

Open Data Streams

Published: 2020-02-21

An open data stream (ODS) is a channel through which you can send specified metric data to an external, third-party system. For example, you might want to store or analyze metric data with a remote tool, such as Splunk, MongoDB, or Amazon Web Services (AWS).

Sending data through an open data stream is a two-step procedure. First, you configure a connection to the target system that will receive the data. Second, you write a trigger that specifies what data to send to the target system and when to send it.

You configure an ODS target through the ExtraHop Admin UI, which requires a user account with full system privileges. You can configure up to 16 ODS connections for each of the following target types:

- HTTP
- Kafka
- MongoDB
- Syslog
- Raw data

The Open Data Streams page in the ExtraHop Admin UI displays configured targets organized by type similar to the following figure:

Open Data Streams

[Add Target](#)

HTTP ▾


Name	Host	Port	Type	Pipelining	Additional Header	Authentication	Status
default	0.0.0.0	80	http	—	—	none	● OK Edit ✕
CloudWatch	monitoring.us-west-2.amazonaws.com	80	http	—	—	aws	● OK Edit ✕

Each target displays the specified configuration settings and enables you to check the connection status, and edit or delete the target.

Status

The Status column displays "OK" if the ExtraHop appliance is able to connect to the target. Otherwise the status is "Error" or "Offline".

Hover over the status to view details about the connection. If there is no correlating ODS trigger or the trigger is not running, the window displays an inactive status. If the connection is active, the window displays metrics such as the number of messages and bytes sent and received and the number of connection attempts, similar to the following figure:



Requests from an ODS trigger to transmit data to the target are asynchronous; there is no support for processing responses to verify the success of a request. The connection details can help you check that the ODS trigger is running and that data is successfully transmitted.

Edit

You can edit any of the target settings except the target name. Because the target name is referenced by the trigger that specifies what data to send through the open data stream, this restriction reduces errors.

Delete

When you delete a target, it is important to also delete or inactivate the correlating ODS trigger; otherwise, the trigger continues to run and consume resources.

After you configure an ODS target, you must create a trigger that specifies what data to manage through the stream. For example, after you create an HTTP target, you might create a trigger that collects the processing time of HTTP responses on a handful of devices, and then sends the metric data to the target.

For more information, see [Open data stream classes](#) in the [ExtraHop Trigger API Reference](#).

Related topics

- [Configure an HTTP target for an open data stream](#)
- [Configure a Kafka target for an open data stream](#)
- [Configure a MongoDB target for an open data stream](#)
- [Configure a raw data target for an open data stream](#)
- [Configure a syslog target for an open data stream](#)
- [Try the ODS walkthrough: Configure an open data stream to send metric data to AWS Cloudwatch](#)