

IBM Spectrum Scale

The IBM Spectrum Scale software based on GPFS is a scale-out parallel file system. Typical large scale environments include data analytics with Hadoop Media & Entertainment, geo-sciences, life sciences, backup and archive. Spectrum scale is also offered as the ESS system, an integration of elements by IBM into a multi-featured NAS system targeted at enterprises that also supports object access, which is covered in the Evaluator Research separately.

In addition to the General Parallel File System, Spectrum Scale includes the Linux operating system as the base system and Spectrum Protect (previously named Tivoli Storage Manager or TSM) integrated for additional advanced features. OpenStack Swift open source software is installed on nodes to provide object storage access with objects stored as files. Additional software support S3 access as well.

IBM Spectrum Scale software can be installed on different servers and utilize multiple storage systems including IBM NVMe-based FlashSystem 9100 in the case of ESS. Additionally, Spectrum Scale can be installed on AWS for customers that run or develop analytics software in the public cloud.

There are multiple editions of Spectrum Scale that are offered with variations of licensing, included functionality, and operating systems supported. Details of the versions is included after the product highlights section.

Highlights

- Large scaling system – up to 16,384 nodes and virtually unlimited number of files and file systems
- Automated file placement
- Automated file migration – to another file system, AWS S3, and IBM Cloud
- Transparent Cloud Tiering to Cloud Object Storage or public clouds under policy control
- All flash or disk-based storage enclosures for external attach
- Advanced features implemented with integrated software
 - Asynchronous replication to remote Spectrum Scale system
 - Asynchronous replication for z Linux for DR and cluster support
 - Synchronous replication within a cluster – local copy
 - Snapshots – redirect on write, up to 256 per file system
 - Thin provisioning
 - Support for anti-virus software
 - NDMP
 - Encryption and secure erase
 - Compression
 - Compliance features to meet SEC 17a-4 requirements
- Global namespace across all file systems
- Advanced GUI
- Active File Manager for file migration and WAN clustering with global namespace

- Integrated management node
- Snapshot manager for file-level snapshots
- OpenStack Swift support for object storage and Cinder for block access
- S3 support for object access
- QoS on a file basis
- SAMBA-based SMB and NFS
- IBM Storage Insights – type GUI for simplified administration
- Certification as target storage for HortonWorks Data Platform
- File audit logging
- Ability to install Spectrum Scale software on AWS
- Data Acceleration for AI
- Containerized client for use with container applications to access Spectrum Scale file system.
- CSI support as Container Ready Storage
- IPV6 support

Evaluator Group Comment: The scale-out capabilities are what the value of the system is rather than advanced features or simplicity of integration. It should be noted that ESS is an all-inclusive system where the capabilities such as the Active File Manager are included with the licensing by node of Spectrum Scale.

The following section is from IBM regarding the different versions of Spectrum Scale available. Spectrum Scale is also included in the Spectrum Software Suite and IBM Storage Suite for IBM Cloud Paks. Consult the [IBM Spectrum Scale FAQ in IBM® Knowledge Center \(www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html\)](http://www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html) for the latest features included in each edition. Note that all nodes in a configuration must be running the same edition of the Spectrum Scale software.

IBM Spectrum Scale Standard Edition

Available on AIX®, Linux, and Windows. This edition provides base GPFS functions. On AIX and Linux, the available features include Information Lifecycle Management (ILM), Active File Management (AFM), and Clustered NFS (CNFS). CNFS is not available on Linux on Z. On Windows, the available features include limited Information Lifecycle Management (ILM).

On Red Hat Enterprise Linux 7.x and 8.0, SLES 12, and Ubuntu 16.04 LTS and 18.04 LTS, the available features include the ability to enable and use the additional protocol functionality integration (NFS, SMB, and object). On Red Hat Enterprise Linux 8.0 and SLES 12, the object protocol functionality is not supported. For information about the features available in respective IBM Spectrum Scale editions, see the following table.

IBM Spectrum Scale Data Access Edition

This edition provides identical functionality as IBM Spectrum Scale Standard Edition under capacity-based licensing.

IBM Spectrum Scale Advanced Edition

Available on AIX and Linux. This edition provides all the features of the Standard Edition and certain additional features. For information about the features available in respective IBM Spectrum Scale editions, see the following table.

IBM Spectrum Scale Data Management Edition

This edition provides identical functionality as IBM Spectrum Scale Advanced Edition under capacity-based licensing.

IBM Spectrum Scale Developer Edition

Available on Red Hat Enterprise Linux on x86_64. This edition provides all the features of the Data Management Edition and it is limited to 12 TB per cluster. You can check your licensed usage on IBM Spectrum Scale Developer Edition by using the `mmlslicense --licensed-usage` command.

Note This edition is made available free of cost for customers to try IBM Spectrum Scale features in test setups. Its use in production is not allowed. There is no support from IBM for IBM Spectrum Scale Developer Edition.

IBM Spectrum Scale Erasure Code Edition

This edition provides identical functionality to IBM Spectrum Scale Data Management Edition plus support for storage rich servers. IBM Spectrum Scale Erasure Code Edition provides network-dispersed erasure coding, distributing data and metadata across the internal disks of a cluster of servers. This allows IBM Spectrum Scale to use internal disks as reliable storage with low overhead and high performance.

IBM Spectrum Scale Standard Edition and IBM Spectrum Scale Advanced Edition can be licensed by IBM Spectrum Scale Client license, IBM Spectrum Scale Server license, and IBM Spectrum Scale FPO (File Placement Option) license.

The following table from IBM lists the availability of key features in the IBM Spectrum Scale editions.

Feature	Data Access	Data Management ¹	Erasure Code Edition
Multi-protocol scalable file service with simultaneous access to a common set of data	✓	✓	✓
Facilitate data access with a global namespace, massively scalable file system, quotas and snapshots, data integrity and availability, and filesets	✓	✓	✓
Simplify management with GUI	✓	✓	✓
Improved efficiency with QoS and compression	✓	✓	✓
Create optimized tiered storage pools based on performance, locality, or cost	✓	✓	✓
Simplify data management with Information Lifecycle Management (ILM) tools that include policy based data placement and migration	✓	✓	✓
Enable worldwide data access using AFM asynchronous replication	✓	✓	✓
Asynchronous multi-site Disaster Recovery		✓	✓
Hybrid cloud (TCT)		✓	✓
Protect data with native software encryption and secure erase, NIST compliant and FIPS certified		✓	✓
File audit logging		✓	✓
Watch folder		✓	✓
Erasure coding	ESS only	ESS only	✓

¹IBM Spectrum Scale Developer Edition provides the same features as Data Management Edition and it is limited to 12 TB per cluster.

Overview of System

IBM Spectrum Scale is offered by IBM to meet large-scale file storage requirements. The key capability with the underlying GPFS software is the very large scaling of up to 16,384 nodes and the ability to create a global namespace for access to large numbers of files. Standard hardware and software are integrated to provide the systems and advanced features. The scaling capability lends itself to adapting the system to meet differing usage environments by configuring dual node pairs using the clustering capabilities of GPFS.

Advanced features are enhanced with the integration of Spectrum Protect (Tivoli Storage Manager). Replication, snapshots, and thin provisioning are all features available with Spectrum Scale. Object access support is achieved with addition of OpenStack Swift open source software on GPFS nodes. In addition to the Swift protocol, S3 is supported with the enhanced S3 API software added. An iSCSI client is also available.




Figure 1: IBM Spectrum Scale Overview

Software Architecture

The Spectrum Scale software includes the following:

- General Parallel File System – The clustered file system and file management software. GPFS includes multiple functions as part of the base capabilities.
 - Asynchronous replication to a remote system
 - Synchronous replication to a local node
 - Storage pools for tiering within the file systems
 - Failover between nodes
 - Snapshots (redirect on write) with up to 256 per file system
 - Reporting of space utilization and consumption
 - Lock management
 - Thin provisioning – allocate on write only
 - File audit logging
- Protocol support for NFS
- HDFS - Hadoop MapReduce Connector
- Certified as target for HortonWorks Data Platform
- Management software for GUI and CLI access and control
- Spectrum Protect for automated migration of data
- Containerized client for container application access
- CSI support as Container Ready Storage

Spectrum Scale uses standard protocols to communicate with existing client computers. Client services are provided for UNIX clients using Network File System (NFS) versions 2, 3 and 4 (including 4.1) over TCP and UDP. UNIX clients and network clients that have implemented NFS can access data on the NAS devices. Spectrum Scale also supports industry-standard network connection of storage using the file transfer protocol (FTP) allowing direct file transfer between clients and the NAS device as well as Hypertext Transfer Protocol (HTTP) version 1.0 and 1.1. SMB is supported with a SAMBA implementation. An iSCSI client is also available.

Object storage access is provided the installation of Swift open source software loaded on each GPFS node with four daemons executing. Each storage node exposes the entire object database. Objects are stored as GPFS files. The software provides full Swift interface. The Keystone authentication service on each node.

Extended S3 object API support was added with IBM modified software to support S3.

Massive parallelism for data transfer with low latency internode communication provides the high performance for data access. Metadata is not transferred in parallel.

Spectrum Scale services

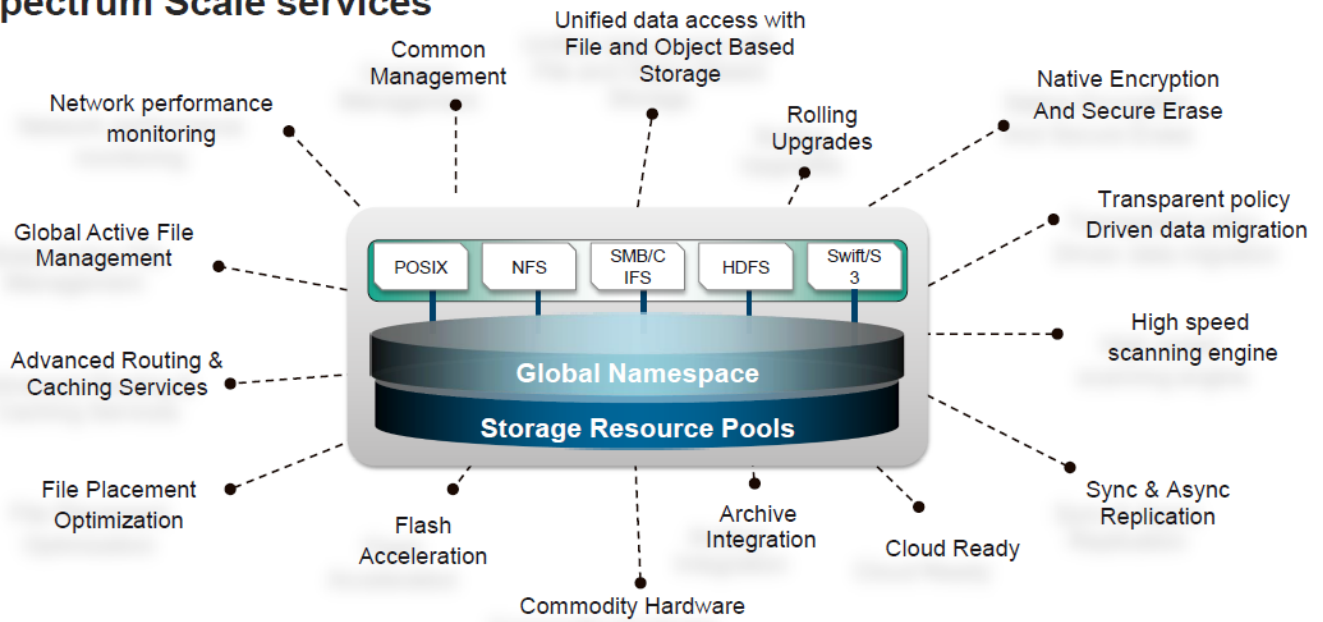


Figure 2: Spectrum Scale Features Overview (Source: IBM)

Reliability, Availability, and Serviceability

The availability of the scale-out file system using Spectrum Scale is enabled by having clustering scaling of nodes where one node can detect the failure of another and take over operation.

Clustering consists of up to 16,384 nodes. A single global namespace may be created across the file systems within a cluster.

The following summarizes the RAS capabilities:

- Spectrum Scale has a clustering failover mechanism. A node may take over the workload of a failed node. The takeover node in the system will take over the failed controller's IP address. NFS accesses will be retried.
- Redundant hot-plug cooling fans and power supplies are standard.
- Ethernet connections provide the service interface. 100GigE for cluster interconnect.
- For administration, Spectrum Scale utilizes a web interface to maintenance software or a command line interface.
- Spectrum Scale Erasure Code Edition provides protection of disk and node storage with Spectrum Scale RAID. Spectrum Scale RAID is software RAID that uses Reed-Solomon declustered erasure coding. Spectrum Scale RAID also provides a Disk Hospital functionality to identify issues, as well as end-to-end checksums.

Evaluator Group Comment: As an integration of different elements, it should be expected that features found in block storage systems should be available to be used. However, no advanced features such as replication in the attached storage systems have been exploited at this point. These could be introduced at some future point to capitalize on already available capabilities.

Software Features and Functions

Software features include:

- **General Parallel File System** – GPFS is the underlying element for Spectrum Scale. Multiple features are provided as part of the GPFS software. An unlimited number of file systems may be created. A global namespace may be created across the file systems. Included with GPFS is the ability to automate file placement into selected storage pools. The allocation of space used by GPFS allows for thin provisioning where space is only committed as required at the file level. The block sizes used in the file systems can be configured at the time of creation to match application usage.
- **Snapshots** – GPFS implements snapshots as a redirect-on-write function to make file system copies. The redirect-on-write function only targets modified data to a new location. A succession of snapshots is maintained and individual snaps can be backed up, restored, or deleted. Users can restore individual files from any snapshot. Up to 256 snaps per file system are supported. For backup to tape, an administrator can configure the system to take a snapshot of the file system prior to the start of the backup. The backup is performed with a backup server connected over the IP network. Microsoft VSS is supported with the snapshots.
- **Replication** – Two types of replication are supported: replicate within the Spectrum Scale system to create a mirror copy and remote replication to another system. The local replication is a synchronous replication, which is stateful, meaning the mirrored copy can be resynced at any time. The remote replication is an asynchronous replication that is stateless – a resynchronization of changed data only cannot be done.
- **IBM Spectrum Software (formerly Tivoli Storage Manager or TSM)** – Spectrum Protect is installed with the Spectrum Scale software but is separately licensed. The features included with Spectrum Protect can be exploited on the Spectrum Scale system. Migration of files according to policy settings. This provides what is termed “integrated ILM (Information Lifecycle Management).”

Evaluator Group Comment: Spectrum Protect is a very complicated application used to protect data. It takes special education to be able to effectively manage Spectrum Protect. While it may have been effective in time and development costs to add Spectrum Protect, it may have added extensive complications in management and overall complexity. A function developed as a basic part of the system would be more effective.

- **Quotas** – Quotas are enabled for users, user groups, or groups of files. Quotas are managed on a file system basis.
- **Reports** – Several built-in reports provide information about the system. The reports available include Node CPU Utilization, Memory Utilization, Network Utilization, and Storage Node Disk Utilization.

- **Anti-virus software** – McAfee and Symantec anti-virus software is supported with Spectrum Scale. Scanning can be done manually or through a scheduled process. Scan on access is support depended on policy settings. Spectrum Scale will load balance across all scan nodes.
- **NDMP Support** – Backup of Spectrum Scale may be done using NDMP, allowing the Spectrum Scale to directly move the data to the backup device.
- **Active File Manager** - The Active File Manager is included, which provides two advanced features for Spectrum Scale. The first feature is a policy-based file management function that supports policy controls for initial file placement and for file migration across tiered storage. Active File Manager can be used to move files from the Spectrum Scale system to other file-based storage systems, to AWS S3, to IBM Cloud, or to tape when used with an installed Spectrum Protect or Spectrum Archive client. The Active File Manager controls movement and deletion of data based on administrator defined policies and an extensive set of pre-defined metadata. The policies are defined in the GUI with very easy to understand graphics and settings.

The second Active File Manager feature is an extension of the distributed file clustering capability to multiple sites (geographically dispersed) with a global namespace across the sites. Unique to this implementation is the ability to cache files at remote sites and manage serialization of updates at remote sites (IBM calls them a cache cluster) or at the local site (called the home cluster). The caching and management is done on a file in a fileset where the fileset is a logical subtree of a file system and includes the files and file information including metadata. Spectrum Scale interface nodes are used as gateway nodes in addition to host interface control for the cache cluster to access files in the home cluster using NFSv3. Requested files have access permissions checked by the home cluster and are transferred to the cache cluster. The files are not encrypted by the Active File Manager for the transfer. A modification to a file at a remote site is can be propagated back to the home site based upon an administrator set of rules regarding exclusive write controls and if there was a loss of network access and multiple updates were made from remote sites or the local site.

Evaluator Group Comment: The number of advanced features is limited compared to competing offerings from companies such as NetApp and Dell EMC for their enterprise NAS systems. Some of the advanced features may be important to some customers. The comparative of advanced features is available in the Evaluator Group matrices for NAS. The Active File Manager is a new development that is very advanced and adds a significant capability to Spectrum Scale that is not available in competitive products currently. There can be a great number of value-add capabilities to Active Cloud Engine including compliance features such as WORM, encryption, and retention controls and the ability to support inexpensive NAS systems as the migration target – especially for archiving. The Active File Manager was developed as extensions to GPFS and available in the earlier SONAS storage system.

Transparent Cloud Tiering – Spectrum Scale can tier file data to object storage, on-premises or in public clouds. Currently, Amazon S3, OpenStack Swift, and IBM Cloud are supported with IBM Cloud Object Storage as a target for on-premises storage.

Data tiered to the cloud is compressed and encrypted and transferred over secure TLS sockets. The encryption key management is the same as used in the local system. Data integrity is assured with hash codes, checksums, and CRCs added to the data. IBM has software available that executes in a virtual machine that will validate data in a cloud. An audit log of the backup and restore operations is maintained.

Compression – Beginning with Spectrum Scale version 5, there are two libraries available for compression:

- Zlib – is meant to be use for cold data with greater reduction but at a cost of speed on read access.
- LZ4 – is used for active data and is optimized for performance on data retrieval while still providing space savings.

Data Acceleration for AI – Spectrum Scale’s Data Acceleration for AI functionality allows for access and movement of data between Spectrum Scale and on-premises or cloud object storage within a single namespace.

Administrators must create policies to select the compression library based on expected access characteristics of a file, requiring a file-level granularity.

Container Access – Spectrum Scale has a containerized client that can be used directly by container applications to access the Spectrum Scale file system. Spectrum Scale also supports the CSI offered by IBM for use as Container Ready Storage.

Performance

IBM has not released results from a SPECsfs2008 nfs.v3 benchmark for a specific configuration yet. IBM states that a pair of nodes can transfer data at 40 GB/s per storage enclosure.

Some performance accelerators include:

- Highly Available write cache that has write data in DRAM and data written to disk in a log action.
- Read data is cached in SSDs in each node.
- Metadata for the GPFS file system is stored in SSDs.

Block allocation

The block allocation segment size was changed with version 5 from 256KB to 4MB. The subblock size is a variable that can be set to match the I/O size expected, which will improve performance and space utilization for small files.

Evaluator Group Comments

Evaluator Group believes that the movement to commercialize the scale-out NAS space is validated by the significant amount of recent activity, and multiple entrants, including Weka.IO, a new owner for Lustre, Qumulo, and others. IBM's Spectrum Scale, scale-out parallel file system, is built on the cluster filesystem GPFS similar to their earlier efforts in this space, SONAS and Scale out File Services, or SoFS. Like other offerings, IBM has been in the HPC market for several years with customer configuration offerings. Spectrum Scale will compete against more integrated approaches such as Dell EMC Isilon, Panasas, HPE with Qumulo and Hitachi HNAS, Weka.IO, and Lustre in the large, scale-out file storage market.

IBM is addressing the market with a solution that contains elements that IBM already had in its arsenal. The ESS is the collection of those parts sold as an integrated solution and Spectrum Scale is a software only version.

Positives:

Performance for not only NFS, but also other file access protocols including CIFS/SMB is impressive. All of this comes with a single namespace, and the ability to manage data placement intelligently. Built-in HSM capabilities along with Spectrum Protect integration provide for data migration, data backup and archiving capabilities on a large scale.

Another positive aspect is that IBM has been developing and shipping GPFS for nearly a decade, providing ample time to work out any potential issues. Spectrum Scale continues to be advanced but many of the components are not new, which should help to lower customers concerns and reduce the possibility for significant problems for adopters.

The addition of the Storage Insights GUI and management should greatly improve the administration of the system, which had many different elements.

The Active File Manager with the ability to do file migration and deletion as well as remote caching with a global namespace is a significant advancement and will require other NAS vendors to add similar capabilities to be competitive.

Potential Concerns:

The integration of the different elements may be overly complex and not have the economies to scale to lower capacity points to meet some customer requirements. A complex product or a complex installation can create problems for customers. Spectrum Scale is clearly designed and targeted at very large-scale NAS deployments, with typical configurations beginning at nearly one PB.

The addition of the open source Swift software adds support for object access but also adds more complexity to an already complex system. More complexity makes a system more difficult to administer and for problem determination.

Massive parallelism for data transfer gives high bandwidth. When multiple files are accessed, the metadata is not transferred in parallel, which is different than some competitive products and may be an issue in some cases.

More detailed information is available at <http://evaluatorgroup.com>

Copyright 2020 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.