

**OPEN**

Compute Summit

January 28-29, 2014

San Jose

# Kinetic Open Storage Platform

## Dis-intermediating applications from storage

James Hughes, Principal Technologist, Seagate Technology  
...and many others







DATA CENTER > STORAGE

## Seagate: Fibre Channel? RAID? SATA? File System? All RUBBISH

**App to disk via ethernet, baby. The rest of you, clear out your desks**

By Chris Mellor, 22nd October 2013

Follow 5,083 followers

45

[Free Regcast : Managing Multi-Vendor Devices with System Centre 2012](#)

Seagate is building hard disk drives with a direct Ethernet interface and object-style API access for scalable object stores, a plan which - if it works - would destroy much of the existing, typical storage stack.

Drives would become native key/value stores that manage their own space mapping with accessing applications simply dealing at the object level with gets and puts instead of using file abstractions.

Seagate says it has developed its Kinetic technology because the existing app-to-drive storage stack is clumsy, inefficient and delays data access. Put an Ethernet interface

### RELATED STORIES

What's the first Kinetic Ethernet hard drive? Psst, it's the 4TB Terascale

Seagate to storage beds:

## Enterprise Backup and Recovery

Gartner Best Practices for Repairing the Broken State of Backup Report.



### MOST READ

### MOST COMMENTED

LG G Flex: A new cheeky curvy mobe with a 'SELF-HEALING' bottom

Everything's going to be all white: Google Nexus 5 mobe expected Friday

In a meeting with a woman? For pity's sake

[Managing Multi-Vendor Devices with System Centre 2010](#)

## RELATED STORIES

What's the first Ethernet

Seagate is building hard disk drives with a direct Ethernet interface and object-style API access for scalable object stores, a plan which - if it works - would destroy much of the existing, typical storage stack.

Drives would become native key/value stores that manage their own space mapping with accessing applications simply dealing at the object level with gets and puts instead of using file abstractions.

[Managing Multi-Vendor Devices with System Centre 2010](#)



# Seagate Kinetic Open Storage Platform

Dis-intermediates applications to drive

- Goes around file systems, volume managers, drivers

Enable ecosystem of value added software

- Partners (like Basho) can create their own system value

Lower TCO

- Eliminates complexity



# Seagate Kinetic Open Storage Platform

Existing Hard Drive Technology

Ethernet Interface

Key Value API





# Seagate Kinetic Open Storage Platform

Existing Hard Drive Technology

and

Ethernet Interface


and

Key Value API



# Connector

- SAS tx, rx is now Ethernet
- Electrically SGMII
- “Backplane Ethernet”





