

# Quarkus - Amazon DynamoDB Client

DynamoDB is a scalable AWS managed NoSQL database. It supports both key-value and document data models, that enables to have a flexible schema for your data. This extension provides functionality that allows the client to communicate with the service when running in Quarkus. You can find more information about DynamoDB at [the Amazon DynamoDB website](#).



The DynamoDB extension is based on [AWS Java SDK 2.x](#). It's a major rewrite of the 1.x code base that offers two programming models (Blocking & Async). Keep in mind it's actively developed and does not support yet all the features available in SDK 1.x such as [Document APIs](#) or [Object Mappers](#)



This technology is considered preview.

In *preview*, backward compatibility and presence in the ecosystem is not guaranteed. Specific improvements might require to change configuration or APIs and plans to become *stable* are under way. Feedback is welcome on our [mailing list](#) or as issues in our [GitHub issue tracker](#).

For a full list of possible extension statuses, check our [FAQ entry](#).

The Quarkus extension supports two programming models:

- Blocking access using URL Connection HTTP client (by default) or the Apache HTTP Client
- [Asynchronous programming](#) based on JDK's [CompletableFuture](#) objects and the Netty HTTP client.

In this guide, we see how you can get your REST services to use the DynamoDB locally and on AWS.

## Prerequisites

To complete this guide, you need:

- JDK 1.8+ installed with [JAVA\\_HOME](#) configured appropriately
- an IDE
- Apache Maven 3.6.3
- An AWS Account to access the DynamoDB service
- Optionally, Docker for your system to run DynamoDB locally for testing purposes

## Setup DynamoDB locally

The easiest way to start working with DynamoDB is to run a local instance as a container.

```
docker run --publish 8000:8000 amazon/dynamodb-local:1.11.477 -jar
DynamoDBLocal.jar -inMemory -sharedDb
```

This starts a DynamoDB instance that is accessible on port **8000**. You can check it's running by accessing the web shell on <http://localhost:8000/shell>.

Have a look at the [Setting Up DynamoDB Local guide](#) for other options to run DynamoDB.

Open <http://localhost:8000/shell> in your browser.

Copy and paste the following code to the shell and run it:

```
var params = {
  TableName: 'QuarkusFruits',
  KeySchema: [{ AttributeName: 'fruitName', KeyType: 'HASH' }],
  AttributeDefinitions: [{ AttributeName: 'fruitName',
AttributeType: 'S', }],
  ProvisionedThroughput: { ReadCapacityUnits: 1,
WriteCapacityUnits: 1, }
};

dynamodb.createTable(params, function(err, data) {
  if (err) ppJson(err);
  else ppJson(data);

});
```

## Set up Dynamodb on AWS

Before you can use the AWS SDKs with DynamoDB, you must get an AWS access key ID and secret access key. For more information, see [Setting Up DynamoDB \(Web Service\)](#).

We recommend to use the AWS CLI to provision the table:

```
aws dynamodb create-table --table-name QuarkusFruits \
                           --attribute-definitions
AttributeName=fruitName,AttributeType=S \
                           --key-schema
AttributeName=fruitName,KeyType=HASH \
                           --provisioned-throughput
ReadCapacityUnits=1,WriteCapacityUnits=1
```

# Solution

The application built here allows to manage elements (fruits) stored in Amazon DynamoDB.

We recommend that you follow the instructions in the next sections and create the application step by step. However, you can go right to the completed example.

Clone the Git repository: `git clone https://github.com/quarkusio/quarkus-quickstarts.git`, or download an [archive](#).

The solution is located in the `amazon-dynamodb-quickstart` directory.

