Article on thatDot.com

Quine vs Quine Enterprise

Quine combines the depth of graph analytics with the speed of event stream processing, enabling you to detect complex patterns in real-time across massive data streams. For many teams, the open-source edition provides everything needed to build sophisticated real-time pattern detection systems.

When You Should Consider Upgrading

Enterprise adds operational features for teams running at scale or in production environments requiring enhanced security, governance, and support:

Horizontal Scaling

Scale effortlessly by adding cluster members to handle massive data volumes and query loads. No architectural changes needed, just add more machines as your data grows.

High Availability

Keep your mission-critical pipelines running with hot standby nodes and instant failover when hardware and networks fail. Zero data loss, maximum uptime.

Enterprise Controls

Maintain security and governance at scale with audit logging, multi-tenant graph isolation, and role-based access control with your existing identity providers

Professional Support

Get direct access to thatDot's engineering experts for resolving critical issues, guidance on data modeling, query optimization and complex use cases, and roadmap input.

…along with powerful features such as extensible user-defined functions, and industry-first supernode optimization for handling billions of edges per node. Explore the detailed feature comparison below for the complete picture.

Your Work Transfers Seamlessly

Trying Enterprise is straightforward: your existing OSS recipes and API integrations work unchanged. If Enterprise features don't match your current needs, you can return to your OSS workflows seamlessly.

Feature Comparison

Feature	Open Source	Enterprise
GRAPH ENGINE & EXTENSIBILITY		
Unlimited Graph Size Both versions can handle an infinite stream of data and build a graph with any number of nodes and edges, limited only by the storage capacity. Gone are the days of trying to keep your node-counts low: you can analyze the complete history of your data, no matter how large, without needing to downsample or archive.		
Historical Graph Versioning All changes to nodes and edges are timestamped and saved, allowing for queries against the graph's state at any point in the past. This allows for powerful "as-of" analysis, letting you “travel back in time” to audit changes or replay event sequences exactly as they occurred.		
Cypher Query Language Hit the ground running by leveraging a widely-adopted, expressive language, making it easier for your team to build and maintain complex data pipelines. Users interact with the graph using the declarative Cypher query language for ingest, analysis, and standing queries.		
Customizable Storage Layer with Data Replication Deploy Quine anywhere, from a local development machine to a large-scale, production cluster, without changing your application logic. Quine supports various storage backends like RocksDB and Cassandra to durably store graph data. Enterprise customers have access to additional storage options like Clickhouse, and can also take advantage of expert consultation for custom implementations.		
Flexible Output Destinations Trigger downstream workflows, alerting your team, or feeding other real-time systems with high-value, processed events. Standing query results can be sent to multiple downstream destinations, including streaming pipelines (Kafka, Kinesis, etc), static collections (output to files), interactive webhooks, and even human notifications (Slack Integration). This ensures Quine fits seamlessly into your existing architecture,		
Streaming & Static Data Ingest Data can be loaded in real-time from streams like Kafka or from static files like CSV and JSON. This allows you to enrich live event streams with historical context from batch files, creating a comprehensive, up-to-the-moment view of your data.		
User Defined Functions Extend Quine's Cypher engine with custom logic written in any JVM-compatible language like Java or Scala. This empowers you to embed complex, proprietary business logic or specialized algorithms directly into your queries, tackling unique challenges that go beyond standard Cypher.		

