

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.



Node Configuration

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

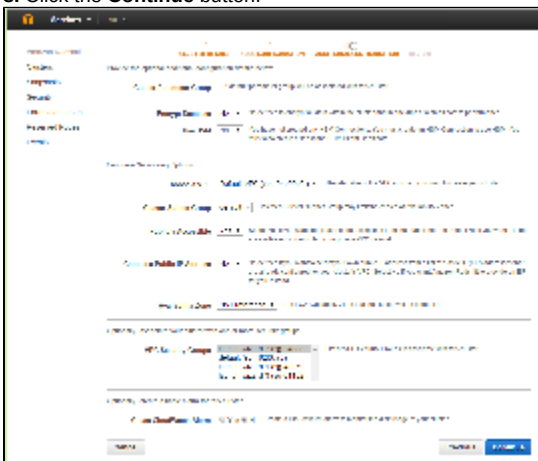
7. Define the **Type and **Number of Nodes**. Click **Continue**.**



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.



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.

CREATE REDSHIFT CLUSTER **PROVISION CLUSTER** NEXT CLUSTER NAME: **my-aws-ds-instance** PUBLIC

These properties specify the name of your cluster, the type of visual workflow engine or framework, roles, public access, and the availability zone in which it will be located.

Cluster Identifier: my-aws-ds-instance
 Name type: **Identifier**
 The name of your cluster. It must be unique across all clusters in the same region and availability zone.

Database Configuration
 These properties specify the database name, type, and parameters that will be set for the database. The parameter group controls configuration values used by the database.

Database Name: yellowfin
 Database type: **PostgreSQL**
 Master user: **yellowfin_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 encrypted VPC to allow for integration with other AWS services, and the security groups that define access rules to your cluster.

Virtual Private Cloud: **Use default**
 Choose a VPC to create your cluster in.
 Choose a security group to create your cluster in.
 Choose a parameter group to create your cluster in.
 Choose a role to create your cluster in.

CloudWatch Alarms
 CloudWatch alarms are used to notify if issues for your cluster arise when a certain threshold is reached. Amazon Redshift VPC logs enabled for your cluster will receive notifications once an error is reported.


Warning: If you are a light or free-tier user, you will start incurring charges as soon as your cluster is active.

Buttons: CANCEL, PROVISION, 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.



Cluster Name

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


Cluster	Cluster Status	DB Instance	In Maintenance	Recent Events
my-aws-ds-instance	creating	unknown	unknown	0

Cluster Details

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

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

Make a note of the above values.




Capacity Details		SSH Ingestion Settings	
Current Node Type	4x4248g	Cluster Public Key	
CPU	7 EC2 Compute Units (2 virtual cores) per node	Cluster Public Key	
Memory	15 GB per node	Cluster Public Key	
Storage	100GB S3 storage per node	Cluster Public Key	
I/O Performance	Medium	Cluster Public Key	
Platform	64-bit	Cluster Public Key	

Yellowfin

top

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

A screenshot showing three stacked configuration panels. The top panel, 'Data Source Details', includes fields for 'Source Name' (Redshift), 'Description' (Redshift Connection), 'Max Rows Returned' (10000), 'Writable' (No), 'Region' (Australia), and 'Time Zone' (Sydney). The middle panel, 'Security', shows 'Access Level' (Unsecure), 'Permissions' (Broadcast and Subscribe checked), and 'Source Filters' (No). The bottom panel, 'Primary Connection Pool Management', includes 'Min Connections' (1), 'Max Connections' (5), 'Refresh Time' (3 hours), 'Timeout' (180 seconds), and 'Use secondary pool' (No). Below these is an 'Availability' section showing a successful connection status. At the bottom are 'Save' and 'Cancel' buttons.

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**.



The screenshot displays the 'Connection' configuration window in the Amazon Redshift console. The configuration is as follows:

- Connection Method: JDBC
- Authentication: Standard Authentication
- Database: Amazon Redshift
- Include schema in SQL: ☐
- Host: yf.cqapc1sgdr1z
- Port: 5439
- Database: yf
- User Name: yf
- Password: [masked]
- Schema: None

Below the configuration fields, a 'Connection Succeeded' message is shown with a green checkmark icon. The details provided are:

- Database: PostgreSQL
- Product Version: 8.6.2
- Driver: PostgreSQL Native Driver
- Driver Version: PostgreSQL 8.3 JDBC4 (build 1101)
- Database Version: 8.6
- Source Name: Amazon Redshift

A link 'Click here to test the connection again.' is present. At the bottom, the 'Database Schema' section shows 'No tables found.'

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

[top](#)