

AllegroGraph v7.2 – Now Available with GNNs, Virtual Graphs, Spark, and Kafka

AllegroGraph 7.2, provides organizations with essential Data Fabric tools, including Graph Neural Networks, Graph Virtualization, Apache Spark graph analytics, and streaming graph pipelines. These new capabilities exemplify AllegroGraph's leadership in empowering data analytics professionals to derive business value out of Knowledge Graphs.

Graph Neural Networks

With AllegroGraph 7.2, users can create Graph Neural Networks (GNNs) and take advantage of a mature AI approach for Knowledge Graph enrichment via text processing for news classification, question and answer, search result organization, event prediction, and more. GNNs created in AllegroGraph enhance neural network methods by processing the graph data through rounds of message passing, as such, the nodes know more about their own features as well as neighbor nodes. This creates an even more accurate representation of the entire graph network. AllegroGraph GNNs advance text classification and relationship extraction for enhancing enterprise-wide Data Fabrics.

Graph Virtualization

AllegroGraph 7.2 allows users to easily virtualize data as part of their AllegroGraph Knowledge Graph solution. When graphs are virtual, the data remains in the source system and is easily linked and queried with other data stored directly in AllegroGraph.

Any data source with a supported JDBC driver can be integrated

into an AllegroGraph Knowledge Graph, including Databases (i.e. Apache Cassandra, AWS Athena, Microsoft SQL Server, MongoDB, MySQL, Oracle Database); BI Tools (i.e. IBM Cognos, Microsoft PowerBI, RapidMiner, Tableau); CRM Systems (i.e. Dynamics CRM, Netsuite, Salesforce, SugarCRM); Cloud Services (i.e. Active Directory, AWS Management, Facebook, Marketo, Microsoft Teams, SAP, ServiceNow) and Shared Data Files (i.e. Box, Gmail, Google Drive, Office365).

Streaming Graph Pipelines using Kafka

Enterprises that need real-time experiences are starting to adopt streaming pipelines to provide insights that adapt to new data in real-time rather than processing data in batches. AllegroGraph is often used as an Entity Event Knowledge Graph platform in diverse settings such as call centers, hospitals, insurance companies, aviation organizations and financial firms.

AllegroGraph 7.2 can be used seamlessly with Apache Kafka, an open-source distributed event streaming platform for high-performance data pipelines, streaming analytics, data integration and mission-critical applications. By coupling AllegroGraph with Apache Kafka, users can create a real-time decision engine that produces real-time event streams based on computations that trigger specific actions. AllegroGraph accepts incoming events, executes instant queries and analytics on the new data and then stores events and results.

Graph Analytics with Apache Spark

AllegroGraph 7.2 enables users to export data out of the Knowledge Graph and then perform graph analytics with Apache Spark, one of the most popular platforms for large-scale data processing. Users immediately gain machine learning and SQL database solutions as well as GraphX and GraphFrames, two frameworks for running graph compute operations on data.

A key benefit of using Apache Spark for graph analytics within

AllegroGraph is that it is built on top of Hadoop MapReduce and extends the MapReduce model to efficiently use more types of computations. Users can access interfaces (including interactive shells) for programming entire clusters with implicit data parallelism and fault-tolerance.

Availability of AllegroGraph 7.2

AllegroGraph 7.2 is immediately available directly from Franz Inc. For more information, visit the AllegroGraph Quick Start page for cloud and download options.

Examples

Visit our Github AllegroGraph Examples page.