

Ensuring that Your SAS® Infrastructure Is Able to Meet Your SAS Users' Demands

Margaret Crevar, SAS Institute Inc., Cary, NC

ABSTRACT

Now that you have deployed SAS®, what should you do to ensure it continues to meet your SAS users' performance expectations? This paper discusses how to proactively monitor your SAS infrastructure, with tools that should be used on a regular basis to keep tabs on infrastructure performance and housekeeping. Basic SAS administration concepts are discussed, with references for blogs, communities, and the new visual learning sites.

INTRODUCTION

Guidelines in this paper are useful for administrators who install, manage, or upgrade SAS environments. These guidelines apply to both simple and complex SAS deployments. Simple deployments typically have fewer direct interactions between machines, whereas the infrastructure for complex deployments typically spans more than one machine. In both cases, ensuring that the infrastructure is set up to support the SAS workloads and I/O profiles is a critical first step in managing the deployment. Simple deployments include single “desktop” machines that run stand-alone SAS applications such as SAS® Enterprise Guide® and SAS® Enterprise Miner™. Simple deployments also include Base SAS® programming environments used to create and run custom SAS programs, and batch environments that run those custom programs on a regular production schedule. Complex deployments include SAS Grid configurations for sharing data, and SAS solutions that leverage SAS metadata and a shared SAS middle tier.

REVIEW

First, review the [Checklist of SAS® Platform Administration Tasks](#). This document was produced as the result of extensive collaboration with many experienced SAS consultants in SAS UK, EMEA, and North America; it is a must for every SAS administrator's toolkit. The information is divided into two checklists:

- The first checklist is a set of one-time tasks that should have been performed with SAS.
- The second checklist describes tasks you should consider scheduling as a regular, recurring activity to keep your environment running in good condition.

Important one-time tasks are also covered in the Hardware Infrastructure section of the Pre-Installation Checklist that you received prior to receiving your SAS Software Depot. (A copy is in Appendix A.) In this section, we list many recommended settings for your infrastructure, with tuning guidelines for specific environments. This Hardware Infrastructure section also identifies settings and infrastructure options that you should avoid using with SAS deployments. Key points include:

- operating system on which you plan to deploy SAS
- file systems you will use for your SAS deployment, permanent SAS data files, and temporary SAS files (both WORK and UTILLOC)
- storage arrays you plan to use
- virtualization system you might use
- clustered file system that is required by SAS® Grid Manager

Ensuring that your complete hardware infrastructure (including the operating system and storage) is properly configured when you initially deploy SAS is of the utmost importance, but it is just the beginning. Continued monitoring and administration of your hardware infrastructure ensures you continue to meet the needs of your SAS users.

ANTICIPATE GROWTH

Over time, the number of SAS users will grow, the size of data being managed and analyzed will increase, and the types of analysis will change. Some growth, however, is not easily predicted and could be triggered by events such as these:

- Expanded use—Other departments within your company want the same reports that you are producing with SAS.
- Regulatory demands—The government changes the requirements of source data for reports to include more years of historical data and/or a different level of granularity.
- Company growth—Due to a merger or acquisition, the volume of data needed to project future growth of your company just increased.

When growth happens—whether predicted or not—the initial hardware put in place will probably not be robust enough to support the extra demand. Keeping the hardware infrastructure robust enough to meet the changing demands—including unpredicted growth—is the ongoing task of the administrator. Because administrators are often the last to hear about changes in business goals that expand the use of SAS, monitoring the SAS environment is critical.

The following sections identify some key proactive steps for anticipating and supporting growth by monitoring the infrastructure.

MONITOR YOUR SAS INFRASTRUCTURE

After setting up your SAS infrastructure, you should take some baselines during normal business hours (preferably on a busy day) so that you have values to compare to your additional monitoring. After you establish these baselines, monitor your hardware infrastructure on a regular basis. By gathering and reviewing general metrics, you can keep a pulse on the behavior of your SAS deployment.

There are many tools that can be used to monitor the SAS environment, including using SAS® Environment Manager that is shipped with SAS 9.4. SAS Environment Manager has predefined processes for gathering information about how SAS is being used, along with the hardware resources where SAS is running. All of the information is placed in SAS data files so that you can analyze data from your system to perform daily and weekly reporting and trend analysis. By collecting and reviewing key metrics, you can see trends that describe how your hardware resources are being used, and you can use this information to plan when to add additional resources.

