

Veritas NetBackup™ Flex Scale Release Notes

3.0

Veritas NetBackup Flex Scale Release Notes

Last updated: 2023-05-21

Legal Notice

Copyright © 2023 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

VERITAS TECHNOLOGIES LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

NB.docs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	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
	Configuring NetBackup Flex Scale in a non-DNS environment	12
	Managing NetBackup and cluster infrastructure from NetBackup Flex Scale UI	12
	Configuring AD/LDAP servers on clusters deployed with only media servers	13
	Restricting access to IPMI	13
	Running a preupgrade check	13
	Support for NetBackup Client	13
Chapter 3	Limitations	14
	Software limitations	14
	Unsupported features of NetBackup in NetBackup Flex Scale	15
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
	DNS servers that are added after initial configuration are not present in the /etc/resolv.conf file	17
	Empty log directories are created in the downloaded log file	17
	For the private network, if you use the default IPv4 IP address but specify an IPv6 IP other than the default, the specified IPv6 IP address is ignored	18
	Node discovery fails during initial configuration if the default password is changed	18

When NetBackup Flex Scale is configured, the size of NetBackup logs might exceed the /log partition size	18
Error message is not displayed when NTP server is added as FQDN during initial configuration in a non-DNS environment	19
Disaster recovery issues	19
Backup data present on the primary site before the time Storage Lifecycle Policies (SLP) was applied is not replicated to the secondary site	19
When disaster recovery gets configured on the secondary site, the catalog storage usage may be displayed as zero	20
Catalog backup policy may fail or use the remote media server for backup	20
Takeover to a secondary cluster fails even after the primary cluster is completely powered off	20
Catalog replication may fail to resume automatically after recovering from node fault that exceeds fault tolerance limit	21
If the replication link is down on a node, the replication IP does not fail over to another node	22
If both primary and secondary clusters are down and are brought online again, it may happen that the replication is in error state	22
Disaster recovery configuration fails if the lockdown mode on the secondary cluster is enterprise or compliance	22
Unable to perform a takeover operation from the new site acting as the secondary	23
Enabling compliance mode for the first time on the secondary cluster may fail if disaster recovery is configured	23
If disaster recovery is configured and an upgrade from NetBackup Flex Scale 2.1 to 3.0, the upgrade operation hangs	23
On a NetBackup Flex Scale cluster with disaster recovery configuration, the replication state shows Primary-Primary on the faulted primary cluster after takeover	24
Unable to create universal shares on a cluster on which disaster recovery is configured or only media server is deployed	24
Miscellaneous issues	25
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	25
The file systems offline operation gets stuck for more than 2hrs after a reboot all operation	25

cvmvoldg agent causes resource faults because the database not updated	25
SQLite, MySQL, MariaDB, PostgreSQL database backups fail in pure IPv6 network configuration	26
Exchange GRT browse of Exchange-aware VMware policy backups may fail with a database error	26
Call Home test fails if a proxy server is configured without specifying a user	27
In a non-DNS NetBackup Flex Scale setup, performing a backup from a snapshot operation fails for NAS-Data-Protection policy	27
In a non-DNS environment, the CRL check does not work if CDP URL is not accessible	27
Unable to add multiple host entries against the same IP address and vice versa in a non-DNS IPv4 environment	28
NetBackup issues	28
The NetBackup web GUI does not list media or storage hosts in Security > Hosts page	28
Media hosts do not appear in the search icon for Recovery host/target host during Nutanix AHV agentless files and folders restore	28
On the NetBackup media server, the ECA health check shows the warning, 'hostname missing'	29
If NetBackup Flex Scale is configured, the storage paths are not displayed under MSDP storage	29
Failure may be observed on STU if the Only use the following media servers is selected for Media server under Storage > Storage unit	29
NetBackup primary server services fail if an nfs share is mounted at <code>/mnt</code> mount path inside the primary server container	
3	0
NetBackup primary container goes into unhealthy state	30
NetBackup fails to discover VMware workloads in an IPv6 environment	30
Networking issues	31
DNS of container does not get updated when the DNS on the network is changed	31
Cluster configuration workflow may get stuck	31
Bond modify operation fails when you modify some bond mode options such as <code>xmit_hash_policy</code>	32
Node panics when <code>eth4</code> and <code>eth6</code> network interfaces are disconnected	32
Node and disk management issues	33

Storage-related logs are not written to the designated log files	33
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	34
An NVMe disk is wrongly selected as a target disk while replacing a SAS SSD	34
Disk replacement might fail in certain situations	35
Replacing an NVMe disk fails with a data movement from source disk to destination disk error	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	36
Add node fails because of memory fragmentation	36
Replace node may fail if the new node is not reachable	37
Node is displayed as unhealthy if the node on which the management console is running is stopped	37
Unable to collect logs from the node if the node where the management console is running is stopped	37
Log rotation does not work for files and directories in /log/VRTSnas/log	38
After replacing a node, the AutoSupport settings are not synchronized to the replacement node	38
Unable to start or stop a cluster node	39
The Add nodes to the cluster button remains disabled even after providing all the inputs	39
Unable to add more than seven nodes simultaneously to the cluster	39
Backup jobs of the workload which uses SSL certificate fail during or post Add node operation	39
Security and authentication issues	40
NetBackup certificates tab and the External certificates tab in the Certificate management page on the NetBackup UI show different hosts list	40
Replicated images do not have retention lock after the lockdown mode is changed from normal to any other mode	40
Unable to switch the lockdown mode from normal to enterprise or compliance for a cluster that is deployed with only media servers and with lockdown mode set to normal	41
CRL mode does not get updated on the secondary site after ECA is renewed on a cluster on which disaster recovery is configured	41

