

Building a Distributed Data Access Layer for Analytics on Any Cloud

Bin Fan | Founding Engineer & VP Open Source | Alluxio
binfan@alluxio.com

About Me



@binfan



binfan@alluxio.com

The journey to a fragmented data world



More data
generated every day



More people & teams need
access to this data



New storage technologies
created every 3-8 years

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



Big data journey and innovation options for enterprises



Challenges with the transition

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

Support more frameworks

- Copying data to multiple compute clouds time consuming and error prone
- Migrating applications for new storage systems is complex & time consuming
- Storing and managing multiple copies of the data becomes expensive

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

Independent scaling of compute & data



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




Use Cases Data Orchestration 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

FUSE 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

Abstract & orchestrate data across data silos

COMPUTE SPREAD ACROSS MANY DIFFERENT FRAMEWORKS



DATA IN DISPARATE STORAGE SYSTEMS

Demos in Office Hour:

- Spark + Alluxio + S3 & Azure
- TPC-DS on Spark+S3 vs Spark+Alluxio+S3

Interacting with data in Alluxio – variety of APIs

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

Spark

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

Hadoop

```
$ hadoop fs -cat alluxio://localhost:19998/myInput
```


POSIX

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


Java

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

Deployment Approaches




Co-locate Alluxio Workers with Spark for optimal I/O performance



Deploy Alluxio as standalone cluster between Spark and Storage

Alluxio Reference Architecture



Interacting with data in Alluxio – flexible app patterns

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

Reading Data


- From under store
- From a co-located Alluxio node
- From a different Alluxio node

Writing Data

- Write only to Alluxio
- Write only to Under Store
- Write synchronously to Alluxio and Under Store
- Write to Alluxio and asynchronously write to Under Store
- Write to Alluxio and replicate to N other workers
- Write to Alluxio and async write to multiple Under stores


Read data in Alluxio, on same node as client

Memory Speed Read of Data




Read data not in Alluxio

Network / Disk Speed Read of Data




Write data only to Alluxio on same node as client

Memory Speed Write of Data



Write data to Alluxio and Under Store synchronously



Interacting with data in Alluxio – data management

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

Data Management


- Pinning
- Prefetch/free
- Cross storage copy and move operations
- TTL



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



Two Sigma

Fastest growing big hedge fund managing \$46 billion for investors

Use case | Cloud bursting on-premise data



- Compute scales elastically independent of storage
- Faster time to insights with seamless data orchestration
- Accelerated workloads with memory-first data approach

Enterprises moving towards independent compute & storage



Join the Alluxio Open Source Community

www.alluxio.org/slack