NetBackup™ Flex Scale Administrator's Guide

3.5.100

April 2026

VERITAS
now part of Cohesity

NetBackup Flex Scale Administrator's Guide

Last updated: 2026-04-30

Legal Notice

Copyright © 2026 VERITAS TECHNOLOGIES LLC All rights reserved.

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

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

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

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

Support/Service Assistance

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

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

Cohesity Software Running on Partner Hardware

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

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

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

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

Contents

Chapter 1

Product overview

This chapter includes the following topics:

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

Chapter 2

Viewing information about the NetBackup Flex Scale cluster environment

This chapter includes the following topics:

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. ■ View and manage NetBackup Flex Scale cluster. To view the cluster dashboard, click **Cluster Management > Cluster dashboard**. To view details about the NetBackup and NetBackup Flex Scale services, click **Cluster Management > Services**. To manage cluster settings, click **Cluster Management > Cluster settings**. ■ To view the NetBackup Flex Scale infrastructure such as the nodes, disks, and hardware details, in the left navigation pane click **Cluster Management > 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 Management > 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.

Table 2-1 *(continued)*

Interface and URL	To access

NetBackup Flex Scale infrastructure management UI
`IPv4: https://ManagementServerIPorFQDN:14161/`
`IPv6: https://[ManagementServerIP]:14161/`
`IPv6:https://ManagementServerFQDN:14161/`
`IPV4: https://console_ip_ipv4:14161/`
`IPv6: https://[console_ip_ipv6]: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 infrastructure management UI" on page 17.

NetBackup UI `https://primaryserverfqdn/webui/`	View and manage NetBackup. You can open NetBackup Flex Scale from the NetBackup UI when you click Appliance management in the left navigation pane.

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

Table 2-2

Interface and URL	To access

NetBackup Flex Scale infrastructure management UI
`IPv4: https://console_ip_ipv4:14161/`
`IPv6: https://[console_ip_ipv6]:14161/`
	You can view and manage NetBackup Flex Scale.
	See "About the NetBackup Flex Scale infrastructure management UI" on page 17.

Accessing the NetBackup web user interface on the appliance

When you deploy a cluster with both NetBackup primary and media servers, you can configure and manage your NetBackup environment from NetBackup Flex Scale or by accessing the NetBackup web UI directly.

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, the FQDN, or the short host name that you specified for the management server during the cluster configuration.

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

1. Use any one of the following options to sign in to the NetBackup Flex Scale infrastructure management UI using a user account with both Appliance administrator and NetBackup administrator role:

   - `https://ManagementServerIPorFQDN:14161` where *ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the management server during the cluster configuration.

     Note: If you access the infrastructure management UI by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

   - `https://console_ip:14161` where *console_ip* is the public IPv4 or IPv6 address that you specified during the cluster configuration.

2. In the left navigation pane, click **NetBackup**.

   This launches the NetBackup Flex Scale web UI in a separate browser tab and you are automatically signed out from the infrastructure management UI.

3. On the NetBackup Flex Scale web UI sign in page, specify the NetBackup administrator user account and password and then click **Sign in**.

   For more information about how to manage NetBackup, 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.

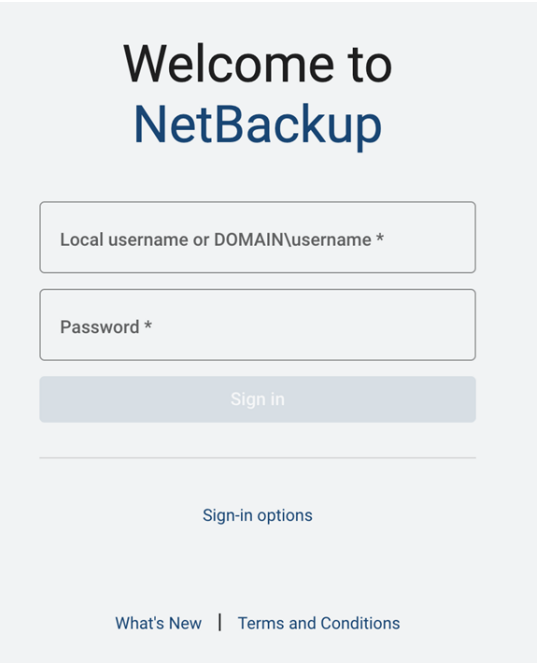

About the NetBackup Flex Scale web UI

The NetBackup Flex Scale web UI provides a centralized view to manage NetBackup and NetBackup Flex Scale. When you deploy primary and media servers in the cluster, you can use this UI to manage NetBackup tasks such as monitoring jobs and events, protecting workloads, and detecting malware and anomalies and to manage the NetBackup Flex Scale cluster and its components.

To access the NetBackup Flex Scale web UI, use the `https://ManagementServerIPorFQDN/webui` URL. Use a user account with NetBackup Administrator and Appliance Administrator role to sign in to the NetBackup Flex Scale web UI.

The following figure shows the NetBackup Flex Scale web UI:

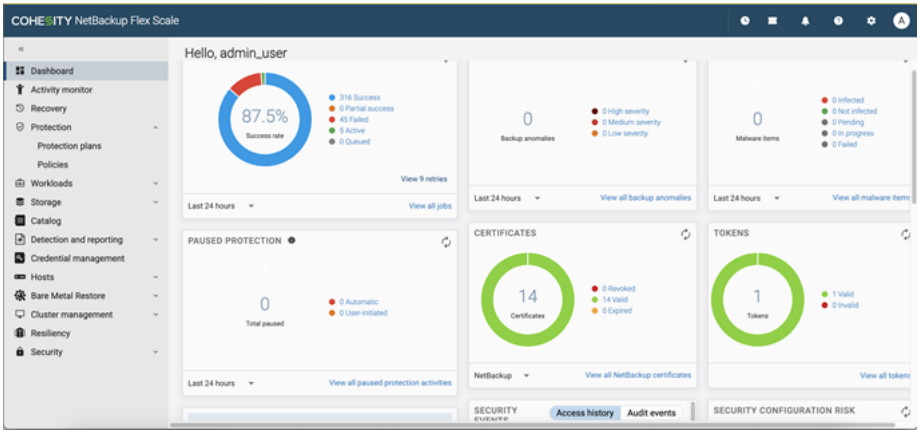

The dashboard displays an overview of NetBackup and NetBackup Flex Scale information. You can manage the cluster using the options that are displayed under **Cluster Management**.

To view cluster details such as storage utilization, infrastructure details, and security features, click **Cluster dashboard**:

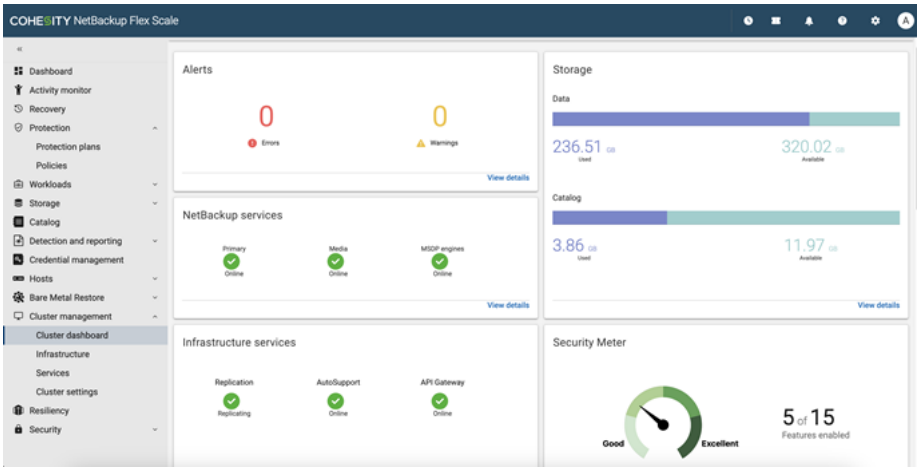

To view and manage cluster infrastructure such as nodes and disks, click **Infrastructure**:

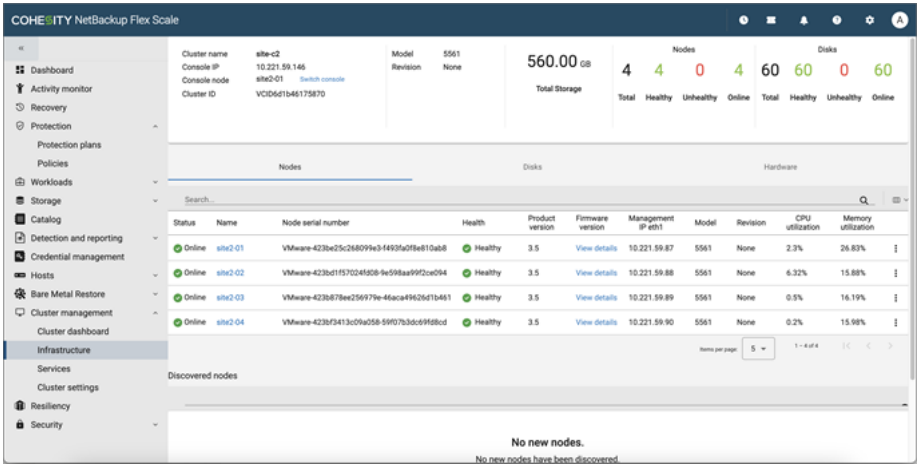

To view details about the services running in the cluster, click **Services**.

To configure and manage cluster settings such as security, primary service replication, and AutoSupport, click **Cluster settings**.

About the NetBackup Flex Scale infrastructure management UI

You can administer your NetBackup Flex Scale appliance cluster configuration using the NetBackup Flex Scale infrastructure management UI. This UI provides a centralized and a consolidate graphical view of your entire appliance configuration. You can use the infrastructure management UI to view all the cluster details such as the cluster nodes and their storage, alerts, and the services running on the appliance.

To access the NetBackup Flex Scale infrastructure management UI:

1. Use any one of the following options to sign in to the infrastructure management UI:
   - If both primary and media servers are deployed, use a user account with both Appliance administrator and NetBackup administrator role or a user account with only Appliance administrator role to sign in to one of the following interfaces:
     - `https://`*`ManagementServerIPorFQDN`*`:14161` where *ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the management server during the cluster configuration.

Note: If you access the infrastructure management UI by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

- `https://console_ip:14161` where *console_ip* is the public IPv4 or IPv6 address that you specified during the cluster configuration.
- Use the console FQDN if it was set during initial configuration.

- If only media servers are deployed, use a user account with an Appliance Administrator role to sign in to the following interface: `https://console_ip:14161` where *console_ip* is the public IPv4 or IPv6 address that you specified during the cluster configuration.

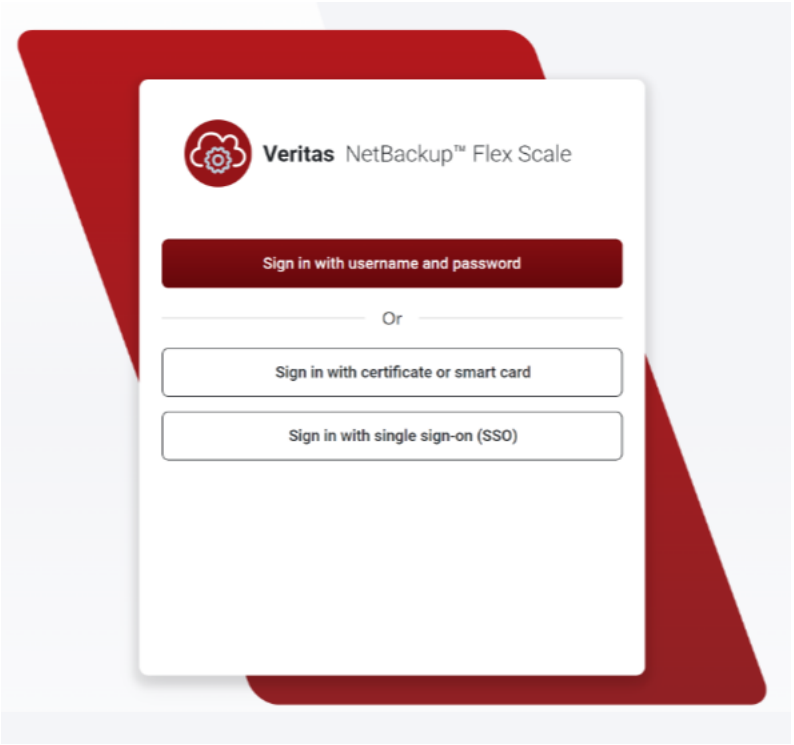

The **Dashboard** tab displays the various elements of your cluster configuration.

You can use the infrastructure management 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.
- Managing security for cluster data.
- Monitoring the appliance performance and data reduction levels.

2. If both the primary and the media servers are deployed in the cluster, click **NetBackup** in the left navigation pane to view and manage NetBackup.

About the Dashboard view

The dashboard in NetBackup Flex Scale infrastructure management console provides an overview of the appliance cluster configuration and its components and services. Use the infrastructure management 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:

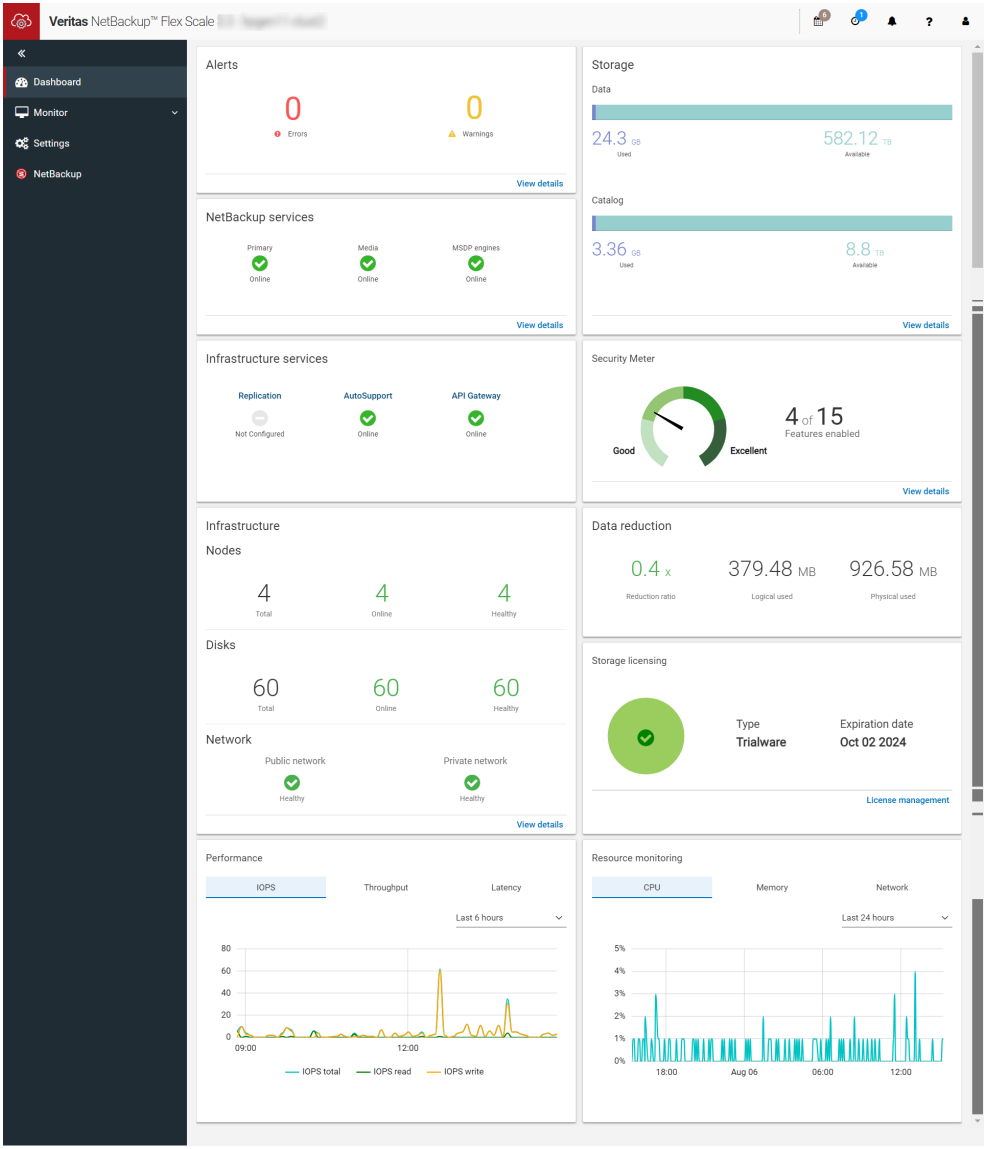

The following figure shows the dashboard for a media only cluster:

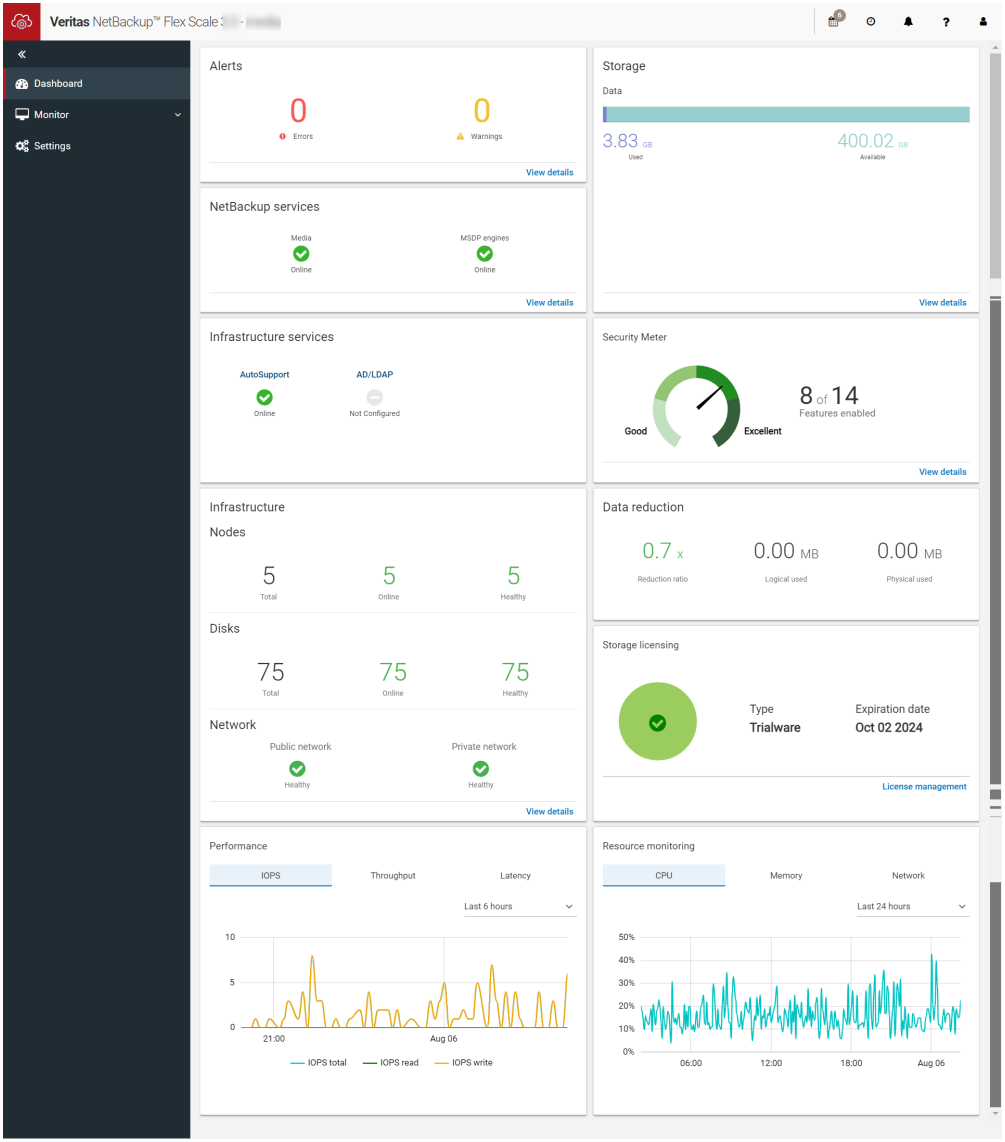

The Dashboard view displays the following information:

- 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 173.

- The Security meter shows the security status of the appliance and offers recommendations to improve the appliance security.
- The NetBackup 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 primary server, media servers, and MSDP engines. If only media servers are deployed, the NetBackup services area includes services such as the media servers and MSDP engines.
- The Infrastructure services area displays the status of the cluster services such as replication, AutoSupport, and API gateway.
- 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 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration.

Note: If you access the NetBackup Flex Scale infrastructure management UI by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration.

  Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

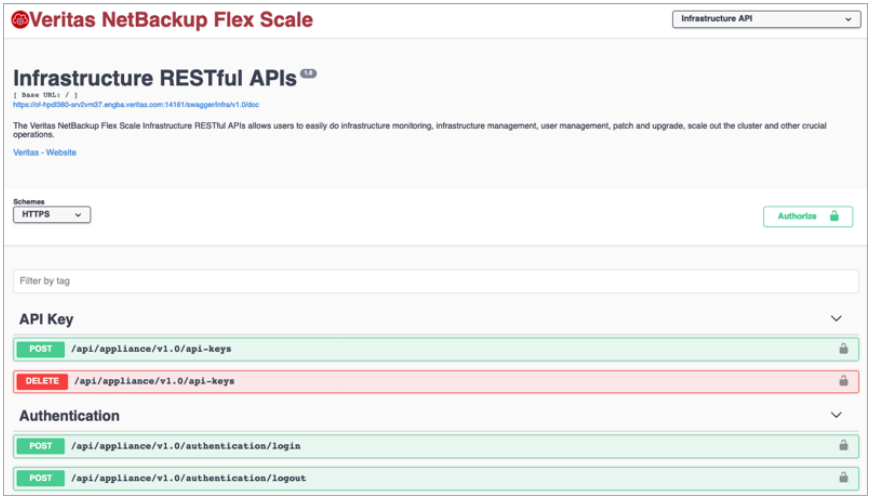

- For a quick introduction on how to use the APIs, refer to the getting started readme available at the following URL:

  `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

Chapter 3

NetBackup Flex Scale infrastructure management

This chapter includes the following topics:

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 *(user role)*	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 *(user role)*	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.
Universal share user *(user role)*	The Universal share user role has the permissions to manage Universal shares/Instant Access in user mode. This role is applicable to only local users except restricted users such as *maintenance*.

Table 3-1 NetBackup Flex Scale users *(continued)*

User role / user	Description
maintenance *(user)*	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: `maintenance` ■ Password: `P@ssw0rd` This is the default password and you must change the password to a non-dictionary password during cluster configuration.

See "Considerations for managing NetBackup Flex Scale users" on page 27.

See "Adding users" on page 30.

See "Removing users" on page 34.

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 during cluster configuration. 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 both appliance administrator and universal share user roles are available post cluster configuration.

Note: When disaster recovery is configured between two NetBackup Flex Scale clusters, Veritas recommends that you add local user accounts on both clusters using the same credentials. This is to ensure that the same credentials work when the NetBckup primary is failed over between the clusters.

- 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 in a NetBackup Flex Scale cluster on which both primary and media servers are deployed, the NetBackup primary server service must be running and healthy in the cluster.

- If both primary server and media servers are deployed on the cluster, AD and LDAP users cannot access the appliance SSH, cluster-level CLI and REST

APIs. If only media servers, are deployed, AD/LDAP users having appliance administrator role can access SSH, GUI, cluster-level CLI and REST APIs.

Note: If you have deployed the cluster with only media servers, the appliance administrator role is assigned to AD/LDAP users using the NetBackup Flex Scale Infrastructure Web UI.

- Local users that have the NetBackup Flex Scale appliance administrator role assigned have access to the cluster-level CLI 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 the NetBackup Administrator and the NetBackup Flex Scale Appliance Administrator role should 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.
  If domain name was provided during AD/LDAP configuration, then you can also use the domain names for assigning roles.
- 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 25.

See “Adding users” on page 30.

See “Removing users” on page 34.

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 27.

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 infrastructure management UI” on page 17.

   See “About the NetBackup Flex Scale web UI” on page 15.

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:

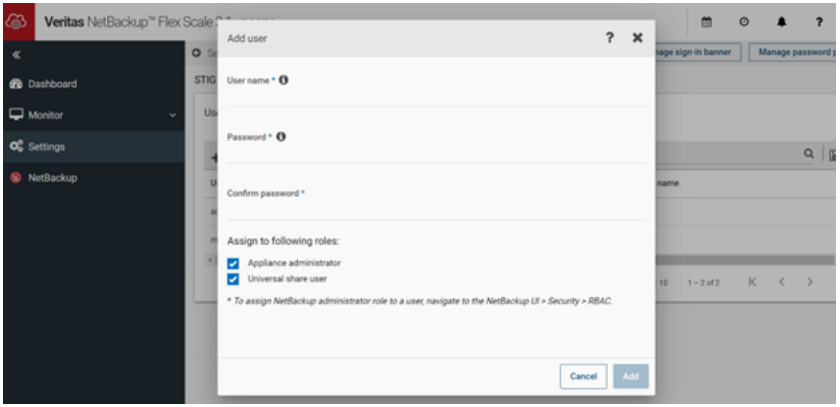

Parameter	Description
User name	Enter the username. If the user has only Appliance administrator role, the following criteria apply: ■ The username can contain a maximum of 31 characters. ■ The username must not start with a number. ■ The username can include only underscore (_) and dash (-) as special characters. ■ The username must not be "root", "primary", or "maintenance". If the user has both Appliance administrator and Universal share user roles or only Universal share user role, the following criteria apply: ■ The username can contain a minimum of 4 characters. ■ The username can contain a maximum of 30 characters. ■ The username can contain only lowercase characters. ■ The username must not start with a number. ■ The username can include only underscore (_) and dash (-) as special characters. ■ The username must not be "root", "primary", or "maintenance".
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 !#$%&+,-.<=>@[]^_{|}~ ■ 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 219.
Confirm password	Enter the password again to confirm.

Parameter	Description
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.
Assign 'Universal share user' role	Select this option to assign the Universal share user role to the user account. This user role has the permissions to manage Universal shares/Instant Access in user mode. **Note:** The 'Universal share user' role is applicable to only local users.

Note: You cannot assign the NetBackup Administrator role from this UI. You must use the NetBackup web UI to manage NetBackup role assignments to user accounts.

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.

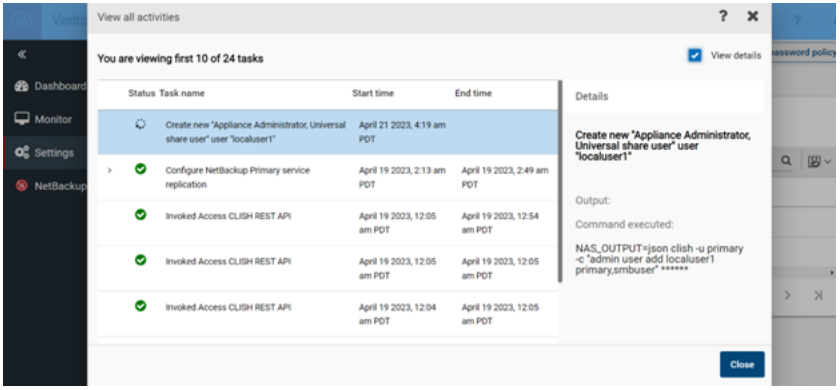

Alternatively, you can also click the **Recent Activity** icon from the top right of the UI to see the status of the most recent operations occurring in the cluster.

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.

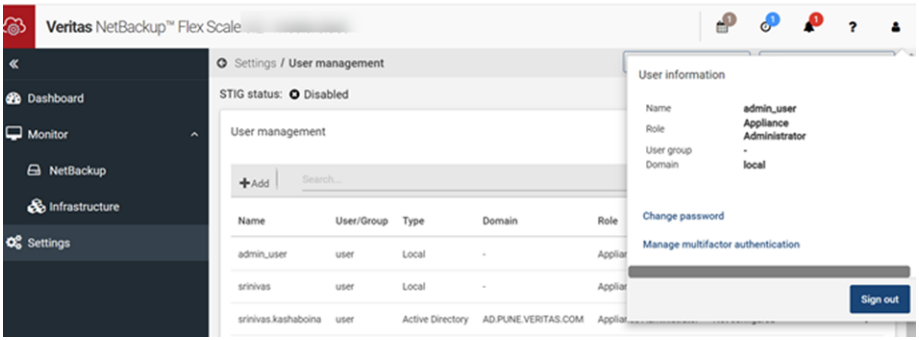

- The password of users with *Appliance administrator* role can also be changed by navigating to **Settings > User management**. Navigate to the user row and click on the vertical ellipsis button from the right side of the UI and then select **Change password**..

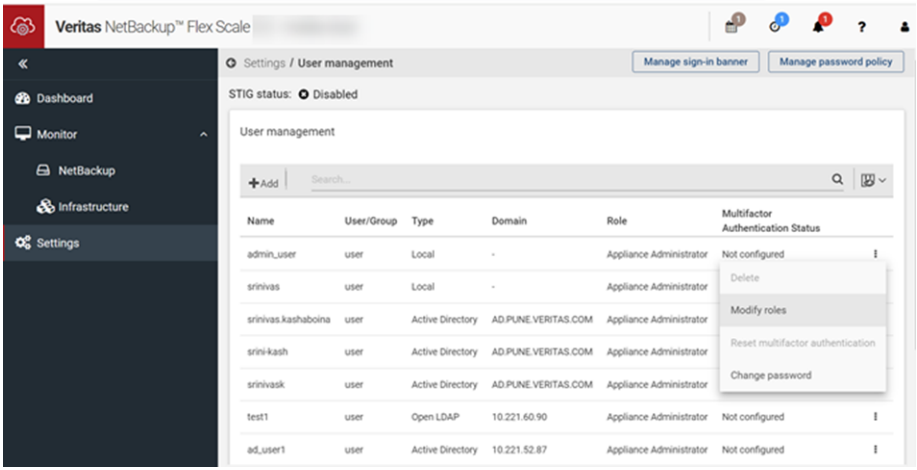

- The password of users who have no role or have only *NetBackup administrator* role or only *Universal share user* role can be changed by logging on to the NetBackup Flex Scale UI. As the user does not have the *Appliance administrator* role, the user is directed to the **Change password** tab of the NetBackup Flex Scale UI where the password can be changed.
- To change the maintenance user password, sign-in with the maintenance user credentials to the NetBackup Flex Scale UI. Click **Change password** option.
- 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 27.

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 infrastructure management UI” on page 17.

   See “About the NetBackup Flex Scale web UI” on page 15.

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**.

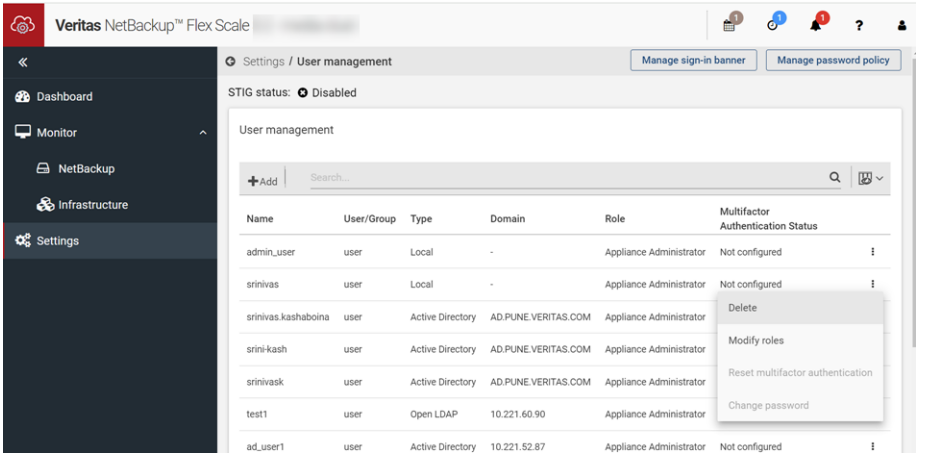

4 On the **Delete user** dialog box, click **Delete** to confirm the user removal.

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.

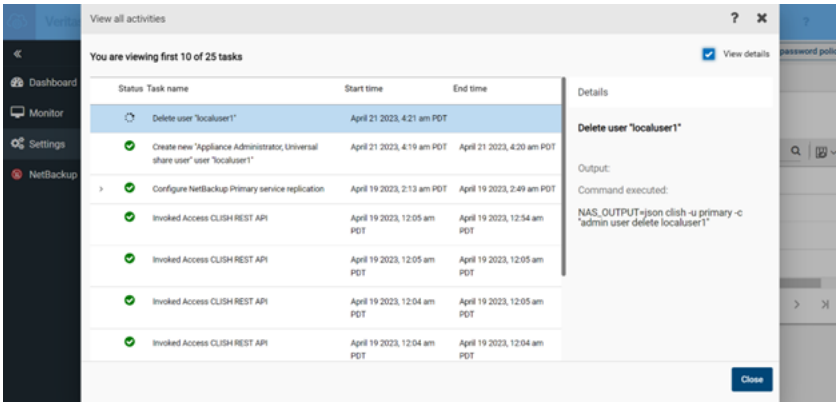

Alternatively, you can also click the **Recent Activity** icon from the top right of the UI to see the status of the most recent operations occurring in the cluster.

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.

Modifying user roles

Perform the following steps to modify user roles 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.

Note: *Maintenance* users cannot modify user roles.

See “Considerations for managing NetBackup Flex Scale users” on page 27.

To modify a user role from the NetBackup Flex Scale cluster

1. Sign in to the NetBackup Flex Scale infrastructure management console UI.

   See “About the NetBackup Flex Scale infrastructure management UI” on page 17.

   See “About the NetBackup Flex Scale web UI” on page 15.

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 modify user roles from a particular user account, in the user row, click on the vertical ellipsis button from the right side of the UI and then select **Modify roles**.

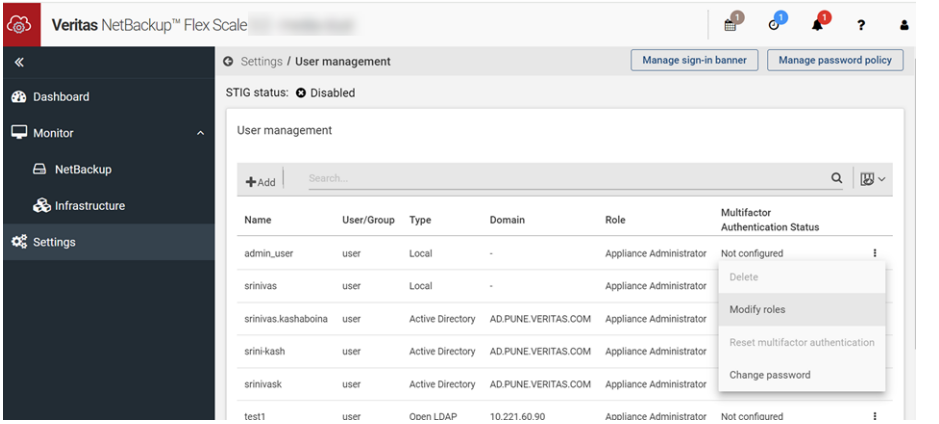

4 On the **Modify roles** dialog box, select the check box to assign the required role or clear the check box to unassign the role. You are required to provide a password if you want to assign the Universal user role to a user. Click **Save** to save the changes.

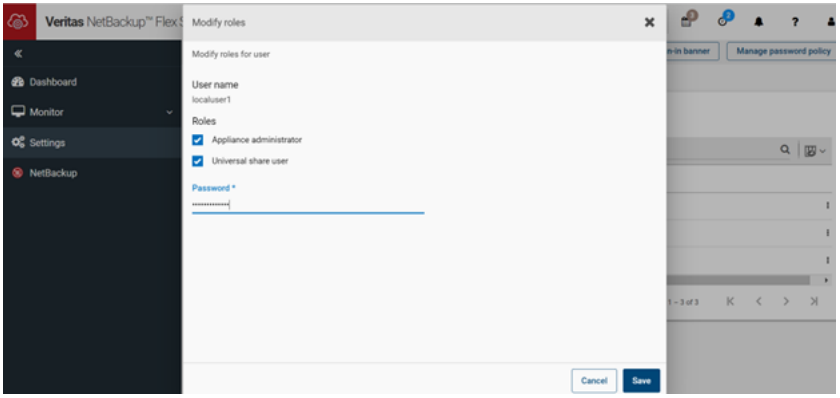

The UI displays a message that confirms that the user role modification operation is triggered. You can click **View Details** to see the progress of the operation.

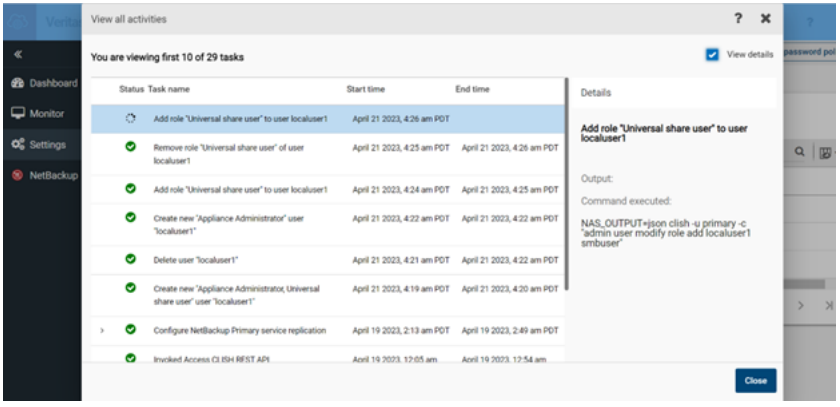

Alternatively, you can also click the **Recent Activity** icon from the top right of the UI to see the status of the most recent operations occurring in the cluster.

5 Once the operation completes, verify that the user role you modified appears correctly in the list of users displayed in the **User management** pane.

Considerations for configuring AD/LDAP

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 disaster recovery is configured between two NetBackup Flex Scale clusters, the AD/LDAP configuration and management must be done from the cluster on which NetBackup primary service is running

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 3.5.100 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*

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

- In a deployment with both primary and media servers:

- Role assignment for AD/LDAP users should be done from the NetBackup web UI.
- Domain users should use the [<AD/LDAP_serverIP|FQDN|domainname>\username] to login to NetBackup UI. Domain name can be used only if it was provided during AD/LDAP configuration.
- 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.
- Do not create domain user with same name as local user as creating duplicate users may cause ambiguity.
- Do not configure more than one AD/LDAP server with the same domain as it may cause ambiguity in fetching and displaying information about domain users.
- AD/LDAP users and groups should have UID and GID assigned in the AD/LDAP server. Role assignment does not work for domain users who do not have UID or GID. If multiple AD/LDAP servers are configured, make sure that none of the servers have conflicting UIDs or GIDs.

- In a deployment with only media servers:
  - Domain users should use the [<AD/LDAP_serverIP|FQDN|domainname>\username] to login to NetBackup Flex Scale UI. Domain name can be used only if it was provided during AD/LDAP configuration.
  - 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.
  - Do not create domain user with same name as local user as creating duplicate users may cause ambiguity.
  - Role assignment is not allowed for a domain user who has a space character in the username. For an AD user, the username is considered as the logon name.

- Do not configure more than one AD/LDAP server with the same domain as it may cause ambiguity in fetching and displaying information about domain users.

- 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 Services management** tab.
2. Click **Configure**.

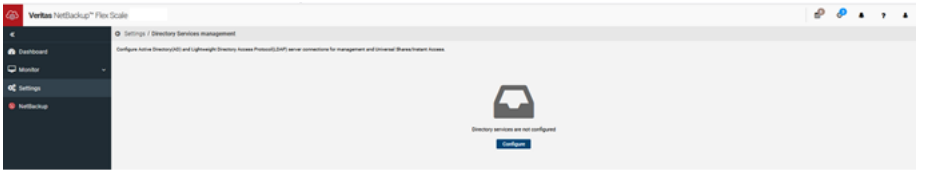

3. Select **Manage universal shares/Instant Access**. Provide the AD server details. Click **Save**.

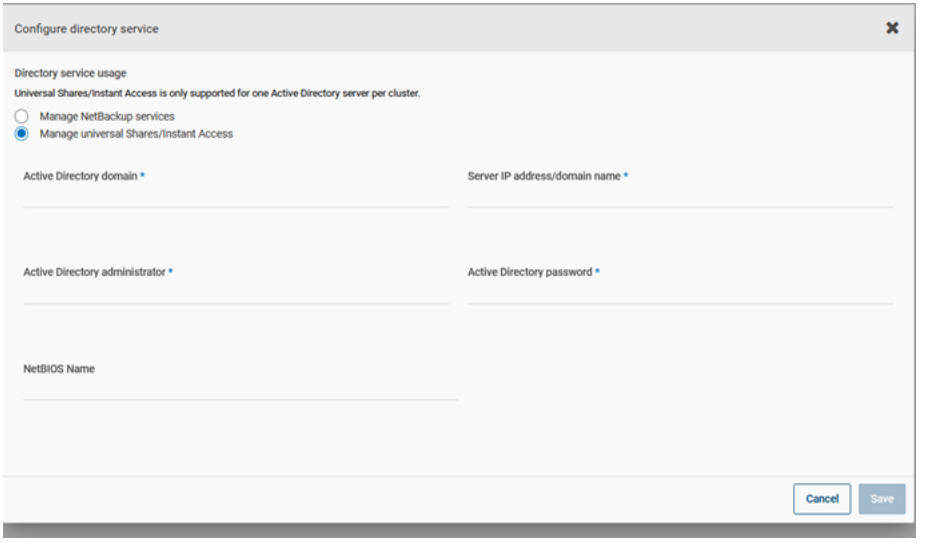

- The administrator of the AD server should have Domain Admin privilege.
- If the AD domain is sub domain, make sure that the domains can be trusted, or configure the AD domains with the type as **Manage NetBackup services** in NetBackup Flex Scale.

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.

Note: If the AD server is not available and you want to access the Universal Shares, you can access the Universal Shares if a local user has Universal share user role. To add an Universal share user role, See “Adding users” on page 30.

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 Services management** tab. Click **Configure**.

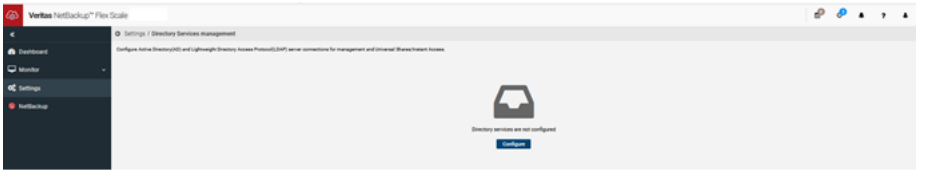

2 Select **Manage NetBackup services**.

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

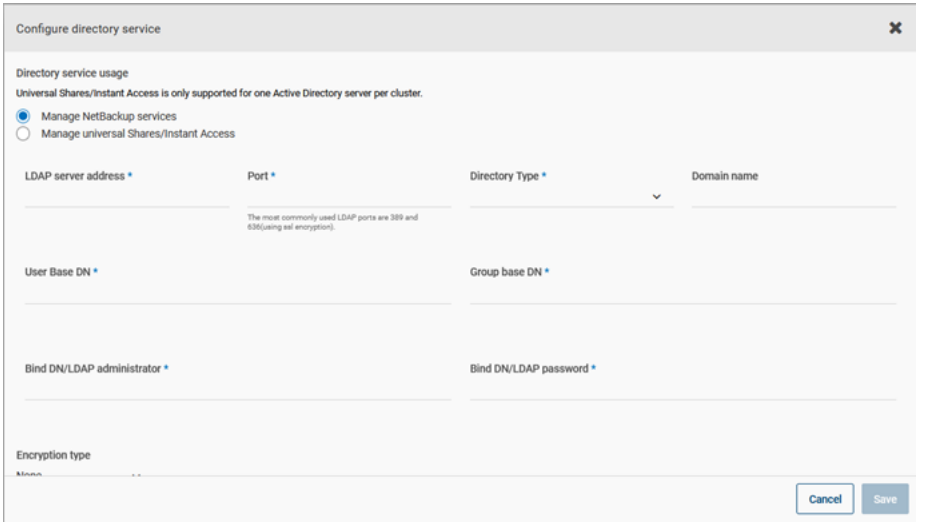

Bind DN/LDAP administrator is the domain name of the administrative user or any user that has search permission to the user container, or user subtree as specified by **UserBaseDN**.

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.

Note: In a NetBackup Flex Scale deployment with both primary and media servers, role assignment for AD/LDAP users should be done from the NetBackup web UI. For more details, see the *NetBackup™ Web UI Administrator's Guide* on SORT.

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 Services 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 cluster-level CLI or the node-level CLI is launched. If you SSH to eth1 (management) IP of any node, the node-level CLI gets launched. If you SSH to console IP, the cluster-level CLI 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 Services 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**.

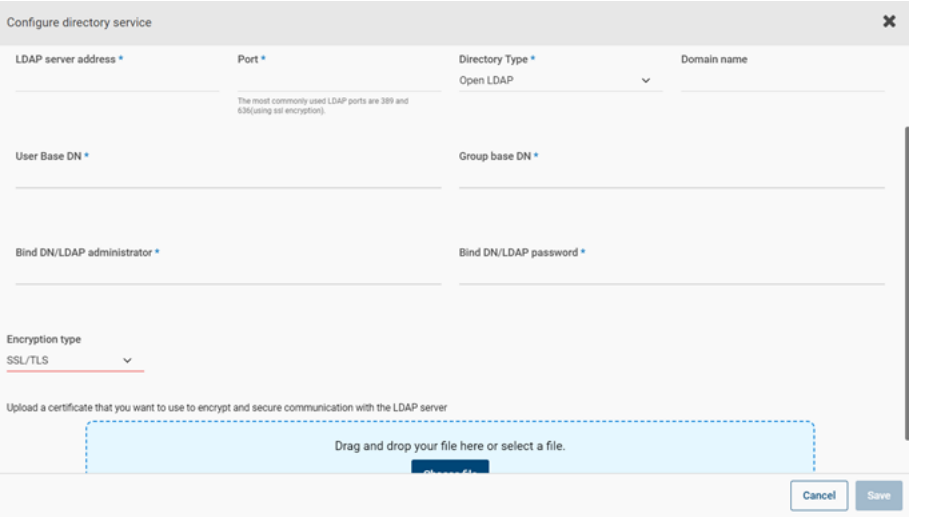

3. Once the AD/LDAP server configuration operation is successful, the AD/LDAP server(s) appears in the list with usage as Appliance management.

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

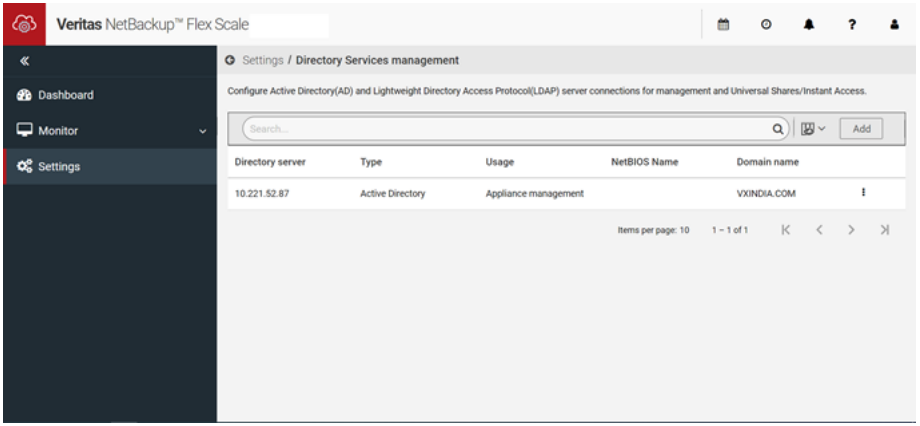

On the same page, there is an option to remove the domain configuration.

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@IP|FQDN|domainname*. For example, *testuser1@10.1.1.10* or *testgroup@mydomain.example.com*. Domain name can be used only if it was provided during AD/LDAP configuration.

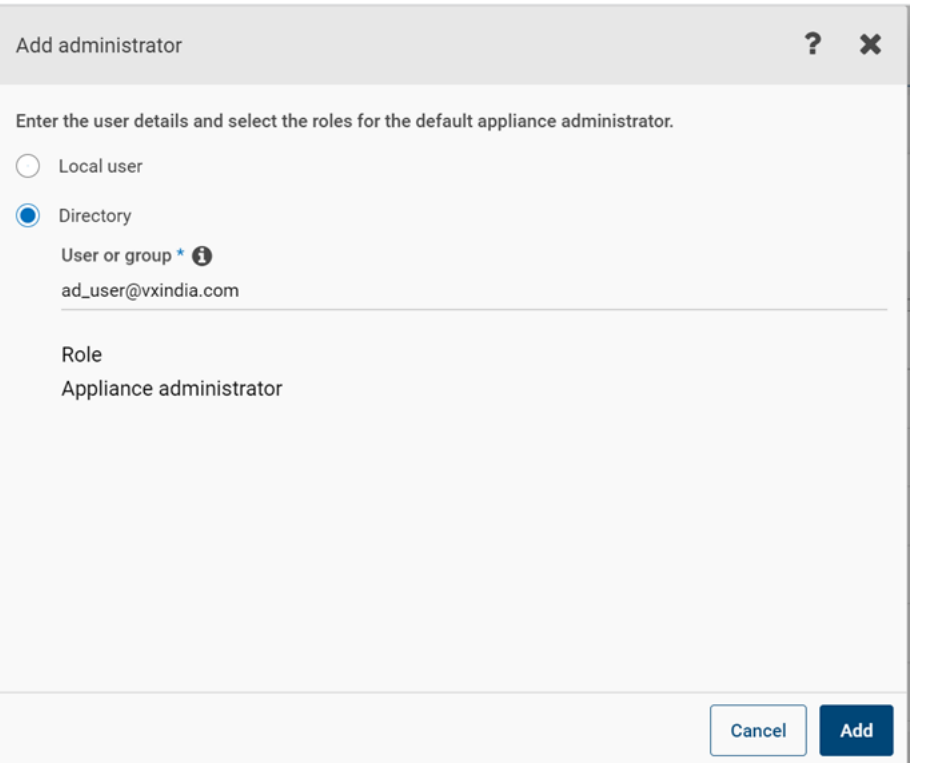

4 In the **Add administrator** dialog box, enter the details. Click **Add**.

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.

The domain name information is also provided.

To remove a particular user role, See "Modifying user roles" on page 36.

Directory services and certificate management

The LDAP communication between NetBackup Flex Scale and AD/LDAP server is not secure by default. You can secure this traffic by using SSL/TLS.

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 39.

Note: To configure AD/LDAP server with SSL, use unencrypted certificate file. Encrypted certificate is not allowed.

Note: If AD/LDAP is configured with `SSL=on`, only FIPS-compliant encryption is supported for a deployment with only media servers.

To configure AD/LDAP

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

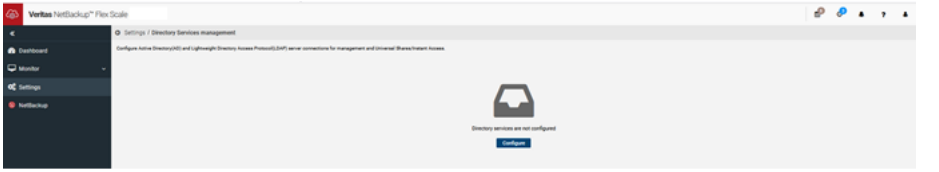

2 Enter the inputs required.

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)
Domain name	Specify the domain name.
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 theAD/ 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.

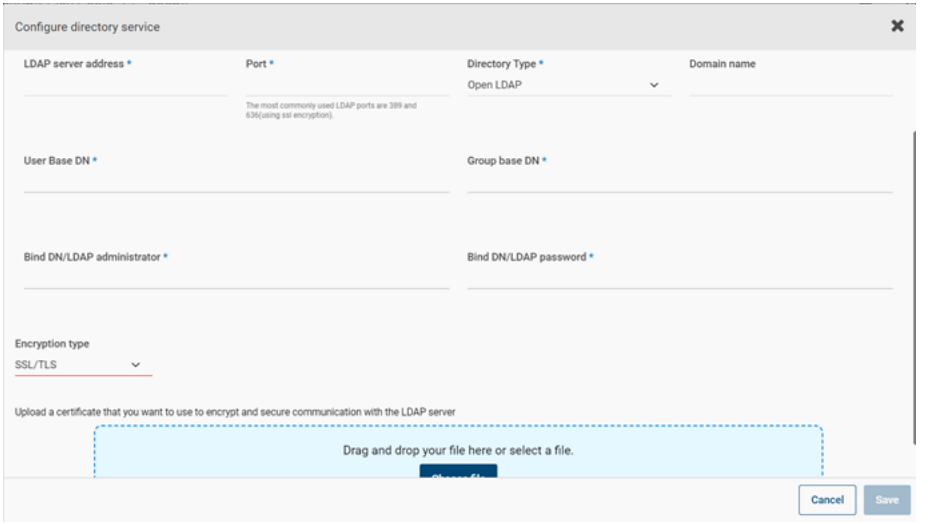

3 Click **Configure**.

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

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 16TB and 20TB	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 6^{th}, 9^{th}, 12^{th} or the 15^{th} 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 is the public IP address, the FQDN, or the short host name 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.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

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

Data storage and catalog storage level alerts are generated if usage exceeds a duration of 6 hrs.

Alerts are generated:

- If the data storage usage (including all data file systems) is more than 90% of usable data storage capacity.
- If the catalog storage usage (including NetBackup catalog file systems) is more than 90% of usable catalog storage capacity.
- If NetBackup file system usage is above 90% for all other file systems including CAT_FS and SCRATCH_FS file systems.

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 > Disk storage > Universal Share** and then select the universal share to view its details.

3. Copy your application data to the Universal Share.

4. Use the NetBackup Administration Console or the NetBackup web UI to 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.

Cloud bucket support for NetBackup Flex Scale

A single Flex Scale cluster supports 2 PB cloud bucket capacity with a maximum of 1 PB of data in a single cloud bucket. To optimize data transfer to a cloud tier, you can set the **UsableMemoryLimit** parameter to **85**. The **UsableMemoryLimit** parameter specifies the maximum usable memory size in percentage. The parameter is included in the `/usr/openv/pdde/pdcr/etc/contentrouter.cfg` file.

To update the parameter:

1. Login to engine container using the `docker exec -it container-id bash` command, where *container-id* is the container ID for image **uss-engine**.
2. Open the /usr/openv/pdde/pdcr/etc/contentrouter.cfg file and set the value of the **UsableMemoryLimit** parameter to **85**.
3. Restart the deduplication (MSDP) services from the MSDP shell. SSH to the engine IP address by using the appliance admin user. Stop and start the MSDP services by using the following commands:

   ```
   dedupe MSDP stop
   ```

   ```
   dedupe MSDP start
   ```

   You can verify the status of the services by using the `dedupe MSDP status` command.

For assistance, you can contact Veritas Technical Support.

Backing up data to Data Domain storage

You can back up data to Data Domain storage using the OpenStorage (OST) plug-in. To enable Data Domain as a storage server, you need to install the Data Domain OST as an EEB on the NetBackup Flex Scale cluster.

Contact Veritas Technical Support if you want to enable Data Domain storage for NetBackup Flex Scale.

To use Data Domain as a storage server:

1. Install the Data domain OST as an EEB.

   For information about how to install the EEB, see the *Installing EEBs using GUI* section of the *Veritas NetBackup™ Flex Scale Installation and Configuration Guide*.
2. Ensure that the EEB is installed successfully and all the required services such as primary, media, and storage services are running. You can view the status of the services on the NetBackup Flex Scale UI **Dashboard**.

3. Add IP and FQDN entries for the Data Domain storage server in the NetBackup Flex Scale GUI. Navigate to **Settings > Network > Custom Hosts** to add the entry.
4. From the NetBackup Administration Console, you can now create a storage unit to identify the Data Domain storage. For information about creating storage units in NetBackup, see the *NetBackup™ Administrator's Guide*.

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.

Firmware compatibility

The new node that you want to add must have the same firmware version for all the components as the cluster nodes. The firmware version of the cluster node components and the new node must be the same. For details about updating the firmware version, refer the *To upgrade the firmware using the rolling method* procedure in the *Updating the firmware in NetBackup Flex Scale clusters* section of the *Veritas NetBackup™ Flex Scale Installation and Configuration Guide*.

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 62.

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.

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 either a fully qualified domain name (FQDN) or a short host name 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 either an FQDN or a short host name for the MSDP engine. 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 MSDP engine.
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 69. 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, the FQDN, or the short host name 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 Management > 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, the FQDN, or the short host name 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**.

   Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

   ---

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

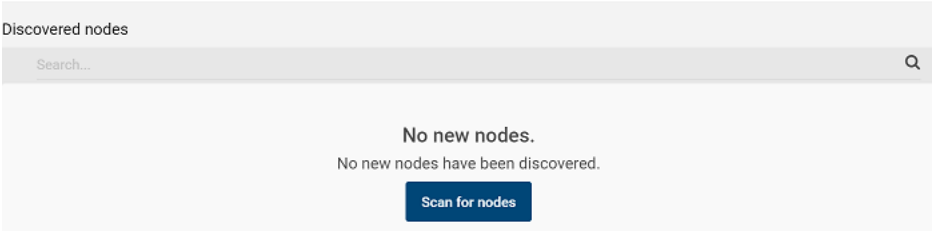

   The nodes connected to the private network of the cluster are discovered and displayed under the **Discovered nodes** section with details such as the node name, storage capacity, serial number, status, and the data network and the management network of the node.

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.

If a node is unhealthy because a vendor package RPM is missing, you can add that node. The vendor package is installed during the add node process.

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 may affect the performance of 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.
  The serial number of the node is displayed, which helps to identify which node is assigned the specified node name.
  By default, the auto-generated node name based on cluster and domain name is displayed.
- **Data-Network-<device-name>**:
  You can specify the IP address and either a fully qualified domain name (FQDN) or a short host name, or you can specify only the IP address for the media server and the MSDP engine.
  The FQDN or the short host name of the media server and the MSDP engine can contain a maximum of 64 characters, including the domain name.

  Note: The node name and serial number are displayed so you can identify the nodes and assign a specific IP address to each of the media server and MSDP engine.

  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 or the short host name for the media server and the MSDP engine. The FQDN is calculated based on the specified IP addresses.

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

- **Management Network**:

  You can 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.

  The serial number of the node is displayed so you can identify the node and assign a specific IP address to each of the management interfaces of the node.

  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 or the short host name. The FQDN is calculated based on the specified IP addresses.

  Click the **Custom** tab to specify the IP address and either a fully qualified domain name or a short host name 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 reimage 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.
- Add the new node(s) independently on each cluster.
- Catalog file system is not grown as part of add node operation.
- The NetBackup catalog filesystem is grown on both the sites after `add node` operation completes on both clusters. Both clusters have the same number of nodes and can 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration.

If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

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 23. 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 unit 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. To erase all the data from the disks, use the `system storage erase-disks configure` command and specify the pass algorithm to use to overwrite the disks with a digital pattern. For more details about this command, see the *Veritas NetBackup Flex Scale Command Reference Guide*.

Firmware compatibility

The new node that you want to use as a replacement node must have the same firmware version for all the components as the cluster nodes. The firmware version of the cluster node components and the new node must be the same. For details about updating the firmware version, refer the *To upgrade the firmware using the rolling method* procedure in the *Updating the firmware in NetBackup Flex Scale clusters* section of the *Veritas NetBackup™ Flex Scale Installation and Configuration Guide*.

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 62.

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, the FQDN, or the short host name 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 Management > 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, the FQDN, or the short host name 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**.

   Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

   ---

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**.

Note: If the node is stopped or excluded, use the iLO remote console to power off the 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Cluster Management > Infrastructure**.
   - Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management UI `https://`*`ManagementServerIPorFQDN`*`:14161` where *ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.

   Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

   ---

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 services**.

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

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 services**.

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

The node status changes to healthy.

Rebooting a node

You can reboot a cluster node. If you reboot a management console node (the node where the NetBackup Flex Scale UI is running), the UI fails over to another cluster node that is online and healthy. It can take a few minutes for the UI to be up on another node.

To identify the management console node, in the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure**. On the**Infrastructure** page, **Console node** shows the management console node.

To reboot a node:

1 Use any one of the following options to log on 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 on 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, and then in the left pane click **Cluster Management > Infrastructure**.
- Use a user account with an Appliance Administrator role to log on 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, and then in the left pane click **Monitor > Infrastructure**.

2 Click the **Nodes** tab.

The list of cluster nodes is displayed.

3 To reboot the node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI, and click **Reboot Node**. When prompted for confirmation, click **Reboot Node**.

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

Adding an excluded node to the cluster

A node is excluded from the cluster if the node reboots or panics multiple times. NetBackup Flex Scale monitors the cluster nodes and if a node reboots or panics five times, the node is excluded from the cluster. For a node that is excluded, the health of the node changes to unhealthy and node status is shown **Excluded** on the **Infrastructure > Nodes** page of the NetBackup Flex Scale UI. All cluster services are stopped on an excluded node, similar to a stop node operation.

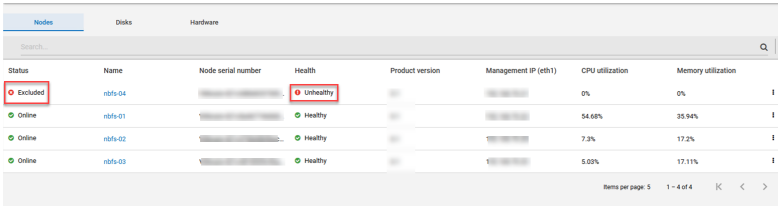

Note: NetBackup Flex Scale automatically detects and sets the node status to **Excluded** if there are multiple occurrences of a node rebooting or going in a panic state. You do not have the option to set the node status to **Excluded**.

You can include the excluded node back to the cluster. When you include an excluded node, all services that are stopped on the node are restarted.

To include an excluded node:

1 Use any one of the following options to log on 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 on to the NetBackup Flex Scale web UI`https://`*`ManagementServerIPorFQDN`*`/webui` where *ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Cluster Management > Infrastructure > Nodes**.
- Use a user account with an Appliance Administrator role to log on to the NetBackup Flex Scale infrastructure management UI`https://`*`ManagementServerIPorFQDN`*`:14161` where *ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Monitor > Infrastructure > Nodes**.

2. To include the excluded node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Include node**.

   A notification is displayed on top of the page.

3. To monitor the progress, click **View details**. 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.

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 156. 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.

   If you want to replace a hard disk located in the mid-bay, you must power off the node. To power off the node, navigate to **Infrastructure > Nodes** page of the UI and click **Shutdown node**.

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, the FQDN, or the short host name 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 Management > 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, the FQDN, or the short host name 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**.

     Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

3 Click **Disks**.

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

To identify a disk on the appliance, you can turn on the beacon from the **Infrastructure > Disks** or the **Infrastructure > Hardware > Hard Disk**. The beacon flashes blue for a minute to help you easily identify the disk in your appliance.

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.

Support for mixed disks

NetBackup Flex Scale Gen10 and Gen11 systems support disk replacements using different capacities. This enhancement simplifies field replacements and enables smoother hardware refresh operations.

NetBackup Flex Scale supports mixed disks in the following scenarios:

- Mixed disk capacities are supported within a node and across the cluster.
- Disk replacements can use different firmware versions.
- You can replace a 14-TB disk with a 16-TB or 20-TB disk in the current Flex Scale release.

To maintain cluster consistency, Flex Scale performs logical capacity normalization when disks of different sizes are used.

Disk replacement behavior

- A failed disk can be replaced with a higher-capacity disk using the **Replace disk** operation in the GUI.
- After the replacement completes:
  - The new disk appears in the GUI with its full physical capacity displayed.

- The usable capacity of the disk is automatically aligned with the capacity used by the other disks in the node or cluster.

This behavior ensures consistent storage utilization and prevents imbalance across the system.

- Strict disk model and vendor qualification checks have been relaxed. This change prevents unnecessary blocks during the following operations:
  - Adding a node to a cluster
  - Replacing a failed disk
  - Performing upgrade operations

  As a result, supported disk replacements and expansions can proceed without requiring identical disk models.

Limitations

- Replacing an NVMe or SSD with a higher-capacity disk is not supported.
- This feature applies only to data disks; catalog disks are not supported.
- Replacing a disk with one from a different vendor is not supported.
- Any additional capacity on a higher-capacity replacement disk remains unused and cannot be utilized by the cluster.

Adding an excluded disk to the cluster

A disk is excluded from a cluster if there are frequent transient errors for the disk. A disk gets excluded if the time difference between the last failure and the current failure is less than six hours and this happens at least five times or more. For a disk that is excluded, the disk status is shown as **excluded** on the **Infrastructure > Disks** page of the NetBackup Flex Scale UI.

Note: NetBackup Flex Scale automatically detects and sets the disk status to excluded if there are frequent disk errors. You do not have the option to set the disk status to excluded.

You can perform the following operations for an excluded disk:

- **Include the excluded disk**: You can include the same disk that is set as excluded. If multiple disks are excluded, include the disks one by one.
- **Replace the excluded disk**: To replace an excluded disk, first physically remove the old disk and insert a new one. Perform a replace disk operation on the new disk.

To include an excluded disk:

1. Use any one of the following options to log on 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 on to the NetBackup Flex Scale web UI
     `https://ManagementServerIPorFQDN/webui`
     where*ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Cluster Management > Infrastructure > Disks**.
   - Use a user account with an Appliance Administrator role to log on to the NetBackup Flex Scale infrastructure management UI
     `https://ManagementServerIPorFQDN:14161`
     where*ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Monitor > Infrastructure > Disks**.

2. To include the excluded disk, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Include disk**.

   If multiple disks are excluded, include the disks one by one.

To replace an excluded disk:

1. Physically remove the old disk and insert a new disk.
2. Use any one of the following options to log on 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 on to the NetBackup Flex Scale web UI
     `https://ManagementServerIPorFQDN/webui`
     where*ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Cluster Management > Infrastructure > Disks**
   - Use a user account with an Appliance Administrator role to log on to the NetBackup Flex Scale infrastructure management UI
     `https://ManagementServerIPorFQDN:14161`
     where*ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the NetBackup Flex Scale

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

3 Perform a replace disk operation for the new disk. See “Replacing a disk” on page 80.

Viewing the disk sync status

When data rebalancing is ongoing or a recovery task is in progress, the disk is marked as faulted and the progress for the disk sync operation is shown in the UI instead of the disk status. You can monitor the progress on the **Infrastructure > Disks** page. The node to which the disk belongs is shown unhealthy in the UI.

A disk can be in recovery task during the following operations:

- Replace disk
- Add node
- Replace node
- Reboot node (shutting down and starting the node manually)

The following figure is an example of the disk sync status shown in the UI:

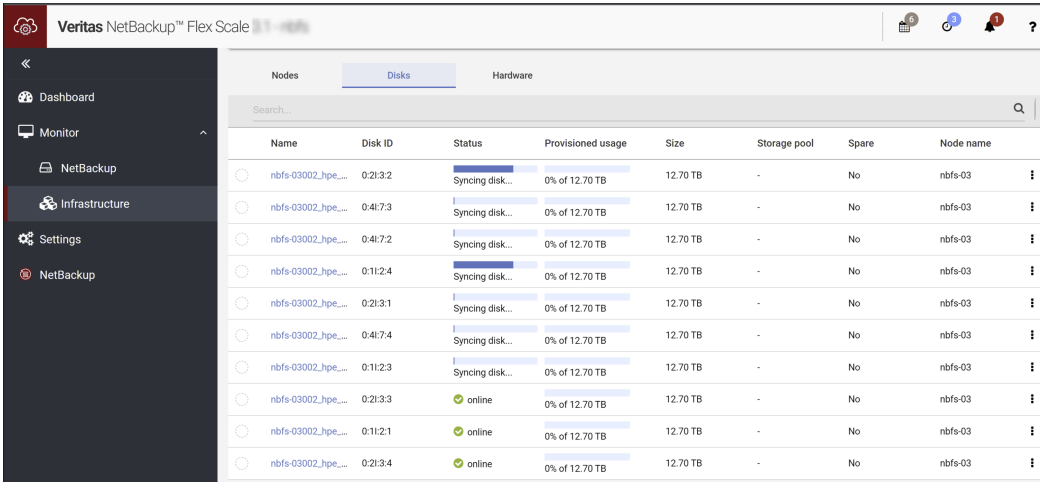

Viewing disk details

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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Cluster Management > 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

You can view the following information:

- Name of the disk.
- Status of the disk:
  - Online: Disk functions normally.
  - Faulted: Disk has hardware issues or is faulty.
  - Excluded: Disk repeatedly goes in a faulted state.
  - Syncing disk: Progress of an ongoing data rebuild operation.
- 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.

- ID of the physical hard disk that corresponds to the logical disk.
- 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Cluster Management > 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.

  Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

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. If the node repeatedly goes in a faulted state, the node status is shown as Excluded. 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.
- Firmware version of the components.
- IP address assigned to the eth1 network interface of the node.
- Appliance model and revision
- 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.

Switching management console to another cluster node

You can switch the management console to the next available cluster node that is online and healthy. The NetBackup Flex Scale UI runs on the node where the management console service runs. When you switch the management console to another node, the UI also fails over to that node. You need to log on to the UI again after the management console fails overs and is brought online on another node. It can take a few minutes for the UI to be up on another node. If you attempt to switch the management console when no other healthy node is available in the cluster, the management console remains online on the current node.

To identify the management console node (node where the management console is running), in the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure**. On the **Infrastructure** page, **Console node** shows the management console node.

To switch the management console:

1. Use any one of the following options to log on 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 on 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, and then in the left pane click **Cluster Management > Infrastructure**.
   - Use a user account with an Appliance Administrator role to log on 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, and then in the left pane click **Monitor > Infrastructure**.
2. Click **Switch console**. When prompted for confirmation, click **OK**.

   A notification is displayed on top of the page.
3. To monitor the progress, click **View details**. 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. After the switch console operation is complete, the management console is brought online on another node. You are required to log on to the UI again.

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 for NetBackup Flex Scale storage and for NetBackup. Thereafter, you can use the NetBackup Flex Scale infrastructure management UI to manage your storage licenses. You must use the NetBackup UI to manage the NetBackup licenses. A valid license is required to maintain a working cluster. Veritas recommends that you add proper licenses before you start protecting production workloads with your appliance.

NetBackup Flex Scale supports the following types of licenses:

- **Trialware**: A built-in evaluation license, which is automatically activated after the initial configuration of the cluster is completed and is valid for 60 days.
- **Subscription**: A time-bound license that must be renewed at the end of the expiration date. Support and maintenance services are included in the time period.
- **Perpetual**: An unlimited validity license with a one-time cost for using the appliance. Support and maintenance services are included for a limited time period.

The following terminology is used to define the storage capacity:

- **Licensed**: Storage capacity for which a license is purchased. You can purchase a subset of the total cluster storage.
- **Used**: Storage capacity that is used for data, excluding the storage required for the catalog file system. Alerts are generated if the used storage is more than the licensed storage.
- **Supported**: Licensed capacity for which you are entitled to Veritas support and maintenance services. The support and maintenance services for the supported storage capacity are available for a limited time period and must be renewed after the expiration.

Based on your storage requirements, you can also license a subset of the total storage capacity of the cluster. With partial licensing, you can purchase a license based on your used capacity rather than the total usable capacity of the cluster. If you need additional storage capacity, you can add additional licenses as per your requirement.

You can add multiple subscription licenses to increase the total cumulative licensed capacity and the subscription period. When you add multiple subscription licenses, the total licensed capacity is the cumulative licensed capacity of all the active subscription licenses. Stacking of licenses is supported only for subscription and perpetual licenses. Stacking multiple trialware licenses is not supported.

The following table shows how the licenses can be stacked:

Existing license	Add additional license
Trialware	Perpetual or Subscription
Perpetual	Perpetual
Subscription	Subscription

Monitoring license compliance

Alerts are generated if you no longer meet the compliance standards for licensing. If AutoSupport is configured for the cluster, alerts are also emailed to the administrators. Alerts are generated in the following situations:

- If used capacity is more than the licensed or supported capacity for a period of six hours an alert is generated and critical operations such as add node and upgrade are blocked.
- If used capacity is more than 90% of the licensed or supported capacity for a period of six hours.
- If one or more licenses expire, which results in not meeting the compliance standards.

See "Adding and removing storage licenses" on page 91.

Adding and removing storage licenses

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

To manage storage licenses:

1. Sign in to the NetBackup Flex Scale infrastructure management console UI.

   See "About the NetBackup Flex Scale infrastructure management UI" on page 17.

2. Click **Settings > Licensing management**.
3. To add a storage license, click **Add license**, on the Add storage licenses page click **Choose file**, and browse to the location of the `.slf` license file and select it. Alternatively, you can also drag and drop the license file directly on the page.

   The page shows the licensed capacity that will be available after the license is added and also indicates if you comply with the licensing model. You can select multiple license files; the licensed capacity reflects the cumulative licensed capacity that will be available.
4. Click **Add** to configure the license into the cluster.

   A notification is displayed on the top of the page. To monitor the progress and view the task details click **View Details**. The ongoing and completed tasks for the add license operation are also displayed in **Recent activity**.

   The **Storage Licenses** tab shows the used, licensed, and supported storage.

5 To remove a particular storage license, in the license row, from the Actions menu (vertical ellipsis) on the right side of the UI and then select **Delete**.

6 When prompted for confirmation, click **Delete**. To monitor the progress of the delete operation, click **View details**.

Once the operation completes, verify that the license you removed does not appear in the list of licenses that are 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, the FQDN, or the short host name that you specified for the management server during the cluster configuration. For more details about the APIs, see the NetBackup Flex Scale APIs on SORT.

Stopping NetBackup service containers

You can start and stop NetBackup services from the UI. You can perform the following operations from the UI:

- Stop each NetBackup service individually.
- Stop all NetBackup services for a specific node or the cluster without taking the file systems offline and unmounting the file systems.
- Stop all NetBackup services for a specific node or the cluster and take the file systems offline and unmount the file systems that are used by the services.
- Start the stopped services.

If you stop the catalog engine (MSDP engine service that maintains the catalog data), all the other MSDP engines become unhealthy and cause the NetBackup jobs to fail.

Note: When you stop a NetBackup service from the UI, its state is persistent across operations such as node reboot or cluster shut down. The service remains stopped until you start the service from the UI.

To stop containers:

1 Use any one of the following options to log on 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 on 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, and then in the left pane click **Cluster Management > Services**.

- Use a user account with an Appliance Administrator role to log on to the NetBackup Flex Scale infrastructure management UI `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, and then in the left pane click **Monitor > NetBackup > Services**.

2. To stop a specific service, select the service container and click **Stop**. When prompted for confirmation, click **Stop container(s)**.

   A notification about the operation is displayed on the top of the page. To monitor the progress, click **View details**.

