


Consolidate high performance file and S3 object storage

Give your applications the choice of
interfaces to the data it needs

Highlights

- Innovation with a single copy of data with multiple interfaces
- Included at no charge with IBM Spectrum Scale and ESS
- Scale out S3 performance as requirements grow
- Start with TB of data and scale capacity of S3 data to YB
- Access same data with either a file or S3 interface
- Integrated lifecycle management for object and file data
- Easily integrate S3 operator into a Red Hat OpenShift cluster
- Access S3 data with TB /s high performance interface



Driving requirements for storage and data access

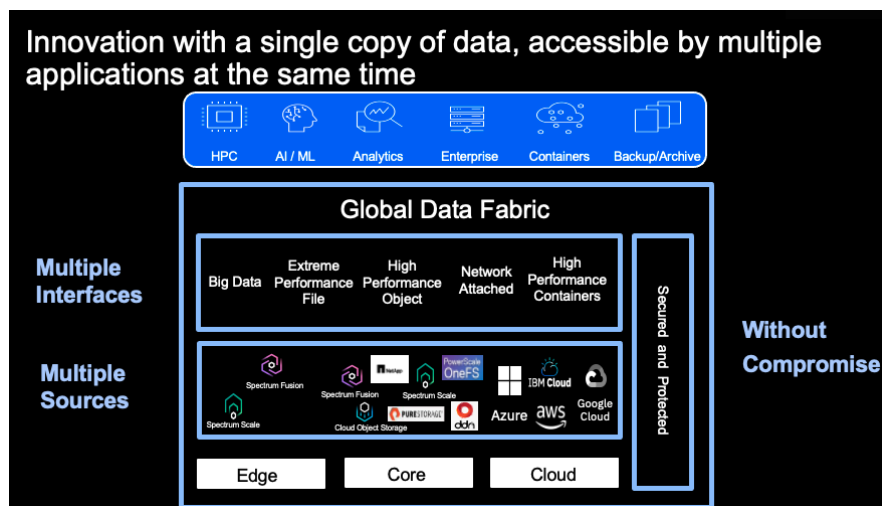
Application modernization and AI is driving requirements for storage. We see 3 main areas that are driving these requirements, agility, resilience and security and finally flexibility.

An agile environment is critical to rapidly respond to application requirements and to support DevOps practices.

An agile environment also brings application context and awareness along with data management between clusters and allows for a free-flowing application environment with free-flowing data.

Resiliency and security maintains business continuity and helps to combat cyber threats. Enterprise infrastructure must also be ready for disasters and provide a high available environment so users can access the data when they need them without losing access. The data must also be kept safe with encryption and backup to ensure there is no data loss.

Flexibility is also important as integration of existing infrastructure and hardware is constant with cloud deployments. Software provides the flexibility for different hardware infrastructure solutions located on the edge, the core or the data center and provides the consistent interface to place applications where they are best deployed.



A Global Data Fabric

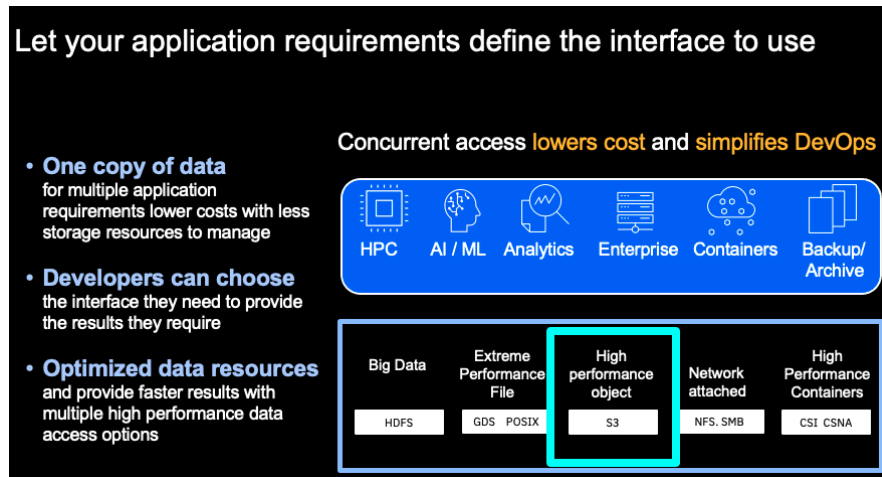
IBM Storage is providing more than storage solutions and capabilities but a data and AI strategy built on innovation and a Global Data Fabric. We provide the innovation with a single copy of data that is accessible by any application at the same time. We provide multiple interfaces from multiple sources that is all secured and protected from edge to core to cloud.

