AmazonDynamoDB Adapter for SAP Integration Suite

Version 1.0.1 – November 2025

THE BEST RUN SAP

Contents

1 Introduction

1.1 Introduction

This is the official guide for the AmazonDynamoDB Adapter for SAP Integration Suite. This guide covers all relevant information for integration developers to start working with the AmazonDynamoDB Adapter. Read this guide carefully before using the Adapter.

1.2 Coding Samples

Any software coding and/or code lines/strings ("Code") included in this documentation are only examples and are not intended to be used in a productive system environment. The Code is only intended to better explain and visualize the syntax and phrasing rules of certain coding. The correctness and completeness of the Code given herein are not guaranteed.

1.3 Internet Hyperlinks

The documentation may contain hyperlinks to the Internet. These hyperlinks are intended to serve as a hint about where to find related information. The availability and the correctness of this related information or the ability of this information to serve a particular purpose are not warranted.

1.4 Overview

Amazon DynamoDB is a fully managed NoSQL database service that provides reliable and fast performance with seamless scalability. DynamoDB provides an easily accessible persistent key–value datastore which facilitates scaling, data replication, encryption at rest, and on-demand backup among other features.

Amazon DynamoDB use case includes:

- Store semi-structured schema items
- Automatic Data Backups
- Reliable Data Access for Streaming Platforms

1.5 Features

- Simple and Flexible Schema**:** Our adapter allows you to leverage the schema-less design of Amazon DynamoDB. You can now utilise a flexible schema design and enhanced querying capability for your data.
- Predictable Performance: Avoid worries about managing workloads, with AmazonDynamoDB Adapter, latencies are predictable and reliable performance can be expected.
- Consistent Reads: The AmazonDynamoDB Adapter provides you the option to ensure that data read from the database is consistent. This option will return the latest data while reflecting changes from successful write operations.
- Dynamic configuration with headers and properties: Assigning dynamic values to different properties allows enhanced flexibility to your integration flows. You can also refer to dynamic parameters using SAP Cloud Integration exchange headers and properties.
- Ease of Access and Connectivity**:** The adapter facilitates and accelerates the implementation time of connecting to Amazon DynamoDB.

2 Installation and Configuration

This section details the prerequisites to install and configure the AmazonDynamoDB Adapter. The AmazonDynamoDB Adapter is available as part of your SAP Integration Suite license.

2.1 Adapter Installation on Cloud Foundry

Before the AmazonDynamoDB Adapter can be used in the Cloud Foundry environment, it must be deployed to the SAP Integration Suite tenant.

2.1.1 Prerequisite

To deploy the AmazonDynamoDB Adapter, you must have access to "*AmazonDynamoDB Adapter for SAP Integration Suite*" as part of your SAP Integration Suite license.

2.1.2 Installation

> The below installation procedure is compatible with Apache Camel 2, Apache Camel 3, and Edge Integration Cell (EIC) platform.

You can deploy the adapter using the following methods:

2.1.2.1 Adapter Installation by Creating a New Integration Flow

The AmazonDynamoDB Adapter is available for selection in the receiver adapter list and can be deployed in the **Design** tab directly as you use it in an Integration flow.

Purpose

To install an adapter for use in your Integration flow.

Procedure

Go to **Design** workspace and select the integration package where you want to create a new Integration flow.

1. Click **Edit** to make the package editable.
2. Go to the **Artifacts** tab. Click **Add** and select **Integration Flow**.
3. Enter the **Name** and **ID** for your flow. Additionally, select **Runtime Profile** from the drop-down and choose **Sender** and **Receiver** systems from the list . Finally, click **Add** to create the integration flow.
4. Go to the newly created integration flow and click **Edit** to make it editable.

5. In the integration flow, click **End** to add a **Connector** between the **End** and the **Receiver Box**. A drop-down with the available adapters appears. The **AmazonDynamoDB** adapter should show up in the list.
6. Select the **AmazonDynamoDB** adapter from the list. The adapter is now imported which *triggers* an adapter deployment. Once the AmazonDynamoDB Adapter is deployed, a success message is *displayed*.

After the above steps are done, the AmazonDynamoDB Adapter is successfully deployed in your Design workspace of the SAP Integration Suite tenant.

2.1.2.2 Adapter Installation without Creating a New Integration Flow

> The following procedure describes how the AmazonDynamoDB Adapter is migrated from the Discover workspace to the Design workspace of the SAP Integration tenant.

This method is useful for scenarios where integration flow packages are migrated from development to a higher environment such as Production.

The AmazonDynamoDB Adapter can be imported into the Design workspace without creating an integration flow. Use the Transport Management Service (TMS) to import/transport the AmazonDynamoDB Adapter to a higher environment. Alternatively, if the TMS is not available in the landscape, the adapter package can be imported into the Design workspace by copying it from the Discover workspace.

Purpose

To import the AmazonDynamoDB Adapter to **Design** workspace by copying the integration package from **Discover** workspace.

Procedure

1. Go to **Discover** workspace.
2. In the search box, search for **AmazonDynamoDB Adapter for SAP Integration Suite** package.
3. Select the package and click **Copy**.
   This copies the package from the Discover workspace to Design workspace.
4. Go to Design workspace and select the copied **AmazonDynamoDB Adapter for SAP Integration Suite** package.
5. In the **Actions** tab of the selected package, click **Deploy**.
   This completes the adapter deployment to Design workspace.

2.1.3 Monitor the Deployment Status

After the adapter deployment is complete, you can check the status in the **Monitor** section.

Purpose

To check the status of the deployed adapter.

Procedure

1. Under the **Monitor** tab, click **Integrations and APIs**. This opens the **Overview** page.
2. On the **Overview** page, go to **Manage Integration Content** section and click **All**. This opens **Integration Content** page with a list of all the deployed adapters.

Here, you can check and confirm the deployment status of your adapter.

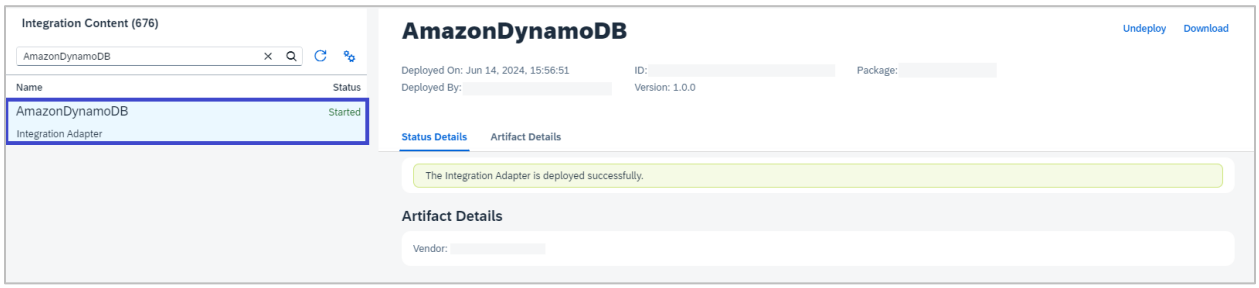

3 Getting Started: AmazonDynamoDB Adapter

This section explains how you can configure the AmazonDynamoDB Adapter for SAP Integration Suite. You can find information about adapter architecture, application configuration, and adapter configuration for AmazonDynamoDB Adapter. Finally, you can find additional information and sample scenarios under References and troubleshooting under Support.

3.1 Architecture Overview

The AmazonDynamoDB Adapter is designed to be employed as a receiver adapter. In such a scenario, SAP Cloud Integration acts as the initiator of the calls. There are various CRUD operations that can be performed at an **item** level in your Amazon DynamoDB tables.

The image below gives a high-level representation of how the AmazonDynamoDB Adapter works. The AmazonDynamoDB Adapter receives operation request from SAP Cloud Integration and interacts with the Amazon DynamoDB target table to execute the same.

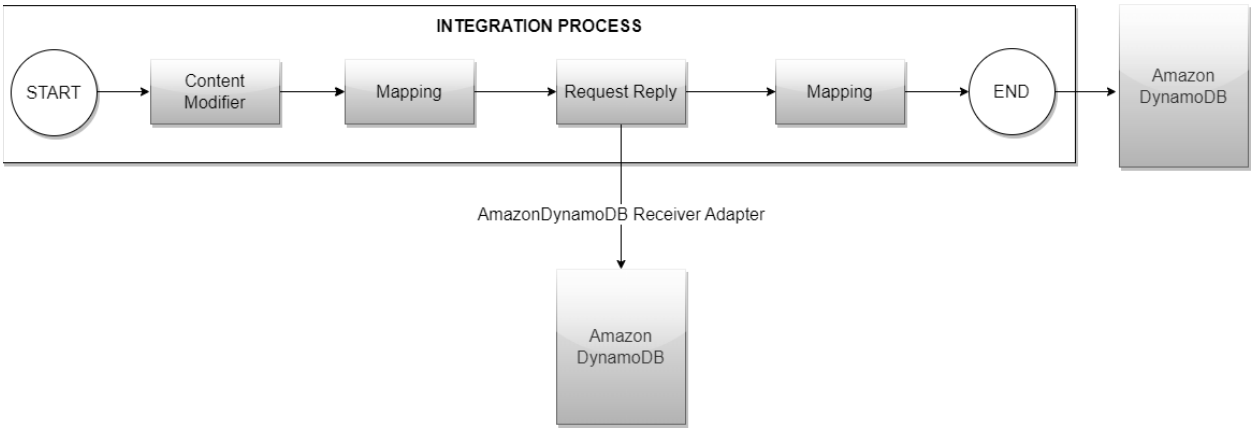

3.2 Application Configuration

You can connect to Amazon DynamoDB via the AmazonDynamoDB Adapter using your user credentials f

Before you begin, the following links can be helpful to set up your Amazon DynamoDB application.

- To create a User Account, see Create an Amazon account.
- For more information on how to set up credentials for programmatic access, see Configure your credentials.

3.3 Authentication

The AmazonDynamoDB Adapter allows creation Amazon credentials for your AWS user. For each user, you can create an Access Key and a Secret Key.

The adapter uses common security artifacts stored in SAP Cloud Integration. You must use **Secure Parameter** artifact to safely store Access Key and Secret Key. These security artifacts can then be accessed in the adapter using aliases.

Before setting up the authentication, you must create the Credentials in **Security Material** in the SAP Integration Suite as described in the next section.

3.3.1 Creating Security Parameter in Security Material

Purpose

To create credentials for authentication.

