



Enduring with persistence to reach the summit

Martin Traverso

*A community and project update
for December 2024*

Hosted by  Starburst

Trino in 2024 by the numbers



- Trino releases 436-467
- 5,000+ additional commits, 40,000+ total
- 225+ unique contributors in 2024, 925+ total
- 10.5k+ stars on GitHub
- 13,500+ Slack members
- [Trino Community Broadcast](#) episodes 54-67

2024 has been a record year for Trino

Hosted by  Starburst



What's happened in 2024?

Hosted by  Starburst

The text "Hosted by" is in a white, sans-serif font. To its right is the Starburst logo, which is a white, multi-pointed starburst or sunburst icon. To the right of the icon is the word "Starburst" in a white, sans-serif font.

Major Updates in 2024

- Improved access control integrations
 - Open Policy Agent (OPA) access control
 - Apache Ranger access control
- Observability
 - Support for OpenTelemetry and OpenMetrics
 - Kafka query event listener
 - OpenLineage query event listener



Major Updates in 2024

- Client protocol improvements
- Python user defined functions
- Software supply chain security



Major Updates in 2024

- Major improvements for Trino Gateway
- New JavaScript client - [trino-js-client](#)
- New C# and ADO.net driver
- Numerous Helm chart improvements

Client protocol improvements

- Challenges
 - Single stream, roundtrip latency
 - Slow clients tie up resources
- Parallel protocol
 - Stage data in object storage
 - Automatic fallback
 - JDBC driver and CLI ready, Python client in the works

Python user-defined functions



```
WITH FUNCTION xor(a boolean, b boolean)
RETURNS boolean
LANGUAGE PYTHON
WITH (handler = 'bool_xor')
AS $$
import operator
def bool_xor(a, b):
    return operator.xor(a, b)
$$
SELECT xor(true, false), xor(false, true);
```

Note the use of \$\$ to enclose Python

Hosted by  Starburst

Python user-defined functions



- CREATE/DROP/SHOW FUNCTION, WITH FUNCTION
- Execution inline with query processing
- Limitations
 - Experimental
 - Python language and core libraries only
 - Lots of room to improve performance

One more



```
WITH FUNCTION reverse_words(s varchar)
RETURNS varchar
LANGUAGE PYTHON
WITH (handler = 'reverse_words')
AS $$
import re

def reverse(s):
    str = ""
    for i in s:
        str = i + str
    return str

pattern = re.compile(r"\w+[,!?\"]\w*")

def process_word(word):
    # Reverse only words without non-letter signs
    return word if pattern.match(word) else reverse(word)

def reverse_words(payload):
    text_words = payload.split(' ')
    return ' '.join([process_word(w) for w in text_words])
$$
SELECT comment, reverse_words(comment)
FROM nation
WHERE nationkey IN (5, 6, 12)
```

Connectors



- File system caching for Iceberg and Delta Lake
- Apache Polaris and Unity Catalog integration
- Lots of quality of life improvements
- New connectors: Exasol, Faker, OpenSearch, Snowflake, Vertica
- Removed legacy connectors: Accumulo, Atop, Localfile, Raptor
- Remove [Kudu](#), [Kinesis](#), and [Phoenix](#) connectors in 2025?

Roadmap



- Secrets management and dynamic catalogs
- Client protocol improvements for all client drivers
- Packaging improvements
- More connectors such as DuckDB, LanceDB, HsqlDB, Loki, ...
- Continued and even ***increased work*** on performance improvements
- Research and prototype next generation optimizer
- SQL language improvements such as PIVOT, ASOF joins, ...

Trino community



- Monthly maintainer and contributor calls
- [Trino Fest](#) in June in Boston
- Japanese edition of [Trino: The Definitive Guide](#)
- New [SQL training material](#)
- Improved [website](#) and [contribution](#) docs



Questions?