

Veritas™ HyperScale 1.1 for OpenStack Release Notes

Ubuntu

Veritas HyperScale for OpenStack Release Notes

Last updated: 2017-09-13

Document version: 1.1 Rev 0

Legal Notice

Copyright © 2017 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Veritas InfoScale, 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
500 E Middlefield Road
Mountain View, CA 94043

<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 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:

doc.feedback@veritas.com

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

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

Release Notes

This document includes the following topics:

- [About this document](#)
- [Software limitations](#)
- [Hardware limitations](#)
- [Known issues](#)

About this document

This document provides important information about Veritas HyperScale for OpenStack version 1.1. Review this entire document before you install the product.

The information in this document supersedes the information provided in the product-specific user documentation.

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

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

For the latest information about updates, patches, and software issues regarding this release, see the following Late Breaking News (LBN):

<http://www.veritas.com/docs/000127549>

Software limitations

This section describes the software limitations in the current HyperScale release.

OpenStack deployments containing multiple domains not supported

HyperScale does not support OpenStack environments where the Keystone Identity service is configured for multiple domains with domain-specific configuration files.

"Double failures" are not supported

Sometimes, two failures can occur in the same workflow. For example, if the storage service on the compute node fails, and the `dml-d` process fails during the failback/failover operation, the failback/failover does not complete. It remains in the `Pending` state. Currently, HyperScale cannot recover from double failures like these.

After you install HyperScale, do not change the OpenStack administrator password

If you install HyperScale and then change the OpenStack administrator password, HyperScale does not function as expected. After you install HyperScale, do not change the OpenStack password. This operation is not supported.

The storage service cannot be restarted

The following storage restart operation is not supported:

```
systemctl hyperscale-mq-storage restart
```

Required OpenStack configuration changes (1906)

When you create a volume snapshot in OpenStack, the size of the snapshot is equal to the value that is specified in the volume size quota. This quota is referred to as the Gigabytes quota in the OpenStack API and command line interface.

For example, if the volume quota is 1000 GB and you create a 100 GB volume, you can create 10 snapshots before you reach the quota limit ($100 \times 10 = 1000$). If you try to create more volumes or snapshots, the operation fails.

To override this limitation, do the following:

- 1 On a node that has the cinder volume service configured, set the following variable in the `cinder.conf` file:

```
no_snapshot_gb_quota=true
```

- 2 Restart cinder block storage services. Enter the following:

```
openstack-service restart cinder
```

Specifying a destination host name for live migration is not supported using the command line interface (CLI) (4017)

In the current HyperScale release, when you live migrate a virtual machine, you cannot specify a destination host name for the `nova live-migration` CLI. HyperScale selects a host from its list of reflection targets.

Controller reboot not supported (6975)

Controller reboot is not fully supported in this release. If the controller node is restarted, the controller may go in to an unpredictable state.

The HyperScale services may not start in the correct order and the HyperScale storage driver may fail to register with the Cinder service. As a result, instance creation may fail due to storage errors.

Veritas recommends that you use the following command for handling planned controller reboot requirements.

Run the following command on the controller node and observe the command output:

```
# hyperscale get-workflows-in-progress
```

If the command output is `TRUE`, it indicates that there is at least one workflow in progress and the controller should not be restarted.

If the command output is `FALSE`, it indicates that there is no active workflow in the pending state and the controller can safely be restarted.

Hardware limitations

This section describes hardware limitations in the current HyperScale release.

Data node disk limitation

Disks on the data node server cannot be part of a logical volume manager (LVM) or any existing disk group.

System with only SSD storage cannot be a data node

A HyperScale data node uses Solid-state drive (SSD) for writing metadata and hard disk drive (HDD) for storing the data. A data node with only SSD and no HDD storage cannot be configured in a HyperScale environment.

Note: A data node with only HDD storage can be used. However, performance benefits using an SSD storage will not be seen in an all-HDD storage configuration.

Known issues

This section describes HyperScale known issues and provides workarounds when they are available.

General issues

This sections describes some general issues related to HyperScale.

VxExplorer report does not include required HyperScale crash directories (6638)

Veritas provides a Data Collector that includes a VxExplorer utility that is used to collect various details such as product information, logs, and debugging information to investigate and troubleshoot issues with your systems. The Data Collector contains a functionality that allows you to collect and submit a VxExplorer archive to technical support for further analysis.

When you run the data collector to gather a VxExplorer report on a HyperScale system, it is observed that VxExplorer archive does not include directories that contain the required crash dump data. This crash data is necessary for performing an accurate root cause analysis about the product issues on the system.

Workaround: You must manually copy all the crash data directories to the VxExplorer archive before submitting it to technical support.

The directory that contains the crash data depends on the Apport configuration settings on the system.

OpenStack features disabled in the current HyperScale release

In the current release of HyperScale, the following OpenStack operations are disabled. They are grayed-out on the Horizon dashboard.

- Evacuate host (**Admin > System > Hypervisors > Compute hosts**)
Currently, if a compute node fails, you cannot evacuate the entire host in one operation. Instead, you must evacuate each instance one at a time using the `nova evacuate` command from the command line interface.

Note: In this release of HyperScale, the `--on-shared-storage` flag is not supported.

- Resize Instance, Rebuild Instance (**Project > Compute > Instances**)
- Manage Attachments (**Project > Compute > Volumes**)

In addition, you cannot create a HyperScale volume from an existing OpenStack volume.

OpenStack features not supported in HyperScale (3283)

The following OpenStack features are not supported in the current HyperScale release, either from the command line or from the graphical user interface:

- Resizing an instance
- Rebuilding an instance
- Live migrating instances is not supported for the VRTSBronze flavor; this feature is supported for the VRTSGold and VRTSilver nova flavors
- Detaching a volume
- Resizing a volume
- Extending a volume (**Project > Compute > Volumes > Volumes > Extend Volume**)
- Modifying the `cinderquota` from the command line is not reflected on the graphical user interface (and vice versa)

OpenStack security issue

OpenStack allows users to upload malware-infected files. This is a high severity issue; however, is it beyond the scope of HyperScale to address.