Our single source of truth provides value to the organization beyond the basic storage capabilities. The Global Data Fabric:

*eliminates duplicate data, simplify workflows, and lower cost as global access means no need to copy data from external sources

*improves application agility as global data access enables applications to run anywhere

*speeds AI results and accuracy because applications can be connected to more data with a single copy whether on the edge, core, or cloud




Multiple interfaces are required but difficult

IBM makes it easy by letting your application requirements define the interface to use and not restricted but the storage that IT defines. Our high performance S3 interface lowers costs for ITOps and simplifies DevOps. We provide one copy of data for multiple application requirements which lower costs with less storage resources to manage. Developers can choose the interface they need to provide the results they require. Our S3 interface also optimizes data resources and provide faster results with multiple high performance data access options.

We let your application requirements define the interface to use:

- one copy of data for multiple application requirements lower costs with less storage resources to manage
- developers can choose the interface they need to provide the results they require
- optimized data resources and provide faster results with multiple high performance data access options



IBM high performance S3 interface

The IBM Spectrum Scale high performance S3 interface is included with IBM Spectrum Scale and ESS at no additional charge. The software interface is a simple Red Hat OpenShift operator download and also provides a high performance OpenShift interface for non S3 workloads. S3 workloads can be any S3 application located anywhere that has access to the high performance interface. We challenge you to find a faster S3 interface. Our S3 interface provides:

- fast AI results for S3 cloud native applications
- scalable solution for ingesting high performance S3 object data from remote locations
- scale performance and capacity as needed
- applications can now optimize with the interface they need to access all the data they require (example: ingest S3 and access file)

A Global Data Fabric with an S3 Interface

A fortune 100 customer used the global data fabric with a high performance S3 interface and was able to access data anywhere for better AI collaboration. Data is accessed with S3 API from one source with multiple concurrent parallel paths. Multiple applications ingesting data with a file interface. Accessing data without movement meant customer did not need 2 copies of data and meant faster results without additional resources.

<h3>Multi-site industrial example</h3> <p>Global data fabric helps store once and access anywhere for better AI collaboration</p>	<h4>Faster results</h4> <p>Data is accessed directly from one source with multiple concurrent parallel paths for multiple concurrent solutions.</p>	<h4>Lower cost</h4> <p>Accessing data without movement meant customer did not need 2 copies of data.</p>
<div data-bbox="359 495 807 775"><h3>Multiple Applications</h3><p>Big Data, HPC, ADAS, CAE, Kubernetes applications</p><h3>200PB+</h3><p>shared data across 2 remote design centers</p><h3>High Performance S3</h3><p>required but data was already ingesting to a file source</p></div> <div data-bbox="807 548 1244 775"><p>The diagram illustrates a 'Global data fabric' as a central green oval. Arrows point to it from 'Ingest File Data' (top left), 'Tape and cloud backup' (bottom left), and 'Big Data' (top right). An arrow points from the fabric to 'Access S3' (top right), which is linked to 'Site C access' (bottom right). The IBM Spectrum Scale and ESS logo is at the bottom.</p></div>		

Customer example

Why IBM?

IBM Storage for Data and AI provides a high-performance global data platform that solves modern data challenges with a single source of truth and a global data fabric. IBM innovation and industry-leading software-defined storage management and data services provides workload agility, faster insights for AI and business results. Our storage connects multiple applications that each may contain different interfaces with multiple data sources in a single managed solution. Finally, you modernize without compromise as efficiency, protection and security are all designed to meet global enterprise requirements and current investments with your infrastructure.

Next steps

→ [IBM Spectrum Scale information](#)

For more information

For more information about our Global Data Fabric products

[IBM Spectrum Scale](#)

[IBM Elastic Storage System](#)

[IBM Cloud Object Storage](#)

[IBM Spectrum Discover](#)

[IBM Spectrum Fusion](#)

© Copyright IBM Corporation 2022.

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at <https://www.ibm.com/legal/us/en/copytrade.shtml>, and select third party trademarks that might be referenced in this document is available at https://www.ibm.com/legal/us/en/copytrade.shtml#section_4.

This document contains information pertaining to the following IBM products which are trademarks and/or registered trademarks of IBM Corporation:
IBM Spectrum Scale



All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.