

# Object Drives


## A New Architectural Positioning?

Mark Carlson

Toshiba

Principal Engineer, Industry Standards

Co-chair SNIA Object Drive TWG



# What are Object Drives?

Interface changed from SCSI protocol based to IP protocol (TCP/IP, HTTP) based

Channel (FC/SAS/SATA) interconnect moves to Ethernet network

Key/Value semantics (Object store)

Hosted software in some cases



# What is driving the market for these devices?

A number of scale out storage solutions expand by adding identical storage nodes incrementally

- Typically use an Ethernet interface and may be connected directly to the internet.

Open source examples include:

- Hadoop's HDFS
- CEPH
- Swift (OpenStack object storage)
- Kinetic Open Source Project

Commercial examples also exist





# Current Solutions

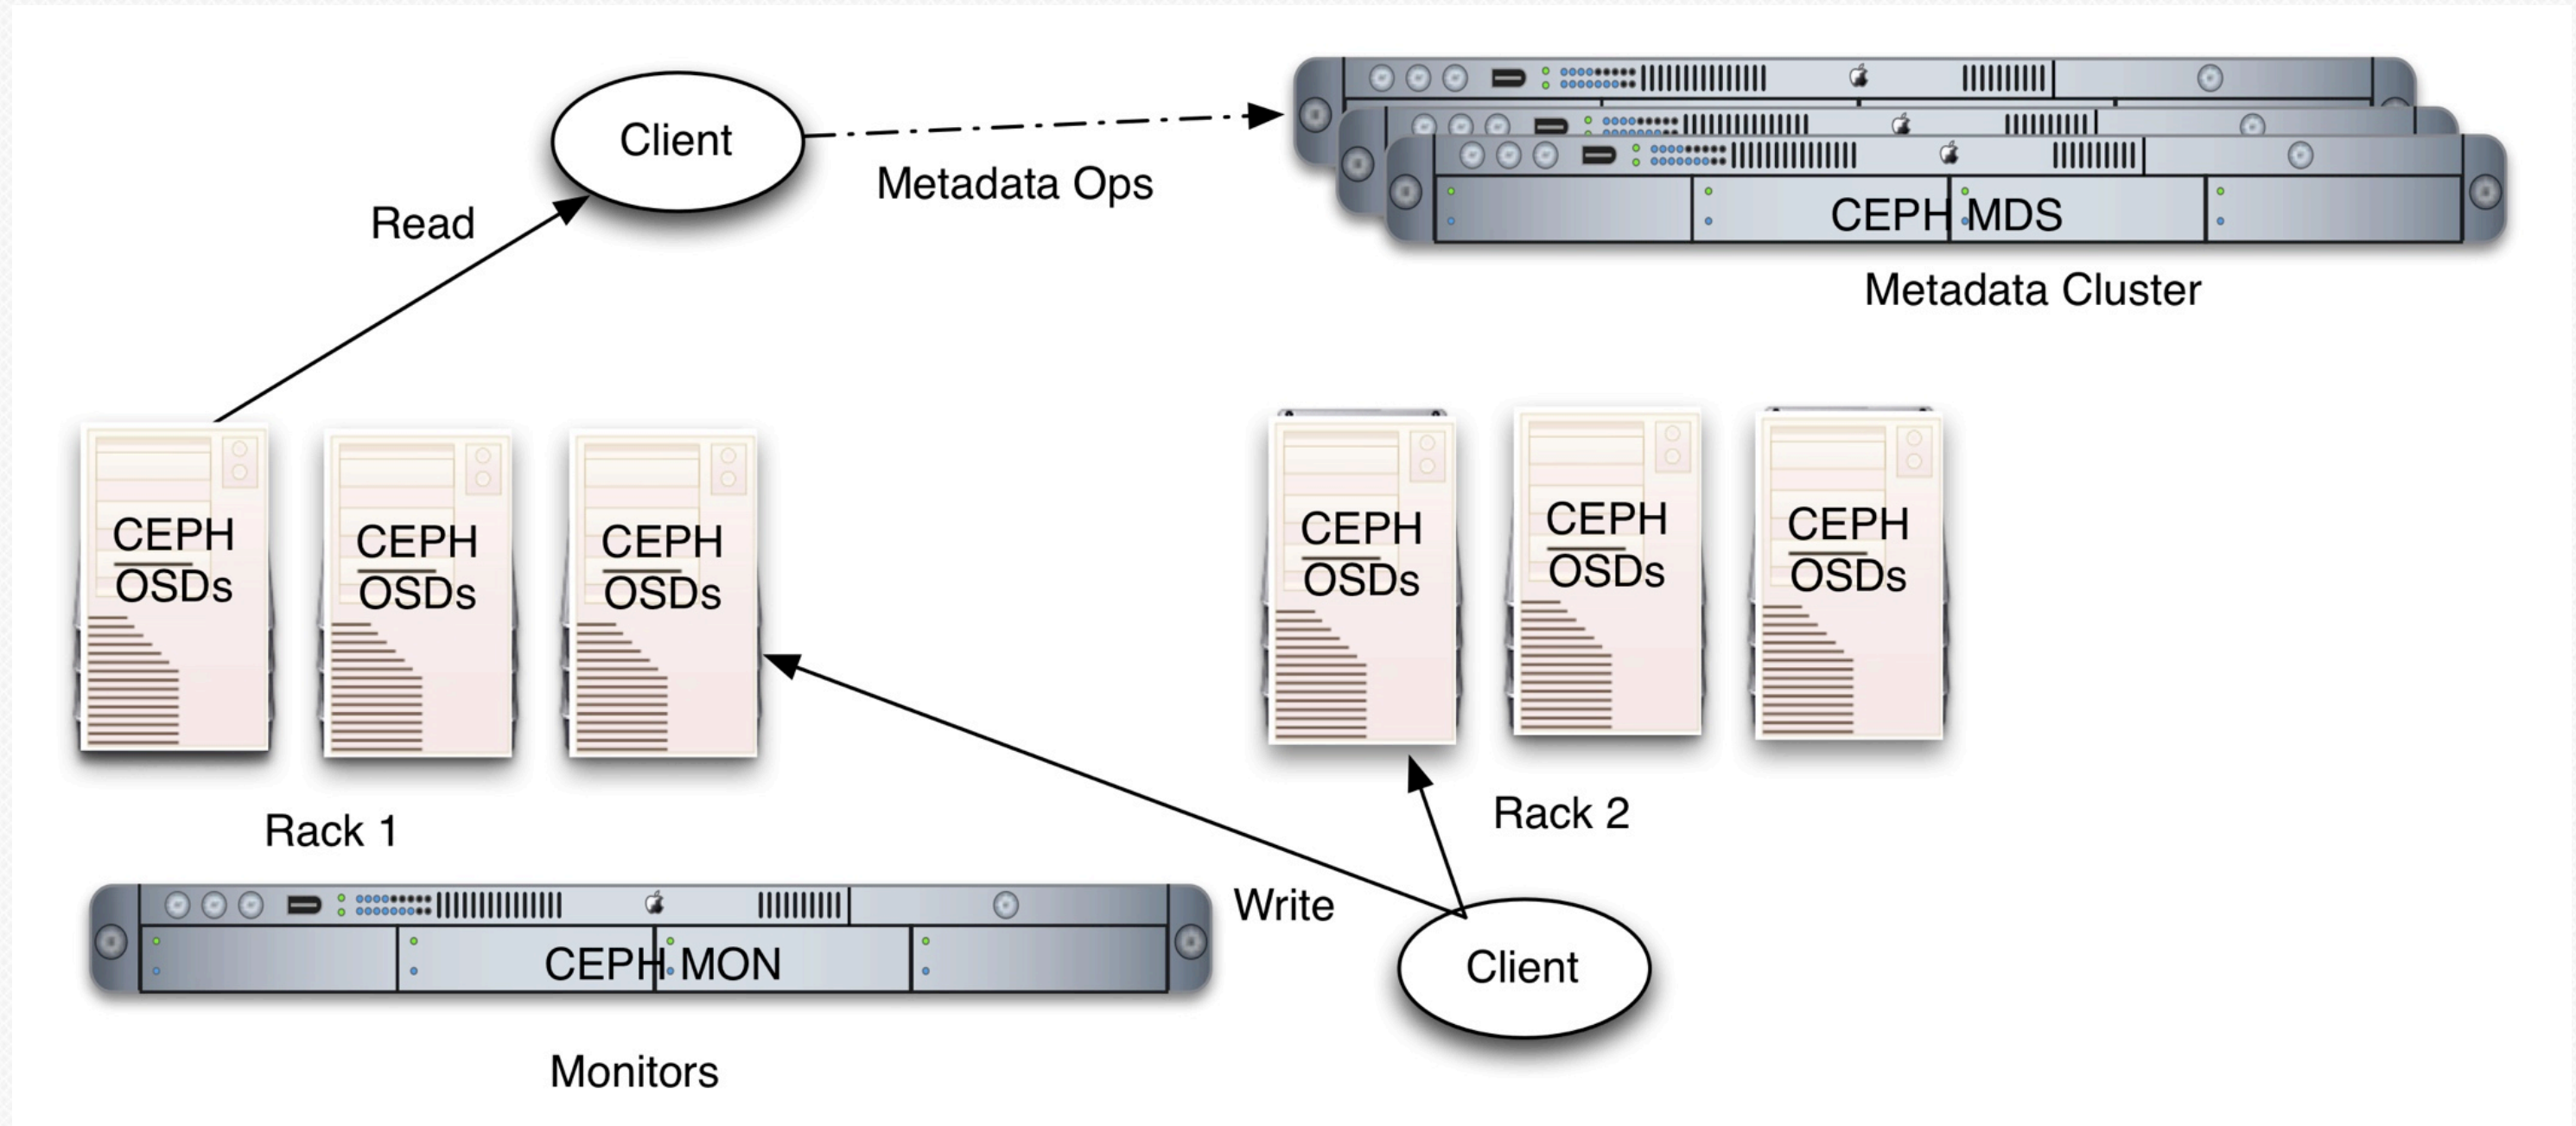
For these solutions, typically a white box server is used for a storage node with DAS storage, CPU, memory, networking

This generalized solution for this specific use case is inefficient, power hungry and adds to management complexity

