Veritas NetBackup™ Flex Scale Release Notes

3.2

April 2025

VERITAS™

Veritas NetBackup Flex Scale Release Notes

Last updated: 2025-04-28

Legal Notice

VERITAS TECHNOLOGIES LLC
2625 Augustine Drive
Santa Clara, CA 95054

http://www.veritas.com

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

NB.docs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1

Getting help

This chapter includes the following topics:

About this document

This document provides information specific to the Veritas NetBackup Flex Scale 3.2 release. Review this document before using the product.

The information in this document supersedes all the information provided in other product-specific documents.

For information about the operating system, hardware, and other general requirements, refer to the *Veritas NetBackup Flex Scale Installation and Configuration Guide*.

You can download the latest version of this document from the Veritas Service and Operations Readiness Tools (SORT) web site at:

https://sort.veritas.com/documents

NetBackup Flex Scale resources

For information about NetBackup Flex Scale features, use cases, data sheets, white papers, and videos, refer to the following product page:

https://www.veritas.com/protection/netbackup/netbackup-flex-scale

User documentation

For information on supported platforms, software and hardware requirements, and installation and administration instructions, refer to the NetBackup Flex Scale documentation here:

- Veritas Support
  https://www.veritas.com/support/en_US.html
  Click the Documentation link, choose Appliances from under the Product filter, and then choose NetBackup Flex Scale to display the latest documentation.
- Veritas Services and Operations Readiness Tools (SORT)
  https://sort.veritas.com/documents
  Select the product and the platform and apply other filters to display the desired documentation.

Chapter 2

Features, enhancements, and changes

This chapter includes the following topics:

What's new in this release

This section lists the major new features and enhancements added in the 3.2 version of NetBackup Flex Scale.

Simplifying cluster deployment with an integrated precheck

Before you start the NetBackup Flex Scale cluster configuration, you can run a precheck from the GUI to validate the specified cluster configuration details. The precheck validates details such as:

- Network connectivity for the management, data, and private networks.
- Configuration details such as IP addresses, FQDNs, DNS, and NTP details.
- Syntax of the provided configuration details.
- Ports are open and accessible.
- For media-only deployment, details such as connectivity to the external primary server and NetBackup version compatibility.
- System hardware health.
- System consistency.

Running the precheck helps you to identify potential issues early on, enabling you to fix the issues before you start the cluster configuration.

Configuring multifactor authentication

NetBackup Flex Scale supports multifactor authentication, which 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. For more details, refer to the *Configuring multifactor authentication* chapter in the*Veritas NetBackup Flex Scale Administrator's Guide*.

Configuring single sign-on (SSO)

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. For more details, refer to the *Single Sign-On (SSO)* chapter in the *Veritas NetBacku Flex Scale Administrator's Guide*.

Bonding management interfaces during initial configuration

During initial configuration, you can create a bond for eth1 and eth2 management interfaces. Management network bonding provides increased network resilience. You can use the setup wizard or a YML-based configuration file to configure bonding during initial configuration.

For more information, see the *Configuring a bond device and assigning an IP address to the bonded device*, *Configuring NetBackup Flex Scale using the setup wizard*, and *Configuring NetBackup Flex Scale using a configuration file* sections of the *Veritas NetBackup Flex Scale Installation and Configuration Guide*.

Bonding operations on management network

After the initial configuration is complete, you can perform bonding operations on the management network. You can create, modify, display, and remove bonds. For more details, refer to the *Bonding operations on data network* section in the *Veritas NetBackup Flex Scale Administrator's Guide*.

Connecting eth7 network interface optional during initial configuration

If you do not plan to configure bonding for the data network during the initial configuration, you do not need to connect the eth7 interface to the public network switch; you only need to ensure that eth5 is in the connected state. Previously, during initial configuration, you were required to connect eth7 irrespective of whether

you planned to use. If you want to configure bonding or create a secondary data network using eth7, ensure that eth7 is connected before performing the operation.

For more information, see the *NetBackup Flex Scale configuration requirements* section in the *Veritas NetBackup™ Flex Scale Best Practices and Troubleshooting Guide* and the *NetBackup Flex Scale network management* chapter of the *Veritas NetBackup™ Flex Scale Administrator's Guide*.

Support for VMware and Tape out backups over Fibre Channel

A dual-port single Fibre Channel card in initiator mode can be installed on NetBackup Flex Scale cluster nodes. The Fibre Channel ports can be configured with specific workloads.

For more information see the *Managing the Fibre Channel ports* section of the *Veritas NetBackup™ Flex Scale Administrator's Guide*.

Support for Data Domain storage

You can enable Data Domain as a backup storage by installing an EEB on the NetBackup Flex Scale.

For more information, see the *Backing up data to Data Domain storage* section of the *Veritas NetBackup™ Flex Scale Administrator's Guide*.

Collecting time-based logs

You can collect logs for a specific duration for the **Infrastructure**, **OS**, **NetBackup primary service**, and **NetBackup media service** components. time-based logs reduce the time required for log collection and the size of the generated log package.

For more information, see the *Collecting logs for cluster nodes* section of the *Veritas NetBackup™ Flex Scale Administrator's Guide*.

Forwarding logs to an external server

You can forward syslogs to an external log server for further analysis and troubleshooting.

For more information, see the *Forwarding logs to an external server* section of the *Veritas NetBackup™ Flex Scale Administrator's Guide*.

Changes to the licensing model

The following changes are made to the licensing model:

- You can license a subset of the total usable cluster capacity. With partial licensing, you can purchase licenses based on your storage requirements.
- To be license-compliant, the used capacity must not be greater than the licensed capacity.
- Starting with version 10.3, NetBackup uses `.slf` license files instead of license keys. After the NetBackup Flex Scale cluster is configured, you must use the NetBackup web UI to add NetBackup licenses.

For more information, see the *License management* section of the*Veritas NetBackup™ Flex Scale Administrator's Guide*.

Support for parallel installation of EEBs

You can choose a parallel or a rolling method to install NetBackup EEBs. Installing multiple EEBs in parallel on all the nodes can help reduce the service downtime.

For more information, see the *Installing EEBs using GUI* section of the *Veritas NetBackup Flex Scale Installation and Configuration Guide*..

Support for only parallel upgrades

Only the parallel upgrade option is supported for upgrading to version 3.2. Previously, both parallel and rolling options were supported.

For more information, see the *NetBackup Flex Scale upgrades and patch management* section of the *Veritas NetBackup Flex Scale Installation and Configuration Guide*.

Including vendor packages on appliance nodes

The vendor packages are pre-installed on out of the box appliance nodes. In the previous release, the vendor packages were not installed with the ISO and you were required to install the vendor packages separately.

Configuring the console FQDN

You can now log in to the NetBackup Flex Scale UI using the console FQDN. For more details, refer to the *Configuring the console FQDN* section in the *Veritas NetBackup Flex Scale Administrator's Guide*.

Configuring MTU on public interfaces

The MTU property controls the maximum transmission unit size for an Ethernet frame.NetBackup Flex Scale allows you to configure MTU on public interfaces. For

more details, refer to the *Configuring MTU on public interfaces* section in the *Veritas NetBackup Flex Scale Administrator's Guide*.

Support for NetBackup Client

Table 2-1 lists the NetBackup Client support for NetBackup Flex Scale.

Table 2-1

Client support	Standard Client	Client Direct
NetBackup 7.7.3 Client	Supported	Not supported
NetBackup 8.0 Client	Supported	Not supported
NetBackup 8.1 Client (and later versions)	Supported	Supported

Chapter 3

Limitations

This chapter includes the following topics:

Software limitations

This section describes the software limitations in NetBackup Flex Scale.

- Swagger does not support downloading of large files.
- Veritas Call Home supports uploading of files with a maximum size of 2 GB. Larger file uploads may fail.
- The NetBackup Flex Scale load balancer feature does not work for VMware Continuous Data Protection (CDP) as the protection plan for VMware CDP needs a specific continuous data protection gateway.
- A single NetBackup Flex Scale cluster supports up to 16 secondary data networks.
- 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.

Unsupported features of NetBackup in NetBackup Flex Scale

The following features of NetBackup are not supported in NetBackup Flex Scale 3.2 release:

- Advanced Disk/Basic Disk storage units
- Client Direct Restore

- DNAS backup host pool with ECA configured
- Primary server only deployment
- MSDP FC OptDup and Replication
- MSDP Cloud and Universal Share on multi-domain
- S3 interface for MSDP
- Replication Director
- SAN Client
- Universal Share with MSDP-Cloud
- 3rd party OST device
- IPv4 and IPv6 mixed mode configuration
- Replicating backup images between NetBackup Flex Scale and NetBackup 8.3 or older MSDP server that is configured as a CloudCatalyst storage server.

Chapter 4

Known issues

This chapter includes the following topics:

Cluster configuration issues

The following known issues are related to cluster configuration.

Cluster configuration fails if there is a conflict between the cluster private network and any other network

The NetBackup Flex Scale cluster uses a private network for inter-node communication and this network should not be reachable or pingable from the nodes outside the cluster. The subnet used for the private network should not conflict with the IP address of any other node. Even if a second NetBackup Flex Scale cluster is present in the data center, it should not be reachable using the private network. (IA-22967)

Workaround:

There is no workaround for this issue.

Cluster configuration process may hang due to an ssh connection failure

The NetBackup Flex Scale cluster configuration process may sometimes get stuck for a long time and eventually fail. This issue occurs due to an ssh connection failure between the nodes. (IA-29939)

Workaround:

There is no workaround for this issue. In such a scenario you may have to initiate the cluster configuration workflow wizard once again.

Node discovery fails during initial configuration if the default password is changed

If you change the default maintenance account password and you click **Rescan** on the Select nodes panel in the NetBackup Flex Scale setup wizard, the node discovery operation hangs. The default password is required on all the nodes. (IA-38247)

Workaround:

If you already changed the password, you must reset the password to the default password using the following command:

```
# usermod -p
'$6$MQFQv7x8IMxW981P$HE0lj8R1HS8BZzomtLCUKDverksLWNouiUuRjBYVNrMVa9M
h1CGDoNu5cvN5lVj7ArpkSVdJHPKk5UlInWw1b1' maintenance
```

When NetBackup Flex Scale is configured, the size of NetBackup logs might exceed the /log partition size

NetBackup hosts have the capability to manage their log retention by configuring the **Keep logs up to GB** option. This option specifies the size of the NetBackup logs that you want to retain. When the log size grows to this value, the older logs are deleted.

When NetBackup Flex Scale cluster is configured, one of the cluster nodes always has both the NetBackup primary server and media roles. Additionally, the NetBackup primary is highly available and can failover to another node in the cluster. So either of the cluster nodes can have both primary and media roles. The cluster nodes share the logging storage with NetBackup hosts. However, the cluster nodes have

their own logging configuration and the log retention configured for the NetBackup hosts is not enforced. (4008252)

Workaround:

The combined size configured for the hosts with NetBackup primary and media role must be less than the maximum storage set aside for the log partition. Set the **Keep logs up to GB** option of these hosts accordingly. The **Keep logs up to GB** option is available on the **NetBackup Administration Console** > **NetBackup Management** > **Host Properties** > **Logging**dialog box (corresponds to the **KEEP_LOGS_SIZE_GB** property in the `bp.conf` file).

Error message is not displayed when NTP server is added as FQDN during initial configuration in a non-DNS environment

NTP server as FQDN is not supported in a non DNS configuration. NTP server should be given as IP address during configuration. (IA-40061)

Workaround:

NTP server should be given as IP address during configuration.

During cluster configuration, sysadmin user is not detected on some of the nodes

During cluster configuration, if a wrong Administrator password is provided or an incorrect updated Administrator password is provided, the 'sysadmin user present on only some of the selected nodes, you will not be able to change the sysadmin user password' error is displayed. (APPSOL-179178)

Workaround:

Update the correct Administrator password using any one of Appliance Node CLI commands:

- To store the default pull tab sticker password: `system store-default-BMC-credentials`
- To store the updated Administrator password: `system store-updated-admin-BMC-credentials`

Disaster recovery issues

The following known issues are related to the NetBackup Flex Scale disaster recovery configuration.

Backup data present on the primary site before the time Storage Lifecycle Policies (SLP) was applied is not replicated to the secondary site

Once you configure an SLP with a backup policy, replication of backup data starts only from that point onwards, so any backup data residing on the primary site before the time that the SLP was applied is not replicated to the secondary site. (IA-27334)

Note: A full client restore or recovery is possible from the secondary cluster only after a full backup schedule is run after the SLP is applied to a policy.

Workaround:

To restore any previous versions of the backup data (data which was present before the SLP was set) from the secondary site, you have to duplicate the backup images manually to the secondary site.

If the replication link is down on a node, the replication IP does not fail over to another node

When you perform disaster recovery, if the replication link is down on the node on which replication IP is residing, the replication IP should fail over to the other node as it is a failover group. But that does not happen and replication is paused and goes into error state. (IA-37024)

Workaround:

In the GUI , go to **Settings** > **Services Management**. Select **Run auto fix**. The IP will become available.

Or

Run the `shutdown -r` command from the node-level CLI on the CVM master node. Restart the node so that CVM master and replication group, GRP_VVR_REP_VIP can failover.

Disaster recovery configuration may take around 2.5 hours to complete when data-collect task runs in the backend

When you start disaster recovery configuration from the NetBackup Flex Scale cluster, it may take up to 2.5 hours for the configuration to complete. This may happen due to `data-collect` task running in the backend which tries to disable it from timing out after 2 hours.

(APPSOL-173301)

Workaround:

There is no workaround for this issue.

After disaster recovery takeover operation, the old recovery points or checkpoints for the primary server catalog file system are not visible in the GUI on the new primary site

The recovery points for the primary server catalog file systems are visible in the GUI under **NetBackup Catalog Management** on the primary site. The information about the recovery points is maintained in the database on the primary site. After a disaster recovery takeover, the database is not synced from the primary site to secondary site and the older recovery points for the primary server catalog file system are not visible on the new primary site. New recovery points are created every 2 hours on the new primary site and then these recovery points become visible in the GUI on the new primary site.

(IA-47712)

Workaround:

There is no workaround for displaying the old recovery points after disaster recovery takeover.

Disaster recovery configuration hangs when eth5/bond0 interface is down on the node where management console and CVM services are online on one or both sites

During disaster recovery configuration, if the eth5/bond0 interface is down on the node where the management console and CVM master are online on one or both sites, then the configuration hangs.

(IA-46084)

Workaround:

To resolve this issue, contact Veritas Technical Support and ask them to refer to article 100056273.

Miscellaneous issues

The following known issues are miscellaneous issues related to NetBackup Flex Scale.

Red Hat Virtualization (RHV) VM discovery and backup and restore jobs fail if the Media server node that is selected as the discovery host, backup host, or recovery host is replaced

If the Media server is replaced with another node, secure communication between the new node and the NetBackup host is lost because the options configured for a secure connection in the `bp.conf` file are deleted. If secure communication is not established, RHV discovery, backup and recovery jobs start failing. (4005637)

Workaround:

Reconfigure the secure connection between the new node and the NetBackup host by configuring the security options that were set earlier in the `bp.conf` file.

The file systems offline operation gets stuck for more than 2hrs after a reboot all operation

After you perform a reboot all operation, the file systems offline operation gets stuck. Hence, few file systems are not available and the respective containers also become offline. The backup/restore operation also get stuck as the primary/media containers are offline or unavailable. (4027460)

Workaround:

Force reboot all the nodes using the `echo b > /proc/sysrq-trigger` command.

SQLite, MySQL, MariaDB, PostgreSQL database backups fail in pure IPv6 network configuration

In an IPv6 environment when multiple IP addresses are configured for the client, the client tries to connect to NetBackup Flex Scale using an IP address that is chosen at random. If the IP address is not recognized as a trusted client, the backup job fails. (4031494)

Workaround:

On the client, disable route discovery for the ethernet interface. Use the `netsh` command to set the `routediscovery` parameter to **disabled**.

Exchange GRT browse of Exchange-aware VMware policy backups may fail with a database error

When browsing for VMware Exchange images, the `nblbc.exe` service crashes and you may see a "Database system error or file read failed." error message. (4031473)

The NCFLBC debug log may contain the following messages:

```
VDDK-Warn: VixDiskLib: Failed to load vddkVimAccess.dll : ErrorCode = 0x7e.!
(../BEDSContext.cpp:159),20:[fsys\shared]
Initial VirtApi DLL load check failed.  Will try again later.
... failed to load bedstrace.dll.
VDDK-Panic: Failed to load vixMntapi (../BEDSContext.cpp:159)
Failed to initialize the VDDK sub system on this thread.
It may have been already initialized. (../BEDSContext.cpp:159)
```

This issue occurs because of a missing Microsoft Visual C++ redistribution package on the system. In this case, the `nblbc.exe` service crashes because of a missing `vcruntime140_1.dll` file.