## Creating the Maven project

First, we need a new project. Create a new project with the following command:

```
mvn io.quarkus:quarkus-maven-plugin:1.5.0.Final:create \
  -DprojectId=org.acme \
  -DprojectArtifactId=amazon-dynamodb-quickstart \
  -DclassName="org.acme.dynamodb.FruitResource" \
  -Dpath="/fruits" \
  -Dextensions="resteasy-jsonb,amazon-dynamodb,resteasy-mutiny"
cd amazon-dynamodb-quickstart
```

This command generates a Maven structure importing the RESTEasy/JAX-RS, Mutiny and DynamoDB Client extensions. After this, the `amazon-dynamodb` extension has been added to your `pom.xml` as well as the Mutiny support for RESTEasy.

## Creating JSON REST service

In this example, we will create an application to manage a list of fruits. The example application will demonstrate the two programming models supported by the extension.

First, let's create the `Fruit` bean as follows:

```
package org.acme.dynamodb;

import java.util.Map;
import java.util.Objects;

import io.quarkus.runtime.annotations.RegisterForReflection;
import software.amazon.awssdk.services.dynamodb.model.AttributeValue;

@RegisterForReflection
public class Fruit {
```

```

private String name;
private String description;

public Fruit() {
}

public static Fruit from(Map<String, AttributeValue> item) {
    Fruit fruit = new Fruit();
    if (item != null && !item.isEmpty()) {
        fruit.setName(item.get(AbstractService.FRUIT_NAME_COL)
.s());
        fruit.setDescription(item.get(AbstractService
.FRUIT_DESC_COL).s());
    }
    return fruit;
}

public String getName() {
    return name;
}

public void setName(String name) {
    this.name = name;
}

public String getDescription() {
    return description;
}

public void setDescription(String description) {
    this.description = description;
}

@Override
public boolean equals(Object obj) {
    if (!(obj instanceof Fruit)) {
        return false;
    }

    Fruit other = (Fruit) obj;

    return Objects.equals(other.name, this.name);
}

@Override
public int hashCode() {
    return Objects.hash(this.name);
}
}

```

Nothing fancy. One important thing to note is that having a default constructor is required by the JSON serialization layer. The static `from` method creates a bean based on the `Map` object provided by the DynamoDB client response.

Now create a `org.acme.dynamodb.AbstractService` that will consist of helper methods that prepare DynamoDB request objects for reading and adding items to the table.

```
package org.acme.dynamodb;

import java.util.HashMap;
import java.util.Map;

import software.amazon.awssdk.services.dynamodb.model.AttributeValue;
import software.amazon.awssdk.services.dynamodb.model.GetItemRequest;
import software.amazon.awssdk.services.dynamodb.model.PutItemRequest;
import software.amazon.awssdk.services.dynamodb.model.ScanRequest;

public abstract class AbstractService {

    public final static String FRUIT_NAME_COL = "fruitName";
    public final static String FRUIT_DESC_COL = "fruitDescription";

    public String getTableName() {
        return "QuarkusFruits";
    }

    protected ScanRequest scanRequest() {
        return ScanRequest.builder().tableName(getTableName())
            .attributesToGet(FRUIT_NAME_COL, FRUIT_DESC_COL)
            .build();
    }

    protected PutItemRequest putRequest(Fruit fruit) {
        Map<String, AttributeValue> item = new HashMap<>();
        item.put(FRUIT_NAME_COL, AttributeValue.builder().s(fruit
            .getName()).build());
        item.put(FRUIT_DESC_COL, AttributeValue.builder().s(fruit
            .getDescription()).build());

        return PutItemRequest.builder()
            .tableName(getTableName())
            .item(item)
            .build();
    }

    protected GetItemRequest getRequest(String name) {
```

```

        Map<String, AttributeValue> key = new HashMap<>();
        key.put(FRUIT_NAME_COL, AttributeValue.builder().s(name)
            .build());

        return GetItemRequest.builder()
            .tableName(getTableName())
            .key(key)
            .attributesToGet(FRUIT_NAME_COL, FRUIT_DESC_COL)
            .build();
    }
}

