

Troubleshoot Hyperflex Datastore Mount Issues in ESXi

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information:](#)

[Problem](#)

[Datastores Inaccessible in vCenter](#)

[Datastores Not Available At All In vCenter/CLI](#)

[Solution](#)

[Network Reachability Check](#)

[IOvisor/ SCVMclient/ NFS Proxy Status Check](#)

[Cluster UUID Resolvable To The ESXI Loopback IP](#)

[Stale Datastore Entries In /etc/vmware/esx.conf](#)

[Check Firewall Rules in ESXi](#)

[Check iptable Rules On The SCVM](#)

[Related Information](#)

Introduction

This document describes the steps you can use to troubleshoot Hyperflex Datastore mount issues.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information:

By default, Hyperflex datastores are mounted in NFS v3.

NFS (Network File System) is a file-sharing protocol used by the hypervisor to communicate with a NAS (Network Attached Storage) server over a standard TCP/IP network.

Here is a description of NFS components used in a vSphere environment:

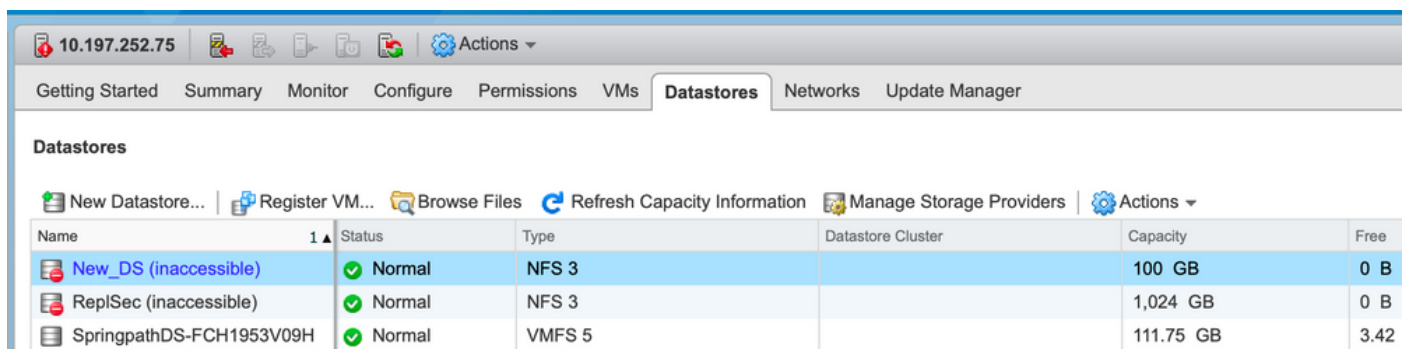
- **NFS server** – a storage device or a server that uses the NFS protocol to make files available over the network. In the Hyperflex world, each controller VM runs an NFS server instance. The NFS server IP for the datastores is the eth1:0 interface IP.
- **NFS datastore** – a shared partition on the NFS server that can be used to hold virtual machine files.
- **NFS client** – ESXi includes a built-in NFS client used to access NFS devices.

In addition to the regular NFS components, there is a VIB installed on the ESXi called the IOVisor. This VIB provides a network file system (NFS) mount point so that the ESXi hypervisor can access the virtual disk drives that are attached to individual virtual machines. From the hypervisor's perspective, it is simply attached to a network file system.

Problem

The symptoms of mount issues can show up in the ESXi host as datastore inaccessible.

Datastores Inaccessible in vCenter



The screenshot shows the vCenter Datastores view. The top navigation bar includes 'Getting Started', 'Summary', 'Monitor', 'Configure', 'Permissions', 'VMs', 'Datastores', 'Networks', and 'Update Manager'. The 'Datastores' tab is active. Below the navigation bar, there are several action buttons: 'New Datastore...', 'Register VM...', 'Browse Files', 'Refresh Capacity Information', 'Manage Storage Providers', and 'Actions'. The main area displays a table of datastores:

Name	Status	Type	Datastore Cluster	Capacity	Free
New_DS (inaccessible)	Normal	NFS 3		100 GB	0 B
ReplSec (inaccessible)	Normal	NFS 3		1,024 GB	0 B
SpringpathDS-FCH1953V09H	Normal	VMFS 5		111.75 GB	3.42

Note: When your Datastores show up as inaccessible in vCenter, they are seen as mounted unavailable in the ESX CLI. This means the datastores were previously mounted on the host.

