Hammerspace Installation and Licensing Guide

Version 5.3

Table of Contents

About This Documentation

This guide covers the installation, initial configuration, and licensing of Hammerspace software. It includes system requirements and sizing guidelines for Anvil (metadata services) and DSX (data services) nodes, firewall port requirements, step-by-step installation procedures for bare-metal and virtual deployments, quorum server configuration, Terraform-based deployment, and upgrade procedures. Licensing activation through the Hammerspace Licensing and Delivery Portal is also covered.

This guide is intended for infrastructure engineers and system administrators performing a new Hammerspace installation or upgrade. Readers should be familiar with Linux system administration, networking, and hypervisor environments.

For ongoing administration after installation, 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. Installing and Deploying Hammerspace

Hammerspace consists of a set of software services on an appliance-style OS, and can be deployed in the cloud, on virtual machines, or on bare-metal hardware. One or more NAS storage volumes are needed to hold user data. For details on configuring storage, please see the Configuration Guide.

1.1. System Overview

The Hammerspace software platform contains the following main components:

- **Anvil:** Metadata server deployed as a highly available (HA) service and is responsible for managing all the metadata for the global file system. It also hosts the API and Management GUI services.
- **DSX (Data Services nodes):** DSX nodes are the workhorses of the Hammerspace Global Data Platform, and include several different functional capabilities:
  - **Portal:** Enables protocol access for NFS, SMB and S3 clients. The portal is primarily stateless; file locking, such as SMB file locks, are appropriately clustered.
  - **Mover:** The mover service is stateless and runs on each DSX instance. It moves files non-disruptively between file storage volumes.
  - **Cloud Mover:** The cloud mover service is responsible for uploading and downloading data from cloud and object storage. See the hardware compatibility list for the complete support matrix.
  - **Store:** The 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.

If you are deploying Hammerspace in the cloud, see the deployment guide specific to each respective cloud platform. {kb-cloud-topic-url}[{kb-cloud-topic-title}] {kb-login-note} are available in the Hammerspace Support Hub.

Figure 1 illustrates an example of a Hammerspace instance deployed in an enterprise infrastructure. A deployment can be as small as an HA pair of two Anvils and two DSX nodes. However, DSX nodes can scale out to 60 nodes per cluster to support enterprise capacity and performance requirements.

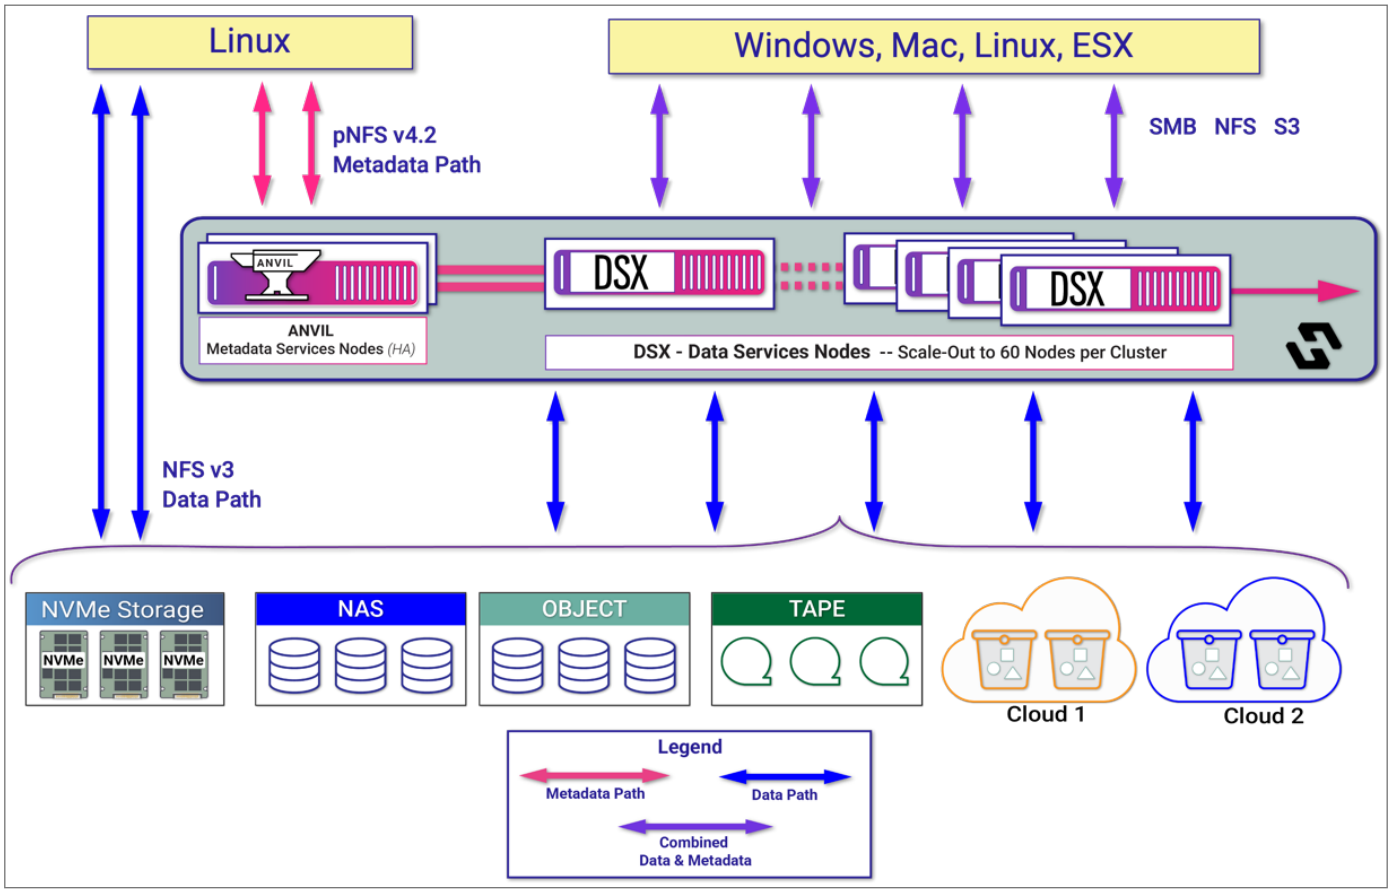

Figure 1. Client protocol data path

1.2. System Requirements

This section provides information on the system requirements for all system elements.

1.2.1. Anvil (Metadata Services Node)

Each Anvil must be equipped with sufficient resources to deliver low-latency metadata operations to clients. The Anvil will make full use of the assigned resources as it runs continuous optimizations of metadata and data placement. It is normal for the Anvil to operate at 20-40% CPU utilization or higher, even with a modest metadata load, in large-file environments. The Anvil automatically scales resources based on available CPU and memory.

Production deployments should not use fewer resources than specified in Table 1.

The figures in this table are production recommendations. The installer enforces its own, lower floor and reports anything below the recommendation as a caution rather than an error, so a system can be installed on less than this table specifies. Size to this table for production; expect the validation screen to be the place where a smaller system announces itself.

The size of the metadata drive depends on the number of files and directories in the system. Please see Sizing Guidelines for more specifics.

Table 1. Anvil requirements

	Virtual machine	Bare metal
CPU[1]	24 vCPU	• Intel Xeon E5 v3 2680 or better. 36 cores minimum. • BIOS set to High Performance or equivalent setting. • It is recommended to disable HyperThreading for maximum performance Any power-saving, performance-throttling CPU settings must be disabled on any bare-metal system that is running Hammerspace software.
Memory	>=64 GB	>=128 GB
Boot drive[2]	200 GB (Hardware RAID recommended). If memory is larger than 160 GB, then the minimum boot drive size is memory x 1.3.	
Additional drives for metadata storage	Minimum (2x) 400 GB, backed by SSD or NVMe.	(2x) 512 GB. NVMe recommended. Hardware RAID is not required.
Network adapter	Minimum 10 GbE	

Notes:

- (1) Only two NUMA nodes are supported. On systems where the BIOS can configure additional NUMA nodes, ensure it is configured for a maximum of two.
- (2) Intel VROC is supported. Details are documented here: Intel VROC support.

Only a 512-byte sector size is supported for boot drives.

1.2.2. DSX (Data Services Node)

The resources required for DSX nodes depend on which services are in use. For example, if a single DSX is used for all the Portal, Store, Cloud Connector, and Mover functions, then it will need more resources than if multiple DSX nodes are deployed. The specifications in Table 2 outline the requirements to get started.

DSX nodes can be easily added (and removed) after initial installation to support a scale-as-you-grow deployment.

The installer's own CPU floor for a DSX node is higher than the minimum shown in this table. A node built to the CPU minimum below is therefore still flagged with a caution on the Virtual Hardware Validation screen. The caution does not block the installation, but if you want a clean validation screen, allocate more CPU than the table's minimum.

Table 2. DSX requirements

	Virtual machine	Bare metal
CPU[1]	>= 4	• Intel Xeon E5 v3 2620v3 or better. 16 cores minimum. • BIOS set to High Performance or equivalent setting. Any power-saving, performance-throttling CPU settings must be disabled on any bare-metal system that is running Hammerspace software.
Memory	>=16GB	>=64 GB
Boot drive[2]	(1x) 100 GB (Recommended)	512 GB (Hardware RAID recommended)
Additional drives	Only required for DSX Store	
Network adapter	Minimum 10 GbE	

Notes:

- (1) Only two NUMA nodes are supported. On systems where the BIOS can configure additional NUMA nodes, ensure it is set to a maximum of two.
- (2) Intel VROC is supported. Details are documented here: Intel VROC support.

Only a 512-byte sector size is supported for boot drives.

See the *Hardware Compatibility List* for additional information regarding specific hardware that has been tested.

The *Hardware Compatibility List* and other Product guides are available as PDF downloads on the Download Packages page of the *Hammerspace License and Delivery Portal*.

1.2.3. Intel VROC Support

Intel VROC (Virtual RAID On a Chip) is an Intel technology that allows systems to create RAID arrays using NVMe SSDs connected directly to the CPU's PCIe lanes, without needing a traditional RAID controller card. Additional Intel VROC product details can be found here: https://www.intel.com/content/www/us/en/software/virtual-raid-on-cpu-vroc.html

Hammerspace supports booting from Intel VROC-configured storage arrays.

Requirements:

- Only supported with UEFI. Traditional BIOS is not supported with Intel VROC.
- Only a single VROC device set is supported. Any additional arrays must be deleted.
- VROC-configured storage is only supported for the boot drive. It is not supported for metadata or data drive.
- It is not supported to use software mirroring on top of VROC arrays. It is expected that the VROC array can provide the required mirroring.

Installation notes:

When booting the ISO, the drive selection screen will look like the screenshot below. Select the device as the boot drive and proceed with the installation.

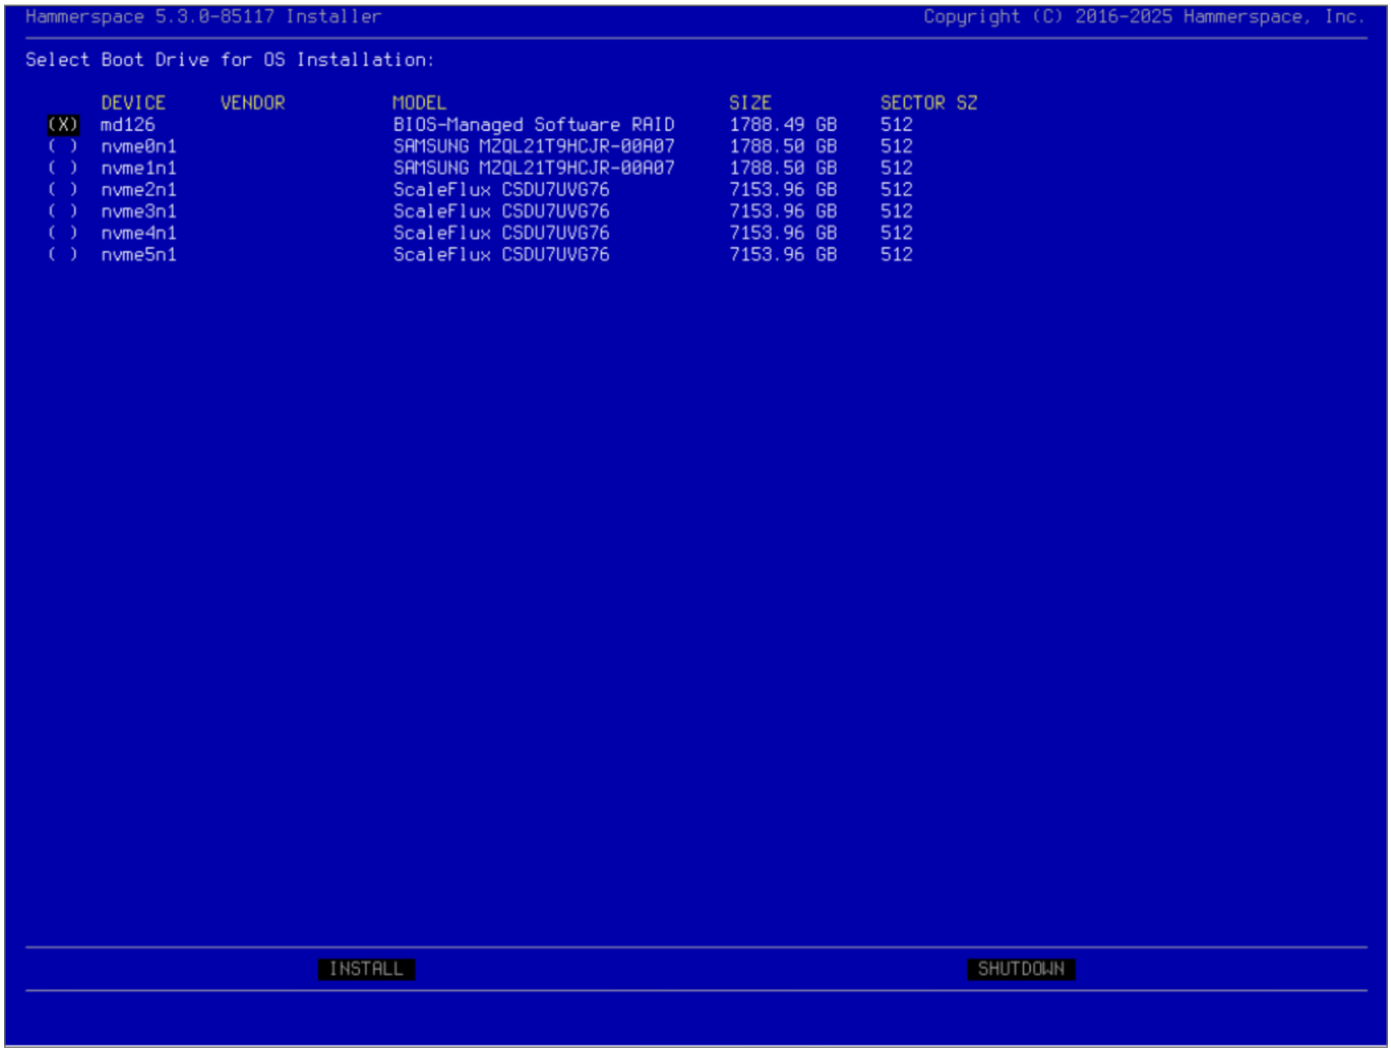

Figure 2. Boot drive selection example using Intel VROC

A RAID1 Intel VROC device will sync in the background on boot. If an unclean shutdown has just occurred, it will resync the entire device. Shutdown gracefully whenever possible.

1.2.4. Sizing Guidelines

Both the Anvil and DSX nodes will auto-size internal tunables based on the resources available at boot time.

This sizing section provides generic recommendations; however, these are provided as guidelines since it is difficult to accurately rely on generic sizing to accommodate the many possible variations found in enterprise environments. The integrated Prometheus exporters provide detailed visibility into the systems utilization and should be used for high-performance environments.

Anvil Sizing Guidelines

To deliver fast metadata operations, it is highly recommended to provide Anvil's metadata drive with NVMe storage.

The size of the metadata drive determines how many files (versions of files, snapshots, etc.) can be stored in the system.

- It is recommended to size the system using **4,096 bytes per unique file and directory**. That covers the most common configurations and regular snapshot usage (e.g., daily/weekly snapshots).
- If many snapshots (256+) are kept, it is recommended to double the size to **8,192 bytes per unique file and directory** to ensure sufficient metadata space.

When the file count scales into the 10s of millions, optimizations built into the metadata storage will reduce the amount of actual capacity used per file. This reduction is not considered in Table 3, as it can be difficult to programmatically predict.

Table 3. Simple sizing table using 4,096 bytes as a reference

Number of files and directories	Recommended metadata capacity
10 million	41 GB
100 million	410 GB
500 million	2 TB
2 billion	8 TB

The other vector for sizing is memory consumption on the Anvil. Each file that is opened and actively used by a client utilizes a small amount of memory, approximately 12k bytes per file.

Metadata Replication Sizing Guidelines

Replication (used to deliver the Global File System functionality) increases the CPU usage. **The exact amount is not static as it depends on the activity in the system and the activity of the remote sites.** It is recommended to allocate approximately 20% of the CPU capacity to handle replication traffic in an active environment; however, in some configurations, it could be necessary to go to 40%.

DSX Sizing Guidelines

DSX nodes have multiple roles, and depending on the function, the sizing approach will be different.

Network throughput is critical as DSX nodes are either serving data, storing data, or moving data to the cloud or between NAS volumes. It is recommended to use at least 10 Gb networking.

Each open SMB connection to a DSX node consumes approximately 1 MB of memory. One thousand SMB clients would consume 1 GB of memory while connected.

CPU consumption on DSX depends on the task. Moving data to and from the cloud is the most CPU-intensive task due to compression.

For environments where >100 GBs of data will be moved to and from object/cloud every day, it is recommended to use a minimum of 8 CPUs / 32 GB of RAM for optimal performance.

VMware Settings

For VMware deployments, it is strongly suggested to reserve 100% of CPU and memory. In vSphere this may be done by setting the Latency Sensitivity setting to **High**. The following table shows settings for installing the cluster on VMware ESXi.

Table 4. VMware ESXi settings

VMware ESXi specific settings	Value
CPU reservation	CPU cores * CPU core MHz
Memory reservation	Reserve all memory
Network adapter	VMXNET 3
SCSI adapter	Para virtualized SCSI controller
Guest OS version	Red Hat Enterprise Linux 7 (64-bit)
VM hardware version	VM version 9 or higher

Time Synchronization in VMware Deployments

It is critical that Anvil and DSX nodes be time synchronized. Hammerspace uses Network Time Protocol (**chrony** implementation), and the Anvil acts as a time server for all DSX nodes.

In VMware environments, vSphere can sometimes adjust the clocks on the virtual machines, which can cause the Hammerspace nodes to lose time sync, potentially causing issues. For this reason, it is necessary to follow the guidance in the VMware knowledge base article, Disabling Time Synchronization (1189).

Before installing Anvil and DSX software in VMware environments, prevent loss of time synchronization by performing the following procedure:

The following instructions are for an ESX 7.x environment.

1. Log in to vCenter vSphere client.
2. Select the virtual machine in the client inventory and choose **Edit Settings…**.
3. Click the **VM Options** tab.
4. Expand the **VMware Tools** section.
5. Make sure **Synchronize guest time with host** is UNSELECTED.

Hyper-V Settings

The following settings are recommended in a Hyper-V environment.

- Use Static MAC addresses for the Anvil and DSX nodes
- Disable Time synchronization using the Hyper-V management tool

1.2.5. Firewall Port Requirements

All Anvil and DSX nodes are protected by their own firewalls for increased security. However, if any other firewalls are used between the nodes, the following ports must be configured correctly for proper operations. The following tables identify the required ports between Anvil and DSX nodes, required connectivity to other NFS systems and object and cloud storage. Note that if certain functionality is not used —for example, syslog or SNMP —then it is not necessary to open those ports.

Global File System Port Requirements

- Ports **443**, **9097**, **9298**, and **9299** must be open between Anvils across all sites.

Anvil Firewall Port Settings

Table 5. Anvil firewall port settings (NIC Roles: Mgmt, Data, HA — Anvil to: DSX, NFS storage, Object / Cloud)

Source	Port	Protocol	Service description	Mgmt	Data	HA	DSX	NFS storage	Object / Cloud
Any	22	TCP	Secure Shell for Management Console	✔		✔	✔		
Any	80	TCP	GUI, Port redirect to https	✔					✔
Any	111	TCP/UDP	Storage, RPC (rpcbind)		✔		✔	✔	
NTP	123	UDP	NTP, Network Time Protocol	✔	✔	✔			
Any	161	TCP/UDP	SNMP	✔	✔				
Any	389, 636	TCP	LDAP		✔				

Source	Port	Protocol	Service description	Mgmt	Data	HA	DSX	NFS storage	Object / Cloud
Any	443	TCP	GUI and REST end point, HTTPS, redirect to 8443	✔					✔
Any	445	TCP	SMB, used for assimilation		✔			✔	
External Syslog	514	UDP	syslogd	✔	✔				
NFS	635	TCP	NFS - mountd	✔	✔			✔	
NFS client	662	TCP/UDP	NFS - statd		✔		✔		
NFS	2049	TCP	NFS		✔		✔	✔	
Anvil	2224	TCP	High Availability			✔			
Anvil	3049	TCP/UDP	HS Store				✔		
Anvil	3121	TCP	High Availability			✔			
DSX	4379	TCP/UDP	SMB High Availability, CTDB		✔		✔		
Anvil	4505 - 4506	TCP	Node management (salt)		✔		✔		
Anvil	5403 - 5405	TCP	High Availability			✔			
External Syslog-TLS	6514	TCP/UDP	Syslog with TLS encryption	✔	✔				
Anvil	7789 - 7812	TCP	High Availability, DRDB			✔			
Anvil	8443	TCP	GUI and REST end point, HTTPS	✔	✔				✔
Anvil	9093	TCP	Services monitoring, Kafka		✔				
DSX	9094	TCP	Data Mover Engine		✔				
Anvil	9095 - 9096	TCP	Data Mover Engine		✔		✔		
Anvil	9097	TCP	Metadata replication		✔				

Source	Port	Protocol	Service description	Mgmt	Data	HA	DSX	NFS storage	Object / Cloud
Anvil	9098	TCP	Cloud mobility		✔		✔		
Anvil	9099	TCP	Cloud mobility, HMDB		✔		✔		
Prometheus server	9100 - 9103	TCP	Prometheus exporters (Anvil)	✔	✔		✔		
Anvil	9298	TCP	HMDB		✔		✔		
Anvil	9299	TCP	Data Manager		✔		✔		
Anvil	9399	TCP	KMS proxy		✔		✔		
Anvil	9929	TCP	High Availability			✔			
NFS	20048	TCP/UDP	NFS - Mount		✔		✔	✔	
NFS	20490 - 20492	TCP	NFS		✔		✔	✔	
Anvil	21064	TCP/UDP	High Availability			✔			
Anvil	30048	TCP	HS Store				✔		
NFS	32803	TCP/UDP	NFS File Locking		✔		✔	✔	

DSX Firewall Port Settings

Table 6. DSX firewall port settings (NIC Roles: Mgmt, Data, Portal — DSX to: Anvil, Object / Cloud)

Source	Port	Protocol	Service description	Mgmt	Data	Portal	Anvil	Object / Cloud
Any	22	TCP	Secure Shell for Management Console	✔			✔	
Any	80	TCP	S3 (HTTP), Object Storage		✔	✔	✔	✔
Any	111	TCP/UDP	Storage, RPC (rpcbind)		✔	✔	✔	
NTP	123	UDP	NTP, Network Time Protocol	✔			✔	
Clients	137 - 139	TCP	SMB data access (NetBIOS)		✔	✔		

Source	Port	Protocol	Service description	Mgmt	Data	Portal	Anvil	Object / Cloud
Any	161	TCP/UDP	SNMP	✔	✔	✔		
Any	389, 636	TCP	LDAP		✔			
Any	443	TCP	S3 (HTTPS), Object Storage		✔	✔	✔	
Clients	445	TCP	SMB data access		✔	✔		
NFS client	662	TCP/UDP	NFS - statd		✔	✔		
NFS	2049	TCP	NFS		✔	✔	✔	
Anvil	3049	TCP/UDP	HS Store		✔			
DSX	4379	TCP/UDP	SMB High Availability, CTDB		✔	✔	✔	
Anvil	4505 - 4506	TCP	Node management (salt)		✔		✔	
Anvil	8443	TCP	HTTPS		✔		✔	✔
Anvil	9093	TCP	Services monitoring, Kafka		✔		✔	
DSX	9094	TCP	Data Mover Engine		✔		✔	
Anvil	9095 - 9096	TCP	Data Mover Engine		✔			
Anvil	9098	TCP	Cloud mobility		✔			
Anvil	9099	TCP/UDP	Cloud mobility, HMDB		✔		✔	
Prometheus server	9100, 9105	TCP	Prometheus exporters (DSX)	✔	✔		✔	
Anvil	9298	TCP	HMDB		✔		✔	
Anvil	9299	TCP	Data Manager				✔	
Anvil	9399	TCP	KMS proxy				✔	
NFS	20048	TCP/UDP	NFS - Mount		✔	✔	✔	

Source	Port	Protocol	Service description	Mgmt	Data	Portal	Anvil	Object / Cloud
NFS	20490 - 20492	TCP	NFS		✓	✓	✓	
Anvil	30048	TCP	HS Store		✓			
NFS	32803	TCP/UDP	NFS File Locking		✓	✓	✓	

1.2.6. Remote Support Services

To simplify the support experience, the product can phone home and send support bundles with product logs. This functionality requires secure connectivity (HTTPS) to **v1proto.sc.hammerspace.com** . It currently resolves to:

- 138.68.228.165
- 67.205.142.28

The IP addresses may change in the future. If this happens, customers will be notified with a support advisory.

1.2.7. Prometheus (and Grafana)

Remote monitoring using Prometheus is supported with integrated collectors. The collectors are turned off by default.

The two things needed during installation are the firewall ports and the command that turns the collectors on. Both are covered in the sections below.

For what the exporters collect, and for storing and visualizing the metrics afterward, see Monitoring Cluster Health in the Administration Guide.

Prometheus Exporter Ports

Remote monitoring with Prometheus requires network connectivity between the Prometheus server and the Hammerspace Anvil and DSX nodes. The exporters listen on the following TCP ports:

Node	Protocol	Ports
Anvil	TCP	9100 – 9103
DSX	TCP	9100 and 9105

Port 9104 is not used by the Hammerspace exporters.

Enabling or Disabling Prometheus Exporters

Use the following command and options to either enable or disable Prometheus exporters in Hammerspace.

