Oracle® Fusion Middleware

Using Oracle GoldenGate Veridata

12*c* (12.2.1.2.0)
E78472-03
December 2017

Oracle Fusion Middleware Using Oracle GoldenGate Veridata, 12*c* (12.2.1.2.0)

E78472-03

Copyright © 2015, 2017, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience vi
Documentation Accessibility vi
Related Information vi
Conventions vii

1 Introduction to Oracle GoldenGate Veridata

About Oracle GoldenGate Veridata 1-1
Comparing Data Using Oracle GoldenGate Veridata 1-1
Accessing Oracle GoldenGate Veridata Web User Interface 1-2
Managing the Page View 1-2

2 Configuring Workflow Objects

Overview 2-1
Configuring Connections 2-2
- Creating a Connection 2-2
- Edit a connection 2-3
- Delete a connection 2-3
- Filter the Existing Connections list 2-3
- Editing Connection Settings 2-4
- Editing Connection Details 2-5
- Editing connection properties 2-6

Configuring Groups 2-6
- Filter the Existing Groups list 2-7
- Create a group 2-7
- Edit a group 2-7
- Delete a group 2-7

Configuring Compare Pairs 2-8
- Existing Compare Pairs List Column Details 2-9
- To filter the list 2-10
- To control delta processing 2-10

Editing a compare pair 2-10
Configuring row partitions 2-11
How to use partitions 2-12
Creating the partitions 2-12
Deleting a Compare Pair 2-14
Mapping Objects Using a Pattern 2-14
Confirm group information 2-15
Select the datasource 2-15
Select a Pattern Mapping Method 2-15
Select and Specify Compare Pair Details 2-18
Generate the mappings 2-20
Save the compare pairs to the repository 2-20
Mapping Objects Manually 2-20
Confirm group information 2-20
Select the datasource 2-21
(Optional) Filter the lists 2-21
Map a table or file 2-22
Specify compare pair details 2-22
Generate the mappings 2-23
Save the compare pairs to the repository 2-23
Reviewing and Saving Compare Pairs 2-23
Viewing details for a compare pair that is finished 2-24
Configuring Column Mappings 2-25
Choosing a column mapping method 2-26
Validating column mappings 2-27
Using Delta Processing 2-28
What is Delta Processing? 2-28
Configuring Profiles 2-34
Editing Profile Settings 2-35
Using the default profile 2-37
Configuring Jobs 2-37

3 Working with Jobs

Overview of Jobs 3-1
Estimating Comparison Time 3-1
Running A Job 3-2
Estimating Comparison Time 3-4
How groups are shown on the Run/Execute Job page 3-5
Overriding Row Partitions 3-5
Configuring run options 3-6

Parameter details 3-6
Viewing jobs that are running 3-7
Managing common tasks 3-9
Using Favorite Jobs 3-9
Using Shortcuts 3-10
Editing Jobs and Shortcuts 3-10
Editing a job 3-11
Viewing jobs that are finished 3-12
Viewing details of a job that is finished 3-14
Viewing details of a group that is finished 3-15
Viewing details for a compare pair that is finished 3-16
Using the charts 3-17
Repairing Out-Of-Sync Jobs 3-18
Viewing Details of an Out-Of-Sync Job 3-19
Viewing the Repair Report 3-20
Using the Comparison Report 3-20

4 Customizing Your Workspace

Working with the Home Page 4-1
Setting home page preferences 4-1
Using the Favorites Manager 4-2
Setting user preferences 4-4
Setting general user preferences 4-4
Setting home page preferences 4-5

5 Tuning Performance

Improving the Performance of Oracle GoldenGate Veridata 5-1
Performance Statistics 5-4

A Column Mapping

Oracle Data Type to Hive Mapping A-1
SQLServer to Hive Mapping A-1

Index

Preface

This document explains the features of the Oracle GoldenGate Veridata Web User Interface.

Audience

This document is intended for developers using Oracle GoldenGate Veridata Web User Interface.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc`.

Accessible Access to Oracle Support

Oracle customers who have purchased support have access to electronic support through My Oracle Support. For information, visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` or visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` if you are hearing impaired.

Related Information

The Oracle GoldenGate Product Documentation Libraries are found at

Oracle GoldenGate

Oracle GoldenGate Application Adapters

Oracle GoldenGate for Big Data

Oracle GoldenGate Plug-in for EMCC

Oracle GoldenGate Monitor

Oracle GoldenGate for HP NonStop (Guardian)

Oracle GoldenGate Veridata

Oracle GoldenGate Studio

Additional Oracle GoldenGate information, including best practices, articles, and solutions, is found at:

Oracle GoldenGate A-Team Chronicles

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1

Introduction to Oracle GoldenGate Veridata

Oracle GoldenGate Veridata uses a two-step process to maintain data accuracy. Compares a source (original, or trusted) set of data with a target (secondary) set of data to identify discrepancies.

- About Oracle GoldenGate Veridata
- Comparing Data Using Oracle GoldenGate Veridata
- Accessing Oracle GoldenGate Veridata Web User Interface
- Managing the Page View

About Oracle GoldenGate Veridata

Oracle GoldenGate Veridata compares a source (original, or trusted) set of data with a target (secondary) set of data to identify discrepancies. By accounting for data that is being replicated while a comparison takes place, Oracle GoldenGate Veridata can run concurrently with data transactions and replication, while still producing an accurate comparison report. You can use Oracle GoldenGate Veridata to compare data sets that reside in identical or different types of databases; the workflow is essentially the same for all supported databases, with only very minor differences.

For more information about the Oracle GoldenGate Veridata components, see Introduction to Oracle GoldenGate Veridata.

Comparing Data Using Oracle GoldenGate Veridata

To compare data accurately while transactional and replication operations are taking place, Oracle GoldenGate Veridata uses a two-step process:

In the initial comparison (or row hash) step, rows are retrieved from the source and target tables with a query. If the source and target databases are of different types, the columns are converted to a standardized data type format for accurate comparison. The rows are compressed by calculating a unique digital signature for non-key columns, and then they are compared. Rows that appear to be out-of-sync are stored in a maybe out-of-sync (MOOS) queue in memory, because at this point the comparison is inconclusive. When replication is working concurrently with a comparison, especially if there is replication latency, rows can appear to be out-of-sync when, in fact, the current data is in flight (somewhere in the replication flow) and replication will synchronize them soon.

The confirmation, or confirm-out-of-sync (COOS), step ensures accurate results by confirming the status of each row in a changing environment. In this step, rows are extracted from the MOOS queue, and their synchronization status is evaluated.

By default, confirmation processing occurs in a thread that is parallel with the initial comparison step, but the confirmation of each row waits until after a specified replication latency threshold has expired. For example, if latency is 60 seconds and the initial comparison step revealed an out-of-sync row at 9:30, then the confirmation

step for that row is not performed until 9:31, allowing replication to post any change that was in-flight.

After latency is accounted for, rows can be confirmed as persistently out-of-sync and are stored in an out-of-sync (OOS) file. This file contains all of the information that is needed to select rows for resynchronization, including the metadata for each table. Optionally, you can use the previous OOS file as input to another run to verify which rows remain out-of-sync.

Accessing Oracle GoldenGate Veridata Web User Interface

To protect data and comparison configurations, Oracle GoldenGate Veridata has security roles. Before attempting to use the Oracle GoldenGate Veridata Web User Interface, you should confirm which role has been granted to you by the Oracle GoldenGate Veridata Administrator. For more information about the security roles, see Configuring Oracle GoldenGate Veridata Security.

To connect to the Veridata Web User Interface, open a web browser and type the following address:

```
http://hostname:port/veridata
```

Where:

hostname is the name of the system where Oracle GoldenGate Veridata Server is installed and *port* is the port number where it is running (default is 8830).

Managing the Page View

Most pages of Oracle GoldenGate Veridata Web User Interface provide options that help you to control the size of lists, to navigate through lists, and to control the objects that are displayed on a page.

To filter a list

You can filter a list of objects wherever you see an expandable area like this:

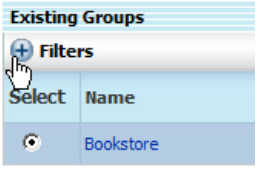

Click **Filters** to expand the options. The filter typically, but not always, contains text boxes that correspond to the columns in the list. For example, if a list includes Name and Description columns, then the filter boxes would be Name Like and Description Like. In some cases, a drop-down selection list is provided instead of a text box.

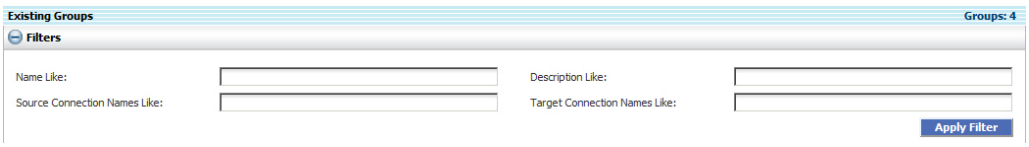

Wildcards are not necessary when using filters based on a string. Any object that contains the string is returned. For example, the string "test" might return names like "test1," "latest," "already_tested," and so forth. The search is not case-sensitive.

Click **Apply Filter** to apply a filter specification.

The total number of items that are contained in a list is displayed on the right-hand side of the list header.

To sort the list columns, click on the column header.

To constrain the number of items that appear in a list on the current page, select a value from **Show <n> items per page**:

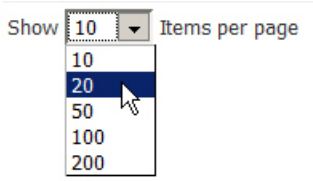

When a list spans multiple pages, use the navigation buttons to move through it:

Click the forward and backward arrows to scroll from page to page sequentially.

Click the double forward or backward arrows to scroll through pages two at a time.

Click the outermost sets of arrows to go to the first or last page in the list.

To drill down from a list

The names of jobs, groups, and compare pairs in lists are hyperlinks, even though they are not underlined. Click them as you would any link to drill down to other, related pages.

To navigate forward and backward through pages

Because of the underlying technology that is used in the Oracle GoldenGate Veridata Web User Interface, clicking the Back button in the browser does not take you back to the previous screen. You should use the menus, the navigation bar, and the breadcrumbs at the top of each page for navigation.

Other page view options

You can customize the appearance and content of some pages by setting user preferences.

2
Configuring Workflow Objects

Oracle GoldenGate Veridata supports connections, groups, compare pairs, profiles, and jobs. You can create, modify, edit, delete, and use all the features by using the User Interface.

- Overview
- Configuring Connections
- Configuring Groups
- Configuring Compare Pairs
- Configuring Column Mappings
- Configuring Profiles
- Configuring Jobs

Overview

To begin using Oracle GoldenGate Veridata, you need to create some objects that identify the data that you want to compare and which help you to manage your work. Create these objects in the following order:

1. Configure datasource connections: Oracle GoldenGate Veridata Server must be able to connect to an Oracle GoldenGate Veridata Agent for each database that contains source and target data that you want to compare. A connection is defined by a host, the port number of an Oracle GoldenGate Veridata Agent (or Manager, if a C-agent), and the datasource that is accessed by the agent. Connections must be created before any other objects are created.
2. Configure groups: You must configure at least one compare group that is linked to a set of source and target datasource connections. A group is a logical container for organizing the objects that you want to compare.
3. Configure compare pairs: You must configure one or more compare pairs for each group that you create. A compare pair is a set of corresponding source and target tables or files. Compare pairs can be created right away after you create a group, or you can edit the group later to add them. Configuring compare pairs will be the most time-consuming task, but you can spread the work across multiple sessions and save it as you go along. You can create a link to your work-in-progress in the Favorites Manager, so that you can return to it quickly.
4. Configure profiles: A profile contains settings for runtime parameters and can be applied globally to a job or to a specific compare pair as an override to the job profile. Profile parameters control such attributes as the sorting method to be used, thread and memory usage, report output, and so forth. Defining run profiles is optional, because Oracle GoldenGate Veridata includes a default profile that contains settings that apply to most usage scenarios. However, as you gain experience with Oracle GoldenGate Veridata, you may want to customize the default profile or create your own custom profiles.

5. Configure jobs: A job is a logical container for one or more compare groups and is the unit of work by which comparison processing is executed. Within one or more jobs, you can manage and run large volumes of compare groups across numerous databases and systems, and you can control the timing of those comparisons.

Configuring Connections

To get started with Oracle GoldenGate Veridata, you must define a connection to the source and target databases that contain the data that you want to compare. Oracle GoldenGate Veridata Server uses the connection information to communicate with Oracle GoldenGate Veridata Agent.

A connection is defined by:

- A host where Oracle GoldenGate Veridata Agent is running
- The port number for Oracle GoldenGate Veridata Agent on that host
- The datasource that is associated with this agent

Connections are managed from the Connection Configuration page. To access this page, click **Connection Configuration** under Configuration in the navigation pane.

All connections that exist within the Oracle GoldenGate Veridata repository are shown in the Existing Connections list on this page. You must have the Administrator or Power User role to do the following from this page:

- Creating a Connection
- Edit a connection
- Delete a connection

Other configuration tasks

- Filter the Existing Connections list
- Editing Connection Settings
- Editing Connection Details
- Editing connection properties

Creating a Connection

1. Click **New**. The New Connection Assistant prompts for the following:
   - A name for the connection, and an optional description.
   - **Host Name or IP Address**: The DNS (Domain Name Server) host name or IP address of the host where the Oracle GoldenGate Veridata Agent is installed. See the Oracle GoldenGate Veridata Agent System Requirements for installation instructions, if needed.
   - **Port**: The port number that is assigned to the agent (or the Manager process, if a C-agent). To find out the port number of a Java agent, view the `server.port` parameter in the `agent.properties` file within the agent's installation directory. To find out the port number for a C-agent Manager, run the GGSCI program from the agent's installation directory, and then use the `INFO MANAGER` command.

- **DataSource Type**: Select the type of datasource from the Datasource Type drop-down list. The list contains the following datasource types:
  - Oracle
  - NSK
  - SQL Server
  - Sybase
  - DB2
  - Informix
  - Hive
  - Teradata
- **Use SSL for communication**: Select the **Use SSL for communication** checkbox for secure communication between the Veridata agent and the server.
- A user name and password for connecting to the datasource (if required by the database).
- (Optional) A separate user can be configured for executing repair operations at the target database. This user needs permission to update as well as to query the tables.

2. Click **Test Connection** to confirm the supplied information is correct. If you have selected the Use SSL for communication checkbox in the previous screen, SSL will be used for verifying the datasource connection.

Edit a connection

1. Click the name of the connection in the Existing Connections list, or select it in the Select column and then click **Edit**.
2. The Edit Connection page is displayed.

Delete a connection

1. Before deleting a connection, you must un-link it from any groups and jobs to which it is linked, or delete the group or job if appropriate.
2. To delete a connection, select it in the Select column of the Existing Connections list, and then click the **Delete** button. Only one connection can be deleted at a time.

Filter the Existing Connections list

Expand Filters to reveal filter options. Strings are matched character for character.

- Name Like: Use a string to filter by connection name.
- Description Like: Use a string to filter by matching strings in the description.
- Datasource Type: Select from the drop-down list to filter by the datasource type.

Click **Apply Filter** to display the selected connections.

Editing Connection Settings

Use the Edit Connection page to edit connection information. To access this page:

1. In the navigation pane, click **Connection Configuration**.
2. In the Existing Connections list, click the name of the connection that you want to edit. The Edit Connection page is displayed.

You must have the Administrator or Power User ***role*** to edit a connection.

Connection Settings tab

This tab modifies the behavior of the agent.

A check mark under Use Default indicates that a parameter is set to the default value.

The current setting is shown under Value.

If the parameter is a toggle, a check mark under Value indicates that it is enabled.

To change a setting

1. Clear the Use Default box.
2. Make your change under Value.
3. Click **Save**.

Parameter Details

- **Agent Message Timeout**: Specifies a time interval, in seconds, after which the Oracle GoldenGate Veridata Server abends if it has not received a message from the Oracle GoldenGate Veridata Agent.
- **Truncate Trailing Spaces When Comparing Values**: Truncates trailing spaces when comparing string-type columns. If a column only contains spaces, then the first space is kept. Truncating trailing spaces is useful when comparing columns that have different lengths. Trailing spaces are truncated from VARCHAR columns automatically, regardless of how this parameter is set. Trailing spaces on LOB data are not trimmed.
- **Initial Compare Fetch Batch Size**: (Oracle database only) Sets the number of rows that are fetched at once for the initial comparison. Increasing the batch size may increase throughput, as compared to standard database access. The default of 0 fetches a batch size of 1000 rows. Any other value that you supply will be the actual number of rows fetched. If used, a value of at least 100 rows is recommended. Values greater than 1000 usually are not productive and waste memory.
- **Use Source or Target Columns as Key Columns When Generating Compare Pairs**: If the source table has primary keys or unique index defined, but the target doesn't, then Oracle GoldenGate Veridata uses the same columns (as the source) as unique identifiers for the target. Similarly if the source doesn't have the primary key or index, but target has them, then key columns of target are used for source. If either (source or target) of the columns of the Primary Key or index is not present on the target side, then the primary key or index isn't considered.

