Informatica®

Informatica® Cloud Data Integration

# Microsoft Azure Cosmos DB SQL API Connector

Informatica Cloud Data Integration Microsoft Azure Cosmos DB SQL API Connector
April 2024

© Copyright Informatica LLC 2018, 2024

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: 2024-04-07

# 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 Microsoft Azure Cosmos DB SQL API Connector** . . . . 6

Microsoft Azure Cosmos DB SQL API Connector assets . . . . . 7

**Chapter 2: Connections for Microsoft Azure Cosmos DB SQL API** . . . . . 8

Connect to Microsoft Azure Cosmos DB SQL API . . . . . 8

- Before you begin . . . . . 8
- Connection details . . . . . 8

**Chapter 3: Mappings for Microsoft Azure Cosmos DB SQL API** . . . . . 10

Microsoft Azure Cosmos DB SQL API source properties . . . . . 10

Microsoft Azure Cosmos DB SQL API target properties . . . . . 12

- Specify a target . . . . . 14

Parameterization . . . . . 15

Edit metadata . . . . . 16

Rules and guidelines for mappings and mapping tasks . . . . . 16

Mappings in advanced mode . . . . . 17

**Chapter 4: Data type reference** . . . . . 18

Microsoft Azure Cosmos DB SQL API and transformation data types . . . . . 18

Data type parsing for Microsoft Azure Cosmos DB SQL API . . . . . 19

**Index** . . . . . 20

# Preface

Use *Microsoft Azure Cosmos DB SQL API Connector* to learn how to read from or write to Microsoft Azure Cosmos DB SQL API. Learn to create a connection, develop and run mappings, mapping tasks, and data transfer 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, on the Informatica Intelligent Cloud Services Status page, click **SUBSCRIBE TO UPDATES**. You can 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 Global Support Center through the Informatica Network or by telephone.

To find online support resources on the Informatica Network, click **Contact Support** in the Informatica Intelligent Cloud Services Help menu to go to the **Cloud Support** page. The **Cloud Support** page includes system status information and community discussions. Log in to Informatica Network and click **Need Help** to find additional resources and to contact Informatica Global Customer Support through email.

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 Microsoft Azure Cosmos DB SQL API Connector

You can use Microsoft Azure Cosmos DB SQL API Connector to securely read data from or write data to Microsoft Azure Cosmos DB SQL API .

Use Microsoft Azure Cosmos DB SQL API Connector to read JSON documents from and write JSON documents to a container in the Cosmos DB database. You can use Microsoft Azure Cosmos DB SQL API Connector to process large volumes of data.

You can use Microsoft Azure Cosmos DB SQL API Connector for the following data integration scenarios:

- Create a Cosmos DB data warehouse. You can aggregate data from Cosmos DB and other source systems, transform the data, and write the data to Cosmos DB.
- Migrate data from a relational database or other data sources to Cosmos DB. For example, you want to migrate data from a relational database to Cosmos DB. You can write data from multiple relational database tables with different schemas to the same Cosmos DB container. A Cosmos DB container contains the data in a Cosmos DB database.
- Move data between operational data stores to synchronize data. For example, an online marketplace uses a relational database as the operational data store. You want to use Cosmos DB instead of the relational database. However, you want to maintain the relational database along with Cosmos DB for a period of time. You can use Cosmos DB SQL API to synchronize data between the relational data store and the Cosmos DB data store.
- Migrate data from Cosmos DB to a data warehouse for reporting. For example, your organization uses a business intelligence tool that does not support Cosmos DB. You must migrate the data from Cosmos DB to a data warehouse so that the business intelligence tool can use the data to generate reports.

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

# Microsoft Azure Cosmos DB SQL API Connector assets

Create assets in Data Integration to integrate data using Microsoft Azure Cosmos DB SQL API Connector.

When you use Microsoft Azure Cosmos DB SQL API Connector, you can include the following Data Integration assets:

- Data transfer task
- Mapping
- Mapping task

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

CHAPTER 2

# Connections for Microsoft Azure Cosmos DB SQL API

Create a Microsoft Azure Cosmos DB SQL API connection to securely read data from or write data to Microsoft Azure Cosmos DB SQL API.

