

Hammerspace Configuration Guide

Version 5.2



Table of Contents

About This Documentation	1
1. Adding Storage to Hammerspace	2
1.1. High-Level Workflow	3
1.2. NetApp	5
1.2.1. Volume Settings	5
1.2.2. Volume Export Settings	5
1.2.3. ONTAP Cluster-Mode Requirements	5
ONTAP 7-Mode Requirements	6
1.2.4. Adding a NetApp Storage System	7
Adding NetApp Storage Volumes	7
Step 1: Select Volumes	7
Step 2: Assimilate Data	8
Step 3: Set Manage-to Capacity for the Storage Volume	9
Step 4: Configure the Volume Capabilities	10
Step 5: Configure Additional IPs for the Data Volume (Optional)	11
Step 6: Review the Configuration Summary	12
1.2.5. Troubleshooting NetApp	13
RPCINFO Failed or "No Root Filehandle"	13
Troubleshooting Steps	15
Code 13: Mount Error	15
Troubleshooting Steps	16
Access Denied by NFS	16
Troubleshooting Actions	16
Code 2: Mount Error	17
Troubleshooting Steps	19
Failed to Add Node	19
Solution	19
Vserver API Missing Vserver Parameter	19
Solution	19
1.3. Dell-EMC Isilon	19
1.3.1. Re-enable Basic Authentication	20
1.3.2. Export Settings	20
1.3.3. Isilon ACL Requirements	21
1.3.4. Adding an Isilon Storage System	21
Adding Isilon Storage Volumes	21
Step 1: Select Volumes	22
Step 2: Assimilate Data	22
Step 3: Set Manage-to Capacity for the Storage Volume	23
Step 4: Configure the Volume Capabilities	24
Step 5: Configure Additional IPs for the Data Volume (Optional)	25
Step 6: Review the Configuration Summary	26
1.4. Adding a DSX Storage System	27
1.4.1. Manually Adding a DSX Node	27

1.5. Adding DSX Storage Volumes	28
1.5.1. Step 1: Select Volumes	28
1.5.2. Step 2: Assimilate Data	29
1.5.3. Step 3: Set Manage-to Capacity for the Storage Volume	30
1.5.4. Step 4: Configure the Volume Capabilities	31
1.5.5. Step 5: Configure Additional IPs for the Data Volume (Optional)	32
1.5.6. Step 6: Review the Configuration Summary	33
1.5.7. NFS-Other Storage Type	34
Required Export Settings	34
Adding NFS Other Storage Volumes	35
Step 1: Select Volumes	36
Step 2: Assimilate Data	36
Step 3: Set the Manage-to Capacity for the Storage Volume	37
Step 4: Configure the Volume Capabilities	38
Step 5: Configure Additional IPs for the Data Volume (Optional)	39
Step 6: Review the Configuration Summary	40
2. Adding Cloud or Object Storage	42
2.1. Adding a Bucket as a Shared Volume	44
3. Troubleshooting	47
3.1. Permission Denied When Adding Isilon Storage System	47
3.1.1. Solution	47
3.2. Adding an Isilon Storage System in the GUI Silently Fails	47
3.2.1. Solution	47
3.3. Volumes from the Added Storage System Are Not Visible	48
3.3.1. Solution	48
3.4. Rescanning Storage Due to Newly Added Share or Export	49
3.4.1. Solution	49
Appendices	50
Appendix A: Glossary of Terms	51
Appendix B: Additional Documentation Resources	54
B.1. Licensing and Delivery Portal	54
B.2. Hammerspace Support Hub	54
Appendix C: Creating a Restricted Administrator for ONTAP Cluster Mode	55
C.1. Before You Begin	55
C.2. Create the Role	55
C.3. Create the User	56
C.4. Web Service Access	57
C.5. Privilege Reference	57
C.6. Offloaded Cloning Requires the console Application	58
C.7. Security Considerations	58

About This Documentation

This guide describes how to add and configure storage in a Hammerspace environment. It covers the high-level workflow for bringing storage under Hammerspace management and provides detailed procedures for adding NFS storage systems — including NetApp (ONTAP Cluster-Mode and 7-Mode) and Dell-EMC Isilon — as well as cloud and object storage. DSX store configuration and general troubleshooting are also included.

This guide is intended for storage administrators and infrastructure engineers responsible for expanding or reconfiguring a Hammerspace deployment. Readers should be familiar with NFS storage concepts and the storage platforms being integrated.

For initial installation, see the *Hammerspace Installation and Licensing Guide*. For day-to-day administration, see the *Hammerspace Administration Guide*.

For support resources, see the [Hammerspace Support Hub](#) and the [Hammerspace Licensing and Delivery Portal](#).



Do not install additional or third-party software, agents, or packages directly on Hammerspace Anvil (metadata) or DSX (data) nodes. These nodes run a managed appliance software stack; anything installed outside that stack is unsupported and is removed during a software upgrade. Such software can also interfere with node operation. If you need monitoring or integration with an external tool, contact Hammerspace Support for a supported approach.

Chapter 1. Adding Storage to Hammerspace

Hammerspace supports storage from Network Attached Storage (NAS) systems, block storage—such as direct-attached storage (SSD, NVMe, HDD, etc.)—object storage, and cloud storage.



SAN and iSCSI storage with multi-pathing is only supported in virtual environments. In bare-metal environments, SAN storage is currently supported, with single-path, as technology preview and is not supported for production workloads.

Prior to adding storage resources to Hammerspace, storage volumes must be created on those storage systems using the vendor's administrative tools. Once new storage resources have been added to the product, Hammerspace can place data on them as part of the global file system.

NAS storage systems can be added to Hammerspace either as platform integrated storage systems or generic storage systems. Block storage is added by installing Hammerspace Data Services nodes (DSX) in a virtual machine or bare-metal server and attaching the block storage to the DSX nodes. This uses the DSX Store function, which is one of several functional capabilities included in the DSX.

In Hammerspace, storage is called a storage volume, which is created and hosted by a storage system. Hammerspace identifies each storage system by a unique hostname and IP address. For platform-integrated storage, this identifier is typically the storage system's management interface. For generic storage systems, this identifier is typically the storage system's IP address.





Figure 1. Client protocol data path

1.1. High-Level Workflow

The high-level flow for adding storage is to first add the storage system(s) and then add volumes. This guide is developed using the graphical user interface, notations have been made for the relevant command line reference.

1. Select **Infrastructure** from the left panel, select **Storage Systems** tab and then click **Add Storage System**. For a complete list of what storage systems are supported, please see the Hardware and Software Compatibility List.
2. Enter a Storage System Name. The name is used within the product as a reference to this storage system. It is not used for data access or to resolve IP addresses. After entering a name, select the **Type**.



The image shows a dialog box titled "Add Storage System" with a close button (X) in the top right corner. Inside the dialog, there is a "Name:" label followed by a text input field containing "Storage System Name". Below this, the "Storage Type:" is set to "NAS Storage" (indicated by a selected radio button), with "Object Storage" as an alternative. Under "Storage Type:", there is a "Type:" label and a dropdown menu. The dropdown menu is open, showing a list of "NAS Storage Systems" with "Dell Embedded NAS" selected and highlighted. Other options in the list include Dell EMC Isilon, Dell EMC Unity, Dell EMC VNX, Google Cloud Filestore, NetApp 7-Mode, and NetApp C-Mode. At the bottom right of the dialog, there are two buttons: "Cancel" and "Add Storage System".

Figure 2. Add Storage System dialog

3. Enter the name and press **Add Storage System**. The next section describes this in detail for each vendor.

Admin CLI

```
node-add
```

4. Adding a volume is the final step before storage can be used within the Hammerspace namespace.

Admin CLI`volume-add`

1.2. NetApp

For NetApp (both 7-mode and Cluster mode), use the NetApp Management IP address when adding the storage node.



Hammerspace requires Read-only, Administrator credentials.

Hammerspace uses legacy ZAPI APIs or the recently introduced ONTAP REST APIs to perform management operations (for example, to discover exported volumes to present to the user during an *Add Storage* workflow). Hammerspace will automatically detect when the REST API is available and switch to it when it is detected.



Hammerspace does **NOT** make any changes to the NetApp configuration settings.

Vservers and qtrees are supported.

