

Veritas NetBackup™ Flex Scale Administrator's Guide

3.0

Veritas NetBackup Flex Scale Administrator's Guide

Last updated: 2023-05-09

Legal Notice

Copyright © 2023 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Veritas Alta, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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

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

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

VERITAS TECHNOLOGIES LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

NB.docs@veritas.com

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

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Product overview	9
	About Veritas NetBackup™ Flex Scale	9
Chapter 2	Viewing information about the NetBackup Flex Scale cluster environment	10
	Accessing NetBackup Flex Scale and NetBackup	10
	Accessing the NetBackup web user interface on the appliance	12
	About the NetBackup Flex Scale management console	13
	About the Dashboard view	16
	Viewing all the activities	19
	Working with NetBackup Flex Scale APIs	19
Chapter 3	NetBackup Flex Scale infrastructure management	21
	User management	21
	Considerations for managing NetBackup Flex Scale users	23
	Adding users	25
	Changing user password	28
	Removing users	29
	Considerations for configuring AD/LDAP	31
	Configuring AD server for Universal shares and Instant Access	32
	Configuring AD/LDAP servers for NetBackup services	34
	Configuring additional AD/LDAP servers for managing NetBackup services/Universal Shares/Instant Access	38
	Configuring AD/LDAP servers on clusters deployed with only media servers	38
	Directory services and certificate management	41
	Region settings management	45
	About NetBackup Flex Scale storage	47
	About Universal Shares	49
	Creating a Protection Point for a Universal Share	51
	Node and disk management	53
	NetBackup Flex Scale network cabling	54

Adding a node to the cluster using the NetBackup Flex Scale web interface	55
Adding a node using the REST APIs	60
Replacing a node in a cluster	63
Starting and stopping nodes	66
Replacing a disk	68
Viewing disk details	70
Viewing node details	71
License management	72
Adding or removing storage licenses	72
Adding NetBackup licenses	75

Chapter 4 NetBackup Flex Scale network management 77

About network management	77
Modifying DNS settings	79
About bonding Ethernet interfaces	80
Bonding operations	81
Creating a bond	81
Modifying a bond	86
Removing a bond	87
Configuring NetBackup Flex Scale in a non-DNS environment	89
Data network configurations	93
Network configuration on plain device (eth5)	93
Network configuration on VLAN (eth5)	98
Network configuration on bonded interfaces (bond0 on eth5 and eth7)	99
VLAN on bond of eth5 and eth7 (bond0)	101
Support for multiple VLAN when disaster recovery is configured	101

Chapter 5 NetBackup Flex Scale infrastructure monitoring 102

About alert management	102
Viewing information about alerts	103
Managing alerts	103
About event notification	104
Purging events	104
About AutoSupport and Call Home	105
Setting up email alerts	105
Setting up SNMP alerts	108
Configuring Call Home settings	109
Monitoring hardware components	111

	Monitoring usage and licensed capacity using Veritas NetInsights Console	113
Chapter 6	Resiliency in NetBackup Flex Scale	114
	Erasure coding in NetBackup Flex Scale	114
	Handling split-brain scenario in NetBackup Flex Scale	115
	High availability of the NetBackup primary service	116
	High availability of NetBackup services	118
	NetBackup catalog protection	118
	NetBackup primary service catalog protection using checkpoints	118
	Performing a recovery of the catalog file system using GUI	119
	Performing a recovery of the catalog file system using REST APIs	122
Chapter 7	EMS server configuration	125
	Configuring an external BYOS media server	125
	Configuring an external NBA media server	126
Chapter 8	Site-based disaster recovery in NetBackup Flex Scale	128
	About site-based disaster recovery in NetBackup Flex Scale	128
	Configuring disaster recovery using GUI	130
	Clearing the host cache	134
	Managing disaster recovery using GUI	134
	Performing disaster recovery using RESTful APIs	138
	Establishing trust and setting up authentication	139
	Configuring disaster recovery	140
	Managing disaster recovery	141
	Active-Active disaster recovery configuration	143
	NetBackup optimized duplication using Storage Lifecycle Policies	145
Chapter 9	NetBackup Flex Scale security	146
	STIG overview for NetBackup Flex Scale	146
	STIG-compliant password policy rules	147
	Enabling STIG for NetBackup Flex Scale	148
	Viewing the NetBackup Flex Scale STIG status	151
	FIPS overview for NetBackup Flex Scale	152
	Viewing the NetBackup Flex Scale FIPS status	153
	Managing the login banner	154

Changing the password policy	155
Support for immutability in NetBackup Flex Scale	158
About lockdown modes	158
Selecting or changing the lockdown mode	159
Configuring immutability using GUI	161
Deploying external certificates on NetBackup Flex Scale	163
Deploying ECA using the GUI	165
Log locations	168
Considerations for performing other operations when ECA is deployed	169
Chapter 10 Troubleshooting	170
Services management	170
Collecting logs for cluster nodes	171
Uploading logs to Veritas Support	174
Downloading logs	175
Checking and repairing storage	176
Troubleshooting NetBackup Flex Scale issues	177
If cluster configuration fails (for example because an IP address that was already in use is specified) and you try to reconfigure the cluster, the UI displays an error but the configuration process continues to run	177
Validation error while adding VMware credentials to NetBackup	178
NetBackup Web UI incorrectly displays some NetBackup Flex Scale processes as failed	178
Unable to create BMR Shared Resource Tree (SRT) on NetBackup Flex Scale Appliance	179
NetBackup configuration files are not persistent across operations that require restarting the system	180
Appendix A Configuring NetBackup optimized duplication 	182
Configuring a Storage Lifecycle Policy for optimized duplication	182
Creating a Storage Lifecycle Policy for optimized duplication	182
Configuring a policy to use an SLP	189
Updating the policy to reverse the replication direction	190
Appendix B Disaster recovery terminologies	192
VVR technology in disaster recovery	192
About response fields in the GET disaster recovery API	193

Appendix C	Configuring Auto Image Replication	199
	Auto Image Replication configuration	199

Product overview

This chapter includes the following topics:

- [About Veritas NetBackup™ Flex Scale](#)

About Veritas NetBackup™ Flex Scale

Veritas NetBackup™ Flex Scale is the next generation, hyper-converged, scale-smart data protection solution that is based on Veritas NetBackup, the industry-leading backup and recovery software.

Key benefits include:

- Simplified NetBackup deployment experience
- Containerized microservices-based distributed architecture
- Automated node discovery and storage provisioning
- Web-based interface for configuration and infrastructure management
- Built-in high availability and resiliency
- Non-disruptive and dynamic node (compute) and storage expansion based on growing data protection needs

Viewing information about the NetBackup Flex Scale cluster environment

This chapter includes the following topics:

- [Accessing NetBackup Flex Scale and NetBackup](#)
- [Accessing the NetBackup web user interface on the appliance](#)
- [About the NetBackup Flex Scale management console](#)
- [About the Dashboard view](#)
- [Working with NetBackup Flex Scale APIs](#)

Accessing NetBackup Flex Scale and NetBackup

The following table describes how to access and manage the NetBackup Flex Scale cluster and NetBackup when both NetBackup primary and media servers are deployed in the cluster:

Table 2-1

Interface and URL	To access
NetBackup Flex Scale web UI https://ManagementServerIPorFQDN/webui	You can use the NetBackup Flex Scale web UI to manage both NetBackup and NetBackup Flex Scale infrastructure. Using the Appliance administrator and NetBackup administrator role: ■ View and manage NetBackup. ■ To view the NetBackup Flex Scale infrastructure such as the nodes, disks, and hardware details, in the left navigation pane click Cluster Monitor > Infrastructure. From the Infrastructure page, you can also open the NetBackup Flex Scale infrastructure management console by clicking Cluster dashboard, which is in the upper-right corner of the UI. To view the cluster details in the infrastructure management console, click Monitor > Infrastructure. Using the Appliance administrator role: ■ View and manage NetBackup Flex Scale infrastructure. To view the NetBackup Flex Scale infrastructure click Cluster Monitor > Infrastructure. From the Infrastructure page, you can also open the NetBackup Flex Scale infrastructure management console by clicking Cluster dashboard, which is in the upper-right corner of the UI. To view the cluster details in the infrastructure management console, click Monitor > Infrastructure. ■ You cannot view or manage NetBackup using this role. Using the NetBackup administrator role: ■ View and manage NetBackup. ■ You cannot access or manage NetBackup Flex Scale infrastructure using this role.
NetBackup Flex Scale infrastructure management console IPv4: https://ManagementServerIPorFQDN:14161/ IPv6: https://[ManagementServerIP]:14161/ IPv6:https://ManagementServerFQDN:14161/	You can view and manage NetBackup Flex Scale. To access NetBackup, in the left navigation pane click NetBackup. This action launches the NetBackup UI in the same browser tab. See “About the NetBackup Flex Scale management console” on page 13.

Table 2-1 (continued)

Interface and URL	To access
NetBackup UI <code>https://primaryserverfqdn/webui/</code>	View and manage NetBackup. You can open NetBackup Flex Scale from the NetBackup UI when you click Appliance management in the left navigation pane.

Accessing the NetBackup web user interface on the appliance

The NetBackup web user interface (UI) allows you to configure and manage your NetBackup environment. You can access the NetBackup web UI directly and also from within the NetBackup Flex Scale infrastructure management console UI.

To access the NetBackup web UI from the NetBackup Flex Scale web UI:

Use a user account with only the NetBackup administrator role or a user account with both Appliance administrator and NetBackup administrator role to log in to the NetBackup Flex Scale web UI:

`https://ManagementServerIPorFQDN/webui`

where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.

To access the NetBackup web UI from the NetBackup Flex Scale infrastructure management console UI

- 1 Sign in to the NetBackup Flex Scale infrastructure management console.
See [“About the NetBackup Flex Scale management console”](#) on page 13.
- 2 From the navigation pane on the left, click **NetBackup**.
This launches the NetBackup web UI in the same browser pane and you are automatically signed out from the infrastructure console.
- 3 On the NetBackup web UI sign in page, specify the NetBackup administrator user account and password and then click **Sign in**.


For more information about the NetBackup web UI, refer to the *NetBackup Web UI Administrator's Guide*.

To access the NetBackup UI directly

- ◆ To access the NetBackup web UI, use the following URL:

`https://primaryserverfqdn/webui`

where *primaryserverfqdn* is the FQDN of the NetBackup primary server.



About the NetBackup Flex Scale management console

You can administer your NetBackup Flex Scale appliance cluster configuration using the NetBackup Flex Scale infrastructure management console. This console provides a centralized consolidated graphical view of your entire appliance configuration. You can use the infrastructure management console to view all the cluster details such as the cluster nodes, the storage that is available on those nodes, and the services running on the appliance.

To access the infrastructure management console

- 1 Open a web browser and type the following URL in the address bar:

`https://ManagementServerIPorFQDN:14161/`


If you are using IPv6 addresses, use the following URL syntax:

`https://[ManagementServerIPorFQDN]:14161/`


Here, *ManagementServerIP* is the public IP address or FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.

Note: Your browser may display a security certificate validation message. You may have to resolve the error by adding a security exception for the specified URL.

- 2 On the sign in page, specify the administrative user name with the Appliance administrator role and the password that you created during the cluster configuration and then click **Sign In**.



- 3 The **Dashboard** tab displays the various elements of your cluster configuration.



You can use the infrastructure UI for performing administrative tasks such as:


- Monitoring and managing the appliance services and storage utilization
- Monitoring and managing the infrastructure components such as nodes, storage disks, and network
- Monitoring the appliance performance and data reduction levels

About the Dashboard view


The NetBackup Flex Scale Dashboard provides an overview of the appliance cluster configuration and its components and services. Use the infrastructure console Dashboard view to monitor the overall health of your cluster.

You can log in to the Dashboard only if you have been assigned the appliance administrator role. If a user does not have the Appliance administrator role, then the user is directed to the **Change password** screen in the NetBackup Flex Scale GUI. Other views are not visible to such a user.

The following figure shows the dashboard for a cluster deployed with NetBackup primary and media servers:



The following figure shows the dashboard for a cluster deployed with only media servers:



The Dashboard view displays the following information:

- The Alerts area displays information about the most recent operations including issues that you may need to take action on. To view a more detailed list of all the alerts, click **View details** in the Alerts box.
- The Storage area displays an aggregated view of the storage utilization on the appliance. If the cluster is deployed with NetBackup primary server and media

servers, the Storage area provides information about the total storage available to all the cluster nodes and the disk space that is consumed by the workload backups and the NetBackup catalog. If the cluster is deployed with only media servers, the Storage area provides information about the total storage available to all the cluster nodes and the disk space that is consumed by the workload backups. To view a more detailed list of all the disk space utilization, click **View details** in the Storage box.

Note: You may see a difference between the total available storage capacity and the actual usable storage space. This happens because NetBackup Flex Scale uses some of that space to store erasure coded data to protect against disk and node failures.

See [“Erasure coding in NetBackup Flex Scale”](#) on page 114.

- The Services area displays the status of the NetBackup services that are running on the appliance. If the cluster is deployed with NetBackup primary server and media servers, the Services area includes services such as NetBackup master server, media servers, storage servers, and replication services. If only media servers are deployed, the Services area includes services such as the media servers and storage servers.
- The Security area displays the status of the security features that can be configured for the cluster. This includes the status of the lockdown mode and WORM storage support, the STIG mode, the FIPS mode, and the sign-in banner.
- The Infrastructure area displays an aggregated view of the infrastructure components on the appliance. This includes the number of cluster nodes and their health, the total number of storage disks and their health, and the status of the public and private network that is configured on the appliance.
- The Data reduction area displays information about the data space optimization that is achieved by the NetBackup deduplication engine.
- The Storage licensing area displays the license type and the expiration date of all the added licenses.
- The Performance area displays the details about the IOPS, throughput that is currently being achieved in the cluster, and latency for read and write operations.
- The Resource monitoring area displays the aggregate data about the CPU usage, the memory usage, and the network usage in the cluster.

Viewing all the activities

The **View all activities** option is available in the top right corner of the NetBackup Flex Scale UI. You can get information about the most recent operations including issues that you may need to take action on. To view a more detailed list of all the activities, click **View details** option.

Working with NetBackup Flex Scale APIs

NetBackup Flex Scale employs a dedicated nginx-based web server that runs the API gateway and the management server. All UI access and API requests are sent to this web server. A single IP address and FQDN is used for accessing the NetBackup Flex Scale infrastructure UI as well as the NetBackup Web UI.

- You can access the NetBackup Flex Scale infrastructure management console UI using the following URL:

```
https://ManagementServerIPorFQDN:14161
```

If you are using IPv6 addresses, use the following URL syntax:

```
https://[ManagementServerIP]:14161
```

- You can access the NetBackup Web UI using the following URL:

```
https://ManagementServerIPorFQDN/webui
```

If you are using IPv6 addresses, use the following URL syntax:

```
https://[ManagementServerIP]/webui
```

Here, *ManagementServerIPorFQDN* is the public IP address or FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.

NetBackup Flex Scale provides APIs that are built using the standard Representational State Transfer (REST) architecture. You can use the RESTful APIs to control all aspects of your NetBackup Flex Scale configuration. There are APIs for infrastructure monitoring and management, user management, and for performing several other operations in your cluster.

Accessing the Swagger-based REST APIs

You can access the NetBackup Flex Scale APIs using the Swagger interface.

To access the APIs from a web browser

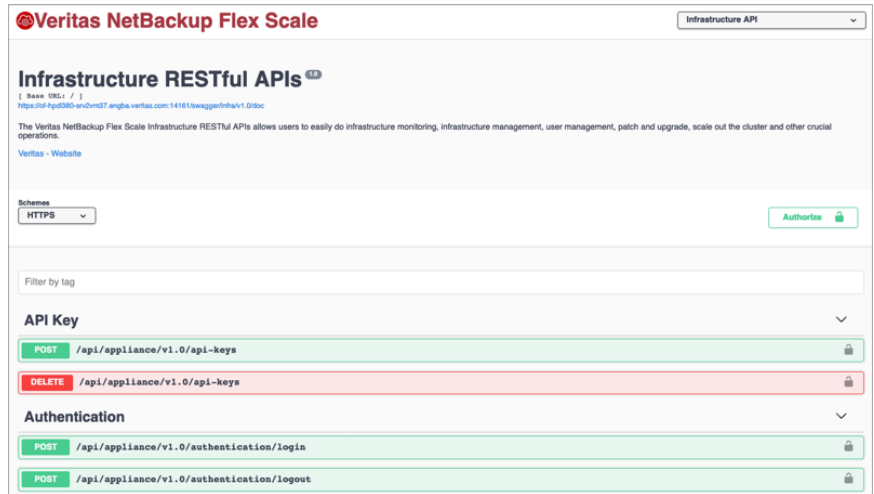
- Open your browser and enter the following URL in the address bar:

```
https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0
```

If you are using IPv6 addresses, use the following URL syntax:

```
https://[ManagementServerIP]:14161/swagger/infra/v1.0
```

Here, *ManagementServerIPorFQDN* is the public IP address or FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration.



- For a quick introduction on how to use the APIs, refer to the getting started readme available at the following URL:

`https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0/readme`

If you are using IPv6 addresses, use the following URL syntax:

`https://[ManagementServerIP]:14161/swagger/infra/v1.0/readme`

For more information on APIs, refer to the Veritas NetBackup Flex Scale API documentation on SORT.

- You can access the sample SDKs here:
<https://github.com/VeritasOS/NetBackup-Flex-Scale-REST-API-nuggets>

NetBackup Flex Scale infrastructure management

This chapter includes the following topics:

- [User management](#)
- [Region settings management](#)
- [About NetBackup Flex Scale storage](#)
- [About Universal Shares](#)
- [Node and disk management](#)
- [License management](#)

User management

NetBackup Flex Scale enables you to add users for administering your appliance. You can add local users as well as users from an Active Directory (AD) server and an Lightweight Directory Access Protocol (LDAP) server. Registering remote users lets you leverage your existing directory service for user management and authentication. Each user account must authenticate itself with a user name and password to access the appliance. For a local user, the user name and password are managed on the appliance. For a registered remote user, the user name and password are managed by the remote directory service.

NetBackup Flex Scale provides the following types of users:

Table 3-1 NetBackup Flex Scale users

User role / user	Description
Appliance Administrator <i>(user role)</i>	The Appliance Administrator user role is the primary admin user for the NetBackup Flex Scale appliance and has the privileges to manage and monitor all the infrastructure components in the cluster. This user role can perform tasks such as managing the cluster nodes, the cluster settings, storage and networking. This user role can also monitor the alerts and notifications related to the infrastructure components and the cluster operations.
NetBackup Administrator <i>(user role)</i>	The NetBackup Administrator user role has the permissions to manage and monitor all the NetBackup services and backup operations in the cluster. This user role can perform tasks such as NetBackup protection policy management, scheduling of backup and restore jobs, and managing NetBackup operations.

Table 3-1 NetBackup Flex Scale users (continued)

User role / user	Description
<code>maintenance</code> <i>(user)</i>	In addition to the user roles specified earlier, NetBackup Flex Scale also provides a built-in user account that has administrative access to the operating system root on the nodes. This user is automatically created during the cluster configuration. To use this account, you must first log on to the node using the infrastructure admin role and then access the OS root by typing the maintenance user password. The maintenance account is used by Veritas Support through the appliance Shell Menu (after an administrative log-on). This account is used specifically to perform maintenance activity or to troubleshoot the appliance. Unlike the other types of users, this user is a single user account and is not a user role. The maintenance user account details are as follows: ■ User name: <code>maintenance</code> ■ Password: <code>P@ssw0rd</code> This is the default password and you must change the password immediately after the cluster is configured. Sign in to the NetBackup Flex Scale infrastructure UI console using the default user credentials and set a new password.

See [“Considerations for managing NetBackup Flex Scale users”](#) on page 23.

See [“Adding users”](#) on page 25.

See [“Removing users”](#) on page 29.

For details about security features, see the *Veritas NetBackup™ Security and Encryption Guide*.

Considerations for managing NetBackup Flex Scale users

Consider the following while managing users in your NetBackup Flex Scale cluster:

- You can add up to 10 user accounts to the NetBackup Flex Scale cluster. You must add at least one user for each user role during the cluster configuration. You can add additional users at any time after the cluster is configured.

You can also add users without assigning a role. You can assign the required role to such user accounts later.

Note: If you have deployed the cluster with only media servers, only the appliance administrator role option is available during cluster configuration and post cluster configuration.

- You can use a single user account and assign both NetBackup Flex Scale and NetBackup administrator roles to the same account. However, two separate user accounts are recommended.
- You can assign the NetBackup admin role to a user account only during the cluster configuration. After the cluster is configured, you must use the NetBackup Web UI to manage NetBackup admin role assignment to user accounts. You cannot use the NetBackup Flex Scale infrastructure UI console to manage NetBackup roles post cluster configuration.
- Veritas recommends that you use the NetBackup Web UI to manage assignment of NetBackup roles. Use the NetBackup Flex Scale infrastructure UI console to manage assignment of NetBackup Flex Scale roles.

Note: NetBackup roles assigned from the NetBackup Web UI are not visible in the NetBackup Flex Scale infrastructure UI console.

- You can add local as well as AD and LDAP user accounts to the user roles. NetBackup Flex Scale supports AD and LDAP in Secure Sockets Layer (SSL) and Non-SSL mode.

Note: For AD/LDAP user account access to remain active, the NetBackup primary server service must be running and healthy in the cluster.

- If the cluster is configured for disaster recovery (DR), then ensure that you manually add the users to the secondary cluster in the same order as they are in the primary cluster.
- If both primary server and media servers are deployed on the cluster, AD and LDAP users cannot access the appliance SSH, CLISH and REST APIs. If only media servers, are deployed, AD/LDAP users having appliance administrator role can access SSH, GUI, CLISH and REST APIs.
- Local users that have the NetBackup Flex Scale appliance administrator role assigned have access to CLISH and the infrastructure REST APIs. Users that

have the NetBackup administrator role assigned have access to the NetBackup REST APIs.

- While removing user accounts from the cluster, you cannot delete all the users from the users list. At least one user with either the NetBackup Administrator or the NetBackup Flex Scale Appliance Administrator role always remains in the users list.
- While assigning roles from the NetBackup UI, you must use the IP or the FQDN of the server that was used during the configuration instead of the "NBU_LDAP_DOMAIN" string.
- Veritas recommends that LDAP and AD user UIDs start from 10000. Otherwise, when you assign a role to the AD/LDAP user, the UIDs of some of the local user may conflict with the UID of a user from the directory server
- Nested LDAP group for role assignment is not supported.
- You can configure multiple AD/LDAP servers.

Consider the following while managing users in your NetBackup Flex Scale cluster after upgrade:

- You must use the "LDAP_Server_FQDN/IP" string for the AD/LDAP server that was configured before upgrade.
- You should use the server name to assign role to AD/LDAP server that is configured after upgrade
- You can use the NetBackup REST API to get the ID value and use the ID value while assigning a role to configured AD/LDAP server users.

See [“User management”](#) on page 21.

See [“Adding users”](#) on page 25.

See [“Removing users”](#) on page 29.

Adding users


Perform the following steps to add users to the NetBackup Flex Scale cluster. Before you proceed, ensure that you have reviewed the considerations for user management in a NetBackup Flex Scale cluster.

See [“Considerations for managing NetBackup Flex Scale users”](#) on page 23.

To add a user to the NetBackup Flex Scale cluster

- 1
- Sign in to the NetBackup Flex Scale infrastructure management console UI.
See [“About the NetBackup Flex Scale management console”](#) on page 13.
- 2
- From the navigation menu on the left, click **Settings** and then click **User management**.

The User management pane displays all the users that currently exist in the cluster configuration.
- 3
- To add a new user, click **Add**.
- 4
- In the **Add new user** dialog box, specify the following parameters:




Parameter	Description
User name	Enter the username. The following criteria apply: ■ The username can contain a maximum of 31 characters. ■ The username must not start with a number. ■ The username can include only underscore (_) and dash (-) special characters. ■ The username must not be "admin", "root", or "primary".

Parameter	Description
Password	Enter a password for the user account. The following criteria apply: ■ The password should contain a minimum of 8 characters. ■ The password should include an uppercase, a lowercase, a number and a special character. ■ The password should include one of the following special characters <code>!#\$%&+,-.<=>@[]^_{}~</code> ■ No white spaces are allowed. ■ Dictionary words are invalid and not accepted. If the STIG option is enabled, NetBackup Flex Scale automatically enforces a higher security password policy. See “STIG-compliant password policy rules” on page 147.
Confirm password	Enter the password again to confirm.
Assign 'Appliance administrator' role	Select this option to assign the NetBackup Flex Scale Appliance Administrator user role to the user account. You must assign a role for the user to be able to sign in to the cluster using the UI. This user role has the permissions to manage all the infrastructure components in the NetBackup Flex Scale cluster.

5 Click **Add**.

The UI displays a message that confirms that the user addition operation is triggered. You can click **View Details** to see the progress of the operation.



Alternatively, you can also click the **Recent Activity** icon from the top right of the UI to see the status of the most recent operations occurring in the cluster.


6 Once the operation completes verify that the new user you added appears in the list of users displayed in the **User management** pane.

Changing user password

NetBackup Flex Scale lets you change the user password.

Considerations when you want to change the user password:

- You can only change your password. You are not permitted to change any other user's password.
- A local user can change the password from the NetBackup Flex Scale UI.
- The passwords of AD/LDAP users are managed by respective AD/LDAP servers and cannot be changed from NetBackup Flex Scale UI.
- The password of users with *Appliance administrator* role can be changed using the user information drop-down at the top-right corner of the NetBackup Flex Scale UI.



- The password of users who have no role or have only *NetBackup administrator* role can be changed by logging on to the NetBackup Flex Scale UI. As the user does not have the *Appliance administrator* role, the user is directed to the **Change password** tab of the NetBackup Flex Scale UI where the password can be changed.
- Passwords can also be changed using respective REST APIs.

Removing users

Perform the following steps to remove users from the NetBackup Flex Scale cluster. Before you proceed, ensure that you have reviewed the considerations for user management in a NetBackup Flex Scale cluster.


See [“Considerations for managing NetBackup Flex Scale users”](#) on page 23.

To remove a user from the NetBackup Flex Scale cluster


- 1 Sign in to the NetBackup Flex Scale infrastructure management console UI.
See [“About the NetBackup Flex Scale management console”](#) on page 13.
- 2 From the navigation menu on the left, click **Settings** and then click **User management**.

The User management pane displays all the users that currently exist in the cluster configuration.

- 3 To remove a particular user account, in the user row, click on the vertical ellipsis button from the right side of the UI and then select **Delete**.



- 4 On the **Delete user** dialog box, click **Delete** to confirm the user removal.
- The UI displays a message that confirms that the user removal operation is triggered. You can click **View Details** to see the progress of the operation.



- Alternatively, you can also click the **Recent Activity** icon from the top right of the UI to see the status of the most recent operations occurring in the cluster.
- 5 Once the operation completes, verify that the user you removed does not appear in the list of users displayed in the **User management** pane.

