LDBC TUC

June 24, 2023

David Haglin, Trovares Inc. *david@trovares.com*

info@trovares.com

TROVARES
seriously scalable

Agenda

- Trovares app Architecture
- LDBC Features for Trovares
- Example Usage
- Graphalytics for “the commercial space”
- Discussion

Trovares xGT Architecture

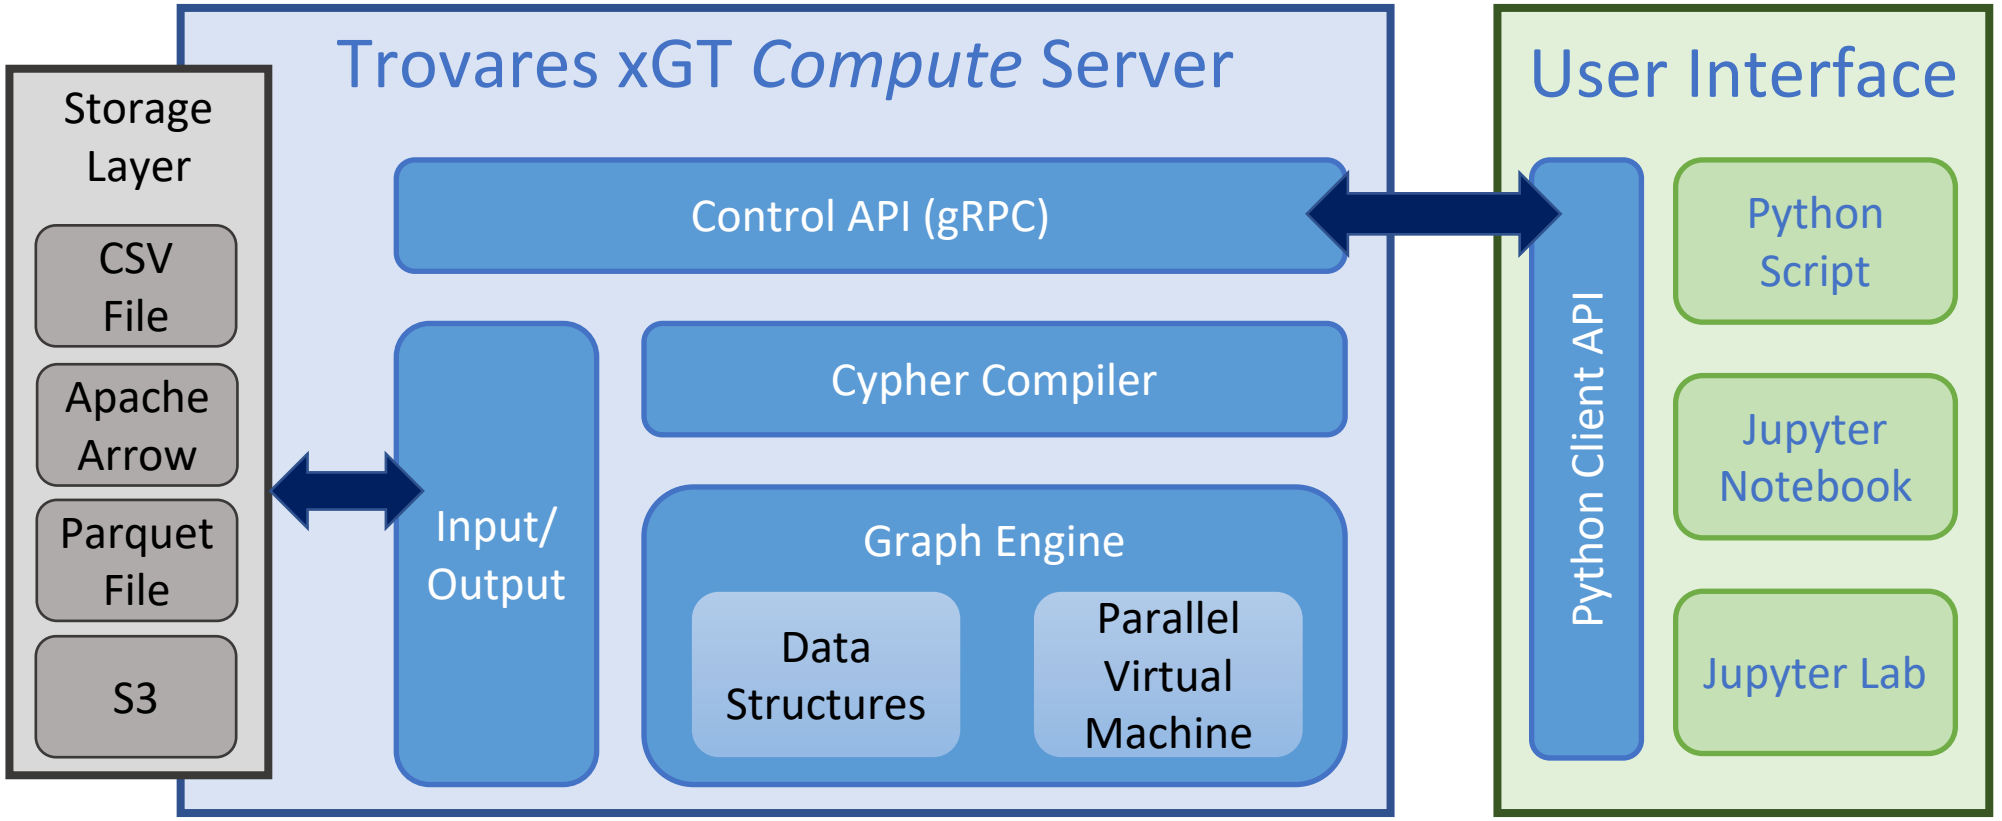

- xGT Server runs a a single-node server, 1GB - 1 Core to 24TB - 572 Cores
- Server designed and implemented by seasoned HPC programmers

TROVARES seriously scalable

Trovares xGT Server Features

- Property Graph analytics engine
- Runs open Cypher
- Includes calls to some whole-graph algorithms:
  - BFS
  - WCC / SCC
  - PageRank, Personalized PageRank
- Not a Database, but provides ACI of the ACID database guarantees
- Ingest is a transaction
- Focus has been computationally-hard queries on large datasets

LDBC Features for Trovares

Datasets

- “real-world synthetic”
- Multi-typed
- Reasonably large

BI Queries

- Interesting starting suite of queries
- Perhaps too many “seeded” queries

Usage Example: Hardware platform comparison

- Use four different size of the Linked Data Benchmark Council (LDBC) synthetic Social Network Benchmark (SNB) data.
- This is not an LDBC Benchmark results: it uses only the LDBC data.
- Statistics about the four dataset sizes:

Scale Factor	Num Vertices	Num Edges	RAM Size (GiBs)
1	3,181,734	17,298,797	2
10	29,987,835	176,908,026	23
100	282,637,871	1,777,459,231	222
1000	2,686,781,095	17,802,784,542	2,422

Usage Example: Analytic Queries

- Trovares developed several “Analytic Queries” (AQ) against the LDBC SNB data. https://datasets.trovares.com/synthetic/ldbc/index.html
  - **AQ1: Three people meeting at some event such as a conference**
  - **AQ2: 10 most introverted countries**
