

Tutorial 6 - Introduction to Lambda III: Serverless Databases

Disclaimer: Subject to updates as corrections are found

Version 0.11

Scoring: 20 pts maximum

The purpose of this tutorial is to introduce the use of relational databases from AWS Lambda. This tutorial will demonstrate the use of SQLite, an in-memory or file-based database that runs inside a Lambda function to provide a “temporary” relational database that lives for the lifetime of the Lambda container. Secondly, the tutorial demonstrates the use of the Amazon Relational Database Service (RDS) to create a persistent relational database using Serverless Aurora MySQL 5.6 for data storage and query support for Lambda functions.

Goals of this tutorial include:

1. Introduce the SQLite database using the command line “sqlite3” tool.
2. Deploy a Lambda Function that uses a file-based SQLite3 database in the “/tmp” directory of the Lambda container that persists between client function invocations
3. Compare the difference between using file-based and in-memory SQLite DBs on Lambda.
4. Create an Amazon RDS Aurora MySQL Serverless database cluster
5. Launch an EC2 instance and install the mysql command line client to interface with the Aurora serverless database.
6. Deploy an AWS Lambda function that uses the MySQL Serverless database.

1. Using the SQLite Command Line

To begin, create a directory called “saaf_sqlite”.

Then clone the git repository under the new directory:

```
git clone https://github.com/wlloyduw/saaf\_sqlite.git
```

If using Windows or Mac, download the “Precompiled binaries” as a zip file from:

<https://www.sqlite.org/download.html>

On Windows/Mac, unzip the zip file, and then run the **sqlite3** program

On Linux, navigate to **saaf_sqlite/sqlite-tools-linux-x86-3250200/** from the github repository and launch the sqlite3 commandline:

```
./sqlite3
```

Check out the available commands using “.help”

```
SQLite version 3.25.2 2018-09-25 19:08:10
Enter ".help" for usage hints.
Connected to a transient in-memory database.
Use ".open FILENAME" to reopen on a persistent database.
sqlite> .help
```

Start by saving a new database file, and then exit the tool:

```
sqlite> .save new.db
sqlite> .quit
```

Then, check the size of an empty sqlite db file:

```
$ ls -l new.db
total 3848
-rw-r--r-- 1 wlloyd wlloyd 4096 Nov  1 19:05 new.db
```

It is only 4096 bytes, very small!

Next, work with data in the database:

```
$ ./sqlite3 new.db
SQLite version 3.25.2 2018-09-25 19:08:10
Enter ".help" for usage hints.
sqlite> .databases
main: /home/wlloyd/git/saaf_sqlite/sqlite-tools-linux-x86-3250200/new.db
sqlite> .tables
```

There are initially no tables.

Create a table and insert some data:

```
sqlite> create table newtable (name text, city text, state text);
sqlite> .tables
newtable
sqlite> insert into newtable values('Susan Smith','Tacoma','Washington');
sqlite> insert into newtable values('Bill Gates','Redmond','Washington');
sqlite> select * from newtable;
Susan Smith|Tacoma|Washington
Bill Gates|Redmond|Washington
```

Now check how the database file has grown after adding a table and a few rows:

```
sqlite> .quit
$ ls -l new.db
```

Question 1. After creating the table ‘newtable’ and loading data to sqlite, what is the size of the new.db database file?

The sqlite3 command line tool can be used to perform common **C**reate **R**ead **U**ppdate and **D**eleate queries on a sqlite database. This allows the database to be preloaded with data and bundled with a Lambda function for deployment to the cloud as needed.

If you're unfamiliar with SQL, consider completing an online tutorial:

SQLite Tutorial:

<http://www.sqlitetutorial.net/>

Follow the steps for "Getting started with SQLite" (1, 2, and 3), and then complete the Basic SQLite tutorial to review performing different types of queries using the sample Chinook database with 11 tables downloaded from step 3.

2. Combining SQLite with AWS Lambda

SQLite can be leveraged directly from programming languages such as Java, Python, and Node.JS. SQLite provides an alternative to basic CSV and text file storage with a SQL-compatible query-able file format. SQLite does not replace a full-fledged enterprise relational database management system (dbms) in terms of scalability, etc. But given the small footprint of SQLite, it provides an excellent database alternative for serverless environments and Internet of Things devices.