Setting lockdown mode to enterprise or compliance fails on the secondary cluster of a NetBackup Flex Scale cluster on which disaster recovery is configured	41
User account gets locked on a management or non-management console node	42
The changed password is not synchronized across the cluster	42
Upgrade issues	43
After an upgrade, if checkpoint is restored, backup and restore jobs may stop working	43
Upgrade fails during pre-flight in VCS service group checks even if the failover service group is ONLINE on a node, but FAULTED on another node	43
Upgrade from version 2.1 to 3.0 fails if the cluster is configured with an external certificate	44
During EEB installation, a hang is observed during the installation of the fourth EEB and the proxy log reports "Internal Server Error"	44
EEB installation may fail if some of the NetBackup services are busy	44
During an upgrade the NetBackup Flex Scale UI shows incorrect status for some of the components	45
After an upgrade Call Home does not work	46
After an upgrade, the proxy server configured for Call Home is disabled but is displayed as enabled in the UI	46
Unable to view the login banner after an upgrade	46
After an upgrade to NetBackup Flex Scale 2.1, the metadata format in cloud storage of MSDP cloud volume is changed	46
Rollback fails after a failed upgrade	47
Add node operation hangs on the secondary site after an upgrade	47
Alerts about inconsistent login banner and password policy appear after an upgrade from NetBackup Flex Scale version 2.1 to NetBackup Flex Scale 3.0	48
Alerts about node being down are generated during an upgrade	48
GUI takes a long time to update the status of the upgrade task	48
In a disaster recovery environment, upgrade gets stuck during node evacuation stage as VVRInfra_Grp cannot be brought down	48

Upgrade may fail after node evacuation if a VCS parallel service group is OFFLINE on a partial set of nodes at the beginning of the upgrade	49
Upgrade may fail if operations such as OS reboot, cluster restart, and node stop and shutdown are used during the upgrade	51
UI issues	51
In-progress user creation tasks disappear from the infrastructure UI if the management console node restarts abruptly	51
During the replace node operation, the UI wrongly shows that the replace operation failed because the data rebuild operation failed	52
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	52
Mozilla Firefox browser may display a security issue while accessing the infrastructure UI	53
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	53
Previously generated log packages are not displayed if the infrastructure management console fails over to another node.	53
User management issues	53
AD server test connection fails due to incorrect username on the IPV6 media only cluster	54
AD/LDAP domain unreachable alerts do not get cleared after the AD/LDAP server is deleted	54
GUI login fails with LDAP user if the domain is configured with SSL	54
Assigning role to correct AD/LDAP user/group with wrong domain causes the user listing to fail	54
After cluster reboot all/shutdown all operation, AD/LDAP domains become unreachable from one or more nodes on a NetBackup Flex Scale cluster on which only media servers are deployed	55
Chapter 5	
Fixed issues	56
AD/LDAP configuration may fail for IPV6 addresses	56
Fixed issues in version 3.0	56

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 3.0 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.
- Late Breaking News (LBN)
https://www.veritas.com/support/en_US.html
View the latest information about updates, patches, and software issues for this release.

VOX community forum

You can use the Veritas Open eXchange (VOX) community forum to connect directly with the NetBackup Flex Scale product development team:

<https://vox.veritas.com>

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 3.0 version of NetBackup Flex Scale.

Configuring NetBackup Flex Scale in a non-DNS environment

In NetBackup Flex Scale, enabling the DNS server is optional. Cluster nodes can be configured without the DNS server IP. You can specify the IP/FQDN details to configure NetBackup Flex Scale in a non-DNS environment.

For more information, see the *Veritas NetBackup™ Flex Scale Administrator's Guide* and the *Veritas NetBackup™ Flex Scale Installation and Configuration Guide*.

Managing NetBackup and cluster infrastructure from NetBackup Flex Scale UI

Using role-based access, you can access and manage both NetBackup and cluster infrastructure from NetBackup Flex Scale UI. You do not need to log in to NetBackup Flex Scale infrastructure management console to view and manage the cluster details.

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

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

NetBackup Flex Scale supports multiple AD/LDAP server configuration on clusters deployed with only media servers. You can configure AD/LDAP servers using the Appliance GUI.

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

Restricting access to IPMI

When the cluster is in compliance or enterprise mode, you can configure restricted access to IPMI to provide higher level of data security and to limit the operations you can perform. For more information, see the following article:

https://www.veritas.com/support/en_US/article.100052289

Running a preupgrade check

A preupgrade check now runs automatically when you start a cluster upgrade. You can also run this check independently after you download an upgrade package from the Veritas website or when you upload an upgrade package from a local system.

For more information, see the *Veritas NetBackup™ Flex Scale Installation and Configuration 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

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.

- The host name of any cluster node should not be the same as the FQDN of the IP that is assigned to the management interface or any publicly resolvable name.
- 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.
- Add data network operation does not work when ECA is deployed.
- Addition of an Auto Image Replication target to an MSDP disk pool on a secondary cluster is not supported.
- You must enable immutability on the NetBackup Flex Scale cluster before adding MSDP cloud LSUs to NetBackup Flex Scale. Attempting to enable immutability for a NetBackup Flex Scale cluster after adding an MSDP Cloud LSU may result in a failure if there are a significant number of files in the MSDP cloud cache.

Unsupported features of NetBackup in NetBackup Flex Scale

The following features of NetBackup are not supported in NetBackup Flex Scale 3.0 release:

- Advanced Disk/Basic Disk storage units
- Client Direct Restore
- Primary server only deployment
- MSDP FC replication
- Replication Director
- VMware Instant rollback
- 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](#)
- [User management 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.

DNS servers that are added after initial configuration are not present in the /etc/resolv.conf file

During initial configuration, only one DNS server can be added. If you add multiple DNS servers using the Appliance GUI after initial configuration, the newly added DNS servers are not synced with the `/etc/resolv.conf` file present in the NetBackup primary container. The `/etc/resolv.conf` file has only one DNS entry done during initial configuration. (IA-31849)

Workaround:

There is no workaround for this issue

Empty log directories are created in the downloaded log file

If the initial cluster configuration fails, you can download the logs by clicking **Download Logs** on the **Configuration Status** page of the NetBackup Flex Scale web interface. The downloaded file is a compressed log file in `.tar` format. After extracting the contents of this file, empty directories are generated, which is misleading and incorrectly indicates that some logs were not collected or downloaded. (IA-32032)

Workaround:

No workaround is required as all the logs are collected and downloaded. You can ignore the empty directories.

