



MORRIS
LABS

Step by step guide on how to setup
and implement Oracle tasks on DMS

September 10th, 2020.

Index

1. Introduction	3
2. Endpoints Setup	3
2.1 Origin Endpoint	3
2.2 Destination of Endpoints	6
3. Task creation	9
3.1. Tasks rules	9
3.1.1. Selecting boards	10
3.1.2. Rename schemas	10
3.1.3. Add columns	10
3.1.4. Practical example	11
3.2. Task setup	14
4. Task management	19
4.1. Resume task	19
5. Monitoring alerts	20
5.1 Set up SNS alert	21
5.2 Confirm subscription to topic	23

1. Introduction

The objective of this document is to serve as a guide to configure a replication task in DMS, considering its origin and destination endpoints, specific rules for each task, such as adding columns, renaming schemes. Alert setup using Amazon SNS is also covered.

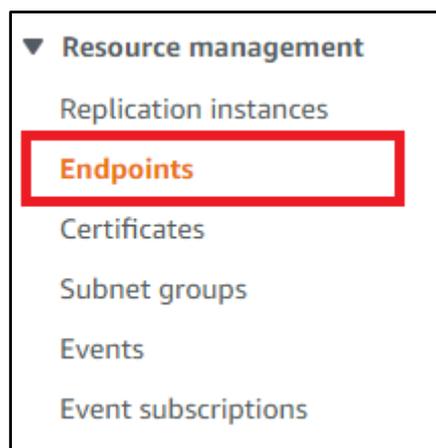
2. Endpoints Setup

An endpoint provides connection, data store type, and location information about its data store. AWS Database Migration Service uses this information to connect to a data warehouse and migrate data from a source endpoint to a destination endpoint.

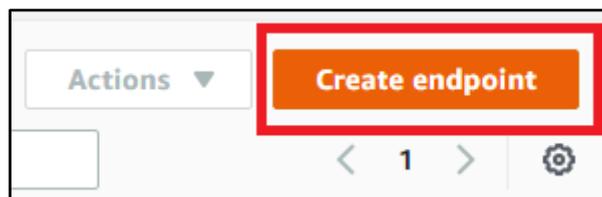
2.1 Origin Endpoint

To create a data source endpoint, follow these steps:

1. Enter the AWS DMS console
<https://console.aws.amazon.com/dms/v2/home?region=us-east-1#dashboard>
2. Enter the option "Endpoints" from the side menu of the console



3. Later, to create a new data source endpoint, press the option "create endpoint" in the upper right corner of the screen.



4. Select the option "source endpoint"

Endpoint type [Info](#)

Source endpoint
A source endpoint allows AWS DMS to read data from a database (on-premises or in the cloud), or from other data source such as Amazon S3.

Target endpoint
A target endpoint allows AWS DMS to write data to a database, or to other data source.

Select RDS DB instance

- For the purposes of this document, a data source endpoint is created for an Oracle database. To do this, the following configuration must be entered:

Endpoint configuration

Endpoint identifier [Info](#)
A label for the endpoint to help you identify it.

NOMBRE-ENDPOINT

Source engine
The type of database engine this endpoint is connected to.

oracle

Server name

IP O HOST DE LA DB ORACLE

Port
The port the database runs on for this endpoint.

PUERTO

Secure Socket Layer (SSL) mode
The type of Secure Socket Layer enforcement

none

 Input valid number only

User name [Info](#)

USUARIO_DB

Password [Info](#)

•

SID/Service name
Use Service name or SID as applicable. Use DB name in case of RDS for Oracle

NOMBRE DE LA DB A REPLICAR

Endpoint configuration oracle

- Subsequently, the extra attributes of the connection are specified, as follows:

▼ Endpoint-specific settings

Extra connection attributes
Type any additional connection parameters here. See the documentation for more information.

useLogminerReader=N;useBfile=Y

Endpoint-specific settings

This extra configuration is exclusive to Oracle endpoints and allows access to Oracle redo logs using the binary reader method for the CDC.

