

Hadoop Integration – Deep Dive

Piyush Chaudhary
Spectrum Scale BD&A Architect

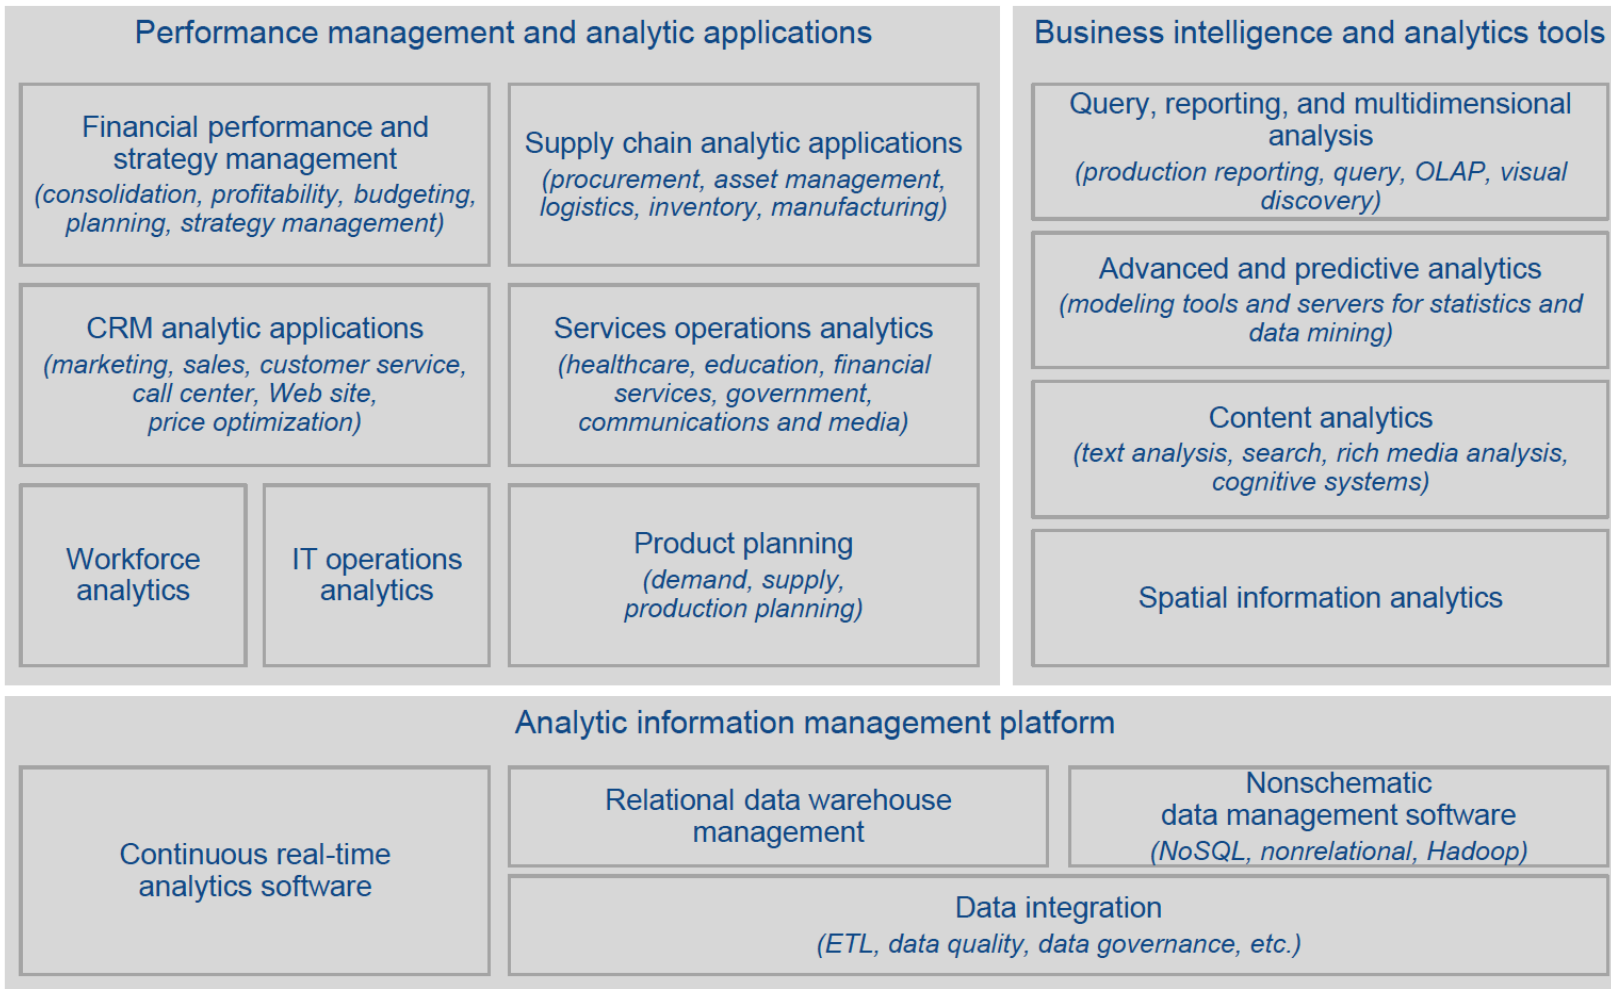


Agenda

- Analytics Market overview
- Spectrum Scale Analytics strategy
- Spectrum Scale Hadoop Integration
 - A tale of two connectors
 - Old GPFS Hadoop connector
 - New Spectrum Scale HDFS Transparency connector

IDC's Business Analytics – Broader Workload Perspective

IDC's Business Analytics Software Taxonomy, 2015

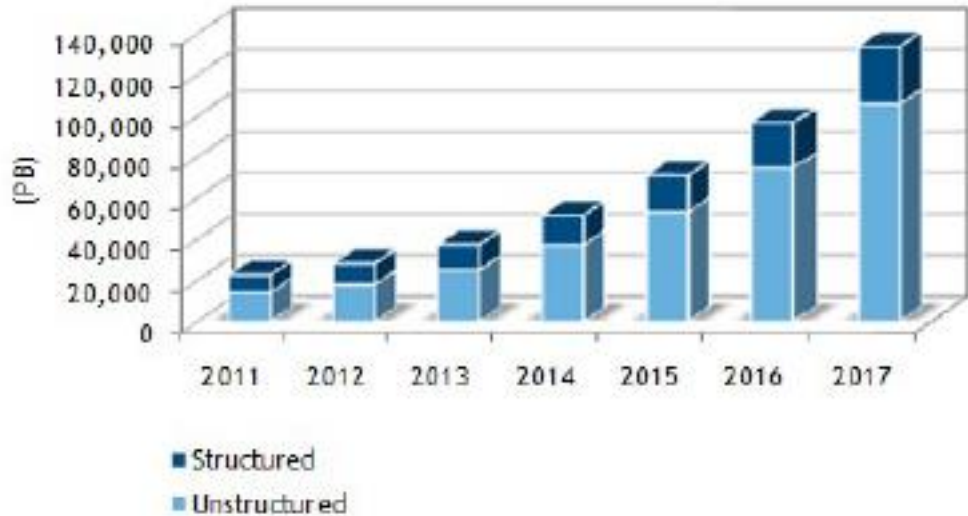


Conclusions

- Analytics market opportunity for Spectrum Scale is broader than just the Hadoop Big Data Market
- IDC definition beginning to account for new workloads
- **NEED** to consider a wide range of Analytics workloads in both Shared nothing and shared storage environments


