S3 Guide

Table of Contents

About This Documentation 1
1. Concepts and Terminology 2
1.1. Hammerspace and S3 Terminology 2
1.2. Multi-protocol File Access Support 3
1.3. S3 Limitations 3
1.4. S3 and Multi-tenancy 4
1.5. Transport Security 4
1.5.1. X.509 Server Certificate 4
2. Planning Your Hammerspace S3 Deployment 5
2.1. Prepare for S3 Configuration 5
2.2. Prerequisites 5
2.3. Supported S3 API Operations 6
3. Create a Local User for S3 Access 7
4. Manage S3 Servers 9
4.1. Create an S3 Server 9
4.2. Edit an S3 Server Configuration 12
4.3. Delete an S3 Server 13
5. Manage S3 Buckets and Bucket Containers 15
5.1. Create an S3 Bucket Container 15
5.2. Delete an S3 Bucket Container 16
5.3. Create an S3 Bucket 17
5.3.1. S3 Bucket Creation Overview 17
5.3.2. Create an S3 Bucket Using the GUI 18
5.3.3. Create an S3 Bucket Using the S3 API 21
5.3.4. Create an S3 Bucket Using the Bucket Container 22
5.4. Delete an S3 Bucket 23
Appendices **26**
Appendix A: Appendix 27
A.1. Unix Permissions - Octal Values 27
A.2. Creating Your S3 Umask Value for a Local Identity Provider 27
Appendix B: Additional Resources 29

About This Documentation

Data managed by Hammerspace may be stored and retrieved as files or objects using NFS, SMB, and S3, regardless of which storage the data is on today or moves to later. Any share or directory can be exposed as an S3 bucket, and buckets and objects created using S3 can be viewed as directories and files through standard NFS and SMB shares. Data is visible and accessible in a unified namespace via any protocol, reducing complexity instead of amplifying it. "File" and "Object" become different lenses through which data may be viewed and accessed.

Hammerspace's highly performant and scalable S3 object interface enables users, applications, and compute to access data via Hammerspace using standard NFS, SMB, and S3 file and object protocols interchangeably.

- In Hammerspace, S3 objects and S3 buckets are represented 1:1 as files and directories.
- Any Hammerspace share or folder can be exposed as a bucket.
- Hammerspace shares can be presented as NFS mount points, SMB shares, and S3 buckets. Any share or directory may be exposed as an S3 bucket. Said differently, any data can now be accessed via NFS, SMB, and S3.
- Security policy is uniform across all protocols, including S3.
- Once objects are written into a Hammerspace bucket or share, you may assume that all standard functionality applies, including:
  - Setting Objectives for data mobility and data tiering
  - Concurrent file access across different protocols
  - Global file system capabilities, i.e., a file can be written to a share via NFS in one site and accessed via S3 at another site so long as permissions are correctly set.

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. Concepts and Terminology

This section introduces the terminology used throughout this guide, and the concepts to understand before planning a Hammerspace S3 deployment: multi-protocol file access, tested S3 limitations, multi-tenancy options, and transport security.

1.1. Hammerspace and S3 Terminology

Term	Definition
Access key	A key used to access and authenticate with the Hammerspace S3 server.
Anvil	A component in the Hammerspace architecture that provides a global namespace and metadata management.
Bucket	A container for storing data on the S3 server.
Bucket Container	An S3 endpoint that clients can authenticate to for the purpose of creating S3 buckets.
DSX (Data Service Exchange)	A Hammerspace node that provides data access and management services. S3 service runs on DSX nodes.
Endpoints	One or more DNS endpoints with prefixes. The endpoint identifies the S3 server and provides a way for clients to connect to different S3 servers. The endpoint must be unique to this S3 server. If no endpoint is configured, the default server option must be checked. Prefix allows virtual hosted style bucket access. The prefix must be the same for all endpoints.
FQDN (Fully Qualified Domain Name)	A unique identifier for the S3 server used by clients to connect.
Global Data Platform	The core platform provided by Hammerspace unifies data across different storage systems and locations.
List Objects Limit	Sets the default number of keys returned by the ListObjects or ListObjectsV2 API call. Note that the ListObjectsV2 API is allowed to override this value.
Local user	A user created on the Hammerspace system to map to Linux identities.
Name	Management name for the S3 server.
Objective	A policy that defines how data should be placed and managed within the Hammerspace system.
RFC2307 / RFC2307bis	A method for mapping Unix users to their associated Microsoft Active Directory accounts. Hammerspace leverages these standards to maintain permissions in multi-protocol environments.
S3 (Simple Storage Service)	An object storage service supported by Hammerspace.

Term	Definition
S3 Access Permissions	Access permissions for objects created in buckets.
s3cmd Command-line Client	A command-line client that is a widely used tool for interacting with S3-compatible storage.
S3 Interface	Tool within the Hammerspace GUI to configure S3 storage in a Hammerspace infrastructure.
S3 Server	A server that provides S3 storage services.
Strict Bucket Lists	When enabled, only return the buckets owned by the user.
Strict Bucket Names	Enables strict bucket naming to conform with Amazon S3 bucket naming rules.
Virtual-hosted Bucket Names	Check to enable virtual hosted style bucket access; if checked, a prefix is required in the endpoint address (Ex: `https://bucket-name.<Prefix>.foo.bar.com/`), and DNS needs to be configured for all buckets that will be addressed using virtual hosted style bucket names.

1.2. Multi-protocol File Access Support

To support multi-protocol access that maintains permission consistency between NFS, SMB, and S3 clients, Hammerspace leverages RFC2307 to map S3 access keys to AD users and Unix UIDs/GIDs. These user mappings control the identity of the objects within the file system.

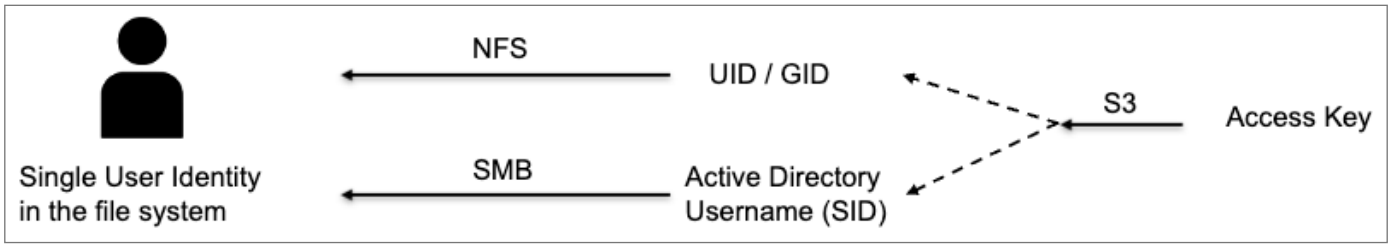

Figure 1. S3 access keys map to a single user identity in the file system

RFC2307 should be implemented in Active Directory and enabled on the Hammerspace cluster if you intend to provide multi-protocol access to files.

1.3. S3 Limitations

The following tested limitations currently apply to Hammerspace support for the S3 service:

- Maximum number of S3 servers: 2048
- Maximum number of buckets: 20,000 per S3 server
- Maximum number of objects (total): No limit; the total number of objects is limited by metadata capacity.
- Maximum number of objects in a bucket: No limit; the total number of objects is limited by metadata capacity.

- Maximum number of Access Keys per S3 Server: 10,000
- Minimum object size: 0 bytes
- Maximum object size (multi-part upload): Chunk Size * maximum number of chunks (10,000); the chunk size is controlled by the S3 client.
- File locking support: S3 does not have the concept of file locking, though locking is supported across SMB and NFS v4.2 protocols.

1.4. S3 and Multi-tenancy

Hammerspace allows you to achieve multi-tenancy with S3 in a few ways. By combining these methods, you can create a secure and isolated environment for each tenant within your Hammerspace infrastructure.

- Support for Multiple S3 Servers: Each S3 server is configured with a unique name (FQDN)/endpoint.
- Buckets: Each S3 server can be assigned a bucket or set of buckets, ensuring that data is logically separated and tenants cannot access each other's data.
- Access Control: Each S3 server can be configured with unique access keys, ensuring that S3 clients and their data can be isolated.
- Objectives: Use objectives to control data placement and access for each tenant, allowing you to set different policies for tenants, such as where their data is stored and who can access it.

1.5. Transport Security

An S3 server can be configured to allow HTTP and/or HTTPS connections, though it is recommended to use HTTPS to ensure transport-level security between the S3 client and the S3 server.

- HTTPS connectivity is enabled by default and is supported on port 443.
- HTTPS connectivity is only supported with TLS (Transport Layer Security) versions 1.2 and 1.3. Support for 1.0 and 1.1 has been removed as those versions are no longer considered safe. If there is a choice, using TLS 1.3 on your client is recommended.

The encryption ciphers have also been reduced to comply with the latest security errata.

1.5.1. X.509 Server Certificate

By default, the product comes with a self-generated security certificate. However, this will often generate "untrusted certificate warnings" that some S3 clients may not be able to bypass.

To avoid these warnings, you can install a certificate in the product using the `cluster-config` CLI command.

For more information about managing server certificates, consult the Hammerspace Administration Guide.

Chapter 2. Planning Your Hammerspace S3 Deployment

Hammerspace supports running multiple S3 servers, each with a unique configuration. The servers are hosted on all the nodes running the S3 Service, which allows each S3 Server to be a scale-out entity.

The following tasks describe a generic workflow for configuring S3 storage for a Hammerspace infrastructure.

1. Identify AD S3 users or create local Hammerspace S3 data accounts
2. Create an S3 server
3. Setup the S3 server DNS records
4. Setup S3 bucket containers and buckets
5. Manage users and access
6. Configure S3 clients

2.1. Prepare for S3 Configuration

Several data points should be decided upon before configuring the S3 server.

- **Endpoint Name:** Each S3 server is usually addressed using an endpoint name, which determines which S3 server the traffic is routed to. This name is represented by a DNS record, which is recommended to be configured before setting up the S3 server.
- **Data Access:** S3 buckets have two methods for data access, traditional path-based or virtual-host style. It is essential to know that the virtual-host style requires additional DNS entries for each bucket accessed using virtual-host addressing.
- **Bucket Management:** Hammerspace supports managing buckets with the administrative API, creating and deleting buckets via the S3 API, creating or deleting buckets within bucket containers by SMB, NFS, or S3 clients, or a combination of all three.
- **Existing Data Access:** If sharing existing data (an entire share or a sub-folder within a share), determine what kind of access should be provided via S3, as permissions are standard across all protocols. However, access can be limited via S3 using an ACL.

2.2. Prerequisites

Hammerspace S3 servers support connections using the same DSX floating IP addresses (FIPs) used by SMB and NFS clients. To help balance S3 traffic and ensure that clients can maintain S3 connections during maintenance or unexpected outages of DSX servers, you should create one DNS A record for each DSX FIP, for each S3 server that you create, just as you did for the cluster itself.

The following image shows a typical DNS configuration for a Hammerspace cluster with two DSX nodes and configured with two S3 servers.

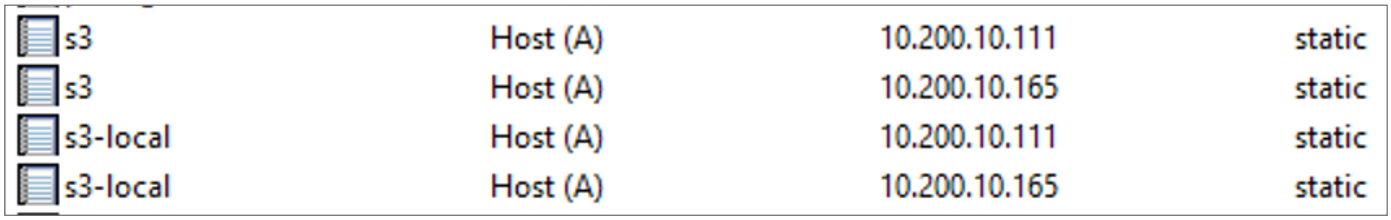

Figure 2. Prerequisites

While the S3 servers also support connections over the DSX static IP addresses, those IPs should not be used since they cannot migrate to other DSX nodes in the event of DSX node maintenance or an outage.