For the private network, if you use the default IPv4 IP address but specify an IPv6 IP other than the default, the specified IPv6 IP address is ignored

When you configure the private network, you have the option to retain the default IPv4 and IPv6 IP addresses or specify a different IP for the IPv4 and IPv6 network. If you choose to use the default IPv4 IP but specify a different IPv6 address, the specified IPv6 IP is ignored and the private network for both IPv4 and IPv6 is configured using the default IP addresses.

However, if you specify an IPv4 address other than the default value, the IPv6 network is configured correctly based on the specified IPv6 IP. (IA-32201)

Workaround:

There is no workaround for this issue.

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

Before configuring the cluster if you run the `support elevate` command from the NetBackup Flex Scale shell menu you are prompted to change the default password. After the password is changed and you click **Rescan** on the Select nodes panel in the NetBackup Flex Scale setup wizard, the node discovery operation hangs. (IA-38247)

Workaround:

Do not run the `support elevate` command before you configure the cluster as you will always be prompted to change the default password on your first login attempt.

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

```
# usermod -p  
'$6$MQFQv7x8IMxW981P$HE01j8R1HS8BZzomtLCUKDverksLWNouiUuRjBYVNrMva9M  
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).

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.

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.

When disaster recovery gets configured on the secondary site, the catalog storage usage may be displayed as zero

During the configuration of disaster recovery on the secondary site, the catalog storage usage may get displayed initially. But later, the catalog storage usage does not get displayed. This happens because the primary file system remains offline on the secondary site. (IA-30046)

Workaround:

No workaround is required. After the disaster recovery configuration is complete and the secondary site becomes the primary, the catalog storage usage gets displayed correctly.

Catalog backup policy may fail or use the remote media server for backup

When disaster recovery is configured, the catalog backup policy is configured with the storage unit that is local to the cluster. After migration, it is possible that the catalog backup policy is configured with the storage unit on the remote site and in such cases, the media server present on the remote site is used for catalog backup. After takeover, the catalog backup can fail as the catalog backup policy is configured to use the storage unit on the remote site which was not available at the time of backup. (IA-30251)

Workaround:

It is recommended that SLPs should be used for catalog backup policy in disaster recovery configuration and catalog backup policy should be updated with reverse SLPs after migration or takeover process.

Takeover to a secondary cluster fails even after the primary cluster is completely powered off

The takeover to a secondary cluster fails even after primary cluster is completely powered off with the following error:

```
<cluster_name> cluster is either running or not completely down.  
Takeover of replication role is not permitted
```

In a rare scenario, the status of a remote cluster cannot be determined. The remote cluster status can be obtained from the console using the following command.

```
# haclus -display <remote_cluster_name> | grep ClusState
```

The takeover operation is permitted only when the remote cluster status is faulted or exited. But due to a bug, the remote cluster status shows the status as unknown even after the remote cluster is powered off. Also, if the remote cluster is powered off when the local cluster is down, the remote cluster status shows init. The takeover is not permitted in these scenarios. (4012004)

Workaround:

Run the following commands from the terminal of any node of the surviving cluster.

1. # /opt/VRTSnas/pysnas/bin/nso_replication.py --command online_nbu
2. # /opt/VRTSnas/pysnas/bin/nso_replication.py --command clear_host_cache

Catalog replication may fail to resume automatically after recovering from node fault that exceeds fault tolerance limit

Catalog replication stops when a fault exceeds fault tolerance limit and is resumed automatically after the faults are restored. But in some cases, the NetBackup Flex Scale cluster may fail to resume the catalog replication.

The following disaster-recovery status API can be used to get the status of catalog replication and the replication status shows that it is paused due to network disconnection.

API:

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

Response:

```
replicationStatus: paused due to network disconnection
```

(IA-32203)

Workaround:

1. Reboot all nodes of the NetBackup Flex Scale cluster on the site where the node fault tolerance exceeded the tolerance limit.
2. Run the following command to restart the cluster nodes.

```
# echo b > /proc/sysrq-trigger
```

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.

If both primary and secondary clusters are down and are brought online again, it may happen that the replication is in error state

If both primary and secondary clusters are down and are brought online again, it may happen that the replication state is displayed as Error in spite of other services on the clusters being in healthy state. This may be caused if the vxnetd start daemon is in hung state on any node of the cluster. (IA-36540)

Workaround:

Restart the nodes on which 'vxnetd start' daemon is in hung state.

Disaster recovery configuration fails if the lockdown mode on the secondary cluster is enterprise or compliance

If you start disaster recovery configuration on the primary cluster and the lockdown mode on the secondary cluster is not normal mode, then the configuration fails. (IA-36129)

Workaround:

Reconfigure the secondary cluster with normal mode and then perform disaster recovery configuration.

Unable to perform a takeover operation from the new site acting as the secondary

When the user performs the first takeover operation and brings the old primary site up again, the fbsync process gets started. If there is a situation where the new primary site goes down while the fbsync is in progress and if the user tries to perform the second takeover operation from the acting secondary, the takeover operations fails. (IA-36090)

Note: When the fbsync is in progress, the secondary cluster is called as acting secondary.

Workaround:

There is no workaround for this issue.

Enabling compliance mode for the first time on the secondary cluster may fail if disaster recovery is configured

When disaster recovery is configured between two NetBackup Flex Scale clusters, the compliance mode can be configured on each cluster independently to provide flexibility of compliance attributes for the disk pools on each cluster. But enabling compliance mode for first time may result in an error on the secondary cluster. This error occurs when the NetBackup primary server tries to re-fetch the compliance attributes from the disk pool that is created on secondary cluster. The disk pool on secondary storage server may provide stale information and give the `WORM Capable` setting as `No`. (IA-39351)

Workaround:

Re-fetch the compliance attributes manually. On the NetBackup Administration console, go to by **Media and Device Management > Disk Pools**. In the **Change Disk Pool** form, click **Refresh**.

If disaster recovery is configured and an upgrade from NetBackup Flex Scale 2.1 to 3.0, the upgrade operation hangs

If you perform an upgrade from NetBackup Flex Scale 2.1 to 3.0 and disaster recovery is configured on the cluster, the upgrade hangs. This happens because the `vxnetd start` daemon hangs on the secondary cluster. (IA-39899)

Workaround:

Restart the node on which `vxnetd start` daemon hangs. But this will cause the upgrade to fail and rollback. Once rollback is completed successfully, you can restart the upgrade process.

On a NetBackup Flex Scale cluster with disaster recovery configuration, the replication state shows Primary-Primary on the faulted primary cluster after takeover

Consider the following scenario on a NetBackup Flex Scale cluster on which disaster recovery is configured.

- The primary cluster (siteA) has faulted.
- A takeover operation is performed to make siteB as the primary cluster.
- The takeover operation is successful and the user continues to use siteB.
- The fault at siteA is restored and siteA is brought back online.
- The replication state on the primary catalog replication shows as `Primary-Primary`.
- So, the media and storage servers on siteA remain in Error state.

(IA-40346)

Workaround:

- Login to any node of siteA and elevate.
- Run the following command from the shell menu:

```
# /opt/VRTSnas/scripts/rep/nso_replication.sh prepare secondary
```

This step converts the siteA as the new secondary and starts replication resync.

- The replication state shows as `Resync` and moves to `Replicating` once resync has finished.
- The media and storage servers will become healthy after a few minutes.

Unable to create universal shares on a cluster on which disaster recovery is configured or only media server is deployed

If you try to create universal shares on a cluster on which disaster recovery is configured or only media server is deployed, the following error message is seen:

```
Failed to get version from the Storage Platform Web Service (SPWS).  
Ensure that Nginx is running and configured correctly on the selected  
MSDP storage server.
```

Workaround:

Contact Veritas Technical Support for troubleshooting this issue.

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.

cvmvoldg agent causes resource faults because the database not updated

As the file system goes offline, the deduplication engine associated with this file system became inactive and there is no backup running on this engine. (4027462)

Workaround:

Clear the faulted resource of the file system, and then offline the respective file system's service group. Once this is done online the file system service group.

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>

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.

NetBackup issues

The following known issues are related to NetBackup.

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

When the external CA certificate is deployed on NetBackup Flex Scale, and you go to **Security > Hosts** page in the NetBackup web UI, the **Host Management** 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**.
- 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/MASTER_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:1
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_master`) forcefully and wait for few minutes till the master container starts on its own.

```
# docker stop nb_master
```

Note: This needs the involvement of the Veritas Support.

NetBackup fails to discover VMware workloads in an IPv6 environment

This issue is applicable for VMware workloads and when the NetBackup Flex Scale cluster is configured to use IPv6 addresses. (4019408)

This issue occurs while assigning VMware workloads to a policy using the **Select automatically through Intelligent Policy query** option in the Java Administration Console UI (Policy properties > Clients tab). In the Java UI, when you click **Advanced Mode** and try to discover workloads using the VMware vApp ID in the search query, the discovery fails.

This happens because NetBackup services are unable to process IPv6 addresses in the expected syntax. As a result, the search query does not display any results in the UI.

Workaround:

In an IPv6 environment, use the **Basic Mode** option in the Java UI to discover VMware workloads using the vApp name in the search query.

Networking issues

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

DNS of container does not get updated when the DNS on the network is changed

When a new DNS server is added to the configuration, the containers do not become aware of it. Only the hosts know about the new DNS server. (IA-24663)

Workaround:

Stop and start each node one by one. This way the containers are restarted and they pick the updated `resolve.conf` file from the host. This causes job failures and jobs may or may not get automatically restarted based on the associated policy.

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.

Bond modify operation fails when you modify some bond mode options such as `xmit_hash_policy`

Some bond modes offer extra settings like `xmit_hash_policy`. During a `bond modify` operation, when such extra settings are updated without changing the bond mode, the operation fails with the *same bond mode* error.

This is because the extra setting is not considered during the modify operation. (IA-26730)

Workaround:

Perform the bond modify operation with a different bond mode. Then, perform the bond modify operation again with the original bond mode and new hash policy.

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`, 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>] vxfsn_process_client_msg+0x6d1/0xb30
[vxfsn]
[19742.723779] [<ffffffffffc1055d23>] vxfsn_vrfsn_cback+0x323/0x1750 [vxfsn]
[19742.723947] [<ffffffffffc1055a00>] ? vxfsn_reconfig_msg+0x840/0x840 [vxfsn]
[19742.724117] [<ffffffffffc1073be8>] vrfsn_step+0x1c8/0x3a0 [vxfsn]
[19742.724280] [<ffffffffffc1055a00>] ? vxfsn_reconfig_msg+0x840/0x840 [vxfsn]
[19742.724448] [<ffffffffffc1075521>] vrfsn_rcv_thread+0x401/0x9b0 [vxfsn]
[19742.724613] [<ffffffffffc1075120>] ? vrfsn_defer_message+0x140/0x140 [vxfsn]
[19742.724782] [<ffffffffffc10761ee>] vxplat_lx_thread_base+0x9e/0xf0 [vxfsn]
[19742.724947] [<ffffffffffc1076150>] ? vxplat_assert+0x20/0x20 [vxfsn]
[19742.725123] [<ffffffffffa6ec61f1>] kthread+0xd1/0xe0
[19742.725282] [<ffffffffffa6ec6120>] ? insert_kthread_work+0x40/0x40
[19742.725449] [<ffffffffffa758dd1d>] ret_from_fork_nospec_begin+0x7/0x21
[19742.725611] [<ffffffffffa6ec6120>] ? insert_kthread_work+0x40/0x40
```

Workaround:

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

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

<code>/log/VRTSns/log/storage_snapshot_destroy.log</code>	<code>/log/VRTSns/log/storage_snapshot.log</code>
<code>/log/VRTSns/log/storage_snapshot_create.log</code>	<code>/log/VRTSns/log/storage_snapshot.log</code>

Designated log file**Logs written to**

/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.

