

ACCELERATED GRAPH MACHINE LEARNING



OVERVIEW

Data scientists and data analysts are constantly searching for correlations in data to solve real-world business problems at enterprise scale. Relational Databases (tables) and Key Value Store (unstructured) are not architected to efficiently include relationships – making it difficult to understand behaviors and unlock predictions; and creating need for disparate workflows for data scientists searching for answers in siloed data. In contrast, graph analytics is architected to exploit data relationships in the most seamless and efficient fashion.

AMD and TigerGraph have pioneered hardware-accelerated graph analytics via massively parallelized Alveo™ hardware accelerators to deliver real-time data insights—now powered by machine learning. Scaling in the cloud can easily be achieved in the AWS cloud via EC2 F1 instances in three different sizes that include either one, two, or eight accelerator cards per instance powering multi-instance clusters.

HIGHLIGHTS

Hardware Accelerated, Real-Time Analytics for Faster Insights

- > Accelerates a breadth of frequently used algorithms
- > Delivers orders of magnitude greater performance vs. CPUs

Seamless Integration of TigerGraph and Custom ML Workflows

- > Support for GSQL, Python, and C++
- > Jupyter Notebook examples for faster ML development

Rapid Scale-Out and Deployment in the Cloud or On-Premise

- > Deploy via AWS EC2 F1 instance for quick time to market
- > Supported by Dell, HPE, and Lenovo servers for on-premise deployment

UP TO

45X

RUNTIME IMPROVEMENT
VS. CPU IMPLEMENTATION¹

1. TigerGraph accelerating Louvain Modularity (Fraud Detection) by Alveo U55C vs CPU only with AMD EPYC 128C/256T CPU



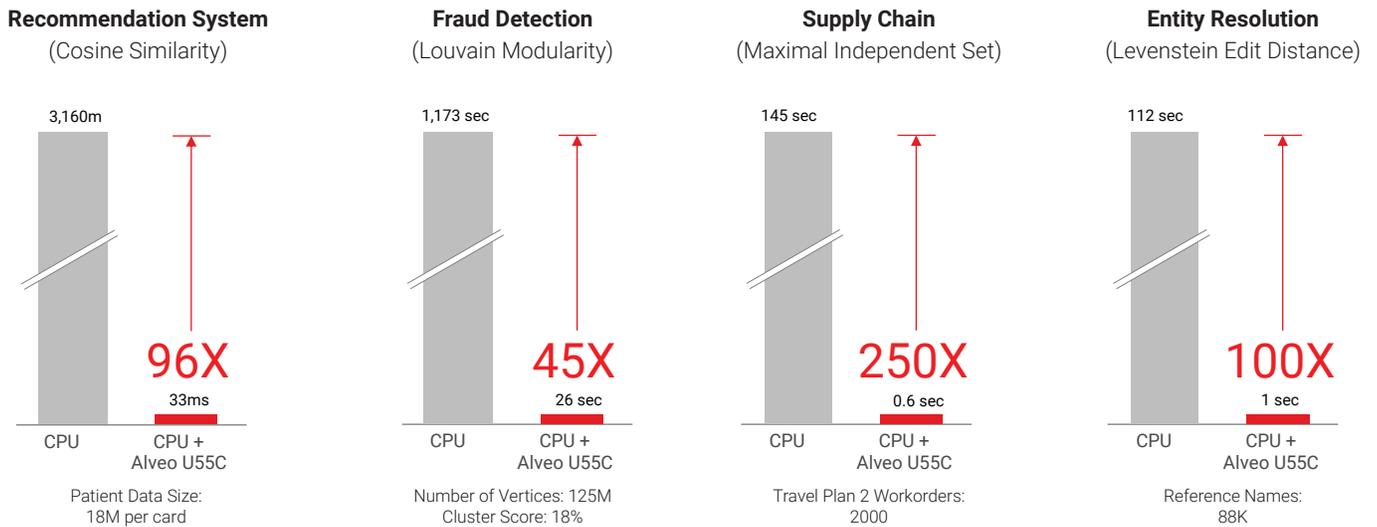
END-APPLICATIONS

- > Fraud Detection
- > Recommendation Systems
- > Supply Chain Analysis
- > Public Sector Data Analysis
- > Healthcare analytics
- > Risk Assessment
- > Product and Service Marketing
- > Customer 360

HARDWARE ACCELERATED, REAL-TIME ANALYTICS FOR FASTER INSIGHTS

Graph analytics is architected to exploit data relationships in the most seamless and efficient fashion. Connections are important as the data itself and finding hidden correlation within the graph is being explored with Machine Learning. AMD’s Accelerated Graph Machine Learning (AGML) library is architected to scale out with multiple accelerator cards running enterprise graph database workload on-premise as well as in the cloud for graphs with billions of vertices and edges.

