

# ARCHITECTURE OVERVIEW

Why TigerGraph is 100x  
faster than competition

July 22, 2020

---



# Today's Presenters



**Victor Lee**

**Head of Product Strategy &  
Developer Relations**

- BS in Electrical Engineering and Computer Science from UC Berkeley
- MS in Electrical Engineering from Stanford University
- PhD in Computer Science from Kent State University focused on graph data mining
- 20+ years in tech industry



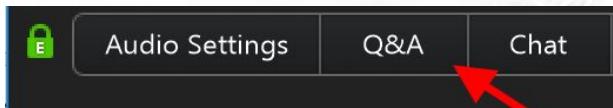
**Rayees Pasha**

**Principal Product Manager**

- MS in Computer Science from University of Memphis
- Prior Lead PM and ENG positions at Workday, Hitachi and HP
- Expertise in Database Management and Big Data Technologies

# Some Housekeeping Items

- Although your phone is muted you can ask questions at any time using the Q&A tab in the Zoom menu
- The webinar is being recorded and will be emailed you with the slides
- If you have any issues with Zoom please contact the organizer via chat



**PLEASE SEND QUESTIONS VIA ZOOM Q&A**



# Today's Outline

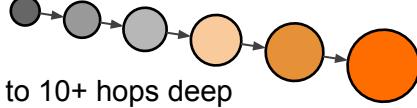


- 1 System Architecture Overview
- 2 Data Ingestion and Storage
- 3 Data Processing
- 4 Non-functional Features
  - HA
  - Transaction Management
  - Security