Although inexpensive, this does not reduce long term ownership costs

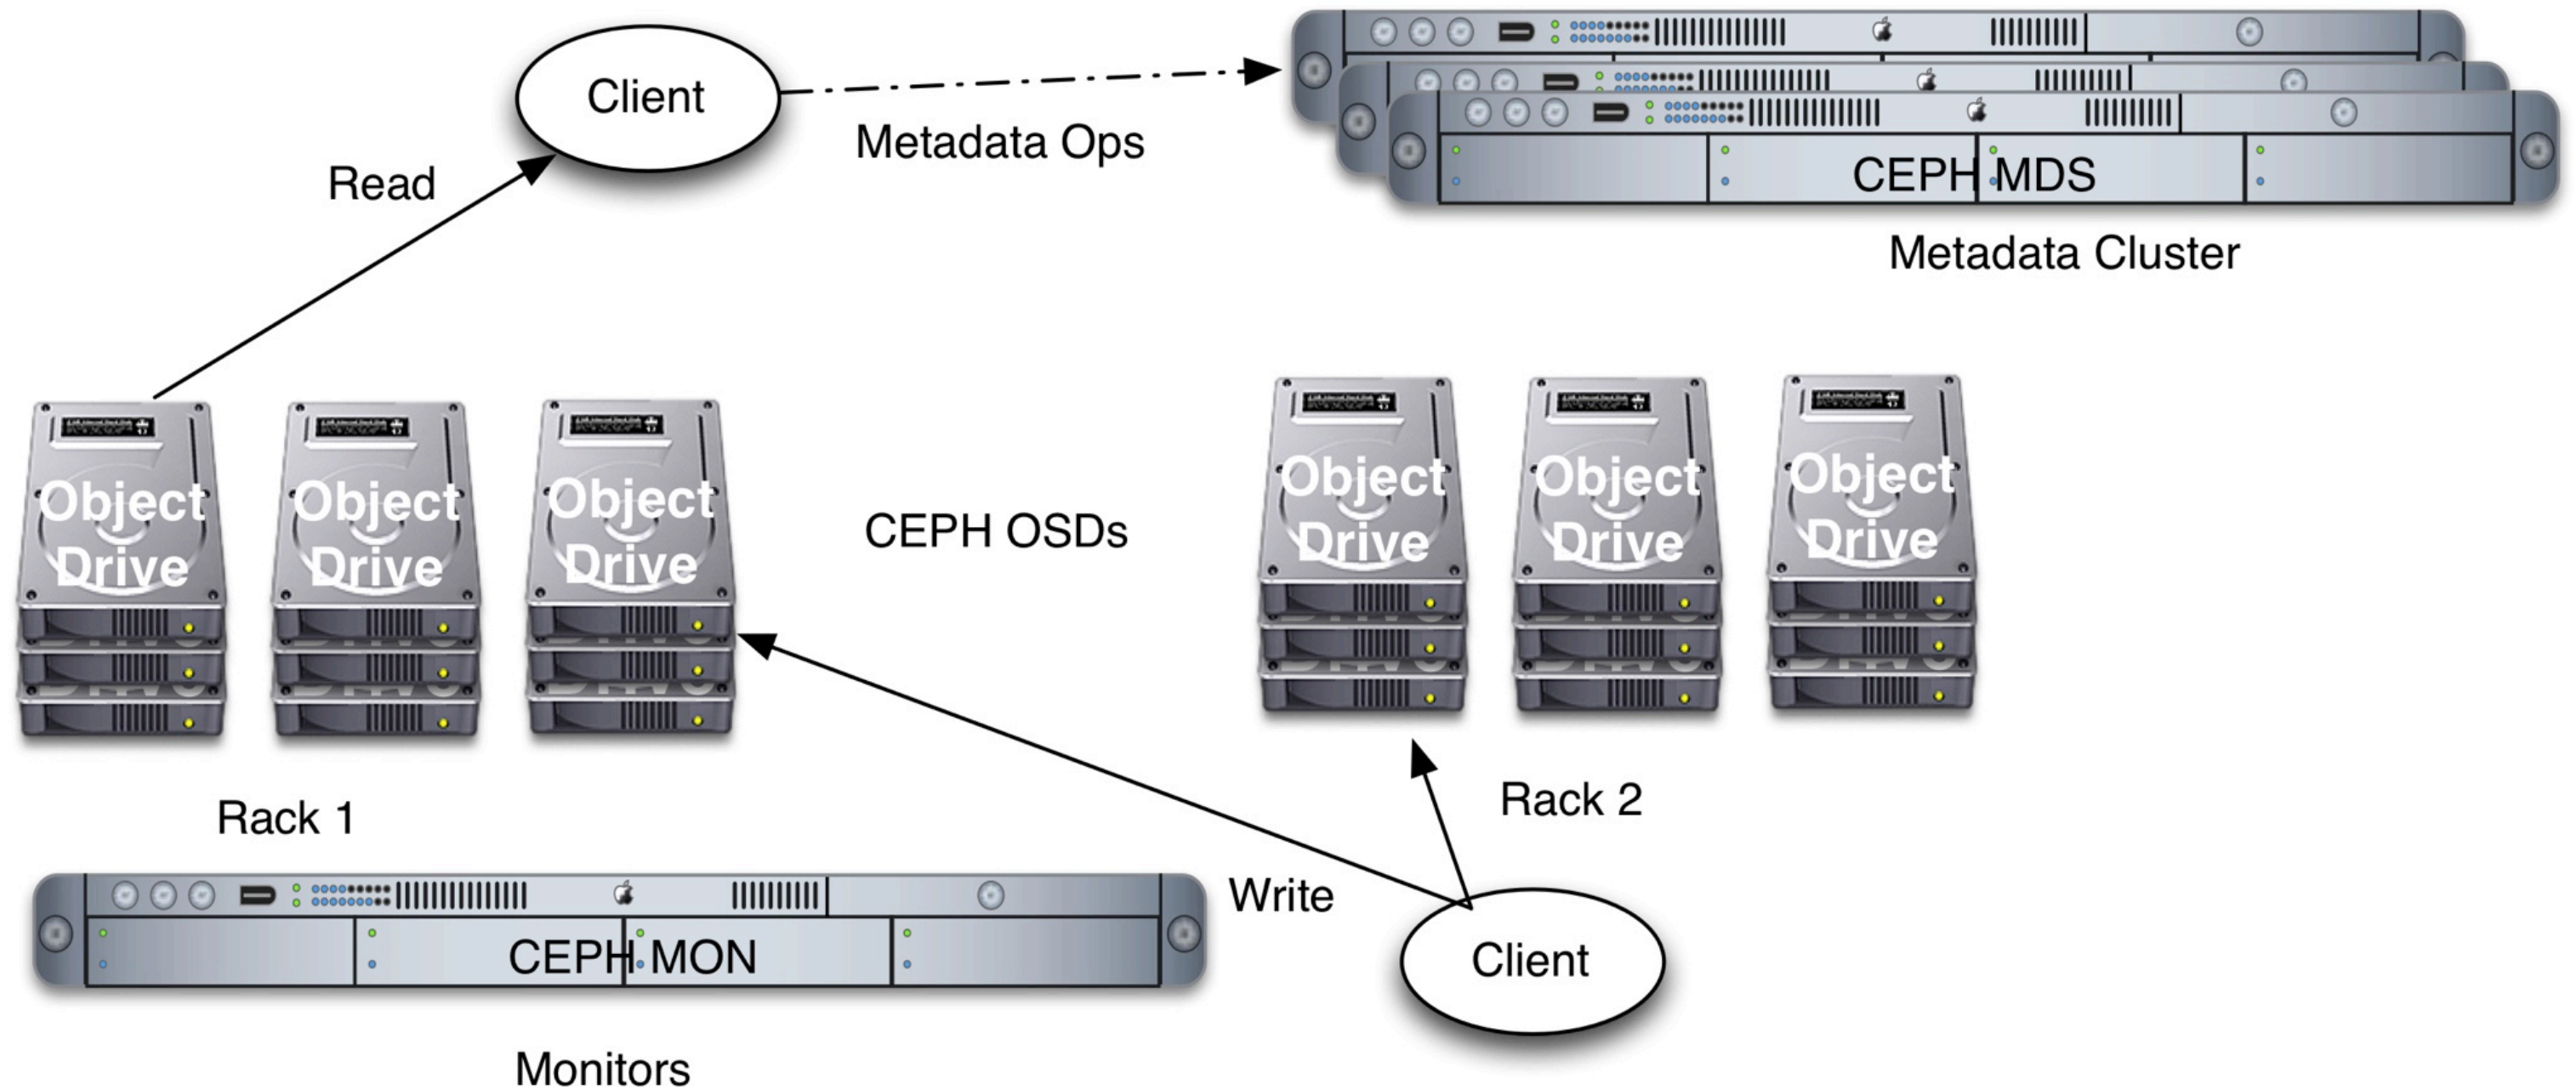


# Example: CEPH Architecture

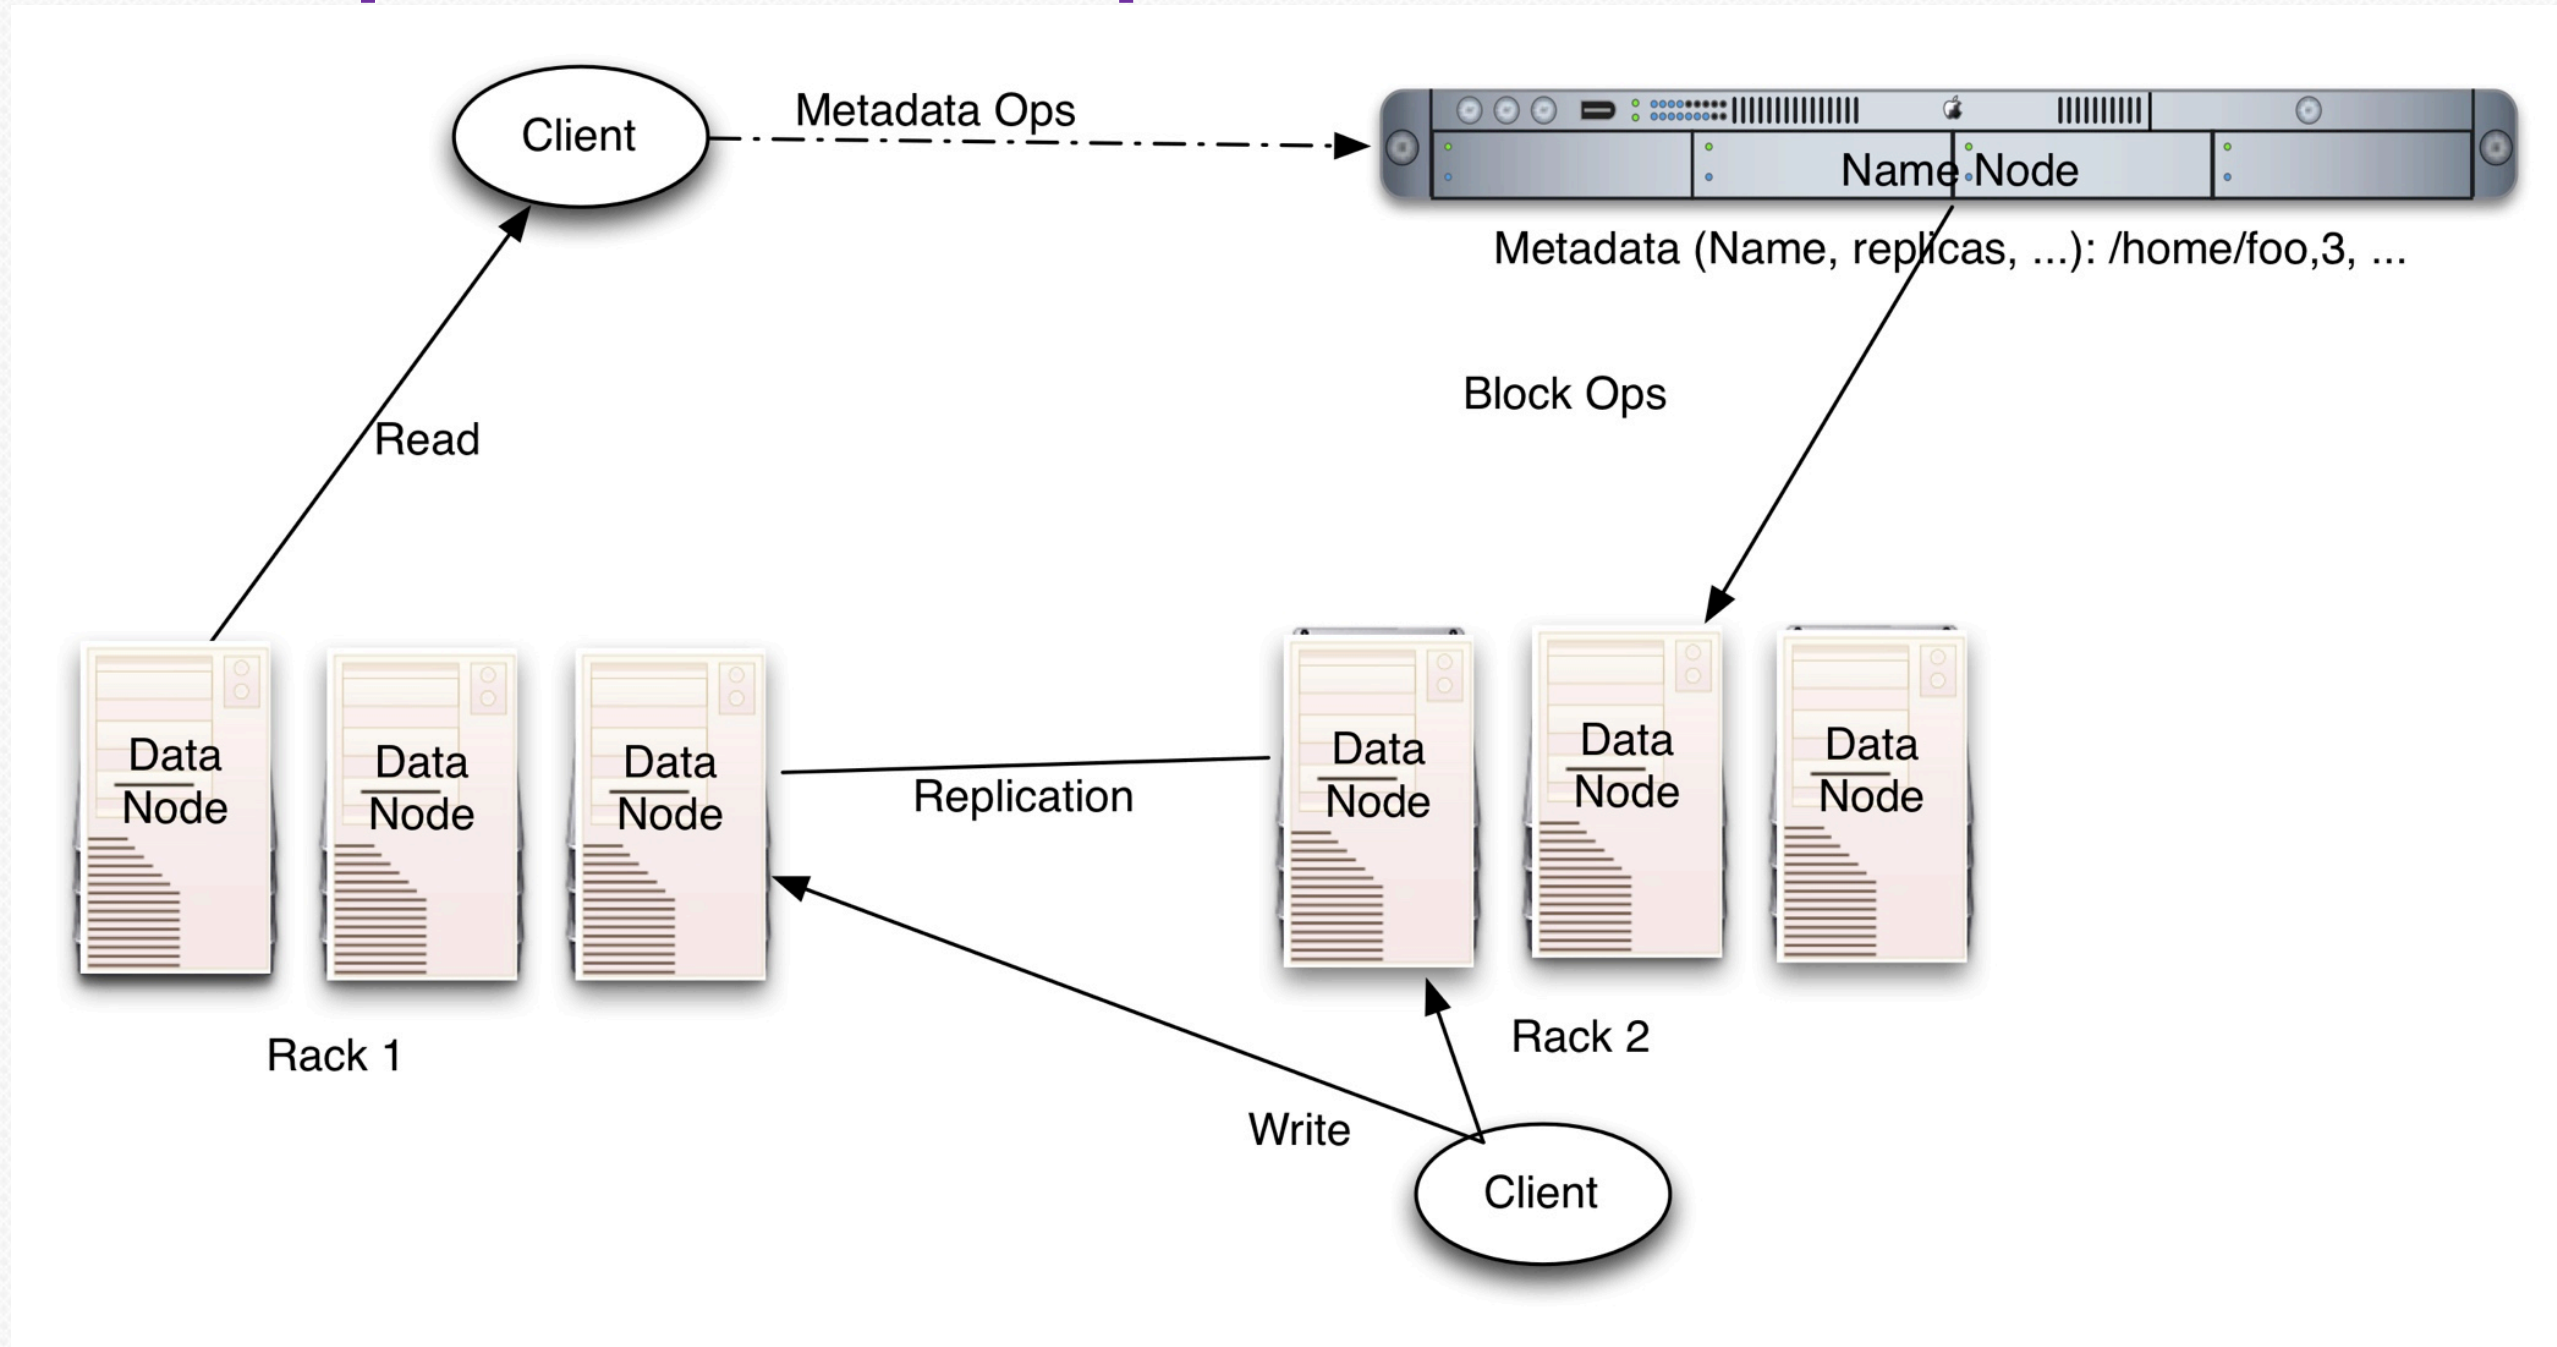




# CEPH with Object Drives

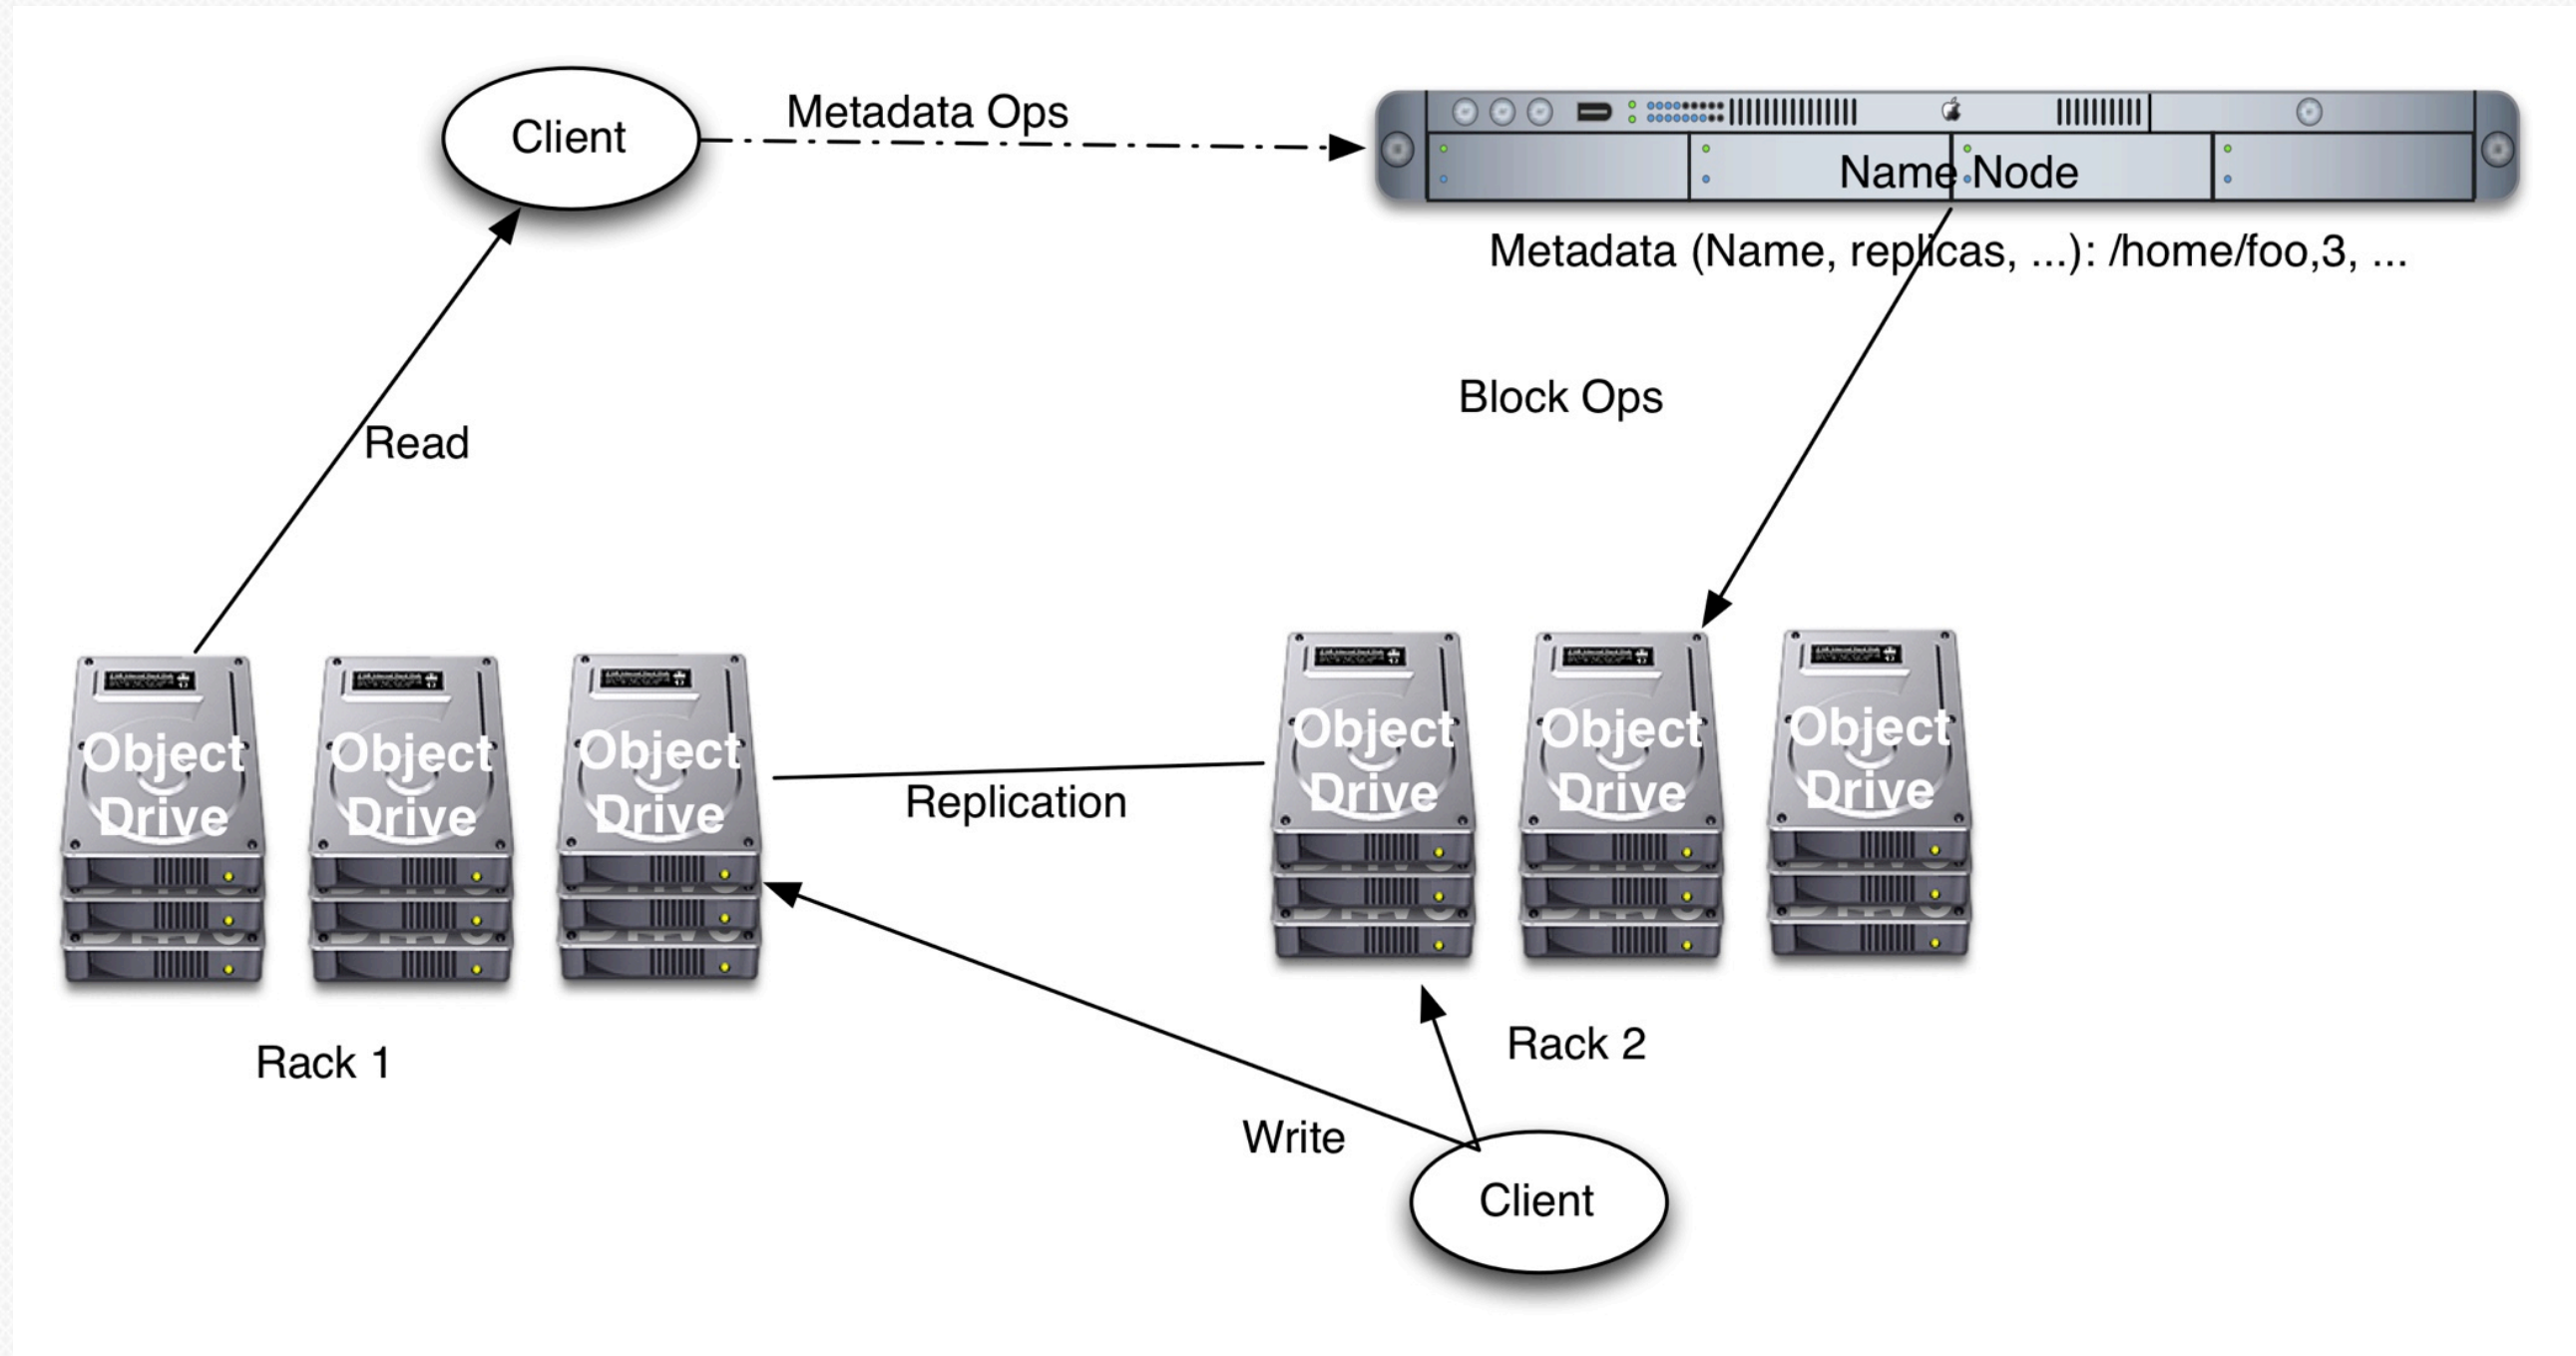


# Example: Hadoop Architecture



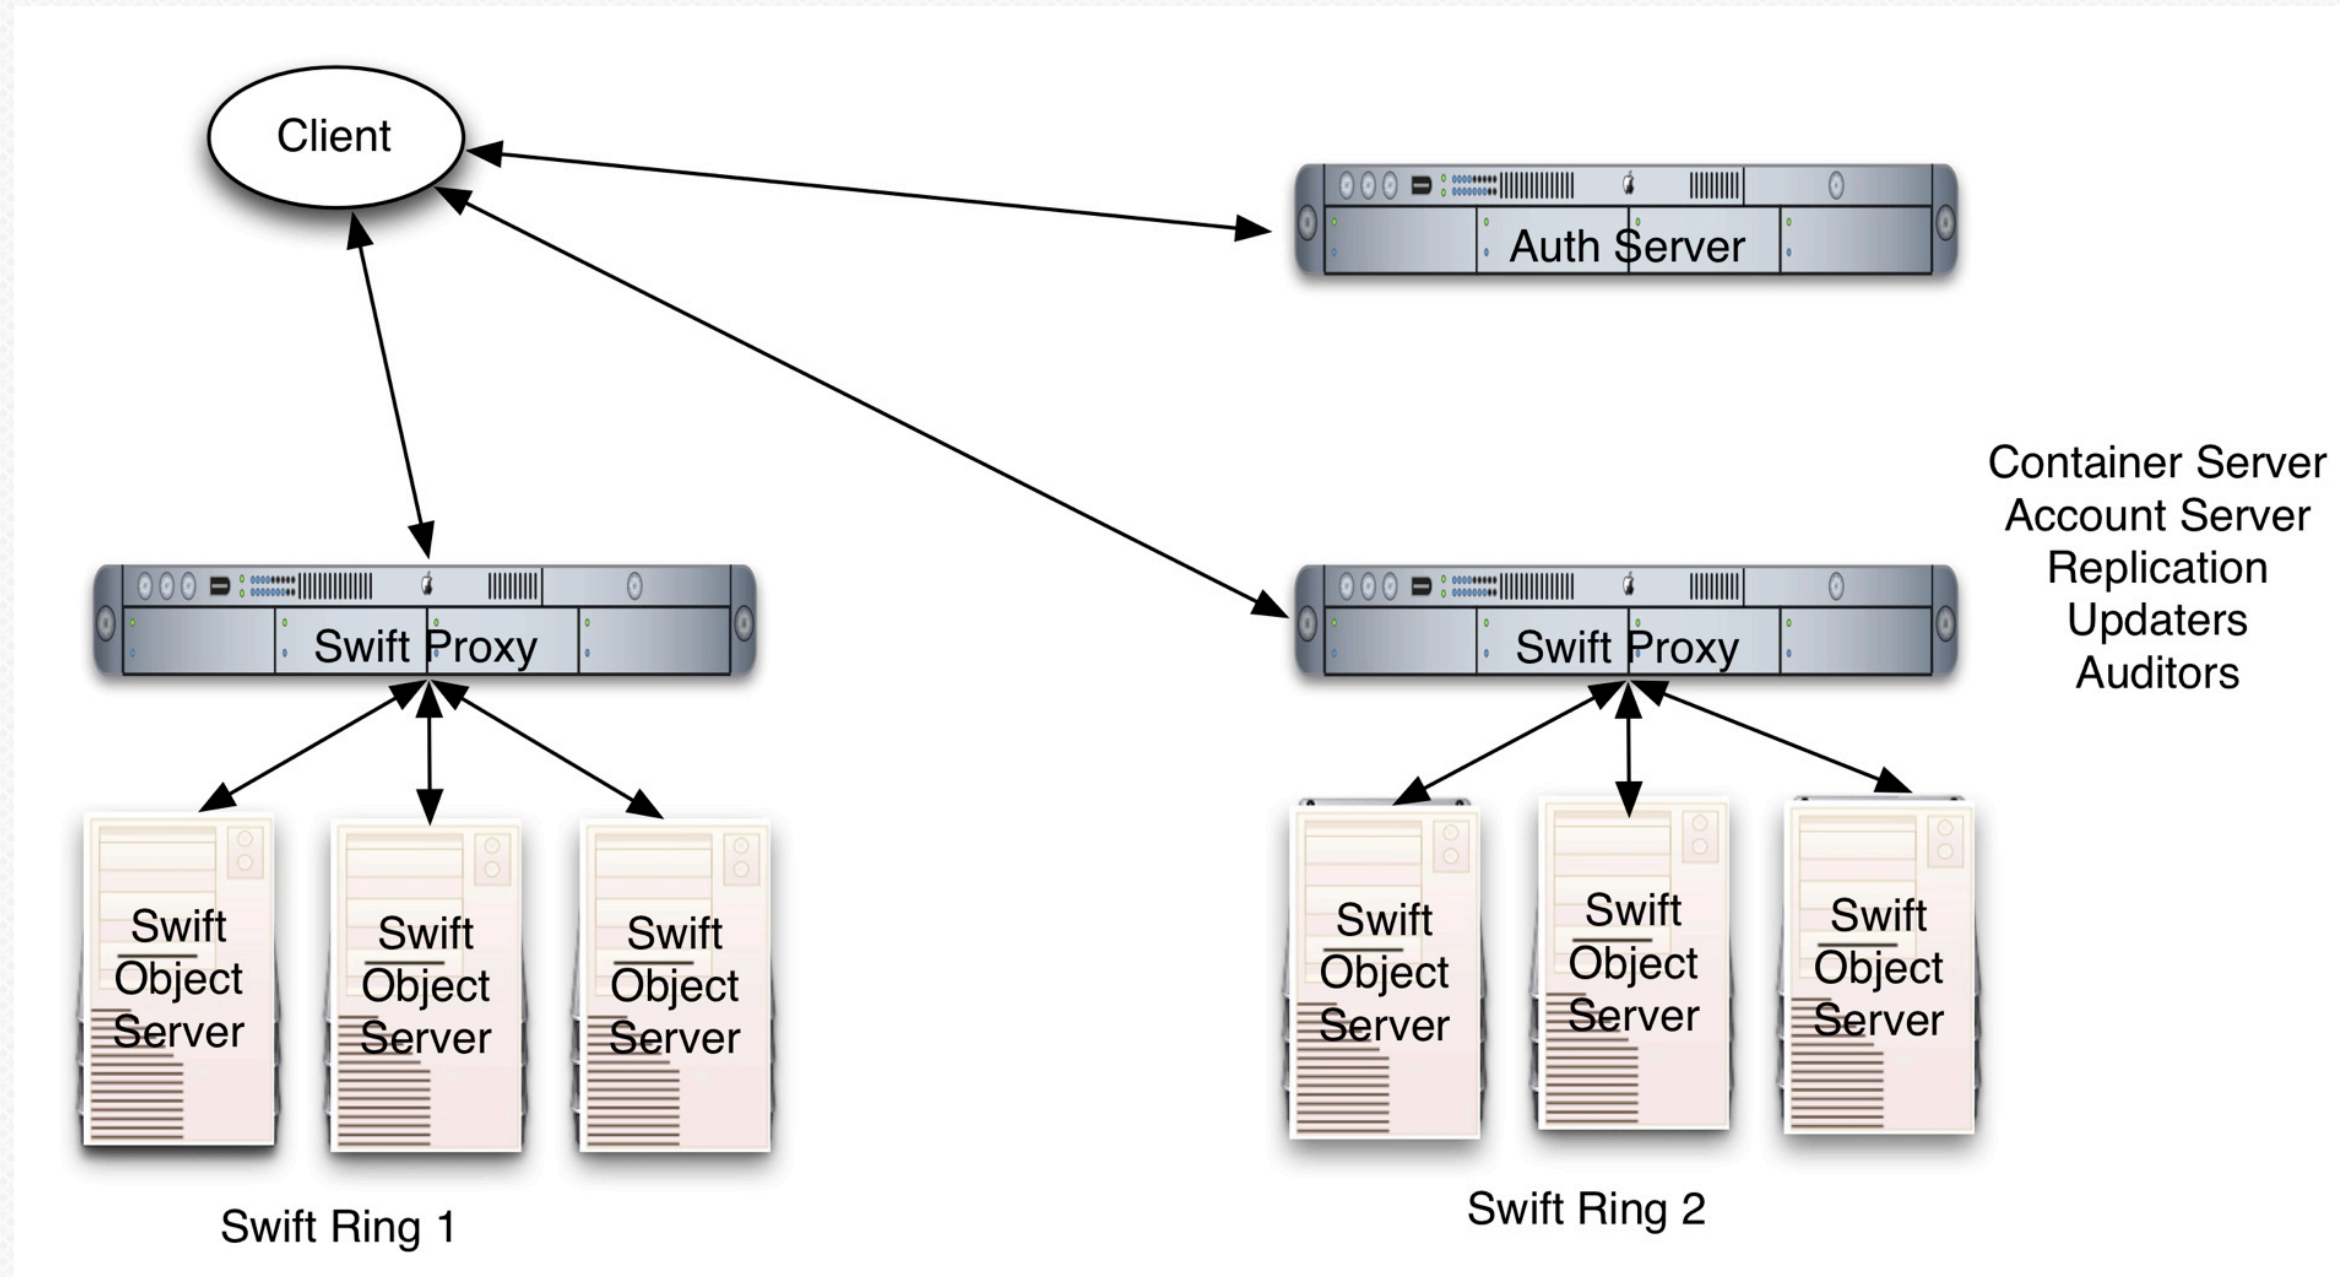


# Hadoop with Object Drives

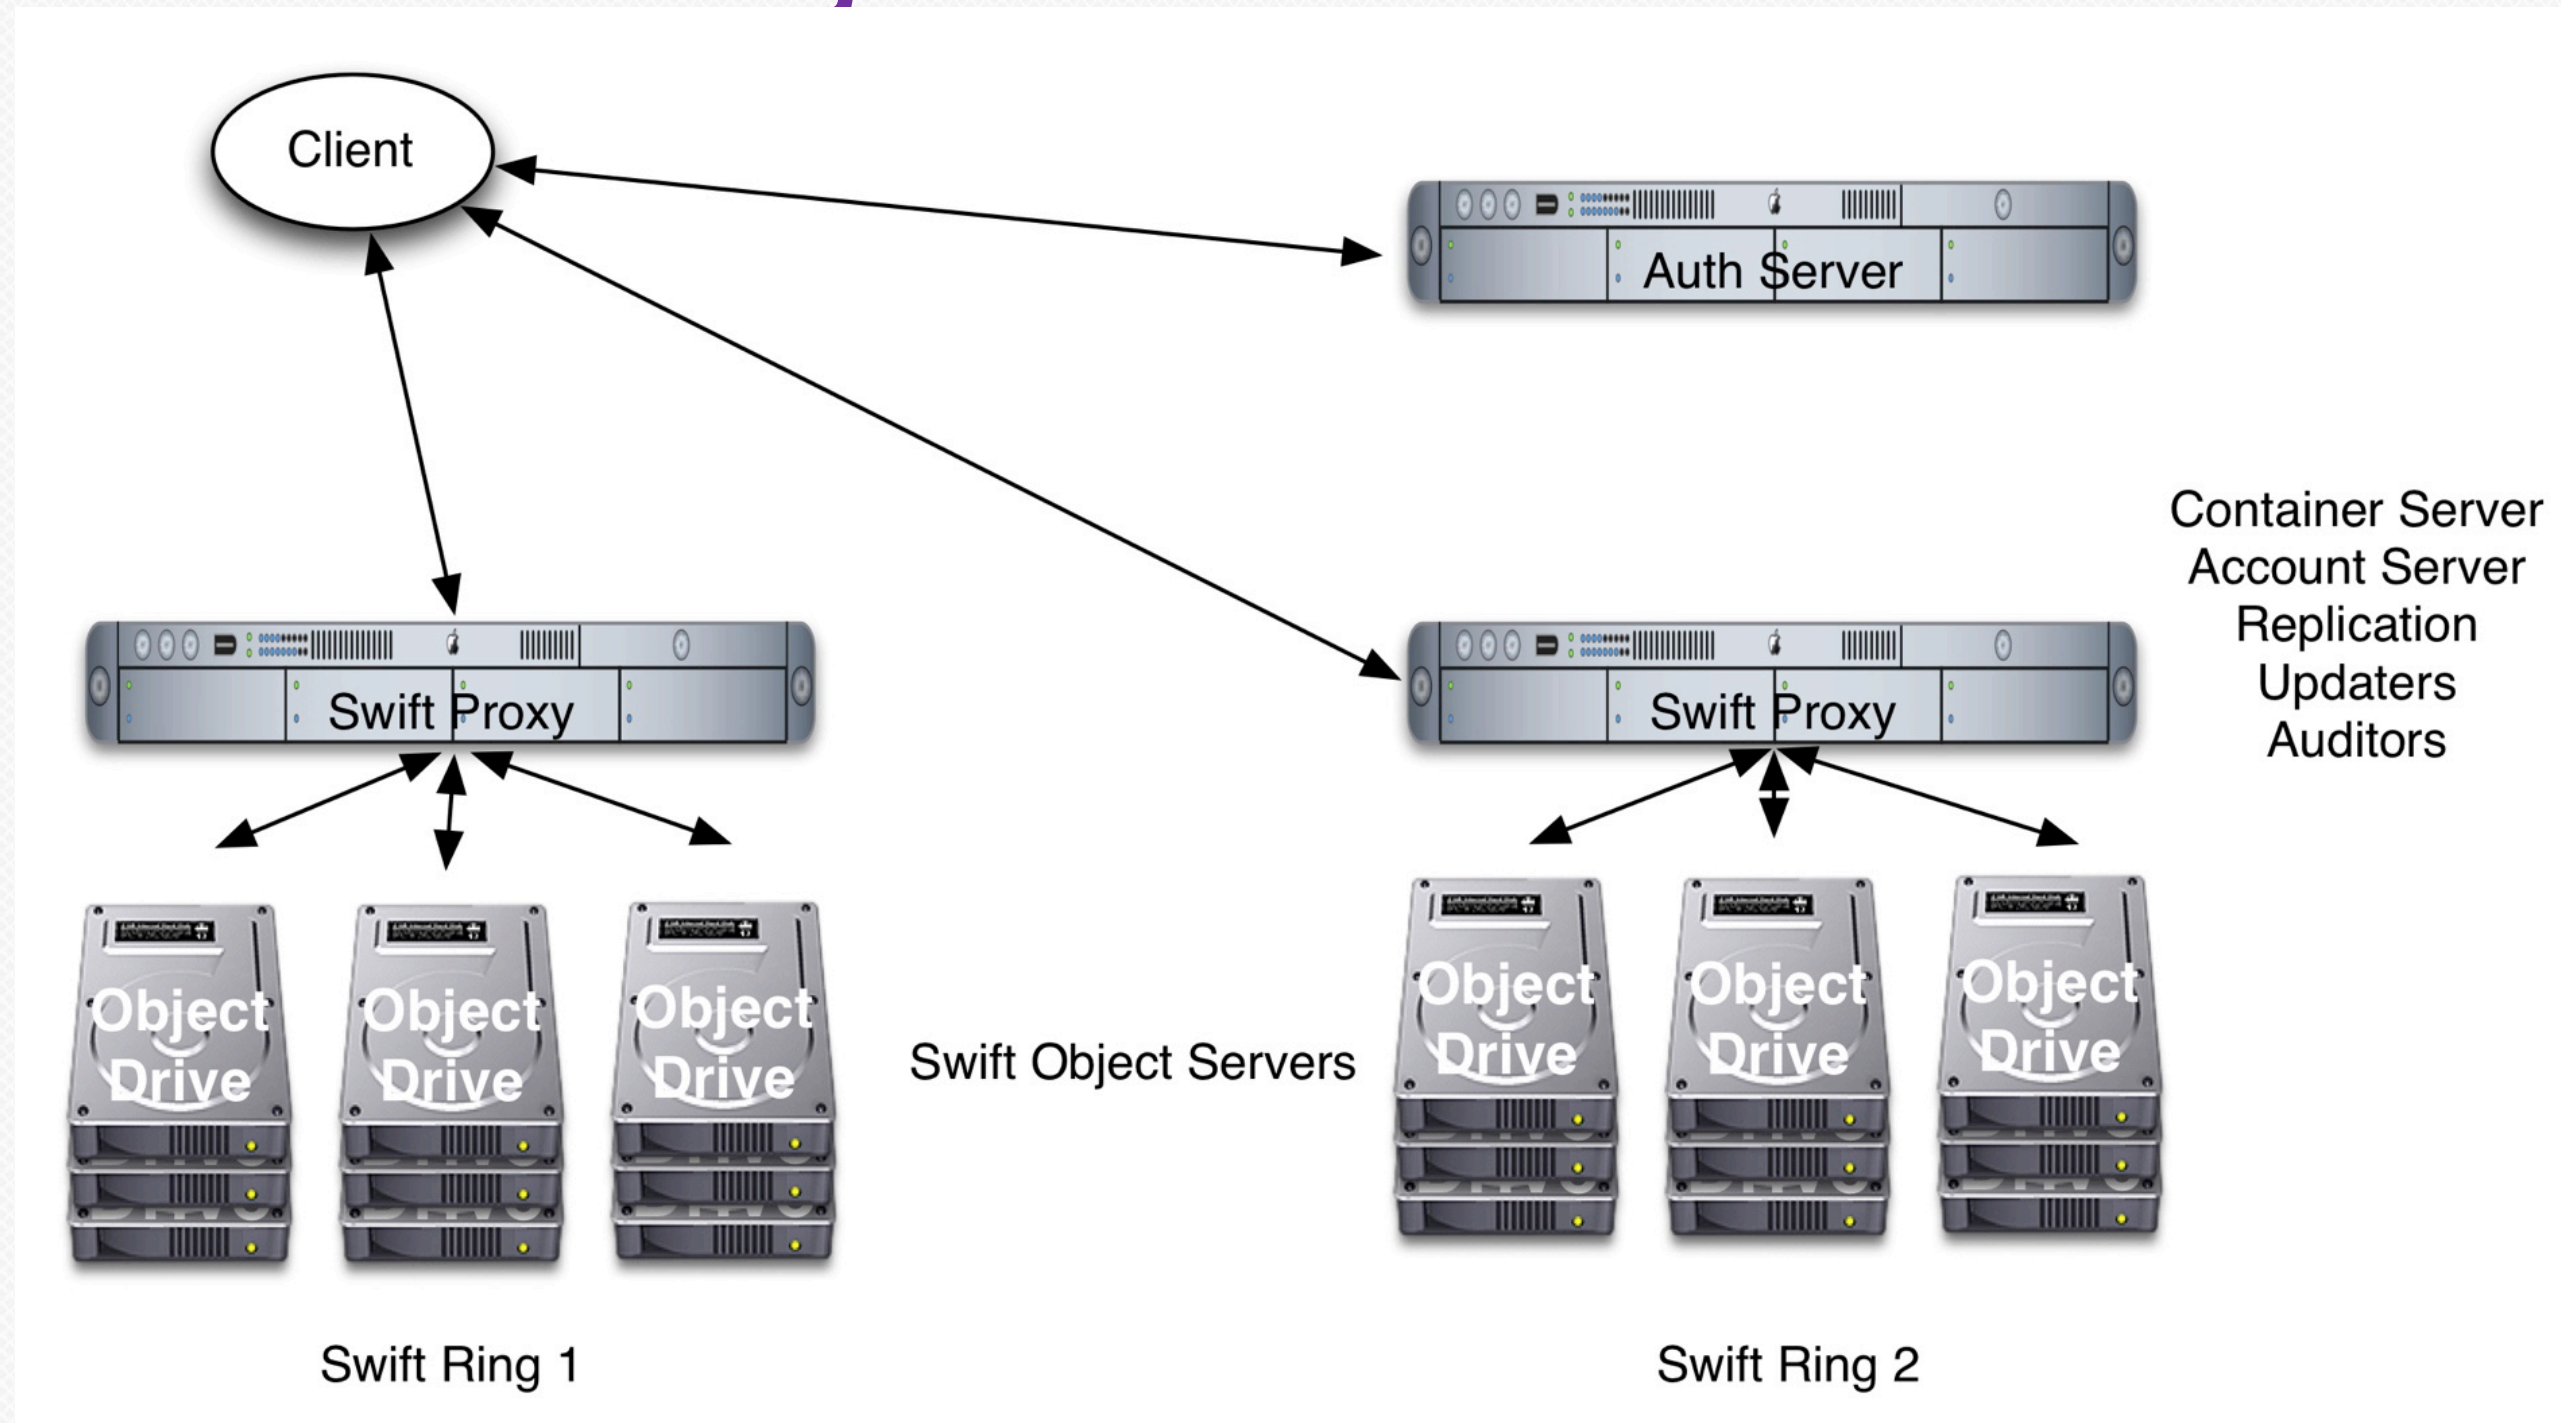




