Hammerspace Azure Deployment Guide

Version 5.3

Table of Contents

About This Documentation

This guide describes how to deploy and configure Hammerspace in Microsoft Azure. It walks through deploying the Hammerspace Data Platform from the Azure Marketplace, validating the deployment, configuring the system through the Management GUI and Admin CLI, accessing shares over SMB, NFS, and Parallel NFS, and configuring DSX high availability with an Azure Load Balancer.

The guide is organized into four areas: deploying Hammerspace in the Azure Marketplace, configuration examples, client access over SMB and NFS, and a reference appendix covering terminology, high availability, and architecture.

This guide is intended for infrastructure engineers and storage architects deploying Hammerspace in Microsoft Azure. Readers should be familiar with Azure concepts such as resource groups, virtual machines, virtual networks, and load balancers, as well as general storage and file-sharing concepts.

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.

Part 1 — Deployment

Chapter 1. Introduction

This document describes how to deploy Hammerspace in the Microsoft Azure Marketplace. Hammerspace software is available in two consumption models: Pay-As-You-Go (PAYGO) and Bring Your Own License (BYOL).

1.1. Hammerspace BYOL (Bring Your Own License)

BYOL licensing is a software subscription model that is activated by contacting Hammerspace. This is often a suitable choice for customers who have a mix of cloud and on-premises deployments and a desire to consolidate licensing.

1.2. Hammerspace PAYGO (Pay-As-You-Go)

PAYGO licensing is a standard cloud consumption model in which usage is metered and billed to the cloud provider on an hourly, monthly, or yearly basis.

1.3. Solution Overview

The Hammerspace Data Platform consists of two main software components: a metadata node (Anvil) and a data services node (DSX).

1.3.1. Anvil

The following services are provided by the Anvil instances:

- Management GUI and CLI.
- Metadata management.
- Mount point for Parallel NFS (NFSv4.2).

1.3.2. DSX

The Data Services instances provide the following services:

- Global namespace.
- Global file system.
- Mount point for NFSv3, SMB3, and S3.
- Data mobility.
- Storage and volume management.

1.4. Hammerspace System Requirements

Anvil must be given sufficient resources to deliver low-latency performance for applications and end users. Anvil metadata nodes take full advantage of allocated resources. It is not unusual for Anvil to operate at 40% or

higher CPU utilization even with a relatively modest metadata workload in a large file system. Metadata performance may suffer if the system becomes resource constrained. Therefore, it is crucial to ensure that Anvil metadata nodes have sufficient system resources.

Listed below are the minimum system resource recommendations for Anvil and DSX in a bare minimum deployment. Production environments or environments requiring high performance need substantially more than the listed minimum recommendation.

Table 1. Minimum system resource recommendations

	Anvil metadata node	DSX data node
Azure virtual machine	2 x Standard E4as v4	1 x Standard F2s v2
CPU cores	16 vCPU	8 vCPU
RAM	32 GB (more recommended)	4 GB (more recommended)
Boot disk	200 GB (if memory is larger than 160 GB, multiply by 1.3 for sizing)	200 GB minimum
Additional disk drives for metadata storage	2 x 400 GB (minimum) backed by SSD or NVMe	As needed for DSX data storage
Network adapter	2 x 10 Gbps Ethernet. One link is dedicated to HA services	1 x 10 Gbps Ethernet

The choices of virtual machine sizes, as well as disk types and sizes, require careful planning. A metadata disk size of 200 GB accommodates approximately 25–50 million files. For small-to-medium environments or lab/demo scenarios, the default settings are typically adequate. For large or workload-intensive environments, consult with Hammerspace before deployment. This is particularly important where metadata performance is critical.

For environments with more than 500 clients using the SMB protocol to access the solution, a larger-sized virtual machine type and/or additional DSX nodes are strongly recommended. Additional Data Store disks can be added non-disruptively and without downtime post-deployment.

Chapter 2. Deploying Hammerspace in Azure

This section describes how to deploy the Hammerspace Data Platform from the Azure Marketplace, complete the post-deployment tasks, and validate the deployment using the Management GUI and the Admin CLI.

2.1. Deployment Steps

1. Open a web browser (Chrome recommended) and navigate to https://portal.azure.com to log in with your Azure account.
2. Navigate to the Marketplace.
3. Search for Hammerspace.
4. Select either the BYOL or PAYGO product listing.

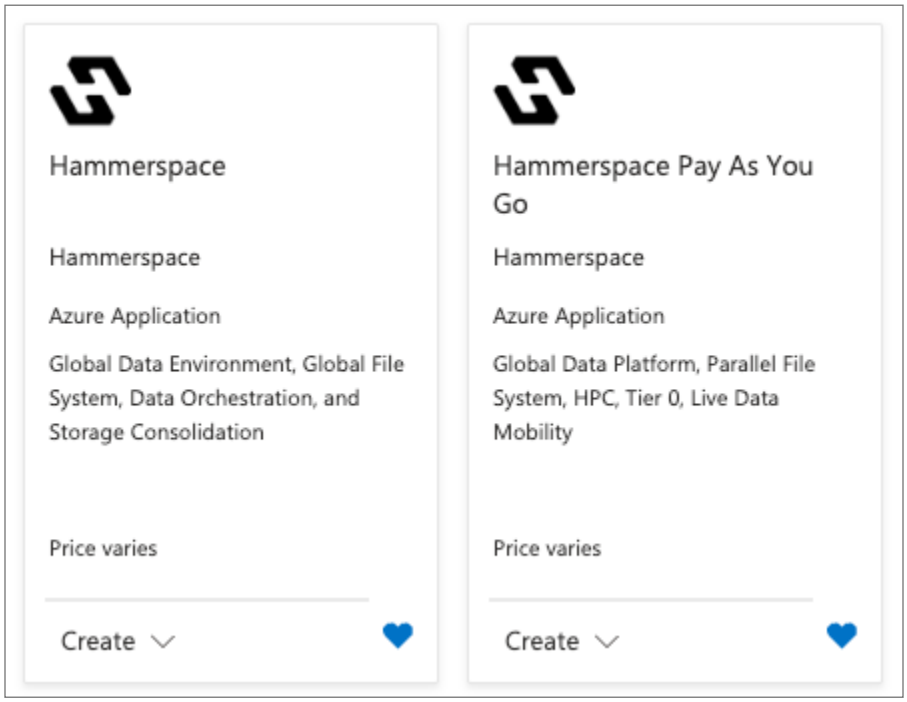

Figure 1. BYOL and PAYGO product listings in the Azure Marketplace

5. Select the subscription to use for the deployment.
6. Select **Hammerspace 5.2**.

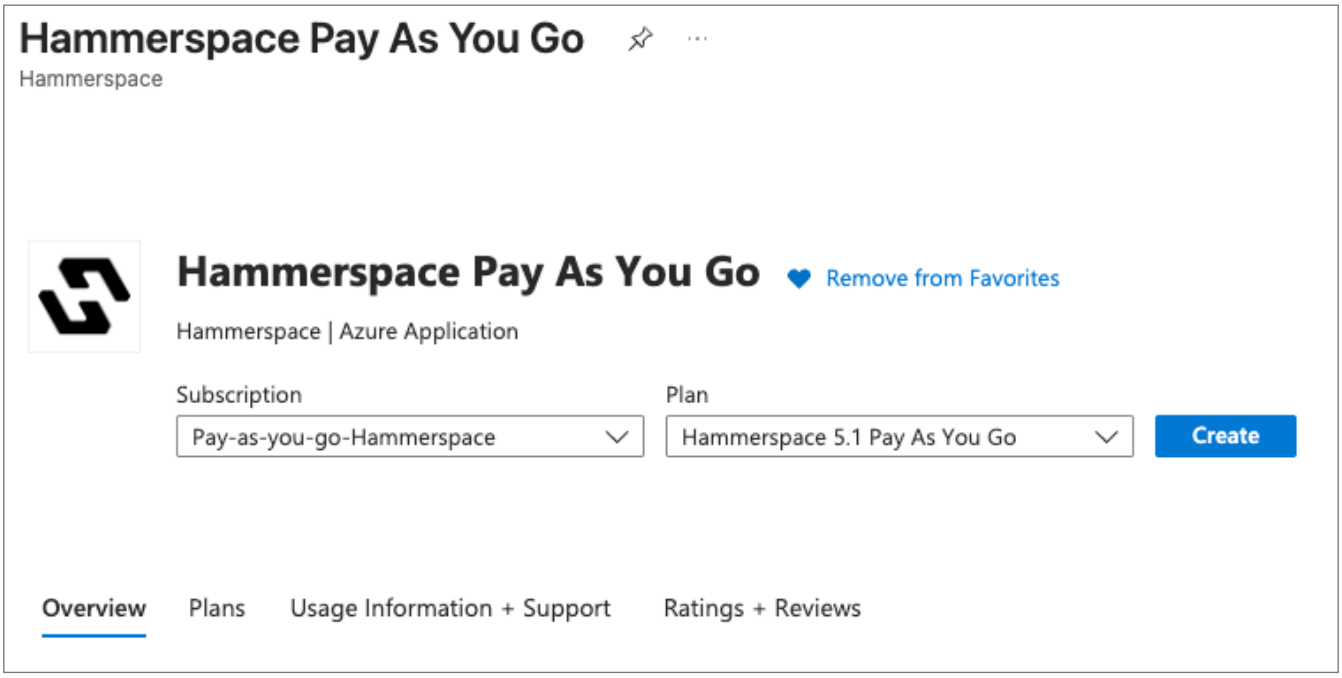

Figure 2. Selecting the Hammerspace 5.2 offer

7. Click **Create** to begin the five-step deployment.

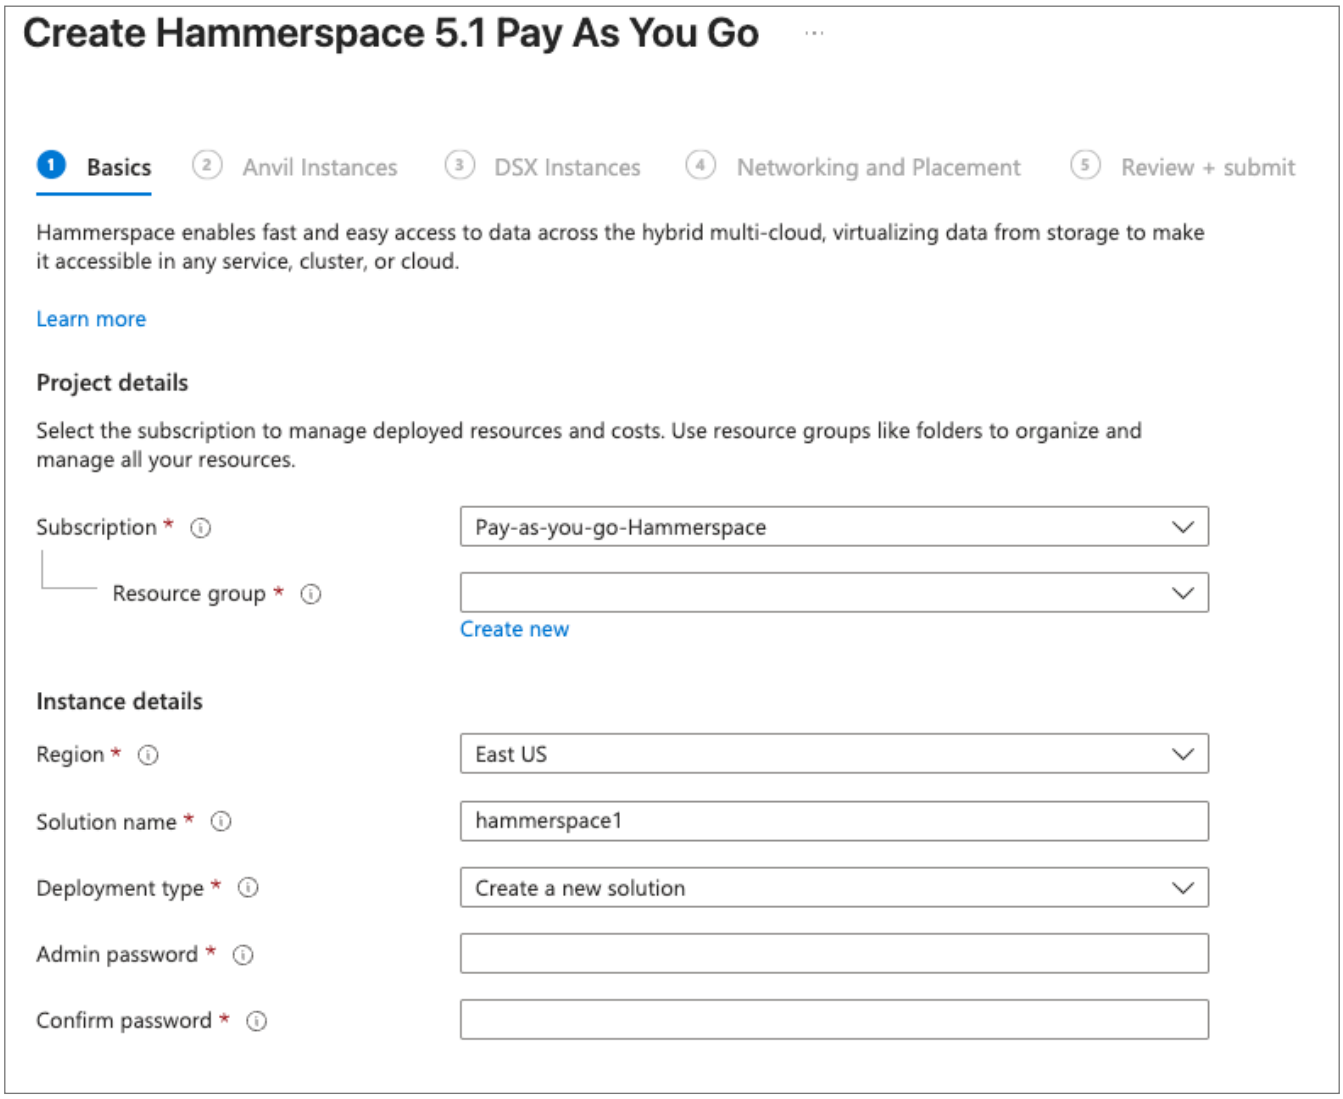

Figure 3. Starting the deployment

Table 2. Settings for the Basics tab

Setting	Value
Project details	
Subscription	Select the desired subscription.
Resource group	Click **Create new** and provide a name in the pop-up menu.
Instance details	
Region	Select a region from the dropdown menu.
Solution name	Give the deployment a name.
Deployment type	Select **Create a new solution** (unless you are adding DSX nodes).
Admin password	Create a password.

Setting	Value
Confirm password	Confirm the password.

8. Click **Next** at the bottom of the page.

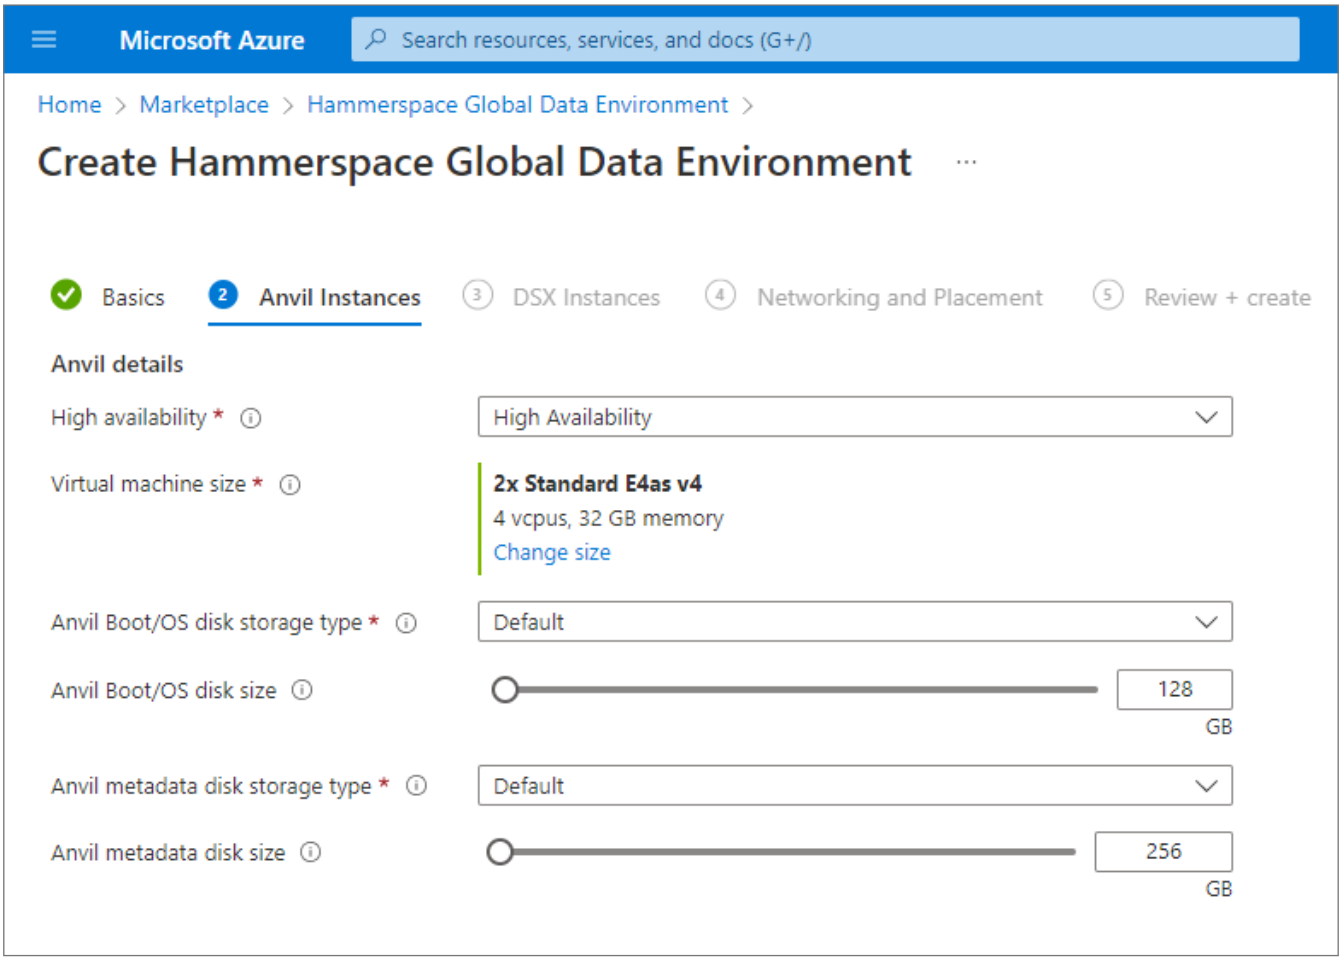

Figure 4. Settings for the Anvil Instances tab

Table 3. Settings for the Anvil Instances tab

Setting	Value
High availability	Select **High Availability** or **Standalone** from the dropdown menu.
Virtual machines	Accept the default selection or click **Change size**.
Boot/OS disk type	Select the disk type for the system binaries and libraries.
Boot/OS disk size	200 GB recommended minimum.
Metadata disk storage type	Select the disk type to suit performance and cost requirements.
Metadata disk size	200 GB recommended minimum.

