EYWA: Building a Distributed Graph Engine on Huawei Cloud

Yinglong Xia
Huawei Cloud America
09/01/2017

The 10th LDBC TUC in Conjunction to VLDB 2017, Munich, Germany

About Huawei

170,000
employees

170
countries

$75B
revenue

14
regional HQ

16
R&D centers

32%
YoY growth

HUAWEI

https://mp.weixin.qq.com/s/RSNfQqSmUYYTB3jTLVkfbw

About Huawei - 2

8/28/2017


华为技术有限公司

公 司 文 件

华为司发组织变动【2017】303号　　签发人：任正非

关于Cloud BU组织变动的通知

公司各部门：

为支撑业务发展，经总裁办讨论，同意Cloud BU组织结构调整方案，现发布如下。

一、本次组织结构调整主要内容

1、Cloud BU迁移至华为集团下，作为一层组织。

2、保持基础设施工程部的组织定位及职责。

二、优化后的组织结构图

- Cloud BU在公司一层组织中的汇报关系

运营商 BG

企业 BG

消费者 BG

Cloud BU

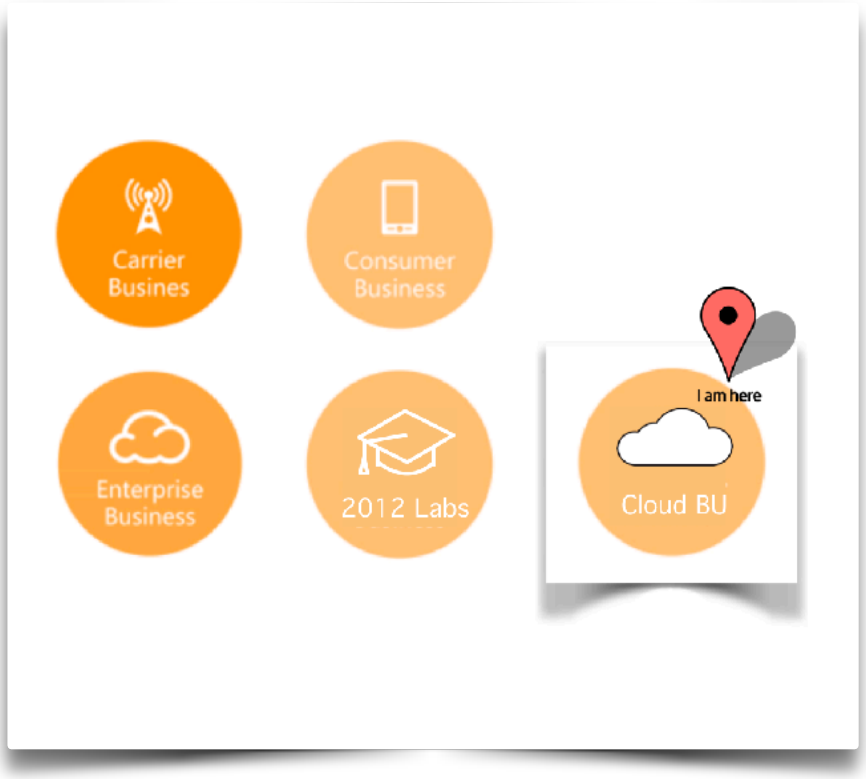

HUAWEI

Eywa

The Na'vi believe that Eywa acts to keep the ecosystem of Pandora in perfect equilibrium. Eywa emphasizes connection, the connections among the creatures on Pandora and the connections between Na'vi and their ancestors.

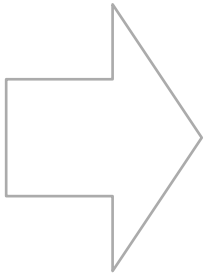

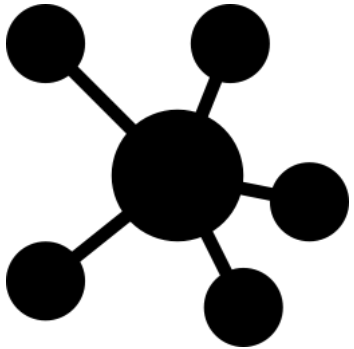

Edge-set sYnergistic Workload Architecture (EYWA) as the in-house graph engine techniques in Huawei

Connection is the essence of graph

Graphs in AI: Intelligent Service

Use Case:

AI assists Technical Services

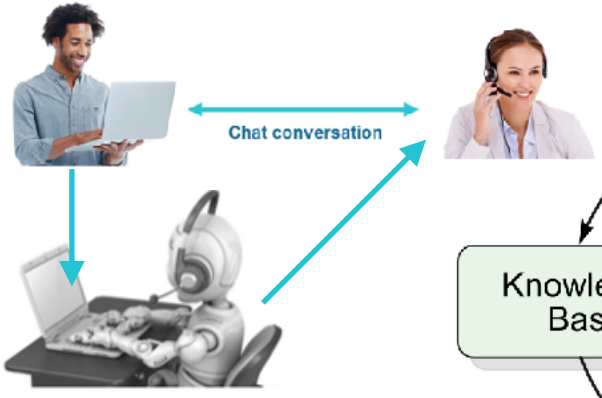

Private Knowledge Graph

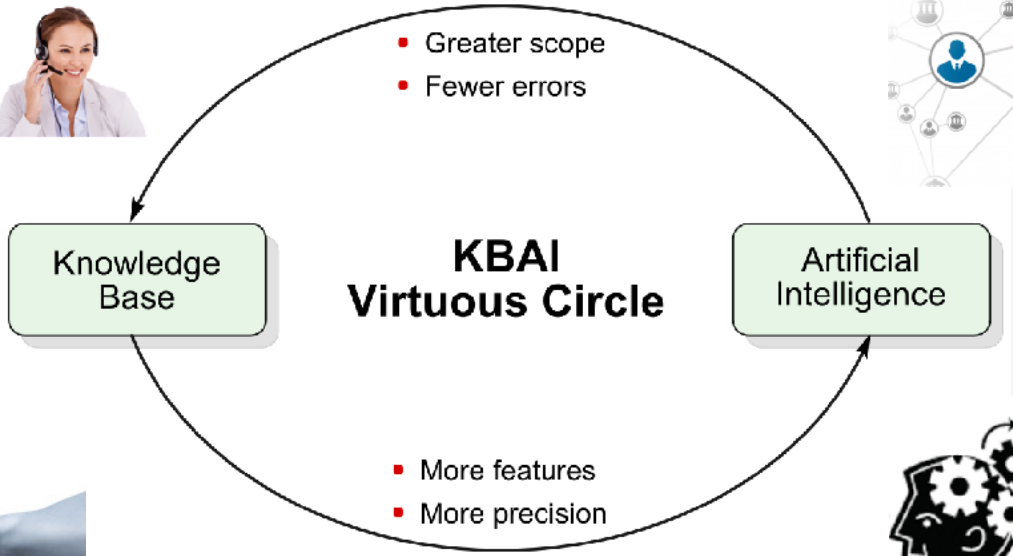

Human-Machine Collaboration

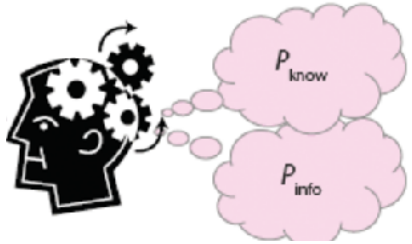

Inference & Reasoning

Eywa Graph Engine Service

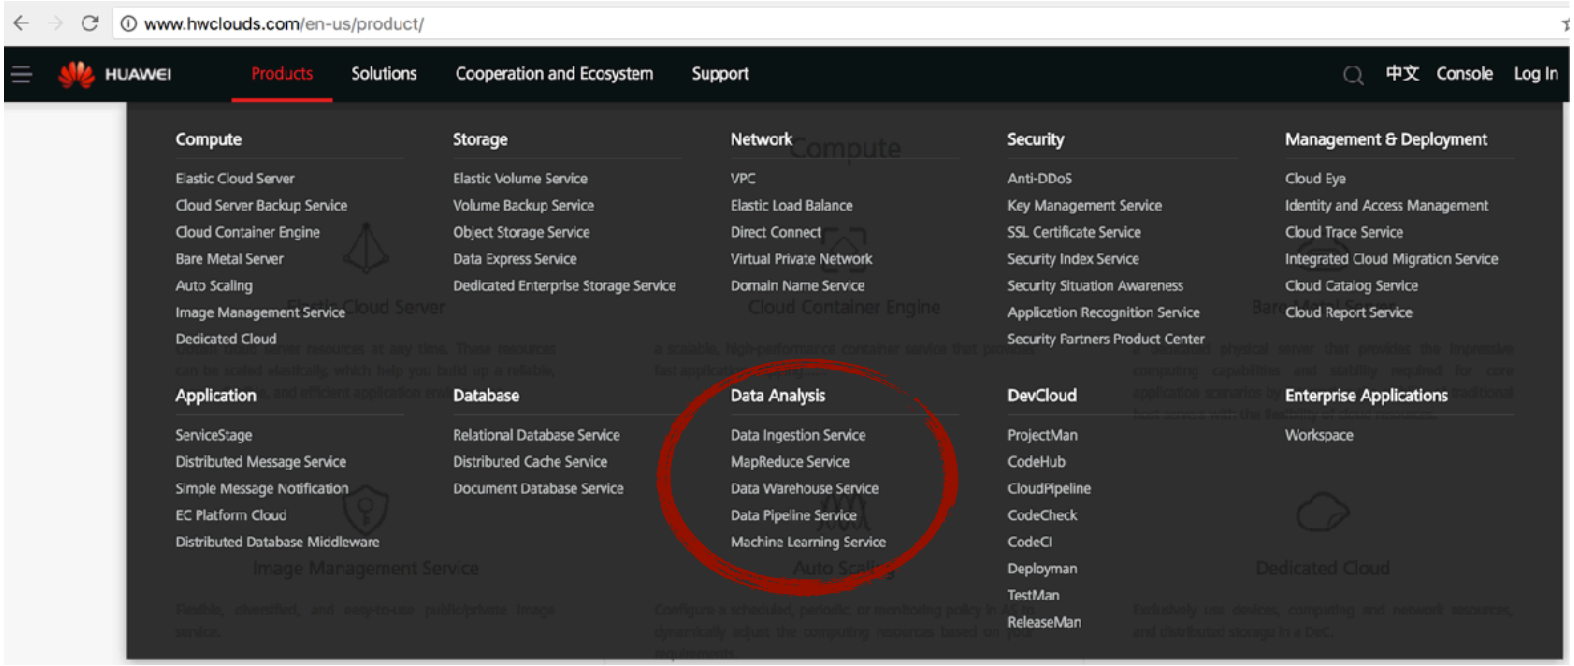

HUAWEI

From Product to Techniques

High usability through clicks

Import graph data

Display result in JSON and visualize it

Usability - Easy to use

