Informatica®

Informatica® Cloud Data Integration

Redis Connector

Informatica Cloud Data Integration Redis Connector
February 2023

© Copyright Informatica LLC 2018, 2023

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

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

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

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

See patents at https://www.informatica.com/legal/patents.html.

DISCLAIMER: Informatica LLC provides this documentation "as is" without warranty of any kind, either express or implied, including, but not limited to, the implied warranties of noninfringement, merchantability, or use for a particular purpose. Informatica LLC does not warrant that this software or documentation is error free. The information provided in this software or documentation may include technical inaccuracies or typographical errors. The information in this software and documentation is subject to change at any time without notice.

NOTICES

This Informatica product (the "Software") includes certain drivers (the "DataDirect Drivers") from DataDirect Technologies, an operating company of Progress Software Corporation ("DataDirect") which are subject to the following terms and conditions:

1. THE DATADIRECT DRIVERS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.
2. IN NO EVENT WILL DATADIRECT OR ITS THIRD PARTY SUPPLIERS BE LIABLE TO THE END-USER CUSTOMER FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR OTHER DAMAGES ARISING OUT OF THE USE OF THE ODBC DRIVERS, WHETHER OR NOT INFORMED OF THE POSSIBILITIES OF DAMAGES IN ADVANCE. THESE LIMITATIONS APPLY TO ALL CAUSES OF ACTION, INCLUDING, WITHOUT LIMITATION, BREACH OF CONTRACT, BREACH OF WARRANTY, NEGLIGENCE, STRICT LIABILITY, MISREPRESENTATION AND OTHER TORTS.

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

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

Publication Date: 2023-04-04

Table of Contents

Preface 4

Informatica Resources 4

- Informatica Documentation 4
- Informatica Intelligent Cloud Services web site 4
- Informatica Intelligent Cloud Services Communities 4
- Informatica Intelligent Cloud Services Marketplace 4
- Data Integration connector documentation 5
- Informatica Knowledge Base 5
- Informatica Intelligent Cloud Services Trust Center 5
- Informatica Global Customer Support 5

Chapter 1: Introduction to Redis Connector 6

Redis Connector assets 6

Administration of Redis Connector 6

Chapter 2: Connections for Redis 7

Redis connection properties 7

Chapter 3: Mappings for Redis 9

Redis sources in mappings 9

- Adding and removing source fields 10

Redis targets in mappings 10

Redis key types 11

Edit metadata in a Redis source and target 12

Reading and writing hierarchical data in mappings 12

Partitioning 12

- Fixed partitioning for Redis sources 13

Rules and guidelines for Redis mappings 13

Chapter 4: Data type reference 14

Redis and transformation data types 14

Index 15

Preface

Use Redis Connector to learn how to read from or write to Redis by using Data Integration. Learn to create a Redis connection, run mappings in advanced mode and tasks in Data Integration.

Informatica Resources

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

Informatica Documentation

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

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

Informatica Intelligent Cloud Services web site

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

Informatica Intelligent Cloud Services Communities

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

Access the Informatica Intelligent Cloud Services Community at:

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

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

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

Informatica Intelligent Cloud Services Marketplace

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

https://marketplace.informatica.com/

Data Integration connector documentation

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

Informatica Knowledge Base

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

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

Informatica Intelligent Cloud Services Trust Center

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

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

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

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

Informatica Global Customer Support

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

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

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

CHAPTER 1

Introduction to Redis Connector

You can use Redis Connector to read data from or write data to Redis.

You can use Redis Connector to read and write Redis key types HASH, LIST, and STRING. You can configure Redis Connector to read HASH and LIST keys as one row per key or as multiple rows per key to enable reading and writing Redis keys with large values.

You can use a Redis connection as a source and as a target only in mappings in advanced mode. When you run a mapping in advanced mode or a task based on a mapping in advanced mode, Data Integration analyzes the data logic in the mapping to automatically assign data logic to run on an advanced cluster.

Redis Connector assets

Create assets in Data Integration to integrate data using Redis Connector.

When you use Redis Connector, you can include the following Data Integration assets:

- Mapping
- Mapping task

Important: You can use Redis Connector only in mappings in advanced mode. You can run a task based on a mapping in advanced mode.

