Informatica®

Informatica® Cloud Data Integration

Amazon DynamoDB V2 Connector

Informatica Cloud Data Integration Amazon DynamoDB V2 Connector
April 2023

© Copyright Informatica LLC 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: 2023-10-10

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 Amazon DynamoDB V2 Connector 6

Amazon DynamoDB V2 Connector assets 6

Authentication 7

- IAM authentication 7

Temporary security credentials using AssumeRole 7

- Temporary security credentials policy 8
- External ID 8

Chapter 2: Connections for Amazon DynamoDB V2 10

Amazon DynamoDB V2 connection properties 10

Chapter 3: Mappings for Amazon DynamoDB V2 12

Amazon DynamoDB V2 sources in mappings 12

- Schema inference for Amazon DynamoDB V2 sources 13
- Adding and removing source fields 13

Amazon DynamoDB V2 targets in mappings 14

Amazon DynamoDB V2 lookups in mappings 16

- Amazon DynamoDB V2 lookup properties 16

Reading and writing hierarchical data in mappings 17

Rules and guidelines for Amazon DynamoDB V2 mappings 17

Chapter 4: Data type reference 18

Amazon DynamoDB V2 and transformation data types 18

Index 20

Preface

Use *Amazon DynamoDB V2 Connector* to learn how to read from or write to Amazon DynamoDB tables by using Cloud Data Integration. Learn to create an Amazon DynamoDB V2 connection, develop and run mappings, and 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 Amazon DynamoDB V2 Connector

You can use Amazon DynamoDB V2 Connector to connect to Amazon DynamoDB from Data Integration. You can use Amazon DynamoDB V2 Connector to read data from and write data to Amazon DynamoDB tables using mappings in advanced mode. You can also use Amazon DynamoDB V2 Connector to read and write hierarchical data types.

Amazon DynamoDB is a fully managed NoSQL database service that you can use to create database tables where you can store structured and unstructured data.

When you read data from or write data to Amazon DynamoDB tables using Amazon DynamoDB V2 Connector, you can specify the Secure Agent. Create an Amazon DynamoDB V2 connection to specify the location of Amazon DynamoDB V2 sources and targets you want to include in a task.

You can use an Amazon DynamoDB V2 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.

Amazon DynamoDB V2 Connector assets

Create assets in Data Integration to integrate data using Amazon DynamoDB V2 Connector.

When you use Amazon DynamoDB V2 Connector, you can include the following Data Integration assets:

- Mapping
- Mapping task

Important: You can use Amazon DynamoDB V2 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.

Authentication

Amazon DynamoDB V2 Connector uses the standard Amazon AWS SDK authentication mechanisms.

If you do not provide the access key and the secret key in the connection, Amazon DynamoDB V2 Connector uses AWS credentials provider chain that looks for credentials in the following order:

1. The **AWS_ACCESS_KEY_ID** and **AWS_SECRET_ACCESS_KEY** or **AWS_ACCESS_KEY** and **AWS_SECRET_KEY** environment variables.
2. The **aws.accessKeyId** and **aws.secretKey** java system properties.
3. The credential profiles file at the default location, `~/.aws/credentials`.
4. The instance profile credentials, such as the master instance profile and the worker instance profile in the advanced cluster, delivered through the Amazon EC2 metadata service.

For more information, see the AWS documentation.

IAM authentication

Use IAM authentication for a secure and controlled access to Amazon DynamoDB tables when you run a task.

You can configure IAM authentication when the Secure Agent runs on an Amazon Elastic Compute Cloud (EC2) system.

1. Create a policy with read and write access to Amazon DynamoDB tables.
2. Create the Amazon EC2 role.

   The Amazon EC2 role is used when you create an EC2 system. For more information about creating the Amazon EC2 role, see the AWS documentation.
3. Link the policy with the Amazon EC2 role.
4. Create an EC2 instance and then assign the Amazon EC2 role that you created in step 2 to the EC2 instance.
5. Install the Secure Agent on the EC2 system.

