



## AtlasGraph: A High-Performance Graph Database Solution for Complex Data Challenges

Haizhi Technology Group Co., Ltd. is a leading big data technology and service company in China with over 1000 employees. Haizhi contributed over a decade into big data mining and analysis. Haizhi is well-equipped to provide comprehensive services, including 24/7 product deployment, training, implementation to its valuable clients.



Beijing Headquarter



Wuhan Support Center



Shenzhen R&D Center



Shanghai R&D Center



### ***Facilitate the digital transformation across all Industries***

Advanced enterprise-level big data solutions designed for practical applications, combining professionalism with user-friendly interfaces.



### ***Expert in enterprise-level data solutions***

Delivering comprehensive data intelligence product solutions and long-term services tailored to over 70 banks, securities firms, and insurance enterprises.

## ***Academician-led HPC Expert Workstation Launched***

Xinhua  
[http://www.xinhuanet.com/tech/2021-03/26/c\\_1127257174.htm](http://www.xinhuanet.com/tech/2021-03/26/c_1127257174.htm)  
**Haizhi High-Performance Graph Computing Academician Expert Workstation Prepares for Launch - Xinhuanet**  
Web: 26 Mar 2021 - On March 25, Haizhi High-Performance Graph Computing Academician Expert Workstation was launched, and Zheng Weimin, academician of the Chinese Academy of...

China Education Network  
[https://www.edu.cn/ke\\_yan\\_yu\\_fa\\_zhan/gao\\_xiao\\_cheng\\_guo/](https://www.edu.cn/ke_yan_yu_fa_zhan/gao_xiao_cheng_guo/)  
**Academician Zheng Weimin of Tsinghua University led the establishment of an academician expert workstation to cooperate with enterprise research and development.**  
Web: 13 Feb 2023 - In March 2021, the first high-performance graph computing academician expert workstation team settled in Haizhi Technology, Zheng Weimin served as the chief scienti...



On 25 March 2021, Academician Zheng Weimin, in collaboration with Haizhi Corporation, launched the "Haizhi HPC Academician Expert Workstation" to advance cutting-edge HPC technology, deepen basic technology, and innovate storage, computing, operating systems, and chips. We aim to establish a data intelligence ecosystem emphasizing hardware-software collaboration and cloud-edge integration, positioning themselves as a technology company with core expertise.

Haizhi, in partnership with Tsinghua University, received the prestigious "2021 CCF Science and Technology Award for Technological Progress - Outstanding Achievement" for our innovative AtlasGraph platform. This award celebrates groundbreaking achievements in computer science with significant international impact.

AtlasGraph tackles the growing challenge of analyzing massive graph datasets, a critical need as knowledge graph applications expand. Our platform incorporates breakthroughs in efficient graph data storage, indexing, and replication, enabling high-performance analysis through innovative graph reduction techniques.

This prestigious recognition underscores AtlasGraph's technological leadership and innovation. As a powerful domestic solution, AtlasGraph offers a compelling alternative to foreign products, promoting technological independence in this critical field.



# AtlasGraph Earns Top Honor: 2022 CIE Science and Technology Award

Haizhi's AtlasGraph platform has received the **prestigious First Prize for Scientific and Technological Progress** in the 2022 China Institute of Electronics (CIE) Science and Technology Awards, the highest honor in China's electronics and information technology sector.

Developed in collaboration with Beijing University of Posts and Telecommunications, Ant Group, and China Mobile Information Technology Co., Ltd., AtlasGraph was recognized for its groundbreaking advancements in large-scale complex heterogeneous graph data analysis. The judging committee, composed of top academics, highlighted **the project's international leadership in key areas like heterogeneous graph modeling and large-scale graph learning.**

科技进步一等奖：17项				
序号	项目名称	主要完成人	主要完成单位	提名者
12	大规模复杂异构图数据智能分析技术与规模化应用	石川, 周俊, 杜军平, 陶涛, 王啸, 张志强, 杨娟, 朱俊雄, 王艺霏, 吴斌, 杨成, 尚晶, 李东, 白婷, 何慧梅	北京邮电大学, 蚂蚁科技集团股份有限公司, 中移动信息技术有限公司, 北京海致星图科技有限公司	北京邮电大学
13	电力基础设施信息物理跨域攻击防御关键技术及应用	张涛, 徐文渊, 石聪聪, 费稼轩, 宋宇波, 孙彦斌, 周小明, 陈雪鸿, 王琦, 李霞, 高鹏, 潘锴, 陈牧, 黄秀丽, 乐翔	国网智能电网研究院有限公司, 浙江大学, 南京南瑞信息通信科技有限公司, 广州大学, 国家工业信息安全发展研究中心, 国网辽宁省电力有限公司, 宁波和利时信息安全研究院有限公司, 东南大学, 国家信息技术安全研究中心	国网智能电网研究院有限公司
14	智能语音信息处理关键技术及应用	杨俊安, 方四安, 刘辉, 殷兵, 王一, 柳林, 刘海波, 王龙, 刘方正, 方磊, 陈雷	中国人民解放军国防科技大学, 合肥讯飞数码科技有限公司	中国人民解放军国防科技大学

AtlasGraph addresses critical challenges in processing and analyzing massive graph datasets. The platform introduces innovative solutions for representation learning, semantic recommendations, and risk management, effectively enabling knowledge discovery from complex data.

This prestigious award underscores AtlasGraph's exceptional innovation and impact. The platform's achievements include 43 granted patents, 51 top-tier publications, 2 international competition wins, and contributions to graph computing standards.



**Gartner**, a globally recognized authority in technology research and advisory, consistently shapes the global tech landscape with its rigorous analysis, forward-looking predictions, and expert insights. Haizhi is proud to be recognized by Gartner in multiple market guides, highlighting our position as a leading provider of data analytics solutions:

- **May 2023:** Named a Representative Vendor in the "Market Guide for Analytics Platforms, China" for our contributions to the evolving data analysis market.
- **June 2023:** Featured in the "Market Guide for Artificial Intelligence Software, China," demonstrating our expertise in this rapidly growing field.
- **August 2023:** Highlighted in the "Market Guide for Graph Database Management Systems" for our self-developed, high-performance graph database, AtlasGraph DB, showcasing Chinese innovation on a global stage.

**These recognitions underscore AtlasGraph commitment to delivering cutting-edge data analytics solutions that empower businesses worldwide.**

# AtlasGraph Achieves World Record Performance in LDBC SNB Interactive Benchmark



□ The Linked Data Benchmark Council (LDBC) is the leading authority in graph database benchmarking and standardization.




□ LDBC SNB Interactive benchmark, **comparable to the renowned TPC-C test for relational databases**, evaluates online transaction processing (OLTP) performance using social network data.

