Veritas™ HyperScale 1.0 for OpenStack Installation Guide

Ubuntu

June 2017

VERITAS™

Veritas HyperScale for OpenStack Installation Guide

Last updated: 2017-06-22

Document version: 1.0 Rev 1

Legal Notice

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/

Contents

Chapter 1

Preparing for installation

This chapter includes the following topics:

HyperScale product environments

This chapter lists the HyperScale installation requirements for the minimum and preferred production environments.

Table 1-1 HyperScale production environments

Environment	Configuration
Minimum production environment	■ Two data nodes; the OpenStack controller can be used as one of the data nodes ■ Three compute nodes

Table 1-1 HyperScale production environments *(continued)*

Environment	Configuration
Preferred production environment	■ Three Cinder nodes, two of which will be used as data nodes ■ Ten compute nodes ■ The OpenStack controller is hosted separately from either data node

You can also deploy HyperScale in OpenStack director-installed environments.

Data node requirements

In OpenStack terminology, a HyperScale data node is a Cinder node.

If you install HyperScale in the minimum production environment, the first data node is installed on the same physical node as the OpenStack controller. The installation turns the controller into a HyperScale controller. This controller displays the HyperScale dashboard and enables you to manage and monitor HyperScale compute nodes and data nodes. You can add the second data node to any Cinder node in your environment that meets the installation requirements.

In the preferred production environment, the OpenStack controller is hosted separately from the first date node. Each data node can be installed on any Cinder node. In a preferred production environment, the third Cinder node is used in case of an irrecoverable data node failure. After the failure, it is added as a new data node in HyperScale deployment, replacing the data node that failed.

Regardless of the production environment you use, before you install HyperScale, the OpenStack deployment must meet the following requirements:

- It must not have any projects or users, except for the projects and the users it creates by default.
- No OpenStack instances should be running.
- The network manager service (`NetworkManager`) must be disabled on the controller.

Table 1-2 Data node hardware requirements

Hardware	Minimum production environment	Preferred production environment
Number of nodes	2	2
CPU	Dual socket (16 core)	Dual socket (16 core)

Table 1-2 Data node hardware requirements *(continued)*

Hardware	Minimum production environment	Preferred production environment
CPU type	Intel	Intel
Memory	64 GB	128 GB
Network interface cards	3 10 Gbps (1 public, 2 private)	2 10 Gbps (1 public, 1 private) 1 40 Gbps (private)
Storage		
Solid-state drive (SSD)	0.8 TB HyperScale supports the following SSDs: ■ Fusion-io Supported kernel version is 4.4.0.21. ■ Intel SATA ■ Intel PCIe For Intel SSDs, the supported kernel version is 3.10.0-327.el7.x86_64. The SSD should not have any partitions and should not be mounted.	1 or 2 TB Same supported drives as the minimum environment.
Hard disk drive (HDD)	1 boot disk along with multiple disks amounting to total storage in the range of 20 TB - 40TB The HDDs should not have any partitions. Disks must be labeled `msdos`; otherwise, HyperScale does not acknowledge the devices.	1 boot disk along with multiple disks amounting to total storage of 64 TB or more Same disk labeling requirements as the minimum environment.

Storage notes:

- The storage in the data nodes must be equal or greater than the total storage across all compute nodes.
- RAID6 is recommended in case of disk failures.

Table 1-3 Data node software requirements

Software	Requirements (same for minimum and preferred environments)
Operating System	Ubuntu 16.04 LTS (Server Edition) Kernel versions: ■ 4.4.0-21 ■ 4.4.0-28 ■ 4.4.0-31 ■ 4.4.0-57 ■ 4.4.0-62 ■ 4.4.0-64 ■ 4.4.0-66 ■ 4.4.0-70 ■ 4.4.0-72 **Note:** Veritas recommends that you disable kernel auto-upgrades on all systems in the HyperScale configuration. An unsupported kernel version may cause a failure and lead to configuration issues. Refer to the operating system documentation for information on how to disable auto-upgrades.
OpenStack distribution	Canonical Ubuntu OpenStack Ocata (15.0)
OpenStack services	The following Cinder services must be enabled: ■ `openstack-cinder-api` ■ `openstack-cinder-scheduler` ■ `openstack-cinder-volume` To verify that the services are enabled, log into the controller and enter the following: `# openstack-status` The status of each service should be `active`.
Network interface	■ The network interface card (NIC), which is configured for the data network, must be enabled. If it is not, enable it with the `ifup` command. ■ You must assign IP addresses to all data interfaces for all nodes. ■ HyperScale supports NIC bonding. Veritas recommends using the active backup mode (mode=1).

Table 1-3 Data node software requirements *(continued)*

Software	Requirements (same for minimum and preferred environments)
Packages	The following packages are required on the data node: ■ `coreutils` ■ `ipcalc` ■ `java` ■ `openjdk-8-jre-headless` Ensure that `/usr/bin/java` is created after installing java packages. ■ `libcurl4-openssl-dev` ■ `lvm2` ■ `python-amqp` ■ `python-anyjson` ■ `python-importlib` ■ `python-kazoo` ■ `python-kombu` ■ `python-sqlalchemy` If the controller is also deployed on the data node, the following are required in addition to the packages mentioned earlier: ■ `bc` ■ `crudini` ■ `dpkg-dev` ■ `python-mysqldb` ■ `rabbitmq-server` ■ `software-properties-common`
Security	To configure the data node in a director-installed environment, it must support Secure Shell (SSH) and have a root password.