Installation and update issues

This section describes issues related to installing and updating HyperScale.

Automatic subscription updates may cause issues with HyperScale packages (3766379)

HyperScale ships its own modified versions of the `libvirt` package and the `qemu-kvm` package. If your system is registered for Linux updates (such as a RHEL subscription), these packages may be updated to another version and HyperScale cannot function properly.

Workaround: Make sure that your system is not set up to receive subscription updates.

Note: As an additional resolution to this issue, HyperScale has disabled `yum` plug ins. In the `/etc/yum.conf` file, the value for `plug-ins` is set to `0`. This change prevents HyperScale packages from being automatically updated. This change is maintained whether you add or remove compute nodes from your HyperScale deployment.

After you install HyperScale, some OpenStack commands do not work (3803)

After you install HyperScale, the following OpenStack command line interface (CLI) commands do not work:

```
cinder snapshot-reset-state [--state <state>] <snapshot> [<snapshot> ...]
cinder snapshot-delete <snapshot> [<snapshot> ...] --force [<True|False>]
cinder --os-volume-api-version 2 snapshot-reset-state [--state <state>]
  <snapshot> [<snapshot> ...]
cinder --os-volume-api-version 2 snapshot-delete <snapshot>
  [<snapshot> ...] --force [<True|False>]
cinder snapshot-metadata <snapshot> <action> <key=value> [<key=value> ...]
cinder snapshot-metadata-update-all <snapshot> <key=value> [<key=value> ...]
cinder --os-volume-api-version 2 snapshot-metadata <snapshot> <action>
  <key=value> [<key=value> ...]
cinder --os-volume-api-version 2 snapshot-metadata-update-all <snapshot>
  <key=value> [<key=value> ...]
cinder extend <volume> <new_size>
cinder --os-volume-api-version 2 extend <volume> <new_size>
cinder snapshot-list [--all-tenants [<0|1>]] [--name <name>]
  [--status <status>] [--volume-id <volume-id>] [--marker <marker>]
  [--limit <limit>]
cinder --os-volume-api-version 2 snapshot-list [-all-tenants [<0|1>]]
  [--name <name>] [--status <status>] [--volume-id <volume-id>]
  [--marker <marker>] [--limit <limit>]
```

Workaround: None.

User interface issues

This section describes issues related to the user interface.

Issue while creating an image from the OpenStack UI (6783)

This issue occurs when you try to create an image from the OpenStack Horizon user interface (**Project > Images > Create Image**).

On the Create An Image dialog, when you specify all the details and click Create Image, the UI repeatedly displays an error asking you to choose the image format type. This occurs even if you have specified the correct image format type.

In the process, the UI also loses the image file you selected, and you have to choose the image file each time there is an error.

This error is related with the order in which you specify the image details, and is mainly associated to the two fields—**Image File** and **Format**, on the **Create An Image** dialog box.

If you specify the image file first and then select the image file format, the UI fails to process the details. As a result, the image creation task remains incomplete.

Workaround:

To successfully create an image from the Horizon UI, specify the details in the following order:

- 1 First, select a random format from the **Format** drop-down list (for example, Raw).
- 2 Then, click **Choose File** and specify the desired image file (for example, an .iso or .qcow2 file).

You can also specify other details such as Name, Description, Minimum Disk size, as per your requirement.

- 3 Then, from the **Format** drop-down list, select the image format type again. This time, ensure that the format matches the image file that you specified earlier.
- 4 Click **Create Image**.

The UI now processes the image file details and creates the image successfully.

Horizon incorrectly allows editing the default HyperScale metadata associated with a volume (6017)

When you create a volume in a HyperScale environment, it is automatically associated with a default set of HyperScale metadata. The metadata includes attributes such as the primary data node IP address, the current data node IP address and ID, and the image ID associated with the vdisk.

The OpenStack Horizon user interface (UI) currently allows editing the default HyperScale metadata values (**Project > Volumes | Actions > Update Metadata**).

However, changing this metadata can cause configuration issues and may eventually result in a failure in the HyperScale functionality.

Workaround: There is no workaround available at this time. Changing default HyperScale metadata values from the UI is not recommended.

On the user interface, a snapshot remains in the Creating state

This issue occurs if an I/O failover/failback workflow is triggered for a vDisk whose episodic data synchronization (EDS) is in progress. The snapshot remains in the **Creating** state on the user interface.

Workaround: None.

The user interface does not notify you if live migration fails (3753305, 1697)

HyperScale uses the Liberty version of OpenStack. In this version, there is no way to alert the controller if live migration fails. Therefore, the interface does not display an alert if migration fails. Instead, it displays the following:

- The virtual machine's **Status** column displays **Error**.
- The **Task** column displays **Migrating**.

To determine whether or not live migration was successful, note the following:

- If the migration is successful, the virtual machine appears on the target hypervisor.
- If the migration fails, the virtual machine appears on the source hypervisor.

Workaround: None. This limitation is consistent with OpenStack Liberty behavior.

An SSD storage alert may be displayed on the dashboard even though storage use is less than 90% (2654)

HyperScale displays an alert on the dashboard when solid-state drive (SSD) use exceeds 90%. However, an alert can appear even if storage use is less than 90%. This issue can occur because HyperScale only uses the amount of write-back cache for displaying alerts, while the doughnut chart on the dashboard calculates both write-back and read cache.

Workaround: None.

The storage dashboard volume count may not be accurate (2655)

The volume count on the storage dashboard includes volumes that were intended to be deleted but are in the `error_deleting` state. These volumes should not be counted in the volume count.

Workaround: None.

Some storage media failures are not displayed as alerts (3622)

If the episodic data synchronization (EDS) to the data node fails due to a failure in the solid-state drive, HyperScale does not display an alert about the storage media failure. The EDS fails until I/O from a virtual machine or application triggers the storage media failover process. After the failover completes successfully, EDS starts working properly.