9. Click **Next** at the bottom of the page to proceed to the DSX Instances tab.

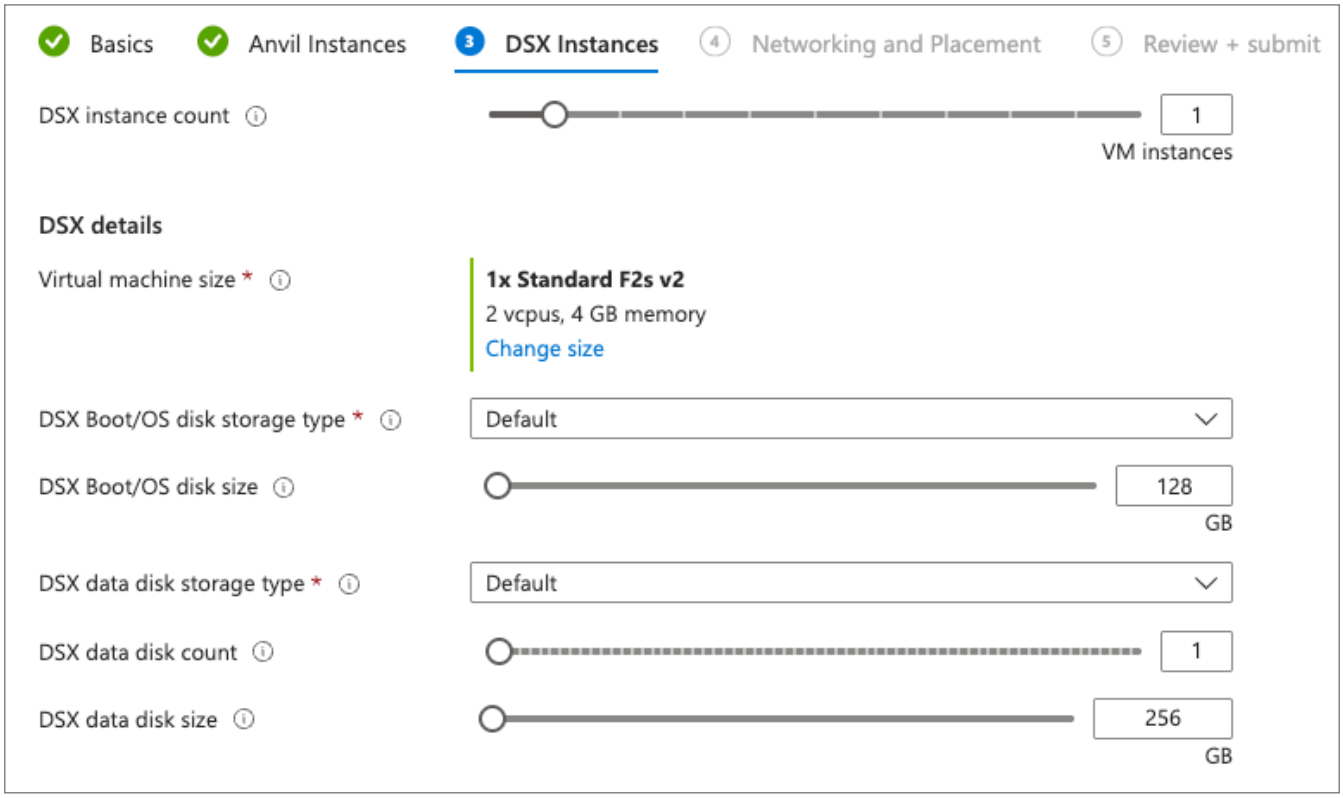

Figure 5. Settings for the DSX Instances tab

Table 4. Settings for the DSX Instances tab

Setting	Value
DSX instance count	Select the number of DSX VM instances to deploy.
DSX details	
Virtual machine size	Accept the default selection or click **Change size**.
Boot/OS disk type	Select the disk type for the system binaries and libraries.
Boot/OS disk size	100 GB recommended minimum.
Data disk storage type	Select the disk type to suit performance and cost requirements.
DSX disk size	500 GB recommended minimum.

10. Click **Next** to proceed to the Networking and Placement tab.

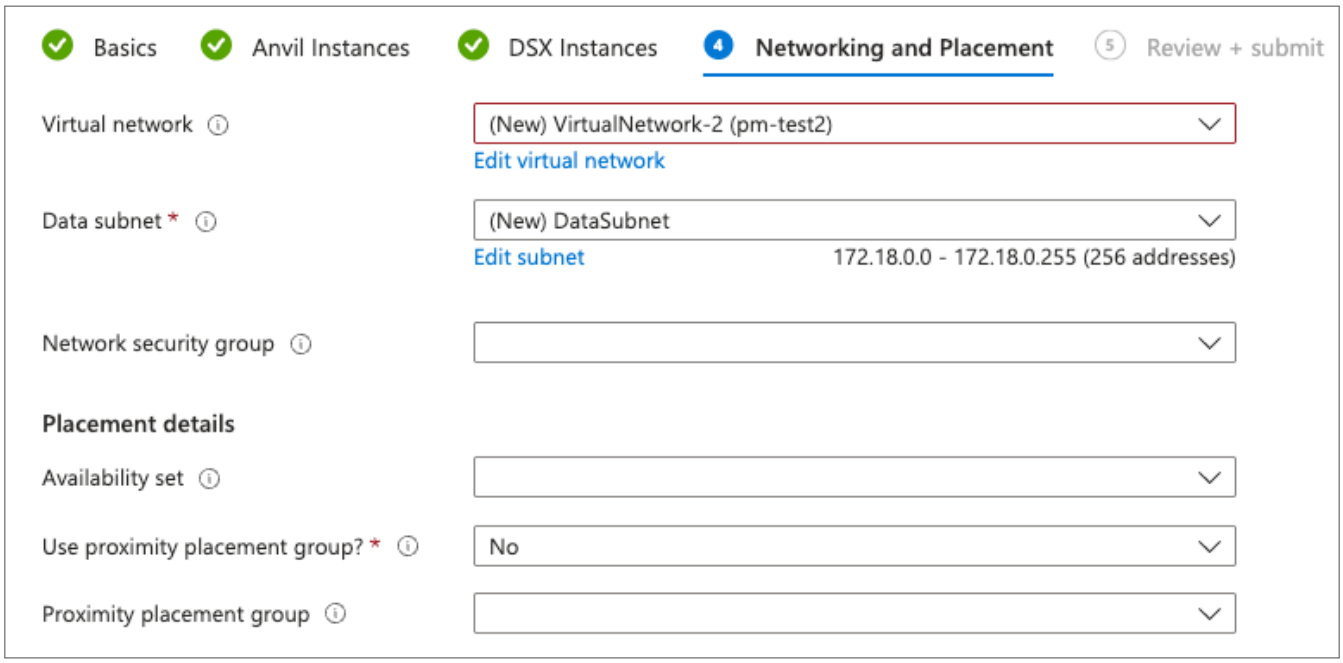

Figure 6. Settings for the Networking and Placement tab

Table 5. Settings for the Networking and Placement tab

Setting	Value
Virtual network	Select from the dropdown menu.
Data subnet	Select from the dropdown menu.
Network security group	Select from the dropdown menu or leave blank.
Placement details (optional)	
Availability set	Select an existing availability set or leave blank to create a new one.
Use proximity placement group	Optional.
Proximity placement group	Optional.

11. Click **Next** to review the configuration.
12. Click **Create** (or use the **Previous** button to make edits).

2.2. Post-deployment Tasks

Once the deployment completes, you should see the following screen.

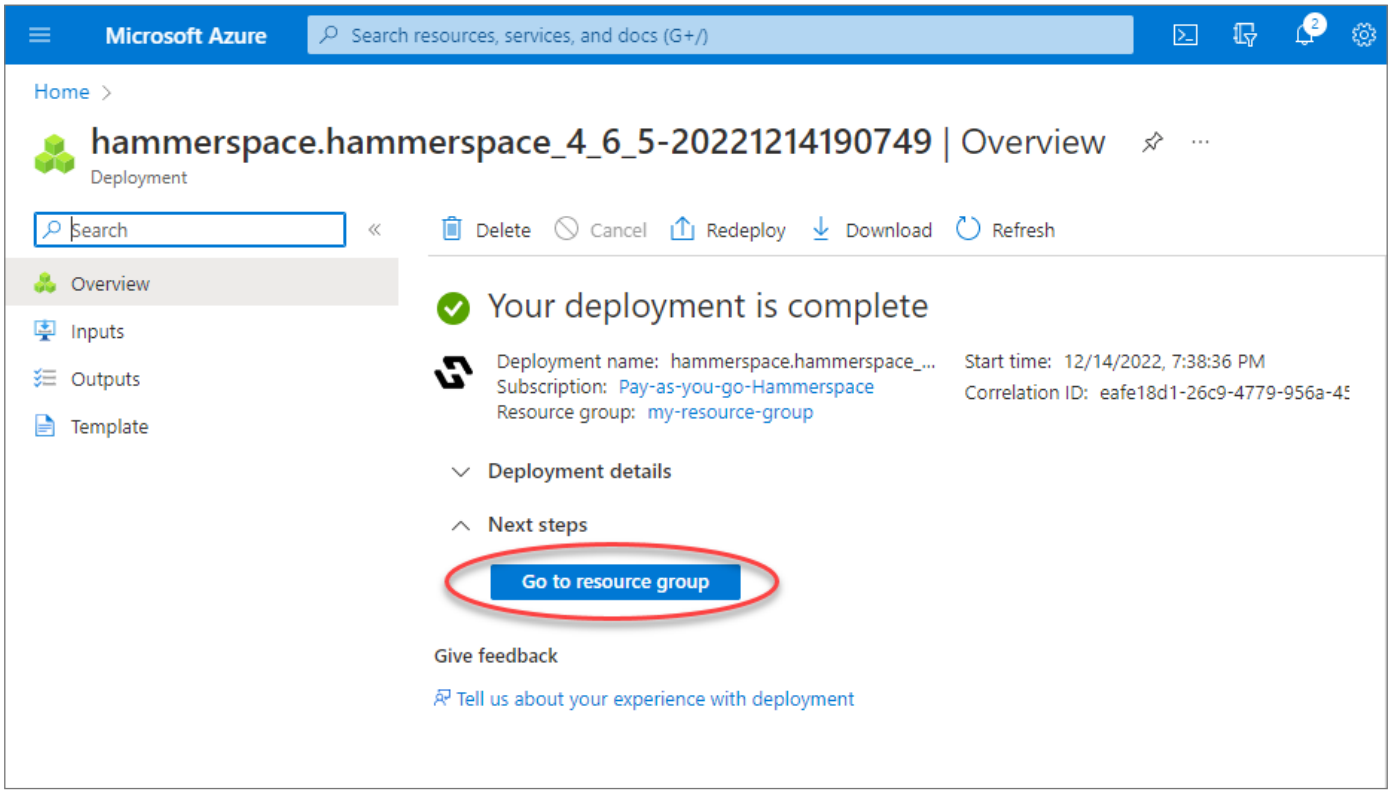

Figure 7. Deployment complete

1. Click **Go to resource group** to proceed.

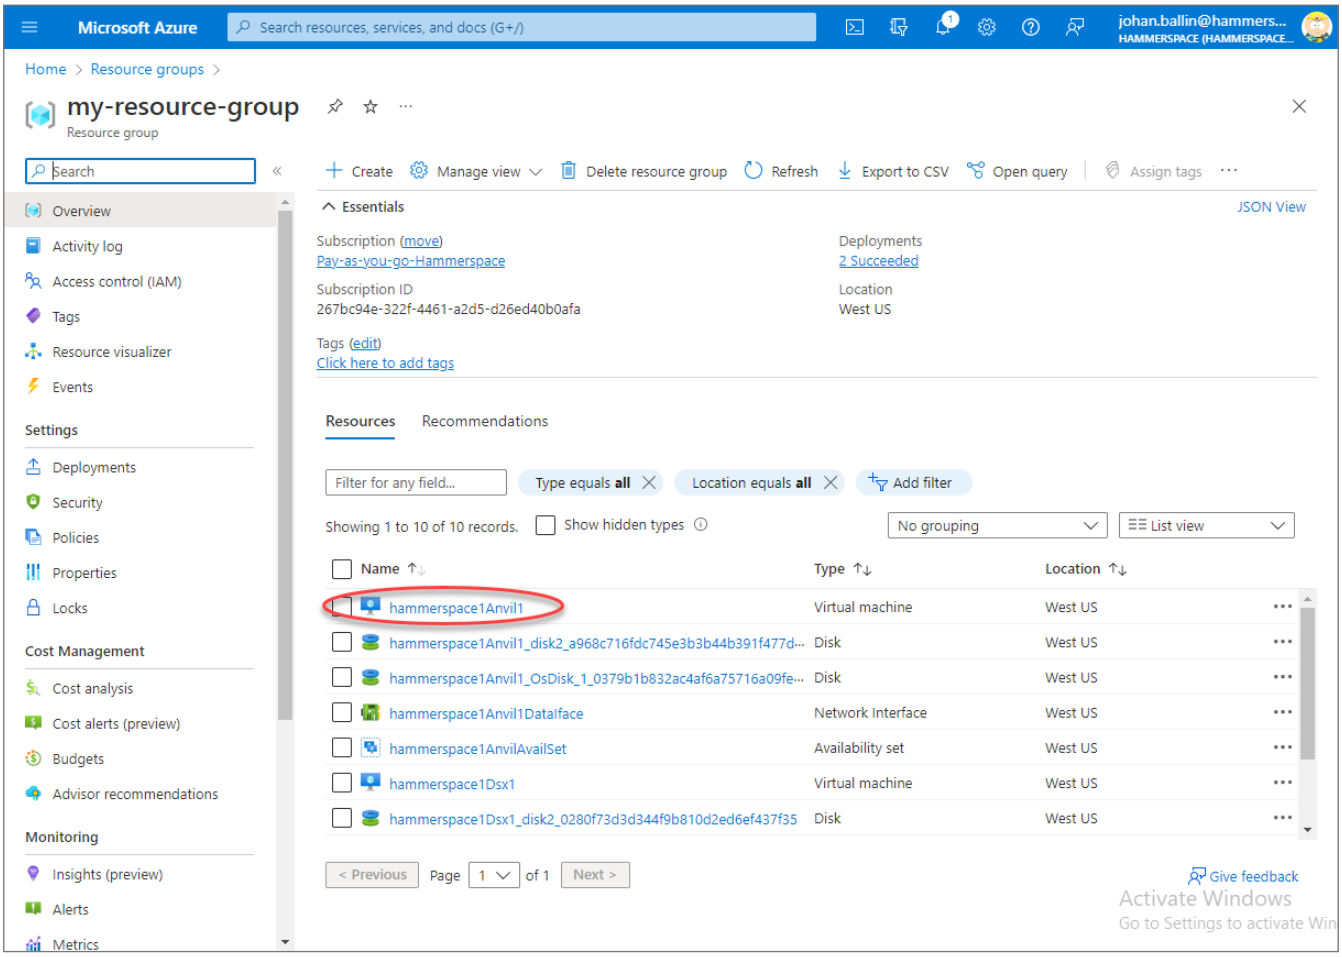

Figure 8. Resource group

2. Click the Anvil hyperlink in your resource group to obtain the IP address of the Management GUI and CLI.

2.3. Hammerspace Management GUI

Point your web browser (Chrome recommended) to the IP address obtained from the Anvil virtual machine in your resource group. The IP address alone is usually sufficient, as Hammerspace redirects to HTTPS on TCP port 8443 — for example, `https://10.41.1.6:8443`.

