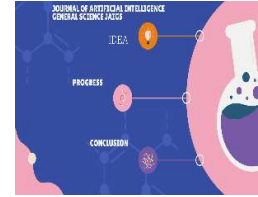




Vol. 02, Issue 01, February 2024
Journal of Artificial Intelligence General Science JAIGS

<https://ojs.boulibrary.com/index.php/JAIGS>



Automate Amazon Aurora Global Database Using Cloud Formation

Mr. Kondala Rao Patibandla

Senior Software Engineer, Aircraft Technical & Operations, Southwest Airlines,
Irving, TX 75038

ABSTRACT

ARTICLEINFO

Article History:

Received:

10.02.2024

Accepted:

15.02.2024

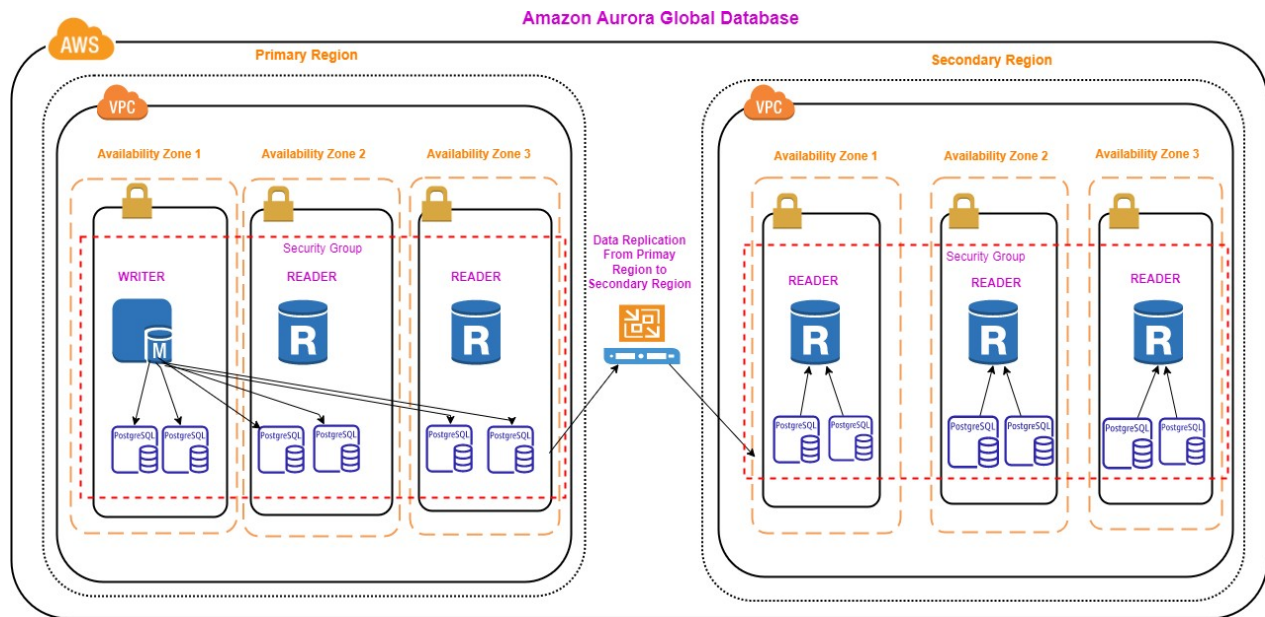
Online: 29.02.2024

Keyword: AWS RDS,
Aurora RDS, Global
Database, RDS cluster,
Aurora Global Database

This article provides an in-depth guide on deploying an AWS Aurora Global Database using AWS CloudFormation. It covers the benefits of Aurora Global Databases, such as high availability, low-latency global reads, and disaster recovery capabilities. The article details the CloudFormation template structure and key parameters needed to set up the Aurora Global Database, enabling seamless data replication across multiple AWS regions. Practical examples and best practices are included to ensure a robust and efficient deployment process.

