

SAS Grid Computing Technical FAQ

July 31, 2007

Application Planning/Configuration

Q. What are the main uses for running SAS on a grid?

A. There are three major functions of a SAS grid, all of which make use of intelligent resource-based load balancing of the SAS workload:

- Workload balancing enables multiple users in a SAS environment to distribute workloads to a shared pool of resources. Better performance is achieved by allowing the grid middleware to handle allocation to jobs to the appropriate machines. This provides greater scalability & manageability for environments with numerous SAS users and SAS jobs.
- Accelerate processing for a long running SAS job. Allows users to distribute parallelized SAS workloads to a shared pool of resources to automatically find and use the best available resource. The ability to have pieces of work run in parallel on different nodes in the grid allows the over all job to complete much more quickly than if all the pieces were run in sequence.
- Scheduling jobs to a grid enables users to distribute jobs within workflows to range of hosts & automatically find and use the best available resource for each job.

Q. What SAS products and solutions take advantage of a SAS Grid automatically?

A. SAS Data Integration Server: Data Integration Studio includes a Loop Transform which allows portions of a job flow to be submitted to multiple processors either on an SMP box or on a grid.

SAS Enterprise Miner: Allows parallel execution of independent nodes either on an SMP box or on a grid.

Several other SAS products can take advantage of a grid with some minimal customization.

Ideally any SAS program (with proper modifications) including stored processes can be deployed in a grid environment

Q. How does the SAS Grid Manager environment compare to using the MP CONNECT feature of SAS/CONNECT.

A. The capabilities available with SAS Grid Manager are an enhancement to the parallel processing capabilities of MP CONNECT. We have integrated SAS/CONNECT with the grid expertise of Platform Computing in order to bring dynamic resource based load balancing and better monitoring capabilities into the SAS environment. Therefore SAS/CONNECT is required along with SAS Grid Manager.

Q. What products are needed on the grid nodes?

A. SAS Grid Manager must be licensed for the number of nodes in the grid. SAS Grid Manager includes the Platform Suite for SAS. LSF, which is part of the Platform Suite for SAS, needs to be installed on all the nodes in the grid. Base SAS and SAS/CONNECT are also required, as well any SAS products required by the execution that is being distributed. For example, if you are running SAS/STAT in a grid environment, you would need SAS/STAT on any of the grid nodes where the SAS/STAT programs can run.

Q. Does the installation and configuration of SAS Grid Manager require a plan file?

A. Yes. There are sample plan files for Data Integration, Enterprise Miner or Foundation SAS grid installations, or a custom plan file can be used.

Q. Can grid computing be done with SAS Grid Manager via Foundation SAS without a full Business Intelligence (BI) implementation?

A. Yes. SAS Grid Manager does require the SAS Metadata Server and use of a plan file during installation and configuration, however SAS BI Server or SAS Enterprise BI Server are not required.

You do have to define a SAS Grid Logical Server in SAS Management Console in order for the SAS application to “see” the grid. You must also add the appropriate grdsvc_enable calls to your SAS program unless the program is generated from a grid-automated solution such as Enterprise Miner or Data Integration Server.

Any SAS program that can be segmented into sub-tasks of execution and/or data can be run in a grid, and/or multiple SAS programs can be submitted to the grid for load balancing.

Q. How hard is it to grid enable a SAS program?

A. The answer depends on the stage of development of the SAS program. For newly developed SAS applications both SAS Enterprise Miner and SAS Data Integration Studio are now grid automated and can generate SAS programs that can automatically leverage a grid environment.

For existing SAS applications that have incorporated the parallel processing capabilities of SAS/CONNECT, simply adding a grdsvc_enable for all of the connections and defining the appropriate metadata for the grid environment enable these applications to leverage the grid environment (see sample code at <http://support.sas.com/rnd/scalability/grid/gridfunc.html>)

For SAS applications being developed from scratch, the appropriate SAS/CONNECT statements must be included to identify the parallel code segments, the appropriate grdsvc_enable calls added and the appropriate metadata defined in order for the application to leverage the grid.

Q. Can SAS Business Intelligence components benefit from a grid environment?

A. Absolutely. Anything that launches a stored process (Add-in for Microsoft Office for example) can benefit from the grid if that stored process was written with the SAS/CONNECT grid enabled statements. You could either write this by hand or save a SAS Data Integration Studio and SAS Enterprise Miner generated program as a stored process.