1. Log in to the Admin CLI.
2. Run the following command to enable Prometheus exporters:

   Command:

   ```
   cluster-config --prometheus-exporters-enable
   ```

   Expected output (the first lines of the cluster configuration; the **Prometheus exporters** field shows the new state):

   ```
   ID:                                aeeadb7f-7258-41e3-aad9-bc57ec176769
   Name:                              AEL640TQCSS160
   State:                             Standalone
   Management IPs:                    [192.0.2.10/20]
   Data IPs:                          [192.0.2.10/20]
   Cluster floating IPs:              [192.0.2.10/20]
   Since:                             2026-09-22 13:11:34 UTC
   Created:                           2026-09-22 13:11:15 UTC
   Timezone:                          UTC
   NFS transport policy:              TLS forbidden
   GFS participant max suspected time: 30 minutes
   Metered billing:                   Disabled
   Prometheus exporters:              Enabled
   <Additional output removed for brevity>
   ```

Alternatively, to disable the exporters, use the same command but with the `prometheus-exporters-disable` option:

```
cluster-config --prometheus-exporters-disable
```

Either command prints the cluster configuration, and the **Prometheus exporters** field in that output shows the new state. To confirm the current state later, run `cluster-config` with no options and check the same field.

1.3. Pre-Installation Worksheet

Gather the following values and complete the following preparation steps before starting the installation. Having everything ready avoids stopping partway through the installer to go find an IP address or generate a compliant password.

1.3.1. VM or Host Preparation

- Confirm CPU, RAM, and disk meet the Anvil or DSX minimum requirements described earlier in this chapter.
- For virtualized deployments, set CPU and memory reservations to 100%. See the VMware Settings section earlier in this chapter.
- Download the installation ISO and verify its checksum matches the filename.
- Attach the ISO and confirm it will be the boot target. This is especially important when reinstalling onto a system that already has an operating system on disk, since an existing installation can take priority over the ISO during boot without warning.

The Administrator Password set during installation must contain a mix of letters, numbers, and symbols, and must meet a minimum length requirement. Prepare a password that meets this policy in advance to avoid repeated rejected attempts during installation.

1.3.2. Anvil: Per-Node Values

Value	First Node	Second Node
Host Name		
Peer Host Name	(second node's hostname)	(first node's hostname)
Local Domain Name		
DNS Server(s)		
NTP Server(s)		
Timezone		

The installer actively tests each configured DNS server for reachability during setup. If a server cannot be reached, you can select CONTINUE to proceed anyway or BACK to correct the value.

1.3.3. Anvil: Cluster-Wide Values

Value	
Cluster Name	The installer supplies a default; override it if you want something more descriptive.
Administrator Password	
Default Gateway	

1.3.4. Anvil: Per Network Interface Values

Value	Data/Management Interface	HA Interface
IP Address		(not required)
Peer IP Address		(not required)
Data Cluster IP Address (floating, client-facing)		(not required)
Subnet Mask		(not required)
VLAN ID (optional)		
MTU	1500 (or 9000 for jumbo frames)	1500

The Peer IP Address for the second Anvil node is automatically discovered over the private HA interconnect during Part 3 of the installation and does not need to be gathered in advance.

1.3.5. DSX: Per-Node Values

Value	
Host Name	
Local Domain Name	
IP Address (per interface; each on a different subnet)	
Subnet Mask or CIDR	
VLAN ID (optional)	
MTU	
Default Gateway	
Anvil Data Cluster IP (from the Anvil installation)	
Anvil Administrator Password (for auto-join)	

Chapter 2. Installing the Metadata Server (Anvil)

Anvil is installed from a bootable ISO image and contains its own appliance operating system. No other OS license is required. The installation process has been designed to request all required information up front, even before an additional node is installed.

Two-node HA cluster

The ISO allows you to set up several architectural variations; however, the basic workflow described in the following sections installs a two-node HA deployment.

It is highly recommended to use a high availability (HA) deployment for the Anvil metadata server. That includes a minimum of two Anvil nodes, a resilient network configuration, and mirrored metadata drives. A two-Anvil HA deployment also requires an external quorum server; see Installing and Configuring a Quorum Server.

It is supported to use two or three Anvil nodes in a high-availability architecture. When deploying a three-node HA configuration, the third node is added after the initial installation and configuration of the first two nodes.

The installation of a two-node Anvil, high-availability (HA) cluster takes approximately 20 minutes. Most of that time is spent waiting for the system image to be installed and the network parameters to be configured.

The installation of DSX node(s) can be done in parallel and may start before the Anvil installation is complete, but it will not finish until the Anvil nodes are fully installed.

Example installation workflow

In the initial part of the two-node HA installation, you will select a block device to boot from the ISO, determine the installation product (Anvil) to install, validate the hardware, and assign roles to network interfaces. The second part of this installation configures the first Anvil node of the HA cluster. You will repeat these steps to install the second Anvil node of the HA cluster.

2.1. Part 1: Initial Part of the Installation

1. Confirm that the Virtual Machine (VM) or bare-metal server meets the minimum system requirements.

The installer re-checks these requirements later, on the Virtual Hardware Validation screen, and reports anything undersized as a caution rather than an error. Cautions do not stop the installation, so it is worth confirming sizing now rather than deciding at the validation screen whether a caution matters.

2. Load the ISO image and power on the VM or bare-metal server.
3. Ensure that the server boots from the ISO image rather than from its own drives.

Before the first boot, confirm that the CD/DVD device is listed **above** the hard drive in the server's firmware boot order. Many virtual machines place the hard drive first by default.

When that happens the server boots whatever is already on its disks and the installer never appears, which looks like the ISO failing to load even though the ISO was never reached.

On a virtual machine, also confirm that the virtual CD/DVD drive is set to connect at power on. A drive that has the ISO attached but is not connected behaves exactly like an empty drive.

4. At the Hammerspace boot menu, select **Install Hammerspace** and press `Enter`.

The highlighted default on this menu is **Boot from next device in boot order**, and the menu acts on that default automatically after a short countdown. Press an arrow key as soon as the menu appears to stop the countdown, then select **Install Hammerspace**. If the countdown expires, the server boots from its drives instead; reset the server and try again.

Installing Hammerspace is destructive to the existing data on the local drives.

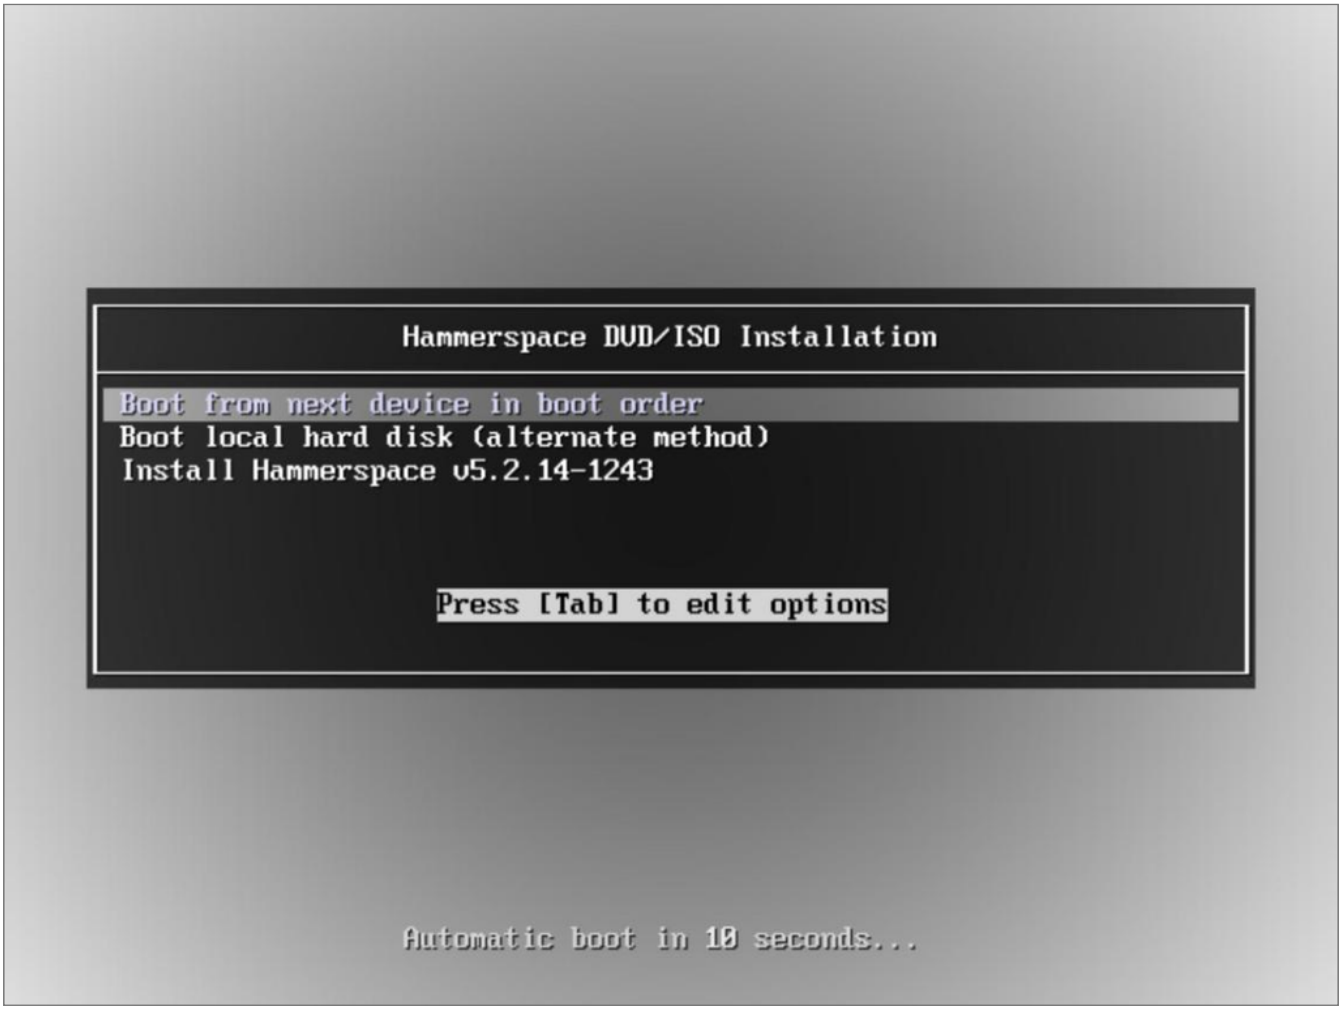

Figure 3. Hammerspace boot menu

5. The installation program automatically discovers local storage devices. Select the boot device for installation and then **Install** to continue.

Boot drives should be at least 200GB. See the Anvil requirements table for full sizing guidance.

Choose the boot device by its size rather than accepting whichever device the installer arrives with selected. The pre-selection is not always the device you want: the same physical disk can enumerate as `sda` on one node and `sdb` on an identically built one, so the pre-selection lands correctly on one server and on the wrong disk on the next.

Selecting a device other than the first one listed adds a caution that the installation may fail to boot unless the firmware can boot from that device. The caution names the device concerned. This is a general warning rather than a detected fault.

All discovered block devices will be overwritten during the installation. Any existing data on these devices will be erased.

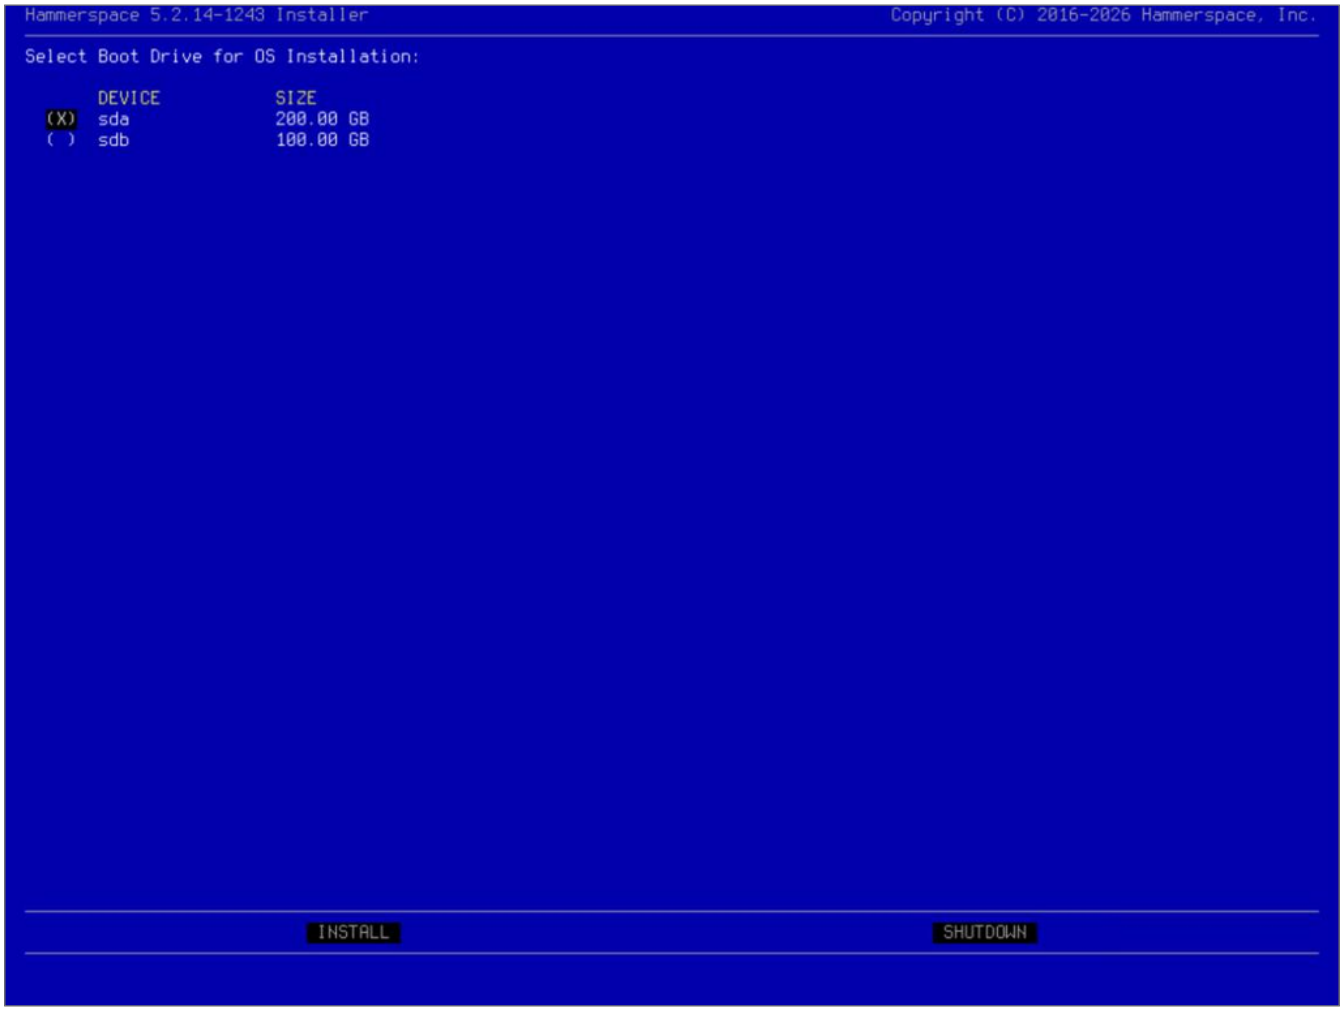

Figure 4. Boot drive selection screen

6. A confirmation dialog asks you to approve overwriting the drives before anything is written. Select **Yes** and press `Enter`.

This dialog defaults to **No**. It names the device the boot image will be written to, so it is also the last chance to catch the wrong disk having been selected on the previous

screen. Read the device name before confirming.

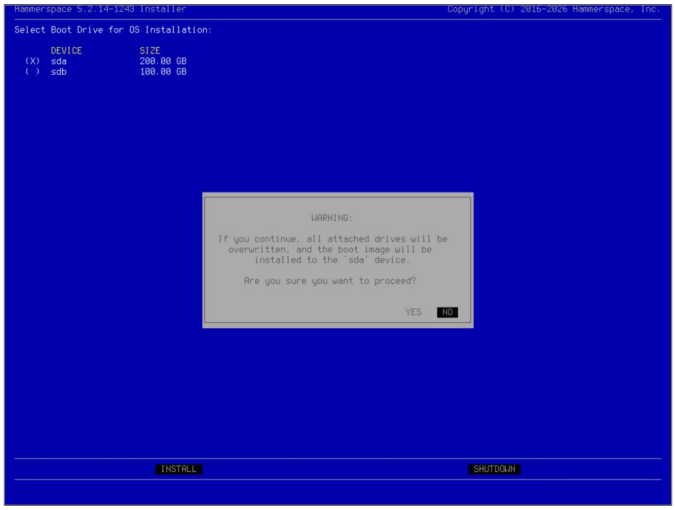

Figure 5. Overwrite confirmation dialog

7. After you confirm, the installer writes the boot image and then reboots using the installed image before any further input is required.

While the boot image is copied, the progress screen reports the size being copied, the target device and its capacity, the current throughput, and the bytes remaining. A single progress bar runs once from 0 to 100 percent.

After the reboot, the server displays **Hammerspace system booting…** for around 90 seconds before the next screen appears. This pause is expected and is not a failure.

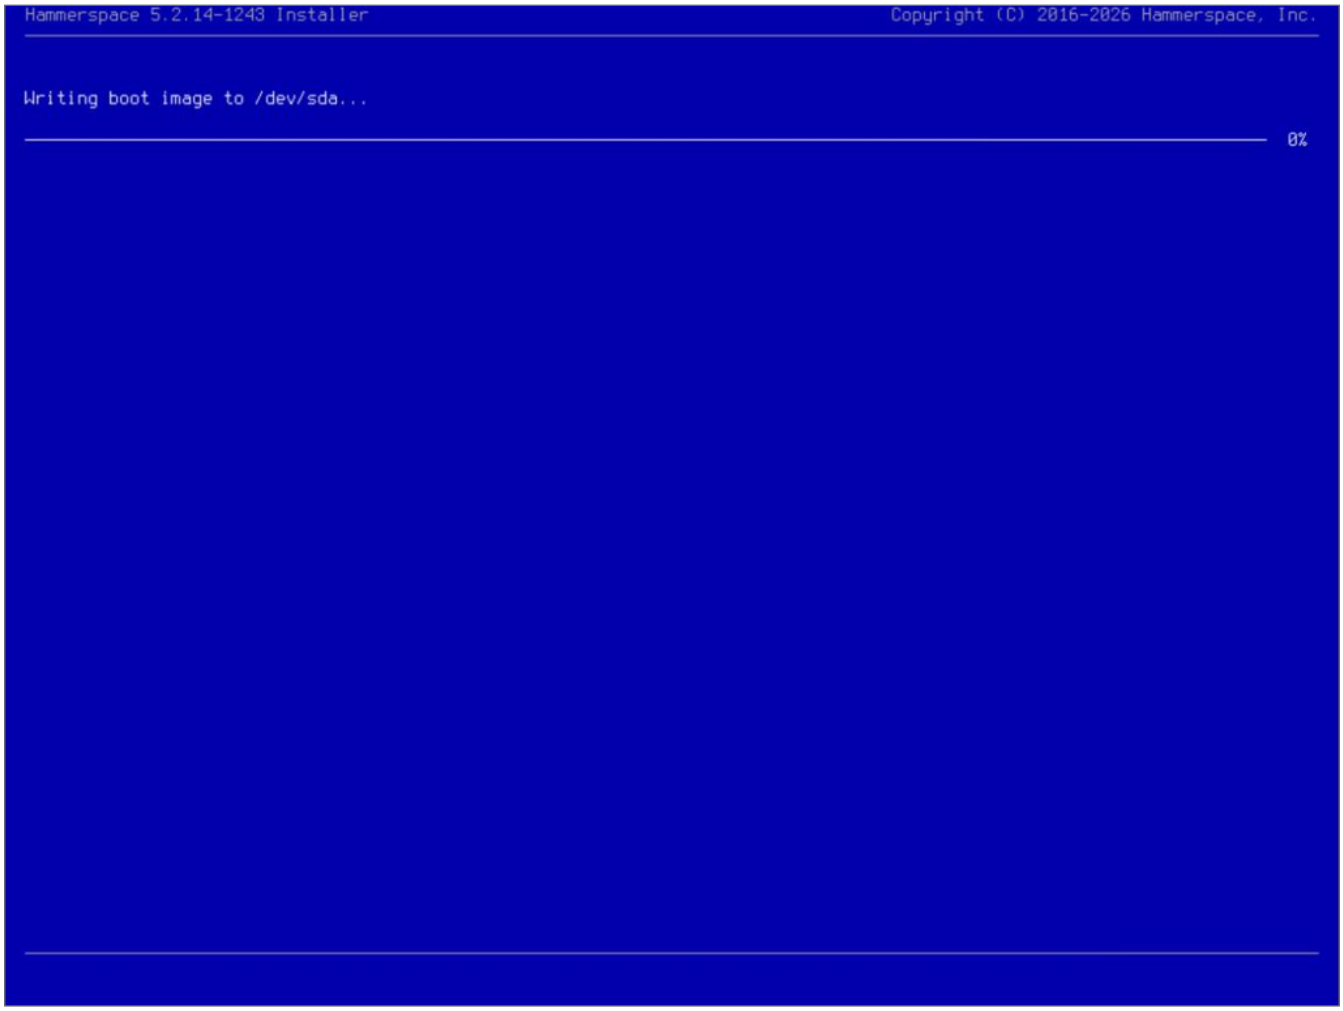

Figure 6. Installation progress screen

8. Review the End User License Agreement (EULA), then select **Accept**.

The EULA is also available online on the Hammerspace End User License Agreement page of the Hammerspace website.

Figure 7. EULA acceptance screen

9. Select the **Anvil** product type, then continue configuring the installation. The Virtual Hardware Validation screen displays.

On a fresh installation the installer arrives at this screen with **Anvil** selected. Where a previous Hammerspace configuration is present on the drives, it arrives with that configuration's product type selected instead. Confirm that the selection is the one you intend before continuing.

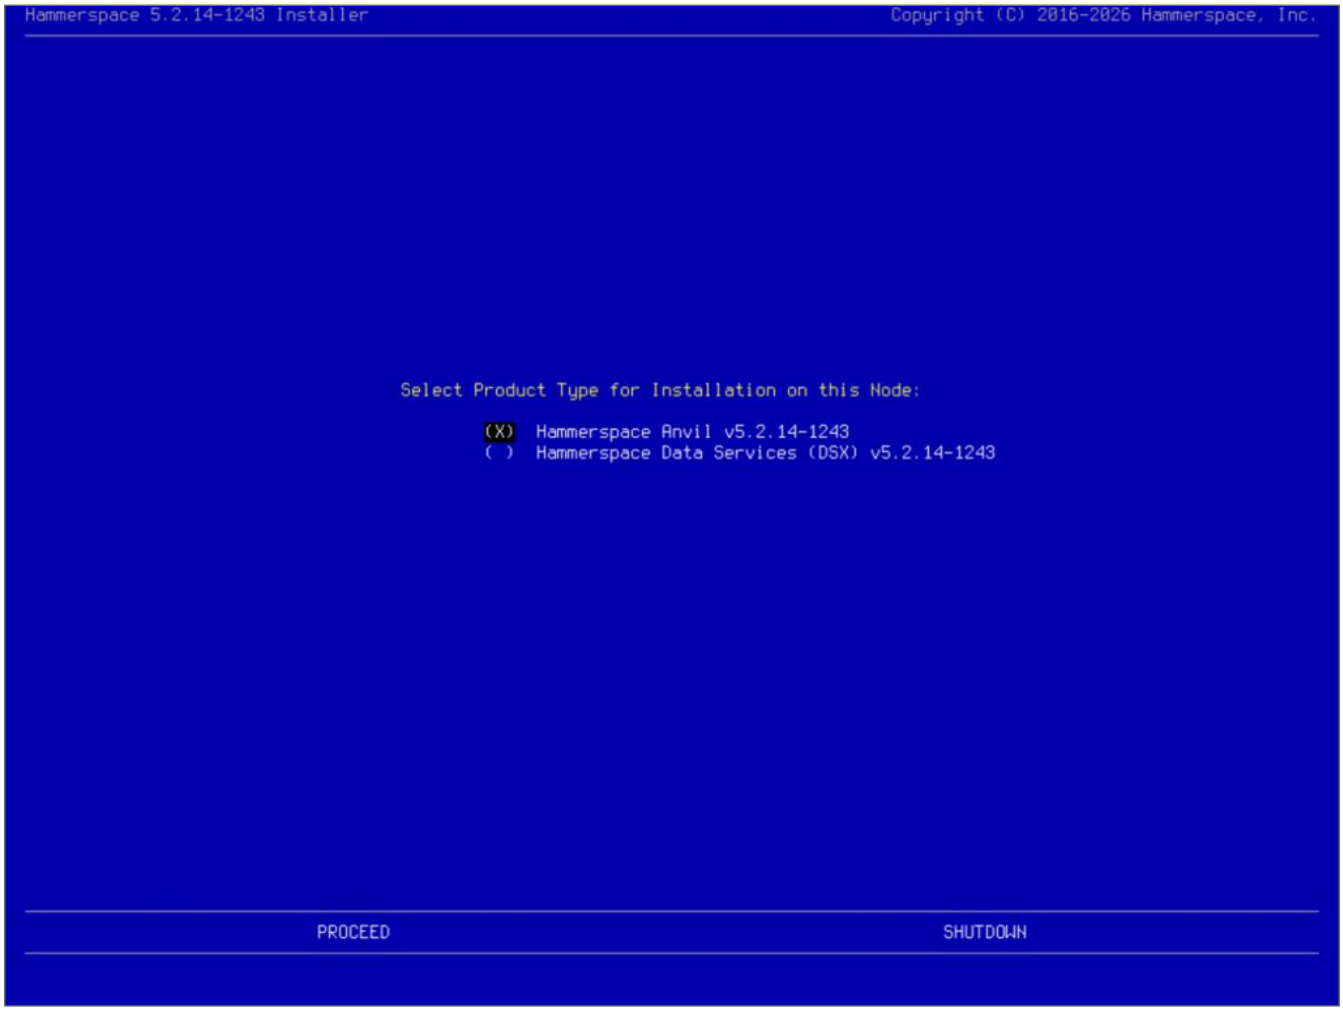

Figure 8. Select the product type

10. Confirm that the discovered hardware has been validated, then select **PROCEED**. In the example below, the installer has flagged cautions for the missing CPU and memory reservations, and for local storage that is smaller and fewer in number than recommended. It is strongly recommended that a production installation in a virtual environment always use CPU and memory reservations.

    Everything on this screen is reported as a caution, never as an error, and **PROCEED** stays available regardless. The installer will let you continue with an undersized system, so the decision about whether a caution is acceptable is yours to make here.

    Two cautions are worth expecting rather than troubleshooting:

    - **Memory allocated at exactly the minimum is still flagged.** A guest sees slightly less memory than the hypervisor allocates to it, so a VM built to the documented minimum presents just under the bar. Allocate somewhat above the minimum rather than exactly at it.
    - **Reservations are checked separately from size.** A correctly sized VM still shows CPU-reservation and memory-reservation cautions until the reservations themselves are set on the hypervisor.