Workaround:

Install the latest version of Microsoft Visual C++ Redistributable for Visual Studio 2015, 2017 and 2019 packages to resolve the issue.

Refer to the following page for the latest installers:

https://support.microsoft.com/en-us/topic/the-latest-supported-visual-c-downloads-2647da03-1eea-4433-9aff-95f26a218cc0

Call Home test fails if a proxy server is configured without specifying a user

If a proxy server is configured for Call Home but the user that can log in to the proxy server is not specified in the Call home settings, the Call Home test fails. (APPSOL-155443)

Workaround:

Set the proxy server user in the Call Home settings. The user name must contain a minimum of two characters.

In a non-DNS NetBackup Flex Scale setup, performing a backup from a snapshot operation fails for NAS-Data-Protection policy

The issue occurs in a NetBackup Flex Scale environment when DNS entries are not present for the NetBackup Flex Scale nodes and there is only one backup host (NetBackup Flex Scale node) specified in the backup host pool. In such a scenario, snapshot operation fails for a NAS-Data-Protection policy. (4065603)

Workaround:

While creating a backup host pool, select at least two NetBackup Flex Scale nodes as part of the backup host pool.

In a non-DNS environment, the CRL check does not work if CDP URL is not accessible

The CRL check does not work if the CDP URL is not accessible in a non-DNS environment. (4063696)

Workaround:

If CRL is enabled, use one of the following options based on your configuration:

- Using CDP
  If `ECA_CRL_PATH` is not specified, NetBackup uses the CRLs from CRL distribution point (CDP) of the peer host's certificate. Ensure that the URLs that are available in the CDP are accessible and are added in the `/etc/host` file.
- Using existing downloaded CRLs
  If you have downloaded all the CRLs for all the required certificates in the trust chain (based on the `ECA_CRL_CHECK` setting), mention the CRL directory in the `ECA_CRL_PATH` in the NetBackup configuration file.

Else, you can also disable CRL by adding `ECA_CRL_CHECK= DISABLE` entry in `bp.conf/registry`.

Unable to add multiple host entries against the same IP address and vice versa in a non-DNS IPv4 environment

In a non-DNS IPv4 environment, if you try to add custom host entries with a different FQDN against the same IP address or if you add custom host entries with a different IP address against the same FQDN, the operation is not permitted.

Also, manually added entries get deleted from the `/etc/hosts` file in case of any resiliency-related operations (such as, failover) occur. (4066927)

Workaround:

If duplicate IP/ FQDN needs to be added, add the entries to `/etc/hosts` on each node.

You have to add the entries again in `/etc/hosts` of each node if any resiliency-related operations occur.

Incorrect information is displayed for the support health check command in an IPv6 environment

The `support health check` command displays incorrect output in an IPv6 NetBackup Flex Scale setup. This command works as expected in an IPv4 setup. (IA-46517)

Workaround:

There is no workaround for this issue.

Change in host time zone is not reflected within containers

The containers are unable to detect that the `/etc/localtime` is updated, the time zone is updated only on the hosts. (IA-35936)

Workaround:

Stop and start each node successively using the **Infrastructure > Nodes > Stop services** and **Infrastructure > Nodes > Start services** command. When you stop and start the nodes, the containers are restarted and can detect the updated `/etc/localtime` on the host. This operation will cause job failures, and jobs may or may not get automatically restarted based on the associated policy.

Failed to sync certificate from NetBackup primary server

This issue is seen when ECA is configured on a NetBackup Flex Scale cluster on which only media servers are deployed.

IRE requires the target MSDP server to enroll ECA certificate on the source primary server, so that AIR can work. The `setting certificate enroll-external-certificates` command is run to help the MSDP engine to enroll the ECA certificate on a primary server. But this command fails as some of the parameters required for this command are not available to the command. This leads to the ECA certificates not getting enrolled for the IRE primary server.

(4153980)

Workaround:

There is no workaround for this issue.

Password change for sysadmin user is partially successful

Changing the password for sysadmin user completes successfully only on some of the nodes and fails on the other nodes. (IA-54456)

Workaround:

Retry the operation by specifying the password that was set successfully on some of the nodes as the old password, and then reset the password by providing a new password.

To reset the password:

1. In the GUI, navigate to **Settings > User management**.
2. For the sysadmin user, click the Actions menu (vertical ellipsis) on the right side of the row and click **Change password**.
3. In the **Old password** field, specify the password that was set successfully on some of the nodes.
4. In the **New password** field, enter a new password and confirm the new password in the **Re-enter password** field.

Frequent alerts related to sysadmin users are generated and resolved automatically

The following alerts are generated frequently and email notifications about the alerts are sent. These alerts are resolved automatically after approximately 10 minutes and email notifications about resolved alerts are also received.

- The IPMI sysadmin user is not present on the node. Please add "sysadmin" user using the IPMI console.
- Failed to create sysadmin account for IPMI. The user creating the sysadmin account does not have sufficient privileges or the password does not meet the password complexity requirements.

(IA-54609)

Workaround:

These alerts are automatically resolved after approximately 10 minutes. You can ignore these alerts or suppress the alerts.

Universal Share backup does not happen when NetBackup Flex Scale is configured with IPv6

Universal Share is not supported on NetBackup Flex Scale if the NetBackup Flex Scale cluster is configured with IPv6. Hence, the Universal Share backup and Copilot backup with Universal Share fails.

(4153805)

Workaround:

There is no workaround for this issue.

Unable to generate kernel crash

After a Gen10/Gen11 physical cluster is configured, the kernel crash does not get generated.

(APPSOL-180671)

Workaround:

For Gen-10 platform:

- Update `/etc/kdump.conf` file to remove "=" from line (on each node)

  ```
  extra_modules = overlay ext4 xfs
  ```

- Set the kernel kdump to 2G (on each node)

  ```
  grubby --update-kernel=ALL --args="crashkernel=512M-2G:64M,2G-:2G"

  reboot
  ```

For Gen-11 platform

- Update `/etc/kdump.conf` file to remove "=" from line (on each node)
- `extra_modules = overlay ext4 xfs`

NetBackup issues

The following known issues are related to NetBackup.

The NetBackup web GUI does not list media or storage hosts in Security > Host mappings page

When the external certificate is deployed on NetBackup Flex Scale, and you go to **Security** > **Host mappings** page in the NetBackup web UI, the **Host mappings** page does not list media or storage hosts. It only has details on the NetBackup primary server and clients. (IA-35048)

Workaround:

You can go to **Settings** > **Network** > **Data-Network** in the NetBackup Flex Scale GUI to get the list of primary, media and storage servers.

Media hosts do not appear in the search icon for Recovery host/target host during Nutanix AHV agentless files and folders restore

When the external certificate is deployed on NetBackup Flex Scale, during the Nutanix AHV agentless files and folders restore, media hosts do not appear in the search icon for Recovery host. (IA-35048)

Workaround:

You can use any NetBackup server or a client as a recovery host. If you want to use the NetBackup Flex Scale media servers as the recovery host, you can go to **Settings** > **Network** > **Data-Network** of the NetBackup Flex Scale GUI to get the list of media servers and manually copy the media server names as the recovery host in the search icon.

On the NetBackup media server, the ECA health check shows the warning, 'hostname missing'

This issue occurs because media server FQDNs are not added during CSR generation. Hence, during ECA health check, the CERTIFICATE_SAN_HOSTNAME_VALIDATION check returns a WARN status on media servers. (IA-35366)

Workaround:

This issue can be ignored as there is no loss in functionality.

If NetBackup Flex Scale is configured, the storage paths are not displayed under MSDP storage

If you have configured NetBackup Flex Scale, the storage paths do not appear under **MSDP storage**. (4001518)

Workaround:

- Log on to the web UI.
- Click **Storage > Disk storage**.
- Click **Available storage on Storage servers** to see the details.

Failure may be observed on STU if the Only use the following media servers is selected for Media server under Storage > Storage unit

If NetBackup Flex Scale is configured and under **Storage** > **Storage unit**, the **Only use the following media servers** is selected for **Media server**, failure may be

observed on STU. This occurs if any of the media servers selected are not active. (4001652)

Workaround:

- Log on to the Java admin console.
- Click **Storage** > **Storage unit**.
- In the **Change Storage Unit** window, select **Use any available media server** option for **Media server**.