For more information about configuring assets and transformations, see *Mappings*, *Transformations*, and *Tasks* in the Data Integration documentation.

Administration of Redis Connector

Before you configure a Redis connection to connect to Redis database, you must complete the following prerequisites:

- Create a user with credentials to access the Redis server.
- Provide read and write access permissions to the user.
- If you want to use TLS for secure communication with the Redis server, enable TLS on the Redis server. Choose the TLS port and configure the TLS server certificates, key, and CA certificates.

CHAPTER 2

Connections for Redis

Use a Redis connection to access the top-level keys in Redis.

Use a Redis connection to read data from or write data to Redis database. You can create a connection on the Connections page or when you create a mapping in advanced mode or a task. Create a connection and associate it with a mapping in advanced mode. Use a Redis connection to import collections, create data objects, preview data, and run mappings in advanced mode.

When you create a Redis connection, you define the connection attributes that Data Integration uses to connect to the Redis database.

You can also parameterize a Redis connection with values specified in a parameter file when you configure a mapping in advanced mode.

Redis connection properties

When you create a Redis connection, you must configure the connection properties.

The following table describes the Redis connection properties:

Property	Description
Connection Name	The name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Redis connection type.
Runtime Environment	Name of the runtime environment where you want to run the tasks. You can specify a Secure Agent or serverless runtime environment.
Host	Host name or IP address of the Redis server.

Property	Description
Port	Redis server port number.
User	Username to access the Redis server.
Password	Password to access the Redis server.
Max Clients Per Worker	The maximum number of Redis client connections used by each worker node.
Flat Hierarchy	Enable this property to perform the following actions based on the data that you read: - Read top-level HASH keys as multiple rows with one row for each key-value pair in the hash. - Read top-level LIST keys as multiple rows with one row for each string value in the list.
Use TLS	Uses TLS to secure the communication with Redis server.
KeyStore File Path	Absolute path of the KeyStore file in the Secure Agent machine that contains private keys and certificates for the Redis server.
KeyStore Passphrase	Passphrase for the KeyStore file.
TrustStore File Path	Absolute path of the TrustStore file that contains certificates for the Redis server.
TrustStore Passphrase	Passphrase for the TrustStore file.

CHAPTER 3

Mappings for Redis

You can configure a mapping in advanced mode to read from and write to Redis.

Use the Data Integration Mapping Designer to define and configure a mapping in advanced mode for Redis sources and targets. In advanced mode, the Mapping Designer updates the mapping canvas to include transformations and functions that enable advanced functionality.

You can configure a Source or Target transformation to represent a Redis object. Add the Source and Target transformations in the mapping canvas and configure the Redis source and target properties. You cannot configure a filter in the Source transformation. However, you can configure a Filter transformation to filter records from the source. Configure the Filter transformation right after the Source transformation to push the filter to the Redis source.

When you run a mappings in advanced mode, Data Integration analyzes the data logic defined in the mapping to automatically assign data logic to run on an advanced cluster.

Redis sources in mappings

To read data from Redis, configure a Redis object as a Source transformation in a mapping in advanced mode.

A Redis source reads all top-level Redis keys that correspond to the Redis object name. Redis object names are converted to Redis binary-string through UTF-8 encoding. Each Redis source in a mapping is configured to read one Redis key type, that is, HASH, LIST, or STRING. Keys of other types are ignored.

You can parameterize the Redis source object with values specified in a parameter file when you configure a mapping in advanced mode.

