



ANNEXURE

Features	Amazon Redshift	Google BigQuery
Service Models	Provisioned (PaaS) for full cluster control Serverless (SaaS) for automatic scaling and management	Fully serverless (SaaS) with no clusters and provisioning, completely managed by Google.
Deployment Models	Supports Provisioned Clusters (VPC-based) and Serverless within AWS Regions.	Supports regional or multi-regional deployments and cross-region queries via BigQuery Omni.
Regions & Availability	Available in most AWS Regions, supporting multi-AZ deployments and cross-region snapshots.	Stores data in regional or multi-regional datasets, ensuring automatic data protection within selected geography. No AZ management required.
Data Storage Options	For provisioned clusters, internal storage options include cluster or RMS. Serverless mode decouples compute and storage. External access to data in S3 is provided by Redshift Spectrum in both.	Data is stored in Colossus , a fully managed, columnar database. External access is possible via BigLake and GCS .
Scalability	Provisioned: Offers horizontal scaling by adding nodes and vertical scaling by upgrading nodes. Serverless: Automatically scales compute (RPUs) and storage (RMS).	Automatically scales computing and storage resources for handling data and concurrent queries.
Versioning & Updates	Engine updates are manual for provisioned and automatic for serverless. Additionally, data snapshots capture table/data state, with schema versioning managed externally.	Engine updates are fully automated. Versioning includes built-in table time travel for 7 days and snapshots for longer-term table versioning.
Communication Gateways	JDBC/ODBC, APIs/SDKs, PrivateLink & VPC endpoints for connectivity.	Connectivity is provided through JDBC/ODBC, REST APIs, and client libraries. Private Service Connect and VPC Service Controls offer secure connections.
Security, Backup, and Recovery	Security includes AES encryption at rest, SSL/TLS in transit, IAM roles, column/row-level security, etc. Automated snapshots and point-in-time recovery are available for backup.	Security includes default encryption at rest and in transit, IAM roles, row/column-level security, etc. Recovery is enhanced by multi-region replication, time travel for 7 days, and snapshots for long-term backup.
Pricing & Cost Optimization	Provisioned clusters use a pay-per-node-hour model with reserved node options for up to a 75% discount. Serverless charges per RPU-second with separate billing for storage by GB/month. Optimization is attained by RA3 nodes separating storage and compute. Using Redshift Spectrum allows queries on S3 data, reducing managed storage usage and overall cost.	On-demand queries are billed at approximately \$5 per terabyte of data scanned. Flat-rate slot reservations are available for predictable or high-volume workloads. Storage costs around \$0.02 per gigabyte per month, with lower rates for long-term data. Using partitioned or clustered tables and query caching helps minimize overall costs.
ELT vs ETL	Supports both ETL (pre-load transformation) and ELT (in-warehouse transformation).	It is optimized for ELT workflows, where raw data is loaded first and transformed within the warehouse. ETL is possible but used less frequently.
BI Tool Support	Integrates with Amazon QuickSight natively. It also supports third-party tools such as Tableau, Power BI, Qlik, and Looker.	Natively integrates with Looker Studio. But it also supports third-party BI tools like Tableau, Power BI, and Qlik.