HITACHI
Inspire the Next

Hitachi File Services Driver for OpenStack Version 1.1
Installation and Configuration Guide

Hitachi Data Systems

About this guide

This document contains the installation and user guide of the Hitachi File Services Driver for OpenStack. Although some Manila operations are mentioned in this guide, describing OpenStack operations is out of the scope of this document.

Product version

This guide is applicable for Hitachi File Services Driver for OpenStack 1.1 or later.

Who should use this guide

This guide is intended for anyone who installs and configures Hitachi File Services Driver for OpenStack. This document assumes that the readers have basic knowledge of Linux operating system.

Related information and publications

OpenStack documentation

- OpenStack Cloud Administrator Guide
- OpenStack Command-Line interface Reference
- OpenStack Configuration Reference
- Red Hat Enterprise Linux OpenStack Platform Product Manual
- SUSE OpenStack Cloud Product Manual

Hitachi NAS Platform and Virtual Storage Platform Gx00 documentation

- Hitachi NAS Platform 3080 and 3090 G1 Hardware Reference
- Hitachi NAS Platform 3080 and 3090 G2 Hardware Reference
- Hitachi NAS Platform Series 4000 Hardware Reference
- Hitachi NAS Platform System Manager Unit (SMU) Hardware Reference
- Hitachi NAS Platform and Hitachi Virtual Storage Platform Gx00 Virtual SMU Administration Guide
- Hitachi NAS Platform and Hitachi Unified Storage System Installation Guide
- Hitachi NAS Platform and Hitachi Virtual Storage Platform Gx00 System Access Guide

- Hitachi NAS Platform and Hitachi Virtual Storage Platform Gx00 File Service Administration Guide
- Hitachi NAS Platform and Hitachi Virtual Storage Platform Gx00 Server Cluster and Administration Guide
- Hitachi NAS Platform and Hitachi Virtual Storage Platform Gx00 Storage Subsystem Administration Guide
- Hitachi NAS Platform and Hitachi Virtual Storage Platform Gx00 Backup Administration Guide
- Hitachi NAS Platform and Hitachi Virtual Storage Platform Gx00 User Administration Guide
- Hitachi NAS Platform and Hitachi Virtual Storage Platform Gx00 Network Administration Guide
- Hitachi NAS Platform and Hitachi Virtual Storage Platform Gx00 Antivirus Administration Guide

Getting help

The Hitachi Data Systems Support Connect is the destination for technical support of products and solutions sold by Hitachi Data Systems. To contact technical support, log on to Hitachi Data Systems Support Connect for contact information: https://support.hds.com/en_us/contact-us.html.

Hitachi Data Systems Community is a global online community for HDS customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to https://community.hds.com, register, and complete your profile.

Comments

Please send us your comments on this document: doc.comments@hds.com. Include the document title, number, and revision, and refer to specific section(s) and paragraph(s) whenever possible.

Thank you! (All comments become the property of Hitachi Data Systems Corporation.)

Chapter 1. About Hitachi File Services Driver for OpenStack

This driver provides support for Hitachi NAS Platform models 3080, 3090, 4040, 4060, 4080 and 4100.

The following operations are supported in the driver release 1.1:

- Create and delete NFS shares

- Manage rules to NFS shares (allow/deny access)

- Create and delete snapshots

- Create shares from snapshots

- Manage and Unmanage Shares

- Extend and Shrink NFS Shares

Chapter 2. System requirements

Hitachi NAS Platform requirements

Hitachi File Services Driver 1.1 provides support for HNAS models 3080, 3090, 4040, 4060, 4080 and 4100 with NAS OS and SMU software version 12.2 or higher. Before using the driver, use the HNAS configuration and management utilities, such as GUI (SMU) or SSC CLI, to create a storage pool (span) and an EVS.

OS and platform support

Hitachi File Services Driver 1.1 is supported for Red Hat Enterprise Linux OpenStack Platform, SUSE OpenStack Cloud, and Ubuntu Cloud in the versions compatible with the OpenStack Liberty release. See Hitachi OpenStack driver support matrix on the Hitachi Data System website: https://www.hds.com/en-us/products-solutions/application-solutions/hitachi-and-openstack.html.

