GRAPH+AI SUMMIT organized by TigerGraph

Graph + AI: What is Next ?
Democratizing Graph Technology Adoption
- The TigerGraph Perspective

Jay Yu, Ph.D. - VP Product & Innovation, TigerGraph

Victor Lee, Ph.D. - VP ML & AI, TigerGraph

October 2021

Speaker Bio

Jay Yu Ph.D.

Vice President of Product & Innovation

TigerGraph

Proven hands-on full-stack innovator, strategic thinker, leader, and evangelist for new technology and product, with 25+ years of industry experience in big tech companies like Intuit and Teradata.

Victor Lee Ph.D.

Vice President of ML / AI

TigerGraph

Product leader, author, and educator, with a passion for algorithms, languages, user experience, and ethics. 7 years at TigerGraph, 3 years as university professor, 20+ years in tech industry.

Outline

1. Graph Overview & Adoption Challenges
2. Graph Adoption Successes So Far
3. New Capabilities to Expand Adoption
4. Innovation to Accelerate Adoption
5. Graph + ML Advancement
6. Summary

Graph Overview & Adoption Challenges

By 2025, graph technologies will be used in 80% of data and analytics innovations, up from 10% in 2021, facilitating rapid decision making across the enterprise.[1]

"To Graph or Not to Graph? That Is *Not* the Question — **You Will Graph.**"[2]

Gartner®

[1]*Gartner, Top Trends in Data and Analytics for 2021, 16 February 2021*
[2]*Gartner, Graph Steps Onto the Main Stage of Data and Analytics: A Gartner Trend Insight Report, 14, December 2021*

Innovative Business Leaders Agree …

- "**Graphs clarify relationships**. Put simply, relying upon entire paths or sets of relationships in a graph can **provide deeper insights** than just looking at nearest neighbors, which is what a typical database will afford you." - Brad Spiers, JP Morgan Chase

- "Advanced analytics has **de-siloed** our business, contributing to **$100M+ in incremental annual profit**." - Harry Powell, *Jaguar Land Rover*

- "The value of graph, with AI algorithms on top, is that not only can we monitor an enterprise, **we can predict** - and that allows us to **avoid problems before they happen.**" - Ed Sverdlin, *UnitedHealth Group*

Why Graph is becoming more and more Important ?

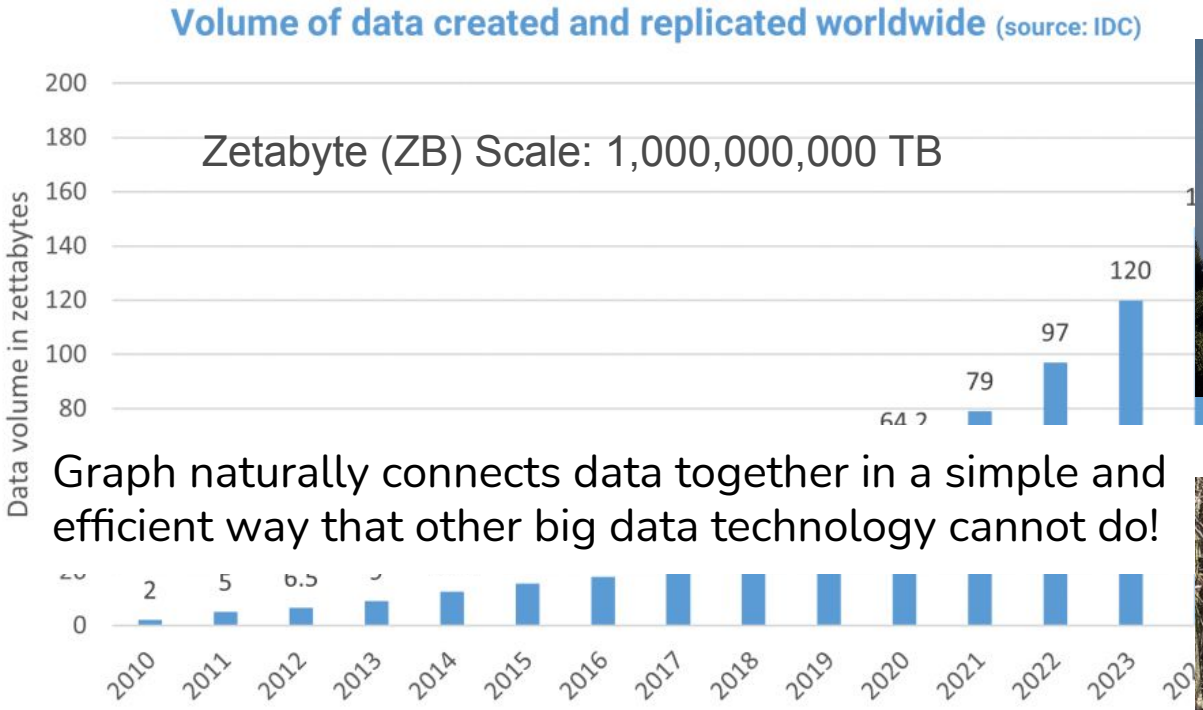

Graph naturally connects data together in a simple and efficient way that other big data technology cannot do!

Data Silo

Data Swamp

What is Graph?

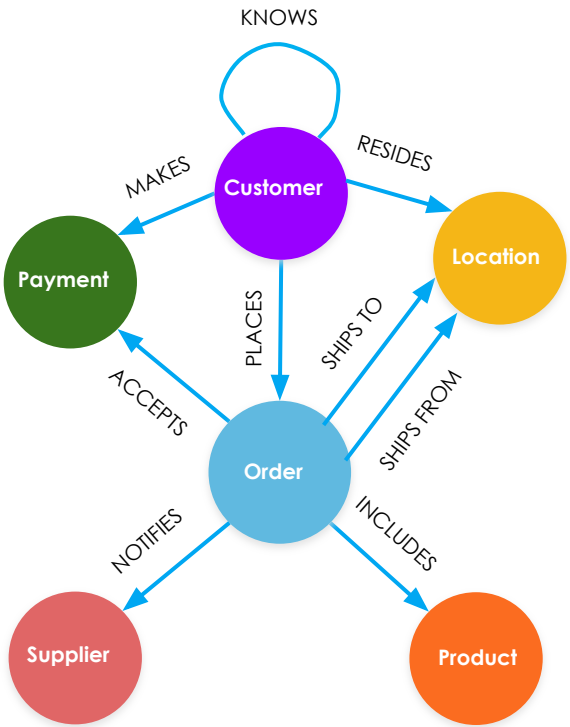

- Graph is Intuitive
- Graph is Meaningful

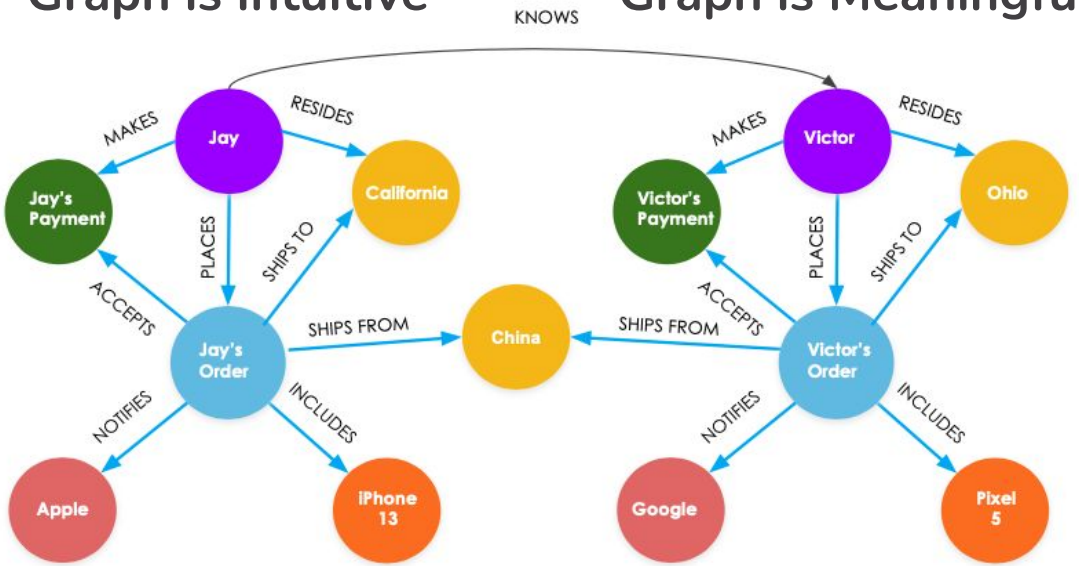

- Graph is Insightful
- Graph is Explainable

Graph Connects Data Together and Elevate **Relationships to be the First Class Citizens**!

