



SAS SYSTEM FOR WINDOWS®: INTEGRATING WITH A NETWORK APPLIANCE™ FILER

Mark Hayakawa, Senior Technical Marketing Engineer, Network Appliance, Inc.

TR-3420

ABSTRACT

The goal of this paper is to help customers maximize SAS application performance running in a Windows environment using Fibre Channel Protocol and network-attached storage and to support maximum I/O throughput for large, sequential, I/O intensive SAS applications.

Specifically, this is a guide for performance tuning in an enterprise-class environment containing the SAS application and using a Network Appliance FAServer® subsystem (The terms FAServer and filer are used interchangeably within this document). The recommendations include configuration and parameter changes to the SAS environment, Windows 2000 Server operating system, CIFS client, and NetApp filer subsystem.

TABLE OF CONTENTS

1. General Issues	3
1.1. Windows.....	3
1.2. SAS	3
1.2.1. Performance	3
1.2.2. SAS I/O Model.....	3
1.2.3. Environments	4
1.2.4. Resources	5
1.3. Network Appliance Filers	5
1.3.1 SnapDrive for Windows	5
1.3.2. Filer FCP LUNs	6
1.3.3. CIFS Support.....	6
2. Performance Tuning Recommendations.....	6
2.1 The NetApp Filer	6
2.1.1. Filer object considerations	6
2.1.2. Filer Data ONTAP considerations.....	6
2.1.3. CIFS considerations	7
2.1.4. FCP LUNs considerations	7
2.1.5. SnapDrive considerations.....	8
2.1.6. Network considerations	8
2.2 SAS Application Environment.....	9
2.3 Windows.....	10
3. SAS Configurations	10
3.1. SAS Windows Workstation Connected to Shared Department Filers	11
3.1.1. Infrastructure	11
3.1.2. NetApp Benefits to SAS users.....	12
3.2. Data Center or Department Server	12
3.2.1. Infrastructure	12
3.2.2. Fibre Channel Host Attach Kit	13
3.2.3 Fibre Channel Interfaces	13
3.2.4. NetApp Benefits to SAS users.....	14
3.3. Data Consolidation	14
3.3.1. Consolidation Using CIFS File Sharing.....	15
3.3.2. Consolidation Using FCP or iSCSI LUNS	15
4. Conclusions.....	15
5. Caveats	16
6. References	16
9. Revision History	Error! Bookmark not defined.

Table of Figures

Figure 1 - Hardware configuration and Ethernet network in a department environment	11
Figure 2 - Logical volume configuration in a departmental environment.....	12
Figure 3 - Hardware configuration and Ethernet network in a data center or departmental server environment.....	13
Figure 4 - Logical volume configuration in a data center or departmental server environment	14
Figure 5 – Data consolidation of single-user workstations to a Windows SAS server.....	15

1. GENERAL ISSUES

1.1. WINDOWS

Local file systems versus shared file systems. There are two file system paradigms of interest for SAS deployment. Modern Windows platforms offer two paradigms for file system deployment: local file systems (FAT, FAT32, and NTFS) and shared file systems (CIFS). These options are also referred to as direct-attached or SAN (local file system) or network-attached storage, or NAS (shared file system). Commercially these paradigms lead to very different storage subsystem products with different costs, capabilities, and scalability. Conceptually, they have only a couple of differences.

Local file systems. This paper will discuss the use of NTFS as the local file system. The local file system interacts with locally attached devices. More recently, Fiber Channel Protocol (FCP) attached devices and SAN devices have come into prominence. The FCP devices are treated as local SCSI devices, but are accessed across a storage network using the Fiber Channel Protocol.

Shared file systems. Shared file systems allow more efficient use of storage resources by enabling multiple hosts to access data while keeping the data management responsibilities (and costs) off the host system. CIFS scales to multiple hosts in a simple fashion that introduces some additional CPU overhead to process the data through the network stack. The CIFS file system provides flexibility and increased storage efficiency for SAS environments.

1.2. SAS

1.2.1. PERFORMANCE

SAS application performance is a complex topic. SAS is a powerful and diverse application. Specifying a "canonical" workload generated by SAS is impossible. However, at the core, SAS's main function is to load, manipulate, and store data. In any given SAS deployment, the underlying compute platform, operating system, and I/O infrastructure work together to service these data manipulation operations.

SAS performance is typically measured in terms of run time or "wall-clock time." Stated simply, the goal is to complete each SAS job in as little time as possible. There are two main components to performance in a SAS deployment: computational capability and I/O subsystem performance. An optimally tuned deployment requires a balanced consumption of these two resources.

SAS workloads are different from database workloads. SAS performs most operations by doing large block (mostly) sequential I/O. This workload has very different characteristics than a typical database on-line transaction processing (OLTP) workload. As a result, typical performance tuning techniques that improve performance for OLTP (and other database) workloads are not necessarily applicable (or even positive) in SAS environments.

The second component in SAS performance is I/O subsystem performance. SAS application performance is very dependent upon I/O subsystem performance. SAS workloads that manipulate large amounts of data are typically constrained more by I/O performance than by computational capabilities. Inspecting SAS I/O performance closely reveals several important performance characteristics.

1.2.2. SAS I/O MODEL

Megabytes per second. The most obvious characteristic of I/O performance is bandwidth. Bandwidth is typically stated in terms of megabytes per second (MB/sec) the I/O subsystem is capable of delivering to the SAS application. Obviously the more MB/sec a subsystem can deliver, the faster a SAS job can complete. The MB/sec capability of an I/O subsystem is a result of many factors, including the connection technology and speed (e.g., SCSI, Fibre Channel, or Ethernet), the storage hardware platform (e.g., JBOD (just a bunch of disks), small RAID boxes, or NetApp filers), and the number and type of disk drives.

CPU cost of I/O. Another characteristic of I/O performance is the CPU cost of doing I/O. The goal is to achieve high bandwidth levels at very low CPU utilizations. Although the absolute cost of I/O varies among host platforms, there are some general characteristics.