- To test that the endpoint connection is correct, apply “run test” as follows:

▼ Test endpoint connection (optional)

VPC
vpc-l

Replication instance
A replication instance performs the database migration
-dms-replicacion-1

⚠ Your endpoint will always be created even if the connection fails
After clicking 'Run test', DMS creates the endpoint with the details you provided and attempts to connect to it. If the connection fails, you can edit the endpoint definition and test the connection again. You can also delete the endpoint manually.

Run test

Endpoint identifier	Replication instance	Status	Message
No records found			

Test endpoint connection

- If the process is done satisfactorily, the following screen will be displayed:

Test endpoint connection

Replication instance
A replication instance performs the database migration
-dms-replicacion-1

Run test

Endpoint identifier	Replication instance	Status	Message
i-dms-endpoint-aspen-1	-dms-replicacion-1	successful	

Test connection successful

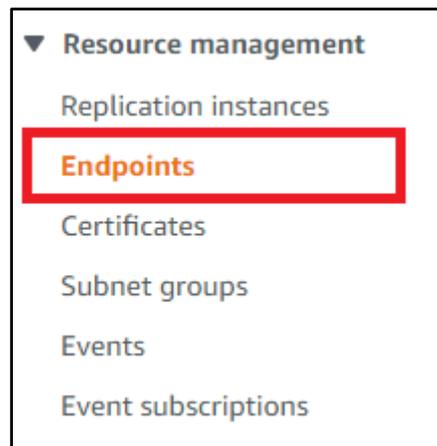
9. Finally press the option "Create endpoint" in the lower right corner of the form.



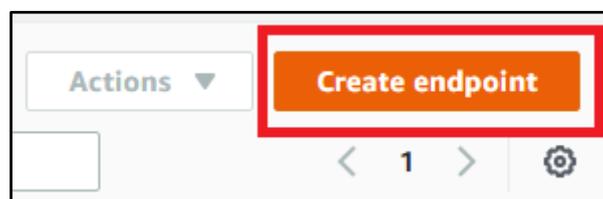
2.2 Destination of Endpoints

To create a data destination (source) endpoint, follow these steps:

1. Enter the AWS DMS console
<https://console.aws.amazon.com/dms/v2/home?region=us-east-1#dashboard>
2. Enter the option "**Endpoints**" from the side menu of the console



3. Later, to create a new data destination endpoint (target), press the option "**Create endpoint**" in the upper right corner of the screen.



4. Select the option "**Target endpoint**"

Endpoint type [Info](#)

Source endpoint
A source endpoint allows AWS DMS to read data from a database (on-premises or in the cloud), or from other data source such as Amazon S3.

Target endpoint
A target endpoint allows AWS DMS to write data to a database, or to other data source.

Select RDS DB instance

5. For the purposes of this document, a data destination (source) endpoint is created for a Redshift database. To do this, the following configuration must be entered:

Endpoint configuration

Endpoint identifier [Info](#)
A label for the endpoint to help you identify it.

NOMBRE-ENDPOINT

Target engine
The type of database engine this endpoint is connected to.

redshift

Server name

HOST_REDSHIFT_CLUSTER

Port
The port the database runs on for this endpoint.

PUERTO

Secure Socket Layer (SSL) mode
The type of Secure Socket Layer enforcement

none

User name [Info](#)

USUARIO_REDSHIFT

Password [Info](#)

....

Database name

DB_REDSHIFT

Endpoint configuration redshift

6. To test that the endpoint connection is correct, apply "run test" as follows:

▼ **Test endpoint connection (optional)**

VPC
 -vpc-1

Replication instance
 A replication instance performs the database migration
 -dms-replicacion-1

⚠ Your endpoint will always be created even if the connection fails ✕

After clicking 'Run test', DMS creates the endpoint with the details you provided and attempts to connect to it. If the connection fails, you can edit the endpoint definition and test the connection again. You can also delete the endpoint manually.