Structured and Unstructured Data Market

Worldwide Structured and Unstructured Enterprise Storage Systems Capacity Shipped, 2011-2017



Source: IDC, 2014

Worldwide Structured and Unstructured Enterprise Storage Systems Customer Revenue, 2011-2017



Source: IDC, 2014

Worldwide Storage in Big Data Revenue by System Type, 2012-2018 (\$B)

	2012	2013	2014	2015	2016	2017	2018	2013-2018 CAGR (%)
Capacity optimized	0.51	0.86	1.47	2.07	2.86	3.69	4.30	38.0
Performance optimized	0.49	0.61	0.72	0.84	0.94	1.06	1.10	12.5
I/O Intensive	0.04	0.21	0.13	0.21	0.33	0.45	0.49	18.5
Total	1.04	1.68	2.32	3.12	4.13	5.20	5.89	28.5

Note: See Table 6 for top 3 assumptions and Table 7 for key forecast assumptions.

Source: IDC, 2014

Conclusion

- Spectrum Scale addresses the sweet spot of the growth segment for Enterprise Storage Systems.

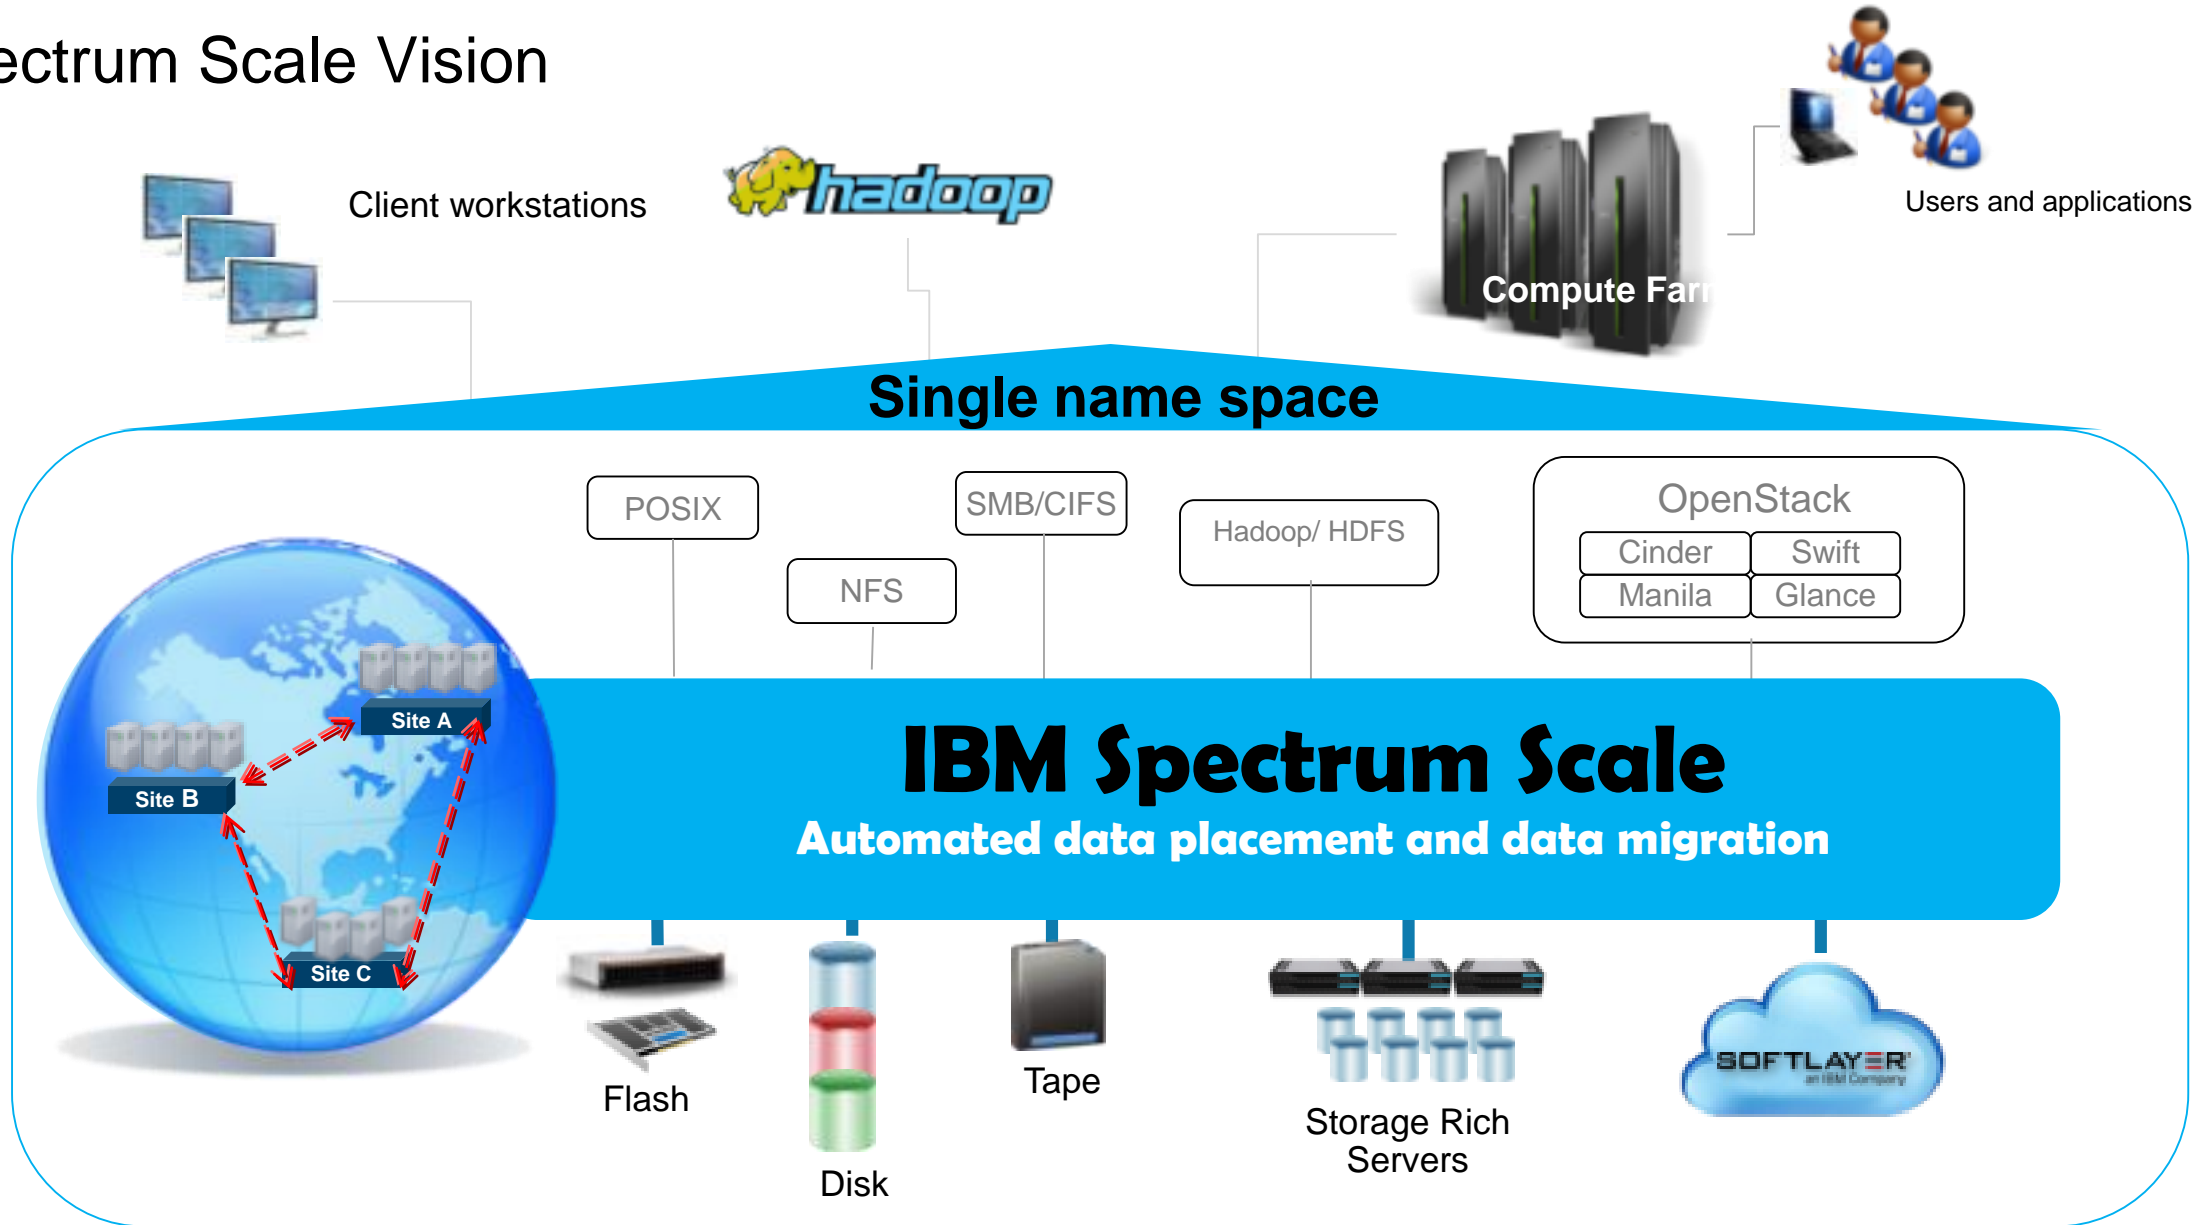
The background of the slide is a photograph of an underwater scene. Sunlight rays are visible, filtering down from the surface, creating a shimmering effect on the water. The water is a deep blue color, and the overall atmosphere is serene and vast.