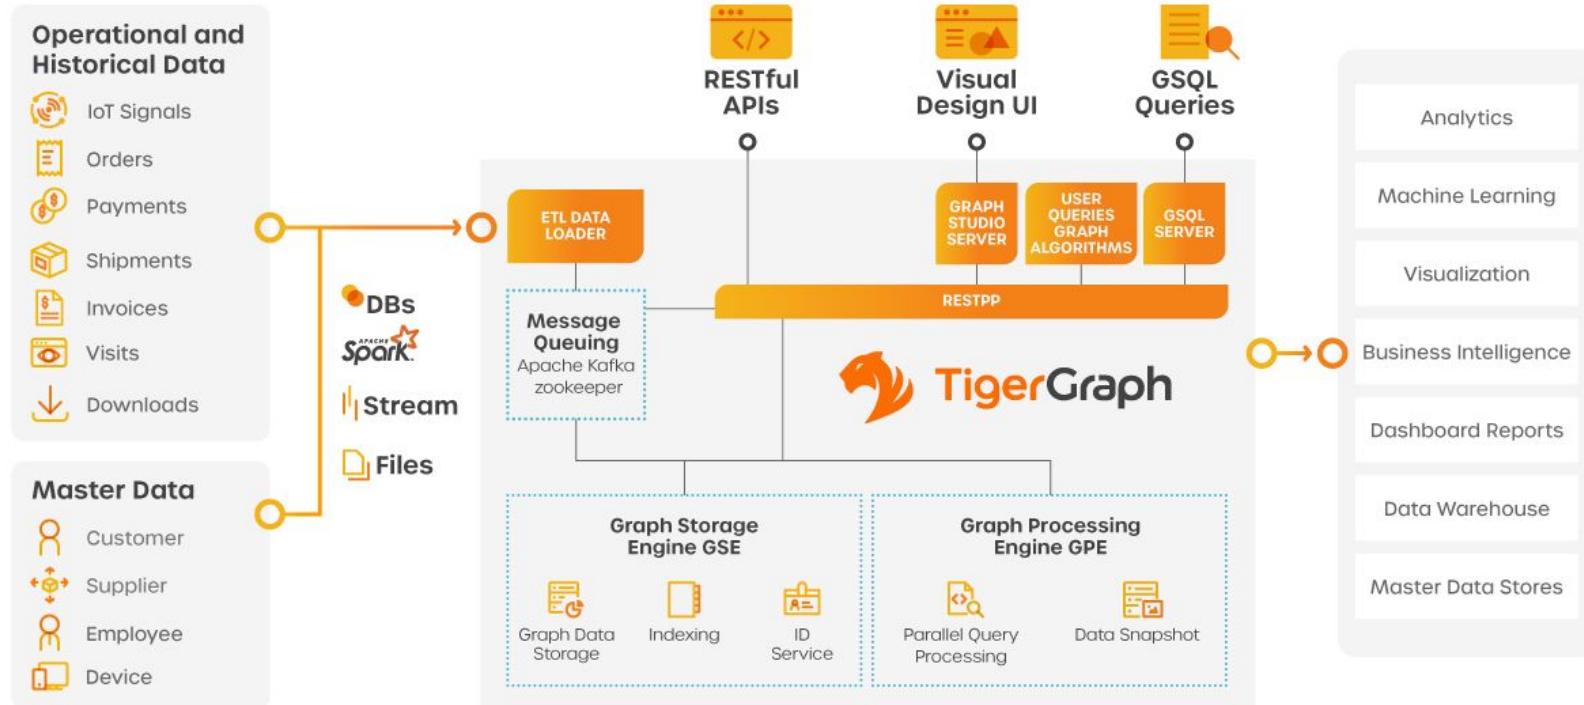
# SYSTEM ARCHITECTURE OVERVIEW



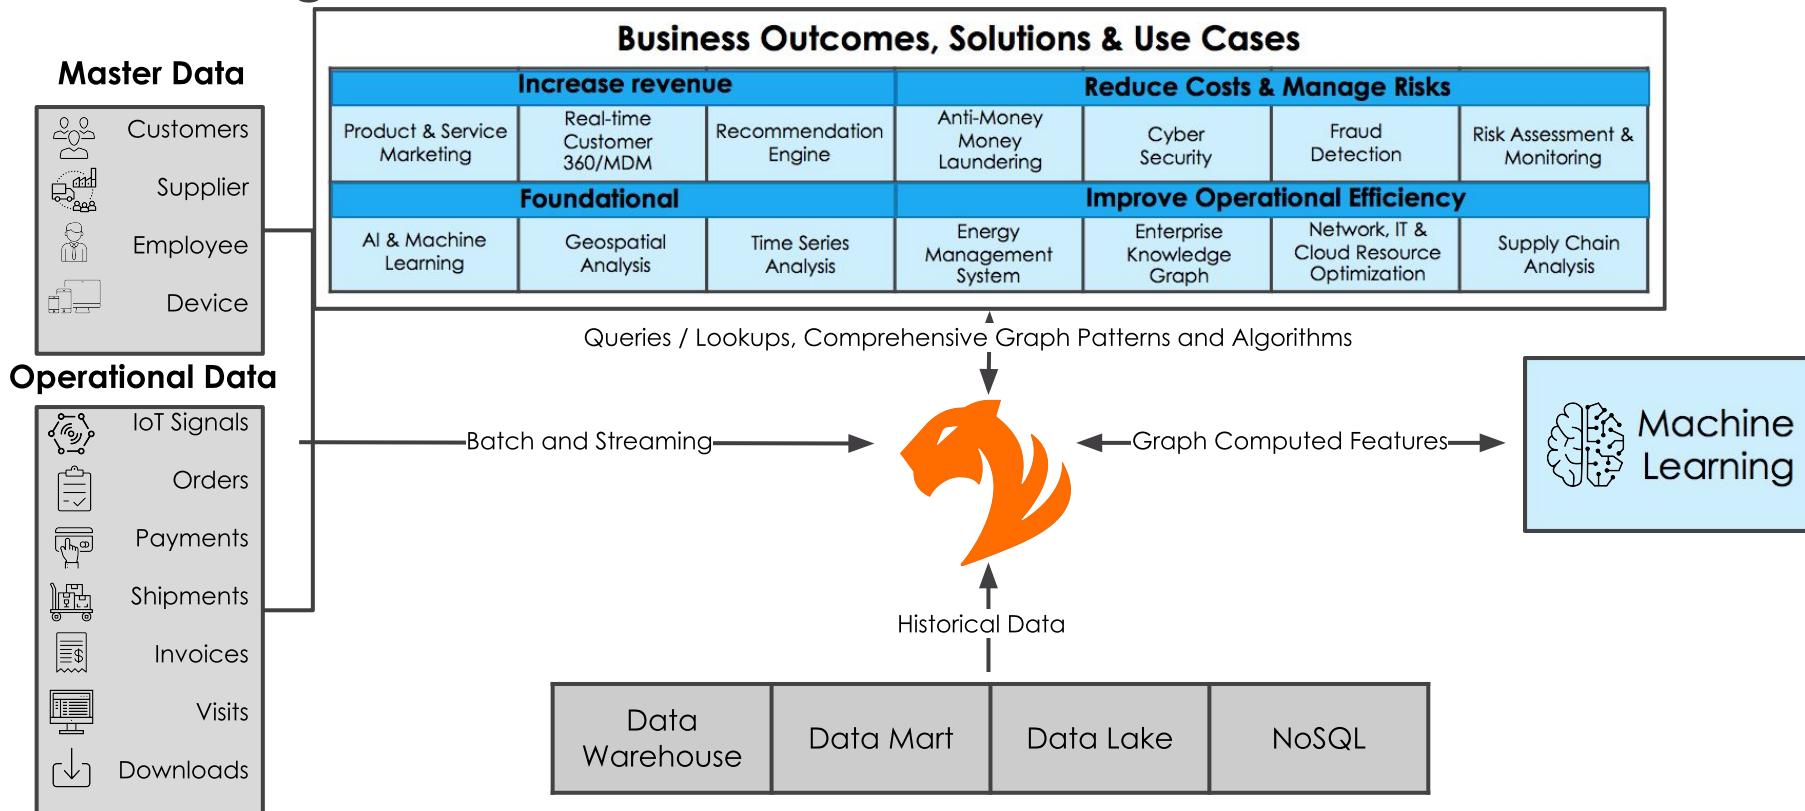
# The TigerGraph Difference

Feature	Design Difference	Benefit
<p><b>Real-Time Deep-Link Querying</b></p>  <p>5 to 10+ hops deep</p>	<ul style="list-style-type: none"><li>Native Graph design</li><li>C++ engine, for high performance</li><li>Storage Architecture</li></ul>	<ul style="list-style-type: none"><li>Uncovers hard-to-find patterns</li><li>Operational, real-time</li><li>HTAP: Transactions+Analytics</li></ul>
<p><b>Handling Massive Scale</b></p> 	<ul style="list-style-type: none"><li>Distributed DB architecture</li><li>Massively parallel processing</li><li>Compressed storage reduces footprint and messaging</li></ul>	<ul style="list-style-type: none"><li>Integrates all your data</li><li>Automatic partitioning</li><li>Elastic scaling of resource usage</li></ul>
<p><b>In-Database Analytics</b></p> 	<ul style="list-style-type: none"><li>GSQL: High-level yet Turing-complete language</li><li>User-extensible graph algorithm library, runs in-DB</li><li>ACID (OLTP) and Accumulators (OLAP)</li></ul>	<ul style="list-style-type: none"><li>Avoids transferring data</li><li>Richer graph context</li><li>In-DB machine learning</li></ul>

# TigerGraph Architecture



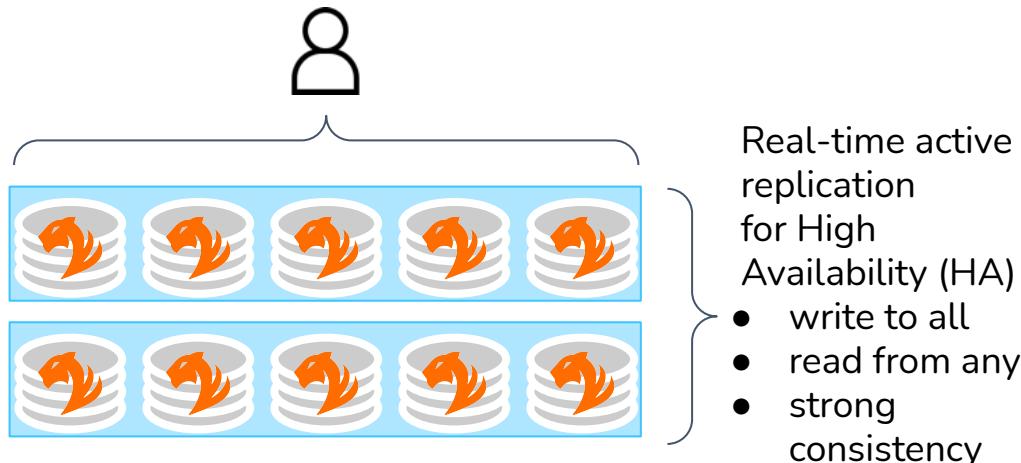
# Seamless integration enables businesses to **accomplish more** with existing investments



# TigerGraph Distributed Database Architecture

## Simple setup, Performant design

- **Setup:** Just tell TigerGraph how many servers.
- TigerGraph seamlessly distributes data.
- Users see a single database, not shards.



## Advantages:

- **Simple to setup and manage**
- **Unlimited scale-out;** simple to expand
- **Scalable OLAP:** massively parallel processing
- **Scalable OLTP:** concurrent ACID transactions
- **Economical**

# DATA INGESTION AND STORAGE



# Data Ingestion

## Step 1

Loaders take in user source data.

- Bulk load of data files or a Kafka stream in CSV or JSON format
- HTTP POSTs via REST services (JSON)
- GSQL Insert commands

## Step 2

Dispatcher takes in the data ingestion requests in the form of updates to the database.