# Example: Swift Architecture



# Swift with Object Drives





# What is needed?

With a new market segment like this, if every vendor pursues their own interface and integration with Client software, the entire market takes a slow path to adoption

By standardizing the points of interoperability, drive manufacturers second source each other, reducing the risk for clients to adopt

SNIA has a good history of solving this by creating consensus standards



# What is the solution?

A new Technical Working Groups (TWG) has been formed to create a specification defining these points of interoperability

- Object Store spec (Key/Value?)
- Drive hosted software

Composed of stakeholders from both drive manufacturers and scale out storage software





# Layers of Requirements

Physical (form factor)

Electrical/Environmental – connectors, power profile, cooling profile

Discovery/Management (including capabilities) , “loading”, security, assignment of ownership, etc. and Configuration (may be loaded software) and setup

Use of the drive (perhaps bifurcation here)



# Physical Layer Requirements

Group doesn't want to be constrained by any physical form factor

Some vendors will deliver a drive in current (i.e. 3.5") form factors

The group will not work in this space ...





# Electrical/Environmental

A large part of the market will be Ethernet based

Some PCIe connections are likely

C) A standardized way that “drive-like” devices connect via Ethernet or other transport

Group may work on this (someday)



# Discovery/Management

B) There is value in having a standardized way to discover and manage providers of the Object Protocol

D) A standardized way to discover, manage and load and run arbitrary software on general compute/storage devices





# Use

A) There is value in a common Object Protocol for storage

Inter-drive object protocol (on the wire)