Procedure

1. In SAP Integration Suite, navigate to **Monitor** > **Integrations and APIs**. This opens the **Overview** page.

2. On the **Overview** page, go to **Manage Security** section and click **Security Material**.

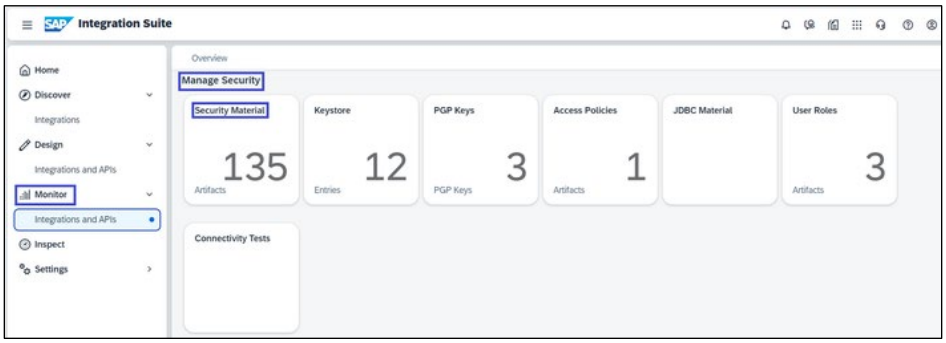

3. On the **Manage Security Material** page, click **Create** to select **Secure Parameter** from the dropdown.

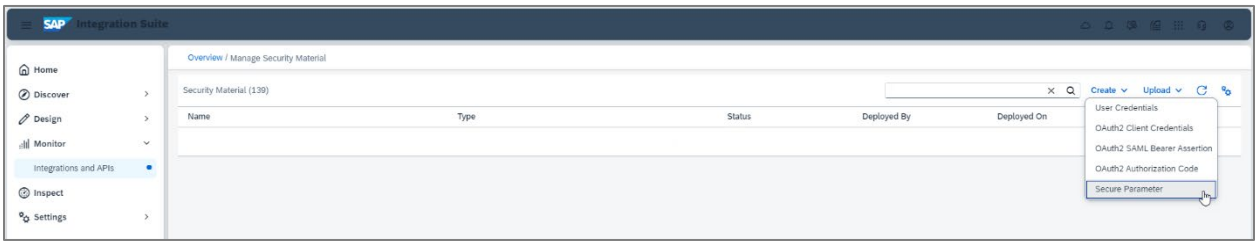

4. In the **Create Secure Parameter** popup, provide the below details.

Field	Description
Name	Specify the name of the security artifact. The artifact name is used as an alias for the confidential data.
Description	Enter a description for the artifact (optional).
Secure Parameter	Enter the confidential value of the attribute. The permissible length of the secure parameter for Cloud Foundry is a maximum of 4096 characters.
Repeat Secure Parameter	Repeat the confidential value of the attribute.

5. Click **Deploy** to complete the process.

The Security artifact created above is used to connect to the Amazon DynamoDB Application by configuring the **Connection tab** of the Adapter. For more details, see Connection.

3.4 Supported Operations

The operations supported by the AmazonDynamoDB Adapter are listed below:

- Put
- Update
- Delete
- Scan
- Query
- Execute Statement

AmazonDynamoDB Adapter currently supports operations for **item** components only.

4 AmazonDynamoDB Adapter Configuration

This section describes the parameters to be configured for your AmazonDynamoDB Adapter. You need to configure the **Connection** and **Processing** tabs. A description and example usage for every field has been added.

4.1 General Tab

The **General** tab provides an overview of basic adapter information including **Channel** and **Adapter** details.

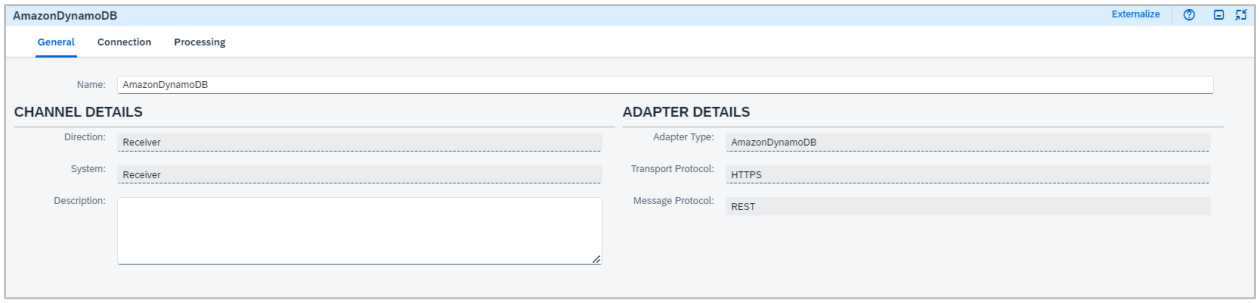

Field	Description
Name	Specify the name of the adapter flow.
Description	Specify the description of the adapter.

4.2 Connection Tab

The Connection tab contains connection and authentication parameters for Amazon DynamoDB.

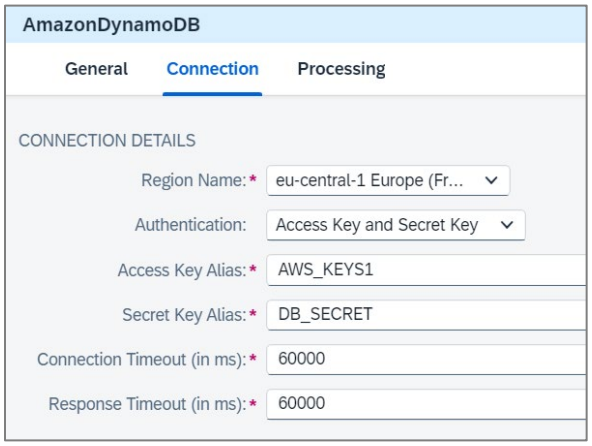

Field	Description
Region Name	Select the AWS Region where the S3 Bucket resides. If you are manually specifying this value, use `eu-central-1` instead of `eu-central-1 Europe (Frankfurt)`.
Authentication	The authentication method used is **Access Key** and **Secret Key.**
Access Key Alias	Specify the name of the Secure Parameter which stores the AWS Access Key.
Secret Key Alias	Specify the name of the Secure Parameter which stores the AWS Secret Key.
Connection Timeout (in ms)	Specify the maximum waiting time (in milliseconds) for the connection to be established. Example: `3000`
Response Timeout (in ms)	Specify the maximum waiting time (in milliseconds) for a response message. Example: `3000`

4.3 Processing Tab

The Processing tab contains all the operational configurations for the AmazonDynamoDB Adapter. A sample configuration for the **Delete** operation is shown in the image below.