Compute node requirements

Table 1-4 Compute node hardware requirements

Hardware	Minimum production environment	Preferred production environment
Number of nodes	3	10
CPU	Dual socket (16 core)	Dual socket (16 core)

Table 1-4 Compute node hardware requirements *(continued)*

Hardware	Minimum production environment	Preferred production environment
CPU type	Intel	Intel
Memory	64 GB	128 GB
Network interface card	3 10 Gbps base-T NIC (1 public, 2 private)	4 10 Gbps base-T NIC (2 public, 2 private). This configuration does not assume port aggregation.
Solid-state drive (SSD)	800 GB HyperScale supports the following SSDs: ■ Fusion-io Supported kernel version is 4.4.0.21. ■ Intel SATA ■ Intel PCIe For Intel SSDs, the supported kernel version is 3.10.0-327.el7.x86_64. The SSD should not have any partitions and should not be mounted.	1 TB Same supported drives as the minimum environment.
Hard disk drive (HDD)	1 boot disk along with multiple disks amounting to total storage in the range of 20 TB - 40TB (No SAN) Disks must be labeled `msdos`; otherwise, HyperScale does not acknowledge the devices.	1 boot disk along with multiple disks amounting to total storage of 64 TB or more Same disk labeling requirements as the minimum environment.

Table 1-5 Compute node software requirements

Software	Requirements (same for both minimum and preferred environments)
Operating system	Ubuntu 16.04 LTS (Server Edition) "Virtual Machine host" software option must be installed. Kernel versions: ■ 4.4.0-21 ■ 4.4.0-28 ■ 4.4.0-31 ■ 4.4.0-57 ■ 4.4.0-62 ■ 4.4.0-64 ■ 4.4.0-66 ■ 4.4.0-70 ■ 4.4.0-72 **Note:** Veritas recommends that you disable kernel auto-upgrades on all systems in the HyperScale configuration. An unsupported kernel version may cause a failure and lead to configuration issues. Refer to the operating system documentation for information on how to disable auto-upgrades.
Network interface	■ The network interface must be enabled. If it is not, enable it with the `ifup` command. ■ You must assign IP addresses to all data interfaces for all nodes. ■ HyperScale supports NIC bonding. Veritas recommends using the active backup mode (mode=1).
Packages	The following packages are required on the compute node: ■ `bc` ■ `cgroup-bin` ■ `dnsmasq-base` ■ `expect` ■ `lvm2` ■ `python 2.6 or later` ■ `python-ceilometerclient` ■ `xfsprogs`

Table 1-5 Compute node software requirements *(continued)*

Software	Requirements (same for both minimum and preferred environments)
Other software	XFS journaling file system installed. The `expect` package must be installed.
Security	To configure a compute node in a director-installed environment, it must support Secure Shell (SSH) and a root password.

Network switch requirements

An appropriate 10 GB base-T mid-rack switch to provide interconnectivity between the HyperScale data interfaces is preferred. However you can also use a top-of-rack switch with the appropriate VLAN for the HyperScale data network.

Table 1-6 Mid-rack network switch requirements

Category	Minimum production environment	Preferred production environment
Number of ports	10	26 (2 40 GB, 24 10GB) This configuration does not assume port aggregation.
Speed	10 Gbps	10 Gbps, 1 40 Gbps

Network port requirements

If you have configured a firewall, then ensure that the firewall settings allow access to the services and ports used by HyperScale. Enable both inbound and outbound communication for these ports and services.

Note: These ports are not user-configurable. Ensure that these are not already in use.

Table 1-7 HyperScale ports and services

Port	Description
42181 / TCP	Used by the ZooKeeper server (`hyperscale-zookeeper`)

Table 1-7 HyperScale ports and services *(continued)*

Port	Description
9898 / TCP	Used by the `hyperscale-mq-storage` service for: ■ I/O redirection from the storage driver to the HyperScale daemon on the compute node. This is required for live migration and to ensure resiliency in case of storage failures.
9999 / TCP	Used by `hyperscale-mq-storage` service for: ■ I/O reflection between compute nodes. ■ Snapshot data transfer between the HyperScale daemon on the compute node and the HyperScale dmld service on the data node.

External storage provider requirements

HyperScale's external storage provider lets infrastructure administrators create and export HyperScale instance snapshots to an external storage device. The exported snapshot can then be used to launch a new instance if the original becomes corrupted, is accidentally deleted, or another data loss or disaster recovery situation occurs. Making sure that data is securely and efficiently transferred between the HyperScale nodes (compute and data nodes) and external storage provider requires some advance planning.

You should connect the external storage provider to the same management and data subnets as the rest of the HyperScale deployment.

If the external storage provider is located in the same rack as the other HyperScale nodes, connect its two interfaces to the appropriate ports of the middle of the rack switch. Each of the subnets should have its own virtual LAN (VLAN). The same holds true even if the external storage provider is located in a different rack. However, the middle of the rack switches must be properly connected and trunked to enable intra-VLAN traffic between the two racks.