NetBackup primary server services fail if an nfs share is mounted at `/mnt` mount path inside the primary server container

This issue occurs if an external nfs share is mounted at the path `/mnt` inside the NetBackup primary server container running on the NetBackup Flex Scale appliance. (4010143)

The NetBackup primary server file system data (`/vx/PRIMARY_FS/data`) is mounted on the /mnt path (as `/mnt/nbdata`) inside the container. If the `/mnt` mount point is used by another entity, the NetBackup services are unable to access the NetBackup file system data and fail.

Workaround:

The `/mnt` path is reserved for NetBackup. You must unmount any shares that are mounted on the `/mnt` path inside the container. Veritas recommends that you do not mount any external shares directly inside the NetBackup containers on the appliance.

NetBackup primary container goes into unhealthy state

It may happen that the NetBackup primary container goes into unhealthy state on its own and the following error message gets displayed:

```
bashrpc error: code = 2 desc = oci runtime error: exec failed:
container_linux.go:235: starting container process caused
"process_linux.go:110:
decoding init error from pipe caused \"read parent: connection reset
 by peer\""
```

This also causes the ongoing backup and restore jobs to fail. (APPSOL-148171)

Workaround:

Stop the NetBackup primary container (nb_primary) forcefully and wait for few minutes till the primary container starts on its own.

```
# docker stop nb_primary
```

Note: This needs the involvement of the Veritas Support.

User login fails from the NetBackup GUI with authentication failed error

In the NetBackup UI, the user login fails and you get the following error:

```
Authentication failed.
```

This occurs when the user tries to access the NetBackup UI. If the password has expired, the NetBackup login API returns authentication failed error and does not specify the reason for failure such as password expiration.

(IA-47930)

Workaround:

Log on to the NetBackup Flex Scale UI. A validation is performed and if your password has expired, you are directed to the **Change password** tab.

MSDP engine and media server fail to come up

When a container is stopped in response to some infrastructure failure, the docker network configuration continues to hold the stale entry for the old container IDs due to an issue in the docker. When the containers are restarted after the infrastructure comes up, the MSDP engines and media server fails to come up due to the stale entries found in the configuration for the same endpoints.

(IA-48582)

Workaround:

To resolve this issue, contact Veritas Technical Support and ask them to refer to article 100060410.

Oracle Snapshot backup for Oracle workload fails with an error

In Oracle IA (Instant access), when Universal share is created on BYO enable media and mounted on client and Storage life cycle policies are created for snapshot and backup, the Oracle Snapshot backup for Oracle workload fails with the following error:

```
 error 156:snapshot creation failed.
```

Only backup jobs are partially successful.

(4153805)

Workaround:

There is no workaround for this issue.

MSDP storage server may go down in a multi-domain scenario if the same user is used by two or more NetBackup domain

When two domains access the same MSDP storage server using the same user id, the storage server rejects the request. In such a situation, the storage server appears as down.

(4139218)

Workaround:

Run following NetBackup Deduplication Shell command to restart MSDP service:

```
dedupe MSDP stop
dedupe MSDP start
```

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

Backup for Nutanix fails with status code 6

When a NetBackup Flex scale is upgraded, the NetBackup version also gets upgraded. After upgrade, the AHV backup job fails. This may happen because after upgrade, Nutanix AHV backup job attempts creation of VM snapshot and tries to mount it on the backup host by accessing the `/usr/openv/tmp/ntxmnt` folder. But this causes backup failure due to insufficient permissions while accessing the folder.

(4156222)

Workaround:

Delete the `/usr/openv/tmp/ntxmnt` folder on the backup host. The folder is re-created with desired permissions during AHV backup job and backup is successful.

Operation to start NetBackup services fails on a NetBackup Flex Scale cluster with media only deployment

On a NetBackup Flex Scale cluster with media only deployment, if NetBackup services are stopped cluster wide, and then NetBackup services are started again, the start operation fails with the following error:

```
error on GUI: Failed to start NetBackup service(s).
Operation on container sacatonvm28 failed with error
Container sacatonvm28 with type master is not stopped.
```

(IA-47689)

Workaround:

This error can be ignored.

Sometimes the NetBackup media services in a container cannot start

NetBackup media services in a container sometimes cannot start because the permissions of the `/var/VRTSpbx/root/` directory are incorrect. (4157827)

Workaround:

To resolve the issue, refer to the 100063348 article.

All the added media servers are not reflected in the Data Domain OST STU

During Data Domain storage server configuration, when you add additional media servers, the storage server is created successfully, but when you create a disk pool using this storage unit, the STU shows only a single media server that was selected on the first configuration screen (4149823)

Workaround:

Add the other media servers after creating the storage server. Use JAVA UI to create Data Domain OST STU.

Networking issues

The following known issues are related to the NetBackup Flex Scale networking module.

Cluster configuration workflow may get stuck

The NetBackup Flex Scale initial cluster configuration workflow wizard may hang and remain stuck at the configuration stage forever. The wizard UI does not display any error message or indicate if a failure has occurred. (IA-26240)

This issue may occur whenever the cluster configuration internal processes become defunct or do not get terminated properly.

Workaround:

There is no workaround for this issue. Contact Veritas Technical Support to help troubleshoot this issue.

Node panics when eth4 and eth6 network interfaces are disconnected

When the network interfaces corresponding to eth4 and eth6 go offline or manually made offline using commands such as, `ifconfig ethx down`, `ip link set down ethx`, or `ifdown ethx`, the node panics, and restarts. This is because when the private network links used for LLT heartbeat messaging are disconnected, the node gets isolated from the other nodes in the cluster and to avoid network split brain, the `vxfencing` module performs node membership arbitration and deliberately panics the node to avoid data corruption.

Network interfaces corresponding to eth4 and eth6 should never be disconnected as they are used as private heartbeat links among cluster nodes. (IA-26984)

The following are sample messages in the crash dump of the node that panics:

```
[19737.900357] LLT INFO V-14-1-10032 link 0 (eth4) node 2 inactive
15 sec (16250505)
[19737.950354] LLT INFO V-14-1-10509 link 0 (eth4) node 2 expired
[19738.050361] LLT INFO V-14-1-10032 link 0 (eth4) node 3 inactive
15 sec (16250505)
[19738.100361] LLT INFO V-14-1-10509 link 0 (eth4) node 3 expired
[19742.720979] VXFEN INFO V-11-1-80 RACER Node is: 0
[19742.720998] VXFEN INFO V-11-1-87 Initiating VxFen Race
[19742.720999] VXFEN INFO V-11-1-111 VxFen Pre-Race Delay: 0
[19742.721012] VXFEN INFO V-11-1-119 LEADER Node : 0 is in current
sub-cluster
[19742.721018] VXFEN CRITICAL V-11-1-89 RACER Node lost the VxFen
race
[19742.721019] VXFEN INFO V-11-1-112 VxFen Post-Race Delay: 0
[19742.721023] VXFEN NOTICE V-11-1-92 Sending LOST_RACE
[19742.721075] Kernel panic - not syncing: VXFEN CRITICAL V-11-1-20

Local cluster node ejected from cluster to prevent potential data
corruption.
[19742.722157] CPU: 0 PID: 8953 Comm: vxfen Kdump: loaded Tainted:
P OE
 ------------ T 3.10.0-1062.9.1.el7.x86_64 #1
[19742.722486] Hardware name: Veritas NetBackup Archive 3420/X11DPU,

BIOS 3.0c 03/27/2019
```

```
[19742.722808] Call Trace:
[19742.722965] [<ffffffffa757ac23>] dump_stack+0x19/0x1b
[19742.723129] [<ffffffffa7574967>] panic+0xe8/0x21f
[19742.723300] [<ffffffffc10668f2>] vxfen_plat_panic+0xc2/0xd0 [vxfen]
[19742.723467] [<ffffffffc1054d61>]
vxfen_process_client_msg+0x6d1/0xb30
[vxfen]
[19742.723779] [<ffffffffc1055d23>] vxfen_vrfsm_cback+0x323/0x1750
[vxfen]
[19742.723947] [<ffffffffc1055a00>] ? vxfen_reconfig_msg+0x840/0x840
 [vxfen]
[19742.724117] [<ffffffffc1073be8>] vrfsm_step+0x1c8/0x3a0 [vxfen]
[19742.724280] [<ffffffffc1055a00>] ? vxfen_reconfig_msg+0x840/0x840
 [vxfen]
[19742.724448] [<ffffffffc1075521>] vrfsm_recv_thread+0x401/0x9b0
[vxfen]
[19742.724613] [<ffffffffc1075120>] ? vrfsm_defer_message+0x140/0x140
 [vxfen]
[19742.724782] [<ffffffffc10761ee>] vxplat_lx_thread_base+0x9e/0xf0
 [vxfen]
[19742.724947] [<ffffffffc1076150>] ? vxplat_assert+0x20/0x20 [vxfen]
[19742.725123] [<ffffffffa6ec61f1>] kthread+0xd1/0xe0
[19742.725282] [<ffffffffa6ec6120>] ? insert_kthread_work+0x40/0x40
[19742.725449] [<ffffffffa758dd1d>]
ret_from_fork_nospec_begin+0x7/0x21
[19742.725611] [<ffffffffa6ec6120>] ? insert_kthread_work+0x40/0x40
```

Workaround:

Bring eth4 and eth6 online to allow the node to join the cluster properly.