Temporary security credentials using AssumeRole

You can use the temporary security credentials using AssumeRole to access the Amazon DynamoDB resources from the same or different AWS accounts.

Ensure that you have the **sts:AssumeRole** permission and a trust relationship established within the AWS accounts to use the temporary security credentials. The trust relationship is defined in the trust policy of the IAM role when you create the role. The IAM role adds the IAM user as a trusted entity and allows IAM users to use the temporary security credentials and access AWS accounts. For more information about how to establish the trust relationship, see the AWS documentation.

When the trusted IAM user requests for the temporary security credentials, the AWS Security Token Service dynamically generates the temporary security credentials that are valid for a specified period and provides the credentials to the trusted IAM users. The temporary security credentials consist of an access key, secret key, and secret token.

To use the dynamically generated temporary security credentials, provide the value of the **IAM Role ARN** connection property when you create an Amazon DynamoDB V2 connection. The IAM Role ARN uniquely identifies the Amazon DynamoDB resources.

Temporary security credentials policy

To use the temporary security credentials to access Amazon DynamoDB resources, both the IAM user and IAM role require policies.

The following section lists the policies required for the IAM user and IAM role:

IAM user

An IAM user must have the `sts:AssumeRole` policy to use the temporary security credentials in the same or different AWS account.

The following sample policy allows an IAM user to use the temporary security credentials in an AWS account:

```
{
"Version":"2012-10-17", "Statement":{ "Effect":"Allow", "Action":"sts:AssumeRole",
"Resource":"arn:aws:iam::<ACCOUNT-HYPHENS>:role/<ROLE-NAME>" }
}
```

IAM role

An IAM role must have a `sts:AssumeRole` policy and a trust policy attached with the IAM role to allow the IAM user to access Amazon DynamoDB tables using the temporary security credentials. The policy specifies the Amazon DynamoDB tables that the IAM user can access and the actions that the IAM user can perform. The trust policy specifies the IAM user from the AWS account that can access the Amazon DynamoDB tables.

The following policy is a sample trust policy:

```
{
"Version":"2012-10-17", "Statement":[{ "Effect":"Allow", "Principal":
{ "AWS":"arn:aws:iam::AWS-account-ID:root" },
"Action":"sts:AssumeRole" }
 ]
  }
   }
```

Here, in the `Principal` attribute, you can also provide the ARN of the IAM user who can use the dynamically generated temporary security credentials and to restrict further access. For example,

```
"Principal" : { "AWS" : "arn:aws:iam:: AWS-account-ID :user/ user-name " }
```

External ID

You can optionally specify the external ID in the AssumeRole request to the AWS Security Token Service for a more secure access to the Amazon DynamoDB tables. The external ID must be a string.

The following sample shows an external Id condition in the trust policy of the assumed IAM role:

```
"Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::AWS_Account_ID : user/user_name"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:ExternalId": "dummy_external_id"
        }
      }
```

```
    }
]
```

CHAPTER 2

Connections for Amazon DynamoDB V2

You can use an Amazon DynamoDB V2 connection to read data from or write data to Amazon DynamoDB tables. You can use the Amazon DynamoDB V2 connection to specify sources and targets in mappings in advanced mode.

Create a connection and associate it with mappings in advanced mode. Define the source and target properties to read data from or write data to Amazon DynamoDB tables.

You can create an Amazon DynamoDB V2 connection on the **Connections** page. Use the connection in the Mapping Designer when you create mappings in advanced mode. You can specify the Amazon DynamoDB V2 connection as a source or target in a mapping task.

Amazon DynamoDB V2 connection properties

When you set up an Amazon DynamoDB V2 connection, you must configure the connection properties.

The following table describes the Amazon DynamoDB V2 connection properties:

Connection 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	Amazon DynamoDB V2.
Runtime Environment	The name of the runtime environment where you want to run tasks. Specify a Secure Agent or serverless runtime environment.
Access Key	The access key to access Amazon DynamoDB. You can optionally enter the access key when you use assume role for an IAM user.