In addition to SAS Environment Manager, there are numerous free hardware monitoring tools such as *nmon* for UNIX and Linux and *perfmon* for Windows, along with very sophisticated monitoring applications from Tivoli and Computer Associates.

The question becomes, what needs to be monitored. While you need to work with your SAS users to determine what matters most to them, areas to monitor include these:

- How busy are the CPUs? If they are above 80% for long periods of time and future growth is anticipated, you need to start planning to increase the number of cores.
- How much free memory is available? If the amount of free memory goes down to 0 *and* you are seeing paging or swapping *and* your paging daemons are showing pressure (or your SAS users are expressing concerns about how long their SAS applications are taking), then the system has run out of memory. If this happens for only a few minutes, you should start planning to add more RAM. If this happens for longer periods of time and these times are during peak SAS usage, you need to increase the amount of RAM in your system.
- How much free space is there in the various file systems being used by SAS? When file systems start to reach 90+% full, there is overhead to finding free space for new writes. You should work with your SAS users to determine whether there are any SAS data files that can be archived to slower storage or tape. The SAS WORK and UTILLOC file systems should be empty if there are no SAS sessions running. There are tools that you can run on a regular basis to ensure that old

temporary SAS data files are removed. These tools are discussed in the Tools to Run on a Regular Basis section of the paper.

- What is the I/O throughput? SAS recommends a minimum of 100 MB/second/core I/O throughput. Ensuring that you can achieve this I/O throughput is part of the initial hardware infrastructure setup and configuration. Note that this I/O throughput rate is a minimum; your SAS users (or at least your SAS production applications) might need more I/O throughput. Monitoring the values of reads and writes to *all* SAS file systems is a must. You should keep the numbers from the initial configuration as baseline values. Over time, compare your new I/O throughput values to determine whether you are seeing performance degradation in any of the file systems. From SAS customer experience, the temporary SAS file systems (WORK and UTILLOC) are the ones that see the most performance degradation.
- Are you experiencing any WAIT states? WAIT states means that there are not enough computer resources to support the SAS sessions. If you see WAIT states of 10+% for a process and there are plenty of CPU and memory resources, then you need to monitor the I/O subsystem infrastructure. Note that this infrastructure is more than the disks in the storage device that stores the SAS data files. It includes the fiber adapters out the front of the storage device; any switches or HBAs (host bus adapters); the network connecting the storage device to the CPU; and the network adapters that are part of the CPUs. Your I/O throughput is only as fast as the slowest component of your I/O subsystem infrastructure.
- Monitor SAS logs? If your SAS users are seeing different behavior with their current SAS jobs compared to several weeks or months ago, then you need to review the SAS logs from the runs in question to see why the times have changed. Have the data volumes being processed increased? Have more SAS steps been added to recent versions of the SAS job, thereby making them run longer? Are there any ERRORS in the SAS jobs (file system full or out of memory ERRORS)? You can monitor the SAS logs by using a SAS job that parses the SAS log and put the information you would like into a SAS data set (requires an expert SAS programmer) or with SAS Environment Manager that is shipped with SAS 9.4. Here is a link to the types of monitoring you can perform with this tool: (http://support.sas.com/rnd/emi/SASEnvMgr/EV_SGF_Paper_Monitoring_101.pdf). In addition, simply turning on FULLSTIMER can give you much information about the computer resources that your SAS job is using and whether there are any potential bottlenecks.
- How long are your production runs taking? Production jobs generally have an SLA (service level agreement) associated with them. You need to monitor how long these SAS jobs are taking. If you are reaching the SLA upper limit, determine what can be done to improve the performance of the SAS jobs. If you are not seeing any potential bottlenecks from your hardware infrastructure, you might need to work with your SAS application developer to determine whether the SAS jobs can be reengineered to meet the SLA requirements. If the reengineering means running more SAS jobs simultaneously, then you need to validate that you have enough hardware resources (CPU, memory, I/O throughput, and disk space) to support the additional SAS jobs.
- Monitor user and network application audit reports to see what users and applications are connecting to the system. This can be done with the help of SAS Environment Manager and a quick review of the SAS Metadata Repository. Remove any users who have not connected to the system recently.

TOOLS TO RUN ON A REGULAR BASIS

We outlined key expectations of the hardware infrastructure in the previous section. In this section, we recommend you run two important administrative tasks on a regular basis (daily, weekly, or monthly). These tasks help you from keeping unneeded data and processes on your system, which results in increased efficiency and improved overall performance.