People and devices are creating oceans of data

Data scientists and LoBs want to fish this ocean for insights

The ocean of data must be efficiently stored, managed, protected, and exploited by the right applications at the right time

Spectrum Scale Vision



Multi-structured Data Mashups provide the Greatest Enterprise Value

Systems of Record

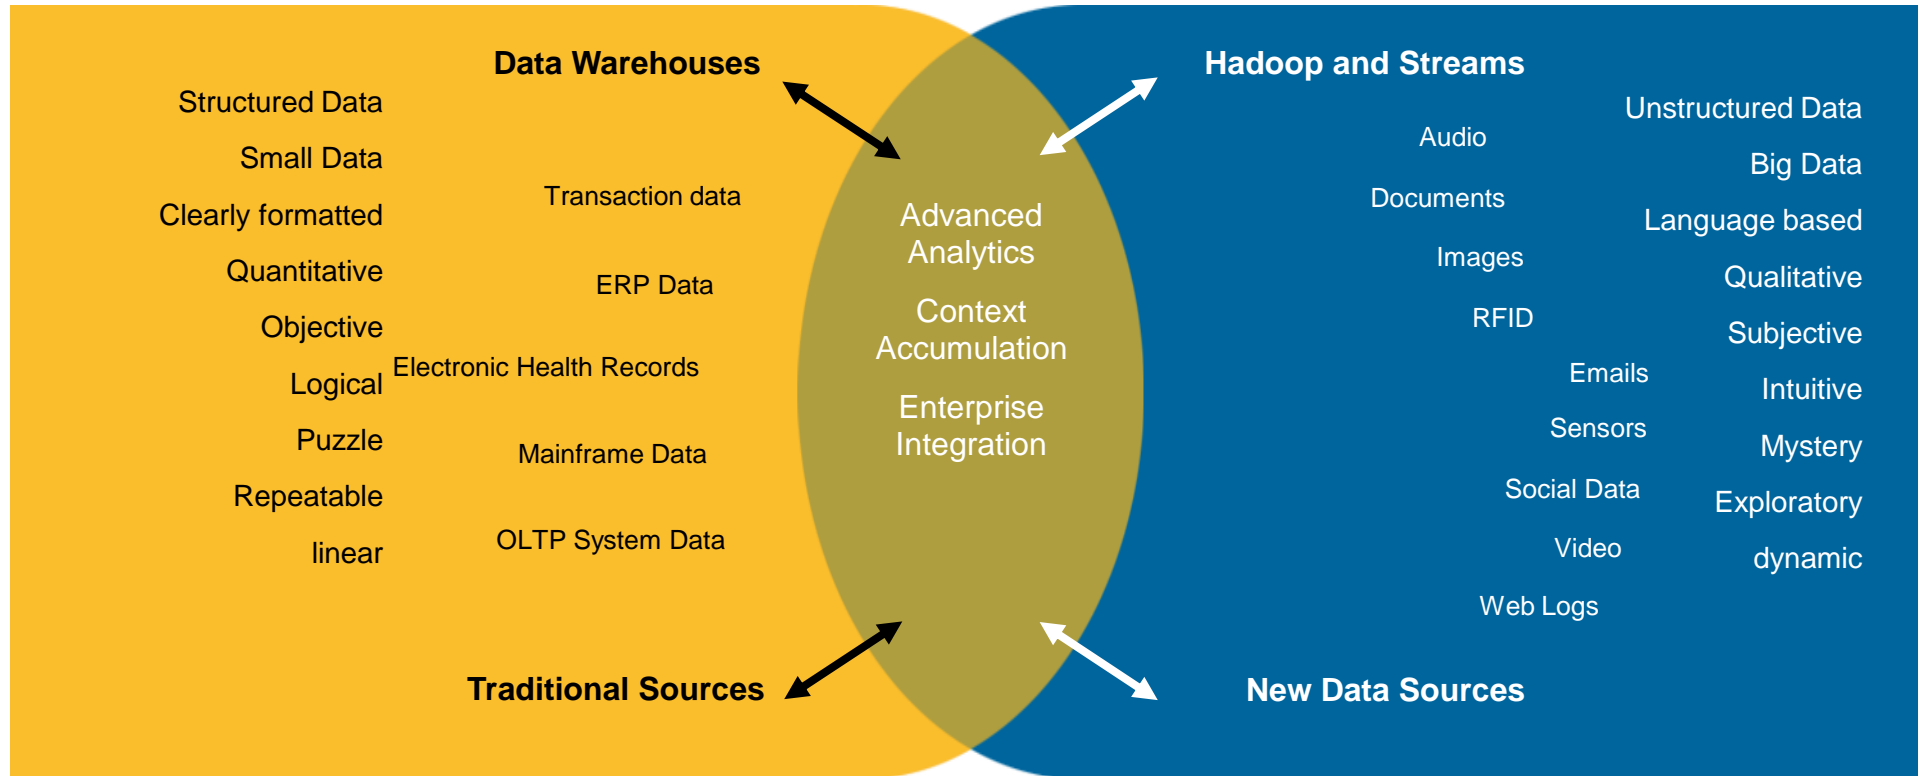
Structured data from operational systems
20% of all data generated