The lowest cost I/O typically comes from the native NTFS file system or SAN platforms. NTFS is tightly integrated with Windows and is highly tuned to minimize I/O costs. Additionally, NTFS provides a single system paradigm (i.e., no shared data among multiple hosts) that minimizes the protocol overhead associated with managing data coherency among multiple hosts.

The next lowest cost I/O typically comes from the implementation of CIFS. Like native NTFS file systems, CIFS files on network drives are mapped to a Windows drive letter. Unlike NTFS, CIFS implementations provide mechanisms for sharing data among multiple host machines. This functionality has network stack overhead not associated with native NTFS file systems.

The least cost effective I/O typically comes from third-party file system implementations layered on top of either direct-attached JBODs or SAN-attached platforms. These file systems are not as tightly integrated with the operating system as native file systems or CIFS implementations. These third-party file systems also provide advanced data sharing features (such as clustered file systems) but ultimately pay a very high CPU cost for this functionality.

I/O caching. The final and most subtle characteristic of I/O performance involves I/O buffer caching on the Windows platform. The highest performing I/O solution is one that minimizes the actual amount of physical I/O traffic. For instance, a SAS job with a dataset that is smaller than the amount of host memory needs to fetch the data from the storage subsystem only once. All subsequent access to that data can be serviced from the host buffer cache, therefore minimizing I/O traffic and maximizing performance.

1.2.3. ENVIRONMENTS

SAS performance tuning techniques depend upon the specific deployment. SAS CIFS deployments have several important characteristics; single host versus multiple host configuration, number of SAS sessions per host, and host memory size versus SAS data set size. This section outlines each of these characteristics.

SASWORK. Most SAS programs depend heavily on the ability to create and access temporary data files quickly and effectively. The location of this temporary data (referred to typically by the directory name *SASWORK*) is configurable via the SAS configuration file. The data files created and accessed in *SASWORK* are not shared among multiple SAS sessions or multiple hosts. Specifically, the data files in *SASWORK* are specific to a given SAS user session. Given this property of *SASWORK*, the specifics of where *SASWORK* is located strongly affects overall SAS performance.

Single versus multiple SAS sessions. On a given host running the SAS application, there can be either a single SAS session (or user) or multiple SAS sessions. The performance and scalability of multiple sessions is dependent upon the underlying platform's capabilities. A full exploration of how many SAS sessions a given platform can support, while important, is beyond the scope of this discussion.

There are however several aspects of multiple SAS sessions that are important for CIFS deployments. Specifically, whether the multiple sessions share the same or different *SASWORK* space is an important issue. Additionally, the amount of memory each SAS session uses and as a result, the amount of memory left for operating system caching is also an important issue. The optimal resolution of these issues is specific to the overall deployment architecture and is discussed in a later section.

Single versus multiple host deployments. SAS deployments are sometimes associated with very large data sets. Accessing and processing this data through a single host can be cost prohibitive and sometimes impossible with today's Windows platform technology. As a result, when shared file systems are used, these deployments often involve multiple host computers processing the data simultaneously. There are several

techniques for deploying and accessing or processing the data simultaneously. A full discussion of these techniques is beyond the scope of this discussion.

There are, however, several aspects of multiple SAS host deployments that are important for CIFS deployments. The main issue is the applicability of locking and caching techniques that enhance single host performance. A secondary issue is the placement and usage of SASWORK directories for each individual host. The optimal resolution is again specific to the overall deployment architecture.

1.2.4. RESOURCES

There are several objects that can be considered resources to a SAS program. These resources can have efficiency tradeoffs between them, and reducing the amount of one may increase the use of another.

CPU time is the amount of time that the processor uses to perform work. This work can be executing the SAS program and its underlying functions; reading and writing of data to local storage; or network traffic to a network-attached device.

I/O describes the unit of measure to read or write data. The data can be on local or network storage, a display, a printer, or any one of a number of different types of input or output media.

Memory is the size of the work area required to hold executable programs and data for processing by the CPU. This is the amount of memory used for the application. System memory is typically measured outside the purview of the application.

Data storage space is the amount of space needed to store and manipulate data. Data storage is usually defined as disk storage, but it can also be tape storage.

Programmer time is not simply the time it takes a programmer to code a program. In addition to the initial development, a programmer will sometimes have to maintain and upgrade the program, as well as recover the program in the event of a media failure or corrupted data.

1.3. NETWORK APPLIANCE FILERS

Network Appliance filers are designed to participate in a wide range of operating system environments that may include Windows, UNIX®, Linux®, or Mac®. The filer also has the ability to serve up data as a CIFS/NFS file server or participate in an FCP SAN or iSCSI SAN.

This paper will concentrate on a Windows environment and the access methods that are supported by the NetApp filer and Windows.

1.3.1 SNAPDRIVE™ FOR WINDOWS

SnapDrive for Windows is a NetApp product that simplifies and automates some of the more difficult tasks associated with using FCP in a Windows environment. SnapDrive installs as a snap-in to the Microsoft® Management Console. All management interactions with NetApp filer LUNs attached to your Windows server must be performed through SnapDrive.

The tasks associated with managing LUNs on the NetApp filer, such as creating or deleting LUNs, or creating or deleting Snapshot™ copies of filer volumes that contain LUNs, are performed from the SnapDrive interface. There is also an option of using a command-line interface to automate SnapDrive commands directed to NetApp filer volumes that hold LUNs.

For more information on SnapDrive, please see NetApp Technical report *Technical Overview of SnapDrive™ 2.0*, at <http://www.netapp.com/library/tr//3197.pdf>. For information on installing SnapDrive on your system, please refer to the documentation that comes with the Windows Attachment Kit that is either shipped with the filer or ordered as a separate product.

1.3.2. FILER FCP LUNS

FCP support is automatically shipped with a NetApp filer and enabled with a license code. FCP LUNs that are built on a NetApp filer are special files created on filer volumes that can be accessed as local storage on the Windows server. The LUNs are formatted as NTFS drives.

