

Hello!

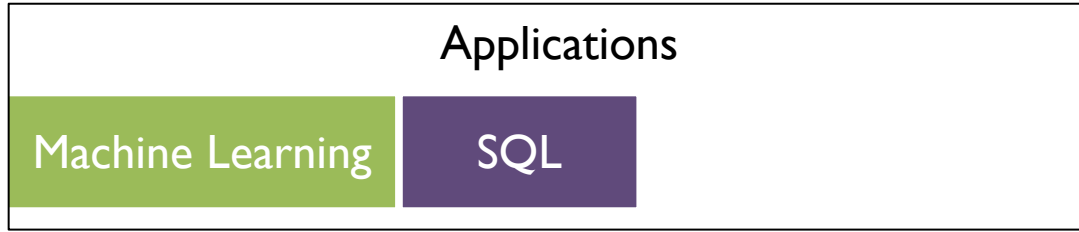
CS 744: SNOWFLAKE

Shivaram Venkataraman

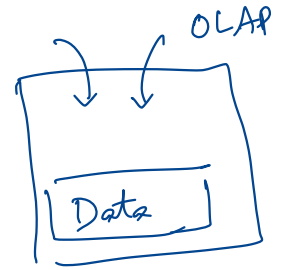
Fall 2022

ADMINISTRIVIA

- Project Proposals – Due 11am tomorrow! → Highly recommend submitting today!
- Midterm on Thursday! Seating layout?
 - ↳ Open book, open notes
 - ↳ send email for extensions



Data
ware
house



SparkSQL/Scope: Given a query how do you run it efficiently?
↳ Optimizer, Schema

Snowflake: How do you build an elastic data warehouse?
↳ OLAP

↓
adjust amount of resources used
at runtime

AWS / Azure

CLOUD COMPUTING STACK

PyTorch

Machine Learning

SCOPE

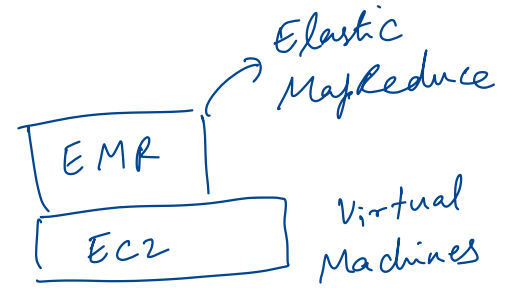
SQL

Analytics as a Service!

Snowflake

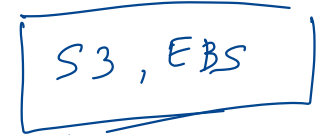
MapReduce
Spark

Computational Engines



GFS

Scalable Storage Systems



Storage as a service



SNOWFLAKE: GOALS

Software-as-a-Service → Users do not install or deploy anything.
Connecting on your browser

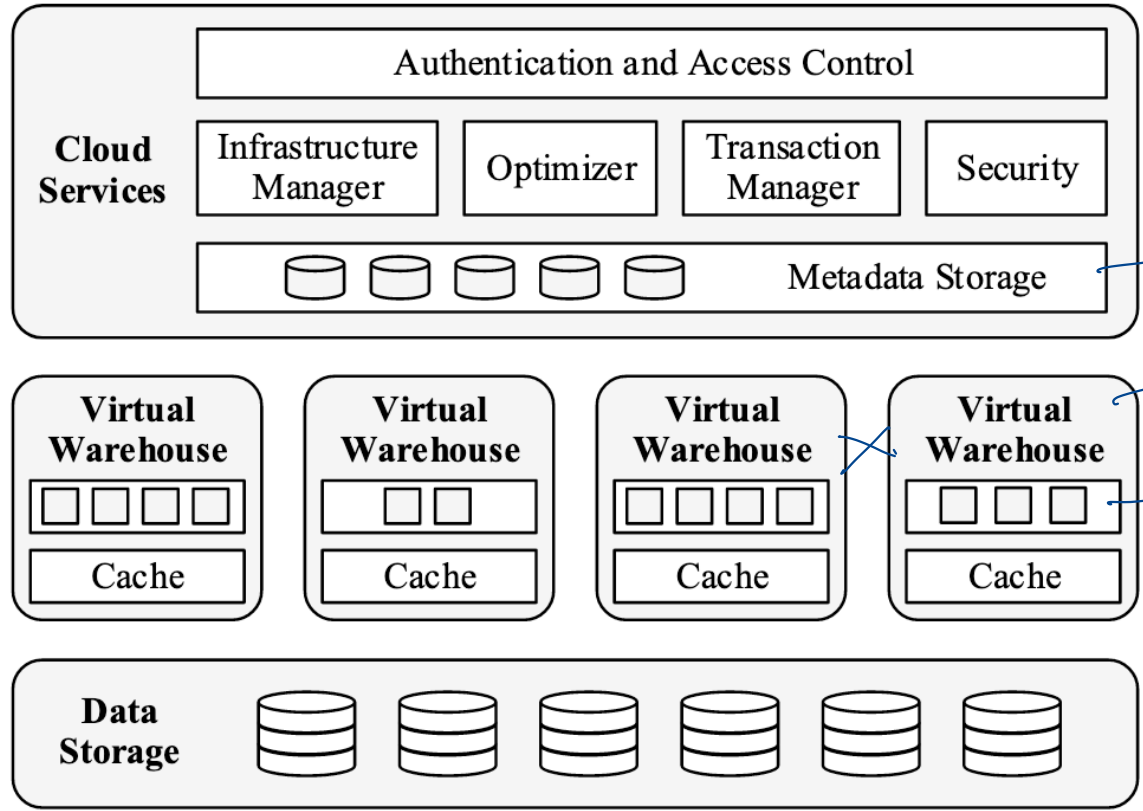
Elastic → Use resources as required.
↳ Use CPU / mem when you have work
shut down otherwise

Highly Available
↳ Fault Tolerant

Semi-Structured Data

↳ Real data arrives as JSON/XML formats. May not have a schema

SNOWFLAKE DESIGN



Control plane
(GFS Master
Mesos...)

Data plane

Compute

shared across
VWs
Storage

Metadata
Storage
separate module

separate for
user
workers
Isolation

STORAGE VS COMPUTE

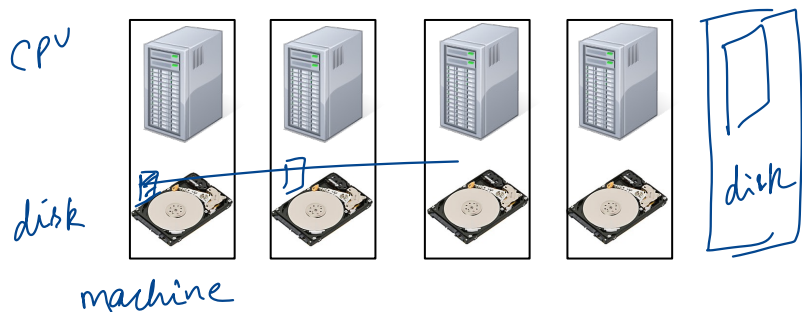
- ties together CPU and disk
cannot independently scale them

- Scale independently
- lose locality!

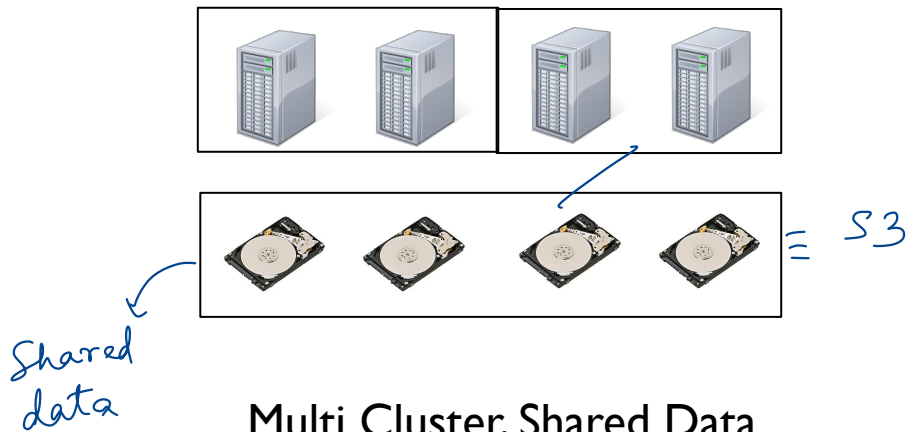
- load balancing

- Data locality \leftrightarrow rack / machine locality

Every storage access is over the network



Shared Nothing



Multi Cluster, Shared Data

When is shared nothing

- Locality is very important

- Data isolation

↳ Storage: Policies on data types

↳ Access policies

Shared data

- Network usage is low

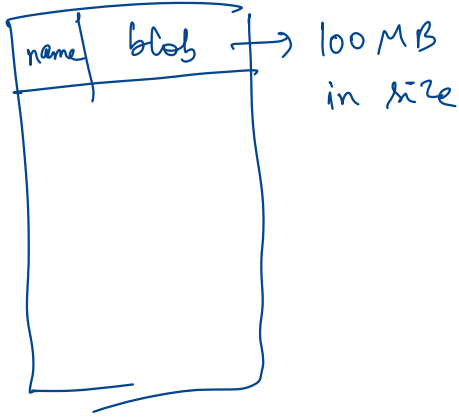
- Irregular arrival of queries

but data is stored permanently

- Replication across data centers make utilization even more important

S3
KV store

STORAGE: HYBRID COLUMNAR



Alice	32
Bob	22
Eve	24
Victor	27

chunk size = 2 rows

Alice,32,Bob,22

Eve,24,Victor,27

Row-oriented

Alice, Bob, 32,22

Eve,Victor,24,27

Hybrid Columnar

Access only part
of file for a
query

Compress much better

Within a file / partition
use columnar
storage

VIRTUAL WAREHOUSES

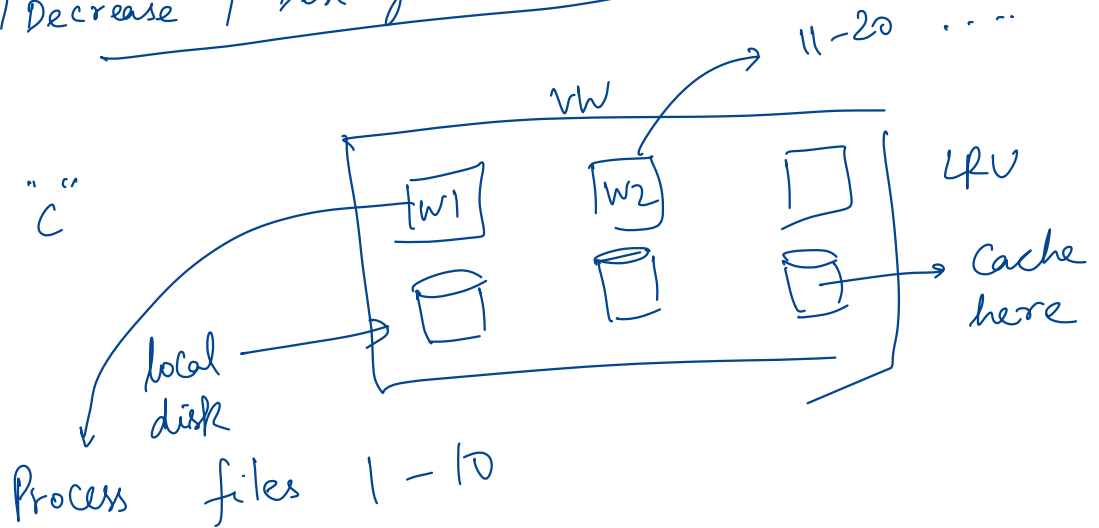
Elasticity, Isolation → VW: collection of worker VMs
↳ Each user can launch a new VW. Diff set of machines
→ Increase / Decrease / Destroy without affecting DB.

Local caching, Stragglers

select * from users
where user.name starts with "C"

Read dominated

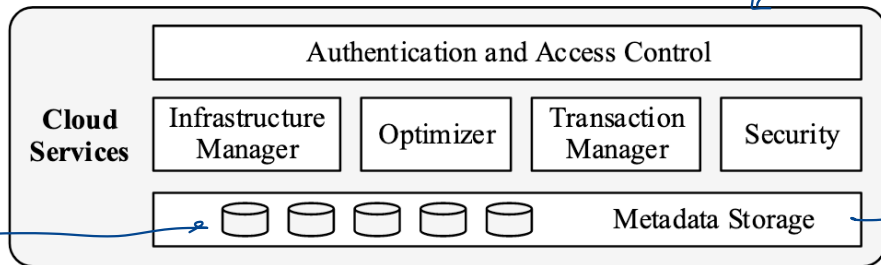
Work Stealing to mitigate stragglers



CLOUD SERVICES

Redis / Memcached.

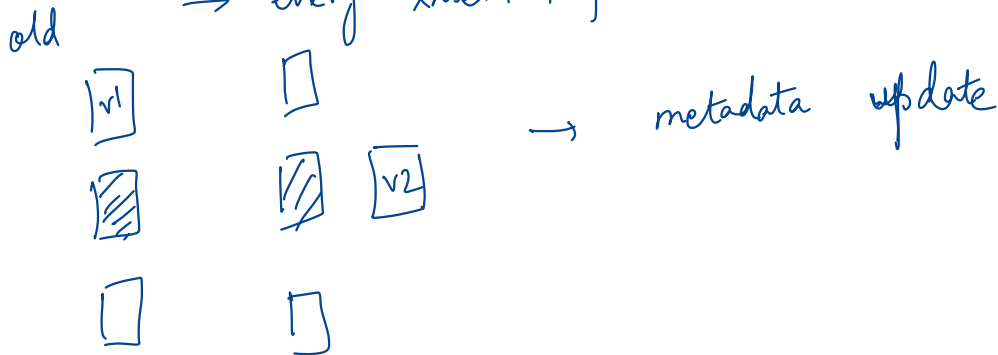
Fast KV
Store



which S3 files
belong to
which table

Concurrency Control → MVCC

Tables have versions
→ every insert / update creates new version



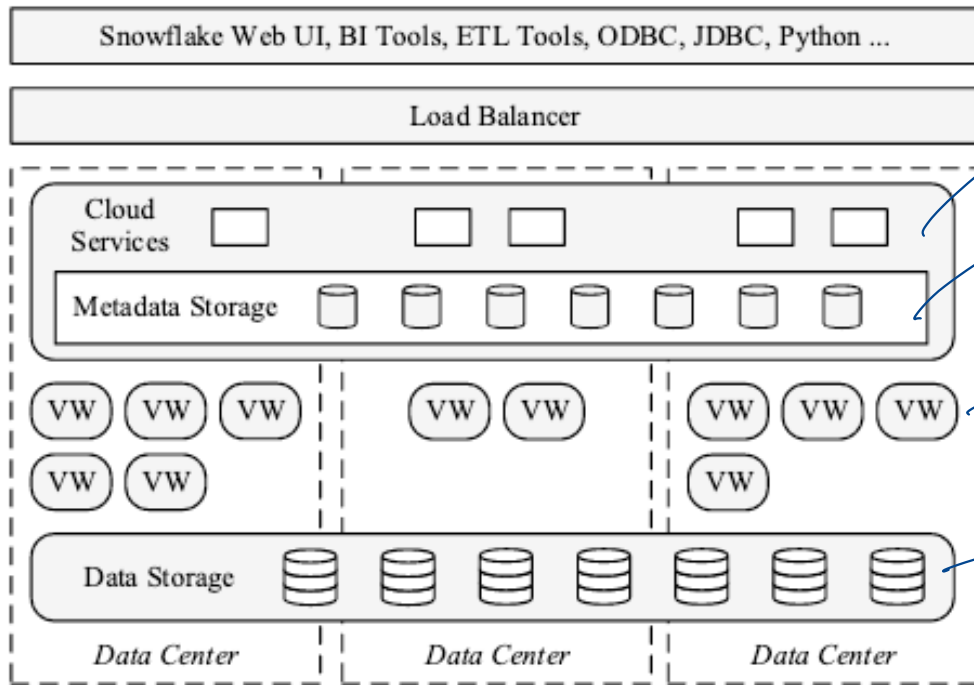
Pruning

↳ limit number of files
read during query
execution

→ min, max for every
col. in the header of
every file

FAULT TOLERANCE

general: design approach
to separate stateful
parts of
your
stack



stateless service

replicated across
DCs.

on failure, retry the
query.

replicated

SEMI STRUCTURED DATA

JSON data

value

```
{  
  first_name: "john",  
  last_name: "doe",  
  order_id: "1234",  
}
```

Integer type

```
{  
  first_name: "bucky",  
  last_name: "badger",  
  order_id: "52342",  
  order_date: "3/3/2020",  
}
```

↳ is not present in all the records.

Extraction, Flattening operations

select value . order_id from
table

Infer types, Pruning

- Type inference on every file
- min/max values for pruning

TIME TRAVEL?

```
SELECT * FROM my_table AT(TIMESTAMP =>
    'Mon, 01 May 2015 16:20:00 -0700'::timestamp);
SELECT * FROM my_table AT(OFFSET => -60*5); -- 5 min ago
SELECT * FROM my_table BEFORE(STATEMENT =>
    '8e5d0ca9-005e-44e6-b858-a8f5b37c5726');
```

Multiple versions of table (MVCC)

Undo accidental deletes

Cheap to clone / snapshot a table

SUMMARY, TAKEAWAYS

Snowflake

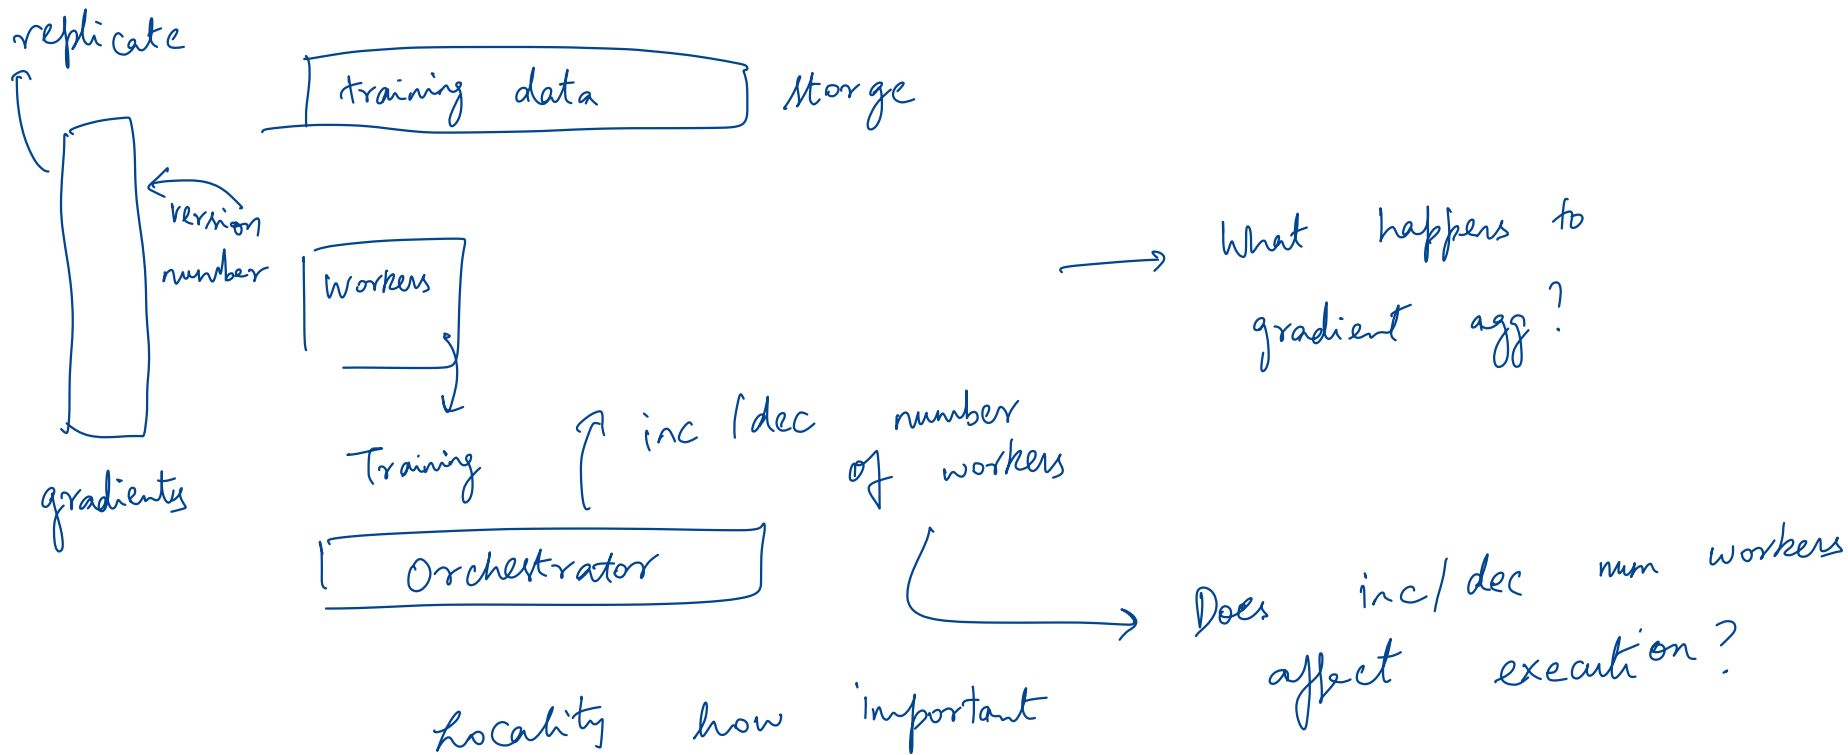
- Cloud computing → Elastic data warehouse
- Key idea: Separation of compute and storage!

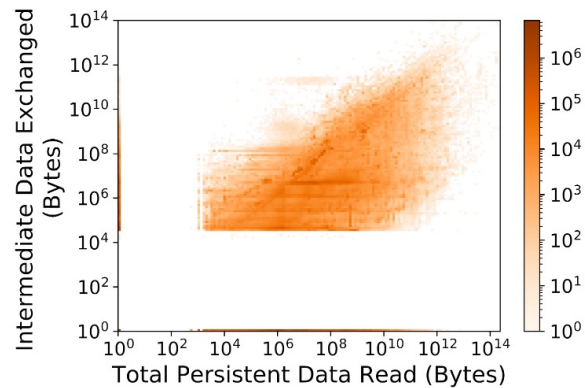
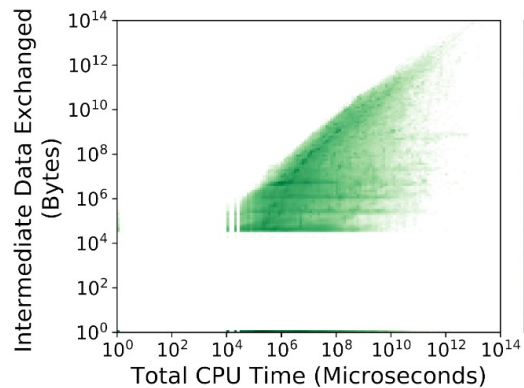
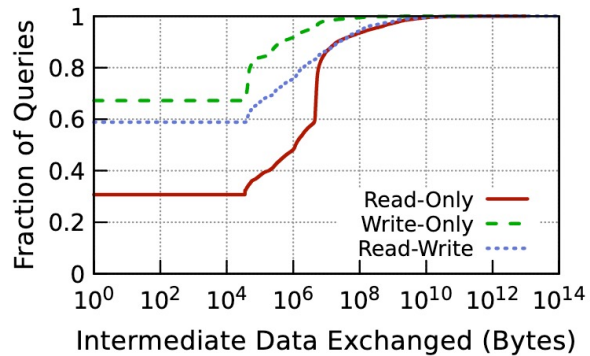
- Hybrid columnar storage format
- Elastic compute with virtual warehouses
- Pruning, semi-structured optimizations, fault tolerant

DISCUSSION

<https://forms.gle/Trfe62jEpI ZUocJk6>

We see how Snowflake leads to the design of an elastic data warehouse. If we were to similarly design an Elastic PyTorch for training how would the design look? What are some design trade-offs compared to existing PyTorch?





NEXT STEPS

Next class: Midterm!