Systems of Insight

Diverse data types that combine structured and unstructured data for business insight

Systems of Engagement

Data that "connects" companies with their customers, partners and employees
80% of all data generated



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 4.1.1 and 4.2
- Currently supported with IOP 4.0 and 4.1
- Integrated with Ambari (IOP 4.1)
- Also supported with Open Source Apache Hadoop

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 4.1.1 and 4.2
- Planned support for IOP 4.1 and 4.2 (6 weeks after GA)
- Ambari integration being developed for both IOP 4.1 and 4.2 (coming soon)

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`

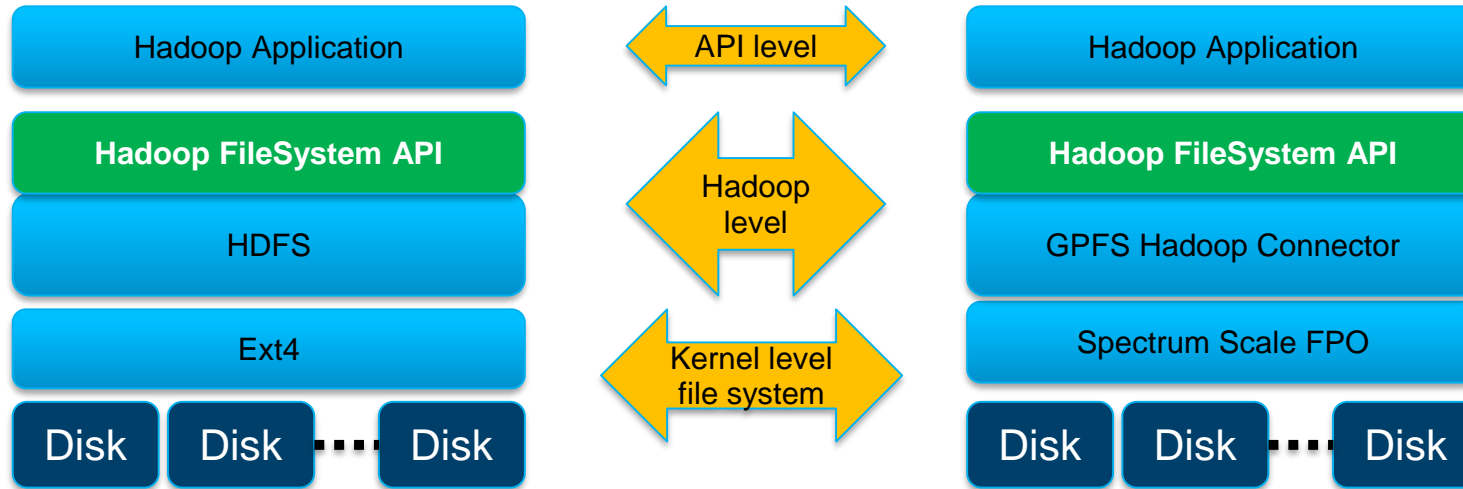
	Optimized for
HDFS	General Hadoop
GlusterFS	file-based scale-out NAS
OrangeFS	high end computing (HEC) systems
SwiftFS	write directly to containers in an OpenStack Swift object store
GridGain	In-Memory Data Fabric
Lustre	
MapR FileSystem	
Quantcast File System	
▪etc...	

