

A Global Data Fabric with a single source of truth


Highlights

- Eliminate duplicate data
- Improve application agility with Kubernetes integration
- Speed AI results and accuracy
- Consolidate global data on a single storage system
- Developers can choose the interface to access any data
- Optimized data resources for performance and costs
- Connect multiple data sources into a single data lake
- Create an open ecosystem and avoid vendor lock-in
- Quickly expand storage as needed
- High performance options from edge to core to cloud

Solving AI data challenges with innovation and simplicity without compromise

Our story begins with with a race. The race to become agile and for organizations to adjust the the constantly changing business requirements that they face. CIOs are constantly looking for ways to improve IT and digital assets to help achieve business goals. There are many reasons that CIOs pursue infrastructure modernization. The top reasons are agility and helping the organization speed its time to market with new applications or by pursuing new opportunities. Agile organizations can quickly redirect their people and priorities toward value-creating opportunities. Quality of Service or reliability and cost efficiencies follow closely as the 2nd and 3rd reasons for pursuing infrastructure modernization.

Organizations need to unlock the value from their data

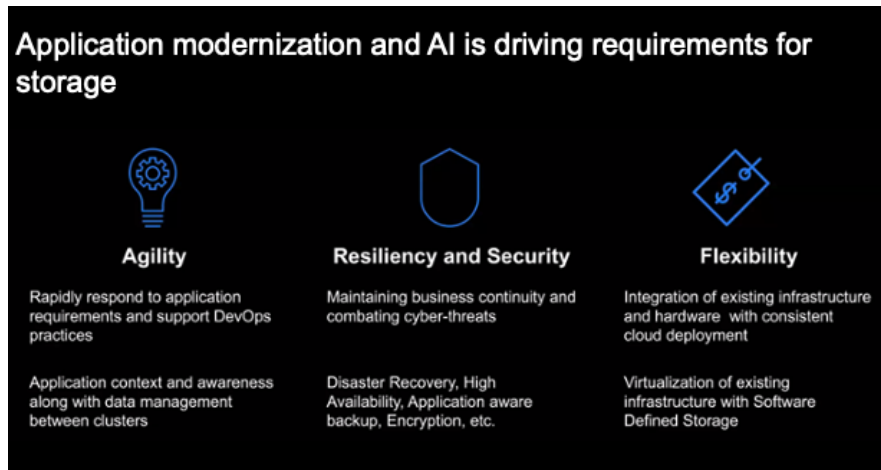


“ We fundamentally believe that core to the competitiveness of every company going forward will be their ability to use AI to unlock real-time value from their data wherever the data resides.¹ ”

¹ IBM (IBM) Q2 2021 Earnings Call Transcript

Unlock the value of your data

Organizations need to unlock the value from their data. As Arvind Krishna, CEO of IBM has so well stated on the Q2 2021 Earnings call “We fundamentally believe that core to the competitiveness of every company going forward will be their ability to use AI to unlock real-time value from their data wherever the data resides.”



Application modernization is driving storage requirements

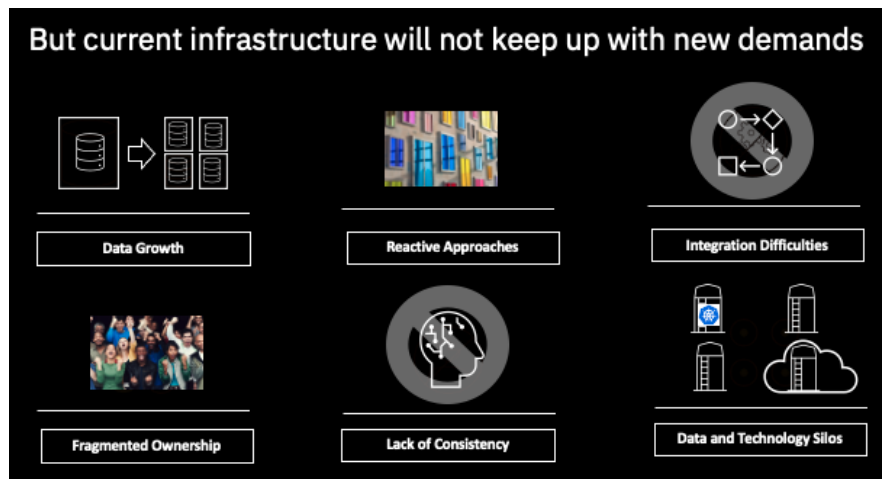
Application modernization and AI is driving requirements for storage. We see 3 main areas that are driving these requirements, agility, resilience and security and finally flexibility.

