

Connecting to Redshift


- [Setup](#)
 - [Redshift Settings](#)
 - [Location](#)
 - [Create Cluster](#)
 - [Cluster Details](#)
 - [Node Configuration](#)
 - [Additional Configuration](#)
 - [Review & Launch](#)
 - [View Clusters](#)
 - [Cluster Name](#)
 - [Cluster Details](#)
- [Yellowfin](#)
 - [Login](#)
 - [Add Data Source](#)
 - [Data Source Details](#)
 - [Connection Details](#)

Setup

[top](#)

Redshift Settings


1. Log in to Amazon Web Services (AWS).
2. Open the **Services** menu at the top of the page, click on **Redshift** within the **Database** category.



Location

While on the Amazon Redshift page it's important to select the closest location to your users. This will generally result in the fastest response times.


3. Click on the **Location** currently in use next to your name at the top right of the page.
4. Select the location closest to you.



Create Cluster

5. While still on the Amazon Redshift page, click on the **Launch Cluster** button.

From here you will be walked through a setup process where you define and configure the cluster.



Cluster Details

On this step you will need to define the following:

- Cluster Identifier
- Database Name
- Database Port
- Master User Name
- Master Password

Each option is documented on the page.

6. Configure the cluster details and click the **Continue** button.

Services ▾ **Edit** ▾

Amazon Redshift

Clusters

Snapshots

Security

Parameter Groups

Reserved Nodes

Events

CLUSTER DETAILS **NODE CONFIGURATION** **ADDITIONAL CONFIGURATION** **REVIEW**

Provide the details of your cluster. Fields marked with * are required.

Cluster Identifier* This is the unique key that identifies a cluster. This parameter is stored as a lowercase string. (e.g. my-dw-instance)

Database Name Optional. A default database named dev is created for the cluster. Optionally, specify a custom database name (e.g. mydb) to create an additional database.

Database Port* Port number on which the database accepts connections.

Master User Name* Name of master user for your cluster. (e.g. awsuser)

Master User Password* Password must contain 8 to 64 printable ASCII characters excluding /, ", ', \, and @. It must contain 1 uppercase letter, 1 lowercase letter, and 1 number.

Confirm Password* Confirm Master User Password.

Node Configuration

Here you will need to define the number and types of nodes. Each option is documented on the page.

7. Define the Type and Number of Nodes. Click Continue.

Services ▾ **Edit** ▾

Amazon Redshift

Clusters

Snapshots

Security

Parameter Groups

Reserved Nodes

Events

CLUSTER DETAILS **NODE CONFIGURATION** **ADDITIONAL CONFIGURATION** **REVIEW**

Choose a number of nodes and Node Type below. Number of Compute Nodes is required for multi-node clusters.

Node Type Specifies the compute, memory, storage, and I/O capacity of the cluster's nodes.

CPU 7 EC2 Compute Units (2 virtual cores) per node

Memory 15 GiB per node

Storage 160GB SSD storage per node

I/O Performance Moderate

Cluster Type

Number of Compute Nodes* Single Node clusters consist of a single node which performs both leader and compute functions.

Maximum 1

Minimum 1

Additional Configuration

Finally, there are some additional configuration items you can define, depending on the your network and security requirements. Each option is documented on the page.

8. Click the Continue button.

Services ▾ | Edit ▾

Amazon Redshift

Clusters

Snapshots

Security

Parameter Groups

Reserved Nodes

Events

✓

✓

○

▬

CLUSTER DETAILS

NODE CONFIGURATION

ADDITIONAL CONFIGURATION

REVIEW

Provide the optional additional configuration details below.

Cluster Parameter Group

A default parameter group will be associated with this cluster.

Encrypt Database

No ▾

Select Yes to encrypt all data within the cluster and in backups at a small cost to performance.

Use HSM

No ▾

You have not created any HSM Connections. You must create an HSM Connection to use HSM. You must also create at least one HSM Client Certificate.

Configure Networking Options:

Choose a VPC

Default VPC (vpc-94b90bf1) ▾

The identifier of the VPC in which you want to create your cluster

Cluster Subnet Group

default ▾

Selected Cluster Subnet Group may limit the choice of Availability Zones

Publicly Accessible

Yes ▾

Select Yes if you want the cluster to be accessible from the public internet. Select No if you want it to be accessible only from within your private VPC network

Choose a Public IP Address

No ▾

Select Yes if you want to select your own public IP address from a list of elastic IP (EIP) addresses that are already configured for your cluster's VPC. Select No if you want Amazon Redshift to provide an EIP for you instead.

Availability Zone

No Preference ▾

The EC2 Availability Zone that the cluster will be created in.

Optionally, associate your cluster with one or more security groups.

VPC Security Groups

launch-wizard-2 (sg-b4eec... ▾

default (sg-162f0c73)

launch-wizard-1 (sg-4eb79... ▾

launch-wizard-3 (sg-01182... ▾

List of VPC Security Groups to associate with this cluster.

Optionally, create a basic alarm for this cluster.

Create CloudWatch Alarm

☐ Yes ☒ No

Create a CloudWatch alarm to monitor the disk usage of your cluster.

Cancel

Previous

Continue

Review & Launch

On the final step you are able to review all the configuration options you've applied in the previous steps.

9. Click the **Launch Cluster** option to finish.

CLUSTER DETAILS

NODE CONFIGURATION

ADDITIONAL CONFIGURATION

REVIEW

You are about to launch a cluster with the following specifications:

Cluster Properties

These attributes specify the name of your cluster, what type of virtual hardware it will run on, how many nodes it will contain, and the availability zone in which it will be located.

Cluster Identifier: my-yf-aws-db-instance

Node Type: dw2.large

Number of Compute Nodes: 1 (leader and compute run on a single node)

Availability Zone: No Preference

Database Configuration

These properties specify the database name, port, and username you will use to connect to the database. The parameter group contains configuration values used by the database.

Database Name: yellowfindb

Database Port: 5439

Master User Name: master_user

A default parameter group will be created when the cluster is launched.

Security, Access, and Encryption

These settings control whether your cluster will be created in an existing VPC to allow for simpler integration with other AWS Services, and the security groups which define access rules to your cluster.

Virtual Private Cloud: vpc-94b90bf1

Cluster Subnet Group:

Publicly Accessible: Yes

Elastic IP: Not used

VPC Security Groups: sg-b4eec4d1

Encrypt Database: No

Use HSM: No

CloudWatch Alarms

CloudWatch alarms are used to notify if metrics for your cluster are within a certain threshold. All recipients under the SNS topic specified for your alarm will receive notifications once an alarm is triggered.

Basic alarms will not be created for this cluster.

⚠ Unless you are eligible for the free trial, you will start accruing charges as soon as your cluster is active.

Applicable charges:
The on-demand hourly rate for this cluster will be **\$0.25**, or **\$0.25/node**. If you have purchased reserved nodes in this region for this node type that are active, your costs will be discounted. Additional nodes will be billed at the on-demand rate.

If you are eligible for a free trial, you will receive 750 hours of free usage for each month of the trial, applied across all running dw2.large nodes across all regions. Regardless of when you start your trial, you will receive two full months of free usage. Once your trial expires or your usage exceeds 750 hours/month, you can shut down your cluster, avoiding any charges, or keep it running at our standard [On-Demand Rate](#).

For more information, see [Amazon Redshift Free Trial FAQ](#), [Amazon Redshift Pricing](#), and [Reserved Nodes Documentation](#).

Cancel

Previous

Launch Cluster

View Clusters

You will now be presented with a message informing you that your cluster is being created.

10. Click the **Close** button to return to a list of available clusters.

Services

Edit

Amazon Redshift

Clusters

Snapshots

Security

Parameter Groups

Reserved Nodes

Events

✓

Cluster **my-yf-aws-db-instance** is being created.


Note: Your cluster may take a few minutes to launch.

[View your cluster on the Clusters dashboard.](#)

Close

Cluster Name

Click on the name of your cluster in order to view its properties.





Cluster Details

Here you will be able to view all the details related to your cluster, including items important for Yellowfin use:

- Custer Name
- End Point
- Port
- Database Name
- Master User Name
- Public IP


Make a note of the above values.



Yellowfin


Login

- 1. Login to your Yellowfin instance.
- 2. Click on the **Administration** link in the main navigation bar and select **Admin Console**.



Add Data Source

- 3. Expand the **Data Sources** list in the centre panel of the Admin Console.
- 4. Click on the **Add** button to create a new connection.



Data Source Details

Here you will need to provide a range of information to tell Yellowfin what the connection is, and how to access the database.

- 5. Define the **Data Source Details**, **Security**, and **Connection Pool** options based on your requirements.

See [Understanding Source Connection Parameters](#) for more information.

Data Source Details	
Source Name:	Redshift
Description:	Redshift Connection
Max Rows Returned:	☐ Unlimited ☒ 10000
Writable:	☐ Yes ☒ No
Region:	Australia
Time Zone:	Sydney
Security	
Access Level:	☒ Unsecure ☐ Secure
Permissions:	☒ Broadcast ☒ Subscribe
Source Filters:	☐ Yes ☒ No
Primary Connection Pool Management	
Min Connections:	1 i
Max Connections:	5 i
Refresh Time:	3 hours i
Timeout:	180 seconds i
Use secondary pool:	☐ Yes ☒ No i
Availability	
Availability:	☒ Available. The connection to this Data Source was successful. i [Set Unavailable]
Save Cancel	

Connection Details

6. In order to connection to your Redshift cluster, define the following:

- **Connection Method:** JDBC
- **Authentication:** Standard Authentication
- **Database:** Amazon Redshift
- **Include schema in SQL:** (ticked)
- **Host:** *this is the **Public IP** or **Endpoint** information (either can be used, depending on the security settings you applied) found on your Cluster Details page earlier.*
- **Port:** *this is the **Port** found on your Cluster Details page earlier.*
- **Database:** *this is the Database Name found on your Cluster Details page earlier.*
- **User Name:** *this is the **Master User Name** found on your Cluster Details page earlier.*
- **Password:** *this is the password you entered while configuring your cluster earlier.*
- **Schema:** None

7. Test the connection and click **Save**.

Connection

Connection Method:

JDBC

Authentication:

Standard Authentication

Database:

Amazon Redshift

Include schema in SQL:

☒

Host:

yf.ccjapojk5x0n.a

Port:

5439

Database:

yf

User Name:

yf

Password:

Schema:

None

Connection Succeeded

Database:

PostgreSQL

Product Version:

8.0.2

Driver:

PostgreSQL Native Driver

Driver Version:

PostgreSQL 9.3 JDBC4 (build 1101)

Database Version:

8.0

Source Name:

Amazon Redshift

Click here to test the connection again.

Database Schema:

No tables found.

You are now ready to add data to your Cluster and build reports.

[top](#)