



Solving Big Data Problems: Storage to the Rescue?

John Webster
Evaluator Group



- Big Data Analytics Storage Maxims
 - The Fundamental JBOD and DAS Architecture
 - Overview of Disk-based Alternatives
 - What are the Advantages and Disadvantages?
 - The Solid State and In-memory Alternatives
 - Summary and Q&A
-
- ◆ Note: References to specific vendors and products are used as real-world examples and do not imply an endorsement

Deliver storage performance at large scale
and at low cost, and all at the same time
(Think early stage Google, Facebook, Twitter)



Minimize the “distance” between processing
and data storage

Big Data analytics is dominated by open
source



Big Data Storage Maxim #4

Big Data analytics software developers manage data at the clustered server level. Storage vendors manage data at the storage system level.



Shared Nothing, Asymmetrical Distributed Computing

Only the Ethernet network is shared

Network Layer


1 Gb Ethernet

Compute Layer

Commodity Servers

Storage Layer

6-12 disks in each server typically JBOD


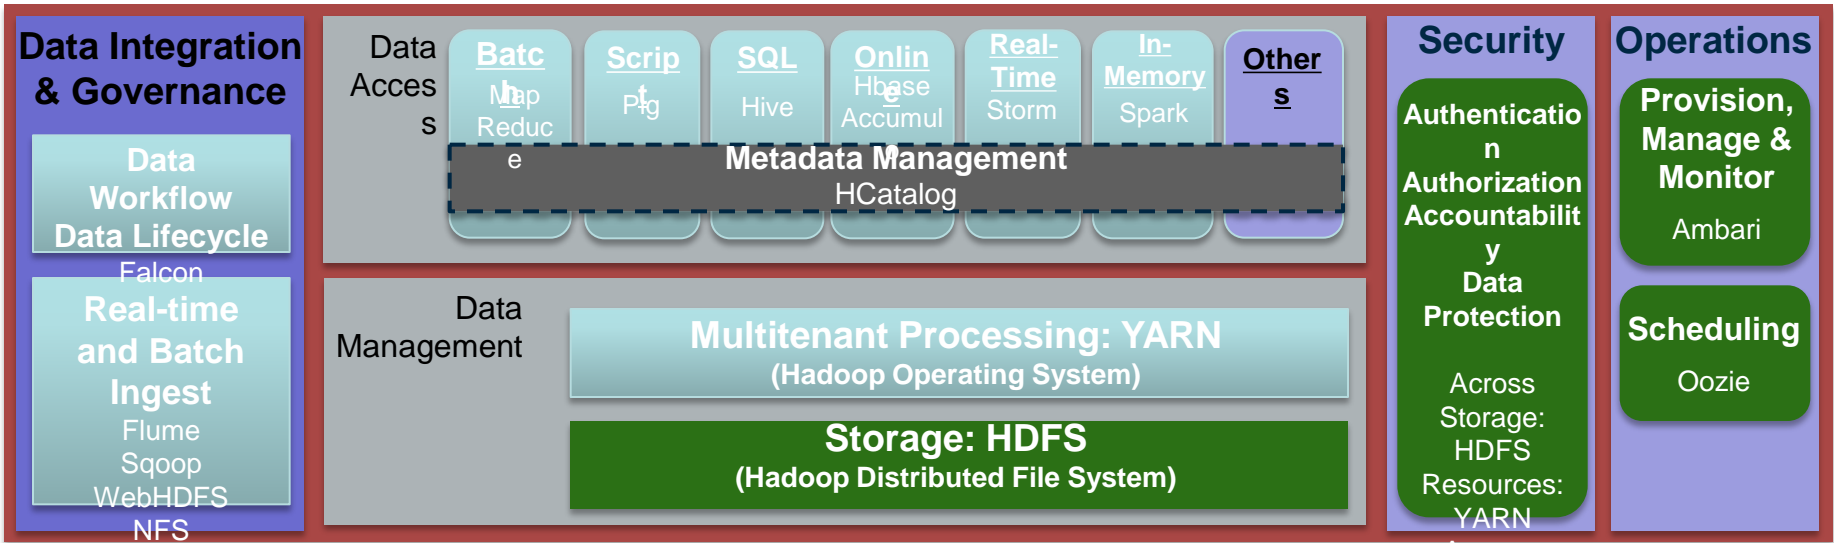


In Hadoop, Control = Name Node; Node 1,2... = Data Node

Apache Hadoop: A Platform for All Applications?

