

# Deploy the ExtraHop packetstore in AWS

Published: 2025-01-23

In this guide, you will learn how to launch the ExtraHop packetstore AMI in your Amazon Web Services (AWS) environment.

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

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

| Packetstore | Supported Instance Types |
|-------------|--------------------------|
| ETA 1150v   | m5.xlarge, m5.2xlarge    |



**Tip:** You can resize your instance without redeploying the packetstore. See the [AWS documentation](#) for details.

### 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 your ExtraHop representative. Your account ID will be linked to the ExtraHop AMI.

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 packetstore AMI and click **Launch**.
7. Select one of the following supported instance types:

| Instance Type | Details                                                                                                                   |
|---------------|---------------------------------------------------------------------------------------------------------------------------|
| m5.xlarge     | Recommended for most installations.                                                                                       |
| m5.2xlarge    | Select <b>m5.2xlarge</b> if you need greater throughput. The cost for this instance is higher than for <b>m5.xlarge</b> . |

8. Click the **Network** drop-down menu and select the default setting or one of the VPCs for your organization.
9. Optional: Click the **IAM role** drop-down menu and select an IAM role.
10. From the **Shutdown behavior** drop-down menu, select **Stop**.
11. Select the **Protect against accidental termination** checkbox.
12. Click **Next: Add Storage**.
13. In the Size (GiB) field for the root volume, type the size of the storage volume. The minimum packetstore size is 1000 GiB (1 TB) and the maximum datastore size is 2047 GiB (2 TB).
14. From the Volume Type drop-down menu, select either **Magnetic** or **General Purpose SSD (GP2)**. If you specify a size greater than 1024 GiB, you must select **General Purpose SSD (GP2)**. GP2 provides better storage performance, although at a higher cost.
15. Click **Next: Add Tags**.
16. Click **Add Tag**.
17. In the Value field, type a name for the instance.

18. Click **Next: Configure Security Group**.
19. Select an existing security group or create a new security group with the required ports.
20. Click **Add Rule** and add the following ports:

| Type       | Port Range |
|------------|------------|
| SSH        | 22         |
| Custom TCP | 443        |
| Custom TCP | 2003       |
| Custom UDP | 2003       |

TCP ports 22, and 443 are required to administer the ExtraHop system. TCP and UDP port 2003 is required for the packet forwarder.

21. Click **Review and Launch**.
22. Select the boot volume option you selected in step 14 and then click **Next**.



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

23. 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 menu and select **Proceed without a key pair**.
25. Click the **I acknowledge...** checkbox and then click **Launch Instances**.
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 address or hostname for the ExtraHop system that is accessible from your environment.

#### Next steps

- [Register your ExtraHop system](#)
- Review the [Trace Appliance Post-deployment Checklist](#).
- [Connect the Command and Discover appliances to the Trace appliance](#).
- Configure remote packet capture (RPCAP) to forward traffic from remote devices to your virtual packetstore. For more information, see [Configure RPCAP for an ExtraHop packetstore](#).
- (Recommended) Configure [AWS traffic mirroring](#) to copy network traffic from your EC2 instances to a RPCAP/ERSPAN/VXLAN/GENEVE interface on your packetstore.

## Create a traffic mirror target

Complete these steps for each Elastic network interface (ENI) you created.

1. In the AWS Management Console, in the top menu, click **Services**.
2. Click **Networking & Content Delivery > VPC**.
3. In the left pane, under Traffic Mirroring, click **Mirror Targets**.
4. Click **Create traffic mirror target**.
5. Optional: In the Name tag field, type a descriptive name for the target.
6. Optional: In the Description field, type a description for the target.
7. From the Target type drop-down menu, select Network Load Balancer.
8. From the Target drop-down menu, select the Elastic Load Balancer you previously created.
9. Click **Create**.

Note the Target ID for each elastic load balancer. You will need the ID when you create a traffic mirror session.

