



Informatica® Data Integration - Free & PayGo
April 2023

Tasks

© Copyright Informatica LLC 2022, 2023

This software and documentation are provided only under a separate license agreement containing restrictions on use and disclosure. No part of this document may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording or otherwise) without prior consent of Informatica LLC.

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation is subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License.

Informatica, Informatica Cloud, Informatica Intelligent Cloud Services, PowerCenter, PowerExchange, and the Informatica logo are trademarks or registered trademarks of Informatica LLC in the United States and many jurisdictions throughout the world. A current list of Informatica trademarks is available on the web at <https://www.informatica.com/trademarks.html>. Other company and product names may be trade names or trademarks of their respective owners.

Portions of this software and/or documentation are subject to copyright held by third parties. Required third party notices are included with the product.

The information in this documentation is subject to change without notice. If you find any problems in this documentation, report them to us at infa_documentation@informatica.com.

Informatica products are warranted according to the terms and conditions of the agreements under which they are provided. INFORMATICA PROVIDES THE INFORMATION IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

Publication Date: 2023-04-04

Table of Contents

Preface	6
Informatica Resources.	6
Informatica Documentation.	6
Informatica Intelligent Cloud Services web site.	6
Informatica Intelligent Cloud Services Communities.	6
Informatica Intelligent Cloud Services Marketplace.	6
Data Integration connector documentation.	7
Informatica Knowledge Base.	7
Informatica Intelligent Cloud Services Trust Center.	7
Informatica Global Customer Support.	7
 Chapter 1: Data integration tasks.....	8
Data filters.	9
Simple data filters.	9
Advanced data filters.	9
Data filter operators.	10
Data filter variables.	11
Rules and guidelines for data filters.	11
Field expressions.	12
Creating a field expression.	12
Transformation language components for expressions.	13
Expression syntax.	13
String and numeric literals.	13
Rules and guidelines for expressions.	14
Adding comments to expressions.	14
Reserved words.	15
Advanced session properties.	16
Parameter files.	22
Schedules.	23
Repeat frequency.	23
Time zones and schedules.	24
Daylight Savings Time changes and schedules.	25
Creating a schedule.	25
Running a task on a schedule.	26
Email notification.	26
Preprocessing and postprocessing commands.	27
Preprocessing and postprocessing SQL commands.	27
Preprocessing and postprocessing operating system commands.	27
Monitoring a job.	28
Stopping a job.	28

Guidelines for sources and targets in data integration tasks.	28
Rules and guidelines for flat file sources.	29
Rules and guidelines for database sources and targets.	29
Chapter 2: Mapping tasks.	30
Mapping task templates.	30
Related objects.	31
Advanced relationships.	32
Pushdown optimization.	32
Simultaneous task runs.	33
Field metadata.	33
Schema change handling.	33
Dynamic schema handling options.	34
Dynamic schema change handling rules and guidelines.	35
Schema mismatch handling.	36
Mapping task configuration.	36
Defining a mapping task.	36
Configuring sources.	37
Configuring targets.	39
Configuring parameters.	41
Configuring a schedule and advanced options.	41
Viewing and editing mapping task details.	45
Sequence Generator values.	45
Running a mapping task.	46
Reprocessing incrementally-loaded source files	46
Chapter 3: Data transfer tasks.	47
Task operations.	47
Data transfer task sources.	48
Source filters.	48
Sort conditions.	48
Second sources.	48
Lookup condition.	49
Second source filters.	50
Data transfer task targets.	50
Database target truncation.	50
Update columns.	50
Field mapping.	51
Field data types.	52
Configuring a data transfer task.	52
Defining the data transfer task.	52
Configuring the source.	52
Configuring a second source.	53

Configuring the target.	54
Configuring the field mapping.	55
Configuring runtime options.	55
Running a data transfer task.	55
Chapter 4: Data loader tasks.	56
Step 1. Connect to your source.	57
Guidelines for customizing your source.	58
Selecting source objects.	59
Defining the source format.	59
Selecting source fields.	60
Configuring filters.	60
Selecting primary key fields.	61
Selecting watermark fields.	62
Step 2. Connect to your target.	63
Incremental loading.	63
Dropping and re-creating target tables.	64
Step 3. Run the task.	64
Running the task on a schedule.	65
Configuring email notifications.	65
Handling runtime errors.	65
Configuring the task location.	65
Runtime environment.	66
Index.	67

Preface

Use *Tasks* to learn how to set up and run Data Integration tasks manually or on a schedule.

Informatica Resources

Informatica provides you with a range of product resources through the Informatica Network and other online portals. Use the resources to get the most from your Informatica products and solutions and to learn from other Informatica users and subject matter experts.

Informatica Documentation

Use the Informatica Documentation Portal to explore an extensive library of documentation for current and recent product releases. To explore the Documentation Portal, visit <https://docs.informatica.com>.

If you have questions, comments, or ideas about the product documentation, contact the Informatica Documentation team at infa_documentation@informatica.com.

Informatica Intelligent Cloud Services web site

You can access the Informatica Intelligent Cloud Services web site at <http://www.informatica.com/cloud>. This site contains information about Informatica Cloud integration services.

Informatica Intelligent Cloud Services Communities

Use the Informatica Intelligent Cloud Services Community to discuss and resolve technical issues. You can also find technical tips, documentation updates, and answers to frequently asked questions.

Access the Informatica Intelligent Cloud Services Community at:

<https://network.informatica.com/community/informatica-network/products/cloud-integration>

Developers can learn more and share tips at the Cloud Developer community:

<https://network.informatica.com/community/informatica-network/products/cloud-integration/cloud-developers>

Informatica Intelligent Cloud Services Marketplace

Visit the Informatica Marketplace to try and buy Data Integration Connectors, templates, and mapplets:

<https://marketplace.informatica.com/>

Data Integration connector documentation

You can access documentation for Data Integration Connectors at the Documentation Portal. To explore the Documentation Portal, visit <https://docs.informatica.com>.

Informatica Knowledge Base

Use the Informatica Knowledge Base to find product resources such as how-to articles, best practices, video tutorials, and answers to frequently asked questions.

To search the Knowledge Base, visit <https://search.informatica.com>. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Intelligent Cloud Services Trust Center

The Informatica Intelligent Cloud Services Trust Center provides information about Informatica security policies and real-time system availability.

You can access the trust center at <https://www.informatica.com/trust-center.html>.

Subscribe to the Informatica Intelligent Cloud Services Trust Center to receive upgrade, maintenance, and incident notifications. The [Informatica Intelligent Cloud Services Status](#) page displays the production status of all the Informatica cloud products. All maintenance updates are posted to this page, and during an outage, it will have the most current information. To ensure you are notified of updates and outages, you can subscribe to receive updates for a single component or all Informatica Intelligent Cloud Services components. Subscribing to all components is the best way to be certain you never miss an update.

To subscribe, go to <https://status.informatica.com/> and click **SUBSCRIBE TO UPDATES**. You can then choose to receive notifications sent as emails, SMS text messages, webhooks, RSS feeds, or any combination of the four.

Informatica Global Customer Support

You can contact a Customer Support Center by telephone or online.

For online support, click **Submit Support Request** in Informatica Intelligent Cloud Services. You can also use Online Support to log a case. Online Support requires a login. You can request a login at <https://network.informatica.com/welcome>.

The telephone numbers for Informatica Global Customer Support are available from the Informatica web site at <https://www.informatica.com/services-and-training/support-services/contact-us.html>.

CHAPTER 1

Data integration tasks

A data integration task is a process that you configure to analyze, extract, transform, and load data. You can run individual tasks manually or set tasks to run on a schedule.

You can use the following tasks to integrate data:

Mapping tasks

Mapping tasks process data based on the data flow logic defined in a mapping.

A mapping reads data from one or more sources, transforms the data based on logic that you define, and writes it to one or more targets. Create a mapping when you need to augment or manipulate your data before you load it to a target. For example, if you need to aggregate data, calculate values, perform complex joins, normalize data, or route data to different targets, you can create a mapping to do this.

A mapping task runs the data flow logic that you've defined in the mapping. Choose this task type after you've created a mapping so that you can run the data flow logic defined in the mapping.

Data transfer tasks

Data transfer tasks move data from one or two sources to a target. You can also choose to sort and filter the data before you load it to the target.

Choose this task type when you want to transfer data from a source object, optionally add fields from a second source object, and write the data to a new or existing target object without changing the source data. For example, if you want to move customer records from an on-premises database table to a table in your cloud data warehouse, create a data transfer task.

Data loader tasks

Data loader tasks provide secure data loading from multi-object sources to corresponding objects in your cloud data warehouse. They can load data incrementally and provide support for schema drift.

To optimize performance, data loading occurs in parallel batches. To fine-tune the data, you can exclude certain objects and fields and also apply some simple filters. If your source data changes frequently, you can load only new and changed records each time the task runs.

Choose this task type when you need to ingest data as-is from multiple objects into your cloud data warehouse. For example, if you need to repeatedly load all the data from files in an Amazon S3 bucket to corresponding tables in Snowflake Data Cloud, create a data loader task.

When you create a task, Data Integration walks you through the required steps. The options and properties that display depend on the task type.

You can create a workflow of multiple tasks by linking the tasks in taskflows. For more information, see *Taskflows*.

Data filters

You can create the following type of data filters in a data transfer task:

- Simple
- Advanced

You can create a set of data filters for each object included in the task. Each set of data filters acts independently of the other sets.

Simple data filters

You can create one or more simple data filters.

When you create multiple simple data filters, the associated task creates an AND operator between the filters and loads rows that apply to all simple data filters.

For example, you load rows from the Account Salesforce object to a database table. However, you want to load only accounts that have greater than or equal to \$100,000 in annual revenue and that have more than 500 employees. You configure the following simple data filters:

Field	Operator	Field Value
AnnualRevenue	>=	100000
NumberOfEmployees	>	500

Configuring simple data filters

You configure simple data filters in the task wizard.

1. To create a simple data filter, perform one of the following actions:
 - On the **Data Filters** page, click **Simple**, and then click **New**. The **Data Filter** dialog box appears.
 - In a data transfer task, in the **Filters** area on the **Source** or **Second Source** page, select **Simple** and then click **New**.
2. Specify the object on which to create the data filter.
You create separate data filters for each source object included in the task.
3. Enter the filter condition based on the field, operator, and field value.
4. Click **OK**.
5. Create additional simple data filters as needed.
To delete a data filter, click the **Delete** icon next to the data filter.
6. Click **Next**.

Advanced data filters

Create an advanced data filter to create complex expressions that use AND, OR, or nested conditions.

When you create an advanced data filter, you enter one expression that contains all filters. The expression that you enter becomes the WHERE clause in the query used to retrieve records from the source.

For example, you load rows from the Account Salesforce object to a database table. However, you want to load records where the billing state is California or New York and the annual revenue is greater than or equal to \$100,000. You configure the following advanced filter expression:

```
(BillingState = 'CA' OR BillingState = 'NY') AND (AnnualRevenue >= 100000)
```

When you create a data filter on a Salesforce object, the corresponding task generates a SOQL query with a WHERE clause. The WHERE clause represents the data filter. The SOQL query must be less than 20,000 characters. If the query exceeds the character limit, the following error appears:

```
Salesforce SOQL limit of 5000 characters has been exceeded for the object: <Salesforce object>. Please exclude more fields or decrease the filters.
```

Note: Filter conditions are not validated until runtime.

Configuring advanced data filters

Configure advanced data filters in the task wizard. You can convert simple data filters to an advanced data filter, but you cannot convert an advanced data filter to simple data filters.

1. In a data transfer task, in the **Filters** area on the **Source** or **Second Source** page, select **Advanced**.
To convert all simple data filters to one advanced data filter, select **Advanced**.
2. If necessary, specify the object on which to create the data filter.
You create separate data filters for each source object included in the task.
3. Enter the filter expression.
Click the field name to add the field to the expression.
4. Click **OK**.
To delete a data filter, click the **Delete** icon next to the data filter.
5. Click **Next**.

Data filter operators

You can use specific operators with each field type.

The following table shows the operators you can use for each field type:

Field type	Operators
Boolean	=, !=, Is Null, Is Not Null
Currency	=, !=, <, <=, >, >=, Is Null, Is Not Null
Date	=, !=, <, <=, >, >=, Is Null, Is Not Null
Datetime	=, !=, <, <=, >, >=, Is Null, Is Not Null
Double	=, !=, <, <=, >, >=, Is Null, Is Not Null
ID	=, !=, Is Null, Is Not Null
Int	=, !=, <, <=, >, >=, Is Null, Is Not Null
Reference	=, !=, Is Null, Is Not Null

Field type	Operators
String	=, !=, LIKE'_%', LIKE'%_', LIKE'%_%', Is Null, Is Not Null, <, <=, >, >=
Textarea	=, !=, LIKE'_%', LIKE'%_', LIKE'%_%', Is Null, Is Not Null, <, <=, >, >=
All other field types	=, !=, Is Null, Is Not Null

Data filter variables

Data filter variables represent the date or time that a task previously ran. Use data filter variables to help capture the source data that changed since the last task run. You can use data filter variables in simple and advanced data filter conditions.

You can use the following data filter variables:

Variable	Description
\$LastRunDate	The start date in GMT time zone of the last task run that was successful or ended with a warning. Does not include time. For example, 2018-09-24. Can be used as a value for filter where the field type is DATE.
\$LastRunTime	The start date and time in GMT time zone of the last task run that was successful or ended with a warning. For example, 2018-09-24 15:23:23. Can be used as a value for filter where the field type is DATETIME.

For example, you can include the following simple filter condition:

```
LastModifiedDate > $LastRunTime
```

Note: Consider time zone differences when comparing dates across time zones. The date and time of the \$LastRunDate and \$LastRunTime variables are based on the time zone set in Informatica Intelligent Cloud Services. The date and time of the actual job is based on the GMT time zone for Salesforce sources and the database server for database sources. The difference in the time zones may yield unexpected results.

Rules and guidelines for data filters

Use the following rules and guidelines for data filters:

- Data filters must contain valid SQL or SOQL operators.
- You cannot include simple and advanced data filters in the same task.
- When you convert a simple data filter to an advanced data filter, you cannot convert the advanced data filter back to a simple data filter.
- A task fails if the fields included in the data filter no longer exist or if the data types of the fields change. If a data type changes, edit the task.
- You can select Equals, Not Equals, Is Null, or Is Not Null operators on fields of the Other data type.
- Applications do not apply filters with Equals, Starts With, or Ends With operators and string fields that contain data that starts or ends with a single quotation mark. To filter these records, use the Contains operator.
- You can only use IS NULL and LIKE operators in data filters for fields of the Text, Ntext, and Image data types.

- If you specify a date and no time for a date/time filter, Data Integration uses 00:00:00 (12:00:00 a.m.) as the time.
- The list of available operators in a simple data filter depends on the data type of the field included in the data filter. Some operators do not apply to all fields included in data filters.
- When you enter more than one simple data filter, applications filter rows that meet the requirements of all data filters.
- When you use a parameter in a data filter, start the data filter with the parameter. For example, use `$Sales=100000` instead of `100000=$Sales`.

Field expressions

When you configure a data transfer task, you can configure the field mapping. The field mapping defines how source fields are mapped to target fields. You can specify an expression for each field mapping.

You can map multiple source fields to the same target field. For example, you can map `SourceFieldA` and `SourceFieldB` to `TargetFieldC`.

Data Integration might suggest operations when you map multiple source fields to a single target field. For example, if you map multiple text fields to a target text field, Data Integration concatenates the source text fields by default. You can change the default expression.

Data Integration provides a transformation language that includes SQL-like functions to transform source data. Use these functions to write expressions, which modify data or test whether data matches the conditions that you specify.

For more information about functions and the Data Integration transformation language, see *Function Reference*.

Creating a field expression

Create a field expression in a task wizard.

1. In the **Field Mappings** page, select the target field for which you want to add an expression.
2. Click **Add or Edit Expression**.

By default, the **Field Expression** dialog box shows the source field as the expression, which indicates that the target contains the same value as the source.

3. Enter the new field expression.

To include source fields and system variables in the expression, you can select them from the **Source Fields** and **System Variables** tabs to insert them into the expression or you can add them to the expression manually.

4. Click **Validate Mapping** to validate the field mappings.
5. Click **Save**.

Validating expressions in field mappings

Use the following rules and guidelines when you validate an expression in a field mapping:

- When you validate mappings, Data Integration performs the following validations:
 - Verifies that the source and target fields in the task exist in the source or target. If the field does not exist, an error appears.
 - Verifies that the correct parameters are used for each function and that the function is valid.
- The expression validator does not perform case-sensitive checks on field names.
- The expression validator verifies that the data type of a field in an expression matches the data type expected by the containing function. However, the expression validator does not check for incompatible data types between the following sets of objects:
 - Source and target fields of tasks.
 - Source field in a lookup condition and the lookup field.
 - Output of an expression or lookup and the target field.The expression or lookup with these incompatible data types may validate successfully, but, at runtime, the task fails and an error appears.
- If you map a string source field to a number target field, the validation succeeds. Data Integration tries to convert the string to a number using the `atoi` (ASCII to Integer) C function.
- The expression validator does not validate lookups.

Transformation language components for expressions

The transformation language includes the following components to create simple or complex expressions:

- Fields. Use the name of a source field to refer to the value of the field.
- Literals. Use numeric or string literals to refer to specific values.
- Functions. Use these SQL-like functions to change data in a task.
- Operators. Use transformation operators to create expressions to perform mathematical computations, combine data, or compare data.
- Constants. Use the predefined constants to reference values that remain constant, such as `TRUE`.

Expression syntax

You can create a simple expression that only contains a field, such as `ORDERS`, or a numeric literal, such as `10`. You can also write complex expressions that include functions nested within functions, or combine different fields using the transformation language operators.

Note: Although the transformation language is based on standard SQL, there are differences between the two languages.

String and numeric literals

You can include numeric or string literals.

Enclose string literals within single quotation marks. For example:

```
'Alice Davis'
```

String literals are case sensitive and can contain any character except a single quotation mark. For example, the following string is not allowed:

```
'Joan's car'
```

To return a string containing a single quotation mark, use the CHR function:

```
'Joan' || CHR(39) || 's car'
```

Do not use single quotation marks with numeric literals. Just enter the number you want to include. For example:

```
.05
```

or

```
$$Sales_Tax
```

Rules and guidelines for expressions

Use the following rules and guidelines when you write expressions:

- For each source field, you can perform a lookup or create an expression. You cannot do both.
- You cannot use strings in numeric expressions.
For example, the expression `1 + '1'` is not valid because you can only perform addition on numeric data types. You cannot add an integer and a string.
- You cannot use strings as numeric parameters.
For example, the expression `SUBSTR(TEXT_VAL, '1', 10)` is not valid because the SUBSTR function requires an integer value, not a string, as the start position.
- You cannot mix data types when using comparison operators.
For example, the expression `123.4 = '123.4'` is not valid because it compares a decimal value with a string.
- You can pass a value from a field, literal string or number, or the results of another expression.
- Separate each argument in a function with a comma.
- Except for literals, the transformation language is not case sensitive.
- The colon (:), comma (,), and period (.) have special meaning and should be used only to specify syntax.
- Data integration tasks treat a dash (-) as a minus operator.
- If you pass a literal value to a function, enclose literal strings within single quotation marks. Do not use quotation marks for literal numbers. Data integration tasks treat any string value enclosed in single quotation marks as a character string.
- Do not use quotation marks to designate fields.
- You can nest multiple functions within an expression. Data integration tasks evaluate the expression starting with the innermost function.
- When you use a parameter in an expression, use the appropriate function to convert the value to the necessary data type. For example, you might use the following expression to define a quarterly bonus for employees:

```
IIF((EMP_SALES < TO_INTEGER($$SalesQuota), 200, 0)
```

Adding comments to expressions

You can use the following comment specifiers to insert comments in expressions:

- Two dashes:

```
-- These are comments
```

- Two forward slashes:

```
// These are comments
```

Data integration tasks ignore all text on a line preceded by comment specifiers. For example, to concatenate two strings, enter the following expression with comments in the middle of the expression:

```
-- This expression concatenates first and last names for customers:
FIRST_NAME -- First names from the CUST table
|| // Concat symbol
LAST_NAME // Last names from the CUST table
// Joe Smith Aug 18 1998
```

Data integration tasks ignore the comments and evaluates the expression as follows:

```
FIRST_NAME || LAST_NAME
```

You cannot continue a comment to a new line:

```
-- This expression concatenates first and last names for customers:
FIRST_NAME -- First names from the CUST table
|| // Concat symbol
LAST_NAME // Last names from the CUST table
Joe Smith Aug 18 1998
```

In this case, data integration tasks do not validate the expression because the last line is not a valid expression.

Reserved words

Some keywords, such as constants, operators, and system variables, are reserved for specific functions. These include:

- :EXT
- :INFA
- :LKP
- :MCR
- :SD
- :SEQ
- :SP
- :TD
- AND
- DD_DELETE
- DD_INSERT
- DD_REJECT
- DD_UPDATE
- FALSE
- NOT
- NULL
- OR
- PROC_RESULT
- SPOUTPUT
- TRUE
- WORKFLOWSTARTTIME

The following words are reserved for Informatica Intelligent Cloud Services:

- ABORTED
- DISABLED
- FAILED
- NOTSTARTED
- STARTED
- STOPPED
- SUCCEEDED

Note: You cannot use a reserved word to name a field. Reserved words have predefined meanings in expressions.

Advanced session properties

Advanced session properties are optional properties that you can configure in mapping tasks. Use caution when you configure advanced session properties. The properties might not be appropriate for use with all tasks.

You can configure the following types of advanced session properties:

- General
- Performance
- Advanced
- Error handling

General options

The following table describes the general options:

General options	Description
Write Backward Compatible Session Log File	Writes the session log to a file.
Session Log File Name	<p>Name for the session log. Use any valid file name.</p> <p>You can customize the session log file name in one of the following ways:</p> <ul style="list-style-type: none">- Using a static name. A static log file name is a simple static string with or without a file extension. <p>If you use a static name, the log file name is appended with a sequence number each time the task runs, for example samplelog.1, samplelog.2. When the maximum number of log files is reached, the numbering sequence begins a new cycle.</p> <ul style="list-style-type: none">- Using a dynamic name. A log file name is dynamic when it includes a parameter defined in a parameter file or a system variable. You can include any of the following system variables: <ul style="list-style-type: none">- \$CurrentTaskName. Replaced with the task name.- \$CurrentTime. Replaced with the current time.- \$CurrentRunId. Replaced with the run ID for the current job. <p>If you use a dynamic name, the file name is unique for every task run. The Maximum Number of Log Files property is not applied. To purge old log files, delete the files manually.</p>
Session Log File Directory	<p>Directory where the session log is saved. Use a directory local to the Secure Agent to run the task.</p> <p>By default, the session log is saved to the following directory:</p> <p><Secure Agent installation directory>/apps/Data_Integration_Server/logs</p>
\$Source Connection Value	Source connection name.
\$Target Connection Value	Target connection name.
Source File Directory	Source file directory path. Use for flat file connections only.
Target File Directory	This option is not used.
Treat Source Rows as	<p>When the task reads source data, it marks each row with an indicator that specifies the target operation to perform when the row reaches the target. Use one of the following options:</p> <ul style="list-style-type: none">- Insert. All rows are marked for insert into the target.- Update. All rows are marked for update in the target.- Delete. All rows are marked for delete from the target.- Data Driven. The task uses the Update Strategy object in the data flow to mark the operation for each source row.
Commit Type	<p>Commit type to use. Use one of the following options.</p> <ul style="list-style-type: none">- Source. The task performs commits based on the number of source rows.- Target. The task performs commits based on the number of target rows. <p>When you do not configure a commit type, the task performs a target commit.</p>

General options	Description
Commit Interval	Interval in rows between commits. When you do not configure a commit interval, the task commits every 10,000 rows.
Commit on End of File	Commits data at the end of the file.
Rollback Transactions on Errors	Rolls back the transaction at the next commit point when the task encounters a non-fatal error. When the task encounters a transformation error, it rolls back the transaction if the error occurs after the effective transaction generator for the target.
Java Classpath	This option is not used.

Performance settings

The following table describes the performance settings:

Performance settings	Description
DTM Buffer Size	Amount of memory allocated to the task from the DTM process. By default, a minimum of 12 MB is allocated to the buffer at run time. Use one of the following options: <ul style="list-style-type: none"> - Auto. Enter Auto to use automatic memory settings. When you use Auto, configure Maximum Memory Allowed for Auto Memory Attributes. - A numeric value. Enter the numeric value that you want to use. The default unit of measure is bytes. Append KB, MB, or GB to the value to specify a different unit of measure. For example, 512MB. You might increase the DTM buffer size in the following circumstances: <ul style="list-style-type: none"> - When a task contains large amounts of character data, increase the DTM buffer size to 24 MB. - When a source contains a large binary object with a precision larger than the allocated DTM buffer size, increase the DTM buffer size so that the task does not fail.
Incremental Aggregation	Performs incremental aggregation for tasks.
Reinitialize Aggregate Cache	This option is not used.
Enable High Precision	Processes the Decimal data type to a precision of 28.
Session Retry on Deadlock	The task retries a write on the target when a deadlock occurs.

Performance settings	Description
Pushdown Optimization	<p>Type of pushdown optimization. Use one of the following options:</p> <ul style="list-style-type: none"> - None. The task processes all transformation logic for the task. - To Source. The task pushes as much of the transformation logic to the source database as possible. - To Target. The task pushes as much of the transformation logic to the target database as possible. - Full. The task pushes as much of the transformation logic to the source and target databases as possible. The task processes any transformation logic that it cannot push to a database. - \$\$PushdownConfig. The task uses the pushdown optimization type specified in the user-defined parameter file for the task. <p>When you use \$\$PushdownConfig, ensure that the user-defined parameter is configured in the parameter file.</p> <p>When you use pushdown optimization, do not use the Error Log Type property.</p> <p>For more information, see the help for the appropriate connector.</p> <p>The pushdown optimization functionality varies depending on the support available for the connector. For more information, see the help for the appropriate connector.</p>
Create Temporary View	<p>Allows the task to create temporary view objects in the database when it pushes the task to the database.</p> <p>Use when the task includes an SQL override in the Source Qualifier transformation or Lookup transformation.</p>
Create Temporary Sequence	This option is not used.
Enable cross-schema pushdown optimization	<p>Enables pushdown optimization for tasks that use source or target objects associated with different schemas within the same database.</p> <p>To see if cross-schema pushdown optimization is applicable to the connector you use, see the help for the relevant connector.</p> <p>This property is enabled by default.</p>
Allow Pushdown for User Incompatible Connections	<p>Indicates that the database user of the active database has read permission on idle databases.</p> <p>If you indicate that the database user of the active database has read permission on idle databases, and it does not, the task fails.</p> <p>If you do not indicate that the database user of the active database has read permission on idle databases, the task does not push transformation logic to the idle databases.</p>
Session Sort Order	Order to use to sort character data for the task.

Advanced options

The following table describes the advanced options:

Advanced options	Description
Constraint Based Load Ordering	This option is not used.
Cache Lookup() Function	This option is not used.

Advanced options	Description
Default Buffer Block Size	<p>Size of buffer blocks used to move data and index caches from sources to targets. By default, the task determines this value at run time.</p> <p>Use one of the following options:</p> <ul style="list-style-type: none"> - Auto. Enter Auto to use automatic memory settings. When you use Auto, configure Maximum Memory Allowed for Auto Memory Attributes. - A numeric value. Enter the numeric value that you want to use. The default unit of measure is bytes. Append KB, MB, or GB to the value to specify a different unit of measure. For example, 512MB. <p>The task must have enough buffer blocks to initialize. The minimum number of buffer blocks must be greater than the total number of Source Qualifiers, Normalizers for COBOL sources, and targets.</p> <p>The number of buffer blocks in a task = DTM Buffer Size / Buffer Block Size. Default settings create enough buffer blocks for 83 sources and targets. If the task contains more than 83, you might need to increase DTM Buffer Size or decrease Default Buffer Block Size.</p>
Line Sequential Buffer Length	<p>Number of bytes that the task reads for each line. Increase this setting from the default of 1024 bytes if source flat file records are larger than 1024 bytes.</p>
Maximum Memory Allowed for Auto Memory Attributes	<p>Maximum memory allocated for automatic cache when you configure the task to determine the cache size at run time.</p> <p>You enable automatic memory settings by configuring a value for this attribute. Enter a numeric value. The default unit is bytes. Append KB, MB, or GB to the value to specify a different unit of measure. For example, 512MB.</p> <p>If the value is set to zero, the task uses default values for memory attributes that you set to auto.</p>
Maximum Percentage of Total Memory Allowed for Auto Memory Attributes	<p>Maximum percentage of memory allocated for automatic cache when you configure the task to determine the cache size at run time. If the value is set to zero, the task uses default values for memory attributes that you set to auto.</p>
Additional Concurrent Pipelines for Lookup Cache Creation	<p>Restricts the number of pipelines that the task can create concurrently to pre-build lookup caches. You can configure this property when the Pre-build Lookup Cache property is enabled for a task or transformation.</p> <p>When the Pre-build Lookup Cache property is enabled, the task creates a lookup cache before the Lookup receives the data. If the task has multiple Lookups, the task creates an additional pipeline for each lookup cache that it builds.</p> <p>To configure the number of pipelines that the task can create concurrently, select one of the following options:</p> <ul style="list-style-type: none"> - Auto. The task determines the number of pipelines it can create at run time. - Numeric value. The task can create the specified number of pipelines to create lookup caches.
Custom Properties	<p>Configure custom properties for the task. You can override the custom properties that the task uses after the job has started. The task also writes the override value of the property to the session log.</p>

Advanced options	Description
Pre-build Lookup Cache	<p>Allows the task to build the lookup cache before the Lookup receives the data. The task can build multiple lookup cache files at the same time to improve performance.</p> <p>Configure one of the following options:</p> <ul style="list-style-type: none"> - Always allowed. The task can build the lookup cache before the Lookup receives the first source row. The task creates an additional pipeline to build the cache. - Always disallowed. The task cannot build the lookup cache before the Lookup receives the first row. <p>When you use this option, configure the Configure the Additional Concurrent Pipelines for Lookup Cache Creation property. The task can pre-build the lookup cache if this property is greater than zero.</p>
DateTime Format String	<p>Date time format for the task. You can specify seconds, milliseconds, or nanoseconds.</p> <p>To specify seconds, enter <code>MM/DD/YYYY HH24:MI:SS</code>.</p> <p>To specify milliseconds, enter <code>MM/DD/YYYY HH24:MI:SS.MS</code>.</p> <p>To specify microseconds, enter <code>MM/DD/YYYY HH24:MI:SS.US</code>.</p> <p>To specify nanoseconds, enter <code>MM/DD/YYYY HH24:MI:SS.NS</code>.</p> <p>By default, the format specifies microseconds, as follows: <code>MM/DD/YYYY HH24:MI:SS.US</code>.</p>
Pre 85 Timestamp Compatibility	This option is not used.

Error handling

The following table describes the error handling options:

Error handling options	Description
Stop on Errors	<p>Indicates how many non-fatal errors the task can encounter before it stops the session. Non-fatal errors include reader, writer, and DTM errors.</p> <p>Enter the number of non-fatal errors you want to allow before stopping the session. The task maintains an independent error count for each source, target, and transformation. If you specify 0, non-fatal errors do not cause the session to stop.</p>
Override Tracing	Overrides tracing levels set on an object level.
On Stored Procedure Error	This option is not used.
On Pre-Session Command Task Error	<p>Determines the behavior when a task that includes pre-session shell commands encounters errors. Use one of the following options:</p> <ul style="list-style-type: none"> - Stop Session. The task stops when errors occur while executing pre-session shell commands. - Continue Session. The task continues regardless of errors. <p>By default, the task stops.</p>
On Pre-Post SQL Error	<p>Determines the behavior when a task that includes pre-session or post-session SQL encounters errors:</p> <ul style="list-style-type: none"> - Stop Session. The task stops when errors occur while executing pre-session or post-session SQL. - Continue. The task continues regardless of errors. <p>By default, the task stops.</p>

Error handling options	Description
Error Log Type	Specifies the type of error log to create. You can specify flat file or no log. Default is none. You cannot log row errors from XML file sources. You can view the XML source errors in the session log. Do not use this property when you use the Pushdown Optimization property.
Error Log File Directory	Specifies the directory where errors are logged. By default, the error log file directory is \$PMBadFilesDir\.
Error Log File Name	Specifies error log file name. By default, the error log file name is PMError.log.
Log Row Data	Specifies whether or not to log transformation row data. When you enable error logging, the task logs transformation row data by default. If you disable this property, n/a or -1 appears in transformation row data fields.
Log Source Row Data	Specifies whether or not to log source row data. By default, the check box is clear and source row data is not logged.
Data Column Delimiter	Delimiter for string type source row data and transformation group row data. By default, the task uses a pipe () delimiter. Tip: Verify that you do not use the same delimiter for the row data as the error logging columns. If you use the same delimiter, you may find it difficult to read the error log file.

Parameter files

A parameter file is a list of user-defined parameters and their associated values.

Use a parameter file to define values that you want to update without having to edit the task. You update the values in the parameter file instead of updating values in a task. The parameter values are applied when the task runs.

You can use a parameter file to define parameter values in mapping tasks.

Define parameter values for connections in the following transformations:

- Source
- Target
- Lookup
- SQL

Define parameter values for objects in the following transformations:

- Source
- Target
- Lookup

Also, define values for parameters in data filters, expressions, and lookup expressions.

Note: Not all connectors support parameter files. To see if a connector supports runtime override of connections and data objects, see the help for the appropriate connector.

You enter the parameter file name and location when you configure the task.

Schedules

You can run tasks manually or you can use schedules to run them at a specific time or interval such as hourly, daily, or weekly.

To use a schedule, you associate the task with a schedule when you configure the task. You can use an existing schedule or create a new schedule. If you want to create a schedule, you can create the schedule from the task's **Schedule** page during task configuration.

When you create a schedule, you specify the date and time. You can configure a schedule to run associated assets throughout the day between 12:00 a.m. and 11:55 p.m. Informatica Intelligent Cloud Services might add a small schedule offset to the start time, end time, and all other time configurations. As a result, scheduled tasks and taskflows might start later than expected. For example, you configure a schedule to run hourly until noon, and the schedule offset for your organization is 10 seconds. Informatica Intelligent Cloud Services extends the end time for the schedule to 12:00:10 p.m., and the last hourly task or taskflow starts at 12:00:10 p.m. To see the schedule offset for your organization, check the **Schedule Offset** organization property.

You can monitor scheduled tasks from the **All Jobs** page in Monitor. Scheduled tasks do not appear on the **My Jobs** page.

When you copy a task that includes a schedule, the schedule is not associated with the new task. To associate a schedule with the new task, edit the task.

If you remove a task from a schedule as the task runs, the job completes. Data Integration cancels any additional runs associated with the schedule.

Repeat frequency

The repeat frequency determines how often tasks run. You can set the repeat frequency to every N minutes, hourly, daily, weekly, biweekly, or monthly.

The following table describes the repeat frequency options:

Option	Description
Does not repeat	Tasks run as scheduled and do not repeat.
Every N minutes	Tasks run on an interval based on a specified number of minutes. You can configure the following options: <ul style="list-style-type: none">- Repeat frequency. Select a frequency in minutes. Options are 5, 10, 15, 20, 30, 45.- Days. Days of the week when you want tasks to run. You can select one or more days of the week.- Time range. Hours of the day when you want tasks to start. Select All Day or configure a time range. You can configure a time range between 00:00-23:55.- Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time.

Option	Description
Hourly	<p>Tasks run on an hourly interval based on the start time of the schedule.</p> <p>You can configure the following options:</p> <ul style="list-style-type: none"> - Repeat frequency. Select a frequency in hours. Options are 1, 2, 3, 4, 6, 8, 12. - Days. Days of the week when you want tasks to run. You can select one or more days of the week. - Time range. Hours of the day when you want tasks to start. Select All Day or configure a time range. You can configure a time range between 00:00-23:55. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time.
Daily	<p>Tasks run daily at the start time configured for the schedule.</p> <p>You can configure the following options:</p> <ul style="list-style-type: none"> - Repeat frequency. The frequency at which you want tasks to run. Select Every Day or Every Weekday. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time.
Weekly	<p>Tasks run on a weekly interval based on the start time of the schedule.</p> <p>You can configure the following options:</p> <ul style="list-style-type: none"> - Days. Days of the week when you want tasks to run. You can select one or more days of the week. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time. <p>If you do not specify a day, the schedule runs regularly on the same day of the week as the start date.</p>
Biweekly	<p>Tasks run every two weeks based on the start time of the schedule.</p> <p>You can configure the following options:</p> <ul style="list-style-type: none"> - Days. Days of the week when you want tasks to run. You can select one or more days of the week. You must select at least one day. - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time. <p>If you configure a biweekly schedule to start at 5 p.m. on a Tuesday and run tasks every two weeks on Mondays, the schedule begins running tasks on the following Monday.</p>
Monthly	<p>Tasks run on a monthly interval based on the start time of the schedule.</p> <p>You can configure the following options:</p> <ul style="list-style-type: none"> - Day. Day of the month when you want tasks to run. You can configure one of the following options: <ul style="list-style-type: none"> - Select the exact date of the month, between 1-28. If you want the task to run on days later in the month, use the <n> <day of the week> option. - Select the <n> <day of the week>. Options for <n> include First, Second, Third, Fourth, and Last. Options for <day of the week> includes Day, and Sunday-Saturday. <p>Tip: With the Day option, you can configure tasks to run on the First Day or the Last Day of the month.</p> - Repeat option. The range of days when you want tasks to run. You can select Repeat Indefinitely or configure an end date and time.

Time zones and schedules

Informatica Intelligent Cloud Services stores time in Coordinated Universal Time (UTC). When you log in, Informatica Intelligent Cloud Services converts the time and displays it in the time zone associated with your user profile.

When you create a schedule, you select the time zone for the scheduler to use. You can select a time zone that is different from your time zone or your organization time zone.

The time zone for system-configured schedules is UTC.

Daylight Savings Time changes and schedules

Informatica Intelligent Cloud Services applies Daylight Savings Time changes to all tasks except biweekly tasks.

When Daylight Savings time goes into effect, tasks scheduled to run between 2:00 a.m. and 2:59 a.m., do not run the day that the time changes from 2:00 a.m. to 3:00 a.m. If a task is scheduled to run biweekly at 2 a.m., it will run at 3 a.m. the day of the time change and at 2 a.m. for the next run.

Daylight Savings Time does not trigger additional runs for tasks that are scheduled to run between 1:00 a.m. - 1:59 a.m. when Standard Time begins. For example, a task is scheduled to run every day at 1:30 a.m. When the time changes from 2 a.m. to 1 a.m., the task does not run again at 1:30 a.m.

Tip: To ensure that Informatica Intelligent Cloud Services does not skip any scheduled runs near the 2 a.m. time change, do not schedule jobs to run between 12:59 a.m. and 3:01 a.m.

Creating a schedule

You can create a schedule in Data Integration when you configure a task. You can also create a schedule in Administrator.

The following procedure describes how to create a schedule when you access the **Schedule** page from Data Integration during task configuration.

1. Select **Run this task on a schedule**, and then click **New**.
2. Configure the following properties:

Property	Description
Schedule Name	Name of the schedule. Each schedule name must be unique within the organization. Schedule names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Maximum length is 100 characters. Schedule names are not case sensitive.
Description	Description of the schedule. Maximum length is 255 characters.
Starts	Date and time when the schedule takes effect. The date format is MM/DD/YYYY. Time appears in the 24-hour format. Click the calendar button to select the start date. The start date and time can affect the repeat frequency for tasks and taskflow jobs that repeat at regular intervals. For example, if the start date is November 10 and the repeat frequency is monthly, the schedule runs associated assets on the tenth day of each month. If the start time is 3:10 and the repeat frequency is hourly, the assets run every hour at 10 minutes past the hour. Default is the current date, current time, and time zone of the user that creates the schedule.

Property	Description
Time Zone	Select the time zone for the schedule to use. The time zone can differ from the organization time zone or user time zone.
Repeats	Repeat frequency for the schedule. Select one of the following options: <ul style="list-style-type: none"> - Does Not Repeat - Every N Minutes - Hourly - Daily - Weekly - Monthly Default is Does Not Repeat.

3. Click **Save** to save the schedule and return to the task configuration page.

Running a task on a schedule

Associate a task with a schedule on the **Schedule** page when you configure the task. You can use an existing schedule or create a schedule.

1. On the **Schedule** page for the task, select **Run this task on a schedule**.
2. To specify whether to use an existing schedule or a new schedule, perform one of the following tasks:
 - To use an existing schedule, select the schedule that you want to use.
 - To create a schedule to use for the task, click **New**, and then configure the schedule properties.
3. Click **Save**.

Email notification

You can configure email notification for a task. When you configure custom email notification, Data Integration uses the custom email notification instead of the email notification options configured for the organization.

To configure email notification options, perform the following steps in the task wizard:

1. Specify whether to use the default email notification options that have been set for your organization or create custom email notification for the task. Configure email notification using the following options:

Field	Description
Use Default Email Notification Options for my Organization	Use the email notification options configured for the organization.
Use Custom Email Notification Options for this Task	Use the email notification options configured for the task. You can send email to different addresses based on whether the task failed, completed with errors, or completed successfully. Use commas to separate a list of email addresses. When you select this option, email notification options configured for the organization are not used.

2. Click **Save**.

Preprocessing and postprocessing commands

You can run preprocessing and postprocessing commands to perform additional jobs. The task runs preprocessing commands before it reads the source. It runs postprocessing commands after it writes to the target.

You can use the following types of commands:

- SQL commands. Use SQL commands to perform database tasks.
- Operating system commands. Use shell and DOS commands to perform operating system tasks.

If any command in the preprocessing or postprocessing scripts fail, the task fails.

Preprocessing and postprocessing SQL commands

You can run SQL commands before or after a task. For example, you can use SQL commands to drop indexes on the target before the task runs, and then recreate them when the task completes. Data Integration does not validate the SQL.

Use the following rules and guidelines when creating the SQL commands:

- Use any command that is valid for the database type. However, Data Integration does not allow nested comments, even if the database allows them.
- Use a semicolon (;) to separate multiple statements. Data Integration issues a commit after each statement.
- Data Integration ignores semicolons within comments. If you need to use a semicolon outside of comments, you can escape it with a backslash (\).

Preprocessing and postprocessing operating system commands

Data Integration can perform operating system commands before or after the task runs.

You can use the following types of operating system commands:

- UNIX. Any valid UNIX command or shell script.
- Windows. Any valid DOS or batch file.

Enter multiple preprocessing or postprocessing commands as a single line without spaces.

If the Secure Agent is on a Windows machine, separate commands with an ampersand (&). If the Secure Agent is on a Linux machine, separate commands with a semicolon (;).

Monitoring a job

You can monitor tasks or taskflows that are currently running, have completed, or have stopped.

Monitor jobs on the following pages:


- Monitor the jobs that you initiated on the **My Jobs** page in Data Integration.
- Monitor running jobs in your organization on the **Running Jobs** page in Monitor.
- Monitor all jobs in your organization on the **All Jobs** page in Monitor.

For more information about monitoring jobs, see *Monitor*.

Stopping a job

A job is an instance of a mapping, task, or taskflow. You can stop a running job on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

1. Open Monitor and select **All Jobs** or **Running Jobs**, or open Data Integration and select **My Jobs**.
2. In the row that contains the job that you want to stop, click the **Stop** icon.



To view details about the stopped job, click the job name.

Guidelines for sources and targets in data integration tasks

Use the following rules and guidelines for sources and targets in data integration tasks:

- Field names must contain 65 characters or less.
- Field names must contain only alphanumeric or underscore characters. Spaces are not allowed.
- Field names cannot start with a number.
- Each field name must be unique within each source and target object.
- Data Integration truncates data if the scale or precision of a numeric target column is less than the scale or precision of the corresponding source column.

Rules and guidelines for flat file sources

Use the following rules and guidelines for flat file sources:

- All date columns in a flat file source must have the same date format.
- The flat file cannot contain empty column names. If a file contains an empty column name, the following error appears:

```
Invalid header line: Empty column name found.
```

- Column names in a flat file must contain printable tab or ASCII characters (ASCII code 32-126). If the file contains a character that is not valid, the following error appears:

```
Invalid header line: Non-printable character found. The file might be binary or might have invalid characters in the header line.
```

- You can use a tab, space, or any printable special character as a delimiter. The delimiter can have a maximum of 10 characters. The delimiter must be different from the escape character and text qualifier.
- For flat file sources with multibyte data on Linux, the default locale must be UTF-8.

Rules and guidelines for database sources and targets

Use the following rules and guidelines for database sources and targets:

- You can use database tables as targets. You can use database tables, aliases, and views as sources.
- Relational targets must meet the minimum system requirements.
- The database user account for each database target connection must have DELETE, INSERT, SELECT, and UPDATE privileges.

CHAPTER 2

Mapping tasks

Use the mapping task to process data based on the data flow logic defined in a mapping.

When you create a mapping task, you select the mapping for the task to use. The mapping must already exist before you can create a mapping task for it. Alternatively, you can create a mapping task using a template.

If the mapping includes parameters, you can define the parameters when you configure the task or define the parameters when you run the task. You can use user-defined parameters for data filters, expressions, and lookup expressions in a mapping task. You define user-defined parameters in a parameter file associated with the task.


At run time, a mapping task processes task data based on the data flow logic from the mapping, the parameters defined in the task, and the user-defined parameters defined in a parameter file, when available.

Mapping task templates

Use a mapping task template to run a mapping task without creating a mapping beforehand.

Each mapping task template is based upon a mapping template. Use a mapping task template when the mapping on which the mapping task template is based suits your needs. When you select a mapping task template, Data Integration creates a copy of the template for you to use. When you define the mapping task in the task wizard, you save a copy of the mapping template on which the mapping task template is based.

Templates are divided into three categories: Integration, Cleansing, and Warehousing, as shown in the following image:



The templates range from simple templates that you can use to copy data from one source to another, to complex templates that you can use for data warehousing-related tasks.

Related objects

When a mapping includes a source that is a parameter and is configured for multiple objects, you can join related objects in the task.

You can join related objects based on existing relationships or custom relationships. Data Integration restricts the type of relationships that you can create based on the connection type.

Use the following relationships to join related objects:

Existing relationships

You can use relationships defined in the source system to join related objects. You can join objects with existing relationships for Salesforce, database, and some Data Integration Connectors connection types.

After you select a primary object, you select a related object from a list of related objects.

Custom relationships

You can use custom relationships to join multiple source objects. You can create custom relationships for the database connection type.

When you create a custom relationship for database objects, you create an inner, left outer, or right outer join on the source fields that you select.

To join source objects, you add the primary source object in the **Objects and Relationships** table. Then you add related objects, specify keys for the primary and related objects, and configure the join type and operator.

For more information about related source objects, see the Source Transformation section in *Transformations*.

New MappingTask7

1 Definition 2 Sources 3 Schedule

Source Parameter Details

NewSource Connection: * mysql View...

Source Type: Multiple

Objects and Relationships

Source Object	Join Condition
a_r104_mar24_al_df_account	
a_r10_2_account	a_r10_2_account.ID = a_r104_mar24_al_df_account.ID
a_r10_c_lookup__c	a_r10_c_lookup__c.ID = a_r104_mar24_al_df_account.ID

Query Options

Filter: [Configure...](#)

Sort: [Configure...](#)

☐ Display source fields in alphabetical order

Data Preview

a_r104_mar24_al_df_account Preview All Columns (Total columns: 45)

Save < Back Next > Finish Cancel

Advanced relationships

You can create an advanced relationship for database sources when the source object in the mapping is a parameter and configured for multiple sources. You cannot create an advanced relationship between source objects that have been joined using a custom relationship.

When you create an advanced relationship, the wizard converts any relationships that you defined to an SQL statement that you can edit.

To create an advanced relationship, you add the primary source object in the Objects and Relationships table. Then you select fields and write the SQL statement that you want to use. Use an SQL statement that is valid for the source database. You can also add additional objects from the source.

Pushdown optimization

Pushdown optimization pushes some of the transformation logic in a mapping to source or target databases for execution, which can improve task performance. By default, pushdown optimization is enabled for all mapping tasks.

When you run a task configured for pushdown optimization, the task converts the transformation logic to an SQL query. The task sends the query to the database, and the database executes the query.

The amount of transformation logic that you can push to the database depends on the database, transformation logic, and task configuration. The task processes all transformation logic that it cannot push to a database.

Use the **Pushdown Optimization** advanced session properties to configure pushdown optimization for a task. You can configure the pushdown optimization type, whether to cancel the task if pushdown optimization is not possible, whether to allow the task to create temporary views in the database, and other options.

Note: Pushdown optimization functionality varies depending on the support available for the connector. For more information, see the help for the appropriate connector.

Simultaneous task runs

You can run multiple instances of a mapping task at the same time.

You can use multiple task instances in a Parallel Paths step of a taskflow or in two different taskflows that run in parallel.

To enable simultaneous task runs, select the **Allow the mapping task to be executed simultaneously** on the **Schedule** tab when you configure the task.

Use caution when you configure mapping tasks to run simultaneously. Mapping features that change each time the task runs, such as in-out parameters and sequence generator values, might produce unexpected results when you run the task instances simultaneously.

Field metadata

You can view and edit field metadata such as the type, precision, and scale for parameterized source and lookup objects with certain connection types.

To view and edit field metadata, use the **Edit Types** option on the appropriate page of the mapping task wizard. You configure source field metadata on the **Sources** page and lookup field metadata on the **Input Parameters** page.

If you edit field metadata in the task and the field metadata changes after the task is saved, Data Integration uses the updated metadata. Typically, this is the desired behavior. However, if the task uses a flat file connection and you want to retain the metadata used at design time, enable the **Retain existing fields at runtime option**.

To see if a connector supports field metadata configuration, see the help for the appropriate connector.

Schema change handling

You can choose how Data Integration handles changes that you make to the schema of some data object types.

By default, if you make changes to the schema, Data Integration does not pick up the changes automatically. If you want Data Integration to refresh the data object schema every time the mapping task runs, you can enable dynamic schema handling.

A schema change includes one or more of the following changes to the data object:

- Fields are added.
- Fields are deleted.
- Fields are renamed.
- Field data type, precision, or scale is updated.

Data Integration automatically refreshes the schema for relational objects every time the task runs. If you want to dynamically refresh the schema for other object types, enable dynamic schema change handling on the **Schedule** page when you configure the task.

The following table describes the schema change handling options:

Option	Description
Asynchronous	Default. Data Integration refreshes the schema when you edit the mapping or mapping task, and when Informatica Intelligent Cloud Services is upgraded.
Dynamic	Data Integration refreshes the schema every time the task runs. Applicable for source, target, and lookup objects of certain connector types. For some connector types, Data Integration can only refresh the schema if the data object is a flat file. If you select this option, the file object format must be delimited. Not applicable to hierarchical data. To see if a connector supports dynamic schema change handling, see the help for the appropriate connector.

If you update fields in the source object and you enable dynamic schema handling, be sure to update the Target transformation field mapping. Data Integration writes Null to the target fields that were previously mapped to the renamed or deleted source fields. If you use a target created at run time, update the target object name so that Data Integration creates a new target when the task runs. The task fails if Data Integration tries to alter a target created in a previous task run.

Dynamic schema handling options

When you enable dynamic schema change handling, you can select how Data Integration applies schema changes from upstream transformations to the target object. If the mapping contains more than one target, select the schema change handling for each target.

To select target schema options, the target field mapping must be automatic.

When you configure target schema options for objects that are created at runtime, Data Integration creates the target the first time you run the task. In subsequent task runs, Data Integration updates the target based on the schema change option that you select.

The schema change handling options available are based on the target connection. To see if a connector supports dynamic schema change handling, see the help for the appropriate connector.

The following table describes the schema handling options:

Schema handling option	Target type	Description
Drop Current and Recreate	Database	Data Integration drops the existing target table and creates a new target table with the schema from the upstream transformations on every run.
Alter and Apply Changes	Database	Data Integration updates the target schema with additive changes to match the schema from the upstream transformations. It does not delete columns from the target.
Don't Apply DDL Changes	Database	Data Integration fetches the target schema at runtime and does not apply upstream schema changes to the target table.

Data Integration does not pass field constraints to the target. For example, the source contains fields S1 and S2 configured with the NOT NULL constraint. The target contains fields T1 and T2 also configured with the NOT NULL constraint. You select the **Alter and Apply Changes** schema handling option. When you run the task, fields S1 and S2 are written to the target with no constraints.

Dynamic schema change handling rules and guidelines

Enable dynamic schema change handling so that Data Integration refreshes the data object schema every time the mapping task runs.

Consider the following rules and guidelines when you enable dynamic schema change handling:

- Changes to the object schema take precedence over changes to the field metadata in the mapping. For example, you add a field to the source object and then edit the metadata of an existing field in the mapping. At run time, Data Integration adds the new field and does not edit the existing field.
- Data Integration resolves parameters before picking up the object schema.
- Data Integration treats renamed fields as deleted and added columns. If you rename a field, you might need to update transformations that reference the renamed field. For example, if you rename a field that is used in the lookup condition, the lookup cannot find the new field and the task fails.
- When you rename, add, or delete fields, you might need to update the field mapping. For example, if you delete all the previously mapped fields in a target object, you must remap at least one field or the task fails.
- Data Integration writes Null values to a target field in the following situations:
 - You rename a target field with automatic field mapping, and the field name does not match a source field.
 - You rename a source field with manual field mapping, and you do not remap the field to the target.
- If you delete a field from a source or lookup object and a downstream transformation references the field, the task fails.
- If you change a source or lookup field type, the task might fail if the new field type results in errors downstream. For example, if you change an integer field in an arithmetic expression to a string field, the expression is not valid and the task fails.
- If you change a target field type, Data Integration converts the data from the incoming field to the new target field type. If the conversion results in an error, Data Integration drops the row. For example if you change a string type to a date type where the string does not contain a date, Data Integration drops the row.

Schema mismatch handling

You can choose how Data Integration responds when it encounters a mismatch between a file and the selected schema.

When you run a mapping through a mapping task, you can configure an advanced option that specifies how Data Integration handles mismatched schemas:

- Skip mismatched files and continue. When Data Integration finds a schema mismatch, it stops searching for other errors in the same file and writes that error to the log. Data Integration doesn't process any other records from that file and continues with the next file.
- Stop on first mismatched file. Data Integration stops all processing when it finds a schema mismatch error, and it writes the error to the log. Data Integration does not roll backfiles processed before the error was found, and it doesn't process the file containing the error.

When you run a mapping in the Mapping Designer, Data Integration evaluates every file and skips the entire file when it encounters a schema mismatch.

Mapping task configuration

Use the mapping task wizard to create a mapping task.

Complete the following steps in the wizard to create a mapping task:

1. Create the mapping task.
2. Configure the source.
3. Configure the target.
4. Define other parameters.
5. Optionally, configure a schedule and advanced options.

As you work through the task wizard, you can click **Save** to save your work at any time. When you have completed the wizard, you can click **Finish** to save and close the task wizard.

Defining a mapping task

1. To create a mapping task, click **New > Tasks** and then complete one of the following steps:
 - To create a mapping task based on a mapping, select **Mapping Task** and click **Create**.
 - To create a mapping task using a template, expand the appropriate template category and select the template you want to use, and then click **Create**.

To edit a mapping task, on the **Explore** page, navigate to the mapping task. In the row that contains the task, click **Actions** and select **Edit**.

2. Configure the following fields:

Field	Description
Task Name	Name of the task. Task names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Maximum length is 100 characters. Task names are not case sensitive.
Location	Project folder in which the task resides. If the Explore page is currently active and a project or folder is selected, the default location for the asset is the selected project or folder. Otherwise, the default location is the location of the most recently saved asset.
Description	Description of the task. Maximum length is 4000 characters.
Runtime Environment	Runtime environment that contains the Secure Agent to run the task.
Task Based On	The basis for the task. Select Mapping.
Mapping	Mapping associated with the task. If you are creating a mapping task based on a mapping, perform the following steps: 1. Click Select and navigate to the mapping that you want to use. 2. Select the mapping and click Select . If you are using a mapping task template, the name of the mapping template on which the mapping task template is based displays in the Mapping field. By default, the mapping is saved in the same folder as the mapping task. Optionally, perform the following steps: 1. In the Mapping field, change the name of the mapping. 2. To save the mapping in a different folder, click Browse .

3. Click **Next**.

Configuring sources

The **Sources** page displays differently based on the basis for the task. If the mapping does not include source parameters, the **Sources** page might not appear.

You can add a single source object or multiple source objects based on the connection type and the mapping configuration. You can also configure a source filter.

After you configure the source, validate the configuration.

If the mapping specifies a connection parameter and after you create the mapping task, you edit the mapping to change the source object, you might need to edit the task. Check the specific connection in the mapping task to determine if you need to reset it and then validate the task again

1. On the **Sources** page, configure the following details as required:

Source parameter detail	Description
Connection	Select a connection.
Object	<p>Select a source object.</p> <p>If a list of objects does not appear, click Select.</p> <p>The Select Source Object dialog box displays up to 200 objects. If the object you want to use does not appear, enter a search string to reduce the number of objects that appear.</p> <p>To configure advanced properties for the source, click Advanced. Not available for all connection types. The properties available vary based on the connection type.</p> <p>When you edit advanced properties, you can restore default values for individual properties or for all advanced properties. Default values are defined in the mapping that the task is based on.</p> <p>For information about a particular connector's properties, see the help for the appropriate connector.</p> <p>Object field names cannot start with a number.</p>
Add Currently Processed File Name	<p>Adds the source file name to each row. Data Integration adds the <code>CurrentlyProcessedFileName</code> field to the source at run time.</p> <p>Available for parameterized source objects with flat file connections.</p>
Display Technical Field Names Instead of Labels	<p>Displays technical names instead of business names.</p> <p>Not available for all connection types.</p>
Display Source Fields in Alphabetical Order	<p>Displays source fields in alphabetical order.</p> <p>By default, fields appear in the order returned by the source system.</p>

2. For a parameterized source object, configure field metadata if required. You can configure field metadata for sources with certain connection types. To see if a connector supports field metadata configuration, see the help for the appropriate connector.

To configure field metadata, click **Edit Types**. In the **Edit Field datatypes** dialog box, configure the following attributes and click **OK**:

Data type attribute	Description
Retain existing fields at runtime	<p>When enabled, the task uses the field metadata that is configured in the task.</p> <p>If field metadata changes after a task is saved, Data Integration uses the updated field metadata. Typically, this is the desired behavior. However, if the task uses a flat file connection and you want to retain the metadata used at design time, enable this option.</p>
Datatype	The data type of the field.

Data type attribute	Description
Precision	Total number of digits in a number. For example, the number 123.45 has a precision of 5. The precision must be greater than or equal to 1.
Scale	Number of digits to the right of the decimal point of a number. For example, the number 123.45 has a scale of 2. Scale must be greater than or equal to 0. The scale of a number must be less than its precision. The maximum scale for a numeric data type is 65535. Not editable for all data types.

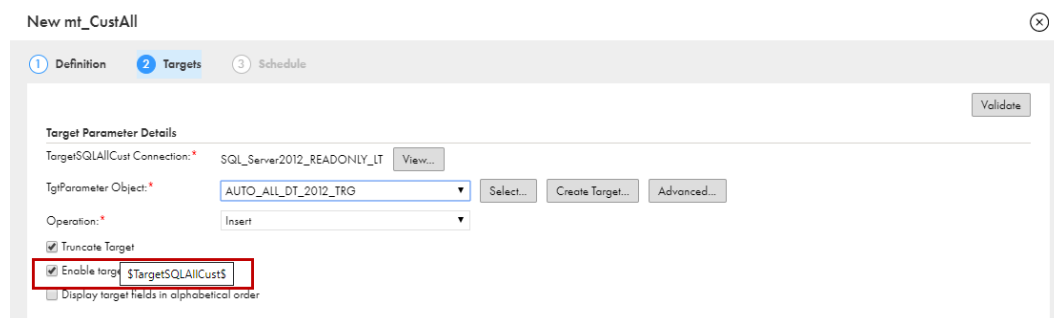
- If necessary, configure query options.
- If necessary, configure maplet details.
This appears when a maplet requires a connection, and the maplet is configured to display on this page.
For some connection types, you can select **Display technical field names instead of labels** to display technical names instead of business names.
To display fields in alphabetical order, click **Display maplet fields in alphabetical order**.
- If necessary, configure stored procedure details.
This displays when a stored procedure requires a connection, and the stored procedure is configured to appear on this page.
- To test the source connection, click **Validate**.
- Click **Next**.

Configuring targets

The **Targets** page displays differently depending on the basis for the task.

For mappings, the **Targets** page displays when the mapping includes parameters for target connections or target objects. The properties that you need to specify are based on the type of parameter. For example, if the target is parameterized but the connection is not, you must specify a target and you can optionally change the connection.

Tip: All of the properties display in one list, even when the task includes multiple objects. Place your the cursor over the properties to determine which objects the properties apply to. For example, in the following image, the **Enable target bulk load** parameter applies to the target named `TargetSQLAllCust`.



Note the following additional information about template parameters:

- For a mapplet, you select a connection. You do not select objects for mapplets.
- When a connection name displays without surrounding dollar signs, it is a logical connection. If the logical connection is associated with multiple objects on the **Targets** page, you select the logical connection once, and then select each object.
- If the logical connection is associated with objects on other pages of the task wizard, be sure to use the same connection for logical connections with the same name.

When you select an object, the **Data Preview** area displays a portion of the data in the object. Data preview displays the first ten rows of the first five columns in the object. It also displays the total number of columns in the object.

If the page has more than one object, you can select the object in the **Data Preview** area to display its data.

Data preview does not display the following types of data:

- Mapplet data.
- Certain Unicode characters.
- Binary data. If the object contains binary data, data preview shows the following text:

BINARY DATA

1. On the **Targets** page, configure the following details as required:

Target parameter detail	Description
Connection	Select a connection. To create a connection, click New . To edit a connection, click View , and in the View Connection dialog box, click Edit .
Object	Select a target object. If a list of objects does not appear, click Select . The Select Target Object dialog box displays up to 200 objects. If the object you want to use does not display, enter a search string to reduce the number of objects that display. To configure advanced properties for the target, click Advanced . Not available for all connection types. The properties available vary based on the connection type. When you edit advanced properties, you can restore the default values for individual properties or for all properties. Default values are defined in the mapping that the task is based on. For information about a particular connector's properties, see the help for the appropriate connector. Object field names cannot start with a number.
Display Technical Field Names Instead of Labels	Displays technical names instead of business names. Not available for all connection types.
Display Target Fields in Alphabetical Order	Displays target fields in alphabetical order. By default, fields appear in the order returned by the target system.

2. Click **Next**.

Configuring parameters

The **Input Parameters** page displays parameters that are not in the Source transformation or Target transformation. Depending on the mapping data flow, you might need to configure some parameters before the task wizard allows you to configure other parameters. For more information, see *Mappings*.

Configuring input or in-out parameters

1. For a task based on a mapping, on the **Input Parameters** or **In-Out Parameters** page, configure the parameters that display.

Depending on the data flow of the mapping, you might need to configure certain parameters before the task wizard allows you to configure other parameters.

For more information on using in-out parameters and input parameters in mappings, see *Mappings*.
2. If necessary, configure lookup details.

This displays when a lookup requires a connection or object and the lookup is configured to display on this page. If you need to select a lookup object, select the object from the list. If a list of objects does not appear, click **Select**.

For some connection types, click **Advanced** to configure advanced properties for the connection. For more information, see *Connections*.

For some connection types, you can select **Display technical field names instead of labels** to display technical names instead of business names.

To display fields in alphabetical order, click **Display lookup fields in alphabetical order**.
3. If necessary, configure mapplet details.

This displays when a mapplet requires a connection and the mapplet is configured to display on this page.

For some connection types, you can select **Display technical field names instead of labels** to display technical names instead of business names.

To display fields in alphabetical order, click **Display mapplet fields in alphabetical order**.
4. If necessary, configure stored procedure details.

This displays when a stored procedure requires a connection and the stored procedure is configured to display on this page.
5. Define the remaining template parameters, as needed.
6. Click **Next**.

Configuring a schedule and advanced options

On the **Schedule** page, you specify whether to run a mapping task manually or schedule it to run at a specific time or interval. You can create a schedule or use an existing schedule.

You can also configure email notification and advanced options for the task on the **Schedule** page.

1. To specify whether to run the task on a schedule or without a schedule, choose one of the following options:
 - If you want to run the task on a schedule, click **Run this task on schedule**. Select the schedule you want to use or click **New** to create a schedule.
 - If you want to run the task without a schedule, click **Do not run this task on a schedule**.
2. Configure email notification options for the task.

3. Optionally, enter the following advanced options:

Field	Description
Pre-Processing Commands	Commands to run before the task.
Post-Processing Commands	Commands to run after the task completes.
Maximum Number of Log Files	<p>Number of session log files to retain. By default, Data Integration stores each type of log file for 10 runs before it overwrites the log files for new runs.</p> <p>Note: If a dollar sign (\$) is present in a custom session log file name, for example, MyLog_ \$CurrentTime, the file name is dynamic. If you customize the session log file name using a dynamic name, the Maximum Number of Log Files property doesn't apply. To purge old log files, delete the files manually.</p>
Schema Change Handling	<p>Determines how Data Integration picks up changes to the object schema. Select one of the following options:</p> <ul style="list-style-type: none">- Asynchronous. Data Integration refreshes the schema when you update the mapping or mapping task, and after an upgrade.- Dynamic. Data Integration refreshes the schema every time the task runs. <p>Default is Asynchronous.</p>
Dynamic Schema Handling	<p>Determines how Data Integration applies schema changes from upstream transformations to the target object. Available when the schema change handling is dynamic and the field mapping is automatic.</p> <p>For each target, select how Data Integration updates the target schema. The options available are based on the target connection.</p> <p>For more information, see "Schema change handling" on page 33 or the help for the appropriate connector.</p>

4. Optionally, if the mapping task contains parameters, you can use parameter values from a parameter file. Choose one of the following options:

- To use a parameter file on a local machine, select **Local**. Enter the following information:

Field	Description
Parameter File Directory	<p>Path for the directory that contains the parameter file, excluding the parameter file name. The Secure Agent must be able to access the directory.</p> <p>You can use an absolute file path or a path relative to one of the following \$PM system variables:</p> <ul style="list-style-type: none"> - \$PMRootDir - \$PMSourceFileDir - \$PMLookupFileDir - \$PMCacheDir - \$PMSessionLogDir - \$PMExtProcDir - \$PMTempDir <p>By default, Data Integration uses the following parameter file directory:</p> <p><Secure Agent installation directory>/apps/Data_Integration_Server/data/userparameters</p>
Parameter File Name	<p>Name of the file that contains the definitions and values of user-defined parameters used in the task.</p> <p>You can provide the file name or the relative path and file name in this field.</p>

- To use a cloud-hosted file, select **Cloud Hosted**. Enter the following information about the file:

Field	Description
Connection	<p>Connection where the parameter file is stored. You can use the following connection types:</p> <ul style="list-style-type: none"> - Amazon S3 - Google Storage V2 - Azure Data Lake Store Gen2
Object	<p>Name of the file that contains the definitions and values of user-defined parameters used in the task.</p>

- Optionally, if you want to create a parameter file based on the parameters and default values specified in the mapping on which the task is based, click **Download Parameter File Template**.

For more information about parameter file templates, see *Mappings*.

- Choose whether to run the task in standard or verbose execution mode.

If you select verbose mode, the mapping generates additional data in the logs that you can use for troubleshooting. Select verbose execution mode only for troubleshooting purposes. Verbose execution mode impacts performance because of the amount of data it generates.

7. Optionally, configure the following pushdown optimization properties:

Property	Description
Pushdown Optimization	<p>Type of pushdown optimization. Use one of the following options:</p> <ul style="list-style-type: none"> - None. The task processes all transformation logic for the task. - To Source. The task pushes as much of the transformation logic to the source database as possible. - To Target. The task pushes as much of the transformation logic to the target database as possible. - Full. The task pushes as much of the transformation logic to the source and target databases as possible. The task processes any transformation logic that it cannot push to a database. - \$\$PushdownConfig. The task uses the pushdown optimization type specified in the user-defined parameter file for the task. <p>When you use \$\$PushdownConfig, ensure that the user-defined parameter is configured in the parameter file.</p> <p>When you use pushdown optimization, do not use the Error Log Type advanced session property.</p> <p>Default is Full.</p> <p>Note: Pushdown optimization functionality varies depending on the support available for the connector. For more information, see the help for the appropriate connector.</p>
Optimization Context Type	<p>Provides context about the mapping configuration for pushdown optimization. If you select an option other than None, Data Integration constructs a single query for pushdown optimization by combining multiple targets in the mapping based on the target configurations. If you select None, the query is not optimized.</p> <p>If Data Integration cannot apply the selected context, Data Integration uses the default pushdown optimization behavior.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - None - SCD Type 2 merge - Multi-insert <p>Default is None.</p> <p>For more information, see the help for the appropriate connector.</p>
If pushdown mode is not possible, cancel the task	<p>Cancels the task if the selected pushdown optimization type is unavailable for the connection.</p> <p>Default is disabled.</p>
Create Temporary View	<p>Allows the task to create temporary view objects in the database when it pushes the task to the database.</p> <p>Use when the task includes an SQL override in the Source Qualifier transformation or Lookup transformation.</p> <p>Default is enabled.</p> <p>Disabled when the pushdown optimization type is None.</p>
Create Temporary Sequence	<p>This option is not used.</p>

8. Choose to enable cross-schema pushdown optimization.
9. If you want to run multiple instances of the task at the same time, enable simultaneous runs of the mapping task.
Some mapping features might produce unexpected results in simultaneous task runs.
10. Click **Finish**.

Viewing and editing mapping task details

You can view details about a mapping task, such as the mapping or runtime environment used by the task.

The **Task Details** page includes the following information:

- The runtime environment used to run the mapping task.
- The date the task was created and the user who created the task.
- The last time the task was updated and the user who updated the task.
- The date of the last run.
- The name and image of the mapping on which the task is based and the date the mapping was last updated.
- Pre-processing and post-processing commands.

To view details for a mapping task, perform the following steps:

1. On the **Explore** page, navigate to the task.
2. In the row that contains the task, click **Actions** and select **View**.

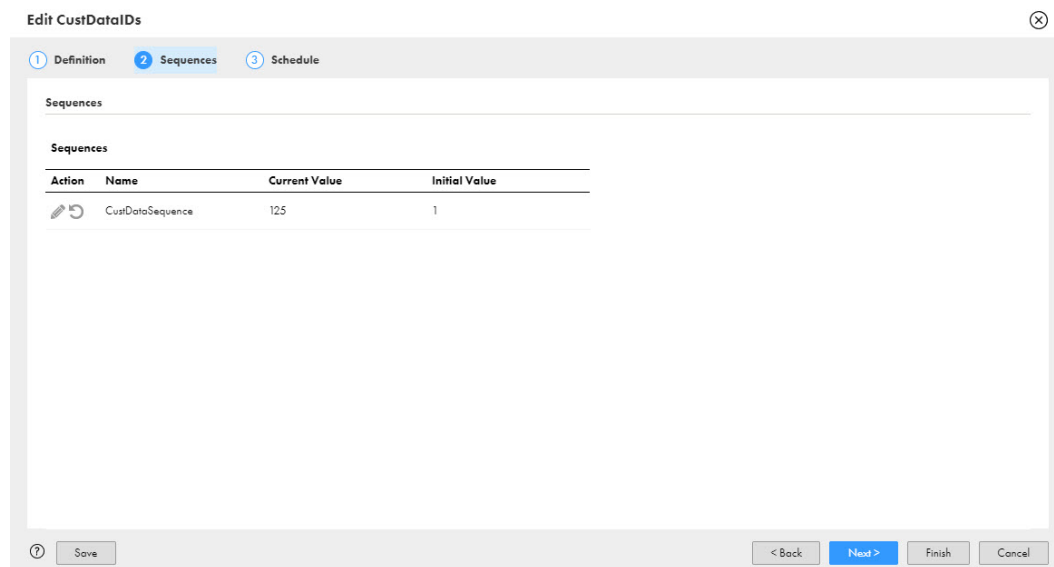
On the **Task Details** page, you can click **Edit** to modify the mapping task.

Sequence Generator values

When you run a mapping task that includes a Sequence Generator transformation in the mapping, you can change the beginning value for the sequence.


To change the beginning value, you change the **Current Value** field in the **Sequences** page in the mapping task wizard. The **Current Value** field shows the first value the task will generate in the sequence, based on the last value generated in the last task execution.

For example, the last time you ran the CustDataIDs task, the last value generated was 124. The next time the task runs, the first number in the sequence is 125 because the Sequence Generator transformation is configured to increment by 1. If you want the sequence to begin with 200, you change the **Current Value** to 200.



Dialog box titled "Edit CustDataIDs" with a close button (X) in the top right corner. The dialog has three tabs: "1 Definition", "2 Sequences" (selected), and "3 Schedule".

Under the "Sequences" tab, there is a section labeled "Sequences" containing a table:

Action	Name	Current Value	Initial Value
	CustDataSequence	125	1

At the bottom of the dialog, there are four buttons: "Save" (disabled), "< Back" (disabled), "Next >" (active/highlighted), "Finish" (disabled), and "Cancel" (disabled).

Running a mapping task

Run a mapping task after you configure all of the required settings.

You can run a mapping task in the following ways:

- **Manually.** To run a mapping task manually, on the **Explore** page, navigate to the task. In the row that contains the task, click **Actions** and select **Run**.
You can also run a mapping task manually, with or without advanced options, from the **Task Details** page. To access the **Task Details** page, click **Actions** and select **View**.
- **On a schedule.** To run a mapping task on a schedule, edit the task in the mapping task wizard to associate the task with a schedule.

Reprocessing incrementally-loaded source files

In the advanced options, you can configure properties to reprocess incrementally-loaded source files. Reprocess source files to create a snapshot of the data from a specified time interval, to debug and discover the source of bad data in your target, and to restore deleted data.

Jobs that reprocess incrementally-loaded files read and process the incrementally-loaded files that were modified during the time that you specify. You can choose to reprocess files that were modified in a given time interval or all files modified after a given start time.

You can reprocess source files when at least one source is configured to incrementally load files and the files have been loaded at least once. The advanced options are not available if the mapping does not have sources configured to incrementally load files, the mapping task has never run, or the last load time was reset.

Reprocessing does not change the mapping task's last load time, so future jobs created from the mapping task run and load data as usual.

Note: Processing is unpredictable when the start or end time of a reprocessing job is within the repeated hour on the day that Daylight Saving Time ends.

CHAPTER 3

Data transfer tasks

Use a data transfer task to transfer data from a source to a target. For example, you might use a data transfer task to transfer data from an on-premises database to a cloud data warehouse.

When you configure a data transfer task, you can augment the source data with data from a lookup source. Based on the source connection that you use, you can also sort and filter the data before you load it to the target.

To see if a data transfer task is applicable to the connectors you are using, see the help for the relevant connectors.

Task operations

When you configure a data transfer task, you specify the task operation. The operations available are based on the target that you select.

You can select the following operations:

Insert

Inserts all source rows into the target. If Data Integration finds a source row that exists in the target, the row fails.

Update

Updates rows in the target that exist in the source. If Data Integration finds a row in the source that does not exist in the target, the row fails.

Upsert

Updates all rows in the target that also exist in the source and inserts all new source rows in to the target.

If a source field contains a NULL value and the corresponding target field contains a value, Data Integration retains the existing value in the target field.

Delete

Deletes all rows from the target that exist in the source.

Data transfer task sources

You can select a single source to transfer data from.

The formatting and advanced options that you can configure for the source depend on the source connection that you select. For example, for a flat file source, you can configure formatting options such as the formatting type. For Salesforce sources, you can configure advanced options such as the SOQL filter condition, row limit, and bulk query.

For information about the options that you can configure for a source connection, see the help for the appropriate connector.

You can preview the data in the source. The preview returns the first 10 rows. To preview rows in alphabetical order, select **Display fields in alphabetical order**. Data Integration does not change the order of rows in the actual source. You can also download the preview results as a CSV file.

Source filters

Apply filter conditions to filter the source data that you transfer to the target.

You can apply the following types of filter conditions:

Simple

Select the source field, and configure the operator and value to use in the filter.

When you define more than one filter condition, the task evaluates them in the order that you specify.

The task evaluates the filter conditions using the AND logical operator to join the conditions. It returns rows that match all the filter conditions.

Advanced

Create a filter expression using the expression editor. You enter one expression that contains all filters. You can use source fields and built-in functions in the expression.

For more information about advanced filters, see [“Advanced data filters” on page 9](#).

Sort conditions

For certain source types, you can sort the source data to provided sorted data to the target.

When you sort data, you select one or more source fields to sort by. If you apply more than one sort condition, Data Integration sorts fields in the listed order.

To see if a connector supports sorting, see the help for the appropriate connector.

Second sources

When you configure a data transfer task, you can add a second source to use as a lookup source. Configure the lookup source on the **Second Source** page.

The task queries the lookup source based on the lookup condition that you specify and returns the result of the lookup to the target.

Select a second source when you want to augment the source data with a related value or values from the lookup source. For example, the source is an orders table that contains a customer ID field. You might

retrieve the customer name and address from the lookup source so that you can include them in the target object. The task returns all fields from the lookup source.

To optimize performance, the task caches the lookup source. The cache remains static and does not change as the task runs. The task deletes the cache files after the task completes.

You can preview the data in the lookup source. The preview returns the first 10 rows. To preview rows in alphabetical order, select **Display fields in alphabetical order**. Data Integration does not change the order of rows in the actual source. You can also download the preview results to a CSV file.

You can also filter the data from both sources before writing it to the target.

Lookup condition

When you select a second source to use as a lookup source, you must configure one or more lookup conditions.

A lookup condition defines when the lookup returns values from the lookup source. When you configure a lookup condition, you compare the value of one or more fields from the original source with values in the lookup source.

A lookup condition includes an incoming field from the original source, a field from the lookup source, and an operator. To avoid possible naming conflicts, the data transfer task applies the prefix `SRC_` to the fields from the original source. If this results in a naming conflict for any field from the original source, the task applies the prefix `IN_SRC_` to the field from the original source.

For example, you might configure the following lookup condition when the original source contains the `CustID` field, the lookup source contains the `CustomerID` field, and you want to return values from the lookup source when the customer IDs match:

Lookup Field	Operator	Incoming Field
CustomerID	Equals	SRC_CustID

You can use the following operators in a lookup condition:

- Equals
- Not Equals
- Less Than
- Less Than or Equals
- Greater Than
- Greater Than or Equals

When you enter multiple conditions, the task evaluates the lookup conditions using the AND logical operator to join the conditions. It returns rows that match all the lookup conditions.

When you include multiple conditions, to optimize performance, enter the conditions in the following order:

1. Equals
2. Less Than, Less Than or Equals, Greater Than, Greater Than or Equals
3. Not Equals

The lookup condition matches null values. When an input field is NULL, the task evaluates the NULL equal to null values in the lookup.

If the lookup condition has multiple matches, the task returns any row.

Second source filters

You can apply filter conditions to filter the combined data.

You can configure the following types of filters:

Simple

To configure a filter condition, select a source field and configure the operator and value to use in the filter. You can select a field from either source. Fields from the original source are prefixed with the characters `SRC_` or `IN_SRC_`.

When you define more than one filter condition, the task evaluates them in the order that you specify. The task evaluates the filter conditions using the AND logical operator to join the conditions. It returns rows that match all the filter conditions.

Advanced

Create a filter expression using the expression editor. You enter one expression that contains all filters. You can use source fields and built-in functions in the expression.

For more information about advanced filters, see [“Advanced data filters” on page 9](#).

Data transfer task targets

You can use a single object as a target for a data transfer task. Select a target object or create a new target object at run time.

The task operations that you can select depend on the target connection that you use. For more information about task operations for different target types, see the help for the appropriate connector.

You can preview the data in the target. The preview returns the first 10 rows. You can download the preview results to a CSV file. To preview rows in alphabetical order, select **Display fields in alphabetical order**. Data Integration does not change the order of rows in the actual target.

Database target truncation

You can configure a data transfer task to truncate a database target table before writing new data to the table when the task uses an Insert operation. By default, Data Integration inserts new rows without truncating the target table.

Update columns

Update columns are columns that uniquely identify rows in the target table. Add update columns when the database target table does not contain a primary key and the data transfer task uses an update, upsert, or delete operation.

When you run the data transfer task, the task uses the field mapping to match rows in the source to the database table. If the data transfer task matches a source row to multiple target rows, it performs the specified task operation on all matched target rows.

Field mapping

Configure field mapping in a data transfer task to map source fields to target fields. Configure field mapping on the **Field Mapping** page of the data transfer task wizard.

You must map at least one source field to a target field. If the task uses multiple sources, fields from the original source are prefixed with the characters `SRC_` or `IN_SRC_`.

You can configure the following field mapping options:

Options

Configure which fields to display. Click **Options** and select from the following display options:

- Show All
- Show Mapped
- Show Unmapped

Automap

Data Integration automatically links fields with the same name or similar name. Click **Automap** and select from the following mapping options:

- Exact Field Name. Data Integration matches fields of the same name.
- Smart Map. Data Integration matches fields with similar names. For example, if you have an incoming field `Cust_Name` and a target field `Customer_Name`, Data Integration automatically links the `Cust_Name` field with the `Customer_Name` field.
- Undo Automap. Data Integration clears fields mapping with Smart Map or Exact Field Name but does not clear manually mapped fields.

Actions

Additional field link options. Provides the following options:

- Map Selected. Links the selected incoming field with the selected target field.
- Unmap Selected. Clears the link for the selected field.
- Clear Mapping. Clears all field mappings.

After you map a field, if you want to configure a field expression, click the mapped field name. You can include fields and built-in functions in the expression but not user-defined functions.

When you create a target at run time, Data Integration maps the source fields to the target fields. You cannot unmap or edit the source fields mapped to the target but you can add fields to the target. You can also edit the mapped field expression and metadata and reorder the added fields. You cannot reorder the fields copied from the source.

Field data types

When you create a data transfer task, Data Integration assigns a data type to each field in the source and target. When you add a field to a target that you create at run time, you select the data type.

Configuring a data transfer task

To configure a data transfer task, perform the following steps:

1. Define the data transfer task.
2. Configure the source.
3. Optionally, configure a second source to use as a lookup source.
4. Configure the target.
5. Configure field mappings.
6. Configure runtime options.

As you work through the data transfer task wizard, you can click **Save** to save your work at any time. Use the **Validation** panel to validate the task. When you have completed the wizard, click **Exit** to close the task wizard.

Defining the data transfer task

1. To create a new data transfer task click **New > Tasks**. Select **Data Transfer Task** and click **Create**.
To edit a data transfer task, on the **Explore** page, navigate to the data transfer task. In the row that contains the task, click **Actions** and select **Edit**.
2. Configure the following properties:

Property	Description
Name	Name of the data transfer task.
Location	Project and folder in which the task resides.
Description	Description of the data transfer task. Maximum length is 4000 characters.
Runtime Environment	Runtime environment that contains the Secure Agent to run the task.

3. Click **Next**.

Configuring the source

1. On the **Source** page, select the source connection.
To create a connection, click **New**. To edit a connection, click **View**, and the in **View Connection** dialog box, click **Edit**.

2. Select the source object.
3. For file sources, configure formatting options.
4. Configure data filters.
5. Configure sort conditions.
Not all connectors support sorting. The **Sort** area appears if the source connector supports sorting.
6. If preview data does not appear automatically, expand the **Data Preview** area to preview the source data.
7. Click **Next**.

Configuring a second source

Optionally, configure a second source to use as a lookup source.

1. On the **Second Source** page, select **Yes** to add a second source to the task.
If you do not want to configure a second source, select **No**.
2. If you add a second source, perform the following steps to configure the source:
 - a. In the **Source Details** area, select **Augment Data with Lookup**.
 - b. Select the source connection and source object.
To create a connection, click **New**. To edit a connection, click **View**, and in the **View Connection** dialog box, click **Edit**.
 - c. For file sources, configure formatting options.
 - d. If preview data does not appear automatically, expand the **Data Preview** area to preview the source data.
 - e. Configure one or more lookup conditions.
 - f. Optionally, configure data filters for the combined sources.
3. Click **Next**.

Configuring the target

You can write data to a single target. You can select an existing target object or create a new object at run time. If you create a target at run time, the task operation is Insert.

1. On the **Target** page, configure the following properties as required:

Property	Description
Connection	Select a connection. To create a connection, click New . To edit a connection, click View , and then in View Connection dialog box, click Edit . To configure advanced properties for the connection, click Advanced Options . Not available for all connections types. For information about a particular connector's properties, see the help for the appropriate connector.
Object	Select a target object. Click Select . The Select Target Object dialog box displays up to 200 objects. If the object you want to use does not display, enter a search string to reduce the number of objects that display.
Target operation	Select one of the following task operations: <ul style="list-style-type: none">- Insert- Update- Upsert- Delete- Data Driven The operations available depend on the target connection.
Formatting Options	For flat file connections only. Select a delimiter and text qualifier. Optionally, select an escape character. If you choose Other for the delimiter, the delimiter cannot be an alphanumeric character or a double quotation mark.
Truncate Target	Database targets with the Insert task operation only. Truncates a database target table before inserting new rows. <ul style="list-style-type: none">- True. Truncates the target table before inserting all rows.- False. Inserts new rows without truncating the target table. Default is False.
Enable Target Bulk Load	Select this option to write data in bulk mode. The default value is false.

2. For database targets, configure update columns if necessary.
3. If preview data does not appear automatically, expand the Data Preview area to preview the target data.
4. Click **Next**.

Configuring the field mapping

Configure field mappings to define the data that the data transfer task writes to the target. Configure field mappings on the **Field Mapping** page.

1. To match fields with the same name, click **Automap > Exact Field Name**. Or, to match fields with similar names, click **Automap > Smart Map**.

You can also select and drag the source fields to the applicable target fields.

2. To configure the field expression, click the mapped field. In the **Field Expression** window, enter the expression you want to use and click **OK**.
3. If you create a target at run time and want to add target fields, click **Add**. Configure the following field properties:

Property	Description
Name	Name of the field.
Type	The data type of the data in the column.
Precision	Total number of digits in a number. For example, the number 123.45 has a precision of 5. The precision must be greater than or equal to 1.
Scale	The number of digits to the right of the decimal point of a number. For example, the number 123.45 has a scale of 2. Scale must be greater than or equal to 0. The scale of a number must be less than its precision. The maximum scale for a numeric data type is 65535.

4. Click **Next**.

Configuring runtime options

Configure runtime options on the **Runtime Options** page.

You can run a data transfer task manually or on a schedule.

1. To run a data transfer task on a schedule, select **Run on a schedule** and then select the schedule.
Note: You must create the schedule in Administrator before you can select it in the task.
2. Configure email notification options for the task.
3. Click **Save**.

Running a data transfer task

You can run a data transfer task in one of the following ways:

- **Manually.** To run a data transfer task manually, on the **Explore** page, navigate to the task. In the row that contains the task, click **Actions** and select **Run**.
You can also run a data transfer task manually from the **Task Details** page. To access the **Task Details** page, click **Actions** and select **View**.
- **On a schedule.** To run a data transfer task on a schedule, edit the task in the data transfer task wizard to associate the task with a schedule.

CHAPTER 4

Data loader tasks

Create a data loader task to load large data sets to a cloud data warehouse quickly and securely.

To start creating a task, click **New** in navigation menu on the left or **Get Started** on the **Welcome** page.

You can create a data loader task in just a few steps:

1. Connect to your source.

Choose the location you want to read data from. You can refine your source data by choosing the objects to read and the fields and records you want to process. You can also select primary key fields and the fields which indicate whether data has changed since the last task run.

2. Connect to your target.

Choose the location you want to load data to. You can also decide whether to keep or re-create the target objects when you re-run the task.

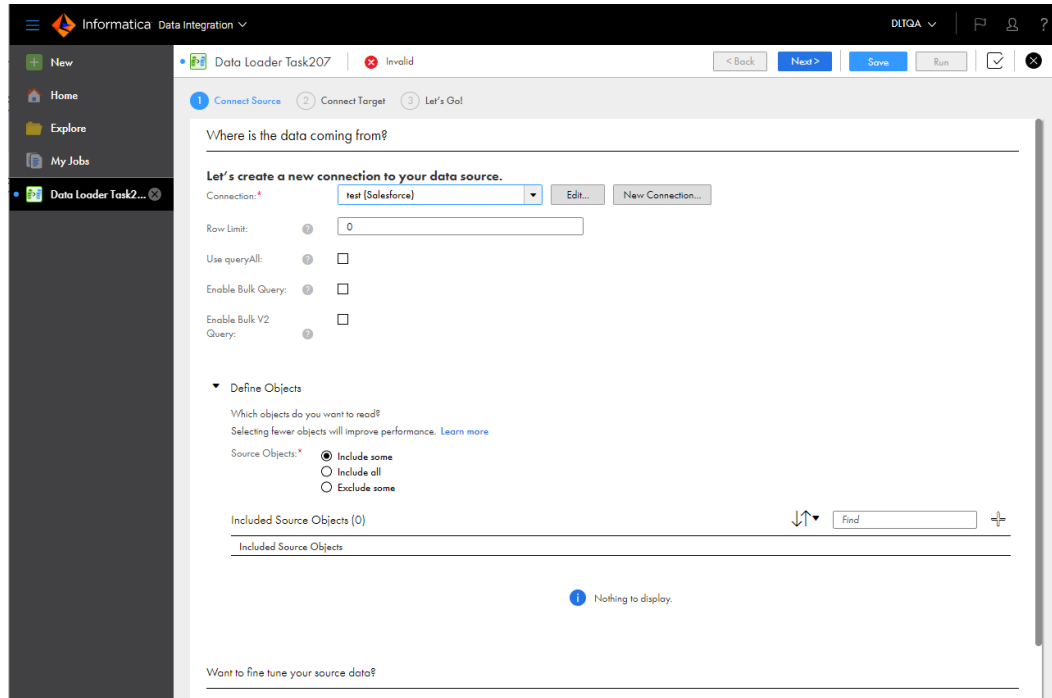
3. Save and run the task.

You can run your task when you finish creating it, save it and run it later, or run it on a schedule. You can also choose whether to send notification emails related to the task run.

Step 1. Connect to your source

Configure a source connection to connect to your source. Configure the source connection on the **Connect Source** page. You can create a new connection or select an existing connection.

The following image shows the **Connect Source** page:



To create a source connection, click **New Connection**. You can also select an existing connection. After you choose the connection, you might have to configure additional properties like **Row Limit** and **Use queryAll** that vary based on the connection type. For more information about connection properties, see *Connections*.

By default, Data Integration adds only the source objects you want to include. You can configure this feature using options in the Source Objects field.

For some connection types, Data Integration automatically detects primary key fields and watermark fields. Watermark fields identify the records that were added or changed since the last task run.

Click **Next** to customize the target connection.

You can customize the following source options:

- Which source objects to read
- The source format (file sources only)
- Which source fields to exclude
- Which records to include
- Primary key fields
- Watermark fields

For guidelines about customizing the source, see [“Guidelines for customizing your source” on page 58](#).

Guidelines for customizing your source

To increase performance by removing unnecessary data, by default Data Integration adds only the source objects you include in the source location.

Use the following guidelines when customizing your source:

Remove unnecessary data from the data flow.

When you create a task, it's best to remove unnecessary source objects, fields, and records from the data flow. Removing unnecessary data decreases the time it takes to run a task and helps to minimize rejected records.

You can choose which source objects to read on the **Connect Source** page. You can also select the fields to exclude and configure filters to exclude unnecessary records.

To prevent duplicate target rows, configure primary key fields.

Primary key fields uniquely identify records in the source and target objects. When you re-run a task, the task uses the primary key fields so that it can update existing rows and insert new rows into the target tables. If the source objects don't have primary key fields defined, the task inserts rows into the target tables, but it cannot update existing rows, which can lead to duplicate rows in the target tables.

Data loader tasks can automatically detect primary key fields for most connection types. You can also select the primary key fields manually. Configure primary key field options on the **Connect Source** page.

To process only new and changed data, configure watermark fields.

Watermark fields are date/time or numeric fields that identify which records were added or changed. If the source objects don't have watermark fields defined, the task must process all records in the source objects each time the task runs, which increases the task processing time.

Data loader tasks can automatically detect watermark fields for most connection types. You can also select the watermark fields manually. Configure watermark field options on the **Connect Source** page.

The following table lists the possible primary key and watermark field configurations and the expected results:

Primary key and watermark field configuration	Result
Primary key fields configured. Watermark fields configured.	Changed records updated, new records inserted into the target tables (upsert). Recommended configuration for best performance.
Primary key fields configured. Watermark fields not required.	Changed records updated, new records inserted into the target tables (upsert). Impacts task performance because the task must perform a full scan on the source.
Primary key fields not required. Watermark fields configured.	All records inserted into the target tables. Records that already exist are duplicated.
Primary key fields not required. Watermark fields not required.	All records inserted into the target tables. Records that already exist are duplicated. Impacts task performance because the task must perform a full scan on the source. This is the least recommended configuration.

Selecting source objects

After you configure a new connection or select an existing connection, you select the objects to read.

To load the data as quickly as possible, configure the task to read only the source objects that you need to process. Reading data from fewer source objects decreases the time it takes to run the task.

You can configure the objects to read on the **Connect Source** page under **Define Objects**.

By default, the task reads data only from the objects you include in the source location. If you want to read data from most objects in the source location, choose **Exclude some**, and then select the objects that you don't want to read. If you want to read data from all objects, choose **Include all**. If you choose **Include all** a data loader task can read up to 2000 objects.

If you choose to read all objects or exclude some objects, you'll need to enter the source path for some source types. If you entered a source path when you created the connection, the default source path is the same as the one you entered when you created the connection. Otherwise, the source path is empty by default. If you plan to read all objects or exclude objects in different locations, enter the path to the parent container.

To configure the objects to exclude or include, click the plus sign (+) icon in the **Excluded Source Objects** or **Included Source Objects** area. After you select objects, the **Excluded Source Objects** or **Included Source Objects** area displays the objects you excluded or included. To delete an object, click the **Delete** icon in the row that contains the object.

Source path syntax

The source path syntax varies based on the connection type. To find the source path syntax, click **Edit**.

For most source types, you enter the source path in the format `<container or bucket name>/<folder name>/<subfolder name>` or `<database name>/<schema name>`. If you enter a source path for an Amazon S3, Azure Blob Storage, or Azure Data Lake Storage Gen 2 source, you'll need to append a slash character at the end of each folder name to distinguish the folder from a file. So, for these sources, the source path format is `<container or bucket name>/<folder name>/<subfolder name>/`, for example, `bucket1/folder1//folder2//folder3/`.

For more information about configuring the source path for a certain connection type, see the help for the appropriate connector.

Defining the source format

If you connect to a file source, you'll need to configure formatting options such as the file format, schema source, or delimiter.

Configure the formatting options on the **Connect Source** page under **Define Source Format**. If you don't see the **Define Source Format** area, then you don't have to configure formatting options for your source.

When you configure formatting options, you can preview the data to ensure that the formatting options are correct. You can select the fields to use in the data preview.

Selecting source fields

By default, a data loader task reads the fields only from the sources you include and loads them to each corresponding target. You can exclude unnecessary fields for one or more source objects so that they are not written to the target. Excluding unnecessary fields from the data flow helps to increase task performance.

Configure the fields to exclude on the **Connect Source** page under **Exclude Fields**. To exclude fields, click the plus sign (+) icon in the **Excluded Fields** area, and then select the source object and the fields to exclude.

The **Excluded Fields** area displays the fields you excluded for each source object. To update the excluded fields for a source object, click the excluded fields for that source object. To delete all excluded fields for a source object, click the **Delete** icon in the row that contains the source object.

Configuring filters

You can configure filters for one or more source objects so that only records that match the filter conditions are written to the target. The task processes the records that match all of the filter conditions.

Configure filters on the **Connect Source** page under **Define Filters**. To add filters, click + in the **Filter Conditions** area. In the **Configure Filters** dialog box, select the source object and configure one or more filter conditions.

This image shows the **Configure Filter** dialog box:

Specify the filter condition. Only rows that match this condition will be transferred. When multiple conditions are specified, they will be ANDed.

Source Object:

Filter Conditions (2) +

Field	Operator	Value
CompanyName	=	Informatica
Country	=	USA

? OK Cancel

To add a filter condition, click the plus sign (+) icon. Each filter condition contains a field, operator, and value. You can use the following operators in a filter condition:

= (equals)

- < (less than)
- > (greater than)
- <= (less than or equal to)
- >= (greater than or equal to)
- != (not equals)

Enter date values in the format YYYY-MM-DD HH24:MI:SS.

If you add multiple filter conditions, the task evaluates the filter conditions in the order in which they appear in the **Configure Filters** dialog box.

For example, you want to extract records from the Orders table only when the city is New York and the order date is later than January 1, 1990. Select "Orders" as the source object, and configure the following filter conditions:

```
City = New York
OrderDate > 1990-01-01 00:00:00
```

The **Filter Conditions** area displays the filters you added for each source object. To update the filters for a source object, click the filters in the row that contains the source object. To delete all filters for a source object, click the **Delete** icon in the row that contains the source object.

Selecting primary key fields

Data loader tasks can automatically detect primary key fields for many source types. You can also choose your own primary key fields, or set them as not required.

Primary key fields uniquely identify each record in a source object. For example, in a CUSTOMERS table, the values in the CUSTOMER_ID column uniquely identify each customer. A source object can have one primary key field, a composite key that consists of several primary key fields, or no primary key fields.

Primary key fields are needed if you want to update rows when you re-run the task. If there are no primary key fields, the task inserts records into the target tables, but it cannot update rows.

Select primary key fields on the **Connect Source** page under **Define Primary Keys**.

You can select primary key fields in the following ways:

Automatically detect from source

The task identifies the primary key fields automatically in each source object that has them. This is the default option for many types of sources. This option is not available for file-based sources.

If you choose this option but there are no primary key fields in the source objects or the task can't detect the primary key fields, then the task inserts records into the target tables.

To find out whether a data loader task can detect primary key fields for a specific source type, see *Connections*.

Primary key fields not required

Select this option if you want to perform a full load every time the task runs. For example, you want to read data from database tables that have primary key fields, but you don't want to update rows in the target when you re-run the task.

When you select this option, the task inserts records into the target tables.

Enter manually

Select this option to choose the primary key fields in each source object. If you select this option, you must also specify whether the source objects have the same primary key fields or different primary key fields.

If the source objects have the same primary key fields, select **The same across all sources**. Then, click **Choose** to select the primary key fields.

If the source objects have different primary key fields, select **Different across sources**. Then, click the plus sign (+) icon in the **Primary Key Fields** area and select the source object and the primary key fields. Repeat this process to select primary key fields for other source objects.

The **Primary Key Fields** area displays the primary key fields for each source object. To update the primary key fields for a source object, click the primary key fields for that source object. To delete all primary key fields for a source object, click the **Delete** icon in the row that contains the source object.

Selecting watermark fields

Data loader tasks can automatically detect watermark fields for many source types. You can also choose your own watermark fields, or set them as not required.

Watermark fields are date/time or numeric fields that identify the records that were added or changed since the last task run. For example, you might select the MODIFIED_DATE column as the watermark field for a source table. A source object can have one or no watermark fields.

If you want the task to load only new and changed data, you'll need to make sure that watermark fields are configured. If there are no watermark fields, the task processes all fields in the source objects.

Select watermark fields on the **Connect Source** page under **Define Watermark Fields**.

You can select watermark fields in the following ways:

Automatically detect from source

The task identifies the watermark field automatically in each source object that has one. This is the default option.

If you choose this option but there are no watermark fields in the source objects or the task can't detect the watermark fields, then the task processes all records each time the task runs.

To find out whether a data loader task can detect watermark fields for a specific source type, see *Connections*.

Watermark field not required

Select this option if you want the task to processes all records each time the task runs.

Enter manually

Select this option to choose the watermark field in each source object. If you select this option, you must also specify whether the source objects have the same watermark field or different watermark fields.

If the source objects have the same watermark field, select **The same across all sources**. Then, click **Choose** to select the watermark field.

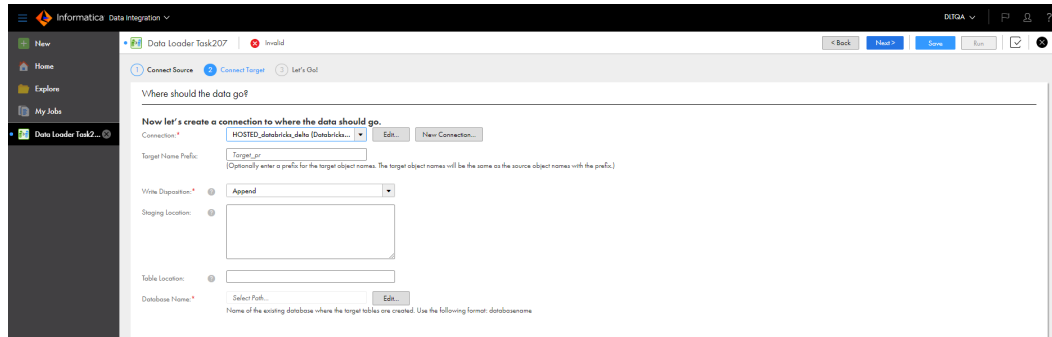
If the source objects have different watermark fields, select **Different across sources**. Then, click the plus sign (+) icon in the **Watermark Fields** area and select the source object and the watermark field. Repeat this process to select watermark fields for other source objects.

The **Watermark Fields** area displays the watermark field for each source object. To update the watermark field for a source object, click the watermark field for that source object. To delete the watermark field for a source object, click the **Delete** icon in the row that contains the source object.

Step 2. Connect to your target

Configure a target connection to connect to your target. Configure the target connection on the **Connect Target** page. You can create a new connection or select an existing connection.

The following image shows the **Connect Target** page:



To create a target connection, click **New Connection**. You can also select an existing connection. After you choose the connection, you might have to configure additional properties like **Write Disposition** and **Staging Location** that vary based on the connection type. For more information about connection properties, see *Connections*.

The first time you run a data loader task, the task creates one target table for each source object. By default, target table names are the same as the source object names. However, special characters in the source table names will be replaced with underscore characters. For example, if the source table name is `Orders$`, the corresponding target table will be named `Orders_`.

If you enter a target name prefix, the target object names will be the same as the source object names, but they'll be preceded with the prefix. For example, if a source table name is `Account`, and you enter the prefix `tgt_`, the target table name will be `tgt_Account`.

When you re-run a task, you can either load data to the target incrementally or drop and re-create the target tables.

Incremental loading

You can configure a data loader task to load data to the target incrementally. When the task loads data incrementally, only new and changed data is loaded to the target each time you re-run the task. Incremental loading increases task performance since fewer rows are loaded.

To configure the task to load data incrementally, select **Yes** under **Load to existing tables** on the **Connect Target** page.

When you load data incrementally, the task runs most efficiently when you configure the following options in the source:

Primary key fields are detected automatically or entered manually.

The task uses the primary key fields to identify rows when updating or inserting data into the target tables. If there are no primary key fields, the task can insert data into the target tables, but it can't update existing rows. This can create duplicate rows in the target tables.

Watermark fields are detected automatically or entered manually.

To load data incrementally, the source objects must contain watermark fields. The task uses the watermark fields to determine which records have been added or changed since the last task run. If

there are no watermark fields, the task can't determine which records are new or changed. So, it processes all records in the source objects each time the task runs.

If the source object structure changes between task runs, the data loader task detects the changes and alters the target tables to match the source objects. Therefore, if you add a field to a source object or change a field's data type, the task adds the new column or updates the column data type in the corresponding target table. The task does not delete any existing target fields, however. If you want to delete unnecessary target fields, you can either delete the fields manually or configure the task to drop and re-create the target tables.

If the name or data type of the watermark column changes between task runs, or if you change the watermark column between task runs, the task performs a full load even if you've configured the target to load to existing tables.

Dropping and re-creating target tables

You can configure a data loader task to drop the target tables and re-create them every time the task runs. When you select this option, the task deletes the target tables created by the original task run and creates new tables using the schema of the source objects.


To configure the task to drop and re-create the target tables, select **No, create new tables every time** under **Load to existing tables** on the **Connect Target** page. To ensure that the task processes all records each time it runs, you also need to select **Watermark field not required** under **Define Watermark Fields** on the **Connect Source** page.

When you select these options, the task performs a full target load each time it runs.

Step 3. Run the task

After you configure the target, you give the task a name and decide how to run it. Configure these options on the **Let's Go** page.

This image shows the **Let's Go** page:



1. Validation icon

Before you run the task, you can check it for errors using the validation panel. To open the validation panel, click the **Validation** icon. If there are validation errors, you can save the task, but you can't run it until you fix the errors. If there are no errors, you can run the task now, save it and run it later, or run it on a schedule.

To run the task now, click **Save** to save the task, and then click **Run**. To save it and run it later, click **Save**.

You can also configure the schedule, email notification options, error handling, and task location. The runtime environment is always the Informatica Cloud Hosted Agent, which is maintained by Informatica.

When you run the task, data loading occurs in parallel batches of five objects each. So, the first batch loads data from the first five source objects to the first five target objects. The next batch loads data from the second set of five source objects, and so on. You can monitor the task progress on the **My Jobs** page. For more information, see *Monitor*.

Running the task on a schedule

You can run your task on a schedule. For example, you might want to run the task every 12 hours or every Friday at noon.

Select a schedule for your task on the **Let's Go** page under **Schedule**. Click **Run on a schedule**, and then select the schedule.

If you want to create a new schedule, click **New Schedule**. To make the schedule easy to identify, give the schedule a meaningful name such as "Fridays at 12PM." For the start time, the date format is MM-DD-YYYY, and the time appears in the 12-hour format.

You can also create, view, edit, and delete schedules in the Administrator service. For more information, see *Organization Administration*.

Configuring email notifications

You can configure the task to send email notifications to users when the task completes successfully, completes with warnings, or fails. If you don't configure email notifications, the task doesn't send any.

Configure email notification options for your task on the **Let's Go** page under **Notifications**. You can send notifications to different addresses based on whether the task completed successfully, completed with warnings, or failed.

You can enter multiple email addresses in each field. If you enter multiple addresses, separate them with commas.

Handling runtime errors

If a task encounters a runtime error, you can stop the task or skip the error and continue. The default is to stop the task.

To determine how Data Loader handles runtime errors, set one of the following options:

- Stop the task.
- Skip the error object and continue the task.

Configuring the task location

You can select the project where your data loader task is stored.

Configure the task location on the **Let's Go** page under **Task Location**. By default, the task is stored in the currently open project, or, if no project is open, in the project where you saved your last task. If you haven't previously created a task, the task is stored in the Default project. To select a different location, click the project name and choose a different project.

If you're editing a task you saved before, the task location is read-only. You can move the task to a different project on the **Explore** page.

Runtime environment

The runtime environment is the execution platform for data loader tasks.

You can view the runtime environment on the **Let's Go** page under **Runtime Environment**. The runtime environment for all data loader tasks is the Informatica Cloud Hosted Agent.

The Informatica Cloud Hosted Agent connects to your sources and targets securely and does the data processing when you run a task. The hosted agent is maintained by Informatica. There is nothing to download, and there are no settings to configure.

INDEX

A

- advanced data filters
 - configuring [9, 10](#)
 - deleting [9, 10](#)
- advanced relationships
 - in mapping tasks [32](#)
- AND
 - reserved word [15](#)

C

- CHR function
 - inserting single quotation mark [13](#)
- Cloud Application Integration community
 - URL [6](#)
- Cloud Developer community
 - URL [6](#)
- comments
 - adding to field expressions [14](#)
- constants
 - definition [13](#)

D

- data filters
 - configuring [9](#)
 - configuring advanced filters [9, 10](#)
 - configuring simple filters [9](#)
 - deleting [9](#)
 - operators [10](#)
 - rules for configuring [11](#)
 - variables [11](#)
- Data Integration community
 - URL [6](#)
- data integration tasks
 - rules for flat file sources [29](#)
- data loader tasks
 - configuring filters [60](#)
 - connecting to the source [57](#)
 - connecting to the target [63](#)
 - creating [56](#)
 - email notifications [64, 65](#)
 - handling runtime errors [65](#)
 - incremental loading [63](#)
 - primary keys [61](#)
 - running [64](#)
 - running on a schedule [65](#)
 - runtime environment [66](#)
 - scheduling [64](#)
 - selecting source fields [60](#)
 - selecting source objects [59](#)
 - source customization guidelines [58](#)
 - source format [59](#)

- data loader tasks (*continued*)
 - source path syntax [59](#)
 - task location [65](#)
 - validating [64](#)
 - watermark fields [62](#)
- data transfer task
 - configuration [52](#)
 - configuring a lookup source [53](#)
 - configuring a second source [53](#)
 - configuring the source [52](#)
 - configuring the target [54](#)
 - creation [52](#)
 - description [47](#)
 - field mapping [55](#)
 - lookup condition [49](#)
 - lookup source [48](#)
 - operations [47](#)
 - running [55](#)
 - runtime options [55](#)
 - scheduling [55](#)
 - second source [48](#)
 - second source filters [50](#)
 - sort conditions [48](#)
 - source filters [48](#)
 - targets [50](#)
- data transfer task sources [48](#)
- database sources and targets
 - rules in synchronization tasks [29](#)
- Daylight Savings Time
 - schedules [25](#)
- DD_DELETE constant
 - reserved word [15](#)
- DD_INSERT constant
 - reserved word [15](#)
- DD_REJECT constant
 - reserved word [15](#)
- DD_UPDATE constant
 - reserved word [15](#)
- deleting
 - advanced data filters [9, 10](#)
 - data filters [9](#)
- dynamic schema handling [33, 36](#)

E

- email notification
 - configuring for tasks [26](#)
- email notifications
 - configuring [65](#)
- expressions
 - rules [14](#)
- :EXT reference qualifier
 - reserved word [15](#)

F

- FALSE constant
 - reserved word [15](#)
- field data type [52](#)
- field expressions
 - comments, adding [14](#)
 - components [13](#)
 - creating [12](#)
 - literals [13](#)
 - reserved words [15](#)
 - rules [14](#)
 - rules for validating [13](#)
 - syntax [13](#)
- field mapping
 - configuration [55](#)
- field metadata [33](#)
- filters
 - configuring [60](#)
- flat files
 - editing metadata in mapping tasks [33](#)
 - rules for sources in data integration tasks [29](#)
- functions
 - definition [13](#)

I

- in-out parameters
 - in mapping tasks [41](#)
- :INFA reference qualifier
 - reserved word [15](#)
- Informatica Cloud Hosted Agent [66](#)
- Informatica Global Customer Support
 - contact information [7](#)
- Informatica Intelligent Cloud Services
 - web site [6](#)
- input parameters
 - in mapping tasks [41](#)

J

- jobs
 - monitoring [28](#)
 - scheduling [23](#)

L

- \$LastRunDate
 - data filter variable [11](#)
- \$LastRunTime
 - data filter variable [11](#)
- literals
 - definition [13](#)
 - single quotation mark requirement [13](#)
 - string and numeric [13](#)
- :LKP reference qualifier
 - reserved word [15](#)

M

- maintenance outages [7](#)
- mapping tasks
 - advanced options [46](#)
 - advanced session properties for mappings [16](#)

mapping tasks (*continued*)

- configuring [36](#)
- configuring schedules, email, and advanced options [41](#)
- configuring sources [37](#)
- creating [36](#)
- editing [36](#)
- input and in-out parameters [41](#)
- joining related sources [31](#)
- overview [30](#)
- pushdown optimization [32](#)
- reprocessing incrementally-loaded files [46](#)
- running [46](#)
- Sequence Generator configuration [45](#)
- templates [30](#)
- using parameter files [22](#)
- viewing details [45](#)

Masking tasks

- using parameter files [22](#)

:MCR reference qualifier

- reserved word [15](#)

monitoring jobs [28](#)

N

- NOT
 - reserved word [15](#)
- NULL constant
 - reserved word [15](#)

O

- operators
 - definition [13](#)
- OR
 - reserved word [15](#)

P

- parameter file
 - configuring for mapping tasks [41](#)
- parameter files
 - overview [22](#)
- parameters
 - configuring [41](#)
 - user defined [22](#)
- postprocessing commands
 - SQL commands [27](#)
 - tasks [27](#)
- preprocessing and postprocessing commands
 - configuring for mapping tasks [41](#)
 - rules for operating system commands [27](#)
 - rules for SQL commands [27](#)
- preprocessing commands
 - SQL commands [27](#)
 - tasks [27](#)
- primary keys
 - configuring [61](#)
- PROC_RESULT variable
 - reserved word [15](#)
- pushdown optimization
 - mapping tasks [32](#)

Q

quotation marks
 inserting single using CHR function [13](#)

R

repeat frequency
 schedules [23](#)
reserved words
 list [15](#)
rules for sources and targets [28](#)
runtime environment
 data loader tasks [66](#)

S

schedules
 configuring [26](#)
 creating [25](#)
 Daylight Savings Time [25](#)
 repeat frequency [23](#)
 time zones [24](#)
scheduling
 mapping tasks [41](#)
 taskflows [23](#)
 tasks [23](#)
:SD reference qualifier
 reserved word [15](#)
:SEQ reference qualifier
 reserved word [15](#)
simple data filters
 configuring [9](#)
sort conditions
 data transfer task [48](#)
source filters
 data transfer tasks [48](#), [50](#)
sources
 configuring filters [60](#)
 configuring for mapping tasks [37](#)
 configuring primary keys [61](#)
 configuring watermark fields [62](#)
 customization guidelines [58](#)
 formatting objects [59](#)
 selecting fields [60](#)
 selecting objects [59](#)
 source path [59](#)
 source path syntax [59](#)
:SP reference qualifier
 reserved word [15](#)
SPOUTPUT
 reserved word [15](#)
status
 Informatica Intelligent Cloud Services [7](#)
stopping a running job [28](#)
string literals
 single quotation mark requirement [13](#)
synchronization tasks
 rules for database sources and targets [29](#)

synchronization tasks (*continued*)
 using parameter files [22](#)
syntax
 for field expressions [13](#)
system status [7](#)

T

targets
 configuring [63](#)
 configuring for mapping tasks [39](#)
 dropping and re-creating [64](#)
 incremental loading [63](#)
task operations
 data transfer tasks [47](#)
taskflows
 monitoring [28](#)
 scheduling [23](#)
tasks
 configuring email notification [26](#)
 configuring field expressions [12](#)
 creating schedules [25](#)
 monitoring [28](#)
 scheduling [23](#)
 stopping a running job [28](#)
:TD reference qualifier
 reserved word [15](#)
time zones
 description [24](#)
TRUE constant
 reserved word [15](#)
trust site
 description [7](#)

U

update columns
 data transfer task [50](#)
upgrade notifications [7](#)
user parameters [22](#)

V

validation
 rules [13](#)
variables
 for data filters [11](#)

W

watermark fields
 configuring [62](#)
web site [6](#)
WORKFLOWSTARTTIME variable
 reserved word [15](#)