Before you integrate an external storage provider with HyperScale, the provider must be installed on a separate node that is not part of the HyperScale environment. The node must be a server-grade system with high I/O throughput. In addition, it must meet the following requirements:

Table 1-8 Hardware requirements

Hardware	Requirement
RAM	32 GB minimum, 64 GB or higher recommended
Virtual CPUs	4 cores minimum, 8 cores recommended
Network speed	1 Gbps minimum, 10 Gbps recommended
Hardware type	Physical machine
Network interface cards (NICs)	3 required ■ A 1 Gbps management NIC. Configure this NIC to be on the same management subnet as the compute node and data node. ■ A 1 Gbps (minimum) data NIC. 10 Gbps is recommended. Configure the other NIC to be on the same data subnet as the compute node and data node. ■ A public network NIC to access the node and check logs and perform troubleshooting.
Physical storage	Enough physical storage to maintain instance snapshots. HyperScale maintains a full copy of an instance's vDisks and metadata when it is exported to the external storage provider node. For example, if you export five instances that have a total of 5 TB of data written to their vDisks, the external storage provider node requires 5 TB of storage.
Network	The firewall on the external storage provider node is disabled, or port 8334 is open. (This must be verified before you try to register the external storage provider from the HyperScale user interface.)

Note: When the external storage provider imports or exports a snapshot, it communicates with the primary and secondary data nodes using the management and data subnets. (This is the same method that compute nodes and data nodes use to communicate with each other.) Therefore, the external storage provider node should either be in the same physical rack as the HyperScale nodes, or, if it is outside the rack, it should be able to communicate to other HyperScale nodes using the management and data subnets.

Figure 1-1 Controller node and data node communication with the external storage provider node

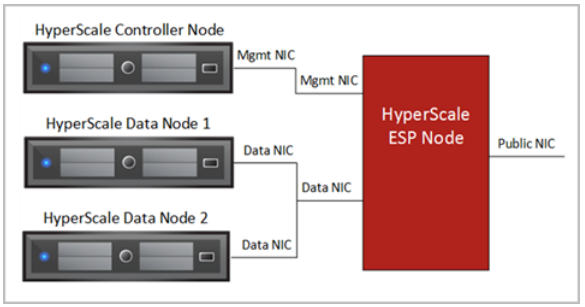

Table 1-9 Software requirements

Software	Requirement
Operating system	Red Hat Enterprise Linux (RHEL) 7.2 Infrastructure Edition (minimum)
Web framework	Flask 0.10.1 or later
Python	2.7 or later

Setting up NetBackup

The NetBackup administrator must set up a dedicated Media Server in the production environment. The Media Server is used as a Backup Proxy.

Table 1-10 NetBackup Media Server requirements

Hardware	Requirement
RAM	32 GB minimum, 64 GB or higher recommended
Virtual CPUs	4 cores minimum, 8 cores recommended
Network speed	1 Gbps minimum, 10 Gbps recommended
Hardware type	Physical machine

Table 1-10 NetBackup Media Server requirements *(continued)*

Hardware	Requirement
Network interface cards (NICs)	3 required ■ A 1 Gbps management NIC. Configure this NIC to be on the same management subnet as the controller node. ■ A 1 Gbps (minimum) data NIC. 10 Gbps is recommended. Configure this NIC to be on the same data subnet as the data nodes. ■ A public network NIC to communicate with the NetBackup Master Server.
Physical storage	Enough physical storage to maintain instance snapshots. HyperScale maintains a full copy of an instance's vDisks and metadata when it is exported to the external storage provider node. For example, if you export five instances that have a total of 5 TB of data written to their vDisks, the external storage provider node requires 5 TB of storage.
Network	The firewall on the external storage provider node is disabled, or port 8334 is open. (This must be verified before you try to register the external storage provider from the HyperScale user interface.)

When the NetBackup administrator completes the configuration, HyperScale, the NetBackup Media Server, and NetBackup Master Server should have the following communication path.

Figure 1-2 Communications between the HyperScale controller and data nodes and NetBackup Media and Master servers

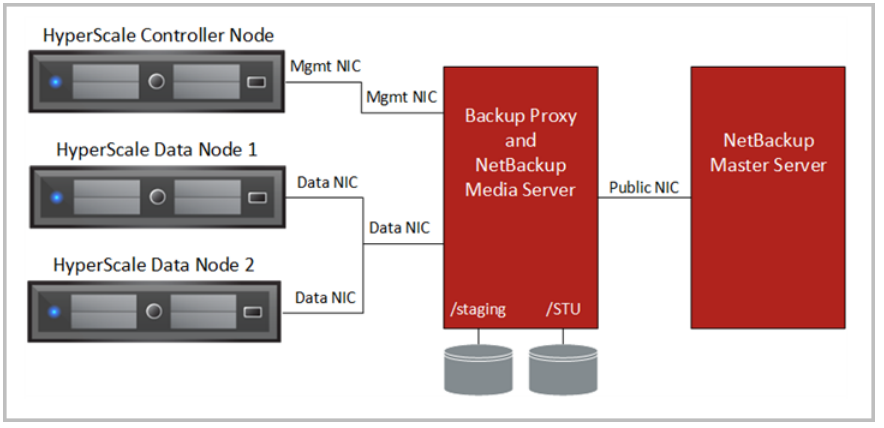

In addition, the NetBackup administrator must configure the directory path on which HyperScale backups are stored; this is known as the staging location. This location needs to be on a separate set of disks than that of the storage unit (STU) disks (in case of a hard-disk backed STU).

NetBackup integration requirements

Note: This HyperScale release is only qualified with NetBackup Version 8.0.

To integrate NetBackup with HyperScale, the NetBackup administrator must verify the following:

- The NetBackup Media Server and Master Server must be installed. When you install the servers, use the fully qualified domain names. Also use the fully qualified domain name when you register the Media Server with the Master Server.
- The Media Server must be running Red Hat Enterprise Linux (RHEL) 7.2. The "Minimal Install" is not supported; however, the RHEL software "Infrastructure Server" and above are supported.
- The firewall on the Backup Proxy node is disabled, or port 8443 is open. (This must be verified before you try to register NetBackup from the HyperScale user interface.)
- The date, time, and zone on the Backup Proxy node must be in sync with the date, time, and zone on the Master Server. If the dates and times are not sync, backups still occur, but they are not displayed on the HyperScale user interface.