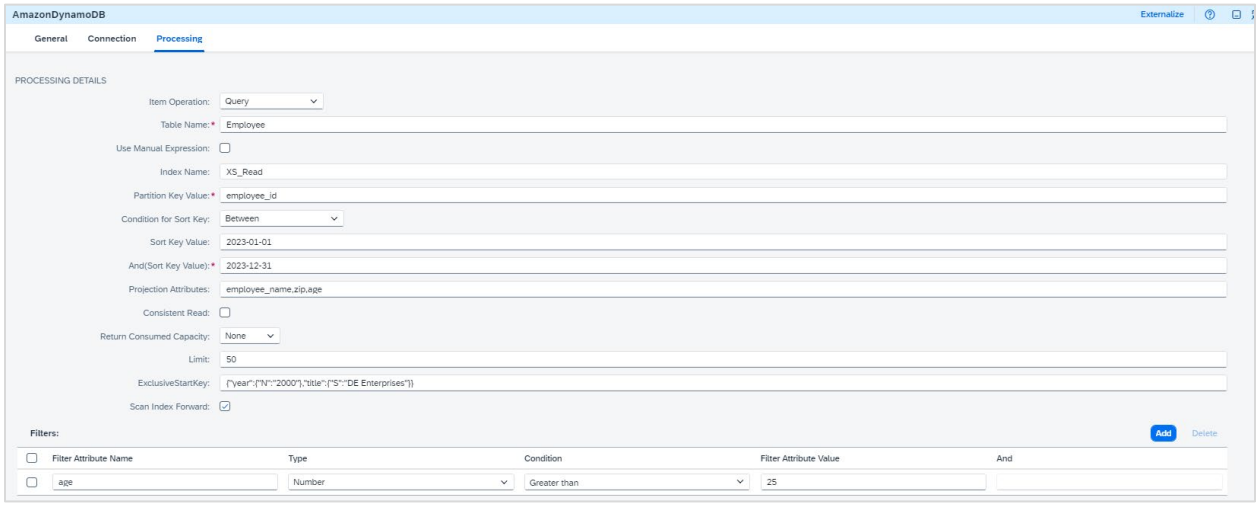

The **Processing tab** contains all the operational configurations for the AmazonDynamoDB Adapter.

The processing tab contains the following fields:

Field	Description
Item Operation	Select the operation to be performed on the item(s).
Table Name	Specify the table in Amazon DynamoDB.
Statement	Specify the PartiQL statement representing the operation to be executed.
Next Token	Specify value to get the remaining results, if **Next Token** was returned in the statement response.
Consistent Read	Enable to opt for the read consistency model. If enabled the operation uses strongly consistent reads otherwise it uses eventually consistent reads.
Return Consumed Capacity	Select the level of detail about the provisioned throughput consumption to be returned in the response. Acceptable Values: `Indexes, None, Total`
Index Name	Specify the secondary index attribute. Example: `XS_Read`
Use Manual Expression	Enable to use manual expression for filter and projection.
Partition Key Value	Specify the value for partition key. Amazon DynamoDB uses the partition key's value to determine the partition for storing an item. Example: `product_id`
Condition for Sort Key	Select the conditional operator to be applied on the Sort Key: For more information, see Condition Expressions for Sort Key.
Sort Key Value	Specify the value for Sort Key. Sort Key determines how data with the same partition is ordered. Example: `timestamp`
And (Sort Key)	Specify the upper limit for Sort Key if the **Condition for Sort Key** is selected as `between`. Example: `timestamp`
Projection Attributes	Specify the attributes that you want to project. The attributes must be comma separated. Example: `Account_ID,creationDate,CustomerName`

Field	Description
Limit	Specify the maximum number of items to be evaluated before the matching condition returns the result set of rows. Note: **Limit** applies to the result set fetched from the table before **Filter** is applied. Example: `50`
ExclusiveStartKey	Specify the value that was returned for the LastEvaluatedKey in the previous operation for retrieving the next page. Example: `{"year":{"N":"2000"},"title":{"S":"DE Enterprises"}}`
Scan Index Forward	Select the order for index traversal. Enable to perform traversal in ascending order; disable to perform traversal in descending order.
Filters	Set a filter condition on an attribute as a post-processing action on the records fetched. Filter Attribute Name `Create_Date` Datatype `Number` Filter Attribute Value `Between` Condition `2023-01-01` And `2023-12-12`
Request Headers	Enter a list of custom headers, separated by a pipe (|), to send to the target system. By default, no custom headers are sent. Use an asterisk (*) to send all custom headers to the target system.
Response Headers	Enter a list of headers coming from the target system's response, separated by a pipe (|), to be received in the message. Use an asterisk (*) to receive all the headers from the target system, which is also the default value. Default: `*`

5 AmazonDynamoDB Adapter Operations

This section contains a comprehensive description of all the operations supported by the AmazonDynamoDB Adapter.

5.1 Put

`Put` allows the creation of any new item. The request payload should be provided with all the required properties and field attributes to be created in the request.

