Opening the

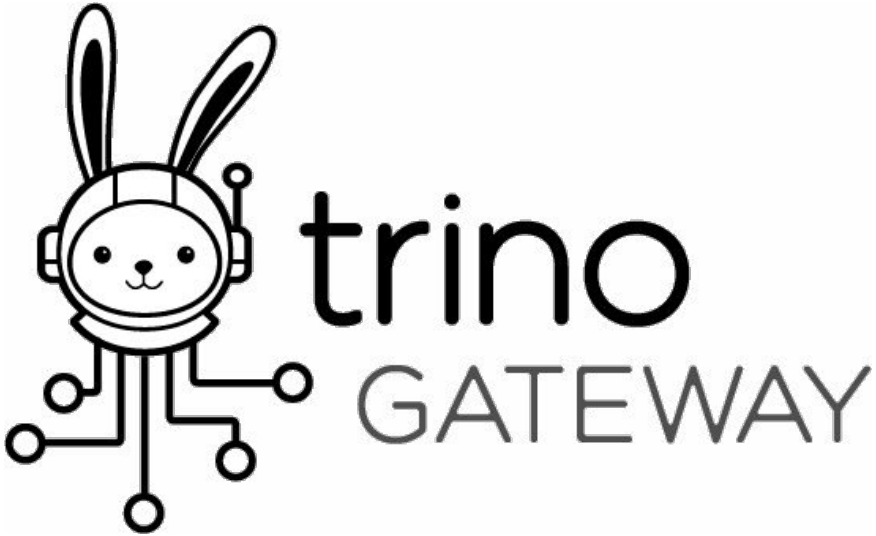

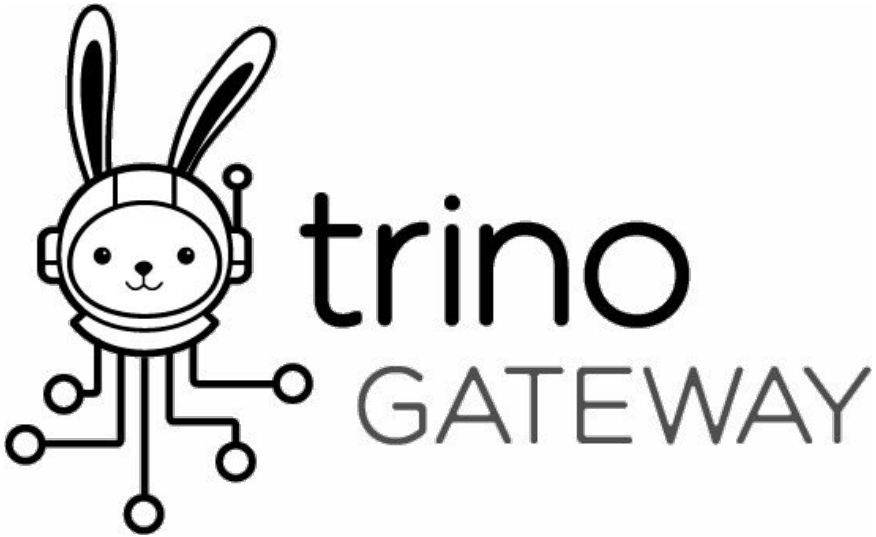

A project update at Trino Summit Dec 2024

Speakers

Trino Gateway subproject maintainers

- Jaeho Yoo
- Manfred Moser
- Star Poon
- Vishal Jadhav
- Will Morrison

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 uses HTTP.

What is Trino Gateway?

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

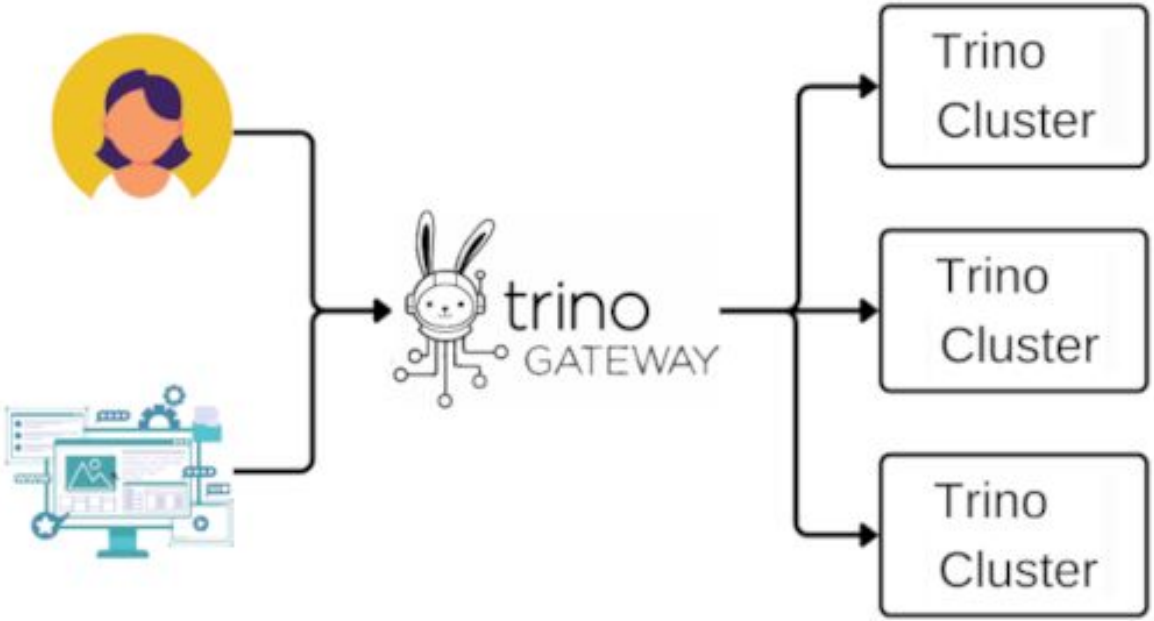

Quick project overview

- Evolved from lyft/presto-gateway
- Contributed by Bloomberg, and indirectly Lyft in July 2023
- First presented at Trino Summit 2023
- Active team of subproject maintainers and contributors
- Public developer sync **every two weeks**
- Over **500** commits of varying size
- **Thirteen** releases so far