Considerations for configuring AD/LDAP

Consider the following while configuring AD/LDAP in your NetBackup Flex Scale cluster:

- In a deployment with both primary and media servers:
 - Domain users should use the [<AD/LDAP_serverIP/FQDN>\username] to login to NetBackup UI.
 - Domain user login to the management IP of the node and console IP over SSH is not supported.
 - Domain user login to the public Appliance/Infrastructure APIs is not supported.
 - AD/LDAP domains can be added using FQDN as well as IP addresses.
 - AD/LDAP servers should be reachable through the data network.
- In a deployment with only media servers:
 - Domain users should use the [<AD/LDAP_serverIP/FQDN>\username] to login to Appliance UI.
 - AD/LDAP users should use only the [username@<AD/LDAP_serverIP/FQDN>] to login to public Appliance/Infrastructure APIs.
 - AD/LDAP domains can be added using FQDN as well as IP addresses.
 - AD/LDAP users should use only the AD/LDAP username to log on to the management IP of the node and console IP using SSH.
 - AD/LDAP servers should be reachable through the management network.
- In a non-DNS environment:
 - If you want to add a domain using FQDN, the IP to FQDN mapping for that domain should be added using the **Network > Custom Hosts** option in the Appliance UI. If multiple VLANs are configured on the cluster, AD/LDAP servers should be configured with the correct VLAN. Only then the AD/LDAP servers can communicate with the multi-VLAN cluster.
 - When configuring AD/LDAP servers for NetBackup Services on a deployment with both NetBackup primary and media servers, the custom hosts entries must be applied to **NetBackup services**.
 - When configuring AD/LDAP servers for NetBackup Services on a deployment with only media servers, the custom hosts entries must be applied to **Cluster Nodes**.
 - When configuring AD/LDAP servers for Universal share/Instant Access, the custom hosts entries must be applied to **NetBackup services**.


Configuring AD server for Universal shares and Instant Access

You can configure an AD server for Universal Shares and Instant Access from the NetBackup Flex Scale GUI. You can configure only one AD server with the type as Universal shares/Instant Access.

The AD server should already be configured before you use the Instant Access MSSQL.

To configure AD server for Universal Shares/Instant Access

- 1 Go to the **Appliance GUI > Settings > Directory service management** tab.
- 2 Click **Configure**.



- 3 Select **Manage universal shares/Instant Access**. Provide the AD server details. Click **Save**.

Configure directory service

Directory service usage

Universal Shares/Instant access is only supported for one Active Directory server per cluster.

☐

 Manage NetBackup services

☒

 Manage Universal Shares/Instant access

Active Directory server domain name *

Active Directory server IP address/domain name *

Active Directory administrator *

Active Directory password *

Cancel

Save

- The administrator of the AD server should have Domain Admin privilege.
 - If the AD domain is sub domain, make sure that the domains can be trusted, or configure the AD domains with the type as **Manage NetBackup services** in NetBackup Flex Scale.
- 4 Once the AD server configuration operation is successful, the AD server appears in the list with usage as Universal Share/Instant Access.

You can perform a test connection to validate the AD configuration. On the same page, there is an option to remove the AD configuration for Universal Shares.

Veritas NetBackup™ Flex Scale - appliance-clust

Settings / Directory Services

Configure Active Directory(AD) and Lightweight Directory Access Protocol(LDAP) server connections for management and universal shares/instant access.

Search...

Add

Directory server	Type	Usage
10.221.52.88	Active Directory	Universal Shares/Instant Access

Items per page: 101 - 1 of 1

You can also use RESTful APIs to configure AD servers for Universal Shares and Instant Access.

To configure AD servers for Universal Shares and Instant Access using RESTful APIs

- 1** Configure an AD server for Universal Share with AD server details as payload.

```
POST /api/appliance/v1.0/settings/winbind
```

- 2** Test the configured server's connections on the cluster with AD server domain name and password as payload.

```
POST /api/appliance/v1.0/settings/winbind{server}/test
```

- 3** Get the successfully configured AD server domain names on the cluster.

```
GET /api/appliance/v1.0/settings/winbind
```

- 4** Get the server details such as, server name, BIND DN, and FQDN with AD server domain name as payload.

```
GET /api/appliance/v1.0/settings/winbind/{server}
```

- 5** Delete, remove or unconfigure the configured AD servers from the cluster with AD server domain name as payload.


```
Delete /api/appliance/v1.0/settings/winbind{server}
```

Configuring AD/LDAP servers for NetBackup services

You can configure AD/LDAP servers for using the NetBackup services from the NetBackup Flex Scale GUI. You can configure only one AD/LDAP server at a time but you can add multiple AD/LDAP servers to the NetBackup Flex Scale cluster. .

To configure AD/LDAP servers for NetBackup services using GUI

- 1 Go to the **Appliance GUI** > **Settings** > **Directory service management** tab. Click **Configure**.



2 Select **Manage NetBackup services**.

Provide the AD/LDAP server details. Click **Save**.

Configure directory service

Directory service usage

Universal Shares/Instant access is only supported for one Active Directory server per cluster.

☒ Manage NetBackup services
 ☐ Manage Universal Shares/Instant access

LDAP server address *

Port *

Directory Type *

Open LDAP

The most commonly used LDAP ports are 389 and 636(using ssl encryption).

User Base DN *

Group base DN *

Bind DN/LDAP administrator *


Bind DN/LDAP password *

Cancel

Save

3 Once the AD/LDAP server configuration operation is successful, the AD/LDAP server(s) appears in the list with usage as NetBackup services.

You can perform a test connection to validate the AD/LDAP configuration. On the same page, there is an option to remove the AD/LDAP configuration.



You can also use RESTful APIs to configure AD/LDAP servers for NetBackup services.

To configure AD/LDAP servers for NetBackup services using RESTful APIs

- 1 Configure an AD/LDAP server to the cluster with AD/LDAP server details as payload.

```
POST /api/appliance/v1.0/settings/ldap
```

- 2 Test the configured server's connections on the cluster with AD/LDAP server domain name and password as payload.

```
POST /api/appliance/v1.0/settings/ldap{server}/test
```

- 3 Get the successfully configured AD/LDAP server domain names on the cluster.

```
GET /api/appliance/v1.0/settings/ldap
```

- 4 Get the server information such as whether, the server is AD/LDAP, BIND DN, and port with AD/LDAP server domain name as payload.

```
GET /api/appliance/v1.0/settings/ldap/{server}
```

- 5 Delete, remove or unconfigure the configured AD/LDAP servers from the cluster with AD/LDAP server domain name as payload.

```
Delete /api/appliance/v1.0/settings/ldap{server}
```

Configuring additional AD/LDAP servers for managing NetBackup services/Universal Shares/Instant Access

You can configure only one AD server with the type as Universal Shares/Instant Access. But if you want to configure multiple AD/LDAP servers, use the following procedure.

To configure additional AD/LDAP servers

- 1 Go to the **Appliance GUI > Settings > Directory service management** tab.
- 2 Click **Add**. Select one of the following options as per your requirement:
 - **Manage universal shares/Instant Access**
 - **Manage NetBackup services**
- 3 Provide the server details. Click **Save**.
- 4 Once the server configuration operation is successful, the server appears in the list with the specified usage.

Configuring AD/LDAP servers on clusters deployed with only media servers

NetBackup Flex Scale supports multiple AD/LDAP server configuration on clusters deployed with only media servers. You can configure AD/LDAP servers using the Appliance GUI. In a NetBackup Flex Scale deployment with only media servers, the AD/ LDAP servers need to be reachable through the management network.

There is an option in the **User management** tab of the Appliance GUI to add role to domain user/groups. Once a domain user has the appliance administrator role, the user can SSH to the host. Based on the IP used to SSH, the Access or the Appliance CLISH is launched. If you SSH to eth1 (management) IP of any node, the Appliance CLISH gets launched. If you SSH to console IP, the Access CLISH gets launched. Users with appliance administrator role and users that are part of groups with appliance administrator role can access the cluster through various interfaces such as SSH, GUI and REST APIs. Veritas recommends that you assign role to domain users/groups using the GUI.

The domain users and groups must have an UID and GID. This must be configured in the domain server so that you can SSH to the nodes. The domain users UID should start from 10000 onwards, else the UID will conflict with the local user UID.

To configure AD/LDAP servers for clusters deployed with only media servers

- 1 You have to first add a domain. Go to the **Appliance GUI > Settings > Directory service management** tab. Click **Configure**.
- 2 Select **Appliance management**.

Provide the AD/LDAP server details. The directory type can be Open LDAP or Active Directory. Click **Save**.


Configure directory service

LDAP server address *	Port *	Directory Type *
WIN-C4QHEF7964F.r7515-149vip27.v	636	Active Directory
The most commonly used LDAP ports are 389 and 636(using ssl encryption).		
User Base DN *	Group base DN *	
cn=users,dc=r7515-149vip27,dc=vxindia,dc=veritas,dc=coi	cn=users,dc=r7515-149vip27,dc=vxindia,dc=veritas,dc=coi	
Bind DN/LDAP administrator *	Bind DN/LDAP password *	
cn=administrator,cn=users,dc=r7515-149vip27,dc=vxindia,		
Encryption type		
SSL/TLS		
Certificate ☒ Uploaded Replace certificate		

Cancel
Save

- 3 Once the AD/LDAP server configuration operation is successful, the AD/LDAP server(s) appears in the list with usage as Appliance management.

You can perform a test connection to validate the newly added domain configuration. Click **Test connection**. For performing a test connection, you have to provide AD/LDAP password. You will receive notifications after the successful completion/failure of the test connection.



On the same page, there is an option to remove the domain configuration.

Roles can be assigned to a directory user as well as group.

To assign a role to AD/LDAP users and groups

- 1 From the navigation menu on the left, go to **Settings > User management**. The **User management** pane displays all the users that currently exist in the cluster configuration.
- 2 To add a new user/group, click **Add**.

- 3 Click on **Directory** for adding external users/group. The user/group should be specified in the format *username@domain*. For example, *testuser1@10.1.1.10* or *testgroup@mydomain.example.com*.

- 4 In the **Add administrator** dialog box, enter the details. Click **Add**.
- 5 Once the operation completes verify that the new user/group you added appears in the list of users displayed in the **User management** pane.

For local users, the user type appears as *Local*. For external users (AD/LDAP users/group), the user type appears as *Active Directory* and *Open LDAP*. It is specified whether the user is an individual user or group.

You can remove a particular user role by going to the user row and clicking on the vertical ellipsis button from the right side of the UI. Select **Delete Appliance Administrator role**.

Directory services and certificate management

The Lightweight Directory Access Protocol (LDAP) is the protocol used to communicate with LDAP servers. LDAP can be used as a directory service for user management. The LDAP server that is present outside the NetBackup Flex Scale cluster is responsible for authentication of users. For sites that use an LDAP server

for access or authentication, NetBackup Flex Scale provides a simple LDAP client configuration interface. The NetBackup Flex Scale cluster acts as an LDAP client talking to the LDAP server.

Active Directory (AD) is a technology created by Microsoft that provides a variety of network services including LDAP directory services, Kerberos-based authentication, Domain Name System (DNS) naming, secure access to resources, and more.

- You can configure AD/LDAP using the NetBackup Flex Scale GUI.
- You can configure or add only one AD/LDAP server at a time. But you can add multiple AD/LDAP servers to the NetBackup Flex Scale cluster.
- You can delete an existing AD/LDAP configuration.
- You cannot modify an existing AD/LDAP configuration. To modify any aspect of the AD/LDAP configuration, you have to delete the existing configuration and add it back with the updated parameters.

When you configure LDAP from the GUI, the domain ID is added as *LDAP_Server_FQDN/IP*. So, if you add LDAP user from the NetBackup GUI, you have to add the domain name as *<user_name>@LDAP_Server_FQDN/IP*.


If you have upgraded to NetBackup Flex Scale 3.0 from an earlier version in which the cluster was deployed with both primary server and media servers, the AD/LDAP servers which were already configured in the previous version will have the same name. For AD/LDAP servers which are newly configured on the upgraded cluster can have one of the following domain names:

- *<user name>@AD/LDAP server IP/FQDN*
- *<user name>@10.221.xx.xx*
- *<user name>@fqdn.domain*

When an AD/LDAP domain is configured using an SSL connection, you must provide the SSL certificate while adding the domain. The server's name used while adding the domain must exactly match with the server CN used for generating the certificate. If you used an FQDN to generate the certificate in a non-DNS environment, then you have to add the IP to FQDN mapping using the **Network > Custom Hosts** option in the Appliance UI. See [“Considerations for configuring AD/LDAP”](#) on page 31.

To configure AD/LDAP


- 1 Navigate to **Settings > Directory Services** and click **Configure**.



2 Enter the inputs required.

LDAP server address	IP address or FQDN of the AD/LDAP server
Port	Port number on which the AD/LDAP server is listening. If AD/LDAP is configured without SSL certificate, then port number should be 389. If AD/LDAP is configured with SSL certificate, then port number should be 636.
Directory type	Specify the directory type. It can be Open LDAP or Active Directory (when using Microsoft Active Directory)
User Base DN	Base DN subtree that is used when searching for user entries on the AD/LDAP server.
Group Base DN	Base DN subtree that is used when searching for group entries on the AD/LDAP server.
Bind DN/LDAP administrator	Distinguished name of the AD/ LDAP user who can search the AD/LDAP directory. Typically, it is the user name of the AD/LDAP server administrator.
Bind DN/LDAP password	Password for the given AD/LDAP administrator user
Encryption type	Specify the encryption type as secure or non-secure. In secure method, SSL/TLS is the encryption method.


If you choose the Encryption type to be SSL/TLS, you have to upload the certificate that you want to use to encrypt and secure the connection with the AD/LDAP server. Click **Choose file** and upload the certificate.



- 3 Click **Configure**.
- 4 You can test the connection after you configure it. Click **Test connection**. You receive notifications after the successful completion/failure of the test connection.

The AD/LDAP details appear in the **Directory Services** tab.

You can perform a test connection to validate the AD/LDAP configuration. On the same page, there is an option to remove the AD configuration.



Region settings management

You can use the **Settings > Region settings** to set the time zone for the cluster and change the Network Time Protocol (NTP) settings.

Cluster time management

You can manually set the time zone for the cluster to avoid problems caused due to cluster time inaccuracy.

To set up the time zone for the cluster

- 1 Go to **Settings > Region settings**.
- 2 Click **Set time zone**. The **Set cluster time zone** dialog box is displayed.
- 3 Under **Region Name**, click a region.

A list of city names is displayed under **City Name**. Choose a city name from the list or type the city name in the Search box.

Click **Next**.
- 4 The task to set the time zone is initiated. Click **Finish**.

View the **Recent activity** icon in the top navigation bar for updates.

NTP server management

NTP is a networking protocol for clock synchronization between computer systems.

You can use the IP address that you specified during installation to help sync up the clock in the cluster, but you need to specify the correct time in the applicable time zone. You can change the NTP server, if needed.

You can perform the following operations:

- Add one or more NTP servers to an existing list of NTP servers. NetBackup Flex Scale starts the synchronization process from the top of the list.
- Remove an NTP server, if multiple NTP servers are configured.
- You can see the status of the NTP server under **NTP server management > NTP server > Added NTP servers**.

To add an NTP server

- 1 **Settings > Region settings**. Under the **NTP server management** pane, click **Add**.
- 2 Specify the name or the IP address for the NTP server. Click **Next**.
- 3 The task to add the NTP server is initiated. Click **Finish**.

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

To remove an NTP server

- 1 **Settings > Region settings**. Under the **NTP server management** pane, click **Remove**.
- 2 Specify the name or the IP address for the NTP server that you want to remove. Click **Next**.
- 3 The task to remove the NTP server is initiated. Click **Finish**.

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

About NetBackup Flex Scale storage

NetBackup Flex Scale supports the following types of storage for each node. The type of storage might differ based on the hardware vendor platform.

The following table shows the storage for HPE vendor platform:

Device type	Number of devices per node	Used for
SATA SSD	2	Operating system and logs
SATA SSD	1	Certificates
NVMe SSD	2	NetBackup primary server catalog, which contains the internal databases that include information about NetBackup backups and configuration, Media Server Deduplication Pool (MSDP) catalog, and Data Change Object (DCO) for erasure-coded volumes
14-TB HDDs	12	Deduplicated backup data, Media file system, primary log file system, and scratch space for the MSDP catalog

The following table shows the storage for Dell vendor platform:

Device type	Number of devices per node	Used for
NVMe SSD	4	Operating system and NetBackup primary server catalog, which contains the internal databases that include information about NetBackup backups and configuration, Media Server Deduplication Pool (MSDP) catalog, and Data Change Object (DCO) for erasure-coded volumes
8-TB and 12-TB HDDs	12	Deduplicated backup data, Media file system, primary log file system, and scratch space for the MSDP catalog

File system layout

The following table shows the file systems and their layout:

File system	Layout	Device type
NetBackup primary catalog	Striped-mirror or mirrored	SSD
MSDP catalog	Striped-mirror or mirrored	SSD
MSDP data	Erasure-coded	HDD
Scratch space	Erasure-coded	HDD
Media	Erasure-coded	HDD
primary log	Erasure-coded	HDD

The scratch file system grows each time you add a node. The NetBackup primary server and the MSDP catalog file systems grow to double the size when you add the sixth node to the cluster. The NetBackup primary server and the MSDP catalog file systems grow when you add the 6th, 9th, 12th or the 15th node to the cluster.

A new MSDP data file system is created for each new node that is added to the cluster. The number of MSDP data file systems is equal to the number of nodes in the cluster.

Viewing storage utilization

You can view the total usable capacity and used storage on the NetBackup Flex Scale dashboard. The dashboard is displayed when you sign in to the NetBackup Flex Scale web interface using the `https://ManagementServerIPorFQDN:14161` URL where *ManagementServerIPorFQDN* is the public IP address that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration; or when you click **Dashboard** in the left navigation pane of the UI. From the **Dashboard** you can navigate to more detailed information. For example, under **Storage**, click **View details** to view the used storage, storage pool the disk is allocated to, and the disk size for each of the disks.

You can also use the following REST APIs to view the storage details:

- `GET /api/appliance/v1.0/storage/filesystems`
- `GET /api/appliance/v1.0/storage/filesystems/{fsName}`
- `GET /api/appliance/v1.0/storage/clusters/{clusterId}/storage-utilised`

Viewing alerts

If the storage usage exceeds 90% of the total available storage, the following alert is generated:

Netbackup file system [*fs name*] usage is above 90%

To clear or respond to the alert, you can free the used space or add more nodes to increase the storage.

To view alerts, do one of the following:

- Click **Dashboard** in the left pane of the NetBackup Flex Scale UI. In the **Alerts** area, click **View details** to see a complete list of alerts.
- At the top of any screen, click the **Bell** icon.
- Click **Settings > Alerts and notifications**. On the **Alerts and notifications** page, use the filters to locate specific types of alerts.

About Universal Shares

The Universal Share feature provides data ingest into a NetBackup Flex Scale appliance using an NFS or a CIFS (SMB) share. Space efficiency is achieved by storing this data directly into an existing NetBackup-based deduplication pool.

You can configure and manage the Universal Share from NetBackup. For more information, see the chapter "Configuring and using universal shares" in the NetBackup Deduplication Guide.

Advantages of Universal Shares

The following information provides a brief description of the advantages for using Universal Shares:

- As a NAS-based storage target
Unlike traditional NAS-based storage targets, Universal Shares offer all of the data protection and management capabilities that are provided by NetBackup.
- As a DB dump location
Universal Shares offer a space saving (deduplicated) dump location, along with direct integration with NetBackup technologies including data retention, replication, and direct integration with cloud technologies.
- Financial and time savings
Universal Shares eliminate the need to purchase and maintain third-party intermediary storage, which typically doubles the required I/O throughput since the data must be moved twice. Universal Shares also cut in half the time it takes to protect valuable application or DB data.
- Protection Points
The Universal Share Protection Point offers a fast point in time copy of all data that exists in the share. This copy of the data can be retained like any other data that is protected within NetBackup. All advanced NetBackup data management

facilities such as Auto Image Replication, Storage Lifecycle Policies, Optimized Duplication, cloud, and tape are all available with any data in the Universal Share.

- **Copy Data Management (CDM)**
The Universal Share Protection Point also offers powerful CDM tools. A read/write copy of any Protection Point can be "provisioned" or made available through a NAS (CIFS/NFS) based share. A provisioned copy of any Protection Point can be used for common CPD activities, including instant recovery or access of data in the provisioned Protection Point. For example, a DB that has been previously dumped to the Universal Share can be run directly from the provisioned Protection Point.
- **Backup without client software**
Client software is not required for Universal Share backups. Universal Shares work with any POSIX-compliant operating system that supports NFS or CIFS.

How it works

The Universal Share feature provides a network-attached storage (NAS) option for NetBackup Flex Scale appliances. Traditional NAS offerings store data in conventional, non-deduplicated disk locations. Data in a Universal Share is placed on highly redundant storage in a space efficient, deduplicated state. The deduplication technology that is used for this repository is the same MSDP location used by standard client-based backups.

Any data that is stored in a Universal Share is automatically placed in the MSDP, where it is deduplicated automatically. This data is then deduplicated against all other data that was previously ingested into the media server's MSDP location. Since a typical MSDP location stores data across a broad scope of data types, the Universal Share offers significant deduplication efficiency. The Protection Point feature lets you create a point in time copy of the data that exists in the specified Universal Share. Once a Protection Point is created, NetBackup automatically catalogs the data in that point and manages it like any other data that is ingested into NetBackup. Since the Protection Point only catalogs the Universal Share data that already resides in the MSDP, no data movement occurs. Therefore, the process of creating a Protection Point can be extremely fast.

Limitations:

- A Universal Share is presented through a single MSDP engine, and its size is limited by the capacity of the single node. If the Universal Share quota is set, its size is limited by its quota and cannot exceed the capacity of a single node.
- The maximum number of Universal Shares that can be created on an MSDP engine is limited, which is specified by the `MaxAllowedLivemounts` parameter

in the `spws.cfg` file and the default value of this parameter in NetBackup Flex Scale is 100.

Protection Point - cataloging and protecting Universal Share data

Any data that is initially ingested into a Universal Share resides in the MSDP located on the appliance-based media server that hosts the Universal Share. This data is not referenced in the NetBackup Catalog and no retention enforcement is enabled. Therefore, the data that resides in the Universal Share is not searchable and cannot be restored using NetBackup. Control of the data in the share is managed only by the host where that share is mounted.

A Protection Point is a point in time copy of the data that exists in a Universal Share. Creation and management of a Protection Point is accomplished through a NetBackup policy, which defines all scheduling and retention of the Protection Point. The Protection Point uses the “Universal-Share” policy type. Once a Protection Point for the data in the Universal Share is created, that point in time copy of the Universal Share data can be managed like any other protected data in NetBackup. Protection Point data can be replicated to other NetBackup Domains or migrated to other storage types like tape or cloud, using Storage Lifecycle Policies. Each Protection Point copy is referenced to the name of the associated Universal Share.

Protection Point restores

Restoring data from a Protection Point is exactly the same as restoring data from a standard client backup. The standard Backup Archive and Restore interface is used. The client name that is referenced for the restore is the Universal Share name that was used when creating the Universal-share policy type. Alternate client restores are fully supported. Make sure that NetBackup Client is installed on the destination client before restore. However, to restore to the system where the Universal Share was originally mounted, NetBackup Client software must be installed on that system. This is necessary since a NetBackup Client is not required to initially place data into the Universal Share.

Creating a Protection Point for a Universal Share

You can create a Protection Point for the data in a Universal Share that lets you manage and protect the data in the share. Creating a Protection Point is accomplished by creating a Universal-Share backup policy.

If an MSDP storage server is configured with multiple Universal Shares, a single policy can be created for some or all of the shares, and enable multi-stream option in policy. You can also create individual policies, one for each share. A universal share backup policy supports only one MSDP storage server.

To create a Protection Point policy for a Universal Share

- 1 Create a Universal Share on an existing MSDP storage server.
For details, see *Create a Universal Share* in [NetBackup Web UI Administrator's Guide](#).
- 2 Mount the exported path of the Universal Share on the storage server.
The **Export path** is found on the details page of the universal share in the NetBackup web UI: click **Storage > Universal Share** and then select the universal share to view its details.
- 3 Copy your application data to the Universal Share.
- 4 Using the NetBackup Administration Console (Java GUI), create a policy.
- 5 On the **Attributes** tab, select **Universal-Share**.
- 6 On the **Schedules** tab, select the type of backup. It can be **Full Backup**, **Differential Incremental Backup** or **Cumulative Incremental Backup**.

Note: **Accelerator** backups are not supported or necessary for Universal Shares.

- 7 On the **Clients** tab, enter the name of the desired client.
Universal share is an agentless technology, so the client name that is specified is used only for cataloging purposes. You can enter a NetBackup Appliance, NetBackup Virtual Appliance, Flex Appliance media server application instance, or MSDP server name or a host where universal share is mounted. The client name can be a short name, Fully Qualified Domain Name (FQDN), or IP address.

8 On **Backup Selections** tab, enter the NetBackup path of the universal share.

You can find the export path from the Universal share details page web UI: **Storage > Storage Configuration > Universal Share**. For example:

```
/mnt/vpfs_shares/3cc7/3cc77559-64f8-4ceb-be90-3e242b89f5e9
```

You can also use the `BACKUP X USING Y` directive, which allows cataloging under a different directory than the universal share path. For example: `BACKUP`

```
/demo/database1 USING
```

```
/mnt/vpfs_shares/3cc7/3cc77559-64f8-4ceb-be90-3e242b89f5e9. In this example, the backup will be cataloged under /demo/database1.
```

9 Run the **Universal-Share** policy.

After the backups are created, you can manage the backups with NetBackup features, such as restore, duplication, Auto Image Replication, and others.

You can instantly access the backups with NetBackup Instant Access APIs.

For information about NetBackup APIs, see the following website:

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

Select NetBackup and then the version at the bottom of the page.

Node and disk management

At a minimum, you must deploy four nodes in a cluster and can scale up to a maximum of sixteen nodes. All the nodes must have the same hardware and the same software version.

A cluster with up to five nodes provides resiliency of one node and a disk. The cluster can tolerate a loss of one node and a disk and continue to perform NetBackup jobs. A cluster with six or more nodes provides resiliency of two nodes, or two SSDs, or four HDDs. If the loss of cluster nodes is more than the supported resiliency, either due to node failure or nodes are stopped or shut down, the cluster goes in an inconsistent state. If a cluster is in an inconsistent state, check the NetBackup Flex Scale Release Notes to see if a related issue and the resolution is mentioned; else call Veritas Technical Support for assistance with resolving the inconsistent cluster state.

Replace a faulted node or disk in the cluster to recover from a node or a disk failure. Add a new node to the cluster to increase the usable cluster storage and compute resources. When you add a node or replace a faulted node, the NetBackup operations continue to run on the cluster nodes.

NetBackup Flex Scale network cabling

Review the following cabling information and ensure that the nodes are cabled correctly before you add a new node to the cluster or replace an unhealthy node in the cluster with a new node.

Management network

The eth1 1 GbE network interface on each node provides access to the management network. Connect the eth1 network interface to a switch.

Private network

The eth4 and eth6 25 GbE network interfaces on each node provide access to the private network and are used by the nodes in the cluster to communicate with each other. Veritas recommends using two separate Ethernet switches, which are capable of supporting 25 Gb link to isolate the eth4 and eth6 interfaces. If you use a single Ethernet switch, connect the eth4 and eth6 interfaces of each node to switch ports in a separate VLAN. Network bonding cannot be configured for these interfaces.

Public or the data network

The eth5 and eth7 10/25 GbE network interfaces on each node provide access to the public data network, which is used to run production workloads. If the public network and the management network (eth1) are in the same network, connect the eth5, eth7, and eth1 network interfaces to a single Ethernet switch capable of supporting 10/25 Gb link.

If the public network and the management network are in a different network, connect the eth5, eth7, and eth1 network interfaces to separate Ethernet switches capable of supporting 10/25 Gb link.

The eth5 and eth7 network interfaces can be used individually or can be bonded. If you configure a network bond, the required settings must be configured on the switch.

Power cabling

Use two power cables to connect the dual power supply modules of the node to dual Power Distribution Units (PDUs).

For more details about the cabling information, refer to the NetBackup Flex Scale Hardware Cabling poster that is included in the Open Me First kit that is shipped with your appliance.

Adding a node to the cluster using the NetBackup Flex Scale web interface

Add a new node to the cluster to increase the usable storage of the cluster and its compute resources. When the add node operation is in progress, the backup and restore jobs continue to run on the cluster nodes. You can add multiple nodes at a time. The new node must have the same hardware configuration as the existing cluster nodes. When you add a node to the cluster, data is rebalanced to distribute it evenly across the cluster nodes. The time it takes to rebalance the data depends on the amount of data being moved between nodes and whether you choose to prioritize NetBackup jobs or data rebalancing.

Note: The fd00:200/120 network is reserved and used internally by NetBackup Flex Scale, and it should not be used anywhere.

You cannot perform the following operations when the add node operation is in progress:

- Add another node to the cluster.
- Replace an existing node in the cluster.
- Upgrade a node in the cluster.
- Add or modify existing data networks.
- Create, edit, or delete a network bond.

Note: The add node operation is irreversible; once you start the add node operation, you cannot cancel it midway through or revert the operation. Contact Veritas Support if you face any issues.

Before you add a new node to the cluster, ensure that the following requirements are met.

Health of the cluster nodes

Verify that all the existing nodes in the cluster are in a healthy state. You cannot add a node if a cluster node is not healthy.

ISO image

Power on the node and install the NetBackup Flex Scale ISO image on the node that you want to add to the cluster. For more details, see the *NetBackup Flex Scale Installation and Configuration Guide*. After the ISO is installed, perform a factory reset to reset the node to its default factory settings.

You must format the disks and ensure that no data is present on the disks; else the add node operation will fail.

Network cabling

The node that you want to add must be connected to the same private and public network as that of the existing cluster. For details about the cabling information, See [“NetBackup Flex Scale network cabling”](#) on page 54.

Jumbo frames

Set the maximum transmission unit (MTU) property, which controls the maximum transmission unit size for an Ethernet frame to 9000 bytes. By default the MTU is set to 1500 bytes. For optimal performance, it is recommended to set a larger frame size to enable jumbo frames for the eth4, eth5, eth6, and eth7 network interfaces. To take advantage of jumbo frames, the Ethernet cards, drivers, and switching must all support jumbo frames.

System clock

Ensure that the system clock is in sync with the cluster nodes. If the system clock is not synchronized, the add operation fails.

IP address requirements

Ensure that you have the following details for each node that you want to add:

Table 3-2

Number of IP addresses	Description
1	Public IP address and a fully qualified domain name (FQDN) for the media server. If multiple data networks are configured for the cluster, you require one public IP address for each data network. For example, for n data networks, you need n number of public IP addresses for the media server.
1	Virtual IP address and an FQDN for the storage server. If multiple data networks are configured for the cluster, you require one virtual IP address for each data network. For example, for n data networks, you need n number of virtual IP addresses for the storage server.

Table 3-2 (continued)

Number of IP addresses	Description
1	Public IP address for the management network interface. Note: If the management network is on a VLAN and the primary data network is on a different VLAN, the management gateway cannot communicate with NetBackup web services.
1	Public IP address for the IPMI interface if the IPMI network is configured for the cluster.


Note: If disaster recovery is configured, See [“Considerations for adding a node when disaster recovery is configured”](#) on page 60. before you add a new node.

To add a new node the cluster, complete the following steps:

- 1 Use any one of the following options to log in using the user account that you created when you configured the cluster:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
<https://ManagementServerIPorFQDN/webui> where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Cluster Monitor > Infrastructure**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
<https://ManagementServerIPorFQDN:14161> where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API

gateway during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.

- 2 On the **Nodes** tab, click **Scan for nodes** to discover the nodes that can be added to the cluster.



The nodes connected to the private network of the cluster are discovered and displayed under the **Discovered nodes** section with details such as the node name, storage capacity, serial number, status, and the data network and the management network of the node.

- 3 Under the **Discovered nodes** section, select the nodes that you want to add and click **Add nodes to the cluster**.

Note: You can only add healthy nodes to the cluster.

The node and storage details of the existing cluster and the projected storage summary after you add the nodes is displayed.

- 4 Specify whether you want to set the priority for NetBackup backup and recovery jobs or data rebalancing. By default, the priority is set to backup and recovery jobs while data rebalancing runs in the background at a lower priority (option **Overall system performance**). Click **Faster reconfiguration** to set the priority to rebalancing data across the cluster. This option increases the speed of node addition operations so that the new node can begin participating in backup and recovery operations faster but can affect the backup and recovery jobs.
- 5 Enter the following information for the selected nodes:
 - **Host name:** Name for the node. The node name can contain a maximum of 63 characters, excluding the domain name.
By default, the auto-generated node name based on cluster and domain name is displayed.
 - **Data-Network-<device-name>:**

You have the option to specify both the IP address and a fully qualified domain name (FQDN), or you can specify only the IP address for the media and the storage server.

The FQDN of the media server and the storage server can contain a maximum of 64 characters, including the domain name.

On the **Automatic** tab, specify a single IP range, multiple IP ranges separated by a comma, comma-separated individual IP addresses, a combination of individual IP addresses and IP ranges separated by a comma, or IP addresses in CIDR format, and then click **Add**. You are not required to enter the FQDN for the media and the storage server. The FQDN is calculated based on the specified IP addresses.

Click the **Custom** tab to specify a fully qualified domain name and the IP address of the media server and the storage server for each node. The media server and the storage server names can contain a maximum of 64 characters, including the domain name.

- **Management Network:**

You have the option to specify both the IP address and an FQDN, or you can specify only the IP address to assign to the management network of the nodes that you want to add.

On the **Automatic** tab, specify a single IP range, multiple IP ranges separated by a comma, comma-separated individual IP addresses, a combination of individual IP addresses and IP ranges separated by a comma, or IP addresses in CIDR format to assign to the eth1 management network interface, and then click **Add**. You are not required to enter the FQDN. The FQDN is calculated based on the specified IP addresses.

Click the **Custom** tab to specify a fully qualified domain name and the IP address to assign to the eth1 management network interface. If you have already configured the IP address for eth1, the IP address is displayed.

- **IPMI Network:** If an IPMI network is configured for the cluster, specify a public IP address to assign to the IPMI interface.

On the **Automatic** tab, you can specify a single IP range, multiple IP ranges separated by a comma, comma-separated individual IP addresses, a combination of individual IP addresses and IP ranges separated by a comma, or IP addresses in CIDR format to assign to the IPMI network interface, and then click **Add**.

Click the **Custom** tab to specify a public IP address to assign to the IPMI network interface of each of the selected nodes.

6 Click Add nodes to the cluster.

The progress of the add node operation is displayed on the **Infrastructure** page.

7 To monitor the status of each of the tasks, click View details on the Infrastructure page. The ongoing and completed tasks for the add node operation are also displayed in Recent activity.

Before adding the selected nodes to the cluster, the add node operation synchronizes the patches or add-ons that are installed on the cluster with the nodes. However if the version on the selected nodes is later than the cluster nodes, the add operation cannot proceed. If AutoSupport and Call Home settings are configured for the cluster, these settings are synchronized with the nodes.

Note: If the add node operation fails during rebalancing of data, contact Veritas Support to resolve the issue. Do not reimagine the node or perform a factory reset.

After the node is added, you can view the new node and its details on the **Nodes** tab. The STIG status of cluster is synchronized with the newly added nodes. If the STIG option is enabled for the cluster, the STIG option is also enabled for the newly added nodes and vice versa. The FIPS option is enabled for the added nodes as this option is enabled with the default factory settings for the Veritas Operating System (VxOS) and for NetBackup MSDP when you create a NetBackup Flex Scale cluster.

Considerations for adding a node when disaster recovery is configured

If disaster recovery is configured and you want to add a new node to the cluster, review the following guidelines:

- Add the new node independently on each cluster.
- The NetBackup catalog file system is grown on both the sites only if the secondary site can also grow the catalog file system.

Adding a node using the REST APIs

You can use REST APIs to add a node to an existing cluster.

You can find the RESTful APIs at

<https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0/> where *ManagementServerIPorFQDN* is the public IP address or FQDN that you specified

for the NetBackup Flex Scale management server and API gateway during the cluster configuration.

If you are using IPv6 addresses, use the following URL

syntax:`https://[ManagementServerIP]:14161/swagger/infra/v1.0`

See [“Working with NetBackup Flex Scale APIs”](#) on page 19. for more details.

Before you use the APIs, power on the node and install the NetBackup Flex Scale ISO image on the node that you want to add to the cluster. For more details, see the *NetBackup Flex Scale Installation and Configuration Guide*. After the ISO is installed, perform a factory reset to reset the node to its default factory settings. Ensure that the management network interface (eth1) of the node is not configured with the management IP address.

You must format the disks and ensure that no data is present on the disks; else the add node operation will fail.

Use the following APIs in the sequence listed below to add a node. For more details about the APIs, see the *Veritas NetBackup Flex Scale APIs on SORT*.

Note: You can add only one node at a time using REST APIs. Use the NetBackup Flex Scale web interface to add multiple nodes simultaneously.

- 1 Authenticate the user. The user must be assigned an Appliance Administrator role.

```
POST /api/appliance/v1.0/authentication/login
```

Specify the user name and password. The API returns a token. Copy the generated token and specify in the following format:

```
Authorize: Bearer generated token
```

You can also generate an API key that is valid for 10 years after the user is authenticated successfully:

```
POST /api/appliance/v1.0/api-keys
```

After the API key is generated, generate the token again using the following API and by specifying the user name and the generated API key:

```
POST /api/appliance/v1.0/authentication/login
```

Note: If both the password and the API key is provided with the user name to generate the token, the API considers only the password and generates the token.

- 2 Discover all the spare nodes that are in the same private network as the existing cluster.

```
GET /api/appliance/v1.0/storage/new-nodes
```

- 3 Get details of a particular node that was discovered in step 1 by using its Avahi IP address.

```
GET /api/appliance/v1.0/storage/new-nodes/{ipAddress}
```

- 4 Synchronize the patch version on the cluster with the new node. All the patches that are installed on the cluster nodes are installed on the new node.

```
POST /api/appliance/v1.0/upgrade
```

With {"eebType": "normal"}: Payload for synchronizing patch upgrade version on new node with the cluster.

- 5 Add the node to the cluster.

```
POST /api/appliance/v1.0/storage/nodes
```

- 6 Rebalance the data across all the nodes to free space on the existing cluster nodes. Rebalancing spreads data evenly across the nodes by moving the data from existing nodes to the new node.

```
POST /api/appliance/v1.0/storage/rebalance
```

- 7 Configure NetBackup on the newly added node.

```
POST /api/appliance/v1.0/netbackup/add
```

- 8 Synchronize data EEBs on the cluster with the newly added node.

```
POST /api/appliance/v1.0/upgrade
```

With {"eebType": "data"}: Payload for synchronizing patch upgrade version on new node with the cluster.

Replacing a node in a cluster

You can replace a cluster node if the node is in an unhealthy state because of a hardware failure that cannot be repaired, such as a boot disk or a power supply failure.

When you replace the node, data from the faulted node is rebuilt on the new node. During this time, the backup and recovery operations continue to run on the remaining healthy nodes in the cluster.

The network settings and the node name of the faulted node are assigned to the new node. After the NetBackup services fail over to the new node, it can run backup and recovery jobs.

Note: After a node is deleted from the cluster, its private IP addresses are kept reserved in the cluster and are not used when a new node is added in the cluster.

Review the following guidelines for replacing a node:

- You can replace only a single node at a time.
- The new node must have the same hardware configuration as the existing cluster nodes.
- You cannot perform any other node operations such as adding a new node when the node is being replaced. A node cannot be replaced if any of the previous cluster reconfiguration tasks are in progress or have failed.

- All the private and public NICs of the node where the management server is running must be up. If any of the NICs are down, the replace node operation fails.
- You cannot replace a node if there is more than one faulty node in a cluster of up to five nodes. A 5-node cluster provides resiliency of one node and a disk. A cluster with six or more nodes provides resiliency of two nodes, or two SSDs, or four HDDs. For a cluster with six or more nodes, if two or more nodes are faulty, contact Veritas Support. You must contact Veritas Support if failures exceed the supported fault tolerance.
- If you repair the faulty node and plan to use the same node as a replacement node, Veritas recommends that you reinstall the NetBackup Flex Scale ISO image on the node and perform a factory reset to reset the node to its default factory settings.
- If site-based disaster recovery is configured, wait at least two hours after a node failure before attempting to replace the node.

Before you replace the node, ensure that the following requirements are met.

ISO image

Power on the new node and install the NetBackup Flex Scale ISO image on the node that you want to add to the cluster. Ensure that EEB 4067542 is installed on the node. For more details, see the *NetBackup Flex Scale Installation and Configuration Guide*.

You must format the disks and ensure that no data is present on the disks; else the replace node operation will fail.

Network cabling

The new node that you want to use to replace the existing node must be connected to the same private and public network as that of the existing cluster. For details about the cabling information, See [“NetBackup Flex Scale network cabling”](#) on page 54.

System clock

Ensure that the system clock of the new node is in sync with the cluster nodes. If the system clock is not synchronized, the replace node operation fails.

To replace a faulted node from the cluster, complete the following steps:

- 1 Use any one of the following options to log in using the user account that you created when you configured the cluster:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator

role to log in to the NetBackup Flex Scale web interface

`https://ManagementServerIPorFQDN/webui` where

ManagementServerIPorFQDN is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Cluster Monitor > Infrastructure**.

- Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where
ManagementServerIPorFQDN is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.

2 Click the **Nodes** tab.

The list of cluster nodes is displayed with the unhealthy node marked in faulted state.

3 If the node that you want to replace is up but in an unhealthy state because of a hardware failure, power off the node. To power off the node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI, and click **Shutdown node**.

If you are in the NetBackup Flex Scale UI, when prompted, click **Open cluster console** to open the NetBackup Flex Scale infrastructure management console in a new browser tab. In the NetBackup Flex Scale infrastructure management console, when prompted to confirm, click **Shut down node**.

4 Click **Scan for nodes** to discover the nodes that are connected to the private network and can be used to replace the unhealthy node.

The nodes in the private network are discovered and displayed under the **Discovered nodes** section with details such as the node name, serial number, and status.

5 For the unhealthy node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI, and click **Replace node**.

If you are in the NetBackup Flex Scale UI, when prompted, click **Open cluster console** to open the NetBackup Flex Scale infrastructure management console in a new browser tab. In the NetBackup Flex Scale infrastructure management console, when prompted to confirm, click **Replace node**.

6 In the **Replace node** dialog box, complete the following steps:

- Set the priority to backup and recovery jobs or data rebalancing. Click **Overall system performance** to set the priority to NetBackup backup and

recovery jobs while data rebalancing runs in the background at a lower priority. Click **Faster reconfiguration** to set the priority to rebalancing data across the cluster. This option increases the speed of node replacement operations so that the new node can begin participating in backup and recovery operations faster but can affect the backup and recovery jobs. The time it takes to rebalance is dependent upon the amount of data being moved between nodes and the priority being set. By default the priority is set to **Overall system performance**.

- Select the node that you want to use to replace the unhealthy node.
 - Click **Replace node**.
- 7 To monitor the status of each of the tasks and the progress of the replace node operation click **View details** on the pop-up window that is displayed on the **Infrastructure** page. The ongoing and completed tasks for the replace node operation are also displayed in **Recent activity** or you can go to the **Dashboard** to check the health of all the nodes.

Before adding the selected node to the cluster, the replace node operation synchronizes the patches or add-ons that are installed on the cluster with the node. However if the version on the selected node is later than the cluster nodes, the replace node operation cannot proceed. If AutoSupport and Call Home settings are configured for the cluster, these settings are synchronized with the node.

After the unhealthy node is replaced, you can view the new node and its details on the **Nodes** tab. The STIG status of cluster is synchronized with this new node. If the STIG option is enabled for the cluster, the STIG option is also enabled for the node and vice versa. The FIPS option is enabled for the node as this option is enabled with the default factory settings for the Veritas Operating System (VxOS) and for NetBackup MSDP when you create a NetBackup Flex Scale cluster.

Starting and stopping nodes

You might need to stop a cluster node for hardware maintenance such as replacing some of the hardware components. For a cluster with less than six nodes, only a single node can be down or you can stop only a single node at any given point in time. For a larger cluster of up to 16 nodes, a maximum of two nodes can be down or can be stopped at any given point in time. When you stop a node, the NetBackup jobs running on the node fail over to other nodes in the cluster and the cluster services running on the node are stopped. After the hardware maintenance is complete, you need to start the node. When you start a node, the cluster services

are started on the node, the node joins the cluster and can start running backup jobs.

If you stop or shut down the node where the NetBackup Flex Scale infrastructure management console (UI) is running, the infrastructure management console fails over to another node. It can take a few minutes for the management console to be up on another node.

Note: If the loss of nodes exceeds the supported fault tolerance, either due to node failures or nodes are stopped or shut down, the cluster goes in an inconsistent state.

To stop and start a node:

- 1 Use any one of the following options to log in using the user account that you created when you configured the cluster:
 - Use a user account with both Appliance administrator and NetBackup administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Cluster Monitor > Infrastructure**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.
- 2 Click the **Nodes** tab.
The list of cluster nodes is displayed.

- 3 To stop the node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI, and click **Stop node**.

If you are in the NetBackup Flex Scale UI, when prompted, click **Open cluster console** to open the NetBackup Flex Scale infrastructure management console in a new browser tab. In the NetBackup Flex Scale infrastructure management console, when prompted to confirm, click **Stop node**.

The node status changes to unhealthy.

- 4 To start the node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI, and click **Start node**.

If you are in the NetBackup Flex Scale UI, when prompted, click **Open cluster console** to open the NetBackup Flex Scale infrastructure management console in a new browser tab. In the NetBackup Flex Scale infrastructure management console, when prompted to confirm, click **Start node**.

The node status changes to healthy.

Replacing a disk

If a disk of a node is in a faulted state, replace the disk to maintain the resiliency of the cluster. A cluster with up to five nodes can tolerate a loss of one node and a disk. A cluster with six or more nodes provides resiliency of two nodes, or two SSDs, or four HDDs. The NetBackup jobs continue to run as long as the fault tolerance is not exceeded. If the failures exceed the fault tolerance, the cluster runs in a degraded state where NetBackup and MSDP services might not be running on all the nodes.

Alerts are generated for faulty disks. See [“Viewing information about alerts”](#) on page 103. If Call Home is configured for your setup, diagnostic information is sent to the AutoSupport server.

Before you begin the disk replacement operation, ensure that replacement disk is the correct size and is formatted.

Warning: Ensure that the following steps are followed correctly; else it might lead to data loss.

To replace a disk:

- 1 Remove the faulty disk from the node and replace the faulty disk with a replacement disk.

Note: On HPE ProLiant Server setup, you must power off the node before physically replacing the faulty NVMe SSD disk.

- 2 Use any one of the following options to log in using the user account that you created when you configured the cluster:

- Use a user account with both Appliance administrator and NetBackup administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface

`https://ManagementServerIPorFQDN/webui` where

ManagementServerIPorFQDN is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Cluster Monitor > Infrastructure**.

- Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console

`https://ManagementServerIPorFQDN:14161` where

ManagementServerIPorFQDN is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.

- 3 Click **Disks**.

The disks of all the nodes are displayed. In the **Status** column, the failed disks are marked as faulted.

- 4 Click the faulty disk that you want to replace, and then click **Replace disk**.

If you are in the NetBackup Flex Scale UI, when prompted, click **Open cluster console** to open the NetBackup Flex Scale infrastructure management console in a new browser tab. In the NetBackup Flex Scale infrastructure management console, when prompted to confirm, click **Replace disk**.

The replacement disk is detected and added to the node.

After the disk is replaced, data rebuild process begins and you can monitor the progress at the bottom of the screen when you click **Infrastructure > Monitor > Disks**. Time taken to rebuild the data depends on the amount of data.

Viewing disk details

You can view the details about the disks on the **Infrastructure** page. To go to the Infrastructure page, use any one of the following options:

- Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
<https://ManagementServerIPorFQDN/webui> where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Cluster Monitor > Infrastructure**.
- Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
<https://ManagementServerIPorFQDN:14161> where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.

You can view the following information:

- Name of the disk.
- Status of the disk. The disk status is online if the disk functions normally and faulted if the disk has hardware issues or is faulty.
- Disk space allocated to the storage pool. The file systems are created on this disk space during the initial configuration or when the disks are allocated to the storage pool when nodes are added to the cluster.

Note: The disk space does not reflect the size of the backup data. The total storage that is used for the data is displayed under **Storage** on the **Dashboard**.

- Total disk size.
- Storage pool the disk is allocated to.
- Whether the disk is designated as a spare disk.
- Enclosure the disk belongs to.
- Nodes that can access the disk.

To view detailed information about a disk, click the disk name. You can view the following additional information:

- Serial number.
- Vendor ID.
- File system that is created on the disk.
- Node the disk resides on and the path of the disk on the node. The path is displayed in *node@path* format. For example, *pbnso01@/dev/sdh* where *pbnso01* is the name of the node and *sdh* is the name of the disk.

Viewing node details

You can view the details about the disks on the **Infrastructure** page. To go to the Infrastructure page, use any one of the following options:

- Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
<https://ManagementServerIPorFQDN/webui> where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Cluster Monitor > Infrastructure**.
- Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
<https://ManagementServerIPorFQDN:14161> where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.

You can view the following information:

- Name of the node.
- Status of the node. The node status is healthy if the node is operational. The node status is unhealthy if the node is shut down or stopped. A green checkmark next to the status indicates that the node is running. A red exclamation mark icon indicates that the node is in a faulted state or has exited from the cluster.
- Serial number of the node.
- Health of the node.
- Product version that is installed on the node.
- IP address assigned to the eth1 network interface of the node.
- CPU and memory utilization.

To view detailed information about the node, click the node name. You can view additional information such as the hardware details and the IP addresses allocated to the node.

On the **Utilization statistics** tab, you can view the graphs for CPU, memory, and network usage. Place the cursor anywhere on the horizontal axis to view the usage trends.

On the **Disks** tab, you can view all the disks the node can access, the usage in terms of percentage of the total available disk capacity, the storage pool the disk is allocated to, and the vendor.

License management

You can add NetBackup Flex Scale storage and NetBackup licenses while configuring the cluster. The initial configuration wizard displays a license page where you can specify the license details. However, entering the licenses during the cluster configuration itself is not mandatory. In case you skip the licensing page, the cluster is automatically configured with an in-built trial license. Thereafter, you can use the NetBackup Flex Scale infrastructure management console UI to manage your appliance licenses. A valid license is required to maintain a working cluster configuration. Veritas recommends that you add proper licenses before you start protecting production workloads with your appliance.

See [“Adding or removing storage licenses”](#) on page 72.

See [“Adding NetBackup licenses”](#) on page 75.

Adding or removing storage licenses

You can use the NetBackup Flex Scale UI or the REST APIs to add or remove storage licenses.


Managing storage licenses using the UI

Perform the following steps to manage NetBackup Flex Scale storage licenses in your cluster using the infrastructure management console UI.


To manage storage licenses

- 1 Sign in to the NetBackup Flex Scale infrastructure management console UI.
See [“About the NetBackup Flex Scale management console”](#) on page 13.
- 2 From the navigation menu on the left, click **Settings** and then click **Licensing**.
The Storage Licenses tab displays all the licenses that currently exist in the cluster configuration.

- 3
- To add a storage license, click **Add license**, and on the Add storage licenses page click **Choose file** and browse and select a license file. Alternatively, you can also drag and drop the license file directly on the page.



The page shows the licensed capacity that will be available for use after the license is added and also indicates if you comply with the licensing model. You can select multiple license files; the license capacity reflects the total licensed capacity that will be available.



Current license capacity

Per terabytes (TiB)

Type

Trialware

Expiration date

Nov 02 2021

New license capacity

Per terabytes (TiB)

0.3125 TiB Available

1782 TiB Licensed

You will be meeting compliance standards

Next upcoming support period expiration date: Sep 26 2024

Drag and drop your license file here or select your file

Choose file

Search...

Q

☐ License	Type	Capacity	Description	Effective date	Support period expiration	
☐ A5936408149_QTY1374_NBFS_1_0_LIC... PERPETUAL		1374 TiB	VERITAS NBFS LICENSE	Sep 27 2019	Sep 26 2024	
☐ A9034143784_QTY408_NBFS_1_0_LIC... PERPETUAL		408 TiB	VERITAS NBFS LICENSE	Nov 06 2020	Nov 05 2026	

- 4 Click **Add** to configure the license into the cluster.
- A notification is displayed on top of the page. To monitor the progress and view task details click **View Details**. Details about the license such as the license type, capacity, status, and the duration for which Veritas Support services are included are listed. If you added multiple licenses, the details of all the added licenses are listed.
- 5 To remove a particular storage license, in the license row, click on the vertical ellipsis button on the right side of the UI and then select **Delete**.

Licensing management

Reference

B

A

Jan 2020

Jul

Jan 2021

Jul

Jan 2022

Jul

Jan 2023

Jul

Jan 2024

Jul

Jan 2025

Jul

Jan 2026

Jul

1 Selected

Delete

Search...

Q

+ Add license

☐ Ref	License	Type	Capacity	Description	Effective date	Support period expiration	Status	
☐ B	A5936408149_QTY1374_NBFS...PERPETUAL		1374 TiB	NETBACKUP FLEX SCALE 1 TB...Sep 27 2019	Sep 26 2024		Active	
☒ A	A9034143784_QTY408_NBFS...PERPETUAL		408 TiB	NETBACKUP FLEX SCALE 1 TB...Nov 06 2020	Nov 05 2026		Active	

Items per page: 4 1 - 2 of 2

- 6 When prompted for confirmation, click Delete. To monitor the progress of the delete operation, click **View details**.
- 7 Once the operation completes, verify that the license you removed does not appear in the list of licenses displayed on the Storage Licenses tab.