I. Introduction

Amazon Aurora Global Database is designed for globally distributed cloud applications in AWS. It provides high availability and database resiliency by way of its ability to fail over to another AWS region. It allows a database to span multiple regions (AWS limits regions to a maximum of six), and it consists of one primary and up to five secondary regions in a global database cluster. Primary region can perform read and write operations, whereas the second region can perform read operations only. The way AWS facilitates this feature is by activating writer endpoints in the primary region and deactivating writer endpoints in secondary regions. Furthermore, Aurora replicates data from primary region to secondary regions, usually under a second.



A high-level illustration of Aurora Global Database.

II. Prerequisites

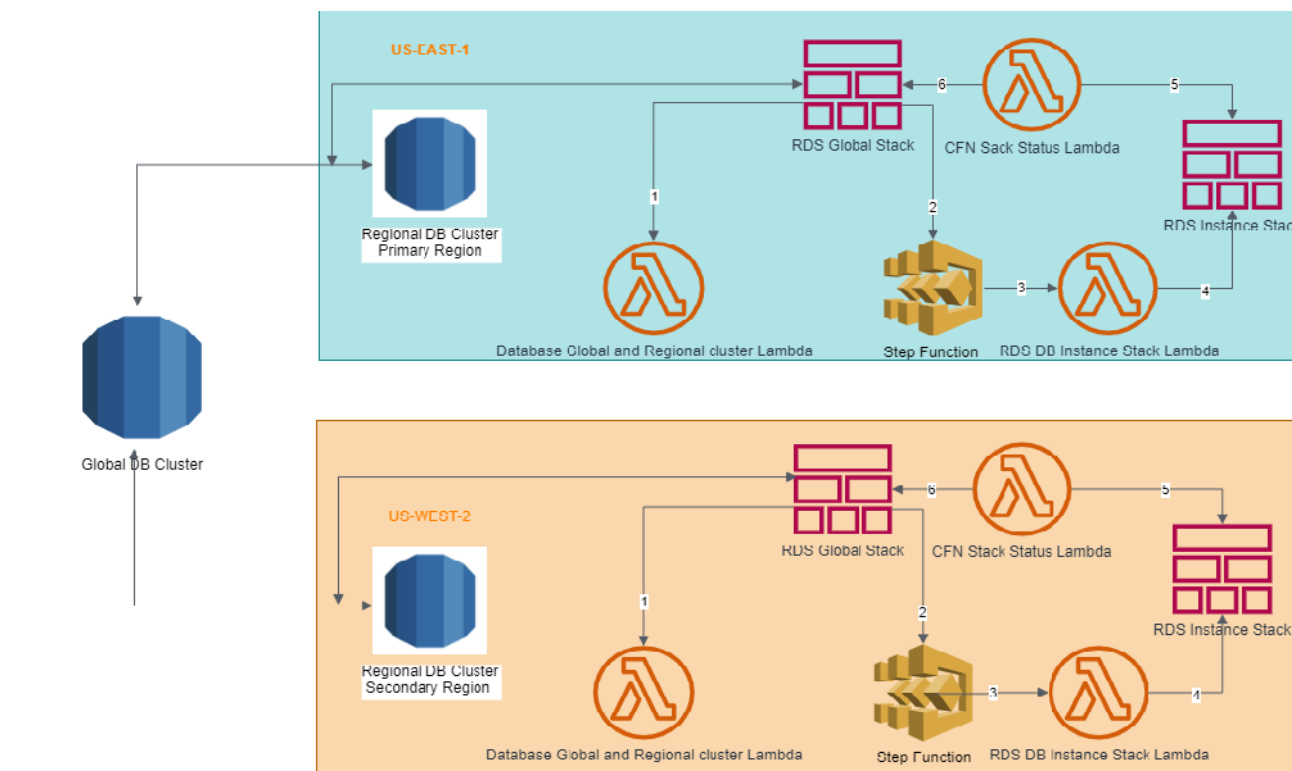
To deploy this solution, you must have the following prerequisites:

- An AWS account.
- AWS CLI with administrator permissions.
- Python 3, preferably the latest version.
- Basic knowledge of AWS SDK for Python (boto3).
- Basic knowledge of CloudFormation templates.
- Basic knowledge of Lambda and Step functions.

III. Creating an RDS Global Database

In order to create an RDS global database, we need to define global and regional database clusters. We then need to define database instances in each regional cluster.

