



# TD SYNnex AWS Redshift Step by Step Guide



**Below is a list of action items as part of the deployment process and post deployment recommendations to customize the Cloud Environment.**

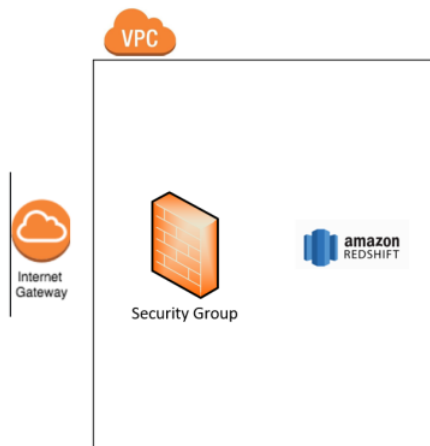
#### Technical Requirements – Customer Inputs

- Which availability zone will the solution be deployed to?
- What will you name the the first database after Redshift cluster is created?  
*(Must contain 1 to 64 alphanumeric characters and must contain only lowercase letters)*
- What size node would you like provisioned in the Redshift cluster?  
*(Current options: “ds2.xlarge”, “ds2.8xlarge”, “dc1.large”, “dc1.8xlarge”)*
- What username would you like associated with the master user account for the Redshift cluster?  
*(Must contain Must be 1 - 128 alphanumeric characters. The first character must be a letter and the username can't be “PUBLIC”)*
- What password would you like associated with the master user account for the Redshift cluster?  
*(Must be between 8 and 64 characters in length, contain at least one uppercase letter, at least one lowercase letter, and at least one number. No ‘ (single quote), “ (double quote), \, /, @, or spaces.)*

#### Benefits of AWS Redshift

AWS Redshift gives the ability to give new insights into your data at a fraction of the cost of other cloud data warehouses. With Redshift, you can query and combine exabytes of structured and semi-structured data across your data warehouse, operational database, and data lake using standard SQL. You can choose the size of your Redshift cluster based on your performance requirements. Redshift managed storage automatically scales your data warehouse storage capacity without you having to add and pay for additional compute instances. This helps you to scale and pay for storage and compute separately and get the optimal amount of storage and compute for diverse workloads.

## Architecture Diagram



## Resources Deployed

The following resources are deployed to build and configure this bundle.

- VPC
- Internet Gateway
- Security Group
- Redshift Cluster

## Cost Breakdown

You can choose the size of your Redshift cluster based on your performance requirements with prices starting as small as \$0.25 per hour and the ability to scale up to petabytes of data and thousands of concurrent users. In the example US availability zone, you will notice the ability to tailor your needs for the pricing model that is desired based on need for compute vs storage.

	vCPU	Memory	Addressable storage capacity	I/O	Price
<b>Dense Compute DC2</b>					
dc2.large	2	15 GiB	0.16TB SSD	0.60 GB/s	\$0.25 per Hour
dc2.8xlarge	32	244 GiB	2.56TB SSD	7.50 GB/s	\$4.80 per Hour
<b>Dense Storage DS2</b>					
ds2.xlarge	4	31 GiB	2TB HDD	0.40 GB/s	\$0.85 per Hour
ds2.8xlarge	36	244 GiB	16TB HDD	3.30 GB/s	\$6.80 per Hour