Managing storage licenses using REST APIs

You can use REST APIs to upload, add, and delete storage licenses in the cluster. You can find the REST APIs at

`https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0/` where *ManagementServerIPorFQDN* is the public IP address or FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration. For more details about the APIs, see the NetBackup Flex Scale APIs on SORT.

Adding NetBackup licenses

You can use the NetBackup Flex Scale UI or the REST APIs to add NetBackup licenses.

Managing NetBackup licenses using the UI


Perform the following steps to add NetBackup licenses to your cluster using the infrastructure management console UI. You can only add NetBackup licenses from this UI. After adding the NetBackup licenses to the cluster, you can manage them using the NetBackup Java Administration Console.

To add NetBackup licenses

- 1 Sign in to the NetBackup Flex Scale infrastructure management console UI. See [“About the NetBackup Flex Scale management console”](#) on page 13.
- 2 From the navigation menu on the left, click **Settings** and then click **Licensing management**.
- 3 Click **NetBackup Licenses**.

The NetBackup Licenses tab displays all the NetBackup licenses that currently exist in the cluster configuration.

- 4 To add a new NetBackup license, click **Add a license** and on the Add a license dialog, enter a NetBackup license key in the License key field.



- 5 Click **Save and add another** if you wish to add another license.
Repeat this step if you want to add multiple license keys simultaneously.
- 6 Click **Add** to configure the license into the cluster.

Managing NetBackup licenses using REST APIs

You can use REST APIs to upload, add, and delete NetBackup licenses in the cluster. You can find the REST APIs at

`https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0/` where *ManagementServerIPorFQDN* is the public IP address or FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration. For more details about the APIs, see the NetBackup Flex Scale APIs on SORT.

NetBackup Flex Scale network management

This chapter includes the following topics:

- [About network management](#)
- [Modifying DNS settings](#)
- [About bonding Ethernet interfaces](#)
- [Bonding operations](#)
- [Configuring NetBackup Flex Scale in a non-DNS environment](#)
- [Data network configurations](#)

About network management

You can perform network-related operations in the **Settings > Network** panel. You can find the details of the network settings for the data network, management network, appliance management console, DNS, and the IPMI network for the cluster.

The data network that is created during the initial configuration is the primary network and the new data networks added by the user after the initial configuration are the secondary data networks.

Network > Data Network displays the data networks configured in the cluster. You can add new secondary data networks and modify or delete existing secondary data networks. You can perform interface bonding at the time of initial configuration or after cluster configuration. You can also create a bond of interfaces, modify the existing bond modes and Maximum Transmission Value (MTU), display current bond mode and remove bond(s). There may be network disruptions when you perform bonding operations.

See [“Data network configurations”](#) on page 93.

See [“Bonding operations”](#) on page 81.

Network > DNS displays the details about the Domain Name Server (DNS).

The **Data network** and **Management network** display the details of the data network and management network.

- Domain name
- DNS servers
- Search Domains

You can only edit the DNS details under the **Management network**.

See [“Modifying DNS settings”](#) on page 79.

Enabling the DNS server is optional. You can configure the cluster such that both the data and management network do not have DNS.

See [“Configuring NetBackup Flex Scale in a non-DNS environment”](#) on page 89.

Network > Management Network displays about the management network.

It includes information such as:

- Routing settings
This includes details such as the subnet mask and gateway of IPv4 addresses and details such as prefix length and router address of IPv6 addresses.
- Management interfaces
Specifies the management interface assignment for each node in the cluster.

Network > IPMI Network displays the details about the IPMI network, if configured.

It includes information such as:

- Routing settings
This includes details such as the subnet mask and gateway of IPv4 addresses and details such as prefix length and router address of IPv6 addresses.
- IPMI interfaces
Specifies the IPMI interface assignment for each node in the cluster.

Network > Appliance specifies the network settings for the appliance management console. If you did not configure the management network earlier, these settings apply only on the data network.

The following details are displayed:

- Console IP
- API Gateway IP

- API Gateway node
- API Gateway FQDN

Note: If you deploy a cluster with only media servers, only the console IP is displayed.


Network > Custom Hosts lets you configure custom host files for cluster nodes or NetBackup services using a hosts file to map host names and domains to IP addresses.

Modifying DNS settings

Network > DNS displays the details about the Domain Name Server (DNS). The **DNS** screen has two tabs. The **Data network** and **Management network** tabs display the DNS server details for the respective networks.

You can only edit the DNS servers that you have specified under the **Management network**.

- Go to **Network > DNS > Management network**.
- Click **Edit** to modify the DNS settings.
- Click **Submit** to save your changes.
- You can add or remove the DNS server name and search domain.
- If STIG is enabled on the cluster, then it is recommended that you configure at least two DNS servers.



About bonding Ethernet interfaces

Bonding associates a set of two or more Ethernet interfaces with one IP address. The association improves network performance on each NetBackup Flex Scale cluster node by increasing the potential bandwidth available on an IP address beyond the limits of a single Ethernet interface. Bonding also provides redundancy for higher availability.

For example, you can bond two 1 gigabit Ethernet interfaces together to provide up to 2 gigabits per second of throughput to a single IP address. Moreover, if one of the interfaces fails, communication continues using the single Ethernet interface.

When you create a bond, you need to specify a bonding mode. In addition, for the following bonding modes: `802.3ad`, `active-backup`, `balance-rr`, `balance-xor`, `broadcast`, `balance-tlb`, and `balance-alb`, make sure that the base network interface driver is configured correctly for the bond type. For type `802.3ad`, the switch must be configured for link aggregation.

The `802.3ad` and `balance-xor` bond types have an option of sub-types, `layer2` and `layer(3+4)`.

Consult your vendor-specific documentation for port aggregation and switch set-up. You can use the `-s` option in the Linux `ethtool` command to check if the base driver supports the link speed retrieval option. The `balance-alb` bond mode type works only if the underlying interface network driver enables you to set a link address.

Note: An added IPv6 address may go into a TENTATIVE state while bonding Ethernet interfaces with `balance-rr`, `balance-xor`, or `broadcast` bond modes. While bonding with those modes, NetBackup Flex Scale requires the switch to balance incoming traffic across the ports, and not deliver looped back packets or duplicates. To work around this issue, enable EtherChannel on your switch, or avoid using these bond modes.

Table 4-1 Bonding mode

Index	Bonding mode	Fault tolerance	Load balancing	Switch setup	Ethtool/base driver support
1	<code>balance-rr</code>	Yes	Yes	Yes	No
2	<code>active-backup</code>	Yes	No	No	No
3	<code>balance-xor</code>	Yes	Yes	Yes	No
4	<code>broadcast</code>	Yes	No	Yes	No

Table 4-1 Bonding mode *(continued)*

Index	Bonding mode	Fault tolerance	Load balancing	Switch setup	Ethtool/base driver support
5	802.3ad	Yes	Yes	Yes	Yes (to retrieve speed)
6	balance-tlb	Yes	Yes	No	Yes (to retrieve speed)
7	balance-alb	Yes	Yes	No	Yes (to retrieve speed)

Note: When you create or remove a bond, SSH connections with Ethernet interfaces involved in that bond may be dropped. When the operation is complete, you must restore the SSH connections to continue administering the appliance.

Bonding operations

After the initial configuration is complete, you can perform operations on bonding. You can create, modify, display, and remove bonds.

See [“Creating a bond”](#) on page 81.

See [“Modifying a bond”](#) on page 86.

See [“Removing a bond”](#) on page 87.

You can also use RESTful APIs to create, modify and remove a bond on the cluster.


For more details about the APIs, see the *Veritas NetBackup FlexScale APIs* on SORT.

Creating a bond


The interfaces eth5 and eth7 can be bonded together after the initial configuration is complete.

To create a bond

- 1 Navigate to **Settings > Network**.
- 2 Click on the **Actions** button (vertical ellipsis) on the right side of **Data Network**. Select **Create Bond**.



3 The **Create Bond** form appears.



The image shows a 'Create Bond' form with a title bar containing a close button (X). The form contains the following text and input fields:

- Warning: "These values effect all data networks and it will temporarily take networks offline"
- Label: "Interface bonding type (eth5, eth7)"
- Input field: "balance-alb"
- Label: "MTU"
- Text: "Enter Value between 1500 to 10000"
- Input field: "1500"
- Buttons: "Cancel" and "Create"

By default, the bond mode here is `balance-alb` and MTU is 1500. You can use the drop-down box to change the bond mode.

Create Bond

These values effect all data networks. Making this change is desruptive. Temporarily take networks offline and may take upto 20 min.

Interface bonding type (eth5, eth7)

- balance-rr
- active-backup
- balance-xor >
- broadcast
- 802.3ad >
- balance-alb
- balance-tlb

een 1500 to 10000

Cancel

Create

See “About bonding Ethernet interfaces” on page 80.

Click **Create**. View the **Recent Activity** panel in the top navigation bar for the status of the task.

[View all activities](#)

You are viewing first 10 of 25 tasks

[View details](#)

	Status	Task name	Start time	End time
>	✓	Create bond corresponding to network devices	March 9 2022, 10:49 pm PST	March 9 2022, 11:21 pm PST
	✓	Create the bond corresponding to devices.	March 9 2022, 10:49 pm PST	March 9 2022, 11:17 pm PST
	✓	Set Maximum Transmission Unit of the network device.	March 9 2022, 11:18 pm PST	March 9 2022, 11:19 pm PST
>	✓	Configure data network	March 2 2022, 1:16 am PST	March 2 2022, 1:37 am PST
>	✓	Configure data network	March 2 2022, 12:46 am PST	March 2 2022, 1:05 am PST
	✓	Invoked Access CLISH REST API	March 1 2022, 10:09 pm PST	March 1 2022, 10:46 pm PST
	✓	Invoked Access CLISH REST API	March 1 2022, 10:09 pm PST	March 1 2022, 10:09 pm PST
	✓	Invoked Access CLISH REST API	March 1 2022, 10:08 pm PST	March 1 2022, 10:09 pm PST
	✓	Invoked Access CLISH REST API	March 1 2022, 10:08 pm PST	March 1 2022, 10:09 pm PST

Close

The display name of the data network changes to **Data-Network-bond0**.

Settings / Network

Data Network (bonded)
Data-Network-bond0
DNS
Management Network
IPMI Network
Appliance

Routing settings

At least one of IPv4 or IPv6 is required.

IPv4		IPv6	
Subnet Mask	Gateway	Prefix Length	Router
255.255.255.0	192.168.50.1		

Media servers

Media server assignment for each node in the cluster. At least one of IPv4 or IPv6 is required.

Nodes	FQDN	IPv4 Address	IPv6 Address
1	media3.vxindia.veritas.com	192.168.50.56	
2	media4.vxindia.veritas.com	192.168.50.57	
3	media1.vxindia.veritas.com	192.168.50.54	
4	media2.vxindia.veritas.com	192.168.50.55	

Modifying a bond

To modify a bond

- 1 Navigate to **Settings > Network**.
- 2 Click on the **Actions** button (vertical ellipsis) on the right side of **Data Network (bonded)**. Select **Modify Bond**.

Settings / Network

Data Network (bonded)

Data-Network-bond0

DNS

Management Network

IPMI Network

Appliance

Add Data Network

Data Network Bonding

Modify Bond

Remove Bond

Subnet Mask

Gateway

Prefix Length

255.255.255.0

192.168.50.1

IPv6

Media servers

Media server assignment for each node in the cluster. At least one of IPv4 or IPv6 is required.

Nodes	FQDN	IPv4 Address
1	media3.vxindia.veritas.com	192.168.50.56
2	media4.vxindia.veritas.com	192.168.50.57
3	media1.vxindia.veritas.com	192.168.50.54
4	media2.vxindia.veritas.com	192.168.50.55

3 The **Modify Bond** form appears.

You can modify the bond mode and the MTU fields. You can use the drop-down box to change the bond mode.

Modify Bond

These values effect all data networks. Making this change is disruptive. Temporarily take networks offline and may take upto 20 min.

Interface bonding type (eth5, eth7)

balance-rr

active-backup

balance-xor >

broadcast

802.3ad >

balance-alb

balance-tlb

1500 to 10000

Cancel Modify

See [“About bonding Ethernet interfaces”](#) on page 80.


Click **Modify**. View the **Recent Activity** panel in the top navigation bar for the status of the task

Removing a bond

You can remove existing bonds.

To remove a bond

- 1 Navigate to **Settings > Network**.
- 2 Click on the **Actions** button (vertical ellipsis) on the right side of **Data Network (bonded)**. Select **Remove Bond**.
- 3 A confirmation screen appears. Click **Remove** to delete the bond.



After you confirm, the bond removal is initiated. View the **Recent Activity** panel in the top navigation bar for the status of the task. On successful removal of the bond, the data network is moved to eth5. The data network name is changed to **Data-Network-eth5**.

Settings / Network

Data Network ⓘ
Data-Network-eth5
DNS
Management Network
IPMI Network
Appliance

Routing settings

At least one of IPv4 or IPv6 is required.


IPv4		IPv6	
Subnet Mask	Gateway	Prefix Length	Router
255.255.255.0	192.168.50.1		

Media servers

Media server assignment for each node in the cluster. At least one of IPv4 or IPv6 is required.

Nodes	FQDN	IPv4 Address
1	media3.vxindia.veritas.com	192.168.50.56
2	media4.vxindia.veritas.com	192.168.50.57
3	media1.vxindia.veritas.com	192.168.50.54
4	media2.vxindia.veritas.com	192.168.50.55

During bond removal, if a secondary network is present on eth7, it is moved to eth5.



Configuring NetBackup Flex Scale in a non-DNS environment

Starting with NetBackup Flex Scale 3.0, enabling the DNS server is optional. Cluster nodes can be configured without the DNS server IP. If you do not want to specify a DNS server, you can specify the IP address, or the FQDN, or the short host name details and these details are communicated to the NetBackup services and all the cluster nodes using multiple set of APIs.

You can use the **Network > Custom Hosts** menu to add IP/FQDN details for clients and IPs for which FQDN mapping is not present in the DNS server. This information

can be added to the cluster nodes or NetBackup services or both. A hosts file is used to map IP addresses and its FQDN so that the system can resolve addresses quickly without querying the DNS. You can select cluster nodes or NetBackup services as an option to add IP/FQDN information.

You can perform the following operations using the NetBackup Flex Scale GUI:

- Add IP/FQDN to cluster nodes. Go to **Network > Custom Hosts**. Click **Add** to add additional IP/FQDN entries to the configuration. In the **Add custom hosts entries** form, add the entry. In the **Apply to** field, select **Cluster nodes**.

Add custom hosts entries

Add additional entries to prioritize over the DNS configuration.

Simple Advanced

IP Address (IPv4 or IPv6) * Fully qualified domain name (FQDN) * Apply to *

157.176.101.44 pub101-n6pip.mvlandom.com Cluster nodes

Cancel Save

You can use the **Advanced** tab to add multiple IP/FQDNs entries at the same time.

- Add IP/FQDN on NetBackup services containers. Go to **Network > Custom Hosts**. Click **Add** to add additional IP/FQDN entries to the configuration. In the **Add custom hosts entries** form, add the entry. In the **Apply to** field, select **NetBackup services**.

The screenshot shows the 'Add custom hosts entries' dialog box with the 'Simple' tab selected. The dialog has a title bar with a close button (X). Below the title bar, it says 'Add additional entries to prioritize over the DNS configuration.' There are two tabs: 'Simple' (selected) and 'Advanced'. The form contains three fields: 'IP Address (IPv4 or IPv6) *' with the value '175.176.101.46', 'Fully qualified domain name (FQDN) *' with the value 'pub101-n8pip.mvlandom.com', and 'Apply to *' with a dropdown menu showing 'NetBackup services'. To the right of the dropdown is a plus icon (+). At the bottom right, there are 'Cancel' and 'Save' buttons.


You can use the **Advanced** tab to add multiple IP/FQDNs entries at the same time.

- Add IP/FQDN on cluster nodes and NetBackup services containers. Go to **Network > Custom Hosts**. Click **Add** to add additional IP/FQDN entries to the configuration. In the **Add custom hosts entries** form, add the entry. In the **Apply to** field, select both **Cluster nodes** and **NetBackup services**.


This screenshot shows the same 'Add custom hosts entries' dialog box, but with the 'Apply to *' dropdown menu open. The dropdown menu shows two options: 'Cluster nodes' and 'NetBackup services', both of which are checked with blue checkmarks. The rest of the dialog, including the 'Simple' tab, the input fields, and the 'Cancel'/'Save' buttons, remains the same as in the previous screenshot.

You can use the **Advanced** tab to add multiple IP/FQDNs entries at the same time.


- Delete IP/FQDN information from cluster nodes. Click on the **Actions** button (vertical ellipsis) on the right side of the IP address and select **Delete**.



- Delete IP/FQDN information from NetBackup services. Click on the **Actions** button (vertical ellipsis) on the right side of the IP address and select **Delete**.



- Delete IP/FQDN information from cluster nodes and NetBackup services. Click on the **Actions** button (vertical ellipsis) on the right side of the IP address and select **Delete**.



Data network configurations

All the NetBackup operational data traffic, including communications with external hosts and services, is routed on this network. A data network is required to set up the cluster.

Note: Ensure that the IP addresses and FQDN that you specify are added to the DNS server that you specify here and are resolvable on the network.

Note: The `fd00:200:120` network is reserved and used internally by NetBackup Flex Scale, and it should not be used anywhere.

You can configure a cluster without a DNS server.

See [“Configuring NetBackup Flex Scale in a non-DNS environment”](#) on page 89.

The following types of data network configurations can be done on NetBackup Flex Scale

- Plain device (eth5)
See [“Network configuration on plain device \(eth5\)”](#) on page 93.
- VLAN on eth5
See [“Network configuration on VLAN \(eth5\)”](#) on page 98.
- Bonded interfaces (bond0 using eth5 and eth7)
See [“Network configuration on bonded interfaces \(bond0 on eth5 and eth7\)”](#) on page 99.
- VLAN on bond (VLAN on bond0)
See [“VLAN on bond of eth5 and eth7 \(bond0\)”](#) on page 101.


To list data networks

- ◆ Navigate to **Settings > Network > Data Network**.

Note: If you have configured disaster recovery, See [“Support for multiple VLAN when disaster recovery is configured”](#) on page 101.

Network configuration on plain device (eth5)

If configuration is done on plain device (eth5), the data network which has been created has the display name as **Data-Network-eth5**.



The following operations are supported:

- Adding a new data network
See [“Adding a data network”](#) on page 99.
- Modifying the secondary data network
See [“Modifying a data network”](#) on page 97.
- Deleting the secondary data network
See [“Deleting a data network”](#) on page 97.

Adding a data network

You can add a new data network.

To add a new data network

- 1 Navigate to **Settings > Network > Data Network**.
- 2 Click on the **Actions** button (vertical ellipsis) on the right side of **Data-Network**. Select **Add Data Network**.

3 The **Add data network** form appears. Provide the necessary details.

You can choose to have IPv4 or IPv6 addresses in your data network. But mixed mode of IPv4 and IPv6 configuration is not supported.

Note: You can configure the data network with the IP protocol version used for the management or the data network during initial cluster configuration. You also have an option to select either automatic or custom mode for the new data network to be created.

The NetBackup primary server, media server and the storage server names can contain a maximum of 64 characters, including the domain name.

Interface	Choose eth5 or eth7.
Routing settings	Enter the subnet mask and gateway for the network.
VLAN ID	Enter the VLAN ID.
Media servers	Enter the FQDN and IP address for each media server as per the input method selected.
Storage servers	Enter the FQDN and IP addresses for each storage server as per the input method selected. Enter the FQDN and IP address of primary server. Note: If you deploy a cluster with only media servers, you do not have to enter these details.

Note: If the cluster is configured in a non-DNS environment, the **Automatic** option is disabled and you have to specify the FQDNs and the IP addresses.

4 Click **Add Data Network**. The process for adding the network is initiated.

- View the **Recent Activity** panel in the top navigation bar for the status of the task. You can monitor the status of the operation by clicking **View details** on the taskbar.

View all activities ? ✕

You are viewing first 10 of 24 tasks ☐ View details

Status	Task name	Start time	End time
✓	Configure data network	March 2, 2022, 1:16 am PST	March 2, 2022, 1:37 am PST
✓	Validating input and prerequisite.	March 2, 2022, 1:16 am PST	March 2, 2022, 1:17 am PST
✓	Configuring VLAN 104 on device eth5.	March 2, 2022, 1:17 am PST	March 2, 2022, 1:19 am PST
✓	Adding FQDN entries.	March 2, 2022, 1:19 am PST	March 2, 2022, 1:20 am PST
✓	Adding physical IPs into cluster.	March 2, 2022, 1:20 am PST	March 2, 2022, 1:24 am PST
✓	Adding virtual IPs into cluster.	March 2, 2022, 1:24 am PST	March 2, 2022, 1:31 am PST
✓	Adding gateway into cluster.	March 2, 2022, 1:31 am PST	March 2, 2022, 1:34 am PST
✓	Configuring network for netbackup.	March 2, 2022, 1:34 am PST	March 2, 2022, 1:35 am PST
✓	Assigning IP to master server.	March 2, 2022, 1:34 am PST	March 2, 2022, 1:34 am PST

Close

- Once the data network is added successfully, the newly added data network appears under **Data Network**.

Veritas NetBackup™ Flex Scale - xylan 🏠 🔍 👤 ? ⌵

Settings / Network

Data Network	
Data Network-eth7-103	
Data Network-eth5-104	
Data Network-eth5-101	
DNS	
Management Network	
IPMI Network	
Appliance	
Custom Hosts	

Edit Delete

Routing settings

IPv4

Subnet Mask	Gateway
255.255.255.0	175.176.104.55

Virtual Local Area Network (VLAN)

VLAN ID

104

Media servers

Nodes	FQDN	IPv4 Address
1	pub104-n1tp.mv1andom.com	175.176.104.39
2	pub104-n2tp.mv1andom.com	175.176.104.40
3	pub104-n3tp.mv1andom.com	175.176.104.41
4	pub104-n4tp.mv1andom.com	175.176.104.42

Storage Server

Nodes	FQDN	IPv4 Address
1	pub104-n1vp.mv1andom.com	175.176.104.47
2	pub104-n2vp.mv1andom.com	175.176.104.48
3	pub104-n3vp.mv1andom.com	175.176.104.49
4	pub104-n4vp.mv1andom.com	175.176.104.50

NetBackup master server

FQDN	IPv4 Address
pub104-mastervip.mv1andom.com	175.176.104.55

Activate Windows
Go to Settings to activate Windows.

Modifying a data network

The attributes of the primary data network cannot be modified.


You can modify the following attributes of the secondary data network:

- VLAN ID
- IP address of the primary, media, and storage servers

Note: If you deploy a cluster with only media servers, the primary server details cannot be modified.

To modify a data network:

- Navigate to **Settings > Network > Data Network**.
- Go to the data network that you want to modify. Click **Edit**. View the **Recent Activity** panel in the top navigation bar for the status of the task.




Deleting a data network

Primary data network cannot be deleted. You can delete secondary data network(s).


To delete a data network:

- Navigate to **Settings > Network > Data Network**.
- Go to the data network that you want to delete. Click **Delete**. View the **Recent Activity** panel in the top navigation bar for the status of the task.



Network configuration on VLAN (eth5)

If configuration is done on VLAN on eth5, the data network which has been created has the display name as **Data-Network--eth5-<vlan-id>**.



Note: The primary data network cannot be modified or deleted.

The following operations are supported on the secondary data network:

- Adding a new data network
See [“Adding a data network”](#) on page 94.

- Modifying a data network
See [“Modifying a data network”](#) on page 97.
- Deleting a data network
See [“Deleting a data network”](#) on page 97.

Network configuration on bonded interfaces (bond0 on eth5 and eth7)

If configuration is done on bond created on eth5 and eth7, the data network which has been created has the display name as **Data-Network-bond0**.

Nodes	FQDN	IPv4 Address
1	pub101-n2pjp.mvlandom.com	175.176.101.40
2	pub101-n3pjp.mvlandom.com	175.176.101.41
3	pub101-n1pjp.mvlandom.com	175.176.101.39
4	pub101-n4pjp.mvlandom.com	175.176.101.42

Nodes	FQDN	IPv4 Address
1	pub101-n1vip.mvlandom.com	175.176.101.47
2	pub101-n2vip.mvlandom.com	175.176.101.48
3	pub101-n4vip.mvlandom.com	175.176.101.50
4	pub101-n3vip.mvlandom.com	175.176.101.49

Note: The primary data network cannot be modified or deleted.

The following operations are supported on the secondary data network:

- Adding a new data network
See [“Adding a data network”](#) on page 99.
- Modifying a data network
See [“Modifying a data network”](#) on page 97.
- Deleting a data network
See [“Deleting a data network”](#) on page 97.

Adding a data network

You can add a new data network.

To add a new data network

- 1 Navigate to **Settings > Network > Data Network (bonded)**.
- 2 Click on the **Actions** button (vertical ellipsis) on the right side of **Data-Network**. Select **Add Data Network**.
- 3 The **Add data network** form appears. Provide the necessary details.

The NetBackup primary server, media server and the storage server names can contain a maximum of 64 characters, including the domain name.

Interface	Choose bond0
Routing settings	Enter the subnet mask and gateway for the network
VLAN ID	Enter the VLAN ID
Media servers	Enter the FQDN and IP address for each media server as per the input method selected
Storage servers	Enter the FQDN and IP addresses for each storage server as per the input method selected
	Enter the FQDN and IP address of primary server
	Note: If you deploy a cluster with only media servers, you do not have to enter these details.

Note: If the cluster is configured in a non-DNS environment, the **Automatic** option is disabled and you have to specify the FQDNs and the IP addresses.

- 4 Click **Add Data Network**. The process for adding the network is initiated.
- 5 View the **Recent Activity** panel in the top navigation bar for the status of the task. You can monitor the status of the operation by clicking **View details** on the taskbar.
- 6 Once the data network is added successfully, the newly added data network appears under **Data Network (bonded)**.

VLAN on bond of eth5 and eth7 (bond0)

If configuration is done on bond created on eth5 and eth7, the data network which has been created has the display name as **Data-Network-bond0-*<vlan-id>***.

The screenshot shows the 'Settings / Network' configuration page. The left sidebar lists 'Data Network (bonded)', 'DNS', 'Management Network', 'IPMI Network', and 'Appliance'. The main panel is titled 'Data-Network-bond0-874' and includes an 'Edit' button. It contains sections for 'Routing settings', 'Virtual Local Area Network (VLAN)', and 'Media servers'.

Routing settings

At least one of IPv4 or IPv6 is required.

IPv4		IPv6	
Subnet Mask	Gateway	Prefix Length	Router
255.255.255.0	192.168.51.1		

Virtual Local Area Network (VLAN)

VLAN ID
874

Media servers

Media server assignment for each node in the cluster. At least one of IPv4 or IPv6 is required.