# Typical High Density

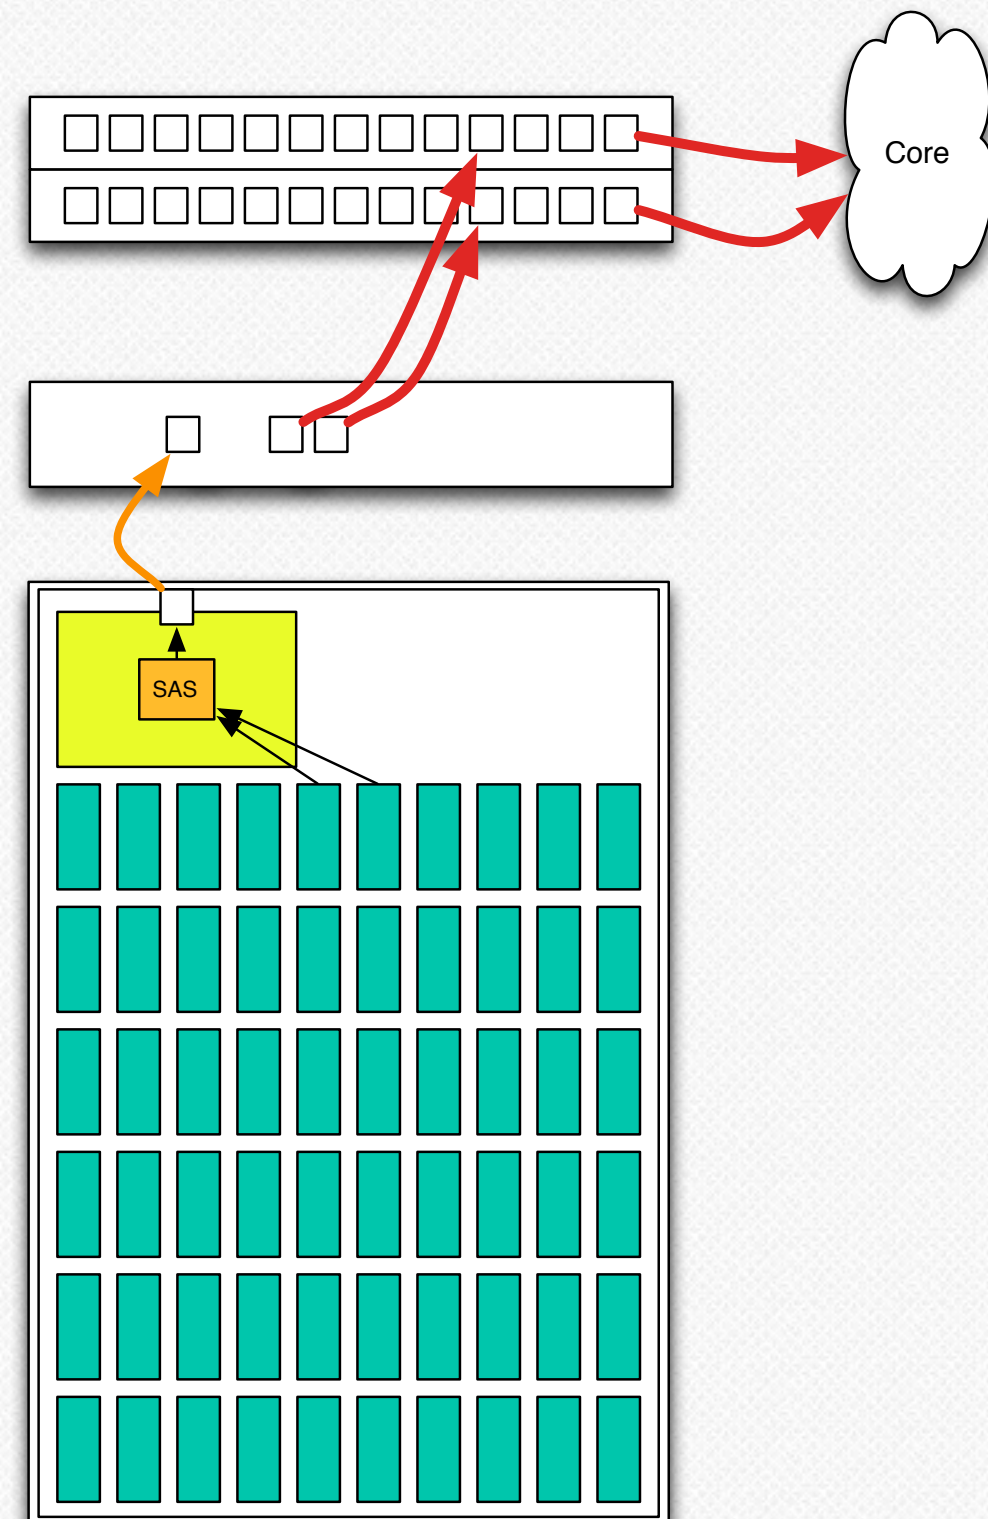
## Intel server

- Double Socket
- 48GB Ram
- 500w

## SAS tray


- Connected to the server

Server fails, all drives are offline until fixed



# Typical HA High Density

Double the server cost

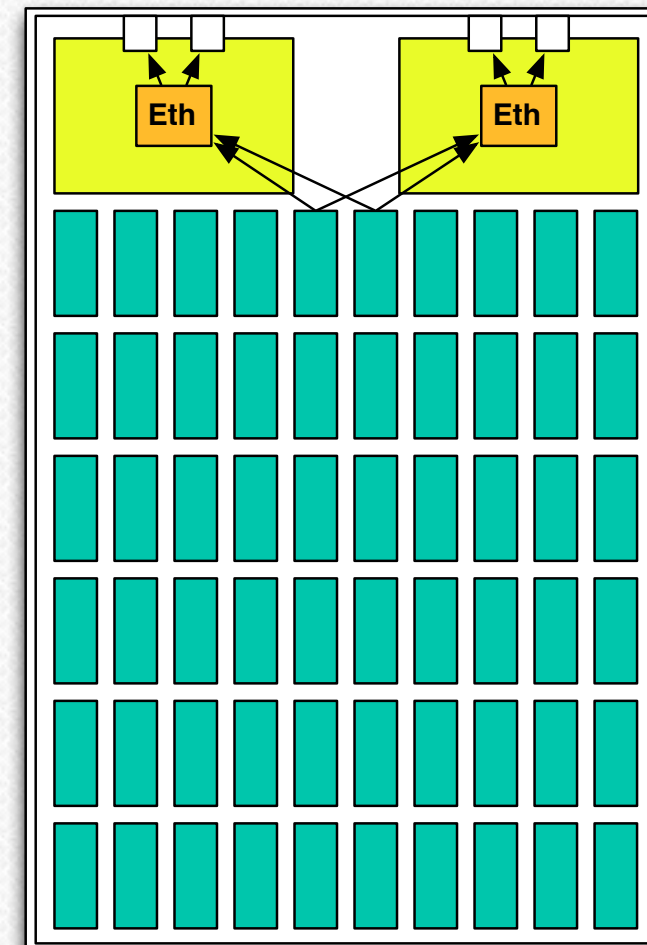




# 60 drive Chassis


Each drive talks to both switches

Each switch has 2 by 10Gb/s Ethernet



# Low cost HA Configuration

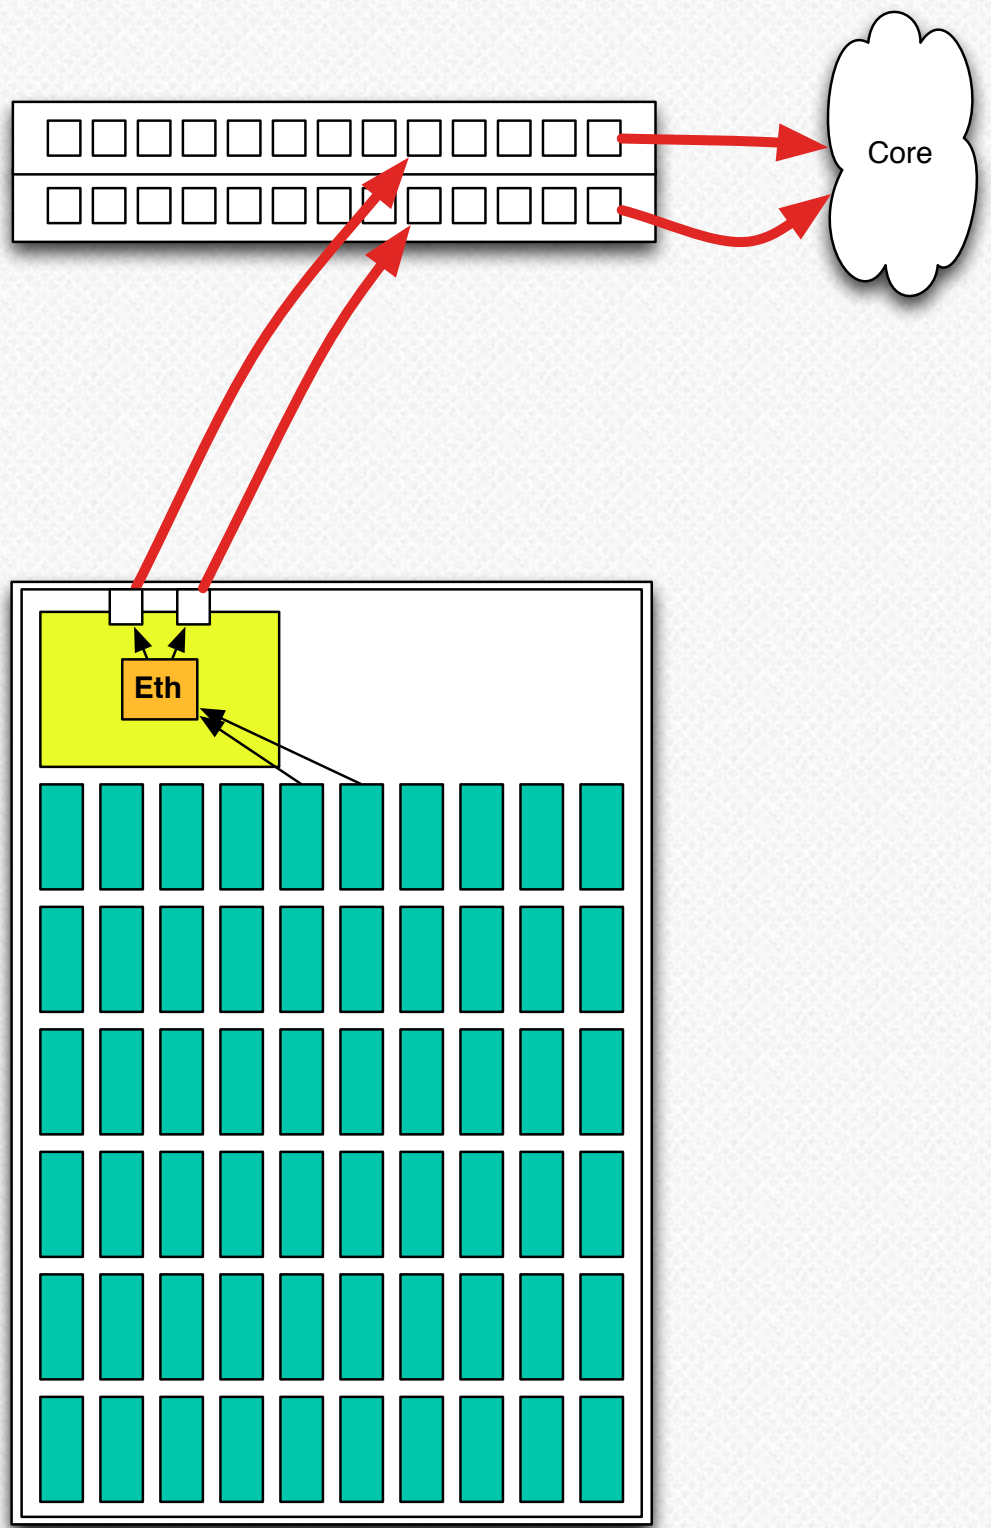
Kinetic Tray talks directly to ToR  
No servers



