

Determining the Best Shared File System to use with SAS® 9.4

Updated: 26MAR2025

There are many reasons why SAS® 9.4 would need access to a shared file system - SAS Grid is the main one, but if you need to share data between your DEV/TEST/PROD/UAT system or between SAS 9.4 and Viya, you will need a shared file system as well. To determine the best shared file system for your given deployment, it is important to understand how the file system is used, the SAS I/O workload characteristics performed on it, and the stressors that SAS Foundation applications produce on the file system. For the purposes of this paper, we use the term "shared file system" to mean both a clustered file system and shared file system, even though "shared" can denote a network file system and a distributed file system – not clustered.

Before we get into information about different shared file systems, let's briefly discuss the SAS I/O characteristics.

GENERAL SAS I/O CHARACTERISTICS

SAS 9.4 creates a high volume of predominately large-block, sequential access I/O, generally at block sizes of 64K, 128K, or 256K, and the interactions with data storage are significantly different from typical interactive applications and RDBMSs.

- SAS tends to perform large block sequential IO – both Reads and Writes.
- SAS does not pre-allocate storage when initializing or when performing Writes to a file. It initially writes to pages of information to the file, then extends the file with a page of data until all the data is written.
- Reading and writing of data is done via an API that utilizes the operating system's (OS) file cache.
- Many temporary files can be created during long-running SAS jobs in SAS temporary file space referred to as SAS WORK and UTILLOC.
- SAS creates standard operating system (OS) files for its data store.
- When executing Writes, there is a single writer thread per SAS session.

There are some portions of the SAS software portfolio that can render IOPs oriented activity:

- Heavily indexed files traversed randomly
- SAS OLAP cubes
- Data manipulation and modeling of some SAS vertical solutions

However, the above tend to be small components in most SAS shops, but they cannot be ignored and they need to be provisioned on separate physical file systems. In summary, the SAS workload can be characterized as predominately large, sequential I/O requests with high volumes of data. It is very important to predetermine SAS usage patterns because this will guide optimal architecture and setup of the individual, underlying file systems and their respective, physical I/O provisioning.

SAS programs often read the same data multiple times. The program can perform logical reads from the local file cache much faster and more efficiently than re-reading it from the physical device. Operations that access small files tend to get a substantial performance benefit from the local file cache. Files that are larger than the size of the host system file cache cannot be cached in their entirety and might not be cached at all.

SAS COMPUTING REQUIREMENTS FOR SHARING DATA

In the SAS infrastructure, there might be several nodes that need access to shared data. These files all need to be accessible from any of the SAS compute nodes via the same physical pathname.

These areas include:

- Permanent SAS files – these include all SAS files (programs, catalogs, data sets, indexes, and so on) that need to persist between SAS sessions and can be shared between SAS users.
- SAS deployment and configuration files – these include all SAS binaries, configuration files, logs

from SAS servers and SAS server repositories.

- SAS WORK - if the system uses [SAS Check Point and Label Restart](#) technology.

RETAINING FILE SYSTEM DATA IN MEMORY

Data read from storage (such as a solid-state device or disk) is referred to as a physical read. When the data is retained in and read from the local file cache, the reads are called logical reads. Logical reads perform substantially better than physical reads (no hard faults are done to physical storage).

Most non-shared file systems will retain data in the local host system file cache until the memory is needed for some other purpose. In comparison, shared file systems implement a wide variety of rules that govern the retention of data in the local node's host system file cache. For example, some file systems allow tuning of the host system file cache, providing options for limiting either the amount of memory that can be used for file cache or the number of files kept in the local file cache.

MAINTENANCE OF FILE SYSTEM METADATA

File system metadata includes such information as lists of files in a directory, file attributes (permissions, creation date, and so on.) and other information about the physical data.

Various shared file systems differ in the maintenance of the file system metadata. Shared file systems are required to make information about file system metadata and file locking available to all systems participating in the shared file system. File system metadata is updated whenever a file is created, modified, deleted or extended, when a lock is obtained or dropped, and on some file systems, when a file is accessed.

The sizes of finished SAS files are not known as file creation begins. Therefore, storage space is not pre-allocated at the beginning of a file WRITE operation. As the file is written, additional file system blocks, or extents, are requested. Each extent request results in an update to file system metadata. In a shared file system, file system metadata changes are coordinated across systems. Because of this, more overhead is associated with shared file system metadata operations as compared to non-shared file systems.

Due to the differences in local file cache behavior and the additional overhead required to maintain file system metadata, programs that performed well with a non-shared file system might behave quite differently with a shared file system. To minimize overhead from the shared file system, a best practice is to maintain SAS WORK and UTILLOC directories in a non-shared file system and place only permanent data in a shared file system. A shared file system might be necessary for GRIDWORK with certain failover and high availability GRID architectures.

