


Building your Data Stack

Evaluating the cost of Cube Cloud
against self-hosting Cube OSS



While it may seem like self-hosting Cube OSS (**Open Source Software**) is the obvious choice (It's free after all, right?), it's important to recognize the cost of infrastructure, resources, and maintenance that comes with this "free" choice. Here we attempt to break down those costs to help with your evaluation.

Three example deployment scenarios for Cube:

1

Just Testing
(small)

2

**Going into
Production**
(medium)

3

**Scaling to
Enterprise**
(large)

But First: Why is Total Cost of Ownership Important?

The Total Cost of Ownership (TCO) is a pivotal metric in any technology decision-making process for several compelling reasons. First, it's crucial to dispel the misconception that self-hosting and internally supporting Open Source Software is entirely free. While self-hosting OSS may not require upfront licensing fees, a TCO analysis accounts for all associated costs, including deployment, maintenance, and support, to give a more comprehensive view of what you should expect to spend.

Additionally, a TCO analysis captures long-term savings or expenses by looking at more than just the initial cost of deployment. Oftentimes, streamlined operations can translate into enhanced efficiency and reduced overhead over a longer time horizon.

Finally, most OSS products are lacking important features that can deliver productivity enhancements to the business – further enhancing the benefit of a TCO calculation.

Let's take a closer look at the three deployment scenarios to understand how Cube Cloud's TCO stacks up against self-hosting Cube OSS.



When you're initially exploring and developing with Cube, your query volume is low, your resource requirements are low and you are focused on testing functionality and building a data model that works. Your front-end team will want to test integrations with their tooling and connections from their data apps. It's also common to utilize a few pre-aggregations at this phase to improve performance for sluggish queries.

On self-hosted Cube OSS, you may choose to deploy a single container locally or put it on a small VM so that you can collaborate and your team can connect to it. **That would typically take a team several hours or days to learn the ins and outs of Cube configuration, then spin up and manage the instance.**

On Cube Cloud, you can sign-in to the free tier in seconds which most likely has what you need for an evaluation. Not only does this save you 100% of hosting, setup, and maintenance costs, but you can also get started much faster. **When you are ready to go into production, it's just a click to get a production cluster spun up on the Starter plan.** And, even on the Starter plan paying monthly, the TCO would still be less – about half as much as Cube OSS.

Note from the Cube Team:

If you want to evaluate additional functionality from our higher tier plans, just [contact sales](#) and we'll set you up with a limited time trial.

Cube OSS	Cube Cloud Starter
Assumptions: <ul style="list-style-type: none">• Single instance of Cube• Testing small pre-aggregations• Two vCPU, 4GB RAM	Assumptions: <ul style="list-style-type: none">• Starter Production Cluster• One Development Instance• Up to 50k requests/day• Online 40 hrs/wk (auto-suspend enabled)
\$1,200 Setup: 8 Hrs @ \$150/hr	
\$753 Hosting: AWS EC2 on-demand	
\$600 Maintenance: 1 Hr/Q @ \$150/hr	
\$2,553 Yr 1 TCO	\$1,188 On Starter: \$99/month minimum
	\$1,188 Yr 1 TCO

53% savings on Cube Cloud: \$1,365

Going to Production



Scenario 2

Launching your first Cube deployment is an exciting and crucial milestone that must be executed correctly to establish trust with your users. You'll need a stable Cube deployment with high-availability/redundancy for the API and Cube Store worker nodes. You'll also need a process and tooling for testing and pushing changes to production without taking the deployment offline (CI/CD and blue/green). A Cube Store cluster is necessary to run pre-aggregations and it needs to be sized large enough to execute queries against your largest rollups. The hosting cost is higher than a Cube cluster that only does query pass-through.

On self-hosted Cube OSS, you may set up a Kubernetes cluster with 2 deployments – one for Blue and one for Green. Your testing and

development work happens on whichever deployment is not serving the production workload. Your CI/CD orchestration pipeline would need to deploy new commits to the non-prod deployment, build or refresh any pre-aggregations in Cube Store that need to be updated, then cutover the traffic routing to the newly updated cluster to complete the code deployment. You likely also need small dev deployments to work collaboratively on feature and dev branches, which requires additional lightweight Cube instances.

On Cube Cloud, you have built-in redundancy and there's no Kubernetes to setup or manage. CI/CD orchestration runs transparently while you simply merge/commit changes to the main branch. You also have access to chat and email support to help you move faster and more confidently.

Self-Hosted Cube

Assumptions:

- Using recommended prod specs for a distributed production cluster with Cube Store
- 4 nodes: 8vCPU, 32GB RAM each
- 2 dev nodes: 2 vCPU, 4GB RAM each
- CI/CD or Blue/green deployment process

\$14,795 Hosting: AWS EC2 on-demand

\$40,000 Setup: 200 Hrs @ \$200/Hr

\$19,200 Maintenance: 12 Hrs/Q @ \$200/Hr

\$6,400 Support: 4 Hrs/Q @ \$200/Hr

\$95,191 2 Yr TCO

Yr 1: \$67,595 Yr 2: \$27,595

Cube Cloud

Assumptions:

- Cube Cloud Premium plan (No VPC)
- Single production cluster running 24/7
- 4 Cube Store Workers (CSWs) active 50%
- 2 Dev Instances @ 40 hrs/wk
- Chat Support
- 99.95% uptime SLA

\$13,988 On Premium with 55,954 CCUs @ \$0.25

\$27,977 Yr 2 TCO

Yr 1/2: \$13,988

71% savings on Cube Cloud: \$67,214

Production at Enterprise Scale



Things are going well and Semantic Layer adoption is taking off!

Cube has become mission-critical to your data teams and products. Data residency requirements also dictate the need for 2 independent regional deployments. Fortunately, Cube can scale to meet high-traffic demands, insulating your data warehouse from costly overruns and providing consistent data models, access control, cache, and APIs from multiple data sources to many data consumer tools.

With this increased demand comes increased scrutiny. Quick resolution of any data pipeline issues is critical. Infrastructure costs must be both predictable and minimized, but performance and availability must remain high. The stakes are high, and Cube is up to the challenge.

On Cube Cloud, your deployments are ready to go on day one with all the scaling up functionality built into Cube Cloud already.

Clusters can auto-scale and auto-suspend with just a click, and the billing dashboard provides immediate clarity on consumption tracking. Request monitoring is built in so you can see the performance of your API responses as they stream in, and logs can be directed to your sink of choice. Cube Cloud comes with a performant custom-built APM tailored to Cube development needs that's simply not available in OSS. This APM implementation helps Cube Cloud customers dramatically improve their development velocity. You have the full support of Cube's infrastructure management team and a dedicated Customer Success Manager to keep things running smoothly.

On self-hosted Cube OSS, while the baseline functionality scales up well, additional tooling is required to best serve the mission-critical use case. These features would need to be developed and managed alongside a self-hosted Cube deployment:

- **CI/CD orchestration pipeline**
- **Application Performance Monitoring stack**
- **Cluster auto-scaling and auto-recovery**
- **Responsive, Cube-knowledgeable DevOps support**

Production at Enterprise Scale



Scenario 3

Self-Hosted Cube

Assumptions:

- Using recommended prod specs for a distributed production cluster
- 2 deployments autoscaling 10 to 90 vCPUs and 22 to 182 GB of RAM (avg 46 vCPU and 94GB RAM)
- 6 dev nodes: 2 vCPU, 4GB RAM each
- Feature build-outs:
 - Autoscaling on stateless and stateful K8s clusters
 - Application Performance Monitoring stack
 - CI/CD or Blue/green deployment process w/pre-aggregation warmup
- 6 month feature build-out delay

\$41,103 Hosting: AWS EKS Fargate on-demand \$82,053 without auto-scaling

\$6,728 Monitoring service: self-hosted Grafana

\$240,000 Setup: CI/CD, Monitoring, Auto-scaling 1,200 hours @ \$200/hr

\$64,000 Maintenance: 80 Hrs/Q @ \$200/Hr

\$9,600 Support: 12 Hrs/Q @ \$200/Hr

\$604,292 3 Yr TCO

Yr 1: \$361,431 Yr 2/3: \$121,431

Cube Cloud

Assumptions:

- Cube Cloud Enterprise plan (VPC)
- 2 autoscaling Production Clusters
- 2-10 API instances: Avg 4 API Inst
- 16 CSWs: Avg 8 CSWs
- 6 Dev Instances at 55 hrs/wk each
- All enterprise features including autoscaling, monitoring, CI/CD
- Dedicated CSM and Chat Support
- 99.99% uptime SLA

\$103,680

On Enterprise with 259,200 CCUs @ \$0.40

\$311,040 3 Yr TCO

Yr 1/2/3: \$103,680

49% savings on Cube Cloud: \$293,252

Plus 6 months of build

* Self-hosting costs are estimated using EC2 Linux on-demand rates for Scenario 1 and 2 (non-autoscaling). Scenario 3 uses AWS Fargate on-demand pricing for the average cluster size.

** Hourly services rates (\$150/hr for small, \$200/hr for medium/large) taken from an informal poll of Devops Consultants.

Comparing Apples to Apple Pie



Below is a brief comparison of how we turn the raw ingredient of Cube OSS (like apples) into the feature rich, production-ready Cube Cloud (akin to a delectable apple pie).