You can also use the `ifup ethx` command if any of the above mentioned commands are used to bring down the interface. If the physical cable is disconnected, then reconnect it.

Add node fails during precheck when a secondary data network is configured over the management interface and the Automatic tab is used for providing input IPs for the new node to be added for the secondary data network over management interface

If the management network (*eth1* or *eth1.*<*VLANID*>) is configured to use a DNS name server, then during the add node operation, the **Automatic** tab is enabled to provide input for the secondary data network IPs for the management interface. If the **Automatic** tab is used to provide inputs for the required secondary data network

IPs for the new node, then during IP/FQDN validation, the FQDN resolution for management network IPs fail. This happens when the data DNS name server is different from the management DNS name server and the input IPs provided for the management network do not exist in the data DNS.

(IA-47639)

Workaround:

Use the **Custom** tab for providing input IP/FQDN for secondary data network on management interface during add node operation.

Static route does not get added if any node of the cluster is powered off or not up

When any one or more nodes are not up in the cluster, static route operation fails.

(4108770)

Workaround:

There is no workaround for adding static routes when one or more nodes in the cluster are down or offline. Static routes can only be added when all the nodes are up and running.

Add secondary data network operation fails on the management interface of the secondary site of a cluster when the management network on the secondary site is not the same as the management network on primary site and disaster recovery is configured using a single virtual IP

If primary and secondary sites have different management networks, then add secondary data network operation on the management interface fails in case disaster recovery is configured using a single virtual IP. This happens as the primary server IP for the secondary data network should be the same on both the primary site as well as the secondary site for secondary data network operation to work.

(IA-47766)

Workaround:

There is no workaround for this issue if disaster recovery is configured in single virtual IP mode. Disaster recovery should be configured using two virtual IPs for add secondary data network operation to work on the management interface, when the management networks on primary and secondary sites are different.

Data network details are not visible on NetBackup Flex Scale UI after console IP change

If you change the console IP and restart the nodes, both the console IP change operation and the node reboot operation are required to be completed before full discovery is triggered. Even if one of the operations is not complete, the full discovery operation fails. As a result, the data network details are not visible on the NetBackup Flex Scale GUI.

(IA-54031)

Workaround:

Run full discovery after the console IP change operation and the node reboot operation are complete.

Adding secondary data network at the secondary site using automatic mode fails

If you try to add a secondary data network on the secondary site using `automatic` mode to a cluster on which disaster recovery is configured, the operation fails in the validation step.

(IA-54221)

Workaround:

Instead of `automatic` mode, use the `custom` mode to add the secondary data network on the secondary site.

Adding secondary data network to the secondary site fails if ECA is configured on the cluster

Add data network operation fails on a disaster recovery configured secondary cluster when the host mapping gets updated if ECA is deployed on the cluster.

(IA-54523)

Workaround:

There is no workaround for this issue

If IP address is assigned to a VLAN over a bond on a node using the set network bond command, SSH using admin user does not work

Prior to cluster configuration, SSH is allowed for a list of admin devices which consists of eth0 and eth1. In case of VLAN over bond of eth1 and eth2, the bonded interface is not part of the pre configuration list of allowed interfaces for SSH. Hence, if you assign an IP address to a VLAN over a bond on a node using the `set network bond` command, SSH using admin user does not work.

(IA-54625)

Workaround:

ILO can be used for logging in with admin user credentials. The node can also be accessed by doing SSH using the root user and default password.

Create/remove bond operations are possible on both data and management networks from the GUI when secondary data/management network is present

NetBackup Flex Scale does not support `create bond` and `remove bond` operations on data and management networks when secondary data network or management network is configured on the cluster. But if you try to create and remove bonds on data/management network from the GUI when secondary data/management network is present, the GUI does not display an error.

(IA-54692)

Workaround:

If you want to perform bond create/remove bond operations when secondary data/management networks are present, then delete all the secondary data networks and then perform the create/remove bond operation.

Add secondary data network operation does not fail on the secondary site if the same NetBackup primary IP/FQDN is used to add the secondary data network in both the primary and secondary site

On a NetBackup Flex cluster configured with disaster recovery, you are not allowed to use the same IP/FQDN to add a secondary data network in both the primary and secondary site. But in the GUI, if you add a secondary data network in the secondary site using the same NetBackup primary IP/FQDN that was used to add the secondary data network in the primary site, the operation is successful. But if you perform

disaster recovery takeover operation, the operation fails as the IP is already online on the primary site.

(IA-47529)

Workaround:

There is no workaround for this issue.

Node and disk management issues

The following known issues are related to the NetBackup Flex Scale node and disk management.

Storage-related logs are not written to the designated log files

When you collect logs from the **Settings > Diagnostics** option of the NetBackup Flex Scale UI, and you select the **NAS** option on the **Generate log package** page, the generated logs are written to the `storage_snapshot.log` file instead of the designated log files. (IA-24755)

Designated log file	Logs written to
`/log/VRTSnas/log/storage_snapshot_destroy.log`	`/log/VRTSnas/log/storage_snapshot.log`
`/log/VRTSnas/log/storage_snapshot_create.log`	`/log/VRTSnas/log/storage_snapshot.log`
`/log/VRTSnas/log/storage_snapshot_delete.log`	`/log/VRTSnas/log/storage_snapshot.log`

Workaround:

There is no workaround for this issue.

Arrival or recovery of the volume does not bring the file system back into online state making the file system unusable

A disk may fail or a connection to a disk may fail. In such cases, if storage tolerance is exceeded, the volume that is constituted from that disk becomes disabled. The disabled volume causes the file system to go to an offline or faulted state making it unavailable for usage. After the underlying problem is corrected, the disk recovers and the volume also becomes enabled automatically. However, the file system does not come online on its own. This issue applies to all the file systems in the NetBackup Flex Scale cluster. (IA-25435)

Workaround:

1. Run AutoFix service from the GUI.

Settings> **Service management**> **Run auto fix**

2. Run the RESTful API for AutoFix.

```
POST /api/appliance/v1.0/management/autofix
```

Unable to replace a stopped node

If a cluster node is stopped for maintenance by using the **Stop node** option on the **Monitor > Infrastructure > Nodes** tab in the NetBackup Flex Scale UI, the node is marked as unhealthy and the **Replace node** option is not disabled. If you now attempt to replace this node, the replace node operation fails. (IA-26268)

Workaround:

There is no workaround for this issue.

Disk replacement might fail in certain situations

When you physically replace a faulty disk on a cluster node and start the disk replacement operation by using the **Replace disk** option on the **Monitor > Infrastructure > Disks** tab in the NetBackup Flex Scale UI, RAID 0 volume is created on the newly added disk and the operating system is queried for the new disks. However, the newly added disks are not discovered immediately by the operating system. There is a delay between RAID 0 creation and disks being available at the operating system level. (IA-27649)

Workaround

Retry the Replace disk operation by clicking the **Replace disk** option on the **Monitor > Infrastructure > Disks** tab in the NetBackup Flex Scale UI.

Replacing an NVMe disk fails with a data movement from source disk to destination disk error

When you physically replace a faulty disk on a cluster node and start the disk replacement operation by using the Replace disk option on the **Monitor>Infrastructure>Disks** tab in the NetBackup Flex Scale UI, an error is displayed in the **Disk replacement details** area even though data rebuild operation is in progress. (IA-30204)

Workaround:

Contact Veritas Support to resolve this issue.

Unable to detect a faulted disk that is brought online after some time

A disk that fails temporarily and is brought online later is not detected by the operating system as the logical device for that disk is still in a failed state. (IA-31660)

Workaround:

To recover the disk, bring the logical device online.