SAS WORK files are not shared and are considered to be private to the process that created the directory; there is no benefit for SAS WORK files to reside in a shared file system. SAS WORK files also tend to generate a higher number of file system metadata updates. This is because multiple copies of data segments might be created to support sort, tree, and other data manipulations by SAS PROCEDURES. This data amplification can be 3 – 4x the size of the incoming file and it can significantly increase the file creation, extension and deletion, and activity. This, in turn, increases metadata operations. In addition, a non-shared file system is more likely to retain SAS WORK files in the local file cache

IMPLICATIONS FOR PHYSICAL RESOURCES

Shared file systems require coordination between multiple SAS compute nodes and that places more demand on the network and storage devices. This means that the network and storage devices must be provisioned appropriately. The storage subsystem might need to support both a higher number of I/O operations for journaling and coherency and a higher throughput rate when compared to the needs of a non-shared file system.

FILE SYSTEM CONFIGURATION

There are some file system options that are available on most UNIX or Windows systems that enable or disable recording when a file is accessed. Unless the workload specifically requires information about when a file was accessed or modified, disabling the metadata update with shared file systems reduces

some of the load on the file system. On UNIX systems, mounting file systems with the option NOATIME= disables the recording of file access times. The mount option NOATIME= improves the performance of all workloads and is a recommended setting if accurate file access time is not required for correct program execution.

SHARED FILE SYSTEM CONSIDERATIONS

In this section, we discuss the shared file systems that SAS has the most experience with from either testing in-house at SAS or with a storage partner; and/or SAS customer deployments.

IBM SPECTRUM SCALE (FORMERLY CALLED GPFS™)

SAS and IBM® Spectrum Scale (version 4 or later) perform well together on IBM AIX 7 and higher, Red Hat® Enterprise Linux® RHEL 7 and higher, and Microsoft® Windows operating systems. Both permanent and temporary SAS data files are managed by Spectrum Scale with excellent throughput and low overhead for file system metadata management.

Spectrum Scale is a very mature shared file system as it has been around since 1996. The maximum capacity limitations are in the 100s of PetaBytes. It can have up to 256 file systems per clustered file system, and there are no limitations to the number of nodes that can be attached to the clustered file system.

Notable Characteristics:

- Uses Spectrum Scale Paged Pool Space to replace Host System File Cache for SAS large block IO paging to/from storage. PPS is in pinned host memory, managed by Spectrum Scale.
- Page management performance has been shown to be superior to local host file cache performance. It has performed better on heavier loaded systems where cache paging services were overwhelmed.
- Has separate volumes for file system metadata improves performance.
- Uses Native Encryption that is supported via IBM Security Key Lifecycle Manager (ISKLVM) version 2.5.01 or higher. There is no encryption for Windows. Encryption is data at rest.

Performance at Scale:

- Excellent when properly provisioned on storage with appropriate CPU resources and node memory for Spectrum Scale Paged Pool Space.
- Very hardy cluster. Node failures or cluster-wide issues are exceedingly rare. If found, they are generally caused by either improper tuning or poor storage resources/array-side provisioning for paged pool space.

Lock Management:

- Token Management via metadata (one MetaNode per open file) to handle all POSIX locking requests.

Anecdotal Negative Field Experiences:

- Highly recommend IBM Support for implementation to avoid issues.
- Rare occurrences of bugs.
- Improper provisioning or tuning by customer can cause serious issues.
- Requires extra cores and memory for file system management to perform ideally.
- Expensive.

Further details about configuring IBM's Spectrum Scale on IBM Power servers for use with SAS can be found in this paper: <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102255>

RED HAT Resilient Storage (FORMERLY CALLED GFS2)

SAS and Red Hat Resilient Storage perform well together on the RHEL 7 and higher operating system. *Please note that Red Hat is deprecating GFS2 starting with RHEL 10 [Resilient Storage Add-On will be discontinued with RHEL 10 - Red Hat Customer Portal](#).* Both permanent data and SAS temporary files are managed by Resilient Storage with excellent throughput and low overhead for file system metadata management.

Resilient Storage is a mature shared file system as it has been around since 1999. The theoretical maximum capacity limitations are 100 TeraBytes and 16 host nodes per clustered file system. Red Hat has no plans to increase their 16-node limitation for Resilient Storage shared file systems.