   Note: On the NetBackup Flex Scale web UI if you choose to stop the NetBackup primary server service, you are redirected to the NetBackup Flex Scale infrastructure management UI.

3. To stop all services for a specific node or to stop all services on the cluster, select the services and click **Stop**. To unmount the file systems that the services use, when prompted, click **Stop and unmount**; else click **Stop**.

   A notification about the operation is displayed on the top of the page. To monitor the progress, click **View details**.

Starting NetBackup service containers

You can start and stop NetBackup services from the UI. You can perform the following operations from the UI:

- Stop each NetBackup service individually.
- Stop all NetBackup services for a specific node or the cluster without taking the file systems offline and unmounting the file systems.
- Stop all NetBackup services for a specific node or the cluster and take the file systems offline and unmount the file systems that are used by the services.
- Start the stopped services.

To start containers:

1. Use any one of the following options to log on 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 on 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, and then in the left pane click **Cluster Management > Services**.
   - Use a user account with an Appliance Administrator role to log on to the NetBackup Flex Scale infrastructure management UI
     `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, and then in the left pane click **Monitor > NetBackup > Services**.

2. To start services, select the services and click **Start**.

   The file systems are mounted automatically when you start the services if you had opted to unmount the file systems when you stopped the services.

Managing the Fibre Channel ports

If the Fibre Channel card is installed on the cluster nodes, you can configure the Fibre Channel ports for a specific workload. If you install the Fibre Channel card, ensure that it is installed on all the nodes. Either all the cluster nodes must have the Fibre Channel card or it must not be installed on any of the cluster nodes.

The Fibre Channel card must be installed in the following slots:

- For HPE Gen11 (Model number: 5561): Slot-5
- For HPE Gen10 (Model number: 5551): Slot-4

Supported Fibre Channel card configuration:

- Brand: QLogic
- Configuration: QLE2772, 32G(auto), 2 ports, hpe pn=SN1610Q Dual-Port FC card

The following types of backups over Fibre Channel are supported:

- VMware SAN transport (initiator)

- Tape out (initiator)

See “Requirements” on page 95.

See “Assigning Fibre Channel ports” on page 99.

See “Unassigning Fibre Channel ports” on page 101.

See “Rescanning Fibre Channel cards” on page 100.

See “Viewing details about the Fibre Channel ports” on page 101.

See “Discovering attached devices” on page 100.

See “Cleaning Fibre Channel ports” on page 101.

Requirements

This section describes the requirements for using Fibre Channel for backup jobs. Review the following requirements if you plan to perform backups over Fibre Channel after configuring the cluster.

Table 3-3

Requirement	Details
Install the Fibre Channel card.	Ensure that the Fibre Channel card is installed on all the nodes of the cluster in the supported slot. ■ For HPE Gen11 (Model number: 5561) install the card in Slot 5. ■ For HPE Gen10 (Model number: 5551) install the card in Slot 4.
Enable BOM (Bill of Materials) configuration for Fibre Channel.	Enable the BOM so the cluster can detect the installed Fibre Channel cards and display the Fibre Channel details. You must enable the BOM before you can assign Fibre Channel ports to workloads. See “Enabling BOM (Bill of Materials) configuration for Fibre Channel” on page 97.
Review the connectivity requirements.	See “Connectivity requirements for Fibre Channel” on page 96.
Review storage unit requirements	See “Storage unit (STU) requirements” on page 96.

Note the following considerations while adding a node to the cluster:

To use Fibre Channel feature after adding a node, add the Fibre Channel on the new node in the supported slot. Ensure that all the nodes in the cluster have the card. Set up zoning for the Fibre Channel ports. Enable BOM configuration for Fibre

Channel before starting the add node operation. Recreate Tape STU from the NetBackup JAVA UI and add the newly added node's media server in the tape STU so the new media server will be available for backup and restore operations.

Note the following considerations while replacing a node in the cluster:

To use Fibre Channel feature after replacing a node, add the Fibre Channel on the new node in the supported slot. Ensure that all the nodes in the cluster have the card. Set up zoning for the Fibre Channel ports. Enable BOM configuration for Fibre Channel before starting the replace node operation. Recreate Tape STU from the NetBackup JAVA UI and add the newly added node's media server in the tape STU.

Note the following considerations for performing factory reset:

If you factory reset all the nodes, add the Fibre Channel card on all the nodes of the cluster in the supported slot. Set up zoning for the Fibre Channel ports. Enable BOM configuration for Fibre Channel before starting the cluster configuration. Before enabling the BOM make sure all the nodes are connected with FC cards. Create Tape STU from the NetBackup JAVA UI.

If you factory reset a specific node, add the Fibre Channel on the node in the supported slot. Ensure that all the nodes in the cluster have the card. Set up zoning for the Fibre Channel ports. Enable BOM configuration for Fibre Channel before starting the factory reset operation. Recreate Tape STU from the NetBackup JAVA UI and add the node's media server in the tape STU.

Note the following considerations when disaster recovery is configured:

If you plan to use Fibre Channel for backups on the primary/secondary cluster, configure disaster recovery first, and then configure Fibre Channel on primary and secondary clusters as per requirement.

Connectivity requirements for Fibre Channel

Note the following connectivity requirements:

- Veritas recommends that every media server, and hence every port which is assigned for tape out option must be connected to all the tape drives. Configuration where all the media servers are not connected to all the tape drives is supported but not recommended.
- Fibre Channel switches are required and zoning must be set up to access the tape drives.

Storage unit (STU) requirements

Note the following storage unit (STU) requirements for the media servers:

- Tape drives are mapped to media servers and these media servers are mapped to STU.
- Veritas recommends that you create a single STU, which will include all the media servers and hence all the tape drives. This single STU for all tape outs can be used to create policies. Use the NetBackup JAVA UI to create the STU.
- If a robot is used move tape cartridges into and out of tape drives, one of the media servers is assigned a robot controller role. This role is present on only one media server, and doesn't fail over if that media server is down. In such situation, even if the other nodes have connectivity to the tape drives, they can't be used. The media server with the robot control role must be brought online to proceed further.

Enabling BOM (Bill of Materials) configuration for Fibre Channel

Before you can configure Fibre Channel ports for performing backups, you must enable Fibre Channel BOM configuration from the Appliance Node-level CLI.

To enable Fibre Channel BOM:

1. Use SSH Login to Node level CLI using the eth1 IP address of the node
2. Run the following command:

   "support bom-conf get

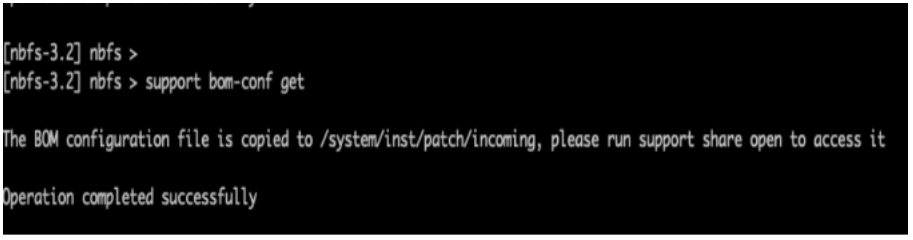

3 To update this `bom-config.json` file open an NFS share by running the following command:

`system software share open`

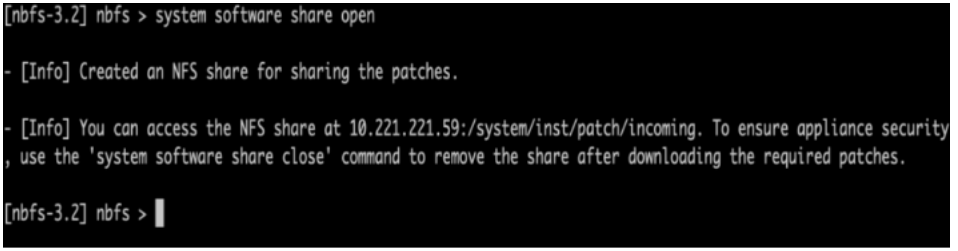

4 Mount the above open NFS share on any available Linux Client by using the Linux command:

`mount -t nfs` *`ipaddressorfqdn`*`:/system/inst/patch/incoming /mnt`

where *ipaddressorfqdn* is the IP address or the FQDN that was displayed when you ran the `system software share open` command.

The `/mnt` location can be changed to the required location for mounting the NFS share.

The `bom-config.json` file is copied to the `/mnt` location.

5 Set the **FC-ENABLE** key value to **1**. Open the `bom-config.json` using the `vim` command:

`vim bom-config.json`

Change the **FC-ENABLE** key to 1 and save the file. the updated file should look similar to shown below:

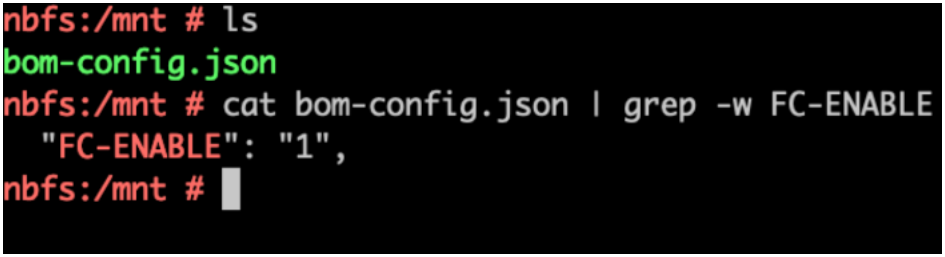

6 Log in to the Appliance Node-level CLI again using the eth1 IP address of the node.

7 Run the following command to update configuration:

```
support bom-conf update
```

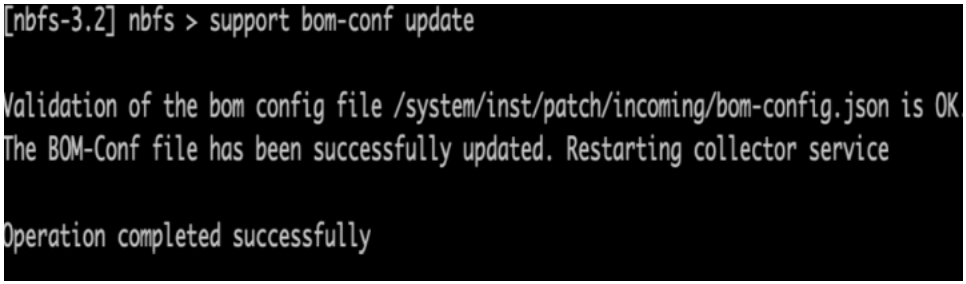

8 After the file is updated, close the open NFS share:

```
system software share close
```

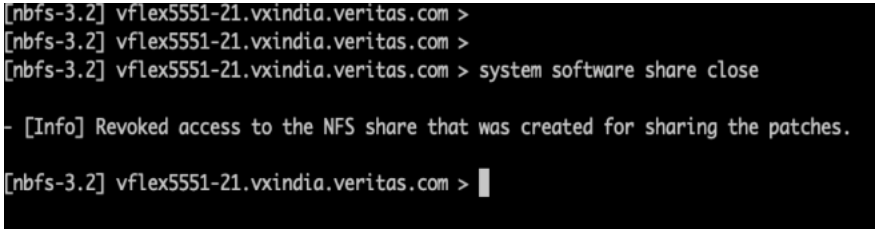

9 Perform the same procedure on all the other nodes.

After you enable the Fibre Channel BOM on all the nodes, navigate to **Monitor > Infrastructure**. You can now view the **Fibre channel** tab.

Assigning Fibre Channel ports

If you want to perform backups over Fibre Channel, you must assign ports for the workloads.

You can configure Fibre Channel ports for a specific workload. Before you begin, ensure that all the media servers are healthy.

To assign Fibre Channel ports to a workload:

1 In the NetBackup Flex Scale management infrastructure UI, navigate to **Monitor > Infrastructure > Fibre channel** or from the NetBackup Flex Scale webui click Cluster Management > Infrastructure > Fibre channel.

2 Select the ports that you want to assign and click **Clean**.

3. Select the ports that you want to assign and click **Discover devices**.
4. Select the ports that you want to assign and click **Assign port**.
5. In the **Selected ports** list, select the ports.
6. In the **Workload** list, select the workload that you want to configure for the port.
7. Click **Assign ports**.

   The port and the workload configured for the port is displayed under **Assigned ports**.
8. Click **Confirm**.

   On the **Fibre Channel** tab, the ports are shown as assigned and the configured workload is displayed.

   After you assign the ports, ensure that the storage device configuration is done from the NetBackup JAVA UI. Create a single Storage Unit (STU) which includes all the media servers and then configure backup policies which include this STU.

Discovering attached devices

The appliance nodes scan for devices when they start. If you connect devices while the appliance nodes are running, use the following procedure to detect the newly connected devices.

To scan for new devices:

1. In the NetBackup Flex Scale management infrastructure UI, navigate to **Monitor > NetBackup > Services** and select the media containers that you want to stop. Click **Stop**.
2. Navigate to **Monitor > Infrastructure > Fibre channel** or from the NetBackup Flex Scale webui click **Cluster Management > Infrastructure > Fibre channel**.
3. Select the ports for which you want to scan for devices and click **Discover devices**.
4. Navigate to **Monitor > NetBackup > Services**. Start the media containers.

Rescanning Fibre Channel cards

Use **Rescan** to detect the connected Fibre Channel cards and the existing or the new Fibre Channel configuration.

To rescan Fibre Channel ports:

1. In the NetBackup Flex Scale management infrastructure UI, navigate to **Monitor > Infrastructure > Fibre channel** or from the NetBackup Flex Scale webui click **Cluster Management > Infrastructure > Fibre channel**.
2. Click **Rescan**.

   Details about the Fibre Channel card are displayed.

Cleaning Fibre Channel ports

The appliance nodes scan for devices when they start. If you remove devices while the appliance nodes are running, use the following procedure to clean stale device information from the system.

To clean device information:

1. In the NetBackup Flex Scale management infrastructure UI, navigate to **Monitor > Infrastructure > Fibre channel** or from the NetBackup Flex Scale webui click **Cluster Management > Infrastructure > Fibre channel**.
2. Select the ports for which you want to rescan and click **Clean**.

Unassigning Fibre Channel ports

To unassign Fibre Channel ports:

1. In the NetBackup Flex Scale management infrastructure UI, navigate to **Monitor > Infrastructure > Fibre channel** or from the NetBackup Flex Scale webui click **Cluster Management > Infrastructure > Fibre channel**.
2. Select the ports that you want to unassign and click **Unassign ports**.

Viewing details about the Fibre Channel ports

You can click **Rescan** to scan the Fibre channel card for displaying the port details. On the Fibre channel tab, you can view details such as the port status, Fibre Channel card mode, and whether the port is assigned to a workload. You can also view additional details on the **Hardware** tab.

To view the details:

1. In the NetBackup Flex Scale management infrastructure UI, navigate to **Monitor > Infrastructure > Hardware** or from the NetBackup Flex Scale webui click **Cluster Management > Infrastructure > Hardware**.
2. Click **Fibre channel HBA**.

   Details such as physical port to logical port mapping, port status, port speed, and vendor details are displayed.

Disabling BOM (Bill of Materials) configuration for Fibre Channel

If the Fibre Channel BOM is enabled on a node, but you don't want the Fibre channel functionality and don't intend to configure Fibre Channel ports for workloads, ensure that you disable the BOM. If the BOM is not disabled, the nodes become unhealthy and you start receiving AutoSupport alerts.

To disable Fibre Channel BOM:

1. Use SSH Login to Node level CLI using the eth1 IP address of the node
2. Run the following command:

   `"support bom-conf get`

```
[nbfs-3.2] nbfs >
[nbfs-3.2] nbfs > support bom-conf get

The BOM configuration file is copied to /system/inst/patch/incoming, please run support share open to access it

Operation completed successfully
```

3. To update this `bom-config.json` file open an NFS share by running the following command:

   `system software share open`

```
[nbfs-3.2] nbfs > system software share open

- [Info] Created an NFS share for sharing the patches.

- [Info] You can access the NFS share at 10.221.221.59:/system/inst/patch/incoming. To ensure appliance security
, use the 'system software share close' command to remove the share after downloading the required patches.

[nbfs-3.2] nbfs >
```

4. Mount the above open NFS share on any available Linux Client by using the Linux command:

   `mount -t nfs` *`ipaddressorfqdn`*`:/system/inst/patch/incoming /mnt`

   where *ipaddressorfqdn* is the IP address or the FQDN that was displayed when you ran the `system software share open` command.

   The `/mnt` location can be changed to the required location for mounting the NFS share.

   The `bom-config.json` file is copied to the `/mnt` location.

5. Set the **FC-ENABLE** key value to **0**. Open the `bom-config.json` using the `vim` command:

   `vim bom-config.json`

   Change the **FC-ENABLE** key to **0** and save the file. The updated file should look similar to shown below:

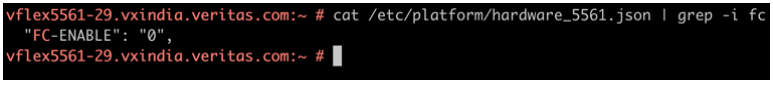

6. Log in to the Appliance Node-level CLI again using the eth1 IP address of the node.

7. Run the following command to update configuration:

   `support bom-conf update`

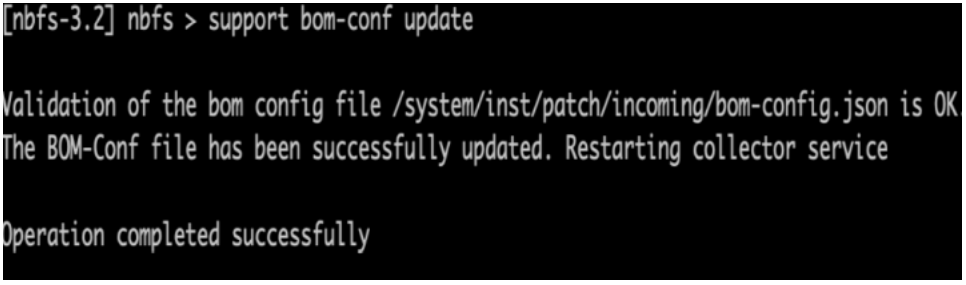

8 After the file is updated, close the open NFS share:

```
system software share close
```

```
[nbfs-3.2] vflex5551-21.vxindia.veritas.com >
[nbfs-3.2] vflex5551-21.vxindia.veritas.com >
[nbfs-3.2] vflex5551-21.vxindia.veritas.com > system software share close

- [Info] Revoked access to the NFS share that was created for sharing the patches.

[nbfs-3.2] vflex5551-21.vxindia.veritas.com >
```

9 Perform the same procedure on all the other nodes.

After you disable the Fibre Channel BOM, the nodes will turn healthy and you will no longer receive alerts.

Managing hardware vendor packages

The vendor packages are preinstalled on out of the box appliance nodes. You do not have to install anything on the raw nodes after you take them out of the box. Use the following procedure if you want to install third-party vendor packages. If you want to install third-party vendor packages, download the required vendor packages from the hardware vendor site. You can uninstall or upgrade the vendor packages after they are installed.

To install vendor packages after the cluster is configured:

1 Use any one of the following options to log on 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 on 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, and then in the left pane click **Cluster Management > Cluster settings > Software management**.
- Use a user account with an Appliance Administrator role to log on to the NetBackup Flex Scale infrastructure management UI `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, and then in the left pane click **Settings > Software management**.

Note: If you access the NetBackup Flex Scale infrastructure management UI by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

---

2. On the Software management page, click **Vendor packages**.
3. On the **Vendor packages** tab click **Verify**.
4. Do one of the following:
   - If there is a change in required vendor packages, download the existing `vendor_utilities.json` file, edit the file to update the vendor package list, and then upload the updated JSON file. The uploaded JSON file replaces the existing file and is now used a reference for the list of vendor packages that must be present on the node. To download the JSON file, on the **Vendor packages** tab click **Download JSON file**. Edit the downloaded file. To upload the updated file, click **Next > Upload JSON file**. Click **Verify**.
   - If there is no change in the required vendor packages, click **Next > Verify**
5. If the required vendor packages are missing or the installed package version is earlier than what is listed in the JSON file, the node is marked as unhealthy and the details about the missing packages are displayed for each node. Click **Add**.
6. Download the required vendor packages from the vendor site.
7. Upload the vendor packages to the driver node. On the **Vendor packages** tab, click **Add** and select the vendor packages that you downloaded. The selected vendor packages are uploaded to the `/system/inst/patch/incoming` location of the driver node and their status in the UI is shown as **Available**.
8. Optionally, to remove the uploaded vendor packages, select the vendor packages and click **Remove**.

9 If the vendor packages are required to be installed in a specific sequence, click the arrows in the **Reorder** column to change the installation sequence.

10 Install the uploaded vendor packages. Select all the vendor packages and click **Install**. Each package is installed on all the nodes. If the packages are installed successfully on all the nodes, a notification is displayed on the top of the page and the package status changes to **Installed**. If the installation fails on one of the nodes, the vendor package is rolled back on all the nodes and the package status is shown as **Available**.

See "Upgrading vendor packages" on page 106.

See "Uninstalling vendor packages" on page 107.

Upgrading vendor packages

Use the following procedure to upgrade the vendor packages to a later version. If you try to upgrade a vendor package to an earlier version or to its current version that is same or earlier an error is displayed.. You can upgrade a vendor package to a later version than the current version.

To upgrade vendor packages:

1 Use any one of the following options to log on 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 on 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, and then in the left pane click **Cluster Management > Cluster settings > Software management**.
- Use a user account with an Appliance Administrator role to log on to the NetBackup Flex Scale infrastructure management UI `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, and then in the left pane click **Settings > Software management**.

Note: If you access the NetBackup Flex Scale infrastructure management UI by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

2 On the Software management page, click **Vendor packages**.

The installed vendor packages are displayed.

3 Download the vendor packages from the vendor site. Ensure that the version of the downloaded packages is not the same or earlier than the current version of the package.

4 Select the installed vendor package that needs to be upgraded, click **Upgrade** and upload the latest package to start the upgrade.

The selected package is upgraded on all the nodes and is displayed as **Installed**, whereas the old package is shown **Available**.

Uninstalling vendor packages

Use the following procedure to uninstall vendor packages from the cluster nodes.

To uninstall vendor packages:

1 Use any one of the following options to log on 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 on 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, and then in the left pane click **Cluster Management > Cluster settings > Software management**.
- Use a user account with an Appliance Administrator role to log on to the NetBackup Flex Scale infrastructure management UI `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, and then in the left pane click **Settings > Software management**.

Note: If you access the NetBackup Flex Scale infrastructure management UI by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

2 On the Software management page, click **Vendor packages**.

The installed vendor packages are displayed.

3 Select the vendor packages that you want to delete and from the Actions menu (vertical ellipsis) click **Rollback**.

The selected vendor packages are uninstalled from all the nodes. To monitor the progress, click **View details** on the Software management page.

Updating credentials for HPE iLO administrator users

The default password of the iLO Administrator user is printed on the appliance pull tab. If you change this default password, ensure that you run the following command to save the updated password:

```
system store-updated-admin-BMC-credentials
```

This command validates the credentials and prompts you to try again if the authentication fails.

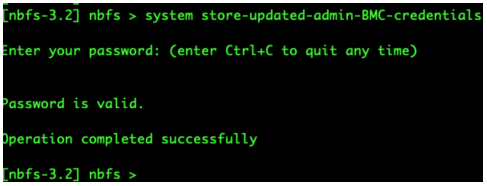

The following figure shows the output that is displayed when the specified password is incorrect:

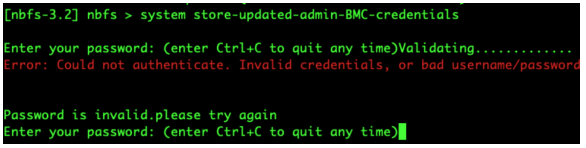

Chapter 4

NetBackup Flex Scale network management

This chapter includes the following topics:

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 134.

See “Bonding operations” on page 115.

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 modify the DNS details for both data and management network.

See “Modifying DNS settings” on page 111.

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 130.

Network > **Management Network** displays the details 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 > Management 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
- Console FQDN, if it is configured on the cluster.

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.

Network > **Static routes** lets you configure static routes on a cluster to communicate with remote clients who are on a different subnet, if the default route is not suitable.

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** and **Data network**.

- Go to **Network** > **DNS** > **Management network** or **Data 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.

Configuring MTU on public interfaces

The MTU property controls the maximum transmission unit size for an Ethernet frame. The standard maximum transmission unit size for Ethernet is 1500 bytes (without headers). MTU defines the largest data packet that a network-connected device accepts. The network routers check the size of each IP packet that they receive against the MTU of the next router that will receive the packet. If the packet exceeds the MTU of the next router, the first router breaks the payload into two or more packets, each with its own headers.

You can manage the MTU values using both the GUI and RESTful APIs. In a NetBackup Flex Scale appliance, the range of MTU is between 1500 to 9000.The MTU value changes the payload size of the packets which modifies the data transfer rate. If the MTU value is edited, the changes are reflected across the cluster.

To edit the value that has been set for MTU using the NetBackup Flex Scale GUI:

- Navigate to **Settings > Network > Data Network** or **Settings > Network > Management Network** .
- Go to the network that you want to modify. Either click on the kebab menu and choose **Edit MTU** or click **Edit** button from the subsection of the chosen interface.
- In the **Edit MTU size** pop-up, enter the new value and click **Modify**.

Notification events are raised when the MTU values are modified.

Note: When you use the GUI to set the MTU value on management NICs, capacity validation of the NIC is not performed because there is no guarantee that all the nodes will have the same direct public switch connections to their management interfaces on a public network. Hence, ensure the MTU value you want to set is actually supported at NIC and switch level on all the nodes. Else, SSH commands between the nodes may start failing intermittently or operations such as add node may hang.

The following RESTful APIs are available to configure and manage MTU values:

- `GET /api/appliance/v1.0/network/mtu`: Fetches the MTU values for all interfaces.
- `GET /api/appliance/v1.0/network/mtu/{interfaceName}`: Fetches the MTU value for the specified interface.
- `PATCH /api/appliance/v1.0/network/mtu/{interfaceName}`: Modifies the MTU value for the specified interface and sets it to the value mentioned in the payload of the API.

Configuring the console FQDN

You can now log in to the NetBackup Flex Scale UI using the console FQDN.

To configure the console FQDN

1. Go to **Settings > Network > Management Network**.
2. Click **Add console FQDN**.
3. For a cluster which is configured with DNS settings, enter the FQDN or short name which resolves to the console IP.

   For a non-DNS cluster, you can add any FQDN which is free or an IP that is not mapped to any FQDN.

   Note: You cannot add FQDN short name.

   Click **Add**.
4. The **Adding Console FQDN** task is initiated. You can click **View Details** to see the progress of the task. 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.
5. Once the configuration is complete, you can access the NetBackup Flex Scale UI using the console FQDN.

- If ECA is configured, then you must reconfigure the external certificate, so that the console FQDN is added in SAN entries.
- If ECA is not configured, then the internal certificate is regenerated, so that the console FQDN is added in the SAN entries,. The GUI server is automatically restarted aftert he internal certificate is re generated.
- On a NetBackup Flex Scale cluster on which both primary and media servers are deployed, the NetBackup SAN and CORS entries are updated with the console FQDN.

Limitations

This feature has the following limitations:

- You cannot edit or remove console FQDN once added.
- ECA must be reconfigured after adding the console FQDN.

Log locations

You can find the logs for troubleshooting at:

- /log/VRTSnas/ console_fqdn_oper.log
- /log/VRTSnas/ isagui_webserver.log
- /log/VRTSnas/ isagui_cert.log

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	balance-rr	Yes	Yes	Yes	No
2	active-backup	Yes	No	No	No
3	balance-xor	Yes	Yes	Yes	No
4	broadcast	Yes	No	Yes	No
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.

Note: When you create or remove a bond, the NetBackup services go offline as the device name for all the IPs of the data network change . You have to bring the IPs online on the new device name (either bond or base device). The NetBackup services come up after the `bond create` or `bond remove` operation is completed.

Bonding operations

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

You can perform bonding operations on both the data network and management network.

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

You can find the RESTful APIs at `https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0/` where *ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name 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 *Veritas NetBackup FlexScale APIs* on SORT.

Bonding operations on data network

You can perform the following operations on the data network:

Create a bond: See “Creating a bond” on page 116.

Modify a bond: See “Modifying a bond” on page 120.

Remove a bond: See “Removing a bond” on page 121.

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**.

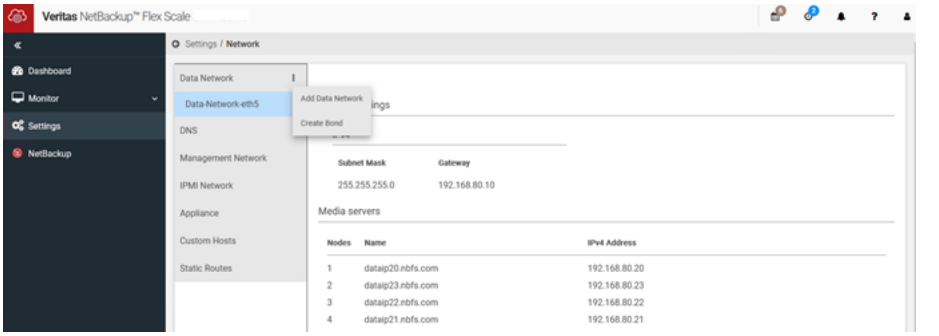

3 The **Create Bond** form appears.

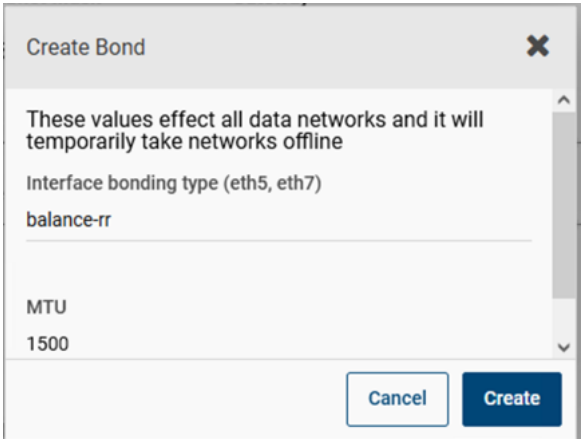

By default, the bond mode here is `balance-alb` and MTU is 1500. You can use the drop-down box to change the bond mode.

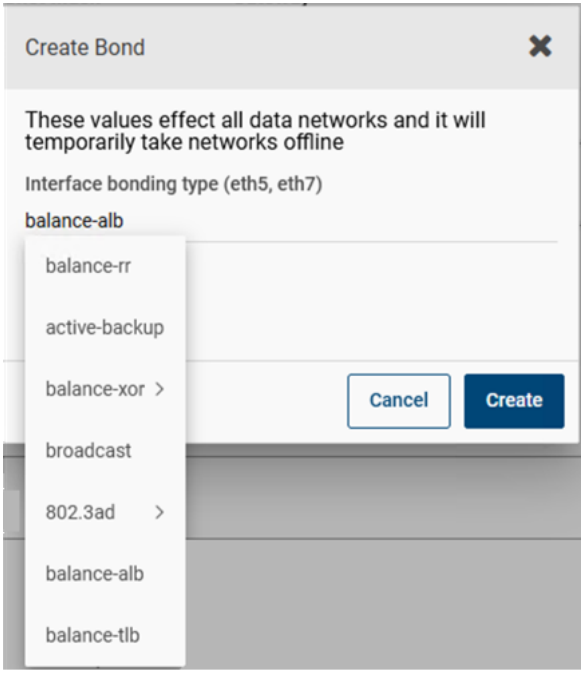

See “About bonding Ethernet interfaces” on page 114.

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

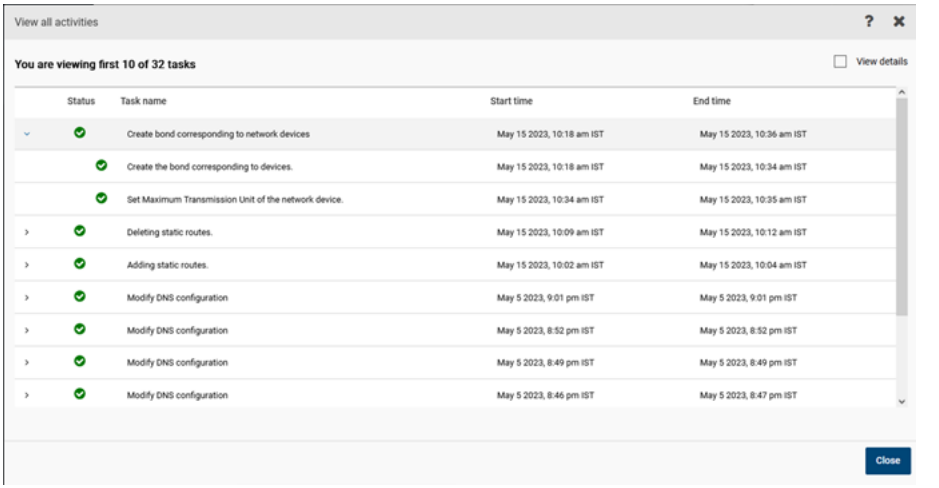

The display name of the data network changes to **Data-Network-bond0**.

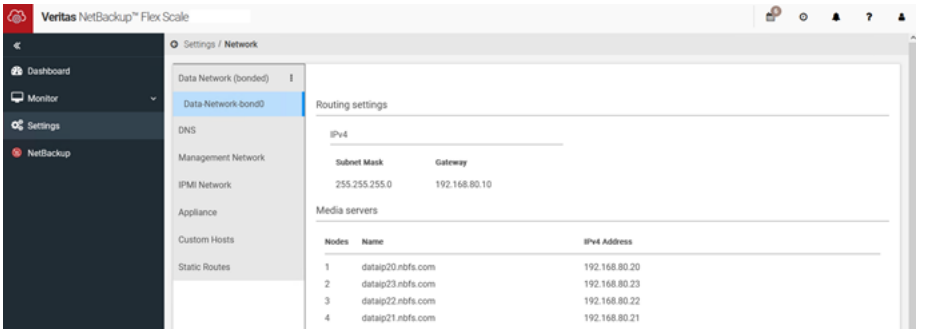

Modifying a bond

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**.

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.

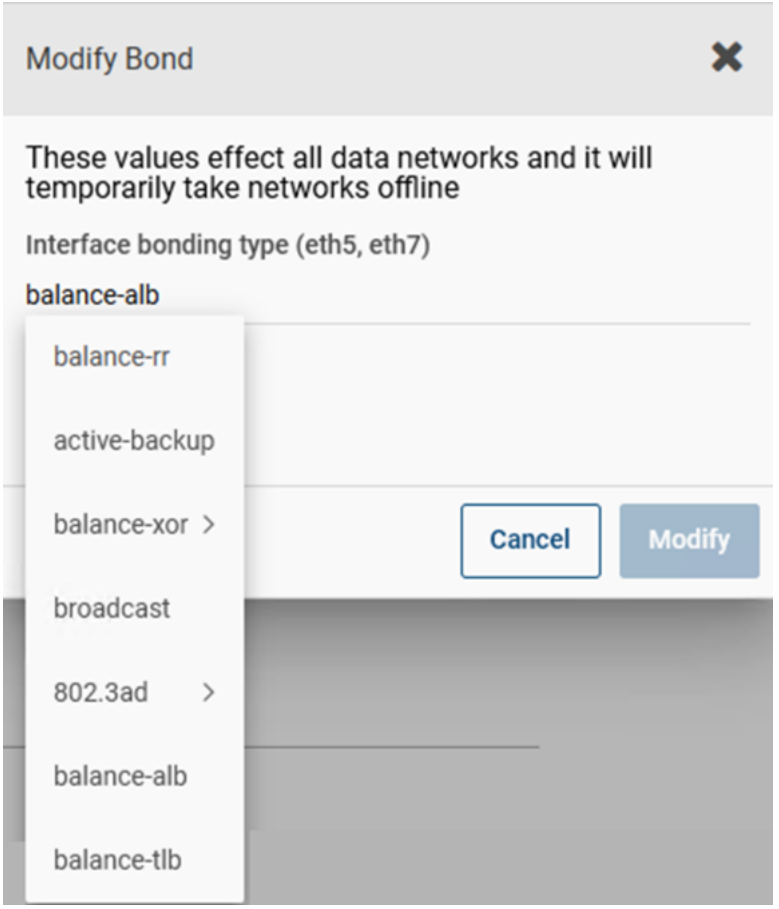

See “About bonding Ethernet interfaces” on page 114.

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.

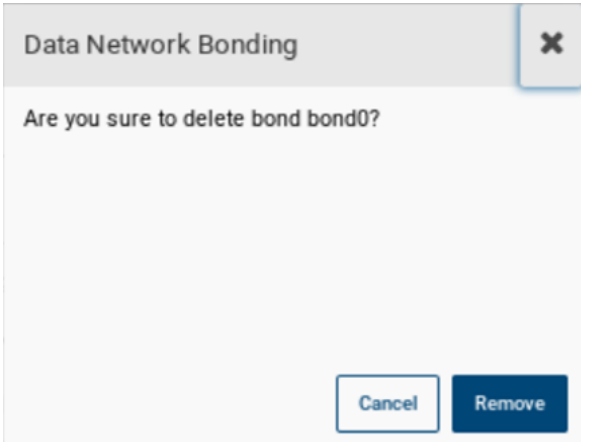

After you confirm, the bond removal is initiated. View the **Recent Activity** panel in the top navigation bar for the status of the task. On successful removal of the bond, the data network is moved to eth5. The data network name is changed to **Data-Network-eth5**.

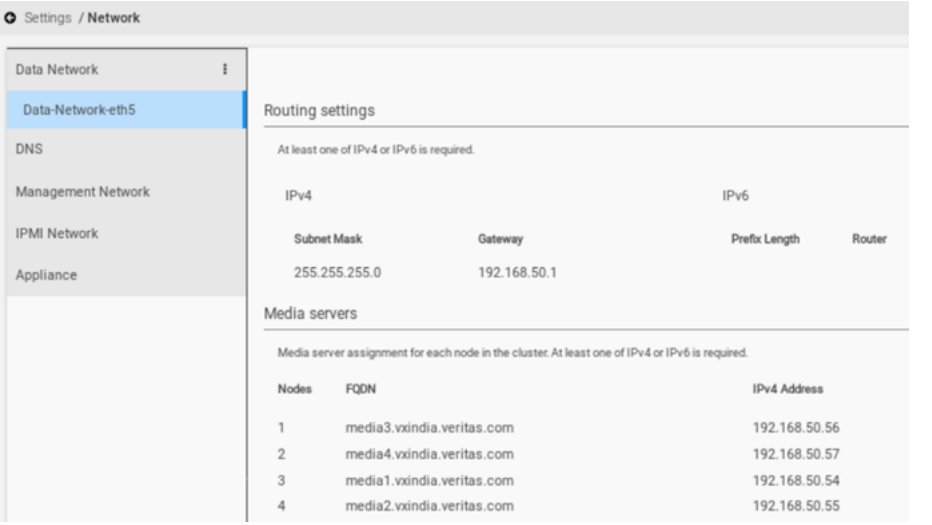

During bond removal, if a secondary network exists, it is moved to eth5. Even if the secondary network was present on eth7 before bond creation, it is moved to eth5 after bond removal.

In the following example, **Data-Network-bond0-874** is moved to eth5 after bond removal.

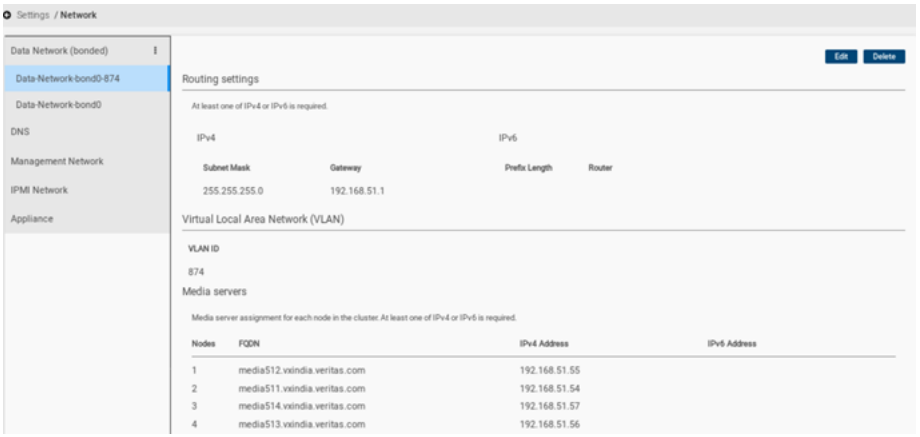

It now appears as **Data-Network-eth5-874**

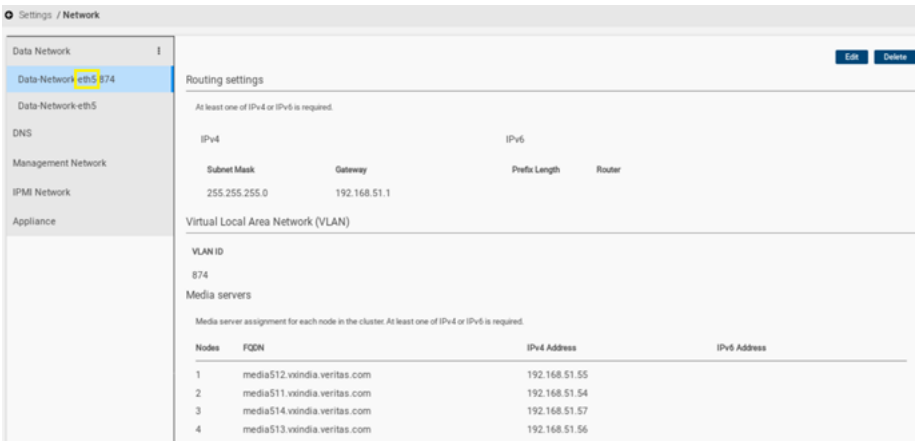

Bonding operations on management network

You can perform the following operations on the management network:

Create a bond: See “Creating a bond” on page 124.

Modify a bond: See “Modifying a bond” on page 127.

Remove a bond: See “Removing a bond” on page 128.

Creating a bond

The interfaces eth1 and eth2 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 **Management Network**. Select **Create Bond**.

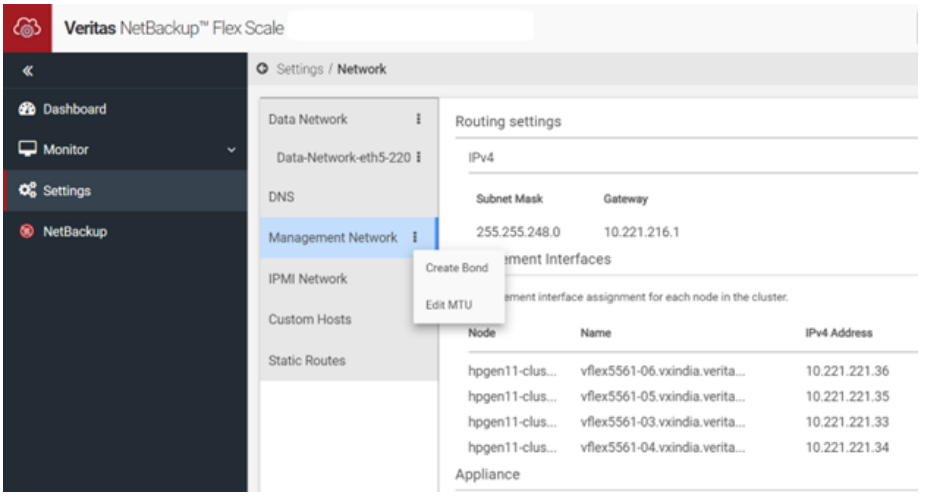

3 The **Create Bond** form appears.

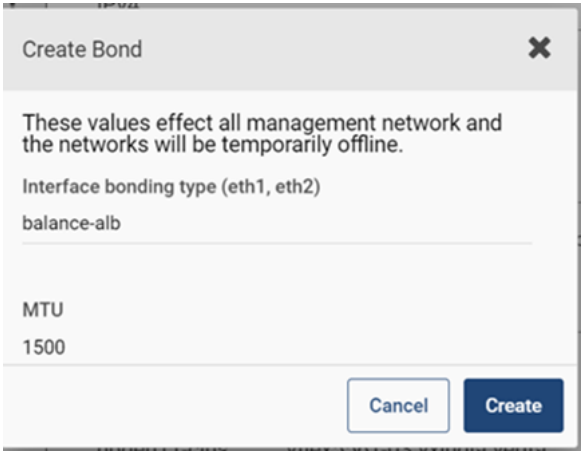

By default, the bond mode here is `balance-alb` and MTU is 1500. You can use the drop-down box to change the bond mode.

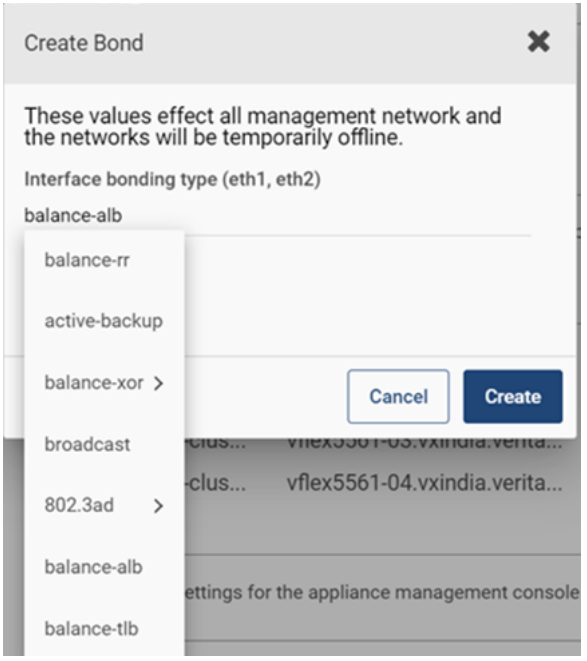

See “About bonding Ethernet interfaces” on page 114.

Click **Create**.

4 A popup appears to inform that the bonding operation brings down the eth1 device and GUI will not be available for some time. Click **OK**.

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

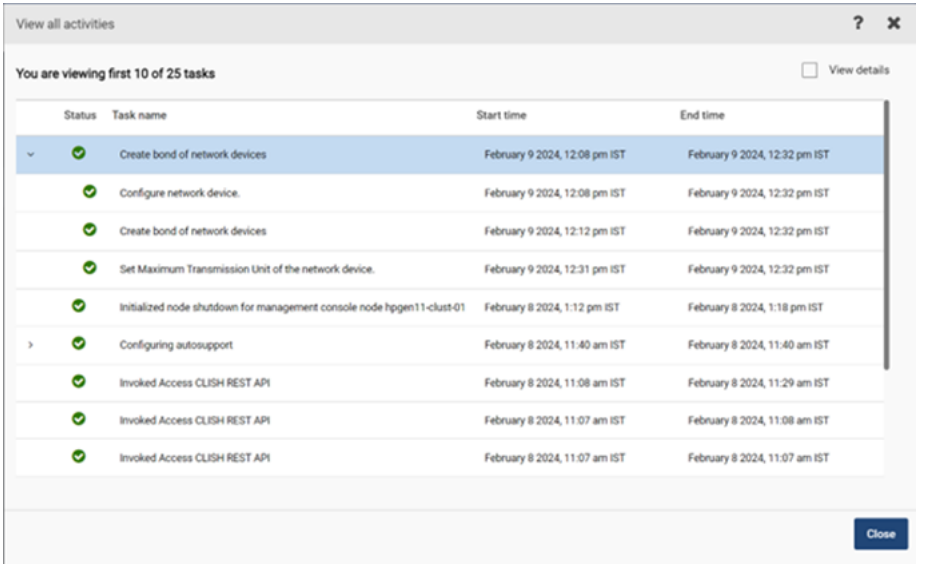

After the successful completion of the task, the display name of the management network changes to **Management Network (bonded)**.

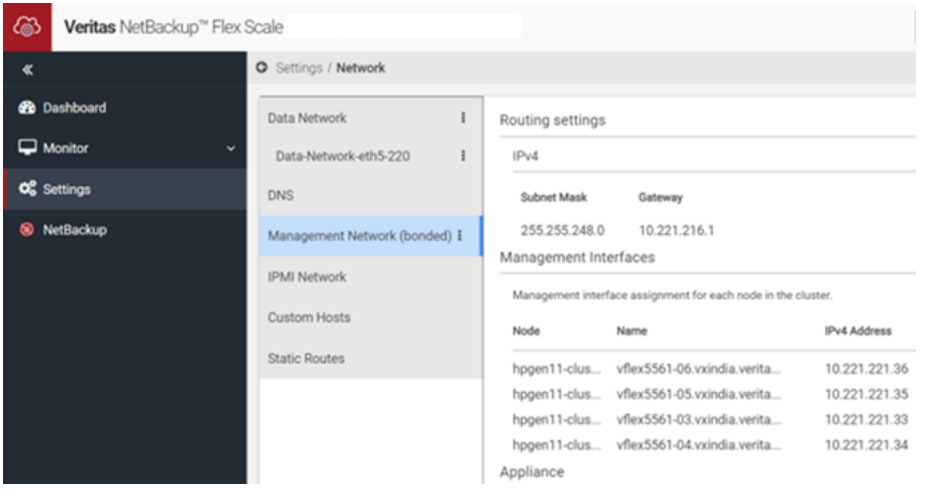

Modifying a bond

To modify a bond

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

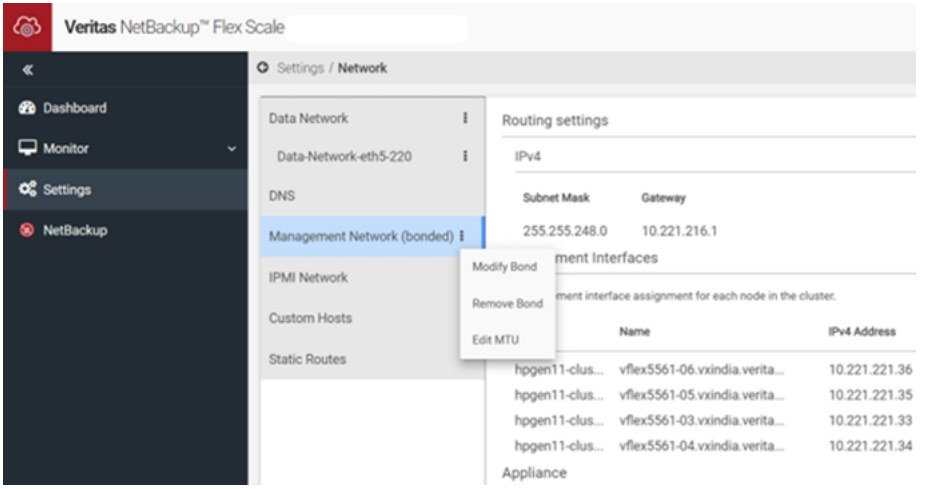

3 The **Modify Bond** form appears.

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

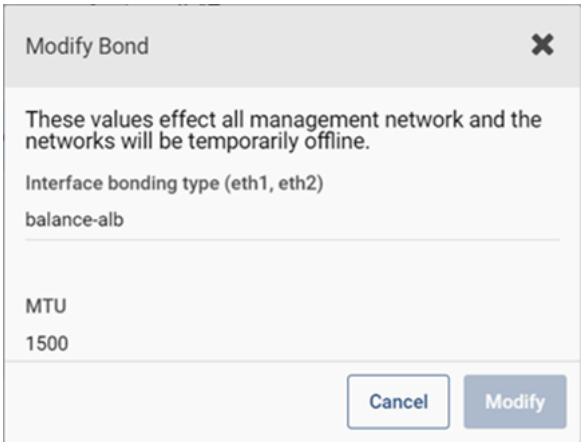

See “About bonding Ethernet interfaces” on page 114.

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 **Management Network (bonded)**. Select **Remove Bond**.

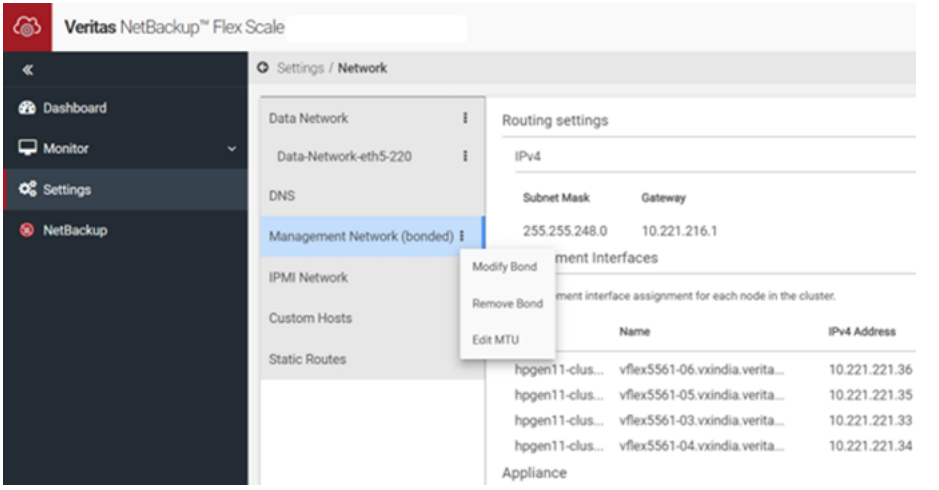

3. A confirmation screen appears. Click **Remove** to delete the bond.

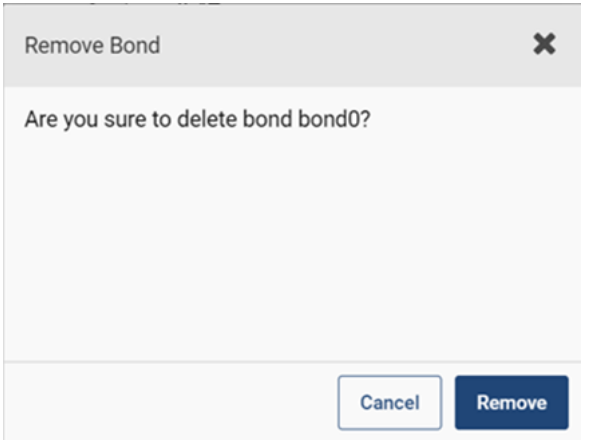

4 After you confirm, the bond removal is initiated. A pop-up is displayed to inform that this operation brings down the bond device and GUI is not available for some time. Click **OK**.

The server unreachable message can appear when the operation is in progress.

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

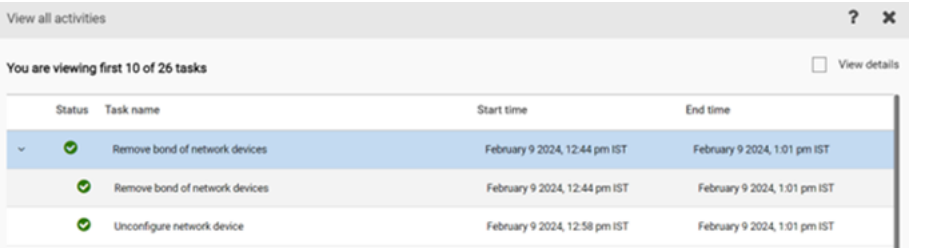

On successful removal of the bond, the management network is moved to eth1. The management network name is changed to **Management Network**.

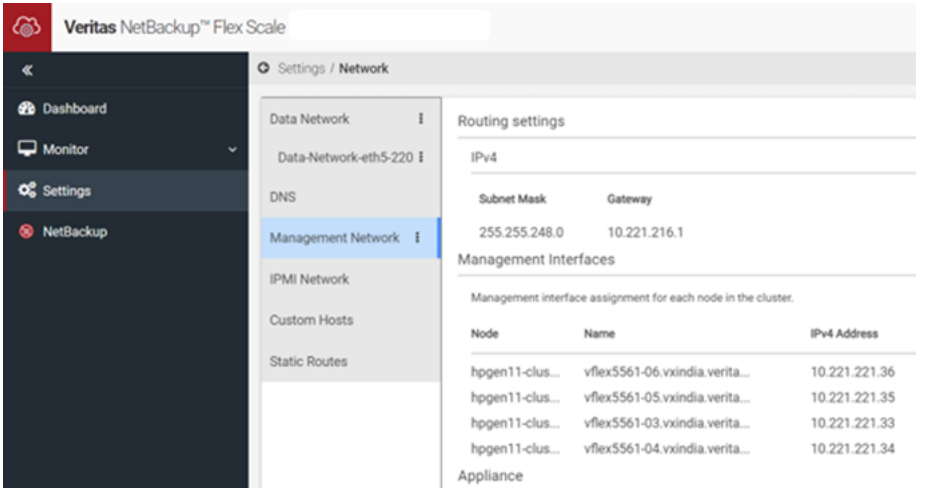

Configuring NetBackup Flex Scale in a non-DNS environment

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**.

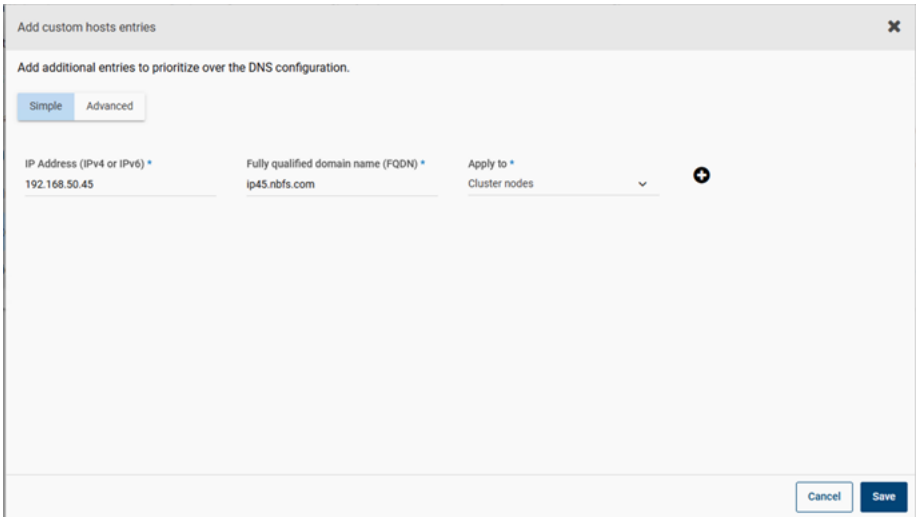

  You can use the **Advanced** tab to add multiple IP/FQDNs entries at the same time.

- 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**.

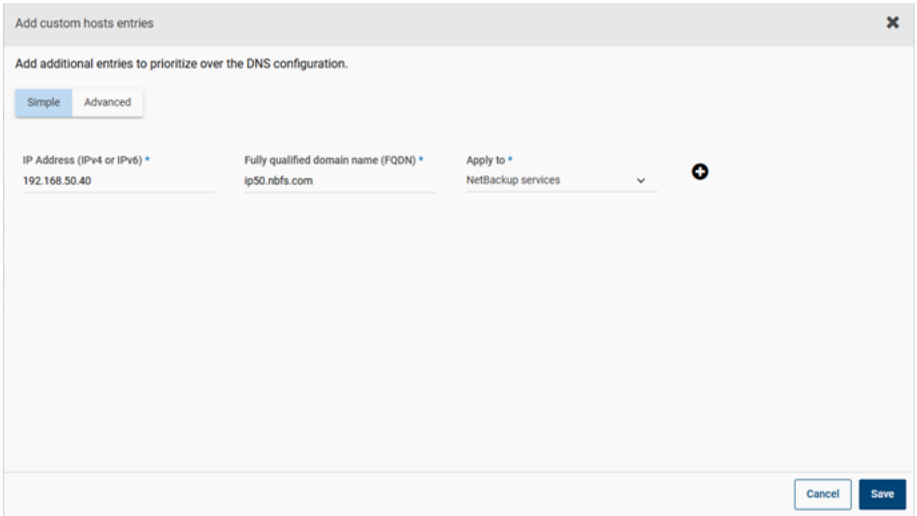

You can use the **Advanced** tab to add multiple IP/FQDNs entries at the same time.

- 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**.

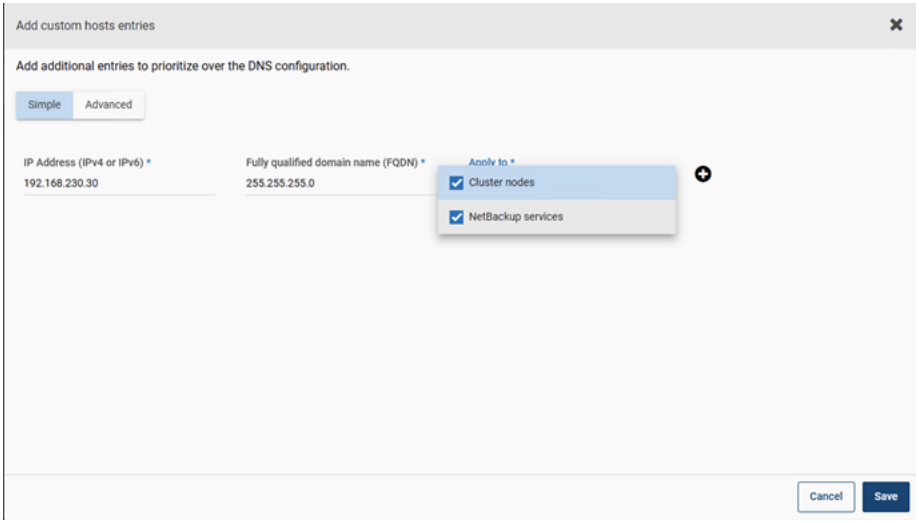

You can use the **Advanced** tab to add multiple IP/FQDNs entries at the same time.

- 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**.

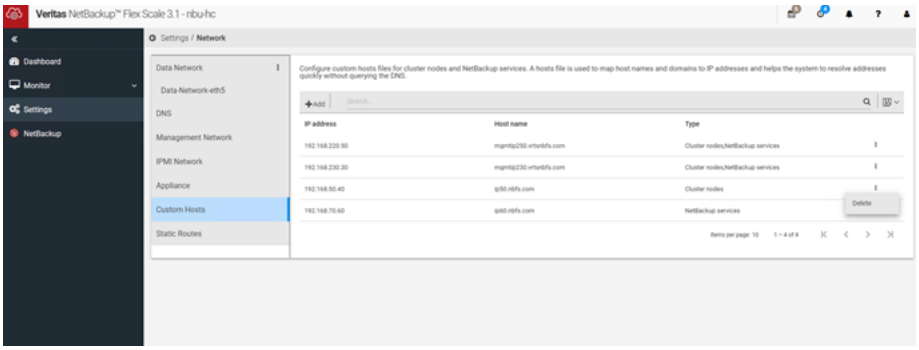

- Delete IP/FQDN information from NetBackup services. Click on the **Actions** button (vertical ellipsis) on the right side of the IP address and select **Delete**.

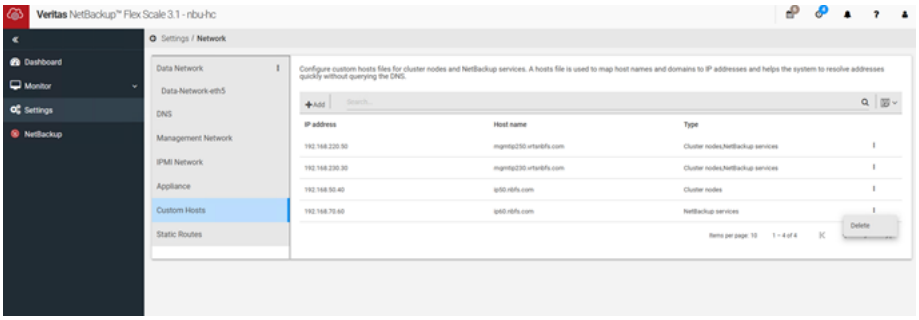

- Delete IP/FQDN information from cluster nodes and NetBackup services. Click on the **Actions** button (vertical ellipsis) on the right side of the IP address and select **Delete**.

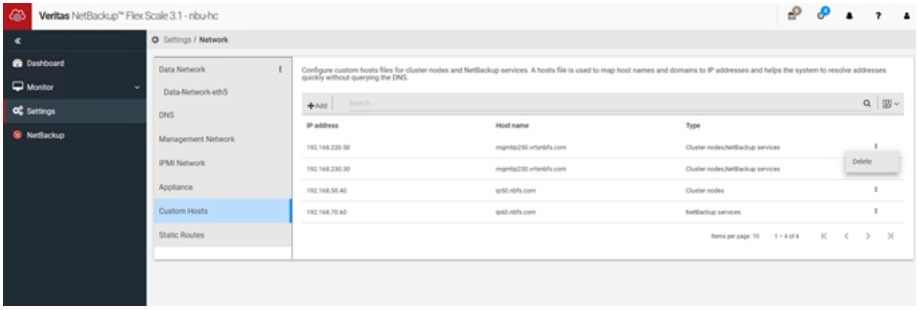

Data network configurations

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.
- The `fd00:200/120` network is reserved and used internally by NetBackup Flex Scale, and it should not be used anywhere.
- A single NetBackup Flex Scale cluster supports up to 16 secondary data networks.
- If disaster recovery is configured using NetBackup catalog replication, additional data networks can be added and managed independently on each NetBackup Flex Scale cluster. But when an additional data network is configured with a new NetBackup primary server FQDN, it must be first configured on the cluster where the NetBackup primary server is running (primary cluster). You can add an additional data network on the secondary cluster only if the FQDN is not provided or if it already exists on the primary cluster. Else, the operation to add the additional data network fails.
- In a multi-VLAN environment, AIR target domain is supported only on the primary VLAN network. It is not supported on the secondary VLAN network.

You can configure a cluster without a DNS server.

See “Configuring NetBackup Flex Scale in a non-DNS environment” on page 130.

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 136.

- VLAN on eth5
  See “Network configuration on VLAN (eth5)” on page 143.
- Bonded interfaces (bond0 using eth5 and eth7)
  See “Network configuration on bonded interfaces (bond0 on eth5 and eth7)” on page 144.
- VLAN on bond (VLAN on bond0)
  See “ VLAN on bond of eth5 and eth7 (bond0)” on page 145.

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 151.

Choosing the correct input method for data network configuration

You can provide inputs for IP/FQDN either using the **Automatic** tab or **Custom** tab from the GUI for operations such as initial configuration, add data network and add node.

Automatic tab:

The **Automatic** tab accepts only IP address. 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. The FQDN is retrieved using `nslookup` depending on the DNS server and search domain entry provided for that network. For `nslookup` to work correctly, the DNS server provided should be valid. In case only the short host name is specified as the FQDN, the correct search domain strings should be specified for `nslookup` to work correctly. The IP/FQDN validation check fails in case there is no information about the IP address in the DNS server provided.

Custom tab:

The **Custom** tab accepts both IP address and FQDN and does validation using forward and reverse lookup of the IP address. When **Custom** tab is used for providing input, the IP/FQDN validation check only fails when there is a conflict between the IP address and FQDN. There can be a conflict:

- If the forward lookup of an FQDN returns an IP address which is different than the input IP provided.

- If the reverse lookup of an IP address returns an FQDN which is different than the input FQDN provided.

In a non-DNS environment, if the **Custom** tab is used, the IP/FQDN validation check is skipped.

Consideration while choosing **Automatic** or **Custom** tab for providing input for IP/FQDN.

- During initial configuration, if the IP/FQDN entries exist in the DNS, use the **Automatic** tab for providing input IP.
- During initial configuration, if the DNS server is not specified or if the IP/FQDN entries do not exist in the DNS, then **Custom** tab should be used to provide the inputs for both IP address and FQDN.
- During data network addition, if DNS server in configured for that network, then **Automatic** tab is enabled by default. **Automatic** tab can be used if the IP/FQDN entries are already added in DNS.
- During data network addition, if DNS server is not configured or if the IP/FQDN entries do not exist in the DNS server, then the **Custom** tab should be used for providing input for both IP/FQDN.
- During node addition operation, if DNS is configured for the data network, the **Automatic** tab should be used for providing input for data network IPs.
- During node addition operation, if DNS is not configured or if the IP/FQDN entries do not exist in the DNS server then the **Custom** tab should be used for providing input for both IP/FQDN.
- During node addition operation, when a secondary data network exists on the management interface (*eth1* or *eth1.<VLANID>*), the **Custom** tab should be used for providing input for both IP/FQDN for the secondary data network IPs.

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**.

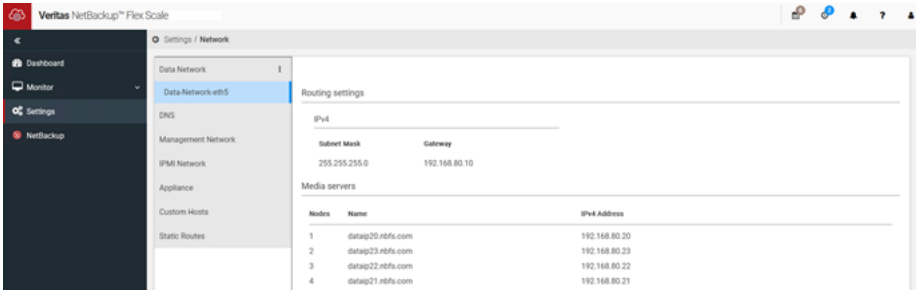

The following operations are supported:

- Adding a new data network
  See “Adding a data network” on page 144.
- Modifying the secondary data network
  See “Modifying a data network” on page 141.
- Deleting the secondary data network
  See “Deleting a data network” on page 142.

Adding a data network

You can add a new data network.

Note: A single NetBackup Flex Scale cluster supports up to 16 secondary data networks.

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**.

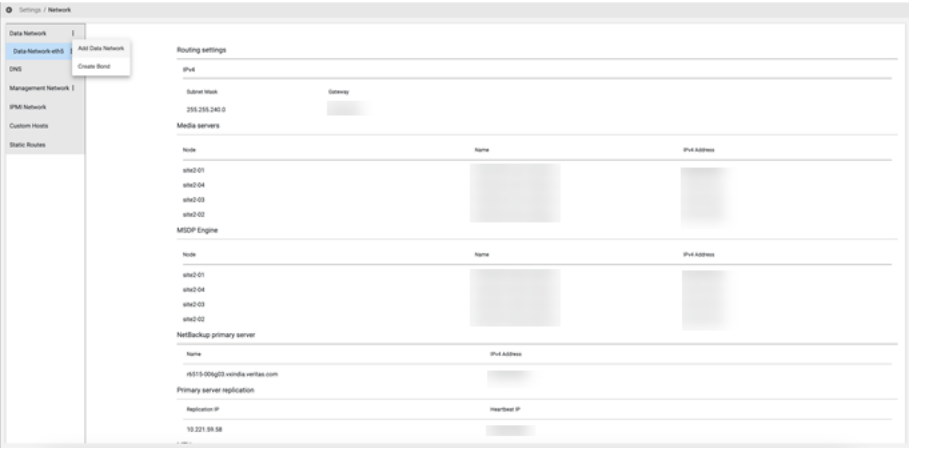

3 The **Add data network** form appears. Provide the necessary details.

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. This is an optional parameter.
Media servers	Enter the FQDN and IP address for each media server as per the input method selected.
MSDP engines	Enter the FQDN and IP addresses for each MSDP engine as per the input method selected.
Primary server	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.

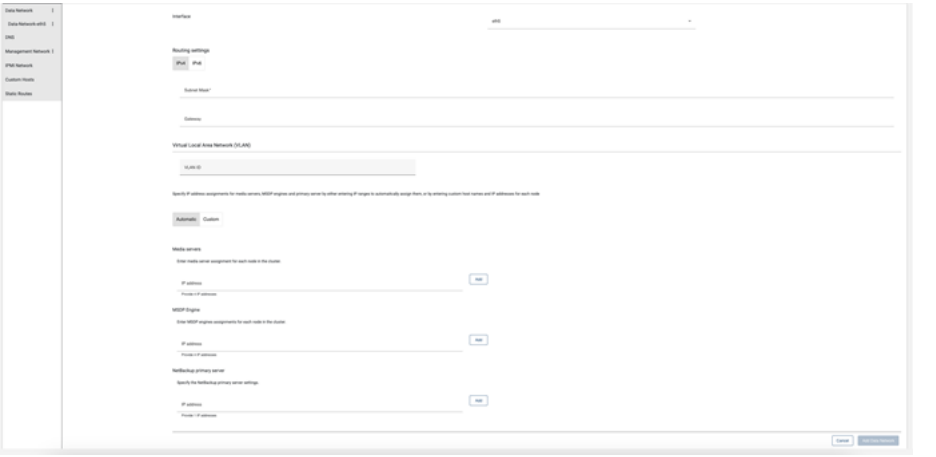

4 Click **Add Data Network**. The process for adding the network is initiated.

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.

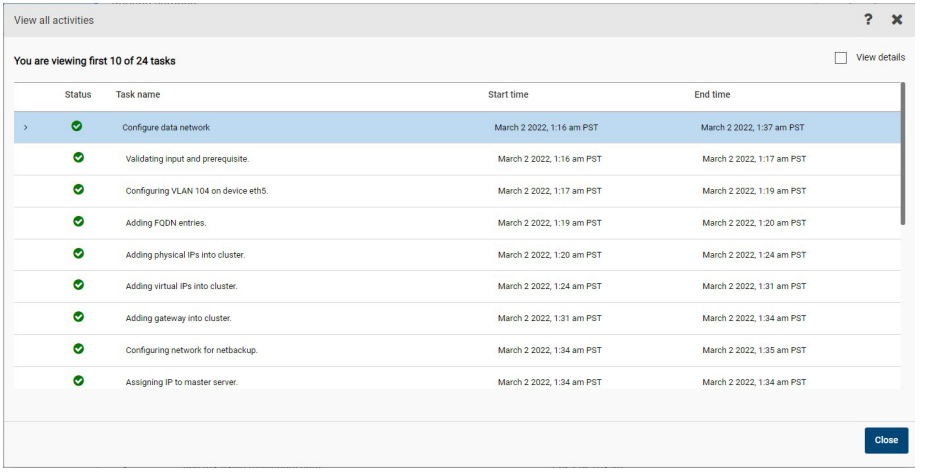

6 Once the data network is added successfully, the newly added data network appears under **Data Network**.

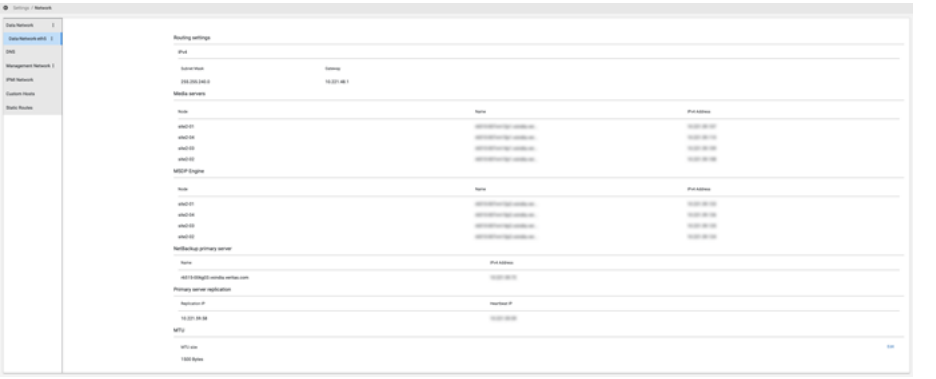

Modifying a data network

You cannot modify any of the attributes of the primary data network.

