

Accelerating Cypher queries with Trovares xGT on IBM Power10

Combining Analytics with IBM's AI Server technology.



Accelerating Cypher queries with Trovares xGT on IBM Power10

Trovares xGT is a graph analytics search engine that handles the most significant and complex graph problems. Trovares' patented xGT graph technology allows data to be unified from multiple sources throughout the organization at ingest speeds that exceed conventional graph tools.¹ Trovares xGT search engine ingests data from Oracle, MongoDB, Databricks, Snowflake, and many other sources like TigerGraph or Neo4j™, allowing organizations to add enterprise-wide knowledge capabilities through scalable graph analytics without replacing or altering their existing data analytics workflow.

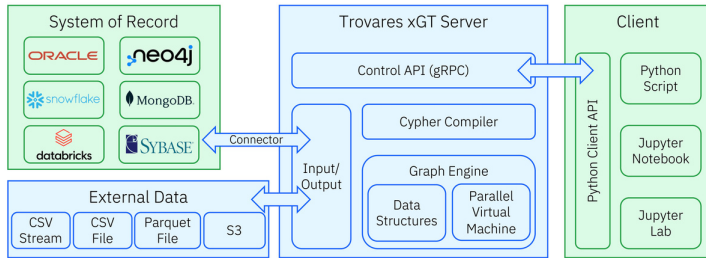


Figure 1. Trovares xGT Architecture

Leveraging Trovares xGT Search Engine for AI Applications IBM Power10

The Trovares xGT search engine enhances the training of AI models. Leveraging the search engine, clients can begin identifying significant motifs and patterns from substantial datasets that span extended periods. The data findings are then passed to AI models for training, helping them know what to look for. Conversely, if AI is used to identify abnormalities, clients can leverage the Trovares xGT search engine to determine if such patterns transpired in the past.

Running on IBM Power10

IBM® Power10 servers are superscalar, multithreaded, multi-core servers with embedded AI acceleration technology. Power10 is the ideal platform for high performance workloads like graph analytics. Running Trovares xGT search engine on Power10 offers many computational performance advantages over x86-based systems,² scaling up to 64TB of shared memory.

Power10 vs x86

Quantifying the capabilities of IBM Power®, the Trovares engineering team conducted several performance tests comparing xGT Accelerator on x86 and the Power E1050. The results revealed that Power10 delivered the fastest results, 60% faster when traversing a 500 million edge graph, 92% faster when traversing a 1 billion edge graph, and 81% faster when traversing a 4 billion edge graph.

Comparing Trovares xGT Accelerator on x86 (8 cores) and Power E1050 (8 cores)

Dataset Size	Edge Count	Trovares xGT Accelerator on x86 (8 cores)		Trovares xGT Accelerator on Power E1050 (8 cores)		% Faster on Power10
		Ingest (seconds)	Query (seconds)	Ingest (seconds)	Query (seconds)	
Scale 25	~ 500 million	119	77	113	48	60%
Scale 26	~ 1 billion	292	123	191	64	92%
Scale 28	~ 4 billion	2611	912	912	358	81%

1. Based on Trovares internal testing of Trovares xGT Accelerator on AWS x1e.4xlarge with 8 Intel Xeon cores (16 vCPUs) and Trovares xGT Accelerator on a Power E1050 LPAR with 8 cores and 2TB of RAM. A 500 million-edge graph, 1 billion-edge graph, and 4 billion-edge graph were used to run the queries. "Percent Faster on Power10" quantifies the difference in "Query seconds" rounded to the nearest percent. Individual results may vary.

Quantifying the capabilities of Power10 even further, the Trovares engineering team conducted additional performance tests. Doubling the number of x86 cores, Trovares xGT Accelerator ran on 16 cores of x86 and compared to Trovares xGT Accelerator running on 8 cores of Power E1050. The results revealed that running on 8 cores of Power E1050 was comparable to 16 cores of x86 when traversing a 500 million edge and a 1 billion edge graph.

Comparing Trovares xGT Accelerator on x86 (16 cores) and Power E1050 (8 cores)

Dataset Size	Edge Count	Trovares xGT Accelerator on x86 (16 cores)		Trovares xGT Accelerator on Power E1050 (8 cores)	
		Ingest (seconds)	Query (seconds)	Ingest (seconds)	Query (seconds)
Scale 25	~ 500 million	90	45	113	48
Scale 26	~ 1 billion	196	68	191	64

1. Based on Trovares internal testing of Trovares xGT Accelerator on AWS x1e.4xlarge with 16 Intel Xeon cores (32 vCPUs) and Trovares xGT Accelerator on a Power E1050 LPAR with 8 cores and 2TB of RAM. A 500 million-edge graph, 1 billion-edge graph, and 4 billion-edge graph were used to run the queries. Individual results may vary.

Reducing an extensive dataset of billions to a small subset of answers is the value of Trovares xGT Accelerator. Affording Trovares the ability to rapidly search the datasets is the value of Power10.



Cross-Industry Benefits

Open Source Connector

xGT supports the Cypher™ query language, commonly used across industries and major commercial enterprises. xGT offers industries, using the Cypher query language, an immediate path to run complex queries at a planetary scale.

Support for Hybrid Environments

Power10 can scale to run multiple databases on the same on premises server and access servers in the cloud. Organizations can run Cypher workloads faster and more efficiently by co-locating the Trovares search engine with existing databases. Extensive datasets can be searched locally for speed and efficiency, while the cloud can be leveraged for less time-sensitive searches.

Cross-Industry Use Cases Complementing AI with Trovares

Cybersecurity

Network traffic is analyzed to identify potential security threats, modeling the relationships between various entities, such as devices, users, and applications, to detect anomalies and unusual patterns in network behavior.

Healthcare Fraud Detection

Connections and relationships are analyzed within healthcare networks to detect fraudulent activities, such as billing fraud or identity theft, and to identify anomalies in healthcare data that may indicate fraudulent claims or abuse of the system.

Infrastructure and Facility Management

Model relationships between different components of critical infrastructure (such as power grids, transportation systems, and communication networks) to optimize maintenance, identify vulnerabilities, and plan for resilience.

Customer Relationship Management (CRM)

Build customer graphs to model customer relationships, purchase history, preferences, and interactions. CRM helps create targeted marketing campaigns and personalized recommendations and improves customer satisfaction.

Supply Chain Optimization

Analyze suppliers, distributors, and logistics relationships to optimize the supply chain. Graph analytics can help identify bottlenecks, streamline inventory management, and improve efficiency.

Fraud Detection and Prevention

Create graphs representing relationships between transactions, customers, and payment methods. Graph analytics can help identify unusual patterns indicative of fraud and enhance security measures.

Customer Journey Analysis

Map the customer journey by analyzing interactions across different touchpoints, such as online and offline channels, and helping retailers understand the complete customer experience and improve accordingly.

Price Optimization

Analyze relationships between pricing strategies, competitor prices, and customer behavior to optimize pricing and promotions. Graph analytics can aid in identifying optimal price points for products.

Quality Control and Defect Analysis

Model relationships between manufacturing processes, components, and quality inspection data to identify patterns related to defects. Graph analytics can help in quality control and proactive defect prevention.

Supply Chain Visibility

Analyze relationships within the supply chain to enhance visibility and traceability of products and components. This is particularly important for compliance, quality assurance, and responsiveness to disruptions.

For more information

To learn more about the solution, contact your IBM representative or IBM Business Partner, or visit <https://www.trovares.com>.



1. [Graph Search](#), Trovares xGT, 2024

2. Based on Trovares internal testing comparing query speeds. Individual results may vary.

