

# Aerospike Graph

## SOLUTION BRIEF

### A scalable, high-performance graph database that's developer-ready

Graph databases are an important method of querying and processing increasingly connected data and the myriad underlying relationships. As graph datasets grow in volume, existing graph databases have struggled to keep up with increasing demands of real-time property graph use cases such as identity graphs in Ad Tech, Customer 360, and real-time fraud prevention. Aerospike Graph was created to address this new generation of graph workloads.

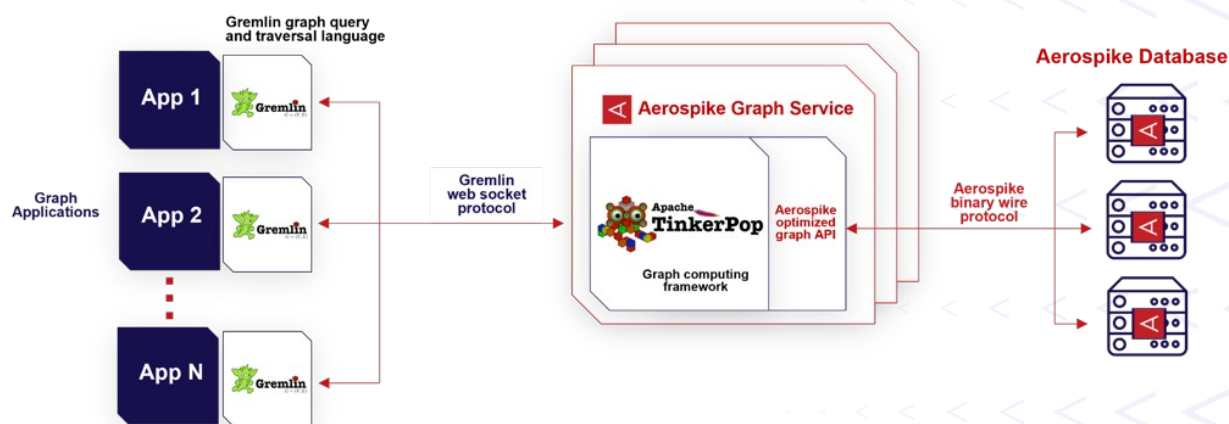


Figure 1. Aerospike Graph architected for scalability and cost efficiency.

Aerospike Graph is a highly scalable, low-latency property graph database built on Aerospike's proven real-time data platform. Aerospike Graph combines the enterprise capabilities of the Aerospike Database - the most scalable real-time NoSQL database - with the property graph data model via the Apache Tinkerpop graph compute engine. Developers will enjoy native support for the Gremlin query language, which enables them to write powerful business processes directly.

## Key Features

Aerospike Graph is an engineered graph database that enables low-latency, multihop graph queries at a very large scale based on widely adopted graph technologies and practices. Key features include:

- **Unlimited scale** of a native graph data model implemented for low-latency access beyond terabytes of data
- **Real-time performance** for wider and deeper transactional (OLTP) graph queries with extreme throughput (>100K QPS) with sub-5ms latency for 2-3 hop queries
- **Independently scaled compute and storage** for accurate and efficient resource utilization and best TCO
- **Power graph query language** through support for the Gremlin query language enables developers to easily and efficiently query and process interconnected graph data
- **Based on a popular open-source framework** implemented as a scale-out graph computing service based on the long-established Apache Tinkerpop framework
- **Bulk data loading** quickly and efficiently loads large volumes of graph data using Apache Spark
- **Leverages Aerospike performance and scalability features** such as secondary indexes and expressions
- **Enterprise-grade** availability, security, and observability

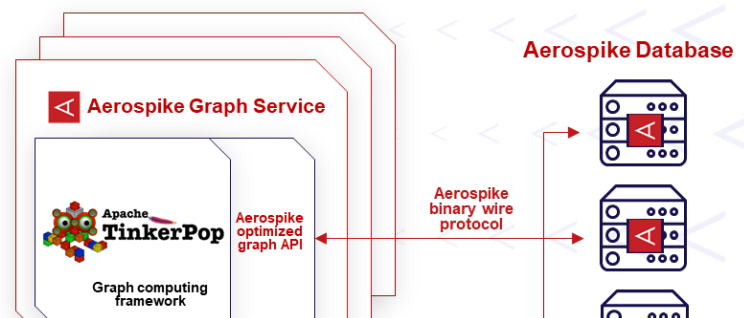
## Key Benefits

- **Removes graph database bottlenecks** through a modern scale-out graph processing architecture that eliminates traditional performance and scalability challenges
- **Operational efficiency** through single vendor support for multiple data models (graph, document, key-value) on a single data platform
- **New operational graph use cases are now possible** through dramatically lower graph query latencies, even at petabyte scale
- **Moves graph model beyond static analytics** via a high transaction throughput data platform to power operational solutions
- **Taps into an established and growing developer base** that uses the popular Gremlin query language

## Aerospike Graph Overview

### Aerospike Graph Service

When Aerospike Graph is deployed, it is run via the Aerospike Graph Service, which is a highly optimized implementation of [Apache TinkerPop](#), which is an open-source property graph computing framework for online transaction processing (OLTP) and online analytical processing (OLAP) graph queries. It's a mature codebase that has been tested and deployed in production since 2009. The Aerospike Graph Service also provides the graph API implementation that provides deep integration between TinkerPop and Aerospike.



For graph operations to work effectively at Aerospike scale, we have engineered the Aerospike Graph Service to exploit many of the unique performance and scalability features of the database. Aerospike Database acts as the scalable persistence layer for Aerospike Graph with features such as secondary indexes and Aerospike Expressions providing the horsepower to deliver real-time results at scale.

Aerospike has created a highly optimized data model to represent graph elements, such as vertices and edges, that map to the native Aerospike data model using records, bins, and other Aerospike features. Aerospike has also implemented numerous traversal strategies and step optimizations by taking advantage of Aerospike Database scalability and performance features like secondary indexes and expressions.

## Gremlin graph query language

Aerospike Graph supports the [Gremlin](#) interface - a popular open-source graph query language - out of the box such that existing Gremlin applications and Gremlin queries can easily be run in an Aerospike Graph environment. From the Apache Tinkerpop website:

“Gremlin is the graph traversal language of Apache TinkerPop. Gremlin is a functional, data-flow language that enables users to succinctly express complex traversals on (or queries of) their application’s property graph.”

To learn more about Gremlin, see the online book [Practical Gremlin](#).

## Operational graph use cases

### Identity Graphs for Ad Tech

The Ad Tech industry is undergoing a major shift in response to changing market conditions and technical challenges resulting from the impending “cookie apocalypse.” Many Ad Tech firms are looking to graph databases to meet these new realities and Aerospike Graph will play a part.

### Fraud Detection

Cyber criminals are getting more and more inventive in their fraud strategies, creating a constantly evolving attack surface for cyber security teams to deal with. Many are looking to real-time OLTP graph databases like Aerospike Graph to better track the relationship between the ocean of relevant data points that effect security.

### Customer 360

As retailers and e-commerce companies become more sophisticated about reaching and influencing their customers, the supporting technology must evolve. New Customer 360 solutions will look to scalable, real-time graph databases like Aerospike Graph to respond instantly to customer behavior

### AI/ML

The popularity of AI and machine learning has skyrocketed as Generative AI has generated headlines. Companies racing to adopt more AI/ML projects will discover that scalable, real-time graph databases like Aerospike Graph may play a major part in supplying the right data to their AI solutions.