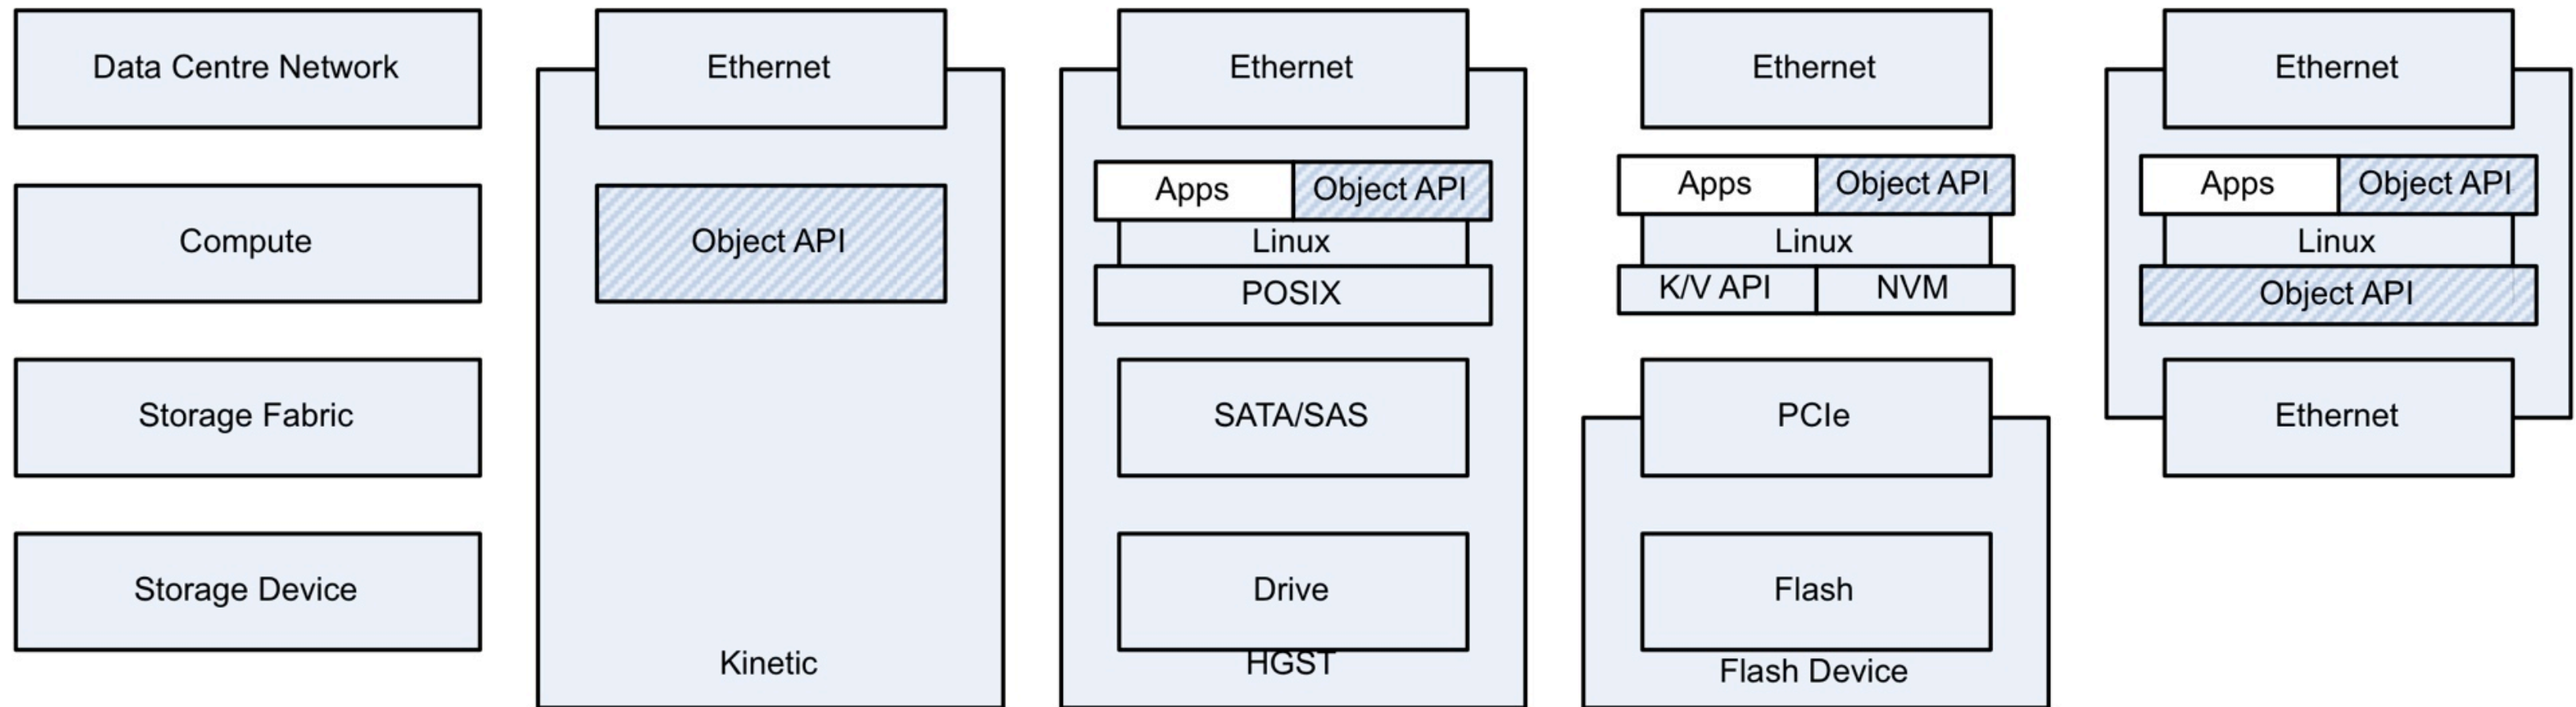
Intra-drive object API (inside drive)

Load your own interface code (works on top of object protocol, not Posix, not blocks)

Do we go beyond key/value?