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

- VLAN ID

- IP address/FQDN of the primary, media, and storage servers
- VLAN ID of the secondary data network.

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

Modify data network functionality can also be used to add or remove individual components such as, IP/FQDN for media server, storage server, or primary server.

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.

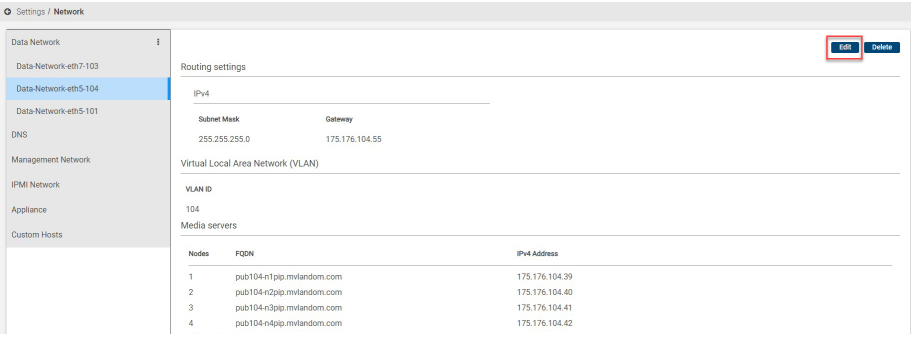

Deleting a data network

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.

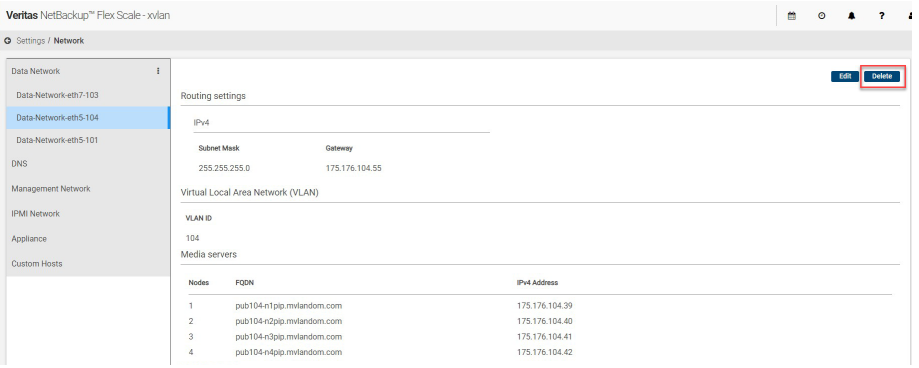

Network configuration on VLAN (eth5)

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*>**.

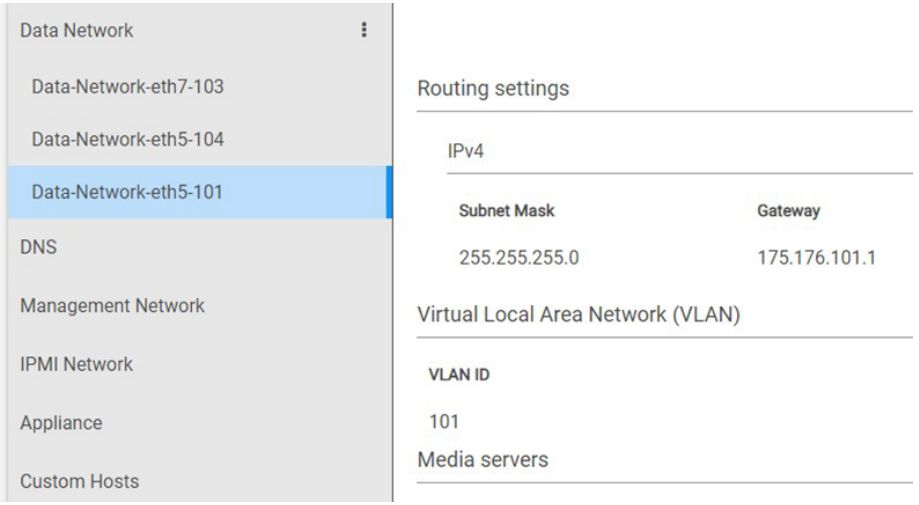

Note: The primary data network cannot be modified or deleted.

The following operations are supported on the secondary data network:

- Adding a new data network
  See “Adding a data network” on page 137.
- Modifying a data network
  See “Modifying a data network” on page 141.

- Deleting a data network
  See “Deleting a data network” on page 142.

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**.

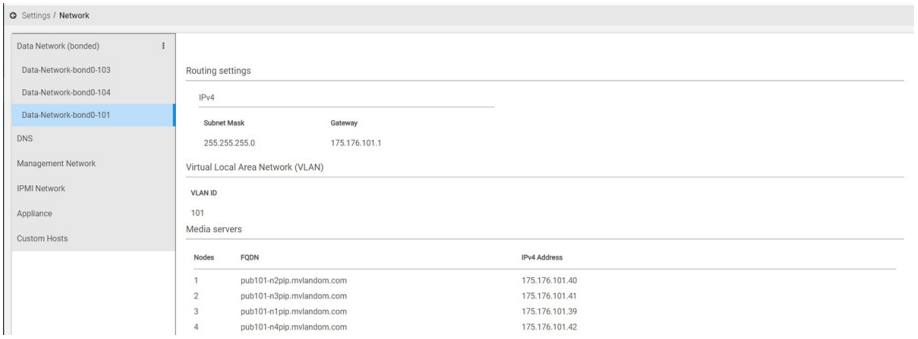

Note: The primary data network cannot be modified or deleted.

The following operations are supported on the secondary data network:

- Adding a new data network
  See “Adding a data network” on page 144.
- Modifying a data network
  See “Modifying a data network” on page 141.
- Deleting a data network
  See “Deleting a data network” on page 142.

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 This is an optional parameter.
Media servers	Enter the FQDN and IP address for each media server as per the input method selected
MSDP engines	Enter the FQDN and IP addresses for each MSDP engine as per the input method selected
Primary server	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*>**.

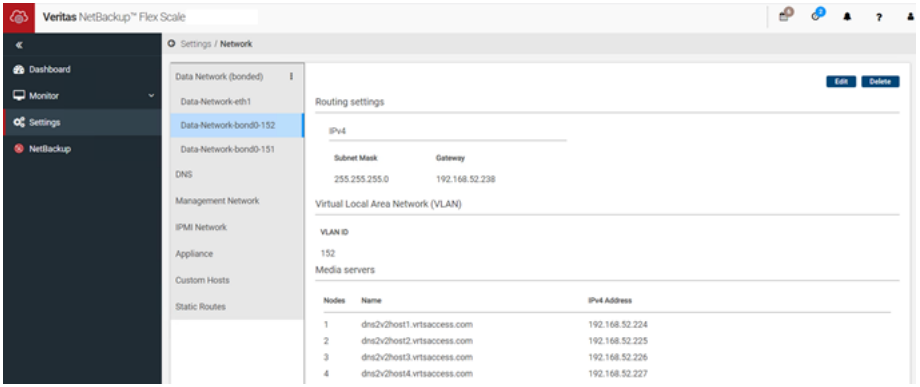

Note: The primary data network cannot be modified or deleted.

The following operations are supported on the secondary data network:

- Adding a new data network
  See "Adding a data network" on page 137.
- Modifying a data network
  See "Modifying a data network" on page 141.
- Deleting a data network
  See "Deleting a data network" on page 142.

Network configuration on management interface (eth1)

You can add a new data network on the management interface (eth1).

To add a new data network on eth1

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 Provide the required 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.

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

Interface	Choose eth1.
Media servers	Enter the FQDN and IP address for each media server as per the input method selected.
MSDP engines	Enter the FQDN and IP addresses for each MSDP engine as per the input method selected.
Primary server	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: When you select eth1 from the drop-down list, the routing details for subnet, gateway and VLAN ID values are auto populated based on the existing management network. These values cannot be modified.

The input IPs for media/MSDP engine/primary server should belong to same subnet as the existing management network.

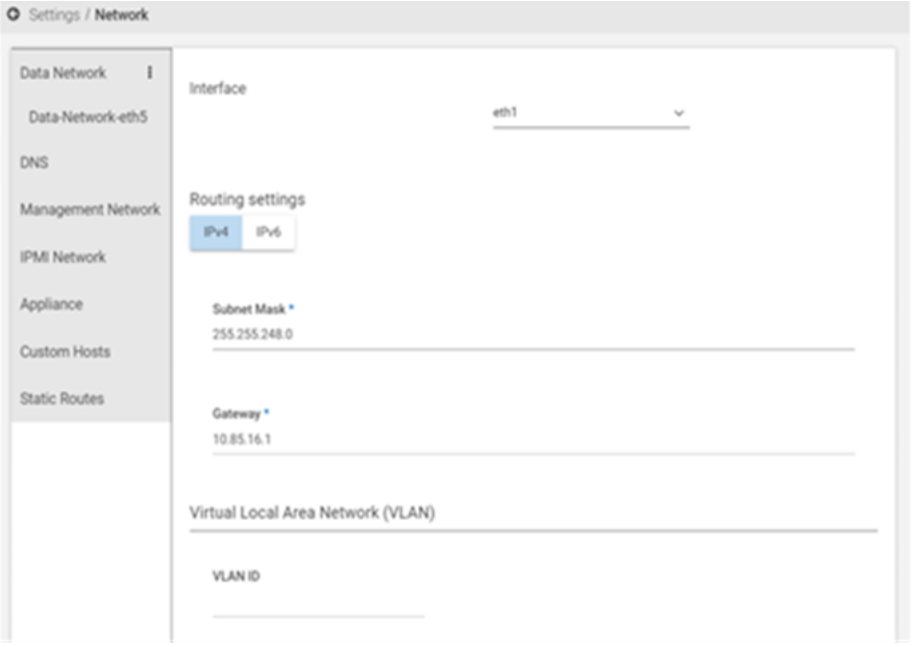

Note: The secondary data network to be added on the management interface must be in the same network as the existing network of the management interface. The secondary data network to be added on the management interface cannot be the same as any existing data networks (primary or secondary).

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**.

Network configurations for adding a partial data network

A secondary data network can also be added with individual components, either only with media server or storage server or primary server. For example, you can add a secondary data network with only media server IPs on data network interface eth7. The secondary data network can be added for all possible combinations between media, storage, and primary server.

To add a new partial 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. Provide the required 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.

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

Interface	Choose eth1 or eth5 or eth7.
Routing settings	Enter the subnet mask and gateway for the network. This field is auto populated for eth1.
VLAN ID	Enter the VLAN ID. This field is auto populated for eth1. It is optional for other interfaces.
Media servers	Enter the FQDN and IP address for each media server as per the input method selected.
MSDP engines	Enter the FQDN and IP addresses for each MSDP engine as per the input method selected.
Primary server	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: After a secondary data network is added with partial components, you can perform modify data network and delete data network operations.

Note: For adding a partial secondary network for media and storage server, the input for IPs for all the nodes is required. There is no option available to add a single media server IP or a single storage server IP to one of the nodes.

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**.

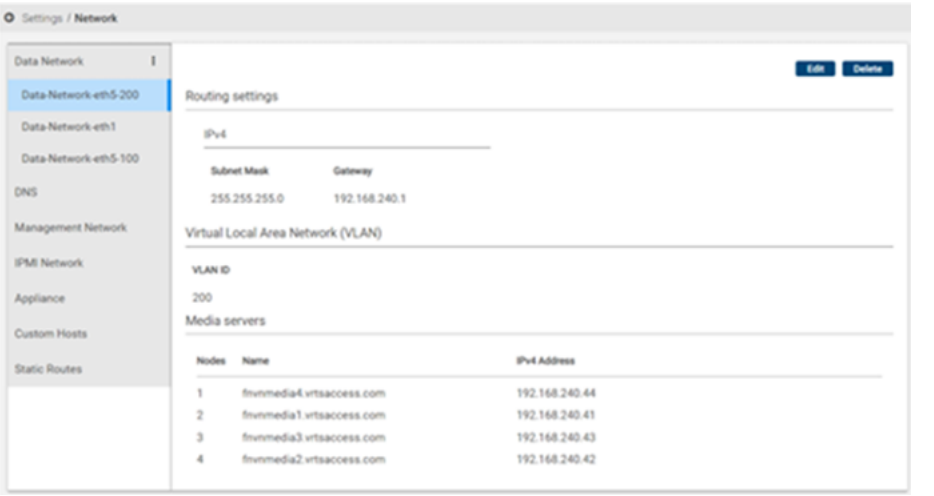

You can add multiple secondary data networks over a device, VLAN, or bond. All the data networks are displayed in the GUI under **Data Network** in the left pane. The first secondary data network on the device is named as **Data-Network-eth5** and the subsequent secondary data networks are named as **Data-Network-eth5_1**, **Data-Network-eth5_2** and so on.

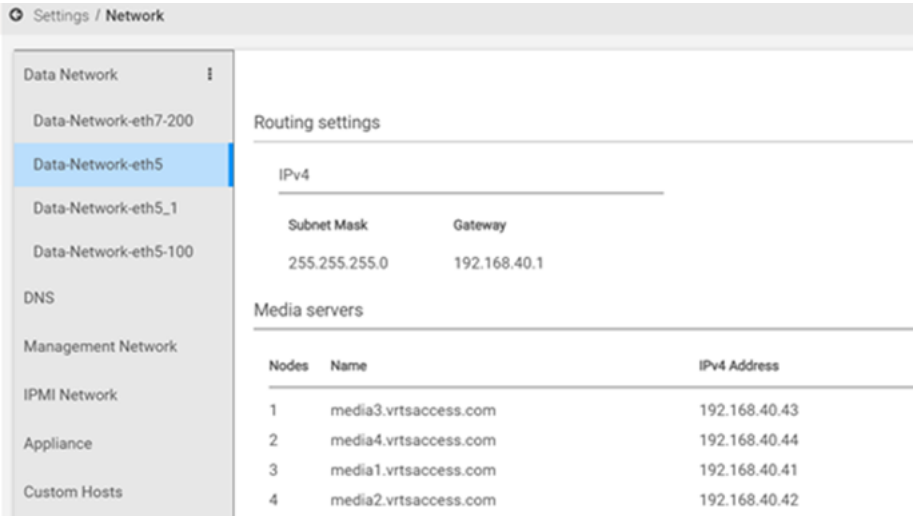

Support for multiple VLAN when disaster recovery is configured

When disaster recovery is configured between two sites using catalog replication between two NetBackup Flex Scale clusters, the primary service is active on only one cluster at any given time. However, the primary service manages the media services, storage servers, and storage units across both clusters.

- Each cluster requires independent configuration and management of additional data networks, with the number of data networks varying depending on network needs.
- Before setting up catalog replication, only the primary cluster can have additional data networks configured. The secondary cluster cannot have any additional data networks.
- If disaster recovery is configured to use a single virtual IP for the NetBackup primary service FQDN, all additional data networks must share the same virtual IP configuration for the primary service FQDN.
- You cannot add data networks on both the primary and secondary clusters simultaneously.
- When performing a replication role change (moving the primary service between clusters), ensure that the DNS server (if configured) is updated with the correct virtual IP for the NetBackup primary service on all data networks of the target cluster.
- If an additional data network introduces a new primary service FQDN, you must configure it first on the cluster where the primary service is active. Only then

can the same FQDN be used when configuring the additional data network on the secondary cluster.

- The **Automatic** option in the **Additional Data Network Configuration** page relies on DNS resolution to map the virtual IP to the correct FQDN. If the primary service uses different virtual IPs on each cluster, the DNS server might not resolve the FQDN correctly when adding it to the secondary cluster, as it may already be resolving to the primary cluster's virtual IP. To avoid this, either use the **Custom** option or update the DNS to correctly resolve the FQDN before adding the data network.

Configuring static routes on a NetBackup Flex Scale cluster

You can configure static routes on a NetBackup Flex Scale cluster to communicate with remote clients who are on a different subnet, if the default route is not suitable.

A static route is a normal Linux route which appears in the output of the `ip route show` command on the Linux server. You must specify the destination IP address, netmask, gateway IP and device to configure the new route. Once you configure static routes for the cluster nodes and NetBackup services, the NetBackup services can reach clients who are part of another data network.

You can view, add and delete static routes using both the GUI and RESTful APIs.

You can manage static routes using the NBFS GUI by navigating to **Settings** > **Network** > **Static routes**.

To add a new static route

1. In the **Network** > **Static routes** page, Click **Add**.
2. In the **Add static route** window, specify the destination IP address, netmask, gateway IP and device. Click **Save**.

   Note: You can only add one static route at a time.

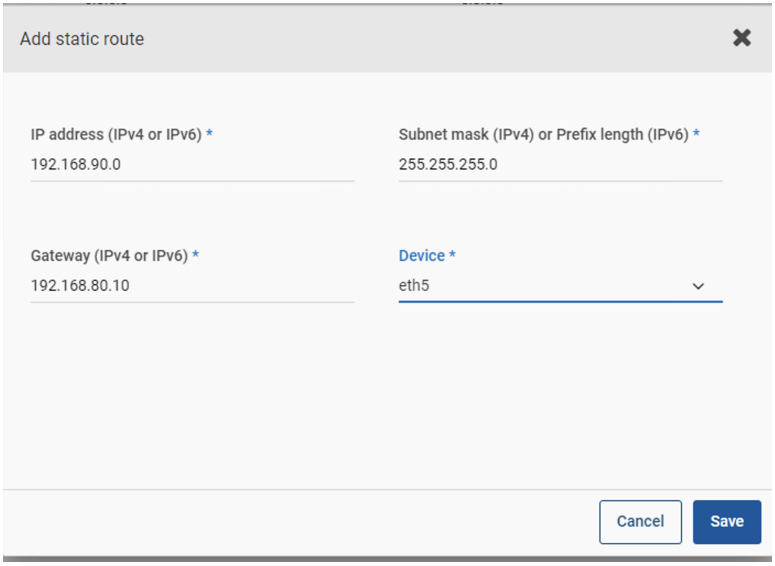

3. The newly added static route will be visible in the list of static routes displayed.

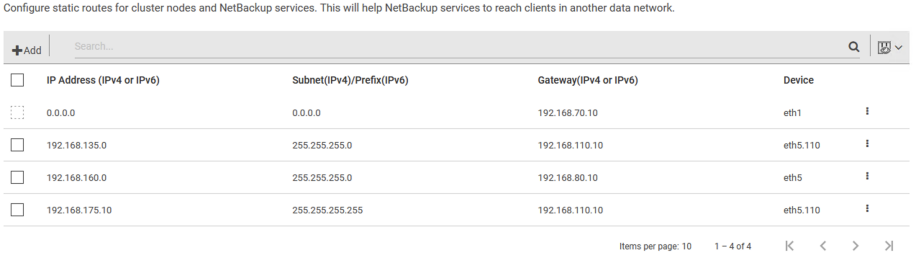

You can either delete a specific static route or delete multiple static routes at the same time.

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

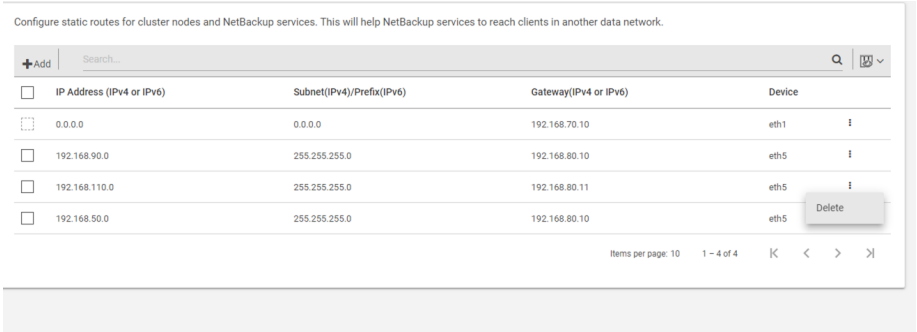

- To remove multiple static routes, select all the routes that you want to delete and select **Delete**. A confirmation dialog box appears. Select **Delete**.

Notification events are raised when new static routes are added to the cluster or when static routes are deleted from the cluster.

The following RESTful APIs are available to configure and manage the static routes:

- `GET /api/appliance/v1.0/network/static-routes`: Get the available static route.
- `PATCH /api/appliance/v1.0/network/static-routes`: Payload for updating static routes to network configurations.
- `GET /api/appliance/v1.0/tasks/{taskId}`: Get the progress of add/delete static routes operations
- `GET /api/appliance/v1.0/network/interfaces`: Get the list of public device interfaces

Chapter 5

NetBackup Flex Scale infrastructure monitoring

This chapter includes the following topics:

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 156.

See "Managing alerts" on page 156.

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.

You can resolve software alerts from the GUI.

To resolve the alerts

1. Go to **Settings** > **Alert management**.
2. For the selected alert, click the **Actions** menu (vertical ellipsis) from the right side of the row and select **Resolve**.
3. The alert no longer appears in the GUI and an AutoSupport resolution mail is sent to both the user and NetInsights. If the issue persists, then a new alert is generated again.

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 158.

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 164.
- To receive email notifications for the generated alerts, See "Setting up email alerts" on page 159.
- To receive SNMP notifications, See "Setting up SNMP alerts" on page 161.

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, the FQDN, or the short host name 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 Management > Cluster settings > AutoSupport**.
   - 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Settings > AutoSupport**.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

2 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.

3 Click **Save**.

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

4 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 and SNMPv3 application protocol.

To configure the SNMP settings:

1 **Cluster dashboard**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, the FQDN, or the short host name 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 Management > Cluster settings > AutoSupport**.
- 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Settings > AutoSupport**.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

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

Select **Enable SNMP** to enable the SNMP alerts.

Specify the **SNMP version**.NetBackup Flex Scale supports SNMP-v2 and SNMP-v3 prototcols.

If you select SNMP-v2, specify the following details:

Parameter	Description
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**.

If you select SNMP-v3 specify the following details:

Parameter	Description
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.

Parameter	Description
SNMP port	Port number of the SNMP server. The default port is 162. **Note:** Your firewall must allow access from the appliance to the SNMP server through the configured port.
SNMP username	SNMP user name
Authentication protocol	Specify the authentication protocol. It provides authentication based on the HMAC-SHA algorithms. Choose one of the following options: ■ None ■ SHA256 ■ SHA512 This is a mandatory field.
SNMP password	Enter the password.
Encryption protocol	Specify the encryption protocol. If provides DES 56-bit encryption in addition to authentication based on the AES standard. ■ None ■ AES128 ■ AES192 ■ AES256 ■ AES512 This is a mandatory field.

Parameter	Description
Encryption passphrase	Enter the passphrase that you want to use for encryption.

3 Click **Save**.

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

The MIB file is located in the `/opt/autosupport/VERITAS-APPLIANCE-MONITORING.mib` location on a NetBackup Flex Scale node. After configuring SNMP successfully, copy the MIB file from the `/opt/autosupport/VERITAS-APPLIANCE-MONITORING.mib` location to your SNMP manager to receive SNMP traps. SNMP traps can only be send out to the configured destination, inbound SNMP queries and walks are not possible as there is no SNMP service on the NetBackup Flex Scale nodes.

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 159.

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, the FQDN, or the short host name 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 Management > Cluster settings > AutoSupport**.
- 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Settings > AutoSupport**.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

---

2 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. (Required if you enable the proxy server)
Proxy password	Password to authenticate the user name that is used to log in to the proxy server. (Required if you enable the proxy server)

3 Click **Save**.

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

4 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 infrastructure management UI" on page 17.
  See "About alert management" on page 155.
- Send email alerts.
  See "Setting up email alerts" on page 159.
- Send SNMP alerts.
  See "Setting up SNMP alerts" on page 161.
- Share hardware monitoring information for analysis and troubleshooting with Veritas Support using Call Home.
  See "Configuring Call Home settings" on page 164.
- Locate disks.
  See "Locating the disks" on page 171.
- Monitor deviations.
  See "Monitoring deviations in firmware, driver, and utilities" on page 169.

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:

- Adapters
- CPU
- DIMM
- Fan
- Fibre Channel HBA (if installed)
- Hard disk
- Network card

- PCI
- Power supply
- RAID
- SSD
- Temperature

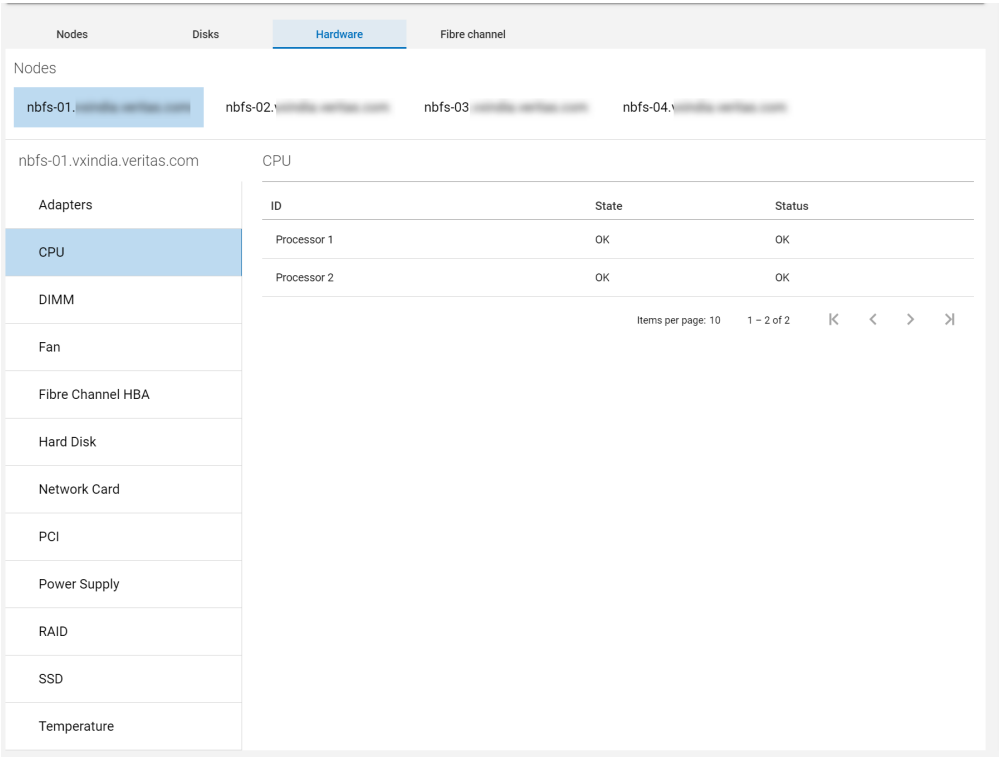

If a hardware component fails or reports an error, an alert icon is displayed next to the component and the node it belongs to.

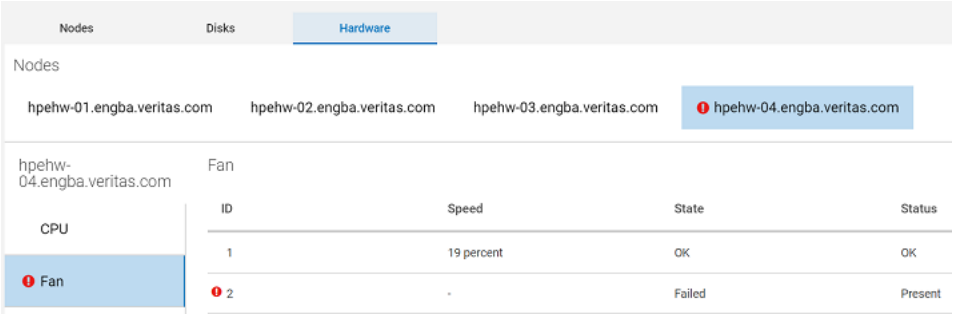

Using the shell menu

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

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

Monitoring deviations in firmware, driver, and utilities

If you plan to leverage your own hardware instead of the out of the box appliance, you can monitor the firmware, driver, and utility versions that are installed on your hardware and report deviations, if any, with Veritas-supported versions.

You can use the following commands to monitor the deviation:

- `show hardware-health node component=Firmware`
- `show hardware-health node component=Driver`
- `show hardware-health node component=Utility`

The following example shows the output that is displayed when you run the `show hardware-health node component=Driver` command on an HPE 5551 model:

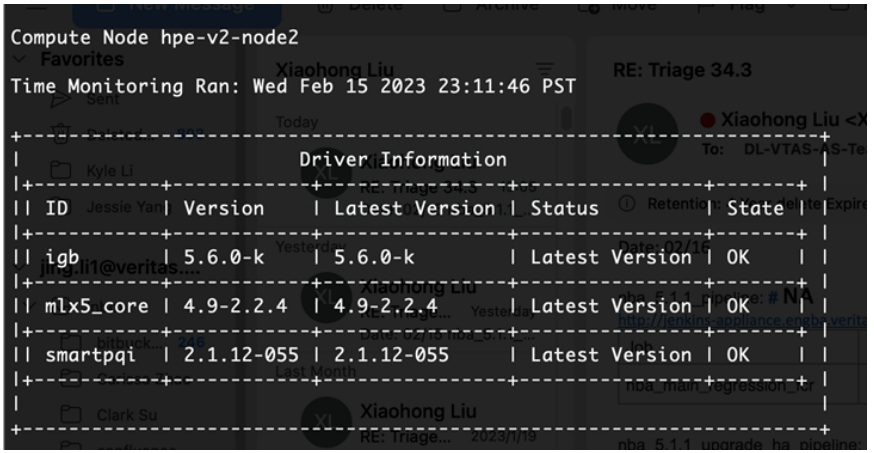

The **Status** column of the firmware, driver or utility output shows one of the following values:

- **Latest Version**: The latest supported version is installed. No action is required.
- **Update Available**: You must contact Veritas Support to get the required update package.
- **Unsupported Version**: The detected version is not supported by Veritas and you must contact Veritas Support. An alert is generated if an unsupported version is detected.
- **Fail**: The service that detects the versions did not run. Wait for a maximum of 24 hours for the next service cycle to run.

Performing health check for the cluster

You can use the `support health check area=`*`area`* `nodename=`*`nodename`* command to validate the network, OS, security, and protocol configuration for the cluster. You can run this command from the cluster-level CLI for a configured NetBackup Flex Scale cluster. The NetBackup Flex Scale cluster-level CLI can be accessed using the console IP address. You can access it by logging into the system using the console IP address with your credentials. For example: `ssh admin_user@console_ip`.

The command checks for any discrepancies in the following areas:

- **all**: All health validations (default).
- **network**: Validate network configuration.
- **os**: Validate OS configuration.

- **security**: Validate STIG, FIPS, lockdown mode configuration.
- **protocols**: Validate protocol configuration.
- **va_config**: Validate fencing configuration.

The output of the health check is saved to a file in JSON format at `/log/VRTSnas/log/`. The naming convention for the JSON file is `/opt/VRTSnas/log/health-check-`*`area-nodenames-time`*`.json`. For example, `/log/VRTSnas/log/health-check-all-nbfs-01-nbfs-02-nbfs-03-nbfs-04-20230417021610.1681722970.json`

You can perform the validations for all the cluster nodes or for specific nodes by specifying a comma-separated list of nodes.

For more details about this command, see the *Veritas NetBackup Flex Scale Command Reference Guide*.

Locating the disks

You can turn on the beacon to identify the box and the bay in which the disk is inserted. The beacon on the disk flashes blue for a duration of one minute. This option is not available for NVMe disks, disks that don't have physical disk ID mapping, or if one logical disk is mapped to multiple physical disks. The **Infrastructure > Disks** page shows the logical disk name in the **Name** column and the corresponding physical disk in the **Disk ID** column.

To locate a disk:

1. Use any one of the following options to log on to NetBackup Flex Scale:
   - 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Cluster Management > 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration, and then in the left pane click **Monitor > Infrastructure**.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

2 To identify the disk, click **Disks** or **Hardware > Hard Disk**, and then click the Actions menu (vertical ellipsis) from the right side of the row and click **Turn on beacon**.

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*.

Chapter 6

Resiliency in NetBackup Flex Scale

This chapter includes the following topics:

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.

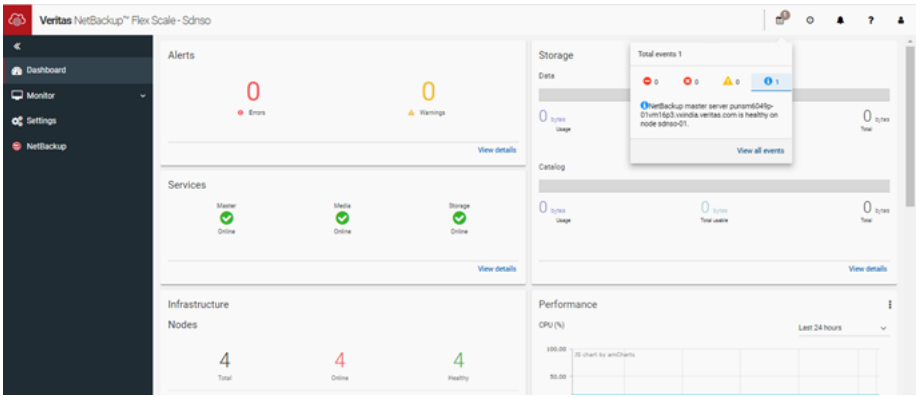

The bell icon in the top navigation bar shows the status of current alerts and alerts that were recently completed. An AutoSupport alert is raised if the primary service is offline.

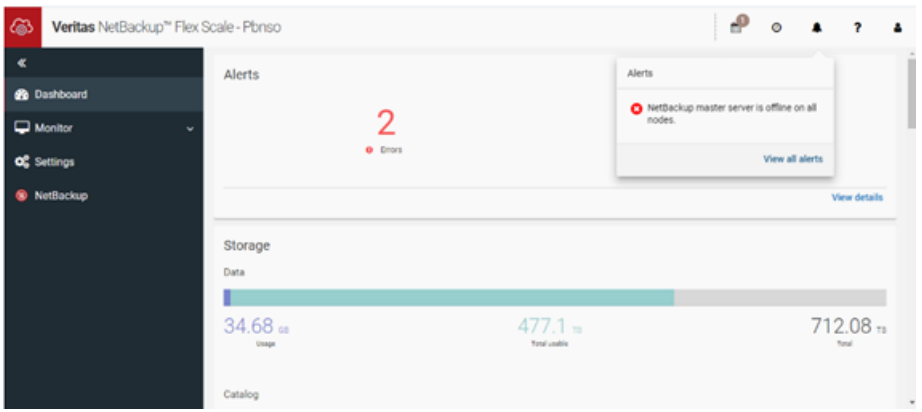

If the primary service goes offline, it is failed over and comes online on another cluster node. For details on the behavior of the NetBackup primary service in clustered environments and handling the jobs, see the *General notes on clustered primary service administration* chapter in the *Veritas NetBackup™ 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 54.

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 178.

See "Performing a recovery of the catalog file system using REST APIs" on page 182.

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.

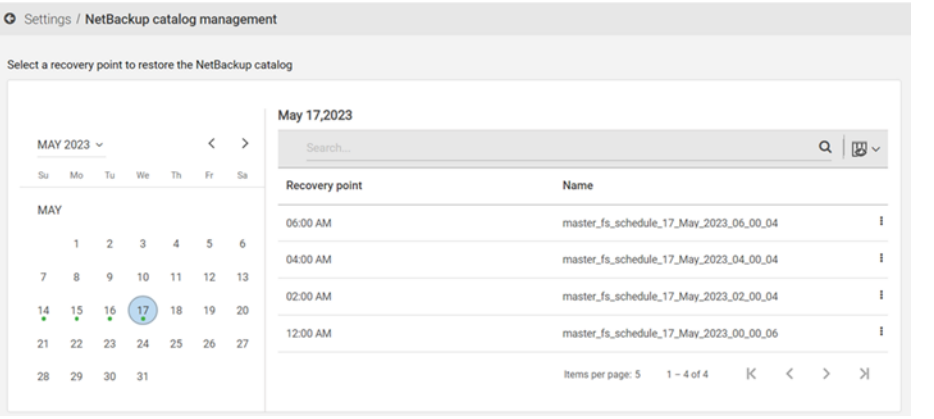

Note: The green dot below a date in the calendar indicates that there are valid checkpoints available for that date.

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 **Actions** menu (vertical ellipsis) from the right side of the row in the UI for a particular checkpoint.
2. The **Validate NetBackup catalog** pop-up opens. Click **Validate** to validate the checkpoint.

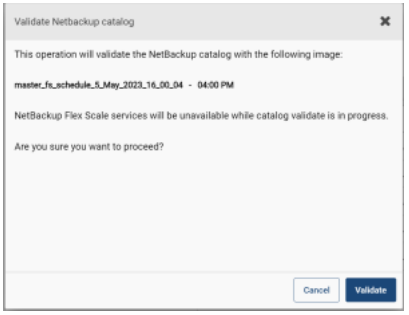

3. A task is initiated in the GUI which can be used to track the status of the validate operation. Once the validate operation is completed, the status of the task changes accordingly.

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.

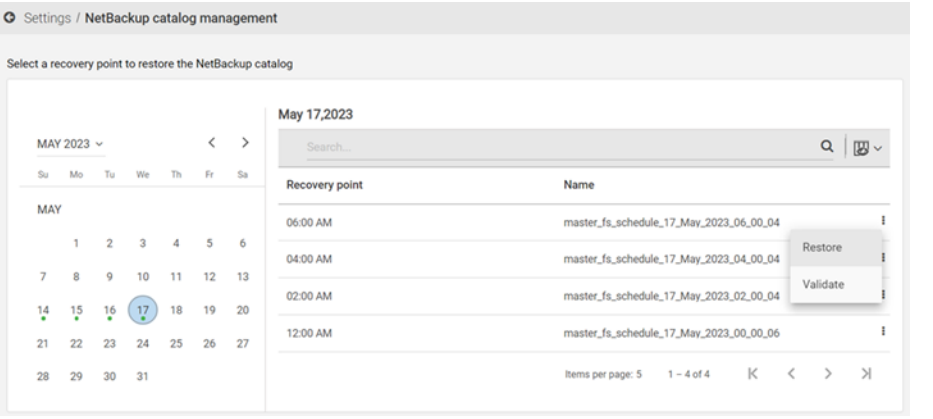

2 A pop-up window appears to ask for confirmation. Click **Restore**.

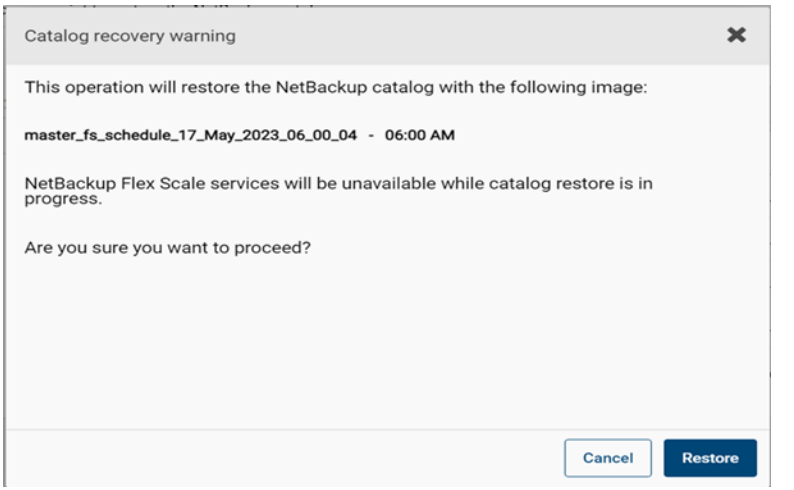

3 If a validation checkpoint exists for the selected checkpoint, the restore operation is performed immediately. If a validation checkpoint does not exist, the restore operation first validates the checkpoint and creates a validation checkpoint. Then, it performs a restore from the validation checkpoint.

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.

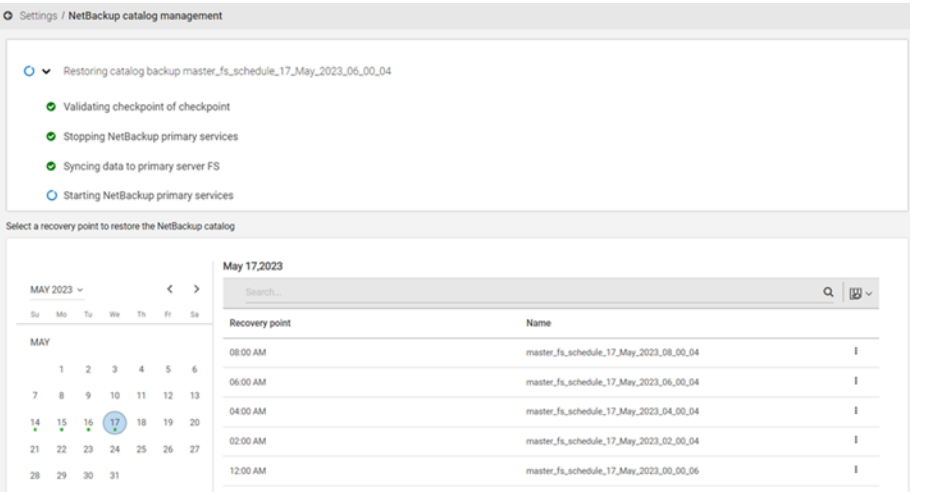

Performing a recovery of the catalog file system using REST APIs

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.

Chapter 7

EMS server configuration

This chapter includes the following topics:

Configuring an external BYOS media server

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

If you have an external media server writing to NetBackup Flex Scale storage server:

- Add the external media server to list of media servers who can write to that storage server.
- Do not use **Allow NetBackup to automatically select** option when selecting the media server. You must explicitly specify the list of media servers to be used.
- When you add a new node (which adds a new media server), you must add that media server explicitly.

To configure an external BYOS media server with NetBackup Flex Scale

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. Go to `/usr/openv/netbackup/bin/admincmd/` and run the following command:

```
# ./nbemmcmd -addhost -machinename ems_fqdn -machinetype
media -primaryserver primary_server_fqdn -operatingsystem
Linux -netbackupversion 10.0

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

   where *ems_fqdn* is the external media server FQDN and *primary_server_fqdn* is the primary server FQDN.
6. Restart the EMS server services.
7. 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

```
bpclntcmd -clear_host_cache
```

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

```
# bptestnetconn
# bptestbpcd -M primary FQDN
# bpclntcmd -hn primary FQDN
```

- After clearing the host cache, *bptestnetconn* paramenter 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 NBFS_Primary_Hostname
     ```

   - 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 Polices** 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.

Chapter 8

Site-based disaster recovery in NetBackup Flex Scale

This chapter includes the following topics:

About site-based disaster recovery in NetBackup Flex Scale

The NetBackup Flex Scale provides a disaster recovery solution by linking two NetBackup Flex Scale clusters from separate remote sites. 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

- Both the NetBackup Flex Scale clusters are configured to be part of single NetBackup domain. The NetBackup domain of the primary site is used once the disaster recovery is configured.
- The NetBackup service is active on only one site at a time and can failover between sites, if required.
- Clients on the primary site backup to the primary site and then the backup images are duplicated to the secondary site.
- You can configure disaster recovery on an active NetBackup Flex Scale Appliance (which has existing backup data and acts as the primary site).
- The NetBackup Flex Scale cluster to be configured as the secondary must be a freshly installed and configured cluster. It should not have any user data. It should not have lockdown mode configured. It should not have any additional data networks configured. It should not have node names or cluster name conflicting with the NetBackup Flex Scale cluster on the primary site. It should have the same number of nodes as the NetBackup Flex Scale cluster on the primary site.
- The NetBackup primary service can be configured to use either the same virtual IP or a different virtual IP when failing over the service between two sites. When using a configuration with a different virtual IP, DNS server(s) must be updated to correctly resolve the NetBackup primary service virtual IP depending on the target site to which the service fails over.
- You can choose to automatically update the NetBackup primary service host name mapping on the DNS servers that are configured with valid credentials. If you have not configured the DNS servers to be updated automatically then ensure that you have updated the DNS entries of the NetBackup primary service FQDNs to resolve to the secondary site.
- Primary site can have the secondary data network configured before disaster recovery configuration.

- While configuring disaster recovery, ensure that you configure ECA on the site that is going to act as the primary site.
- If ECA is deployed on the primary site before adding the secondary site, then you must redeploy ECA from the primary site after disaster recovery configuration is complete.
- 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.
- If you want to add a node to the cluster, See "Considerations for adding a node when disaster recovery is configured" on page 69.
- If you want to configure VLAN, See "Support for multiple VLAN when disaster recovery is configured" on page 151.
- 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*.

Note: Disaster recovery can be configured between a cluster with HPE model 5551 and a cluster with HPE model 5561.

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.
2. Go to **Settings** > **Primary service replication**. Click **Configure**.

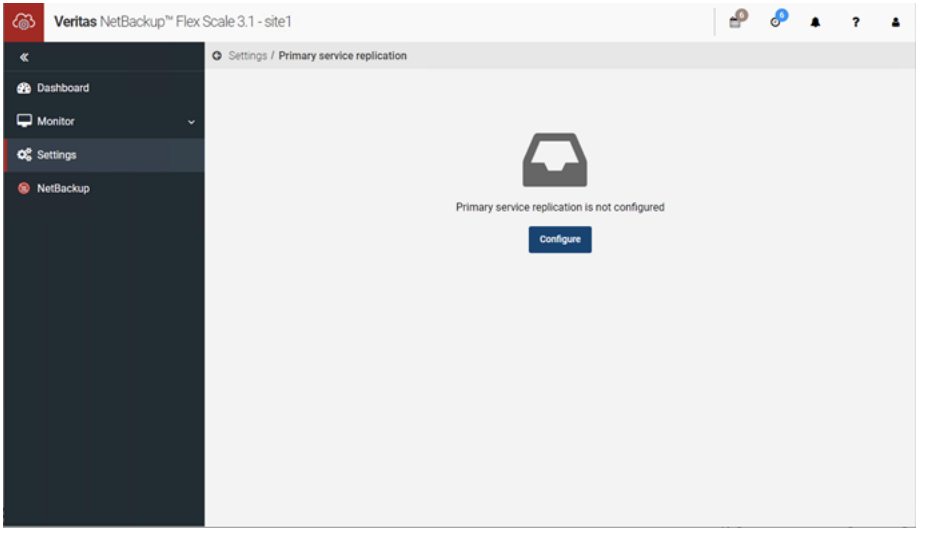

3. Enter the configuration details for the primary and secondary site.
   - Replication and heartbeat IP address for the primary and secondary sites. Provide IPv4 or IPv6 address according to the network configuration.
   - Management server IPv4/IPv6 address or FQDN.

     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.
   - Read the details provided for configuring NetBackup primary service virtual IP(s) and select or clear the check box as per your requirement.
     - If this option is selected, the NetBackup primary service uses the same virtual IP(s) on both clusters. Use this option if the same network or

VLAN is available and the same virtual IP(s) can be used on both the clusters.

- If this option is not selected, the secondary cluster uses the previously configured virtual IP for NetBackup primary service and DNS server(s) may have to be updated while migrating or failing over the NetBackup primary service between the two clusters.
- If this option is selected during disaster recovery configuration and the clusters are in different networks (different subnet and gateway) then the configuration fails as the clusters cannot use the same virtual IP.
- If this option is selected, automatic DNS update is not required and will not be available in **Settings -> Primary service replication**.
- This selection cannot be changed after disaster recovery configuration.

Click **Configure**.

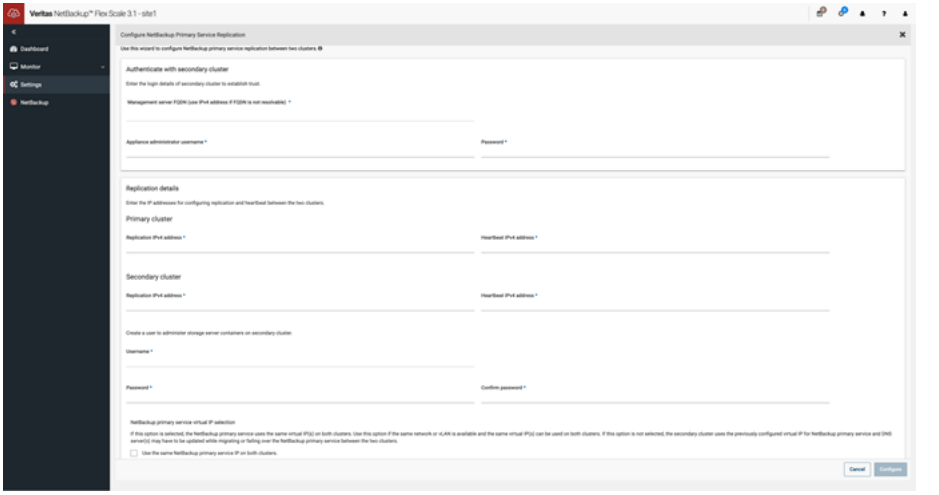

4 Verify the security certificate information of the secondary site that is displayed. Click **Trust**.

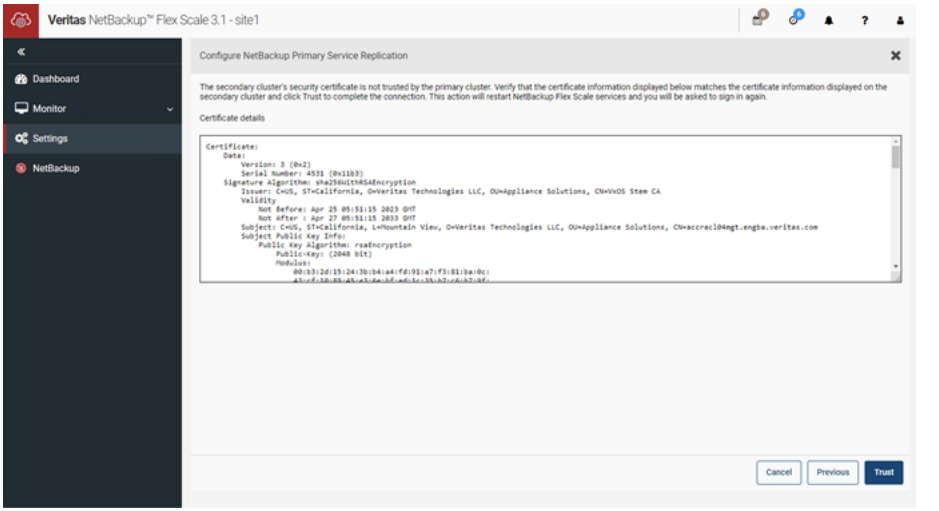

5 A warning pop-up appears that existing images, if any, on secondary site will be erased. Click **Erase images** to continue with disaster recovery configuration.

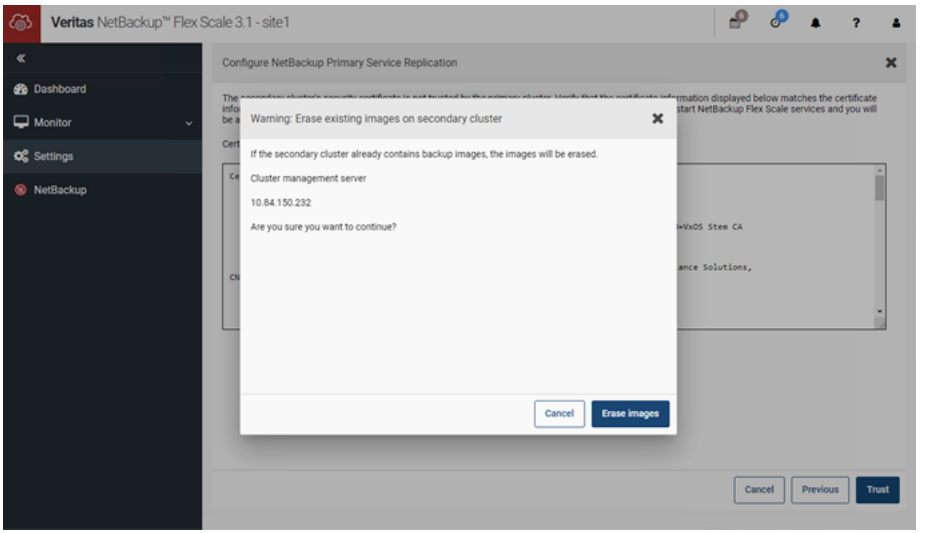

6 Disaster recovery configuration is initiated.

Cluster certificates are installed on both the clusters and the web server is restarted.

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

You can also monitor the configuration from the GUI of the primary site by navigating to **Cluster Management > Cluster settings > Primary service replication**.

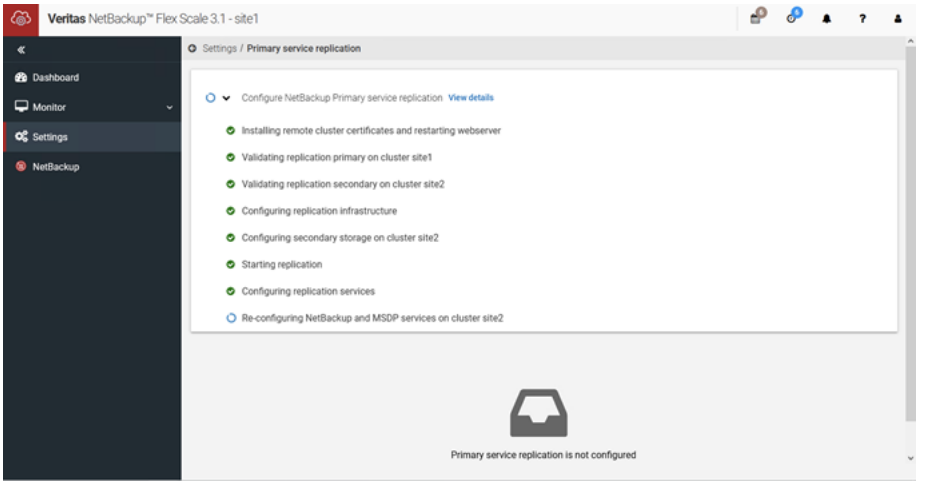

On the second cluster only 'Configure replication' task appears.

7 After the configuration is complete, you can also check the status by navigating to **Settings** > **Primary service replication** 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

Automated NetBackup SLP management

When you configure primary service replication between two NetBackup Flex Scale clusters, the NetBackup primary service catalog data is replicated between the clusters. You have to configure NetBackup Storage Lifecycle Policies (SLPs) for backup policies and protection plans to duplicate data between the two NetBackup Flex Scale clusters. You also must update the SLPs, backup policies, and protection plans if there is a cluster fault and the replication primary role needs to be changed to failover the NetBackup primary service.

NetBackup Flex Scale supports a mechanism to update the SLPs, backup policies, and protection plans automatically when changing the replication roles after a cluster fault. However, this functionality works only if the SLPs, backup policies and protection plans are using the default Media Server Deduplication Pool storage unit (STU) that is created with the NetBackup Flex Scale cluster configuration. This STU is named `stu_<storage_server_fqdn>`. If you use the default STU as storage target for your SLPs, backup policies and protection plans, you can select the **Update SLP configuration** option when changing replication roles.

Note: The **Update SLP configuration** option is available only when the current replication primary cluster is down and the **Make Primary** operation is performed on the replication secondary.

If you select the Update SLP configuration option, the following changes are performed on the NetBackup primary server:

1. Backup policies and protection plans that are not using SLPs and are using default STU as backup target:

   If the STU is from the faulted NetBackup Flex Scalecluster, it is updated to use the default STU on the new replication primary.

2. Backup policies and protection plans that are using SLPs with default STU as backup and duplication target:

   If the STU used by the backup target is from the faulted NetBackup Flex Scale cluster, it is updated to use the default STU on the new replication primary. The duplication target is updated to use the STU from the faulted NetBackup Flex Scale cluster.

3. Backup policies and protection plans that are using replication template SLPs:

   - The following SLP templates are created while configuring primary service replication.
     - <cluster1>_7days_to_<cluster2>_7days

- <cluster1>_30days_to_<cluster2>_30days
- <cluster2>_7days_to_<cluster1>_7days
- <cluster2>_30days_to_<cluster1>_30days

- Here, the first cluster name in the SLP is the backup target and the second cluster name in the SLP is the duplication target. The number of days reflects the retention period of the backup copy on each cluster.
- If the backup policies and protection plans are using the template SLPs from *cluster1* to *cluster2*, they are updated to use the corresponding template SLP in the reverse direction if cluster1 is faulted and vice versa.

4. All of these changes are applied only to newly created NetBackup jobs. Existing, running, and queued backup jobs continue to use the previous backup and duplication targets.
5. The behavior of the automatic **Update SLP configuration** option is to make sure that backup jobs continue to function after a cluster fault. The duplication is reversed to let the backup images duplicated to the remote cluster when the cluster fault is restored.
6. All the changes to SLPs, backup policies, and protection plans are reverted back to their original configuration when the cluster fault is restored and the original replication primary automatically gets converted as the replication secondary.

DNS key management

Before changing the replication roles, the DNS server(s) configured with NetBackup Flex Scale may need to be updated to correctly resolve NetBackup primary server FQDN(s). NetBackup Flex Scale gives you the option to update the DNS server(s) automatically when changing the replication role of the cluster.

The **DNS key management** wizard in the **Primary service replication** page can be used to configure TSIG (Transaction signature) key for DNS servers configured on the NetBackup Flex Scale cluster. TSIG as specified in RFC 2845 is a shared key message authentication mechanism that is available in BIND DNS. A TSIG key provides the means to authenticate and verify the validity of exchanged DNS data. It uses a shared secret key between a resolver and either one or two servers to provide security. For TSIG authentication to work correctly, the clock has to be in sync between the NetBackup Flex Scale cluster(s) and the DNS server(s).

NetBackup Flex Scale 3.5.100 supports Bind 9 compatible DNS servers for automatic updates. Windows DNS server is not supported for automatic DNS update. If the primary service replication is configured to use same virtual IP for NetBackup primary

service on both clusters, then DNS update is not required before changing the replication role. So the **DNS key management** wizard and automatic DNS update option are not available in this configuration.

You can manage the DNS keys using the GUI.

To manage DNS secret key using GUI

1. Log on to NetBackup Flex Scale GUI of the primary or secondary site.
2. Go to **Settings > Primary service replication**. Click **View details** under **DNS key management**.

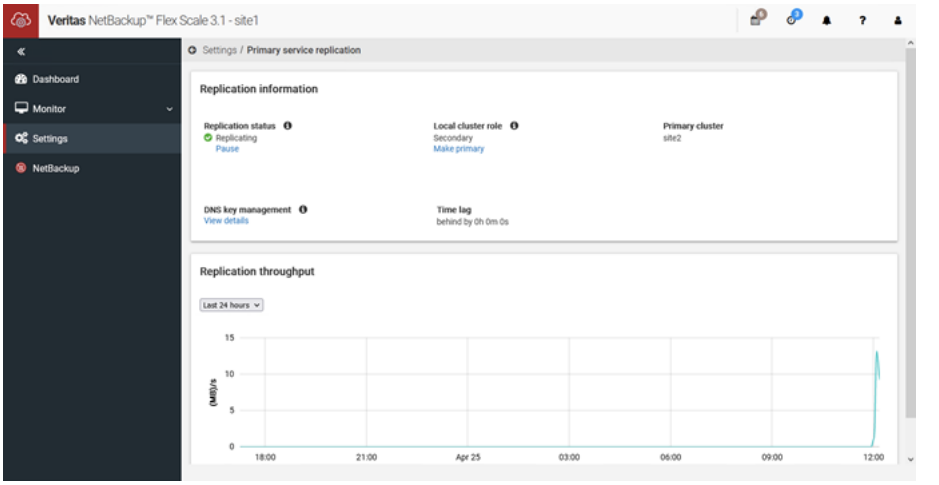

3 The **DNS key management** screen appears. It lists all the configured DNS server that you can manage.

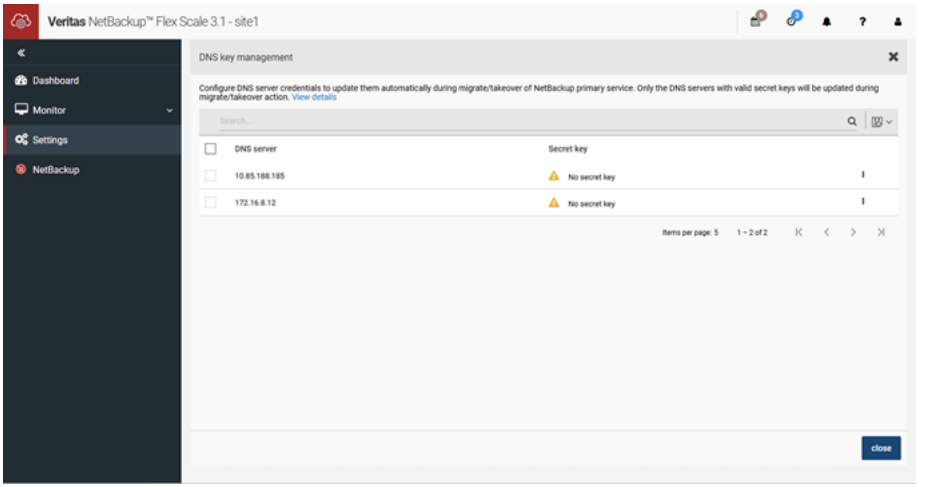

Note: DNS key management tab does not appear if NetBackup primary service configuration has been done using a single virtual IP.

You can also upload the secret key using GUI.

To upload secret key using GUI

1 Ensure that the DNS server secret key file (generated using tsig-keygen on the DNS server) is ready for upload.

Only DNS servers with valid secret keys are allowed to be updated automatically during role change.

2 Click the **Actions** menu (vertical ellipsis) from the right side of the row in the GUI to open the additional options for a selected DNS server.

3 Select **Upload secret key** option.

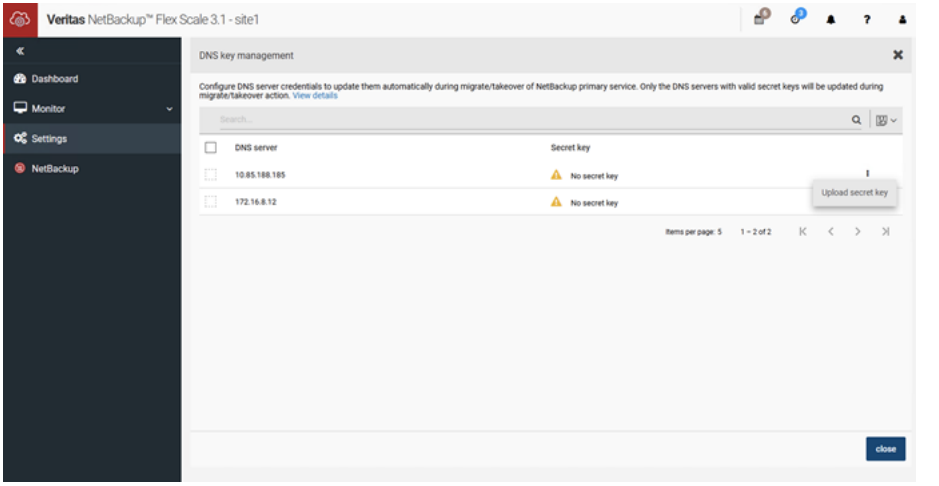

4 Using the **File select** dialog box, select the secret key file to upload. The secret key file is uploaded and then validated against the DNS server. Once the key is validated, the screen shows the key as valid or invalid.

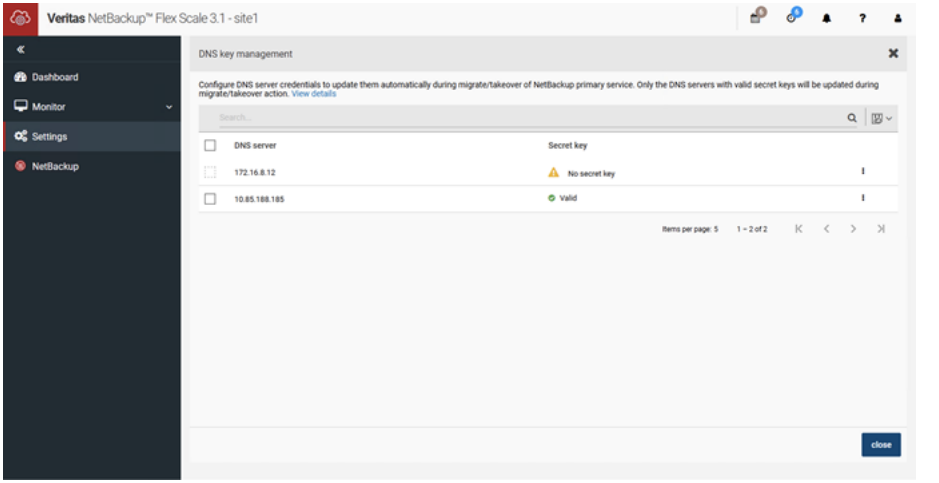

You can also remove the secret key using GUI.

To remove secret key using GUI

1. Click the **Actions** menu (vertical ellipsis) from the right side of the row in the GUI to open the additional options for a selected DNS server.
2. Select **Remove secret key** option.

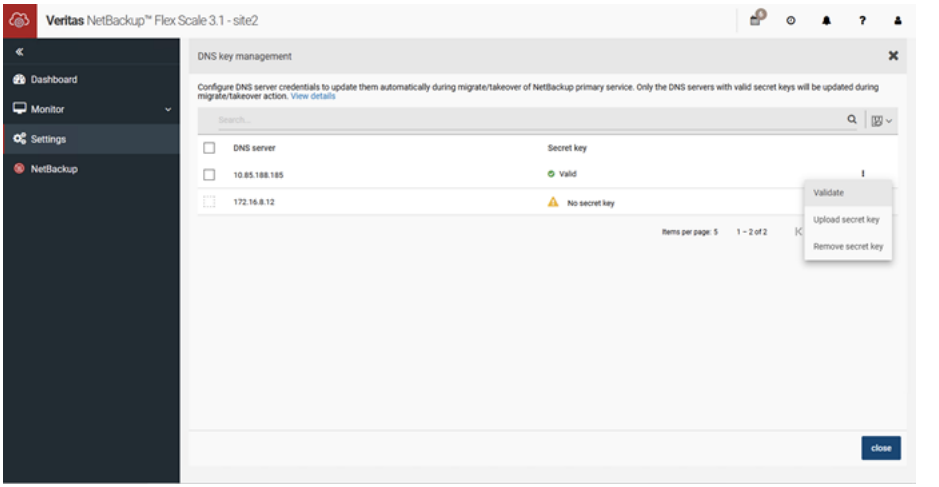

You can also manage the DNS keys using RESTful APIs.

- To retrieve the list of DNS servers to configure automatic DNS update.

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

- To upload a secret key file for a given DNS server.

  ```
  PUT api/appliance/v1.0/disaster-recovery/dns/{serverIP}
  ```

- To configure a DNS server for DNS automatic update.

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

- To validate a secret key file for a given DNS server.

  ```
  POST /api/appliance/v1.0/disaster-recovery/dns/{serverIP}/validate
  ```

- To unconfigure a DNS server from DNS automatic update.