Namespacing Host multiple, isolated graphs within a single cluster to support safe multi-tenancy or sandbox different use cases. This enables you to serve multiple teams or use cases on the same hardware footprint.		
Supernode Mitigation Supernodes are the Achilles heel of every graph system. If you work with real-world data, you will definitely have supernodes. While this has been a deal-breaking for all other graph systems, Quine Enterprise includes revolutionary new technology that solves a critical bottleneck that paralyzes traditional graph systems, allowing you to work with *billions* of edges per node.		
SCALE & RESILIENCE		
Clustering Run as a multi-member cluster to scale horizontally, increasing processing throughput and capacity. With Quine Enterprise you can handle massive ingest volumes and query loads by simply adding more machines, ensuring your system grows with your data.		
Enterprise-Grade Resilience Deploy standby cluster members (hot spares) that automatically take over for failed nodes to minimize downtime, and automatically detect failed members and network partitions to maintain operational stability. Automated cluster management in Quine Enterprise ensures no data is lost even if cluster members die unexpectedly. This provides high availability for mission-critical applications, keeping your data pipelines running 24/7.		
Kubernetes Helm Chart Deployment Deploy and manage Quine clusters in Kubernetes environments using production-ready Helm charts. Includes rolling upgrades/downgrades, resource management, and integration with Cassandra-compatible databases. Helm charts provide declarative configuration management, making it easy to deploy consistent environments across development, staging, and production Kubernetes clusters.		
Integrated Monitoring Stack Built-in integration with enterprise monitoring tools including Prometheus metrics export and pre-configured Grafana dashboards. Quine enables you to set up Grafana backed by InfluxDB to monitor Quine instances, and exposes JVM and application metrics via Prometheus JMX exporter with Kubernetes service discovery annotations. Use the included Grafana dashboards for monitoring ingest rates, node/edge counts, persistor latency, cluster health, and memory usage.		
SECURITY & GOVERNANCE		
Role-based Access Control Connect to your enterprise directory service to provide fine-grained access control. Integrate with standard identity providers (like Okta, Microsoft, KeyCloak, AWS Cognito, and more) for secure user authentication. This allows you to enforce the principle of least privilege, ensuring users and applications only have access to the data they need.		
Audit Logging Provides a tamper-evident API endpoint to get a digest of configurations for auditing purposes.		

SUPPORT & PARTNERSHIP		
Discord Community Support Users can ask questions and get help from a helpful and supportive community that includes thatDot developers on the public Discord server.		
Enterprise Support Receive professional support from the thatDot team. Ensure that when critical production issues arise, you have an expert on hand to resolve them quickly.		
Roadmap Priority Turn your subscription into a partnership, allowing your organization's unique needs to directly influence the product's evolution.		
Expert Consultation Gain access to thatDot's experts for guidance on data modeling, query optimization, and complex use cases. Accelerate your time-to-value and build more robust solutions by leveraging the deep expertise of the engineers who create the product.		

Community-Driven Development

Quine Open Source remains in active development with regular releases and community contributions. Our Enterprise users directly support continued OSS innovation through feedback, testing, and development resources. We're grateful for their partnership.

Quine OSS is developed openly on GitHub.

Next Steps

- **Haven't tried Quine OSS yet?** Download it here and join our Discord community for support from engineers and community members.
- **Considering Enterprise?** Get started with a trial, or request a consultation demo to see Enterprise features in action and discuss your use case with our solution architects.
- **Questions about which edition fits your needs?** We're happy to help you evaluate options based on your specific requirements and scale.

Quine.io Content

Notes

- *Recommend **not** adding anything related to Enterprise in the **Getting Started** section.*
- *For everything below, link to Compare Editions here: https://www.thatdot.com/quine-open-source-vs-enterprise/*

Site wide

Footer

Add

"Running Quine in production at scale? Compare Quine Open Source vs Enterprise."

or

"Running Quine in production at scale? Quine Enterprise provides clustering, hot-standby failovers, role-based access, and production support. Compare editions."

Specific pages

Add a note/tip block at the end:

Operational Considerations

https://quine.io/core-concepts/operational-considerations/

Quine Enterprise adds clustering, failover, and Kubernetes Helm deployment. Compare editions.

Persistors

https://quine.io/learn/persistors/

Quine Enterprise adds operational controls and high availability clustering around Cassandra-backed deployments. Compare editions.

Webserver config

https://quine.io/reference/config/quine-webserver-advertise/

Quine Enterprise adds role based access control and other security & governance features. Compare editions.

Exploration UI

For multi-tenant role based access control, see Quine Enterprise. Compare editions

Ingest Streams

For even higher ingest rates and high availability, see Enterprise clustering. Compare editions.

Quine Github Readme

At the end, after “Contributing to Quine” heading:

Quine Enterprise

The team behind Quine also provides Quine Enterprise, which adds operational features like clustering, hot-standby failovers, role-based access, and production support. Compare editions here.

Other copy you can use

General, longer

Production at scale

For horizontally scaled, highly available deployments, Quine Enterprise adds operational features like clustering, hot-standby failovers, role-based access, and production support. Your OSS projects run unchanged. Compare Quine Open Source and Quine Enterprise.

Shorter

General: For production scale, HA, and RBAC, see Quine Enterprise. Compare editions.
Scale: Need clustering and failover for 24/7 pipelines? Compare editions.
Security/Governance: Need SSO/RBAC and audit logging? Compare editions.
Support: Need expert support and guidance in production? Compare editions.