- cleanwork: This SAS script is delivered with UNIX and Windows deployments of SAS. This script is part of the UNIX and Linux SAS installation in the `!SASROOT/utilities/bin` directory. (For more information, see the [UNIX documentation](#) and the [Windows documentation](#).) The *cleanwork*

script removes all WORK and UTILLOC directories associated with a SAS process that is no longer running. The script takes the hexadecimal process ID (PID) number from the name of the SAS WORK and UTILLOC directories that was created when the SAS process started. This ensures that the *cleanwork* utility will only remove old SASWORK and UTILLOC directories which are no longer in use – those directories currently being used by active SAS sessions will not be affected. Most of the time, SAS sessions will remove their own work directories automatically. However, if SAS is terminated abnormally and leaves its work directory behind, then the *cleanwork* utility will ensure you can reclaim that disk space.

- Delete orphaned SAS Workspace Servers: The current way SAS users are submitting their SAS code is through SAS clients (SAS Enterprise Guide, SAS Enterprise Miner, SAS Studio, and so on). This is great, but these users think that by “killing” the client GUI on their desktop system, it also terminates the back-end SAS processing. Prior to SAS 9.4, the backend SAS servers continued to run, but now there are tools that will terminate idle SAS servers. For more information about this tool, see [SAS 9.4 Intelligence Platform: Application Server Application Guide](#). And there is another tool that can be run on a regular basis to detect (and terminate) these orphaned sessions for you. For more information about this tool, see [SAS Usage Note 45627](#).

SAS FILES TO BACK UP AND WHEN

In addition to making sure you have enough computer resources to meet your SAS users' demands, you need to ensure that critical SAS files (SAS data sets, SAS catalogs, SAS executables, and so on) are backed up in a format from which they can be restored. The key to the previous sentence is *critical SAS files*. The reason to mention this is because the amount of time it takes to back up hundreds of TBs of SAS data files might exceed your downtime window for each day. It is very easy to back up all the SAS data files, but there comes a time when your backup window is not long enough to back up all the SAS files, so you need to determine what the critical SAS files are.

Before we get into the discussion of what SAS files need to be backed up, read this paper to understand the types of data that needs to be backed up: [Backing Up SAS Content in Your SAS®9 Enterprise Intelligence Platform](#). Remember that SAS data files should be closed for update before you back them up, especially the data files associated with the SAS Metadata Server and other SAS servers. For more information about how to back up data files associated with SAS servers, review this webinar available at <http://www.sas.com/apps/webnet/video-sharing.html?bcid=4232011724001>.

To help you understand what I mean about critical SAS files, let me review a recent discussion with a SAS customer. This customer does a complete backup of their entire system (operating system, all SAS files, and so on) every Saturday, and this full backup takes more than 30 hours. However, they do a *daily* ETL process, which means their actual downtime window to do this full backup is only 10 hours. We are working with them to determine the following information so that they back up critical data only:

If there have been no changes to the operating system, then they do not need to back up the operating system files every week.

If there have been no changes to the SAS installation (executables), then they do not need to back up this directory every week. They can make a copy and keep it until they make an update to the SAS installation.

Of their hundreds of TBs of SAS data files, how many are critical (the data is stored only in SAS) compared to other SAS data files that are actually extracts from data in an external data store? When a backup takes several hours to create, restoring a single file from this large backup might take longer than re-creating the SAS data file from the external database where the original data is stored.

Of the critical data, do these files change daily or weekly? If not, consider a quarterly backup of the unchanged data and keep that backup forever.

As you can see, determining what the critical SAS files are requires a detailed technical assessment of your SAS applications. This is more than a simple discussion; you need to talk to all your SAS users, or at least the owners of your production jobs. In addition, you can monitor all the SAS data files using the

following UNIX commands or a SAS macro to see how often they are accessed and updated. (See Appendix B.) Here are the UNIX commands you can run:

- `find /xxx -atime 90` This command finds files accessed in the last 90 days in the /xxx directory.
- `find /xxx -mtime 90` This command finds files modified in the last 90 days in the /xxx directory.

This is a good way to monitor how often the files are accessed and modified so that you can understand if they need to be backed up and on what schedule.

In addition, you should periodically test the process of restoring the SAS data files from the backup. In this way, you can determine how long it takes to restore the entire backup (in case of a disaster) or to restore a single SAS data file (that was deleted or corrupted by accident). The time to discover if this process is working correctly is not after a disaster occurs.

IS FAILOVER OR HIGH AVAILABILITY REQUIRED?

The next discussion that my team often has is how to maintain high availability (HA) of the SAS servers on the various systems they run on and what is needed to support failover in the case of a disaster. I have this discussion after the backup discussion because it is required to know the most critical files for your production and the ad hoc SAS applications before you put together a failover or HA design.

The first question to answer: Is this a requirement of the SAS application or of your IT shop in general? The reason for this question is that it helps determine how to design your complete SAS infrastructure.

What is the longest amount of time that your SAS applications can be down? This might have many answers, such as 1—2 days for your ad hoc SAS applications, but 1—2 hours for your production SAS jobs.

Are you failing over to hardware in your existing data center? Or to a data center in a different city? The latter seems to be the more popular answer here. This goes back to how are you going to keep the various SAS data in sync between the two data centers? My team has been involved in lots of conversations with the IT staff at SAS customer sites who have purchased various storage arrays that can perform real-time copying of data from one storage array to another. Remember that SAS data files need to be “closed” before you copy them. Real-time updates of files that are opened do not give you the results you would like. There are several SAS data files (such as the SAS Metadata Repository) that need to have the SAS server “stopped” and a backup macro run before you can do the copy. We are letting you know this to ensure that the data you are copying to the other data center is usable.

Another area where my team frequently sees an issue is with a hardware design that uses two NIC cards in the system and the cards are in a failover mode. (One is active; the other is passive, waiting to go active when the first one fails.) If your SAS I/O throughput demands can be met with a single NIC card, then this is okay, but most SAS users need multiple NIC cards and “wasting” a card as a hot backup does not make sense. Use multipathing to stripe across the two cards. They will both be used and if one fails, the other one will work.

Note that HA has been built into many of the SAS 9.4 products, but you need to understand how it works and how not to misuse it. For example, the SAS Metadata Server has a clustering availability to it with SAS 9.4. Also note that all of the SAS Metadata Servers need to be in the same data center as they share a single repository. We have had discussions with customers who would like to have one of the three mandatory SAS Metadata Servers in a different data center.

SAS Grid Manager and the required shared file system really help with most HA needs. A question that does come up frequently is, will SAS Grid Manager automatically restart SAS jobs that are running on a failed SAS Grid node? The answer is yes it can, if you have first enabled the Checkpoint and Restart capability in your SAS code and in your Grid queue configuration.

SUPPORT AND MAINTENANCE

Ensure that all your SAS platform administration staff is familiar with how to contact SAS Technical Support for help, along with the SAS Technical Support website, support.sas.com, which has answers to

frequently asked questions. There is an [Ask the Expert](#) section on this website that includes lots of information about SAS and also a video about the [Top 10 SAS Support Resources](#).

There is a SAS Administration community that has lots of information. This community is a way for SAS and other SAS users to share tips and tricks with regard to administering SAS.

In addition to the SAS Administration community, there is a SAS Administration blog; go ahead and make this site a favorite: <http://blogs.sas.com/content/sgf/tag/sas-administrators/>.

CONCLUSION

From the [Checklist of SAS Platform Administration Tasks](#), I have highlighted the tasks that address the issues that my team encounters most frequently when investigating problems in a SAS environment. Note that the information in this paper works for all SAS® Foundation applications, including SAS Grid Manager.

Ensuring that your complete hardware infrastructure (including the operating system and storage) is properly configured when you initially deploy SAS is of the utmost importance. Taking a few weeks to make sure you have the complete hardware infrastructure set up correctly at the start of your SAS application is a must. A phrase I seem to say frequently is, "Failure to plan on your part does not constitute an emergency on my part." But as you can see, planning for the initial deployment is not the only critical action for you; you need to monitor the hardware infrastructure to ensure it remains adequate as the SAS application data and users grow over time.

ACKNOWLEDGMENTS

Many thanks to the SAS Global Enablement & Learning team for putting together the Checklist of SAS Platform Administration Tasks. Many thanks to Ken Gahagan for sharing his CLEANOLDDATAFILES SAS macro. And thank you to Donna Bennett, Tony Brown, Rob Collum, Gretel Easter, Jim Kuell, and BJ Mulligan for reviewing the paper and advising on its content.

