

Send records from ExtraHop to Google BigQuery

Published: 2024-07-16

You can configure your ExtraHop system to send transaction-level records to a Google BigQuery server for long-term storage, and then query those records from the ExtraHop system and the ExtraHop REST API. Records on BigQuery recordstores expire after 90 days.

Before you begin

- Any console and all connected sensors must be running the same ExtraHop firmware version.
- You need the BigQuery project ID
- You need the credential file (JSON) from your BigQuery service account. The service account requires the BigQuery Data Editor, BigQuery Data Viewer, and BigQuery User roles.
- For access to the cloud-based recordstore included with RevealX Standard Investigation, your sensors must be able to access outbound TCP 443 (HTTPS) to these fully qualified domain names:
 - `bigquery.googleapis.com`
 - `bigquerystorage.googleapis.com`
 - `oauth2.googleapis.com`
 - `www.googleapis.com`
 - `www.mtls.googleapis.com`
 - `iamcredentials.googleapis.com`

You can also review the public guidance from Google about [computing possible IP address ranges](#) for googleapis.com.

- If you want to configure the BigQuery recordstore settings with Google Cloud workload identity federation authentication, you need the configuration file from your workload identity pool.



Note: The workload identity provider must be set up to provide a fully valid OIDC ID Token in response to a Client Credentials request. For more information about workload identity federation, see <https://cloud.google.com/iam/docs/workload-identity-federation>.

Enable BigQuery as the recordstore

Complete this procedure on all connected ExtraHop sensors and console.



Note: Any triggers configured to send records through `commitRecord` to an ExtraHop recordstore are automatically redirected to BigQuery. No further configuration is required.



Important: If your ExtraHop system includes a console, configure all appliances with the same recordstore settings or transfer management to manage settings from the console.

1. Log in to the Administration settings on the ExtraHop system through `https://<extrahop-hostname-or-IP-address>/admin`.
2. In the Records section, click **Recordstore**.
3. Select **Enable BigQuery as the recordstore**.



Important: If you are migrating to BigQuery from a connected ExtraHop recordstore, you will no longer be able to access records stored on the recordstore.


4. In the Project ID field, type the ID for your BigQuery project.
You can find the project ID in the BigQuery API console.
5. In the JSON Credential File field, click **Choose File** and select one of the following files:
 - The credential file saved from your [BigQuery service account](#).

See the Google Cloud documentation on how to create a service account and generate a service account key.

 **Important:** Create your service account with the following BigQuery roles:

- BigQuery Data Editor
 - BigQuery Data Viewer
 - BigQuery User
- The configuration file from your workload identity pool.
6. Optional: If you chose the configuration file from your workload identity pool in the previous step, select **Authenticate through local Identity Provider for Workload Identity Federation** and type the credentials from your identity provider in the following fields:
 - **Token URL**
 - **Client ID**
 - **Client Secret**
 7. Click **Test Connection** to verify that your sensor can communicate with the BigQuery server.
 8. Click **Save**.

After your configuration is complete, you can query for stored records in the ExtraHop system by clicking **Records**.

 **Important:** Do not modify or delete the table in BigQuery where the records are stored. Deleting the table deletes all stored records.

Transfer recordstore settings

If you have an ExtraHop console connected to your ExtraHop sensors, you can configure and manage the recordstore settings on the sensor, or transfer the management of the settings to the console. Transferring and managing the recordstore settings on the console enables you to keep the recordstore settings up to date across multiple sensors.

Recordstore settings are configured for connected third-party recordstores and do not apply to the ExtraHop recordstore.

1. Log in to the Administration settings on the ExtraHop system through `https://<extrahop-hostname-or-IP-address>/admin`.
2. In the Records section, click **Recordstore**.
3. From the **Recordstore settings** drop-down list, select the console and then click **Transfer Ownership**.
If you later decide to manage the settings on the sensor, select **this sensor** from the Recordstore settings drop-down list and then click **Transfer Ownership**.