How is Graph Model different than other Data Models ?

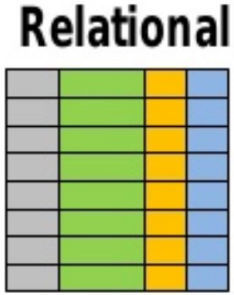

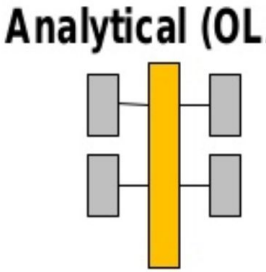

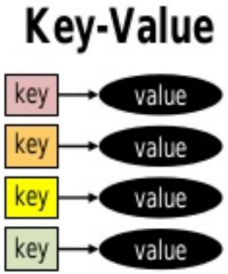

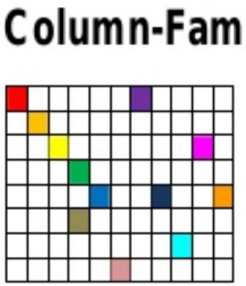

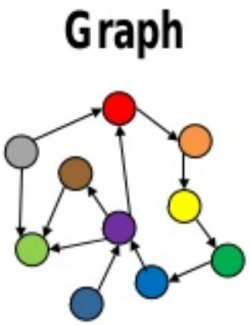

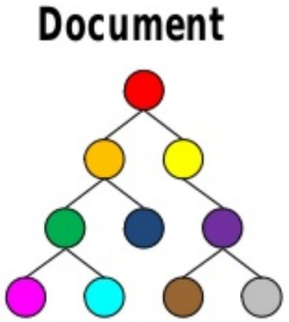

Graph Model can represent all other Data Models **Naturally** !

Author: Dan McCreary: https://www.slideshare.net/Dataversity/nosql-now-nosql-architecture-patterns-23589170

What are the Barriers to Adopt Graph Tech?

- Is it valuable to business? customer benefit / business value?
- Is it real to scale? enterprise-scale ready?
- Is it hard to learn? by engineers / data scientists ...?
- Is it long to build? time to value?
- Is it costly to operate? huge effort to manage?

Graph Adoption Successes So Far

Fortune-500/1000 Companies' Graph Adoption

CONNECT ALL DATASETS AND PIPELINES

Friction-free scale up from GB to TB to Petabyte with lowest cost of ownership

UNITEDHEALTH GROUP®

Customer 360 connecting 200+ datasets and pipelines

Fortune 50 Retailer

Item 360 for eCommerce across 100+ datasets

xandr AT&T

Identity graph connecting multiple data pipelines

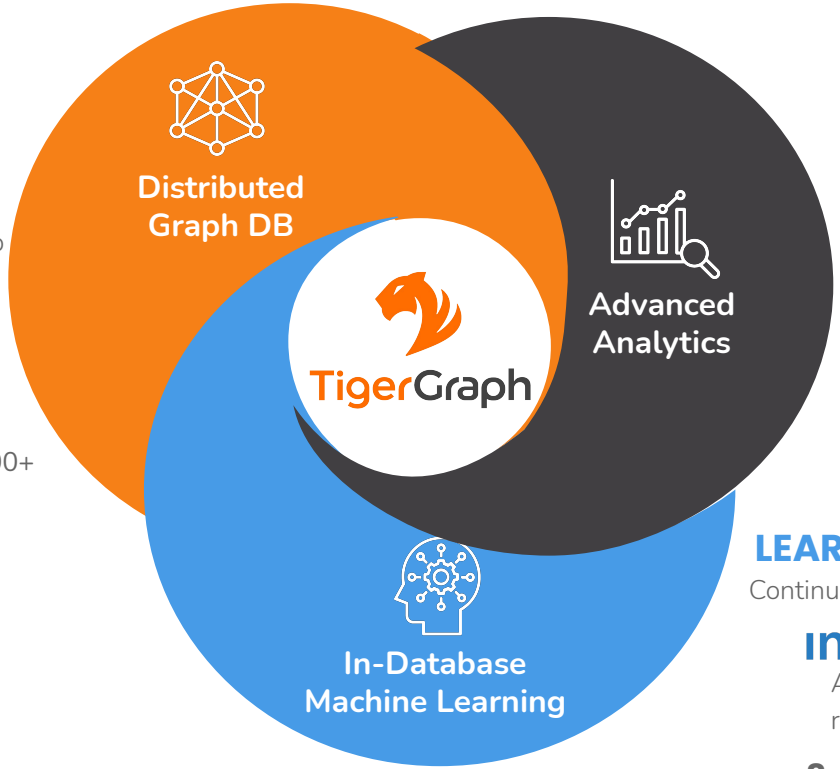

ANALYZE CONNECTED DATA

10-100X faster than current solutions

Jaguar Land Rover

Supply chain planning accelerated from 3 weeks to 45 minutes

China Mobile International

Fraud Detection - batch to real-time for 750 million calls/day

LEARN FROM CONNECTED DATA

Continuous graph-based feature generation and training

intuit

AI-based Customer 360 for entity resolution, recommendation engine, fraud detection

8 out of top 10 global banks

Real-time fraud detection and credit risk assessment

Rapid Growing Customer Adoption - Use Cases

Supply Chain Management	Customer 360	Anti-Money Laundry	IT Network Analysis
Data 360	Entity Resolution	Fraud Detection	Social Network Analysis
Investment Opportunity Analysis	Recommendation Systems	Cybersecurity	Knowledge Graphs

