

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.

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.



Capacity Details

Current Node Type:

dw2.large

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

Platform:

64-bit

SSH Ingestion Settings

Cluster Public Key:

ssh-rsa
AAAAAB3NzaC1yc2EAAAADAQABAAQDnu7VqEBecf8xaiwRLP11
nuqWSoM2dnH4FE+5d1wTVJ6coA//vRodPIuf4XE27P4VvWGCf
ZsoFW609E8LHDq5G9kUpL0DuHbT3ZuxIMkBTsq8wV5NE8qh6IP
0dQWmF8UA0do04jpfQDkuFQhC5hcrR2CdQudPLPBTjQf7mkXE
uM0hecBF5axbj52vncF1FACQLJGgQ5iwoGAKpDdk2P4CP+F6F1
hdINDq+pdatk2Vqy2q+ht9hL3KBXGTQmxVcIc5k7EtIH9RhPm
Hsfn7NYeX/buCylpeNa1Km1T57Lho4jdLh50MyI1+0KNpgcA01
agir6RimzLL02xMfLncKd3 Amazon-Redshift


Node IP Addresses:

Node	Public IP	Private IP
Shared	54.66.138.222	172.31.22.15

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 ⓘ
Max Connections:	5 ⓘ
Refresh Time:	3 hours ⓘ
Timeout:	180 seconds ⓘ
Use secondary pool:	☐ Yes ☒ No ⓘ
Availability	
Availability:	☒ Available. The connection to this Data Source was successful. ⓘ [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](#)