

NetBackup™ Flex Scale Release Notes

3.5.100

NetBackup Flex Scale Release Notes

Last updated: 2026-04-22

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	Getting help	10
	About this document	10
	NetBackup Flex Scale resources	10
Chapter 2	Features, enhancements, and changes	12
	What's new in this release	12
	Support for mixed disks	12
	Enhancements to ACL configuration	12
	IPMI user update	13
	Support for NetBackup Client	13
Chapter 3	Limitations	14
	Software limitations	14
	Unsupported features of NetBackup in NetBackup Flex Scale	14
Chapter 4	Known issues	16
	Cluster configuration issues	16
	Cluster configuration fails if there is a conflict between the cluster private network and any other network	16
	Cluster configuration process may hang due to an ssh connection failure	17
	Node discovery fails during initial configuration if the default password is changed	17
	When NetBackup Flex Scale is configured, the size of NetBackup logs might exceed the /log partition size	17
	Cluster configuration fails during NetBackup configuration stage	18
	If the console node reboots, the install operation fails for inconsistent EEBs	18
	Disaster recovery issues	18
	Backup data present on the primary site before the time Storage Lifecycle Policies (SLP) was applied is not replicated to the secondary site	19

If the replication link is down on a node, the replication IP does not fail over to another node	19
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	19
Miscellaneous issues	20
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	20
The file systems offline operation gets stuck for more than 2hrs after a reboot all operation	20
SQLite, MySQL, MariaDB, PostgreSQL database backups fail in pure IPv6 network configuration	20
Exchange GRT browse of Exchange-aware VMware policy backups may fail with a database error	21
In a non-DNS NetBackup Flex Scale setup, performing a backup from a snapshot operation fails for NAS-Data-Protection policy	21
In a non-DNS environment, the CRL check does not work if CDP URL is not accessible	22
Unable to add multiple host entries against the same IP address and vice versa in a non-DNS IPv4 environment	22
Incorrect information is displayed for the support health check command in an IPv6 environment	22
Change in host time zone is not reflected within containers	23
Unable to access the cluster-level CLI on a media only deployment if a non-primary AD/LDAP user tries to SSH using the console IP address	23
SSH to eth1 IP of a node fails with an incorrect error message for a non-primary AD/LDAP user	23
Python libraries generate core dump	24
SSH does not work with containers	24
NetBackup issues	24
The NetBackup web GUI does not list media or storage hosts in Security > Host mappings page	24
Media hosts do not appear in the search icon for Recovery host/target host during Nutanix AHV agentless files and folders restore	25
On the NetBackup media server, the ECA health check shows the warning, 'hostname missing'	25
If NetBackup Flex Scale is configured, the storage paths are not displayed under MSDP storage	25

Failure may be observed on STU if the Only use the following media servers is selected for Media server under Storage > Storage unit	26
NetBackup primary server services fail if an nfs share is mounted at /mnt mount path inside the primary server container	26
NetBackup primary container goes into unhealthy state	26
User login fails from the NetBackup GUI with authentication failed error	27
MSDP engine and media server fail to come up	27
Oracle Snapshot backup for Oracle workload fails with an error	27
MSDP storage server may go down in a multi-domain scenario if the same user is used by two or more NetBackup domain	28
Backup for Nutanix fails with status code 6	28
Sometimes the NetBackup media services in a container cannot start	29
All the added media servers are not reflected in the Data Domain OST STU	29
Networking issues	29
Node panics when eth4 and eth6 network interfaces are disconnected	29
Static route does not get added if any node of the cluster is powered off or not up	31
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	31
Data network details are not visible on NetBackup Flex Scale UI after console IP change	32
Create/remove bond operations are possible on both data and management networks from the GUI when secondary data/management network is present	32
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	32
Edit secondary data network operation fails with incorrect error when only FQDN or IP address is changed	33
Management bond and data bond creation on a 15-node cluster takes more than an hour	33

