Design of Highly Scalable Graph Database Systems

Ricky and Jamie

Ultipa

BiDEDE '23:https://doi.org/10.1145/3579142.3594293

Node To Tree To Graph

From Single Instance to Horizontal Scalability

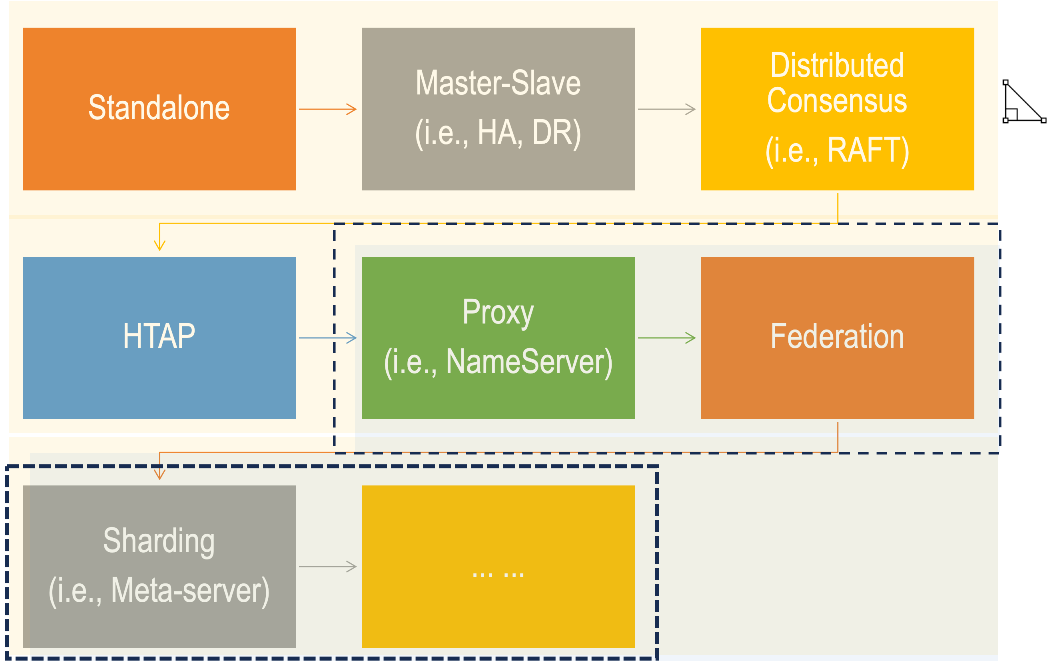

- Stability ↑ ↑ ↑
- Storage ↑ ↑ ↑
- Integration ↑ ↑ ↑ ☆
- Sharding ↑ ↑ ↑ +