Workaround: If you see alerts for a virtual machine the **Alert/Events** area of the HyperScale dashboard, you can live migrate the virtual machine.

The user interface shows that a volume snapshot is available when it is still in progress (3785)

On the user interface, the **Images** page (path **Project > Compute > Images**) lists the image names and their status. The **Images** page may indicate that a volume snapshot is available when the snapshot is still in progress.

Workaround: To display the correct status of the snapshot, go to the instance version time line, where the virtual machine snapshot status is **Creating**.

When volume or volume snapshot quotas need to be raised, alerts are not displayed (4241)

Alerts are not displayed when a volume quota or volume snapshot quota nears its limit. When the volume quota or volume snapshot quota limit is reached, the following can occur:

- Episodic data synchronization (EDS) stops.
- Live migration, user snapshots, and backups are affected.
- Other workflows that involve creating volumes or volume snapshots fail.

Workaround: Increase the quota values to a sufficiently large number.

The user interface may not display the correct format and size for an image (4312)

If you take an instance-level snapshot, and then view the image information on the **Project > Compute > Images** page, you may notice that the image's format is not displayed and its size is 0 bytes.

This behavior is expected based on the way OpenStack and HyperScale store their images.

OpenStack stores image in qcow2 format in its glance store for boot volumes. The user interface shows the number of bytes that this snapshot image occupies in the glance store. However, HyperScale does not store its images in glance. It stores them in cinder. Therefore, the image's glance size is displayed as 0.

Workaround: None.

Solid-state disk (SSD) failures are not cleared from the Alerts/Events area after the SSD is reactivated (4441)

If an SSD fails a on compute node, the following occurs:

- If a virtual machine on the compute node has a reflection factor of 1 or above, its storage is served from the reflection compute node.
- No new virtual machines are provisioned on the compute node that has the failed SSD.
- An alert is displayed in the **Alerts/Events** area of the HyperScale dashboard indicating that the SSD failed.

However, if you correct the issue, either by recovering the SSD or removing the compute node from the configuration, the alert still displays.

Workaround: Remove the alert manually from the HyperScale database.

To remove the alert manually from the HyperScale database

- 1 On the compute node, log in as `root`.
- 2 Log in to the database monitor.

```
# mysql
Welcome to the MariaDB monitor. Commands end with ; or \g.
MariaDB[(none)
```

3 Access the HyperScale database.

```
MariaDB [(none)]> use HyperScale
Database changed
MariaDB [HyperScale]>
```

4 Determine the globally unique identifier (GUID) of the hypervisor running on the compute node. Enter the following query:

```
MariaDB [HyperScale]> select * from hyperscale_nodes;
```

In the output, locate the `hypervisor_id` column and note the GUID.

5 Remove the alert from the HyperScale database `hyperscale_alerts` table manually. Enter the following command, substituting the GUID you found in step 4 for `hypervisor_guid`.

```
delete from hyperscale_alerts where topic='alert.compute.ssd.failed'
and assoc_object_id='hypervisor_guid'
```

The storage dashboard may not display the correct number of snapshots (4463)

The OpenStack `cinder.conf` file includes the parameter `osapi_max_limit`. This parameter determines the number of items that are returned for the API. If the actual number of snapshots exceeds this value, only 1000 are accounted for, and any others are excluded. This results in inaccurate values for snapshots on the HyperScale storage dashboard (**HyperScale > System > Overview**). Therefore, the dashboard may show fewer snapshots than you have.

Workaround: To resolve this issue do the following:

- 1 Edit the `cinder.conf` file and increase the `osapi_max_limit` value.
- 2 Restart the cinder storage service to put your change into effect. Enter:

```
openstack-service restart cinder
```

Alerts for VRTSBronze virtual machines are not cleared from the user interface (4546)

This issue applies to virtual machine with the HyperScale VRTSBronze flavor. Some alerts continue to display after they are no longer relevant. They are displayed in the **Alerts/Events** area of the following pages:

- **HyperScale > System > Overview**

■ HyperScale > System > Primary

For example, if an alert appears because a service has been restarted on a VRTSBronze virtual machine, it is still displayed after the VRTSBronze virtual machine is deleted. The alert should be cleared.

Workaround: None.

The dashboard shows incorrect primary persistent storage data (4585)

The persistent storage data for the **Overall** and **Primary** views is incorrect. This issue occurs because the dashboard displays the amount of storage the hard disk drive has allocated rather than the actual disk usage.

Workaround: None.

Disabling HyperScale on a compute node is considered an alert (4661)

When you disable HyperScale on a compute node, a **Service is down** alert is displayed in the **Alerts/Events** area of the HyperScale dashboard (**HyperScale > System > Overview**). In this case, however, HyperScale is disabled intentionally, so it is not an issue, and the alert is not appropriate.

Workaround: None.

A backup snapshot remains on the Images page, even though the backup failed (6484, 6548)

In this scenario, a backup operation fails because the staging location on the NetBackup server is completely filled. When the operation fails, the backup snapshot that was in progress should be deleted from the Images page on the HyperScale user interface. However, the snapshot remains.

Workaround: None.

Compute node issues

This section describes issues related to compute nodes.

Enable HyperScale operation on a compute node might fail due to vrtsveki package removal issue (5397)

This issue occurs when you try to enable HyperScale on a compute node where HyperScale was earlier disabled. The enable HyperScale operation may sometimes fail due to the vrtsveki package.

Messages similar to the following may appear in the compute node log file:

```
'Removing vrtsveki (7.2.0.000-1) ...', 'Unable to unload the veki module.  
A reboot may be required to unload it', 'dpkg: error processing package  
vrtsveki (--purge):', ' subprocess installed pre-removal script  
returned error exit status 1  
, 'Starting veki ', 'Errors were encountered while processing:  
, ' vrtsveki', '  
<drive>: Sub-process /usr/bin/dpkg returned an error code (1)']
```

vrtsveki is one of the several packages that get removed when you disable HyperScale on a compute node. In some cases, the vrtsveki package removal may fail and the stale entry leads to an issue while enabling HyperScale on that compute node.