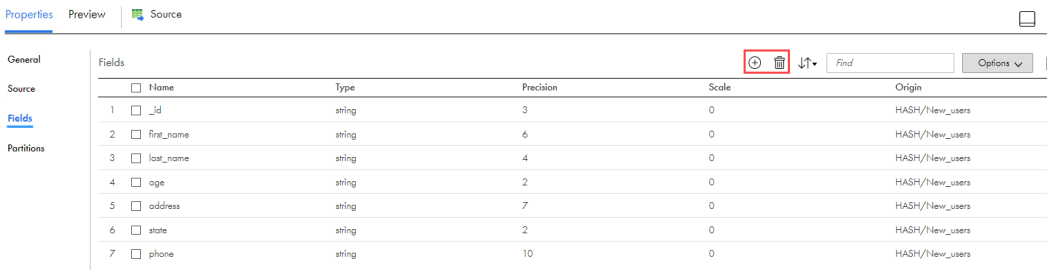
The following table describes the source properties that you can configure for a Redis source:

Property	Description
Connection	Name of the Redis source connection. You can select an existing connection or create a new connection. You can also click **New Parameter** to define a parameter for the source connection. If you want to overwrite the source connection properties at runtime, select the **Allow parameter to be overridden at run time** option. Specify the parameter file directory and name in the advanced session properties.
Source Type	Type of the source object. Select Single Object.
Object	Name of the Redis source object based on the selected source type. Select the source object for a single source.

Adding and removing source fields

You can configure the source fields that you want to use in the data flow.

You can add or remove source fields of primitive data types. You can add fields to a mapping source and retrieve a field from the source object that is not displayed in the list.



For more information on how to add or remove source fields, see the topic "Add and remove fields" in Source transformation in the Data Integration documentation.

Note: You cannot add or remove source fields for Redis STRING data type. Otherwise, the mapping throws a runtime exception.

Redis targets in mappings

To write data to a Redis target, configure a Redis object as a Target transformation in a mapping in advanced mode.

A Redis target writes incoming rows as top-level Redis keys. Each Redis target in a mapping is configured to write all rows as one Redis key type HASH, LIST, or STRING. Redis Connector creates keys with names in ${OBJECT_NAME}:${ID_COLUMN_VALUE} format. Redis object names are converted to Redis binary-string through UTF-8 encoding.

You can parameterize the Redis target object with values specified in a parameter file when you configure a mapping in advanced mode.

The following table describes the target properties that you can configure for a Redis target:

Property	Description
Connection	Name of the Redis connection that is associated with a dataset. You can select an existing connection or create a new connection. You can also click **New Parameter** to define a parameter for the target connection. To overwrite the target connection properties at runtime, select the **Allow parameter to be overridden at run time** option.
Target Type	Type of target object. Select Single Object or Parameter. You can add multiple Redis targets in a mapping in advanced mode.
Object	Name of the Redis target object based on the target type selected. You can select an existing target object or create a new target object at runtime.
Create New at Runtime	Creates a Redis target table at runtime based on the target table you specify. To create a target table at runtime, specify the following details: - Enter the Redis object name to which you want to write in the **Object Name** field. When you write a row to a Redis object, the name of the top-level Redis key is in ${OBJECT_NAME}:${ID_COLUMN_VALUE} format. - Enter the row type in the **Path** field. You can specify row types, HASH, LIST, or STRING based on the row type in the target object. Row types must be in uppercase. The advanced cluster inserts Redis keys into a Redis database that does not already have Redis keys that match the target key pattern. You can edit the metadata of the source fields before creating the target.
Operation	The target operation. You can select the **Insert** operation. Other operations are not applicable.

Redis key types

Redis is a key-value store that contains a list of keys and values. The value of each key is a particular type, including HASH, LIST, and STRING.

Redis Connector supports HASH, LIST, and STRING key types. The supported key types can have the following values:

- HASH value type is a list of binary-string-key and binary-string-value pairs.
- LIST value type is a list of binary-string values.
- STRING value type is a single binary-string value.

The fields in Redis sources match the corresponding schema as shown in the following table. For Redis targets, the mapped incoming fields must exactly match the corresponding schema as shown in the following table:

Key types	If flat hierarchy is not enabled	If flat hierarchy is enabled
HASH	Schema at source is inferred from sampling the data object. Schema at target is inferred from mapped incoming fields. Targets must have an "_id" field.	_id string hkey string value string
LIST	_id string value array of string	_id string value string
STRING	_id string value string	_id string value string

You can use an Expression or Hierarchy Processor transformation to obtain the required schema for the target.

Edit metadata in a Redis source and target

You can edit data types to customize the precision, scale, and data types of fields. You can edit the field metadata in a Source or Target transformation to change data types, precision, and scale. To edit metadata, use the **Edit Metadata** option for the source and target fields on the **Fields** tab of the Properties panel.