Evolution of Distributed Graph Systems

From Single Instance to Horizontal Scalability

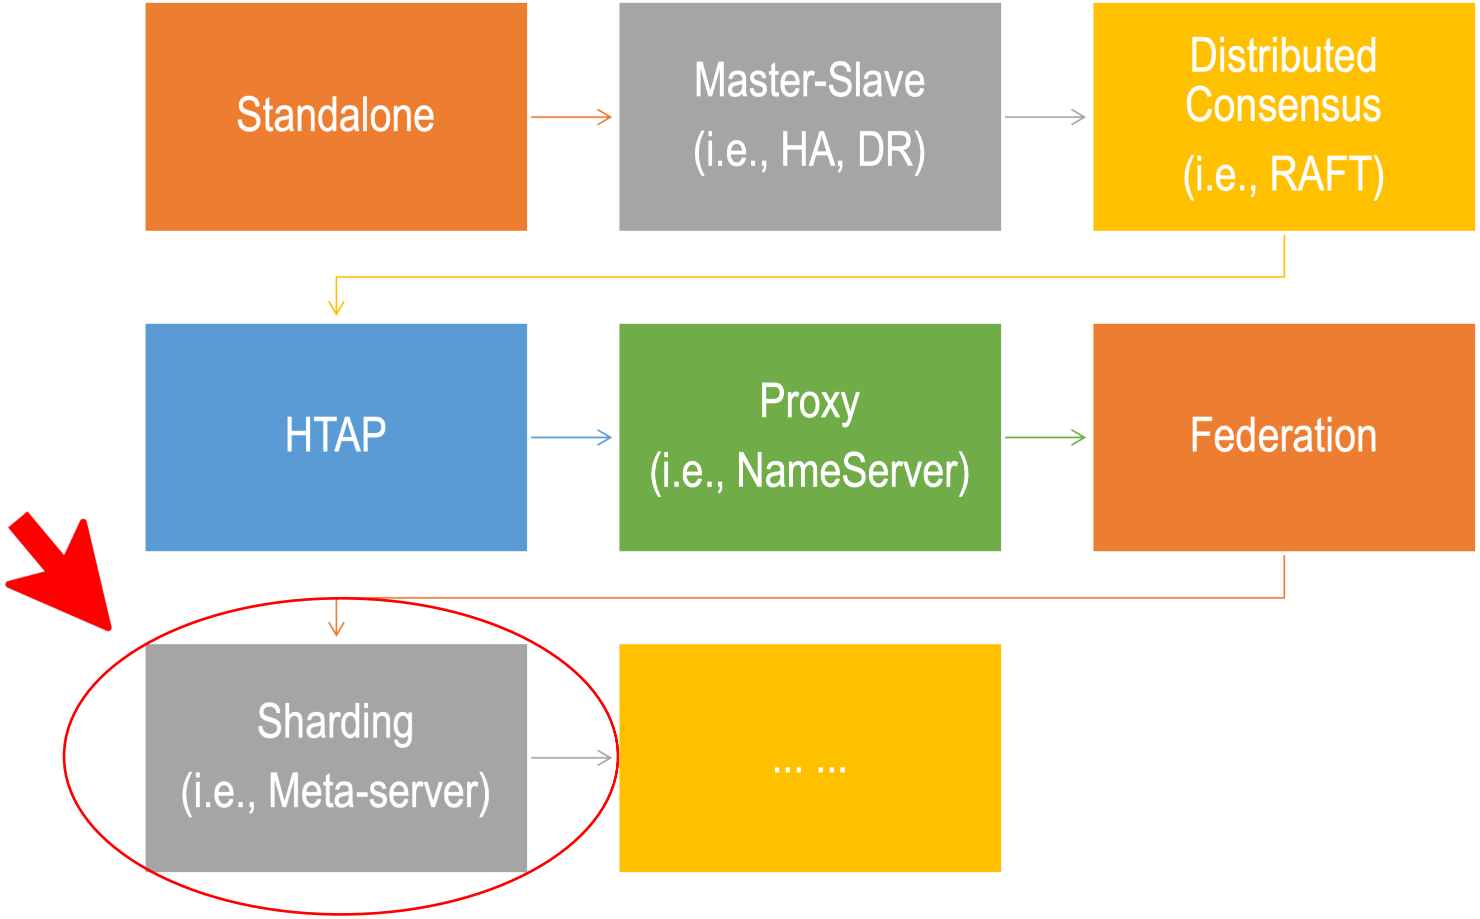

The Magic Quadrant of Graph DBMS

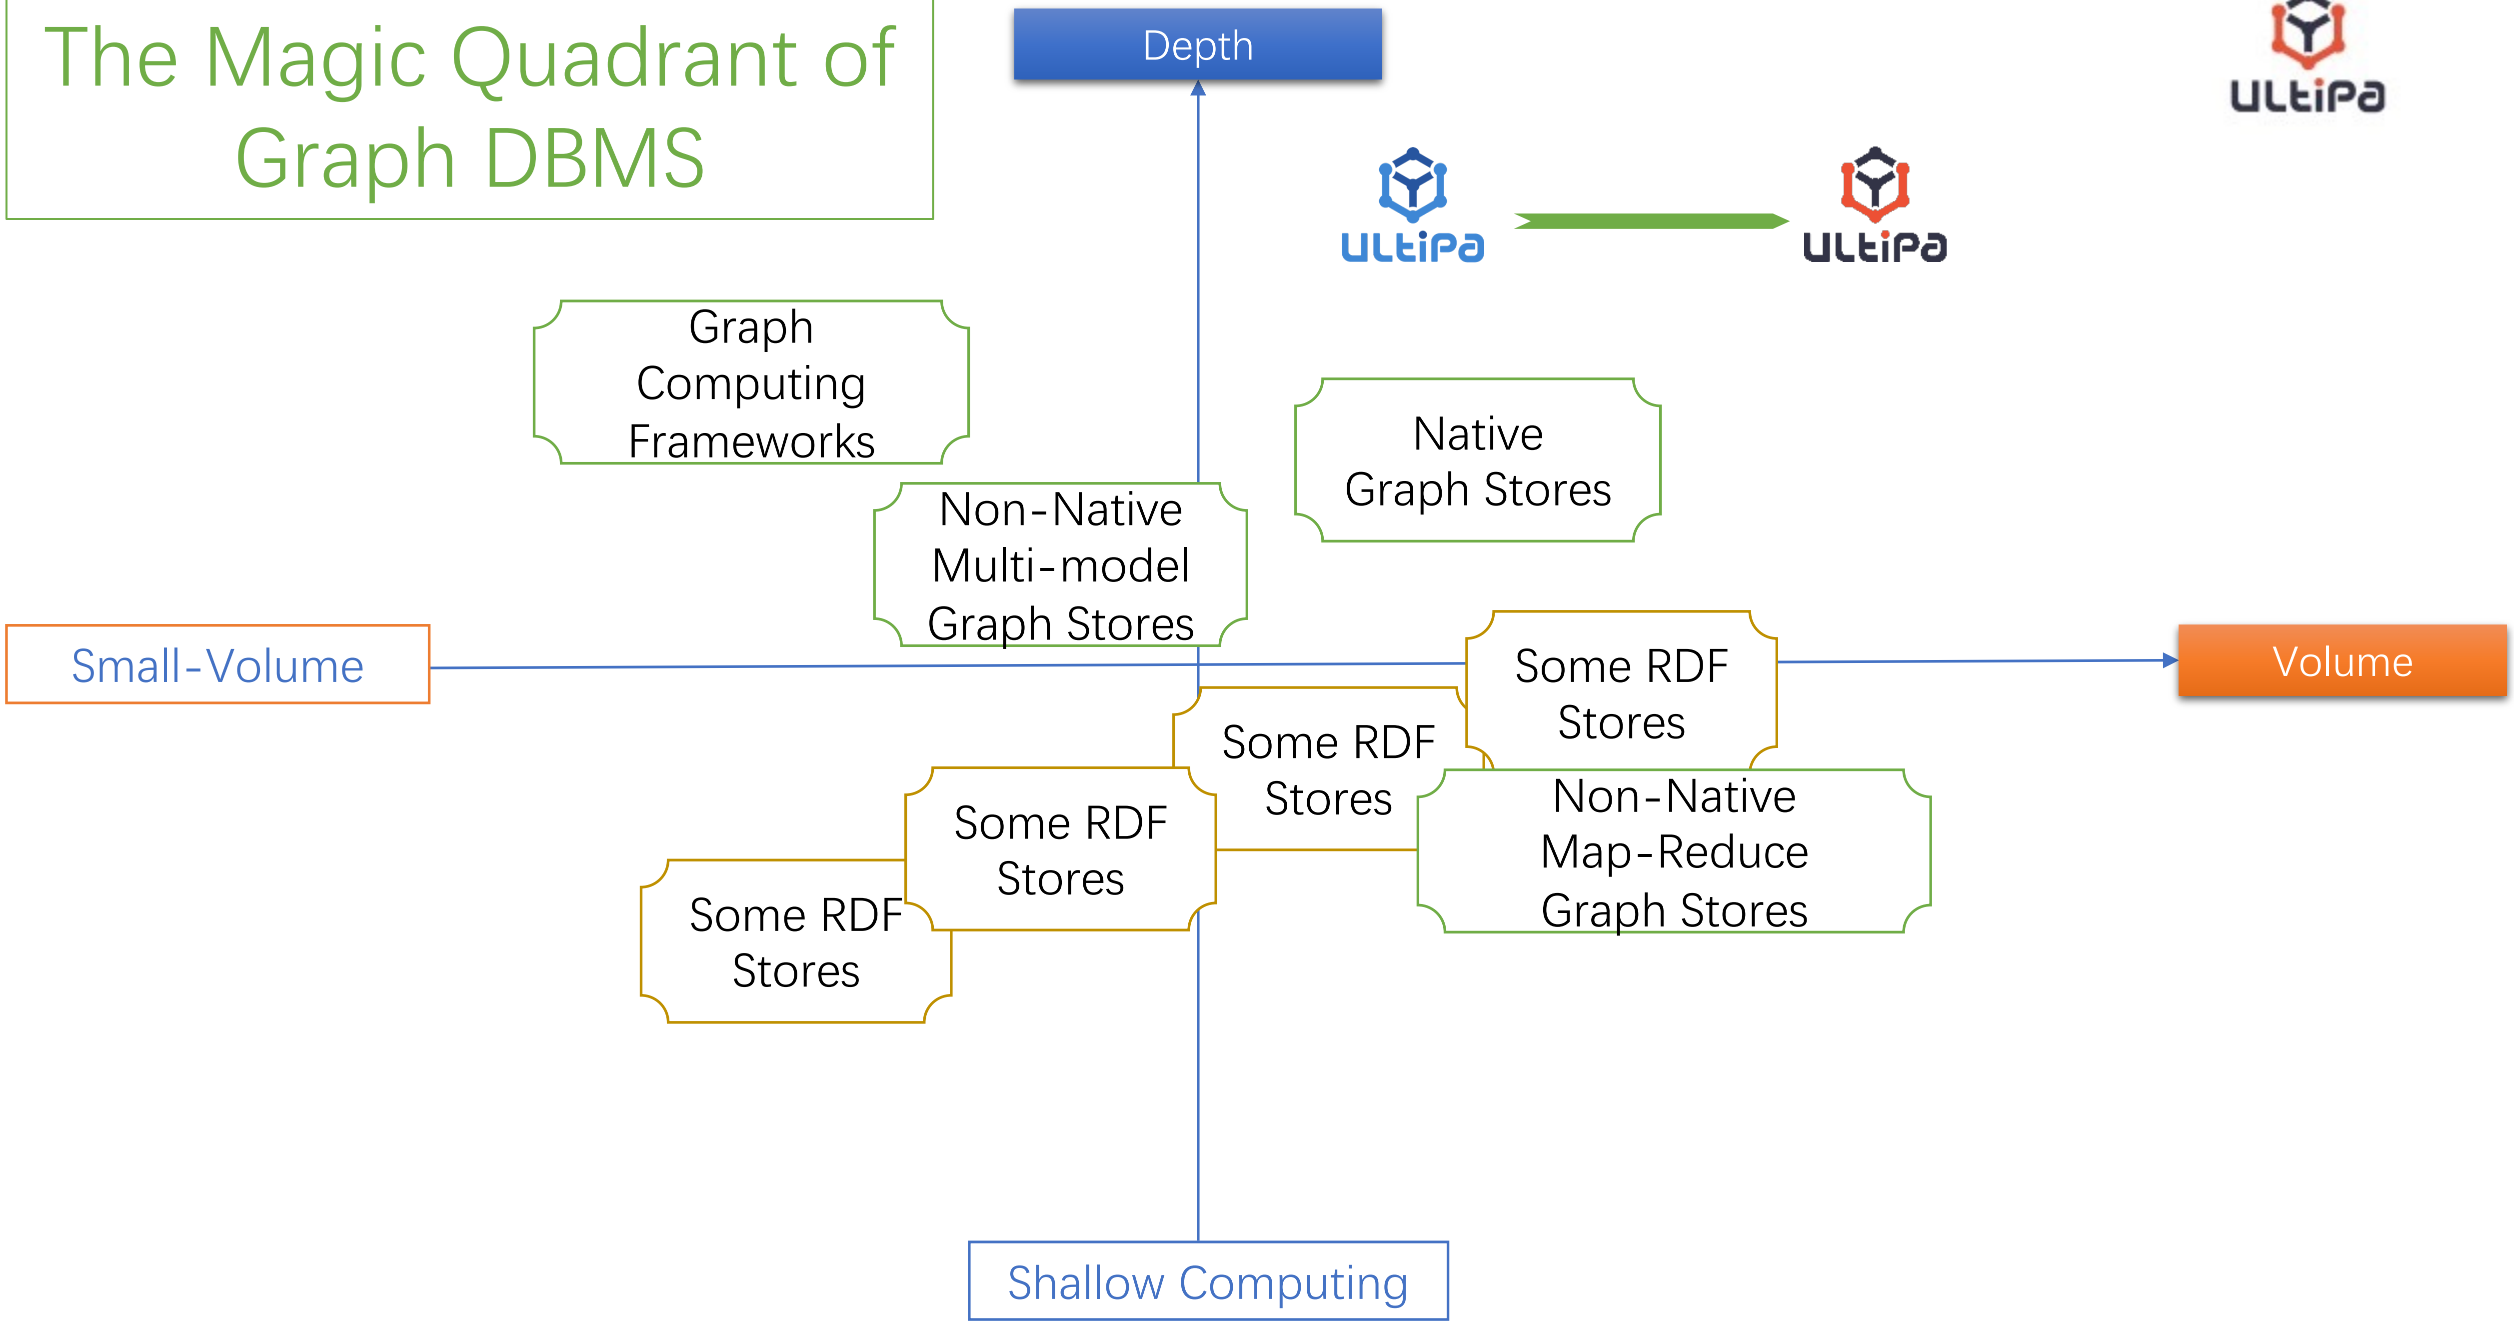

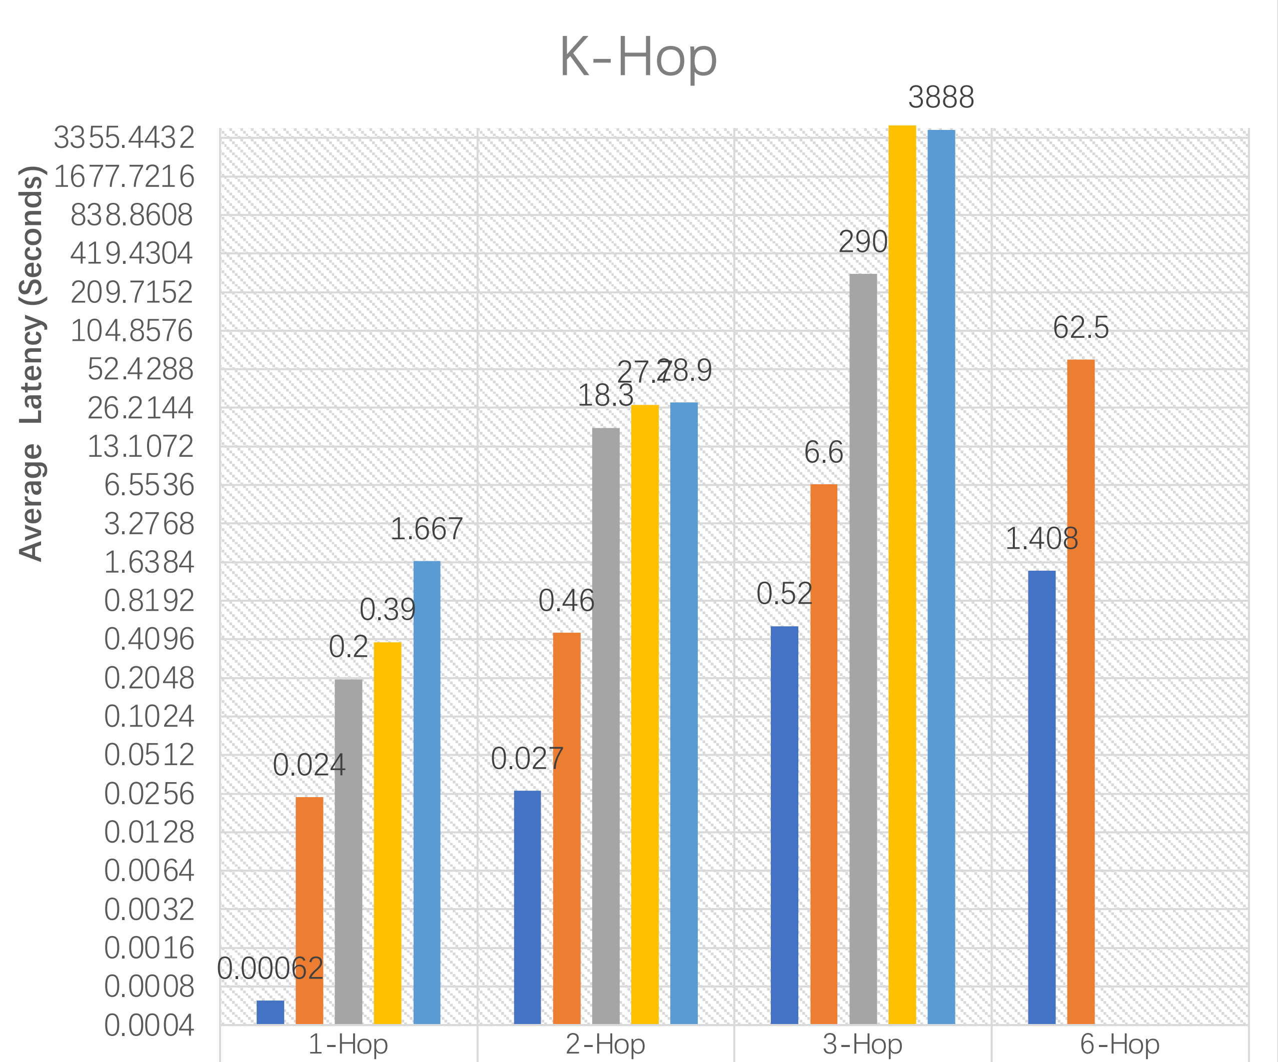

	1-Hop	2-Hop	3-Hop	6-Hop
Ultipa	0.00062	0.027	0.52	1.408
TigerGraph	0.024	0.46	6.6	62.5