Category	Feature	Available on Cube OSS	Available on a Cube Cloud Plan
Fully-Managed & Optimized Deployments	Fully managed runtime, Cube Store, Caching, & APIs	No	Yes
	Host on AWS, GCP, Azure	Yes	Yes
	Host on-premise	Yes	No
	Host on a VPC or BYOC	Yes	Yes
	Development instances and production clusters	Yes	Yes
	Multi-cluster deployments	No	Yes
	Distributed file storage failovers and consistency checking	No	Yes
Developer Toolkit: Collaboration and Observability	Cloud IDE	No	Yes
	Management console	No	Yes
	Auto-suspend, auto-scaling	No	Yes
	Monitoring integrations (logs export)	No	Yes
	Performance monitoring & query Tracing	No	Yes
	Automatic testing endpoints	No	Yes
Security	SSO (SAML 2.0), Role-based Access Control	No	Yes
	Cube Store encryption at-resy	No	Yes
	SSL connection to SQL API	No	Yes
Compliance	SOC2 Type II, HIPAA, GDPR	No	Yes
	Annual penetration testing	No	Yes
Integrations	Semantic Layer Sync	No	Yes
	Streaming pre-aggregations: ksqlDB	No	Yes
	Integration with monitoring solutions	No	Yes
	Cube Store support for Azure Blob Storage	No	Yes
Support	Online resources (Webinars, Docs, Slack)	Yes	Yes
	Chat/email w/support engineers	No	Starter+
	Dedicated CSM	No	Enterprise+
	Support availability- weekdays 8am-8pm ET	No	Yes



What If Your Data Must Stay on Your Network?

If your contractual obligation or corporate preference is to retain your data in your own infrastructure, then we have you covered with "Bring Your Own Cloud" (BYOC).

Before jumping straight to this option, however, most customers in this scenario choose to list Cube as a subprocessor and can still comply with data residency requirements

(e.g. European data can't leave Europe) with dedicated regional deployments and VPC Peering.

BYOC is simply Cube Cloud deployed in a public cloud account you own and control. The Cube team still manages the infrastructure and supports the deployment, and you get the full benefit of Cube Cloud without any data being stored off premises.



If you are interested in exploring BYOC options and understanding all the ways we can deploy Cube Cloud, please [contact sales](#) to learn more.



When we first found Cube, we considered hosting it internally. However, we realized doing so would require quite the overhead from our DevOps team. If you plan on using Cube but don't want to (or can't) increase your infrastructure OPEX, then go for Cube Cloud – you won't regret it.



Alessandro Lollo
Senior Data Engineer at Cloud Academy

Start Building Your Data Stack

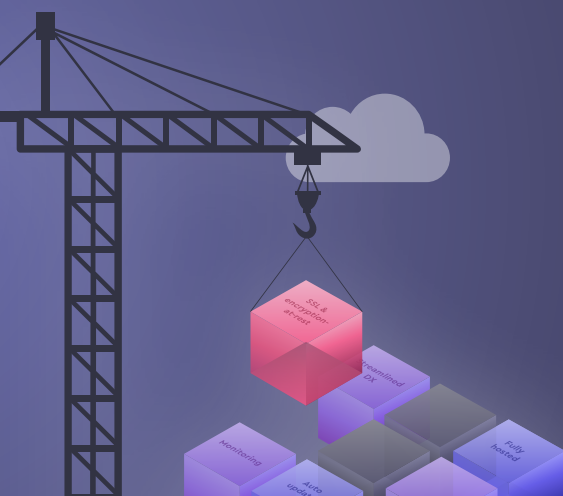


Regardless of stage or size of deployment, running your Semantic Layer on Cube Cloud costs less than Self-Hosting Cube OSS

For dev and test workloads, small production use cases, and enterprise scale deployments, Cube Cloud consistently saves you both time and money. Cube Cloud is supported, fully-managed Cube plus key features to accelerate your development workflow while providing a predictable cost structure. Instead of building out your own self-hosted deployment, Cube Cloud is ready to power your production workload immediately.

We fully believe in the power of open source software and are proud to provide Cube as an open source product. We've also created Cube Cloud to provide an even better value to an organization looking to increase speed-to-value, reduce cost overall, and reduce overhead maintaining bespoke integrations. While these are just 3 common deployment scenarios, there are many other configurations that could be a better fit for your organization. We'd be happy to chat with you about your use case and show you how Cube Cloud is so much more capable and even more affordable than self-hosting Cube OSS.

Scenarios	Self-Hosted Cube OSS	Cube Cloud	Savings
Just Testing: 1 Yr TCO	\$2,553	\$1,188	53%
Going into Production: 2 Yr TCO	\$95,191	\$27,977	72%
Production at Enterprise Scale: 3 Yr TCO	\$604,292	\$311,040	49%



Visit cube.dev for more information or to contact your account manager.

TRY CUBE FOR FREE