Presentation & Application
 Enable both existing and new applications to provide value to the organization

Operations
 Empower existing operations and security tools to manage



Source:
 Hortonworks

HDFS as a Persistent Storage Layer

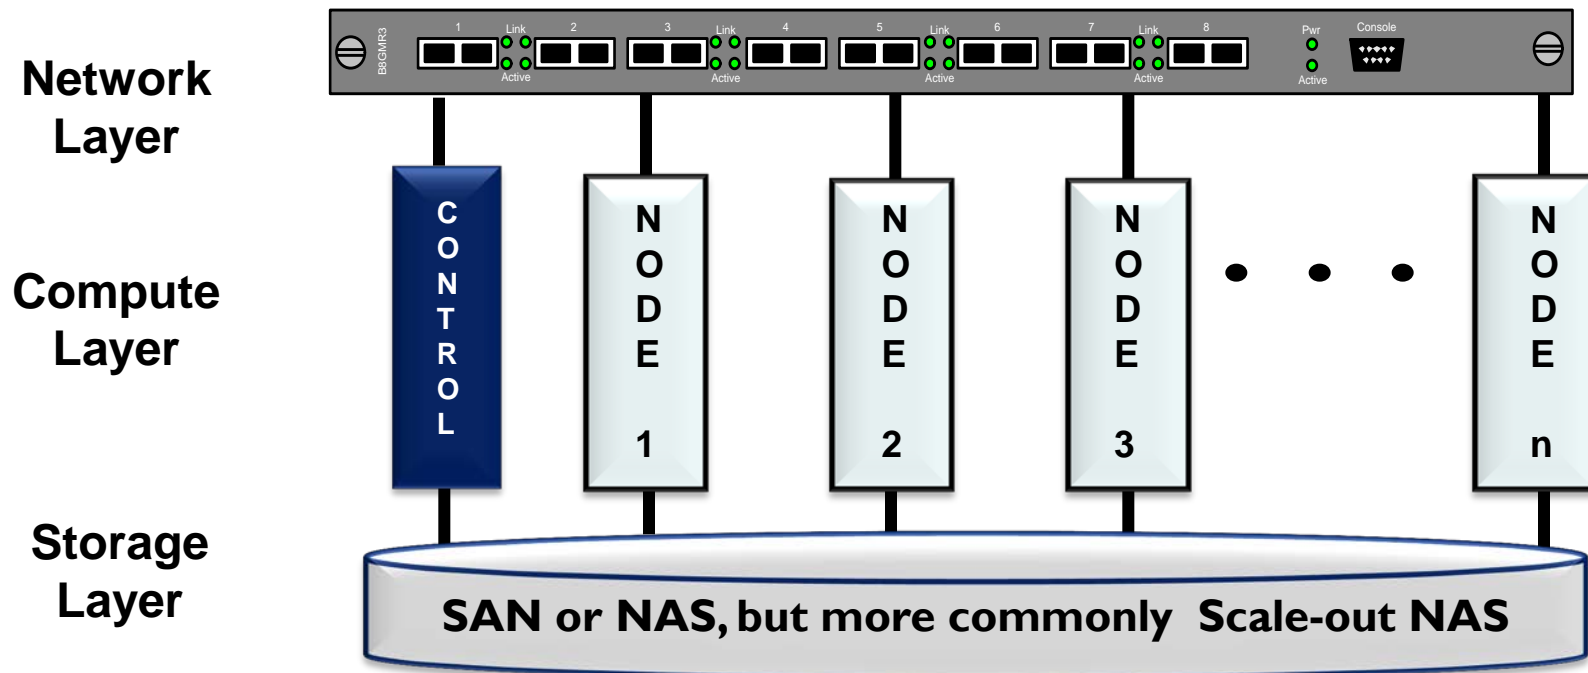
Advantages

- ◆ **Storage performance at large scale and low cost**
- ◆ **Minimize distance between data and compute**
- ◆ **Data node failures tolerated**
- ◆ **Open Source**

