

Flow Optimization Settings Dialog

Contents:

- *Enable optimization for jobs from this flow*
 - *General Optimizations*
 - *Column pruning optimization*
 - *Filter optimization*
 - *Databases that Support Pushdown*
 - *Column pruning from source*
 - *Filter pushdown*
 - *Other Databases*
 - *Column pruning from source*
-

In the Flow Optimization Settings dialog, you can configure the following settings, which provide finer-grained control and performance tuning over your flow and its job executions. From the Flow View menu, select **Optimization settings**.

This feature must be enabled at the workspace level. When enabled, the settings in this dialog are applied to the current flow.

- See *Workspace Settings Page*.

These optimizations are designed to improve performance by pre-filtering the volume of data by reducing the columns and rows to the ones that are actually used.

When these filters are enabled, the number of filters successfully applied to a job execution is listed in the Optimization summary in the Job Details page. See *Job Details Page*.

Enable optimization for jobs from this flow

When enabled, the Trifacta application attempts to apply any of the listed optimizations that are enabled to jobs that are executed for this flow.

NOTE: When this option is disabled, then no optimization settings are available.

NOTE: If two consecutive job executions of a flow fail, then optimizations are skipped for the flow. If the job execution then succeeds, optimizations are automatically disabled for the flow. They can be re-enabled if needed.

General Optimizations

The following optimizations can be enabled or disabled in general. For individual data sources, you may be able to enable or disable these settings based on your environment and its requirements .

Tip: These optimizations are applied at the recipe level. They can be applied on any flow and may improve performance within the Transformer page.

Column pruning optimization

When enabled, job execution performance is improved by removing any unused or redundant columns based on the recipe that is selected.

Filter optimization

When this setting is enabled, the Trifacta application optimizes job performance on this flow by pushing data filters to recipes.

Databases that Support Pushdown

Individual types of databases may support one or more of the following pushdowns. Additional restrictions may apply for your specific database.

Tip: These optimizations are applied to queries of your relational datasources that support pushdown. These optimizations are applied within the source, which limits the volume of data that is transferred during job execution.

NOTE: For each relational connection, you can enable the optimization capabilities to improve the flow and its job execution performance. The optimization settings may vary based on the type of relational connections.

Column pruning from source

When enabled, job execution performance is improved by removing any unused or redundant columns from the source database.

Limitations:

- Column pruning optimizations cannot be applied to imported datasets generated with custom SQL.

Filter pushdown

When this setting is enabled, the Trifacta application optimizes job performance on this flow by pushing data filters directly on the source database.

Limitations:

- Filter pushdown optimizations cannot be applied to imported datasets generated with custom SQL.
- Pushdown filters cannot be applied to dates in your relational sources.

NOTE: SQL-based filtering is performed on a best-effort basis. When these optimizations are enabled for your flow, there is no guarantee that they will be applied during job execution.

NOTE: The connection types may or may not be available in your product edition. For more information, see *Connection Types*.

Other Databases

Databases that do not support pushdown may support the following optimization settings.

Column pruning from source

When enabled, job execution performance is improved by removing any unused or redundant columns from the source database.