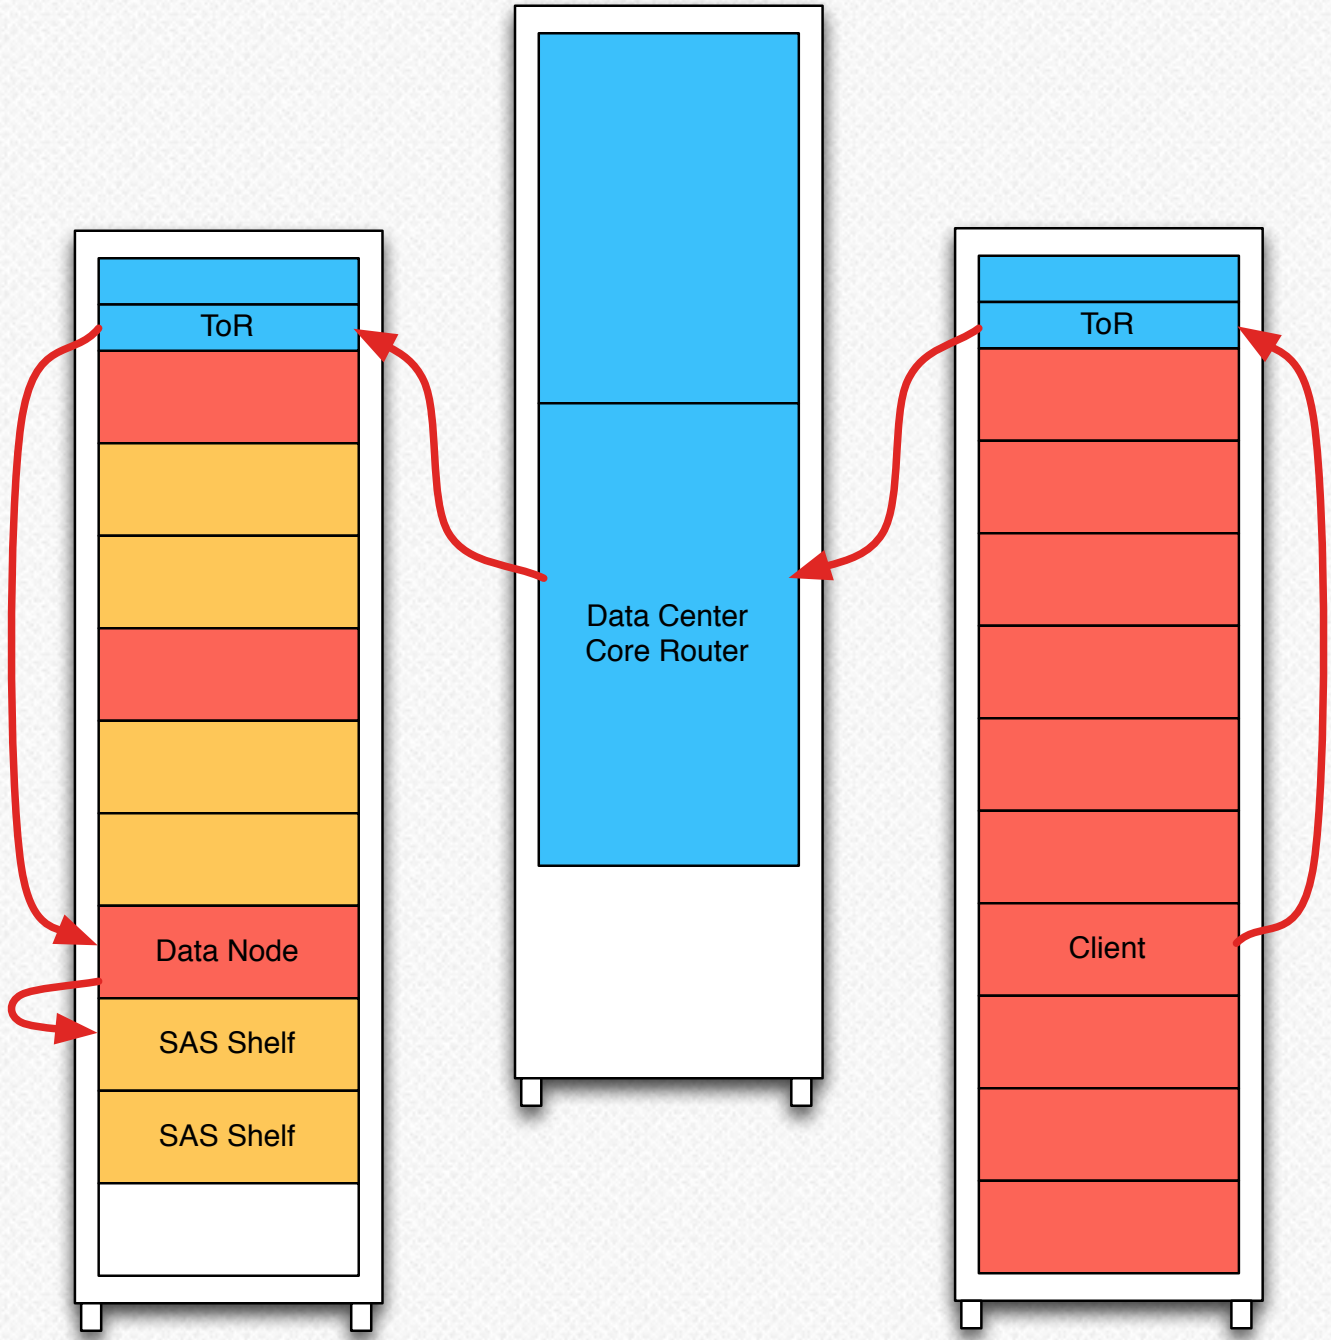


# Even lower cost configuration

Valuable for cold storage replicas





# Existing Traffic Flow






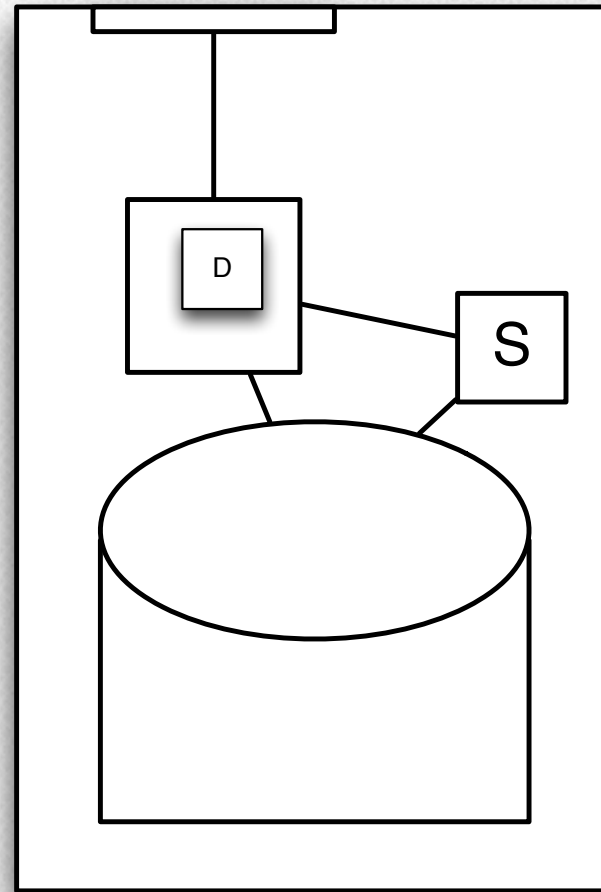
# Kinetic Traffic Flow















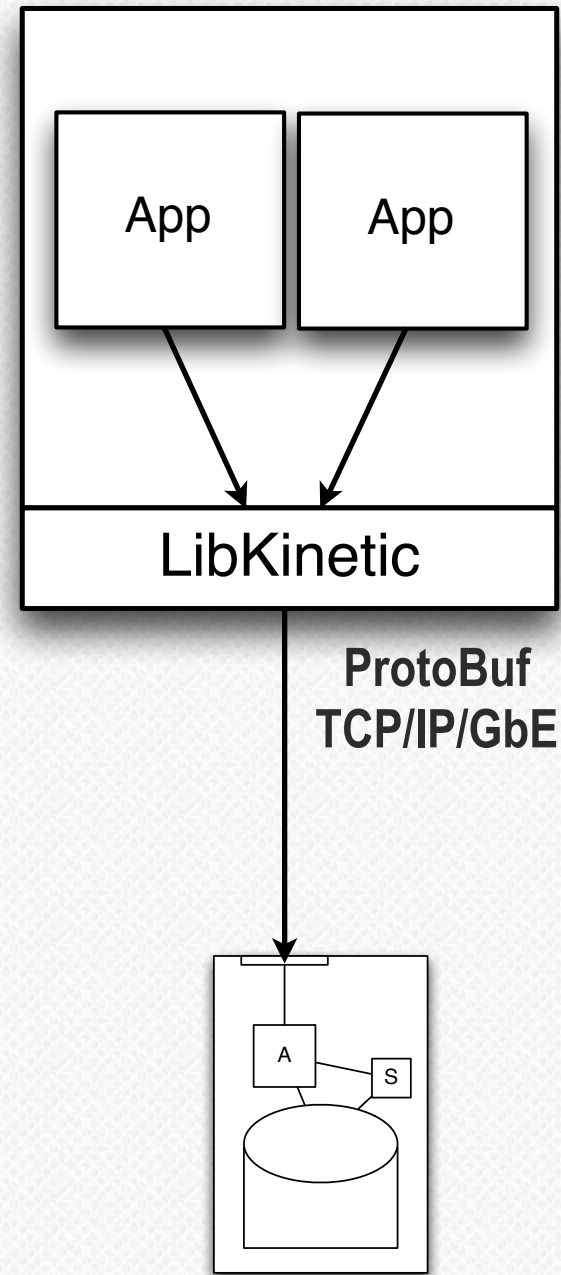



















C++, Java, Python, Erlang, DIY

ProtoBuf  