1. Query IDS to get internal IDs
2. Convert data to internal format
3. Send data to one or more corresponding GPEs

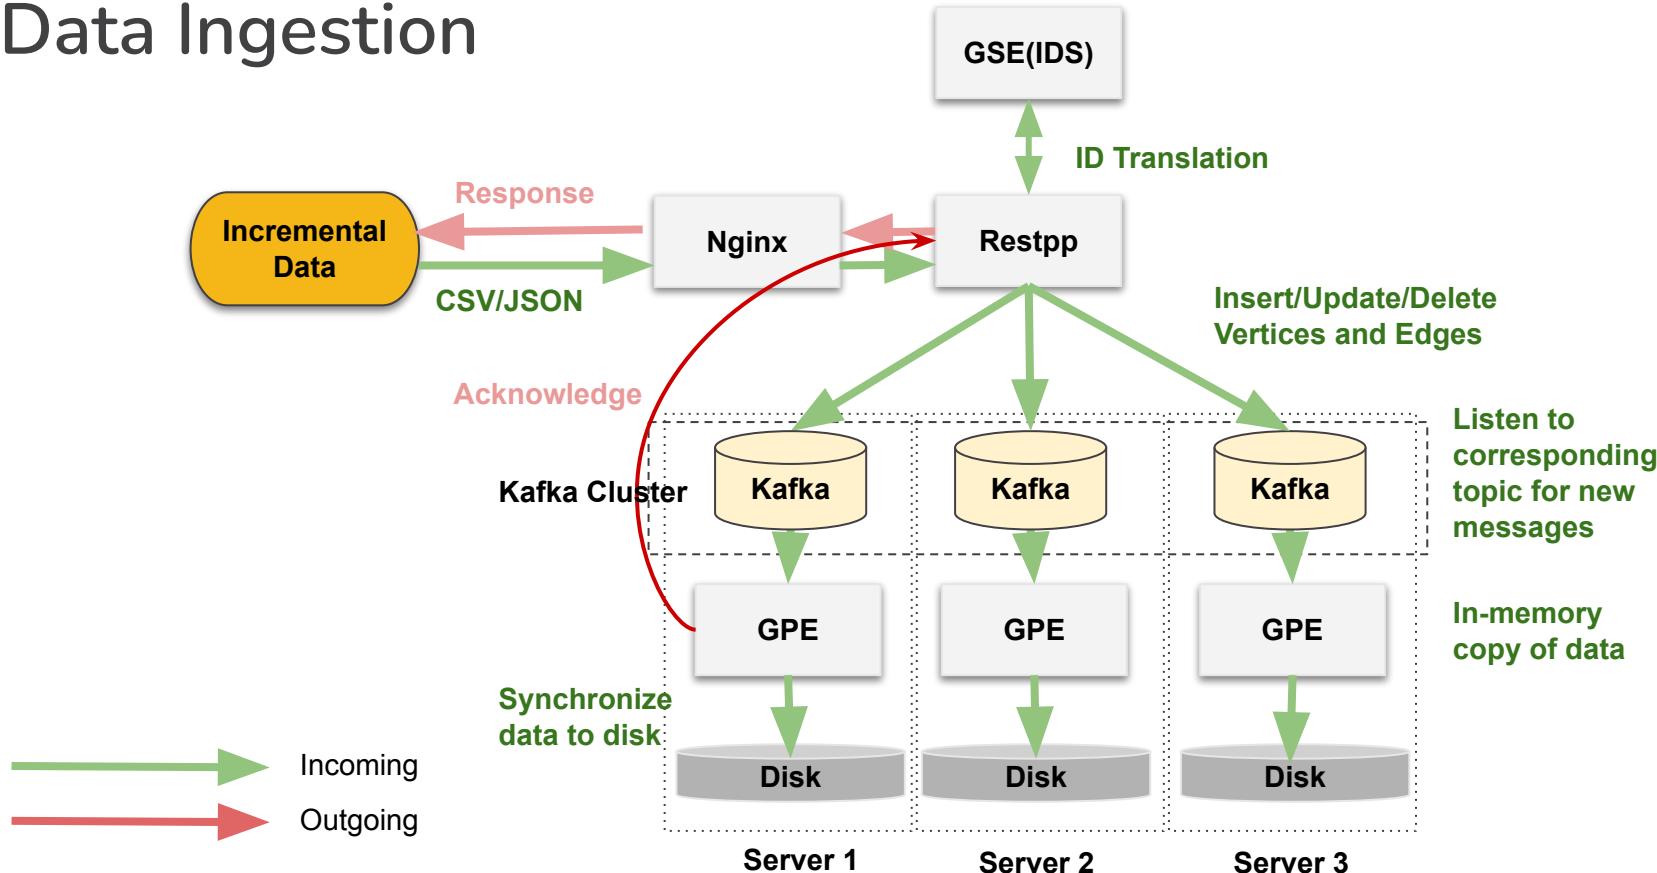
## Step 3

Each GPE consumes the partial data updates, processes it and puts it on disk.

Loading Jobs and POST use UPSERT semantics:

- If vertex/edge doesn't yet exist, create it.
- If vertex/edge already exists, update it.
- Idempotent

# Data Ingestion



# TigerGraph Distributed Native Graph Storage

“USER123” <---> 1234321

**IDS:** Bidirectional external ID to Internal ID mapping

1234321, John, 33, [john@abc.com](mailto:john@abc.com)  
1234322, Tom, 27, [tom@abc.com](mailto:tom@abc.com)

...

**Vertex Partitions:** Vertex internal ID and attributes

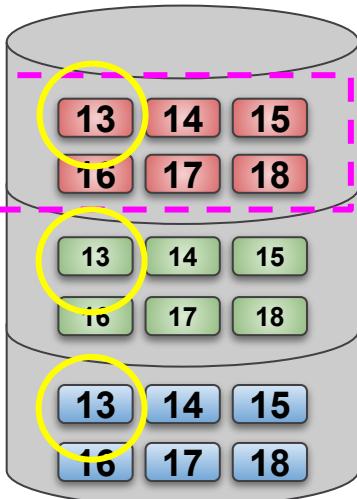
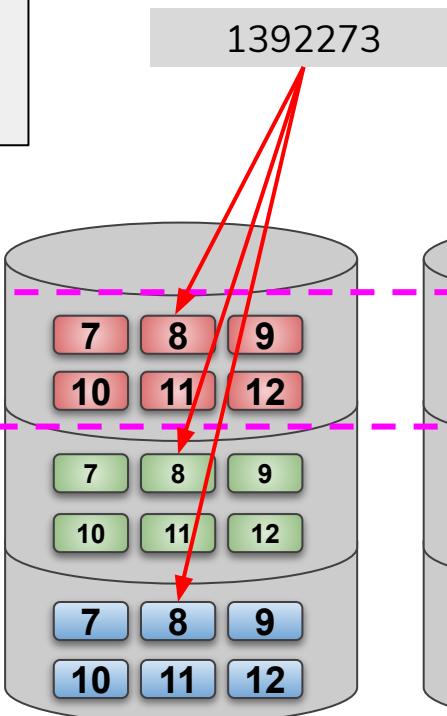
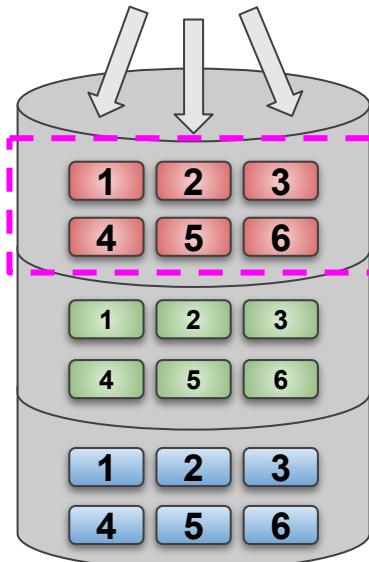
1234321, 1234322, 2020-04-23, 3.3  
1234321, 1234324, 2020-02-13, 2.3

...

**Edge Partitions:** Source vertex internal ID, target vertex internal ID, edge attributes

# TigerGraph Distributed Native Graph Storage

Data is split into segments.



IDS  
VERTEX  
EDGE

Data in segments is distributed across the cluster.

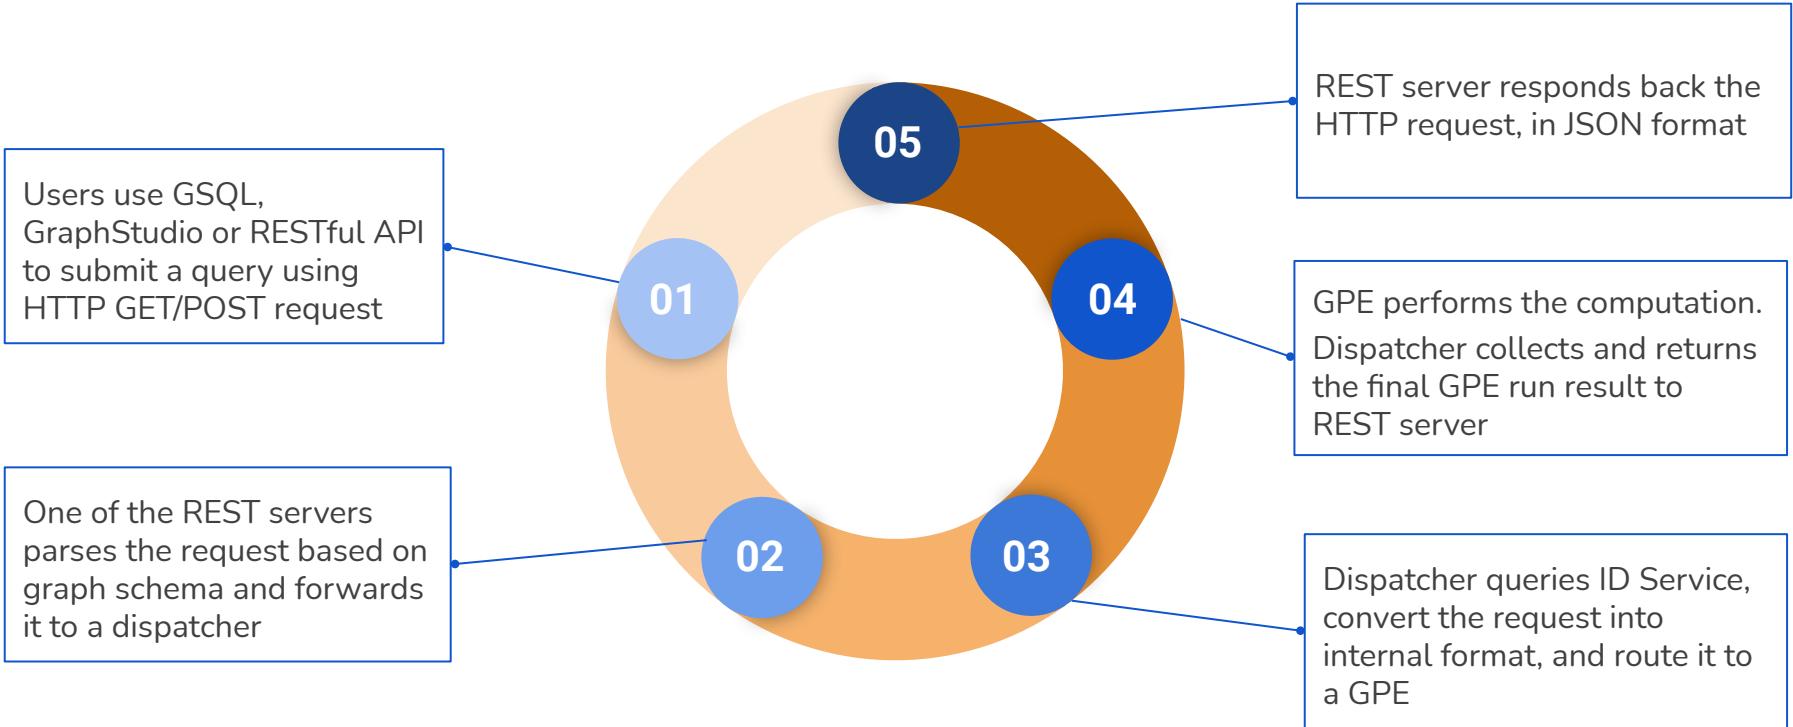
The segments with same ID store data for the same set of vertices.