Workaround: Reboot the compute node and then enable HyperScale again. A reboot will successfully remove the vrtsveki package from the system.

A virtual disk may lose its resiliency, but the application stays online (2667)

This issue occurs when the storage service on a compute node fails, and the failure triggers the failover and failback workflow. The compute node stops serving as a reflection target for the current running virtual machines.

For example:

- Compute node A owns virtual disk A. Compute node B owns virtual disk B.
- Compute node A is the reflection target for virtual disk B. Compute node B is the reflection target for virtual disk A.
- If the storage service on compute node A fails, two things occur. Virtual disk A fails over to compute node B, and compute node B loses its reflection target. Compute node B loses its resiliency but the application continues to run.
- When compute node A comes back up, virtual disk A fails back.

Workaround: None.

In certain circumstances, manual restoration of reflection on an available compute node fails (4428)

This issue occurs in the following circumstances:

- A HyperScale deployment includes three compute nodes: Node A, Node B, and Node C.
- Bring down Node C's management port.

- On Node A, launch a virtual machine with the flavor VRTSSilver. Node A is the reflection source, and Node B is the reflection target.
- Add storage to Node A, and start I/O on it.
- Kill the storage service on Node B (the reflection target). This triggers restoration of reflection (RoR), but the RoR workflow fails.
- Bring up Node C's management port to make it to the other compute nodes. Although the port comes up, the RabbitMQ messaging service does not restart. This is RabbitMQ's expected behavior.
- Now that Node C is available, trigger RoR manually.
- Manual RoR fails.

Workaround: None. This issue is caused by the current behavior of the RabbitMQ service.

Disabling a compute node removes packages and breaks the nova service (4595)

When you disable a compute node, HyperScale removes the `libvirt` and `qemu` packages that were installed when the compute node was enabled. If the `yum` repository is not set or accessible, the HyperScale-specific `libvirt` and `qemu` packages are not installed. The nova service fails to start.

Workaround:

To resolve this issue

- 1 On the compute node, install the `libvirt` and `qemu` packages:

```
# /usr/bin/yum -y install libvirt libvirt-client \  
libvirt-python libvirt-debuginfo libvirt-devel libvirt-java \  
libvirt-java-devel  
# /usr/bin/yum -y install qemu-kvm-rhev qemu-img-rhev \  
qemu-kvm-rhev-debuginfo qemu-kvm-rhev-tools qemu-guest-agent
```

- 2 Start `libvirt`.

```
# systemctl start libvirtd
```

- 3 Restart the OpenStack service.

```
# openstack-service restart
```

Disabling a compute node may hang due to a storage media failure (6355)

If the storage media fails on a compute node, and then you try to disable the node, the operation may hang. You can verify this by reviewing the disable compute logs on the controller node. The `storage_unsetup.sh` operation is hung.

Workaround: Live migrate the virtual machines on the compute node, restart the node, and retry the disable operation.

The host evacuate feature is not supported (6539)

Currently, if a compute node fails, you cannot evacuate the entire host in one operation. Instead, you must evacuate each instance one at a time using the `nova evacuate` command from the command line interface.

Manual restoration of reflection may fail (6941)

This issue can occur in the following circumstances:

1. In a two compute node environment, one of the nodes runs a virtual machine with the `VRTSSilver` flavor.
2. The node running the virtual machine goes down, and evacuation is triggered.
3. The virtual machine is evacuated successfully to the other node.
4. While the original node is down, a vDisk is attached to the virtual machine.
5. The failed node is brought up.
6. Restoration of reflection is triggered manually. It fails, and the following exception is included in the log file:

```
2017-05-29 12:37:11,119 - Tid-139709716100864 -  
ERROR - Failed for external function call with exception  
'unicode' object has no attribute 'iteritems'
```

Workaround: Manually restore reflection to the node before you add the vDisk to the virtual machine whose reflection is compromised.

Data node issues

This section describes issues related to data nodes.

Data node appears as down after the data network interface is taken down and brought up again (6824)

On a data node, when you take the data network interface down and then bring it up again, the status of the data node appears as down. Even though the HyperScale services are active and running, the data node appears as down.

This issue occurs when the controller and the data node are set up on the same system.

Workaround: Restart the `hyperscale serengeti` service on the data node.

Run the following command:

```
# systemctl start hyperscale-serengeti.service
```

After the service restarts successfully, the status of the data node should reflect as up.

Adding more than one SSD to the data node may cause a loss of meta data (6967)

While adding a data node if you specify multiple SSD disks for the meta disks, HyperScale automatically creates mirrored volumes on them. However, it is observed that SSDs used for mirroring inadvertently get assigned to the data disks pool. This may sometimes result in a loss of meta data.

Workaround: Veritas recommends that you do not use SSD mirroring for meta data. While specifying storage for meta disks, do not select multiple SSD disks.

Issue due to data network failure on the reflection target data node (5596)

This issue occurs in an environment with two data nodes where one data node (primary) serves all the instance data and the other data node (secondary) acts only as a reflection target.

If there is a data network failure on the reflection target node (secondary data node), reflection from the primary gets disabled and the vdisks associated with the instances are marked as degraded.

The following reflection target failure alert is displayed in Horizon (**HyperScale > System > Overview | Workload**):

```
Datanode reflection failure for vdisk <vdisk_name> in VM <vm_name>.
```

To restore the reflection, the target data node must be removed and then added back in to the configuration.

However, as the status of the reflection target data node appears as up and running, you cannot remove that data node from the configuration.

Workaround:

Perform the following steps for the data node that serves as the reflection target:

- 1 Stop the HyperScale storage service on the reflection target data node:

```
# service hyperscale-dmld stop
```

The dmld service may restart automatically. Run this command multiple times to ensure that the service has stopped.