1.2.1. Volume Settings

- Volumes only need to be exported over NFSv3 (even if client access is over SMB or NFSv4.2).
- Exports must be exported read/write to all NFS clients that will be using the shares for read/write activity.
- Volumes must be exported with **root=<ANVIL IPs>,<DSX IPs>,<Floating IPs>** to ensure proper access by the product. Make sure cluster IPs, node IPs and floating IPs are included. This must be done for all volumes in the Junction path including the root volume.
- Generally, it is required to export the root volume (/) as RO to be able to mount any shares under /.
- Security style set to UNIX for R/W usage (required for R/W assimilation).
- Security style set to **Mixed** or **NTFS** is supported using RO assimilation. For RO assimilation, no RW flags are required.

1.2.2. Volume Export Settings

Volumes must be exported with the following settings to all NFS client IPs that will use the storage directly. This includes NFSv4.2 clients, and Anvil and DSX nodes.

Volume export settings`sec=sys,rw,nosuid`

1.2.3. ONTAP Cluster-Mode Requirements

Ensure that ONTAPI is enabled by entering the **vservers services web show** command in the ONTAP command line.

ONTAP CLI

Command:

```
# vserver services web show
```

Expected output:

```
Vserver Type Service Name Description Enabled
-----
NAS_1 data ontapi Remote Administrative API true
Support
```

If the ONTAPI service is disabled, enter the following command to enable it:

ONTAP CLI

```
# vserver services web modify --enabled true --vserver <vserver-name> -name ontapi
```

Change `<vserver-name>` to match your ONTAP cluster admin vServer. Consult the NetApp product documentation for more information.



For a restricted account instead of the default `admin` account, see [Creating a Restricted Administrator for ONTAP Cluster Mode](#).

ONTAP 7-Mode Requirements

Ensure that the required settings: **httpd.admin.enable**, **httpd.admin.ssl.enable** and **tls.enable** are set to **ON**. To set these values, enter the following commands in the ONTAP command line:

ONTAP CLI

```
# options httpd.admin.enable on
# options httpd.admin.ssl enable on
# options tls.enable on
```



If volume snapshots are turned on for the volume under Hammerspace management, it is highly recommended that they are turned off. Share-level snapshots are available in Hammerspace. Hammerspace snapshots can also be stored in the cloud and on other storage volumes to increase the resiliency of the snapshot data.

1.2.4. Adding a NetApp Storage System

1. Log in to the Management GUI as admin.
2. Navigate to **Infrastructure** > **Storage System** and click **Add Storage System**.
3. Provide the required information and click **Add Storage System**.

Add Storage System

1 Details

Storage System Details

i Provide storage system details.

* Name: NetApp System

Description: Description

* Storage Type: ☒ NAS Storage ☐ Object Storage

* Type: NetApp C-Mode

Access details

* IP Address: Management IP address

* User name: User name

* Password: Password

Close Add Storage System

Figure 3. Adding NetApp C-Mode Storage System

Adding NetApp Storage Volumes

Perform the following workflow to add NetApp storage volumes.

Step 1: Select Volumes

Select the volumes to add from the added NetApp storage system. Note that for C-Mode configurations, the root volume (sometimes exported as /) is not supported as a valid volume, as it is most used to store internal files for ONTAP.

Add Volumes

1 Selections
2 Data
3 Capacity
4 Capability
5 IP
6 Review & Add

Storage System: NetApp

☐ Type
Name
Total
Used
Free
Percentage Used
Storage Type
Browse

☐ Volume
NetApp::svm0:data
Add volume to view information

Showing 1 - 1 of 1 Storage systems

Close
Previous Step
Next Step

Figure 4. Step 1: Select storage volumes to add

Alternatively, you can use the following command to add the volume:

Admin CLI

```
volume-add
```

Step 2: Assimilate Data

The second step will detect existing data on the volume and prompt the user to go through the assimilation process.

Add Volumes

1 Selections 2 Data 3 Capacity 4 Capability 5 IP 6 Review & Add

Specify access and whether to assimilate existing data.

Search

Volume Name	Read Only	Assimilate
NetApp::svm0:data	☐	☐

Showing 1 - 1 of 1 Storage volume

Close Previous Step Next Step

Figure 5. Step 2: Assimilate

Step 3: Set Manage-to Capacity for the Storage Volume

Hammerspace enables live data mobility for all data stored on Hammerspace volumes. With this core functionality, the system can now automatically load-balance across storage volumes without disrupting data access.

The manage-to capacity for volumes allows administrators to set a high threshold for a volume, after which the system will no longer place new data on this volume unless it determines that all other volumes are full and that no object storage volume is available. The data placement logic is driven by a machine-learning engine that automatically accounts for capacity, placement, and user-driven needs when placing data.

The **Low Threshold** setting determines how much data moves from a volume once the high threshold has been reached. The system will schedule mobility from the volume until the low threshold has been met.

The high and low percentage values can be changed on the volume settings screen after the volume has been added.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

Adjust capacities and mobility thresholds.

Q Search

Name	Capacity	Usage	Compression	Threshold	Low Threshold	High Threshold
NetApp::svm0:data	Auto (17.18 GB)	303 KB	—		60 %	75 %

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Next Step

Figure 6. Step 3: Setting the manage-to capacity

Step 4: Configure the Volume Capabilities

Each volume type is configured with a default set of capabilities. It is recommended to leave the Durability and Availability settings to the default values.

Objectives such as `place-on-<volume-name>`, `exclude-from-<volume-name>`, and `confine-to-<volume-name>` are auto-created when following the default GUI workflow. Unchecking the Auto Create box will require the admin to manually create those objectives after the volume is added.

The performance test will do a quick IOPS and throughput test of the volume; these values are used as the basis for performance-oriented objectives. The performance test will take approximately 1-2 minutes per volume.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

Capability

Search

Volume Name	Durability	Availability	Objectives	Performance Test
NetApp::svm0:data	Auto	Auto	☒ Auto Create	☐ Skip

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Next Step

Figure 7. Configuring volume capability

Step 5: Configure Additional IPs for the Data Volume (Optional)

A storage volume can have one or more IP addresses. Add or remove additional IPs as needed. The IP addresses can also be edited after the volume is added using the GUI or the admin CLI. IP addresses are generally auto-discovered.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

IP

Search

Volume Name	Discovered IPs	Additional IPs	Excluded IPs
NetApp::svm0:data	—	—	—

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Next Step

Figure 8. Step 5: Configure additional IPs

Step 6: Review the Configuration Summary

The final screen in the volume add step gives a summary of what will be done when clicking **Add Volume**.

Adding a volume is a background process and may take a few minutes to complete per volume. Progress can be viewed in the task window or on the CLI.

Add Volumes

1 Selections
2 Data
3 Capacity
4 Capability
5 IP
6 Review & Add

Review

Q Search

Volume Name	Shared	Assimilation			Threshold 
		Read Only	Share	Destination Path	
NetApp::svm0:data	—	No	—	—	60% - 75%

Showing 1 - 1 of 1 Storage volume

Close
Previous Step
Add Volume

Figure 9. Adding NetApp storage volumes

Alternatively, you can use the following command to add the volume:

Admin CLI

```
volume-add
```

1.2.5. Troubleshooting NetApp

The following errors can occur when the configuration does not meet expectations. Also, check the general troubleshooting section if needed.

RPCINFO Failed or "No Root Filehandle"

When a NetApp storage system is added to Anvil, the native NetApp APIs are used to detect the storage volumes and the parameters for them.

The following table shows an example of the output from `node-list --full`:

Admin CLI

Command:

```
node-list --full
```

Expected output:

```

Name: NTAP
Type: NetApp C-Mode
Vendor: NetApp
Product name: NetApp Release 9.5P6: Wed Jul 10 18:43:50 UTC 2019
Platform services:
Type: LOGICAL_VOLUME
ID: 20dd5e43-e0d1-11ea-8d2c-005056a83b78
Name: pm-svm-3:nfs_svm3_1
Exported path: /nfs_svm3_1
#IPs: [10.200.79.6/22]#
Node: NTAP
Usage: DS
Reserved: false
Capacity: [Total: 4.2GB, Used: 421.8KB (<1%), Free: 4.2GB]
Extended info: volume: nfs_svm3_1; vserver: pm-svm-3
Type: LOGICAL_VOLUME
ID: 07fdea15-ac1e-11ea-8855-005056a83b78
Name: pm-svm-1:qa_vol
Exported path: /qa_vol
#IPs: [10.200.79.3/22]#
Node: NTAP
Usage: DS
Reserved: false
Capacity: [Total: 20.4GB, Used: 5.2GB (26%), Free: 15.1GB]
Extended info: volume: qa_vol; vserver: pm-svm-1
Try 'license-add --help' for more information.