How? - Performant and Scalable Graph DB

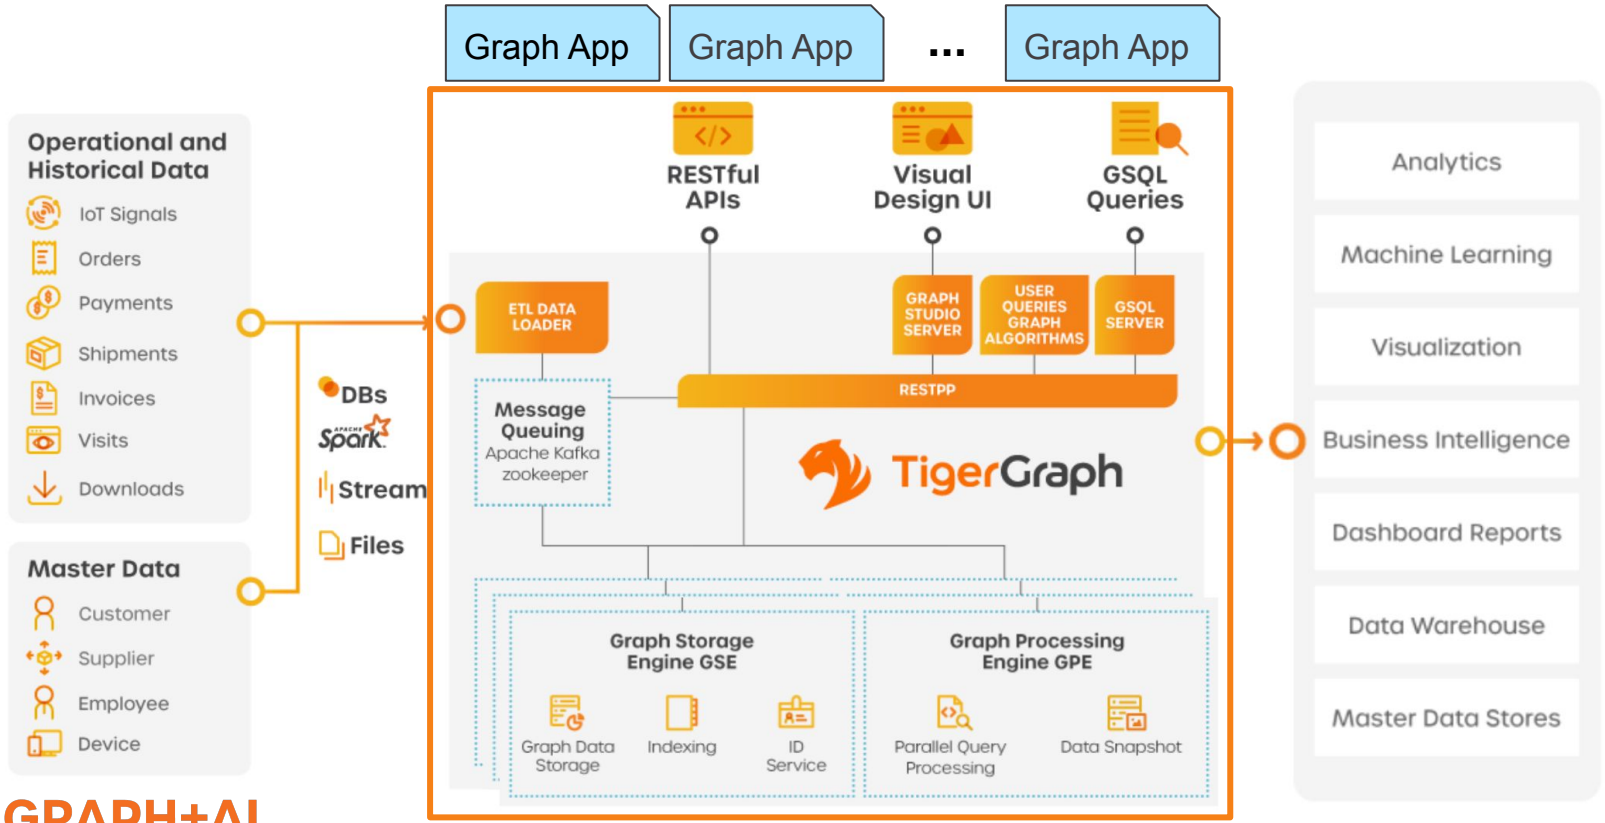

How? - Familiar, Powerful Query Language (GSQL)

Schema-Based (Strong Typing)

SQL-Like

Parameterized Queries

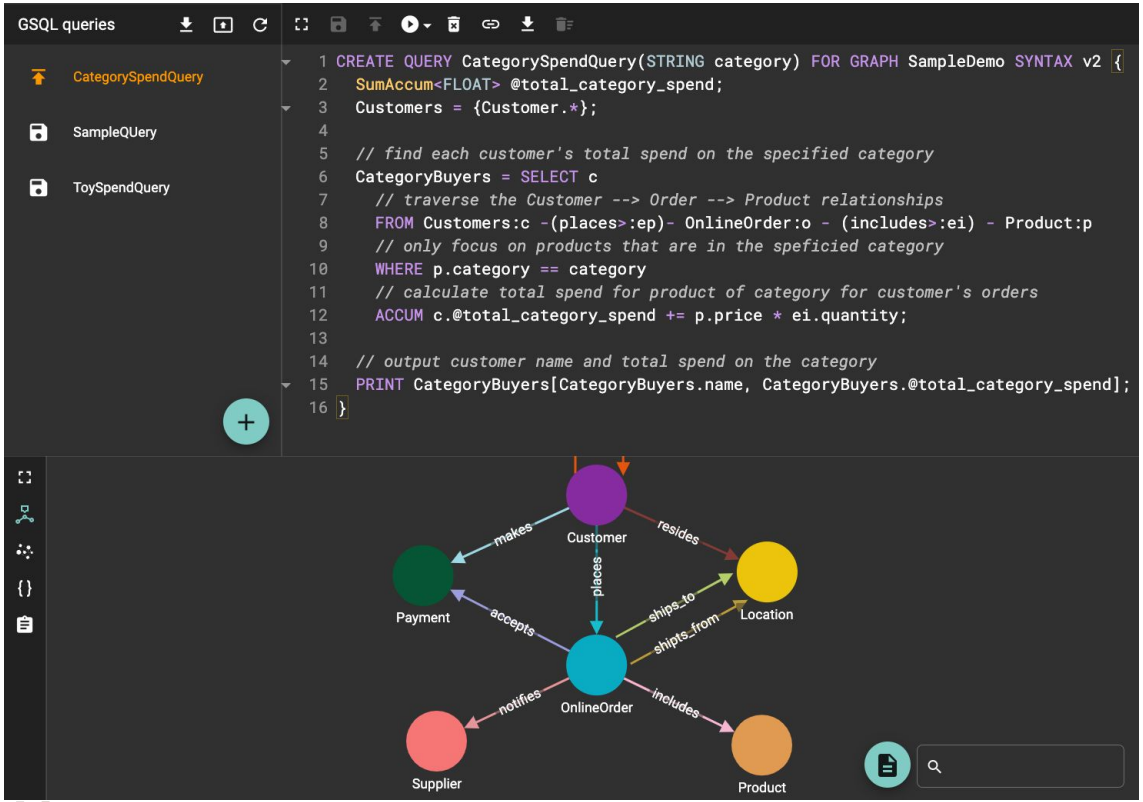

Control Flow (FOR, WHILE, IF/ELSE)