```

Then, create a `org.acme.dynamodb.FruitSyncService` that will be the business layer of our application and stores/loads the fruits from DynamoDB using the synchronous client.

```

package org.acme.dynamodb;

import java.util.List;
import java.util.stream.Collectors;

import javax.enterprise.context.ApplicationScoped;
import javax.inject.Inject;

import software.amazon.awssdk.services.dynamodb.DynamoDbClient;

@ApplicationScoped
public class FruitSyncService extends AbstractService {

    @Inject
    DynamoDbClient dynamoDB;

    public List<Fruit> findAll() {
        return dynamoDB.scanPaginator(scanRequest()).items().
            stream()
                .map(Fruit::from)
                .collect(Collectors.toList());
    }

    public List<Fruit> add(Fruit fruit) {
        dynamoDB.putItem(putRequest(fruit));
        return findAll();
    }

    public Fruit get(String name) {
        return Fruit.from(dynamoDB.getItem(getRequest(name)).item(
        ));
    }
}

```

Now, edit the `org.acme.dynamodb.FruitResource` class as follows:

```
package org.acme.dynamodb;

import java.util.List;

import javax.inject.Inject;
import javax.ws.rs.Consumes;
import javax.ws.rs.GET;
import javax.ws.rs.POST;
import javax.ws.rs.Path;
import javax.ws.rs.PathParam;
import javax.ws.rs.Produces;
import javax.ws.rs.core.MediaType;

@Path("/fruits")
@Produces(MediaType.APPLICATION_JSON)
@Consumes(MediaType.APPLICATION_JSON)
public class FruitResource {

    @Inject
    FruitSyncService service;

    @GET
    public List<Fruit> getAll() {
        return service.findAll();
    }

    @GET
    @Path("/{name}")
    public Fruit getSingle(@PathParam("name") String name) {
        return service.get(name);
    }

    @POST
    public List<Fruit> add(Fruit fruit) {
        service.add(fruit);
        return getAll();
    }
}
```

The implementation is pretty straightforward and you just need to define your endpoints using the JAX-RS annotations and use the `FruitSyncService` to list/add new fruits.

## Configuring DynamoDB clients

Both DynamoDB clients (sync and async) are configurable via the `application.properties` file that can be provided in the `src/main/resources` directory. Additionally, you need to added to the

classpath a proper implementation of the sync client. By default the extension uses URL connection HTTP client, so you need to add a URL connection client dependency to the `pom.xml` file:

```
<dependency>
  <groupId>software.amazon.awssdk</groupId>
  <artifactId>url-connection-client</artifactId>
</dependency>
```

If you want to use Apache HTTP client instead, configure it as follows:

```
quarkus.dynamodb.sync-client.type=apache
```

And add following dependency to the application `pom.xml`:

```
<dependency>
  <groupId>software.amazon.awssdk</groupId>
  <artifactId>apache-client</artifactId>
</dependency>
```

If you're going to use a local DynamoDB instance, configure it as follows:

```
quarkus.dynamodb.endpoint-override=http://localhost:8000

quarkus.dynamodb.aws.region=eu-central-1
quarkus.dynamodb.aws.credentials.type=static
quarkus.dynamodb.aws.credentials.static-provider.access-key-id=test-key
quarkus.dynamodb.aws.credentials.static-provider.secret-access-key=test-secret
```

- `quarkus.dynamodb.aws.region` - It's required by the client, but since you're using a local DynamoDB instance you can pick any valid AWS region.
- `quarkus.dynamodb.aws.credentials.type` - Set `static` credentials provider with any values for `access-key-id` and `secret-access-key`
- `quarkus.dynamodb.endpoint-override` - Override the DynamoDB client to use a local instance instead of an AWS service

If you want to work with an AWS account, you'd need to set it with:

```
quarkus.dynamodb.aws.region=<YOUR_REGION>
quarkus.dynamodb.aws.credentials.type=default
```



- `quarkus.dynamodb.aws.region` you should set it to the region where you provisioned the DynamoDB table,
- `quarkus.dynamodb.aws.credentials.type` - use the `default` credentials provider chain that looks for credentials in this order:
  - Java System Properties - `aws.accessKeyId` and `aws.secretKey`
  - Environment Variables - `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`
  - Credential profiles file at the default location (`~/.aws/credentials`) shared by all AWS SDKs and the AWS CLI
  - Credentials delivered through the Amazon EC2 container service if the `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` environment variable is set and the security manager has permission to access the variable,
  - Instance profile credentials delivered through the Amazon EC2 metadata service

## Next steps

### Packaging

Packaging your application is as simple as `./mvnw clean package`. It can be run with `java -jar target/amazon-dynamodb-quickstart-1.0-SNAPSHOT-runner.jar`.

With GraalVM installed, you can also create a native executable binary: `./mvnw clean package -Dnative`. Depending on your system, that will take some time.

### Going asynchronous

Thanks to the AWS SDK v2.x used by the Quarkus extension, you can use the asynchronous programming model out of the box.

Create a `org.acme.dynamodb.FruitAsyncService` that will be similar to our `FruitSyncService` but using an asynchronous programming model.

```

package org.acme.dynamodb;

import java.util.List;
import java.util.concurrent.CompletableFuture;
import java.util.stream.Collectors;

import javax.enterprise.context.ApplicationScoped;
import javax.inject.Inject;

import io.smallrye.mutiny.Uni;
import software.amazon.awssdk.services.dynamodb.DynamoDbAsyncClient;

@ApplicationScoped
public class FruitAsyncService extends AbstractService {

    @Inject
    DynamoDbAsyncClient dynamoDB;

    public Uni<List<Fruit>> findAll() {
        return Uni.createFrom().completionStage(() -> dynamoDB.
scan(scanRequest()))
            .onItem().apply(res -> res.items().stream().map
(Fruit::from).collect(Collectors.toList()));
    }

    public Uni<List<Fruit>> add(Fruit fruit) {
        return Uni.createFrom().completionStage(() -> dynamoDB
.putItem(putRequest(fruit)))
            .onItem().ignore().andSwitchTo(this::findAll);
    }

    public Uni<Fruit> get(String name) {
        return Uni.createFrom().completionStage(() -> dynamoDB
.getItem(getRequest(name)))
            .onItem().apply(resp -> Fruit.from(resp.item()));
    }
}

```

In the previous code, we create **Uni** instances from the **CompletionStage** objects returned by the asynchronous DynamoDB client, and then transform the emitted item.

Then, create an asynchronous REST resource that consumes this async service:

```

package org.acme.dynamodb;

import io.smallrye.mutiny.Uni;

import javax.inject.Inject;
import javax.ws.rs.*;
import javax.ws.rs.core.MediaType;
import java.util.List;

@Path("/async-fruits")
@Produces(MediaType.APPLICATION_JSON)
@Consumes(MediaType.APPLICATION_JSON)
public class FruitAsyncResource {

    @Inject
    FruitAsyncService service;

    @GET
    public Uni<List<Fruit>> getAll() {
        return service.findAll();
    }

    @GET
    @Path("/{name}")
    public Uni<Fruit> getSingle(@PathParam("name") String name) {
        return service.get(name);
    }

    @POST
    public Uni<List<Fruit>> add(Fruit fruit) {
        return service.add(fruit)
            .onItem().ignore().andSwitchTo(this::getAll);
    }
}

```


And add Netty HTTP client dependency to the `pom.xml`:



```

<dependency>
  <groupId>software.amazon.awssdk</groupId>
  <artifactId>netty-nio-client</artifactId>
</dependency>

```

## Configuration Reference

 Configuration property fixed at build time - All other configuration properties are overridable at runtime

Configuration property	Type	Default
 <code>quarkus.dynamodb.interceptors</code>  List of execution interceptors that will have access to read and modify the request and response objects as they are processed by the AWS SDK. The list should consists of class names which implements <code>software.amazon.awssdk.core.interceptor.ExecutionInterceptor</code> interface.	list of class name	
 <code>quarkus.dynamodb.sync-client.type</code>  Type of the sync HTTP client implementation	url, apache	url
<code>quarkus.dynamodb.enable-endpoint-discovery</code>  Enable DynamoDB service endpoint discovery.	boolean	false
AWS SDK client configurations	Type	Default
<code>quarkus.dynamodb.endpoint-override</code>  The endpoint URI with which the SDK should communicate. If not specified, an appropriate endpoint to be used for the given service and region.	URI	
<code>quarkus.dynamodb.api-call-timeout</code>  The amount of time to allow the client to complete the execution of an API call. This timeout covers the entire client execution except for marshalling. This includes request handler execution, all HTTP requests including retries, unmarshalling, etc. This value should always be positive, if present.	Duration ?	
<code>quarkus.dynamodb.api-call-attempt-timeout</code>  The amount of time to wait for the HTTP request to complete before giving up and timing out. This value should always be positive, if present.	Duration ?	
AWS services configurations	Type	Default

<p><code>quarkus.dynamodb.aws.region</code></p> <p>An Amazon Web Services region that hosts the given service.</p> <p>It overrides region provider chain with static value of region with which the service client should communicate.</p> <p>If not set, region is retrieved via the default providers chain in the following order:</p> <ul style="list-style-type: none"> <li>• <code>aws.region</code> system property</li> <li>• <code>region</code> property from the profile file</li> <li>• Instance profile file</li> </ul> <p>See <code>software.amazon.awssdk.regions.Region</code> for available regions.</p>	Region	
---	--------	--

`quarkus.dynamodb.aws.credentials.type`

Configure the credentials provider that should be used to authenticate with AWS.

Available values:

- **default** - the provider will attempt to identify the credentials automatically using the following checks:
  - Java System Properties - `aws.accessKeyId` and `aws.secretKey`
  - Environment Variables - `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`
  - Credential profiles file at the default location (`~/.aws/credentials`) shared by all AWS SDKs and the AWS CLI
  - Credentials delivered through the Amazon EC2 container service if `AWS_CONTAINER_CREDENTIALS_RELATIVE_URI` environment variable is set and security manager has permission to access the variable.
  - Instance profile credentials delivered through the Amazon EC2 metadata service
- **static** - the provider that uses the access key and secret access key specified in the `static-provider` section of the config.
- **system-property** - it loads credentials from the `aws.accessKeyId`, `aws.secretAccessKey` and `aws.sessionToken` system properties.
- **env-variable** - it loads credentials from the `AWS_ACCESS_KEY_ID`, `AWS_SECRET_ACCESS_KEY` and `AWS_SESSION_TOKEN` environment variables.
- **profile** - credentials are based on AWS configuration profiles. This loads credentials from a [profile file](#), allowing you to share multiple sets of AWS security credentials between different tools like the AWS SDK for Java and the AWS CLI.
- **container** - It loads credentials from a local metadata service. Containers currently supported by the AWS SDK are **Amazon Elastic Container Service (ECS)** and **AWS Greengrass**
- **instance-profile** - It loads credentials from the Amazon EC2 Instance Metadata Service.
- **process** - Credentials are loaded from an external process. This is used to support the `credential_process` setting in the profile credentials file. See [Sourcing Credentials From External Processes](#) for more information.
- **anonymous** - It always returns anonymous AWS credentials. Anonymous AWS credentials result in un-authenticated requests and will fail unless the resource or API's policy has been configured to specifically allow anonymous access.

default,  
static,  
system-  
prop-  
erty,  
env-  
variab-  
le,  
profil-  
e,  
contai-  
ner,  
instan-  
ce-  
profil-  
e,  
proces-  
s,  
anonym-  
ous

default

Default credentials provider configuration	Type	Default
<code>quarkus.dynamodb.aws.credentials.default-provider.async-credential-update-enabled</code>  Whether this provider should fetch credentials asynchronously in the background. If this is <code>true</code> , threads are less likely to block, but additional resources are used to maintain the provider.	boolean	<code>false</code>
<code>quarkus.dynamodb.aws.credentials.default-provider.reuse-last-provider-enabled</code>  Whether the provider should reuse the last successful credentials provider in the chain. Reusing the last successful credentials provider will typically return credentials faster than searching through the chain.	boolean	<code>true</code>
Static credentials provider configuration	Type	Default
<code>quarkus.dynamodb.aws.credentials.static-provider.access-key-id</code>  AWS Access key id	string	
<code>quarkus.dynamodb.aws.credentials.static-provider.secret-access-key</code>  AWS Secret access key	string	
AWS Profile credentials provider configuration	Type	Default
<code>quarkus.dynamodb.aws.credentials.profile-provider.profile-name</code>  The name of the profile that should be used by this credentials provider. If not specified, the value in <code>AWS_PROFILE</code> environment variable or <code>aws.profile</code> system property is used and defaults to <code>default</code> name.	string	
Process credentials provider configuration	Type	Default
<code>quarkus.dynamodb.aws.credentials.process-provider.async-credential-update-enabled</code>  Whether the provider should fetch credentials asynchronously in the background. If this is true, threads are less likely to block when credentials are loaded, but additional resources are used to maintain the provider.	boolean	<code>false</code>

<code>quarkus.dynamodb.aws.credentials.process-provider.credential-refresh-threshold</code>  The amount of time between when the credentials expire and when the credentials should start to be refreshed. This allows the credentials to be refreshed <b>before</b> they are reported to expire.	Duration ?	15S
<code>quarkus.dynamodb.aws.credentials.process-provider.process-output-limit</code>  The maximum size of the output that can be returned by the external process before an exception is raised.	Memory Size ?	1024
<code>quarkus.dynamodb.aws.credentials.process-provider.command</code>  The command that should be executed to retrieve credentials.	string	
<b>Sync HTTP transport configurations</b>	<b>Type</b>	<b>Default</b>
<code>quarkus.dynamodb.sync-client.connection-timeout</code>  The maximum amount of time to establish a connection before timing out.	Duration ?	2S
<code>quarkus.dynamodb.sync-client.socket-timeout</code>  The amount of time to wait for data to be transferred over an established, open connection before the connection is timed out.	Duration ?	30S
<b>Apache HTTP client specific configurations</b>	<b>Type</b>	<b>Default</b>
<code>quarkus.dynamodb.sync-client.apache.connection-acquisition-timeout</code>  The amount of time to wait when acquiring a connection from the pool before giving up and timing out.	Duration ?	10S
<code>quarkus.dynamodb.sync-client.apache.connection-max-idle-time</code>  The maximum amount of time that a connection should be allowed to remain open while idle.	Duration ?	60S
<code>quarkus.dynamodb.sync-client.apache.connection-time-to-live</code>  The maximum amount of time that a connection should be allowed to remain open, regardless of usage frequency.	Duration ?	



<code>quarkus.dynamodb.sync-client.apache.max-connections</code>		
The maximum number of connections allowed in the connection pool. Each built HTTP client has its own private connection pool.	int	50
<code>quarkus.dynamodb.sync-client.apache.expect-continue-enabled</code>		
Whether the client should send an HTTP expect-continue handshake before each request.	boolean	true
<code>quarkus.dynamodb.sync-client.apache.use-idle-connection-reaper</code>		
Whether the idle connections in the connection pool should be closed asynchronously. When enabled, connections left idling for longer than <code>quarkus..sync-client.connection-max-idle-time</code> will be closed. This will not close connections currently in use.	boolean	true
<code>quarkus.dynamodb.sync-client.apache.proxy.enabled</code>		
Enable HTTP proxy	boolean	false
<code>quarkus.dynamodb.sync-client.apache.proxy.endpoint</code>		
The endpoint of the proxy server that the SDK should connect through. Currently, the endpoint is limited to a host and port. Any other URI components will result in an exception being raised.	URI	
<code>quarkus.dynamodb.sync-client.apache.proxy.username</code>		
The username to use when connecting through a proxy.	string	
<code>quarkus.dynamodb.sync-client.apache.proxy.password</code>		
The password to use when connecting through a proxy.	string	
<code>quarkus.dynamodb.sync-client.apache.proxy.ntlm-domain</code>		
For NTLM proxies - the Windows domain name to use when authenticating with the proxy.	string	
<code>quarkus.dynamodb.sync-client.apache.proxy.ntlm-workstation</code>		
For NTLM proxies - the Windows workstation name to use when authenticating with the proxy.	string	

<code>quarkus.dynamodb.sync-client.apache.proxy.preemptive-basic-authentication-enabled</code>  Whether to attempt to authenticate preemptively against the proxy server using basic authentication.	boolean	
<code>quarkus.dynamodb.sync-client.apache.proxy.non-proxy-hosts</code>  The hosts that the client is allowed to access without going through the proxy.	list of string	
<code>quarkus.dynamodb.sync-client.apache.tls-managers-provider.type</code>  TLS managers provider type.  Available providers: <ul style="list-style-type: none"> <li>• <b>none</b> - Use this provider if you don't want the client to present any certificates to the remote TLS host.</li> <li>• <b>system-property</b> - Provider checks the standard <code>javax.net.ssl.keyStore</code>, <code>javax.net.ssl.keyStorePassword</code>, and <code>javax.net.ssl.keyStoreType</code> properties defined by the <a href="#">JSSE</a>.</li> <li>• <b>file-store</b> - Provider that loads a the key store from a file.</li> </ul>	none, system-property, file-store	system-property
<code>quarkus.dynamodb.sync-client.apache.tls-managers-provider.file-store.path</code>  Path to the key store.	path	
<code>quarkus.dynamodb.sync-client.apache.tls-managers-provider.file-store.type</code>  Key store type. See the KeyStore section in the <a href="#">Java Cryptography Architecture Standard Algorithm Name Documentation</a> for information about standard keystore types.	string	
<code>quarkus.dynamodb.sync-client.apache.tls-managers-provider.file-store.password</code>  Key store password	string	
<b>Netty HTTP transport configurations</b>	<b>Type</b>	<b>Default</b>
<code>quarkus.dynamodb.async-client.max-concurrency</code>  The maximum number of allowed concurrent requests. For HTTP/1.1 this is the same as max connections. For HTTP/2 the number of connections that will be used depends on the max streams allowed per connection.	int	50

<code>quarkus.dynamodb.async-client.max-pending-connection-acquires</code>		
The maximum number of pending acquires allowed. Once this exceeds, acquire tries will be failed.	int	10000
<code>quarkus.dynamodb.async-client.read-timeout</code>		
The amount of time to wait for a read on a socket before an exception is thrown. Specify 0 to disable.	Duration ?	30S
<code>quarkus.dynamodb.async-client.write-timeout</code>		
The amount of time to wait for a write on a socket before an exception is thrown. Specify 0 to disable.	Duration ?	30S
<code>quarkus.dynamodb.async-client.connection-timeout</code>		
The amount of time to wait when initially establishing a connection before giving up and timing out.	Duration ?	10S
<code>quarkus.dynamodb.async-client.connection-acquisition-timeout</code>		
The amount of time to wait when acquiring a connection from the pool before giving up and timing out.	Duration ?	2S
<code>quarkus.dynamodb.async-client.connection-time-to-live</code>		
The maximum amount of time that a connection should be allowed to remain open, regardless of usage frequency.	Duration ?	
<code>quarkus.dynamodb.async-client.connection-max-idle-time</code>		
The maximum amount of time that a connection should be allowed to remain open while idle. Currently has no effect if <code>quarkus..async-client.use-idle-connection-reaper</code> is false.	Duration ?	60S
<code>quarkus.dynamodb.async-client.use-idle-connection-reaper</code>		
Whether the idle connections in the connection pool should be closed. When enabled, connections left idling for longer than <code>quarkus..async-client.connection-max-idle-time</code> will be closed. This will not close connections currently in use.	boolean	true
<code>quarkus.dynamodb.async-client.protocol</code>		
The HTTP protocol to use.	http1-1, http2	http1-1

<code>quarkus.dynamodb.async-client.ssl-provider</code>		
The SSL Provider to be used in the Netty client. Default is <b>OPENSSL</b> if available, <b>JDK</b> otherwise.	jdk, openssl, openssl-refcnt	
<code>quarkus.dynamodb.async-client.http2.max-streams</code>		
The maximum number of concurrent streams for an HTTP/2 connection. This setting is only respected when the HTTP/2 protocol is used.	long	4294967295
<code>quarkus.dynamodb.async-client.http2.initial-window-size</code>		
The initial window size for an HTTP/2 stream. This setting is only respected when the HTTP/2 protocol is used.	int	1048576
<code>quarkus.dynamodb.async-client.proxy.enabled</code>		
Enable HTTP proxy.	boolean	false
<code>quarkus.dynamodb.async-client.proxy.endpoint</code>		
The endpoint of the proxy server that the SDK should connect through. Currently, the endpoint is limited to a host and port. Any other URI components will result in an exception being raised.	URI	
<code>quarkus.dynamodb.async-client.proxy.non-proxy-hosts</code>		
The hosts that the client is allowed to access without going through the proxy.	list of string	
<code>quarkus.dynamodb.async-client.tls-managers-provider.type</code>		
<p>TLS managers provider type.</p> <p>Available providers:</p> <ul style="list-style-type: none"> <li><b>none</b> - Use this provider if you don't want the client to present any certificates to the remote TLS host.</li> <li><b>system-property</b> - Provider checks the standard <code>javax.net.ssl.keyStore</code>, <code>javax.net.ssl.keyStorePassword</code>, and <code>javax.net.ssl.keyStoreType</code> properties defined by the <b>JSSE</b>.</li> <li><b>file-store</b> - Provider that loads a the key store from a file.</li> </ul>	none, system-property, file-store	system-property
<code>quarkus.dynamodb.async-client.tls-managers-provider.file-store.path</code>		
Path to the key store.	path	

<code>quarkus.dynamodb.async-client.tls-managers-provider.file-store.type</code>		
Key store type. See the KeyStore section in the <a href="#">Java Cryptography Architecture Standard Algorithm Name Documentation</a> for information about standard keystore types.	string	
<code>quarkus.dynamodb.async-client.tls-managers-provider.file-store.password</code>		
Key store password	string	
<code>quarkus.dynamodb.async-client.event-loop.override</code>		
Enable the custom configuration of the Netty event loop group.	boolean	<code>false</code>
<code>quarkus.dynamodb.async-client.event-loop.number-of-threads</code>		
Number of threads to use for the event loop group. If not set, the default Netty thread count is used (which is double the number of available processors unless the <code>io.netty.eventLoopThreads</code> system property is set.	int	
<code>quarkus.dynamodb.async-client.event-loop.thread-name-prefix</code>		
The thread name prefix for threads created by this thread factory used by event loop group. The prefix will be appended with a number unique to the thread factory and a number unique to the thread. If not specified it defaults to <code>aws-java-sdk-NettyEventLoop</code>	string	



#### About the Duration format

The format for durations uses the standard `java.time.Duration` format. You can learn more about it in the [Duration#parse\(\) javadoc](#).

You can also provide duration values starting with a number. In this case, if the value consists only of a number, the converter treats the value as seconds. Otherwise, `PT` is implicitly prepended to the value to obtain a standard `java.time.Duration` format.



#### About the MemorySize format

A size configuration option recognises string in this format (shown as a regular expression): `[0-9]+[KkMmGgTtPpEeZzYy]?.` If no suffix is given, assume bytes.