Let us keep in mind that in order to define an RDS global database, we need to Subnet Group, RDS Security group & DBParameters group.



A sample Amazon Aurora Global Database topology.

The sample representation of an Amazon Aurora Global Database topology depicted above involves the following components and resources in its setup:

1. RDS Global Stack

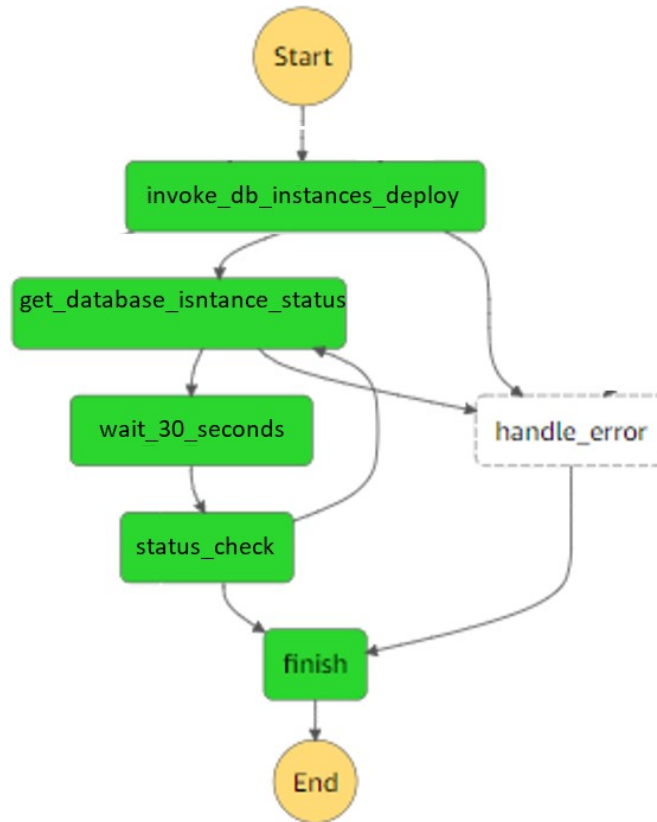
This is the base CloudFormation (CFN) stack to create RDS Aurora Global, regional database clusters, and instances in each regional cluster. This stack defines RDS subnet, Database Global and Regional cluster Lambda, StepFunction, RDS DB instances stack Lambda & CFN stack status Lambda as resources to be created.

2. Database Global and Regional Cluster Lambda

This Lambda creates regional database clusters first, and it then creates a global database cluster by assigning the newly created regional clusters to the global cluster.

3. Step Function

This state machine is responsible for creating database instances stack as a task, waiting and checking the status of this task until completion.



4. RDS DB Instance Stack Lambda

This Lambda is responsible for creating a CloudFormation stack that creates database instances.

5. CFN Stack Status Lambda

This Lambda is responsible for checking the RDS instances stack's status and returning the status to the Step Function.

All of the above resources are defined in the 'global-rds.yaml' CFN template. Code snippets for these resources are given below. For ease of reference, the individual code snippets carry the same number as the resources explained above.

AWSCLI command to deploy cloudformation template:

```

# Deploy database cluster in primary region
aws cloudformation create-stack --region=us-east-1
--stack-name global-db-east-1 --template-body global-rds.yaml
--parameters pPrivateSubnetId1=<your private subnet1>
pPrivateSubnetId2=<your private subnet2> pPrivateSubnetId3=<your private subnet3>
pDatabaseInstanceClass=db.r5.large pDatabaseEngineType=aurora-postgresql
pDatabaseEngineVersion=14.x

# Deploy database cluster in secondary region
aws cloudformation create-stack --region=us-west-2
--stack-name global-db-east-1 --template-body global-rds.yaml
--parameters pPrivateSubnetId1=<your private subnet1>
pPrivateSubnetId2=<your private subnet2> pPrivateSubnetId3=<your private subnet3>
pDatabaseInstanceClass=db.r5.large pDatabaseEngineType=aurora-postgresql
pDatabaseEngineVersion=14.x