HTAP - Graph Updates & Analytics Queries

Built-in High Performance Parallelism

How? - Delightful Developer Experience

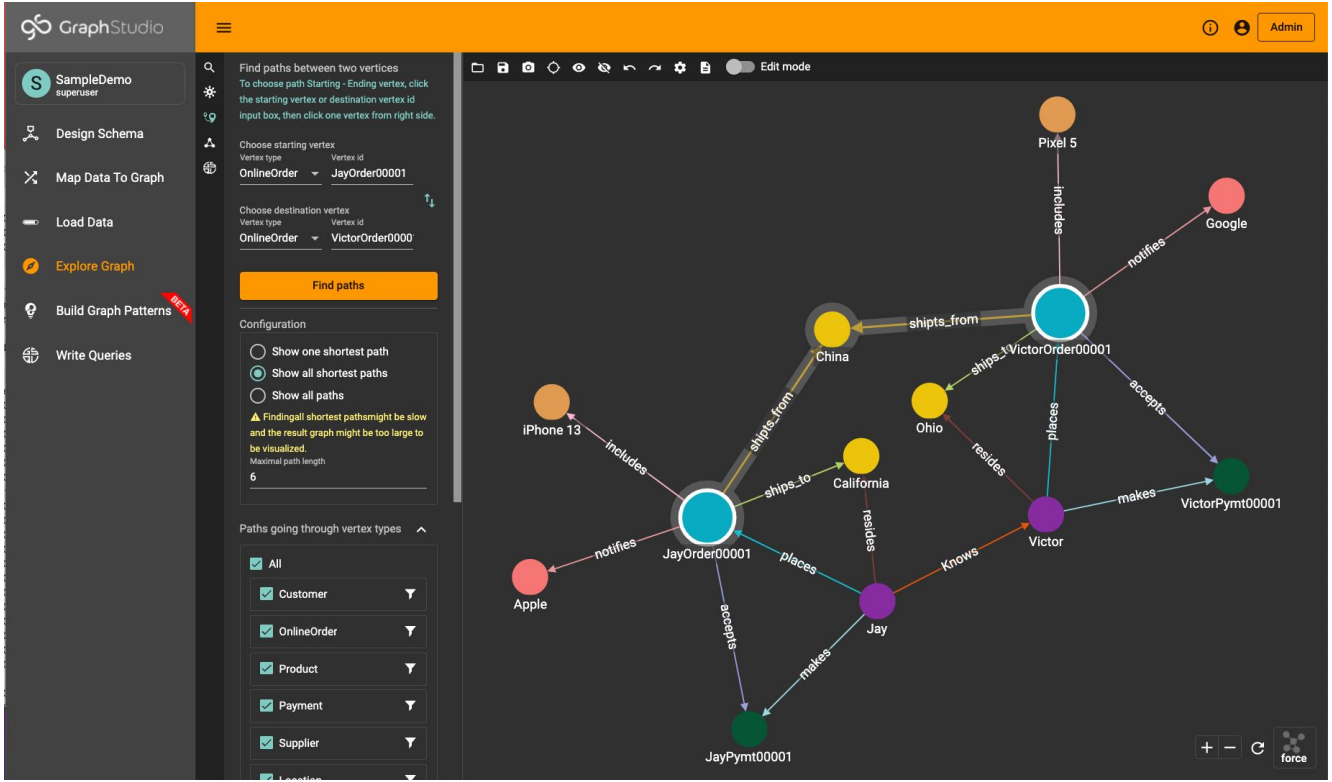

How? - In-Database ML / Graph Data Science

- **Battle-tested by Production Use Cases at Scale**

  Fraud Detection, Recommendation, Market Segmentation, Influence/Risk Scoring

- **Fast and Scalable**

  Native MPP design means faster execution

- **Familiar**

  Closely mirrors algorithmic pseudocode

- **Open-source**

  Free, accessible and modifiable

- **Comprehensive**

  Over 40 algorithms, in several categories

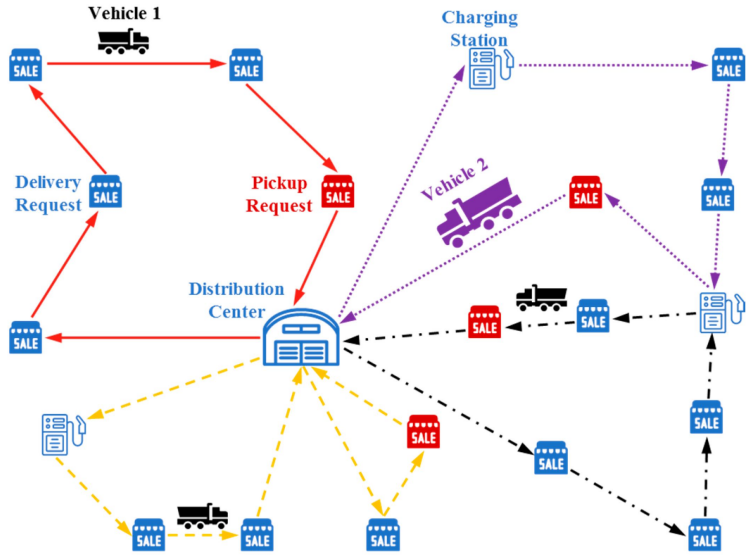

... to be covered in detail later by Dr. Lee

How? - Graph DB as a Service

Amazon Web Services

Microsoft Azure

Google Cloud

Also launchable from the Cloud Marketplaces

How? - Rapid Growing and Engaging Community

- Graph + AI Summit Attendees
  - 9,000+ registered, 300% YoY growth
- Thriving Developer Community
  - 340% YoY growth in weekly visits
- Active Open Source Contribution
  - GitHub projects 700% YoY growth / pyTigerGraph
- Engaged Developers at All Levels, including Students …

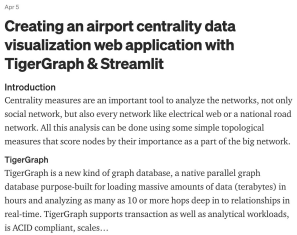
Apr 5

Creating an airport centrality data visualization web application with TigerGraph & Streamlit

Introduction

Centrality measures are an important tool to analyze the networks, not only social network, but also every network like electrical web or a national road network. All this analysis can be done using some simple topological measures that score nodes by their importance as a part of the big network.

TigerGraph

TigerGraph is a new kind of graph database, a native parallel graph database purpose-built for loading massive amounts of data (terabytes) in hours and analyzing as many as 10 or more hops deep in to relationships in real-time. TigerGraph supports transaction as well as analytical workloads, is ACID compliant, scales…

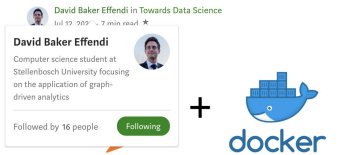

Efficient Use of TigerGraph and Docker

TigerGraph can be combined with Docker to run on nearly any OS…

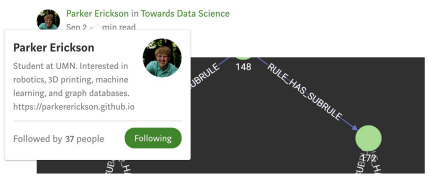

Decision Trees in TigerGraph

Make ML Classifications Inside Your Database

How? - Ecosystem of Solution and HW Partners

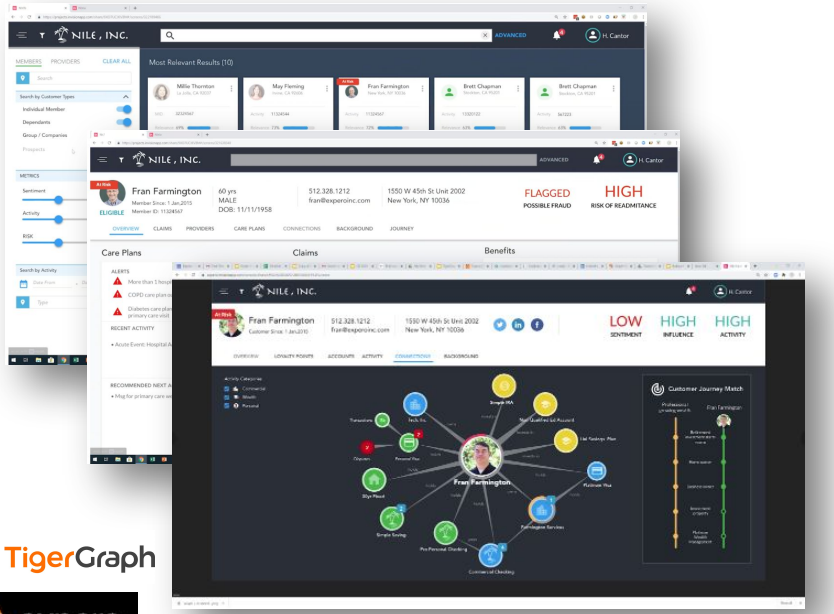

TigerGraph

expero

Turn-key Graph Solutions built on TigerGraph reducing time-to-market and dev cost by **2-3X +** !

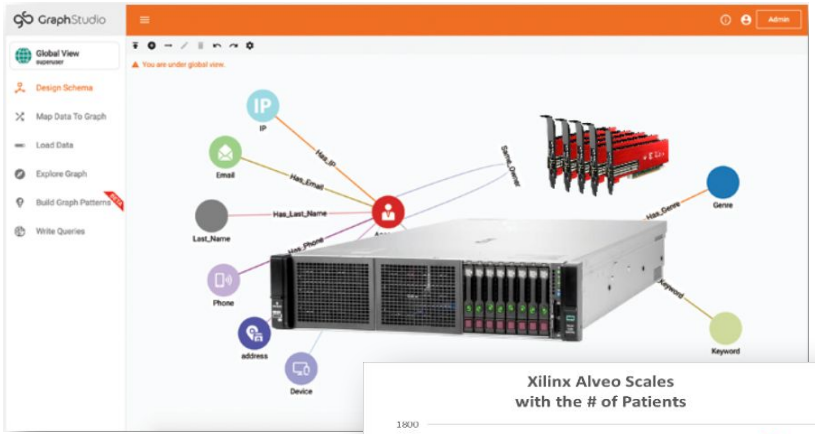

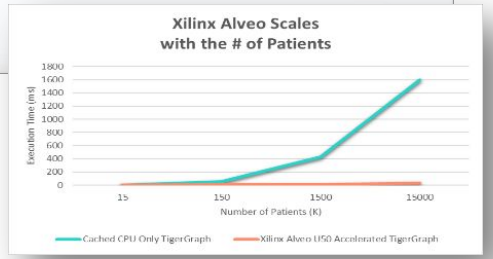

TigerGraph

Hewlett Packard Enterprise

XILINX

Deep HW Integration accelerating compute intensive graph algorithms by up to **20X !**

More to come ...

How? - Strong Foundation for Business Outcomes

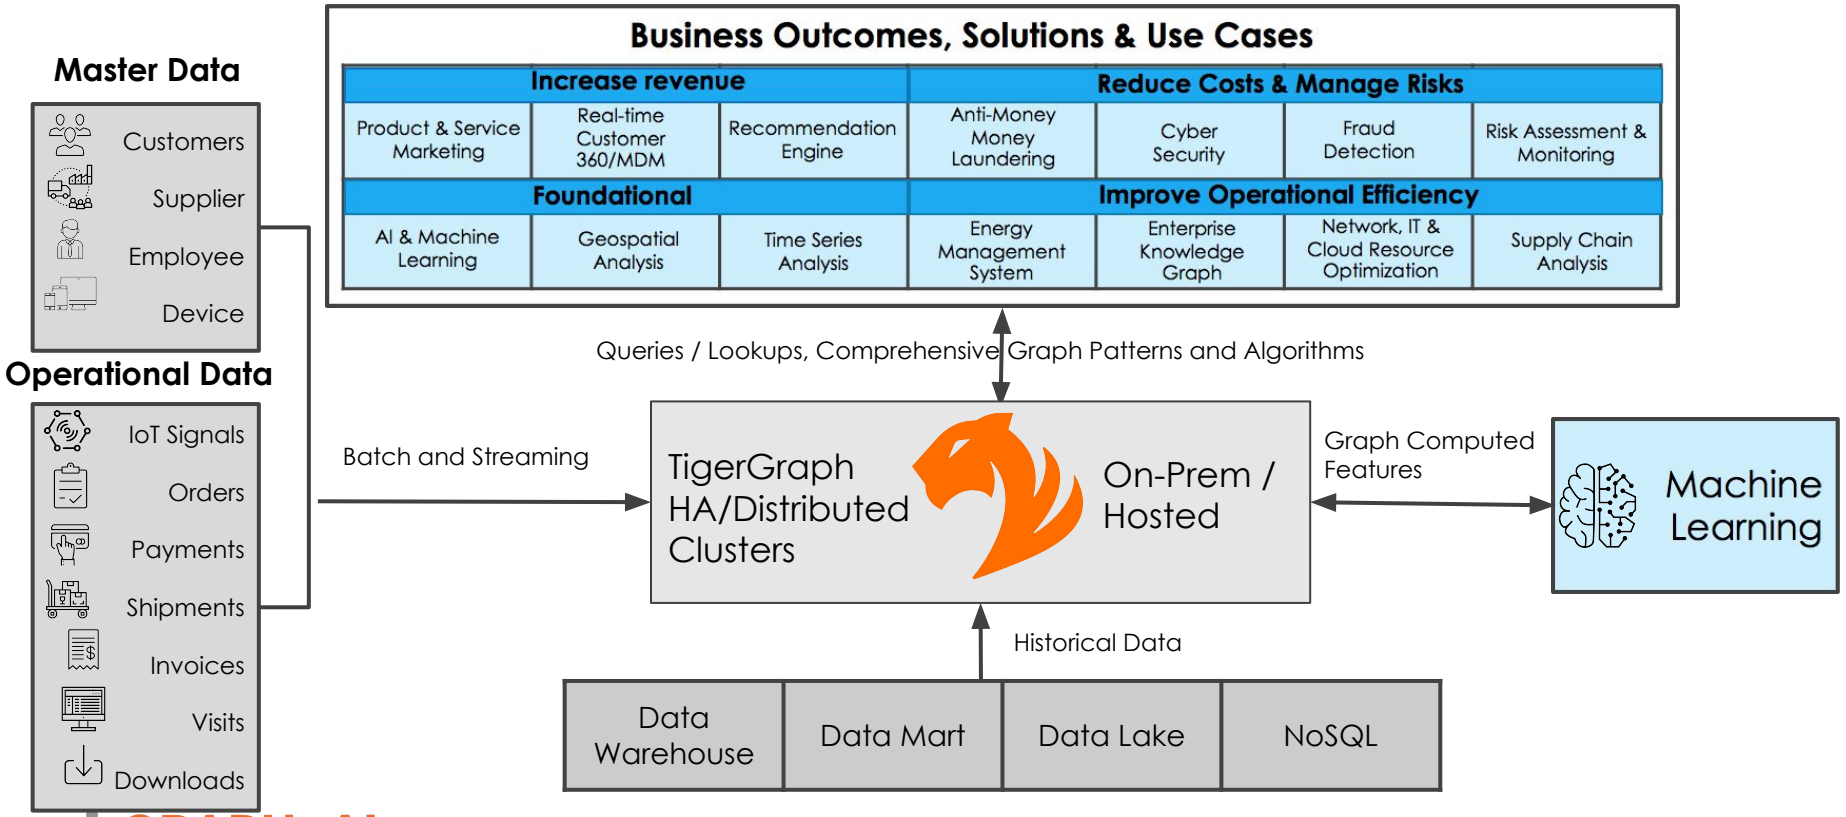

New Capabilities to Expand Adoption

Announcing TigerGraph Enterprise 3.2!

- Released on: 9/30/2021
- 7 months of hard work since 3.1
- 40+ New Features / Capabilities
- 400,000+ lines of code added / modified

3.2 Enterprise-grade Capability Highlights

Cross-Region Replication

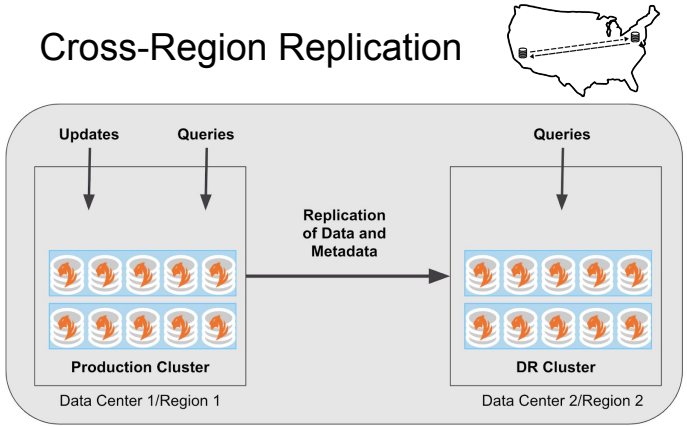

Full Kubernetes Support

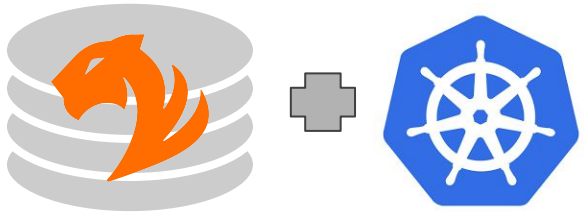

Query Workload Management

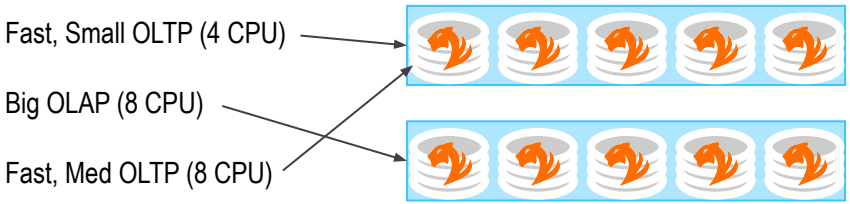

Role Based Access Control

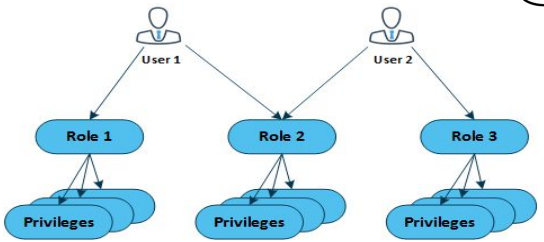

3.2 Query Language Capability Highlights

- **Running and Installing Queries**
  - Query argument flexibility: specify only what you care about
  - Query installation faster and more efficient: up to 3x - 10x improvement

- **Enhanced GSQL Capabilities**
  - Expanded library of built-in functions: 29 new string / math functions
  - Relaxed variable declaration rules: declare anywhere before you use
  - Dynamic subquery result: allow anonymous tuples in response
  - Function overloading: support object-oriented concepts

3.2 Developer Tool Capability Highlights

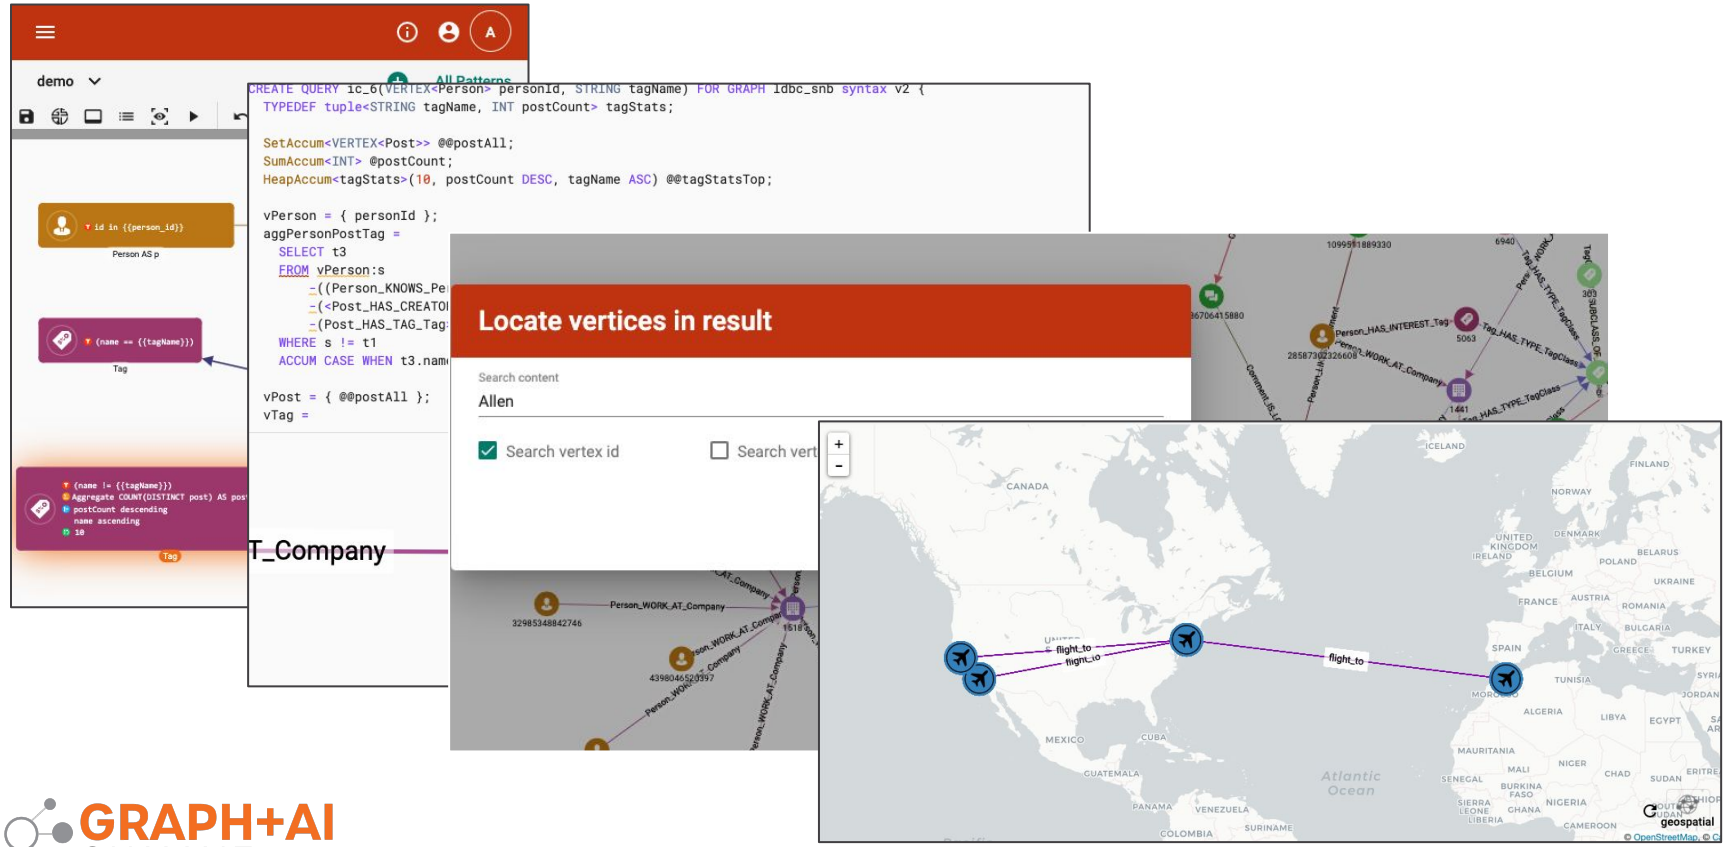

3.2 Scalability: New Record in LDBC Benchmark

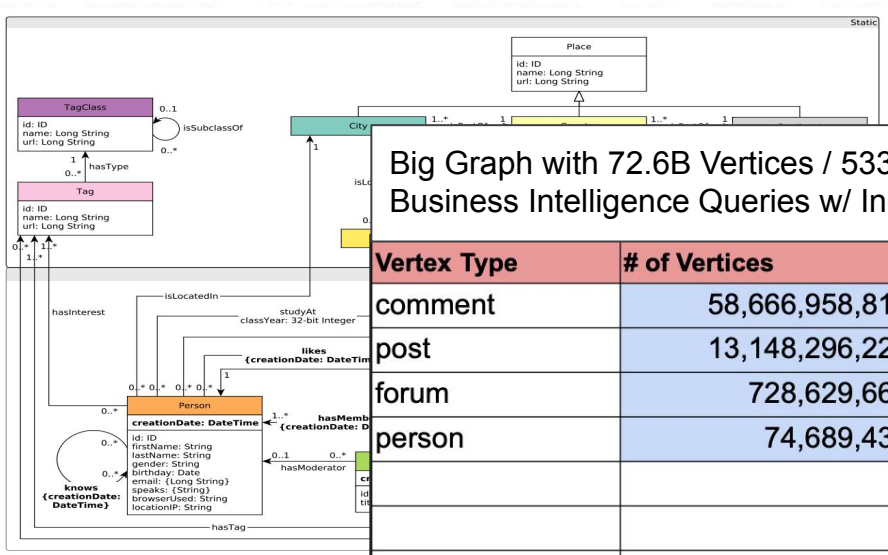

Big Graph with 72.6B Vertices / 533.5B Edges / 36.1TB
Business Intelligence Queries w/ Incremental Batch Updates

Vertex Type	# of Vertices	Vertex Type	# of vertices
comment	58,666,958,815	company	1,575
post	13,148,296,221	university	6,380
forum	728,629,666	city	1,343
person	74,689,437	country	111
		continent	6
		tag	16,080
		tagClass	71

TigerGraph: ***the first and only vendor*** completed LDBC SNB-SF30K BI at 36TB!

Innovation to Accelerate Adoption

A Glimpse into Work-in-Progress Innovations

- What if we can get 100,000+ application / service developers to adopt graph easily ?
- What if we can get business users to get deep graph insights without the need to code ?
- What if we can dramatically reduce the time and effort to develop a complete e2e graph solution ?

Demo: GraphQL Query API on TigerGraph

For Developers:

Renchu Song

Sr. Dev Manager, Tools and Applications

TigerGraph

100,000+ developers can use **familiar and state-of-art GraphQL Data API** language to query data directly in TigerGraph !

Demo: No-Code Visual Query Builder

For Domain Experts:

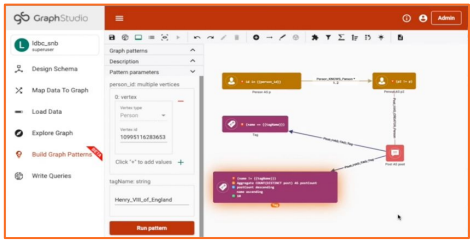

Renchu Song
Sr. Dev Manager, Tools and Applications
TigerGraph

Domain Experts / Data Analysts / Data Scientists … can **visually describe their queries** without the need to learn and code GSQL!

Demo: Graph Solutions and Solution Accelerator

For Business Leaders on e2e Development:

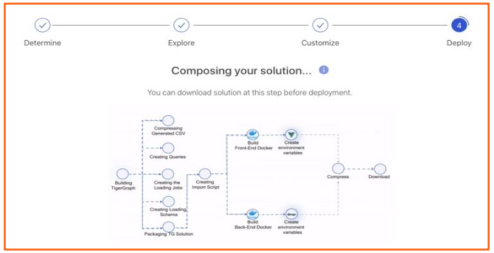

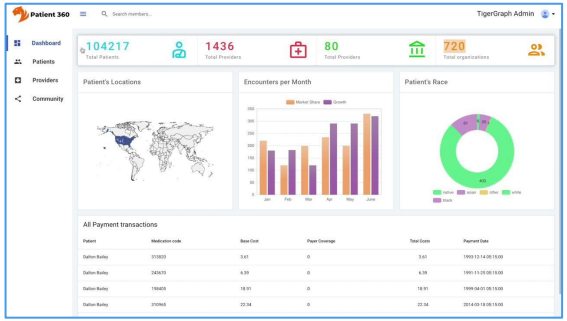

Jon Herke
Sr. Dev Manager, Dev Relations
TigerGraph

Renchu Song
Sr. Dev Manager, Tools and Applications
TigerGraph

Accelerate end-to-end Graph Solution / App Building and Deployment:
3X+ Reduction in Development Cost, and **4X+ Speedup** for Time-to-value !

Graph + ML Advancement

Why Graph + ML?

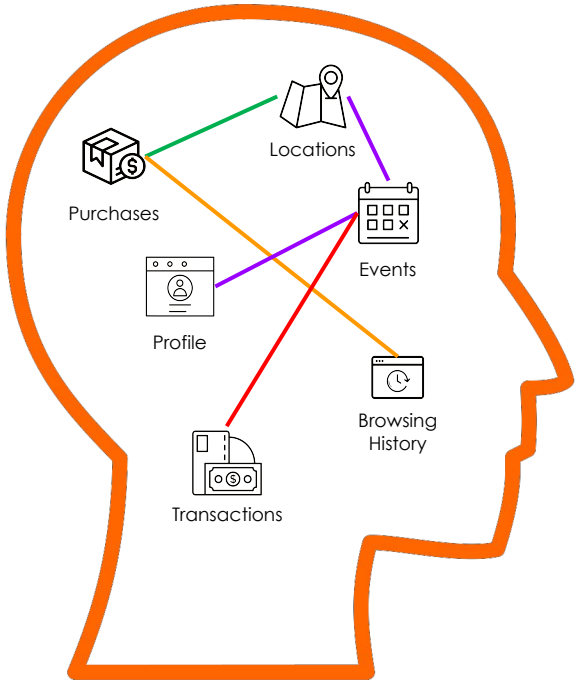

- Natural Merging of Data + Learning: Relationships are Fundamental
- Meaningful Patterns
- Enhances Conventional ML/AI
- Explainable Results

TigerGraph ML Today

In-Database Graph Algorithms

- 40+ algorithms for similarity, clustering, best paths, community
- No export needed. Live, updatable data
- Open-source, customizable
- Scaleable, Ultra-fast MPP graph engine
- Toolkit, building blocks for problem-solving

Graph Features Enriching ML Quality

- Graph enriches ML training data: quality in ⇒ quality out
- Use graph algorithms, pattern-matching queries, or graph embedding

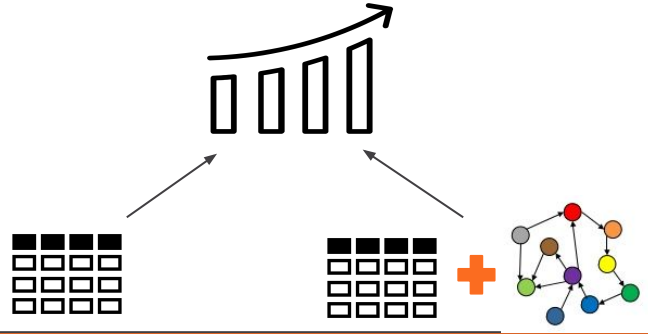

Solving Enterprise-Scale Business Problems

Today: Industry-leading Business Value

Healthcare:
Real-time recommendations

- 1.3TB graph brain
- Real-time care recommendations
- Improving healthcare, lowering cost

Industrial Supply Chain:
Analytics for decisions

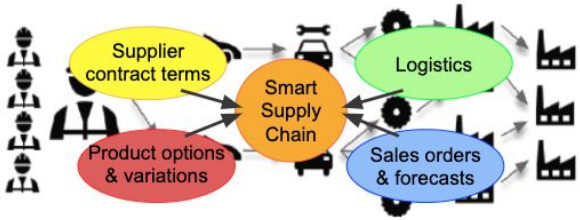

- Analytics: weeks→ minutes
- Reveal opportunities, optimize tactical & strategic decisions
- Saving $100M+/yr

Financial Services:
Real-time fraud detection

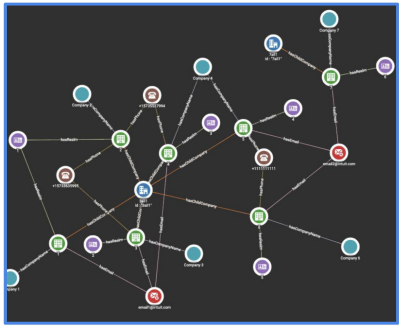

- Integrates multiple tools
- "Magical" real-time visual results for investigators
- Scalable for growth

Looking Forward: Enterprise Graph ML

- In-Database Graph Data Science Library

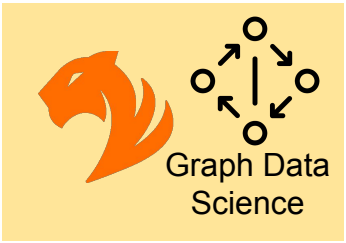

- In-Database Machine Learning

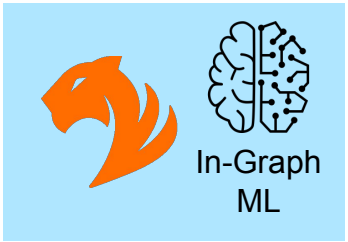

- Graph+ML Workbench

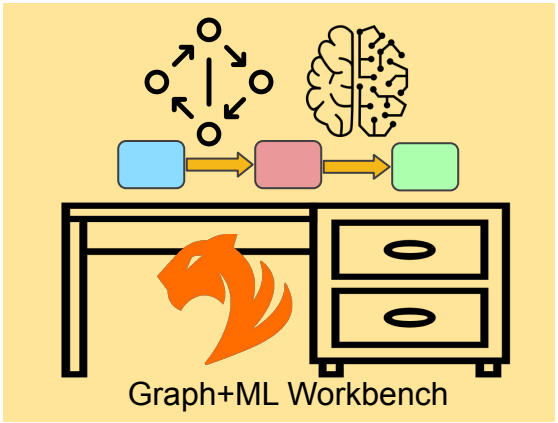

- **TigerGraph's Scale & Performance**

TigerGraph In-Database *Graph Data Science Library*

Signals our commitment to serving the needs of data scientists

- More algorithms (15 released this week)
  - Graph Embedding
  - Link prediction
  - Similarity
  - Centrality
- More than just algorithms

In-Database
- No export needed
- Live, updatable data
- Scaleable, Ultra-fast engine
- GSQL query language

For Data Scientists
- easier and faster to run
- include ML, such as graph embeddings
- will integrate with feature & model management
- will integrate with Graph+ML Workbench

Graph Neural Networks

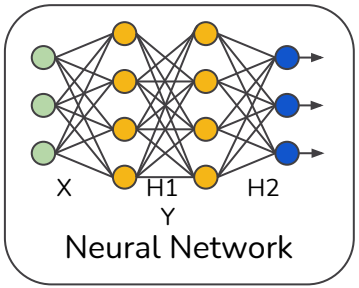

Powerful Machine Learning to Predict and Classify

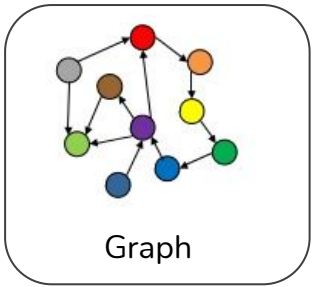

Insight from Connected Data

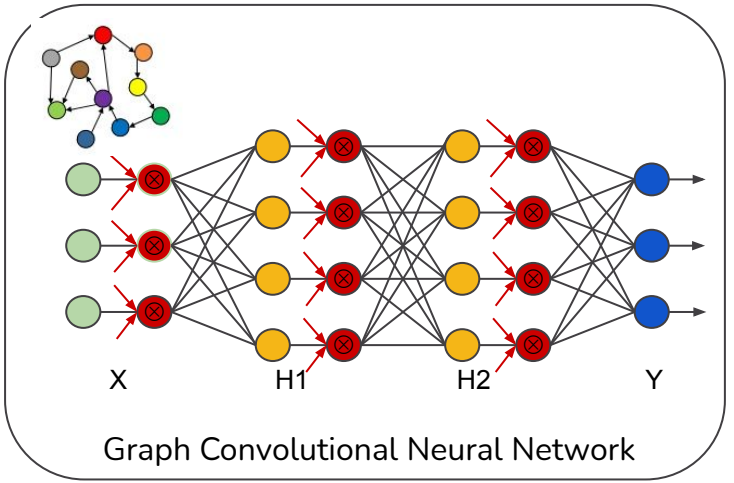

- Combines the added insight from connected data with the modeling power of neural networks
- Uses the graph during training; in-database training is the ideal.

Graph + ML Workbench

- The added power of graph
- The same convenience of conventional MLOps

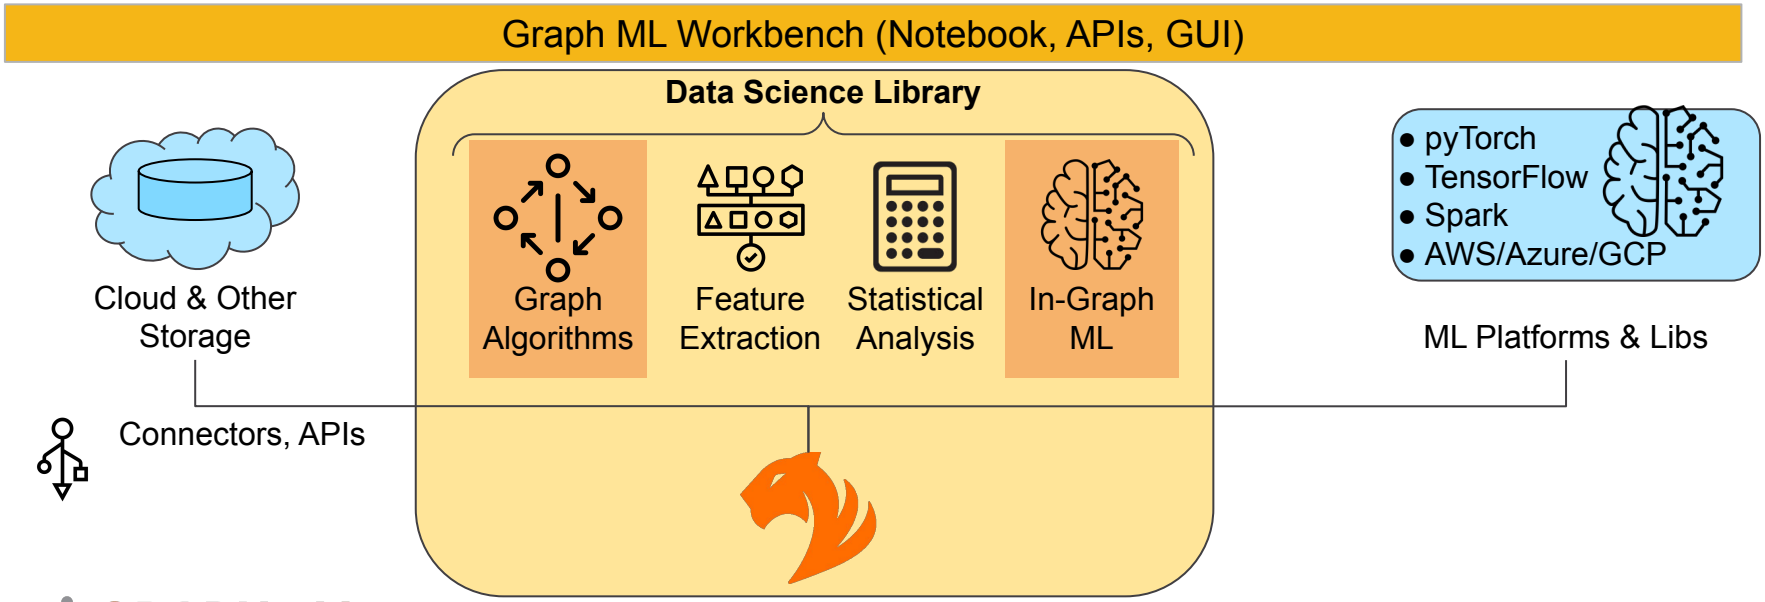

Summary

Key Takeaways

Graph technology is **mature, scalable and available** for you to adopt …

- Graph is **beautiful** - the most natural way to model all data
- Graph is **valuable** - backed by many success stories from our customers
- Graph is **real** - world's leading scalable graph technology
- Graph is **easy** - delightful developer and business user experience
- Graph is **quick** - rapid development by turn-key solutions and toolkits
- Graph is **peaceful** - operation delegated to cloud
- Graph is **intelligent** - AI/ML made easy by in-db graph data science libraries

Any excuses for not to graph now ?

TigerGraph

Question and Answer

- Try TigerGraph Cloud (tgcloud.io) - Free to Start
- Download TigerGraph's Developer Edition
- Take a Test Drive - Online Demo
- Get TigerGraph Certification
- Join the Community


@TigerGraphDB

/tigergraph

/TigerGraphDB

/company/TigerGraph