Sample configuration for `Put` operation:

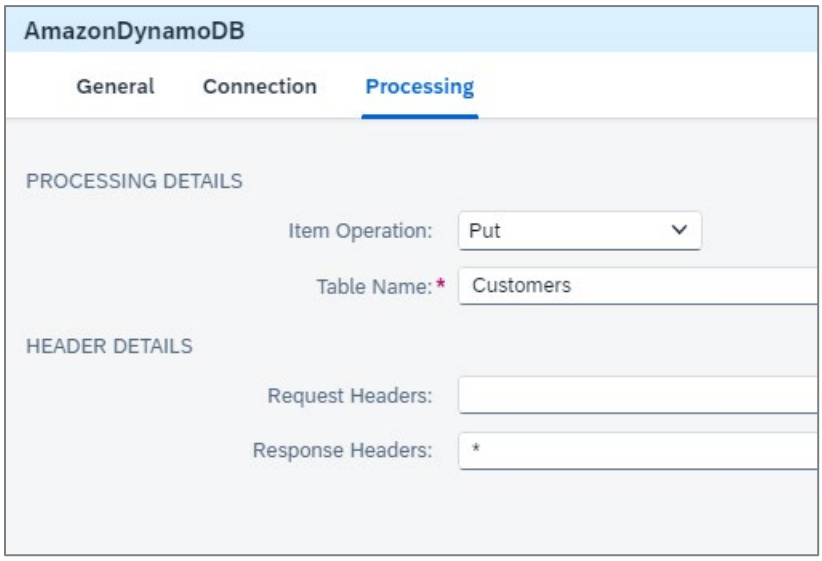

Field	Description
Item Operation	Select the operation to be performed on the above item: `Put`
Table Name	Specify the target table in Amazon DynamoDB. Example: `Customers` This is a mandatory field.

Sample request payload for `Put` operation:

```
{
    "Item": {
        "customerName": {
            "S": "BZD"
        },
        "origin_country": {
            "S": "France"
        },
        "numberOfCustomers": {
            "N": "12304"
        },
        "product_description": {
            "M": {
                "available_colors": {
                    "SS": [
                        "Red",
                        "Blue",
                        "Yellow"
                    ]
                },
                "product_id": {
                    "N": "11"
                },
                "purchase_amount_in_last2years": {
                    "NS": [
                        "22000",
                        "23000"
                    ]
                }
            }
        },
        "Accountid": {
            "S": "1007"
        },
        "creation_date": {
            "S": "2024-06-07"
        }
    },
    "ReturnValues": "ALL_OLD",
    "ReturnConsumedCapacity": "TOTAL"
}
```

5.2 Update

This operation allows the update of any new item. The values should be provided in the payload with all the required properties and field attributes to be created in the request.

Sample Configuration for `Update` operation:

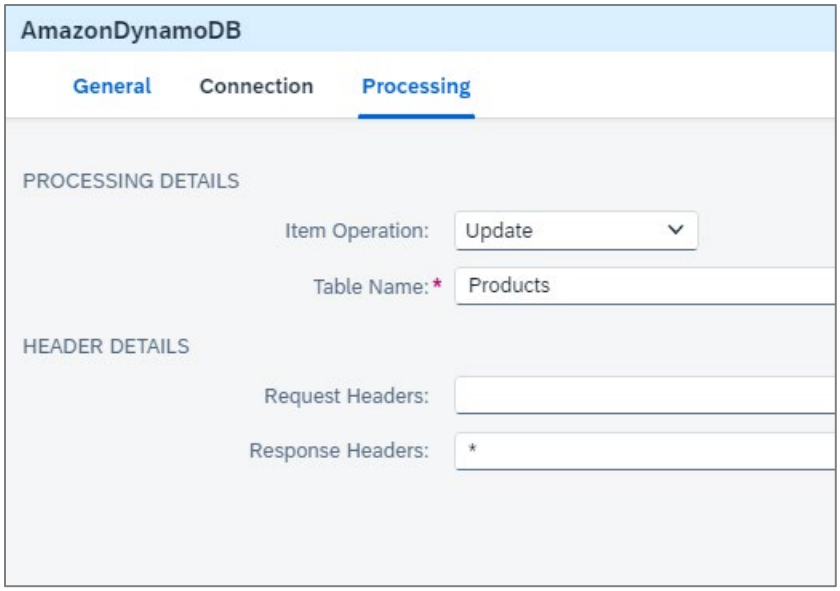

The configuration details are as follows:

Field	Description
Item Operation	Select the operation to be performed on the above item: `Update`.
Table Name	Specify the target table in Amazon DynamoDB. Example: `Products` This is a mandatory field.

Sample request payload for `Update` operation:

```
{
    "Key": {
        "CustomerID": {
            "N": 3
        },
        "OrderID": {
            "N": 103
        }
    },
    "UpdateExpression": "SET Active = :active",
    "ConditionExpression": "Age >= :age",
    "ExpressionAttributeValues": {
        ":age": {
            "N": 0
        },
        ":active": {
            "BOOL": false
        }
    },
    "ReturnValues": "UPDATED_NEW",
    "ReturnConsumedCapacity": "TOTAL"
}
```

5.3 Delete

This operation allows the deletion of any item. The request payload should provide the details of the item to be deleted.

Sample configuration for `Delete` operation:

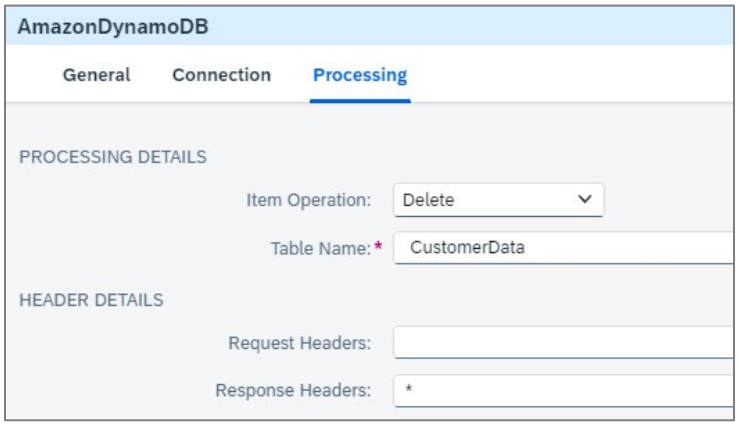

Sample payload for `Delete` operation.

```
{
   "Key": {
      "Account_id": {
         "S": "3433"
      },
      "creation_date": {
         "S": "2022-11-18"
      }
   },
   "ConditionExpression": "qr_code = :qrCodeValue AND item_ordered = :itemOrderedValue",
   "ExpressionAttributeValues": {
      ":qrCodeValue": {
         "B": "U29tZSBkYXRhIGluIFhNTA == "
      },
      ": itemOrderedValue": {
         "BOOL": true
      }
   },
   "ReturnValues": "ALL_OLD"
}
```

The configuration details are as follows:

Field	Description
Item Operation	Select the operation to be performed on the above item: `Delete`
Table Name	Specify the target table in Amazon DynamoDB. Example: `CustomerData` This is a mandatory field.

5.4 Scan

The operation allows you to search through the entire table. You must specify the **Table Name** and attributes under **Projection Attributes**. You can also receive customized responses by using **Filters** and **Limit**.

Sample configuration for `Scan` operation:

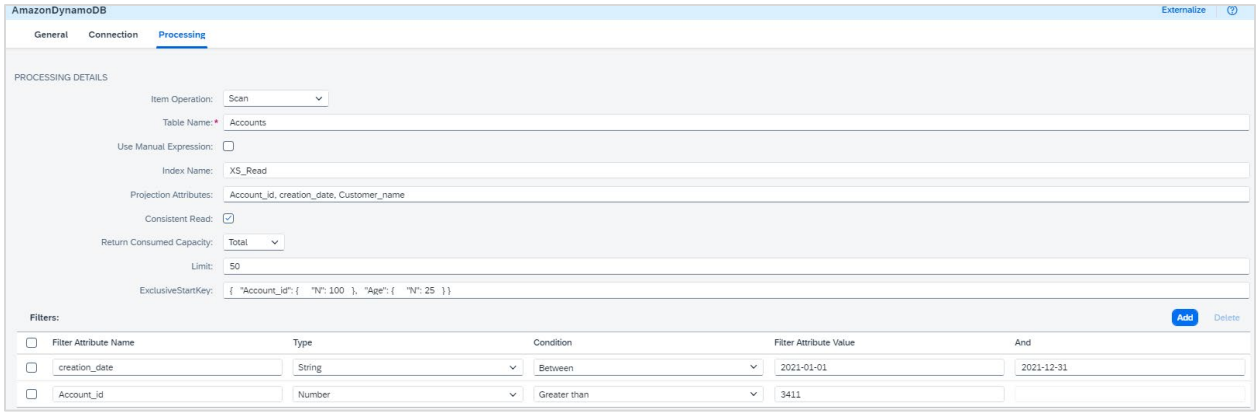

Alternatively, you can enable **Use Manual Expression** to Scan using the payload.

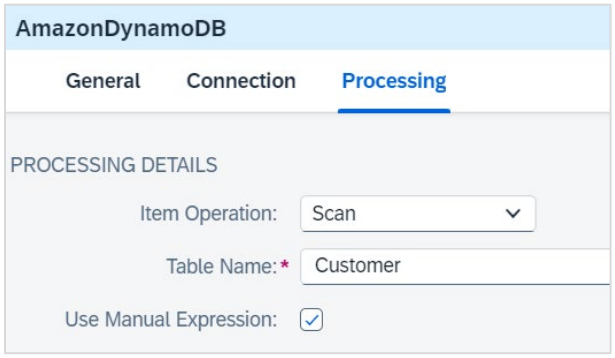

Sample payload for Scan operation.

```
{
    "ProjectionExpression": "CustomerID,OrderID,City,FirstName,LastName,Age,Active",
    "FilterExpression": "#Age >= :age AND #Active = :flag",
    "ExpressionAttributeNames": {
        "#Age": "Age",
        "#Active": "Active"
    },
    "ExpressionAttributeValues": {":age": {"N": "0"},
        ":flag": { "BOOL": false }
    },
    "ConsistentRead": true,
    "ReturnConsumedCapacity": "TOTAL"
}
```

Field	Description
Item Operation	Select the operation to be performed: `Scan`
Table Name	Specify the table in Amazon DynamoDB. Example: `Accounts` This is a mandatory field.
Index Name	Specify the name of the secondary index to scan. Example: `XS_Read`
Use Manual Expression	Enable to use manual expression for filter and projection. For more information, see sample payload for Scan.
Projection Attributes	Specify the attributes that you want to project. The attributes must be comma separated. Example: `Employee_Name,EID`
Consistent Read	Select the read consistency model. If enabled the operation uses strongly consistent reads otherwise it uses eventually consistent reads.
Return Consumed Capacity	Select the level of detail about the provisioned throughput consumption to be returned in the response. Example: `Total`
Limit	Specify the maximum number of items to be evaluated before the matching condition returns the result set of rows. **Limit** applies to the result set fetched from the table before **Filter** is applied. Example: `50`
ExclusiveStartKey	Specify the value that was returned for the `LastEvaluatedKey` in the previous operation, in order to retrieve the next page. Example: `{"year":{"N":"2000"},"title":{"S":"DE Enterprises"}}`

Field	Description

Filters	Set a filter condition on an attribute as a post-processing action on the records fetched.
	
	`Filter Attribute Name` `Create_Date`
	`Datatype` `Number`
	`Filter Attribute Value` `Between`
	`Condition` `2023-01-01`
	`And` `2023-12-12`

5.5 Query

You can use the **Query** to search for data records in the same partition. This operation is used to find items based on primary key values. Specify the **Partition Key Value** for that partition key attribute. Query returns all items with the input partition key value.

