Hewlett Packard Enterprise

Reference Architecture

HPE Reference Architecture for accelerated Graph Analytics optimized for Fraud Analytics

Using HPE ProLiant DL385 Gen10 Plus v2 server, TigerGraph, and Xilinx Alveo U50 data center accelerator card

CONTENTS

Executive summary 3
Introduction 3
Solution overview 7
Solution components 8
HPE ProLiant DL385 Gen10 Plus v2 server 8
Xilinx Alveo U50 Data Center Accelerator Card 9
Aruba 8325 switch series 9
Hardware 11
Software 11
Application software 11
Best practices and configuration guidance for the Louvain Fraud Detection solution 14
Getting started 15
User guide 15
Capacity and sizing 16
Workload description 16
Analysis and recommendations 17
Summary 19
Implementing a proof-of-concept 19
Implementing a Pilot 19
Appendix A: Bill of materials 20
Resources and additional links 22

EXECUTIVE SUMMARY

Businesses will lose more than US$200 billion due to online fraud between 2020 and 2024, according to Juniper Research. This stunning loss has been driven by increased fraud attempt sophistication and a rising number of attack vectors. While banks have been fighting back, fraudsters have adjusted their techniques to remain below the radar.

Fortunately, banks have a powerful weapon in the war against fraud: graph analytics. Advanced analytics in graph databases can uncover suspicious patterns of online payment activity in ways that other approaches cannot, helping to stop fraud before it's committed.

Graph analytics provides organizations with an unprecedented ability to uncover relationships between connected data, at scale and in real-time. These insights are being used to increase revenues, lower costs, increase customer satisfaction, and improve employee productivity.

Graph is transforming the data analytics landscape. Gartner forecasts that "By 2025, graph technologies will be used in 80% of data and analytics innovations, up from 10% in 2021, facilitating rapid decision making across Enterprise."

Graph analytics offers several advantages to overcome large and complex data challenges, compared to other analytical solutions for relational databases. TigerGraph is the only native parallel graph database purpose-built for real-time analysis of data at an enterprise scale. It supports transactional as well as analytical workloads, is ACID compliant, and scales up and out with automatic data partitioning. Many of the world's top companies are using TigerGraph for fraud detection, customer360, supply chain optimization, and other applications.

Hewlett Packard Enterprise, Xilinx®, and TigerGraph have partnered to obtain faster, deeper, and wider insights on connected data using HPE ProLiant DL385 Gen10 Plus v2 servers with Xilinx Alveo U50 Data Center PCIe accelerator cards. The joint solution provides at least a 15% improvement in accuracy up to 25%, 20x faster results than CPU-only solutions while reducing the system RAM requirements by 66%.

Document purpose: This Reference Architecture provides comprehensive architectural guidelines and implementation of accelerated graph analytics, optimized for fraud analytics, on HPE ProLiant DL385 Gen10 Plus v2 servers using TigerGraph and Xilinx Alveo U50 Data Center Accelerator Cards. Organizations can readily apply the blueprint tested in this document to implement a graph analytics solution quickly. This document also provides details on how to get started for top use cases such as fraud detection.

Target audience: This document is intended for subject matter experts, domain experts, data engineers, data scientists, architects, pre-sales engineers, service consultants, partner engineers, and line of managers who need to answer complex questions in business.

This Reference Architecture describes solution testing performed in August 2021.

INTRODUCTION

Customers in key vertical markets such as financial services, cyber security, healthcare, retail, and manufacturing want faster, deeper, and wider insights on connected data. Gaining these insights is transformative for these businesses to better understand customers and product features or requirements to reduce fraud or risk, and optimize global supply chains. Critical business questions are easy to ask but hard to answer without scalable graph technology.

In this Reference Architecture, we will focus on fraud detection for financial services, where TigerGraph has been deployed into production at eight of the top ten banks, globally.

Business use cases include:

- Internal Fraud - Entitlements
- AML and Fraud Rings
- Credit Card and Transaction Fraud
- ID falsification and theft
- Cyber - Malware
- Surveillance - IoT and Asset Fraud
- Audit and Compliance
- Claims, Dispute, and Misappropriation

- Law Enforcement

Figure 1 shows the example of an Anti-Money Laundering use case.

FIGURE 1. Example of Anti-Money Laundering use case

In graph database and analytics, relationships are first-class citizens, and datasets and silos are all connected. TigerGraph allows users to look for semantic patterns of relationships and to search far and wide more easily.

TigerGraph's Fraud Detection and Investigation Toolkit helps customers increase accuracy, decrease false positives, increase the speed of automated and human outcomes, adapt to ever-evolving fraud, continuously improve and integrates with and improves current tech stacks.

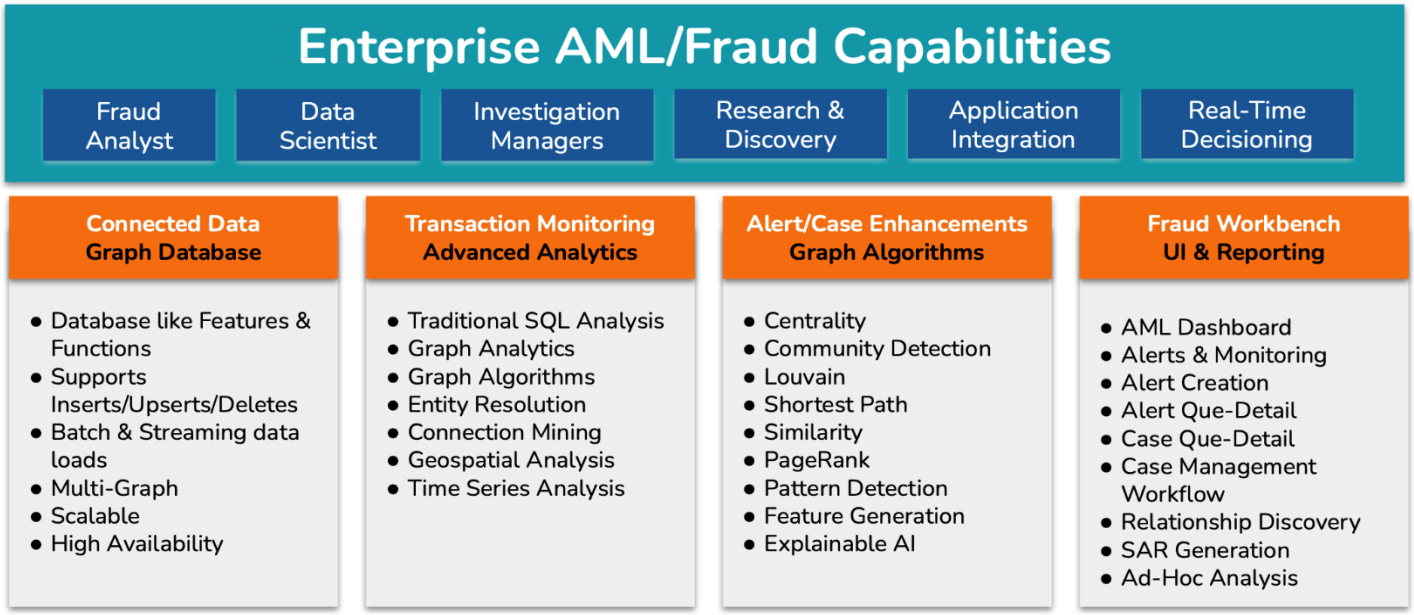

FIGURE 2. TigerGraph Fraud Detection and Investigation Toolkit capabilities

TigerGraph's Fraud Detection and Investigation Toolkit features built-in graph analytics, visualization with explainable ML/AI, no-code investigation, pattern and community algorithms, and deep learning model interpretability.

The Fraud Detection and Investigation Toolkit has functionality for every user. It can accelerate the steps it takes to detect fraud and analyze deep links and relationships in real-time and deliver more accurate results than alternative solutions.