Disadvantages

- ◆ **Hadoop NameNode lacks active/active failover (i.e. it's a SPOF)**
- ◆ **For data integrity and protection, HDFS creates three full clone copies of data**
 - ◆ 3x the storage for each file – slow and inefficient
 - ◆ If all three copies are corrupted, you're still hosed (reload and start over)
- ◆ **No storage tiering (now available in 2.3)**
- ◆ **Limited ways to respond to corporate security and data governance policies**
- ◆ **Data in/out processes can take longer than the actual query process**
- ◆ **Inability to dis-aggregate storage from compute so that the two can be scaled independently**

Shared Storage as Primary Storage

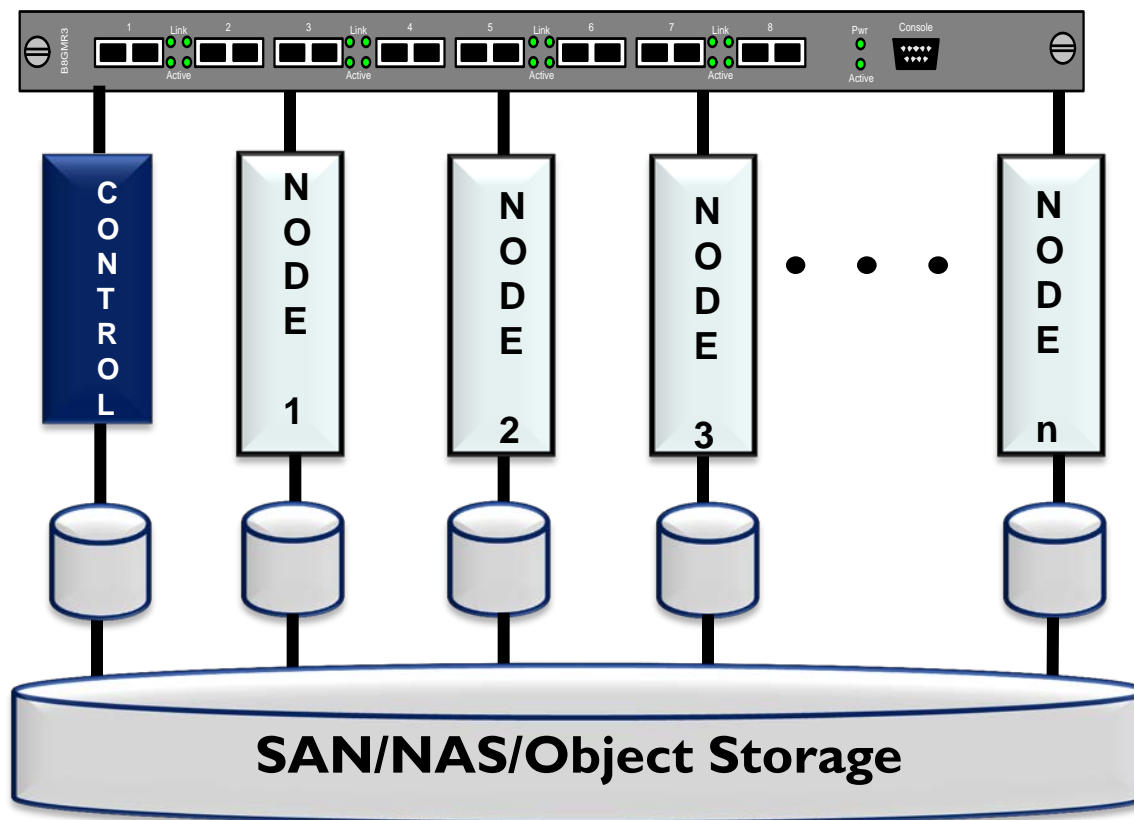


Shared Storage as Secondary Storage


Network Layer

Compute Layer

Storage Layer



Hadoop On Scale-out Storage



- Scale-out storage replaces node-level DAS
- HDFS implemented as “over the wire” protocol or CDMI interface to underlying FS
- NameNode SPOF eliminated
- Decoupled storage and compute layers
- Data services, data protection, and DR by storage-resident services
- Examples include EMC Isilon, IBM Elastic Storage, Ceph