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.



Current Infrastructure can't keep up with demands

Many challenges prevent organizations from achieving the levels of business and operational agility they need to foster innovation and remain competitive in today's markets. You hear them mentioned all the time – growth, speed, variety, and veracity of data. But there is so much more that is creating a sense of urgency for organizations just like yours - it's the underlying services that your IT infrastructure delivers, consistently, to enable that agility and enhance workforce productivity without compromising business requirements. To highlight a few that are affecting the way our customers are doing business today, consider the following:

1. The growth of data is constant – leaving organization's continually facing the challenges of aggregating, managing, governing, and creating new value from their data
2. Organizations don't often realize there is a problem with their data in advance of a situation. Reacting creates in-the-moment data management solutions that fail to scale as the business grows, applications change, and new data formats are adopted
3. A lack of processes, data management systems, and inadequate data strategies contribute toward data quality issues and untrustworthy data sets.
4. Lack of data ownership is a key shortfall for most organizations. This fragmentation is often across different departments and stakeholders, versus the business as a whole.
5. Many organizations are unable to invoke enough support to improve their data culture. This is often driven by a lack of knowledge or skill set on the proper techniques to Manage data effectively.
6. Lastly, reactive approaches, incomplete processes, fragmented data ownership, and insufficient skills are catalysts that foster single-purpose solution adoption that ultimately traps data in technology silos that ultimately increase compliance and productivity risks for the business

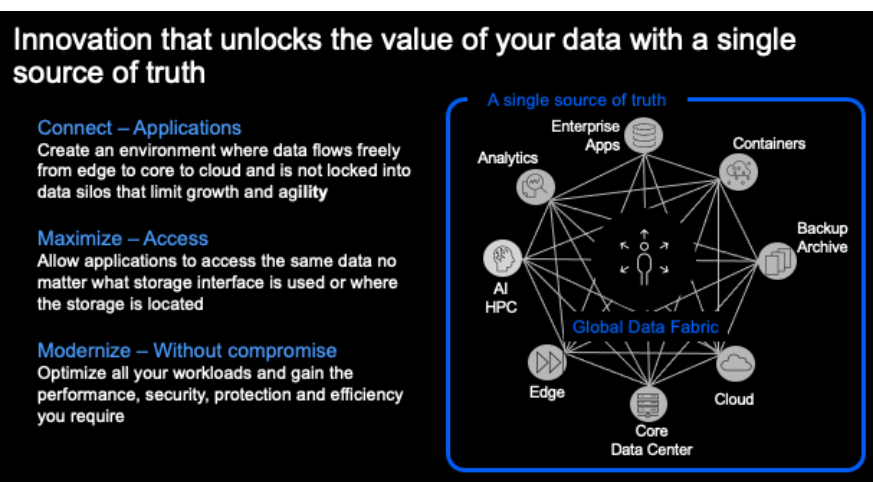
Time to RAISE the bar for your data requirements

We think its time to raise the bar with what you expect for your AI storage. Our storage solutions enable application agility, brings faster insights for AI, and protects data while optimizing costs.

Are you ready to...

- connect your applications to the data they need with the performance they require?
- maximize your resources and provides faster results with an open ecosystem and investment protection?
- optimize your data with the efficiency to help lower costs without compromising the security and protection required?

A Global Data Fabric



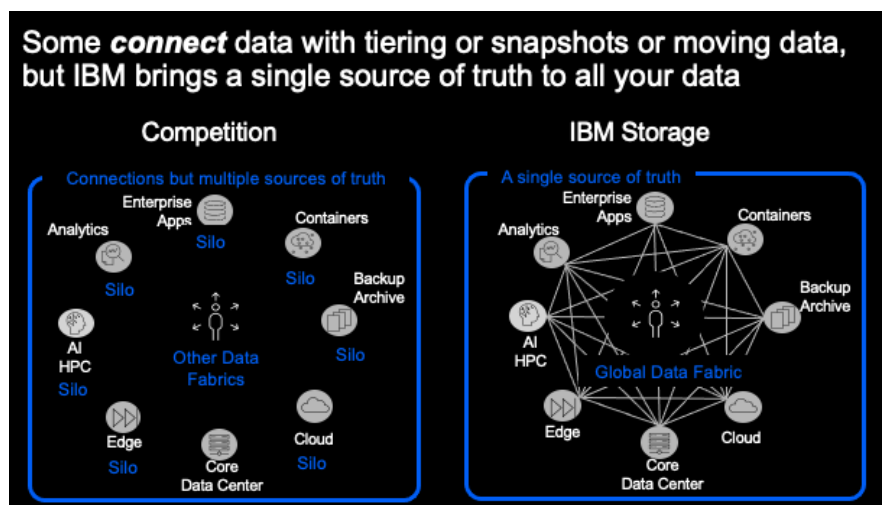
Innovation that unlocks the value of your data