- 2 Confirm that the service has stopped and the status of the data node in Horizon is down.
- 3 Remove the data node using the dn-remove command or from Horizon (**HyperScale > System > Data Plane**).
- 4 Add the data node back in to the configuration again.
- 5 After the data node is added successfully, synchronize the data with the primary data node that hosts the instances.

After the synchronization is complete, the degraded state of the vdisks is cleared and the reflection is enabled. The secondary data node starts serving as the reflection target.

Refer to the *HyperScale for OpenStack User's Guide* for more information on how to add or remove data nodes and synchronize data node reflection after a failure.

If a network failure occurs on the primary data node, recovery reads and application reads do not complete (3617)

If a network failure occurs on the primary data node, recovery reads and application reads remain stuck. A series of messages is displayed on the compute node, beginning with the following:

```
[i:volume] vol_recover_req_cb failed from ird {ID}
```

```
-BS ret -1
```

Workaround: Configure NIC bonding to achieve redundancy at the network level. Combining network interfaces in a bonded configuration helps prevent a network failure due to a loss of a single physical link. Veritas recommends using the active backup mode (mode=1).

The storage service may crash because of outdated metadata (3634)

The disk header of a data node disk includes metadata for identification. In the following situations, the disk header is not cleaned, the HyperScale storage service (dmLd) crashes:

- If an external LUN is attached to the data node, and the LUN was used in an earlier HyperScale deployment.
- If a physical disk was used in an earlier HyperScale deployment, and it is also used in the current deployment.

Workaround: Do one or both of the following:

- Only add newly-created LUNs or disks to the data node.
- Clean the disk headers. Use the following command. Replace the variable *vdisk_name* with the name of the actual vDisk.

```
# dd if=/dev/zero of=vdisk_name bs=4M count=1
```

Recovering a data node may freeze (3761)

This issue occurs in a two data node environment, in the following circumstances:

- 1 The primary virtual machine has any flavor with reflection enabled, and it has a data disk attached.
- 2 A snapshot of the data disk is attached to the primary virtual machine.

- 3 While data is being recovered, a critical failure occurs, and the data node storage agent fails and cannot come up. This triggers failover to the secondary virtual machine.
- 4 When failover completes and the data is being recovered from the second data node, the data node storage agent on the second data node is killed. The data node restarts as part of the auto-start process.

The recovery freezes, and messages similar to the following are written to the compute logs:

```
2016-05-11 10:37:31 [o:ird] Trying reconnect for connection #0
2016-05-11 10:37:31 [o:ird] Reopening connection
2016-05-11 10:37:31 [o:ird] Connect: failed [111]
2016-05-11 10:37:31 [o:ird] Connection reopen failed 111
2016-05-11 10:37:31 [o:ird] Channel reconnect failed
2016-05-11 10:37:31 [i:volume] vol_recover_req_cb failed from ird
{a893e7a3-b735-4bef-bab3-04c8f403c64f}-BS ret -1
2016-05-11 10:37:32 [o:ird] Trying to reconnect channel
2016-05-11 10:37:32 [o:ird] Reconnecting channel
2016-05-11 10:37:32 [o:ird] Trying reconnect for connection #0
2016-05-11 10:37:32 [o:ird] Reopening connection
2016-05-11 10:37:32 [o:ird] Connect: failed [111]
2016-05-11 10:37:32 [o:ird] Connection reopen failed 111
2016-05-11 10:37:32 [o:ird] Channel reconnect failed
2016-05-11 10:37:32 [i:volume] vol_recover_req_cb failed from ird
{a893e7a3-b735-4bef-bab3-04c8f403c64f}-BS ret -1
2016-05-11 10:37:33 [o:ird] Trying to reconnect channel
2016-05-11 10:37:33 [o:ird] Reconnecting channel
2016-05-11 10:37:33 [o:ird] Trying reconnect for connection #0
```

Workaround: After the data node service agent recovers, attach the same snapshot again.

To bring up a failed data node, you must manually rename the binary file (3763)

When a data node fails over, the `dmld` binary file is renamed `dmld.disabled`. To bring the data node up again, do the following:

- 1 On the data node, navigate to the path `/opt/VRTSofdn/bin/`.
- 2 Rename the file `dmld.disabled` to `dmld`.
- 3 Restart the `hyperscale-dmld` service.

A vDisk may go into the Error state even when one active data node is available (4098)

Due to a series of data node failures, you may create a vDisk that goes into the `Error` state even though there is one active data node.

Workaround: For the current virtual machine, there is no workaround for this issue. Create a new virtual machine and restore the resiliency of the data nodes.

When you remove a data node, logical volume manager (LVM) objects are not removed (4113)

This issue occurs when a data node uses multiple solid-state drives (SSDs) as meta disks, and you remove the data node. The LVM objects are not removed.

Workaround: To resolve this issue, use the following Linux commands to remove the LVM objects. In this example, two SSDs were used as metadisks: `ssd1` and `ssd2`.

1 Remove the logical volume.

```
# lvremove /dev/ofvg/ofmetavol
```

2 Reduce the volume group.

```
# vgreduce ofvg ssd1
```

3 Remove the physical volume that is associated with `ssd1`.

```
# pvremove ssd1
```

4 Remove the volume group.

```
# vgremove ofvg
```

5 Remove the physical volume that is associated with `ssd2`.

```
# pvremove ssd2
```

For more information on these commands, see the Linux man pages.

If you remove a failed data node and re-add it with the same `data_ip` address, some snapshots are corrupted (4148)

In this scenario, a data node fails and is re-added with the same `data_ip` address. Then, if you take a snapshot of a virtual machine with the VRTSBronze flavor which is associated with the re-added data node, the snapshot is corrupted.

Workaround: Re-add the data node with a different `data_ip` address.

