


Configure a Kafka target for an open data stream

Published: 2017-11-17

You can export data on an ExtraHop Discover appliance to any Kafka server for long-term archiving and comparison with other sources.

1. Log into the Admin UI on the ExtraHop Discover appliance.
2. In the System Configuration section, click **Open Data Streams**.
3. Click **Add Target**.
4. From the Target Type drop-down menu, select **Kafka**.
5. In the Name field, type a name to identify the target.
6. From the Compression drop-down list, select one of the following compression methods that will be applied to the transmitted data:
 - **None**
 - **GZIP**
 - **Snappy**
7. From the Partition strategy drop-down list, select one of the following partitioning methods that will be applied to the transmitted data:
 - **Default (Hash Key)**
 - **Manual**
 - **Random**
 - **Round Robin**
8. Specify at least one Kafka broker, also referred to as a node in a Kafka cluster, that can receive transmitted data.



Note: You can add multiple brokers that are part of the same Kafka cluster to ensure connectivity in case a single broker is unavailable. All brokers must be part of the same cluster.

 - a) In the Host field, type the hostname or IP address of the Kafka broker.
 - b) In the Port field, type the port number of the Kafka broker.
 - c) Click the plus (+) icon.
9. Optional: Click **Test** to establish a connection between the Discover appliance and the remote Kafka server and send a test message to the server. The dialog box displays a message that indicates whether the connection succeeded or failed. If the test fails, edit the target configuration and test the connection again.
10. Click **Save**.

Next steps

Create a trigger that specifies what Kafka message data to send and initiates the transmission of data to the target. For more information, see the [Remote.Kafka](#) class in the [ExtraHop Trigger API Reference](#).