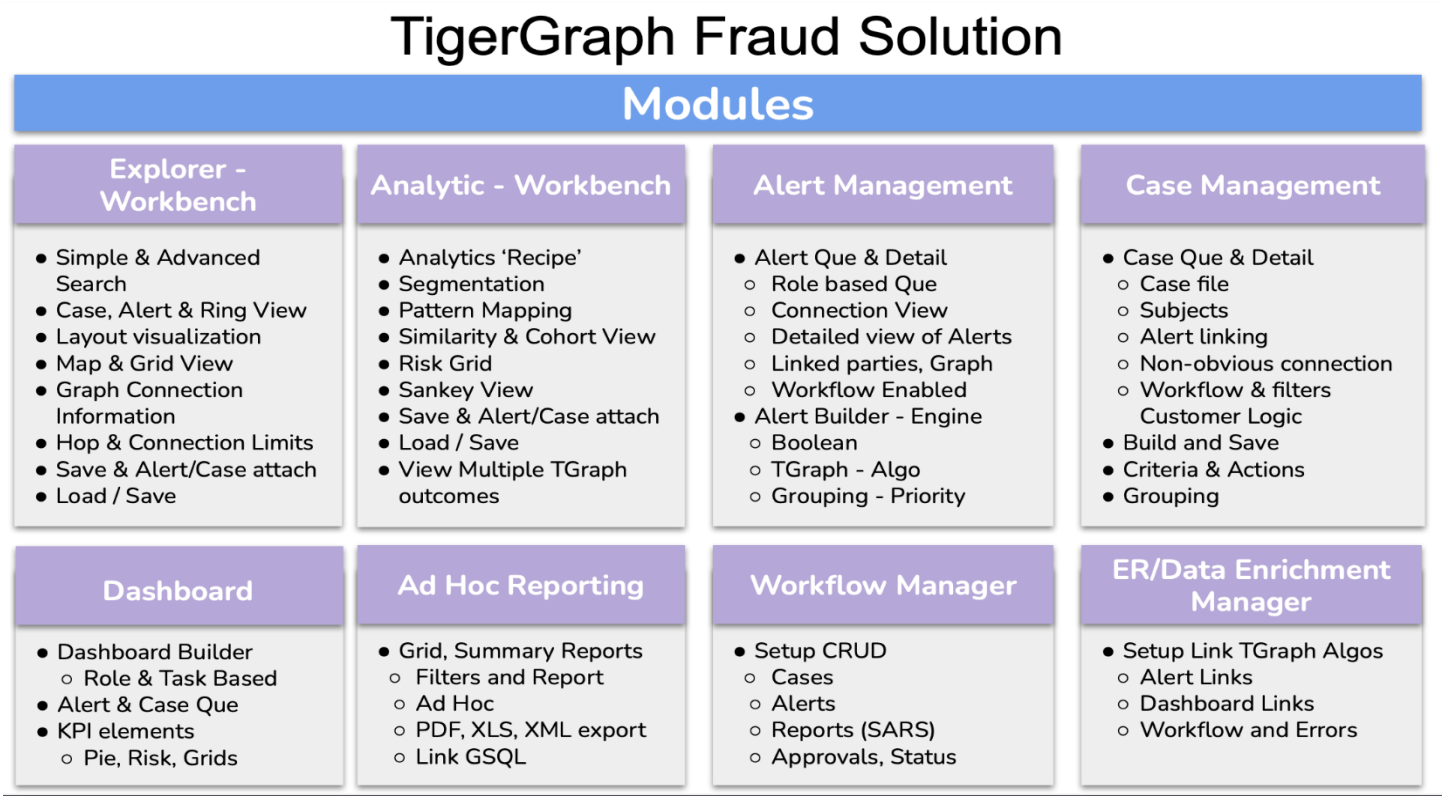

FIGURE 3. TigerGraph Fraud Detection and Investigation Toolkit Modules

TigerGraph represents a new generation, distributed computational model, and offers the following advantages:

- Faster data loading to build graphs quickly
- Faster execution of graph algorithms
- Real-time capability for streaming updates and insertions
- Ability to unify real-time analytics with large-scale offline data processing
- Ability to scale up and scale out for distributed applications

SOLUTION OVERVIEW

The solution in this Reference Architecture combines TigerGraph's native scale-out graph analytics software integrated with Xilinx Graph Analytics libraries to accelerate the Louvain algorithms all running on HPE ProLiant DL385 Gen10 Plus v2 servers with the AMD® EPYC® 7713 Processors and Xilinx Alveo U50 Data Center Accelerator Cards.

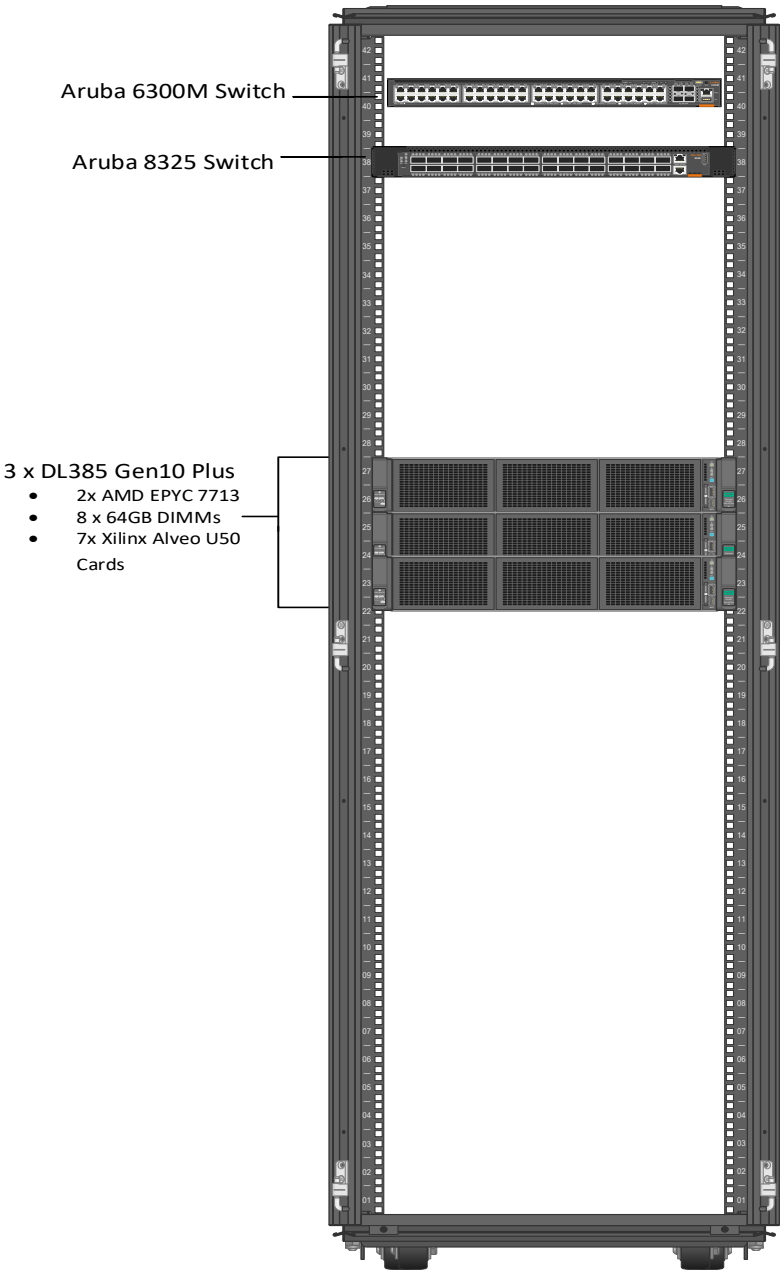

FIGURE 4. Architecture diagram using the HPE ProLiant DL385 Gen10 Plus v2 servers with Xilinx Alveo U50 Data center Accelerator cards

The solution enables fraud detection. In a fraud detection scenario, there will be multiple nodes and in this Reference Architecture, we have three nodes connected via an Aruba 8325 100GbE switch for cluster communication and Aruba 6300 for the management network.

Figure 5 demonstrates a high-level architectural view for the Multi-node configuration for TigerGraph using Xilinx U50 accelerator cards for Louvain fraud detection.