Run test

Endpoint identifier	Replication instance	Status	Message
No records found			

Test endpoint connection

7. If the process is done satisfactorily, the following screen will be displayed:

Test endpoint connection

Replication instance
 A replication instance performs the database migration
 -dms-replicacion-1

Run test

Endpoint identifier	Replication instance	Status	Message
<input type="text" value=""/> -dms-endpoint-redshift	<input type="text" value=""/> -dms-replicacion-1	successful	

Test connection successful

8. Finally press the option "**Create endpoint**" in the lower right corner of the form.

Cancel
Create endpoint

3. Task creation

3.1. Tasks rules

Table mapping uses various types of rules to specify the data source, source schema, data, and transformations to occur during the task. You can use table mapping to specify the individual boards in a database to be migrated and the schema to use for migration.

In the context of this project, the rules that allow obtaining certain boards from a specific origin, renaming the replicated schemas and adding different columns to the different selected boards will be configured.

These types of rules are represented in a file in JSON format which must be added to the corresponding task when creating them.

3.1.1. Selecting boards

The task to select one or more boards for replication is as follows:

```
{
  "rule-type": "selection",
  "rule-id": "<ID_REGLA>",
  "rule-name": "<ID_REGLA>",
  "object-locator": {
    "schema-name": "<ESQUEMA_ORIGEN> ",
    "table-name": "<TABLA_A_REPLICAR>"
  },
  "rule-action": "include"
}
```

It should be considered that both the board and the schema indicated in the rule must exist in the database specified in the endpoint source of the corresponding task, otherwise DMS will ignore the newly created rule because it cannot find the table or schema indicated.

3.1.2. Rename schemas

The task that allows renaming one or more schemas for replication is as follows:

```
{
  "rule-type": "transformation",
  "rule-id": "<ID_REGLA>",
```

```

"rule-name": "<ID_REGLA>",
"rule-target": "schema",
"object-locator": {
  "schema-name": "<ESQUEMA_ORIGEN>"
},
"rule-action": "rename",
"value": "<NUEVO NOMBRE ESQUEMA>",
"old-value": null
}

```

It should be noted that this schema renown is only visible in the data destination and is not required if it is not required. If this task is omitted, DMS will use the schema name as it is in the source.

3.1.3.Add columns

In the context of this project, two columns were added to the different boards to be replicated. The first, a column that shows the date the replication was performed, and the second, a column that shows the type of operation performed (INSERT, UPDATE). The tasks that allow you to add one or more columns to a specific table are the following:

```

{
  "rule-type": "transformation",
  "rule-id": "<ID_REGLA>",
  "rule-name": "<ID_REGLA>",
  "rule-target": "column",
  "object-locator": {
    "schema-name": "<ESQUEMA_ORIGEN>",
    "table-name": "<TABLA_A_REPLICAR>"
  },
  "rule-action": "add-column",
  "value": "fecha_replicacion",
  "expression": "$AR_H_TIMESTAMP",
  "data-type": {
    "type": "datetime",
    "precision": 6
  }
},

```

Add replication date column

```

{
  "rule-type": "transformation",

```

```

"rule-id": "<ID_REGLA>",
"rule-name": "<ID_REGLA>",
"rule-target": "column",
"object-locator": {
  "schema-name": "<ESQUEMA_ORIGEN>",
  "table-name": "<TABLA_A_REPLICAR>"
},
"rule-action": "add-column",
"value": "tipo_operacion",
"expression": "$AR_H_OPERATION",
"data-type": {
  "type": "string",
  "length": 50
}
},

```

Add operation type column

It should be considered that both the table and the schema indicated in the rule must exist in the database specified in the endpoint source of the corresponding task in addition to having the data selection rule for the table and schema indicated in the task, in otherwise DMS will skip the newly created rule because it cannot find the designated table or schema or because it was not selected for replication.

3.1.4. Practical example

To exemplify the rules file for the task to be created, a practical case is presented to better understand how to form this JSON file.

As an example, it is required to replicate the following boards, which belong to the same source system and must be stored in the destination system called "Destination_1".