*Spectrum
Scale
(GPFS) is no
different*

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

Old GPFS Hadoop Connector Approach

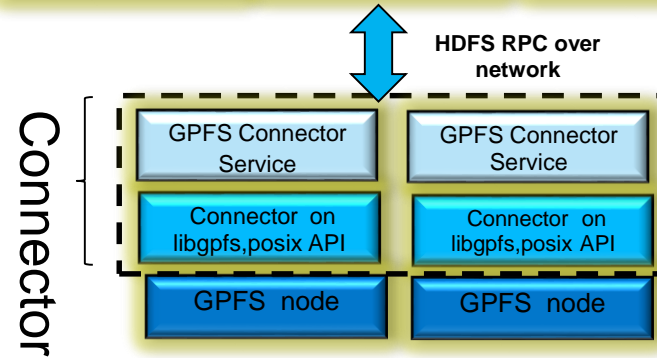

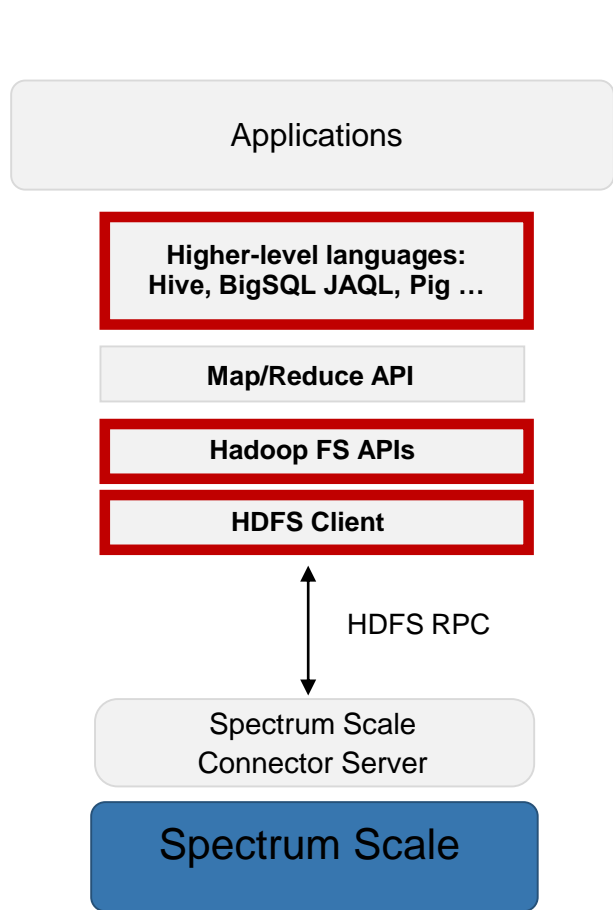
*Applications communicate with Hadoop using FileSystem API.
Therefore, transparency is preserved.*




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

Source: hadoop.apache.org

New Spectrum Scale HDFS Transparency Design



Commodity hardware

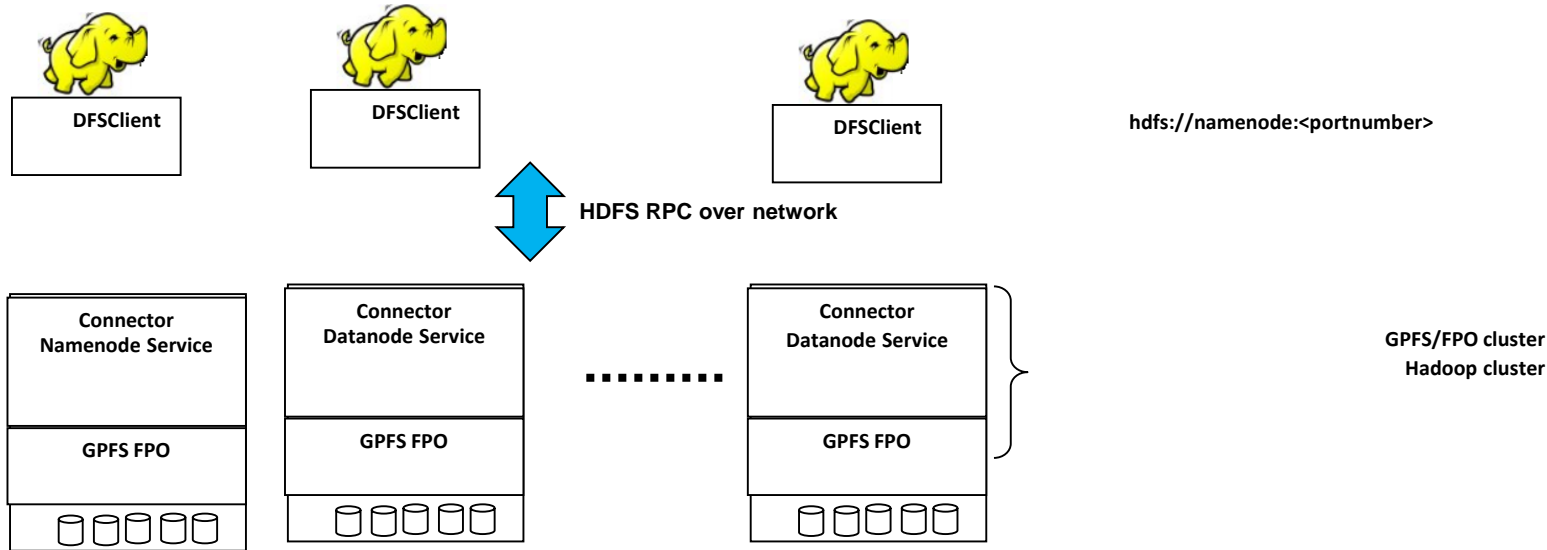


Shared storage

Supported Hadoop versions: 2.7.1

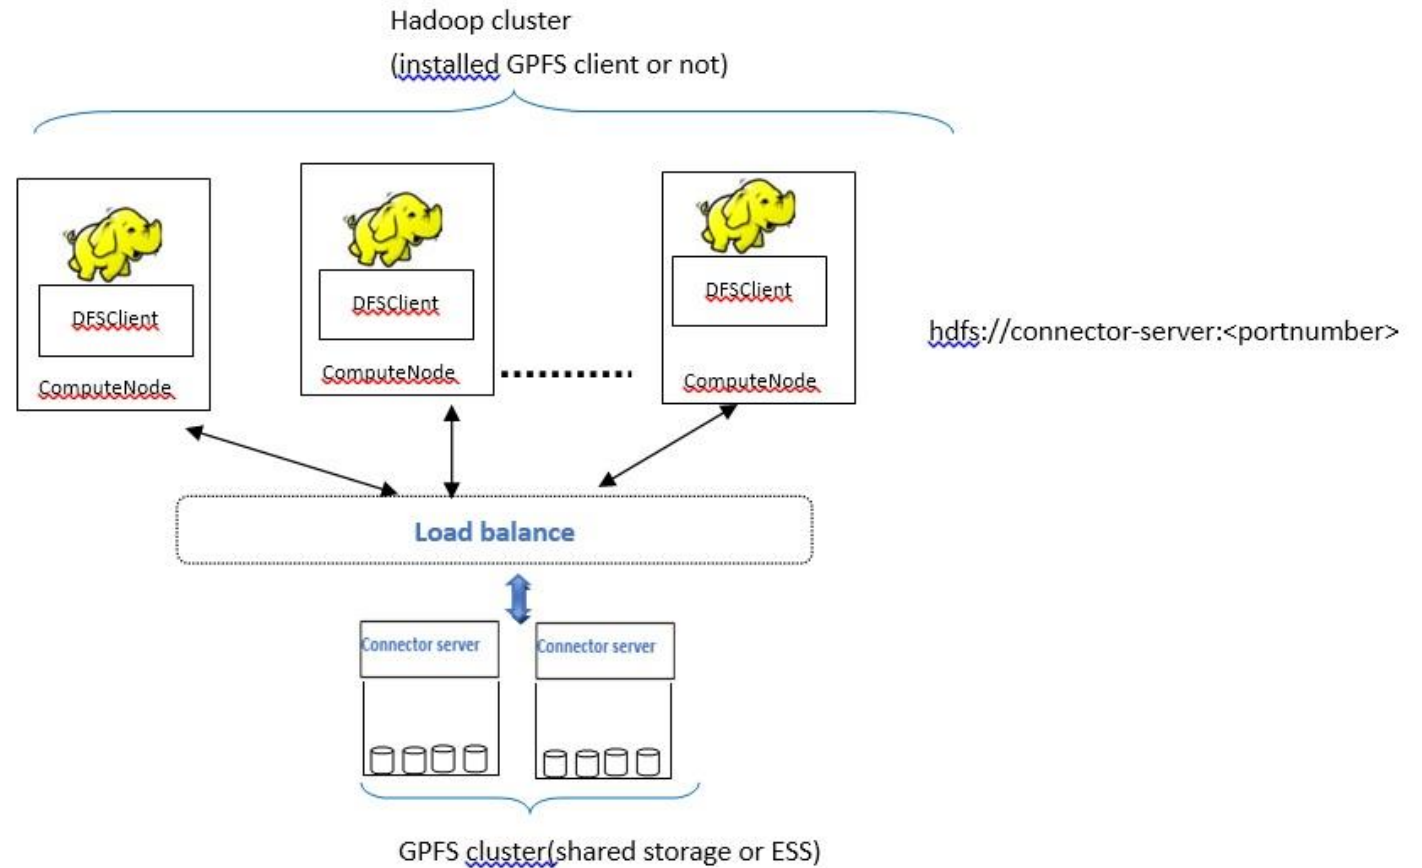
New Spectrum Scale HDFS Transparency Design