If add data network operation fails on bond0, the physical IP does not get removed from the restarted node	33
Node and disk management issues	34
Storage-related logs are not written to the designated log files	34
Arrival or recovery of the volume does not bring the file system back into online state making the file system unusable	34
Unable to replace a stopped node	35
Disk replacement might fail in certain situations	35
Unable to detect a faulted disk that is brought online after some time	35
Nodes may go into an irrecoverable state if shut down and reboot operations are performed using IPMI-based commands	3
Replace node may fail if the new node is not reachable	36
Unable to collect logs from the node if the node where the management console is running is stopped	37
Unable to start or stop a cluster node	37
Backup jobs of the workload which uses SSL certificate fail during or post Add node operation	37
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	37
Incorrect error message shown when a node to be added restarts or panics	38
Unhealthy disk are seen on the Infrastructure page after you delete a node from the cluster	38
After the management console node reboots, rollback of any running operations doesn't happen automatically	38
Add node operation fails at NetBackup configuring stage	39
After upgrade from a prior release to 3.5.100, add node fails if MSDP has at least one Cloud LSU configured	39
Replace node fails with error "Failed to check connectivity with the gateway"	39
Node discovery for add/replace node operation fails while performing a passwordless SSH	39
Discrepancy in health of the disk in iLO and NetBackup Flex Scale web UI	40
Recovery does not happen in a NetBackup Flex Scale cluster after a reboot operation	40
Task for installing the data EEB while performing replace node fails when a data EEB is installed on the new node	40

Task for installing the data EEB while performing add node fails when a data EEB is installed on the new node	41
After management console group fails over to a new node, some of the CLISH commands fail	41
Security and authentication issues	41
NetBackup certificates tab and the External certificates tab in the Certificate management page on the NetBackup UI show different hosts list	41
User account gets locked on a management or non-management console node	42
The changed password is not synchronized across the cluster	42
Certificate renewal alert is not generated automatically during deployment	42
During IPMI restriction enable/disable operation, some of the nodes operations may fail	43
After switching to FIPS security mode, all the users are deleted and the Administrator password is reset to default pull tag password	43
Upgrade issues	43
Incorrect status is shown for the appliance firmware components after a firmware upgrade	43
EEB installation fails if you attempt to install an EEB that is downloaded using MFT	44
Upgrade precheck fails during the NetBackup license validation step even though valid NetBackup licenses are present	44
The EEB install subtask of the add node operation fails when the setup is upgraded, all mandatory EEBs are installed, and a new node with an earlier version of NetBackup Flex Scale is added	45
If EEB installation or rollback is performed using CLISH and any node's EEB installation or rollback is missed, the GUI does not display an EEB inconsistency alert	45
Unable to see progress card during EEB installation and rollback operation	45
Upgrade operations may fail on a NetBackup Flex Scale cluster on which only media servers are deployed when an upgrade is performed from version 3.2 or 3.2.100 to version 3.5.100	46
Gateway IP not visible in the GUI if a cluster is upgraded from version 3.2 or 3.2.100 to version 3.5.100 after an add gateway operation	46
UI issues	46

During the replace node operation, the UI wrongly shows that the replace operation failed because the data rebuild operation failed	46
Previously generated log packages are not displayed if the infrastructure management console fails over to another node.	47
Smart card authentication fails for a cluster that includes both primary and media servers with IPv6 configuration	47
Incorrect search results are displayed when you search for EEBs on the Software management > Add-ons tab	47
Upgrade progress is not updated on the View details page of the GUI	47
Only three IPMI IP addresses are shown in the GUI post configuration for a four node iLO-FIPS enabled cluster	48
NetBackup Flex Scale One UI does not work with 2FA on latest Chrome versions	48

Chapter 5	Fixed issues	49
	Fixed issues in version 3.5.100	49

Getting help

This chapter includes the following topics:

- [About this document](#)
- [NetBackup Flex Scale resources](#)

About this document

This document provides information specific to the Veritas NetBackup Flex Scale 3.5.100 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.

Features, enhancements, and changes

This chapter includes the following topics:

- [What's new in this release](#)
- [Support for NetBackup Client](#)

What's new in this release

This section lists the major new features and enhancements added in the 3.5.100 version of NetBackup Flex Scale.

