Informatica®

Informatica® Cloud Data Integration

Couchbase Connector

Informatica Cloud Data Integration Couchbase Connector
October 2022

© Copyright Informatica LLC 2018, 2022

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: 2022-11-08

Table of Contents

Preface 5

Informatica Resources 5

- Informatica Documentation 5
- Informatica Intelligent Cloud Services web site 5
- Informatica Intelligent Cloud Services Communities 5
- Informatica Intelligent Cloud Services Marketplace 5
- Data Integration connector documentation 6
- Informatica Knowledge Base 6
- Informatica Intelligent Cloud Services Trust Center 6
- Informatica Global Customer Support 6

Chapter 1: Introduction to Couchbase Connector 7

Couchbase Connector overview 7

Couchbase supported task and object types 8

Introduction to Couchbase 8

Chapter 2: Couchbase Connections 9

Couchbase connections overview 9

Couchbase connection properties 9

Rules and guidelines for Couchbase connections 10

Chapter 3: Mappings and Mapping Tasks with Couchbase 11

Mappings and mapping tasks with Couchbase overview 11

Configure preSQL and postSQL commands for Couchbase sources and targets 11

Couchbase sources in mappings 12

Couchbase targets in mappings 13

Partitioning 14

- Key Range Partitioning 14
- Passthrough Partitioning 15

Chapter 4: Virtual Table Operations 16

Virtual table operations 16

Virtual table operations overview 16

Example for base table and virtual table 17

Appendix A: Data Type Reference 19

Data type reference overview 19

Couchbase and transformation data types 20

Couchbase complex data types 20

Index **21**

Preface

Use *Couchbase Connector* to learn how to read from or write to Couchbase by using Cloud Data Integration. Learn to create a connection, develop mappings, and run mapping tasks in Cloud 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 Couchbase Connector

This chapter includes the following topics:

- Couchbase Connector overview, 7
- Couchbase supported task and object types, 8
- Introduction to Couchbase , 8

Couchbase Connector overview

You can use Couchbase Connector to connect to a Couchbase database from Data Integration. Use Couchbase Connector to read data from or write data to Couchbase.

You can use Couchbase collections as sources and targets in mappings, and then run the associated mapping tasks to read from or write to Couchbase.

You can switch mappings to advanced mode to include transformations and functions that enable advanced functionality.

The following examples are data integration scenarios in which you can use Couchbase Connector:

- Your organization needs to extract data from Couchbase and aggregate data into a data warehouse or data mart for interactive data exploration and analysis.
- Your organization needs to load data from flat files, relational database, or message queues to Couchbase .
- Your organization needs to ensure consistency between your SQL data store and Couchbase data store.

Couchbase supported task and object types

The following table lists the Couchbase Connector object types that you can include in Data Integration tasks:

Task Type	Source	Target	Lookup
Synchronization	No	No	No
Mapping	Yes	Yes	No

Introduction to Couchbase

Couchbase is an open source, distributed NoSQL document-oriented database engine that is highly scalable, provides high availability and requires a simple administration. You can use Couchbase to store large amounts of data spread across data centers or when your applications require high write access speed.

Couchbase uses JSON documents, which has a self-explanatory format and proficient in representing various structures and relationships.

In Couchbase, documents are similar to rows in a relational table. Each row is uniquely identified by a row key. The column name uniquely identifies each column in the column family. The number of columns in each row can vary, and client applications can determine the number of columns in each row.

For example, in a relational data mode for an online flight booking application, you can use three separate tables for flights, schedules and airlines. This allows you to search for flights by dates. However, in Couchbase the document model is a single route object that embeds an array of schedules for all the flights in between the two airports. So, each route document is limited to itself. One request to a single Couchbase node give you access to all the information required to complete the booking.

CHAPTER 2

Couchbase Connections

This chapter includes the following topics:

- Couchbase connections overview, 9
- Couchbase connection properties, 9
- Rules and guidelines for Couchbase connections, 10

Couchbase connections overview

Use a Couchbase connection to access a Couchbase collection, import Couchbase collections, create data objects, preview data, and run mappings.