Next, explore the saaf_sqlite project in Netbeans or another Java IDE.

"saaf_sqlite" provides a Java-based Lambda "Hello" function based on SAAF from Tutorial #4. Look at the code inside: **saaf_sqlite/SAAF/java_template/src/main/java/lambda/HelloSqlite.java**.

```
setCurrentDirectory("/tmp");
try
{
    // Connection string an in-memory SQLite DB
    //Connection con = DriverManager.getConnection("jdbc:sqlite:");

    // Connection string for a file-based SQLite DB
    Connection con = DriverManager.getConnection("jdbc:sqlite:mytest.db");
```

The first line of code (LOC) calls a helper function to set the working directory to "/tmp" inside the Lambda function.

"/tmp" provides a read/write 512MB filesystem on Lambda.

As a security precaution, code deployed to Lambda has only limited permission to write to the filesystem, as /tmp is enabled for read/write.

SQLite can work with databases entirely in memory, or on disk. The first database connection string is commented out, but could be used if wanting to work with a database only in memory.

The advantage of creating the database on disk is that data persists beyond the runtime of the Java code. On Lambda, this means as long as the original runtime container is preserved, the data is preserved. If containers are kept WARM, they can last up to 6 to 8 hours. After 6-8 hours, it will be necessary to save any SQLite databases to S3 to persist the data for longer.

Perform a clean build of the `saaf_sqlite` project to create a jar file.

Following instructions from tutorial #4, deploy a new lambda function called “helloSqlite”. Choose one method for invoking “helloSqlite”.

If wanting to use a HTTP/REST URL, configure the API Gateway to provide a URL for access via Curl. Otherwise use the “helloSqlite” Lambda function name and the AWS Lambda CLI.

Under your new project, modify the `callservice.sh` script to invoke your newly deployed SQLite Lambda function: `saaf_sqlite/SAAF/java_template/test/callservice.sh`

Then run the script. Below the API Gateway invocation code in BASH has been commented out using a “#” in front of each line.

```
$ ./callservice.sh
Invoking Lambda function using AWS CLI
real 0m11.985s
user 0m0.288s
sys 0m0.064s

AWS CLI RESULT:
{
  .... // some attributes removed from brevity...
  "uuid": "8c321d18-d16e-4cd8-acac-cbc8d65fe138",
  "error": "",
  "vmuptime": 1541129227,
  "newcontainer": 1,
  "value": "Hello Fred Smith",
  "names": [
    "Fred Smith"
  ]
}
```

Using a file, each time the service is called and the same runtime container is used, a name is appended to the temporary file-based SQLite database. We see the “names” array in the JSON grow with each subsequent call.

Try running the `./callservice.sh` script now several times (3x-5x) to watch the names array grow.

Now, try out what happens when two clients call the Lambda function at the same time.

Inspect the simple `calltwice.sh` script:

```
cat calltwice.sh
```

Now, try running calltwice.sh:

```
./calltwice.sh
```

Invoking a Lambda with two clients in parallel forces Lambda to create additional server infrastructure.

Question 2. When the second client calls the helloSqlite Lambda function, how is the data different in the second container environment compared to the initial/first container?

Now, try out a memory-only SQLite database. Modify your Lambda code to swap out the type of database. Comment out the file-based database in favor of memory only:

```
setCurrentDirectory("/tmp");
try
{
    // Connection string an in-memory SQLite DB
    Connection con = DriverManager.getConnection("jdbc:sqlite:");

    // Connection string for a file-based SQLite DB
    // Connection con = DriverManager.getConnection("jdbc:sqlite:mytest.db");
}
```

Build a new JAR file, and redeploy it to Lambda for the helloSqlite Lambda function.

Using callservice.sh, try calling the Lambda several times in succession.

Question 3. For Lambda calls that execute in the same runtime container identified by the UUID returned in JSON, does the data persist between client Lambda calls with an in-memory DB? (YES or NO)

Next, let's modify the code for helloSqlite to add a static int counter that tracks the total number of calls to the container.

