

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

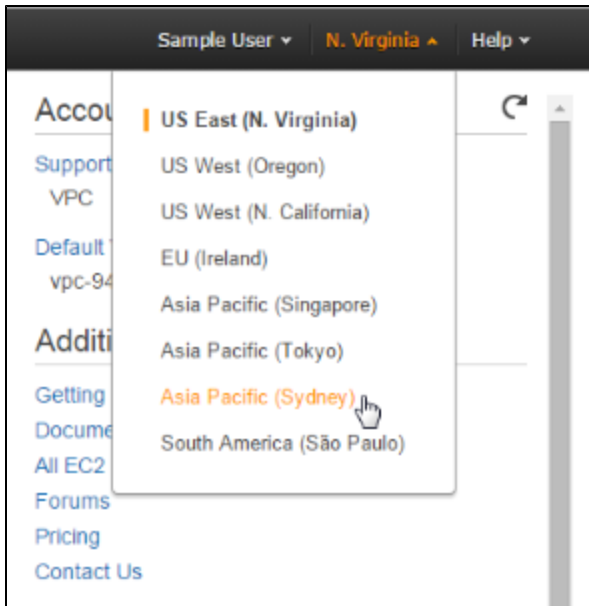
1. Log in to Amazon Web Services (AWS).
2. Open the **Services** menu at the top of the page, click on **Redshift** within the **Database** category.



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.

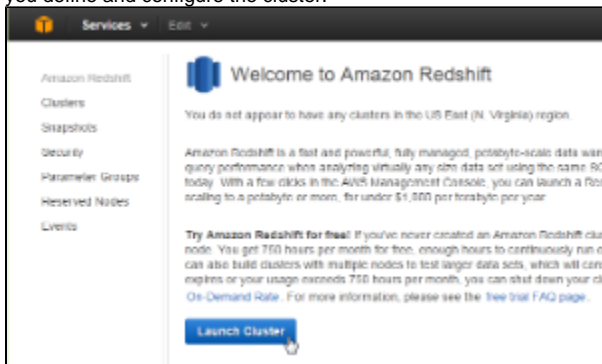
3. Click on the **Location** currently in use next to your name at the top right of the page.
4. Select the location closest to you.



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.



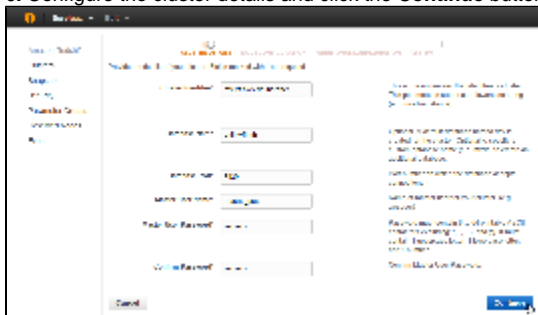
Cluster Details

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.

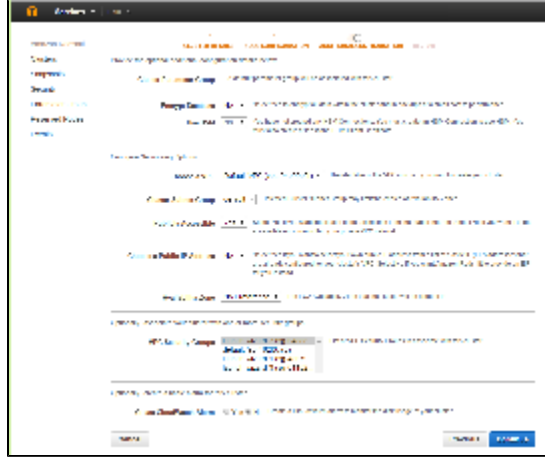
6. Configure the cluster details and click the **Continue** button.



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

[illegible]

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.

[illegible]

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

Cluster Properties

These properties specify the name of your cluster, the type of node group, the number of nodes, the number of nodes in each availability zone, and the availability zone in which it will be located.

Database Configuration

These properties specify the database name, user, and password. The database name must be unique within the cluster. The password you specify will be used to connect to the database.

Security, Access, and Credentials

These settings control whether your cluster will be created in an existing VPC or a new VPC, the number of subnets, and the security groups that will be created for your cluster.

CloudWatch Alarms

CloudWatch alarms are used to notify you when your cluster is in a certain state. All alarms are created by default. You can modify the alarms to suit your needs.

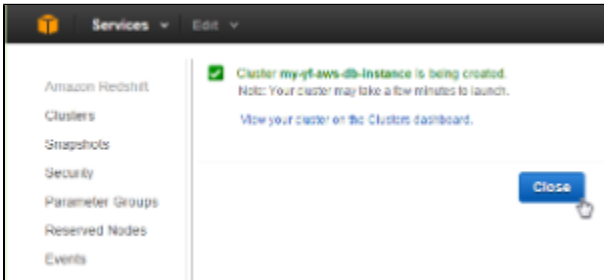
Warning: If you are eligible for the free trial, you will earn credits as soon as your cluster is active. If you are not eligible for the free trial, you will be charged for the cluster as soon as it is launched. The charges will be applied to your AWS account. For more information, see the Amazon Redshift pricing page.

Buttons: 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.



