

# Alluxio: Open Source Data Orchestration for Analytics and AI in the Cloud

Haoyuan (H.Y.) Li | Founder, Chairman & CTO | [haoyuan@alluxio.com](mailto:haoyuan@alluxio.com)


2019-11-18 @ PDSW 2019



# The Alluxio Story

–amplab  
2013

Originated as Tachyon project, at the UC Berkley's AMP Lab by then Ph.D. student & now Alluxio CTO, Haoyuan (H.Y.) Li.

 ALLUXIO  
2015

Open Source project established & company to commercialize Alluxio founded

ANDREESSEN  
HOROWITZ

Goal: **Orchestrate Data at Memory Speed for the Cloud** for data driven apps such as Big Data Analytics, ML and AI.

CRN

2018




2018



2019

# Early Days Contributors Growth





# Open Source Started From UC Berkeley AMPLab



1000+ contributors &  
growing



4000+ Git Stars

Join the  
conversation on  
**Slack**  
[slackin.alluxio.io](https://slackin.alluxio.io)



Apache 2.0 Licensed



GitHub's Top 100 Most  
Valuable Repositories  
Out of 96 Million



# Companies Running Alluxio ([Learn More](#))

## Financial Services



## Retail & Entertainment



## Data & Analytics Services



## Technology




## Consumer



## Telco & Media

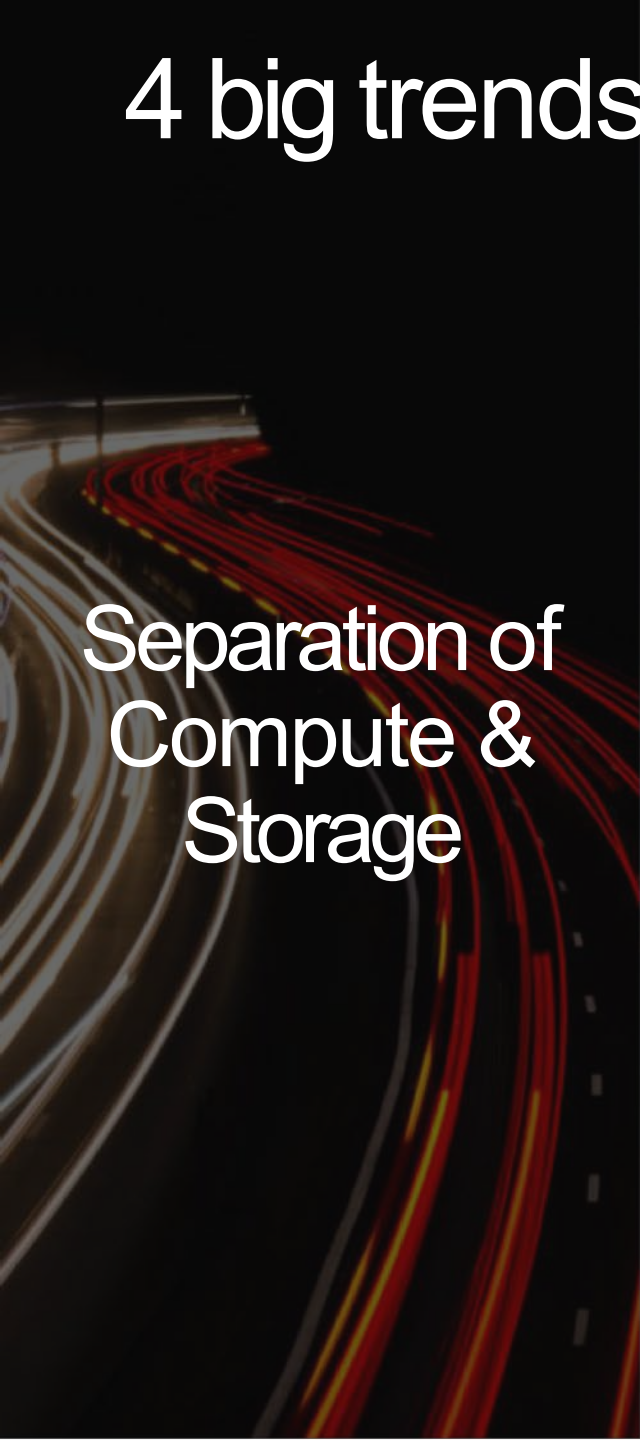


## Travel & Transportation





# 4 big trends driving the need for a new architecture



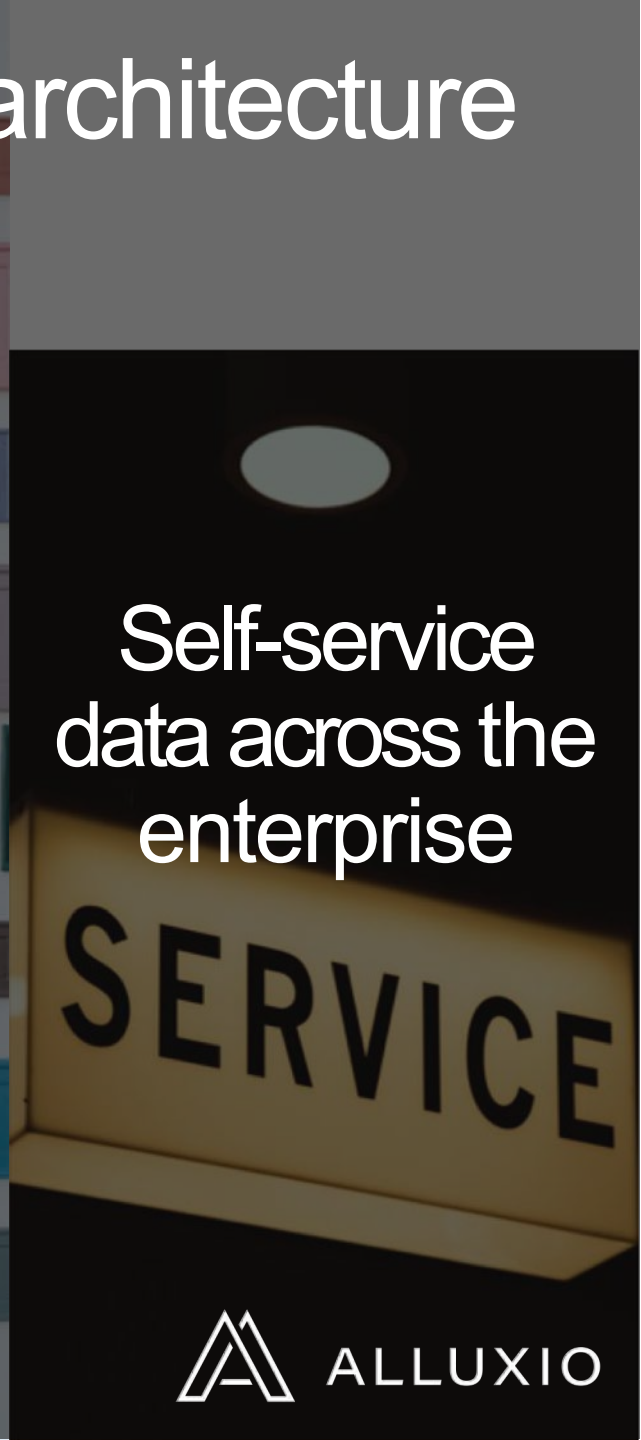
Separation of  
Compute &  
Storage



Hybrid – Multi  
cloud  
environments



Rise  
of the object  
store



Self-service  
data across the  
enterprise

# Data Ecosystem - Beta

COMPUTE



STORAGE

# Data Ecosystem 1.0

COMPUTE



STORAGE





# Data Ecosystem 1.0 – The Challenges



Complex





Low performance



Expensive

COMPUTE



STORAGE



# Data stack journey and innovation paths

## Co-located


Co-located  
compute & HDFS  
on the same cluster



- Typically compute-bound clusters over 100% capacity
- Compute & I/O need to be scaled together even when not needed

## Disaggregated

Disaggregated  
compute & HDFS  
on the same cluster



- Compute & I/O can be scaled independently but I/O still needed on HDFS which is expensive

## Support more frameworks

Support Presto, Spark  
across DCs without  
app changes

## HDFS for Hybrid Cloud

Burst HDFS data in  
the cloud,  
public or private

## Transition to Object store

Enable & accelerate  
big data on  
object stores

# Independent scaling of compute & storage



Java File API

HDFS Interface

S3 Interface

POSIX Interface

REST API



ALLUXIO

Data Orchestration for the Cloud

HDFS Driver

Swift Driver

S3 Driver

NFS Driver





# APIs to Interact with data in Alluxio

Applications have great flexibility to read / write data with many options

## Spark

```
> rdd = sc.textFile("alluxio://localhost:19998/myInput")
```

## Presto

```
CREATE SCHEMA hive.web  
WITH (location = 'alluxio://master:port/my-table/')
```

## POSIX

```
$ cat /mnt/alluxio/myInput
```

## Java

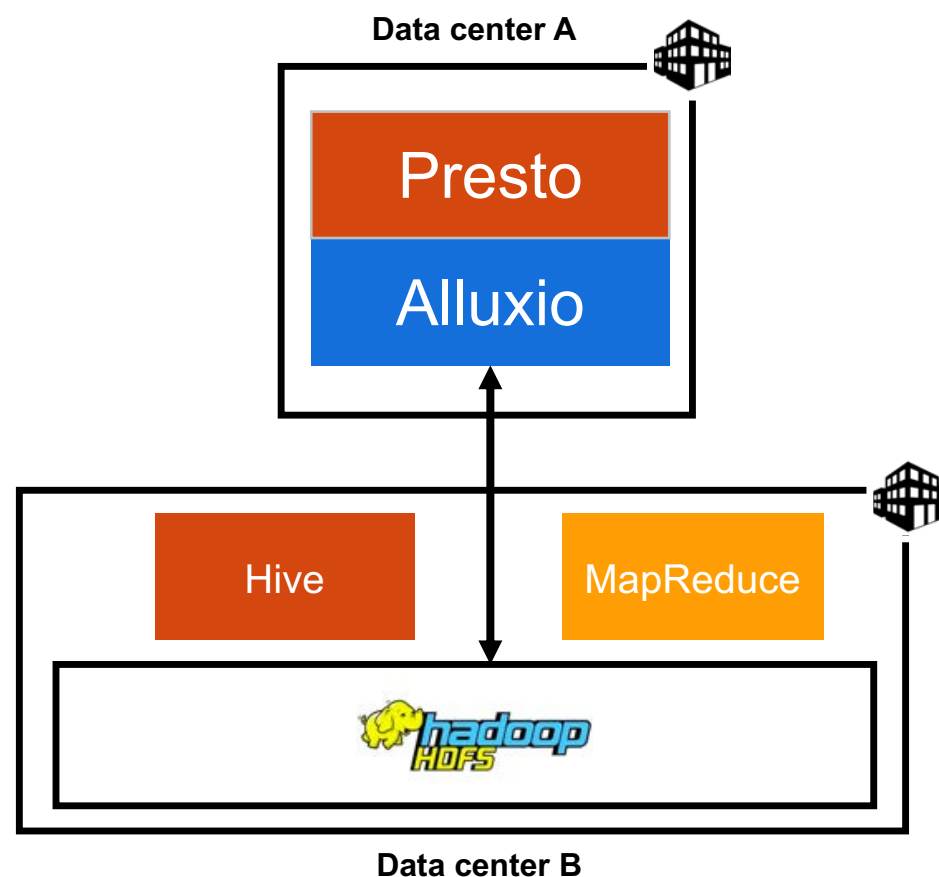
```
FileSystem fs = FileSystem.Factory.get();  
FileInputStream in = fs.openFile(new AlluxioURI("/myInput"));
```

# Challenges with supporting more frameworks across data centers

## Support more frameworks

- Running new frameworks on existing an HDFS cluster can dramatically affect performance of existing workloads
- Orchestrating data to compute clusters in another data center is typically a manual effort and time consuming
- Storing and managing multiple copies of the data becomes expensive

## On-premise satellite compute clusters across data centers



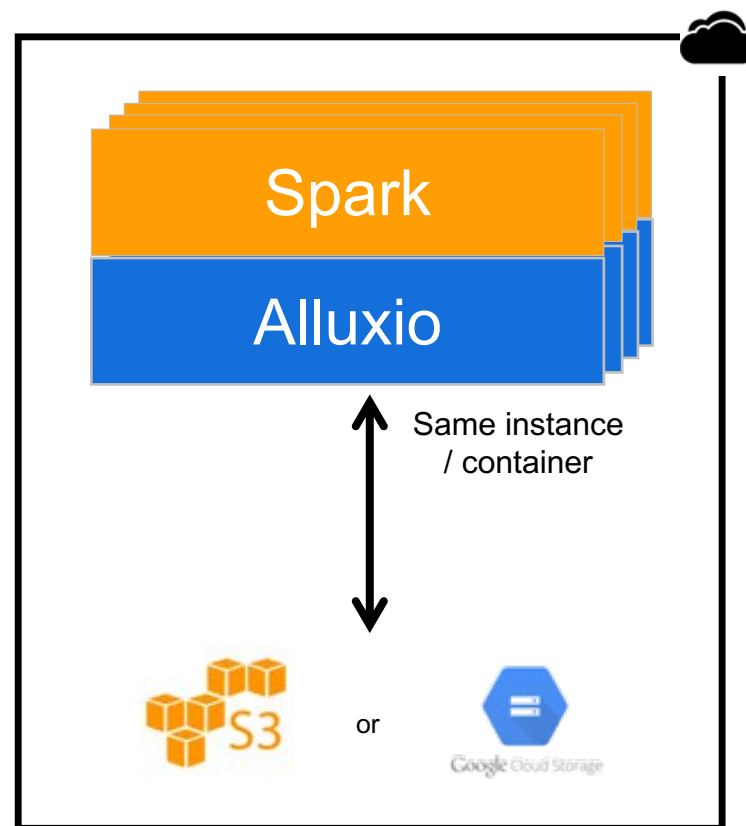


# Challenges with running workloads on cloud storage

## Compute caching for S3 / GCS

- S3 performance is variable and consistent query SLAs are hard to achieve
- S3 metadata operations are expensive making workloads run longer
- S3 egress costs add up making the solution expensive
- S3 is eventually consistent making it hard to predict query results

## Accelerate analytical frameworks on the public cloud

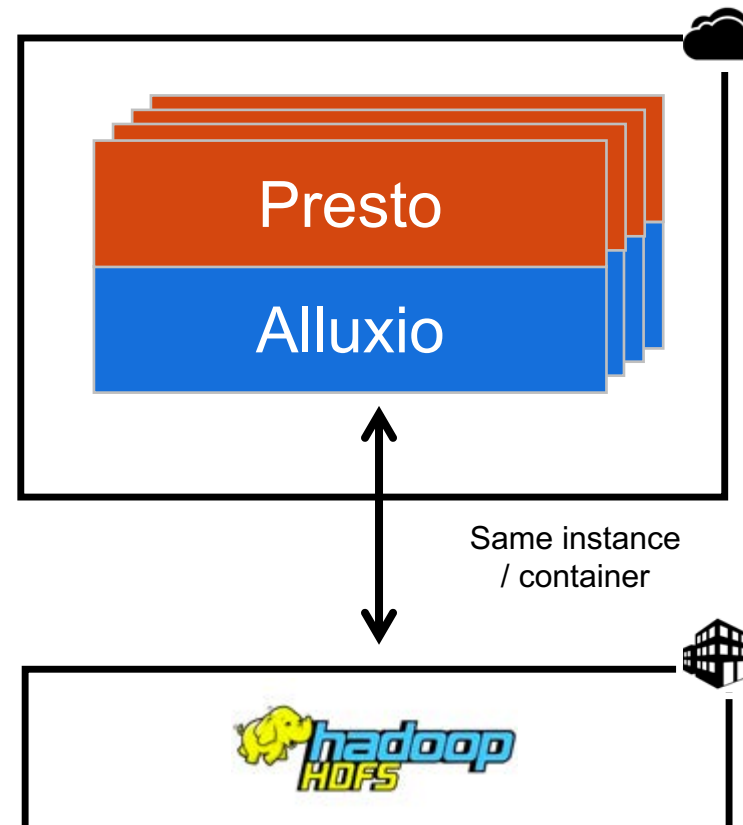


# Challenges with Hybrid Cloud

## HDFS for Hybrid Cloud

- Accessing data over WAN too slow
- Copying data to compute cloud time consuming and complex
- Using another storage system like S3 means expensive application changes
- Using S3 via HDFS connector leads to extremely low performance

## Burst big data workloads in hybrid cloud environments



### Solution Benefits

- Same performance as local
  - Same end-user experience
- 
- 100% of I/O is offloaded

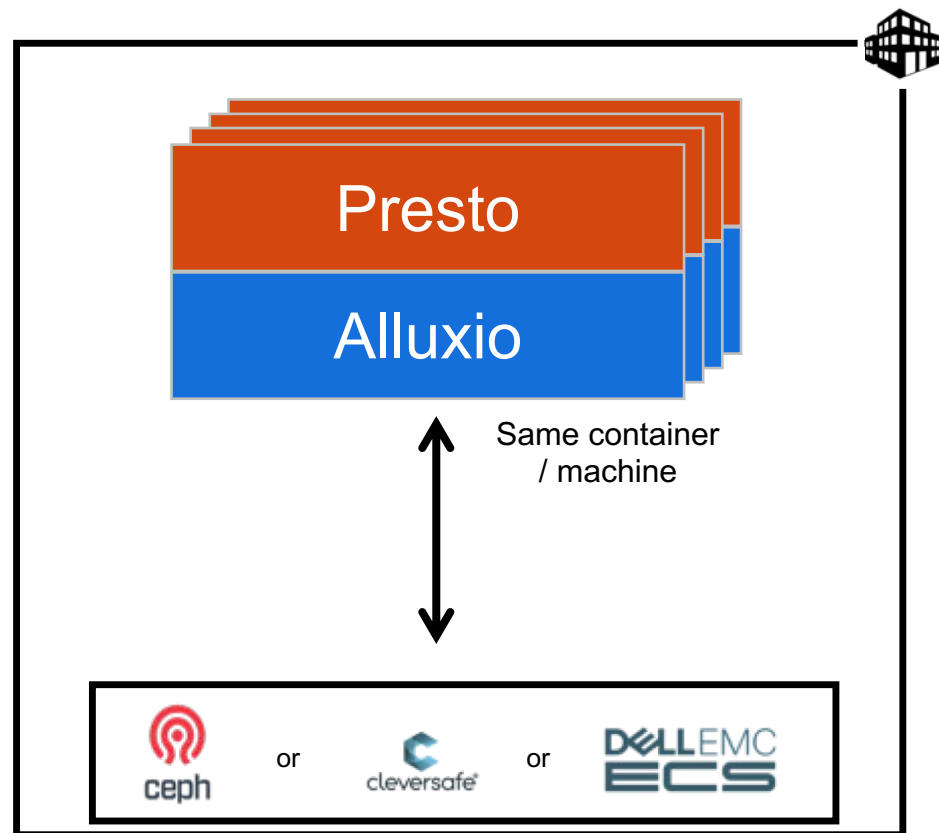


# Challenges running Big Data on Object Stores & Alluxio Solution

## Transition to Object store

- Object stores performance for big data workloads can be very poor
- No native support for popular frameworks
- Expensive metadata operations reduce performance even more
- No support for hybrid environments directly

## Dramatically speed-up big data on object stores on premise




## Solution Benefits

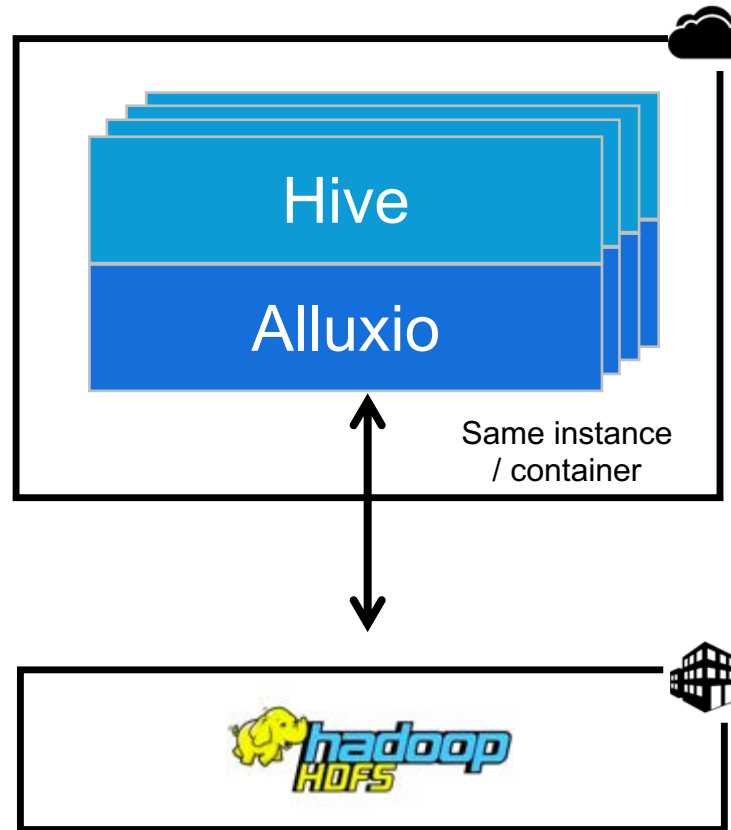
- Same performance as HDFS
- Uses HDFS APIs
- Same end-user experience
- Storage at fraction of the cost of HDFS

# Use Cases Alluxio Enables

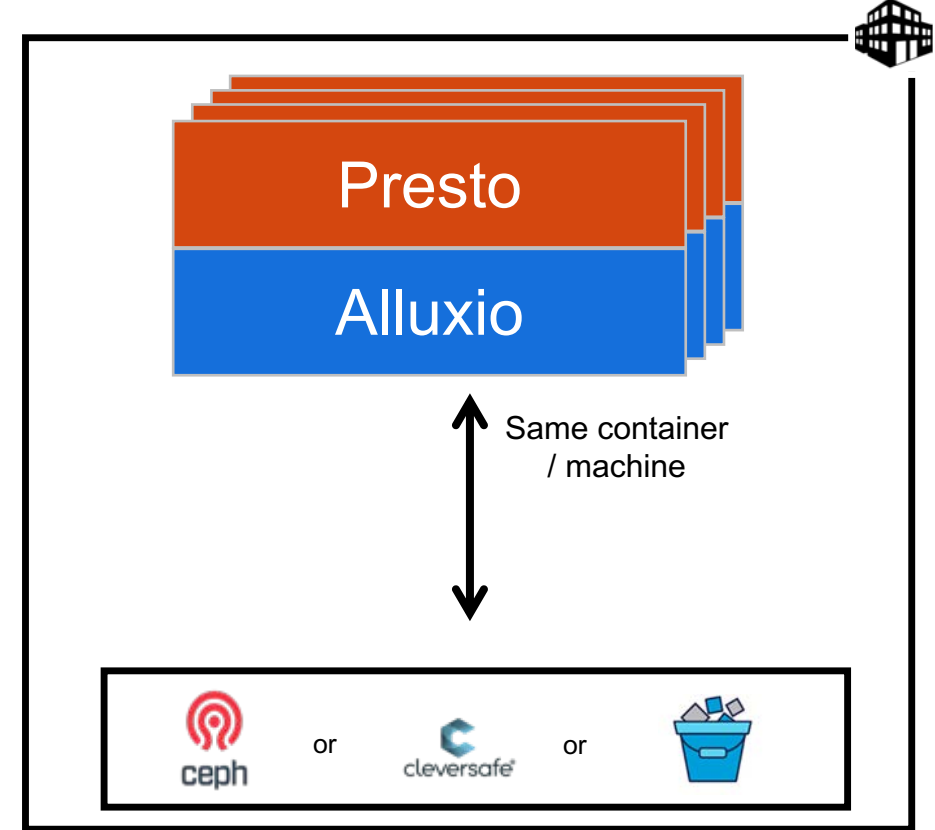
Accelerate big data frameworks  
on the public cloud



Burst big data workloads in  
hybrid cloud environments

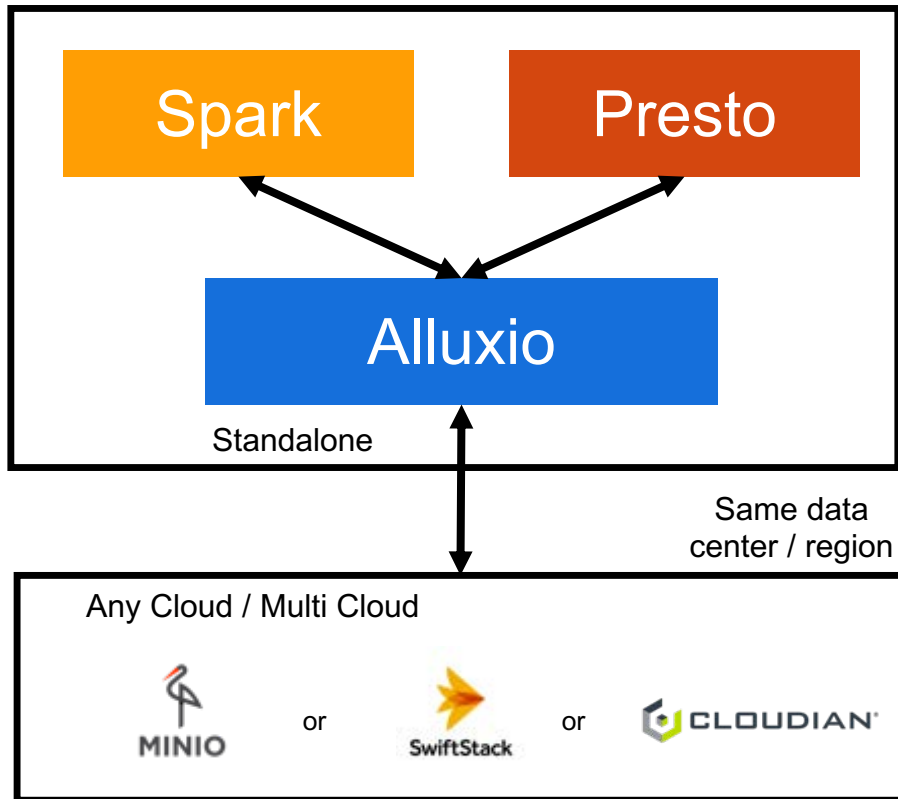


Dramatically speed-up big data  
on object stores on premise

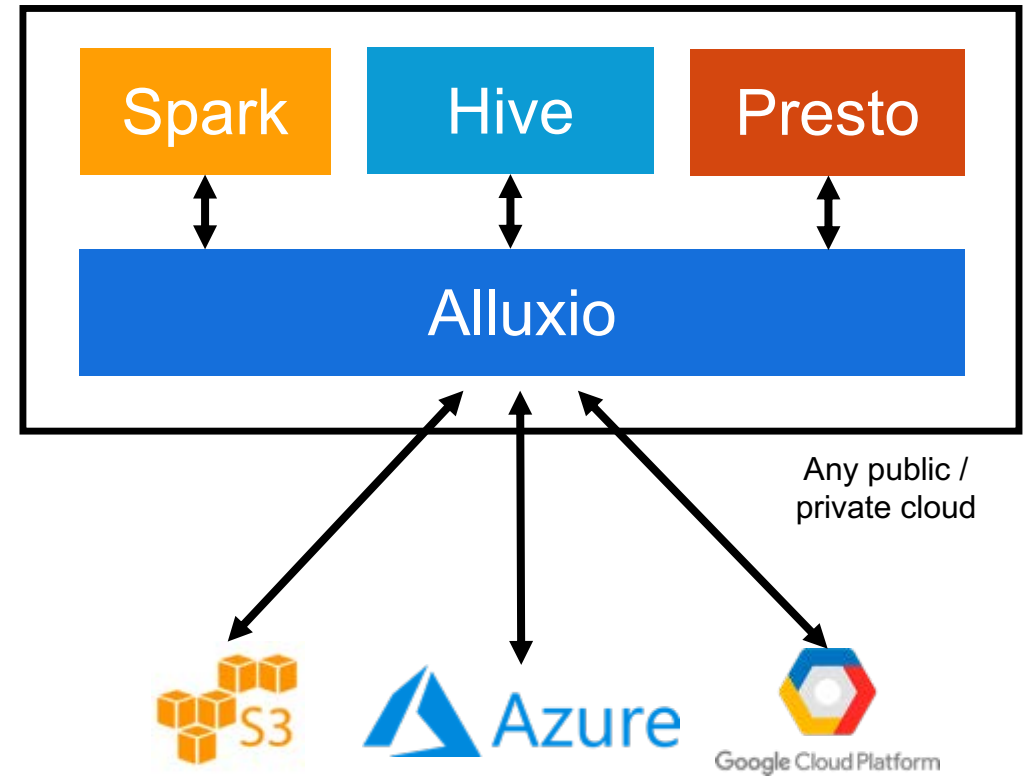




# Advanced Use Cases



**Enable big data on object stores across single or multiple clouds**



**Orchestrate data frameworks on the public cloud**

# Alluxio – Key innovations

## **Data Locality** with Intelligent Multi-tiering

Accelerate big data  
workloads with transparent  
tiered local data

## **Data Accessibility** for popular APIs & API translation

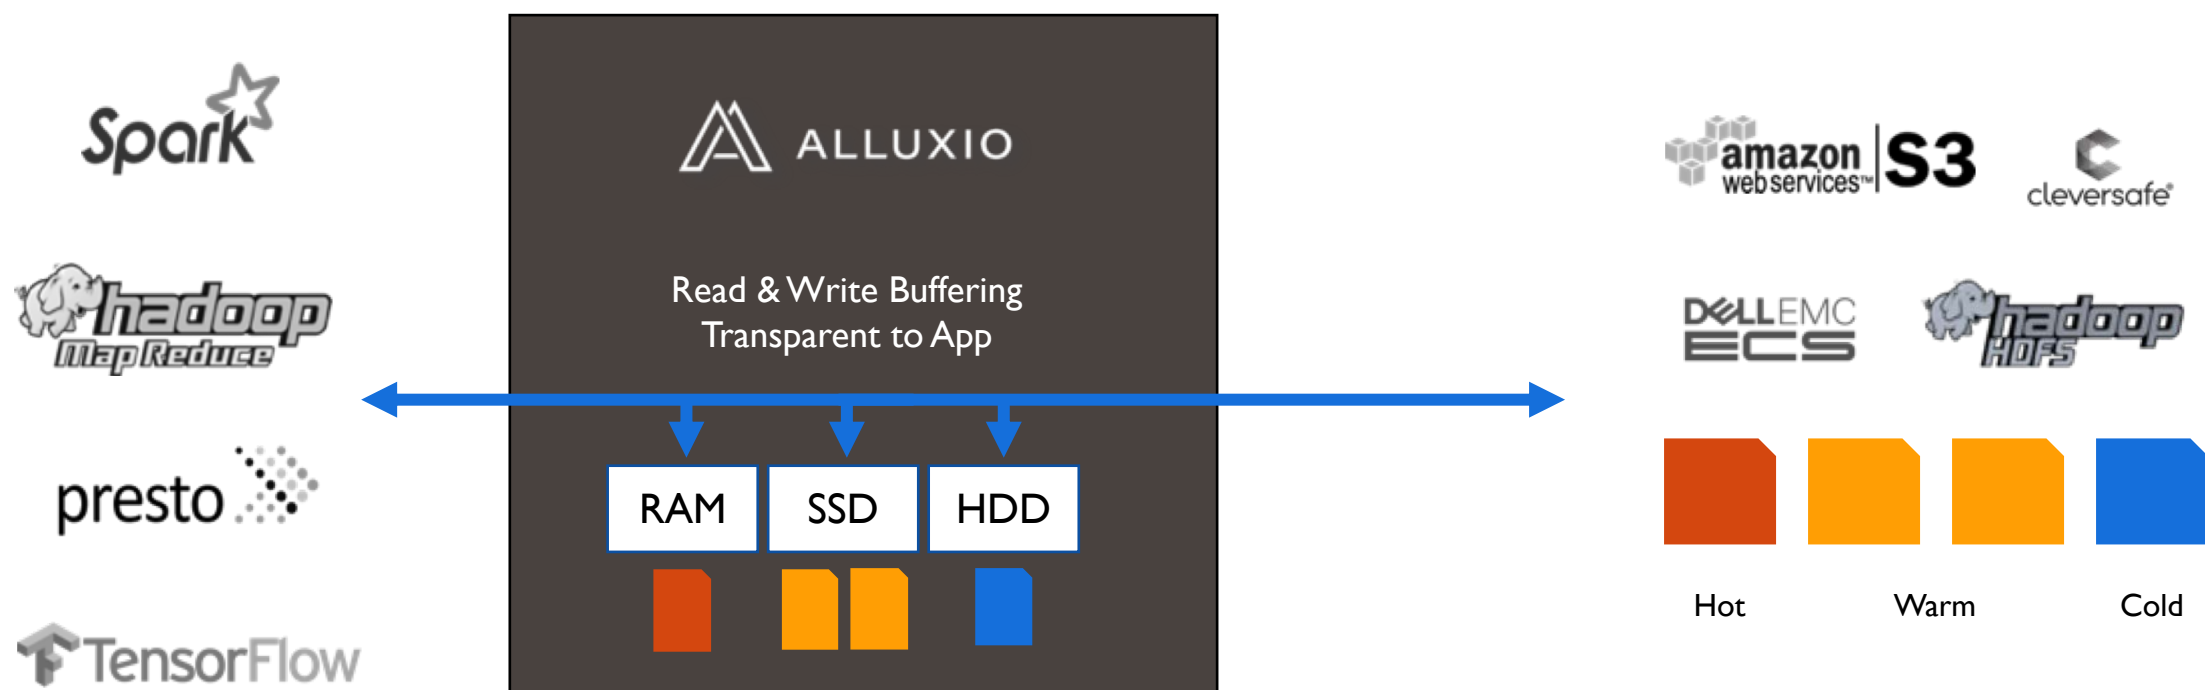
Run Spark, Hive, Presto, ML  
workloads on your data  
located anywhere

## **Data Elasticity** with a unified namespace

Abstract data silos & storage  
systems to independently scale  
data on-demand with compute

# Data Locality with Intelligent Multi-tiering

Local performance from remote data using multi-tier storage





# Data Accessibility via popular APIs and API Translation

Convert from Client-side Interface to native Storage Interface



Java File API

HDFS Interface

S3 Interface

POSIX Interface

REST API



HDFS Driver

S3 Driver

Swift Driver

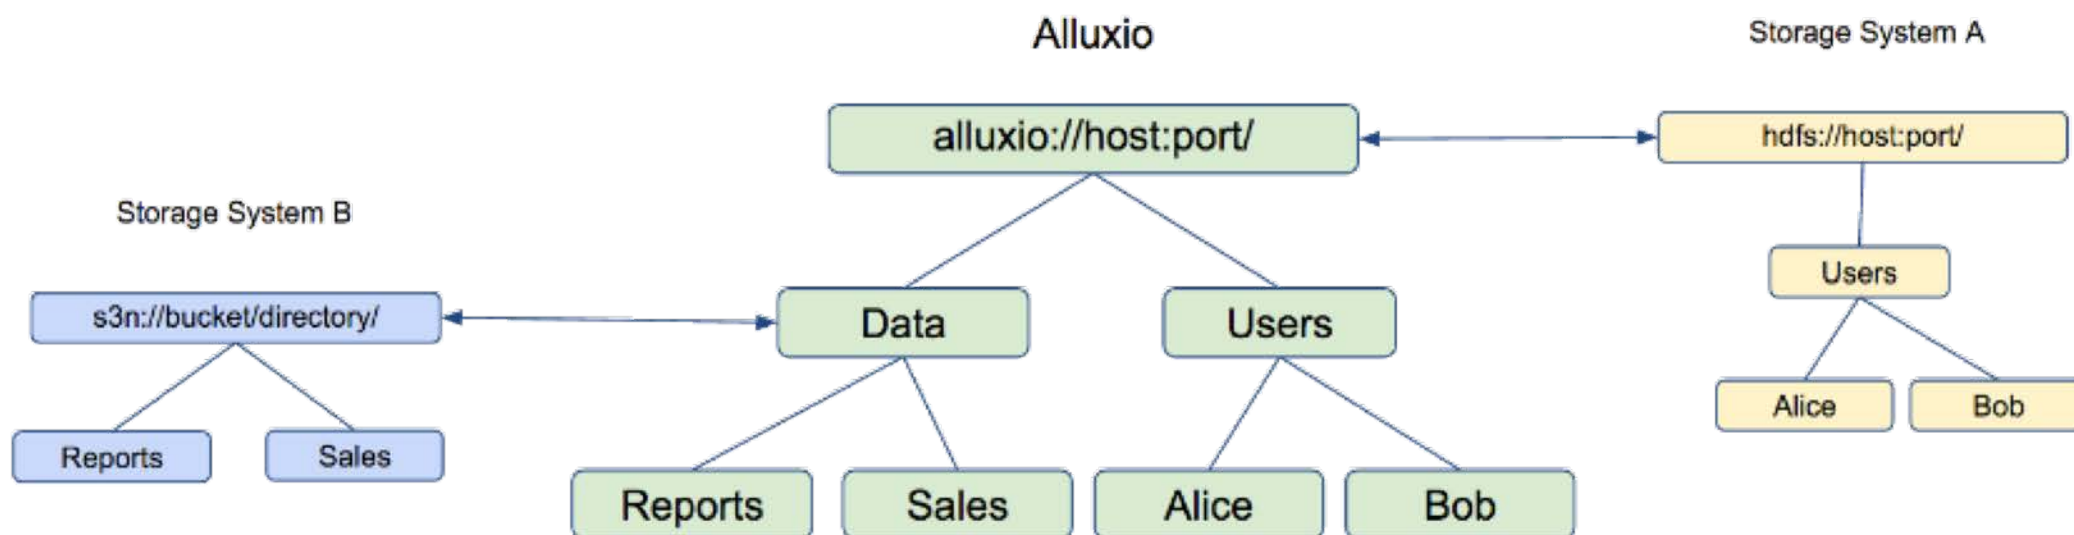
NFS Driver