After this information is verified, the NetBackup administrator must communicate the following information to the HyperScale infrastructure administrator:

- The Media Server fully qualified domain name.
- The Media Server password.
- The OpenStack password.
- The directory path on which HyperScale backups are staged temporarily; this is known as the staging location. This location needs to be on a separate set of disks than that of the storage unit (STU) disks (in case of a hard-disk backed STU).

Configuring OpenStack to support HyperScale live migration

To perform live migrations in HyperScale, you must make the following changes to the OpenStack configuration.

Table 1-11 OpenStack configuration changes to support HyperScale live migration

For this item ...	Make this change ...
`nova.conf`	`rpc_response_timeout = 180` (The default is 60.)
`cinder.conf`	`rpc_response_timeout = 180` (The default is 60.)
MySQL	`Max Connections = 2048`

Table 1-11 OpenStack configuration changes to support HyperScale live migration *(continued)*

For this item ...	Make this change ...

RabbitMQ configuration file	Add the following lines:
	`[binary,`
	`{packet, raw}`
	
	`,`
	`{reuseaddr, true}`
	
	`,`
	`{backlog, 128}`
	
	`,`
	`{nodelay, true}`
	
	`,{linger,{true,0}},`
	`{exit_on_close, false}`
	
	`,`
	`{sndbuf,32768}`
	
	`,`
	`{recbuf,32768}`
	
	`,`
	`{keepalive, true}`
	
	`]}`

Chapter 2

Installing HyperScale in an OpenStack private cloud environment

This chapter includes the following topics:

Installing HyperScale on a OpenStack controller

The first installation task is to install the HyperScale software on a OpenStack controller. The steps in this section enable HyperScale on the controller.

Before you perform these steps, make sure of the following:

- The controller meets the HyperScale installation requirements.
- On all systems where you wish to install HyperScale, verify that the system clocks are synchronized and they belong to the same time zone.
- Have the following information ready:
  - The OpenStack admin password
  - The MySQL root password
  - The OpenStack controller's management IP address

- Enable the RabbitMQ management plug-in and restart the RabbitMQ server. On the controller, enter the following:

```
# /usr/sbin/rabbitmq-plugins enable rabbitmq_management
# systemctl restart rabbitmq-server
```

To install HyperScale on a OpenStack controller

1. Download the HyperScale installer bin package on the controller node.
2. Make sure that the bin file has execute permissions on the node.

   ```
   # chmod 777 <hyperscale bin file name>
   ```

3. Run the HyperScale binary file.

   ```
   # ./HyperScale.bin
   ```

   The bin package file name can vary depending on the product release version.

4. At the prompts, enter the following credentials:

   ```
   OpenStack Controller's Management IP:
   OpenStack Administrator Password:
   MySQL Root Password:
   ```

5. Respond to the installer prompts to proceed:
   - Review the software license agreement and type **y**.
   - To participate in the Product Improvement Program, type **y**.
     If you choose to participate in the program, system and product usage information details are collected anonymously.

   The installer starts the controller deployment process. Messages similar to the following indicate the installation status:

   ```
   Verifying archive integrity of Veritas HyperScale Installer...Successful
   Uncompressing Veritas HyperScale Installer
   Installing
   Starting HyperScale Controller deployment...
   Verify readiness for deployment                [ DONE ]
   Save configuration                             [ DONE ]
   Install packages                               [ DONE ]
   Set up database                                [ DONE ]
   Register Veritas HyperScale with OpenStack     [ DONE ]
   Start services                                 [ DONE ]
   Collect and upload telemetry data              [ DONE ]
   ```

```
Veritas HyperScale Controller deployed successfully.
```

Note: If there are any failures during this step, note the errors and then refer to the installer logs for more details. The HyperScale installation log file is located at `/var/opt/VRTSofcore/logs/ofengine.log`.

6 At the end of the installation status messages, the installer displays a URL similar to the following:

```
http://<FQDN>/horizon/storage/dataplane
```

Here, `<FQDN>` is the fully qualified domain name of the node where you just installed the controller.

This URL points to the HyperScale dashboard. Make a note of the URL. This is needed in the next step.

This completes the HyperScale installation on the controller. The next task is to add the first data node to the configuration.

Note: HyperScale does not support the Internet Explorer or Safari browsers. If you like, you can also participate in telemetry data collection.

Adding a data node

Adding a data node is the first task you perform after installing HyperScale on the OpenStack controller. You can add a data node from the HyperScale graphical user interface (GUI) or from the command line.

Before you proceed

Check for the following:

- Ensure that you have installed HyperScale on the OpenStack controller
- Keep the following information ready. You will need to specify these details while adding a data node.
  - The management IP address of the Cinder node on which you are adding the data node.
    In the minimum production environment, this Cinder node is the controller system.
  - The Cinder node password.

- The OpenStack admin password.

- Make sure that the disks you intend to use for data node storage have the `msdos` label type.

About adding data nodes

HyperScale supports up to two data nodes. In the minimum production environment, the first data node you add must be on the same node as the HyperScale controller. In the preferred production environment, both data nodes are hosted on different physical nodes.

Data nodes are used to store point in time snapshots of virtual machine volumes. Every 15 minutes, the compute nodes write new data to the data plane. This operation is known as Episodic Data Synchronization.