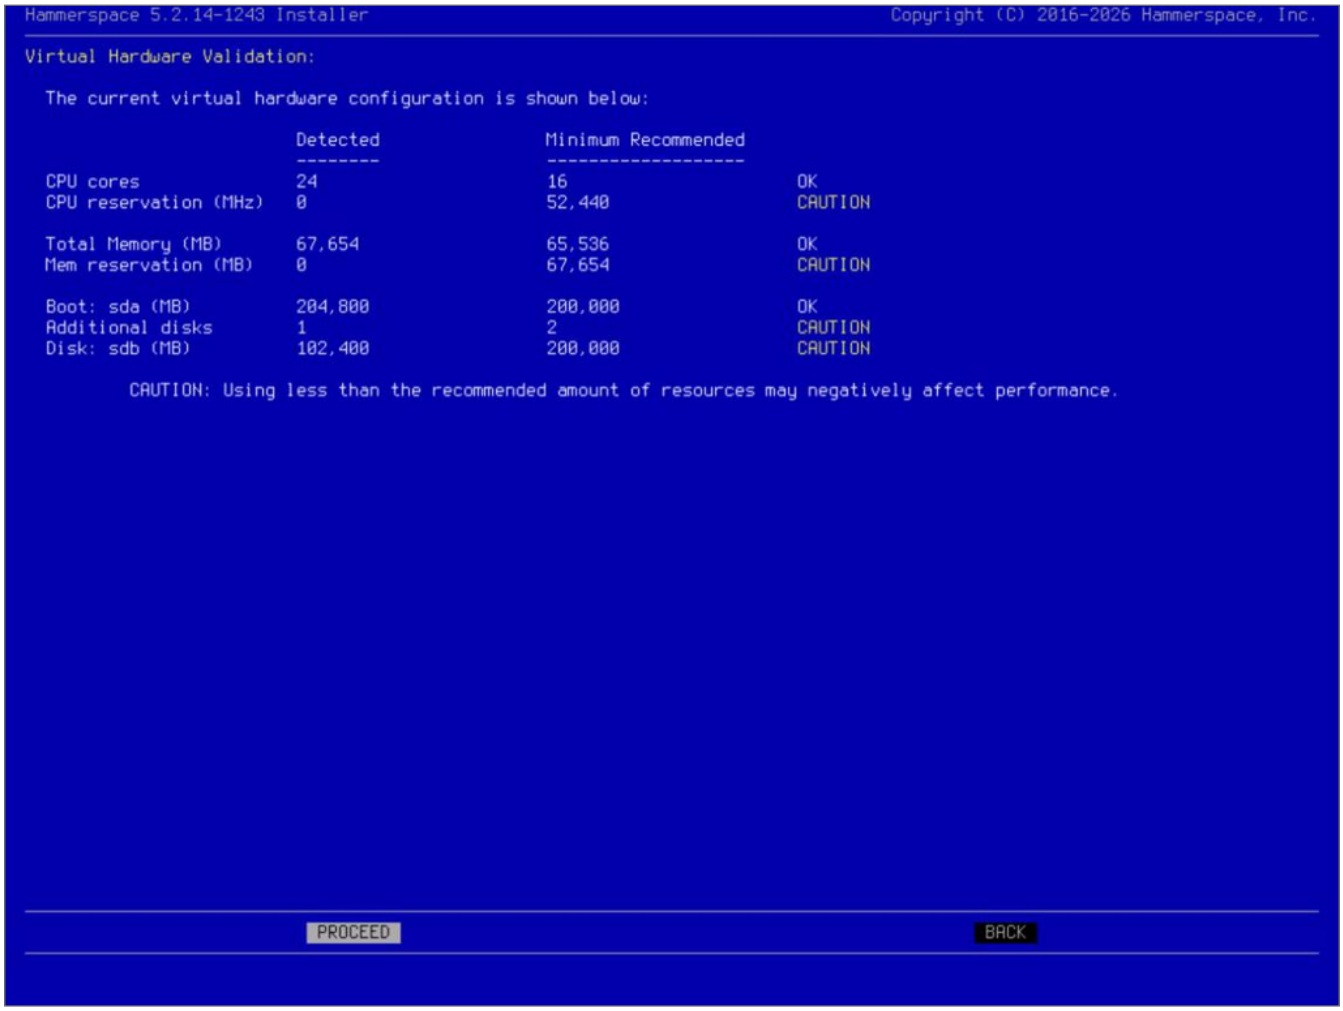

Figure 9. Resource validation check

11. If the installer finds an earlier Hammerspace installation on the drives, it asks whether to reuse the previous settings. Select **No** to configure the node from scratch.

This prompt appears only when the drives already hold a Hammerspace installation, so it is not seen on a new server. It defaults to **No**. Answering **Yes** pre-fills the configuration form from the previous installation, which is useful when you are deliberately rebuilding a node to the same specification and misleading when you are not.

In this release the prompt appears at the start of the configuration screen, after the networking steps that follow, rather than at this point in the sequence.

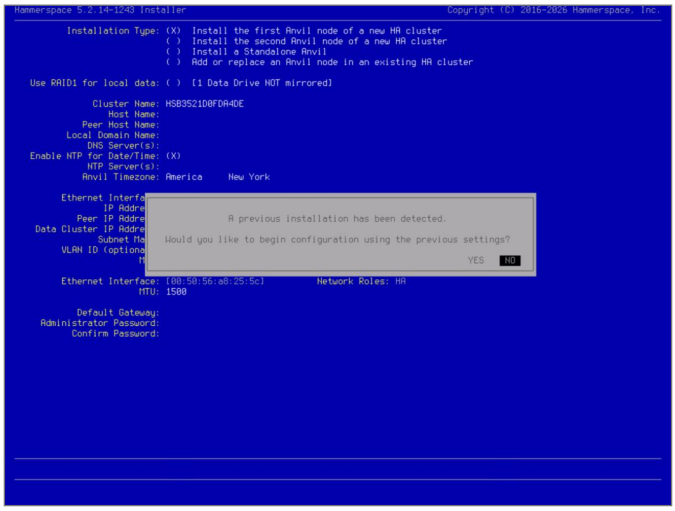

Figure 10. Previous installation detected

12. Configure networking by optionally bonding interfaces together and then assigning roles to the available network interfaces.

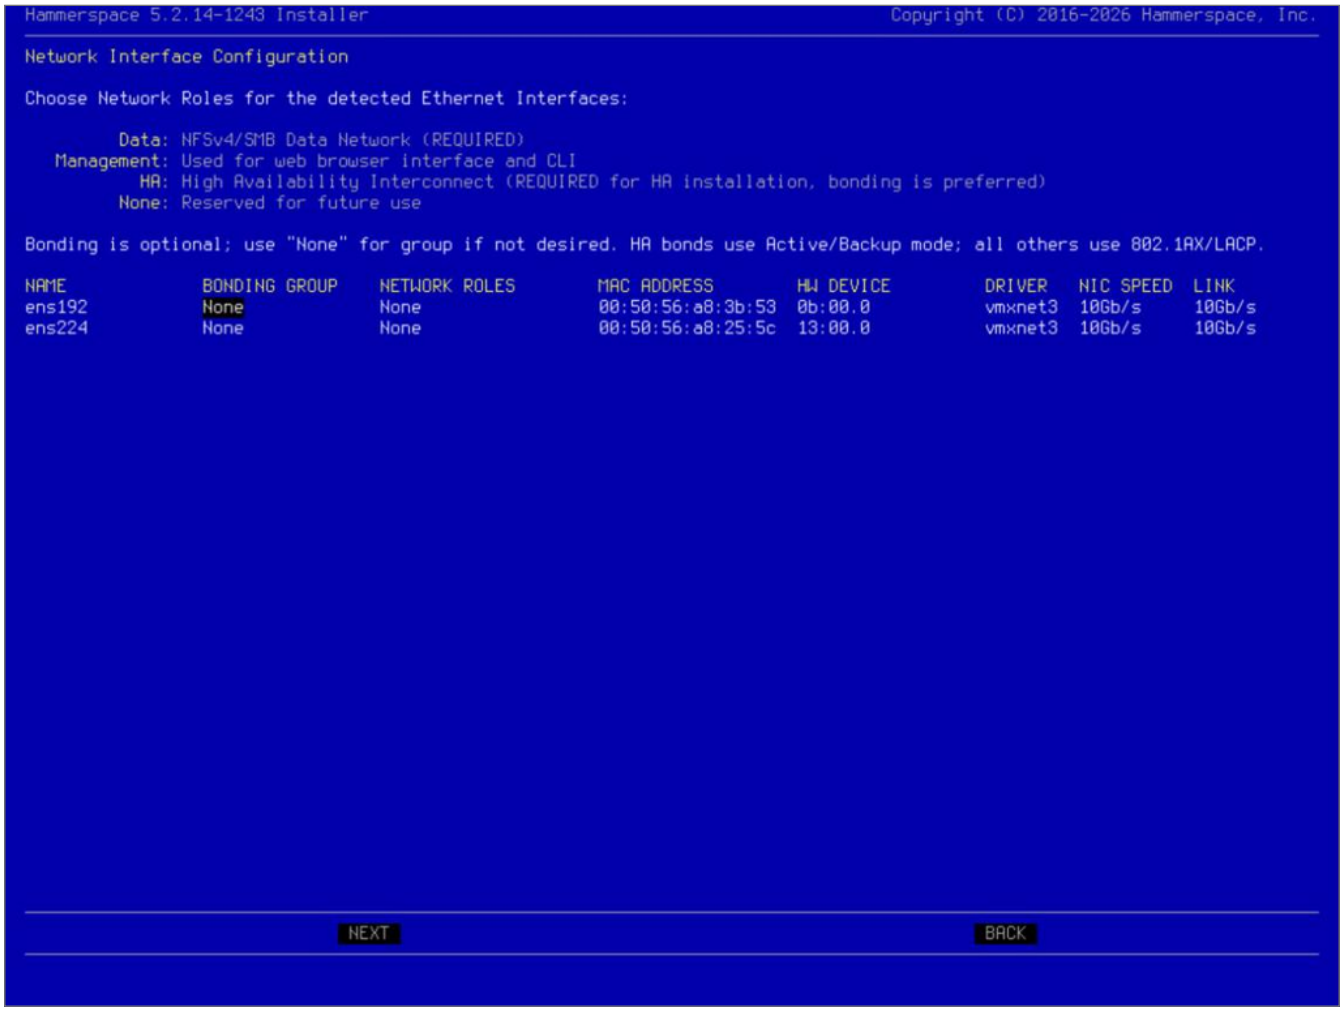

Figure 11. Configure networking

Hammerspace only supports LACP (Link Aggregation Control Protocol) bonding for non-HA roles. HA bonding is active-backup only.

Whether to bond is a decision about your environment, not a Hammerspace requirement:

- **Consider bonding** when you have two or more physical NICs cabled to switch ports that already support LACP, and you want the network to keep working if one link or cable fails. This matters most for the High Availability role: if the HA interconnect between Anvil nodes has no fallback path, losing that single link can prevent the nodes from communicating during a failure.
- **Selecting `None` is a reasonable choice** when you only have one interface available for a role, when redundancy is already handled below the guest (for example, NIC teaming on the hypervisor's virtual switch), or when your switch ports aren't configured for LACP yet and you'd rather complete the install now and revisit networking later.

Bonding is only offered when two or more interfaces could actually be bonded together. If each interface is attached to a different network, no bonding peer is available and the column stays empty.

This choice isn't final. Interface and bonding configuration can be changed after installation from the Admin CLI.

13. To create a bond, use the tab or arrow keys to select an entry in the **BONDING PEER** column and then press `Enter`. Then use the arrow keys to select an interface to bond with and press `Enter`.
14. To assign roles, use the tab or arrow keys to select an entry in the **NETWORK ROLES** column and then press `Enter`. Use the arrow keys to choose from the list of role combinations, and press `Enter` to assign the highlighted combination. When all interfaces have been assigned, select **NEXT** and press `Enter`.

Roles are assigned to an interface as a **predefined combination**, not as individual roles that you switch on and off. The list offers combinations such as `Data, Management` or `Data, Management, HA` as single choices.

The list also narrows as you work. Once one interface has taken a role, the combinations offered for the remaining interfaces drop that role, so the choices available on the second interface are fewer than those on the first. This is the installer preventing a conflicting assignment rather than a fault.

On a server with a single interface, the installer may arrive with a combination already assigned.

The network roles provide the following functionality:

See the Firewall Port Requirements section for port-level detail per role.

- **Data:** Used to connect all storage clients and servers (including DSX nodes). NFS and SMB data traffic flows across the data Network, as well as some Inter-node communication.
- **Management:** Used for management and monitoring tasks. This role is not used for file sharing, data movement, or HA cluster communications.
- **High Availability:** Used for peer-to-peer communications between the nodes in an HA deployment.

Select the interface on this network that is configured to be connected to the second node on a private interconnect. For bare-metal deployments, a back-to-back network cable connection is preferred (no switch needed).

All Anvil installations require at least one network interface with management, data, and high-availability roles. It is supported to share all roles on one interface, including the HA role. When sharing all roles on an interface, ensure the interface speed is at least 10 Gbit/s.

It is also possible to configure additional interfaces once the system is installed. Hammerspace supports floating IP addresses and VLAN assignment to interfaces. All these configuration settings are accessible in the Admin CLI interface after installation is complete.

After the installation has completed successfully, eject the CD-ROM or remove the ISO image from the virtual machine. This matters most when you have moved the CD/DVD device above the hard drive in the boot order, because the server will otherwise return to the installer on its next restart instead of booting the system you just installed.

2.2. Part 2: Installing the First Anvil Node in the HA Cluster

Continue the installation by selecting the installation type and providing the required details. For this example, we are installing the first Anvil of a new HA cluster.

1. Select the Installation type: **Install the first node of a new HA cluster**, then enter the other installation and configuration details on the page.

In the Installation type list, the arrow keys move the highlight but do not commit the choice. Press `Space` to select the highlighted installation type before moving on to the rest of the form.

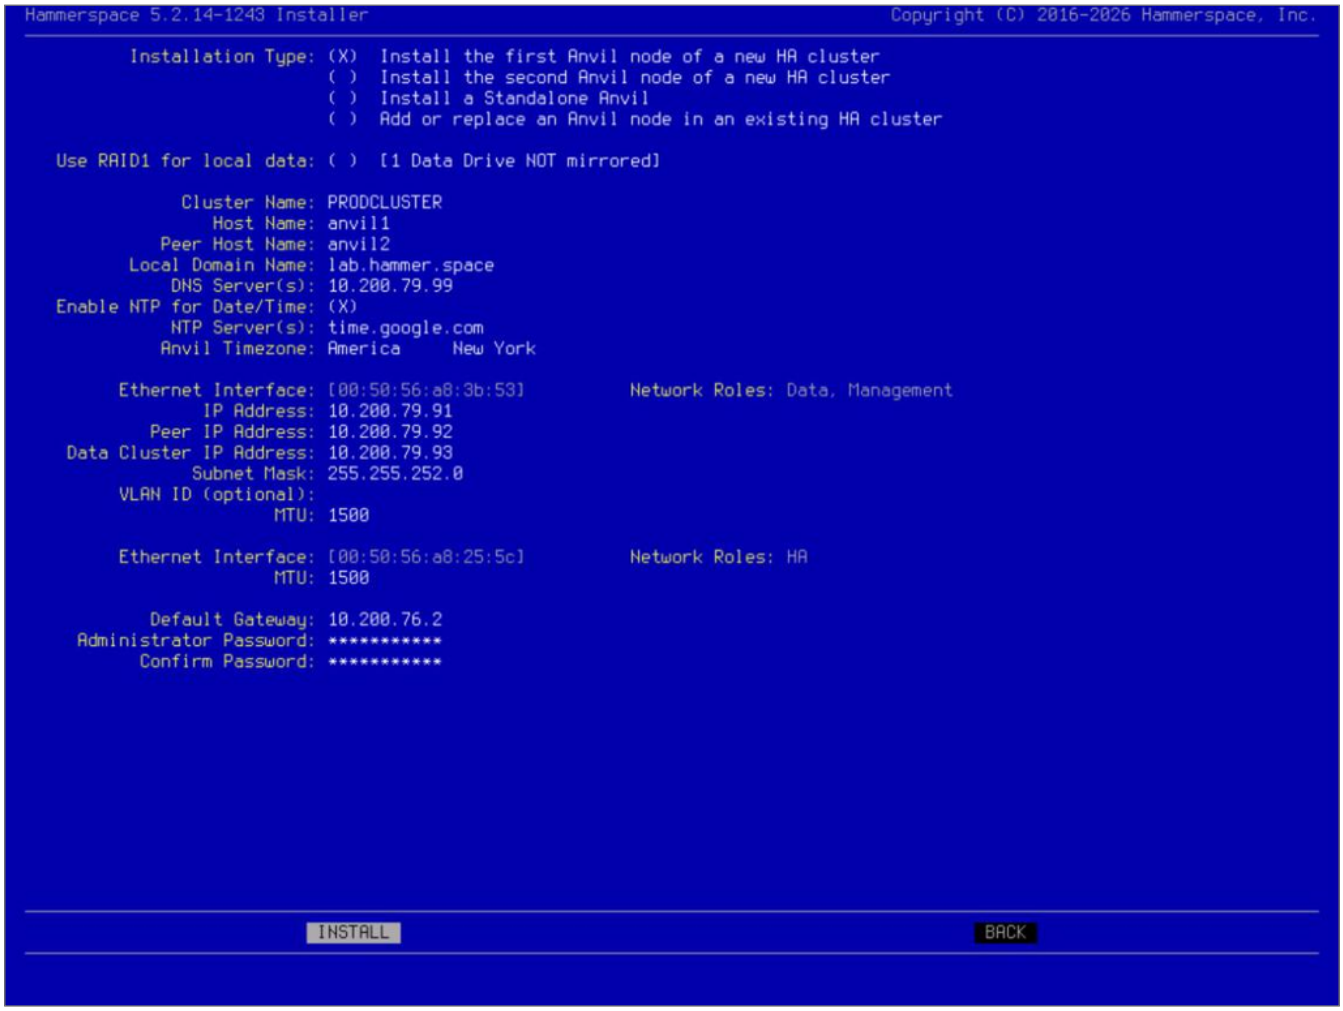

Figure 12. Installation type and first-node configuration details

Where:

- **Cluster Name**: This is the SMB Server Name and the Cluster Name. The installer supplies a generated name; replace it with a name that means something in your environment.

  The SMB Server Name can be changed when joining Active Directory.

- **Use RAID1 for local data**: Mirrors the local metadata drives. This option requires two or more data drives. With a single data drive the field reports that the drive is not mirrored and cannot be selected.
- **Host Name**: Enter the host name of the first node.
- **Peer Host Name**: Enter the hostname of the second node.
- **Local Domain Name**: Enter the name of the local domain.
- **DNS Server(s):** Use comma-separated entries for multiple DNS servers. Required for NTP and Active Directory support. The Anvil and DSX nodes communicate with each other using only IP addresses; however, any communication outside of the nodes may require DNS.

  The installer does not test whether the DNS servers are reachable while you are filling in the form. That test happens later, after the configuration is committed, during the network-start phase. An unreachable DNS server therefore surfaces partway through

configuration rather than at the field itself.

- **NTP Server(s):** Use comma-separated entries for multiple NTP servers. This is required for Active Directory support and strongly recommended for all other configurations. The Anvil servers will act as NTP servers for the DSX nodes.

  

  Identifying an NTP server by hostname requires a DNS server to be configured. To use NTP without DNS, enter the NTP server's IP address instead.

- **Timezone:** Set the time zone for the installation.
- **Ethernet Interface**: For the Management and Data Network Interfaces, configure the following:
- **IP Address**
- **Peer IP Address**
- **Data Cluster IP Address**: The floating IP address that clients and the Management GUI use to reach the cluster. This address automatically moves to whichever Anvil node is currently active, so it stays reachable regardless of which node is primary at any given time.
- **Subnet Mask (or Prefix)**

  

  Data and management interfaces must be on different subnets if configured at separate interfaces.

- **VLAN ID (optional):** Enter a value here only if VLAN tagging must be performed by the Anvil or DSX node.
- **MTU size:** The default MTU size is 1500 bytes. To use Jumbo frames, change the MTU size to match your network infrastructure. Typically, the correct value is 9000.
- **Default Gateway:** This is the IP address of the Default Gateway.
- **Administrator password:** Sets the password for the **admin** user. The password must be 8 to 20 characters long and must contain at least one upper-case letter, one lower-case letter, one number, and one non-alphanumeric character. Spaces are not allowed; any other printable character may be used.

  

  The interface carrying the High Availability role is configured differently from the data and management interfaces: it has an MTU setting only. It does not ask for an IP address, peer IP address, cluster IP address, or subnet mask, because the HA interconnect addresses itself automatically. An empty-looking HA block is correct, not an incomplete form.

2. Review the configuration for accuracy. If all settings are correct, select **INSTALL** and press `Enter` to continue.

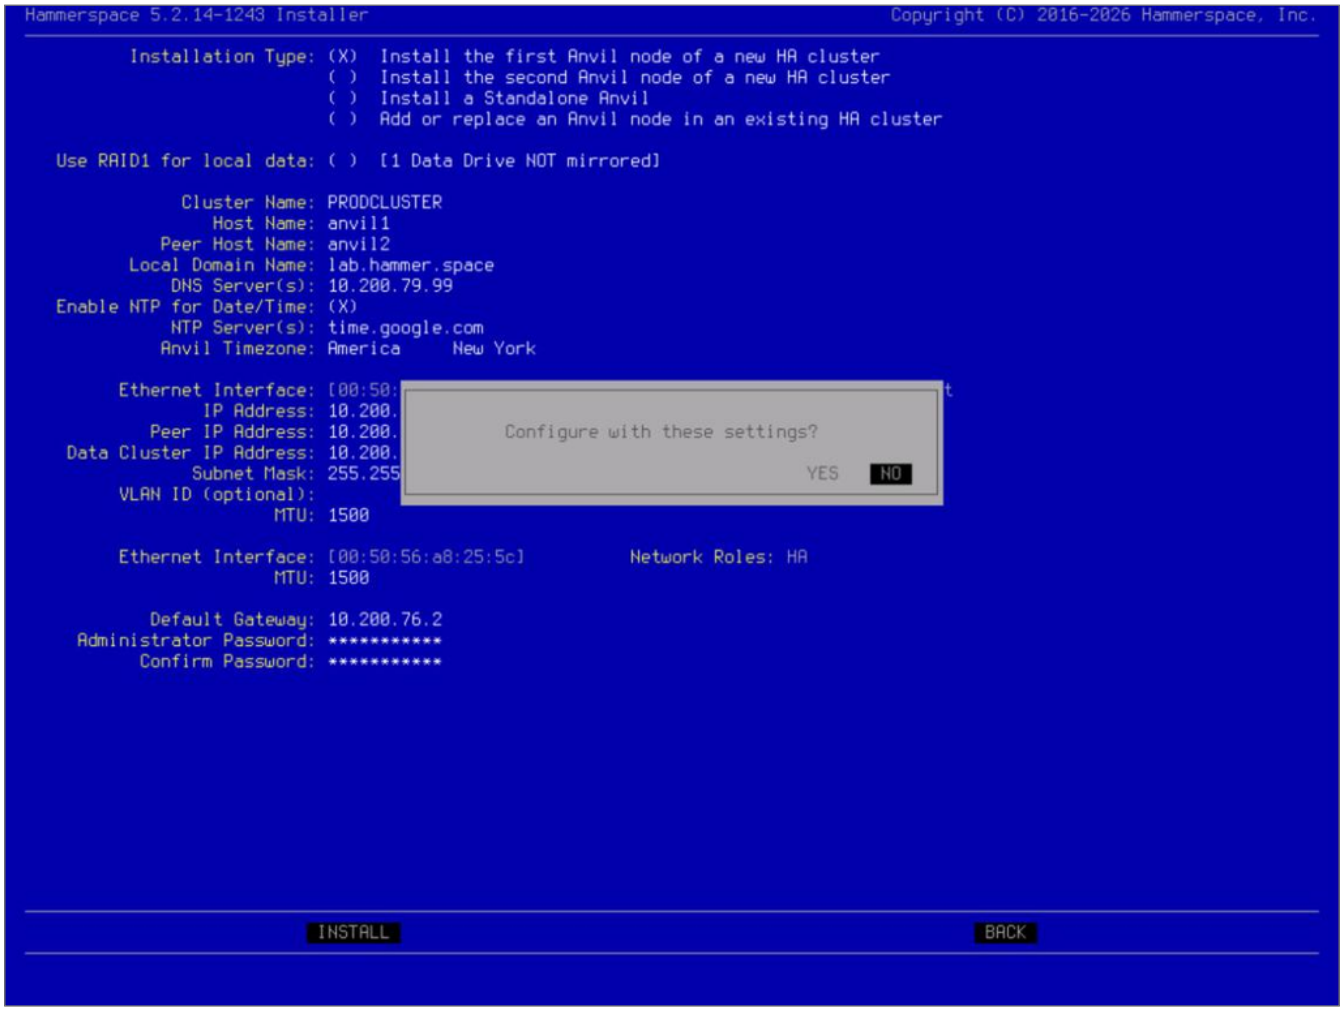

Figure 13. Confirming the configuration before installing

3. Select **Yes** and press `Enter` to confirm.

Once the installation of the first node completes, the node waits for its peer and displays the address that the second node needs.