□ AtlasGraph, leveraging its innovative in-memory graph compression and dynamic graph analysis optimization techniques, **achieved the highest score ever recorded** in the LDBC SNB Interactive benchmark. This achievement demonstrates AtlasGraph's unmatched speed and efficiency in handling complex, real-time queries on massive datasets.


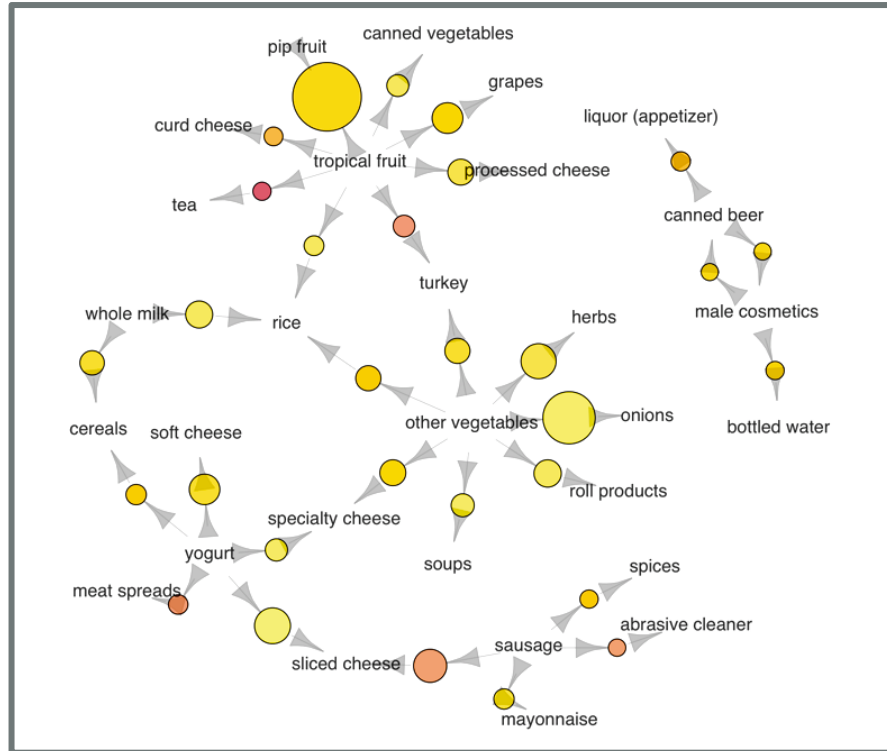
## AUDITED RESULTS USING AN IMPERATIVE LANGUAGE

Benchmark setup	SF	Hardware	Throughput	Documents
System: <b>AtlasGraph 3.0.0</b>	30	AWS r5d.12xlarge, 384GB RAM, 48×Intel Xeon Platinum 8259CL vCPUs	37,631.25 ops/s	Full disclosure report
Test sponsor: <b>StarGraph</b>				Executive summary
Date: <b>2023-12-25</b>	100	AWS r5d.12xlarge, 384GB RAM, 48×Intel Xeon Platinum 8259CL vCPUs	48,764.08 ops/s	Signatures
Queries implemented in: <b>Rust stored procedures</b>				Supplementary package
Systems cost: <b>2,032,461 RMB</b>	300	AWS r5d.12xlarge, 384GB RAM, 48×Intel Xeon Platinum 8259CL vCPUs	48,311.63 ops/s	
System: <b>GraphScope Flex 0.23.0</b>	30	AWS r5d.12xlarge, 374GB RAM, 48×Intel Xeon Platinum 8175M vCPUs	33,180.87 ops/s	Full disclosure report
Test sponsor: <b>Alibaba DAMO Academy</b>				Executive summary
Date: <b>2023-07-13</b>	100	AWS r5d.12xlarge, 374GB RAM, 48×Intel Xeon Platinum 8175M vCPUs	33,625.36 ops/s	Signatures
Queries implemented in: <b>C++ stored procedures</b>				Supplementary package
Systems cost: <b>99,236.04 USD</b>	300	AWS r5d.12xlarge, 374GB RAM, 48×Intel Xeon Platinum 8175M vCPUs	33,261.38 ops/s	

# Graph Databases: Trends Shaping the Future



## Real World



## Relational Table

Transaction 1	
Transaction 2	
Transaction 3	
Transaction 4	
Transaction 5	
Transaction 6	
Transaction 7	
Transaction 8	

Multi-Context is Preserved with Graph Analytics





## Cloud-Native Architecture

Elastic scalability for efficient resource utilization, high availability, fault tolerance, and low-cost maintenance.



## HTAP Capabilities

Hybrid Transactional/Analytical Processing with a high-performance graph compute engine. 30+ pre-built algorithms and an extensible engine empower complex data mining and machine learning.



## Ease of Use

Designed with open standards in mind, AtlasGraph aligns with openCypher and the evolving ISO GQL standard.



## Real-Time, Large-Scale Graph Processing

Supports trillions of nodes and integrates with a streaming engine for real-time data ingestion and analysis, enabling timely, data-driven decisions.



## High Performance

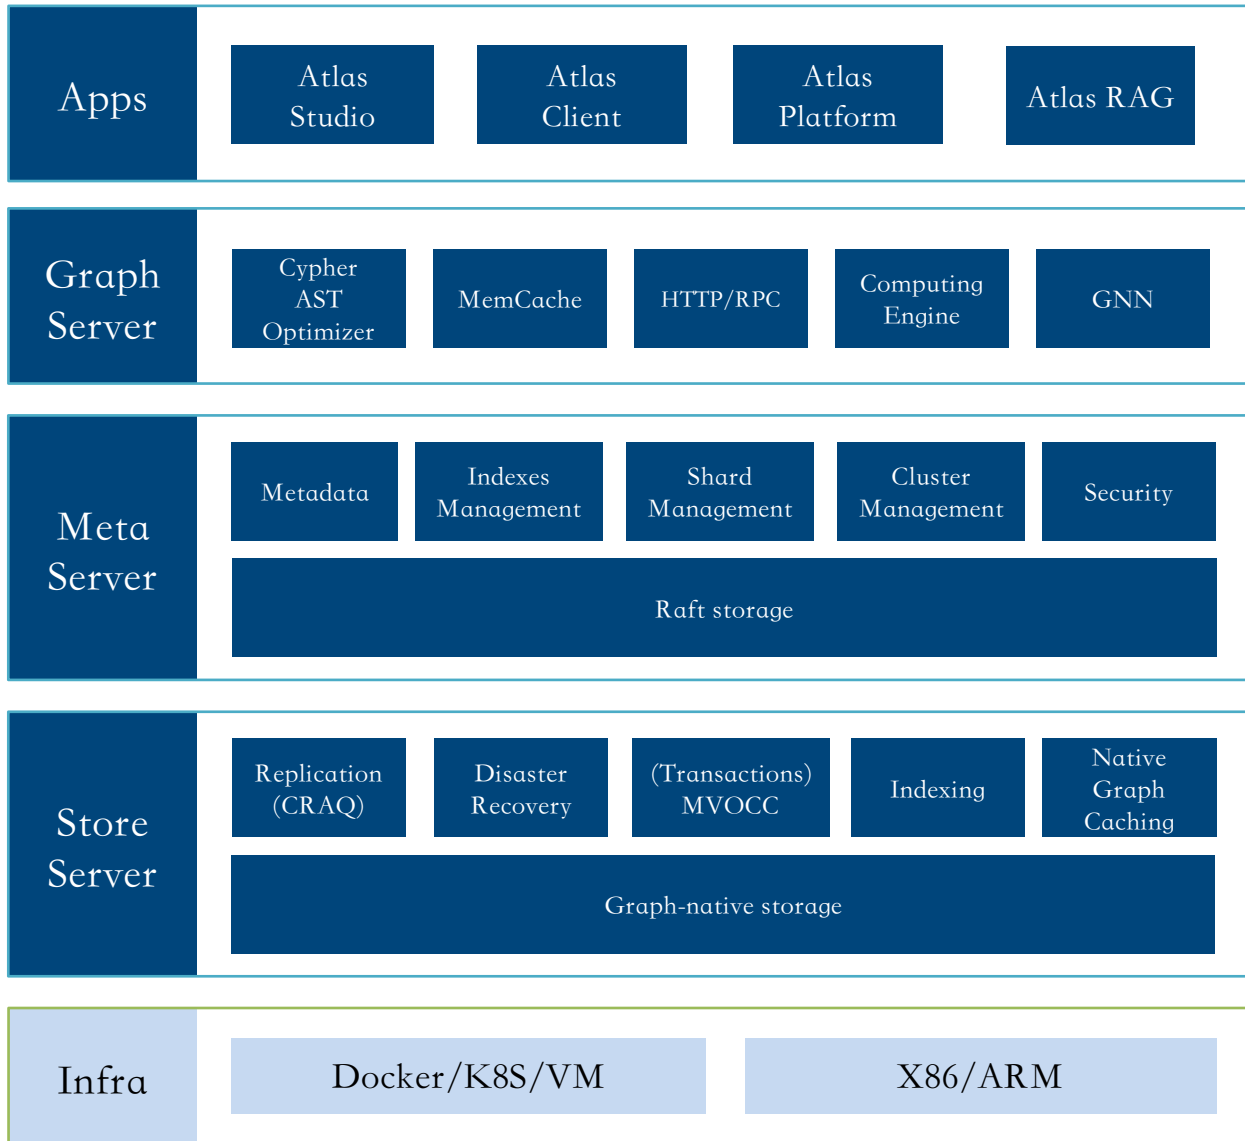
Built on a distributed storage and compute engine using Rust, AtlasGraph incorporates refined memory management and a built-in indexing system for millisecond-level query response times.



## MPP

Distributed storage and parallel computing across large-scale clusters. Shared-nothing architecture ensures separation of storage and compute for optimal performance.

# AtlasGraph: Architecture Overview



Scalability

- Rust-based engine for optimal speed
- Distributed architecture for linear scalability
- Optimized storage and compute for large graphs



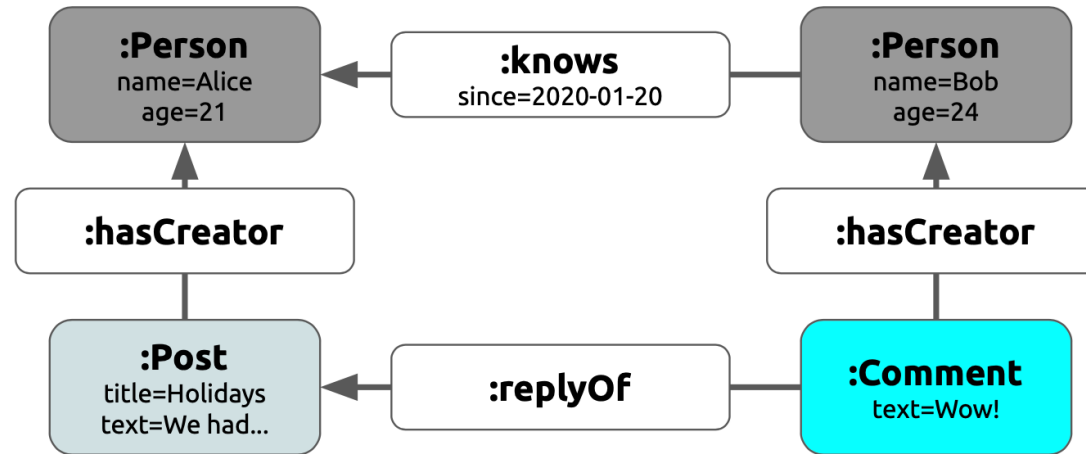
Reliability

- MVCC for transactional consistency
- Multi-replica management for high availability
- Online backups for disaster recovery



Usability

- Code-free analysis with Atlas Platform integration
- Query performance analysis & optimization hints
- Rich library of built-in analytical functions