Sample configuration for `Query` operation:

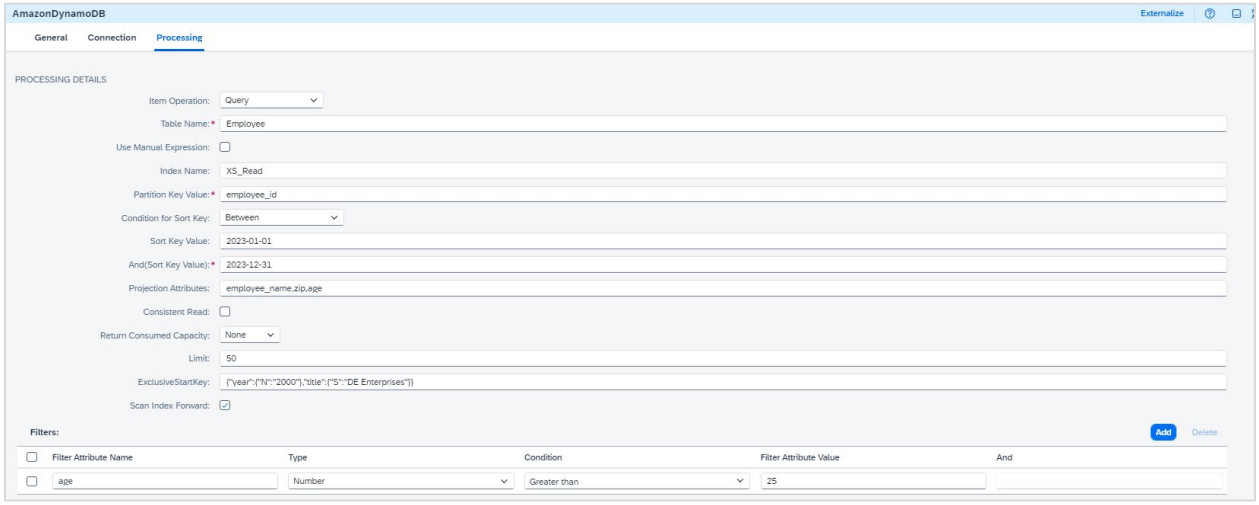

Alternatively, you can enable **Use Manual Expression** to Query using the payload.

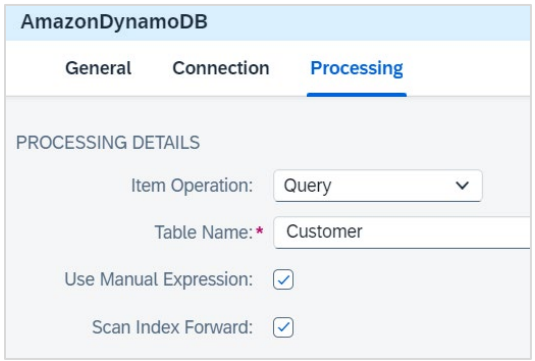

Sample payload for `Query` operation.

```
{
    "ProjectionExpression": "CustomerID,OrderID,City,FirstName,LastName,Age,Active",
    "FilterExpression": "#Age >= :age AND #Active = :flag",
    "KeyConditionExpression": "#partitionKeyName = :partitionkeyval and #sortkey =
:sortkeyval",
    "ExpressionAttributeNames": {
        "#Age": "Age",
        "#Active": "Active",
        "#partitionKeyName": "CustomerID",
        "#sortkey": "OrderID"
    },
    "ExpressionAttributeValues": {
        ":partitionkeyval": {
            "N": "1"
        },
        ":sortkeyval": {
            "N": "1"
        }
    }
}
```

A description of the fields is given in the table below.

Field	Description
Item Operation	Select the operation to be performed on the above item: `Query`
Table Name	Specify the target table in Amazon DynamoDB. Example: `Customer` This is a mandatory field.
Index Name	Specify the name of the secondary index to query. Example: `XS_Get`
Use Manual Expression	Manually specify the attributes for filter and projection. For more information, see sample payload for Query.
Partition Key Value	Specify the value for partition key. Amazon DynamoDB uses the partition key's value to determine the partition for storing an item. Example: `product_id`
Condition for Sort Key	Select the conditional operator to be applied on the Sort Key: For more information, see Condition Expressions for Sort Key.

Field	Description
Sort Key Value	Specify the value for Sort Key. Sort Key determines how data with the same partition is ordered. Example: `timestamp` Sort key is part of both Table's composite key as well as Index's composite key. a) If `Index Name` is provided, it will utilise Index's composite key. Example: Index's Composite key is (CustomerID + OrderID) where OrderID is the sort key for local or global index and Partition Key is CustomerID. b) `Index Name` it is blank, it will utilize Table's composite key. Example: Table's Composite key is (CustomerID + DepartmentID) where DepartmentID is the sort key for the Table and Partition Key is CustomerID.
And (Sort Key)	Specify the upper limit for Sort Key if the **Condition for Sort Key** is selected as `between`. Example: `timestamp`
Projection Attributes	Specify the attributes that you want to project. Example: `Account_ID,creationDate,CustomerName`
Consistent Read	Select the read consistency model. If enabled the operation uses strongly consistent reads otherwise it uses eventually consistent reads.
Return Consumed Capacity	Select the level of detail about the provisioned throughput consumption to be returned in the response. Example: `Total`
Limit	Specify the maximum number of items to be returned. Example: `50`

Field	Description
ExclusiveStartKey	Specify the value that was returned for the `LastEvaluatedKey` in the previous operation, in order to retrieve the next page. Example: `{"year":{"N":"2000"},"title":{"S":"DE Enterprises"}}`
Scan Index Forward	Enable to specify the order of index traversal. If enabled, traversal is performed in ascending order. If unchecked, traversal is performed in descending order.

