

Deploy the ExtraHop Explore Appliance in AWS

Published: 2019-01-08

In this guide, you will learn how to launch the ExtraHop Explore appliance AMI in your Amazon Web Services (AWS) environment, and join multiple Explore appliances to create an Explore cluster.


System requirements

Your environment must meet the following requirements to deploy a virtual Explore appliance in AWS:

- An AWS account
- Access to the Amazon Machine Image (AMI) of the ExtraHop Explore appliance
- An Explore appliance product key
- An AWS instance type that most closely matches the Discover appliance VM size, as follows:

Appliance	Recommended Instance Types
EXA 5100v	m4.2xlarge (8 vCPU and 32 GB RAM) m4.4xlarge (16 vCPU and 64 GB RAM)

- A datastore size between 200 GB and 1.2 TB for the m4.2xlarge instance type or between 200 GB and 2.5 TB for the m4.4xlarge instance type.

 **Important:** If you want to deploy more than one ExtraHop virtual appliance, create the new instance with the original deployment package or clone an existing instance that has never been started.


Create the Explore instance in AWS

Before you begin


The Amazon Machine Images (AMIs) of ExtraHop appliances are not publicly shared. Before you can start the deployment procedure, you must send your AWS account ID to support@extrahop.com. Your account ID will be linked to the ExtraHop AMIs.

1. Sign in to AWS with your user name and password.
2. Click **EC2**.
3. In the left navigation panel, under Images, click **AMIs**.
4. Above the table of AMIs, change the **Filter** from **Owned by Me** to **Private Images**.
5. In the filter box, type `EXTRAHOP` and then press ENTER.
6. Select the checkbox next to the ExtraHop Explore appliance AMI and click **Launch**.
7. On the Choose an Instance Type page, select **m4.2xlarge**, and then click **Next: Configure Instance Details**.
8. In the **Number of instances** text box, type the number nodes in your Explore cluster.
9. Click the **Network** drop-down list and select the default setting or one of the VPCs for your organization.
10. From the **Shutdown behavior** drop-down list, select **Stop**.
11. Click the **Protect against accidental termination** checkbox.
12. Optional: Click the **IAM role** drop-down list and select an IAM role.

13. Optional: If you launched into a VPC and want to add more than one interface, scroll down to the Network Interfaces section and click **Add Device** to add additional interfaces to the instance.

 **Note:** If you add more than one interface, make sure that each interface is on a different subnet.

14. Click **Next: Add Storage**.
15. In the Size (GiB) field for the root volume, type the size of the storage volume. The minimum datastore size is 186 GiB (200 GB) and the maximum datastore size is 2328 GiB (2.5 TB). If you specify a storage volume greater than 2047 GiB, you can safely ignore any warning messages stating that root volumes of this size might result in the instance not booting successfully.

 **Note:** Consult with your ExtraHop sales representative or Technical Support to determine the datastore disk size that is best for your needs.

16. From the Volume Type drop-down menu, select either **Magnetic** or **General Purpose SSD (GP2)**. You must select **General Purpose SSD (GP2)** if you specify a size greater than 1024 GiB. GP2 provides better storage performance, although at a higher cost.

17. Click **Next: Tag Instance**.

18. In the Value field, type a name for the instance.

19. Click **Next: Configure Security Group**.

20. On the Configure Security Group page, create a new security group or add ports to an existing group. If you already have a security group with the required ports for ExtraHop, you can skip this step.

- a) Select either **Create a new security group** or **Select an existing security group**. If you choose to edit an existing group, select the group you want to edit. If you choose to create a new group, type a name for the Security group and type a Description.
- b) Click the **Type** drop-down list, and select a protocol. Type the port number in the **Port Range** field.
- c) For each additional port needed, click the **Add Rule** button. Then click the **Type** drop-down list, select a protocol, and type the port number in the **Port Range** field.

The following ports must be open for the Explore appliance AWS instance:

- TCP port 443: Enables you to administer the Explore appliance through the Web UI. Requests sent to port 80 are automatically redirected to HTTPS port 443.
- TCP port 9443: Enables Explore nodes to communicate with other Explore nodes in the same cluster.

21. Click **Review and Launch**.

22. Select **Make General Purpose (SSD)...(recommended)** and click **Next**.

 **Note:** If you select **Make General Purpose (SSD)...(recommended)**, you will not see this step on subsequent instance launches.

23. Scroll down to review the AMI details, instance type, and security group information, and then click **Launch**.

24. In the pop-up window, click the first drop-down list and select **Proceed without a key pair**.

25. Click the **I acknowledge...** checkbox and then click **Launch Instance**.




26. Click **View Instances** to return to the AWS Management Console.

From the AWS Management Console, you can view your instance on the Initializing screen.

Under the table, on the **Description** tab, you can find an IP or hostname for the Explore appliance that is accessible from your environment.

Configure the Explore appliance

After you obtain the IP address for the Explore appliance, log into the Explore Admin UI through the following URL: `https://<explore_ip_address>/admin` and complete the following recommended procedures.

- [Register the Explore appliance](#) 
- [Create an Explore cluster](#)
- [Configure the system time](#) 
- [Configure email notifications](#) 
- [Pair the Explore appliance to all Discover and Command appliances](#)
- [Send record data to the Explore appliance](#)


Create an Explore cluster

If you are deploying more than one Explore appliance, join the appliances together to create a cluster. For the best performance, data redundancy, and stability, you must configure at least three Explore appliances in an Explore cluster.


In the following example, the Explore appliances have the following IP addresses:

- Node 1: 10.20.227.177
- Node 2: 10.20.227.178
- Node 3: 10.20.227.179


You will join nodes 2 and 3 to node 1 to create the Explore cluster.

