



Highlights

- Consolidate storage across traditional file and new-era workloads for object, Hadoop and analytics use cases
 - Achieve new operational efficiency and cost effectiveness—deliver up to 10 times higher performance on the same hardware¹
 - Help lower the cost of data retention up to 90 percent through cognitive and policy-driven automation²
 - Improve application performance with scale-out and flash-based acceleration
 - Enable collaboration and efficient sharing of resources among global, distributed teams
 - Transparently tier to and from cloud object storage on-premises or to the public cloud
-

IBM Spectrum Scale

Cognitive storage manages unstructured data for cloud, big data, analytics, objects and more

Enterprises and organizations are creating, analyzing and keeping more data than ever before. Those that can deliver insights faster while managing rapid infrastructure growth are the leaders in their industry. And to deliver those insights, an organization's underlying storage must support both new-era big data and traditional applications with security, reliability and high-performance. To handle massive unstructured data growth, the solution must scale seamlessly while matching data value to the capabilities and costs of different storage tiers and types. IBM® Spectrum Scale™ meets these challenges and more. It is a high-performance parallel file system for managing data at scale with the distinctive ability to perform archive and analytics in place.

IBM Spectrum Scale enables the unification of virtualization, analytics, file and object use cases into a single scale-out storage solution. IBM Spectrum Scale can provide a single namespace for all of this data, offering a single point of management with an intuitive graphical user interface. Using storage policies transparent to end users, data can be compressed or tiered to the tape or cloud to help cut costs; data can also be tiered to high-performance media, including server cache, based on a heat map of data 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).

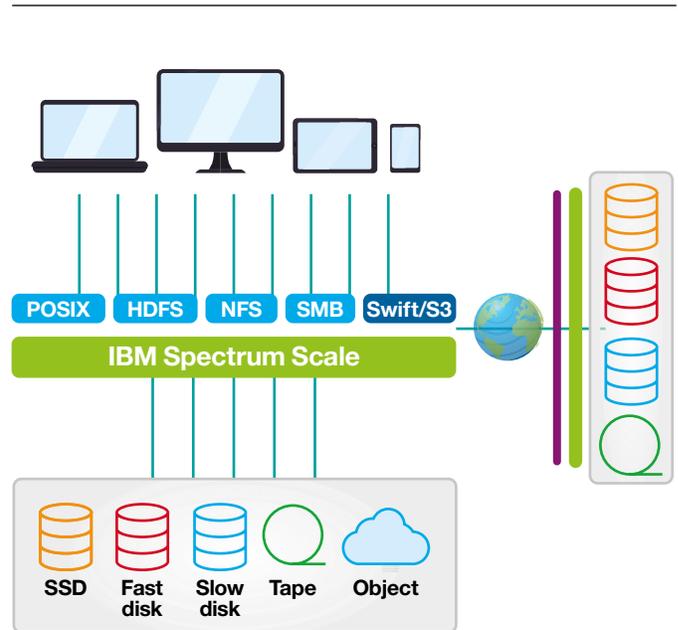


IBM Spectrum Scale is an enterprise-grade parallel file system that provides superior resiliency, scalability and control. Based on IBM General Parallel File System (GPFS™), IBM Spectrum Scale delivers scalable capacity and performance to handle demanding data analytics, content repositories and technical computing workloads. Storage administrators can combine flash, disk, cloud, and tape storage into a unified system that is higher performing and lower cost than traditional approaches. With thousands of customers and more than 15 years of demanding production deployments, IBM Spectrum Scale is a file system that can adapt to both application performance and capacity needs across the enterprise. By including IBM Spectrum Scale in the software-defined infrastructure, organizations can streamline data workflows, help improve service, reduce costs, manage risk and deliver business results today while positioning the enterprise for future growth.

Proven technology for high-performance data management

IBM Spectrum Scale is full-featured, software-defined storage with management tools for advanced storage virtualization, integrated high availability, automated tiered storage and the performance to effectively manage very large quantities of file or object data.