In addition, if there are two data nodes in the environment, for the VRTSSilver flavor and above, one data node is the primary data node and the other is the secondary data node. In this configuration, data is reflected from the primary data node to the secondary data node. If one of the data nodes fail, all requests are served from the other data node.

For virtual machines with the VRTSBronze flavor, there is no reflection at the data node level. If the data node that is associated with a Bronze virtual machine fails, you cannot access all the volume versions.

Adding a data node from the HyperScale GUI

To add a data node from the graphic user interface

1. Launch a web browser and type the URL that you noted during the HyperScale installation.

   The URL format is as follows:

   `http://<FQDN>/horizon/storage`

   Here, `<FQDN>` is the fully qualified domain name of the controller node.

   Note: You can also use the public IP address of the controller node in the URL:

   `http://<publicIPaddressofcontroller>/horizon/storage`

   ---

2. Specify the OpenStack admin user name and password on the login page and click **Connect**.

3 Navigate to **HyperScale > System > Data Nodes**.

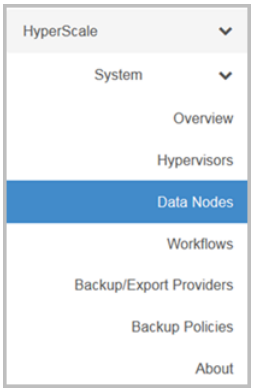

4 On the **Data Nodes** page, click **Add Data Node**.

5 Perform the following actions on the Add Data Node panel:

- Specify the requested details as described in the following table:

Field name	Description
Management IP address	Specify the IP address assigned to the management network interface on the node that you want to add as a data node. Typically, the first data node is the same as the controller node. **Note:** Do not specify the public IP address of the node.
Hypervisor Password	Specify the administrator password of the node that you want to add as a data node.
OpenStack Password	Specify the OpenStack administrator password.

- Click **Discover**.

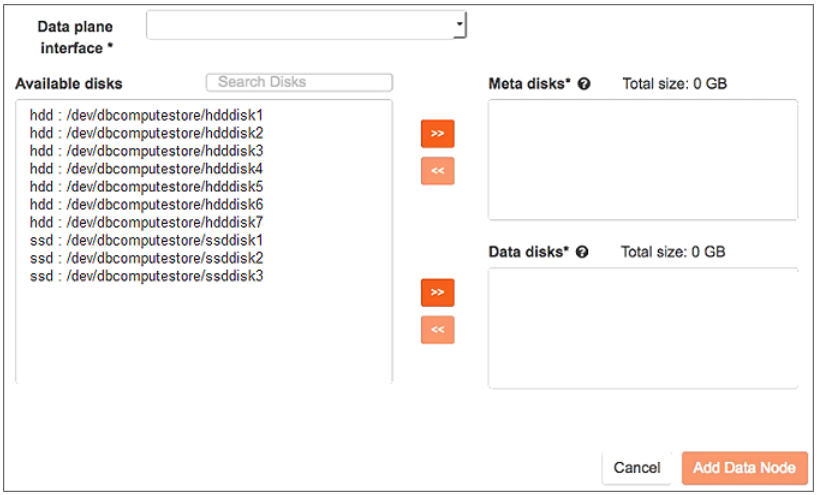

HyperScale discovers the network interfaces and storage on the node.

- From the **Data plane interface** drop-down list, select the network interface associated with the HyperScale data network.
  The drop-down list displays a network interface only if it is enabled and an IP address is assigned to it.

  Note: Do not select any other interface.

- Assign storage for the meta disk and data disks.
  Select disks from the **Available disks** area and use the right arrow icons to move them to the **Meta disk** area and **Data disks** area.
  At a minimum you need two disks—one for the meta disk and one for the data disk. For better performance, choose a solid-state drive (SSD) for the meta disks. If you specify multiple SSDs for the meta disks, HyperScale automatically creates mirrored volumes on them. Mirrored metadata ensures that even if one SSD fails, the data node continues to work and does not have to switch to an alternate data node.

- Click **Add Data Node**.
  HyperScale begins to configure the data node. The dashboard displays various messages that indicate the progress.
  After the process is complete, the name of the data node that you just added appears in the GUI under **HyperScale > System > Data Nodes**.
  If there are any failures during the add node operation, refer to the log file to troubleshoot the issues. To access the log file for a specific data node,

on the Data Nodes tab, under the Logs column, click the **Download** link for that node.

Adding a data node using HyperScale commands

Perform the following steps on the controller node to add a data node using HyperScale commands.

To add a data node from the command line

1. On the controller node, source the OpenStack RC file. This operation sets the required environment variables.

   ```
   # source keystonerc_admin
   ```

2. Add the data node using the `dn-add` command. It has the following syntax:

   ```
   # hyperscale dn-add [--data_disks <data_disks>]
   [--meta_disk <meta_disk>] datanode-mgmt-ip
   <datanode-pwd> <openstack-pwd>
   <dndata-itface> <data_cidr_range>
   ```

   Here, <*data-mgmt-ip*> is the IP address associated with the management network interface and <*dndata-itface*> is the name of the interface associated with the data network.

   Use commas to specify multiple disks.

   Example:

   ```
   # hyperscale dn-add --data_disks /dev/mapper/mpathb,/dev/mapper/mpathc
   --meta_disk /dev/fioa
   172.202.202.1 root123 admin123
   eth5 172.202.202.0/24
   ```

3 Monitor the add operation as it progresses.

Check the following log file on the controller node:

`/var/opt/VRTSofcore/log/<datanode_mgmtip>_config.log`

4 Verify that the data node was created. The nodes table has an entry for the new data node and the data node's status appears as `up`.