An NVMe disk is wrongly selected as a target disk while replacing a SAS SSD

If both NVMe disks and SAS SSDs fail and are physically replaced simultaneously on the cluster node, and then the SAS SSD is replaced first using the **Replace disk** option on the **Monitor > Infrastructure > Disks** tab in the NetBackup Flex Scale

UI, the NVMe disk might get selected as target disk. The data on the SAS SSD is rebuilt on the NVMe disk. Now, when you attempt to replace an NVMe disk, the **Replace disk** operation fails as a disk of similar type is no longer available on the node. (IA-27647)

Workaround

Replace the NVMe SSD first in such scenarios. After all the NVMe disks are replaced, replace the faulted SAS SSD with the newly added SAS SSDs.

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
```

Add node fails because of memory fragmentation

The add node operation fails as contiguous memory cannot be allocated because the memory is too fragmented. (1A-32612)

Workaround:

Contact Veritas Support to resolve the issue.

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.

Node is displayed as unhealthy if the node on which the management console is running is stopped

If you stop the node where the management console is running, the node is displayed as unhealthy and the option to start the node is no longer displayed in the UI. (IA-37062)

Workaround:

- Switch the management console to some other node before stopping the node.
- If the node was stopped before switching the management console, log in using the management console IP address and in the shell menu run the `run cluster start nodename` command where *nodename* is the name of the node that was stopped.

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:

When the node on which the management console is running is stopped, the nodes goes out cluster and the management console fails over to another node in the cluster. However, all the logs before the management console fails over cannot be collected using the UI; instead you need to use the following script to collect the logs:

- 1 Using SSH log in to the node and run the `support elevate` command.
- 2 Run the `/opt/VRTSnas/scripts/support/collect_debuginfo.sh` script and specify the following parameters:
 - `-o upload`: Upload the collected debug information for the specified modules to the provided location

- `-n nodename`: name of the node from where logs are to be collected
- `-d debug-URL`: URL in (ftp|scp|file)://url format
- `-m module`: comma separated supported module names (nas,os,explorer,install,nas-procstacks,netbackup,appliance,sds,api_gateway,upgrade,backup,vdd)|all
- `-t tar_name`: optional tar file name to create the file with tar_name_timestamp.tar.xz naming convention

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

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

Workaround:

Complete the following steps for all the cluster nodes. If you add or replace a node, ensure that you complete the following steps for the newly added node or the replacement node.

- 1 Using SSH log in to each cluster node and enter the `support elevate` command to access the root shell.
- 2 Edit the `/etc/logrotate.d/veritasaccess` file and add "**su root accessuser**" for each logrotate rule in the file.
- 3 Edit crontab for primary user and remove the entry "**@daily sudo /opt/VRTSnas/scripts/misc/key_check.sh**"

After replacing a node, the AutoSupport settings are not synchronized to the replacement node

When you replace a node, the email, SNMP, and Call Home settings are not synchronized to the replacement node, which is added to the cluster. This node is unable to send emails or alerts for the events that occur on the node. (IA-37656)

Workaround:

To synchronize the AutoSupport cluster settings with the replacement node, in the NetBackup Flex Scale UI, click **Settings > AutoSupport** and reconfigure the email, SNMP, and Call Home and proxy settings.

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.

The Add nodes to the cluster button remains disabled even after providing all the inputs

In the add node wizard, if you switch between the Custom and Automatic tabs in the NetBackup Flex Scale UI, **Add nodes to cluster the cluster** remains disabled even though all the required values are specified. (IA-40140)

Workaround:

On the **Data-Network-eth5** tab, delete the first entry, which is the IP address specified for the media server and specify the value again, or use the NetBackup Flex Scale infrastructure console instead of the NetBackup Flex Scale UI.

Unable to add more than seven nodes simultaneously to the cluster

The available IP addresses are less than the required IP addresses, which results in a failure during the add node operation. (IA-40368)

Workaround:

Add the nodes to the cluster in batches of seven or less.

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.

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/opensv/netbackup/bin/nbstserv -terminate
```

- Restart the nbstserv service.

```
/usr/opensv/netbackup/bin/nbstserv
```

Unable to switch the lockdown mode from normal to enterprise or compliance for a cluster that is deployed with only media servers and with lockdown mode set to normal

If you click **Settings > Security management > Immutability** in the UI and attempt to change the lockdown mode from normal to enterprise or from normal to compliance for a cluster that is deployed with only media servers, the WORM storage is not enabled. (4049911)

Workaround:

- 1 In the NetBackup Flex Scale UI, change the lockdown mode from normal to enterprise or compliance by clicking **Settings > Security management > Immutability > Edit**.
- 2 Run following command manually on the primary server:

```
# /usr/opensv/netbackup/bin/admincmd/nbdevconfig -updatedp -stype PureDisk -dp disk_pool
```
- 3 Go to the NetBackup Administration Console and in the left pane, expand **NetBackup Management > Storage > Storage Units**, and then select the **Use WORM** checkbox. Save the settings.

CRL mode does not get updated on the secondary site after ECA is renewed on a cluster on which disaster recovery is configured

When adding ECA certificate on a cluster on which disaster recovery is configured, you must not set the CRL mode as **No Change**. If you have set the CRL mode as **No Change** and you renew ECA, it fails as the secondary site does not have CRL mode entries. (IA-40121)

Workaround:

Choose the same CRL mode as what was selected on the primary before disaster recovery configuration.

Setting lockdown mode to enterprise or compliance fails on the secondary cluster of a NetBackup Flex Scale cluster on which disaster recovery is configured

This issue is applicable only to the secondary cluster of a NetBackup Flex Scale cluster on which disaster recovery is configured. On the secondary cluster, when the lockdown mode is set for the first time, it fails due to a bug. This leaves the system in such a state that all the successive attempts continue to fail unless at

least one of the retention values (minimum retention or maximum retention) is changed.