- Each node is installed with connector datanode server
- Only one node is installed with connector namenode server
- Connector namenode server can be configured with HA similar to HDFS
- First released in November 2015



New Spectrum Scale HDFS Transparency Design

- Shared storage support
- Connector servers are installed over limited nodes (ex. GPFS NSD servers)
- GPFS client is not needed on the Hadoop computing nodes
- First released in January 2016



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
- Recently released
 - HDFS + Spectrum Scale Federation
 - Federate multiple Spectrum Scale clusters
 - Isolate multiple Hadoop clusters on the same filesystem (restrict to sub-directory)
- Coming Soon
 - BigInsights 4.2 support (additional components)

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_64/1.1"/>
<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.C"/>
<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) ?

[← Back](#) [Next →](#)

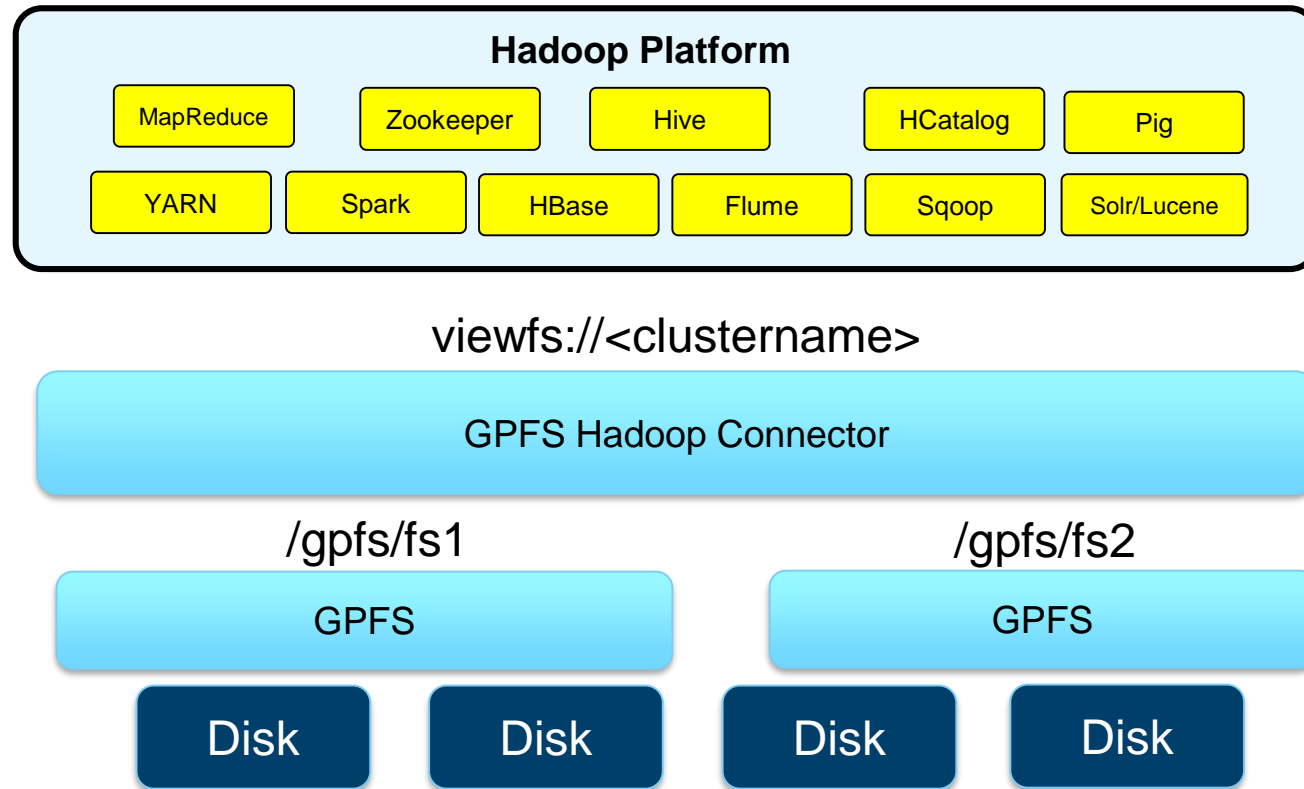
Current Ambari Integration

The screenshot displays the Ambari web interface for the Spectrum Scale service. The top navigation bar shows 'Ambari' with '0 ops' and '1 alert' indicators. The main navigation includes 'Dashboard', 'Services', 'Hosts 1', 'Alerts', and 'Admin'. The user is logged in as 'admin'. The left sidebar lists various services, with 'Spectrum Scale' highlighted. The main content area shows the 'Summary' and 'Metrics' sections. A 'Service Actions' dropdown menu is open, listing actions such as Start, Stop, Restart All, Restart GPFS Hadoop Connectors, Restart GPFS Nodes, Run Service Check, Turn On Maintenance Mode, Collect_Snap_Data, Upgrade_SpectrumScale, and Upgrade_Connector. The 'Metrics' section shows 'Filesystem Utilization' at 0%, 'Inode Utilization' at 0%, 'Active Quorum Nodes' at 3/3, and 'Active NSD Nodes' at 5/5.