IBM has taken the challenges and created innovation that unlocks the value of your data with a single source of truth. We have created a Global Data Fabric optimized for the next generation of enterprise solutions and AI workloads. What is a global data fabric? A global data fabric is an architecture and a data strategy that brings fluid data and a single source of truth with multiple file and object access choices from multiple data sources to users and applications accessible at the edge, the core, and the cloud.

How can the global data fabric help bring business results : The

The global data fabric breaks down silos and current data limitations to bring faster AI insights and modernizes data infrastructure for optimized business results. It is globally accessible with active file management (AFM), secured with audit logging and immutable data, protected with comprehensive resiliency options, and optimized with lifecycle management for costs and data caching for performance.

IBM provides a global data fabric that brings fluid data with multiple file and object access choices from multiple data sources accessible at the edge, the core, and the cloud. Only IBM provides a global data fabric that is globally accessible with active file management (AFM), secured with audit logging and immutable data, protected with comprehensive resiliency options, and optimized with lifecycle management for costs and data caching for performance.

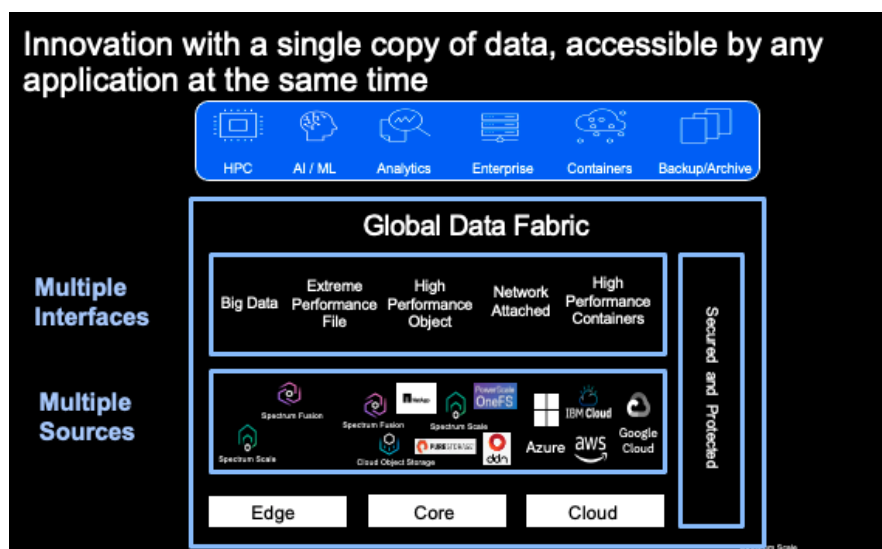


Other competitive data fabrics

We know there are others who are building a data fabric but what they are not creating is a single source of truth or a transparent global data solution.

IBM eliminates the complexity of managing data for ITOps and accessing data for DevOps and global applications.

Some connect a namespace with tiering or snapshots or moving data, but IBM brings a single source of truth to all your data.



IBM 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 multipole interfaces from multiple sources that is all secured and protected from edge to core to cloud.

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

Concurrent access lowers cost and simplifies DevOps

Big Data	Extreme Performance File	High performance object	Network attached	High Performance Containers
HDFS	GFS, POSIX	S3	NFS, SMB	CSI, CSNA

Application requirements determine the interface

We let your DevOps teams, or your applications define the interface they need to access the data they require.

- 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

Maximize resources and connect to what your applications need

- **An open ecosystem** of storage options leveraging multi-vendor and multi-cloud resources
- **Connecting data sources** from edge to core to cloud brings more data to applications for faster results
- **Quickly expand data** from resources you choose with performance you require