2.3. Supported S3 API Operations

Hammerspace supports a subset of the S3 API covering bucket, object and multipart-upload operations, along with path-style and virtual-hosted bucket addressing and presigned URLs using AWS Signature Version 4 query parameters.

The operation list is release-specific, so it is maintained in one place rather than duplicated here. For the authoritative list for your release — including the operations that are accepted but behave in a fixed way, such as the Intelligent-Tiering calls — see the **S3 API Overview** section of the Hammerspace Administration Guide.

Operations that do not appear in that list are not supported.

Chapter 3. Create a Local User for S3 Access

Hammerspace S3 supports both a local identity provider using cluster local accounts and Active Directory (AD) for AD accounts. Access keys and secret keys will need to be created for each one regardless of which account is used.

The following procedure outlines how to create a local Hammerspace user for the purpose of S3 access, a requirement for using a Hammerspace S3 server that uses a local identity provider.

- Navigate to Users from the cluster Administration menu and click Create User.

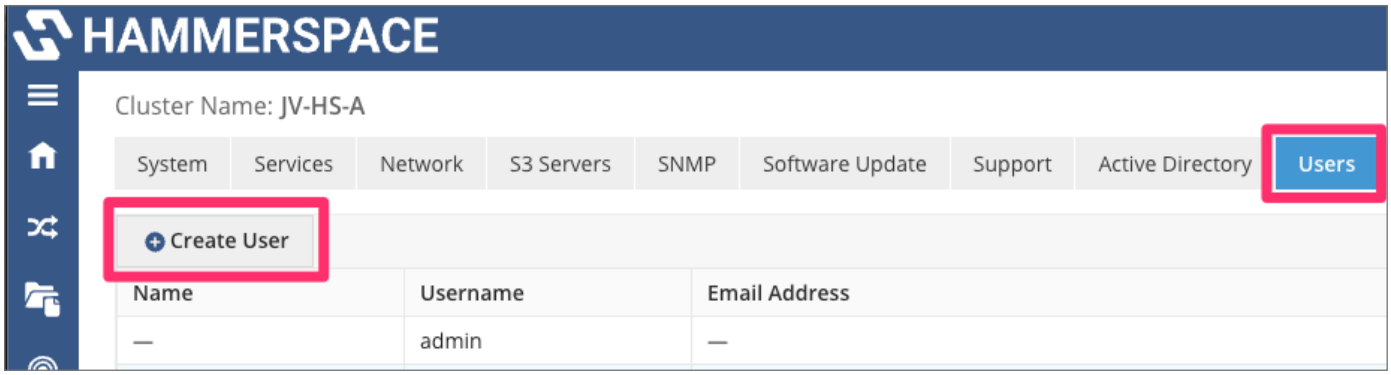

Figure 3. Create a local user for S3 access

- Provide a unique username and password, select Data access for the Role type, click the Data Access Role S3 check box, and click Create.

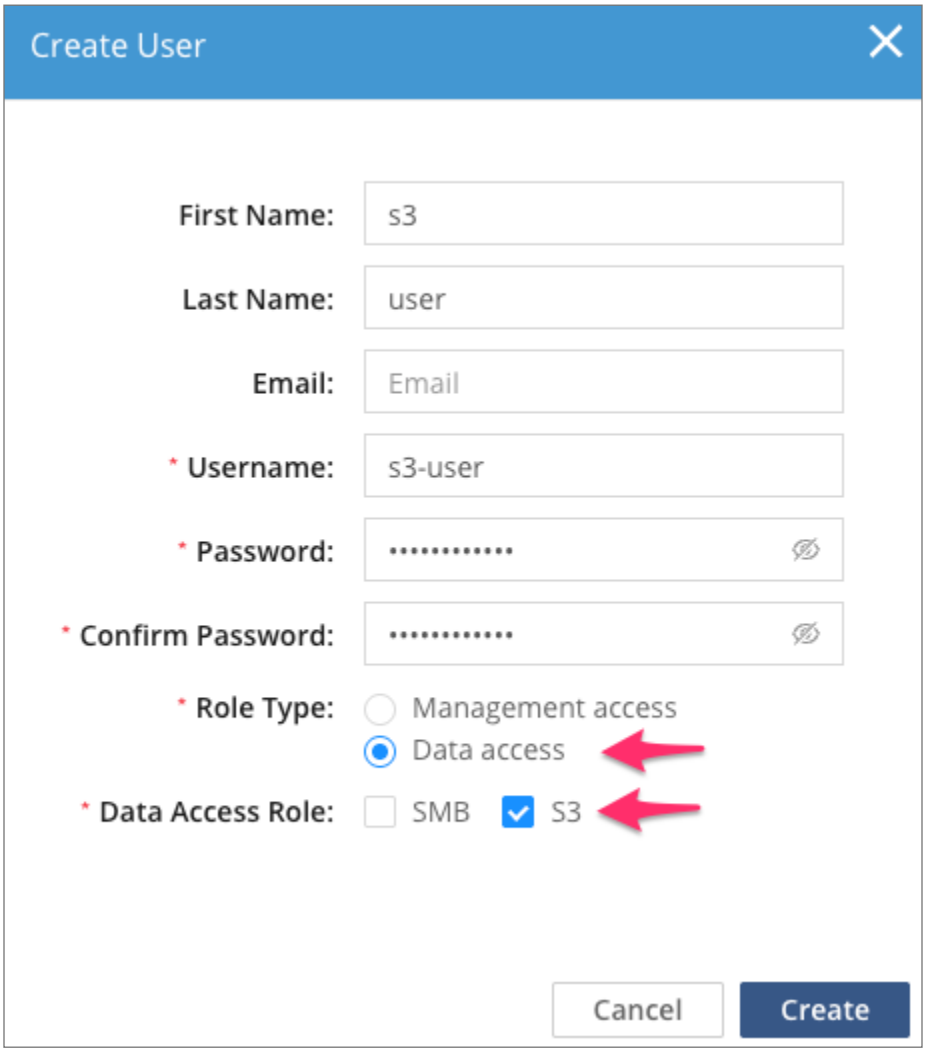

Figure 4. Create a local user for S3 access

- Confirm the new S3 user is listed among the local Hammerspace users. Note that while you did set a username and password, access to S3 requires a separate access key and secret key.

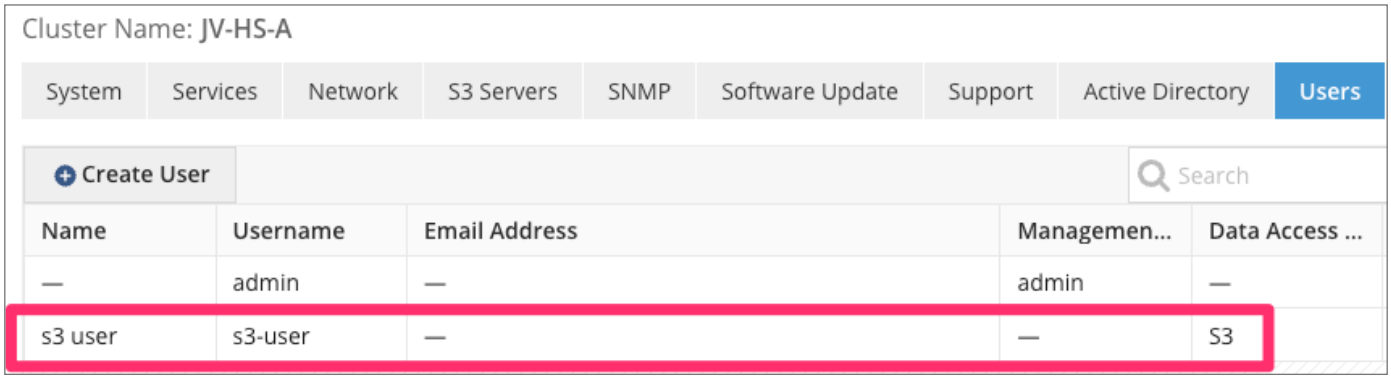

Figure 5. Create a local user for S3 access

Remember the S3 username, as you will need it when creating S3 buckets. If more local users are required, create them at this time.

Chapter 4. Manage S3 Servers

The following section will outline the procedures used to create, edit, and delete Hammerspace S3 servers.

4.1. Create an S3 Server

The following steps outline how to create an S3 server using the Create S3 Server wizard.

- Navigate to S3 Servers from the cluster Administration menu and click Create S3 Server.

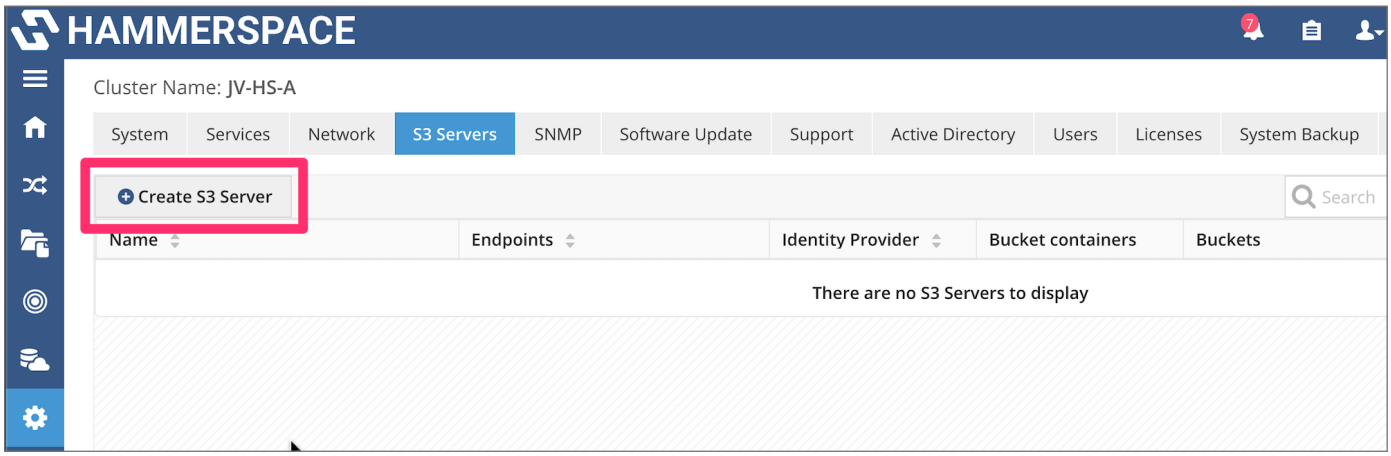

Figure 6. Create an S3 server

- Provide a unique name for the S3 server in the Name field; note that the Endpoint field will automatically be populated, but can be edited with whatever endpoint FQDN you prefer. Complete the following steps to complete step 1:
  - Select an Identity Provider, either Local for local Hammerspace users or Active Directory, to specify an AD user.
  - Set the S3 access permissions at this time, which can be public, group-based (based on the Unix group membership of the uploading user), or private to limit access to only the uploading user.
  - Configure the other optional settings as desired. Click Create to create only the S3 server without any buckets or access keys, or select Next to move to step 3.

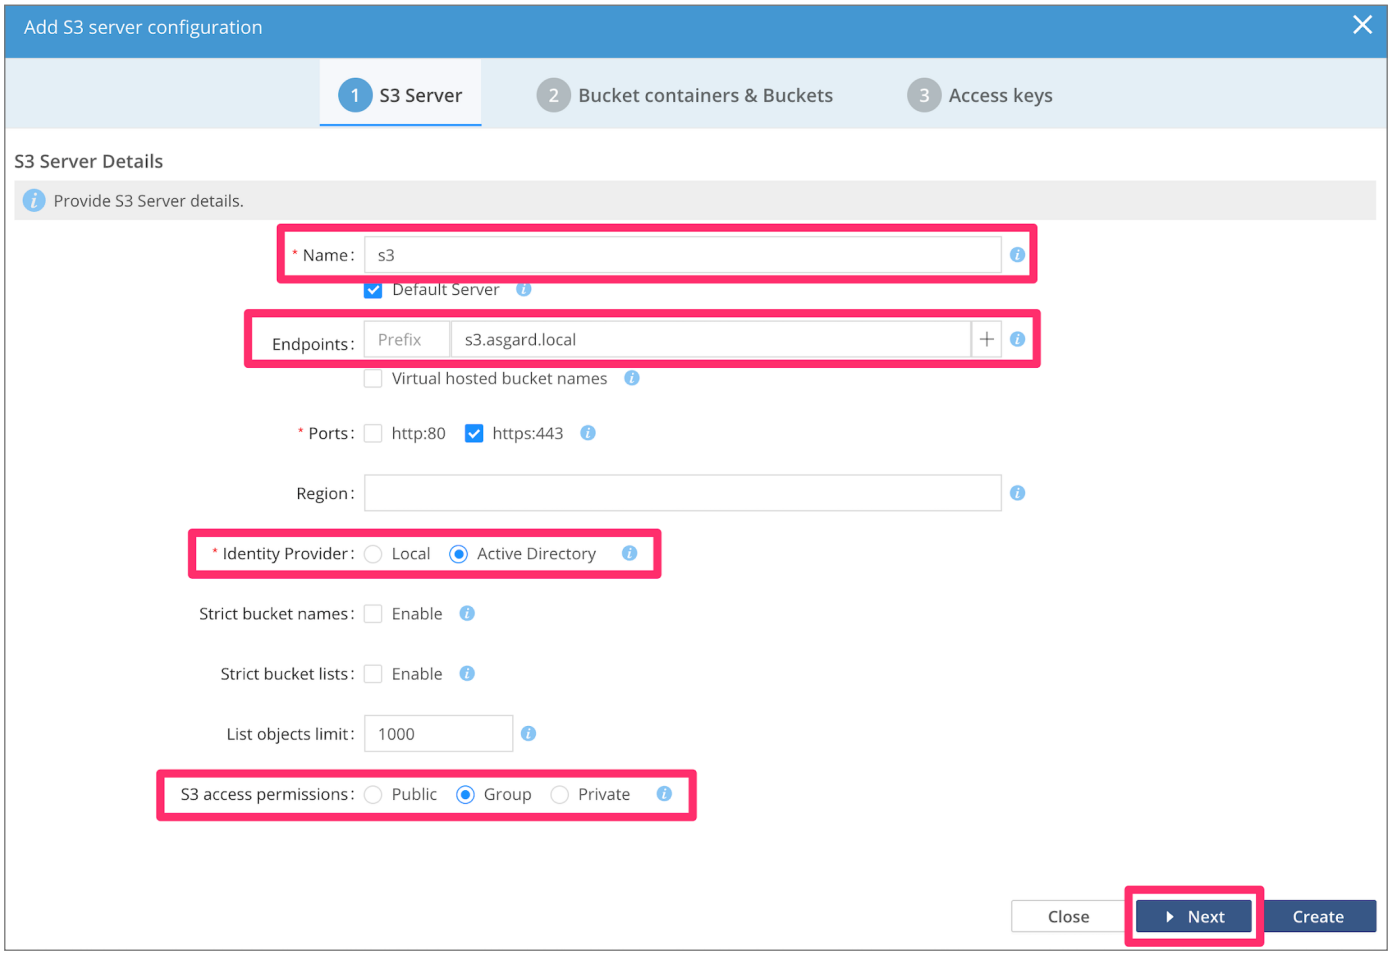

Figure 7. Create an S3 server

- If the S3 access permissions are set to Group, and the Identity Provider is Active Directory, then the AD user who uploaded the file will be the file owner, and the members of the user's primary AD group will be able to edit those files.
- Additionally, if the Identity Provider is Local, then a default umask of 0002 will be set. Refer to the Appendix section Creating your S3 Umask Value for a Local Identity Provider and Unix Permissions - Octal Values for an explanation of the syntax of umask values so that you can customize one as needed to set file and folder permissions for objects uploaded using S3.
- In step 2, you have the option of creating an S3 bucket or a bucket container, which allows clients to create buckets using S3 commands. Note that you must pre-create the share that you will use to host the bucket or bucket container. Click Next to move to step 3.

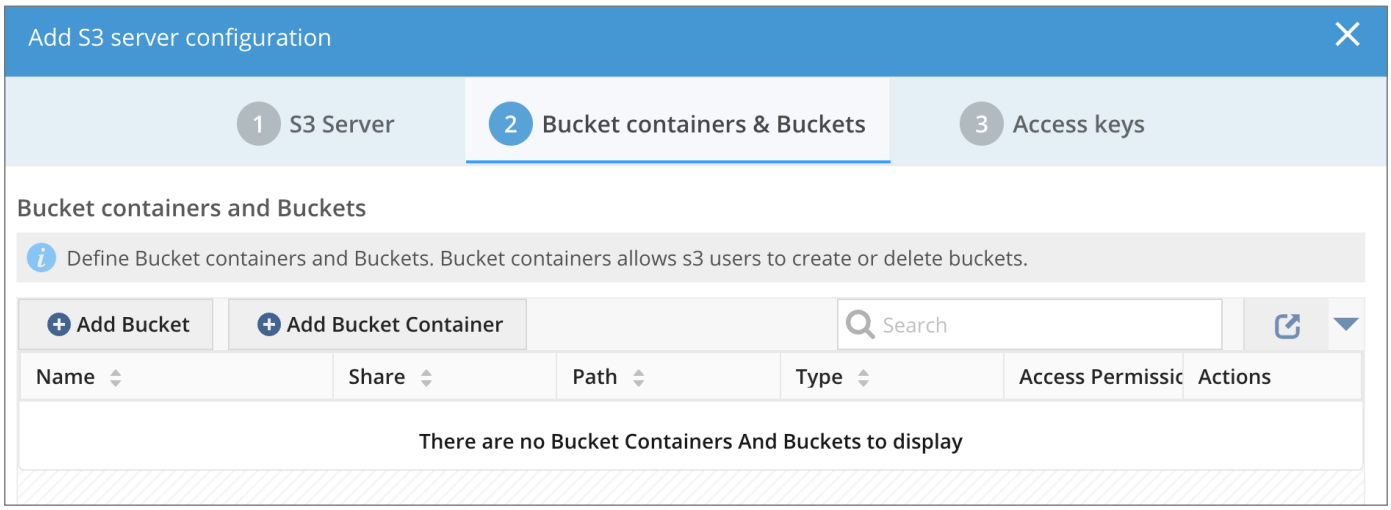

Figure 8. Create an S3 server

- Refer to the Create an S3 Bucket and Create an S3 Bucket Container sections of this document for an explanation of the options used when creating those objects.
- In step 3, you have the option of creating or uploading one or more access keys to provide access to the bucket.
  - Note that regardless of which identity provider you selected in step 1, the users you specify must be created before adding them here.
  - In this example, the S3 server identity provider is Active Directory, so our S3 user is an AD account in the format user@domain. It is not necessary to create access keys now; they can always be added later, as described in the Managing S3 Users section of this document. Once you are finished with this step, click Create.

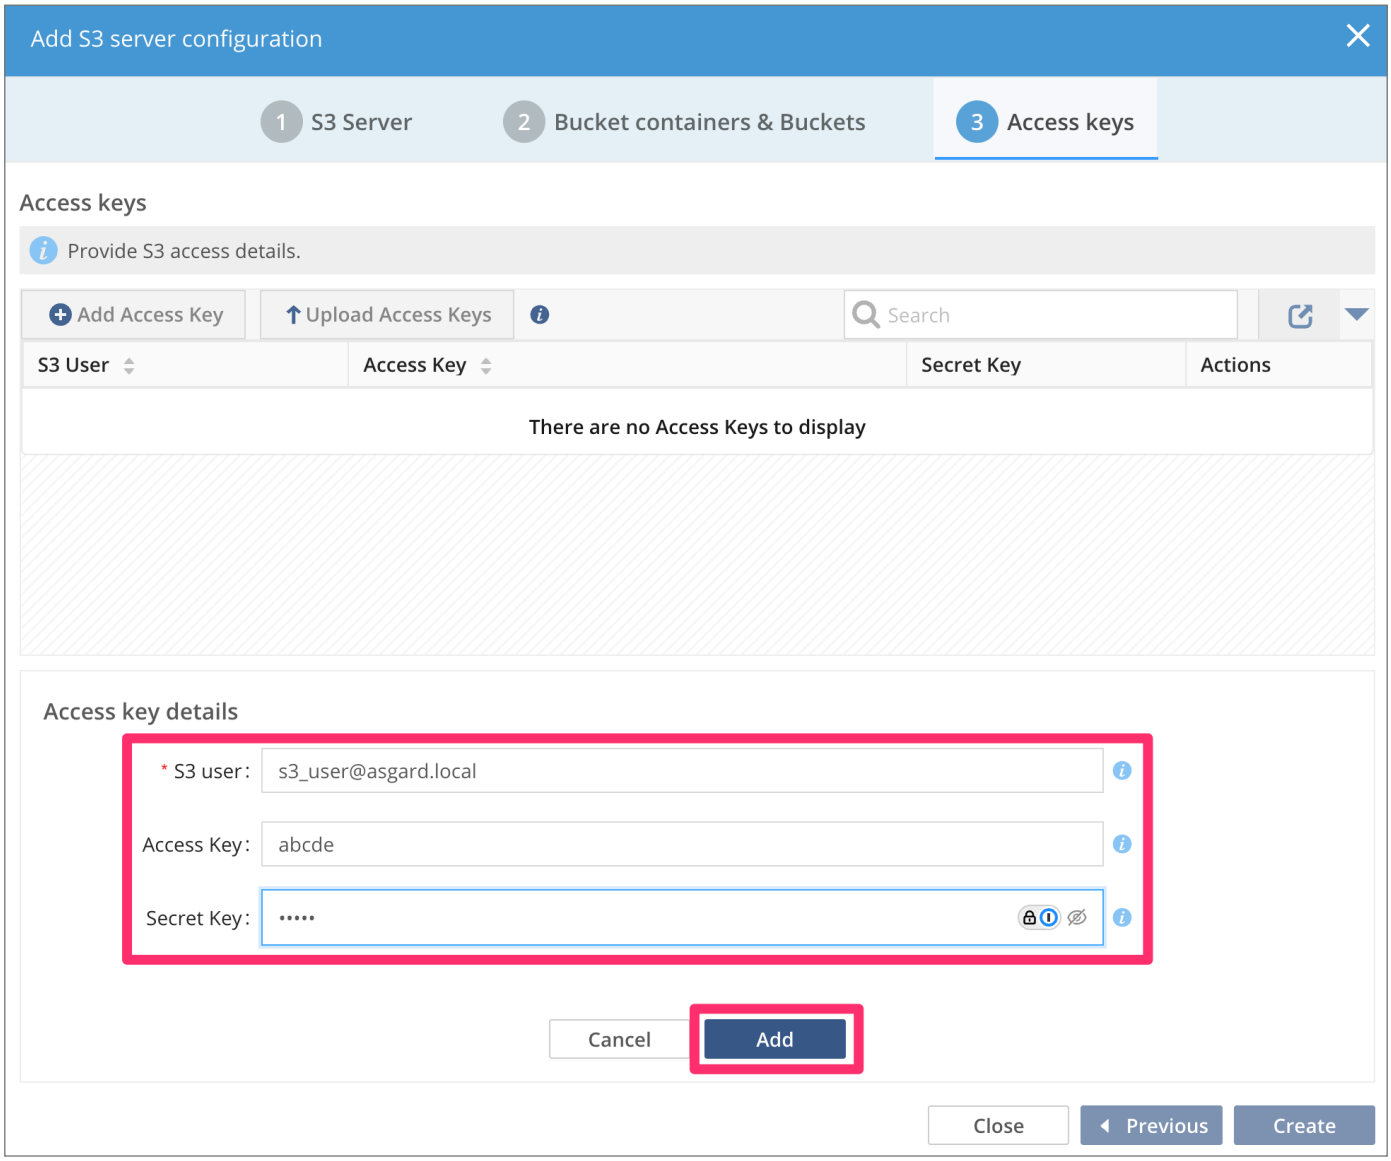

Figure 9. Create an S3 server

Your S3 server has been created. You are ready to create buckets or, if needed, make changes to the S3 server using the pencil icon to the far right of the S3 server name.

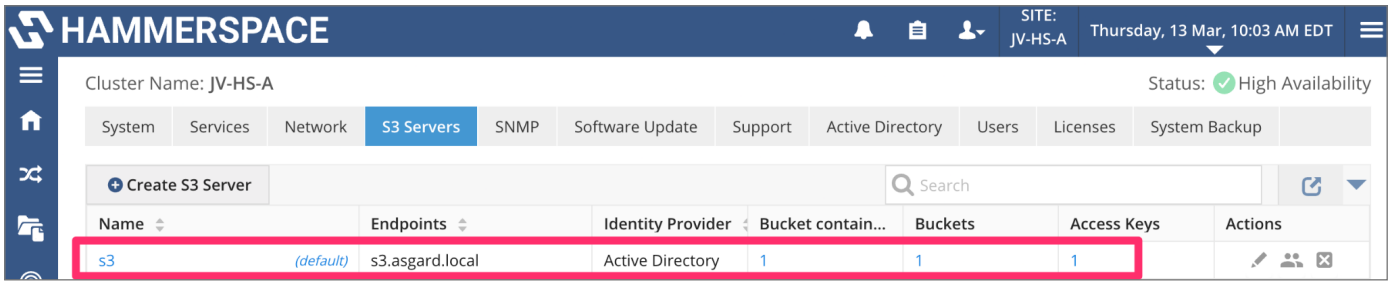

Figure 10. Create an S3 server

4.2. Edit an S3 Server Configuration

To edit the configuration of an existing S3 server, navigate to the S3 Servers tab in the cluster Administration

menu.

Click the Pencil icon in the Actions column for the S3 server.

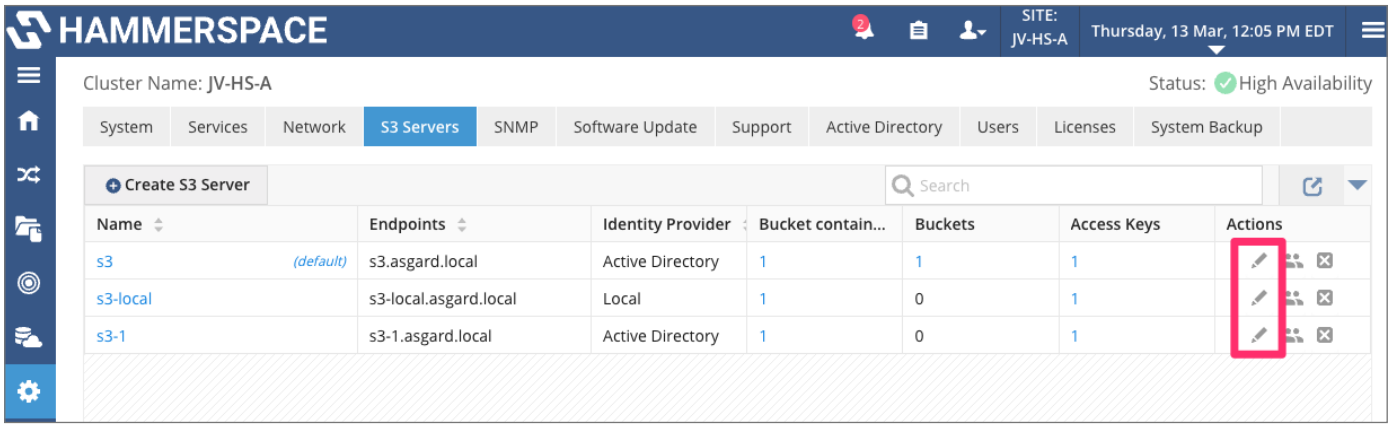

Figure 11. Edit an S3 server configuration

Most of the settings can be changed in the Edit S3 server configuration window, excluding the bucket name or identity provider.

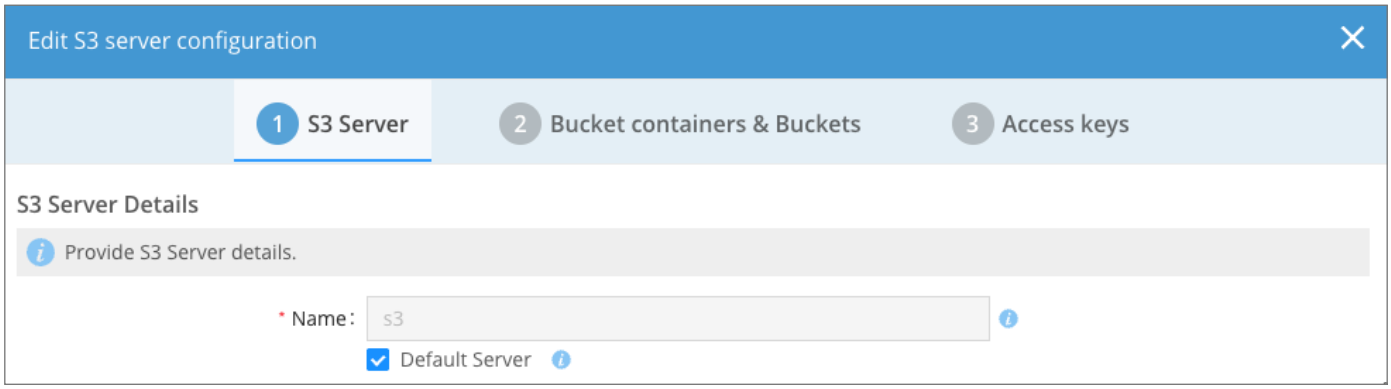

Figure 12. Edit an S3 server configuration

4.3. Delete an S3 Server

In this section, we will discuss how to delete a Hammerspace S3 server. Keep in mind that deleting a server does not delete any data, including S3 folders and their contents. It only deletes the server and S3 bucket configuration. Additionally, bucket folders and their contents can be removed manually (if desired), but this should only be done if the bucket endpoint is deleted.

- To delete an S3 server, navigate to the S3 Servers tab in the cluster Administration menu and click on the X icon under the Actions column for the S3 server.

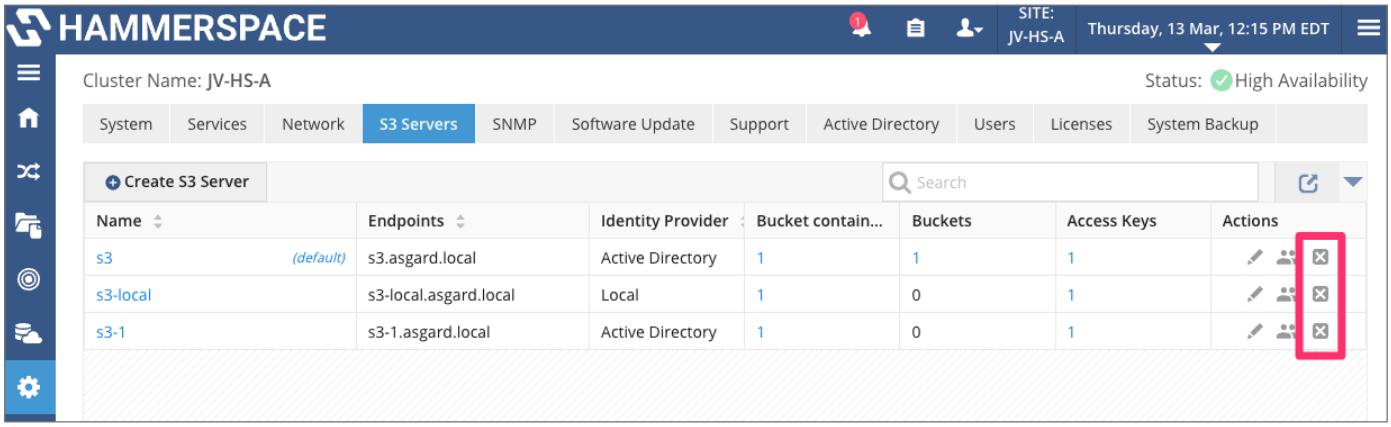

Figure 13. Delete an S3 server

- Type the S3 server name and click the Delete S3 server configuration button to complete the deletion process.

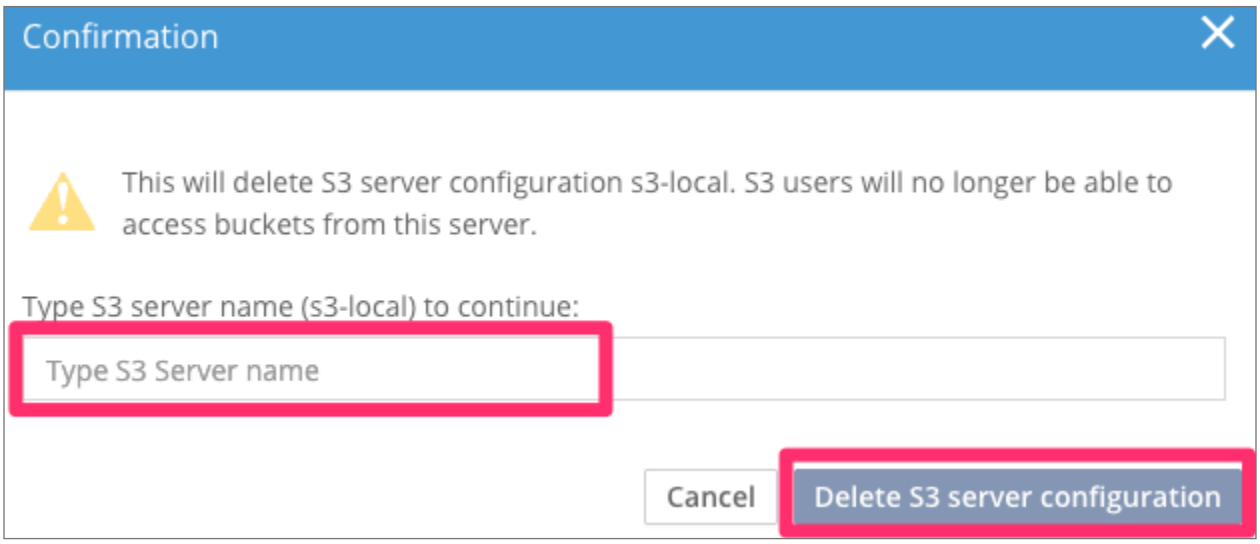

Figure 14. Delete an S3 server

Chapter 5. Manage S3 Buckets and Bucket Containers

In this section, we will review the following procedures:

- Create and delete S3 bucket containers
- Create, edit, and delete S3 buckets

5.1. Create an S3 Bucket Container

Bucket Containers can be created either during the S3 server creation process or after the S3 server has been created using the Data menu - Bucket Containers tab - Add Bucket Container wizard.

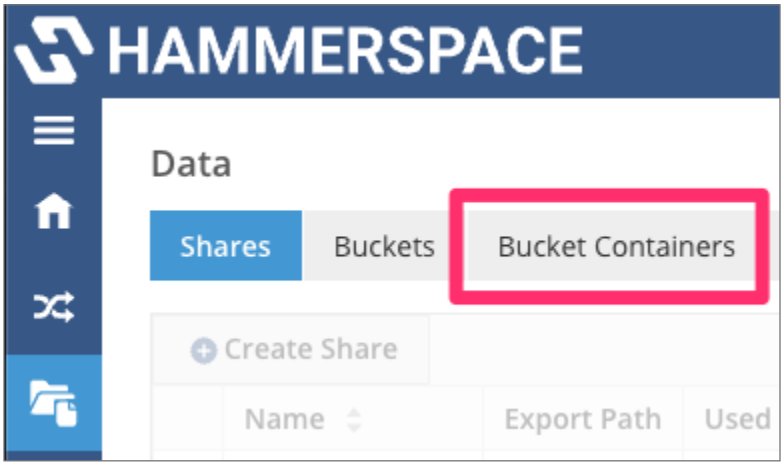

Figure 15. Create an S3 bucket container

To create a bucket container, you need to specify three values:

1. Bucket container name
2. Share where the bucket container will be created
3. Bucket container path within the share

As mentioned previously, the bucket Share *must* be created and the bucket owner identified before creating the bucket.

The following demonstrates a bucket container that is being created as part of the Create S3 Server wizard, though the procedure is similar if you create it using the Data menu | Bucket Containers | Add Bucket Container wizard.

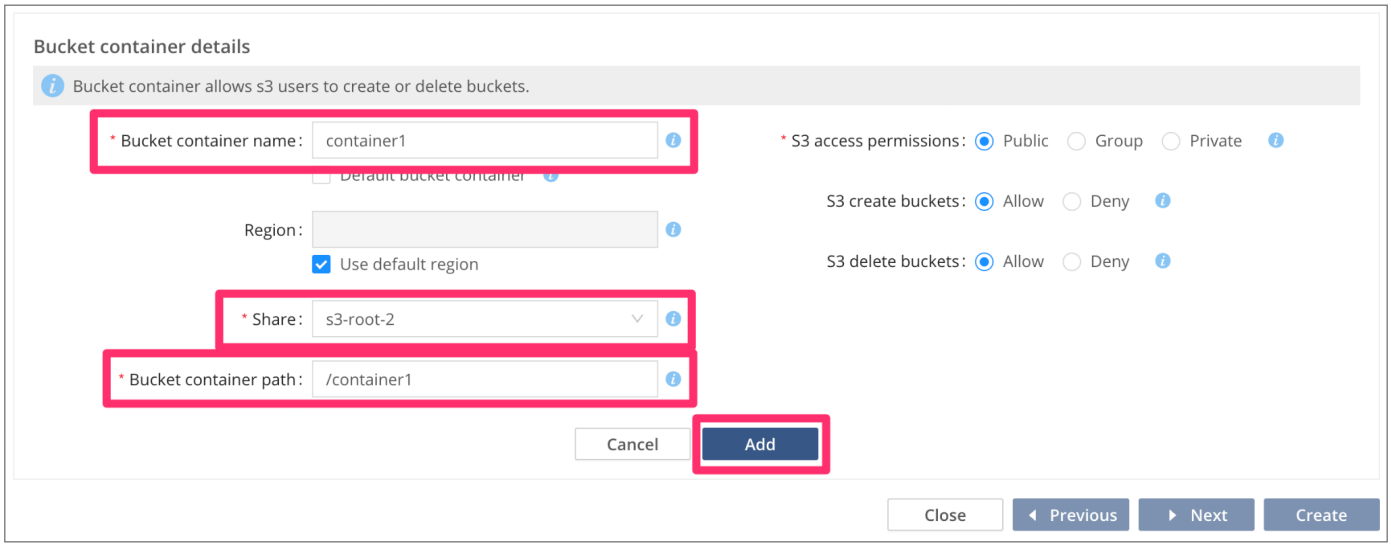

Figure 16. Create an S3 bucket container

- Make sure you configure the Bucket Container options per your requirements.
- The S3 access permissions setting controls only access to the bucket itself, not the file system ACLs for the bucket.
- Buckets and Bucket Containers cannot share the same path.

5.2. Delete an S3 Bucket Container

- To delete an S3 bucket container, navigate to the Bucket Containers tab in the Data Menu Administration menu and click on the X in the Actions column for the bucket container.

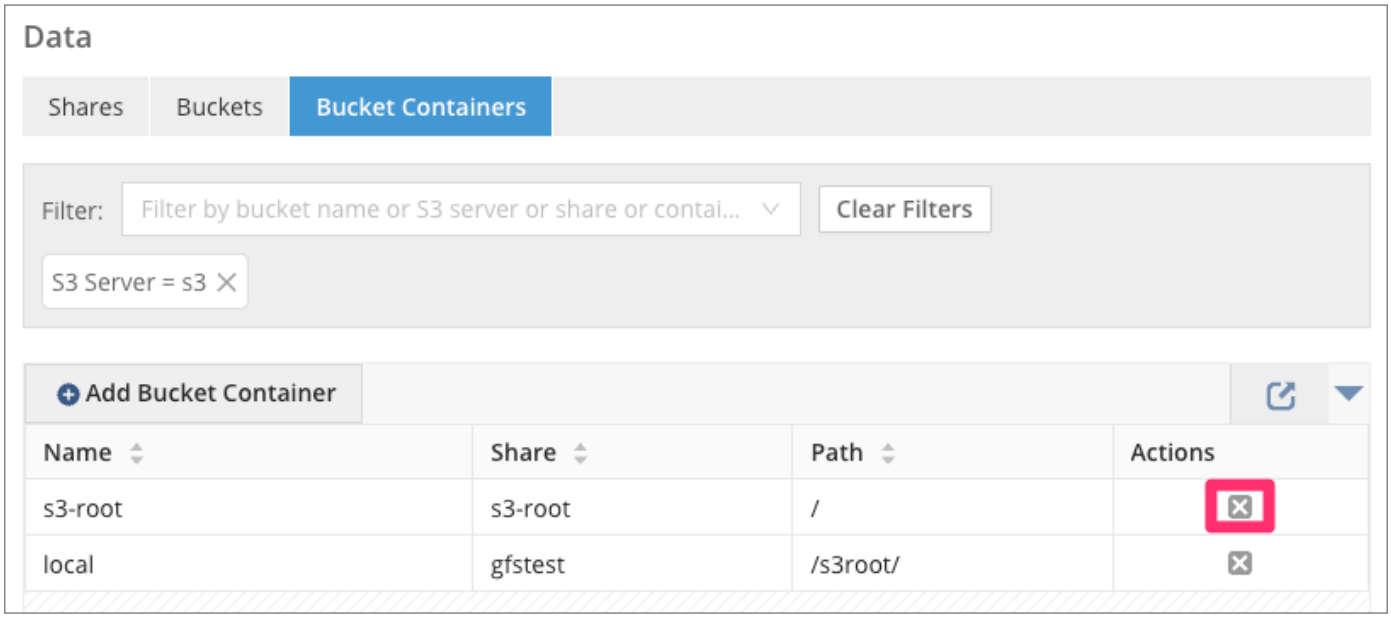

Figure 17. Delete an S3 bucket container

- Type the bucket container name as indicated in the Confirmation window. Review the warning that specifies how deleting a bucket container will cause users to lose access to all the buckets it contains. If acceptable, click the Delete bucket container configuration button to complete the deletion process.

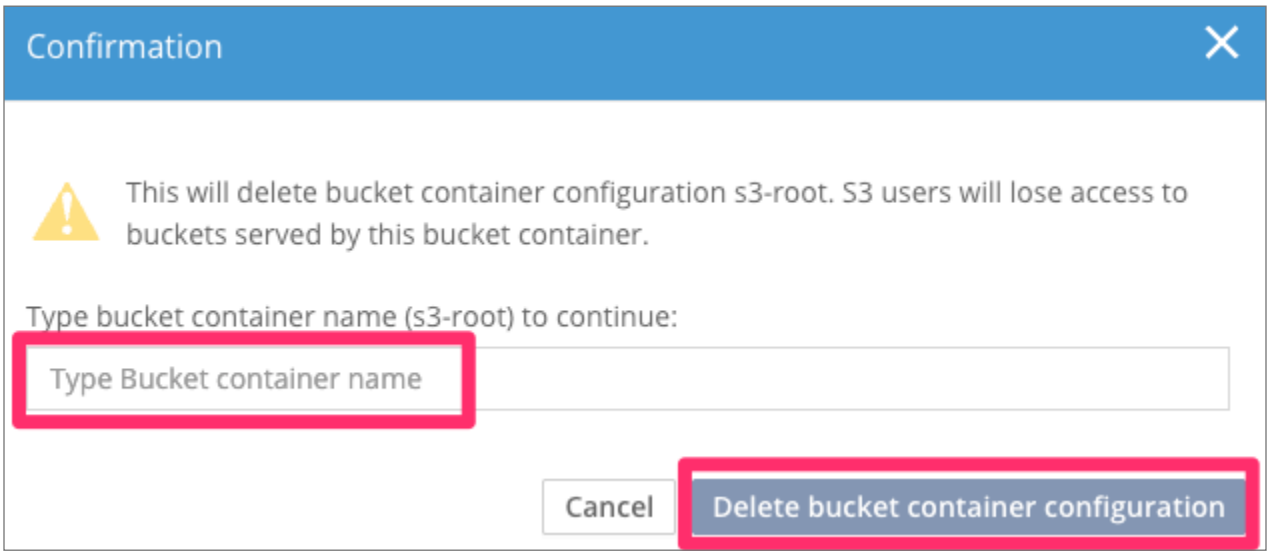

Figure 18. Delete an S3 bucket container

5.3. Create an S3 Bucket

Buckets can be created either during the S3 server creation process after the S3 server has been created using the Data menu | Buckets tab | Add Bucket wizard or using S3 commands within a properly configured bucket container.

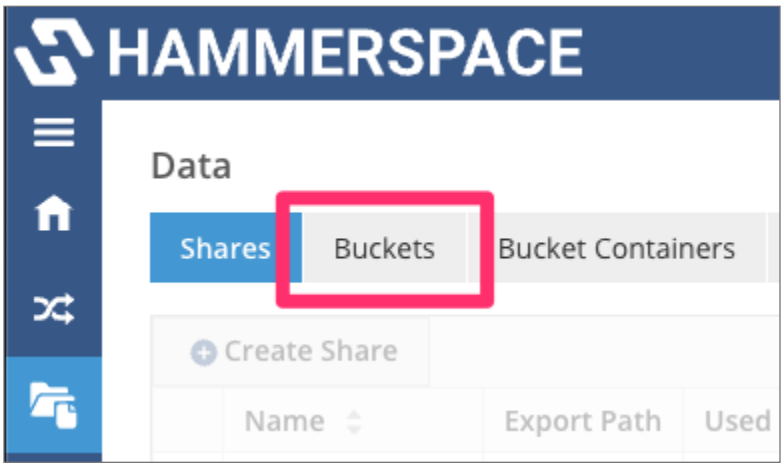

Figure 19. Create an S3 bucket

5.3.1. S3 Bucket Creation Overview

To create a bucket, you need to specify at least three values:

- The bucket name (what the clients see)
- The share where the bucket will be created
- The bucket path within the share
- The bucket owner
  - An AD user in the user@domain format (if using Active Directory as the identity provider)

- A local cluster username for a user that has S3 data access (if using a local identity provider)

If a new folder is created as part of the bucket creation process, based on the S3 access permissions (ACL) selected, it will have the following permissions set:

S3 ACL	NFS v3	NFS 4 (nfs4_getfacl)	SMB
Public	0777/drwxrwxrwx	A:fd:OWNER@:rwaDdxtTnNcCoy A:fdg:GROUP@:rwadxtTnNcy A:fdg:EVERYONE@:rwadxtTnNcy	
Group	0770/drwxrwx---	A:fd:OWNER@:rwaDdxtTnNcCoy A:fdg:GROUP@:rwadxtTnNcy	
Private	0700/drwx------	A:fd:OWNER@:rwaDdxtTnNcCoy	

If the folder that will be used for the bucket has already been created, it will retain its existing ACLs or POSIX permissions.

5.3.2. Create an S3 Bucket Using the GUI

The following steps outline how to create a bucket using the cluster Data menu | Buckets tab | Add Bucket wizard, though the procedure is similar to creating the bucket during the S3 server creation process. As mentioned previously, the bucket Share must be created and the bucket owner identified before creating the bucket.

- Navigate to the Buckets tab in the Data Menu, select your target S3 server, and click Add Bucket.

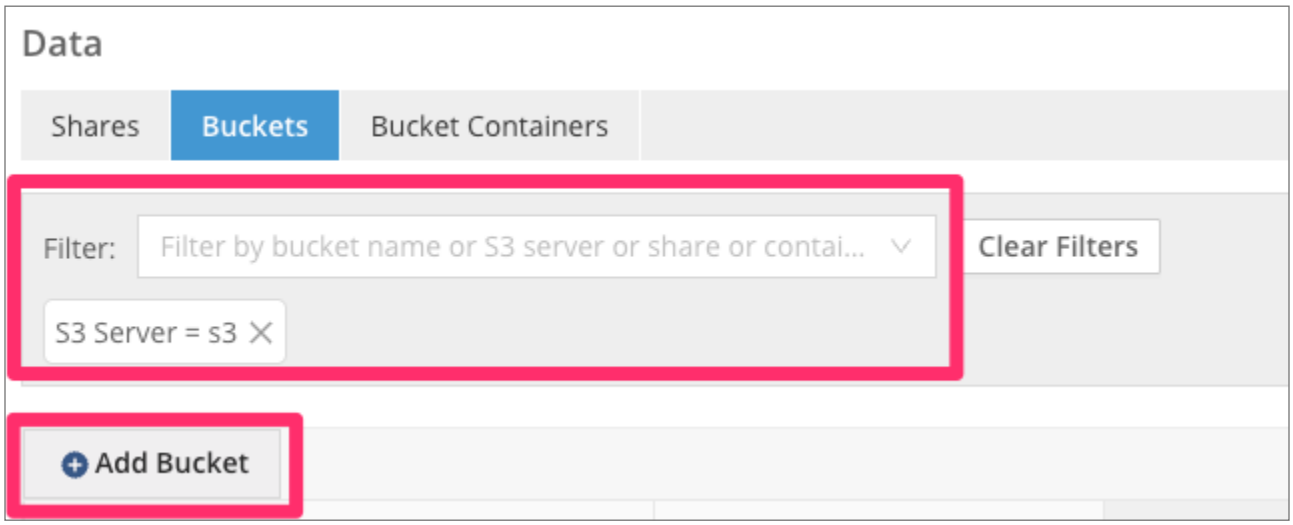

Figure 20. Create an S3 bucket using the GUI

- Provide a Bucket name, change the options as desired, and then click Next to move to step 2.

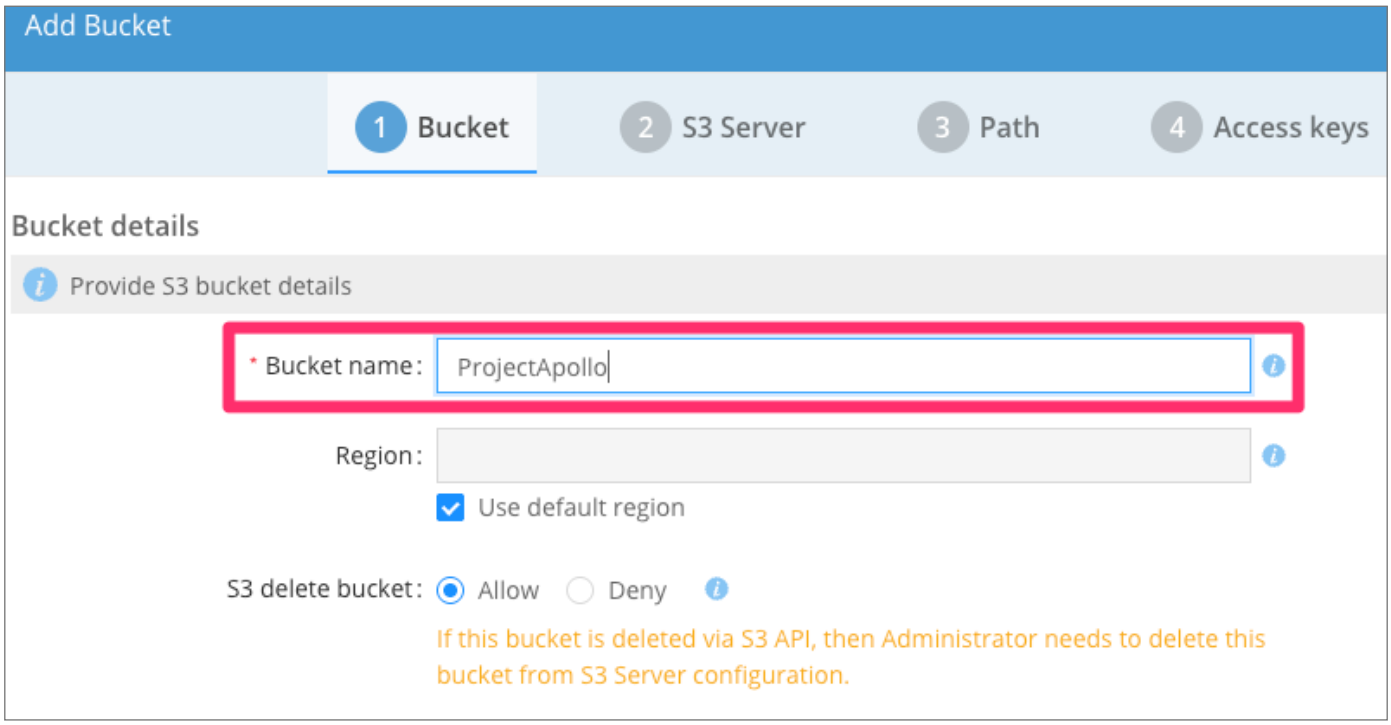

Figure 21. Create an S3 bucket using the GUI

- Since we are using an existing S3 server, there are no options to select in Step 2; click next to move to Step 3.

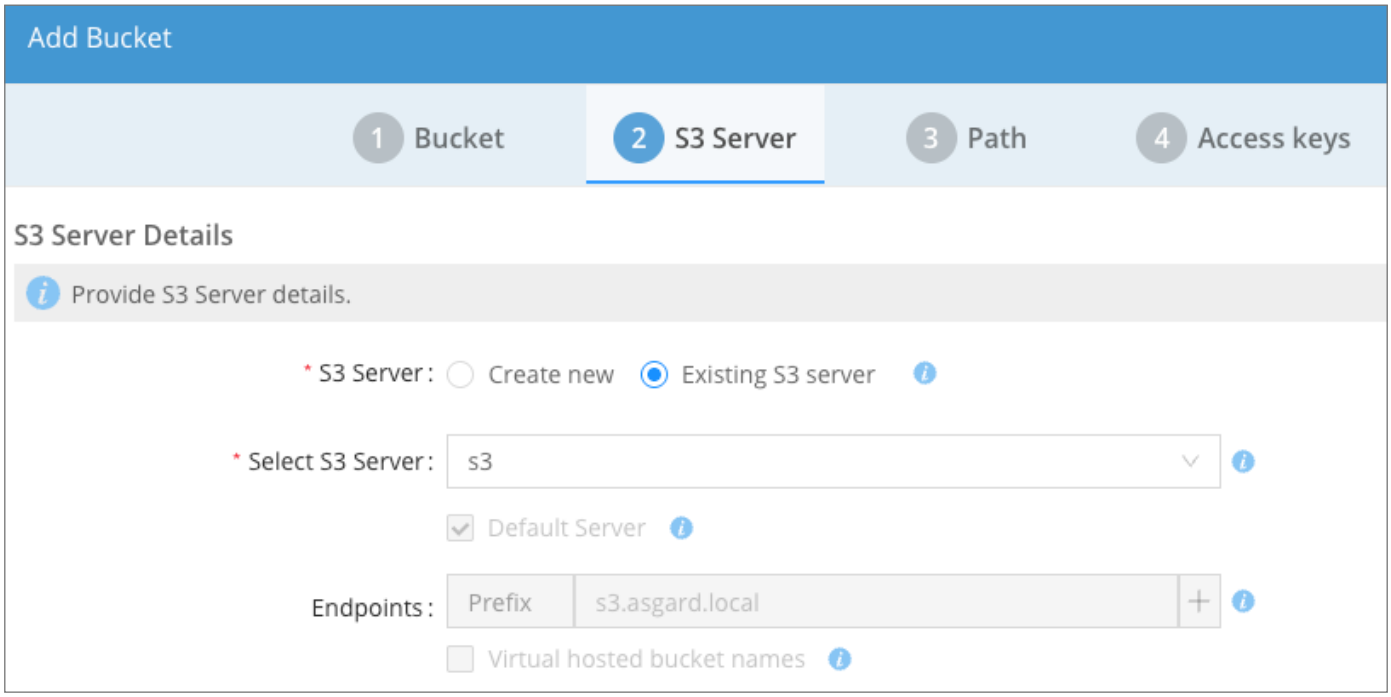

Figure 22. Create an S3 bucket using the GUI

- Provide a share where the bucket will reside, a path to the bucket folder in the share (if the folder does not exist, it will be created), set the bucket owner, modify the S3 access permissions (if needed), and then click next to move to step 4.

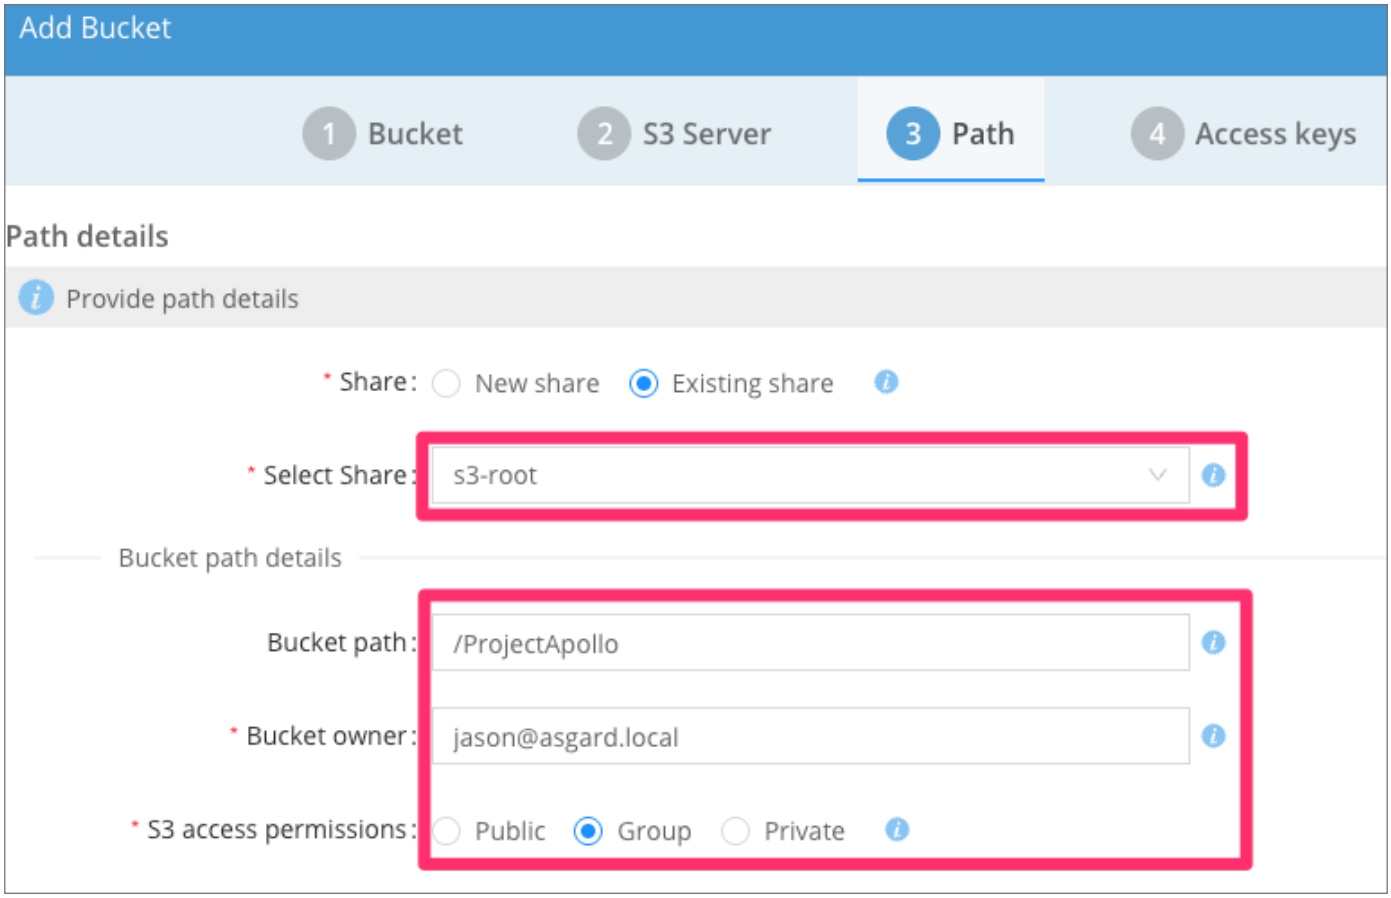

Figure 23. Create an S3 bucket using the GUI

- If AD is used as the identity provider, the owner will be an account in the format user@domain; for a local identity provider, it will be a local Hammerspace user with S3 access.
- Buckets and Bucket Containers cannot share the same path.

In this example, we see that multiple access keys are already present in the bucket configuration; do not remove them here, or they will be removed from the S3 server. If the bucket S3 access permissions allow it, these S3 users will automatically inherit access to this new bucket. Once done, click the Add button to create the bucket.

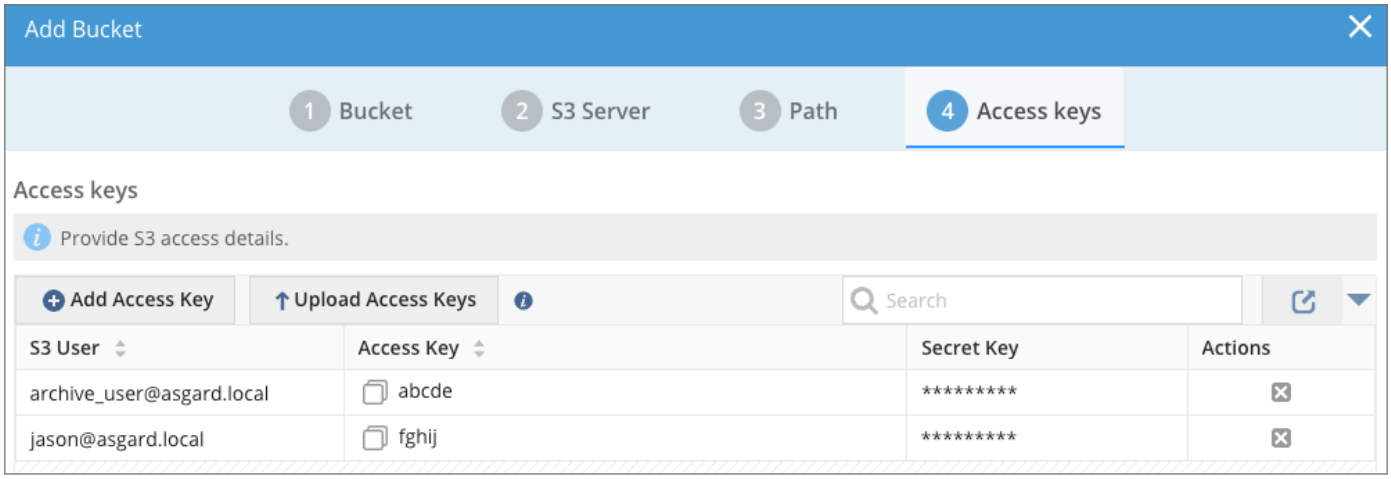

Figure 24. Create an S3 bucket using the GUI

- The S3 bucket can now be seen in the Buckets list.

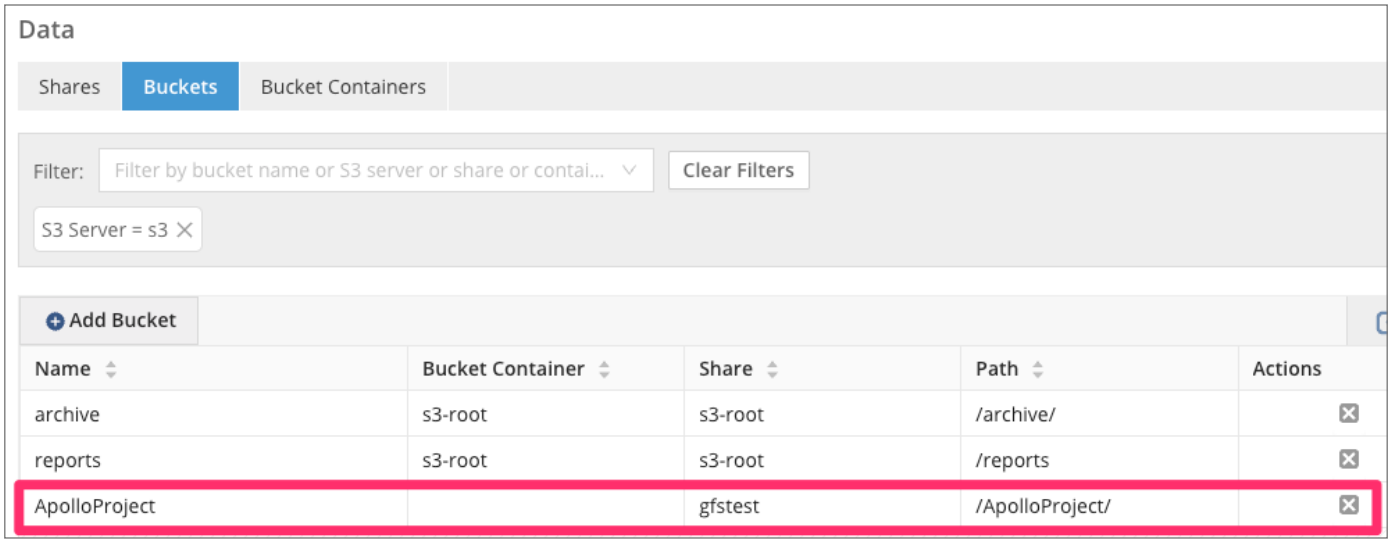

Figure 25. Create an S3 bucket using the GUI

5.3.3. Create an S3 Bucket Using the S3 API

Hammerspace supports using the S3 API to create buckets within S3 bucket containers. The following example shows this operation being performed using the Amazon S3 CLI:

```
# s3cmd mb s3://ApolloProject4
Bucket 's3://ApolloProject4/' created
# s3cmd ls
2025-03-24 20:23 s3://ApolloProject
2025-03-25 01:30 s3://ApolloProject2
2025-03-25 02:00 s3://ApolloProject4
2025-03-12 23:16 s3://archive
2025-03-10 20:22 s3://reports
```

The newly created bucket will also appear in the Hammerspace Data menu | Buckets tab shortly after being created using the S3 API:

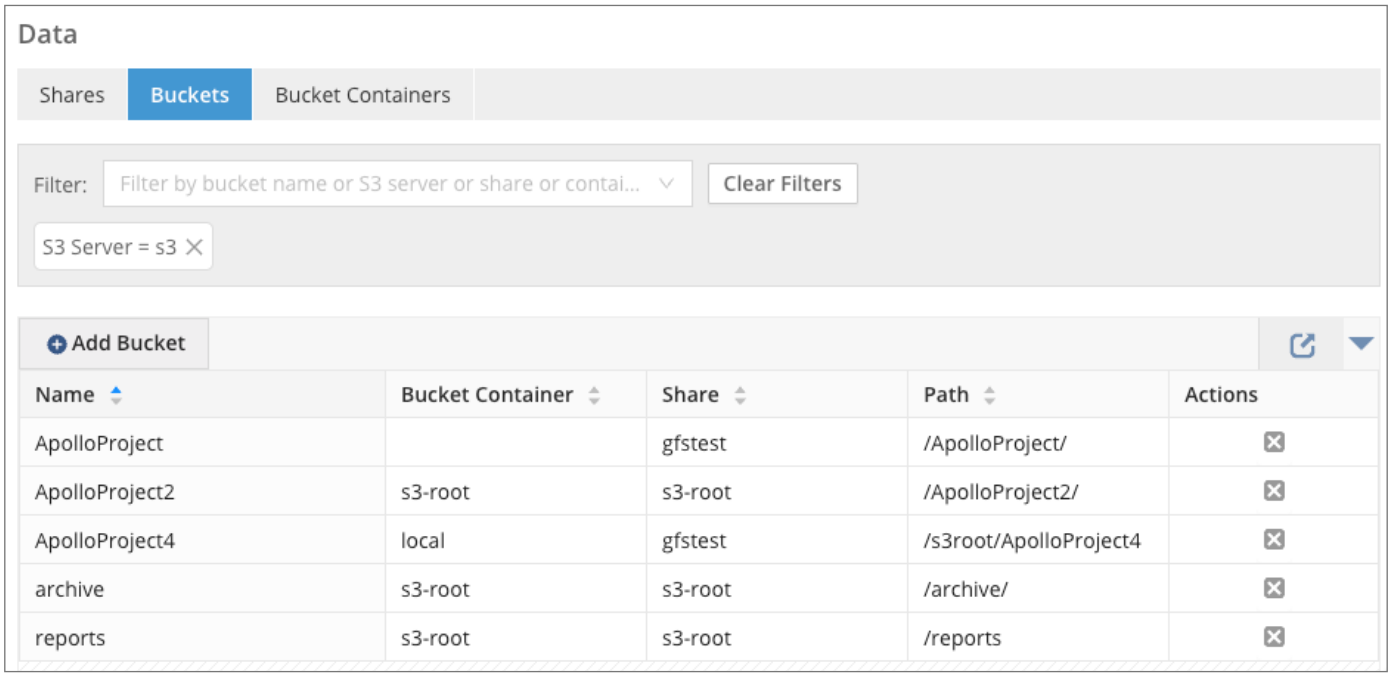

Figure 26. Create an S3 bucket using the S3 API

5.3.4. Create an S3 Bucket Using the Bucket Container

Hammerspace S3 Bucket Containers will automatically create S3 buckets for folders that are created within their root folder.