```

The output shows that the detected IP addresses are different between the two volumes. In this example the two different volumes belong to different SVMs with different data interfaces, hence the different IP addresses.

When adding a volume to an Anvil node, the system first checks for NFS access to that volume and uses any IP addresses detected for it.

If an Anvil or DSX cannot reach the IP address(es) detected for the volume, the following message will show during the add-volume wizard in the GUI:




Figure 10. RPCINFO failure alert in the volume-add wizard

Admin CLI

Command:

```
volume-add --node-name NTAP --logical-volume-name pm-svm-3:nfs_svm3_1
```

Expected output:

```
volume-add: No root filehandle found for pm-svm-3:nfs_svm3_1:/nfs_svm3_1.
```

Troubleshooting Steps

Verify that the NFS service is turned on for the NetApp SVM.

Verify that a regular Linux client can see the NFS export by issuing the `showmount` command. Note that the IP address used is from the volume.

Linux shell command

```
# showmount -e 10.200.79.6
clnt_create: RPC: Program not registered
```

In this example, the NFS service was not turned on in the SVM. By turning on the NFS service, the `showmount` command results in the following output:

Linux shell command

```
# showmount -e 10.200.79.6
Export list for 10.200.79.6:
/nfs_svm3_1 (everyone)
/ (everyone)
```

Code 13: Mount Error

The following error is seen when adding a volume:

Error message

```
The logical-volume 'svm0:data' failed diagnostic with reason: '10.200.79.4:/data failed I/O test with
error: Exception('Exception: Mount error. Error code 13',)'.

```

The diagnostic starts with the mount command issued against `10.200.79.4:/data` and it is failing with error code 13. Error code 13 in NFS is permission denied.

The error experienced in a Linux shell is:

Linux shell command

```
# mount -t nfs -o vers=3 10.200.79.6:/nfs_svm3_1 /test
```

```
mount.nfs: access denied by server while mounting 10.200.79.6:/nfs_svm3_1
```

Troubleshooting Steps

This is most likely due to the export policy setting on the volume and the root junction. Both the junction path and in this example, / must allow for the directory to be mounted as root from both Anvil and DSX.

Access Denied by NFS

When the export policies are set correctly, and the NFS service is turned on and/or reachable by the auto-detected IP address and the error **"Access was denied by NFS"** is shown during the volume-add wizard, then it is likely the security style is set to NTFS/mixed. NTFS or mixed security styles are not supported as Read-Write volumes but rather must be added as Read-Only volumes.




Figure 11. Access denied by NFS alert

This can be further validated by looking in the NetApp event logs for a message like the following:

Event message

```
Event:
secd.nfsAuth.noNameMap: vserver (pm-svm-3) Cannot map UNIX name to CIFS name. Error: Get user credentials
procedure failed [ 13 ms] Determined UNIX id 0 is UNIX user 'root' [ 13] Trying to map 'root' to Windows
user 'root' using implicit mapping [ 13] Could not find Windows name 'root' [ 14] Unable to map 'root'.
No default Windows user defined. **[ 14] FAILURE: Name mapping for UNIX user 'root' failed. No mapping
found
```

From a Linux shell, the mount will succeed but no files can be listed or created.

Linux shell commands

```
# mount -t nfs -o vers=3 10.200.79.6:/nfs_svm3_1 /test
# cd /test
-bash: cd: /test: Permission denied

# date >> /test/file
-bash: /test/file: Permission denied
```

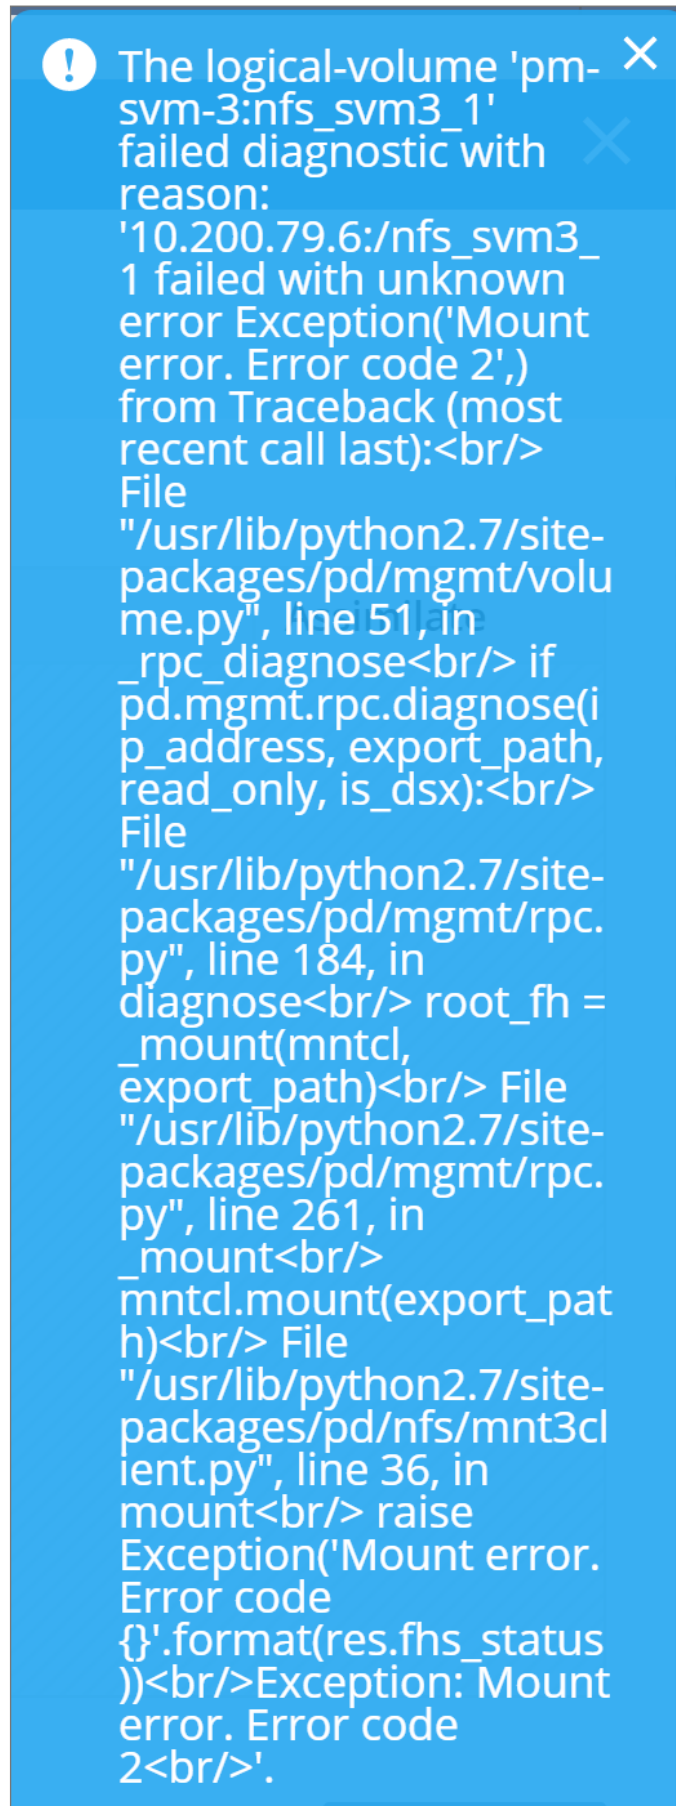
Troubleshooting Actions

- If this is a new volume, change the security style from NTFS/Mixed to UNIX.

- If this is a volume with existing data, add the volume as a Read-Only volume, this will allow the volume to be a source for assimilation jobs.
- If the issue is seen on an NTFS or mixed-mode volume, make sure that a 1:1 mapping exists between `win\administrator` and `root` in the `/etc/usermap.cfg` file.

Code 2: Mount Error

During the volume-add wizard, the error message *Mount error code 2* is displayed. This is most likely because the volume was unmounted, or the junction path changed in the SVM namespace after the Storage System was added to Anvil.



! The logical-volume 'pm-svm-3:nfs_svm3_1' failed diagnostic with reason:

'10.200.79.6:/nfs_svm3_1 failed with unknown error Exception('Mount error. Error code 2',) from Traceback (most recent call last):
File

"/usr/lib/python2.7/site-packages/pd/mgmt/volume.py", line 51, in _rpc_diagnose
 if pd.mgmt.rpc.diagnose(ip_address, export_path, read_only, is_dsx):
File

"/usr/lib/python2.7/site-packages/pd/mgmt/rpc.py", line 184, in diagnose
 root_fh = _mount(mntcl, export_path)
File

"/usr/lib/python2.7/site-packages/pd/mgmt/rpc.py", line 261, in _mount
 mntcl.mount(export_path)
File

"/usr/lib/python2.7/site-packages/pd/nfs/mnt3client.py", line 36, in mount
 raise Exception('Mount error. Error code {0}'.format(res.fhs_status))
Exception: Mount error. Error code 2
'.

Figure 12. Code 2: Mount error

Troubleshooting Steps

1. Verify that the volume is mounted and exported in the SVM namespace.
2. Correct the problem by mounting the volume in the SVM namespace.
3. Re-scan the storage system from the Anvil to detect junction path changes.

Failed to Add Node

When trying to add a NetApp storage node, a message like this is displayed:

Error message

```
Failed to add node <netappname> - Authentication to the NetApp storage system at <IP> failed with the supplied credentials
```

Solution

This is most likely because the system was detected to have the REST APIs enabled but Hammerspace was not allowed access to those APIs.

If you are using a restricted administrator account rather than the built-in `admin` account, see [Creating a Restricted Administrator for ONTAP Cluster Mode](#) to confirm the account has the `http` and `ontapi` applications and the required read-only privileges.

Vserver API Missing Vserver Parameter

When trying to add a NetApp storage node, a message like this is displayed:

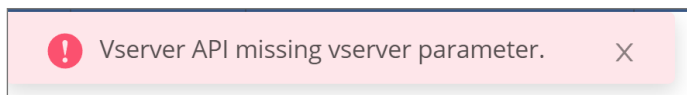


Figure 13. Vserver API missing vservers parameter

Solution

This is most likely because the wrong NetApp endpoint was specified, and the API could not be executed.

1.3. Dell-EMC Isilon

For Isilon, the Management IP is a cluster IP address. Hammerspace will use the credentials provided and discover the exported volumes to present to the user during the *add-storage* workflow.

The credentials provided will also be used when performing file cloning operations via the Isilon API. A dedicated user on the Isilon can be created to monitor and audit the activity if desired.



Hammerspace does **NOT** make any changes to the Isilon configuration settings.

The credentials must, at minimum, have the following privileges:

Table 1. Access privileges on Dell Isilon storage

Privilege	Description	Access	Function
ISI_PRIV_LOGIN_PAPI	Platform API	Read-Only	Discovery
ISI_PRIV_NETWORK	Network	Read-Only	Discovery
ISI_PRIV_NFS	NFS	Read-Only	Discovery
ISI_PRIV_QUOTA	Quota	Read-Only	Discovery
ISI_PRIV_SNAPSHOT	Snapshot	Read-Only	Cloning
ISI_PRIV_IFS_RESTORE	Restore	Read-Only	Cloning
ISI_PRIV_IFS_BACKUP	Backup	Read-Only	Cloning



Starting with Isilon OneFS 9.2 (and later), HTTP basic authentication is turned off by default, causing the storage system-add operation to fail.

1.3.1. Re-enable Basic Authentication

HTTP basic authentication must be re-enabled using the following instructions:

1. Log in to the WebUI as a user with administrative privileges.
2. Click **Protocols** > **HTTP settings**.
3. In the Authentication settings section, click the Authentication drop-down and select **HTTP basic authentication**.

The ISI CLI can also be used to change this setting:

Isilon CLI

```
# isi_config -t web-config auth_basic=true
```

1.3.2. Export Settings

- Exports only needs to be exported over NFS v3 (even if client access to shares and exports are over SMB or NFS 4.1/4.2). The exception for that is if data needs to be assimilated using SMB.
- Enable SNMP and allow SNMP v2 access and set the Community Name Read-only string. This is used to authenticate SNMP API access.
- Exports must be exported with the option **Root Clients** set to all Anvil and DSX IP addresses. Make sure cluster IPs, node IPs, and floating IPs are included.
- Exports must be exported read/write to all NFS 4.2 clients that will use the shares.
- The advanced Export option, **Map lookup UID**, must be set to **Off**.

1.3.3. Isilon ACL Requirements

Verify the following settings:

- **ACL policies environment is balanced:** **Protocols > ACLs > ACL Policies** Environment “Balanced” is configured (default setting). If this value is changed to other than Balanced, adding exported share to the Hammerspace could fail. Currently, “Balanced” is the only validated ACL Policies setting.
- **Group owner inheritance is set:** Ensure that Linux and Windows semantics is checked.

Group owner inheritance

☐ When an ACL exists, use Linux and Windows semantics, otherwise use BSD semantics

☐ BSD semantics - Inherit group owner from the parent folder

☒ Linux and Windows semantics - Inherit group owner from the creator's primary group

Figure 14. Isilon group owner inheritance settings

1.3.4. Adding an Isilon Storage System

Complete the following steps to add an Isilon storage system:

1. Navigate to **Infrastructure > Storage System** and click **Add Storage System**.
2. Provide the required input to add the storage system.

Add Storage System
×

1 Details

Storage System Details

i Provide storage system details.

* Name:

Description:

* Storage Type: ☒ NAS Storage ☐ Object Storage

* Type:

Access details

* IP Address:

* User name:

* Password: 🔍

* SNMP:

Close
Add Storage System

Figure 15. Adding an Isilon storage system

Adding Isilon Storage Volumes

Step 1: Select Volumes

Select the Isilon volumes to add.

Add Volumes [X]

1 Selections 2 Data 3 Capacity 4 Capability 5 IP 6 Review & Add

Storage System: Isilon System

Search

☐	Type	Name	Total	Used	Free	Percentage Used	Storage Type	Browse
☐	Volume	Isilon System::ifs/HMS	Add volume to view information					<
☐	Volume	Isilon System::ifs/Woz	Add volume to view information					<
☐	Volume	Isilon System::ifs/apps	Add volume to view information					<
☐	Volume	Isilon System::ifs/cagle	Add volume to view information					<
☐	Volume	Isilon System::ifs/data	Add volume to view information					<
☐	Volume	Isilon System::ifs/df-new-share	Add volume to view information					<
☐	Volume	Isilon System::ifs/df-test	Add volume to view information					<
☐	Volume	Isilon System::ifs/images	Add volume to view information					<
☐	Volume	Isilon System::ifs/movie	Add volume to view information					<
☐	Volume	Isilon System::ifs/pm	Add volume to view information					<
☐	Volume	Isilon System::ifs/prodzone/pr...	Add volume to view information					<
☐	Volume	Isilon System::ifs/prodzone/pr...	Add volume to view information					<
☐	Volume	Isilon System::ifs/roni	Add volume to view information					<

Showing 1 - 13 of 13 Storage systems

Close Previous Step Next Step

Figure 16. Step 1: Select storage volumes to add

Alternatively, you can use the following command to add volumes:

Admin CLI

```
volume-add
```

Step 2: Assimilate Data

The second step will detect existing data on the volume and prompt the user to go through the assimilation process.

Add Volumes

1 Selections 2 Data 3 Capacity 4 Capability 5 IP 6 Review & Add

Specify access and whether to assimilate existing data.

Search

Volume Name	Read Only	Assimilate
Isilon System::/ifs/images	☐	☐

Showing 1 - 1 of 1 Storage volume

Close Previous Step Next Step

Figure 17. Step 2: Assimilate

Step 3: Set Manage-to Capacity for the Storage Volume

Hammerspace enables live data mobility for all data stored on managed volumes. With this core functionality, it is now possible to have the system automatically move data across storage volumes without any disruption to data access. The description of the thresholds below assumes that objectives are not influencing the data placement behavior.

