

Marketplace Deployment Guide

Hammerspace 4.5

Amazon AWS

INTRODUCTION	3
SOLUTION OVERVIEW	3
DEPLOYING HAMMERSPACE IN AWS MARKETPLACE	4
ADDING ADDITIONAL DSX NODES.....	10
ACTIVATING LICENSES	11

Introduction

This guide details step by step instructions for deploying Hammerspace 4.5 in Amazon AWS Marketplace.

Solution Overview

Hammerspace is a set of software services that are deployed as virtual machines in order to manage data in a single global namespace.

The Hammerspace solution consists of the following main components:

Metadata Server (Anvil) deployed as a service and is responsible for managing all the metadata for the namespace.

Data Services (DSX) includes several different components:

1. **Portal** - Enables protocol access for NFS v3 and SMB clients. The portal is stateless except for SMB file locks that are maintained across portals.
2. **Mover** - The mover is stateless and runs on each DSX instance. It moves files non-disruptively between storage volumes.
3. **Cloud Connector** - The cloud connector provides connectivity to S3 compatible object storage and cloud storage from Amazon AWS, Microsoft Azure and Google Cloud Platform.
4. **Store** - The store service makes block storage available for use in the global namespace. Using DSX Store is a simple way to include local NVMe drives and block volumes.

Deploying Hammerspace in AWS Marketplace

This document describes how to deploy Hammerspace using AWS Marketplace.

Hammerspace can be deployed as a fully resilient, highly available solution as well as with flexible resilience for testing environment. Data resilience is driven by easy to use objectives after the product is deployed. For further details, see the Getting Start guide and also the Administration guide.

Pre-requisites

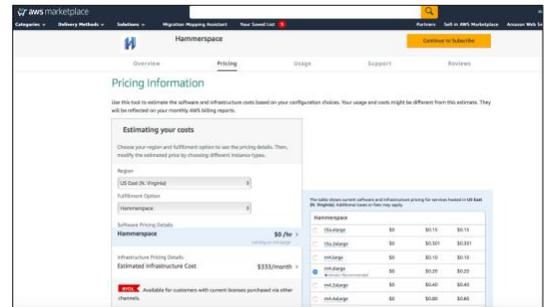
- Permissions to execute CloudFormation template and to create resources
- For a successful High Availability deployment, the HA nodes must be deployed in a network that can reach `ec2.[region].amazonaws.com`

1. Go to AWS Marketplace. Search for Hammerspace and select the solution you want to deploy.



Click on the solution you want to deploy. Currently there is only one Hammerspace solution listed in AWS Marketplace which requires purchasing a license through Hammerspace (BYOL).

2. On the “Product Overview” screen, there is more information available for the solution.

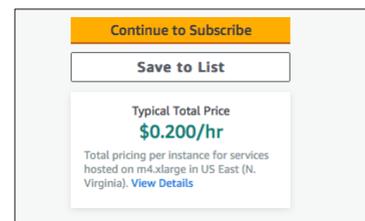


Note that several instance types have been qualified with Hammerspace. Which type you choose to use depends on what workload and the size of the environment it is serving. One great feature of Amazon AWS is that you can go in and change to a bigger/smaller version of the Instance Type later if you get it wrong.

Hammerspace has set the default instance type to m4.2xlarge.

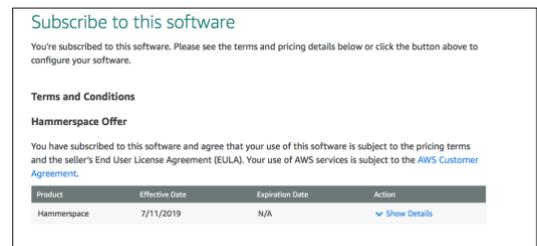
You can estimate the AWS costs for different instance types on the ‘Pricing Information’ screen.

3. Click on ‘Continue to Subscribe’.

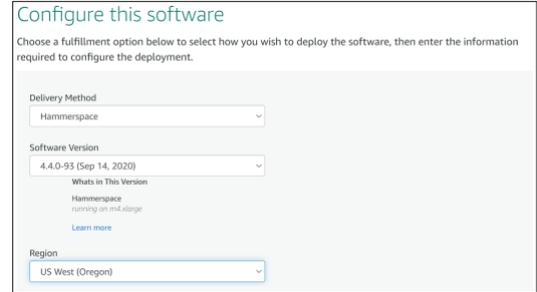


Solutions deployed through Marketplace can be used for up to 30 days before requiring a valid license.

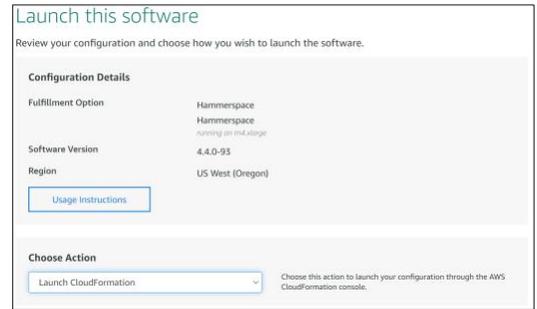
4. Subscribe to the software by accepting the terms of EULA, then click on ‘Continue to Configuration’ to go to the next screen.



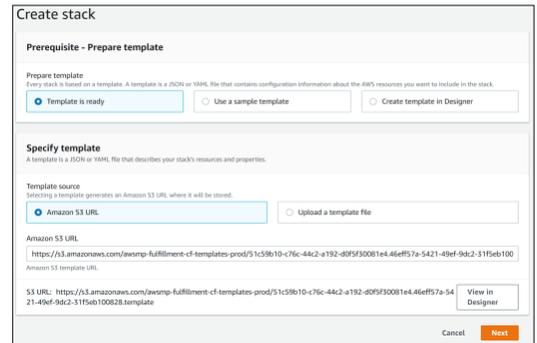
- Select the version you want to deploy and the region for deployment and then click on 'Continue to Launch'.



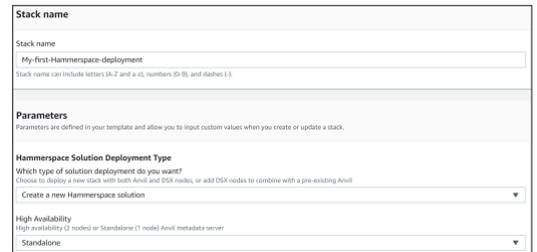
On the next screen, choose 'Launch CloudFormation' under 'Choose Action' and then hit 'Launch'



- On the Create Stack screen, click Next
All the required fields are already filled in and selected.



- To create a Hammerspace solution, choose a name for the stack and select 'Deployment type' as 'Create a new Hammerspace solution'.
Select for the metadata server to be deployed as a single instance or a highly available instance.



- Anvil**
Anvil can be deployed in small configurations on systems with as little as 16 GB of memory however for larger installations, or very active installations where metadata performance is critical, it is recommended to deploy a larger instance.
Enter a boot disk size and a metadata disk size. The metadata disk holds all the metadata for the Hammerspace file system and needs to be of sufficient size. A 100 GB metadata disk is large enough for approximately 25-100 million files.



9. Data Services – DSX

The DSX is responsible for storing data, moving data, deduplication, compression and serving data, this can typically all be easily done with a 4 CPU/16 GB environment. If the environment has more than 500 SMB clients, it is recommended to move to a larger instance size or additional DSX nodes.

The data disk is completely dependent on the initial selected size. Additional data disks can be added without downtime after the DSX is deployed.

Up to 60 DSX nodes are supported however the CloudFormation template only allows for 10 nodes at a time to be deployed.

The Anvil Cluster IP is only required when adding additional DSX nodes to an already deployed solution.



DSX Data Services Nodes

EC2 Instance Type
Choose a supported EC2 instance type
m4.xlarge (4 vCPUs, 16 GB Mem)

Root/OS Disk Size
Size in GB (minimum 100, maximum 1000)
100

Data Store Size
Size in GB (minimum 100, maximum 16384)
200

DSX Instance Count
How many DSX instances to create as part of this stack? (minimum 1, maximum 10)
1

Anvil Cluster IP (leave blank unless adding DSX nodes to an existing installation)
Internal IP of existing Anvil instance

10. Under 'Networking and Security', select network details according to your configuration. This section of the CloudFormation template is specific to the environment where Hammerspace is being deployed:

- Availability Zone
- VPC ID
- Subnet



Networking and Security

Availability Zone
Location where resources will be created

VPC ID
Choose which Virtual Private Cloud to deploy your instances in

Data/Management Subnet ID
Choose a subnet within your Virtual Private Cloud

Data/Management Subnet Gateway IP (leave blank unless deploying a new High Availability solution)
Gateway IP Address of the subnet chosen for Data/Management

High Availability Subnet CIDR (leave blank unless deploying a new High Availability solution)
Choose an unused CIDR range within the VPC to be used only for HA. (Example: 10.0.254.240/28)

Security Group ID/CIDR
Use 0.0.0.0/0 for open access

KeyPair Name (optional)
Name of an existing EC2 KeyPair to enable SSH access to the instances (optional)

High Availability Subnet is required when deploying a HA deployment. The HA subnet is only used for Anvil to Anvil traffic related to resilience. This subnet will be automatically configured with a Security Group for increased security.

Security Group CIDR must include the Anvil, DSX nodes as well as all clients that may access data.

AWS Marketplace does not allow password authentication for SSH access. Instead, a key pair is used to securely access your Anvil instance using SSH. Choose an existing key pair or create a new key pair to allow SSH access.

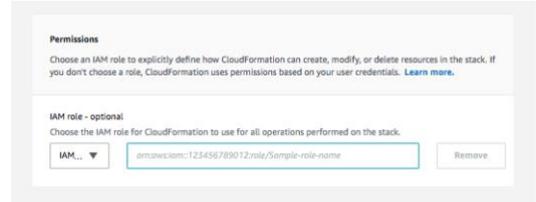
While selecting a Key Pair is strictly speaking optional, it is highly recommended to do at this stage as adding key pair information in AWS is not possible to do later. In that case or to allow SSH access to additional users, you can

add IAM users using the IAMAdminGroupUrl once the stack has been created.

Press Next to continue

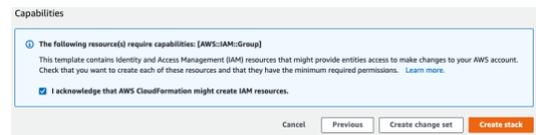
- It is recommended to add tags to identify your deployments.

If you don't have all the permissions needed to create the stack, you can provide access to CloudFormation to make those calls on your behalf. In that case use the 'Permissions' section to set the IAM role for CloudFormation to use.



For the rest of the 'Advanced options', make any changes as needed for your setup.

- Go over the details of the deployment and accept the disclaimer at the end, then click 'Create stack'.



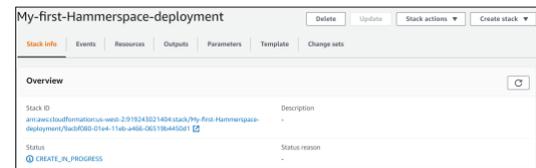
You will need to wait for about 5-15 minutes to allow the nodes and networking to be setup and all the nodes to come online.

High Availability deployments take longer to deploy and initialize the first time.

- Once the instances have been created, this is what you will see.

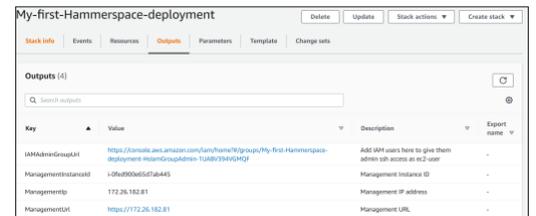
Navigate to the Resources tab to monitor the Resource creation steps.

Keep in mind that the first time Hammerspace starts, several initialization steps are executed. Sub-sequent reboots or stop/start instance operations are significantly faster.



- Click on the 'Outputs' tab at the top of the page to get information on the instances created.

Follow the HTTPS link Management URL to login to the GUI as 'admin'. Use the provided "Management Instance ID" (starting with i-) as the password.



For SSH access, ssh in as `ec2-user@<ManagementIP>` using your key pair credentials to get admin CLI access. No password is requested. E.g.
`ssh -i "janedoe.pem" ec2-user@ManagementIp`

With Hammerspace successfully deployed, use the Getting Started Guide to get up and running in a few minutes.

Adding additional DSX nodes

A Hammerspace deployment can be expanded with capacity leveraging the scale-out deployment of DSX. Up to 60 DSX nodes can be configured in a single Hammerspace deployment.

Use the Marketplace and CloudFormation template to deploy additional nodes. Only a few of the fields are needed when expanding capacity

- Anvil Management IP address (same as the GUI management address)
- Networking (recommended to deploy in the same region)

All the Anvil specific fields will be ignored when selecting Adding DSX nodes

1. Create a new stack deployment

Choose 'Add Data Services (DSX) to a Hammerspace solution' from the 'Deployment Type' drop down menu.

Follow the instructions for filling in the remaining fields as outlined earlier and click Deploy.

It is recommended to use tags on deployments that might be sharing a Hammerspace solution.