Neo4j	0.2	18.3	290	Timeout
JanusGraph	0.39	27.7	4300	Timeout
ArangoDB	1.667	28.9	3888	Timeout

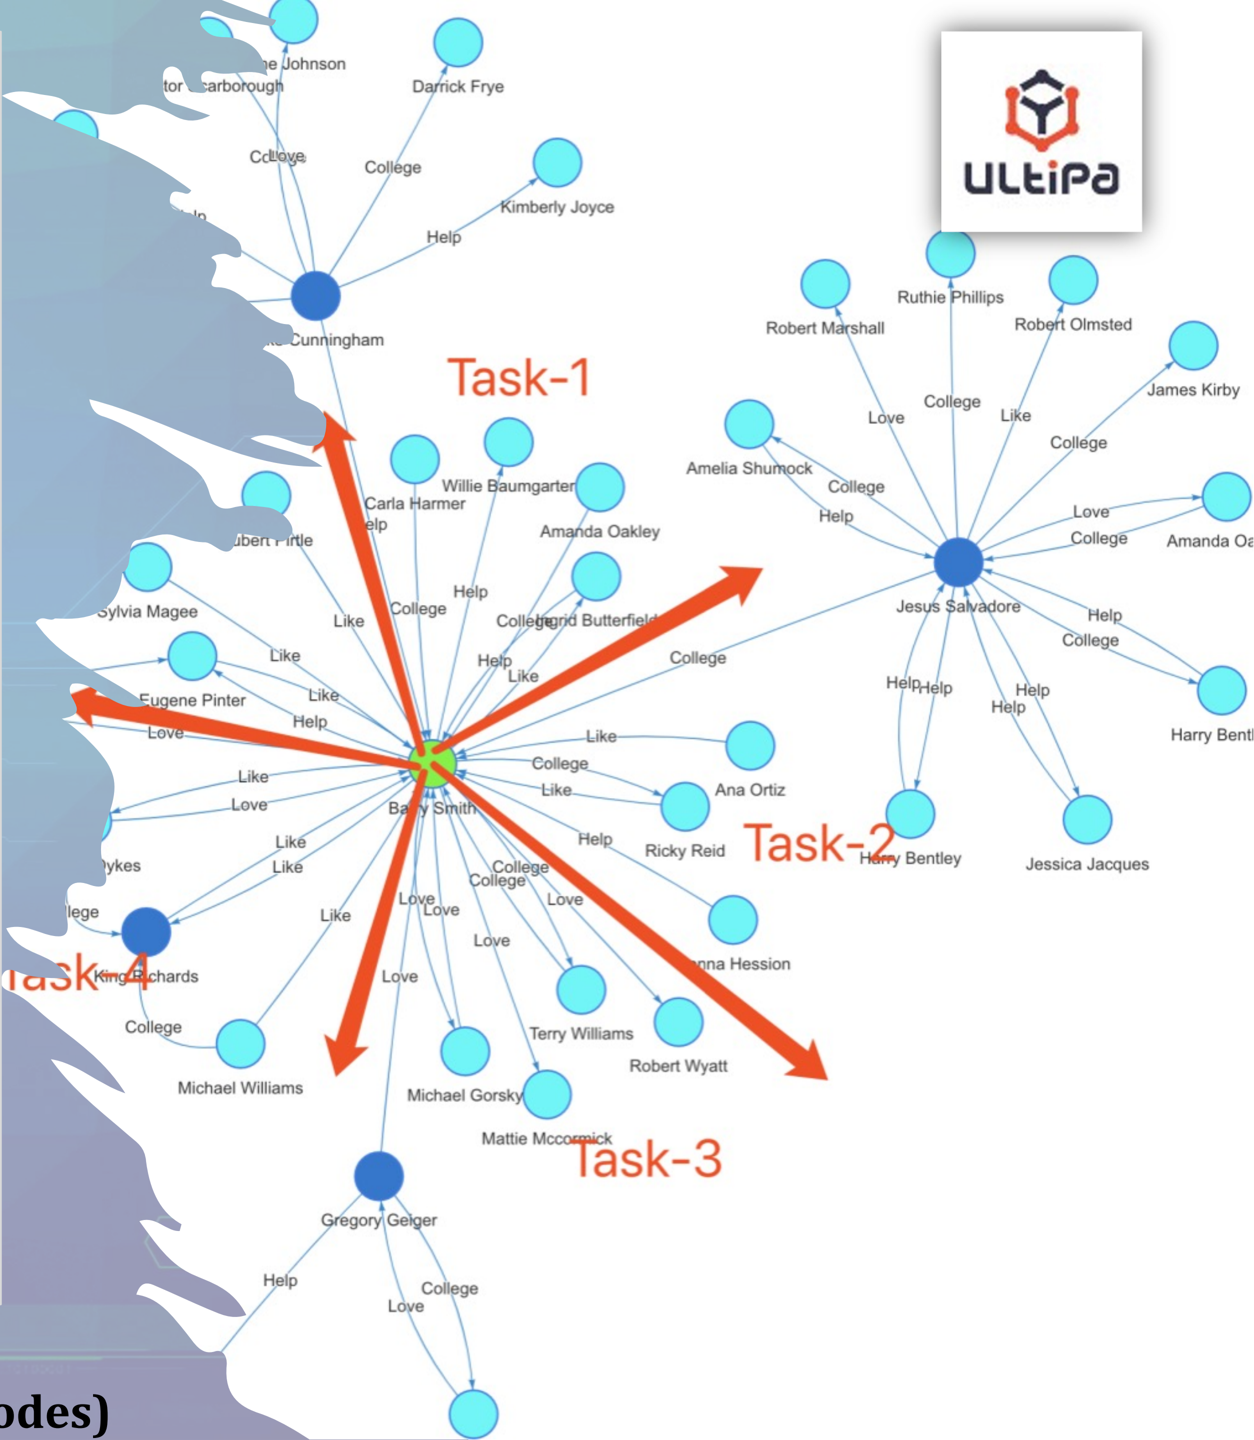

K-hop on Twitter-2010 Dataset
(42 Million Nodes & 1.47 Billion Edges with many supernodes)

Examine a graph's density

$$\frac{|E|}{|V|(|V|-1)} \longrightarrow \left(2\times\frac{|E|}{|V|}\right)^{k}$$

