



## **TigerGraph Unveils Next Generation Hybrid Search to its Graph Database to Power AI at Scale; Also Introduces a Game-Changing Community Edition**

*The multi-modal combination delivers deeper learning and insights, Graph RAG for smarter AI, and super fast response times*

**Redwood City, Calif. – March 4, 2025** – TigerGraph, the enterprise AI infrastructure and graph database leader, today announced its next generation graph and vector hybrid search delivering the industry's most advanced solution for detecting data anomalies through sophisticated pattern analysis, identifying critical deviations from expected norms, and providing actionable recommendations. The integration of graph and vector search on a single high-performance, scalable platform offers businesses a comprehensive solution for developing significantly more accurate AI systems for fraud and anti-money laundering detection, real-time personalized recommendations, and image and multimedia matches among others. Simultaneously, TigerGraph is releasing a Community Edition of its graph database that offers significant compute power and storage capacity.

TigerGraph is revolutionizing vector search with unmatched speed, accuracy, and scalability - essential components for advanced AI and ML systems. By leveraging graphs to represent proprietary local knowledge and real-time data, including their interrelationships, graph-enhanced AI and GraphRAG deliver superior personalization and explainability. This multi-modal approach simplifies the design and operation of complex AI use cases, dramatically reducing infrastructure complexity and code requirements while providing enterprise-grade security, access controls, and reliability.

TigerGraph vector search benefits include:

- **5.2x faster vector searches with 23% higher recall than competitors** to rapidly uncover the most similar items while using 22.4x fewer resources and reducing operational costs.
- **6x faster indexing** – blazing fast loading and automatic incremental updating of search indexes, ensuring accurate searches even with large datasets and rapid changes.
- **Advanced hybrid search of structured and unstructured data** – enhances discoverability and contextually rich understanding for ML and AI systems, significantly improving their analytical capabilities.

- **Rich relationship modeling** – delivers support for complex relationships between entities and creates sophisticated knowledge graphs.
- **Integrated query language** – express hybrid graph+vector queries in GSQL, harmonically achieve structured and unstructured query composition. Our Python library also supports vector database operations.

TigerGraph's Community Edition is the most powerful graph database that's free to use, even in production:

- **16 CPUs** of compute power for significantly higher performance.
- **200 GB graph storage and 100 GB of vector storage** to enable AI-driven applications.
- **Extensive AI/ML open-source library**, simplifying the development of graph + vector applications, including GraphRAG.
- **GSQL, OpenCypher, and ISO GQL** for the widest and most powerful query language support.

“We’re continuing to lead the way in delivering the industry’s fastest, most scalable analytics for AI and machine learning users,” said Rajeev Shrivastava, CEO of TigerGraph. “The engineer in me is excited to put these solutions directly into the hands of developers who are building mission critical, AI dependent products that improve their customers’ lives.”

Start using [Hybrid Search with Community Edition](#) today. Read TigerGraph’s blog and [SIGMOD 2025 paper](#) for a technical deep dive.

On the heels of the release of TigerGraph [Savanna](#), the most innovative cloud native graph database platform for supercharging AI systems, TigerGraph continues to lead the market as the enterprise-grade graph database for customers that need to quickly discover relationships, analyze complex patterns, and make mission critical decisions.

## About TigerGraph

TigerGraph, the enterprise AI infrastructure and graph database leader, delivers massively parallel storage and computation that scales independently and without size limits, to meet the changing workloads and growing data volumes required for crucial business needs and AI adoption within companies. By providing visibility into the multidimensional data connections and relationships, TigerGraph has become a trusted partner to leading companies including JPMC, Intuit, United Healthcare, and Unilever successfully solving fraud detection, entity resolution, customer 360, supply chain management, and many other problems. Headquartered in Silicon Valley, California and with offices around the world TigerGraph is backed by Tiger Global Management, Softbank, Susquehanna International Group (SIG), Oceanpine Capital, Celesta Capital, Nvidia, Blackopal Ventures, and Qiming Venture Partners.

**Contact:** [paige.leidig@tigergraph.com](mailto:paige.leidig@tigergraph.com)