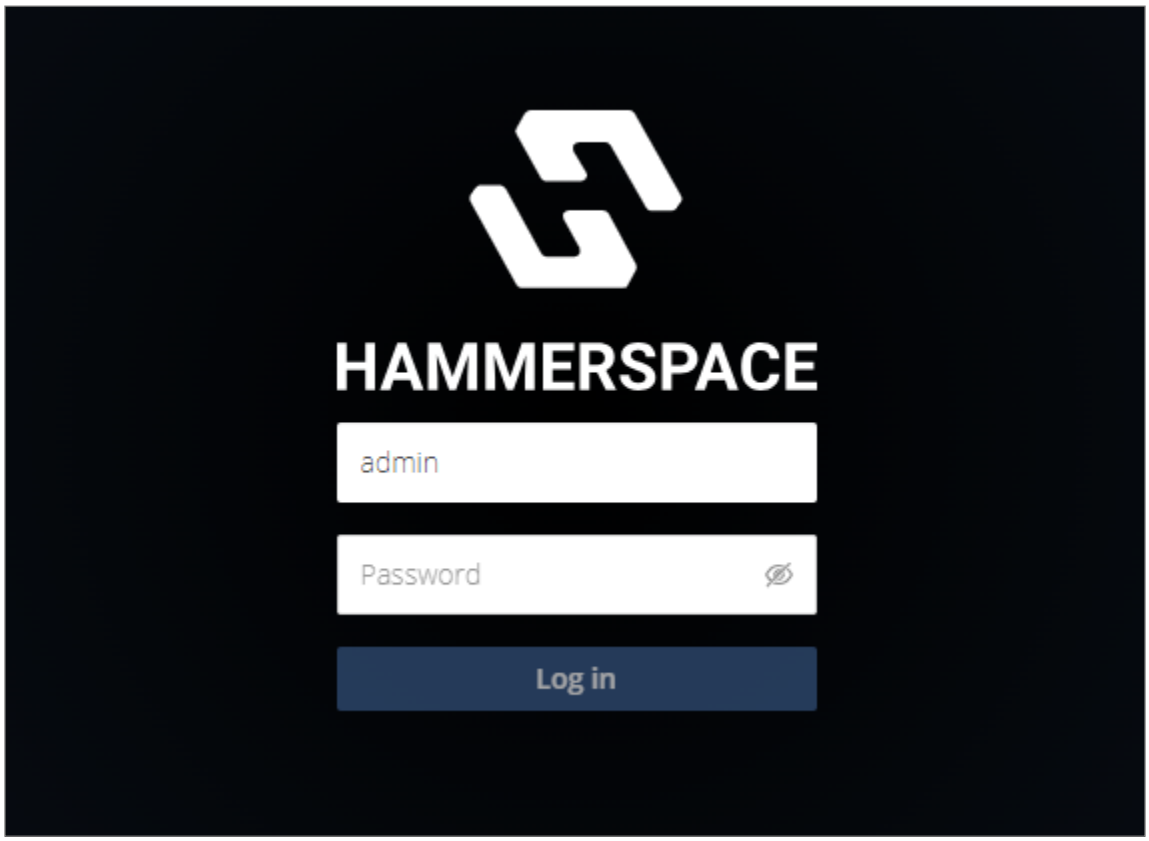

Figure 9. Management GUI login screen

Enter the username and password in the form fields:

- **Username**: `admin`
- **Password**: the password you created in the **Basics** tab during deployment.

When you initially log in, a pop-up window prompts you to accept the Hammerspace End User License Agreement (EULA). This occurs only once.

2.3.1. Validating the Deployment with the Management GUI

From the gear icon in the left navigation pane, select the **Network** tab. All Anvil and DSX nodes provisioned through the deployment template should be visible here.

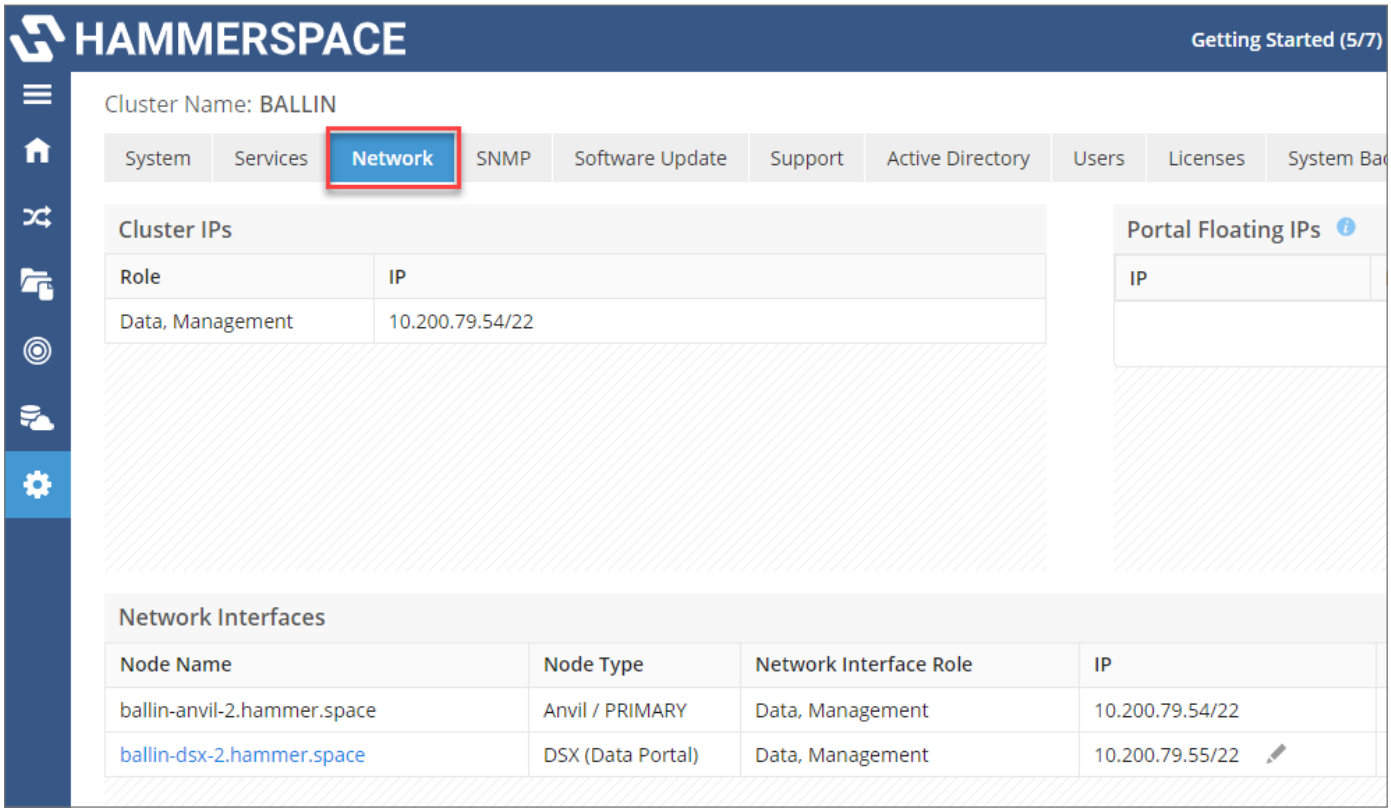

Figure 10. Network tab showing provisioned Anvil and DSX nodes

Take note of the DSX IP addresses and/or hostnames. The DSX nodes should contain the key term **Data Portal**. This is the mount point for all front-end protocols, such as SMB (v1, 2, and 3), NFS (v3.0 and 4.2), and the Kubernetes Container Storage Interface (CSI driver).

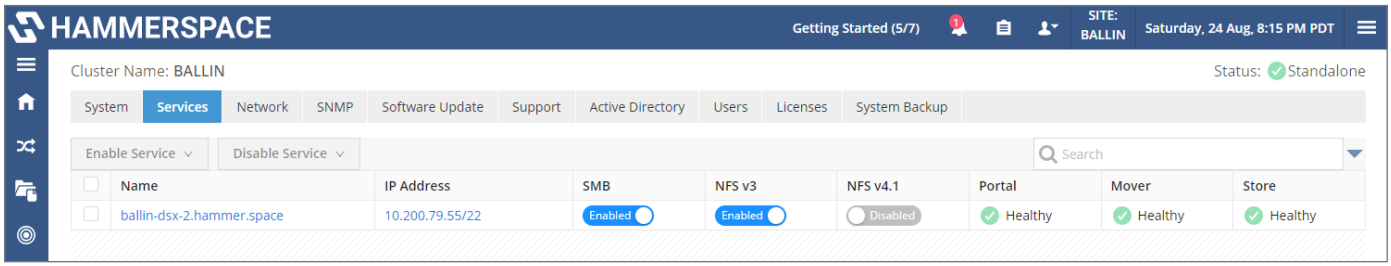

Figure 11. DSX Data Portal entry

The **Services** tab validates that the Data Portal, Data Mover, and Store, as well as the SMB and NFS protocols, are up and running and ready to use.

2.4. Hammerspace Admin CLI

The Admin CLI provides command-line control of system configuration and provides parity with the Management GUI. It is a restricted shell, also known as a CLISH. To access the Admin CLI, open a terminal session (or use an SSH client) to connect over SSH.

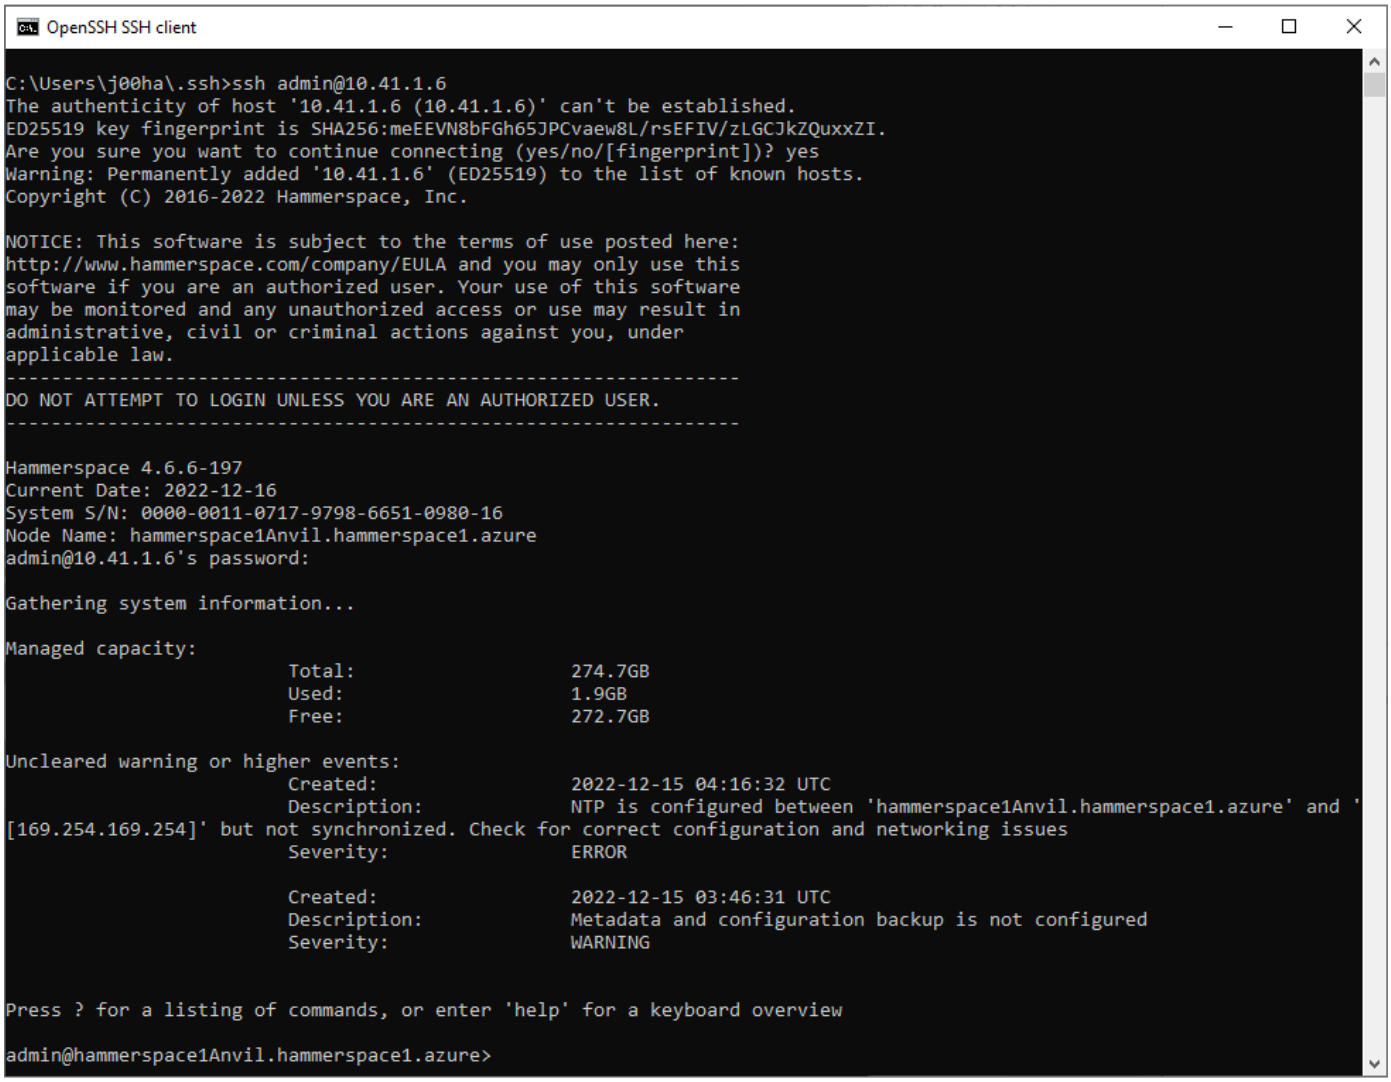

Figure 12. Admin CLI accessed over SSH

If using a terminal client, the command is:

```
$ ssh admin@<anvil-hostname-or-ip-address>
```

The terminal window also notifies you if the NTP configuration is erroneous or missing. Correct it by using the Hammerspace Admin CLI to enter the following commands:

```
ntp-config --servers-clear
ntp-config --servers <ntp-server>
```

You can replace the value with an NTP server or server pool of your choice. The examples use the NTP Pool Project. Choose the option that best suits your organization.

Part 2 — Configuration

Chapter 3. Configuring Hammerspace in Microsoft Azure

The following examples describe common configuration tasks after deploying Hammerspace in Microsoft Azure, using both the Management GUI and the Admin CLI.

3.1. Changing the Admin User Password

Use the following procedure to reset the **admin** user password.

1. Navigate to **Administration**.
2. Select the **Users** tab.
3. Click **Reset Password**.

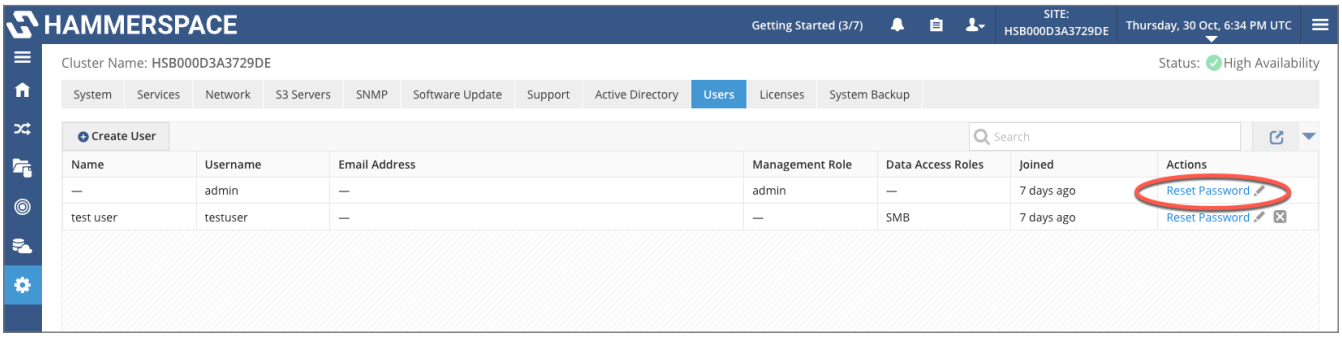

Figure 13. User management

3.2. Configuring DNS Server Settings Using the Admin CLI

View the DNS server settings:

```
dns-config --view
```

Remove the DNS server configuration:

```
dns-config --servers-clear
```

Add the DNS server configuration:

```
dns-config --servers 192.168.1.101,192.168.2.102
```

3.3. Configuring NTP Server Settings Using the Admin CLI

View the NTP server settings:

```
ntp-config --view
```

Remove the NTP server configuration:

```
ntp-config --servers-clear
```

Add the NTP server configuration:

```
ntp-config --servers 0.north-america.pool.ntp.org
```

3.4. Creating a Share

Use the following procedure to create a share.

1. From the left navigation pane, select **Data** (the file-folder icon).

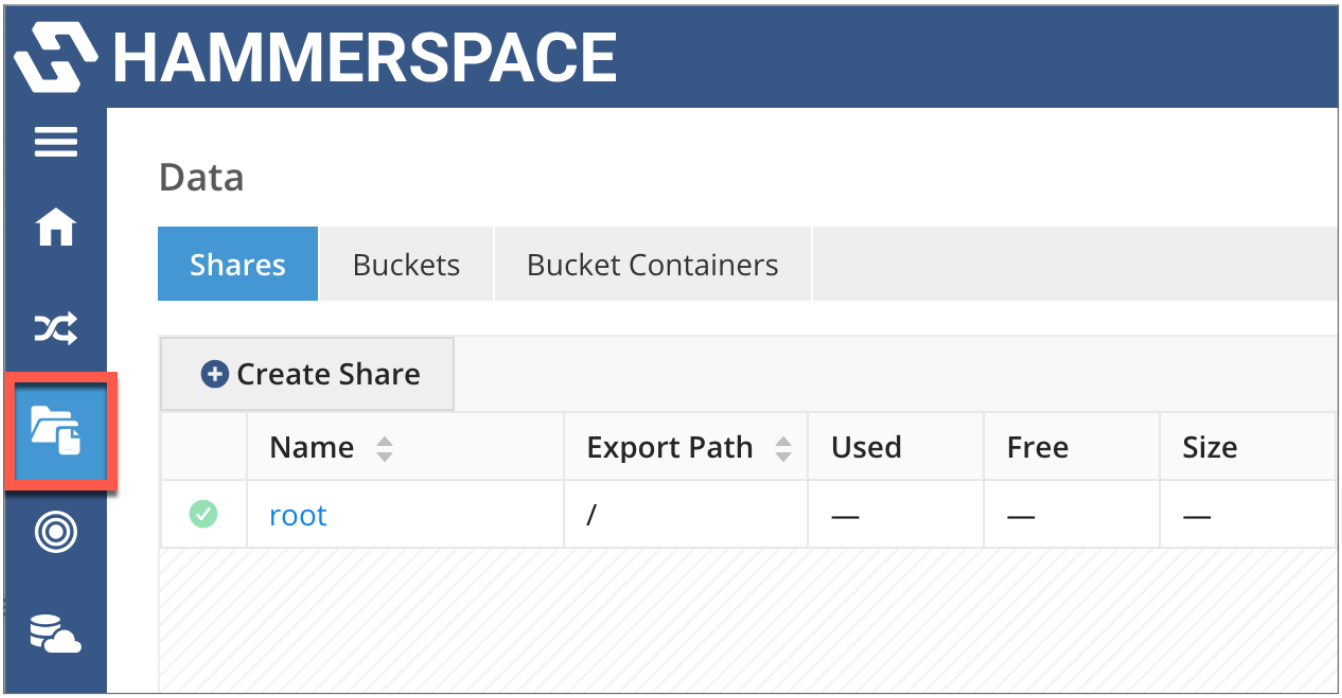

Figure 14. Location of the Data icon

2. Enter share details in the Create Share wizard.

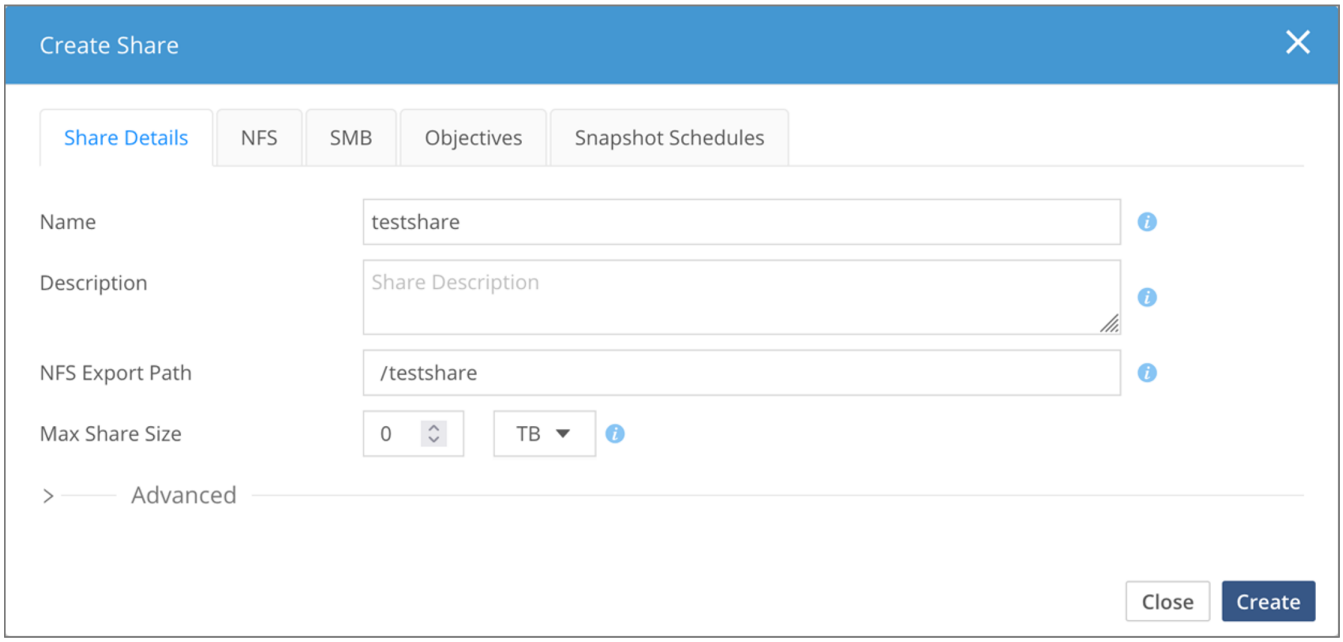

Figure 15. Create Share wizard

3. In the **Share Details** tab, enter the information requested.

Table 6. Share details

Item	Value
Name	Name the share.
Description	Optional description of the share.
NFS Export Path	The NFS export path for the root of this share. Must start with `/`.
Max Share Size	Implements a soft quota. A zero (0) denotes no quota.
Advanced (optional)	See the Hammerspace Licensing and Delivery Portal for the Administration Guide.
ABE	Lists files and directories based on the user's permissions.
Access time update	Unchecking disables access-time updates.
Retain deleted files	Check to retain deleted files across the share.
Create versioned files	Check to implement event-based file retention.
File conflict resolution	See the Administration Guide for more information.
WORM	Implements Write Once Read Many on the share.

4. If your share is an NFS share, enter the information requested in the **NFS** tab.

Table 7. NFS details

Item	Value
Client Specification	IP, CIDR subnet, hostname, subdomain, or `@netgroup`.
Permissions	Read-Only or Read-Write.
Root Squash	If checked, maps the NFS client `root` (uid 0) to the anonymous uid.
Secure Port	Unselecting this option disables the requirement that NFS client connections originate from a privileged source port (< 1024).

5. If your share is an SMB share, enter the information requested in the **SMB** tab.

Table 8. SMB details

Item	Value
Alias	SMB share alias name.
Path	SMB alias path.
Browsable	Controls whether the share is browsable by clients.

6. Add objectives. Objectives provide powerful declarative control over your data throughout its lifecycle. Key aspects of data management — such as availability, durability, placement, movement, protection, workflow, and custom criteria — can be controlled with ease at a file-granular level. Refer to the *Hammerspace Administration Guide*, available in the Hammerspace Licensing and Delivery Portal, and the *Hammerspace Objectives Guide*, available in the Hammerspace Support Hub, to learn more about how to use objectives.
7. Add snapshot schedules. Snapshots can be scheduled directly upon share creation from this tab, by either creating a new schedule, adding a pre-defined schedule, or copying an existing snapshot schedule from another share.

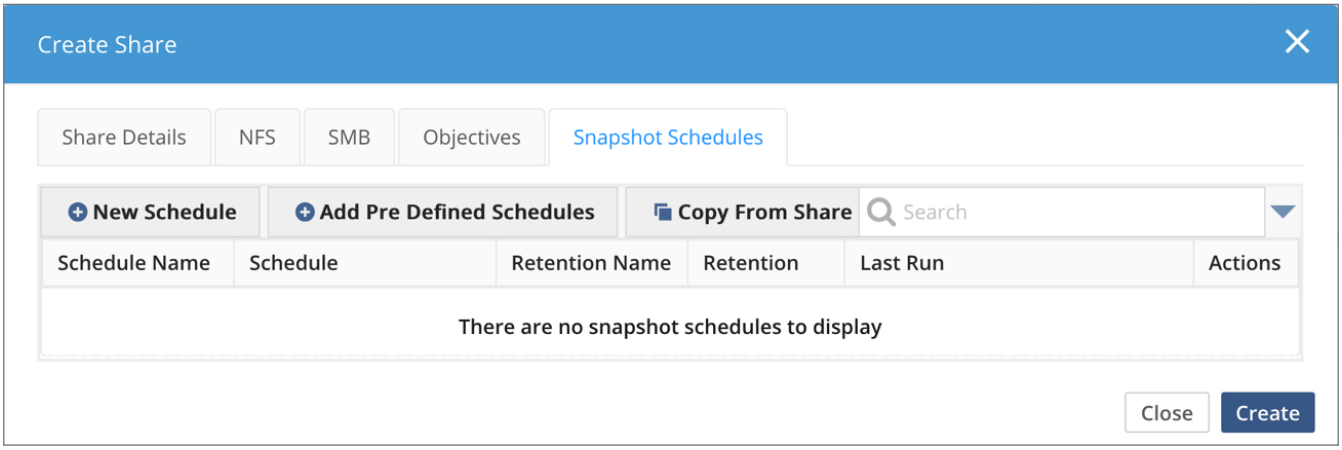

Figure 16. Adding a snapshot schedule for a share

3.5. Adding Azure Blob as a Storage System

1. Go to the **Infrastructure** > **Storage Systems** tab in the Hammerspace Management GUI.

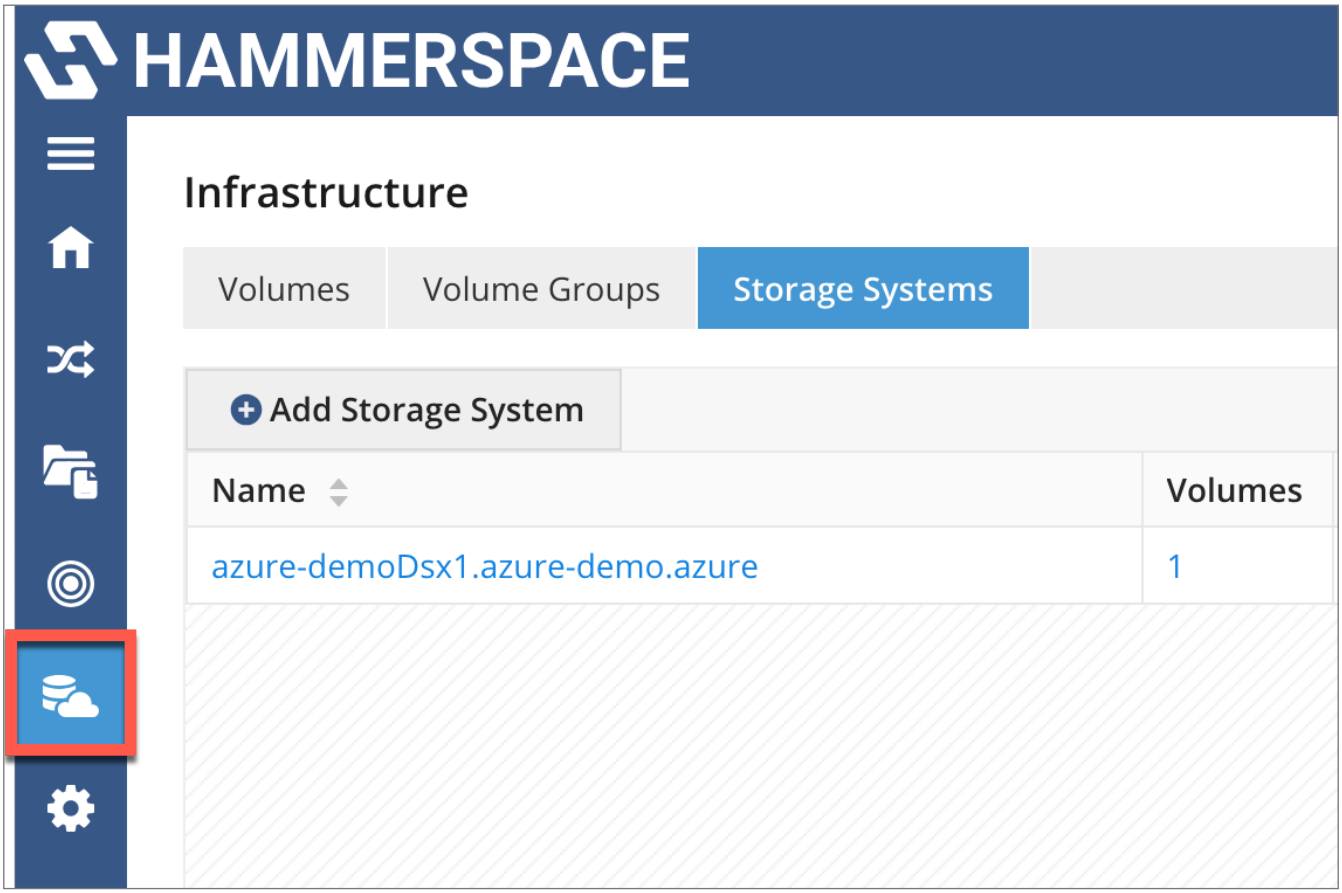

Figure 17. Storage Systems tab

2. Select **Storage Type** and enter the requested information.

Table 9. Storage-type values

Item	Value
Storage Type	Object Storage.
Type	Azure.
Storage Account Name	Enter the storage account name.
Secret Key	Enter the secret key.
Endpoint	Optionally, enter endpoint and port information.
Proxy Host	If a proxy is used.
Proxy Port	If a proxy is used.
Proxy Username	If a proxy is used.
Proxy Password	If a proxy is used.

3.6. Creating Volumes

1. Go to the **Infrastructure** > **Storage Systems** tab in the Hammerspace Management GUI. Click the **+ Volume** button to open the **Add Volumes** wizard.

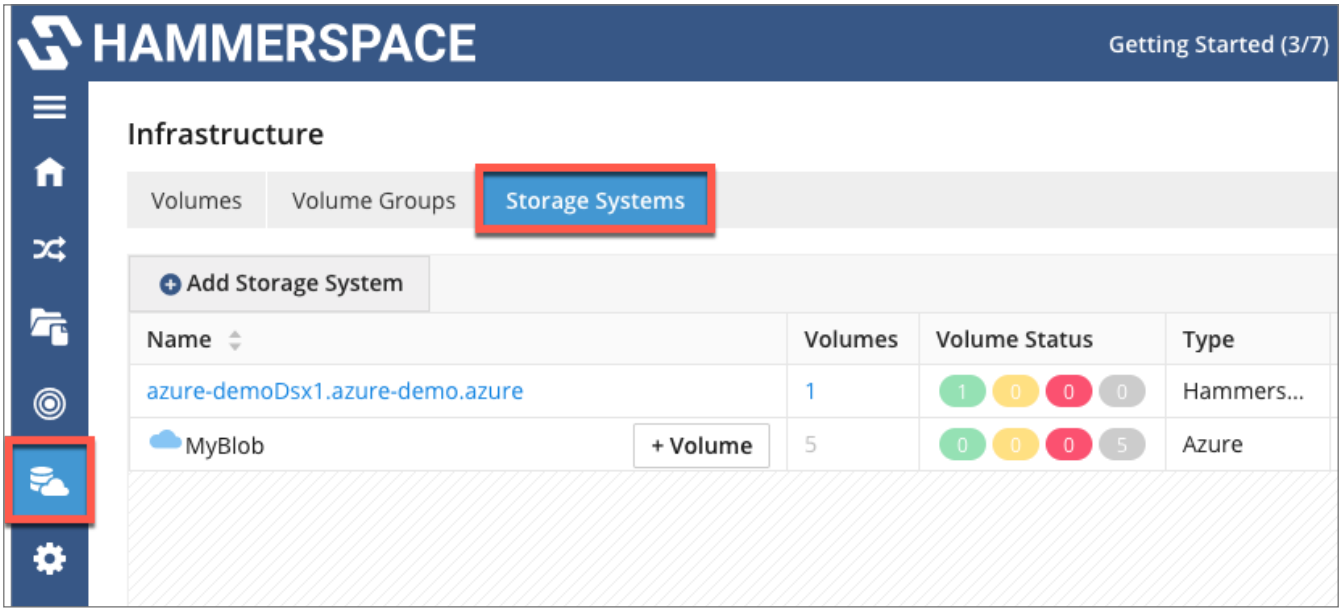

Figure 18. Location of the + Volume button

2. Select the object storage containers you wish to add as Hammerspace volumes.

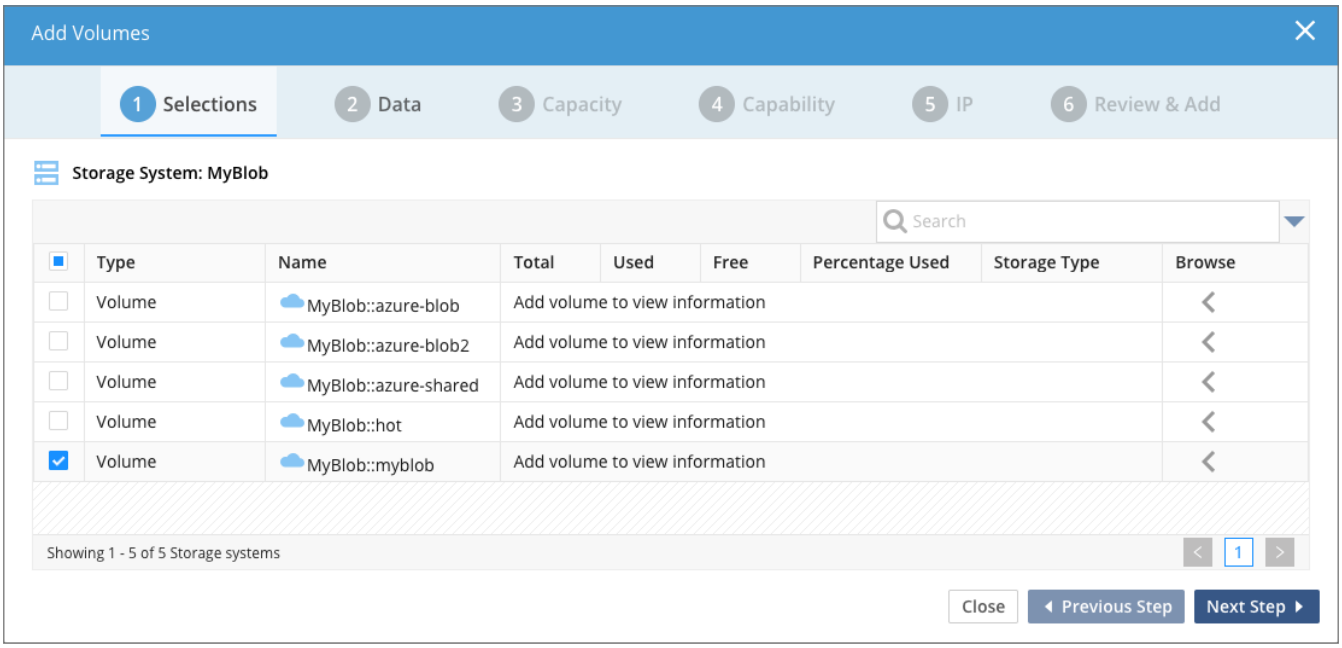

Figure 19. Selecting volumes

3. Click through each tab of the wizard, using the volume-configuration settings below.

Table 10. Volume configuration settings

Item	Value
Shared Volume	Uncheck (leave the checkbox empty).
Capacity	Default: Auto (choose a size to apply a quota).
Durability	Default: Auto.
Availability	Default: Auto.
Objectives	Check to automatically create associated objectives.

4. Review your selections and, if satisfied, click **Add Volume** to finalize.
5. Navigate to the **Volumes** tab to view the newly created volume.