## Labels

Types (or classes) of vertices and edges

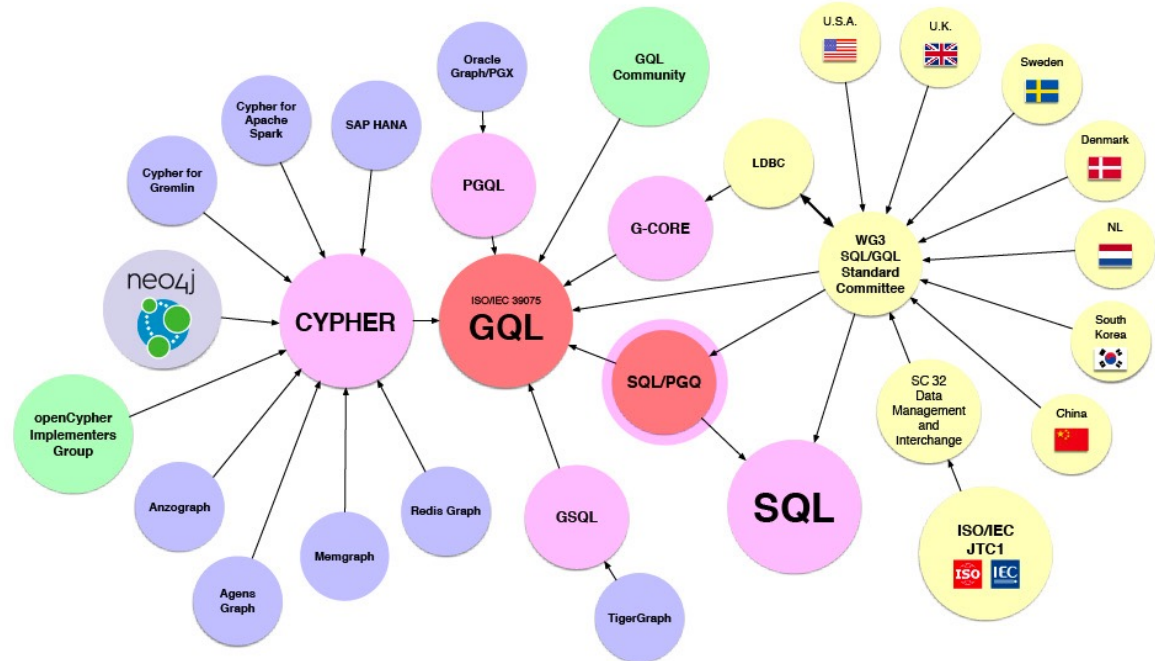
## Properties

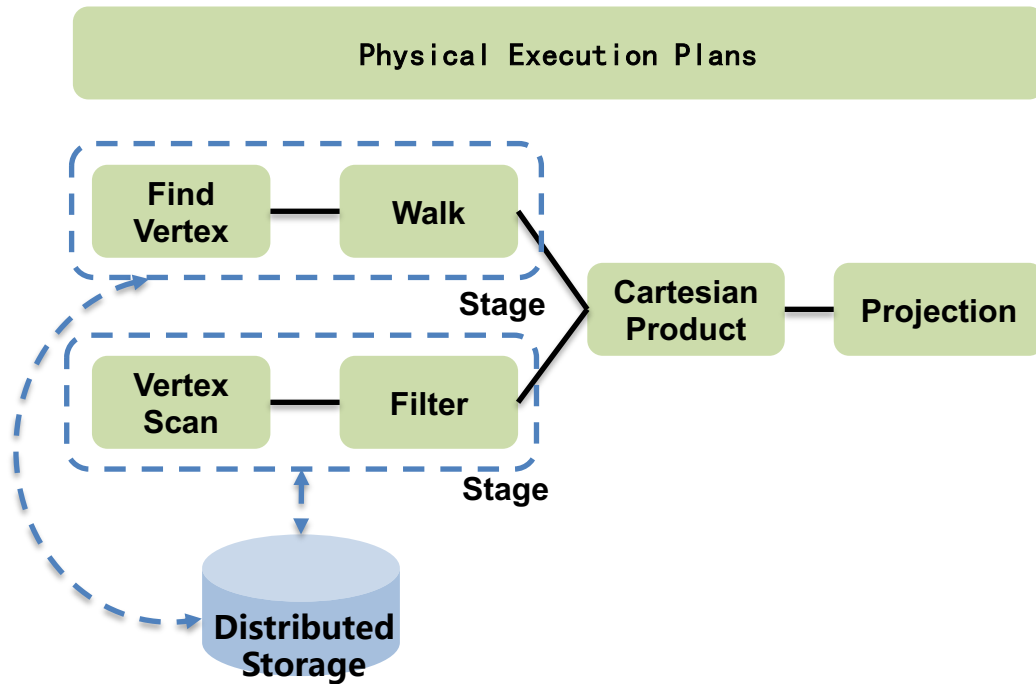
Arbitrary (key,value) pairs where key identifies a property and value is the corresponding value of this property

## SQL-like, Analyst-Friendly, Easy to Use

- Incorporate by reference specifications from SQL/Framework and SQL/Foundation
- Capabilities needed by both SQL/PGQ (Property Graph Queries in SQL) and the GQL standard
- GQL Specific Capabilities

- 2019-09 – 39075 Database Language GQL project approved - this is the start.
- 2021-11 – CD Ballot started.
- 2022-02 – CD Ballot ended and comment resolution started.
- 2023-06 – DIS (Draft International Standard) Ballot starts.
- 2023-11 – DIS ballot ends.
- 2024-04 – International Standard.dard.





01

**Segmented Execution Plans:** Breaks down queries for targeted optimization.

02

**Pushdown to Storage Engine:** Pushes computation closer to the data, minimizing data movement.

03

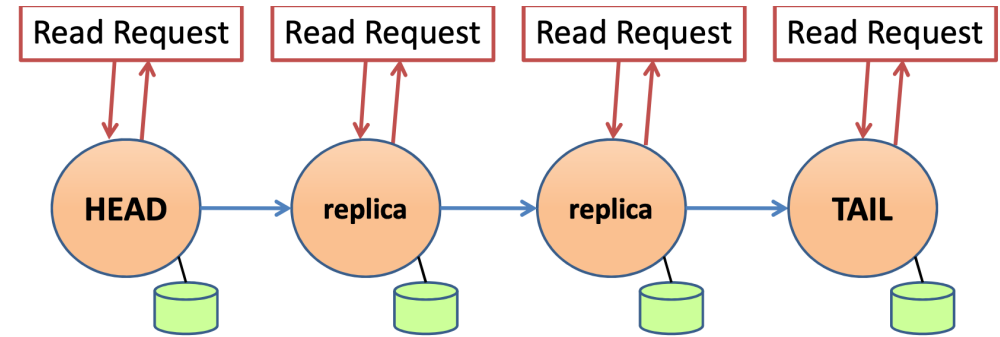
**Pipeline Execution:** Enables parallel processing for optimal resource use and faster results.

## Data High Availability:

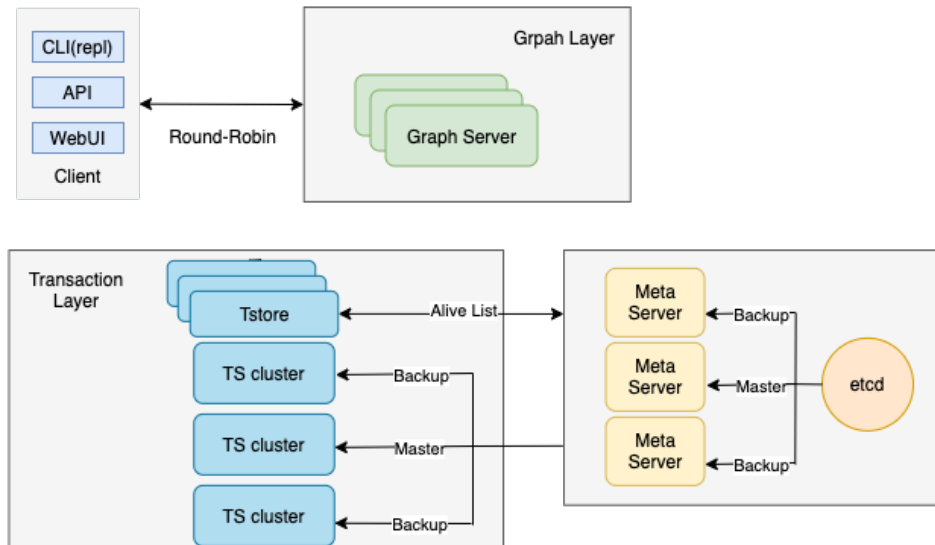
- Chain Replication (CRAQ) for robust data copies.
- Read-write separation for better performance.

## Service High Availability:

- Active-standby architecture for critical services.
- Raft consensus for seamless failover.



Chain Replication



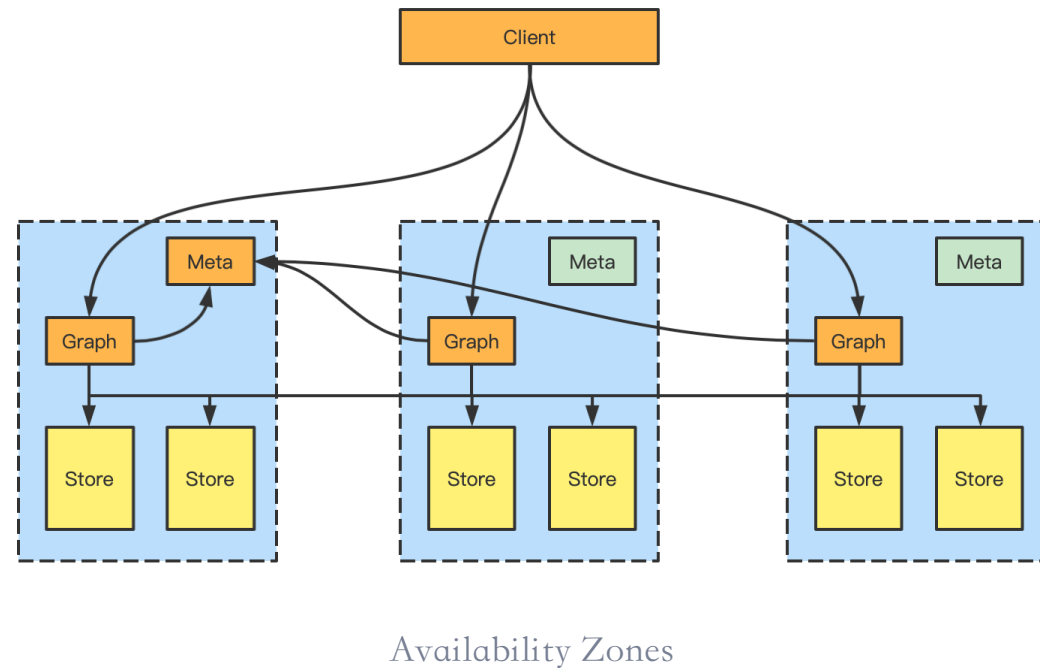
Service High Availability:

## Data Durability:

- Leverages Chain Replication for data replication across availability zones (AZs).
- Ensures data safety and service continuity even in the event of an AZ failure.

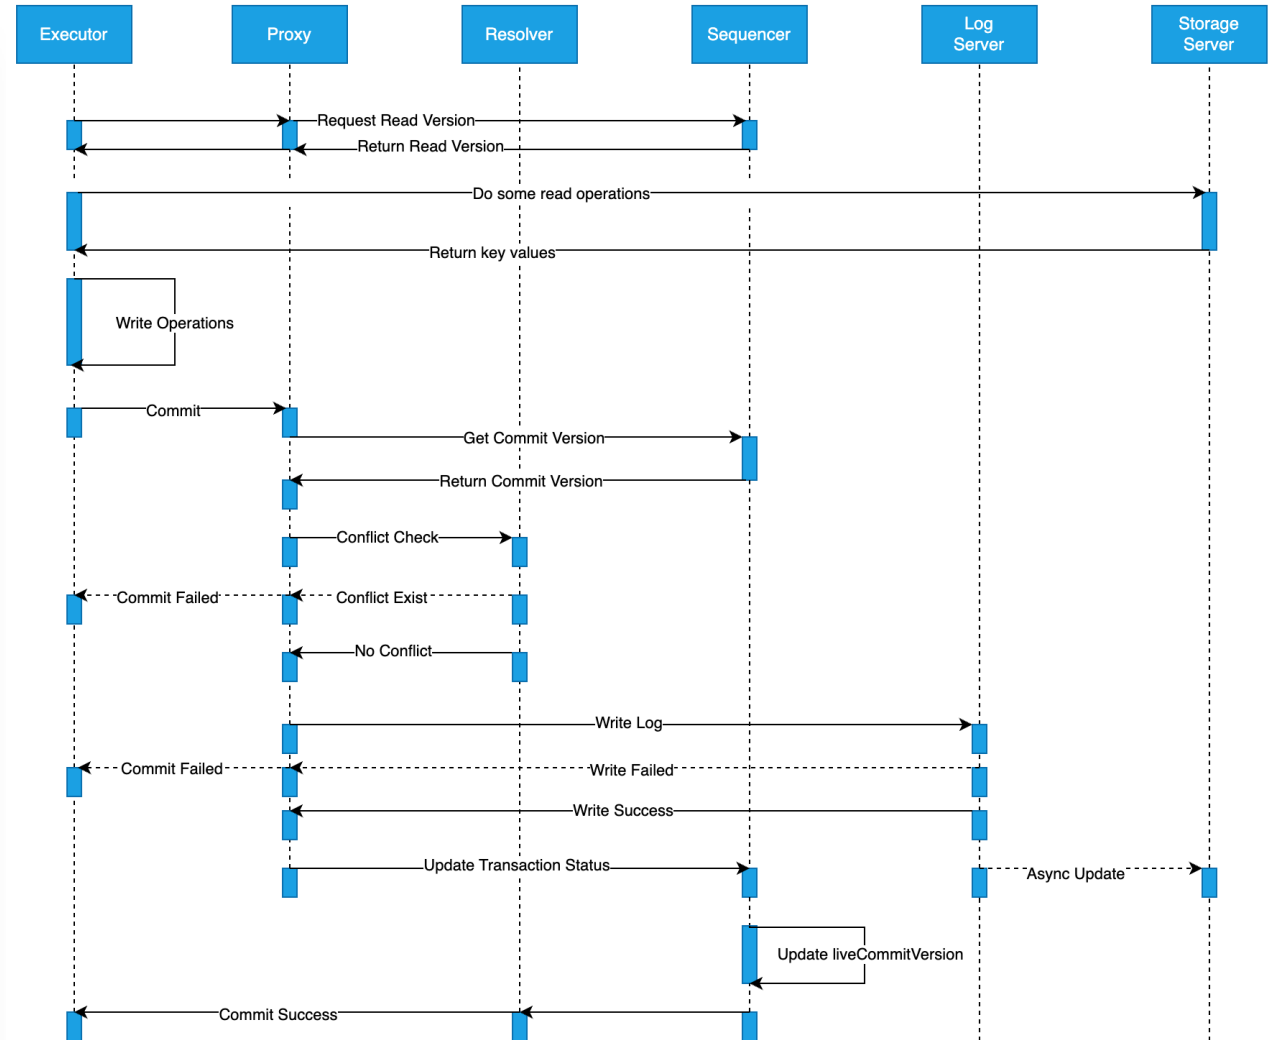
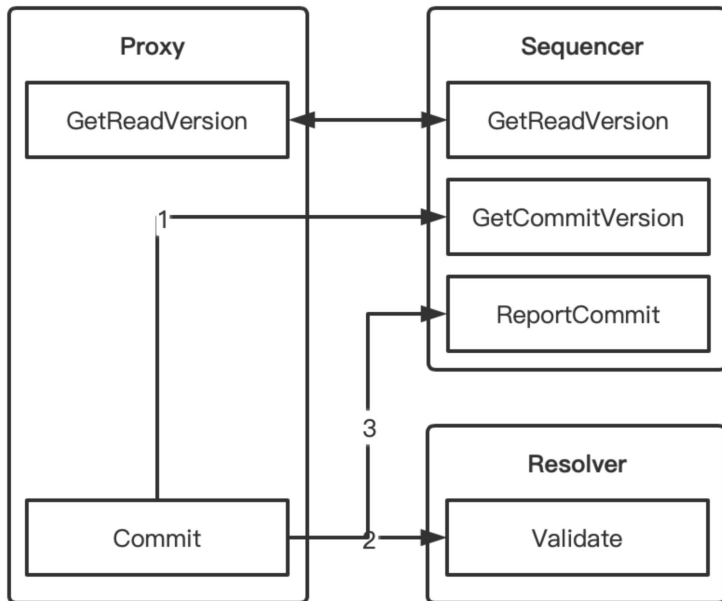
## Automated Failover:

- The system automatically recovers from AZ failures without manual intervention.
- Data consistency is maintained throughout the failover process.



## MVCC-based Transaction Management:

Ensures data consistency without sacrificing analytical performance.



MVOCC Execution Process



## Secure & Efficient Resource Sharing

- Resource Groups: Isolate tenants while sharing a single database instance.
- Granular Control: Allocate CPU priority and memory per resource group.
- Benefits: Predictable performance, efficient resource utilization, enhanced security.



### Resource Isolation

root user

Cluster: atlasgraph

Please enter the resource group name... Q

Add a resource group

Resource group name	CPU Priority	Memory quota (%)	Number of users	operate
default	high	 70%		
high	middle	 30%	1	<a href="#">Edit</a> <a href="#">Delete</a>

A total of 2 resource groups

< 1 > 10 items/page

## Extensive Algorithm Library:

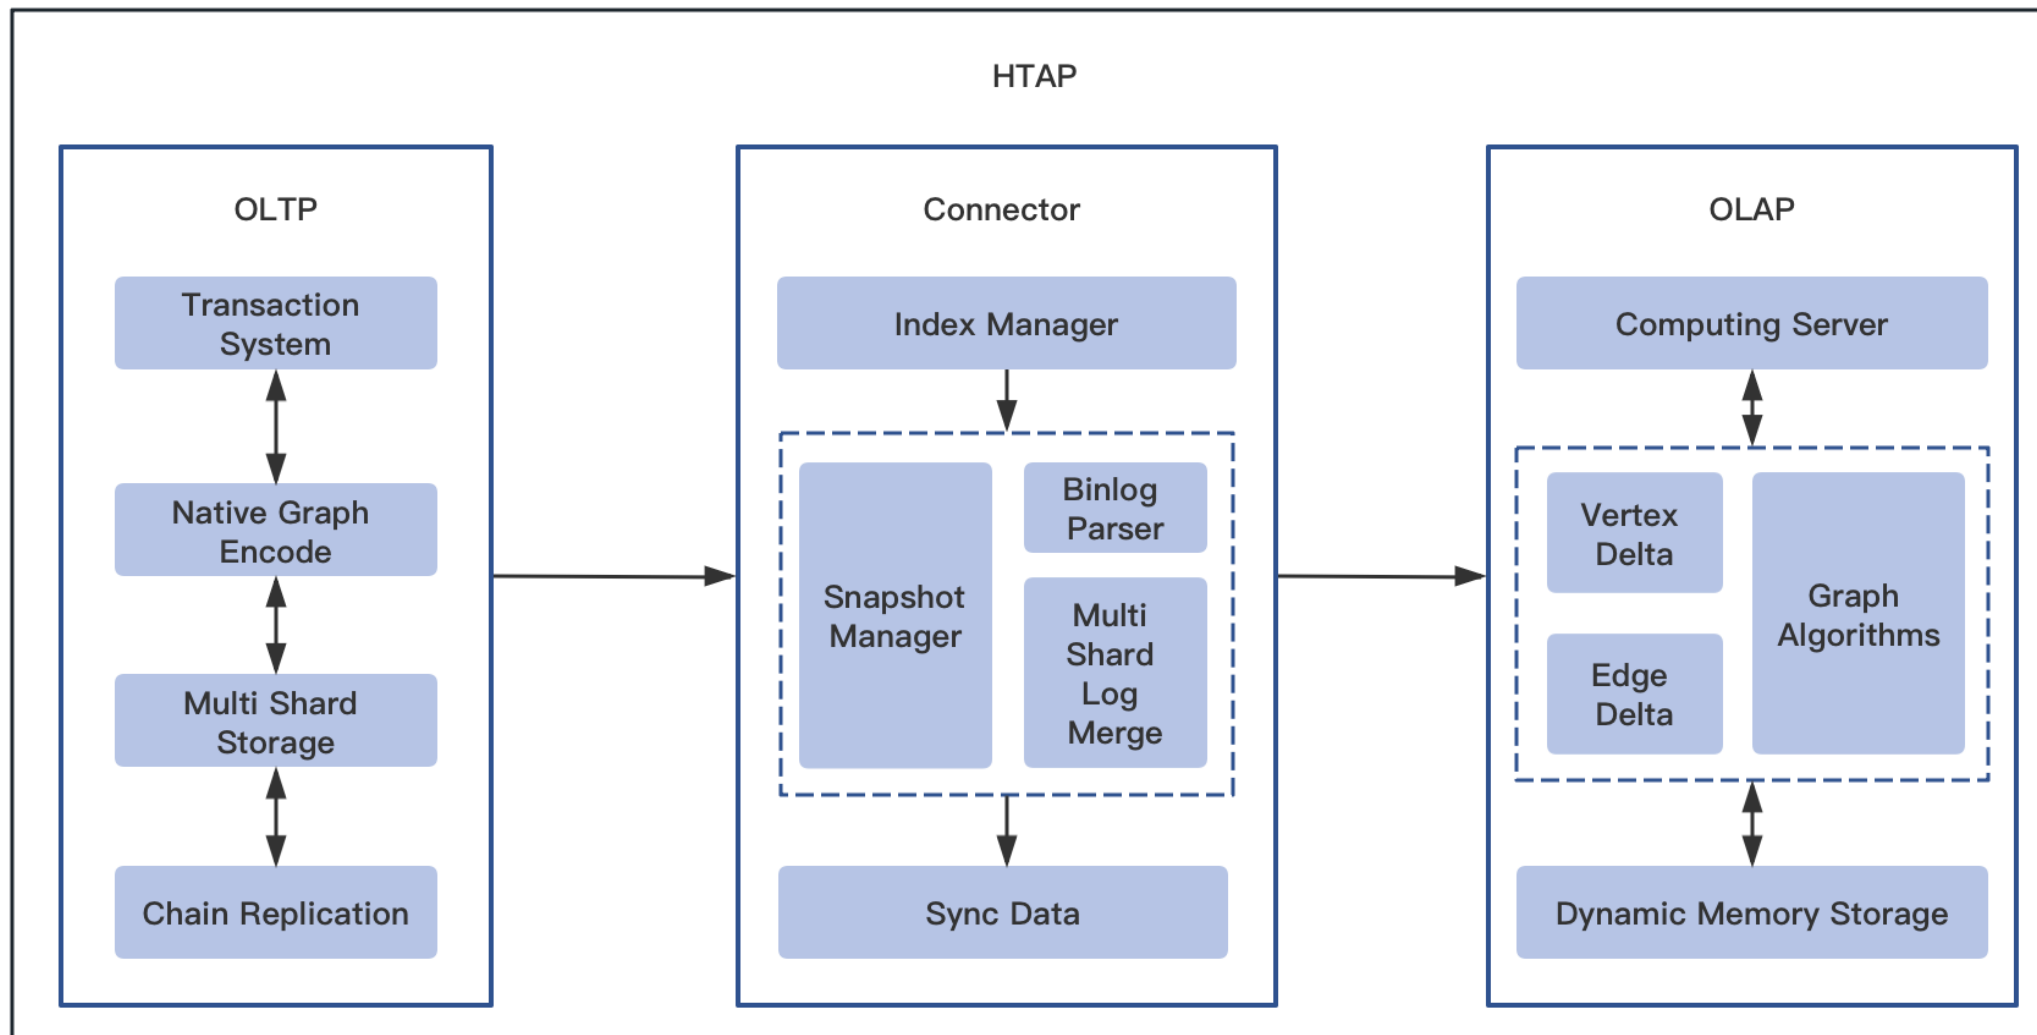
- **30+ Built-in Algorithms:** Covering all common graph algorithms, including pathfinding, community detection, similarity calculations, and more.
- **Advanced Algorithms:** Featuring advanced graph analysis capabilities such as loop detection, chain identification, all-pairs shortest paths, and subgraph pattern recognition.