1. To view the failed logical device, use the `ssacli ctrl slot=0 ld all show` command.
2. To bring the failed logical device online, run the `ssacli ctrl slot=0 ld` *`number_of_failed_ld`* `modify reenable forced` command where *number_of_failed_ld* is the ID of the failed logical device.

Nodes may go into an irrecoverable state if shut down and reboot operations are performed using IPMI-based commands

If you use IPMI-based commands such as `ipmitool` and `ipmipower` to power off and power on NetBackup Flex Scale cluster nodes, it may cause the nodes to go into an irrecoverable state. (4019742)

This issue occurs because IPMI-based power commands do not perform a graceful shutdown of the operating system before powering off the node. The file systems on the nodes may fail to unmount before the power off, and may fail to mount when the node is powered back on. The file systems eventually appear in a partial or a faulted state. As a result, the NetBackup services containers fail to start and the cluster appears in an inconsistent state.

Workaround:

Do not use IPMI power utility commands to perform shut down and reboot operations on the NetBackup Flex Scale cluster nodes. If you wish to perform maintenance on the nodes, Veritas recommends that you perform a graceful shutdown of the nodes, one node at a time. Use the NetBackup Flex Scale infrastructure management console UI to stop, start, or shutdown the nodes.

For emergency scenarios or in situations where the system is unresponsive and you do not have physical access to the nodes, you can use the SysRq key to force a reboot on the nodes.

Run the following command to reboot the nodes without corrupting the file system:

```
echo b > /proc/sysrq_trigger
```

Replace node may fail if the new node is not reachable

Replace node operation may fail if the new node is not reachable due to network issues. (IA-30473)

Workaround:

There is no workaround for this issue. Contact Veritas Technical Support to help troubleshoot this issue.

Unable to collect logs from the node if the node where the management console is running is stopped

If you stop the node where the management console is running, the node goes out of cluster and you cannot collect logs from the node using the **Settings > Diagnostics** option in UI. (IA-37068)

Workaround:

To resolve this issue, contact Veritas Technical Support and ask them to refer to article 100056211.

Log rotation does not work for files and directories in /log/VRTSnas

Log files present in `/log/VRTSnas` are not rotated on daily basis due to change in ownership of the files present in `/log/VRTSnas`. Ownership of the files and directories is changed from root:root to root:accessuser. (IA-37405)

Workaround:

To resolve this issue, contact Veritas Technical Support and ask tem to refer to article 100056212.

Unable to start or stop a cluster node

Instead of SSH, hacli protocol is set for cluster node communication. Starting or stopping of cluster nodes is not supported for hacli mode of communication. (IA-37087)

Workaround:

Delete the `/opt/VRTSnas/conf/force_hacli` file and run the `cluster start` *`nodename`* or `cluster stop` *`nodename`* command.

Backup jobs of the workload which uses SSL certificate fail during or post Add node operation

Backup jobs of the workload which uses SSL certificate fail during or post Add node operation due to the renewal of the ECA certificates on the NetBackup Flex Scale nodes. This happens because the renewal of certificates causes the Nutanix and other workloads SSL certificates to be removed from the `ca_file.pem` file due to which the backup jobs fail. (4053617)

Workaround:

After add node operation is complete, re-append the SSL certificates into the `ca_file.pem` file and trigger the backup job.

During an add node operation, the error shown on the Infrastructure page is not identical to the error seen when you view the task details

When an add node operation fails, the failure shown on the **Infrastructure** page and the task details shown when you click **View details** might not match. The **View all activities** page does not create a parent task for the failure shown on the **Infrastructure** page, which is misleading. (IA-46303)

Workaround:

There is no workaround for this issue.

After a replace node operation is performed on a deployment in which ECA is enabled, the universal share is not mounted on the new node's engine

If you perform a replace node operation and update ECA on that node, the universal shares do not get mounted on the engine of the new node.

(IA-47731)

Workaround:

After replace node operation is completed and ECA is deployed on the new node, perform the following steps:

1. Log on to the NetBackup Flex Scale web UI.
2. Go to **Monitor > NetBackup > Services** tab.
3. Select the engine of the newly replaced node and perform a stop container operation.
4. After stop operation is completed, select the same engine and perform start container operation.

Health of the node does not change to unhealthy when disks are physically replaced

If a disk is physically replaced and you do not perform a replace disk operation from the UI, the UI shows the old disk as faulted but the health of the corresponding node is not shown unhealthy. (IA-47734)

Workaround:

Perform a replace disk operation from the UI. Navigate to the **Infrastructure > Disks** page, click the faulty disk that you want to replace, and then click **Replace disk**. Don't perform add node or upgrade operation in such state.

Incorrect error message shown when a node to be added restarts or panics

During an add node operation, if the new node being added restarts or panics, the add node operation fails with an error for configuring the network. This failure message can be misleading as it does not mention the actual reason for the failure. (IA-46222)

Workaround:

There is no workaround for this issue.

Add node operation shows NetBackup configuration failure when a newly added node restarts during rebalancing of data

During an add node operation, if the node being added restarts when rebalancing of data is ongoing and you click **View details** to monitor the progress, the UI incorrectly shows that **Rebalancing data across cluster** completed successfully and the failure occurred during **Configuring NetBackup**. (IA-46302)

Workaround:

There is no workaround for this issue.

Unhealthy disk are seen on the Infrastructure page after you delete a node from the cluster

If an add node operation fails and you delete the new node from the cluster, stale disk entries for the deleted node are shown in the UI. The **Infrastructure** page shows unhealthy disks for the deleted node. After the add node operation fails, the recovery tasks run partial discovery, which does not remove these entries from the UI. (IA-46378)

Workaround:

Run full discovery by navigating to **Settings > Services management > Run full discovery** or wait for the full discovery to run automatically as scheduled. The status is updated automatically in the UI after the full discovery is completed.

Proxy and etcd services do not come online when node shutdown fails

Proxy and `etcd` may remain in stopped state if the shutdown node operation fails. Even if the node is manually restarted after the shutdown failure, the proxy, and `etcd` services may continue to remain stopped. As a result, the media server does not start on the given node. The MSDP engine which had failed over to another node does not failback.

(IA-47340)

Workaround:

To resolve this issue, contact Veritas Technical Support and ask them to refer to article100056227.

After the management console node reboots, rollback of any running operations doesn't happen automatically

When any operation is triggered from the GUI and the node where the management console is online goes down or reboots, running operations are not rolled back automatically. This is because the process was running on the management console node and when the node goes down the management console switches to a new node but the process ends. (IA-46011)

Workaround:

Manual steps need to be performed to do the cleanup of any such failed operations. Contact Veritas Technical Support as the steps to be performed depend on the type of failed operation.

No value is displayed in the Used for column for faulted or excluded data disks

If a hard disk drive is faulted or excluded and you navigate to **Infrastructure > Disks**, the **Used for** column is empty. No value is displayed in this column for a faulted or an excluded data disk, but the corresponding value is shown for SSDs. For displaying details in the GUI, when the `storage fs list fsname` command is run, it fails to retrieve excluded and faulted data disks. (IA-54193)

Workaround:

There is no workaround for this issue.

ECA deployment fails for a newly added node on a cluster on which both primary and media servers are configured

If ECA is deployed when one or more nodes are down on a cluster on which both primary and media servers are configured or if a node is added to a cluster on which ECA is deployed, then an attempt is made to deploy ECA on the node when it comes online. But the ECA deployment may fail on the node again.

(IA-54470)

Workaround:

There is no workaround for this issue. You can run Autofix but it may not result in successful ECA deployment

Sysadmin user password is not synced to a newly added node

After a successful add node operation, the sysadmin password is not updated to be consistent with the sysadmin password on the other cluster nodes. As the password is not synced to newly added node, you cannot enable IPMI restriction on the newly added node. (APPSOL-178887)

Workaround:

Change the sysadmin user password on the newly added node by navigating to **Settings > User management**. In the sysadmin user row, click the Actions menu (vertical ellipsis) from the right side of the row in the UI and click **Change password**.

Add node operation fails while rebalancing data

The add node operation fails during data rebalance with error "Filesystems rebalance failed". (IA-54624 and 4149088)

Workaround:

To resolve this issue, contact Veritas Technical Support and ask them to refer to article 100064184.