Cluster Name

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

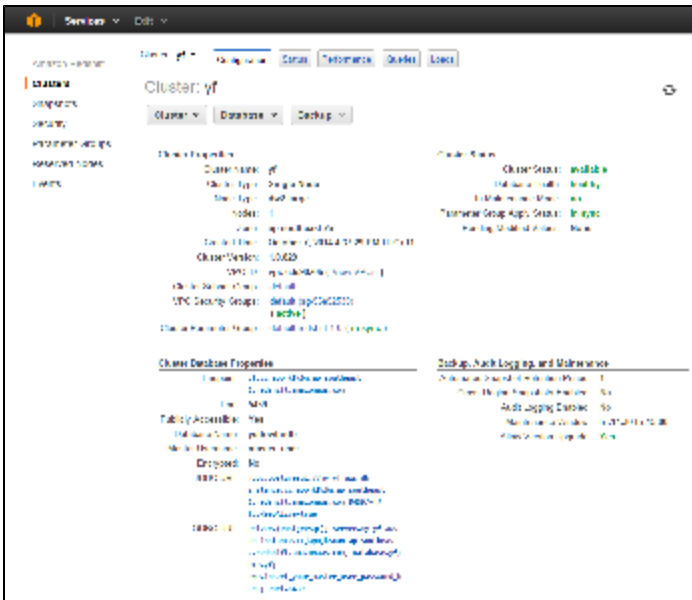


Cluster Details

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

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

Make a note of the above values.



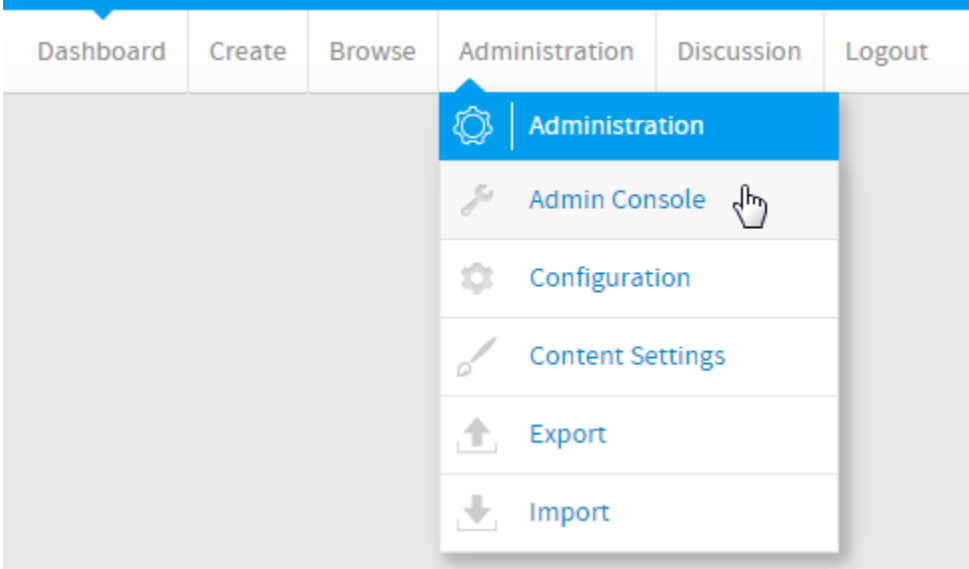
Capacity Details		SSD Ingestion Settings	
Current Node Type	4x4248g	Cluster Public Key	
CPU	7 EC2 Compute Units (2 virtual cores) per node	Cluster Private Key	
Memory	15 GB per node	Node IP Address:	
Storage	100GB S3 storage per node		
I/O Performance	Medium		
Platform	64-bit		

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 of the 'Data Source Details' configuration form. The form is divided into several sections. The 'Data Source Details' section includes fields for 'Source Name' (Redshift), 'Description' (Redshift Connection), 'Max Rows Returned' (Unlimited/10000), 'Writable' (Yes/No), 'Region' (Australia), and 'Time Zone' (Sydney). The 'Security' section includes 'Access Level' (Unsecure/Secure), 'Permissions' (Broadcast, Subscribe), and 'Source Filters' (Yes/No). The 'Primary Connection Pool Management' section includes 'Min Connections' (1), 'Max Connections' (5), 'Refresh Time' (3 hours), 'Timeout' (180 seconds), and 'Use secondary pool' (Yes/No). The 'Availability' section shows a status of 'Available' with a message 'The connection to this Data Source was successful.' and a '[Set Unavailable]' link. 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 shows the 'Test Connection' dialog box in the AWS Glue console. The 'Connection Method' is set to 'JDBC'. The 'Authentication' is 'Standard Authentication'. The 'Database' is 'Amazon Redshift'. The 'Include schema in SQL' checkbox is checked. The 'Host' is 'yf-cjapc1k5dncz'. The 'Port' is '5439'. The 'Database' is 'yf'. The 'User Name' is 'yf'. The 'Password' is masked with asterisks. The 'Schema' is 'None'. Below the form, a green checkmark indicates 'Connection Succeeded'. The details shown are: Database: PostgreSQL, Product Version: 8.6.2, Driver: PostgreSQL Native Driver, Driver Version: PostgreSQL 8.3 JDBC4 (build 1101), Database Version: 8.6, and Source Name: Amazon Redshift. A link 'Click here to test the connection again.' is provided. 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](#)