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

Couchbase connection properties

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

The following table describes the Couchbase connection properties:

Property	Description
Connection Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! $ % ^ & * ( ) - + = { [}] | \ : ; " ' < , > . ? /
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Type	The connection type. Select **Couchbase**.
Runtime Environment	Name of the runtime environment where you want to run the tasks.
Host Name	Host name or IP address of the Couchbase server.

Property	Description
Port	Couchbase server port number. Default is 9042.
Username	User name to access the Couchbase server.
Password	Password corresponding to the user name to access the Couchbase server.
SSL Mode	Not applicable for Couchbase Connector. Select **disabled**.
SSL Certificate Path	Not applicable for Couchbase Connector.
Additional Connection Properties	Enter one or more JDBC connection parameters in the following format: `<param1>=<value>;<param2>=<value>;<param3>=<value>` Couchbase Connector supports the following connection parameters: **QueryMode** It is used to send queries to Couchbase Server. **LogLevel** Species whether the Secure Agent logs error messages in the session log. **LogPath** The complete path to the folder where the driver saves log files when logging is enabled. **AuthMech** The authentication mechanism that the driver uses to connect to the Couchbase server.

Rules and guidelines for Couchbase connections

Consider the following rules and guidelines when you create mapping and mapping tasks:

- If the input value to the Couchbase string contains more than 255 characters, the task fails.
- If you specify the host name, port number, user name, and password of the Couchbase server in the **Additional Connection Properties**, the values specified in the **Additional Connection Properties** takes precedence.
- When you read data from or write data to a Couchbase virtual table, you must not specify **1** as the value of the **QueryMode** parameter in the additional connection properties.

CHAPTER 3

Mappings and Mapping Tasks with Couchbase

This chapter includes the following topics:

- Mappings and mapping tasks with Couchbase overview, 11
- Configure preSQL and postSQL commands for Couchbase sources and targets, 11
- Couchbase sources in mappings, 12
- Couchbase targets in mappings, 13
- Partitioning, 14

Mappings and mapping tasks with Couchbase overview

Use the Data Integration Mapping Designer to create a mapping. When you create a mapping, you configure a source or target to represent a Couchbase collection. In advanced mode, the Mapping Designer updates the mapping canvas to include transformations and functions that enable advanced functionality.

Describe the flow of data from source and target along with the required transformations before the agent writes data to the target. When you create a mapping task, select the mapping that you want to use. Use the Mapping Task wizard to create a mapping task. The mapping task processes data based on the data flow logic you define in the mapping.

Configure preSQL and postSQL commands for Couchbase sources and targets

You can specify **preSQL** and **postSQL** advanced properties for Couchbase sources and targets.

You can perform the following operations by using preSQL and postSQL commands:

- INSERT
- UPDATE
- DELETE

- TRUNCATE
- DROP

Couchbase sources in mappings

To read data from Couchbase, configure a Couchbase object as the Source transformation in a mapping.

Specify the name and description of the Couchbase source. Configure the source, query options, and advanced properties for the source object.

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

Property	Description
Connection	Name of the Couchbase source connection.
Source Type	Type of the Couchbase source objects available.
Object	Name of the Couchbase source object based on the source type selected.

The following table describes the query options that you can configure in a Source transformation:

Property	Description
Filter	Click **Configure** to add conditions to filter records and reduce the number of rows that the Secure Agent reads from the source. You can specify the following filter conditions: - **Not parameterized**. Use a basic filter to specify the object, field, operator, and value to select specific records. - **Completely parameterized**. Use a parameter to represent the field mapping. - **Advanced**. Use an advanced filter to define a more complex filter condition on an object. You can configure filter conditions for columns of the following data types: - Date - Double - Float - Integer - String
Sort	Not applicable for Couchbase Connector.

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

Property	Description
preSQL	SQL statement that you want to run before reading data from the source.
postSQL	SQL statement that you want to run after reading data from the source.

You can set the tracing level in the advanced properties session to determine the amount of details that logs contain.

The following table describes the tracing levels that you can configure:

Property	Description
Terse	The Secure Agent logs initialization information, error messages, and notification of rejected data.
Normal	The Secure Agent logs initialization and status information, errors encountered, and skipped rows due to transformation row errors. Summarizes session results, but not at the level of individual rows.
Verbose Initialization	In addition to normal tracing, the Secure Agent logs additional initialization details, names of index and data files used, and detailed transformation statistics.
Verbose Data	In addition to verbose initialization tracing, the Secure Agent logs each row that passes into the mapping. Also notes where the Secure Agent truncates string data to fit the precision of a column and provides detailed transformation statistics. When you configure the tracing level to verbose data, the Secure Agent writes row data for all rows in a block when it processes a transformation.