Check the Datastores via CLI:

- SSH to the ESXi host, and enter the command:

```
[root@node1:~] esxcfg-nas -l
test1 is 10.197.252.106:test1 from 3203172317343203629-5043383143428344954 mounted unavailable
test2 is 10.197.252.106:test2 from 3203172317343203629-5043383143428344954 mounted unavailable
```

Datastores Not Available At All In vCenter/CLI

Name	Status	Type	Datastore Cluster	Capacity	Free
SpringpathDS-FCH1953V09H	Normal	VMFS 5		111.75 GB	3.42

Note: When your Datastores are not present in vCenter or CLI. This indicates that the Datastore was never successfully mounted on the host previously.

- Check the Datastores via CLI

SSH to the ESXi host and enter the command:

```
[root@node1:~] esxcfg-nas -l
[root@node1:~]
```

Solution

The reasons for the mount issue can be different, check the list of checks to validate & correct if any.

Network Reachability Check

The first thing to check in case of any datastore issues is whether the host is able to reach the NFS server IP.

The NFS server IP in the case of Hyperflex is the IP assigned to the virtual interface **eth1:0**, which is present on one of the SCVMs.

If the ESXi hosts are unable to ping the NFS server IP it causes the datastores to become inaccessible.

Find the **eth1:0** IP with the **ifconfig** command on all SCVMs.

Note: The **Eth1:0** is a virtual interface and is present on only one of the SCVMs.

```
root@SpringpathControllerGDAKPUCJLE:~# ifconfig eth1:0
eth1:0 Link encap:Ethernet HWaddr 00:50:56:8b:62:d5
inet addr:10.197.252.106 Bcast:10.197.252.127 Mask:255.255.255.224
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
```

In order to the ESXi host with datastore mount issues and check if it is able to reach the NFS server IP.

```
[root@node1:~] ping 10.197.252.106
PING 10.197.252.106 (10.197.252.106): 56 data bytes
```

64 bytes from 10.197.252.106: icmp_seq=0 ttl=64 time=0.312 ms
64 bytes from 10.197.252.106: icmp_seq=1 ttl=64 time=0.166 m

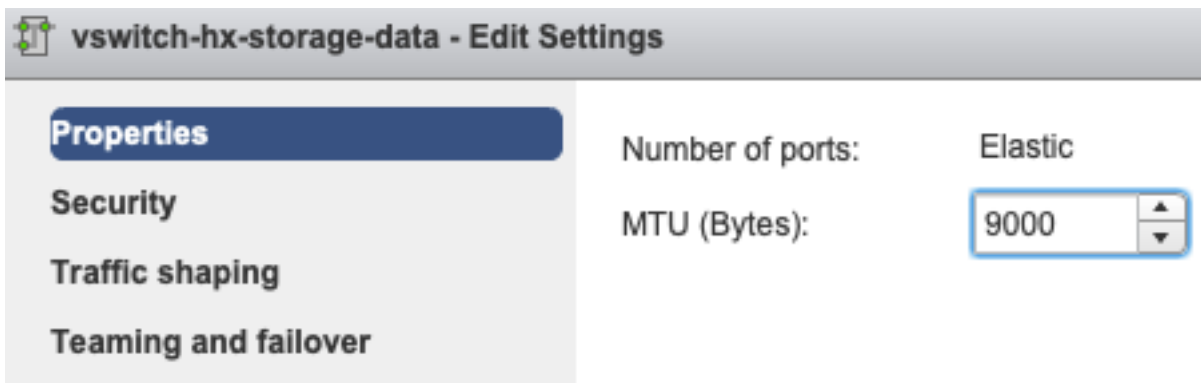
If you are able to ping, proceed with the steps to troubleshoot in the next section.