You can use a Microsoft Azure Cosmos DB SQL API connection to specify sources and targets in mappings and mapping tasks.

## Connect to Microsoft Azure Cosmos DB SQL API

Let's configure the Microsoft Azure Cosmos DB SQL API connection properties to connect to Microsoft Azure Cosmos DB SQL API.

### Before you begin

Before you configure the connection properties, you'll need to get the Cosmos DB URI, database name, and key values from your Microsoft Azure Cosmos DB SQL API account. You can find the key details on the **Keys** tab of your Microsoft Azure Cosmos DB SQL API settings.

### Connection details

The following table describes the basic connection properties:

| Property | Description |
|---|---|
| Connection Name | Name of the connection.<br>Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -,<br>Maximum length is 255 characters. |
| Description | Description of the connection. Maximum length is 4000 characters. |
| Type | Microsoft Azure Cosmos DB SQL API |

| Property | Description |
|---|---|
| Use Secret Vault | Stores sensitive credentials for this connection in the secrets manager that is configured for your organization.<br>This property appears only if secrets manager is set up for your organization.<br>When you enable the secret vault in the connection, you can select which credentials that the Secure Agent retrieves from the secrets manager. If you don't enable this option, the credentials are stored in the repository or on a local Secure Agent, depending on how your organization is configured.<br>For information about how to configure and use a secrets manager, see "Secrets manager configuration" in the Administrator help. |
| Runtime Environment | The name of the runtime environment where you want to run tasks.<br>Select a Secure Agent, Hosted Agent, or serverless runtime environment.<br>Do not use a Hosted Agent if you use the connection in mappings in advanced mode. |
| Cosmos DB URI | The URI to access the Microsoft Azure Cosmos DB account. |
| Key | The primary or secondary key that grants you complete administrative access to the resources in the Microsoft Azure Cosmos DB account. |
| Database | Name of the database that contains the container from which you want to read or write JSON documents. |

CHAPTER 3

# Mappings for Microsoft Azure Cosmos DB SQL API

When you configure a mapping, you describe the flow of data from the source to the target.

A mapping defines reusable data flow logic that you can use in mapping tasks.

When you create a mapping, you define the Source transformation and Target transformation to represent a Microsoft Azure Cosmos DB SQL API object. Use the Mapping Designer in Data Integration to add the Source or Target transformations in the mapping canvas and configure the Microsoft Azure Cosmos DB SQL API source and target properties.

In advanced mode, the Mapping Designer updates the mapping canvas to include transformations and functions that enable advanced functionality.

You can use Monitor to monitor the jobs.

## Microsoft Azure Cosmos DB SQL API source properties

To read data from Microsoft Azure Cosmos DB SQL API, configure a Microsoft Azure Cosmos DB SQL API object as the Source transformation in a mapping.

The following table describes the Microsoft Azure Cosmos DB SQL API source properties that you can configure in a Source transformation:

| Property | Description |
|---|---|
| Connection | Name of the source connection.<br>You can select an existing connection, create a new connection, or define parameter values for the source connection property.<br>If you want to overwrite the source connection properties at runtime, select the **Allow parameter to be overridden at run time** option.<br>Specify the parameter file directory and name in the advanced session properties. |
| Source Type | Type of source object. Select **Single Object** or **Parameter**. |
| Object | Name of the source object. |

| Property | Description |
|---|---|
| Parameter | The parameter that holds the value of the object that you want to define at run-time.<br>Select an existing parameter for the source object, or click **New Parameter** to define a new parameter for the source object.<br>The **Parameter** property appears only if you select parameter as the source type.<br>If you want to overwrite the parameter at runtime, select the **Allow parameter to be overridden at run time** option when you create a parameter. When the task runs, the agent uses the parameters from the file that you specify in the advanced session properties. |
| Format | Mandatory. Microsoft Azure Cosmos DB format options.<br>Select one of the following formats:<br>- None.[1] To read multilevel-hierarchical data in the string format.<br>- Json. To read or write JSON documents that have one level of hierarchy.<br>Default is None. |
| [1]Doesn't apply to mappings in advanced mode. | |

The following table describes the advanced properties that you can configure in a Source transformation:

| Property | Description |
|---|---|
| Throughput[1] | Request units processed per second. Overrides the throughput that you specify for the collection in Microsoft Azure Cosmos DB SQL API at runtime.<br>You cannot override the throughput of the database. Before you override the throughput for the collection, ensure that the throughput is set for the collection in Microsoft Azure Cosmos DB SQL API.<br>Specify the throughput value based on the throughput type:<br>- Autoscale throughput. Specify the value in multiples of 1000. The minimum value is 1000 and maximum value is 1000000.<br>- Manual throughput. Specify the value in multiples of 100. The minimum value for manual throughput is 400 and maximum value is 1000000.<br>Default is -1. |
| Partition Key | Mandatory. Value for the partition field key. You can add more than one value to the field.<br>You can specify the following values:<br>- <All>. Reads data from all partitions.<br>- Field name. Reads data from the field name partition. For example, you can read data from a partition named on the City field, Boston. You can specify comma-separated multiple field names.<br>- <null>. Reads data from the null partition.<br>Default is <All>.<br>When you use the partition key for mappings in advanced mode, ensure that the key is of string data type. |
| Page Size[1] | Number of documents to read per request. Default is 50. |

| Property | Description |
|---|---|
| Filter Query | A case-sensitive filter query with conditional and logical operators to filter the source data.<br>Use the following syntax:<br>`<objectName>.<columnName>="conditionValue"`<br>For example, `Address.City="Boston"`<br>To use the array data type for mappings in advanced mode, use the following syntax:<br>`<objectName>.<arrayName>[rowNumber][colNumber]="conditionValue"`<br>To use the struct data type for mappings in advanced mode, use the following syntax:<br>`<objectName>.<structLevel1>.<structLevel2>.<structLevel<n>>.<variableName>= "conditionValue"` |
| Tracing Level | Not applicable. |
| [1]Doesn't apply to mappings in advanced mode. | |

# Microsoft Azure Cosmos DB SQL API target properties

You can configure a Target transformation to represent a Microsoft Azure Cosmos DB SQL API target.

The following table describes the Microsoft Azure Cosmos DB SQL API target properties that you can configure in a Target transformation:

| Property | Description |
|---|---|
| Connection | Name of the target connection.<br>You can select an existing connection, create a new connection, or define parameter values for the target connection property.<br>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**. |
| Object | The target object for the task. Select the target object.<br>You can write data by selecting an existing target object or creating a new target by using the **Create New at Runtime** option. |

| Property | Description |
|---|---|
| Create New at Runtime | Creates a Microsoft Azure Cosmos DB SQL API target at runtime.<br>To create a target table at runtime, provide the following parameters:<br>- Object Name. Name of the collection to be created in Microsoft Azure Cosmos DB SQL API.<br>Name must not start or end with a space or contain the following characters: \ / # ?<br>- Partition Key. The partition key for the collection.<br>Specify the key in the following format: `/ID`. Ensure that the partition key field is an existing field from the Source transformation or any previous transformation.<br>Partition key must not contain only numbers, spaces, unicode characters, and special characters except underscore.<br>- Collection Throughput. Request units processed per second.<br>Specify the value for Autoscale throughput in multiples of 1000. The minimum value for autoscale throughput is 1000 and maximum value is 1000000.<br>Specify the value for Manual throughput in multiples of 100. The minimum value for manual throughput is 400 and maximum value is 1000000.<br>- Collection Throughput Scale. Type of throughput. Select Autoscale or Manual.<br>- Formatting options. Not applicable.<br>- Handle Special Characters. Not applicable. |
| Parameter | The parameter that holds the value of the object that you want to define at run-time.<br>Select an existing parameter for the target object or click **New Parameter** to define a new parameter for the target object.<br>The **Parameter** property appears only if you select parameter as the target type.<br>If you want to overwrite the target object at runtime, select the **Allow parameter to be overridden at run time** option. When the task runs, the agent uses the parameters from the file that you specify in the advanced session properties. |
| Format | Mandatory. Microsoft Azure Cosmos DB format options.<br>Select one of the following formats:<br>- None.[1] To write multilevel-hierarchical data in the string format.<br>- Json. To read or write JSON documents that have one level of hierarchy.<br>Default is None.<br>When you create a Microsoft Azure Cosmos DB SQL API target at runtime, you cannot use the None option. |
| Operation | Not applicable. You must configure the target operation in the target advanced properties. |

[1]Doesn't apply to mappings in advanced mode.

The following table describes the advanced properties that you can configure in a Target transformation:

| Advanced Property | Description |
|---|---|
| Throughput[1] | Request units processed per second. Overrides the throughput that you specify for the collection in Microsoft Azure Cosmos DB SQL API at runtime.<br>You cannot override the throughput of the database. Before you override the throughput for the collection, ensure that the throughput is set for the collection in Microsoft Azure Cosmos DB SQL API.<br>Specify the throughput value based on the throughput type:<br>- Autoscale throughput. Specify the value in multiples of 1000. The minimum value is 1000 and maximum value is 1000000.<br>- Manual throughput. Specify the value in multiples of 100. The minimum value for manual throughput is 400 and maximum value is 1000000.<br>Default is -1. |
| Automatic Id Generation[1] | Generate ID for the documents written to the target. Specify any of the following values:<br>- Enabled. Cosmos DB generates IDs for the documents.<br>- Disabled. The source object provides IDs for the documents. Data Integration rejects the rows if a Null value is provided to the ID port in the target or if you do not connect a port to the ID port in the target. |
| Treat Source Rows As | Operation to perform.<br>Select one of the following options:<br>- Insert<br>- Update[1]<br>- Upsert<br>- Delete<br>You must connect the ID port and Partition Key port to perform update, upsert, or delete operations. Data Integration displays an exception if you do not connect the ID port or the partition key for upsert and update operations, whereas, the delete operation fails.<br>For upsert and update operations, in addition to the ports you want to update, you must connect all other ports a document contains. The values for the unconnected ports get deleted during upsert and update operations. |
| Execution Priority[1] | Prioritizes bulk write requests in Microsoft Azure Cosmos DB SQL API.<br>Set the priority to low to throttle throughput of low priority requests and allow the execution of only the higher priority requests. Low priority requests are throttled only when there are no other high priority requests.<br>Default is High. |
| Is Bulk Write[1] | Indicates whether to write bulk data. By default, the Secure Agent does not write bulk data. |
| Forward Rejected Rows[1] | Determines whether the transformation passes rejected rows to the next transformation or drops rejected rows.<br>By default, the mapping task forwards rejected rows to the next transformation. |
| [1]Doesn't apply to mappings in advanced mode. | |

## Specify a target

You can use an existing target or create a new target to write data to Microsoft Azure Cosmos DB SQL API . If you choose to create a new target, the agent creates the target when it runs the task.

When you use an existing target, you can infer schema from the selected target object or generate schema from the incoming fields.

Consider the following rules and guidelines for an existing target:

- You can infer schema from the selected target object only when you read primitive data types.
- You can generate schema from the incoming fields when you read primitive or hierarchical data types for mappings in advanced mode.

When you create a new target, specify the object name, partition key, collection throughput, and collection throughput scale.

Consider the following rules and guidelines when you create a new target:

- If you specify the collection throughput, you must also specify the collection throughput scale. Ensure that the value for the collection throughput is a valid value.
- If you want to inherit the value of collection throughput and collection throughput scale from the database, leave the fields empty. Ensure that the collection throughput and collection throughput scale are specified for the database.

# Parameterization

You can parameterize the connection, objects, and the advanced runtime properties in mappings.

Consider the following rules and guidelines when you use parameterization:

**General guidelines**

- When you create a mapping with a parameterized target that you want to create at runtime, set the target field mapping to automatic.
- You cannot parameterize the field mapping.