## Create a traffic mirror filter

You must create a filter to allow or restrict traffic from your ENI traffic mirror sources to your ExtraHop system.

We recommend the following filtering rules to help avoid mirroring duplicate frames from peer EC2 instances that are in a single VPC to the sensor.

- All outbound traffic is mirrored to the sensor, whether the traffic is sent from one peer device to another on the subnet or if the traffic is sent to a device outside of the subnet.
- Inbound traffic is only mirrored to the sensor when the traffic is from an external device. For example, this rule ensures that an app server request is not mirrored twice: once from the sending app server and once from the database that received the request.
- Rule numbers determine the order in which the filters are applied. Rules with lower numbers, such as 100, are applied first.

 **Important:** These filters should only be applied when mirroring all of the instances in a CIDR block.

1. In the AWS Management Console, in the left pane under Traffic Mirroring, click **Mirror Filters**.
2. Click **Create traffic mirror filter**.
3. In the Name tag field, type a name for the filter.
4. In the Description field, type a description for the filter.
5. Under Network services, select the **amazon-dns** checkbox.
6. In the Inbound rules section, click **Add rule**.
7. Configure an inbound rule:
  - a) In the Number field, type a number for the rule, such as 100.
  - b) From the Rule action drop-down menu, select **reject**.
  - c) From the Protocol drop-down menu, select **All protocols**.
  - d) In the Source CIDR block field, type the CIDR block for the subnet.
  - e) In the Destination CIDR block field, type the CIDR block for the subnet.
  - f) In the Description field, type a description for the rule.
8. In the Inbound rules sections, click **Add rule**.
9. Configure an additional inbound rule:
  - a) In the Number field, type a number for the rule, such as 200.
  - b) From the Rule action drop-down menu, select **accept**.
  - c) From the Protocol drop-down menu, select **All protocols**.
  - d) In the Source CIDR block field, type 0 . 0 . 0 . 0 / 0.
  - e) In the Destination CIDR block field, type 0 . 0 . 0 . 0 / 0.
  - f) In the Description field, type a description for the rule.
10. In the Outbound rules section, click **Add rule**.
11. Configure an outbound rule:
  - a) In the Number field, type a number for the rule, such as 100.
  - b) From the Rule action drop-down menu, select **accept**.
  - c) From the Protocol drop-down menu, select **All protocols**.
  - d) In the Source CIDR block field, type 0 . 0 . 0 . 0 / 0.
  - e) In the Destination CIDR block field, type 0 . 0 . 0 . 0 / 0.
  - f) In the Description field, type a description for the rule.
12. Click **Create**.

## Create a traffic mirror session

You must create a session for the AWS resource that you want to monitor. You can create a maximum of 500 traffic mirror sessions per sensor. You only need to create one inbound and one outbound rule, and both must allow all traffic from any protocol.

 **Important:** To prevent mirror packets from being truncated, set the traffic mirror source interface MTU value to 54 bytes less than the traffic mirror target MTU value for IPv4 and 74 bytes less than the traffic mirror target MTU value for IPv6. For more information about configuring the network MTU value, see the following AWS documentation: [Network Maximum Transmission Unit \(MTU\) for Your EC2 Instance](#).

1. In the AWS Management Console, in the left pane, under Traffic Mirroring, click **Mirror Sessions**.
2. Click **Create traffic mirror session**.
3. In the Name tag field, type a descriptive name for the session.
4. In the Description field, type a description for the session.
5. From the Mirror source drop-down menu, select the source ENI.  
The source ENI is typically attached to the EC2 instance that you want to monitor.
6. From the Mirror target drop-down menu, select the traffic mirror target ID generated for the target ENI.
7. In the Session number field, type 1.
8. For the VNI field, leave this field empty.  
The system assigns a random unique VNI.
9. For the Packet length field, leave this field empty.  
This mirrors the entire packet.
10. From the Filter drop-down menu, select the ID for the traffic mirror filter you created.
11. Click **Create**.