The location of a vertex can be calculated based on its internal ID

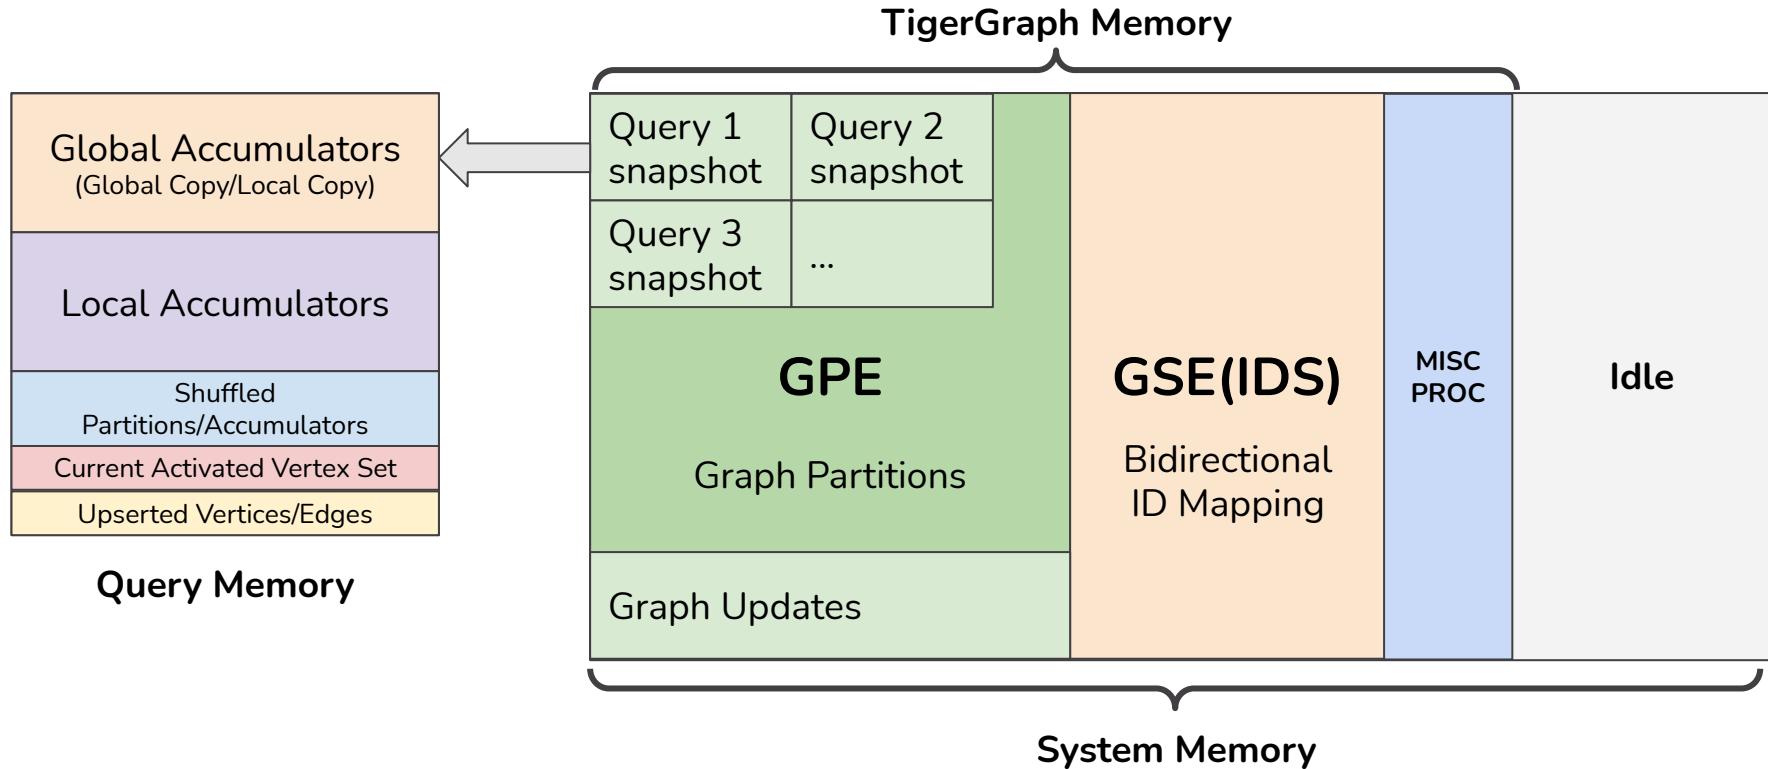
# DATA PROCESSING



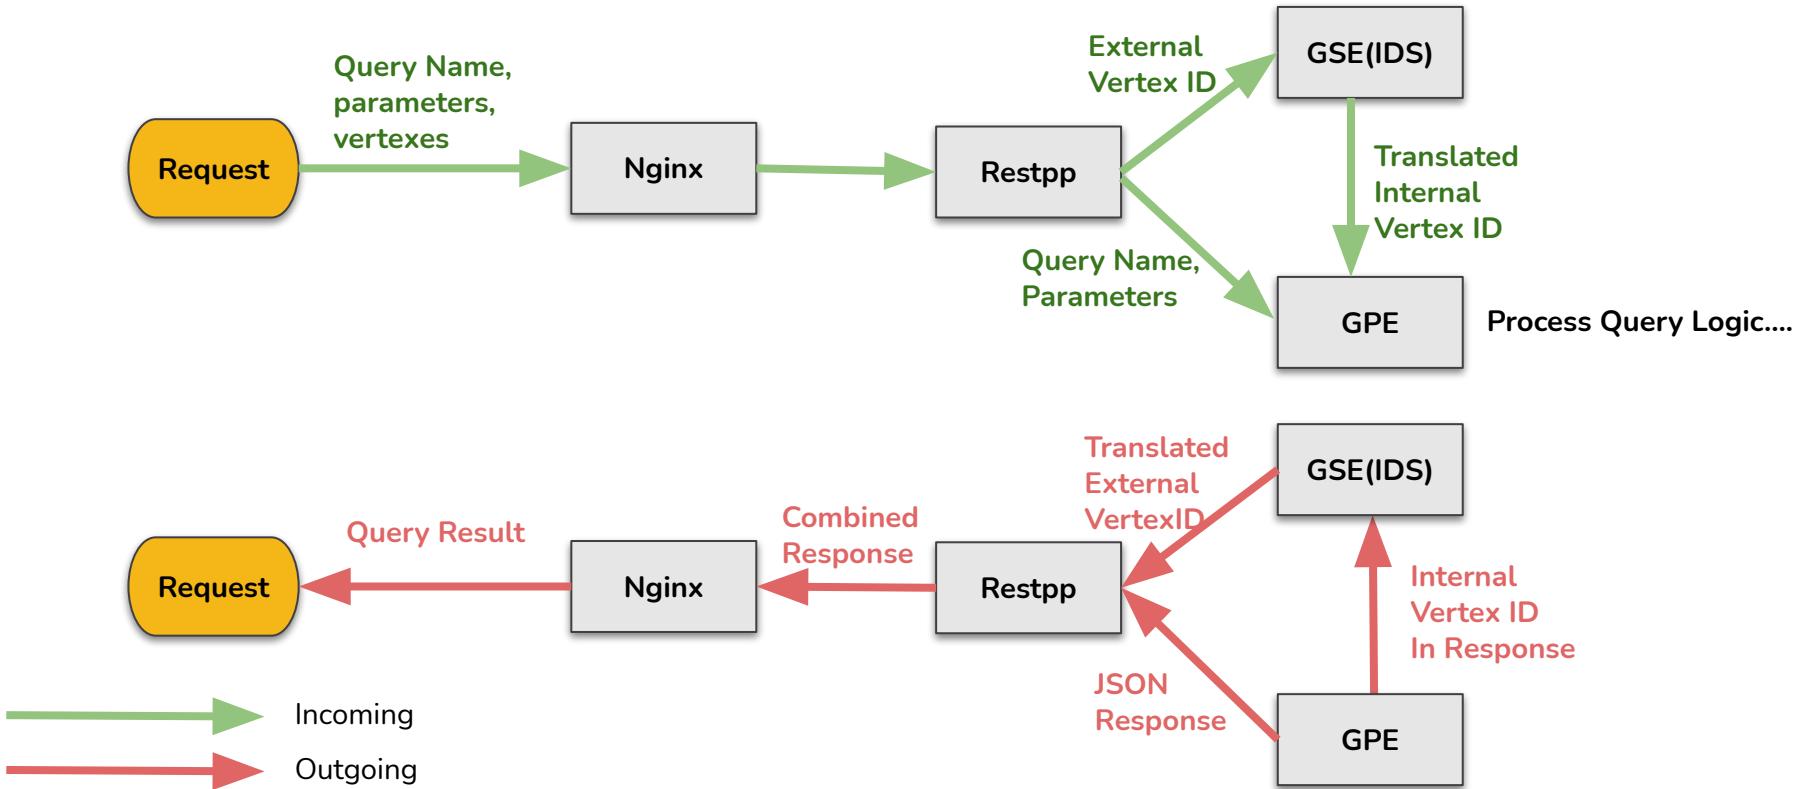
# Query Processing Workflow Overview



# TigerGraph Memory Usage Overview



# Query Processing workflow



# GSQL Queries



## Schema-Based

Optimizes storage efficiency and query speed. Supports data-independent app/query development.



## Conventional Control Flow (FOR, WHILE, IF/ELSE)

Makes it easy to implement conventional algorithms



## Built-in High Performance Parallelism

Achieves fast results while being easy to code. Accumulators turbocharge parallel computation.



## SQL-Like

Familiar to 1 million users



## Procedural Queries

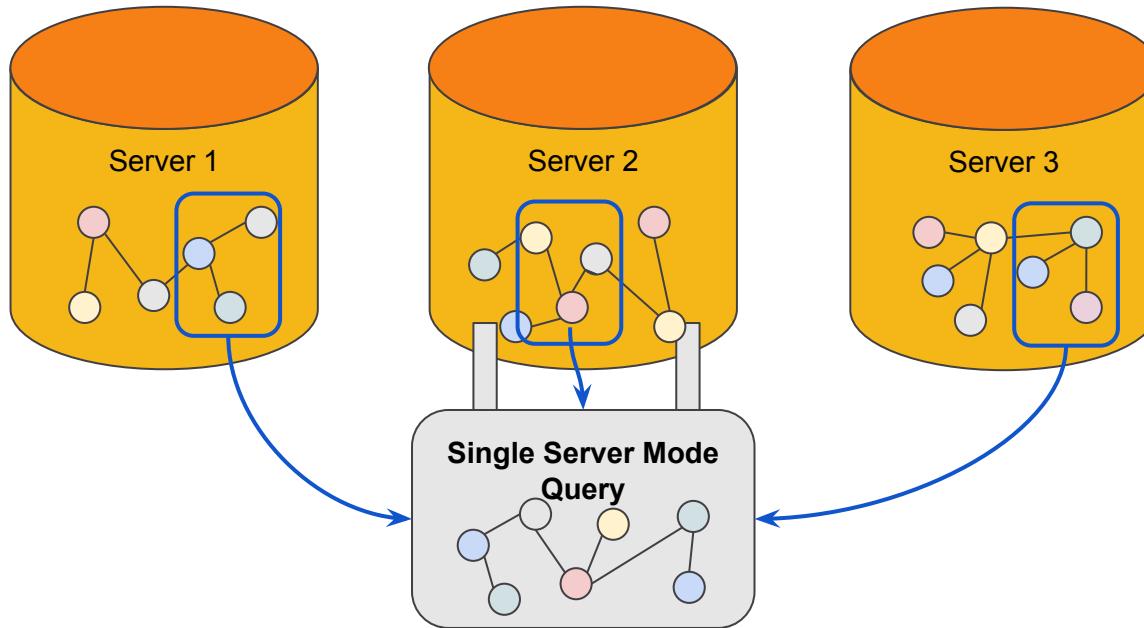
Parameterized queries are flexible and can be used to build more complex queries



## Transactional Graph Updates

HTAP - Hybrid Transactional / Analytical Processing with real-time data updates

# MPP - Distributed Cluster in Single Server mode



## Single Server Mode

- The cluster elects one server to be the master for that query.
- All query computations take place on query master.
- Vertex and edge data are copied to the query master as needed.
- Best for queries with one or a few starting vertices.

# HTAP: Hybrid Transaction/Analytical Processing

## OLTP - Transactional

- Real-time read and write
- ACID properties  
(guarantee that transaction is correct)
- Concurrency  
(many transactions at the same time)



## OLAP - Analytical

- Multi-dimensional Analysis
- Compute-intensive
- Data-intensive
- Aggregation



**TigerGraph**

- + Real-time read & write
- + Real-time, compute-intensive, multi-dimensional analysis
- + Real-time aggregation