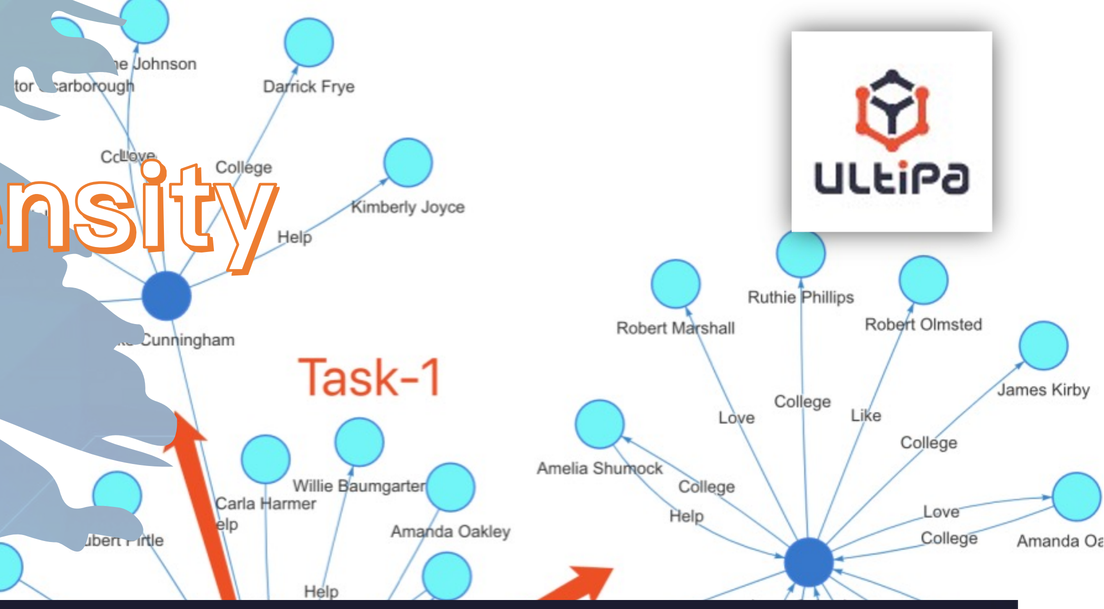

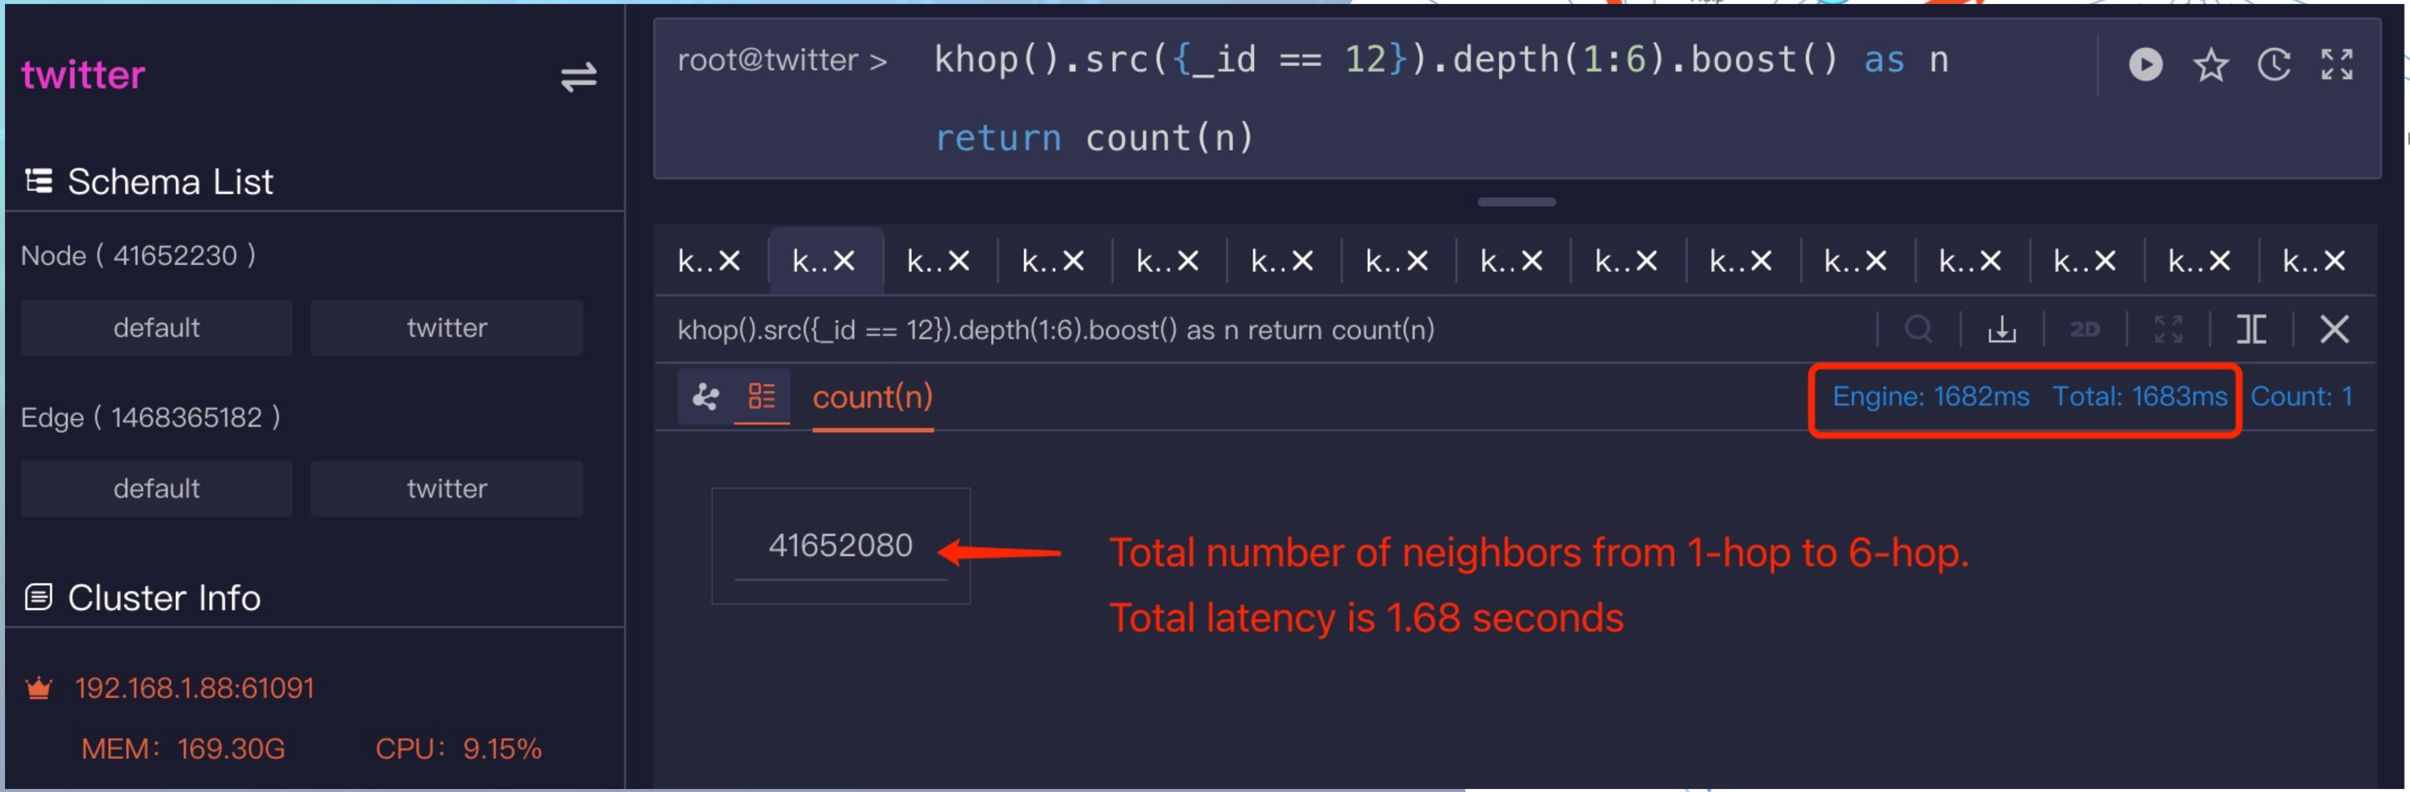

3 schools of Distributed Graph Systems

ULTIPA

Distributed Graph System	Pros	Cons
Distributed Consensus with HTAP [3]	● High Performance ● Better ACID support ● Small H/W footprint	● Vertical Scalability ● Difficult to handle 10 billion plus nodes and edges
Proxy/Name-server/Grid or Federation	● Balanced approach to scalability& performance ● No data migration	● Non-transparent graph partitioning (human-logic based)
Automated Shard	● Unlimited Scalability ● Great meta-data query and ingestion performance ● Sophisticated Cluster Management	● Degraded graph query performance ● Sophisticated Cluster Management ● Large H/W footprint

Scenarios of Distributed Graph Systems

ULTIPA

Type	Characteristics	Business Scenarios
High Density Parallel Graph Computing (HDPC)	• Real-time read/write data, online processing & calculation • Ideal for deep range queries	• Transaction interception • Online Anti-fraud • Anomaly detection • Real-time recommendation • AI/ML Augmentation • Other real-time scenarios
HDPC & Shard	• Separation of read/write operations • Elastic compute nodes [19] for shard/offline data	• Knowledge Graph • LLM Augmentation • Indicator calculation • Audit • Cloud Data Center • Graph at the core of IT Infra.
Shard	• Meta-data oriented • Shallow neighborhood calculation (1-2 hop) only	• Archive • Data Warehouse • Data Science

Hierarchy

	Peer	127.0.0.1:40061	HDPC / Computing & Storage Server/Instance
	Shard	[Peer1…3]	HTAP Cluster
	NameServer	[Shard1, Shard2]	Management &Computing Server
	NameServer Cluster	[N1,N2]	Multiple NameServer
	Elastic Compute Node	[Peer1, Peer2]	Dynamically allocate compute nodes
	Meta	[127.0.0.1:50061]	Configuration&Listener