If you are not able to ping, you have to check your environment to fix the reachability. there are a few pointers that can be looked upon:

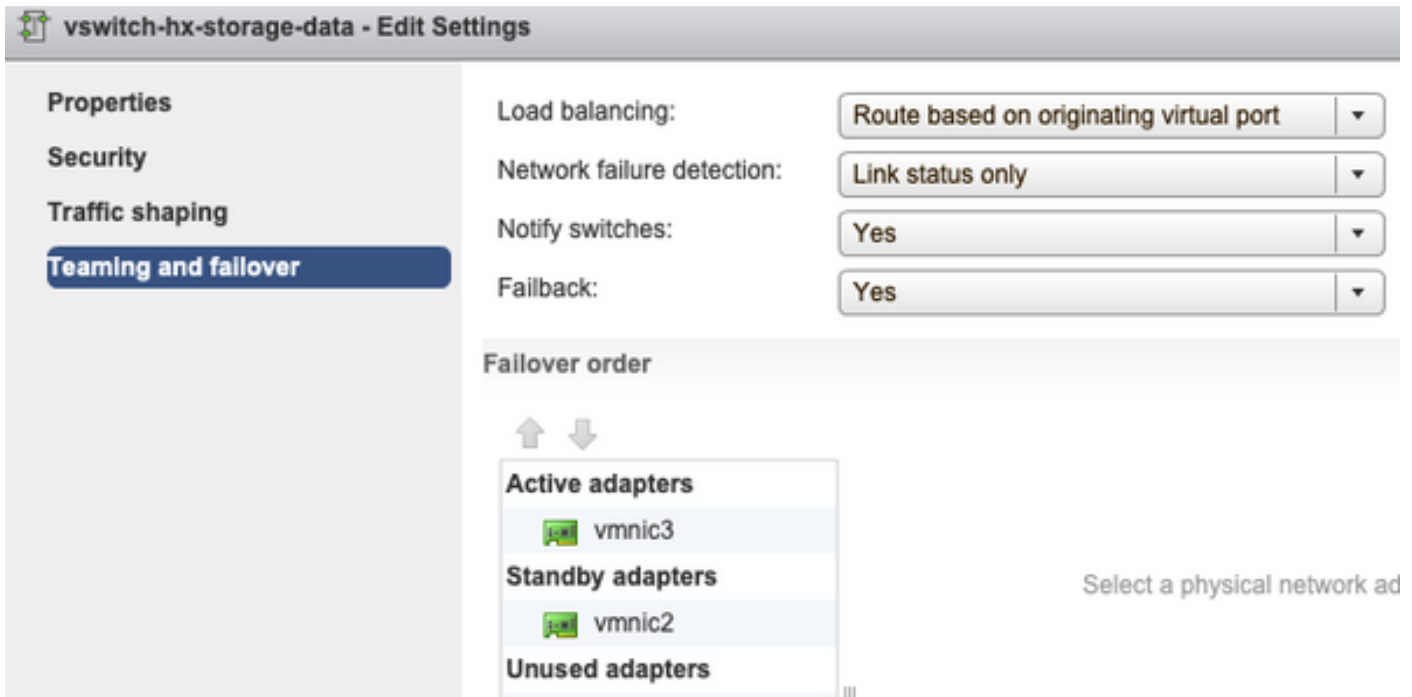
- **hx-storage-data vSwitch Settings:**

Note: By default, all the config is done by the installer during the cluster deployment. If it has been changed manually after that, please verify the settings

MTU Settings - If you have enabled jumbo MTU during cluster deployment, the MTU on the vSwitch must also be 9000. In case you do not use jumbo MTU this must be 1500.