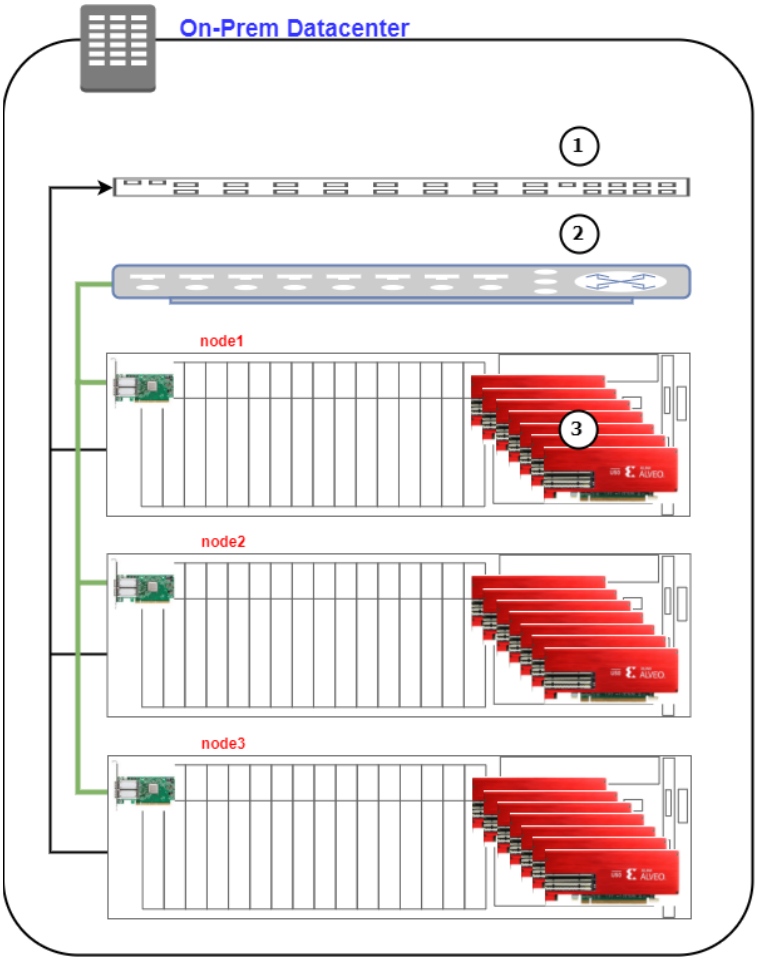

1 **10g Switch**

This can be dedicated switch or part of corp network with isolation used for server access.

2 **100g Switch**

This is used for Graph DB cluster communication and U50 Card to Card communication

3 **U50 Cards for acceleration**

This is used as acceleration card for graphDB

FIGURE 5. Multi-Node Architecture for TigerGraph with Xilinx Alveo U50 Data center Accelerator cards

SOLUTION COMPONENTS

HPE ProLiant DL385 Gen10 Plus v2 server

HPE ProLiant DL385 Gen10 Plus v2 server is redefining price or performance with the new math for virtualized compute. HPE ProLiant DL385 Gen10 Plus v2 server offers the 3rd generation AMD EPYC 7003 Series processors delivering increased performance over the prior generation. With 64 cores per processor for a total of 128 cores (per 2-socket configuration), 32 DIMMs for memory up to 3200 MT/s, and increased data transfer rates with PCIe Gen4 capabilities. The HPE ProLiant DL385 Gen10 Plus v2 server delivers low-cost virtual machines (VMs) with unprecedented security. Equipped with PCIe Gen4 capabilities, the HPE ProLiant DL385 Gen10 Plus v2 server offers improved data transfer rates and higher networking speeds. HPE ProLiant DL385 Gen10 Plus v2 server is a purpose-built platform for virtualization, high-performance computing, and memory-centric workloads.

Key-value points

- Greater VM density and VM price-performance with up to 64 cores per processor and 32 DIMM slots.
- 4 TB of memory capacity per CPU and 410 GB/s of memory bandwidth for faster in-memory performance.
- Increased memory-centric application performance with HPE DDR4 SmartMemory running at memory speeds up to 3200 MT/s.

- Increased security with silicon root of trust, secure memory encryption, and secure encrypted virtualization.
- Up to 61 TB of NVMe SSDs coupled with PCIe 4 with four lanes to each SSD delivering greater I/O bandwidth for fast access to large datastores.

Latest features

- The 3rd Generation AMD EPYC processors offer up to 64 cores.
- 7 nm at 225W process enables enterprise data centers to attain better power efficiency.
- Improved NUMA optimization eliminates memory performance reductions when multiple processors simultaneously attempt to access memory.

FIGURE 6. HPE ProLiant DL385 Gen10 Plus v2 server

For more details, refer to HPE ProLiant DL385 Gen10 Plus v2 server QuickSpecs.

Xilinx Alveo U50 Data Center Accelerator Card

The Xilinx® Alveo™ U50 Data Center accelerator card provides optimized acceleration for workloads in graph analytics and machine learning. Built on Xilinx UltraScale+ architecture and packaged up in an efficient 75-watt, small form factor, and armed with 100 Gbps networking I/O, PCIe Gen4, and HBM, Alveo U50 is designed for deployment in any server.

Alveo accelerator cards are adaptable to changing acceleration requirements and algorithm standards, capable of accelerating any workload without changing hardware, and reduces the overall cost of ownership.

FIGURE 7. Xilinx Alveo U50 Data Center Accelerator card

Aruba 8325 switch series

The Aruba 8325 switch series provides industry-leading line rate 1/10/25 GbE (SFP/SFP+/SFP28) and 40/100 GbE (QSFP+/QSFP28) connectivity in a compact 1U form factor. The Aruba 8325 rounds out Aruba's campus and data center switching portfolio with an enterprise

core and aggregation solution that ensures higher performance and higher uptime. The Aruba 8325 switch series is based on the ArubaOS-CX, a modern software system for the enterprise core that automates and simplifies many critical and complex network tasks, delivers enhanced fault tolerance and facilitates zero service disruption during planned or unplanned control-plane events.

FIGURE 8. Aruba 8325-32C 32-port 100G

Key features

- High-performance 6.4Tbps with 1905Mpps throughput
- High availability with redundant power supplies and fans
- Suitable for core or aggregation in the campus or Top of Rack (ToR) in the data center
- ArubaOS-CX enables automation and programmability using built-in REST APIs and Python scripts
- Intelligent monitoring and visibility with Aruba Network Analytics Engine
- Advanced Layer 2/3 feature set includes BGP, OSPF, VRF, and IPv6
- Compact 1U switches with 1/10/25GbE and 40/100GbE connectivity

Design principles

The design decisions are dependent on the size of the databases that need to be supported, the number of variables or dimensions that need to be vectorized in the database, and how many databases need to be supported at the same time.

For example, when storing 1TB of data in the DDR4 Host RAM, a rough estimate is that you will need 3x that amount in SSD capacity. This storage is used for partitions of persistent data, which is done to ensure that there is enough space for data such as meta-data, reliability replications, backup, and Xilinx partitions. For example, to support 5 Graphs on one node or server, and each is 1TB, then you would want to provision at least 15TB of SSD storage.

The number of Graphs that can run on a particular node or server depends on the size of the Graph databases and how often you want to run each database. Generally, there will only be 1 Graph running at a time, so depending on how often you need to access the graph database on the server will determine how many Graphs you store on a server. This is not a hard requirement, just a recommendation.

The Xilinx U50 Alveo card will cache only the relevant small subset of data in a compact vectorized format in the 8GB of High Bandwidth Memory (HBM). Using seven Xilinx U50 Alveo cards in a server, 56GB of HBM would be available for data. The size of each vertex determines how many vertices can be supported on each card. As more vertices are added, you can linearly scale support using additional Alveo cards. For this benchmark, 450M vertices were processed across three HPE ProLiant DL385 v2 nodes integrated with twenty-one Xilinx Alveo U50 cards.

System cluster specifications should be discussed with your HPE, TigerGraph, and Xilinx support teams to define the appropriate configuration for your particular use case.