## Seamless Integration and Flexibility:

- **Cypher Integration:** Easily invoke algorithms directly within your Cypher queries.
- **Subgraph Analysis:** Perform targeted analysis on user-defined subgraphs.
- **Flexible Parameterization:** Fine-tune algorithm behavior to suit your specific needs.

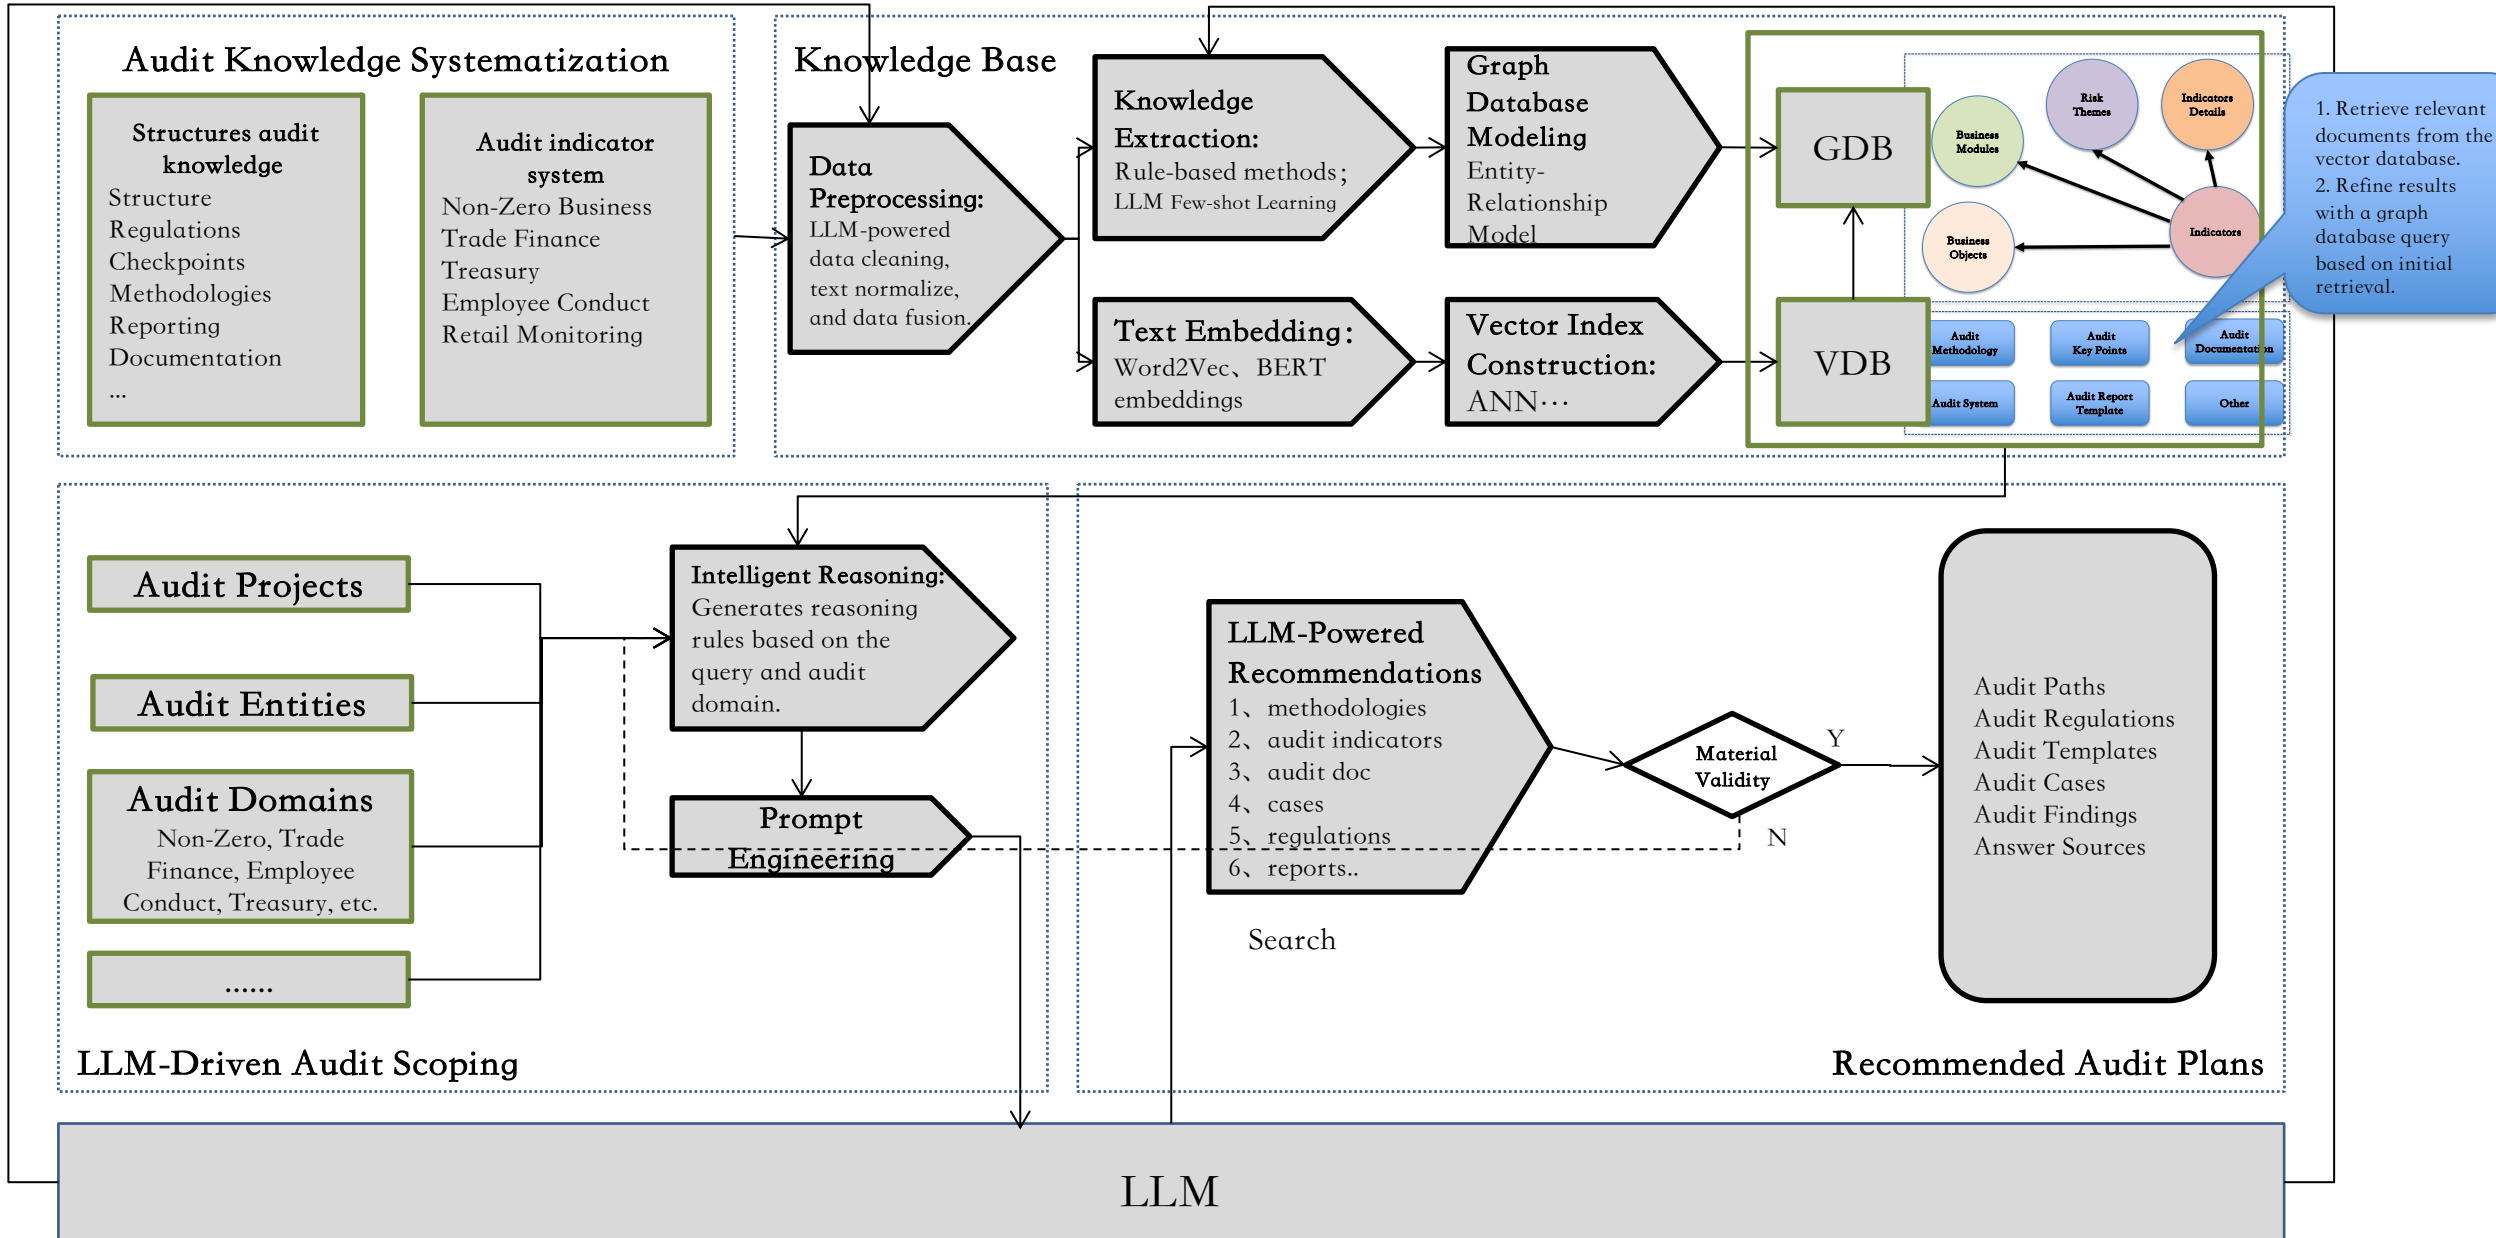
## Optimized for Real-World Deployments:

- **Hybrid Execution Modes:** Seamlessly switch between single-machine and distributed computing for optimal performance in diverse environments.
- **Algorithm-Specific Optimizations:** Hand-tuned implementations for common algorithms deliver superior performance compared to generic approaches.
- **Resource-Efficient Design:** Designed to excel in environments with limited compute resources and network bandwidth.





# GraphAgent: Smarter Auditing with LLM & Graphs





Thank you!  
Q & A