Couchbase targets in mappings

Configure a Couchbase object as the Target transformation in a mapping to write data to a Couchbase target.

Specify the name and description of Couchbase target. Configure the target and advanced properties for the target object.

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

Property	Description
Connection	Name of the active Couchbase connection that is associated with a dataset.
Target Type	Type of the Couchbase target objects available. You can write data to a single Couchbase target object or parameterize the object. You cannot write data to multiple objects.
Object	Name of the Couchbase target object based on the target type selected.

Property	Description
Operation	You can select one the following operations: - Insert - Update - Delete
Update columns	Specifies the default primary key column and temporary key columns to update, upsert, or delete target data. If the Couchbase target does not include a primary key column, and the mapping performs an update, upsert, or delete task operation, click **Add** to add the primary key column and temporary key columns. You must select the default primary key column and the temporary key columns. By default, no columns are specified. Select the key columns to upsert or update data to or delete data from the Couchbase target table.

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

Property	Description
UpdateMode	Specifies the mode to write data to Couchbase target. You can specify the following modes: **Update As Update** Updates all rows flagged for update if the entries exist. **Update Else Insert** Updates all rows flagged for update if the entries exist in the target. If the entries do not exist, the Secure Agent inserts the entries. **Note:** Use **Update Else Insert** option to perform upsert operation on a Couchbase target table.
preSQL	SQL statement that you want to run before writing data to the target.
postSQL	SQL statement that you want to run after writing the data into the target.
Forward Rejected Rows	Not applicable for Couchbase Connector.

Partitioning

When you read data from a Couchbase source, you can configure key range partitioning and passthrough partitioning to optimize the mapping performance at run time.

Key Range Partitioning

You can configure key range partitioning when you use a mapping task to read data from Couchbase sources and use simple or hybrid connection mode. With key range partitioning, the Secure Agent distributes rows of

source data based on the fields that you define as partition keys. The Secure Agent compares the field value to the range values for each partition and sends rows to the appropriate partitions.

Use key range partitioning for columns that have an even distribution of data values. Otherwise, the partitions might have unequal size. For example, a column might have 10 rows between key values 1 and 1000 and the column might have 999 rows between key values 1001 and 2000. If the mapping includes multiple sources, use the same number of key ranges for each source.

When you define key range partitioning for a column, the Secure Agent reads the rows that are within the specified partition range. For example, if you configure two partitions for a column with the ranges as 10 through 20 and 30 through 40, the Secure Agent does not read the rows 20 through 30 because these rows are not within the specified partition range.

You can configure a partition key for fields of the following data types:

- String
- Integer

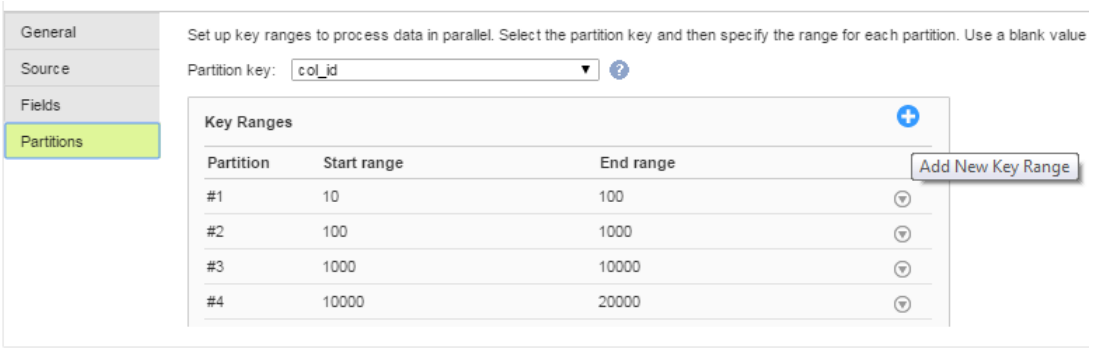
Configuring key range partitioning

Perform the following steps to configure key range partitioning for Couchbase sources:

1. In the Source Properties, click the **Partitions** tab.
2. Select the required partition key from the list.
3. Click **Add New Key Range** to define the number of partitions and the key ranges based on which the Secure Agent must partition data.