TCP/IP/GbE




- Application
- Clustering
- Management



C++, Java, Python, Erlang, DIY

- Interconnect

ProtoBuf  
TCP/IP/GbE



- Storage

Proprietary  
to System Vendor


GPL  
Standard

Proprietary  
to Seagate





- Application
- Clustering
- Management




Proprietary  
to System Vendor


- Interconnect

ProtoBuf  
TCP/IP/GbE

GPL  
Standard




- Storage



Proprietary  
to Seagate

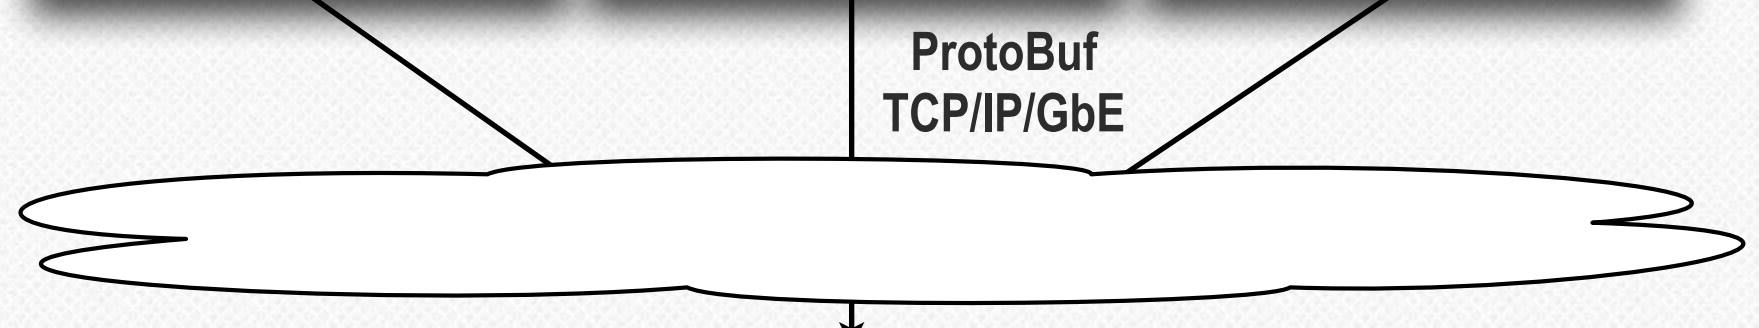


- Application
- Clustering
- Management




Proprietary  
to System Vendor

- Interconnect



GPL  
Standard

- Storage

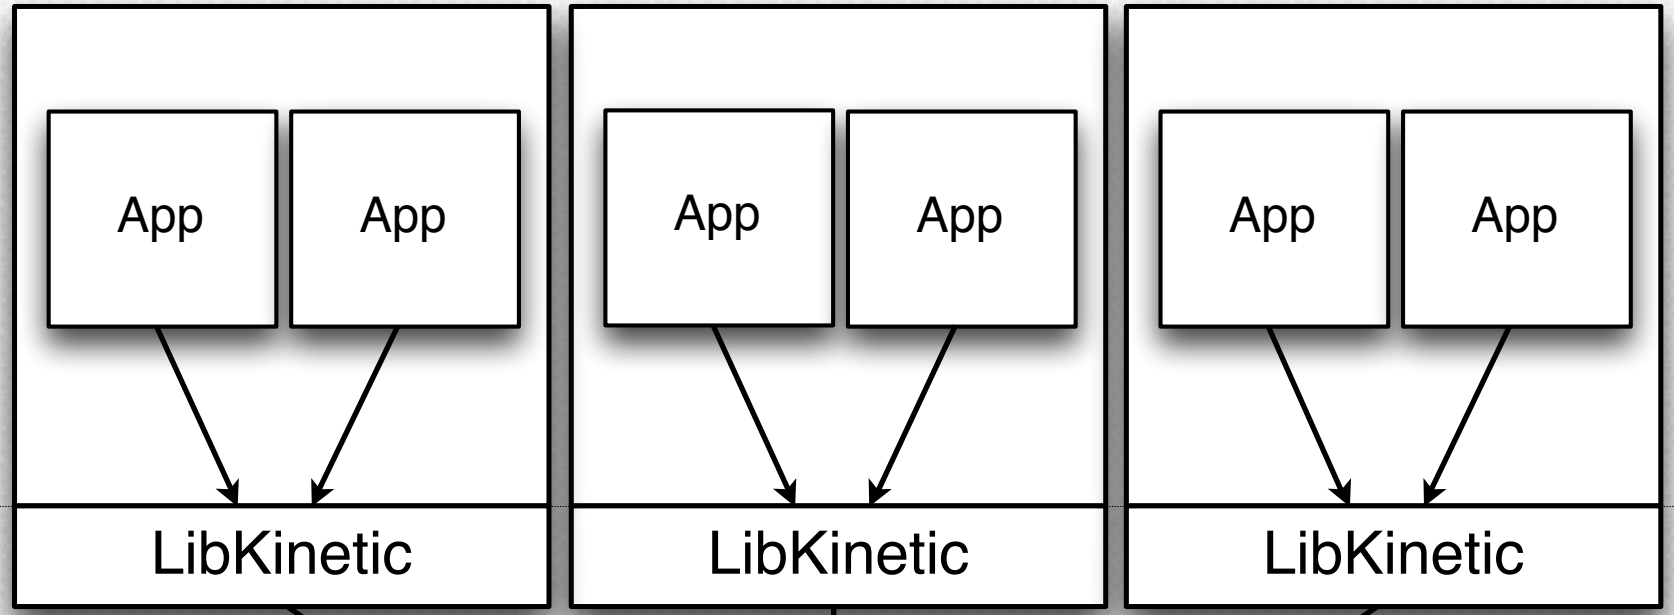


Proprietary  
to Seagate



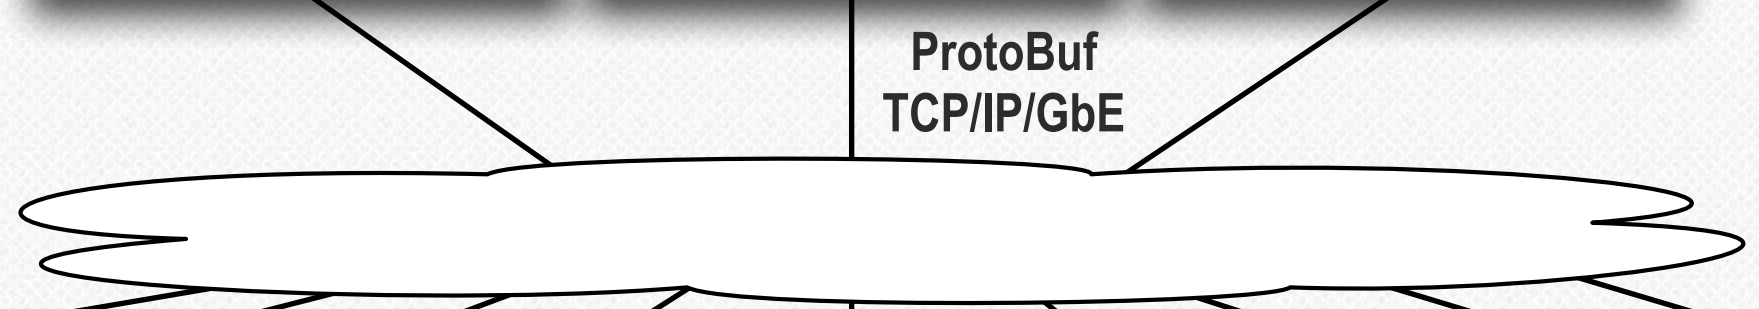


- Application
- Clustering
- Management




Proprietary  
to System Vendor

- Interconnect



GPL  
Standard

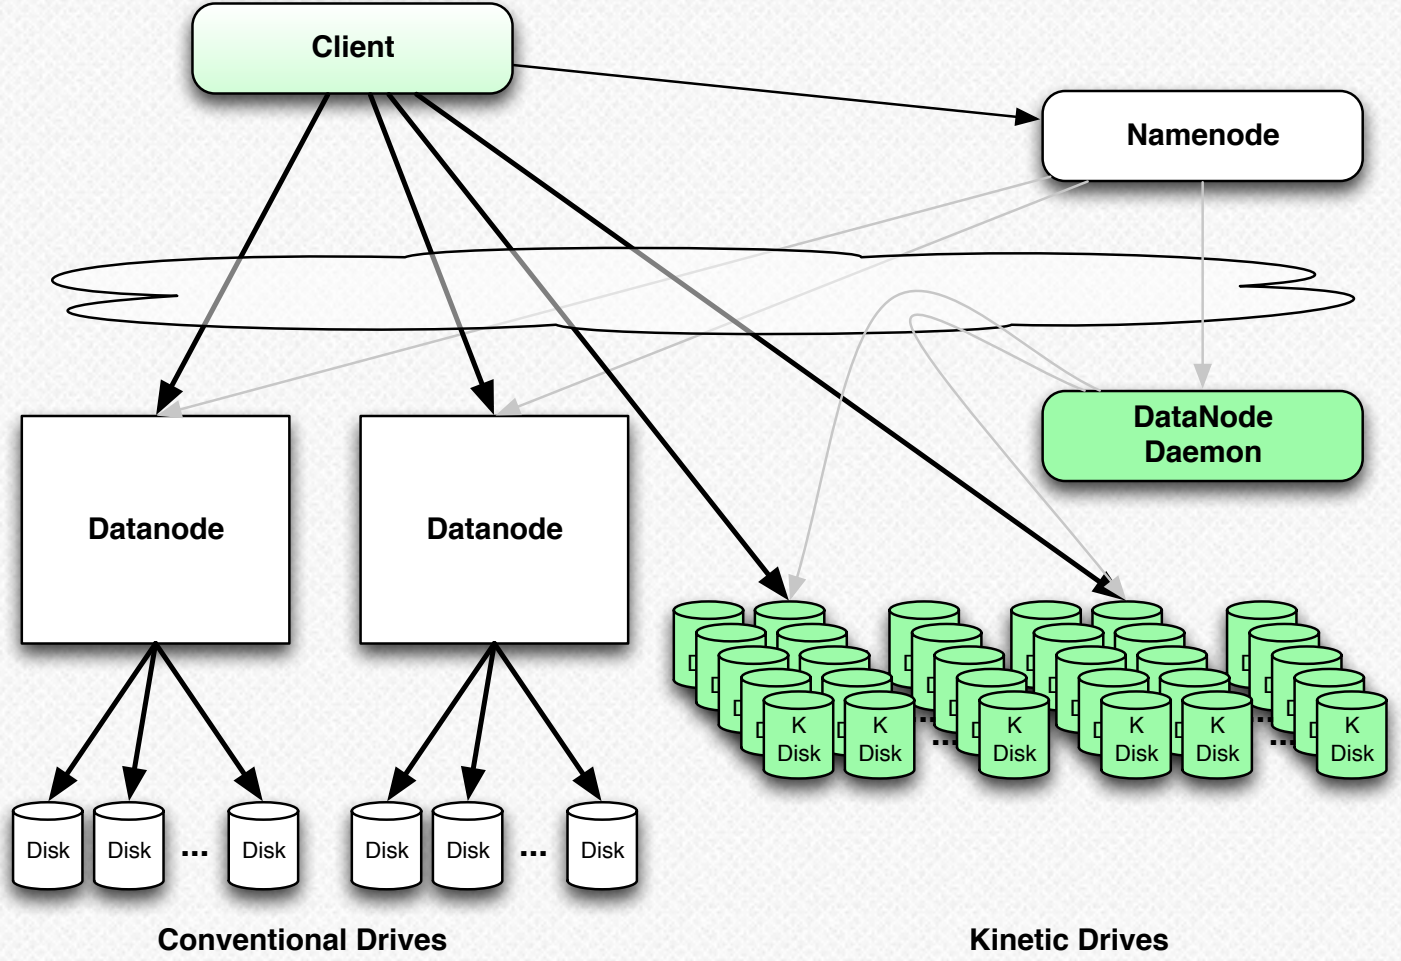
- Storage



Proprietary  
to Seagate

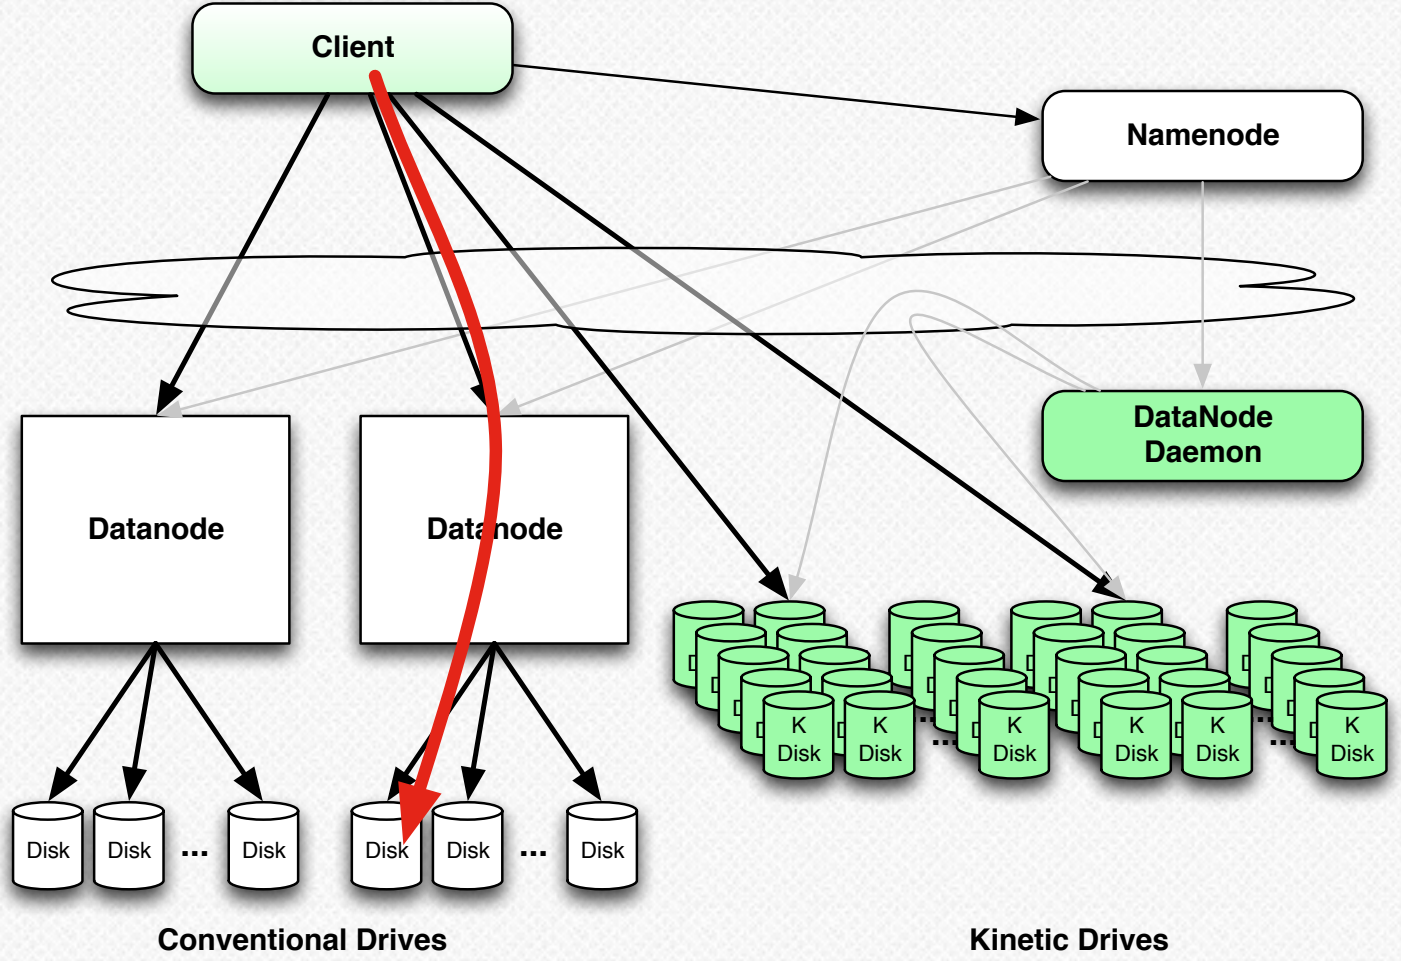


# HDFS



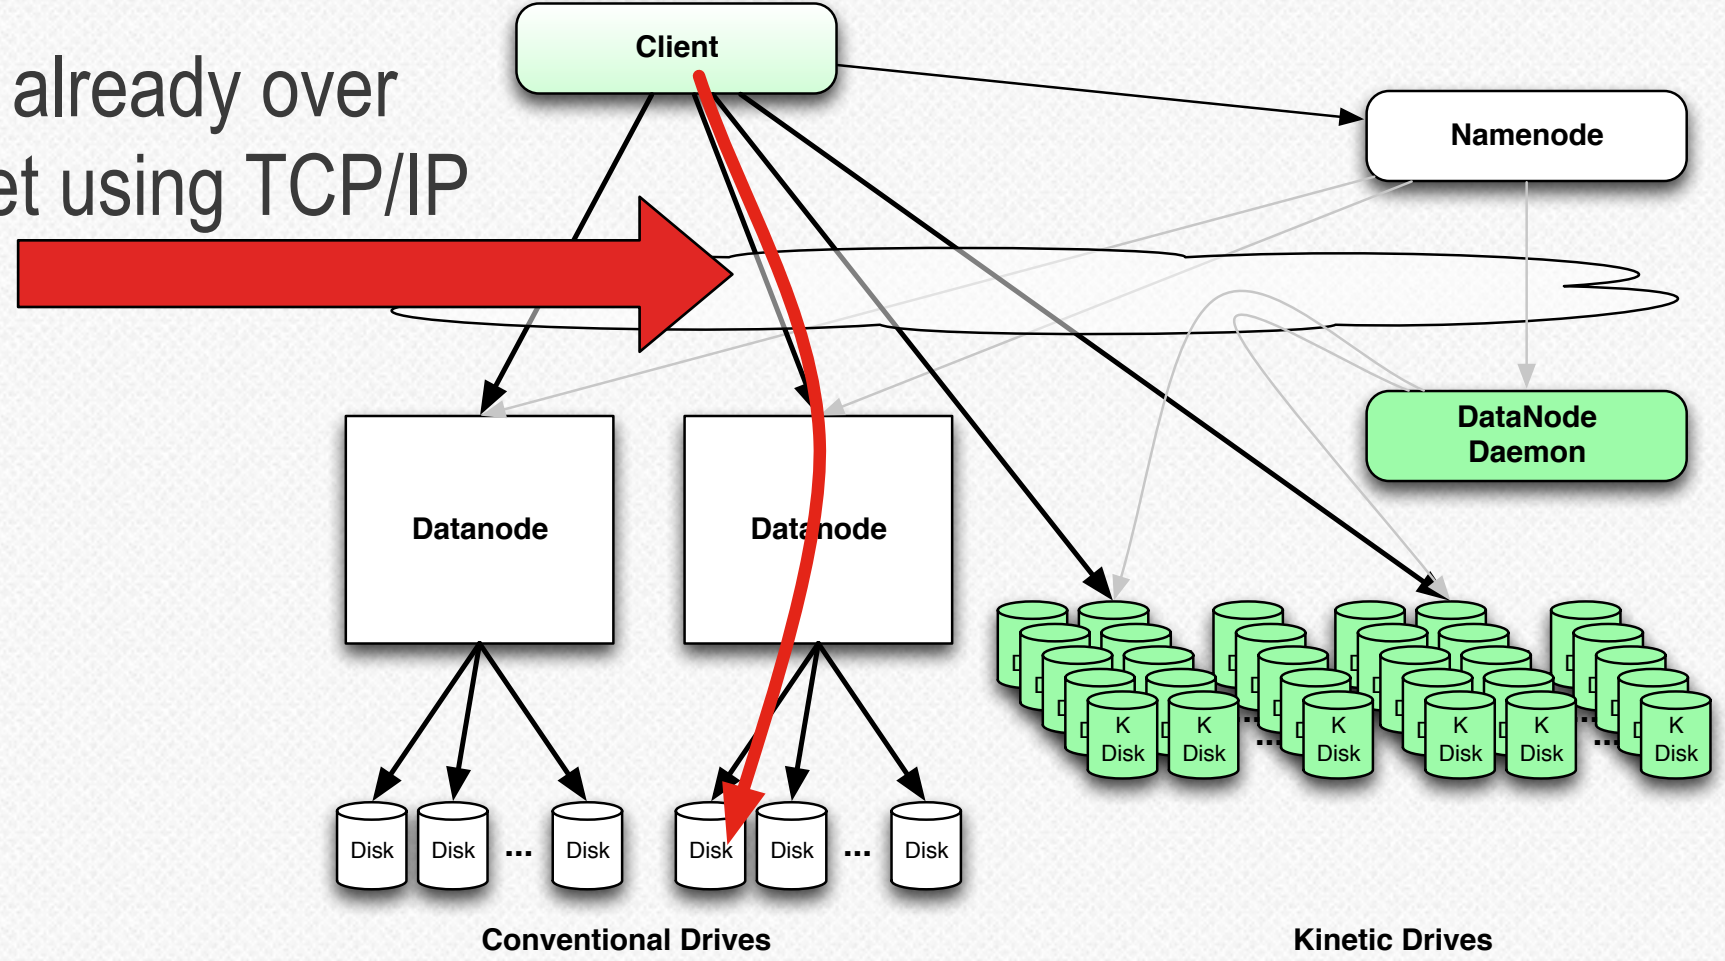