> **Note:**
>
> The behavior of this option in the source column differs from its behavior in the target column.

- **Use All Columns as Key Columns When Generating Compare Pairs**:

  Enables Automatic mapping for the source and target connections. If this check box is disabled either at the source or the target connection, then automatic mapping is also disabled for that group. If you enable this option to map all columns from source and target, then the mapping is considered only when both source and target table don't't have primary or unique keys.

Editing Connection Details

Use the Edit Connection page to edit connection information. The settings on this page are applied globally to all comparisons that are run from the associated data source, unless an override is permitted elsewhere.

You must have the Administrator or Power User role to edit a connection.

Connection Details tab

This tab captures the information that is required for the Oracle GoldenGate Veridata Agent to connect to a database.

Parameter Details

- **Name**: The name cannot be changed.
- **Description**: A description is optional and can be edited as needed.
- **Host Name or IP address**: Either the Domain Name Server (DNS) host name or IP address of the system. This information can be changed.
- **Port**: The port number of the Oracle GoldenGate Veridata Agent on the system. The port number can be changed if the one listed is not the correct port that was specified when the agent was installed.
- **Datasource Type**: The type of database that is accessed by this connection. This attribute cannot be changed.
- (Optional) Click **Verify** to verify that the connection will be successful. The connection will also be verified by Oracle GoldenGate Veridata at runtime.
- **Catalog**: (Does not apply to all databases) The database that contains the objects that are to be compared.
- **User**: The name of an existing user that has connection access to the database. The user can be changed. For required permissions, see Securing Access to Oracle GoldenGate Veridata by Defining User Roles.
- **Password**: The user's password.
- (Optional) If you configured a separate user for executing repair operations at the target database then you can change this user's name and password.
- (Optional) Click **Test Connection** to verify that the credentials are valid.

When you are finished making your changes, click **Save** to save them to the repository.

Editing connection properties

Use the Edit Connection page to edit connection information. To access this page:

1. In the navigation pane, click **Connection Configuration**.
2. In the Existing Connections list, click the name of the connection that you want to edit. The Edit Connection page is displayed.

You must have the Administrator or Power User ***role*** to edit a connection.

Connection Properties tab

This tab defines rules for how each data type in the underlying database is interpreted and mapped if compared to data from a different type of database. This tab sets global values for all instances of a data type. To override the format mapping of a specific column in any given table, go to the Column Mapping Configuration page for the compare pair and then use the User Defined column mapping method.

The supported data types are displayed with default mappings to Oracle GoldenGate comparison formats. In cases where the automatic mapping is not sufficient, you can select another supported format.

To change a format setting:

1. Clear the Use Default box.
2. Make a selection under Comparison Formats.
3. Depending on the data type and format that you specified for Comparison Formats, you might need to supply or select additional information in the Precision, Scale, and Timezone columns.
4. Click **Save**.

Other connection configuration tabs

Editing Connection Details

Editing Connection Settings

Configuring Groups

Groups are logical containers for one or more compare pairs. They help you to organize and partition large or diverse sets of data into more manageable units. Groups are linked to jobs when jobs are created. Any group can be linked to one or more jobs, allowing you complete control over how and when data is compared.

> **Note:**
>
> A group is associated with a set of connections to the source and target data. Before creating a group, you must create these connections.

Groups are managed from the Group Configuration page. To access this page, click **Group Configuration** in the navigation pane.

All groups that are defined within the Oracle GoldenGate Veridata repository are shown in the Existing Groups list on this page. You can do the following tasks from this page:

- Create a group
- Edit a group
- Delete a group

You must have the Administrator or Power User role to create, edit, or delete a group.

Other configuration tasks

Configuring Compare Pairs

Configuring Connections

Configuring Profiles

Configuring Jobs

Filter the Existing Groups list

Expand Filters and type a search string in one of the following boxes. Strings are matched character for character.

- **Name Like**: filters by group name.
- **Description Like**: filters based on a matching string in the description.
- **Source Connection Names Like**: filters by the name of the source connection.
- **Target Connection Names Like**: filters by the name of the target connection.

Click **Apply Filter** to display the selected groups.

Create a group

To create a group, click **New**. The New Group Assistant is displayed. The assistant will prompt for:

- A name and description.
- Connection information

Edit a group

To edit a group, click the name of the group in the Existing Groups list, or select it in the Select column and then click **Edit**. The Edit Group page is displayed.

Delete a group

Before deleting a group, you must remove it from any jobs to which it was linked (or delete the job if appropriate.) To delete a group, select it in the Select column of the Existing Groups list, and then click the **Delete** button. Only one group at a time can be deleted.

Configuring Compare Pairs

A compare pair is the logical relationship between a source table or file and a target table or file for the purpose of comparing their data. Compare pairs are linked to groups. Because of this, all of the source and target objects that you configure into compare pairs for any given group must be accessible from the datasource connections that are associated with that group.

> **Note:**
>
> Note: Before creating compare pairs, you must create a group to contain them.

Use the Compare Pair Configuration page to view, create, modify, and save compare pairs and their column mappings. To access this page:

1. Click the name of a group within any page that displays groups. To access this page from the navigation pane, click **Group Configuration**, then the name of the group for which you want to configure compare pairs. This displays the Edit Group page.
2. Click **Go to Compare Pair Configuration...**
   - The Group Information section at the top of this page displays information about the group that you are working with.
   - The tabs on this page contain the tools that you need to work with compare pairs.

The Existing Compare Pairs tab is your start point. It displays any compare pairs that currently exist in the selected group. The columns in the Existing Compare Pairs list show details about those compare pairs.

Modifying Existing Compare Pairs

To control delta processing (For more information about delta processing, see Using Delta Processing).

Editing a compare pair

Deleting a Compare Pair

To create new compare pairs

You must have the Administrator or Power User role to create compare pairs.

You can use the following methods to create compare pairs.

- **Pattern Mapping**

  Use the Pattern Mapping tab to map numerous source and target objects at once by using:

  – An exact name match (for example TAB1=TAB1).

  – A SQL percent (%) wildcard or an asterisk (*) wildcard to map differently named source and targets whose naming conventions support wildcarding.

To create compare pairs using pattern mapping see Mapping Objects Using a Pattern.

- **Manual Mapping**

  Use the Manual Mapping tab to map source objects to target objects one at a time. This method is useful when source and target names differ vastly and cannot be mapped by using wildcards.

  > **Note:**
  >
  > The manual mapping tab is also where you configure row partitions for new or existing compare pairs.

  To create compare pairs using manual mapping, see Mapping Objects Manually.

- **Automatic Key Column Mapping**

  In the absence of unique identifiers, such as the primary keys and unique indices, which can uniquely identify each row, you need to manually define any column or columns to be used as a unique identifier. These mappings are automated by options in the **Manual Mapping** tab and the **Pattern Mapping** tab. You can enable automatic key column mapping either at the Connection Level or at the Compare Pair Level. If you have enabled it at the Connection Level, then any group using this Connection has the feature enabled. You need to enable the**Use All Columns as Key Columns When Generating Compare Pairs** checkbox at both the source as well as the target levels. If the checkbox is disabled at any one level, then the automatic key column mapping is disabled for the corresponding group. See Editing Connection Settings.

- **Combination Mapping**

  You can use a combination of both mapping methods, if needed. Use the pattern method first, and then use the manual method for any objects that remain unmapped, or to make revisions such as to add row partitions.

Existing Compare Pairs List Column Details

- Compare Pair Name: This is either the default name format of `<source>=<target>` or a user-defined name. In either case, if you hover the mouse cursor over a compare pair name, the actual source and target object names are displayed.
- Column Mapping: Links to the Column Mapping Configuration page, where you can create or edit the mappings between the source and target columns of a compare pair.
- Source Catalog (Some platforms): The source metadata catalog or database.
- Source Schema: The owner of the source database objects that are to be compared.
- Target Catalog (Some platforms): The target metadata catalog or database.
- Target Schema: The owner of the target database objects that are to be compared.
- Row Partition: If any row partitions (subsets) are defined, there is a notification here. Otherwise, this field is blank.

- Profile: If a run profile exists for a compare pair, it is shown here. Otherwise, this field is blank and the default profile will be used during comparisons. (A profile can be specified for an individual compare pair on the Manual Mapping tab.)
- Validation Status: Shows whether or not the columns of the source and target objects are suitable for being compared, based on the results of any previous validation that was performed.
- Key Mapping Method: Can be either System Generated or User Defined, depending on the method that was chosen to map the key columns for this compare pair.
- Column Mapping When No Keys are Specified: Can be either System Generated or User Defined, depending on the method that was chosen to map the non-key columns for this compare pair.

To filter the list

1. Expand Filters to expose filter options.
   - Compare Pairs with Status filters on one of the following:
     - Validated means that the source and target columns are compatible and suitable for comparison.
     - Preliminary Validation Failed means that a preliminary validation (done from a configuration page and not by a runtime process) failed.
     - Runtime Validation Failed means that the compare pair failed the runtime validation when a job was started.
   - The remaining options filter based on a string that matches a string in:
     - a compare pair name
     - a table or file name
     - a profile name
2. Click **Apply Filter** to display the selected compare pairs.

To control delta processing

On any platform, you can enable or disable delta processing for any or all compare pairs from the Existing Compare Pairs tab. Select their names in the Select column (or select Page to select all pairs on the current page), and then click **Enable Delta Processing** or **Disable Delta Processing**. When delta processing is enabled, a yellow triangle delta symbol appears next to the name of the affected compare pair whenever the pair is displayed. For more information about delta processing, see Using Delta Processing.

Editing a compare pair

Use the Compare Pair Configuration page to edit compare pairs. To access this page:

1. Click the name of a group on any page that displays groups. This displays the Edit Group page.
2. In the Existing Compare Pairs list, click the name of the compare pair that you want to edit. This activates the Manual Mapping tab of the Compare Pair Configuration page.

You also can go directly to this page by clicking the name of any compare pair on any page that displays them.

You must have the Administrator or Power User role to edit compare pairs.

To edit compare pair properties

You can change any of the following. After you make a change, click **Save**.

- To change the name of the compare pair

  Type the new name in the Compare Pair Name field. The name cannot contain spaces but can contain underscores and equal signs. It is case-sensitive.

- To control delta processing

  From the drop-down list, select either **Enabled** or **Disabled**. To use delta processing, server-side sorting must be enabled.

To change the profile that is associated with this compare pair

In the Profile field, select the name of the profile. Alternatively, you can click the Browse button next to Profile to browse for a profile using the Select a Profile page, which contains filtering options.

To add or change a row partition

Next to Row Partitions, click **Configure**.

To edit the column mapping:

1. Make the Existing Compare Pairs tab active.
2. In the Column Mapping column, click **Edit** next to the compare pair whose column mapping you want to change.
3. Make your changes on the Column Mapping Configuration page.

Configuring row partitions

Oracle GoldenGate Veridata supports selecting a subset, or partition, of rows for comparison by means of a SQL predicate statement or Enscribe partition range. Using partitions allows you to compare source and target tables or files that have the same structure but a different number of rows. For example, you could compare a production table to a data warehouse table that may contain more rows because of historical data. Using row partitions also speeds throughput by splitting the load into multiple processing streams.

As an example of how partitions can be used, suppose a target SQL table FIN2 has an ID column that contains values 1 through 2000 (so 2,000 rows), but the source table FIN only has 1,000 rows with ID values 1 through 1000. A possible configuration for the target table row partition could be:

Compare Pair Name	Source Table/Target Table	Partition criteria
FinPart1	Fin/Fin2	ID < 500
FinPart2	Fin/Fin2	ID between 500 and 1000

How to use partitions

1. For any given job run, one compare pair configuration can support:
   - One row partition for the source table or file, and one row partition for the target table or file.

     or
   - One row partition for either of the source objects or the target objects.
2. You can add any number of row partitions to a compare pair, but only one can be active at runtime.

To compare multiple partitions concurrently

To compare multiple row partitions for the same source and target objects during the same job run, create a new compare pair for each partition and include all of those compare pairs in the run.

Once partitions are defined by means of unique compare pairs, multiple independent comparisons are possible in parallel or over the course of time, for example one row partition per night.

Creating the partitions

Partitions are created in the Compare Pair Row Partitions editor. The appearance of this editor changes slightly depending on whether you are configuring a table-based or file-based compare pair.

To open the row partition editor

1. Navigate to the Manual Mapping tab of the Compare Pair Configuration page.
2. Go to the Row Partitions field.
3. Click the **Configure** button to display the Compare Pair Row Partitions editor.
4. Select one of the following tasks to work with the partitions.
5. Click **OK** to close the Compare Pair Row Partitions editor. The Row Partitions box of the Compare Pair Configuration page now contains the word "Specified."

To create a row partition from a copy

You can create partitions for one source or target object, and then apply those same partition(s) to the corresponding object with one mouse click. To allow for minor differences in the copied partition, you can edit it as needed.

1. Create the initial row partition for the source or target object.
2. Click **Copy All From Target** or **Copy All From Source**, depending on where you created the initial partition. The copied partition inherits the Use At Runtime setting of the original partition.
3. Click **OK** to close the Compare Pair Row Partitions editor. The Row Partitions box of the Compare Pair Configuration page now contains the word "Specified."

To create a new row partition

1. Under Source or Target, click **New**. The editor expands to display the Row Partition Configuration area.

> **Note:**
>
> You can cancel your work and start over at any time by clicking the Cancel button that is within this area of the editor. To exit the editor completely, click the **Cancel** button that is at the bottom-most portion of the editor page.

2. In the **Name** box, type a name for this partition. Use one word that can include underscores, hyphens, and other standard keyboard special characters.
3. Do one of the following:
4. Click **Use At Runtime** to make this statement the active statement when the next comparison is run.
5. Click **Apply**. The name of the new partition appears in the Row Partitions list above the work area, and the partition criteria is displayed.
6. Repeat these steps for any additional partitions that you want to create.
7. Click **OK** to close the Compare Pair Row Partitions editor. The Row Partitions box of the Compare Pair Configuration page now contains the word "Specified."

To edit a row partition

1. In the Select column under Source or Target, select the partition that you want to change.
2. Click **Edit**.
3. Under Row Partition Configuration, change the name or partition information as needed.

> **Note:**
>
> You can cancel your work and start over at any time by clicking the **Cancel** button that is within this area of the editor. To exit the editor completely, click the **Cancel** button that is at the bottom-most portion of the editor page.

4. (Optional) Click **Use At Runtime** to make this partition the active one for the compare pair.
5. Click **Apply**.
6. Click **OK** to save the partition statement and return to the Compare Pair Configuration page.

To delete a row partition:

1. In the Select column, select the partition that you want to delete.

2. Click **Remove**.
3. Remove other partitions as needed.
4. Click **OK** to return to the Compare Pair Configuration page.

Deleting a Compare Pair

Use the Compare Pair Configuration page to delete compare pairs.

You must have the Administrator or Power User role to delete compare pairs.

To delete compare pairs

1. Make the Existing Compare Pairs tab active.
2. (Optional) Expand Filters to filter the Existing Compare Pairs list as needed, and then click **Apply Filter**. You can filter by:
   - Compare Pairs with status
     - **Validated** means that the source and target columns are compatible and suitable for comparison.
     - **Preliminary Validation Failed** means that the compare pair has not been processed in a job yet, but a preliminary validation failed.
     - **Runtime Validation Failed** means that the compare pair failed the runtime validation when a job was started.

     More on validation...
   - Compare Pair Like

     The default name format is <*source_table*/*file*>=<*target_table*/*file*> but a compare pair can have a user-defined name instead. The filter returns names containing the specified string.
   - Source Table Name Like and Target Table Name Like

     To determine the source and target table or file names if a user-defined name is shown, hover the mouse cursor over the name of the compare pair in the list. A hint appears showing the actual names.
   - Profile Name Like

     There will be a profile name shown if one was specified to override the default profiles for jobs that contain this compare pair.
3. In the Select column, select the compare pairs that you want to delete, or select the Page check box to delete all of the compare pairs on the page. Use the page management tools to make selections on other list pages as needed.
4. Click **Delete**. You will be prompted to confirm your request.