**Mappings**

- When you use input parameters, specify the parameter name in the following format:
  - Format in a mapping task: $name$
  - Format in a parameter file: $name
- When you use in-out parameters, specify the parameter name in the following format in a mapping task or a parameter file: $$name.

**Mappings in advanced mode**

- When you use input parameters, specify the parameter name in the following format in a mapping task: $name$.
- When you use input parameters in a parameter file to parameterize the connection or object, the parameter value is stored in the following format: $$name.
- You cannot use input parameters in a parameter file to parameterize the advanced properties at runtime.
- You cannot parameterize hierarchical data types.

# Edit metadata

You can edit field data types to customize the precision and data types of fields. You can edit the field metadata in a Source or Target transformation to change simple data types and precision.

To edit metadata, use the **Edit Metadata** option for the source and target fields on the **Fields** tab of the **Properties** panel.

## Rules and guidelines

- When you read and write values of integer, double, bigint, and boolean data types in a mapping, you cannot edit the precision. You can edit the precision only when you create a mapping task.
- When you read and write values of string data type in a mapping and edit the precision, the mapping does not consider the change in the precision. You can edit the precision when you create a mapping task.

# Rules and guidelines for mappings and mapping tasks

Consider the following rules and guidelines for mappings and mapping tasks:

- When two or more mappings run concurrently to update the throughput of the same collection, the mappings fail to update the throughput.
- The target operation that you configure in the target advanced properties takes precedence over the target operation that you configure in the target properties.
- When you run a mapping to read data that contains unicode characters with comma and write to a Microsoft Azure Cosmos DB SQL API target, the data is incorrectly written to the target.
- When you run a mapping to read data that contains unicode characters with comma and write to a Microsoft Azure Cosmos DB SQL API target, the data is incorrectly written to the target.
- When you select a connection and a source object, the first document is read and a schema is created from it. The schema chosen is applied for all the documents and the additional columns in the documents are ignored.
- If the value of the partition key is null in the source data, specify the value of the partition key as null within angular brackets in the advanced source properties.
  For example, `<null>`.
- When you read from or write to Microsoft Azure Cosmos DB SQL API, ensure that the collection contains at least one document to infer the schema.
- When you read large amount of data in Microsoft Azure Cosmos DB SQL API and use the None format option, to improve performance, specify the precision of the Document field according to the size of the data to be processed using the Edit Metadata option.
- When you update or delete data in a partitioned collection in a Microsoft Azure Cosmos DB SQL API target and use the bulk API, ensure that the source data has more than one row.

# Mappings in advanced mode

You can create a mapping in advanced mode to read data from or write data to Microsoft Azure Cosmos DB SQL API.

Consider the following rules and guidelines for mappings in advanced mode:

- You cannot read data in binary format.
- You cannot read data from non-partitioned collections.
- When you read boolean fields in a mapping, Boolean fields in an array are written as null in the target.
- When you read string data, you cannot use in-out parameters in the advanced properties.
- You cannot add multiple pipelines in a mapping.
- The unique key that you specify when you create a container is not considered in a mapping.
- When you read from Microsoft Azure Cosmos DB SQL API and write to Microsoft Azure Data Lake Storage Gen2, a zero-byte file is created in the target with other target files.
- When you read from or write to Microsoft Azure Cosmos DB SQL API, ensure that the collection contains at least one document to infer the schema.
- When you read data from a Microsoft Azure Synapse SQL and write data to Boolean fields in Microsoft Azure Cosmos DB SQL API, ensure that the data type of the source fields is Bit.
- When you read from or write to Microsoft Azure Cosmos DB SQL API, ensure that the ID column in the source is mapped to the ID column in the target.
  If the source does not have an ID column, use the Expression transformation to add an ID column and then run the mapping.
- When you run a task that reads from Microsoft Azure Cosmos DB SQL API and writes to Microsoft Azure Data Lake Storage Gen2 in the verbose mode, the execution log shows the following error even though the task is successful:
  `java.lang.IllegalArgumentException: Failed to load any of the given libraries: [netty_tcnative_linux_x86_64, netty_tcnative_linux_x86_64_fedora, netty_tcnative_x86_64, netty_tcnative]`

  Ignore the message.
- When you read from and write to Microsoft Azure Cosmos DB SQL API, any values of double data type in the source are written in exponential form in the target.
  For example, 1234567 is written as 1.234567E6.
- You cannot partially parameterize the advanced properties. For example, do not specify `Add100.$$filtercolumn=8`.