Usability - Flexible to use

Perform user queries through query language

HUAWEI

Eywa on Cloud - the Architecture

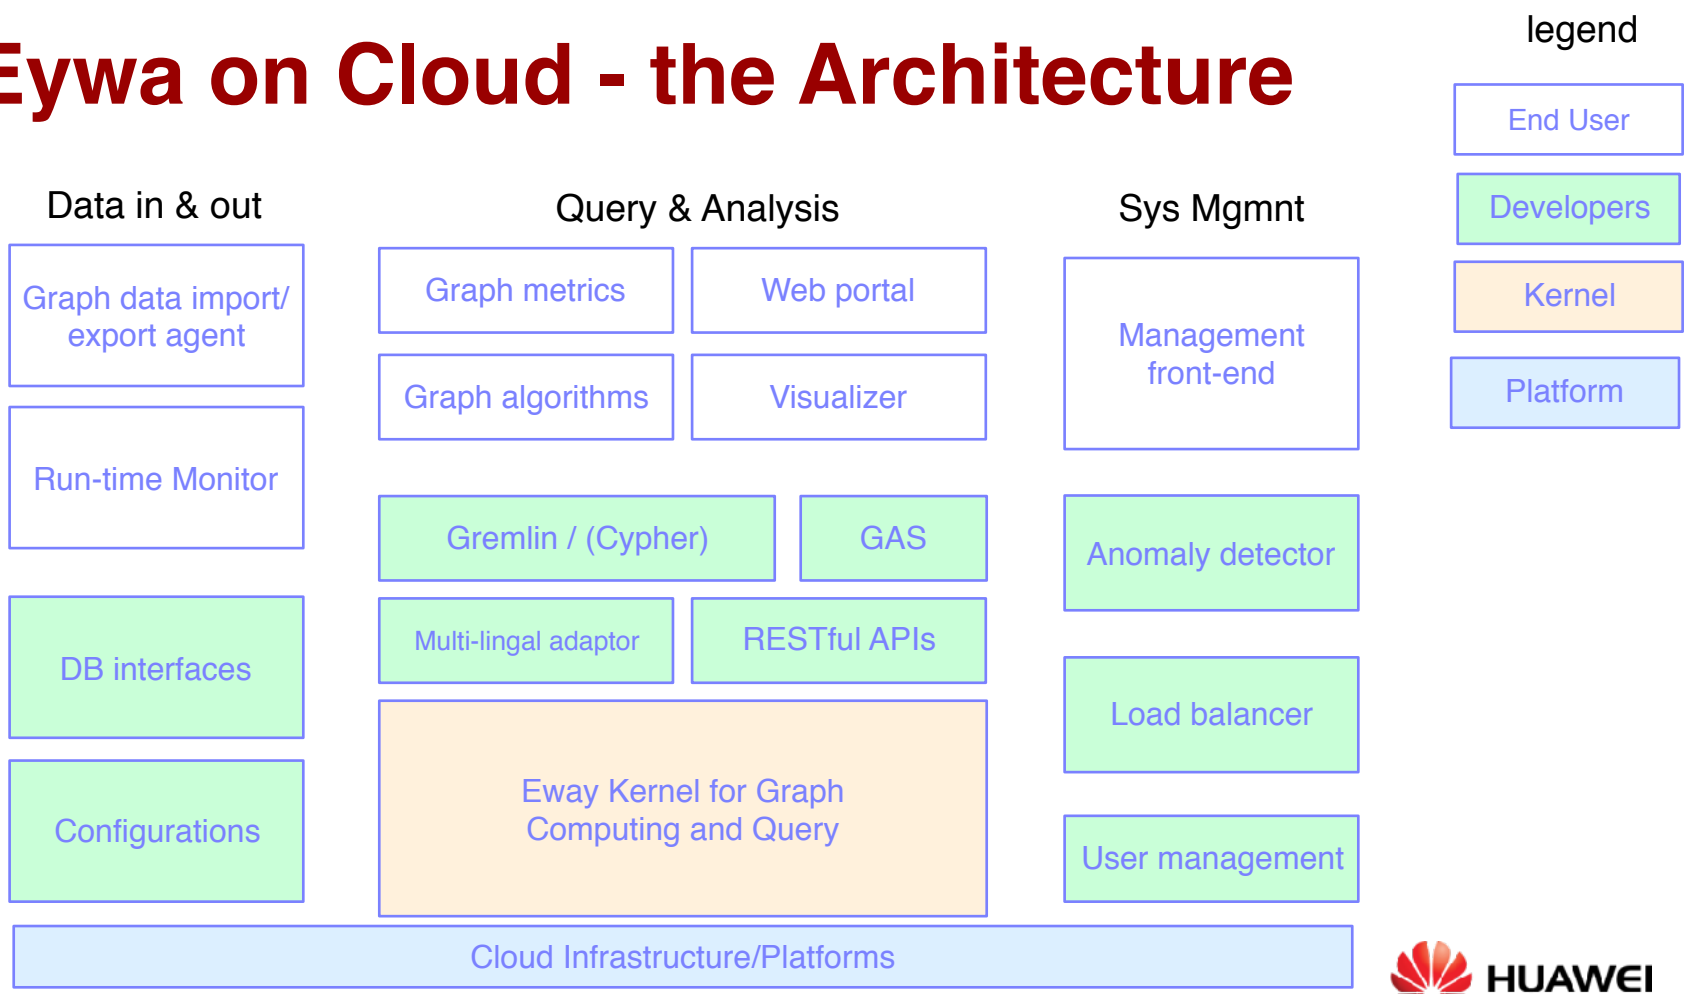

HUAWEI

Trade-off Between Analytics and Query

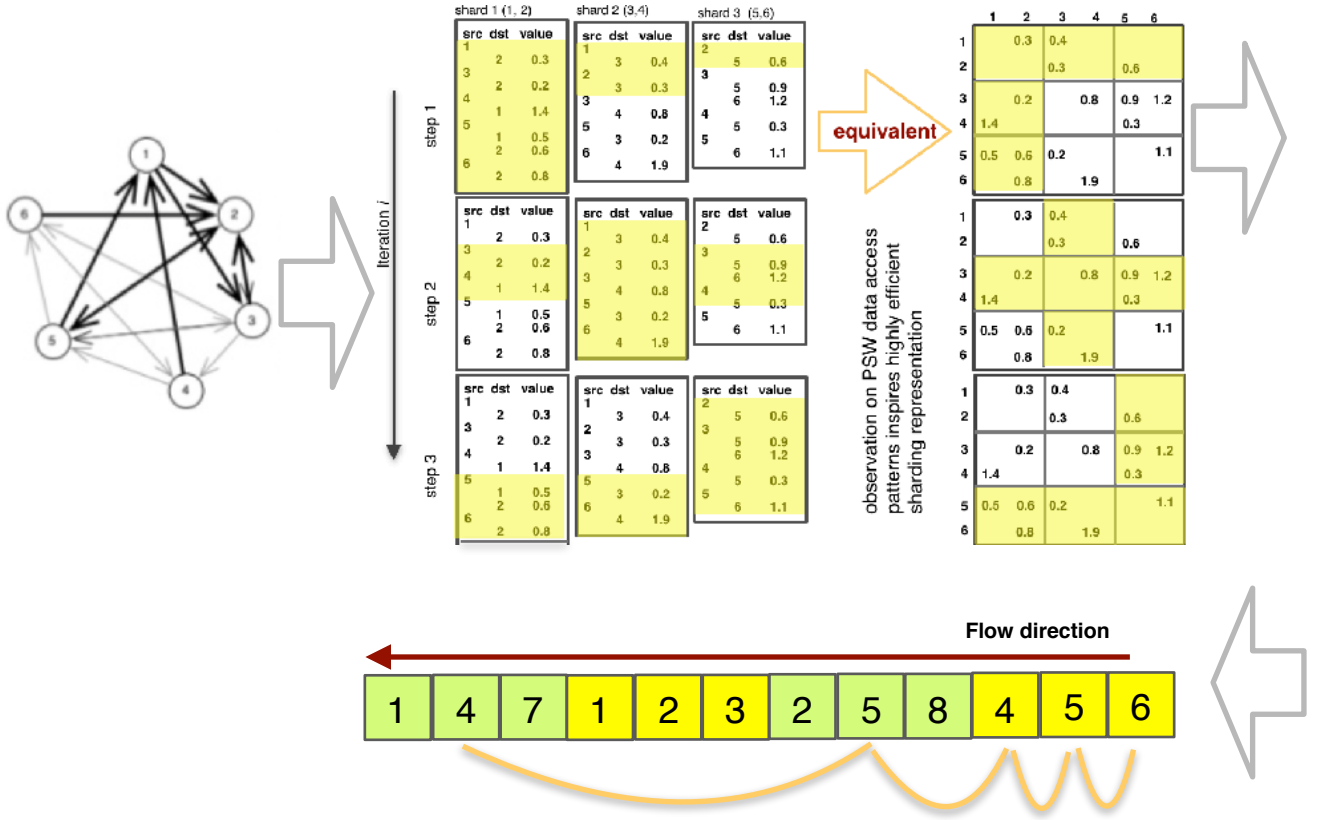

HUAWEI

Labeled Property Graph Model

- **Definition**
  - **graph + property + label**
- **Industry adoption**
  - **de facto model for GraphDB**
  - **Noe4j, Titan, DSE Graph etc**

