

Deploy an ExtraHop Flow Sensor with AWS

Published: 2024-01-19

This guide explains how to deploy the ExtraHop flow sensor virtual appliance (EFC 1291v) on the Amazon Web Services (AWS) platform.

The EFC 1291v is designed to connect to Reveal(x) 360 and collect flow-based traffic from your network. Packet analysis is not available.

Your environment must meet the following requirements to deploy an EFC 1291v appliance in AWS:

- An AWS account
- Access to the Amazon Machine Image (AMI) of the ExtraHop 1100v appliance
- An EFC 1291v appliance product key
- An AWS instance type that most closely matches the EFC appliance VM size, as follows:

Appliance	Supported Instance Type
Reveal(x) EFC 1291v	c5.xlarge (4 vCPU and 8 GB RAM)

Deployment overview

Collecting flow logs requires the following configuration setup.

1. Configure an IAM policy and IAM role.
2. Deploy the ExtraHop flow sensor instance in AWS.
3. Download and configure an ExtraHop-supplied Lambda function. The Lambda function runs whenever new flow logs become available and then relays any new events to your sensor. See the following AWS documentation for more information: [Using AWS Lambda with Amazon Kinesis Firehose](#).
4. Enable VPC Flow Logs publishing for a set of VPCs in your environment.
5. Add a Lambda trigger.
6. Optional: Configure Route 53.

Configure an IAM permission policy and IAM role

1. Create an [IAM policy](#) through the JSON tab with the following parameters:

```

{
  "Statement": [
    {
      "Action": [
        "ec2:Describe*"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ],
  "Version": "2012-10-17"
}

```

2. [Create an IAM role](#) and attach the permission policy.

The ExtraHop system requires an instance IAM role to correlate IP addresses from flow logs to instances, gateways, and Lambdas.

3. Click the **Trust relationships** tab and edit the trust policy to appear as follows:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```


Deploy the sensor AMI

1. Deploy a Reveal(x) EDA 1100V by following the [Deploy an ExtraHop sensor in AWS](#) guide.

The EDA 1100V is a packet sensor that becomes a flow logs sensor when the license is entered. The sensor will no longer process packets.



Tip: You can subscribe to the Reveal(x) 1100v (BYOL) software through the AWS Marketplace.

2. Log in to the ExtraHop system through `https://<extrahop-hostname-or-IP-address>`. The username is setup and the password is the string of numbers after the i- in the instance ID.
3. Follow the prompts to accept the license agreement, enter the product key, change the default setup and shell user account passwords, connect to ExtraHop Cloud Services, and connect to Reveal(x) 360.
4. Click the System Settings icon , and then click **All Administration**.
5. In the Access Settings section, click **API Access**.
6. In the Generate an API Key section, type a description for the new key and then click **Generate**.
7. Generate the Flow Log secret from the REST API Explorer.
 - a) Click **Open ExtraHop API Explorer**.
 - b) Click **Enter API Key** and then paste or type your API key into the API Key field.
 - c) Click **Authorize** and then click **Close**.
 - d) Click **ExtraHop** and then click **POST /extrahop/flowlogs/secret**.
 - e) Click **Try it out** and then click **Send Request**.
 - f) In the Response Body section, view and record the `secret` value. You will need the secret for the `EXTRAHOP_SECRET_KEY` environment variable in the next procedure.

Configure the Lambda function

An ExtraHop-supplied lambda function routes new flow log events to the ExtraHop flow sensor whenever called by a lambda trigger.

For more information about creating Lambda functions, see the [AWS documentation](#).




Important: The Lambda function must be on the same VPC and subnet as the flow log sensor. The function must also be part of a security group that allows outbound TCP 443 traffic to the management interface of the collector.

1. Download the `exflowlogs-lambda.zip` file from the [ExtraHop downloads](#) page.
2. In AWS, create a Lambda function.
 - The function must have the Go1.x runtime.

- The function must have an execution role with the following permissions:
 - **CloudWatch Logs:**
 - CreateLogGroup
 - CreateLogStream
 - PutLogEvents
 - **EC2:**
 - CreateNetworkInterface
 - DeleteNetworkInterface
 - DescribeNetworkInterfaces
- You must enable connectivity between your Lambda function and the VPC and subnet that your collector is on. The function must also be part of a security group that allows traffic between the function and the collector.
- Upload the `exflowlogs-lambda.zip` file.