Nodes	FQDN	IPv4 Address	IPv6 Address
1	media512.vxindia.veritas.com	192.168.51.55	
2	media511.vxindia.veritas.com	192.168.51.54	
3	media514.vxindia.veritas.com	192.168.51.57	
4	media513.vxindia.veritas.com	192.168.51.56	

Note: The primary data network cannot be modified or deleted.

The following operations are supported on the secondary data network:

- Adding a new data network
See [“Adding a data network”](#) on page 94.
- Modifying a data network
See [“Modifying a data network”](#) on page 97.
- Deleting a data network
See [“Deleting a data network”](#) on page 97.

Support for multiple VLAN when disaster recovery is configured

If disaster recovery is configured in a single domain configuration, the primary service is operational from one site at any point of time but the primary service can manage the media services, storage servers and storage units on both sites.

- You cannot perform add data network operations on both the sites (primary and secondary) at the same time.
- Add data network operation should be performed first on the primary site and then on the secondary site for the same data network.

NetBackup Flex Scale infrastructure monitoring

This chapter includes the following topics:

- [About alert management](#)
- [About event notification](#)
- [About AutoSupport and Call Home](#)
- [Monitoring hardware components](#)
- [Monitoring usage and licensed capacity using Veritas NetInsights Console](#)

About alert management

NetBackup Flex Scale displays alerts for current problems or critical conditions that take place in the system. You can use **Alert management** to view, enable, or disable the alerts as required. The alerts are displayed or shared in the following ways:

- Email
- SNMP

Alerts and events are visible on the NetBackup Flex Scale management UI dashboard.

See [“Viewing information about alerts”](#) on page 103.

See [“Managing alerts”](#) on page 103.

Viewing information about alerts

From the **Settings > Alert management** page, you can view all the alerts and manage them. You can view the following details for the alerts:

- Name
- Object type
- Severity level
- Date and time at which the alert is generated

[Table 5-1](#) describes the valid NetBackup Flex Scale severity levels.

Table 5-1 Severity levels

Valid value	Description
crit	Indicates a critical condition
err	Indicates an error condition
info	Indicates an informational message
warn	Indicates a warning condition

The bell icon in the top navigation bar shows the status of current alerts and alerts that were recently completed. For more details, select **View all alerts**.

Managing alerts

You can disable or enable the alerts.

To manage the alerts

- 1** Go to **Settings > Alert management**, and then click **Manage Alerts**.
The **Manage Alerts** dialog box is displayed.
- 2** Select the check box for the alert that you want to disable, and click **Next**.
- 3** The task is initiated to update the alert suppression. Click **Finish** to complete the task.
- 4** View the **Recent Activity** panel in the top navigation bar for the status of the task.

Note: To enable the alerts, you can clear the check boxes for the alerts.

Note: Hardware alerts cannot be suppressed.

About event notification

NetBackup Flex Scale shows the status of current events and events that were recently completed

From the **Settings > Events** page, you can view all the events. You can view the following details for the events:

- Date and time
- Severity level
- Object type
- Source
- Message

You can filter the type of event information that you want to see based on severity levels.

[Table 5-2](#) describes the valid NetBackup Flex Scale severity levels.

Table 5-2 Severity levels

Valid value	Description
critical	Indicates a critical condition
error	Indicates an error condition
information	Indicates an informational message
warning	Indicates a warning condition

The **Events** icon in the top navigation bar shows the status of current events and events that were recently completed. For more details, select **View all events**.

See [“Purging events”](#) on page 104.

Purging events

You can purge event records greater than a specified number of days.

To purge events

- 1 Prerequisites:
You must have some events to purge.
- 2 In **Settings > Events**, select **Purge**.
- 3 Indicate the event records that you want to purge.
- 4 View the **Recent Activity** panel for the status of the task.

About AutoSupport and Call Home

Veritas AutoSupport is a framework that provides improved support experience through proactive monitoring of the appliance, automated error reporting, and support case creation. AutoSupport uses the Call Home service to automatically upload diagnostic and heartbeat data over SSL-encrypted channels to a Veritas secure operations center for further processing. The AutoSupport framework analyzes the Call Home data and correlates it with other site configuration data held by Veritas to provide proactive customer support and incident response for hardware failures.

Starting with version 2.1, Call Home uses a highly available centralized AutoSupport client service, which is used for communicating with the AutoSupport server. Each node no longer communicates with the AutoSupport server; instead Call Home is at cluster level.

- To configure Call Home, See [“Configuring Call Home settings”](#) on page 109.
- To receive email notifications for the generated alerts, See [“Setting up email alerts”](#) on page 105.
- To receive SNMP notifications, See [“Setting up SNMP alerts”](#) on page 108.

You can enable email notifications for the alerts that the appliance generates, monitor the appliance using the SNMP alerts, and upload the monitored hardware and software details to the Veritas AutoSupport server.

You can configure these settings from the **Settings > AutoSupport** option. These settings are configured at cluster level.

Setting up email alerts

The appliance can be configured to send email notifications when hardware and software components fail or encounter errors. The email notifications are sent using the Simple Mail Transfer protocol (SMTP).

To set up email alerts:

- 1 Use any one of the following options to log in using the user account that you created:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.
- 2 In the left navigation click **Settings**, and then click **AutoSupport**.

3 Click **Email service settings** and specify the following details:

Parameter	Description
Notification interval	Time interval in minutes between subsequent notifications. The time interval must be greater than zero and a multiple of 15.
SMTP server	Host name or the IP address of the SMTP server that is used to send email notifications for the alerts generated by the appliance.
Server port	Port number for the SMTP server. The default port is 25.
Software administrator email	Email address of the administrators who are the recipients of the software-related email alerts. Use a comma to separate multiple email addresses.
Hardware administrator email	Email address of the administrators who are the recipients of the hardware-related email alerts. Use a comma to separate multiple email addresses.
Sender email	Source email address that is used to send email alerts.
SMTP account	User name to access the SMTP account.
Password	Password for the user name if authentication is required to access the SMTP account.
Encryption Enabled	Turn on to use a secure connection and to encrypt communication with the SMTP server. By default, the communication is not encrypted.

4 Click **Save**.

A notification is displayed. Click **View details** to view the status and progress of each of the tasks.

5 Optionally, click **Test configuration** to verify the settings. A test email is sent to the configured email addresses. If the test email is not received, contact your system administrator for assistance.

Setting up SNMP alerts

You can configure the appliance to generate and send Simple Network Management Protocol (SNMP) traps to your SNMP server to monitor the hardware. The appliance uses the SNMPv2-SMI application protocol.

To configure the SNMP settings:

- 1 Use any one of the following options to log in using the user account that you created:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.
- 2 In the left navigation click **Settings**, and then click **AutoSupport**.

- 3 On the **AutoSupport** page, click **SNMP settings** and specify the following details:

Parameter	Description
Enable SNMP	Turn on to enable the SNMP alerts.
SNMP server	Host name or the IP address of the SNMP server. The IP address can be an IPv4 or an IPv6 address. Alert notifications that are generated by the appliance are sent to this server.
SNMP port	Port number of the SNMP server. The default port is 161. Note: Your firewall must allow access from the appliance to the SNMP server through the configured port.
Community	Community to which the alerts are sent. The default value is public .

- 4 Click **Save**.

A notification is displayed. To view the status and progress of each of the tasks, click **View details**.

Configuring Call Home settings

If Call Home is configured, the appliance uploads hardware and software information to the Veritas AutoSupport server and sends email alerts to administrators when hardware errors are detected. Veritas Support uses this information to troubleshoot and resolve the issues. The appliance uses the HTTPS protocol and uses port 443 to connect to the AutoSupport server. You can configure the email addresses that you want to use for failure notifications. See [“Setting up email alerts”](#) on page 105.

To configure Call Home settings:

- 1 Use any one of the following options to log in using the user account that you created:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
<https://ManagementServerIPorFQDN/webui> where *ManagementServerIPorFQDN* is the public IP address or the FQDN that

you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.

- Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
<https://ManagementServerIPorFQDN:14161> where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.

- 2 In the left navigation click **Settings**, and then click **AutoSupport**.
- 3 Click **Call Home and proxy settings** and specify the following details:

Parameter	Description
Enable Call Home transmission	Turn on to upload the appliance health information to the Veritas AutoSupport server.
Enable the proxy server	If the appliance connects to the AutoSupport server through a proxy server, turn on to configure a proxy server.
Enable proxy tunneling	If the proxy server supports SSL tunneling, turn on to enable SSL tunneling.
Proxy server	Name of the proxy server. (Required if you enable the proxy server)
Proxy port	Proxy server port. (Required if you enable the proxy server)
Proxy username	User name to log in to the proxy server.
Proxy password	Password to authenticate the user name that is used to log in to the proxy server.

- 4 Click **Save**.
A notification is displayed. To view the status and progress of each of the tasks, click **View details**.
- 5 Optionally, click **Test configuration** to verify that the configured settings are valid and the appliance can communicate with the AutoSupport server.

Monitoring hardware components

After you configure the cluster, you can monitor the hardware components using the NetBackup Flex Scale UI or the Appliance Shell Menu. If a problem is detected, the following notification mechanisms can be used to view and report the issue:

- View alerts on the **Dashboard** or the **Settings > Alerts and notifications** option of the web interface.
See [“About the NetBackup Flex Scale management console”](#) on page 13.
See [“About alert management”](#) on page 102.
- Send email alerts.
See [“Setting up email alerts”](#) on page 105.
- Send SNMP alerts.
See [“Setting up SNMP alerts”](#) on page 108.
- Share hardware monitoring information for analysis and troubleshooting with Veritas Support using Call Home.
See [“Configuring Call Home settings”](#) on page 109.

Using the NetBackup Flex Scale web interface

The **Hardware** tab on the **Monitor > Infrastructure** page shows the status of the hardware components for each of the nodes. If a problem occurs and Call Home is enabled, the diagnostic data is automatically sent to Veritas Support and a support case is automatically created. This option helps to provide proactive service and leads to a faster resolution of any hardware issues.

You can monitor the following hardware components in the web interface:

- CPU
- Fan
- Hard disk
- Network card
- PCI
- Power supply
- RAID

Nodes	Disks	Hardware
Nodes		
vrts-01.engba.veritas.com	vrts-02.engba.veritas.com	vrts-03.engba.veritas.com vrts-04.engba.veritas.com
vrts-01.engba.veritas.com CPU		
CPU	ID	State Status
Fan	1	OK OK
Hard Disk	2	OK OK
Network Card	Showing 1-2 of 2 (0 selected)	
PCI		
Power Supply		
RAID		

If a hardware component fails or reports an error, an alert icon is displayed next to the component and the node it belongs to.

Nodes	Disks	Hardware
Nodes		
hpehw-01.engba.veritas.com	hpehw-02.engba.veritas.com	hpehw-03.engba.veritas.com  hpehw-04.engba.veritas.com
hpehw-04.engba.veritas.com Fan		
CPU	ID	Speed State Status
 Fan	1	19 percent OK OK
	 2	- Failed Present

Using the shell menu

You can also use the Veritas Appliance shell menu to check the health and status of the hardware components. After logging in to the shell menu, you can use the following commands to monitor the status:

Command	Description
show hardware-errors	View hardware component faults and errors
show hardware-health node [component=]	View the status of a specific hardware component. The <code>Item</code> parameter specifies the component for which the data is queried. The following options are available. The default option is <code>All</code>. Note: Some of the options might vary based on the hardware platform. ■ All ■ Product ■ Fan ■ CPU ■ RAID ■ Power ■ PCI ■ Network ■ DIMM ■ SSD ■ Firmware ■ Array (Only for Hewlett Packard Enterprise hardware) ■ PhysicalDrive ■ LogicalDrive

Monitoring usage and licensed capacity using Veritas NetInsights Console

NetBackup Flex Scale integrates with Veritas NetInsights Console to manage your usage and license entitlements. You must register your appliance by signing in to the NetInsights portal (<https://netInsights.veritas.com>) with your Veritas Account Manager credentials.

Veritas Usage Insights, which is a part of Veritas NetInsights Console provides near real-time reporting of usage against licensed capacity, shows usage trends, and helps predict your future capacity requirements. For more details, see the *Veritas Usage Insights Getting Started Guide*.

Resiliency in NetBackup Flex Scale

This chapter includes the following topics:

- [Erasure coding in NetBackup Flex Scale](#)
- [Handling split-brain scenario in NetBackup Flex Scale](#)
- [High availability of the NetBackup primary service](#)
- [High availability of NetBackup services](#)
- [NetBackup catalog protection](#)
- [NetBackup primary service catalog protection using checkpoints](#)

Erasure coding in NetBackup Flex Scale

Erasure coding (EC) offers a more robust solution in redundancy and fault tolerance for critical storage archives. As storage systems expand and become more complex, traditional data protection mechanisms prove to be inadequate against failures. In erasure coding, data is broken into fragments, expanded, and encoded with redundant data pieces and stored across different locations or storage media. When one or more disks fail, the data on failed disks is reconstructed using the parity information in the encoded disks and data in the surviving disks. Multiple nodes can read or write concurrently to an erasure coded volume without data corruption.

Erasure coding allows setting ratios of original data and coding data. With a ratio of m parts of original data to n parts of coding data, the code can tolerate the loss of any n parts and regenerate the original m parts. In NetBackup Flex Scale, the erasure coding ratio is 8:4. This means that every 8 parts of data are enriched with 4 parts of coding data and the data is spread across 12 ($8 + 4$) disks.

For example, when the deduplication engine receives 2MB of data from a backup image, the 2M stripe is ready to be erasure-coded. The deduplication engine uses the erasure-coded file system underneath to store this data. The data is divided into eight 256K data fragments. Four 256K parity fragments are created for resiliency. These fragments are then sent to the underlying storage which is distributed across nodes and disks.

Benefits of erasure coding:

- Improved performance
NetBackup Flex Scale uses the I/O resources of the entire disk pool to ingest backup data, and hence, eliminates bottlenecks by limiting the data protection jobs to a single node's disks. Also, during data restore or when a failed disk is rebuilt, many disks contribute to the workload simultaneously by reading erasure coded data fragments.
- High level of data protection
Erasure coding's distributed nature also provides a high level of data protection. Deployments with less than 6 nodes can protect against failure of a node and a disk or any two disks. Deployments with six or more nodes can protect against simultaneous failure of two nodes or any four disks (HDDs) in the disk pool.
- High usable capacity
Advantages of using the erasure-coded data layout include a usable capacity of 67% for large capacity HDD on top of high performance and resiliency.

NetBackup Flex Scale requires a four-node base configuration for a deployment. When a new node is added, NetBackup Flex Scale ensures data is balanced across the nodes by intelligently moving the smallest possible amount of data to maintain or increase its existing resiliency characteristics.

The Reed-Solomon algorithm is used to build the erasure coding solution.

Handling split-brain scenario in NetBackup Flex Scale

A split-brain occurs when the cluster membership view differs among the cluster nodes, increasing the chance of data corruption. With majority-based I/O fencing, the potential for data corruption is eliminated as it provides a reliable arbitration mechanism which does not require any extra hardware. In a split-brain scenario, arbitration is done based on `majority` number of nodes among the sub-clusters. The node with the lowest node ID in the cluster is called the leader node and it a role in case of a tie.

Deciding cluster majority for majority-based I/O fencing mechanism:

- If N is defined as the total number of nodes in the cluster, then majority is equal to $N/2 + 1$.
- If there are even number of cluster nodes and both the sub-clusters have $N/2$ number of nodes, the partition with the leader node is treated as majority and that partition survives.

How majority-based I/O fencing works

An algorithm is used to decide the winner sub-cluster in the following way:

- The node with the lowest node ID in the current cluster membership is designated as the leader node in the fencing race.
- When a network partition occurs, each racer sub-cluster computes the number of nodes in its partition and compares it with the majority value.
- If a racer finds that its partition does not have majority, it sends a LOST_RACE message to all the nodes in its partition including itself and all the nodes panic.
- If the racer finds that it does have majority, it sends a WON_RACE message to all the nodes. Thus, the partition with majority nodes survives.

High availability of the NetBackup primary service

The NetBackup primary service manages backups, archives, and restores. The primary service is responsible for media and device selection for NetBackup. Typically, the primary service contains the NetBackup catalog. The catalog contains the internal databases that contain information about NetBackup backups and configuration.


Note: This section is not applicable if you deploy a cluster with only media servers.

NetBackup Flex Scale is a clustered solution and hence, it ensures availability of the NetBackup primary service. The primary service is guarded against failures that include but are not limited to the following scenarios:


- Failure of the NIC where the NetBackup primary service IP address is hosted
- Failure of the node where the NetBackup primary service is running
- Failure of the NetBackup primary service process
- Failure of the docker container
- Failure of the docker service

The file system that is used as the NetBackup catalog file system is a three-way mirrored file system.

The **Events** icon in the top navigation bar shows the status of the current events. The status of the health of the primary service is displayed.



The bell icon in the top navigation bar shows the status of current alerts and alerts that were recently completed. An AutoSupport alert is raised if the primary service is offline.



If the primary service goes offline, it is failed over and comes online on another cluster node. For details on the behavior of the NetBackup primary service in clustered environments and handling the jobs, see the *General notes on clustered primary service administration* chapter in the *Veritas NetBackup™ Clustered primary service Administrator's Guide* on SORT.

For more details on the file system layout and disk selection, See [“About NetBackup Flex Scale storage”](#) on page 47.

High availability of NetBackup services

NetBackup Flex Scale uses docker containers to run different roles required for serving backup/restore jobs.

The containers are guarded against the following failures:

- Failure of the NIC hosting the public and private IPs that are required for communicating with other NetBackup services.
- Failure of the file systems which are used to store the deduplicated data and catalog file system.
- Failure of the node on which these containers are running.

If any of the above resources are found to be faulted, NetBackup Flex Scale fails over necessary containers to suitable nodes in the cluster.

NetBackup catalog protection

NetBackup Flex Scale automatically creates a catalog backup policy to protect the NetBackup catalog from any type of corruption. Catalog backups are created using the default retention levels. Veritas recommends that you review the default retention level and settings to ensure that it meets your data retention guidelines.

For additional information on configuring catalog backup, see the *Configuring backups* and *Protecting the NetBackup catalog* sections in the *NetBackup Administrator's Guide, Volume I*.

NetBackup primary service catalog protection using checkpoints

In NetBackup Flex Scale, you can protect your primary service from software failures or from being corrupted using checkpoints. The checkpoints are created for the primary service catalog file system every 2 hours according to a schedule. A maximum of 36 checkpoints can be created. Once the total number of checkpoints exceed 36, the oldest checkpoint is deleted and a new checkpoint is created.

For the restore operation, NetBackup Flex Scale uses the checkpoint of checkpoint which is provided by the underlying file system. Checkpoint of checkpoint is a new checkpoint that is created by mounting the checkpoint that is selected for validation. After the new checkpoint is created, it is used for validation and restore.

All the checkpoints consume storage from the same volume set. As the checkpoints are copy-on-write, only modified data gets pushed to the checkpoints. But depending on data change rate, the checkpoints can consume considerable storage. If the

primary file set cannot allocate storage during the write or file creation operation, instead of returning an ENOSPC error to the user, the oldest checkpoint is automatically deleted to make space. An email notification is sent if checkpoints are deleted to make free space.

All other operations such as adding a node and replacing a node are blocked when catalog restore is in progress.

Note: This section is not applicable if you deploy a cluster with only media servers.

See [“Performing a recovery of the catalog file system using GUI”](#) on page 119.

See [“Performing a recovery of the catalog file system using REST APIs”](#) on page 122.

Note: When the primary service catalog file system is recovered using an old checkpoint, the primary service goes back to the point in time of the checkpoint. The images created after that point in time will remain in the system and cannot be accessed unless they are reimported.

For more details on import, see the *Importing backup images, Phase I* and *Importing backup images, Phase II* sections in the *NetBackup Administrator's Guide, Volume I*.

Performing a recovery of the catalog file system using GUI

You can use the NetBackup Flex Scale GUI for the recovery of primary service catalog file system from the checkpoints.


The following operations can be performed using the GUI:

- List the valid checkpoints for restore
- Validate a checkpoint
- Restore a checkpoint

You can list all the valid checkpoints for the primary service.

To list all the valid checkpoints

- 1 Go to **Settings > NetBackup catalog management**.
- 2 A screen with a calendar view gets displayed. Select a date in the calendar. The list of valid checkpoints for that date is displayed in tabular form.



Note: The green dot below a date in the calendar indicates that there are valid checkpoints available for that date.

You can validate a specific checkpoint. For validation operation to be completed, the primary server services must be running in the cluster.


To validate a checkpoint

- 1 From the list of checkpoints, click the kebab menu for a particular checkpoint. Click **Validate** to validate the checkpoint.
- 2 A task is initiated in the GUI which can be used to track the status of the validate operation. Once the validate operation is completed, the status of the task changes accordingly.


You can restore a checkpoint. For the restore operation to be completed, the primary server services must be in the ONLINE state in the cluster.

To restore a checkpoint


- 1 From the list of checkpoints, click on the kebab menu for a particular checkpoint. Click **Restore** to restore the checkpoint.



- 2 A pop up window appears asking for confirmation. Click **Restore**.



- 3 If a validation checkpoint exists for the selected checkpoint, the restore operation is performed immediately. If a validation checkpoint does not exist, the restore operation first validates the checkpoint and creates a validation checkpoint. Then, it performs a restore from the validation checkpoint.
- 4 The progress of the restore operation is displayed on the progress card that gets created with the restore operation. A task is also initiated in the GUI which can be used to track the status of the restore operation. Once the restore operation is completed, the status of the task changes accordingly.



Performing a recovery of the catalog file system using REST APIs

You can use REST APIs for the recovery of primary service catalog file system from the checkpoints. The REST API calls should be made in the following order:

1. Get the list of checkpoints for the primary service catalog file system. This API returns a list of all the available checkpoints for the primary catalog file system. The checkpoints can be used to recover the data.

```
GET /api/appliance/v1.0/netbackup/checkpoints
```

```
URI : /api/appliance/v1.0/netbackup/checkpoints
```

```
Type : GET
```

```
Response Body:
```

```
{
  "links": {
    "self": {
      "href": "/api/appliance/v1.0/netbackup/checkpoints"
    }
  },
}
```

```

"data": [
  {
    "type": "netbackup",
    "id": "1",
    "links": {
      "self": {
        "href": "/api/appliance/v1.0/netbackup/checkpoints/
        checkpoint1"
      }
    }
  }
]
}

```

The checkpoint name which is returned by this API should be passed as input to the API which is called to restore the catalog.

2. Validate a checkpoint. This API validates the checkpoint by creating a new checkpoint from the selected checkpoint and runs the NetBackup database validation on the newly created checkpoint. After validation completes the new checkpoint can be used for restore.

```

POST /api/appliance/v1.0/netbackup/checkpoints/restore-catalog/
{checkpointName}
URI: /api/appliance/v1.0/netbackup/checkpoints/restore-catalog/{
checkpointName}
Type : POST
Input Parameter for Request : Checkpoint name
Response:
{
  "taskId": "string",
  "message": "string"
}

```

This is an asynchronous API and it returns a task ID when API execution is successful. You can find the execution status and details by providing this taskID as input in the GET {taskId} API. The task is also visible in the **Recent activity** icon in the top navigation bar in the GUI.

3. Restore the data to the primary service catalog file system. This API restores the data (from the checkpoint that is created in the previous step) to the primary service catalog file system.

```

POST /api/appliance/v1.0/netbackup/checkpoints/sync-catalog
URI : /api/appliance/v1.0/netbackup/checkpoints/sync-catalog
Type : POST

```

Input Parameters : None

Response :

```
{  
    "taskId": "string",  
    "message": "string"  
}
```

This is an asynchronous API and it returns a task ID when API execution is successful. You can find the execution status and details by providing this taskID as input in the GET {taskId} API. The task is also visible in the **Recent activity** icon in the top navigation bar in the GUI.

EMS server configuration

This chapter includes the following topics:

- [Configuring an external BYOS media server](#)
- [Configuring an external NBA media server](#)

Configuring an external BYOS media server

You can configure an external BYOS media server (EMS) with NetBackup Flex Scale.

To configure an external BYOS media server with NetBackup Flex Scale 3.0

- 1 Deploy mapping BYO on the external media server.
 - Enter the NetBackup Flex Scale primary server details..
 - Enter the token that you get from the primary server.
- 2 Go to **Host Properties > Primary Server > Servers > Additional server** tab. Add an entry for the external media server on the primary server.
- 3 Go to **Host Properties > Media Server > Servers -> Additional server** tab. Add an entry for the external media server (all media cluster nodes).
- 4 Go to **Host Properties > Clients > Servers > Additional server** tab. Add an entry for the external media server on the client server.
- 5 Run the following command:

```
NBEMMCMD, Version: 10.0  
Command completed successfully.
```

where *ems_fqdn* is the external media server FQDN and is the primary server FQDN.

- 6 Restart the EMS server services.

- 7 Add the DNSMASQ entry on the EMS server. If the EMS server is Linux, you have to add the server's IP and FQDN entry as the nameserver. If the EMS server is a Windows server, give the preferred DNS as DNSMASQ server's IP. This entry is required if you want the EMS to get connected to the secondary site primary server after disaster recovery.
- 8 Perform a robot inventory. Go to **Configure Storage devices** and confirm that the external media server is detected.

To perform a robot inventory, refer to the *Veritas NetBackup™ Administrator's Guide, Volume I*.

If you have configured disaster recovery and want to configure an external media server, then perform the following steps:

- Clear the host cache on the primary server node

```
bpcIntcmd -clear_host_cache
```

- Check for the correct IP resolved by EMS after disaster recovery is configured:

```
# bptestnetconn
# bptestbpcd -M
# bpcIntcmd -hn
```

- After clearing the host cache, *bptestnetconn* parameter should not show the entry, *Stale entry message for EMS server*.

Configuring an external NBA media server

This section describes how to configure an external NetBackup Appliance (NBA) media server with NetBackup Flex Scale.

To configure an external NBA media server with NetBackup Flex Scale

- 1 Deploy the NetBackup Appliance ISO on a physical 5330/5340 Appliance.
- 2 Perform the following pre-configuration steps before configuring the media role.
 - Configure the network for NBA. Set the hostname, DNS, and NTP server for NBA.
 - On the Java console, go to **Host Properties > Master Server > Servers > Additional server** tab. Add an entry for EMS on the primary server.
 - On the Java console, go to **Host Properties > Media Server > Servers -> Additional server** tab. Add an entry for EMS.

- 3** Configure media role for NBA as EMS on the NetBackup Appliance CLISH.
 - In the Appliance CLISH, enter the following command:

Media
 - Choose to trust the certificate during media role configuration.
 - Configure storage for AdvancedDisk pool (ADP) during media role configuration.
- 4** Create SLP on the Java console. Go to **Storage > Storage Lifecycle Policies** and right click and choose **New Storage Lifecycle Policy** to create a new SLP.
- 5** Create a backup policy. Go to **Policies > Attribute**. Select the SLP created in the previous step for **Policy storage**.
- 6** Run the backup job.
- 7** Run the restore job.