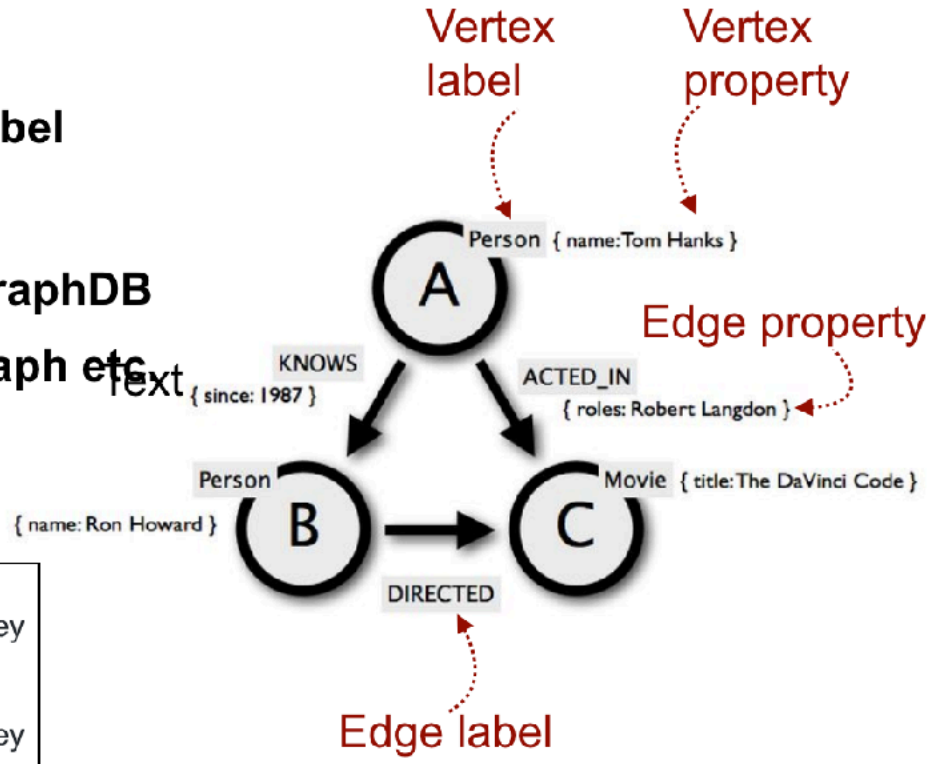

- each vertex has a collection of properties defined by a map from key to value.
- each edge has a collection of properties defined by a map from key to value.

Challenges - Label Ambiguity

Books, authors, and readers
Managers, likes, and collaborators

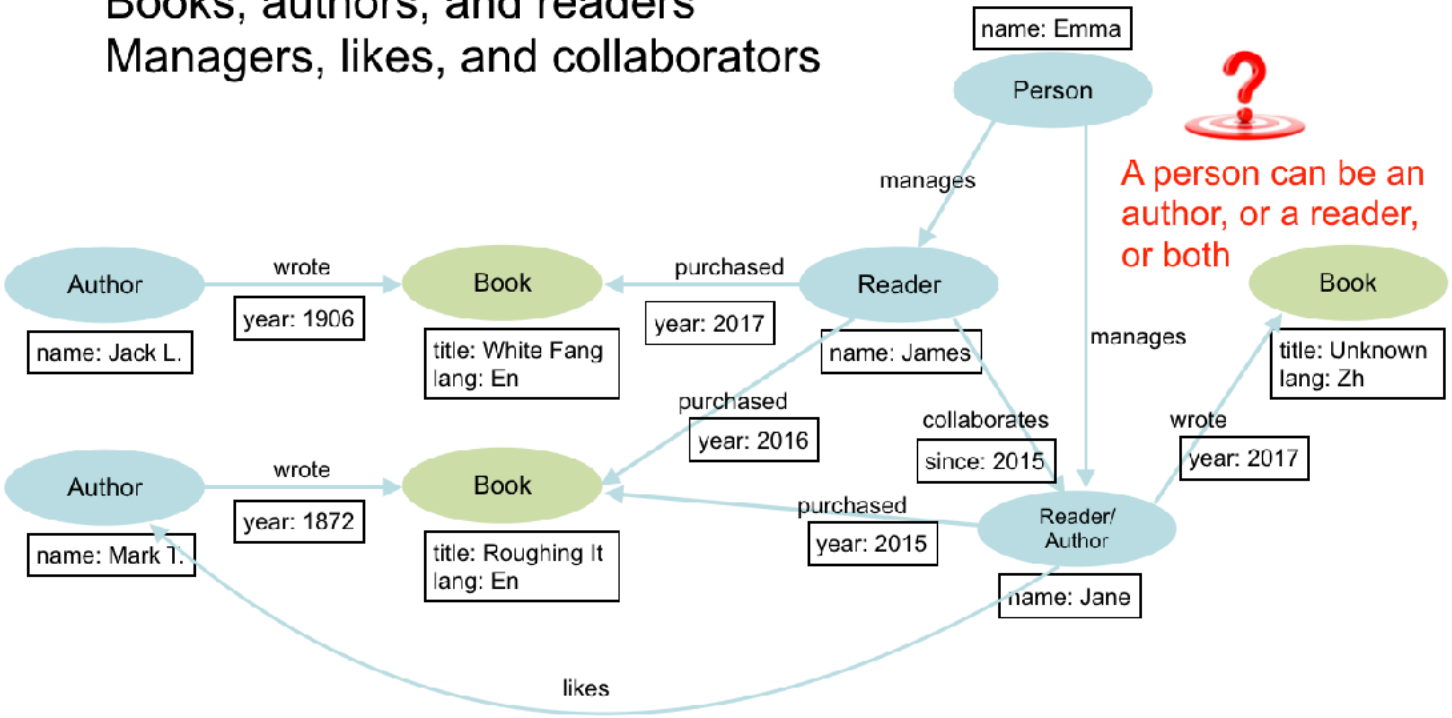

HUAWEI

Graph-based Property Management

- **Schema graph of a property graph**
  - **Vertex — a label representing a vertex type**
  - **Edge — a label representing an edge type**
- **Schema graph ontology**
  - **Relevant labels are connected using directed edge**