➤ Advantages

- ◆ Addresses the enterprise storage management requirements
 - Data protection/disaster recovery/business continuance
 - Data governance and compliance
 - Digital records management and archiving

➤ Disadvantages

- ◆ Additional cost
- ◆ Potential performance impact
- ◆ Using a vendor specific solution introduces proprietary data/storage management software

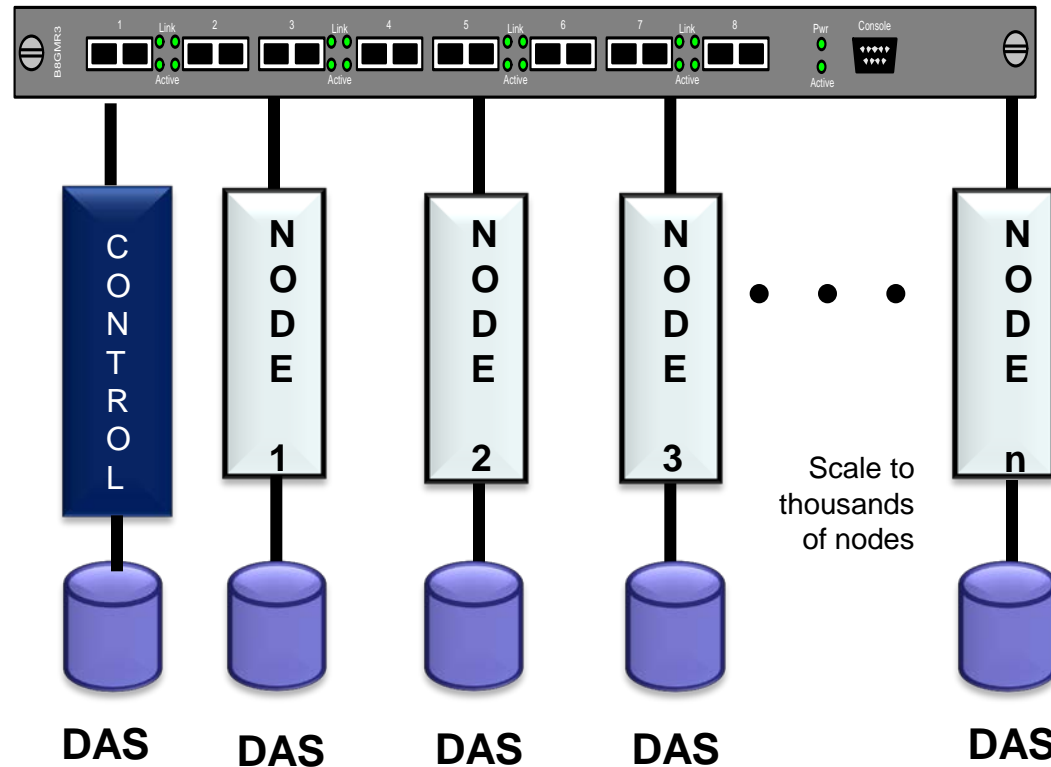
What About SSD?