Chapter 3. Driver configuration

Overview

To configure the Hitachi File Services Driver, make sure that the controller and compute nodes have access to the NAS server management port, and compute and neutron nodes have access to the data ports (EVS IPs or aggregations). If manila-share service is not running on controller node, it must have access to the management port. The driver configuration can be summarized in the following steps:

1. Create a file system on HNAS for consumption by Manila. Make sure that the file system is not created as a replication target.
2. Install and configure an OpenStack environment with default Manila parameters and services.
3. Install historwrapper and manila-hnas packages.
4. Configure the NAS server parameters on manila.conf.
5. Prepare the network.
6. Configure/create share type.
7. Restart the services.
8. Configure the network.

The following sections detail the summary steps 3 through 8 above. The steps 1 and 2 are outside the scope of this guide. Refer to the documentation provided with Hitachi NAS Platform and the OpenStack Manila for detailed instructions.

Package Installation

Before you start the installation, make sure you have access to the following two packages:

- `historwrapper-1.1-1.x86_64.rpm`

- `openstack-manila-hnas-1.1-1.x86_64.rpm`

As a prerequisite, you must first install the package HistorWrapper:

> $ rpm -ivh historwrapper-1.1-1.x86_64.rpm

Once you have installed HistorWrapper, install Manila-HNAS package:

> $ rpm -ivh openstack-manila-hnas-1.1-1.x86_64.rpm

Hitachi File Services Driver parameters configuration

The following parameters need to be configured in the **[DEFAULT]** section of

`/etc/manila/manila.conf`:

[DEFAULT]	
Option	**Description**
enabled_share_backends	Name of the section on manila.conf used to specify a backend. E.g. *enabled_share_backends = hnas1*
enabled_share_protocols	Specify a list of protocols to be allowed for share creation. For Hitachi driver this must be: *NFS*

The following parameter need to be configured in the **[oslo_concurrency]** section of

`/etc/manila/manila.conf`:

[oslo_concurrency]	
Option	**Description**
lock_path	A directory to manila create lock files. Make sure this path exists and manila has read/write permission to it. E.g. *lock_path = /opt/stack/manila/manila_locks*

The following parameters need to be configured in the **[backend_name]** section of

`/etc/manila/manila.conf`:

[hnas1]	
Option	**Description**
share_backend_name	A name for the backend. E.g. *share_backend_name = HDS1*
share_driver	Python module path. For Hitachi driver this must be: *manila.share.drivers.hitachi.hds_hnas.HDSHNASDriver*
driver_handles_share_servers	DHSS, Driver working mode. For Hitachi driver **this must be**: *False*
hds_hnas_ip	HNAS management interface IP for communication between Manila node and HNAS. E.g. *hds_hnas_ip = 172.24.44.15*
hds_hnas_user	This field is used to provide username credential to HNAS.
hds_hnas_password	This field is used to provide password credential to HNAS.
hds_hnas_evs_id	ID or Label from EVS which this backend is assigned to (ID and Label can be listed by CLI "evs list" or EVS Management in HNAS Interface).
hds_hnas_evs_ip	EVS IP for mounting shares (this can be listed by CLI "evs list" or EVS Management in HNAS Interface).
hds_hnas_file_system_name	Name of the file system in HNAS, located in the specified EVS.
hds_hnas_cluster_admin_ip0*	If HNAS is in a multi-node cluster, set this parameter with the IP of the cluster's admin node.
hds_hnas_driver_helper	Python class to be used for driver helper. For HiStor backend, this must be: *manila.share.drivers.hitachi.histor.HNASHISTORBACKEND*

- Optional parameters

The following settings show an example of a valid driver configuration:

> **[DEFAULT]**
> enabled_share_backends = hds1
> enabled_share_protocols = NFS
>
> **[hds1]**
> share_backend_name = HDS1
> share_driver = manila.share.drivers.hitachi.hds_hnas.HDSHNASDriver
> driver_handles_share_servers = False
> hds_hnas_ip = 172.24.44.15
> hds_hnas_user = supervisor
> hds_hnas_password = supervisor
> hds_hnas_evs_id = 3
> hds_hnas_evs_ip = 172.24.3.200
> hds_hnas_file_system_name = FS-ManilaDev1
> hds_hnas_driver_helper = manila.share.drivers.hitachi.histor.HNASHISTORBACKEND

Prepare the network

In the driver mode used by Hitachi File Services Driver (DHSS = False), the driver does not handle network configuration, it is up to the administrator to configure the network. It is required that HNAS management interface can be reached from Manila-Share node through Admin Network, while the selected EVS data interface is accessible from OpenStack Cloud, such as through Neutron Flat networking. The following is the sample configuration scenario used to describe the step-by-step instructions later in this section:

Manila-Share Node:

eth0: Admin Network, can reach the NAS server management interface.

eth1: Data Network, can reach the NAS server EVS IP (data interface). This interface is only required if you plan to use Share Migration.
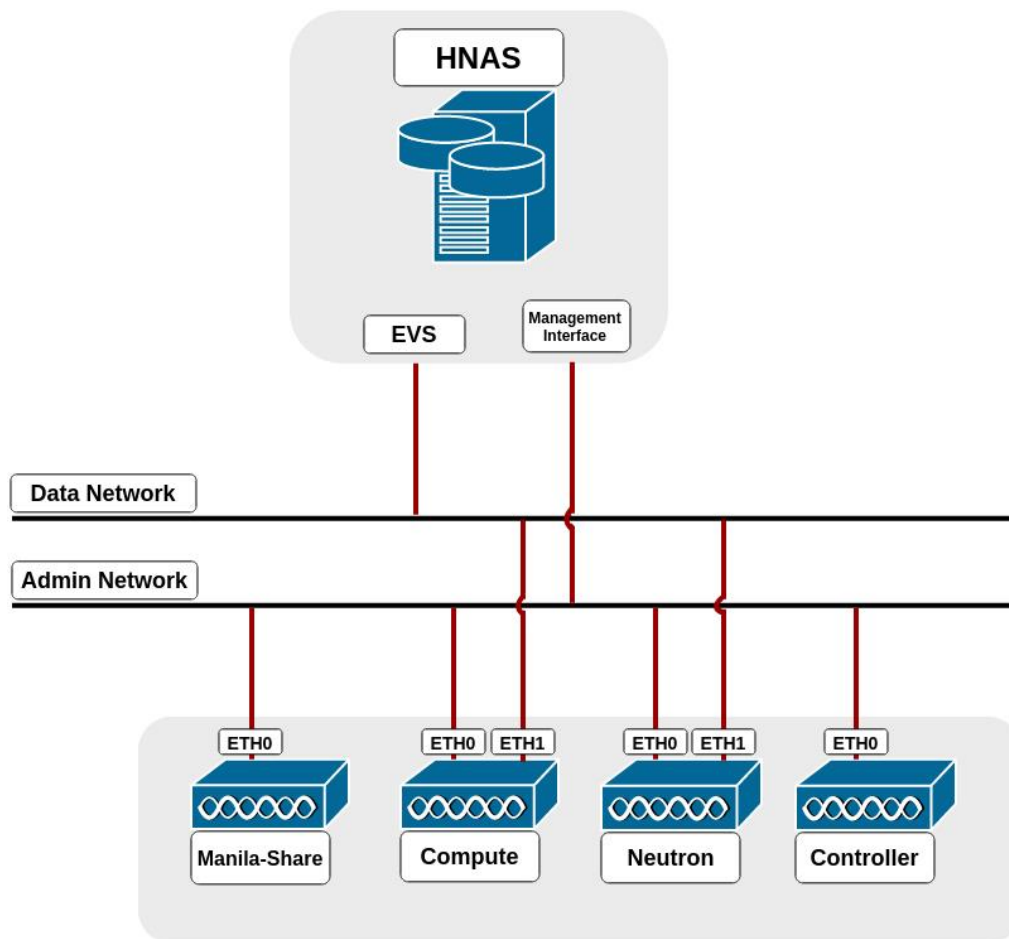
Neutron Node and Compute Nodes:

eth0: Admin Network, can reach the NAS server management interface.

eth1: Data Network, can reach the NAS server EVS IP (data interface).