Teaming and Failover - By default, the system tries to ensure that the storage data traffic is switched locally by the FI. Hence the active & standby adapters across all hosts must be the same.



Port Group Vlan settings - The storage-data VLAN must be specified on both **Storage Controller Data Network & Storage Hypervisor Data Network** port groups.

Storage Controller Data Network - Edit Settings

Properties

Security

Traffic shaping

Network label: Storage Controller Data Network

VLAN ID: 953

Storage Hypervisor Data Network - Edit Settings

Properties

Security

Traffic shaping

Network label: Storage Hypervisor Data Network

VLAN ID: 953

No overrides on Port Group level- The **Teaming & Failover** settings done on the vSwitch level get applied to the port groups by default, hence it is recommended to not override the settings on the port-group level.

Storage Hypervisor Data Network - Edit Settings

Properties

Security

Traffic shaping

Teaming and failover

Load balancing: Override Route based on originating virtual port

Network failure detection: Override Link status only

Notify switches: Override Yes

Failback: Override Yes

Failover order

Override

Active adapters

- vmnic3

Standby adapters

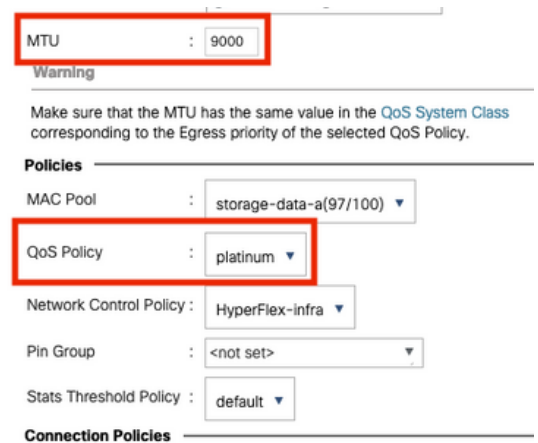
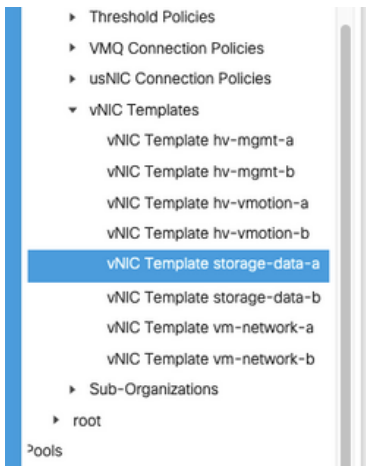
- vmnic2

Select a physical network adapter from the list to view its details.

• **UCS vNIC Settings:**


Note: By default, all the config is done by the installer during the cluster deployment. If it has been changed manually after that, please verify the settings

MTU Settings- ensure the MTU size and QoS policy are configured correctly in the storage-data vnic template. The storage-data vnics use Platinum QoS policy and the MTU must be configured as per your environment.



VLAN Settings - The hx-storage-data VLAN created during the cluster deployment must be allowed in the vnic template. ensure it is not marked as native

LAN / Policies / root / Sub-Organizations / hxcluster1 / vNIC Templates / vNIC Template storage-da...



Iovisor/ SCVMclient/ NFS Proxy Status Check

The SCVMclient vib in the ESXI acts as the NFS Proxy. It intercepts the Virtual Machine IO, sends it to the respective SCVM, and serves them back with the needed info.