1.3.3. CIFS SUPPORT

CIFS support is automatically shipped with a NetApp filer and enabled with a license code. The filer can be configured to accommodate CIFS only files, UNIX only, or mixed.

2. PERFORMANCE TUNING RECOMMENDATIONS

When running SAS applications whose data and executable files reside on a NetApp filer, there are three areas that can affect SAS application performance; the NetApp filer, the SAS environment, and the Windows environment.

2.1 THE NETAPP FILER

The NetApp filer can be configured to optimize SAS workload performance. Because the filer is a multiprotocol appliance, there are different groups of filer parameters and objects that can be modified to improve performance.

2.1.1. FILER OBJECT CONSIDERATIONS

Efficient volume and RAID definitions. Volumes defined with too few disks will have lower performance than volumes with larger numbers of disks. The more disks in a volume, the better the performance will be. The data will be spread across more physical accessors, allowing better parallelism during reads. RAID groups with a small ratio of data disks to the parity disk will make less data storage available, but the chance of double disk failure will be lower. Conversely, a high ratio of data disks to the parity disk will make more data storage available, but the chance of double disk failure will be higher. The new RAID-DP™ filer volume option can mitigate the double disk failure issue. It is up to the owner of the data to determine the balance between safety and disk economy.

With the advent of Data ONTAP™ 7G, efficient volume sizing is now an unneeded exercise. A new storage object called an aggregate abstracts the physical nature of a traditional volume and allows volumes of any size to be defined in an aggregate with many disks. The data residing on this new volume type called a flexible volume or FlexVol™ volume is spread across all of the disks that compose the aggregate. Now a FlexVol volume can be sized from 20MB up to 16TB, and can share all of the spindles in the aggregate.

Number of disks in SASWORK volume. The SASWORK volume typically gets the most activity. In a data center implementation, the NetApp filer can easily be connected to the Windows host that is processing SAS data with FCP adapters and cables. The SASWORK data can be placed on LUNs defined on filer volumes with many disks defined to it. More physical accessors will dramatically improve performance. Even though all of the storage is not being used, the key is the number of accessors working in parallel to satisfy I/O requests.

Effective qtree definitions. Qtrees can be used to segregate data for business purposes, as well as dividing the data up into easily managed groups that can be backed up to tape using the filer's NDMP backup processes. If all the data is in one qtree, only one tape job can be run against the data at a time. By segregating the data into multiple qtrees, more tape backup jobs can be run against individual qtrees.

Proper placement of data for effective Snapshot copies. Place SAS data sets with similar backup requirements or associations on the same filer volume. One Snapshot copy of the filer volume will capture all the data.

2.1.2. FILER DATA ONTAP CONSIDERATIONS

Some of the options that can be configured will have to be set from a privileged mode. To set privileged mode, enter the following command on the filer console:

```
sasfiler> priv set advanced
```

Note: the prompt will now have an asterisk after the filer name.

To return to normal filer mode issue the following command:

```
sasfiler*> priv set admin
```

2.1.3. CIFS CONSIDERATIONS

The following filer CIFS parameters need to be changed from the filer console while in privileged mode.

Enable oplocks. Make sure that `cifs.oplocks.enable` is on. Oplocks (opportunistic locks) enable a CIFS client in certain file-sharing scenarios to perform client-side caching of read-ahead, write-behind, and lock information. A client can then read from or write to a file without regularly reminding the server that it needs access to the file in question. This improves performance by reducing network traffic. CIFS oplocks on the filer are on by default. You might turn CIFS oplocks off under the following circumstances:

You are using a database application whose documentation recommends that oplocks be turned off.

The CIFS clients are on an unreliable network.

You are handling critical data and you cannot afford even the slightest data loss.

Otherwise, you can leave CIFS oplocks on.

```
sasfiler*> options cifs.oplocks.enable on
```

Change the `cifs.neg_buf_size` setting. `cifs.neg_buf_size` is one of the CIFS options on the filer (under advanced mode) that can help increase write performance in releases later than 6.1. Setting this value to 32K+260 (33028) can help improve write performance for large transfers. The default value of this option is 16K+260. Make sure that this option is set before the clients connect to the filer. If they are already connected, do a `cifs terminate` and `cifs restart` on the filer.

```
sasfiler*> options cifs.neg_buf_size 33028
```

Change the `cifs.oplocks_opendelta` setting. Under certain workloads setting `cifs.oplocks_opendelta` to 0 can improve CIFS throughput performance by 3% to 5%. Again, if setting this value to 0 results in client disconnects, reset it to 8 (the default value).

```
sasfiler*> options cifs.oplocks_opendelta 0
```

Change the `cifs.tcp_window_size` setting on the filer to 64240. Setting this option will allow more messages to be in-flight and will allow for larger data transfers.

```
sasfiler*> options cifs.tcp_window_size 64240
```

The following filer parameters can be changed from the filer console while in privileged or normal mode.

Turn `minra` off. SAS programming generally performs sequential reads and writes. Turning minimum read-ahead off will allow for maximum read-ahead caching.

```
sasfiler> vol options volid minra off
```

2.1.4. FCP LUNS CONSIDERATIONS

Group SASWORK data on separate filer volumes. SASWORK data does not need to be backed up. SASWORK data sets can be lost and recovered with very little difficulty. Resubmitting the SAS job will recreate the data. Snapshot copies should be turned off, and `space_reservations` should also be turned

off to free more space on the SASWORK volume. The `space_reservations` parameter can be set from the MMC SnapDrive window during LUN creation or it can be set using the following command:

```
sasfiler> lun set reservation lun_path disable
```

Snapshot copies can be turned off using the following command:

```
sasfiler> vol options volid nosnap on
```

Group user data on separate filer volumes. User data sets usually require backups and high availability. Turn Snapshot copies on to ensure that the data can be recovered. Turn `space_reservations` on to ensure that there will be enough space on the filer volume to hold the entire allocated data set. `Space_reservations` is normally set to off. To set `space_reservations` on, use the following command:

```
sasfiler> lun set reservation lun_path enable
```

Consider separating NAS and SAN workloads. There are configurations in which it is advantageous to use an FC-enabled filer to share files via NAS protocols like NFS and CIFS. Filer configurations with limited numbers of disks would be one example of combining SAN and NAS workloads.