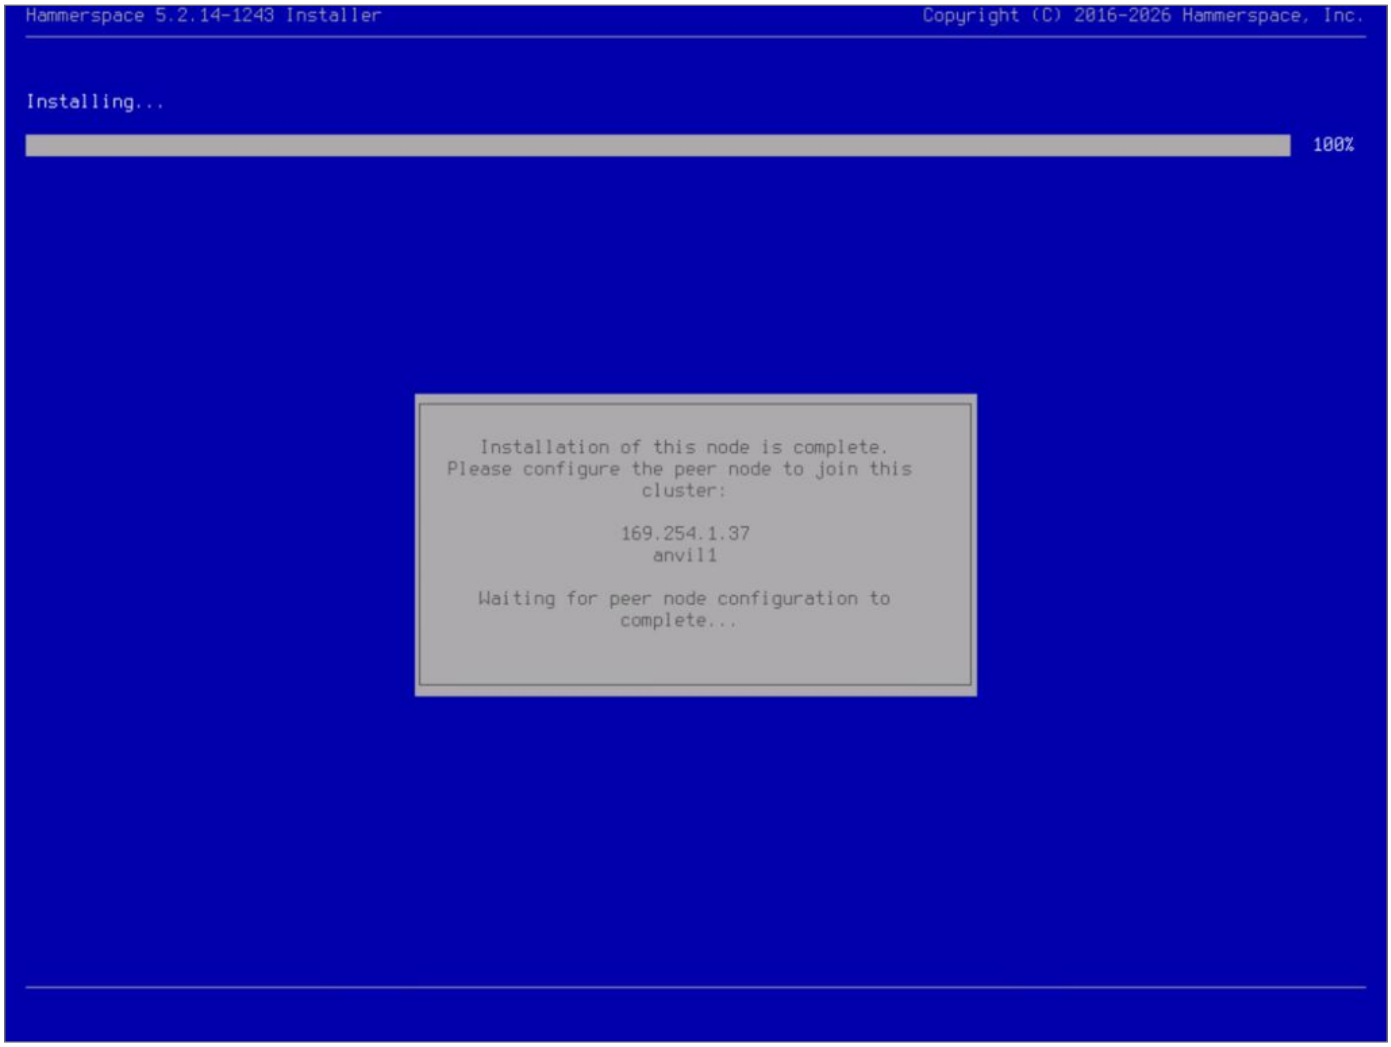

Figure 14. Waiting for the peer node

Make a note of the address displayed on this screen. It is a link-local address on the private interconnect, in the `169.254.x.x` range, rather than the data-network address you assigned to the node. You are asked to confirm it when installing the second Anvil node, and the second node's discovery may offer several similar addresses without labelling which is which, so the exact value matters.

The installation is not complete, and the cluster cannot be used until the second Anvil node of the cluster has been successfully added.

Continue to Part 3: Installing Second Anvil Node in HA Cluster below to install the second Anvil node.

2.3. Part 3: Installing Second Anvil Node in HA Cluster

Complete the following steps to create the second node on the high-availability cluster:

1. Repeat all of Part 1: Initial Part of the Installation (Steps 1–14) on the second node, through the network role assignment.
2. Next, select the Installation type: **Install the second Anvil node of a new HA cluster**.

As on the first node, the arrow keys move the highlight but do not commit the choice. Press `Space` to select this installation type.

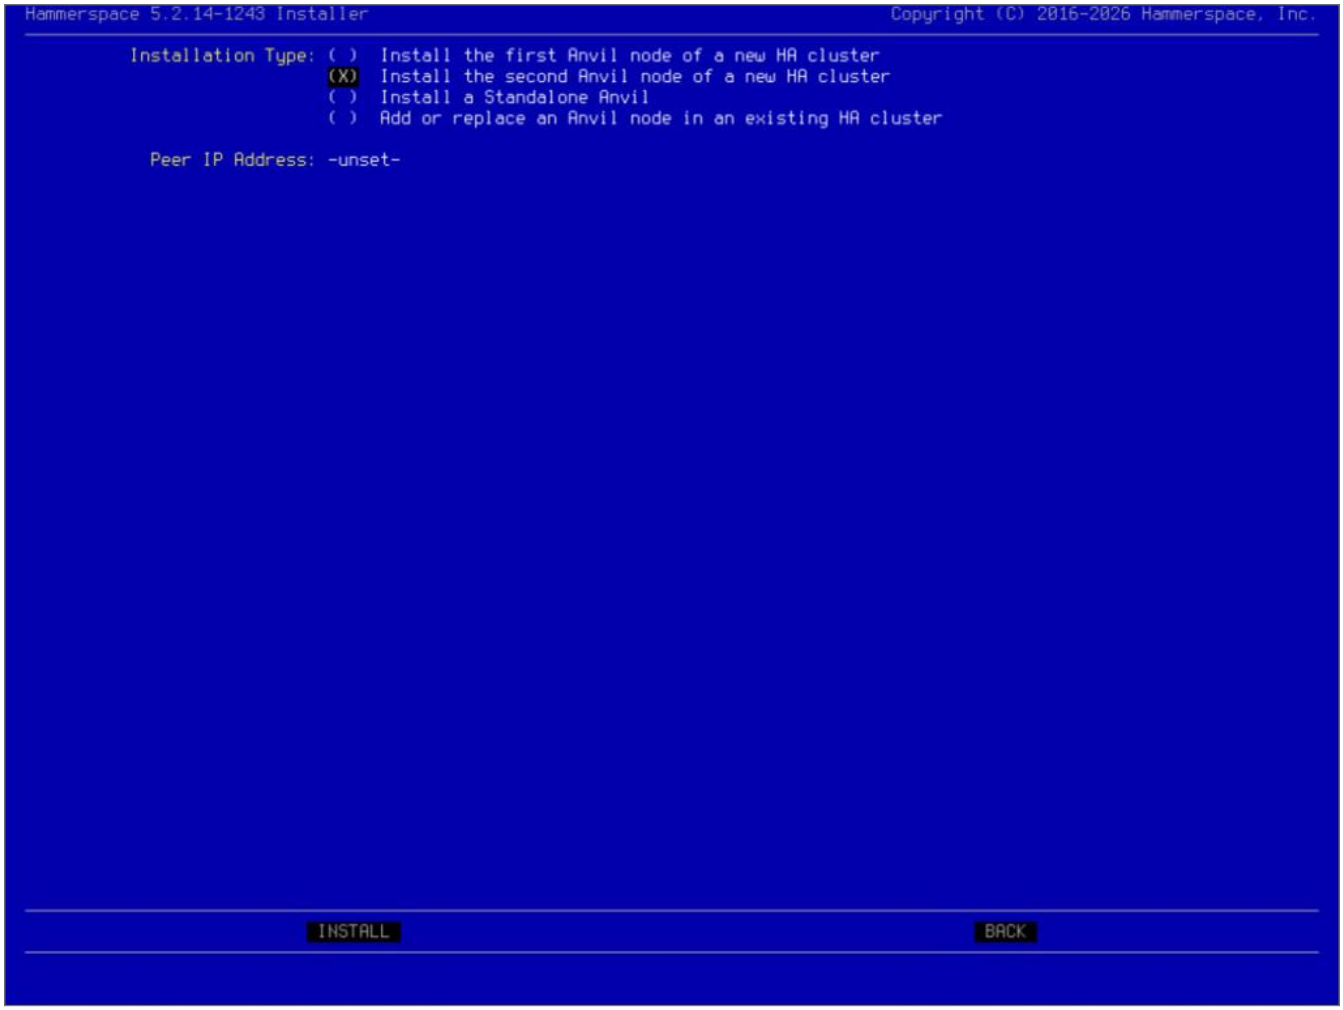

Figure 15. Selecting the second-node installation type

Selecting this installation type reduces the form to a single field, **Peer IP Address**. The second node does not ask for a host name, IP address, subnet mask, gateway, DNS or NTP servers, timezone, or administrator password. It inherits all of them from the first node across the private interconnect once the two nodes have found each other. A form with only one field on it is the expected result, not a screen that has failed to load.

3. Press Tab on your keyboard to move the cursor to **Peer IP Address** and press `Enter`. The peer IP address will be automatically discovered using the private interconnect.

Discovery may return several addresses, listed without labels to identify which node each belongs to. Select the exact address that the first node displayed on its **waiting for peer** screen. If you have multiple clusters sharing the same heartbeat network, this is also how you make sure you are joining the right one.

4. Once the Peer IP address is selected, select **Install** and press `Enter`. Always double-check that the peer IP address matches the one displayed on the console of the first node.

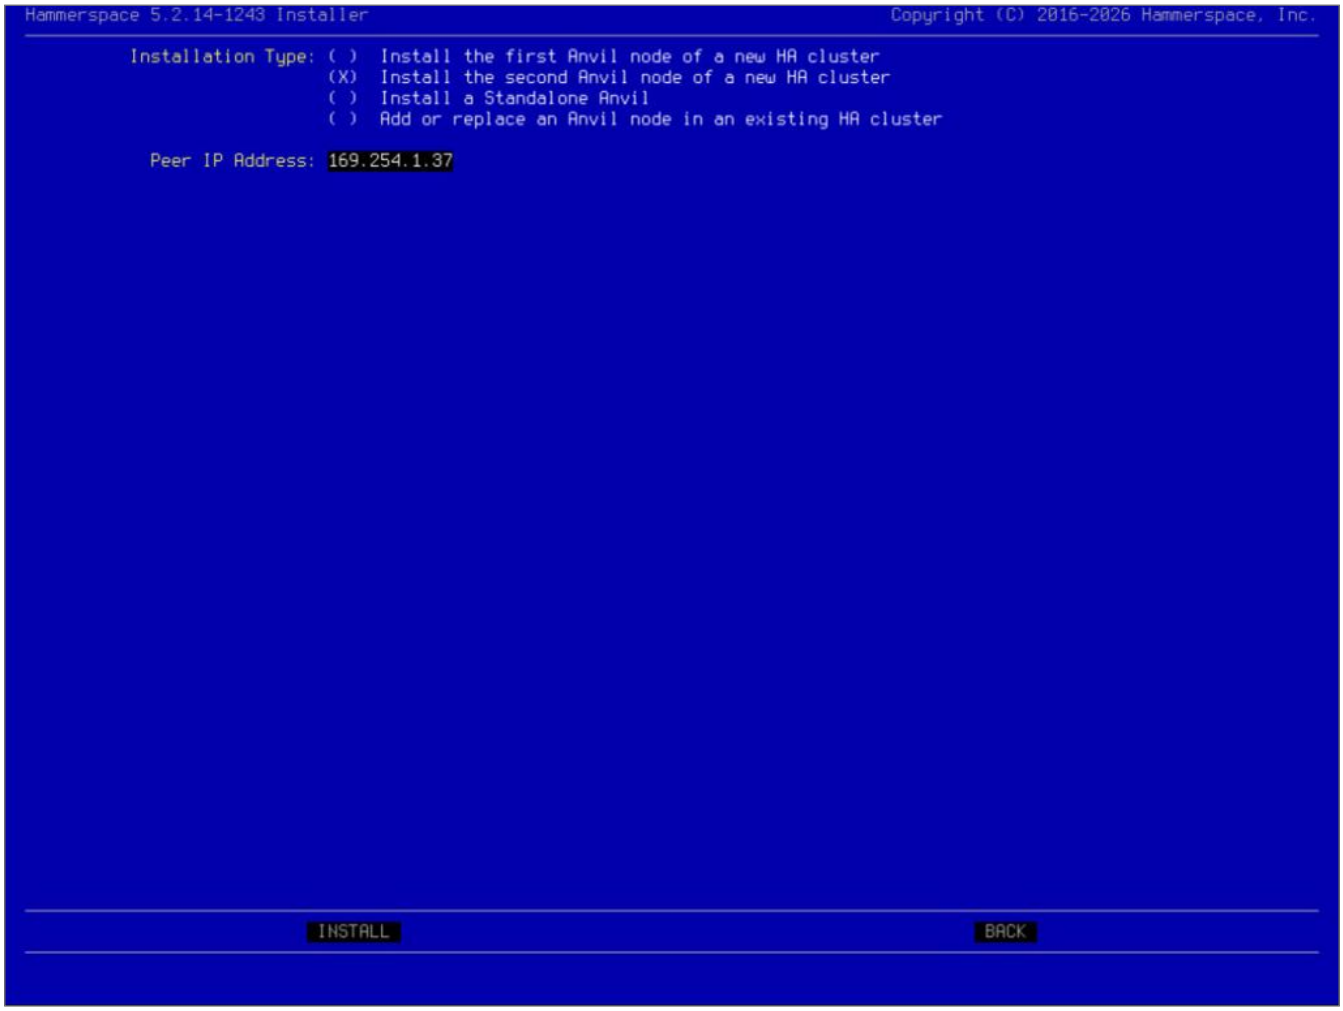

Figure 16. Peer address selected, ready to install

5. Select **Yes** and press `Enter`.
6. The HA cluster and network settings are automatically configured, and the installation starts.
7. The following screen is displayed after the installation has completed:

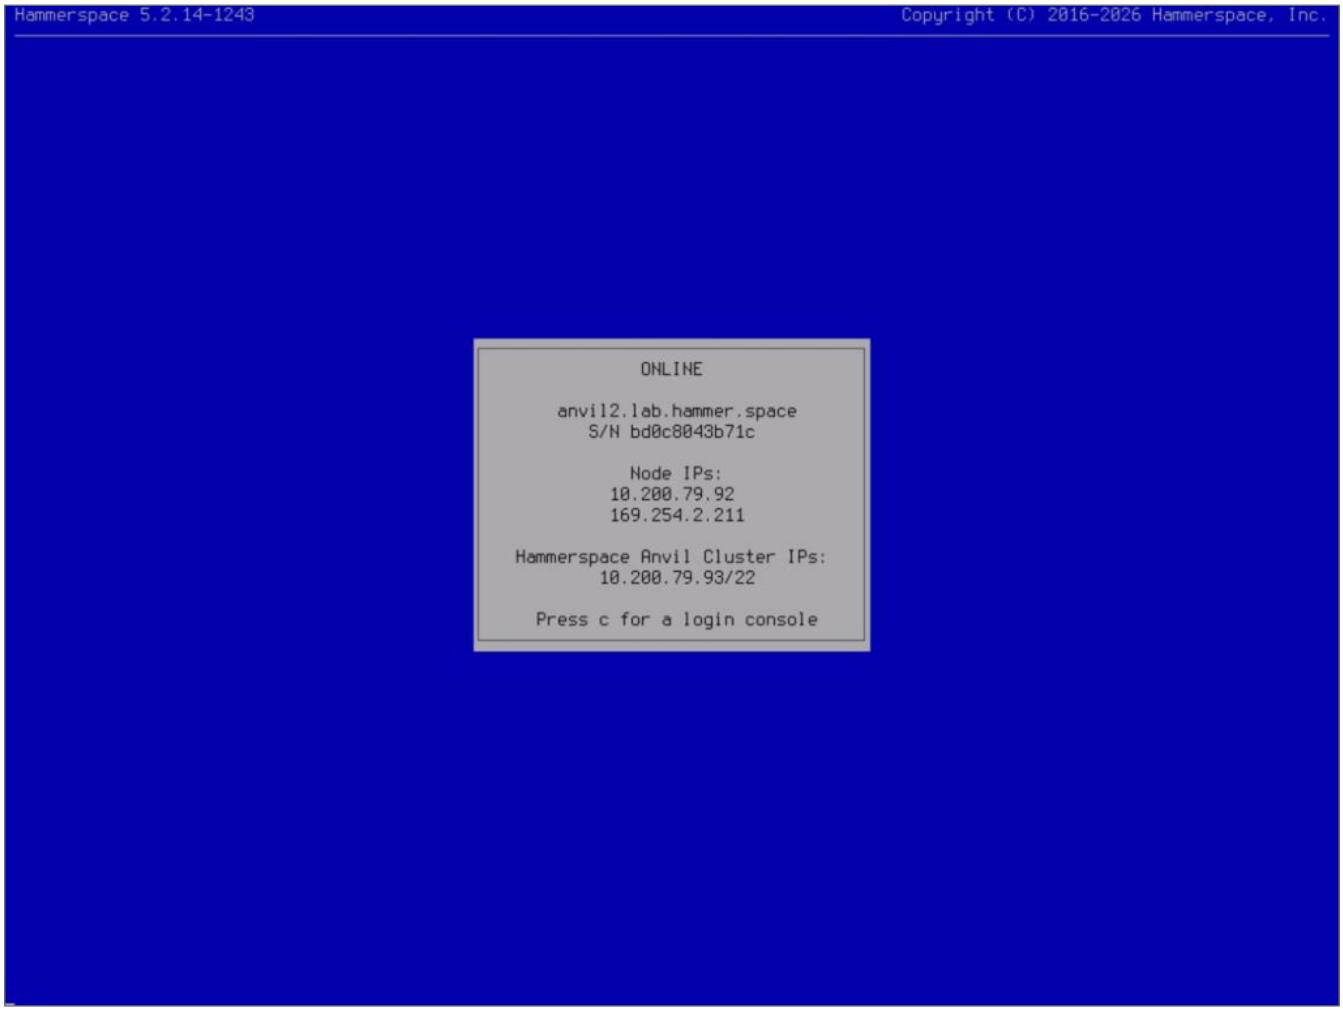

Figure 17. Second node installation complete

Both nodes should now report the same Hammerspace Anvil Cluster IP. If they do not, the two nodes have not joined the same cluster and the peer address selected in the previous step is the first thing to check.

8. Validate the installation by opening a web browser and going to the URL for the Management Graphical User Interface (Management GUI).

   The URL for the Management GUI is the Anvil Data Cluster IP — the floating address you assigned during the first node's installation, not the address of either individual node. In this example, it is `https://192.0.2.10:8443`.

The Management GUI answers on port 8443 and on the standard HTTPS port, so <code><a href="https://<cluster-ip&gt" class="bare">https://<cluster-ip&gt;</code> reaches it as well. A plain <code><a href="http://<cluster-ip&gt" class="bare">http://<cluster-ip&gt;</code> address is redirected to <code><a href="https://<cluster-ip>:8443" class="bare">https://<cluster-ip>:8443</code>.

9. When the web interface displays, log in using the **admin** credentials specified during installation to confirm successful installation.

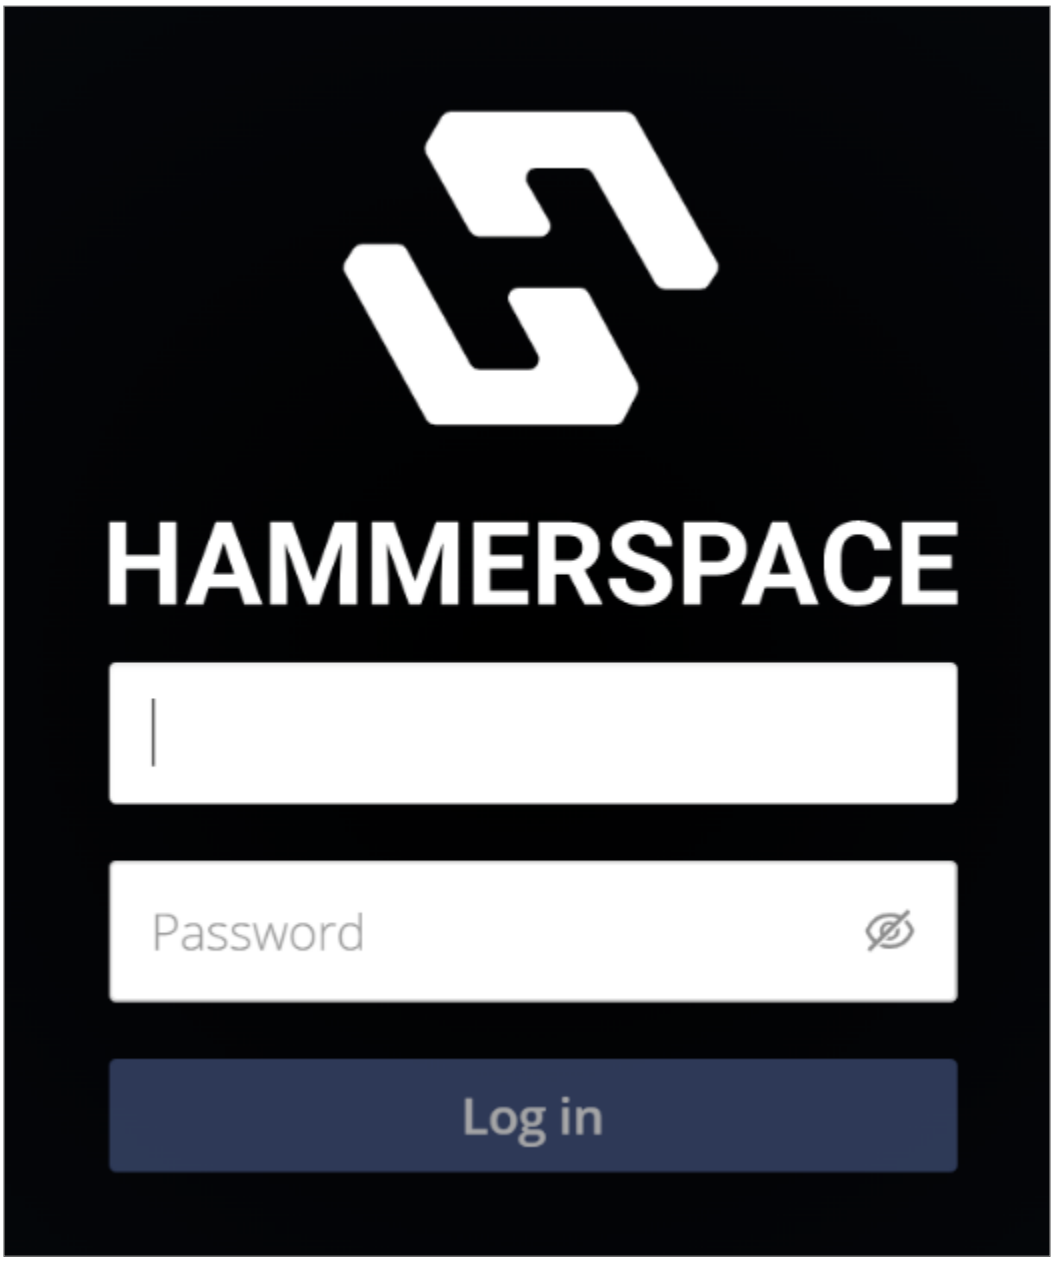

Figure 18. Logging in to the Management GUI

If you installed from an ISO image in a vSphere environment, you must disable time synchronization settings for each VM. See the time synchronization section for information.

2.4. Optional: Adding a Third Node to a Two-Node HA Cluster

Complete the following steps to add an optional third Anvil node to a two-node HA cluster:

1. Ensure the two-node cluster is installed and configured. The third node is added after the initial installation using the console installer.
2. Repeat all of Part 1: Initial Part of the Installation (Steps 1–14) on the new node, through the network role assignment.

3. On the **Installation type** configuration screen, select **Add or replace an Anvil node in an existing HA cluster.**

> Do *NOT* check the **Recover configuration** option. That workflow replaces a failed Anvil node.

The example below shows how to add a third Anvil node to a configuration with a dedicated HA network. In this configuration, the HA interface IP address must be manually specified.

To view the existing IP addresses that are in use by Hammerspace product nodes, navigate to **Administration > Network** in the Management GUI.

Pick an address that does not conflict with an existing address on the HA interconnect network.

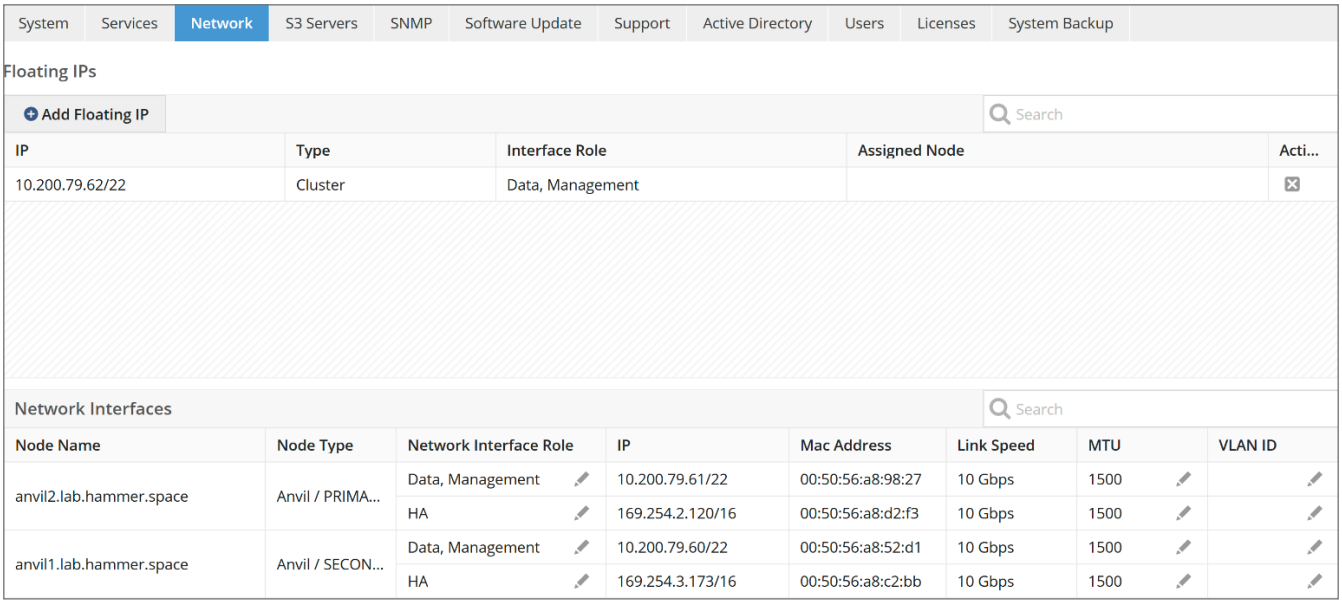

Fill in the configuration information in the console installer.