Hardware

Three (3) HPE ProLiant DL385 Gen10 Plus v2 servers was used with the below configuration in the tested environment:

- **Processors**: 2x AMD EPYC 7713, 64 cores/processor, Base Frequency 2.0 GHz, 256 MB L3 Cache
- **Memory**: 512 GB RAM (8 x 64GB DIMMs)
- **Accelerator**: 7x Xilinx Alveo U50 Cards (depends on the size of the database)
- **100GbE Card**
- **Storage:** 6 TB Minimum up to 60 TB (depends on the size of the database). SATA or NVME

Software

The software components used for the Reference Architecture:

- **OS**: Ubuntu 20.04.2 LTS
- **TigerGraph 3.1.5 Enterprise Edition**
- **HPE System ROM**: A42 v2.40 1/21/2021 or later
- **HPE iLO 5 version**: 2.41 Mar 08 2021 or later

Refer to https://xilinx.github.io/graphanalytics/index.html for the recommendations for the Xilinx Run Time shell (XRT) and Xilinx Vitis Graph Analytics Library.

Xilinx Vitis Library Graph Analytics Library is designed to work with Vitis 2020.2 and later, and therefore inherits the system requirements of Vitis and Xilinx Run Time (XRT).

Application software

TigerGraph platform

TigerGraph is a native parallel graph database purpose-built for loading massive amounts of data (terabytes) in only hours and analyzing as many as 10 or more hops deep into relationships in real-time. It supports transactional as well as analytical workloads, is ACID compliant, scales up and out with automatic data partitioning. TigerGraph's proven technology has been effective in use cases such as detecting fraud, customer360, and supply chain optimization to make sense of ever-changing big data.

- **Speed** – Massively parallel processing provides sub-second response for queries with tens of millions of entities or relationships.
- **Scale-out** – Scales with your growing needs. Trillion-edge graphs are running real-time analytics in production.
- **Deep-link analytics** – Gain deeper insights through queries that can traverse 10 or more hops and perform complex analytics.
- **Graph Query Language** – The GSQL™ query language is the choice for high-performance graph operations and analytics. High-level syntax, Turing completeness, and built-in parallelism mean faster performance and development.
- **MultiGraph** – Multiple groups can share the same master database while retaining local control and security. This helps enterprises break down data silos, improving transparency and access to data.
- **Visual Interface** – TigerGraph GraphStudio™ is a simple yet powerful graphical user interface. GraphStudio integrates all the phases of graph data analytics into one easy-to-use graphical user interface.

Figure 9 shows the TigerGraph Analytics Platform with Xilinx Alveo U50 Data Center Accelerator Card.

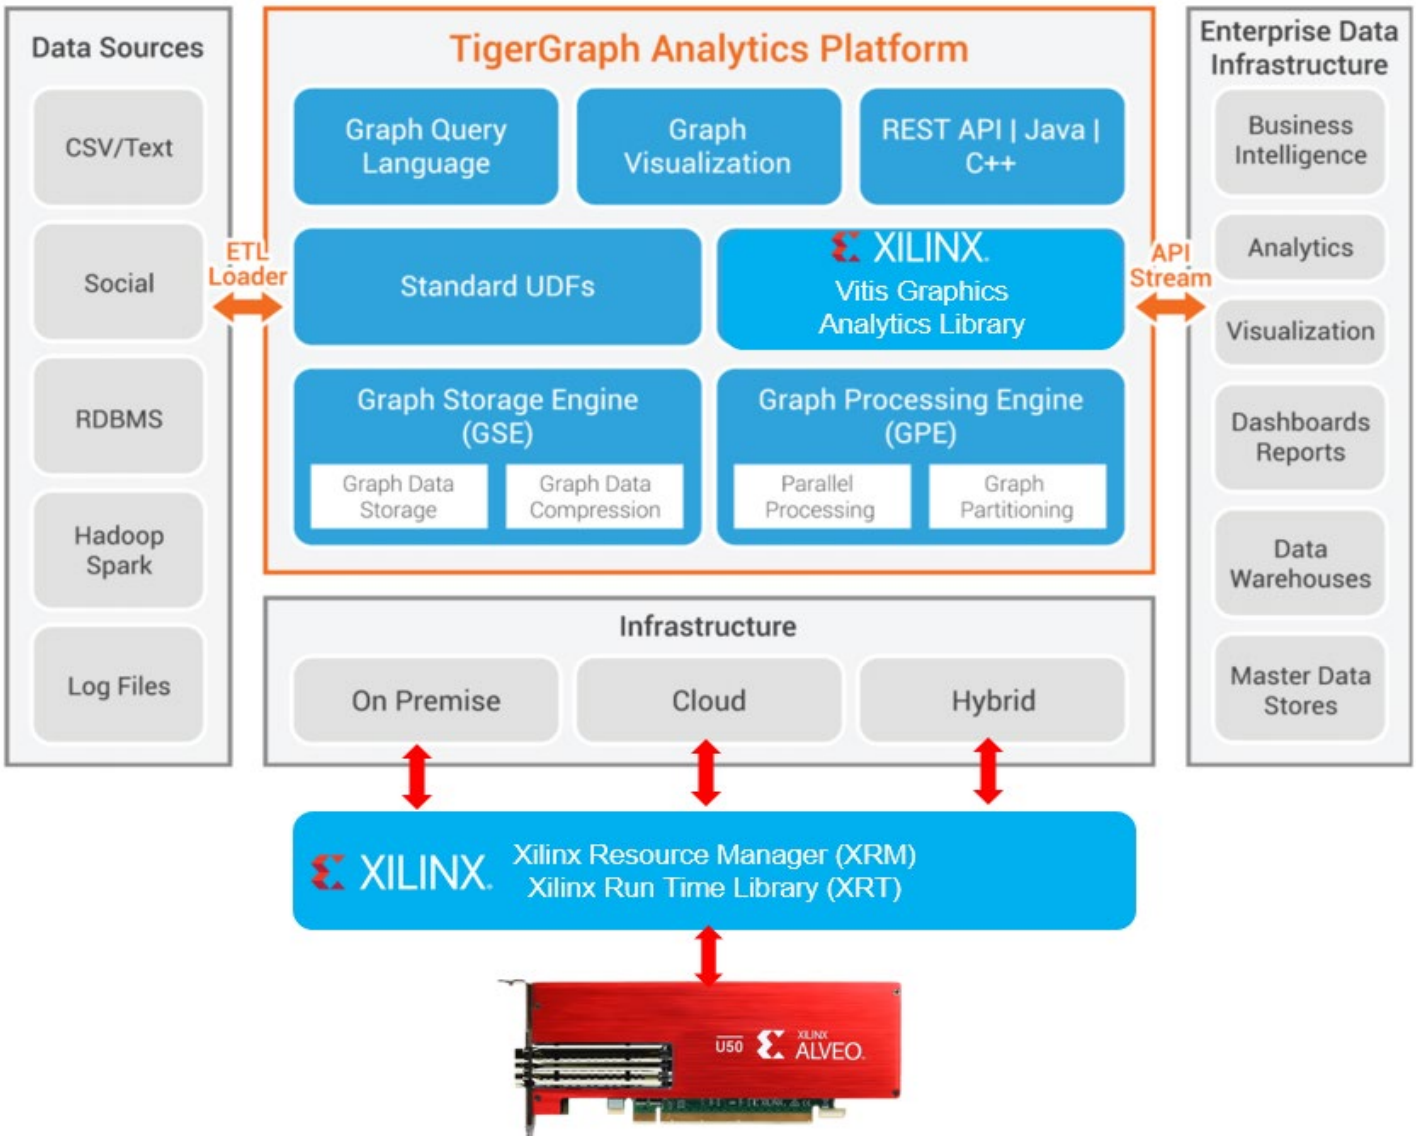

FIGURE 9. TigerGraph Analytics Platform with Alveo

It's easy to develop a solution on top of the TigerGraph platform.

For data scientists, analytics and Machine Learning engineers TigerGraph offers GraphStudio. GraphStudio integrates all the phases of graph analytics into one easy-to-use user interface. GraphStudio is great for ad-hoc, interactive analytics, and for learning to use the TigerGraph platform.

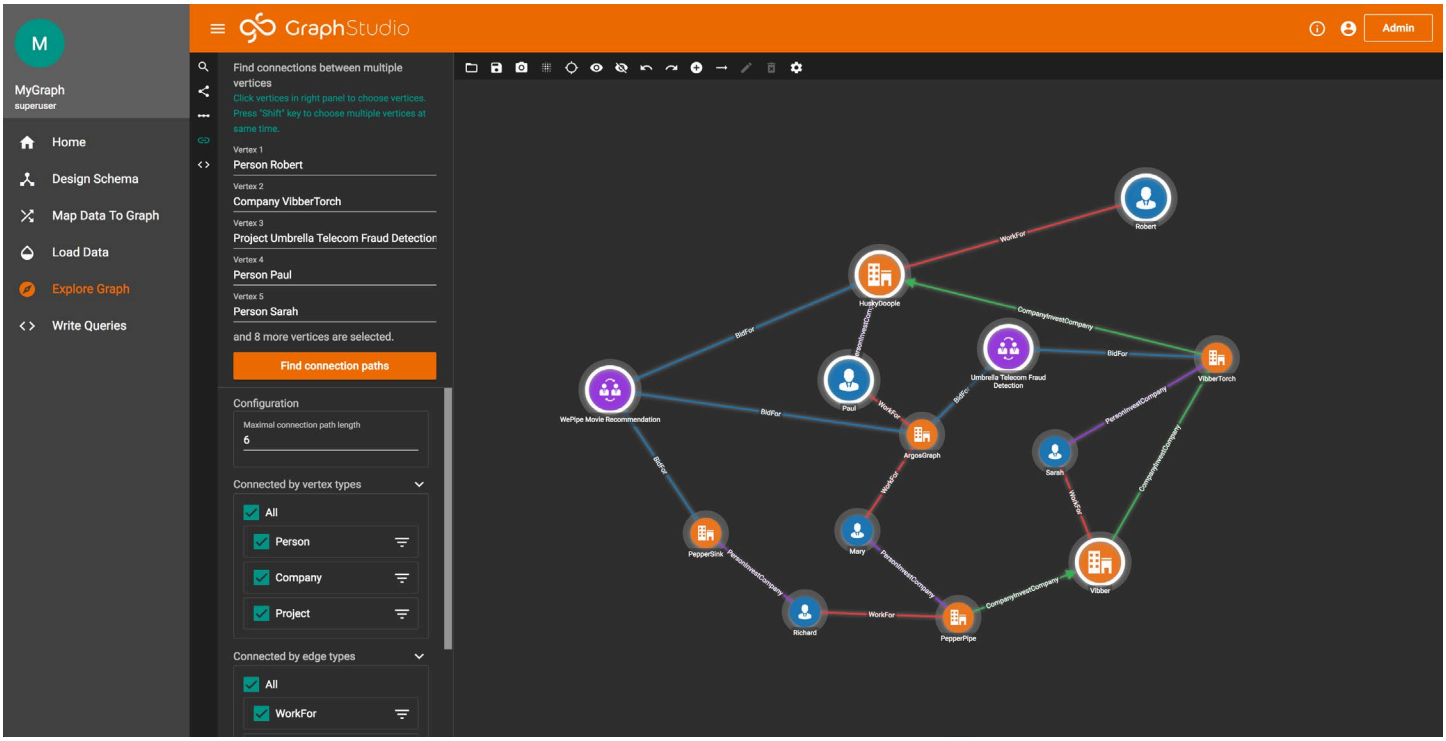

FIGURE 10. TigerGraph GraphStudio

Free Xilinx, Vitis Graph Analytics Library

The Vitis Graph Library is an open-sourced Vitis library written in C++ for accelerating graph applications for a variety of use-cases. The plug-ins consist of two major components:

- Xilinx Vitis Graph Analytics Library is provided as a custom TigerGraph User Defined Function (UDF) that can be called directly from GSQL.
- Xilinx Resource Manager (XRM) and Xilinx Run Time (XRT) libraries are installed on all processing nodes to manage resource allocation on Alveo Acceleration cards and data movement between the CPU and the FPGA.

See the Vitis Graph Analytics Libraries at, https://github.com/Xilinx/Vitis_Libraries/tree/master/graph.

Xilinx also offers starter kits for data scientists to get started quickly for top use cases:

- Detecting communities of fraud in payments
- Product recommendation engine or Customer 360

Xilinx Vitis Graph Analytics Library includes:

- **Similarity analysis**: Cosine Similarity, Jaccard Similarity, k-nearest neighbor
- **Centrality analysis**: PageRank
- **Pathfinding**: Single Source Shortest Path (SSSP), Multi-Sources Shortest Path (MSSP)
- **Connectivity analysis**: Weakly connected components and strongly connected components

- **Community Detection**: Louvain Modularity (New for 2021.1), Label Propagation, and Triangle Count
- **Search**: Breadth-First Search and 2-Hop Search (New for 2021.1)
- **Graph Format**: Calculate Degree and Format Convert between CSR and CSC
- Fuzzy Search
- Random forest
- Multi-Layer Perceptron (MLP)
- Conjugate Gradient Solvers (CG)
- Maximum Independent Set (coming soon)

BEST PRACTICES AND CONFIGURATION GUIDANCE FOR THE LOUVAIN FRAUD DETECTION SOLUTION

TigerGraph uses GSQL query language for fast and scalable graph operations and analytics. A GSQL script of detecting fraud using Louvain modularity is created to baseline the functionalities and computation complexity of the algorithm. It is used to verify and validate the functional correctness and performance improvement of the Alveo accelerated design. Below is the block diagram of Louvain modularity computation in GSQL, which is executed entirely on the CPU.

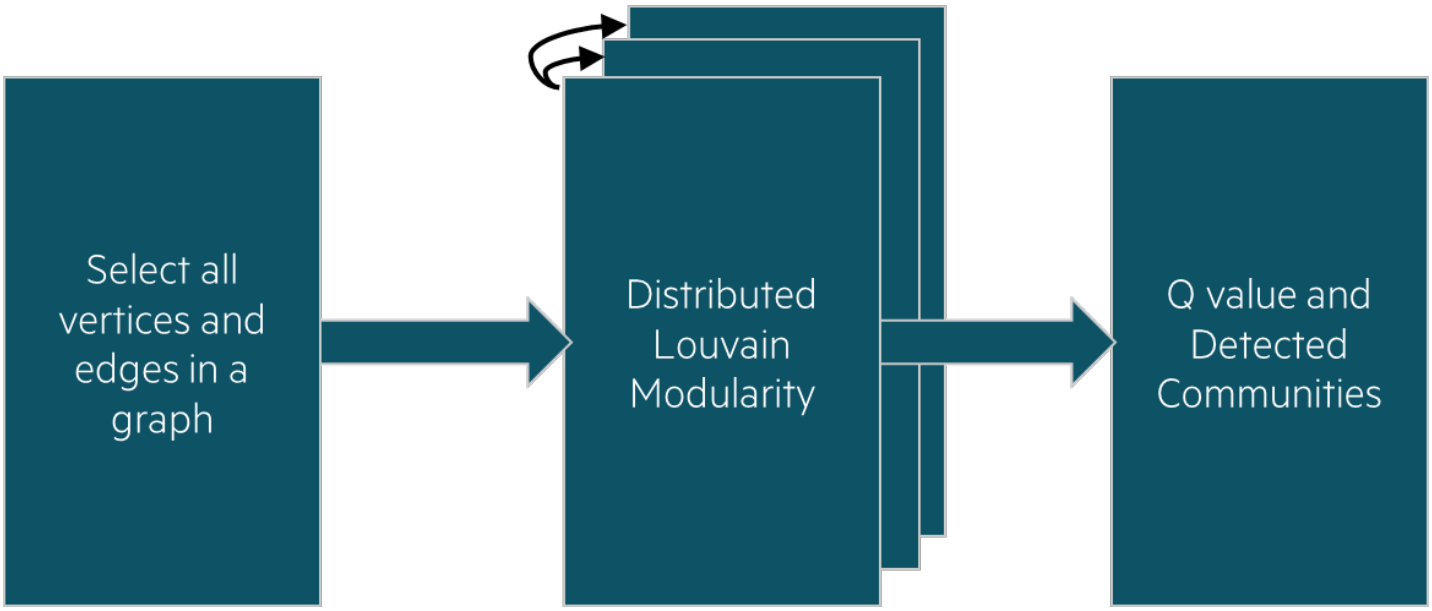

FIGURE 11. Block diagram of Louvain Modularity in GSQL on CPU only

Xilinx Vitis Graph Analytics Library plugin provides a user-defined function that offloads the computation of Louvain Modularity to the FPGA:

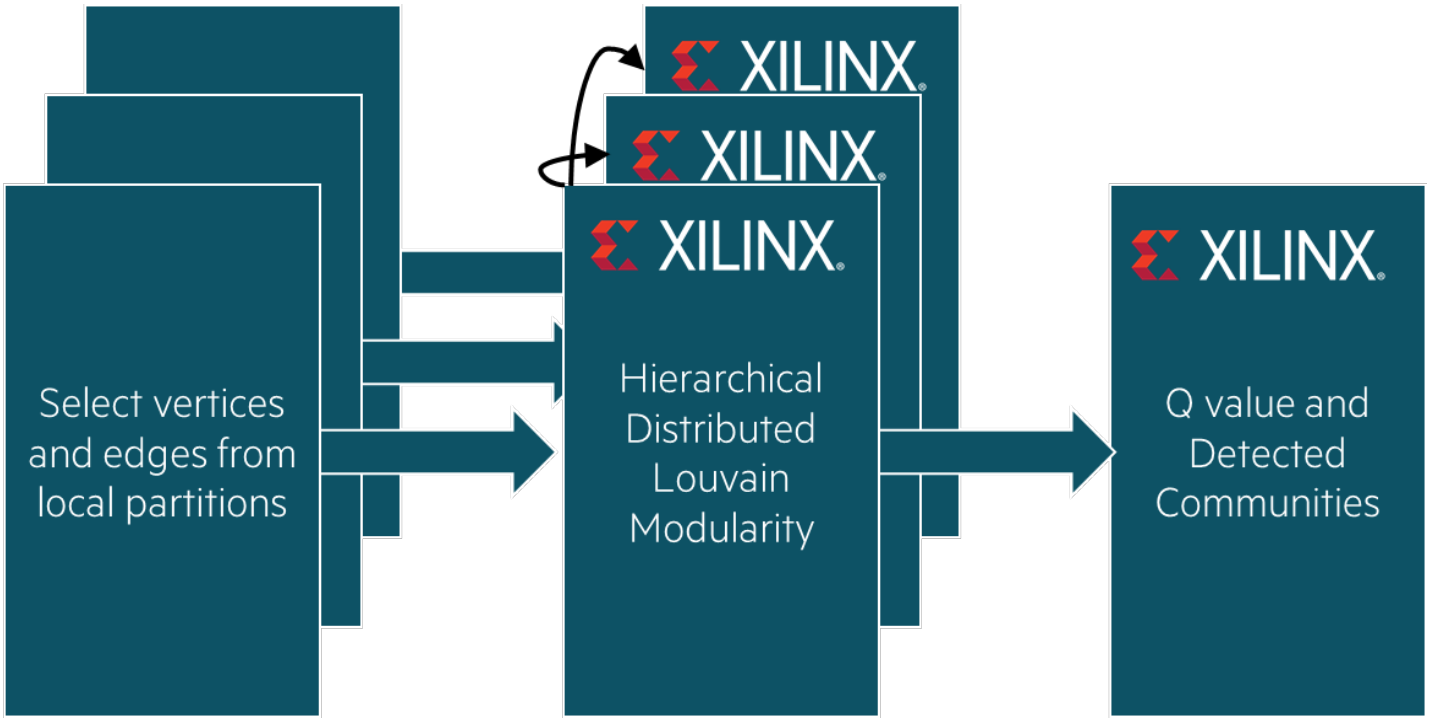

FIGURE 12. Block diagram of Louvain Modularity computation in GSQL on CPU + Xilinx Alveo

The details of how this is accomplished are described in the "Getting started" and "User guide" sections below:

Getting started

An overview of the Xilinx Alveo Graph Analytics products and how to get started is at the following link:

https://xilinx.github.io/graphanalytics/

This link has the information about the Louvain Modularity overview, installation, and demo. This page also has the links for TigerGraph Plugin Methodology and Vitis Development Flow as well.

User guide

For overview and installation instructions go to the following links:

- Overview of Louvain Modularity
- Install Xilinx LouvainMod and Community Detection

After the installation has been completed, you can run the demonstration via the link below:

- Xilinx Louvain Modularity Alveo Product Demos

Once the demos have been successfully completed, the next step is to integrate Xilinx's Graph Library IP with TigerGraph's software. The URL below explains how the L3 Vitis Graph Library function is ported into TigerGraph SW framework, and finally running the application on Alveo by calling GSQL query from TigerGraph SW.

- Integration with TigerGraph

The Vitis Development Flow can be used to create custom accelerator functions targeting Xilinx U50 Alveo cards. Below is an overview and GitHub URL for reference.

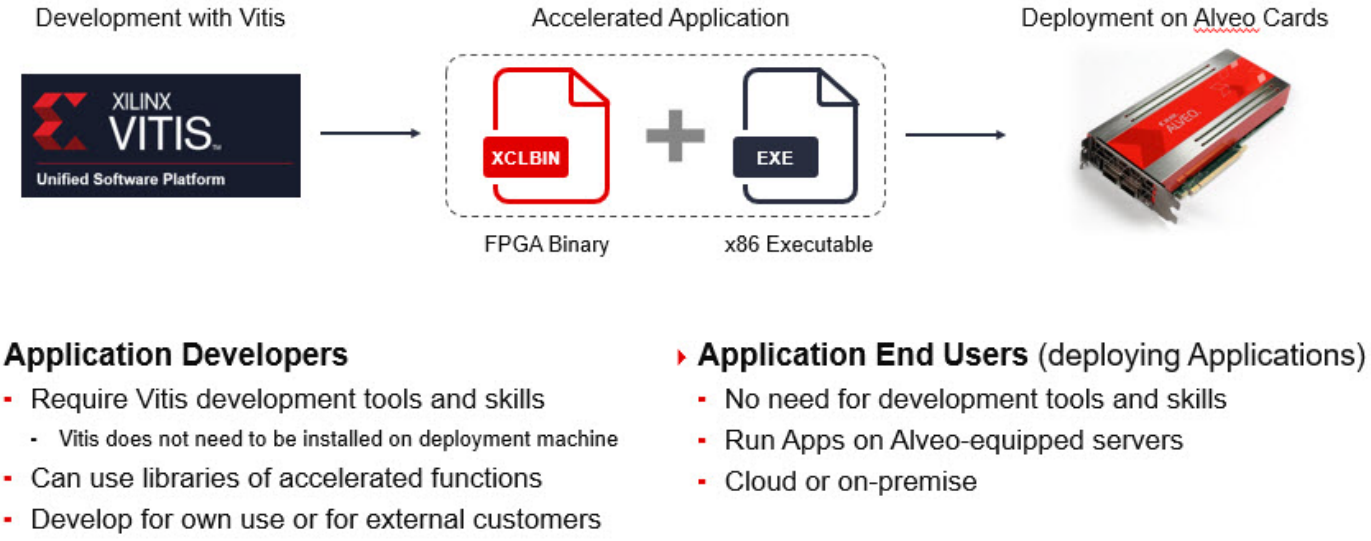

FIGURE 13. Development Flow for Users who want to create their own application to run on Alveo

Instructions for this can be found at the following link:

- Vitis Development Flow

CAPACITY AND SIZING

Workload description

Graph techniques can analyze a large volume of data points – and the crucial relationships between them – to deliver fraud alert scores. This can be done in real-time with the addition of Xilinx Alveo accelerators. Graph can be used for fighting financial fraud by analyzing the links between people, phones, and bank accounts (among other things) to reveal indicators of fraudulent behavior, not only helping banks pinpoint suspicious activity in a sea of data but also giving them the tools to explain what's going on.