Site-based disaster recovery in NetBackup Flex Scale

This chapter includes the following topics:

- [About site-based disaster recovery in NetBackup Flex Scale](#)
- [Configuring disaster recovery using GUI](#)
- [Clearing the host cache](#)
- [Managing disaster recovery using GUI](#)
- [Performing disaster recovery using RESTful APIs](#)
- [Active-Active disaster recovery configuration](#)
- [NetBackup optimized duplication using Storage Lifecycle Policies](#)

About site-based disaster recovery in NetBackup Flex Scale

The NetBackup Flex Scale provides a disaster recovery solution. The NetBackup catalog is replicated using Veritas Volume Replicator (VVR). NetBackup optimized duplication is used to replicate the backup images.

Features of disaster recovery

- Protect NetBackup catalog and backup data against a site failure.
- Support for dual site and single NetBackup domain configuration.

- Mechanism to takeover and migrate the primary service from one site to another.
- Configuration, management, failover, and migration using RESTful APIs and GUI.
- Asynchronous continuous replication for NetBackup catalog and backup data.

Considerations before configuring disaster recovery

- The NetBackup Flex Scale primary site and the secondary site should be in the same NetBackup domain.
- At any given time, the NetBackup primary service is active on one site only. The media service is active on both the sites.
- Clients on the primary site backup to the primary site and then the backup images are duplicated to the secondary site.
- The node names of the clusters in the primary and secondary site should not be the same.
- You can configure disaster recovery on an active NetBackup Flex Scale Appliance (which has existing backup data and acts as the primary site).
- The secondary site that is added must be a freshly installed or configured NetBackup Flex Scale Appliance. Else, any data present on the secondary site is lost when you set up disaster recovery.
- As part of the initial configuration of the secondary site, only the primary data network should be configured on the secondary site. Secondary data network should not be configured on the secondary site before disaster recovery configuration.
- As part of the initial configuration on the secondary site (whether in the YAML file or in the GUI), users should be created on the secondary site in the same order as they are present in the primary site.
- As part of initial configuration pre-requisites, the secondary site must be in Normal lockdown mode only. The primary site can have any lockdown mode configured. After disaster recovery configuration, you can configure any lockdown mode on any of the clusters as per the lockdown mode guidelines.
- As part of disaster recovery migration and takeover, the FQDN and IP mapping for NetBackup primary service needs to be updated manually in the DNS. The NetBackup primary service FQDN should point to the virtual IP present on the new primary site.
- Primary site can have the secondary data network configured before disaster recovery configuration.

- After disaster recovery is configured between two sites, disaster recovery configuration cannot be removed.

Note: Disaster recovery can also be configured if one or both the clusters are in non-DNS configurations.

If disaster recovery is configured:

- The cluster from which the disaster recovery configuration is initiated is called the primary site and other cluster is called the secondary site.
- Users have to be manually added on the secondary site in the same order as they are present in the primary site.
- If you want to add a node to the cluster, See [“Considerations for adding a node when disaster recovery is configured”](#) on page 60.
- If you want to configure VLAN, See [“Support for multiple VLAN when disaster recovery is configured”](#) on page 101.
- If disaster recovery is configured in a non-DNS environment, ensure that you add the NetBackup Client host entries in **Settings > Network > Custom hosts** at both sites.


Note: If a cloud LSU is present before the entire cluster on which the NetBackup Flex Scale software resides is destroyed, refer to the *About the disaster recovery for cloud LSU* section in the *Veritas NetBackup Deduplication Guide*.

Configuring disaster recovery using GUI

You can set up and configure disaster recovery using the NetBackup Flex Scale GUI.

To configure disaster recovery using GUI

- 1 Log on to the NetBackup Flex Scale GUI of the cluster from which you want to configure disaster recovery. After the disaster recovery configuration is complete, this cluster acts as the primary site.




- 3 Enter the configuration details for the primary and secondary site.
 - Replication and heartbeat IP address for the primary and secondary sites. Provide IPv4 or IPv6 addresses according to the network configuration.
 - Management server IPv4/IPv6 address or FQDN.

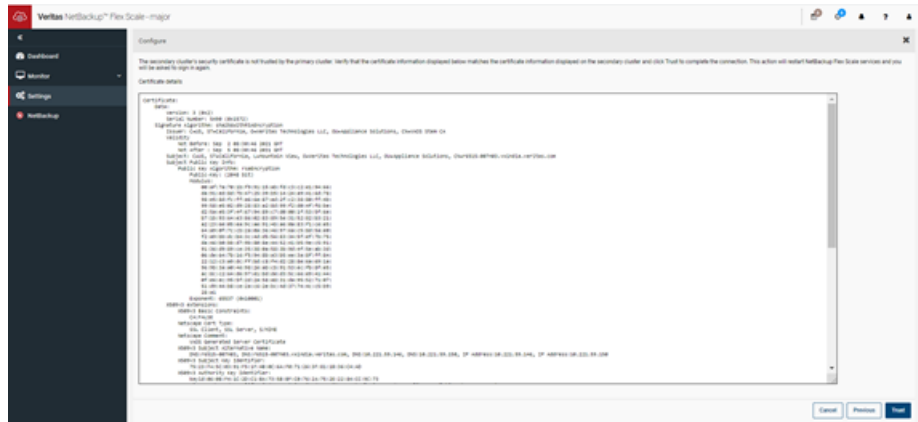
Note: If you want to configure disaster recovery in a non-DNS environment, provide IPv4/IPv6 address for the management server.

- Appliance administrator username and password of the secondary site.
- Secondary storage credentials which are used to create a new storage user for the secondary site.

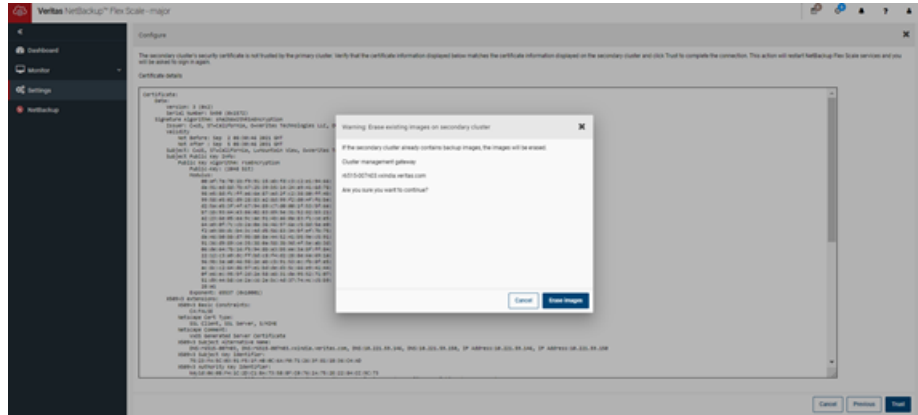
Click **Configure**.



- 4 Verify the security certificate information of the secondary site that is displayed. Click **Trust**.




- 5 A warning pop-up appears that existing images, if any, on secondary site will be erased. Click **Erase images** to continue with disaster recovery configuration.



6 Disaster recovery configuration is initiated.

You are required to log on to the clusters again after the certificates are exchanged. This takes a few minutes.

You can monitor the configuration from **Settings** > for both the clusters. On the primary site, all the sub-tasks can be seen in progress.



On the second cluster only 'Configure replication' task appears.

7 After the configuration is complete, you can also check the status by navigating to **Settings** > for both the clusters. This shows the information specific to primary and secondary site with some common information. You can also check the replication status in the **Services** section of the dashboard.

Clearing the host cache

There is a time window in which the DNS entries on the clients are cached. Clients can continue to use the earlier primary service's IP address even after failover and backup and restore jobs may fail till the new DNS entry is fetched from DNS server.

For additional information, see

https://www.veritas.com/support/en_US/article.100025822


Managing disaster recovery using GUI

You can manage disaster recovery using the NetBackup Flex Scale GUI. You can perform both takeover and migrate operations using the GUI. It is also possible to pause the replication and resume it later.

Performing migration


To perform migration using GUI

- 1 Log in to the NetBackup Flex Scale GUI of the secondary site.
- 2 Go to **Settings > Master service replication**. Click **Change Replication Roles**.



3 The **Change Replication Roles** window appears.

If you want to perform a migration, select the **Switch roles** option. Ensure that you have updated the DNS entries of the NetBackup primary service FQDNs to resolve to the secondary site. Click **Change roles**.



Note: Only the **Switch roles** option is enabled when both the primary and secondary sites are up. The **Make primary** option is disabled.

Note: In case of non-DNS sites (either one or both), updating the DNS entry for the NetBackup primary service IP is not applicable.


Performing takeover

You can perform a takeover operation only if the primary site is fully down. The Dashboard of the secondary site should show that the replication status as `Disconnected`.

To perform takeover using GUI

- 1 Log in to the NetBackup Flex Scale GUI of the secondary site.
- 2 Go to **Settings > Master service replication**. Click **Change Replication Roles**.
- 3 The **Change Replication Roles** window appears.

If you want to perform a takeover, select the **Make primary** option. Ensure that you have updated the DNS entries of the NetBackup primary service FQDNs to resolve to the secondary site. Click **Change roles**.



Note: The **Make primary** option is enabled only when the primary site is fully down. The **Switch roles** option is disabled.

Note: In case of non-DNS sites (either one or both), updating the DNS entry for the NetBackup primary service IP is not applicable.

Pause replication

You can pause an ongoing replication operation.

To pause replication

- 1 Go to **Settings** > page of the primary or the secondary site.
- 2 Click **Pause**. A task is initiated in the GUI. After task completion, replication is paused on both the clusters and the **Resume** button gets enabled on both the clusters.

Resume replication

You can resume a replication operation which is in paused state.

To resume replication

- 1 Go to **Settings** > page of the primary or the secondary site.
- 2 Click **Resume** to resume replication. A task is initiated in the GUI. After task completion, replication is resumed on both the clusters and the **Pause** button gets enabled on both the clusters.

Note: You can perform a resume operation only if the replication is in paused state. The pause and resume operations can be performed from either of the primary or secondary sites.

Performing disaster recovery using RESTful APIs

You can set up and configure disaster recovery using RESTful APIs

High level steps to configure disaster recovery

- Deploy the secondary site like a regular NetBackup Flex Scale cluster. The secondary site configuration should be equivalent to the primary site in terms of number of nodes, hardware and software revision.
- Create trust between the two clusters through the exchange of Appliance web server and its CA certificates. Establish an authentication mechanism between the two clusters.
See [“Establishing trust and setting up authentication”](#) on page 139.
- Configure the primary and secondary site for disaster recovery.
See [“Configuring disaster recovery”](#) on page 140.
- Configure NetBackup optimized duplication SLPs.
See [“Configuring a Storage Lifecycle Policy for optimized duplication”](#) on page 182.

Establishing trust and setting up authentication

In NetBackup Flex Scale Appliance, the Appliance web server creates a self-signed CA certificate and an Appliance web server certificate (signed by that CA) for every cluster. As the CA certificates are different, you have to ensure that both the clusters trust the CA of the other. This is done by adding one cluster's gateway CA certificate to the trusted certificate store of the other. Each cluster should be able to trust the secondary site and perform the required operations to configure and manage disaster recovery. Appliance web server certificates are exchanged between both the clusters to enable authentication.

You can use the following RESTful APIs to setup trust and authentication. The secondary site can be added as a disaster recovery cluster to the primary site anytime after the primary site is up and running. The secondary site has to be a freshly installed and configured NetBackup Flex Scale cluster.

The RESTful API calls must be made in the following order on the management server of the clusters. The API calls do not use SSH between the clusters.

You can get the list of available certificates on the cluster using the API:

```
GET /api/appliance/v1.0/certificates
```

This API returns the URI of *certificateName* as Appliance web services certificate, Appliance web services CA certificate and the root certificate. The *certificateName* should be passed as input to the GET specific certificate API. The certificate is in base64 encoded format.

To set up trust between both the clusters

- 1 Get the appliance web services certificate on the primary site by providing the certificate name (appliance-webservice) as input in the GET specific certificate API.

```
GET /api/appliance/v1.0/certificates/{certificateName}
```

- 2 Get the appliance web services CA certificate on the primary site by providing the certificate name (appliance-webservice-ca) as input in the GET specific certificate API.

```
GET /api/appliance/v1.0/certificates/{certificateName}
```

- 3 Establish trust by passing the certificates obtained from the primary site to the secondary site. Execute the following API on the secondary site:

```
POST /api/appliance/v1.0/certificates
```

The API imports the certificates of one cluster and exports the certificates on the other cluster to establish trust and enable certificate authentication.

Set `type` as `appliance-web-service` and `purpose` as `remote-cluster-trust-auth`. Use the management server FQDN of the remote cluster for `gateway`.

- 4 Get the appliance web services certificate on the secondary site by providing the certificate name (`appliance-web-service`) as input in the GET specific certificate API.

```
GET /api/appliance/v1.0/certificates{certificateName}
```

- 5 Get the appliance web services CA certificate on the secondary site.

```
GET /api/appliance/v1.0/certificates/appliance-web-service-ca
```

- 6 Establish trust by passing the certificates obtained from the secondary site to the primary site. Execute the following API on the primary site:

```
POST /api/appliance/v1.0/certificates
```

The API imports the certificates of one cluster and exports the certificates on the other cluster to establish trust and enable certificate authentication.

Set `type` as `appliance-web-service` and `purpose` as `remote-cluster-trust-auth`. Use the management server FQDN of the remote cluster for `gateway`.

For more information, see the *Veritas NetBackup Flex Scale APIs* on SORT.

Configuring disaster recovery

The user has to invoke RESTful API calls to configure disaster recovery between the clusters. The NetBackup Flex Scale cluster that is to be used as the primary may already have backup data. Before you run the API, ensure that the NetBackup Flex Scale cluster is in healthy state. The NetBackup Flex Scale cluster that is used as secondary must be a freshly configured cluster. Any changes made after the initial configuration is lost, including any backup images.

For the list of ports that should be open, see *Firewall and network port requirements* section in the *Veritas NetBackup™ Flex Scale Installation and Configuration Guide* on SORT.

To configure disaster recovery:

- Four virtual IP addresses are required, two on each cluster in the data network. One set of IP addresses are used for VVR replication and another set of IP addresses are used for heartbeat between clusters.
- The API call to set up disaster recovery is asynchronous. Once the primary site receives the call, it creates a task and returns a task ID. You can find the status of the replication setup by querying the task with the task ID using the following API:

```
GET /api/appliance/v1.0/tasks/{taskId}
```

- Run the API to configure disaster recovery by specifying the primary and secondary site parameters. This API should be executed on the primary site.

```
POST /api/appliance/v1.0/disaster-recovery
```

This API can be retried if you encounter any network or timeout failures during disaster recovery configuration. But, do not attempt to retry it if disaster recovery is already configured.

For more information, see the *Veritas NetBackup Flex Scale APIs* on SORT.

The API sets up replication of the NetBackup catalog using VVR. For replicating the NetBackup policies, use NetBackup optimized duplication (SLP).

See [“Configuring a Storage Lifecycle Policy for optimized duplication”](#) on page 182.

Managing disaster recovery

You can manage disaster recovery using RESTful APIs.

1. Perform takeover:

Run the API to takeover a secondary site as primary. Set the value of `operation` field as `takeover`.

```
PATCH /api/appliance/v1.0/disaster-recovery
```

This API is run to give the primary role to the cluster on which this API is executed. Takeover is a user initiated operation and does not happen automatically. The API should be called on the secondary after the current primary site is down and not reachable. The primary site is completely down and the nodes of the primary site are not reachable.

Before you make this API call:

- The FQDN and IP mapping for NetBackup primary service needs to be updated manually in the DNS. The NetBackup primary service FQDN should point to the virtual IP present on the new primary site.
- In case multiple data networks are configured, the FQDN and IP mapping of the NetBackup primary service for each data network needs to be updated manually in the DNS such that the NetBackup primary service FQDN for each data network points to the virtual IP present on the new primary site.

To clear the host cache on clients, See [“Clearing the host cache”](#) on page 134.

Current backup and restore jobs fail after you make this call and you have to restart the jobs after the takeover is complete. There is a time window before the NetBackup primary service is brought online on the secondary during which the backup and restore jobs fail.

2. Perform migration:

Run the API to migrate the primary role to the cluster on which the API is executed. Set the value of `operation` response field as `migrate`.

```
PATCH /api/appliance/v1.0/disaster-recovery
```

This API can be called only on the secondary gateway.

Before you make this API call:

- The FQDN and IP mapping for NetBackup primary service needs to be updated manually in the DNS. The NetBackup primary service FQDN should point to the virtual IP present on the new primary site.
- In case multiple data networks are configured, the FQDN and IP mapping of the NetBackup primary service for each data network needs to be updated manually in the DNS such that the NetBackup primary service FQDN for each data network points to the virtual IP present on the new primary site.
- The secondary should be completely in sync with the primary.

To clear the host cache on clients, See [“Clearing the host cache”](#) on page 134.

Current backup and restore jobs fail after you make this call and you have to restart the jobs after the migration is complete. There is a time window before the NetBackup primary service is brought online on the secondary during which the backup and restore jobs fail.

To change the SLP of each of policies used in the new primary, See [“Updating the policy to reverse the replication direction”](#) on page 190.

3. Get the replication status:

You can get the state of the primary service catalog and the details of how the cluster lags with respect to the primary using the API:

```
GET /api/appliance/v1.0/disaster-recovery
```

For details on the response fields, See [“About response fields in the GET disaster recovery API”](#) on page 193.

4. Pause replication:

Run the API to pause the NetBackup catalog replication between the primary and secondary sites. Set the value of the `operation` response field as `pauseReplication`.

```
PATCH /api/appliance/v1.0/disaster-recovery
```

5. Resume replication:

Run the API to resume the NetBackup catalog replication between the primary and secondary sites. Set the value of the `operation` response field as `resumeReplication`.

```
PATCH /api/appliance/v1.0/disaster-recovery
```

To restore the backup image from the disaster recovery cluster after a takeover or migrate operation, see the *Promoting a copy to a primary copy* section in the *Managing backup images* chapter in the *Veritas NetBackup™ Administrator's Guide, Volume I* on SORT.

For more information, see the *Veritas NetBackup Flex Scale APIs* on SORT.

Active-Active disaster recovery configuration

Once the catalog replication is configured, you can use site specific backup policies to perform local backup and use duplication policies to replicate backup images to remote cluster for protection against site outage. The clients communicate with the primary server on the primary site for sending metadata and communicate with media servers available on local cluster for sending backup data. The clients that are located at the primary site should use the local STU available at the primary site for backup and use the remote site's STU for duplication. Similarly, the clients that are located at the secondary site should use the local STU available at the secondary site for backup and use the remote site's STU for duplication. This way, it is ensured that duplications are performed across the sites for local backups. After disaster recovery configuration, the STUs of both the sites are always available for backups and duplications. Default SLPs are created as a template to configure backups and duplications at each site. The NetBackup administrator can use the

default SLPs or can create new SLPs based on the template to perform backup and duplication.

Active-Active disaster configuration can be done from the NetBackup primary server using the web UI or the NetBackup Java GUI.

Takeover operation in Active-Active disaster recovery sites:

When any disaster happens, such as the primary site goes down for some reason, the clients at the primary site will not be able to continue the backups at the primary site as the local STU will not be available. The clients at the secondary site will also not be able to perform duplications as the remote STU will not be available for duplications. The backup jobs for primary site clients will fail or remain in waiting state for the STU to become available. The duplications jobs from secondary site clients will also fail or remain in waiting state for the STU to become available. The same situation arises when the secondary site is down due to any disaster.

When any disaster happens at a primary site, the secondary site is the only available site and must become the primary site so that the NetBackup primary server is available for the clients and pending jobs can be resumed or restarted. The takeover operation needs to be performed to make the secondary site as the primary site. Once the site becomes primary, the duplication jobs of old primary site are resumed/restarted and completed. The clients at secondary site continue to perform backup as the local STU is available for backups, but duplication jobs will remain in failed/waiting state.

Note: The NetBackup administrator can make the backup storage as the new primary site's STU until the old primary site is available so that the client backup jobs of the old primary site can be completed. This can be done by making the primary site's SLPs to point to the new primary site's STU as backup storage.

Important: When the site role changes, the clients residing in domains other than the NetBackup Flex Scale disaster recovery domain may be required to update the host entries in the NetBackup configuration to connect to the media servers of the site.

Migrate operation in Active-Active disaster recovery sites:

If disaster recovery is configured and the primary site requires maintenance or updates, it can be switched to the role of secondary site using migrate operation. When the migrate operation is performed, the primary site becomes the secondary site and the secondary site becomes the primary site without affecting the backup and duplication jobs of the client. The NetBackup administrator does not have to change the SLPs since both the sites are up and the NetBackup primary service

is also available. The clients at both sites will continue to do backups (to the local site STU) and duplication (to the remote site STU).

NetBackup optimized duplication using Storage Lifecycle Policies

Storage Lifecycle Policy (SLP) is a way of creating a storage plan for backup data. Each SLP can have operations which define how data is stored, retained, and replicated. Backup data can be duplicated from one storage unit to another storage unit. During disaster recovery configuration, the storage unit for backup comes from the primary site and the storage unit for duplication comes from the secondary site. During primary service failover and migration, Veritas recommends reversing the direction of duplication by changing the SLP under the backup policy.

As part of disaster recovery configuration, four SLP templates are created using the NetBackup RESTful APIs for SLP management.

Format of the SLP template:

```
<source_cluster_name>_<retention>_to_<destination_cluster_name>_retention
```

- Weekly retention: `clusterA_7days_to_clusterB_7days`
- Monthly retention: `clusterA_30days_to_clusterB_30days`
- Weekly retention: `clusterB_7days_to_clusterA_7days`
- Monthly retention: `clusterB_30days_to_clusterA_30days`

For example, the SLP template `clusterA_7days_to_clusterB_7days` backups the images to cluster A, duplicates the images from clusterA to clusterB, and retains the data for 7 days in both the clusters.

For more details on the NetBackup RESTful APIs used for SLP configuration, refer the *NetBackup Configuration API* document on SORT.

To promote the secondary (copy #2) as primary, see the *Promoting a copy to a primary copy* section in the *Managing backup images* chapter in the *Veritas NetBackup™ Administrator's Guide, Volume I* on SORT.

NetBackup Flex Scale security

This chapter includes the following topics:

- [STIG overview for NetBackup Flex Scale](#)
- [FIPS overview for NetBackup Flex Scale](#)
- [Managing the login banner](#)
- [Changing the password policy](#)
- [Support for immutability in NetBackup Flex Scale](#)
- [Deploying external certificates on NetBackup Flex Scale](#)

STIG overview for NetBackup Flex Scale

The Security Technical Implementation Guide (STIG) provides technical guidance for increasing the security of information systems and software to help prevent malicious computer attacks. This type of security is also referred to as hardening.

NetBackup Flex Scale uses STIG to meet security requirements as per the Defense Information Systems Agency (DISA) profile:

STIG for Red Hat Enterprise Linux 7 Security Technical Implementation Guide - Version 3, Release 3.

The STIG option is enabled at cluster level. If the STIG option is enabled, the STIG rules are enforced on all the nodes in a cluster.

STIG-compliant password policy rules

To comply with the Security Technical Implementation Guide (STIG), NetBackup Flex Scale automatically enforces a higher security password policy when the STIG option is enabled. After the STIG option is enabled, all current user passwords that were created under the default policy remain valid. When you change any user passwords, the STIG-compliant policy rules must be followed.

The STIG-compliant password policy rules are listed below:

Password complexity

- Minimum characters: 15
- Minimum numbers: 1
- Minimum lowercase characters: 1
- Minimum uppercase characters: 1
- Minimum special characters: 1
The permitted special characters are: ~!@#\$%^&_+~[]{}.,<>|
- Minimum character classes: 4
- Maximum consecutive repeating characters: 2
- Maximum consecutive repeating characters of the same type: 4
- Minimum number of different characters: 8
- No whitespaces.
- Dictionary words are not allowed

Password age

- Days after which a password expires: 60
- Minimum days before a password can be changed: 1
- Days before a password must be changed: 60
- The previous seven passwords cannot be reused.

Password lockout

- Number of incorrect login attempts before lockout: 3
- Time before locked account is reenabled (seconds): 604800
- Time between login failures before account lockout (seconds): 900

Enabling STIG for NetBackup Flex Scale

With NetBackup Flex Scale version 3.0, you can enable STIG hardening rules for increased security. These rules are based on the following profile from the Defense Information Systems Agency (DISA):

STIG for Red Hat Enterprise Linux 7 Security Technical Implementation Guide - Version 3, Release 3.

After the STIG option is enabled:

- A STIG-compliant password policy is automatically enforced. All current user passwords that were created under the default password policy remain valid. Once a password expires, you must follow the STIG-compliant policy rules when you change the password.
See [“STIG-compliant password policy rules”](#) on page 147.
- The STIG default login banner is displayed when you log in to the NetBackup Flex Scale UI and the NetBackup Administration Console. View the **Alert! Accessing Information System** window and click **Continue** to proceed.

Review the following guidelines before enabling STIG:

- When you enable STIG, the STIG option is configured for all the nodes in a cluster. The cluster must be configured before you enable the STIG option.
- The STIG option does not allow individual rule control.
- Before you enable STIG, it is recommended that you complete the following prerequisites. However, not completing the prerequisites does not prevent you from enabling STIG. You can complete these requirements after you enable the STIG option.
 - Configure at least two NTP servers for the cluster.
 - Configure at least two DNS servers for the cluster.
 - Configure an SMTP server to enable notifications.
- After the STIG option is enabled, a factory reset is required to disable the associated rules. You cannot disable the option using the UI or the REST APIs.
- Veritas recommends that you do not perform any other tasks while the STIG operation is in progress.
- If site-based disaster recovery is configured, ensure that both the primary and the secondary clusters have similar STIG configuration. If STIG is enabled for the primary cluster, the STIG option must be enabled for the secondary cluster. Similarly, if STIG is not enabled for the primary cluster, do not enable STIG for the secondary cluster.

Enabling STIG using the NetBackup Flex Scale web interface


To enable the STIG hardening rules, complete the following steps:

- 1 Use any one of the following options to log in using the user account that you created:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.
- 2 In the navigation pane, click **Settings**.
- 3 Click **Security management**.

4 On the **STIG** tab, click **Enable STIG**.

If the prerequisites are not met, you are prompted to resolve the errors. However you can choose to ignore these errors and proceed by clicking **Continue**. You can complete the prerequisites later after you enable the STIG option. If the requirements are met, review the displayed guidelines and click **Enable**.

Note: Do not perform any other tasks until the STIG enable operation is complete.



5 To monitor the progress, click **View details** on the **Security** page. The ongoing and completed tasks for the operation are also displayed in **Recent activity**.