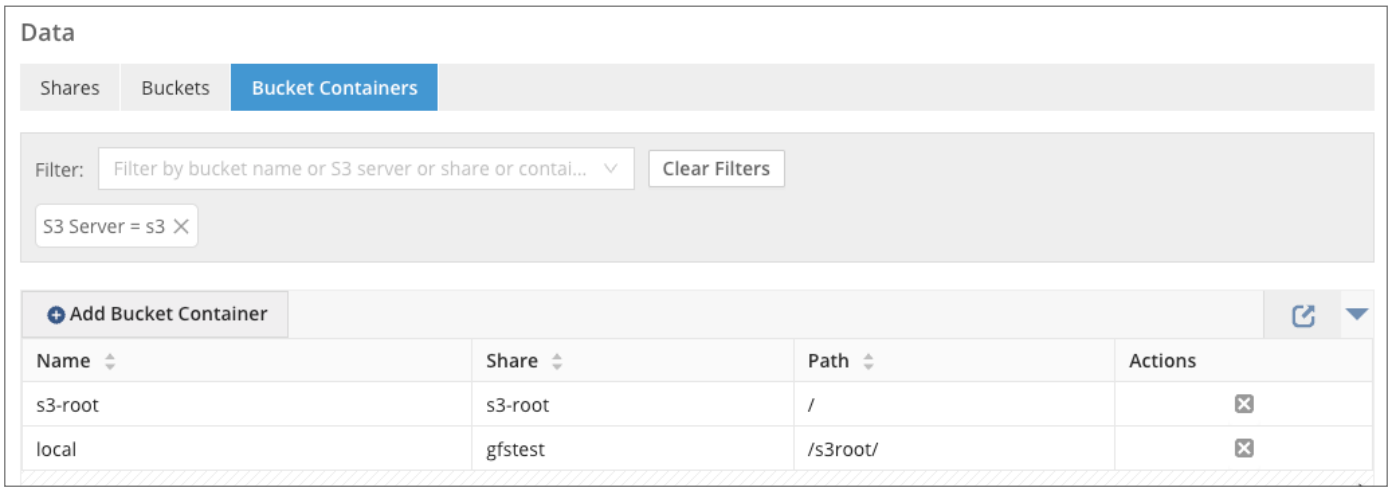

Figure 27. Create an S3 bucket using the bucket container

To create a bucket, simply create a folder within one of the shares and paths shown:

- \\gfstest\s3root\
- \\s3-root\

The bucket permissions and access keys will be inherited from the S3 Bucket Container.

5.4. Delete an S3 Bucket

There are three methods you can use to delete an S3 bucket. In this section, we will provide an overview of each one.

- **Method 1:** Navigate to the Buckets tab in the Data Menu and click the X button in the actions column shown for the bucket.

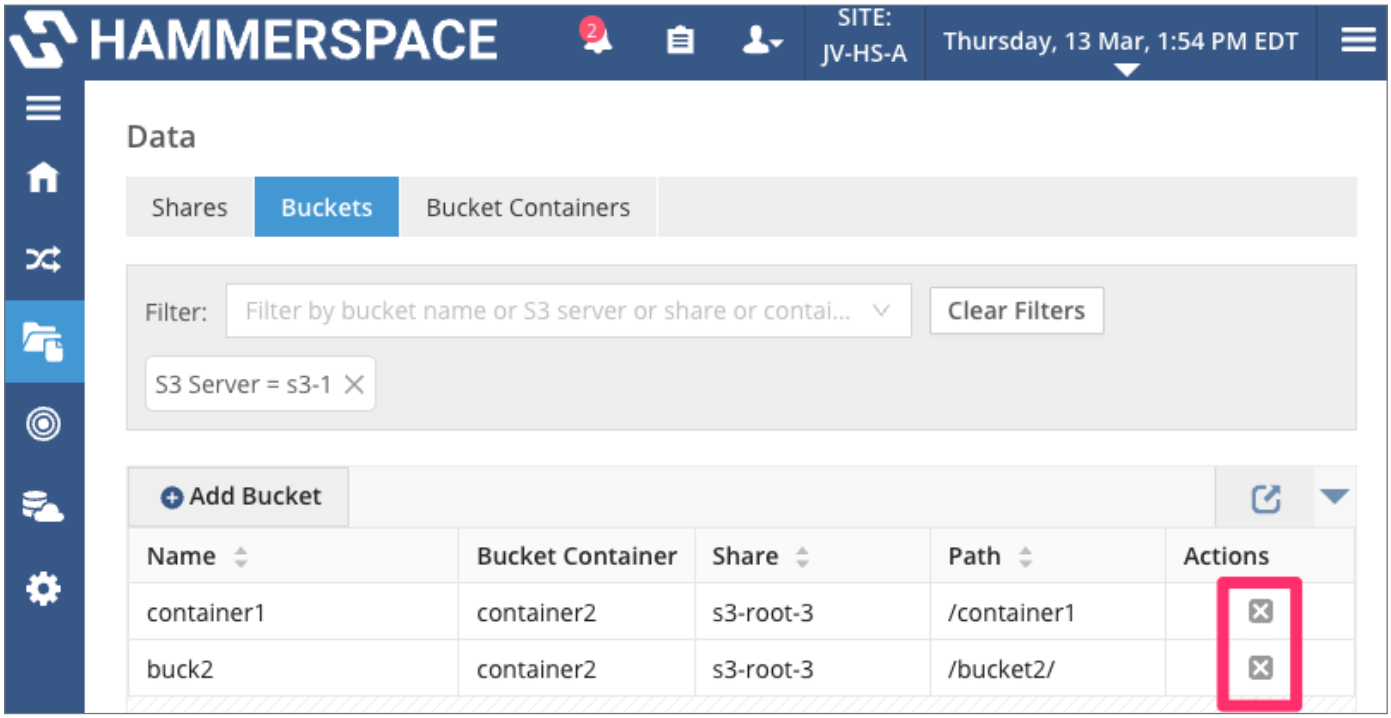

Figure 28. Delete an S3 bucket

- If the bucket is empty, you only need to confirm the bucket name and then click Delete bucket configuration. Note that the bucket folder will be removed.
- If the bucket has files in it, you need to confirm the bucket name, check the box to Delete data from the bucket, enter yes, and then click Delete bucket configuration. Note that the bucket folder and its contents will be removed. This action is shown in the following screenshot.

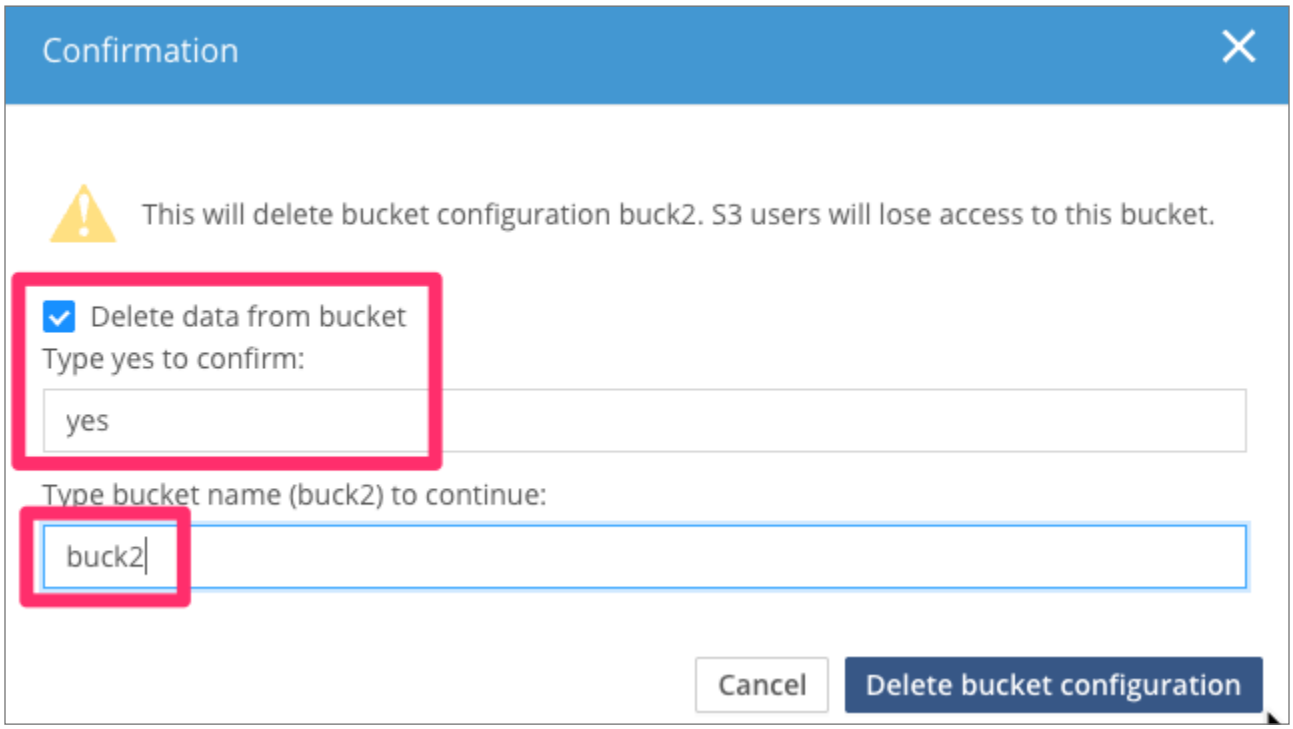

Figure 29. Delete an S3 bucket

- **Method 2:** For buckets within Bucket Containers that are configured to allow bucket deletion, you can use S3 commands to delete a bucket. Note that the bucket must be emptied first, as shown in this example:

```
# s3cmd --recursive --force rm s3://ApolloProject3
delete: 's3://ApolloProject3/launch7.txt'
# s3cmd rb s3://ApolloProject3
Bucket 's3://ApolloProject3/' removed
```

- **Method 3:** You can delete the S3 server, which will delete the bucket configuration but leave the bucket folders and files untouched. This is the only way to delete a Hammerspace S3 bucket endpoint without deleting the bucket and bucket contents, though these could be deleted by NFS or SMB clients that have the appropriate permissions.

It is important to remember that if your bucket folders are continued within a share that also serves NFS or SMB clients, then you should carefully review folder permissions to prevent bucket folders from accidental deletion. If an NFS or SMB client deletes the folder that backs an S3 bucket, S3 clients will see errors when they attempt to connect to it.

Appendices

Appendix A: Appendix

This appendix contains additional reference material for various topics discussed in this document.

A.1. Unix Permissions - Octal Values

The following table can be used as a reference for converting Unix permissions from their binary value to an octal value, which is needed to determine the desired umask setting for your S3 server that uses a local identity provider.

Permissions	Octal Value	Binary Value	Description
`---`	0	000	No permissions
`--x`	1	001	Execute
`-w-`	2	010	Write
`-wx`	3	011	Write and execute
`r--`	4	100	Read
`r-x`	5	101	Read and execute
`rw-`	6	110	Read and write
`rwx`	7	111	Read, write, and execute

To determine what value to use in your umask, you simply determine what you would need to subtract from 7 to reach your desired permission set.

For example, if we want to grant only read permissions for new files/folders (octal value **4** above), we would use a umask value of **3** as 7 - 3 = 4.

A.2. Creating Your S3 Umask Value for a Local Identity Provider

A Hammerspace S3 server that uses the local identity provider relies on a umask setting to determine the permissions for new files and folders that were created using the S3 protocol. The default setting is 0002, which can be explained as follows:

- The first value is 0, which sets the sticky bit to off. The sticky bit is a special file permission flag in Unix-like systems that, when set on a directory, prevents users (other than the owner or root) from deleting or moving files within that directory, even if they have write access.
- The second value sets the permissions for the file owner; a **0** (7 - 0 = octal permission value of 7) means that the owner will have full access to all files and folders that they create (read/write for files, read/write/execute for folders).

- The third value sets the permissions for the groups that the file owner is a member of; a **0** (7 - 0 = octal permission value of 7) means that members of those groups will have full access to all files and folders the owner creates (read/write for files, read/write/execute for folders).
- The fourth value sets the permissions for all other users; a **2** (7 - 2 = octal permission value of 5, the Hammerspace default) means that all other users will have only read/execute permissions for folders and read permissions for files.

- A umask of **0** (octal permission value of 7) grants read/write/execute permissions to folders and read/write permissions to files; this is expected.
- Hammerspace does not alter the existing permissions of any files or folders that were or are created by non-S3 clients.
- Umask is only used with S3 servers that use a local identity provider, not those that use Active Directory.

Appendix B: Additional Resources

The following resources will provide additional information about Hammerspace S3 servers or the S3 protocol in general.

- **Hammerspace Administration Guide**
- **Amazon S3 Tools**
- **Amazon S3 API**