The following abbreviated output shows the hostid of the new data node and its status:

```
# hyperscale nodes
+--------------------------------------+ ...  +--------+ ...
| hostid                               | ...  | status | ...
+--------------------------------------+ ...  +--------+ ...
{0001a036-9f33-3ea0-0000-000000000000}   | ...  | up     | ...
```

If an error occurs when you add a data node, information is written to the `/var/opt/VRTSofcore/log/<datanode_mgmtip>_config.log` file on the controller node.

Adding a second data node for resiliency

Add a second data node to your environment to ensure that the data written to the primary data node is resilient. This ensures that if an error occurs on the first data node, data is available on the second data node, and there is no data loss.

See "Adding a data node" on page 22.

Enabling HyperScale on a compute node

You can enable HyperScale on a compute node from the graphical user interface (GUI) or from the command line.

Before you proceed

The following information is required to enable HyperScale on a compute node:

- The hypervisor password
- The OpenStack password

About enabling HyperScale on compute nodes

The HyperScale compute node hosts the workloads. Although you can run all services on a single compute node, you can add additional compute nodes to scale out and expand your HyperScale cloud environment.

To make sure that the data on your compute nodes is resilient, you must enable HyperScale on enough compute nodes to account for the reflection factor (the number of nodes to which the data is written) plus one. You need the additional compute node in case one of the reflection targets goes down. The reflection factor is based on the HyperScale flavor you assign to the virtual machines.

The following table shows the minimum number of compute nodes you need for each virtual machine flavor.

Table 2-1 Minimum number of compute nodes to provide resiliency

Flavor	Reflection factor	Compute nodes (including the source compute node)
VRTSSilver	1	3
VRTSGold	2	4

If an error occurs when you enable HyperScale on a compute node, information is written to the following log file on the controller node:

`/var/opt/VRTSofcore/logs/<computenode_mgmtip>_config.log`

Enabling HyperScale on a compute node using the GUI

To enable HyperScale from the GUI

1. Launch a web browser and type the HyperScale dashboard URL.

   The URL format is as follows:

   `http://<FQDN>/horizon/storage`

   Here, `<FQDN>` is the fully qualified domain name of the controller node.

   You can also use the public IP address of the controller node in the URL:
   `http://<publicIPaddressofcontroller>/horizon/storage`

2. Specify the OpenStack admin user name and password on the login page and click **Connect**.

3. Navigate to **HyperScale > System > Hypervisors**.

4 In the **All Hypervisors** table, locate the hypervisor host on which you want to enable HyperScale and on the right side of the table, under the **Actions** column, click **Enable HyperScale**.

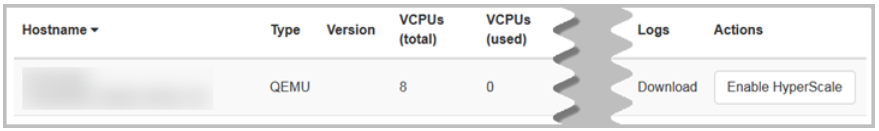

5 Perform the following actions on the Enable HyperScale panel:

- Specify the requested details as described in the following table:

Field name	Description
Hypervisor Password	Specify the administrator password of the node on which you want to enable HyperScale.
OpenStack Password	Specify the OpenStack administrator password.

- Click **Discover**.

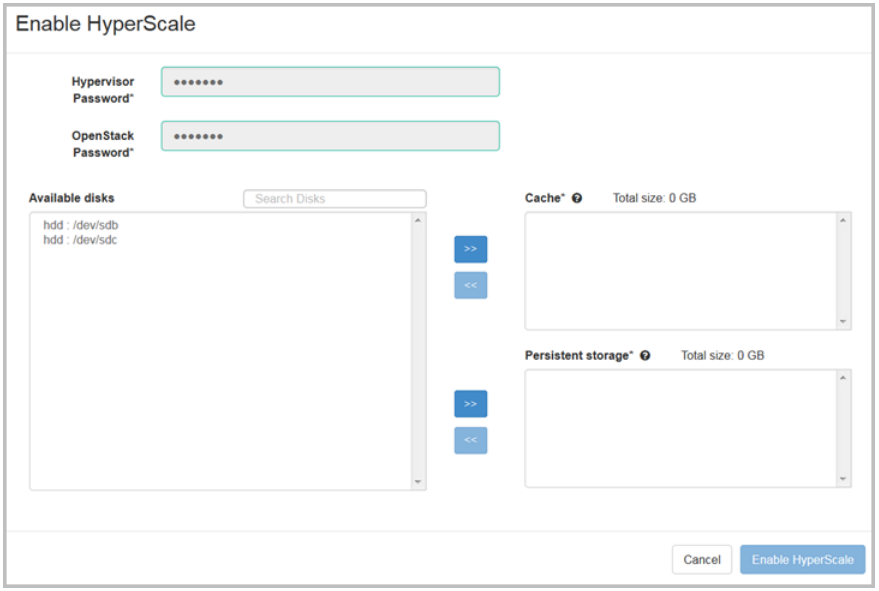

HyperScale discovers the storage for the specified hypervisor.

- Assign disks for the Cache and Persistent storage.
  Select disks from the **Available disks** area and use the right arrow icons to move them to the **Cache** area and **Persistent storage** area.

At a minimum you need two disks—one for the cache and one for persistent storage. For better performance, choose a solid-state drive (SSD) for the cache.

- Click **Enable HyperScale**.
  HyperScale begins to enable HyperScale on the hypervisor. The dashboard displays various messages that indicate the progress.
  After the process is complete, the status under the HyperScale column displays as **Enabled/Running**.
  If there are any failures during the operation, refer to the log file to troubleshoot the issues. To access the log file for a specific hypervisor, on the Hypervisors tab, under the Logs column, click the **Download** link for that hypervisor.

Enabling HyperScale from the command line

Perform the following steps on the controller node to enable HyperScale on a compute node using commands.

To enable HyperScale from the command line

1. On the controller node, source the OpenStack RC file to set the required environment variables.

   ```
   # source keystonerc_admin
   ```

2. Enable HyperScale on the hypervisor using the `compute-enable` command.

   ```
   # hyperscale compute-enable [--data_intf <data-intf>]
   <compute-ip> <compute-pwd> <openstack-pwd>
   <compute-meta-disks>
   <compute-data-disks>
   ```

   Here, *<data-intf>* is the name of the interface associated with the data network and *<compute-ip>* is the IP address associated with the management network interface.

   Use commas to specify multiple disks.

   Example:

   ```
   # hyperscale compute-enable --data_intf ens224
   172.101.101.7 root123 admin@123
   /dev/dbcomputestore/ssddisk3,/dev/dbcomputestore/ssddisk4
   /dev/dbcomputestore/hdddisk3,/dev/dbcomputestore/hdddisk4
   ```

3 Monitor the operation as it progresses.

Check the following log file on the controller:

`/var/opt/VRTSofcore/log/<computenode_mgmtIPaddress>_config.log`

4 Verify that HyperScale was enabled on the compute node.

```
# hyperscale nodes
```

If there are any failures during the operation, refer to the following log file on the controller node to troubleshoot the issues:

`/var/opt/VRTSofcore/log/<computenode_mgmtIPaddress>_config.log`

Chapter 3

Uninstalling HyperScale

This chapter includes the following topics:

Disabling HyperScale on a compute node

Before you perform the steps in this section, terminate all instances running on the compute node. Otherwise, you cannot disable HyperScale.

To disable HyperScale on a compute node

1. Navigate to **Storage > System > Hypervisors**.
2. In the **Hypervisors** table, locate the hypervisor on which you want to disable HyperScale.

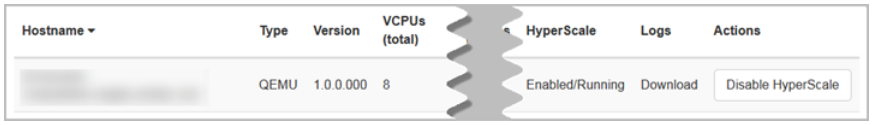

3. In the **Actions** column, select **Disable HyperScale**.

4 On the **Disable HyperScale** dialog box, enter the compute node password and the OpenStack admin password.

5 Click **Disable HyperScale**.

If there are any failures, refer to the log file to troubleshoot the issues. To access the log file for a specific compute node, on the Hypervisors tab, under the Logs column, click the **Download** link for that node.

Removing a data node

Before you remove a data node, make sure of the following:

- HyperScale must be disabled on the compute nodes.
- You have the hypervisor password and OpenStack password of the data node you want to remove.

To remove a data node

1 Navigate to **HyperScale > System > Data Nodes**.

2 On the **All Data Nodes** table, locate the data node, and in the **Actions** column click **Remove Data Node**.

3 On the **Remove Data** node dialog box, enter the hypervisor password and the OpenStack password.

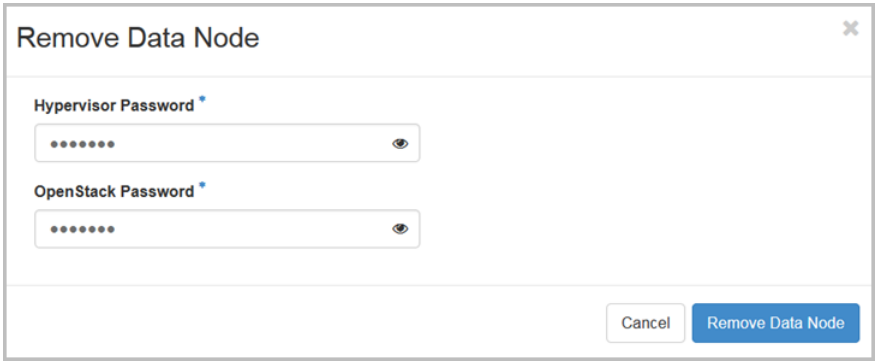

4 Click **Remove Data node**.

5 On the **All Data Nodes** page, the **Status** column is updated as the data node is removed.

If any issues occur during this task, you can download the data node log from the **All Data Nodes** page, or view the log on the following path:

`/var/opt/VRTSofcore/log/<datanode_mgmtip>_config.log`

Removing a data node that has computes enabled is possible only in the following cases:

- A network failure on the data node has triggered a datanode failover and that has caused the vdisks to be evacuated to a peer data node
- A permanent data node service failure where the service does not start even after three restart attempts
- If there is a storage failure

Uninstalling HyperScale

Before you uninstall HyperScale, make sure that there are no data nodes in the environment. In addition, have the following information ready:

- Openstack admin password
- MySql root user password

To uninstall HyperScale

1. Log in to the HyperScale controller as the root user.
2. Enter the following:

```
# ./HyperScale.bin --uninstall
```

3. At the prompts, enter the following credentials.

```
OpenStack Controller's Management IP:
OpenStack Administrator Password:
MySQL Root Password:
```

4. On the confirmation prompt, type **y** to proceed with the uninstallation.
5. Note the system output as HyperScale is uninstalled.

```
Unconfigure current configuration

Verify readiness prior to removing Veritas HyperScale [ DONE ]
Stopping Veritas HyperScale Services                  [ DONE ]
Deleting Veritas HyperScale Database                  [ DONE ]
Reinstating original configuration on the controller  [ DONE ]
Removing Veritas HyperScale packages                  [ DONE ]

Veritas HyperScale Controller unconfigured successfully .
```

This completes the uninstallation process.

