

Spectrum Scale Hadoop Integration and Support for HortonWorks HDP

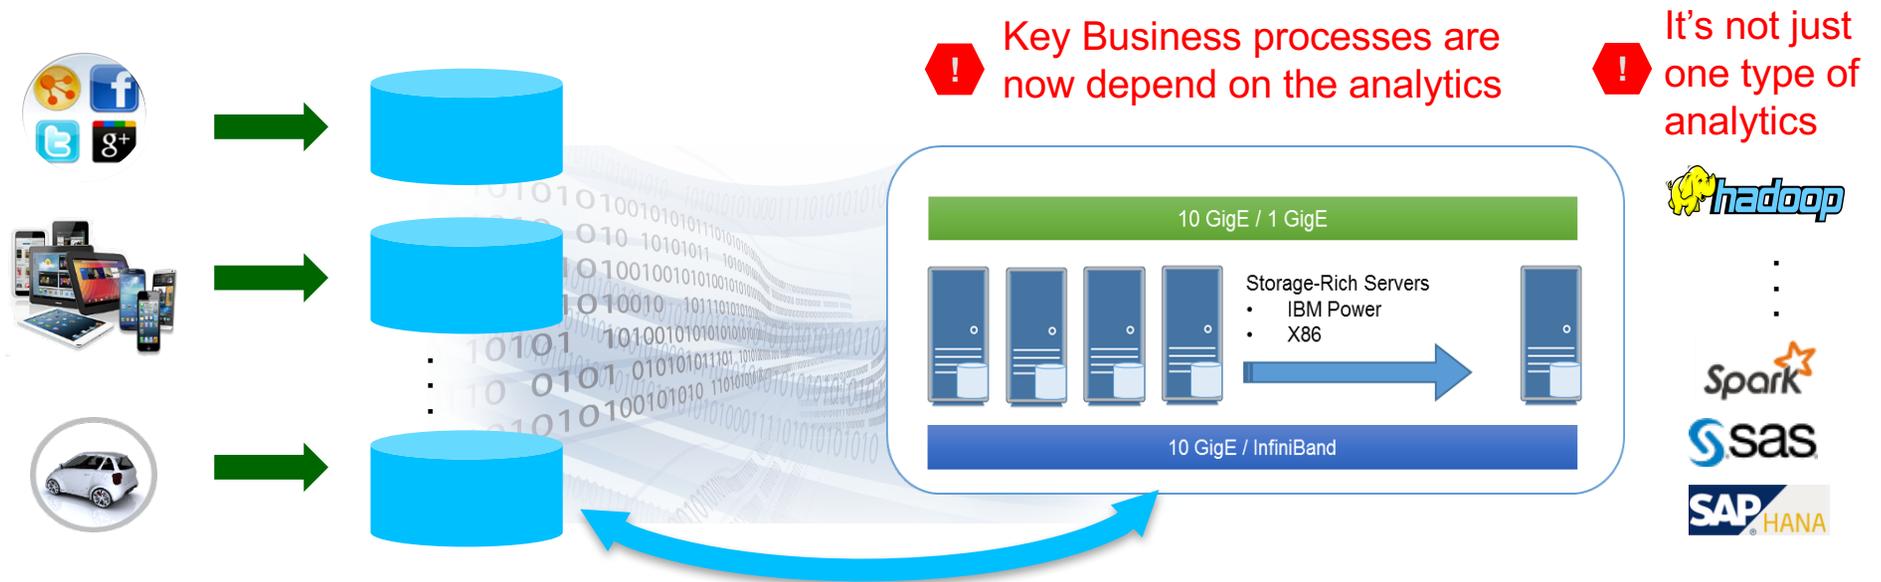
Ted Hoover

Manager

Spectrum Scale and High Performance Analytics
Solutions

Challenges with the Early Big Data Storage Models

IBM Storage & SDI



Ingest data at various end points

More data source than ever before, not just data you own, but public or rented data

Move data to the analytics engine

It takes hours or days to move the data!

Perform analytics

Can't just throw away metadata due to regulations or business requirement

Repeat!