  ```
  DELETE /api/appliance/v1.0/disaster-recovery/dns/{serverIP}
  ```

Managing disaster recovery using GUI

You can manage disaster recovery using the NetBackup Flex Scale GUI. You can switch the role of the primary and secondary site using the GUI. It is also possible to pause the replication and resume it later.

You can perform a role change to reverse the roles of the primary and secondary sites when both the sites are online.

To perform a role change when both the sites are online

1. Log on to the NetBackup Flex Scale GUI of the secondary site.
2. Go to **Settings** > **Primary service replication**. Click **Make primary**.

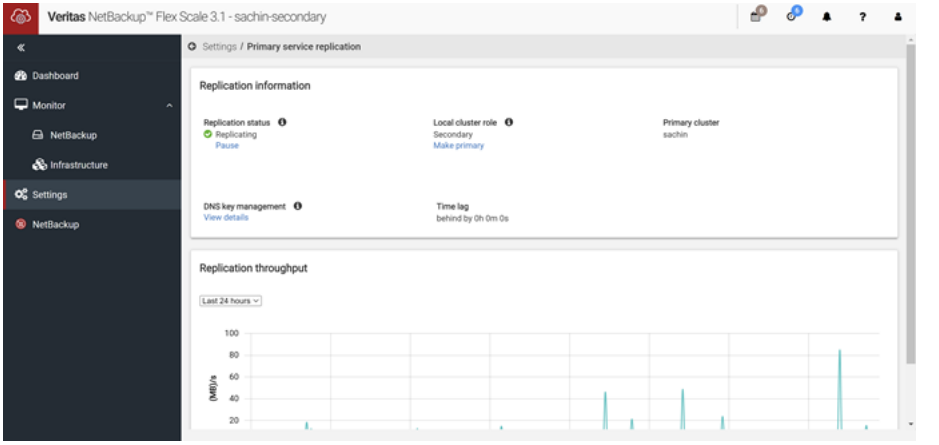

3 The **Make primary** window appears.

You can perform a role change and make the secondary site as the primary site. Select the **Update DNS server** check box if you want to automatically update the NetBackup primary service host name mapping on the DNS servers that are configured with valid credentials. If you have not configured the DNS servers to be updated automatically then ensure that you have updated the DNS entries of the NetBackup primary service FQDNs to resolve to the secondary site.

Click **Confirm**.

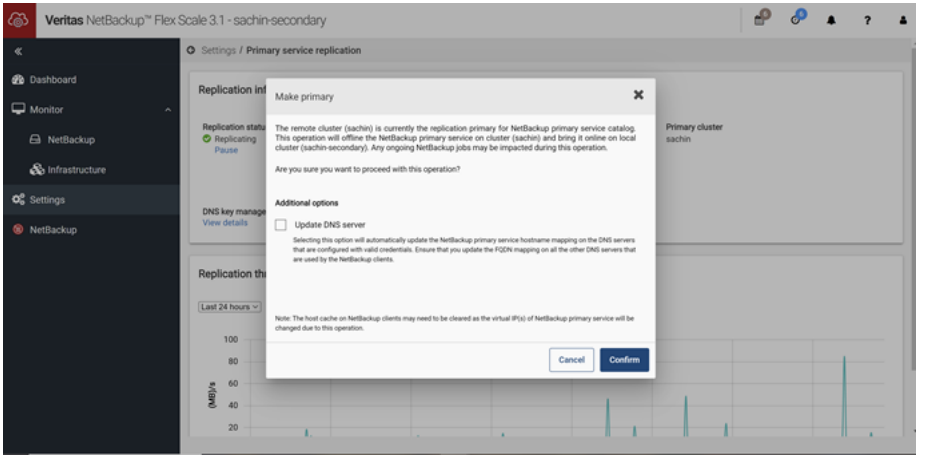

Note: In case of non-DNS sites (both primary and secondary), updating the DNS entry for the NetBackup primary service IP is not applicable.

In case NetBackup primary service configuration is done using single virtual IP, updating DNS server is not applicable.

You can perform a role change when the primary cluster is down.

To perform a role change when the primary cluster is down

1. Log on to the NetBackup Flex Scale GUI of the secondary site.
2. Go to **Settings** > **Primary service replication**. Click **Make primary**.

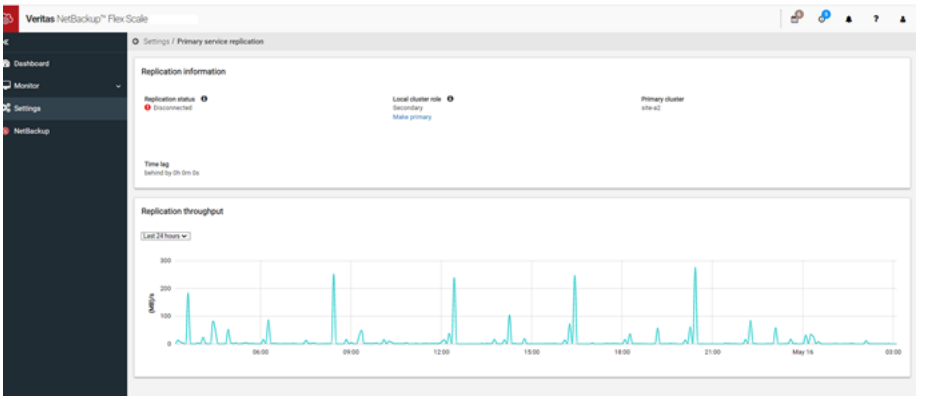

3 The **Make primary** window appears.

You can perform a role change and bring up the NetBackup primary service on the local cluster. Select the **Update DNS server** check box if you want to automatically update the NetBackup primary service host name mapping on the DNS servers that are configured with valid credentials. If you have not configured the DNS servers to be updated automatically then ensure that you have updated the DNS entries of the NetBackup primary service FQDNs to resolve to the secondary site.

Select the **Update SLP configuration** check box if you want to update NetBackup policies and SLPs after this operation. Read the note provided below this check box for more details.

See "Automated NetBackup SLP management" on page 195.

Click **Confirm**.

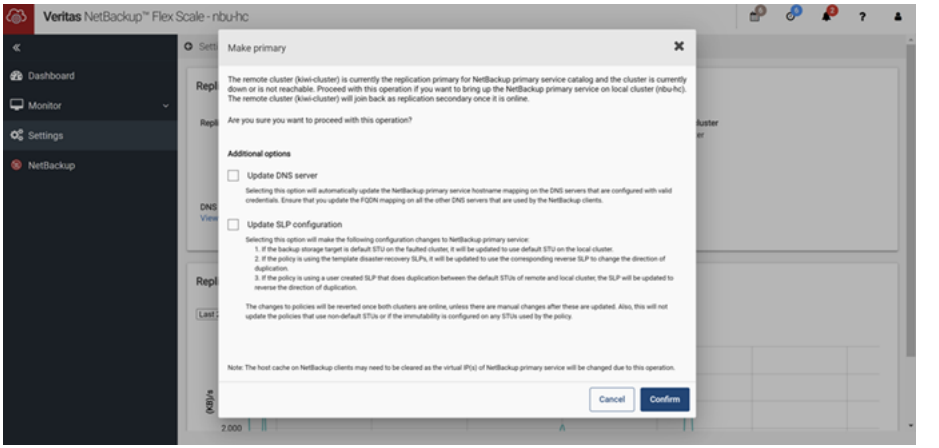

Note: In case of non-DNS sites (both primary and secondary), updating the DNS entry for the NetBackup primary service IP is not applicable.

In case NetBackup primary service configuration is done using single virtual IP, updating DNS server is not applicable.

You can perform a role change when the replication is in paused, resync or error state and both sites are online.

To perform a role change when the replication is in paused, resync or in error state

1. Log on to the NetBackup Flex Scale GUI of the secondary site.
2. Go to **Settings** > **Primary service replication**. Click **Make primary**.

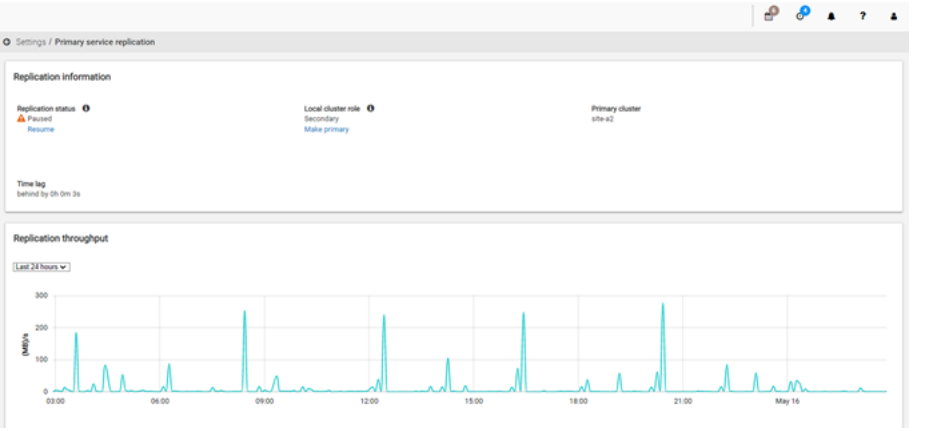

3 The **Make primary** window appears.

You can perform a role change and bring up the NetBackup primary service on the local cluster. Select the **Update DNS server** check box if you want to automatically update the NetBackup primary service host name mapping on the DNS servers that are configured with valid credentials. If you have not configured the DNS servers to be updated automatically then ensure that you have updated the DNS entries of the NetBackup primary service FQDNs to resolve to the secondary site.

Click **Confirm**.

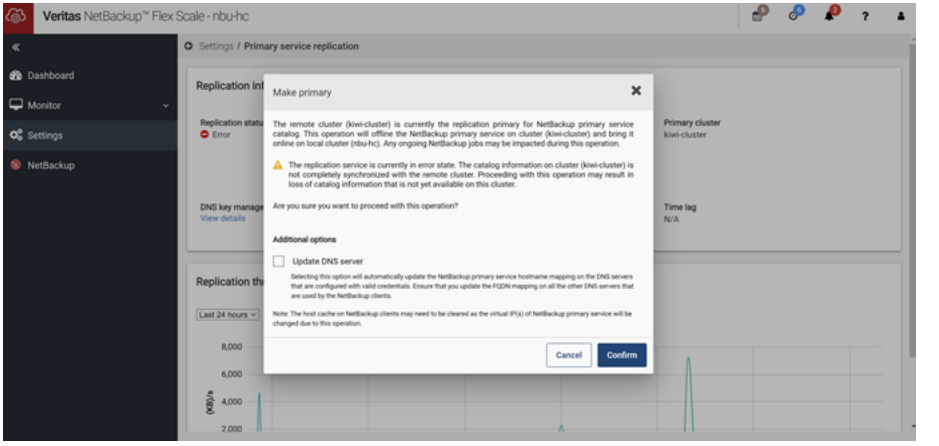

You can pause an ongoing replication operation.

To pause replication

1 Go to **Settings** > **Primary service replication** 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** link gets enabled on both the clusters.

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

To resume replication

1 Go to **Settings** > **Primary service replication** 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** link 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 209.
- Configure the primary and secondary site for disaster recovery.
  See "Configuring disaster recovery" on page 211.
- Configure NetBackup optimized duplication SLPs.
  See "Configuring a Storage Lifecycle Policy for optimized duplication" on page 398.

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-webservice` 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-webservice) 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-webservice-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-webservice` 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 Best Practices and Troubleshooting 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 398.

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 194.

   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 194.

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 406.

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 409.

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.

Chapter 9

NetBackup Flex Scale security

This chapter includes the following topics:

About the security meter

NetBackup Flex Scale includes a security meter to view and configure the appliance security settings from one location. The security meter provides security insights and recommendations to improve the appliance security. The security meter keeps a track of the security settings and shows you a list of available security features with quick links to configure them. The security meter displays the current security index from good to excellent based on how many security features are turned on. Built-in security features are turned on and shown as **Enabled** in the security meter.

The security meter can be found on the appliance dashboard. Only a user with an Appliance administrator role can view and manage the settings from the security meter.

The following figure shows the security meter that is displayed on the appliance dashboard:

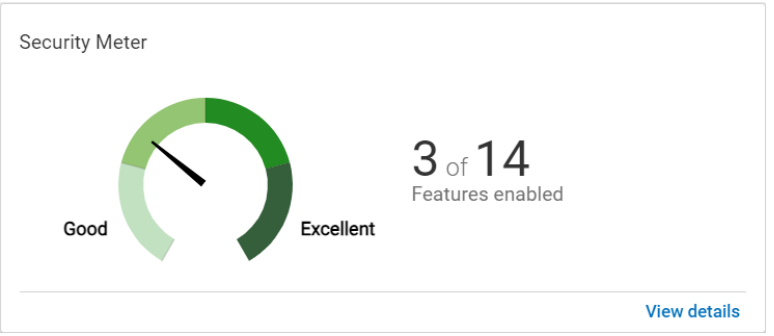

The following figure shows all the appliance security settings that you can track and manage:

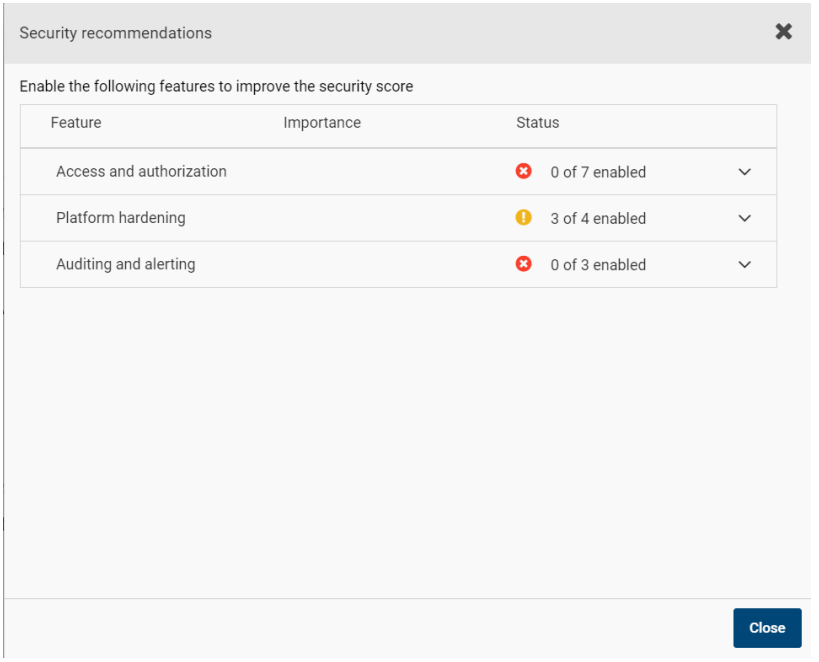

STIG overview for NetBackup Flex Scale

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 8 Security Technical Implementation Guide - Version 1, Release 10`

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.

NetBackup Flex Scale also supports DISA's Application Security and Development STIG Version 5, Release 3.

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 3.5.100, 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 8 Security Technical Implementation Guide - Version 1, Release 10`

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 219.
- 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, the FQDN, or the short host name 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 Management > Cluster settings**.
   - 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Settings**.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

2 Click **Security management**.

3 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.

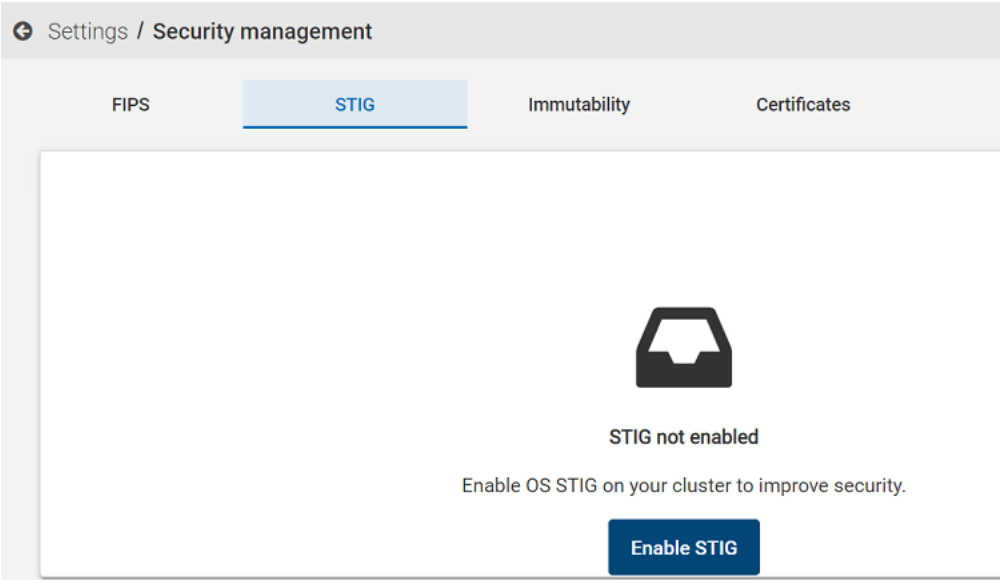

4 To monitor the progress, click **View details** on the **Security** page. The ongoing and completed tasks for the operation are also displayed in **Recent activity**.

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, the FQDN, or the short host name 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.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

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.

Viewing the status using the REST APIs

You can find the RESTful APIs at `https://ManagementServerIPorFQDN:14161/swagger/infra/v1.0/` where *ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

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, the FQDN, or the short host name 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 Management > Cluster settings > Security management**.
   - 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Settings > Security management**.

     Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

2. 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 3.5.100, 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.

Starting with NetBackup Flex Scale version 3.2, the application layer is FIPS-compliant.

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.

Viewing the status using the REST APIs

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.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

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, the FQDN, or the short host name 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 Management > Cluster settings > Security management**.
   - 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Settings > Security management**.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

2 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**.

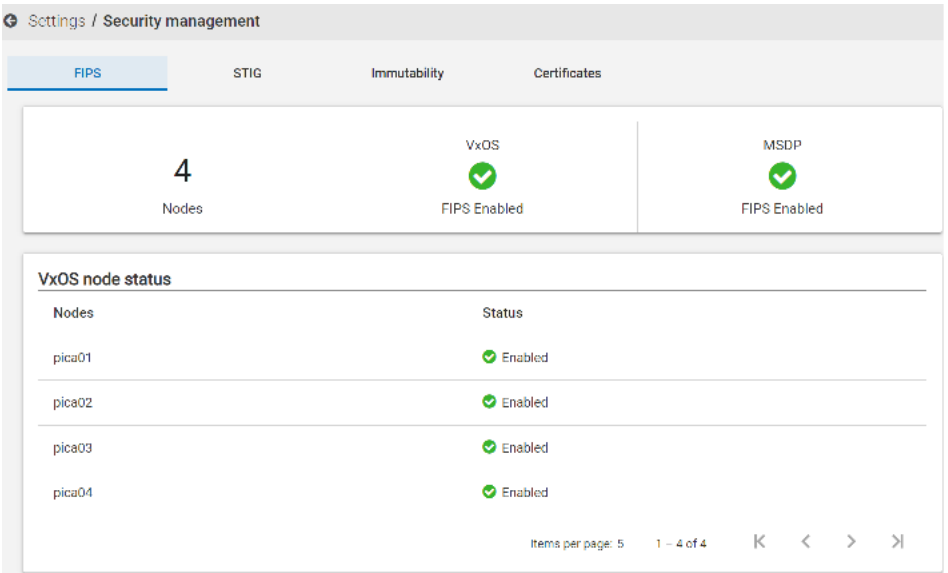

Managing the login banner

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.

Note: Ensure that you change the banner from the NetBackup Flex Scale infrastructure management UI or the NetBackup Flex Scale web UI. If you change the banner in the NetBackup UI, the changes are not reflected in NetBackup Flex Scale.

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, the FQDN, or the short host name 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 Management > Cluster settings > User management**.
   - 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Settings > User management**.

     Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

2. Click **Manage sign-in banner**.

   If a banner is already set, the Manage sign-in banner page displays the current banner.

3. On the Manage sign-in banner page, click **Edit**.

4. (Optional) Under **Sign-in banner heading**, enter the banner heading. The heading can be a maximum of 250 characters.

5 Under **Sign in banner text**, enter the text for the banner message. The message can be a maximum of 4000 characters.

6 To review the changes, click **Preview**.

7 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.

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

Note: Ensure that you change the password policy from the NetBackup Flex Scale infrastructure management UI or the NetBackup Flex Scale web UI. If you change the password policy in the NetBackup UI, the changes are not reflected in NetBackup Flex Scale.

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, the FQDN, or the short host name 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 Management > Cluster settings > User management**.

- 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, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server and API gateway during the cluster configuration, and then in the left pane click **Settings > User management**.

Note: If you access the NetBackup Flex Scale infrastructure management console by using the short host name from a node, set the DNS settings (name server, domain name, and search domain) or ensure that the entry for mapping the short host name to an IP address exists in the hosts file of the node.

---

2 Click **Manage password policy**.

3 On the Manage password policy page, click **Edit**.

4 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.

5 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

Table 9-1 *(continued)*

Parameter	Description
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
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

- Restricted access to Remote Management Platform (HPE 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.
  - User has the option to remove the retention locks and expire image data.
  - 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/openv/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 mode is 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).
- If you use the NetBackup Flex Scale UI to change the retention period without changing the lockdown mode, you have to manually update the disk volume in NetBackup, either through the NetBackup CLI or the NetBackup Web UI. This is necessary to synchronize the information in the NetBackup database.

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 `show cluster-id` command to get the cluster ID for the entire cluster or the `show serial-number` command to get the serial number for a specific node.

   You can also get the node serial number from the NetBackup Flex Scale web UI by navigating to **Infrastructure > Nodes > Node serial number**. You can get the cluster ID from the NetBackup Flex Scale web UI by navigating to **Infrastructure > Cluster Id**.

3. If you have permission to generate access key on the SHI portal, then go to the portal to generate access key with the cluster ID or serial number that you got in step 2. For more information, refer to the *System Health Insights User Guide*.

   If you do not have the permission to generate the access key on the SHI portal, open a ticket with Veritas Support to generate an access key. Set an access passphrase which will be required later 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 the maintenance password. Enter the maintenance password and press **Enter**. You are prompted to enter an access key. Enter the access key that you got in step 3. You are prompted to enter an access passphrase. Enter the access passphrase that you set 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. Enter the maintenance password. You are prompted to enter an access passphrase. Enter the access passphrase set in step 3. Press **Enter** to get into the root shell.

7. Repeat steps 3 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.

9. Login to the management console IP and run the `system lock enable` to lock all the nodes in the cluster.

Note: If an access key is generated with cluster ID, then the access key can be used to unlock all the nodes in the same cluster. If an access key is generated with serial number, then the access key can only be used to unlock the appliance which have the serial number.

The access key expires in 2 hours.

Restricted access to Remote Management Platform (HPe iLO)

If you select enterprise or compliance mode, you can restrict remote management access to the node by selecting the **Restrict remote management access** check box. This option is not available for normal lockdown mode. Restricting remote management access to nodes provides an additional level of data security and limits the privileges and operations that you can perform.

After you enable this restriction, an IPMI Administrator user on an HPE platform has only **Login** and **Virtual Power and Reset** privileges. With these privileges, the user can only view settings in iLO and perform power-related operations.

After you enable restricted remote access, remember that:

- In enterprise lockdown mode, you can enable or disable restricted remote management access.
- In compliance lockdown mode, you can only enable restricted remote management access, but cannot disable the remote management access restriction.
- You can also choose to enable or disable restricted remote management access after the initial configuration is complete.

Warning: Once you enable restricted remote management access, all destructive operations are disabled for all the IPMI users. Users can view and perform limited operations in the IPMI web GUI but cannot access the remote console. Physical access to the system is required to logon to the console.

Table 9-2 lists the privileges given for a local account in iLO.

Table 9-2 HPE iLO

Privileges	Description
Login	Enables a user to log on to iLO.
Remote Console	Enables a user to access the host system remote console, including video, keyboard, and mouse control. Users with this privilege can access the BIOS, and therefore may be able to perform host-based BIOS, iLO, storage, and network tasks.
User Config	Enables a user to add, edit, and delete local iLO user accounts. A user with this privilege can change privileges for all users. If you are not assigned this privilege, you can view your own settings and change your own password.

Table 9-2 HPE iLO *(continued)*

Privileges	Description
iLO Config	Enables a user to configure most iLO settings, including security settings, and to update the iLO firmware. This privilege does not enable local user account administration. After iLO is configured, revoking this privilege from all users prevents reconfiguration from the following interfaces: ■ iLO web interface ■ iLO RESTful API ■ CLI ■ HPQLOCFG Users who have access to the following interfaces can still reconfigure iLO: ■ UEFI System Utilities ■ HPONCFG Only a user who has the Administer User Accounts privilege can enable or disable this privilege.
Virtual Media	Enables a user to use the virtual media feature on the host system.
Virtual Power and Reset	Enables a user to power-cycle or reset the host system. These activities interrupt the system availability. A user with this privilege can diagnose the system by using the **Generate NMI to System** button.
Host NIC Config	Enables a user to configure the host NIC settings. This privilege does not affect configuration through host-based utilities.
Host Bios Config	Allows configuration of the host BIOS settings by using the UEFI System Utilities. This privilege is required for replacing the active system ROM with the redundant system ROM. This privilege does not affect configuration through host-based utilities.
Host Storage Config	Enables a user to configure the host storage settings. This privilege does not affect configuration through host-based utilities.

Table 9-2 HPE iLO *(continued)*

Privileges	Description
Recovery Set	Enables a user to manage the System Recovery Set. By default, the Recovery Set privilege is assigned to the default administrator account. This privilege can be added to a user account only by creating or editing the account with an account that already has this privilege. If there is no user account with the Recovery Set privilege, and an account with this privilege is required, reset the management processor to the factory default settings. The factory default reset creates a default Administrator account with the Recovery Set privilege. This privilege is not available when iLO security is disabled with the system maintenance switch. For information about the default account credentials and how to configure this privilege without access to an account that has this privilege, see the *iLO User Guide*.

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**.

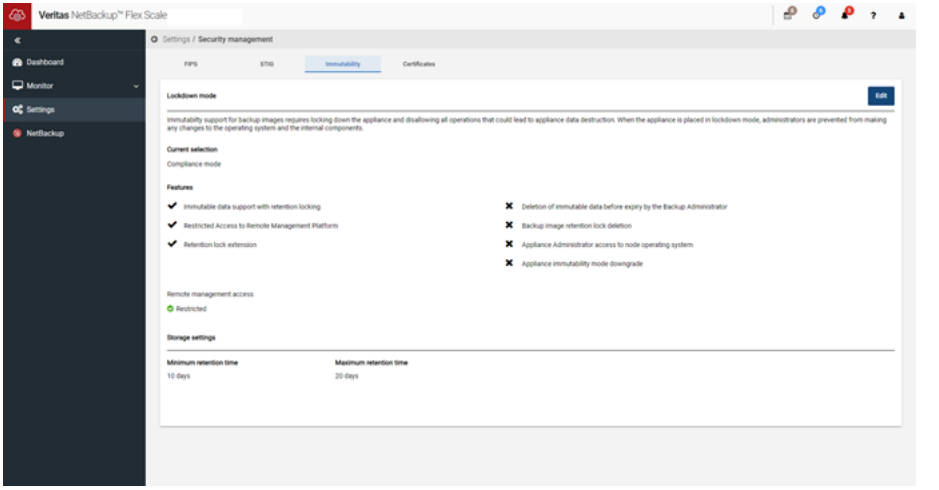

2 Choose any of the lockdown modes under **Mode Selection**. You can choose from Normal, Enterprise, and Compliance. The supported features for each mode are listed.

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.

If you select enterprise or compliance mode, you can restrict remote management access to the node by selecting the **Restrict remote management access** check box.

Warning: Once you enable restricted remote management access, all destructive operations are disabled for all the IPMI users. Users can view and perform limited operations in the IPMI web GUI but cannot access the remote console. Physical access to the system is required to logon to the console.

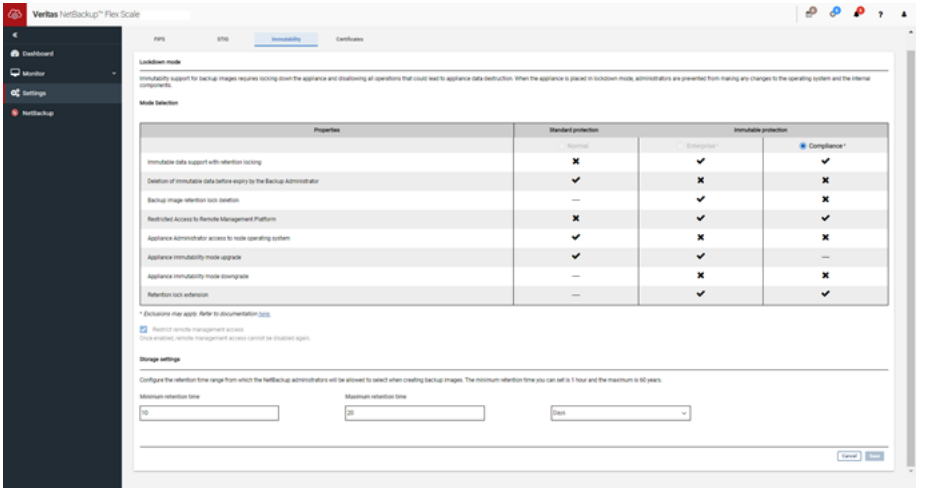

Authenticating users using digital certificates or smart cards

You can configure NetBackup Flex Scale to authenticate users with a smart card or a digital certificate. After configuration, the users can use the **Sign in with**

certificate or smart card option to sign in to NetBackup Flex Scale UI using smart cards or digital certificates.

Before you configure user authentication using smart cards or digital certificates, note the following:

- Digital certificate or smart card authentication can be configured for LDAP, AD, and local users.
- Ensure that LDAP is configured if you want to authenticate LDAP users by using digital certificate or smart card.
- Ensure that AD is configured if you want to authenticate AD users by using digital or smart.
- Ensure that you create a local user if you want to authenticate local users by using digital or smart card.
- Smart card authentication requires a list of trusted root or intermediate CA certificates. You must add the CA certificates that are associated with the user digital certificates or the user smart cards.

To authenticate users with a certificate or smart card for media server only deployment:

1. Use a user account with an Appliance Administrator role to log in to the NetBackup Flex Scale infrastructure management UI `https://consoleIP:14161` where *consoleIP* is the public IP address that you specified for the infrastructure management UI during the cluster configuration.
2. In the left pane click **Settings > Security > Smart card authentication**.
3. Use the slider to turn on smart card authentication.
4. In the **Configure smart card authentication** dialog box, specify the following options:
   - In the user authentication domain list, specify the following information:
     - None, which is a default option, indicates that only local users can be authenticated using smart card.
     - For an AD user, select the configured AD server.
     - For an LDAP user, select the configured LDAP server.
   - Under **Certificate mapping attribute**, to specify the user using the username format, click **Common name**. To specify the user using the username and domain format (for example, username@test.com) click **User principal name**.
   - Optionally, enter the Online Certificate Status Protocol (OCSP) URI. OSCP is used for checking the validity of the certificate. The OCSP responder is

a remote independent entity (certificate vendor authority). If you do not provide the OCSP URI, the URI in the user certificate is used.

- Click **Save**.

5 To the right of CA certificates click **Add**.

You can upload a CA certificate or a chain certificate. The leaf certificate can be created directly from root certificate or from an intermediate certificate. Chain certificate is a concatenation of root and intermediate certificate.

6 Click **Browse** to select the CA certificate or drag and drop the CA certificate and click **Add**.

Certificates must be in PEM format, with certificate file type as `.pem`. Only one certificate can be added at a time. The web server is restarted after you add the certificate and the certificate is added to the web server trust store `/shared/cluster_certs/cac/`. The selected CA certificate is displayed under CA certificates.

7 Upload the client certificate to the browser's certificate store. See the browser documentation for importing client certificates.

8 Add Appliance administrator user role to a smart card user. To add the Appliance administrator role to an AD, LDAP, or a local user, navigate to **Settings > User management**.

9 To log on using the smart card, when you enter the URL for the UI, you are prompted to select the certificate that you added to the browser trust store. Select the certificate to authenticate. Selecting the certificate is a one-time activity. You can now use the **Sign in with certificate or smart card** option to sign in to the UI.

To authenticate users with a certificate or smart card for a cluster where both the primary server and media servers are deployed:

1 Use a user account with both Appliance Administrator and NetBackup Administrator role to log in to the NetBackup Flex Scale web interface `https://`*`ManagementServerIPorFQDN`*`/webui` where *ManagementServerIPorFQDN* is the public IP address, the FQDN, or the short host name that you specified for the NetBackup Flex Scale management server during the cluster configuration.

2 At the top right, click **Settings > Smart card authentication** and follow the steps mentioned in the "Configuring authentication options" section of the *NetBackup™ Web UI Administrator's Guide*. After configuring smart card authentication, you need to perform additional steps before you can log in using a smart card.

3 Restart the `nbwmc` service.

- Perform an SSH to the primary server.
- Run the `nbwmc -terminate` command.
- Run the `nbwmc start` command. This may take a few minutes.
- Check if the 13731 port is open.

  ```
  # netstat -tnlp | grep 13731
  tcp 0 0 0.0.0.0:13731 0.0.0.0:* LISTEN 35895/java
  ```

  If the port is open, it means that the `nbwmc` service has been restarted successfully.

4 Get the NetBackup Flex Scale web UI's root CA and intermediate certificates and upload it to browser:

- Navigate to **Cluster Management > Cluster settings > Security > Certificates** and click **Download root certificate**.
- The downloaded certificate contains two certificate keys in a single file. Separate the downloaded certificate in two files:
  - `root_ca.pem`: Upload to the browser's trusted root certificate store.
  - `stem_ca.pem`: Upload to the browser's intermediate certificate store.

5 Get NetBackup web root CA certificates and upload it to the browser. To get the NetBackup web certificate:

- Get the NetBackup web root CA certificate using swagger or using CURL API:

  ```
  curl -X 'GET'
  \'https://primary-server-FQDN/netbackup/security/cacert' \ -H
   'accept: application/vnd.netbackup+json;version=9.0'
  ```

- Copy the web root certificate from the received response to a file. Ensure that you replace the \n character with newline.
- Upload the web root CA certificate to update the SAN entries in the NetBackup web certificate.
- If you use Mozilla Firefox browser, enable **network.cors_preflight.allow_client_cert** to set it to **true**.

6 Log in to the NetBackup Flex Scale UI by clicking **Sign in with certificate or smart card** on login screen and when prompted select the certificate that you uploaded to the browser trust store.

About system certificates on NetBackup Flex Scale

NetBackup Flex Scale supports one certificate for all services. The certificate can be internal or external. The Appliance CA creates the internal certificate. Any CA can create the external certificate using a CSR generated using the NetBackup Flex Scale GUI. The CA certificate can also be downloaded using the GUI by navigating to **Settings > Security Management > Certificates**. The admin user can change the certificate mode by navigating to **Settings > Security Management > Certificates** in the GUI. After changing the certificate mode, the admin user should restart all the services to start using the new certificate.

You can navigate to **Settings > Security > Certificates > Download root certificates** to download the VxOS certificates which contain the intermediate and root CA certificates (both VxOS root and VxOS stem).

Note: Ensure that you have the latest Chrome update for the secure connection to work with VxOS certificates.

If the system certificate mode is set to internal, the certificate that is signed by the Appliance CA (internal certificate) is used for all the GUI services.

You must update the clients trust-store with CA certificate to secure connection.

Deploying external certificates on NetBackup Flex Scale

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.

- You do not get any alert for NBCA expiry or renewal. An event is raised when NBCA is about to expire and renewed in the background.
- NBCA is auto renewed 60 days before expiration.
  If NBCA renewal fails, a failed task can be seen on NetBackup Flex Scale GUI.
- 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.
- 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.
- If you want to deploy ECA on a cluster on which disaster recovery is already configured, ensure that you configure ECA on the primary cluster.
- If ECA is deployed on the primary cluster before adding the secondary cluster, then you must redeploy ECA from the primary cluster after disaster recovery configuration is complete. This is to ensure proper connectivity between the primary server, media server, and storage services.
- If CRL mode is selected as CRL URL during ECA deployment, ensure that the CRL URL host name is resolvable by the existing DNS servers. If there are no DNS servers or if the DNS server cannot resolve the CRL URL host name, you must add the CRL URL as a custom host entry for the NetBackup container and the cluster nodes. This is also applicable if a DNS server is present during ECA deployment but is removed later.
- If you do not want to generate CSR from the GUI, then you can use your own certificate for ECA deployment. In such a scenario, you must upload your own unencrypted private key.
- If ECA is configured with the CRL as an URL, and if the CRL server becomes unreachable or unavailable for more than 24 hours for any reason, the NetBackup services on the NetBackup Flex Scale cluster appears as degraded. Once the

connectivity to the CRL server is established again, the NetBackup services appear as healthy.

Considerations while deploying ECA on a cluster on which only media server is deployed:

There are some additional considerations that you need to keep in mind when you deploy ECA on a media server only cluster.

- If you have deployed media server only clusters with external NetBackup primary server on BYO:

  If ECA deployment is done after media server only configuration:

  - The primary server should be configured in ECA + NBCA mode before starting ECA deployment on the cluster.
  - The CA chain (Root + Intermediate) used should be same trusted certificate chain for both primary and media server only cluster.

  If media server only deployment is done after ECA configuration on NetBackup BYO:

  - Pure ECA mode is not supported.
  - If the primary server is deployed in NBCA + ECA mode then media server can be deployed using it and ECA can be configured on media server only cluster.
  - The CA chain (Root + Intermediate) used should be same trusted certificate chain for both primary and media server only cluster.

- If you have deployed media server only cluster with external NetBackup primary server in a NetBackup Flex Scale cluster:

  If ECA deployment is done after media server only configuration:

  - Primary server should be configured in ECA + NBCA mode before starting ECA deployment on the media server only cluster.
  - This can be done using the NetBackup Flex Scale ECA deployment workflow.
  - The CA chain (Root + Intermediate) used should be same trusted certificate chain for the cluster on which both primary and media servers are deployed and media server only cluster.

  If media server only deployment is done after ECA configuration on a NetBackup Flex Scale cluster on which both primary and media server are deployed

  - Pure ECA mode is not supported a NetBackup Flex Scale cluster on which both primary and media server are deployed.

- If the cluster is deployed in NBCA +ECA mode, then media server only cluster can be deployed using it and ECA can be configured on media server only cluster.
- The CA chain (Root + Intermediate) used should be same trusted certificate chain for the cluster on which both primary and media servers are deployed and media server only cluster.

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.

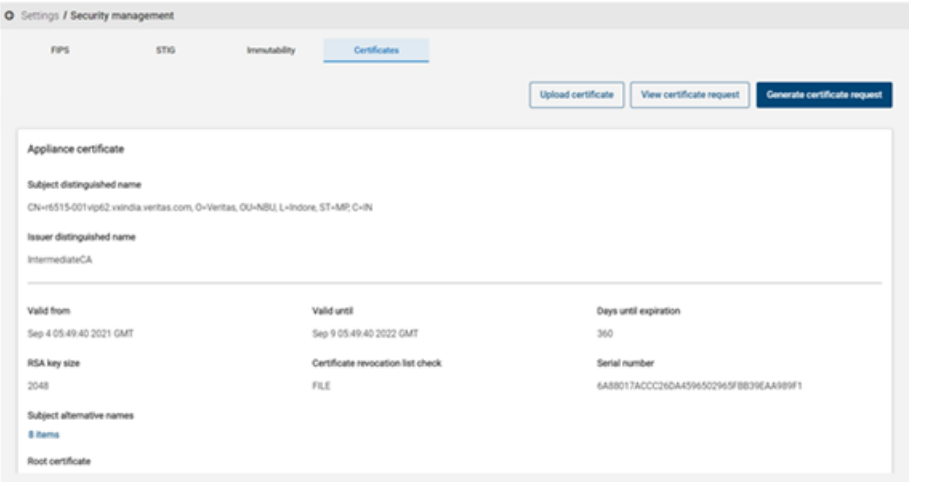

3 Click **Generate** to generate a certificate request.

The CSR is used to generate certificate. This certificate along with the CA bundle (Root CA + Intermediate CA) should be uploaded on the upload page.

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.

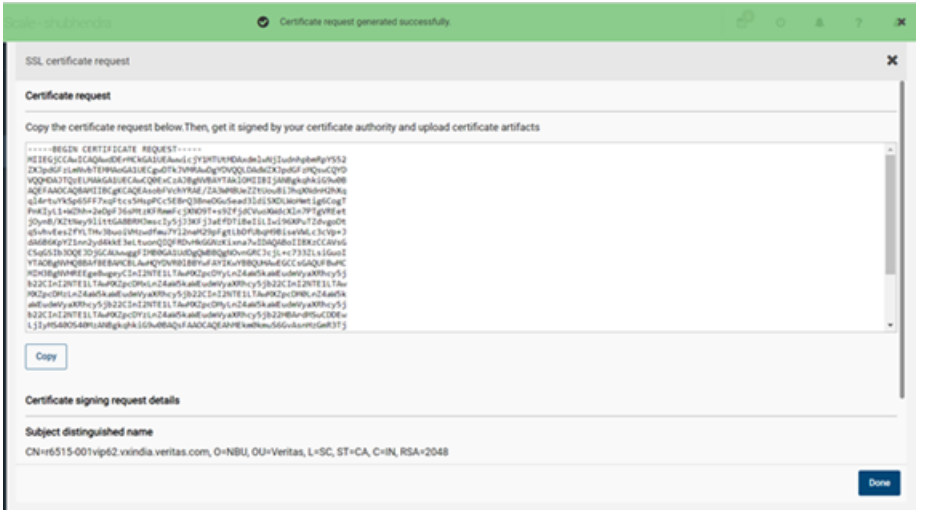

5 Go to **Certificates** > **Upload certificate** to upload the certificate artifacts. For a fresh deployment, both the certificate and CA certificate bundle are mandatory. CRL is optional. For renewal, any one of the three are required.

Both the certificate and CA bundle must be a .PEM file. The CA bundle .PEM file should contain entries for all the CA certificates from root to intermediate till the immediate parent of the leaf certificate.

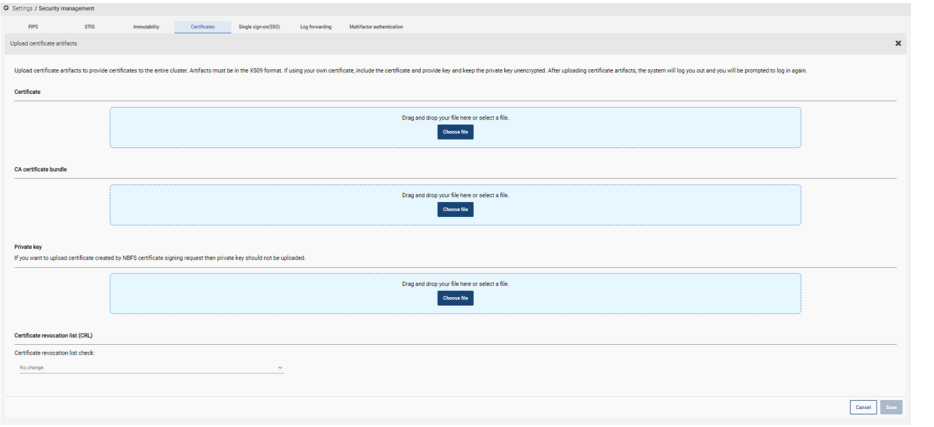

6 After the certificate is uploaded, **Save** gets enabled. Click **Save**. The deployment is initiated.

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/security_artifacts/`.

All the logs pertaining to ECA deployment are present in the `/log/VRTSnas` 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 server FQDN as an additional SAN entry.
- Do not add the media server FQDN in the 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.
- Do not add the media server FQDN in the SAN entry.

If ECA is deployed and you want to add a new data network:

- Add the new primary and storage server FQDNs to the CSR and deploy the new certificate before adding the new data network.
- Do not add the media server FQDN in the SAN entry.

Configuring isolated recovery environment (IRE)

Organizations can minimize ransom demands by using encryption and creating a stringent security perimeter. In addition, they need to isolate, analyze, and preserve a copy of data to ensure business continuity. IRE enables organizations to meet these needs and satisfy strict regulatory and retention requirements. Veritas customers can easily deploy an IRE using their existing Veritas NetBackup infrastructure as part of a multi-layered resiliency strategy.

NetBackup Flex Scale uses the Pull model to pull the replication request from the IRE domain through a specific window as defined in the IRE air gap schedule. By initiating a data transfer request from inside the IRE domain, there is better control over data flow to secure the environment further both logically and physically. You can determine the Service Lifecycle Policy (SLP) windows and configure the air-gapped schedule for maximum protection.

The NetBackup Flex Scale IRE solution optimizes data movement whereby the request to send data comes from the IRE side, the MSDP reverse connection. You can deploy another the tertiary copy of the backup images behind a firewall to an isolated environment without opening any inbound firewall ports to NetBackup. This keeps the environment secure, allowing a sandbox approach to perform malware scans or test recovery procedures before recovering at a larger scale. You can optionally add a physical air gap as an additional layer of protection. By empowering the destination environment to request the data from the source environment (by invitation only), it is possible to support 24x7 data movement while isolating the stored data from any potential threats.

The IRE solution also supports multiple configurations. Hence, you can have a single IRE domain for multiple production domains. Another key feature of this solution is that the IRE domain is not required to have the same configuration as your production domain. You can configure an IRE domain as per your requirements and use it to securely transfer backups from production to IRE.

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

- NetBackup Flex scale Appliance: 3.2 or later
- Storage server: 19.0.1 or later
- NetBackup: 10.3.0.1 or later

For more information, refer to the NetBackup Deduplication Guide on SORT.

Chapter 10

Configuring multifactor authentication

This chapter includes the following topics:

About multifactor authentication

Multifactor authentication is a robust security measure widely used for adding an additional layer of security to the authentication process by requiring users to provide a unique, time-limited code along with their regular login credentials. It is a multiple-step account login process that requires you to enter a 6-digit one-time password along with your password.

It is strongly recommended that you configure multifactor authentication to protect the security of your account.

See “Configuring multifactor authentication for your user account” on page 256.

If multifactor authentication is enforced in the NetBackup Flex Scale cluster, all users must configure multifactor authentication for their user accounts for successful sign-in.

See "Configuring multifactor authentication for your user account when it is enforced in the cluster" on page 258.

Considerations before configuring multifactor authentication

Some considerations that you need to remember before you configure multifactor authentication:

- The Appliance administrator can see the status of all the users on the **Settings > User management** page.
- If AD/LDAP server configuration is removed from the cluster without removing the AD/LDAP user's MFA configuration, the Appliance administrator may see stale entries for AD/LDAP users.
- If you are an AD/LDAP user with no role, you cannot login to the appliance.
- A local administrator is a non AD/LDAP user.
- If NetBackup Flex Scale has been deployed with both primary and media servers, and if the user does not have the Appliance administrator role and has only NetBackup administrator role, the user is directed to the home screen.
- Local administrator users' roles must be assigned from the NetBackup Flex Scale GUI.
- When catalog replication for disaster recovery is configured between two NetBackup FlexScale clusters, users are managed independently on each cluster and the corresponding multifactor authentication configuration should be done separately on each cluster. Veritas recommends that you use the following guidelines when making user configuration changes in a NetBackup Flex Scale cluster on which disaster recovery is configured:
  - When adding local users, both the clusters should use the same credentials.
  - AD/LDAP configuration must be performed only on the primary cluster on which disaster recovery is configured.
  - When configuring multifactor authentication for a user, the same multifactor authentication secret key must be used for both clusters.
  - When enforcing multifactor authentication, it should be enforced with the same start date on both the clusters.

Configuring multifactor authentication for your user account

You must first install and configure authenticator application on your smart device that provides you with the one-time password.

Supported authenticator applications

If the NetBackup Flex Scale administrator has enforced multifactor authentication in the NetBackup Flex Scale cluster, you must configure it for your user account for successful sign-in. You must configure multifactor authentication before the start date of enforcement. Else, you will lose access to the appliance and your automation workflow (if using login API) will also be impacted.

Even if multifactor authentication is not enforced, it is recommended that you configure it for enhanced security.

To configure multifactor authentication for your user account

1. Sign in to the NetBackup Flex Scale UI.
2. On the top right, click the profile icon and click **Manage multifactor authentication**.
3. On the **Manage multifactor authentication** screen, click **Configure**.
4. On the next screen, follow the given steps.

   Install and configure authenticator application on your smart device. It generates one-time password and sends it on your smart device.
5. Scan the QR code with the authenticator application or enter the key manually.

   The manual key should be base32 encoded and can contain between 26 to 208 characters with or without padding.
6. Enter the one-time password that you see in the authenticator application.
7. Click **Configure**.

   At the time of next sign-in, you need to enter the one-time password along with the username and password.

See “Disabling multifactor authentication for your user account” on page 257.

Disabling multifactor authentication for your user account

You can disable MFA for your user account only If multifactor authentication is not enforced. However, it is strongly recommended that you configure multifactor authentication to protect the security of your account.

If multifactor authentication is enforced, and you want to reset it, See "Resetting multifactor authentication for a user" on page 259.

To disable multifactor authentication for your user account

1. Sign in to the NetBackup Flex Scale UI.
2. If you are an Appliance administrator, click the profile icon on the top right, and select **Manage multifactor authentication**.

   If you are not an Appliance administrator, select **Manage multifactor authentication** in the home screen.
3. If you have already configured multifactor authentication for your user account, you can see the **Disable** button.
4. Click **Disable**.
5. Enter the one-time password and click **Submit**.

Enforcing multifactor authentication for all users

Only the NetBackup Flex Scale administrator can enforce multifactor authentication for all NetBackup Flex Scale users.

Before you enforce multifactor authentication:

- Multifactor authentication can be enforced only if at least two local users have configured it.
- You can set a future start date for enforcement so that the users get sufficient time to configure their multifactor authentication.
- If multifactor authentication is not configured by the start date, the user will not have access to the appliance. If the user's automation workflow uses login API, then that will also be impacted.
- Once multifactor authentication is enforced, it cannot be reversed.
- It is not possible to postpone the start date of enforcement after it is set.
- You can prepone the start date for enforcement using the **Reinforce** button on the **Settings > Security management > Multifactor authentication** page.

- The start date for enforcement cannot be more than 90 days from the current date.

To enforce multifactor authentication for all users

1. Sign in to the NetBackup Flex Scale UI.
2. Go to **Settings > Security management > Multifactor authentication**.
3. Click **Enforce** to enforce multifactor authentication for all NetBackup Flex Scale users.

   Notify all users that they must configure multifactor authentication for their user accounts to be able to successfully sign in.

   See "Configuring multifactor authentication for your user account" on page 256.

Configuring multifactor authentication for your user account when it is enforced in the cluster

After multifactor authentication is enforced in the cluster, you must configure it for your user account if you have not already configured it. If you do not configure multifactor authentication for your account after the enforcement, you cannot sign-in to the appliance and any automation workflow using the login API will also be impacted.

To configure multifactor authentication after the enforcement

1. Open a web browser and go to the following URL.

   https://*console-IP*:14161/login

   The *console-IP* is the management console IP address where the web interface is hosted.
2. Enter the **Username** and **Password**.
3. Click **Sign in**. The **Configure multifactor authentication** screen is displayed.
4. On the next screen, follow the given steps.

   Install and configure an authenticator application on your smart device. It generates a one-time password and sends it to your smart device.

   Supported authenticator applications
5. Scan the QR code with the authenticator application or enter the key manually.

   The manual key should be base32 encoded and can contain between 26 to 208 characters with or without padding.

6 Enter the one-time password that you see in the authenticator application.

7 Click **Submit**.

Successful configuration takes you back to the sign-in screen.

Enter the username, password, and one-time password for successful sign-in.

Resetting multifactor authentication for a user

Only the NetBackup Flex Scale administrator can reset multifactor authentication for other NetBackup Flex Scale users.

Before you reset multifactor authentication:

- The logged in administrator cannot reset his own multifactor authentication. It can only be reset by another Appliance administrator.
- If multifactor authentication is not enforced, then it is possible to reset it for any user.
- If multifactor authentication is enforced, then you can reset it for a local (non AD/LDAP) administrator only if at least one other local administrator is present who has multifactor authentication configured.

To reset multifactor authentication for an NetBackup Flex Scale user

1 Sign in to the NetBackup Flex Scale UI.

2 Go to **Settings > User management**.

3 Navigate to the user row and click on the vertical ellipsis button from the right side of the UI and then select **Reset multifactor authentication** .

4 In the **Reset multifactor authentication** pop-up, click **Reset**.

Chapter 11

Single Sign-On (SSO)

This chapter includes the following topics:

About single sign-on (SSO) configuration

You can configure single sign-on (SSO) with any identity provider (IDP) that uses the SAML 2.0 protocol for exchanging authentication and authorization information. Note that you can configure an IDP with more than one Veritas product.

Note the following requirements and limitations:

- To use SSO, you must have a SAML 2.0 compliant identity provider configured in your environment.
- Configuration of the IDP requires the NetBackup Flex Scale GUI.
- Global logout is not supported.

Configuring SSO on a NetBackup Flex Scale cluster on which both primary and media servers are deployed

Configuring SSO on NetBackup Flex Scale cluster on which both primary and media servers are deployed involves the following steps:

Table 11-1

Task	Description
Configuring SSO on an NetBackup Flex Scale cluster	See To configure SSO on a cluster on which both primary and media servers are deployed
Adding users/group	See To add users/group in RBAC
Configuring an identity provider	See To configure an identity provider
Logging into NetBackup Flex Scale with SSO	See Login with SSO

Configuring SSO on an NetBackup Flex Scale cluster

To configure SSO on a cluster on which both primary and media servers are deployed

1. Go to **Settings > Security management > Single sign-on (SSO)**. Click **Add**.
2. Give the IDP name and upload the IDP metadata xml and optionally provide the custom user field and group field values. The user field and group field values should be same as configured on the IDP. Click **Save**.

   The UI displays a message that confirms that the add identity provider task is triggered. You can click **View Details** to see the progress of the task. 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.
3. Once the configuration is complete, the SSO identify provider details are displayed on the screen. Click **Download service provider xml** to download the details and upload it on IDP server, if required.

Adding users/group in RBAC

To add users/group in RBAC

1. Login to the NetBackup web UI. Go to **Security > RBAC**.
2. Select the Appliance Administrator role and select the **Users** tab.
3. Add the user/group name with domain and select the user as SAML user or SAML group. Click **Add to list**.

Configuring an identity provider

To configure an identity provider

- Login with SSO works only if the configuration on the IDP side is done. Each IDP has different steps for configuration.

  Refer to the following links for the configuration steps for each identity provider.

- ADFS: Enrolling NetBackup Flex Scale primary server as a service provider to ADFS
- Azure: Enrolling NetBackup Flex Scale primary server as a service provider to Azure
- Okta: Enrolling NetBackup Flex Scale primary server as a service provider to Okta
- PingFederate: Enrolling NetBackup Flex Scale primary server as a service provider to PingFederate

Logging into NetBackup Flex Scale with SSO

Login with SSO

1. Navigate to infrastructure GUI login page. Click **Sign-in with single sign-on (SSO)**.
2. Enter SSO credentials and click **Sign in**.

Limitations

There are some limitations when you configure SSO on a NetBackup Flex Scale cluster on which both primary and media servers are deployed.

- Identity provider cannot be edited. It can be removed and added again.
- If the same identity provider is removed and added again with a different name, then all the existing SAML users for that IDP will not be able to login. In such cases, either the admin has to remove and add the SAML users and groups in RBAC again or keep the same name when adding the identity provider.
- Single logout is not implemented. If SAML users log out of the application, and try to login with SSO again, the user is not asked for their login credentials unless the SSO session has expired. This applies to any other application using the same IDP.
- If after identity provider configuration, External certificate authority (ECA) is configured, then login with SSO does not work until the identity provider is updated with the latest service provider metadata xml from the NetBackup Flex Scale. This can be done by downloading the service provider metadata xml from **Settings > Security > Single-Sign on > Download service provider metadata**. This metadata needs to be updated on the IDP side.
- AD/IDP server date, time, and time zone should be the same as the NetBackup Flex Scale cluster. Else, the SSO login fails.
- SAML users or the SAML group users cannot login using the NetBackup Flex Scale login screen for a cluster on which both primary and media servers are deployed.

- SAML users or SAML group users cannot configure multifactor authentication option available in the **Security > Multifactor authentication** section.
- If disaster recovery or primary service replication is configured after the SSO is configured on both the primary and secondary clusters, then the identity provider configured on the secondary cluster ceases to exist and the SAML users in its RBAC cannot login using SSO. Only the primary cluster SAML users can login using SSO on both the clusters.
- If disaster recovery or primary service replication is configured after the SSO is configured on only the secondary cluster, then SSO is unconfigured as its NetBackup primary cluster points to the primary cluster.
- If SSO is configured after disaster recovery configuration from either the primary or secondary cluster, then it is configured for both the clusters and users can login with SSO for both clusters.

Log location

The logs can be found by logging into the NetBackup Flex Scale CLISH, elevating to root and accessing the logs at:

- `/log/VRTSnas/ nbu_sso_config.log`
- `/log/VRTSnas/ isagui_webserver.log`
- `/log/VRTSnas/ isagui_sso_config.log`

The Table 11-2 lists the common error messages.

Table 11-2 Common error messages

Error message	Description
`You are not authorized to access this application`	User is a valid AD/LDAP and IDP user but does not have the Appliance administrator role in NBU RBAC or the Identity provider was deleted and added again with a different name after adding the SAML users in NetBackup RBAC.
`Authentication failed, userPrincipalName field not found in response`	SAML response from the IDP does not contain the user field. This can be due to userPrincipalName field attribute mapping not being created on the IDP side or the custom attribute name is different on the IDP side as provided in the NetBackup Flex Scale.

Table 11-2 Common error messages *(continued)*

Error message	Description
`Unable to get response from identity provider`	Date and time of Identity provider does not match with NetBackup Flex Scale cluster, Identity provider certificate is not updated with latest NetBackup primary certificate, or the certificate revocation check is not disabled on the identity provider.

Configuring SSO on a NetBackup Flex Scale cluster on which only media servers are deployed

Configuring SSO on NetBackup Flex Scale cluster on which only media servers are deployed involves the following steps:

Table 11-3

Task	Description
Configuring SSO on an NetBackup Flex Scale cluster	See To configure SSO on cluster on which only media servers are deployed
Adding users/group	See "Directory services and certificate management" on page 50. See "Adding users" on page 30. **Note:** SSO can be configured only for AD/LDAP users for media server only deployment.
Configuring an identity provider	See To configure an identity provider
Logging into NetBackup Flex Scale with SSO	See Login with SSO

Configuring SSO on an NetBackup Flex Scale cluster

To configure SSO on cluster on which only media servers are deployed

1. Go to **Settings > Security management > Single sign-on (SSO)**. Click **Add**.
2. Give the IDP name and upload the IDP metadata xml and optionally provide the custom user field and group field values. The user field and group field values should be same as configured on the IDP. Click **Save**.

   The UI displays a message that confirms that the add identity provider task is triggered. You can click **View Details** to see the progress of the task. 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.
3. Once the configuration is complete, the SSO identify provider details are displayed on the screen. Click **Download service provider xml** to download the details and upload it on IDP server, if required.

Configuring an identity provider

To configure an identity provider

- Login with SSO works only if the configuration on the IDP side is done. Each IDP has different steps for configuration.

  Refer to the following links for the configuration steps for each identity provider.

  - ADFS: Enrolling NetBackup Flex Scale primary server as a service provider to ADFS
  - Azure: Enrolling NetBackup Flex Scale primary server as a service provider to Azure
  - Okta: Enrolling NetBackup Flex Scale primary server as a service provider to Okta
  - PingFederate: Enrolling NetBackup Flex Scale primary server as a service provider to PingFederate

Logging into NetBackup Flex Scale with SSO

Login with SSO

1. Navigate to infrastructure GUI login page. Click **Sign-in with single sign-on (SSO)**.
2. Enter SSO credentials and click **Sign in**.

Limitations

There are some limitations when you configure SSO on a NetBackup Flex Scale cluster on which only media servers are deployed.

- Identity provider cannot be edited. It can be removed and added again.
- Single logout is not implemented. If SAML users log out of the application, and try to login with SSO again, the user is not asked for their login credentials unless the SSO session has expired. This applies to any other application using the same IDP.
- If after identity provider configuration, External certificate authority (ECA) is configured, then login with SSO does not work until the identity provider is updated with the latest service provider metadata xml from the NetBackup Flex Scale. This can be done by downloading the service provider metadata xml from **Settings > Security > Single-Sign on > Download service provider metadata**. This metadata needs to be updated on the IDP side.
- AD/IDP server date, time, and time zone should be the same as the NetBackup Flex Scale cluster. Else, the SSO login fails.

Log location

The logs can be found by logging into the NetBackup Flex Scale CLISH, elevating to root and accessing the logs at:

- `/log/VRTSnas/ nbu_sso_config.log`
- `/log/VRTSnas/ isagui_webserver.log`
- `/log/VRTSnas/ isagui_sso_config.log`

The Table 11-4 lists the common error messages.

Table 11-4 Common error messages

Error message	Description
`User is not authorized`	User is a valid AD/LDAP and IDP user but does not have the Appliance administrator role in NetBackup Flex Scale user management.
`User principal name missing/ Failed to get user details from identity provider`	SAML response from the IDP does not contain the user field. This can be due to userPrincipalName field attribute mapping not being created on the IDP side or the custom attribute name is different on the IDP side as provided in the NetBackup Flex Scale.

Table 11-4 Common error messages *(continued)*

Error message	Description
`Authentication Failed, Invalid document signature`	Date and time of Identity provider does not match with NetBackup Flex Scale cluster, Identity provider certificate is not updated with latest NetBackup primary certificate, or the certificate revocation check is not disabled on the identity provider.
`Authentication Failed, SAML assertion is not yet valid`	Date and time of Identity provider do not match with NetBackup Flex Scale cluster.
`Single sign-on failed due to an internal error`	Processing SAML callback response failed on NetBackup Flex Scale side due to some exception.

Appendix A

Maintenance procedures for HPE servers

This appendix includes the following topics:

Replacement procedure for a chassis fan

This topic describes the process for replacing a chassis fan in a NetBackup Flex Scale node. Each node contains six fans.

Identifying a chassis fan failure (performed by the CHS team)

The following section describes how to identify a chassis fan failure from NetBackup Flex Scale:

An alert is generated for the node where the fan is malfunctioning. To view the alert, do one of the following from the NetBackup Flex Scale infrastructure management UI:

- Click **Dashboard** in the left pane. 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 management**. On the Alerts and notifications page, use the filters to locate specific types of alerts.

Note: If Call Home is configured for your setup, diagnostic information is sent to the AutoSupport server.

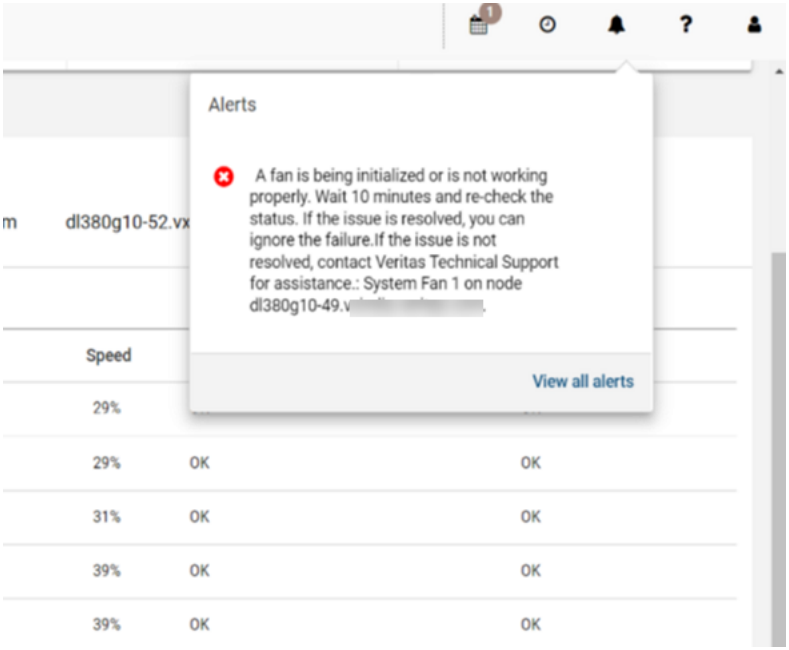

The node where the fan malfunctions is shown as Unhealthy. In the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure > Nodes** to view the health of the nodes.

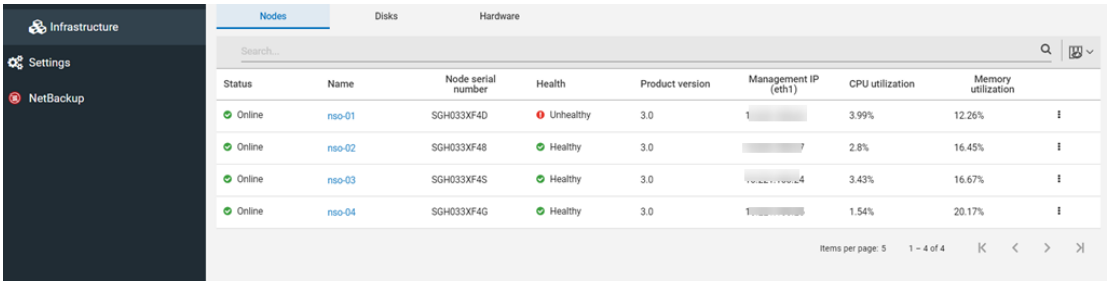

Navigate to **Monitor > Infrastructure > Hardware**, click the node that is unhealthy and for which the alert was generated, and then click **Fan**. The state of the fan is shown as **Failed**.

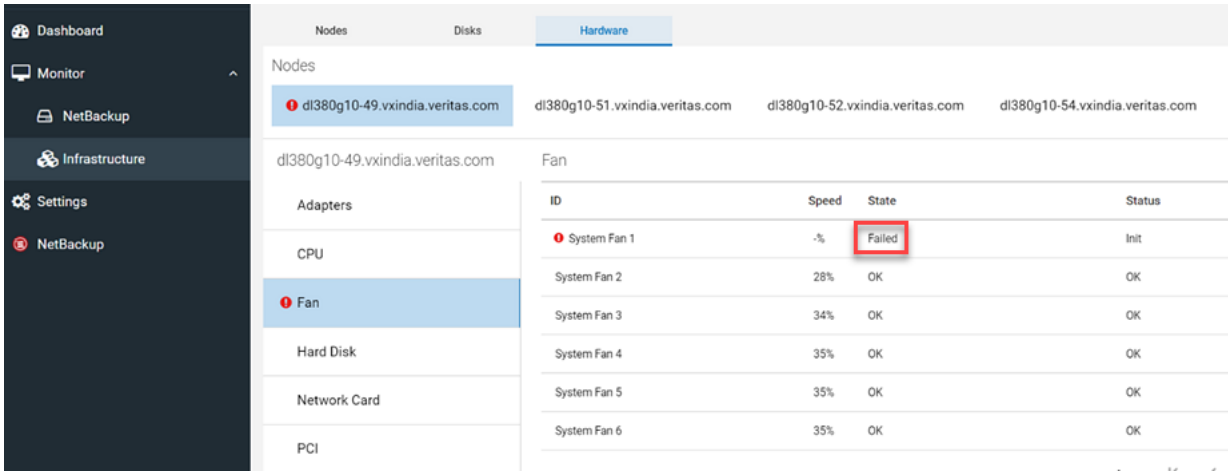

The following section describes how to identify a chassis fan failure from third-party tools:

The HPE Integrated Lights-Out (iLO) remote console shows a failure. Navigate to **Power & Thermal > Fans** and note the details in the **Summary** section for malfunctioning fan.

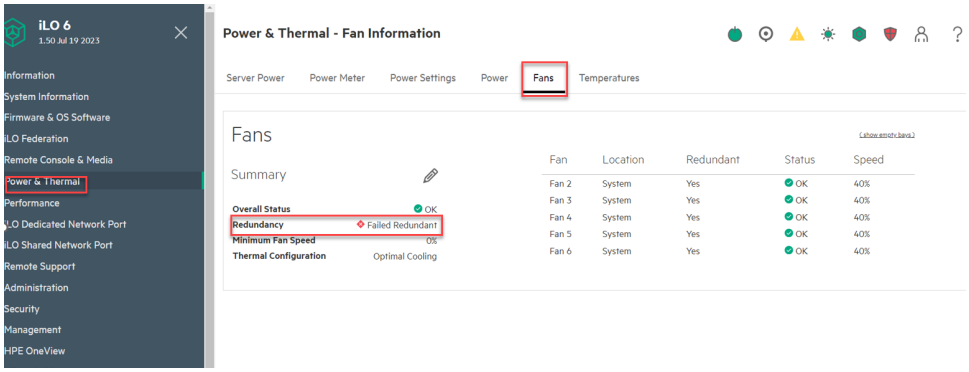

Shutting down the node (performed by Veritas TSE)

You need to shut down the node on which the fan is malfunctioning. For a cluster with fewer than six nodes, only a single node can be down, stopped, or shut down at any given point in time. For a larger cluster of up to 16 nodes, a maximum of two nodes can be down, stopped, or shut down at any given point in time. When you shut down a node, the cluster services running on the node are stopped and the NetBackup jobs running on the node fail over to other cluster nodes. After the hardware maintenance is complete, you need to turn on 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 shut down the node where the NetBackup Flex Scale infrastructure management UI is running, it fails over to another node. It can take a few minutes for the UI to be up on another node.

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

To shut down the node:

1 Shut down the node where the chassis fan failed. Navigate to **Monitor > Infrastructure > Nodes** and click **Shutdown node**.

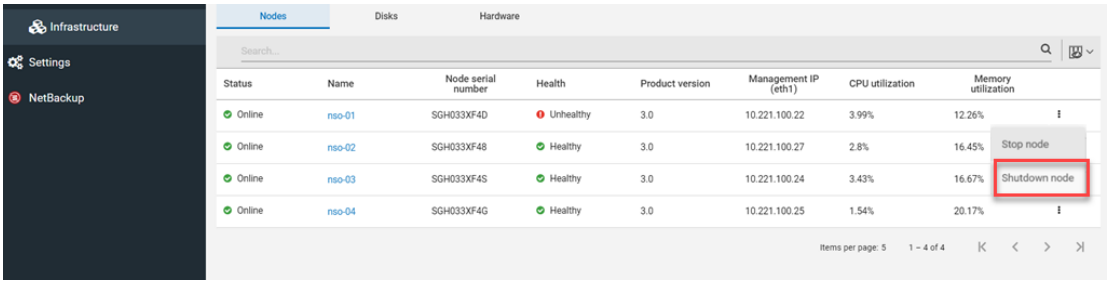

2 Confirm that the node is shut down successfully. In the UI, you can view the notification at the top of the page. In the iLO remote console, wait until the system shows the **System halted** message.

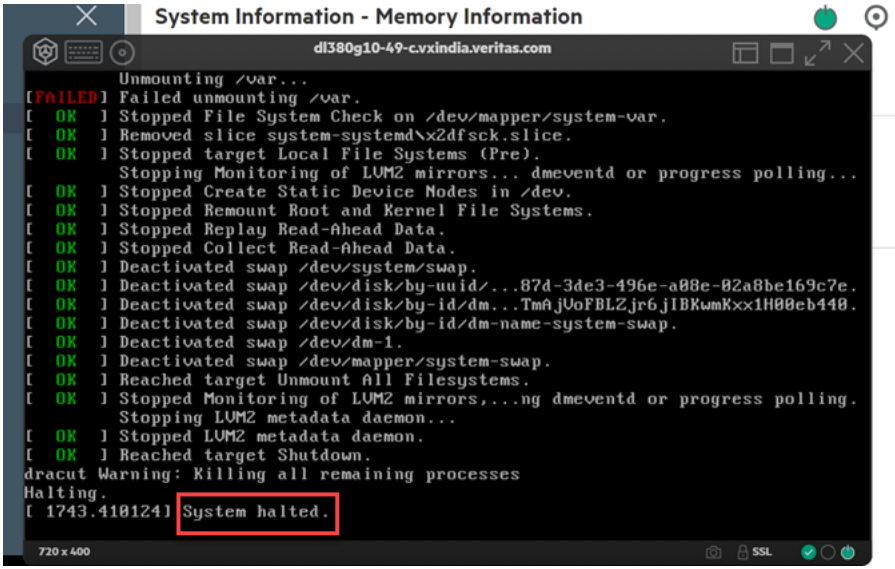

3 Shut down the node. Press the Power button on the front panel of the server or from the iLO remote console use the **Server Power > Press and Hold** option.

4 Contact the hardware vendor to replace the hardware component.

Replacing the chassis fan (performed by the HPE vendor)

The HPE representative replaces the chassis fan.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

Power on the server from the iLO remote console using the **Server Power > Press and Hold** option and wait till it joins the cluster.

Check the node status and the services status in the NetBackup Flex Scale infrastructure management UI:

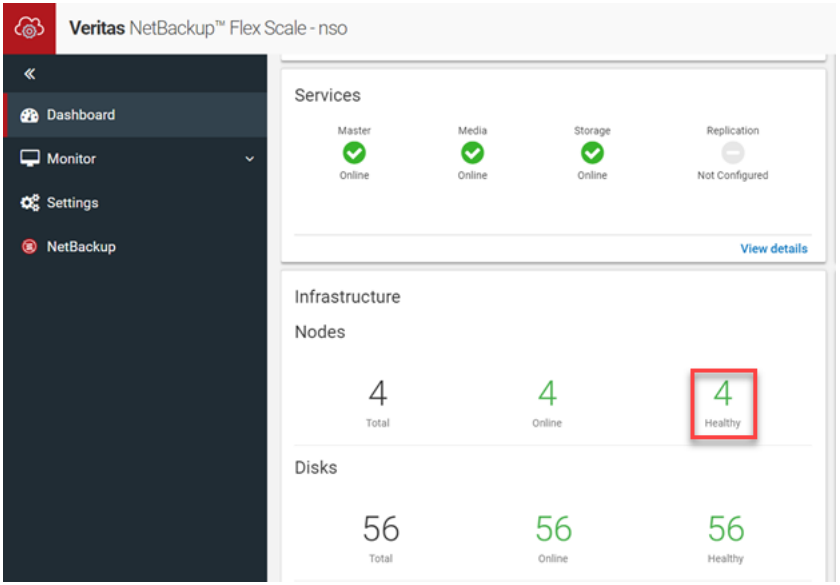

Verify the fan status in the NetBackup Flex Scale infrastructure management UI. Navigate to **Monitor > Infrastructure > Hardware > Fan**.

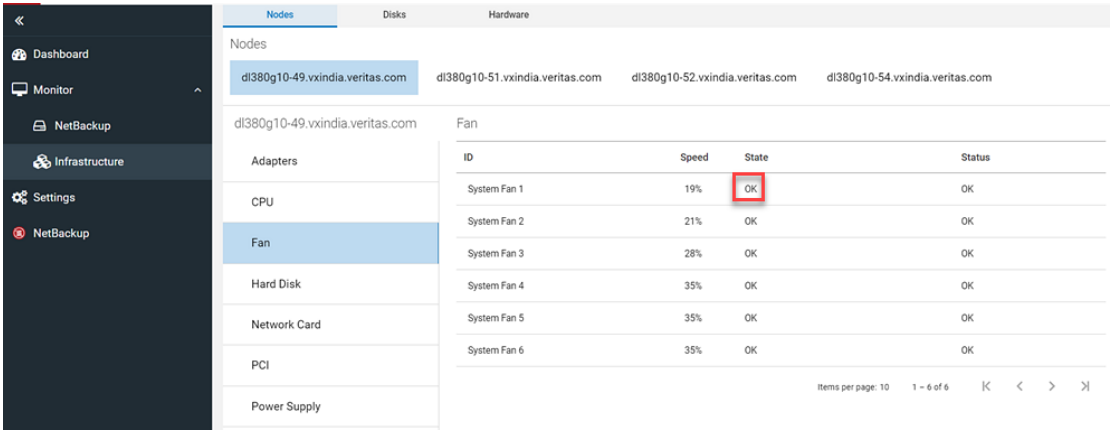

Verify the fan status in the iLO remote console. Navigate to **Power & Thermal > Fans**.

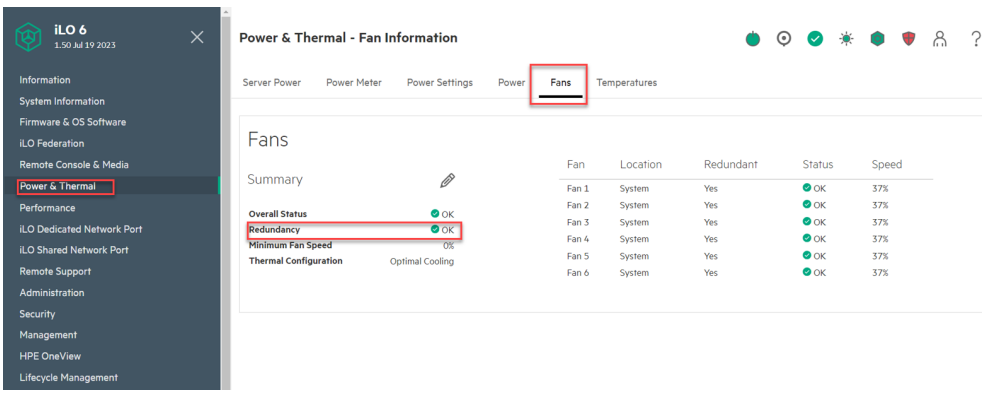

Replacement procedure for power supply

This topic describes the process for replacing a power supply unit in a NetBackup Flex Scale node. Each node contains two power supply units. If one of the PSUs fails, the node can function. If both the units fail, the node becomes unhealthy.

Identifying a power supply failure (performed by the CHS team)

The following section describes how to identify a power supply failure from NetBackup Flex Scale:

A node can tolerate failure of a single power supply unit. However, you are required to replace the failed power supply unit.

If both the power supply units fail, the node is shown as unhealthy because the node is down. In the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure > Nodes** to view the health of the nodes. If the NetBackup Flex Scale infrastructure management UI is running on the node where the failure occurs, the UI fails over to another cluster node.

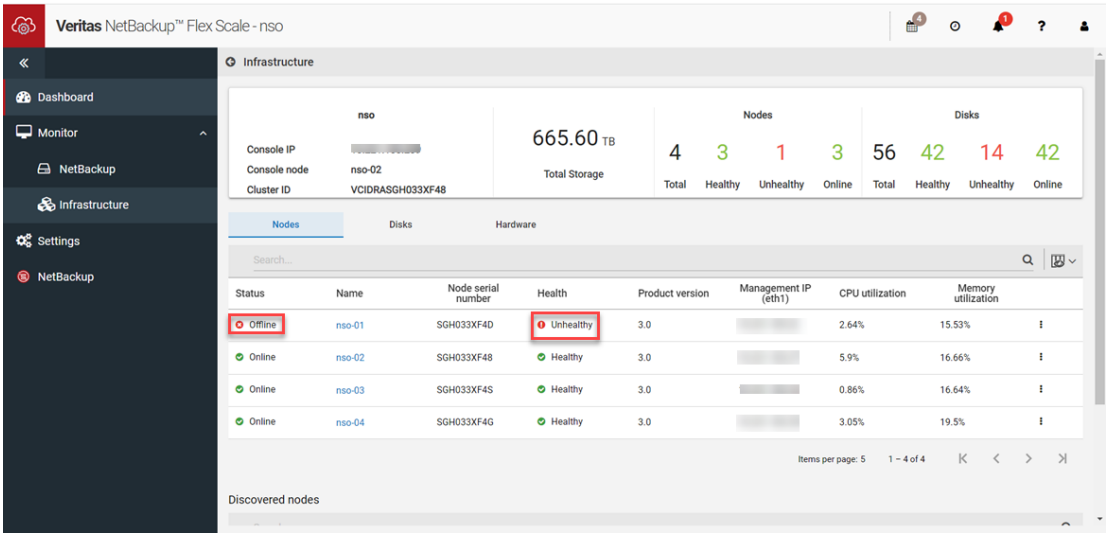

The following section describes how to identify a power supply failure from third-party tools:

The HPE Integrated Lights-Out (iLO) remote console shows a failure. Navigate to **Power & Thermal > Power** and note the status of the malfunctioning power supply unit.

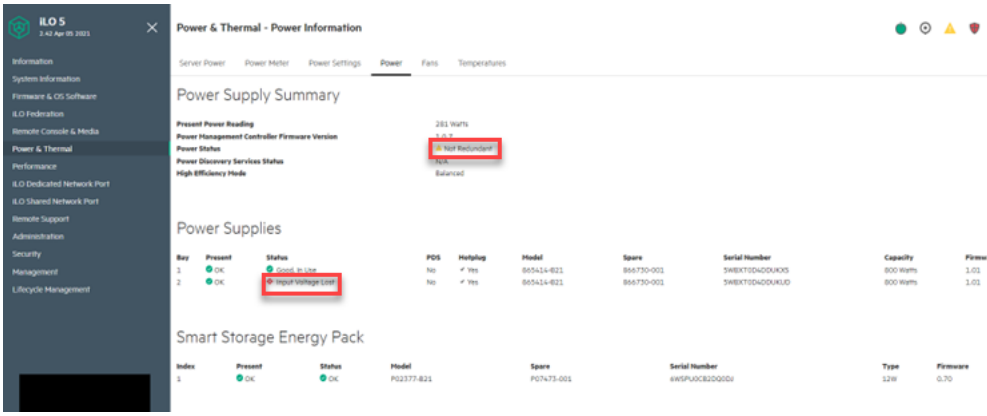

Collecting HPE Active Health System (AHS) logs (performed by Veritas Support)

Before you contact the hardware vendor for replacing the failed component, collect AHS logs. To collect the AHS logs, in the NetBackup Flex Scale infrastructure management UI, navigate to **Settings > Diagnostics > Basic > Appliance**.

Replacing the power supply (performed by the HPE vendor)

Contact the hardware vendor to replace the hardware component. An HPE representative replaces the power supply.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

Reinsert the power cables back into the power supply and the node will be healthy again. Verify the node status in the NetBackup Flex Scale infrastructure management UI:

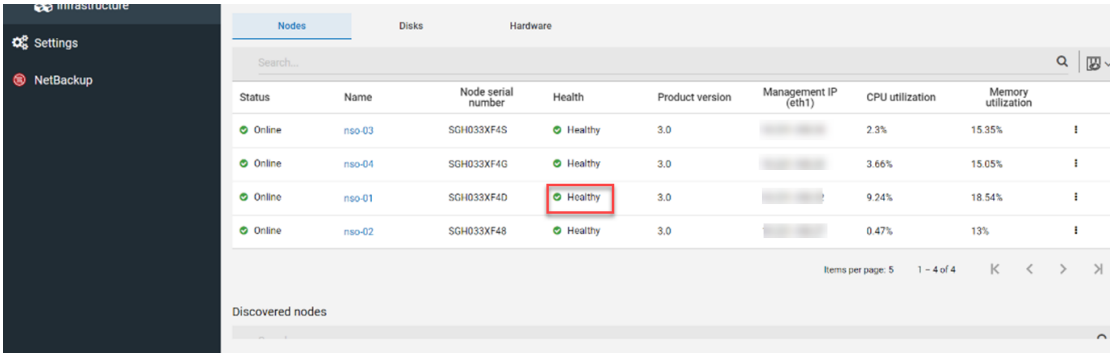

Verify that all the services are up and running:

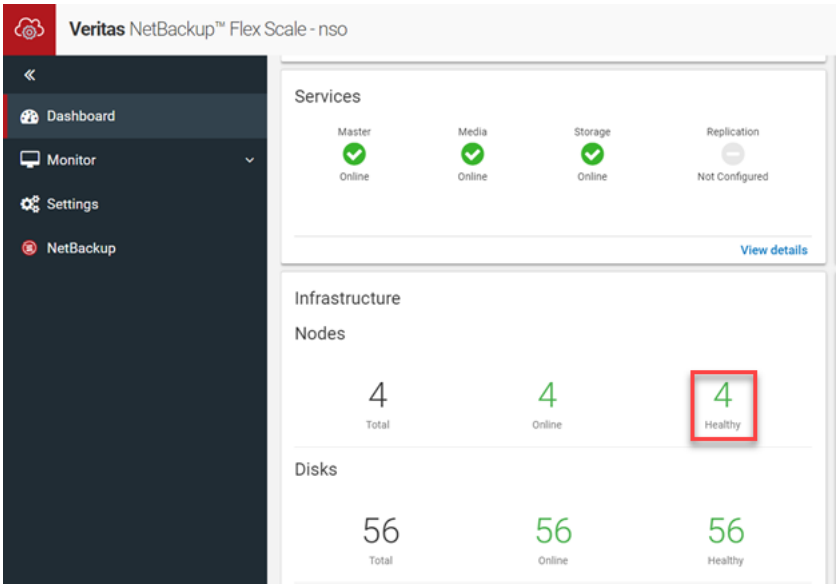

