

Adopting Trino's Fault-tolerant Execution mode at Quora

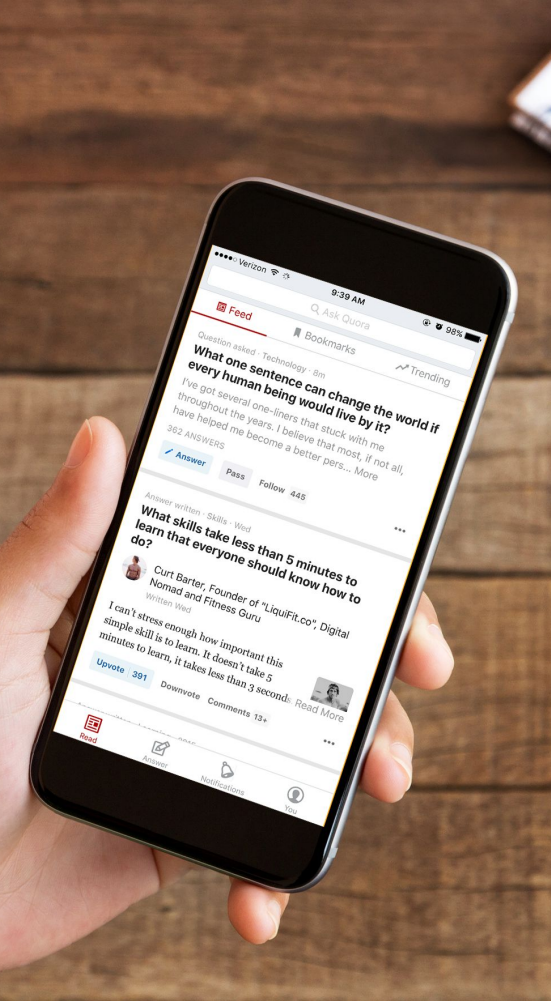
Gabriel Fernandes - Software Engineer (gfdoliveira@quora.com)

Yifan Pan - Software Engineer (epan@quora.com)

The Quora logo, featuring the word "Quora" in a bold, red, serif font.

What is Quora?

Our mission is to share and grow the world's knowledge.



Overview

- How is Trino used at Quora?
- Trino setup before and after fault-tolerant execution
- Benchmark results using fault-tolerant execution

How is Trino used at Quora?

How is Trino used at Quora

Use Case	Query Running Time
Ad-hoc Analysis	Seconds to hours
Time-series Dashboard	Seconds to minutes
Extract-transform-load (ETL)	Minutes to ~ 1, 2 hours
A/B Testing	Tens of minutes ~ 3, 4 hours

Trino setup before and after fault-tolerant execution (FTE)


Trino setup before and after FTE

Use Case	Query Running Time	Run with FTE?		
		~ 10/2022		
Ad-hoc Analysis	Seconds to hours	No		
Time-series Dashboard	Seconds to minutes	No		
ETL	Minutes to ~1, 2 hours	No		
A/B testing	Tens of minutes to ~3,4 hours	No		
		OOM requiring scale up clusters or set up stricter concurrency control		

Trino setup before and after FTE

Use Case	Query Running Time	Run with FTE?		
		~ 10/2022	10/2022 ~ 4/2023	
Ad-hoc Analysis	Seconds to hours	No	No	
Time-series Dashboard	Seconds to minutes	No	No	
ETL	Minutes to ~1, 2 hours	No	Yes	
A/B testing	Tens of minutes to ~3,4 hours	No	Yes	
		OOM requiring scale up clusters or set up stricter concurrency control	-Fewer OOM errors -Higher costs -Slower short queries	

Trino setup before and after FTE

Use Case	Query Running Time	Run with FTE?		
		~ 10/2022	10/2022 ~ 4/2023	4/2023 ~ Now
Ad-hoc Analysis	Seconds to hours	No	No	No
Time-series Dashboard	Seconds to minutes	No	No	No
ETL	Minutes to ~1, 2 hours	No	Yes	Selected queries run with FTE
A/B testing	Tens of minutes to ~3,4 hours	No	Yes	Yes
		OOM requiring scale up clusters or set up stricter concurrency control	-Fewer OOM errors -Higher costs -Slower short queries	

Benchmark results (using Trino 406)

Benchmark - Failure rate

- Reduced query failure rate by 71%
 - Queries survive occasional worker failure
 - No out of memory errors

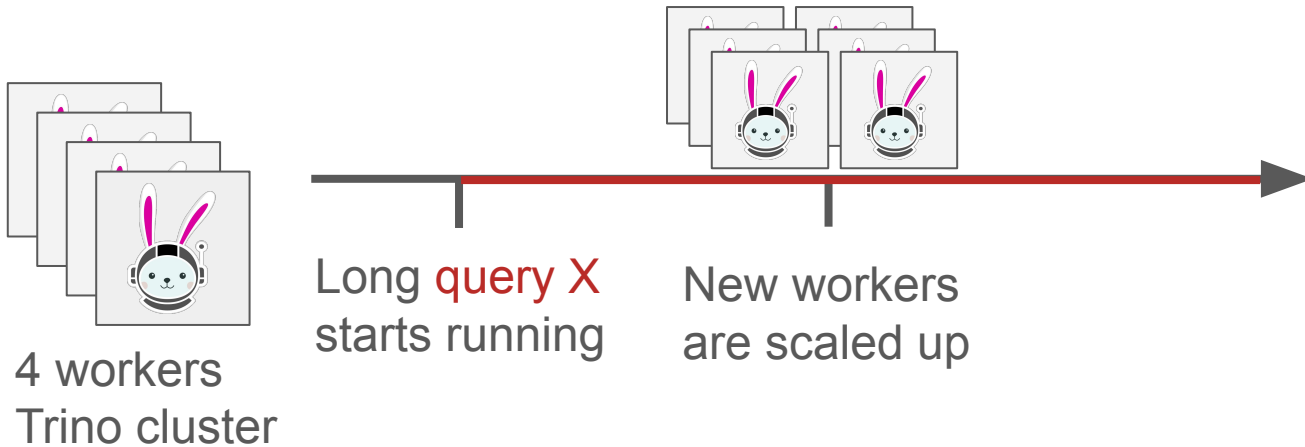
In the 2 weeks before using FTE:
157 EXCEEDED_LOCAL_MEMORY_LIMIT
40 CLUSTER_OUT_OF_MEMORY




5 months into using FTE:
3 CLUSTER_OUT_OF_MEMORY

Benchmark - Autoscaling

- Better autoscaling in FTE
 - Scale up is more efficient



Benchmark - Autoscaling



Traditional mode: New workers likely **don't help run query X**

FTE mode: New workers help running query X right away

Benchmark - Autoscaling

- Better autoscaling in FTE
 - Scale up is more efficient
 - Quicker to gracefully scale down
 - Workers don't have to wait for full query completion

Benchmark - Execution time

Fault Tolerance overhead is significant for short queries*

Execution time for all queries		
Metric	Without FTE	With FTE
Mean	6.06s	8.11s
p50	1.0ms	2.5ms
p90	88.2ms	201.2ms

Benchmark - Execution time

Fault Tolerance overhead is significant for short queries*

Quicker execution of long queries

Execution time for all queries

Metric	Without FTE	With FTE
Mean	6.06s	8.11s
p50	1.0ms	2.5ms
p90	88.2ms	201.2ms

Execution time for queries > 1h CPU Time

Metric	Without FTE	With FTE
Mean	449.6s	246.2s
p50	1.8s	1.3s
p90	15.7s	10.5s

Benchmark - Execution time

Fault Tolerance overhead is significant for short queries*

Quicker execution of long queries

Execution time for all queries

Metric	Without FTE	With FTE
Mean	6.06s	8.11s
p50	1.0ms	2.5ms
p90	88.2ms	201.2ms

Execution time for queries > 1h CPU Time

Metric	Without FTE	With FTE
Mean	449.6s	246.2s
p50	1.8s	1.3s
p90	15.7s	10.5s

*Note: results are from Trino 406, but many optimizations were made in Trino 412 for short queries, with 40%+ latency reduction.



Benchmark - Costs

Using EC2 instances for Trino workers and S3 storage

	Using Fault-Tolerance Execution
Worker Instance costs	May see savings due to improved query execution time.
S3 Request costs	Increased by about 7% of Instance Costs due to query checkpointing.
S3 Storage costs	Unchanged.

Summary

Summary

- Maintaining both a FTE and a non-FTE cluster is ideal
- Proper query routing between clusters is necessary to improve reliability, avoid execution time degradation
- Cost increase is expected

Thank you!

Gabriel Fernandes - Software Engineer
(gfdoliveira@quora.com)

Yifan Pan - Software Engineer
(epan@quora.com)

Quora