Chapter 4

Troubleshooting HyperScale installation

This chapter includes the following topics:

About HyperScale log files

HyperScale includes one or more log files for the following components.

Table 4-1 Controller logs

Log file	Contents
`/var/opt/VRTSofcore/logs/ofengine.log`	Installation information.
`/var/opt/VRTSofcore/logs/<datanode_mgmtIP>_config.log`	Information about add and remove operations performed on a data node.
`/var/opt/VRTSofcore/logs/<computenode_mgmtIP>_config.log`	Information about enable and disable HyperScale operations performed on a compute node.
`/var/opt/VRTSofcore/logs/controller.log`	Information about all HyperScale tasks.

Table 4-1 Controller logs *(continued)*

Log file	Contents
`/var/opt/VRTSofcore/logs/default.log`	Statistics logs.
`/var/opt/VRTSofmn/logs/zookeeper.log`	Zookeeper service logs.

Table 4-2 Compute node logs

Log file	Contents
`var/opt/VRTSofcore/logs/hypervisor.log`	HyperScale MQ service logs on the compute node.
`var/log/hypervisor.log`	Compute storage service logs.

Table 4-3 Data node logs

Log file	Contents
`/var/opt/VRTSofcore/logs/dnhypervisor.log`	HyperScale MQ service logs on the data node.
`/var/log/HyperScale.log`	Data node storage service logs.
`/var/log/ofbud.log`	Information used to troubleshooting the HyperScale bud service.

Table 4-4 External storage provider and NetBackup logs

Log file	Contents
`/var/opt/VRTSofesp/logs/hyperscale-esp.log`	Logs each export or import operation initiation request.
`/var/opt/VRTSofesp/logs/Export_`*`ID`*`/export.log`	Detailed information on export job `Export_`*`ID`*.
`/var/opt/VRTSofesp/logs/Import_`*`ID`*`/export.log`	Detailed information on import job `Import_`*`ID`*.
`media server FQDN``_esp_add.log`	Logs the progress of registering the NetBackup Media Server with HyperScale.

HyperScale installation may fail with a dpkg error

When you run the HyperScale.bin file to install HyperScale, the installation may fail due to a dpkg error.

The following error may appear in the install log:

```
dpkg: error: dpkg status database is locked by another process
```

Workaround:

This error indicates that another installation or update is already in progress. You can run the following to check which process is holding a lock:

```
# fuser /var/lib/dpkg/lock
```

Wait for the process to complete and then reboot the node where the installation failed. This releases the lock held by another process that may be using the dpkg package.

Then run the HyperScale installer again.

Issues with OpenStack Horizon dashboard after HyperScale installation

After installing HyperScale on the controller, you might not be able to logon to the OpenStack Horizon user interface (UI) or the UI might appear distorted or might even stop working altogether.

Workaround:

Perform the following steps:

1. Clear your browser cache and restart the browser, if prompted.
2. Do the following on the controller node:
   - Connect to the controller node using ssh.
   - Navigate to the openstack-dashboard directory:
     `/usr/share/openstack-dashboard`
   - Run the `manage.py` file:
     `# python manage.py compress --force`
   - Restart apache2.service:
     `# systemctl restart apache2.service`

Launch the Horizon UI again and verify that you are able to log on and view the HyperScale dashboard.

Enable HyperScale operation for a compute node may fail due to a system clock drift

After installing HyperScale, when you log on to the Horizon user interface (**Horizon | HyperScale > Hypervisors**) and try to enable HyperScale on a compute node, the operation may fail and the UI may display the following message:

```
Failed to enable compute
```

The following error may appear in the compute node log file:

```
ERROR - Compute node is not in time sync with controller.
DEBUG - compute_precheck.execute for module compute_precheck returned
out = Compute node is not in time sync with controller and err = -1
```

The log message indicates that there may be a time drift between the system clocks on the compute node and the controller node. It may also be the case that the compute node and the controller node are not in the same time zone.

Workaround:

Make sure that the system clocks on all the nodes in the HyperScale configuration are synchronized and the nodes belong to the same time zone.

- To align the time zones, run the following command on all the nodes:

  ```
  # sudo timedatectl set-timezone <desiredtimezone>
  ```

  For example, to set the time zone to America/New_York, the command is:

  ```
  # sudo timedatectl set-timezone America/New_York
  ```

- To synchronize the system clock with the time on the controller node, run the following command on the compute node:

  ```
  # date --set="$(ssh root@<controllernodepublicIPaddress> date)"
  ```

  Note: If there is a time drift on a data node, run this command on the data node.

- After changing the system time, run the following on the compute node:

  ```
  # service nova-compute restart
  ```

  For a data node, run the following command:

  ```
  # service cinder-volume restart
  ```

After fixing the time drift on all the nodes, try to enable HyperScale on the compute node again.