After the operation is complete, you can view the STIG status for all the cluster nodes. If STIG is enabled for a node, the status is displayed as **Enabled**. If the STIG option cannot be enabled for a node, the status is displayed as **Not Enabled**, and if the node status cannot be retrieved because the node is stopped, shut down, or not reachable, the status is displayed as **Unknown**.

For nodes that display **Unknown** status, you can enable the STIG option again or wait for the node to automatically synchronize its status with the cluster after the node is up.

If some of the STIG rules fail or you make any updates to the cluster settings or configuration, you can enforce the STIG rules again on the nodes where the STIG option is already enabled by clicking **Enable STIG**.

Enabling STIG using REST APIs

You can use the following API to enable STIG:

```
PATCH /api/appliance/v1.0/security/stig
```

You can find the REST APIs at

`https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0/` where *ManagementServerIPorFQDN* is the public IP address or FQDN that you specified for the management server and API gateway during the cluster configuration. For more details about the APIs, see the NetBackup Flex Scale APIs on SORT.

Viewing the NetBackup Flex Scale STIG status

You can use the NetBackup Flex Scale web interface or the REST APIs to view the STIG status.

You can find the RESTful APIs at

`https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0/` where *ManagementServerIPorFQDN* is the public IP address or FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration.

If you are using IPv6 addresses, use the following URL syntax:

```
https://[ManagementServerIP]:14161/swagger/infra/v1.0
```

Use the following API to view the STIG status:

```
GET /api/appliance/v1.0/security/stig
```

Viewing the status in the web interface

- 1 Use any one of the following options to log in using the user account that you created:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**,

in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.

- Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.
- 2 In the navigation pane, click **Settings**, and then click **Security management**.
- 3 Click the **STIG** tab.

The STIG status for all the cluster nodes is displayed:

- **Enabled:** The STIG option was successfully enabled for the node.
- **Not Enabled:** The STIG option is not enabled for the node.
- **Unknown:** The node status cannot not be retrieved because the node is stopped, shut down, or not reachable.

FIPS overview for NetBackup Flex Scale

The Federal Information Processing Standards (FIPS) define U.S. and Canadian Government security and interoperability requirements for computer systems. The National Institute of Standards and Technology (NIST) issued the FIPS 140 Publication Series to coordinate the requirements and standards for validating cryptography modules. The FIPS 140-2 standard specifies the security requirements for cryptographic modules and applies to both the hardware and the software components. It also describes the approved security functions for symmetric and asymmetric key encryption, message authentication, and hashing.

For more information about the FIPS 140-2 standard and its validation program, see the following links:

<https://csrc.nist.gov/csrc/media/publications/fips/140/2/final/documents/fips1402.pdf>

<https://csrc.nist.gov/projects/cryptographic-module-validation-program>

The NetBackup Flex Scale Cryptographic Module is FIPS validated. Starting with NetBackup Flex Scale 3.0, FIPS 140-2 standard is enabled with the default factory settings for the Veritas Optimized Operating System (VxOS). After FIPS for VxOS is enabled, the `sshd` uses the following FIPS approved ciphers:

- aes128-ctr
- aes192-ctr

- aes256-ctr

The FIPS 140-2 standard is enabled for NetBackup MSDP when you create a NetBackup Flex Scale cluster.

Note: You cannot disable the FIPS option for VxOS or for NetBackup MSDP.

Viewing the NetBackup Flex Scale FIPS status

You can use the NetBackup Flex Scale web interface or the REST APIs to view the FIPS status. The FIPS 140-2 standard is enabled with the default factory settings for the Veritas Operating System (VxOS) and for NetBackup MSDP when you create a NetBackup Flex Scale cluster.

You can find the RESTful APIs at

`https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0/` where *ManagementServerIPorFQDN* is the public IP address or FQDN that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration.

If you are using IPv6 addresses, use the following URL syntax:

`https://[ManagementServerIP]:14161/swagger/infra/v1.0`

Use the following API to view the FIPS status:

`GET /api/appliance/v1.0/security/fips`

Viewing the status in the web interface

- 1 Use any one of the following options to log in using the user account that you created:

- Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface

`https://ManagementServerIPorFQDN/webui` where

ManagementServerIPorFQDN is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.

- Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console


`https://ManagementServerIPorFQDN:14161` where

ManagementServerIPorFQDN is the public IP address or the FQDN that

you specified for the NetBackup Flex Scale management server during the cluster configuration.

- 2 In the navigation pane, click **Settings**, and then click **Security management**.
- 3 Click the **FIPS** tab.

The FIPS status for VxOS and NetBackup MSDP is displayed. You can also view the FIPS status for VxOS for each node in the cluster. For nodes that are unreachable or are stopped, the status is displayed as **Unknown**.



Managing the login banner

You can create a customized text banner that appears when you sign in to NetBackup Flex Scale UI, system console, or NetBackup UI. You can use the login banner to communicate various kinds of messages to users. Typical uses for the login banner include legal notices, warning messages, and company policy information.

If the Security Technical Implementation Guide (STIG) mode is enabled for the cluster, the login banner cannot be modified.

To set a login banner:

- 1 Use any one of the following options to log in using the user account that you created:

- Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.
- 2 In the left navigation pane, click **Settings** and then click **User management**.
 - 3 Click **Manage sign-in banner**.
 If a banner is already set, the Manage sign-in banner page displays the current banner.
 - 4 On the Manage sign-in banner page, click **Edit**.
-
- Note:** The Edit button is disabled and the banner cannot be edited if the STIG mode is enabled.
-
- 5 (Optional) Under **Sign-in banner heading**, enter the banner heading. The heading can be a maximum of 250 characters.
 - 6 Under **Sign in banner text**, enter the text for the banner message. The message can be a maximum of 4000 characters.
 - 7 To review the changes, click **Preview**.
 - 8 To confirm the changes, click **Save**.

Changing the password policy

You can customize the password policies by setting rules for the passwords that are used by the NetBackup Flex Scale local users. You can set rules for password complexity, password age, and password lockout. Password complexity specifies the number and type of characters a password must include. Password age defines the duration for which the password is valid. Password lockout specifies the number

of failed attempts because of incorrect usage of passwords after which a user is prevented from logging in to the account.

If the Security Technical Implementation Guide (STIG) mode is enabled for the cluster, the password policy cannot be modified.

The default password policy for a local user is as follows:

Password complexity:

- Minimum characters: 8
- Minimum numbers: 1
- Minimum lowercase characters: 1
- Minimum uppercase characters: 1
- Minimum special characters: 1

To edit the password policy:

- 1 Use any one of the following options to log in using the user account that you created:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.
- 2 In the left navigation pane, click **Settings** and then click **User management**.
- 3 Click **Manage password policy**.
- 4 On the Manage password policy page, click **Edit**.

Note: The Edit button is disabled and the password policy cannot be edited if the STIG mode is enabled.

- 5 If you want your password policy to comply with STIG, select **Reset to STIG default values** to fill in the default values for all the parameters.
- 6 Edit the parameters as required. To ignore a rule, leave the corresponding parameter blank. After making the changes click **Save**.

Table 9-1

Parameter	Description
Minimum characters	Minimum number of characters to include in a password
Minimum uppercase characters	Minimum number of uppercase characters to include in a password
Maximum repetitive characters of the same class	Maximum number of consecutive uppercase, lowercase, numeric, and special characters
Minimum numbers	Minimum number of numeric characters
Minimum special characters	Minimum number of special characters in a password
Minimum character classes	Minimum character classes to include in a password. Character classes include uppercase, lowercase, numeric, and special characters.
Minimum lowercase characters	Minimum number of lowercase characters
Maximum repetitive characters	Maximum number of characters that can be repeated in a password.
Character difference with old password	Number of characters the new password must differ by from the previous password
Days after which password can be changed	Number of days after which a password can be changed
Days after which password must be changed	Number of days after which a password must be changed
Days before warning message	Number of days before the password expires to display a warning
Minimum different passwords before allowing reuse	Number of unique passwords before a previous password can be reused
Number of incorrect login attempts before lockout	Number of failed login attempts after which the account gets locked

Table 9-1 (continued)

Parameter	Description
Time before locked account is reenabled	Duration in seconds the account remains locked
Time between login failures before account lockout	Number of seconds between consecutive failed login attempts

Support for immutability in NetBackup Flex Scale

Immutability support for backup images requires locking down the appliance and not permitting any operations that can lead to data destruction. When the appliance is placed in lockdown mode, administrators are prevented from making any changes to the operating system and the internal components.

Important features:

- Immutable data support with retention locking
- Retention lock deletion for backup images
- Access to Remote Management Platform (HP ILO)
- Transition between different modes
- Retention lock extension

About lockdown modes

Lockdown mode is one of the features of ransomware protection. The lockdown mode protects your cluster data from internal and external threats by securing all the external endpoints from unauthorized access. Access to all the services is protected and authenticated.

NetBackup Flex Scale lockdown mode offers additional security levels to protect your appliance and data, in addition to the hardened, secure operating environment that comes out of the box.

Lockdown mode provides the following benefits:

- It prevents unauthorized access or modification to the underlying operating system (OS). Once the lockdown mode is enabled, administrators cannot make changes to the OS or the internal components. If you need access to the OS for emergency operations, you must contact Veritas Technical Support to obtain a Support Key and temporarily unlock the appliance. This functionality prevents unauthorized changes even if a malicious user gains access to stolen credentials.

- It gives the appliance users options for managing WORM (Write Once Read Many) data. Your data is protected from being encrypted, modified, and deleted using WORM properties.

Different lockdown modes provide different level of granularity for WORM and retention. The NetBackup Flex Scale appliance support three lockdown modes.

- **Normal mode:**

This is the default mode of the cluster if the lockdown mode is not specified during installation. In this mode, WORM and retention capabilities are disabled. User cannot create worm STU in this mode.

- **Enterprise mode:**

In this mode, WORM and data retention features are enabled. User can choose to create WORM enabled STU. Also, in this mode user has the option to remove the retention locks and expire image data. The user can extend the retention period but cannot reduce the retention period.

The retention time period can be extended from the NetBackup primary container only if the user has the NetBackup administrator role.

Retention can be disabled or retention lock can be removed using the MSDP Restricted Shell only if the user has the appliance administrator role.

After removing the images retention locks from the MSDP Restricted Shell, the user still cannot expire images from the NetBackup Administration Console, but can expire the images from the NetBackup primary server using the following command:

```
/usr/opensv/netbackup/bin/admincmd/bpexpdate -backupid
n155-h201.cdc.veritas.com_1631842421 -d 0 -copy 1
-try_expire_worm_copy
```

- **Compliance mode:**

In this mode, WORM and data retention features are enabled. The user can extend the retention period. The user does not have the option to remove retention locks and expire image data before the predefined time. Once appliance lockdown mode is set to compliance, user does not have the option to delete data until it is expired.

Veritas strongly recommends that you enable enterprise lockdown mode to prevent unauthorized access to the OS, even if you do not plan to create WORM storage instances.

Selecting or changing the lockdown mode

The user can select the lockdown mode during initial configuration. After cluster configuration, user has the option to see/change the lockdown mode using both

GUI and REST APIs. The lockdown modes can be switched only if the engines are healthy. The user can switch between the following modes without any restriction:

- From normal to enterprise mode
- From normal to compliance mode
- From enterprise to compliance mode

The user can set minimum and maximum retention time for backup images for enterprise and compliance mode only. Creation of images with retention time less than the minimum retention time or greater than the maximum retention time is not allowed. This minimum and maximum retention time should be set by the appliance administrator as per the retention requirement of their use case.

- Once the lockdown mode is set, only Appliance administrators can change the lockdown mode.
- The lockdown modes are maintained during upgrade.
- Only the Appliance administrator can remove the retention locks if the lockdown mode is enterprise.
- Only the users with appliance administrator role can disable retention or remove the retention lock using the MSDP Restricted Shell.
- The user cannot change the mode if any existing operation is in progress.

Restrictions in different modes

- If the mode is set to compliance mode, the administrator cannot change the mode to enterprise or normal mode.
- If lockdown mode is set to compliance or enterprise for any node, it is not available for factory reset.
- During add and replace node operations, the new node is automatically placed in the existing lockdown mode of the cluster. The lockdown mode of the node that got replaced is set to normal and the node is available for factory reset.
- Cluster maintenance shell is enabled with two-factor authentication (2FA).

To access the root shell when lockdown mode is configured

- 1 Log on to the node-level CLI on any node in the cluster.
- 2 Run the `support generate-otp` to get the OTP (valid for 2 hours) for the entire cluster.
- 3 Open a ticket with Veritas Support to generate a security key. Set a Support password which is later used to elevate to root.
- 4 Log on to the NetBackup Flex Scale shell on any node in the cluster.


- 5 Run the `support unlock` command. You are prompted to enter a security key. Enter the security key that you got in step 3. Press **Enter** to unlock the root shell access to the current node (all other nodes remain locked).
- 6 Run the `support elevate` command. You are prompted to enter a Support password. Enter the Support password set in step 3. Press **Enter**. Type the maintenance password to get into the root shell.
- 7 Repeat steps 4 to 6 to get into the root shell of all other nodes.
- 8 Run the `support lock` command on a specific node to lock that node. If no manual lock is issued, the node is locked automatically after 12 hours. All the current users are removed from the root shell in a single node.

Configuring immutability using GUI

You can configure immutability using the NetBackup Flex Scale GUI.

To configure immutability using GUI

- 1 Go to **Settings > Security management > Immutability**. Click **Edit**.




- 2 Choose any of the lockdown modes under **Mode Selection**. You can choose from Normal, Enterprise, and Compliance. The supported features for each mode are listed.

Note: You cannot downgrade the lockdown mode after it is configured. For example, if the lockdown mode is set to compliance, you cannot change the lockdown mode to enterprise or normal.

- 3 If you choose Enterprise or Compliance mode, you also have to configure the retention time range from which the NetBackup administrators will be allowed to select when creating backup images. Specify the maximum and minimum retention time. Click **Save**.

Note: Ensure that the image retention period in backup policies is within the lockdown mode retention time to avoid any error during backup.



Deploying external certificates on NetBackup Flex Scale

Starting from this release, you can generate and use external certificates instead of internal certificates. External Certificate Authority (ECA) certificates are the digital credentials that attest to the certificate owner's identity and affiliation. Once you deploy the external certificates, all the NetBackup Flex Scale components use them. These include the NetBackup primary server, media server, storage engine, management gateway, and the NetBackup Flex Scale web services. One certificate is deployed for all the components. The external certificates also deploy a certificate bundle and (optionally) certificate revocation list. To generate an external certificate, you have to create a certificate request with proper 'Subject Distinguished Name' and 'Subject Alternative Names.' You can generate a certificate request using the GUI. The necessary FQDNs are auto-populated to generate the correct request. You can add additional information as needed. Based on the certificate request, you can create an external certificate. When deploying external certificate for the first time, you have to provide a CA certificate bundle. This is used to validate the

incoming and deployed external certificate. You can also optionally provide a certification revocation list. NetBackup components use the CRL.

Some important terminologies:

- A certificate authority, also known as a certification authority, is a trusted organization that verifies websites (and other entities) so that you know who you are communicating with online. Their objective is to make the internet a more secure place for both organizations and users. Becoming a Certificate Authority (CA) means that you (or your customers) oversee the issuing process of cryptographic pairs of private keys and public certificates.
- Certificate bundle (CA bundle) is a file that contains root and intermediate certificates. The end-entity certificate along with a CA bundle constitutes the certificate chain.
- Certificate Revocation List (CRL) is a list of digital certificates that have been revoked by the issuing Certificate Authority (CA) before their scheduled expiration date and should no longer be trusted. CRL is optional. It may be provided as a file or embedded in certificate as a URL.
- Subject Alternative Name: This field lets you specify additional host names (such as sites, IP addresses, common names) to be protected by a single SSL certificate. They are added to generate certificates for new nodes or additional VLAN IPs to be added in the future.

Considerations while deploying ECA:

- All certificates for communication should be obtained from a common trusted CA. Auto Image Replication (AIR) between MDSPs that uses different external CAs is not supported but you can concatenate the individual root CA certificates into one file and upload them as a CA bundle.
- After ECA is deployed on the cluster, you can renew or update the ECA.
- It is recommended to pause backup/restore operations before starting ECA deployment/renewal.
- The CA bundle and CRL file independent of other security artifacts.
- When you deploy security artifacts, they are validated and if inconsistencies are found, you are notified, and deployment does not proceed. If you provide an external certificate and CA certificate bundle, the EC certificate is validated against the user provided CA certificate bundle. If only one of the items is provided, it is validated against deployed artifacts.
- Only NetBackup Certificate Authority (NBCA) + ECA deployment is supported in this release.
- You cannot revert to NBCA deployment once NBCA + ECA deployment is done.

- NBCA renewal has to be run manually from the management node after NBCA is revoked.
- NBCA is auto renewed 60 days before expiration. You are notified 60 days before the expiration of the ECA certificates. An alert appears on the appliance GUI and an email is also sent. If NBCA renewal fails, user is notified and AutoSupport alert is raised.
- You can deploy external certificate only if all NetBackup Flex Scale components are up and running. These include NetBackup primary and media services, storage engines, management gateway, and NetBackup Flex Scale management web services.
- You cannot deploy security artifacts, if upgrade, add node or VLAN operation is in progress and vice versa.
- If the ECA's subject alternative names have information on new nodes (FQDNs) to be added, add node operation succeeds seamlessly and all services come up after the add node operation. If subject alternative names are not updated, add node operation fails.
- For Nutanix, HBase workloads using SSL certificates, append the respective SSL certificates to the CA bundle after ECA certificates are renewed. If you do not append the SSL certificates to the CA bundle during ECA renewal, backup and restore operations for the workloads may fail.

Deploying ECA using the GUI

You can perform all external certificates related operations from the **Settings > Security Management > Certificates** tab.

- Upload certificate
- View certificate request
- Generate certificate request

To deploy ECA using the GUI

- 1 Go to the **Settings > Security Management > Certificates** tab.
- 2 Click **Generate certificate request** and fill out the form. The SAN field is filled with default SAN entries that are mandatory for the configuration. For standard default configuration, it has the FQDN entries for all the storage servers, NetBackup primary FQDN, console IP and API gateway FQDN.

Settings / Security management

FIPS STIG Immutability **Certificates**

Upload certificate View certificate request Generate certificate request

Appliance certificate

Subject distinguished name
CN=r515-001vip62.vxindia.veritas.com, O=Veritas, OU=NBUI, L=Indore, ST=MP, C=IN

Issuer distinguished name
IntermediateCA

Valid from Sep 4 05:49:40 2021 GMT	Valid until Sep 9 05:49:40 2022 GMT	Days until expiration 360
RSA key size 2048	Certificate revocation list check FILE	Serial number 6A8B017ACCC26DA4596502965F8B39EAA89F1

Subject alternative names
[8 items](#)

Root certificate

- 3 Click **Generate** to generate a certificate request.

Generate external certificate request

Fill out the certificate details below, then send them to your internal certificate authority to get it signed into a certificate. Your certificate authority will return a system-wide certificate or CA bundle.

Certificate subject

Common name (CN) *
r515-001vip62.vxindia.veritas.com

Organization (O)
NBUI

Organizational unit (OU)
Veritas

Locality (L)
SC

State/Province (ST)
CA

Country/Region (C)
IN

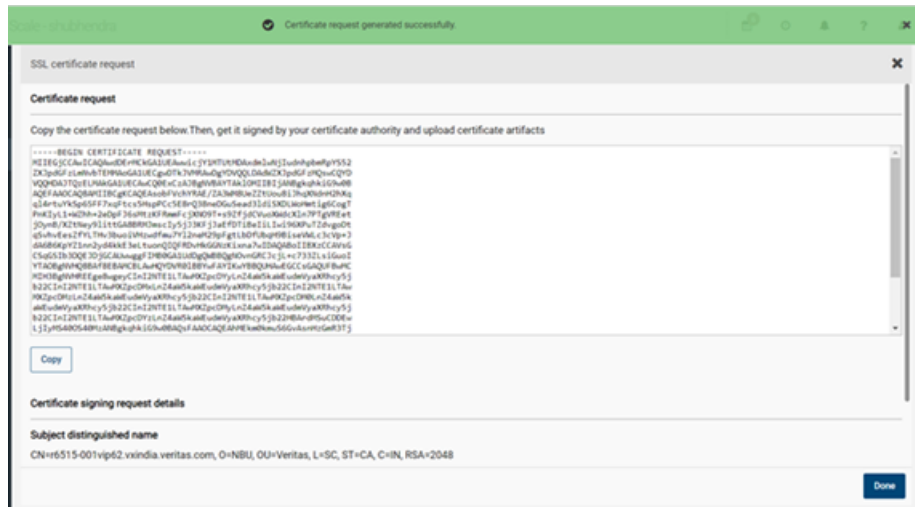
Subject distinguished name *
CN=r515-001vip62.vxindia.veritas.com, O=NBUI, OU=Veritas, L=SC, ST=CA, C=IN

RSA key size *
2048

Subject alternative names

Cancel Generate

- 4 Click **Copy** to copy the certificate request. This has to be given to a CA signing authority who uses this to generate a certificate. CA signing authority will provide a certificate along with the root certificate and intermediate certificate used to generate the certificate. Authority may also provide a CRL file.



- 5 Go to **Certificates > Upload certificate** to upload the certificate artifacts. For a fresh deployment, both the certificate and CA certificate bundle are mandatory. CRL is optional. For renewal, any one of the three are required.

- 6 After the certificate is uploaded, **Save** gets enabled. Click **Save**. The deployment is initiated.

After deployment is successfully completed for all components, the GUI restarts with the uploaded external certificate.

You can verify the browser certificate from the GUI to verify that it is the same one that was used during deployment.

Log locations

The configuration file can be found at location:

`/shared/security_artifacts/config.json`. It contains information about CSR, current key size, and whether ECA is deployed and the details of the nodes on which it is deployed.

The deployed artifacts can be found at the location:

`/shared/ec_certificate-artifact/`.

All the logs pertaining to ECA deployment are present in the `/log/VRTSnas/log` directory:

- nbfs_deploy_certificate.log: It contains the logs of the driver script which is responsible for performing pre-checks and then calling subsequent scripts to deploy ECA on various components.
- ec_validate.log: It contains the logs for the ec_validate script responsible for performing validations on the artifacts.
- nbu_nbca_certificate_deploy.log: It contains the logs for the NBCA deployment on NetBackup components.
- nbu_eca_certificate_deploy.log: It contains the logs for ECA deployment for NetBackup components.
- isagui_deploy_mgmt_cert.log: It contains the logs for deployment of ECA on NetBackup Flex Scale GUI and API gateway.
- nbfs_deploy_cert_on_node.log: It contains the logs when ECA is internally deployed on newly added or replaced node in add/replace node workflow if ECA is already deployed.

Considerations for performing other operations when ECA is deployed

If ECA is deployed and you want to add a new node to the cluster:

- Before starting an add node operation, generate a new certificate request which has the new node's storage FQDN as an additional SAN entry.
- Upload the new certificate and then add the new node. Otherwise, the add node operation will fail.

If ECA is deployed on a cluster where disaster recovery is configured and you want to add a new node to the cluster:

- The ECA has to be renewed on the primary irrespective of whether you add the new node on the primary or secondary cluster.
- The new certificate must contain FQDN of the storage server of the node which is going to be added.
- A new CSR has to be generated which contains the FQDN of the new storage server as a SAN entry. A new certificate is generated with the new CSR which is used for ECA deployment/renewal.


Troubleshooting

This chapter includes the following topics:

- [Services management](#)
- [Collecting logs for cluster nodes](#)
- [Checking and repairing storage](#)
- [Troubleshooting NetBackup Flex Scale issues](#)

Services management

You can use the **Settings > Services management** tab to manage your storage services.



When there is a device failure or a new device is added, the device information may not be consistent. In such a scenario, you can **Run full discovery** to update the database with latest system settings and usage data.

You can fix service-related faults in the nodes of the cluster that are in running state. If a service has faulted because of issues in services such as the volume manager, NLM, deduplication, or the management console, you can **Run auto fix** and bring

the service online after the issue has been resolved. No other operation should be initiated when you run AutoFix.


You can also call the `POST/api/appliance/v1.0/management/autofix` RESTful API to run auto fix.

Collecting logs for cluster nodes

You can collect logs for an individual node in the cluster or for all the cluster nodes for error analysis and troubleshooting. You can choose to collect logs for all the components to get a comprehensive view of the system, or you can collect logs for specific components that have an issue.

- 1 Use any one of the following options to log in using the user account that you created:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.
- 2 In the left navigation pane click **Settings**, and then click **Diagnostics**.

3 Click **Generate log package**.



On the Generate log package page, on the Basic tab you can select components or on the Advanced tab you can select individual log files to include in the log package.

4 To select the components for which you want to collect logs:

- On the **Basic** tab under **Log settings**, select the components. To select all the components, click **All**.

Component	Description
Service status	Status of the services running on the cluster nodes
OS	Kernel and user-level debug logs.
Infrastructure	Generic product information, including stack trace for running daemons, policy-based storage provisioning logs, and API gateway logs.
Explorer	Logs and system environment data collected by the <code>VxExplorer</code> utility.
Initial deployment	Initial cluster configuration logs. Installation logs are located in the <code>/opt/VRTS/install/logs/</code> directory.
NetBackup master service	Logs for NetBackup server services and components. This component is displayed only if both NetBackup primary and media servers are deployed in the cluster.

Component	Description
Appliance	Hardware-specific logs. Includes HPE Active Health System (AHS) logs. Note: If High Security mode is enabled in the HPE iLO Remote Console, AHS logs cannot be downloaded locally and are not included in the log package. To download the AHS logs, log in to the iLO Remote Console. Click Information > Active Health System Log. On the Active Health System Log page click Download.
NetBackup media service	NetBackup media server logs from nodes.
Upgrade	Upgrade logs.
NetBackup MSDP	Deduplication engine logs for all nodes.
NBSU	Information collected using the NetBackup Support Utility tool.

- Under **Package options**, specify a name for the log package and optionally the case number if provided by Veritas Support.
- Under **Nodes**, select the nodes for which you want to collect the logs.

Note: If you switch from Basic to Advanced view (or vice versa), all the options that you specified until that point are lost and you will have to specify them again.

5 To select individual log files:

- On the **Advanced** tab under **Log settings** expand the nodes and select the log files.

- Under **Package options**, specify a name for the log package and optionally the case number if provided by Veritas Support.

6 Click **Generate.**

To view the progress of the operation, click **View details** on the **Diagnostics** page. If you had specified the case number, it is prefixed to the package name. The generated log package is displayed under Packaged logs. You can now download the package to your node or upload it to Veritas Support for error analysis.

See “[Downloading logs](#)” on page 175.

See “[Uploading logs to Veritas Support](#)” on page 174.

Uploading logs to Veritas Support

You can share the logs with the Support personnel to help you diagnose problems.

1 Use any one of the following options to log in using the user account that you created:

- Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface

`https://ManagementServerIPorFQDN/webui` where

ManagementServerIPorFQDN is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.

- Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console

`https://ManagementServerIPorFQDN:14161` where

ManagementServerIPorFQDN is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.