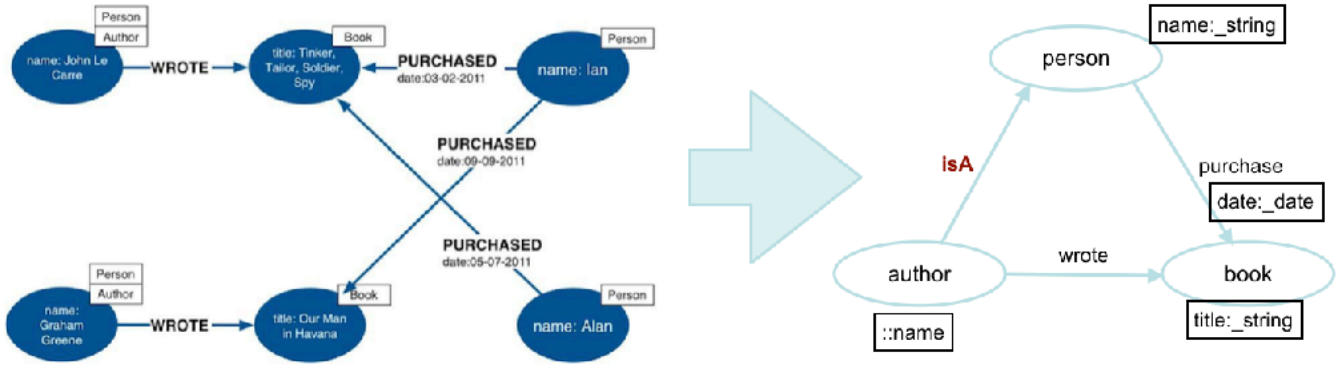

HUAWEI

Graph-based DevOps for Eywa

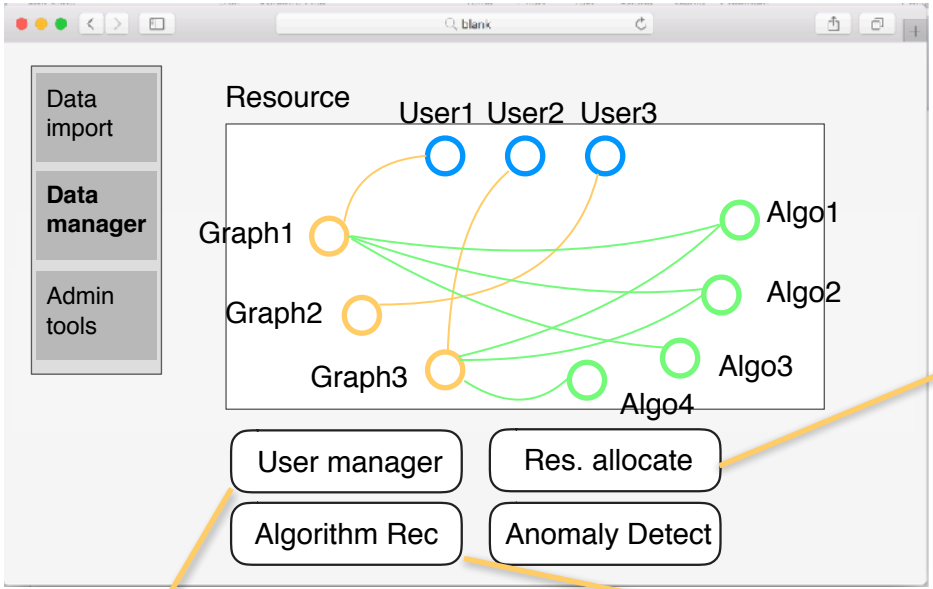

Use a graph to link algorithms, datasets, and users, which DevOps by grouping users, predicting algorithm usage, and offer helps

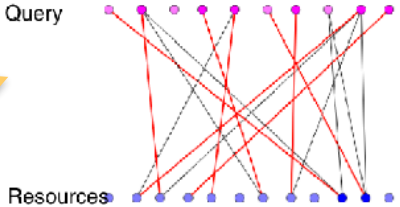

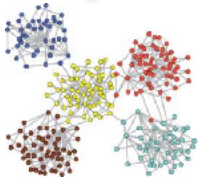

Utilize ML and graph algorithms to analyze the above graphs, Eywa enjoys some basic AI capability in its DevOps

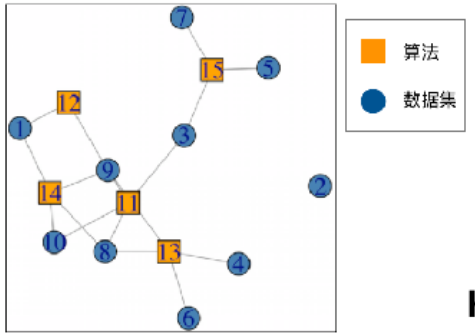

HUAWEI

Experiments

- **Infrastructure**
  - 9 VMs from data center
  - each with 22-core CPU , 125 GB memory
- **Dataset**
  - Orkut: |V|=3.07M, |E|=117M
  - Friendster: |V|=65M, |E|=1.8B
  - Kronecker: |V|=984M, |E|=106.5B
- **Query evaluation**
  - Given a set of vertices, perform 3-hop local traversal