The manage-to capacity for volumes allows administrators to set the desired thresholds of a volume before the system will no longer place new data on this volume unless it determines that all other volumes are above threshold or not available.

Once a volume goes above the low threshold, the system may consider moving data to other volumes that are below threshold. For new files, placement will generally prefer volumes that are below threshold.

The high and low percentage value can be changed on the volume settings screen after the volume has been added.



If volumes are very small sizes, or very large files (as a percentage of the volume size) are used, then the system may exceed thresholds unexpectedly and may take longer to react to capacity changes.

Step 4: Configure the Volume Capabilities

Objectives such as `place-on-<volume name>`, `exclude-from-<volume-name>` and `confine-to-<volume-name>` are auto-created when following the default GUI workflow. Unchecking the Auto Create box will require the admin to manually create those objectives after the volume is added.

The performance test will do a quick IOPS and throughput test of the volume; these values are used as the basis for performance-oriented objectives. The performance test will take approximately 1-2 minutes per volume.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

Capability

Search

Volume Name	Durability	Availability	Objectives	Performance Test
Isilon System::/ifs/images	Auto	Auto	☒ Auto Create	☐ Skip

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Next Step

Figure 19. Configuring volume capabilities

Step 5: Configure Additional IPs for the Data Volume (Optional)

A storage volume can have one or more IP addresses. Add or remove additional IPs as needed. The IP addresses can also be edited after the volume is added using the GUI or the admin CLI. IP addresses are generally auto-discovered.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

IP

Search

Volume Name	Discovered IPs	Additional IPs	Excluded IPs
Isilon System::/ifs/images	—	—	—

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Next Step

Figure 20. Configuring additional IPs

Step 6: Review the Configuration Summary

The final screen in the volume-add workflow gives a summary of what will be done when clicking **Add Volume**.

Adding a volume is a background process and may take a few minutes per volume to complete. The progress of the operation can be viewed in the task window or on the CLI.

Alternatively, you can use the following command to view the operation progress:

The DSX Store service makes block storage available as file storage volumes. Using DSX Store is a simple way to include local NVMe drives and white box storage solutions into the global file system. Block storage is made available using XFS as a file system on top of the block device.

Adding a Hammerspace DSX Store is either fully automated or a single-click operation. During installation, the DSX is configured to connect only to a specific Anvil Cluster IP. If the Anvil administrator credentials were also provided during installation, then the next step of manually adding the DSX has already been completed automatically, and you can proceed to the [Adding DSX Storage Volumes](#) section.

In the **Storage Systems** tab, click on the + (plus sign), on the right side.

Infrastructure									
Volumes	Volume Groups	Storage Systems							
+ Add Storage System		Search						X Share	
Name	Volumes	Volume Status	Type	Used	Free	Size	Percentage Used	Actions	
ssite4-dsx-1.lab.hammer.space								+	

Figure 22. Manually adding a DSX node

The DSX must be added using this method even though it may not have any storage attached to it. Any of the DSX services hosted on the DSX, such as Data Portals, Cloud Mover, and Data Mover, will not function unless the DSX is added to the cluster. It is not required to have storage attached to the DSX node.

1.5. Adding DSX Storage Volumes

To complete the DSX Store configuration, click **+ Volume** button in the **Storage Systems** view or **Add Volume** from the **Volumes** view.

1.5.1. Step 1: Select Volumes

Select the volumes to add from DSX. Each configured block device will automatically be formatted with an XFS file system during the installation process.

See the *Hammerspace 5.2 Administration Guide* for instructions on adding new block storage after the DSX has been installed.



Product guides are available as PDF downloads on the [Hammerspace Licensing and Delivery Portal](#) downloads page.




Figure 23. Step 1: Adding a DSX storage volume

1.5.2. Step 2: Assimilate Data

The second step detects existing data on the volume and prompts the user to go through the data assimilation process.

For a brand-new DSX installation, there will be no existing data to assimilate.

Add Volumes

1 Selections 2 Data 3 Capacity 4 Capability 5 IP 6 Review & Add

Specify access and whether to assimilate existing data.

Search

Volume Name	Read Only	Assimilate
tf-dsx-1.lab.hammerspace://hsvol0	☐	☒

Showing 1 - 1 of 1 Storage volume

Close Previous Step Next Step

Figure 24. Step 2: Assimilating the data

1.5.3. Step 3: Set Manage-to Capacity for the Storage Volume

Hammerspace enables live data mobility for all data stored on Hammerspace volumes. With this core functionality, it is now possible to have the system automatically load-balance across storage volumes without any disruption for data access.

The manage-to capacity for volumes allows administrators to set the desired high threshold of a volume before the system will no longer place new data on this volume unless it determines that all other volumes are full and that there is no object storage volume available. The data placement logic is driven by a machine learning engine that automatically takes capacity, placement and user-driven needs into account for placing data.

The Low Threshold setting determines how much data moves off a volume once the High Threshold has been reached. The system will schedule mobility from the volume until the low threshold has been met.

The high and low percentage value can be changed on the volume settings screen after the volume has been added.

Add Volumes

1 Selections
2 Data
3 Capacity
4 Capability
5 IP
6 Review & Add

Adjust capacities and mobility thresholds.

tf-dsx-1.lab.hammer.space::hsvol0

Auto (107.32 GB)

786 MB

—

60

%

75

%

Showing 1 - 1 of 1 Storage volume

<

1

>

Close

Previous Step

Next Step

Figure 25. Step 3: Setting the manage-to capacity

1.5.4. Step 4: Configure the Volume Capabilities

Each volume type is configured with a default set of capabilities. It is recommended to leave the Durability and Availability setting to the default values.

Objectives such as `place-on-<volume name>`, `exclude-from-<volume-name>` and `confine-to-<volume-name>` are auto-created when following the default GUI workflow. Unchecking the Auto Create box will require the admin to manually create those objectives after the volume is added.

The performance test will do a quick IOPS and throughput test of the volume; these values are used as the basis for performance-oriented objectives. The performance test will take approximately 1-2 minutes per volume.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

Capability

Q Search

Volume Name	Durability <i>?</i>	Availability <i>?</i>	Objectives <i>?</i>	Performance Test <i>?</i>
tf-dsx-1.lab.hammer.space::/hsvol0	Auto	Auto	☒ Auto Create	☐ Skip

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Next Step

Figure 26. Step 4: Configure the volume capabilities

1.5.5. Step 5: Configure Additional IPs for the Data Volume (Optional)

A storage volume can have one or more IP addresses. Add or remove additional IPs as needed. The IP addresses can also be edited after the volume is added using the Management GUI or the Admin CLI. IP addresses are generally auto-discovered.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

IP

Search

Volume Name	Discovered IPs	Additional IPs	Excluded IPs
tf-dsx-1.lab.hammer.space::/hsvol0	—	—	—

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Next Step

Figure 27. Step 5: Configure additional IPs

1.5.6. Step 6: Review the Configuration Summary

This final screen in the volume add step gives a summary of what will be done when clicking on **Add Volumes**.

Adding a volume is a background process and may take a few minutes to complete per volume. The progress of the volume add can be viewed in the task window or on the CLI.

Add Volumes [X]

1 Selections 2 Data 3 Capacity 4 Capability 5 IP 6 Review & Add

Review

Search

Volume Name	Shared	Assimilation		Destination Path	Threshold
		Read Only	Share		
tf-dsx-1.lab.hammerspace://hsvol0	—	No	—	—	60% - 75%

Showing 1 - 1 of 1 Storage volume

Close Previous Step Add Volume

Figure 28. Step 6: Adding DSX storage volumes

Alternatively, you can use the following command to add the volume:

Admin CLI

```
volume-add
```

1.5.7. NFS-Other Storage Type

For NAS storage systems other than those listed in the GUI/CLI, the **NFS Other** type can be used to add the third-party storage into the namespace.

The Management IP address represents the actual NFS export IP address. If the storage supports multiple IP addresses, those can be added later using the `volume-update` command.

Required Export Settings

- Exports must be exported over NFS v3 (even if client access is over SMB or NFS v4).
- Exports must be exported read/write to all client IPs that will use the storage.
- Exports must be configured with the Root Clients option set to the Anvil and DSX IP addresses. Make sure cluster IPs, node IPs, and floating IPs are included.
- The root user cannot be squashed (`no_root_squash` export setting) for the Hammerspace product nodes and floating IPs.