Additionally, you may have to manually mark the relevant STU for **Use WORM** on the NetBackup UI. Usually this is done automatically, but in some cases, manual action is needed. This happens because the API to update the storage server and STU from the Flex Scale cluster fails. You can refresh the storage server from the NetBackup primary server since it polls the attributes from storage server in Flex Scale cluster and gets the WORM attribute. However, there is no such polling for STU. Hence, any STU on such a storage server must be manually marked as WORM.

(IA-46555)

Workaround:

Retry setting the lockdown mode after changing at least one of the retention values (minimum retention or maximum retention). For example, if in the first attempt, the minimum retention was set to 1 days and maximum retention was set to 365 days, in the second attempt, set the minimum retention to 1 days and maximum retention to 366 days. For STU that needs to be WORM-enabled, mark the STU to **Use WORM** from the NetBackup UI.

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.

Upgrade issues

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

After an upgrade, if checkpoint is restored, backup and restore jobs may stop working

When a checkpoint is restored, the configuration files in the primary file system are replaced with old copies of the configuration files from the checkpoint. All the NetBackup processes that are in running state have to be restarted to run with the restored configuration files. (IA-27537)

Workaround:

After the sync-catalog operation is completed, perform the following steps:

1. Log on to the primary server container on the node on which the NetBackup primary worker group is online.

```
docker exec -it nb_master /bin/bash
```

2. Run the `/usr/opensv/netbackup/bin/bp.kill_all` command and wait for all the processes to stop.
3. Run the `/usr/opensv/netbackup/bin/bp.start_all` command.
4. Exit from the primary server container.

Upgrade fails during pre-flight in VCS service group checks even if the failover service group is ONLINE on a node, but FAULTED on another node

For failover service groups, the FAULTED state of a service group on a node is not an issue as long as the service group is online on another node in the cluster. The FAULTED state on the service group on a node can be cleared manually to pass the upgrade prechecks. (IA-30306, IA-40255)

Workaround:

1. Login in maintenance mode.
2. Identify the VCS service groups that are in FAULTED state by executing the following command:

```
hagrp -state | grep FAULTED
```

3. Check if the above service group is ONLINE on at least one node, but FAULTED on other nodes:

```
hagrp -state <service group>
```

4. From the Appliance GUI, go to **Support > Run autofix** to clear all the faults on the cluster.
5. If a service group is still in FAULTED state, login in maintenance mode and run the following command:

```
hagrp -clear <service group>
```

Upgrade from version 2.1 to 3.0 fails if the cluster is configured with an external certificate

Upgrading from version 2.1 to 3.0 fails while upgrading the NetBackup packages. (4064943)

Workaround:

There is no workaround for this issue.

During EEB installation, a hang is observed during the installation of the fourth EEB and the proxy log reports "Internal Server Error"

This issue occurs when the EEB installation is in-progress, and the upload operation gets triggered. This deletes some upgrade data state file and the EEB installation task is not marked done in GUI. (IA-30521)

Workaround:

After cluster configuration, when EEB installation is in-progress, do not trigger any other operation until the installation is completed.

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:

1. Using SSH log in to the appliance with the admin account. The Veritas Appliance Shell is displayed. Enter the following command:

```
support elevate
```

2. Manually disable the NetBackup Flex Scale health check on all the nodes so that NetBackup services do not get restarted automatically

```
hacli -cmd "docker ps -qf label=image.category=netbackup |  
xargs -i docker exec {} /opt/veritas/vxapp-manage/nbu-health disable"
```

3. Stop all NetBackup services on all the nodes. Run the following command multiple times until all the nodes reports 'No NB/MM daemons appear to be running'.

```
hacli -cmd "docker ps -qf label=image.category=netbackup |  
xargs -i docker exec {} /usr/openv/netbackup/bin/bp.kill_all FORCEKILL"
```

4. Install the EEB.
5. After the EEB is installed successfully, re-enable the NetBackup Flex Scale health check.

```
hacli -cmd "docker ps -qf label=image.category=netbackup |  
xargs -i docker exec {} /opt/veritas/vxapp-manage/nbu-health enable"
```

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.

After an upgrade Call Home does not work

In version 1.3.1, the AutoSupport client uses a token generated on each cluster node to communicate with the AutoSupport server. In version 3.0, instead of each node, the communication with the AutoSupport server is at cluster level, making the generated tokens invalid. (APPSOL-154989)

Workaround:

1 Using SSH log in to one of the cluster nodes.

2 Run the following commands:

```
hacli -cmd "rm /usr/opensw/runtime_data/"  
  
/opt/veritas/appliance/autosupport_alerts/scripts/alerts_ctrl.py  
callhome test
```

A message confirming Call Home test passed is displayed.

After an upgrade, the proxy server configured for Call Home is disabled but is displayed as enabled in the UI

After upgrading from version 1.3.1 to 3.0, the proxy server is disabled but wrongly shows as enabled in the UI. (APPSOL-154850)

Workaround:

In the UI, click **Settings > AutoSupport** and on the **Call home and proxy settings** tab disable the proxy server, and then enable the proxy server and specify the parameters again.

Unable to view the login banner after an upgrade

After upgrading from version 1.3.1 to 3.0, the login banner is not displayed in the NetBackup Administration Console and when you log in to the primary, media, and storage containers. (4051797)

Workaround:

There is no workaround for this issue.

After an upgrade to NetBackup Flex Scale 2.1, the metadata format in cloud storage of MSDP cloud volume is changed

If MSDP cloud volumes are configured in NetBackup Flex Scale 1.3/1.3.1 and you upgrade to NetBackup Flex Scale 2.1, the metadata format in cloud storage of MSDP cloud volume gets changed. It is recommended to follow the workaround to convert the old metadata format to new one. (4042920)

Note: Do not downgrade NetBackup Flex Scale 2.1 to 1.3/1.3.1 after the metadata conversion since 1.3/1.3.1 does not recognize the new metadata format in cloud storage.