  - **AQ3: A commenter paying for Likes**
  - **AQ4: Super-fans**
- We ran these 4 queries against all sizes of data sets shown above on:
  - HPE DL385: one socket, 32 AMD cores, 64 vCPUs, 256 GB RAM
  - One socket of Flex280: one socket, 28 Intel cores, 56 vCPUs, 750 GB RAM
  - Flex280: 8 sockets, 224 Intel cores, 448 vCPUs, 6 TB RAM

Benchmark: Results

DL385

Scale	AQ1	TE1	AQ2	TE2	AQ3	TE3	AQ4	TE4
1	0.12	4.44	2.07	177.19	0.08	1.49	0.11	2.94
10	0.77	82.85	35.13	4,678.07	0.52	20.92	1.18	40.58
100	14.36	1,529.94	846.03	93,314.73	5.92	254.51	15.33	490.76

Flex280, 1 socket

Scale	AQ1	TE1	AQ2	TE2	AQ3	TE3	AQ4	TE4
1	0.04	4.44	0.74	177.19	0.05	1.49	0.06	2.94
10	0.33	82.85	8.01	4,678.07	0.30	20.92	0.72	40.58
100	5.68	1,529.94	206.97	93,314.73	3.88	254.51	8.73	490.76

Flex280, 8 socket

Scale	AQ1	TE1	AQ2	TE2	AQ3	TE3	AQ4	TE4
1	0.12	4.44	1.51	177.19	0.15	1.49	0.20	2.94
10	0.13	82.85	2.52	4,678.07	0.91	20.92	1.82	40.58
100	0.97	1,529.94	43.00	93,314.73	8.88	254.51	20.32	490.76
1000	19.42	31,491.88	836.45	1,411,958.68	91.16	2,848.24	210.99	5,634.71

* AQx Time in seconds, TE units 1 million traversed edges

Graphalytics for “the commercial space”:
some ideas to consider …

- Ditch the Java implementation of the test harness
- Adjust required output to align with real-world expectations:
  - Nobody really wants a BFS, they want their own variant of a BFS
  - Nobody cares about the nodes not reachable from the selected root
- Graph products for commercial space require periodic updates to graph data structure. Benchmarks should embrace this. Example:
  - Read in 90% of all edges, run BFS
  - Read in next 5% of the edges, run BFS
  - Read in next 5% of the edges, run BFS
  - Validate all three BFS results (on only reachable nodes)
  - Capture time from just after ingesting 90% to the end.