Only the Ethernet network is shared

Network Layer
10+ Gb Ethernet

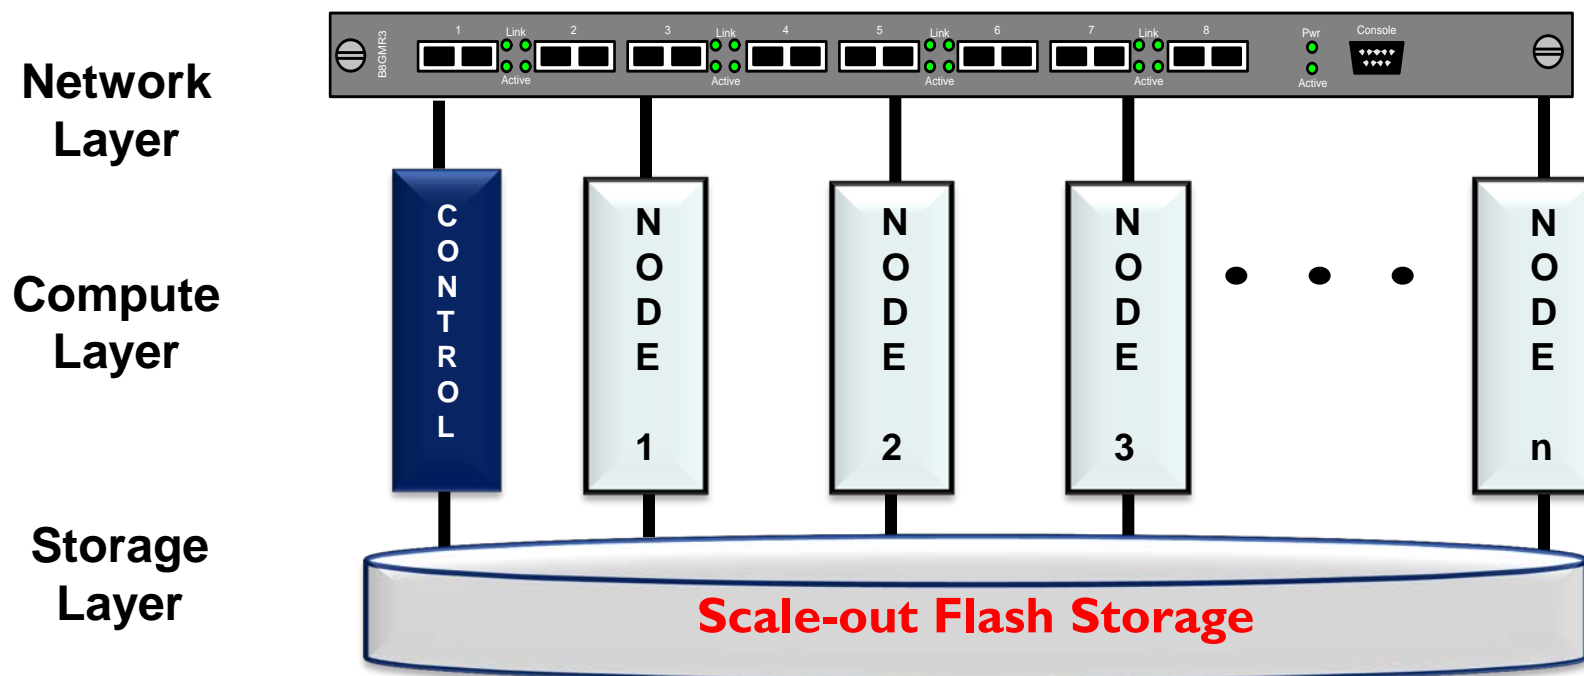
Compute Layer
Commodity Servers

Storage Layer
SSD
in/attached to each server



In Hadoop, Control = Name Node; Node 1,2... = Data Node

What About SSD?





➤ Tachyon

- ◆ UC Berkeley Amp Lab project
- ◆ “Reliable, memory-centric storage for Big Data Analytics clusters” (i.e. memory as persistent data store across cluster nodes)
- ◆ One in-memory data copy inside JVM, use operation “lineage” to re-compute data if failure
- ◆ Initial use in Apache Spark environments

What About In-memory Computing?

➤ Apache Ignite



- ◆ In-memory “data fabric”
- ◆ Distributed in-memory platform for computing and transacting on large-scale data sets in real-time
- ◆ “Orders of magnitude faster than possible with traditional disk-based or flash technologies.”
- ◆ Tier -1 storage?
- ◆ Originated as GridGain Data Fabric

- The need for a longer-term, persistent storage layer is now recognized
- For Hadoop, HDFS may or may not be that storage layer
- Enterprise storage architects and administrators will be more directly involved in managing Big Data analytics storage over time
- Now is the time to research and understand the options