Workaround: After the upgrade to NetBackup Flex Scale 2.1 is successfully completed, perform the following steps:

1. Using SSH log in to the appliance with the admin account. The Veritas Appliance Shell is displayed. Enter the following command:

```
support elevate
```

2. Manually run commands to start the conversion of cloud metadata format.

```
docker ps -qf ancestor=uss-engine | xargs -i docker exec {}  
/usr/opensv/pdde/pdcr/bin/cacontrol --catalog cloudmetadataconvertion
```

3. Run the following command to check the conversion status.

```
docker ps -qf ancestor=uss-engine | xargs -i docker exec {}  
/usr/opensv/pdde/pdcr/bin/cacontrol --catalog cloudmetadataconvertstatus
```

Rollback fails after a failed upgrade

An upgrade from version 1.3.1 to 3.0 fails, but you are unable to roll back to the earlier version if the Docker daemon is not configured correctly. (IA-36990)

Workaround:

Contact Veritas Support to resolve the issue.

Add node operation hangs on the secondary site after an upgrade

After upgrading from version 1.3.1 to 3.0 the add node operation might hang on the secondary site. (IA-37443)

Workaround:

- 1 Using SSH, log in to a node on the secondary site.
- 2 Kill the `cluster_reconfig.py` and `pkg_manager` processes.
- 3 Delete the `/shared/isagui/cluster_config/gui_oper_progress.json` file.

Alerts about inconsistent login banner and password policy appear after an upgrade from NetBackup Flex Scale version 2.1 to NetBackup Flex Scale 3.0

After an upgrade is successfully performed from NetBackup Flex Scale version 2.1 to 3.0, you may see two alerts about inconsistent login banner and inconsistent password policy. (IA-40218)

Workaround:

Go to **Settings > Services management** tab and perform **Run full discovery** to fix this issue.

Alerts about node being down are generated during an upgrade

Node down alerts are wrongly generated during an upgrade. (IA-40370)

Workaround:

There is no workaround for this issue.

GUI takes a long time to update the status of the upgrade task

When you upgrade to version 3.0, the upgrade completes successfully, but the upgrade task in **Recent Activity** is not updated. (IA-40381)

Workaround:

To update the task status in the GUI, run the following command on the management console node:

```
# systemctl restart crond
```

In a disaster recovery environment, upgrade gets stuck during node evacuation stage as VVRInfra_Grp cannot be brought down

During upgrade, all the service groups are brought offline on each node and then the node is evacuated from the cluster. The upgrade process waits for all service groups to be offline for the evacuation to succeed on that node. The VVRInfra_Grp service group may be stuck and can fail to go offline because of the following reasons:

- A resource of type RVGShared is present, which tries to restart vradmind service as soon as it finds the that service is not running.
- `/opt/VRTSnas/scripts/rep/nso_replication.sh auto_resyncd` is in running state and restarts the vradmind service.

(IA-40311)

Workaround:

1. Login to Appliance shell using admin user credentials on any node in the cluster except for the node on which the evacuation is not hung.
2. Enter the `support elevate` command to elevate to maintenance mode.
3. Enter the maintenance user account password to log into the maintenance mode.
4. On the bash shell, execute the following commands in the specified order.

```
■ # /opt/VRTSvcs/bin/haconf -makerw

■ # /opt/VRTSvcs/bin/hares -modify VVRInfra_Grp-vvr_daemon
  StartProgram "/usr/sbin/vxstart_vvr start && /bin/touch
  /var/run/VRTSvcs/vvr_service"

■ # /opt/VRTSvcs/bin/hares -modify VVRInfra_Grp-vvr_daemon
  StopProgram "/usr/sbin/vxstart_vvr stop && /bin/rm -f
  /var/run/VRTSvcs/vvr_service"

■ # /opt/VRTSvcs/bin/hares -modify VVRInfra_Grp-vvr_daemon
  CleanProgram "/usr/sbin/vxstart_vvr stop && /bin/rm -f
  /var/run/VRTSvcs/vvr_service"

■ # /opt/VRTSvcs/bin/haconf -dump -makero
```

5. Log into the maintenance mode and execute the following command on the node on which the evacuation is hung.

```
# rm -f /var/run/VRTSvcs/vvr_service
```

Upgrade may fail after node evacuation if a VCS parallel service group is OFFLINE on a partial set of nodes at the beginning of the upgrade

If a VCS parallel service group is in OFFLINE state on few nodes (partial set of nodes) and ONLINE on other nodes at the start of the upgrade, upgrade may fail after the node evacuation step (`hastop -evacuate`). This happens because as a part of the VCS service group check after the `hastop -evacuate` command is executed, the parallel service group should be ONLINE on all the nodes. But if such a group was in OFFLINE state on a partial set of nodes at the beginning of the

upgrade, it remains in OFFLINE state on those nodes and the VCS service group check fails after waiting for 3-4 hours. (IA-40597)

Workaround:

Perform the following steps and restart upgrade.

1. Identify the parallel service groups.

```
# hagrps -display | grep -i parallel
```

Service groups with 1 in the 4th column are parallel service groups.

For example:

```
# hagrps -display | grep -i parallel
```

CanHostNLM	Parallel	global	1
GLOBAL_API_SERVER	Parallel	global	0
ManagementConsole	Parallel	global	0
NBUMasterBrain	Parallel	global	0
NBUMasterWorker	Parallel	global	0
NFSShareOfflineGrp	Parallel	global	0
NLMGroup	Parallel	global	0
NicMonitorGrp	Parallel	global	1
Phantomgroup_pubeth0	Parallel	global	1
Phantomgroup_pubeth1	Parallel	global	1
Phantomgroup_pubeth2	Parallel	global	1
VVRInfra_Grp	Parallel	global	1

2. Before you start the upgrade, check if any of the parallel service group is OFFLINE on partial set of nodes and online on other nodes.

```
# hagrps -state <SG name>
```

For example:

```
# hagrps -state VVRInfra_Grp
```

#Group	Attribute	System	Value
VVRInfra_Grp	State	dellsitel1-01	ONLINE
VVRInfra_Grp	State	dellsitel1-02	OFFLINE
VVRInfra_Grp	State	dellsitel1-03	OFFLINE
VVRInfra_Grp	State	dellsitel1-04	ONLINE

3. Try to online the parallel service groups on all the nodes of the cluster before upgrade.

```
# hagrps -online <SG Name> -sys <nodename>
```

For example:

```
# hagrps -online VVRInfra_Grp -sys dellsite1-02
```

Upgrade may fail if operations such as OS reboot, cluster restart, and node stop and shutdown are used during the upgrade

Performing any of the following operations during an upgrade can lead to an upgrade failure and the cluster might go in an inconsistent state. It is recommended that these operations should not be performed until the upgrade completes successfully. (IA-40377)

- Logging into maintenance mode and executing OS commands such as reboot and shutdown.
- Logging into a node and executing commands such as cluster reboot, shutdown, and stop from the command-line interface.
- Executing node maintenance operations from the UI such as shutdown node and start node.
- Operations that will disrupt network connections such as unplugging the network cable.