Experiments - 1

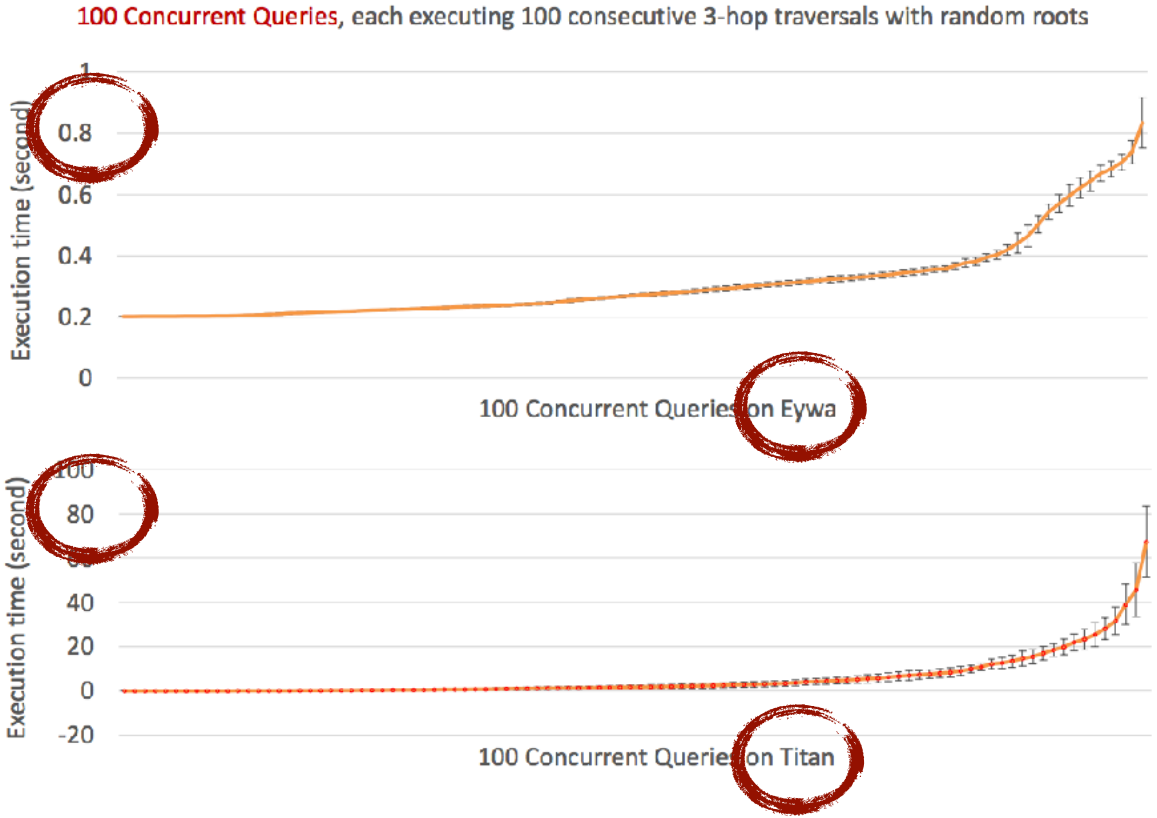

- Dataset: **Orkut**
- Platform: Eywa vs Titan
- Infrastructure: a single machine
- Test bench: local traversal
- Concurrency: 100 queries

HUAWEI

Experiments - 2

100 Concurrent Queries, each executing 100 consecutive 3-hop traversals with random roots

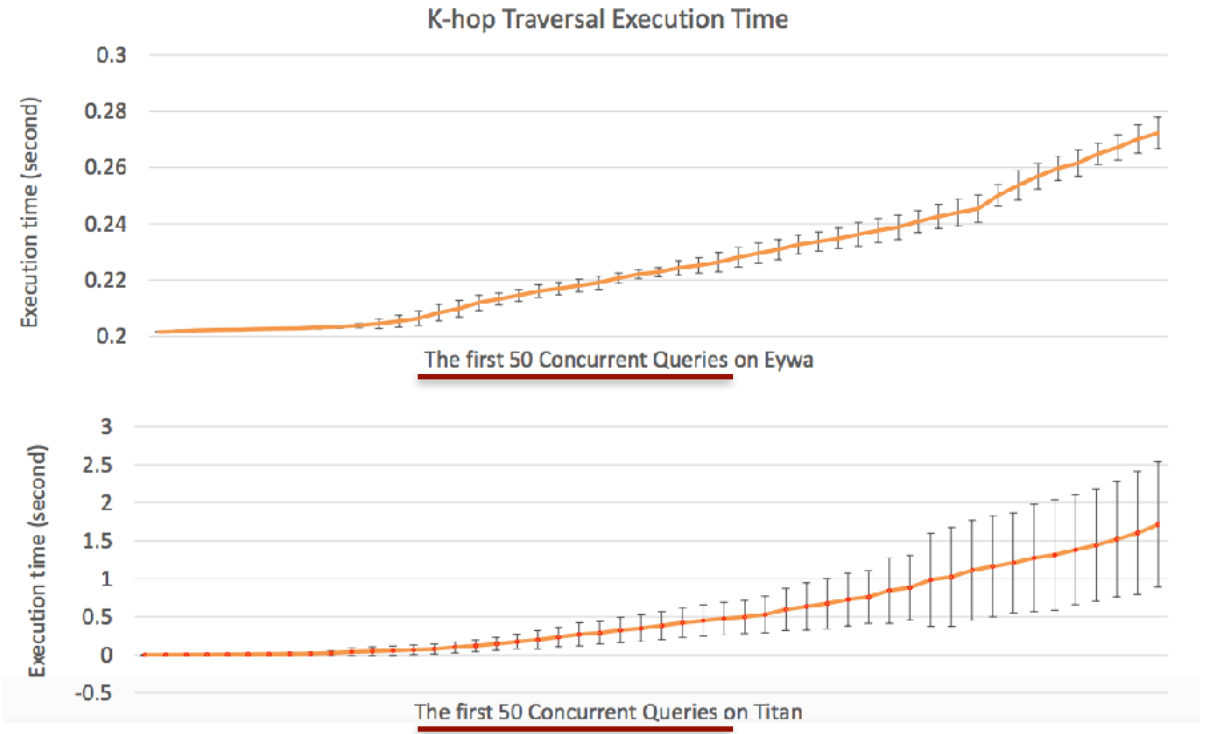

- Dataset: Orkut
- Platform: Eywa vs Titan
- Infrastructure: a single machine
- Test bench: local traversal
- Concurrency: 100 queries

HUAWEI

Experiments - 3

Eywa VS. Titan (|V|: 3072441, |E|:117185083)