Connection property	Description
Secret Key	The secret key to access Amazon DynamoDB. This value is associated with the access key and uniquely identifies the account. You can optionally enter the secret key when you use assume role for an IAM user.
Region Name	The AWS region of Amazon DynamoDB that you want to access.
Assume Role	Enables the IAM entity to assume a role.
Assume Role ARN	The ARN of the IAM role assumed by the IAM user to generate the temporary security credentials.
External Id	The external ID to generate the temporary security credentials.
Additional Options	Optional properties in key-value pairs that you can configure when you read data from or write data to Amazon DynamoDB. To specify more than one property, separate the key-value pairs with an ampersand. For example, propertyName1=<value1>&propertyName2=<value2> You can configure the following parameter: **prefixFieldNames=true**. If you configure this parameter in the source connection, it prefixes an underscore character to all the columns when you import the table. If you configure this parameter in the target connection, it removes the first character from all the target columns.

CHAPTER 3

Mappings for Amazon DynamoDB V2

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

Use the Data Integration Mapping Designer to define and configure a mapping in advanced mode for Amazon DynamoDB V2 sources and targets. In advanced mode, the Mapping Designer updates the mapping canvas to include transformations and functions that enable advanced functionality.

In advanced mode, you can configure a Source or Target transformation to represent an Amazon DynamoDB V2 object. Add the Source and Target transformations in the mapping canvas and configure the Amazon DynamoDB V2 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 Amazon DynamoDB source, configure the Filter transformation right after the Source transformation.

When you run mappings in advanced mode, Data Integration analyzes the data logic defined in the mapping to automatically assign data logic to run on an advanced cluster. You can use Monitor to monitor the job.

Amazon DynamoDB V2 sources in mappings

To read data from an Amazon DynamoDB table, configure an Amazon DynamoDB V2 object as the Source transformation in a mapping.

Specify the name and description of the Amazon DynamoDB V2 source. Configure the source properties for the source object. You can also parameterize the Amazon DynamoDB V2 source object with values specified in a parameter file when you configure a mapping in advanced mode.

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 in a Source transformation:

Property	Description
Connection	Name of the source connection. You can select an existing connection or create a new connection. You can also select **New Parameter** to define a parameter for the source connection. If you want to overwrite the source connection properties at runtime, select the **Allow parameter to be overridden at run time** option. Specify the parameter file directory and name in the advanced session properties.
Source Type	Type of the source object. Select **Single Object** or **Parameter**.
Object	Name of the DynamoDB V2 source object based on the source type selected.

The following table describes the Amazon DynamoDB V2 advanced source property that you can configure in a Source transformation:

Property	Description
Refresh duration for Assume Role Session	Generates a new set of credentials for the assume role after the specified duration and refreshes the assume role session after fixed intervals. You can configure **Refresh duration for Assume Role Session** when you configure assume role in the connection properties. The default value is 900 seconds. The maximum value is 12 hours depending upon the maximum session duration that you set for the role.

Schema inference for Amazon DynamoDB V2 sources

The schema of the Amazon DynamoDB source is inferred by sampling rows from the Amazon DynamoDB table.

The data that you read from an Amazon DynamoDB table using a mapping must have a strict schema. Amazon DynamoDB V2 Connector picks the most common data types from the sample for mixed columns, sub-columns, and lists.

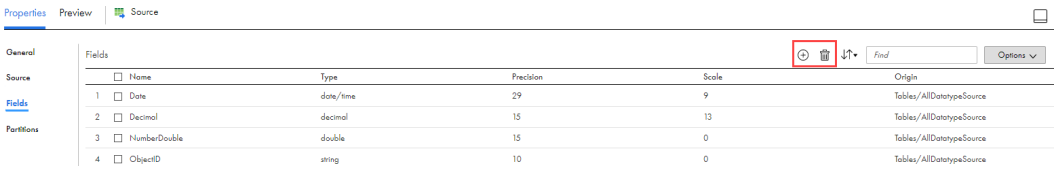
Consider the following rules and guidelines to infer the schema of an Amazon DynamoDB V2 source:

- Amazon DynamoDB V2 Connector chooses the string data type when all the rows in a sample have a null value for a particular column.
- If a row has an extra top-level column or nested column that does not appear in the schema, this extra column is dropped.
- If a row is missing a top-level column or nested column, the source uses a null value for that column.
- If a row has a column with a type that does not match the schema, the source tries to cast the value to the type in the schema. If this cast fails, the source uses a null value for that column.

Adding and removing source fields

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

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.

Amazon DynamoDB V2 targets in mappings

To write data to an Amazon DynamoDB table, configure an Amazon DynamoDB object as the Target transformation in a mapping in advanced mode.

Specify the name and description of the Amazon DynamoDB V2 target. Configure the target properties for the target object. You can also parameterize the Amazon DynamoDB V2 target object with values specified in a parameter file when you configure a mapping in advanced mode.

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

Property	Description
Connection	Name of the target connection. You can select an existing connection or create a new connection. You can also select **New Parameter** to define a parameter for the target connection. To overwrite the target connection properties at runtime, select the **Allow parameter to be overridden at run time** option.
Target Type	Target type. Select **Single Object** or **Parameter**.
Object	Name of the target object. You can select an existing object or create a new target at runtime. When you select an existing target object, the schema is inferred from the incoming fields to the target transformation. When you create a new target at run time and the target already exists, the mapping appends the data to the existing target.
Operation	Target operation. You can only select the insert operation.

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

Property	Description
Object Name	The DynamoDB table name.
Hash Key Name	The name of the hash key of the table.

Property	Description
Hash Key Type	The type of the hash key of the table. Select from the following values: - Binary - Number - String
Range Key Name	Optional. Name of the range key of the table.
Range Key Type	Optional. Type of the range key of the table. Select from the following values: - Binary - Number - String
Billing Mode	The type of billing mode. Select **OnDemand** or **Provisioned**.
Provisioned Read Capacity Units	The number of read capacity units of the table. Required when you select the Provisioned billing mode.
Provisioned Write Capacity Units	The number of write capacity units of the table. Required when you select the Provisioned billing mode.
AWS Tags	The labels to identify DynamoDB target tables. The AWS tags are key-value pairs. You can specify zero or more tags when you create a new target table. Specify the AWS tags in the following JSON format: `{"tag1":"value1", "tag2":"value2"}` Specify a maximum of 255 characters for the AWS tags.
Path	Not applicable.

The following table describes the Amazon DynamoDB V2 advanced target property that you can configure in a Target transformation:

Property	Description
Refresh duration for Assume Role Session	Generates a new set of credentials for the assume role after the specified duration and refreshes the assume role session after fixed intervals. You can configure **Refresh duration for Assume Role Session** when you configure assume role in the connection properties. The default value is 900 seconds. The maximum value is 12 hours depending upon the maximum session duration that you set for the role.

Amazon DynamoDB V2 lookups in mappings

Add a Lookup transformation to retrieve data based on a specified lookup condition. When you add a Lookup transformation to a mapping, you define the lookup connection, lookup objects, and lookup properties related to Amazon DynamoDB.

You can look up Amazon DynamoDB data values based on a condition that you configure. In the Lookup transformation, select the lookup connection and object. Then, define the lookup condition and the outcome for multiple matches.

The mapping queries the lookup source based on the lookup fields and the defined lookup condition. The lookup operation returns the result to the Lookup transformation, which then passes the results downstream.

You can configure connected and unconnected lookups for mappings in advanced mode.

For more information about Lookup transformation, see *Transformations* in the Data Integration documentation.

Amazon DynamoDB V2 lookup properties

The following table describes the Amazon DynamoDB V2 lookup object properties that you can configure in a Lookup transformation:

Property	Description
Connection	Name of the lookup connection. You can select an existing connection, create a new connection, or define parameter values for the lookup connection property. If you want to overwrite the lookup connection properties at run time, select the **Allow parameter to be overridden at run time** option.
Source Type	Type of the source object. Select Single Object or Parameter.
Parameter	A parameter file where you define values that you want to update without having to edit the task. Select an existing parameter for the lookup object or click **New Parameter** to define a new parameter for the lookup object. The **Parameter** property appears only if you select parameter as the lookup type. If you want to overwrite the parameter at run time, select the **Allow parameter to be overridden at run time** option. When the task runs, the Secure Agent uses the parameters from the file that you specify in the advanced session properties.
Lookup Object	Name of the lookup object for the mapping.
Filter	Not applicable.
Sort	Not applicable.