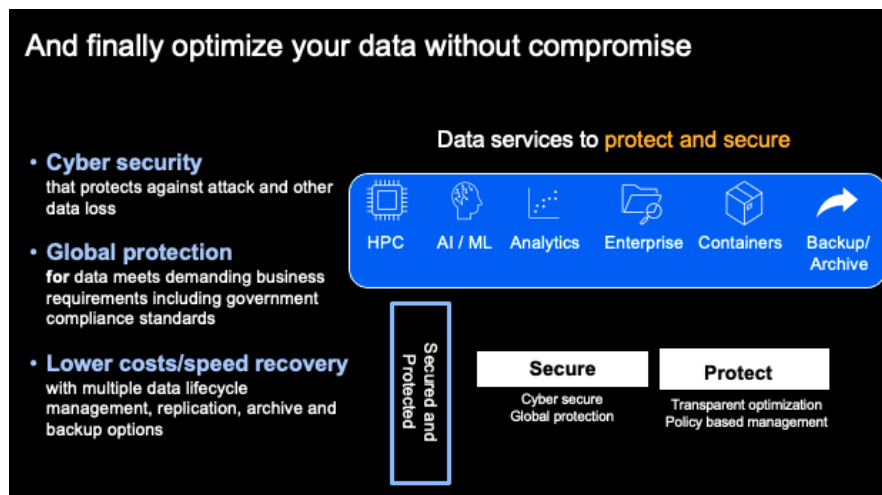
An open ecosystem brings investment protection

Edge	Core	Cloud
------	------	-------

Maximize resources and connect to data your application needs

We maximize resources and connect to what your applications need

We provide an open ecosystem of storage options leveraging multi-vendor and multi-cloud resources. We connect data sources from edge to core to cloud brings more data to applications for faster results. We also allow you to quickly expand data from resources you choose with performance you require.

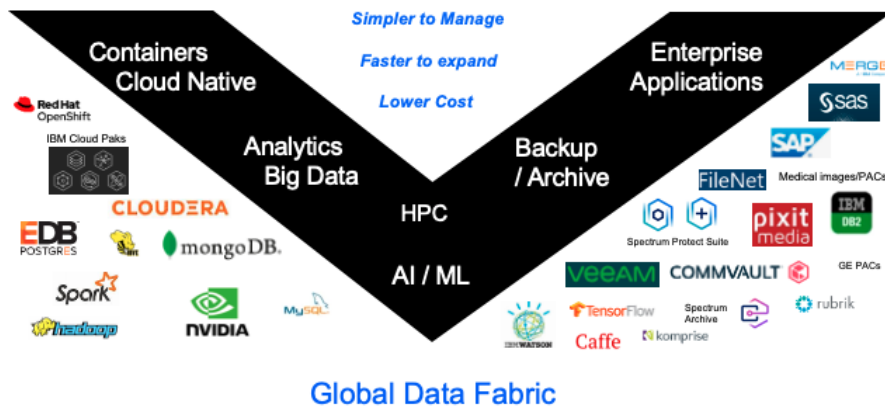


Optimize your data without compromise

And finally, we optimize your data without compromise with the security and protection required for you're your critical data assets.

The data fabric is integrated with cyber security capabilities including encryption, immutable data and secure audit trails that protects against attack and other data loss. The Global Data Fabric provides global protection for demanding business requirements including government compliance standards for financial services. We also lower costs/speed recovery with multiple data lifecycle management, replication, archive and backup options.

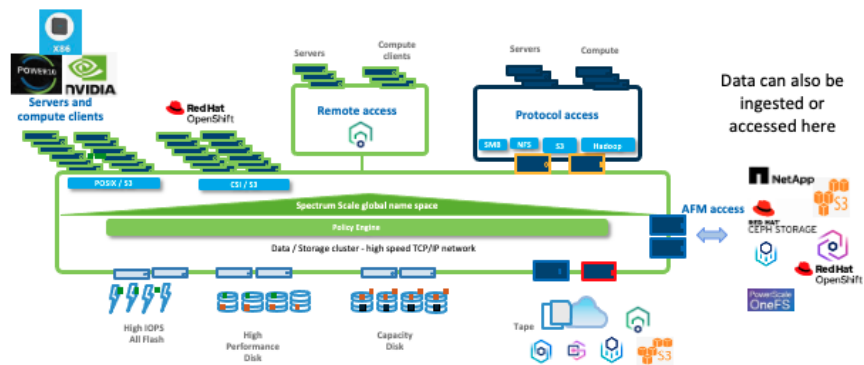
We can consolidate everything on a single global data fabric



We consolidate everything on a global data fabric

Many of the key capabilities and usage cases are in play here. Our Global Data Fabric has production enterprise class customers doing many of the solutions shown. There is much more that is possible, and we optimize your data for your global applications and workloads.

You optimize your workloads and let IBM optimize your data



Let IBM optimize your data

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

→ [Learn more about IBM Data and AI solutions](#)

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 2021.

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.