Big Data Platform requires a “Data Ocean” Approach

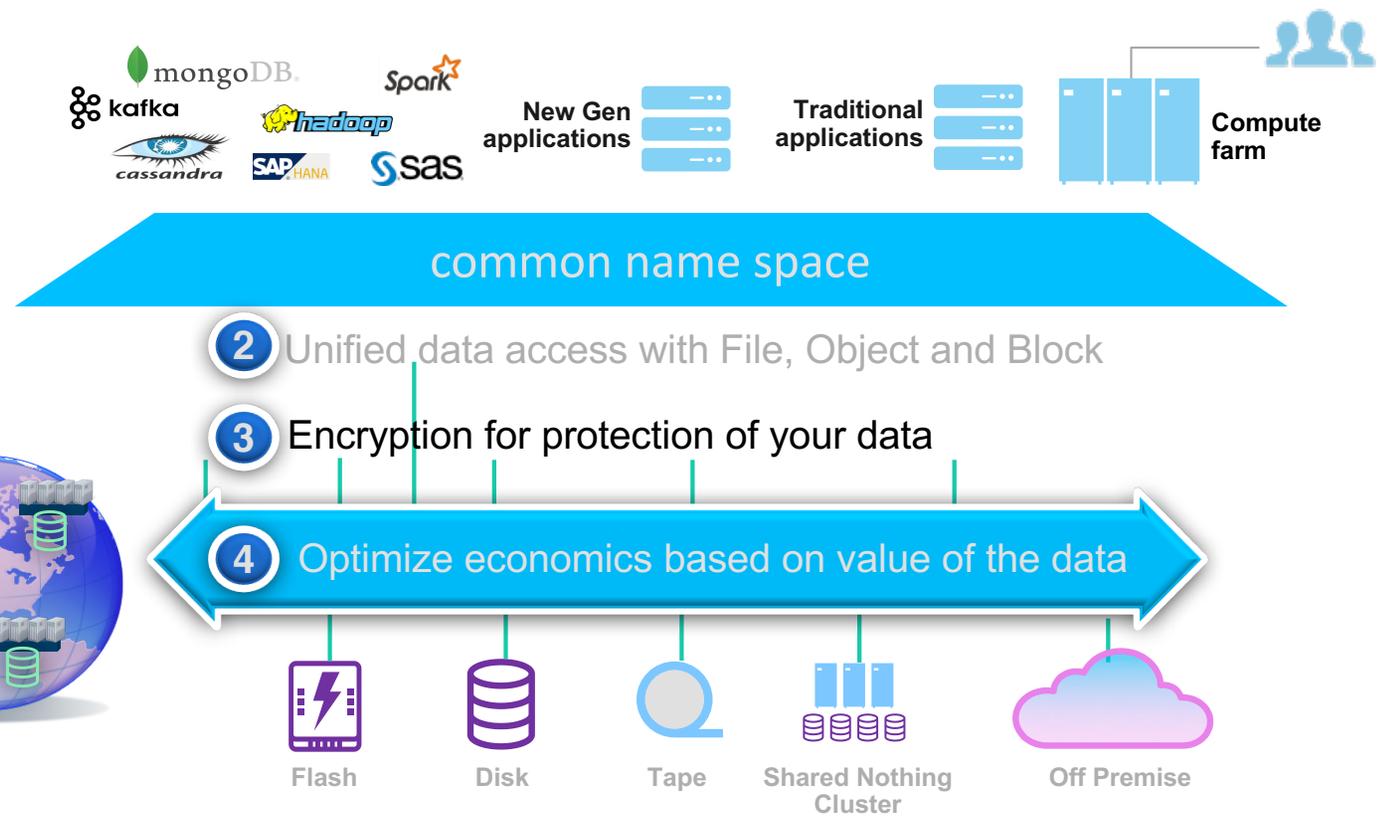
Fortune 2000 Enterprise

IBM Storage & SDI

6 Bring analytics to the data

1 Single Name Space to house all your data

5 Geographically dispersed management of data including disaster recovery



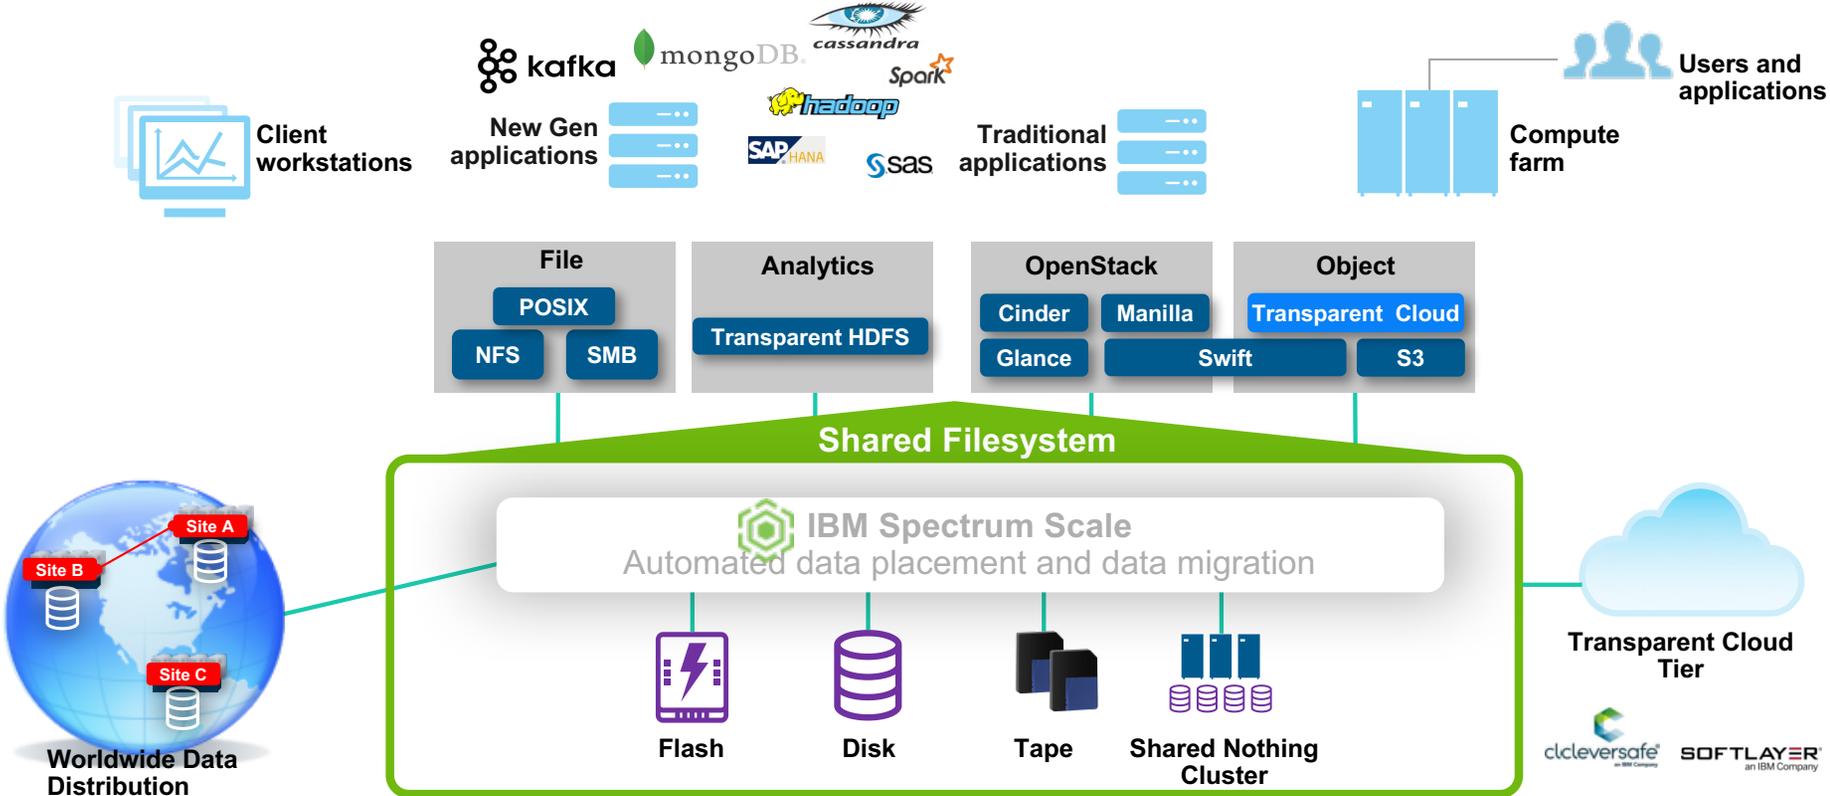
2 Unified data access with File, Object and Block

3 Encryption for protection of your data

4 Optimize economics based on value of the data

- Flash
- Disk
- Tape
- Shared Nothing Cluster
- Off Premise

Spectrum Scale: Unleash new storage economics on a global scale



4000+ customers using Spectrum Scale as data plane for HPC and analytics workload

Hortonworks announced the planned availability of Hortonworks Data Platform (HDP®) for IBM Elastic Storage Server (ESS) and IBM Spectrum Scale IBM Storage & SDI

Las Vegas, NV (IBM PartnerWorld) - 14 Feb 2017: IBM (NYSE: [IBM](#)) and. The agreement with Hortonworks will lead to certification of Hortonworks HDP on Power with IBM Spectrum Scale and Hortonworks HDP on x86 with IBM Spectrum Scale.

With the [recent announcement](#) of availability of Hortonworks HDP on POWER8 platform, IBM clients already have increased choice when selecting the platform for their Hadoop distribution. This new agreement will enable IBM clients to also leverage their existing and future investments in IBM storage in deploying Hadoop based big data applications. Additionally, this agreement will provide existing Hortonworks HDP customers an enterprise-class storage alternative with IBM Spectrum Scale for their Hadoop and Spark workloads.

The agreement represents the first IBM Storage offering and industry first enterprise Software Defined Storage solution certified for Hortonworks.

When completed, the HDP certifications will provide a significant layer of confidence to existing IBM and future clients. Those users can now run Hadoop applications on the leading software-defined storage solutions. IBM clients will have the benefits of enterprise storage to analyze data in place with Hortonworks analytics applications. With the choice of centralized or distributed deployments, organizations can improve business efficiency with the data management, backup, security and hybrid cloud storage.

<http://www-03.ibm.com/press/us/en/pressrelease/51562.wss>

Why are we partnering with Hortonworks?

- Hortonworks is one of the key distributions of Hadoop
 - Pure play 100% open source distribution
 - Hortonworks is #1 Apache Hadoop committer
 - Has 1000+ customers
 - ODPi compliance
 - Apache Spark is part of HDP distribution
- Spectrum Scale is already certified and supported with IBM's Hadoop distribution (IBM IOP/BigInsights)
- Hortonworks certifying Spectrum Scale/ESS will allow clients to confidently deploy Spectrum Scale/ESS for Big Data Analytics use cases powered by Hadoop

What value does Spectrum Scale provide to Hadoop?

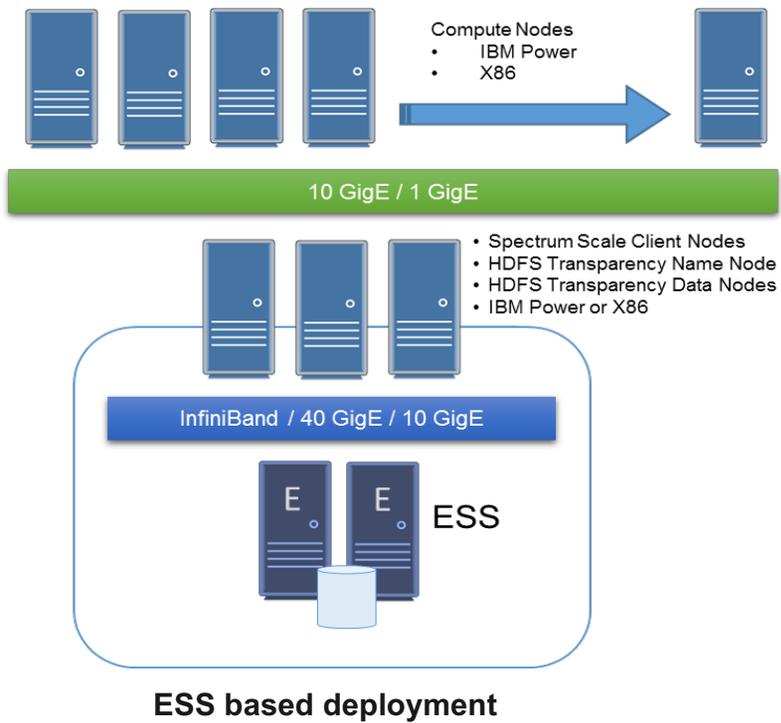
- The default storage for Hadoop is HDFS
- HDFS is Hadoop Distributed File System which runs on storage rich servers (storage internal to servers)
- Spectrum Scale provides a HDFS connector and allows existing Hadoop applications to run directly on Spectrum Scale
- This enables customers to create complex analytics workflows, minimize data movement & copies and speed up time to insight

HDFS Shortcomings addressed by Spectrum Scale

- HDFS is a shared nothing architecture, which is very inefficient for high throughput jobs (disks and cores grow in same ratio)
- Costly data protection:
 - uses 3-way replication; limited RAID/erasure coding
- Works only with Hadoop i.e weak support for File or Object protocols
- Clients have to copy data from enterprise storage to HDFS in order to run Hadoop jobs, this can result in running on stale data.

Key Deployment Models

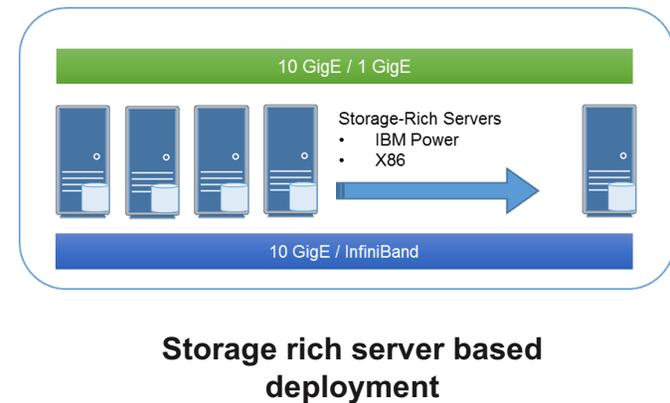
Shared Storage Architecture



2017 enablement focus

IBM Storage & SDI

Shared Nothing Architecture



Why Spectrum Scale

- Reduce the datacenter footprint

Spectrum Scale provides **in-place analytics** for your Hadoop data eliminating the need for multiple copies of the same data or large migrations of data between HDFS and the POSIX file system.

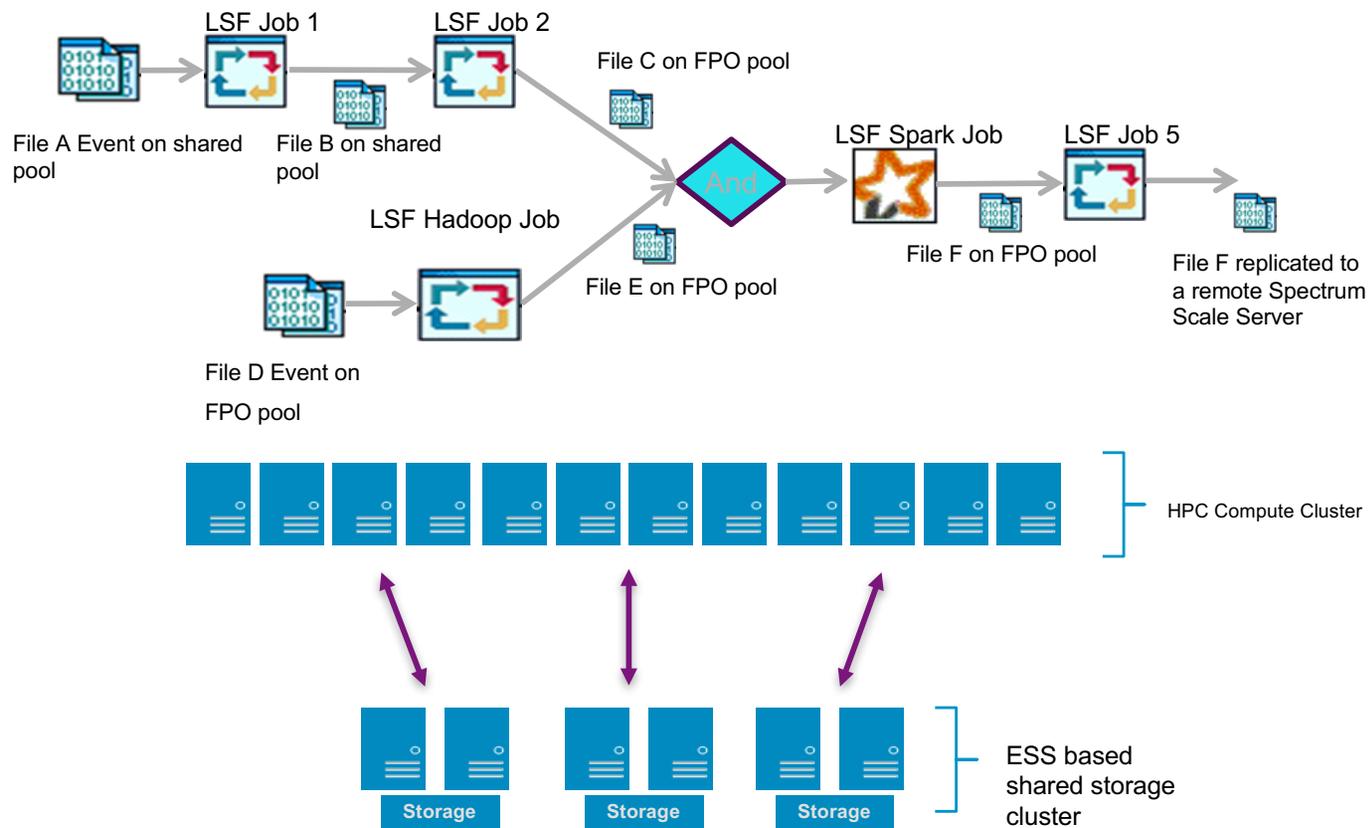
- Control cluster sprawl

Spectrum Scale delivers **federation across clusters**, both Scale and HDFS, turning isolated data lakes into a data ocean while still maintaining needed separations.

- Make HDFS access enterprise-ready

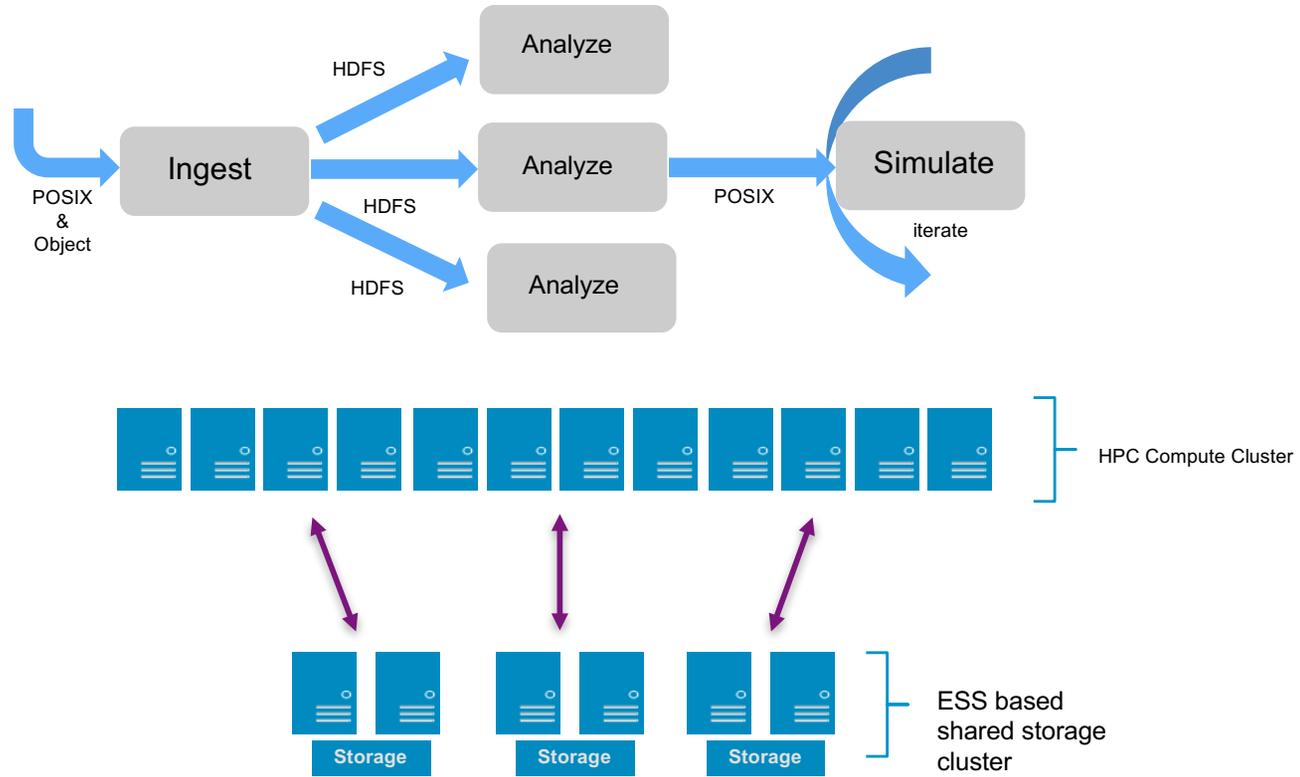
IBM Elastic Storage Server (Pre-integrated Spectrum Scale based system) adds **storage functions necessary to the enterprise** (e.g. Encryption, DR, SW RAID) to your Hadoop setup.

Client Use Case: Life Sciences with HPC and Hadoop/Spark

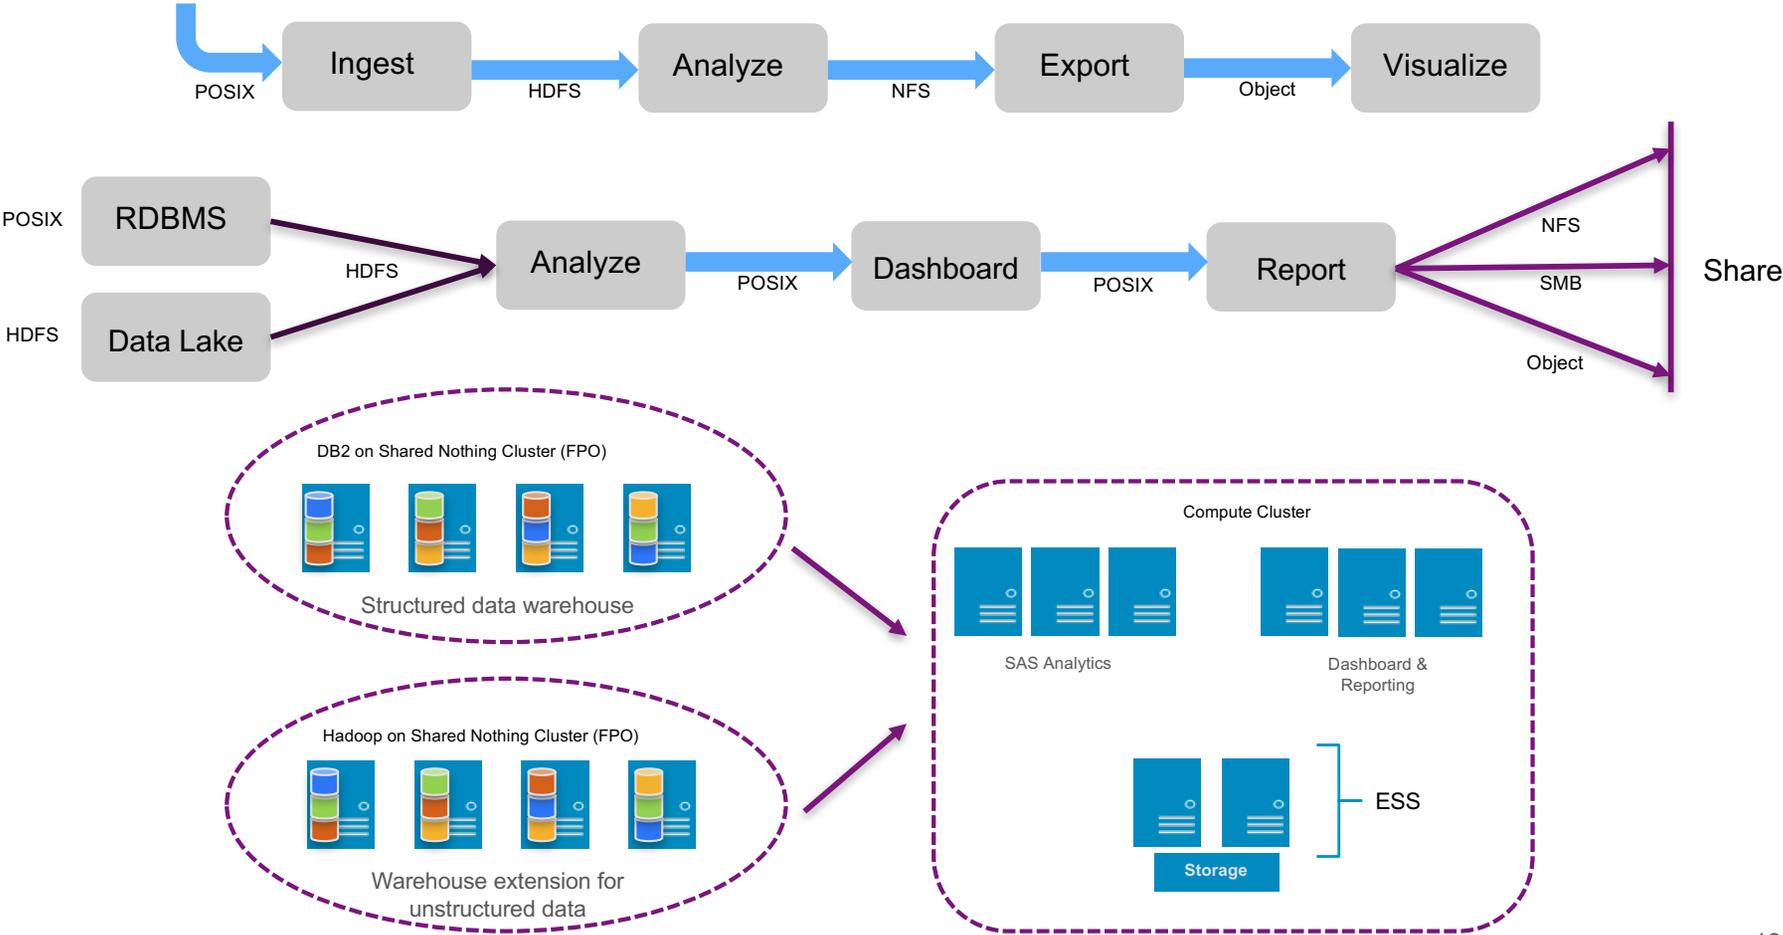


Client Use Case: HPC and Data Analytics

IBM Storage & SDI



Client Use Case : Unified Analytic/Workflow Pipelines



Federate Spectrum Scale Filesystem with Existing HDFS Filesystem

Extending the Filesystem

- Run analytics across multiple HDFS and/or Spectrum Scale clusters
 - No need to move the data
 - Applications can still access data that has been moved seamlessly

viewfs://clusterX:/hadoop/hdfs/file1
viewfs://clusterY:/hadoop/gpfs/file1



hdfs://nn1.node.net:8020/hadoop/hdfs



Cluster X

hdfs://nn2.node.net:8020/hadoop/gpfs



Cluster Y

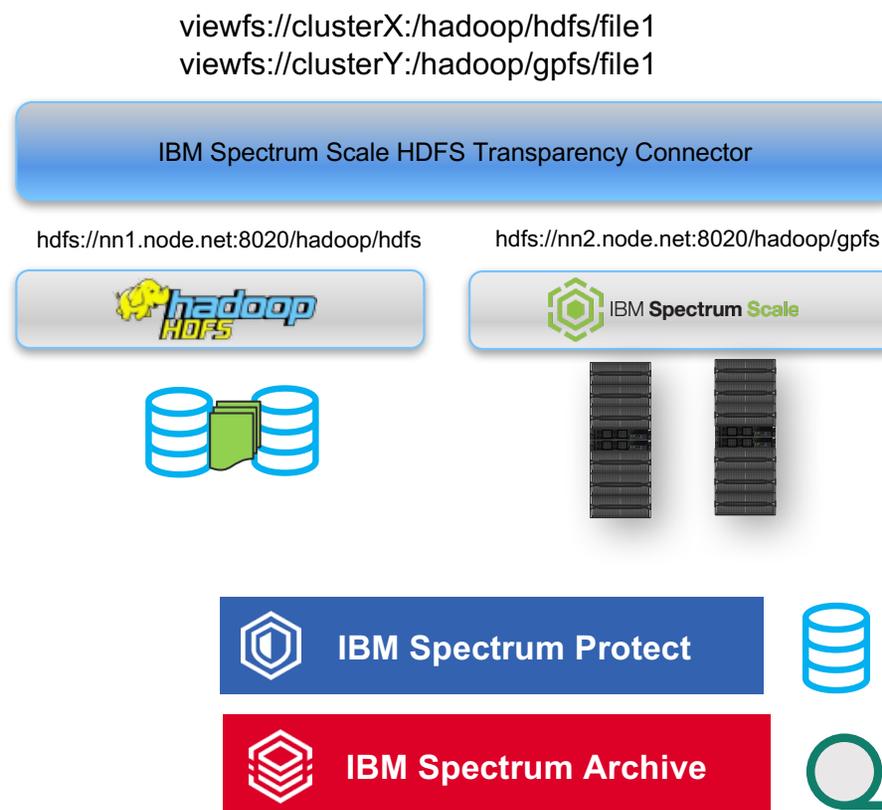
Commands:

```
$ hadoop distcp viewfs://clusterX:/hadoop/hdfs/file1  
viewfs://clusterY:/hadoop/gpfs/file1  
$ Hadoop fs rm viewfs://clusterX:/hadoop/hdfs/file1
```

Spectrum Scale Makes HDFS Enterprise Ready

Protecting Business Data

- Use ESS warm data tier with Spectrum Archive to tape
- Powerful policy engine
 - Information Lifecycle Management
 - Fast metadata 'scanning' and data movement
 - Automated data migration to based on threshold
- Users not affected by data migration
 - Single namespace
- Optionally Spectrum Protect and Spectrum Archive can be used directly with Spectrum Scale FPO



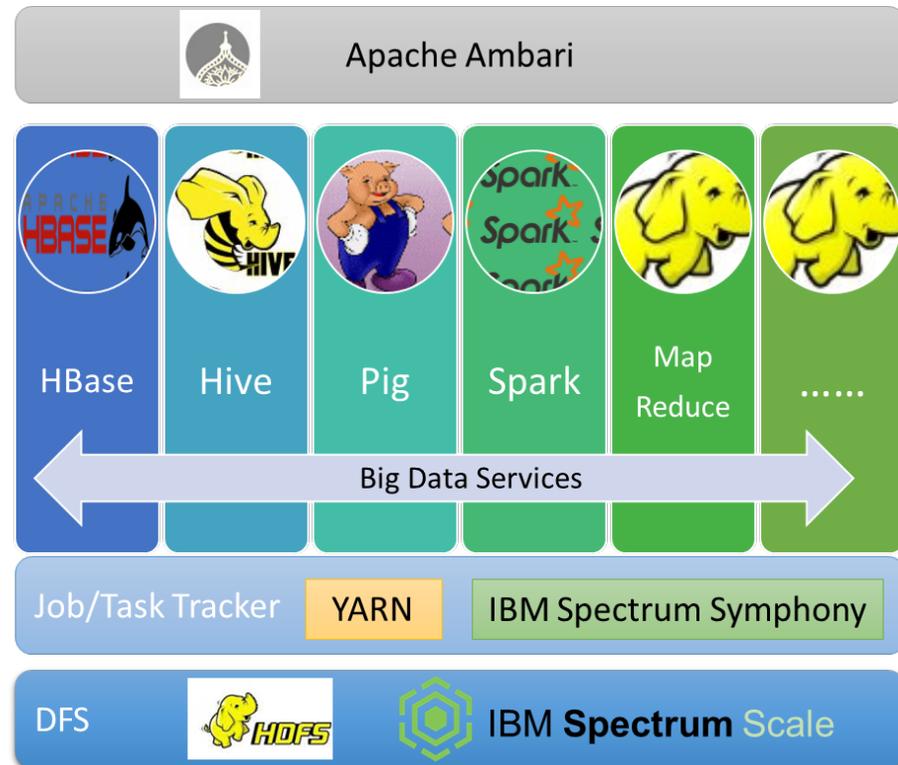
Hadoop Data Analytics with IBM Spectrum Scale

IBM Spectrum Scale is a proven enterprise ready alternative of HDFS for Hadoop data analytics

Completely compatible with HDFS client API and Hadoop applications

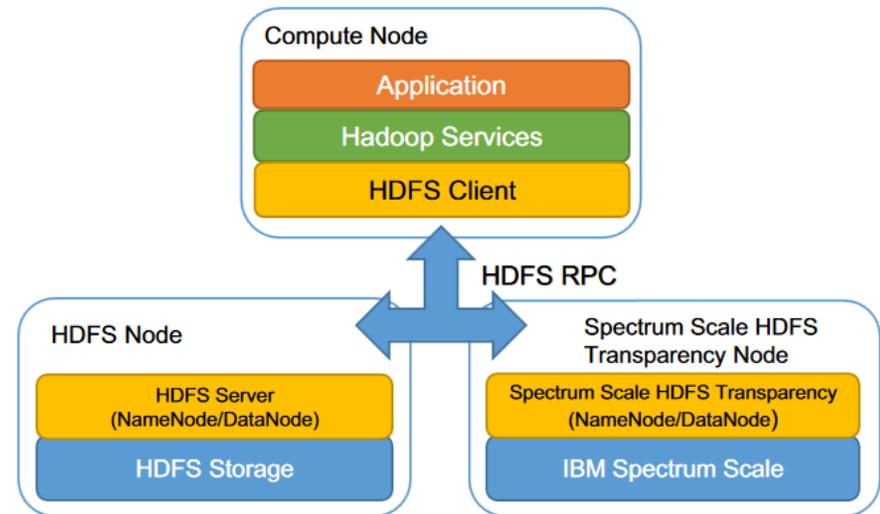
Unified File, Object and HDFS access interface

Mature enterprise features make it easy to manage and protect data



Spectrum Scale HDFS Transparency Connector

- Integrates with HDFS – reuses HDFS client and implements NameNode and DataNode RPCs
- Stateless
- Free download – [link](#)
- Supports Spectrum Scale 4.1.x and 4.2.x
- Currently supported with IOP 4.1 and 4.2
- Ambari integration available for both IOP 4.1 and 4.2



Spectrum Scale Ambari Integration

The screenshot displays the Ambari web interface for a cluster named 'mycluster'. The top navigation bar includes 'Dashboard', 'Services', 'Hosts', 'Alerts', and 'Admin'. The 'Services' tab is active, showing the configuration for 'Spectrum Scale'. The interface is divided into a left sidebar with a list of services and a main content area. The 'Spectrum Scale' service is highlighted in the sidebar. The main content area shows the 'Summary' and 'Metrics' sections. The 'Summary' section indicates that the 'GPFS Master' is 'Started' and '3/3 GPFS Nodes Live'. The 'Metrics' section displays four key metrics: 'Filesystem Utilization' at 0%, 'Inode Utilization' at 0%, 'Active Quorum Nodes' at 3/3, and 'Active NSD Nodes' at 3/3. A 'Service Actions' dropdown menu is open, listing various actions such as 'Start', 'Stop', 'Restart All', and 'Upgrade_SpectrumScale'. The 'Spectrum Scale' service is listed in the sidebar with a green checkmark, indicating it is running.