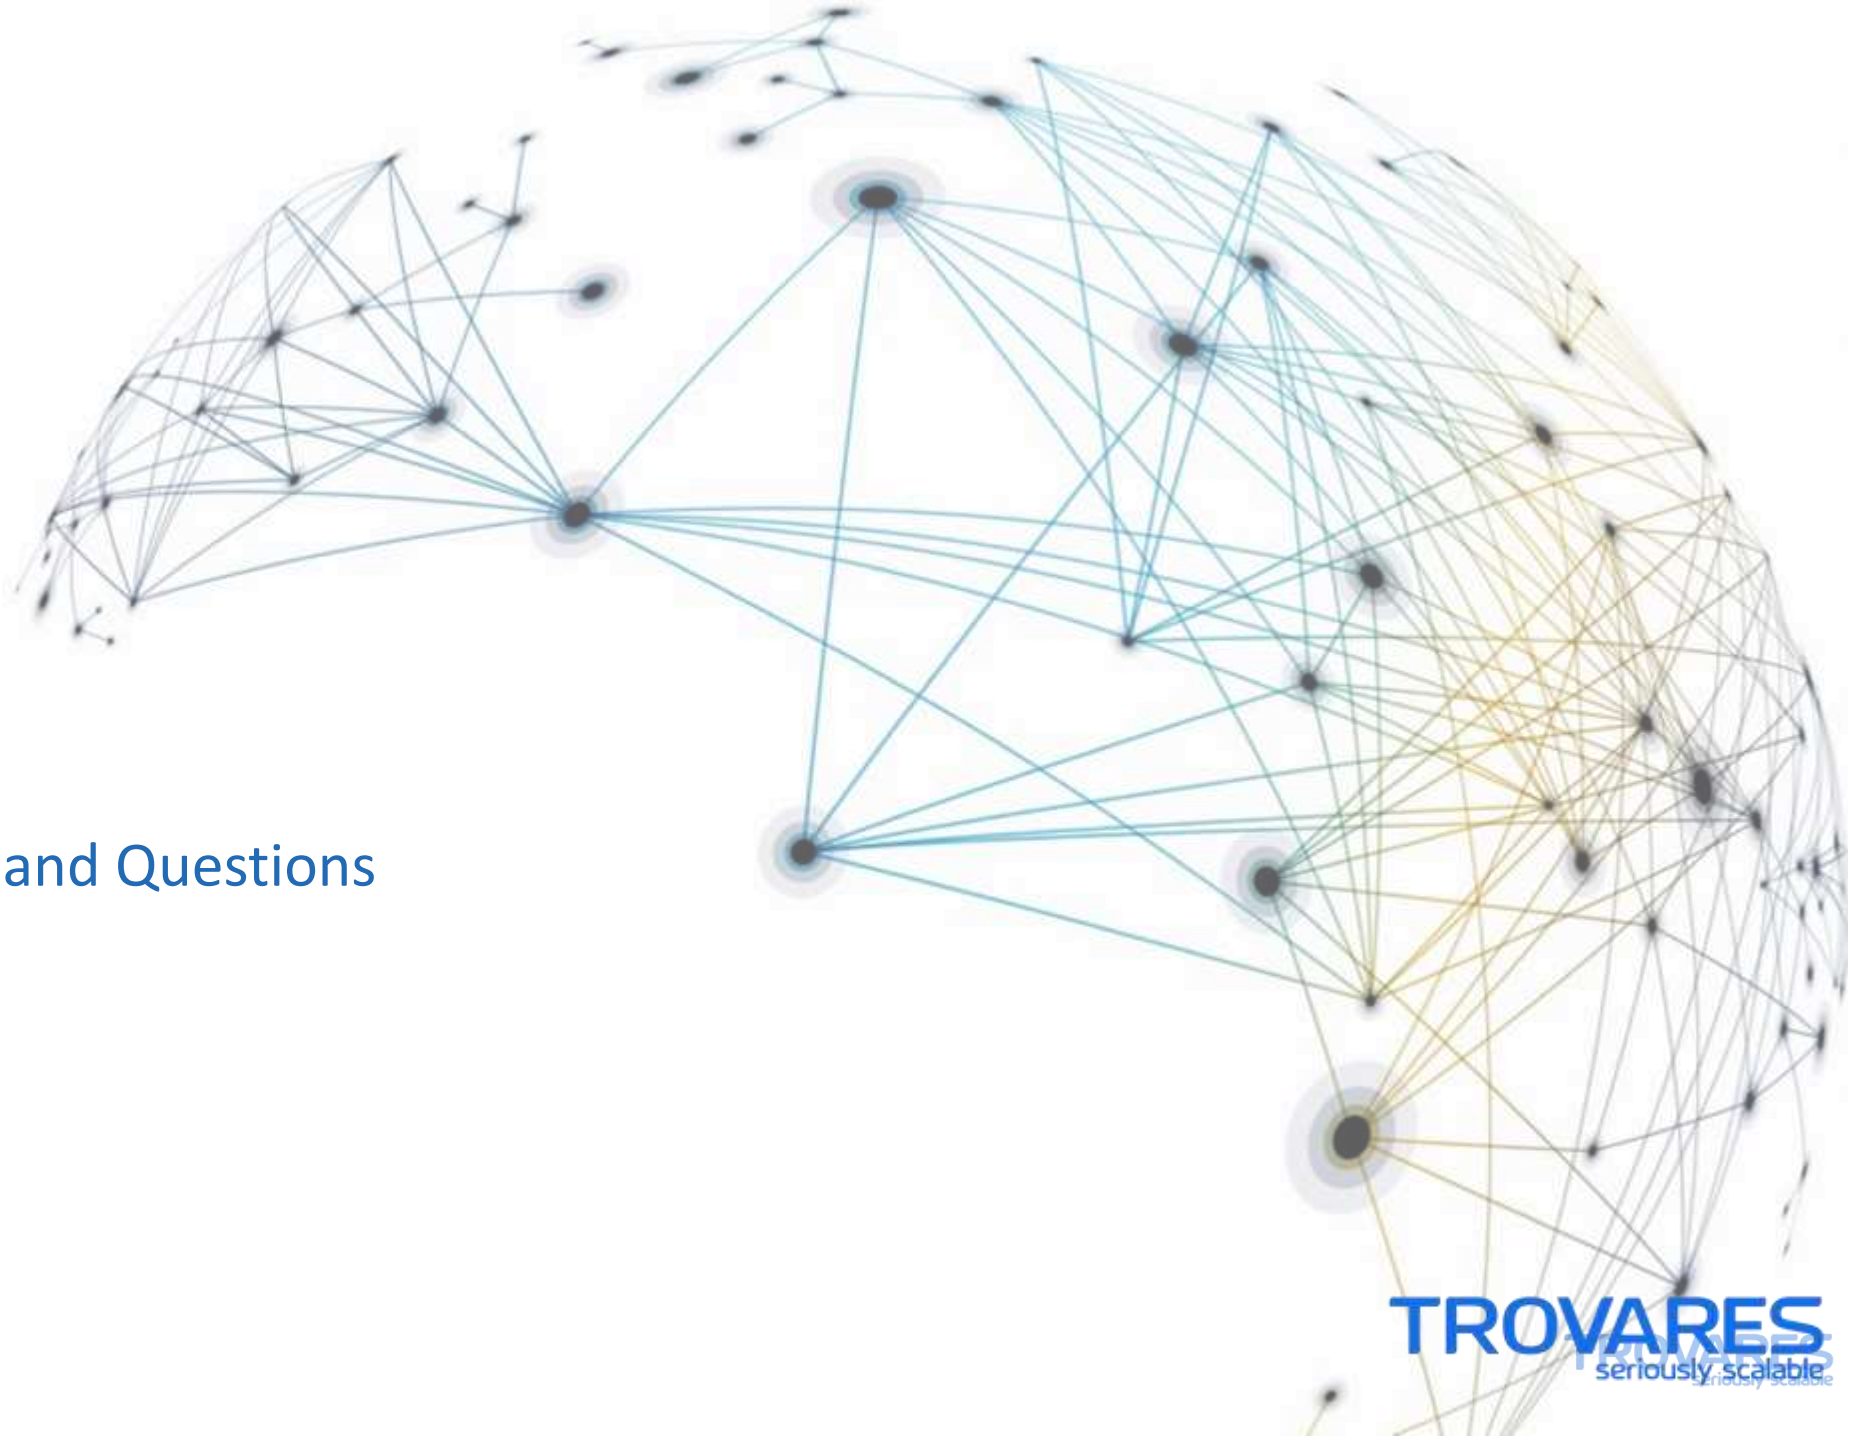

Discussion and Questions

info@trovares.com

TROVARES
seriously scalable