RECOMMENDED READING

- Checklist of SAS Platform Administration Tasks available at <http://support.sas.com/resources/papers/Platform-Administration-Tasks.pdf>

CONTACT INFORMATION

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

Margaret Crevar
100 SAS Campus Drive
Cary, NC 27513
SAS Institute Inc.
Margaret.Crevar@sas.com
<http://www.sas.com>

APPENDIX A

PRE-INSTALLATION RESOURCE DOCUMENTATION

HARDWARE INFRASTRUCTURE

When SAS is used for analysis on large volumes of data (in the GBs), SAS reads and writes the data using large-block sequential I/O. To gain the optimal performance from the hardware when performing these I/Os, we strongly recommend that you review the following information to ensure that the infrastructure (CPUs, memory, and I/O subsystem) is configured as optimally as possible.

OPERATING SYSTEM TUNING GUIDELINES

Tuning guidelines for working with SAS on various operating systems can be found in [SAS Usage Note 53873](#).

CPU GUIDELINES

SAS recommends using current generation processors whenever possible for all systems.

MEMORY GUIDELINES

For each tier of the environment, SAS recommends the following minimum memory guidelines:

- SAS Compute tier: A minimum of 8GB of RAM per core
- SAS Middle tier: A minimum of 24GB or 8GB of RAM per core, whichever is larger
- SAS Metadata tier: A minimum of 8GB of RAM per core

It is also important to understand the amount of virtual memory that is required in the system. SAS recommends that virtual memory be 1.5 to 2 times the amount of physical RAM. If, in monitoring your system, it is evident that the machine is paging a lot, then SAS recommends either adding more memory or moving the paging file to a drive with a more robust I/O throughput rate compared to the default drive. In some cases, both of these steps might be necessary.

BEST PRACTICES OF I/O CONFIGURATION

Configuring the I/O subsystem (disks within the storage, adapters coming out of the storage, interconnect between the storage and processors, and input into the processors) to be able to deliver the I/O throughput recommended by SAS keeps the processor busy, allows the workloads to execute without delays, and makes the SAS users happy. Here is the recommended I/O throughput for the typical file systems required by the SAS Compute tier:

- Overall I/O throughput needs to be a minimum of 100—125 MB/sec/core.
- For SAS WORK, a minimum of 100 MB/sec/core.
- For permanent SAS data files, a minimum of 50—75 MB/sec/core.

For more information about how SAS does I/O, see the [Best Practices for Configuring Your I/O Subsystem for SAS®9 Applications \(revised April 2015\)](#) paper.

In addition, it is a good idea to establish baseline I/O capabilities, before users begin placing demands on the system, as well as monitor the I/O if users begin suggesting changes in performance. To test the I/O throughput, platform-specific scripts are available:

- [Testing Throughput for Your SAS®9 File Systems: UNIX and Linux Platforms](#)
- [Testing Throughput for Your SAS®9 File Systems: Microsoft Windows Platforms](#)

FILE SYSTEM GUIDELINES

The Best Practices for Configuring Your I/O Subsystem for SAS 9 Applications (revised April 2015) paper lists the preferred local file systems for SAS (that is, JFS2 for AIX, XFS for RHEL, and NTFS for Windows). Specific tuning for these file systems can be found in the operating system tuning papers mentioned earlier.

For SAS Grid implementations, a clustered file system is required. SAS has tested SAS Grid with many file systems and the results of that testing along with any available tuning guidelines can be found in the [A Survey of Shared File Systems \(updated October 2014\)](#) paper. In addition to this paper, there are more detailed papers about the Red Hat GFS2 and IBM GPFS clustered file systems in [SAS Usage Note 53875](#).

Because of SAS WORK (the temporary file system for SAS applications), which does large sequential reads and writes and then destroys these files at the termination of the SAS session, SAS does not recommend NFS mounted file systems. These systems have a history of file-locking issues and the network can negatively influence the performance of SAS when accessing files across it, especially when doing writes.

STORAGE ARRAY GUIDELINES

Storage arrays play an important part in the I/O subsystem infrastructure. SAS has several papers about tuning guidelines for various storage arrays that are available in [SAS Usage Note 53874](#).