Define a static int at the start of public class HelloSqlite:

```
public class HelloSqlite implements RequestHandler<Request, HashMap<String, Object>>
{
    static String CONTAINER_ID = "/tmp/container-id";
    static Charset CHARSET = Charset.forName("US-ASCII");

    static int uses = 0;
```

Then modify the definition of String hello near the bottom of the Lambda function to report the uses count:

```
// *****
// Set hello message here
// *****
uses = uses + 1;
String hello = "Hello " + request.getName() + " calls=" + uses;
```

Build a new JAR file, and redeploy it to Lambda for the helloSqlite Lambda function.

Using callservice.sh, try calling the Lambda with the static uses counter several times in succession:

```
./callservice.sh
./callservice.sh
./callservice.sh
```

Question 4. Does the value of the static int persist for Lambda calls that execute in the same runtime container identified by the UUID returned in JSON? (YES or NO)

Now, try running with calltwice.sh.

Question 5. How is the value of the static int different across different runtime containers identified by the UUID returned in JSON?

Next, inspect the SQL code for the helloSqlite Lambda function:

```
// Detect if the table 'mytable' exists in the database
PreparedStatement ps = con.prepareStatement("SELECT name FROM sqlite_master WHERE
type='table' AND name='mytable'");
ResultSet rs = ps.executeQuery();
if (!rs.next())
{
    // 'mytable' does not exist, and should be created
    logger.log("trying to create table 'mytable'");
    ps = con.prepareStatement("CREATE TABLE mytable ( name text, col2 text, col3
text);");
    ps.execute();
}
rs.close();

// Insert row into mytable
ps = con.prepareStatement("insert into mytable values('" + request.getName() +
"', 'b', 'c');");
ps.execute();

// Query mytable to obtain full resultset
ps = con.prepareStatement("select * from mytable;");
rs = ps.executeQuery();
```

The approach of our helloSqlite Lambda is to create a new file (or memory) database each time.

Question 6. Before inserting rows into 'mytable', what has to be done and why in the Java code above?

3. Optional Exercise: Persisting SQLite database files to S3

Leveraging concepts from tutorial #5, modify the file-based version of helloSqlite to always save the database file in /tmp to S3 at the end of the function. Add a key to the request.java to allow the user to specify a database filename.

At the beginning of the function handler, using the user provided database filename from the request, check if the specified file exists in /tmp. If it does not exist, then download the file from S3. Then change the SQLite connection string to open the user provided database name. This way a user could request a specific database from S3 for their Lambda function call.

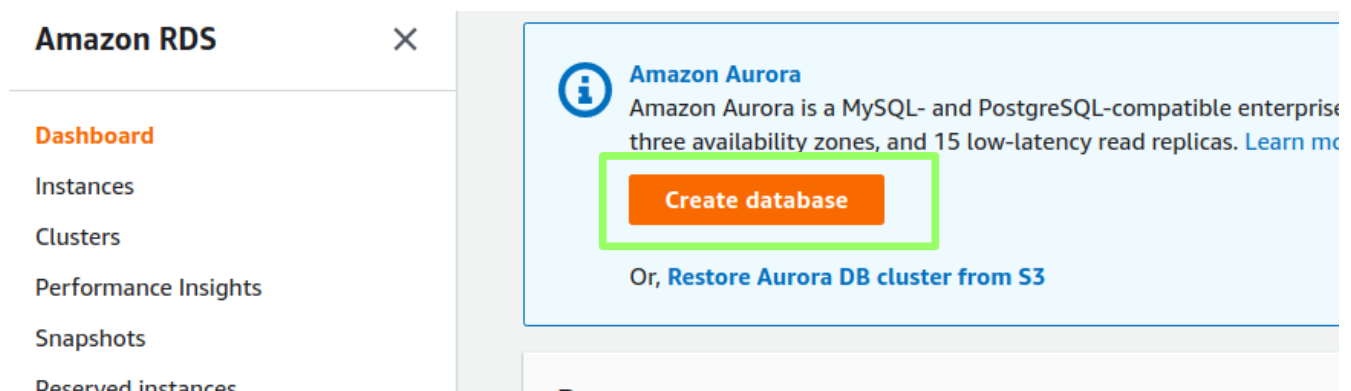
4. Create a AWS RDS Aurora MySQL Serverless Database