   Use a blank value for the start range to indicate the minimum value. Use a blank value for the end range to indicate the maximum value.

   The following image displays the **Partitions** tab:



Passthrough Partitioning

You can configure passthrough partitioning when you use a mapping task to write data to Couchbase targets. With passthrough partitioning, the Secure Agent distributes rows of target data based on the number of threads that you define as partition. The Secure Agent compares the field value to the range values for each partition and sends rows to the appropriate partitions.

When you define passthrough partitioning for a column, the Secure Agent writes the rows that are within the specified partition range. For example, if you configure two partitions for a column with the ranges as 10 through 20 and 30 through 40, the Secure Agent does not write the rows 20 through 30 because these rows are not within the specified partition range.

CHAPTER 4

Virtual Table Operations

This chapter includes the following topics:

- Virtual table operations, 16
- Virtual table operations overview, 16
- Example for base table and virtual table, 17

Virtual table operations

You can use virtual table operations to access, write, or delete rows from a virtual table.

You can perform the following operations in a virtual table:

Read

Use the read operation to retrieve rows from a virtual table.

Insert

Use the insert operation to add one or more rows to a virtual table. Each row that you insert must correspond to a primary key in the source table. If you insert data into a row key that contains data, Couchbase overwrites the original data in the row with the new data.

Update

Use the update operation to update rows in a virtual table.

Note: If the collection type is list, sort the corresponding source or target virtual table data in descending order before you pass data to the Update Strategy transformation.

Virtual table operations overview

Couchbase Connector creates a separate virtual table for each array type collection in the column family. When you use a virtual table in a mapping, you can perform read, insert, upsert, and update operations.

Example for base table and virtual table

The Informatica Couchbase JDBC driver creates virtual tables if the column family contains collections such as set, list, or map.

Base Tables

Base tables represent data from collections of Couchbase documents. Documents appear as rows, and all attributes that are not arrays appear as columns. In each base table, the driver creates a primary key column named PK that identifies which Couchbase document each row comes from.

In the JDBC layer, the name of the base table is the document type that it represents. In the Couchbase layer, the name of the base table is the bucket that the data comes from.

Virtual Tables

Virtual tables provide support for arrays. Each virtual table contains the data from one array, and each row in the table represents an element from the array. If an element contains an array, then the driver creates additional virtual tables as needed to expand the nested data.

In each virtual table, the driver creates a primary key column named PK that identifies the Couchbase document the array comes from and references the PK column from the related base table. The driver also creates an index column (with the suffix _IDX in its name) to indicate the position of the element within the array.

In the JDBC layer, the name of the virtual table is formed using the document type that contains the array, an underscore character (_), and the name of the array. In the Couchbase layer, the name of the virtual table is formed using the name of the bucket that the data comes from, a period (.), and then the name of the array followed a set of closed square brackets ([]) for each hierarchy level in which the array is nested.

Example

The following example shows the base tables and virtual tables that the driver would generate if it connected to a Couchbase database named ExampleDatabase, which contains two documents named Customer_123221 and Order_221354.

The Customer_123221 document is of type Customer and contains the following attributes. The SavedAddresses attribute is an array.

```
{
"Type": "Customer",
"Name": "John Doe",
"SavedAddresses": ["123 Main St.", "456 1st Ave"]
}
```

The Order_221354 document is of type Order and contains the following attributes. The CreditCard attribute is an object, and the Items attribute is an array of objects.

```
{
"Type": "Order",
"CustomerID":"Customer_123221",
"CreditCard":
{
"Type":"Visa",
"CardNumber":"4111 1111 1111 1111",
"Expiry":"12/12",
"CVN":"123"
},
"Items":
[
{"ItemID":89123, "Quantity":1},
{"ItemID":92312, "Quantity":5}
]
}
```

The following table shows that when the driver connects to ExampleDatabase and generates the schema, the driver creates a collection for each document type. The driver exposes these collections as two base tables, which are shown below.