When data node synchronization fails, you should stop episodic data synchronization (EDS) (4286)

This issue occurs in a two data node environment under the following circumstances:

1. One data node (`datanode-1`) runs a virtual machine that is configured with the VRTSSilver flavor.
2. The other data node (`datanode-2`) runs a virtual machine that is configured with the VRTSBronze flavor.
3. The administrator creates a snapshot of the VRTSBronze virtual machine's boot volume and attaches the snapshot to the VRTSSilver virtual machine. The snapshot volume is degraded.
4. On `datanode-2`, data is written to the snapshot volume and the administrator creates another snapshot.
5. The administrator enters the command to synchronize `datanode-1`.
6. The operation fails with the following error:

```
failed to create vdisk {GUID}
error 610.
```

Trying to restore reflection for this volume from `datanode-1` to `datanode-2` fails because the parent volume for the VRTSBronze virtual is not available on `datanode-1`. The parent volume is not synchronized because it is a VRTSBronze volume. Therefore, the data node rejects the synchronization request.

Workaround: Stop EDS of the vDisk if its data node dies. Otherwise, the next vDisk snapshots will be corrupted.

If some vDisks are in a degraded state and others are not, snapshots may be corrupted (4306)

This issue can occur in the following circumstances:

- A data node that is used as a reflection target fails.

- A HyperScale user triggers a snapshot of the data disk. The data disk goes into a degraded state. The boot disk is not degraded because its snapshot has not occurred.
- The administrator re-adds the failed data node and triggers a data node resynchronization. The resynchronization only occurs for the data disk.
- The boot disk snapshots that are reflected to the new data node are corrupted.

Workaround: When you re-add the data node, specify a new data node IP address.

You cannot create and attach a vDisk to a virtual machine during data node failover (4338)

If you try to create and attached a vDisk to a virtual machine while the data node is failing over, the vDisk may go into the `ERROR` state.

Workaround: Before you create and attach the vDisk, wait until the failover completes.

EDS operations to the target data node may fail (4385)

During a data node failover, if edits to the target data node are missed, all subsequent episodic data synchronization (EDS) operations fail.

One of the reasons that the target data node may miss edits is if it killed; for example, if the `dmld` service on the data node is stopped.

Workaround: None.

Virtual machine issues

This section describes issues related to virtual machines.

Instances are not reachable using their floating IP addresses after a compute node reboot (6485)

This issue occurs when you reboot a compute node. Instances running on the compute node are no longer accessible using the floating IP address associated with them. Even though an instance is up and running, it is not accessible via the IP address.

Workaround: Restart the `neutron-linuxbridge-agent.service` on the controller node.

Run the following command on the controller node:

```
# systemctl start neutron-linuxbridge-agent.service
```

After this service restarts successfully, try to access the instance using its associated floating IP address.

Launch instance operation may fail due to boot volume and flavor disk size mismatch issue

While launching a new instance from an instance snapshot, if the disk size defined for the flavor that you associate with the new instance is greater than the size of the snapshot (or the boot volume from which the instance snapshot was created), the launch instance process fails.

OpenStack does not allow you to proceed with the instance creation and displays the following error on the Launch Instance window:

```
Volume is smaller than the minimum size specified in image metadata.  
Volume size is <bootvolumesize> bytes, minimum size is  
<flavordisksize> bytes. (HTTP 400)
```

Note: OpenStack does not enforce such a size condition if you launch the instance directly from a boot volume (rather than from a snapshot) and associate a flavor with a disk size greater than the boot volume size.

Workaround: This is a known OpenStack issue and there is no workaround available at this time.

Data is lost when a virtual machine is terminated, and the data disk is attached to another virtual machine (3728072, 1626)

If you terminate a virtual machine and then attach its data disk to another virtual machine, data from the first virtual machine is lost.

Workaround: None.

If the nova-compute status is down (xxx), a virtual machine may appear to be deleted, but it is not (3730758)

If you want to terminate a virtual machine, make sure that the nova-compute service is up on the hypervisor. If you try to delete a virtual machine when the nova-compute service is down, the user interface indicates that the virtual machine is successfully deleted, but it is still up.

Workaround: Before you delete a virtual machine, verify the state of the compute system using the OpenStack command `nova-manage service list`. If the output of the status check commands shows that a compute system is down (xxx), the

compute system is not in communication with the controller. You cannot delete a virtual machine.

When a virtual machine is launched, it is set to the `ERROR` state (2601)

When a virtual machine is launched, it is incorrectly set to the `ERROR` state. This issue may occur because disk space is reported incorrectly on any of the compute nodes. Reasons for this issue can include the following:

- Disk space is reported incorrectly because the `hyperscale-mq-storage` storage service is either dead or stopped on the HyperScale enabled computes nodes. (See Workaround 1.)
- You run the `nova-manage service list`, and you see that the compute nodes are disabled. (See Workaround 2.)

Workaround 1: On the affected compute node, start the `hyperscale-mq-store`. On the controller node, check the status of the compute node. Enter:

```
hyperscale nodes
```

Workaround 2: On the controller node, enable the compute node. Enter:

```
nova-manage service enable --host=compute_name \  
--service=nova-compute
```

Virtual machines with the Bronze flavor may be associated with the same data node (3152)

The OpenStack cinder scheduler determines which virtual machines are associated with which data nodes. The scheduler does this based on the amount of available storage on the data nodes.

When you have two data nodes with the same amount of available storage, an issue occurs. For example, if you create a two virtual machines with the `VRTSBronze` flavor one after the other, they may both be associated with the same data node. This is because there has been no change in storage capacity (that is, no snapshots) between the creation of the first and second virtual machine.

Workaround: None

A restoration of reflection workflow is not triggered (3742)

If a virtual machine is in the failed over state and its reflection is lost, a restoration of reflection (RoR) workflow is not triggered.

Workaround: None.

You cannot set the `qos_maxiops` or `qos_miniops` metadata values to `-1` (infinity) (3961)

You can edit the metadata for HyperScale flavors. However, do not set the quality of service (QoS) metadata values for maximum IOPS (`qos_maxiops`) or minimum IOPS (`qos_miniops`) to `-1` (infinity).

When either of these values is set to `-1`, you cannot launch the instance.

Virtual machines are scheduled on a compute node after a hard disk drive error or failure (4229)

The nova-scheduler provisions virtual machines on compute nodes. Even if hard disk storage on the compute node has an error or fails, virtual machines can still be provisioned on the node. You need to stop virtual machines from being provisioned.

To resolve this issue

- 1 List all the compute nodes.

```
# nova-manage service list
```

- 2 In the list, identify the compute node whose hard disk has failed. The storage dashboard's **Alerts/Events** area displays compute node hard disk drive I/O errors.

- 3 Disable the compute node from nova. Enter:

```
# nova-manage service disable --host=compute_node_name \  
--service=nova-compute
```

- 4 Verify that the compute node is disabled. Enter:

```
# nova-manage service list
```

Snapshots or exports do not include `qemu-guest-agent` hardware data (4289)

This issue occurs when you have a virtual machine which has the `qemu-guest-agent` installed. When you take a snapshot or export the virtual machine, only the software package information is preserved, not the hardware information. When you restore from the snapshot or export and launch a new virtual machine, it does not have the `qemu-guest-agent` hardware channel.

Workaround: After you create the new virtual machine from the snapshot or export, do the following:

- 1 Make sure that Red Hat Virtual Machine Manager (virt-manager) is installed on the node where the instance is running.
- 2 Add a `qemu-guest-agent` channel named `Controller VitrIO Serial`.
- 3 From the virt-manager, force a shutdown.
- 4 Power on the virtual machine.
- 5 Start the `qemu-guest-agent` service.

After a volume attach operation fails, you may be unable to delete a virtual machine (4445)

The following issue may occur:

- You try attach a vDisk to a virtual machine.
- The operation fails and the vDisk remains in the attaching state.
- You cannot delete the virtual machine, because the virtual machine's vDisks are deleted.

Workaround: Delete all the vDisks from the database, then delete the virtual machine.

<https://ask.openstack.org/en/question/66918/how-to-delete-volume-with-available-status-and-attached-to/>

To create a volume or launch an instance, the image must be public (6332)

When you launch a new instance, or when you create a volume from an image, you must specify the image to use. This image must be public.

When you launch a new instance, the **Launch Instance** wizard displays the visibility of each image (public or private). Before you create a volume, navigate to the **Images** page (**Project > Compute > Images**) to check the visibility of each image.

Live migration issues

This section describes issues related to live migration.

Live migration may fail with an out of memory error

If the stack overflows in the `libvirt` virtualization API, some live migration scenarios may fail with an out of memory error.

Workaround: None.

Issues that occur if a live migration of a virtual machine fails (2615)

If you live migrate a virtual machine and the operation fails in the last stages of the process, the virtual machine does not migrate. However, virtual machine metadata specific to HyperScale is updated with details of the new (intended) host. This issue prevents you from trying to live migrate the virtual machine again.

Workaround: None.

If a live migration is in progress on a compute node, that node cannot participate in a failover/failback operation (2664)

This issue occurs in the following circumstances:

1. On one compute node, node A, a live migration is in progress.
2. On another compute node, node B, a storage failure occurs that requires node B failover to node A.
3. The failover cannot be triggered because node A is in the middle of a live migration.

Workaround: None.

You cannot live migrate a virtual machine whose I/O has failed over (3544)

If any storage media on the compute node goes down, an I/O failover is triggered. After the failover workflow completes, you cannot live migrate the virtual machine to another compute node.

Workaround: None.

Performing a live migration during a storage service failure recovery workflow may lead to a crash of the storage service on the reflection compute node (4030)

Do not live migrate a virtual machine while a storage service failure (SSF) recovery workflow is in progress. The live migration causes the storage service on the reflection compute node to crash.

Workaround: Perform the live migration when an SSF workflow is not progress.

When the migration time out value is exceeded, the live migration aborts and metadata is modified (4317)

This issue has two parts, with a workaround for each one:

- If a live migration takes longer than the live migration progress time out value that is specified in the `nova.conf` file, the OpenStack nova compute service aborts the migration.
- When the operation aborts, virtual metadata is updated. The current host IP address is replaced with the reflection host IP address. If you try the live migration again, it fails.

Workarounds: To resolve these issues, do the following:

- To ensure that the live migration completes, edit the `nova.conf` file, and increase the value (in seconds) of `live_migration_progress_timeout`. The default is 150. To disable timeouts, specify 0.
- To resolve the issue with the virtual machine IP address, delete the virtual machine and recreate it using a recovery snapshot.

Snapshot issues

This section describes issues related to snapshots.

Taking snapshots of virtual machine volumes while the virtual machine is being migrated is not supported

You cannot take snapshots of virtual machine volumes while the virtual machine is being migrated.

Workaround: Wait until the migration completes before you take the snapshot.

Snapshot creation may be delayed (3780374)

If you create virtual machines that collectively use more than 20 disks, you may see a delay when you create snapshots for those disks. This delay is due to a timeout issue in the messaging system.

Workaround: None.

If you take a manual snapshot during the failback process, the virtual machine hangs (2664)

If you take a manual snapshot during the storage agent recovery process, the virtual machine I/O stops, and the virtual machine hangs. Messages similar to the following appear in the virtual machine's `libvirt` log.

- 3 Create a virtual machine with the VRSTSilver flavor. Data node 2 is its primary data node.
- 4 Attach the VRTSBronze snapshot to the VRTSSilver virtual machine.
Note that the snapshot volume that was created does not have reflection enabled at the data node level.

During a storage service failure, a manual snapshot goes into the `Error` state (4586)

If you try to manually create a snapshot during a storage service failure, the snapshot goes into the `error` state. A timeout message is written to the controller log.

Workaround: Wait until the storage service is restored before taking a snapshot.

Failover issues

This section describes issues related to the failover process.

I/O failover may not complete, and the virtual machine retries the operation endlessly (2500)

In rare situations, an I/O failover operation may not complete. The virtual machine retries the operation endlessly. The virtual machine is inaccessible, although the user interface may indicate that it is active.

Workaround: None.