A key feature of a graph is its ability to perform at speed, especially compared to relational database solutions such as SQL. Banks have been doing fraud detection for years, but one of the things that graph brings to the party – apart from depth of analysis – is speed. While SQL depends on bulky table joins, the graph is less memory intensive and able to handle a greater query load.

Fraud detection systems tend to rely on looking at transactions that exceed preset levels or people who try to max out a credit card with no intention of paying it off. These types of suspicious transactions are easy to detect because they rely primarily on the information in the transaction itself, looking at the amount, the destination, or other properties that might generate warning signals.

However, fraud has become more complex than that as fraudsters have learned to work across multiple accounts, including mule accounts not directly controlled by or associated with them, such that individual transactions look ordinary and would not trigger an investigation. This is an especially big problem as banks work to comply with anti-money laundering (AML) regulations, pushing fraudsters to ever greater lengths to obfuscate their activities.

Fraudsters will employ hundreds of accounts to launder money in a technique known as 'smurfing'. In this technique, money that needs to be laundered is disbursed to accounts in quantities small enough to avoid triggering automated reporting limits. The money is then 'layered', that is, mixed, divided, and transferred from these accounts to other accounts, with the history of the money becoming more and more difficult to trace with every transfer. Ultimately it is funneled to the final destination, its origins virtually lost in a long and complex audit trail.

Layering requires a certain amount of setup as the fraudster must create all of the accounts they need with the banks, but once this is in place, it becomes easy to move large amounts of money quickly thanks to electronic banking application programming interfaces (APIs).

This is a network of activity, a pattern that can be traced and matched if you have the computational systems and resources in place to do it.

Many financial institutions have built their fraud detection systems on the backs of their legacy SQL databases because that's where the data is stored. While these SQL databases are optimal for processing transactions, they struggle with relational analysis as the number of connections or hops between data points grows.

To understand fraud, you have to examine a lot of contextual information. This brings in several different types of data – accounts, individuals, transactions, etc. – all of which, in a relational database, would be stored in separate tables. To link this data together, you have to make a table join, a temporary construct created in memory that allows you to read across the associated data and extract the information that you require.

Analysis and recommendations

Benchmark results

Here is a video you can watch that explains the benchmark results based on the setup and use case described in previous sections.

The benchmark is run on a dataset of 450 Million vertices running the Louvain community detection algorithm and compares the results on CPU vs CPU + Xilinx Alveo U50. The results are in the table and graphs below.

	TigerGraph (CPU Only)			TigerGraph (CPU + Xilinx Alveo)			Advantages with Xilinx Alveo		
Vertices (Millions)	Memory (GB)	Cluster Score (Q)	Time (Sec)	Memory (GB)	Cluster Score (Q)	Time (Sec)	Memory	Cluster Score	Time
450	235	0.6529	5281	77.6	0.77325	233.4	-67%	18%	22.6

FIGURE 14. Xilinx Alveo U50 speeds up the solution, improves the accuracy, and reduces the memory required

With the Xilinx Alveo accelerators in the solution, the results are more than 20x faster.

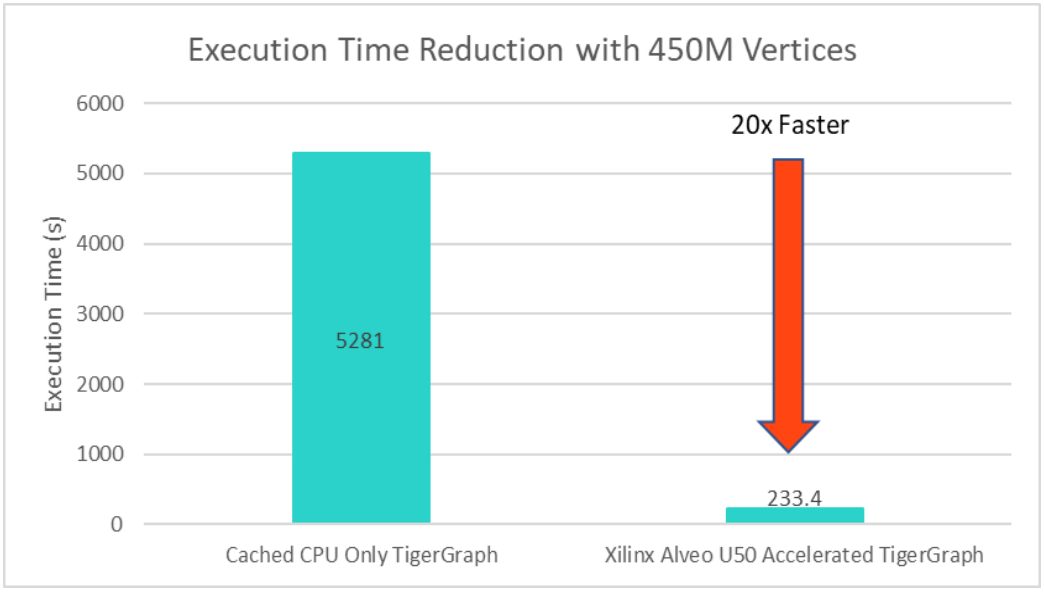

FIGURE 15. TigerGraph Execution Time Reduction with Alveo Accelerators

With the Xilinx Alveo accelerators in the solution, the memory required is 66% lower.

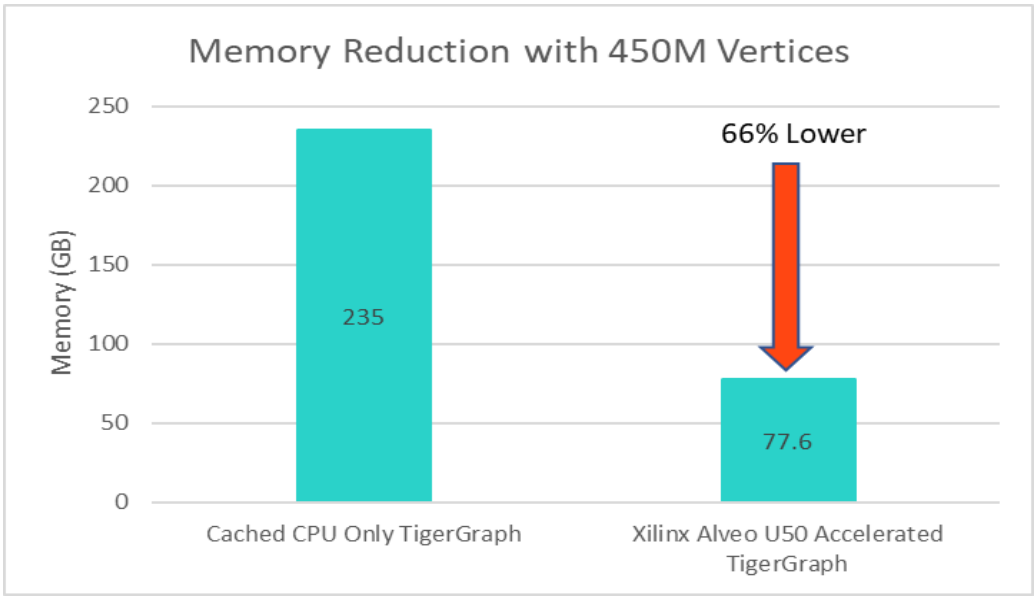

FIGURE 16. TigerGraph Host System Memory Reduction with Alveo Accelerators

With the Xilinx Alveo accelerators in the solution, the accuracy of the results increases by 18%.

