

IA optimized for AI and hybrid cloud

Data drives your business and you generate a lot of it. Challenges related to data may be slowing you down and putting your data — and your business — at risk. You need an information architecture (IA) optimized for faster business results.

Are these challenges slowing you down?

Data locked in silos, making it difficult to access and use AI across the hybrid cloud.

Costly legacy infrastructure preventing you from accessing data when and where you need it.

Increasing risks of outside attacks, inside vulnerabilities and system failures.



IBM Storage for data and AI can help.

Solve difficult data problems and optimize the use of your data to bring greater business value to your organization.

IBM Storage for data and AI can help you to:



Retrieve data faster

Achieve faster performance and scale as you need to.



Access data anywhere

Break down data silos to benefit from hybrid cloud and application flexibility.



Improve cost efficiency

Reduce costs with optimized archiving, data reduction and global data placement.



Innovate faster

Deploy container-native storage with global data access and built-in data protection.



Optimize costs and improve agility in AI and hybrid cloud environments

IBM® Storage for data and AI provides global data access, enterprise availability and storage services that are simple to manage, faster to access and optimized to scale.

Find out how IBM Storage for data and AI can help you modernize your information architecture to bring more value to your organization. Download the ebook **Modernize your IA for AI and Hybrid Cloud** today.

© Copyright IBM Corporation 2021
Produced in the United States of America
August 2021

IBM and the IBM logo are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.
The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

Claim: Retrieve data faster: Achieve faster performance and scale as needed.

Achieve faster performance substantiation: Faster performance is achieved because of IBM General File System technology. [GPFS](#) is a cluster file system. This means that it provides concurrent access to a single file system or set of file systems from multiple nodes. This enables high performance access to this common set of data to support a scale-out solution or provide a high availability platform

Scale as needed: Spectrum Scale can scale up to YB. The current file system size limits are: Current file system size limits GPFS 2.3 or later, file system architectural limit 2^{99} bytes—2 to the 99th power is big; 2 to the 95th power is approximately 39 billion. For comparison, 1 petabyte is 1 million billion bytes—A Yottabyte is 2 to the 80th power.

Claim: Access data anywhere: Break down data silos to benefit from hybrid cloud and application flexibility.

Substantiation: Spectrum Scale supports access of data anywhere because

Scale enables low-latency read and write access to data from anywhere in the world using AFM distributed routing and advanced caching technology. AFM expands the IBM Spectrum Scale global namespace across geographical distances, providing fast read and write performance with automated namespace management. As data is written or modified at one location, all other locations get the same data with minimal delays.

Claim: Improve cost efficiency: Reduce costs with optimized archiving, data reduction and global data placement.

Improve cost efficiency Substantiation: Using storage policies transparent to end users, data can be compressed or tiered to tape or cloud to help cut costs; data can also be tiered to high-performance media, including server cache, to lower latency and improve performance. Intelligent caching of data at remote sites ensures that data is available with local read/write performance across geographically distributed sites using Active File Management (AFM). Additionally, [IBM COS File Access](#) provides a low-cost enterprise software-defined storage solution for consolidating infrequently used (write once/modify never/read rarely) NFS and SMB files from one or more applications or filers (NAS, Windows or Linux File Servers) in IBM Cloud Object Storage located on-premises or in the IBM Cloud. Clients can use IBM COS File Access to discover and migrate existing cold files from multiple geo-dispersed file shares to COS to free up storage space.

Claim: Innovate faster: Deploy container-native storage with global data access and built-in data protection.

Global data access Substantiation: Spectrum Scale supports access of data anywhere because Scale enables low-latency read and write access to data from anywhere in the world using AFM distributed routing and advanced caching technology. AFM expands the IBM Spectrum Scale global namespace across geographical distances, providing fast read and write performance with automated namespace management. As data is written or modified at one location, all other locations get the same data with minimal delays.

Built in data protection Substantiation: IBM Spectrum Scale provides system scalability, very high-availability and reliability with no single point of failure in large-scale storage infrastructures. For additional reliability, IBM Spectrum Scale supports snapshots, synchronous and asynchronous replication, and asynchronous error diagnosis while affected input/output (I/O) operations continue. IBM Spectrum Scale offers protection of data at rest and secure deletion with file-level encryption

Claim: Optimize costs and improve agility in AI and hybrid cloud environments

Optimize costs Substantiation: Costs are optimized by using storage policies transparent to end users, data can be compressed or tiered to tape or cloud to help cut costs. Data can also be tiered to high-performance media, including server cache, to lower latency and improve performance.

Improve agility Substantiation: Agility is improved because AFM expands the IBM Spectrum Scale global namespace across geographical distances, providing fast read and write performance with automated namespace management. As data is written or modified at one location, all other locations get the same data with minimal delays.

Claim: IBM Storage for data and AI provides global data access, enterprise availability and storage services that are simple to manage, faster to access and optimized to scale

Global data access Substantiation: Spectrum Scale supports global data access because Scale enables low-latency read and write access to data from anywhere in the world using AFM distributed routing and advanced caching technology. AFM expands the IBM Spectrum Scale global namespace across geographical distances, providing fast read and write performance with automated namespace management. As data is written or modified at one location, all other locations get the same data with minimal delays.

Enterprise availability:

Substantiation: Administrators can configure the file system so that it automatically remains available if a disk or server fails. The system is designed to transparently fail over metadata operations and other services, which can be distributed throughout the entire cluster. For additional reliability, IBM Spectrum Scale supports snapshots, synchronous and asynchronous replication, and asynchronous error diagnosis while affected input/output (I/O) operations continue.

Substantiation: IBM Cloud Object storage supports always on data with up to 8 nines availability

Storage services that are simple to manage faster to access and optimized to scale Substantiation: IBM Spectrum Scale allows different applications or services to access the same data without moving or altering it. Data can be written and retrieved as either files or objects. Rather than use a copy and change gateway, IBM Spectrum Scale supports both protocols natively for higher performance and simplified administration. The common storage layer enables most IBM Spectrum Scale features, including authentication, encryption, and tiering, for both object and file storage.

