

Configure Optimizer Service

Contents:

- *Pre-requisites*
- *Configure Service*
 - *Additional configuration*
- *Use*
 - *Enable for workspace users*
 - *Enable optimizations for individual flows*
- *Logging*

The Optimizer service manages optimizations of job executions that involve queries to relational sources. Based on the enabled optimizations, the service can tune the SQL to optimize queries for better performance and less data transfer.

Pre-requisites

The Optimizer service database must be installed. This database is installed as part of normal install and upgrade operations. For more information, see *Install Databases*.

Configure Service

1. You can apply this change through the *Admin Settings Page* (recommended) or `trifacta-conf.json`. For more information, see *Platform Configuration Methods*.
2. Locate the following configuration:

```
"optimizer-service.enabled":true,  
"optimizer-service.host":"localhost",  
"optimizer-service.port":41913,  
"optimizer-service.autoRestart":true,
```

3. Verify that the `enabled` parameter is set to `true`.
4. Additional configuration settings are described below.

Parameter	Description
<code>enabled</code>	Set this value to <code>true</code> to enable the Optimizer service.
<code>host</code>	Host of the Optimizer service. Leave this value as <code>localhost</code> .
<code>port</code>	Port number that Optimizer service uses to communicate. Default value is 41922. This port must be opened on the Trifacta node.
<code>autoRestart</code>	When set to <code>true</code> , the Optimizer service automatically restarts and attempts to return to its pre-restart state. This value should be set to <code>false</code> for debugging purposes only.

Additional configuration

The following parameters can be modified as needed.

Parameter	Description
-----------	-------------

optimizer-service. maximumOptimizationTimeInS econds	<p>Maximum time in seconds that the service is permitted to optimize a job. If this time is exceeded, the service errors out, and the job proceeds without optimization. Default value is 15.</p> <p>This value may be increased if optimization is desired for any of the following:</p> <ul style="list-style-type: none"> • Very complex recipes • Extremely wide datasets • Jobs reading from several relational sources and pushing down filters to them
optimizer-service. maximumNumberOfWalksPerB atch	<p>Maximum number of passes the service is permitted to make on a recipe to optimize a job. If this number of passes is exceeded, the job is partially optimized. Default value is 200.</p> <p>This value may be increased if the full suite of optimizations is desired for complex recipes with more than 100 steps.</p>

Use

Enable for workspace users

A workspace administrator must enable optimizations for workspace users.

Steps:

1. You apply this change through the *Workspace Settings Page*. For more information, see *Platform Configuration Methods*.
2. Locate the following setting and set it to `Enabled`:

```
Logical and physical optimizations of jobs
```

3. Users can now optimize jobs at the flow level. See below.

Enable optimizations for individual flows

For an individual flow, you can specify the set of optimizations to apply to the job. In Flow View, select **Optimizati on settings** from the flow context menu. For more information, see *Flow Optimization Settings Dialog*.

Logging

You can configure how logging is managed for the Optimizer service. For more information on configuring logging for the Optimizer service, see *Configure Logging for Services*.