

Apache Gluten (incubating)

Weiting Chen weiting.chen@intel.com




CONTENTS


1. Why Native Execution Engine
2. Introducing Apache Gluten
3. Performance & Best Practice
4. Cross Apache Projects Collaborations
5. Summary

GLUTEN FRAMEWORK

Apache Spark




Apache Gluten



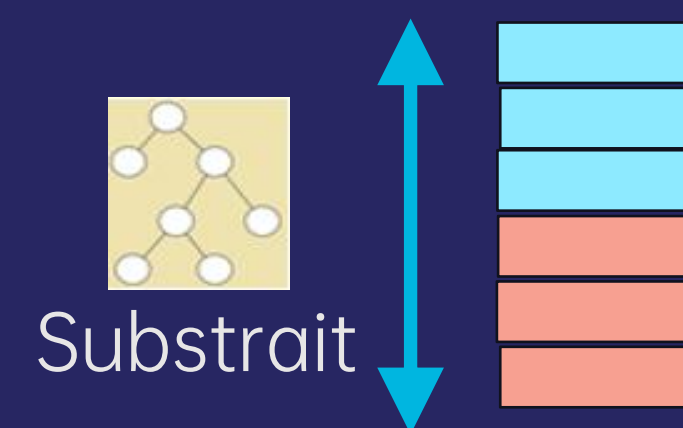
Substrait

Velox




Substrait

ClickHouse



Substrait

Arrow +
Gazelle



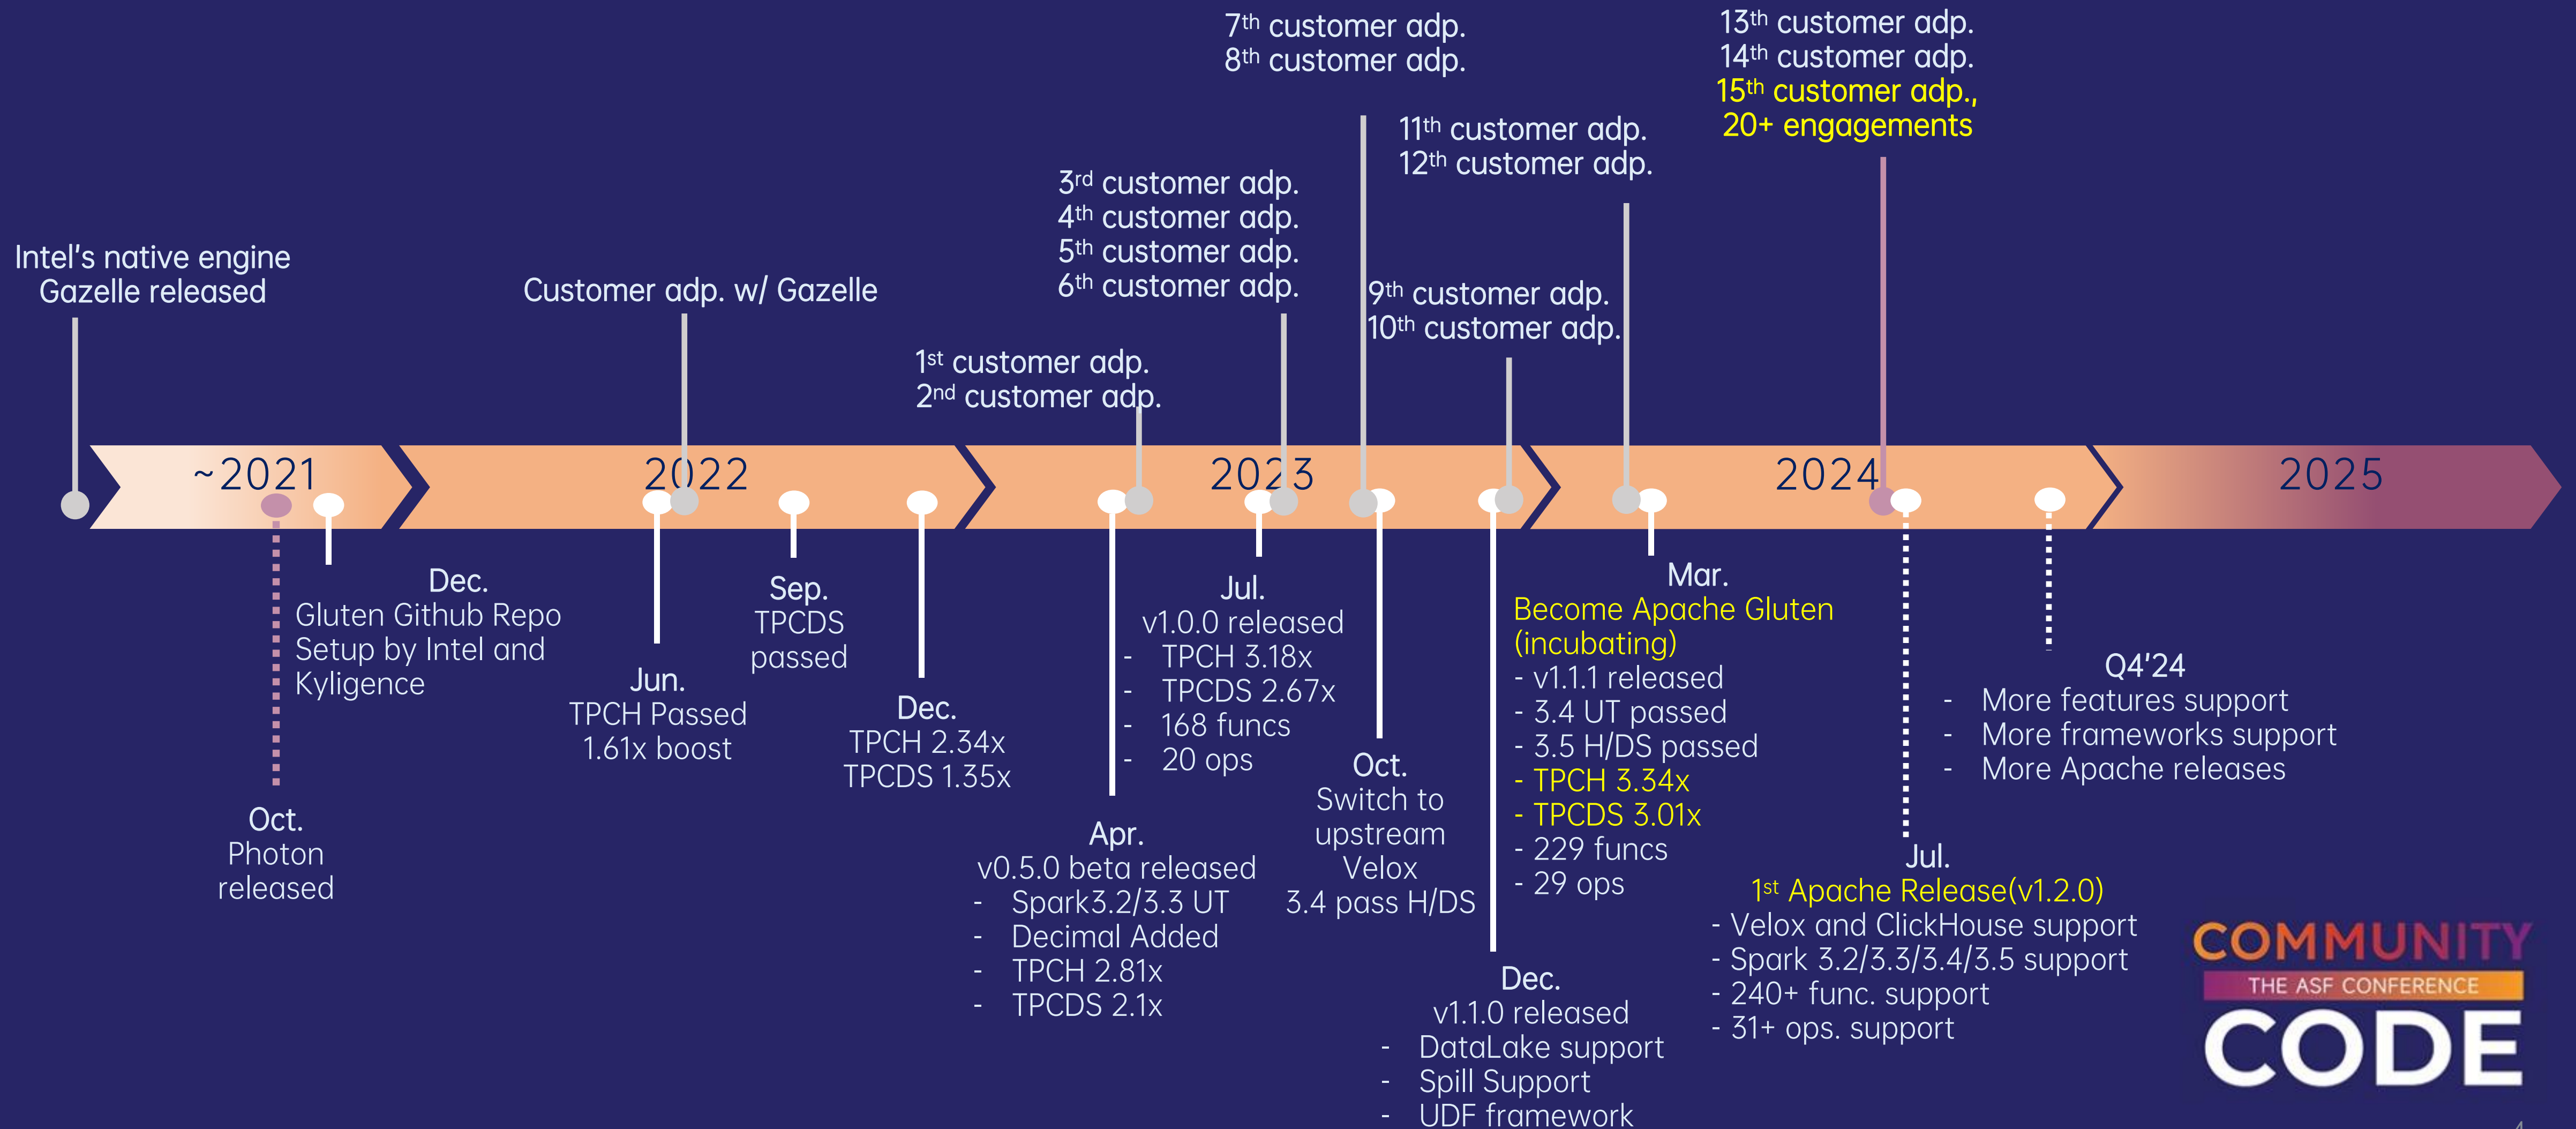
Substrait

FPGA/GPU/ASIC
Accelerators

Reference:

1. **Gluten:** <https://github.com/apache/incubator-gluten>
2. **Velox:** <https://github.com/facebookincubator/velox>
3. **Clickhouse:** <https://github.com/Kyigence/ClickHouse> forked from <https://github.com/ClickHouse/ClickHouse>
4. **Gazelle:** https://github.com/oap-project/gazelle_plugin
5. **Apache Arrow:** <https://github.com/apache/arrow>
6. **Substrait:** <https://substrait.io/>