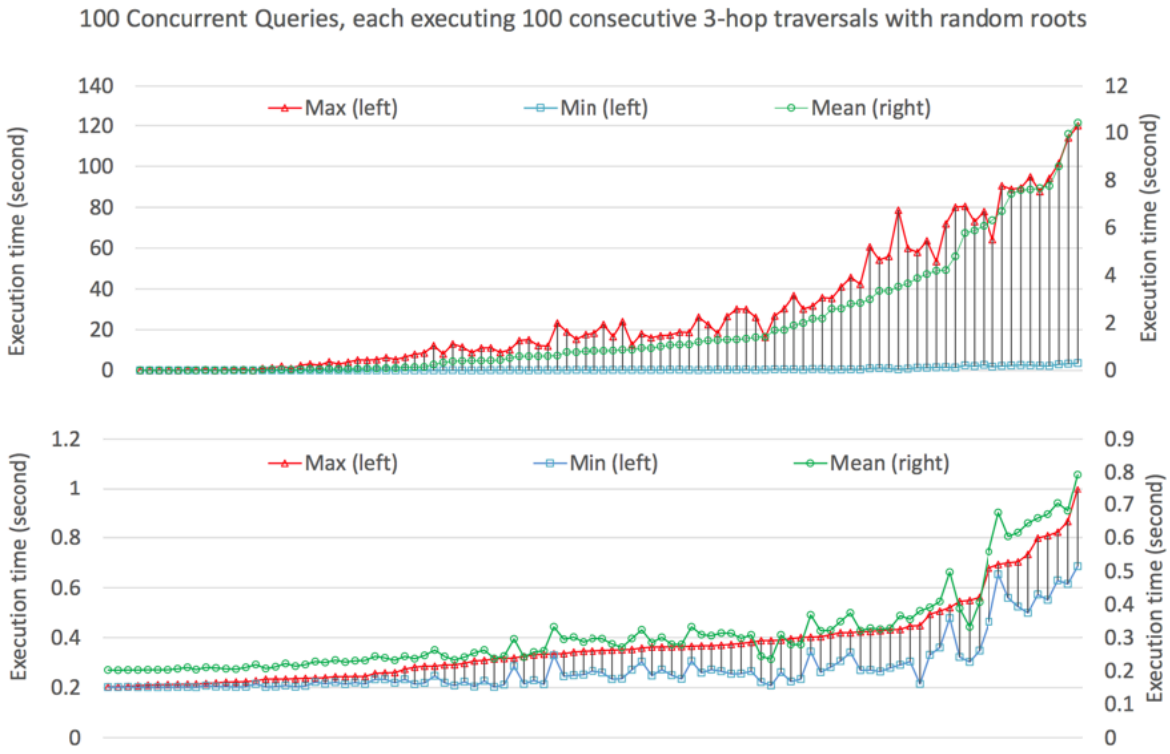

- Eywa outperformed the baseline method
- Eywa shows consistent running time

HUAWEI

Experiment - 4

Eywa (|V|: 984,125,490, |E|:106,579,558,164)

100 Concurrent Queries, each executing consecutive 3-hop traversals with random roots

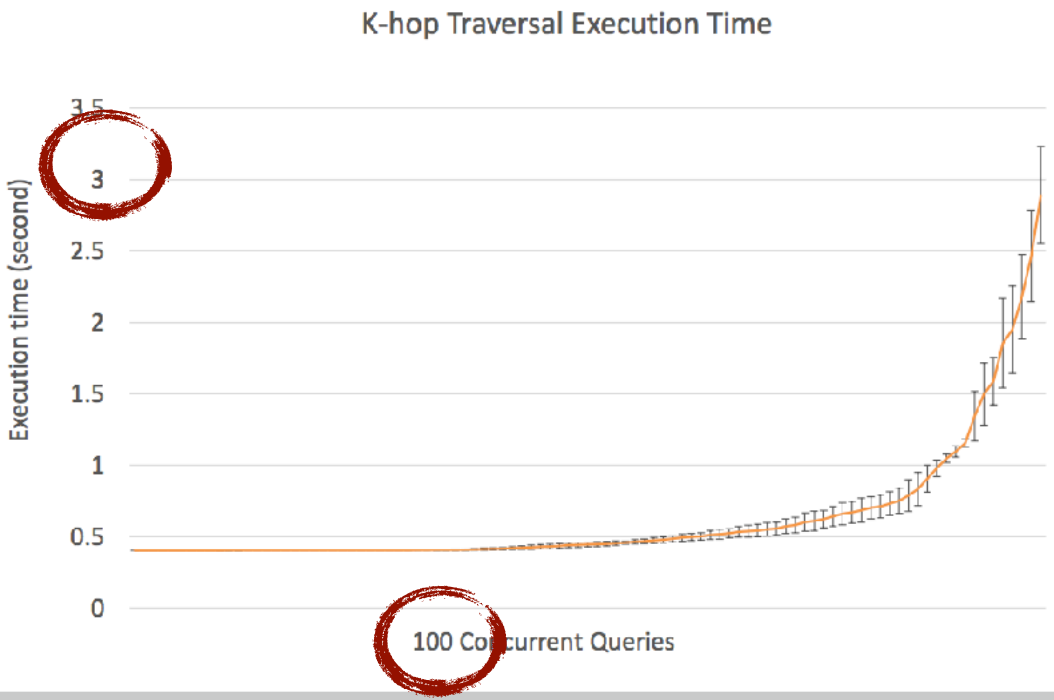

- Dataset: **Kronecker**
- Platform: Eywa
- Infrastructure: 9 VMs
- Test bench: local traversal
- Concurrency: 100 queries