Nombre Tabla	Sistema de Origen	Sistema destino
Tabla_1	Origen_1	Destino_1
Tabla_2	Origen_1	Destino_1

Once the boards to be replicated have been defined, the following JSON file is created:

```

{
  "rules": [
    // Tarea de renombre de esquemas (origen_1 -> destino_1)
    {

```

```

    "rule-type": "transformation",
    "rule-id": "1",
    "rule-name": "1",
    "rule-target": "schema",
    "object-locator": {
        "schema-name": "Origen_1"
    },
    "rule-action": "rename",
    "value": "Destino_1",
    "old-value": null
},
// Tarea de selección de tabla (tabla_1)
{
    "rule-type": "selection",
    "rule-id": "2",
    "rule-name": "2",
    "object-locator": {
        "schema-name": "Origen_1",
        "table-name": "Tabla_1"
    },
    "rule-action": "include"
},
// Tarea de agregación de columna para fecha de replicación (tabla_1)
{
    "rule-type": "transformation",
    "rule-id": "3",
    "rule-name": "3",
    "rule-target": "column",
    "object-locator": {
        "schema-name": "Origen_1",
        "table-name": "Tabla_1"
    },
    "rule-action": "add-column",
    "value": "fecha_replicacion",
    "expression": "$AR_H_TIMESTAMP",
    "data-type": {
        "type": "datetime",
        "precision": 6
    }
},
// Tarea de agregación de columna para tipo de operación (tabla_1)
{
    "rule-type": "transformation",
    "rule-id": "4",
    "rule-name": "4",
    "rule-target": "column",
    "object-locator": {
        "schema-name": "Origen_1",
        "table-name": "Tabla_1"
    }
}

```

```

    },
    "rule-action": "add-column",
    "value": "tipo_operacion",
    "expression": "$AR_H_OPERATION",
    "data-type": {
        "type": "string",
        "length": 50
    }
},
// Tarea de selección de tabla (tabla_2)
{
    "rule-type": "selection",
    "rule-id": "5",
    "rule-name": "5",
    "object-locator": {
        "schema-name": "Origen_1",
        "table-name": "Tabla_2"
    },
    "rule-action": "include"
},
// Tarea de agregación de columna para fecha de replicación (tabla_2)
{
    "rule-type": "transformation",
    "rule-id": "6",
    "rule-name": "6",
    "rule-target": "column",
    "object-locator": {
        "schema-name": "Origen_1",
        "table-name": "Tabla_2"
    },
    "rule-action": "add-column",
    "value": "fecha_replicacion",
    "expression": "$AR_H_TIMESTAMP",
    "data-type": {
        "type": "datetime",
        "precision": 6
    }
},
// Tarea de agregación de columna para tipo de operación (tabla_2)
{
    "rule-type": "transformation",
    "rule-id": "7",
    "rule-name": "7",
    "rule-target": "column",
    "object-locator": {
        "schema-name": "Origen_1",
        "table-name": "Tabla_2"
    },
    "rule-action": "add-column",

```

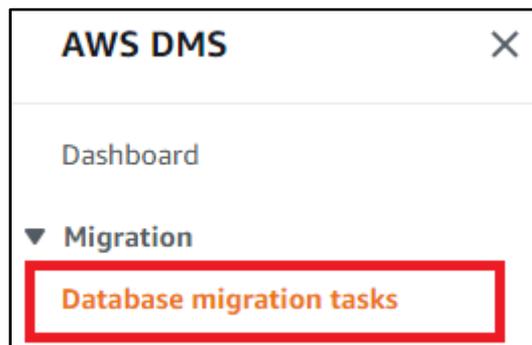
```
    "value": "tipo_operacion",
    "expression": "$AR_H_OPERATION",
    "data-type": {
      "type": "string",
      "length": 50
    }
  }
]
```

3.2. Task setup

Once the source and destination endpoints and the rules for the different tasks have been created, in order to effectively perform the data replication process, a task must be created that can perform this action.

In order to do this, the following steps must be performed:

1. Enter the AWS DMS console
<https://console.aws.amazon.com/dms/v2/home?region=us-east-1#dashboard>
2. Enter the option "**Database migration tasks**" from the side menu of the console



3. Later, to create a new data source endpoint, press the option "**create task**" in the upper right corner of the screen.