CHAPTER 4

# Data type reference

Data Integration uses the following data types in Microsoft Azure Cosmos DB SQL API mappings and mapping tasks:

- Microsoft Azure Cosmos DB SQL API native data types appear in the Source transformation and Target transformation when you choose to edit metadata for the fields.
- Transformation data types. Set of data types that appear in the transformations. These are internal data types based on ANSI SQL-92 generic data types, which the Secure Agent uses to move data across platforms. They appear in all transformations in a mapping.

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

## Microsoft Azure Cosmos DB SQL API and transformation data types

The following table compares the JSON data type to the transformation data type:

| JSON Data Type | Transformation Data Type | Range and Description |
|---|---|---|
| array[1] | array | The size of a document must be less that or equal to 4 MB. |
| boolean | integer | The default transformation type for boolean is integer. You can specify string data type with values of True and False. True is equivalent to the integer 1 and False is equivalent to the integer 0. |
| Number (double) | double | -1.79769313486231570E+308 to +1.79769313486231570E+308.<br>Precision 15. |
| Number (float) | double | -1.79769313486231570E+308 to +1.79769313486231570E+308.<br>Precision 15. |

| JSON Data Type | Transformation Data Type | Range and Description |
|---|---|---|
| Number (int) | integer | -2,147,483,648 to 2,147,483,647 Precision 10, scale 0 |
| Number (long) | bigint | -9,223,372,036,854,775,000 to 9,223,372,036,854,775,000.<br>Precision 19, scale 0. |
| string | string | 1 to 104,857,600 characters. |
| struct[1] | struct | The size of a document must be less that or equal to 4 MB. |
| [1]Applies only to mappings in advanced mode. | | |

### Rules and guidelines for data types

You cannot edit the metadata of the fields of hierarchical data type.

# Data type parsing for Microsoft Azure Cosmos DB SQL API

When you read from or write to Microsoft Azure Cosmos DB SQL API, Data Integration parses data based on the data types defined in the schema. If the data values do not match the data types defined in the schema, Data Integration rejects the document.

The following table lists the data types allowed at run time for the data types specified in the schema:

| Data Type in Schema | Allowed Run-time Data Type |
|---|---|
| Integer | Short, Integer |
| BigInt or Long | Short, Integer, Long (maximum precision 19) |
| Float | Short, Integer, Long, Float |
| Double | Short, Integer, Long, Float, Double |
| Decimal (maximum precision 28) | Short, Integer, Long, Float, Double, Long |
| String | String |

# INDEX


## A

Azure Blob Container Name 12

## C

Cloud Application Integration community
  URL 4
Cloud Developer community
  URL 4
connections
  Microsoft Azure Cosmos DB SQL API 8
Cosmos DB
  data types 18

## D

Data Integration community
  URL 4
data types
  Cosmos DB 18
  parsing 19

## E

error file directory 12

## F

field delimiter 12
filter query 10
forward rejected rows 12

## I

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

## M

maintenance outages 5
mappings
  source properties 10
mappings *(continued)*
  target properties 12
Microsoft Azure Cosmos DB SQL API
  connection properties 8
  overview 6
  rules and guidelines for objects 16
Microsoft Azure Cosmos DB SQL API Connector
  assets 7

## N

native data type 18
number of concurrent sessions to Blob Storage 12

## P

page size 10
partition key 10
properties
  in mappings 10, 12

## S

status
  Informatica Intelligent Cloud Services 5
success file directory 12
system status 5

## T

throughput 10
tracing level 10
transformation data type 18
trust site
  description 5

## U

upgrade notifications 5

## W

web site 4