Ensure that the VIB is installed on our hosts, for this ssh to one of the ESXI and run the commands:

```
[root@node1:~] esxcli software vib list | grep -i spring
scvmclient 3.5.2b-31674 Springpath VMwareAccepted 2019-04-17
stHypervisorSvc 3.5.2b-31674 Springpath VMwareAccepted 2019-05-20
vmware-esx-STFSNasPlugin 1.0.1-21 Springpath VMwareAccepted 2018-11-23
```

Check the status of the scvmclient on the esxi now and ensure its **running**, if it is stopped please start it with the command **/etc/init.d/scvmclient start**

```
[root@node1:~] /etc/init.d/scvmclient status
+ LOGFILE=/var/run/springpath/scvmclient_status
+ mkdir -p /var/run/springpath
+ trap mv /var/run/springpath/scvmclient_status /var/run/springpath/scvmclient_status.old && cat /var/run/springpath/scvmclient_status.old |logger -s EXIT
+ exec
+ exec
Scvmclient is running
```

Cluster UUID Resolvable To The ESXI Loopback IP

Hyperflex maps the UUID of the cluster to the loopback interface of the ESXi, so that the ESXi passes the NFS requests to its own scvmclient. If this is not there, you can face issues with the datastores mount on the host. In order to verify this, ssh to the host which has datastores

mounted, and ssh to the host with issues, and cat the file **/etc/hosts**

If you see the nonfunctional host does not have the entry in **/etc/hosts**, you can copy it from a functional host into the **/etc/hosts** of the nonfunctional host.

Nonfunctional Host

```
[root@node1:~] cat /etc/hosts
# Do not remove these lines, or various programs
# that require network functionality will fail.
127.0.0.1 localhost.localdomain localhost
::1 localhost.localdomain localhost
10.197.252.75 node1
```

Functional Host

```
[root@node2:~] cat /etc/hosts
# Do not remove these lines, or various programs
# that require network functionality will fail.
127.0.0.1 localhost.localdomain localhost
::1 localhost.localdomain localhost
10.197.252.76 node2
127.0.0.1 3203172317343203629-5043383143428344954.springpath 3203172317343203629-
5043383143428344954
```

Stale Datastore Entries In **/etc/vmware/esx.conf**

If the HX cluster has been recreated without the re-install of ESXi, you might have old datastore entries in the **esx.conf** file.

This does not allow you to mount the new datastores with the same name. You can check all the HX datastores in **esx.conf** from the file:

```
[root@node1:~] cat /etc/vmware/esx.conf | grep -I nas
/nas/RepSec/share = "10.197.252.106:RepSec"
/nas/RepSec/enabled = "true"
/nas/RepSec/host = "5983172317343203629-5043383143428344954"
/nas/RepSec/readOnly = "false"
/nas/DS/share = "10.197.252.106:DS"
/nas/DS/enabled = "true"
/nas/DS/host = "3203172317343203629-5043383143428344954"
/nas/DS/readOnly = "false"
```

if in the output, you see that the old datastore that are mapped and use the old cluster UUID, hence ESXi does not allow you to mount the same-named datastore with the new UUID.

In order to resolve this, is needed to remove the old datastore entry with the command - **esxcfg-nas -d RepSec**

Once removed, retry the mount of the datastore from the HX-Connect

Check Firewall Rules in ESXi

Check for firewall enable settings

It is set to **False**, it causes problems.

```
[root@node1:~] esxcli network firewall get
Default Action: DROP
Enabled: false
Loaded: true
```

Enable it with the commands:

```
[root@node1:~] esxcli network firewall set -e true
```

```
[root@node1:~] esxcli network firewall get
Default Action: DROP
Enabled: true
Loaded: true
```

Check for connection rule settings:

It is set to **False**, it causes problems.

```
[root@node1:~] esxcli network firewall ruleset list | grep -i scvm
ScvmClientConnectionRule false
```

Enable it with the commands:

```
[root@node1:~] esxcli network firewall ruleset set -e true -r ScvmClientConnectionRule
```

```
[root@node1:~] esxcli network firewall ruleset list | grep -i scvm
ScvmClientConnectionRule true
```

Check iptable Rules On The SCVM

Check and match the number of rules on all the SCVMs. If they do not match, open a TAC case to get it corrected.

```
root@SpringpathControllerI51U7U6QZX:~# iptables -L | wc -l
48
```

Related Information

- [Technical Support & Documentation - Cisco Systems](#)