Workaround:

There is no workaround for this issue.

UI issues

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

In-progress user creation tasks disappear from the infrastructure UI if the management console node restarts abruptly

This issue occurs if the node where the management console service is running fails or restarts abruptly. (IA-25874)

The in-progress user creation tasks disappear from the tasks list displayed in the NetBackup Flex Scale infrastructure UI. Even though the task itself completes successfully, the UI fails to display it. This issue applies for in-progress tasks, events, and alerts.

Workaround:

There is no workaround for this issue at the moment.

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.

User management issues

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

AD server test connection fails due to incorrect username on the IPV6 media only cluster

During AD server test connection, the pop-up auto-populates the wrong username on the IPV6 media only cluster. It populates `cn=BINDDN+UserBaseDN` whereas it should populate `cn=BindDN` only. Hence, the test connection fails. (IA-39642)

Workaround:

Delete the auto-populated entry and manually enter `BindDN` as the username.

AD/LDAP domain unreachable alerts do not get cleared after the AD/LDAP server is deleted

When a new AD/LDAP server is added, if the domain is not reachable from the node, the following alerts are generated in the GUI:

```
Domain XXX is not reachable from node YYY.
```

If the server which is not reachable from the node is deleted, all the pending alerts should get cleared from the GUI. But this does not happen and the alerts are visible in the **Alert management** screen. (IA-39890)

Workaround:

There is no workaround to clear these alerts from the GUI. You can ignore the alerts based on the time stamps.

GUI login fails with LDAP user if the domain is configured with SSL

Login fails if an LDAP user tries to log in from the GUI when SSL is set. (IA-40332)

Workaround:

There is no workaround for this issue.

Assigning role to correct AD/LDAP user/group with wrong domain causes the user listing to fail

This issue affects NetBackup Flex Scale cluster which are deployed with only media servers. Role assignment takes `<user/groupname@domainid>` as input. Currently, the validation for domain id is not done. Hence, valid domain users/groups with wrong domain are assigned with appliance administrator role. This causes user listing in **User management** page to fail. (IA-40058)

Workaround:

Delete the appliance administrator role for that user/group from CLISH.

To delete appliance administrator role for a user:

```
admin user modify role delete username@doamin primary ldap
```

To delete appliance administrator role for a group:

```
admin user modify role delete groupname@doamin primary ldap
```

You can delete the appliance administrator role for that user/group with APIs.

```
DELETE /api/appliance/v1.0/users/{userName}/{userRole}
```

After cluster reboot all/shutdown all operation, AD/LDAP domains become unreachable from one or more nodes on a NetBackup Flex Scale cluster on which only media servers are deployed

This issue is seen in a NetBackup Flex Scale cluster on which only media servers are deployed. After `cluster reboot all/shutdown all` operation, when all the nodes are up and running, alerts related to AD/LDAP domains being unreachable from one or more nodes can be seen. Hence the SSH login to those nodes using AD/LDAP users does not work. Also, if the Appliance GUI is hosted on one of those nodes, login to Appliance GUI/REST APIs does not work with AD/LDAP users. (IA-40062)

Workaround:

Log in to Appliance GUI using one of the local user admin credentials. Go to **Settings > Services Management**. Select **Run auto fix**.

You can also use REST APIs to restore connectivity to the AD/LDAP domains. Authenticate a local user with appliance administrator role and execute the autofix API:

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

Fixed issues

This chapter includes the following topics:

- [AD/LDAP configuration may fail for IPv6 addresses](#)
- [Fixed issues in version 3.0](#)

AD/LDAP configuration may fail for IPv6 addresses

Setting up an AD/LDAP configuration from the NetBackup Flex Scale infrastructure management UI fails if the specified IP address of the AD/LDAP server is an IPv6 address. (4020899)

This issue occurs because the NetBackup services are not able to process the IPv6 address in the expected syntax. The colon character is not supported in domain name for role assignments from NetBackup. This causes a connection failure with the AD/LDAP server and the AD/LDAP configuration cannot be initialized.

Workaround:

Instead of an IPv6 address, specify the FQDN of the AD/LDAP server in the NetBackup Flex Scale infrastructure management UI to successfully set up the AD/LDAP configuration.

Fixed issues in version 3.0

The following issues are fixed in this release:

Table 5-1

ID	Description
4050207	Self test fails in a media server only deployment.

Table 5-1 (continued)

ID	Description
APPSOL-153513	Cluster configuration fails with error docker createOneFile failed for file /mnt/nblogs/.tpconfig_cred in container nb_media.
IA-25220	Unable to start a node that is shut down.
IA-32031	Initial configuration of the cluster fails if multiple DNS servers are specified in the YML configuration file.
IA-36631	If disaster recovery is configured and an upgrade is performed, an alert appears that the master server is offline on the secondary cluster
IA-36790	After you upgrade to NetBackup Flex Scale 2.1 from 1.3.1 version, the replication is in paused state.
IA-36943	Unable to configure and create AD/LDAP users on the secondary cluster if disaster recovery is configured.
IA-36833	After upgrade is performed, the 'Executing the full discovery' task appears as ongoing under RecentActivity.
IA-36963	During upgrade on a cluster on which disaster recovery is configured, the upgrade task may be triggered again.