

WHITE PAPER

Uncovering Business Value in Connected Data

More and Better Business Insights with Leading-edge
Technologies from AMD, Dell, and TigerGraph

Sponsored by



DELLTechnologies





Table of Contents

Introduction	3
The Problem with Legacy Tools	4
The Solution – A TigerGraph Database Running on Dell PowerEdge Servers Powered by AMD EPYC™ Processors	4
Use Cases	6
Summary	7

Introduction

Data insights are fundamental to creating efficient business operations, discovering new revenue opportunities, and gaining competitive advantages. Enterprises across the world are looking to mature their analytics capabilities and become more data-driven. Graph databases and analytics are at the forefront of advanced analytic technologies. Graph analytics allows for deep exploration of connected data, tracing the complex interrelationships among various entities, such as organizations, people, transactions, records, but most importantly, in context.

Graph analytics is one of the fastest growing trends in data and analytics and will grow at 100 percent annually through 2022 and beyond, [according to Gartner](#). By 2025, graph technologies will be used in 80% of data and analytics innovations, up from 10% in 2021, facilitating rapid decision making across the enterprise.

Graph technologies can be used to rapidly analyze thousands of data sources with billions of elements. Because of the complexity and variety of data, and the number and depth of interconnections between multiple data sets, traditional data structures of relational and NoSQL databases are not always a good fit. Graph analytics excels at answering complex questions about relationships in large data sets, which is not always practical or even possible at scale using Simple Query Language (SQL) queries.

The Problem with Legacy Tools

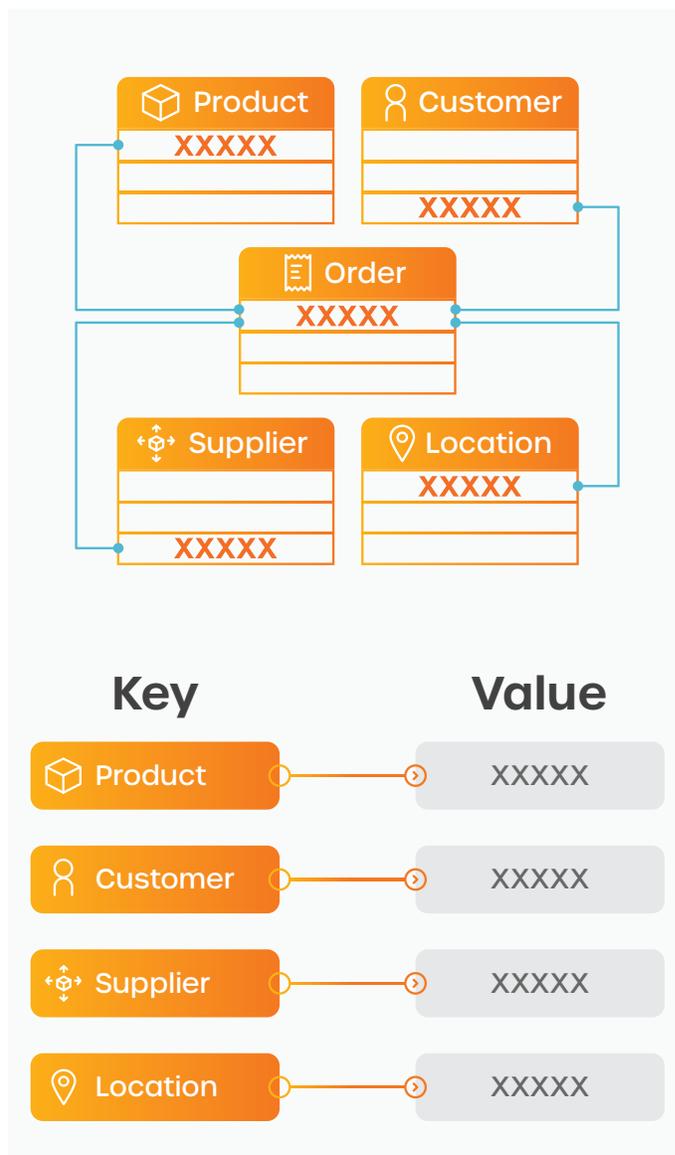
In the world of data analytics and data management, relational databases and non-relational databases have been around for a long time. Relational databases work well for structured data with defined relationships that can be organized in a tabular format and rely on SQL to extract information. Non-relational databases, most often called NoSQL for short, tackle the challenge of storing unstructured data that is fluid and doesn't conform to a predefined data model or relational schema. Examples of unstructured data include text, emails, images, and videos. Both relational and NoSQL databases are good for managing large volumes of data. However, they were not designed to generate new insights from data where relationships between data points matter more than individual points themselves. For this we need another type of database – a graph database.

Relational Databases

Relational databases work well for managing predefined data relationships across multiple linked tables of information and are good for analyzing data with a finite set of pre-defined outcomes. Relational databases are great for transactions, indexing and searching data, and basic analysis. However, the rigid schema of relationship database systems makes it difficult to quickly add new connections or adapt to changing business requirements. Relational databases store data for each business entity such as customer, order, product, and payment data in separate database tables. To understand and analyze relationships across the business entities, relational databases require table joins, which can take hours, even days for the complex joins, and are computationally expensive as the size of the data grows.

NoSQL Databases

The key-value NoSQL databases effectively have no schema but also provide high performance for simple transactions. The main difference between relational and NoSQL databases is that relational databases are normalized. That is, they store data in a tabular form, arranged in a table with rows and columns. A non-relational database, or NoSQL database, stores data as a set of key-value files in a data lake. Each key-value pair constitutes a row. Queries and algorithms to perform relationship analysis require iterative full-table scans across millions or billions of rows, making it very difficult to perform a deeper analysis of the relationships beyond two or three levels.



The Solution – A TigerGraph Database Running on Dell PowerEdge Servers Powered by AMD EPYC™ Processors

Another kind of database is needed to analyze relationships and extract insights from data, and that is a graph database. A graph is composed of two elements: a node and a relationship. Each node represents an entity (a person, place, thing, category, or another piece of data), and each edge represents a relationship of how two or more nodes are connected. This general-purpose structure allows you to model all kinds of scenarios – a network of people or devices, a population's medical history, a company's supply chain, or anything else defined by relationships. Because the graph database has pre-connected entities and a flexible schema, it has high performance both for complex transactions and for deep flexible analytics.



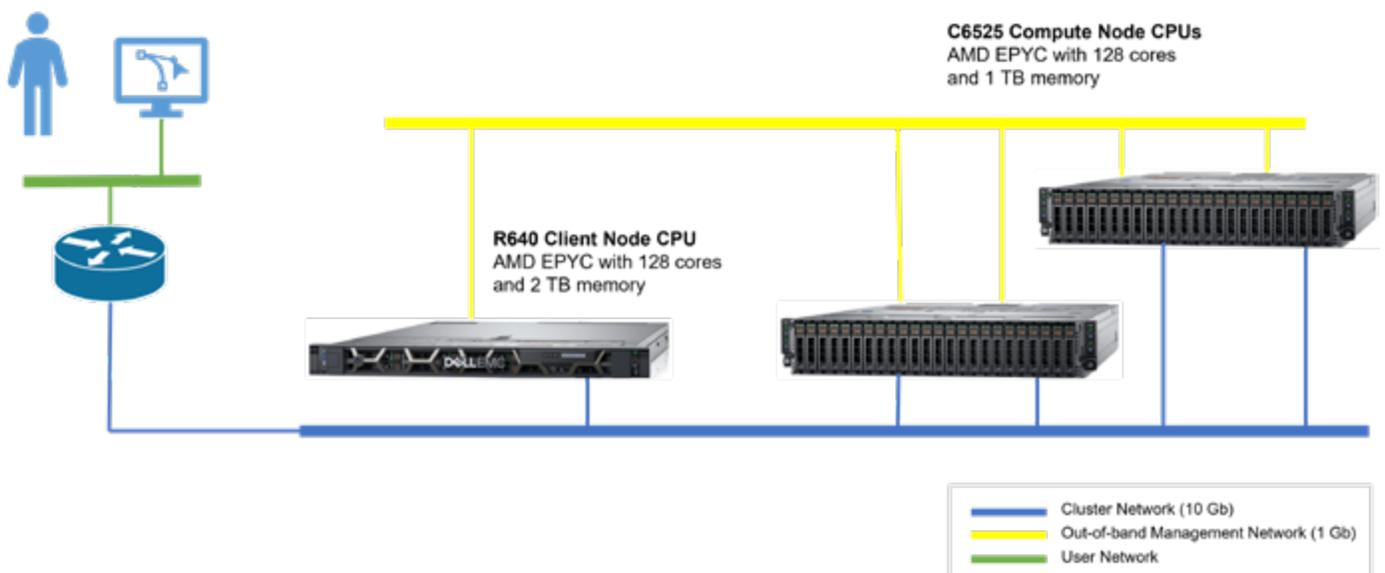
Tech giants like Google, Facebook, LinkedIn, and PayPal all tapped into the power of graph databases to create booming businesses by harnessing the power of data connections. Now your business can leverage graph databases to their advantage with TigerGraph running on Dell PowerEdge servers powered by the AMD EPYC™ family of processors.

TigerGraph is a scalable graph database system built from the ground up to support massively parallel computation of queries and analytics. TigerGraph is a native parallel graph database, with its proprietary storage designed from the ground up to store graph nodes, edges, and their attributes, with an engine that computes queries and analytics in massively parallel processing (MPP) fashion for significant scale-up and scale-out performance. Key benefits of a TigerGraph database running on Dell and AMD infrastructure include:

- [Perform better, faster queries and analytics](#) than relational and NoSQL systems
- [Model data simply and naturally](#), removing the need for data normalization
- [Improve the accuracy of AI/ML models](#) by generating relationship data on links between objects
- Achieve continuous availability and [reduced memory footprint](#) with the [massive scale](#) of distributed architecture
- Simplify deployment and accelerate time to insight thanks to in-database analytics
- Migrate from relational database systems, no code required
- Use any cloud, including private cloud or a traditional data center

The beauty of a TigerGraph / Dell / AMD solution is that it is built for high performance while being efficient, secure, and able to be deployed on-premises with just a handful of servers. Here is an example deployment for a very large healthcare organization with 104 million patients and 25,000 parallel queries:

- 1x Client Node, consisting of a Dell EMC PowerEdge R6525 rack server with AMD EPYC processors, 128 cores, and 2 TB of memory
- 8x Compute Nodes, consisting of Dell EMC PowerEdge C6525 C-series servers with AMD-EPYC 7702 processors, 128 cores, and 1 TB of memory each



With a record retrieval time of 10 ms against a 104 million patient database, the healthcare **provide achieved an improvement of 10 times the prior performance**, resulting in a more cost effective solution and just as importantly a significant improvement in customer service.

Use Cases

Financial Services - Fraud Detection - MasterCard

With over 50 billion transactions annually, MasterCard uses network-level fraud monitoring to reduce cost and safeguard against fraud events. For the last 15 years, MasterCard's data science teams have been developing capabilities for their fraud detection system based on time-series data using legacy tools. This has become unscalable. MasterCard turned to TigerGraph to create new fraud detection capabilities by taking transaction data and representing the transaction data as a graph - consumers and merchants are graph nodes and transactions are graph edges. Thanks to AMD's exceptional per-core performance and AMD Infinity Guard security features, MasterCard is able to speed data analytics with the confidence of leading security features.

With TigerGraph, running on Dell and AMD EPYC infrastructure, MasterCard was able to reduce fraud false positives by 60% and avoid over \$5 billion in losses.

Telecommunications - A Major ISP and Mobile Carrier

A major telecommunications conglomerate in Asia Pacific wanted to provide better customer service on its network. Due to a lack of insight to user-level and grid-level performance, the telecom company's response to customer network complaints was slow, resulting in low customer satisfaction and higher risk of churn. The company's existing technology stack did not support complex and dynamic data joins, nor was it able to traverse into datasets to easily provide performance scoring. The company's existing solution required multiple systems and heavy coding to achieve insights, and it had no near-real-time capabilities for impact analysis. The telecom was unable to combine and correlate customer- and cell-level performance details at an hourly interval. Without such details, it was unable to pinpoint root causes. Using TigerGraph and running on Dell PowerEdge servers powered by AMD EPYC processors, the telecom conglomerate built a digital twin of its network that shows real-time status across 5G and broadband

systems, improving network quality and customer satisfaction.

