



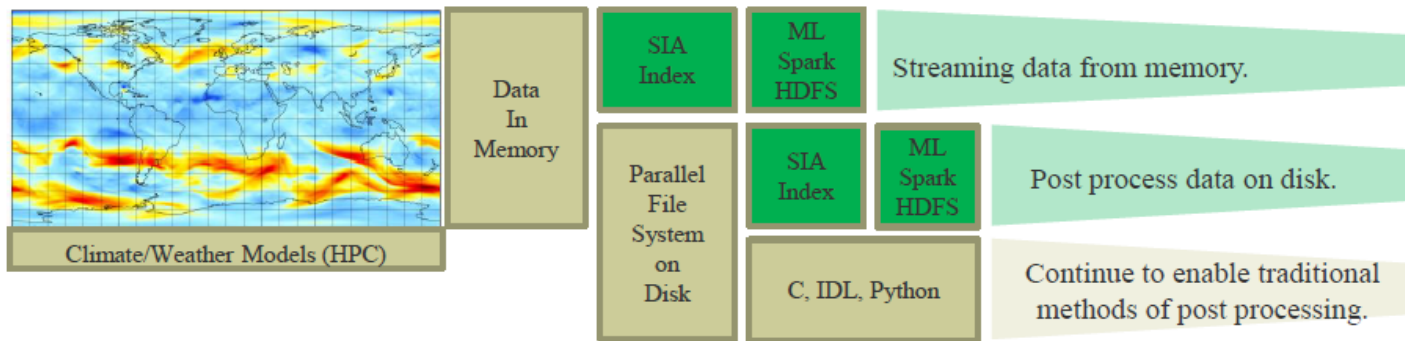
IBM **Spectrum Scale**

IBM Spectrum Scale HDFS Support

IBM Spectrum Scale German User Meeting 2017
Mar 8th+9th, 2017

Ulf Troppens

Future of Data Analytics




- Future HPC systems must be able to efficiently transform information into knowledge using both traditional analytics and emerging *machine learning* techniques.
- Requires the ability to be able to index data in memory and/or on disk and enable analytics to be performed on the data where it resides – even in memory
- All without having to modify the data

Essentials


- 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 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.

→ The Spectrum Scale Transparency Connector brings analytics to the data.

Unleash New Storage Economics on a Global Scale



Single name space




IBM Spectrum Scale

Automated data placement and data migration



Active = Hot+Warm = "Online"

Archive = Cold = "Offline"

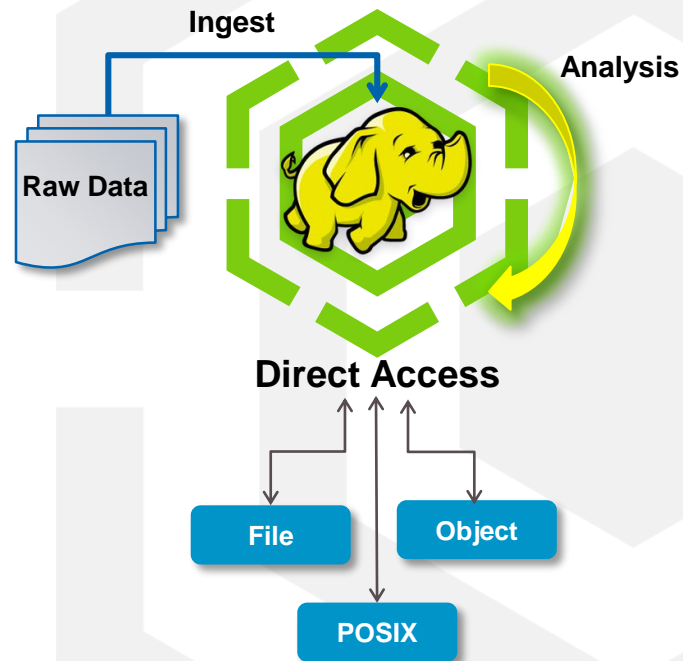


Analytics without complexity

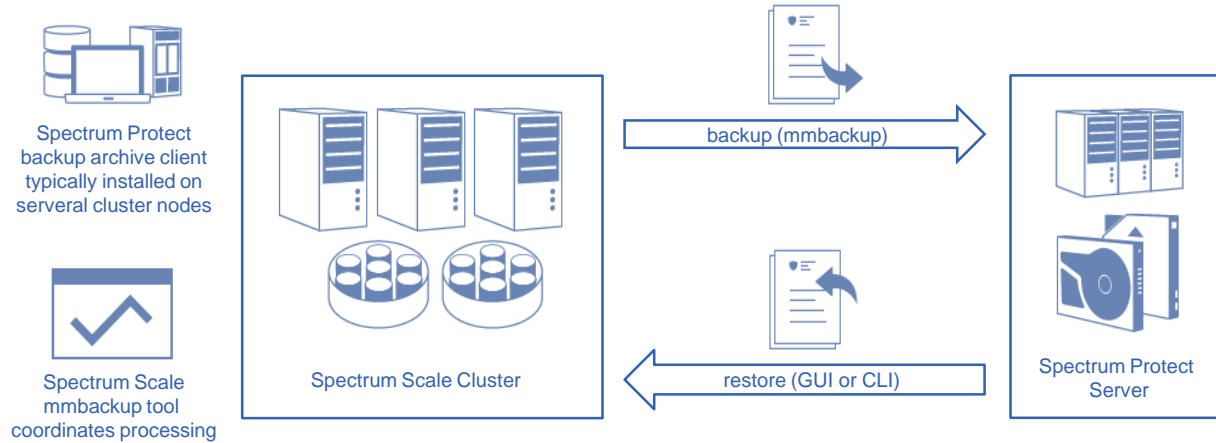
Challenge

- Separate storage systems for ingest, analysis, results
 - HDFS requires locality aware storage (namenode)
 - Data transfer slows time to results
 - Different frameworks & analytics tools use data differently
- **HDFS Transparency**
 - Map/Reduce on shared, or shared nothing storage
 - No waiting for data transfer between storage systems
 - Immediately share results
 - Single 'Data Lake' for all applications
 - Enterprise data management
 - Archive and Analysis in-place

→ Analyze object and file data without copying into HDFS



Backup Of Large Spectrum Scale File Systems



Function
Massive parallel filesystem backup processing
Spectrum Scale mmbackup creates local shadow of Spectrum Protect DB and uses policy engine to identify files for backup
Spectrum Protect backup archive client is used under the hood to backup files to Spectrum Protect Server
Spectrum Protect restore (CLI or GUI) can be used to restore files

- ➔ Use any backup program to backup file, object and Hadoop data
- ➔ Use Spectrum Protect to benefit from mmbackup and SOBAR to backup and restore huge amounts of data



IBM Delivers New Platform to Help Clients Address Storage Challenges at Massive Scale

Las Vegas, NV (IBM PartnerWorld) – 14 Feb 2017: IBM (NYSE: [IBM](#)) and Hortonworks (NASDAQ: HDP) today announced the planned availability of Hortonworks Data Platform (HDP®) for IBM Elastic Storage Server (ESS) and IBM Spectrum Scale. 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.

<https://hortonworks.com/press-releases/ibm-delivers-new-platform-help-clients-address-storage-challenges-massive-scale/>

A Tale of Two Connectors

GPFS Hadoop Connector

- Henceforth known as the “old” connector
- Emulates a Hadoop compatible filesystem – i.e. replaces HDFS
- Stateless
- Free download – [link](#)
- Supports Spectrum Scale 4.1.x, 4.1.1.x and 4.2
- Currently supported with IOP 4.0.x and 4.1.x
- Integrated with Ambari (IOP 4.1.x)

Spectrum Scale HDFS Transparency Connector

- Henceforth known as the “new” connector
- Integrates with HDFS – reuses HDFS client and implements NameNode and DataNode RPCs
- Stateless
- Free download – [link](#)
- Supports Spectrum Scale 4.1.x, 4.1.1.x and 4.2
- Planned for IOP 4.2 (April timeframe)
- Ambari integration supported with IOP 4.2

Old GPFS Hadoop Connector Approach

How can we be sure we're compatible?

Hadoop File System API intended to be open.

```
public abstract class  
org.apache.hadoop.fs.FileSystem
```

Source: *hadoop.apache.org*

“All user code that may potentially use the Hadoop Distributed File System should be written to use a FileSystem object.”