Support for mixed disks

NetBackup Flex Scale Gen10 and Gen11 now support disk replacements with different capacities. Mixed-capacity disks are supported within nodes and across clusters, with logical capacity normalization to ensure consistency.

For more information, see the *Support for mixed disks* section of the *NetBackup™ Flex Scale Administration Guide*.

Enhancements to ACL configuration

This release introduces the ACL Configuration page, which enables you to upload, apply, download, remove, and monitor Appliance Compatibility List (ACL) configuration files on the cluster.

For more information, see the *ACL configuration* section of the *NetBackup™ Flex Scale Installation and Configuration Guide*.

IPMI user update

The IPMI sysadmin user has been removed. All IPMI-related operations are now performed using the existing Administrator user.

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

Limitations

This chapter includes the following topics:

- [Software limitations](#)
- [Unsupported features of NetBackup in NetBackup Flex Scale](#)

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 3.5.100 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.

Known issues

This chapter includes the following topics:

- [Cluster configuration issues](#)
- [Disaster recovery issues](#)
- [Miscellaneous issues](#)
- [NetBackup issues](#)
- [Networking issues](#)
- [Node and disk management issues](#)
- [Security and authentication issues](#)
- [Upgrade issues](#)
- [UI issues](#)

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$HE0lj8RlHS8BZzomtLCUKDverksLWNouiUuRjBYVNrMva9M  
h1CGDoNu5cvN51Vj7ArpkSVdJHPKk5U1InWw1b1' 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).

Cluster configuration fails during NetBackup configuration stage

When the data network IP and management network IP belong to different networks and the NetBackup primary IP is already being used, the precheck gets passed even though it should fail. As a result, the cluster configuration fails during NetBackup configuration stage (after 6 hours) since the primary server IP is already being used.

(IA-56242)

Workaround:

There is no workaround for this issue. Make sure all the IPs are not in use before cluster configuration. If any IP is in use, the cluster configuration may fail.

If the console node reboots, the install operation fails for inconsistent EEBs

While installing EEBs, if the management console node reboots, the EEBs are in an inconsistent state. If you install all the EEBs to make the EEBs consistent across the cluster, the EEB installation operation fails again.

(APPSOL-177924)

Workaround:

Contact Veritas Technical Support to resolve this issue.

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 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.!\n(..\\BEDSContext.cpp:159),20:[fsys\\shared]\nInitial VirtApi DLL load check failed. Will try again later.\n... failed to load bedstrace.dll.\nVDDK-Panic: Failed to load vixMntapi (..\\BEDSContext.cpp:159)\nFailed to initialize the VDDK sub system on this thread.\nIt 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>

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/hosts` 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.

Unable to access the cluster-level CLI on a media only deployment if a non-primary AD/LDAP user tries to SSH using the console IP address

If a non-primary AD/LDAP user tries to log on using SSH to console IP, the user is logged on successfully, but the cluster-level CLI commands (VxOS CLISH commands) are not displayed and the following error message is displayed:

```
/usr/bin/groups: cannot find name for group ID group ID
```

The cluster-level CLI is expected to show a limited set of commands, such as the commands for the man page module. However, no CLISH menu is displayed and instead, the following error message is seen:

```
[nbfs-3.5.100] /bin/grep: /etc/resolv.conf: Permission denied
```

(NBFS-509)

Workaround:

There is no workaround for this issue.

SSH to eth1 IP of a node fails with an incorrect error message for a non-primary AD/LDAP user

If a non-primary AD/LDAP user tries to log on using SSH to eth1 IP of the node, the log in is expected to be restricted with the "Not authorised" error message, but instead fails with the following error messages:

```
sudo: you do not exist in the passwd database
Cannot open netlink socket: Permission denied
Cannot open netlink socket: Permission denied
This IP address is not configured for management. Refer to the
appliance documentation for more information.
```

(NBFSA-509)

Workaround:

There is no workaround for this issue.

Python libraries generate core dump

Python threads are not closed causing a race condition to occur. This issue causes the python threads to generate a core dump in the `/crashdump/core/user/` directory. (IA-52908)

Workaround:

There is no workaround for this issue.