Notable Characteristics:

- Journaled File System
- Compatible with Red Hat multipath software
- Does not support native Red Hat encryption
- Free Red Hat support
- Reasonably priced

Performance at Scale:

- Generally, good, but can be variable based on load placed on the Distributed Lock Manager (DLM).
- DLM performance can be sensitive to a very high number of files per directory structure in a cluster, especially when the files have the extensive ACL rules applied (like Scalable Performance Data Server does). DLM performance also relies on network fabric connectivity bandwidth between host nodes and to storage.
- Many users simultaneously validating metadata-defined library structures via SAS applications (for example, populating the library panes of a SAS EG User Screen), coupled with a high number of files (>20,000) containing extensive ACL rules, can place pressure on the DLM. This is especially notable with Scalable Performance Data Server.
- Until the spin lock issues mentioned below are fixed, the maximum number of cores in individual GRID nodes attached to a Resilient Storage cluster should be 8.
- Our internal testing showed that Resilient Storage has a 5% overhead when compared with XFS on the same storage.

Lock Management:

- Requires separate Red Hat DLM product.
- DLM requires a separate IO-based network for operation.

Anecdotal Negative Field Experiences:

- Rare occurrences of node failures causing the entire cluster to fail. This is generally caused by improper implementation of disaster recovery mechanisms (for example, live backups of SAS files in the cluster file system).
- Scalable Performance Data Server usage – High numbers of files with ACL rules can degrade DLM performance
- DLM does not scale at times. It is suspected that there might be lock request serialization issues when it gets very busy.
- A spin lock issue was recently discovered that greatly impacts the performance of Resilient Storage. More details about this can be found in this Red Hat portal document:
<http://access.redhat.com/solutions/2786031>

Further details about configuring Red Hat's Resilient Storage for use with SAS can be found in this paper:
http://support.sas.com/resources/papers/proceedings11/342794_OptimizingSASonRHEL6and7.pdf.

VERITAS INFOSCALE

SAS and Veritas® Infoscale (version 7.1 or later) perform well together. Both permanent data and SAS temporary files are managed by Veritas with excellent throughput and low overhead for file system metadata management.

Veritas is a very mature shared file system as it has been around since 1991. The maximum capacity limitations are in 18 ExaBytes,

Notable Characteristics:

- Veritas can be installed as a local or a clustered file system.
- SAS testing has shown better performance with Veritas Volume Manager 256K LVM striping on FLASH, despite the underlying array preference of 64K transfers.
- Does not support transparent encryption or compression.

Performance at Scale:

- Very good on our internal test bed with a maximal 12+ GB/second IO throughput.
- Our internal testing showed that Infoscene has a 15% overhead when compared with XFS on the same storage.

Lock Management:

- Untested to date

Anecdotal Negative Field Experiences:

- None to date

Further details about configuring Veritas Infoscene for use with SAS can be found in this paper:

http://support.sas.com/resources/papers/Performance_and_Tuning_Considerations_for_SAS_using_Veritas_Infoscene70.pdf.

INTEL® ENTERPRISE EDITION FOR LUSTRE®

SAS and Intel® Enterprise Edition for Lustre® (Version 2.0 or higher) performs well on RHEL. Intel Enterprise Edition for Lustre (we'll reference as 'Lustre' below) is a commercial version of Intel's parallel file system, Lustre.

Lustre is a shared file system that has been around since 1999 but did not become widely adopted until 2012. The maximum capacity limitation is 100s of PetaBytes per clustered file system.

Lustre is available in many public cloud offerings. AWS FSx for Lustre and Azure Managed Lustre are two examples.

Notable Characteristics:

- Journaled File System.
- Does not have the disaster recovery utilities that other clustered file systems have.
- Does not support transparent encryption.

Performance at Scale:

- Generally, good performance.

Lock Management:

- Lustre Distributed Lock Manager (LDLM).
- Cache coherent. Uses Metadata Servers (MDSs) and associated metadata targets (MDTs) of the inode that owns the file for metadata locks.
- File locks are managed by the object storage servers (OSSs) and their associated object storage targets (OSTs).

Anecdotal Negative Field Experiences:

- Support is mainly via support communities and chat rooms

NFS

NFS is available on all operating systems for free. However, do not let the cost be the reason to consider using it. NFS is best used as a Write once, Read often file system. If this behavior is not how you will be using persistent SAS data files, you should not consider NFS. Also, with the overhead of Writes (see below), it is strongly discouraged using NFS for SAS WORK and UTILLOC.

NFS is a shared file system that has been around for a very long time. The theoretical maximum capacity limitation is 18 ExaBytes per clustered file system.

NFS is the protocol used with NetApp filers. NetApp filers are very popular in many public cloud offerings. AWS FSx for NetApp ONTAP or Azure Netapp files are two examples.