Mapping Objects Using a Pattern

If your source and target object names are identical, or the naming conventions support using wildcards, you can create compare pairs by using the pattern mapping method. Pattern mapping can save a great deal of time compared to manual mapping, because it generates multiple compare pairs at once. You can combine the pattern mapping and manual mapping methods as needed.

To use pattern mapping, use the Pattern Mapping tab of the Compare Pair Configuration page.

You must have the Administrator or Power User role to configure compare pairs.

To use pattern mapping, complete the six tasks in the following sections in order.

Confirm group information

Under Group Information, verify that the group you have chosen to contain the compare pairs is displayed at the top of the page.

Select the datasource

Specify the database structures that contain the tables to be mapped. The selections vary by database type:

Windows, UNIX, and Linux-based databases

Under Datasource Information, select the Source and Target catalogs (if applicable to the database) and schemas that contain the tables that you want to map as compare pairs. You can select from a drop-down list or click the **Browse** button to open a selection dialog that has filtering options. To use the filtering options, do the following:

1. Supply a search string to filter the list, and then click **Apply Filter**. The filter returns a list of names.
2. In the Select column of the list, click the name of the appropriate schema or catalog for this compare pair.
3. Click the **Select** button.

Enscribe databases

1. In the File Pattern box under Datasource Information, type the Source and Target volume, sub-volume, and file patterns of the source and target files that you want to map as compare pairs. Use the format of volume.subvolume.file_pattern (for example, volume.subvolume.*).
2. Click the **Refresh Source Files/Tables and Refresh Target Files/Tables** buttons.

Select a Pattern Mapping Method

Under Pattern Compare Pair Mapping, select one of the following methods for matching source table names to target table names:

- **Table Mapping Method**
  - **Map Source and Target Tables Using Exact Name**s.

    This pattern matches names character-for-character, so each source and target name must be identical. This pattern is useful for comparing production and failover databases, for example.
  - **Map Source and Target Tables Using SQL % Wildcard Pattern**...

    To use this method, supply a wildcard string in Where Source Names Like and Where Target Names Like that includes the percent symbol (%) as the

wildcard. A % in the target matches the text that is matched by the % in the source.Example:

Assume source tables of:

```
SOURCE_TABLE_1
SOURCE_TABLE_2
MY_SOURCE_TABLE_1
MY_SOURCE_TABLE_2
DUMMY_TABLE
```

Assume target tables of:

```
TARGET_TABLE_1
TARGET_TABLE_2
MY_TARGET_TABLE_1
MY_TARGET_TABLE_2
DUMMY_TABLE
MY_DUMMY_TABLE
```

Some possible pattern matches are:

Source pattern: `SOURCE_%`

Target pattern: `TARGET_%`

Matches: `SOURCE_TABLE1=TARGET_TABLE1, SOURCE_TABLE2=TARGET_TABLE2`

Source pattern: `MY_SOURCE_%`

Target pattern: `MY_TARGET_%`

Matches: `MY_SOURCE_TABLE1=MY_TARGET_TABLE1, MY_SOURCE_TABLE2=MYTARGET_TABLE2`

Source pattern: `SOURCE_%`

Target pattern: `MY_TARGET_%`

Matches: `SOURCE_TABLE1=MY_TARGET_TABLE1, SOURCE_TABLE2=MYTARGET_TABLE2`

Source pattern: `DUMMY_TABLE`

Target pattern: `%_TABLE`

Matches: `DUMMY_TABLE=DUMMY_TABLE`

Source pattern: `DUMMY_TABLE`

Target pattern: `TARGET_%_`

Matches: None

To preview the objects that are selected with the wildcards, click the **Preview** link that is shown after Map Source and Target Tables Using SQL % Wildcard Pattern. This opens the Pattern Mapping Preview page.

– **Map Source and Target Tables Using Oracle GoldenGate * Wildcard Pattern**...

```
SOURCE_TABLE_1
SOURCE_TABLE_2
MY_SOURCE_TABLE_1
MY_SOURCE_TABLE_2
DUMMY_TABLE
```

Assume target tables of:

```
TARGET_TABLE_1
TARGET_TABLE_2
MY_TARGET_TABLE_1
MY_TARGET_TABLE_2
DUMMY_TABLE
MY_DUMMY_TABLE
```

Possible Oracle GoldenGate wildcard matches are:

Source pattern: `SOURCE_*`

Target pattern: `TARGET_*`

Matches: None

Source pattern: `*`

Target pattern: `MY_*`

Matches: `DUMMY_TABLE=MY_DUMMY_TABLE`

Source pattern: `DUMMY_TABLE`

Target pattern: `*`

Matches: `DUMMY_TABLE=DUMMY_TABLE`

Source pattern: `*`

Target pattern: `*`

Matches: `DUMMY_TABLE=DUMMY_TABLE`

The wildcard resolution is case-sensitive. To preview the objects that are selected with the wildcards, click the **Preview** link that is shown after Map Source and Target Tables Using GoldenGate * Wildcard Pattern. This opens the Pattern Mapping Preview page.

- **Column Mapping When No Keys are Specified**: You can enable or disable automatic mapping either at the Connection Level (**Connection Settings** tab) or at the Compare Pair level (**Compare Pair Configuration** page). If this mapping is enabled at the Connection Level, then any group using this Connection will have the feature enabled. If the feature is disabled at the Connection level, then it also be disabled at Compare Pair level. You can override automatic mapping at the Compare Pair Generation stage. See Automatic Mapping in Configuring Compare Pairs.
  - **Use Source Key Columns for Target**: Enables automatic mapping for the target table. This option is available on the **Compare Pair Generation** page. If you disable the option for the target side, and if there is no primary key or unique index on the target side and the source has primary key or unique index, then only the mapping of columns happen but, automatic mapping doesn't happen on the target table.
  - **Use Target Key Columns for Source**:

    Enables automatic mapping for the source table. This option is available on the **Compare Pair Generation** page. If you disable the option for Source side, and if there is no primary key or unique index on the source side and target has primary key or unique index, then only the mapping of columns happen but, automatic mapping doesn't happen on the source table.

- **Use All Columns as Key Columns When Generating Compare Pairs**: Enables automatic mapping for all columns. Ensure that this option is selected both at the source as well as the target groups. If the checkbox is disabled either at the source or the target levels, then the automatic mapping is disabled for the corresponding groups.

Select and Specify Compare Pair Details

You can enable or disable automatic mapping either at the Connection Level (**Connection Settings** tab) or at the Compare Pair level (**Compare Pair Configuration** page). If this mapping is enabled at the Connection Level, then any group using this Connection will have the feature enabled. If the feature is disabled at the Connection level, then it also be disabled at Compare Pair level. You can override automatic mapping at the Compare Pair Generation stage.

This topic contains the details of the configurations in the **Manual Mapping** tab.

Selecting a Compare Pair

In the **Selected Compare Pair** section, specify the Source **Table Name** and the Target **Table Name** in the respective fields.

- In the **Column Mapping When No Keys are Specified** section, you can select the following:
  - **Use Source Key Columns for Target**: Enables automatic mapping for the target table. If you enable this option and if the source table contains primary key, but the target table doesn't, then, auto mapping of key columns is done based on source table primary keys.
  - **Use Target Key Columns for Source**:

    Enables automatic mapping for the source table. If you enable this option and if the target table has primary keys, but the source table doesn't, then, auto mapping of key columns is done based on the target table primary keys.
  - **Use All Columns as Key Columns**: Enables automatic mapping for all groups. Ensure that this option is selected both at the source as well as the target group. If you have enabled this option, then all columns from source and target tables get mapped only when both source and target tables don't have primary keys

See Automatic Mapping in Configuring Compare Pairs.

Under Compare Pair Details, specify the following:

Compare Pair Naming Format

Specifies the format for naming compare pairs. The default is *=*. This format uses an equal sign to link source and target object names. Example default names: `customers=customers2`, `\FIN.$DATA.FN.CUST1=\FIN.$DATA.FN.CUST1`. Case sensitivity is observed. Oracle GoldenGate Veridata first looks for a case-sensitive name match. Failing that, it looks for a case-insensitive name match. Objects that fail on both counts are not mapped.

You can change the naming format if desired.

To specify a name:

- A user-defined name cannot contain spaces but can contain underscores and equal signs. It is case-sensitive. If you will be generating a large number of compare pairs, use a naming convention that is intuitive. Oracle GoldenGate Veridata will display the actual object names if you hover the mouse over the name of a compare pair.

> **Tip:**
>
> To condense the width of columns that display the names of compare pairs, you can use one asterisk as the name format if the source and target objects have identical names.

The result is that just one name is displayed under Compare Pair Name columns throughout the application, instead of <*name*>=<*name*>. This is especially useful for long Enscribe names.

Profile

You can leave this field blank (to use the default job profile) or you can specify a profile for this compare pair, which overrides the job profile.

To select a profile:

- Select the profile from the drop-down list or click the **Browse** button to open the Select a Profile page, which has filtering options.
- To use Select a Profile
  - Supply a search string to filter the list by profile name, and then click **Apply Filter**.
  - Select the correct profile from the list, and then click **Select**.

Delta Processing

For more information about delta processing, see Using Delta Processing.

Select Enabled to use the delta processing feature. By default it is disabled. To use delta processing, server-side sorting must be enabled.

Key Mapping Method and Column Mapping Method

System Generated mapping is the default column mapping method for keys and comparison columns. You can work more with column mapping after you generate the mapping and save the compare pair to the repository.

Generate the mappings

When you are finished with the preceding tasks, click **Generate Mappings**. A message at the top of the page indicates whether the operation succeeded or failed.

If there are any duplicate mappings, you are notified after you click **Save**. These can be removed later.

Save the compare pairs to the repository

Until now, the new compare pairs are not saved to the repository. Saving your work frequently is a good practice. Whether your Oracle GoldenGate Veridata Web User Interface session is terminated intentionally or unintentionally, you will lose the compare pairs that you created unless they are saved to the repository. You can always return to any compare pair to perform additional configuration tasks.

The Preview tab displays compare pairs that were generated successfully, but not yet saved. To save your work, go to the Preview tab now.

> **Note:**
>
> To make it easy to return to your work on a compare pair, save the task as a Favorites shortcut.

Mapping Objects Manually

If your source and target naming conventions do not support using the pattern mapping method, or if you would rather do the mapping yourself, use the manual mapping method to create compare pairs. Manual mapping is performed by pairing source and target tables or files one at a time by name. Manual mapping can be used in conjunction with pattern mapping as needed.

Use the Manual Mapping tab of the Compare Pair Configuration page to manually map objects. This tab is also used to configure row subsets.

You must have the Administrator or Power User role to configure compare pairs.

To use manual mapping, complete the six tasks in the following sections in order.

Confirm group information

The group that you have chosen to contain the compare pairs is displayed at the top of the page, along with the source and target connections. Verify that these details are correct before proceeding.

Select the datasource

Specify the database structures that contain the tables to be mapped. The selections vary by database type:

Windows, UNIX, and Linux-based databases

Under Datasource Information, select the Source and Target catalogs (if applicable to the database) and schemas that contain the tables that you want to map as compare pairs. You can select from a drop-down list or click the **Browse** button to open a selection dialog that has filtering options. To use the filtering options, do the following:

1. Supply a search string to filter the list, and then click **Apply Filter**. The filter returns a list of names.
2. In the Select column of the list, click the name of the appropriate schema or catalog for this compare pair.
3. Click the **Select** button.

Enscribe databases

1. In the File Pattern box under Datasource Information, enter the Source and Target volume, sub-volume, and file patterns of the source and target files that you want to map as compare pairs. Use the format of volume.subvolume.file_pattern (for example, volume.subvolume.*).
2. Click the **Refresh Source Files/Tables and Refresh Target Files/Tables** buttons.

(Optional) Filter the lists

By default, the Manual Compare Pair Mapping list shows all of the source and target objects that are contained by the specified datasources. You can filter the list at any time in the process of creating compare pairs. For example, after you finish mapping several compare pairs, it might be helpful to reduce the size of the list by using the filter to show only unmapped objects. This saves you navigation time.

To filter the list:

1. Expand Filters to show filter options.
   - Show All: Shows all tables or files that are contained by the specified datasources. Those already mapped in another compare pair are denoted by a check mark, but you can still select them for this pair if desired. However, make certain to give this compare pair a different name.
   - Show Mapped: Shows only previously mapped tables or files.
   - Show Not Mapped: Shows only tables or files not mapped in any other compare pair.
   - Names Like: Shows specific compare pairs whose names contain the supplied string.
2. Click **Apply Filter** to enact the filter and update the view.

Map a table or file

In the list of objects that can be mapped, there will be a check mark in the Mapped column for any object that is already mapped (unless a filter was used to filter out mapped objects).

Map one source object and one target object at a time:

- Click the **Select** button next to the source object.
- Click the **Select** button next to the target object.

As you select objects, their names are placed in the Compare Pair Name field under Compare Pair Details for further configuration of this pair.

Specify compare pair details

So far, you have mapped object names, but there are still other configuration options that might apply to this compare pair.

Under Compare Pair Details, specify the following:

Compare Pair Name

You can keep the default name format of `<source>=<target>` or use another name of your choosing. Example default names:

```
customers=customers2
```

```
\FIN.$DATA.FN.CUST1=\FIN.$DATA.FN.CUST1
```

To specify a name:

- A user-defined name cannot contain spaces but can contain underscores and equal signs. It is case-sensitive. If you will be generating a large number of compare pairs, use a naming convention that is intuitive. The Oracle GoldenGate Veridata Web User Interface displays the actual object names if you hover the mouse over the name of a compare pair.

Delta Processing

Delta processing can be used for all supported databases. Select **Enable Delta Processing** to use the delta processing feature. By default it is disabled. To use delta processing, server-side sorting must be enabled.

For more information about delta processing, see Using Delta Processing.

Profile

You can leave this field blank (to use the default job profile) or you can specify a profile for this compare pair that will override the job profile.

To select a profile:

1. Select the profile from the drop-down list or, to use filtering options, click the **Browse** button to open the Select a Profile page.
2. To use **Select a Profile**:

a. Supply a search string to filter the list by profile name, and then click **Apply Filter**.

b. Select the correct profile from the list, and then click **Select**.

Row Partitions

You can specify which rows to include or not to include in a comparison by specifying a SQL predicate statement or an Enscribe partition.

For more information on row partitions, see Configuring row partitions.

Key Mapping Method and Column Mapping Method

System Generated mapping is the default column mapping method for keys and comparison columns. You can work more with column mapping after you save the compare pair to the repository

Generate the mappings

When you are finished with the preceding tasks, click **Generate Compare Pair**. A message at the top of the page indicates whether the operation succeeded or failed.

If there are any duplicate mappings, you will be notified after you click **Save**. These can be removed later.

Save the compare pairs to the repository

Until now, the new compare pairs are not saved to the repository. Saving your work frequently is a good practice. Whether your Oracle GoldenGate Veridata Web User Interface session is terminated intentionally or unintentionally, you will lose the compare pairs that you created unless they are saved to the repository. You can always return to any compare pair to perform additional configuration tasks.

The Preview tab displays compare pairs that were generated successfully, but not yet saved. To save your work, select the **Preview** tab now.

> **Note:**
>
> To make it easy to return to your work on a compare pair, save the task as a Favorites shortcut.

Reviewing and Saving Compare Pairs

Use the Preview tab to review, edit, and save or discard the compare pairs that you generate on the pattern-mapping or manual-mapping tab of the Compare Pair Configuration page.

Saving your work frequently is a good practice. If your Oracle GoldenGate Veridata User Interface session is terminated intentionally or unintentionally, you will lose the compare pairs on the Preview tab unless they are saved to the repository. You can always return to any compare pair after saving it, to perform additional configuration tasks.

Viewing details for a compare pair that is finished

Use the Details for Finished Compare Pair page to view near-real-time information about a compare pair that is finished being processed.

To access this page:

1. Under Finished Jobs in the navigation pane, select View By Compare Pair.
2. In the Compare Pair Name column of the Finished Compare Pairs list, click the name of the pair for which you want to view details.

The upper portion of the page shows specific information about the compare pair, including:

- Compare Pair Name: the name of the compare pair.
- Source Table and Target Table: the names of the underlying source and target objects.
- Compare Pair Run ID: the identifier used by Oracle GoldenGate Veridata Server while running a comparison. It is useful for locating messages related to this comparison in the veridata.log file, which is used for troubleshooting purposes when resolving support cases.
- Compare Pair Start Time: the time when processing started for this compare pair.
- Compare Pair Run Duration: how long it took to process this compare pair.
- Run Status: the outcome of this comparison (Waiting, Canceled, Canceling, Running, Finished)
- Comparison Status: the status of the comparison (see Possible comparison status conditions.) If delta processing is enabled for this compare pair, the delta symbol is displayed.
- Overall Rows Compared: the total number of rows in the source and target objects that were compared during this run of the compare pair.
- Overall Rows With OOS: the number of rows in the source and target objects that are out-of-sync.

Using the charts

The charts on this page help you to determine what types of operations are causing the most out-of-sync conditions and which phases of processing are taking the most time to complete.

Overall Rows Compared pie chart

Plots the result from the Overall Rows Compared field according to the possible outcomes of the comparison: In-Sync or Out-Of-sync.

Operations Out-Of-Sync pie chart

Plots the number of out-of-sync rows according to their operation type:

Inserts: source table has the row, but the target does not.

Updates: source table has different row values from the target.

Deletes: target table has the row, but the source does not.

Compare Pair Performance table

Lists performance statistics based on each phase of a run: the sorting phase (for source and target), the initial comparison phase, and the confirmation phase. Note: sorting statistics are only available if the job profile specifies server-side sorting (row sorts are performed by Oracle GoldenGate Veridata, not by the database engines).

Start Time: The time that a run phase started.

Run Duration: The amount of time that it took to complete a run phase.

Rows Processed: The number of rows that were processed in a run phase.

Rows Per Second: The performance rate of a run phase, in terms of rows processed.

Bytes Processed: The overall volume of data, in bytes, that was processed in a run phase.

Rows Per Second: The performance rate of a run phase, in terms of bytes processed.

Compare Pair Performance History chart

Plots the performance of each of the run phases as a bar chart. To use this chart:

From View Step, select a run phase to plot.

Initial comparison step: Select Initial Compare.

Confirmation step: Select Confirm Out-Of-Sync.

Sorting process: Select Sorting.

From View Performance History By, select the criterion for this performance analysis. This sets the units of measure for the vertical axis of the chart. You can measure:

Number of rows or bytes that were processed

Number of rows or bytes per second that were processed

Run duration (the time spent processing the rows)

Other actions from here

To view a report of the results for this run of this compare pair, click **View Comparison Report**.

To view details about out-of-sync rows, including column values, click **View Out-Of-Sync Rows**. This link only appears if the compare pair has out-of-sync rows.

To edit this compare pair, click **Compare Pair Configuration** at the top of the page to go to the Compare Pair Configuration page.

Configuring Column Mappings

This section provides explains how to use column mapping for compare pairs.

> **Note:**
>
> You must create a compare pair before you can do column mapping. See Configuring Compare Pairs.

To access the Column Mapping Configuration page:

1. Click the name of a group on any page that displays groups. To access this page from the navigation pane, click **Group Configuration**, then the name of the group that you want to verify. This displays the Edit Group page.
2. Click **Go to Compare Pair Configuration**.
3. Click the **Edit...** link for the compare pare you are interested in.

To view an existing column mapping

Use the Existing Column Mapping tab to view the current mappings for the key columns and the comparison columns, and which method is being used to map them. This tab also shows the source and target columns that were excluded from being mapped, either explicitly by a user or because Oracle GoldenGate Veridata could not find a corresponding column with the same name.

To test the validity of the current mapping

It is possible for object metadata to change from the time that a column mapping was created to the time that you view it on the Existing Column Mapping tab. To confirm that the mapping is still valid, click Validate Existing Column Mapping.

See Validating column mappings.

To create or change a column mapping

Use the Column Mapping tab to create or change a column mapping.

You must have the Administrator or Power User role to create or edit column mappings.

Choosing a column mapping method

There are two methods for mapping key columns and comparison columns:

- System-generated: Column mappings are configured automatically by Oracle GoldenGate Veridata at runtime based on current object metadata.
- User-defined: Column mappings are configured manually by an Oracle GoldenGate Veridata user who has the Administrator or Power User role.

All new compare pairs default to System Generated for key columns and for comparison columns. You can change to a different mapping method at any time.

How these methods apply to keys

- **System Generated**: If you know that the objects in a compare pair both contain a primary key or a unique index, you can leave the key mapping method set to the default of System Generated. The key columns will be mapped automatically. To map keys, Oracle GoldenGate Veridata finds all indexes on the source and target

objects and tries to find a primary key on each one. If primary keys are not found, Oracle GoldenGate Veridata tries to use the smallest index (least number of columns), and then it maps the columns that have identical names and comparison formats. Any columns that cannot be matched are excluded from the configuration.

- **User Defined**: If an object has neither a primary key or unique key, you can use the User Defined method to map key columns manually, one by one. You can also use the user-defined method to override existing keys or indexes, but the columns that you select to use as a key must ensure the uniqueness of rows. Also avoid using source and target indexes that have different precision levels or other characteristics that can reduce the accuracy of row selection, especially in a heterogeneous environment.

How these methods apply to comparison columns

- **System Generated**: If the source and target comparison columns have the same names and comparison formats, you can leave the comparison column mapping method set to the default of System Generated. Oracle GoldenGate Veridata will map those columns automatically at runtime. Non-matching columns are excluded from the configuration. By default, this method includes all of the columns in a comparison. This method defaults to the hash comparison method. You can change the comparison method later by editing the compare pair.
- **User Defined**: Use this method to map source and target columns manually and to control the comparison method.

You can combine these methods to speed up the mapping process. If most of the column names support system-generated mapping, you can use it and then switch to the user-defined method to map the remaining columns, or to exclude columns from the comparison. For example, you can exclude columns if you know that their values never change or if you expect their values to be out-of-sync.

Validating column mappings

Validation is a preliminary test to determine whether or not the source and target table structures are compatible, and that they both have primary or unique key columns that match. You can perform a manual validation at any time. Oracle GoldenGate Veridata always performs a validation at runtime.

To perform the validation

Click **Validate Column Mapping** on the Existing Compare Pairs tab of the Compare Pair Configuration page. Make certain to place a check mark in the Select column for all of the pairs that you want to validate. The result is displayed under Validation Status.

To access the Compare Pair Configuration page:

1. Click the name of a group on any page that displays groups. To access this page from the navigation pane, click **Group Configuration**, then the name of the group that you want to verify. This displays the Edit Group page.
2. Click **Go to Compare Pair Configuration**.

You can also perform validation when you are configuring column mapping on the Column Mapping Configuration page.

You can also perform validation when you are configuring column mapping on the Column Mapping Configuration page. On the Compare Pair Configuration page, click **Edit** in the Column Mapping column.

What if a compare pair fails the validation test?

A result of **other than** ✓ or ? indicates that the validation failed.

The ? status indicates that a validation was not performed. Check the **Select** column to make sure that you included the compare pair in the validation.

If you are manually mapping columns and a validation fails, click **Reset** to go back to the former mapping and then try the mapping and validation again. It is possible that you selected the wrong source and target columns as a pair.

If a second manual mapping fails validation, or if a system-generated mapping fails validation, use the color of the status icon as your guide to the cause and then check the underlying data for incompatibilities.

If a previous validation of the compare pair succeeded, but the current one failed, this probably means that the structure of the underlying objects changed.

> **Note:**
>
> A successful preliminary validation does not guarantee a successful runtime column validation. It is possible that changes to the structure of a source or target object can invalidate the original test and cause the job to fail at runtime. The validation process should only be used as a preliminary test for current incompatibilities.

Using Delta Processing

This section provides answers to typical questions about the delta processing feature that is available for all Oracle GoldenGate Veridata supported databases.

What is Delta Processing?

How does it work on NonStop Platforms?

Oracle GoldenGate Veridata finds a changed block by detecting a change in its Volume Sequence Number (VSN) since the time of the last comparison. The VSN is a disk-specific change number that increments sequentially with each database operation that is performed on the data. Each time that a row changes, there is a change in the VSN of the disk block where the row resides.

There is no relationship between a VSN in a file on one disk and a VSN on another. Oracle GoldenGate Veridata tracks VSNs on a per-partition basis on the source and target disks and maintains its own correlations to perform accurate delta comparisons. Once you enable delta processing, it is used for all subsequent runs until you disable it again.

> **Note:**
>
> The first run of a compare pair always compares all of the rows in the source and target objects to establish an initial VSN state from which to evaluate deltas in future runs.

How does it work on all other Platforms?

Oracle GoldenGate Veridata compares a source database table to the target database table. The source and target tables are configured using compare pairs, which are grouped and added to a job to run the comparison (see Configuring Compare Pairs). When all the rows in the table are compared, it is a Full Comparison Job.

During the subsequent runs of a comparison job, the comparison of the tables can be performed based on what has changed in the tables from previous job run; these jobs are Delta Processing Jobs. Delta processing is usually performed on tables that contain a large number of rows so it is probable that in these tables there will be columns eligible for delta processing. The delta comparison column must contain a value that is modified every time that the row is modified and this value must always be increasing. Any data type that meets this requirement is supported. By default, the columns of the table that are mapped to Numeric or Timestamp comparison formats are supported. For example, `TIMESTAMP`, `TIMESTAMP_TZ`, and `NUMBER`.

The delta base is the value of the Delta Column on the basis of which the delta comparison was performed. Every time a comparison is run, a delta base value is captured. Depending on the number of delta comparison jobs performed, there can be multiple delta base values so a list of delta base values for the compare pair is generated. For example, the first time a Full Comparison is run and the maximum value of the Delta Column is the delta base, `DeltaBase-1`. A second Delta Processing Job run based on `DeltaBase-1` results in `DeltaBase-2` being captured again as the maximum of Delta Column. In the third run, you can use either `DeltaBase-1` or `DeltaBase-2` for the comparison or run a Full Comparison Job.

When Should I use Delta Processing?

Delta processing is suitable for use with very large Enscribe files and SQL tables that, otherwise, would take a long time to process. It does consume additional overhead, so it is probably not practical for use with smaller sets of data. Try running a test comparison without delta processing first. If, in your opinion, the compare pair takes too long to process, try running it again with delta processing enabled. If the delta-enabled run is significantly shorter than the first test, continue to use it. If there is only marginal improvement, it might be better to disable delta processing to prevent the added overhead. The performance gains of delta processing are in the initial comparison step of the run. Delta processing can cause the confirmation step to be longer if the source and target rows end up on different data blocks.

What Process Performs the Delta Processing?

The delta processing is performed by the Oracle GoldenGate Veridata Agent.

For NonStop platforms, the VSN information is retrieved by a privileged process named `vsnserv`. During the installation of Oracle GoldenGate Veridata Agent on a NonStop platform, PROGID was used for the `vsnserv` program to run as SUPER.SUPER to be able to read the file labels for this purpose.

For all other platforms, the delta processing queries the compare pair to retrieve the delta base values for both the source and target tables, which creates the column mapping.

What sorting method can be used with delta processing?

To use delta processing, you must enable server-side sorting by setting the sorting method to Server within the profile that is associated with the compare pair or the one that is associated with the job when you run it.

> **Tip:**
>
> If you always will be using delta comparisons, consider setting the sorting method to Server within the default Oracle GoldenGate Veridata profile. That way, nobody will forget to select the correct profile when the jobs are run.

What other important things should I know when using delta processing?

The delta processing mechanism can fail to detect an out-of-sync delete, if that delete was the only source row that was modified in a block, and if that delete did not get propagated to the target. In such a case, the block on the target that contains the relevant row does not get modified, so it is skipped by the target Veridata Agent during delta processing.

The delta value field is pre-populated in the column mapping UI, based on the query for retrieving the delta value.

Support Considerations

The following table shows the supported delta column types. The possible list of delta columns for the delta configuration UI is identified by reviewing the compare formats for the corresponding column-pairs in the source and target tables.

Veridata Comparison Format	Can it be Delta Column?
`STRING`	No
`STRING_EN`	No
`TIMESTAMP`	Yes
`TIMESTAMP_TZ`	Yes*
`DATETIME`	Yes
`SMALLDATETIME`	No
`DATE`	No
`TIME`	No
`NUMBER`	Yes
`FLOAT`	No
`BINARY`	Yes*
`BLOB`	No
`CLOB`	No
`INTERVAL`	No

Veridata Comparison Format	Can it be Delta Column?
DEC_FLOAT	No
BINARY_TIMESTAMP	Yes*
SYBTIME	No
CLOB_NFC	No
STRING_NFC	No

> **Note:**
>
> * denotes that Hive doesn't support the delta column. Supported data types for Hive are NUMBER, TIMESTAMP and DATETIME.

For Oracle Database, the `ORA_ROWSCN` pseudo delta column is supported, and selected by default when `ROWDEPENDENCIS` are enabled for that table.

For DB2 for i, z/OS, and LUW, columns having the `GENERATED FOR EACH ROW ON UPDATE AS ROW CHANGE TIMESTAMP` clause is selected by default.

For all supported databases, `TIMESTAMP` columns are shown in order first then followed by `NUMERIC` columns.

Only one column is supported for delta processing and is similar to:

```
SELECT column-names from table name where delta_column delta_condition
supplied_delta_value
```

For example:

```
select * from TableA where startdate >= '12-01-2012 21:24:00'
```

How do I enable or disable delta processing?

You can enable or disable delta processing for database compare pair from the Delta Processing tab and the Existing Compare Pairs tab of the Compare Pair Configuration page. In addition, you can use the Delta Processing Enabled check box on the Run Configuration page to toggle this feature.

How do I configure delta processing?

On the Column Mapping Configuration page, click the **Delta Processin**g tab.By default, the options are automatically populated and the Enable Delta Processing check box is selected. You can disable the feature by clearing the check box, which renders all other options inactive.

You can use the defaults or change any of the following:

- the source or target columns for processing,
- whether to use a source or target query
- the source or target database query

After you have configured the delta processing for the compare pair, click **Save** to apply your changes.

How do I know that delta processing is being used?

When a compare pair is configured for delta processing, a delta symbol is displayed in the lists of the Edit Group and Compare Pair Configuration pages.

If delta processing is enabled for at least one compare pair in a job that is selected on the Run/Execute Job page, there will be a check mark in the Delta Processing Enabled box.

Selecting Delta-Base value

Veridata automatically selects the Delta column value of the latest in-sync comparison as delta base value.

If Veridata could not select the delta base value automatically the “Select …” link is shown in the run-jobs page. And On click of that all the delta base values are listed and users can choose a delta base value from it, it is always advisable to choose in-sync comparison for delta-base values.

Delta base time can also be selected by navigating to the finished compare details for the previous comparison and clicking the run button. Veridata will use the time from that comparison as the delta base time. If the out-of-sync data from the previous comparison has not been fixed, those out-of-sync rows may not be detected again until another complete comparison is run.

When Veridata automatically selects the Delta-Base value the time stamp of the comparison is shown as a link to the user in run-jobs page with tool-tip.

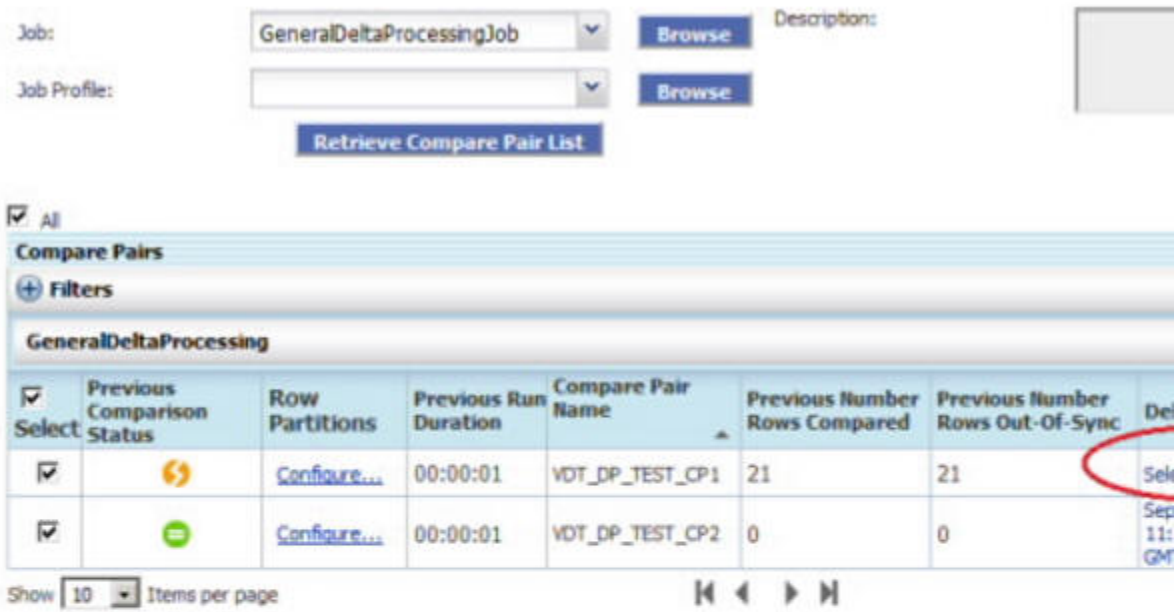

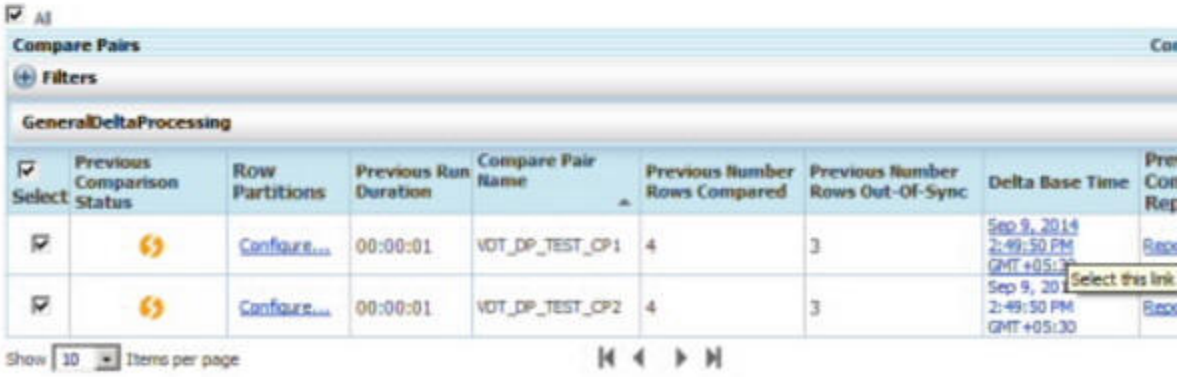

You can select that link to choose and alternate delta base value.

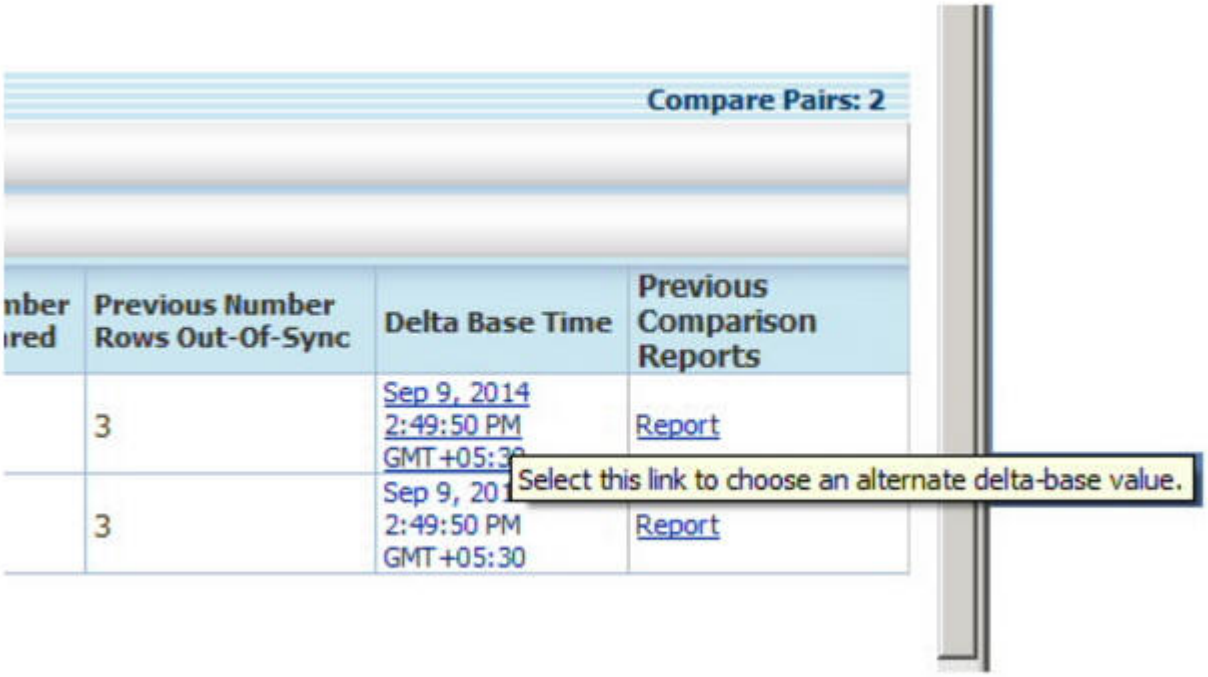

Can I override delta processing when I run a job?

What if I perform maintenance on the tables or files for which I am using delta processing?

When you perform maintenance on objects in a compare pair that has delta processing enabled, the best practice is to disable delta processing for the next run so that Oracle GoldenGate Veridata compares all of the rows. You can disable delta processing at the compare pair level or as a job override. Starting again with a full comparison allows a new delta base state to be established and will make subsequent delta comparisons faster. Otherwise, delta processing could actually take longer than a complete comparison of all rows.

For example, if a 'FUP RELOAD' is performed on the source, but not on the target, it could cause delta processing to return a much larger number of rows from the source than from the target, based on the last delta state. The source rows that are returned would be rows that actually did not change. This happens because FUP moves records around and combines blocks, but does not change the data. However, the VSNs for the affected blocks will change. Oracle GoldenGate Veridata cannot detect that the reload was done since the last delta state. Thus, the next time that the VSN for a block changes, all of the rows in that block will be returned.

Conversely, on the target, no VSNs are changed for the corresponding data (because a reload was not done there), so those rows are not returned for delta processing. This anomaly will be resolved by the confirmation step, but this slows the overall comparison process because that step is much slower than the initial comparison step.

Configuring Profiles

A profile is a set of global processing parameters, each containing unique settings for a specific purpose. Oracle GoldenGate Veridata provides a default profile, but you

probably will want to create your own profiles. too. You can create as many profiles as needed and associate them with any job or compare pair (to override the job profile; see Configuring Compare Pairs). You can override profile assignments at run time.

Profiles are managed from the Profile Configuration page. All profiles that exist within the Oracle GoldenGate Veridata repository are shown in the **Existing Profiles** list on this page.

To filter the Existing Profiles list:

1. **Expand Filters** and type a search string. The filter matches a search string character for character.
   - **Names Like**: filters by profile name.
   - **Description Like**: filters by a string in the job description.
2. Click **Apply Filter** to display the selected profiles.

To configure profiles

You must have the Administrator or Power User role to create, edit, or delete a profile.

- To create a profile:

  To create a profile, click **New**. The New Profile Assistant prompts for a name and description and then directs you to the Edit Profile page so that you can configure profile settings.
- To edit a profile:

  To edit a profile, click its name in the Existing Profiles list, or select it in the Select column and then click **Edit**. The Edit Profile page is displayed. See Editing Profile Settings.
- To delete a profile:

  To delete a profile, select it in the Select column of the Existing Profiles list, and then click the **Delete** button.
- To change the default profile:

  Oracle GoldenGate provides a default profile (see Using the default profile) that is used for all of the compare pairs and jobs that are not linked to a user-defined profile. This profile is used automatically unless another profile is selected when creating a compare pair or a job, or when running a job. You can change the settings of the default profile by editing it.

Other configuration tasks

Configuring Connections

Configuring Groups

Configuring Jobs

Editing Profile Settings

This section describes the parameters that you can edit for a profile.

The profile settings are categorized as follows:

- General: Controls output options.

- Sorting Method: Controls sorting method and memory management. Data is sorted to match keys (or a key specification) so that the correct source and target rows are compared.
- Initial Compare: Controls parameters for the process that performs the initial compare step.
- Confirm-Out-Of-Sync: Controls parameters for the process that performs the confirmation step.
- Repair: Controls parameters for the repair process.

Specifying a Sorting Method

The Sorting Method settings of the Edit Profile page specify whether data sorting will be performed by the database or by Oracle GoldenGate Veridata Server. Specify this method in the Sort Data Using option on the Edit Profile page.

By default, Oracle GoldenGate Veridata uses the database to sort data for comparison. This default is due to historical conditions that are no longer valid. Server-side sorting is the current recommended sorting method. Database sorting should only be considered when the ordering produced by the database is identical to the ordering produced by Oracle GoldenGate Veridata Server. Following is a list of the types of conditions that will produce differing sorted ordering of the rows:

- Character encoding conditions: Oracle GoldenGate Veridata compares character data as UTF-8 encoded bytes. To match server-side ordering, key columns that contain character data must contain only ASCII data or be encoded using UTF-8, and the database must use binary comparisons for character data (no comparisons that are case-insensitive or specific to a locale).
- Some datetime data types, such as Teradata `TIME`, may sort differently in the database and in Oracle GoldenGate Veridata Server.
- To make database ordering consistent with the ordering done by Oracle GoldenGate Veridata Server, the Oracle GoldenGate Veridata Agent may add `ORDER BY` clauses to the initial comparison SELECT statement that will make the database ignore indexes on the columns. An example is `TIMESTAMP` with `TIMEZONE` data, where Oracle GoldenGate Veridata orders the data by the string representation of the data rather than by the absolute time.

When Oracle GoldenGate Veridata Server performs the sort, the Oracle GoldenGate Veridata Agents return data in the natural order that is provided by each database, and then the data is sorted by two server sort processes, one to sort source rows and the other to sort target rows. Server-side sorting supports a maximum row length of 32768 bytes. This limit normally is not exceeded when the hash comparison method is used.

Specifying Temporary Storage Directory for Source and Target Data

Specifies a location on the source disk or target disk to use as temporary storage when there is not enough memory to process all of the data that is being sorted. If no locations are defined, the default is to use a directory under the Oracle GoldenGate Veridata Server home location.

Choosing locations on different physical drives might speed up comparisons in some circumstances. You can specify multiple locations for each process, separating each one with a semicolon (for example /tmp/sort1; /tmp/sort2). All locations specified must already exist. The drives used should have sufficient free disk space. To calculate the approximate amount of space needed, use this formula:

1.5 * (Trows * (Tkey + 20)) * nTables

where:

Trows = the number of rows in table

Tkey = the average size of the table key, in bytes

nTables = the number of tables that are being compared

Using the default profile

The default Oracle GoldenGate Veridata profile is applied when a custom profile is not linked to a job or compare pair.

The name of the default profile is `$default`.

To view and edit the settings of this profile, use the Edit Profile page.

To edit the default profile, you must have the Oracle GoldenGate Veridata administrator role.

To restore the default profile to its original (installed) settings, click the **Reset all values back to Oracle GoldenGate defaults** button on the Edit Profile page.

Configuring Jobs

To run comparisons, you must run a job. The job configuration determines which compare groups are processed.

> **Note:**
>
> Before creating a job, you must create at least one compare group (see Configuring Compare Pairs) to link to the job. To use customized runtime parameter settings, you must also create at least one profile (see Configuring Profiles). Otherwise, the job will use the default profile that is supplied by Oracle GoldenGate Veridata.

Jobs are managed from the Job Configuration page. To access this page, click **Job Configuration** in the navigation pane. All jobs that are defined within the Oracle GoldenGate Veridata repository are shown in the Existing Jobs list on this page.

To filter the Existing Jobs list

1. **Expand Filters** and type a search string. Strings are matched character for character.
   - **Names Like**: filters by job name.
   - **Description Like**: filters based on a string in the description.
2. Click **Apply Filter** to display the selected jobs.

To configure a job

You must have the Administrator or Power User role to create, edit, or delete a job.

- To create a job

  To create a job, click **New**. The New Job Assistant is displayed. The assistant will prompt for:

  - a name and description.
  - groups to be linked to the job.
  - (optional) a profile to link to the job; otherwise the default profile will be used.
  - (optional) connections to link to the job; connections specified will override those linked to each group.

- To edit a job

  To edit a job. click the name of the job in the Existing Jobs list, or select it in the Select column and then click **Edit**.

  The Edit Job page is displayed.

- To delete a job

  To delete a job, select it in the Select column of the Existing Jobs list, and then click the **Delete** button.

- To run a job

  To run a job, select it in the Select column of the Existing Jobs list, and then click the **Run** button to go to the Run/Execute Job page.

Other configuration tasks

Configuring Connections

Configuring Groups

Configuring Profiles

3
Working with Jobs

Learn all about jobs, including how to run a job, monitor jobs that are running, and also how to repair out-of-sync jobs.

- Overview of Jobs
- Estimating Comparison Time
- Running A Job
- Viewing jobs that are running
- Viewing jobs that are finished
- Repairing Out-Of-Sync Jobs
- Using the Comparison Report

Overview of Jobs

Once you have configured groups and compare pairs into jobs, you can run those jobs. Oracle GoldenGate Veridata enables you to control which groups and compare pairs are processed during any given job run, and with which runtime parameters. Once a job is running, you have easy access to views of current and finished runs.

Related Topics

Estimating Comparison Time

Running A Job .

Viewing jobs that are running

Viewing jobs that are finished

Estimating Comparison Time

When tables are large, you might want to estimate the amount of time that a comparison will take before running a full comparison. To get an estimate, run a test comparison of a limited number of rows. For example, if there are 100 million rows in a table, you can run a comparison for the first million rows and then multiply that amount of time by 100.

To specify the number of rows compared

1. On the navigation pane, click **Profile Configuration**.
2. In the list, click the name of the profile that is linked to the job (or select it, and then click **Edit**). The Edit Profile page is displayed.
3. Click **Initial Compare**.
4. Under General, clear the **Use Default** box for the Limit Number of Input Rows setting, and then type the number of rows that you want to compare.

5. After running the test comparison, change Limit Number of Input Rows back to the default for the full comparison run.

Running A Job

Jobs are executed from the Run/Execute Job page. Any job that has been previously configured can be run from this page. Before running a job, you may want to estimate its duration (see Estimating Comparison Time).

Running a job without modification

1. For Job, select a job from the drop-down list or click the **Browse** button to open the Select a Job page, where you can use filters to constrain the selection choices.
   - Optional) Under Filters, type a string in Names Like or Names Not Like to include or exclude jobs by name, and then click **Apply Filter** to display the selected jobs.
   - Select a job from the list by clicking in the Select column. You can only start one job at a time.
   - Click the **Select** button to select the job and close the Select a Job page.
2. (Optional) For Job Profile, select a profile if you want it to override the profile that is linked to this job. You can select from the drop-down list or click the **Browse** button to open the Select a Profile page, where filters are provided.
   - (Optional) Under Filters, type a string in Names Like or Description Like to include profiles based on their names or a string in their descriptions, and then click **Apply Filter** to display the selected profiles.
   - Click in the Select column next to the name of the profile that you want to use.
   - Click the **Select** button to select the profile.
3. Click **Run Job**. By default, all compare pairs that are linked to the selected job will be processed, even though their names do not appear in the Compare Pairs list.

Modifying a job before running it

After selecting the job and (optional) profile, click **Retrieve Compare Pair List**. All of the compare pairs in the job are displayed in the Compare Pairs list, and all are selected by default.

The compare pairs are organized according to the groups that contain them, with the groups listed in alphabetical order; for an example, see How groups are shown on the Run/Execute Job page.

About the list columns:

The following information is displayed in the Compare Pairs list. This information can help you to refine the job configuration.

- **Previous Comparison Status**: Indicates the results of the previous comparison run of the compare pair.
- **Row Partitions**: Enables you to configure or override row partitions (see the steps below).
- **Previous Run Duration**: The length of time that the previous comparison took. Can be used as a basis for estimating the duration of the next comparison.

- **Compare Pair Name**: The name of the compare pair. To view actual object names, hover the mouse over the name.
- **Previous Number Rows Compared**: The number of rows that were compared in the previous run.
- **Previous Number Rows Out Of Sync**: The number of rows that were out of synchronization in the previous run.
- **Delta Base Time**: Shows the end time of the previous comparison, which provides a start time on which to base delta processing for the next run (see the steps below).
- **Previous Comparison Reports**: Links to comparison reports from previous runs.

1. Filter the list, if necessary.

   To filter the list of compare pairs, expand Filters to expose filter options. You can filter the list of compare pairs by the following:

   - **Group Names Like**: Type a search string to show names that contain the string.
   - **Compare Pairs**: You can filter compare pairs based on their previous comparison status.
   - **Compare Pair Name Like**: Type a string to filter by the name of the compare pair.
   - **Source Table Name Like**: Type a string to filter by the name of the source table or file.
   - **Target Table Name Like**: Type a string to filter by the name of the target table or file.

   After you specify your filter criteria, click **Apply Filter**. By default, all items that the filter returns are selected for processing.

2. Select the compare pairs that you want to be processed in this job.

   - To select all of the compare pairs (in all of the groups) in the list, make certain that the All check box displays a check mark. (This box is located directly above the Compare Pairs header.) This action affects only the compare pairs that are in the list, not any that were excluded by means of the filter.
   - To deselect all of the compare pairs in the list, clear All. To deselect all when some pairs are selected and others are not, you might need to click **All** twice, once to select all and once to deselect all.
   - To select all of the compare pairs that are in a given group, click the Select box beneath the name of the group.
   - To include or exclude an individual compare pair, select or clear the Select box next to its name.

   To learn how navigate a large list, see Managing the Page View.

3. To configure or override row partitions, click in the Configure in the Row Partitions column to open the row partitions editor.

4. To change the default job attributes and profile parameters, click the **Set Override Run Options** link to go to the Run Options page (see Configuring run options).

> **Note:**
>
> The **Command Line To Be Used** box shows the underlying Vericom command that will be used to run this job. For more information about Vericom, see the Using the Veridata Command Line Tool.

5. If delta processing is enabled for any of the compare pairs in this job, the **Delta Processing Enabled** check box is shown. To disable delta processing, clear the check box.
6. Click **Run Job**.

What happens when you run a job

The groups that you selected for processing will be scheduled in the order shown in the list, but they may not be processed in that order. The time that it takes to process each group depends on the size of that group, the operating system that is hosting the database, the speed of the database, the system resources that are available and their activity, and various other factors. Because Oracle GoldenGate Veridata uses multi-threaded processing, groups could start running out of order. This is normal. To start a particular group right away, assign it to a dedicated job, and then run that job before the other ones.

Running jobs remain running when you log out of Oracle GoldenGate Veridata Web. When you log in again, you will see current statistics if the job is still running, or you will see results if the job has finished.

Related Topics

Viewing jobs that are running

Viewing jobs that are finished

Using the Comparison Report

Estimating Comparison Time

When tables are large, you might want to estimate the amount of time that a comparison will take before running a full comparison. To get an estimate, run a test comparison of a limited number of rows. For example, if there are 100 million rows in a table, you can run a comparison for the first million rows and then multiply that amount of time by 100.

To specify the number of rows compared

1. On the navigation pane, click **Profile Configuration**.
2. On the list, click the name of the profile that is linked to the job (or select it, and then click **Edit**). The Edit Profile page is displayed.
3. Click **Initial Compare**.
4. Under General, clear the **Use Default box for the Limit Number of Input Rows** setting, and then type the number of rows that you want to compare.
5. After running the test comparison, change Limit Number of Input Rows back to the default for the full comparison run.

How groups are shown on the Run/Execute Job page

In this example, there are two groups: Bookstore and BookstoreReporting:

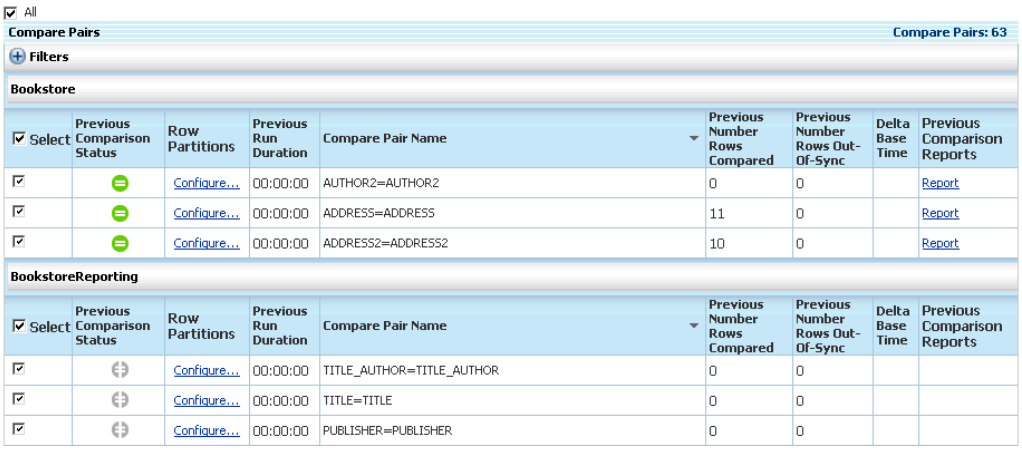

All

Compare Pairs — Compare Pairs: 63

Filters

Bookstore

Select	Previous Comparison Status	Row Partitions	Previous Run Duration	Compare Pair Name	Previous Number Rows Compared	Previous Number Rows Out-Of-Sync	Delta Base Time	Previous Comparison Reports
☑		Configure...	00:00:00	AUTHOR2=AUTHOR2	0	0		Report
☑		Configure...	00:00:00	ADDRESS=ADDRESS	11	0		Report
☑		Configure...	00:00:00	ADDRESS2=ADDRESS2	10	0		Report

BookstoreReporting

Select	Previous Comparison Status	Row Partitions	Previous Run Duration	Compare Pair Name	Previous Number Rows Compared	Previous Number Rows Out-Of-Sync	Delta Base Time	Previous Comparison Reports
☑		Configure...	00:00:00	TITLE_AUTHOR=TITLE_AUTHOR	0	0		
☑		Configure...	00:00:00	TITLE=TITLE	0	0		
☑		Configure...	00:00:00	PUBLISHER=PUBLISHER	0	0		

Overriding Row Partitions

To select an existing row partition

1. There is a check mark in the Use column next to the partition that is currently active. To select a different one, click in the Select column next to its name.
2. Click **Edit**.
3. Under Row Partition Configuration, select **Use At Runtime**.
4. Click **Apply**. A check mark is placed under Use next to this partition.
5. Click **OK** to close the Compare Pair Row Partitions window.

To specify a new row partition

1. Click **New** under Source or Target. In the Name box under Row Partition Configuration, the name "Override_Partition_<n>" appears, where <n> is a sequential number. A partition override name cannot be customized.
2. Do one of the following:
   - SQL tables: In the SQL Predicate Statement field, type the predicate statement. The SQL predicate is the conditional statement that follows the WHERE keyword, for example: LAST_NAME BETWEEN "A" AND "M". Do not type the WHERE keyword; it will be added automatically at runtime by Oracle GoldenGate Veridata.
   - Enscribe: For Enscribe Partition, type a beginning key value after Start Key and an end key value after End Key. If the key is in ASCII format, select the check box next to Is ASCII. The default is hex format.
3. (Optional) Click **Use At Runtime** to make this partition the active one for this job run.
4. Click **Apply**. The name of the new partition appears in the Row Partitions list, and the partition definition is displayed.

5. Repeat these steps to create another partition.

> **Note:**
>
> Only one override partition can be created per source and target object.

6. Click **OK** to close the partition editor.

Configuring run options

Use the Run Options page to override the run settings configured in the job profile. By default, all parameters are set to Use Profile.

To change a run option

1. Clear the **Use Profile** box.
2. Make the change under **Value**.

> **Note:**
>
> If a parameter is a toggle (enabled or disabled), a check mark under Value means that it is enabled.

3. Click **OK**.

> **Note:**
>
> Other processing options can be specified by going to the Profile Configuration page.

Parameter details

- General (No Comparisons Performed)

  Select only one of these items for any given run, as they are mutually exclusive.

  - **Review Previous Out-Of-Sync Results**: Compares only those rows that were out-of-sync in the previous run, based on the information that is stored in the Out-of-Sync file. Oracle GoldenGate Veridata will compare these rows against their current state. The results identify which rows were brought back into synchronization by replication or another method. By default, this parameter is disabled. Check Value to enable it.
  - **Generate Out-Of-Sync XML from Previous Run**: Generates an OOSXML file based on the OOS file from the previous run. It generates XML for every row from the OOS file. You can use the XML from this file for your own purposes, such as to view the out-of-sync information in an XML editor. By default, this parameter is disabled. Check Value to enable it. When this parameter is enabled, Review Previous Out-Of-Sync Results is ignored.

- Profile Setting Overrides
  - **Max Concurrent Comparison Threads**: Specifies the number of concurrent processing threads to use. You can use up to the same number of threads as there are processors on the server system.
  - **Delay Confirm-Out-Of-Sync By (seconds)**: Delays the confirmation step by the specified number of seconds to account for replication lag. Delaying the confirmation step reduces the number of false out-of-sync results that occur because an updated source value was not replicated fast enough.
  - **Perform Confirm-Out-Of-Sync Step**: Controls whether or not the confirmation step is performed. By default, it is performed. Clear the check box under Value to skip the confirmation step and only perform the initial (row hash) comparison. You might skip the confirmation step if, for example, activity on the source tables is quiesced or if replication is not continuously updating the target table(s).
  - **Out-Of-Sync Output Format**: Controls how out-of-sync row information is written to the OOS file on disk. The default is binary format, which is compatible with the Oracle GoldenGate Veridata Web User Interface browser. To change the output option, select one of the following under Value:
    * To output only in XML, select XML.
    * To output in both binary and XML formats, select Both.
    * To suppress out-of-sync output, select None.
  - **Maximum Size of Each Out-Of-Sync XML Chunk (Rows)**: Limits the number of out-of-sync rows that are written to an XML-based OOS report chunk. The OOS file is actually written as several files to prevent one file from becoming so large that it affects system resources. Using chunks enables periodic archiving or purging to remove aged files. The current file is closed when the specified number of rows is written, and a new file is opened.
- Tracing
  - **Trace Agent Rowhash**:Turns on tracing for the initial comparison step.
  - **Trace Agent COOS**: Turns on tracing for the confirmation step.
  - **Trace Server**: Turns on tracing for Oracle GoldenGate Veridata Server. Select or enter a value to set the desired trace level. The higher the level, the more detailed the trace data.

  By default, all tracing is off.

Viewing jobs that are running

Use the All Running Jobs page to view near-real-time information about jobs that are running. To access this page, click **Running Jobs** in the navigation pane.

Using the chart view

The chart view shows an aggregate status for all of the compare pairs that are contained in all of the jobs that are running.

Using the list view

The list view shows statistics and status conditions for the running comparisons. By default, the list is organized within the context of the job as a unit. You can filter the list to organize it by group or compare pair, and by comparison status.

- A quick way to change the context of the list is to click **View By Job** or **View By Group** or **View By Compare Pair** under the Running Jobs node in the navigation pane. The columns change context based on the selected node.
- To further refine the list, click any section of the Compare Pair Status chart to display jobs that contain compare pairs with the selected status. (See **Possible comparison status conditions**.) Click **View All Statuses** to view all jobs.
- To control the view even more, filter the list with Filters.

More on how to filter this list:

To see full filter options for the Running Jobs list, expand **Filters**.

- Basic filter options:
  - You can filter by:
    - Comparison Status: show all jobs (default) or constrain the list based on jobs with a specific comparison status (same as filtering with the Compare Pair Status chart).
    - job name
    - group name
    - compare pair name
- Advanced filter options
  1. Expand Advanced filter options for additional filter options.
  2. Under View List By, you can set the way that the list is displayed: by Job, Group, or Compare Pair. The view determines the context for the information that is in the columns of the list. This is the same as using the context nodes under Running Jobs.

     For example, with Group as the view, the Compare Pairs Processed column shows the number of compare pairs that are being processed for any given group. By contrast, if you change the view to Job, that same column would then show the number of compare pairs that are being processed for any given job.

     View By Job is the default view when you click **Running Jobs** in the navigation pane, but you can change that with user preferences (see Setting user preferences).
  3. Under List Columns to View, you can refine your chosen view by selecting or deselecting the names of columns to display.

     For example, you might want to limit the information to a basic overview of name, status, and number of out-of-sync rows. To get more information, you can always drill down to pages that contain more detail.

     The columns that you can work with depends on the view that you select.

Your column filters only apply to the active page. If you leave this page and then return to it, the column list reverts to the default for the view that is selected.

4. Click **Apply Filter** to display the selected view.

Drilling down

Click a name in the list to view details for a specific job, group, or compare pair, depending on the view that you selected.

Related Topics

Viewing jobs that are finished

Setting user preferences

Managing common tasks

Use the Favorites manager to store links to jobs and tasks that you work on frequently. Your favorites are stored in the Oracle GoldenGate Veridata repository as part of your user profile. Your favorites are available only to you and are visible on any page of the Oracle GoldenGate Veridata Web User Interface application at the lower left-hand side of the screen.

Like the way that favorites or bookmarks work in a web browser, clicking a Favorites link takes you directly to the underlying job or task without requiring navigation or filtering. For example, if you have a large number of tables to map, and you want to complete part of the task today and the rest tomorrow, simply save today's work to the repository and then add the mapping page to the shortcuts list. When you return tomorrow, click the link in Favorites instead of navigating to the task from the navigation pane. Similarly, if you run a job frequently, add it as a favorite job and then click the link in Favorites instead of running it from the Run/Execute Job page.

Using Favorite Jobs

Use the Jobs section of the Favorites manager to store links to jobs that you work with frequently. Only jobs can be added to this section. If you try to add a group or compare pair, the name of the container job is added to Favorites.

A favorite job can be added from the Edit Job page and the Details For Finished Job page.

You can set your Home page User Preference settings to only show these favorite jobs.

You will always see the latest data available for a favorite job. For example, if the job just finished running, you will see the same data that appears on the Finished Jobs page.

To add a favorite job:

You can add a favorite job when (and only when) you see an Add link in the Favorites manager next to Jobs.

To add a job to the Jobs list, click the **Add** link next to Jobs. The name of the job is then displayed in the Jobs section.

Using Shortcuts

Use the Shortcuts section of the Favorites manager to store links to a specific object, a previous job, a group, a report, or other workflow object.

You cannot add a job to the Jobs list and then also add it to the Shortcuts list. However, you can save a different run of that job to the Shortcuts list. When adding a finished job to the Shortcuts list, you can control whether the link goes to the latest run or a specific run from history.

To add a favorite shortcut

You can add a favorite shortcut when (and only when) you see an Add link in the Favorites manager next to Shortcuts. To add a shortcut to the Shortcuts list, take the following steps:

1. Click the **Add** link next to Shortcuts. The Add to Favorites dialog is displayed, with the name of the object appearing in the Task Data box.
2. Type a description in the Description box. Spaces are permissible. Important! This description is what will appear in the Shortcuts list, so the best practice is to keep it short.
3. (For a finished group, job, or compare pair only) Under Task Type, you can select one of the following:
   - To always see the latest status or data of a finished job, group, or compare pair, select **Always Use Latest Data**. Whenever the status of the object changes, the current state is shown (for example, the most recent run of a job). Note that this is the forced default for objects that are not a job, group, or compare pair.
   - To create a static snapshot of the status or data of a finished job, group, or compare pair as of the current date and time, select **Use Data At This Specific Point In Time**. Creating one or more snapshot views is useful for comparing one comparison run against another. You must give each snapshot of the same finished object a different name. For example, you can create one shortcut named ALL_3_31 and another shortcut named ALL_6_30 for the same job that is run at the end of March and at the end of June.

Editing Jobs and Shortcuts

You can edit the properties of a saved job or shortcut. The edit properties may differ slightly depending on which Favorite type it is.

To edit a favorite job or shortcut

1. In the Favorites manager, click **Edit** in the header bar of the Jobs or Shortcuts section that contains the object that you want to edit. The Favorites Task list is displayed. The Favorites column shows which type of Favorite the item is: Job or Shortcut.
2. If needed, expand Filter to filter the list based on one or both of the following attributes:
   - Use **Description** to filter based on the name of a job or shortcut. Do not use a wildcard.

- (Shortcuts only) Use Task Type Like to filter shortcuts based on whether the type is Always Use Latest Data or Use Data at This Specific Point in Time. When combined with the Description filter, the Task Type filter is a powerful way to immediately focus on all of the snapshots of one or more jobs, groups, or compare pairs that you saved as Shortcuts.

3. Do either of the following:
   - To edit a job or shortcut, click its name in the Description column of the list, and then make your changes on the Task Detail page.
   - To delete a job or shortcut, click in the Select column next to its name in the list, and then click **Delete**.

Editing a job

Use the Edit Job page to change the attributes of a job. To access this page,

1. In the navigation pane, click **Job Configuration**.
2. In the Existing Jobs list, click the name of the job, or select it in the Select column and then click **Edit**.

To edit job attributes

- **Name**: The job name cannot be changed.
- **Description**: The description can be changed.
- **Source Connection and Target Connection**: To change the source or target connection that is associated with this job, select the connection from the drop-down list, or click the **Browse** button to find it using the Select a Connection page, which provides filtering options. The connection override applies only to this job.

> **Note:**
>
> The connection section is optional. Setting the connection(s) here will override the connection information that is specified at the group level, and all of the groups in the job will use those connections when the job is run. Do not set connections here if some groups require different metadata connections than others.