After upgrade from a prior release to 3.2, add node fails if MSDP has at least one Cloud LSU configured

From NetBackup 10.3 release and later, all credentials are required to be migrated into Credential Management Service. (4152594)

Workaround:

Refer to the "Migrating or upgrading MSDP Cloud and CMS" section of the *NetBackup™ Web UI Administrator's Guide*.

Replace node fails with error "Failed to check connectivity with the gateway"

If both data network and management network devices are bonded and data network bond is created before management network bond, the data network is assigned the name bond0 and the management network bond is assigned the name bond1. On the new node, which will be used for replace node if the management IP address is already assigned using vxos-shell command `set network bond`, the replace node workflow will fail as the new node is in incorrect configuration with management bond name set as "bond0". (IA-54348)

Workaround:

Log in to the vxos-shell for that node and execute the command `delete network bond`. After the command completes, retry the replace node operation from the GUI.

Security and authentication issues

The following known issues are related to security and authentication.

NetBackup certificates tab and the External certificates tab in the Certificate management page on the NetBackup UI show different hosts list

The hosts lists are displayed under **Certificates management** > **NetBackup certificates** and **Certificates management** > **External certificates**. Both tabs should show a single certificate configured across all NetBackup Flex Scale hosts. But the **External certificates** tab shows single external certificate and all clients with external certificates while the **NetBackup certificates** tab shows multiple NetBackup certificates and no client certificates. (IA-35070)

Workaround:

There is no workaround for this issue.

Replicated images do not have retention lock after the lockdown mode is changed from normal to any other mode

If the lockdown mode of the NetBackup Flex Scale cluster was configured as normal mode (not enterprise or compliance), and you switch the mode from 'normal' to 'enterprise' or 'compliance', all the new backup images should be protected by WORM retention locks. But an exception occurs and all the new images replicated to the cluster still do not have the retention lock. (4050463)

Workaround:

Login to the primary server and restart the `nbstserv` service.

- Stop the nbstserv service.

  ```
  /usr/openv/netbackup/bin/nbstserv -terminate
  ```

- Restart the nbstserv service.

  ```
  /usr/openv/netbackup/bin/nbstserv
  ```

User account gets locked on a management or non-management console node

If the user account is locked on a management console node because of multiple incorrect login attempts, both SSH and GUI sign-in fail on that node till the account lock period is complete.

If the user account is locked on a non-management console node because of multiple incorrect login attempts, SSH to that specific node is blocked. SSH to all the other nodes and sign-in to the GUI continues to work.

Workaround:

Account lock depends on the password policies and STIG rules. Wait for the lock period to get completed.

The changed password is not synchronized across the cluster

If the STIG option is enabled or custom password rules are set for expiry, the user password expires as per the set password policy. After the password expires, the user is prompted to change the password if the user logs on to the system via SSH. If the user changes the password at the prompt, the password is changed only locally and is not synchronized across the cluster.

(IA-35890)

Workaround:

After the password expires, when prompted, do not change the password at the OS prompt. Instead, log on to the GUI with your credentials and change the password from the GUI.

Certificate renewal alert is not generated automatically during deployment

During ECA configuration, `DISABLE_CERT_AUTO_RENEW=1` entry is added in `bp.conf` file of both the primary and media servers. This entry prevents auto renewal of host-ID based certificate. Hence, alerts are not generated while renewing certificate if the `CLIENT_NAME` in `bp.conf` and SAN of host-id certificate are mismatched.

(IA-44935)

Workaround:

NetBackup Flex Scale handles the renewal of host ID-based certificate. So this issue can be ignored.

During IPMI restriction enable/disable operation, some of the nodes operations may fail

When IPMI restriction enable/disable operations are performed, some of the nodes operations may fail due to iLOrest issue. As a result, the status of the nodes on the GUI are not consistent with the actual status of the nodes. The GUI may show the status of the nodes as restricted but actually some of nodes are not restricted. Or, the status of the nodes may be restricted but the status does not appear as restricted in the GUI. Alerts also get generated accordingly.

(APPSOL-178561)

Workaround:

In Enterprise lockdown mode, user is allowed to enable/disable IPMI restriction on GUI. So to resolve the issue, you can try to enable/disable the IPMI restriction multiple times. In Compliance lockdown mode, user is not allowed to retry enable/disable IPMI restriction on GUI. To resolve the issue, contact Veritas Technical Support and ask them to refer to article 100064472.

After switching to FIPS security mode, all the users are deleted including the sysadmin user and the Administrator password is reset to default pull tag password

This behavior is as per the iLO FIPS security mode design. You are required to update the Administrator password and recreate the sysadmin user. (APPSOL-179331)

Workaround:

Update the Administrator password and then set the sysadmin password from the GUI. As part of sysadmin password change, the user is created and the password is updated.

Update the correct Administrator password by using any one of Appliance Node CLI commands:

To store the default pull tab sticker password: `system store-default-BMC-credentials`

To store the updated Administrator password: `system store-updated-admin-BMC-credentials`

Using iLO, create the sysadmin user with all the privileges and set the same password as the cluster-level sysadmin password.

Upgrade issues

The following known issues are related to the NetBackup Flex Scale upgrade.

EEB installation may fail if some of the NetBackup services are busy

During an EEB installation, the installer automatically stops the NetBackup services, patches the binaries, and then restarts the NetBackup services. But it may happen that the NetBackup services are busy. If any of the NetBackup services on any of the nodes fail to be stopped, it causes the EEB installation to fail. (4022006)

Workaround:

To resolve this issue, contact Veritas Technical Support and ask them to refer to article 100056214.

During an upgrade the NetBackup Flex Scale UI shows incorrect status for some of the components

During upgrade, the UI shows incorrect status for nodes and security settings such as STIG and FIPS. After the nodes are upgraded, the node count and the node status is not displayed correctly. The STIG and FIPS status incorrectly shows as disabled for clusters that had STIG and FIPS enabled before the upgrade. However, these issues are transient and the correct status is displayed either as the upgrade progresses or after the upgrade is completed successfully. (IA-36300)

Workaround:

No workaround is required for this issue. The UI shows the correct status for the components as the upgrade progresses.

Unable to upgrade from version 3.0.0.1 to 3.2 when an AD/LDAP server contains a maintenance user account and you attempt to change the maintenance user password

NetBackup Flex Scale contains a built-in maintenance user account. If you attempt to change the password for this user account and an AD/LDAP user with the same name exists, upgrade fails with the following error message:

```
V-493-10-5158 Password change is not allowed for AD/LDAP user.
```

Having the same AD/LDAP username as the built-in maintenance account causes a conflict. Instead of changing the password for the built-in maintenance account, NetBackup Flex Scale tries to change the password for the AD/LDAP maintenance account, which is not allowed.

(IA-46448)

Workaround:

Rename or delete the AD/LDAP maintenance user from the AD/LDAP server, and then try again.

Server busy error is displayed during an upgrade rollback

If an upgrade fails, during rollback, server busy or another operation is in progress error message is displayed. Some lock files acquired during the upgrade are not deleted and the error is displayed because of these stale lock files. (IA-40468)

Workaround:

Contact Veritas Support to delete the stale lock files.

After upgrade, MSDP cloud operations fail on an IPv6 setup if the cluster has MSDP cloud configured with data network in a non-DNS environment

For an IPv6 cluster, during upgrade, the public DNS name servers present in resolv.conf container get removed. Hence, the MSDP cloud operations fail.

(IA-47834)

Workaround:

After upgrade, use the GUI to add IPv6 enterprise DNS name servers that can resolve cloud provider IPs.

Pre-upgrade check fails intermittently on the primary cluster

When you run the pre-upgrade check by clicking **Start pre-check** from the GUI, sometimes the precheck operation fails on the primary cluster when the secondary cluster precheck fails, and at times the operation passes even when it fails on the secondary cluster. (IA-51128)

Workaround:

Run the precheck on the primary cluster first. If the precheck succeeds on the primary cluster, run the precheck on the secondary cluster. Start the upgrade only after the precheck operation succeeds on both the clusters.

Incorrect status is shown for the appliance firmware components after a firmware upgrade

When you upgrade the firmware after upgrading the NetBackup Flex Scale cluster to 3.2, wrong status and state is displayed for the firmware components. When you run the `show hardware-health node component=Firmware` command from the Appliance Node CLI, the status and state of the firmware components is shown as Unsupported/Failed. (APPSOL-179299)

Workaround:

to resolve this issue, contact Veritas Support and ask them to refer to article 100064274.