```
                        Power Supply Information

+---------------+--------+----------+-----------------+--------------+-------+
| ID            | Status | Wattage  | LineInputVoltage | HighWaterMark | State |
+---------------+--------+----------+-----------------+--------------+-------+
| Power Supply 1 | OK    | 90 Watts | -               | 899 Watts    | OK    |
+---------------+--------+----------+-----------------+--------------+-------+
| Power Supply 2 | OK    | 165 Watts | -              | 899 Watts    | OK    |
+---------------+--------+----------+-----------------+--------------+-------+
```

To view the power supply status in the NetBackup Flex Scale management infrastructure UI, navigate to **Monitor > Infrastructure > Hardware > Power Supply**.

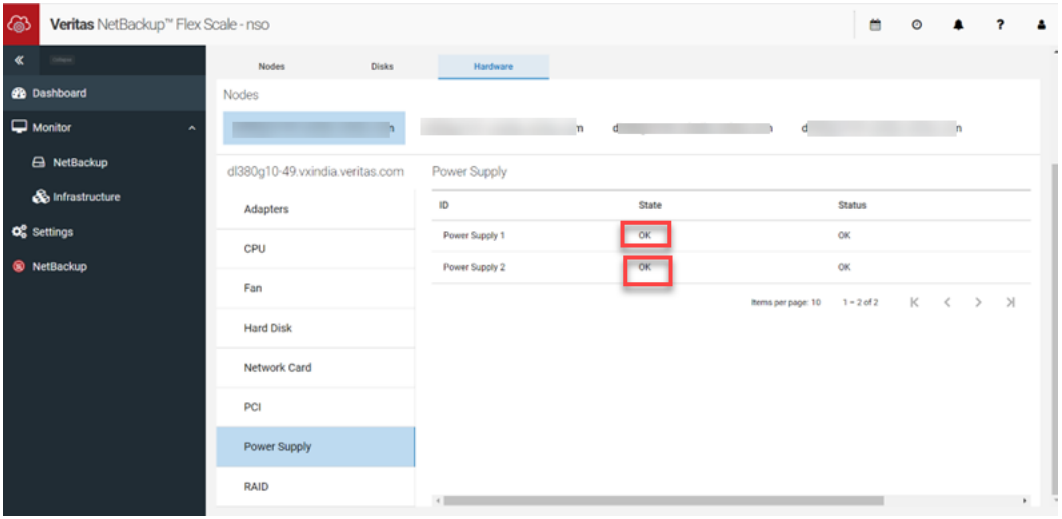

Verify the power supply status in the iLO remote console.

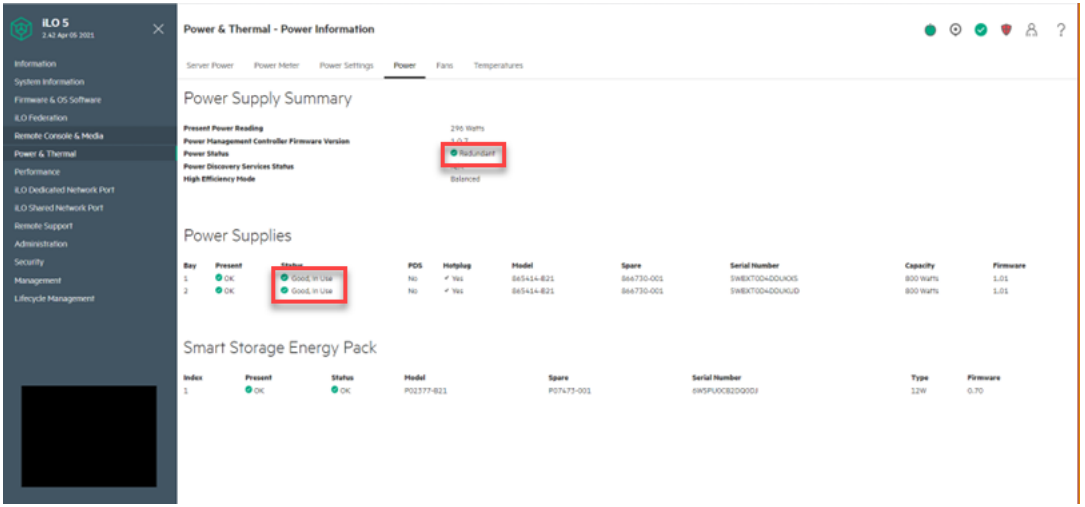

Replacement procedure for a single OS disk

This topic describes the process of replacing a single OS disk that failed or is unreachable. Each node has two OS disks.

Identifying an OS disk failure (performed by the CHS team)

The following section describes how to identify a single OS disk failure from NetBackup Flex Scale:

An alert is generated for an OS disk failure or for an unreachable disk. To view the alert, do one of the following from the NetBackup Flex Scale infrastructure management UI:

- Click **Dashboard** in the left pane. In the **Alerts** area, click **View details** to see a complete list of alerts.
- At the top of any screen, click the **Bell** icon.

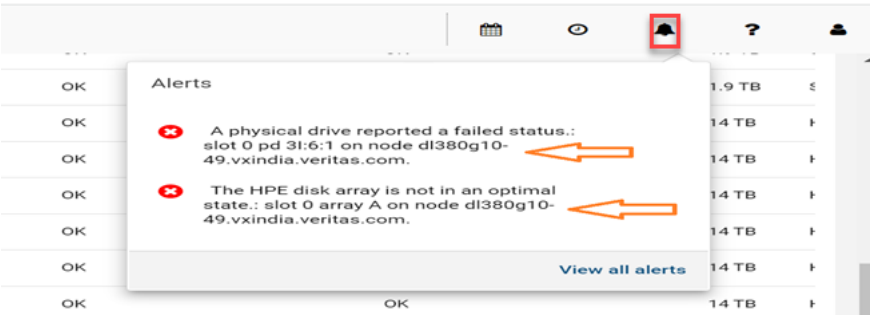

- Click **Settings > Alerts management**. On the Alerts management page, use the filters to locate specific types of alerts.

If SMTP is configured for AutoSupport, you receive email alerts. If Call Home is configured for your setup, diagnostic information is sent to the AutoSupport server.

Navigate to **Monitor > Infrastructure > Hardware** and select the node on which the OS disk went bad, and then click **Hard Disk**. The UI shows the failure for the corresponding OS disk:

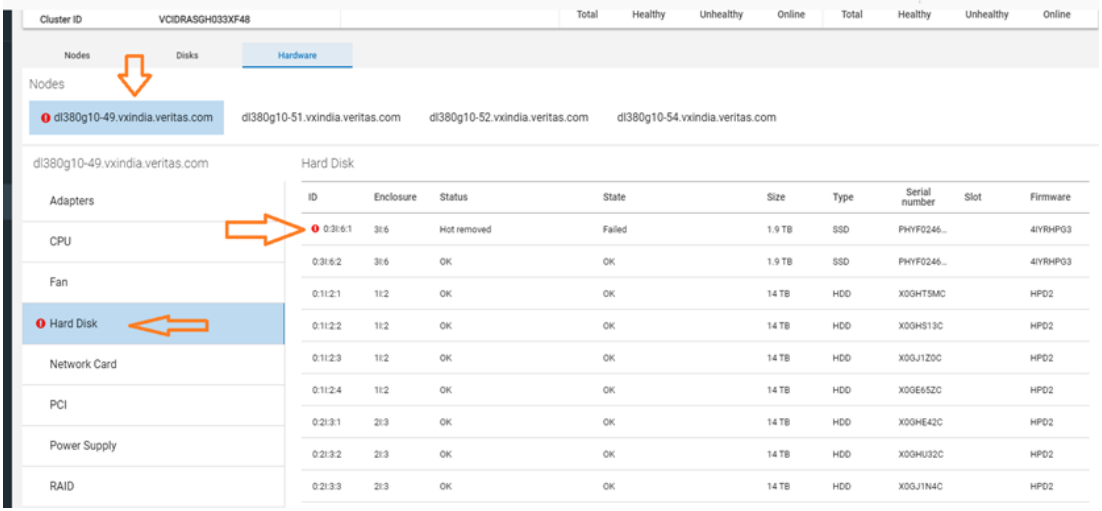

The following section describes how to identify an OS disk failure from third-party tools:

The HPE Integrated Lights-Out (iLO) remote console shows a failure. The Health for the OS disk is shown as **Critical** and **Warning** for the Volume of the RAID 1 in iLO.

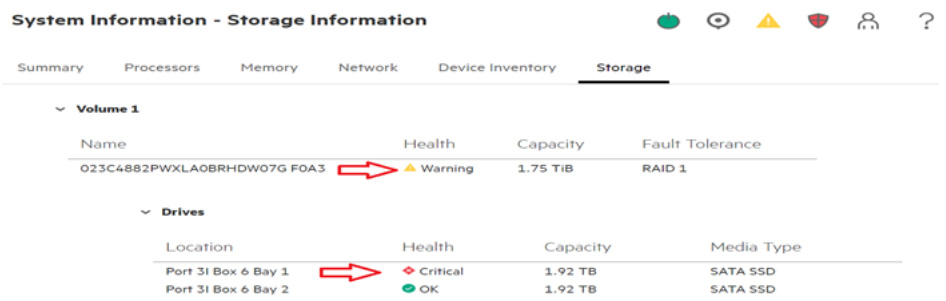

The health of the node is shown as unhealthy for that node in the NetBackup Flex Scale UI. Navigate to **Monitor > Infrastructure > Nodes** to view the node health.

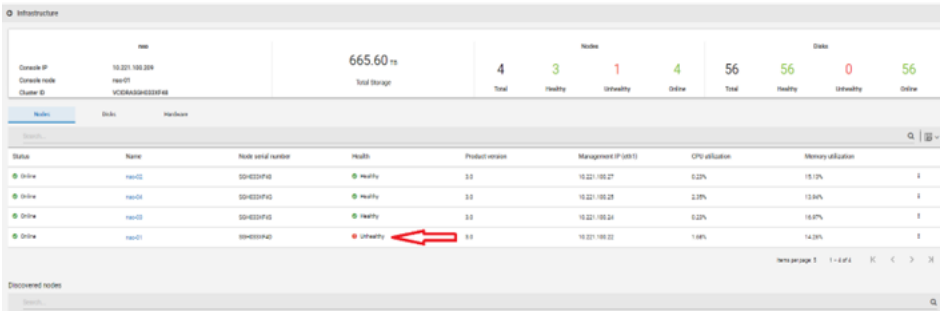

Replacement procedure (performed by the HPE vendor)

An HPE representative identifies the faulty disk, its physical location in the appliance, and replaces the faulty OS disk. You can use the AHS logs to find the required details, and then replace the disk.

Note: With NetBackup Flex Scaleversion 3.1, you can beacon the disk from the UI.

After you get the physical location of the disk on the appliance, replace the OS disk with a new OS disk. Note the model number of the new disk and ensure that it matches with the older one.

To replace the disk, the HPE representative completes the following steps:

1 Check the disk model number from the iLO remote console.

2 Identify the corresponding location of the OS disks in the appliance. In this example, Box6 – Bay 1 and Bay 2.

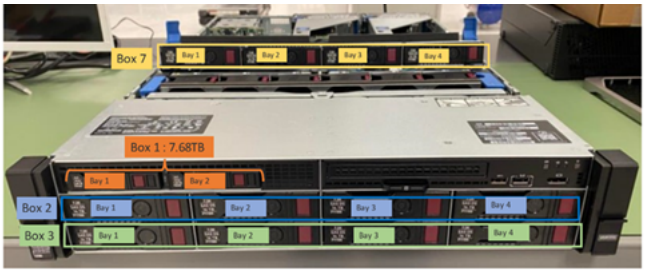

3 Refer to the HPE procedure to replace the disk.

4 In iLO, after the OS disk is replaced, Health for the OS disk is set to **OK** but the Health of the Volume of the RAID 1 is set to **Warning** till the rebuild completes.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

To verify that the issue is resolved, Veritas TSE completes the following steps:

1. Wait till the RAID controller rebuilds the new OS disk. This operation takes approximately two hours. To check the rebuild progress, run the following command after elevating to root access:

```
nbfs3.1> support elevate
#    ssacli ctrl all show config
HPE Smart Array P816i-a SR Gen10 in Slot 0 (Embedded)  (sn:
PWXLA0BRHDW07G)
   Internal Drive Cage at Port 1I, Box 2, OK
   Internal Drive Cage at Port 2I, Box 3, OK
   Internal Drive Cage at Port 3I, Box 6, OK
   Internal Drive Cage at Port 4I, Box 7, OK
   Port Name: 1I (Mixed)
   Port Name: 2I (Mixed)
   Port Name: 3I (Mixed)
   Port Name: 4I (Mixed)
   Array A (Solid State SATA, Unused Space: 0  MB)
     logicaldrive 1 (1.75 TB, RAID 1, Recovering, 4.13% complete)

      physicaldrive 3I:6:1 (port 3I:box 6:bay 1, SATA SSD, 1.9
TB, Rebuilding)
      physicaldrive 3I:6:2 (port 3I:box 6:bay 2, SATA SSD, 1.9
TB, OK)
```

2. After the rebuild completes successfully, verify that all the AutoSupport alerts are resolved and the node state shows healthy in the NetBackup Flex Scale UI. To verify, navigate to **Monitor > Infrastructure > Nodes**.

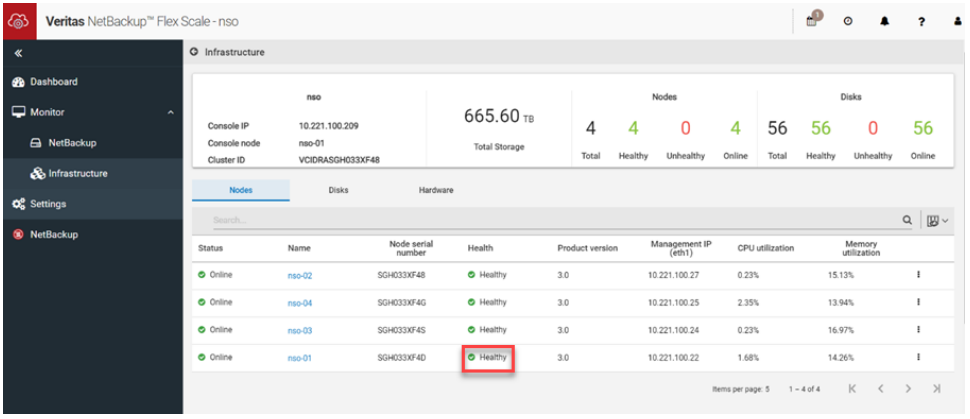

3 In iLO, verify that the Health of the Volume for the RAID 1 is set to **OK**.

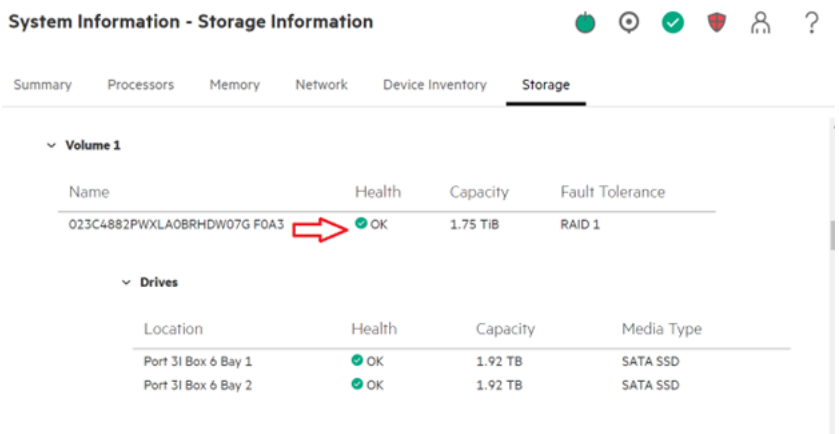

Replacement procedure for both OS disks on a non- management console node

This topic describes the process of replacing both the OS disks that failed or are unreachable on a non- management console node. Each node includes two OS disks.

Identifying the management console node

To identify the management console node, in the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure**. On the **Infrastructure** page, **Console node** shows the management console node and the remaining nodes are the non-management console nodes.

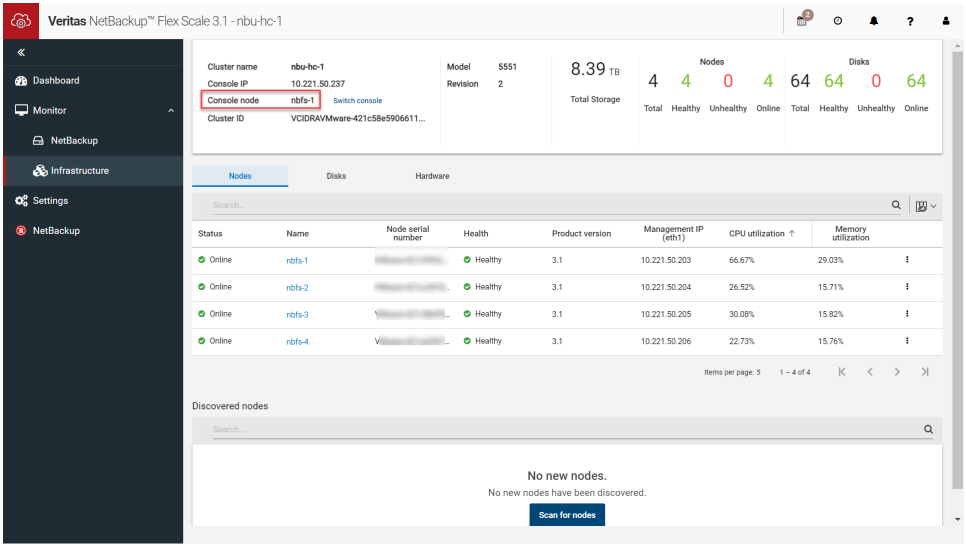

Identifying OS disk failures on a non- management console node of the cluster (performed by the CHS team)

If both OS disks go bad or are unreachable on a non-management console node of the cluster, you receive email alerts if SMTP is configured for AutoSupport. In the NetBackup Flex Scale infrastructure management console UI, the node status is shown unhealthy for the node where the failure occurs. Navigate to **Monitor > Infrastructure > Nodes** to view the node status.

Note: For NetBackup Flex Scale version 3.0, stale information is displayed for the node state. The node is shown healthy.

The HPE Integrated Lights-Out (iLO) remote console shows a failure. The health of the OS disks and the logical drive of RAID 1 is shown as **Critical** in the iLO remote console.

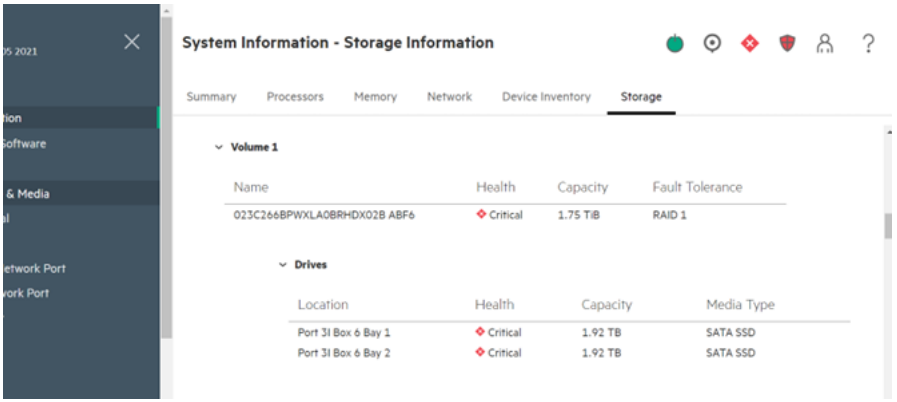

Replacing the faulty OS disks (performed by the HPE vendor)

An HPE representative identifies the faulty disk, its physical location in the appliance, and replaces the OS disks. The representative can use the AHS logs to find the required details, and then replaces both the disks simultaneously so that the logical drive (RAID 1) can be created.

Note: With NetBackup Flex Scale version 3.1, you can beacon the disk from the UI

When replacing the disks, ensure that the new disks have the same model number as the old disks. Check the model number in the iLO remote console. To create the logical drive (RAID 1 configuration), ensure that both the OS disks are replaced at the same time.

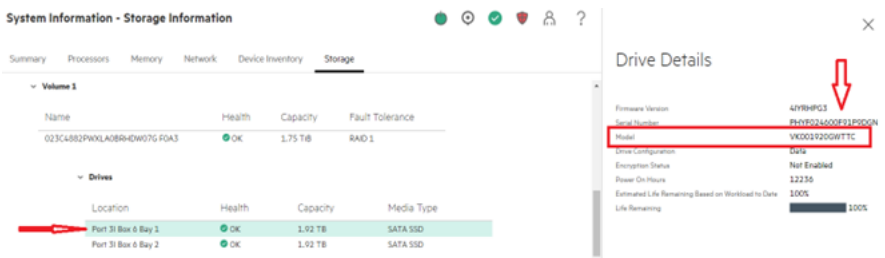

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

After physical disk replacement, check OS disks and the logical drive status in the iLO remote console. If the logical drive status is shown as **Critical**, fix the failed drive; else proceed with node replacement.

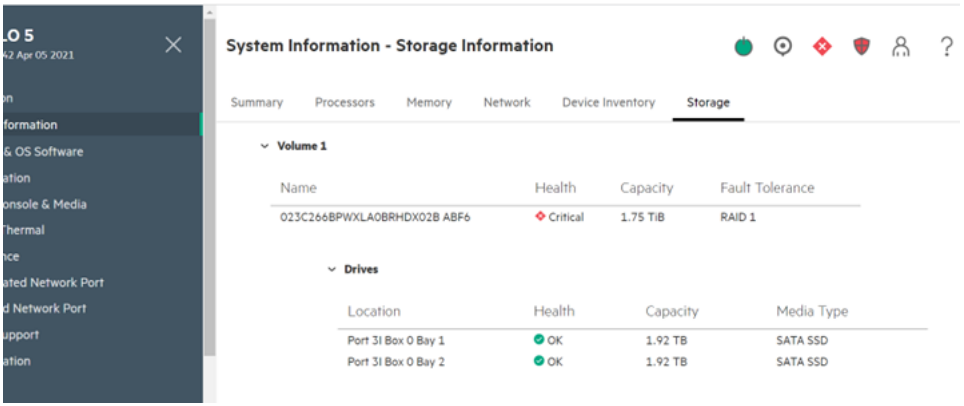

To fix the failed logical drive, identify the logical drive, delete the existing failed logical drive, and create it again. After you recreate the logical drive, replace the node where both the OS disks have failed.

Identifying the OS logical drive

To identify the OS logical drive:

- In the NetBackup Flex Scale UI, you get an alert with slot and physical drive ID for faulted disk. Use this information to identify the volume that has that physical drive.
- The OS disks are of 2TB each. Check the volume that has those disks in the iLO remote console.

Once you identify the failed logical drive, delete the existing failed logical drive, and recreate it.

Deleting the logical drive

To delete the logical drive:

1. Log in to the iLO remote console and launch the console for the node:
   - Click the **Menu** icon, and then click **Power > Cold Boot**.

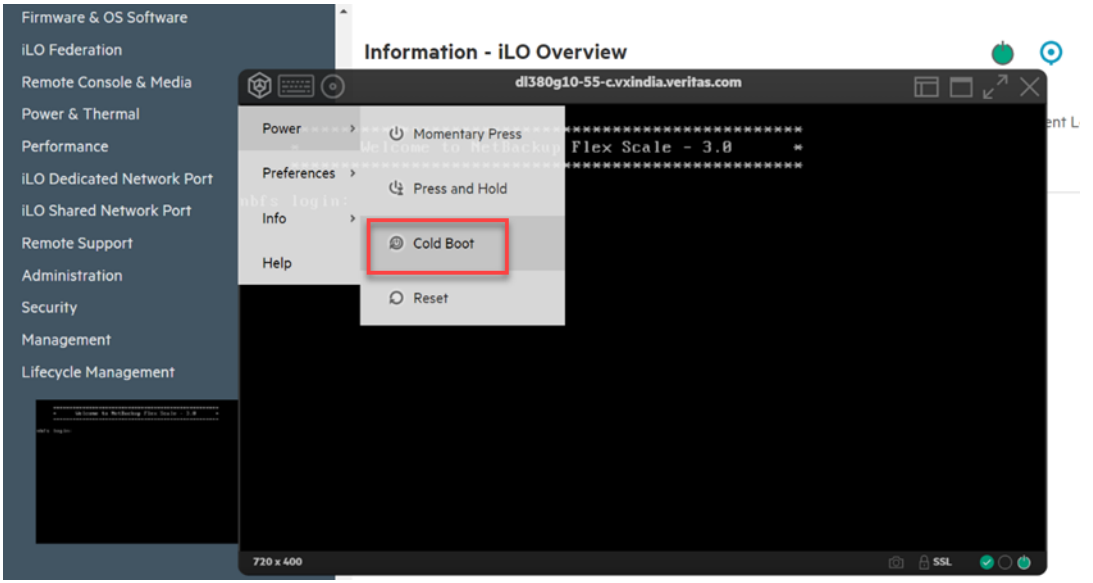

- Press Fn + F10 to select **Intelligent Provisioning**.

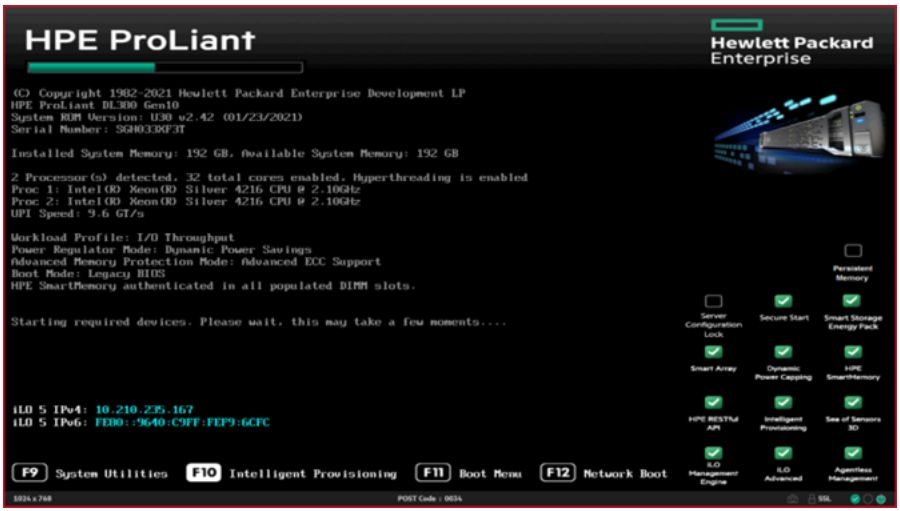

- Click **Smart Storage Administrator** and press Enter.

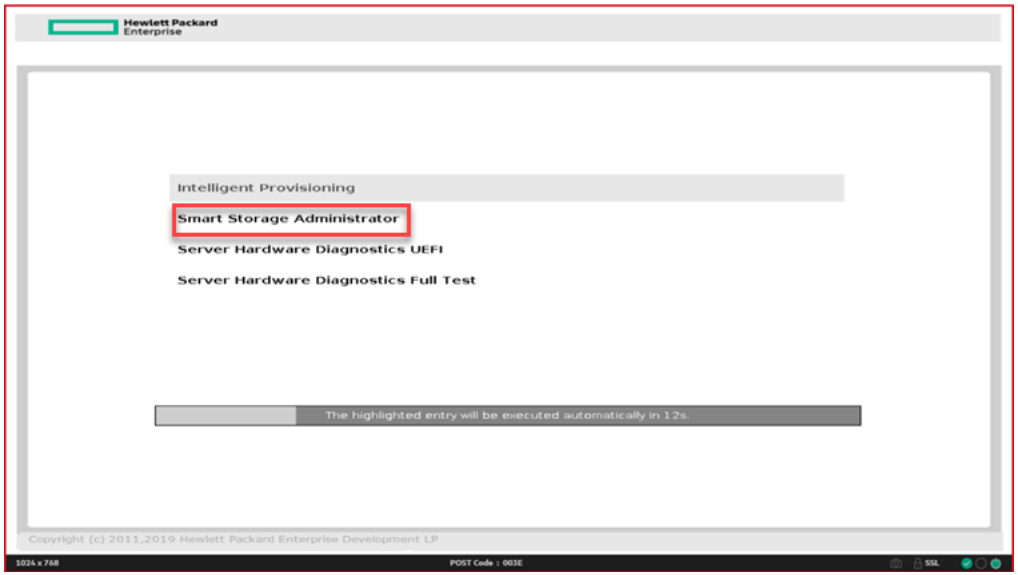

This will take you on the HPE smart array (controller) page.

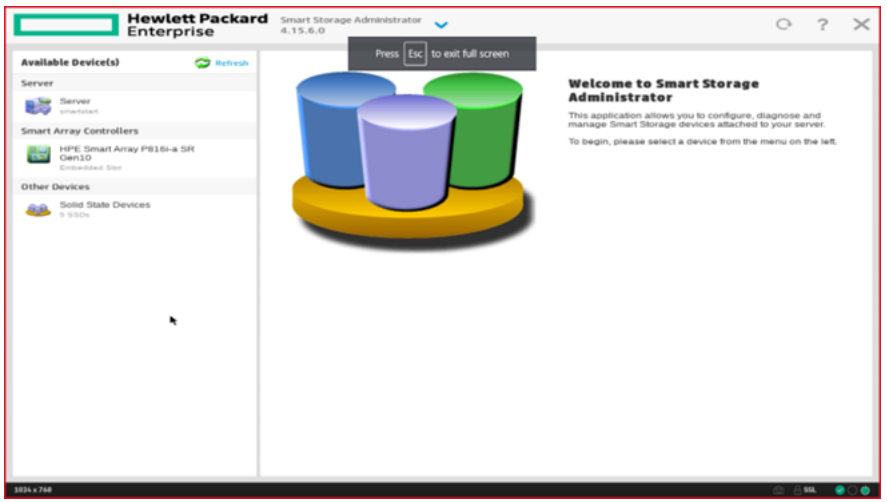

2 In the left pane, click **HPE Smart Array P816i-a SR Gen10**.

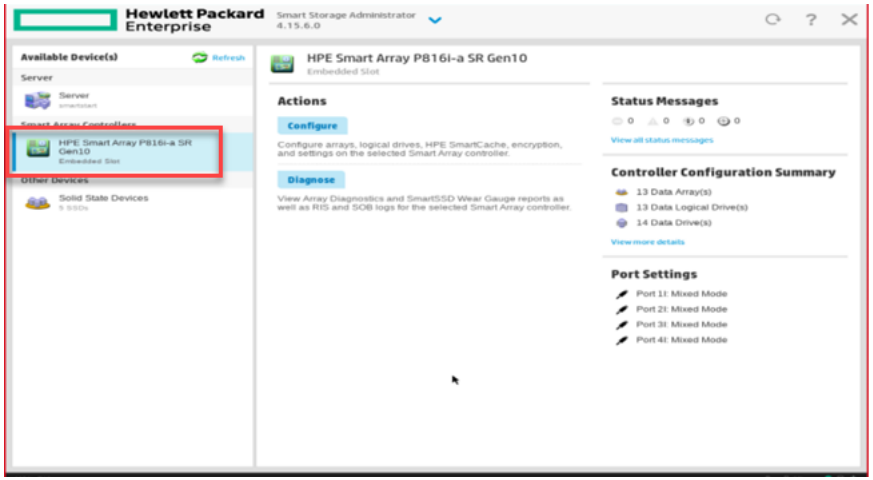

- Click **Configure** to expand the list of **Controller Devices** and **Tools** along with available **Actions** as Storage Administrator.

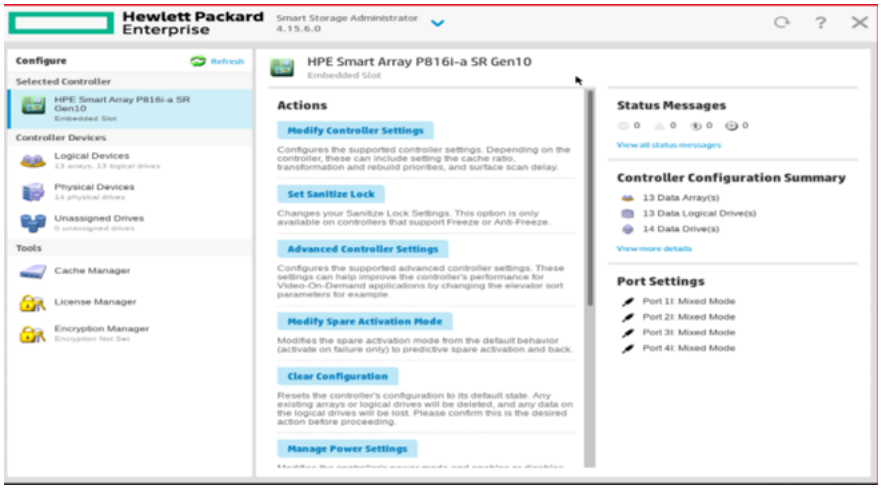

- Click **Logical Drives** in the left pane.
  The Array details are displayed.

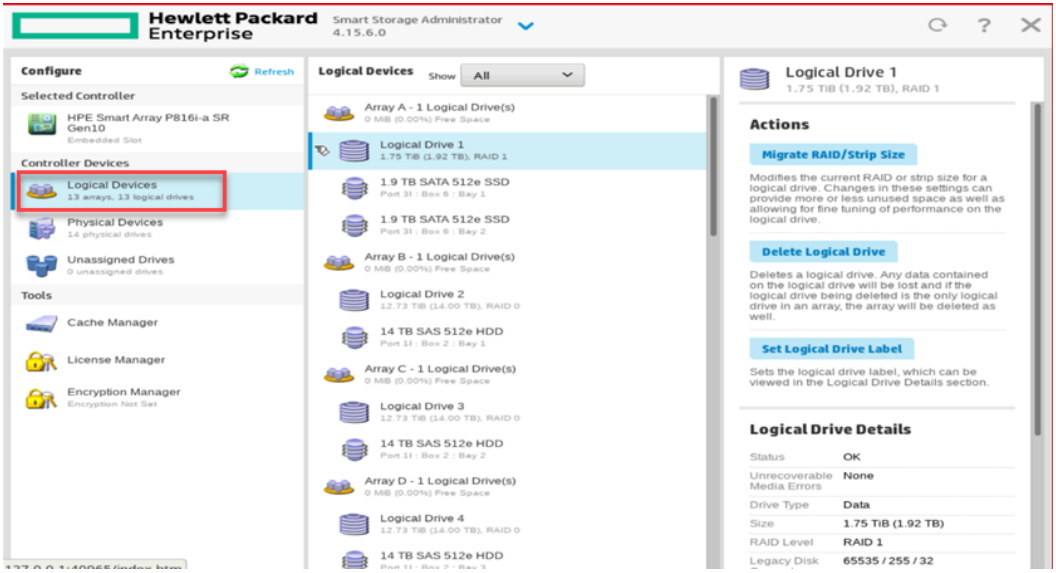

3 Delete the logical drive:

- Select the logical drive and click **Delete Logical Drive**.

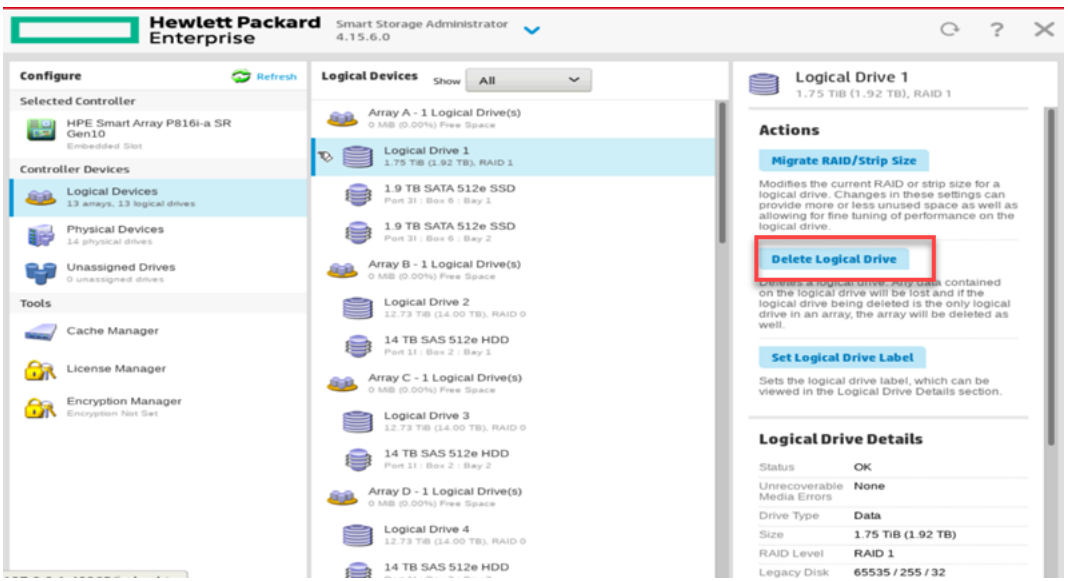

- Click **Yes** to confirm.

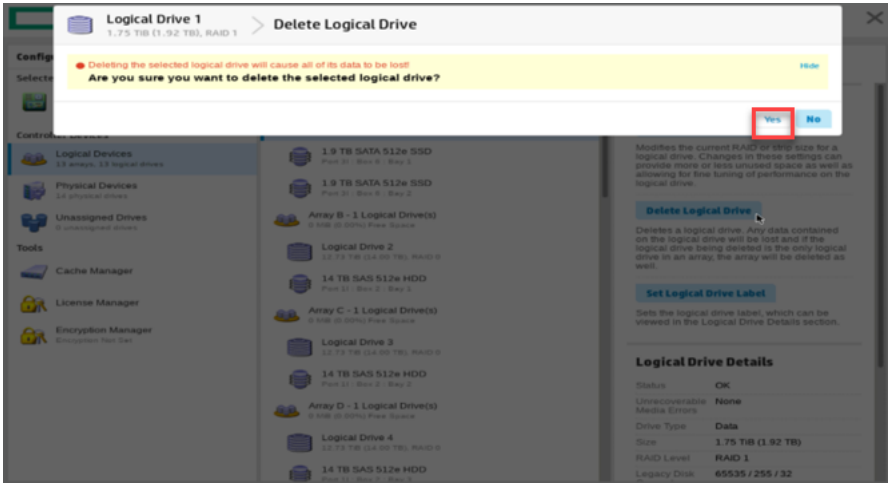

- Click **Finish**.

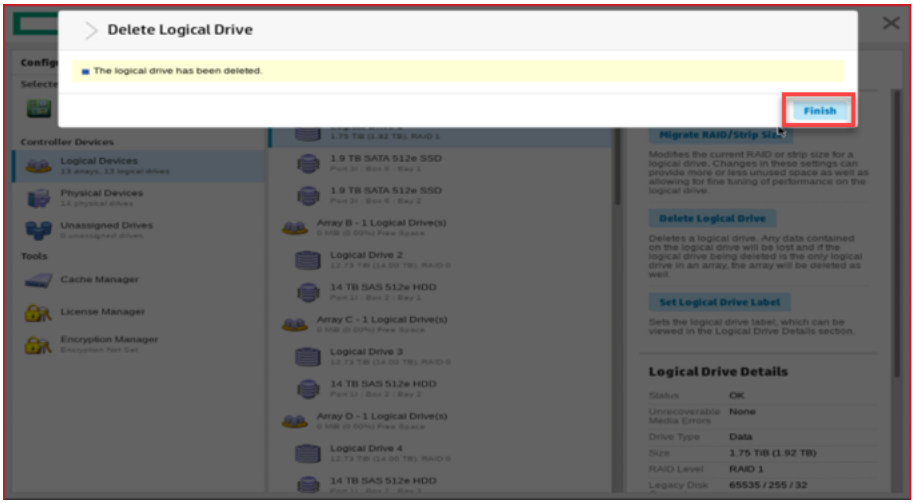

4 The deleted logical drive is no longer listed once you delete it. Ensure that under **Unassigned Drives**, **2 unassigned drives** is displayed.

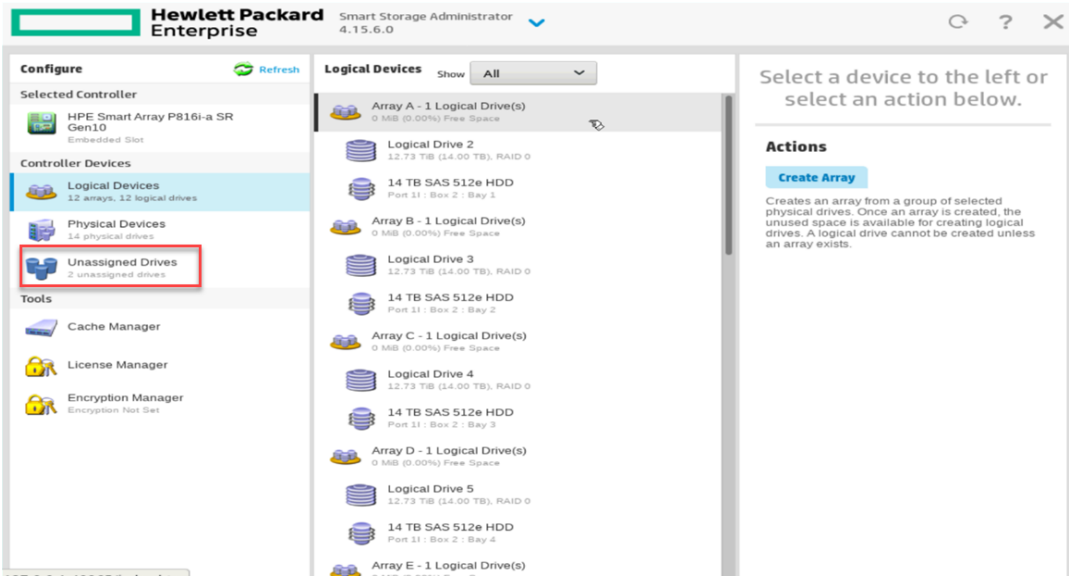

Recreating the logical drive

To create the logical drive again:

1 Log in to the iLO remote console and launch the console for the node:

- Click the **Menu** icon, and then click **Power > Cold Boot**.

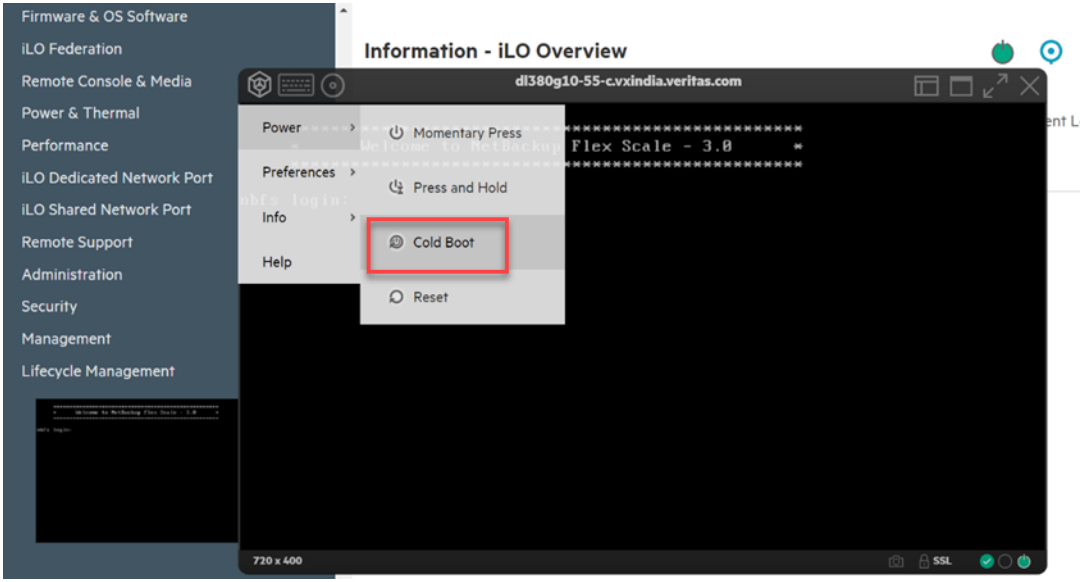

- Press Fn + F10 to select **Intelligent Provisioning**.

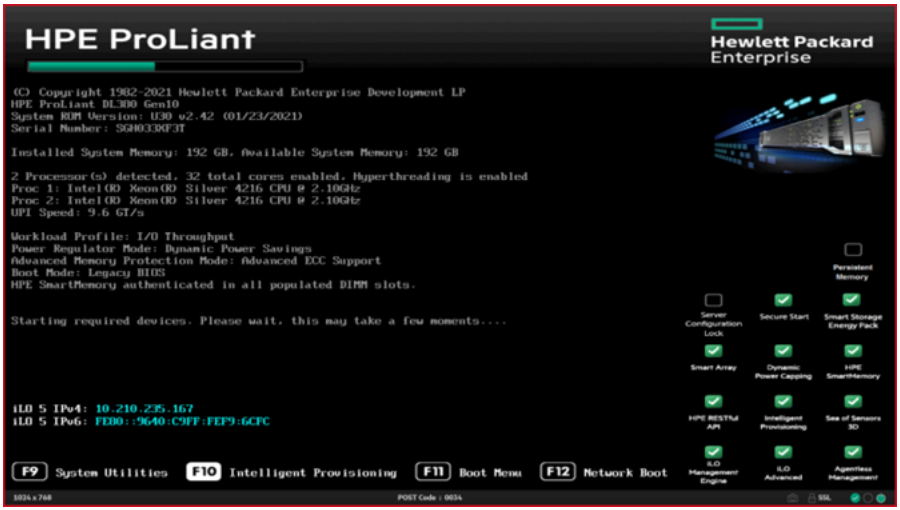

- Click **Smart Storage Administrator** and press Enter.

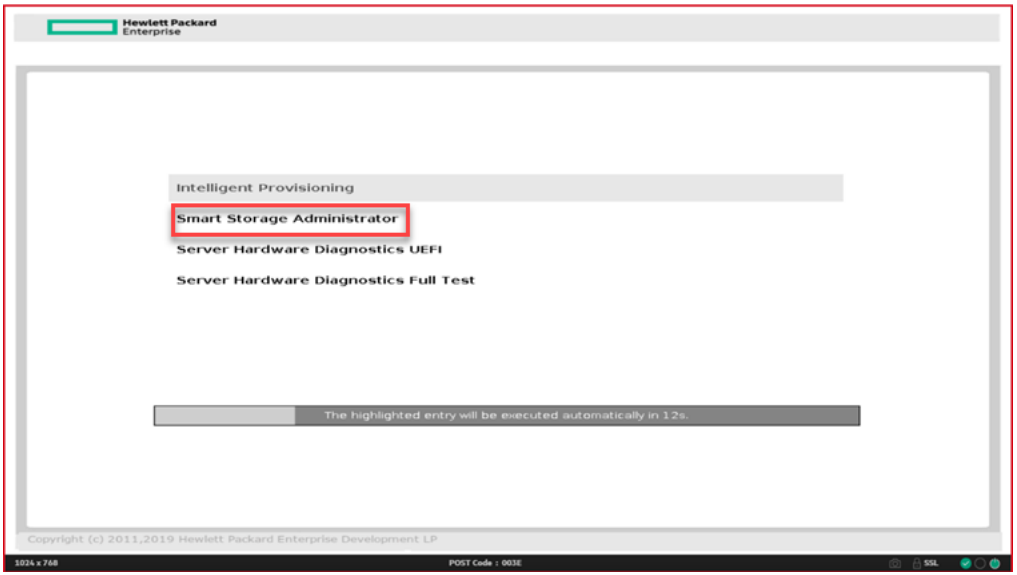

The HPE smart array (controller) page is displayed.

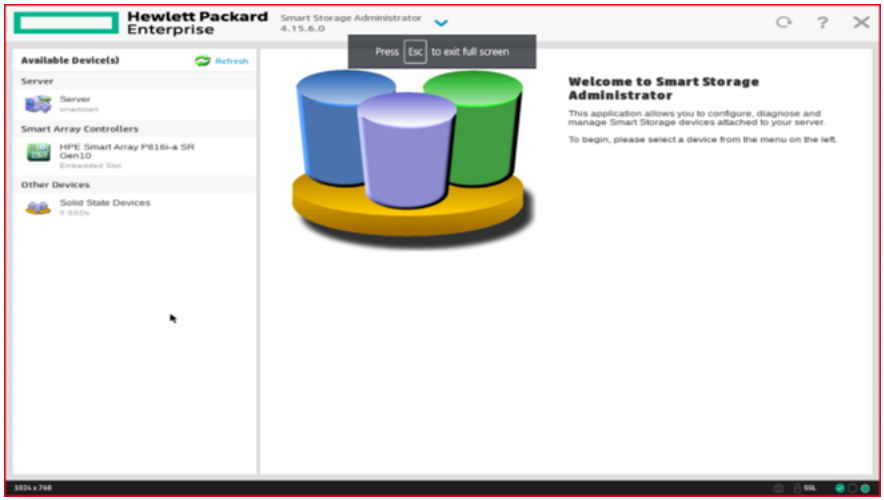
Hewlett Packard Enterprise
Smart Storage Administrator
4.15.6.0
Press Esc to exit full screen
Available Device(s)
Refresh
Server
Server
Smart Array Controllers
HPE Smart Array P816i-a SR Gen10
Other Devices
Solid State Devices
Welcome to Smart Storage Administrator
This application allows you to configure, diagnose and manage Smart Storage devices attached to your server.
To begin, please select a device from the menu on the left.

2 In the left pane, click **HPE Smart Array P816i-a SR Gen10**.

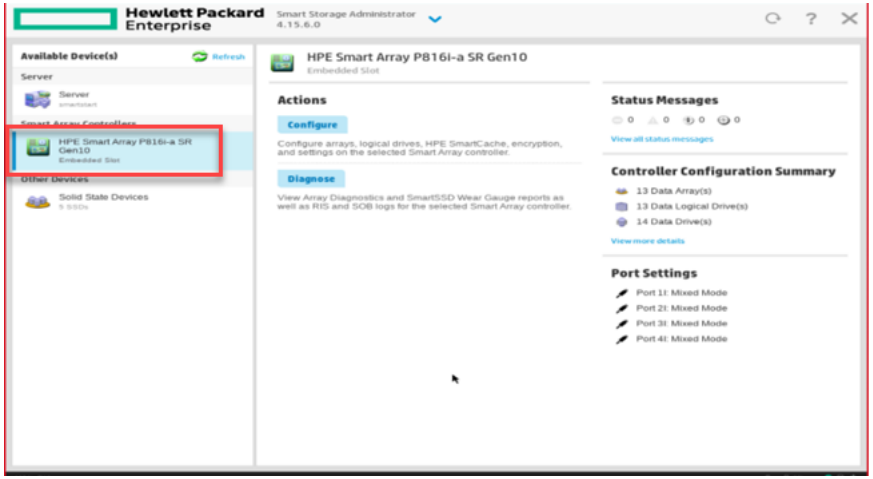

This will expand the list of **Controller Devices** and **Tools** along with available **Actions** as Storage Administrator.

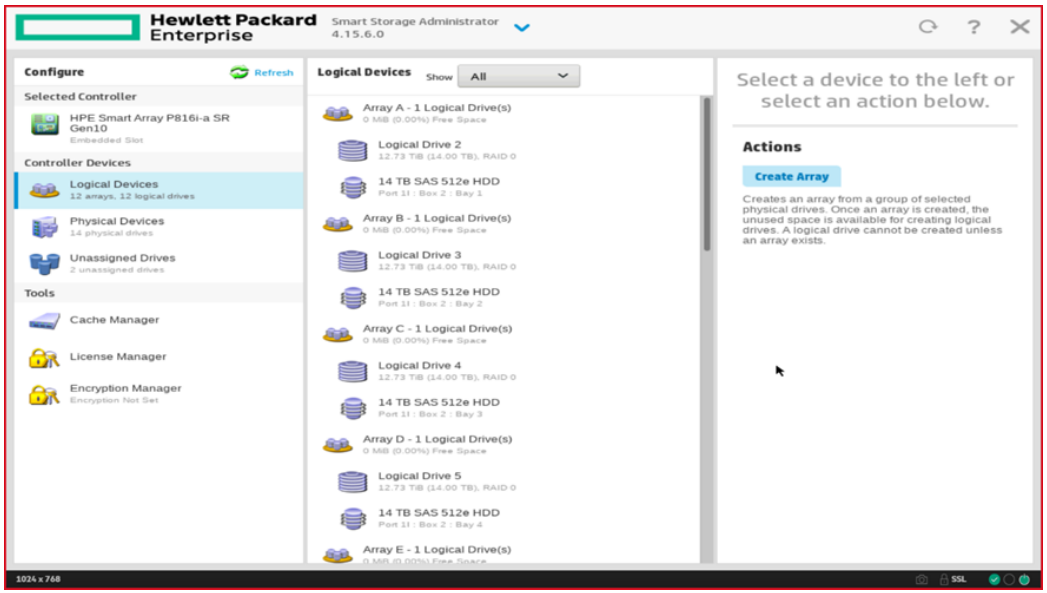

3 To create logical drive, click **Unassigned Drives** in the left pane, and then click **Create Array**.

4 Select OS disks from **Unassigned Drives** and click **Create Array**.

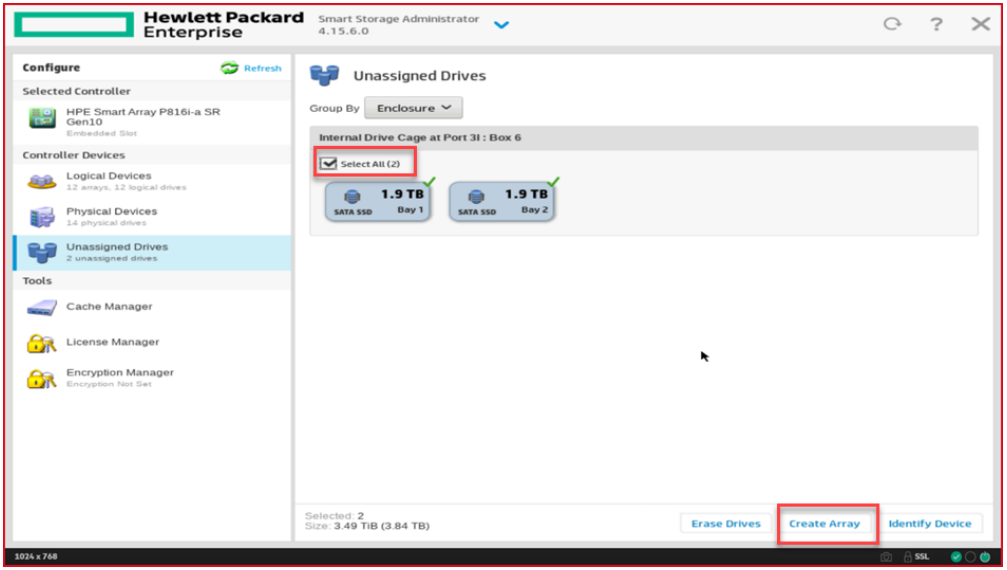

5 Click **Yes** to confirm.

6 Click **RAID 1**, and then click **Create Logical Drive**.

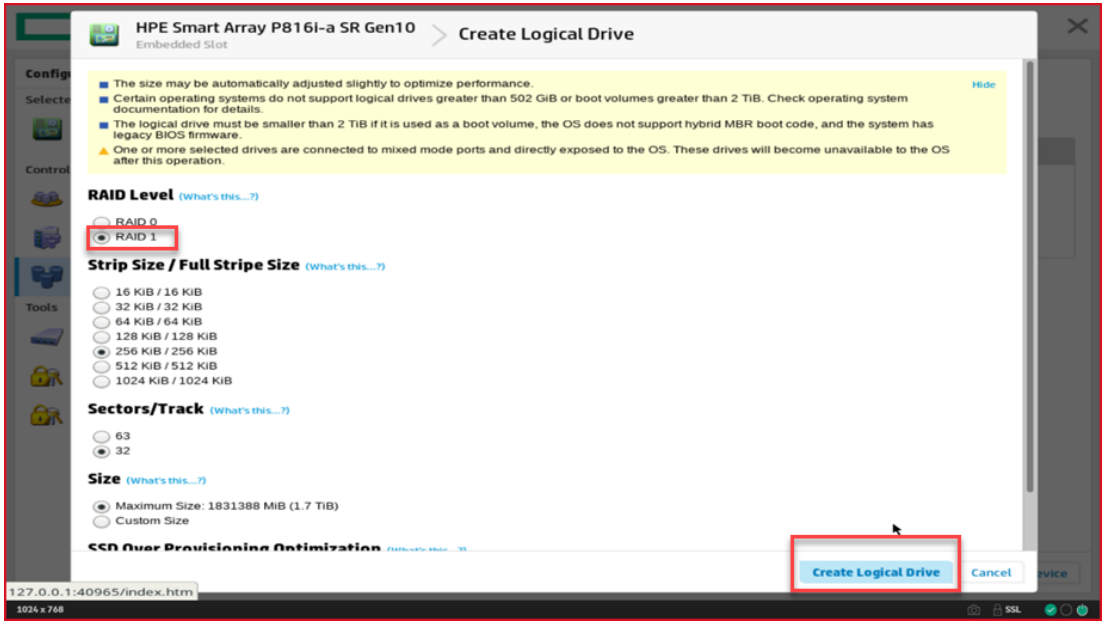

7 Click **Yes** to confirm.

8 Click **Finish**.

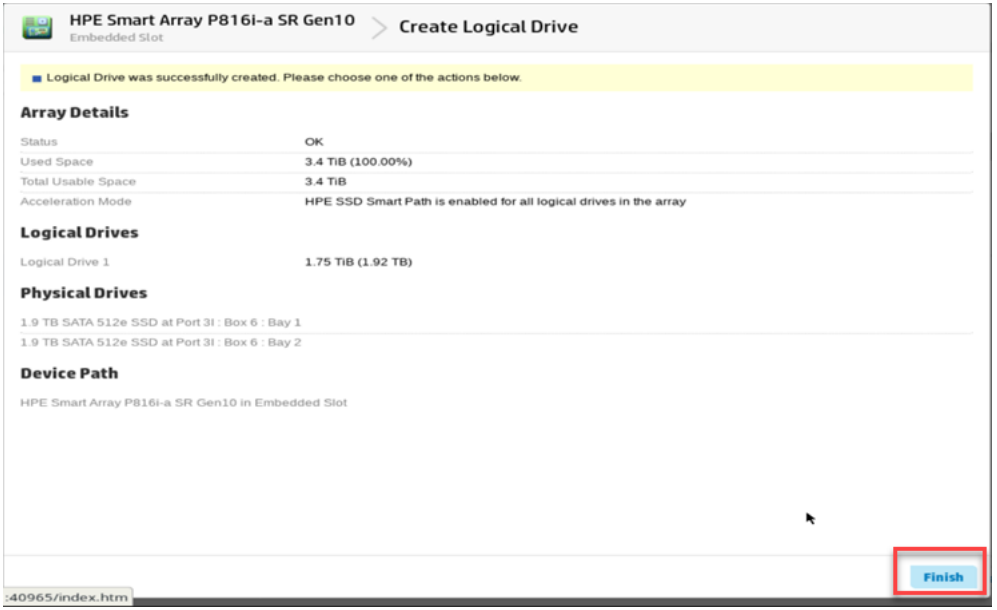

The logical drive is created.

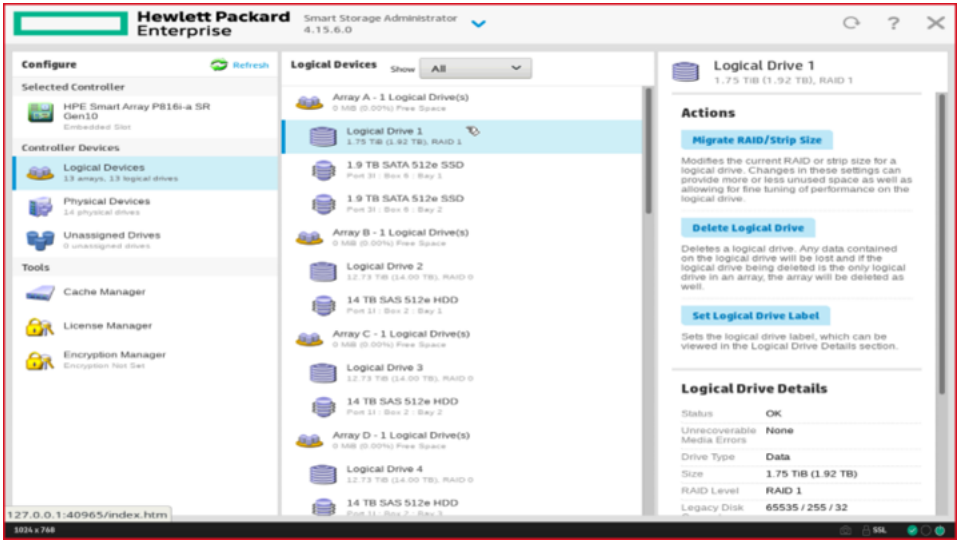

Replacing the node

To replace the node:

1 Check the node status in the NetBackup Flex Scale infrastructure management UI. The node is shown unhealthy and the node status is offline.

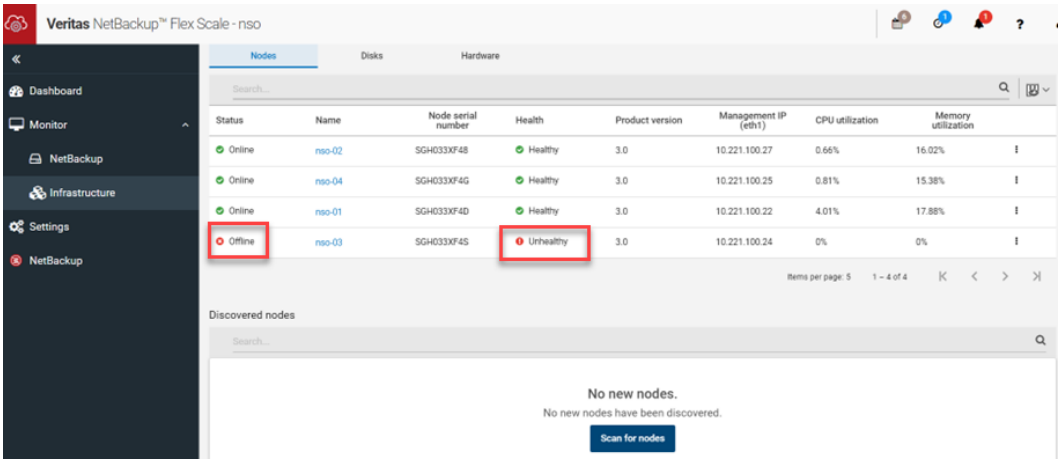

2 Prepare the node for replacement:

- Deploy ISO on the failed node where the OS disks were replaced.
- Perform factory reset on that node from the node CLI by using the `system factory-reset` command.

3 From the NetBackup Flex Scale infrastructure management UI perform the replace node operation.

- Scan for nodes. Click **Scan for nodes** to discover the nodes.
- Select a node from the displayed list and replace the node. For the unhealthy node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI, and click **Replace node**.

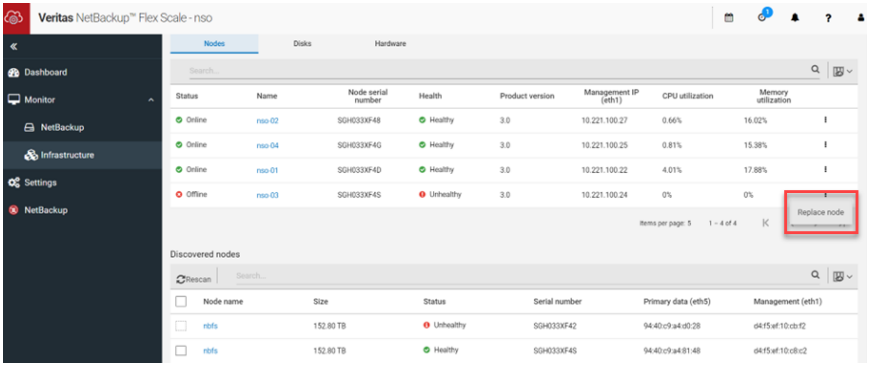

- In the Replace node dialog box, select the node that you want to use to replace the unhealthy node and click **Replace node**.

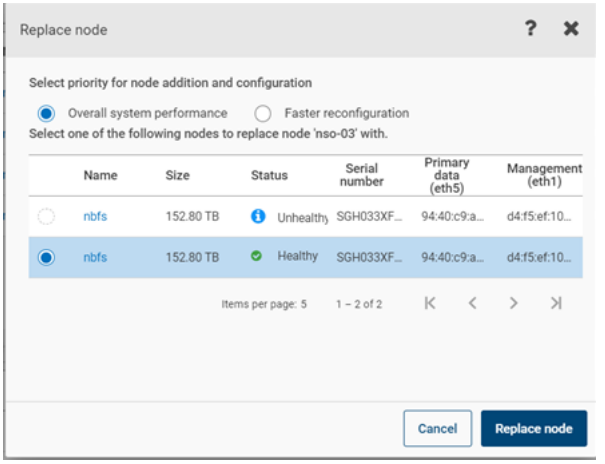
Replace node
Select priority for node addition and configuration
Overall system performance
Faster reconfiguration
Select one of the following nodes to replace node 'nso-03' with.
Name
Size
Status
Serial number
Primary data (eth5)
Management (eth1)
nbfs
152.80 TB
Unhealthy
SGH033XF...
94:40:c9:a...
d4:f5:ef:10...
nbfs
152.80 TB
Healthy
SGH033XF...
94:40:c9:a...
d4:f5:ef:10...
Items per page: 5
1 – 2 of 2
Cancel
Replace node

4 After the replace node operation completes successfully check the node status in the NetBackup Flex Scale infrastructure management UI. Navigate to **Monitor > Infrastructure > Nodes**. The node is online and shown healthy.

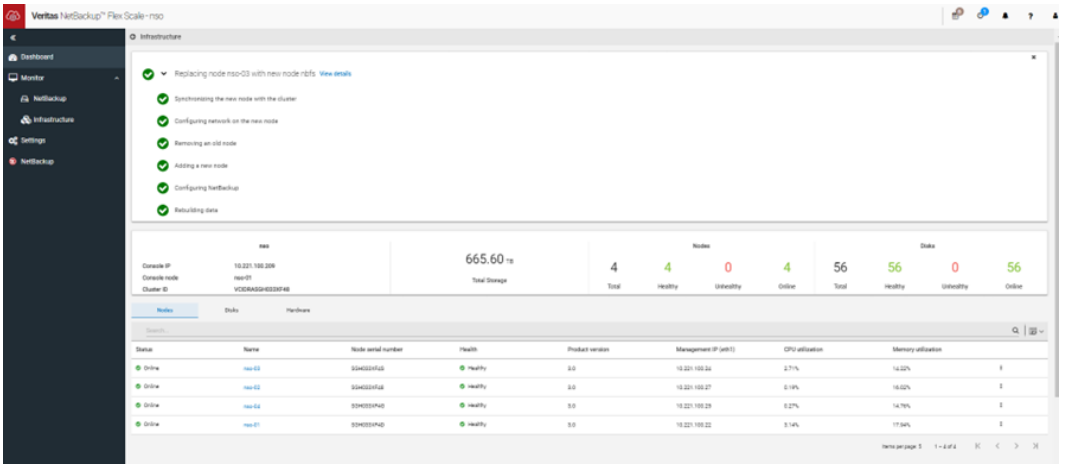

Click **Dashboard** and verify that all the NetBackup services are running.

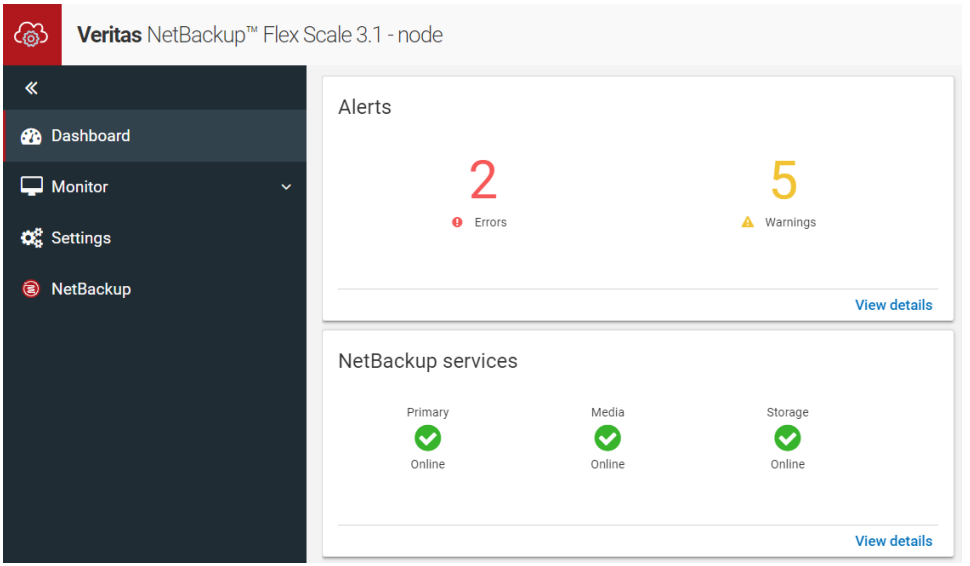

Replacement procedure for NVMe disks (SSDs)

NVMes are solid state drives (SSDs). The 7.68 TB NVMe RI drives are in slots 3 and 4 in the mid-bay.

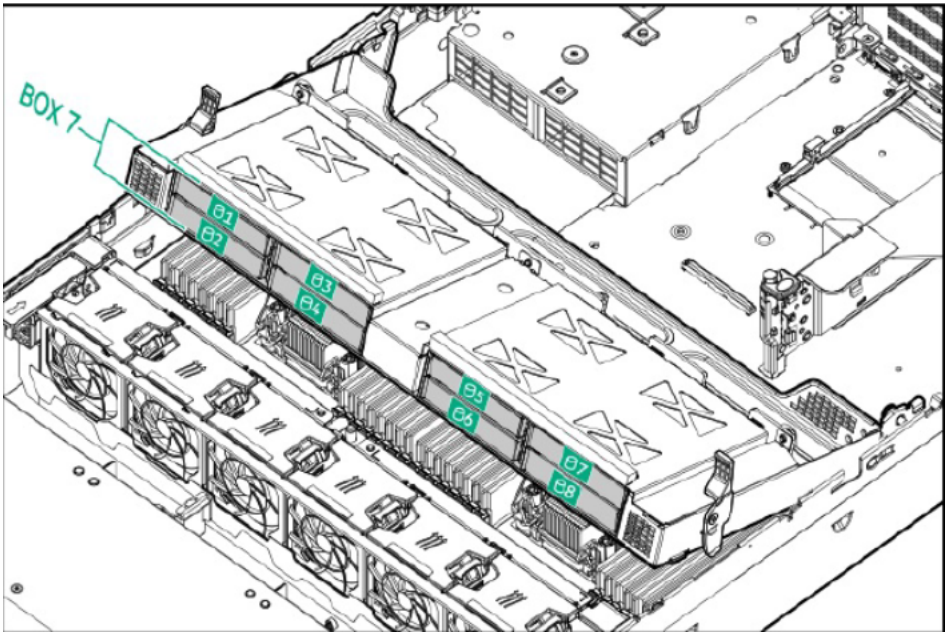

To view the serial number of the NVMe disks, log in to the iLO remote console and click **System Information > Storage**.

The following figures show the serial number for the two NVMe disks. Note the serial number of the NVMe disks. The serial number changes after the faulted disk is replaced.

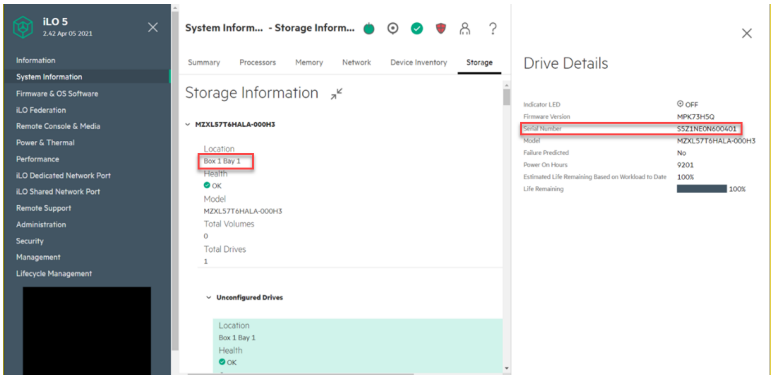

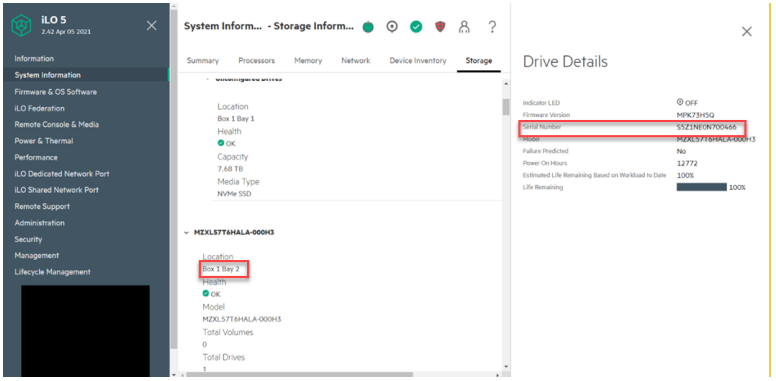

Identifying an NVMe disk failure (performed by the CHS team)

In the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure > Disks**. The disk status is shown **faulted**.

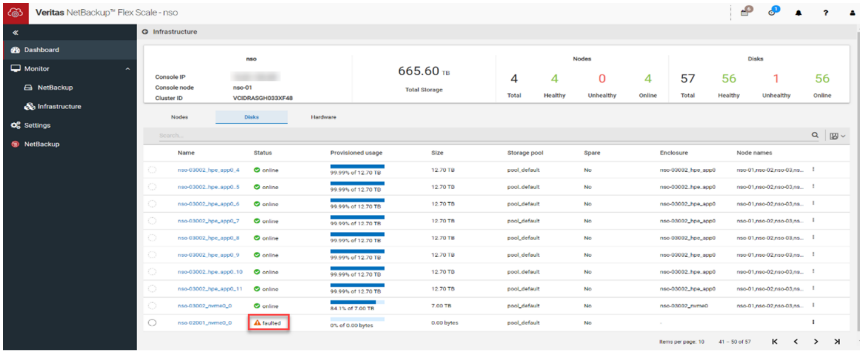

SSH to the the eth1 (management) IP address of the node and in the node-level CLI, run the following command to check the status of the NVMe disks:

```
show hardware-health node component=SSD
```

If there is some problem with the NVMes disks (7.68 TB capacity), in the **Status** column the disk status is shown **NOT-OK**.

Shutting down the node (performed by Veritas TSE)

Before an HPE representative can replace the NVMe disks, you must shut down the node.

To shut down the node:

1. Sign in to the NetBackup Flex Scale infrastructure management UI and navigate to **Monitor > Infrastructure > Nodes**.
2. On the node where the failure occurred, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Shutdown node**.

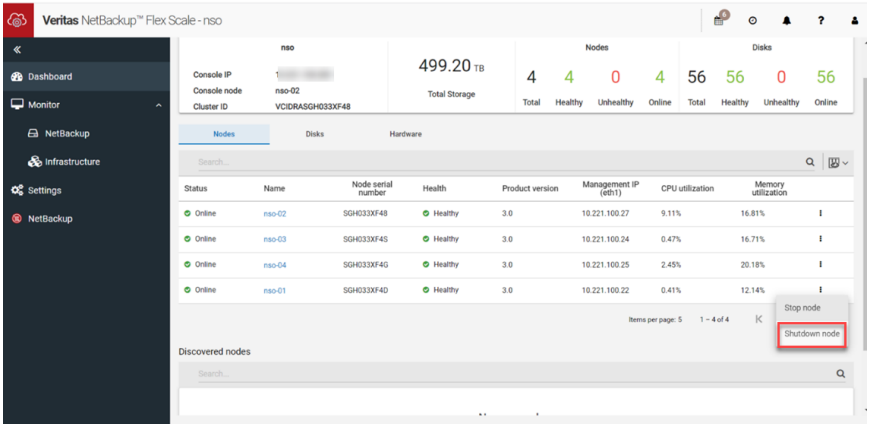

3 Confirm that the node is shut down successfully. In the UI, you can view the notification at the top of the page. In the iLO remote console, wait until the system shows the **System halted** message.

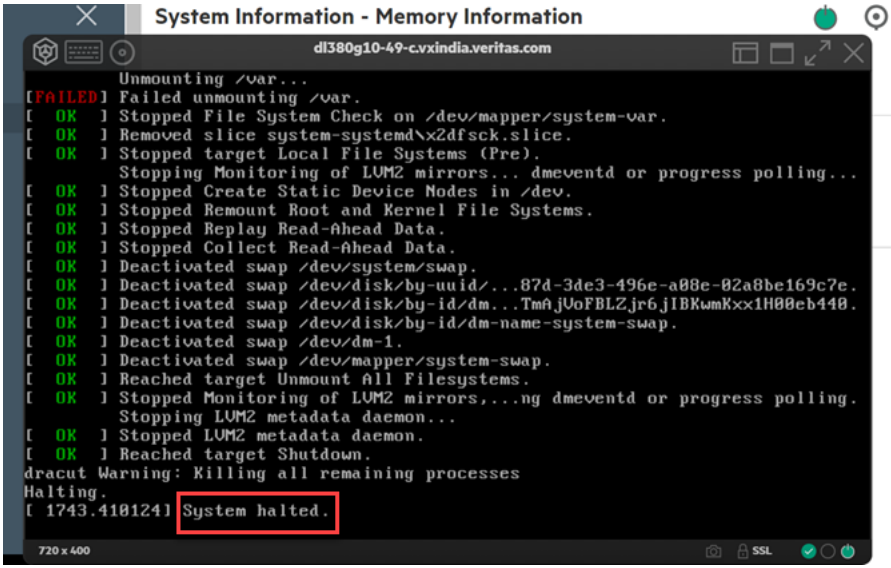

4 Shut down the node. Press the Power button on the front panel of the server or from the iLO remote console use the **Server Power > Press and Hold** option.

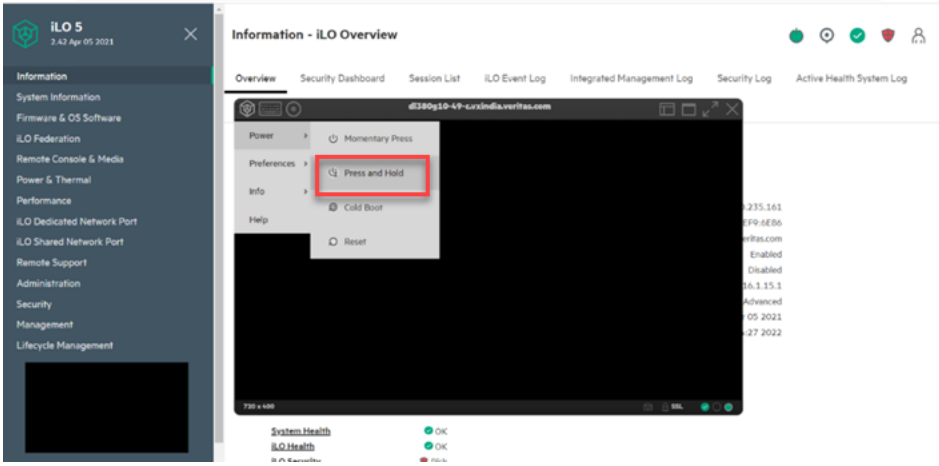

Replacing NVMe disks (performed by the HPE vendor)

The HPE representative replaces the faulted NVMe disk with another NVMe disk.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

To verify that the issue is resolved, complete the following steps:

1. Restart the node from the iLO remote console using the **Power > Momentary Press** option.
2. Confirm that the serial number of the NVMe disk is updated. After an NVMe disk is replaced, the serial number is different from the one that you noted earlier.

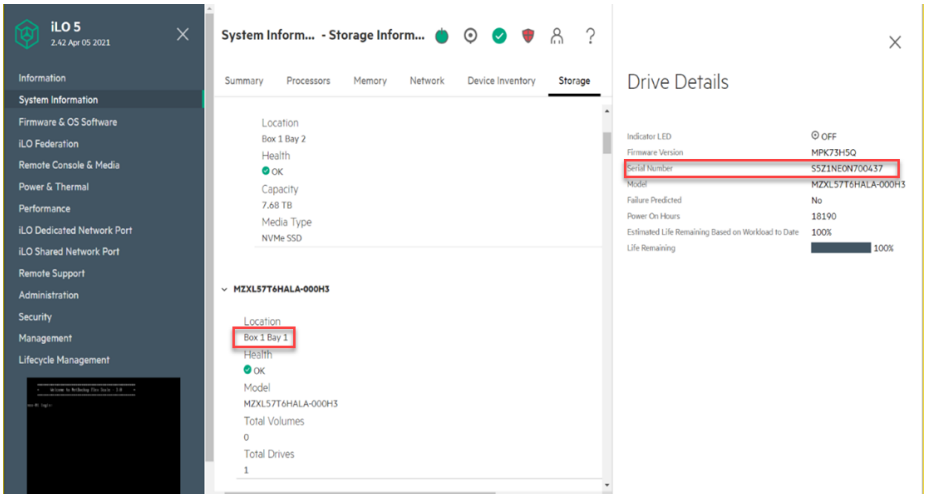

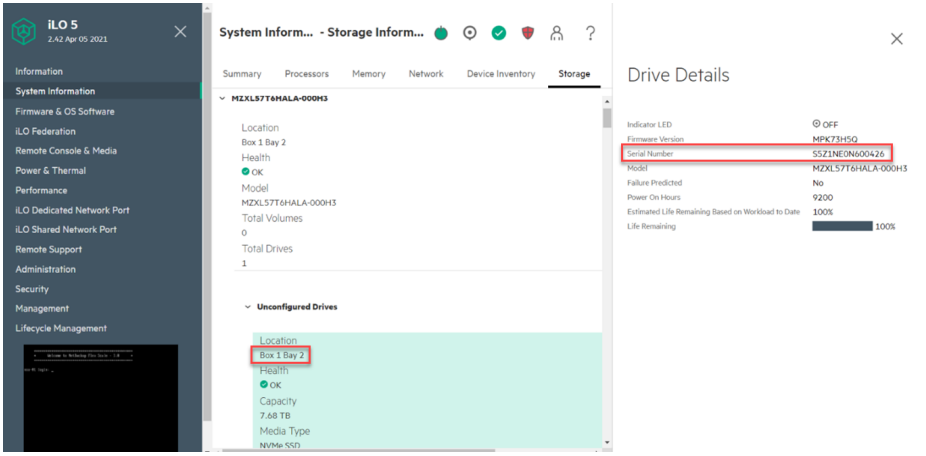

3 SSH to the eth1 (management) IP address of the node and run the `show hardware-health node component=SSD` command in the node-level CLI. In the Status column, the NVMe disk status is shown **OK**.

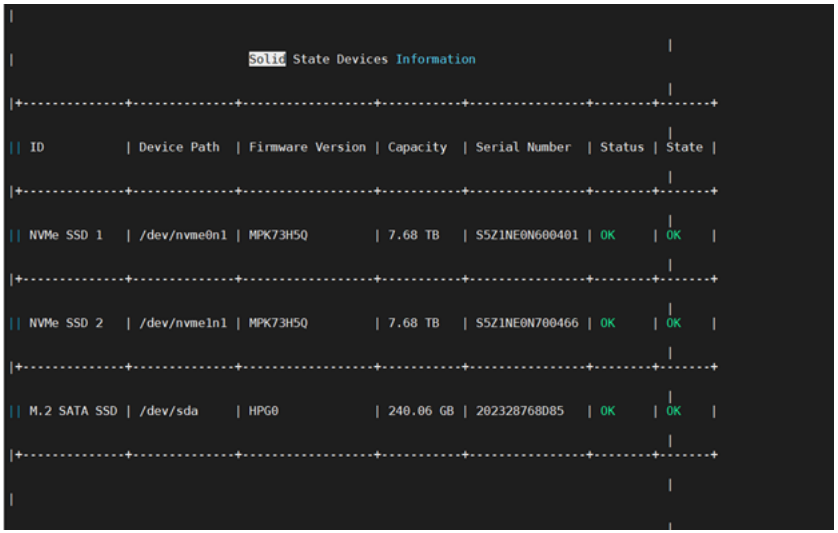

4 In the NetBackup Flex Scale infrastructure management UI, replace the faulty NVMe disk with a new NVMe disk. If both the NMVe disks are faulty, replace with new disks successively. Navigate to **Monitor > Infrastructure > Disks**. Click the faulty disk that you want to replace, and then click Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Replace disk**. The replacement disk is detected and added to the node.

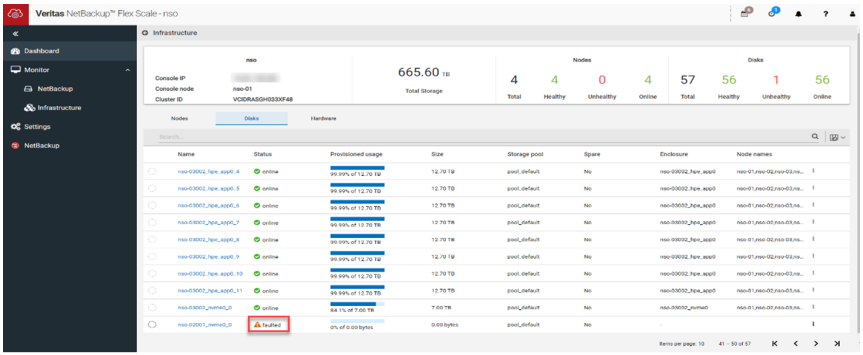

After the disk is replaced, data rebuild process begins and you can monitor the progress at the bottom of the screen when you click **Infrastructure > Monitor > Disks**. Time taken to rebuild the data depends on the amount of data.

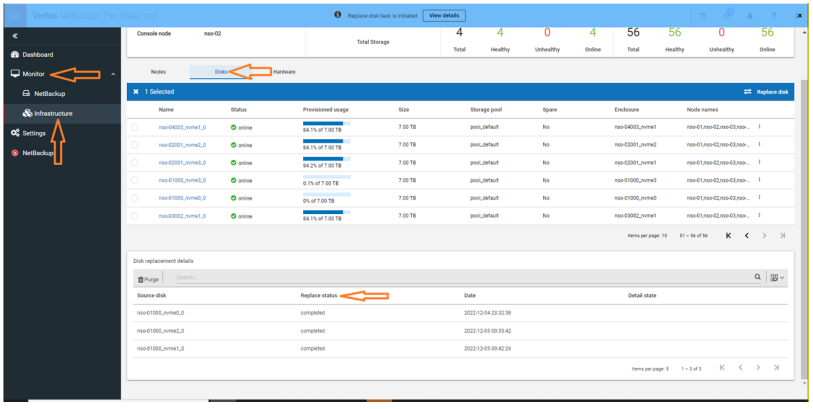

Replacement procedure for RAID controller

The following figures show the serial number for the RAID controller. Note the serial number of the RAID controller before replacement. The serial number changes after the RAID controller is replaced.

To view the serial number of the RAID controller, log in to the iLO remote console and click **System Information > Device Inventory**.

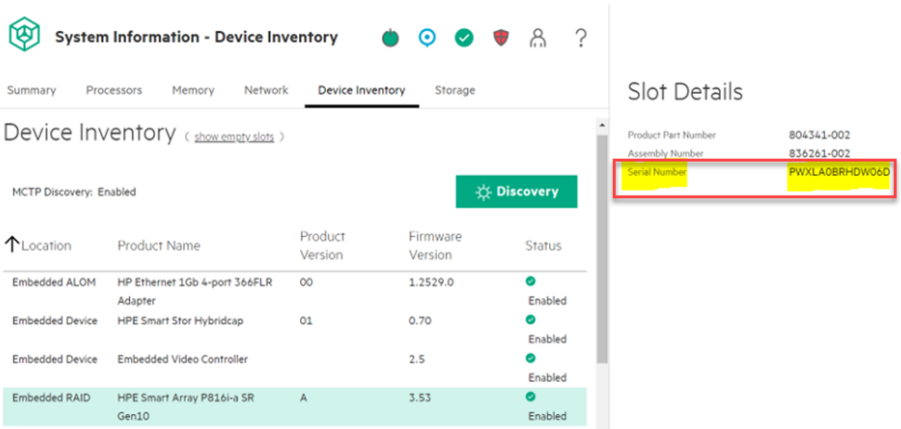

Shutting down the node (performed by Veritas TSE)

Before an HPE representative can replace the RAID controller, you must shut down the node.

To shut down the node:

1. Sign in to the NetBackup Flex Scale infrastructure management UI and navigate to **Monitor > Infrastructure > Nodes**.
2. On the node where the failure occurred, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Shutdown node**.

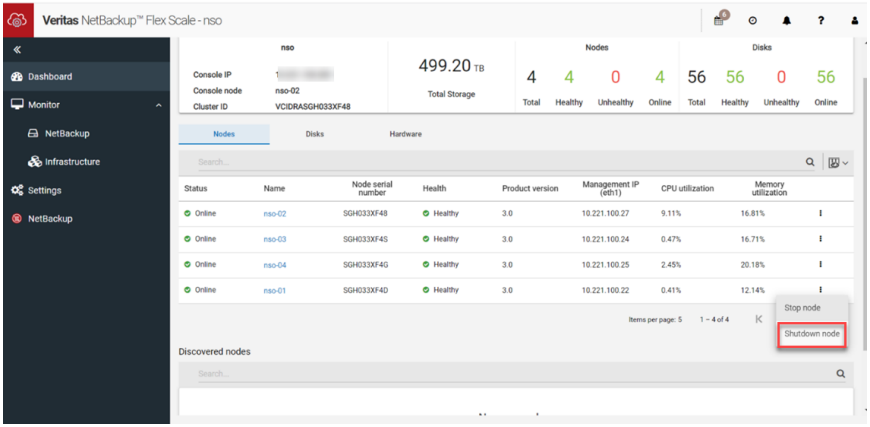

3 Confirm that the node is shut down successfully. In the UI, you can view the notification at the top of the page. In the iLO remote console, wait until the system shows the **System halted** message.

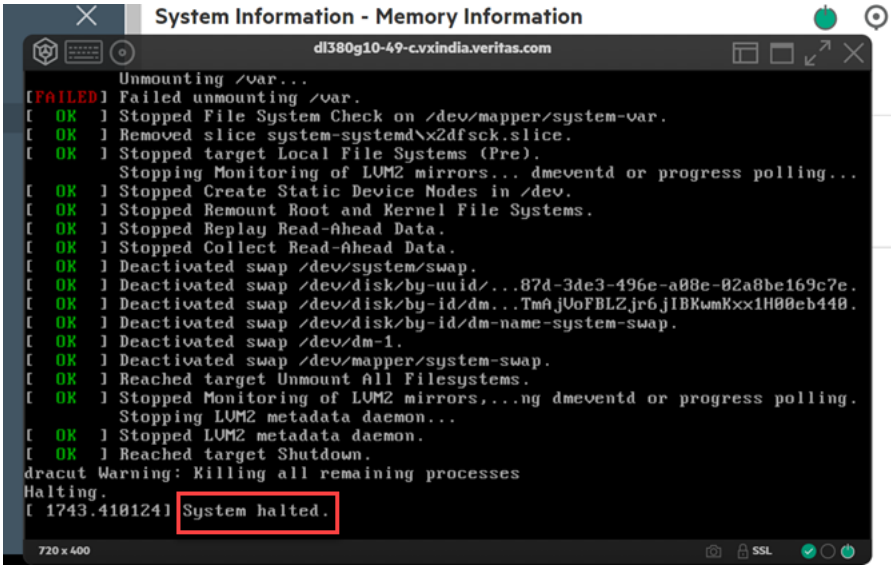

4 Shut down the node. Press the Power button on the front panel of the server or from the iLO remote console use the **Server Power > Press and Hold** option.

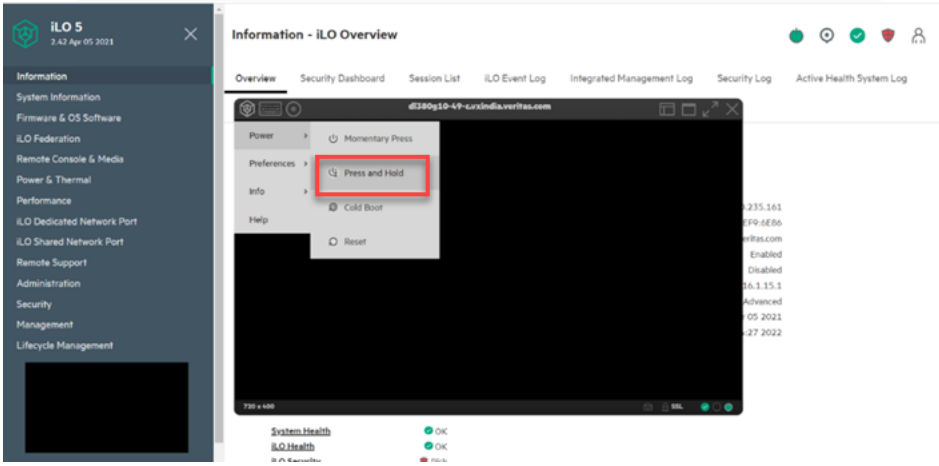

Replacing the RAID controller (performed by HPE)

The HPE representative replaces the RAID controller.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

To verify that the issue is resolved, complete the following steps:

1. Power on the server and wait till the node joins the cluster.
2. Verify the state and serial number of the RAID controller in the iLO remote console.

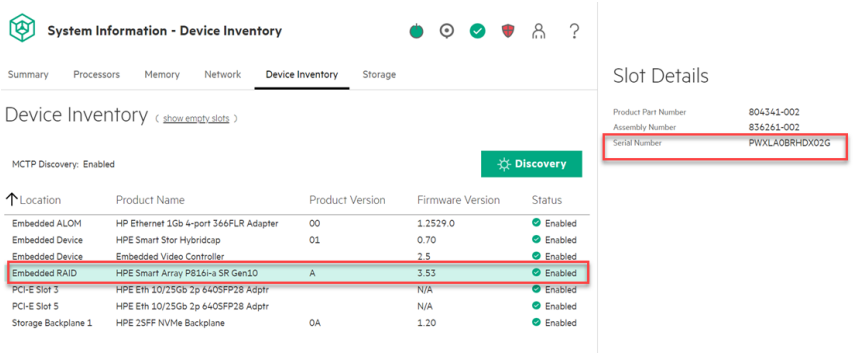

3. Verify that the node status in the NetBackup Flex Scale infrastructure management UI is shown healthy.

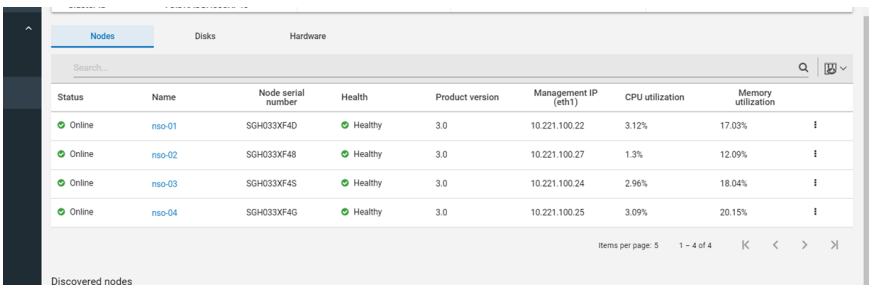

4 After the node joins the cluster, to avoid the data corruption replace that node. When the controller is replaced, the write data is lost as there is no way to move the Flash Backed Write Cache (FBWC) from one controller to another controller.

Shut down the node. Navigate to **Monitor > Infrastructure > Nodes**. Click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Shutdown node**.

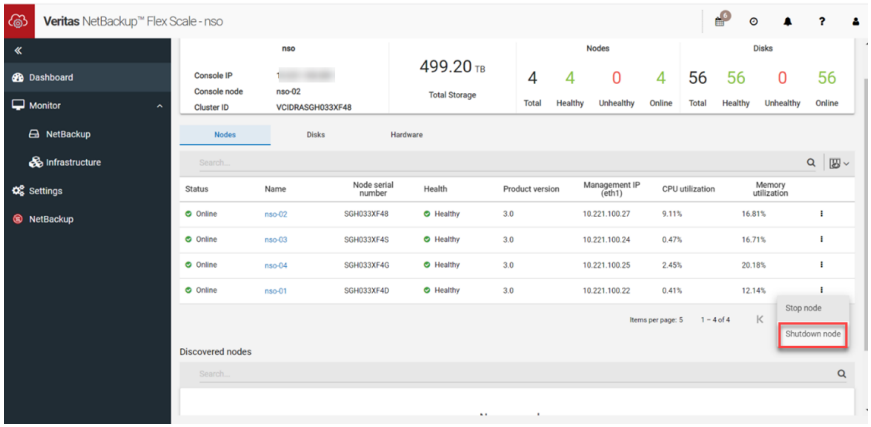

The node is shown unhealthy after you shut down the node.

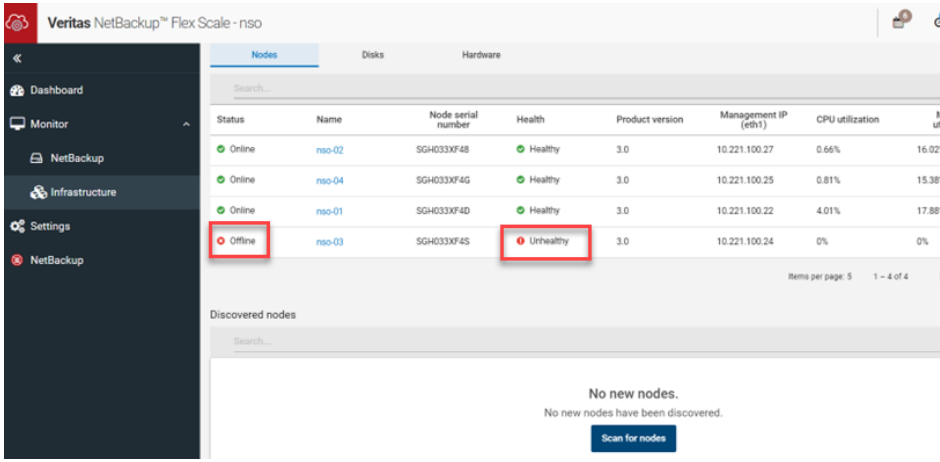

5 Confirm that the node is shut down successfully. In the UI, you can view the notification at the top of the page. In the iLO remote console, wait until the system shows the **System halted** message.

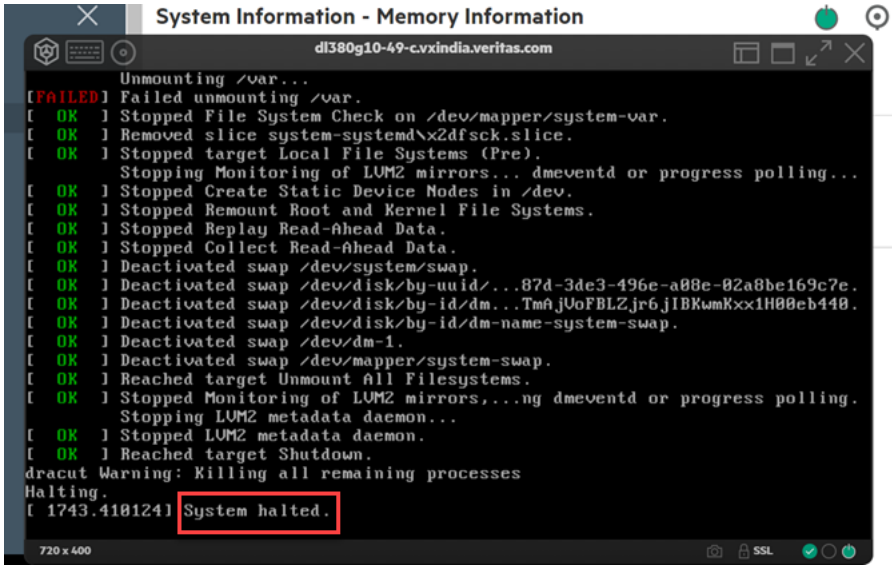

6 Shut down the node. Press the Power button on the front panel of the server or from the iLO remote console use the **Server Power > Press and Hold** option.

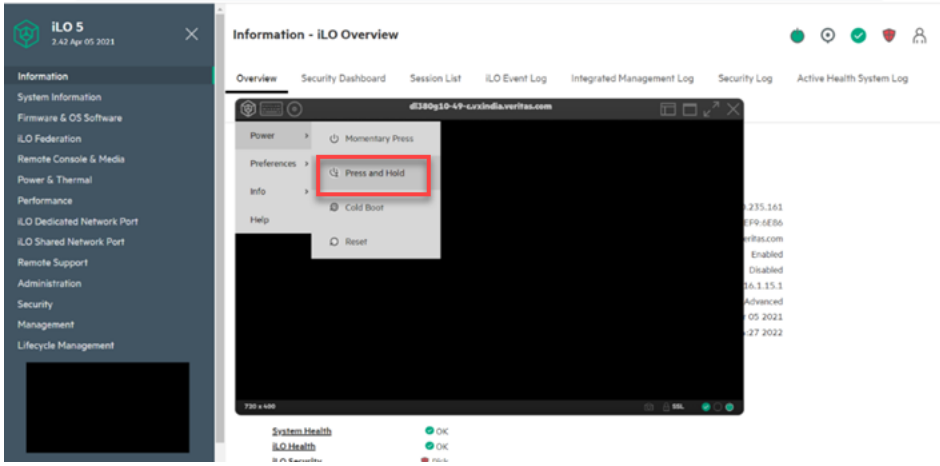

7 Prepare the node for replacement:

- Deploy the ISO on the node that was shut down.
- Perform factory reset on that node from the node-level CLI by using the `system factory-reset` command.

8 From the NetBackup Flex Scale infrastructure management UI perform the replace node operation.

- Scan for nodes. Click **Scan for nodes** to discover the nodes.
- Select a node from the displayed list and replace the node. For the unhealthy node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI, and click **Replace node**.

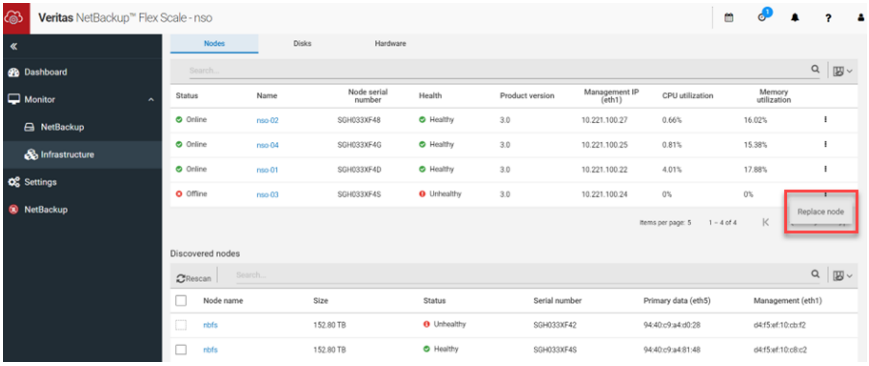

- In the Replace node dialog box, select the node that you want to use to replace the unhealthy node and click **Replace node**.

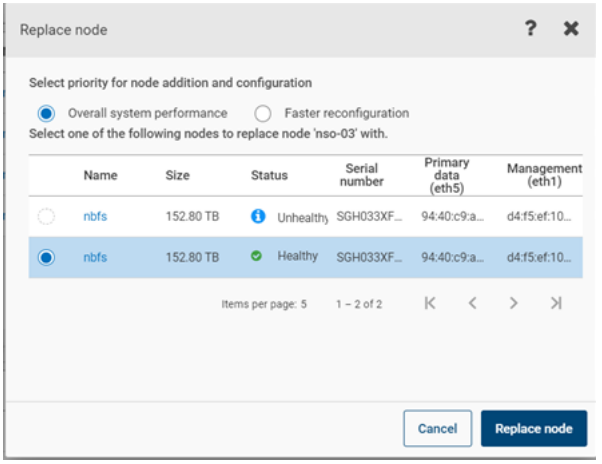
Replace node
Select priority for node addition and configuration
Overall system performance
Faster reconfiguration
Select one of the following nodes to replace node 'nso-03' with.
Name
Size
Status
Serial number
Primary data (eth5)
Management (eth1)
nbfs
152.80 TB
Unhealthy
SGH033XF...
94:40:c9:a...
d4:f5:ef:10...
nbfs
152.80 TB
Healthy
SGH033XF...
94:40:c9:a...
d4:f5:ef:10...
Items per page: 5
1 – 2 of 2
Cancel
Replace node

9 After the replace node operation completes successfully check the node status in the NetBackup Flex Scale infrastructure management UI. Navigate to **Monitor > Infrastructure > Nodes**. The node is online and shown healthy

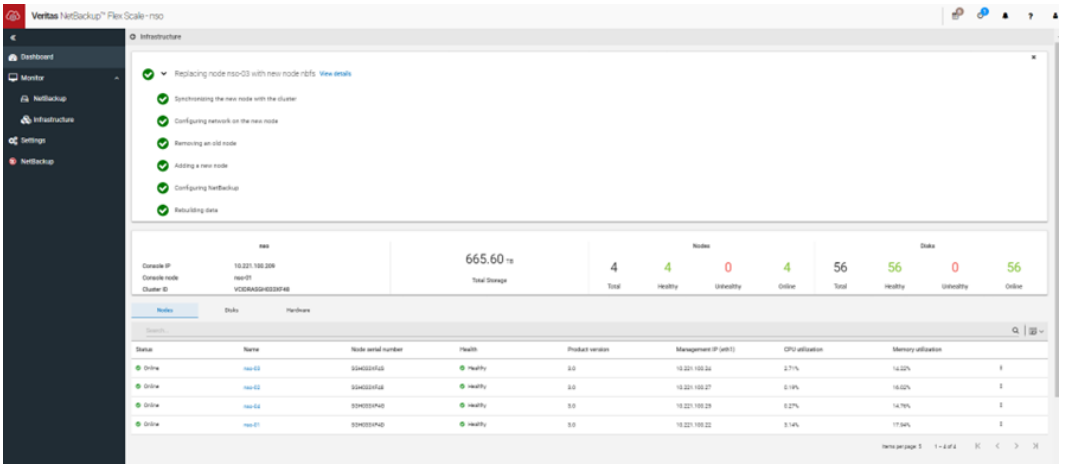

Click **Dashboard** and verify that all the NetBackup services are running.

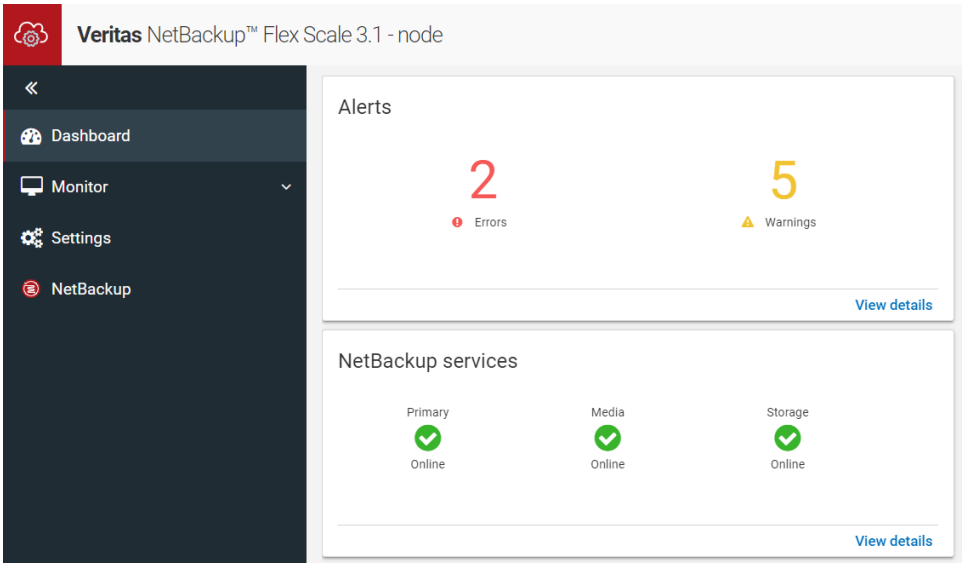

Replacement procedure for an Integrated Lights-Out (iLO) port

This topic describes the process of replacing the iLO port on an HPE server node. Each node has one iLO port.

Identifying an iLO port failure (performed by the customer)

The iLO port is configured with a default user name, password, and DNS name. The default information is on the serial label pull tab attached to the server. Use these values to access iLO remotely by using a web browser.

Verify if the iLO UI is accessible using the iLO host name. The following error is shown if the UI is not accessible.

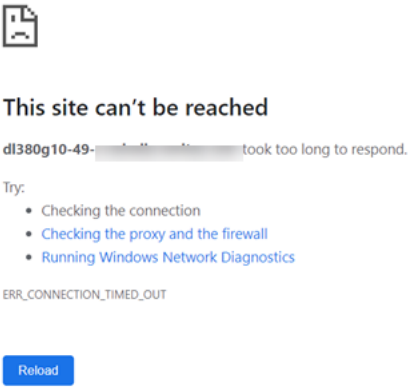

Verify if the IP allocated to the iLO interface is reachable. The following error is shown if the IP is unreachable.

```
nso-01:~ # nslookup dl380g10-49-c
Server:         172.16.8.33
Address:        172.16.8.33#53

Name:   dl380g10-49-c.
Address: 10.210.235.161

nso-01:~ # ping 10.210.235.161
PING 10.210.235.161 (10.210.235.161) 56(84) bytes of data.

--- 10.210.235.161 ping statistics ---
46 packets transmitted, 0 received, 100% packet loss, time 45003ms
```

If you encounter these issues, there might be a problem with the cables connection or the iLO port. To isolate the issue, first tighten the cable connections and try accessing the iLO UI and the IP. If you face the same issue, replace the cables and try again. If the issue persists it implies that the iLO port is faulty. Contact Veritas TSE to replace the node with the faulty iLO port.

Replacing the node with the faulty iLO port (performed by Veritas TSE)

To replace the node:

1 Check the node status in the NetBackup Flex Scale infrastructure management UI. The node is shown unhealthy and the node status is offline.

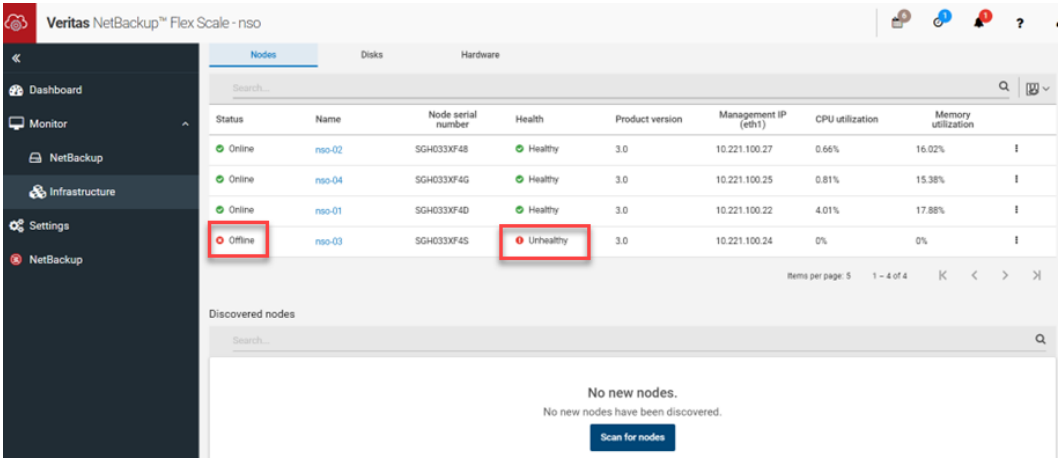

2 Prepare the node for replacement:

- Deploy ISO on the failed node where the OS disks were replaced.
- Perform factory reset on that node from the node CLI by using the `system factory-reset` command.

3 From the NetBackup Flex Scale infrastructure management UI perform the replace node operation.

- Scan for nodes. Click **Scan for nodes** to discover the nodes.
- Select a node from the displayed list and replace the node. For the unhealthy node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI, and click **Replace node**.

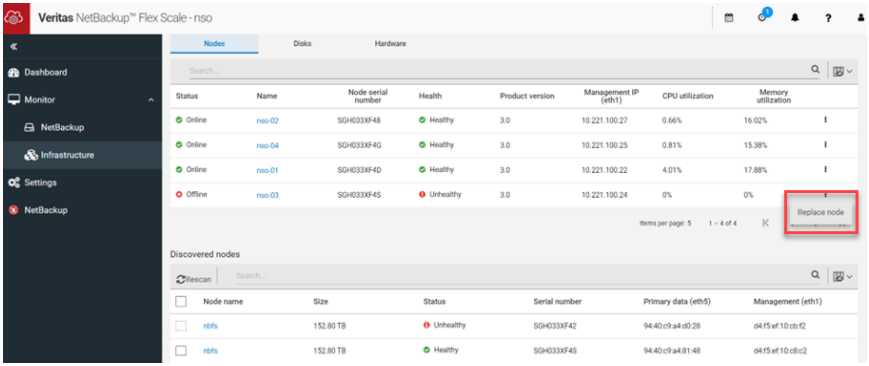

- In the Replace node dialog box, select the node that you want to use to replace the unhealthy node and click **Replace node**.

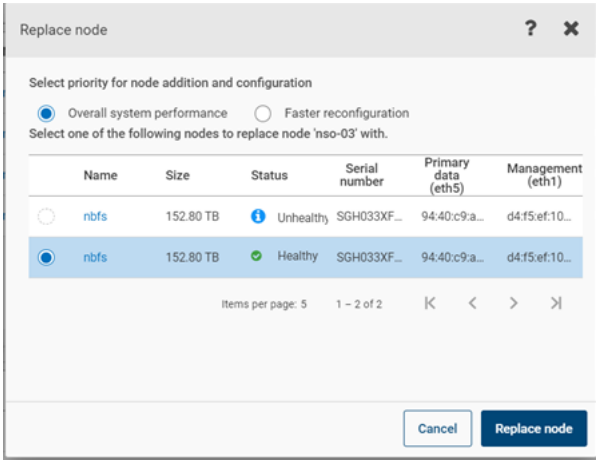
Replace node
Select priority for node addition and configuration
Overall system performance
Faster reconfiguration
Select one of the following nodes to replace node 'nso-03' with.
Name
Size
Status
Serial number
Primary data (eth5)
Management (eth1)
nbfs
152.80 TB
Unhealthy
SGH033XF...
94:40:c9:a...
d4:f5:ef:10...
nbfs
152.80 TB
Healthy
SGH033XF...
94:40:c9:a...
d4:f5:ef:10...
Items per page: 5
1 – 2 of 2
Cancel
Replace node

4 After the replace node operation completes successfully check the node status in the NetBackup Flex Scale infrastructure management UI. Navigate to **Monitor > Infrastructure > Nodes**. The node is online and shown healthy.

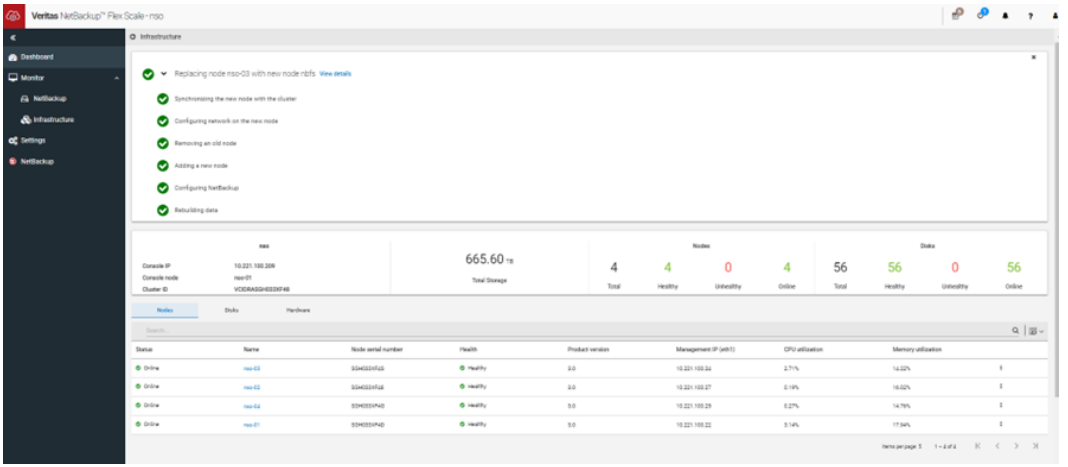

Click **Dashboard** and verify that all the NetBackup services are running.

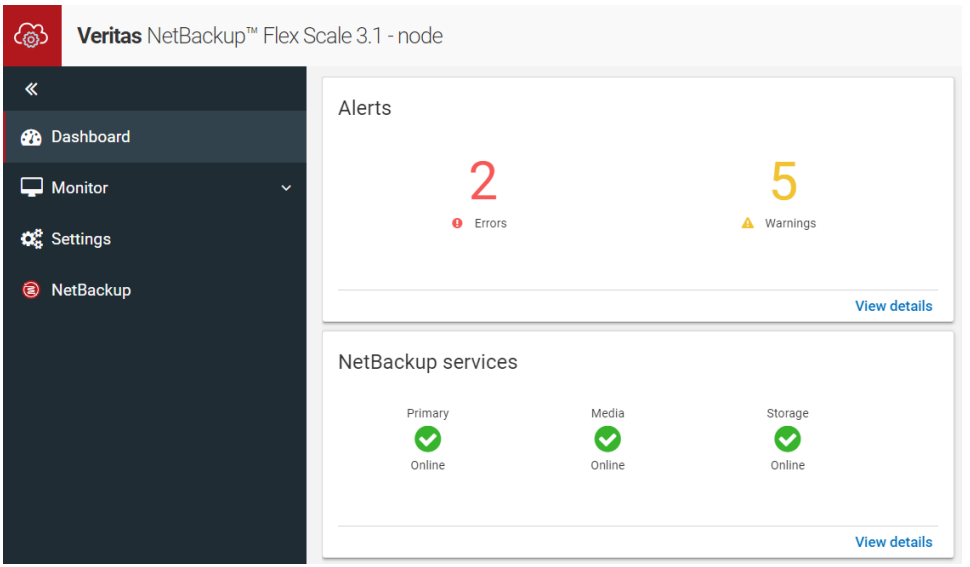

Completing the post-replacement tasks (performed by Veritas TSE)

After you replace the node, verify that the issue is resolved.

Plug in the cable in the iLO port and verify if the iLO UI is accessible using the iLO host name.

Verify that the IP address allocated to the iLO interface is reachable.

```
nso-01:~ # ping 10.210.235.161
PING 10.210.235.161 (10.210.235.161) 56(84) bytes of data.
64 bytes from 10.210.235.161: icmp_seq=305 ttl=254 time=0.209 ms
64 bytes from 10.210.235.161: icmp_seq=306 ttl=254 time=0.138 ms
64 bytes from 10.210.235.161: icmp_seq=307 ttl=254 time=0.155 ms
64 bytes from 10.210.235.161: icmp_seq=308 ttl=254 time=0.240 ms
64 bytes from 10.210.235.161: icmp_seq=309 ttl=254 time=0.173 ms
64 bytes from 10.210.235.161: icmp_seq=310 ttl=254 time=0.199 ms
64 bytes from 10.210.235.161: icmp_seq=311 ttl=254 time=0.171 ms
64 bytes from 10.210.235.161: icmp_seq=312 ttl=254 time=0.186 ms
64 bytes from 10.210.235.161: icmp_seq=313 ttl=254 time=0.261 ms
64 bytes from 10.210.235.161: icmp_seq=314 ttl=254 time=0.183 ms
64 bytes from 10.210.235.161: icmp_seq=315 ttl=254 time=0.220 ms
64 bytes from 10.210.235.161: icmp_seq=316 ttl=254 time=0.213 ms
64 bytes from 10.210.235.161: icmp_seq=317 ttl=254 time=0.146 ms
64 bytes from 10.210.235.161: icmp_seq=318 ttl=254 time=0.196 ms
64 bytes from 10.210.235.161: icmp_seq=319 ttl=254 time=0.180 ms
64 bytes from 10.210.235.161: icmp_seq=320 ttl=254 time=0.138 ms
64 bytes from 10.210.235.161: icmp_seq=321 ttl=254 time=0.124 ms
64 bytes from 10.210.235.161: icmp_seq=322 ttl=254 time=0.132 ms
64 bytes from 10.210.235.161: icmp_seq=323 ttl=254 time=0.156 ms

--- 10.210.235.161 ping statistics ---
323 packets transmitted, 19 received, 94% packet loss, time 322010ms
rtt min/avg/max/mdev = 0.124/0.180/0.261/0.036 ms
nso-01:~ #
```

Replacement procedure for quad-port NIC

This topic describes the process of replacing the NIC quad port on an HPE server node.

Identifying a quad-port NIC failure (performed by the customer)

Verify if the cables are unplugged or faulty. After removing cable from quad-port NIC node state in the UI shows Unhealthy.

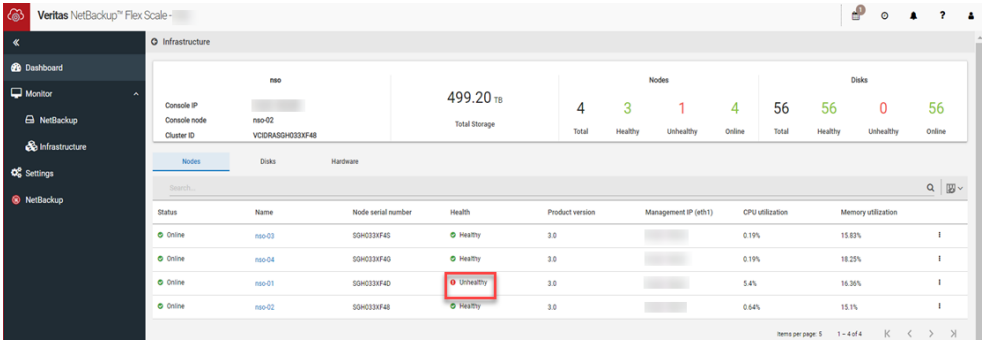

The eth1 network interface state is shown unplugged in the UI. To view the link status, navigate to **Infrastructure > Hardware > Network Card** in the NetBackup Flex Scale infrastructure UI.

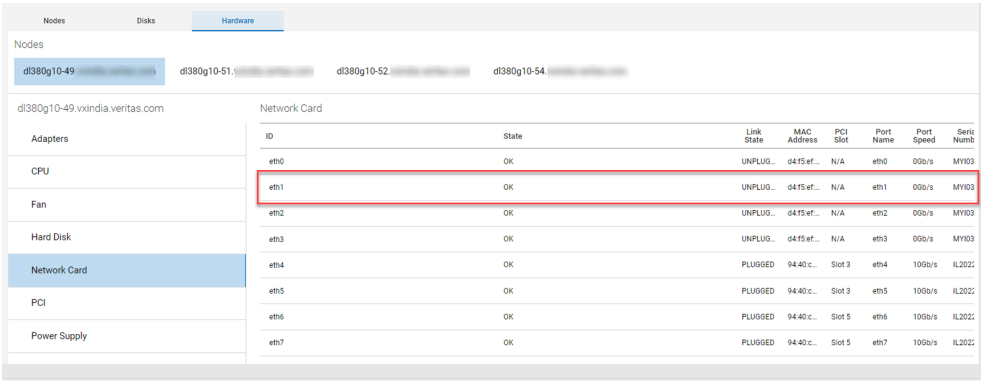

The network status for the public network is shown unhealthy. Navigate to the dashboard page of the UI:

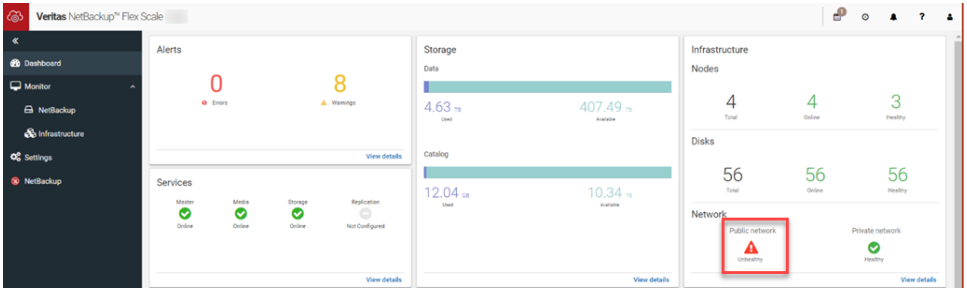

In the iLO remote console, adapter1 slot 3 status is shown as link down.

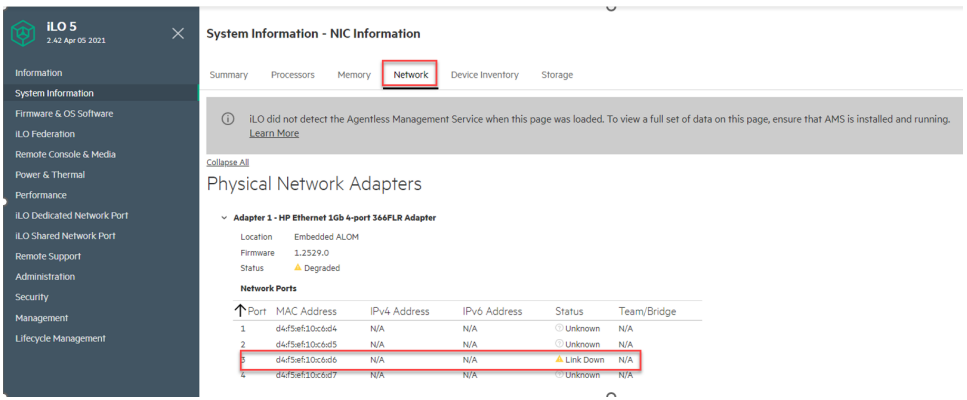

An event notifying that the eth1 network interface is in faulted state is shown on the **Settings > Events** page of the UI:

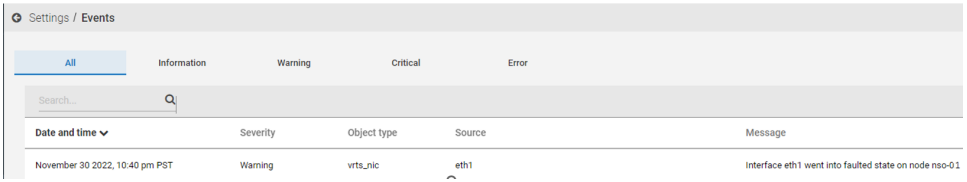

If you encounter these issues, there might be a problem with the cables connection or the quad-port. To isolate the issue, first tighten the cable connections. If you face the same issue, replace the cables and try again. If the issue persists it implies that the port is faulty. Contact Veritas TSE to replace the faulty quad-port NIC.

Note the MAC address of the network interface. After the quad-port NIC is replaced, the MAC address will change:

```
nso-01:~ # ifconfig eth1
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 10.221.100.22  netmask 255.255.240.0  broadcast 10.221.111.255
        inet6 fe80::d6f5:efff:fe10:c6d6  prefixlen 64  scopeid 0x20<link>
        ether d4:f5:ef:10:c6:d6  txqueuelen 1000  (Ethernet)
        RX packets 8047487  bytes 746965790 (712.3 MiB)
        RX errors 0  dropped 415994  overruns 0  frame 0
        TX packets 93640  bytes 99023069 (94.4 MiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
        device memory 0xe1500000-e15fffff
```

Shutting down the node (performed by Veritas TSE)

Before an HPE representative can replace the quad-port NIC, you must shut down the node.

To shut down the node:

1. Sign in to the NetBackup Flex Scale infrastructure management UI and navigate to **Monitor > Infrastructure > Nodes**.
2. On the node where the failure occurred, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Shutdown node**.

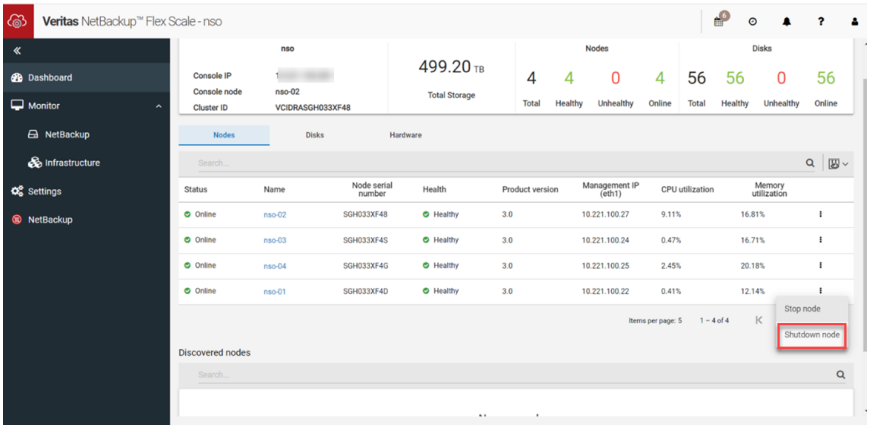

Replacing the quad-port NIC (performed by the HPE vendor)

The HPE representative replaces the quad-port NIC.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

To verify that the issue is resolved, complete the following steps:

1 Restart the node from the iLO remote console using the **Power > Momentary Press** option.

The green color power symbol indicates that the node has started.

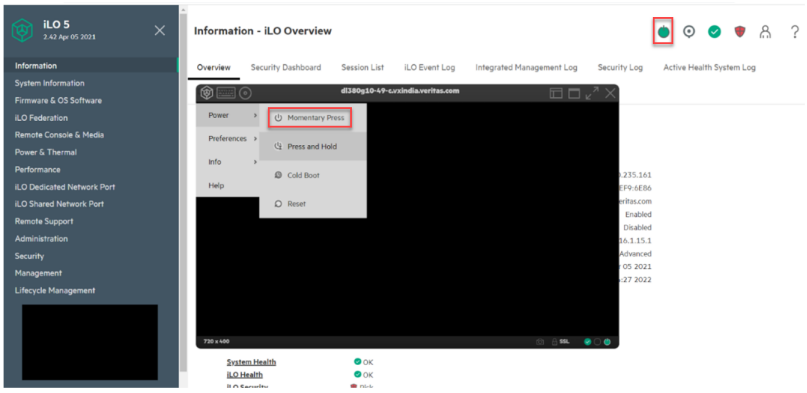

2 As the port is changed physically, the MAC address is also changed, which results in the following cases:

The node status is shown unhealthy in the UI:

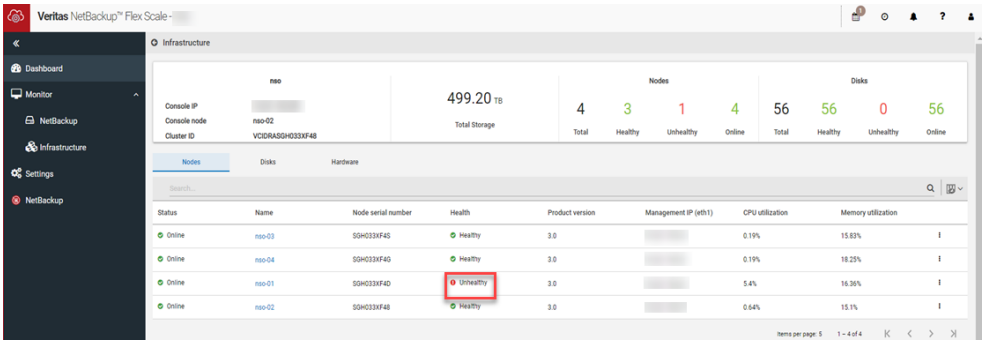

The eth1 network interface state is shown unplugged in the UI. To view the link status, navigate to **Infrastructure > Hardware > Network Card** in the NetBackup Flex Scale infrastructure UI.

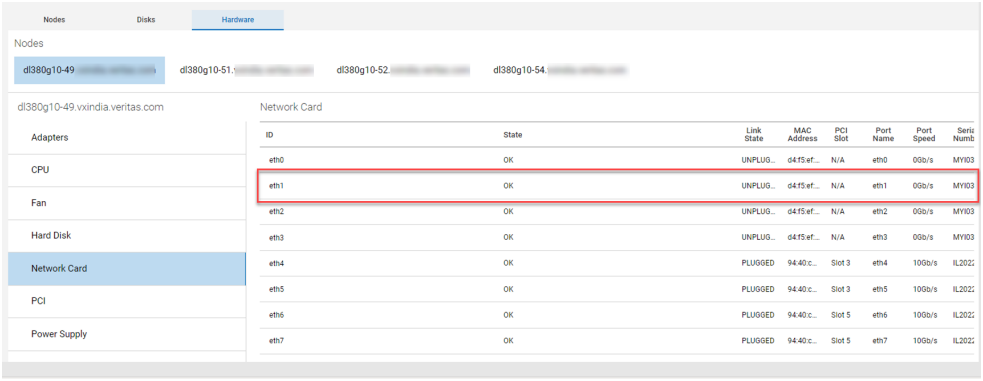

In the iLO remote console, adapter1 slot 3 status is shown as link down:

An event notifying that the eth1 network interface is in faulted state is shown on the **Settings > Events** page of the UI:

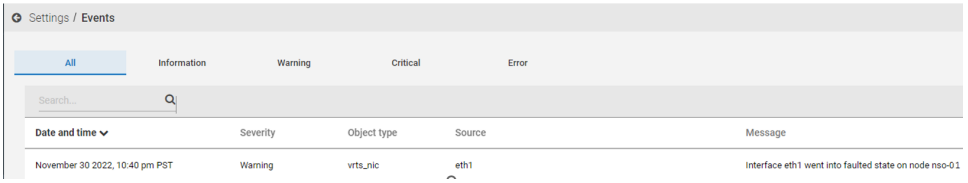

3 Verify that the node status is shown healthy in the UI:

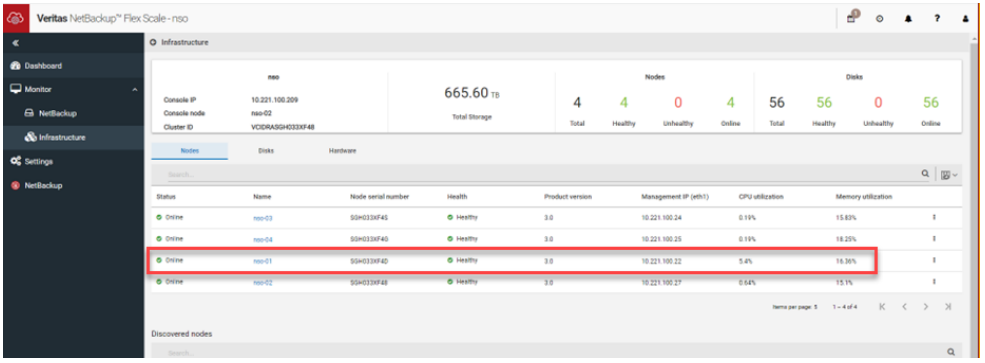

4 Verify that the eth1 link state has changed to **Plugged**:

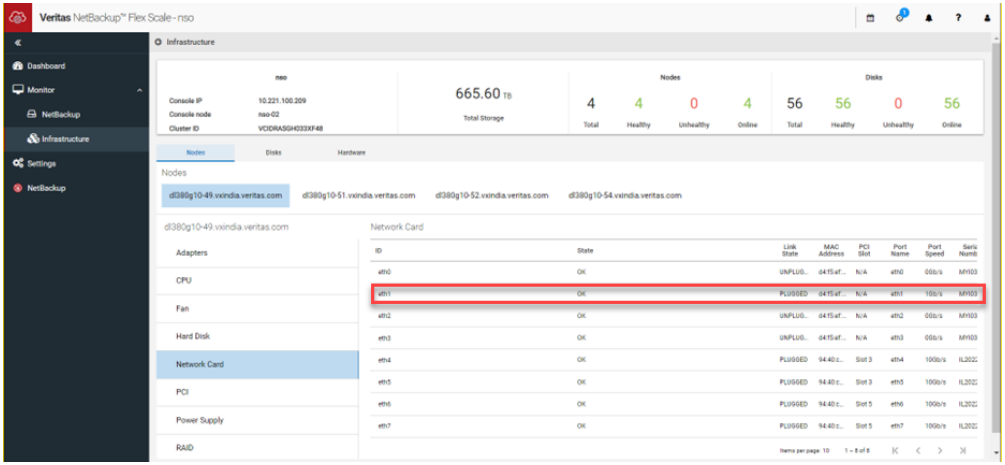

5 The public network status is shown healthy:

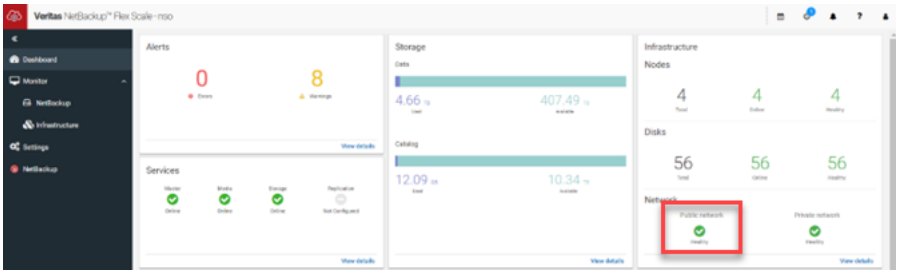

6 Verify that Adapter 1 slot 3 link status is shown **OK**:

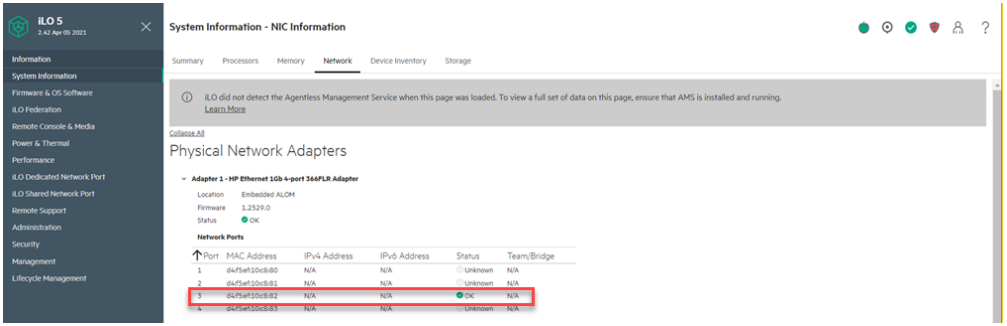

7 Verify that the changed MAC ID can be seen in the `system hardware-health` and in `ifconfig eth1` output:

```
                                    Network Card Information
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
| ID   | PCI Slot | Card Model | Serial Number | Port Mode   | Port Speed | MAC Address       | Link State | State |
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
| eth0 | N/A      | 665238-001 | MYI0300G6B    | Twisted Pair | 0Gb/s      | d4:f5:ef:10:c8:83 | UNPLUGGED  | OK    |
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
| eth1 | N/A      | 665238-001 | MYI0300G6B    | Twisted Pair | 1Gb/s      | d4:f5:ef:10:c8:82 | PLUGGED    | OK    |
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
| eth2 | N/A      | 665238-001 | MYI0300G6B    | Twisted Pair | 0Gb/s      | d4:f5:ef:10:c8:81 | UNPLUGGED  | OK    |
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
| eth3 | N/A      | 665238-001 | MYI0300G6B    | Twisted Pair | 0Gb/s      | d4:f5:ef:10:c8:80 | UNPLUGGED  | OK    |
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
| eth4 | Slot 3   | 817751-001 | IL202202L6    | FIBRE        | 10Gb/s     | 94:40:c9:a3:3f:d9 | PLUGGED    | OK    |
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
| eth5 | Slot 3   | 817751-001 | IL202202L6    | FIBRE        | 10Gb/s     | 94:40:c9:a3:3f:d8 | PLUGGED    | OK    |
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
| eth6 | Slot 5   | 817751-001 | IL202202YX    | FIBRE        | 10Gb/s     | 94:40:c9:a4:a0:89 | PLUGGED    | OK    |
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
| eth7 | Slot 5   | 817751-001 | IL202202YX    | FIBRE        | 10Gb/s     | 94:40:c9:a4:a0:88 | PLUGGED    | OK    |
+------+----------+------------+---------------+-------------+------------+-------------------+------------+-------+
```

Procedure for memory expansion (DIMMs)

This topic describes how to expand the memory by installing additional DIMMs.

Check the total memory before adding DIMMs to the node. You can check the total memory of the node by using the `support shell` command from the node-level CLI, and then using any of the available system commands:

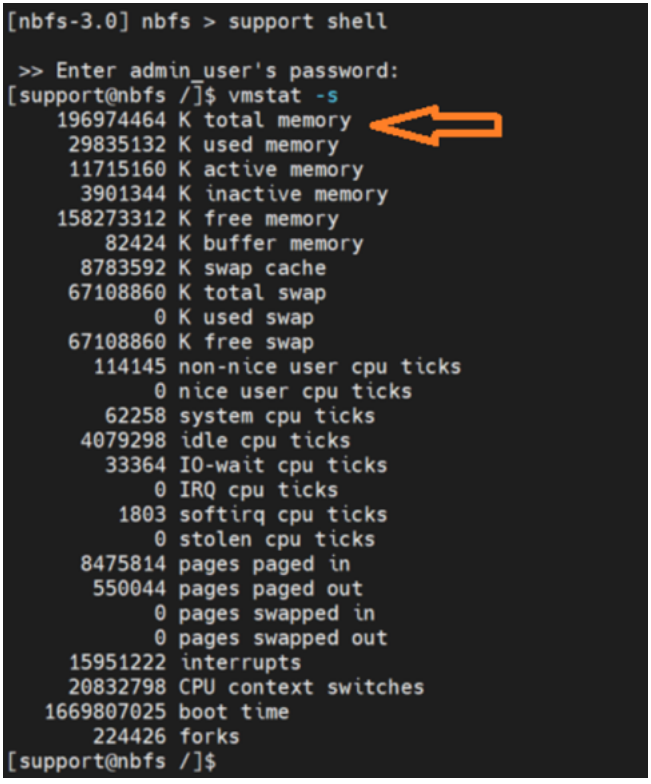

Before you add the DIMMs, in the ILO remote console, view the slots where DIMMs are not present:

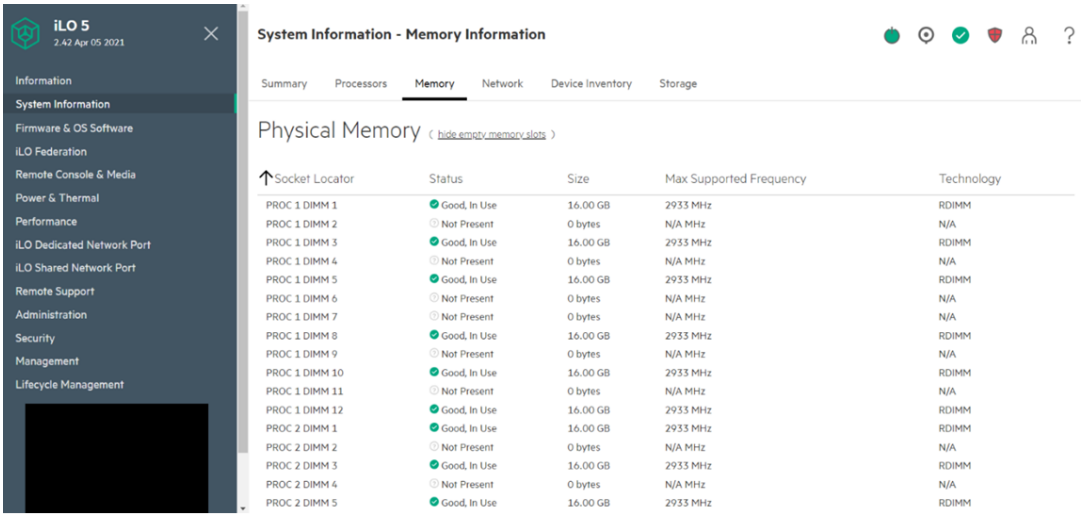

Locating the management console IP address

To identify the management console node, in the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure**. On the **Infrastructure** page, **Console IP** shows the management console IP address.

Shutting down the cluster (performed by Veritas TSE)

To shut down the cluster:

1. SSH to the management console IP.
2. Run the `cluster shutdown all` command and wait till the cluster shuts down.