2. In order for the Anvil to discover and add the DSX node(s), you need to specify the 'Cluster Management IP' of the Anvil in the field 'Cluster IP'. This IP is generally the same as the IP used for the GUI access.

Specify the capacity for the initial DSX data store.

Specify the number of DSX nodes to be added to the Hammerspace Solution in the 'DSX instance count' field.

Stack name

Stack name

Adding-DSX-nodes-to-Hammerspace

Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

Hammerspace Solution Deployment Type

Which type of solution deployment do you want?

Choose to deploy a new stack with both Anvil and DSX nodes, or add DSX nodes to combine with a pre-existing Anvil.

Add Data Services (DSX) to a Hammerspace solution

DSX Data Services Nodes

EC2 Instance Type

Choose a supported EC2 instance type.

m4.large (4 vCPUs, 16 GiB Mem)

Boot/OS Disk Size

100-16384 GiB (minimum 100, maximum 16384)

100

Data Store Size

100-16384 GiB (minimum 100, maximum 16384)

200

DSX Instance Count

How many DSX instances to create as part of this stack? (minimum 1, maximum 10)

2

Anvil Cluster IP (leave blank unless adding DSX nodes to an existing installation)

external IP of existing Anvil instance

172.26.182.81

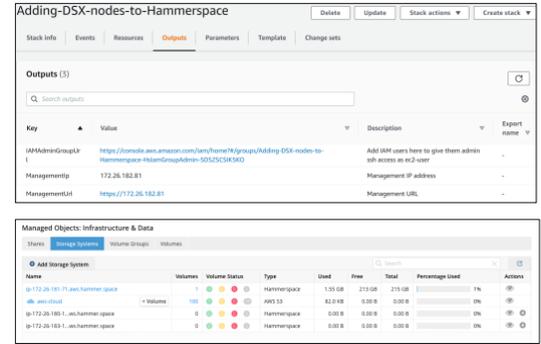
- Once the stack is deployed, the DSX nodes are discovered but not automatically added to the Anvil configuration as an extra step of security.

Click the + icon under the Action column to add the DSX node.

Click +Volume icon to add the new storage volumes to the usable capacity.

Adding and removing capacity is a non-disruptive operation and can be done at any time.

New storage volumes are by default automatically used and no additional configuration steps are required.



Activating Licenses

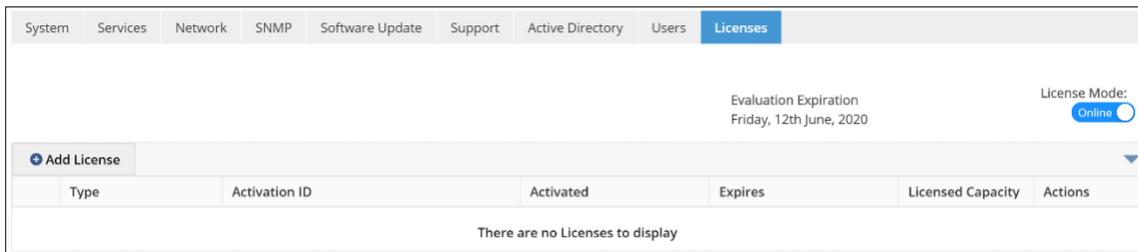
To activate your license please refer to the entitlement email received from Hammerspace, note the Activation ID and follow the steps below.

Please contact support if you did not receive an Activation ID.

Hammerspace licenses can be activated using the GUI or the CLI.

GUI

- Navigate to Administration -> Licenses tab



- Click Add License and add the activation ID. The default license type is Node based.

Add License

Type Node Capacity

Activation ID

Add

It is normal for the Activation to take up to several minutes.

3. After a successful activation, the license is installed.

Type	Activation ID	Activated	Expires	Licensed Capacity	Actions
Node	aaaa-bbbb-cccc-dddd-eeee	Friday, 14th August, 2020	Monday, 14th August, 2023	—	

License Mode: Online

4. To add a capacity license, choose the Capacity option and enter the capacity you want to license.

Add License

Type Node Capacity

Activation ID

Capacity TB

Add

5. Once the license is activated, the following screen will be displayed.

System Capacity Report

File Capacity Licensed 10 TB Used 0 B Lent 0 B Over 0 B

Object Capacity Licensed 0 B Used 0 B Borrowed 0 B Over 0 B

License Mode: Online

Type	Activation ID	Activated	Expires	Licensed Capacity	Actions
File Capacity	aaaa-bbbb-cccc-dddd-eeee	Tuesday, 17th November, 2...	Wednesday, 24th August, 2...	10 TB	