Ambari mycluster 0 ops 0 alerts Dashboard Services Hosts Alerts Admin admin

Summary Configs Quick Links Service Actions

Summary

- [GPFS Master](#) Started
- [GPFS Nodes](#) 3/3 GPFS Nodes Live

Metrics

Filesystem Utilization	Inode Utilization	Active Quorum Nodes	Active NSD Nodes
0%	0%	3/3	3/3

Service Actions:

- Start
- Stop
- Restart All
- Restart GPFS Nodes
- Run Service Check
- Turn On Maintenance Mode
- Collect_Snap_Data
- Upgrade_SpectrumScale
- Upgrade_Transparency
- Integrate_Transparency
- UnIntegrate_Transparency

Spectrum Scale

Actions

Ambari Integration with HDFS Transparency: Overview

- Spectrum Scale is added as a new service after full IOP install with HDFS
- Spectrum Scale service “integrates” with HDFS
- Supports “un-integrate” capability
 - Flip back and forth between HDFS & GPFS
 - Will not move data back and forth between HDFS & GPFS
- Will simplify future upgrades

Add Service Wizard

ADD SERVICE WIZARD

Choose Services

- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Configure Identities
- Review
- Install, Start and Test
- Summary

Choose Services

Choose which services you want to install on your cluster.

<input type="checkbox"/> Service	Version	Description
<input checked="" type="checkbox"/> HDFS	2.7.2	Apache Hadoop Distributed File System
<input checked="" type="checkbox"/> YARN + MapReduce2	2.7.2	Apache Hadoop NextGen MapReduce (YARN)
<input checked="" type="checkbox"/> Hive	1.2.1	Data warehouse system for ad-hoc queries & analysis of large datasets and table & storage management service
<input checked="" type="checkbox"/> HBase	1.2.0	Non-relational distributed database and centralized service for configuration management & synchronization
<input checked="" type="checkbox"/> Pig	0.15.0	Scripting platform for analyzing large datasets
<input checked="" type="checkbox"/> Sqoop	1.4.6	Tool for transferring bulk data between Apache Hadoop and structured data stores such as relational databases
<input checked="" type="checkbox"/> Oozie	4.2.0	System for workflow coordination and execution of Apache Hadoop jobs. This also includes the installation of the optional Oozie Web Console which relies on and will install the ExtJS Library .
<input checked="" type="checkbox"/> ZooKeeper	3.4.6	Centralized service which provides highly reliable distributed coordination
<input checked="" type="checkbox"/> Flume	1.6.0	A distributed service for collecting, aggregating, and moving large amounts of streaming data into HDFS
<input checked="" type="checkbox"/> Titan	1.0.0	Titan is a scalable graph database optimized for storing and querying graphs containing hundreds of billions of vertices and edges distributed across a multi-machine cluster.
<input checked="" type="checkbox"/> Ambari Metrics	0.1.0	A system for metrics collection that provides storage and retrieval capability for metrics collected from the cluster
<input type="checkbox"/> Spectrum Scale	4.2.0	High-performance, scalable storage manages yottabytes of unstructured data (formerly known as General Parallel File System, or GPFS)
<input checked="" type="checkbox"/> Kafka	0.9.0.1	A high-throughput distributed messaging system
<input checked="" type="checkbox"/> Knox	0.7.0	Provides a single point of authentication and access for Apache Hadoop services in a cluster
<input type="checkbox"/> Ranger	0.5.2	Comprehensive security for Hadoop
<input type="checkbox"/> Ranger KMS	0.5.2	Key Management Server

HDFS Transparency NameNode and DataNodes

Ambari mycluster 0 ops 0 alerts Dashboard Services Hosts Alerts Admin admin

Summary Heatmaps Configs Quick Links - Service Actions

Summary

NameNode Started	Disk Remaining	2.0 TB / 2.0 TB
DataNodes 3/3 Started	Blocks (total)	0
DataNodes Status 3 live / 0 dead / 0 decommissioning	Block Errors	0 corrupt / 0 missing
NFSGateways 0/0 Started	Total Files + Directories	0
NameNode Uptime 464.48 secs	Upgrade Status	No pending
NameNode Heap 129.0 MB / 1011.3 MB (12.8% used)	Safe Mode Status	Not in safe mode
Disk Usage (DFS Used) 4.6 GB / 2.0 TB (0.23%)		
Disk Usage (Non DFS Used) 0 Bytes / 2.0 TB (0.00%)		

Service Actions

- ▶ Start
- Stop
- ↻ Restart All
- ↻ Restart DataNodes
- 🔍 Run Service Check
- 🔧 Turn On Maintenance Mode
- 🔄 Rebalance HDFS
- 📄 Download Client Configs

Metrics Actions Last 1 hour

NameNode GC count 	NameNode GC time 	NN Connection Load 	NameNode Heap 	NameNode Host Load
NameNode RPC 	Failed disk volumes 0	Corrupted Blocks 0	Under Replicated Blocks 0	DataNode JVM Heap Memory Used 115.65MB

Ambari Metrics Spectrum Scale Kafka Knox Slider Solr Spark

Big Data Oceans extending HDFS across Clusters

Unified Data Repository, Support Multiple Analytics



Extending the Filesystem

- Run analytics across multiple HDFS and/or Spectrum Scale clusters
 - No need to move the data
 - Build Data Oceans on demand

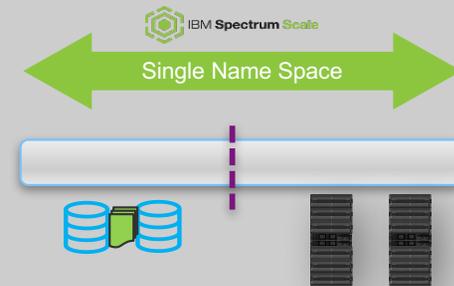
Federate ESS with Existing HDFS Filesystem



Improve Hadoop Cluster Utilization

- Manually move less frequently accessed data to an ESS tier
- Applications can still access data that has been moved seamlessly

Federate ESS with Existing Spectrum Scale Filesystem



Expand use of Shared Nothing Clusters

- Simplicity of Storage Rich Servers with enterprise features
- Advanced Routing (AFM), encryption, QoS, compression
- Mix cluster types
- Backup and Archive Support

Spectrum Scale Wiki Links

[Spectrum Scale wiki – Analytics](#)

[Spectrum Scale wiki - FPO](#)

[Spectrum Scale wiki – Hadoop](#)

[Spectrum Scale wiki – HDFS Transparency connector](#)

Thank You.
IBM Storage & SDI