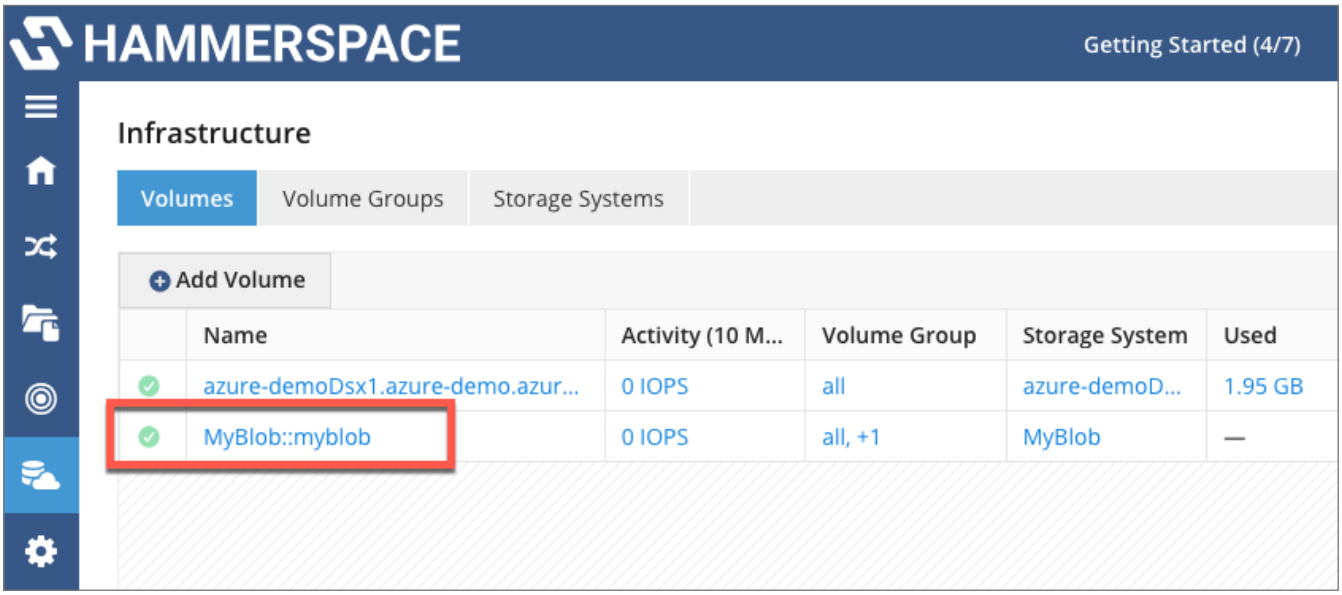

Figure 20. Newly added volume in the Volumes tab

3.7. Local User Configuration

Navigate to **Administration** by clicking the gear icon at the bottom of the left-side navigation pane. Select the **Users** tab, then click **+ Create User** and enter the details described below.

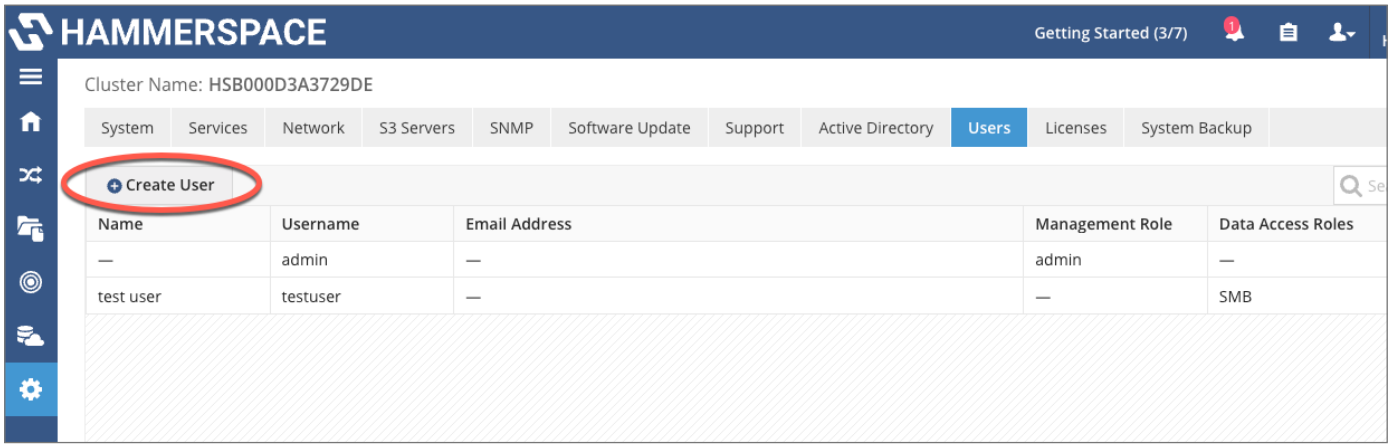

Figure 21. Location of the + Create User button in the Management GUI

Table 11. Create-user details

Item	Value
First Name	Enter the first name (optional).
Last Name	Enter the last name (optional).
Email	Enter an email address (optional).
Username	Create a name for the user (or service).
Password	Enter a password.
Confirm Password	Confirm the password.
Role Type	Select the **Data Access** radio button.
Data Access Role	Select the SMB and/or S3 checkboxes.

3.8. Active Directory Integration

1. Go to **Administration** (gear icon) > **Active Directory** tab.

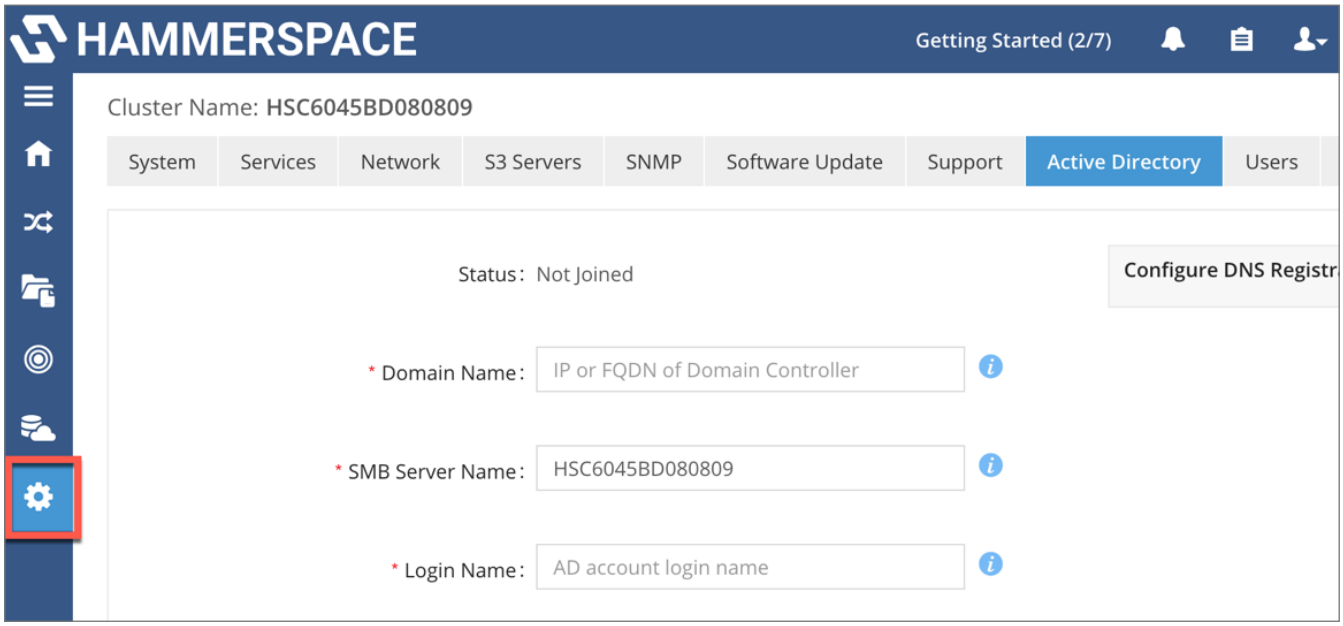

Figure 22. Location of the Active Directory tab in the Management GUI

2. Enter the requested information using the table below as a guide.

Table 12. Active Directory details

Item	Value
Domain Name	Enter the IP address or FQDN of the domain controller.
SMB Server Name	Enter the name desired for AD account registration.
Login Name	Enter the AD account name.
Password	Enter the AD account password.
Organizational Unit	Optionally add the OU to which this machine should be added.
AD Servers	Select Auto or Manual discovery of AD servers.
SMB–NFS User Mapping	Select the desired user-mapping option.

3. Optionally, add and configure Windows DNS registration.

3.9. Upgrading the System

1. Navigate to **Infrastructure** (gear icon), then select the **Software Update** tab.

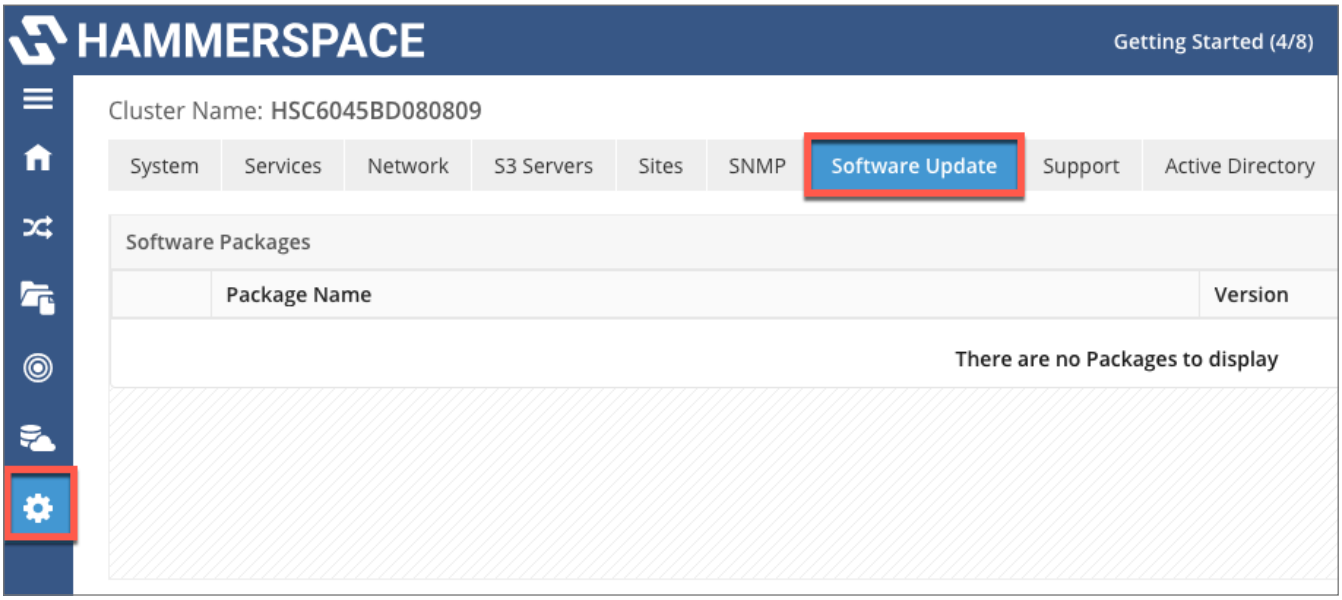

Figure 23. Software Update tab

2. Click the **Upload Software Package** button.

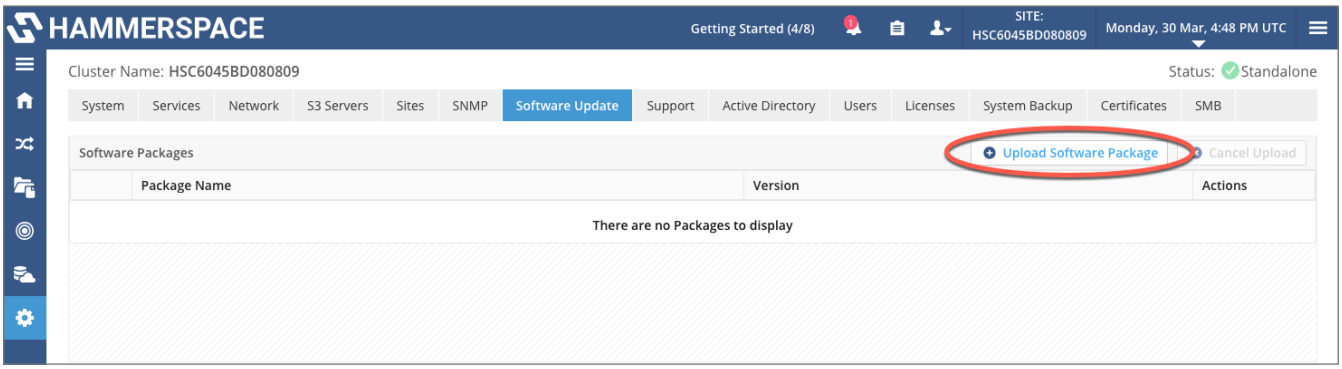

Figure 24. Upload Software Package

3. Use the radio button to select how to upload the software upgrade package to the system.

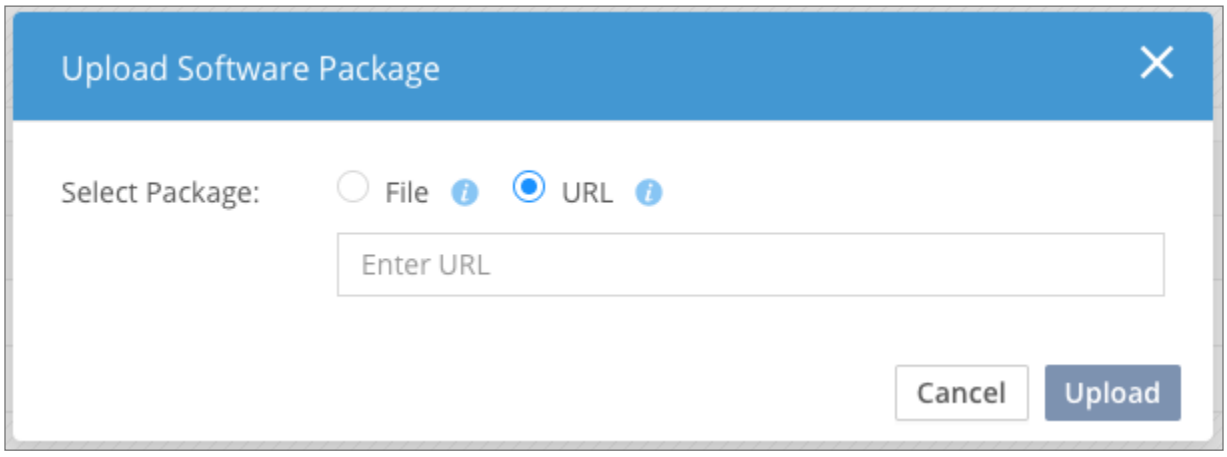

Figure 25. Select Software Package

Software upload packages can be obtained from Hammerspace. They can be identified by the `.upd` file-type extension.

4. Once the software package has uploaded, select its radio button and click the **Update All** button.

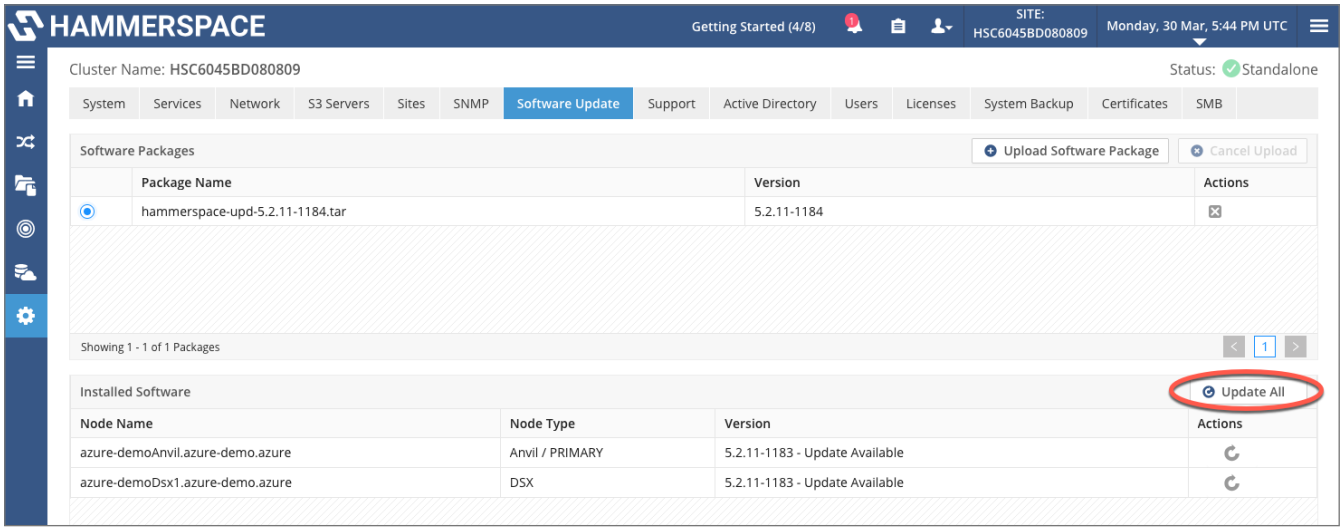

Figure 26. Update Nodes

Part 3 — Client Access

Chapter 4. Accessing Shares from Clients

The Hammerspace Global File System is fully multiprotocol. Applications and end users can simultaneously access shares over SMB and NFS from Windows, macOS, and Linux clients, as described below.