HUAWEI

Experiment - 5

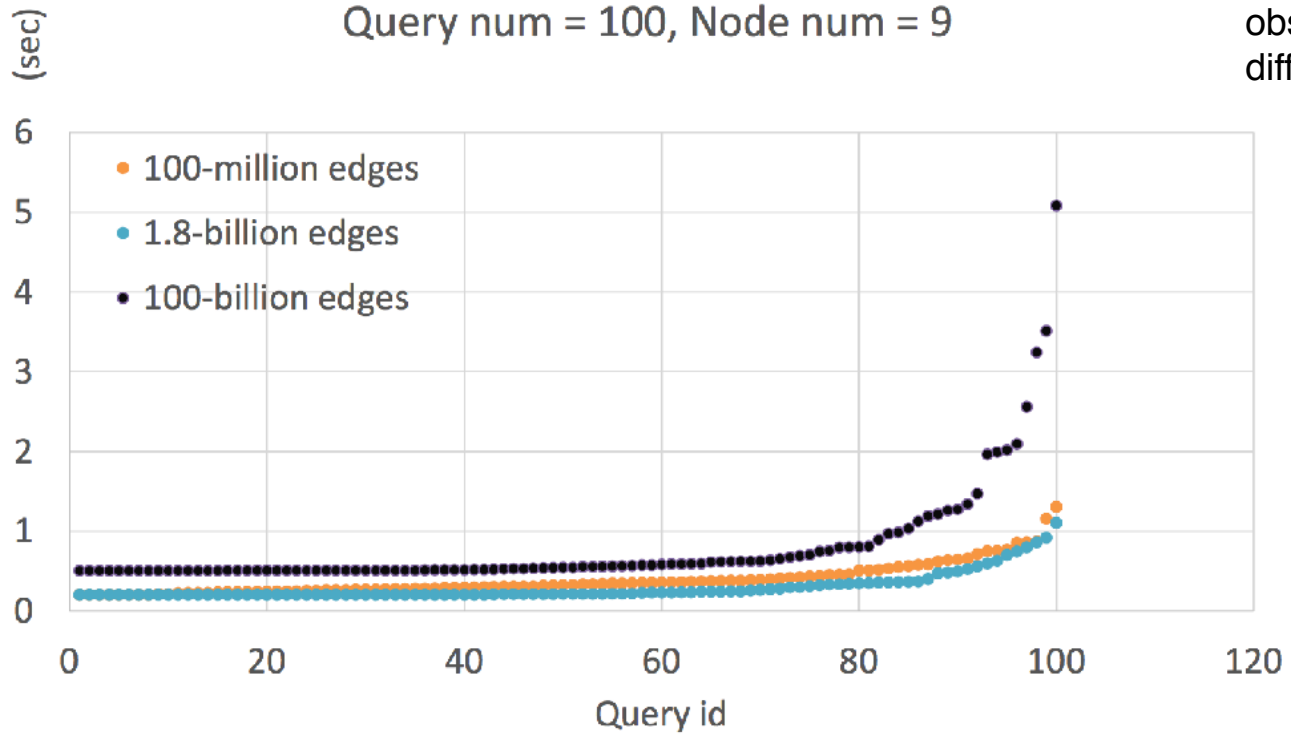

- ***Consistent*** performance observed across graphs of different scales

HUAWEI

Experiment - 6

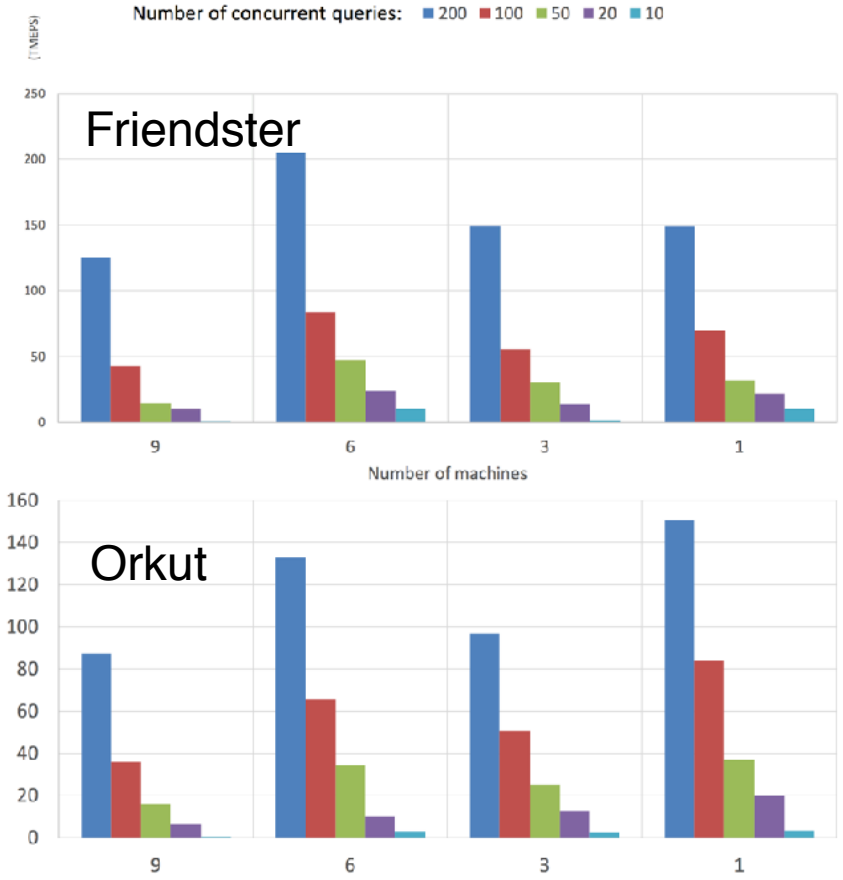

- Consistent performance observed across graphs of different scales
- Better overall throughput for **higher concurrency**
- For **bigger** graph, the **throughput** increases to some extend as the number of VMs increases

HUAWEI

Experiment - 7

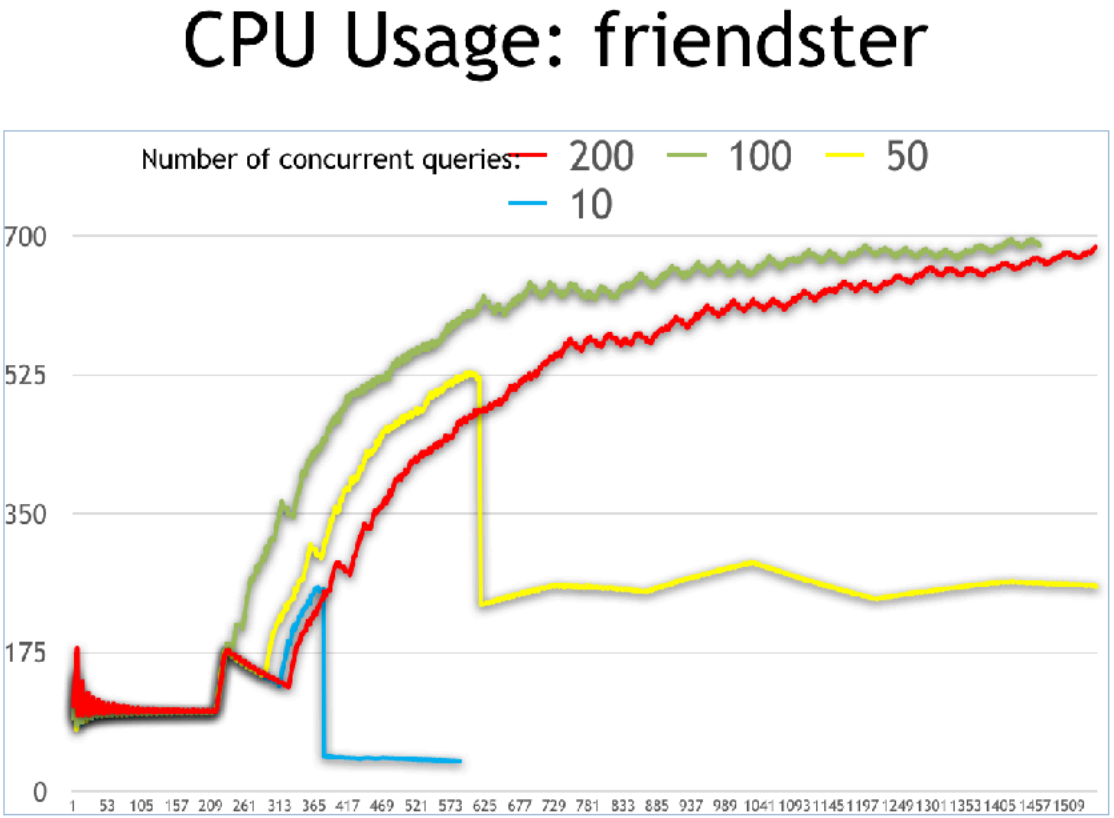

- CPU usages seems not high
- Further concurrent queries can be supported
- Each query may incur multiple threads (workers) for distributed communication reasons, which consumes little CPU resources, but occupy a core (OS can swap them out)

HUAWEI