- The file system ID must be unique for each export from this NFS server.

Example settings from a generic Linux server

```
/space/app *(rw,sec=sys,no_root_squash,fsid=1)
```

Add Storage System

Name:

Storage Type: ☒ NAS Storage ☐ Object Storage

Type:

Access Details

IP Address:

Description:

Cancel

Add Storage System

Figure 29. Adding NFS Other storage system

Adding NFS Other Storage Volumes

Use the following workflow to add *NFS Other* volumes using the Management GUI.

Step 1: Select Volumes

Select the volumes to add from **NFS Other**. These volumes are discovered using the API equivalent of the `showmount` command. If a volume is missing from the list, ensure that it is properly exported over NFSv3 from the NFS server.




Figure 30. Step 1: Selecting storage volumes to add

Step 2: Assimilate Data

The second step will automatically detect existing data on the volume and prompt the user to go through the assimilation process.

Add Volumes

1 Selections 2 Data 3 Capacity 4 Capability 5 IP 6 Review & Add

Specify access and whether to assimilate existing data.

Search

Volume Name	Read Only	Assimilate
NFS Other:::/space/app	☐	☐

Showing 1 - 1 of 1 Storage volume

Close Previous Step Next Step

Figure 31. Step 2: Assimilating data from added volumes

Step 3: Set the Manage-to Capacity for the Storage Volume

Hammerspace enables live data mobility for all data stored on Hammerspace volumes. With this core functionality, the system can now automatically load-balance across storage volumes without disrupting data access.

The manage-to capacity for volumes allows administrators to set a high threshold for a volume, after which the system will no longer place new data on this volume unless it determines that all other volumes are full and that no object storage volume is available.

The data placement logic is driven by a machine-learning engine that automatically accounts for capacity, placement, and user-driven needs when placing data.

The low threshold setting determines how much data moves off a volume once the high threshold has been reached. The system will schedule mobility from the volume until the low threshold has been met.

The high and low percentage values can be changed on the volume settings screen after the volume has been added.

Add Volumes

1 Selections
2 Data
3 Capacity
4 Capability
5 IP
6 Review & Add

Adjust capacities and mobility thresholds.

Q Search

Name	Capacity	Usage	Compression	Threshold ?	Low Threshold	High Threshold
NFS Other:::/space/app	Auto (313.81 GB)	—	—		60 %	75 %

Showing 1 - 1 of 1 Storage volume

Close
Previous Step
Next Step

Figure 32. Step 3: Setting the manage-to capacity

Step 4: Configure the Volume Capabilities

Each volume type is configured with a default set of capabilities. It is recommended to leave the Durability and Availability settings to the default values.

Objectives such as `place-on-<volume name>`, `exclude-from-<volume-name>`, and `confine-to-<volume-name>` are auto-created when following the default GUI workflow. Unchecking the **Auto Create** box will require the admin to manually create those objectives after the volume is added.

The performance test will do a quick IOPS and throughput test of the volume; these values are used as the basis for performance-oriented objectives. The performance test will take approximately 1-2 minutes per volume.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

Capability

Q Search

Volume Name	Durability?	Availability?	Objectives?	Performance Test?
NFS Other:::/space/app	Auto	Auto	☒ Auto Create	☐ Skip

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Next Step

Figure 33. Step 4: Configure volume capabilities

Step 5: Configure Additional IPs for the Data Volume (Optional)

A storage volume can have one or more IP addresses. Add or remove additional IPs as needed. The IP addresses can also be edited after the volume is added using the Management GUI or the Admin CLI. IP addresses are generally auto-discovered.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

IP

Q Search

Volume Name	Discovered IPs	Additional IPs	Excluded IPs
NFS Other:::/space/app	—	—	—

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Next Step

Figure 34. Step 5: Configure additional IPs

Step 6: Review the Configuration Summary

This final screen in the volume add step gives a summary of what will be done when clicking on **Add Volumes**.

Adding a volume is a background process and may take a few minutes to complete per volume. Progress can be viewed in the task window or on the CLI.

Add Volumes

1 Selections

2 Data

3 Capacity

4 Capability

5 IP

6 Review & Add

Review

Q Search

Volume Name	Shared	Assimilation			Threshold
		Read Only	Share	Destination Path	
NFS Other:::/space/app	—	No	—	—	60% - 75%

Showing 1 - 1 of 1 Storage volume

Close

Previous Step

Add Volume

Figure 35. Step 6: Adding NFS Other storage volume

Alternatively, you can use the following command to add **NFS Other** volumes:

Admin CLI

```
volume-add
```

Page 41

© 2026 - Hammerspace

HS3001-USEN-17 | 5.2

Chapter 2. Adding Cloud or Object Storage

Object and cloud storage are fully supported with Hammerspace. Data stored in object storage is treated as inactive. For data to be accessed over NFS or SMB, Hammerspace will automatically rehydrate the data from the object and place it on a file backend while it is being read or written.

The following procedure shows how to add object endpoints and buckets. Amazon S3 is used as a reference example.

1. Select **Infrastructure** from the left panel, then Storage Systems, and click **Add Storage System**.
2. Select **Amazon S3** as the Type and provide an access key and secret key.
3. Selecting an endpoint is optional for Amazon AWS, Microsoft Azure, and Google Cloud. The endpoint selection is automatic, but it can be specified.
4. For better security, it is recommended that the access key and the secret key are specific to an IAM user who can only access data in that bucket.
5. If you need to access Cloud storage via a Proxy, enter the Proxy Host and Port details.
6. Review your settings and click **Add Storage System**.

Add Storage System

Name:

S3 System

Storage Type:

☐ NAS Storage

☒ Object Storage

Type:

Amazon S3

Access Details

Instance profile authentication:

☐ Enable i

Access Key:

Access Key

Secret Key:

Secret Key 🗑

Endpoint:

☒ Auto Select

☐ Select Endpoint

Trust self-signed certificates:

☐

Use virtual-host bucket naming:

☐

Use V4 auth signing:

☒

Proxy Host:

Proxy Host

Proxy Port:

Proxy Port

Proxy Username:

Proxy Username

Proxy Password:

Proxy Password

Cancel

Add Storage System

Figure 36. Add an object storage system

- After Amazon S3 is added, it will be listed under **Storage Systems**.
- To add an object storage volume (represented by a *bucket* in the object storage system), click **+ Volume** next to the source storage system. If there are many buckets, use the search field to narrow down the list.




Figure 37. Selecting an object bucket

- Once a bucket is selected, click **Next Step**. It is supported to add more than one bucket at the same time.

2.1. Adding a Bucket as a Shared Volume

A bucket can be added as a shared volume. This means that other installations will also be able to upload their data in the bucket. This is most commonly used when configuring a Global File System, but it can also be used when multiple sites share a single bucket to leverage global deduplication.

- If desired, set the maximum capacity of the object storage volume. This will restrict how much data can be uploaded into the bucket from this site. If the volume is shared across multiple sites, the same change must be applied to all sites.

Add Volumes

1 Selections 2 Data 3 Capacity 4 Capability 5 IP 6 Review & Add

Specify access and whether to assimilate existing data.

Search

Volume Name	Shared Volume ⁱ	Read Only	Assimilate
S3 System::pm-native-1	☒	—	—

Showing 1 - 1 of 1 Storage volume

Close Previous Step Next Step

Figure 38. Add an S3 bucket as a shared object volume

Add Volumes

1 Selections 2 Data 3 Capacity 4 Capability 5 IP 6 Review & Add

Adjust capacities and mobility thresholds.

Search

Name	Capacity	Usage	Compression	Threshold ⁱ	Low Threshold	High Threshold
 S3 System::pm-native-1	10 TB	—	☐ No Compression		60 %	75 %

Showing 1 - 1 of 1 Storage volume

Close Previous Step Next Step

Figure 39. Setting desired capacity restriction

2. Step through the remaining tabs and add the object storage volume.

Chapter 3. Troubleshooting

This section contains troubleshooting help for adding and configuring storage systems and volumes.

3.1. Permission Denied When Adding Isilon Storage System

Adding storage volumes from an Isilon storage system gives a *permission denied* error message in the wizard or on the CLI.