The AWS Relational Database Service now offers a serverless MySQL database. The advantage of the serverless database is that it automatically goes idle when not actively used. When idle, the only charges are for storage. Serverless Aurora supports automatic horizontal scaling of database servers up to ~15 nodes. The nodes provide read replicas, where a master is used for database writes. (Fall 2018: Multiple R/W masters, called "multi-master" is presently in beta) Vertical scaling allows the CPU and memory resources for the master R/W node to scale from 2 vCPUs and 4GB RAM presumably up to 64 vCPUs and 488 GB RAM, aligning with r4 instance types: <https://www.ec2instances.info/?filter=r4> Aurora additionally provides automatic database backups and replicas.

To get started, lets create a serverless database!

**** This portion of the tutorial requires AWS credits to complete. ****
Aurora serverless does not run in the FREE tier.

Go to "RDS", and click "Create database" to launch the wizard:




First specify the “Amazon Aurora” engine and the following choose options:

Engine options

Engine type [Info](#)

☒ Amazon Aurora



Edition

☒ Amazon Aurora with MySQL compatibility

☐ Amazon Aurora with PostgreSQL compatibility

Version [Info](#)

Aurora (MySQL)-5.6.10a ▼

Database features are supported with specific engine versions. [Info](#)

Database Location

☒ Regional

You provision your Aurora database in a single AWS Region.

Then, below, be sure to select the “Serverless”:

Database features

☐ **One writer and multiple readers**
Supports multiple reader instances connected to the same storage volume as a single writer instance. This is a good general-purpose option for most workloads.

☐ **One writer and multiple readers - Parallel query**
Improves the performance of analytic queries by pushing processing down to the Aurora storage layer. This is a good option for hybrid transactional/analytic workloads.

☐ **Multiple writers**
Supports multiple writer instances connected to the same storage volume. This is a good option for when continuous writer availability is required.

☒ **Serverless**
You specify the minimum and maximum amount of resources needed, and Aurora scales the capacity based on database load. This is a good option for intermittent or unpredictable workloads.

Aurora MySQL-5.6.10a is the only “serverless” MySQL database version supported.

Next, specify settings as below. Set and remember the password:

DB cluster identifier

Type a name for your DB cluster. The name must be unique across all DB clusters owned by your AWS account in the current AWS Region.

tcss562

The DB cluster identifier is a case-sensitive, but is stored as all lowercase(as in "mydbcluster"). Constraints: 1 to 60 alphanumeric characters or hyphens (1 to 15 for SQL Server). First character must be a letter. Can't contain two consecutive hyphens. Can't end with a hyphen.

Master username [Info](#)

Specify an alphanumeric string that defines the login ID for the master user.

tcss562

Master Username must start with a letter. Must contain 1 to 16 alphanumeric characters.

Master password [Info](#)

.....

Master Password must be at least eight characters long, as in "mypassword". Can be any printable ASCII character except "/", "", or "@".

Confirm password [Info](#)

.....

Next, provide scaling specifications. To save cost, specify the minimum possible scaling. Note that running the server for 1 hour with 2 Aurora Capacity Units and 4GB of memory costs 12¢. Aurora Serverless costs 6¢/per ACU/per hour billed to the nearest second with a 5-minute minimum each time the database is activated. To limit the maximum per hour charge, set the maximum Aurora capacity units to the minimum setting **1**. This effectively disables auto-scaling.

Open the "Additional scaling configuration"

Capacity settings

This billing estimate is based on published prices. [Learn more](#)

Minimum Aurora capacity unit [Info](#)

1

2GB RAM

Maximum Aurora capacity unit [Info](#)

1

2GB RAM

▼ Additional scaling configuration

☐ Force scaling the capacity to the specified values when the timeout is reached [Info](#)

Enable to force capacity scaling as soon as possible. Disable to cancel the capacity changes when a timeout is reached

☒ Pause compute capacity after consecutive minutes of inactivity [Info](#)

You are only charged for database storage while the compute capacity is paused

0

hours

5

minutes

0

seconds

Max: 24 hours

ENABLE THE OPTION FOR “PAUSE COMPUTE CAPACITY AFTER CONSECUTIVE MINUTES OF INACTIVITY”... THIS IS CRUCIAL OTHERWISE THE DATABASE WILL HAVE A CONTINUOUS CHARGE...

