

Enable Relational Connections

Contents:

- *Supported Relational Databases*
 - *Ports*
 - *Enable*
 - *Limitations*
 - *Execution at scale*
 - *Encryption Key File*
 - *Configure Relational Query Timeout*
 - *Configure Connections*
 - *Enable SSO Connections*
-

The Trifacta® platform can be configured to access data stored in relational database sources over JDBC protocol. When this connection method is used, individual database tables and views can be imported as datasets.

Supported Relational Databases

The Trifacta platform can connect to these relational database platforms. Supported versions are the following:

- Oracle 12.1.0.2
- SQL Server 12.0.4
- PostgreSQL 9.3.10
- Teradata 14.10+

i NOTE: To enable Teradata connections, you must download and install Teradata drivers first. For more information, see *Enable Teradata Connections*.

Ports

For any relational source to which you are connecting, the Trifacta node must be able to access it through the specified host and port value.

Please contact your database administrator for the host and port information.

Enable

This feature is enabled automatically.

Limitations

i NOTE: For this release, relational connectivity is read-only. Writing to relational databases is not supported.

- You cannot swap relational sources if they are from databases provided by different vendors. See *Flow View Page*.
- There are some differences in behavior between reading tables and views. See *Using Databases*.

Execution at scale

Jobs for large-scale relational sources can be executed on the Hadoop running environment. After the data source has been imported and wrangled, no additional configuration is required to execute at scale.

i **NOTE:** End-to-end performance is likely to be impacted by:

- streaming data volumes over 1 TB from the source,
- streaming from multiple concurrent sources,
- overall network bandwidth.

When the job is completed, any temporary files are automatically removed from HDFS.

For more information, see *Run Job Page*.

Encryption Key File

Before you create relational connections, you must create and reference an encryption key file, see *Create Encryption Key File*.

Configure Relational Query Timeout

If you are experiencing very long query times, particularly for datasets backed by database views, you should consider raising the application timeout settings. These settings apply to the entire application and may have secondary effects. For more information, see *Miscellaneous Configuration*.

Configure Connections

Configure at least one individual connection for any of the supported relational systems. You can configure more than one connection to the same relational system using different credentials. See *Connection Types*.

Enable SSO Connections

If you have enabled Kerberos on the Hadoop cluster, you can leverage the Kerberos global keytab to enable SSO connections to relational sources. See *Enable SSO for Relational Connections*.