- > Recommendation System: Real-time personalized recommendations are produced 96x faster for Cosine Similarity compared to a CPU. Millions of patient records are analyzed to deliver clinical outcome predictions in milliseconds.
- > Fraud Detection: Real-time community detection is executed 45X faster for Louvain Modularity compared to a CPU implementation. Communities of fraud rings and money launderers are uncovered in real-time across hundreds of accounts.
- > Supply Chain Management: Travel plans and transportation flows are analyzed 250x faster for Maximal Independent Set (MIS) compared to a CPU implementation—delivering insights in seconds vs. minutes.
- > Entity Resolution: Identification of unique entities are made 100X faster for Levenstein Edit Distance compared to a CPU running detection of multiple entries that are the same through fuzzy matching and quickly deduplicate the dataset.



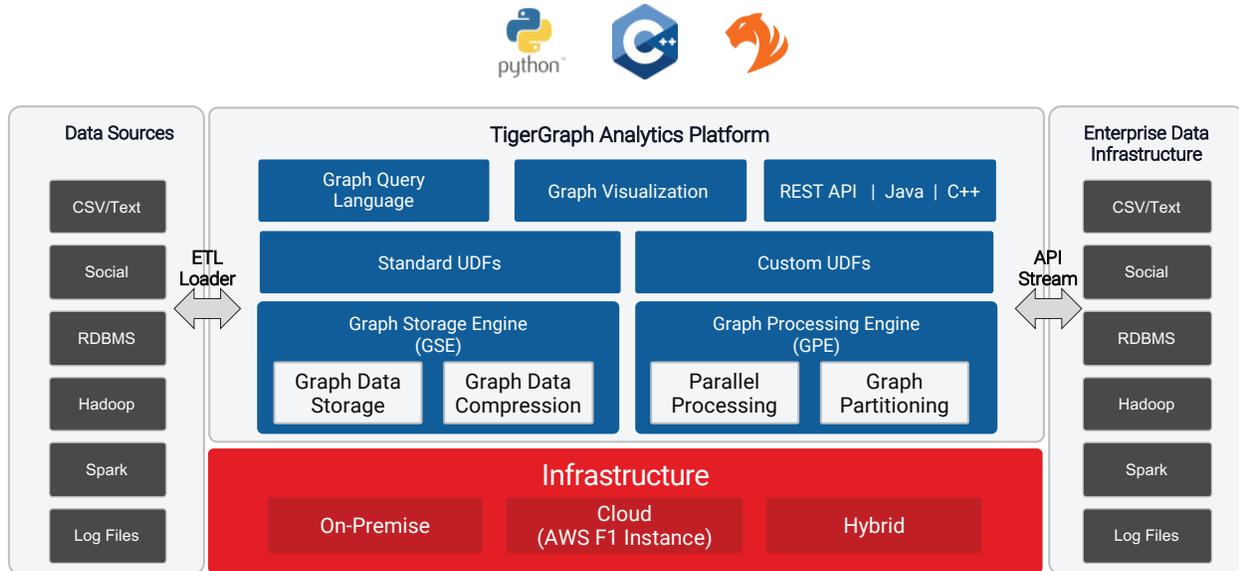
Relative Runtime and Memory Utilization for Cloud and On-Premise

| Use Cases | Algorithms | On-premise (Alveo U55C) | | Cloud (AWS F1) | |
|-----------------------|-------------------------|-------------------------|----------|----------------|----------|
| | | Runtime | Memory | Runtime | Memory |
| Fraud Detection | Louvain Modularity | 45x | 60% less | 10x | 70% less |
| Recommendation Engine | Cosine Similarity | 96x | Same | 19x | Same |
| Supply Chain | Maximal Independent Set | 250x | Same | 70x | Same |
| Entity Resolution | Levenstein Distance | 100x | Same | 45x | Same |

CPU Info: single node with 128 cores, 256 CPUs and 512 GB memory

SEAMLESS INTEGRATION OF TIGERGRAPH AND ML WORKFLOWS

Data scientist do not need to change their existing graph workflows to turbo charge insights. AMD's Accelerated Graph Machine Learning (AGML) library seamless integrates into TigerGraph's GSQL query language.



RAPID SCALE-OUT IN THE CLOUD OR ON-PREMISE

Scaling in the cloud can easily be achieved with AWS offers Amazon EC2 F1 instances in three different sizes that include either one, two, or eight accelerator cards per instance powering multi-instance clusters. For on-premise deployments, AGML is supported by Dell, HPE, and Lenovo servers with Alveo U50 and U55C accelerator cards.



TAKE THE NEXT STEP

- > Learn about Accelerated Graph Machine Learning at <https://xilinx.github.io/graphanalytics/index.html>
- > More details on cloud support with Amazon EC2 F1 Instances at <https://aws.amazon.com/ec2/instance-types/f1/>
- > Reference Architecture published by AMD, TigerGraph, and HPE at <https://www.hpe.com/psnow/doc/a50004506enw>
- > Find out more about supported Alveo accelerator cards at <https://www.xilinx.com/u55c>