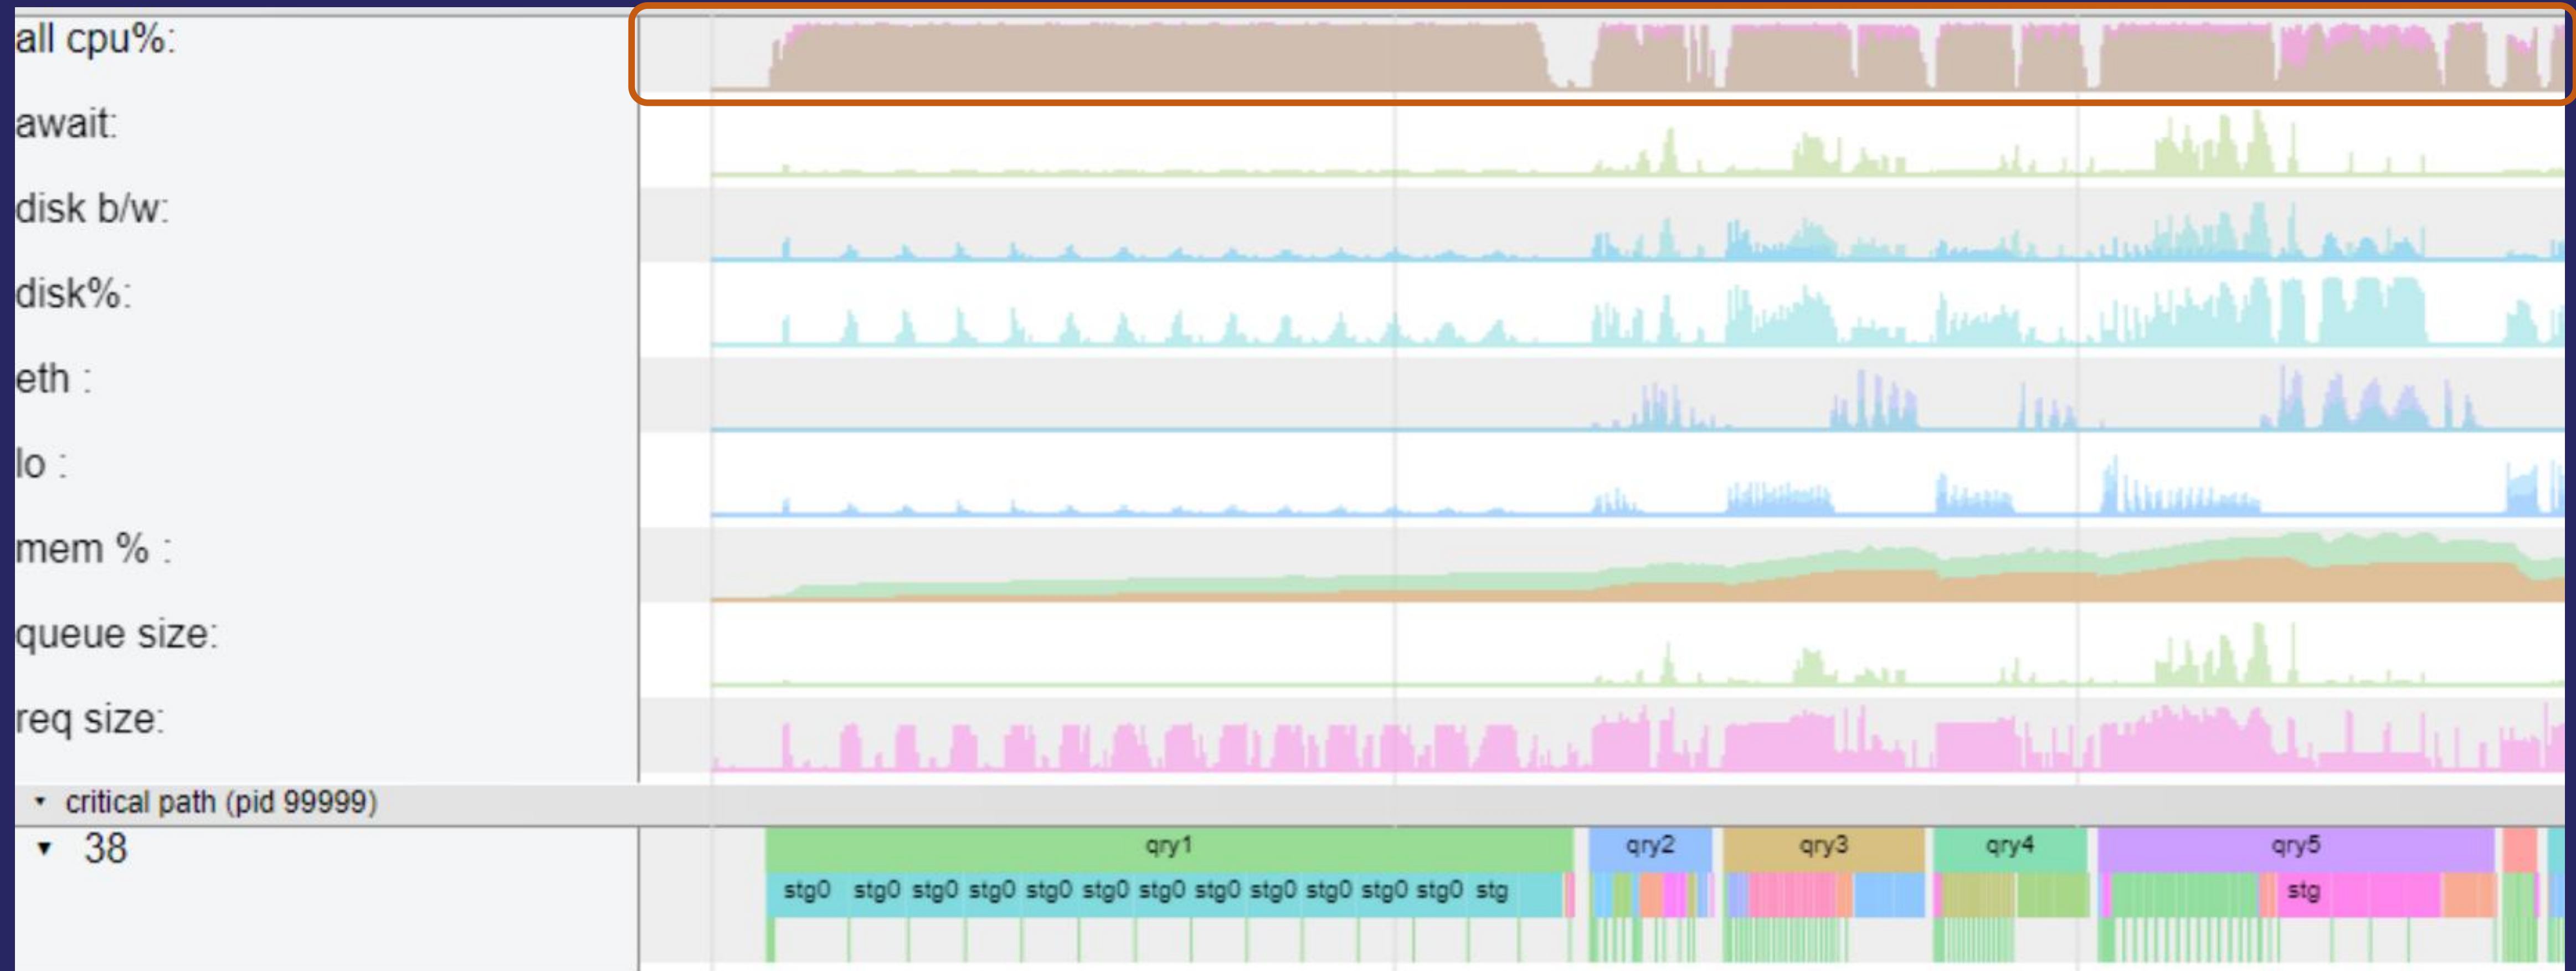
GLUTEN JOURNEY



BASIC OPERATOR PERFORMANCE STOPPED GROW



CPU HAS BECOME THE BOTTLENECK



Node: 3
CPU: 2 x 36C Intel(R) Xeon(R) Platinum 8360Y 3.5GHz
Memory: 512GB DDR4

Disk: 4 x INTEL SSDPE2KE016T8
NIC: 25Gbps
Dataset: 3T

SQL ENGINE DEVELOPED YEARS ...