Failed to replace a node after upgrading from 3.0 to 3.2

The replace operation fails with error "uploading the missing EEBs VRTSnbfsapp_EEB_ET4071200-3.0.0.0-1.x86_64.rpm on *newnode* ". The EEB, which was installed at version 3.0 still exists, and was not rolled backed during the upgrade. (IA-54406)

Workaround:

If the appliance is at version 3.0 or 3.0.0.1, before an upgrade you must roll back EEB `VRTSnbfsapp_EEB_ET4071200-3.0.0.0-1.x86_64.rpm` if it is installed.

While upgrading and post upgrade to 3.2 on an ECA configured cluster, you may see '502 Bad Gateway' error while accessing the NetBackup Flex Scale UI using management gateway FQDN or IP

When ECA is configured on a 3.1 cluster and you start an upgrade to 3.2, at around 80% status, you will encounter error '502 Bad Gateway' while accessing UI using the management gateway FQDN or IP. Post successful upgrade too, this issue persists. (IA-54634)

Workaround:

You can monitor the upgrade progress by accessing NetBackup Flex Scale UI using the console IP. And post upgrade completion, ECA needs to reconfigured. Refer the "Deploying external certificates on NetBackup Flex Scale" section of the *Veritas NetBackup™ Flex Scale Administrator's Guide*.

Upgrade progress continues to remain in the same state at 51% after all the node are rebooted

There are some system services that are dependent on local file systems. If these file systems have less or no free space available, services don't work as expected, which causes the upgrade process to hang. (IA-54653)

Workaround:

Check space usage for the file systems before starting an upgrade by using the following command:

```
hacli -cmd df -kh fs_name
```

where ***fs_name*** corresponds to the following file systems:

- /var/log
- /var/log/audit
- /home
- /repository
- /system/data
- /uss/config

When an upgrade is performed from NetBackup Flex Scale version 3.0.0.1 to 3.2, the 3.0.0.1 EEBs are visible on the Vendor Packages page after upgrade

If you upgrade from NetBackup Flex Scale version 3.0.0.1 to version 3.2, the 3.0.0.1 EEBs are visible on the **Vendor Packages** page after upgrade. This is because to RHEL7 uses md5sum check while RHEL8 uses rpm --checksig(sha256).

(APPSOL-179083)

Workaround:

If vendor packages are listed as available, go to the **Vendor packages** tab in the GUI. Click on the vertical ellipsis button from the right side of the UI and then select **Remove** if add-on packages are not required. The EEBs and old vendor packages will get deleted from the incoming folder.

UI issues

The following known issues are related to the NetBackup Flex ScaleUI.

During the replace node operation, the UI wrongly shows that the replace operation failed because the data rebuild operation failed

If the private network is down and SSH connection is lost between the cluster node where the NetBackup Flex Scale Appliance GUI is running and the replacement node, the UI wrongly shows that the replace operation failed even though the network connectivity was restored between the nodes and the node was replaced successfully. (IA-27044)

Workaround

Contact Veritas Support to resolve this issue if the cluster is in an inconsistent state.

Changes in the local user operations are not reflected correctly in the NetBackup GUI when the failover of the management console and the NetBackup primary occurs at the same time

Changes in the local users such as user addition/deletion/password modification is not reflected correctly in the NetBackup GUI when the management console fails over during or after the NetBackup primary service failover. This behavior is observed when both the management console and NetBackup primary service failover. If the management console comes online while the onlining of NetBackup primary is still in-progress or in post-execution state, a timing issue may occur. If the timing issue

occurs, the changes in the local users such as user addition, deletion and password modification is not reflected correctly in the NetBackup GUI. (IA-27380)

Workaround:

1. Stop the node (on which the NetBackup primary server resides) using the GUI.

   Monitor > **Infrastructure** > **Stop node**

2. Start the node again using the

   Monitor > **Infrastructure** > **Start node**

Note: If the node selected in the first step has the NetBackup primary server and the management console, you may again face this issue depending on the timeline of the failover of these two services. In that case, repeat the workaround.

Mozilla Firefox browser may display a security issue while accessing the infrastructure UI

This issue may occur if you are accessing the NetBackup Flex Scale infrastructure UI using the Mozilla Firefox browser. (IA-29852)

Firefox displays a "Connection not secure" message for the URL of the UI. Even if you add the product certificate to the browser's trusted authorities list, the browser continues to indicate that the connection is insecure.

Workaround:

There is no workaround for this issue at the moment.

Recent operations that were completed successfully are not reflected in the UI if the NetBackup Flex Scale management console fails over to another cluster node

After the management console fails over to another node, the UI does not reflect the current status of the completed operations and displays incorrect status, which can be misleading. The status is updated automatically after the full discovery is completed. (IA-31524)

Workaround:

There is no workaround for this issue. Wait for the full discovery to complete, which is automatically scheduled.

Previously generated log packages are not displayed if the infrastructure management console fails over to another node.

After an upgrade, if the infrastructure management fails over to another node, the previously generated log packages are not displayed under **Packaged logs** when you click **Settings > Diagnostics**. (IA-36280)

Workaround

There is no workaround for this issue.

Smart card authentication fails for a cluster that includes both primary and media servers with IPv6 configuration

Smart card authentication fails for local, AD, and LDAP users if a cluster with primary and media servers is configured using IPv6 addresses. The OCSP verification fails. (IA-47508)

Workaround:

Ensure that you specify an IPv6 OCSP URI when you configure smart card authentication for a cluster with both primary and media servers.

Multiple tasks appear to be running in parallel during an add node operation

During an add node operation, if you click **View details** to monitor the progress, multiple tasks appear to run in parallel in the UI because the sub-tasks are misaligned. (IA-46329)

Workaround:

There is no workaround for this issue.

Incorrect search results are displayed when you search for EEBs on the Software management > Add-ons tab

When you search for EEBs on the **Settings > Software Management > Add-ons** tab, the search results include EEBs where the EEB names match a subset of the text instead of an exact match. (IA-47617)

Workaround:

There is no workaround for this issue.

Upgrade progress is not updated on the View details page of the GUI

If you click **View details** to monitor the upgrade progress, the page might not show the current status and details of the current ongoing tasks. The upgrade task appears to be running for a long time without providing any detailed information about its current status. (IA-54100)

Workaround:

Navigate to **Settings > Software management** and expand **Cluster upgrade progress status** to view the upgrade process.

Only three IPMI IP addresses are shown in the GUI post configuration for a four node iLO-FIPS enabled cluster

After the cluster configuration, when you navigate to **Settings > Network > IPMI network**, instead of four, only three IPMI interfaces with assigned IP addresses are displayed. (APPSOL-179050)

Workaround:

This issue occurs intermittently and the list of IP addresses assigned to IPMI interfaces gets updated correctly after discovery is run automatically. You can run full discovery by navigating to **Settings > Services management > Run full discovery** or wait for the full discovery to run automatically as scheduled.

Chapter 5

Fixed issues

This chapter includes the following topics:

Fixed issues in version 3.2

The following issues are fixed in this release:

Table 5-1

ID	Description
IA-47618	EEB installation fails with Server busy error if you attempt to install another EEB immediately after installing a previous EEB
IA-47724	Number of online disks shown is incorrect after an OS disk goes offline
IA-47580	All the selected EEBs are not uploaded during the initial configuration
IA-47979	Storage licensing check fails during a pre-upgrade check for a cluster where disaster recovery is configured
IA-47672	Disk size is incorrectly shown as ? when an excluded disk is added back to the cluster
IA-47696	Alerts about faulted disks are not resolved after disk or node replacement
IA-47817	Unhealthy node count is not updated when a node is shut down or stopped
IA-47824	Disk might go in a faulted state after you include the same disk

Table 5-1 *(continued)*

ID	Description
IA-47782	AutoSupport settings are not synchronized on the newly added node
IA-37062	Node is displayed as unhealthy if the node on which the management console is running is stopped
IA-47800	During an add node operation, you might be prompted to enter IPv6addressesfor a cluster with IPv4 addresses
APPSOL-174025	Alert about unsupported smartpqi driver is not resolved immediately after installing the kmod-smartpqi package
IA-47939	Include disk option is not disabled on an offline node
IA-47998	Dashboard shows an unconfigured icon for the NetBackup primary server when its status is offline
IA-47999	NetBackup primary server status is shown online on the dashboard after performing a stop containers operation
IA-27647	An NVMedisk is wrongly selected as a target disk while replacing a SAS SSD
IA-47978	Add data network operation fails after an upgrade