

Effective Usage of SAS® Enterprise Guide® in a SAS® 9.4 Grid Manager Environment

Edoardo Riva, SAS Institute Inc., Cary, NC

ABSTRACT

With the introduction of new features in SAS® 9.4 Grid Manager, administrators of SAS solutions now have even better capabilities for effectively managing the use of SAS® Enterprise Guide® in a grid environment. In this paper, we explain and demonstrate proven practices for configuring the SAS 9.4 Grid Manager environment, leveraging grid options sets and grid-spawned SAS workspace servers. Further, we walk through the options provided by SAS Enterprise Guide that make the most effective use of the grid environment.

INTRODUCTION

In the past few years, a new trend has emerged in the IT community and is now on the frontline of every company. It is called big data. Many software vendors had to develop strategies and tools to process large volumes of data and analytic programs more efficiently. Here at SAS, we have been promoting SAS® Grid Computing for many years now, and it was built from the outset to process large amounts of data quickly for large numbers of concurrent users.

During the lifetime of SAS Grid Computing software, the interfaces to manage and execute parallel processing of data have evolved to become simpler to use and provide an increasing set of capabilities.

In SAS 9.4, there are new capabilities to make both a grid administrator's and a business user's lives easier. SAS Grid Manager has added grid options sets, a collection of options that can be associated with a specific client application for specific users and groups. The SAS object spawner has the capability to leverage the grid to control the starting of workspace servers. And SAS Enterprise Guide is able to seamlessly integrate with a grid environment so that business users can get faster results from the existing IT infrastructure without the requirement to modify their existing SAS Enterprise Guide projects.

This paper discusses all of the above and much more, to let all SAS Enterprise Guide users get the maximum from their SAS Grid Manager infrastructure.

THE BENEFITS OF A SAS GRID MANAGER ENVIRONMENT TO SUPPORT SAS ENTERPRISE GUIDE USERS

In SAS Grid Manager environments, SAS administrators have to set up and support a shared SAS® Analytics Platform for multiple SAS users running any number of SAS applications.

Luckily, behind the well-known benefits of increased availability, parallelized processing, and the ability to support a larger number of users within a single centralized environment, SAS Grid Manager provides built-in governance, which might bring the most benefit of all.

Many SAS Enterprise Guide users are used to self-service analytics run on their personal PCs and, when they are moved to a centralized shared environment, they might not see this as a great improvement. They are used to a great deal of independence on their SAS environment, unrestrained by the bonds of accountability. That is, until they are found to be the one responsible for consuming all the available resources and crashing a shared server. The governance features will definitely please the IT department, reassuring them that their servers are being used to their full potential, and that resources can be dynamically allocated according to ever-changing business priorities.

GRID OPTIONS SETS

In almost all SAS Grid Manager environments, a common solution to effectively manage resources is to set up separate configurations based on a mix of requirements for departments, client applications, and user roles. To accomplish this in previous SAS releases, administrators had to define multiple SAS Application Server contexts, each with its own Grid Server definition and associated options. We can see this type of setup in Figure 1, where the administrator created specific configurations for users of SAS® Data Integration Studio, SAS Enterprise Guide, and SAS® Enterprise Miner™.

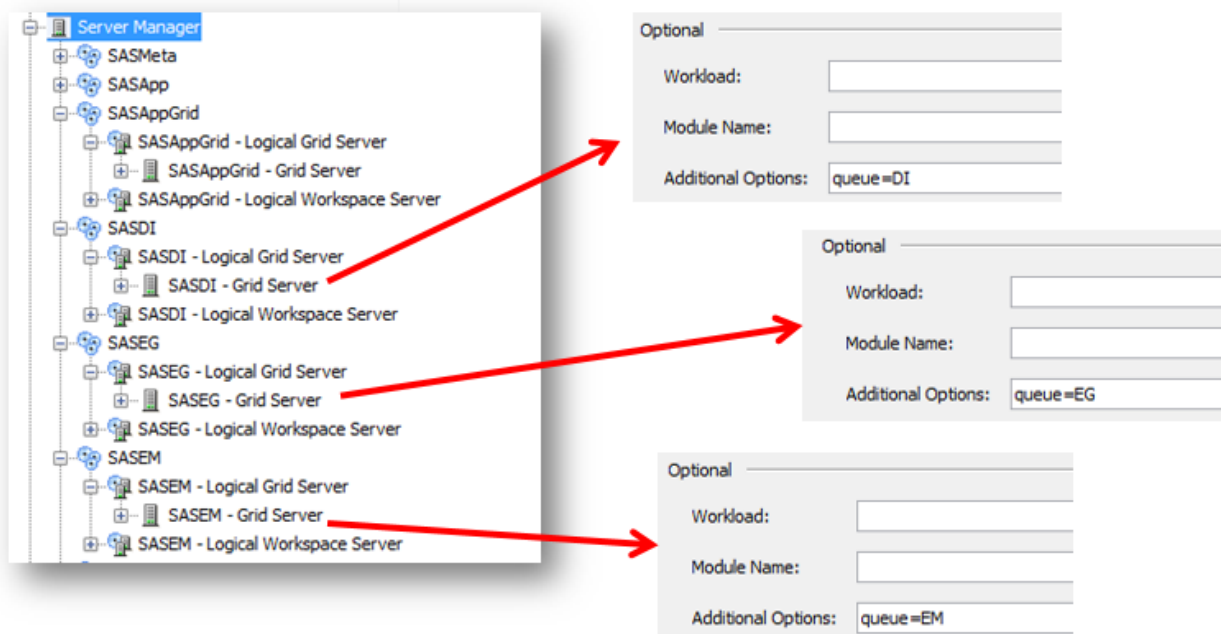


Figure 1 - Metadata structure before SAS 9.4

With this basic setup, it is up to the SAS Enterprise Guide users to choose the SASEG context to direct their jobs to a queue called “EG”; the same holds true for the other clients. If they choose the wrong server context, their session runs with the wrong settings. If an administrator wanted to limit access to a specific queue (for example, a high priority queue), he had to use metadata permissions to grant Read metadata permission on that context/server only to the desired users/groups. All the other users will not even see that grid server definition. Thus, they will not have access to that queue.

A drawback of this type of configuration is an unnecessary proliferation of SAS Application Servers that might be difficult to use and maintain.

SAS delivered the grid options set feature in SAS 9.4 grid computing to solve these issues.

A grid options set is a convenient way to name a collection of SAS system options, grid options, and required grid resources that are stored in metadata and are automatically applied when jobs are submitted to the grid. The good news is that most of SAS software is aware of this feature, so that the right set can be selected whenever an identified user accesses the client application and submits a job to the grid.

Options sets can be shared across different server contexts, client applications, or users, leaving maximum flexibility to SAS administrators. Two sample configuration scenarios are illustrated in Figure 2 and Figure 3.