4.1. Accessing Shares from Windows

1. Open Windows File Explorer, right-click the computer icon, and select **Map network drive...**.

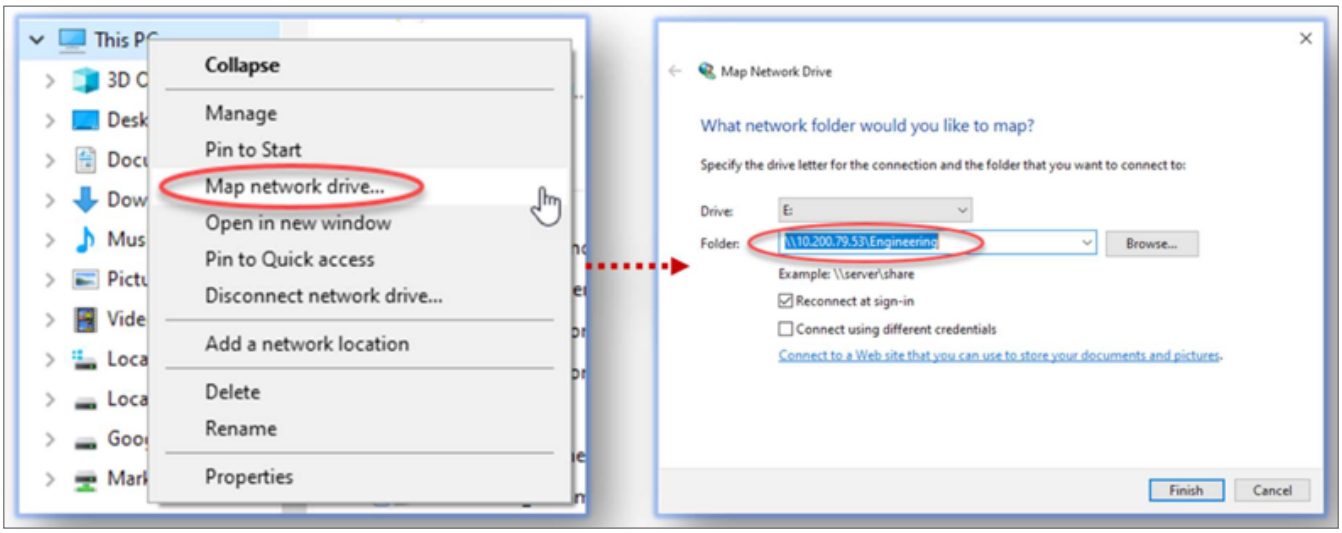

Figure 27. Location of Map network drive

2. Enter the IP address or hostname followed by the share name, and select the appropriate options.

4.2. Accessing Shares from macOS

On macOS, use the Finder **Connect to Server** dialog to mount a share from a DSX node.

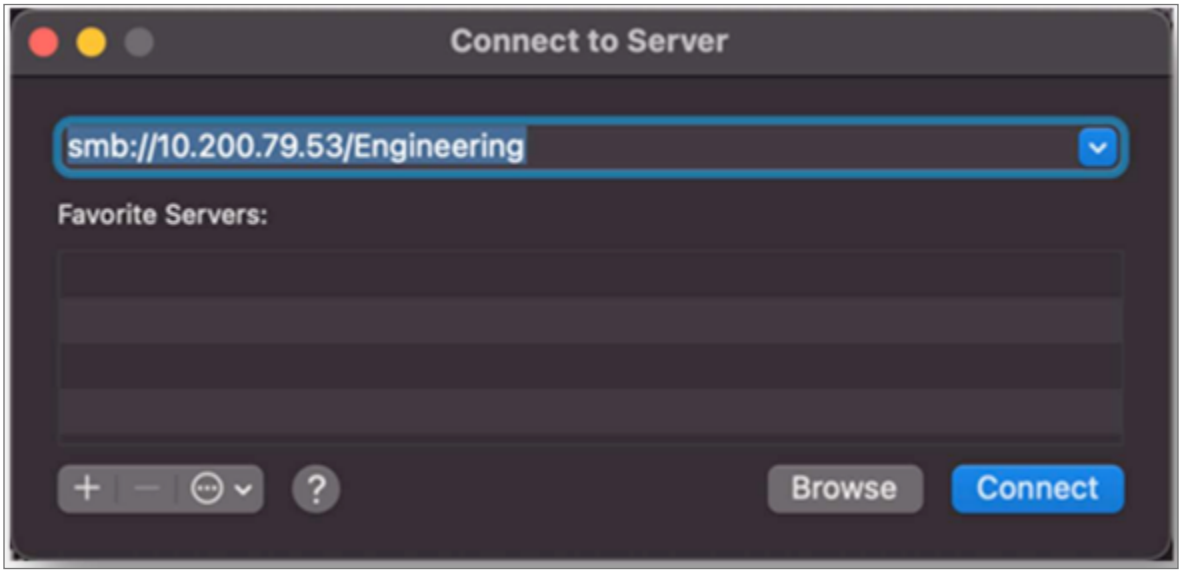

Figure 28. macOS Connect to Server dialog

4.3. Mounting NFS Exports from Linux

Depending on the Linux distribution, there may be a preference for using the `mount` command with options or the `mount.nfs` command, as shown below.

```
# mount.nfs <DSX-node>:/<export-path> <mount-path>
```

Or:

```
# mount -t nfs -o vers=3 <DSX-node>:/<export-path> <mount-path>
```

Some client systems may require additional NFS options, such as `nolock` and/or `rw`. Consult the documentation for each distribution.

4.4. Using Parallel NFS (v4.2) for High-Performance Workloads

Mounting shares using NFSv4.2 is slightly different from NFSv3 and SMB. NFSv4.2 exports are available from the Hammerspace metadata server (Anvil), while NFSv3 and SMB exports are available from any of the DSX nodes.

4.4.1. How to Mount Anvil Using NFSv4.2

```
# mount -t nfs -o vers=4.2 <anvil-node>:/<export-path> <mount-path>
```

NFSv4.2 is a more modern protocol with significant performance improvements. Only relatively recent Linux versions (for example, RHEL/CentOS 7.5 and later) have the most recent fixes. For some older Red Hat/CentOS Linux versions, Hammerspace has done the work to backport the upstream fixes.

Most Linux versions available in the cloud are capable of mounting NFSv4.2. For more information, see RFC 5661 and RFC 8434, or contact Hammerspace at support@hammerspace.com.

Appendices

Appendix A: Definition of Terms

The following terms may be helpful to users of this documentation.

Term	Definition
Anvil	Hammerspace metadata node. This virtual machine contains the administrative GUI and CLI as well as metadata, and performs multiple critical functions in a Hammerspace deployment.
Availability set	A logical grouping of VMs that allows Azure to understand how your application is built to provide redundancy and availability. Each availability set can be configured with up to three fault domains and twenty update domains. These configurations cannot be changed once the availability set has been created. Update domains indicate groups of virtual machines and underlying physical hardware that can be rebooted at the same time. See the Microsoft documentation on Availability Sets for more information.
Availability zone	Azure availability zones are connected by a high-performance network with a round-trip latency of less than 2 ms. They help your data stay synchronized and accessible when things go wrong. Each zone is composed of one or more datacenters equipped with independent power, cooling, and networking infrastructure. Availability zones are designed so that if one zone is affected, regional services, capacity, and high availability are supported by the remaining two zones.
DSX	Hammerspace data services node (Data Service eXtensions). This virtual machine provides the protocol front-end services, such as SMB, NFS, and Kubernetes CSI. It also controls backend connections and services to third-party storage volumes and data mobility services, and can optionally provide directly attached storage.
Global Data Environment	Provides a parallel global file system with file-granular declarative control of data. Hammerspace objectives provide highly customizable metadata tagging and descriptors which, in turn, can be used in declarative actions that allow organizations to control the business outcomes of their data. See the Hammerspace website for more information.
Mover	Also known as the Data Mover. Provides non-disruptive and end-user-transparent data mobility. Hammerspace administrators can move live data declaratively between backend storage systems and volumes without any disruption to productivity.
Portal	Hosted by the DSX node, this is the mount point for all front-end protocols (SMB, NFS, and the Kubernetes CSI driver) that make up the Hammerspace Global File System.
Proximity placement group	A logical grouping used to make sure that Azure compute resources are physically located close to each other. Proximity placement groups are useful for workloads where low latency is a requirement. See the Microsoft documentation for additional information.

Term	Definition
Region	Each Azure region features datacenters deployed within a latency-defined perimeter. They are connected through a dedicated regional low-latency network. This design ensures that Azure services within any region offer the best possible performance and security.
Resource group	A logical container that holds related resources for an Azure solution. The resource group can include all the resources for the solution, or only those resources that you want to manage as a group. Generally, add resources that share the same lifecycle to the same resource group so you can easily deploy, update, and delete them as a group.
Store	The directly attached storage disks on DSX, such as the ones provisioned during the deployment process.

Appendix B: DSX High Availability Configuration

Microsoft Azure networking does not support Gratuitous ARP. Therefore, you must configure an Azure Load Balancer to support high availability. DSX high availability in Microsoft Azure requires configuring the following items at a minimum:

- An Azure Load Balancer, including:
  - A frontend IP address.
  - A backend pool.
  - A health probe.
  - A load-balancing rule.

B.1. RBAC Permissions Required

To configure an Azure Load Balancer, the primary Azure Role-Based Access Control (RBAC) role required is **Network Contributor**. This role provides the permissions needed to manage network resources, including creating, configuring, and managing Azure Load Balancers.

For a more granular approach, you can create a custom role with only the specific permissions needed for load-balancer operations, adhering to the principle of least privilege. This can be achieved by including actions such as `Microsoft.Network/loadBalancers/*`.

B.2. Azure Load Balancer Configuration

This section requires the following configurations:

- Frontend IP address.
- Backend IP pool.
- Inbound load-balancing rules.

1. Log in to the Azure Portal at https://portal.azure.com.
2. Search for **load balancers** in the top-center search bar.

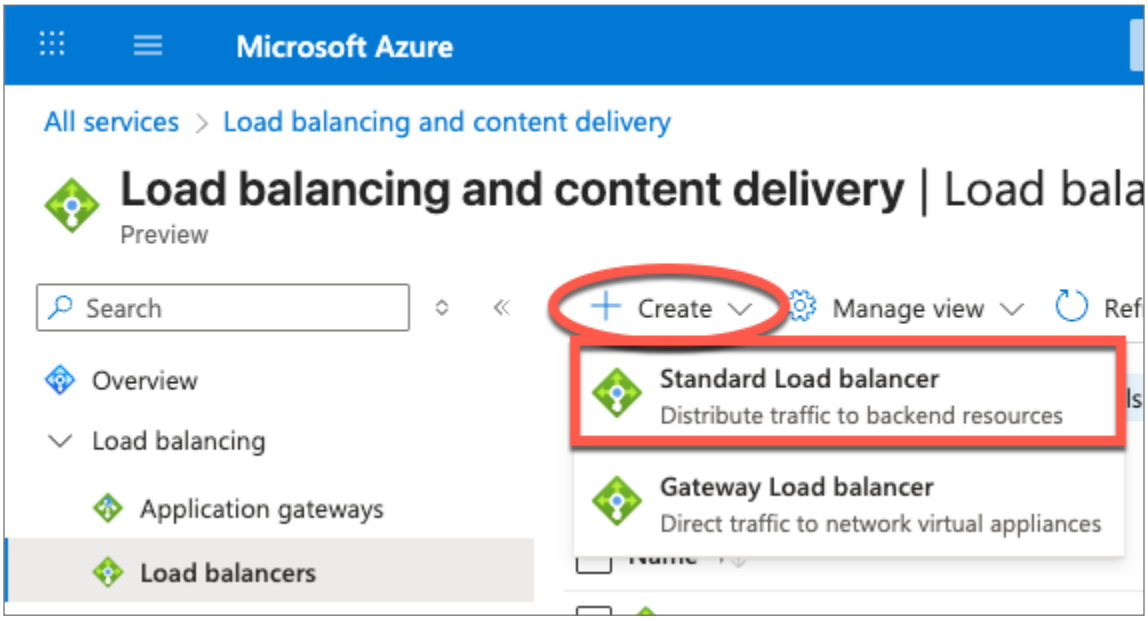

Figure 29. Searching for load balancers

3. Select **Load balancers** from the dropdown search menu.
4. Select **Standard Load balancer**.

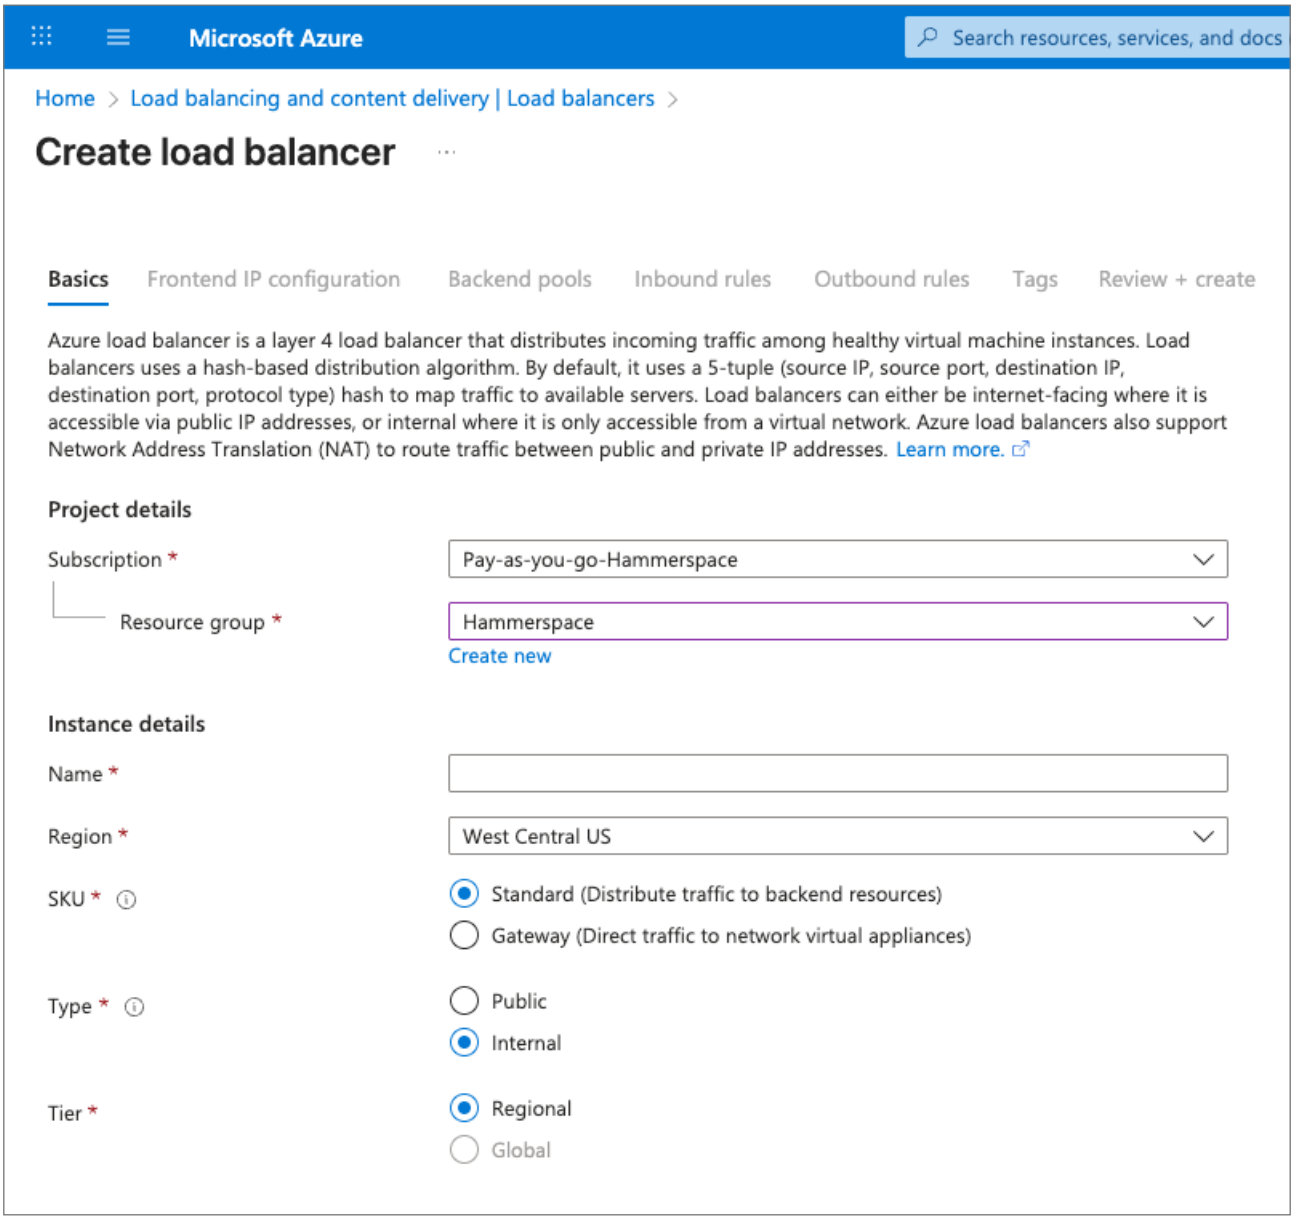

Figure 30. Load Balancer configuration

Table 13. Load-balancer settings

Setting	Value
Project Details	
Subscription	Select your subscription.
Resource group	Select the target resource group containing the Hammerspace nodes.
Instance Details	
Name	Give the load balancer a name.
Region	Select the Azure region.

Setting	Value
SKU	Standard.
Type	Internal.
Tier	Regional.

5. Select **Next: Frontend IP configuration** at the bottom of the page.
6. Click **+ Add a frontend IP configuration**.

Table 14. Front-end configuration settings

Setting	Value
Name	Provide a name.
IP version	Select the IPv4 radio button.
Virtual network	Select the target network.
Subnet	Select the target subnet.
Assignment	Static.
IP address	Enter the frontend IP address for your load balancer.

7. Click **Save** at the bottom of the **Add frontend IP configuration** panel.
8. Select **Next: Backend pools** at the bottom of the page.
9. Click **+ Add a backend pool**.

Table 15. Back-end pool settings

Setting	Value
Name	Provide a name for the backend pool.
Backend Pool Configuration	Select the NIC radio button.