Khop Source Parallelization

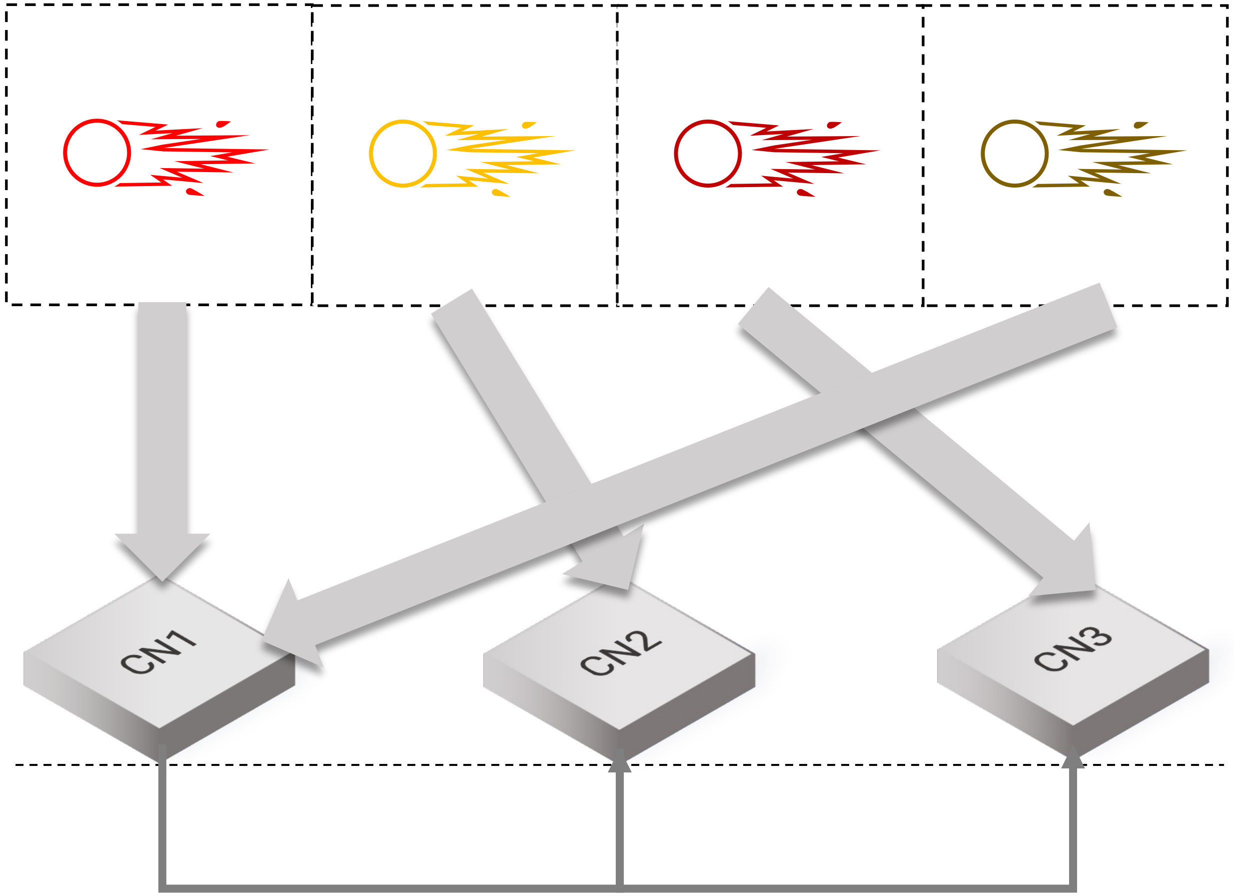

find().nodes() as nodes => [1,2,3,4]
khop().src().depth(3)

uncollect [“1” , “2” , “3” , “4”]
as nodes
khop().src({_id == nodes}).depth(3)

Khop(1) on CN1
Khop(2) on CN2
Khop(3) on CN3
Khop(4) on CN1

A-to-B Path Parallelization

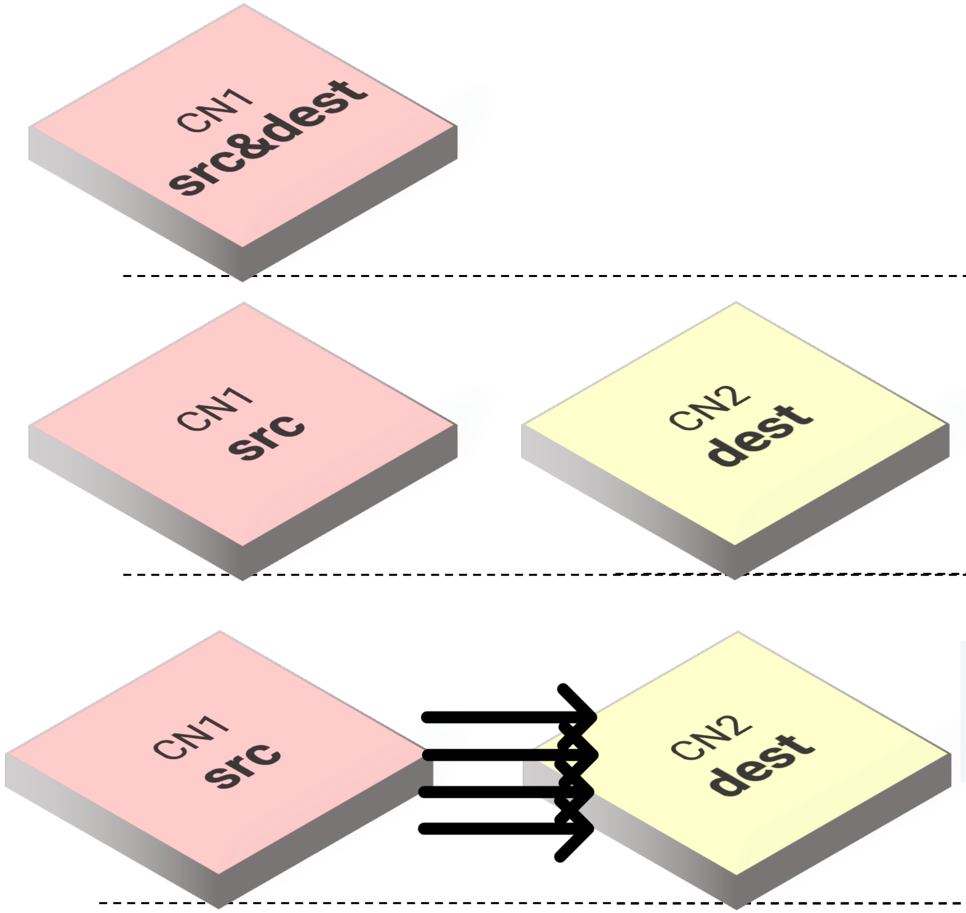


```
find().nodes() as src
find().nodes() as dest
ab().src(src).dest(dest).depth(1:2)
```

Dynamic CN

```
find().nodes({_id == 1}) as src
find().nodes({_id == 2}) as dest
ab().src(src).dest(dest).depth(1:2)
```

Chains/Broadcast

```
find().nodes({_id == 1}) as src
 ↘find().nodes({_id == 2}) as dest
    ↘ab().src(src).dest(dest).depth(1:2)
```

Cluster Hierarchy (as a Graph)