Who uses Trino Gateway

Or is migrating from a legacy Presto Gateway fork

- BestSecret
- Bloomberg
- Expedia
- LY
- Naver
- Slack
- LinkedIn
- Wise

And there are many more.

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

Progress update

Goals from 2023 and current status:

- Short-term goals
  - ✅Modernizing dependencies continued
  - ✅Improving test suite
  - ✅Aligning code styles and patterns with Trino
  - ✅Supporting modern deployment tools
  - Tracing and metrics - to do: OpenTelemetry, ✅OpenMetrics
  - ✅Containerization
- Medium-term and long-term goals for capabilities and improvements
  - To do: Automatic query retry to other clusters if backend goes offline (memory)
  - ✅Integration with group providers for routing
  - To do: Supporting for parallel testing (e.g. duplicate a query to a second cluster and consume result)

Major new features from 2024

- Add Docker container
- Add Helm chart
- Update dependencies and removed security issues
- Require Java 23
- Add modernized and enriched user interface
- Improve access control, authentication, and other security

Routing

- Add support for configurable router modules
- More user, query, and HTTP request details available for routing logic
- New router modules
  - Based on query count per cluster
  - Based on cookie content
  - Using an external service for routing implementation

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 are directed to different clusters based on various factors

- Load balancing:
  - Random distribution (stochastic)
  - Based on Trino cluster load (query count)

Trino Gateway at Naver

- Naver Corporation is a global ICT company, providing South Korea's number 1 search portal "NAVER"

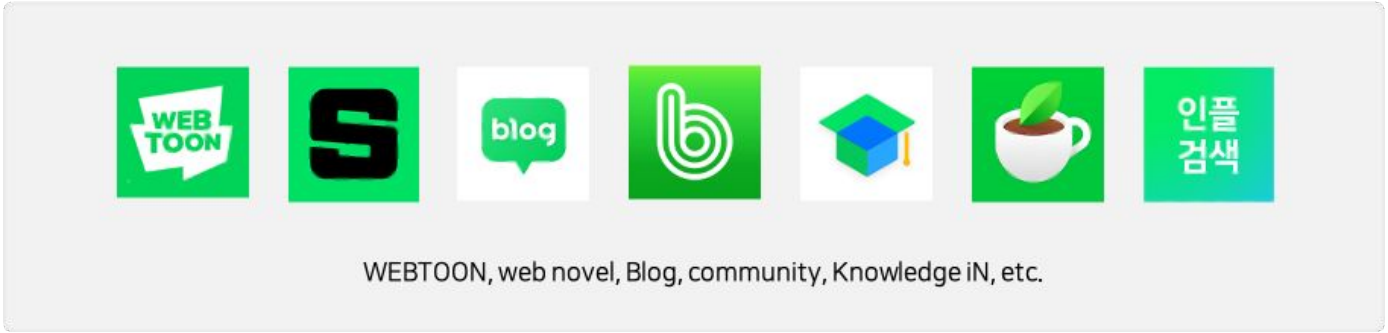

- AIDA Data Platform Team: CQuery
- 2024: 60K queries, 800 TB data, 25 Trillion rows every day.
- 12 PB business data stored in our data centers

Use an external routing service

Configuration in gateway-ha-config.yaml

```yaml
routingRules:
  rulesEngineEnabled: true
  rulesType: EXTERNAL
  rulesExternalConfiguration:
    urlPath: https://custom-routing-service.com/api/route
  excludeHeaders:
    - 'Authorization'
    - 'Accept-Encoding'
```

Custom routing service

- Handles POST requests from Trino Gateway
- Receives query information, headers, and user data
- Implements advanced routing logic

Benefits

- Enables dynamic, context-aware routing decisions
- Allows for easy updates to routing logic without Trino Gateway changes
- Facilitates integration with existing systems and data sources

Details for logical routing

Implemented in external service:

- Custom Routing Manager
  - Easily configurable by adding modules
  - Enables complex, dynamic routing decisions
- Routing Cateria
  - Query statistics
  - Cluster load and performance metrics
  - User privileges (e.g., VIP status)
- Resource Optimization
  - Identify and block “failure-prone” queries
  - Prevent resource waste on likely-to-fail operations
- Metadata Query Handling
  - Route catalog and table info requests to a dedicated metadata cluster (1 worker)
  - Efficiently manage high-volume requests from IDEs (e.g., PyCharm) and BI tools (e.g., Superset)
  - Reduce load on main clusters and Hive Metastore (HMS)
  - Mitigate impact of users without full table access permissions

Project future - For developers

- Plugin framework for routing modules
- Add more routing modules
- Update configuration to types
- Improve performance and scalability
- Remove easy-rules
- Test and support new Trino client protocol
- Improve logic for cluster status changes and monitoring

Project future - For users

- Improve getting started info and process
- Add routing and monitoring tools
- Improve Helm chart and support database management
- Richer UI
- Automatic query retry
- Add support to manage Trino cluster - longer term, big task!

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.