# Data Elasticity via Unified Namespace

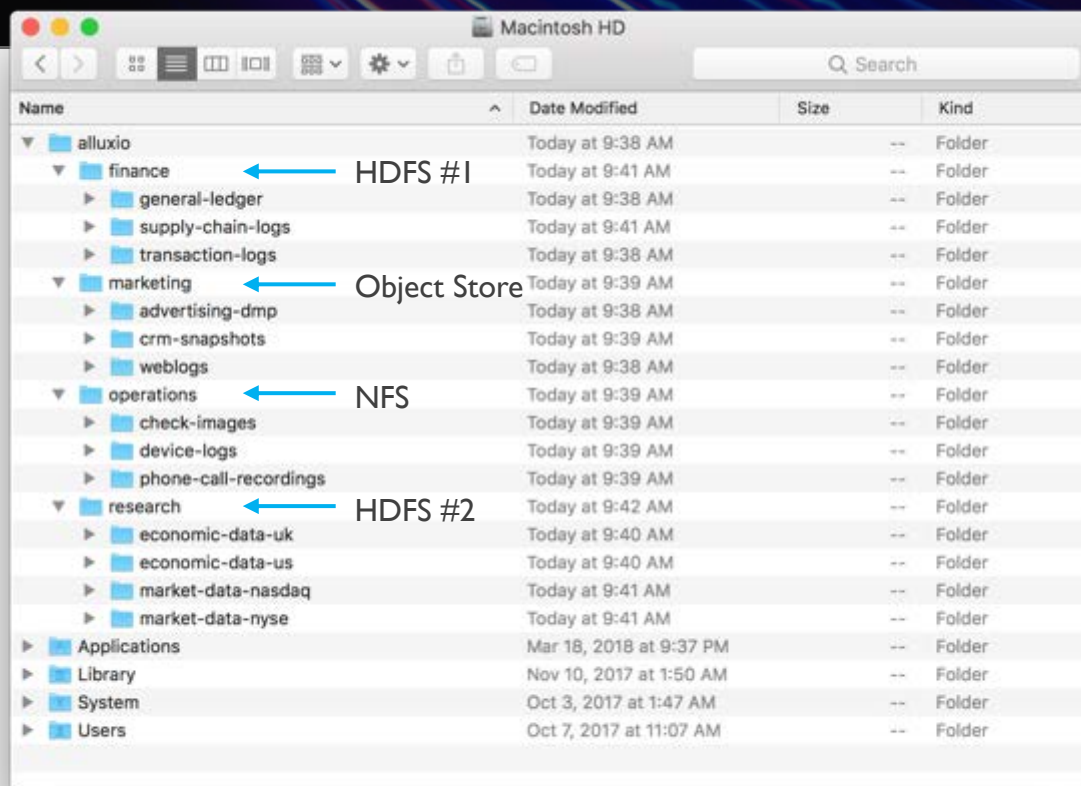
Enables effective data management across different Under Store

- Uses Mounting with Transparent Naming



# Unified Namespace: Global Data Accessibility

Transparent access to understorage makes all enterprise data available locally



## SUPPORTS

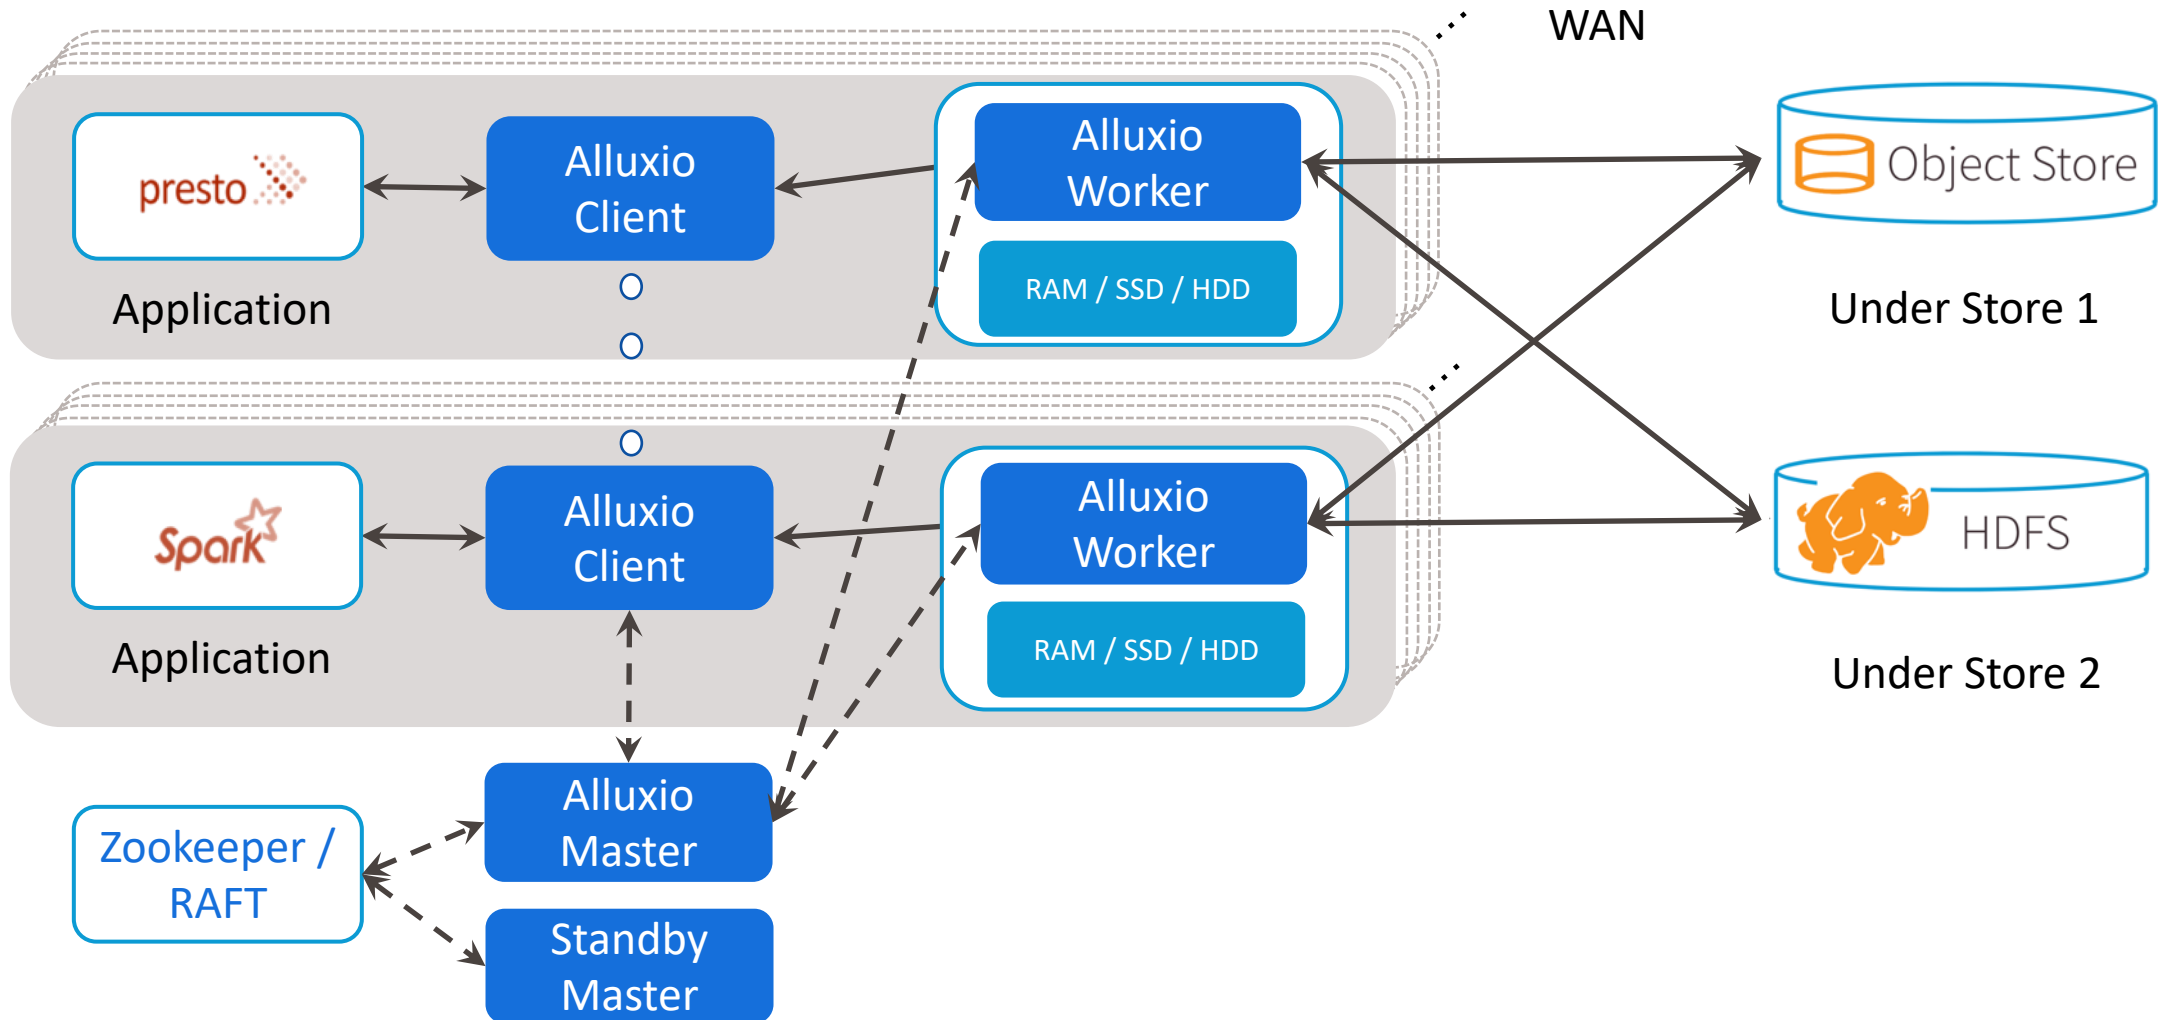
- HDFS
- NFS
- OpenStack
- Ceph
- Amazon S3
- Azure
- Google Cloud

## IT OPS FRIENDLY

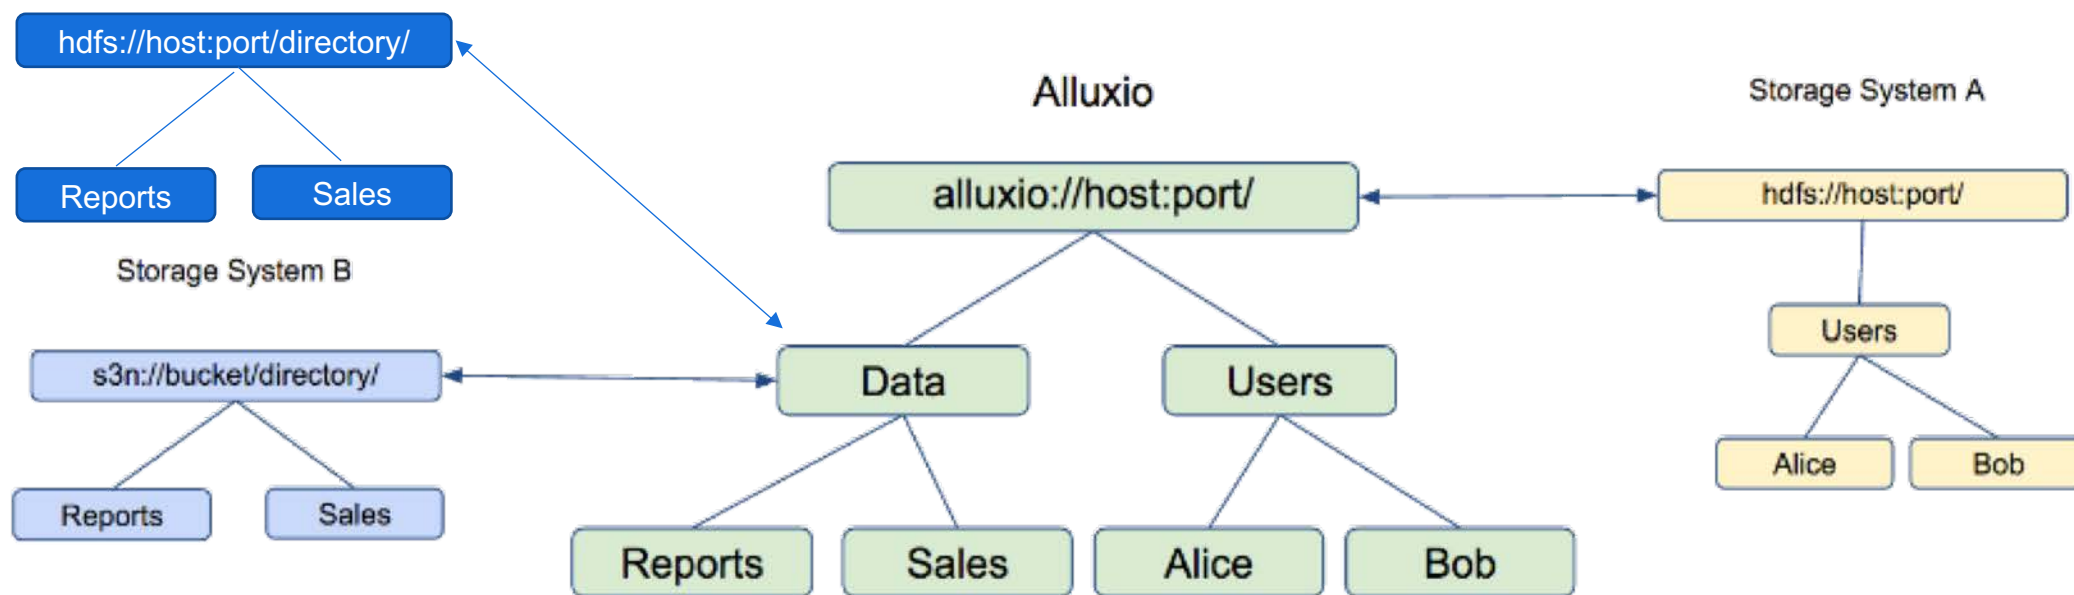
- Storage mounted into Alluxio by central IT
- Security in Alluxio mirrors source data
- Authentication through LDAP/AD
- Wireline encryption



# Alluxio Reference Architecture



# Policy Driven under File System Migration



# Research Directions

---

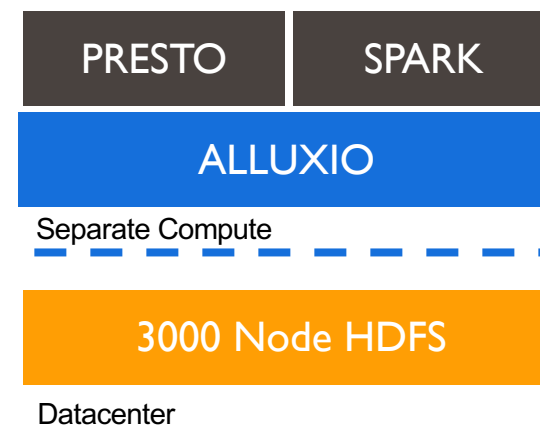
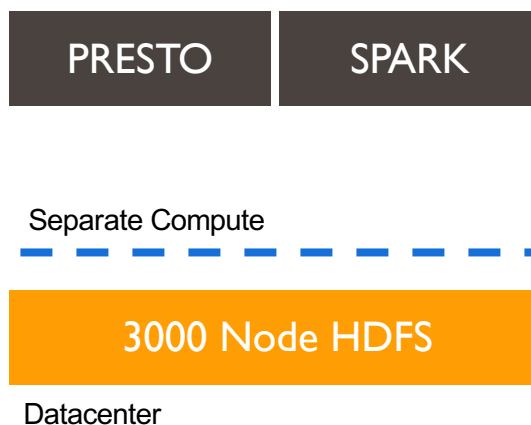
- ▲ Machine-learning based Data Orchestration Policies
- ▲ Scalable and High-performance File System Metadata service
- ▲ Optimization for in-memory data partition / format
- ▲ Cross-layer optimization for distributed compute and storage systems



# JD.com | Performance Use Case in DC



\$70B e-commerce retailer



## Project:

- Offload HDFS with separate clusters of Presto and Spark

## Problem:

- HDFS cluster is compute and network bound
- Performance is inconsistent

## Alluxio solution:

- Alluxio offloads the network I/O as well as the compute

## Result:

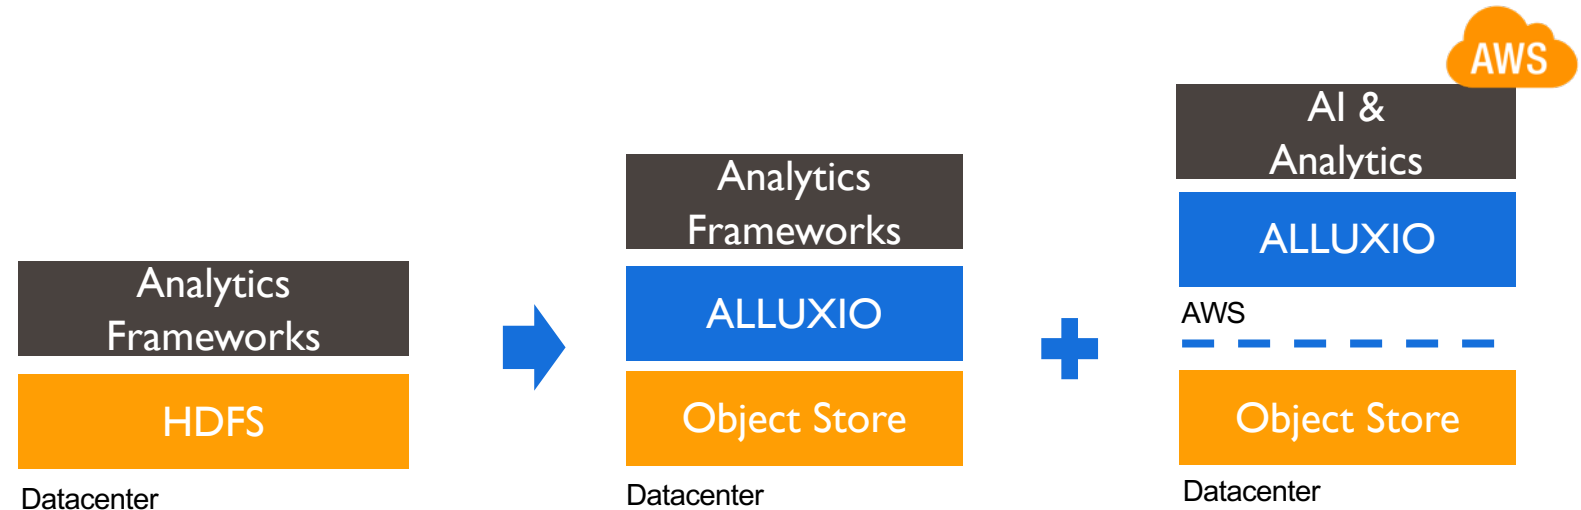
- Teams can run additional workloads without taxing the existing HDFS cluster



# DBS Bank | Performance & Hybrid



*Largest bank in Southeast Asia*



## Initial Project:

- Digital Bank Initiative
- Solve scaling challenges by separating compute and using object storage

## Problem:

- Coupled systems were not flexible to scale

## Alluxio solution:


1. Alluxio provides intelligent caching layer for object storage
2. Burst workloads to hybrid cloud

## Result:

- Enables data on-demand, Alluxio now considered mature layer in stack



# Walmart | Performance Use Case in Cloud



## Project:

- Utilize Presto for interactive queries on cloud object store compute

## Problem:

- Low performance of queries too slow to be usable
- Inconsistent performance of queries

## Alluxio solution:

- Alluxio provides intelligent distributed caching layer for object storage

## Result:

- High performance queries
- Consistent performance
- Interactive query performance for analysts

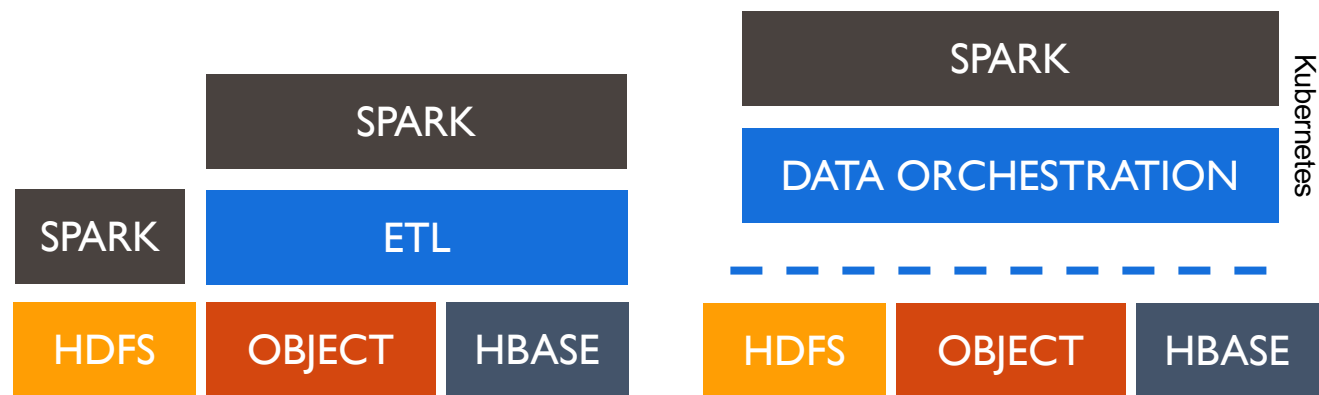




# China Unicom

Leading Chinese Telco serving 320 million subscribers

## Use case | Data orchestration for agility




- Single namespace to access & address all data
- Data local to compute accelerates workloads

# NetEase Games

Leading Online Game Company in China

## Use Case | On-premise Caching for Presto



- Large query variance during peak hours before
- Alluxio brings data local to Presto to reduce the latency during peak hours

<https://www.alluxio.io/blog/presto-on-alluxio-how-netease-games-leveraged-alluxio-to-boost-ad-hoc-sql-on-hdfs/>

# Next steps - Try it out!

---

- [Getting Started](#)
- [Try 10 Minutes Alluxio & Presto Tutorial on Laptop](#)
- [Try 10 Minutes Alluxio & Presto Tutorial on AWS](#)
- [Spark and Alluxio in 5 minutes](#)

Questions or Suggestions? [Engage with our Community in Slack!](#)





# Questions?

Welcome to join the Alluxio Open Source Community!

[www.alluxio.io](http://www.alluxio.io) | [slackin.alluxio.io](https://slackin.alluxio.io) | [@alluxio](https://twitter.com/alluxio)