Rockset Pricing Guide

Abstract

Rockset is a serverless search and analytics engine. This document describes Rockset's pricing.

March 2019

contact@rockset.com

[ROCKSET]

Rockset Pricing Guide

Rockset is a serverless search and analytics engine with full SQL compatibility that provides unprecedented speed and flexibility in building data-driven applications, microservices, data science notebooks and live dashboards.

Simple, usage-based pricing

Active GB

Your usage is calculated based on the Active GB: the total amount of data you have in Rockset. All active data, measured in uncompressed bytes, is automatically indexed, replicated, and distributed in the cloud. The associated compute required to ingest and query your data is included in your plan. All Pro and Enterprise accounts have scalable query compute, meaning you have the ability to purchase additional compute in case your application requires additional QPS.

Rockset charges based solely on Active GB; the additional storage requirements of indexes and replicas are not counted towards your usage. In contrast, other database systems must be provisioned with enough storage to account for indexes, backups, replicas, and other internal operations (such as periodically compacting the database to improve performance), and often require additional compute for tools that transform and load data

Pricing includes software, cloud hardware, and support.

Pricing

	Free	Basic	Pro	Enterprise
	$0	$6 per Active GB/month	$9 per Active GB/month	Contact Us
Billing	Free, 2 Active GB	No monthly minimum	Monthly minimum $2000	Contact Us
What's included	Unlimited queries, inserts and users Unlimited data source integrations Streaming ingest upto 50 KB/s Query compute Suitable for 10 QPS* 3-way indexes Replicas 1 online, 1 offline	Unlimited queries, inserts and users Unlimited data source integrations Streaming ingest upto 100 KB/s Query compute Suitable for 100 QPS* 3-way indexes Replicas 1 online, 1 offline	Unlimited queries, inserts and users Unlimited data source integrations Streaming ingest upto 100 KB/s Query compute Suitable for 100 QPS* 3-way indexes Replicas 2 online, 1 offline	Unlimited queries, inserts and users Unlimited data source integrations Streaming ingest upto 100 KB/s Query compute Suitable for 100 QPS* 3-way indexes Replicas 2 online, 1 offline
Performance & Reliability	Fixed query compute Fixed streaming ingest Multi-tenant	Fixed query compute Fixed streaming ingest Multi-tenant	Scalable query compute for additional QPS Scalable streaming ingest for additional KB/s Dedicated compute	Scalable query compute for additional QPS Scalable streaming ingest for additional KB/s Dedicated compute
Security	Standard encryption GSuite SSO	Standard encryption GSuite SSO	Standard encryption GSuite SSO	Advanced encryption with custom keys GSuite, Okta & Custom SSO Role-based access control

			Role-based access control	Two factor authentication AWS private link IP whitelisting
Support	Community Support	Ticket Based support Business hours coverage (M-F, 9-5pm PST)	4-Hour SEV1 Response Time Business hours coverage (M-F, 9-5pm PST)	1-Hour SEV1 Response Time 24x7 Coverage

*QPS calculated based on point lookup. QPS will vary based on the type of query.

Prices listed here are for US regions only. All prices are in USD, exclusive of taxes. Prices subject to change

Calculating active GB size

Active GB is measured as the total number of uncompressed bytes required to represent the data stored in the document. Total data stored in Rockset is determined by summing the size of all active data stored in Rockset. One way to think of this size is the total number of bytes returned by the “SELECT * FROM my_collection” query.

For most data formats such as JSON, CSVs, TSVs or data sets imported from other structured data management systems, the Active GB size will be comparable to the size of the input data set. For more compressed data formats such as Parquet, the Active GB size will usually be several times larger than the size of the input data set.

Scalable query compute

Each Rockset account is auto-provisioned with the necessary compute resources required for querying active data. The estimated 100 QPS that the auto-provisioned compute supports for Basic, Pro and Enterprise accounts is based on a point lookup `select _id from collection where _id=':id'` which was run over a 250GB collection of JSON records, where each record has a unique _id. Actual QPS will vary based on your data and your queries.

For your application's scale and performance needs you might choose to purchase additional query compute. Scalable query compute is a feature available in Pro and Enterprise accounts. The common reasons why you should consider purchasing more query compute are:

- **Faster queries**: If your workload involves queries that scan large volumes of data or high cardinality aggregations, then adding more compute will automatically increase the parallelism of all your query processing there by making them faster.

- **Higher throughput**: If your workload involves running a large number of concurrent queries, then adding more compute will provide more compute capacity to increase the overall throughput to run more concurrent queries.

- **Multiple workloads**: If you run multiple workloads and want to provision some head room to avoid resource contention.

- **Insufficient resources**: If you run complex queries against giant data sets,, you might run into "Resource Exhaustion" or "Query Deadline Exceeded" errors. Increasing compute could help resolve those issues when encountered.

EXAMPLE SCENARIOS

Application Development Example 1

Say you ingest 4GB JSON data every day, only retaining the past 3 months' data as active documents in Rockset, resulting in a total of 360GB of data.

Assuming you choose the Pro edition:

Cost of 360 Active GB at $9 per Active GB/month = 360*$9= $3,240
Total monthly bill = $3,240 per month

Application Development Example 2

Say you ingest 2GB JSON every day, only retaining the past 3 months' data as active documents in Rockset, resulting in a total of 180GB of data.

Assuming you choose the Pro edition:

Cost of 180 Active GB at $9 per Active GB/month = 180*$9= $1,620
Total monthly bill = $2,000 per month (minimum for Pro edition)

Data Science Example 1

Say you ingest a 500GB CSV data set for analysis.

Assuming you choose the Basic edition:

Cost of 500 Active GB at $6 per Active GB/month = 500*$6 = $3,000
Total monthly bill = $3,000 per month

Data Science Example 2

Say you ingest a 50GB Parquet data set for analysis. Rockset prices based on uncompressed data size, which is 400GB in this case because Parquet supports very efficient compression.

Assuming you choose the Basic edition:

Cost of 400 Active GB at $6 per Active GB/month = 400*$6 = $2,400
Total monthly bill = $2,400 per month