4. Once inside the DMS task creation panel, the task name is specified

Task configuration

Task identifier

5. The replication instance to use is selected

Replication instance

Choose a replication instance ▲

 |

6. The endpoint source of the task is selected

Source database endpoint

Choose a source database endpoint ▲

Q |

- .dms-endpoint-

- The target endpoint of the task is selected

Target database endpoint

Choose a target database endpoint ▲

Q |

dms-endpoint-redshift

- Subsequently, the type of migration is chosen, in this case the type that allows full load and CDC will be selected.

Migration type [Info](#)

Migrate existing data and replicate ongoing changes ▼

- To activate the monitoring of the task logs, the CloudWatch logs are activated as follows.

Enable CloudWatch logs [Info](#)

DMS task logging uses Amazon CloudWatch to log information during the migration process. You can change the component activities logged and the amount of information logged for each one.

 CloudWatch logs usage will be charged at standard rates. See [here](#) for more details.

Increasing message severity can fill up the available disk space. Do so only when needed for debugging. You can delete task logs to free up disk space after debugging is complete. [Learn more](#) 

Source Unload

Message severity logged when data is unloaded from the source database.

Default ▼

Target Load

Message severity logged when data is loaded into the target database.

Default ▼

Task manager

Message severity logged when the task manager triggers an event.

Default ▼

Source capture

Message severity logged when data is captured from the source database.

Default ▼

Target apply

Message severity logged when data and data definition language (DDL) statements are applied to the target database.

Default ▼

10. To associate the rules created in step 3.1 of this document, you must select the **JSON editor** option in the **Table mapping** section and paste the content of the generated JSON file.

Table mappings

Editing mode

Guided UI
Set up your table mapping rules using a step-by-step guided interface.

JSON editor [Learn more](#) 
Enter your table mapping rules directly, in JSON format.

Specify at least one selection rule with an include action. After you do this, you can add one or more transformation rules.

```
1 {
2   "rules": [
3     // Tarea de renombre de esquemas (origen_1 -> destino_1)
4     {
5       "rule-type": "transformation",
6       "rule-id": "1",
7       "rule-name": "1",
8       "rule-target": "schema",
9       "object-locator": {
10        "schema-name": "Origen_1"
11      },
12       "rule-action": "rename",
13       "value": "Destino_1",
14       "old-value": null
15     },
16   ],
17 }
```

11. So that the task to be created does not start when it is created but when the user requires it, the following option is selected.

Migration task startup configuration

Start migration task

Automatically on create
Available only if the premigration assessment is not enabled.

Manually later

12. It is also configured that the maximum number of tables that are replicated in parallel in the same task to 5, to do this, the **Maximum number of boards to load in parallel** option must be changed in the **Advanced task settings** section as follows.

Full load tuning settings

<p>Maximum number of tables to load in parallel</p> <input style="width: 90%; border: 1px solid #ccc;" type="text" value="5"/> <p><small>Should be an integer range from 1 to 49</small></p>	<p>Transaction consistency timeout (seconds)</p> <input style="width: 90%; border: 1px solid #ccc;" type="text" value="600"/> <p><small>Should be an integer range from 0 to 2147483647</small></p>
<p>Commit rate during full load</p> <input style="width: 90%; border: 1px solid #ccc;" type="text" value="10000"/> <p><small>Should be an integer range from 1 to 50000</small></p>	

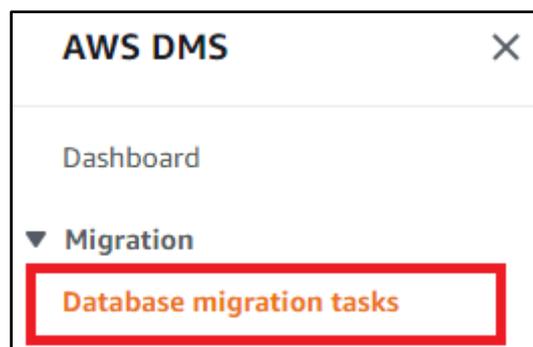
13. Finally, click on the Create task option in the lower right corner of the screen.

4. Task management

4.1. Resume task

If for any reason it is required to pause a task created by DMS, it allows to resume it again. In order to do this, the following steps must be followed.