```
nso>
nso> cluster shutdown all
 12% [/] Stopping cluster processes on all
NetBackup Flex Scale cluster INFO V-493-10-125 SSH sessions to all nodes may terminate.
```

3 In the iLO remote console, wait until the system shows the **System halted** message. for each node in the cluster.

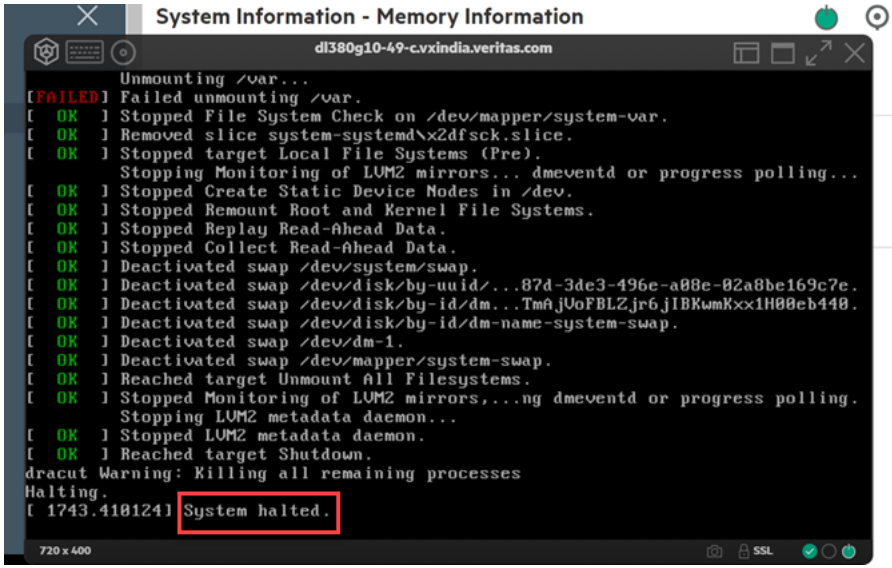

4 Press the Power button on the front panel of the server or from the iLO remote console use the **Server Power > Press and Hold** option.

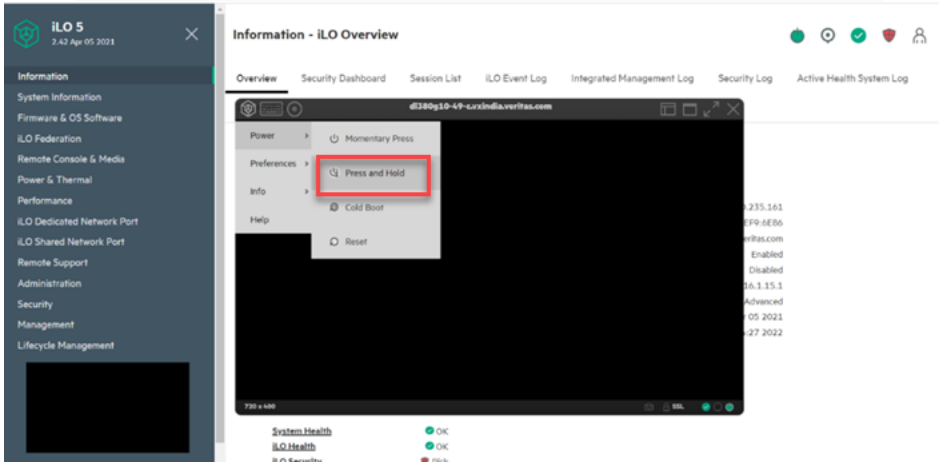

Installing additional DIMMs (performed by HPE)

The HPE representative installs the additional DIMMs.

Below are the recommendations from HPE for DIMMs slots:

HPE ProLiant Gen10 12 slot per CPU												
DIMM population order												
1 DIMM					8							
2 DIMM s					8				10			
3 DIMM s					8				10			12
4 DIMM s			3		5			8		10		
5 DIMM s*			3		5			8		10		12
6 DIMM s	1		3		5			8		10		12
7 DIMM s*	1		3		5		7	8		10		12
8 DIMM s			3	4	5	6	7	8	9	10		
9 DIMM s*	1		3		5		7	8	9	10	11	12
10 DIMM s*	1		3	4	5	6	7	8	9	10		12
11 DIMM s*	1		3	4	5	6	7	8	9	10	11	12
12 DIMM s	1	2	3	4	5	6	7	8	9	10	11	12

Notes: *Unbalanced, not recommended

Table 2. HPE SmartMemory DIMM population guidelines for HPE Gen11 servers with 16 DIMM slots per CPU

HPE ProLiant Gen11 servers 16 slots per CPU DIMM population order																
1 DIMM										10						
2 DIMMs[1]			3							10						
4 DIMMs[1]			3				7			10				14		
6 DIMMs			3		5		7			10				14		16
8 DIMMs[1, 2]	1		3		5		7			10		12		14		16
12 DIMMs	1	2	3		5	6	7			10	11	12		14	15	16
16 DIMMs[1, 2]	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16

[1] Support Hemi (hemisphere mode)
[2] Support Software Guard Extensions (SGX)

Completing the post-installation tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is installed, verify the memory on each node.

To verify that the DIMMs are installed, complete the following steps:

1 Power on the all the servers together and wait till the cluster is up.

Note: As soon as the power cable is plugged in the node, it starts automatically. So wait for DIMM replacement on all the nodes and attach power cable together.

---

2 Run the `show hardware-health` command:

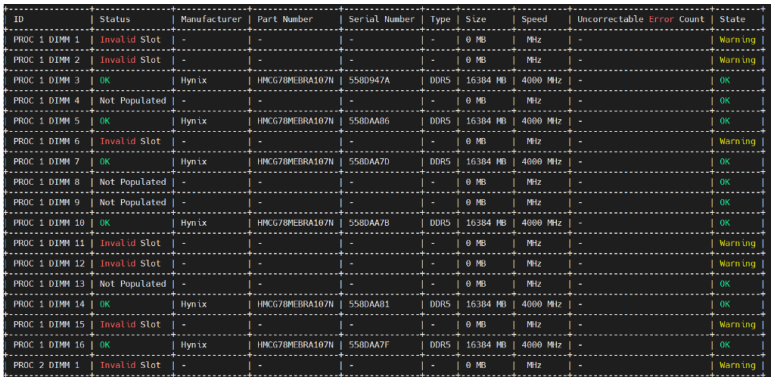

ID	Status	Manufacturer	Part Number	Serial Number	Type	Size	Speed	Uncorrectable Error Count	State
PROC 1 DIMM 1	Invalid Slot	-	-	-	-	0 MB	MHz	-	Warning
PROC 1 DIMM 2	Invalid Slot	-	-	-	-	0 MB	MHz	-	Warning
PROC 1 DIMM 3	OK	Hynix	HMCG78MEBRA107N	558D947A	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 4	Not Populated	-	-	-	-	0 MB	MHz	-	OK
PROC 1 DIMM 5	OK	Hynix	HMCG78MEBRA107N	558DAA86	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 6	Invalid Slot	-	-	-	-	0 MB	MHz	-	Warning
PROC 1 DIMM 7	OK	Hynix	HMCG78MEBRA107N	558DAA7D	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 8	Not Populated	-	-	-	-	0 MB	MHz	-	OK
PROC 1 DIMM 9	Not Populated	-	-	-	-	0 MB	MHz	-	OK
PROC 1 DIMM 10	OK	Hynix	HMCG78MEBRA107N	558DAA7B	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 11	Invalid Slot	-	-	-	-	0 MB	MHz	-	Warning
PROC 1 DIMM 12	Invalid Slot	-	-	-	-	0 MB	MHz	-	Warning
PROC 1 DIMM 13	Not Populated	-	-	-	-	0 MB	MHz	-	OK
PROC 1 DIMM 14	OK	Hynix	HMCG78MEBRA107N	558DAA81	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 15	Invalid Slot	-	-	-	-	0 MB	MHz	-	Warning
PROC 1 DIMM 16	OK	Hynix	HMCG78MEBRA107N	558DAA7F	DDR5	16384 MB	4000 MHz	-	OK
PROC 2 DIMM 1	Invalid Slot	-	-	-	-	0 MB	MHz	-	Warning

3 After adding the DIMMs, modify bom-conf. Either 12 or 24 DIMMS are supported. By default there are 12 DIMMs per node.

- Log in to node's appliance shell.
- Run command `support bom-conf get` command.

```
[nbfs-3.2] vflex5561-40.vxindia.veritas.com > support bom-conf get

The BOM configuration file is copied to /system/inst/patch/incoming, please run support share open to access it

Operation completed successfully
```

  The bom-conf file is stored in the `/system/inst/patch/incoming` location

- Run the `support elevate` command and provide the required credentials.
- Go to `/system/inst/patch/incoming` and edit the `bom-config.json` file.

```
vflex5561-40.vxindia.veritas.com:~ # cd /system/inst/patch/incoming
vflex5561-40.vxindia.veritas.com:/system/inst/patch/incoming # ls
bom-config.json
```

- Set DIMM to DIMM:24 and save the file.

```
    }
  }
},
"DIMM": "24",
"FC-ENABLE": "0",
"Packages" : [
  "VRTSvxos-mft",
  "kernel-mft",
  "nvme-cli",
  "flashupdt",
  "syscfg.x86_64",
  "sysinfo"
],
"BlackListModules" : [
  "ahci",
  "qla2xxx",
  "isci",
  "skx_edac"
]
```

- Go to the appliance shell and run the `support bom-conf update` command.
- Run the `show hardware-health` command:

DIMM Information

ID	Status	Manufacturer	Part Number	Serial Number	Type	Size	Speed	Uncorrectable Error Count	State
PROC 1 DIMM 1	OK	Hynix	HMCG78MEBRA107N	558D9479	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 2	OK	Hynix	HMCG78MEBRA107N	558DAA88	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 3	OK	Hynix	HMCG78MEBRA107N	558D947A	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 4	Not Populated	-	-	-	-	0 MB	MHz	-	OK
PROC 1 DIMM 5	OK	Hynix	HMCG78MEBRA107N	558DAA86	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 6	OK	Hynix	HMCG78MEBRA107N	558D947B	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 7	OK	Hynix	HMCG78MEBRA107N	558DAA7D	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 8	Not Populated	-	-	-	-	0 MB	MHz	-	OK
PROC 1 DIMM 9	Not Populated	-	-	-	-	0 MB	MHz	-	OK
PROC 1 DIMM 10	OK	Hynix	HMCG78MEBRA107N	558DAA7B	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 11	OK	Hynix	HMCG78MEBRA107N	558D948D	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 12	OK	Hynix	HMCG78MEBRA107N	558D9483	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 13	Not Populated	-	-	-	-	0 MB	MHz	-	OK
PROC 1 DIMM 14	OK	Hynix	HMCG78MEBRA107N	558DAA81	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 15	OK	Hynix	HMCG78MEBRA107N	558D9472	DDR5	16384 MB	4000 MHz	-	OK
PROC 1 DIMM 16	OK	Hynix	HMCG78MEBRA107N	558DAA7F	DDR5	16384 MB	4000 MHz	-	OK

4 Verify the total memory of each node. You can check the total memory of the node by using the `support shell` command from the node-level CLI, and then using any of the available system commands:

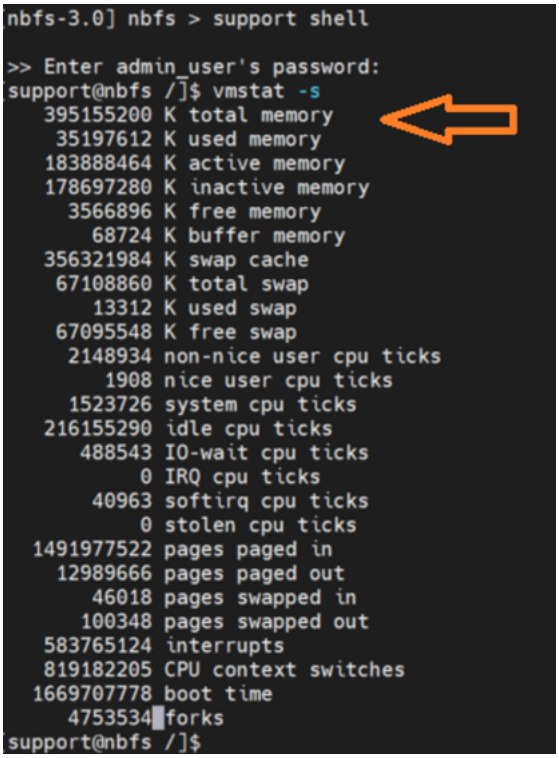

5 Check the node state. Navigate to **Monitor > Infrastructure > Nodes**.

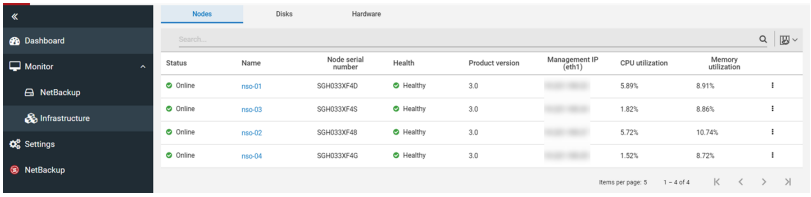

6 Click **Dashboard** and check the status of the services.

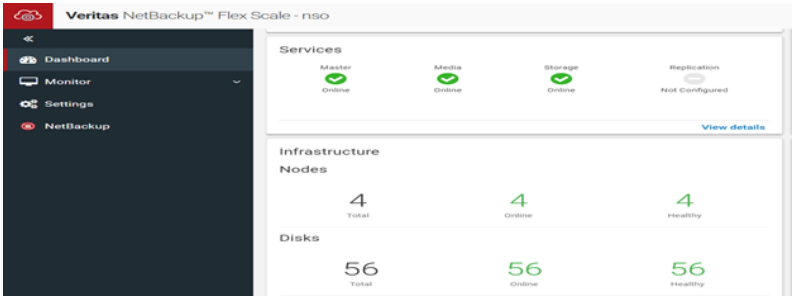

7 Verify that the added DIMMs show **Good, in use** status in the iLO remote console.

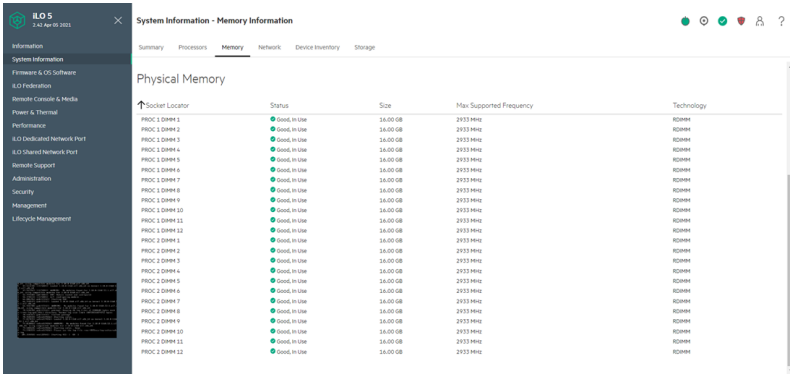

Replacement procedure for memory (DIMMs)

This topic describes the process for replacing DIMMs that are faulty.

Identifying DIMM failure (performed by the CHS team)

Email alerts are generated for DIMM failures if SMTP is configured for AutoSupport. The alert includes the CPU and the slot the DIMM is missing or faulty.

The HPE Integrated Lights-Out (iLO) remote console shows the missing DIMM slots:

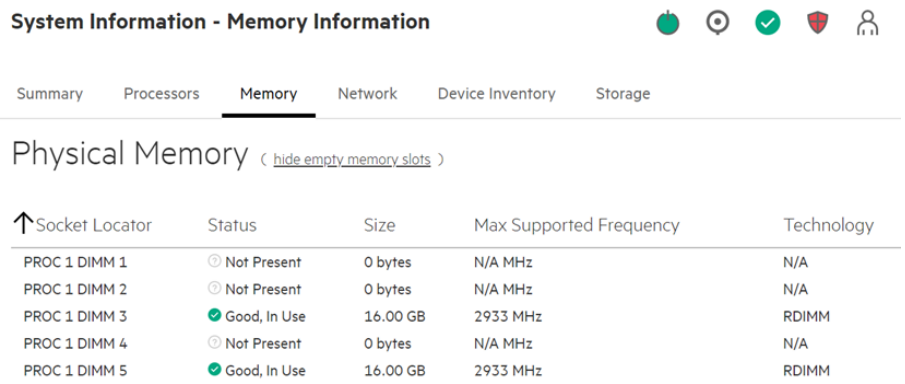

Replacing the DIMMs (performed by Veritas TSE)

To replace the DIMMs, complete the following steps:

1. Sign in to the NetBackup Flex Scale infrastructure management UI and navigate to **Monitor > Infrastructure > Nodes**.
2. On the node where the failure occurred, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Shutdown node**.

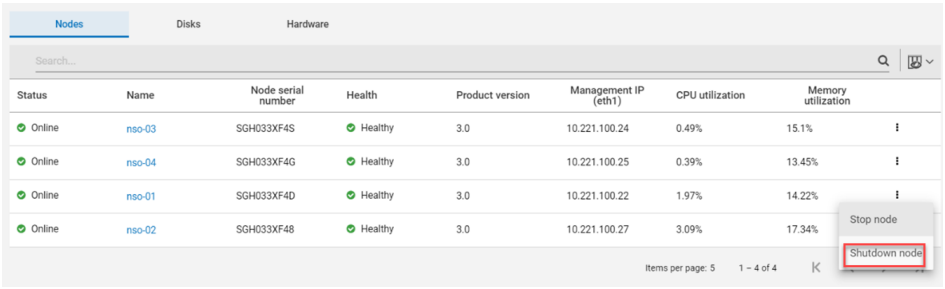

3. Confirm that the node is shut down successfully. In the UI, you can view the notification at the top of the page.

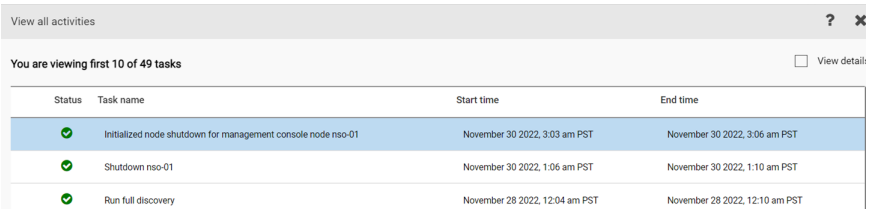

4. Replace the DIMM.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware component is replaced, verify that the issue is resolved.

To verify that the issue is resolved, complete the following steps:

1 Power on the server and wait till the node joins the cluster.

2 Check the node state. Navigate to **Monitor > Infrastructure > Nodes**.

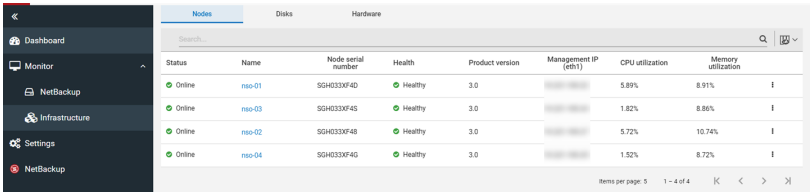

3 Click **Dashboard** and check the status of the services.

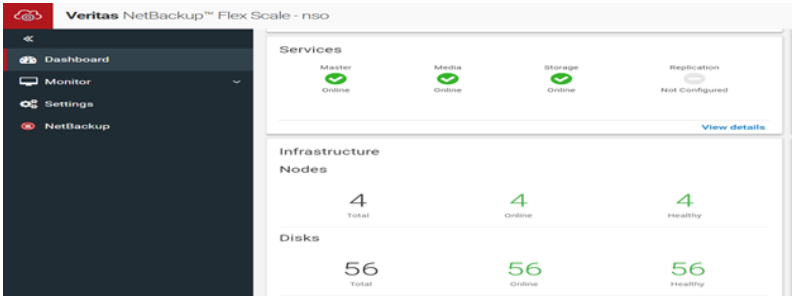

4 Verify that the replaced DIMMs show **Good, in use** status in the iLO remote console.

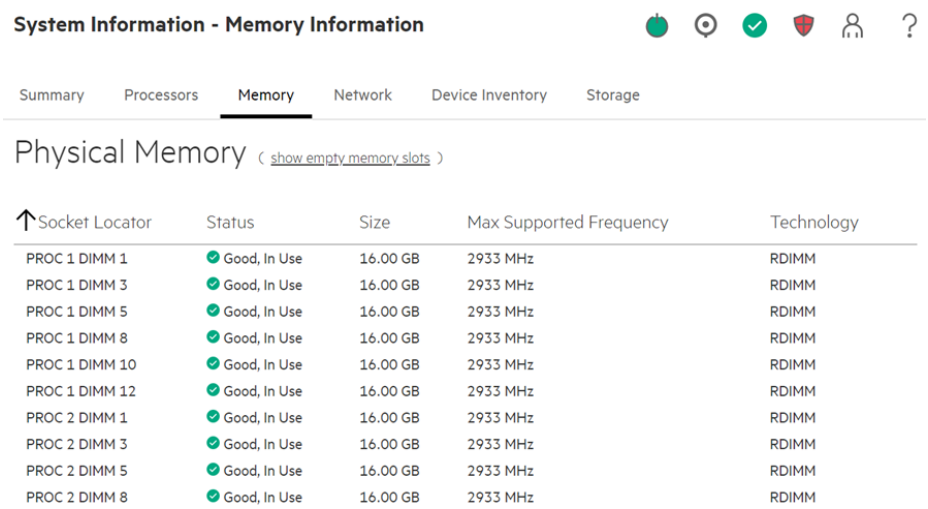

Replacement procedure for Mellanox port

This topic describes the process of replacing the Mellanox port on an HPE server node.

Identifying a Mellanox port failure (performed by the customer)

There might be a problem with the cables connection or the Mellanox port. To isolate the issue, first disconnect and reconnect the cables. If you face the same issue, replace the cables and try again.

Mellanox private port (eth4)

After you unplug the cable from the Mellanox private port (eth4), the node status shows unhealthy. In the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure > Nodes**.

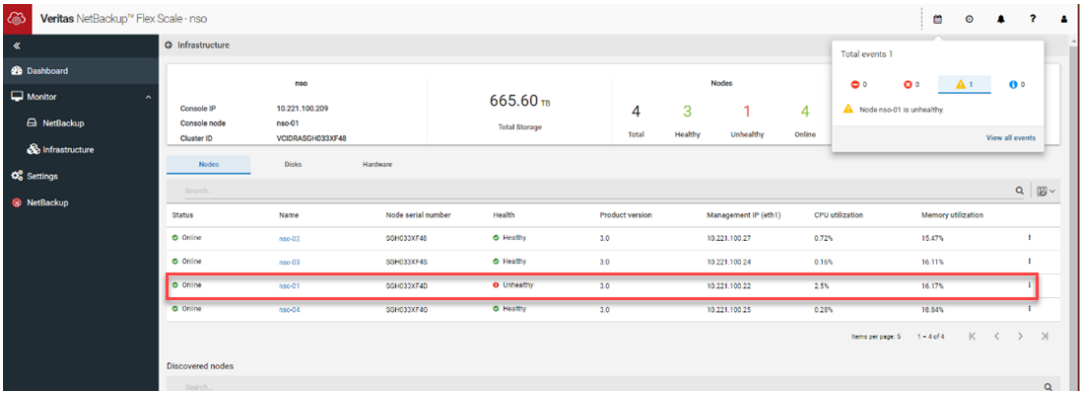

Navigate to **Infrastructure > Hardware > Network Card**. On the **Hardware** tab, the eth4 status is shown **Unplugged**.

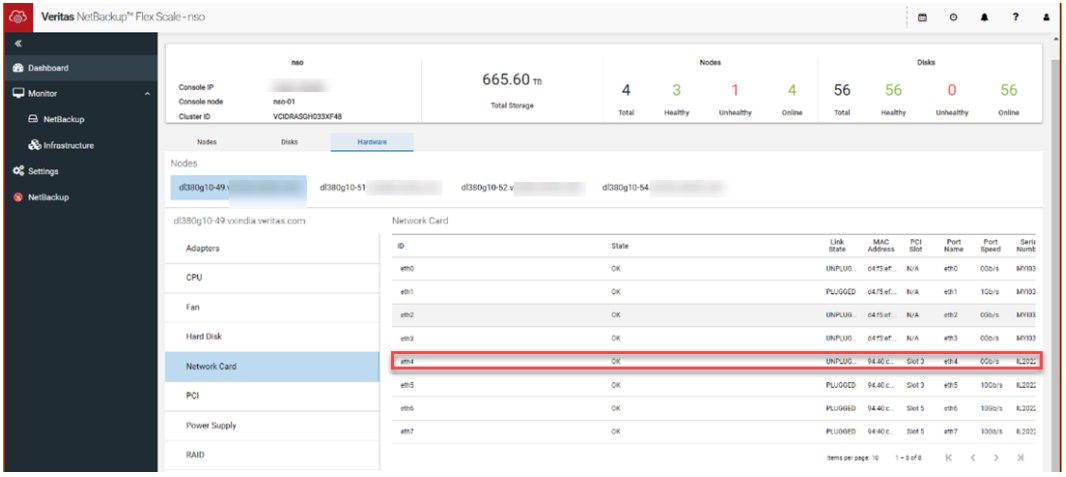

Click **Dashboard**. The private network status is shown unhealthy:

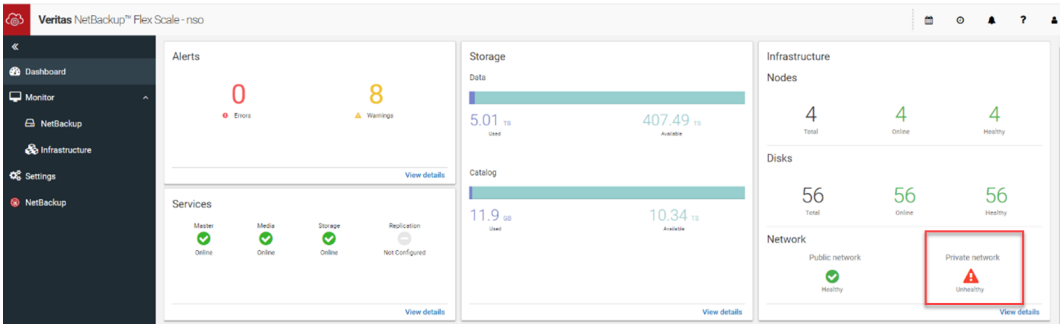

In the iLO remote console, navigate to **System information > Network**. The status for the affected port is shown **Link down**.

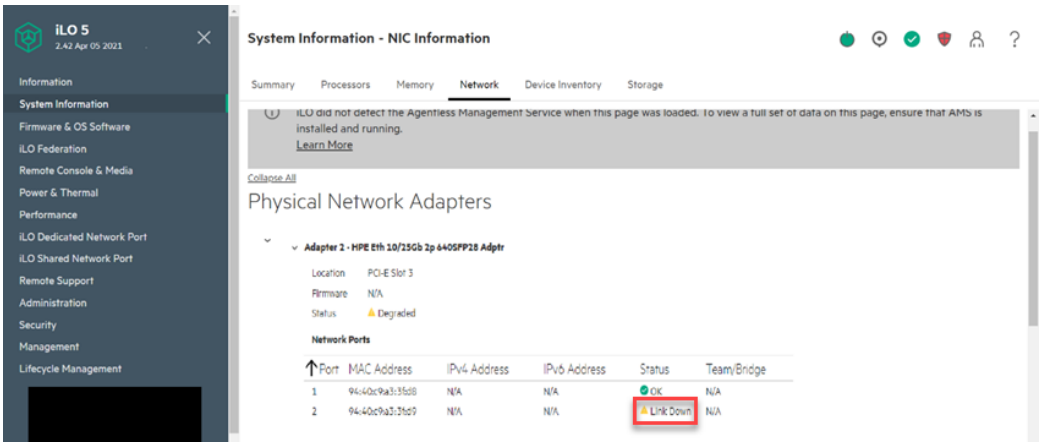

An event is generated stating that eth4 is down. In the NetBackup Flex Scale infrastructure management UI, navigate to **Settings > Events**:

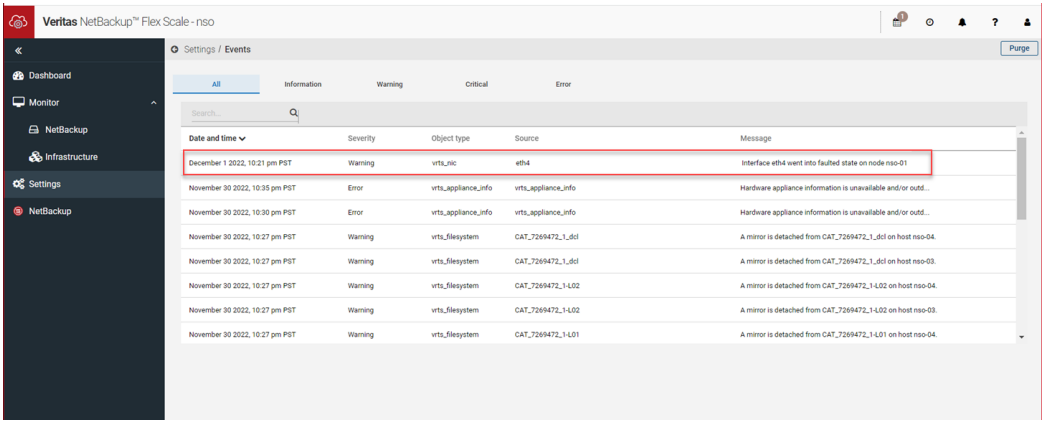

Plug in the cable in the affected port. Review the cabling guidelines before you plug in the cable.

After plugging in the eth4 cable, the node status is shown healthy. Navigate to **Monitor > Infrastructure > Nodes**:

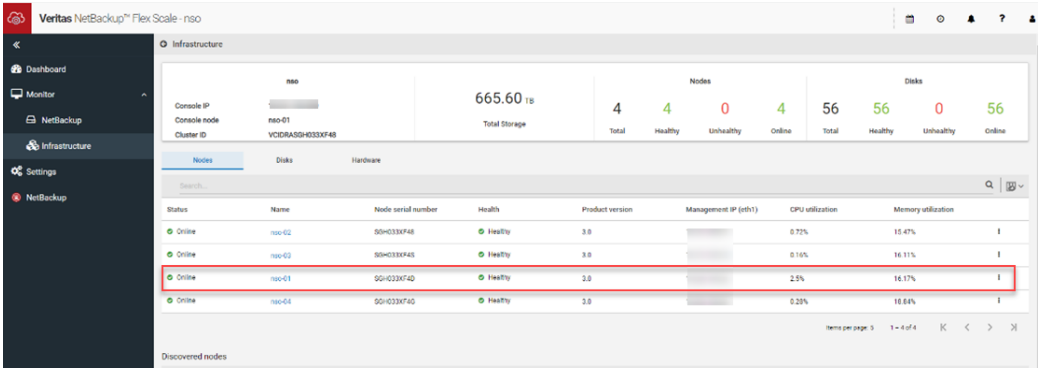

The eth4 status changes to Plugged on the Hardware tab. Navigate to **Monitor > Infrastructure > Hardware > Network Card**:

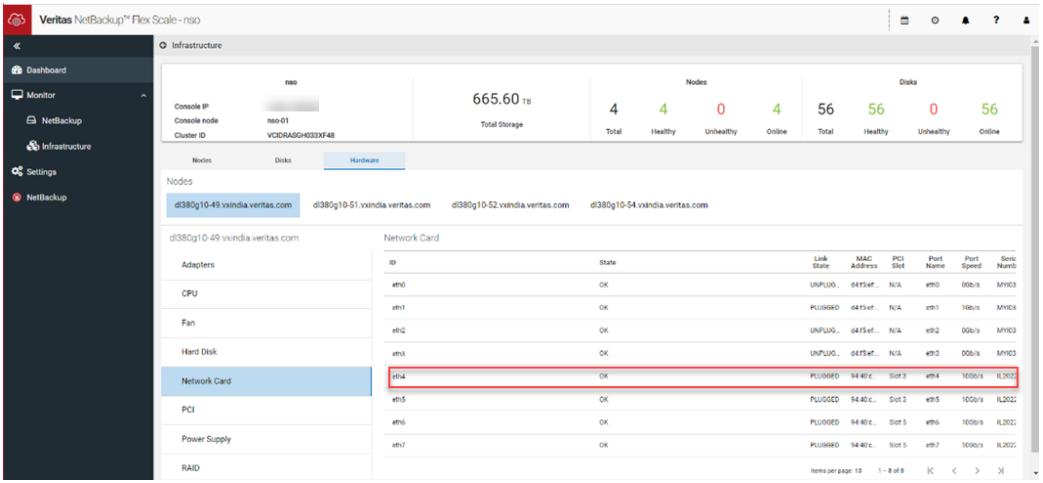

In the iLO remote console, navigate to **System information > Network**. The status for the affected port is shown **OK**.

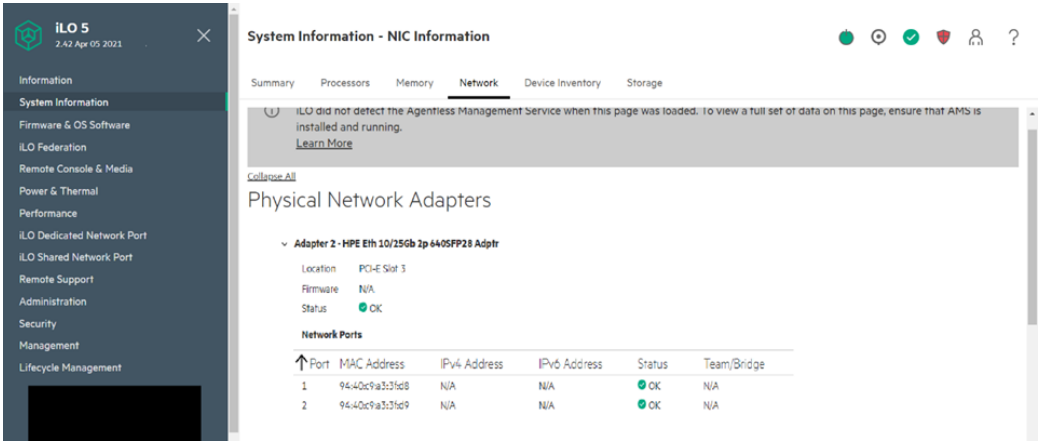

An event is generated, which shows that eth4 is online. navigate to **Settings > Events**:

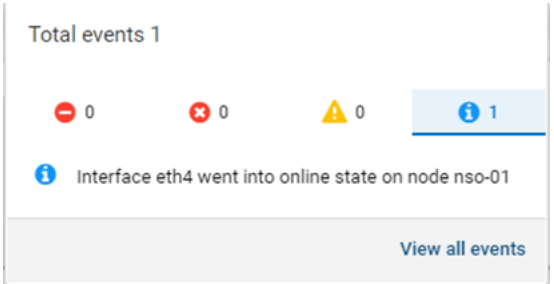

On the Dashboard, the private network status is shown healthy. It might take approximately five minutes to refresh the status in the UI.

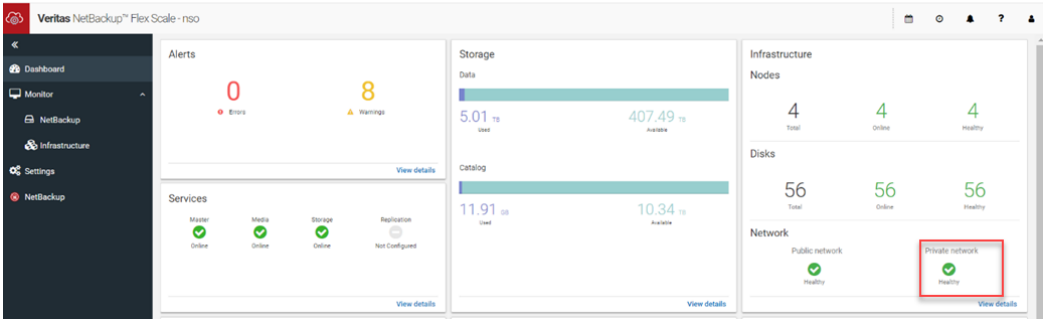

Mellanox public port (eth5)

After you unplug the cable from the Mellanox public port (eth5), the node status shows unhealthy. In the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure > Nodes**.

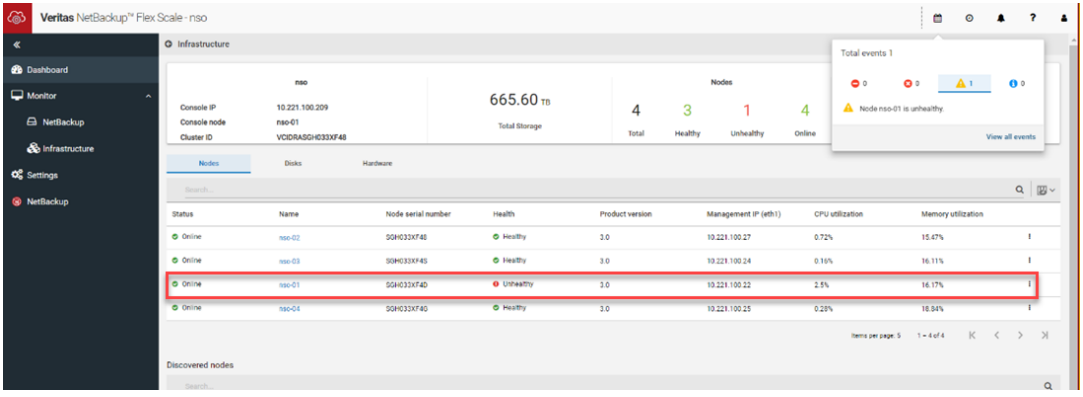

Navigate to **Infrastructure > Hardware > Network Card**. On the **Hardware** tab, the eth5 status is shown **Unplugged**.

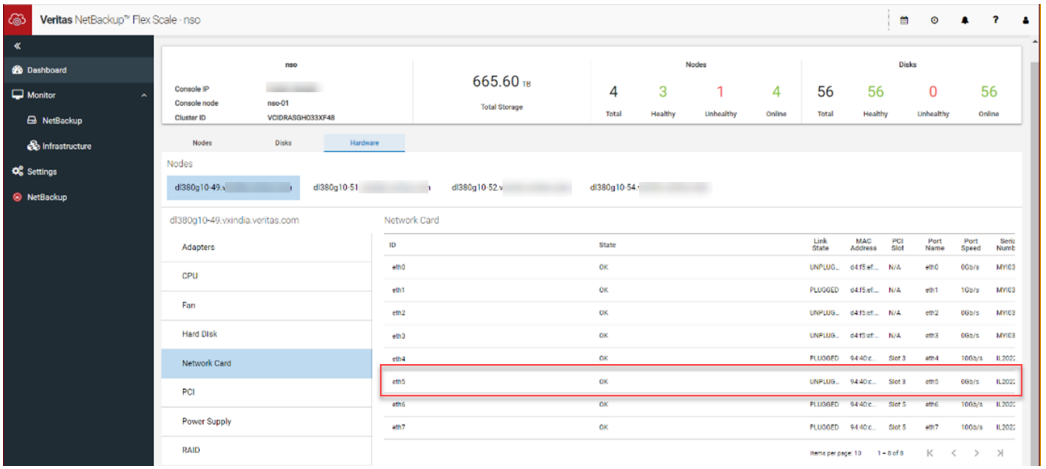

Click **Dashboard**. The public network status is shown unhealthy:

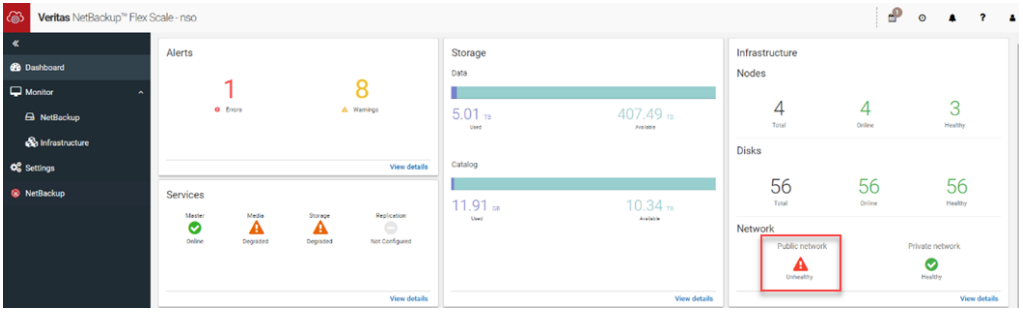

In the iLO remote console, navigate to **System information > Network**. The status for the affected port is shown **Link down**.

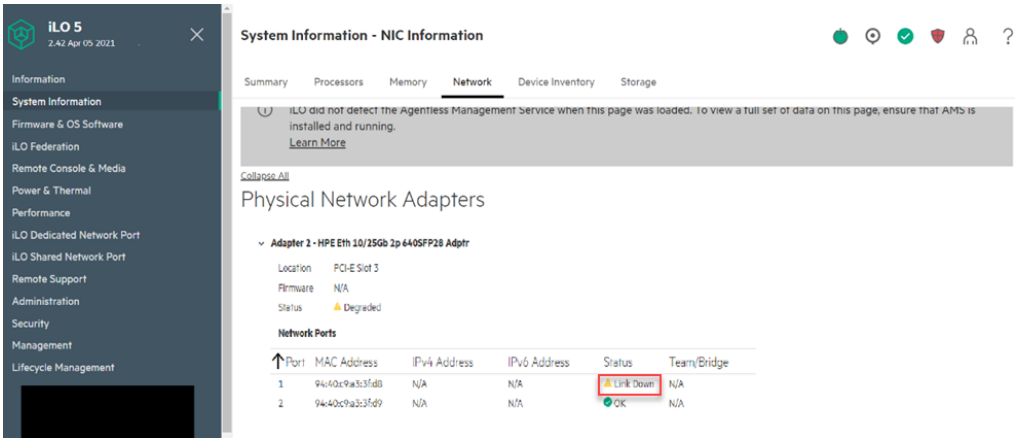

Plug in the cable in the affected port. Review the cabling guidelines before you plug in the cable. After plugging in the eth5 cable, the node status is shown healthy. Navigate to**Monitor > Infrastructure > Nodes**:

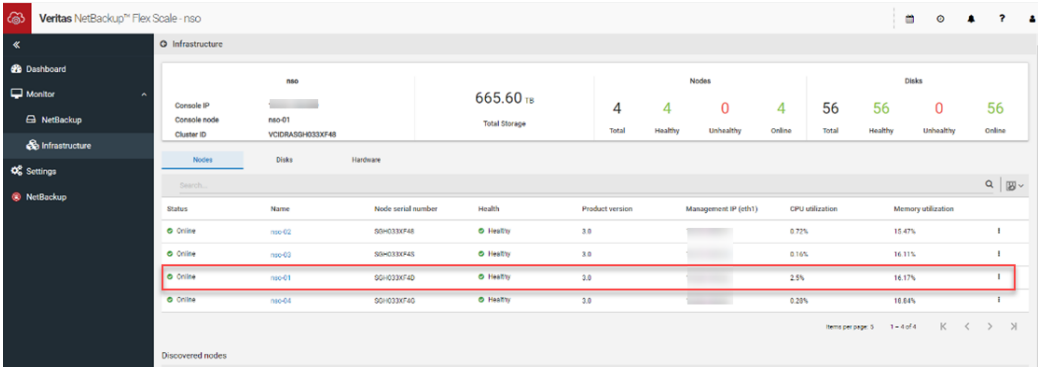

The eth5 status changes to **Plugged** on the **Hardware** tab. Navigate to **Monitor > Infrastructure > Hardware > Network Card**:

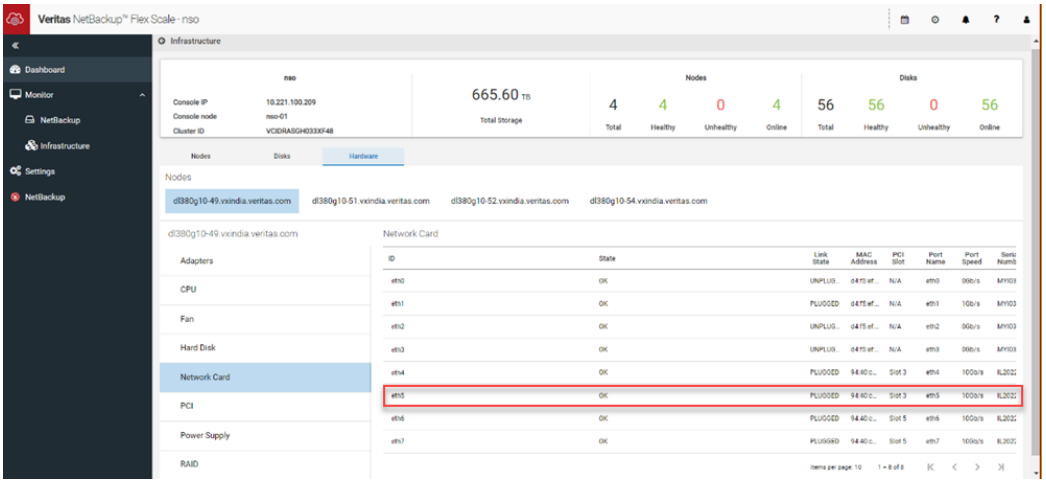

In the iLO remote console, navigate to **System information > Network**. The status for the affected port is shown **OK**.

Adapter 2 - HPE Eth 10/25Gb 2p 640SFP28 Adptr

Location PCI-E Slot 3
Firmware N/A
Status OK

Network Ports

Port	MAC Address	IPv4 Address	IPv6 Address	Status	Team/Bridge
1	94:40:c9:a3:3f:d8	N/A	N/A	OK	N/A
2	94:40:c9:a3:3f:d9	N/A	N/A	OK	N/A

On the **Dashboard**, the public network status is shown healthy. It might take approximately five minutes to refresh the status in the UI.

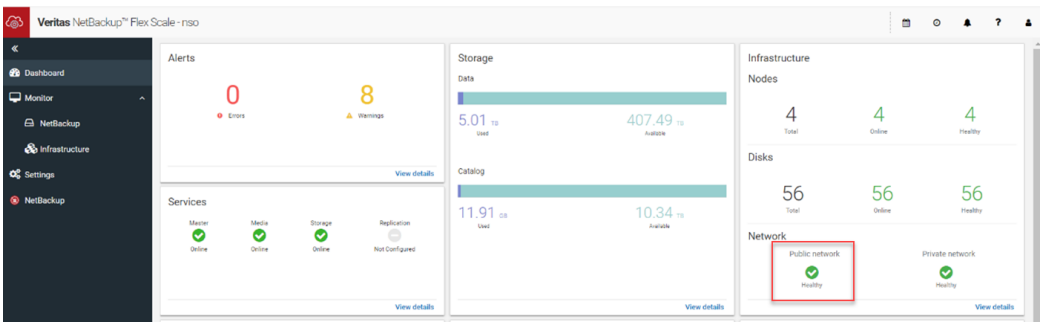

If the issue persists it implies that the Mellanox port is faulty. The CHS team will need to involve HPE for replacement of the port. Contact Veritas TSE to replace the node with the faulty iLO port.

Collecting HPE Active Health System (AHS) logs (performed by Veritas TSE)

Before you contact the hardware vendor for replacing the failed component, collect AHS logs. To collect the AHS logs, in the NetBackup Flex Scale infrastructure management UI, navigate to **Settings > Diagnostics > Basic > Appliance**.

Noting the MAC IDs (performed by Veritas TSE)

From the node-level CLI use the `system hardware-health` command to note the eth4 and eth5 MAC addresses.

```
| eth4 | Slot 3   | 817751-001 | IL202202L6    | FIBRE        | 10Gb/s     | 94:40:c9:a3:3f:d9
+------+----------+------------+---------------+--------------+------------+------------------
| eth5 | Slot 3   | 817751-001 | IL202202L6    | FIBRE        | 0Gb/s      | 94:40:c9:a3:3f:d8
```

Shutting down the node (performed by Veritas TSE)

Before an HPE representative can replace the Mellanox port, you must shut down the node.

To shut down the node:

1. Sign in to the NetBackup Flex Scale infrastructure management UI and navigate to **Monitor > Infrastructure > Nodes**.
2. On the node where the failure occurred, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Shutdown node**.

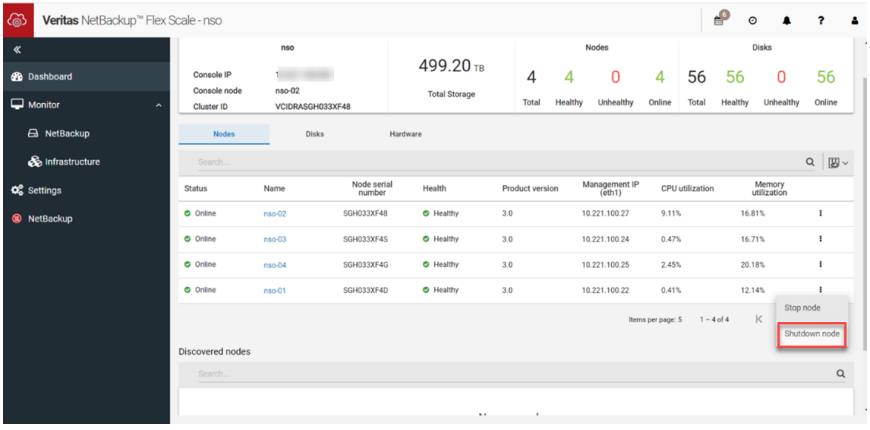

Replacing the Mellanox port (performed by the HPE vendor)

The HPE representative replaces the Mellanox port.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

To verify that the issue is resolved, complete the following steps:

1 Restart the node from the iLO remote console using the **Power > Momentary Press** option.

The green color power symbol indicates that the node has started.

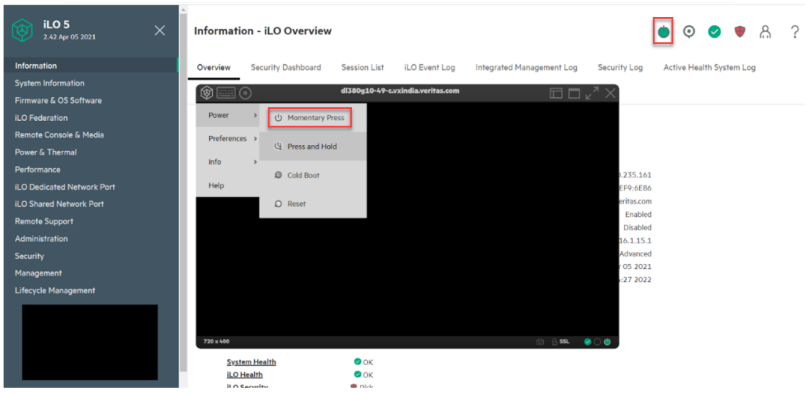

2 As the port is changed physically, the MAC address is also changed, which results in the following cases:

The node status is shown unhealthy in the UI:

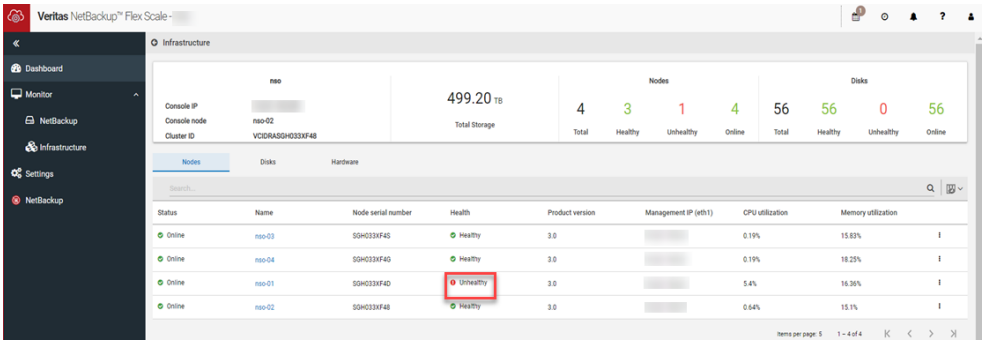

An event is shown on the **Settings > Events** page of the UI:

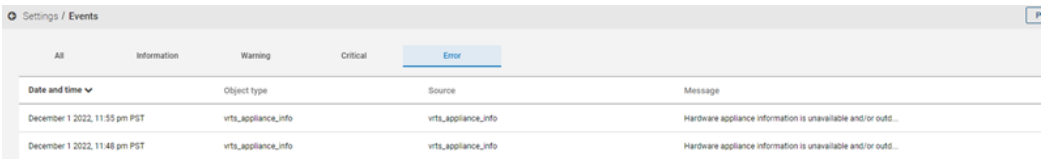

From the node-level CLI, the `system hardware-health` command shows that eth4 was plugged and eth5 was unplugged:

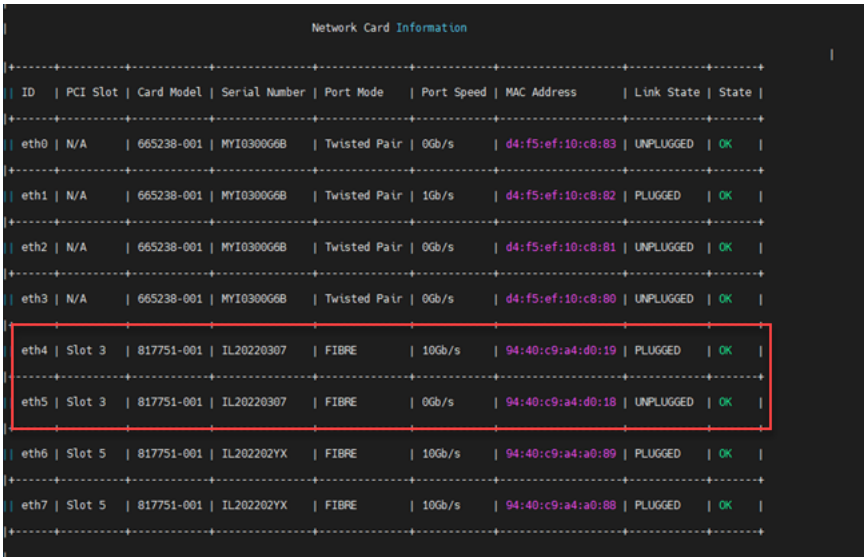

Network Card Information

ID	PCI Slot	Card Model	Serial Number	Port Mode	Port Speed	MAC Address	Link State	State
eth0	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:83	UNPLUGGED	OK
eth1	N/A	665238-001	MYI0300G6B	Twisted Pair	1Gb/s	d4:f5:ef:10:c8:82	PLUGGED	OK
eth2	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:81	UNPLUGGED	OK
eth3	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:80	UNPLUGGED	OK
eth4	Slot 3	817751-001	IL20220307	FIBRE	10Gb/s	94:40:c9:a4:d0:19	PLUGGED	OK
eth5	Slot 3	817751-001	IL20220307	FIBRE	0Gb/s	94:40:c9:a4:d0:18	UNPLUGGED	OK
eth6	Slot 5	817751-001	IL202202YX	FIBRE	10Gb/s	94:40:c9:a4:a0:89	PLUGGED	OK
eth7	Slot 5	817751-001	IL202202YX	FIBRE	10Gb/s	94:40:c9:a4:a0:88	PLUGGED	OK

Click **Dashboard**. The public network status is shown unhealthy:

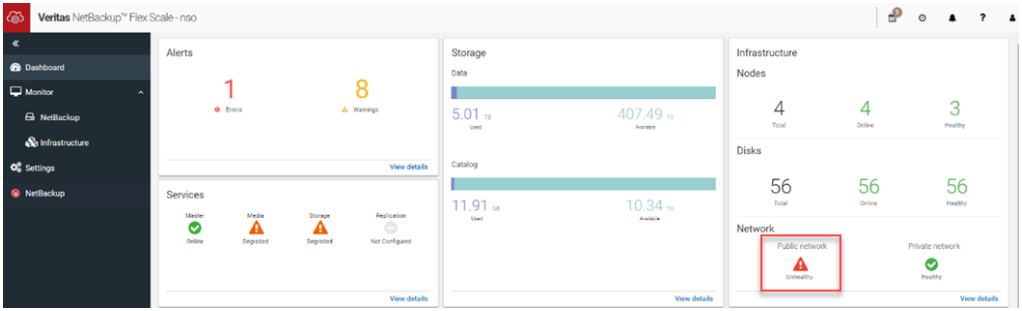

In the iLO remote console go to **System information > Network**:

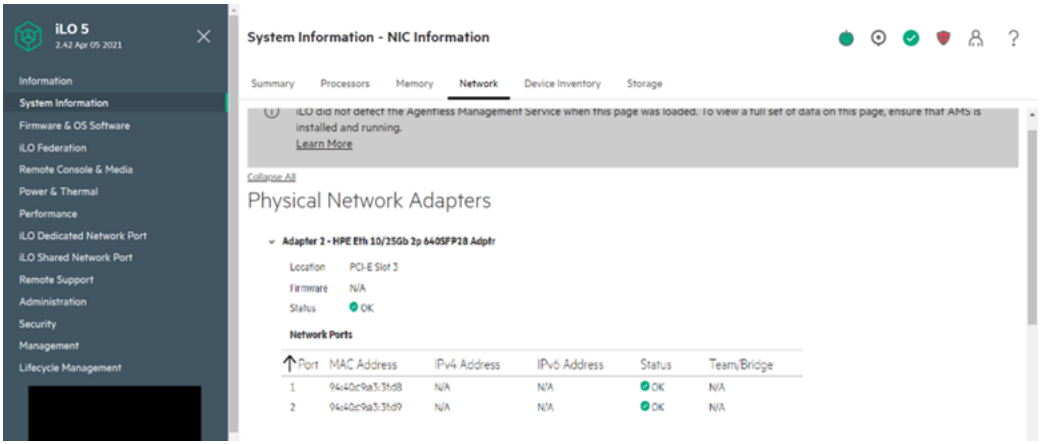

3 Verify that the node status is shown healthy in the UI:

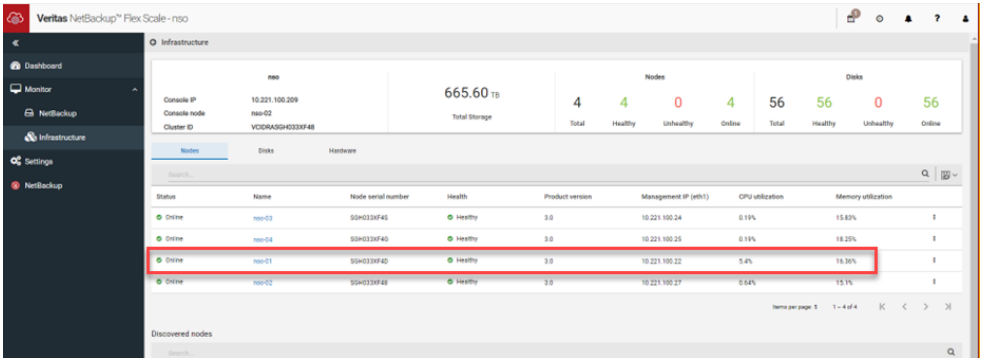

4 The public network status is shown healthy:

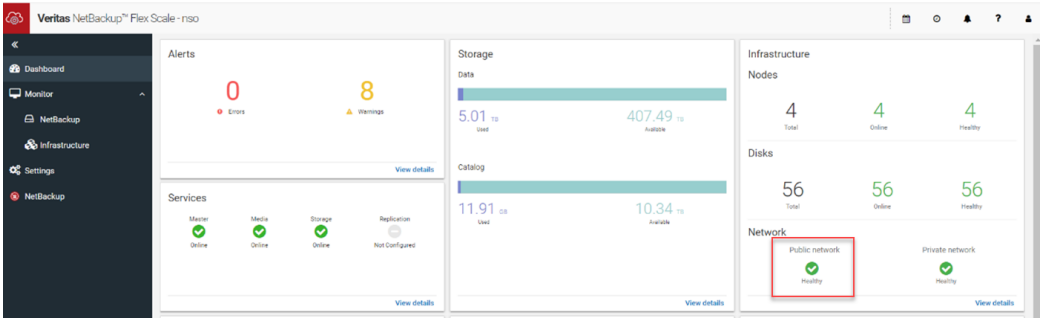

5 Verify that eth4 and eth5 both are shown **PLUGGED** on the **Hardware** tab.

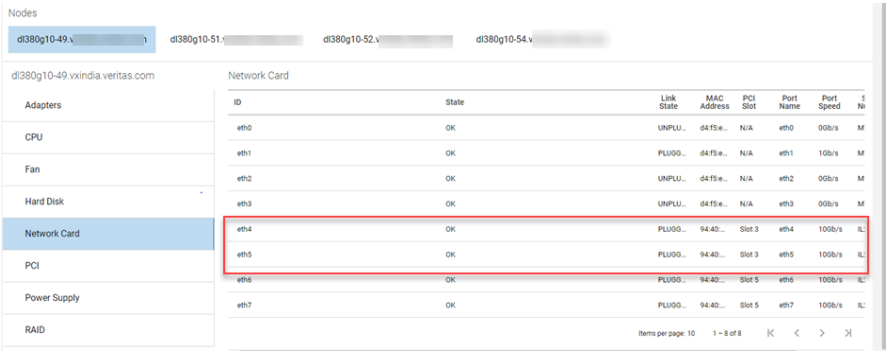

6 Verify that the changed MAC ID can be seen in the `system hardware-health`and eth4 and eth5 are shown plugged:

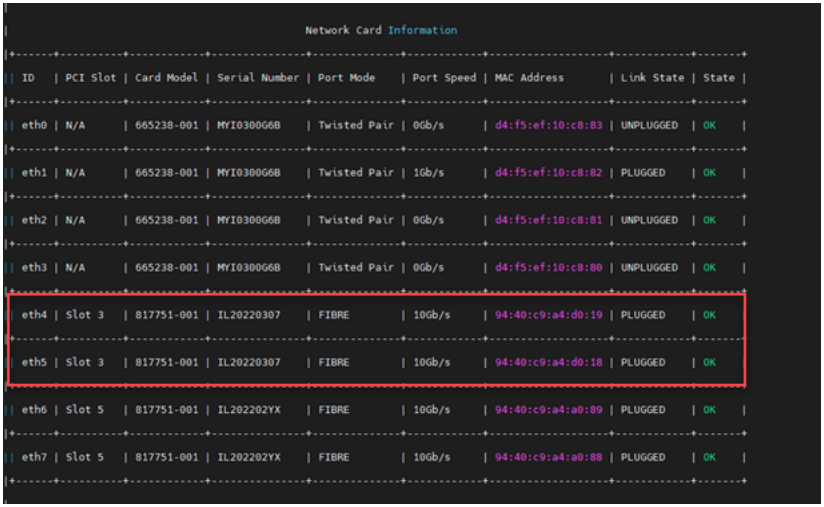

Network Card Information

ID	PCI Slot	Card Model	Serial Number	Port Mode	Port Speed	MAC Address	Link State	State
eth0	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:83	UNPLUGGED	OK
eth1	N/A	665238-001	MYI0300G6B	Twisted Pair	1Gb/s	d4:f5:ef:10:c8:82	PLUGGED	OK
eth2	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:81	UNPLUGGED	OK
eth3	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:80	UNPLUGGED	OK
eth4	Slot 3	817751-001	IL20220307	FIBRE	10Gb/s	94:40:c9:a4:d0:19	PLUGGED	OK
eth5	Slot 3	817751-001	IL20220307	FIBRE	10Gb/s	94:40:c9:a4:d0:18	PLUGGED	OK
eth6	Slot 5	817751-001	IL202202YX	FIBRE	10Gb/s	94:40:c9:a4:a0:89	PLUGGED	OK
eth7	Slot 5	817751-001	IL202202YX	FIBRE	10Gb/s	94:40:c9:a4:a0:88	PLUGGED	OK

7 An event is generated notifying that eth4 and eth5 are online:

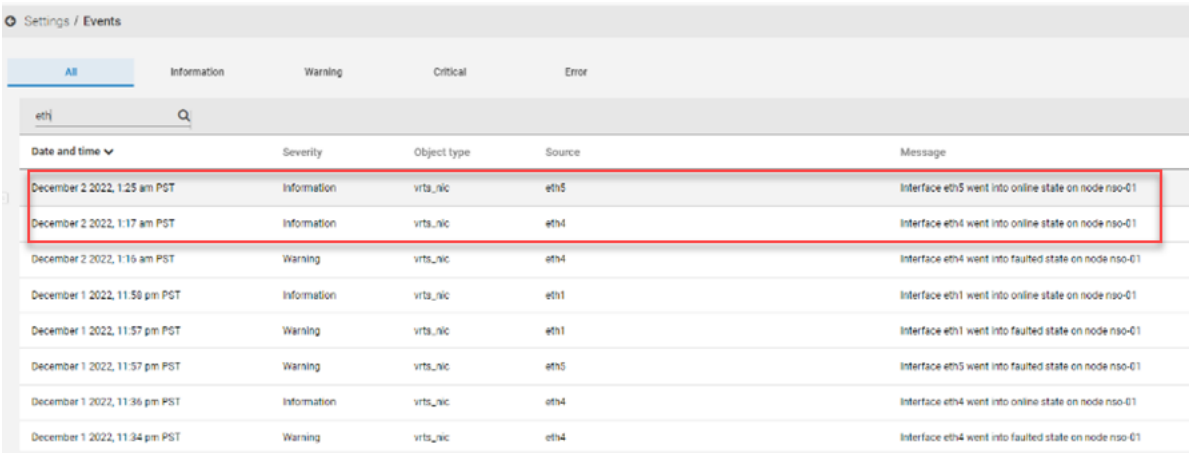

Settings / Events

All | Information | Warning | Critical | Error

eth

Date and time	Severity	Object type	Source	Message
December 2 2022, 1:25 am PST	Information	vrts_nic	eth5	Interface eth5 went into online state on node nso-01
December 2 2022, 1:17 am PST	Information	vrts_nic	eth4	Interface eth4 went into online state on node nso-01
December 2 2022, 1:16 am PST	Warning	vrts_nic	eth4	Interface eth4 went into faulted state on node nso-01
December 1 2022, 11:58 pm PST	Information	vrts_nic	eth1	Interface eth1 went into online state on node nso-01
December 1 2022, 11:57 pm PST	Warning	vrts_nic	eth1	Interface eth1 went into faulted state on node nso-01
December 1 2022, 11:57 pm PST	Warning	vrts_nic	eth5	Interface eth5 went into faulted state on node nso-01
December 1 2022, 11:36 pm PST	Information	vrts_nic	eth4	Interface eth4 went into online state on node nso-01
December 1 2022, 11:34 pm PST	Warning	vrts_nic	eth4	Interface eth4 went into faulted state on node nso-01

Replacement procedure for SFP port

This topic describes the process of replacing the SFP port on an HPE server node. The SFPs are hot-swappable, but you need to disconnect the network cable from the node.

Identifying an SFP failure (performed by the customer)

There might be a problem with the cable connection or the SFP port. To isolate the issue, first disconnect and reconnect the cables. If you face the same issue, replace the cables and try again.

SFP private port (eth4)

After you unplug the cable from the SFP private port (eth4), the node status shows unhealthy. In the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure > Nodes**.

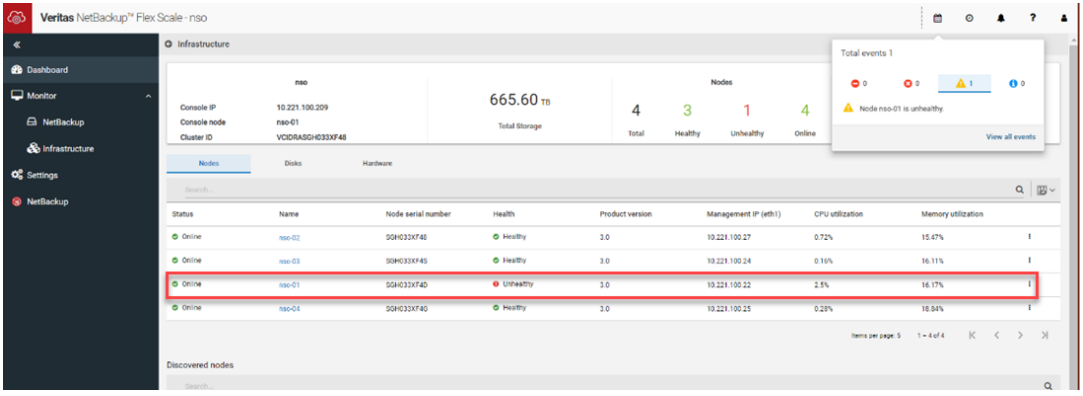

Navigate to **Infrastructure > Hardware > Network Card**. On the **Hardware** tab, the eth4 status is shown **Unplugged**.

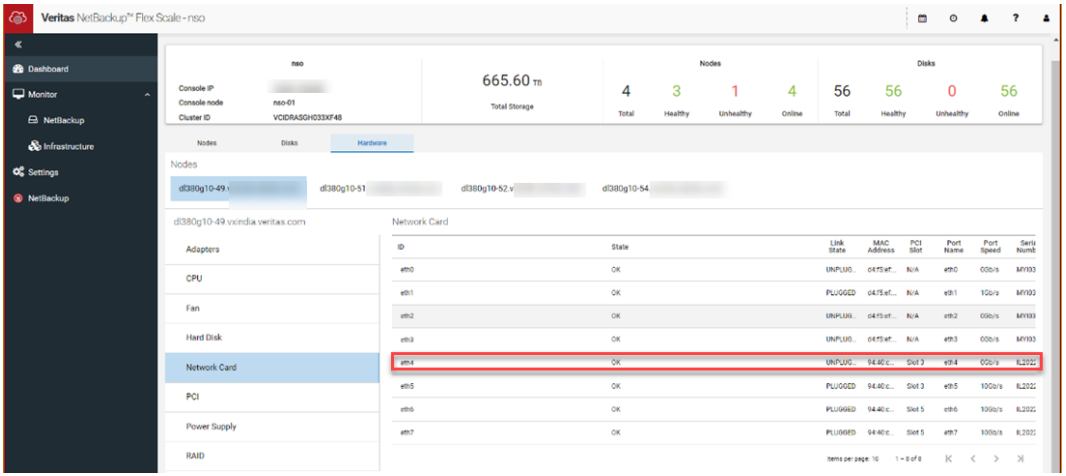

Click **Dashboard**. The private network status is shown unhealthy:

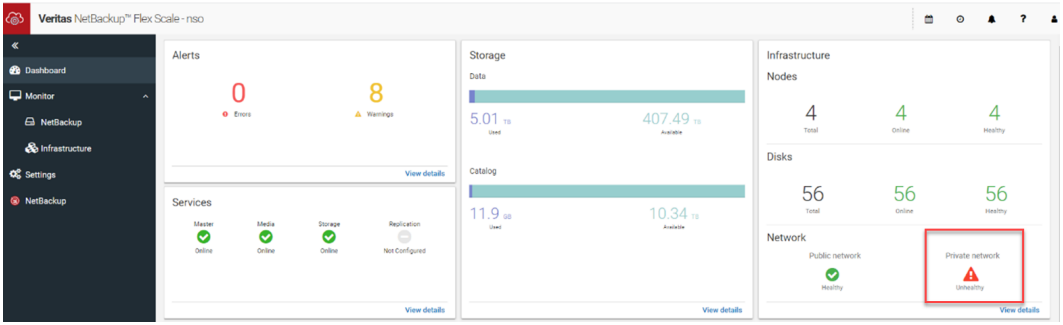

In the iLO remote console, navigate to **System information > Network**. The status for the affected port is shown **Link down**.

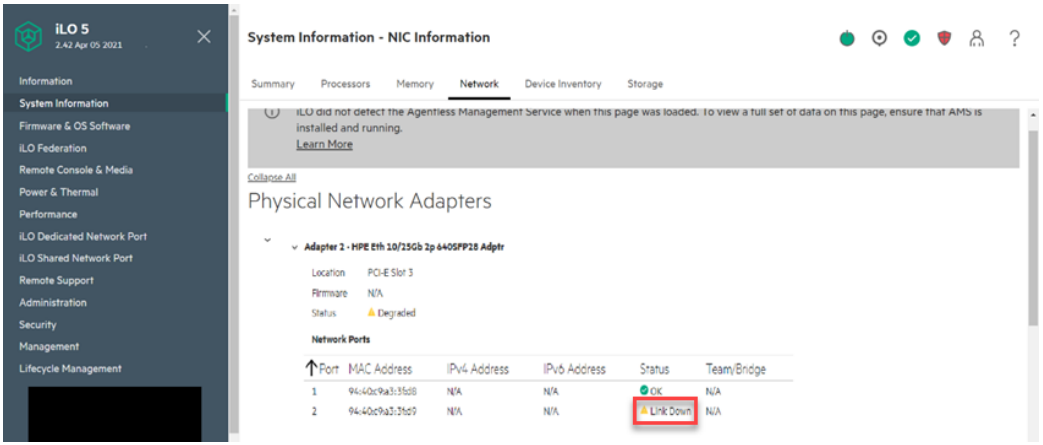

An event is generated stating that eth4 is down. In the NetBackup Flex Scale infrastructure management UI, navigate to **Settings > Events**:

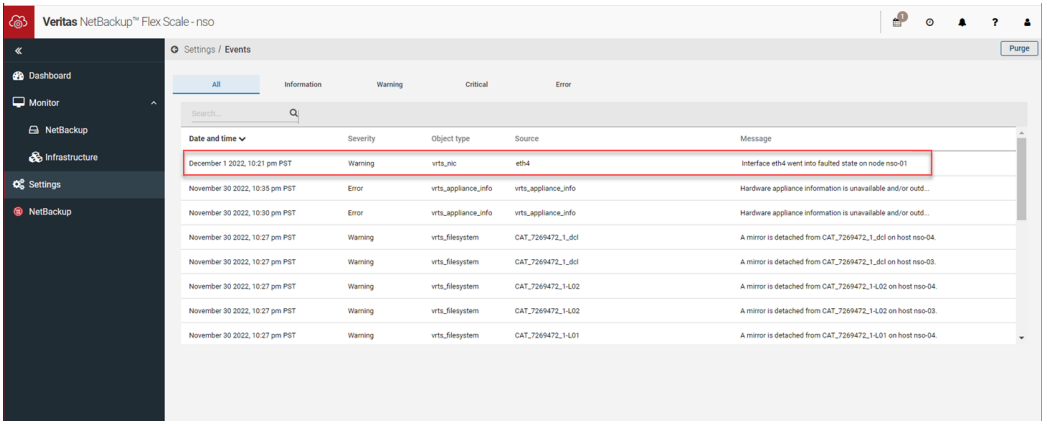

Plug in the cable in the affected port. Review the cabling guidelines before you plug in the cable.

After plugging in the eth4 cable, the node status is shown healthy. Navigate to **Monitor > Infrastructure > Nodes**:

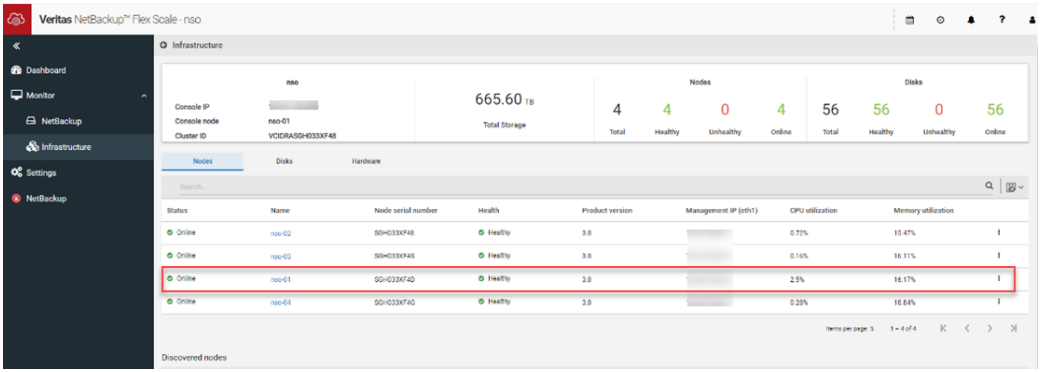

The eth4 status changes to Plugged on the Hardware tab. Navigate to **Monitor > Infrastructure > Hardware > Network Card**:

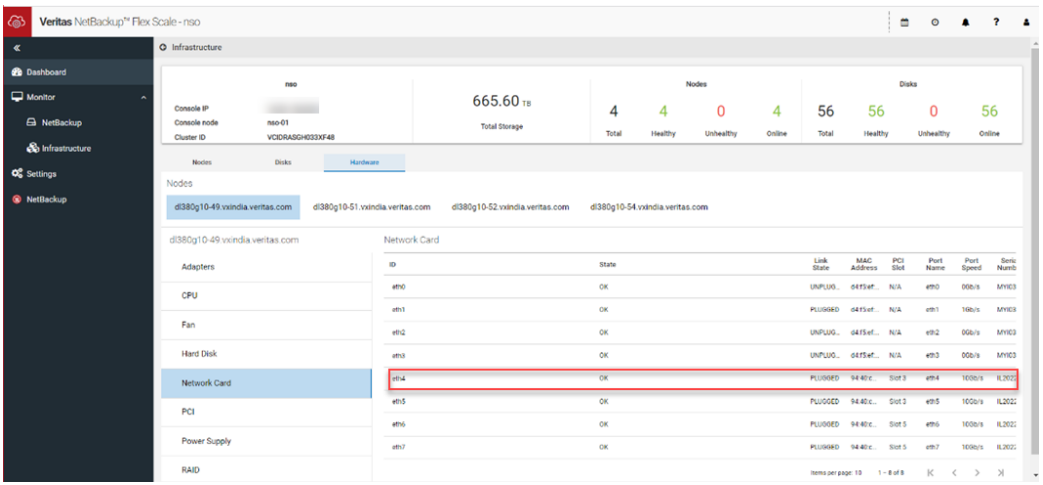

In the iLO remote console, navigate to **System information > Network**. The status for the affected port is shown **OK**.

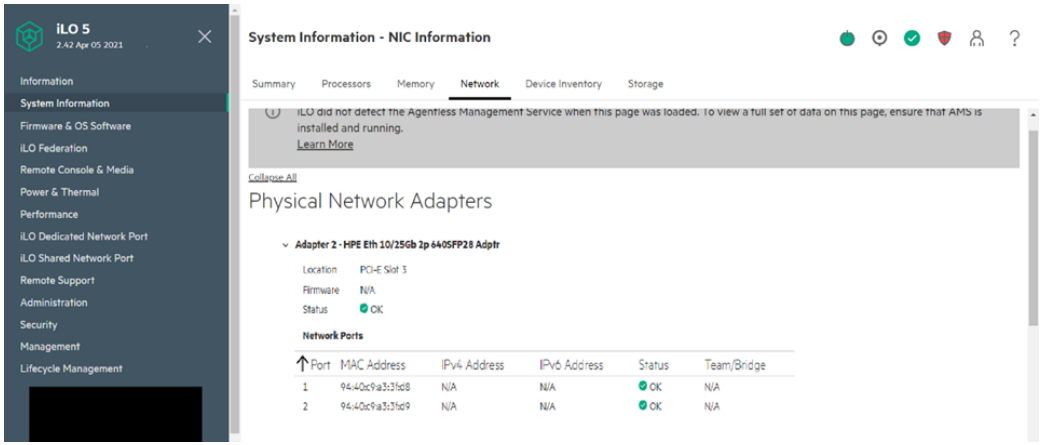

An event is generated, which shows that eth4 is online. navigate to **Settings > Events**:

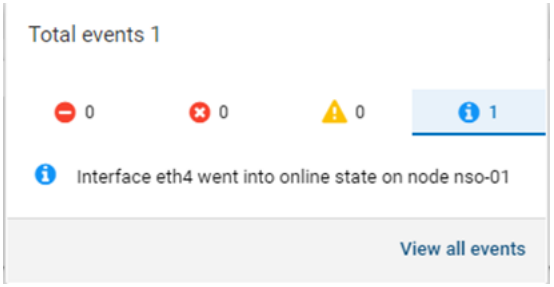

On the Dashboard, the private network status is shown healthy. It might take approximately five minutes to refresh the status in the UI.

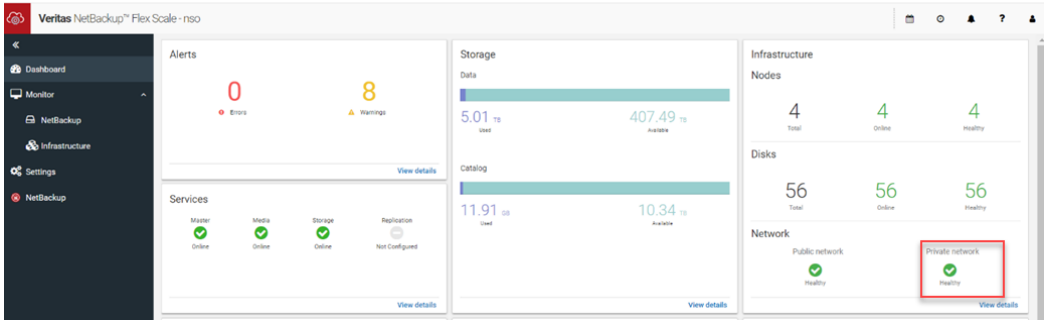

SFP public port (eth5)

After you unplug the cable from the SFP public port (eth5), the node status shows unhealthy. In the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure > Nodes**.

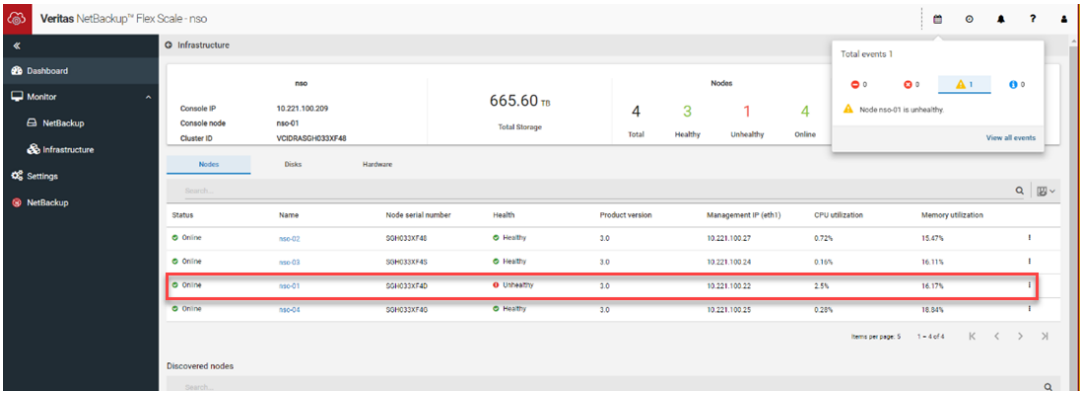

Navigate to **Infrastructure > Hardware > Network Card**. On the **Hardware** tab, the eth5 status is shown **Unplugged**.

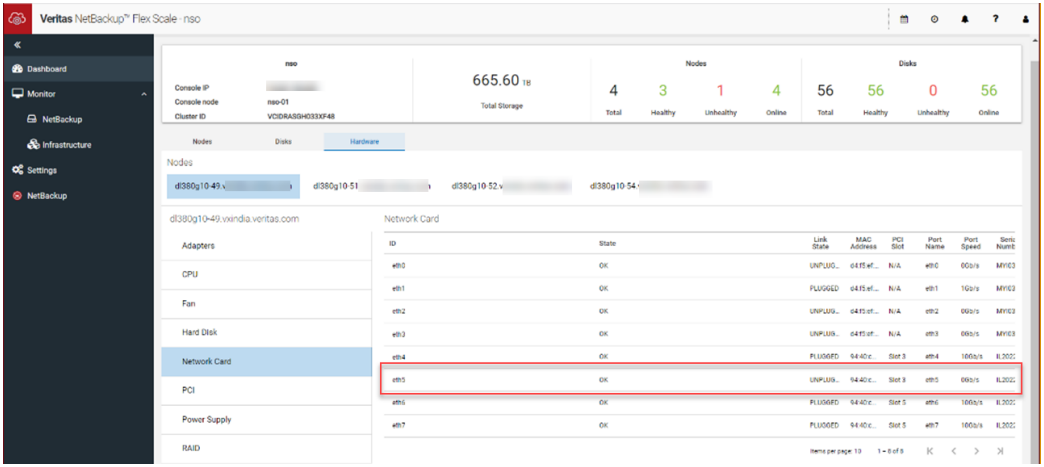

Click **Dashboard**. The public network status is shown unhealthy:

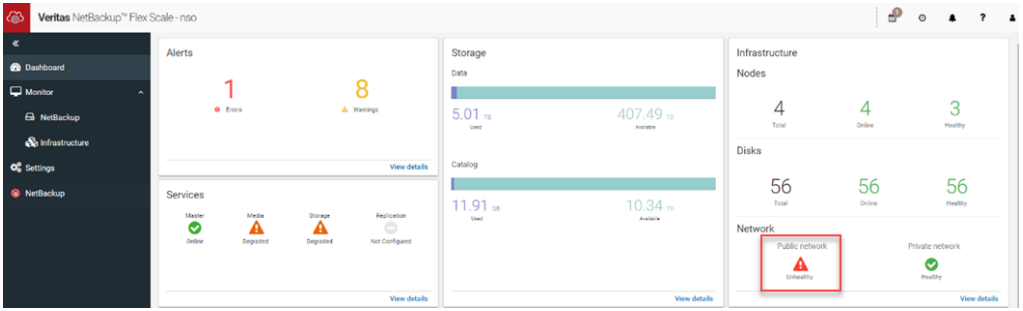

In the iLO remote console, navigate to **System information > Network**. The status for the affected port is shown **Link down**.

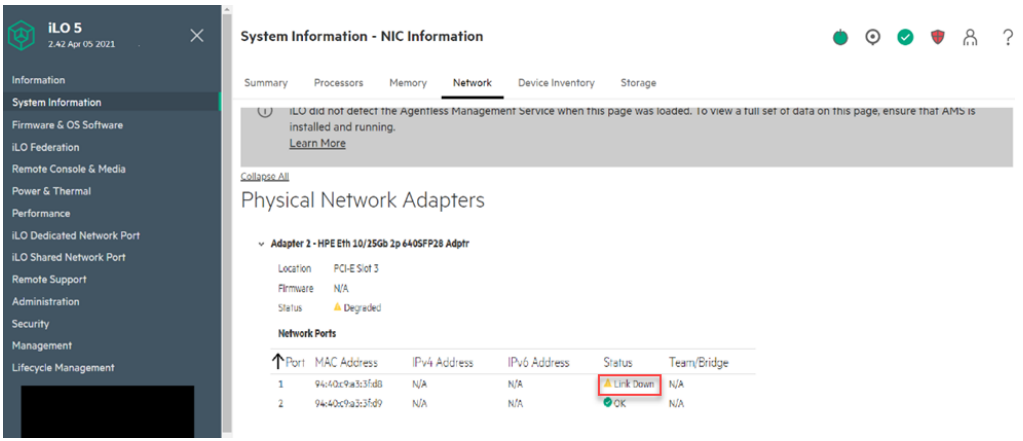

Plug in the cable in the affected port. Review the cabling guidelines before you plug in the cable. After plugging in the eth5 cable, the node status is shown healthy. Navigate to**Monitor > Infrastructure > Nodes**:

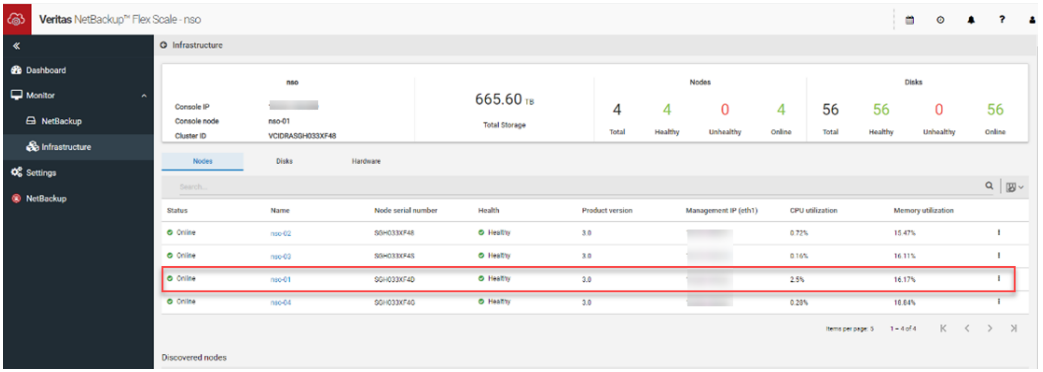

The eth5 status changes to **Plugged** on the **Hardware** tab. Navigate to **Monitor > Infrastructure > Hardware > Network Card**:

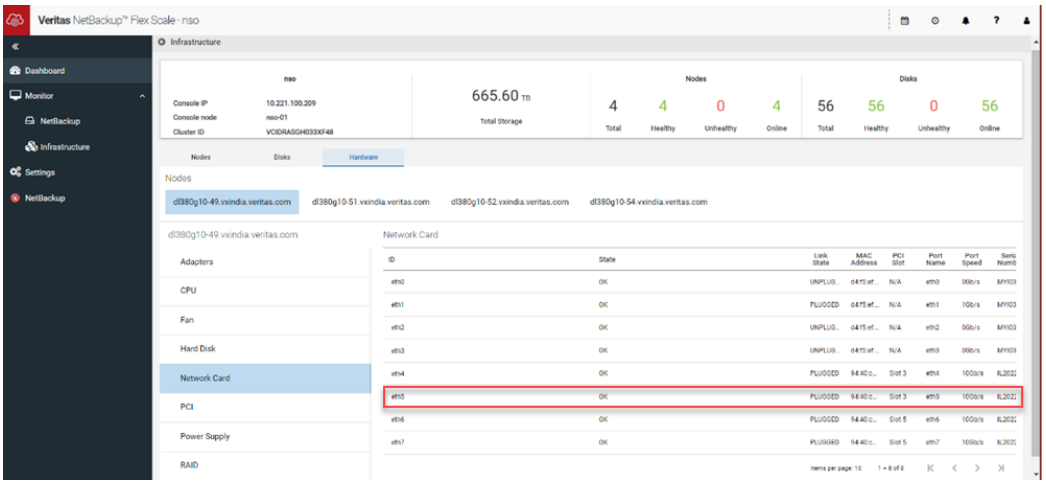

In the iLO remote console, navigate to **System information > Network**. The status for the affected port is shown **OK**.

Adapter 2 - HPE Eth 10/25Gb 2p 640SFP28 Adptr

Location	PCI-E Slot 3
Firmware	N/A
Status	OK

Network Ports

Port	MAC Address	IPv4 Address	IPv6 Address	Status	Team/Bridge
1	94:40:c9:a3:3f:d8	N/A	N/A	OK	N/A
2	94:40:c9:a3:3f:d9	N/A	N/A	OK	N/A

On the **Dashboard**, the public network status is shown healthy. It might take approximately five minutes to refresh the status in the UI.

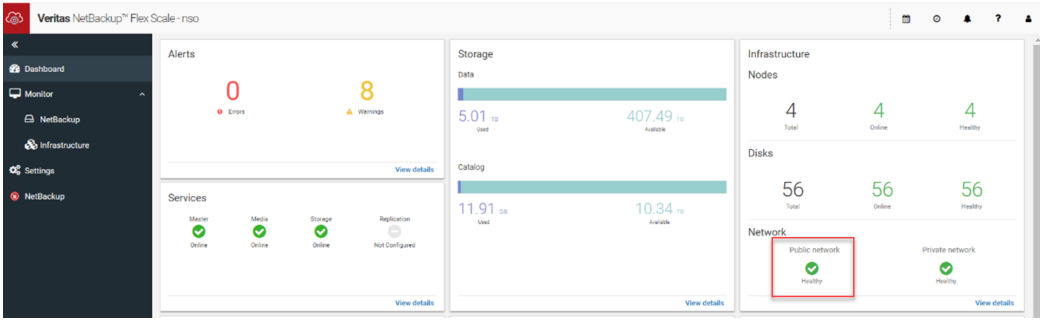

If the issue persists it implies that the SFP port is faulty. The CHS team will need to involve HPE for replacement of the port. Contact Veritas TSE to replace the node with the faulty iLO port.

Collecting HPE Active Health System (AHS) logs (performed by Veritas TSE)

Before you contact the hardware vendor for replacing the failed component, collect AHS logs. To collect the AHS logs, in the NetBackup Flex Scale infrastructure management UI, navigate to **Settings > Diagnostics > Basic > Appliance**.

Noting the MAC IDs (performed by Veritas TSE)

From the node-level CLI use the `system hardware-health` command to note the eth4 and eth5 MAC addresses.