## OLTP and OLAP Together in Graph

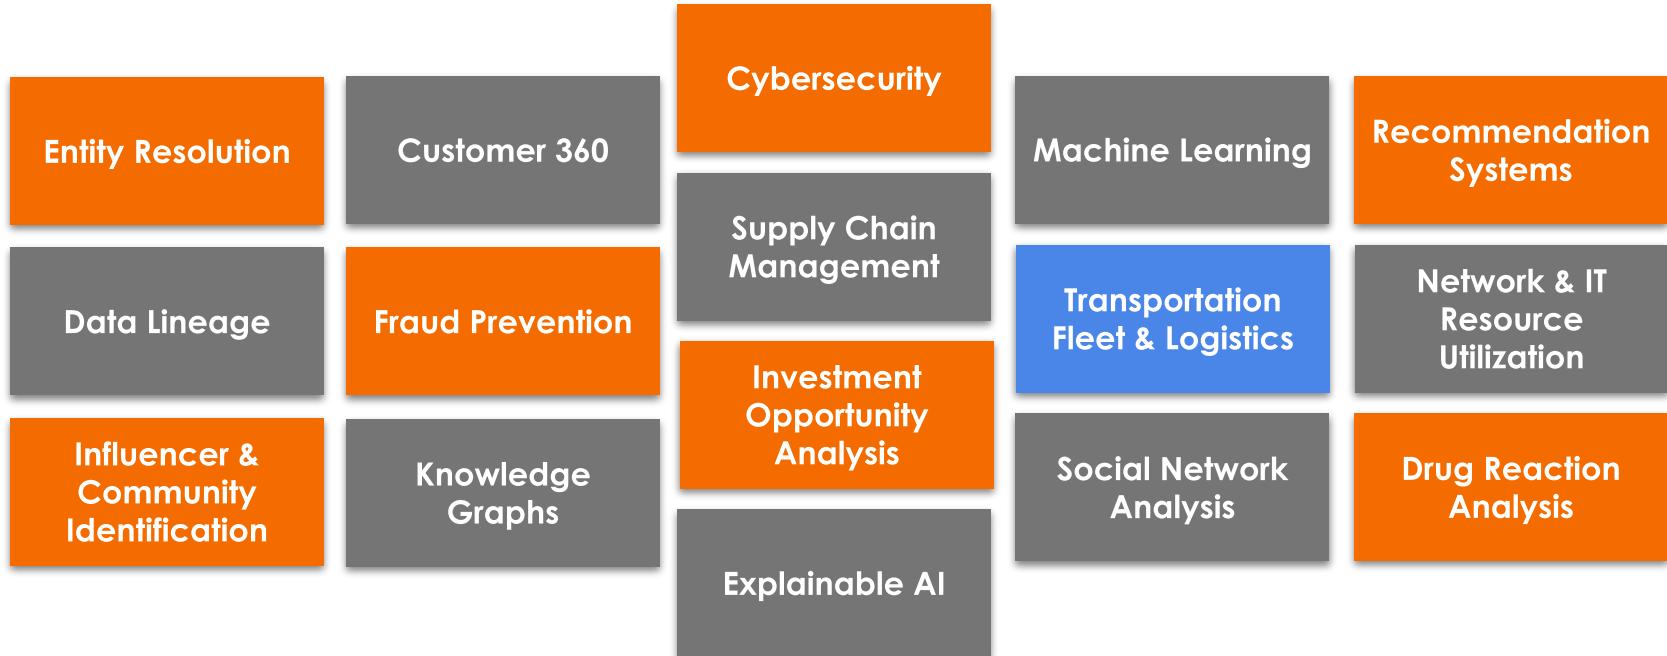
### Some Graph Databases

- + ACID
- + Concurrency

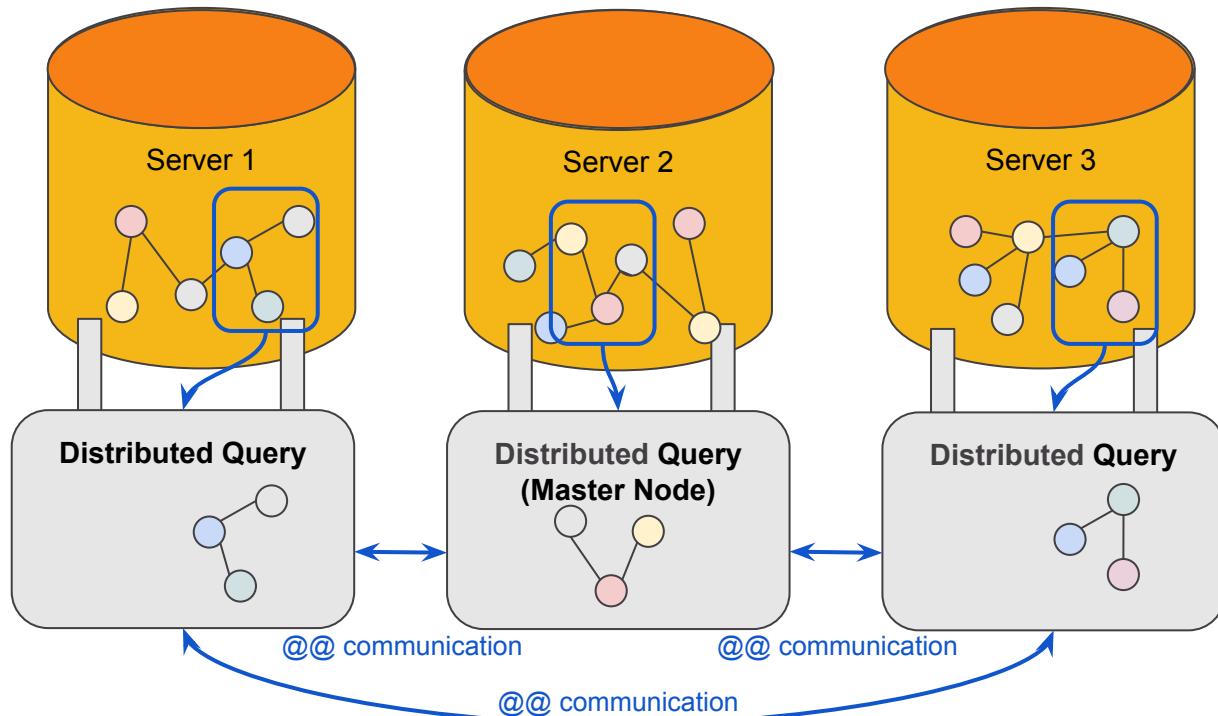
### All Graph Databases

- o Multi-dimensional data
- o Real-time read

# There's No Limit To Where Graph Can Help



# MPP - Distributed Cluster in Distributed mode

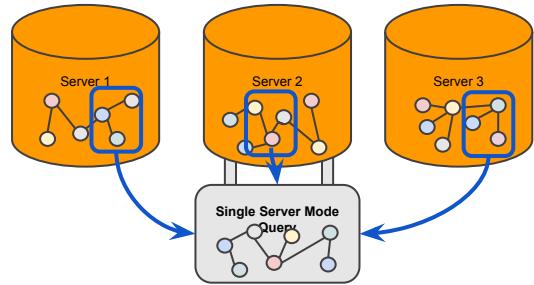


## Distributed Mode

- The server that receives the query becomes the master.
- Computations execute on **all** servers in parallel.
- Global accumulators are transferred across the cluster.
- If your query starts from all or most vertices, use this mode.

# MPP mechanism

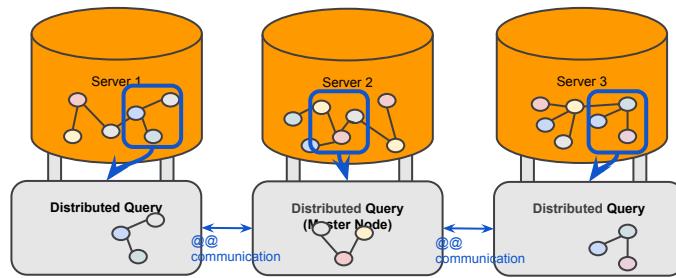
## Single Server mode VERSUS Distributed mode



### Single Server Mode is better when

1. Starting from a single or small number of vertices.
2. Modest number of vertices/edges are traversed.
3. Heavy usage of global accumulators.

Ex: Point query, single entity-based transaction/update



### Distributed Mode is better when

1. Starting from all or a large number of vertices.
2. Very large number vertices/edges are traversed.

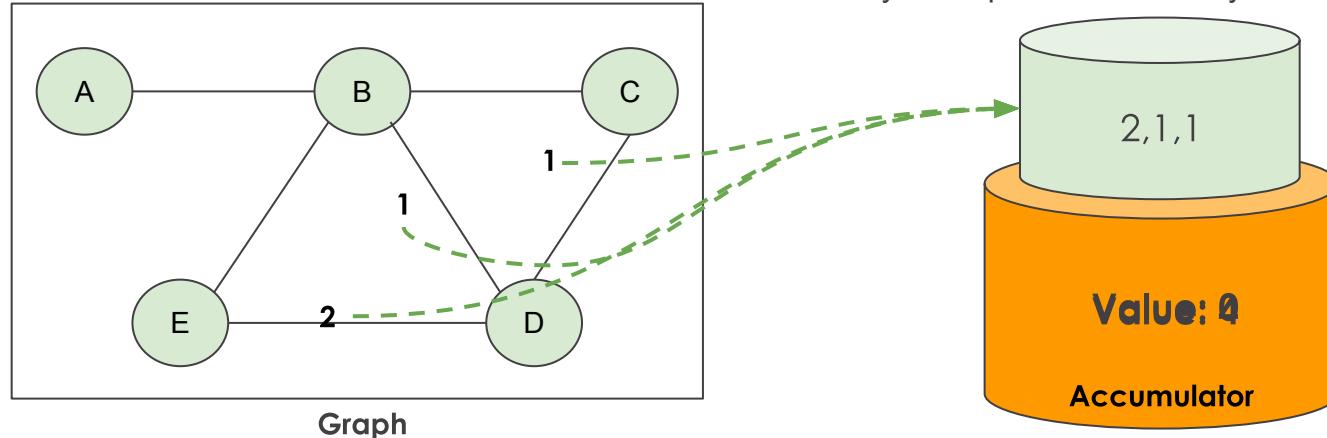
Ex: Most graph algorithms & global analytics  
(PageRank, Closeness Centrality, Louvain Community, etc.)

# Accumulators

**Accumulators** are a special type of variables that accumulate information from multiple workers/threads during a query. Workers/threads work asynchronously.

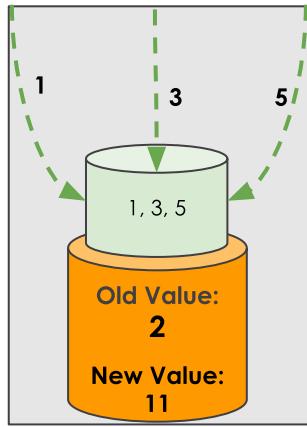
**Accumulating phase 1:** Receive messages and store them temporarily in a bucket that belongs to the accumulator.

**Accumulating phase 2:** Aggregate the messages it received based on its accumulator type. The aggregated value will become the accumulator's value which can be accessed by other parts of the Query.



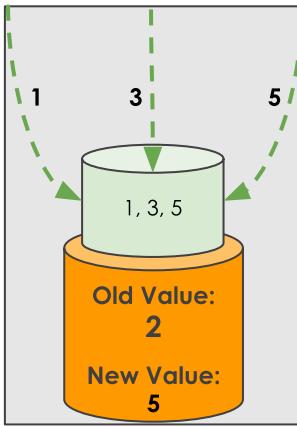
# Accumulators

The GSQL language provides many different accumulators, which follow the same rules for receiving and accessing data. However, each of them has their unique way of **aggregating values**.



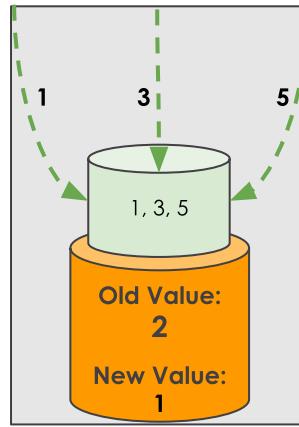
`SumAccum<int>`

Computes and stores the cumulative sum of numeric values or the cumulative concatenation of text values.



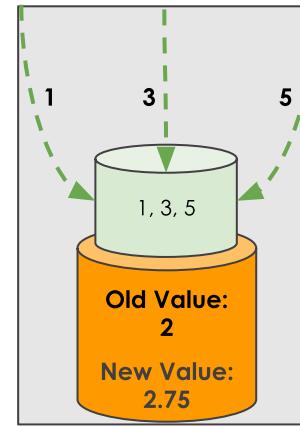
`MaxAccum<int>`

Computes and stores the cumulative maximum of a series of values.



`MinAccum<int>`

Computes and stores the cumulative minimum of a series of values.



`AvgAccum`

Computes and stores the cumulative mean of a series of numeric values.

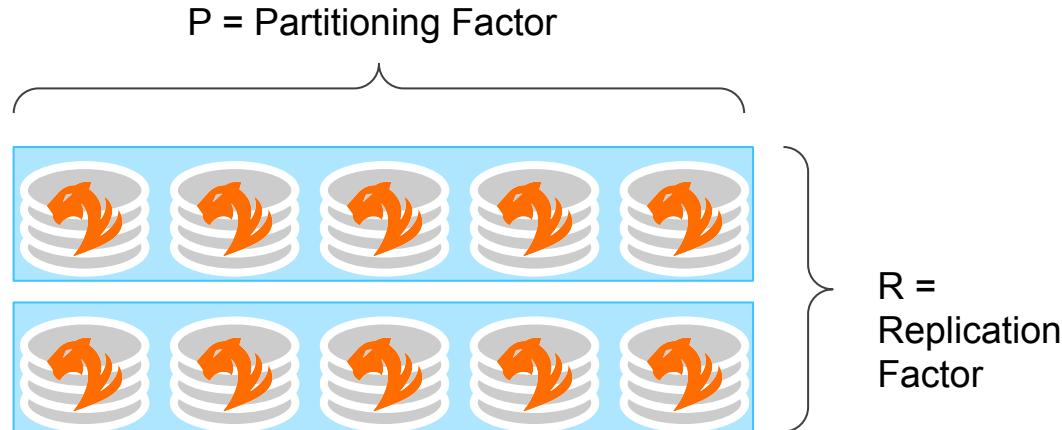
# NON-FUNCTIONAL FEATURES

- High Availability
- Access Control & Security
- Transaction Management



# High Availability

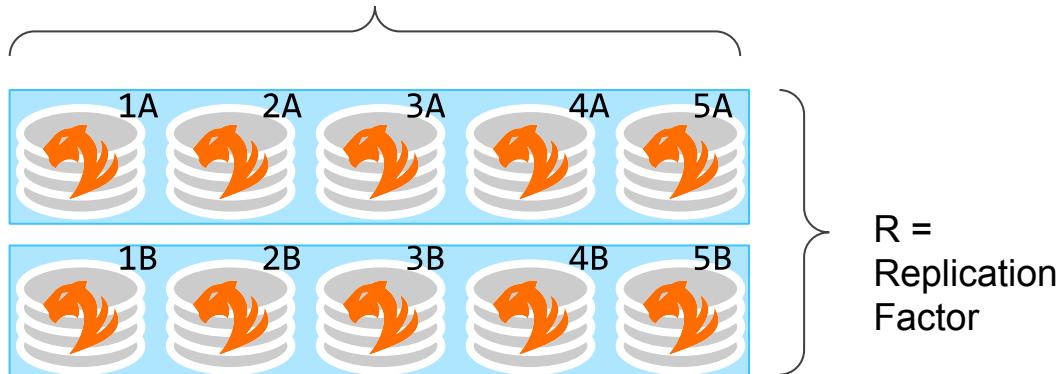
- TigerGraph HA Replication provides both Increased Throughput and Continuous Operation
- Cluster size =  $P \times R$  (Partitions x Replicas)
- Any cluster size is allowed, except 1x2



# HA and Concurrency

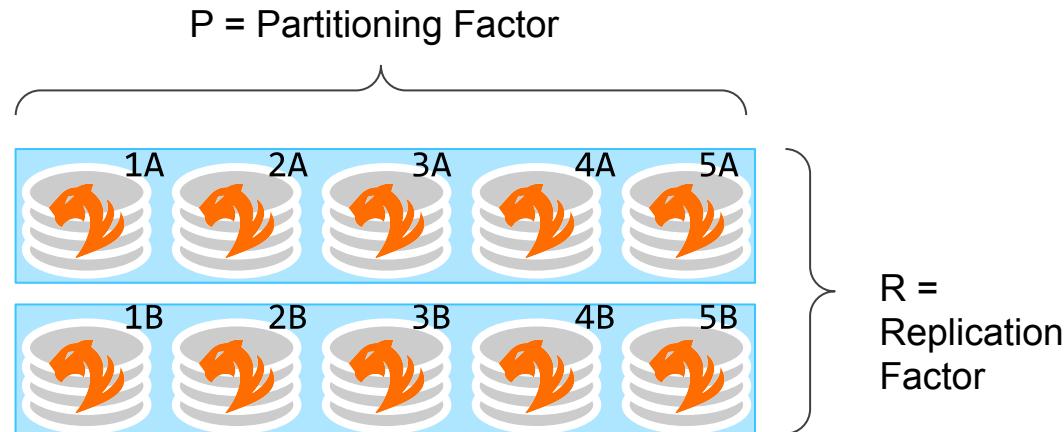
- Each server has  $T$  available workers for serving requests (GSQL query, REST POST, etc.)  
 $T$  is a system configuration parameter, defaults to 8. Consider number of CPU cores.
- Cluster's total number of workers =  $T \times P \times R$ , e.g.  $8 \times 5 \times 2 = 80$ 
  - A point mode query uses 1 worker.
  - A distributed mode query use  $P$  workers.

$P$  = Partitioning Factor



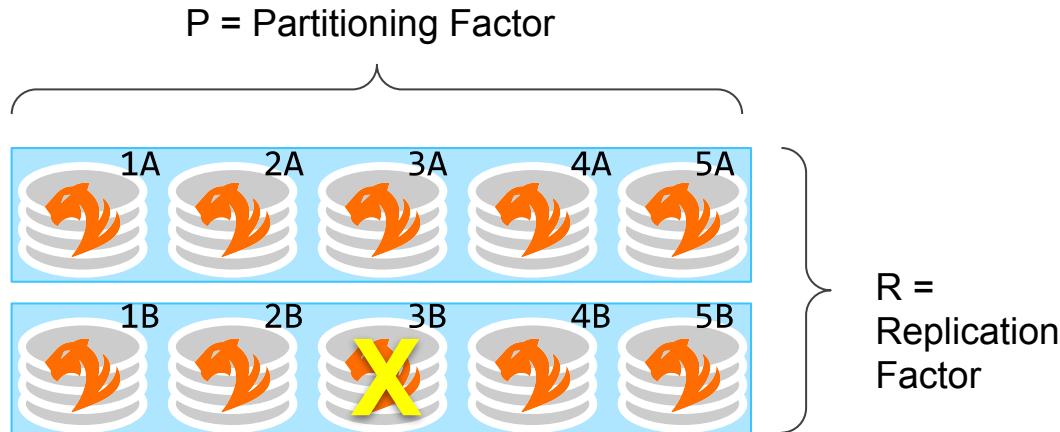
# HA Read and Write Behavior

- All Replicas are Read/Write, always in sync with the latest updates
- Writes go to all replicas (e.g. both 1A and 1B).
- Reads can be from any one replica (e.g. either 1A or 1B).
- Distributed queries can mix replicas (e.g. {1A, 2B, 3B, 4A, 5B} is a valid active set for a request.)



# HA Continuous Operation

- If any single server is unavailable (expected or unexpected):
  - When it fails to respond after a certain number of tries, requests will automatically divert to another replica (e.g. 3B is unavailable, so use 3A)
  - If it fails in the middle of a transaction, that transaction might be aborted.
- System continues to operation, with reduced throughput, until server is restored.



# NON-FUNCTIONAL FEATURES

- High Availability
- Access Control & Security
- Transaction Management



# Role-Based Access Control

- Follows SQL approach for roles.
  - GSQL:  
`GRANT <role> ON GRAPH <graph> TO <user1, user2, ... >`  
`REVOKE <role> ON GRAPH <graph> FROM <user1, user2, ... >`
- Can map TigerGraph roles to external LDAP roles and groups.

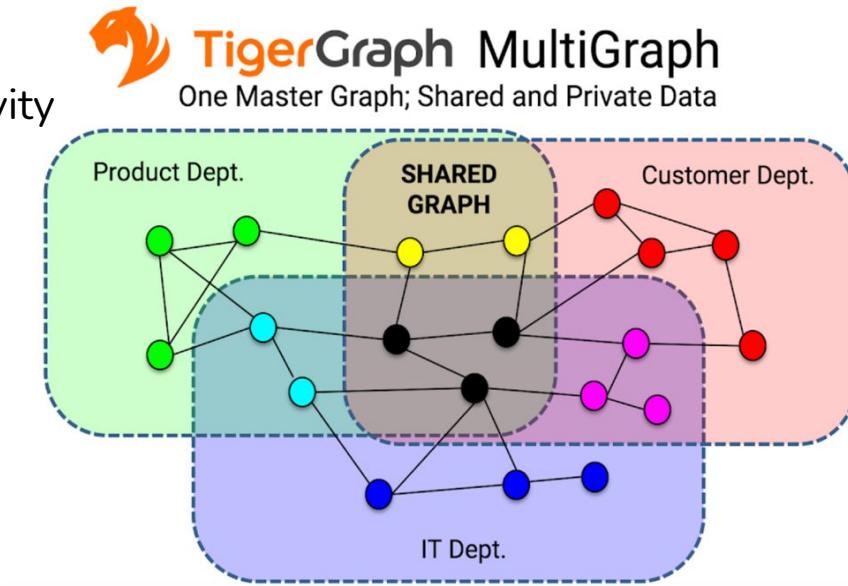
# Admin Portal UI for Managing User Privileges

The Admin Portal UI for Managing User Privileges is a web-based application designed to facilitate the management of user roles and privileges. The interface is organized into several key sections:

- Left Sidebar:** Contains navigation links for Dashboard, User Management (which is the currently selected option), and License.
- Header:** Includes a logo for "AdminPortal", a three-line menu icon, a notification bell icon, a user profile icon, and a "GraphStudio" button.
- Main Content Area:** Features a "Role Management" section with tabs for "My Profile" and "All Users". The "All Users" tab is selected, showing a list of roles and their respective users:
  - Super User (2):** tigergraph, Duc
  - Admin (2):** Emily, Lily
  - Designer (1):** James
  - Query Writer (5):** Lily, Dan, Amanda, Linda, Luke
  - Query Reader (1):** Lily
  - Observer (1):** Tom
- Search and Filter:** A "Search user" input field and a "C" refresh/clear button.
- Bottom Status Bar:** Displays a "Cluster Service Status" indicator (green circle) and a power button icon.

# MultiGraph for RBAC and Data Sharing

- **Share & Collaborate**
  - Multiple groups share one master database  
⇒ data integration, insights, productivity
- **Real-time, Updatable**  
Shared updates, no copying  
⇒ cleaner, faster, cheaper, safer
- **Fine-Grained Security**
  - Each group is granted its own view
  - Each group has its own admin user, who manages local users' privileges.



# Roles and Privileges

## Built-In Roles:

- **Queryreader:** run existing loading jobs & queries for its assigned graph.
- **Querywriter:** Queryreader privileges + create queries and run data-manipulation commands on its assigned graph.
- **Designer:** Querywriter privileges + modify the schema, create loading jobs for its assigned graph.
- **Globaldesigner:** Designer privileges + create global schema, create objects. Also, delete graphs which they created.
- **Admin:** Designer privileges, + create/drop users, grant/revoke roles for its assigned graph. That is, control existence & privileges of its local users.
- **Superuser:** admin privileges on all graphs. Create global vertex & edge types, create multiple graphs, and clear the database.

## User-Defined Roles: in development

Command Type	Operations	super-user	admin	global-designer	designer	query-writer	query-reader
Status	Ls	x	x	x	x	x	x
User Management	Create/Drop User	x	x	-	-	-	-
	Show User	x	x	x	x	x	x
	Alter (Change) Password	x	x	x	x	x	x
	Grant/Revoke Role	x	x	-	-	-	-
	Create/Drop/Show Secret	x	x	x	x	x	x
	Create/Drop/Show/ Refresh Token (Deprecated)	x	x	x	x	x	x
Schema Design	Create/Drop Vertex/Edge/Graph	x	-	x	-	-	-
	Clear Graph Store	x	-	-	-	-	-
	Drop All	x	-	-	-	-	-
	Use Graph	x	x	x	x	x	x
	Use Global	x	x	x	x	x	x
	Create/Run Global Schema_Change Job	x	-	x	-	-	-
	Create/Run Schema_Change Job	x	x	x	x	-	-
Loading and Querying	Create/Drop Loading Job	x	x	x	x	-	-

# Data Encryption

- Encrypted Data at Rest
  - Choice of encryption levels (file, volume, partition, disk)
    - Kernel level: dm-crypt / cryptsetup
    - User level: FUSE (Filesystem in User Space)
  - Automatically encrypted in TigerGraph Cloud
- Encrypted Data in Transit
  - Can set up SSL/TLS for HTTPS protocol
  - Automatically encrypted in TigerGraph Cloud

# NON-FUNCTIONAL FEATURES

- High Availability
- Access Control & Security
- Transaction Management



# Transactional Model

- The TigerGraph distributed database provides full ACID transactions with sequential consistency
- Transactions definition:
  - Each GSQL Query procedure is a transaction. Each query may have multiple SELECT, INSERT, or UPDATE statements.
  - Each REST++ GET, POST, or DELETE operation (which may have multiple update operations within it) is a transaction.

# ACID Compliance

Atomicity	Consistency	Isolation Level	Durability
<p>GSQL query w/ or w/o updates = Transaction</p> <p>Transactions are “all or nothing”: either all changes are successful, or none is successful.</p>	<p><b>Single-server Consistency:</b> A transaction obeys data validation rules.</p> <p><b>Distributed System Sequential Consistency:</b> Every replica of the data performs the same operations in the same order.</p>	<p><b>Repeatable Read:</b></p> <ul style="list-style-type: none"><li>Each transaction sees the same data.</li></ul> <p><b>No Dirty/Phantom Read:</b></p> <ul style="list-style-type: none"><li>A transaction's updates are not visible to other transactions until the update is committed.</li></ul>	<p>The TigerGraph platform implements write-ahead logging (WAL) to disk to provide durability.</p> <p>Logs are consumed periodically to update the database on disk.</p>

# Resources

TigerGraph Platform Overview:

<https://docs.tigergraph.com/intro/tigergraph-platform-overview>

HA Cluster Configuration:

<https://docs.tigergraph.com/admin/admin-guide/installation-and-configuration/ha-cluster>

Transaction Processing and ACID Support:

<https://docs.tigergraph.com/dev/transaction-and-acid>

MultiGraph:

<https://docs.tigergraph.com/intro/multigraph-overview>



# Get Started for Free

- Try [TigerGraph Cloud](#)
- Download [TigerGraph's Developer Edition](#)
- Take a [Test Drive - Online Demo](#)
- Get TigerGraph [Certified](#)
- Join the [Community](#)



@TigerGraphDB



/tigergraph



/TigerGraphDB



/company/TigerGraph



# BACKUP