 **Important:** Each node that you join must have the same configuration (physical or virtual) and ExtraHop firmware version.

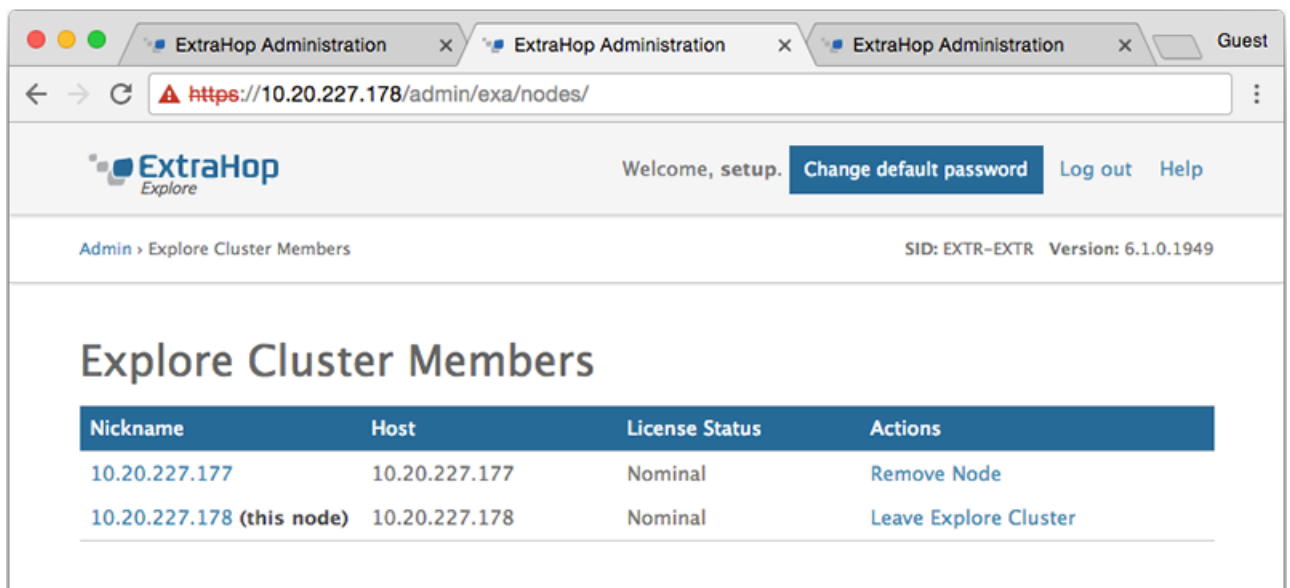
1. Log into the Admin UI of all three Explore appliances with the setup user account in three separate browser windows or tabs.
2. Select the browser window of node 1.
3. In the Status and Diagnostics section, click **Fingerprint** and note the fingerprint value. You will later confirm that the fingerprint for node 1 matches when you join the remaining two nodes.
4. Select the browser window of node 2.
5. In the Explore Cluster Settings section, click **Join Cluster**.
6. In the Host field, type the hostname or IP address of node 1 and then click **Continue**.

 **Note:** For cloud-based deployments, be sure to type the IP address listed in the Interfaces table on the Connectivity page.


7. Confirm that the fingerprint on this page matches the fingerprint you noted in step 3.



8. In the Setup Password field, type the password for the node 1 `setup` user account and then click **Join**.
When the join is complete, the Explore Cluster Settings section has two new entries: **Explore Cluster Members** and **Data Management**.
9. Click Explore Cluster Members. You should see node 1 and node 2 in the list.



10. In the Status and Diagnostics section, click **Explore Cluster Status**. Wait for the Status field to change to `Green` before adding the next node.
11. Repeat steps 5 - 11 to join each additional node to the new cluster.
 - Tip:** To avoid creating multiple clusters, always join a new node to the existing cluster and not to another single appliance.
12. When you have added all of your Explore appliances to the cluster, click **Explore Cluster Members** in the Explore Cluster Settings section. You should see all of the joined nodes in the list, similar to the following figure.



13. In the Explore Cluster Settings section, click **Data Management** and make sure that **Replication Level** is set to **1** and **Shard Reallocation** is **ON**.


Connect the Explore appliance to Discover and Command appliances

After you deploy the Explore appliance, you must establish a connection from all ExtraHop Discover and Command appliances to the Explore appliance before you can query records.

Important: If you have an Explore cluster of three or more Explore nodes, connect the Discover appliance to each Explore node so that the Discover appliance can distribute the workload across the entire Explore cluster.

Note: If you manage all of your Discover appliances from a Command appliance, you only need to perform this procedure from the Command appliance.

1. Log into the Admin UI of the Discover or Command appliance .
2. In the ExtraHop Explore Settings section, click **Connect Explore Appliances**.
3. Click **Add New**.
4. In the Explore node field, type the hostname or IP address of any Explore appliance in the Explore cluster.
5. For each additional Explore appliance in the cluster, click **Add New** and enter the individual hostname or IP address in the corresponding Explore node field.



6. Click **Save**.
7. Confirm that the fingerprint on this page matches the fingerprint of node 1 of the Explore cluster.
8. In the Explore Setup Password field, type the password for the Explore node 1 `setup` user account and then click **Connect**.
9. When the Explore Cluster settings are saved, click **Done**.

Next steps

Important: If you only deployed a single Explore appliance, after you connect to your Discover or Command appliance, you must log into the Admin UI on the Explore appliance and set the **Explore Cluster Settings > Data Management > Replication Level** to **0**.

Send record data to the Explore appliance

After your Explore appliance is connected to all of your Discover and Command appliances, you must configure the type of records you want to store.

See [Records](#) for more information about Explore configuration settings, how to generate and store records, and how to create record queries.