

Graph OLAP database built for hyperfast analytics at scale



AT A GLANCE

High performance graph query execution can scale to billions or even trillions of triples for analyzing more data with deeper insights

Supports both labeled property graphs (using W3C RDF and SPARQL* proposed standard) and coming soon OpenCypher*

Fast parallel data loads help you get data analyzed fast without taking the database offline

All-in-one analytics with graph algorithms, graph views, named queries, aggregates, built-in data science functions, data warehouse-style BI and reporting functions

Priced for affordability and scale

Easy to download and get started

AnzoGraph® is a native, Massively Parallel Processing (MPP) distributed Graph OLAP (GOLAP) database, providing hyperfast advanced analytics at big data scale.

Graph Database and Scale

Most solutions in the graph database market are geared toward Online Transactional Processing (Graph OLTP) and can't scale when you need to perform deep analysis of billions or even trillions of triples in an acceptable timeframe. Such Graph OLTP systems add database servers to replicate data for improved access, address more users or to provide geographic coverage.

The Graph OLAP Database

Graph OLAP databases like AnzoGraph add database servers to enhance load & query speeds and to handle more data. Developers can use AnzoGraph to build and execute data warehouse-style BI analytics, graph algorithms, data science functions and inferencing, all in one award-winning database.

AnzoGraph offers this combined analytic functionality in a single enterprise-grade Graph OLAP database, to enable unprecedented new levels of data-driven insight to help drive new business opportunities, minimize costs and increase competitiveness. Graph OLAP databases complement various OLTP databases and can be used with OLTP databases or on a stand-alone basis.

AnzoGraph is the same enterprise graph database embedded within Anzo®, in production use by the leading global organizations to solve their most complex business challenges. Now AnzoGraph is available on a standalone basis for use on bare metal, on VMs and containers and in the cloud.



ANZOGRAPH USE CASES

- Execute graph analytics, data science analytics and interactive business intelligent and reporting in one system
- Build out inferencing use cases
- Build and use knowledge graphs for metadata management, data discovery and analytics
- Solve analytics problems such as buyer intent analysis, key opinion leader analysis, predictive analytics, market basket analysis, recommendation engines, and more!
- Natively use industry standard data models such as FIBO in financial services, HL7/FHIR in healthcare, and CDISC in pharma
- Tackle machine learning insight in both training the algorithms and deployment of them



DEPLOY ANALYTICS ANYWHERE

*Standards-based
multi-graph and federating
services support offers
greater access for
your users*

*Ensure data integrity with
ACID transaction support*

*Use memory more
efficiently with
compressed data store*

*Supports standard UI and
visualization tools (eg:
Zeppelin Notebooks) and
third party visualization
applications*

*Deploy anywhere, both
behind the firewall and in
the cloud (eg: AWS, GCP,
Azure) with ondemand
scale-up and scaledown*

Features and Benefits

Ease and value of graph performance comparable to a best-in-class database

Scalable MPP Platform

High performance query execution can scale to billions or even trillions of triples. Use AnzoGraph with 1 node or scale up to multiple nodes for faster and deeper insights.

Hyperfast Load and Query

AnzoGraph supports parallel data loads so you can get data analyzed fast. There's no need to take the database offline while loading. Shared-nothing architecture provides for superior horizontal scalability and super-fast queries. Aggregates and functions that consider all the data are hyperfast when compared with other databases.

Conforms to W3C Standards PLUS

AnzoGraph uses standards from W3C regarding RDF data formats and the SPARQL language. In addition, extensions support labeled property graphs (LPG) under the new proposed standards (RDF* and SPARQL*)

Data Warehouse-style BI Analytics

Graph views, named queries, analytic/data mining /reporting functions help support the graph data warehouse

Powerful Graph Algorithms

Graph algorithms including pagerank, shortest path, all paths and many others help you solve your data analytics challenges at scale with ease

Easy Integration

REST API and full SPARQL 1.1 compliance to connect AnzoGraph to your applications and secure it via SSL.

Coming soon: OpenCypher

OpenCypher support is planned to be released soon, including Bolt API.

Free Trial!

Visit AnzoGraph.com for a free 60-day full feature trial or try on AWS Marketplace.




ANALYTICS SUPPORTED

- **Graph Algorithms**
Page Rank, Connected Components, Triangle Enumeration, Triangle Counting, Clustering Coefficient, Label Propagation, Shortest Path, and more!
- **Data Warehouse-style BI Analytics**
Graph Views, Named Queries, Window Aggregates, Advanced Grouping Sets, 100+ Excel type of functions such as MEAN and AVG, Conditional Expressions, and more!
- **Inferencing**
RDFS+
- **Built-in Data Science Analytics**
Normal Distribution, Bernoulli Distribution, Geometric Mean Metric and more!

COST EFFECTIVE DEPLOYMENT

- **Cloud**
Cloud vendor agnostic (e.g. Amazon AWS, Google Cloud, Microsoft Azure)
 - **On-premises**
Linux bare metal and virtual machines. Use commodity hardware
 - **Other**
Docker, Kubernetes
- Persist data in NFS, HDFS or S3 for flexibility and low-cost storage.

Fastest and Most Scalable Graph Database Download Benchmarks at AnzoGraph.com



- LUBM – 110 times faster than any previous results
- TPC-H (GHIB) – 217 times faster than a leading OLTP solution for load and analysis
- Graph 500 - Load 41.6 million vertices and 1.47 billion edges in 4 ½ Minutes