10. Click **+ Add** below the **IP configurations** heading.
11. When the side panel opens, check the DSX instances you wish to load balance.
12. Click **Add** at the bottom of the panel.
13. When the side panel closes, click **Save** at the bottom of the page.
14. Click **Next: Inbound rules**.
15. Select **+ Add a load balancing rule**.
16. Select the following when the **Add load balancing rule** panel opens.

Table 16. Load-balancing rule settings

Setting	Value
Name	Provide a name for the load-balancing rule.
IP version	Select the IPv4 radio button.
Frontend IP address	Select the frontend IP address you previously configured.
Backend pool	Select the backend pool you previously configured.
HA ports	Check.
Protocol	TCP.
Port	Turns off when selecting High Availability ports.
Backend port	Turns off when selecting High Availability ports.
Health probe	Click to create a health probe (required).
Session persistence	Source IP.
Idle timeout	30.
Enable TCP Reset	Check.
Enable Floating IP	Unchecked.

Table 17. Health-probe settings

Setting	Value
Name	Give the health probe a name.
Protocol	Select TCP.
Port	445 (suggested), but it can be any open port.
Interval	5.

17. Click **Save** at the bottom of the panel.
18. Click **Next: Outbound rule**.
19. Click **Next: Tags**.
20. Optionally, create tags for the load balancer.
21. Select **Next: Review + create**.
22. After reviewing that the configuration is correct, select **Create**.

B.3. Example Topology Diagram

The following simplified diagram illustrates a client mounting the virtual IP address of the Azure Load Balancer. The Anvil HA configuration is omitted to simplify the illustration and highlight how DSX high availability works with Microsoft Azure.

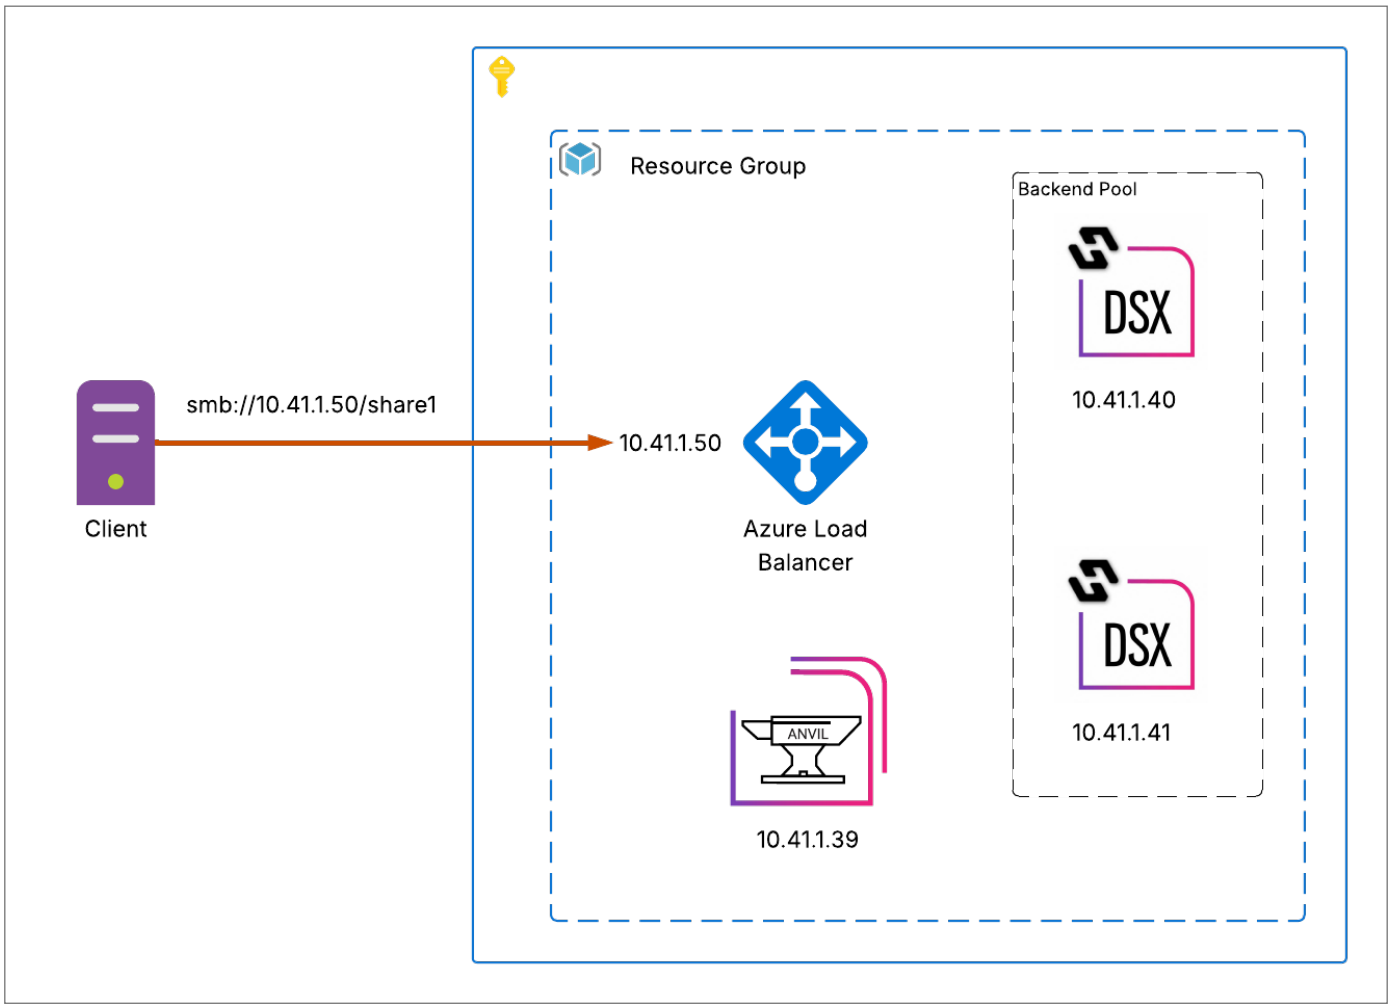

Figure 31. Example topology with an Azure Load Balancer

Appendix C: Architectural Reference Diagrams

C.1. Standalone Anvil and DSX

The minimum configuration of a Hammerspace instance consists of a single Anvil and a single DSX. Each instance requires at least two disks: one for the boot partition and the second for data or metadata.

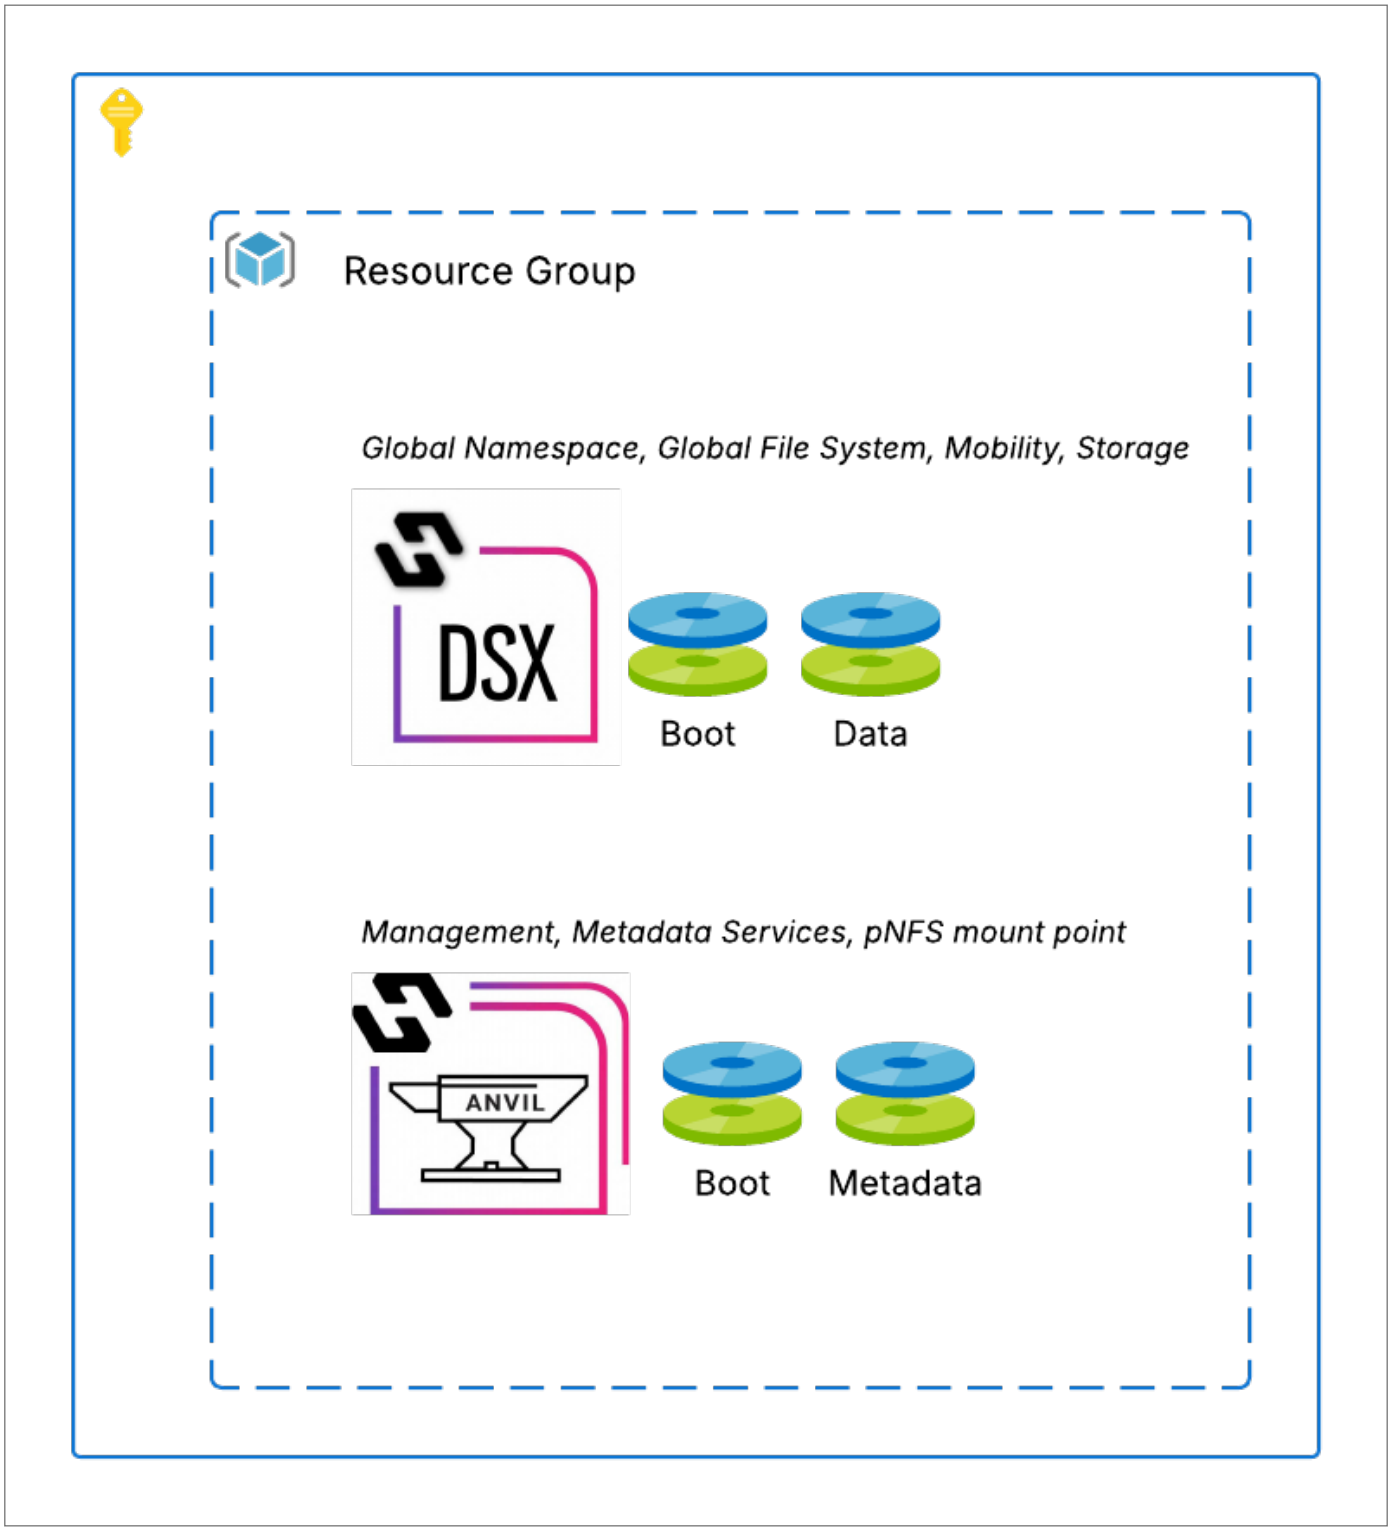

Figure 32. Standalone Anvil and DSX

C.2. High Availability Anvil and Standalone DSX

The following diagram illustrates a single Anvil and a highly available configuration of DSX nodes. Microsoft Azure requires configuring an Azure Load Balancer to support high availability.

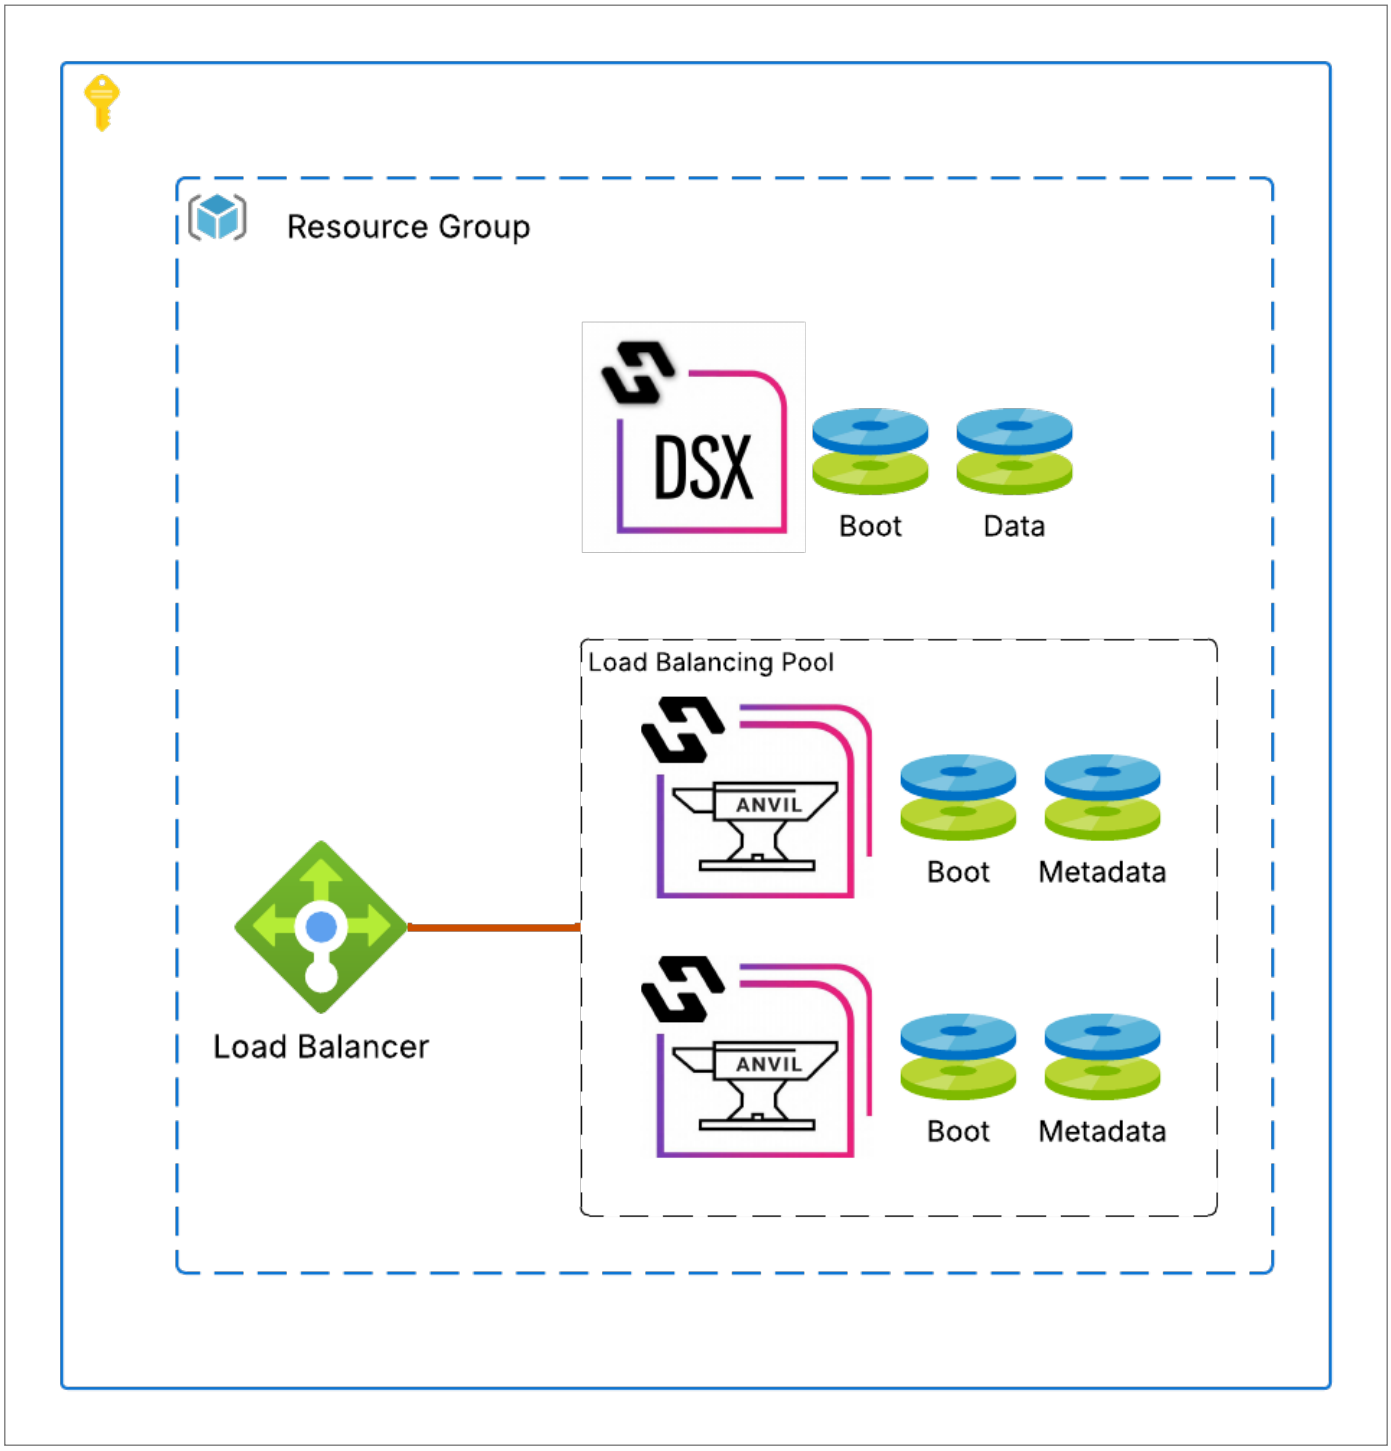

Figure 33. High availability Anvil and standalone DSX

C.3. High Availability Anvil and DSX

Microsoft Azure networking does not support Gratuitous ARP and thus cannot support floating IP addresses or dynamic high-availability solutions such as VRRP. Therefore, high availability requires configuring an Azure Load Balancer, as illustrated in the following diagram.

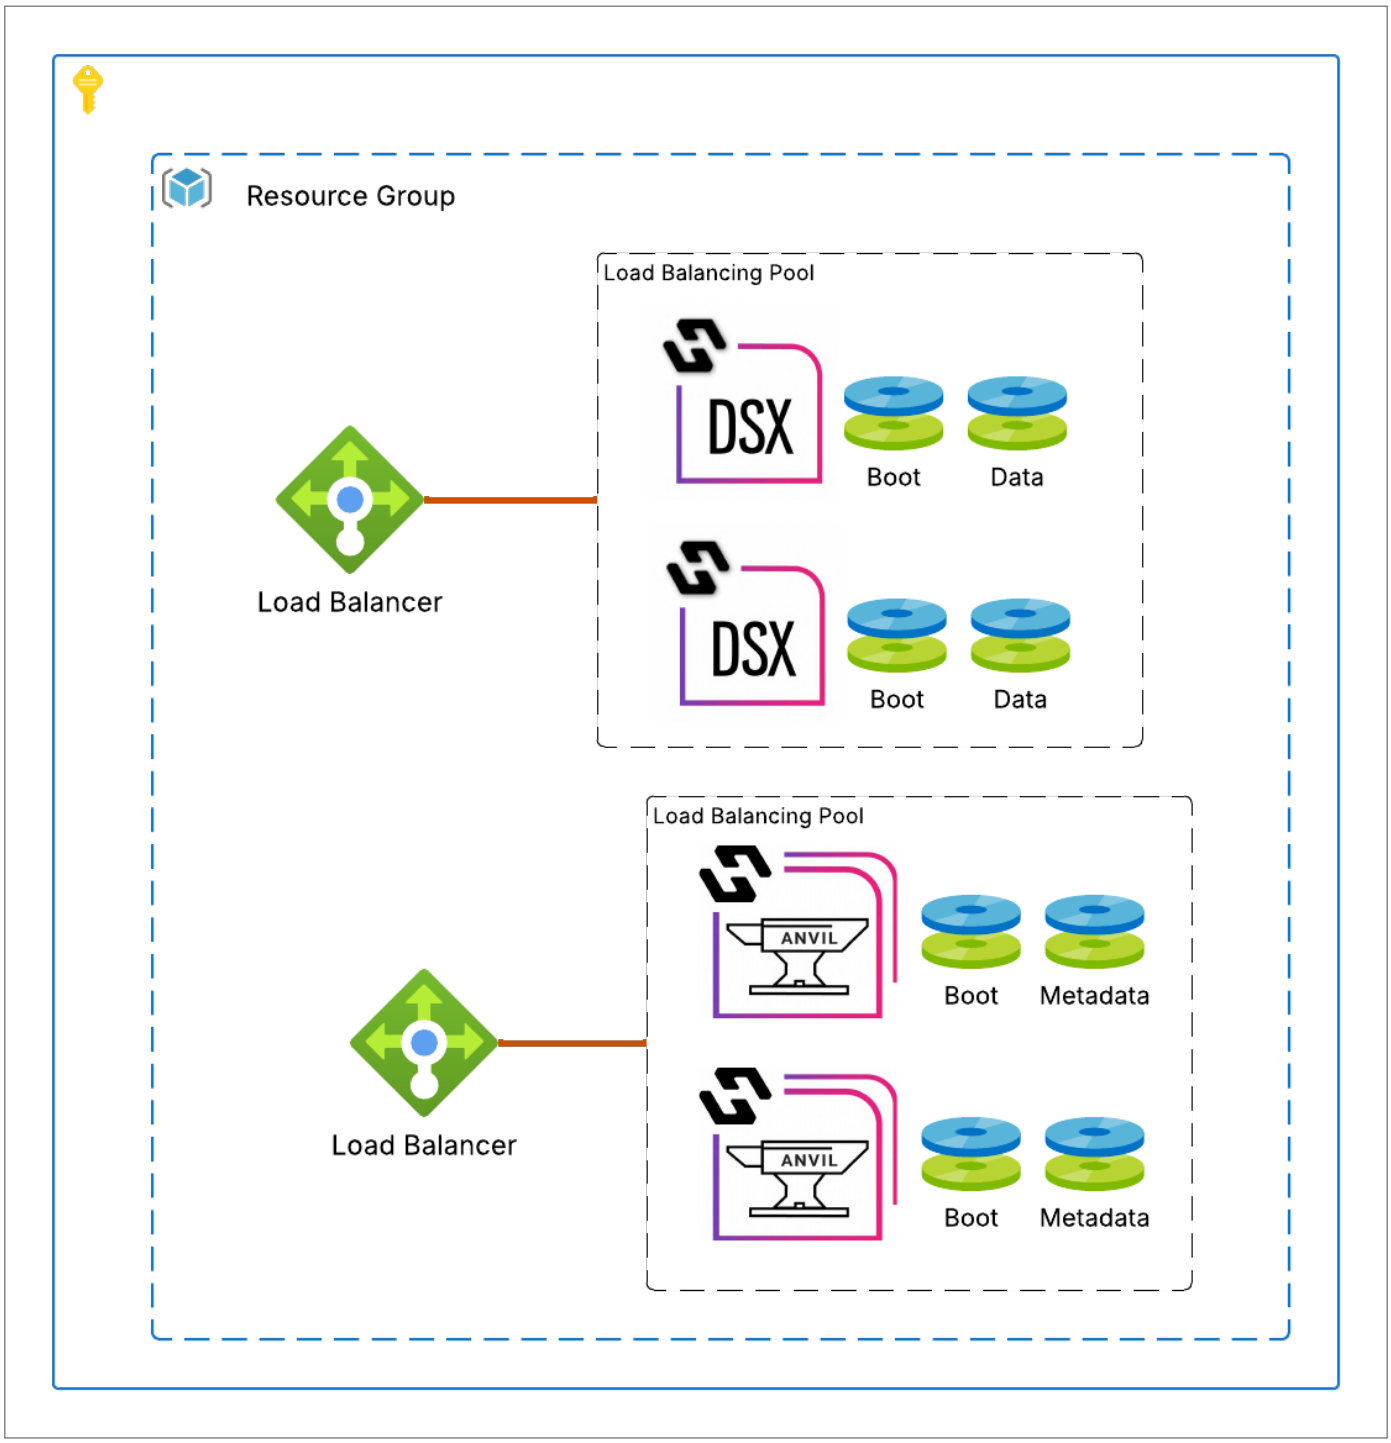

Figure 34. High availability Anvil and DSX

C.4. Data Path: SMB and NFSv3

The following diagram illustrates the read/write data path between an NFSv3 or SMB client and a minimal Hammerspace instance. The DSX node is the mount point and delivers the global namespace and global file system. The Anvil node controls metadata and data management, data orchestration, and all administrative functions. It also provides the Hammerspace Management GUI and Admin CLI.

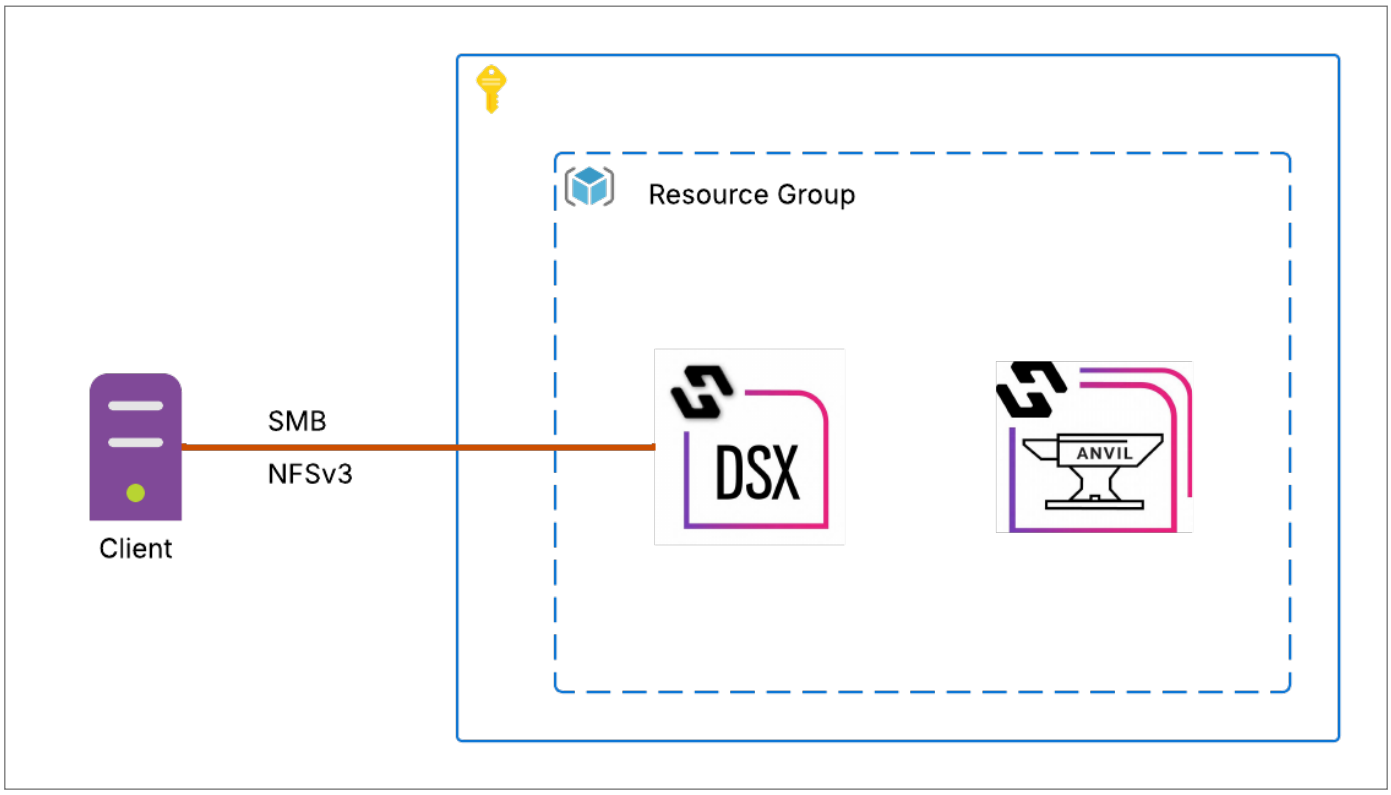

Figure 35. Data path for SMB and NFSv3

C.5. Data Path: Parallel NFS (NFSv4.2)

Parallel NFS is ideal for high-performance use cases. The data and control connections are separated for greater parallelization. They are suitable for the most demanding workloads, including AI training and inference, genomic sequencing, rendering, and other HPC scenarios requiring extreme performance.

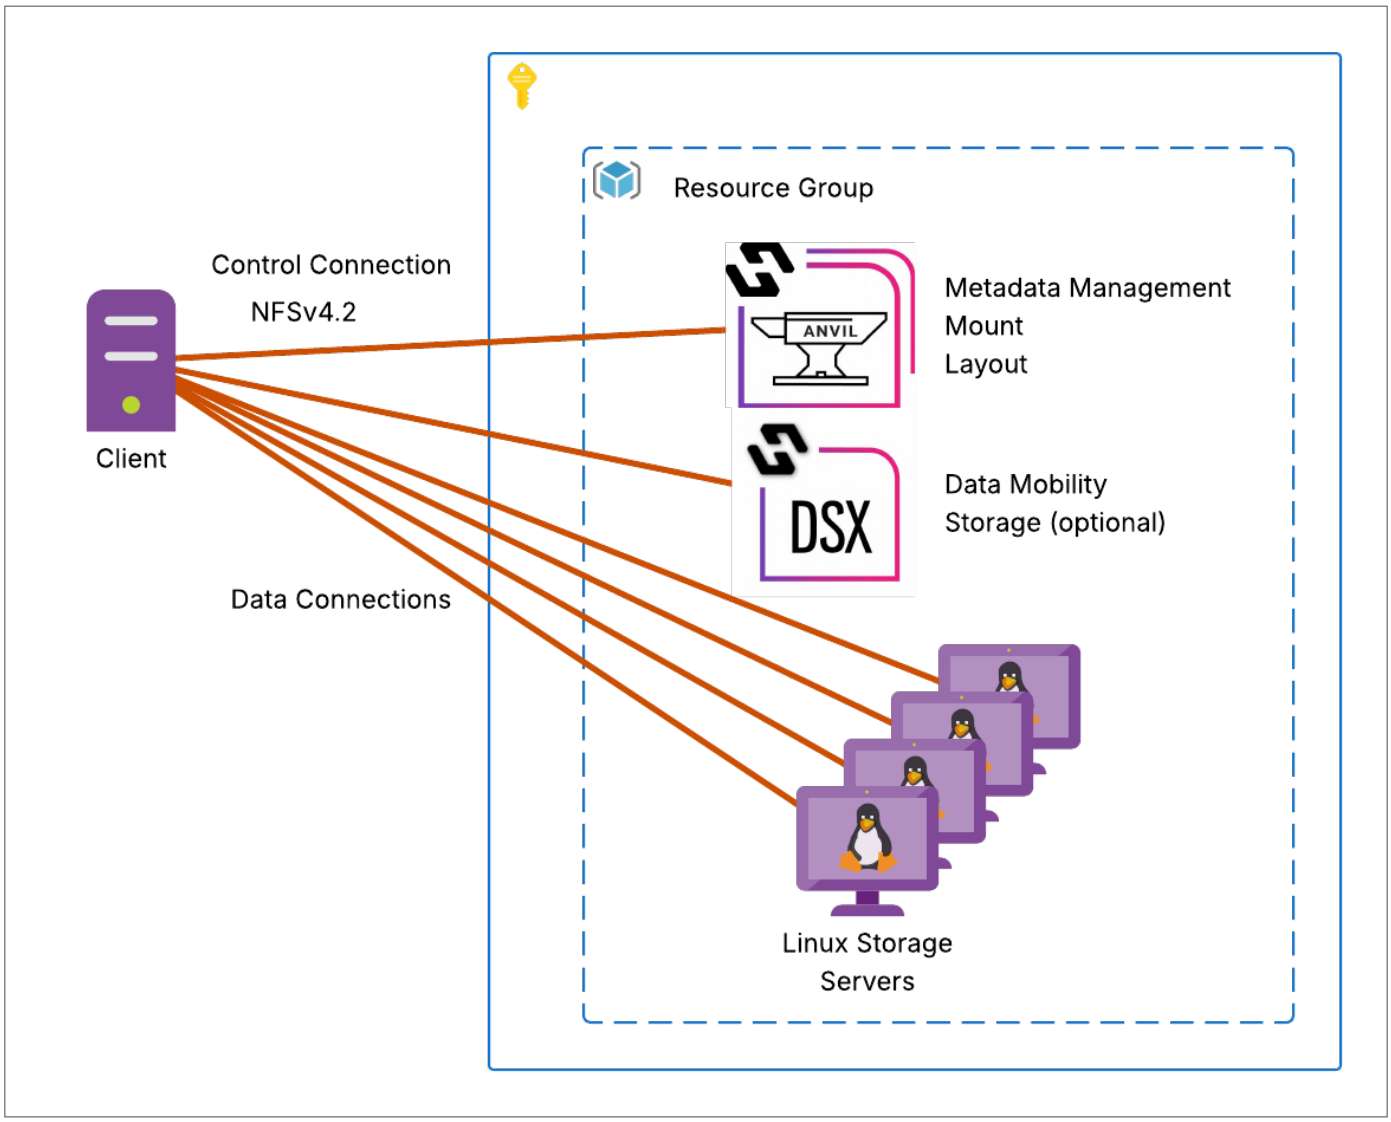

Figure 36. Data path for Parallel NFS (NFSv4.2)

NFSv4.2 is used for the highest performance requirements, including AI/ML workloads. Connections are parallelized from the client to multiple backend Linux servers. For more information on Parallel NFS, see the following documentation:

- RFC 5661
- RFC 8434

Appendix D: Documentation and Support

D.1. Azure References

- Get started with Azure
- Azure networking documentation
- Azure virtual machines documentation
- Azure Load Balancer documentation
- Create an internal load balancer (Azure Portal)
- Azure resource provider operations

D.2. Hammerspace Resources

- Hammerspace on Azure
- Hammerspace resources
- Hammerspace support
- Email: support@hammerspace.com

For additional support resources, see the Hammerspace Support Hub and the Hammerspace Licensing and Delivery Portal.