Notable Characteristics:

- Cheap/free
- Easy to implement

Performance at Scale:

- Large-block READ performance is acceptable aside from network affectations.
- Large-block WRITE performance is poor due to close-to-open cache consistency issues in NFS3 and NFS4.
- Private LAN segments only. From a performance perspective, it is not recommended for WAN.
- File system performance is largely dependent on the underlying network fabric's available bandwidth.

Lock Management:

- Built-in.
- NFS, in particular NFS4, can exhibit problems with file available upon just-released file locks. Other nodes might not "see" these files immediately and this leads to an unreliable infrastructure.

Anecdotal Negative Field Experiences:

- Cache Coherency issues, metadata refresh on newly released locks, and some suspected serialization issues in IO queues render this as a negative choice for WRITE intensive SAS file systems. Details on these issues can be found in the paper: [Best Practices for Configuring your IO Subsystem for SAS 9 Applications](#).
- DO NOT USE for SAS WORK/UTILLOC and other heavy WRITE permanent SAS data files systems.
- These comments hold true for all storage that uses NFS as their only file system. In particular, EMC Isilon, whose OneFS file system is NFSbased. If you plan to use this storage, please review this paper: <http://support.sas.com/resources/papers/Advisory-Regarding-SAS-Grid-Manager-with-Isilon.pdf>.

COMMON INTERNET FILE SYSTEM (CIFS)

Common Internet File System (CIFS) is the native shared file system provided with Windows operating systems. With recent patches, CIFS can be used for workloads with moderate levels of concurrency and works best for workloads that get limited benefit from the local file cache. The recommended configuration is to place SAS WORK directories on a non-CIFS file system and use CIFS to manage shared, permanent files (both SAS data files and reports/output).

With the release of the Windows 2008 operating system, many improvements were made to address performance issues and connectivity via 10-Gigabit Ethernet (GbE). These greatly improve the throughput and responsiveness of the CIFS file system.

However, workload results showed that there was relatively poor retention of data in the local file cache. Workloads that reuse data from local file cache will not perform nearly as well with CIFS when compared to a local file system.

In order to function properly and perform optimally, the following registry settings, Microsoft patches, and environment variable settings for SAS 9 software should be used:

- [HKLM\SYSTEM\CurrentControlSet\Control\SessionManager\Executive]
- AdditionalDelayedWorkerThreads = 0x28
- AdditionalCriticalWorkerThreads = 0x28
- [HKLM\SYSTEM\CurrentControlSet\Services\lanmanworkstation\parameters]

- DirectoryCacheLifetime=0x0
- FileNotFoundCacheLifetime=0x0
- Hotfix referenced by <http://support.microsoft.com/default.aspx?scid=kb;EN-US;974609>
- Hotfix referenced by <http://support.microsoft.com/kb/2646563>
- Environment variable SAS_NO_RANDOM_ACCESS = 1
- Windows hotfix [2646563](#) created for SAS users disabled file read-ahead under some circumstances

OTHER SHARED FILE SYSTEMS

There are several other shared file systems that SAS has determined should not be used with SAS Grid Manager for performance reasons. These include the following:

- Red Hat Gluster per Red Hat it is not designed to work with large, block sequential IO
- Red Hat CEPH per Red Hat it is not designed to work with large, block sequential IO
- Oracle CFS
- Parallel NFS

We will continue to review the performance of these file systems with SAS as our storage partners release newer versions.

SHARED FILE SYSTEMS IN THE PUBLIC CLOUD

In the public cloud, there are not as many choices available. The two SAS 9.4 has done the most testing with are Lustre and NFS. Details of this testing can be found on this website [62239 - SAS® on the public cloud tuning guidelines](#)

- AWS FSx for Lustre or Azure Managed Lustre
- AWS FSx for NetApp ONTAP or Azure Netapp files

CONCLUSION

This paper contains information gathered over several years from testing in-house, at SAS, with our storage partners and/or from SAS customer deployments using these shared file systems in their SAS infrastructure. It is designed to help you understand how to configure shared file system to achieve the minimum IO throughput recommended for SAS 9.4 – a minimum of 150 MB/sec/physical core.

In addition to meeting the minimum IO throughput, you need to make sure all the security features required of a shared file system are available to you when you select the shared file system for SAS 9.4.

We will update this paper as we work with new versions of these shared file systems and others. For notifications of these updates along with other hardware, storage and operating system tuning guidelines, please subscribe to the [SAS Administration Community](#).

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author:

Margaret Crevar
 SAS Institute Inc
 +1 919.623-5257
Margaret.Crevar@sas.com
<http://www.sas.com>

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies.