



TigerGraph

The Only Scalable Graph
Database for The Enterprise

ENTERPRISE DATASHEET





TigerGraph Enterprise Edition

A graph database is designed to facilitate analysis of relationships in data. A graph database stores data as entities and the relationships between those entities. It is composed of two things: vertices and edges. Vertices represent entities such as a person, product, location, payment, order and so on; edges represent the relationship between these entities, for example, this person initiated this payment to purchase this product with this order.

Graph analytics explores these connections in data and reveals insights about the connected data. These capabilities enable applications such as customer 360, cyber threat mitigation, digital twins, entity resolution, fraud detection, supply chain optimization, and much more.

TigerGraph is the only scalable graph database for the enterprise. TigerGraph's innovative architecture allows siloed data sets to be connected for deeper and wider analysis at scale. Additionally, TigerGraph supports real-time in-place updates for operational analytics use cases.

- Four top tier banks use TigerGraph to improve fraud detection rates by 20% or more.
- Over 300 million consumers receive personalized offers with recommendation engines powered by TigerGraph.
- More than 50 million patients receive care path recommendations to assist them on their wellness journey.
- One billion people depend on the energy infrastructure optimized by TigerGraph to reduce power outages.

The Only Scalable Graph Database for the Enterprise



TigerGraph Benefits

No Limits on Scalability—TigerGraph is designed to scale to Petabytes. Customers can future-proof their investment even if the use case needs grow exponentially.

Richer Data Context for Applications—with TigerGraph, users can perform wider and deeper analysis of data and uncover hard-to-find patterns.

Flexibility to Run AI Workloads—TigerGraph provides extensive in-database machine learning capabilities. If the users prefer, TigerGraph can help with feature extraction and expedite ML training in other downstream systems.

Maximum Price Performance—TigerGraph uses MPP architecture to uniformly utilize resources based on workload needs to deliver ROI.



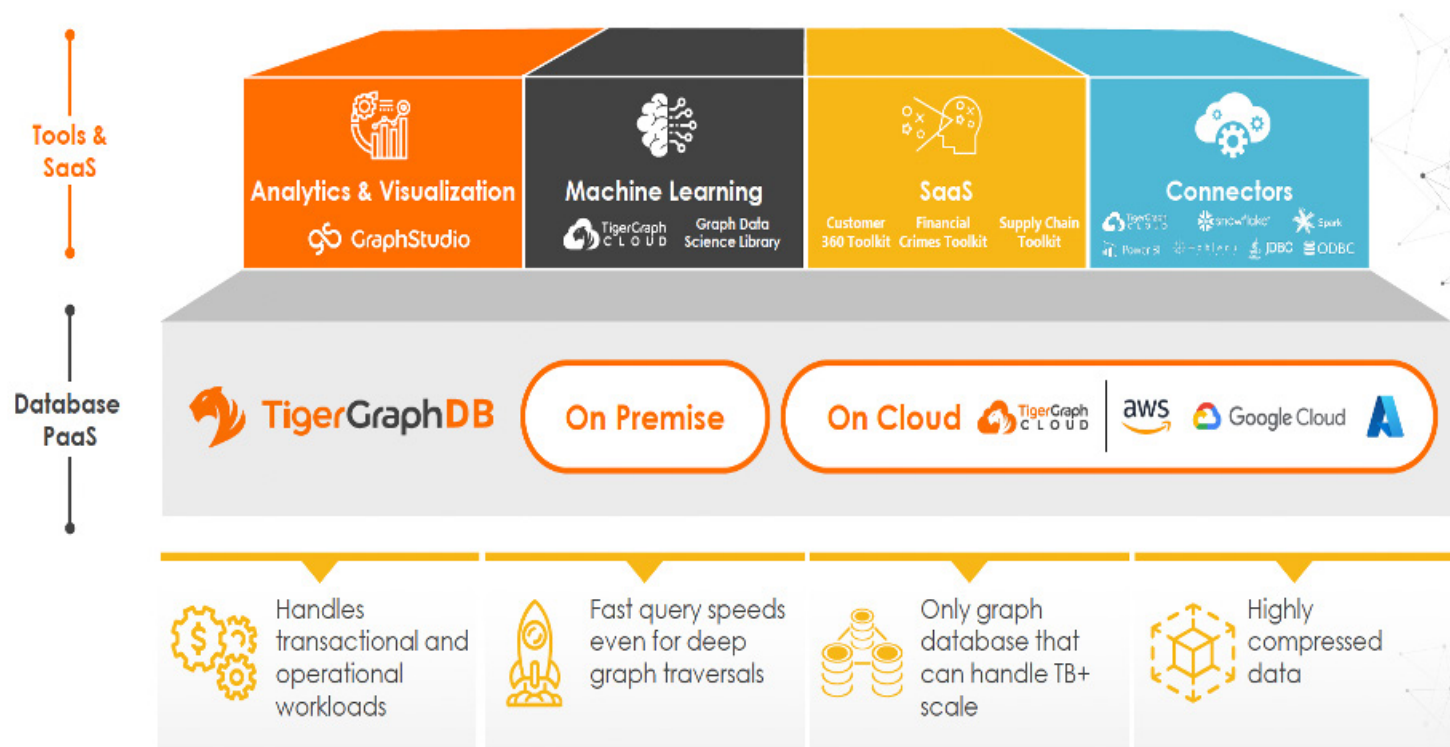
The TigerGraph Advantage

TigerGraph is the only scalable graph database for the enterprise. It is characterized by an ability to handle transactional and analytical workloads and execute queries quickly (even those queries requiring multiple hops across graphs consisting of billions of vertices).



"TigerGraph is an absolutely phenomenal product on which we essentially built everything."

Edward Sverdlin, Vice President,
Advanced Technology Collaborative,
UnitedHealth Group



At its core TigerGraph provides Analytics and Visualization, Machine Learning, Software-as-a-Service, and Connectors capabilities.



The TigerGraph Advantage

TigerGraph provides these advantages:



Speed

Massively parallel processing provides sub-second response for queries with tens of millions of entities/relationships.



Scale-Out

TigerGraph's distributed database scales out with your growing needs (and stays fast, of course). Trillion-edge graphs are running real-time analytics in production.



Real-time Deep-Link Querying

Gain deeper insights through queries which can traverse 10 or more hops and perform complex analytics.



Mixed Workloads

TigerGraph is designed for HTAP (Transactions + Analytics) with strong transactional guarantees to run both operational and analytical workloads on the same cluster.



Total Cost of Ownership

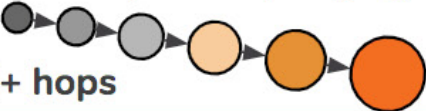



Scale up from a proof of concept into production with the lowest total cost of ownership (due to higher performance, scalability and better compression than other graph databases).



TigerGraph has this highly scalable system that doesn't block reads when transactions are happening and is ideal for extremely high volume transaction processing. When people say, do you support ACID, it's usually yes and they do it on a single server. It's much, much more difficult to do ACID transactions for a distributed cluster and TigerGraph has nailed that one pretty well"

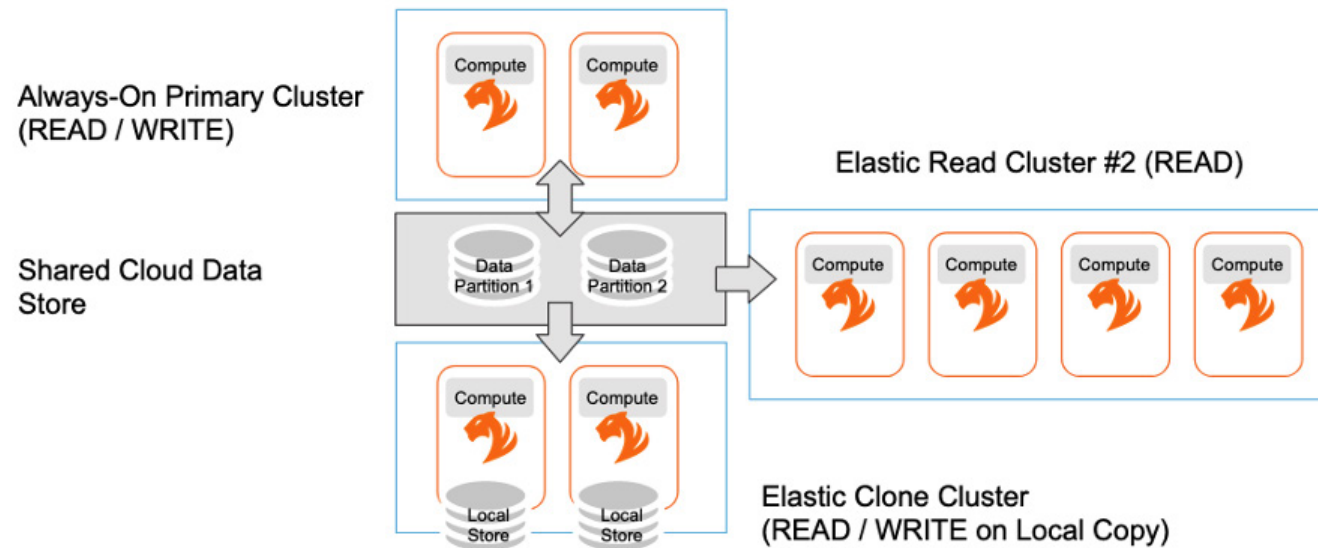
Dan McCreary, Distinguished Engineer, Optum

The key TigerGraph features are summarized in the following table with unique aspects of the design and corresponding business benefits:

Feature	Design Difference	Benefit
Real-Time Deep-Link Querying  5 to 10+ hops	<ul style="list-style-type: none"> • Native Graph design • C++ engine for high performance • Storage Architecture 	<ul style="list-style-type: none"> • Uncovers hard-to-find patterns • Operational, real-time • HTAP: Transactions+Analytics
Handling Massive Scale 	<ul style="list-style-type: none"> • Distributed DB architecture • Massively parallel processing • Compressed storage reduces footprint and messaging 	<ul style="list-style-type: none"> • Integrates all your data • Automatic partitioning • Elastic scaling of resource usage
In-Database Analytics & Machine Learning 	<ul style="list-style-type: none"> • GSQL: High-level yet Turing-complete language • User-extensible graph algorithm library, runs in-DB • ACID (OLTP) & Accumulators (OLAP) 	<ul style="list-style-type: none"> • Avoids transferring data • Richer graph context • Graph-based feature extraction for supervised machine learning • In-DB machine learning training
	<ul style="list-style-type: none"> • No-code migration from RDBMS • No-code Visual Query Builder 	<ul style="list-style-type: none"> • Democratize self-service analytics to derive new-insights from legacy/external data stores