- On the **Code** tab, under **Runtime settings**, set the handler value to `exflowlogs-lambda`.
- On the Configuration tab, click **General configuration**.
 - Set the Memory field to `128 MB`.
 - Set the Timeout field to `10 seconds`.
- Click **Function URL** The function URL is required when you configure Kinesis Firehose.
 - Select **NONE** as the auth type.

 **Note:** Setting the auth type to **NONE** will not allow public access to the lambda because the resource-based policy of the function is always in effect and must grant public access before the function URL can receive requests.

- Click **Environment variables** and add the following values:
 - **EDA_HOST:** The IP address or hostname of the VPC flow logs sensor.
 - **EXTRAHOP_SECRET_KEY:** The secret you generated through the ExtraHop REST API in the previous procedure.
 - **VERIFY_EDA_HOST_CERT:** If the sensor has the default self-signed certificate, specify `0` to disable certificate verification in the Lambda HTTP client. Otherwise, specify `1`.

Environment variables

You can define environment variables as key-value pairs that are accessible from your function code. These are useful to store configuration settings without the need to change function code. [Learn more](#) 

Key	Value	
<input type="text" value="EDA_HOST"/>	<input type="text" value="10.11.12.13"/>	<input type="button" value="Remove"/>
<input type="text" value="EXTRAHOP_SECRET_KEY"/>	<input type="text" value="*****"/>	<input type="button" value="Remove"/>
<input type="text" value="VERIFY_EDA_HOST_CERT"/>	<input type="text" value="1"/>	<input type="button" value="Remove"/>
<input type="button" value="Add environment variable"/>		

 **Note:** It can take up to 10 minutes before devices are discovered and metrics are published from flow logs.

Create a Kinesis Firehose stream with an HTTP endpoint

1. Navigate to the Amazon Kinesis dashboard.
2. In the left pane, click **Delivery streams**.
3. Click **Create delivery stream**.
4. Choose the following source and destinations:
 - **Source: Direct PUT**
 - **Destination: HTTP Endpoint**
5. Type a unique name in the delivery stream name field.
6. Specify the following destination settings:
 - **HTTP endpoint name: HTTP endpoint**
 - **HTTP endpoint URL:** The URL of the Lambda function
 - **Access key :** The secret you generated through the ExtraHop REST API in the previous procedure.
 - **Content encoding: GZIP**
7. In the **Backup settings** section, select an existing S3 backup bucket or create a new bucket.
8. Click **Create delivery stream**.

Create the VPC flow log

Identify the VPCs that you want to monitor with the flow sensor.

- If your ExtraHop AWS deployment includes packet sensors, you should avoid monitoring a particular VPC with both a packet sensor and a flow logs sensor.
- While it is possible to send logs for smaller units like individual subnets or interfaces, sending the entire VPC yields the best discovery of devices.

1. Select your VPC.
2. Click the **Flow logs** tab and then click **Create flow log**
3. Configure the following settings:
 - **Filter:** Accept
 - **Maximum aggregation interval:** 1 Minute
 - **Destination:** Send to Kinesis Firehose in the same account or in a different account
 - **Kinesis Firehose delivery stream name:** Select the stream name you created previously
 - **Log record format:** Select **Custom format** and then select the log format attributes in the following order:
 - **end**
 - **log-status**
 - **vpc-id**
 - **interface-id**
 - **srcaddr**
 - **dstaddr**
 - **srcport**
 - **dstport**
 - **protocol**
 - **tcp-flags**
 - **packets**
 - **bytes**
 - **pkt-srcaddr**
 - **pkt-dstaddr**

The format preview should appear similar to the following figure.

Format preview

```

${end} ${log-status} ${vpc-id} ${interface-id} ${srcaddr} ${dstaddr} ${srcport}
${dstport} ${protocol} ${tcp-flags} ${packets} ${bytes} ${pkt-srcaddr} ${pkt-dstaddr}
  
```

Configure Route 53 logs (optional)

Amazon Route 53 provides DNS query logging, which is not required for the flow log configuration but is strongly recommended when the Amazon DNS server is configured.

To configure Route 53 to log DNS queries that originate in your VPCs, see the following AWS documentation: [Managing Resolver query logging configurations](#).

1. Go to the Route 53 service.
2. In the Resolver section, click **Query logging**.
3. Click **Configure query logging**.
 - a. Type a query logging configuration name.
 - b. Select **Kinesis Data Firehose delivery stream** as the query logs destination.
 - c. Select the Kinesis Data Firehose delivery stream that you created previously.
 - d. In the VPCs to log queries for section, click **Add VPC**.
 - e. Select the VPCs that you want to log queries for and then click **Add**.
 - f. Click **Configure query logging**.