Products

ClickHouse

DuckDB

MongoDB


Libraries

Velox


Apache Arrow



Native Implementation

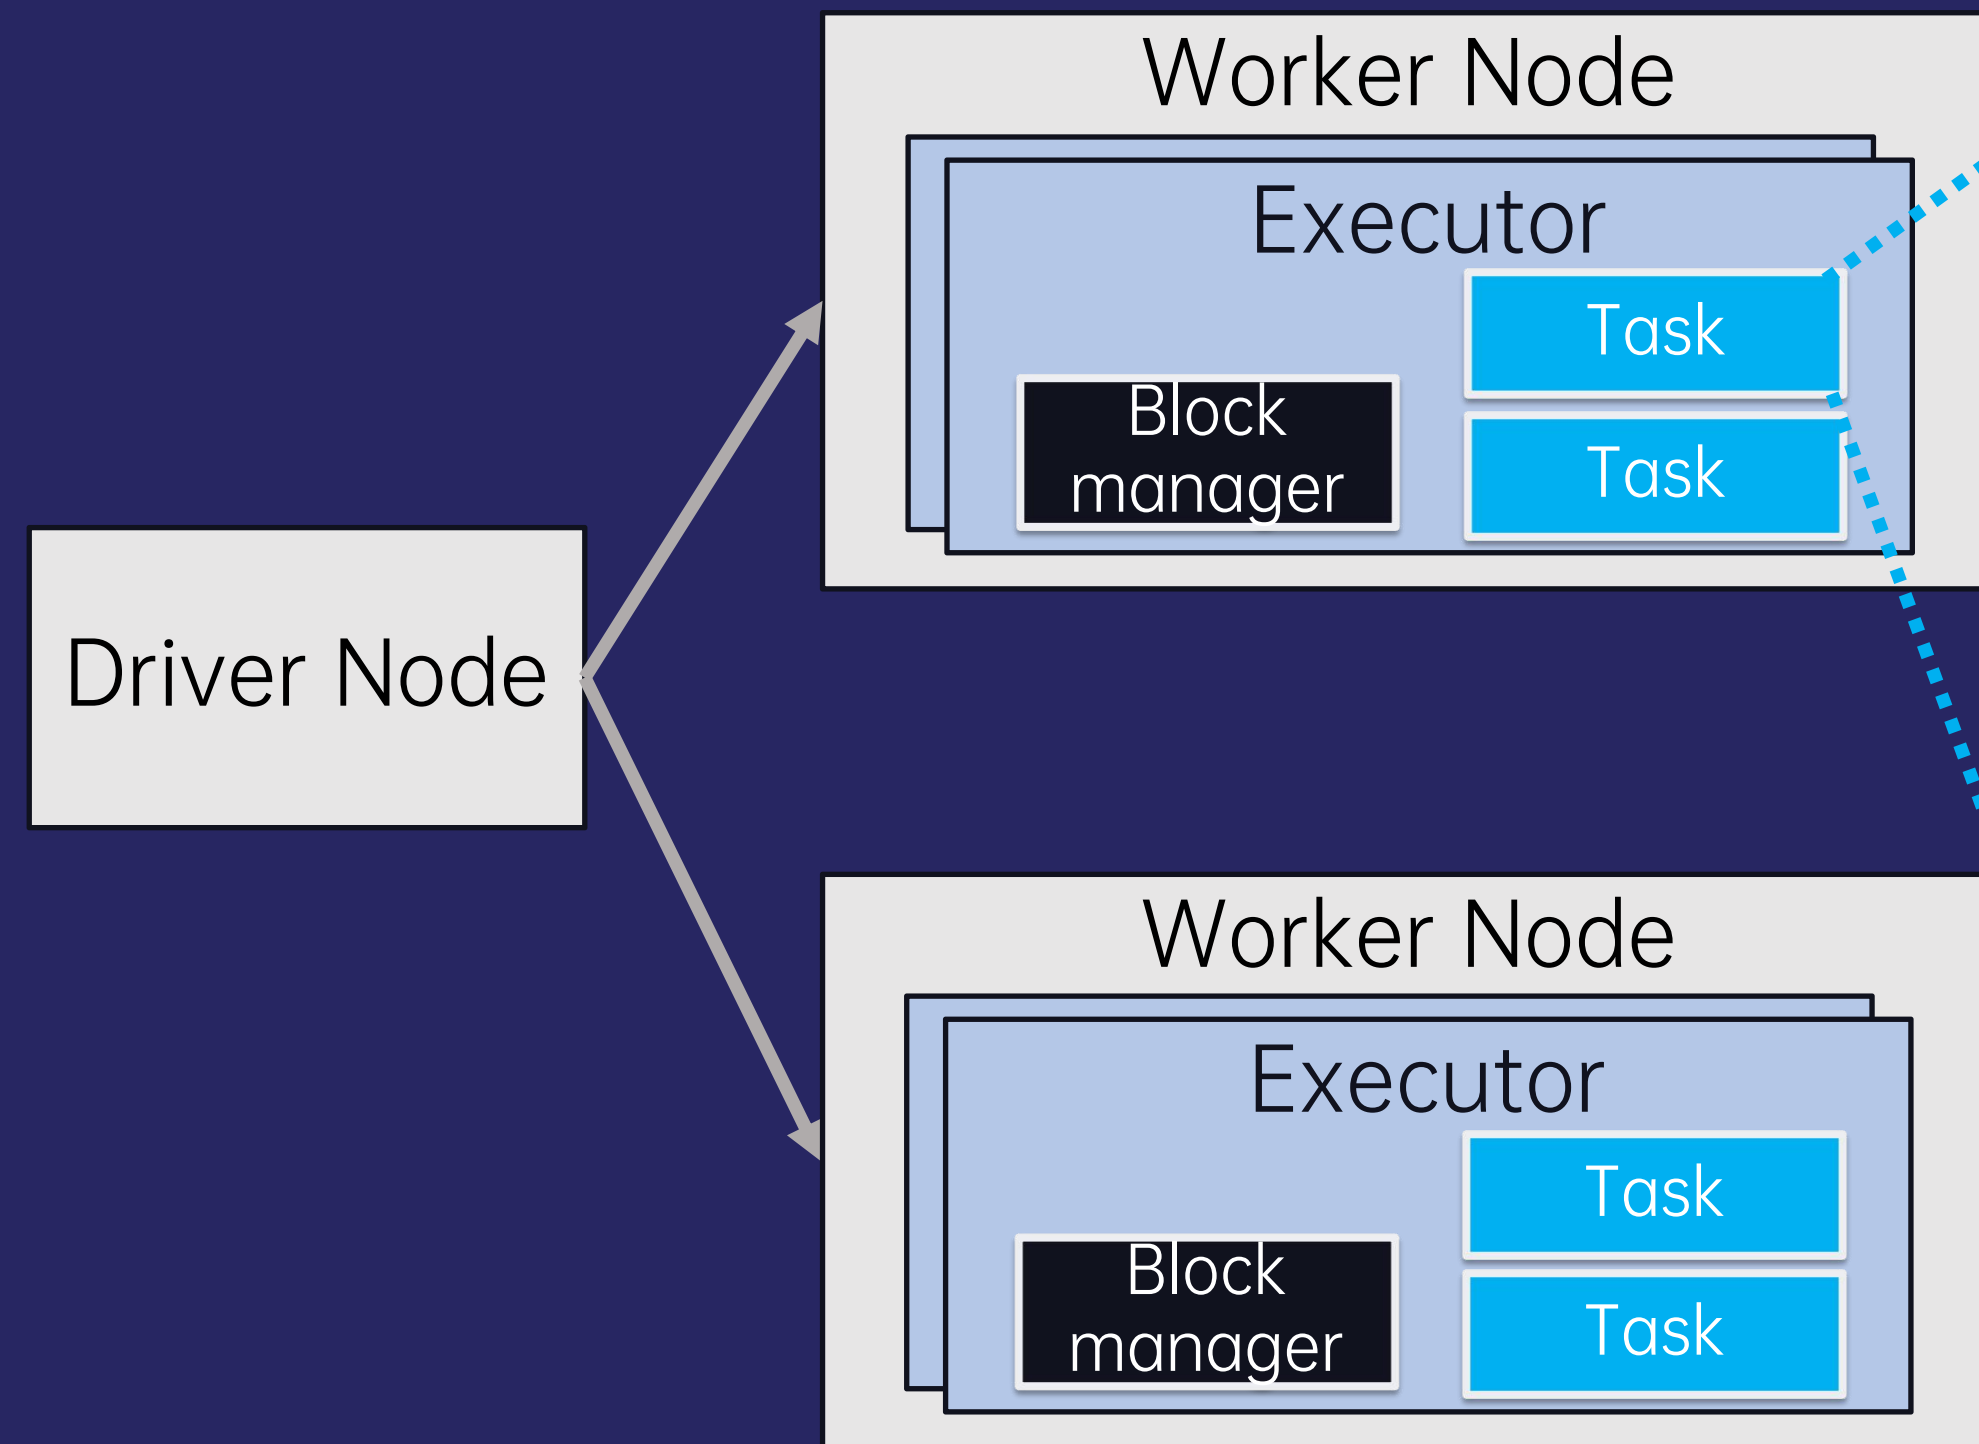


Columnar Store



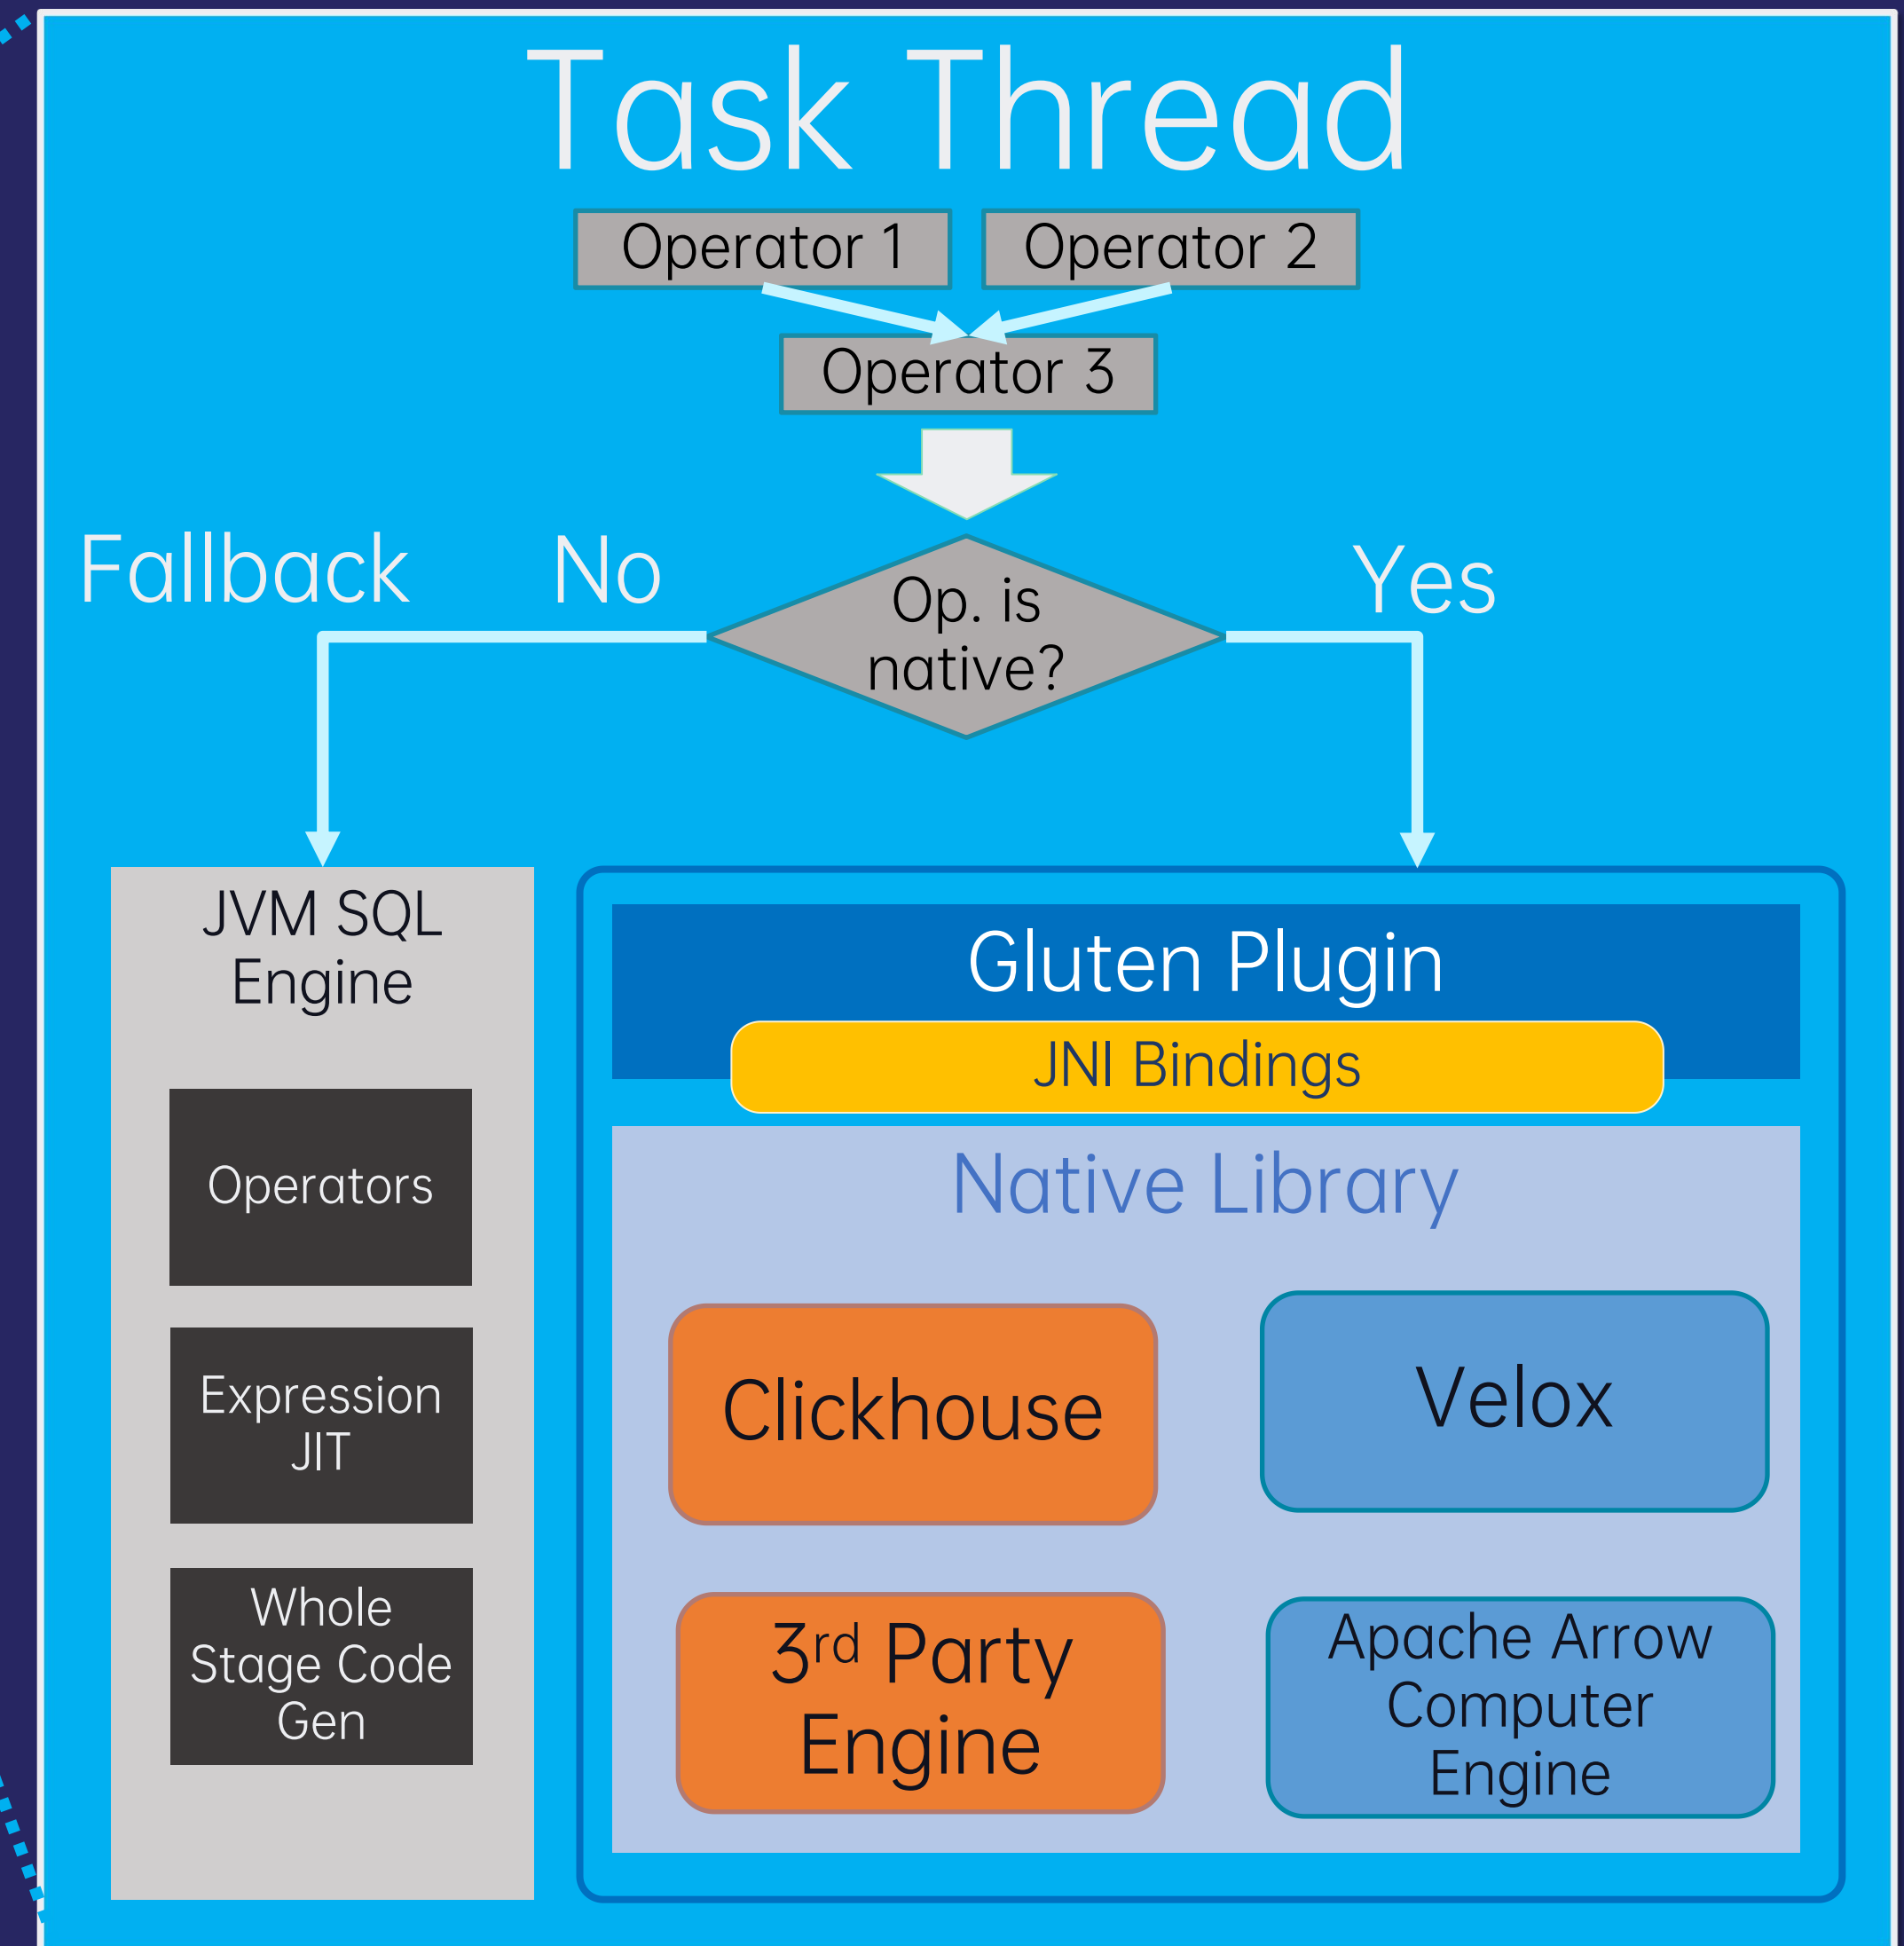
Vectorized Processing

AN EVOLUTION ON THE WAY



Spark Scale Out Framework +
Optimal Native Library

Never ever redesign the wheel !



GLUTEN COMPONENTS

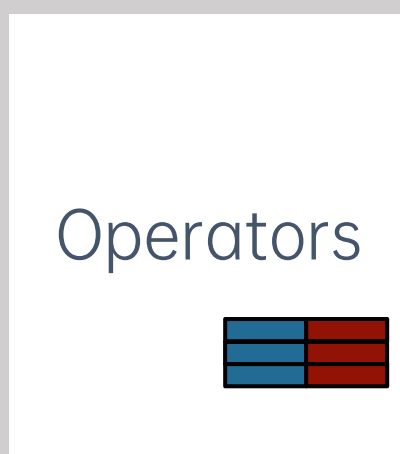


DataFrame

Catalyst Query Plan Optimization

Tungsten Physical Plan Execution

JVM SQL Engine



Expression
JIT

Whole
Stage
Code Gen

Gluten Plugin

Plan
Conversion

Memory Mgr.

Columnar
Shuffle

Shim Layer

Fallback

Metric

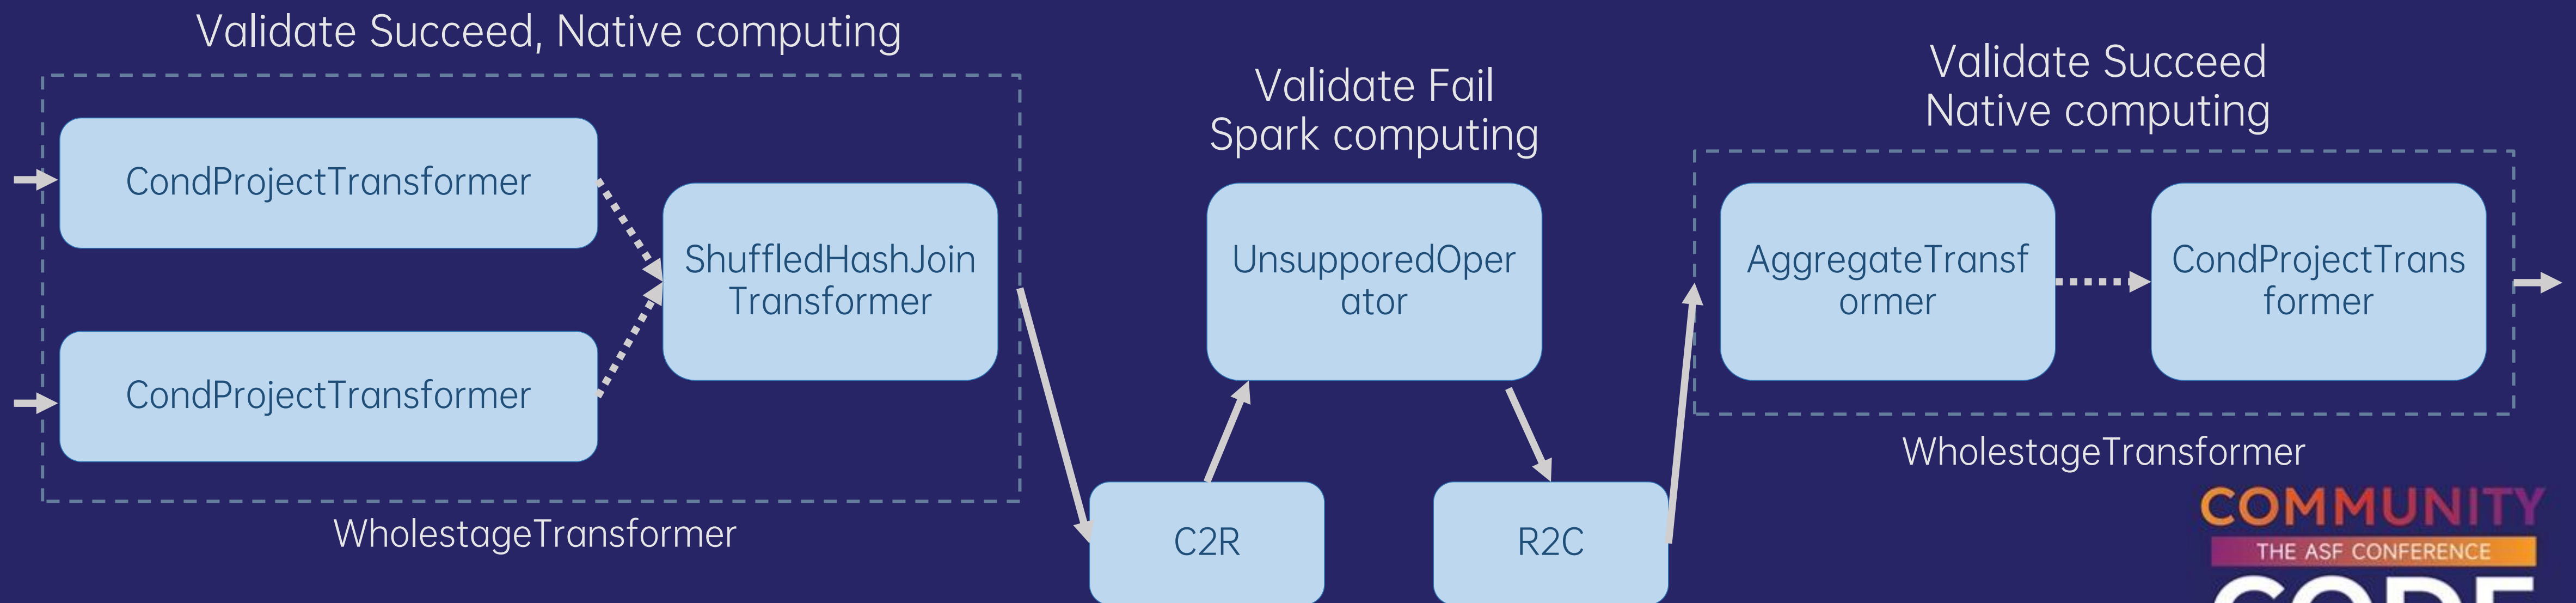
Native Library

GLUTEN FALLBACK PROCESSING

Replace Spark physical plan by Transformers

Unsupported operators with express will fallback to Vanilla Spark

Combining several operators into WholestageTransformer to generate whole stage subtrait plan



GLUTEN SHUFFLE

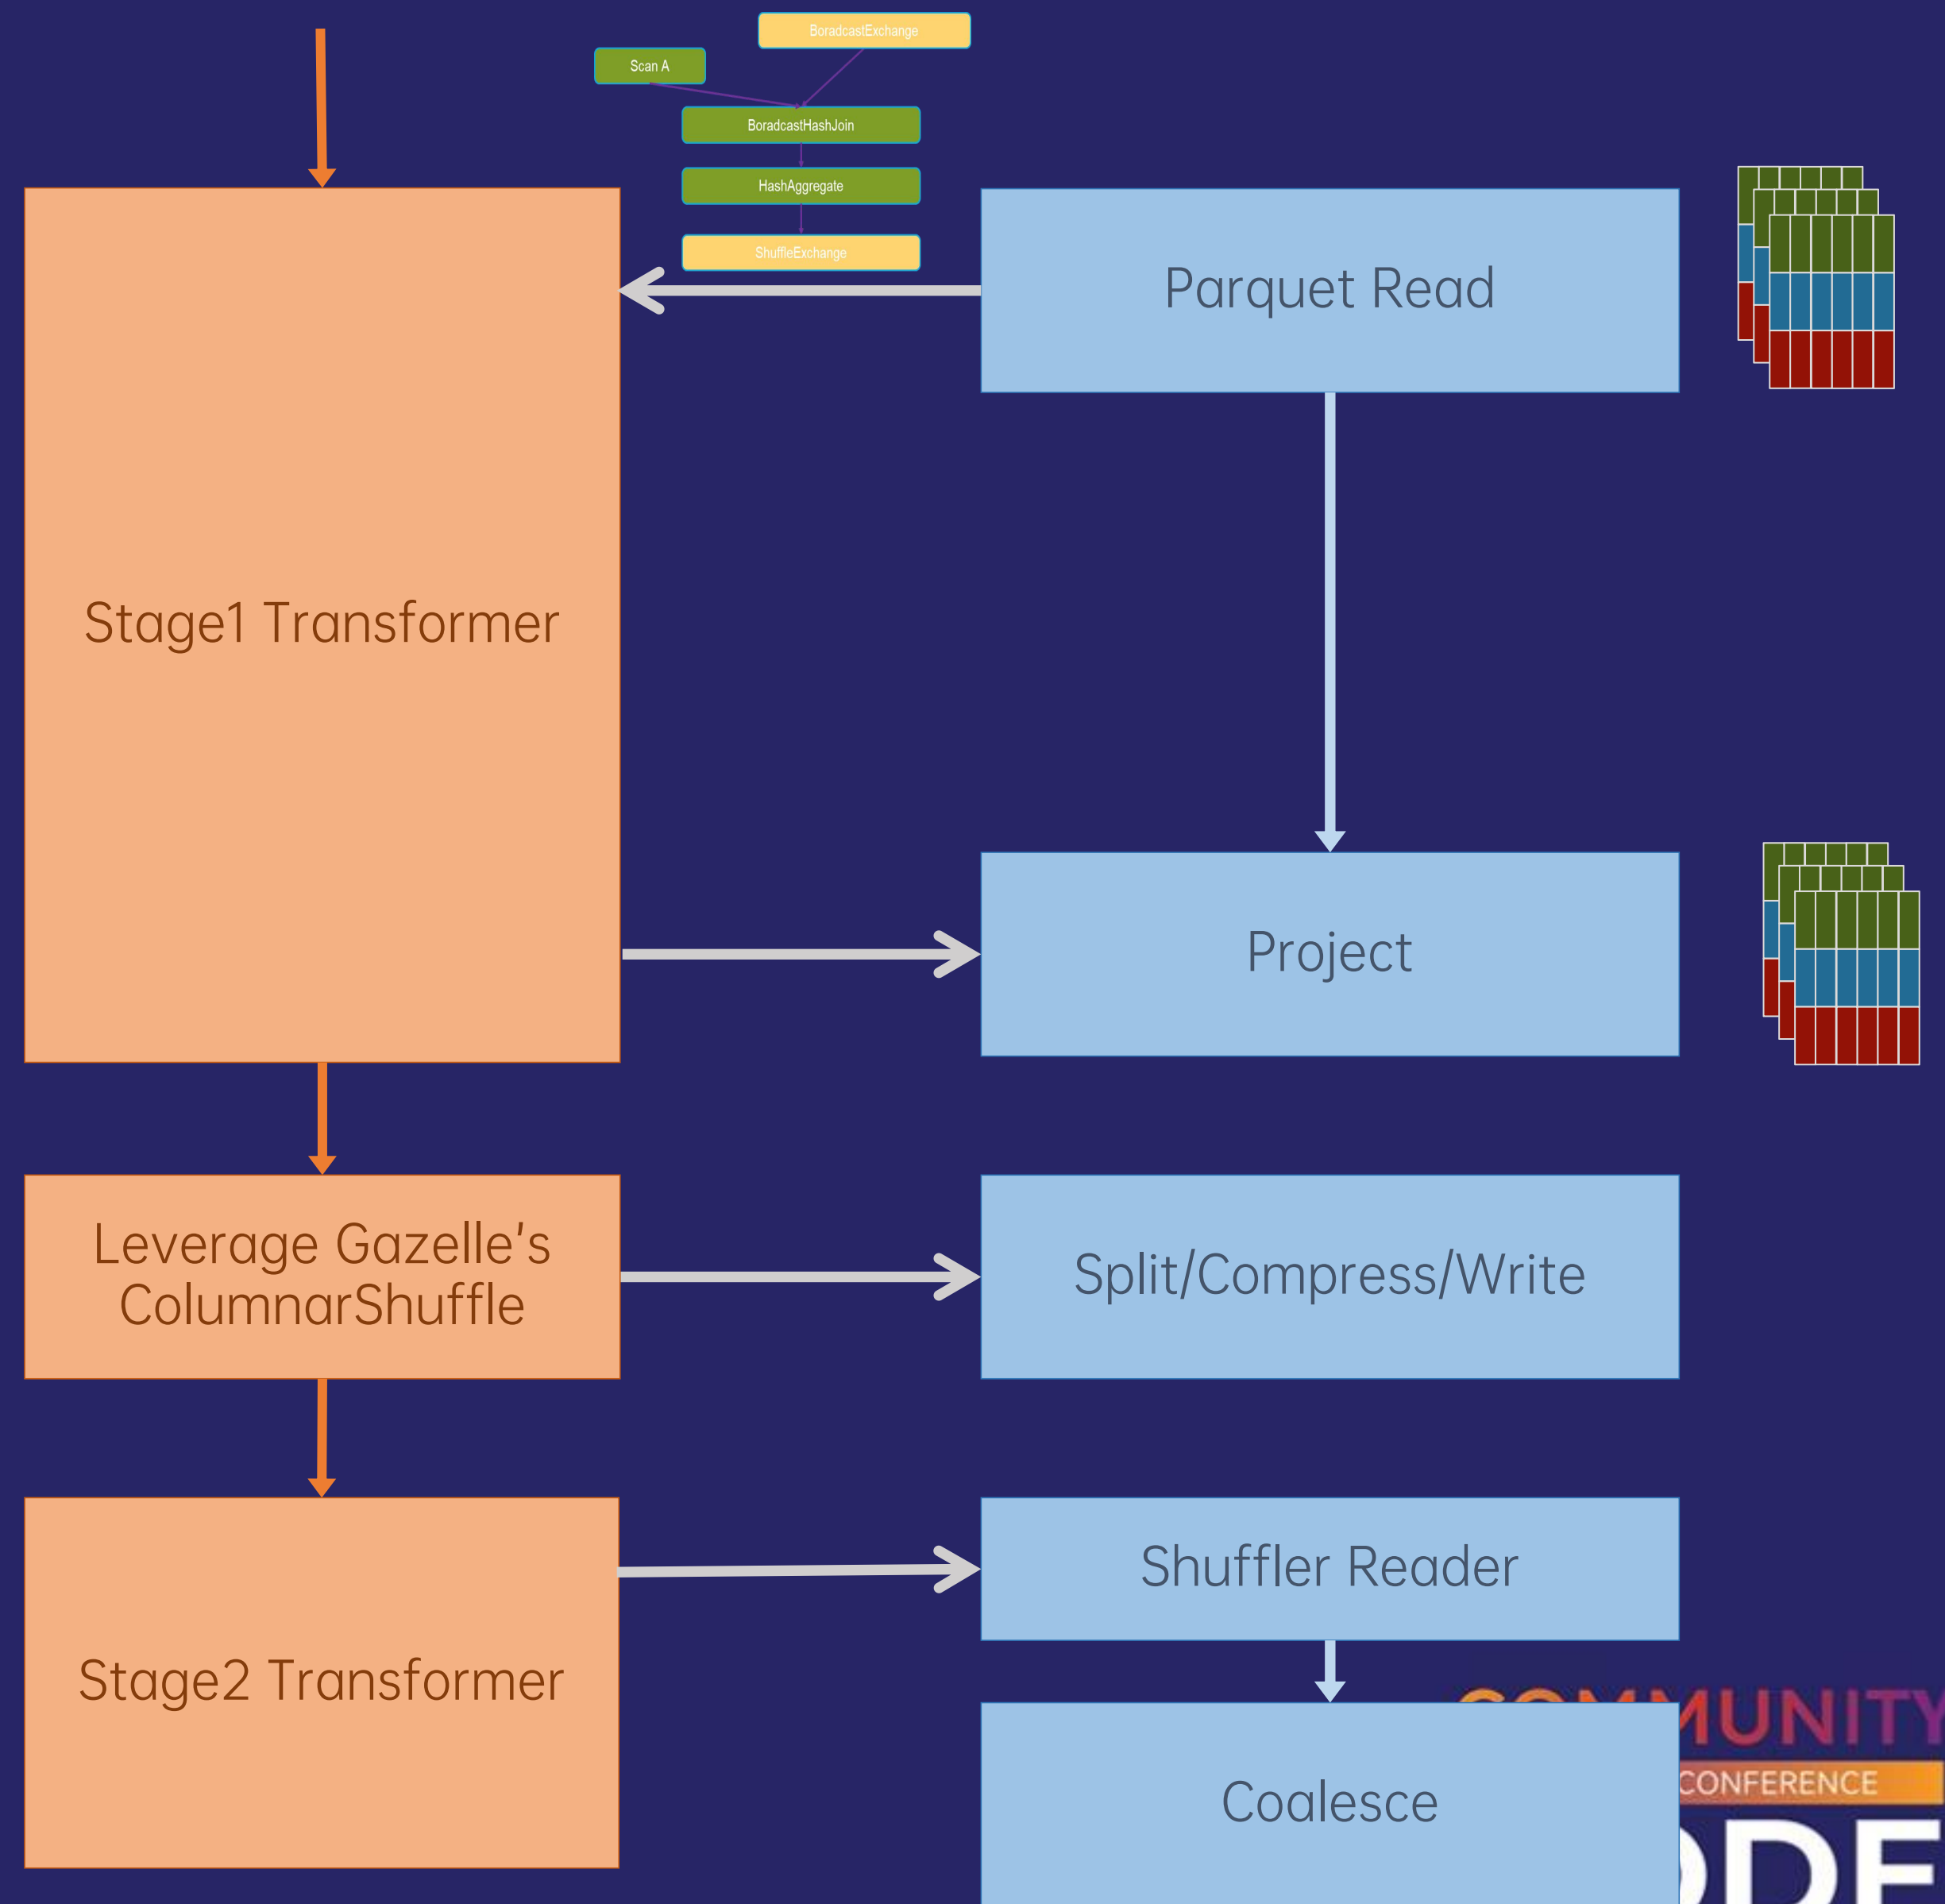
Velox backend implementation reuse
Gazelle's Arrow based shuffle

Clickhouse backend use its internal
columnar format

Cache split batches in memory

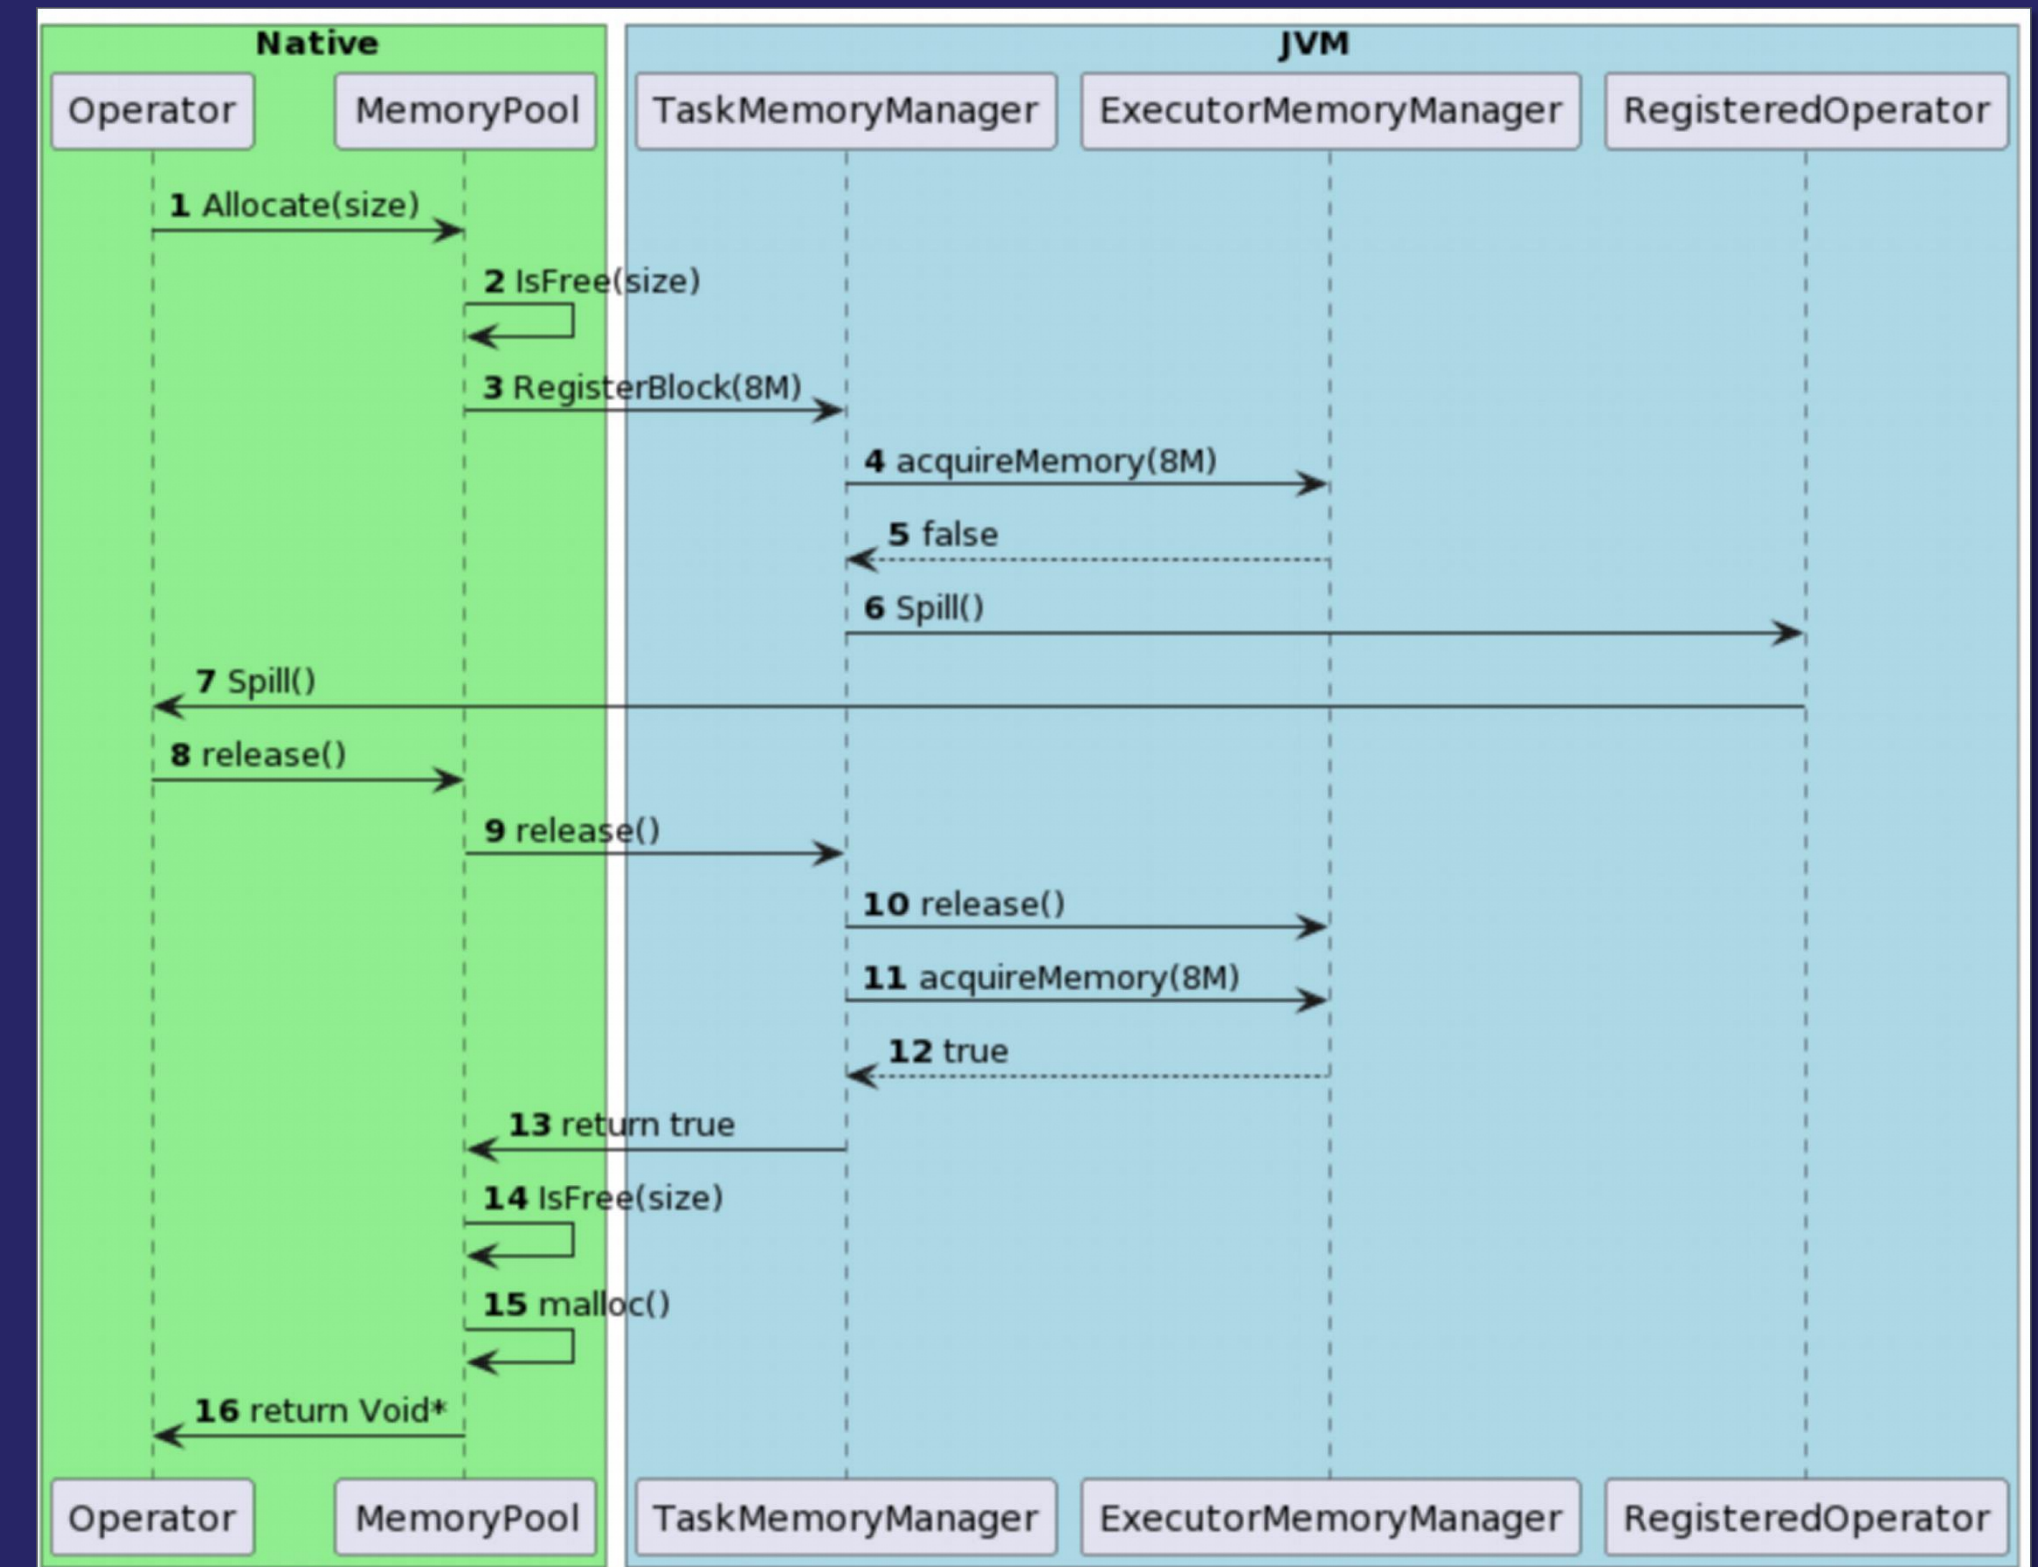
On OOM, spill all batches into single
file with part. id seq.

Merge all spill files into data file at end



GLUTEN MEMORY MANAGEMENT

- Use Spark off-heap memory as unified native memory mgmt.
- MemoryPool hold a block of memory
- Acquire from task memory manager if not enough
- Spill happens if task doesn't have enough memory



GLUTEN CURRENT STATUS

- v1.2.0 as 1st Apache release WIP
- ClickHouse & Velox backend support
- Ubuntu 20/22, CentOS7/8/9 Support
- Spark3.2/3.3/3.4/3.5 Support
- AQE, DPP Support
- Unified Memory management via Spark off-heap memory
- Support most data types
- Support most Datalake
- ~85% Spark common func. support
- TPC-H/TPC-DS fully offloaded
- More customer workloads adoption

Function Coverage:

Spark3.2 total	Spark Common	Gluten
361	289	240+

Operator Coverage:

Spark3.2 total	Spark Common	Gluten
90	37	31+

Data Type Coverage:

BOOL	BYTE	SHORT	INT	LONG	FLOAT	DOUBLE	STRING	NULL	BINARY
ARRAY	MAP	STRUCT	DATE	TIMESTAMP	DECIMAL	CALENDAR	UDT		

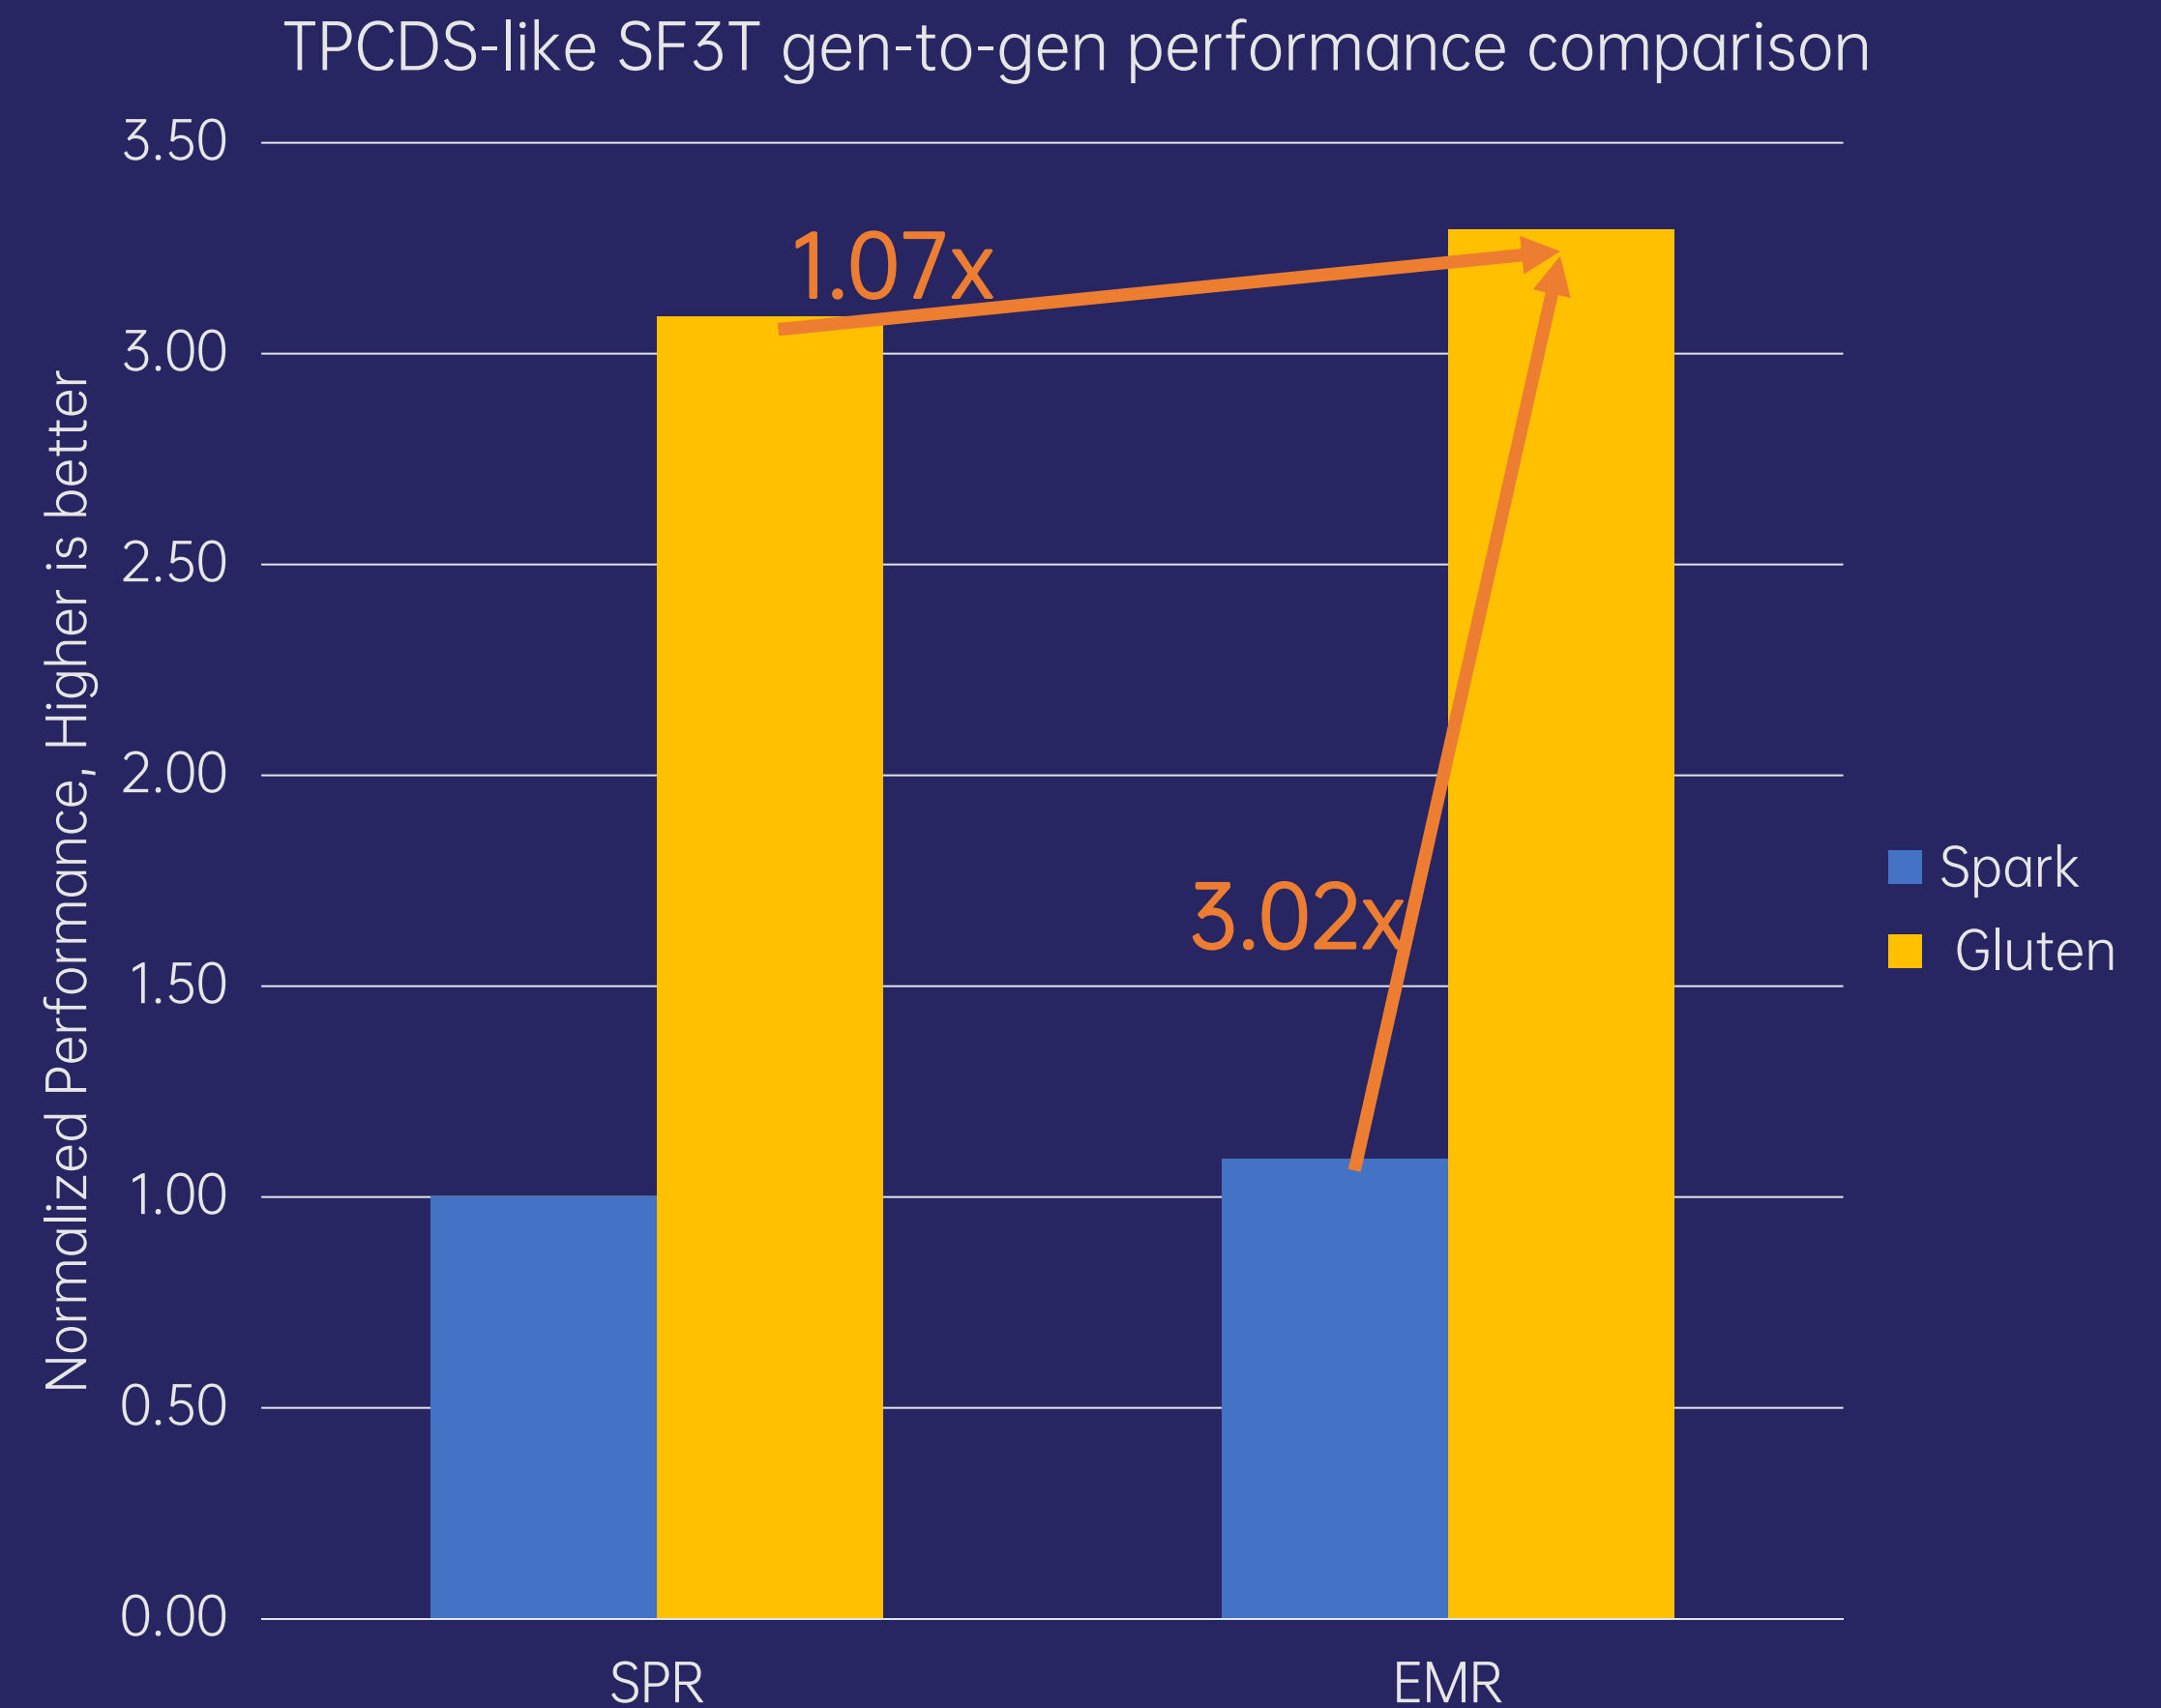
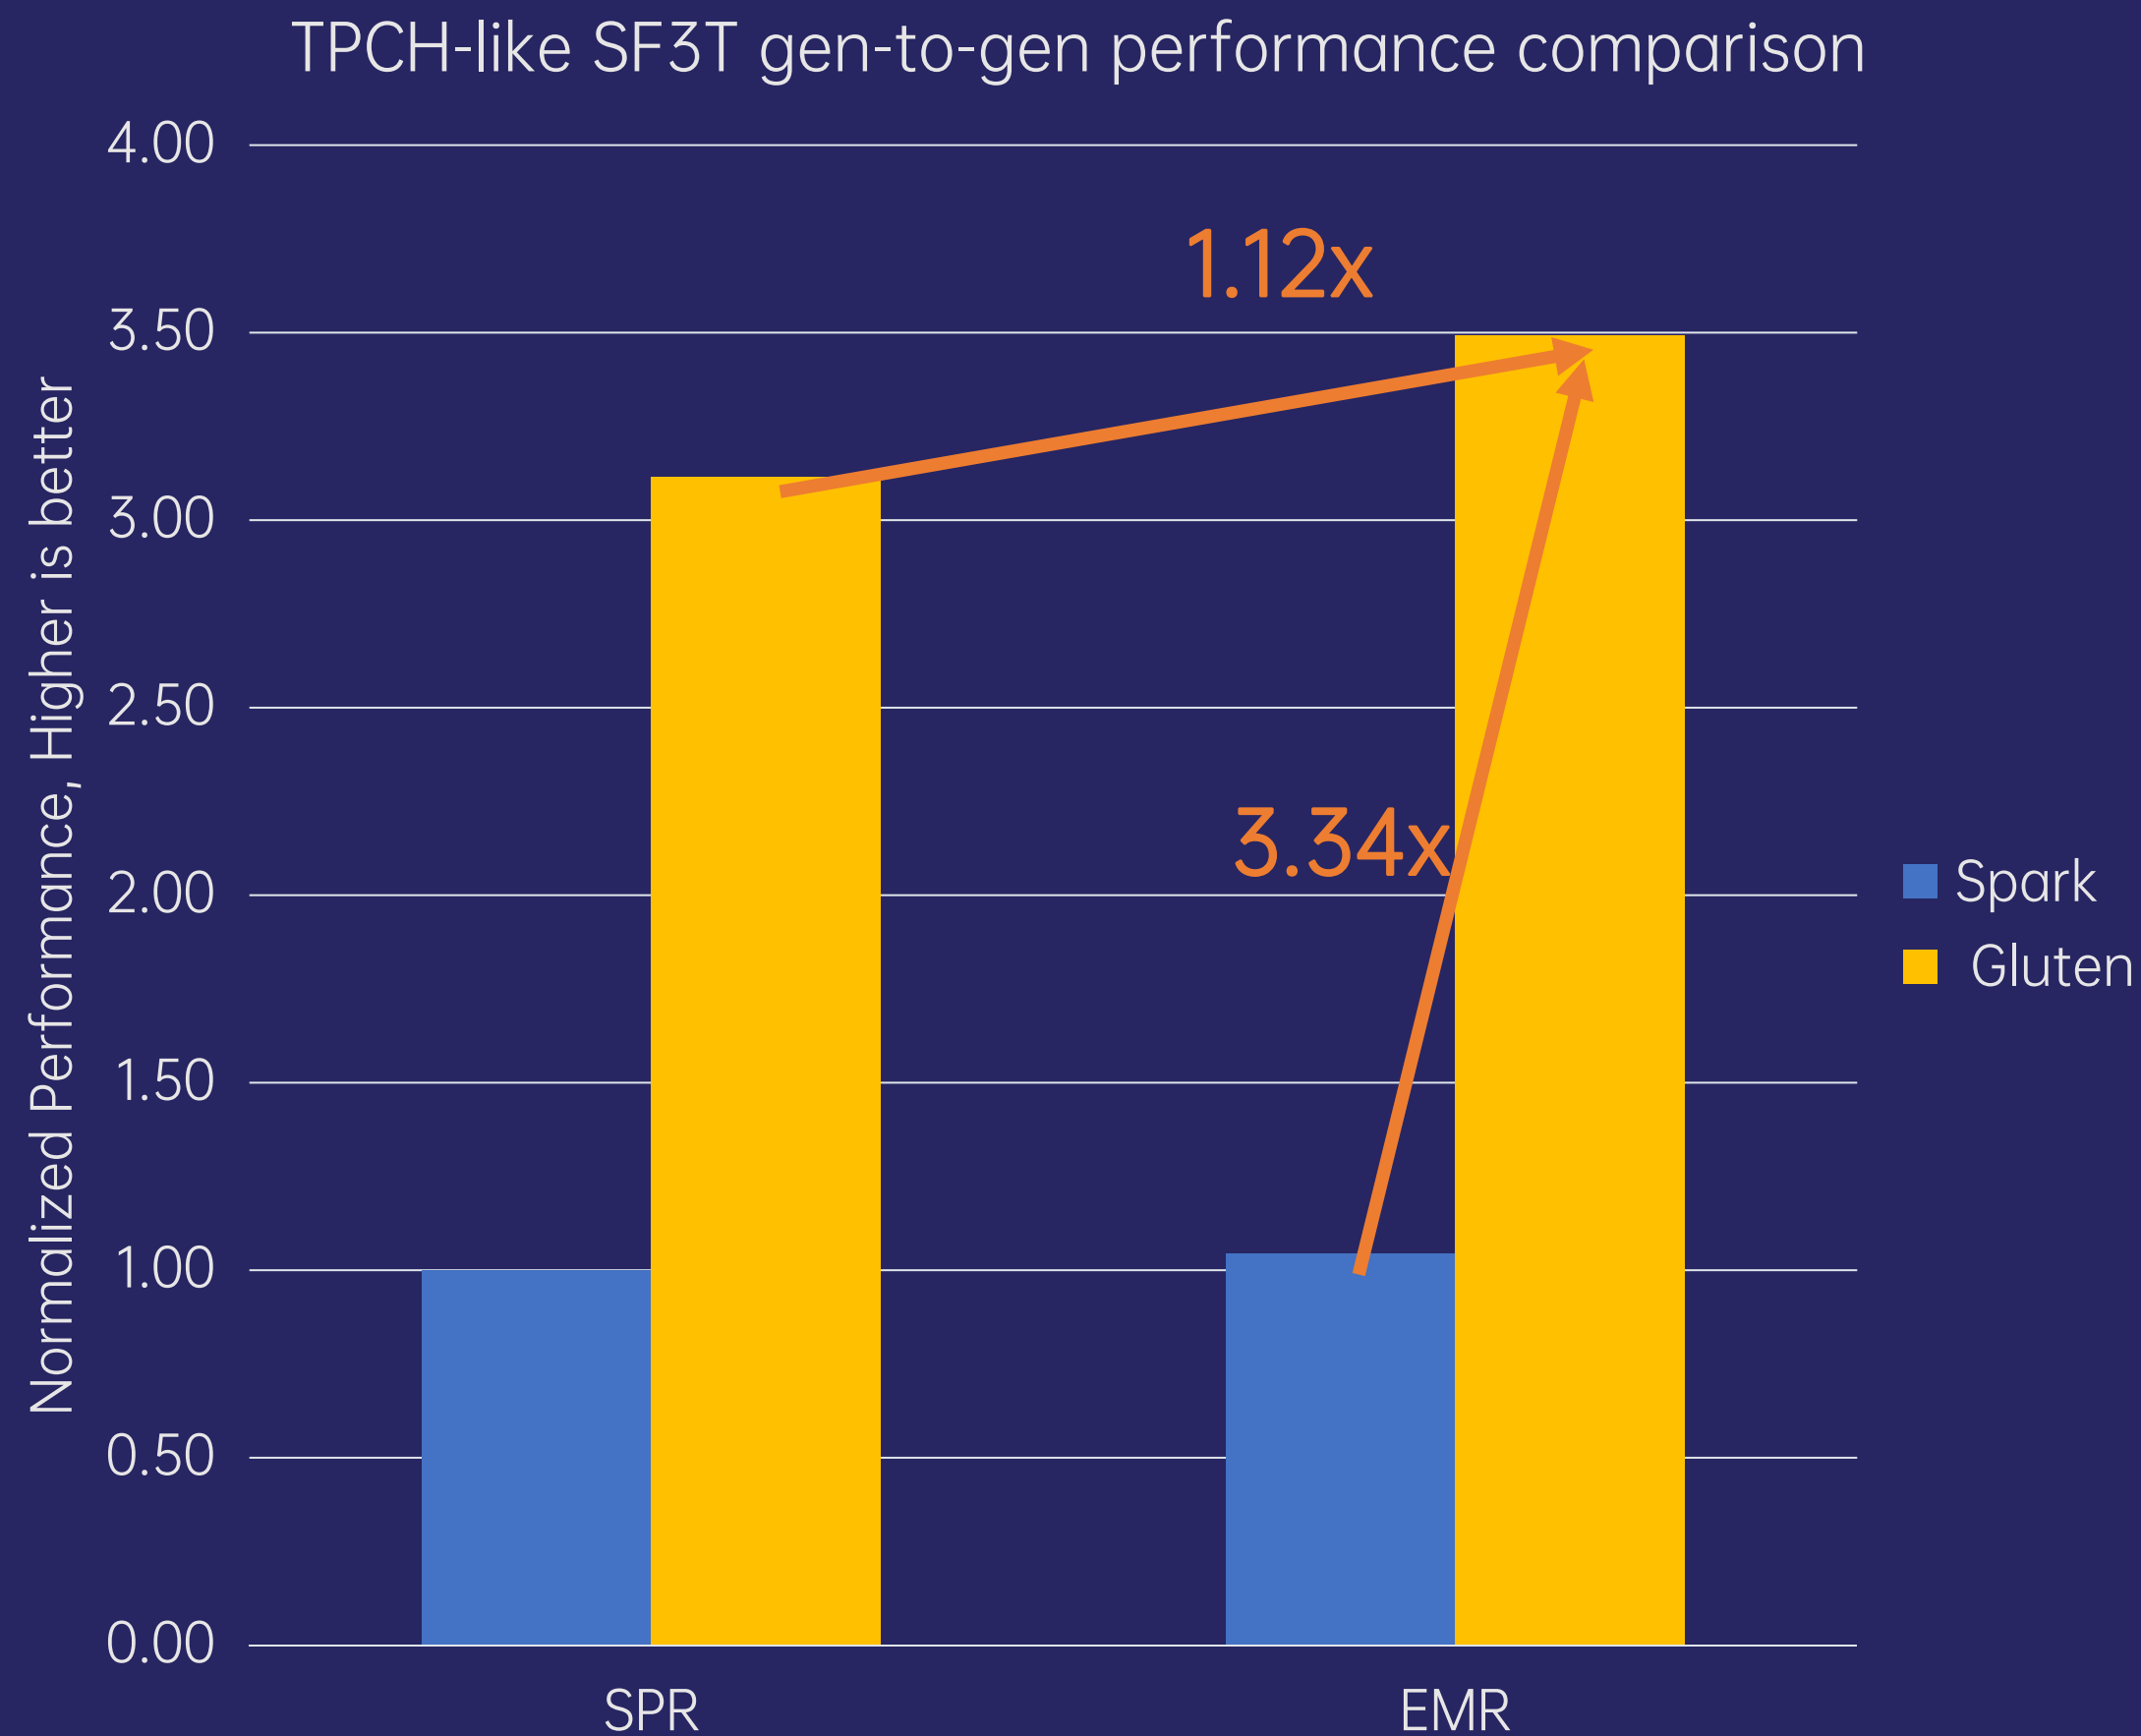
*: Spark Functions use Spark3.2 as the baseline