Ambari Integration with HDFS Transparency

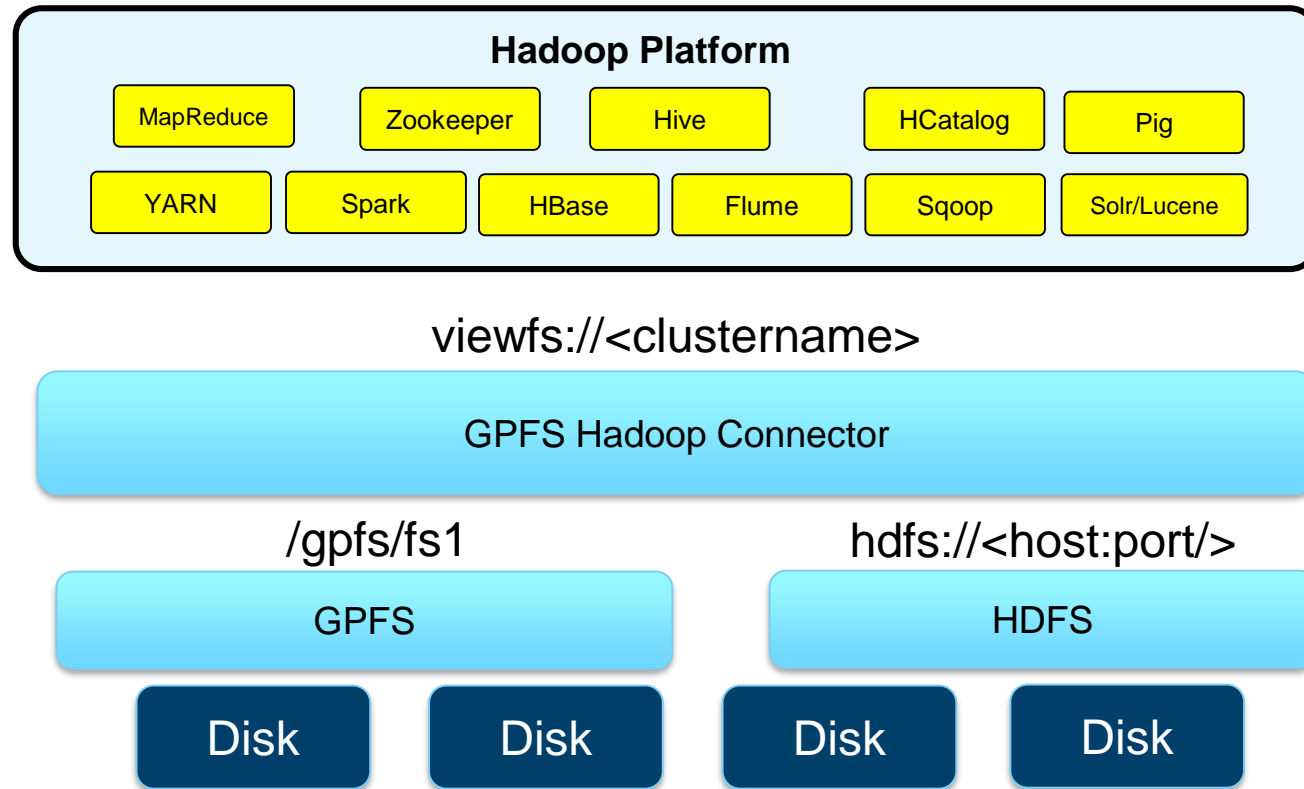
- Biggest change is that there is no new stack
- Spectrum Scale is added as a new service after full IOP install with HDFS (use dummy directory / mount point for HDFS)
- Spectrum Scale service “integrates” with HDFS
- Will support “un-integrate” capability
 - Flip back and forth between HDFS & GPFS
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