With the ability to independently scale in capacity, performance, protocols and resources, IBM Spectrum Scale is a clear leader in large, demanding environments. Organizations that may not have multiple petabytes of data today can start small with the confidence that IBM Spectrum Scale has already been tested in these environments.



Remove data-related bottlenecks

Slow storage negatively impacts applications, delays schedules and wastes expensive infrastructure. IBM Spectrum Scale can speed time-to-results and maximize utilization by providing parallel access to data, shared disks and storage-rich servers, improving scalability for high-performance workloads. IBM Spectrum scale is a parallel file system, where the intelligence is in the client and the client spreads the load across all storage nodes in a cluster even for individual files, while in traditional Scale-Out NAS one file can really only be accessed through one node at a time by an individual client. This parallel file system architecture allows IBM Spectrum Scale to seamlessly handle tens of thousands of clients, billions of files and yottabytes of data.

Simplify data management at scale

Part of the IBM Spectrum Storage™ family of solutions, IBM Spectrum Scale includes integrated management tools and an intuitive graphical user interface to help manage data at scale. IBM Spectrum Scale can span multiple storage environments and data centers to eliminate data silos and “filer sprawl.” IBM Spectrum Scale can cognitively spread data across multiple storage devices—optimizing available storage utilization, reducing administration and delivering high performance where needed. IBM Spectrum Scale has multiple deployment options and configurations to incorporate current NFS filers, block storage and storage-rich servers into a global namespace with universal access. The IBM Spectrum Scale file system supports interfaces for file (POSIX, NFS, CIFS), object (S3, SWIFT) or Hadoop Distributed File System (HDFS) for in-place analytics.

Empower global collaboration

IBM Spectrum 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. AFM leverages the inherent scalability of IBM Spectrum Scale, providing a high-performance, location-independent solution that masks network failures and hides wide-area latencies and outages. These game-changing capabilities accelerate project schedules and improve productivity for globally distributed teams.

Cognitive data management

IBM Spectrum Scale can help improve performance, lower costs, add resiliency or simplify collaboration with algorithmic and policy-driven data movement, copying and caching. IBM Spectrum Scale catalogs data across multiple storage pools, including the cloud. It tracks usage profiles, storage latency and a broad range of standard and custom metadata from which data movement policies can be constructed.

IBM Spectrum Scale is the caretaker of business-critical data with the ability to replicate, encrypt, compress, and distribute data across different hardware platforms, systems and data centers.

Armed with the knowledge of the data usage and the underlying storage, IBM Spectrum Scale curates data across multiple tiers of storage, including tape and cloud. The powerful data-aware intelligence engine can create optimized tiered storage pools by grouping devices—flash, solid-state drive (SSD), disk or tape—based on performance, locality or cost. Migration policies transparently move data from one storage pool to another without changing the file’s location in the directory structure. Cognitive analysis of data usage patterns can help administrators pull data back up to higher performance tiers as needed.

For example, administrators can create a rule that moves files out of the high-performance pool if the pool is more than 80 percent full—reserving premium storage for use by active file data. The information lifecycle management toolset built into IBM Spectrum Scale helps simplify data management by enabling additional control over data placement. The toolset includes storage pooling and a high-performance, scalable, rule-based policy engine.

End-to-end data availability, reliability and integrity

IBM Spectrum Scale provides system scalability, very high availability and reliability with no single point of failure in large-scale storage infrastructures. Administrators can configure the file system so that it automatically remains available if a disk or server fails. IBM Spectrum Scale is designed to transparently fail over metadata operations and other IBM Spectrum Scale 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. IBM Spectrum Scale offers the protection of data at rest and secure deletion with file-level encryption.

Highlights of IBM Spectrum Scale v4.2.2

Transparent cloud tiering: Public cloud or on-premises IBM Cloud Object Storage can be added as a tier of storage in IBM Spectrum Scale. Ideal for adding active archive storage pools or taking advantage of Storage as a Service, the design leaves end users unaffected by the data movement to and from cloud. IBM Spectrum Scale manages the metadata, movement and caching to seamlessly tier to and from any Amazon S3 or OpenStack Swift storage without the inconvenience, complexity and performance hit of adding a separate cloud or object storage silo.