```
| eth4 | Slot 3   | 817751-001 | IL202202L6    | FIBRE        | 10Gb/s     | 94:40:c9:a3:3f:d9
+------+----------+------------+---------------+--------------+------------+------------------
| eth5 | Slot 3   | 817751-001 | IL202202L6    | FIBRE        | 0Gb/s      | 94:40:c9:a3:3f:d8
```

Shutting down the node (performed by Veritas TSE)

Before an HPE representative can replace the SFP port, you must shut down the node.

To shut down the node:

1. Sign in to the NetBackup Flex Scale infrastructure management UI and navigate to **Monitor > Infrastructure > Nodes**.
2. On the node where the failure occurred, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Shutdown node**.

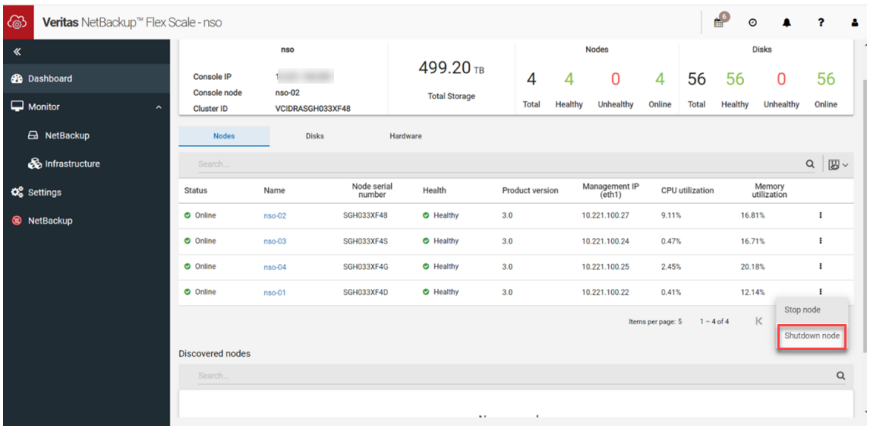

3. Confirm that the node is shut down successfully. In the UI, you can view the notification at the top of the page.

Replacing the SFP port (performed by the HPE vendor)

The HPE representative replaces the SFP port.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

To verify that the issue is resolved, complete the following steps:

1 Restart the node from the iLO remote console using the **Power > Momentary Press** option.

The green color power symbol indicates that the node has started.

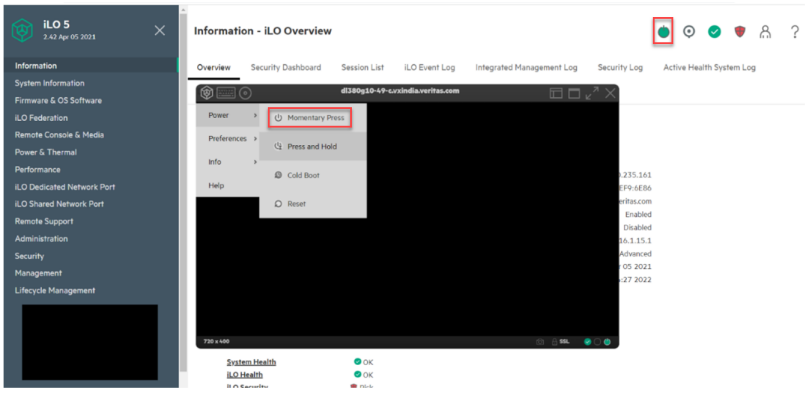

2 As the port is changed physically, the MAC address is also changed, which results in the following cases:

The node status is shown unhealthy in the UI:

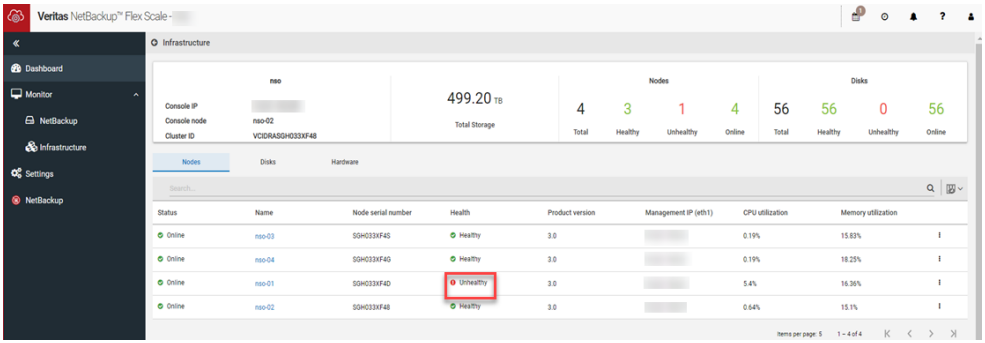

An event is shown on the **Settings > Events** page of the UI:

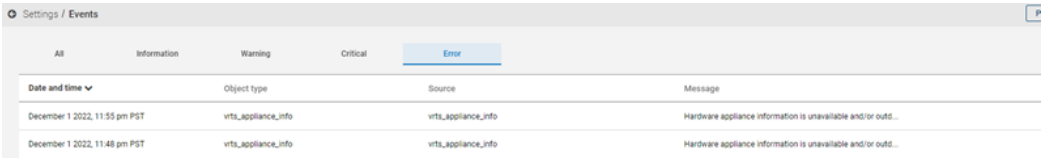

From the node-level CLI, the `system hardware-health` command shows that eth4 was plugged and eth5 was unplugged:

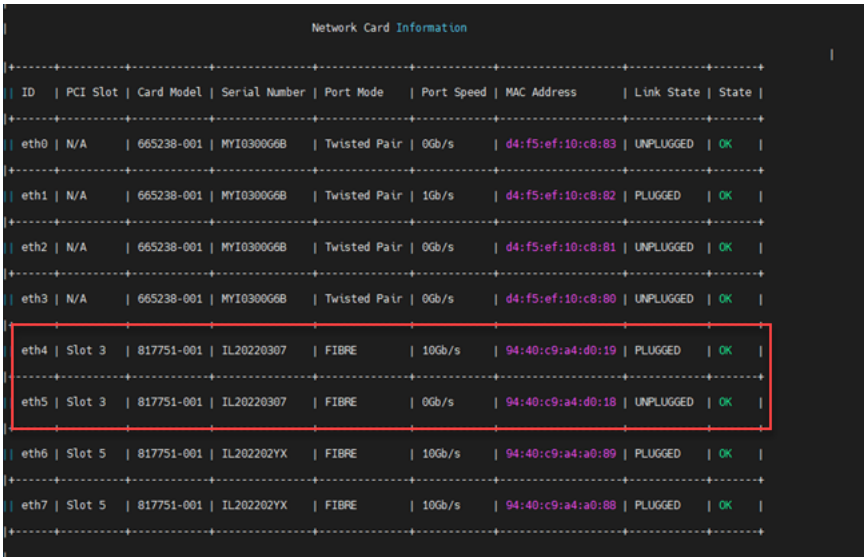

Network Card Information

ID	PCI Slot	Card Model	Serial Number	Port Mode	Port Speed	MAC Address	Link State	State
eth0	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:83	UNPLUGGED	OK
eth1	N/A	665238-001	MYI0300G6B	Twisted Pair	1Gb/s	d4:f5:ef:10:c8:82	PLUGGED	OK
eth2	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:81	UNPLUGGED	OK
eth3	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:80	UNPLUGGED	OK
eth4	Slot 3	817751-001	IL20220307	FIBRE	10Gb/s	94:40:c9:a4:d0:19	PLUGGED	OK
eth5	Slot 3	817751-001	IL20220307	FIBRE	0Gb/s	94:40:c9:a4:d0:18	UNPLUGGED	OK
eth6	Slot 5	817751-001	IL202202YX	FIBRE	10Gb/s	94:40:c9:a4:a0:89	PLUGGED	OK
eth7	Slot 5	817751-001	IL202202YX	FIBRE	10Gb/s	94:40:c9:a4:a0:88	PLUGGED	OK

Click **Dashboard**. The public network status is shown unhealthy:

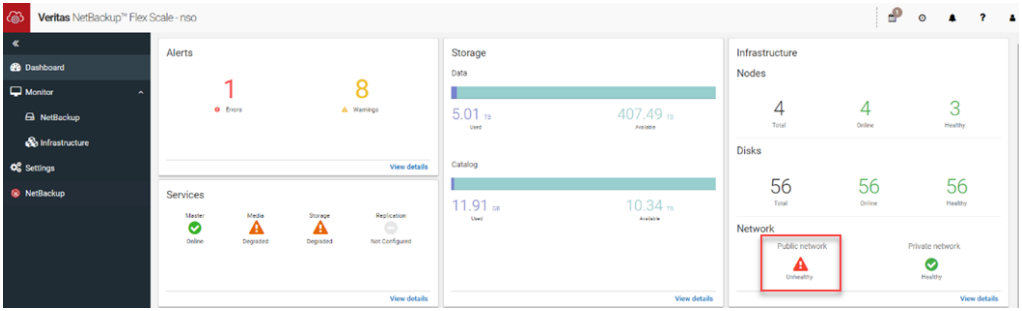

In the iLO remote console go to **System information > Network**:

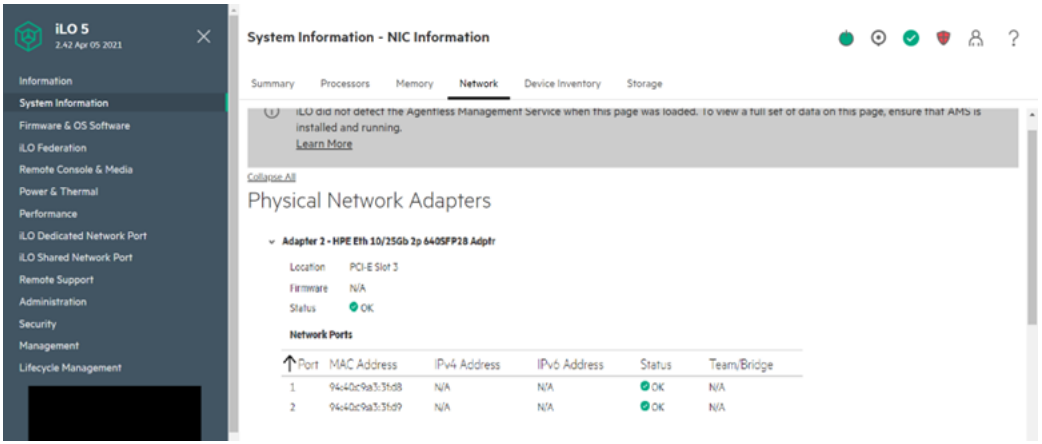

3 If you are on version 3.0 or earlier, complete the following steps to bring up eth4 and eth5. You need to elevate to root access to perform these steps.

- SSH to another cluster node:

  ```
  ssh primary@node_name
  ```

- `su`
- Enter the maintenance password.
- Open the `/etc/sysconfig/network-scripts/ifcfg-eth1` file and you will see old MAC ID in the **HWADDR** field. Change the MAC ID in the **HWADDR** field to the new MAC ID.
- Run following commands:

  ```
  ifdown eth4
  ifup eth4
  ```

Repeat the same steps for eth5.

After completing the steps for eth4 and eth5, navigate to **Settings > Services management** and click **Run full discovery**.

4 Verify that the node status is shown healthy in the UI:

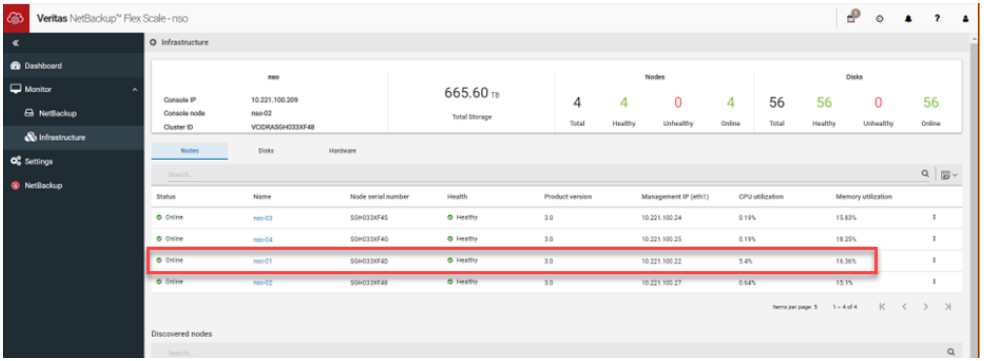

5 The public network status is shown healthy:

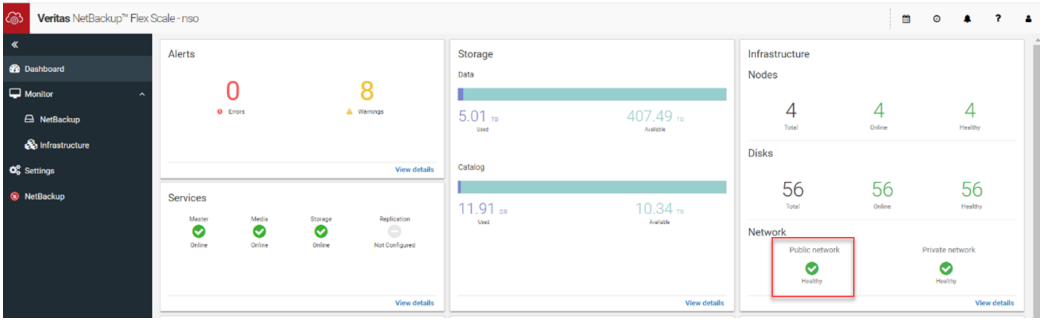

6 Verify that eth4 and eth5 both are shown **PLUGGED** on the **Hardware** tab.

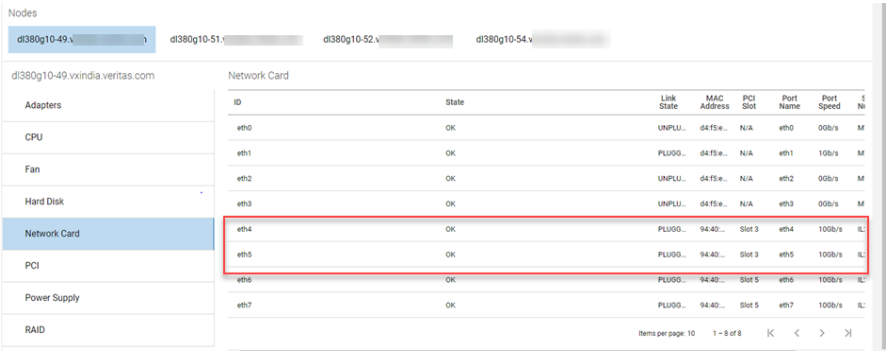

7 Verify that the changed MAC ID can be seen in the `system hardware-health`and eth4 and eth5 are shown plugged:

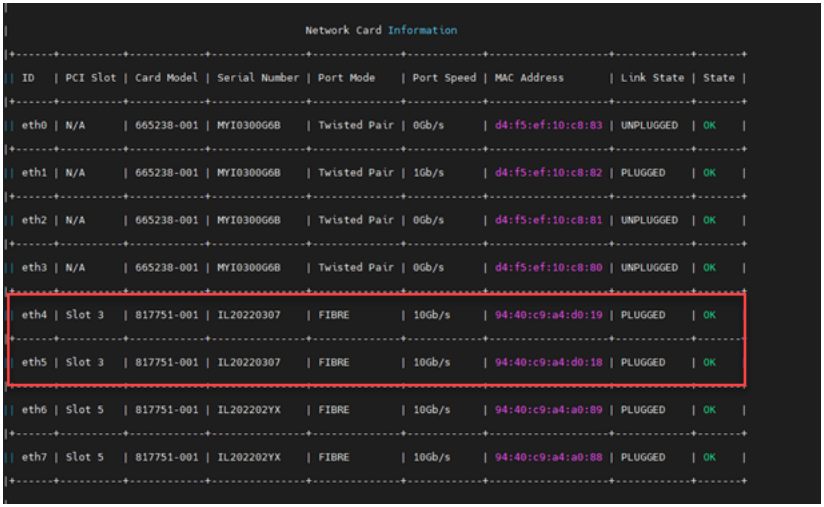

Network Card Information

ID	PCI Slot	Card Model	Serial Number	Port Mode	Port Speed	MAC Address	Link State	State
eth0	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:83	UNPLUGGED	OK
eth1	N/A	665238-001	MYI0300G6B	Twisted Pair	1Gb/s	d4:f5:ef:10:c8:82	PLUGGED	OK
eth2	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:81	UNPLUGGED	OK
eth3	N/A	665238-001	MYI0300G6B	Twisted Pair	0Gb/s	d4:f5:ef:10:c8:80	UNPLUGGED	OK
eth4	Slot 3	817751-001	IL20220307	FIBRE	10Gb/s	94:40:c9:a4:d0:19	PLUGGED	OK
eth5	Slot 3	817751-001	IL20220307	FIBRE	10Gb/s	94:40:c9:a4:d0:18	PLUGGED	OK
eth6	Slot 5	817751-001	IL202202YX	FIBRE	10Gb/s	94:40:c9:a4:a0:89	PLUGGED	OK
eth7	Slot 5	817751-001	IL202202YX	FIBRE	10Gb/s	94:40:c9:a4:a0:88	PLUGGED	OK

8 An event is generated notifying that eth4 and eth5 are online:

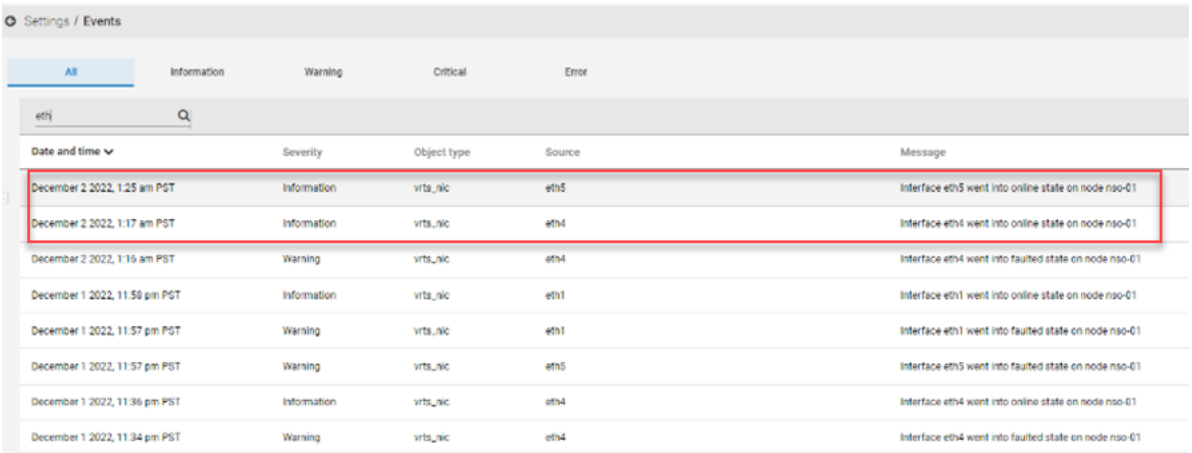

Settings / Events

All | Information | Warning | Critical | Error

eth

Date and time	Severity	Object type	Source	Message
December 2 2022, 1:25 am PST	Information	vrts_nic	eth5	Interface eth5 went into online state on node nso-01
December 2 2022, 1:17 am PST	Information	vrts_nic	eth4	Interface eth4 went into online state on node nso-01
December 2 2022, 1:16 am PST	Warning	vrts_nic	eth4	Interface eth4 went into faulted state on node nso-01
December 1 2022, 11:58 pm PST	Information	vrts_nic	eth1	Interface eth1 went into online state on node nso-01
December 1 2022, 11:57 pm PST	Warning	vrts_nic	eth1	Interface eth1 went into faulted state on node nso-01
December 1 2022, 11:57 pm PST	Warning	vrts_nic	eth5	Interface eth5 went into faulted state on node nso-01
December 1 2022, 11:36 pm PST	Information	vrts_nic	eth4	Interface eth4 went into online state on node nso-01
December 1 2022, 11:34 pm PST	Warning	vrts_nic	eth4	Interface eth4 went into faulted state on node nso-01

Replacement procedure for chassis

This topic describes the process for replacing a chassis in a NetBackup Flex Scale node.

Identifying a chassis failure (performed by the CHS team)

The node where the chassis malfunctions is shown as Unhealthy and the status is shown Offline. In the NetBackup Flex Scale infrastructure management UI, navigate to **Monitor > Infrastructure > Nodes** to view the health of the nodes.

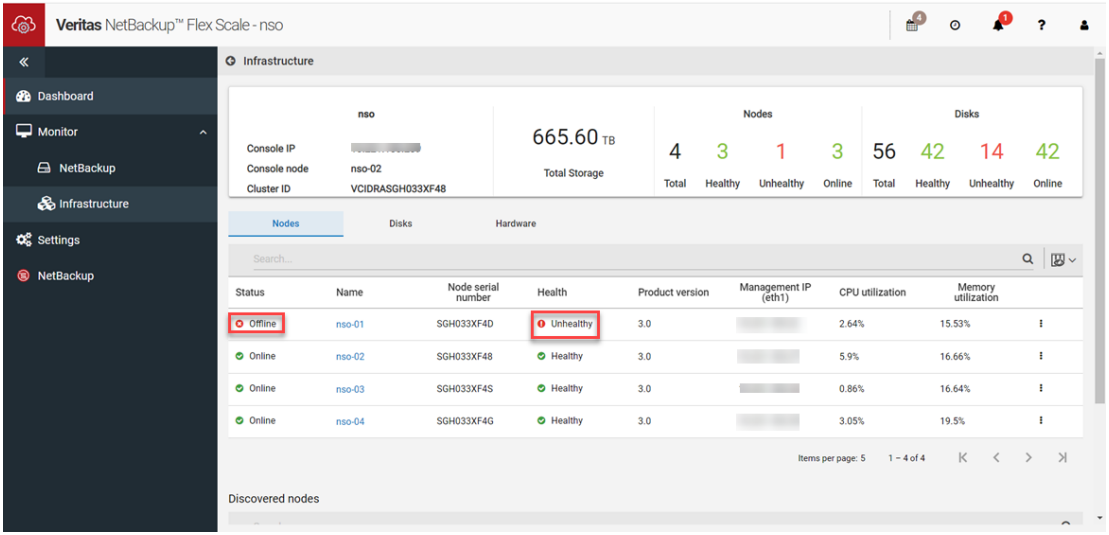

Shutting down the node (performed by Veritas TSE)

You need to shut down the node on which the chassis is malfunctioning. For a cluster with fewer than six nodes, only a single node can be down, stopped, or shut down at any given point in time. For a larger cluster of up to 16 nodes, a maximum of two nodes can be down, stopped, or shut down at any given point in time. When you shut down a node, the cluster services running on the node are stopped and the NetBackup jobs running on the node fail over to other cluster nodes. After the hardware maintenance is complete, you need to turn on 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 shut down the node where the NetBackup Flex Scale infrastructure management UI is running, it fails over to another node. It can take a few minutes for the UI to be up on another node

---

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

To shut down the node:

1. Shut down the node where the chassis failed. Navigate to **Monitor > Infrastructure > Nodes** and click **Shutdown node**.
2. Confirm that the node is shut down successfully. In the UI, you can view the notification at the top of the page.
3. Contact the hardware vendor to replace the hardware component.

Collecting HPE Active Health System (AHS) logs (performed by Veritas Support)

Before you contact the hardware vendor for replacing the failed component, collect AHS logs. To collect the AHS logs, in the NetBackup Flex Scale infrastructure management UI, navigate to **Settings > Diagnostics > Basic > Appliance**.

Replacing the chassis (performed by the HPE vendor)

The HPE representative replaces the chassis.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

Perform the node replacement operation. When you select the node that you want to use to replace the unhealthy node, ensure that you select the node on which the chassis is replaced.

To replace the node:

1. Check the node status in the NetBackup Flex Scale infrastructure management UI. The node is shown unhealthy and the node status is offline.
2. Prepare the node for replacement:
   - Deploy ISO on the failed node where the OS disks were replaced.
   - Perform factory reset on that node from the node CLI by using the `system factory-reset` command.
3. From the NetBackup Flex Scale infrastructure management UI perform the replace node operation.
   - Scan for nodes. Click **Scan for nodes** to discover the nodes.
   - Select a node from the displayed list and replace the node. For the unhealthy node, click the Actions menu (vertical ellipsis) from the right side of the row in the UI, and click **Replace node**.

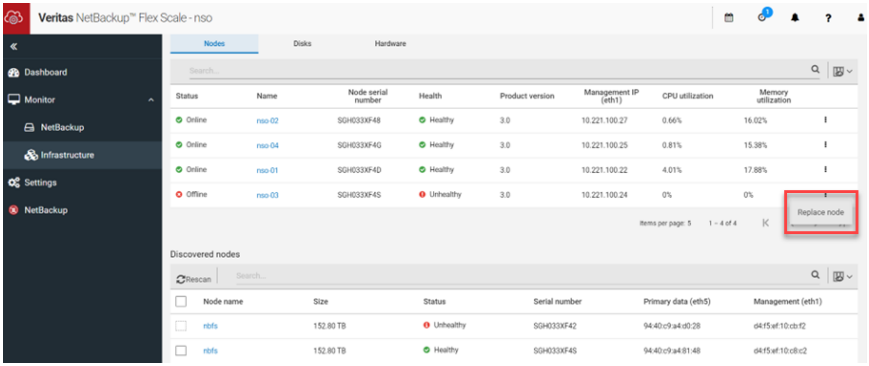

- In the Replace node dialog box, select the node that you want to use to replace the unhealthy node and click **Replace node**.

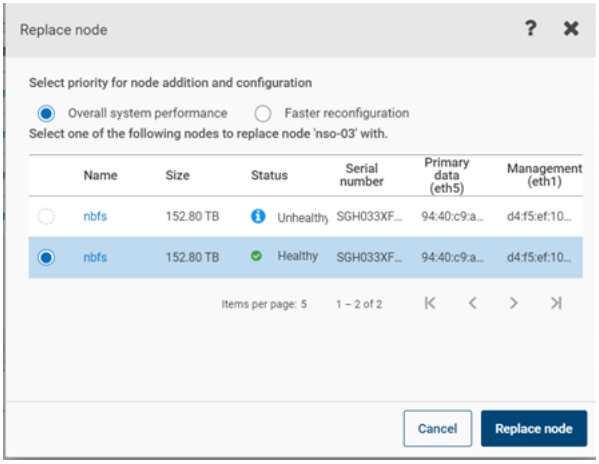
Replace node
Select priority for node addition and configuration
Overall system performance
Faster reconfiguration
Select one of the following nodes to replace node 'nso-03' with.
Name
Size
Status
Serial number
Primary data (eth5)
Management (eth1)
nbfs
152.80 TB
Unhealthy
SGH033XF...
94:40:c9:a...
d4:f5:ef:10...
nbfs
152.80 TB
Healthy
SGH033XF...
94:40:c9:a...
d4:f5:ef:10...
Items per page: 5
1 – 2 of 2
Cancel
Replace node

4 After the replace node operation completes successfully check the node status in the NetBackup Flex Scale infrastructure management UI. Navigate to **Monitor > Infrastructure > Nodes**. The node is online and shown healthy.

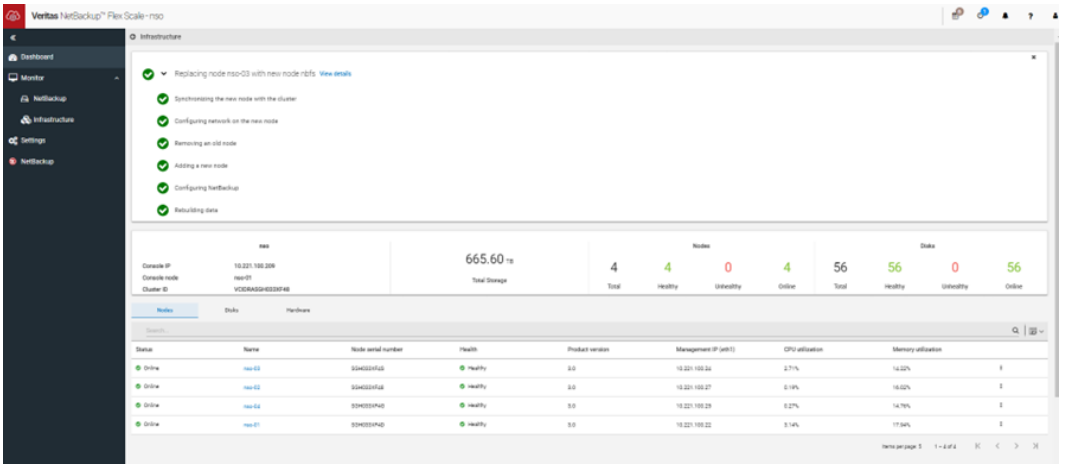

Click **Dashboard** and verify that all the NetBackup services are running.

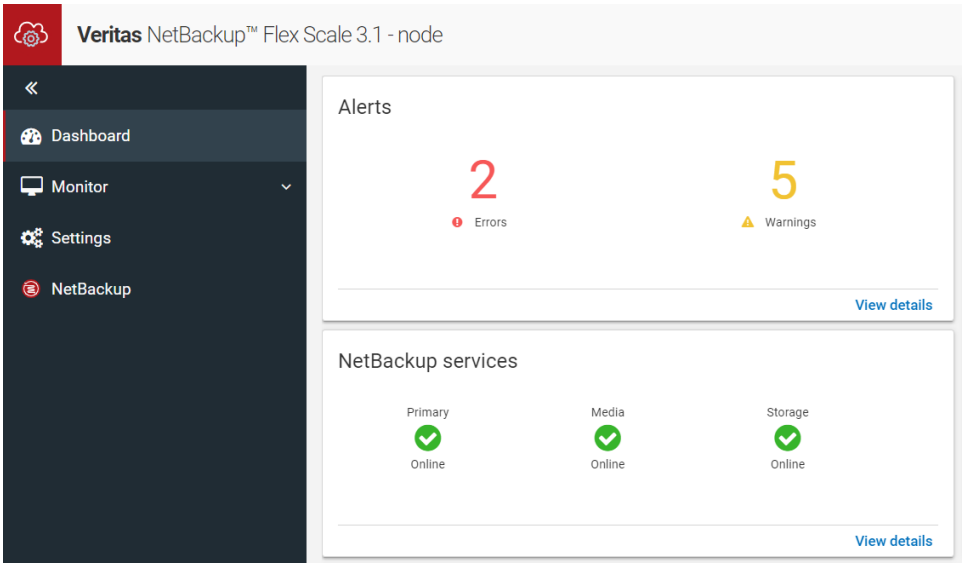

Replacement procedure for a hard disk drive

This topic describes the process for replacing a hard disk drive (HDD) in a NetBackup Flex Scale node. There are 12 hard disks of 16 TB or 20 TB based on your model. If the HDD is located in mid bay, you need to shut down the node where the faulty HDD is located.

Identifying a hard disk drive failure (performed by the CHS team)

From the node-level CLI, use the `show hardware-health` command to check the status of the HDD. If there is some problem with a HDD, the status of that disk is shown **NOT-OK**.

Replacement the HDD (performed by the HPE vendor)

An HPE representative identifies the faulty HDD, its physical location in the appliance, and replaces the faulty HDD.

To replace the HDD that is not located in mid bay, the HPE representative completes the following steps:

1 Check the HDD serial number and the slot from the iLO remote console.

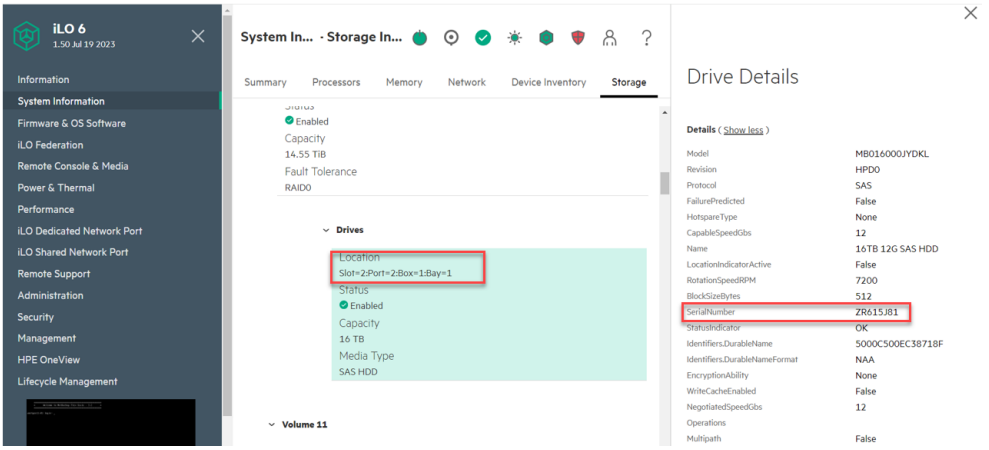

2 Identify the corresponding location of the HDD in the appliance.

3 Check the model number of the new disk to make sure it matches with the older one.

4 Replace the faulted HDD with another HDD.

To replace the HDD that is located in mid bay, the HPE representative completes the following steps:

1 Note the HDD location in the iLO.

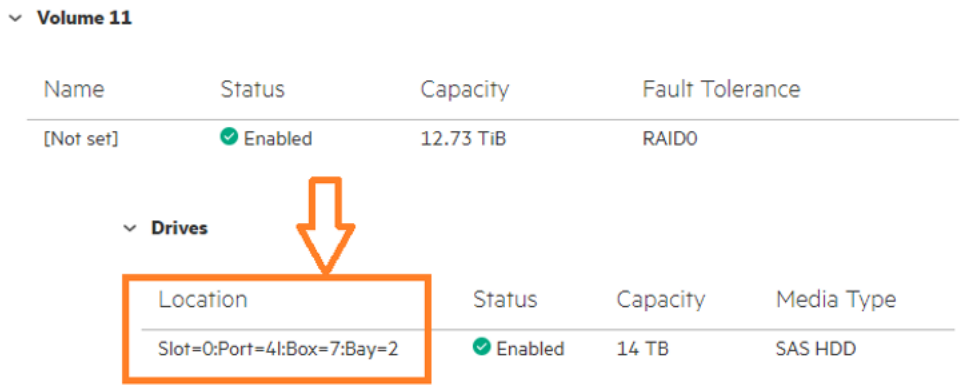

2 Find the corresponding physical location on the appliance.

3 Note the model number in iLO.

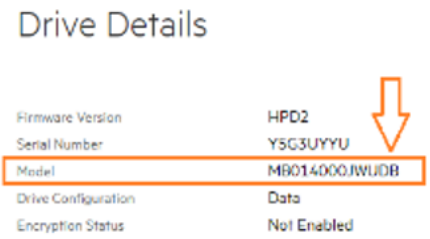

4 Check the model number of the new disk to make sure it matches with the older one.

5 Shut down the node where the faulty disk is present.

6 Confirm that the node is shut down.

A notification is displayed on the top of the page.

7 Power off the node by pressing the power button.

8 Disconnect both the power cables and any network cables that prevents the node from sliding forward. Mark the cables so they can be put back in the same port.

9 Pull the node forward in the rack to expose the mid bay.

10 Replace the faulty disk with the new disk.

11 Connect both the power cables back and any network cables that were unplugged.

12 Power on the node by pressing the power button on the front of the server.

13 Wait till the node joins back the cluster. It takes approximately 20 minutes for the node to join the cluster. The status is shown in the GUI.

14 Elevate to root prompt of the management console node and check the recovery task using the `vxtask list`. Wait till recovery is complete.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

To verify that the issue is resolved, Veritas TSE completes the following steps:

1. Replace the old, faulted disk with new disks. In the NetBackup Flex Scale infrastructure management UI, in the left pane click **Monitor > Infrastructure**.
2. Click **Disks**.
3. Click the faulty disk that you want to replace, and then click **Replace disk**.

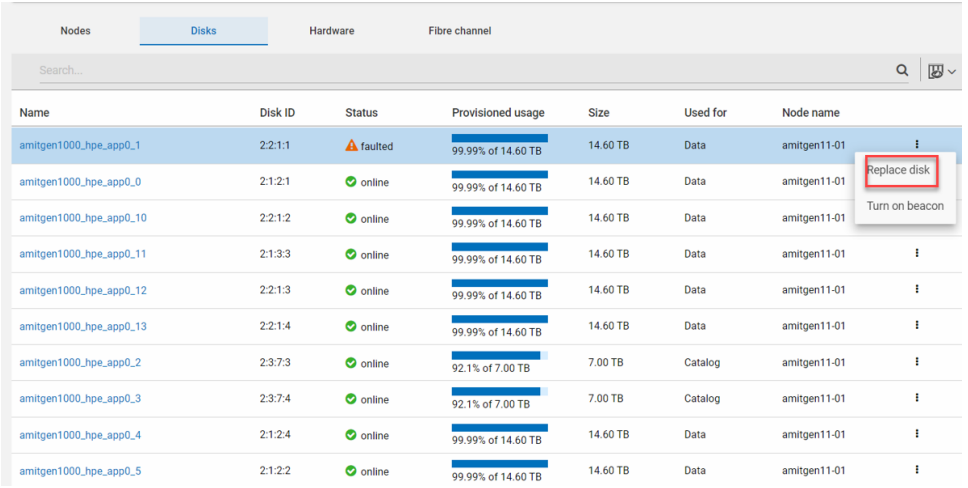

The **Disk replacement details** shows the disk that is being replaced. Wait till the status is shown **Completed**.

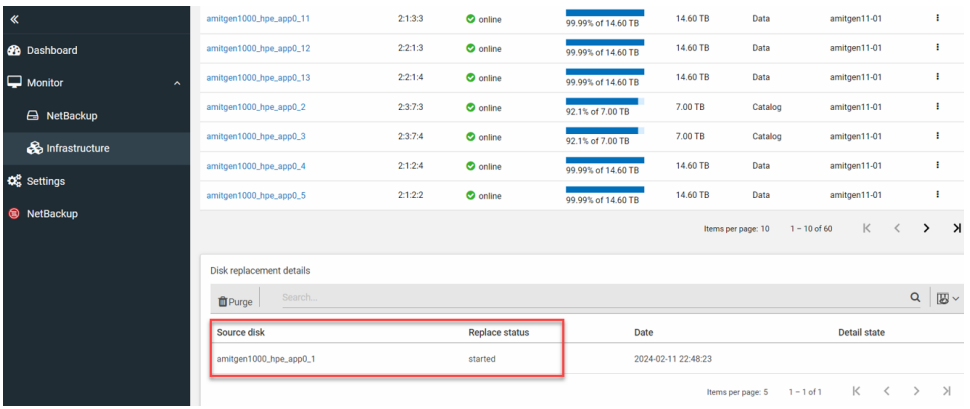

Replacement procedure for a Fibre Channel card for a cluster node

Use the following procedure to replace a Fibre Channel card when the node is a part of the cluster.

Verifying if the Fibre Channel ports are assigned (performed by Veritas TSE)

Log in to the UI and verify if the Fibre Channel ports are assigned to a workload.

To check if the Fibre Channel ports are assigned to a workload:

1. Sign in to the NetBackup Flex Scale infrastructure management UI and navigate to **Monitor > Infrastructure > Fibre channel**.
2. If the port is assigned, the **Assigned** column displays **Yes** and the **Workload** column shows the workload to which port is assigned.

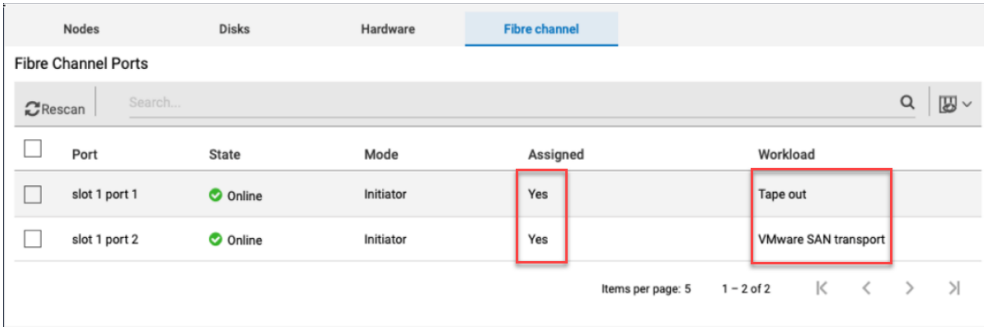

If the ports are unassigned, the **Assigned** column shows **No** and no workload is shown in the **Workload** column.

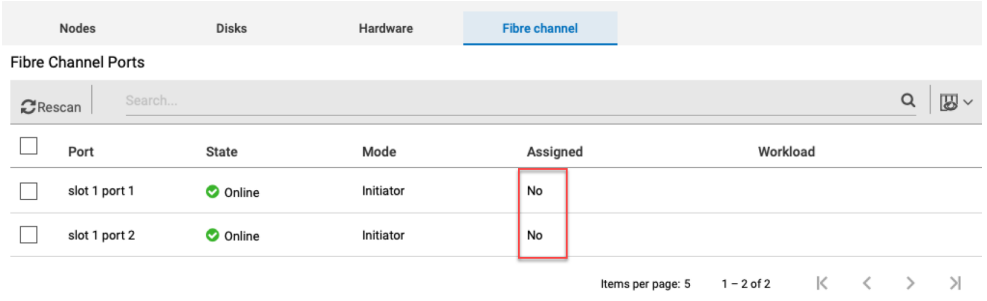

3 If ports are already unassigned continue to the next step of shutting down the node.

4 If the ports are assigned, unassign the ports. Select the assigned ports and click **Unassign port**.

Wait for the unassign operation to complete.

Shutting down the node (performed by Veritas TSE)

Before an HPE representative can replace the Fibre Channel card, you must shut down the node.

To shut down the node:

1 Sign in to the NetBackup Flex Scale infrastructure management UI and navigate to **Monitor > Infrastructure > Nodes**.

2 On the node where the you want to replace the card, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Shutdown node**.

3 Confirm that the node is shut down successfully. In the UI, you can view the notification at the top of the page.

4 Press the Power button on the front panel of the server to power off the node.

Replacing the Fibre Channel (performed by HPE)

The HPE representative replaces the Fibre Channel card for 5561 model.

The HPE representative replaces the Fibre Channel card for 5551 model.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

To verify that the issue is resolved, complete the following steps:

1 Power on the node from corresponding iLO interface and wait till the node joins the cluster.

Wait till the node and all the disks in the node are healthy.

2 Create the zoning again with new Fibre Channel card WWPN (World Wide Port Name) .

3 Ensure that all the cluster nodes have a Fibre Channel card in the correct slot. For a 5561 model, the card must be inserted in Slot-5. For a 5551 model, the card must be inserted in Slot-4.

4 Sign in to the NetBackup Flex Scale infrastructure management UI and navigate to **Monitor > Infrastructure**.

5 Ensure that the **Fibre channel** tab is displayed in the UI.

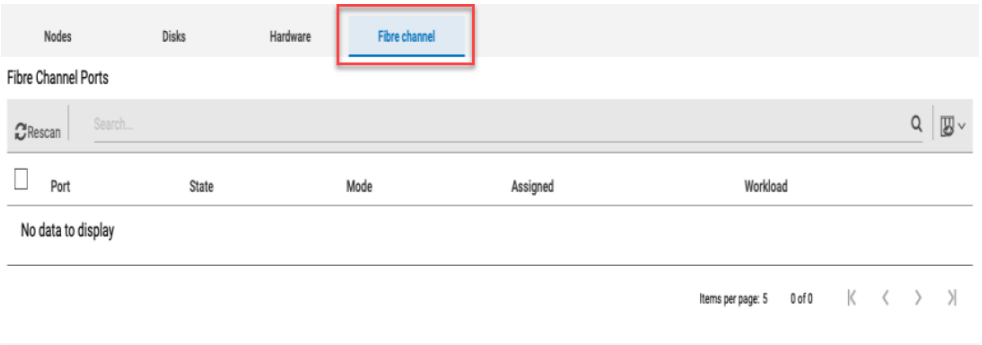

6 Click **Fibre channel** and click **Rescan** to scan connected Fibre Channel cards.

Details about the Fibre Channel ports, such as their state, mode, and assignment are displayed. The state is shown as online, the mode as initiator, and the port is not be assigned to any workload.

You can now assign the Fibre Channel ports to workloads. Before assigning the ports, select the port and click **Discover devices** to discover devices.

Replacement procedure for a Fibre Channel card for a node that is not in a cluster

Use the following procedure to replace a Fibre Channel card when the NetBackup Flex Scale ISO is installed on the node but it is not a part of the cluster.

Shutting down the node (performed by Veritas TSE)

Shut down the node from the iLO remote console of the node by using the **Server Power > Press and Hold** option.

Replacing the Fibre Channel (performed by HPE)

The HPE representative replaces the Fibre Channel card for 5561 model.

The HPE representative replaces the Fibre Channel card for 5551 model.

Completing the post-replacement tasks (performed by Veritas TSE)

After the hardware vendor notifies you that the hardware component is replaced, verify that the issue is resolved.

Power on the node from the iLO remote console by using the **Power > Momentary Press** option. The green color power symbol indicates that the node has started.

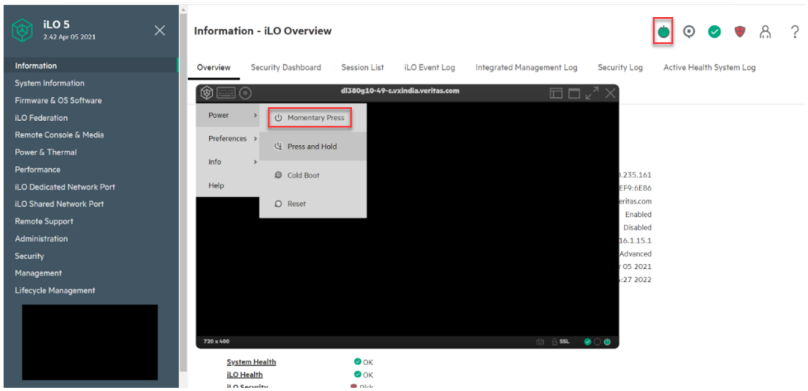

After restarting the node, enable the Fibre Channel BOM configuration from the Appliance Node-level CLI.

To enable Fibre Channel BOM:

1. Use SSH Login to Node level CLI using the eth1 IP address of the node
2. Run the following command:

   ```
   support bom-conf get
   ```

```
[nbfs-3.2] nbfs >
[nbfs-3.2] nbfs > support bom-conf get

The BOM configuration file is copied to /system/inst/patch/incoming, please run support share open to access it

Operation completed successfully
```

3. To update this `bom-config.json` file open an NFS share by running the following command:

   `system software share open`

```
[nbfs-3.2] nbfs > system software share open

- [Info] Created an NFS share for sharing the patches.

- [Info] You can access the NFS share at 10.221.221.59:/system/inst/patch/incoming. To ensure appliance security
, use the 'system software share close' command to remove the share after downloading the required patches.

[nbfs-3.2] nbfs >
```

4. Mount the above open NFS share on any available Linux Client by using the Linux command:

   `mount -t nfs` *`ipaddressorfqdn`*`:/system/inst/patch/incoming /mnt`

   where *ipaddressorfqdn* is the IP address or the FQDN that was displayed when you ran the `system software share open` command.

   The `/mnt` location can be changed to the required location for mounting the NFS share.

   The `bom-config.json` file is copied to the `/mnt` location.

5. Set the **FC-ENABLE** key value to **1**. Open the `bom-config.json` using the `vim` command:

   `vim bom-config.json`

   Change the **FC-ENABLE** key to 1 and save the file. the updated file should look similar to shown below:

```
nbfs:/mnt # ls
bom-config.json
nbfs:/mnt # cat bom-config.json | grep -w FC-ENABLE
  "FC-ENABLE": "1",
nbfs:/mnt #
```

6. Log in to the Appliance Node-level CLI again using the eth1 IP address of the node.

7 Run the following command to update configuration:

```
support bom-conf update
```

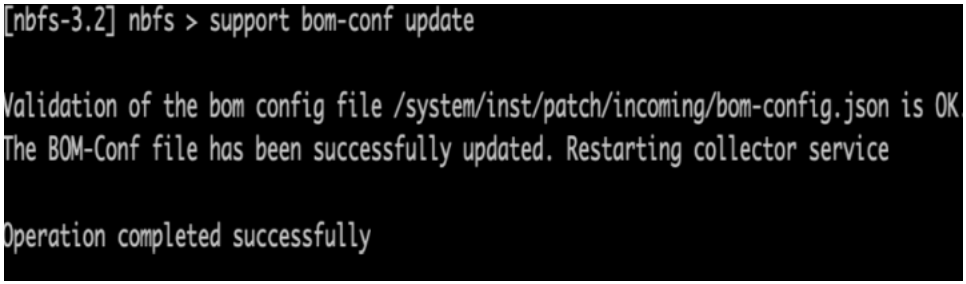

8 After the file is updated, close the open NFS share:

```
system software share close
```

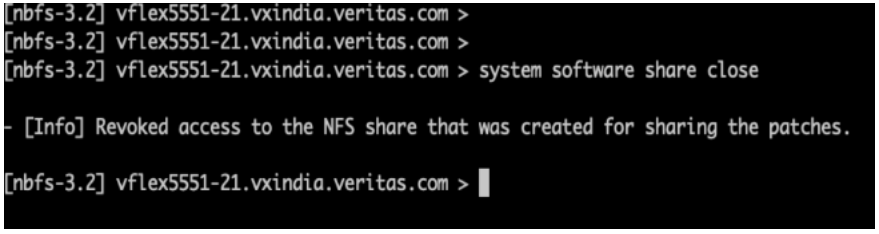

After you enable the Fibre Channel BOM, create the zoning again with new Fibre Channel card WWPN (World Wide Port Name).

Appendix B

Configuring NetBackup optimized duplication

This appendix includes the following topics:

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.

5 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.

6 Click **OK** to save the **New Operation** window parameter settings.

7 Click **Add** to add a second operation to the SLP. The **New Operation** window is displayed.

8 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.

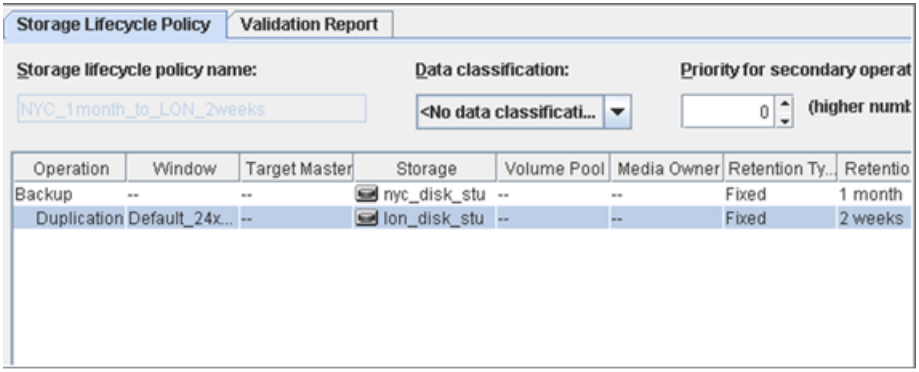

Note: The left margin of the Duplication operation is indented from the entry for the first backup operation which indicates that the Duplication operation is derived from the second backup operation.

10 Click **OK** in the **Change Storage Lifecycle Policy** window to save the changes to the NYC_1month_to_LON_2weeksSLP, 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 <*primary 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

Appendix C

Disaster recovery terminologies

This appendix includes the following topics:

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 `passthru` 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 `vxrlink recover` 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 `needs_recovery` 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
host: Pri or Sec IP not available or vradmind not running	The Primary IP or the Secondary IP address is not available, or the `vradmind` daemon on the host is not running or is running on a different port.
host: disk group missing.	Host *host* does not have any disk group with the same name as that specified in the *remote_dg* attribute, of the Primary RLINK pointing to this host.
host: RLINK missing.	Primary RVG has an RLINK to *host*, but *host* does not have corresponding rlink to this Primary RLINK.
host: RLINK dissociated.	Host *host* does have an RLINK corresponding to the Primary RLINK. However, it is not associated with the Secondary RVG.
host: disk-group mismatch.	The *remote_dg* attribute of either the Primary RLINK or Secondary RLINK is incorrect.
host: RLINK mismatch.	The *remote_rlink* attribute of either the Primary RLINK or Secondary RLINK is incorrect.
host: host mismatch.	The *local_host* and/or *remote_host* attribute of either the Primary RLINK or Secondary RLINK is incorrect.
host: multiple Primary error.	The same Secondary RVG has more than one Primary RVGs.
host: two or more nodes on same host.	Two or more RVGs in the same RDS are located on the same host, *host*. This configuration is not supported.
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.

Message	Message definition
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 `vradmind` server on the Secondary host is unable to establish contact with the `vradmind` server on the Primary. This can be because the Primary RVG of this RDS cannot be found or the `vradmind` 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.
host: 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.
host: Primary SRL missing	This error can occur if the Primary SRL was disassociated from the Primary RVG or is missing.
host: Secondary SRL missing	This error can occur if the Secondary SRL was disassociated from the Secondary RVG or is missing.

Appendix D

Configuring Auto Image Replication

This appendix includes the following topics:

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**.

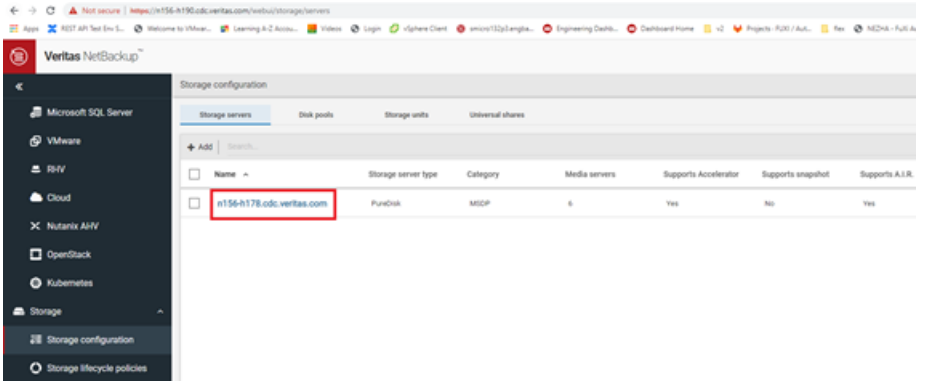

3. Add MSDP_SERVER in target primary server. In the target domain, logon to the target primary server using the following command:

```
ssh <backup_admin user>@<primary ip/hostname>
(appliance_admin password)
sudo bash
echo "MSDP_SERVER = <Source MSDP server name>" >>
/usr/openv/netbackup/bp.conf
```

4 If both domains are configured with ECA, skip this step.

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**.

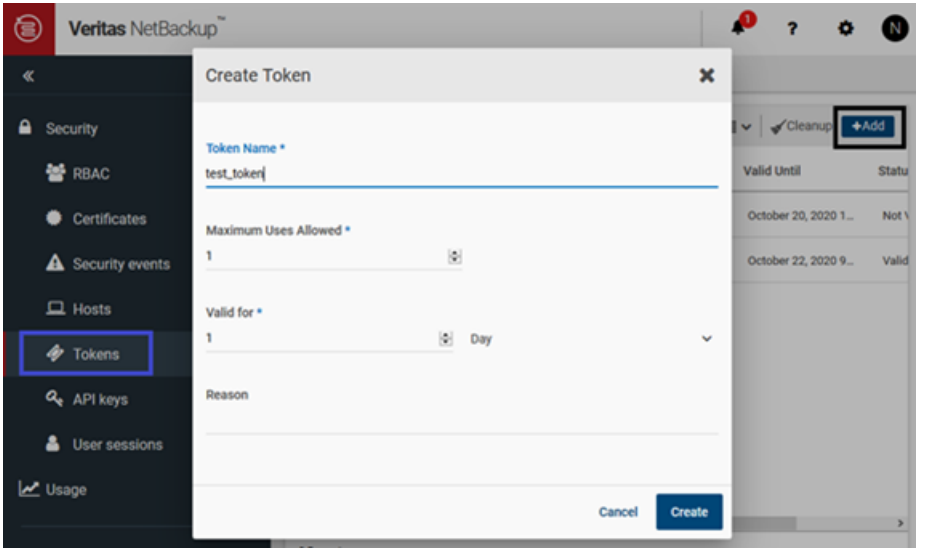

5 Add replication targets for the disk pool.

In the **Disk pools** tab, click on the disk pool link.

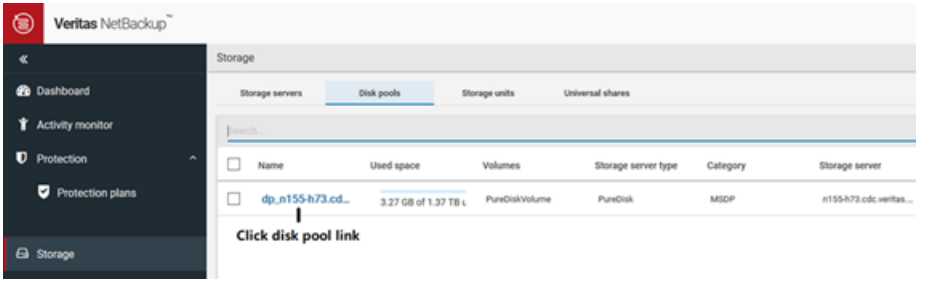

Click **Add** to add the replication target.

6 In the **Add replication targets** window:

- Select the target primary server.
- Input the target domain token.

  Note: If both domains are configured with ECA, the target's token is not required.

- Select the target volume.
- Input the target storage credentials. Refer to the YAML configuration file for details.
  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
```

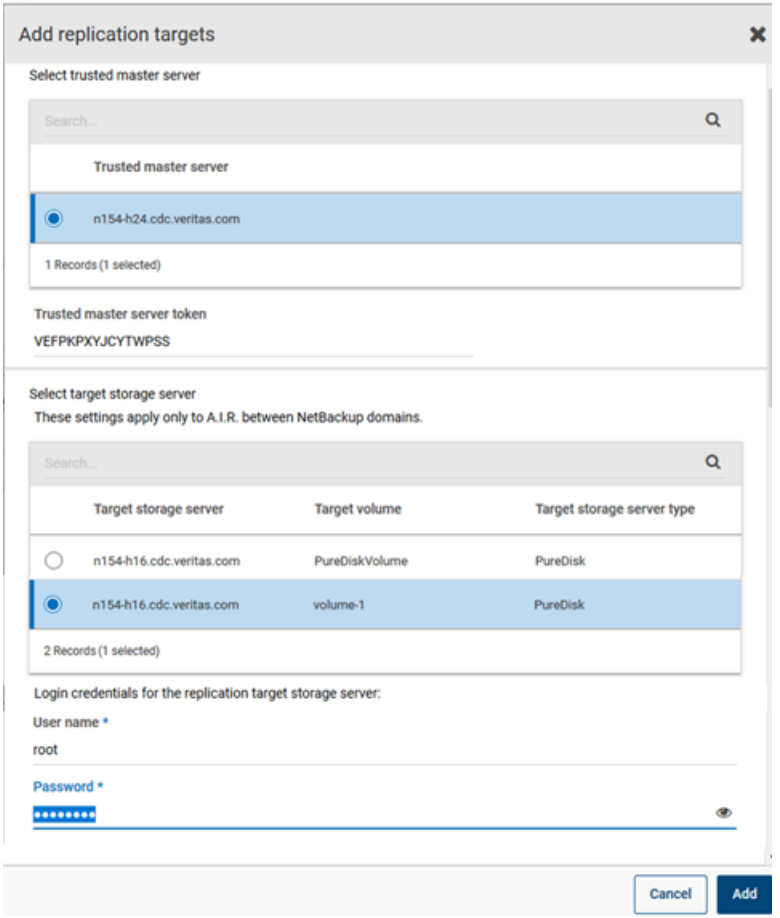

Click **Add**.