Spectrum Scale Federation – Use Case 1



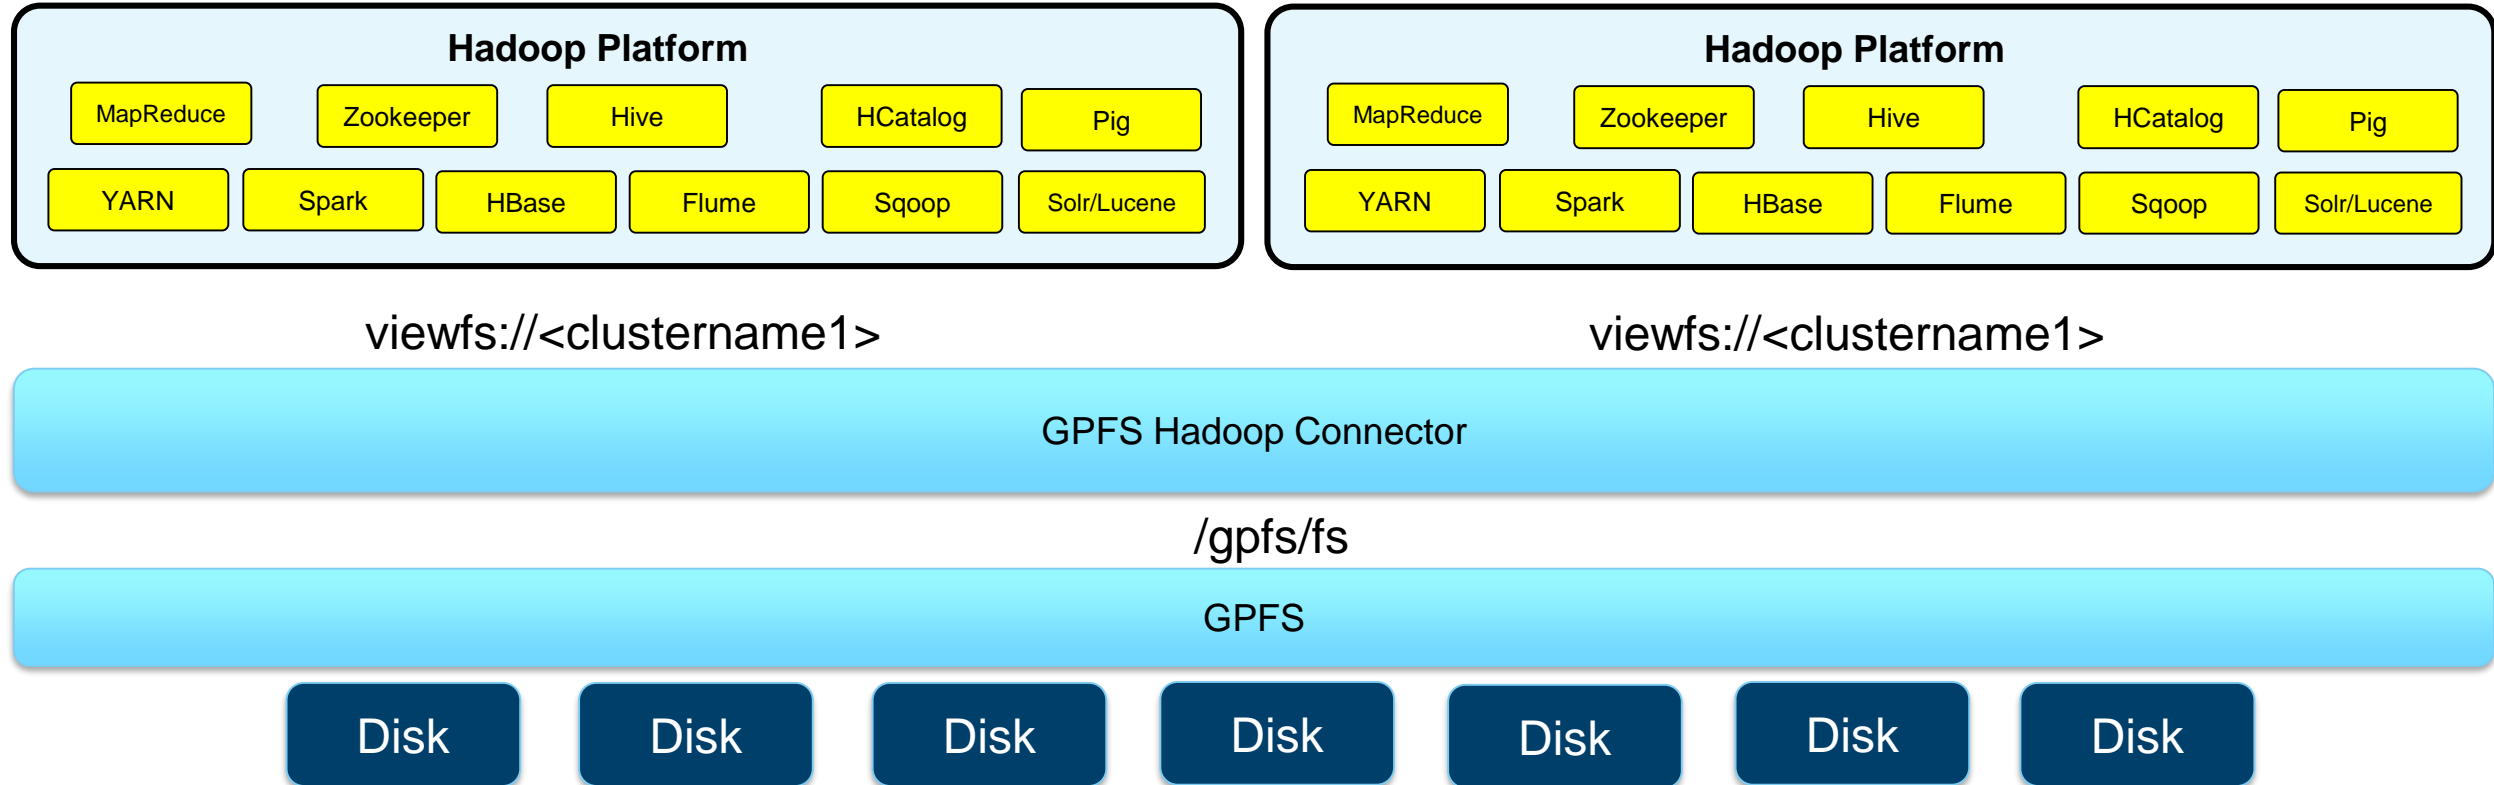
- Make 2+ GPFS file systems appear as one uniform file system for application
- Allow customer to leverage existing GPFS file system and new GPFS file system for data analysis

Spectrum Scale Federation – Use Case 2



- Allow customers to run applications over GPFS & HDFS seamlessly
- Doesn't require customers to move data from HDFS to GPFS

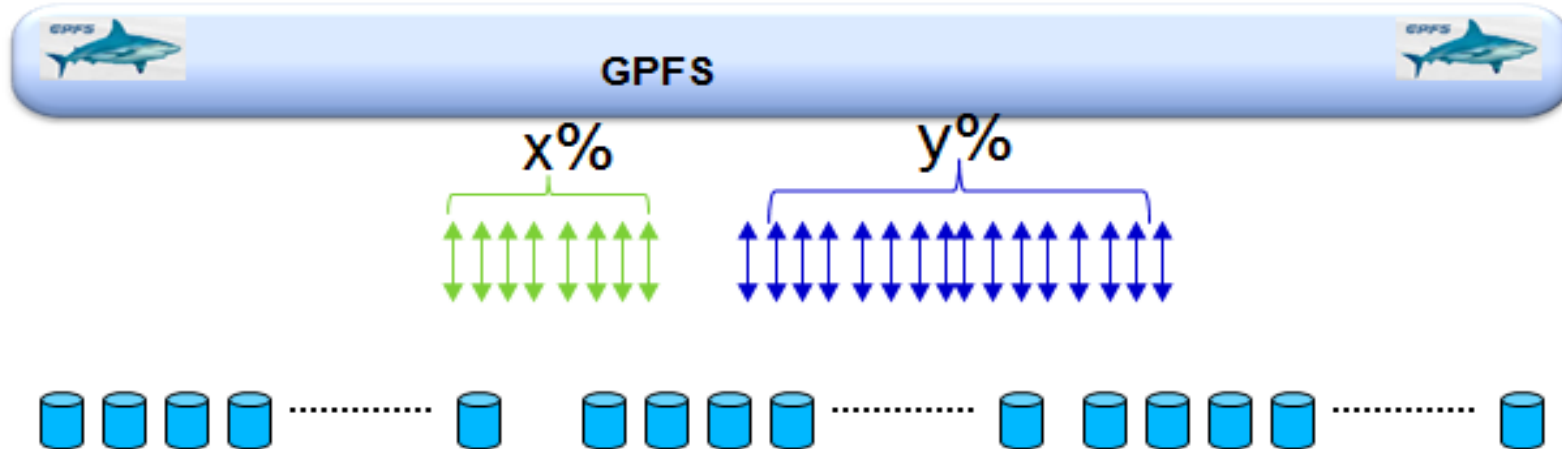
Spectrum Scale – Workload Isolation



Support multiple Hadoop clusters over the same GPFS filesystem for different applications, groups or LOBs

Spectrum Scale FPO – QOS Support

- Node failure is often-seen events for commodity hardware
- Restripe for data protection is painful for customer with large data
- QOS for FPO - Reduce the restripe impact in FPO because of node failure



- Maintenance IO+normal IO
- Cap the io bandwidth %
- Dynamically change

Revamped Spectrum Scale Wiki Links

[Spectrum Scale wiki – Analytics](#)

[Spectrum Scale wiki - FPO](#)

[Spectrum Scale wiki – Hadoop](#)

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