However, there are issues that would lead to separating NAS and SAN workloads. At the volume level, the most important issue is one of host filesystem synchronization. In order to get a consistent, valid, and usable Snapshot copy of a filer LUN, the LUN's host must externalize the filesystem cache before the Snapshot copy is taken. NAS filesystems are owned by the filer, and can, with the exception of database applications, have Snapshot copies taken using the normal filer Snapshot automation tool. Therefore a Snapshot copy taken using the SnapDrive tool would be appropriate for a volume that contained both CIFS and FCP LUNs, but the filer Snapshot scheduler would be appropriate only for volumes that contained only CIFS filesystems.

2.1.5. SNAPDRIVE CONSIDERATIONS

The naming convention used for the initiator groups (igroups) is based on the remote procedure calls (RPCs) from the SnapDrive management, and will not be compatible with user-selected names.

Jumbo frames is not supported on SnapDrive 3.2. However jumbo frames is supported with CIFS functions over Gigabit Ethernet.

If you are planning to use the multipath I/O (MPIO) function of SnapDrive, be aware that MPIO requires Windows 2000 Service Pack 3 or 4. The option to install MPIO will be disabled in the SnapDrive installation application when Windows 2000 Service Pack 3 or 4 is not installed on the Windows system.

There are post SP3 Microsoft Windows 2000 patches (included in SP4) that have to be installed before SnapDrive will operate properly.

Q816990. "FTDisk May Cause a "STOP Error 0x00000D1" Error Message When You Shut Down Your Computer"

Q815198. "WMI Classes Information for Multipath Drivers Is Not Displayed in WBEMTest"

Q815616/Q325040. "Clustered Disk Drive Letter Unexpectedly Changes" / "After You Restart Your Computer, an Error Message Occurs, or a Drive Letter Changes, or a Dynamic Disk Is Displayed As BASIC"

2.1.6. NETWORK CONSIDERATIONS

Gigabit Ethernet connections. Gigabit Ethernet connections will obviously provide better throughput than 10/100 Ethernet connections.

Jumbo frames. Set jumbo frames on if possible. Be aware that jumbo frames is not standardized, and can be different sizes to different vendors. Make sure that all participating hardware, including switches, in your network have jumbo frames specified the same.

XMIT/RECV buffer size. The XMIT and RECV buffer sizes should match between the filer and the Windows client. A multiple of 4K will also align the buffers to the normal filer unit of work.

2.2 SAS APPLICATION ENVIRONMENT

A SAS application environment can be configured to exploit the characteristics of the NetApp filer. The primary SAS tuning points when using a NetApp filer are those dealing with data set placement.

SASFILE usage. Input and output to CIFS, if the files are to be shared. If the files are not shared, then the input and output files can be located on a filer LUN. For recovery purposes, it is recommended that SASFILE data sets be located on a different filer volume than the SASWORK data sets.

SASWORK. Placing SASWORK on a NetApp filer can provide SAS programs with large amounts of storage. Additionally, if the SASWORK filer volume is configured with multiple LUNs spanning multiple physical filer disks, more accessors can be brought to bear for each I/O operation, greatly improving performance. Because SASWORK data sets do not need to be backed up, they should be placed on their own filer volumes, separate from any data that has a backup requirement.

Place SAS logs and list files on separate filer volumes. Placing SAS logs and list on separate volumes will reduce the I/O contention during write operations.

Memory configuration. SAS performance can be strongly affected by host memory size, SAS memory configuration settings, and data set size. Each SAS session specifies (or takes a default) configuration file that specifies several options for the session. The options of particular interest for I/O performance are

memsize: specifies how much host memory a given SAS session is allowed to use

sortsize: specifies how much data to sort at one time (each sort is broken into multiple load, sort, store temporary data phases)

maxmemquery: specifies how much memory a query is allowed to use.

In addition to the SAS settings, performance also depends on the amount of host memory, the amount of memory used by the application, and the amount of memory left for the operating system to allocate to the buffer cache. Note also that multiple SAS sessions can be running on a single host, increasing the total memory used for the SAS application and data sets.

Selecting the optimal setting for memory settings depends on the relationship of the SAS data set size to the host memory size. Consider a simple, single SAS session host configuration. Host memory, SAS memory, and OS buffer cache have the following relationship:

Host memory is divided into three categories: SAS application memory, OS buffer cache, and "other."

The SAS application memory is limited by the configuration variable memsize.

In general, the operating system requires some memory for normal OS function, and allocates the rest as buffer cache. The OS buffer cache is used to hold recently read and prefetched data as well as recently written data.

This simple memory usage description results in two classes of SAS data sets: data sets that are smaller than host memory and data sets that are larger than host memory. For data sets that are smaller than host memory, increasing the SAS memory variables to values larger than the data set size will result in minimal I/O and maximum performance. For data sets larger than host memory, reducing SAS memory consumption

and thereby allowing more memory for the OS buffer cache increases prefetch effectiveness and write caching, maximizing I/O performance and effectiveness.

The same set of variables and considerations apply for host environments with multiple SAS sessions. However, the SAS application and data set size is computed as the sum of all active sessions.

2.3 WINDOWS

There are several Windows changes that can be made to help the NetApp filer perform better. These modifications generally affect the I/O subsystem and the hardware that supports it.

Used only NetApp-approved FCP host bus adapter cards (HBAs). NetApp will certify those FCP HBAs that have been tested and found to work properly on NetApp filers. Use of noncertified FCP HBAs may generate unpredictable results.

Used only NetApp-approved FCP switches. NetApp will certify those FCP switches that have been tested and found to work properly on NetApp filers. Use of noncertified FCP switches may generate unpredictable results.

Make sure that the Windows is optimized for proper network throughput. From the Control Panel, open Network Connections. Right-click Local Area Connection and click Properties. In "Components checked are used by this connection," double-click File and Printer Sharing for Microsoft Networks. Under Optimization, notice that "Maximize data throughput for file sharing" is selected by default. Turn the option off to reduce paging activity. This can improve network performance quite a bit.

Network interface card parameters. Many network interface cards (NIC) support various performance tuning options. Setting these can improve overall performance. The Intel® Pro Ethernet card is one of the better cards to use. On this card, set Adaptive Performance tuning to "maximum bandwidth."

Increase buffer size. Increasing coalesce buffers, receive buffers, and transmit control blocks from default values can improve performance under certain conditions. DO NOT change these if the Windows 2000 machine does not have a large amount of memory. Remember that changing these options would require a restart. Consult the Help menu on the Intel Pro set config menu for more details.

Window size setting. Large window size basically increases the number of messages that can be in flight. The maximum window size that is supported on a NetApp filer is 64240. Increasing this on both the filer and Windows machine can dramatically improve performance for large transfers. You need to set the `cifs.tcp_window_size` option to 64240 on the filer.

Volume mount points. Allow for dynamic data storage growth.

Striped volumes. Multiple filer LUNs defined on different volumes can be combined to form a logical volume with many accessors. LUNs can also be defined on volumes that reside on different filer heads for addition throughput.

Network definitions. Network definitions should match the other elements of the network servicing the Windows machine.

Jumbo frames. Jumbo frames should be enabled if possible, and the size should match the other elements of the network servicing the Windows machine. Jumbo frames is not supported on SnapDrive 2.1. However, jumbo frames is supported with CIFS functions over Gigabit Ethernet.

Transmit and receive buffers. The transmit and receive buffers should match the other elements of the network servicing the Windows machine.

3. SAS CONFIGURATIONS

There are different SAS configurations in which the NetApp filer can participate. The NetApp filer's simplicity, great capacity, and multiprotocol storage capabilities can add value to most SAS configurations.

3.1. SAS WINDOWS WORKSTATION CONNECTED TO SHARED DEPARTMENT FILERS

Some SAS workloads are executed on standalone Windows workstations that process small amounts of SAS data that may be geographically dispersed. These workstations have a requirement to share data and also have the processed data backed up in a timely fashion. The filer or filers can contain a single department's data, or they can be shared among several departments.

3.1.1. INFRASTRUCTURE

The following infrastructure was used to set up a SAS environment over Ethernet using Windows CIFS protocol to files stored on a NetApp filer:

- Intel-based workstation running Windows 2000 and SAS software
- Network Appliance filer licensed for the CIFS protocol
- IP-based Ethernet network

An IP-based Ethernet network is required for the host to communicate to the filer and perform administrative functions. Network connectivity should already exist in environments where data sets over CIFS are in use.

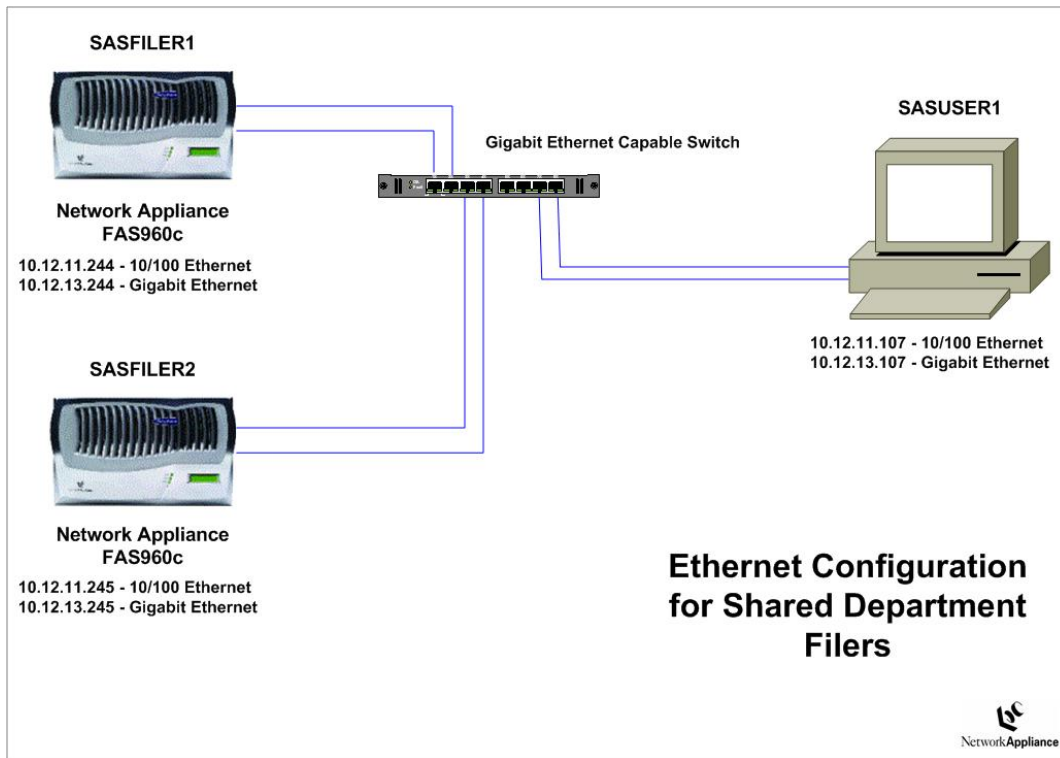


Figure 1) Hardware configuration and Ethernet network in a department environment.

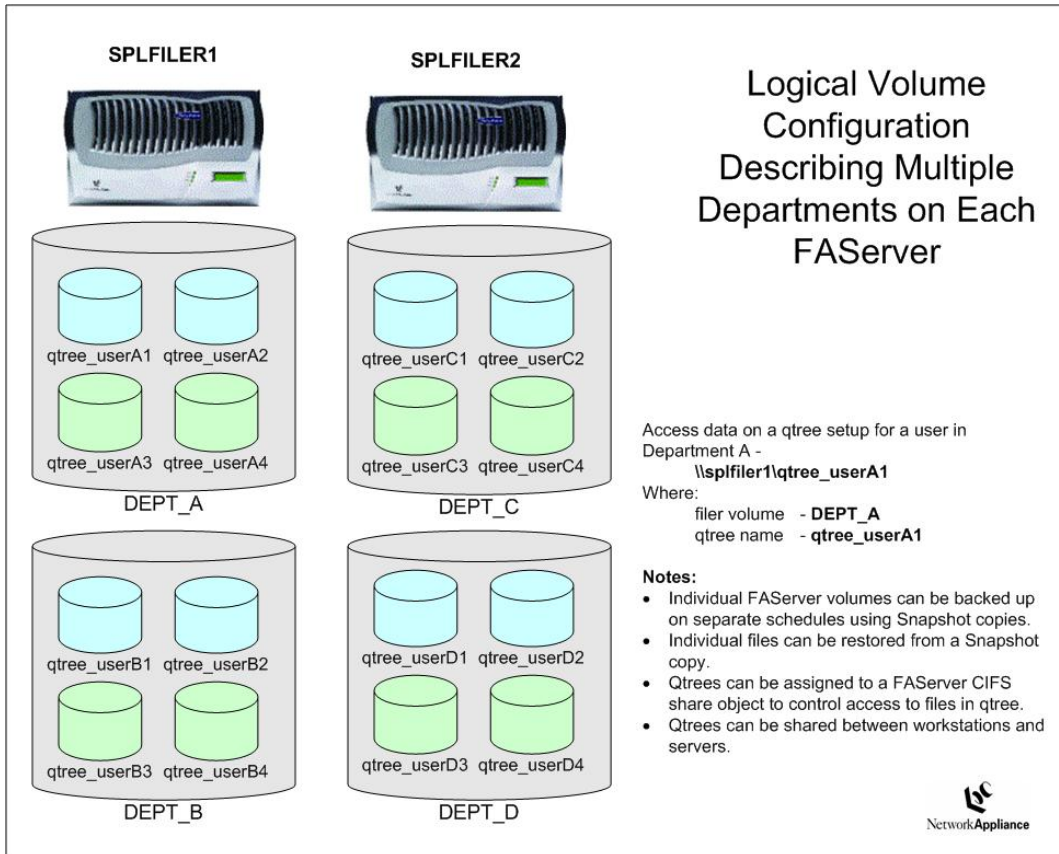


Figure 2) Logical volume configuration in a departmental environment.

3.1.2. NETAPP BENEFITS TO SAS USERS

The primary benefit to SAS users is the ability to share data between various users and groups. The data that is placed on a CIFS share defined to a NetApp filer volume will take advantage of the multiple accessors that make up the RAID group or groups. Additionally, a department's data can be consolidated at one point, making backups and physical control much easier. Finally, new storage can be added on-the-fly as user requirements change.

3.2. DATA CENTER OR DEPARTMENT SERVER

Some SAS environments will have a Windows Server in a data center or a department that services SAS RSUBMIT requests from local and remote workstations.

3.2.1. INFRASTRUCTURE

The following infrastructure was used to set up a SAS environment over Fibre Channel to a NetApp filer:

- Intel-based server running Windows 2000 and SAS software
- Network Appliance filer licensed for the CIFS protocol
- Network Appliance Fibre Channel host attach kit
- Network Appliance filer supporting Fibre Channel interfaces
- IP-based Ethernet network

3.2.2. FIBRE CHANNEL HOST ATTACH KIT

The Network Appliance Fibre Channel attach kit includes the Emulex LP9000L-F2 HBA in initiator mode, the Network Appliance Emulex HBA driver, and Network Appliance SnapDrive software for Windows.

3.2.3 FIBRE CHANNEL INTERFACES

The Network Appliance filer should be capable of Fibre Channel SAN connection via a NetApp-approved HBA in target mode installed in the filer. The Fibre Channel connection can either be directly attached between the Windows server and the NetApp filer, or the connection can be routed through a NetApp-approved Fibre Channel switch.

An IP-based Ethernet network is required for the host to communicate to the filer and perform administrative functions. Network connectivity should already exist in environments where LUNs or data sets over CIFS are in use.

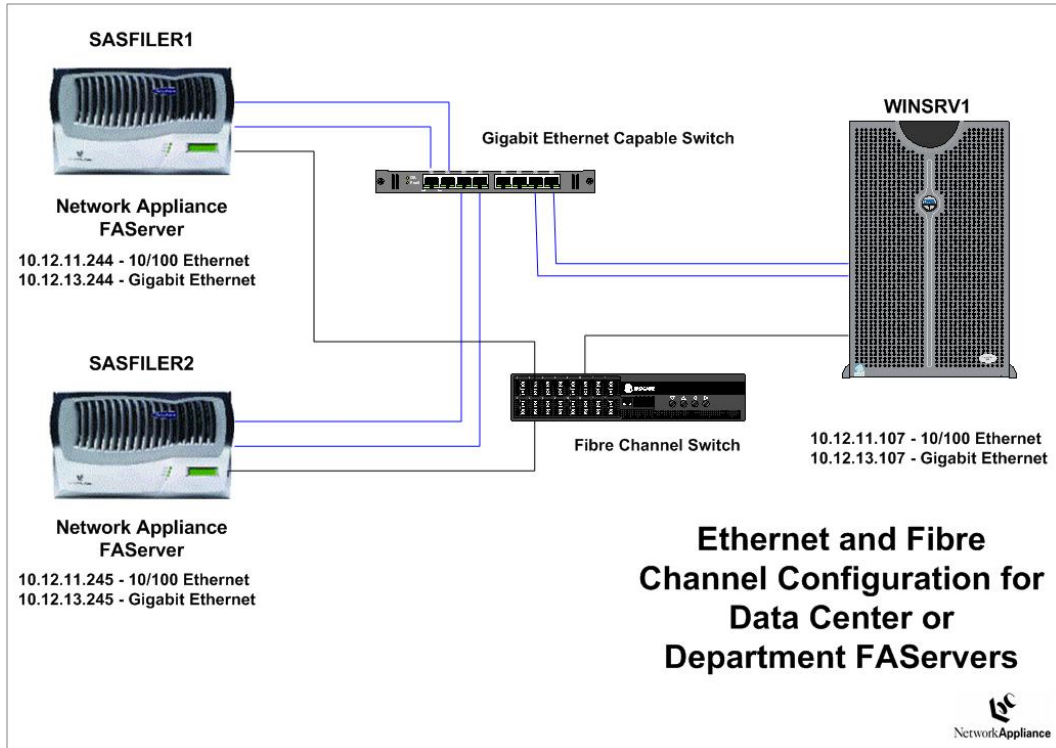


Figure 3) Hardware configuration and Ethernet network in a data center or departmental server environment.

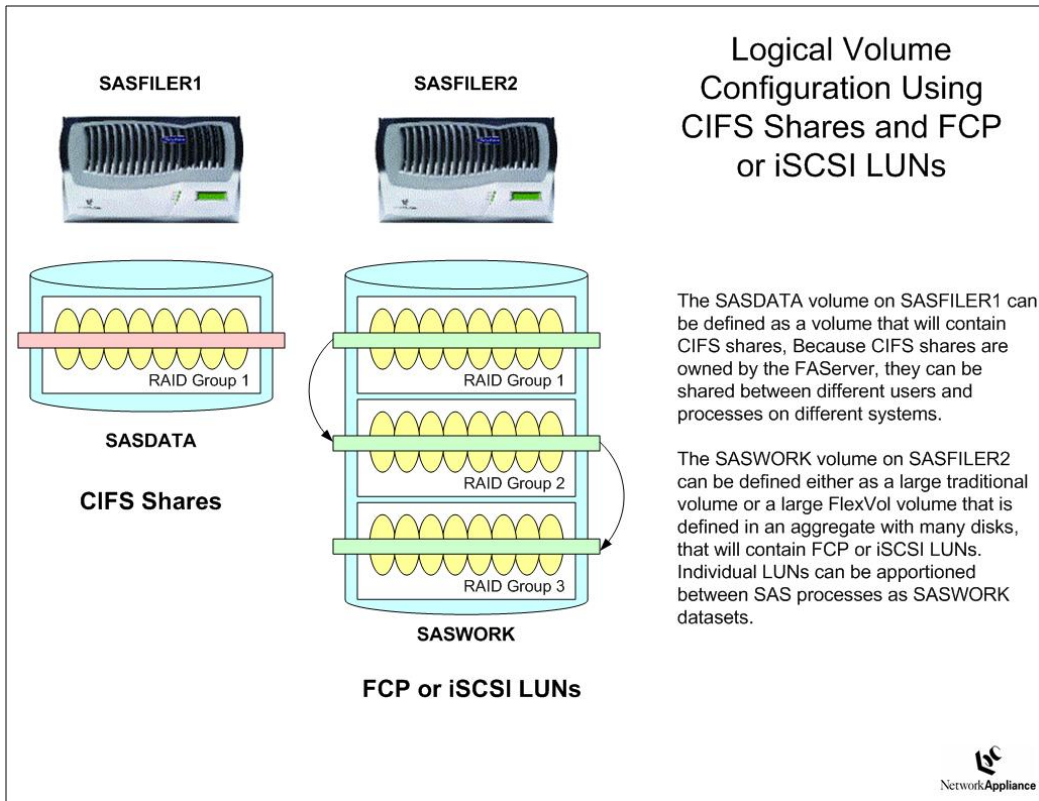


Figure 4) Logical volume configuration in a data center or departmental server environment

3.2.4. NETAPP BENEFITS TO SAS USERS

Like the benefits to SAS users in a departmental setting, the ability to share data between various users and groups is very important. The data that is placed on a CIFS share defined to a NetApp filer volume will take advantage of the multiple accessors that make up the RAID group or groups. Additionally, data centers can consolidate the data from many departments onto one or more filers, making backups and physical control much easier. Finally, NetApp storage management software can be leveraged to control multiple filers at a single control point, making the overall management of storage simpler.

3.3. DATA CONSOLIDATION

Data consolidation can be accomplished using NetApp filers by decoupling data from single servers and consolidating the data on one or more multiprotocol filers. A NetApp filer can participate in both SAN and NAS storage networks. Depending on the workload, a single filer can simultaneously participate in both kinds of storage networks. SAS data that is stored on multiple servers or workstations can be consolidated onto a single NetApp filer.

Once on the filer, the data can be shared among different SAS users accessing the data in several ways. The most basic and simplest method is with the CIFS protocol. CIFS can be used by workstations or servers to access data. Next a single SAS Server utilizing FCP or iSCSI can access SAN-attached LUNs on the filer that appear as local disks. The data residing on filer LUNs can be accessed on the SAS Server by several means; like SAS/CONNECT RSUBMIT, or a remote login session like Windows remote desktop connection.

Data residing on a filer can also be backed up at a central point, utilizing NetApp Snapshot technology. The data backed up can be recovered rapidly in case of data corruption or deletion. Additionally, tape backups of filer objects can be done with native filer NDMP tape drives or OEM backup software running on a server.

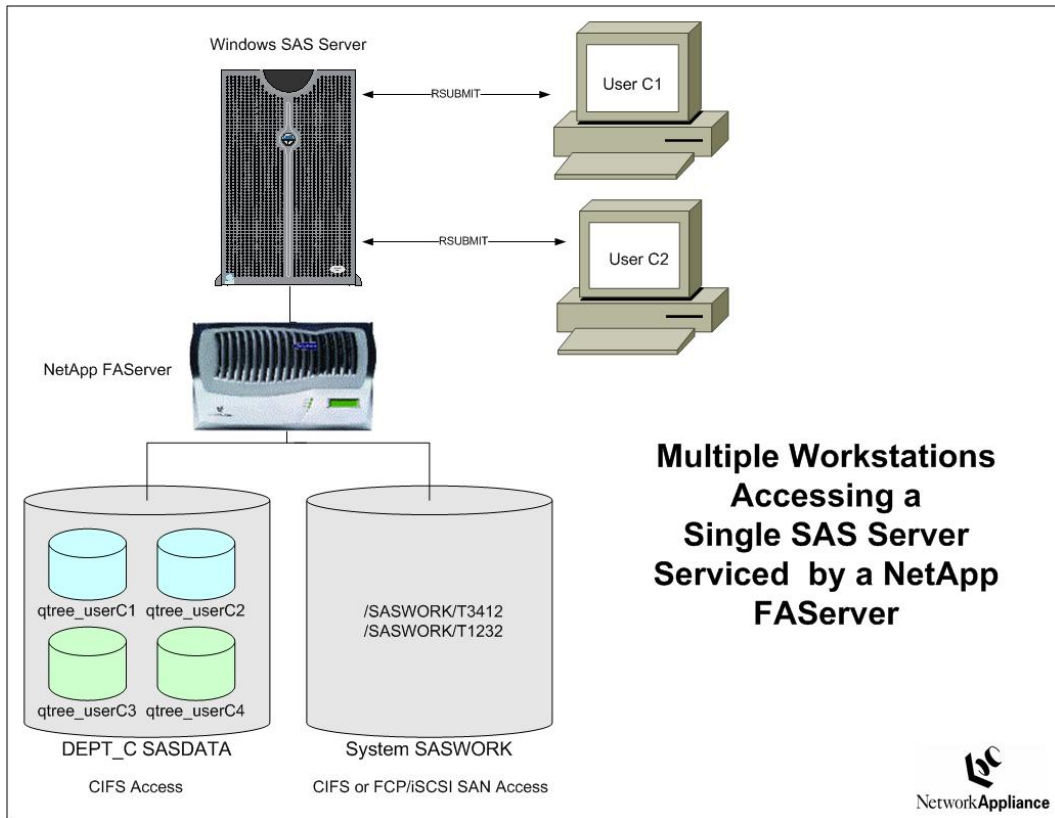


Figure 5) Data consolidation of single user workstations to a Windows SAS server.

3.3.1. CONSOLIDATION USING CIFS FILE SHARING

CIFS file sharing is native to Windows. The NetApp filer can act as a CIFS file server that will provide access to SAS data as a mapped network drive. A NetApp filer CIFS share can be mounted as a mapped network drive on a remote Windows workstation. Normal SAS processing on the remote workstation would be available to the data on the mapped network drive.

A NetApp filer volume can contain multiple CIFS shares, with each CIFS share spanning all of the disks in the RAID groups associated with the filer volume. This allows many smaller groupings of data to benefit from the speed and resilience of a larger RAID.

3.3.2. CONSOLIDATION USING FCP OR ISCSI LUNS

Filer LUNS, whether accessed using FCP or iSCSI, can be used to store SAS data, or used for SASWORK storage. Separate SAS servers can have their own volumes individually defined in separate filer volumes or several SAS servers can have LUNs defined on the same filer volume. The differentiating point could be whether the data has similar backup requirements or similar usage requirements.

4. CONCLUSIONS

The Network Appliance filer can provide large amounts of storage to SAS applications. This storage can be managed in a way that will maximize I/O throughput for SAS applications.

Additionally, the use of FCP or iSCSI LUNs defined on a NetApp filer volume can be used to provide large amounts of SASWORK data space as well as combine multiple LUNs to service different systems or users.

Finally, the use of the NetApp filer as a file-sharing appliance can consolidate data from various SAS data repositories. NetApp storage management software can be leveraged to control multiple filers at a single control point, making the overall management of storage simpler.

5. CAVEATS

The information presented in this document has been tested by Network Appliance using only a limited set of hardware and software options. Your experience may differ from that presented here. If you have any problems with the techniques shown in this paper, please contact the author.

6. REFERENCES

Network Appliance—Data ONTAP Administration Guide—available at now.netapp.com

Network Appliance—Technical report; Technical Overview of SnapDrive:

<http://www.netapp.com/library/tr/3197.pdf>

Network Appliance—Technical report; iSCSI: Accelerating the Transition of Network Storage:

<http://www.netapp.com/library/tr/3241.pdf>

Network Appliance—Technical report; SAS System on Network Appliance Filers—:

<http://www.netapp.com/library/tr/3244.pdf>

Network Appliance—Technical report; Windows 2000 and NetApp Filers: An Overview:

<http://www.netapp.com/library/tr/3113.pdf>

Network Appliance—Technical report; Introduction to Data ONTAP Release 7G:

<http://www.netapp.com/library/tr/3356.pdf>

Network Appliance, Inc.

© 2005 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, DataFabric, FAServer, FilerView, NetCache, NearStore, SecureShare, SnapManager, SnapMirror, SnapRestore, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, and WAFL are registered trademarks and Network Appliance, ApplianceWatch, BareMetal, Camera-to-Viewer, Center-to-Edge, ContentDirector, ContentFabric, Data ONTAP, EdgeFiler, HyperSAN, InfoFabric, MultiStore, NetApp Availability Assurance, NetApp ProTech Expert, NOW, NOW NetApp on the Web, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapMigrator, Snapshot, SnapSuite, SnapVault, SohoCache, SohoFiler, SpinMirror, SpinShot, SpinStor, The evolution of storage, Vfiler, VFM, Virtual File Manager, and Web Filer are trademarks of Network Appliance, Inc. in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.