```

1. RDS Global Stack

```

YAML
1 AWSTemplateFormatVersion: "2010-09-09"
2 Transform: "AWS::Serverless-2016-10-31"
3 Description: AWS
4
5 Parameters:
6   pPrivateSubnetId1:
7     Description: AWS RDS Global DB subnet 1 Goupd Id
8     Type: String
9
10  pPrivateSubnetId2:
11    Description: AWS RDS Global DB subnet 2 Goupd Id
12    Type: String
13
14  pPrivateSubnetId3:
15    Description: AWS RDS Global DB subnet 3 Goupd Id
16    Type: String
17
18  pDatabaseInstanceClass:
19    Description: Database Instance Type
20    Type: String
21
22  pDatabaseEngineType:
23    Description: Database Engine Type
24    Type: String
25
26  pDatabaseEngineVersion:
27    Description: Database Engine Version
28    Type: String
29
30  Resources:
31

```

2. Database Global and Regional Cluster Lambda

```

Python
1 import os
2 import boto3
3
4 def handler(event, context):
5     resource_properties = event.get("ResourceProperties")
6     # Create database regional cluster first
7     cluster_arn = create_db_regional_cluster(resource_properties)
8     # Create database global cluster with regional cluster id
9     create_global_cluster(resource_properties, cluster_arn)
10    return True
11
12 def get_rds_client(region):
13    return boto3.client('rds', region)
14
15 def create_global_cluster(resource_properties, cluster_arn):
16    rds_client = get_rds_client(resource_properties.get('Region'))
17    rds_client.create_global_cluster(
18        GlobalClusterIdentifier=resource_properties.get('GlobalClusterId'),
19        SourceDBClusterIdentifier=cluster_arn
20    )
21
22 def create_db_regional_cluster(resource_properties):
23    rds_client = get_rds_client(resource_properties.get('Region'))
24    response = rds_client.create_db_cluster(
25        DBClusterIdentifier=resource_properties.get('ClusterId'),
26        Engine=resource_properties.get('Engine'),
27        EngineVersion=resource_properties.get('EngineVersion'),
28        Port=resource_properties.get('Port')
29    )
30    return response.get('DBCluster').get('DBClusterArn')
31

```

3. RDS DB Instance Stack Lambda

```

Python
1 import boto3
2
3 def handler(event, context):
4     stack_name = event.get('StackName')
5     region = event.get('Region')
6     params = event.get('Parameters')
7     params['pDatabaseParameterGroup'] = get_rds_params_group(region)
8     params['pDatabaseSubnetGroup'] = get_rds_subnet_group(region)
9
10    def get_cfn_client(region):
11        return boto3.client('cloudformation', region)
12
13    def get_rds_client(region):
14        return boto3.client('rds', region)
15
16    def create_database_instances(stack_name, params, region, template_path):
17        get_cfn_client(region).create_stack(
18            StackName=stack_name,
19            TemplateBody=parse_template(template_path)
20            Parameters=params,
21            Capabilities=['CAPABILITY_AUTO_EXPAND']
22        )
23
24    def parse_template(template_path, region):
25        with open(template) as template_file:
26            data = template_file.read()
27        get_cfn_client(region).validate_template(Template=data)
28        return data
29
30    def get_rds_params_group(region):
31        paras_group = []
32        paginator = get_rds_client(region).get_paginator('descrip_db_cluster_parameter_group')
33        for grouppage in paginator.paginate():
34            paras_group =return_list+ grouppage.get('DBClusterParameterGroup')
35        return paras_group
36
37    def get_rds_subnet_group(region):
38        subnet_group = []
39        paginator = get_rds_client(region).get_paginator('describe_db_subnet_group')
40        for grouppage in paginator.paginate():
41            subnet_group =return_list+ grouppage.get('DBSubnetGroup')
42        return subnet_group

```

4. RDS Instance Stack

```

YAML
1 AWSTemplateFormationVersion: "2010-09-09"
2 Transform: "AWS::Serverless-2016-10-31"
3 Description: AWS
4
5 Parameters:
6   pDatabaseInstanceClass:
7     Description: Database Instance Class
8     Type: String
9
10  pDatabaseSubentGroup:
11    Description: Database Subnet Group
12    Type: String
13
14  pDatabaseParameterGroup:
15    Description: Database Parameter Group
16    Type: String
17
18 Resources:
19 rPrimaryDatabaseInstance:
20   Type: AWS::RDS::DBInstance
21   Properties:
22     DBInstanceIdentifier: !Sub 'db-instance-${AWS::Region}-1'
23     DBClusterIdentifier: !Sub regional-db-cluster-${AWS::Region}
24     DBInstanceClass: !Ref pDatabaseInstanceClass
25     DBSubnetGroupName: !Ref pDatabaseSubentGroup
26     DBParameterGroup: !Ref pDatabaseParameterGroup
27     Engine: aurora-Postgresql
28
29 rReplicationDatabaseInstance1:
30   Type: AWS::RDS::DBInstance
31   Properties:
32     DBInstanceIdentifier: 'db-instance-${AWS::Region}-2'
33     DBClusterIdentifier: !Sub test-cluster-${AWS::Region}
34     DBInstanceClass: !Ref pDatabaseInstanceClass
35     DBSubnetGroupName: !Ref pDatabaseSubentGroup
36     DBParameterGroup: !Ref pDatabaseParameterGroup
37     Engine: aurora-Postgresql
38
39 rReplicationDatabaseInstance2:
40   Type: AWS::RDS::DBInstance
41   Properties:
42     DBInstanceIdentifier: 'db-instance-${AWS::Region}-3'
43     DBClusterIdentifier: !Sub test-cluster-${AWS::Region}
44     DBInstanceClass: !Ref pDatabaseInstanceClass
45     DBSubnetGroupName: !Ref pDatabaseSubentGroup
46     DBParameterGroup: !Ref pDatabaseParameterGroup

```

5. CFN Stack Status Lambda


```

Python
1 import boto3
2
3 def handler(event, context):
4     stack_name = event.get('StackName')
5     region = event.get('Region')
6     stack_status = get_stack_status(stack_name, region)
7     if stack_status == 'CREATE_IN_PROGRESS':
8         return 'WAIT'
9     if stack_status == 'CREATE_COMPLETE':
10        return 'SUCCESS'
11
12 def get_cfn_client(region):
13     return boto3.client('cloudformation', region)
14
15 def get_stack_status(stack_name, region):
16     stack_response = get_cfn_client(region).describe_stacks(
17         StackName=stack_name
18     ).get('Stack')
19
20     if stack_response:
21         stack_status = stack_response[0].get('StackStatus')
22
23     return stack_status

```

When all the steps defined above are completed successfully, one can see the newly created Amazon Aurora Global PostgreSQL Database, as shown below.

<input type="checkbox"/>	DB identifier	DB cluster identifier	Role	Engine	Region & AZ	Size	Status
<input type="radio"/>	<input type="checkbox"/> global-db-cluster	global-db-cluster	Global database	Aurora PostgreSQL	2 regions	2 clusters	✔ Available
<input type="radio"/>	<input type="checkbox"/> regional-db-cluster-us-east-1	regional-db-cluster-us-east-1	Primary cluster	Aurora PostgreSQL	us-east-1	3 instances	✔ Available
<input type="radio"/>	<input type="checkbox"/> db-instance-us-east-1-1	regional-db-cluster-us-east-1	Writer instance	Aurora PostgreSQL	us-east-1b	db.r5.large	✔ Available
<input type="radio"/>	<input type="checkbox"/> db-instance-us-east-1-2	regional-db-cluster-us-east-1	Reader instance	Aurora PostgreSQL	us-east-1c	db.r5.large	✔ Available
<input type="radio"/>	<input type="checkbox"/> db-instance-us-east-1-3	regional-db-cluster-us-east-1	Reader instance	Aurora PostgreSQL	us-east-1a	db.r5.large	✔ Available
<input type="radio"/>	<input type="checkbox"/> regional-db-cluster-us-west-2	regional-db-cluster-us-west-2	Secondary cluster	Aurora PostgreSQL	us-west-2	3 instances	✔ Available
<input type="radio"/>	<input type="checkbox"/> db-instance-us-west-2-1	regional-db-cluster-us-west-2	Reader instance	Aurora PostgreSQL	us-west-2c	db.r5.large	✔ Available
<input type="radio"/>	<input type="checkbox"/> db-instance-us-west-2-2	regional-db-cluster-us-west-2	Reader instance	Aurora PostgreSQL	us-west-2b	db.r5.large	✔ Available
<input type="radio"/>	<input type="checkbox"/> db-instance-us-west-2-3	regional-db-cluster-us-west-2	Reader instance	Aurora PostgreSQL	us-west-2a	db.r5.large	✔ Available

IV. Fail-Over Scenario

With Aurora Global Database, one can expect two failover scenarios – managed planned failover and unplanned failover.

