



AWS Redshift Assessment - Findings & Recommendations Report

Findings Summary

WAF Pillar: Security		High Risk
Definition	Security pillar focuses on the ability to protect a system’s data, clusters and applications while delivering business value through risk assessments and mitigation strategies.	
Findings	<ul style="list-style-type: none"> • Database encryption has not been enabled. • Cluster is publicly accessible. However, the VPC subnet group enabled for IP restriction. • Database is running on default port. • S3 VPC Endpoint is not enabled. 	
Priority	<p style="text-align: center;">Recommendations</p> <ul style="list-style-type: none"> • Your AWS Redshift cluster is at security risk. Though encryption is an optional setting in Amazon Redshift, AWS recommends that it is enabled for clusters that contain sensitive data. • Almost always, a Redshift cluster should be setup in a private subnet with publicly available option set to No and firewalled with a strict security group. Even when the cluster is configured to be in public subnet consciously, and marked publicly accessible, we still need to ensure the cluster is exposed to appropriate IPs and external parties via security group ingress rules. • Ensure that your AWS Redshift database clusters are not using their default endpoint port (i.e. 5439) in order to promote port obfuscation as an additional layer of Défense against non-targeted attack. • Ensure that S3 VPC Endpoint is enabled, Your AWS Redshift instances running in private subnets of a VPC will have controlled access to S3 buckets, objects, and API functions that are in the same region as the VPC. 	<p style="text-align: center;">Benefits/Outcome</p> <p>Improved security for data at rest and in transit</p> <p>Improved security for data at rest and in transit</p> <p>Improved security</p> <p>Improved security for S3</p>

<p>Action Items</p> <ul style="list-style-type: none"> • Encrypt the data using AWS KMS or CMS. • Enable the SSL Connections. <p>https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-db-encryption.html</p> <p>https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html</p>		
<p>WAF Pillar: Operational Excellence</p>		<p>Medium</p>
<p>Definition</p>	<p>Operational excellence pillar focuses on running efficient systems and gaining insights into the effectiveness of the operations to deliver business value. Operational process is continuous process to improve supporting process and procedures.</p>	

Findings	<ul style="list-style-type: none"> • Audit Logging is not enabled. • Cloud Watch Alerts Not Enabled. • Auto update of statistics is not enabled. 	
Priority	<p style="text-align: center;">Recommendations</p> <ul style="list-style-type: none"> • Ensure that your Amazon Redshift Audit Logging feature are enabled so it starts recoding database usage information such as queries performed and connection attempts which is useful for troubleshooting sessions. • Ensure that cloud watch logs are enabled, to track the metrics such as CPU Utilization, Network I/O, Disk Space used. • Ensure the Auto Sort, Auto Vacuum and Auto Analyse are enabled to efficiently sort the data in blocks, reclaim the deleted space and gather the table statistics. 	<p style="text-align: center;">Benefits/Outcome</p> <p>Better Insights</p> <p>Insights on Cluster Performance</p> <p>Better Maintenance</p>
<p>Action Items</p> <ul style="list-style-type: none"> • Enable Audit Logs - https://docs.aws.amazon.com/redshift/latest/mgmt/db-auditing.html • Enable Cloud watch Alerts - https://aws.amazon.com/premiumsupport/knowledge-center/logs-redshift-database-cluster/ • Turn on Auto Management - https://docs.aws.amazon.com/redshift/latest/dg/t_Reclaiming_storage_space202.html 		

WAF Pillar: Performance		Medium
Definition	Performance efficiency pillar focuses on the efficient use of computing resources to meet requirements and how to maintain that efficiency as demand changes and technologies evolve.	
Findings	<ul style="list-style-type: none"> Scan time for _____ is high. Concurrency Scaling not turned ON. 92 Columns found as Large width columns. Interleaved sort Keys Found. 	
Priority	Recommendations	Benefits/Outcome
	<ul style="list-style-type: none"> AWS Recommends to design using Time Series tables when the data is in TB scale and split into multiple year format and use the views to combine all tables for reporting purposes. This might help in improving the performance. Almost always, a Redshift cluster should be setup with concurrency scaling is enable. So that the reporting peak load is taken care. Please note Redshift provides 1hr of free scaling time as credit for every 24 hrs of cluster in use. Large width columns will lead to data spill to disk, avoid larger width columns to effectively manage the memory usage. Ensure that interleaved sort keys are used only when they are at-most required. Concurrency scaling will not work with interleaved tables. 	<p>Improved Performance</p> <p>Faster reporting performance</p> <p>Performance Improvement</p> <p>Performance Improvement</p>

<p>Action Items</p> <ul style="list-style-type: none"> • Implement Time Series for Huge tables : https://docs.aws.amazon.com/redshift/latest/dg/c_best-practices-time-series-tables.html • Enable Concurrency scaling: https://docs.aws.amazon.com/redshift/latest/dg/concurrency-scaling.html • Recheck the Large width columns https://docs.aws.amazon.com/redshift/latest/dg/c_best-practices-smallest-column-size.html • Re-evaluate Interleaved Sort Keys https://docs.aws.amazon.com/redshift/latest/dg/t_Sorting_data.html#t_Sorting_data-interleaved 	
<p>WAF Pillar: Cost</p>	
<p>Very High</p>	
Definition	Cost Optimization pillar includes the ability to run systems to deliver business value at the lowest price point.
Findings	<ul style="list-style-type: none"> • On-Demand Instances are used – DS2.Xlarge – 10 Nodes. • Consider RA3 evaluation instead of DS2 Instances. • Consider moving aged data to data lake and leverage spectrum.

Priority	Recommendations	Benefits/Outcome
	<ul style="list-style-type: none"> • Ensure that your Amazon Redshift On-demand clusters are converted to Reserved Instances for at least 1 yr. /3 yr. This will help in reducing cost by 50-70%. • We have noticed that RA3 is 2X more performant than DS2 cluster, would recommend evaluating. • Ensure the moving aged data to s3 to reduce the cluster size, leverage spectrum to query. 	<p>Lower cost</p> <p>Performance</p> <p>Lower Cost</p>
<p>Action Items</p> <ul style="list-style-type: none"> • Purchase Reserved Instances: https://aws.amazon.com/redshift/pricing/ • Evaluate RA3: https://docs.aws.amazon.com/redshift/latest/mgmt/working-with-clusters.html#rs-ra3-node-types • Leverage Redshift Spectrum: https://docs.aws.amazon.com/redshift/latest/dg/c-getting-started-using-spectrum.html 		

WAF Pillar: Reliability		Low
Definition	Reliability pillar encompasses the ability of a system to recover from infrastructure or service disruptions.	
Findings	<ul style="list-style-type: none"> • Cross Region backup – Not Enabled. • Temporary and Staging tables Found in Backup. 	

Priority	Recommendations	Benefits/Outcome
	<ul style="list-style-type: none"> • Ensure that cross region backup is enabled for outage / Disaster recovery. • Set the table backup flags to 0 for the temporary and staging tables for backups. <ul style="list-style-type: none"> - Helps in reducing the overall time for backup. - Reduce time for restoration. 	<p>Better Resilient</p> <p>Better Resilient</p>
<p>Action Items</p> <ul style="list-style-type: none"> • Enable cross region backup: https://aws.amazon.com/blogs/aws/automated-cross-region-snapshot-copy-for-amazon-redshift/ • Remove temporary tables for backup https://docs.aws.amazon.com/redshift/latest/dg/r_STV_TBL_PERM.html 		