DATABASES IN THE CLOUD

@andy_pavlo
CMU-Q 15-440
December 3rd, 2014

CARNEGIE MELLON
DATABASE GROUP

OLTP vs. OLAP databases.

Source: https://www.flickr.com/photos/adesigna/3237575990

On-line Transaction Processing

- Fast operations that ingest new data and then update state using ACID transactions.
- Only access a small amount of data.
- Volume: 1k to 1m txn/sec
- Latency: >1-50 ms
- Database Size: 100s GB to 10s TB

Example

- Maintain player's state in an on-line game in the OLTP database.

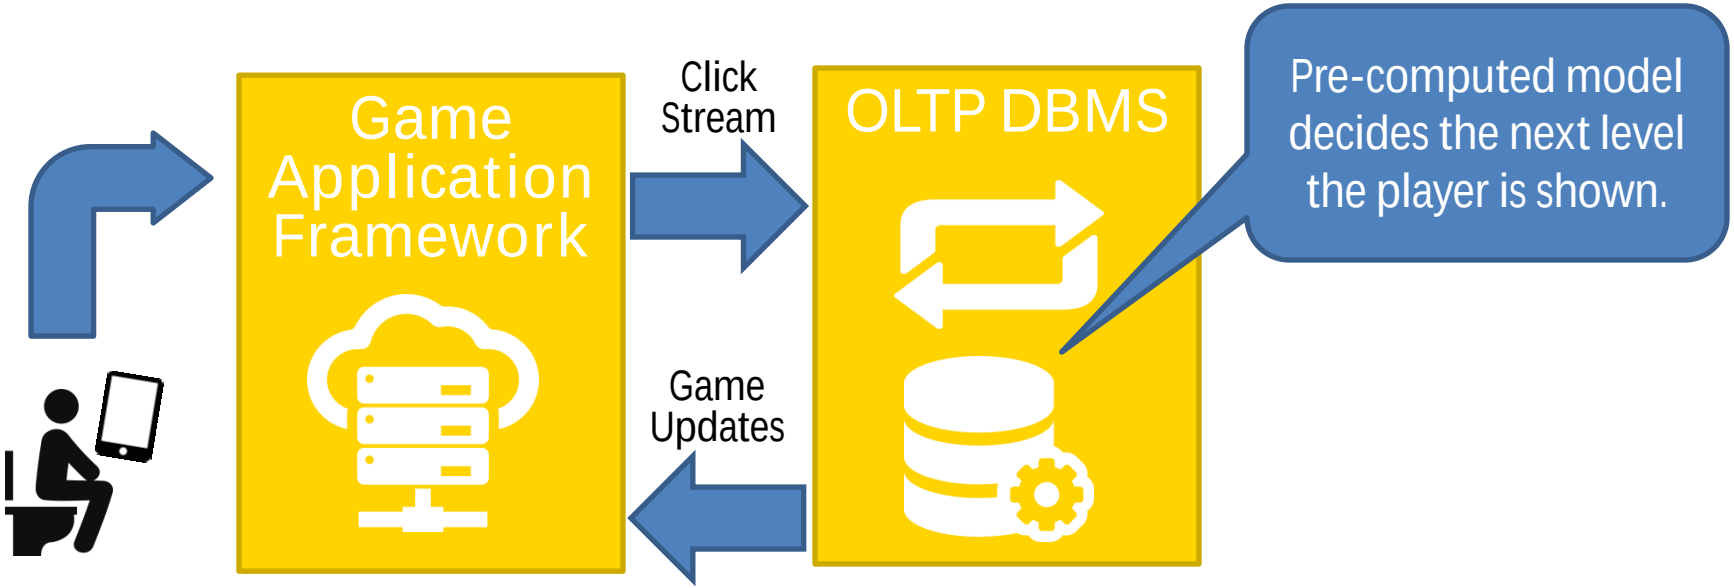

Example

- Maintain player's state in an on-line game in the OLTP database.

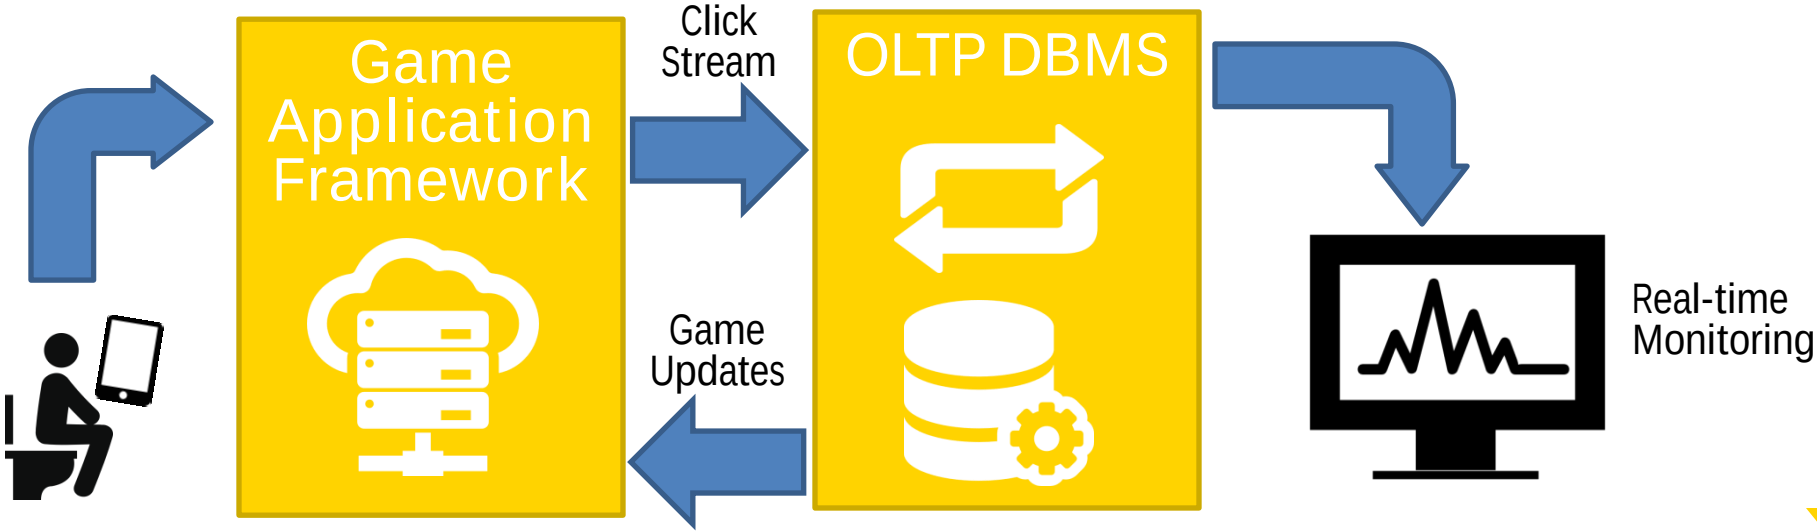

Database Warehouses

- Complete history of OLTP databases.
- Complex queries that analyze large segments of fact tables and combine them with dimension tables.
- Volume: A couple queries per second
- Latency: 1-60 seconds
- Database Size: 100s TB to 10s PB

Example

- Compute model used to guide OLTP DBMS decisions from historical data.

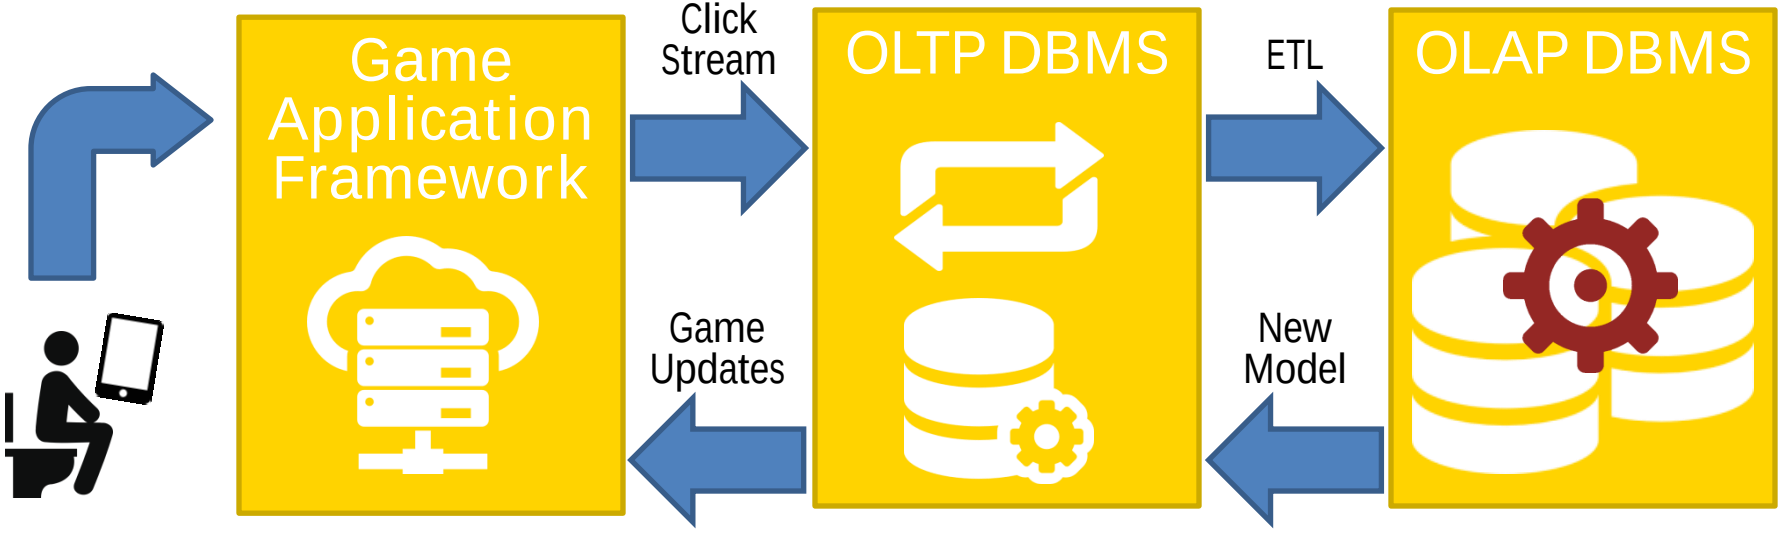

OLTP vs. OLAP

- Storage Format:
  - *OLTP → Row-oriented*
  - *OLAP → Column-oriented*
- Primary Database Location:
  - *OLTP → In-Memory*
  - *OLAP → Disks*
- Workloads:
  - *OLTP → Write-Heavy*
  - *OLAP → Read-Only*

Things to consider with databases in the cloud.

Source: https://www.flickr.com/photos/arvidnn/152491335

Good Things

- Better Resource Utilization
- Elastic Scaling
- Database-as-a-Service Offerings

Better Resource Utilization

- Combine multiple silos onto overprovisioned resources.
- Public platform providers achieve better economies of scale.
- Database machines are (mostly) dead.
- Optimal multi-tenant placement is a difficult problem.

Elastic Scaling

- Automatically provision new resources on the fly as needed.
- Scaling *up* vs. Scaling *out*.
- Difficult for OLTP DBMS to continue processing transactions while data migrates.

OLTP Scale-out Example

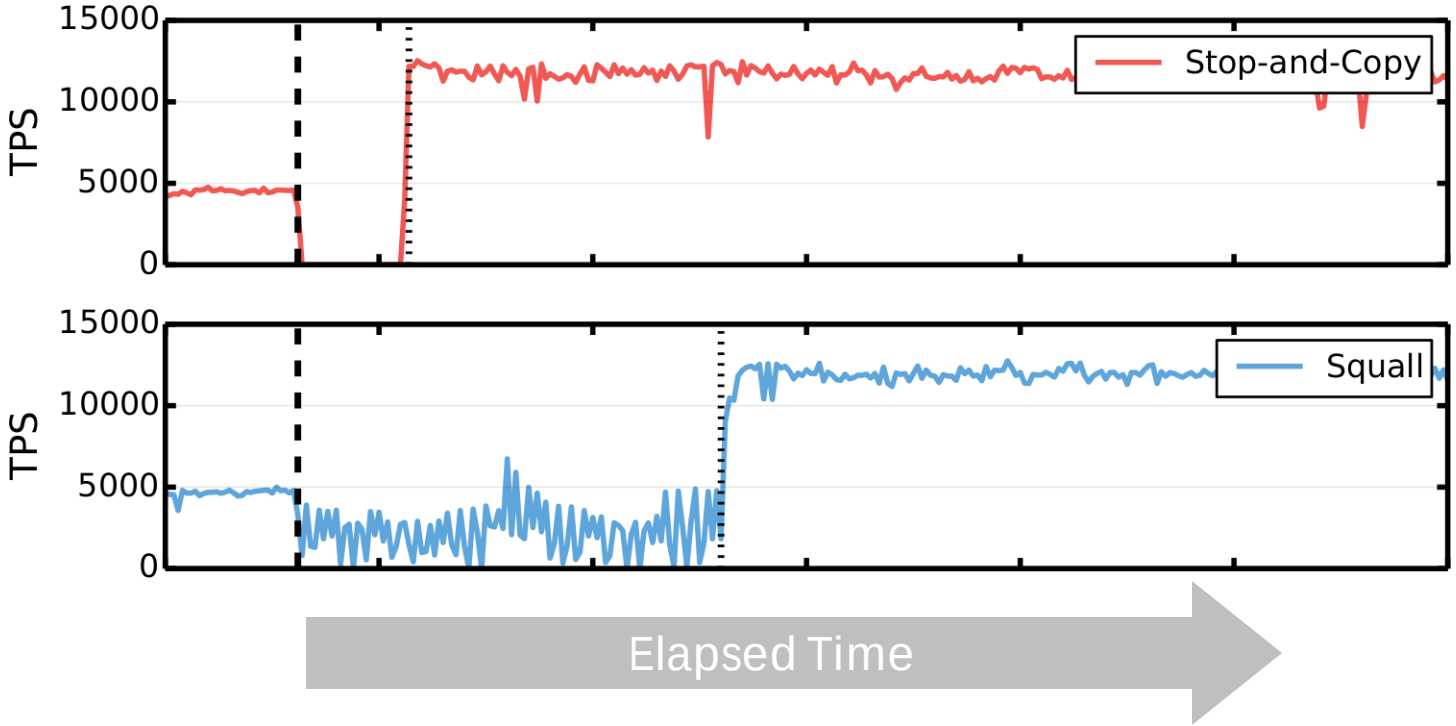

TPC-C Benchmark on H-Store (Fall 2014)
Scaling from 3 to 4 nodes

E-Store: Fine-Grained Elastic Partitioning for Distributed Transaction Processing
R.Taft, E.Mansour, M.Serafini, J.Duggan, A.J. Elmore, A.Aboulnaga, A.Pavlo, M.Stonebraker
Proceedings of the VLDB Endowment, vol. 8, iss. 3, pages. 245—256, November 2014.

Database-as-a-Service

- Cloud provider manages physical configuration of a DBMS.
- Ideal for applications that are co-located in the same provider's compute platform.
- Combine private data with curated databases (i.e., data marts)

Bad Things

- I/O Virtualization
- File system Replication
- Security + Privacy Concerns
- Performance Variance

I/O Virtualization

- Distributed file system stores data transparently across multiple nodes.
- The data is not "local" to the DBMS.
- This causes a DBMS "*pull data to query*" rather than "*push query to data*".

OLAP I/O Virtualization

```
SELECT YEAR(o_date) AS o_year, AVG(o_amount)
FROM orders GROUP BY o_year
ORDER BY o_year ASC
```

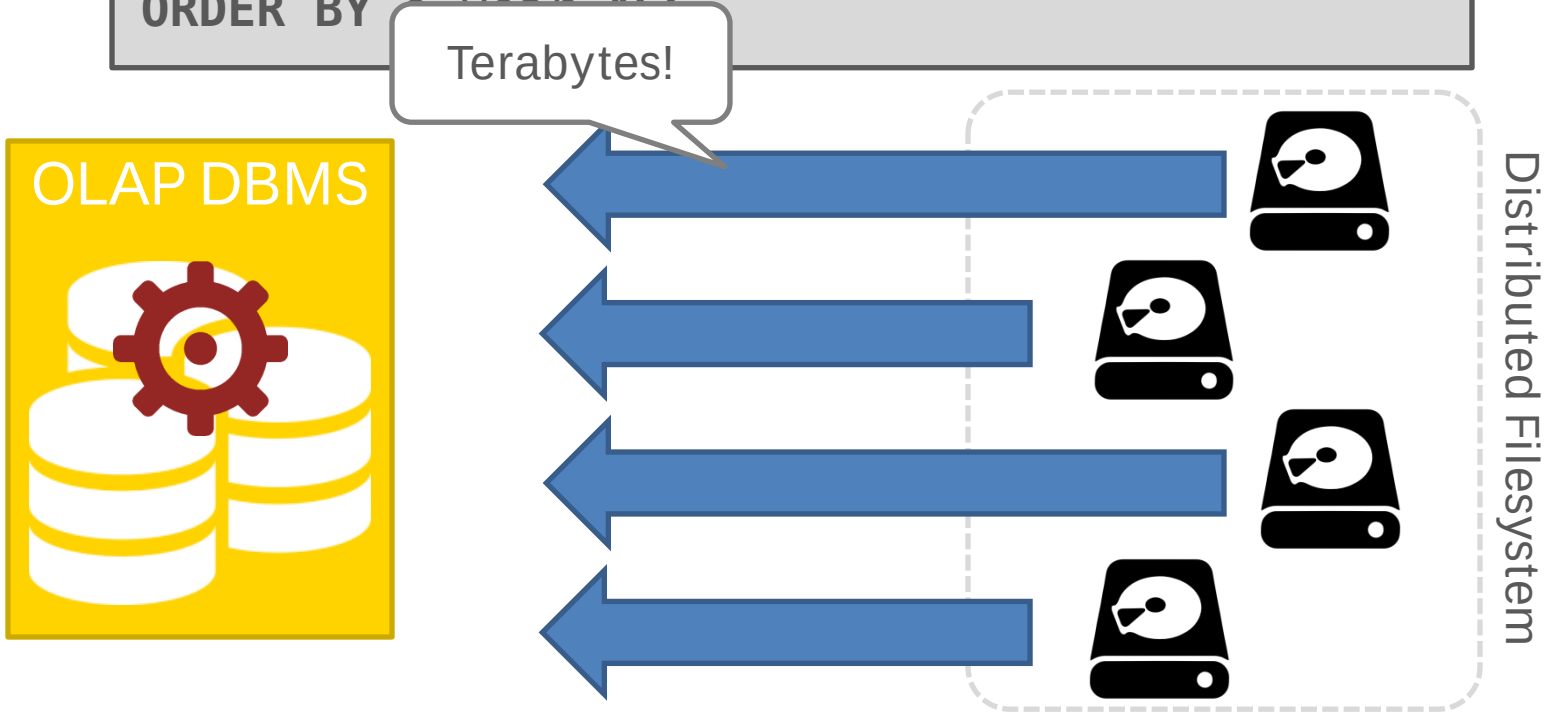

OLAP I/O Virtualization

```
SELECT YEAR(o_date) AS o_year, AVG(o_amount)
FROM orders GROUP BY o_year
ORDER BY o_year ASC
```

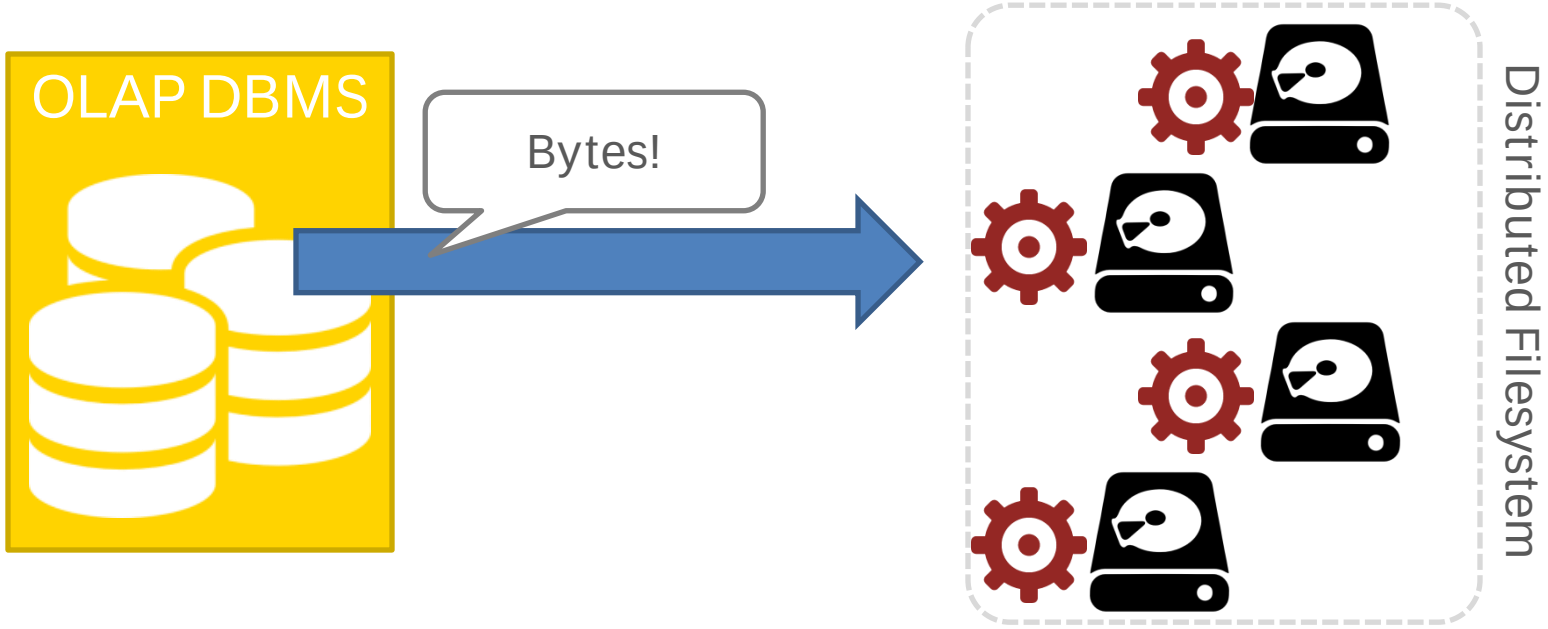

File System Replication

- The DBMS should not rely on file system replication for durability.
- OLTP systems maintain replicas in-memory.
- OLAP systems can store copies of tables in different ways on replica nodes.

OLAP Replication

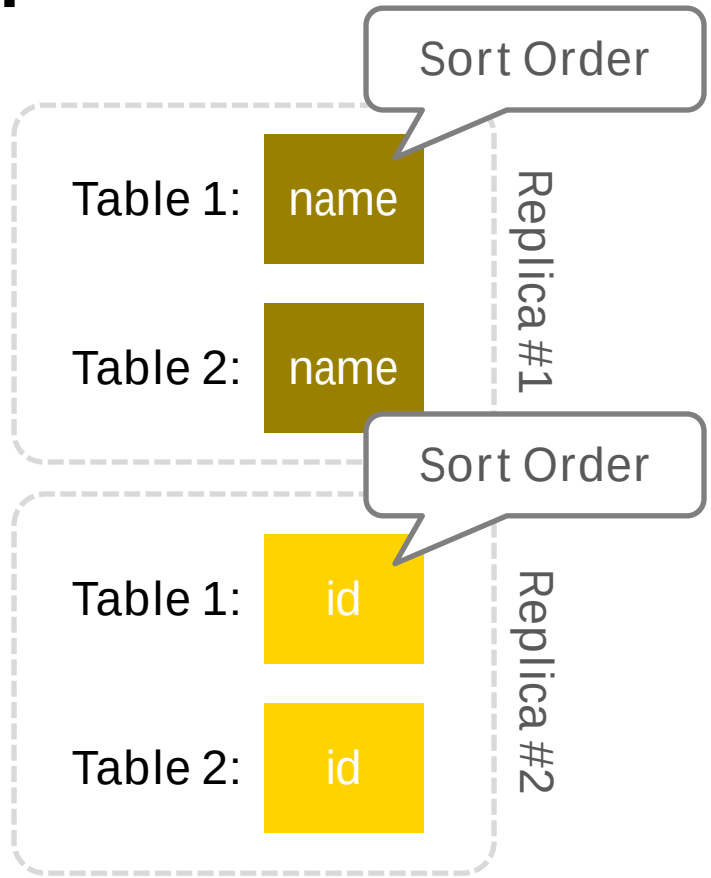

OLAP Replication

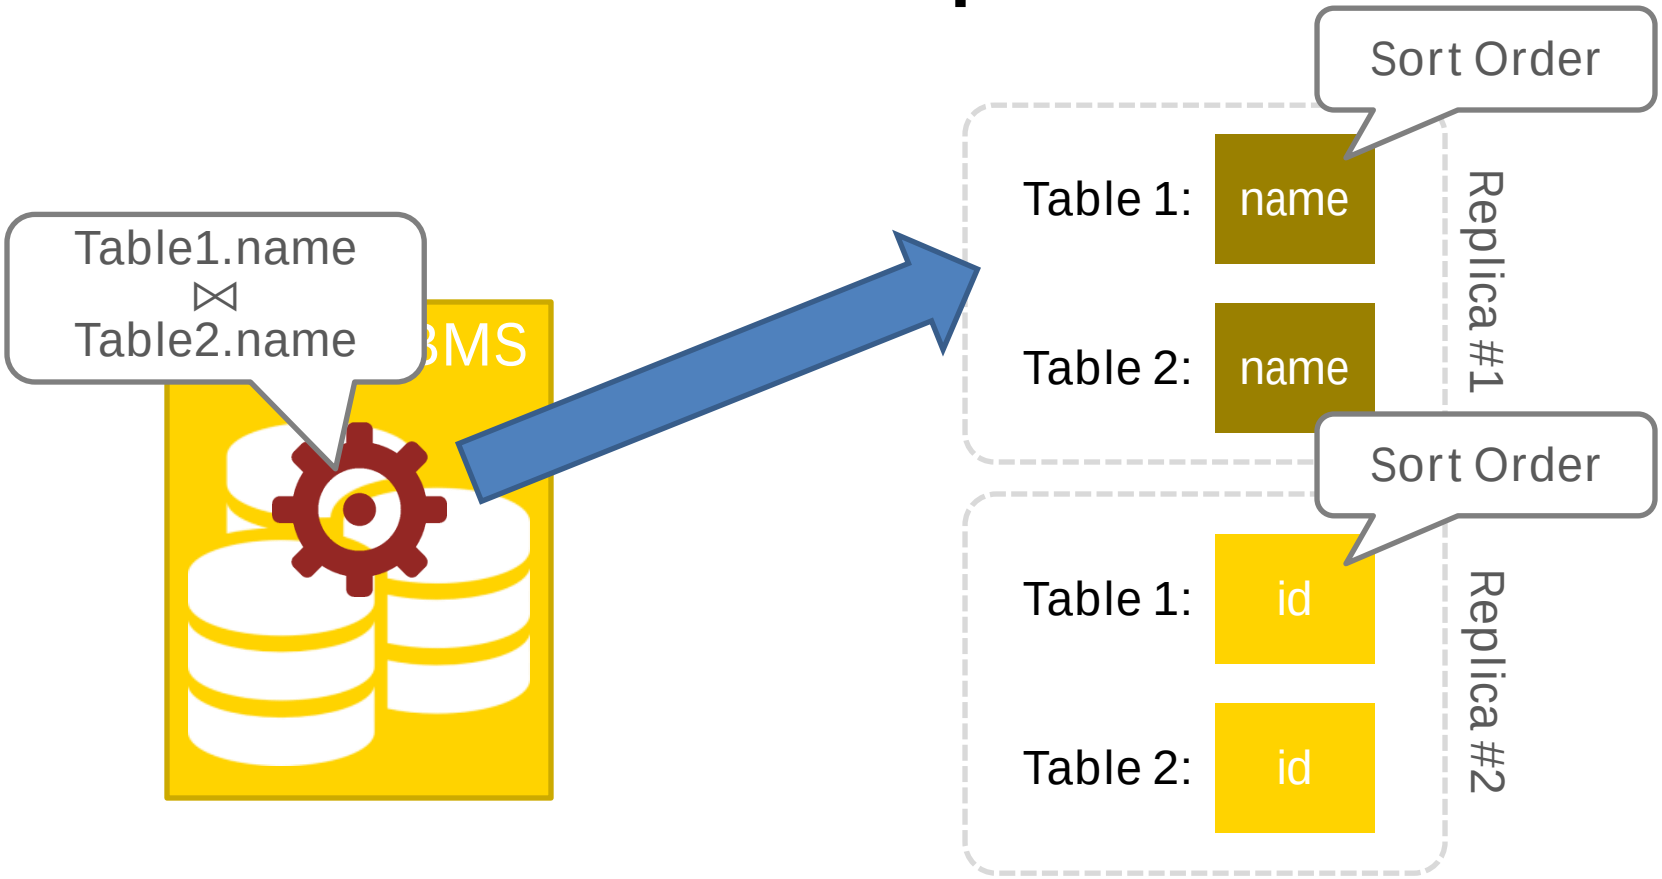

Security + Privacy Concerns

- No truly encrypted solution exists.
- Many companies are unable to use public cloud platforms.

Performance Variance

- DBMSs are sensitive to changes in underlying hardware performance.
- "Noisy" neighbors on instances can cause large fluctuations in performance.

OLTP Performance Variance

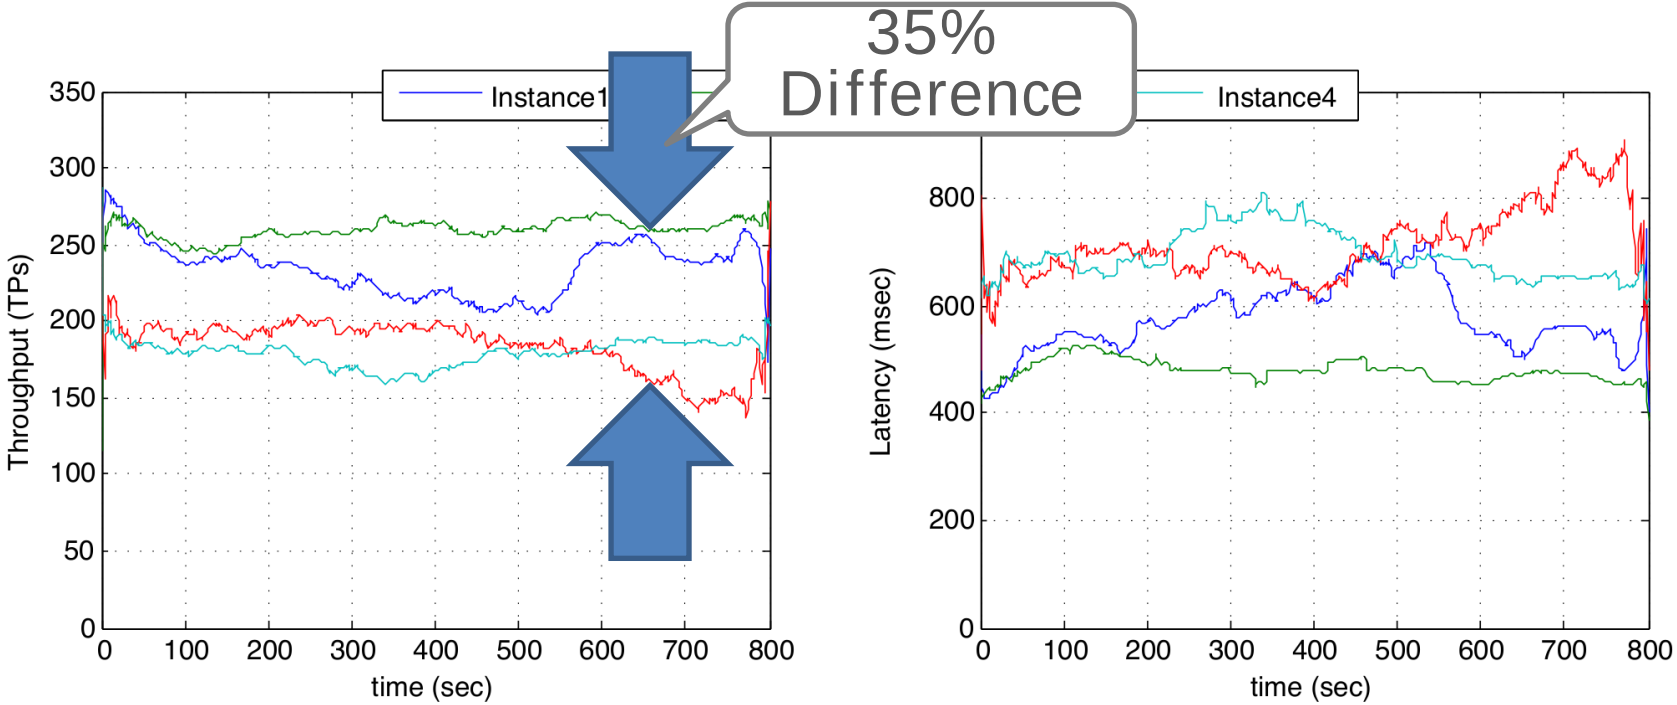

YCSB on MySQL (Winter 2012)
Medium EC2 Instances

OLTP-Bench: An Extensible Testbed for Benchmarking Relational Databases
Djellel Eddine Difallah, Andrew Pavlo, Carlo Curino, Philippe Cudre-Mauroux
Proceedings of the VLDB Endowment, vol. 7, pages. 277—288, December 2013.

Source: https://www.flickr.com/photos/alestra/889158563

Important Features

- Automatic Back-ups
- Geo-replication
- Elasticity / Live Reconfiguration
- Efficient Multi-Tenancy
- Workload Awareness

Cloud Database Vendors

- Cloud-friendly systems
- Database-as-a-Service (DBaaS)

Cloud-friendly DBMSs

- Most DBMS vendors make it easy to deploy on cloud platforms.
- Others provide support for easy scale-out in a cloud environment.
- More than just pre-configured instances.

OLTP DBaaS

- Amazon RDS / Aurora
- Microsoft Azure
- Google Cloud SQL
- Database.com
- ClearDB
- GenieDB
- Clustrix

OLAP DBaaS

- Amazon Redshift 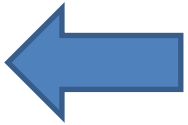
- Google BigQuery
- Microsoft Azure
- Snowflake

Parting Thoughts

- The cloud does not magically make database problems go away.
- It's easier to run an OLTP DBMS than a OLAP DBMS on the cloud.
- AFAIK, there is no truly autonomous DBMS as of yet.

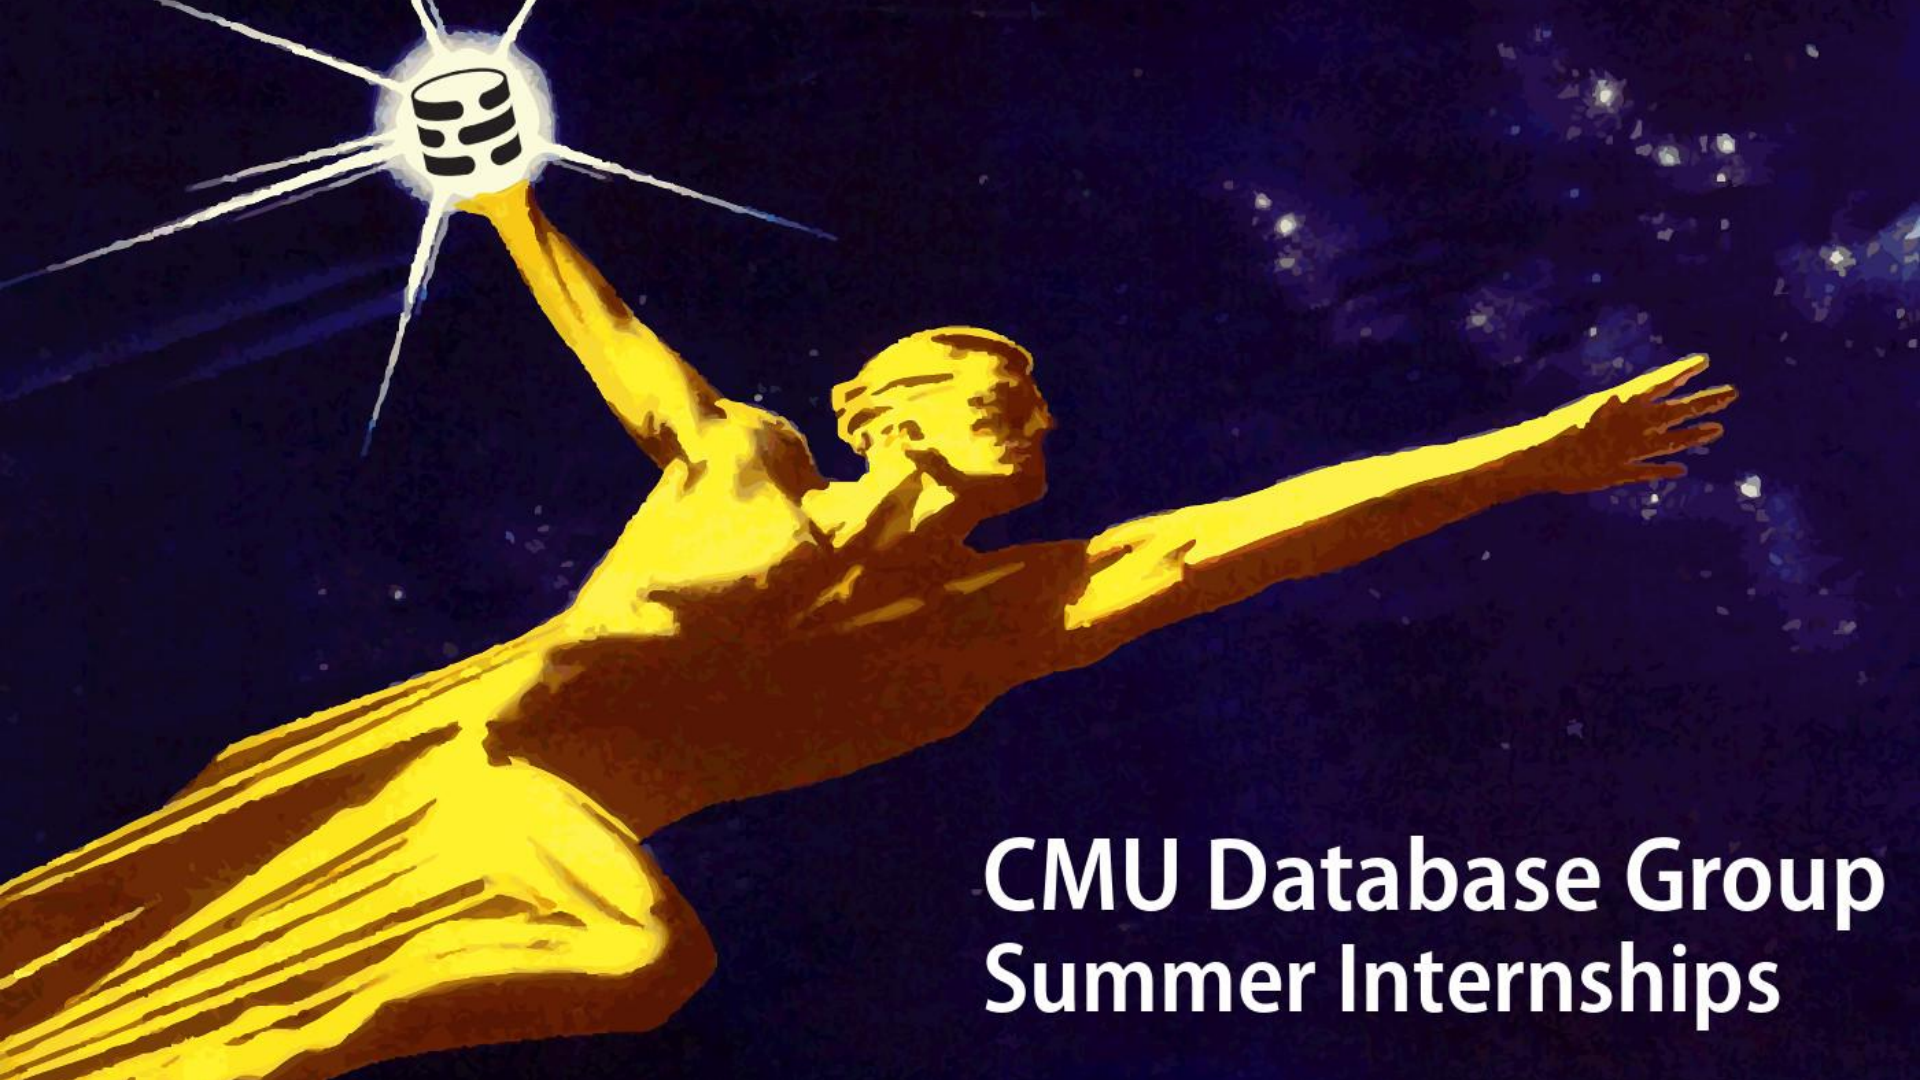
CMU Database Group
Summer Internships

END

@andy_pavlo