Reading and writing hierarchical data in mappings

You can use the Hierarchy Processor transformation in a mapping in advanced mode to process data from complex data sources. You can read complex data types with hierarchical input and write as relational output. You can also read complex data types with relational input and write as hierarchical output.

When you configure a mapping in advanced mode to convert relational to hierarchical output, you can join data sources, group by and order by data fields, filter for specific information, and aggregate both the input and output data. When you configure a mapping in advanced mode to convert hierarchical to relational output, you can read from a single hierarchical data source and write to multiple relational data targets.

For more information about configuring a Hierarchy Processor transformation, see the Hierarchy Processor transformation section in *Transformations* in Data Integration help.

Partitioning

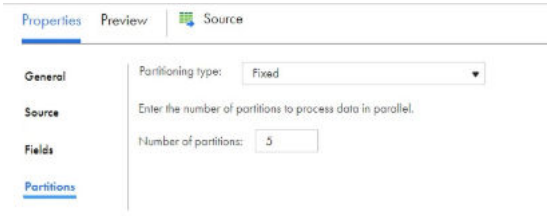
When you read data from Redis, you can configure partitioning to optimize the mapping performance at run time. The partition type controls how the agent distributes data among partitions at partition points.

Fixed partitioning for Redis sources

You can configure fixed partitioning in a mapping that reads data from Redis sources.

You can enable partitioning when you configure the Source transformation in the Mapping Designer.

On the **Partitions** tab for the Source transformation, you select fixed partitioning and enter the number of partitions, as shown in the following image:



The mapping task distributes rows of data based on the number of partitions that you specify.

Rules and guidelines for Redis mappings

Consider the following guidelines before you run a mapping in advanced mode:

- Redis Connector samples rows to populate the object catalog. Use the search field to search for a specific object name if the source object is not displayed in the list.
- The source tables must not be empty. Otherwise, the mapping fails.
- All rows belonging to a HASH-type Redis source must have the same schema. Otherwise, the schema inference might result in incorrect results.
- If the source contains binary fields, the Data Preview section shows jumbled text.

CHAPTER 4

Data type reference

Data Integration uses the following data types in mappings with Redis:

Redis native data types

Binary-string is the only native field data type of Redis. The Redis binary-string data type appears in the **Fields** tab for Source and Target transformations when you choose to edit metadata for the fields.

Transformation data types

Set of data types that appear in the transformations. They are internal data types based on ANSI SQL-92 generic data types, which the Secure Agent uses to move data across platforms. Transformation data types appear in all transformations in a mapping.

When Data Integration reads source data, it converts the native data types to the comparable transformation data types before transforming the data. When Data Integration writes to a target, it converts the transformation data types to the comparable native data types.

Redis and transformation data types

The Redis native type is binary-string for all supported Redis simple values. A Redis binary-string is a variable length array of bytes. You can map Redis binary-string type to the transformation data types, binary and string. Redis binary-string values are converted to and from transformation string data type values through UTF-8 encoding.

The following table describes the Redis data types that Data Integration supports and the corresponding transformation data types:

Redis Data Types	Transformation Data Types	Range and Description
String	Binary	Null values cause row errors in the target.
String	String	Null values cause row errors in the target.

INDEX

C

Cloud Application Integration community
  URL 4
Cloud Developer community
  URL 4

D

Data Integration community
  URL 4

I

Informatica Global Customer Support
  contact information 5
Informatica Intelligent Cloud Services
  web site 4

M

maintenance outages 5
mapping in advanced mode
  Redis sources 9
  Redis targets 10
mapping tasks
  overview 9

P

Partitioning
  Redis sources 13

R

Redis connection
  rules and guidelines 13
Redis connections
  overview 7
  properties 7
Redis connector
  administration 6
Redis Connector
  overview 6
  supported task and object types 6
Redis data types
  mapping to transformation data types 14
  overview 14

S

status
  Informatica Intelligent Cloud Services 5
system status 5

T

trust site
  description 5

U

upgrade notifications 5

W

web site 4