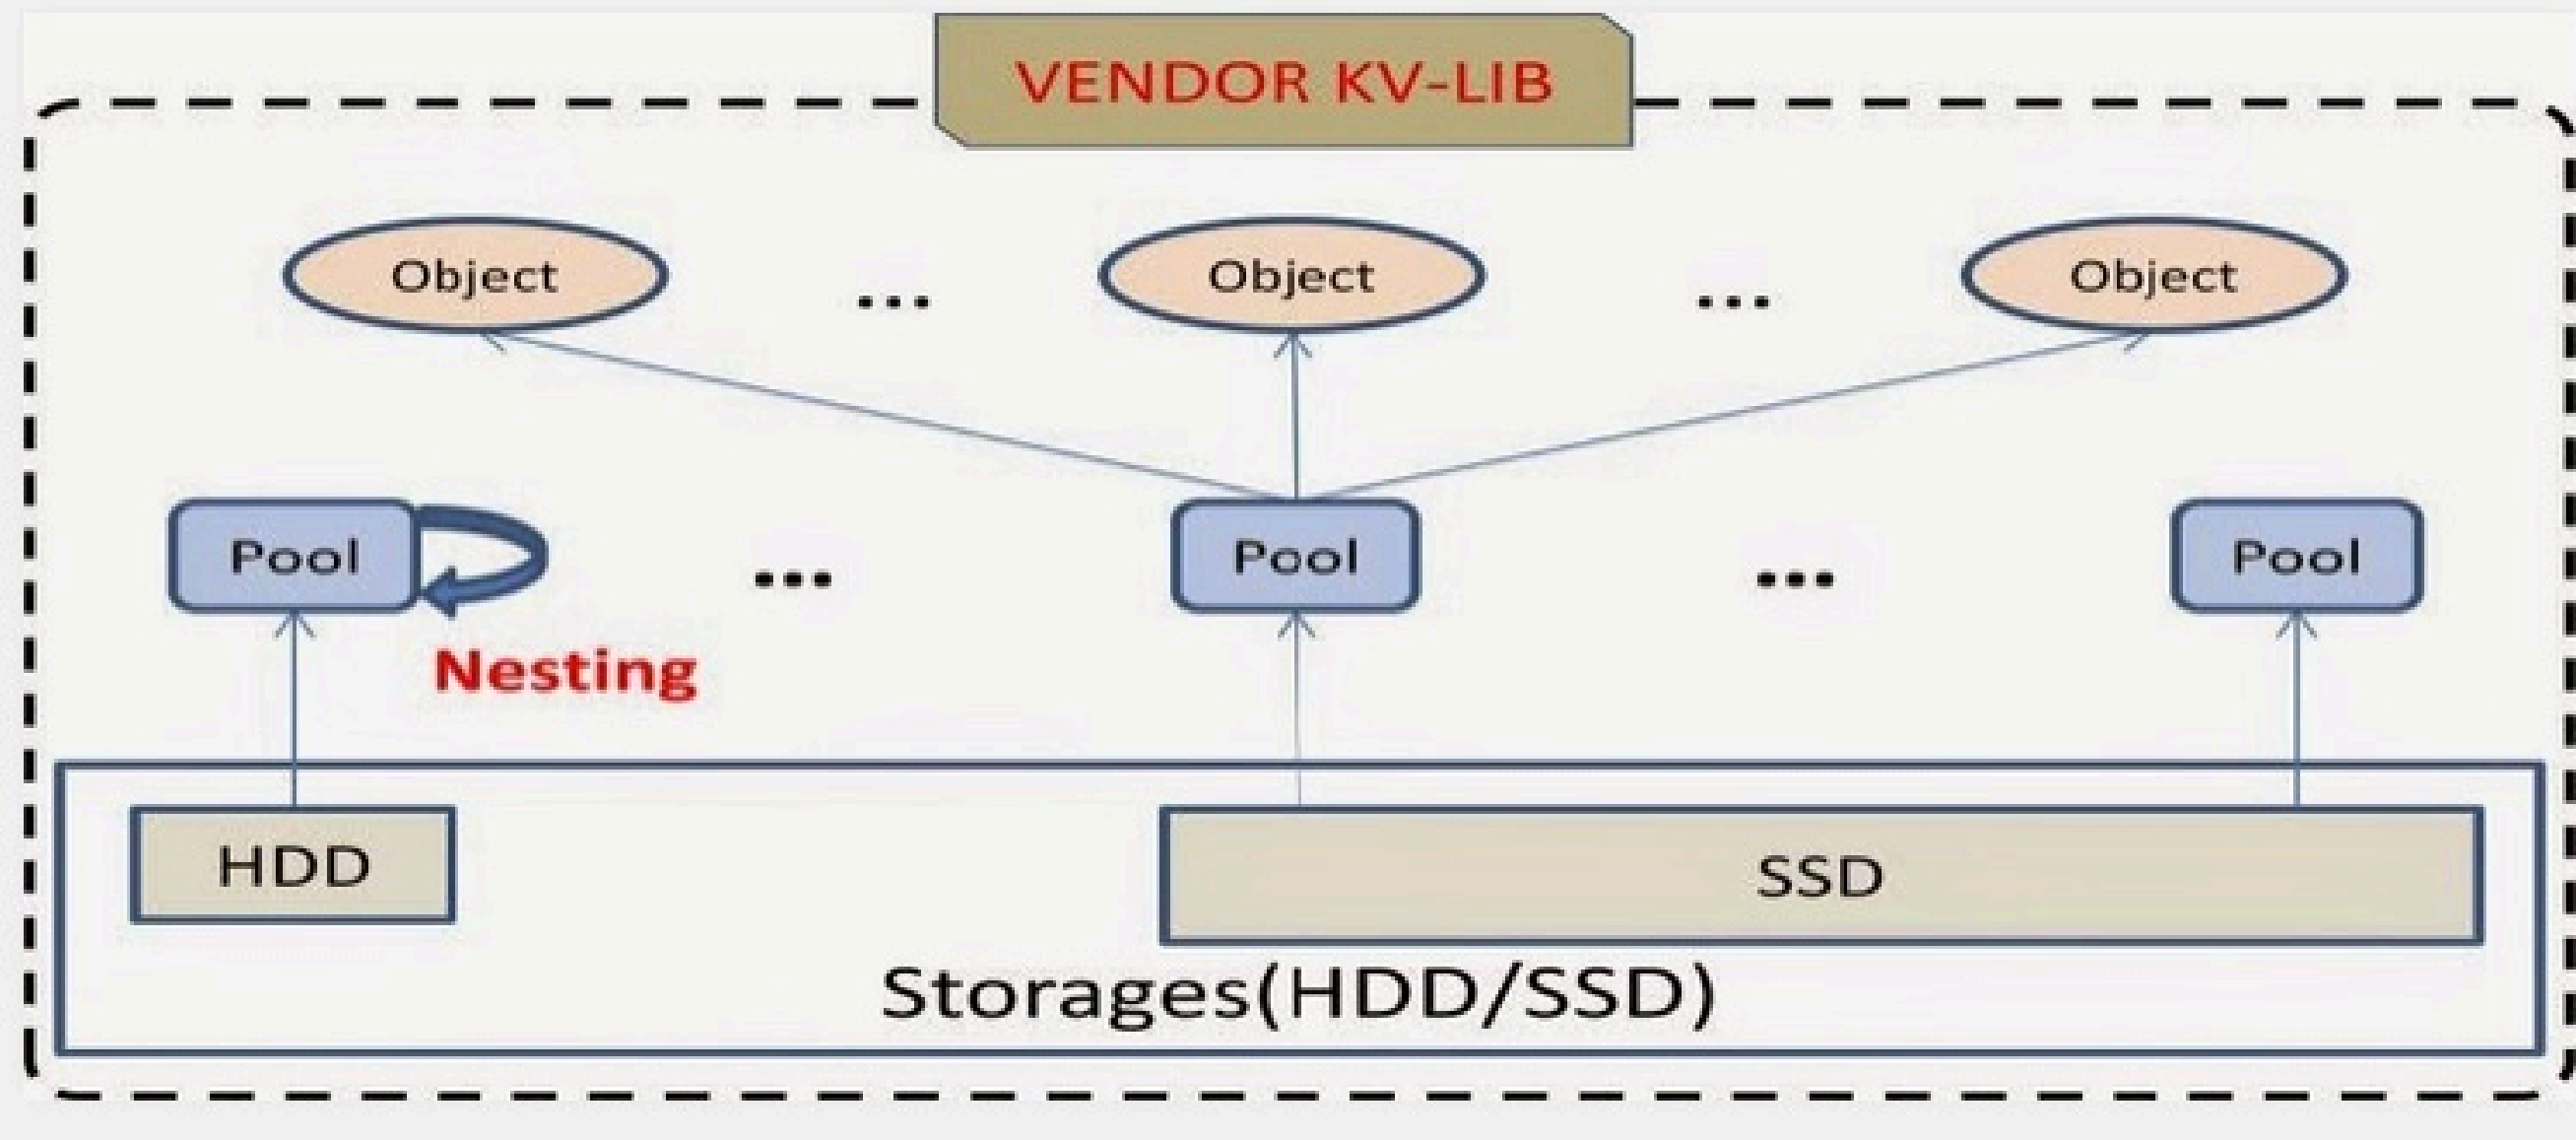
# Some Deployment Options





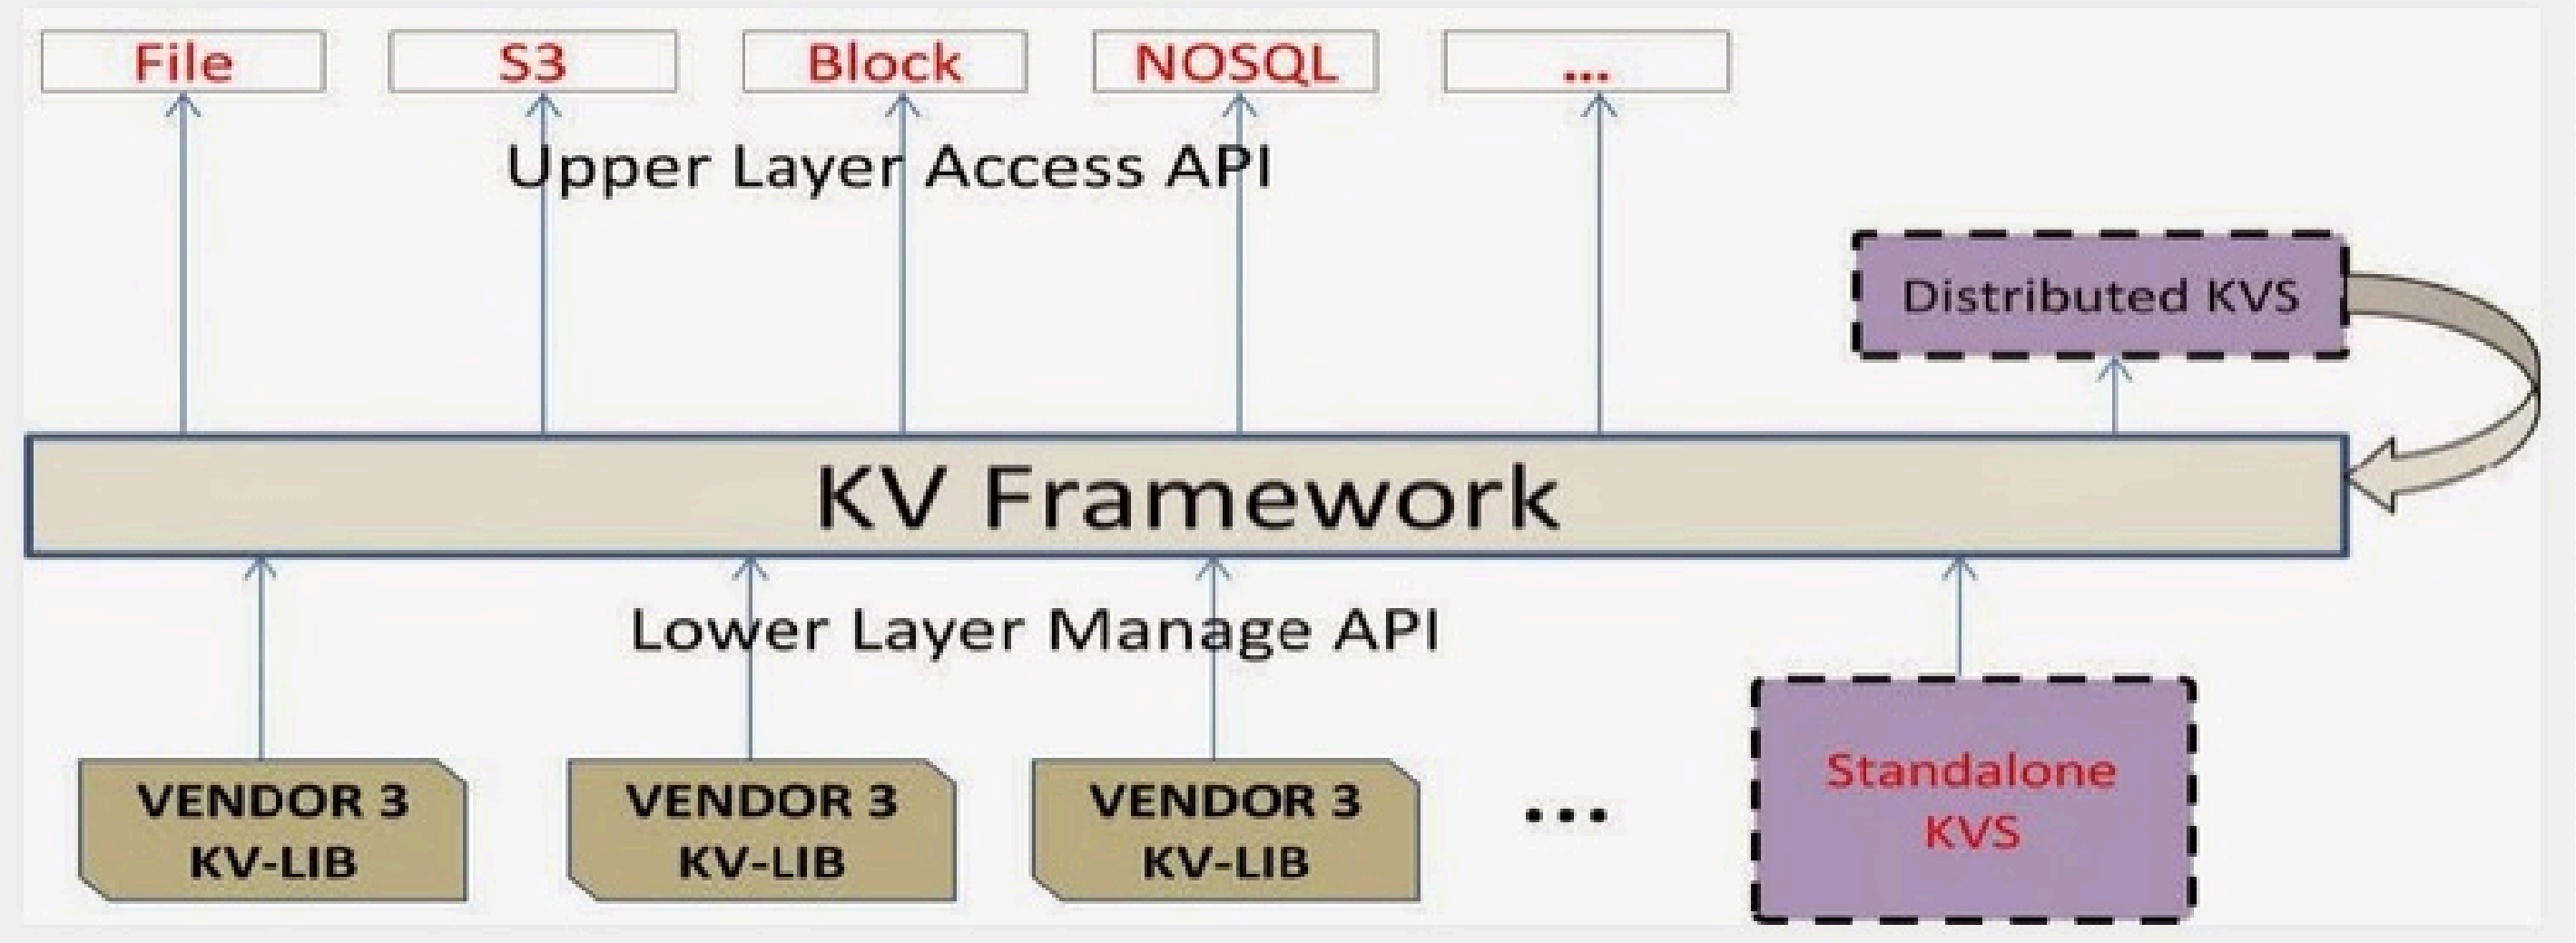
# Key Value Framework (KVF) Data Model

Based on Key Value Framework (KVF), different vendors can register its Key Value Library (KV-LIB) to manage its Key Value Storage (KVS). These KV-LIBs should follow KVF data model, KVS provides pools, and pool provides objects. And pool supports nesting design, which means one pool can be on top of another pool, Shown as Figure\_2.



# Key Value Framework Architecture

Key Value Framework (KVF) provides 2 layer interfaces: Lower Layer Manage (LLM) API and Upper Layer Access (ULA) API, shown as Figure\_3.





# Next steps

Join the Object Drive TWG

<https://members.snia.org/apps/org/workgroup/objecttwg/>

Scale out storage software vendors: what are your requirements for this work?



# Thank You