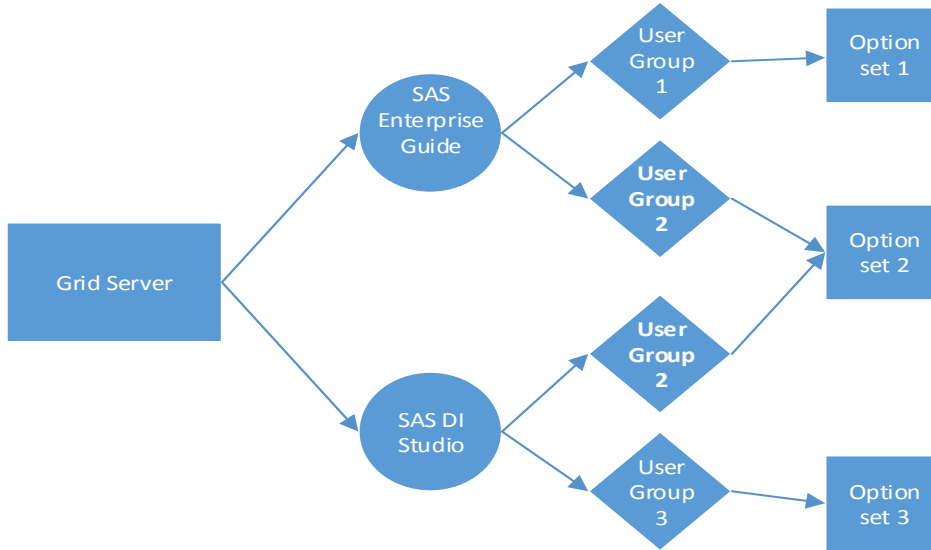


Figure 2 - Single server, multiple application scenario

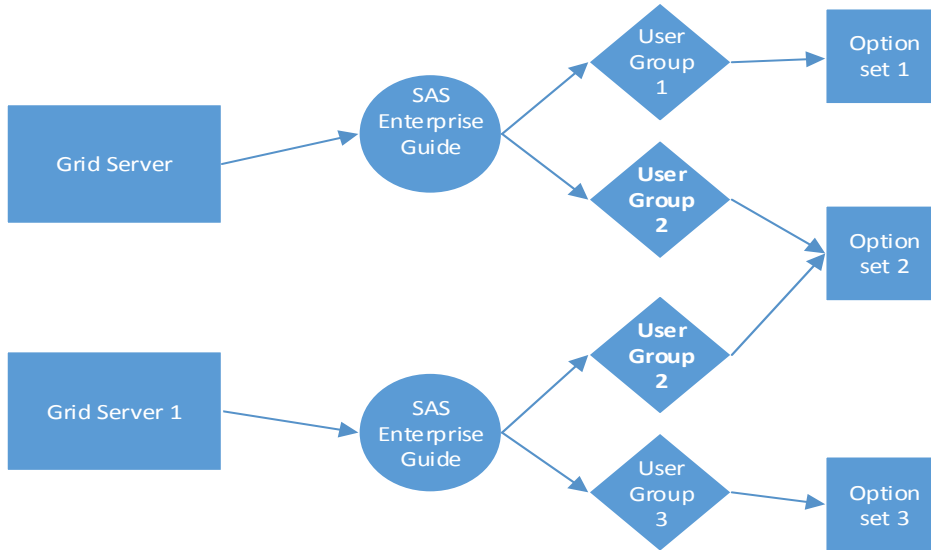
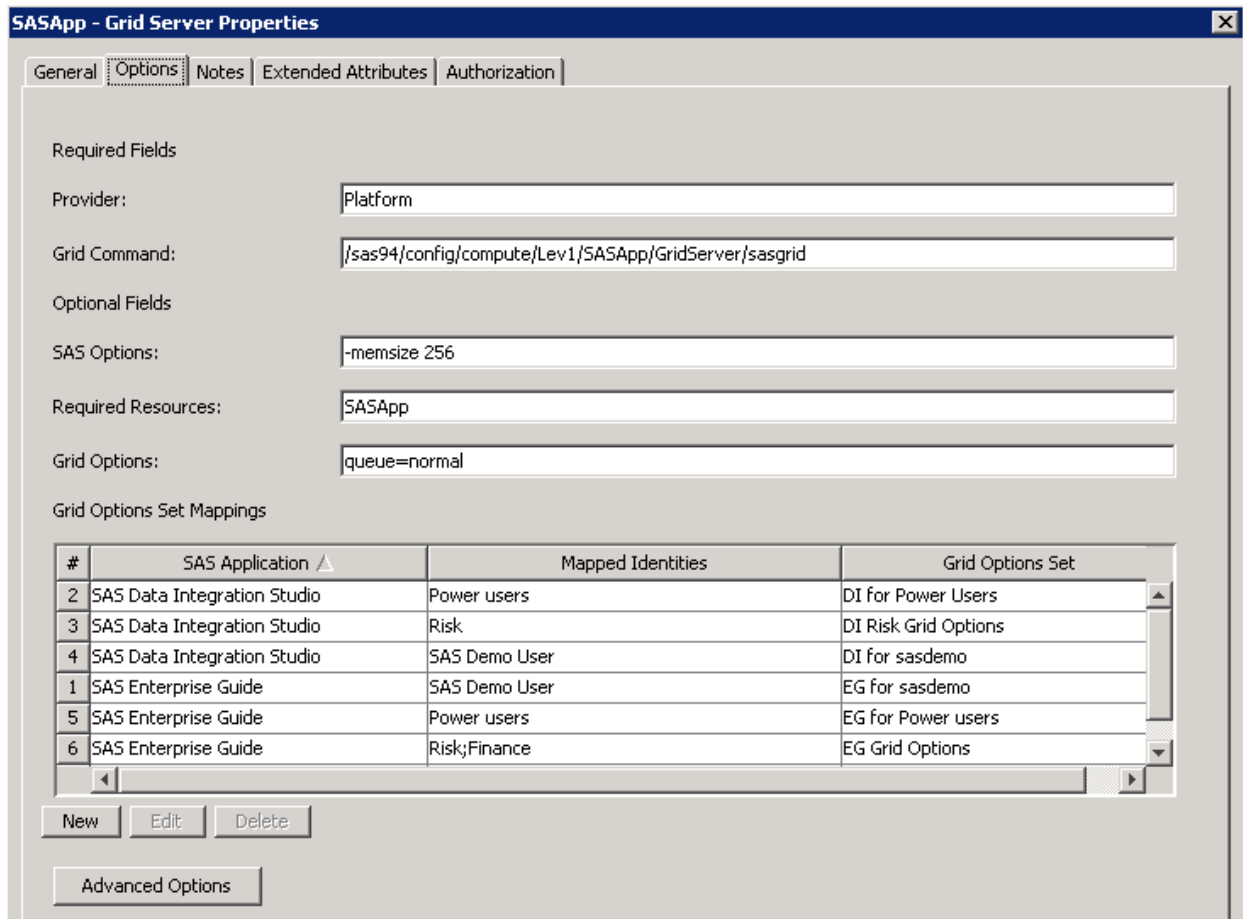


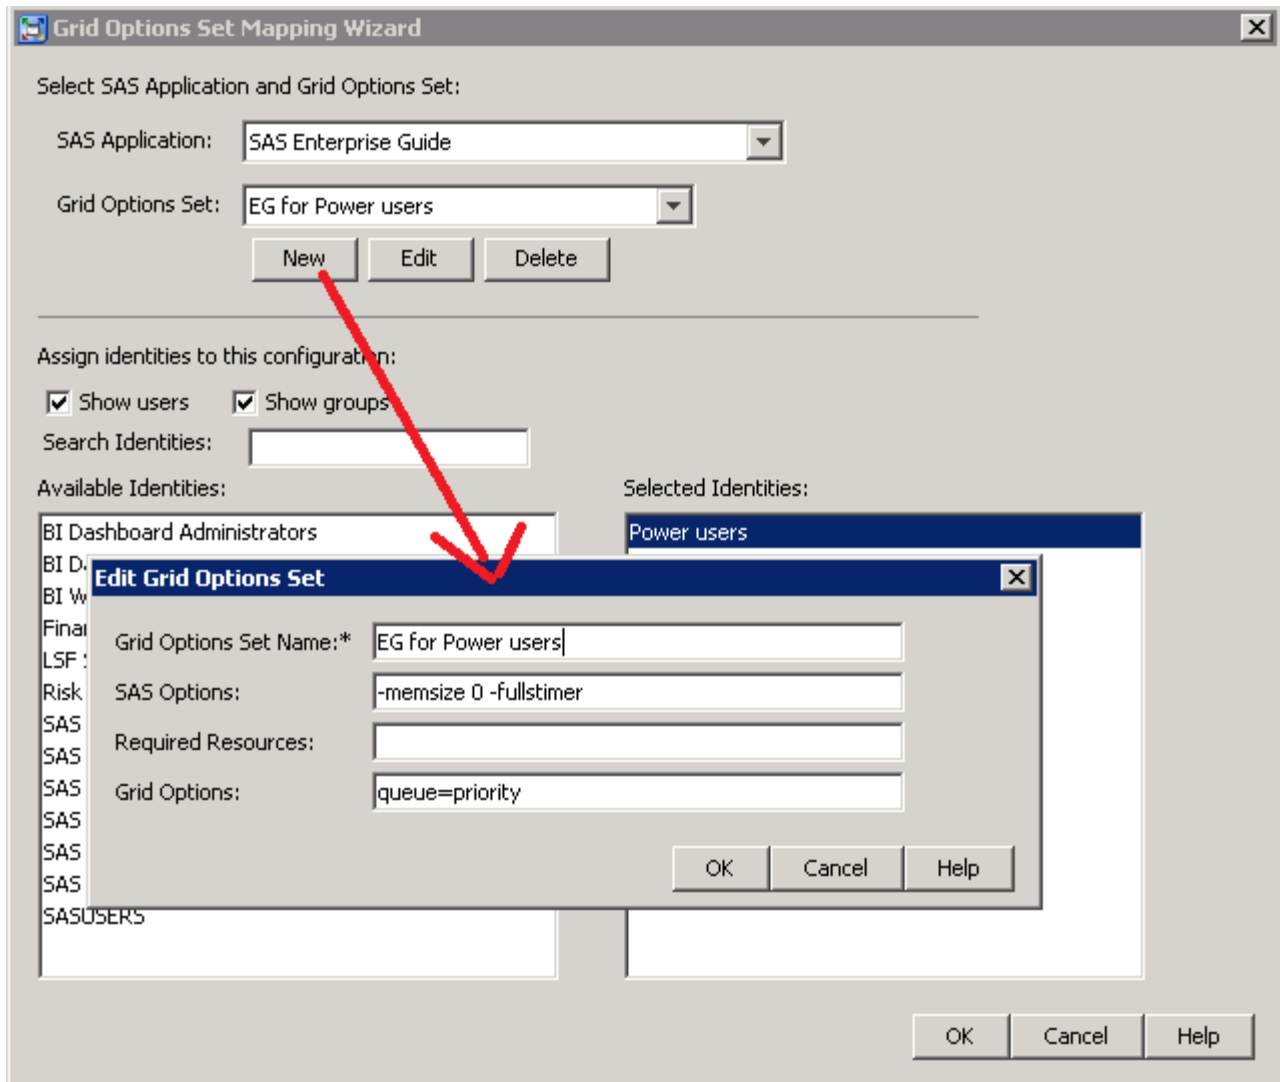
Figure 3 - Multiple server, single application scenario

The most common configuration with SAS Enterprise Guide is to define a dedicated options set for each grid queue (see next paragraph), as shown in Display 1.



Display 1 - A common configuration

When you open a specific grid options set you will find that ... it's not just about grid, it contains SAS options also. This means that you do not have to specify a new workspace server in metadata just to set a different memsize for some users - or any other option you might need. It is all here!



Display 2 - Options set for SAS Enterprise Guide power users

ONE GRID TO RULE THEM ALL

Explaining how SAS Grid Manager works and how to configure every detail is beyond the scope of this paper. Nevertheless, most organizations that use SAS consist of a variety of user categories and each one has its own needs and expectations. It is therefore important that your grid configuration fits the requirements of all of them.

Grid options sets tell SAS Enterprise Guide the options to use when starting a server session, but the detailed SAS Grid Manager configuration is stored in a repository outside of the SAS® Metadata Server.

In this paper we have chosen to show the usage of Platform RTM for SAS as the graphical interface to configure all grid parameters. Although not required, this tool greatly simplifies grid administration.

SAS Enterprise Guide users running interactive programs typically expect the results to be returned within a short amount of time. At present, the current out of the box grid options are set for long-running batch jobs, and put a latency of 20 seconds on the start of every server session. Therefore, the first tuning to perform is usually at the cluster level, to reduce grid services sleep times so that interactive session start faster. In Platform RTM, select *Config* ► *LSF* ► *Batch* Parameters and edit *MBD_SLEEP_TIME*, *SDB_SLEEP_TIME*, *MBD_REFRESH_TIME*, *JOB_SCHEDULING_INTERVAL*. Never put these values to 0. You should tailor the actual values to your grid (consider factors such as number of nodes, number of concurrent users, patterns of utilization ...), probably tuning them in multiple iterations. A starting point we use in our test environment is given below in Display 3.

LSF Configuration Edit [edit: default_1]

LSF Params Configuration

DEFAULT_QUEUE
 default job queue name.

MBD_SLEEP_TIME
 mbatchd scheduling interval (60 secs is default).

SBD_SLEEP_TIME
 sbatchd scheduling interval (30 secs is default).

Advanced Attributes
 Other LSF attributes

Key: Value:

Key: Value:

Display 3 - Cluster parameters

SAS Enterprise Guide, as well as a SAS® Add-In for Microsoft Office, is designed to keep the server session open for the full duration of the client session - unless a user explicitly chooses to disconnect from the server. In grid terms, this means that one job slot on that server is taken.

Therefore, whenever SAS Enterprise Guide is involved, you have to increase the number of job slots for each machine (*MXJ* parameter) from a default of 1 per core up to 5 or even 10 per core, depending on usage. This will increase the number of simultaneous SAS sessions on each grid node. Interactive workloads are usually sporadic, intermittent, with short cpu bursts followed by periods of inactivity when the user is reviewing the results or exploring the data. Because these jobs are not I/O- or compute-intensive like large batch jobs, more jobs can be safely run on each machine.

After this change, it is advisable to implement cpu utilization thresholds for each machine, to prevent servers from being overloaded. With this limit in place, even if many users submit cpu-intensive work at the same time, SAS Grid Manager can manage the workload by suspending some jobs and resuming them when resources are available.

Both changes are done at the host level; in RTM, select *Config* ► *LSF* ► *Batch Hosts* ► *default*, edit *Max Job Slots* value, and add the Advanced Attribute *ut*.

LSF Configuration Edit [edit: default]

LSF Hosts Configuration

Host Name *
 Official name of the host as returned by hostname.

Max Job Slots
 Number of job slots on the host.
 Use "!" to make the number of job slots equal to the number of Cores on a host.

Add Advanced Attributes

Key: Value:

Key: Value:

Advanced Attributes
 Other LSF attributes that are not built in Platform RTM.

Key: Value:

Key: Value:

Key: Value:

Key: Value:

Display 4 - Default batch host parameters

Even with this tuning, one user can easily use up all of the slots of a grid: all they have to do is to start many SAS Enterprise Guide sessions, or write code that uses all the available slots from a single SAS session. When a machine runs out of slots, it is “closed” and, when no machine has a free slot, no user can get another workspace. It does not matter that the user is not actually using resources; he might go for lunch, leaving his session open on a result page: no cpu, no I/O, nothing used on server.

The best way to prevent this is by creating a dedicated queue called EGDefault, with a *UJOB_LIMIT* parameter low enough (for example, 3). After that, each user will be then limited to 3 concurrent server sessions, whether started from the same client or from different SAS Enterprise Guide instances. When using SAS Enterprise Guide parallel features (they are explained later in the paper), the value of *UJOB_LIMIT* should be higher, provided that proper server sizing has been performed to accommodate for the additional resources required.

In RTM, you can create this queue selecting *Config* ► *LSF* ► *Queues* ► *Add*.

LSF Configuration Edit [edit: EGDefault]

LSF Queues Configuration

Queue Name *
The queue name must be specified; all other parameters are optional.

QJOB_LIMIT
Total number of job slots that this queue can use.

UJOB_LIMIT
Maximum number of job slots that each user can use in this queue.

Display 5 - EGDefault queue

At this point, all you have to do is to create a grid options set in SAS Management Console and add this EGDefault queue as a grid option to it.

There will always be ad hoc users or projects that do not fit into default categories (for example, because they might be running jobs that have a high priority or jobs that require a large number of computing resources). For users requiring higher priority for their jobs or require more computing resources, it is just a case of defining a new queue such as EGPower. It is common to limit access to this special queue to selected users, to prevent misuses. In previous releases, this would have been done by defining a special user group and then adding it to the *USERS* parameter in the queue definition. While effective, this has the disadvantage of duplicating user-related management both in metadata and in grid configuration files. With SAS 9.4, it is possible to apply metadata security to grid options sets to keep all in one place—that is, in metadata.

Finally, if you have other queues, dedicated for example to SAS® Data Integration Studio users or to batch processing, put job slot limits there, too, to compensate the large increase to the *Max Job Slots* parameter we made for default hosts. Display 6 shows the Advanced Attribute *PJOB_LIMIT* added to a batch queue, to enforce the limit of one batch job per physical core on every host.

Advanced Attributes
Other LSF attributes that are not built in Platform RTM.

[Add Advanced Attributes](#)

Key: Value:

Display 6 - Detail of the parameters for a batch queue

Cluster Parameters	Cluster Hosts	Resources	Queues	Batch Hosts	Batch Parameters	Batch Users	Resource Allocation	*Apply Settings to Cluster: sas_cluster
sas_cluster_default_2 (2) - sas_cluster (1) - LSF Configuration (Queues)								Add
Records: 30								go clear
<< Previous		Showing Rows 1 to 13 of 13 [1]					Next >>	
Queue Name	QJOB Limit	UJOB Limit	Priority	NICE	RERUNNABLE			
normal	default	default	30	20	default	<input type="checkbox"/>		
owners	default	default	43	10	default	<input type="checkbox"/>		
priority	default	default	43	10	default	<input type="checkbox"/>		
chkpnt_rerun_queue	default	default	40	10	YES	<input type="checkbox"/>		
unicodcmd	default	default	30	20	default	<input type="checkbox"/>		
EGDefault	default	3	30	0	NO	<input type="checkbox"/>		
DIDefault	default	5	30	0	NO	<input type="checkbox"/>		
batch	default	default	30	0	YES	<input type="checkbox"/>		
EGPower	default	1	40	0	NO	<input type="checkbox"/>		
<< Previous		Showing Rows 1 to 13 of 13 [1]					Next >>	

Display 7 - All queues at a glance in RTM

TAKING ADVANTAGE OF GRID-LAUNCHED SERVERS

One of the new features that came with SAS 9.4 is that SAS workspace server processes can be directly launched on the grid. With the first maintenance release for 9.4 (9.4 M1), this extends to stored process servers and pooled workspace servers, too, but would have required another paper to discuss the benefits and configuration for the user community.

What does this mean? Well, in a sentence, the object spawner submits the command to start a workspace server to SAS Grid Manager, which in turn decides which server best fits the request and directly starts the workspace server. This is something new because, in previous SAS releases, the object spawner delegates to SAS Grid Manager the responsibility to choose the best host on the grid, but then keeps the role to start the actual server process.

Table 1 summarizes load-balancing and server processes across releases.

Load balancing type	Who runs the load balancing algorithm	Who starts the server process	Availability
Any non-grid	Object spawner	Object spawner	All IOM servers, all SAS releases
Grid	SAS Grid Manager	Object spawner	All servers starting in 9.3
Grid + grid-launched server	SAS Grid Manager	SAS Grid Manager	Workspace Server only starting in 9.4 Stored Process Server, Pooled Workspace Server starting in 9.4M1

Table 1 - Load balancing who's who

UNDERSTANDING THE IMPACT OF SAS 9.4 GRID-LAUNCHED SERVER PROCESSES

One of the golden rules of SAS grid computing is that SAS Grid Manager has only control of workloads that it directly starts. Any other process is considered an external workload running on the same server. SAS Grid Manager can measure the cpu load and other resource metrics even for external workloads, but it cannot manage (start, suspend, kill, etc.) or even list them. Prior to SAS 9.4, technical architects had to balance between users' need to take full advantage of grid hardware and administrators' requirements to manage it. Let's see how this evolved across different SAS Enterprise Guide releases.

The right configuration for every release

Starting with SAS Enterprise Guide 4.2, pre-code and post-code execution scripts have been provided so that submitted code would run in the grid. See <http://support.sas.com/rnd/scalability/grid/download.html>. The pre-code and post-code execution scripts start a grid server behind the scenes and route the SAS code to the grid server for execution. This implementation allows a grid policy to be applied to a user's SAS code execution, but it comes at the expense of an underused workspace server acting as the grid server client. An administrator that wants to list server processes has to switch through different interfaces: the Server Manager plug-in for SAS Management Console monitors and controls workspace server sessions, while the Grid Manager plug-in does it for Grid Server sessions. This implementation requires planning and manual configuration for the end user to leverage the grid.


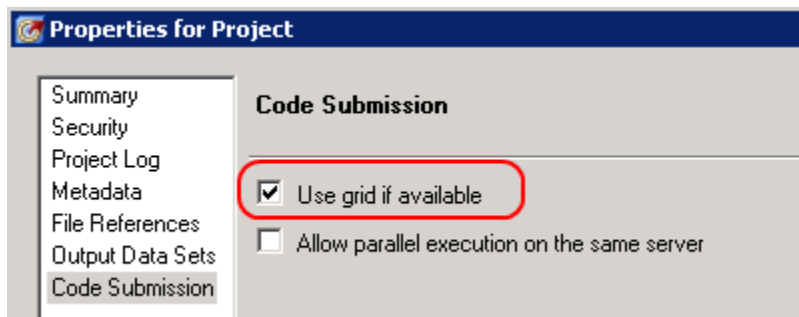


Figure 4 - Process interaction to run jobs on the grid with SAS Enterprise Guide 4.x and 5.1

SAS Enterprise Guide 5.1 solves end-users problems by automatically detecting the presence of a grid. In this case, it provides a simple check box in each project properties dialog box, or even in the properties of individual tasks, so that users can select to execute the code on the grid at the project level or at individual task level.



Display 8 - Using the grid at the project level

SAS Enterprise Guide 5.1 provides even more user friendly integration points, like a different server icon to highlight the connection to a grid server or additional information in the task status panel to display whether a specific task is running on the grid or inside a traditional workspace server.

As with SAS Enterprise Guide 4.x, 5.1 also uses pre-code and post-code files and each client session generates 2 processes at the OS level. The only improvement for administrators is the possibility to override the end-users dialog box shown in Display 8, through a metadata property, to always force the submission to the grid.

For further details about SAS Enterprise Guide 5.1 and grid integration, see the paper [Best Practices for Administering SAS® Enterprise Guide®](#).

SAS Enterprise Guide 6.1 is the first release to fully exploit SAS 9.4. After an administrator configures grid-launched workspace servers, every workspace server session required by SAS Enterprise Guide is started on the least busy node in the grid by SAS Grid Manager. It does not require additional client-side coding and can be managed and monitored as any other grid session from grid tools like RTM or the Grid Manager plug-in.

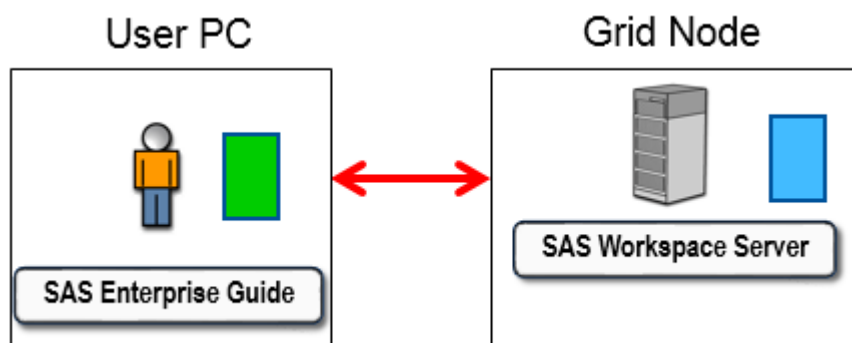
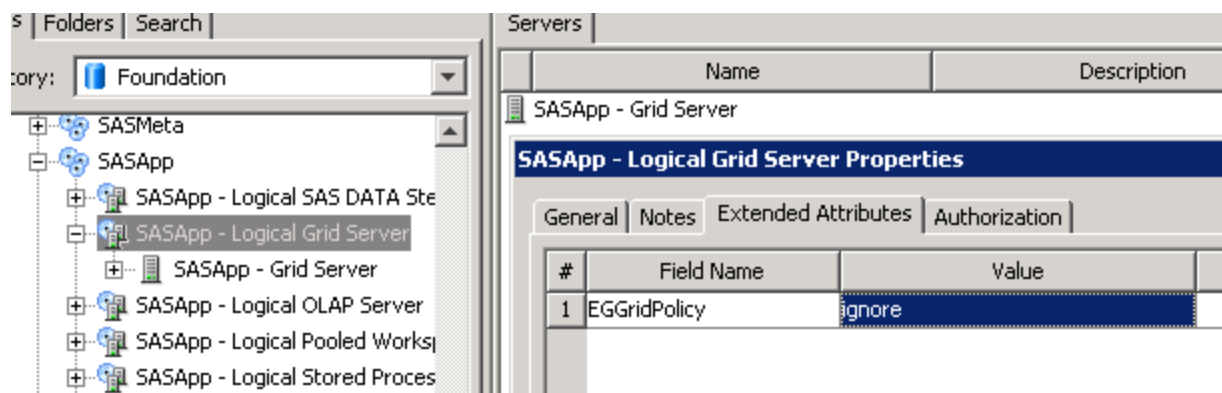


Figure 5 - Process interaction to run jobs on the grid starting with SAS Enterprise Guide 6.x

Even with this configuration, if an end user specifies the option 'Use grid if available', or selects 'Connect to grid', then a workspace server starts, performs a grid-enabled SIGNON, and submits code to the additional grid session. He would still end up with two SAS sessions running on the grid. To prevent this possibility and fully exploit the new grid-launched workspace servers, administrators should set the EGGridPolicy attribute on the logical grid server to a value of *ignore*. This attribute overrides any other settings in the application.



Display 9 - Preventing end users from connecting to the grid 'the old way'

Advantages of grid-launched workspace servers

To recap the advantages of using SAS Grid Manager to launch workspace servers for SAS Enterprise Guide clients, with this new feature you gain the following:

- The ability to apply grid policies to each session.
- Better resource utilization, by eliminating an additional grid server process.
- A simpler environment, with no GRIDWORK and RMTWORK libraries in SAS Enterprise Guide: the only temporary library that is used is the traditional SASWORK.
- Less confusion when trying to monitor the user's SAS process: now the same SAS process can be viewed from both SAS Management Console and Platform RTM for SAS, or any other grid monitoring tool.

```

JOBID  USER  STAT  QUEUE  FROM_HOST  EXEC_HOST  JOB_NAME  SUBMIT_TIME
537    sasdmo  RUN    EGDefault  sasserver02  sasserver04  SASAPP - WORKSPACE SERVER_4F71F000-8547-11E3-BF09-0050568240DB  Jan 24 17:31
538    sasdmo  RUN    EGDefault  sasserver02  sasserver02  SASAPP - WORKSPACE SERVER_5AFD0C00-8547-11E3-BF09-0050568240DB  Jan 24 17:32
539    sasdmo  RUN    EGDefault  sasserver02  sasserver03  SASAPP - WORKSPACE SERVER_6B86C000-8547-11E3-BF09-0050568240DB  Jan 24 17:32
542    sasdmo  PEND  EGDefault  sasserver04  -            SASAPP - WORKSPACE SERVER_C2B59000-8547-11E3-93A0-005056825AD0  Jan 24 17:35
QUEUE_NAME  PRIO  STATUS  MAX_JL/U  JL/P  JL/H  NJOBS  PEND  RUN  SUSP
normal      30    Open:Active  -      -    -    0      0    0    0
EGDefault   30    Open:Active  -      3    -    4      1    3    0

```

Figure 6 - Monitoring workspace servers through the usage of command-line grid tools

Is one downside of grid-launched workspace servers that, if all grid nodes are closed because they are busy or they ran out of job slots, no workspace server will start? This is different from SAS 9.3, where nodes just needed to be in any state other than close/admin. Actually, this might not be a downside: in a managed environment, why would we let another session start on a node that is already overloaded? If this is really what an administrator wants, then he would better reconfigure grid polices and increase the number of permitted sessions. Alternatively, he could empower selected end-users by letting them switch to a high priority queue.

A final consideration on the topic of grid-launched workspace servers: if you have different SAS solutions other than SAS Enterprise Guide leveraging the grid environment, take some time testing whether it is able to take advantage of this new feature. In some particular cases, the overhead imposed by the grid on starting each workspace server session might interfere with how the solution is designed and result in undesired time increase. In these situations, it might be enough to define in metadata multiple application server contexts and configure the instance leveraged by the solution as non-balanced or load balanced with a different algorithm.

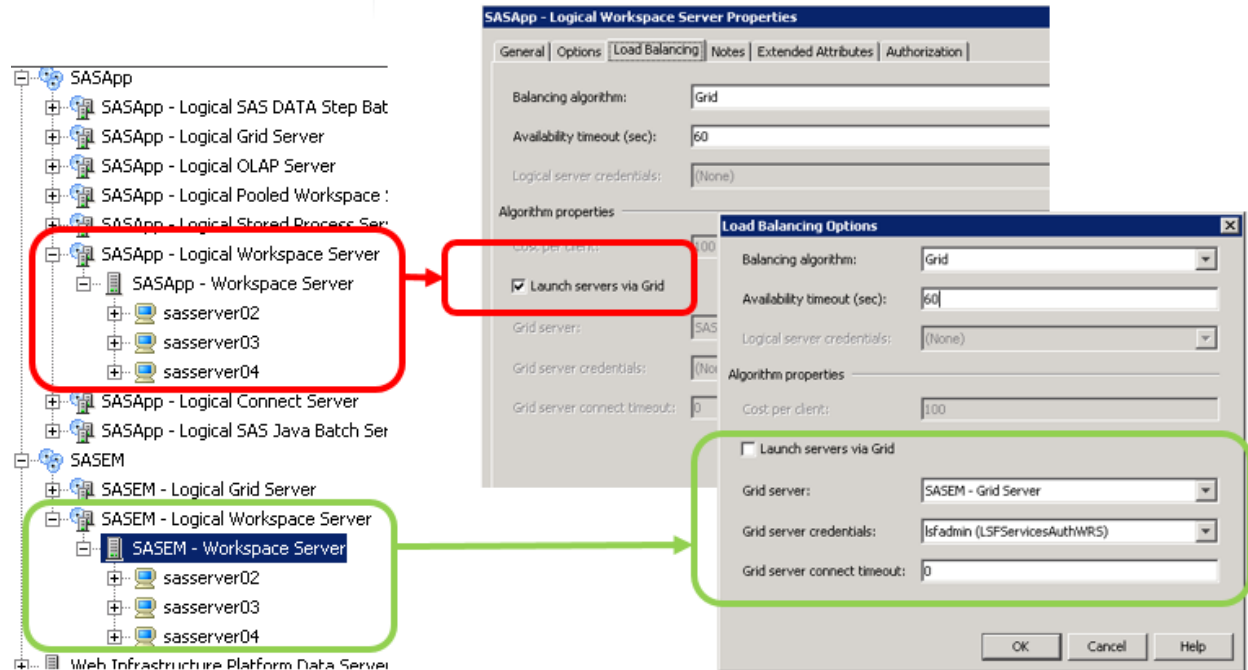


Figure 7 - Two different application server contexts dedicated to specific applications

SAS ENTERPRISE GUIDE PARALLEL PROCESSING

SAS Enterprise Guide 5.1 introduced the possibility to start multiple workspace servers to run several tasks in parallel. This is not limited to grid environments; it works even on a single server. It can happen with any of the two features:

- Parallel execution of process flows
- Data Exploration

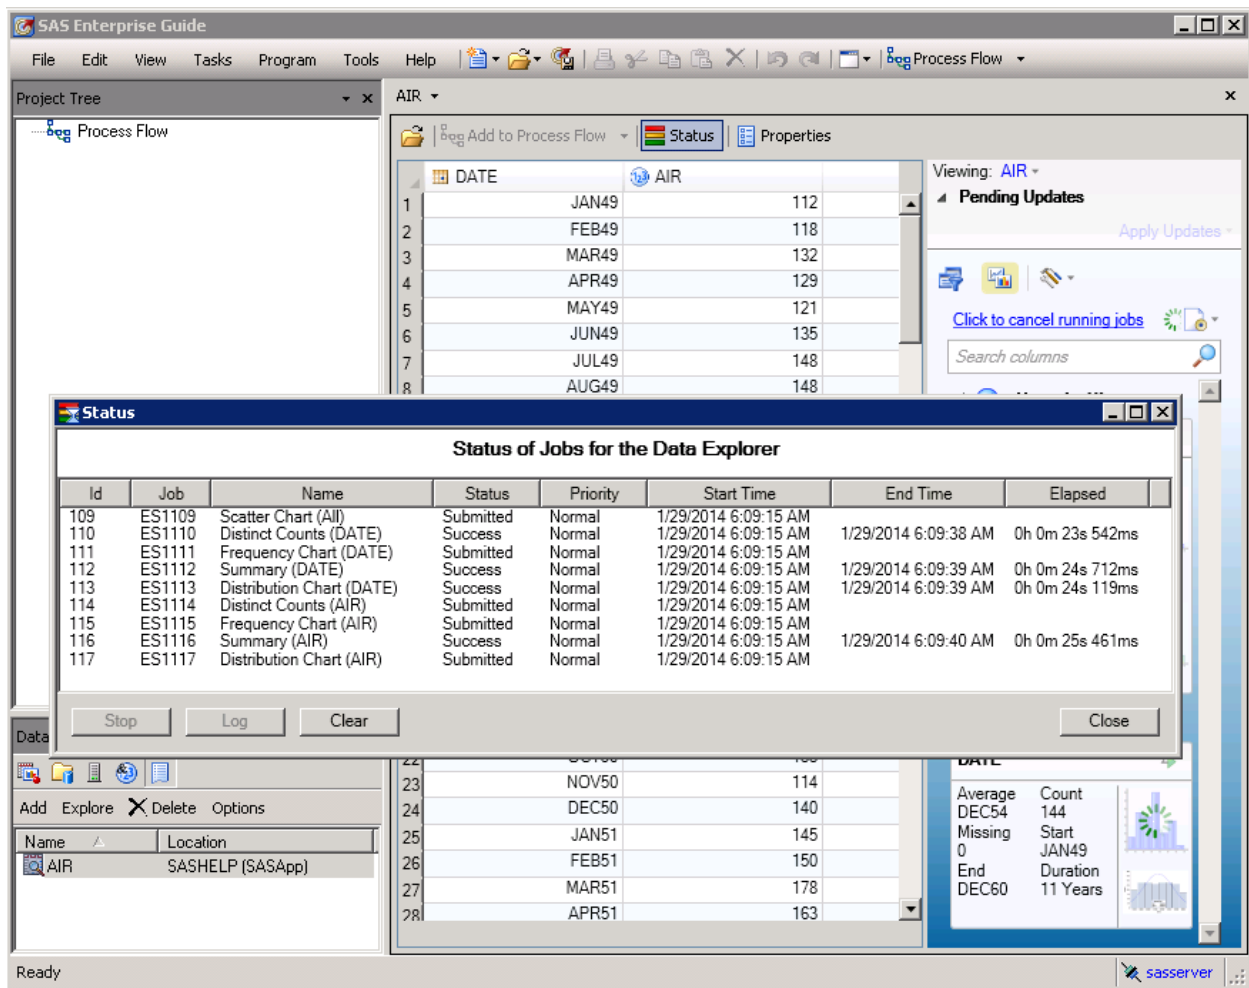
End users can enable the parallel execution of process flows. They do this by selecting the *Allow parallel execution on the same server* option. This can be set at a project level on the project properties dialog box, and can be overridden at a task level using the task properties dialog box. SAS Enterprise Guide determines when it is possible to run tasks in parallel (if there are parallel paths in the process flow), and starts additional workspace servers to run the tasks.

The screenshot displays the SAS Enterprise Guide interface with a process flow diagram and a task status table. The process flow diagram shows a sequence of tasks: 1.1, 1.2, 2.1, 2.2a, 2.2b, 2.3, 3.1, 3.2, 4, 5a, and 5b. Tasks 1.1, 2.1, 2.2a, 2.2b, 3.1, and 3.2 are shown as independent branches that can be executed in parallel. The task status table below the diagram provides a detailed view of the execution progress for each task.

Task	Status	Server	Server Type	Queue
1.1	Running DATA step...	SASApp	Workspace	
3.1	Running...	SASApp	Workspace	
1.2	Waiting on [1.1]...	SASApp	Workspace	
2.2a	Running DATA step...	SASApp	Workspace	
2.2b	Connecting to server...	SASApp	Workspace	
3.2	Waiting on [3.1]...	SASApp	Workspace	
2.3	Waiting on [2.2a, 2.2b]...	SASApp	Workspace	
4	Waiting on [1.2, 2.3, 3.2]...	SASApp	Workspace	
5a	Waiting on [4]...	SASApp	Workspace	

Display 10 - Parallel execution of independent branches

Data Exploration is started by either selecting *File ► Open ► Data Exploration*, or by clicking the Add button at the top of the Data Exploration History panel. Whenever a user runs a task, like requesting statistics in the Quick Start panel, several workspace server sessions are started in parallel to service the requested jobs. Users can monitor and manage these jobs from the Status panel.



Display 11 - Monitoring the status of parallel jobs in Data Explorer

CONFIGURING PARALLEL PROCESSING IN SAS GRID MANAGER ENVIRONMENTS

Exploiting parallel capabilities might result in improved time to value and a better user experience. Yet, when used in a grid environment, SAS Enterprise Guide has to be properly tuned to avoid unexpected behavior.

Let us see with a practical example.

An administrator, following the practice suggested earlier, creates an EGDefault queue with a limit of 3 concurrent slots per user, with the assumption that no one should be using more than 3 concurrent SAS Enterprise Guide instances at the same time. Then he configures workspace servers to be load balanced and selects the option to launch servers via grid.

A smart user reads one of the SAS Global Forum papers referenced at the end of this paper and decides to run some projects in parallel, but gets strange results. He starts with a simple project and he is happy to see two streams running in parallel. As soon as SAS Enterprise Guide tries to start more workspace servers, all tasks keep waiting in a 'Connecting to server' status until a timeout is reached. Finally, only two of them run.

Task Status					
Task	Status	Queue	Server	Server Type	
Setinit	Connecting to server...		SASApp	Workspace	
Long DataStep1	Connecting to server...		SASApp	Workspace	
Long DataStep2	Connecting to server...		SASApp	Workspace	

Display 12 - Parallel tasks waiting to be submitted

Understanding how SAS Enterprise Guide starts and manages workspace servers is fundamental to avoid this undesired behavior.

Whenever one of the parallel features are requested, SAS Enterprise Guide creates a client-managed pool, containing as many workspace server sessions as needed to perform processing for that session, plus one. The first workspace server that is started is the main one the user interacts with, while all the other ones receive and execute the parallel tasks. After a certain amount of inactivity, SAS Enterprise Guide closes all sessions but the first one.

This client-managed pool is different from the server-managed Pooled Workspace Server, and contains no specific code to make it grid-enabled. Yet, when an administrator configures grid-launched workspace servers, all the members of the pool are automatically launched on the grid and consequently become subject to grid policies.

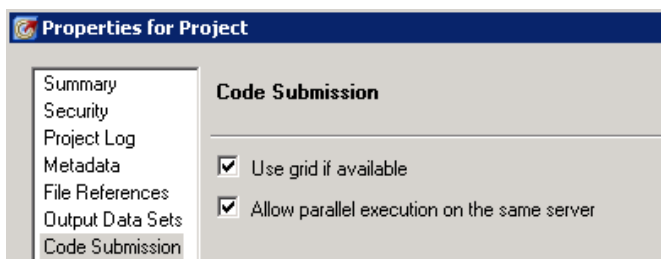
The number of simultaneous workspace sessions that SAS Enterprise Guide can start, as well as other parameters, are controlled by a client configuration file, but grid policies set on the back end always have precedence.

We can now understand that, whenever a limit is enforced on the grid, it is required to configure the same or a lower limit in SAS Enterprise Guide, too, to obtain proper functionality.

The steps are similar to the ones described in [SAS Note 51225](#).

- Open the configuration file `%appdata%\SAS\SharedSettings\6.1\Engine\SystemSettings.xml`
- Modify the parameters `WSPoolSize`, `WSNormalLimit`, and `Userjoblimit` so that they are equal to or lower than the limit imposed on the grid, minus one (one grid job slot is always needed by the master session of the client pool). For example, if a grid queue is limiting 6 job slots per user, set these values to 5 or less.
- If the grid experiences a continuous shortage of job slots, modify the parameter `WSLongLife`. This is the time in seconds that a workspace will be allowed to be idle before being ended. A running session counts toward the total number of used job slots on the grid, even if it is idle. The sooner it stops, the sooner it will release its job slot, making it available for other processes. Conversely, if an idle server is closed too soon and after a short time it is needed again, the user experience will suffer for the delay to continuously open and close sessions. A good balance could be around 10 minutes (600 seconds).

ADDITIONAL PARALLEL PROCESSING IMPLICATIONS



Display 13 - Selecting both code submission options

We have seen that the option *Allow parallel execution on the same server* does not enable processing on the grid, but it can affect how grid-enabled jobs are processed. If a user selects both this option and the *Use grid if available* option, then multiple workspace servers are started. Each server just acts a wrapper to sign on to the grid and submit to the grid each parallel branch in a project.

Figure 8 and Figure 9 show what happens when both options are selected and a project with just three parallel branches is executed. If you find it difficult to understand which session is parent of which other, you are not alone: the environment easily becomes too difficult to monitor, manage, and tune.




Figure 8 - Multiple servers launched by a single client

JOBID	USER	STAT	QUEUE	FROM_HOST	EXEC_HOST	JOB_NAME	SUBMIT_TIME
770	sasdemo	RUN	normal	sasserver02	sasserver04	SASAPP - WORKSPACE	SERVER_EEC46C00-F3D2-11E2-8E59-005056824E1B Jul 23 16:03
771	sasdemo	RUN	normal	sasserver02	sasserver02	SASAPP - WORKSPACE	SERVER_F6258C00-F3D2-11E2-8E59-005056824E1B Jul 23 16:03
772	sasdemo	RUN	normal	sasserver02	sasserver04	SASAPP - WORKSPACE	SERVER_FC555000-F3D2-11E2-8E59-005056824E1B Jul 23 16:03
773	sasdemo	RUN	normal	sasserver02	sasserver02	SASAPP - WORKSPACE	SERVER_03AAD800-F3D3-11E2-8E59-005056824E1B Jul 23 16:04
774	sasdemo	RUN	normal	sasserver04	sasserver01	EG.sasdemo.23822	Jul 23 16:04
775	sasdemo	RUN	normal	sasserver02	sasserver03	EG.sasdemo.9932	Jul 23 16:04
776	sasdemo	PEND	normal	sasserver02	-	EG.sasdemo.10181	Jul 23 16:04

Figure 9 - Trying to monitor the processes on the grid

Selecting both options might cause additional problems, because each workspace server assigns a GRIDWORK library and an RMTWORK library to synchronize code execution with the grid session that it started. When SAS Enterprise Guide tries to read from one of those libraries, the result might be unpredictable.

It should be clear now why, if you are using grid-launched workspace servers, it is advisable to avoid enabling the *Use Grid if available* option: all workspace servers are already running on the grid. This is definitely a situation where an administrator would set the *EGGridPolicy* attribute on the logical grid server to a value of *ignore*, as we have shown earlier in this paper.

CODE ANALYSIS

SAS Enterprise Guide has the capability to analyze legacy SAS code and refactor it. There are different papers that describe the types of code analysis available from the *Program ► Analyze* menu, including *I Didn't Know SAS® Enterprise Guide® Could Do That!*

Which type of analyses is best suited to produce code that will run in a grid?

Despite the name, *Analyze for Grid Computing* is not always the right choice. As in many cases, the answer is that it depends on what is the objective of our analysis.

Let us have a look at what happens to some sample code that we want to refactor:

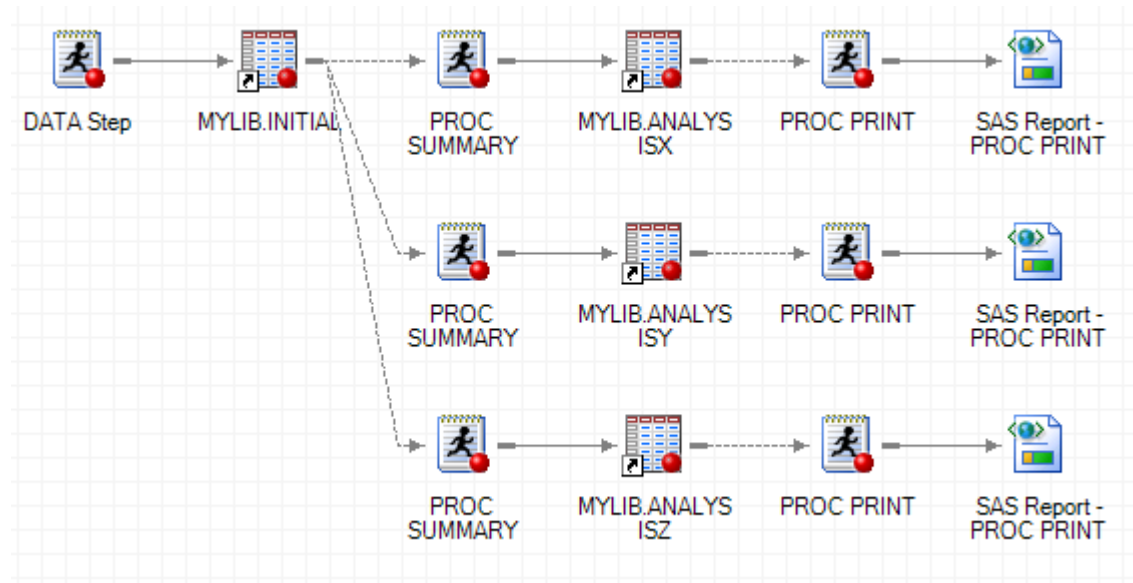
```

data mylib.initial;
  input x y z ;
  cards;
1 2 3
...
...
4 5 6
run;
proc summary data=mylib.initial;
  var x;
  output out=mylib.analysisX mean=meanx;
run;
proc print data=mylib.analysisX; run;
proc summary data=mylib.initial;
  var y;
  output out=mylib.analysisY mean=meany;
run;
proc print data=mylib.analysisY; run;
proc summary data=mylib.initial;
  var z;
  output out=mylib.analysisZ mean=meanz;
run;
proc print data=mylib.analysisZ; run;

```

Analyze for Program Flow

The Analyze for Program Flow function analyzes the program code and divides it into independent program nodes, then it inserts them into a new process flow. The previous code becomes similar to Display 14.

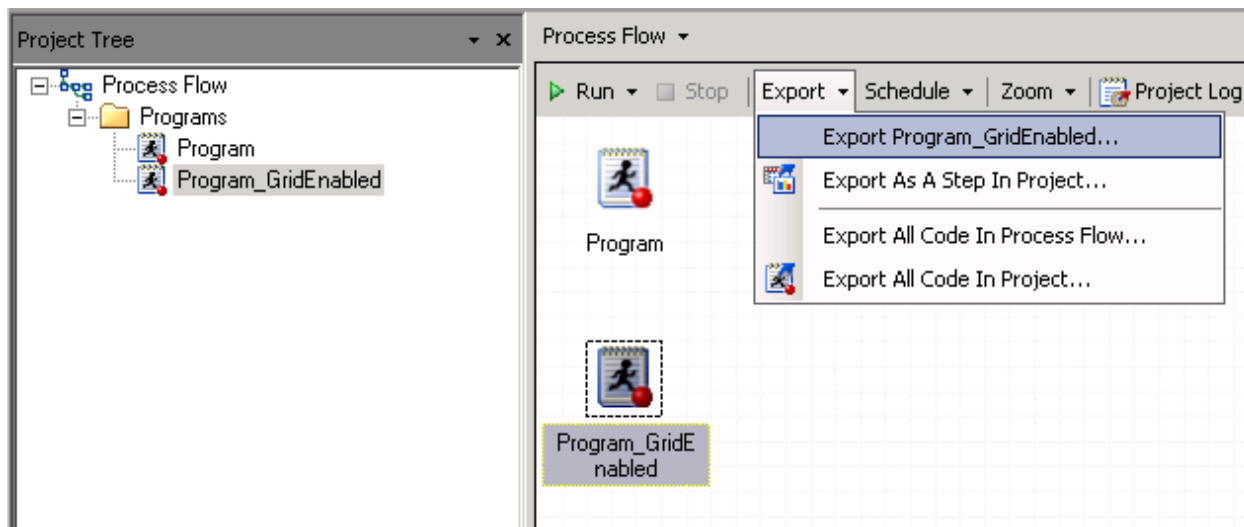


Display 14 - The result of the analyzed and imported code

We can see that the flow has three independent parallel branches. When using grid-launched workspace servers, just selecting the *Allow parallel execution on the same server* option will run our legacy code, in parallel, on the grid, from an interactive session. Additional workspace server connections are made to accomplish this, so the original code has to be as stateless as possible. All libraries should be defined in metadata or assigned in autoexec files; if the code sets a macro variable, it might not be available for a downstream task to use since it could run on another workspace instance.

Analyze for Grid Computing

The Analyze for Grid Computing function modifies the original program code inserting appropriate grid set-up macros into it. After the analysis completes, you can save the grid enabled code to a new code node by clicking the Add to project button. The code can then be exported as new program (for example, to run it in batch).



Display 15 - The result of an analysis for grid computing, ready to be exported

The generated grid-enabled code starts multiple grid sessions for parallel processing, not additional workspace servers. This has several implications; the most important is that all jobs will be submitted without enforcing any grid options set defined in metadata, and will run in the default grid queue. This can be fine in situations where the default queue is configured with parameters suitable for batch processing, and we actually submit the resulting program in batch.

JOBID	USER	STAT	QUEUE	FROM_HOST	EXEC_HOST	JOB_NAME	SUBMIT_TIME
952	sasdemo	RUN	EGDefault	sasserver02	sasserver03	SASAPP - WORKSPACE_SERVER_93D	
956	sasdemo	RUN	normal	sasserver03	sasserver02	SASGrid:1262	Jan 29 20:49
957	sasdemo	RUN	normal	sasserver03	sasserver01	SASGrid:1262	Jan 29 20:49
958	sasdemo	RUN	normal	sasserver03	sasserver03	SASGrid:1262	Jan 29 20:49

Figure 10 - Monitoring the sessions started from the generated grid-enabled code

CONCLUSION

SAS Enterprise Guide users can benefit from many features provided by SAS Grid Manager. Administrators can get a better control on the server environment while providing end users the required resources. In this paper we have shown how easily the latest SAS Grid Manager enhancements can be leveraged by SAS Enterprise Guide users and administrators, to provide them all with ...THE POWER TO KNOW®.

REFERENCES

- Allemang, Mark. 2014. "I Didn't Know SAS® Enterprise Guide® Could Do That!" When they are available, see the *Proceedings of the SAS Global 2014 Conference*. Cary, NC: SAS Institute Inc.
- Haigh, Doug. 2013. "The Top Four User-Requested Grid Features Delivered with SAS® Grid Manager 9.4". *Proceedings of the SAS Global 2013 Conference*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/resources/papers/proceedings13/470-2013.pdf>
- Smith, Casey. 2012. "Best Practices for Administering SAS® Enterprise Guide®." *Proceedings of the SAS Global 2012 Conference*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/resources/papers/proceedings12/297-2012.pdf>
- SAS Institute Inc. 2013. *Grid Computing in SAS® 9.4, Second Edition*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/documentation>
- Platform Computing Corporation. 2011. *Platform LSF Configuration Reference Version 8.01*. Available at <http://support.sas.com/rnd/scalability/platform>
- SAS Institute Inc. 2013. SAS Note 51225. "How to set the number of parallel processes that SAS® Enterprise Guide® can execute." Available at <http://support.sas.com/kb/51/225.html>

RECOMMENDED READING

- Granger, Erwan. 2012. "RTM and SASGSUB, the Power to Know®... what your Grid is doing." *Proceedings of the SAS Global 2012 Conference*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/resources/papers/proceedings12/370-2012.pdf>
- SAS Institute Inc. 2013. SAS® Intelligence Platform: Desktop Application Administration Guide, Second Edition. Cary, NC. SAS Institute Inc.
- SAS Grid Computing focus area at <http://support.sas.com/rnd/scalability/grid/index.html>

CONTACT INFORMATION

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

Edoardo Riva
100 SAS Campus Drive
Cary, NC 27513
SAS Institute Inc.
edoardo.riva@sas.com
<http://www.sas.com>

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

Other brand and product names are trademarks of their respective companies.