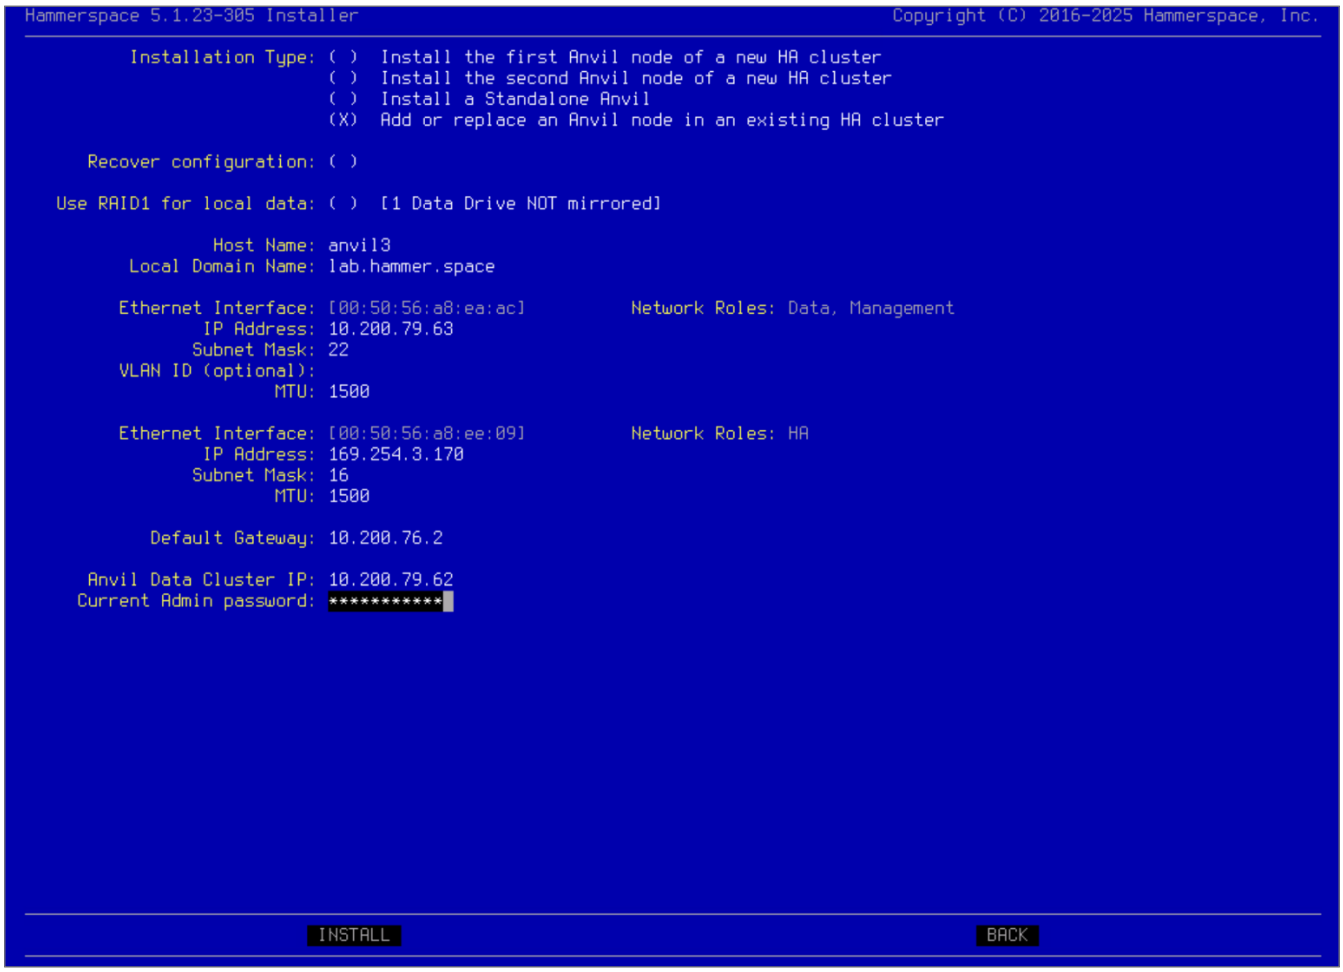

4. Press **Install,** and after a few minutes, the third HA Anvil will be available in the cluster.

If the cluster is configured with very large metadata drives, it will take several minutes for the metadata drives to finish their background synchronization.

5. Validate that the third node has successfully been added by logging into the Management GUI and navigating to **Administration > System**. In the example below, there are three Anvil nodes: one Primary and two Secondary.

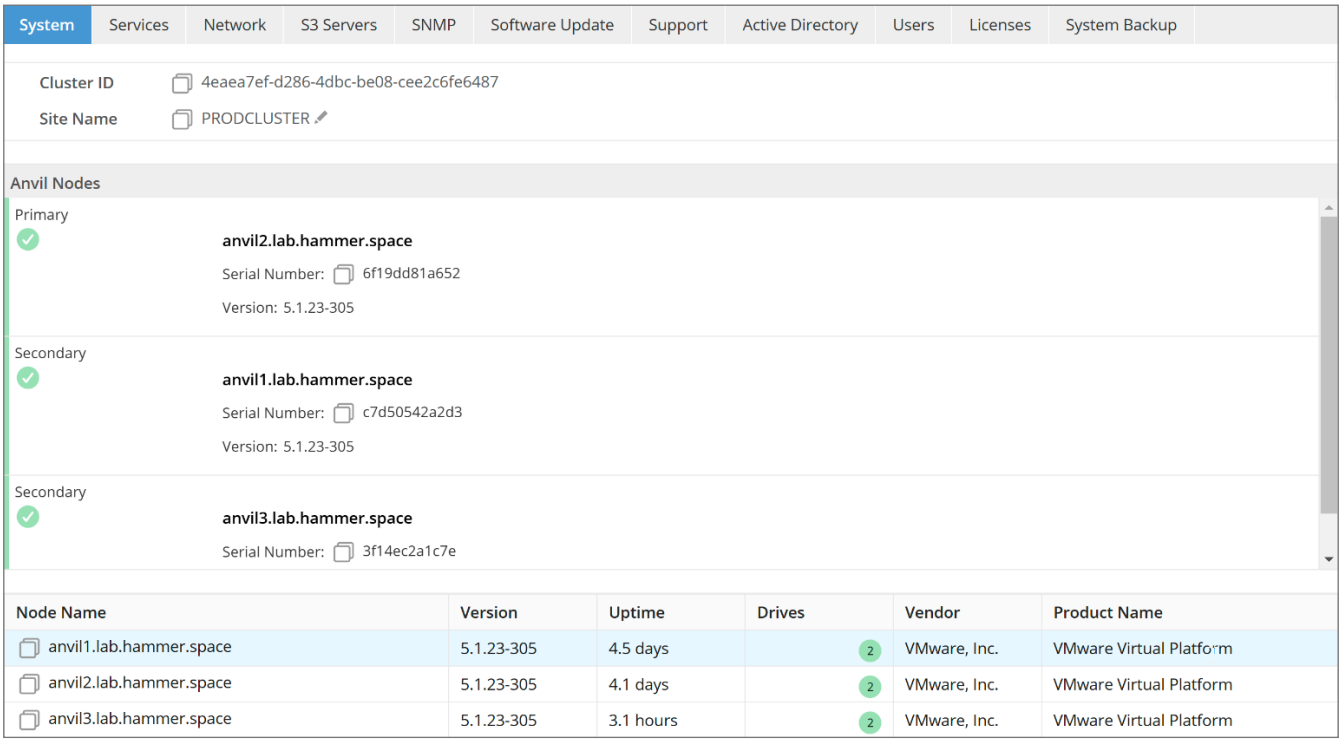

2.5. Recovering a Failed or Missing Anvil Node

Use this procedure when an Anvil node in an HA cluster must be reinstalled — for example, after a failed software update, hardware replacement, or an accidental reinstall.

2.5.1. Cancel Any Stuck Software Update

If the surviving Anvil still shows a software update in progress, cancel it from the Admin CLI on the surviving Anvil.

1. View the in-progress update task:

```
software-update-status
```

2. Cancel the task using its ID:

```
software-update-cancel --task-id <ID>
```

2.5.2. Reinstall the Failed Anvil

1. Install the same software build that the surviving Anvil is running on the replacement node.
2. During installation, choose the installation type **Add or replace an Anvil node in an existing HA cluster**.
3. Use the same network and identity settings the failed node had.
4. At the end of installation, the installer displays a CLI command to run on the surviving (primary) Anvil.

2.5.3. Rejoin the Cluster

Run the command provided by the installer on the surviving Anvil to rejoin the reinstalled node into the HA cluster, then verify cluster health: `node-list` should show both Anvils with **HW state** `OK` and **Node mode** `ONLINE`, and the **Metadata servers** line of `cluster-config` should list both, one with `Role: Primary`.

Chapter 3. Installing Data Services (DSX)

DSX is installed from an ISO image that contains the operating system and DSX software.

1. Confirm that the VM or bare-metal server meets the minimum system requirements.
2. Load the ISO image and power on the VM or bare-metal server.

As with the Anvil nodes, confirm that the CD/DVD device sits above the hard drive in the server's firmware boot order, and that a virtual CD/DVD drive is set to connect at power on. Otherwise the server boots from its own drives and the installer never appears.

3. Select the **Install Hammerspace** option and press `Enter`.

The boot menu defaults to **Boot from next device in boot order** and acts on that default after a short countdown. Press an arrow key as soon as the menu appears to stop the countdown before making your selection.

4. The installation program discovers local storage devices. Select the boot device for installation, then click **INSTALL** to continue.

Boot device capacity should be at least 100GB. See the DSX requirements table for bare-metal sizing guidance.

Choose the boot device by its size rather than accepting whichever device the installer arrives with selected. On a DSX node with a separate data drive this matters more than on an Anvil, because the pre-selected device is sometimes the larger data drive rather than the boot drive.

Selecting a device other than the first one listed adds a caution that the installation may fail to boot unless the firmware can boot from that device. The caution names the device concerned. This is a general warning rather than a detected fault.

All discovered block devices will be overwritten during the installation. Any existing data on these devices will be erased.

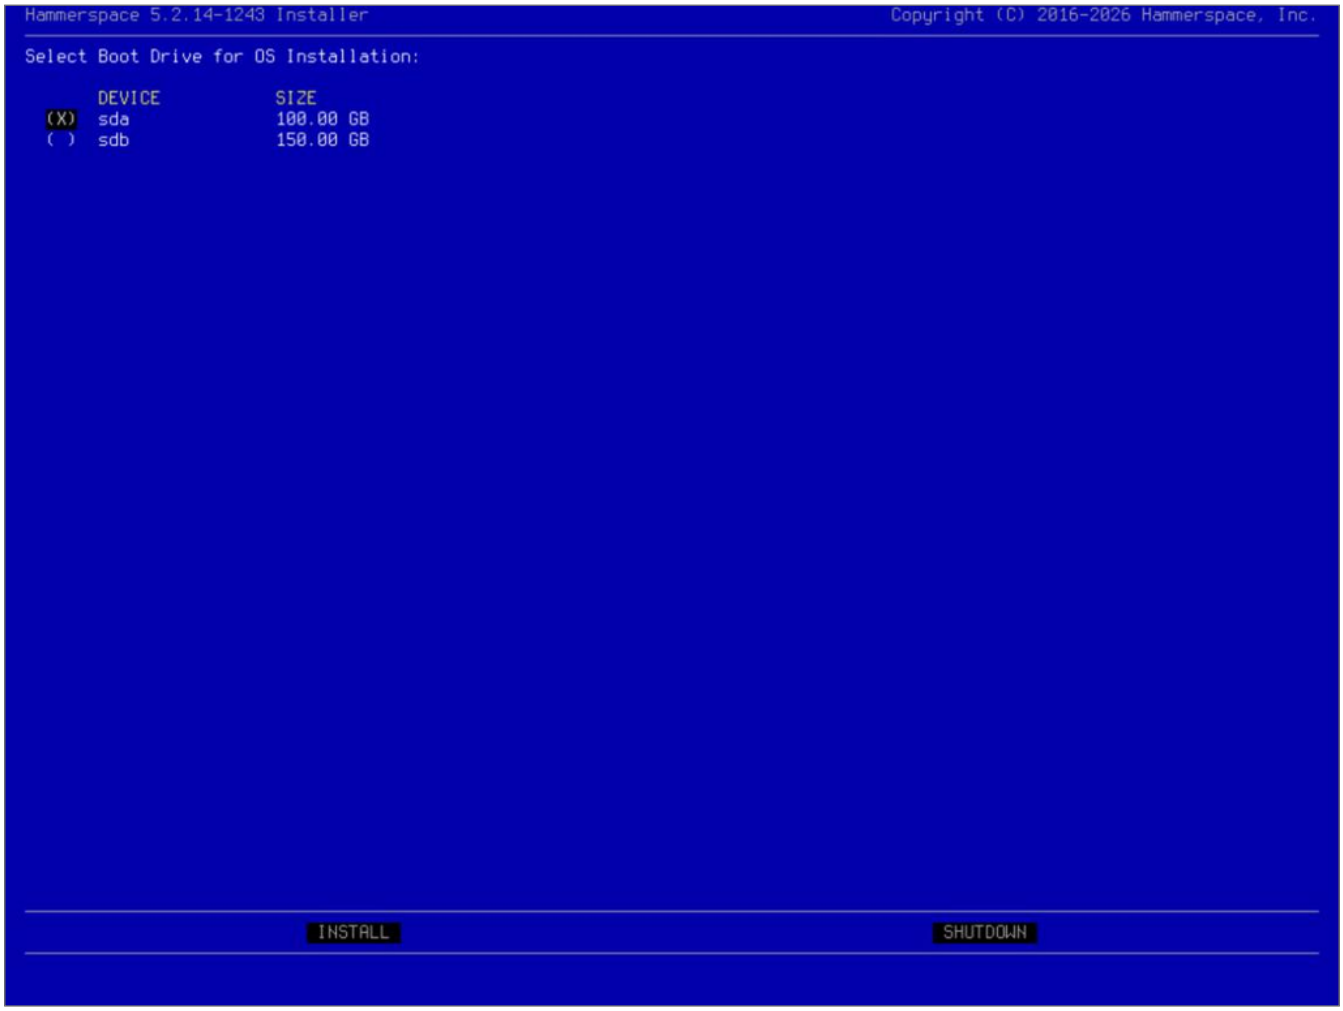

Figure 19. Boot drive selection

5. A confirmation dialog asks you to approve overwriting the drives. It names the device the boot image will be written to and defaults to **No**. Check the device name, then select **Yes** and press `Enter`.
6. Review the End User License Agreement (EULA), then select **Accept** and press `Enter`. The EULA is the same as that which was seen in the Anvil installation.
7. Select for Product type, **Hammerspace Data Services (DSX)**.

The installer usually arrives at this screen with a product type already selected, based on the resources it detects. Confirm the selection before continuing.

8. The Virtual Hardware Validation screen displays. Confirm that the discovered hardware is validated, select **Proceed**, and press `Enter`.

As on the Anvil nodes, undersized resources are reported as cautions rather than errors, and **Proceed** remains available. Allocate memory somewhat above the documented minimum rather than exactly at it: a guest sees slightly less memory than the hypervisor allocates, so a VM built to the exact minimum is still flagged.

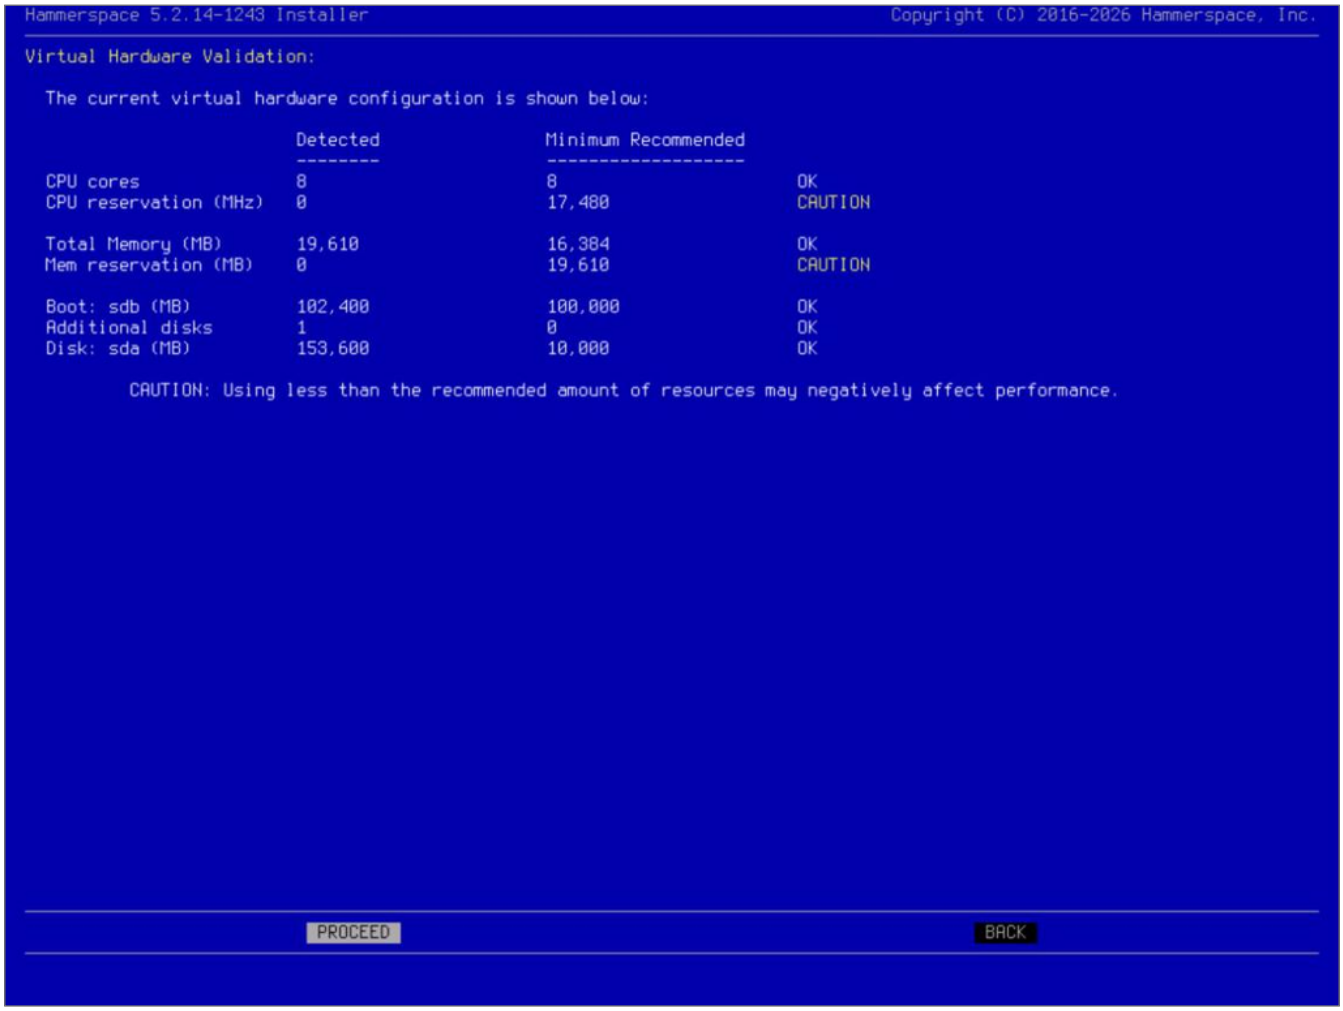

Figure 20. Virtual hardware validation

9. Configure networking by optionally bonding interfaces together and then assigning roles to the available network interfaces.
10. To create a bond, use the tab or arrow keys to select an entry in the **BONDING PEER** column and then press `Enter`. Then use the arrow keys to select an interface to bond with and press `Enter`.
11. To assign roles, use the tab or arrow keys to select an entry in the **NETWORK ROLES** column and then press `Enter`. Use the arrow keys to choose from the list of role combinations, and press `Enter` to assign the highlighted combination. When all interfaces have been assigned, select **NEXT** and press `Enter`.

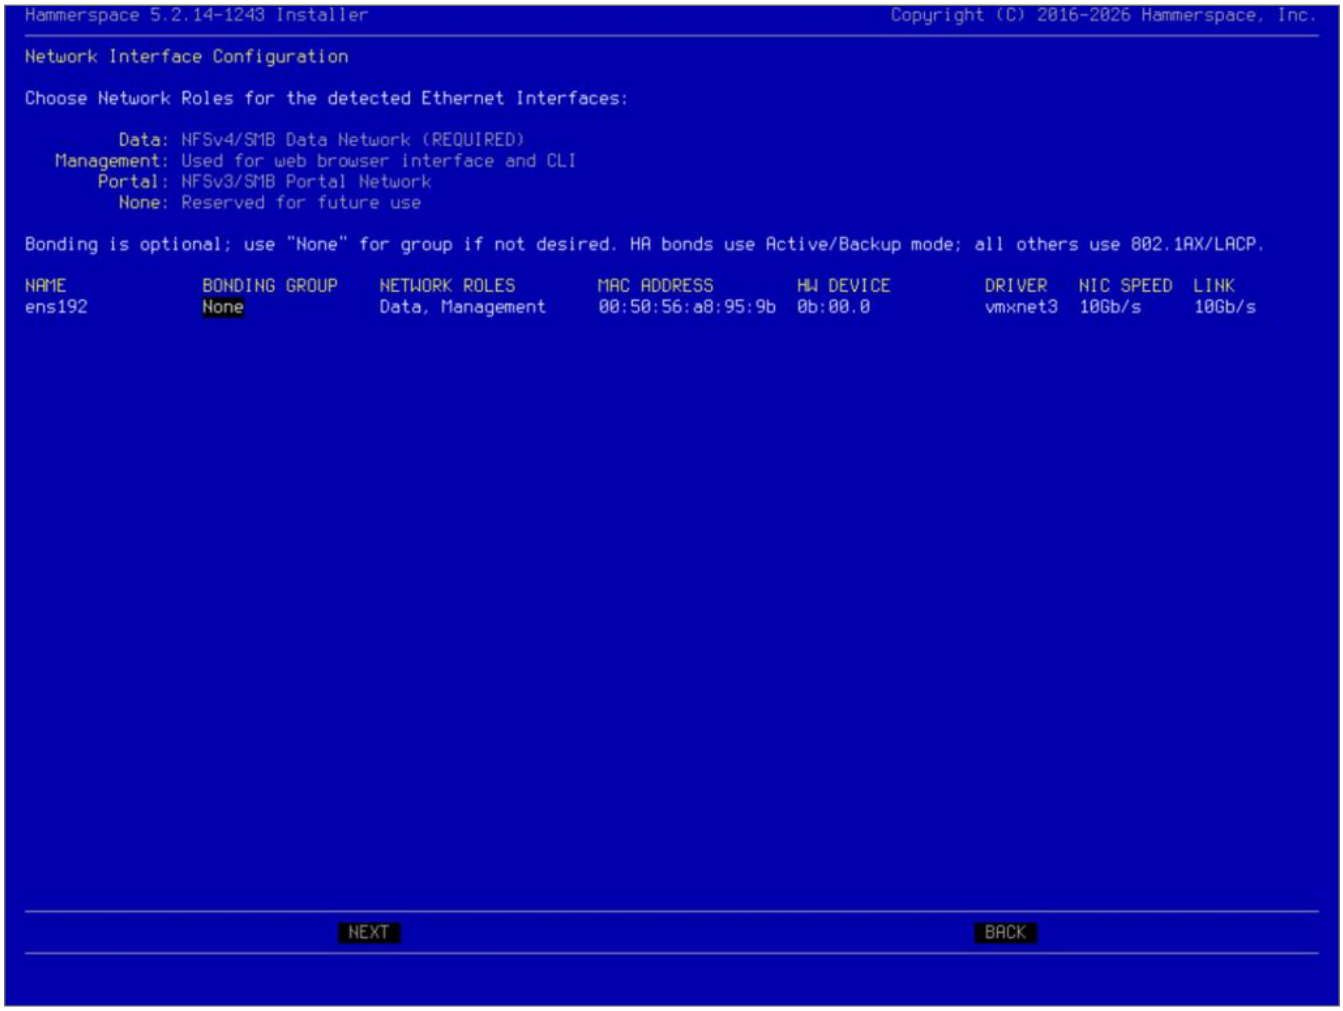

Figure 21. Network role selection

Roles are assigned as predefined combinations rather than as individual roles you switch on and off, and the combinations offered for one interface narrow once another interface has claimed a role. On a DSX node with a single network interface, the installer may arrive with `Data, Management` already assigned, and the **Portal** role is not offered because there is no second interface to carry it.

The network roles provide the following functionality:

- **Data:** Used to connect all storage clients and servers (including DSX nodes). NFS and SMB data traffic flows across the data network, as well as some inter-node communication. Floating IP Addresses and additional Data interfaces can be configured after the installation is complete.
- **Management:** Used for management and monitoring tasks. This role is not used for file sharing, data movement, or HA-cluster communications.
- **Portal:** Optionally used for NFSv3 and SMB network traffic to clients. If it is not configured, then the Data network will be used for client traffic.

  All installations require management and data roles. This can be configured with only one interface if it is configured to share data and management traffic.

It is also possible to configure additional interfaces once the system is installed. DSX supports floating IP addresses for high-availability configurations and VLAN assignment to interfaces.

All configuration settings are accessible from the Admin CLI after installation is complete.