Managed Planned Fail-Over

A managed planned fail-over scenario works best when both the regions of the global cluster are in normal operation. When performing this operation, the writer endpoint in the active region is replaced with a reader endpoint. Vice-versa happens in the passive region, i.e., the reader endpoint in the passive region is replaced with the writer endpoint. This ensures that active and passive regions are flipped after performing the fail-over operation.

Planned fail-over can be performed in multiple ways. Some of the ways are:

- Using AWS console
- AWS CLI
- Scripts that use AWS SDK
- AWS CDK

Using AWS Console

The picture below depicts options to select in the AWS console's 'Databases' section on the 'RDS' page.

The screenshot shows the AWS console 'Databases' section. At the top, there are buttons for 'Group resources', 'Modify', 'Actions', 'Restore from S3', and 'Create database'. Below these is a search bar and a table of database clusters. The 'Fail over global database' option is highlighted in the Actions menu for the selected cluster.

DB identifier	DB cluster identifier	Role	Engine	Regions	Instances	Status	CPU
global-db-cluster	global-db-cluster	Global database	Aurora PostgreSQL	2 regions	2 clusters	Available	-
regional-db-cluster-us-east-1	regional-db-cluster-us-east-1	Primary cluster	Aurora PostgreSQL	us-east-1	3 instances	Available	-
db-instance-us-east-1-1	regional-db-cluster-us-east-1	Writer instance	Aurora PostgreSQL	us-east-1b	db.r5.large	Available	10.74%
db-instance-us-east-1-2	regional-db-cluster-us-east-1	Reader instance	Aurora PostgreSQL	us-east-1c	db.r5.large	Available	8.97%
db-instance-us-east-1-3	regional-db-cluster-us-east-1	Reader instance	Aurora PostgreSQL	us-east-1a	db.r5.large	Available	8.72%
regional-db-cluster-us-west-2	regional-db-cluster-us-west-2	Secondary cluster	Aurora PostgreSQL	us-west-2	3 instances	Available	-
db-instance-us-west-2-1	regional-db-cluster-us-west-2	Reader instance	Aurora PostgreSQL	us-west-2c	db.r5.large	Available	-
db-instance-us-west-2-2	regional-db-cluster-us-west-2	Reader instance	Aurora PostgreSQL	us-west-2b	db.r5.large	Available	-
db-instance-us-west-2-3	regional-db-cluster-us-west-2	Reader instance	Aurora PostgreSQL	us-west-2a	db.r5.large	Available	-

AWS CLI

Execute the command given below to perform managed planned fail-over using AWS CLI.

```
awsrds --region us-east-1 failover-global-cluster
--global-cluster-identifier global-db-cluster
--target-db-cluster-identifier arn:aws:rds:us-west-2:{AWS Account Number}:cluster:db-regional-cluster-us-west-2
```

Unplanned Fail-Over

We perform unplanned fail-over when the current active database cluster goes down. The following steps need to be performed:

- Remove the passive region (secondary region) database cluster from the global cluster. After removing it from the global database cluster, this works as a stand-alone database cluster, and one of the reader instances turns into a writer instance. We can assign it back to the global cluster, allowing us to perform write and read operations on a stand-alone database cluster.
- Delete the affected database cluster, which was running as an active cluster in the global database once the affected AWS region is operational. Then, assign a stand-alone cluster to the global database as an active region cluster. Finally, create a new secondary database cluster in the previously affected region and assign it to the global database cluster as a passive region cluster.

V. Conclusion

I have defined comprehensive steps which would create and configure an Amazon Aurora Global Database setup. This would provide a database with high availability and fault tolerance. This database setup can cater to a multi-regional application setup, making it resilient to failures. We also provided steps to automate and simplify creating a complex global database setup.