Unified file and object storage: 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.

Enhanced graphical user interface: IBM Spectrum Scale has an intuitive interface that simplifies tasks and enhances monitoring. Administrators can quickly and easily monitor and provision capacity—without using the command-line interface. They can also use system health, capacity and performance to identify trends and proactively resolve issues. Multi-cluster installations can benefit from integration with IBM Spectrum Control™ for an enterprise view of their data.

Snapshots and synchronous/asynchronous replication: IBM Spectrum Scale can be part of the enterprise disaster-recovery plans with the ability to quickly back up, copy and restore data as needed. With automatic fail-over and intelligent fail-back, IBM Spectrum Scale keeps businesses and organizations up and running.

Policy-driven compression: IBM Spectrum Scale can help reduce the size of data at rest. Using policy-driven compression, administrators can specify for data to be compressed only where it will be effective and efficient (for example, when it won't slow down applications).

Encryption and Secure File Delete: IBM Spectrum Scale can encrypt data in flight or at rest with independent key management that integrates with leading enterprise key management systems.

Integration with Hadoop workloads and storage: IBM Spectrum Scale supports Hadoop workloads and the HDFS—without requiring any changes to applications. With IBM Spectrum Scale Hadoop connector, multiple IBM Spectrum Scale clusters or another HDFS repository can be federated into a single HDFS instance. IBM reduces the need to move data, simplifying the deployment and workflow of Hadoop, Apache Spark and related packages.

IBM Systems
Data Sheet

IBM Spectrum Scale at a glance

Operating systems supported	IBM AIX®; Linux: Red Hat, SUSE Linux Enterprise Server; Microsoft Windows Server 2012, Microsoft Windows 7; IBM z Systems™
Hardware supported	x86 architecture: Intel EM64T processors or AMD Opteron, minimum 1 GB system memory IBM POWER® architecture: AIX v6.1 or v7.1, Linux on POWER3 (minimum), minimum 1 GB system memory; z Systems (Linux only)
Maximum number of files/file system	2 ⁶⁴ (9 quintillion) files per file system
Maximum file system size	2 ⁹⁹ bytes
Minimum/maximum number of nodes	1 - 16,384
Protocols	POSIX, GPFS, NFS v4.0, SMB v3.0 Big data and analytics: Hadoop MapReduce Cloud: OpenStack Cinder (block), OpenStack Swift (object), S3 (object)
Cloud object storage	IBM Cloud Storage System (Cleversafe), Amazon S3, IBM SoftLayer® Native Object, OpenStack Swift and Amazon S3 compatible providers



Take the next step. [Click here.](#)
➤ See the full list of specifications.

Why IBM?

Innovative technology, open standards, excellent performance, and a broad portfolio of proven storage software and hardware solutions offerings—all backed by recognized industry leadership—are just a few of the reasons to consider storage solutions from IBM. In addition, IBM delivers some of the best storage products, technologies, services and solutions in the industry without the complexity of dealing with different hardware and software vendors.

For more information

To learn more about IBM Spectrum Scale, please contact your IBM representative or IBM Business Partner, or visit: ibm.com/systems/storage/spectrum/scale/



© Copyright IBM Corporation 2017

IBM Corporation
IBM Systems
Route 100
Somers, NY 10589

Produced in the United States of America
April 2017

IBM, the IBM logo, ibm.com, AIX, POWER, IBM Spectrum Control, IBM Spectrum Scale, IBM Spectrum Storage, GPFS, and z Systems 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 “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

SoftLayer is a trademark of SoftLayer, Inc., an IBM Company.

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

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.

¹ IBM analysis based on a list price comparison of a 1 PB IBM TS3500 Tape Library, 1 IBM Spectrum Scale license and 1 IBM Spectrum Archive EE license, compared to the cost of IBM DS5100, including annual maintenance.

² IBM case study on Cypress Semiconductor, May 2014.
<http://public.dhe.ibm.com/common/ssi/ecm/dc/en/dcc03034usen/DCC03034USEN.PDF>



Please Recycle



IBM
Spectrum
Scale