1. Enter the AWS DMS console
<https://console.aws.amazon.com/dms/v2/home?region=us-east-1#dashboard>
2. Enter the option "**Database migration tasks**" from the side menu of the console



3. Subsequently, all the tasks created by the users will be displayed, identifying the tasks that are stopped with the Stopped status.



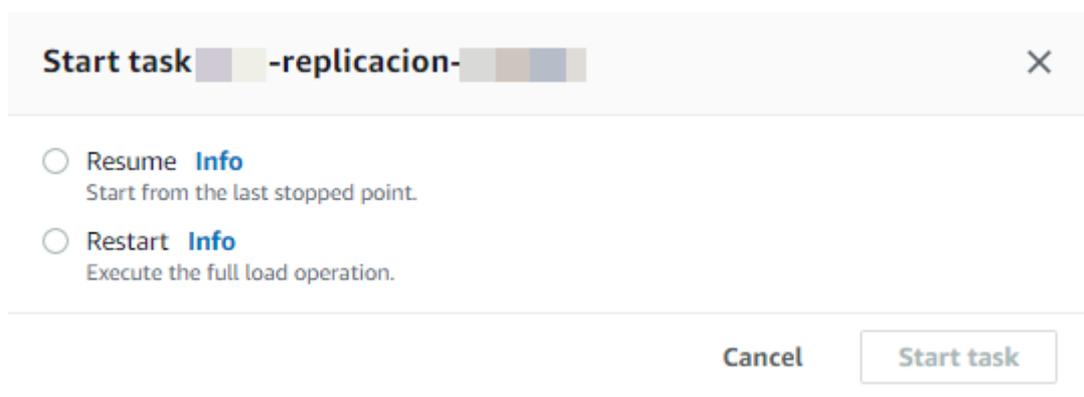
4. To be able to resume this type of task, you must select the name of the paused task by clicking on it.



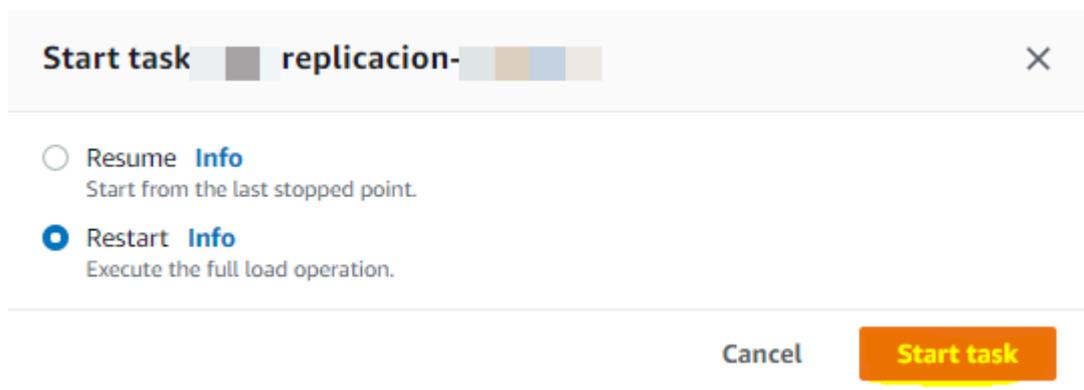
5. Within the task, you must press the **Actions** option in the upper right corner and then the **Restart / Resume** option.



6. Once this is done, two options will be displayed. The first, **Resume**, which allows you to resume the task from the point where it was stopped, and the second **Restart**, which allows you to run the task from the beginning.



7. Finally, the option required by the user is selected and **Start task** is pressed to begin with the execution of the task.