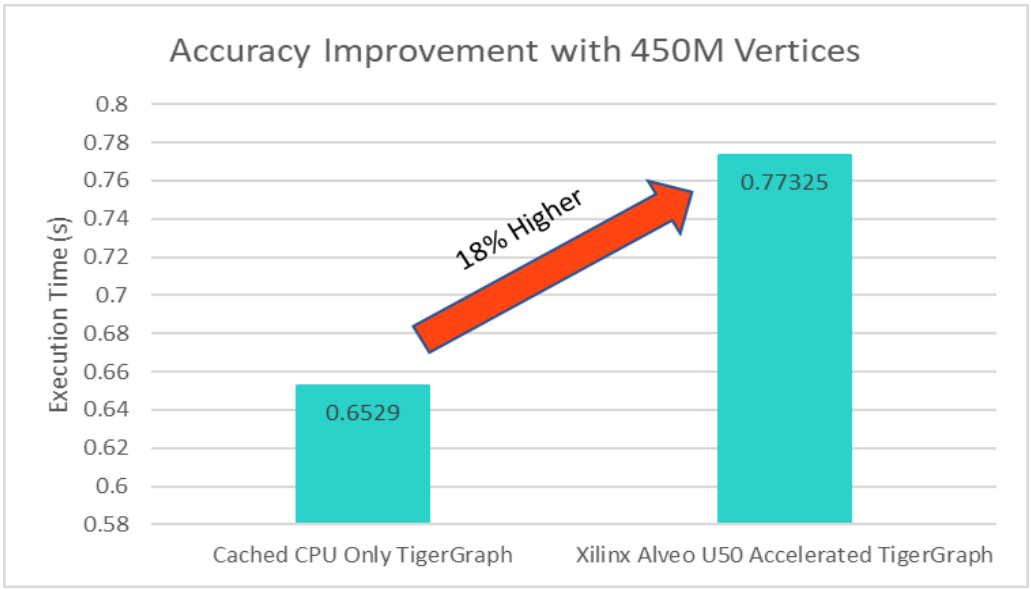

FIGURE 17. TigerGraph Accuracy Improvement with Alveo Accelerators

KEY POINT

- There are 3x HPE ProLiant DL385 Gen10 Plus v2 servers and 7x Xilinx Alveo U50 per node used for testing and illustration purpose in this Reference Architecture, however, the number of servers and number of accelerator cards per node can be scaled up or scaled down according to data size.
- As the database grows larger, the performance on Xilinx Alveo U50 stays almost linearly flat while the time on CPU grows exponentially.
- When results are needed fast, then Alveo accelerators along with the TigerGraph database meet that need.

SUMMARY

Whether for detecting fraud, anti-money laundering, cyber forensics and ransomware, sanction and trade compliance, audit and law enforcement, or another real-world challenge, the ability to quickly and efficiently explore, discover and predict complex relationships is a huge competitive differentiator for businesses today. Getting it done involves more than merely having connected data – it's about real-time and up-to-date correlation, detection, and discovery. Organizations need to be able to transform structured, semi-structured, and unstructured data and massive enterprise data silos into an intelligent, interconnected, and operational data network that can reveal critical patterns and insights to support business goals.

This accelerated graph analytics solution is capable of drawing deep connections between data entities in real-time, it's designed for scale and performance to solve complex problems, and fast insights from interconnected data.

This solution offers customers:

- 20x faster accelerated Louvain community detection when detecting fraud ring activity.
- An 18% improvement in accuracy if not higher.
- Reduces the system RAM requirements by 66%.

IMPLEMENTING A PROOF-OF-CONCEPT

As a matter of best practice for all deployments, Hewlett Packard Enterprise recommends implementing a proof-of-concept using a test environment that matches as closely as possible the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept (PoCs), contact a Hewlett Packard Enterprise Services representative (hpe.com/us/en/services/consulting.html) or TigerGraph or Xilinx (contacts at the end). PoCs are based on the specific use case(s) and will include other enterprise technical requirements with well-defined success criteria and conclude with an Executive Presentation. PoCs usually take between two and four weeks.

As referenced above, you can reproduce the Fraud Detection use case at https://xilinx.github.io/graphanalytics/index.html.

IMPLEMENTING A PILOT

Custom Pilots focus on a single use case such as detecting communities of Fraud. HPE or Xilinx or both and TigerGraph can build a pilot implementation of the solution that can be rolled out for user acceptance testing. Pilots will focus on use case functionality, system integration, and other components that are required to move the solution into production. Pilots typically take about four weeks once customer requirements are clearly understood and documented.

APPENDIX A: BILL OF MATERIALS

The following BOMs contain electronic license to use (E-LTU) parts. Electronic software license delivery is now available in most countries. Hewlett Packard Enterprise recommends purchasing electronic products over physical products (when available) for faster delivery and for the convenience of not tracking and managing confidential paper licenses. For more information, please contact your reseller or a Hewlett Packard Enterprise representative.

NOTE

Part numbers are at the time of publication or testing and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult with your Hewlett Packard Enterprise Reseller or Hewlett Packard Enterprise Sales Representative for more details hpe.com/us/en/services/consulting.html.

TABLE A1. Bill of materials for HPE ProLiant DL385 Gen10 Plus v2 server

Part number	Qty	Description
P38411-B21	3	HPE ProLiant DL385 Gen10 Plus v2 8SFF Configure-to-order Server
P38693-B21	6	HPE processor AMD EPYC 7713 for Gen10 Plus Servers
P07650-B21	24	HPE 64GB (1x64GB) Dual Rank x4 DDR4-3200 CAS-22-22-22 Registered Smart Memory Kit
P40502-B21	3	HPE 480GB SATA 6G Mixed Use SFF BC Multi-Vendor SSD
P40505-B21	6	HPE 3.84TB SATA 6G Mixed Use SFF BC Multi-Vendor SSD
P14581-B21	3	HPE DL38X Gen10 Plus x8/x8 Tertiary Riser Kit
R4B02C	21	Xilinx Alveo U50 Accelerator for HPE
P14587-B21	3	HPE DL38X Gen10 Plus x8/x16/x8 Secondary Riser Kit
P31246-B21	3	HPE Ethernet 100Gb 1-port QSFP28 PCIe3 x16 MCX515A-CCAT Adapter
P01366-B21	3	HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit
804331-B21	3	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
P10115-B21	3	Broadcom BCM57414 Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
P14608-B21	3	HPE DL38X Gen10 Plus Maximum Performance Fan Kit
865438-B21	6	HPE 800W Flex Slot Titanium Hot Plug Low Halogen Power Supply Kit
P13771-B21	3	HPE Trusted Platform Module 2.0 Gen10 Plus Black Rivets Kit
P14610-B21	3	HPE DL38X Gen10 Plus High-Performance Heat Sink Kit
P22018-B21	3	HPE DL38X Gen10 Plus 2U SFF Easy Install Rail Kit

TABLE A2. Bill of materials for Aruba Switches

Part number	Qty	Description
JL627A	1	Aruba 8325-32C 32-port 100G QSFP+/QSFP28 Back-to-Front 6 Fans and 2 Power Supply Bundle
JL307A	6	Aruba 100G QSFP28 to QSFP28 3m Direct Attach Copper Cable
R0Z25A	3	Aruba 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable
JL483B	1	Aruba X474 4-post Rack Kit
J9583B	1	Aruba X414 1U Universal 4-post Rack Mount Kit
JL716A	1	Aruba X544 Universal 4-post Duct Kit (Must order 4-post rack mount kit separately)
JL762A	1	Aruba 6300M 48-port 1GbE and 4-port SFP56 Power-to-Port 2 Fan Trays 1 PSU Bundle
J9283D	2	Aruba 10G SFP+ to SFP+ 3m Direct Attach Copper Cable
R0M46A	2	Aruba 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable

Part number	Qty	Description
JL760A	1	Aruba X371 12VDC 250W 100-240VAC Power-to-Port Power Supply

RESOURCES AND ADDITIONAL LINKS

HPE Reference Architectures, hpe.com/info/ra

HPE Servers, hpe.com/servers

HPE Storage, hpe.com/storage

HPE Networking, hpe.com/networking

HPE Technology Consulting Services, hpe.com/us/en/services/consulting.html

Xilinx Alveo Graph Analytics Products, https://xilinx.github.io/graphanalytics/index.html

TigerGraph Solution on HPE, https://www.tigergraph.com/hpe/

To help us improve our documents, please provide feedback at hpe.com/contact/feedback.

a50004892enw, September 2024

Hewlett Packard Enterprise