reverse-connections
```

- To add an MSDP reverse connection:

```
setting ire-network-control add-reverse-connection
remote_storage_server=<production MSDP server>
[remote_primary_server=<production primary server>]
[local_storage_server=<IRE network interface>]
```

Where:

- *<production MSDP server>* is the FQDN of the MSDP server in your production environment.
- *[remote_primary_server=<production primary server>]* is an optional parameter for the FQDN of the primary server in your production environment. This parameter is required if the IRE domain uses an alternative name to access the production primary server. This scenario usually occurs if the production primary server runs on multiple networks with multiple hostnames.

- [local_storage_server=<IRE network interface>] is an optional parameter for the hostname of the network interface to use for image replication on the IRE storage server. This parameter is required if the network interface for replication is different than the IRE storage server name.
- To verify that a reverse connection works:

```
setting ire-network-control validate-reverse-connection
remote_storage_server=<production MSDP server>
[remote_primary_server=<production primary server>]
[local_storage_server=<IRE network interface>]
```
- To remove an MSDP reverse connection:

```
setting ire-network-control remove-reverse-connection
remote_storage_server=<production MSDP server>
```
- To view the allowed IP addresses and subnets:

```
setting ire-network-control show-allows
```
- To add IP addresses and subnets to the allowed list:

```
setting ire-network-control allow-subnets subnets=<CIDR subnets
or IP addresses>
```

Where <CIDR subnets or IP addresses> is a comma-separated list of the allowed IP addresses and subnets, in CIDR notation.

For example:

```
setting ire-network-control allow-subnets
subnets=10.80.120.208,10.84.48.0/20
```

Note: The IRE primary server, the IRE media servers, and the DNS server for the IRE must be included in the allowed list. If all of these servers are in the same subnet, only the subnet is required to be in the allowed list.

- To remove the IP addresses and subnets from the allowed list:

```
setting ire-network-control allow-subnets subnets=,
```
- To view the daily air gap schedule:

```
setting ire-network-control show-schedule
```
- To change the air gap schedule:

```
setting ire-network-control set-schedule start_time=<time>
duration=<duration>
```

For example:

```
setting ire-network-control set-schedule start_time=10:00:00
duration=03:00:00
```

- To stop the air gap schedule:
`setting ire-network-control delete-schedule`
- To view the current network status and check whether the external network is open or closed:
`setting ire-network-control external-network-status`
- To manually open the external network:
`setting ire-network-control external-network-open`
- To manually close the external network and resume the air gap schedule:
`setting ire-network-control resume-schedule`

Note: The commands may take a few minutes to take effect.

Configuring data transmission between a production environment and an IRE storage server

Once the configuration of an isolated recovery environment (IRE) is completed, the production Access Appliance hosts are no longer able to access the storage server. You need to add MSDP reverse connections to allow data transmission between the production MSDP storage server and the IRE storage server. Then you can add the replication operation.

To configure data transmission between a production environment and an IRE

- 1 Open an SSH session to the IRE storage server. Run the following command to determine if the external network is open:

```
setting ire-network-control external-network-status
```

If it is not, run the following command:

```
setting ire-network-control external-network-open
```

- 2 Run the following command to add an MSDP reverse connection:

```
setting ire-network-control add-reverse-connection
remote_storage_server=<production MSDP server>
[remote_primary_server=<production primary server>]
[local_storage_server=<IRE network interface>]
```

Where:

- *<production MSDP server>* is the fully qualified domain name (FQDN) of the MSDP server in your production environment.

Configuring data transmission between a production environment and an IRE storage server

- [remote_primary_server=<production primary server>] is an optional parameter for the FQDN of the primary server in your production environment. This parameter is required if the IRE domain uses an alternative name to access the production primary server. This scenario usually occurs if the production primary server runs on multiple networks with multiple hostnames.
 - [local_storage_server=<IRE network interface>] is an optional parameter for the hostname of the network interface to use for image replication on the IRE storage server. This parameter is required if the network interface for replication is different than the IRE storage server name.
- 3 If necessary, repeat the previous step to add additional MSDP reverse connections.
 - 4 If Auto Image Replication (AIR) is not already configured on the production domain, run the following command to copy the IRE schedule to the production domain as a storage lifecycle policy (SLP) window:

```
setting ire-network-control sync-ire-window
production_primary_server=<production primary server>
production_primary_server_username=<production username>
[slp_window_name=<SLP window name>]
```

Where:

- <production primary server> is the FQDN of the primary server in your production environment.
 - <production username> is the username of the NetBackup primary user with permission to list SLPs and SLP windows in the production environment. For Windows users, enter the username in the format **<domain name>\<username>**. For other users, enter the username only.
 - [slp_window_name=<SLP window name>] is an optional parameter to give a name for the SLP window. If you do not provide this parameter, the name of the SLP window is IRE_DEFAULT_WINDOW.
- 5 If you do not have them already, create a source SLP on the production primary server and a target import SLP on the IRE primary server.

See [“Support for NetBackup Auto Image Replication”](#) on page 94.

Note: You cannot add the replication operation from the production NetBackup primary when you create the SLPs as the target storage is air gapped. Continue to the next step to add the replication operation.

- 6 Run the following command to add the IRE storage server as a replication target of the production Access Appliance domain and to add the replication operation to the SLP:

```
setting ire-network-control add-replication-op
production_primary_server=<production primary server>
production_primary_server_username=<production username>
production_storage_server=<production storage server>
ire_primary_server_username=<IRE username>
source_slp_name=<production SLP name> target_import_slp_name=<IRE
SLP name> target_storage_server=<target storage server>
target_storage_server_username=<target storage server username>
production_storage_unit=<MSDP storage unit> [slp_window_name=<slp
window name>]
```

Where:

- *<production primary server>* is the FQDN of the primary server in your production environment.
- *<production username>* is the username of the NetBackup primary user with permission to list SLPs and SLP windows in the production environment. For Windows users, enter the username in the format **<domain name>\<username>**. For other users, enter the username only.
- *<production storage server>* is the FQDN of the production storage server in your production environment.
- *<IRE username>* is the username for an administrator on the IRE NetBackup primary server. For Windows users, enter the username in the format **<domain name>\<username>**. For other users, enter the username only.
- *<source SLP name>* is the SLP name from the production primary server to add the replication operation to.
- *<target SLP name>* is the import SLP name from the IRE primary server.
- *<target storage server>* is the FQDN of the target storage server in your IRE environment.
- *<target storage server username>* is the username of the MSDP storage server used while configuring the target storage server .
- *<MSDP storage unit>* is the name of the MSDP storage unit that is the replication source in the source SLP.
- *[slp_window_name=<slp window name>]* is an optional parameter for the name of the SLP window that is synced with the IRE schedule. This

parameter must match the SLP window name from the previous step, if applicable. If you do not provide this parameter, the default name is used.

- 7 If you opened the external network at the beginning of this procedure, run the following command to close it and resume the air gap schedule:

```
setting ire-network-control resume-schedule
```

Note: The `setting ire-network-control allow-devices` is not supported in Access Appliance.

Troubleshooting

This chapter includes the following topics:

- [Additional resources](#)
- [Generating Access Appliance S3 server keys using the helper script](#)

Additional resources

See the following documentation for more information on Access Appliance and Veritas NetBackup:

- *Access Appliance Initial Configuration Guide* for the supported NetBackup clients ports on the [SORT](#) website.
- *Access Appliance Troubleshooting Guide* for setting the NetBackup client log levels and debugging options on the [SORT](#) website.
- *Veritas NetBackup* product documentation on the [SORT](#) website.

Generating Access Appliance S3 server keys using the helper script

Create the access and the secret keys using the Access Appliance helper script in case you do not want to use the Active Directory Domain user to create and own the buckets. This is an alternative way to get the Access Appliance S3 server credential keys.

- Location of the helper script:
`/opt/VRTSnas/scripts/utils/objectaccess/objectaccess_client.py`
- The Access Appliance helper script can be used from any client system that has Python installed.

- To run the script, your S3 client needs to have the `argparse` and `requests` Python modules.
If these modules are missing, install both these modules using `pip` or `easy_install`.
- Add the `ADMIN_URL` name in your `/etc/hosts` file. The `ADMIN_URL` has the format **admin.<cluster_name>:port number**. The port number should be 8144. This url should point to the Access Appliance management console IP address.
- Create the access and the secret key using the Access Appliance helper script by providing the user name, password, and `ADMIN_URL` (check the online Help of the Access Appliance helper script for all of the provided operations like `list key` and `delete key`).

Create a secret key:

```
clus_01:~ # ./objectaccess_client.py --create_key
--server admin.clus:8144 --username localuser1 --password root123
--insecure
UserName                : localuser1
AccessKeyId              : Y2FkODU2NTU2MjVhYzV
Status                   : Active
SecretAccessKey          : ODk0YzQxMDhkMmRjM2M5OTUzNjI5OWIzMdgyNzY
```

The `<localuser1>` is the local user created on both the Access Appliance cluster nodes with same unique ID.

List a secret key for the specified user:

```
clus_01:~ # ./objectaccess_client.py --list_key --server
admin.clus:8144 --username localuser2 --password root123 --insecure
```

Delete a secret key for the specified user:

```
clus_01:~ # ./objectaccess_client.py --delete_key
ZTKyNDdjZTViM2EyMWZ --server admin.clus:8144 --username localuser2
--password root123 --insecure
```

- If the object server is enabled without the `SSL` option, you need to add the `--insecure` option.

```
clus_01 ~# ./objectaccess_client.py --server
admin.clus:8144 --username <uname> --create_key --insecure
```

Index

A

- about
 - backup storage for NetBackup 8
 - Veritas Data Deduplication 14
- additional resources
 - documentation 172

B

- backup services
 - displaying the status of 155
 - starting 155
 - stopping 155
- backup storage for NetBackup 8

C

- configuring
 - backups using NetBackup or other third-party backup applications 156
 - Veritas Data Deduplication 23
- configuring NetBackup
 - prerequisites 145

D

- displaying
 - NetBackup configurations 155
 - status of backup services 155

G

- generating Access Appliance S3 server keys
 - using the helper script 172
- global deduplication
 - Veritas Data Deduplication 89

I

- integration with Access Appliance
 - NetBackup 7

M

- migrating
 - Veritas Data Deduplication 136

N

- NetBackup
 - configuring 156
 - displaying configurations 155
 - integration with Access Appliance 7

P

- prerequisites
 - configuring NetBackup 145

S

- starting
 - backup services 155
- stopping
 - backup services 155
- supported configurations and versions
 - NetBackup with MSDP-C 12
 - NetBackup with Veritas Data Deduplication 11
- system requirements
 - supported configurations and versions for
 - NetBackup with MSDP-C 12
 - supported configurations and versions for NetBackup with Veritas Data Deduplication 11

U

- use case
 - long-term data retention 8
- using the helper script
 - generating Access Appliance S3 server keys 172
- using Veritas Data Deduplication
 - benefits 16

V

- Veritas Data Deduplication
 - about 14

Veritas Data Deduplication *(continued)*

- configuring 23
- global deduplication 89
- migration 136