One Aurora Capacity Unit (ACU) has approximately 2 GB of memory with corresponding CPU and networking, similar to that of Aurora Standard instances “r4” instances (see above link).


For Connectivity specify: Virtual Private Cloud: Default VPC

Then open “Additional connectivity configuration”:

Subnet group: Select an existing DB subnet group if available, otherwise select: “Create new DB Subnet Group”

VPC security groups: Choose existing, then select “default” from the dropdown list.

For the remaining settings, the defaults can be used.
The press:



The RDS Clusters list will appear.

The state of creation can be monitored. Cluster creation takes a couple minutes.

	tcss562	Aurora MySQL	5.6.10a	 creating	0 capacity unit	serverless	none
--	---------	--------------	---------	---	-----------------	------------	------


5. Launch a t2.micro EC2 VM to connect to the Aurora DB cluster

It is not possible to directly connect to the Aurora MySQL Serverless cluster. This is because the cluster lives on a Virtual Private Cloud (VPC) that does not allow inbound traffic from the internet. This provides network isolation and security. Accessing the database requires launching an EC2 instance in the same VPC as the RDS cluster and associating a Public IP address with this EC2 instance. The RDS cluster itself does not have a public IP that is accessible from the outside. Providing direct connectivity to the RDS cluster requires: (1) setting up either a NAT Gateway (4.5¢/hour), (2) configuring a VM to act as a router/gateway which requires special network configuration, or (3) installing a proxy server such as haproxy on the publicly accessible VM to proxy inbound traffic for MySQL to the RDS cluster. A database client could then connect to the VM, not RDS, and the traffic is redirected. Fortunately, if deploying AWS Lambda functions in the same VPC as the RDS cluster, no special networking appliance (e.g. NAT gateway instance or router) is required to access the database. This saves cost and complexity.

Navigate to EC2.

Click on “Launch Instance”

Select the latest version of Ubuntu:

 **Ubuntu Server 18.04 LTS (HVM), SSD Volume Type** - ami-0ac019f4cb7cb7e6

Ubuntu Server 18.04 LTS (HVM),EBS General Purpose (SSD) Volume Type. Support available from Canonical (<http://www.ubuntu.com/cloud/services>).

Free tier eligible

Root device type: ebs Virtualization type: hvm ENA Enabled: Yes

Select

64-bit (x86)

Click [Select].

Select a t2.micro, a free-tier instance. New accounts have 750 hours of free t2.micro time/month.

Click [Next: Configure Instance Details]

Provide the following settings:

Network: vpc (default) - select your default VPC -

Subnet: no preference

Auto-assign Public IP: Enabled

Click [Next: Add Storage], then

Click [Next: Add Tags], then

Click [Next: Configure Security Group], then

Choose "Select an existing security group".

And check the "default".

Then click "Review and Launch".

Choose an existing keypair from a prior tutorial if available.

Otherwise, create a new keypair.

In the EC2 console, select the VM. If you have not enabled SSH access from your network, select the new VM, and click on "default" for Security groups. Click the "Inbound" Tab, and hit the [Edit] button. Click [Add Rule] and add a "SSH" "TCP" "22" rule for "My IP". This should allow SSH access to the t2.micro instance.

Security groups default. 1

Next in the EC2 console copy the public IP:

The right-hand COPY icon makes it easy to copy the IP address.

IPv4 Public IP 34.230.46.211 

Now, using the command line, navigate to the folder where the keypair is stored, and ssh into the newly created t2.micro VM. Paste the address and SSH:

```
$ ssh -i <your key file name> ubuntu@<t2.micro IPv4 public IP>
```

Now on the Ubuntu t2.micro instance launched in the same VPC as the RDS cluster run the commands:

```
sudo apt update
sudo apt upgrade
```

Optionally, install the AWSCLI if wanting to work with AWS directly from the VM. It is not required.

```
# installing the AWSCLI is optional
apt install awscli
sudo apt install awscli
aws configure
# provide ACCESS_KEY and SECRET_KEY
# Find your credentials on your existing VM with: "cat ~/.aws/credentials"
```

Next, install the mysql client to support connecting to the new RDS cluster:

```
# Install mysql client
sudo apt install mysql-client-core-5.7
```

Now customize the following command to point at your RDS cluster.