12. Fill in the following settings for the DSX:

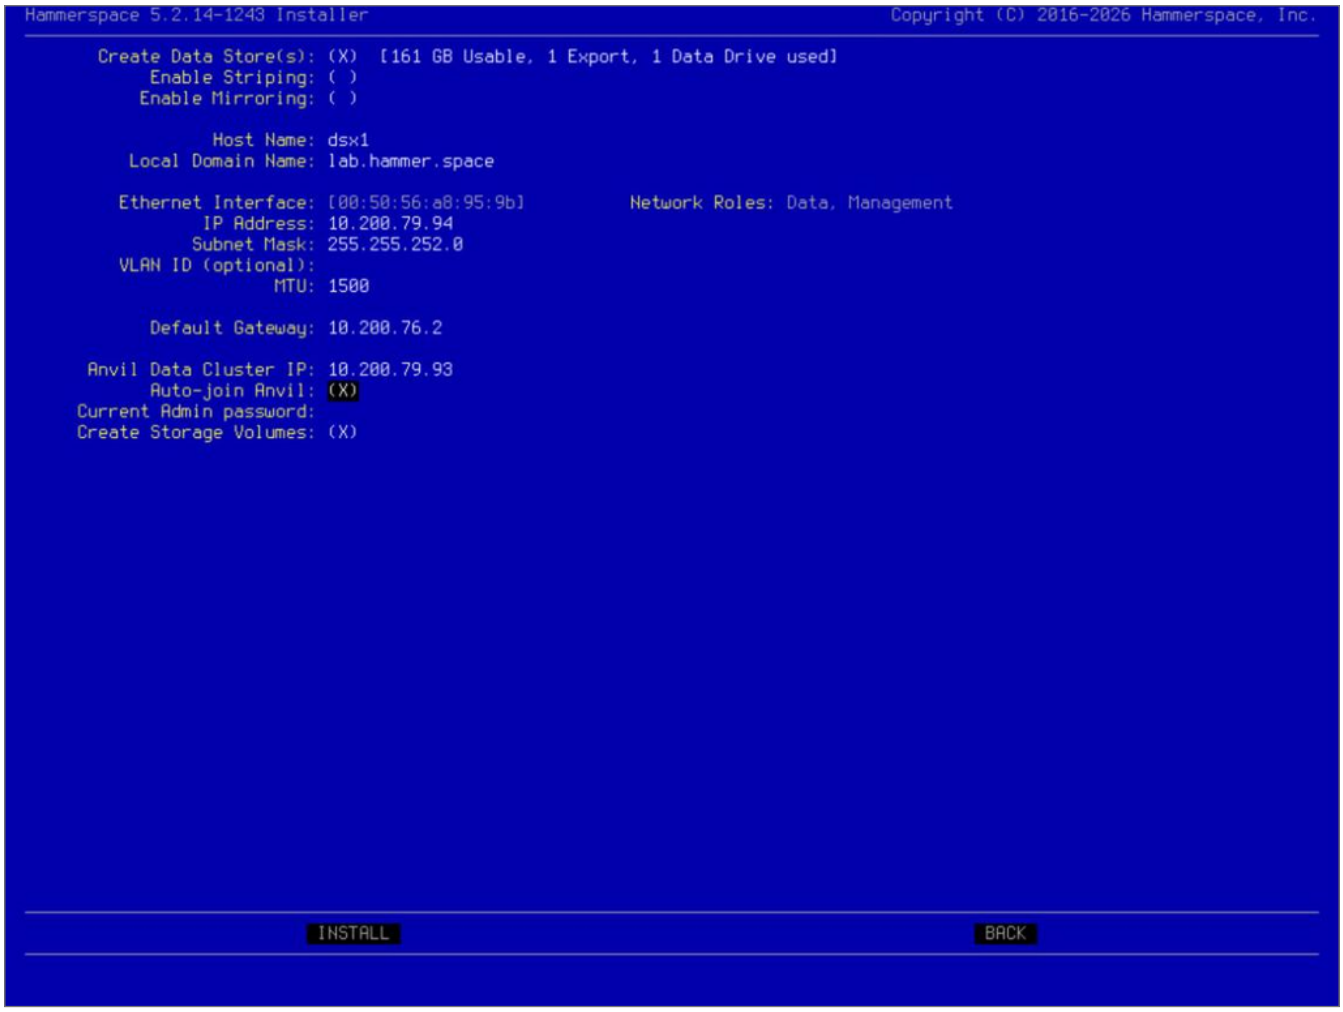

Figure 22. Create Data Store settings

Where:

- **Create Data Store(s):** Enables locally attached block storage to be used as storage volumes.
- **Enable Striping:** (Optional) Enables data striping across locally attached block storage. For local drives, this can only be done during installation.
- **Enable Mirroring:** (Optional) Enables mirroring of data to locally attached block storage. For local drives, this can only be done during installation.

Striping and mirroring each require two or more data drives, and for local drives both are install-time decisions only. Neither can be enabled after the installation completes, so decide before committing the configuration.

- **Host Name**: Enter the DSX host name.
- **Local Domain Name**: Enter the name of the local domain.
- **Ethernet Interface**: For each network Interface, configure the following:
  - IP Address (each interface needs to be on a different subnet)
  - Subnet Mask or CIDR
  - VLAN ID (optional)
  - MTU size (default is 1500), 9000 for Jumbo Frames.
  - IP address of the **Default gateway**
  - IP address of the **Anvil Data Cluster IP** (set during Anvil installation)

The **Anvil Data Cluster IP** is the cluster's floating address, not the address of either individual Anvil node.

- **Auto-join Anvil:** Enables the DSX to join the Anvil automatically. Requires the Anvil admin password.
- **Current Admin Password:** Anvil admin password.
- **Create Storage Volumes:** Let the DSX automatically create volumes on data drives and add them to the configuration. Using this option will speed up and simplify the initial setup of DSX nodes.

The DSX configuration form has no DNS or NTP fields. This is expected, and differs from the Anvil form, where DNS is required. DSX nodes take their time from the Anvil servers, which act as NTP servers for them.

13. Review the configuration for accuracy. If all settings are correct, select **Install** and press `Enter`.
14. Select **Yes** and press `Enter`.

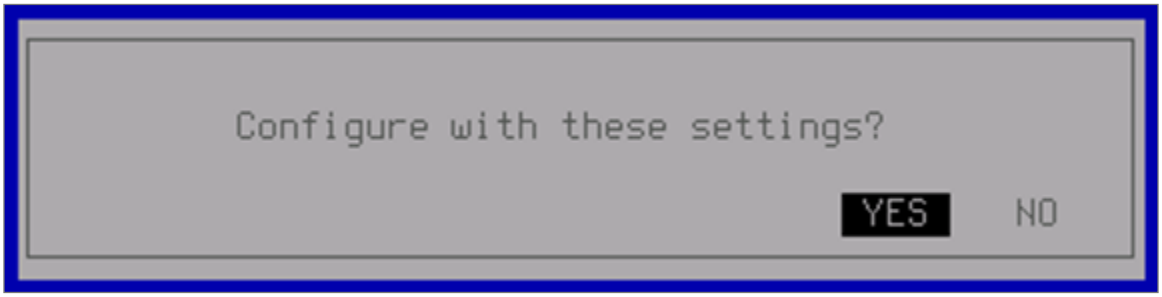

Figure 23. Confirming the DSX configuration

The installer configures the networking and installs the DSX software. Once the installation completes, the **ONLINE** dialog displays.

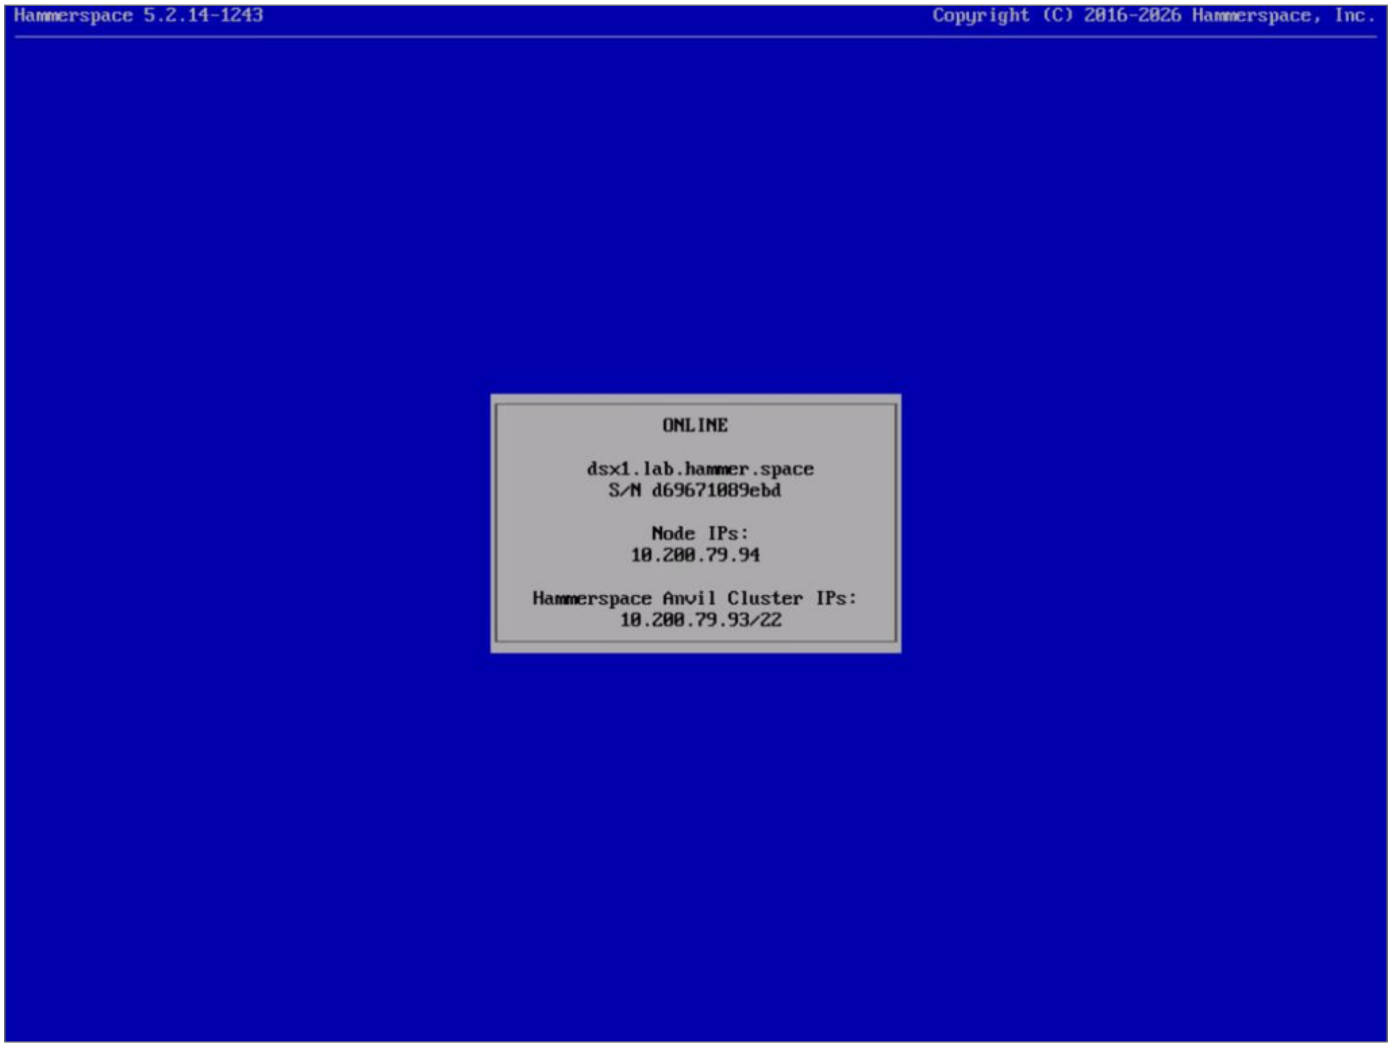

Figure 24. Confirmation that the DSX is online

If you install DSX concurrently with the Anvils and the Anvil installation has not yet completed, the DSX configuration process will pause for 60 seconds and retry until it is successful.

You can also verify the installation was successful by checking the DSX configuration in the Management GUI. Log in to the GUI using a browser, then select **Infrastructure › Storage Systems**. The DSX node is listed under its host name, with **Type** `Hammerspace`.

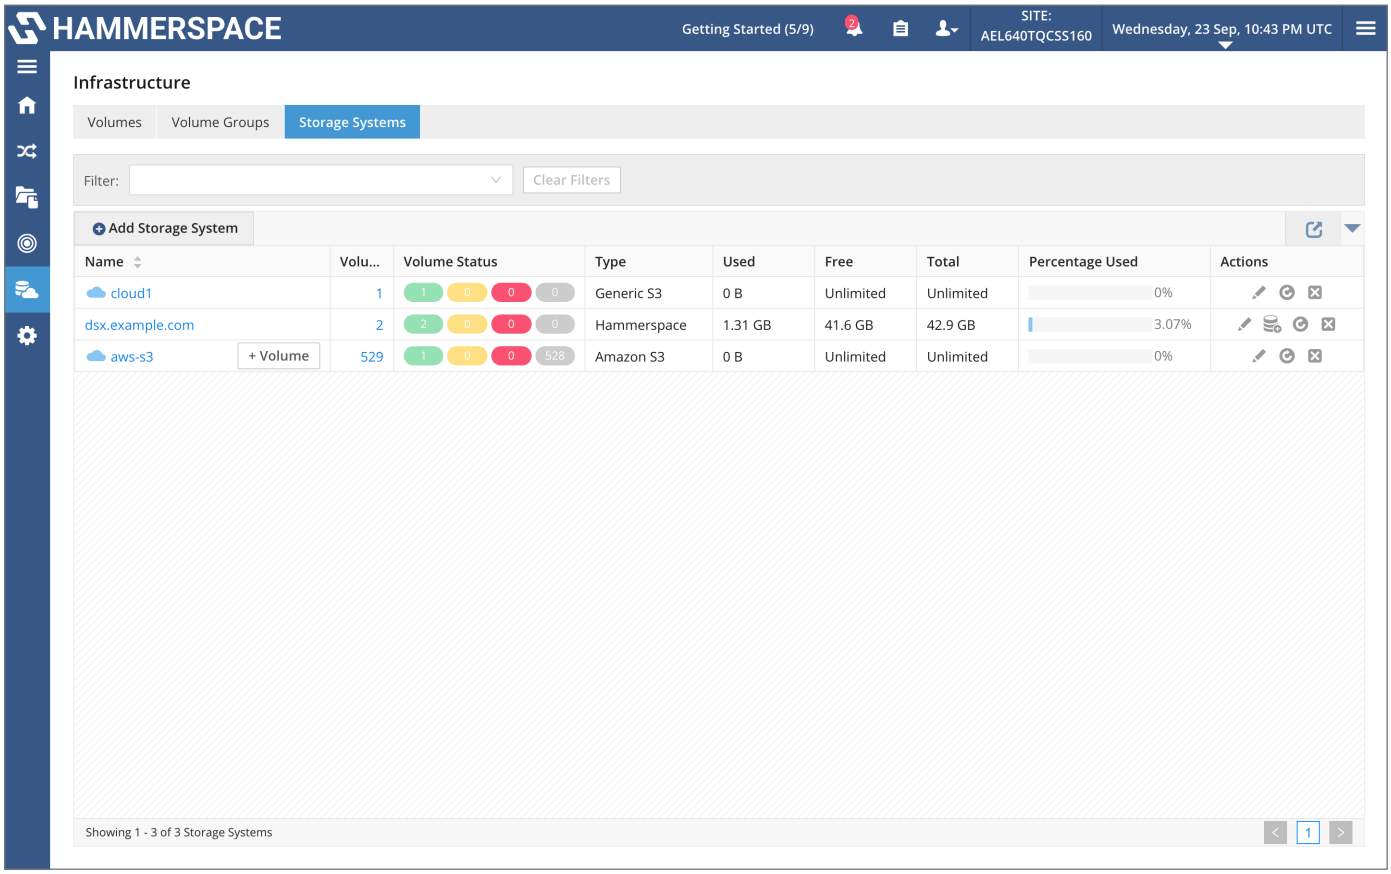

Figure 25. New DSX in the Management GUI

The DSX node is automatically added if the admin credentials are specified correctly during installation.

After the installation has completed successfully, eject the CD-ROM or remove the ISO image from the virtual machine, particularly if you moved the CD/DVD device above the hard drive in the boot order.

Chapter 4. Installing and Configuring a Quorum Server

Hammerspace requires using an external quorum server to improve resilience against certain failure conditions, especially in a two-node HA configuration. The supported external quorum server software is *Qnetd,* and it is available as an open-source package for most Linux distributions.

A single quorum server can serve multiple clusters.

External quorum server requirements:

- Qnetd (version 3.0 or later) is installed.
- Server resources are configured according to the vendor-recommended Linux requirements. A minimal server configuration is sufficient. The Qnetd software does not require significant resources to operate.
- The choice of Linux distribution is up to the customer.
- The CPU architecture can be either x86_64 or ARM.
- Only one quorum server is required and supported per cluster.
- The server must be able to connect to all Anvil nodes in the cluster.
- Must be in a different failure domain than the Anvil nodes to ensure it has the impact that is expected.

The importance of the quorum server is its functionality. A quorum server enables a cluster to use a non-cluster node as a voting member. This is especially helpful when cluster nodes lose connectivity with each other, a condition that can lead to a *split-brain*.

When a cluster membership change occurs, the surviving side with the majority of votes will become (or remain) the primary node and continue to provide services to clients.

The quorum server enables automatic decision-making for the cluster in a split-brain situation, eliminating potential downtime caused by manual steps.

4.1. Installing Qnetd

The following instructions showing how to install Qnetd on a regular Linux node are provided for two different Linux distributions. Other Linux distributions can also be used.

4.1.1. On Red Hat and Red Hat-like Distributions

Use the following commands to install Qnetd on Red Hat or Red Hat-like distributions:

Installing Qnetd on Red Hat or Red Hat-like Linux

```
# dnf install -y corosync-qnetd

Disable TLS:
```

```
# sed -i 's/COROSYNC_QNETD_OPTIONS=""/COROSYNC_QNETD_OPTIONS="-s off"/g' /etc/sysconfig/corosync-qnetd

Enable and start the service:

# systemctl enable corosync-qnetd.service

# systemctl start corosync-qnetd.service

Open firewall ports (if needed):

# firewall-cmd --zone public --add-service=high-availability

# firewall-cmd --zone public --add-service=high-availability --permanent
```

4.1.2. On Ubuntu

Use the following commands to install Qnetd on Ubuntu Linux:

Installing Qnetd on Ubuntu Linux

```
# apt install -y corosync-qnetd

Disable TLS:

# sed -i 's/COROSYNC_QNETD_OPTIONS=""/COROSYNC_QNETD_OPTIONS="-s off"/g' /etc/default/corosync-qnetd

Enable and start the service:

# systemctl enable corosync-qnetd.service

# systemctl start corosync-qnetd.service

Open firewall ports (if enabled; 5403 is the default qnetd port):

# ufw allow 5403
```

Once the Qnetd software is installed on the Linux server, no further configuration is required on the Linux quorum server.

4.2. Configuring an External Quorum Server

After the quorum server is installed, the cluster must be configured to use it.

External quorum server functionality is configured using the sysadmin CLI, not the Hammerspace Admin CLI.

Configuring external quorum server

```
# ssh sysadmin@<ANVIL CLUSTER IP>

...

Run quorum-device-set to configure 192.0.2.50 as the quorum server. Sample output included below.

sysadmin> quorum-device-set --host 192.0.2.50

anvil2.example.com: Succeeded

anvil1.example.com: Succeeded

Corosync configuration has been synchronized, please reload corosync daemon using 'pcs cluster reload
corosync' command.

Corosync reloaded

anvil2.example.com: Cluster Enabled

anvil1.example.com: Cluster Enabled

anvil2.example.com: Starting Cluster...

anvil1.example.com: Starting Cluster...
```

Troubleshooting

If you get an error like this while trying to configure a quorum server, then it is likely that one of the Anvil nodes is not reachable at the moment.

Sysadmin CLI

```
Unable to connect to anvil2.example.com, check if pcsd is running there or try setting higher timeout
with --request-timeout option (Failed to connect to anvil2.example.com port 2224: Connection refused)

Error: Unable to set corosync config: Unable to connect to anvil2.example.com, check if pcsd is running
there or try setting higher timeout with --request-timeout option (Failed to connect to
anvil2.example.com port 2224: Connection refused)

Command '['pcs', 'cluster', 'sync']' returned non-zero exit status 1.
```

An external quorum server is **required** for high-availability (two-Anvil) deployments.

4.3. Checking Status of Configured Quorum Server

You can check the status of the external quorum server using the following command:

Check the status of external quorum server

```
# ssh sysadmin@<ANVIL CLUSTER IP>
...
Run quorum-device-status to verify the configuration.
sysadmin> quorum-device-status

Qdevice information
-------------------
Model: Net
Node ID: 2
Configured node list:
0 Node ID = 1
1 Node ID = 2
Membership node list: 1, 2

Qdevice-net information
----------------------
Cluster name: hs-0f9e8d7c-6b5a-4321-8000-000000000001
QNetd host: 192.0.2.50:5403
Algorithm: Fifty-Fifty split
Tie-breaker: Node with lowest node ID
State: Connected
```

Chapter 5. Installation Using Terraform

The product can be deployed using Terraform automation. Terraform scripts and documentation are available on the Hammerspace Terraform GitHub.

Chapter 6. Completing the Installation

You have installed the components necessary to start using Hammerspace. Before using Hammerspace in production, complete licensing — see Licensing.

6.1. Verifying the Installation

Before handing the cluster over or starting configuration work, confirm that each part of the installation finished the way it was meant to. Every check below is read-only — none of them change the configuration — so they are safe to run at any point after installation, and they can be done in any order.

1. **Confirm the Anvil nodes are healthy.** Log in to the Management GUI in a browser, using the Anvil Data Cluster IP address you set during installation, and open **Administration › System**. The **Anvil Nodes** panel shows one card per Anvil with a green check when the node is healthy; in a two-node HA deployment both Anvil nodes should be present, one card headed **Primary** and the other **Secondary**, and the **Status** at the top right of the page should read **High Availability**. From the Admin CLI, `node-list` lists every node with its state — each Anvil should show **Product node type** `Anvil`, **HW state** `OK` and **Node mode** `ONLINE` — and `cluster-config` shows which Anvil is the primary in its **Metadata servers** line (`Role: Primary`):

   Admin CLI

   ```
   node-list
   ```

   Expected output (two-node HA cluster; a standalone deployment lists one Anvil):

   ```
   total 4
   Name:                   anvil1.example.com
   Type:                   Product
   Product node type:      Anvil
   Internal ID:            1073741831
   ID:                     2df81f5b-8e0f-5077-a0ce-f40e2a1093f0
   HW state:               OK
   Node state:             MANAGED
   Node mode:              ONLINE
   Management IP:          192.0.2.11/20
   SW version:             5.2.14-1266

   Name:                   anvil2.example.com
   Type:                   Product
   Product node type:      Anvil
   Internal ID:            1073741832
   ID:                     8b6d6ff7-43a5-51b2-957e-5062e8f758c6
   HW state:               OK
   Node state:             MANAGED
   Node mode:              ONLINE
   Management IP:          192.0.2.12/20
   SW version:             5.2.14-1266
   ```

```
Name:                    dsx.example.com
Type:                    Product
Product node type:       DSX
Internal ID:             1073741834
ID:                      2c7ac822-2cc2-590e-86cb-f52f71b8882a
HW state:                OK
Node state:              MANAGED
Node mode:               ONLINE
Management IP:           192.0.2.21/20
SW version:              5.2.14-1266
<Additional output removed for brevity>
```

```
cluster-config
```

2. **Confirm the quorum server is connected.** This applies to two-Anvil HA deployments, which require an external quorum server. Connect to the Anvil cluster IP as the `sysadmin` user and run `quorum-device-status`. In the output, **State** should read `Connected`. For the full output and for configuration steps, see Installing and Configuring a Quorum Server.
3. **Confirm the DSX nodes are online.** In the Management GUI, navigate to **Infrastructure › Storage Systems**. Each DSX node you installed should be listed there under its host name, with **Type** `Hammerspace`; the list shows no node state, so for the node's health open **Administration › System**, where each product node has a green check, or run `node-list` from the Admin CLI and confirm **Node mode** `ONLINE`. A DSX node installed with **Auto-join Anvil** enabled is added automatically, provided the Anvil administrator credentials were entered correctly during its installation.

   

   If you installed DSX nodes at the same time as the Anvils, a DSX node may take a little longer to appear. The DSX configuration process pauses for 60 seconds and retries until the Anvil is reachable.

4. **For VMware deployments, confirm time synchronization is disabled for each VM.** This is a per-VM vSphere setting, not something the installer applies, so it is easy to miss — and it has to be done for every Anvil and DSX virtual machine, not just once for the cluster. See the time synchronization section.
5. **Confirm licensing once it is complete.** Licensing is covered in Licensing. After a license has been applied, verify it from the Admin CLI with `license-list`, where the **Status** field should read `Valid`, or open the **Administration › Licenses** tab in the Management GUI, where a valid license is listed with a green check.

6.2. Next Steps

Log in to the Management GUI and use the **Getting Started** checklist — the **Getting Started** button in the top bar — to finish the most common setup steps. Each item is a link to the page where that step is done, and the button's count shows how many are complete.

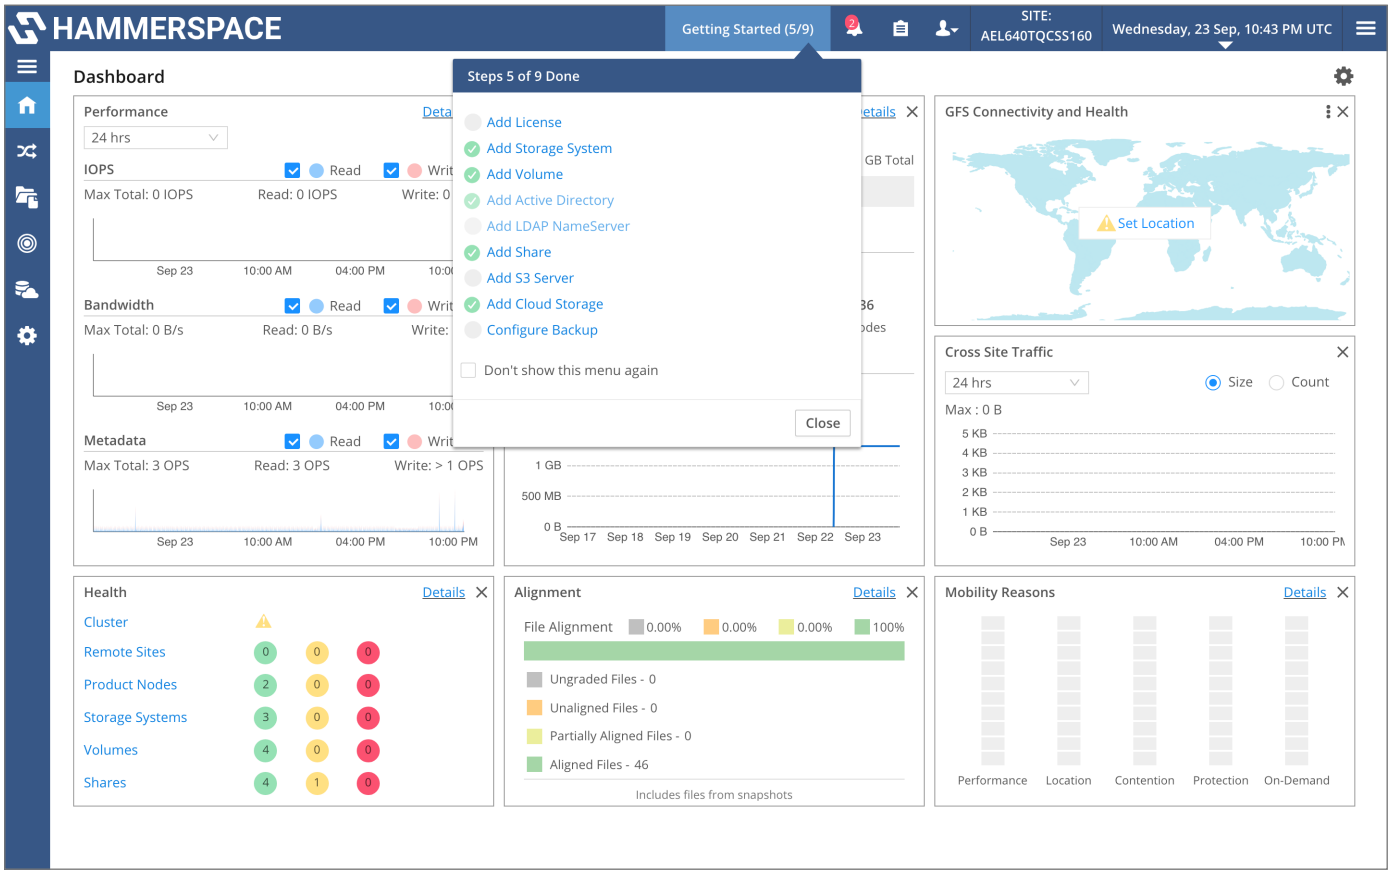

Figure 26. Getting Started checklist

Chapter 7. Upgrading Hammerspace

The Hammerspace product (Product) supports upgrading an existing, running version without scheduled downtime for applications.