2 In the left navigation pane click **Settings, and then click **Diagnostics**.**

The existing log packages are displayed under **Packaged logs**.

3 Click the package that you want to upload to and click **Send.**

On the **Send logs** dialog box, click one of the following options and click **Send**.

- **Send to Veritas using call home:** If Call Home is configured on your appliance, logs are uploaded with TLS-encryption using HTTPS protocol to the Veritas AutoSupport server.

- **Send using SFTP:** Use the SSH File Transfer Protocol to upload the logs. Specify the following details to use SFTP:

Field	Description
Server address	Remote server name
Port	Port number. By default the SFTP server is configured to listen to port 22.
User name	User name to authenticate with the server
Password	Password to authenticate with the server
Path	Location to upload the logs

Downloading logs

If a node is a part of the cluster, you can download log packages from the node. For information about how to collect logs and create log packages, see [Collecting logs for cluster nodes](#)

- 1 Use any one of the following options to log in using the user account that you created:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.
- 2 In the left navigation pane click **Settings** and then click **Diagnostics**.
 The log packages that are already created are displayed under **Packaged logs**.

- 3 Click the log package that you want to download and click **Download**.
- 4 On the **Download logs** dialog box, click **Download**.

Checking and repairing storage

You can check storage for errors and perform consistency checks and repair the storage from the NetBackup Flex Scale UI if the storage is corrupted or degraded. The storage check and repair operations are time consuming operations and can impact the system performance.

Warning: Do not run the repair storage operation without the guidance from Veritas Support. Always contact Veritas Support if you run into any storage related issues. Running the repair storage operation without any guidance from Veritas support may lead to improper system behavior.

- 1 Use any one of the following options to log in using the user account that you created:
 - Use a user account with both Appliance Administrator and NetBackup Administrator role, or a user account with only an Appliance administrator role to log in to the NetBackup Flex Scale web interface
`https://ManagementServerIPorFQDN/webui` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration. In the left pane click **Cluster Monitor > Infrastructure**, in the upper-right corner click **Cluster dashboard**, and when prompted click **Open cluster dashboard**.
 - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management console
`https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address or the FQDN that you specified for the NetBackup Flex Scale management server during the cluster configuration.
- 2 On the Dashboard in the **Services** area click **View details**.
- 3 On the NetBackup page click **Processes**.
- 4 To check storage:
 - Click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Check storage**.
 - Review the message and when prompted for confirmation, click **Continue**.

The progress of the operation is displayed on top of the page and a notification is displayed after the operation is complete.

- A report containing the details and result of the operation is generated and can be viewed as a log package by clicking **Settings > Diagnostics**.

5 To repair storage:

- For the row that shows as unhealthy in the **Status** column, click the Actions menu (vertical ellipsis) from the right side of the row and click **Check storage**.
- Review the message and when prompted for confirmation, click **Continue**. The progress of the operation is displayed on top of the page and a notification is displayed after the operation is complete.
- A report containing the details and result of the operation is generated as a log package and can be viewed by clicking **Settings > Diagnostics**.

Troubleshooting NetBackup Flex Scale issues

This section describes common issues that you might encounter when using NetBackup Flex Scale and possible solutions for those issues.

If cluster configuration fails (for example because an IP address that was already in use is specified) and you try to reconfigure the cluster, the UI displays an error but the configuration process continues to run

Cause

If the license key was already registered when the cluster configuration failed and you try to reconfigure the cluster, an error registering the license key is reported in the UI but the configuration process continues to run. The error is displayed in the UI because the configuration process tries to install the same license key which was installed previously during the cluster configuration. However, the configuration process is not terminated and continues to run.

Solution

Before you reconfigure the cluster, delete the license key from all the nodes using the following steps:

- 1 Log on to each node using the following credentials:

User: admin

Password: P@ssw0rd

The NetBackup Flex Scale Appliance Shell Menu is displayed.

- 2 To access the root shell type the `support elevate` command.

To access the root shell when lockdown mode is configured, See [“Selecting or changing the lockdown mode”](#) on page 159.

- 3 Run the following command to delete the license key:

```
rm /etc/vx/licenses/lic/Access_Perpetual.slf
```

Validation error while adding VMware credentials to NetBackup

For protecting VMware virtual machines, you add the login credentials of the VMware vCenter Server, VMware ESXi server, or VMware vCloud Director server to the NetBackup configuration. While adding the access credentials, you may see the following credential validation error in the NetBackup Java Console UI.

```
VMware credential validation failed. Cannot connect on socket (Status
25)
```

Cause


This issue occurs when the NetBackup Flex Scale appliance cluster is configured using a pure IPv6 address configuration while the VMware system uses IPv4 addresses. NetBackup Flex Scale appliance does not support communication between a pure IPv6 and a pure IPv4 address configuration.

Solution

To resolve this issue, the VMware system must be configured to use either a pure IPv6 address or a mixed mode dual stack IP address configuration.

NetBackup Web UI incorrectly displays some NetBackup Flex Scale processes as failed

The **Activity Monitor > Daemons** tab in the NetBackup Web UI displays the status of the NetBackup processes running in your NetBackup environment. Sometimes, the status of the NetBackup Flex Scale processes such as `spad` and `spoold` may appear as failed even though the processes themselves are running fine in the cluster.



Solution:

There is no solution at the moment. It is safe to ignore the incorrect status displayed in the NetBackup Web UI.

Unable to create BMR Shared Resource Tree (SRT) on NetBackup Flex Scale Appliance

Cause

This issue occurs because configuration of Bare Metal Restore (BMR) boot server is not supported on the NetBackup Flex Scale Appliance. The following error message is displayed:

```
Failed to setup loop device: No such file or directory.
Cannot mount media, please try again.
```

Solution

You can create a platform specific boot server separately based on the client OS platform which needs to be restored.

NetBackup configuration files are not persistent across operations that require restarting the system

The `VmTools.cfg` and `vixDiskLib.ini` files do not persist after a node is restarted as a part of an operation such as upgrading or replacing a node.

Workaround:

Use the `cp-nbu-config` command to copy the NetBackup configuration files from the user's home space to the specified NetBackup configuration destination directory. A NetBackup administrator can use the `cp-nbu-config` command to create and edit a NetBackup touch configuration file in any of the following directories:

- `/usr/opensv/netbackup`
- `/usr/opensv/netbackup/bin`
- `/usr/opensv/java`
- `/usr/opensv/lib/ost-plugins`
- `/usr/opensv/netbackup/bin/snapcfg`
- `/usr/opensv/netbackup/db/cloudSnap/credential`
- `/usr/opensv/netbackup/db/cloudSnap/proxy`
- `/usr/opensv/netbackup/db/config`
- `/usr/opensv/netbackup/db/event`
- `/usr/opensv/netbackup/db/images`
- `/usr/opensv/netbackup/db/media`
- `/usr/opensv/netbackup/ext/db_ext`
- `/usr/opensv/netbackup/ext/db_ext/db2`
- `/usr/opensv/var`
- `/usr/opensv/volmgr`
- `/usr/opensv/volmgr/database`

To create or edit a touch configuration file

- 1 Log in to the node.
- 2 Create a new configuration file in the NetBackup administrator home directory, or use the `cp` command to copy an existing configuration file from its original location to the home directory. For example:

```
cp /usr/opensv/lib/ost-plugins/pd.conf ~/
```

- 3** Make changes to the file in the home directory.
- 4** Run the following command to install the file in its original directory or a supported destination directory:

```
sudo /opt/veritas/vxapp-manage/cp-nbu-config configuration-file destination
```

where *configuration-file* is the file that you created or edited, and *destination* is the directory where it needs to be installed.

```
sudo /opt/veritas/vxapp-manage/cp-nbu-config ~/pd.conf  
/usr/openv/lib/ost-plugins
```

Configuring NetBackup optimized duplication

This appendix includes the following topics:

- [Configuring a Storage Lifecycle Policy for optimized duplication](#)

Configuring a Storage Lifecycle Policy for optimized duplication

This section outlines the process for the Storage Lifecycle Policy (SLP) to perform the following operations.

- A backup to MSDP storage
The backup images created by the backup operation are retained for one month.
- A duplication of the second backup image to a secondary MSDP pool
The copies of the backup images that are written to another MSDP pool (remote) which is retained for three months.

For more information, see *Creating a storage lifecycle policy* section of the *NetBackup Deduplication Guide*.

Creating a Storage Lifecycle Policy for optimized duplication

Locate and double-click the NetBackup Administration Console shortcut located on the desktop of the system, to launch the NetBackup Administration Console. Then, enter the host name, user name, and password in the Administration Console login screen to connect to your primary server.

To create a Storage Lifecycle Policy

- 1 Select **NetBackup Management > Storage > Storage Lifecycle Policies** in the left pane of the Administration Console. The list of currently configured Storage Lifecycle Policies is displayed in the right pane of the Administration Console.

Note: By default, no Storage Lifecycle Policies are configured.

- 2 Select **Actions > New > Storage Lifecycle Policy**. The **New Storage Lifecycle Policy** window is displayed.
- 3 Enter the information provided in the table in the **New Storage Lifecycle Policy** window.

Storage Lifecycle Policy Parameter Value

Storage lifecycle policy name	NYC_1month_to_LON_2weeks
Data classification	<no data classification>
Priority for secondary operations	0 (default)
State of secondary operation processing	Active (default)

Note: Enter a meaningful name for SLP in the format *<source site/location>_<retention>_to_<destination site/location>_retention*. For example, *NYC_1month_to_LON_2weeks*.

- 4 Click **Add** to add the first operation to the SLP. The **New Operation** window is displayed.

- Enter the information provided in the table to complete the configuration of the first storage operation for the SLP.

New Operation properties	Value
Storage source	--- (Default setting)
Operation	Backup (Default setting)
Destination storage	nyc_disk_stu You can change this setting to reflect your STU.
Volume pool	N/A
Media owner	N/A
Retention type	Fixed (Default setting)
Retention period	1 month (Retention level 3) You can change this setting.

- Click **OK** to save the **New Operation** window parameter settings.
- Click **Add** to add a second operation to the SLP. The **New Operation** window is displayed.

- Use the information provided in the table to complete the configuration of the second SLP operation.

New Operation properties	Value
Source storage	nyc_disk_stu (Backup) (Default)
Operation	Duplication (Default)
Destination storage	lon_disk_stu You can change this setting.
Volume pool	Not selected(default)
Media owner	Not selected(default)
Retention type	Fixed (default)
Retention period	2 weeks (level 1) You can change this setting.
Alternate read server	None (default)
Preserve multiplexing	Not selected(default)
Postpone creation of this copy until the source is about to expire	Not selected(default)

- 9 Click **OK** to save the **New Operation** window parameter settings.

Verify that the **Duplication** operation entry is added to the **Operation** list of the *NYC_1month_to_LON_2weeks* Storage Lifecycle Policy.

The figure illustrates the *NYC_1month_to_LON_2weeks* SLP with two operations included.

Operation	Window	Target Master	Storage	Volume Pool	Media Owner	Retention Ty...	Retentio
Backup	--	--	nyc_disk_stu	--	--	Fixed	1 month
Duplication	Default_24x...	--	lon_disk_stu	--	--	Fixed	2 weeks

Note: The left margin of the Duplication operation is indented from the entry for the first backup operation which indicates that the Duplication operation is derived from the second backup operation.

- 10 Click **OK** in the **Change Storage Lifecycle Policy** window to save the changes to the *NYC_1month_to_LON_2weeks*SLP, and to close the **Change Storage Lifecycle Policy** window.

- 11** Verify that entry for the *NYC_1month_to_LON_2weeks* storage lifecycle policy is visible in the list of SLPs that is displayed in the right pane of the Administration Console.

- 12** Repeat steps 4 to 11 to create another SLP called *LON_1month_to_NYC_2weeks* to backup at London and duplicate to NYC. This time replace step 5 with the following for the backup operation.

New Operation properties	Value
Storage source	--- (Default setting)
Operation	Backup (Default setting)
Destination storage	lon_disk_stu You can change this setting to reflect your STU.
Volume pool	N/A
Media owner	N/A
Retention type	Fixed (default)
Retention period	1 month (Retention level 3) You can change this setting.

For the second operation (duplication) in the SLP, choose the following

New Operation properties	Value
Source storage	nyc_disk_stu (Backup)
Operation	Duplication (Default)
Destination storage	nyc_disk_stu You can change this setting.
Volume pool	Not selected(default)
Media owner	Not selected(default)
Retention type	Fixed (default)
Retention period	2 weeks (level 1) You can change this setting.
Alternate read server	None (default)
Preserve multiplexing	Not selected(default)
Postpone creation of this copy until the source is about to expire	Not selected(default)

Configuring a policy to use an SLP

You can configure a policy that uses the *NYC_1month_to_LON_2weeks* SLP to automatically perform the operations specified in the SLP.

To configure a policy to use an SLP

- 1** Select **NetBackup Management > Policies** in the left pane of the Administration Console.
- 2** Select **Actions > New > New Policy**.
- 3** Type the policy name, test-backup-policy in the **Add a New Policy** dialog box, and click **OK**. The **Change Policy** window is displayed.
- 4** Use the information provided in the table to configure the policy Attribute settings.

Attributes parameter	Value
Policy name	test-backup-policy
Policy type	Standard (OS of primary server)
Policy storage	NYC_1month_to_LON_2weeks (scroll to locate this selection)
All other policy attributes	Use default settings – do not change

- 5** Click the **Schedules** tab of the policy.
- 6** Click **New** to add a new schedule to the policy. The **Add Schedule** window is displayed.
- 7** Configure the Schedule Attributes using the information provided in the table.

Schedule parameter	Value
Schedule name	full_slp
Type of Backup	Full Backup (default)
Calendar	Not selected (default)
Frequency	1 week (default)
Override policy storage selection	Not checked (default)
Media multiplexing	1 (default)

- 8 Click **OK** to save the *full_slp* schedule.

Note: No start window has been set up for this schedule. Only manual backups can be run using this schedule.

- 9 Verify that the entry for the *full_slp* schedule is displayed in the **Schedules** tab.
- 10 Click the **Clients** tab of the *test-backup-policy* policy.
- 11 Add a NetBackup Client hostname for the **Client name** and check **Detect client operating system** or select the correct **Hardware and operating system**.

Note: Use the method of your choice to enter the hardware and operating system of the client.

- 12 Verify that an entry for *<master server>* is added to the **Clients** tab.
- 13 Click the **Backup Selections** tab of the policy.
- 14 Add an entry for */etc* to the **Backup Selections** of the policy.
- 15 Verify that the */etc* entry is displayed in the list of **Backup Selections** of the policy.
- 16 Click **OK** to save the *test-backup-policy* policy.
- 17 Verify that the *test-backup-policy* policy is included in the list of policies displayed in the **All Policies** pane of the Administration Console.
- 18 Select **NetBackup Management > Policies** in the left pane of the Administration Console.
- 19 Right-click the entry for the *NYC_1month_to_LON_2weeks* policy in the **All Policies** pane, and then select **Actions > Manual Backup**.
- 20 Click **OK** in the **Manual Backup** window to initiate the operation.
- 21 Immediately click **Activity Monitor** in the left pane of the Administration Console. Note the most recent job entries in the Activity Monitor that use the *NYC_1month_to_LON_2weeks* policy.

Updating the policy to reverse the replication direction

In the event of a disaster or failover, there may be a need to reverse the replication direction by using predefined SLPs, if they are pre-created and available. In this

example, the backup policy called *test-backup-policy* is updated to use the *SLP LON_1month_to_NYC_2weeks* instead of *NYC_1month_to_LON_2weeks* which reverses the direction of replication or duplication.

To update the policy to reverse the replication direction

- 1 Select **NetBackup Management > Policies** in the left pane of the Administration Console.
- 2 Select the policy *test-backup-policy* on the right pane of the Administration Console.
- 3 Right click *test-backup-policy* and select **Change**. Use the information provided in the table to configure the policy Attribute settings. The **Change Policy** window is displayed.

Attributes parameter	Value
Policy type	Standard (OS of primary server)
Policy storage	LON_1month_to_NYC_2weeks (scroll to locate this selection)
All other policy attributes	Use default settings – do not change

Disaster recovery terminologies

This appendix includes the following topics:

- [VVR technology in disaster recovery](#)
- [About response fields in the GET disaster recovery API](#)

VVR technology in disaster recovery

In NetBackup Flex Scale, the NetBackup catalog is replicated using Veritas Volume Replicator (VVR). VVR is a block level replication software. A consistency group of volumes is created on the source location and remote location. In VVR terminology, consistency group is known as Replicated Volume Group (RVG). VVR replicates data from a primary RVG on the primary site to the secondary RVG on the secondary site. VVR maintains write-order fidelity to guarantee consistent and recoverable copy of data on secondary site.

RVG consists of following components:

- RLINK: Establishes the replication link between primary and secondary RVG.
- SRL: Storage Replicator Log is a circular buffer of all the writes for an RVG. SRL enables VVR to maintain write-order fidelity at the secondary site.
- Data Volumes: Group of related volumes.
- DCM: Data Change Map is a component of VVR which is used to track application writes when the SRL overflows. This enables VVR to avoid complete resynchronization of the data on the secondary. DCM is only active on the primary site.

About response fields in the GET disaster recovery API

The important replication fields displayed under attributes in response section for GET : /api/appliance/v1.0/disaster-recovery are:

- **state:** Displays the state of the Primary RVG.
The following table lists the values for the RVG state field and their meanings.

acting_secondary	This Primary RVG is currently the acting Secondary as part of the fast failback process. Writes to the data volumes in this RVG are disabled independent of whether the RVG is started or stopped.
disabled for I/O	Primary RVG is disabled for I/O, that is, the RVG is stopped.
enabled for I/O	Primary RVG is enabled for I/O, that is, RVG is started.
needs recovery	State after an import or reboot.
passthru	The Primary RVG is in <code>passthru</code> mode because the Primary SRL is detached or missing.

- **dataStatus:** Shows the data status of the Secondary.
The following table lists the values for the Data status field and their meanings:

consistent, behind	Secondary data is consistent but not up-to-date with the Primary data.
consistent, stale	The data on this Secondary is consistent. Replication to this Secondary has been stopped; the Primary RLINK is detached.
consistent, up-to-date	The Secondary data is consistent and is current or up-to-date with the Primary data. The Primary role can be migrated to this Secondary.
inconsistent	The data on the Secondary volumes is not consistent and the Secondary cannot take over.
needs recovery	State after an import or reboot. The <code>vxrlink recover</code> command clears this state.

N/A Current state of the Secondary data cannot be determined. This may occur because of a configuration error on this Secondary.

- **replicationMode:** Displays the mode of replication. As the NetBackup catalog of NetBackup Flex Scale is configured for asynchronous replication, the value of this field is always asynchronous.
- **replicationStatus** Displays the status of the replication to the Secondary. The following table lists the values for the Replication status field and their meanings:

Value	Meaning
logging to DCM	DCM is active for this Secondary, that is, new updates on Primary are tracked using DCM for this Secondary. The following information may be displayed: ■ needs dcm resynchronization—To continue replication, resynchronize the Secondary using DCM resynchronization. ■ needs failback synchronization—To continue replication, start failback synchronization to this Secondary.
needs failback synchronization	This Primary RVG is acting as Secondary as part of the fast failback process. To continue replication, start failback resynchronization on the new Primary.
not replicating	Data is not being replicated to Secondary because Primary RLINK is in <code>needs_recovery</code> state.
paused by user	Replication to Secondary is paused because of some administrative action. This results in the following states: primary paused—Primary RLINK is paused. secondary paused—Secondary RLINK is paused.
paused due to error	Replication to Secondary is paused because of the following errors: ■ secondary config error—Secondary has some configuration error. ■ secondary log error—Secondary SRL has an I/O error.
paused due to network disconnection	Replication to Secondary is paused because of some network problem.
replicating	connected—Replication can take place if there are updates on the Primary data volumes

Value	Meaning
resync in progress	Resynchronization to the Secondary is in progress. autosync—Resynchronization type is autosync. dcm resynchronization—Resynchronization after an SRL overflow. failback resynchronization—Resynchronization using failback logging.
resync paused by user	Resynchronization to Secondary is paused because of some administrative action. This results in the following states: - primary paused—Primary RLINK is paused. - secondary paused—Secondary RLINK is paused.
resync paused due to error	Resynchronization to Secondary is paused because of the following errors: - secondary config error—Secondary has some configuration error. - secondary log error—Secondary SRL has an I/O error.
resync paused due to network disconnection	Resynchronization to Secondary is paused because of some network problem.
stopped	Replication to Secondary is stopped because of the following: - Primary detached—Primary RLINK is detached. - Secondary detached—Secondary RLINK is detached.
N/A	The replication status cannot be determined.

- **loggingTo:** Indicates whether updates for this Secondary are tracked on the Primary using the SRL or DCM.
The following table lists the values for the Logging to field and their meanings:

Value	Meaning
DCM (contains xxx Kbytes) (log_type)	DCM is active (in use) for the replication to this Secondary. log_type can be autosync, failback logging, or SRL protection logging.
SRL (xxx Kbytes behind, yyy % full)	Updates to be transferred to Secondary are logged into the SRL and are currently occupying xxx Kbytes or yyy% of the SRL
SRL	SRL is used for logging. Check the Data status field for the status of the Secondary data.

- **timeStampInformation** : Displays a timestamp to indicate the time by which secondary is behind.
- **configErrors** : Describes the configuration related errors for NetBackup catalog replication.

The following table lists the error messages and their meanings:

Message	Message definition
<i>host</i> : Pri or Sec IP not available or vradmind not running	The Primary IP or the Secondary IP address is not available, or the <code>vradmind</code> daemon on the host is not running or is running on a different port.
<i>host</i> : disk group missing.	Host <i>host</i> does not have any disk group with the same name as that specified in the <i>remote_dg</i> attribute, of the Primary RLINK pointing to this host.
<i>host</i> : RLINK missing.	Primary RVG has an RLINK to <i>host</i> , but <i>host</i> does not have corresponding rlink to this Primary RLINK.
<i>host</i> : RLINK dissociated.	Host <i>host</i> does have an RLINK corresponding to the Primary RLINK. However, it is not associated with the Secondary RVG.
<i>host</i> : disk-group mismatch.	The <i>remote_dg</i> attribute of either the Primary RLINK or Secondary RLINK is incorrect.
<i>host</i> : RLINK mismatch.	The <i>remote_rlink</i> attribute of either the Primary RLINK or Secondary RLINK is incorrect.
<i>host</i> : host mismatch.	The <i>local_host</i> and/or <i>remote_host</i> attribute of either the Primary RLINK or Secondary RLINK is incorrect.
<i>host</i> : Primary-Primary configuration.	The two Primary RVG RLINKs are pointing to each other. This situation will arise after the original Primary comes up after a Primary failover.
<i>host</i> : multiple Primary error.	The same Secondary RVG has more than one Primary RVGs.
<i>host</i> : two or more nodes on same host.	Two or more RVGs in the same RDS are located on the same host, <i>host</i> . This configuration is not supported.

Message

Message definition

No Primary RVG.

Ignore this error if it displays only for a few seconds and disappears. If the error persists, then the problem may be due to the Secondary not being able to determine its Primary because it has some configuration error, the Primary is not reachable, or it does not have any RLINK.

host: Primary and Secondary have same disk-group ID.

This condition happens in the cases when the split mirrored plexes of the Primary volumes are exported without using the `Disk Group split` option and then imported on the Secondary with force option.

host: unknown

The configuration status is currently unknown.

host: stale information

The configuration status may be stale.

host: no data volume

This error can occur if one of the RVGs in the RDS does not have a data volume associated with it.

host: network-protocol mismatch

The `protocol` attribute of the Primary RLINK is different from that of the Secondary RLINK.

host: VVR-heartbeat-port mismatch

The `local_port` attribute setting for the Primary RLINK is different from that of the Secondary RLINK.

host: unsupported VVR version in cross-version replication

The cross-version replication in VVR is only supported between two immediate major releases.

host: no contact from Primary

This error can occur if the `vradmin` server on the Secondary host is unable to establish contact with the `vradmin` server on the Primary. This can be because the Primary RVG of this RDS cannot be found or the `vradmin` server is not running or is unreachable.

host: vxconfigd disabled

The `vxconfigd` daemon is currently disabled on the host on which this error occurred.

host: volume-number mismatch

This error can occur if the number of volumes in the Primary and Secondary RVGs of the RDS are different.

host: volume-size mismatch

This error can occur if the sizes of some or all of the data volumes in the Primary and Secondary RVGs of the RDS do not match.

Message	Message definition
<i>host</i> : volume-name mismatch	This error can occur if some or all of the volumes in the Primary and Secondary RVGs of the RDS are not mapped correctly.
<i>host</i> : Primary SRL missing	This error can occur if the Primary SRL was disassociated from the Primary RVG or is missing.
<i>host</i> : Secondary SRL missing	This error can occur if the Secondary SRL was disassociated from the Secondary RVG or is missing.

Configuring Auto Image Replication

This appendix includes the following topics:

- [Auto Image Replication configuration](#)


Auto Image Replication configuration

The backups that are generated in one NetBackup domain can be replicated to storage in one or more target NetBackup domains. This process is referred to as Auto Image Replication.


You can configure Auto Image Replication in NetBackup Flex Scale.

To configure Auto Image Replication

- 1 Go to the NetBackup web UI and add a trusted primary server. Click **Next**.



- 2 In the source domain, get the MSDP_SERVER from the NetBackup web UI. Navigate to **Storage > Storage configuration > Storage servers**.




- 3 Add MSDP_SERVER in target primary server. In the target domain, logon to the target primary server using the following command:

```
ssh <appliance_admin user>@<master ip/hostname>  
(appliance_admin password)  
sudo bash  
echo "MSDP_SERVER = <Source MSDP server name>" >>  
/usr/openv/netbackup/bp.conf
```

YAML file:


```
user_management:  
  storage_server:  
    - password: P@ssw0rd1234  
      user_name: root  
  users:  
    - password: P@ssw0rd1234  
      roles:  
        - appliance_admin  
        - backup_admin  
      user_name: appadmin  
    - password: P@ssw0rd1234  
      roles:  
        - backup_admin  
      user_name: nbuadmin
```

- 4 Get the token from the target domain NetBackup web UI. Go to **Security** > **Token**. In the **Create token** window, enter the token name and other required details. Click **Create**.



5 Add replication targets for the disk pool.

In the **Disk pools** tab, click on the disk pool link.



Click **Add** to add the replication target.



6 In the **Add replication targets** window:

- Select the target primary server.
- Input the target domain token.
- Select the target volume.
- Input the target storage credentials. Refer to the YAML configuration file for details.

Add replication targets

Select trusted master server

Search...

Trusted master server

☒

n154-h24.cdc.veritas.com

1 Records (1 selected)

Trusted master server token

VEFPKPYJCYTWPSS

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
☐ n154-h16.cdc.veritas.com	PureDiskVolume	PureDisk
☒ n154-h16.cdc.veritas.com	volume-1	PureDisk

2 Records (1 selected)

Login credentials for the replication target storage server:

User name *

root

Password *

Cancel

Add

Click **Add**.