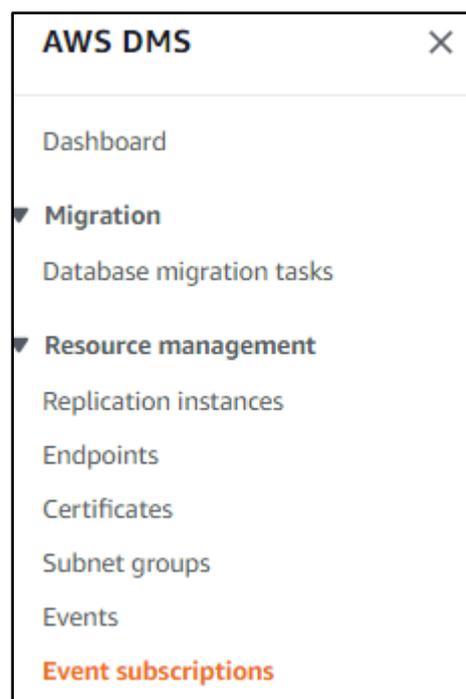
5. Monitoring alerts

5.1 Set up SNS alert

Once the tasks are defined, it is necessary to be able to monitor them and be pending in the event that one of them presents some type of failure. To do this, the Amazon SNS service is used, which allows sending alerts to the emails of the users defined with the failures that the tasks present in the CloudWatch logs.

To do this, the following steps must be followed:

1. Enter the AWS DMS console
<https://console.aws.amazon.com/dms/v2/home?region=us-east-1#dashboard>
2. Enter the **Event subscriptions** option in the left side menu.



3. To start creating the subscription event that allows creating alerts of different types and sending them to different users, select the **Create event** subscription option in the upper right corner of the screen.



4. For starters, the name of the event is specified.

Details

Name
The name for your event subscription

Enabled

5. Then a new topic of the Amazon SNS service is created which is in charge of sending emails to different designated users through their email. If more than one email is associated, they must be divided by space.

Target

Send notification to

Existing topics

Create new email topic

Create new SMS topic

Topic name

With these recipients
Email addresses or phone numbers of SMS enabled devices to send the notifications to

6. After this, the action to be taken to trigger an Amazon SNS alert is configured. First, the **Source type** is selected as **replication-task**, which indicates that the events will be developed at the task level.
- 7.

Event source

Source type
Source Type of resource this subscription will consume events from

8. Then the specific event that will trigger the task-level alerts is selected. In this case, they are configured so that when some type of failure occurs in the tasks, an alert is sent, so the triggering event for the tasks must be **Failure**

Event categories

All event categories

Select specific event categories

Choose event categories ▼

failure ✕

9. Finally, the tasks to be monitored are selected, how all the tasks created and to be created will be configured, the All tasks option is **selected**.

Replication tasks

All tasks

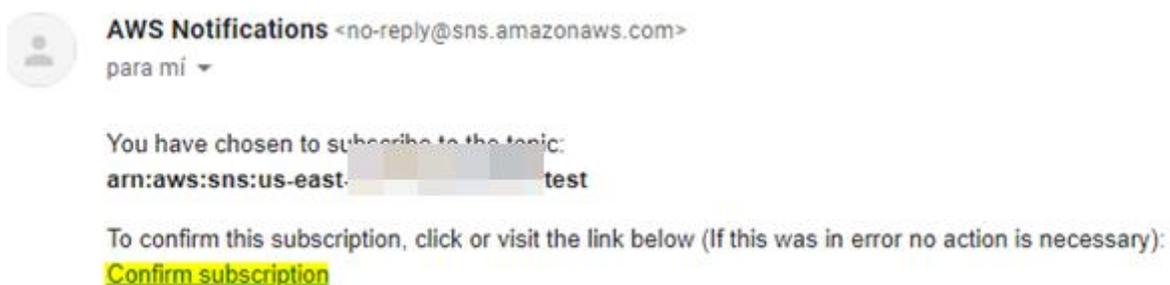
Select specific tasks

10. Finally, the **Create event subscription** option is selected in the lower right corner of the screen, to create the event that monitors the created tasks.

Cancel **Create event subscription**

5.2 Confirm subscription to topic

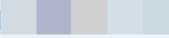
Once the alert is created, the users associated with it will receive a subscription confirmation email which they must open and press the **Confirm subscription** option to subscribe to the topic and be able to receive the generated alerts.



Finally the successful subscription to the topic is displayed.



Subscription confirmed!

You have subscribed  morrisopazo.com to the topic:
test.

Your subscription's id is:

arn:aws:sns:us-east-1: test:

If it was not your intention to subscribe, [click here to unsubscribe](#).