

Because one Trino cluster is not enough

Speakers


Manfred Moser

Trino and Trino Gateway maintainer and expert

Co-author of Trino: The Definitive Guide

Will Morrison

Trino and Trino Gateway maintainer and expert



Starburst

The Trino Company

Both work with Trino and Trino Gateway at

What is Trino?




Trino is a ludicrously fast, **open source, SQL query engine** designed to query **large data sets** from one or more **disparate data sources.**



Why Trino?



- Many clients, many data sources, one analytics platform
- Great support for all lakehouse platforms
- Powerful SQL support
- Performance, performance, performance
- Run on prem or cloud
- Manage yourself on prem or in cloud
- Use cloud provider offering
- Or use managed product like  **Starburst** Galaxy

Clients for data engineers and others



All the data sources





Why multiple Trino clusters?

- On-prem **and** on cloud
- Different data sources in different locations
- Different groups as owners
- Different cluster purpose - production, staging, test, development
- Different workloads
 - Analytics - mostly many small queries, that are okay to fail
 - Batch data processing - fewer, large queries, that can not fail
 - Number of concurrent users and variety of queries
 - Predictability of workload - users vs automated tasks
- Workload isolation
- Default and fault-tolerant execution modes of Trino



Multiple Trino clusters

Are a **reality**, and a **pain** for users.

Every cluster has a different URL.


Requires separate configuration in client applications.

Side note - Trino client protocol used HTTP requests.

What is Trino Gateway?



Trino Gateway is a load balancer, proxy server, and configurable routing gateway for multiple Trino clusters.





Quick project overview

- Evolved from lyft/presto-gateway
- Contributed by Bloomberg, and indirectly Lyft in July 2023
- Presented at [Trino Summit 2023](#)
- Active team of subproject maintainers and contributors
- Significant adoption by large Trino users
- Public developer sync every two weeks
- Over 400 commits of varying size
- Eleven releases so far



Use case: High availability

The simplest deployment mode:

- Two or more identical clusters.
- Random routing is built-in.
- No downtime, if a cluster goes down. (apart from queries it took with it)
- Even distribution among clusters.



Use case: No downtime upgrades

Extension of high availability setup

- Blue/Green deployment
 - Even distribution
 - Blue gets drained of workload, Green manages all workload
 - Blue gets upgraded and tested
 - Workload is shifted to Blue
 - Green is upgraded
 - Back to even distribution
- Can also use other upgrade methods like canary.



Use case: Smart routing

- Different users and queries hit different clusters
- Load balancing:
 - Randomly even distribution (stochastic)
 - Based on Trino cluster load (query count)
- Logical routing:
 - File-based configuration of rules using query, user, and other information in the request
 - External service with custom logic



Get started

- With JAR file anywhere
- Using container with Helm on k8s
- One or more stateless Trino Gateway nodes
- Optionally behind a pure load balancer
- Requires a RDBMS - MySQL or PostgreSQL
- Quickstart guide or Java runner for local testing
- Configuration in yaml file



Demo time

Final words



“If you use multiple Trino clusters, you want to look at Trino Gateway.”

The project has come a long way already, ...
but there is an even longer road ahead. We got big ideas.

Website at <https://trinodb.github.io/trino-gateway/>

Source code at <https://github.com/trinodb/trino-gateway>

Join us on **trino-gateway** and **trino-gateway-dev** on Trino slack.