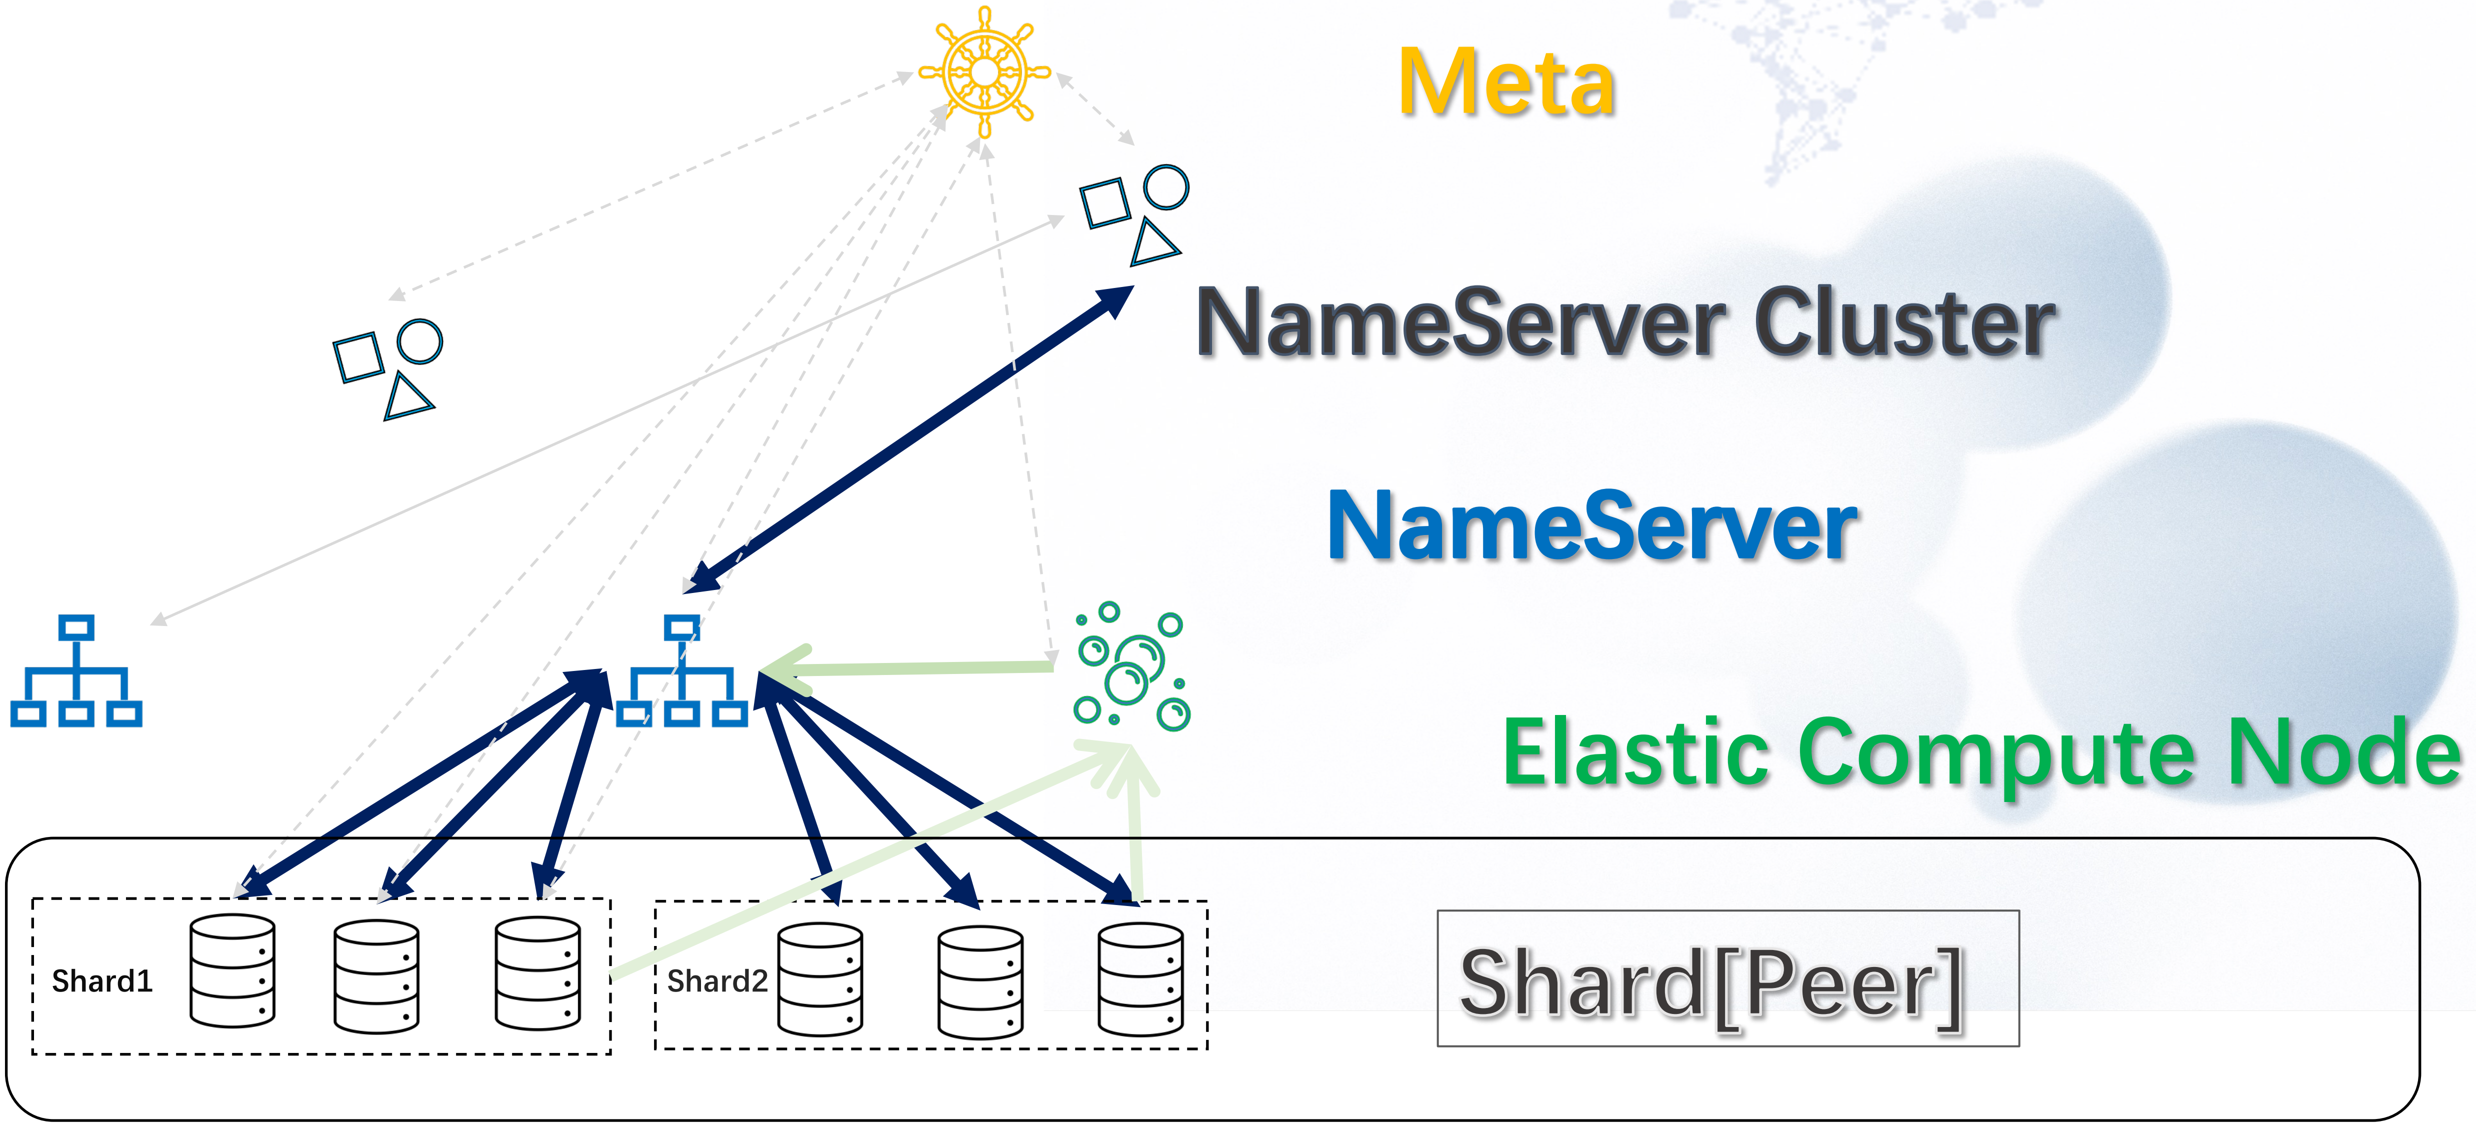

Distributed Ultipa 5.0 Summary

- Type 1: Data are processed on name servers (or proxies)
- Type 2: Data are processed on shard servers and name servers (Peer-to-peer architecture).
- Exchange Operator between Relational Data Stream(s) and Graph Algo
- Optimization of relational data flow (as start node/edge) is necessary in restricted range graph queries.

HTAP/Instance vs. Shard

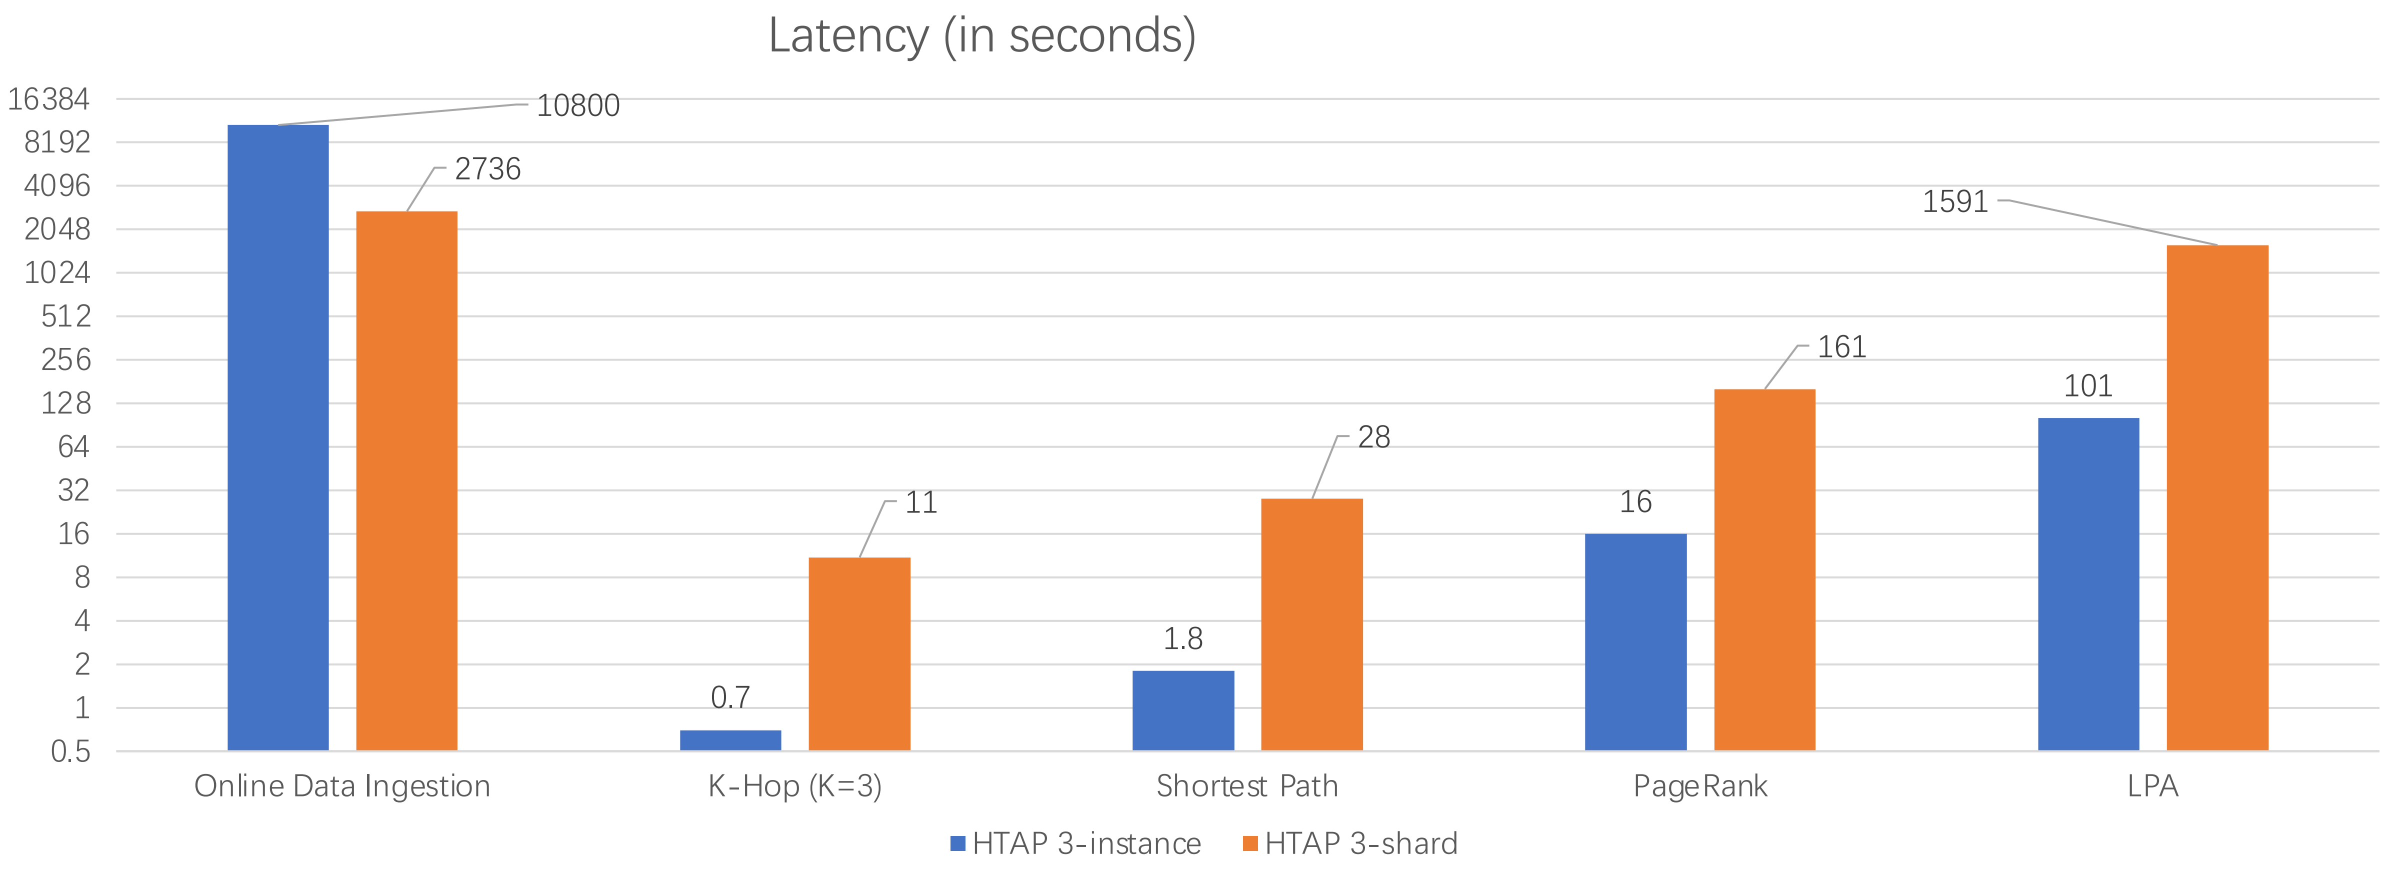

Summary

Three schools of Distributed Graph Systems

Distributed Graph System	Pros	Cons
Distributed Consensus with HTAP [3]	• High Performance • Better ACID support • Small H/W footprint	• Vertical Scalability • Difficult to handle 10 billion plus nodes and edges
Proxy/Name-server/Grid or Federation	• Balanced approach to scalability & performance • No data migration	• Non-transparent graph partitioning (human-logic based)
Automated Shard	• Unlimited Scalability • Great meta-data query and ingestion performance • Sophisticated Cluster Management	• Degraded graph query performance • Sophisticated Cluster Management • Large H/W footprint

ULTIPA

Ultipa builds the world's fastest graph database and killer apps that empower smart enterprise with graph-augmented business data intelligence.

Key Features of Ultipa Graph DBMS:

- Micro-second Query Processing & Ultra-Deep Data Penetration
- Real-time Attribution/Contribution Analysis
- Real-time Stress Testing & Scenario Simulation
- Highly Visualized 3D Interactive Web GUI
- HTAP-MPP Cluster and Fast Deployment & Migration

Ultipa Product Matrix:

- Real-time Decision Making & Anti-Fraud System
- Intraday Liquidity Risk & Cash Management System
- Real-time Asset & Liability Management System
- Smart Data Intelligence Toolkits
- Smart BI & Advanced Analytics

Contact Us

www.ultipa.com

support@ultipa.com

Data Governance

Graph to Relate Everything
Penetrate Everything
Quantify Everything

Agility & Capability

High-dimensional Correlation Analysis w/ Finest Granularity
Real-time RCA (Root-Cause)

Computing Power

Faster by 10,000x
100+ Algorithms
30-Hop Plus

Killer-App

Data Intelligence Toolkits
Asset-Liability Management
Smart BI/RTD Applications

Contact: ricky@ultipa.com
+1-408-917-0675
https://www.linkedin.com/in/rickysun

Web: https://www.ultipa.com

Ultipa V4

HTAP Distributed Consensus

<u>Success Stories</u>:
Deployed with G-SIB banks, stock exchanges and insurance companies. Largest commercial deployment of 100B+ graph size.

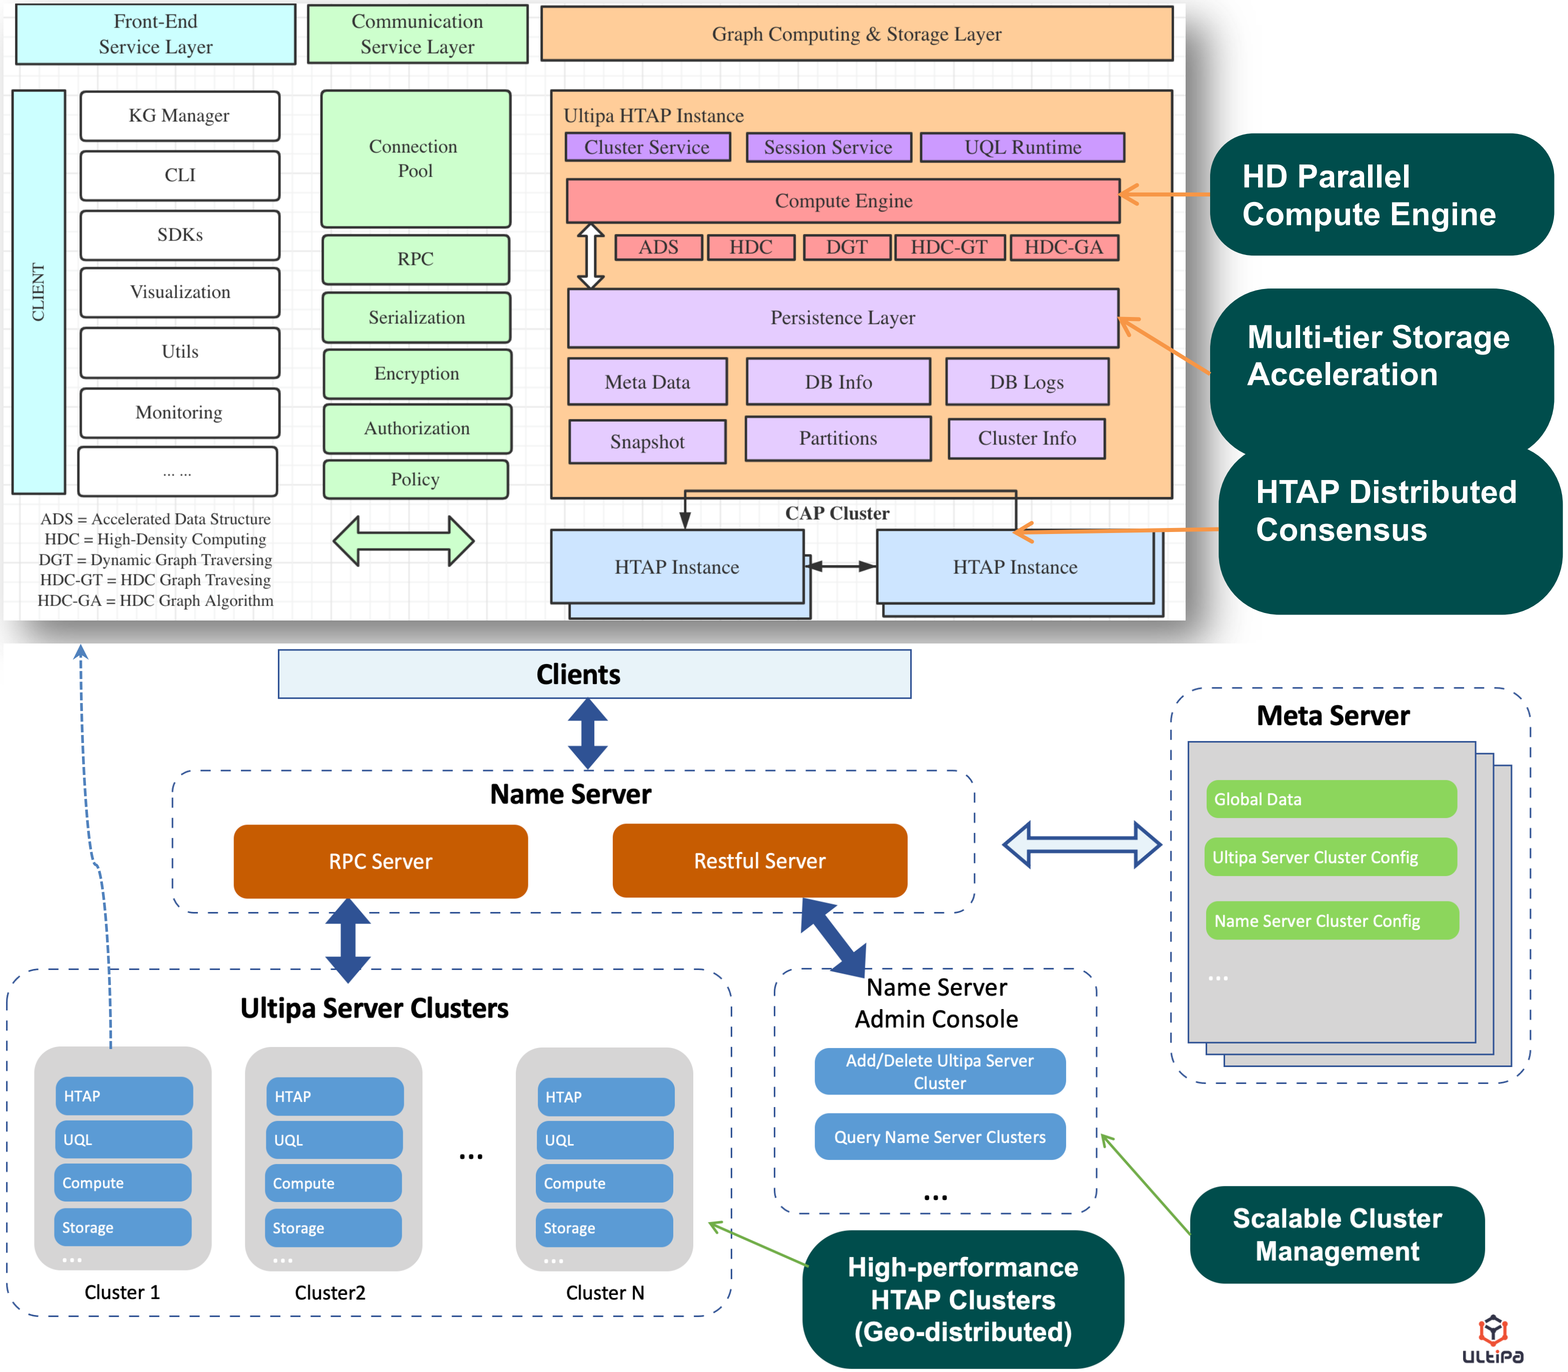

Ultipa V5

Horizontal Scalability
Low-mem Consumption
(GA in 2023)

- Cost-based Query Optimizer
- Self-defined Sharding
- Index-selection Optimizer
- Resource Consumption Estimation
- Push-down and Exchange
- Pregal-basd Distributed Graph Algorithm

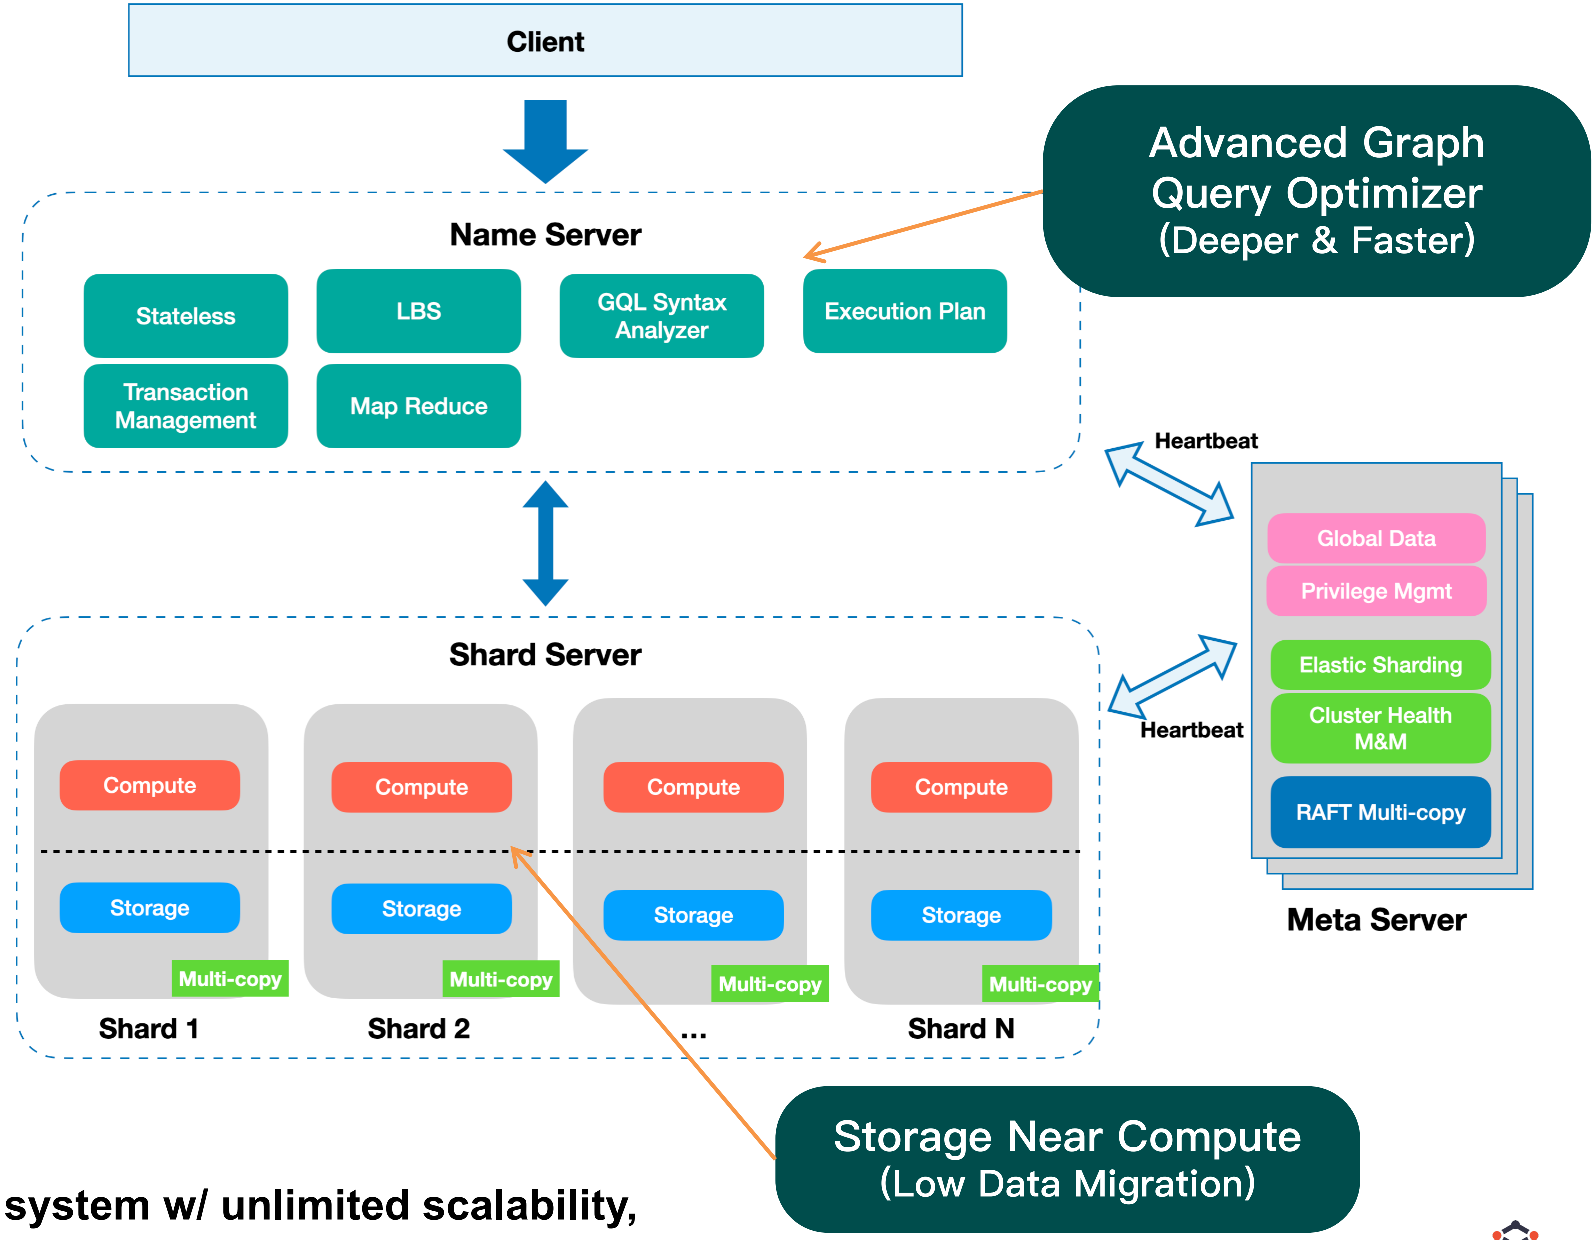

Sophisticated scalable graph database system w/ unlimited scalability, deep-data processing, and elastic computing capabilities.