# Existing Traffic Flow



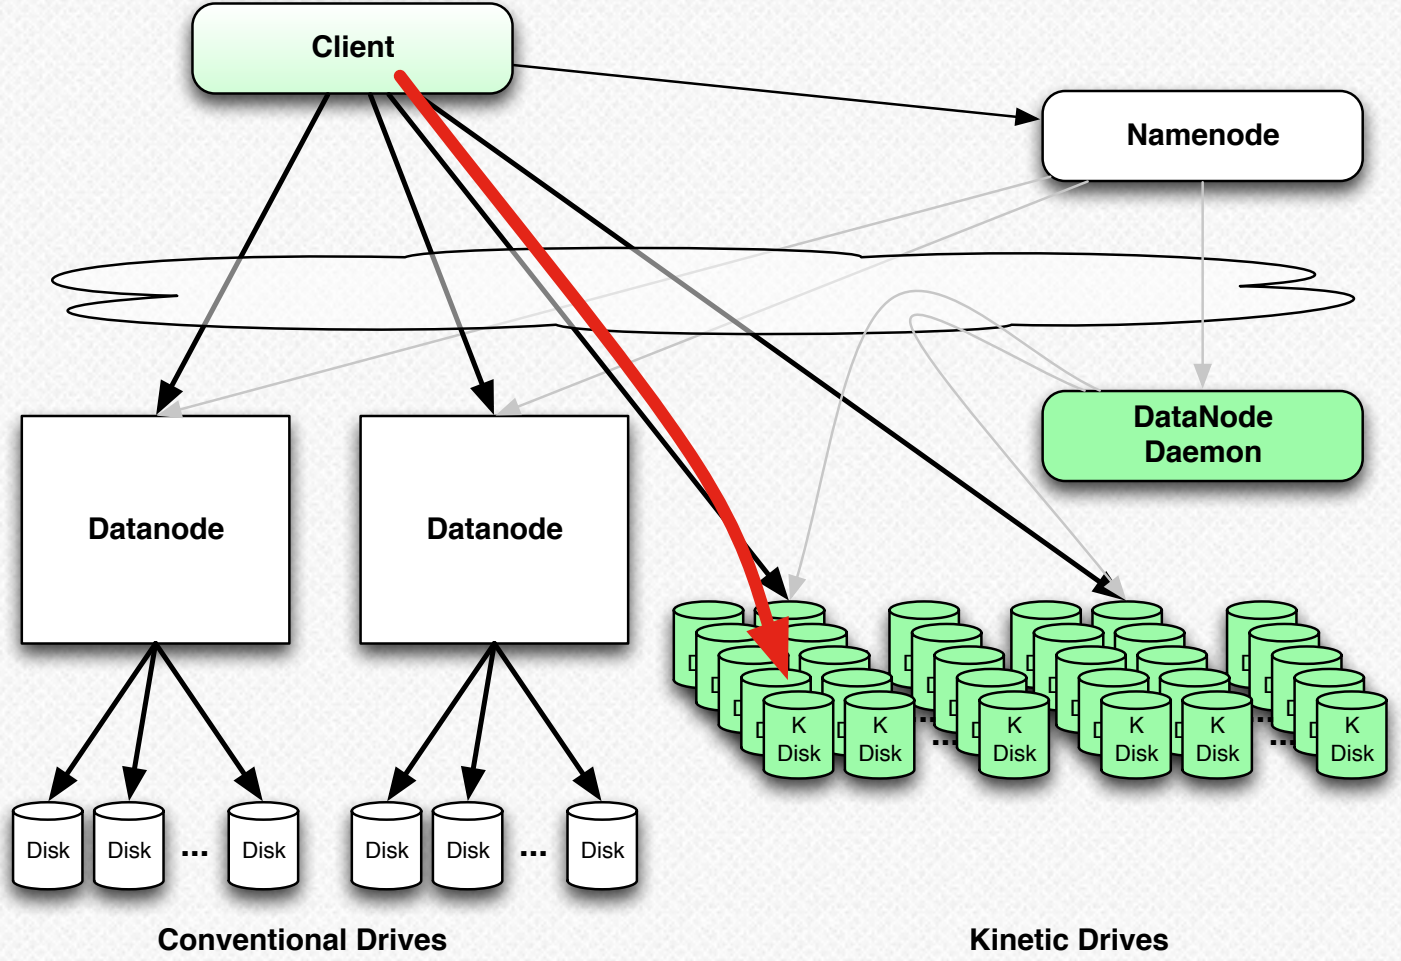
# Existing Traffic Flow

Data is already over Ethernet using TCP/IP

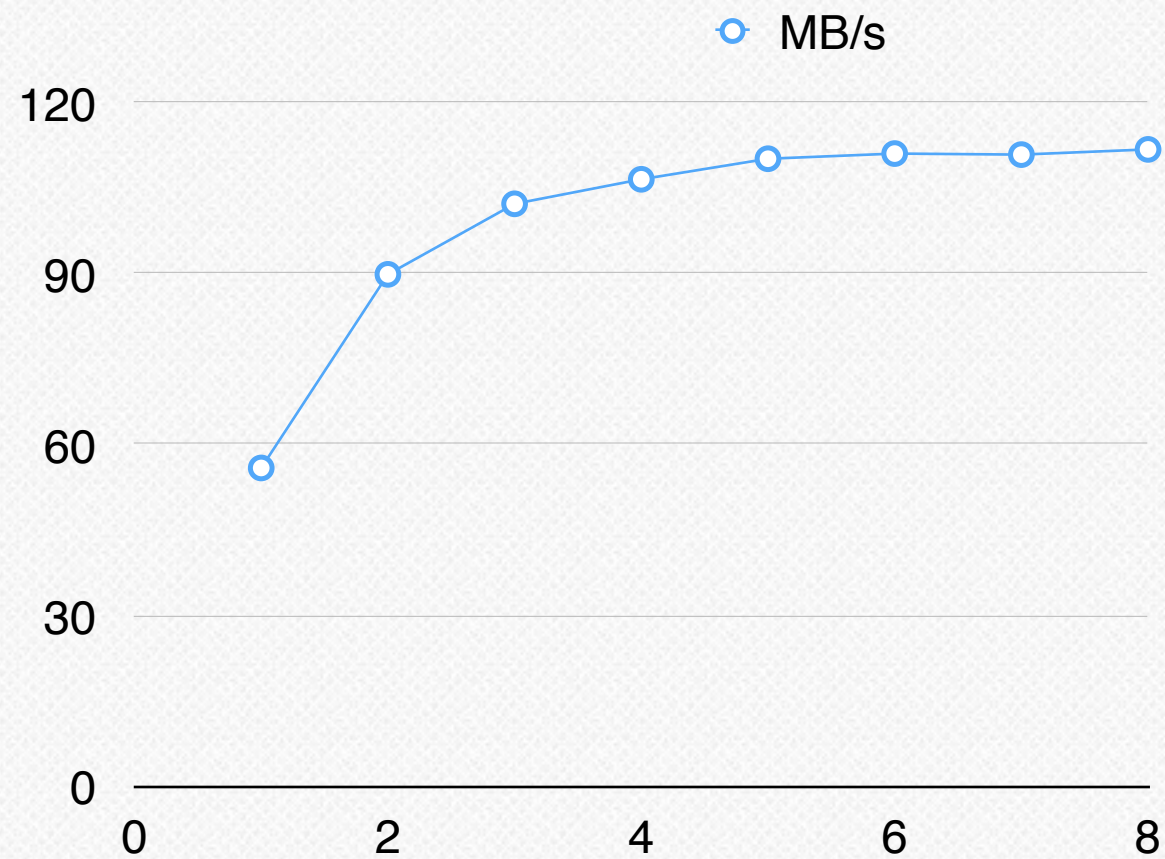




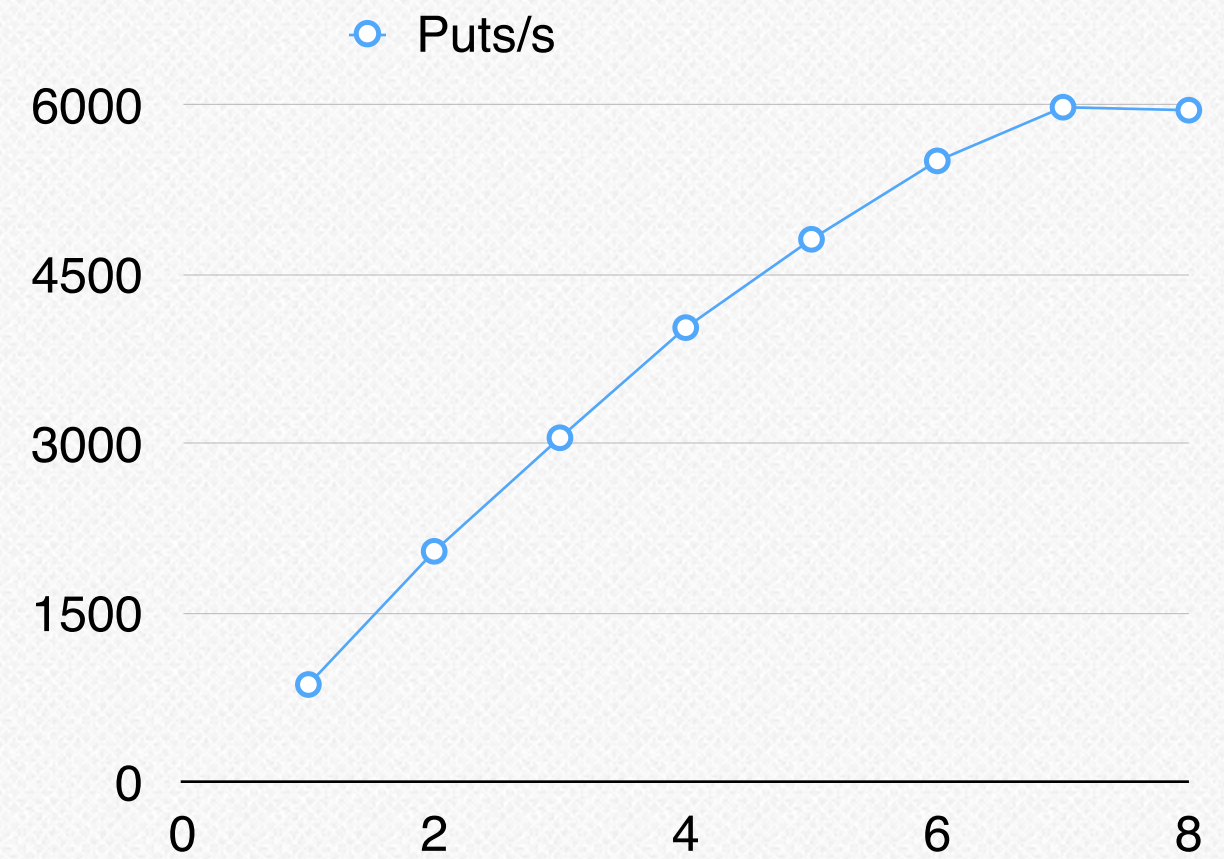
# Kinetic Traffic Flow



# Write Performance Results (preliminary)



1MB values put rate (MB/s)



10 byte value put rate





# Seagate Kinetic Open Storage Platform

Dis-intermediates applications to drive

- Goes around file systems, volume managers, drivers

Enable ecosystem of value added software

- Partners (like Basho) can create their own system value

Lower TCO

- Eliminates complexity



[Developers.Seagate.Com](http://Developers.Seagate.Com)