1. In **Neutron Node**, create a bridge (e.g. `br-hnas`) and add the data interface to this bridge:

```
# ifconfig eth1 down
# ovs-vsctl add-br br-hnas
# ovs-vsctl add-port br-hnas eth1
# ifconfig eth1 up
```

2. Edit `/etc/neutron/plugins/ml2/ml2_conf.ini` (default directory) add (do not replace) this information in the following settings in their respective tags as shown below:

```
[ml2]
type_drivers = flat
mechanism_drivers = openvswitch
```

```
[ml2_type_flat]
flat_networks = physnet2

[ml2_type_vlan]
network_vlan_ranges = physnet2:2000:2500

[ovs]
bridge_mappings = physnet2:br-hnas
```

NOTE: You may have to repeat the last line above in another file in the Compute Node and in Controller Node, if it exists. The file is located in:

`/etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini`

3. In HNAS, create a route to the tenant network. Make sure multi-tenancy is enabled and routes are configured per EVS. Use the command `route-net-add` in HNAS console, where the network parameter should be the tenant's private network, while the gateway parameter should be the FLAT network gateway and the `console-context –evs` parameter should be the ID of EVS in use:

```
$ console-context --evs 3 route-net-add --gateway 192.168.1.1 10.0.0.0/24
```

Share type configuration

Manila requires that the share type include the `driver_handles_share_servers extra-spec`. This ensures that the share will be created on a backend that supports the requested `driver_handles_share_servers` capability. For File Services Driver, this must be set to False.

```
$ manila type-create hitachi False
```

Restart the services

Restart all manila services (manila-share, manila-scheduler and manila-api) and neutron services (neutron-*). This step is specific to your environment. If you are running in devstack for example, you have to log into screen, stop the process and run it again. If you are running it in a distro like RHEL or SUSE, a service command (e.g. `service manila-api restart`) is used to restart the service.

Configure the network

In Neutron Controller it is necessary to create a network, a subnet and to add this subnet interface to a router:

1. Create a network to the given tenant (demo), providing the DEMO_ID (this can be fetched using keystone `tenant-list`), a name for the network, the name of the physical network over which the virtual network is implemented and the type of the physical mechanism by which the virtual network is implemented:

   > $ neutron net-create --tenant-id <DEMO_ID> hnas_network --provider:physical_network=physnet2 --provider:network_type=flat

2. Create a subnet to the same tenant (demo), providing the DEMO_ID, the gateway IP of this subnet, a name for the subnet, the NETWORK_ID created on previously step and CIDR of subnet:

   > $ neutron subnet-create --tenant-id <DEMO_ID> --gateway <GATEWAY> --name hnas_subnet <NETWORK_ID> <SUBNET_CIDR>

3. Add the subnet interface to a router, providing the ROUTER_ID and SUBNET_ID created on previously step (can be fetched using neutron `subnet-list`):

   > $ neutron router-interface-add <ROUTER_ID> <SUBNET_ID>

Additional Notes:

- HNAS has restrictions on the number of EVS, file systems and virtual-volumes. Refer to the documentation provided with Hitachi NAS Platform for specifications for your system.

- Shares and snapshots are thin provisioned. Only the actual used space in HNAS is reported to Manila. In addition, a snapshot does not initially consume any space in HNAS

and stores difference between the share and the snapshot. As such, the space grows as share data is changed.

- Admins should manage the tenant's quota (*manila quota-update*) to control the backend usage.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas
+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa
+44 (0) 1753 618000
info.emea@hds.com

Asia Pacific
+852 3189 7900
hds.marketing.apac@hds.com

Hitachi Data Systems

MK-92ADPTR115-01