Filters	Set a filter condition on an attribute as a post-processing action on the records fetched.
	
	`Filter Attribute Name` `Create_Date`
	`Datatype` `Number`
	`Filter Attribute Value` `Between`
	`Condition` `2023-01-01`
	`And` `2023-12-12`

5.6 Execute Statement

This operation allows you to perform reads and singleton writes using PartiQL on the data stored in Amazon DynamoDB.

Sample payload for `execute` statement

This payload only works for when **Statement** is set to

```
pm.request.headers.add({key: 'PartiQL_SQL', value: 'SELECT CustomerID,OrderID FROM
Test_Exploratory WHERE CustomerID IN (?,?,?,?,?) '});
```

```
{
   "Parameters": [
      {
         "Value1": 10
      },
      {
         " Value2": "20"
      },
      {
         " Value3": 30
      },
      {
         " Value4": "40"
      },
      {
         " Value5": "50"
      }
   ],
   "Limit": 2,
   "ReturnConsumedCapacity": "TOTAL"
}
```

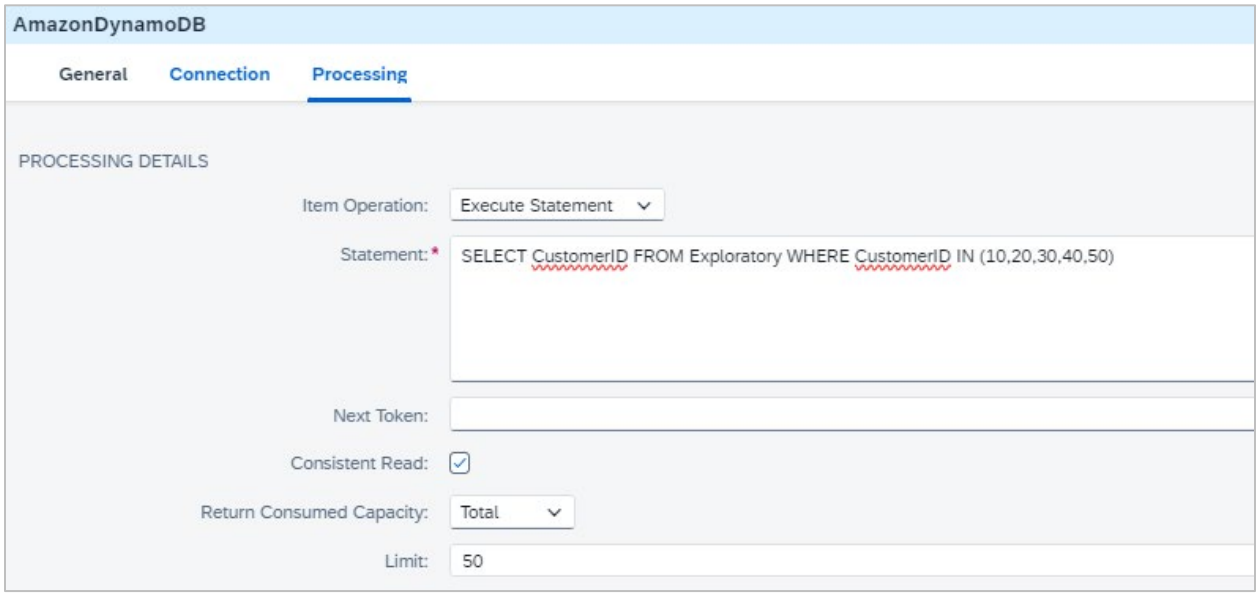

For PartiQL reads (`SELECT` statement), if the total number of processed items exceeds the maximum dataset size limit of 1 MB, the read stops and results are returned to the user as a `LastEvaluatedKey` value to continue the read in a subsequent operation. If the filter criteria in `WHERE` clause does not match any data, the read will return an empty result set.

If `LastEvaluatedKey` is present in the response, you need to paginate the result set. If `NextToken` is present, you need to paginate the result set and include `NextToken`.

The configuration details are as follows:

Field	Description
Item Operation	Select the operation to be performed on the above item: `Execute Statement`

Field	Description
Statement	Specify the partiQL statement representing the action to be performed. Example: `SELECT CustomerID FROM Exploratory WHERE CustomerID IN (10,20,30,40,50)`
Next Token	Specify this value to get the remaining results, if Next Token was returned in the statement response. Example: `QqezbsHJMVZ1ujF6kHDk1QhcWxmAYwRXQZkhbGa2/XIzoMhy11gAj2VQtGgQmBWl4K19MRFg34NVoG5M9PgDt3O37R6XvgV/6UZr7XbDDXDxnr5OrIhshp8X8NHua/QNUIf6yvt0u4KG3iTQWYMIQzqrccILj6qAaLpcmllJakbXOk9Dm98ya6c95yDOfaw/bu7hfoESGVm6LQEP1y01rLLcrD7WcngzRbGAPyyF2oFlKk57qBn/YM1OmbEA/ioRvzXzfqwP7E78d2GpJYVim4jdHRIZZQbkYHEqq1nav/IkKHd/hO9t66bIzyJdWdwOWKcbJroqmuh8nMoNgzihliJKX+hatFT+/UXtkU2erbUHpTYYK03pYvms=`
Consistent Read	Select the read consistency model. If enabled the operation uses strongly consistent reads otherwise it uses eventually consistent reads.
Return Consumed Capacity	Select the level of detail about the provisioned throughput consumption to be returned in the response. Example: `Total`
Limit	Specify the maximum number of items to be returned. Example: `50`