If the UID mapping is set incorrectly or overlooked as part of configuring the Isilon export for use with Hammerspace, then the following error will show in the UI:




Figure 40. Permission-denied error message

The following error message is shown in the CLI:

Error message

```
volume-add: Failed to discover volume [172.27.102.190, /ifs/home]: /tmp/1541116442901-0/admin/fYzUcgadEI  
(Permission denied)
```

3.1.1. Solution

Go into the Isilon Administration Tool and correctly configure root client access for the Anvil and DSX nodes. It is available under NFS export settings.

3.2. Adding an Isilon Storage System in the GUI Silently Fails

When adding an Isilon storage system via the GUI, a brief progress bar appears and completes; however, nothing is added to the Storage Systems list. In the Tasks window, under progress, you see Failed.

3.2.1. Solution

Either the wrong admin credentials or the incorrect SNMP value has been used.

To view additional details available in the CLI, log in as the **admin** user and run the following command:

Example: Incorrect SNMP value

```
task-list
```

```
ID: fd3b3636-a8c6-4129-a38b-9c1a94f23e37
Name: node-add
Params: name: Isilon8.1; force: false; created-by: Uoid [uuid=6b439e52-3361-4db8-81e8-02cd898be011,
objectType=USER]
Status: FAILED
Created: 2018-11-01 23:35:25 UTC
Started: 2018-11-01 23:35:26 UTC
Ended: 2018-11-01 23:35:38 UTC
Exit Value: Failed to get SNMP interfaces
```

Example: Incorrect login credentials

```
task-list
ID: fb8d5dd3-2c9b-4107-9d57-3c245bc534ae
Name: node-add
Params: name: Isilon; force: false; created-by: Uoid [uuid=6b439e52-3361-4db8-81e8-02cd898be011,
objectType=USER]
Status: FAILED
Created: 2018-11-01 22:36:53 UTC
Started: 2018-11-01 22:36:53 UTC
Ended: 2018-11-01 22:36:56 UTC
Exit Value: failed to identify node. Please check that the node is available, and that the node's
credentials/type are correct.
```

Example: Wrong IP/hostname

```
task-list
ID: af69b73a-6ce9-48ec-b8b9-1e4889287862
Name: node-add
Params: name: Isilon8.1; force: false; created-by: Uoid [uuid=6b439e52-3361-4db8-81e8-02cd898be011,
objectType=USER]
Status: FAILED
Created: 2018-11-01 23:41:28 UTC
Started: 2018-11-01 23:41:28 UTC
Ended: 2018-11-01 23:41:33 UTC
Exit Value: failed to identify node. please check the the node is available, and that the node's
credentials/type are correct
```

3.3. Volumes from the Added Storage System Are Not Visible

You have added a storage system, but cannot see the volume(s) on it.

3.3.1. Solution

Verify from the vendor management tools or a regular Linux client that the storage system is exporting the volumes.

3.4. Rescanning Storage Due to Newly Added Share or Export

You have added a share/export to your storage system after adding it to Hammerspace. How do you re-scan it?

3.4.1. Solution

A manual rescan is required for Hammerspace to detect new exports, shares, volumes or buckets in storage systems. Login into Hammerspace using the CLI and execute the following command:

Admin CLI

```
node-refresh <-id | -name> --rescan
```

The output of the command lists all volumes discovered on the storage system.

Appendices

Appendix A: Glossary of Terms

The following list of terms may be helpful to users of this documentation. For more in-depth definitions of terms relevant to Hammerspace Global File Systems and data orchestration, see the [Dictionary of Terms](#) on the Hammerspace website.

Term	Definition
Alignment	Used to report if a file is currently adhering to all, some, or none of the currently applied objectives.
Aligned	Indicates that all applicable objectives have been met for a file.
Anvil	A Hammerspace metadata server. Responsible for managing all the metadata for the Global File System, and hosts the API and Management GUI services.
Assimilation	The ingestion of all existing file and folder metadata required to serve data using the new Hammerspace global namespace. Can be Read-Write or Read-Only.
Attribute	A type of custom metadata that can be applied and values set by the user using the HSTK.
Cloud Mover	A DSX service responsible for uploading and downloading data from cloud and object storage.
Collections	Subsets of files determined by a defined query of metadata updated on each pass of the sweeper process. Accessible in any directory in a Hammerspace share in .collections.
Condition	Rules used to specify when an objective will be applied. If no condition is specified, the objective will be applied as TRUE.
Data Orchestration	The movement of data based on Hammerspace objectives.
DSX (Data Services nodes)	The Hammerspace data services node. Includes functional capabilities like Portal, Mover, Cloud Mover, and Store.
DSX Store	A DSX service that makes block storage available as file storage volumes using XFS.
Global file share	A share added to more than one Hammerspace cluster. Each cluster has the same view of the share contents.
Hammerscript	The proprietary scripting language used to define and work with metadata, modeled after spreadsheet formula syntax. Used for queries, objectives, and reporting.
Hammerspace Toolkit (HSTK)	A collection of Python-based extensions for command-line interfaces (Linux Bash, Windows CMD/PowerShell) used to query and manipulate file system metadata. See HSTK GitHub for more information.

Term	Definition
Instance	A file placed on Hammerspace will have at least one instance written to the underlying storage; more if the objectives require it.
Label	A type of custom metadata; a simple string that uses an administrator-defined schema.
Metadata	Data that provides information about other data. Hammerspace manages POSIX, Internal, and Custom metadata.
Mover	A stateless DSX service that moves files non-disruptively between file storage volumes.
Objective	Rules that define the intent for managing files. Uses metadata criteria to orchestrate data access and placement.
Offline file	A file considered offline when it is only available in an object storage volume.
Online file	A file considered online when it is in either a DSX volume or a connected NFS storage volume.
Partially aligned	Indicates that a file adheres to at least some of the applicable objectives.
Portal	A DSX service enabling protocol access for NFS, SMB, and S3 clients.
Retention policy	Specified length of time to keep a file system snapshot or file clone.
Schedule	Specified time and frequency to take a snapshot or make a file clone.
Shadow namespace	Provides hidden access to Hammerspace metadata and functionality through a mounted file system, using ?. extensions.
Share	A file system presented for file access by clients via NFS, SMB, and/or S3 protocols. Objectives determine where the file instances are placed.
Site	Refers to a specific Hammerspace cluster, which may or may not be at another physical location.
SLO (Service-level objective)	A target value or range of values measured by service-level indicators. Often refers to an objective to place file instances on specific storage.
SMB (Server Message Block)	A network communication protocol for providing shared access to files, printers, and serial ports between nodes on a network.
Snapshot	A record of the state of the file system at a specific point in time.
Tag	A type of custom metadata; a key-value pair that does not use a predefined schema.
Unaligned	Indicates that a file does not currently adhere to any of the applicable objectives.

Term	Definition
Volume	A storage container used by Hammerspace to store instances of files. Can be DSX storage drives, NFS exports, or object storage buckets.
Volume groups	Used to group storage systems, volumes, or other volume groups for easier targeting in objectives.

Appendix B: Additional Documentation Resources

Additional Hammerspace documentation can be viewed or downloaded as PDF files on the *Hammerspace Licensing and Delivery Portal* and the *Hammerspace Support Hub*. A Hammerspace account is required to access both sites.



Because our titles are available as PDF downloads (and not currently viewable as *always-latest-version* HTML), please check the download pages and download often to verify you are working with the latest version of the documentation.

B.1. Licensing and Delivery Portal

The following Product documentation titles are available on the *Hammerspace License and Delivery Portal*, [Hammerspace Licensing and Delivery Portal](#) downloads page:

- [Hammerspace Hardware and Software Compatibility List](#)
- [Hammerspace Global File System Deployment Guide](#)
- [Hammerspace Installation and Licensing Guide](#)
- [Hammerspace Configuration Guide](#)
- [Hammerspace Administration Guide](#)
- [Hammerspace Command Line Reference](#)
- [Hammerspace Release Notes](#)

B.2. Hammerspace Support Hub

Extensive documentation and guidance are available on the [Hammerspace Support Hub, Resources](#) page. The Support Portal library includes Solution Guides, Technical Briefs, and Knowledge-Base Articles, among other content types. The following titles may be of special interest to users of this guide:

- [{kb-assimilation-url}{kb-assimilation-title}](#) {kb-login-note} [Knowledge Article]: This article includes a technical guide that provides an overview of a typical data migration with Hammerspace using assimilation and data orchestration.
- [{kb-objectives-url}{kb-objectives-title}](#) {kb-login-note} [Knowledge article]: This article includes a technical guide that outlines the configuration of Hammerspace Objectives.
- [{kb-hstk-url}{kb-hstk-title}](#) {kb-login-note} [Knowledge article]: This document introduces the concepts of Hammerspace metadata and the tools to manipulate and use it to achieve business objectives for data.
- [Hammerspace S3 Guide](#) [Solution Guide]: This document covers the creation and management of S3 servers, buckets, and bucket containers, as well as user access control and security considerations, ensuring seamless multi-protocol data access and efficient data management.