** : Data source includes localfs, hdfs, S3, GCS, Hive, abfs, deltalake, iceberg, hudi, alluxio

***: File Format includes parquet, orc, text(hive), csv, json

Not supported cases will fallback to Vanilla Spark

GLUTEN + VELOX PERFORMANCE on Intel EMR



- Gluten on Intel[®] Xeon[®] 8592+ vs 8480+: 1.12x perf TPCH Like SF3T, 1.07x perf TPCDS Like SF3T
- Spark on Intel[®] Xeon[®] 8592+ vs 8480+: 1.04x perf TPCH Like SF3T, 1.09x perf TPCDS Like SF3T
- Gluten vs Spark on Intel[®] Xeon[®] 8592+: 3.34x perf TPCH Like SF3T, 3.02x perf TPCDS Like SF3T

*Performance varies by use, configuration and other factors. See backup for configuration details

GLUTEN CUSTOMER SUCCESS STORIES

Alibaba EMR 5.11.1

Alibaba EMR has taken a significant stride by embracing the innovative Gluten + Velox project for their **Alibaba EMR 5.11.1** product, unveiled in April 2023. This remarkable integration marks a pioneering step as the first choice of CSP customers. The incorporation of Gluten yields remarkable benefits, enabling the product to not only achieve a remarkable **2x speedup** in TPC-DS-like workloads when compared to Vanilla Spark's Java engine but also seamlessly synergize with **Apache Celeborn** as 1st Remote Shuffle Service support within Gluten framework.

BONC BEH

BONC has adopted Gluten as a part of their key features in their latest BEH product. With Gluten support, BEH has achieved a **2.91x speedup** comparing to upstream Spark and even more with **Intel® QAT** support, the produce can reach an **extra 8% performance gain** as well. BONC's Big Data solution guide has been published at Intel website.

MeiTuan Datawarehouse

Meituan has adopted Gluten + Velox for their data warehouse solution to achieve data vectorization. They actively contributed to Gluten community and achieved **20000+ ETL** tasks, **40% resource saving**, as well as **13% execution time reduction**. They highlighted the excellent collaboration and results on Intel IPDC Summit 2023, publish one article in weixin and is in the progress to extend the adoption to more tasks.

More customer success stories are coming ...

GLUTEN PUBLICATIONS

No.	Title	Authors	Type	Publish Date	Where	Links	Sort by Publish Date
1	Gluten: A Middle Layer to offload Spark SQL to Native Engines for Execution Acceleration	Weiting Chen	Video	07/20/2022	Data AI Summit 2022	https://www.youtube.com/watch?v=0Q6gHT_N-1U&ab_channel=Databricks	
2	Best Exploration of Columnar Shuffle Design	Binwei Yang Rong Ma	Video	07/23/2022	Data AI Summit 2022	https://www.youtube.com/watch?v=RICMoj00j1A&ab_channel=Databricks	
3	Accelerate Spark SQL Queries with Gluten	Weiting Chen	Article	09/02/2022	medium.com	https://medium.com/intel-analytics-software/accelerate-spark-sql-queries-with-gluten-9000b65d1b4e	
4	Introducing Velox: An open source unified execution engine	Meta	Article	03/09/2023	engineering.fb.com	https://engineering.fb.com/2023/03/09/open-source/velox-open-source-execution-engine/	
5	提速 Spark SQL 2 倍， GLUTEN 向量化引擎原理剖析	Kyigence	Article	03/30/2023	CSDN	https://zhuanlan.zhihu.com/p/617944074	
6	Intel Participation in the Velox Project	Dave Cohen	Video	06/23/2023	Presto Foundation	https://www.youtube.com/watch?v=yZ8F1vWqFXw&ab_channel=PrestoFoundation	
7	Gluten + Celeborn: 让 Native Spark 拥抱 Cloud Native	Alibaba EMR Weiting Chen Yuan Zhou	Article	07/10/2023	CSDN ali developer cloud	https://blog.csdn.net/weixin_45906054/article/details/131651065	
8	The Gluten Open-Source Software Project: Modernizing Java-based Query Engines for the Lakehouse Era	Intel CESG Binwei Yang Weiting Chen	Paper	09/27/2023	CDMS'23	https://ceur-ws.org/Vol-3462/CDMS8.pdf	
9	Apache Spark—大数据计算引擎	BIGO	Article	11/14/2023	CSDN	https://blog.csdn.net/m0_70952941/article/details/134396816	
10	采用 AVX-512 指令和英特尔® QAT 加速器优化的东方国信大数据解决方案	BONC	Article	11/23/2023	intel.cn	https://www.intel.cn/content/www/cn/zh/artificial-intelligence/analytics/bonc-big-data-solutions-optimized-avx512-and-qat.html	
11	Gluten: Double Spark performance	Yuan Zhou	Video	11/18/2023	Data AI Con 2023	https://www.slidestalk.com/slidestalk/71777?video	
12	Spark Committer 深度解读: Apache Spark Native Engine(Chinese version) Apache Spark Native Engine(English version)	Netease	Article	12/04/2023	Zhihu.com	in Chinese: https://zhuanlan.zhihu.com/p/670297787 in English: https://medium.com/@KyububiApache/apache-spark-native-engine-3e1060567ed0	
13	Apache Gluten Status & Plan	Binwei Yang	Video	04/06/2024	VeloxCon	https://www.youtube.com/watch?v=H7L5W6Vio3U&list=PLJvBe8nQAEsEBSOUY0IRFVZr2_YeHYkUR&index=8	
14	Accelerating Spark at Microsoft using Gluten & Velox	Microsoft	Video	04/06/2024	VeloxCon	https://www.youtube.com/watch?v=7pXOAJiSITYs&ab_channel=VeloxCon	
15	Velox at IBM	IBM	Video	04/06/2024	VeloxCon	https://youtu.be/npoEudB5nPo?si=hToh-acObN3miM1Q	
16	Unlocking Data Query Performance	Pinterest	Video	04/06/2024	VeloxCon	https://www.youtube.com/watch?v=pQ4bMyXXLss&list=PLJvBe8nQAEsEBSOUY0IRFVZr2_YeHYkUR&index=10&t=3s&ab_channel=VeloxCon	
17	Native execution engine for Fabric Spark	Microsoft	Article	05/17/2024	Microsoft.com	https://learn.microsoft.com/en-us/fabric/data-engineering/native-execution-engine-overview?tabs=sparksq	
18	Best Practice of Gluten + Velox in Meituan's Production Environment	Meituan	Article	06/21/2024	weixin	https://mp.weixin.qq.com/s/VvmhQi8YMsm0P5xYoiGEZQ	

GLUTEN CROSS PROJECTS COLLABORATIONS

Apache Gluten(incubating) + Velox

Apache Gluten(incubating) + ClickHouse

Apache Gluten(incubating) + Apache Arrow

Apache Gluten(incubating) + Apache Kyuubi in [PR#5800](#)

Apache Gluten(incubating) + Apache Celeborn in [PR#1528](#)

Apache Gluten(incubating) + Apache Uniffle(incubating) in [PR#950](#)

and more collaboration is coming!!!

APACHE GLUTEN (incubating)



Initiated as an open-source project by Intel and Kylligence in 2022, has become an incubating project under Apache Software Foundation(ASF) since Jan. 2024.

- **Offloading** performance-critical data processing tasks to **native library**.
- Converting Spark physical plans into **Substrait** plans for native execution.
- Seamless switching between native backends, with current support for **Velox** and **ClickHouse**.
- Leveraging Spark's robust distributed framework.
- Managing data sharing and optimizing memory mgmt. between JVM and native.
- Defining clear and unified JNI interfaces for native libraries.
- Expanding support to encompass native accelerators.
- Diverse community engagement in Apache Gluten from over **25+** companies' contributions.

RECAP

Gluten fundamentally changes the Spark-SQL performance profile by enabling it to offload computations to accelerator libraries

With Gluten,

- Customer can take performance benefit from native libraries to speedup Spark SQL
- Spark Developers can add various native library's support by implementing the JNI interfaces
- Library Developers can reproduce a stage and optimize for it
- Accelerator Developers can develop toolkit decoupled from Spark

Join with US!

Please visit <https://github.com/apache/incubator-gluten> or <https://gluten.apache.org> for more information.



Thanks

Weiting Chen weiting.chen@intel.com



主题字体 THEME FONTS : OPPOSans

下载链接 Download link : <https://fontmeme.com/fonts/oppo-sans-font/>

主题色彩 THEME COLORS :