Navigate back into RDS in the AWS management console.

On the left hand-side select "Clusters", then select "tcss562".

First, it is necessary to configure the security group to allow the t2.micro to connect to the database. Scroll down to Details, and look under "Security and network".

Click on the "Security groups" link:

Security groups
Security Group sg-0de31f001975ee99b
(active)

In the Security editor, click the [Inbound] tab.

Click [Edit].

Click [Add Rule].

Select "MYSQL/Aurora".

For a Custom address range, provide "0.0.0.0/0".

This enables any VM within the private VPC network to be able to connect to the database.

Hit [Save].

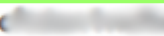
Now navigate back to RDS.

On the left hand-side select "Clusters", then select "tcss562".

Scroll down to Details, and look under "Security and network".

Copy and paste the name of the "Database endpoint".

For example, copy the Database endpoint name:

Database endpoint
tcss562.cluster-.us-east-1.rds.amazonaws.com

Now customize the mysql command to connect to your RDS cluster.

Replace <Database endpoint> and <your database password>.

```
mysql --host=<Database endpoint> --port=3306 --enable-cleartext-plugin  
--user=tcss562 --password=<your database password>
```

When first connecting to your Amazon RDS connection, if getting to this step took longer than 5-minutes from the time the cluster was initially created, then the RDS cluster has gone to sleep. In this case, calling “mysql” will take up to 30 seconds to respond. Please wait. It is working. If more than a couple minutes, stop, and troubleshoot.

After sometime, mysql should connect:

```
mysql: [Warning] Using a password on the command line interface can be insecure.  
Welcome to the MySQL monitor.  Commands end with ; or \g.  
Your MySQL connection id is 10  
Server version: 5.6.10 MySQL Community Server (GPL)  
Copyright (c) 2000, 2018, Oracle and/or its affiliates. All rights reserved.  
Oracle is a registered trademark of Oracle Corporation and/or its  
affiliates. Other names may be trademarks of their respective  
owners.  
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.  
mysql>
```

Try out the following commands.

MySQL can support multiple databases within a single server.

Display the databases on your RDS cluster:

```
show databases;
```

Now, create a new database:

```
create database TEST;
```

And check the list again:

```
show databases;
```

It is necessary to issue a “use” command for mysql to direct SQL queries to the database:

```
use TEST;
```

Next, create “mytable” to store data:

```
CREATE TABLE mytable (name VARCHAR(40), col2 VARCHAR(40), col3 VARCHAR(40));
```

Then display the list of known tables in the database:

```
show tables;
```

And describe the structure of the table:

```
describe mytable;
```

Now, try adding some data:

```
insert into mytable values ('fred','testcol2','testcol3');
```

And then check if it was inserted:

```
select * from mytable;
```

Help is available with the “help” command:

```
help
```

Exit mysql with:

```
\q
```

It may be useful to “stop” and “start” your ec2 instance that has access to the Amazon RDS cluster to support working with mysql. If no longer planning to use the ec2instance, terminate it completely. Note that “stopped” instances incur storage charges. New AWS accounts receive 30GB of disk space for 1 year for free. After 1 year, the charge is 10¢/GB/month. The Ubuntu t2.small requires 8GB of storage. The annual storage cost after the free introductory year goes to \$9.60/year for a stopped instance with 1 - 8GB EBS volume.

6. Accessing Aurora Serverless Database from AWS Lambda

Next, on your development computer, create a directory called “saaf_rds”. Then under the new directory, clone the git repository:

```
git clone https://github.com/wlloyduw/saaf\_rds\_serverless.git
```

This project, provides a Lambda function that will interact with your Amazon RDS cluster. It requires “mytable” to have been created under the “TEST” database. Optionally, it should be possible to create the database and table programmatically from Java if necessary.

Once acquiring the project files, it is necessary to create a file called “db.properties”. There is a template provided. Copy this template to be named “db.properties” and edit this file to specify how to connect to your RDS cluster:

Find and edit this file:

```
cd saaf_rds_serverless/SAAF/java_template/src/main/resources/  
cp db.properties.empty db.properties  
gedit db.properties
```

The **URL** should be specified as follows:

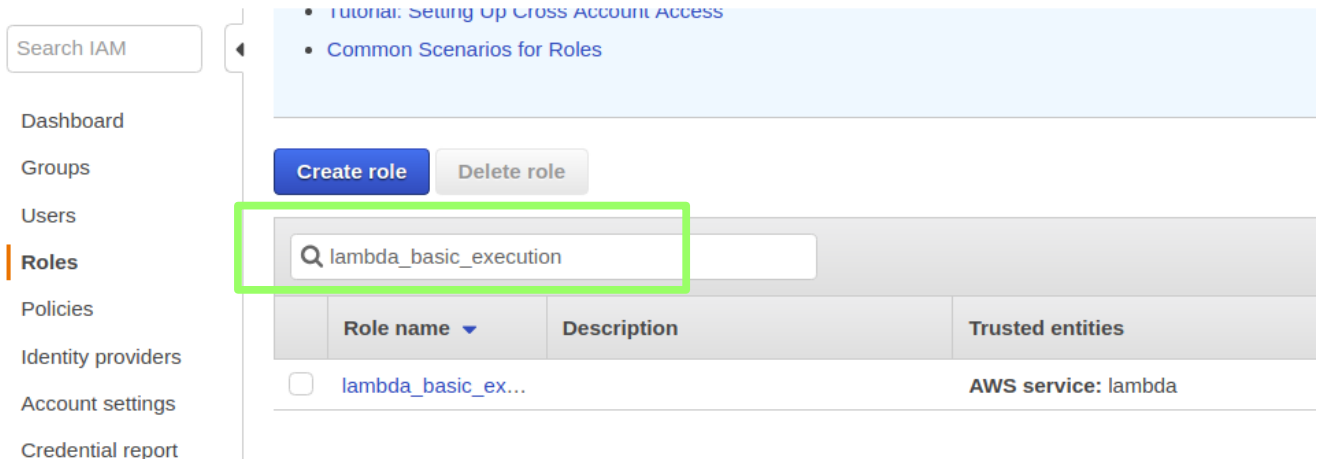
```
jdbc:mysql://<your RDS cluster name>:3306/TEST
```

Replace “<your RDS cluster name>” with the cluster name used to connect with mysql above. Be sure to add values for **password**, and **username** as well based on how your RDS cluster was initially configured

Now, create a new Lambda function.

When configuring the security role for the new Lambda function, new policies need to be added to the security role.

Using Identity Access Manager GUI, navigate to security roles:



For example, the role above is called “lambda_basic_execution”.

Find your Lambda security role, and then select it by clicking on the name.

Attach new policies:

AmazonRDSFullAccess




AWSLambdaVPCLambdaAccessExecutionRole

Maximum CLI/API session duration 1 h

Permissions Trust relationships Access Advisor

▼ Permissions policies (5 policies applied)

Attach policies

Policy name ▼
▶  AmazonRDSFullAccess
▶  AmazonS3FullAccess
▶  AWSLambdaVPCLambdaAccessExecutionRole

Now, finish creating the new Lambda function. Be sure to update the handler name:

Handler [Info](#)

lambda.HelloMySQL::handleRequest

Next, deploy this Lambda function into the same VPC and subnet as your RDS cluster. If not, there will be no connectivity between Lambda and the RDS cluster. First, discover which VPC and Subnets the RDS cluster has been created in. Note these, and configure the Lambda function to use the same VPC and Subnet.

In a separate window, navigate to RDS. On the left hand-side select "Clusters", then select "tcss562". Scroll down to Details, and look under "Security and network".

Security and network

VPC
vpc-4f8db929

Subnet group
dbsubnet_useast1e

Subnets
subnet-1a82f326
subnet-caa9bee7

Here, note the VPC, and any subnet in use.

Configure the Lambda function to match the VPC and subnets. Note the VPC/subnet names here are examples only. Select the default security group.

If using the Ohio region, expect subnets to be us-east-2a, us-east-2b, or us-east-2c.