MISCELLANEOUS

In addition, there are several resources about how to set up the infrastructure to best support SAS that are available for your review:

- [Grand Designs: Why It Pays to Think About Technical Architecture Design Before You Act](#)
- [How to Maintain Happy SAS®9 Users \(revised April 2015\)](#)
- [A Guide to SAS® for the IT Organization](#)
- [Top 10 Resources Every SAS® Administrator Should Know About](#)
- [Guidelines for Preparing Your Computer Systems for SAS®](#)
- [SAS Administrators Blog](#)

Finally, SAS recommends regular monitoring of the environment to ensure ample compute resources for SAS. Additional papers that provide guidelines for appropriate monitoring are available in [SAS Usage Note 53877](#).

APPENDIX B

```
/* ***** */
/* NAME: CLEANOLDDATAFILES */
/* DESCRIPTION: This macro will process the specified directory */
/* tree and delete all non-hidden entities older than the */
/* specified date. */
/* */
/* PRODUCT: SAS */
/* USAGE: %CLEANOLDDATAFILES("/path","datetime") */
/* */
/* The first parameter specifies the location that the macro will */
/* use to begin file system traversal. */
/* */
/* The second parameter specifies the oldest last modified */
/* date that should be retained. All files with a last modified */
/* date/time prior to the specified datetime will be deleted */
/* */
/* This macro will not process hidden files or directories and */
/* any attempted deletion of a directory that contains a hidden */
/* file or directory will fail. */
/* */
/* This macro will only delete entities for which the user */
/* running the macro has the authority to delete. */
/* */
/* WARNING: */
/* If this macro is invoked by a pooled workspace server all */
/* files that match the criteria for which the pooled workspace */
/* server has permissions will be deleted. Use this macro */
/* carefully. SAS accepts no responsibility for the actions */
/* performed by this macro. */
/* */
/* USE AT YOUR OWN RISK! */
/* ***** */
```

```
%macro appendFiles(root, levels);
  %do l = 1 %to &levels;
    %let d= d_&root._&l;
    proc append base=d_&root data=&d;
      run;
    proc delete data=&d;
      run;
  %end;
%mend;
```

```
%macro procDirectory(dname, before, lev_root);
%let dsn=work.d_&lev_root;
```

```

data &dsn (keep=lr before_dt dir fname isdir lastModDT del bytes);
length dir $1024.;
length fname $1024.;
length lr $256.;
lr=length(trim("&lev_root"))||"&lev_root";
before_dt=input(&before,datetime19.);
dir=&dname;
level=1;
rc=filename("startdir",&dname);
did=dopen("startdir");
if did > 0 then do;
  numents = dnum(did);
  do i = 1 to numents;
    fname=dread(did,i);
    rc=filename("dent",&dname||"/"||fname);
    did2=dopen("dent");
    if did2 > 0 then do;
      isdir="Y";
      nextlvl="&lev_root" || "_" || level;
      dirname=&dname || "/" || fname;
      bytes=.;
      lastModDT=input(dinfo(did2,"Last Modified"),datetime19.);
      if lastModDT < before_dt then
        del = "Y";
      else
        del = "N";
      output;
      rc=dclose(did2);
      call execute('%procDirectory("'"||trim(dirname)||"', "'||&before||"',
' || compress(nextlvl) || '');
      level=level + 1;
    end;
  else do;
    isdir="N";
    fid=fopen("dent");
    if fid > 0 then do;
      bytes=input(finfo(fid,"File Size (bytes)"),20.);
      lastModDT=input(finfo(fid,"Last Modified"),datetime19.);
      if lastModDT < before_dt then
        del = "Y";
      else
        del = "N";
    end;
    output;
    rc=fclose(fid);
    rc=filename("dent","");
  end;
end;
end;

```

```

rc=dclose(did);
level=level - 1;
call execute('%appendFiles('||"&lev_root"||', '||level||')');
end;

%mend procDirectory;

%macro CLEANOLDDATAFILES(dname, before);
proc delete data=work.CGW_STATUS;
run;

data _null_ ;
  %procDirectory(&dname, &before, 1);
run;

proc sort data=work.d_1; by descending lr;
run;

data work.CGW_STATUS (drop=todel rc);
set work.d_1;
length status $256.;
if del="Y" then do;
  todel=trim(dir)||"/"||trim(fname);
  rc=filename("f", todel);
  rc=fdelete("f");
  if rc = 0 then status="Deleted" ;
  else status="Could not delete: " || sysmsg() ;
end;
else status="Retained";
run;

%mend CLEANOLDDATAFILES;

%CLEANOLDDATAFILES;

```

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.