High-availability deployment is required with configured floating IPs; IP failover will occur as part of the upgrade.

The basic workflow to upgrade the cluster software includes the following steps:

1. Download the update file.
2. Upload the update file to the cluster.
3. Install the new software on the nodes in the cluster.
4. Monitor the installation progress.

7.1. Downloading the Update File

The update (`.upd`) file can be downloaded from the Download Packages page of the *Software Delivery and Licensing portal*.

1. Go to the linked page and click the link for the software version you are using, for example, **Hammerspace-5.x**.
2. Agree to the export license agreement.
3. Then select the version of the update (`.upd`) file you wish to download.
   - You can download the `.upd` file by clicking on the filename download link on the right side of the page.
   - You can check boxes on the left side of the page to download multiple software packages at once.
   - On the left side of the page, click the **+** button for the software package to see the MD5 signature and SHA256 CheckSum for the file(s) you download.

Continue to the section, Uploading the Update Package to the Cluster, for the next steps.

The product *Release Notes* specify which versions are compatible for upgrade. The release notes are available to download on the *Software Delivery and Licensing portal*.

7.2. Uploading the Update Package to the Cluster

Start by uploading the update file to the cluster. You can upload the file using the Management GUI or the Admin CLI.

It is recommended to use the GUI upload method for most use cases.

7.2.1. Using the GUI

1. In the Management GUI, navigate to **Administration** in the left navigation menu and then select the **Software Update** tab:

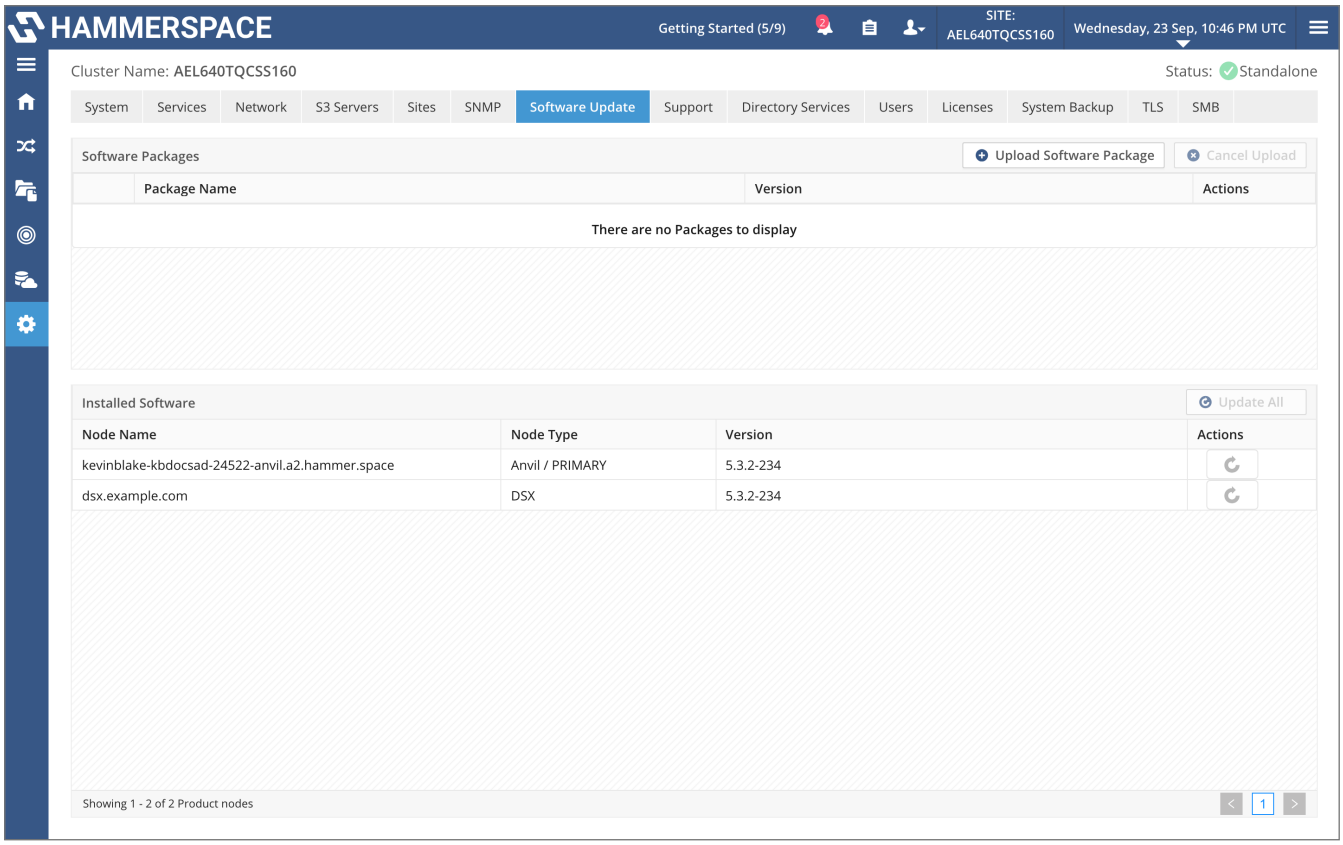

2. Click on **Upload Software Package** to select the update file. Once the file is selected, it will start uploading, and a progress bar will appear.

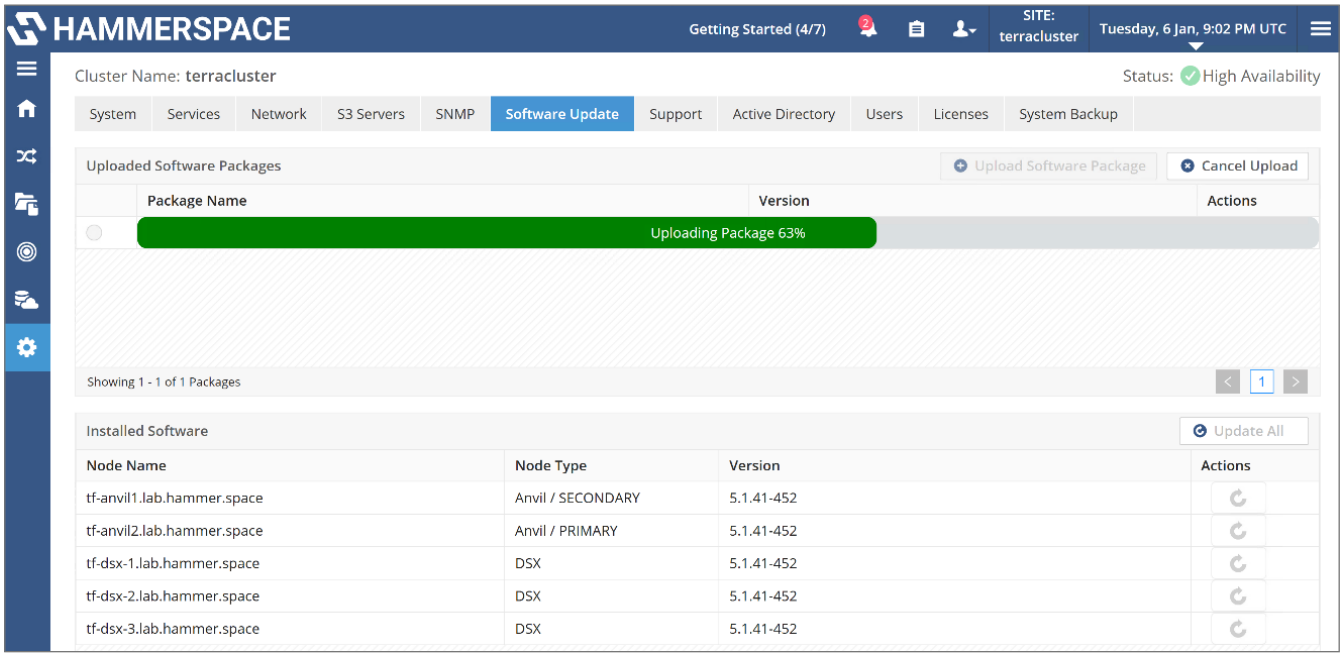

Once the upload is complete, the **Update actions** become enabled.

7.2.2. Using the CLI

To update the Hammerspace cluster using the CLI, you must first upload the UPD software package to the cluster using a supported URI. The cluster does not host a local file browser, so the update package must be accessible via a network-reachable location.

Prerequisites

Ensure the update package (`.upd` file) is stored in a location accessible by the cluster.

The cluster supports the following URI schemes:

- http://
- scp://
- ftp://
- file://

You must ensure that the cluster has network access to the server or file share hosting the package. This may require configuring routing, DNS, or firewall rules to allow access.

Example command

Command:

```
software-upload --package-location <URI>
```

Expected output:

```
success
```

The command returns once the package has been fetched and stored on the cluster, which can take several minutes for a package of a few gigabytes. Replace **<URI>** with the full path to the update package, for example

- HTTP:

  ```
  software-upload --package-location http://updates.example.com/path/to/hammerspace-5.2.8-145.upd
  ```

- SCP:

  ```
  software-upload --package-location scp://user@host:/path/to/hammerspace-5.2.8-145.upd
  ```

- FTP:

```
software-upload --package-location ftp://user:pass@ftp.example.com/pub/hammerspace-5.2.8-145.upd
```

- Local file (when using file://) the file is located in share called data:

```
software-upload --package-location file:///hs/data/hammerspace-5.2.8-145.upd
```

Once the package is uploaded, proceed with the software installation instructions provided in the next section.

Example: Uploading the UPD File from a Local Path on the Cluster

Uploading the file from a local file path on the cluster requires additional steps. The following workflow example uploads the UPD file to a Hammerspace share named '**data**'. This is a user-configured share in your environment and can be named anything.

1. Mount the Hammerspace **data** share on the workstation.
2. Download the UPD from the *Hammerspace Software & Delivery* site and save it to the **data** share.
3. Execute the following command in the Admin CLI:

   Command:

   ```
   software-upload --package-location file:///hs/data/hammerspace-5.2.8-145.upd
   ```

   Expected output:

   ```
   success
   ```

   

   Note the triple `///` and the addition of **hs** to get to the root of the file system. Also, be aware that the permissions on the UPD need to be readable by other (**chmod o+r**) since the management layer reading the update file does not have root access.

Troubleshooting a UPD local-path upload

If the file is not readable by other, the upload fails. In this example the share is mounted on a Linux client at `/mnt/data` and the file is a placeholder named `hammerspace-docs-test.upd`; `ls -l` on the client shows the permissions:

Command:

```
ls -l /mnt/data/hammerspace-docs-test.upd
```

Expected output:

```
-rw-r----- 1 root root 0 Sep 23 22:20 /mnt/data/hammerspace-docs-test.upd
```

The system then returns the following error message:

Command:

```
software-upload --package-location file:///hs/data/hammerspace-docs-test.upd
```

Error:

```
software-upload: URI '/hs/data/hammerspace-docs-test.upd' not found.
```

The same message is returned when the file does not exist at the path given, so check the path and the file name as well as the permissions.

Solution

You must change the permissions on the UPD to make sure that non-root users (other) can read the file. This can be done from a Linux or Windows client.

Linux Client CLI

```
# chmod o+r hammerspace-docs-test.upd
```

7.3. Updating the Nodes

After the UPD file is uploaded to the cluster, you can begin updating the nodes.

7.3.1. Using the GUI

In this example workflow, we will update a reference installation currently running **5.1.41-452** to **5.2.8-145**, which is available for the upgrade.

It is recommended to use the **Update All** action in most configurations, as Hammerspace will automatically determine the correct order for applying updates.

In configurations where Mirroring is used via Objectives, it is recommended to first upgrade the Anvil nodes, then upgrade each DSX individually, and wait between upgrades for any re-silvering operations to finish. Re-silvering can be monitored by checking the alignment status of the share/files.

1. In the **Software Update** tab, press **Update All** (in the **Installed Software** panel; the button is unavailable until a package upload has completed) to start the process. It may take up to a few minutes for the process to show progress in the UI. The reference configuration is an HA config, and it will start by upgrading the secondary node first.

Once the secondary node is upgraded, the primary node upgrade will begin, and the Management GUI will be unavailable for a few minutes.

Close the Management GUI tab or window to minimize the effects of browser cache.

7.3.2. Using the CLI

Use the following commands to initiate the software update on the nodes using the Admin CLI.

1. List the uploaded package to confirm its version:

   Command:

   ```
   software-list
   ```

   Expected output:

   ```
   total 1
   Version:                5.3.1
   Build:                  1209
   Package name:           hammerspace-upd-5.3.1-1209.tar
   Full version:           5.3.1-1209
   Product name:           Hammerspace
   ```

2. Start the update on every node, naming the version shown by `software-list`:

   Command:

   ```
   software-apply --all --version 5.3.1-1209
   ```

   Expected output:

   ```
   Software update started.
   The progress may be monitored using software-update-status --task-id 873a8a88-62ce-4df5-8a19-
   a9d653aa8f4d
   ```

   The task ID identifies this update in `software-update-status` (next section) and in `software-update-cancel`.

7.4. Monitoring the Update Progress

Progress can be monitored from the **Tasks** icon in the top bar of the Management GUI, on the **Software Update** page, or in the Admin CLI using the `software-update-status` command.

The end-to-end, non-disruptive upgrade process takes approximately 60 minutes for a four-node installation; however, this estimate can vary depending on the configuration. For example, bare-metal environments usually

take longer because a physical server boots more slowly than a virtual server.

7.4.1. Using the GUI

After logging back in, navigate to the **Software Update** page and check the status.

Even though the management UI is not available for a few minutes, file operations are minimally affected by the upgrade process.

The table with nodes will show progress; it may take a few minutes for the page to refresh. Progress can also be monitored from the **Tasks** icon (the clipboard) in the top bar, which lists the running and recent tasks and links to **View All Tasks**.

7.4.2. Using the CLI

More detailed results of the upgrade are available in the Admin CLI. The `software-update-status` command lists the update task with one **Items** entry per node. While the update runs the task shows **Status** `ACTIVE`; the node being updated shows `IN_PROGRESS` with its percentage complete and a progress message (for example `Applying package updates`), and the nodes still to come show `PENDING`. The nodes are updated one at a time, in the order described under *Manual Upgrade Order* later in this chapter. When every node has finished, the task shows `COMPLETED` with its end time and each node shows `SUCCESS` with its start and end times:

Command:

```
software-update-status
```

Expected output:

```
total 1
ID:                     873a8a88-62ce-4df5-8a19-a9d653aa8f4d
Version:                5.3.1-1209
Status:                 COMPLETED
Started:                2026-09-23 21:42:25 UTC
Ended:                  2026-09-23 22:10:28 UTC
Items:
                        [Node name: anvil1.example.com, Status: SUCCESS, Completed: 100%, Started: 2026-
09-23 21:42:36 UTC, Ended: 2026-09-23 21:51:37 UTC]
                        [Node name: anvil2.example.com, Status: SUCCESS, Completed: 100%, Started: 2026-
09-23 21:51:37 UTC, Ended: 2026-09-23 22:01:27 UTC]
                        [Node name: dsx.example.com, Status: SUCCESS, Completed: 100%, Started: 2026-09-
23 22:01:27 UTC, Ended: 2026-09-23 22:10:07 UTC]
```

7.5. Manual Upgrade Order

If updating the cluster manually, update one node at a time: in the Management GUI, use the update icon in the **Actions** column of the **Installed Software** table on the **Software Update** tab; from the Admin CLI, run `software-`

`apply --node-name <node name> --version <version>` for each node. The supported order is:

1. Update the Anvil secondary node.
2. Update the Anvil primary node.
3. Update the DSX nodes.

The DSX nodes can generally be updated in any order after Steps 1 and 2. Floating IP addresses will automatically fail over to other DSX nodes during the upgrade.

If you are using the built-in mirror capability in the product, then it is recommended to ensure that the shares reach full alignment before upgrading the next DSX node.

7.6. Upgrading Global File System Environments

When upgrading an environment with replicated shares (Global File System), the version change to 5.2 will stop metadata replication until ALL participating sites have been upgraded to 5.2.

Upgrade all sites simultaneously to minimize the metadata replication delay during the upgrade.

Chapter 8. Licensing

Hammerspace requires a valid license to operate beyond the initial trial period. Without a license, the BYOL version of the software will operate as a 30-day trial. At the end of the trial period, client access to the data will be disabled; however, the data itself will remain unaffected. A license key is required to continue using the software after the trial period.

Each customer is provided an account on the *Software Delivery & License Portal*, where software and documentation can be downloaded. Information about entitlements associated with the customer account is also available.

An entitlement email is sent to the account owner and lists one or more Activation IDs associated with each product when an entitlement is created, updated, or renewed. It will include information like this example:

Example entitlement email

```
Subject: Entitlement ID(s): zzzz-yyyy-xxxx-wwww-vvvv-uuuu-tttt-ssss from Hammerspace

ActivationID: aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111

Product Name: Hammerspace

Version: 5.2.0

Total Quantity: 10

Start Date: Jan 1, 2026

Expiration Date: Jan 1, 2027

Duration: 365 Days
```

On installations with connectivity to the *Software Delivery and Licensing portal*, licenses can be activated using the Management GUI or the Admin CLI. If the Anvil is not connected to the licensing server, see the section on using Offline licensing.

Only a single license can be applied to a system at a time. If changing a license, the old license must be removed first.

8.1. Licensing Connectivity Requirements

Connectivity to the licensing service is required to use online licensing. If that cannot be configured due to networking restrictions, you must use Offline licensing.

8.1.1. Online Licensing Connectivity Requirements

Online licensing requires connectivity to **hammerspace.compliance.flexnetoperations.com**. The port is **443** (TCP/HTTPS).

If using an *IP Allow list*, use a network block instead of individual IP addresses because these may change from one call to the next.

The following network block(s) should be whitelisted:

- IPv4: 64.14.29.0/24 (Deprecated end of January 2026)
- IPv4: 185.146.155.0/24 (Starting end of January 2026)

It is currently not supported to use security appliances that intercept HTTPS requests and replace them with their own connections (e.g., Zscaler). It will cause the session to become invalid, and licensing will not complete. If it is not possible to avoid this, then it is recommended to use
Offline licensing.

8.2. Online License Activation

The activation ID from the entitlement email is required to license the product. The same activation ID is also available in the Software Delivery and Licensing portal and can be copied from there. Please contact support if you did not receive the activation ID or cannot locate it.

Online licensing automatically checks in with the license server every 24 hours. If the expiration date changes (after a software renewal, for example), it will be picked up automatically, and no further action is required.

If the license server cannot be reached, then the existing expiration date will continue to be used.

8.2.1. Using the GUI

Complete the following steps to activate your license using the Management GUI.

1. Navigate to **Administration** › **Licenses** tab.

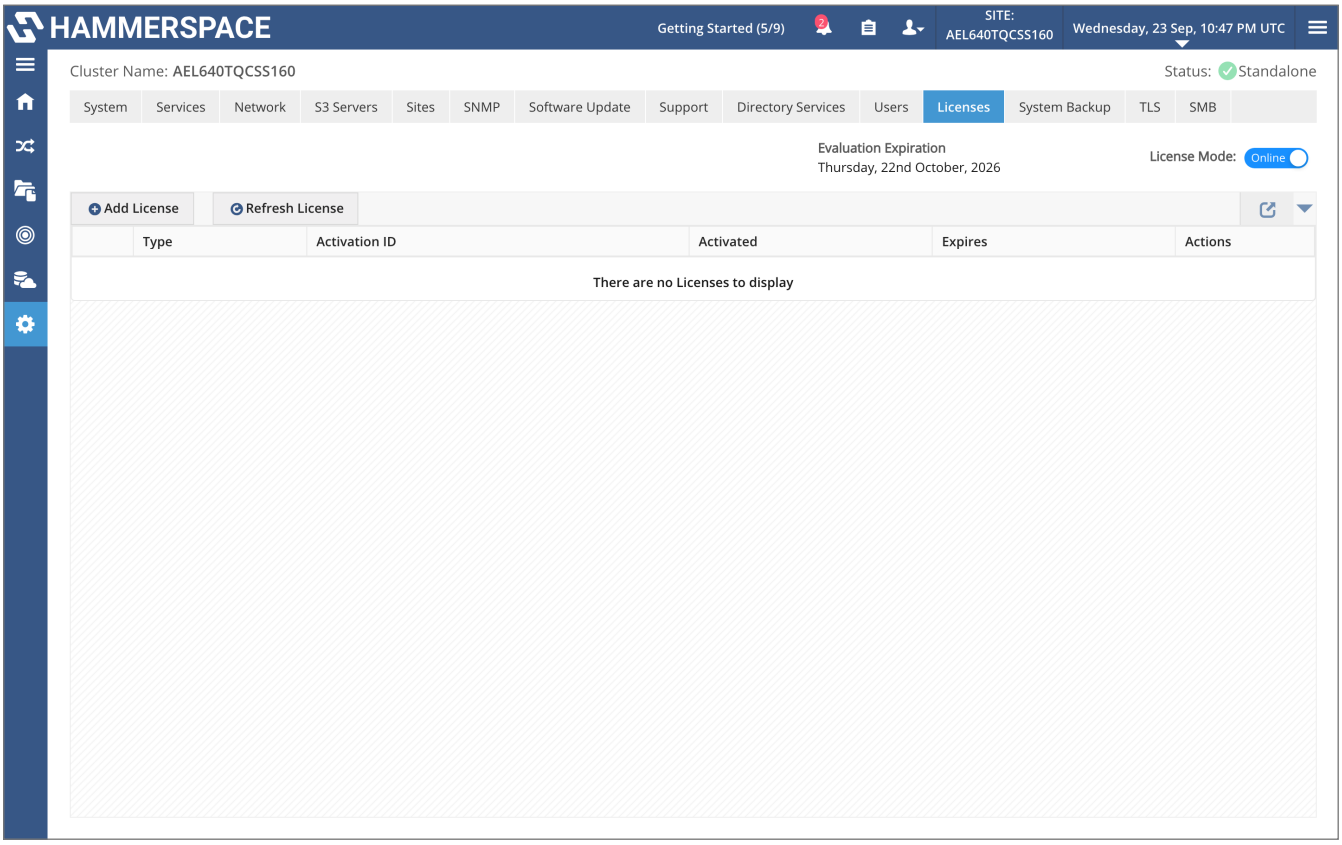

Figure 27. View license

2. Click **Add License** and add the activation ID. The default license type is Node based. It is normal for the Activation to take up to several minutes. It is safe to navigate away from the page after the job has been submitted.

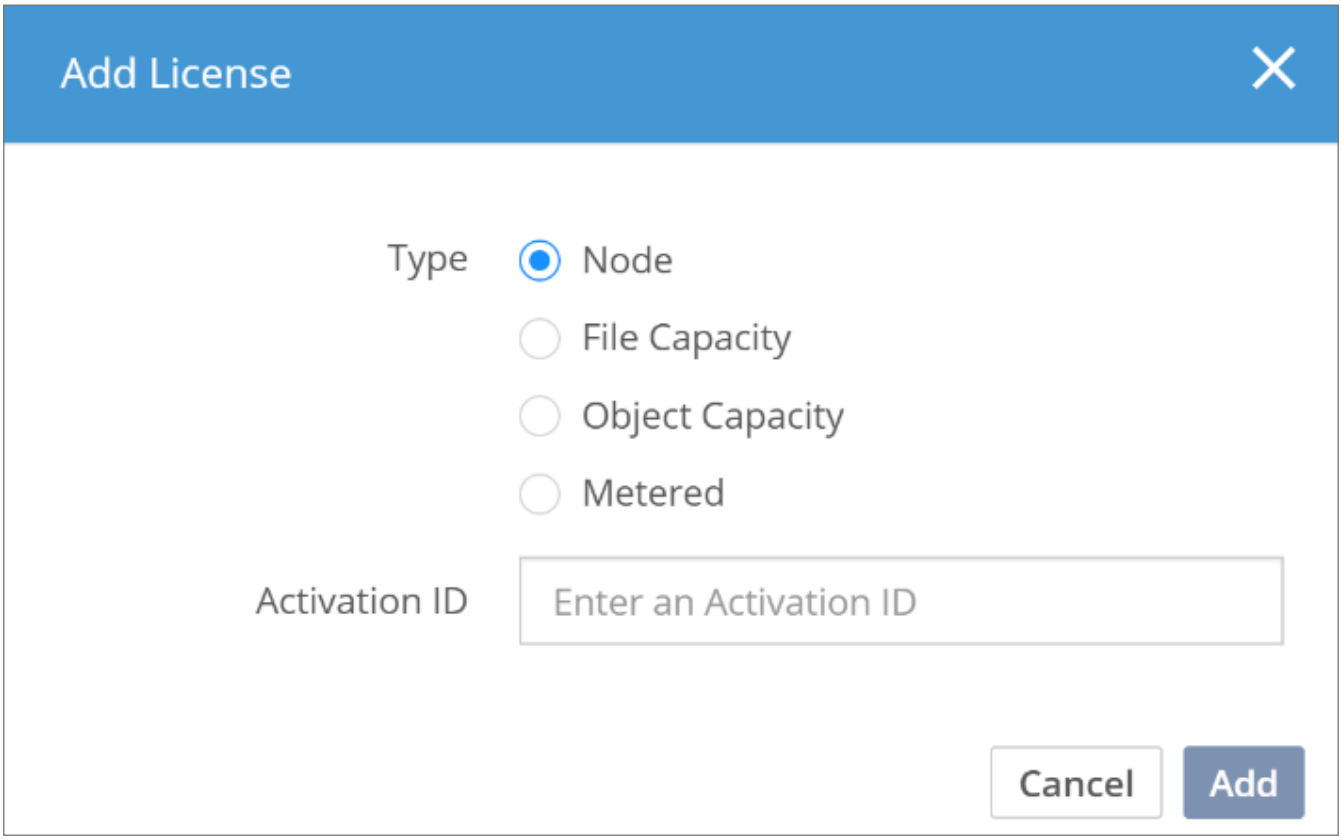

Figure 28. Add license

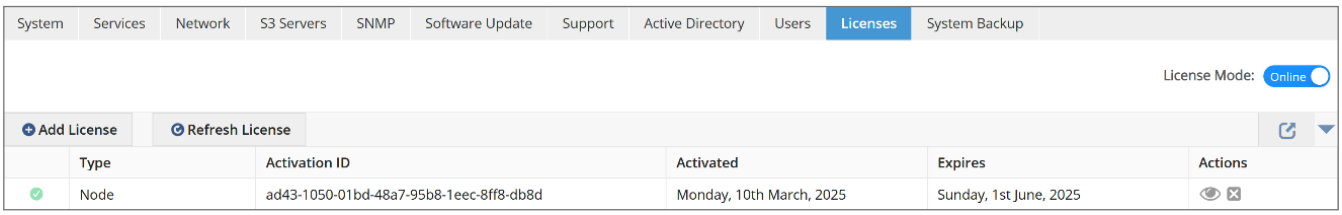

Figure 29. License view with a license installed

After a successful activation, the license is installed and appears in the **Licenses** tab.