Network

VPC [Info](#)
Choose a VPC for your function to access.

Default vpc-4f8db929 (172.31.0.0/16) ▼

Subnets*
Select the VPC subnets for Lambda to use to set up your VPC configuration. Format: "subnet-id (cidr-block) | az name-tag".

▼

subnet-1a82f326 (172.31.48.0/20) | us-east-1e ✕
subnet-caa9bee7 (172.31.64.0/20) | us-east-1d ✕

Security groups*
Choose the VPC security groups for Lambda to use to set up your VPC configuration. Format: "sg-id (sg-name) | name-tag". The table below shows the inbound and outbound rules for the security groups that you chose.

▼

sg-6e5e3911 (default) ✕

Next, increase the Lambda function timeout to maximum: 5 minutes. Working with RDS will require long timeouts greater than 25 seconds:

Basic settings

Description

Memory (MB) [Info](#)
Your function is allocated CPU proportional to the memory configured.

512 MB

Timeout [Info](#)
 min sec


Next configure callservice.sh to use the name of your newly deployed Lambda function. You may choose to configure the API Gateway to have a HTTP/REST endpoint callable by curl, or simply use the AWS CLI to invoke the function directly. **Using the AWS CLI to**

invoke Lambda directly is recommend because of the potential for long timeouts when working with RDS Aurora Serverless.

Using `callservice.sh`, invoke the Lambda function to writes to the database=TEST table=mytable several times (3x-5x).

Each call to HelloMySQL will append a row to the table with the provided name.

It may be necessary to troubleshoot your Lambda function's connectivity to RDS. From the Monitoring tab of Lambda, use the [View logs in CloudWatch] button:



7. For question #7, modify the Lambda service to return the MySQL version as a response object parameter. Add a getter/setter to `Response.java` for "version". Then, add an additional SQL query to obtain the version of MySQL. Use the following SQL query:

```
select version() as version;
```

With a result set, read the value from the column, and add it to the Response object.

Now, using `callservice.sh`, run the service. Capture the complete output from the terminal as the answer for #7.

Aurora Serverless deprovisions itself after ~5 minutes of inactivity. When COLD, the first call to the serverless database takes ~25-30 seconds to bring the database back up.


The serverless freeze/thaw lifecycle is reported at the bottom of the RDS cluster page. It is possible to observe from the log how long it takes to restore the cluster in each instance:

Recent events (8)

Q Filter db events

Time	System notes
Fri Nov 02 00:10:47 GMT-700 2018	The DB cluster is resumed.
Fri Nov 02 00:10:10 GMT-700 2018	The DB cluster is being resumed.
Thu Nov 01 23:14:22 GMT-700 2018	The DB cluster is paused.
Thu Nov 01 23:14:15 GMT-700 2018	The DB cluster is being paused.
Thu Nov 01 22:57:20 GMT-700 2018	The DB cluster is resumed.
Thu Nov 01 22:56:46 GMT-700 2018	The DB cluster is being resumed.
Thu Nov 01 22:20:58 GMT-700 2018	The DB cluster is paused.
Thu Nov 01 22:20:51 GMT-700 2018	The DB cluster is being paused.

CloudWatch graphs show RDS cluster resource utilization:



Submitting Tutorial #6

Create a PDF file using Google Docs, MS Word, or OpenOffice. Capture answers to questions 1-7 and submit the PDF on Canvas.

Be sure to terminate EC2 instances, and destroy your RDS cluster once completing the tutorial.

Related Articles providing additional background:

Article describing use cases for when to use the SQLite database:

<https://www.sqlite.org/whentouse.html>

Amazon RDS Aurora MySQL 5.6 Serverless

<https://aws.amazon.com/blogs/aws/aurora-serverless-ga/>

[Serverless DB Blog Article](#)

<https://medium.com/searce/amazon-aurora-serverless-features-limitations-glitches-d07f0374a2ab>

Serverless Aurora MySQL 5.6

<https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-serverless.how-it-works.html#aurora-serverless.how-it-works.auto-scaling>

Research paper on AWS Aurora – Cloud Native relational database with built in read replication up to 15-nodes:

<https://media.amazonwebservices.com/blog/2017/aurora-design-considerations-paper.pdf>