Latest File System APIs are described here:

<https://hadoop.apache.org/docs/current/api/org/apache/hadoop/fs/FileSystem.html>

Old GPFS Hadoop Connector Approach

All based on
org.apache.hadoop.fs.FileSystem API

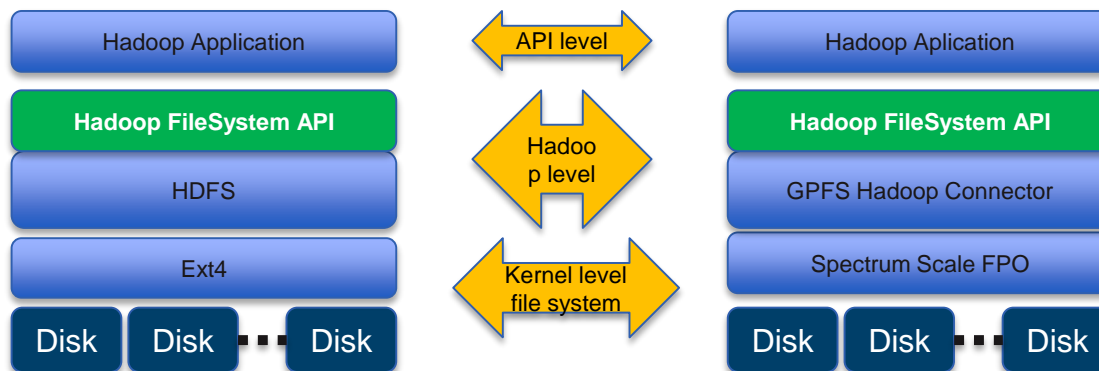
	Optimized for
HDFS	General Hadoop
GlusterFS	file-based scale-out NAS
OrangeFS	high end computing (HEC) systems
SwiftFS	write directly to containers in an OpenStack Swift object store
GridGain	In-Memory Data Fabric
Lustre	
MapR FileSystem	
Quantcast File System	
▪etc...	

*Spectrum
Scale
(GPFS) is no
different*

Source: <https://wiki.apache.org/hadoop/HDFS>

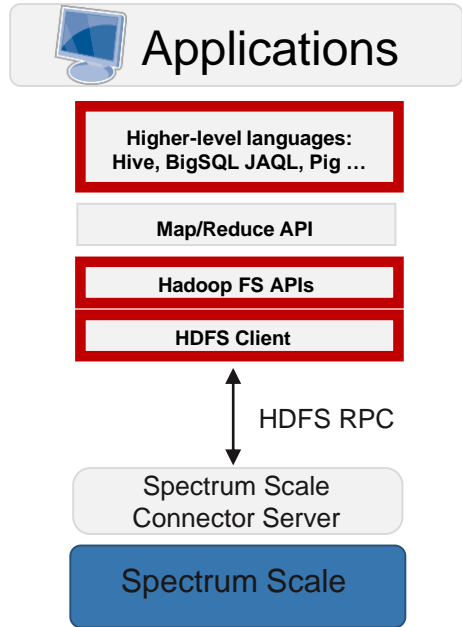
Old GPFS Hadoop Connector Approach

*Applications communicate with Hadoop using FileSystem API.
Therefore, transparency is preserved.*


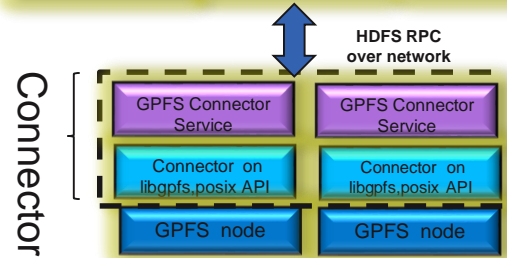



*“All user code that may potentially use the Hadoop Distributed File System should be written to use a **FileSystem** object.”*


New Spectrum Scale HDFS Transparency Design



hdfs://hostnameX:portnumber



Commodity hardware

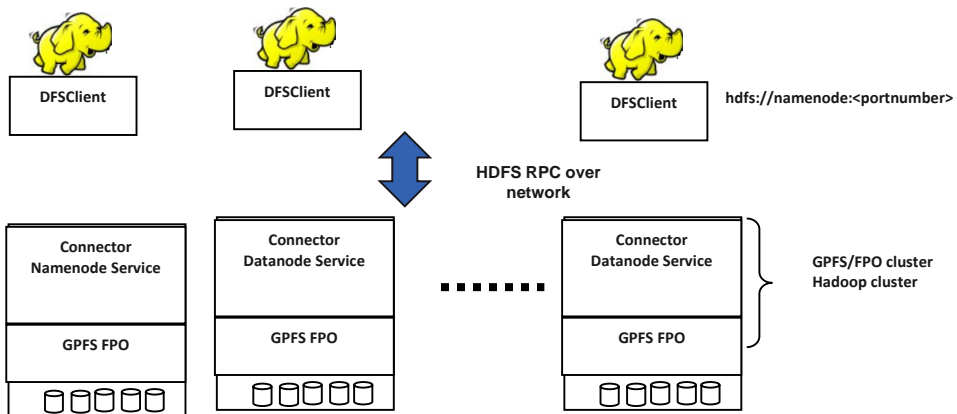


Shared storage

Supported Hadoop versions: 2.7.1

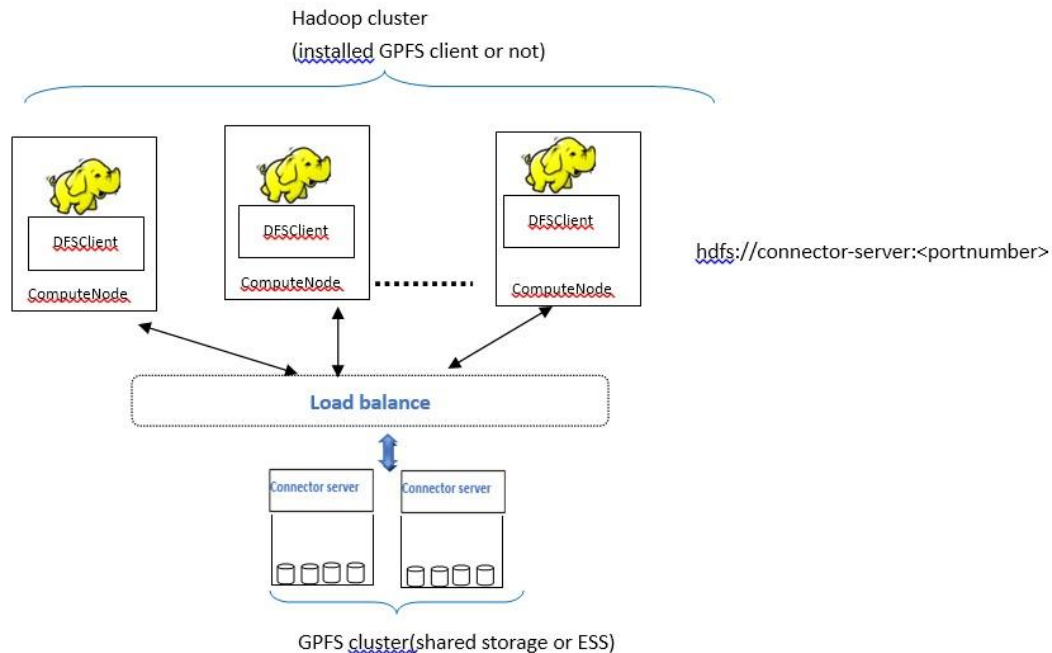
New Spectrum Scale HDFS Transparency Design

- Each node will be installed with connector datanode server
- Only one node will be installed with connector namenode server
- Connector namenode server will be configured with HA, just similar as HDFS
- GA'ed 2015/11/20



New Spectrum Scale HDFS Transparency Design

- Connector servers are installed over limited nodes (ex. GPFS NSD servers)
- GPFS client is not needed over the Hadoop computing nodes
- DNS rotation or CES can be used to load balance for HDFS Client
- GA'ed for 2016/1/22



New Spectrum Scale HDFS Transparency Design

- Key Advantages
 - Support workloads that have hard coded HDFS dependencies
 - Simpler integration for currently compatible workloads & components
 - Leverage HDFS Client cache for better performance
 - No need to install Spectrum Scale clients on all nodes
 - Full Kerberos support for Hadoop ecosystem
- Coming soon
 - BigInsights 4.2 support (additional components)
 - HDFS + Spectrum Scale Federation
 - Federate multiple Spectrum Scale clusters
 - Isolate multiple Hadoop clusters on the same filesystem (restrict to sub-directory)

Current Ambari Integration

- New BigInsights 4.1.SpectrumScale stack
- Inherits from BigInsights 4.1 stack
- Removes HDFS, add Spectrum Scale, change all dependencies
- Can install IOP + Spectrum Scale (either new GPFS filesystem or integrate with existing filesystem)
- Value Add integration
- Basic Spectrum Scale monitoring (AMS)
- Support separate connector control
- Support GPFS and connector upgrades
- Collect GPFS snap
- Change GPFS parameters
- Add new nodes
- Remove nodes
- Provide quick link to Spectrum Scale GUI for full management and monitoring

Select Stack

- Install Options
- Confirm Hosts
- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Review
- Install, Start and Test
- Summary

Please select the service stack that you want to use to install your Hadoop cluster.

Stacks

- BigInsights 4.1.SpectrumScale
- BigInsights 4.1

Advanced Repository Options

Customize the repository Base URLs for downloading the Stack software packages. If your hosts do not have access to the internet, you will have to create a local mirror of the Stack repository that is accessible by all hosts and use those Base URLs here.

Important: When using local mirror repositories, you only need to provide Base URLs for the Operating System you are installing for your Stack. Uncheck all other repositories.

OS	Name	Base URL
<input checked="" type="checkbox"/> redhat6	GPFS-4.1.1	<input type="text" value="http://smn/repos/GPFS/RHEL6/x86_64/4.1.1"/>
	IOP-4.1-mirror	<input type="text" value="http://birepo-build.svl.ibm.com/repos/IOP/RHEL6/x86_64/4.1"/>
	IOP-UTILS-1.1-mirror	<input type="text" value="http://birepo-build.svl.ibm.com/repos/IOP-UTILS/RHEL6/x86_64/1"/>
<input type="checkbox"/> redhat7	GPFS-4.1.1	<input type="text" value="http://c902mnp08/install/repos/IOP-UTILS/rhel/7/ppc64le/1.1/"/>
	IOP-4.1-mirror	<input type="text" value="http://c902mnp08/install/repos/GPFS/rhel/7/ppc64le/4.1.1"/>
	IOP-UTILS-1.1-mirror	<input type="text" value="http://c902mnp08/install/repos/IOP/rhel/7/ppc64le/4.1.x/GA/4.1.0"/>
<input type="checkbox"/> suse11	GPFS-4.1.1	<input type="text" value="http://192.168.9.3/repos/GPFS/SLES/x86_64/4.1.1"/>
	IOP-4.1	<input type="text" value="http://birepo-build.svl.ibm.com/repos/IOP/SLES/x86_64/4.1/"/>
	IOP-UTILS-1.0	<input type="text" value="http://birepo-build.svl.ibm.com/repos/IOP-UTILS/SLES/x86_64/1.1"/>

Skip Repository Base URL validation (Advanced) [?](#)

Current Ambari Integration

Ambari
a 0 ops 1 alert
Dashboard Services Hosts 1 Alerts Admin ☰ admin ▾

- ✔ MapReduce2
- ✔ YARN
- ▲ Hive 1
- ✔ HBase
- ☰ Pig
- ☰ Sqoop
- ✔ Oozie
- ✔ ZooKeeper
- ▲ Flume
- ✔ Ambari Metrics
- ✔ Spectrum Scale
- ✔ Kafka
- ✔ Knox
- ☰ Slider
- ✔ Solr
- ✔ Spark ↻

Actions ▾

Summary
Configs
Quick Links ▾
Service Actions ▾

Summary

[GPFS Master](#) ✔ Started

[GPFS Hadoop Connectors](#) 5/5 GPFS Hadoop Connectors Live

[GPFS Nodes](#) 5/5 GPFS Nodes Live

Metrics

<p>Filesystem Utilization</p> <div style="font-size: 48px; color: green; margin: 10px 0;">0%</div>	<p>Inode Utilization</p> <div style="font-size: 48px; color: green; margin: 10px 0;">0%</div>	<p>Active Quorum Nodes</p> <div style="font-size: 36px; color: green; margin: 10px 0;">3/3</div>	<p>Active NSD Nodes</p> <div style="font-size: 36px; color: green; margin: 10px 0;">5/5</div>	<div style="font-size: 48px; color: #999; margin: 10px 0;">+</div>
--	---	--	---	--

- ▶ Start
- Stop
- ↻ Restart All
- ⊙ Restart GPFS Hadoop Connectors
- ⊙ Restart GPFS Nodes
- 🔍 Run Service Check
- 🔒 Turn On Maintenance Mode
- ⊙ Collect_Snap_Data
- ⊙ Upgrade_SpectrumScale
- ⊙ Upgrade_Connector

Ambari Integration with HDFS Transparency

- Biggest change is that there is no new stack
- Spectrum Scale is added as a new service after full IOP install with HDFS (use dummy directory / mount point for HDFS)
- Spectrum Scale service “integrates” with HDFS
- Will support “un-integrate” capability
 - Flip back and forth between HDFS & GPFS
 - Will not move data back and forth between HDFS & GPFS
- Will simplify future upgrades

References

- Spectrum Scale Knowledge Center
https://www.ibm.com/support/knowledgecenter/STXKQY_4.2.2/com.ibm.spectrum.scale.v4r22.doc/bl1adv_hadoop.htm
- IBM developerWorks
<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20%28GPFS%29/page/IBM%20Spectrum%20Scale%20support%20for%20Hadoop>
- Deployment Guide and other useful information
<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20%28GPFS%29/page/References?section=HDFSTIG>