Informatica

Informatica® Cloud Data Integration

Elasticsearch Connector

Informatica Cloud Data Integration Elasticsearch Connector
April 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 Elasticsearch Connector 6

Elasticsearch Connector assets 6

Administration of Elasticsearch Connector 6

Chapter 2: Connections for Elasticsearch 7

Elasticsearch connection properties 7

Chapter 3: Mappings for Elasticsearch 9

Elasticsearch sources in mappings 9

- Adding and removing source fields 10
- Joining multiple objects 10

Elasticsearch targets in mappings 10

Edit metadata in an Elasticsearch source and target 11

Reading and writing hierarchical data in mappings 11

Rules and guidelines for Elasticsearch mappings 12

Chapter 4: Data type reference 13

Elasticsearch and transformation data types 13

Index 15

Preface

Use Elasticsearch Connector to learn how to read from or write to Elasticsearch by using Data Integration. Learn to create an Elasticsearch 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 Elasticsearch Connector

You can use Elasticsearch Connector to read data from or write data to collections in an Elasticsearch database.

You can use an Elasticsearch 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.

You can use Elasticsearch Connector to integrate and migrate data from diverse data sources to Elasticsearch. You can also use Elasticsearch Connector to read and write complex data types.

Elasticsearch Connector assets

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

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

- Mapping
- Mapping task

Important: You can use Elasticsearch 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 Elasticsearch Connector

Before you configure an Elasticsearch connection to connect to an Elasticsearch database, configure a user with permissions to read from and write to indices in the Elasticsearch Cloud.

CHAPTER 2

Connections for Elasticsearch

Use an Elasticsearch connection to access an Elasticsearch table.

Use an Elasticsearch connection to read data from or write data to an Elasticsearch database. When you create an Elasticsearch connection, you define the connection attributes that Data Integration uses to connect to the Elasticsearch database.

Elasticsearch connection properties

When you create an Elasticsearch connection, you must configure the connection properties.

The following table describes the Elasticsearch 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 Elasticsearch 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 Elasticsearch server.
Port	Elasticsearch server port number. Default is 9243.
Authentication	Authentication method to access the Elasticsearch resources. Basic authentication uses user name and password credentials to connect to the Elasticsearch server.

Property	Description
User Name	User name to access the Elasticsearch server.
Password	Password corresponding to the user name to access the Elasticsearch server.

CHAPTER 3

Mappings for Elasticsearch

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

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

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

You can configure a Source or Target transformation to represent an Elasticsearch object. Add the Source and Target transformations in the mapping canvas and configure the Elasticsearch 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. To push the filter to the Elasticsearch source, configure the Filter transformation right after the Source transformation.

Elasticsearch sources in mappings

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

Specify the name and description of the Elasticsearch source. When you configure a Source transformation, you can add or remove source fields of primitive data types.

The following table describes the source properties that you can configure for an Elasticsearch source:

Property	Description
Connection	Name of the Elasticsearch 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.
Source Type	Type of the source object. Select Single Object, Multiple Objects, or Parameter. Related objects appear based on simple primary key-foreign key relationship. You can use the Advanced Relationship option to define the relationship for objects that you want to join.
Object	Name of the Elasticsearch source object based on the source type selected. Select the source object for a single source. When you select the multiple source option, you can add source objects and configure a relationship between them.

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.

Joining multiple objects

When you create a Source transformation, you can select multiple Elasticsearch objects as the source type and then configure an advanced relationship to combine the tables.

Perform the following steps to join multiple objects in a Source transformation:

1. In the Source transformation properties, select the **Source Type** as **Multiple Objects**.
2. From the **Actions** menu, select **Add Source Object**, and then select the source object that you want to add from the displayed list.
3. From the **Related Objects Actions** menu, select the Advanced Relationship option to define the relationship between the tables.

   Note: You can join only two source tables. You must define only the join condition, not the complete join query.

Elasticsearch targets in mappings

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

Specify the name and description of an Elasticsearch target.