- **Profile Information**: To change the profile that is associated with this job, select the profile from the drop-down list, or click the **Browse** button to find it by using the Select a Profile page, which provides filtering options. If no profile is selected, the default profile will be used.

To link or unlink groups

The Linked Groups list shows the groups that exist in the repository and are available to be linked to this job. It also shows the connections that will be used by each group, and the number of compare pairs that each group contains.

The list header shows the total number of groups that are available.

Use the selection boxes in the Link column to:

- Add a group to this job by checking the box next to its name.

- Remove a group from this job by clearing the box next to its name.

The changes will be active for the next run of this job.

To filter the list:

To filter the Linked Groups list, expand Filters. The list can be filtered to show:

- **All** shows all groups that are available in the repository.
- **Only Enabled** shows only those groups that are linked to this job.
- **Only Disabled** shows only those groups that are not linked to this job.

The list also can be filtered to show or exclude group names that match a specified string.

Click **Apply Filter** to display the selected groups.

Click **Save** to save the changes to the repository.

Related Topics

Working with Jobs

Viewing jobs that are finished

Use the All Finished Jobs page to view near-real-time information about jobs that are finished being processed. To access this page, click **Finished Jobs** in the navigation pane.

By default, this page shows all finished jobs. Use the Time Range Filter to constrain the view as needed.

To use the Time Range Filter

- **Most Recent Comparison Run**: Selects the most recent run for each job that exists.
- **View Last**: Selects all of the runs of all of the jobs that started within the last <n> Hours or Minutes.
- **From**: Selects all of the runs of all of the jobs that finished within a specific date range, with the option to include the time of day. Supply date and time formats in GMT (Greenwich Mean Time).

Click **Apply Filter** to activate your selection. The filter selects both cancelled and completed jobs.

Using the chart view

The chart view shows an aggregate status for all of the compare pairs in all of the finished jobs that you selected with the Time Range Filter.

Using the list view

The list view shows statistics and status conditions for the comparisons that finished within the selected time range. By default, the list is organized within the context of the job as a unit. You can filter the list to organize it by group or compare pair, and by comparison status (See Possible comparison status conditions.).

- A quick way to change the context of the list is to click **View By Job**, **View By Group**, or **View By Compare Pair** under the Finished Jobs node in the navigation pane. The columns change context based on the selected node.
- To further refine the list, click any section of the Compare Pair Status chart to display jobs that contain compare pairs with the selected comparison status, such as Out-Of-Sync. Click **View All Statuses** to view all jobs.
- To control the view even more, filter the list with Filters.

To see full filter options for the Finished Jobs list, expand **Filters**.

Basic filter options

You can filter by:

- **Comparison Status**: show all jobs (default) or constrain the list based on jobs with a specific comparison status (see Possible comparison status conditions)
- **Job Name**
- **Group Name**
- **Compare Pair Name**

Advanced filter options

Expand Advanced filter options for additional filter options.

1. Under **View List By**, you can set the way that the list is displayed: by Job, Group, or Compare Pair. The view determines the context for the information that is in the columns of the list.

   For example, with Group as the view, the Compare Pairs Processed column shows the number of compare pairs that were processed for any given group. By contrast, with Job as the view, that same column shows the number of compare pairs that were processed for any given job.

   You can set any of these views as the default by setting user preferences (see Setting user preferences).

2. Under **List Columns to View**, you can refine your chosen view by selecting or deselecting the names of columns to display.

   For example, you might want to limit the information to a basic overview of name, status, and number of out-of-sync rows. To get more information, you can then drill down to pages that contain more detail.

   The columns that you can work with depends on the view that you select.

   Your column filter only applies to the active page. If you leave this page and then return to it, the column list reverts to the default for the view that is selected.

Click **Apply Filter** to display the selected view.

Possible comparison status conditions

Drilling down

Click a name in the list to view details for a specific job (see Viewing details of a job that is finished), group (see Viewing details of a group that is finished), or compare pair, depending on the view that you selected.

Viewing Reports

Reports show processing results and the causes of errors. Click in the Reports column of the list to view the report for a specific job, group, or compare pair.

The report is a text file that is stored on disk. The default location is `domain_home/veridata/reports`. This location is dependent on the `server.veridata_data` property defined in the `domain_home/config/veridata/veridata.cfg` file.

Related Topics

Setting user preferences

Viewing jobs that are running

Using the Comparison Report

Viewing details of a job that is finished

Use the Details For Finished Job page to view information about a specific run of a finished job. To access this page:

1. Under Finished Jobs in the navigation pane, click **View by Job**.
2. In the Finished Jobs list, click the job run that you want to view.

It shows:

- The name and description of the job
- The start and end times of this run
- The duration of this run (how long it took to be processed)

Using the chart view

The chart view shows the aggregate status of all of the compare pairs in all of the groups that were processed in this job run.

Using the list view

The list view initially shows all of the groups that were processed in this job run. You can filter the list by using the Filters option.

To filter the list:

- To filter the list based on a specific outcome for the compare pairs in this job, click that section of the Compare Pair Status chart. (See ***Possible comparison status conditions***.)
- You can get this same filter by expanding Filters and using Show Groups With Comparison Status.
- Within Filters, you can also filter by group name. Supply a search string for Show Group Name Like. All group names that contain this string are displayed.

Click **Apply Filter** to activate the filter.

Viewing reports

To view the comparison reports for any group in this job run, click **Report** in the row of that group.

Drilling down

Click a group name in the list to view the Details for Finished Group page. This page provides results at the group level.

Other actions from here

To edit this job, click **Job Configuration** at the top of the page to go to the Edit Job page.

To run the job from this page, click **Run...** at the top of the page to go to the Run/Execute Job page.

Purging Reports

To purge old reports, compare results, and obsolete data, use the **Delete** option in the Finished Jobs list view. In the Finished Jobs list, select the job(s) for which you want to purge reports by clicking the Select check box and then click **Delete**.

When you purge job reports, select the optional **Delete Directories completely** check box to delete the entire reports directory content, including files and directories created by the user.

By default, the **Delete Directories completely** check box is not selected and therefore only the files created by Veridata are purged. The directories are purged only if they are empty after the Veridata files are purged. Files created by users within the reports directories will not be deleted.

> **Note:**
>
> Purging can only be done for finished jobs. You cannot purge reports for compare pairs or groups.

Related Topics

Setting user preferences

Viewing details of a group that is finished

Use the Details For Finished Group page to view information about a specific run of a finished group. To access this page:

1. Under Finished Jobs in the navigation pane, select **View By Group**.
2. Under Group Name, click the name of the group.

It shows:

- The name and description of the group
- The start and end times of this run

- The duration of this run (how long it took to be processed)

Host Connection Information shows the source and target connections that were used for this run of the group.

Using the chart view

The chart view shows the aggregate status of all of the compare pairs that were processed for this group in this run.

Using the list view

The list view initially shows all of the compare pairs that were processed for this group in this run. You can filter the list by using the Filters option.

To filter the list:

1. Expand Filters to show filter options. You can filter by the following attributes:
   - **Comparison Status**: select from the Show Compare Pairs With Status list to show all of the compare pairs (default) in the group or to constrain the list based on those with a specific comparison status. You can get this same filter by clicking any status within the Compare Pair Status chart. (See Possible comparison status conditions.)
   - **Show Compare Pair Name Like**: filters based on a string in compare pair names.
   - **Show Compare Pairs Where Amount of OOS Rows**: select from the list to show only those compare pairs that contain an amount of out-of-sync rows that is greater than, smaller than, or equal to a specific value. Using this filter helps you focus on tables or files that have the most out-of-sync rows.
2. Click **Apply Filter** to activate the filter.

Viewing reports

To view the comparison report for a compare pair in this run, click **Report** in the Reports column of the list.

To view the out-of-sync report for a compare pair, click **View** in the OOS Rows column in the list. This link is only shown if a compare pair has out-of-sync rows.

Drilling down

Click a compare pair in the list to view the Details for Finished Compare Pair page (see Viewing details for a compare pair that is finished).

Other actions from here

To edit this group, click **Group Configuration** at the top of the page to go to the Edit Group page.

To run the job that is associated with this group, click **Run** at the top of the page to go to the Run/Execute Job page.

Viewing details for a compare pair that is finished

Use the Details for Finished Compare Pair page to view near-real-time information about a compare pair that is finished being processed.

The upper portion of the page shows specific information about the compare pair, including:

- **Compare Pair Name**: the name of the compare pair.
- **Source Table and Target Table**: the names of the underlying source and target objects.
- **Compare Pair Run ID**: the identifier used by Oracle GoldenGate Veridata Server while running a comparison. It is useful for locating messages related to this comparison in the veridata.log file, which is used for troubleshooting purposes when resolving support cases.
- **Compare Pair Start Time**: the time when processing started for this compare pair.
- **Compare Pair Run Duration**: how long it took to process this compare pair.
- **Run Status**: the outcome of this comparison (Waiting, Canceled, Canceling, Running, Finished)
- **Comparison Status**: the status of the comparison (see Possible comparison status conditions.) If delta processing is enabled for this compare pair, the delta symbol is displayed.
- **Overall Rows Compared**: the total number of rows in the source and target objects that were compared during this run of the compare pair.
- **Overall Rows With OOS**: the number of rows in the source and target objects that are out-of-sync.

Using the charts

The charts on this page help you to determine what types of operations are causing the most out-of-sync conditions and which phases of processing are taking the most time to complete.

Overall Rows Compared pie chart

Plots the result from the Overall Rows Compared field according to the possible outcomes of the comparison: In-Sync or Out-Of-sync.

Operations Out-Of-Sync pie chart

Plots the number of out-of-sync rows according to their operation type:

- **Inserts**: source table has the row, but the target does not.
- **Updates**: source table has different row values from the target.
- **Deletes**: target table has the row, but the source does not.

Compare Pair Performance table

Lists performance statistics based on each phase of a run: the sorting phase (for source and target), the initial comparison phase, and the confirmation phase. Note: sorting statistics are only available if the job profile specifies server-side sorting (row sorts are performed by Oracle GoldenGate Veridata, not by the database engines).

- **Start Time**: The time that a run phase started.
- **Run Duration**: The amount of time that it took to complete a run phase.
- **Rows Processed**: The number of rows that were processed in a run phase.

- **Rows Per Second**: The performance rate of a run phase, in terms of rows processed.
- **Bytes Processed**: The overall volume of data, in bytes, that was processed in a run phase.
- **Rows Per Second**: The performance rate of a run phase, in terms of bytes processed.

Compare Pair Performance History chart

Plots the performance of each of the run phases as a bar chart. To use this chart:

1. From View Step, select a run phase to plot.
   - Initial comparison step: Select Initial Compare.
   - Confirmation step: Select Confirm Out-Of-Sync.
   - Sorting process: Select Sorting.
2. From View Performance History By, select the criterion for this performance analysis. This sets the units of measure for the vertical axis of the chart. You can measure:
   - Number of rows or bytes that were processed
   - Number of rows or bytes per second that were processed
   - Run duration (the time spent processing the rows)

Other actions from here

To view a report of the results for this run of this compare pair, click **View Comparison Report**.

To view details about out-of-sync rows, including column values, click **View Out-Of-Sync Rows**. This link only appears if the compare pair has out-of-sync rows.

To edit this compare pair, click **Compare Pair Configuration** at the top of the page to go to the Compare Pair Configuration page.

Repairing Out-Of-Sync Jobs

From the Finished Jobs page, you can select the out-of-sync comparisons for repair. Jobs, groups, and compare pairs can be selected for repair. To repair an out-of-sync comparison, select the job from the finished jobs list and click **Repair**. The Details for Out-Of-Sync page appears. See Viewing Details of an Out-Of-Sync Job.

The Repair Jobs page displays a summary of all repair jobs. You can use the Filters on this page to display older repair jobs and to filter repair jobs by repair status and job name. Following are the possible repair statuses.

Filter	Description
	Cancelling: Indicates that a cancel has been requested, but the item has not stopped processing.

Filter	Description
	Running: Indicates that an active repair job is running.
	Pending: Indicates that a repair request is pending. This status appears for repair groups and repair pairs within a running repair job.
	Errors: Indicates that the repair job is completed with an error.
	Cancelled: Indicates that the repair job operation was cancelled.
	Warnings: Indicates that one or more rows could not be repaired.
	Out-Of-Sync: Indicates that the row was out-of-sync in the base comparison and no repair has been attempted.
	Successful: Indicates that a repair operation has been successfully applied.

Viewing Details of an Out-Of-Sync Job

Click a finished comparison job from the Repair Jobs list on the Repair Jobs page to view the details of the out-of-sync job. The Details of out-of-sync job page contains all information about a finished comparison job that is finished with one or more out-of-sync compare pairs. When the page is initially displayed, a summary of the out-of-sync job and the list of out-of-sync Groups belonging to the out-of-sync Job are displayed. The out-of-sync Group items can be expanded to display the list of out-of-sync compare pairs. Each out-of-sync compare pair item can be expanded to display the summary information and the list out-of-sync rows.

> **Note:**
>
> Example: If there are duplicate rows in Source and Target tables and if the tables are out-of-sync then, in the **Out-of-Sync** page, the **Operation** column in the **Out-of-Sync-Row Details** table is displayed with an *(asterisk). Oracle GoldenGate Veridata doesn't provide any support to avoid duplicate rows. With Auto Mapping, the chances of occurrences of duplicate rows are more when no primary key exists on Source and Target tables.

Only authorized users will be able to initiate repair jobs from this page. If a repair job is running for the base comparison job, this page displays the current status for the running repair job and will periodically refresh itself. After one or more repair jobs have been run against a base comparison, this page displays the cumulative statistics for all of the repair jobs.

Viewing the Repair Report

To view and download the repair report, click the link under the Report column for a repair job on the Repair Jobs page.

Using the Comparison Report

To view a comparison report, use the Reports page.

A comparison report is generated for each job, group, and compare pair that is finished being processed. It contains summary details about out-of-sync row counts, the number of records processed, performance statistics, errors, and so forth.

The comparison report tells you how extensive an out-of-sync problem is. It also provides performance statistics and, optionally, column details. The comparison report can be viewed by any user role.

To select a time range

1. Under Report Period, select one of the following view choices:
   - **View Most Recent Comparison Run**: Selects the most recent report for each job that exists.
   - **View Last**: Selects reports that were created within the last <n> Days, Hours or Minutes.
   - **View From**: Selects reports that were created within a specific date range, with the option to include the time of day. Supply dates and times in the format that is accepted by the client machine that is running Oracle GoldenGate Veridata Web User Interface.
2. Click **Apply**.

To filter the selection

1. Expand Filter Reports.
2. Type your search criteria in one of the following fields:

- **Job Like**: Retrieves reports for jobs whose names contain a string that matches the one supplied.
- **Group Like**: Retrieves reports for groups whose names contain a string that matches the one supplied.
- **Compare Pair Like**: Retrieves reports for compare pairs whose names contain a string that matches the one supplied.
- **Show Report Names Like**: Retrieves reports whose file names contain a string that match the one supplied.

3. Click **Apply Filter**.

To view a report

To view a report, click the name of the report in the Report Name column of the Existing Reports list.

Alternately, you can click the **Select** button in the list next to the name of the report, and then click **View**. The Report View page will be displayed.

Viewing the report file on disk

The report is a text file that is stored on disk. It can be viewed from its location on disk.

4
Customizing Your Workspace

Learn how to manage your work from the home page, including how to save important links using the favorites manager, and set user preferences.

- Working with the Home Page
- Using the Favorites Manager
- Setting user preferences

Working with the Home Page

The Home page is the first page that you see when you log in to the Oracle GoldenGate Veridata Web User Interface. It is your personalized view of the comparison jobs that are running or finished, from which you can drill down to view more specific information. The information that is displayed on this page is controlled by the Home Page user preferences (see Setting home page preferences).

View options

Under View, you can toggle between **Running Jobs** and **Finished jobs**.

To refresh the page, click the **Refresh** button. You can set the default rate for automatic refreshes by Setting home page preferences

Chart views

If Running Jobs is the Home page view, the chart shows an aggregate status of all compare pairs in jobs that are running. If Finished Jobs is the Home page view, the chart view shows all compare pairs in jobs that finished within the time interval that is specified with home page preferences.

Drilling down from the list view

Depending on whether you select Running Jobs or Finished Jobs, the list view contains:

- Times and statistics for each running or finished job.
- Node markers (+) that can be expanded to view at-a-glance details of a specific group and compare pair.
- Links from each listed object that go to the details page for that object.

To view all running jobs (see Viewing jobs that are running) or all finished jobs (see Viewing jobs that are finished), click the link at the bottom of the page.

Setting home page preferences

The Home user preferences control the appearance of the Home page. This page is the initial view that you see when you log into Oracle GoldenGate Veridata.

Default View

The default view on the Home page can be set to show either finished jobs or running jobs.

The default view is Finished Jobs. When used, this parameter makes the Show Jobs Finished (units) and Show Jobs Finished (unit time range) settings active (see below), allowing you to specify a time range for which finished jobs are shown.

The Running Jobs view shows a list of all running jobs, which cannot be filtered by a time range.

Show Only Favorite Jobs

Filters the Home page to show only the favorite jobs that are configured in the Favorites panel for this user.

Show Jobs Finished (units)

Activates when Finished Jobs is selected for Default View. It sets the unit of measure for Show Jobs Finished (unit time range). Select Minutes, Hours, or Days from the drop-down list.

Show Jobs Finished (unit time range)

Activates when Finished Jobs is selected for Default View. It constrains the number of finished jobs on the Home page to a specific number of the units selected in Show Jobs Finished (units). For example if the unit is Days, a value of 2 shows all of the jobs that finished running in the past two days.

Back to Setting user preferences

Using the Favorites Manager

Use the Favorites manager to store links to jobs and tasks that you work on frequently. Your favorites are stored in the Oracle GoldenGate Veridata repository as part of your user profile. Your favorites are available only to you and are visible on any page of the Oracle GoldenGate Veridata Web User Interface application at the lower left-hand side of the screen.

Like the way that favorites or bookmarks work in a web browser, clicking a Favorites link takes you directly to the underlying job or task without requiring navigation or filtering. For example, if you have a large number of tables to map, and you want to complete part of the task one day and the rest some other day, save that day's work to the repository and then add the mapping page to the shortcuts list. The next day, click the link in Favorites instead of navigating to the task from the navigation pane. Similarly, if you run a job frequently, add it as a favorite job and then click the link in Favorites instead of running it from the Run/Execute Job page.

Using Favorite Jobs

Use the Jobs section of the Favorites manager to store links to jobs that you work with frequently. Only jobs can be added to this section. If you try to add a group or compare pair, the name of the container job is added to Favorites.

A favorite job can be added from the Edit Job page and the Details For Finished Job page.

You can set your Home page User Preference settings to only show these favorite jobs.

You will always see the latest data available for a favorite job. For example, if the job just finished running, you will see the same data that appears on the Finished Jobs page.

To add a favorite job:

You can add a favorite job when (and only when) you see an Add link in the Favorites manager next to Jobs.

To add a job to the Jobs list, click the **Add** link next to Jobs. The name of the job is then displayed in the Jobs section.

Using Shortcuts

Use the Shortcuts section of the Favorites manager to store links to a specific object, a previous job, a group, a report, or other workflow object.

You cannot add a job to the Jobs list and then also add it to the Shortcuts list. However, you can save a different run of that job to the Shortcuts list. When adding a finished job to the Shortcuts list, you can control whether the link goes to the latest run or a specific run from history.

To add a favorite shortcut

You can add a favorite shortcut when (and only when) you see an Add link in the Favorites manager next to Shortcuts. To add a shortcut to the Shortcuts list, take the following steps:

1. Click the **Add** link next to Shortcuts. The Add to Favorites dialog is displayed, with the name of the object appearing in the Task Data box.
2. Type a description in the **Description** box. Spaces are permissible. Important! This description is what will appear in the Shortcuts list, so the best practice is to keep it short.
3. (For a finished group, job, or compare pair only) Under Task Type, you can select one of the following:
   - To always see the latest status or data of a finished job, group, or compare pair, select **Always Use Latest Data**. Whenever the status of the object changes, the current state is shown (for example, the most recent run of a job). Note that this is the forced default for objects that are not a job, group, or compare pair.
   - To create a static snapshot of the status or data of a finished job, group, or compare pair as of the current date and time, select **Use Data At This Specific Point In Time**. Creating one or more snapshot views is useful for comparing one comparison run against another. You must give each snapshot of the same finished object a different name. For example, you can create one shortcut named ALL_3_31 and another shortcut named ALL_6_30 for the same job that is run at the end of March and at the end of June.

Editing Jobs and Shortcuts

You can edit the properties of a saved job or shortcut. The edit properties may differ slightly depending on which Favorite type it is.

To edit a favorite job or shortcut

1. In the Favorites manager, click **Edit** in the header bar of the Jobs or Shortcuts section that contains the object that you want to edit. The Favorites Task list is displayed. The Favorites column shows which type of Favorite the item is: Job or Shortcut.
   - If needed, expand Filter to filter the list based on one or both of the following attributes:
   - Use Description to filter based on the name of a job or shortcut. Do not use a wildcard.
   - (Shortcuts only) Use Task Type Like to filter shortcuts based on whether the type is Always Use Latest Data or Use Data at This Specific Point in Time. When combined with the Description filter, the Task Type filter is a powerful way to immediately focus on all of the snapshots of one or more jobs, groups, or compare pairs that you saved as Shortcuts.
2. Do either of the following:
   - To edit a job or shortcut, click its name in the Description column of the list, and then make your changes on the Task Detail page.

     or
   - To delete a job or shortcut, click in the Select column next to its name in the list, and then click **Delete**.

Setting user preferences

Use the User Preferences page to customize certain aspects of your Oracle GoldenGate Veridata environment by customizing default parameter settings. To access this page, click **User Preferences** under Options/Settings in the navigation pane.

You can change the following User Preference settings:

- General (see Setting general user preferences)
- Home (see Setting home page preferences)

Setting general user preferences

The General user preferences control basic interface settings. To access the General user preferences:

1. Under Options/Settings in the navigation pane, click **User Preferences**.
2. Click **General**.

Poll/Refresh interval

Controls the refresh rate, in seconds, of the statistics on the Running Jobs page. The default is to refresh every 30 seconds.

Default Number of Items to Show Per Page

Controls the number of items that are displayed in the lists of the Oracle GoldenGate Veridata Web User Interface, for example the lists of existing groups. By default, 10 list items are displayed.

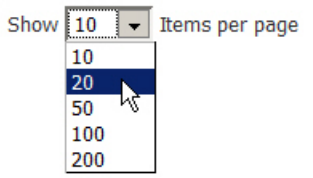

You can override this parameter's setting at any time by using the Show drop-down menu on any page that displays lists.

Show charts

Controls whether or not the pie chart is displayed by default on pages that show comparison status, such as the Finished Jobs pages. By default, charts are displayed.

You can override this parameter's setting at any time by clicking the **Show Chart**

or **Hide Chart**

control on a page that includes a chart option.

Finished Jobs View and Running Jobs View

These parameters control whether the Finished Jobs Overview and Running Jobs Overview pages display their results in terms of jobs, groups, or compare pairs.

- When set to the default of Job, only the job names are displayed. The Start Time and Duration columns apply to each job as a whole.
- When set to Group, job names and group names are displayed. The Start Time and Duration columns apply to each group as a whole.
- When set to Compare Pair, job names, group names, and compare pair names are all displayed. The Start Time and Duration columns apply to each compare pair.

Charts are not affected by these parameters.

Back to Setting user preferences

Setting home page preferences

The Home user preferences control the appearance of the Home page. This page is the initial view that you see when you log into Oracle GoldenGate Veridata.

To access the Home user preferences:

1. Under Options/Settings in the navigation pane, click **User Preferences**.
2. Click **Home**.

Default View

The default view on the Home page can be set to show either finished jobs or running jobs.

The default view is Finished Jobs. When used, this parameter makes the Show Jobs Finished (units) and Show Jobs Finished (unit time range) settings active (see below), allowing you to specify a time range for which finished jobs are shown.

The Running Jobs view shows a list of all running jobs, which cannot be filtered by a time range.

Show Only Favorite Jobs

Filters the Home page to show only the favorite jobs that are configured in the Favorites panel for this user.

Show Jobs Finished (units)

Activates when Finished Jobs is selected for Default View. It sets the unit of measure for Show Jobs Finished (unit time range). Select **Minutes**, **Hours**, or **Days** from the drop-down list.

Show Jobs Finished (unit time range)

Activates when Finished Jobs is selected for Default View. It constrains the number of finished jobs on the Home page to a specific number of the units selected in Show Jobs Finished (units). For example if the unit is Days, a value of 2 shows all of the jobs that finished running in the past two days.

5
Tuning Performance

Learn ways to monitor and improve performance when processing large volumes of data.

- Improving the Performance of Oracle GoldenGate Veridata
- Performance Statistics

Improving the Performance of Oracle GoldenGate Veridata

The following are some of the factors that influence the performance of Oracle GoldenGate Veridata and some ways you can improve its performance when processing large volumes of data.

Database and network use

Two critical performance factors for Oracle GoldenGate Veridata are:

- How data is sorted
- How the data is sent across the network

Performance Statistics for these performance factors are printed to the comparison report for each finished comparison and are recorded for the initial comparison step on the source and target (see Using the Comparison Report).

Network use

Oracle GoldenGate Veridata automatically optimizes its use of the network by using hashing and network message compression for efficient data transfer. The greater the size in bytes of the average row of a table or file (see the bytes/row performance statistic), the greater the rate of compression that is achieved through hashing (see the rh bytes/row and hash comp rate performance statistics). Whether a row is 50 bytes or 1000, the number of bytes that are used to represent its non-key values will be 12. Therefore, as a percentage of table size (in bytes), larger rows tend to use the network more efficiently. For those same reasons, the smaller the key size relative to row size, the more efficient the use of the network.

Additionally, on the NonStop platform check the send and receive TCP/IP buffer sizes. They should be set to 32K for Oracle GoldenGate Veridata.

Database access

By default, Oracle GoldenGate Veridata uses the database to sort data for comparison. The following factors affect the performance of the database sorting mechanism:

- The number of rows in the tables being compared
- The indexes that are defined on the tables
- The keys that are being used

- The way that the database is tuned

After some test runs, if the performance of comparisons is not satisfactory, it might be faster to use server-side sorting, where Oracle GoldenGate Veridata Server itself performs the sorting.

Configuration options

The following are some other things you can consider to improve comparison performance.

Partition large tables

You can divide large source and target tables into row partitions, each partition being associated with a different row subset. Row partitions enable you to process sets of data in parallel and also to control the timing of processing. For example, you can compare one partition today and the other one tomorrow. In addition, the results of a subset comparison can give you an idea of the synchronization status of the entire table.

Exclude columns

If a table contains columns that you know will never change, or if it does not matter whether those columns are in-sync, you can exclude those columns from the comparison to reduce the processing load. You exclude columns when you create or edit compare pairs.

Use delta processing

You can configure compare pairs to use delta processing, a performance feature whereby Oracle GoldenGate Veridata only compares data blocks that have changed, instead of comparing all of the rows in a table or file. For more information about delta processing, see Using Delta Processing.

Change the database transaction isolation level

Each Oracle GoldenGate Veridata agent has an agent.properties file in the root of its installation folder that contains environment parameters. One of those parameters is database.transaction.isolation. This property controls the transaction isolation level that is used during the initial comparison step. The default value for SQL Server and Teradata is READ_UNCOMMITTED. This means that the query that is issued to select rows can read rows that have been modified by other transactions but not yet committed (dirty reads). It does not issue shared locks to prevent other transactions from modifying the data, nor is it blocked by the locks of other transactions.

The advantage of using the READ UNCOMMITTED option is that the initial reads of the data are faster because they are not affected by locking. The negative impact of this is that more records could be labeled as possibly out-of-sync (because the data can change over the course of the transaction) and would need to be compared again during the confirmation step. If you think that there are too many rows being compared in the confirmation step, you can edit the properties file and set database.transaction.isolation to COMMITED, which only permits the fetching of committed records. You must weigh any improvement against the possibility that the initial comparison step becomes slower due to the affect of locks to preserve read consistency.

> **Note:**
>
> (The only value that is supported for Oracle is READ_COMMITTED, and the confirmation step always uses READ_COMMITTED, because at this stage dirty reads are not acceptable.)

Profile options

You can control certain parameters for the initial and confirmation steps that can help increase performance.

Initial comparison step parameters

- **Limit the number of rows that are compared**: By using the Limit Number of Input Rows parameter for a specific job profile, you can constrain the number of rows that are fetched for processing. This enables you to process a smaller number of rows to get an idea of how out-of-sync (or not) the entire table is. Based on the results, you can make a decision about whether to run a complete comparison or just resynchronize the data. The Limit Number of Input Rows parameter is a General parameter for the initial comparison process.
- **Increase process priority (NonStop only)**: Assign the Oracle GoldenGate Veridata Agent the highest process priority possible on a NonStop system. You can assign a priority (as well as a process name and CPU number) by using the NonStop settings of the initial and confirmation steps in the job profile.
- **Increase processing threads**: The default is four threads. If the machine where Oracle GoldenGate Veridata Server is running has more processors, you can change the value of the Max Concurrent Comparison Threads parameter accordingly, keeping all threads busy by executing simultaneous comparisons. This parameter is in the General profile settings of the initial comparison process.

Confirmation step parameters

- **Run each comparison step separately**: By default, Oracle GoldenGate Veridata runs the initial compare and confirmation processes concurrently. If you run them in sequence, fewer system resources are used, but it will take longer to get results. This functionality is controlled by the Run Concurrently with Initial Compare parameter of the comparison step's General profile settings.
- **Skip the confirmation step**: The default is to always perform a confirmation step. You can skip this step if the database is quiesced or if replication is not actively replicating data changes. Use the Perform Confirmation Step parameter of the confirmation step's General profile settings.
- **Increase process priority (NSK)**: Assign the Oracle GoldenGate Veridata Agent the highest process priority possible on a NonStop system. You can assign a priority (as well as a process name and CPU number) by using the NonStop settings of the initial and confirmation steps in the job profile.

Connection options

Try increasing the batch size of fetched rows to increase throughput. To do this, increase the size of the Initial Compare Fetch Batch Size parameter for the initial and confirmation steps.

Performance Statistics

The two most critical aspects of Oracle GoldenGate Veridata performance are database access and network usage. Performance statistics for both of these aspects are printed to the comparison report for each comparison performed and are recorded for the initial comparison step on the source and target systems. The following describes these statistics. Depending on the results of these statistics, there are ways to optimize database access and network usage.

duration

The time spent processing the fetched rows.

rows fetched

The number of rows fetched from the database.

rows/sec

The number of rows processed per second.

bytes fetched

The total number of bytes that were processed.

bytes/sec

The number of rows, in terms of bytes, processed per second.

lob chunks fetched

The number of 32k blocks of LOB data fetched.

batches fetched

The number row batches fetched. The default is 10 rows per batch.

ipc msgs

The number of interprocess messages between the server and agent processes.

ipc bytes

The number of bytes transferred between the server and agent processes.

bytes applied

The number of bytes per message that were applied.

lob chunks applied

The number 32k byte LOB chunks applied at the target database.

lob fetch time duration (secs)

The amount of time spent fetching LOB data.

batches applied

The total number of bytes that were processed.

transaction batches

The number of transactions used to apply the data.

transaction single statements

the number of single row transactions applied during error recovery.

A

Column Mapping

Oracle GoldenGate supports the Hive column mapping to various datatypes, such as Oracle, SQL Server.

This appendix contains the following topics:

- SQL Server to Hive Mapping
- Oracle to Hive Mapping

Oracle Data Type to Hive Mapping

Oracle Datatype	Oracle GoldenGate Veridata Format	Hive Datatype
CHAR, NCHAR	String	CHAR
VARCHAR, NVARCHAR	String	VARCHAR
LONGRAW, BLOB	Blob	BINARY
RAW	Binary	BINARY
FLOAT, BINARY_FLOAT	Float	FLOAT
BINARY_DOUBLE	Binary	DOUBLE
DATE	Date	DATE
TIMESTAMP	Timestamp	TIMESTAMP
LONG, CLOB, NCLOB	Clob	STRING

SQLServer to Hive Mapping

SQLServer Datatype	Oracle GoldenGate Veridata Format	Hive Data type
BIGINT, BIT, INT, SMALLINT TINYINT	numeric	TINYINT, SMALLINT, INT, BIGINT
DECIMAL, MONEY, SMALLMONEY	dec_float	DECIMAL
FLOAT	float	FLOAT

SQLServer Datatype	Oracle GoldenGate Veridata Format	Hive Data type
REAL	float	DOUBLE
CHAR, NCHAR,	string	CHAR
VARCHAR, NVARCHAR,	string	VARCHAR STRING
TEXT, NTEXT	clob	STRING
BINARY, VARBINARY	binary	BINARY
IMAGE	blob	BINARY
DATE	date	DATE
DATETIME, DATETIME2, SMALLDATETIME, DATETIMEOFFSET	timestamp	TIMESTAMP
TIME	time	TIMESTAMP

Index

P

port

port *(continued)*
web component, *1-2*