If you want to load-balance several SAS programs over a pool of resources (as opposed to accelerating a single app by breaking it into parallel pieces) you can wrapper the generated program (from SAS Enterprise Guide for example) with a set of pre-defined grid enabling statements so that the various programs are sent to the grid and SAS Grid Manager/LSF load balances them (see questions below about Enterprise Guide)

If you are scheduling SAS Web Report Studio reports with LSF and you have SAS Grid Manager (which would indicate multiple machines) then when the job gets triggered to run it will be directed to the most appropriate resource in the grid and would obey any of the policies that have been defined and implemented via LSF. SAS Grid Manager includes all the scheduling pieces so there would be no additional software required.

Q. What are the benefits of using grid-enabled Stored Processes?

A. A stored process is just a SAS program packaged in a specific way, so this usage would have the same benefits from grid that apply to any SAS programs. This would include:

- Workload balancing and flexibility
- The benefits of parallelization if the code in the stored process can be broken into parallel chunks of work. This would require modifying the Stored Process to insert the Grid and MPCONNECT stmts.
- From an enterprise perspective, you can establish policies to guarantee service levels so that your stored processes get access to the right resources and right number of resources during certain parts of the day compared to the other batch, Data Integration or other work that is also being submitted to the grid.

Q. Does SAS grid computing work with Enterprise Guide?

A. Yes, SAS programs that are generated by Enterprise Guide can be submitted to a grid to provide load balancing of jobs submitted by multiple EG users. This can be enabled to happen automatically by following the steps outlined in the document: [A How To Guide for Submitting Enterprise Guide Programs and Tasks to a SAS Grid.](#) This guide can be found at http://support.sas.com/rnd/scalability/grid/EG_on_grid.pdf

In order to accelerate a single Enterprise Guide program via parallelization, the program would have to be manually modified to add SAS/CONNECT and grid-specific statements to identify the sub-tasks to run in parallel on the grid.

Q. Does SAS Grid Computing work with Risk Dimensions?

A. If you are interested in running Risk Dimensions on a grid, the consultants in SAS Professional Services and the RiskAdvisory group have tools to enable this.

Q. Can Cron be used instead of LSF scheduler to schedule jobs to the grid?

A. No, Cron is a basic entry level OS scheduling capability available on a single host. If you need enterprise level scheduling capabilities involving load balancing over two or more machines (grid computing) then SAS Grid Manager, which includes the Platform Suite for SAS, is required.

Q. For the SAS Grid Logical Server definition, what if the SAS command differs based on machine? For example, if you have an MVS machine and two UNIX machines in your grid, each with a different command that starts SAS.

A. First of all, MVS is not a supported platform for the SAS Grid Manager environment. However, the command that is specified in the SAS Grid Logical Server definition is actually a script file (bat file for Windows and shell script for UNIX). So if you have different invocation commands on different hosts you simply modify the script file on each host to contain the host specific command. Do not use fully qualified path names for script file specified as the SAS command value unless the path exists on each of the nodes in the grid. For an environment that is a mixture of Windows and UNIX, do not specify an extension on the script file specified in the SAS Grid Logical Server definition.

Q. If SAS is installed in a 4way hard partition within a 16way machine in a grid, will SAS Grid Manager be able to limit the execution of SAS to just those 4 CPUs on that node?

A. SAS would only execute on the 4 CPUs in the partition - not because of SAS Grid Manager but because the partitioning makes it appear as a 4-way machine. There is no way for SAS to see the other 12 CPUs.

Q. Are Sun Solaris "containers" supported in a grid?

A. Yes – in SAS 9.1.3 support Solaris containers requires a hot fix for the Platform software – contact SAS Tech Support to obtain the hotfix. A large SMP box would be much more expensive than several smaller boxes – so using Solaris containers would make sense if the customer already has the box and wants to carve part of it out for the grid and use other containers for other apps.

Q. Can you run SAS applications with SAS Grid Manager in heterogeneous grid environments?

A. In general, yes. SAS programs that are being distributed in a heterogeneous environment (mixture of windows and UNIX) are more complex to configure and do have some restrictions. In such environments SAS:

- cannot share catalogs (including formats) or views across architectures;
- can share data sets across architecture with CEDA but you would want to put in the format that matches the largest number of machines to minimize the performance hit caused by conversion;
- has differing encodings for SAS datasets which affects the application of integrity constraints in SAS jobs;

- programs need to use a universal naming convention for filenames and libnames.
- has considerations for userid mapping;
- has considerations for shared vs. independent grid configurations.

Both Data Integration Studio and Enterprise Miner require homogenous grid environments. Data Integration Studio requires that the grid be either all UNIX or all Windows because it does not use universal names in the libnames/filenames that it generates. Enterprise Miner model training requires all the same OS on all of the grid nodes that it uses because it shares views from one node to the next.

Q. Can Windows 32 bit servers and Windows 64 bit servers be combined in one grid?

A. Yes this can be done. Since 32 bit and 64 bit data sets are stored in different formats CEDA will be used any time these servers access data created by the other and there will be translation done on the fly. There could be a performance implication depending on how often this happens.

Q. What is the level of granularity of the grid? Can groupings be created by processors/cores or just servers?

A. Creating groupings by processors/cores is something controlled by the operating system, not the grid middleware. Therefore, you can group processors/cores on a single SMP box if the OS supports some type of partitioning/containers/LPARs. You can assign a number of job slots to an SMP machine which defines the number of SAS processes SAS Grid Manager would run on this box simultaneously at any given time - however this does not map to specific CPUs.

SAS Grid Manager lets you group multiple boxes/servers either by defining resources in the SAS metadata and LSF configuration files or by defining multiple Logical Grid definitions under different application servers.

Q. Is there any type of reporting showing capacity and utilization?

A. The Host view in the Grid Manager plug-in in SAS Management Console will show this type of information. In addition, the LIM daemon that is part of LSF generates logs with this type of information. The best user information on the LIM and the resources information it collects is located in the "Administering Platform LSF" guide, the chapter on "Working with Resources". The complete Platform documentation is on the first CD of the "Platform Suite for SAS" media.

SAS metadata and security

Q. What grid-specific SAS metadata entities are created when you install and set up a SAS grid?

A. There are two definitions set up in the metadata when you install and configure a SAS grid. One is the Grid Monitoring Server Definition which has to do with communication between the Platform Suite for SAS and the Grid Manager plug-in for the SAS Management Console. The other is the Grid Logical Server Definition which is the "bridge" between a SAS application and the grid, and allows a SAS application to distribute to the resources in the grid.

The metadata related to characteristics of the individual machines in the grid is maintained by the Platform components and therefore completely removed from the SAS application.

Q. What access in the metadata does a user need in order to submit jobs to the SAS grid?

A. If you have read metadata access to the Grid Logical Server definition, then you can submit jobs to the grid. If you do not have access to the server context containing the Grid Logical Server then you won't be able to "see" the grid or submit jobs to it.

Q. Does every user need an account on every node on the grid, or is there a generic "grid server user" id involved?

A. A user submitting a job to the grid is submitting under his/her own credentials, so the user submitting the job will need to be defined on any node in the grid where his/her job may be executed. Note that for Data Integration Studio and Enterprise Miner, where the job submission goes through the Workspace Server. In this case, the SAS sessions on the grid will run with the same userid that the workspace server is running with, which is typically the userid of the client.

Q. How can you restrict certain users from running jobs on certain machines in the grid, or only allow certain machines to be used by certain groups, for certain purposes or at certain times of day?

A. To restrict groups of users to different sets machines in the grid, one best practice is to create different Grid Logical Server Definitions and restrict who can read the metadata. Depending on the complexity of the business rules involved, other scenarios may be accomplished by creating Platform LSF queues. See the SUGI31 paper by Cheryl Doninger and Alan Wong titled “SAS Goes Grid – Managing the Workload Across Your Enterprise” <http://www2.sas.com/proceedings/sugi31/211-31.pdf>

Q. Does LSF store user names and passwords? Does LSF encrypt data or user credentials, and if so what type and bit level of encryption?

A. Yes. See “Use of Credentials by SAS Grid Manager” in the “SAS Grid Manager Security Administration” document at <http://support.sas.com/rnd/scalability/grid/gridinstall.html>

Q. What user rights need to be granted to users in a Windows grid when using Platform LSF?

A. In a Windows grid, there are three different groups of users that need different rights. One group is the install user(s), one is the group of users that will run LSF services, and the third group are the users who will submit SAS jobs to the grid. For details see “LSF User Accounts Under Windows” in the “SAS Grid Manager Security Administration” document at <http://support.sas.com/rnd/scalability/grid/gridinstall.html>

Access to data in a grid

Q. Does data for grid enabled application need to be present on every node in the grid?

A. The data does not have to be stored on every node, but it does need to be accessible by every node. Several options are discussed in the FAQ questions that follow. Note also that input data is just one data issue when running in a grid. You also need to think about how extensively your application uses SASWORK, where to put the partial result sets generated by the processing done on each node, and where and how to re-aggregate those partial result sets.

Q. What options are there for making all input data accessible to all the nodes in a grid?

A. The options include Network Attached Storage (NAS) and Cluster File Systems (CFS). NAS is a server dedicated to file sharing. In an environment like a grid that consists of many servers, storage of data can be centralized, as can the security, management, and backup of the data. A Cluster File System (CFS) is a cluster of servers that all work together to provide high performance service to their clients. To the clients

the cluster is transparent. Generally speaking CFS is a faster but more expensive choice. Some examples (in no particular order) are:

Global File System (GFS) <http://sourceware.org/cluster/gfs/> (Linux)

Sun QFS

http://www.pcs.co.uk/catalogue/SUN_MICROSYSTEMS/Sun_Storage/Storage_Software/SunQFS.html

EMC Celerra Multipath File System (MPFS)

http://www.emc.com/products/software/celerra_mpfs/

IBM SAN FS <http://www-03.ibm.com/servers/storage/san/>

Veritas CFS

http://www.veritas.com/Products/www?c=keymessage&refId=145&psId=1786&ln=en_HK

PVFS <http://www.pvfs.org/pvfs2/> (UNIX and Windows?)

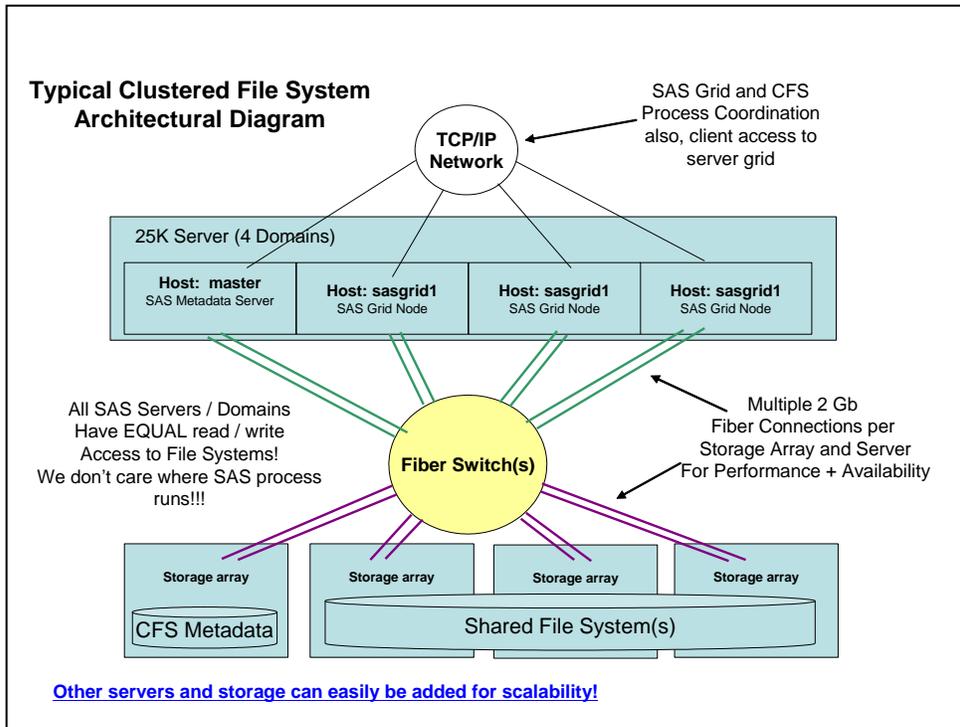
Microsoft (Distributed File System/MS Cluster Server – Windows)

http://en.wikipedia.org/wiki/Windows_Server_2003

Polyserve Matrix Server <http://www.polyserve.com/> (Linux and Windows)

Q. Which Cluster File Systems are supported by or have been tested with SAS Grid?

A. There is no known reason why the various Cluster File Systems would not work with SAS. R&D is pursuing work with different files systems as resources permit. Note this is a general storage-with-SAS issue, not a Grid-specific issue. Here is a link to a general article on Cluster File Systems and below is a diagram of a typical clustered file system architecture <http://www.informit.com/articles/article.asp?p=357689&seqNum=1&rl=1>



Platform LSF Questions

Q. What are the types of LSF scheduling algorithms that are available and how do you change from one to the other?

A. There is not a set of scheduling algorithms per se that you can use interchangeably in order to cause different load balancing behavior. You can however cause different scheduling behavior through queue definitions and other LSF configuration options. This is documented in the “Administering Platform LSF” guide in the “Scheduling Policies” chapter.

Q. If the machines in the grid are not busy and I always run 1 job, will the same machine always be picked?

A. In general, LSF will not pick the same machine every time a job is run.

Q. Is it possible to get the number of max job slots for a particular resource?

A. In SAS 9.2 and in hot fix E9BA05 for SAS 9.1.3SP4 the value returned by the `grdsvc_nnodes` function is the number of job slots per resource. In SAS 9.1.3SP3 and SAS 9.1.3SP4 the value returned by the `grdsvc_nnodes` function is the number of CPUs.

Q. The number of job slots per host defaults to the number of CPUs, is there a way to change that?

A. The number of job slots is set in the `lsb.hosts` file. You can change the default by applying Hotfix E9BA05, changing the `lsb.hosts` file then running the `reconfig` commands. In the `lsb.hosts` file under the `MXJ` heading, doing the following allocates 1 for hostA, 2 for hostB, 4 for hostC and the defaults for all others:

```
Begin Host
HOST_NAME MXJ r1m pg ls tmp DISPATCH_WINDOW # Keywords
#hostA 1 3.5/4.5 15/ 12/15 0 () # Example
#hostB 2 3.5 15/18 12/ 0/ (5:19:00-1:8:30 20:00-8:30)
#hostC 4 3.5/5 18 15 () () # Example
default ! () () () () # Example
End Host
```

Q. Does LSF source the profile when it runs each SAS session as a user?

A. No, however this can be accomplished by sourcing the profile in the `sasgrid.sh` shell script that is used to start SAS.

Q. Does LSF provide a way to notify an administrator that a group of servers on the grid was at a performance threshold?

A. The grid administrator can use the `lshosts`, `lsload` and `bhosts` commands to check load thresholds, current load indices and host status.

Q. Is there a way to automatically shift those resources around given some load/threshold?

A. LSF supports automatic job suspension and job migration. For example, one can configure a suspension threshold (for certain load indices) and a migration threshold (in

minutes). After certain load indices reach the threshold, the job will be suspended first, then after the suspension time reach the migration threshold, job will be automatically migrated to a lightly-loaded machine.

There are certain restrictions and requirements. Note that LSF does this without internal knowledge of the application being run by the job, and SAS jobs don't automatically suspend and resume easily.

For more details, you should refer to LSF 6.0 admin guide

Chapter 27: Load Thresholds
Chapter 23: Migrating jobs
Chapter 37: Tuning the cluster

Q. How does Platform LSF enable failover?

A. The LSF failover information can be found in the "Administering Platform LSF Version 6.0" guide Chapter 2 Failover p. 45. This document can be found on the first CD of the Platform Suite for SAS.

In addition, it's possible to specify which machines are candidates to become the master in the event of failure. If no `LSF_MASTER_LIST` is defined, when failover happens, any slave lim can become the master lim, and `sbd` will spawn `mbd` when it detects that it is now on the master host.

For better performance, you might want to specify `LSF_MASTER_LIST` to name a list of candidate master hosts. This can reduce the time it takes to reconfigure a cluster when the cluster becomes large. For details, refer to chapter 37, page 485.