Marketing & Sales - Entity Resolution / Customer 360 - Xandr

Xandr is the advertising and analytics division of AT&T's Warner Media. The company operates an online platform called Community that connects advertisers, publishers, and consumer media brands. TigerGraph running on Dell servers, powered by AMD EPYC processors, is enabling Xandr to merge data from silos across the Warner Media universe and build the first, and largest, identity graph of its kind in the advertising industry. TigerGraph's ability to integrate datasets and resolve ambiguous entities at scale is allowing Xandr to uncover insights hidden in its data. Specifically, Xandr can identify attributes associated with people, devices, and households that provide marketers the ability to target audiences with customized commercials aligned to their interests and needs, resulting in better advertising performance.

Healthcare - UnitedHealth Group

UnitedHealth Group (UHG) is the largest healthcare company in the world by revenue. UHG has almost 50 million medical members and employs over 181,000 people. UHG is looking to improve the quality of care to its members, while reducing the associated costs. UHG identified graph databases and analytics as the best ways to do this and selected a joint solution of TigerGraph running on Dell hardware, powered by AMD EPYC processors. 'Hardened at the core' with the AMD Secure Processor, the AMD EPYC family of processors helps customers protect their most important assets - their data.

TigerGraph is enabling UHG to build the largest healthcare graph with over 10 billion vertices and 50 billion edges. It currently contains 1.2 terabytes of data that includes information on its members, providers, claims, visits, prescriptions, procedures, and more. Over 23,000 users access the graph each day on a variety of applications. TigerGraph empowers UHG's medical providers to make real-time care-path recommendations using knowledge of 50 million patients. UHG has implemented TigerGraph in its contact centers and estimates \$150 million in savings a year by enabling its medical professionals to provide accurate and effective care path recommendations in real time.

Cybersecurity - Leader in Cyber Resilience

A market leader in cyber resilience, the company is one of the first companies to harness the cloud and AI to stop cyber threats in real-time. The company requires accurate, near real-time threat data and executes thousands of threat classifications per second across massive data sets. The cybersecurity company couldn't scale its threat classification services with existing relational and recognized it needed an entirely new back-end to power its classification services to keep up with the ever-expanding cyber threat landscape. With the power of TigerGraph, the cybersecurity company can now categorize new threats as they emerge and protect customers using AI-based intelligence in the cloud. Through its collaboration with TigerGraph, the company is leveraging next-generation database technology to maintain its position as a market leader in the space.

Summary

A TigerGraph graph database running on Dell PowerEdge servers with AMD EPYC processors is the optimal solution for your data-centric enterprise and provides performance, efficiency, and scale that significantly ex-

ceed anything available elsewhere in the public cloud. TigerGraph is the only scalable graph database for the enterprise. TigerGraph's proven technology connects data silos for deeper, wider, and operational analytics scale.

- Eight out of the top ten global banks use TigerGraph for real-time fraud detection.
- Over 50 million patients receive care path recommendations to assist them on their wellness journey.
- More than 300 million consumers receive personalized offers with recommendation engines powered by TigerGraph.

TigerGraph supports applications such as fraud detection, customer 360, MDM, IoT, AI, and machine learning. We invite you to reach out to us for a personalized benefit realization session that will showcase how TigerGraph, Dell, and AMD allow enterprises to monetize their data by helping enterprises understand their customers, markets, products, and operations.

Find out more at www.tigergraph.com

Dell Technologies:

Dell Technologies (NYSE:DELL) helps organizations and individuals build their digital future and transform how they work, live and play. The company provides customers with the industry's broadest and most innovative technology and services portfolio for the data era.

TigerGraph:

TigerGraph is a platform for advanced analytics and machine learning on connected data. Based on the industry's first and only distributed native graph database, TigerGraph's proven technology supports advanced analytics and machine learning applications such as fraud detection, anti-money laundering (AML), entity resolution, customer 360, recommendations, knowledge graph, cybersecurity, supply chain, IoT, and network analysis.

AMD:

For more than 50 years AMD has driven innovation in high-performance computing, graphics and visualization technologies. Billions of people, leading Fortune 500 businesses and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work and play. AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit amd.com/epyc. AMD, the AMD Arrow logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc.