



# Getting started with TigerGraph on GCP

The Google Cloud logo, with "Google" in its multi-colored font and "Cloud" in grey. A blue horizontal line is positioned below the "Google" text.



# Speaker

Google Cloud



Bruce Hsu  
Google - Customer Engineer

TigerGraph



Shuo Yang  
TigerGraph - Head of Cloud

# Agenda

- Google Cloud Infrastructure Overview
- Google Cloud for Database Hosting
- TigerGraph, a Scale Out Graph Analysis Engine
- Demo - Launch an Instance to Get Started

# Google Cloud Infrastructure Overview

# GCP Infrastructure Numbers

**28** cloud regions (+10 announced regions\*) [url](#)

**85** zones [url](#)

**146** edge locations [url](#)

**100+** CDN locations [url](#)

**100+** Dedicated Interconnect locations [url](#)

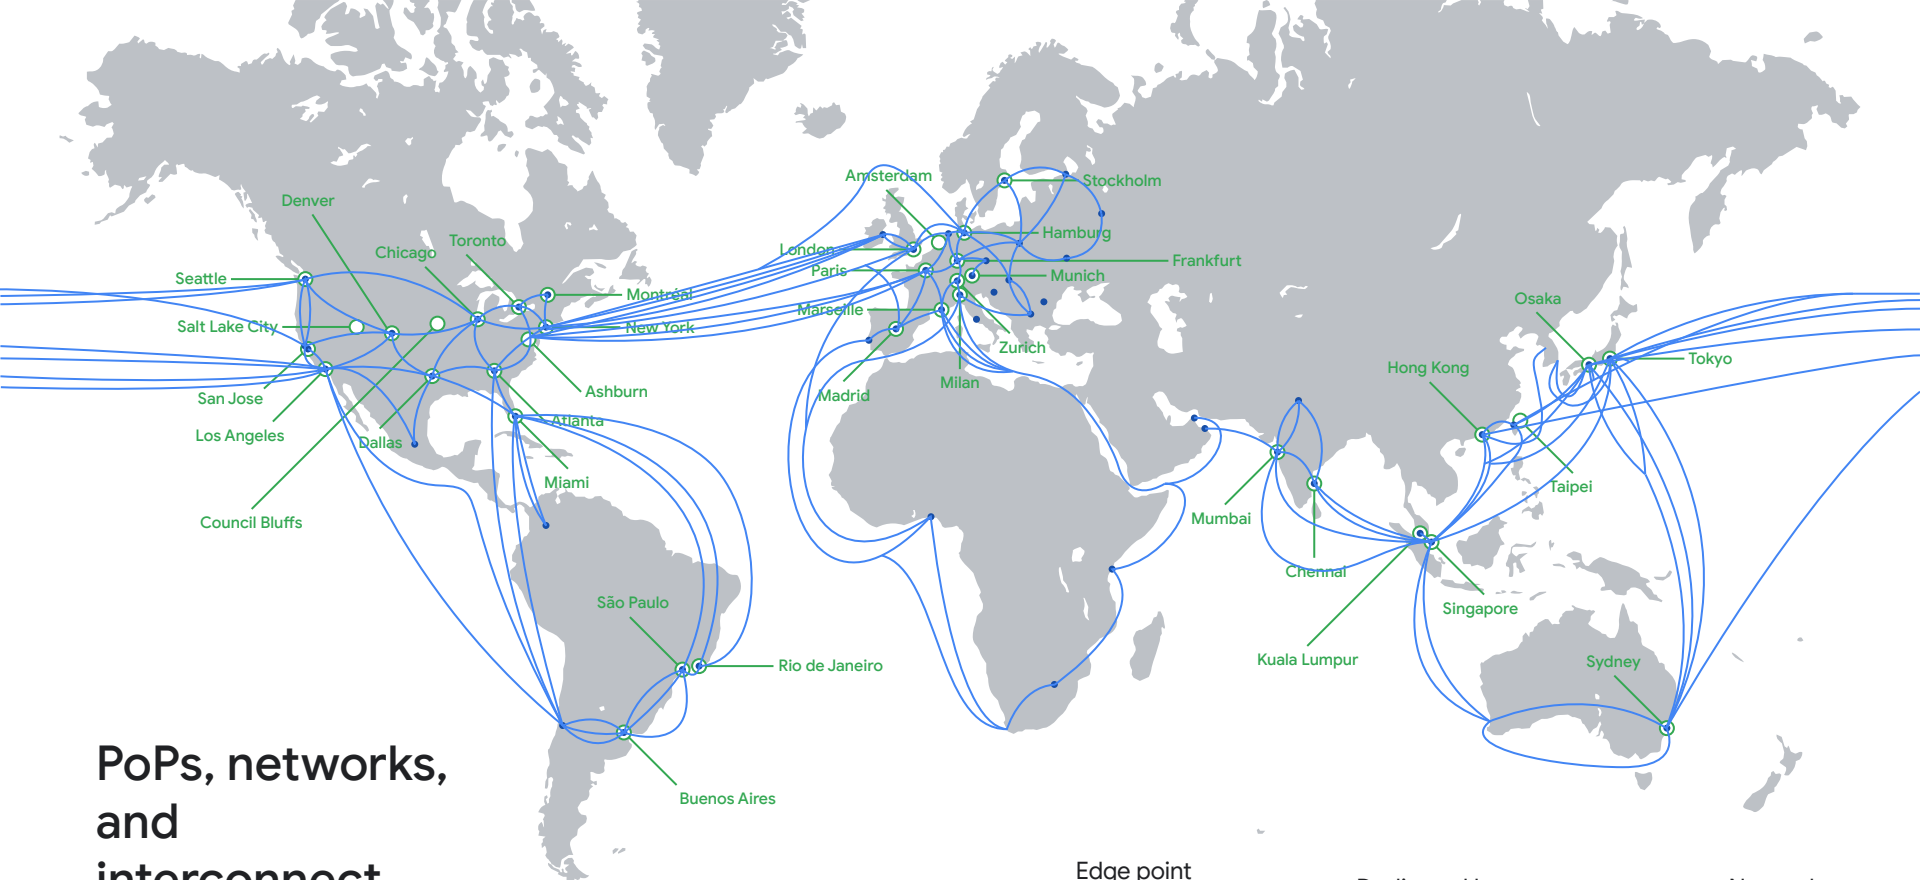
**20+** Google data centers [url](#)

**50+** renewable energy project locations [url](#)

**18** subsea cable investments [url](#)

In **2021** we've launched **4** new regions.





# PoPs, networks, and interconnect

- Edge point of presence
- Dedicated Interconnect
- Network

# Nine cloud products with 5 Billion users each





# GCP Portfolio + Partner Solution




| [GRAPHAISUMMIT.COM](https://graphaisummit.com) | [#GRAPHAISUMMIT](https://twitter.com/graphaisummit)

Smart Business  
Analytics and AI




Data  
Management



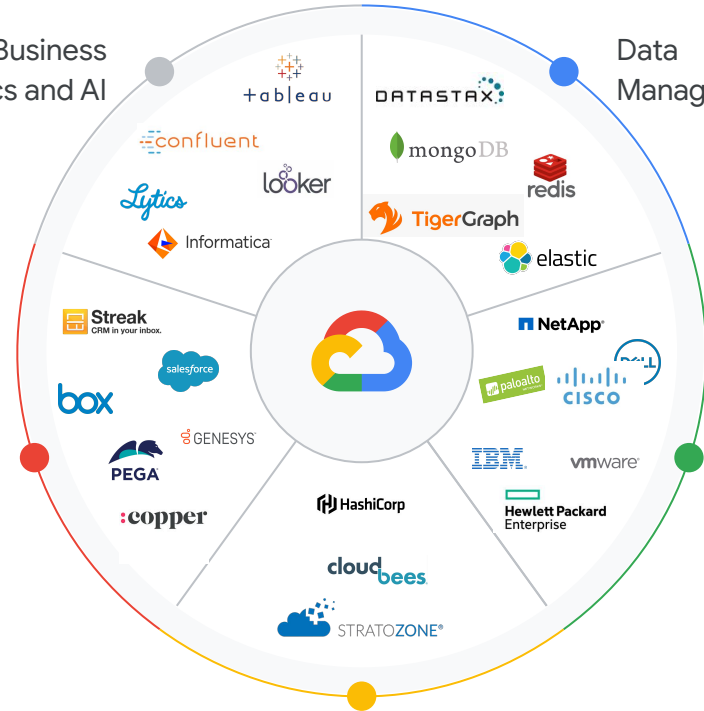
Productivity  
& Collaboration



Infrastructure  
Modernization



Application  
Development



# Google Cloud for Database Hosting

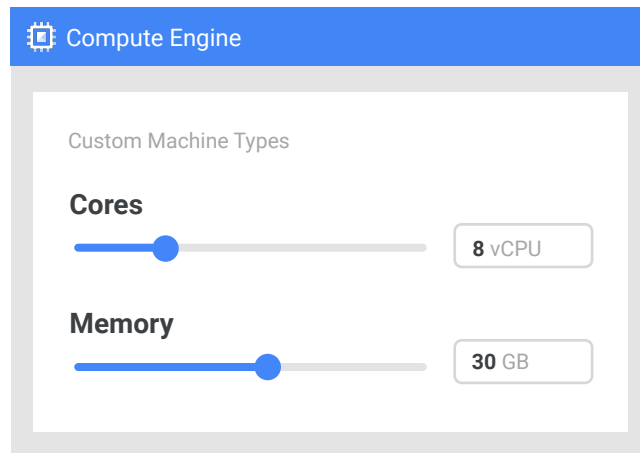
# VM Families

General Purpose			Workload-Optimized		
<ul style="list-style-type: none"> <li>• Web Serving</li> <li>• Steady-state LOB apps</li> <li>• Dev &amp; Test environments</li> </ul>	<ul style="list-style-type: none"> <li>• Enterprise apps</li> <li>• Medium databases</li> <li>• Web &amp; App Serving</li> </ul>	<ul style="list-style-type: none"> <li>• Scale-out workloads</li> <li>• Wide variety of workloads used in cloud-native deployments</li> </ul>	<ul style="list-style-type: none"> <li>• HPC (Compute)               <ul style="list-style-type: none"> <li>• EDA</li> </ul> </li> <li>• High Performance Web Serving</li> <li>• Ad Serving</li> <li>• Gaming</li> <li>• Media Transcode</li> <li>• AI/ML</li> </ul>	<ul style="list-style-type: none"> <li>• SAP HANA</li> <li>• Largest in memory DBs</li> <li>• Real-time data analytics</li> <li>• In-memory cache</li> </ul>	<ul style="list-style-type: none"> <li>• HPC (GPU)</li> <li>• ML</li> <li>• Massive parallelized computation</li> </ul>
Cost savings a priority	Flexible, Complete Feature Set	Best Perf/\$ for scale out workloads	Highest performance CPUs	Most memory on Compute Engine	Highest performance GPUs
Cost Optimized (E2)	Balanced (N2, N2D)	'Tau' Scale-out Optimized (T2D*)	Compute-Optimized (C2, C2D*)	Memory-Optimized (M2)	Accelerator-Optimized (A2)

# Custom Machine Type

## Configure your own machine types, with average 19% savings

- Create a machine type with 1 vCPU and up to 96 vCPU
  - Or any even number of vCPUs in between
- Use Compute Engine stop/start feature
  - Move your workload to a smaller or larger Custom Machine Type instance or predefined configuration
- Priced by the resources they use (i.e. vCPU, memory)



Predictably calculate the cost of Custom Machine Type shapes in any configuration.

# Right Sizing Recommendations

- Get insights into your usage
- Compute Engine recommends the exact fit
- Optimize your cloud spend
- Avoid paying for idle and oversized resources

**Resize instance**

This instance has had low CPU and memory utilization recently. You can save an estimated \$17 per month by switching to the machine type: g1-small (1 vCPU, 1.7 GB memory). [Learn more](#)

**Current machine type**  
n1-standard-1 (1 vCPU, 3.75 GB memory)

**New machine type**  
g1-small (1 vCPU, 1.7 GB memory) Recommended [Customize](#)

**⚠ This instance is running. Changing the machine type will restart the VM.**

[Apply](#) [Dismiss recommendation](#) [Cancel](#)

# Pricing Advantage

## Sustained Use Discounts (SUD):

- No upfront commitments
- Automatic Discounting
- Up to 30% discount

## Committed Use Discounts (CUD):

- Purchase Aggregate Cores & RAM per region
- Flexible and Easy to manage
- Up to 70% discount

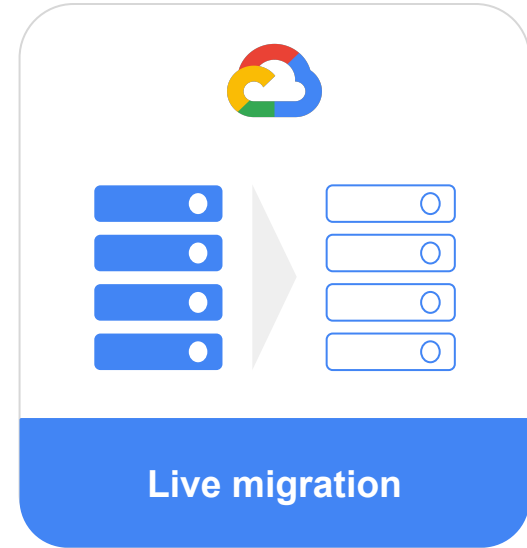
## Rightsizing and Custom Machine Types:

- Deploy custom instances for your needs
- Recommendations based on utilization
- Average 20% savings (up to 40%) .

# Live Migration

## Manage your infrastructure while business continues as usual

- VM instances keep running when a host event occurs
  - Software or hardware update
  - Switch to new configurations
- Live migrates your running instances to another host in the same zone rather than requiring your VMs to be rebooted
- Perform maintenance to keep your infrastructure protected and reliable without interrupting any of your VMs.



# Storage Options



## Google Cloud Storage

Exabyte-scale, feature-rich object storage  
Automatically scaling throughput



## Persistent Disk

High-performance, replicated block storage (HDD/SSD)  
Up to 64TB per PD, Up to 1.2 GB/s per instance



## Local Storage

Local SSD (NVMe), physically attached via PCIe  
Up to 3TB per Local SSD, Up to 2.65 GB/s per instance



## Cloud Filestore

Highly available, durable, POSIX-compliant shared storage  
across tens of thousands of nodes



## Partner, hybrid, and open-source

Storage solutions for NetApp, Elastifile, DDN, Lustre, and more





# Encryption by default at Google



At rest



In transit

# Compliance, Regulations & Certifications

Global	Americas		Europe, Middle East & Africa		Asia Pacific	
ISO 27001	<b>USA</b>	CFTC Rule 1.31(c)-(d)	<b>Europe</b>	<b>Spain</b>	<b>Australia</b>	<b>Singapore</b>
ISO 27017	HIPAA	FINRA Rule 4511(c)	GDPR	Esquema Nacional de Seguridad	Australian Privacy Principles	MTCS Tier 3
ISO 27018	HiTrust	_____	EU Model Contract Clauses	_____	Australian Prudential Regulatory Authority Standards	
SOC 1	FedRAMP	<b>Canada</b>	Privacy Shield	<b>UK</b>	IRAP	
SOC 2	FIPS 140-2	Personal Information & Electronic Documents Act	TISAX	NCSC Cloud Security Principles	_____	
SOC 3	COPPA	_____	_____	NHS IG Toolkit	<b>Japan</b>	
PCI DSS	FERPA	<b>Argentina</b>	<b>Germany</b>	_____	FISC	
CSA STAR	NIST 800-53	Personal Data Protection Law	BSI C5	_____	My Number Act	
MCAA	NIST 800-171		<b>South Africa</b>			
Independent Security Evaluators Audit	Sarbanes-Oxley		POPI			
	SEC Rule 17a-4(f)					

<https://cloud.google.com/security/compliance>

# Unlock More Data Insight



Google  
Big Query



TigerGraph



vertex.ai

**Session:** Transform with Vertex AI



Benazir Fateh  
AI/ML Specialist



Same planet-scale core platform that powers all internal and Google Cloud observability

- 1.5 quadrillion metric points in memory
- 13 million queries per second
- >20 quadrillion metric points on disk
- 2.5 EB / month logs



Built-in with zero config usage for all Google Cloud services



Powered by Google's SRE principles

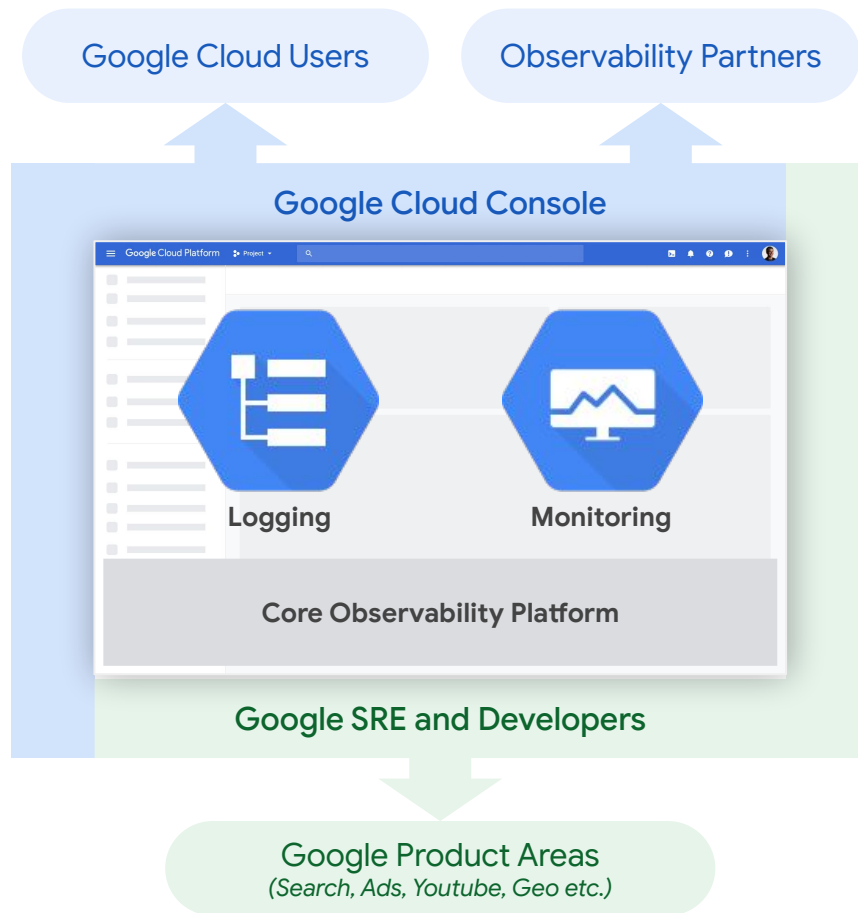


Defacto platform logging + metrics provider

- Powers partner services like Dynatrace, Datadog, NewRelic, etc.




Day-zero observability support for new Google cloud services



# TigerGraph, a Scale Out Graph Analysis Engine

# Distributed Native Graph Storage

Data of different components are split into segments.

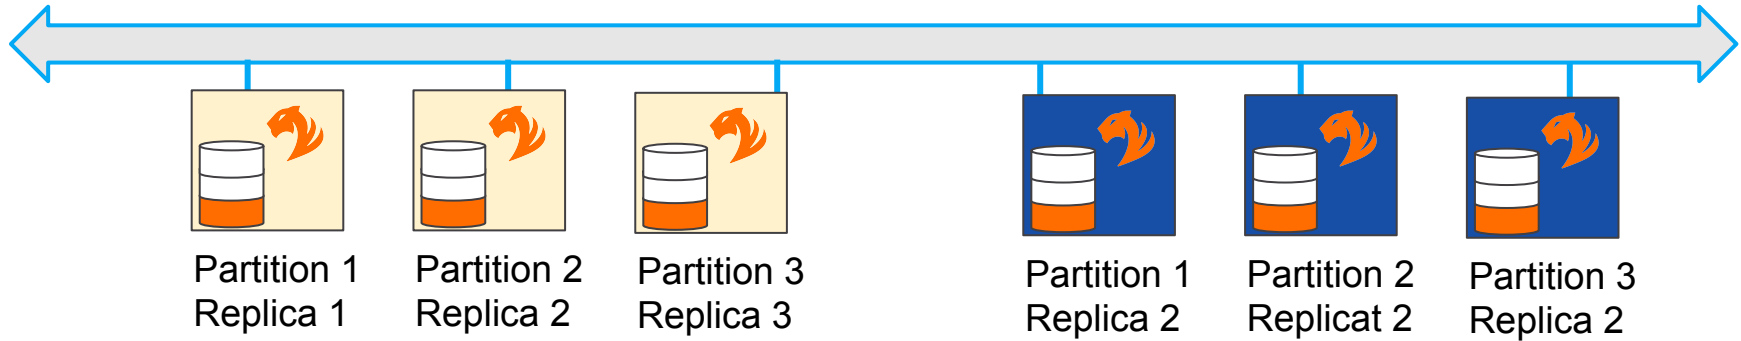



The segments are stored distributedly across the cluster.


The segments of different components with same ID stores data for the same set of vertices under the same vertex type.

The location of a vertex can be calculated based on its internal ID

# Partition and Replication



# Query - Single Server mode

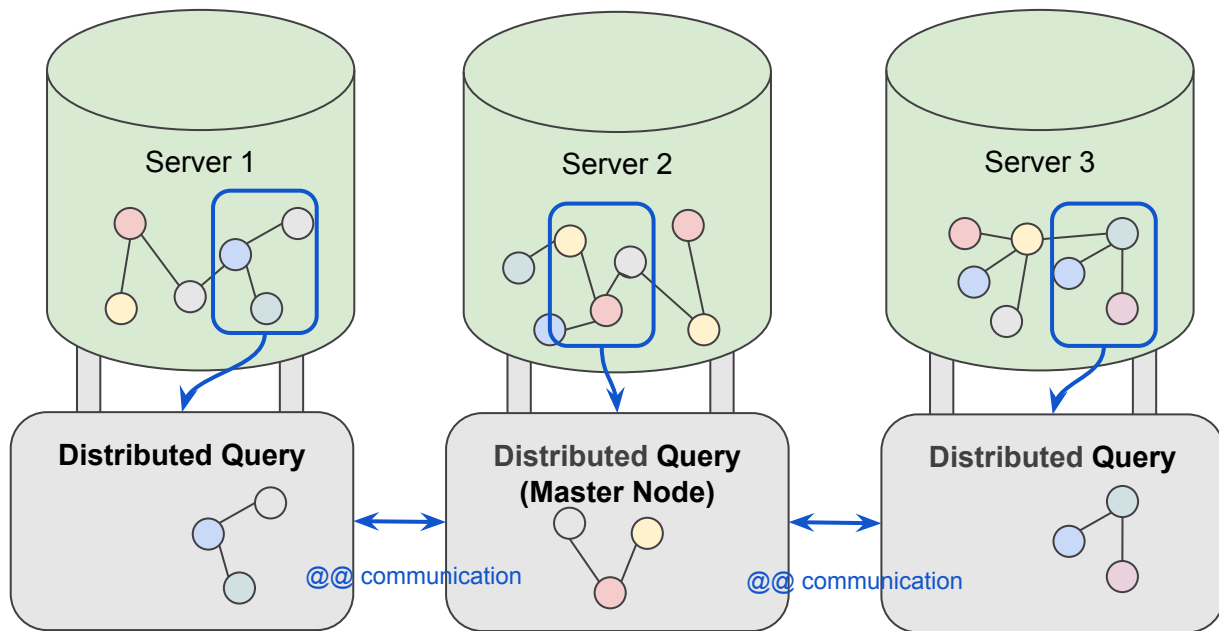


## Single Server Mode

- The cluster elects one server to be master for that query.
- All query computation takes place on query master.
- Vertex and edge data are copied to the query master as needed.
- **Best for queries with one or a few starting vertices.**
- **If your query starts from all vertices, don't use this mode.**



# Query - Distributed mode



## Distributed Mode

- The server that received the query becomes the master.
- Computation executes on **all** servers in parallel.
- Accumulators are transferred across the cluster.
- **If your query starts from all or most vertices, use this mode.**

# Benchmark: LDBC-SNB Schema (twitter like)

- 70B vertices, 533B edges
  - Number of machines: 40
  - Machine: m1-ultramem-40

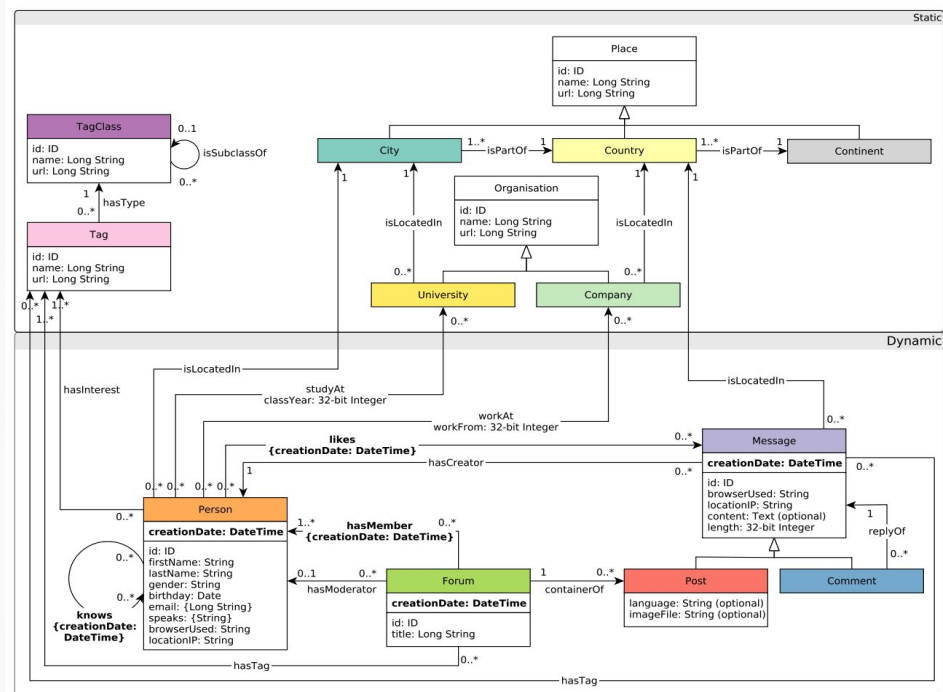


Figure 3.1: UML class diagram-style depiction of the LDBC SNB graph schema. Note that the knows edges should be treated as undirected (but are serialized only in a single direction). The cardinality of the hasModerator edge has changed between version 0.3.x (where it was exactly 1) and version 0.4.x (where it is 0..1).

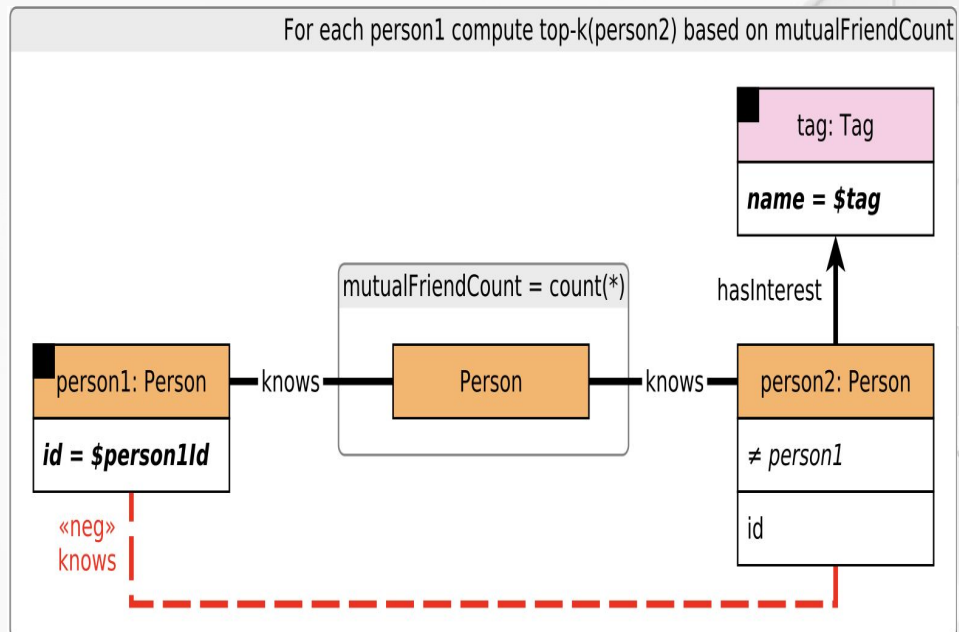
# Benchmark: LDBC-SNB Schema (twitter like)

Recommend new friends to a specific person

- Who is associated to in a specific tag (e.g., living in a specific country)
- Rank by the number of mutual friends

Query time: 7 seconds

[More details to be published here](#)



[LDBC-SNB: Friends recommendation](#)

# Demo - Get started with TigerGraph Cloud

Please try TigerGraph cloud [www.tgcloud.io](http://www.tgcloud.io)

Or visit from [Google Marketplace](#)