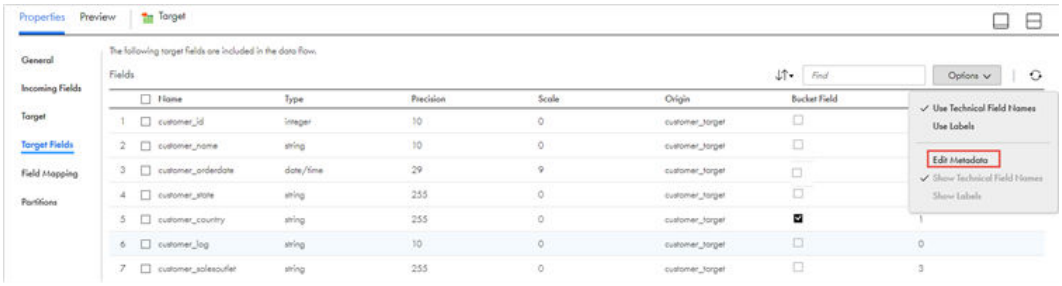
The following table describes the target properties that you can configure for an Elasticsearch target:

Property	Description
Connection	Name of the Elasticsearch 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.
Target Type	Type of target object. Select Single Object or Parameter. You can add multiple Elasticsearch targets in a mappings in advanced mode.
Object	Name of the Elasticsearch target object based on the target type selected. You can select an existing target object or create a new target object at runtime. When you select an existing target object, the schema is inferred from the incoming fields to the target transformation.

Property	Description
Create New at Runtime	Creates an Elasticsearch target table at runtime based on the target table you specify. To create a target table at runtime, enter the name of the target table that you want to create in the **Object Name** field. The agent creates the target table based on the object name and the path you specify. 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.

Edit metadata in an Elasticsearch source and target

You can edit field 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 simple 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 such as structs and arrays. 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 the Data Integration help.

Rules and guidelines for Elasticsearch mappings

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

- The source tables must not be empty. Otherwise, the mapping fails.
- If the source contains binary data, the Data preview section shows jumbled text.
- When the mapping in advanced mode runs, the Secure Agent does not preserve the same order of columns in the Elasticsearch target as in the source.
- You cannot fetch the details of the rejected rows when you write data to an Elasticsearch target.
- You must maintain a uniform schema across fields and sub fields in the Elasticsearch collections.

CHAPTER 4

Data type reference

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

Elasticsearch native data types

Elasticsearch 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.

Elasticsearch and transformation data types

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

Elasticsearch Data Types	Transformation Data Types	Range and Description
Array	Array	-
Binary	Varbinary	Base64 encoded string
Boolean	Integer	true and false
Byte	Integer	A signed 8-bit integer. Minimum value: -128 Maximum value: 127
Date	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D.(precision to second).
Double	Double	Double-precision 64-bit IEEE 754 floating point number, restricted to finite values.
Float	Double	Single-precision 32-bit IEEE 754 floating point number, restricted to finite values.

Elasticsearch Data Types	Transformation Data Types	Range and Description
Integer	Integer	A signed 32-bit integer. Minimum value: -231 Maximum value: 231-1
Long	BigInt	A signed 64-bit integer. Minimum value: -263 Maximum value: 263-1
Null	String	-
Object	Struct	-
Short	Integer	A signed 16-bit integer. Minimum value: -32,768 Maximum value: 32,767
String	String	-

INDEX

C

Cloud Application Integration community
  URL 4
Cloud Developer community
  URL 4

D

Data Integration community
  URL 4

E

Elasticsearch connection
  rules and guidelines 12
Elasticsearch connections
  overview 7
  properties 7
Elasticsearch connector
  administration 6
Elasticsearch Connector
  assets 6
  overview 6
Elasticsearch data types
  mapping to transformation data types 13
  overview 13

I

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

M

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

S

status
  Informatica Intelligent Cloud Services 5
system status 5

T

trust site
  description 5

U

upgrade notifications 5

W

web site 4