If the storage service fails while you create an instance, the I/O failover/failback workflow remains pending (2617)

If the storage service fails while you create an instance, the I/O failover/failback workflow remains in the `Pending` state. Also, the instance fails and goes into the `Error` state.

Workaround: None.

After a successful failover, the primary data node IP address and other information are not changed in the user interface (3284)

When the primary data node fails over, the secondary data node becomes the primary. However, the primary node IP address displayed on the user interface is not changed.

Also, the cinder metadata does not change. The `Primary_datanode_ip` and `Secondary_datanode_ip` values remain the same. This is by design.

Workaround: To determine the actual primary data node, display the cinder metadata and note the value for `current_data_node_owner`. That value indicates which data node owns the vDisk. That data node is the primary.

Virtual machine storage does not failover if there is no I/O activity (3622)

If the storage media on a compute node fails, HyperScale provides resiliency. If any I/O for a virtual machine encounters an error, the I/O is served from its reflection compute node.

However, there is a limitation. If a virtual machine is idle and there is no I/O activity, the storage does not failover. For failures of episodic or user snapshots and backend store flush errors, an alert is raised in the **Alerts/Events** area on the HyperScale storage dashboard. When I/O on that virtual machine occurs, the storage is failed over to the reflection compute node.

Workaround: If you see this error for a virtual machine, you can live migrate the virtual machine.

You cannot perform a hard restart of the virtual machine after its data has failed over to the reflection compute node (3768)

This issue occurs when the storage media on the source compute node fails, and the data fails over to the reflection compute node.

If you need to restart the virtual machine on the source compute node, you must perform a soft restart. In a soft restart, the virtual machine is brought down gracefully and no data is lost. A hard restart, which can damage data, does not bring up the virtual machine.

I/O failover workflow fails for a guest virtual machine with the message `_vdisk_fo_newsource_edit exception: timed out` (3965)

This issue occurs in a four compute node environment, in which a guest virtual machine has a reflection factor of three. The events are as follows:

- Two of the reflection targets fail.
- The `dmld` service on the primary data node fails, triggering a failover.

- When the failover completes, a large file is copied from the outside to the guest virtual machine's data disk using the OpenStack `scp` command.
- While the file is copied, the storage service on the host compute node fails.
- The failover workflow fails.

Workaround: None.

Storage added to a virtual machine during an I/O failover/failback operation is not attached (4554)

If you add storage to a virtual machine during an I/O failover/failback operation, the storage is added, but it is not attached to the virtual machine.

Workaround: Wait until the I/O failover/failback operation completes before you add storage.

Storage issues

This section describes storage issues.

If a volume goes into the `Error_Deleting` state, you cannot delete it or perform any other operations on it (3719191)

If you try to delete a volume, and it goes into the `Error_Deleting` state, a `Permission Denied` message is displayed. You cannot delete the volume or perform any other operation on it.

The `Error_Deleting` state can occur for several reasons including a crash of the data node virtual machine's storage layer.

Workaround: None.

HyperScale may not work with some volume types or backend types (1504)

This issue occurs in the following circumstances:

- In an OpenStack environment, the administrator creates a virtual machine with one of the standard OpenStack flavors.
- A volume is created with its own backend type, and it is attached to the instance.
- HyperScale is installed, and the administrator creates a volume with its own backend type.
- When the administrator tries to attach the volume to the instance, the operation fails.

- Error information is written to the `cinder` and `nova` logs.

Workaround: None.

You cannot add storage to a HyperScale virtual machine through the cinder CLI or API (3073)

You cannot add storage to a HyperScale virtual machine through the `cinder` CLI or API. Storage that you add through the `cinder` CLI or API is not associated with the second data node.

Workaround: To resolve this issue, do one of the following:

- On the HyperScale dashboard, navigate to **Project > Compute > Instances**, and in the **Actions** column for the instance, select **Add Storage**.
- On the HyperScale command line, enter:

```
# hyperscale storage-create-and-attach
```

- Use the following HyperScale API call:

```
hyperscale storage_create_and_attach
```

If you add storage to a virtual machine during a failover, the virtual machine may hang (4558)

When the storage media fails on the compute node, the virtual machines on the compute node fail over. If you add storage to one of these virtual machines during the failover, the virtual machine may hang.

Workaround: None.

External storage provider and backup issues

This section describes issues with the HyperScale external storage provider and other backup vendors.

The backup operation fails after one hour because a token times out (3526)

In this issue, the backup operation fails after one hour, and a message similar to the following is displayed on the backup provider server:

```
backup_c_client: disk_read_api.c:194: resp_hdlr_ex: Assertion  
'raw_bytes_recd + json_bytes_recd == bytes_recd' failed.
```

Workaround:**To resolve this issue**

- 1 On the backup provider server, update the `keystone.conf` file and set the token expiration parameter to a larger value. The default expiration value is 3600 seconds (one hour). For example, increase the default value to 14,400 seconds (four hours).

```
expiration = 14400
```

- 2 Restart the `http` service.

```
# service httpd restart
```

The `auth_token_timeout` value in the external storage provider configuration file must match the keystone token time out value (4020)

In the HyperScale external storage provider configuration file (`opt/VRTSofesp/etc/esp-config`) the value of the `auth_token_timeout` parameter should match the keystone token time out value.

This parameter has been set to match the default keystone token time out of 1 hour (3600 seconds). If you change the keystone token time out value, you must also change the `auth_token_timeout` changed to match.

After you change the time out value, you must restart the external storage provider service for the changes to take effect. On the external storage provider node, enter the following:

```
service hyperscale-esp restart
```

Do not start an export or backup while a live migration is in progress (6445)

If a live migration is in progress, do not start an export or backup. Wait until the live migration completes.

Ceilometer issues

This section describes issues related to the ceilometer.

Ceilometer statistics may include duplicate entries (3839317)

This issue is relevant to those who work with the ceilometer database directly. The ceilometer statistics are updated every 30 seconds and should include no missing entries or duplicate entries. Occasionally, duplicate entries may occur.

Workaround: None.