Appendix C: Creating a Restricted Administrator for ONTAP Cluster Mode

Adding an ONTAP cluster to Hammerspace requires an administrator account. The default cluster admin account works, but grants far more privilege than Hammerspace uses. This section creates a restricted account with read-only access to only the management APIs Hammerspace needs, plus write access to the two clone APIs required for offloaded cloning.

Verified against ONTAP 9.14.1 and 9.17.1 with Hammerspace 5.2.14.

C.1. Before You Begin

Two constraints determine how the account must be built.

The account must be cluster-scoped. Hammerspace calls `system node show`, `storage aggregate show`, and `system license show`, which exist only at cluster scope. Do not pass `-vserver` to `security login create`. Creating a role with `-vserver` produces a separate SVM-scoped role that a cluster-scoped account never consults, so such entries have no effect.

The account needs three applications. Hammerspace uses the ONTAP REST API for discovery, but offloaded cloning runs over the legacy ONTAPI (ZAPI) `system-cli` call, which opens a CLI session. The account therefore requires `http`, `ontapi`, and `console`. Omitting `console` makes cloning fail silently — see [Offloaded Cloning Requires the console Application](#).

Log in to the cluster as a cluster administrator, and use the **cluster** management LIF address throughout.

C.2. Create the Role

Deny everything by default, then grant back only what Hammerspace needs.

```
security login role create -role hammerspace -access none      -cmddirname DEFAULT
security login role create -role hammerspace -access readonly -cmddirname "vserver show"
security login role create -role hammerspace -access readonly -cmddirname "volume show"
security login role create -role hammerspace -access readonly -cmddirname "volume qtree show"
security login role create -role hammerspace -access readonly -cmddirname "network interface show"
security login role create -role hammerspace -access readonly -cmddirname "system node show"
security login role create -role hammerspace -access readonly -cmddirname "storage aggregate show"
security login role create -role hammerspace -access readonly -cmddirname "system license show"
security login role create -role hammerspace -access readonly -cmddirname "version"
security login role create -role hammerspace -access readonly -cmddirname "cluster identity show"
security login role create -role hammerspace -access all       -cmddirname "volume file clone"
security login role create -role hammerspace -access all       -cmddirname "system node run"
```



These commands produce warnings that additional commands are also being affected, for example `Warning: This operation will also affect the following commands: "volume`

create" "volume modify". This is expected. Granting readonly on a command directory exposes its sibling create and modify entries at read-only level; it does not make them executable.

Equivalent REST call, which returns HTTP 201:

```
curl -sk -u admin -X POST "https://<cluster-mgmt-ip>/api/security/roles" \
-H 'Content-Type: application/json' -d '{
  "name": "hammerspace",
  "owner": {"name": "<cluster-name>"},
  "privileges": [
    {"path": "DEFAULT",          "access": "none"},
    {"path": "vserver show",     "access": "readonly"},
    {"path": "volume show",      "access": "readonly"},
    {"path": "volume qtree show", "access": "readonly"},
    {"path": "network interface show", "access": "readonly"},
    {"path": "system node show",  "access": "readonly"},
    {"path": "storage aggregate show", "access": "readonly"},
    {"path": "system license show", "access": "readonly"},
    {"path": "version",           "access": "readonly"},
    {"path": "cluster identity show", "access": "readonly"},
    {"path": "volume file clone",  "access": "all"},
    {"path": "system node run",    "access": "all"}
  ]
}'
```

Verify the role:

```
security login role show -role hammerspace
```

C.3. Create the User

```
security login create -user-or-group-name hsadmin -application http -authentication-method password
-role hammerspace
security login create -user-or-group-name hsadmin -application ontapi -authentication-method password
-role hammerspace
security login create -user-or-group-name hsadmin -application console -authentication-method password
-role hammerspace
```

Enter a password when prompted. The CLI reads the password from the terminal and cannot be scripted; use the REST equivalent for automation:

```
curl -sk -u admin -X POST "https://<cluster-mgmt-ip>/api/security/accounts" \
-H 'Content-Type: application/json' -d '{
  "name": "hsadmin",
  "password": "<password>",
```

```
"role": {"name": "hammerspace"},
"applications": [
  {"application": "http", "authentication_methods": ["password"]},
  {"application": "ontapi", "authentication_methods": ["password"]},
  {"application": "console", "authentication_methods": ["password"]}
]}'
```

Verify. The `Vserver` column must show the **cluster** name, not an SVM:

```
security login show -user-or-group-name hsadmin
```

C.4. Web Service Access

On ONTAP 9.11 and later, creating the role automatically binds it to the `rest`, `ontapi`, `security`, and `sysmgr` web services. No further action is required. Confirm with:

```
vserver services web access show -role hammerspace
```

On releases earlier than 9.11, bind the role manually:

```
vserver services web access create -vserver <cluster-name> -name rest -role hammerspace
vserver services web access create -vserver <cluster-name> -name ontapi -role hammerspace
```

C.5. Privilege Reference

Command Directory	Access	Required For
DEFAULT	none	Deny-by-default baseline
vserver show	readonly	/api/cluster, /api/svm/svms
volume show	readonly	/api/storage/volumes, /api/private/cli/volume
volume qtree show	readonly	/api/storage/qtrees
network interface show	readonly	/api/network/ip/interfaces — data LIF discovery
system node show	readonly	/api/cluster/nodes
storage aggregate show	readonly	/api/storage/aggregates
system license show	readonly	/api/cluster/licensing/licenses, .../flexclone
version	readonly	/api/private/cli/version
cluster identity show	readonly	/api/cluster uuid and name — without it, node-add fails with Invalid UUID string:

Command Directory	Access	Required For
<code>volume file clone</code>	<code>all</code>	<code>volume file clone start</code> , run via <code>ONTAPI system-cli</code> — the actual clone path
<code>system node run</code>	<code>all</code>	<code>ONTAPI system-cli</code> — clone health probe

The following are **not** required and may be omitted: `cluster show`, `system node autosupport show`, and `security login role show-ontapi`.

Verified against ONTAP 9.14.1 and 9.17.1. Note that ONTAP returns 200 with a field silently omitted when the role lacks the privilege behind that field, so verification must check response **content**, not just status codes.

C.6. Offloaded Cloning Requires the console Application

Offloaded cloning is performed by the `ONTAPI system-cli` call, which opens a CLI session to run `volume file clone start`. ONTAP authorizes that session against the account's console application, independently of its role.



An account without console receives `status="passed"`, an empty `cli-output`, and a `cli-result-value` of `0`. Because there is no error status and no `Error` string in the output, the clone is reported as successful although no clone is created. This affects any account other than the built-in `admin` — including an account holding the full built-in `admin` role.

Confirm with:

```
security login show -user-or-group-name hsadmin
```

`console`, `http`, and `ontapi` must all be listed.

C.7. Security Considerations

`system node run -access all` is required for offloaded cloning and cannot be narrowed — `readonly` is insufficient, because the `ONTAPI system-cli` call is classified as a write operation even when the command it carries is a read.

This privilege permits the account to open an ONTAP clustershell session and run commands there. Each command is still authorized individually against the role, so the account cannot exceed its other privileges — but this is real clustershell access, not a restricted nodeshell, and the account should be treated as more than read-only and its password protected accordingly.

`system node run` alone is not sufficient for cloning: without the `console` application the session cannot be opened at all, and the failure is silent (see [above](#)).

If site policy prohibits this privilege, omit both `system node run` and `volume file clone`. Discovery, volume add, read-write data access, and assimilation all continue to work; only offloaded cloning is unavailable, and clone

operations fall back to host-side copy.