PK	Name
"Customer_123221"	John Doe

PK	CustomerID	CreditCard_Type	CreditCard_Number	CreditCard_Expiry	CreditCard_CVN
"Order_221354"	"Customer_123221"	"Visa"	"4111 1111 1111 1111"	"12/12"	"123"

The SavedAddresses array from the Customer_123221 document and the Items array from the Order_221354 document do not appear in these base tables. Instead, the driver generates a virtual table for each array. The following tables show the virtual tables that represent data from the SavedAddresses and Items arrays.

PK	SavedAddresses_IDX	SavedAddresses
"Customer_123221"	0	"123 Main St."
"Customer_123221"	1	"456 1st Ave"

PK	Items_IDX	ItemID	Quantity
"Order_221354"	0	89123	1
"Order_221354"	1	92312	5

You can select, insert, and update data in the base tables and virtual tables as if they were standard relational tables, and the driver will handle the storage details within Couchbase.

For example, to append an item to the Items array, where the ItemID is 78123 and the Quantity is 6, execute the following statement:

```
INSERT INTO "Order_Items" ("PK", "ItemID", "Quantity") VALUES ('Order_221354', '78123',
'6')
```

APPENDIX A

Data Type Reference

This appendix includes the following topics:

- Data type reference overview, 19
- Couchbase and transformation data types, 20
- Couchbase complex data types, 20

Data type reference overview

Data Integration uses the following data types in mappings, synchronization tasks, and mapping tasks with Couchbase:

Couchbase native data types

Couchbase data types appear 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.

Couchbase and transformation data types

The following table describes the data types that Data Integration supports for Couchbase sources and targets:

Couchbase Data Type	Transformation Data Type	Range and Description for the Transformation Data Type
BOOLEAN	String	1 to 104,857,600 characters. Precision 255.
INT	Integer	-2,147,483,648 to 2,147,483,647 Precision of 10.
FLOAT	Decimal	For transformations that support precision up to 38 digits, the precision is 1 to 38 digits, and the scale is 0 to 38. For transformations that support precision up to 28 digits, the precision is 1 to 28 digits, and the scale is 0 to 28. If you specify the precision greater than the maximum number of digits, the Secure Agent decimal values to double in high precision mode.
DOUBLE	Double	Precision 15.
STRING	String	1 to 104,857,600 characters. Precision 255.
NULL	String	1 to 104,857,600 characters. Precision 255.

Couchbase complex data types

Couchbase connector supports the following complex data types:

Couchbase Complex Data Type	Description for the Transformation Data Type
ARRAYS	Arrays are an ordered list of zero or more values. The values can be of any type. You can enclose Arrays in square brackets ([ ]). Use Commas to separate each value. For example: ["one", "two", "three"] and [1, 2, 3], and ["one", 2, "three"].
OBJECTS	Objects consist of name-value pairs. The name must be a string, and the value can be any supported JSON data type. You can enclose Objects in curly braces ({ }). Use Commas to separate each pair. Use the colon (:) character to separate the name from its value within each pair. For example: {"age": 17}, {"fname": "Jane", "lname": "Doe"}

INDEX

C

Cloud Application Integration community
  URL 5
Cloud Developer community
  URL 5
configuring key range
  partitioning 15
Couchbase
  introduction 8
Couchbase Complex data types
  complex data types 20
Couchbase connections
  overview 9
  properties 9
  rules and guidelines for Couchbase connections 10
Couchbase connector
  supported task and object types 8
Couchbase Connector
  overview 7
Couchbase data types
  mapping to transformation data types 20
  overview 19

D

Data Integration community
  URL 5

I

Informatica Global Customer Support
  contact information 6
Informatica Intelligent Cloud Services
  web site 5

M

maintenance outages 6
mapping
  Couchbase sources 12
  Couchbase targets 13
mapping tasks
  overview 11

P

Partitioning
  key range 15
  passthrough 15
post SQL commands
  entering 11
pre SQL commands
  entering 11

S

status
  Informatica Intelligent Cloud Services 6
system status 6

T

trust site
  description 6

U

upgrade notifications 6

V

virtual table
  collection 17
  list 17
  map 17
  operations 16
  set 17

W

web site 5