8.2.2. Using the CLI

You can access your license using the command line:

1. Log in to Hammerspace via SSH as an admin user, then add the license with the `license-add` command.

   Command:

   ```
   license-add --type NODE --activation-id aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111
   ```

   Expected output:

   ```
   Activation ID: aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111
   Type: Node
   ```

```
Status: Valid
Activated: 2026-02-22 19:48:30 UTC
Expires: 2027-08-14 23:59:59 UTC
Number of acquired licenses: 2
```

2. Review the installed license(s).

   Command:

   ```
   license-list
   ```

   Expected output:

   ```
   total 1
   Activation ID: aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111
   Type: Node
   Status: Valid
   Activated: 2026-02-22 19:48:30 UTC
   Expires: 2027-08-14 23:59:59 UTC
   Number of acquired licenses: 2
   ```

8.3. Offline License Activation

Licensing in offline mode is required when a Hammerspace installation cannot reach the licensing service endpoints.

Licensing is a cluster operation and does not need to be executed on each node, even in offline mode.

With offline licensing, the expiration date will not change after the license is installed. When the license expiration changes, such as after a software renewal, a new offline license key must be generated and installed following the *Update Offline License* process.

8.3.1. Using the GUI

You can use the Management GUI to update your license in offline mode. This involves setting offline licensing mode and then uploading a new license from the *License & Delivery Portal*. Complete the following steps to activate a new license in offline mode using the GUI.

1. Start by switching the product to Offline License Mode. Click the License Mode toggle.

Figure 30. Offline licensing - Switching licensing mode to Offline

2. Click **Add License** and input the Activation ID into the dialog.
3. Click **Add** and a request file will be automatically downloaded using your browser action.

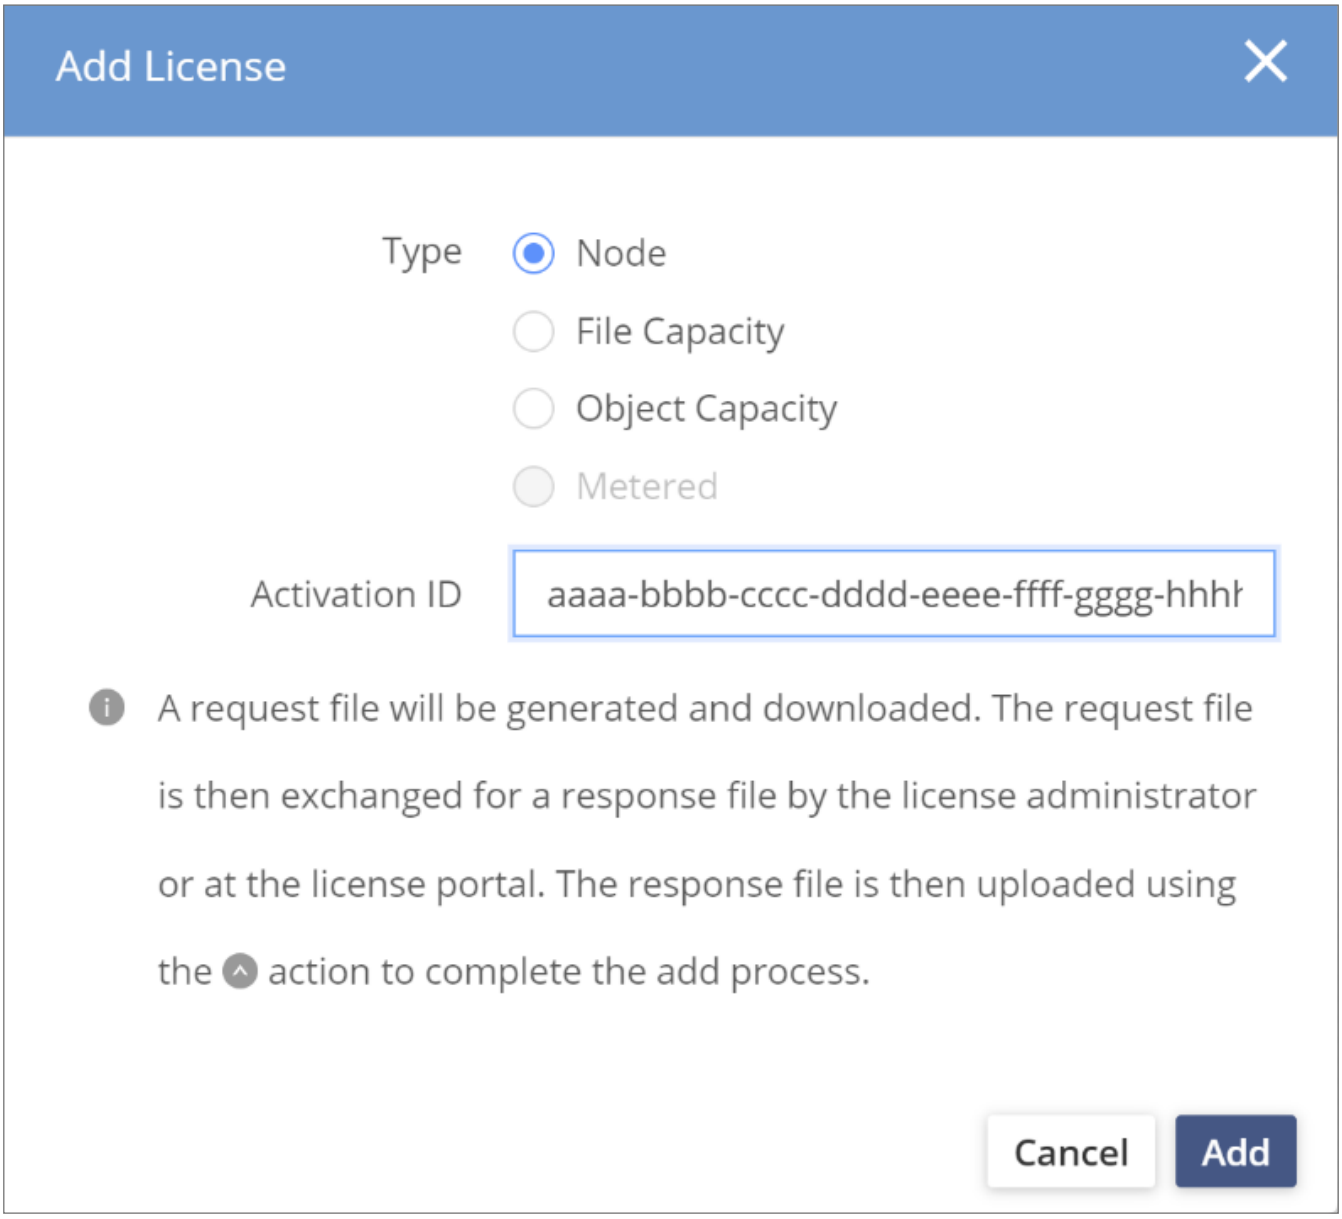

Figure 31. Adding the activation ID

4. Use that request file to generate a license file. Log in to the Software Delivery and Licensing portal (Flexera), then navigate to **Devices** > **Offline device management**.

5. Select the **license-add-request.bin** file from the previous step and click **Upload**.

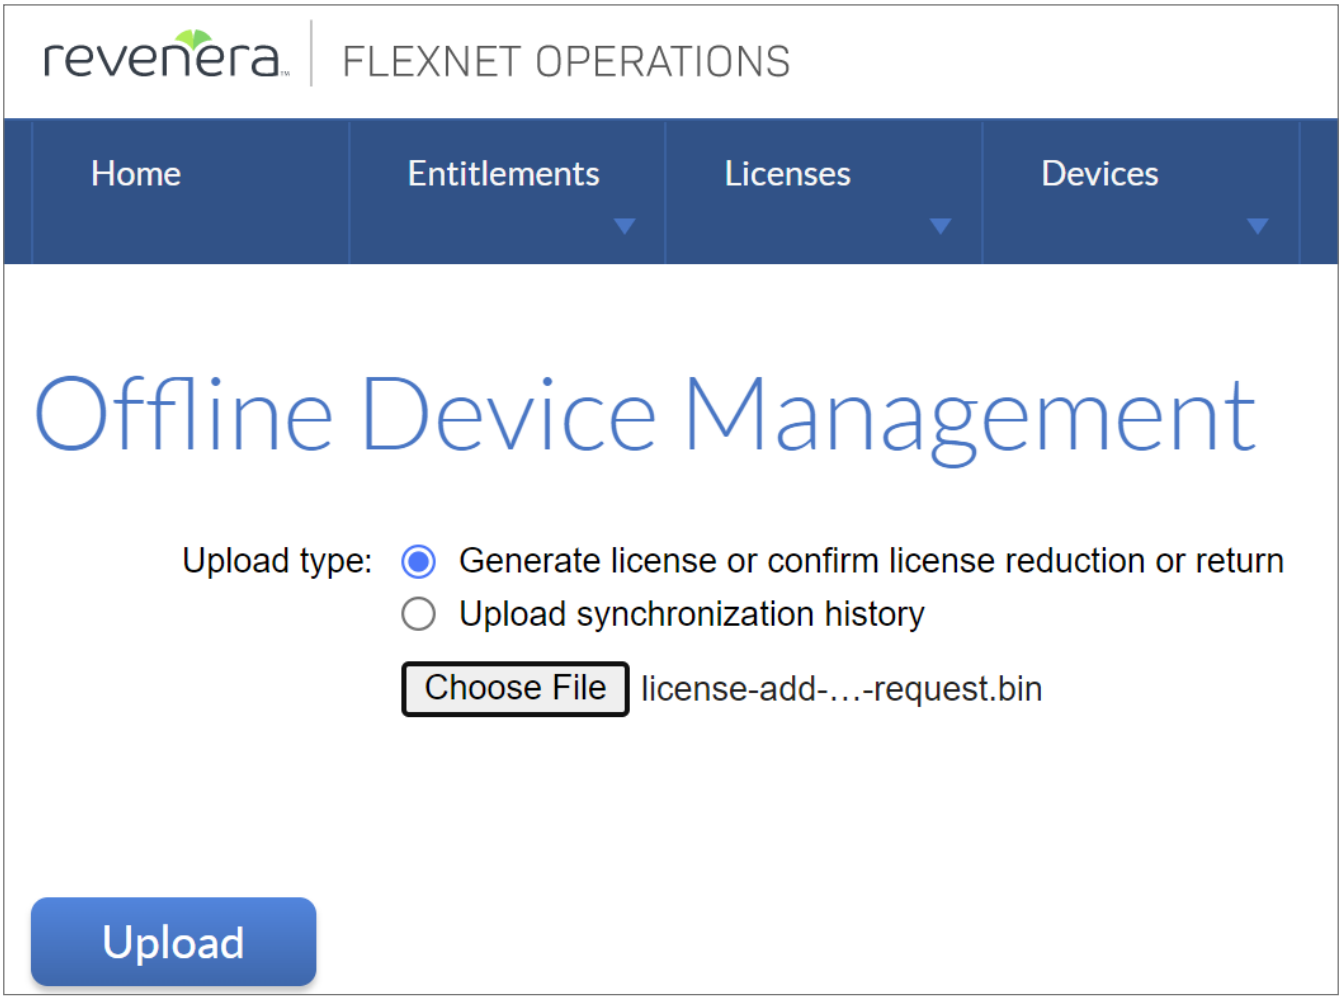

Figure 32. Selecting .bin file to upload

6. After clicking **Upload**, the license file is generated and can be downloaded by clicking on the **here** link.

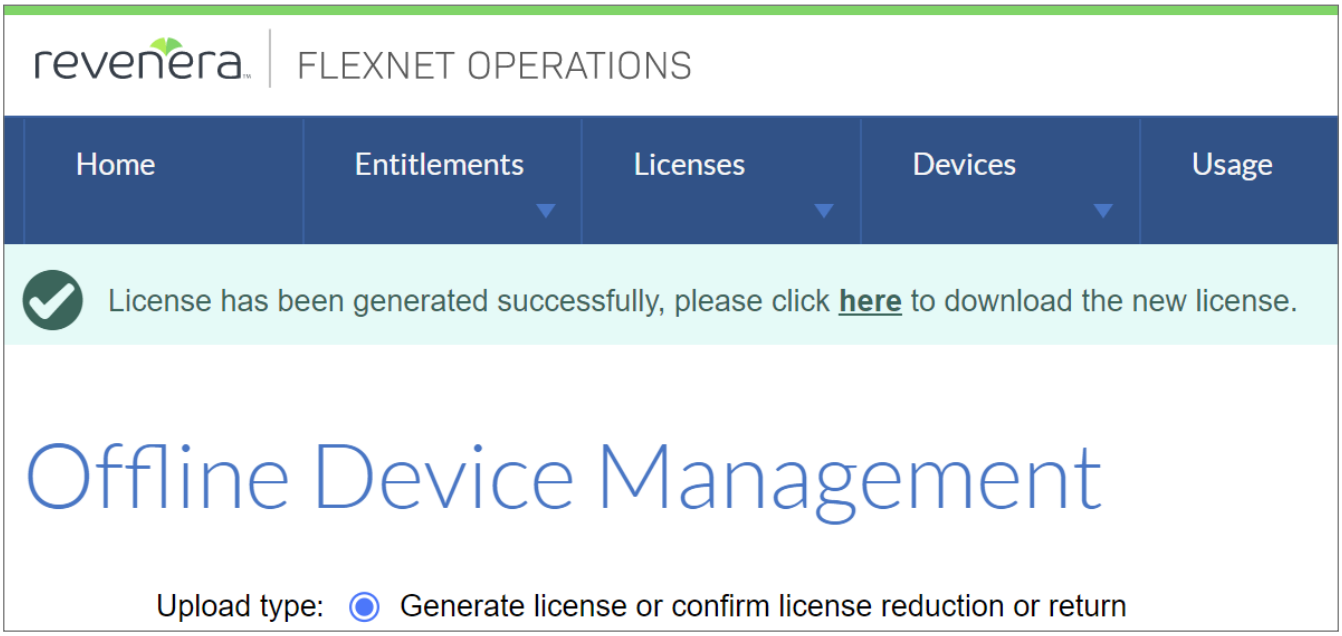

Figure 33. Click "here" link to download license file

7. Navigate back to the Licenses tab in the GUI and upload the license file by clicking on the ^ icon in the **Actions** column.

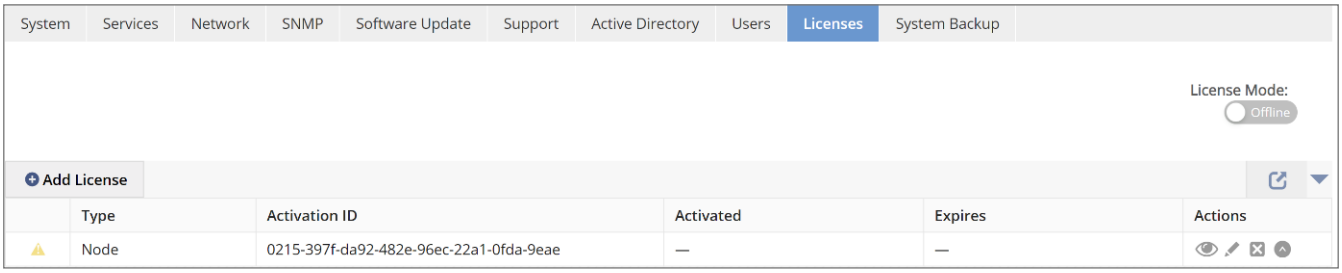

Figure 34. Uploading the license file

8. The system is now licensed until expiration date.

> To avoid disruption to data access, ensure a new license file is generated before the previous one expires.

8.3.2. Using the CLI

You can activate a license for offline deployments using the command line:

1. Mount the share that will be used to save the license information. This can be any share in the system except the root share, which is typically exported read-only.
2. Create a subdirectory for the license information and set its permissions to **777**. On Linux this can be done by issuing `chmod 777 <sub-directory>`.
3. In SMB, ensure everyone has full access to the directory. The license file must be stored in a subdirectory, not at the root of a share.
4. Execute the following command on the Anvil using a user with the **admin** role. The share name in this example is **data**. Make sure you add `/hs/` before the share name if using a file: path.

Command:

```
license-offline-add --type node --activation-id aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111 --export-uri
file:/hs/data/license/request.bin
```

Expected output:

```
The license request file has been exported to file:/hs/data/license/request.bin. Exchange the request
file with your license admin or at your license portal for a license response file, and then upload
the response file to complete the license add process.
```

The request file is written by the management layer, not by your user, which is why the subdirectory needs the permissions set in step 2. Until the response file is imported, `license-list` shows the activation as a pending offline request:

Command:

```
license-list
```

Expected output:

```
total 1
Activation ID:              aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111
Type:                       Node
Pending offline response:   Add
Status:                     LicenseStatus.PENDING
```

If the license has already expired, the share will not be accessible. There are two options:

- Manually enable the protocols for temporary use by going to **Administration** › **Services** tab in the Management GUI.
- After manually enabling the protocols, use the share and start at Step 1 above (RECOMMENDED APPROACH).

  If enabling the protocols fails, follow the workflow below. Use it only if the license is expired and the protocols cannot be enabled:

  Command:

  ```
  license-offline-add --type node --activation-id aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111 --export
  -uri file:/opt/pd/mgmt/webapps/webui/request.bin
  ```

  Expected output:

```
The license request file has been exported to file:/opt/pd/mgmt/webapps/webui/request.bin.
Exchange the request file with your license admin or at your license portal for a license
response file, and then upload the response file to complete the license add process.
```

You can navigate to the Management GUI and access the local directory by going to `https://<anvil-management-ip>/request.bin` in a browser. The **request.bin** file will be saved there and can be downloaded.

5. After the **request.bin** file has been generated and downloaded to the local desktop, log in to the Software Licensing and Delivery Portal and under **Devices**, select **Offline Device Management**. If you do not have access to the portal, contact your license administrator to execute the next steps.
6. Upload the **request.bin** file. This will generate a response file, **capabilityResponse.bin**, that can be downloaded from this page and used to activate the license.

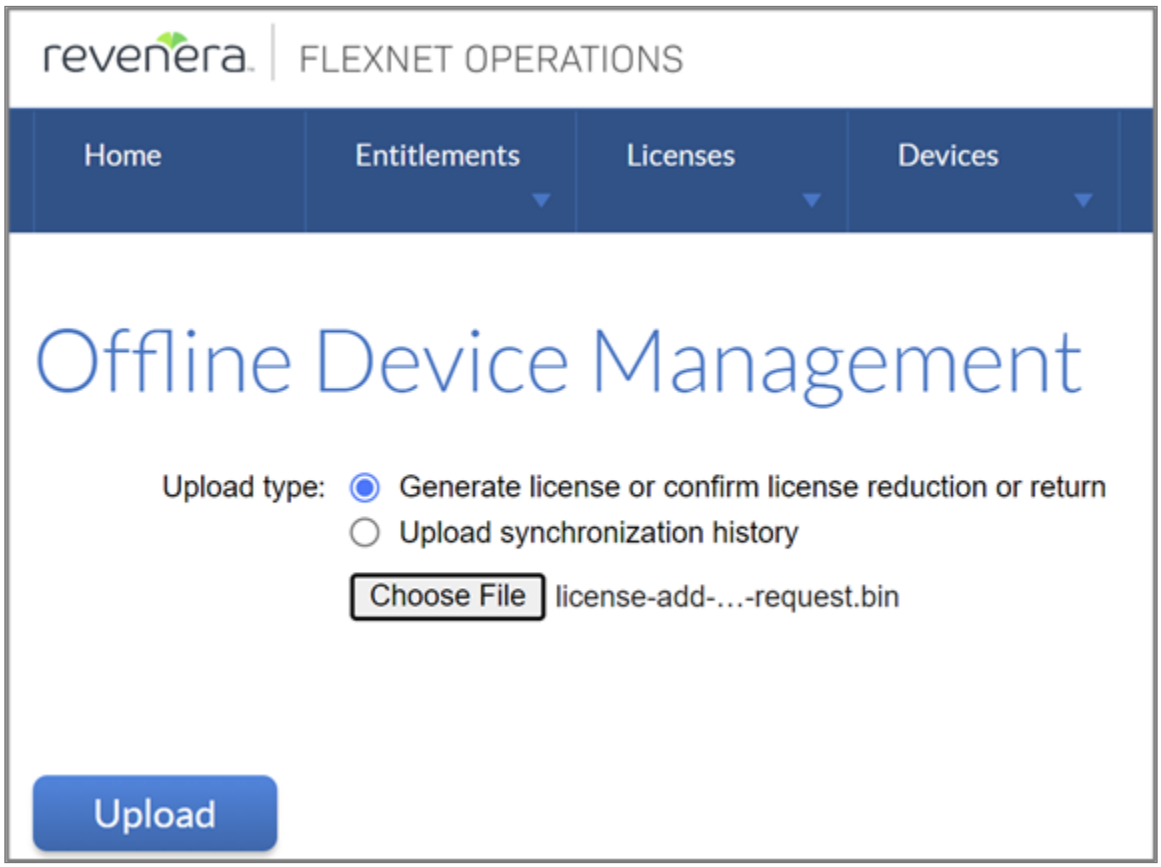

Figure 35. Upload request.bin file

7. Enable NFS/SMB protocols to be able to copy in the response file. Copy the `capabilityResponse.bin` file to the subdirectory under the share. This can be done using NFS or SMB.
8. Run the following command to activate the license (the share named **data** is used in the example below).

   Command:

```
license-offline-add --activation-id aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111 --import-uri
```

```
file:/hs/data/license/capabilityResponse.bin
```

Expected output:

```
Process started. To monitor progress, use: task-list --id bcc076a6-1bbd-4b1f-8a78-c14eaba4d999
```

9. Once the license-add task has completed, validate that the license was installed successfully by running the `license-list` command:

   Command:

   ```
   license-list
   ```

   Expected output:

   ```
   total 1
   Activation ID: aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111
   Type: Node
   Status: Valid
   Activated: 2026-01-03 17:34:03 UTC
   Expires: 2027-08-14 23:59:59 UTC
   Number of acquired licenses: 1
   ```

8.3.3. What Information Is in the Capability Response .bin File?

The capability response file .bin is encoded, making it difficult to read with standard text editors. The contents can be deciphered with the license library SDK (available upon request).

Below is an example of the contents of the **capabilityResponse.bin** file:

Example contents: capabilityResponse.bin

```
MessageType="Capability request"

Vendor=primary

RequestHostID=21173e83-783a-4deb-b09a-be246191b996

ResponseTime="Jul 10, 2025 10:13:10 AM"

ForceResponse=true

IdentityName=datasphere

DesiredRights=(ActivationID=aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111:1)
```

```
RequestHostIDType=String

MessageTime="Jul 10, 2025 10:16:01 AM"

MachineType=Virtual

IsServer=false

SourceIds=21173e83-783a-4deb-b09a-be246191b996

RequestOperation="Normal request"

ServerInstance=1
```

8.4. Updating an Offline License

An offline license installation requires a new license key to update the expired license. This is non-disruptive to the system and can be done at any time.

8.4.1. Using the GUI

1. Navigate to the **Licenses** tab in the Management GUI.
2. Click the **Edit** (pen icon) of the installed license key.
3. Download the request file and follow the process starting at Step 2 of the *Offline License Activation, Using the GUI* section.

8.4.2. Using the CLI

If using the Admin CLI, use the `license-offline-update` command and follow the same process, starting at Step 2 of the *Offline License Activation, Using the CLI* section.

8.5. Removing a License

Be aware that removing a license will disable protocol access to data if the initial grace period has expired.

8.5.1. Removing an Online License

Click the **Remove** icon or use the Admin CLI to remove the license. The product will contact the license service and automatically update the information in the system.

8.5.2. Removing an Offline License

Removing an offline license requires the same process as adding a license. After clicking the **Remove** icon in the Management GUI or running the Admin CLI command, `license-offline-remove`, download the request file and upload it to the Flexera Offline Device Management page.

Download the response file and apply it to the Product. This will remove the license and keep the licensing up to date in the Software *Licensing and Delivery Portal*.

8.6. License Troubleshooting

If you attempt to use a license that is already in use, you will see similar results to those shown below:

Command:

```
license-add --type NODE --activation-id aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111
```

Error:

```
license-add: An entity (license activation) named ' aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111 ' already
exists.
```

If you attempt to use an already expired license, you will see similar results to those shown below:

Command:

```
license-add --type NODE --activation-id aaaa-bbbb-cccc-dddd-eeee-ffff-0000-1111
```

Error:

```
license-add: License activation failed: The activation ID was denied, All features of entitlement, aaaa-
bbbb-cccc-dddd-eeee-ffff-0000-1111, are expired.
```

If you attempt to use a license with the wrong format, you will see similar results to those shown below:

Command:

```
license-add --type NODE --activation-id aaaa-bbbb
```

Error:

```
license-add: activationId invalid activation-id format
Try 'license-add --help' for more information.
```

Appendices

Appendix A: Terminology

The following list of terms and definitions may be helpful to users of this documentation. For more in-depth definitions of terminology relevant to the Hammerspace product, see the Dictionary of Terms on the Hammerspace web site.

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*. Login using Hammerspace account credentials 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.

Licensing and Delivery Portal

The following Product documentation titles are available on the *Hammerspace License and Delivery Portal*, Hammerspace Documentation 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

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 is an introduction to the concepts of Hammerspace metadata, and the tools to manipulate and use metadata and put it to work to achieve business objectives for data.
- Hammerspace S3 Guide [Solution Guide]: This document covers the creation and management of S3 servers, bucket containers, and buckets, as well as user access control and security considerations, ensuring seamless multi-protocol data access and efficient data management.