WORKLOAD TYPES	CUSTOMER EXAMPLES
TRANSACTIONAL	<ul style="list-style-type: none"> • China Construction Bank - CCB uses TigerGraph in-transaction during each payment to check for fraud detection and reject payments that are above the risk threshold set.
ANALYTICAL	<ul style="list-style-type: none"> • Jaguar Land Rover - One of the largest supply chain deployment for graph analytics, reducing the cycle time for planning from 3 weeks to 45 minutes. • HBO - one of the largest deployment of entity resolution and recommendation engine with a graph database in media.
PREDICTIVE ANALYTICS	<ul style="list-style-type: none"> • FinTell - Predicting risk-based on mobile device and payment transaction data. • Multiple Major Banks - Predicting likely fraudster rings based on Louvain Community Detection algorithm. • Webroot - Predicting websites that are likely to contain malware.
OPERATIONAL	<ul style="list-style-type: none"> • US Xpress - Route Optimization for shipping with TigerGraph: The vehicles are routed and dispatched based on the output from TigerGraph for operations every day. • Kickdynamic - Email recommendations inserted into each email sent out by over 200 brands across Europe using TigerGraph (millions per year). • Citrix - TigerGraph embedded into the Citrix platform for employee productivity, analyzes employee behavior for content recommendations every day.
MIXED TRANSACTIONS, DEEP LINK ANALYTICS, OPERATIONAL, PREDICTIVE	<ul style="list-style-type: none"> • Intuit - Every Intuit QuickBooks transaction passes through TigerGraph for real-time fraud check, 50+ fraud investigators use GraphStudio every day for running analytical queries for fraud detection, Intuit uses a knowledge graph powered by TigerGraph for customer 360. • UnitedHealth Group - Largest healthcare graph in the world, used by over 23,000 users every day operations (call center & other employees) in production with customer 360 and similarity matching with transactions + deep analytical queries driving productivity improvement.
TRANSLYTICAL (TRANSACTIONAL + ANALYTICAL)	<ul style="list-style-type: none"> • Pagantis - Real-time credit risk assessment, delivering 500 ms response for each credit request for eCommerce purchases with deep-link analytical queries for each transaction supporting core operations of Pagantis. • China Mobile - Over 300 million calls are run through TigerGraph in real-time every day, with TigerGraph rating the risk for each call for fraud. Call graph is updated in real-time each transaction to re-rate the fraud risk. Deep link analytics based on 5+ connections executed in real-time for each call. • Xandr - World's largest identity graph in production in the media industry.
SEARCH	<ul style="list-style-type: none"> • Ippen Digital - Knowledge graph powering publishing industry in Germany to power content search and recommendations. • OpenCorporates - Knowledge Graph powering search and visualization of corporate entities and their relationships at the world's largest open database of corporate information. • CAS - Knowledge Graph powering search through the abstracts as well as full length research papers and publications for CAS.

Compute on Demand



About TigerGraph

TigerGraph is the only scalable graph database for the enterprise. TigerGraph's proven technology connects data silos for deeper, wider and operational analytics at scale. Four out of the top five global banks use TigerGraph for real-time fraud detection. Over 50 million patients receive care path recommendations to assist them on their wellness journey. 300 million consumers receive personalized offers with recommendation engines powered by TigerGraph. The energy infrastructure for 1 billion people is optimized by TigerGraph for reducing power outages. TigerGraph's proven technology supports applications such as fraud detection, customer 360, MDM, IoT, AI, and machine learning.

For more information visit www.tigergraph.com and follow us at: [Facebook](#) [Twitter](#) [LinkedIn](#)

Contact us at sales@tigergraph.com

Latest Functionality

Platform:

- **Elasticity**—Added support for provisioning Elastic Read-only (ER) clusters with different partitions than the primary cluster.
Note: This feature is currently in the preview stage only available for enterprise TigerGraph Cloud customers. If you are a paid TigerGraph Cloud enterprise customer and want to set up an ER cluster for your environment, please open a support ticket.
- **Manageability**—Introduced TigerGraph Kubernetes Operator. TigerGraph Kubernetes Operator allows you to automate operations such as the creation, status checking and deletion of TigerGraph clusters.
- **Security**—Sensitive data such as user credentials in log and configuration files are now encrypted.
- **Ecosystem integration**—Added the ability to stream data in formats such as AVRO and JSON from external Kafka clusters to Data Streaming Connector.
- **Performance**—Improved data loading speed and reliability. Improved the speed and reliability of database catalog operations such as vertex and edge definition, schema changes, and query installation.
- **Query Language Enhancement**—Reducer function add() now accumulates values when loading to a MAP attribute.
- **GraphStudio Enhancement**—Added the ability to one-click install built-in Graph Data Science Library algorithms. Added the ability to load data from Google Cloud Storage through GraphStudio UI.