SSH does not work with containers

The ssh daemon service inside the containers exits as it finds that the private keys at `/etc/ssh/ssh_host_*_key` have incorrect ownership. Hence, SSH does not work with containers.

(IA-57906)

Workaround:

Switch the management console node to the previous node where GID of the `ssh_keys` group is correct and perform any user management operation so that the container gets the correct GID of `ssh_keys` group which is 997.

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.

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/VRTSpxb/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.

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 `vx fencing` 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] [<ffffffffffa757ac23>] dump_stack+0x19/0x1b
[19742.723129] [<ffffffffffa7574967>] panic+0xe8/0x21f
[19742.723300] [<ffffffffffc10668f2>] vxfen_plat_panic+0xc2/0xd0 [vxfen]
[19742.723467] [<ffffffffffc1054d61>]
vxfen_process_client_msg+0x6d1/0xb30
[vxfen]
[19742.723779] [<ffffffffffc1055d23>] vxfen_vrfsm_cback+0x323/0x1750
[vxfen]
[19742.723947] [<ffffffffffc1055a00>] ? vxfen_reconfig_msg+0x840/0x840
[vxfen]
[19742.724117] [<ffffffffffc1073be8>] vrfsm_step+0x1c8/0x3a0 [vxfen]
[19742.724280] [<ffffffffffc1055a00>] ? vxfen_reconfig_msg+0x840/0x840
[vxfen]
[19742.724448] [<ffffffffffc1075521>] vrfsm_recv_thread+0x401/0x9b0
[vxfen]
[19742.724613] [<ffffffffffc1075120>] ? vrfsm_defer_message+0x140/0x140
[vxfen]
[19742.724782] [<ffffffffffc10761ee>] vxplat_lx_thread_base+0x9e/0xf0
[vxfen]
[19742.724947] [<ffffffffffc1076150>] ? vxplat_assert+0x20/0x20 [vxfen]
```

```
[19742.725123] [<fffffffffa6ec61f1>] kthread+0xd1/0xe0
[19742.725282] [<fffffffffa6ec6120>] ? insert_kthread_work+0x40/0x40
[19742.725449] [<fffffffffa758dd1d>]
ret_from_fork_nospec_begin+0x7/0x21
[19742.725611] [<fffffffffa6ec6120>] ? 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.

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.

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.

Edit secondary data network operation fails with incorrect error when only FQDN or IP address is changed

If an FQDN or IP address is changed using the **Edit** option on the secondary data network page on a NetBackup Flex Scale cluster, the edit data network task fails with the following error:

```
Edit request rejected, None of the parameter is modified in request.
```

This is not the correct error message. This is by design as modifying only FQDNs without modifying IPs for the secondary data network is not supported.

Workaround:

There is no workaround for this issue.

Management bond and data bond creation on a 15-node cluster takes more than an hour

On a NetBackup Flex Scale cluster configured with 15 nodes, the management bond creation and data bond creation operations take a long time (almost an hour).

The time for create bond operation increases as the number of nodes or number of VLANs increase. The reason for this increase in time is because many configuration files are required to be synced up with all the nodes and bond configuration has to be done on each node by moving all the VLANs and IPs to the bond device in a sequential manner.

(IA-54430)

Workaround:

There is not workaround for this issue.

If add data network operation fails on bond0, the physical IP does not get removed from the restarted node

The physical IP does not get removed from the node that is restarted after the failure of an add data network operation.

(IA-47084)

Workaround:

1. Login to the node where unused IP address is present, using VxOS shell and do a `su support` elevate.
2. Find the netmask for the IP using the `ip addr show <interface-name>` command.
3. Delete the IP address using the `ip addr del <IP>/<Netmask> dev <interface-name>` command.

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.

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 `ssaccli ctrl slot=0 ld all show` command.
- 2 To bring the failed logical device online, run the `ssaccli 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.

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.

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.

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.

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.

Add node operation fails at NetBackup configuring stage

This issue occurs if `etcd startup` fails to listen on port 2379.

(4180259)

Workaround:

To resolve this issue, contact Veritas Technical Support and ask them to refer to article 100072143.

After upgrade from a prior release to 3.5.100, 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.

Node discovery for add/replace node operation fails while performing a passwordless SSH

If a passwordless SSH is performed from a cluster node whose MTU is 9000 MTU to a new node whose MTU is set to 1500 on the switch port, the SSH fails. This causes the node discovery for add/replace node to fail.

(IA-55899)

Workaround:

Make sure that switch port's MTU for the new node is set to 9000 before initiating discover node operation for add/replace node.

Discrepancy in health of the disk in iLO and NetBackup Flex Scale web UI

The web UI reports the disk as failed and displays the disk status as faulted. The disk remains in a failed state in web UI for a long time even though the hardware reports it as healthy. (IA-56807)

Workaround:

To resolve this issue, contact Veritas Support and ask them to refer to article 100071656.

Recovery does not happen in a NetBackup Flex Scale cluster after a reboot operation

Rebooting two nodes in a NetBackup Flex Scale cluster causes the recovery to fail. The nodes and disks appear in unhealthy state. If you click on *Unhealthy* on the **Infrastructure** tab, it shows that resync is pending on the disk.

(ASCM-199)

Workaround:

To recover from this state, navigate to **Settings > Services management** and run **Autofix** from the Appliance management UI.

Task for installing the data EEB while performing replace node fails when a data EEB is installed on the new node

If you try to install data EEBs on the new node during a replace node operation, the task for installing the data EEB while performing replace node fails. But the replace node operation continues and there is no functional impact.

(IA-57983)

Workaround:

If EEB is inconsistent then install the EEBs on nodes on which it is not present.

Task for installing the data EEB while performing add node fails when a data EEB is installed on the new node

If you try to install data EEBs on the new node during an add node operation, the task for installing the data EEB while performing add node fails. But the add node operation continues and there is no functional impact.

(IA-58040)

Workaround:

If EEB is inconsistent then install the EEBs on nodes on which it is not present.

After management console group fails over to a new node, some of the CLISH commands fail

During the post-online trigger for the management console group, the `vxctl defaultdg` command fails because the `vxconfigd` daemon is down on that node. This can cause some VxVM commands to not work on the node and may result in errors in CLISH commands.

(IA-55847)

Workaround:

Switch the management console group to a new node from 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.

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 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. (APPSOL-179331)

Workaround:

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`

Upgrade issues

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

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 3.5.100, 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.

EEB installation fails if you attempt to install an EEB that is downloaded using MFT

The naming convention for an EEB file is

`name-version-release.architecture.rpm`. If an EEB is downloaded by using Managed File Transfer (MFT), a 12-digit timestamp is added in the middle of the *version* component of the EEB file name, resulting in a file name check failure. (IA-56204)

Workaround:

Rename the downloaded EEB RPM file by deleting the extra 12-digit timestamp that was added to the file name, upload the renamed EEB RPM file to the UI, and then install the EEB.

Upgrade precheck fails during the NetBackup license validation step even though valid NetBackup licenses are present

This happens in a scenario where multiple NetBackup licenses are present and one or more license(s) have expired. Even if one of the NetBackup license is still valid, the upgrade precheck fails in the NetBackup license validation check step. The upgrade precheck fails with the following error:

```
[Error] V-409-776-30076: The NetBackup license has expired or is not valid.  
Add a valid NetBackup license or contact Veritas Support.
```

(IA-58020)

Workaround:

Remove the expired NetBackup license(s) from the configuration and run the upgrade precheck again. For removing the NetBackup licenses, login to the NetBackup UI and select **License Management**.

The EEB install subtask of the add node operation fails when the setup is upgraded, all mandatory EEBs are installed, and a new node with an earlier version of NetBackup Flex Scale is added

If an upgrade is performed and then all EEBs are installed post-upgrade, and if an add node operation is performed where the newly added node is on an earlier version of NetBackup Flex Scale, the add node operation upgrades the node and attempts to install the EEBs. However, the EEB installation subtask fails because they are already installed, and an error message appears. Despite this, the add node operation still completes successfully.

(IA-58039)

Workaround:

The error message can be ignored. This message has no impact on the add node operation.

If EEB installation or rollback is performed using CLISH and any node's EEB installation or rollback is missed, the GUI does not display an EEB inconsistency alert

If you perform EEB installation or rollback using CLISH and if any node's EEB installation or rollback is missed, the GUI does not display an EEB inconsistency alert.

(IA-58080)

Workaround:

Run **Full discovery**. Full discovery discovers the inconsistent EEBs and raises an alert.

Note: This issue is not observed if EEB installation or rollback is performed using GUI.

Unable to see progress card during EEB installation and rollback operation

When the EEB progress card is closed, an internal flag gets set. This flag does not get cleared when a new installation or rollback begins, leaving the progress card hidden.

(IA-58077)

Workaround:

Refresh the web page when EEB installation or rollback starts.

Upgrade operations may fail on a NetBackup Flex Scale cluster on which only media servers are deployed when an upgrade is performed from version 3.2 or 3.2.100 to version 3.5.100

If the API keys specified during cluster configuration on NetBackup Flex Scale version 3.0.0.1 are no longer valid on a cluster with media only deployment and if the setup has been upgraded to version 3.2 or 3.2.100, then an upgrade to version 3.5 may fail.

(IA-58149)

Workaround:

To resolve this issue, contact Veritas Technical Support and ask them to refer to article 100074130.

Gateway IP not visible in the GUI if a cluster is upgraded from version 3.2 or 3.2.100 to version 3.5.100 after an add gateway operation

If an add gateway operation is performed on a secondary data network on a NetBackup Flex Scale cluster version 3.2 or 3.2.100 and then an upgrade is performed from version 3.2 or 3.2.100 to version 3.5.100, the gateway IP address is not visible in the GUI after the upgrade. This happens because as part of the add gateway operation done on the 3.2 and 3.2.100 cluster, the gateway is not added to `net_info_list.conf` file.

Workaround:

There is no workaround for this issue.

UI issues

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

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.

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

After an upgrade, if the infrastructure management console 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.

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.

NetBackup Flex Scale One UI does not work with 2FA on latest Chrome versions

This issue occurs only when you select **Enhanced protection** in the **Settings > Privacy** tab in Chrome. The issue does not occur with the **Standard Protection** option. A cluster has multiple certifications. If it is an insecure network, the browser does not allow you to proceed if you have opted for **Enhanced protection** option.

Workaround:

1. Login to the NetBackup Flex Scale UI and navigate to **Settings > Security management > Certificates**.
2. Click **Download root certificate**. The `root-cert.pem` file gets downloaded.
3. Open the `root-cert.pem` file. It has 2 certificates (stem file and CA file). Split both certificates and upgrade them to the browser.
4. Upload `stem.pem` to intermediate CA.
5. Upload `ca.pem` to trusted CA
6. Restart browser.

Fixed issues

This chapter includes the following topics:

- [Fixed issues in version 3.5.100](#)

Fixed issues in version 3.5.100

The following issues are fixed in this release:

Table 5-1

ID	Description
IA-56502	No option is displayed to navigate away from the precheck tasks screen after the precheck is completed successfully
APPSOL-173301	Disaster recovery configuration may take around 2.5 hours to complete when data-collect task runs in the backend
IA-47712	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
IA-55671	Redeployment of ECA fails on a NetBackup Flex Scale cluster on which disaster recovery is configured
APPSOL-182270	Kernel crash dump is not getting generated post upgrade
IA-56907	Node addition may fail at the data rebalance stage on a cluster where disaster recovery is configured
IA-57012	Checkpoints are not seen on the NetBackup Flex Scale web UI on the disaster recovery primary site
IA-40061	Error message is not displayed when NTP server is added as FQDN during initial configuration in a non-DNS environment

Table 5-1 (continued)

ID	Description
IA-54221	Adding secondary data network at the secondary site using automatic mode fails
IA-47639	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
4153805	Universal Share backup does not happen when NetBackup Flex Scale is configured with IPv6