Reading and writing hierarchical data in mappings

You can use the Hierarchy Processor transformation in a mapping to process data from complex data sources such as List, Map, and Set. 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 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 to convert hierarchical to relational output, you can read from a single hierarchical data source and write to multiple relational data targets.

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

Rules and guidelines for Amazon DynamoDB V2 mappings

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

- The source tables must not be empty. Otherwise, the mapping fails.
- All rows in the source table must have the same schema. Otherwise, the schema inference might produce incorrect results.
- At runtime, all the rows in the source table must match the schema that was inferred when the mapping was created. Otherwise, the mapping run might produce incorrect results.
- A table that contains columns with null values might treat the columns as string types when the mapping runs.
- If the source contains binary data, the Data Preview section shows jumbled text.
- When you enable data preview for mappings that use an existing Amazon DynamoDB target, the data preview fails.
- You cannot use complex data type columns as lookup conditions.

CHAPTER 4

Data type reference

Data Integration uses the following data types in mappings with Amazon DynamoDB V2:

Amazon DynamoDB V2 native data types

Amazon DynamoDB V2 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.

You can edit the data types of the source and target DynamoDB fields to change the Amazon DynamoDB V2 native data type, transformation data type, precision, and scale of fields.

Amazon DynamoDB V2 and transformation data types

The following table lists the Amazon DynamoDB V2 data types that Data Integration supports and the corresponding transformation data types:

Amazon DynamoDB V2 data type	Description	Transformation data type	Description
Binary	An ordered list of zero or more bytes.	Binary	An ordered list of zero or more bytes.
Boolean	True or false.	Integer	Standard signed int32 integer.
List	An ordered list of zero or more elements of mixed types.	Array	An ordered list of zero or more elements of a single fixed type.

Amazon DynamoDB V2 data type	Description	Transformation data type	Description
Map	A dynamic list of key-value pairs of string key-type and dynamic, mixed value-types.	Struct	A fixed list of key-value pairs of string key-type and mixed but fixed value-types.
Number	A floating point number with 38 decimal digits of precision and a floating scale.	Double	IEEE double-precision floating-point number that has approximately 16 decimal digits of precision.
Set	An unoredered list of zero or more unique elements of a single fixed type (Binary, Number, or String).	Array	An ordered list of zero or more elements of a single fixed type.
String	A Unicode string encoded in UTF-8.	String	A Unicode string.

All DynamoDB columns, except for the hash key and range key columns, can be null. Null values are distinct from zero-length values, including binary, lists, and strings.

When you use certain midstream transformations, such as the Hierarchy-to-Hierarchy (H2H) transformation, the order of the elements might change in an array value.

INDEX

A

administration
  Authentication 7
Amazon DynamoDB V2
  connection properties 10
  rules and guidelines 17
Amazon DynamoDB V2 Connector
  overview 6
  supported task and object types 6
Amazon DynamoDB V2 data types
  overview 18

C

Cloud Application Integration community
  URL 4
Cloud Developer community
  URL 4
connection properties
  overview 10
connections
  Amazon DynamoDB V2 10

D

Data Integration community
  URL 4

I

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

L

lookup
  database 16

M

maintenance outages 5
mapping
  overview 12
mapping in advanced mode
  Amazon DynamoDB V2 targets 14
mapping tasks
  Amazon DynamoDB V2 12
mappings
  lookup overview 16
  lookup properties 16

S

status
  Informatica Intelligent Cloud Services 5
system status 5

T

temporary security credentials
  policy 8
trust site
  description 5

U

upgrade notifications 5

W

web site 4