

## SAS Grid: Grid Enablement of SAS

Adam H. Diaz, IBM Platform Computing, Research Triangle Park, NC

### ABSTRACT

This will be a short discussion of the critical decision making factors for the use of SAS Grid. Many factors including reliability, scalability and performance of an overall SAS infrastructure can have a large influence on the bottom line. The talk will deal with SAS Grid for SAS 9.3. The audience only needs a high level understanding of SAS for this talk.

### INTRODUCTION

The intent of this paper is to simply provide a discussion framework for the benefits of SAS Grid. There are many factors that influence when it is time to begin consideration of an enterprise level installation of SAS. This includes the need for a solution that is scalable, powerful and addresses the needs of not only the individual SAS users but also the entire business their work supports. This paper focuses on introducing SAS Grid and its place in the SAS ecosystem as well as the warning signs consideration of this technology might be warranted.

### WHAT IS SAS GRID?

SAS Grid is a product from SAS that “grid enables” SAS deployments. This means that this product enables an enterprise infrastructure to distribute the compute portion of SAS work to a networked collection of computers. While SAS Grid has a variety of methods for job submission, the user experience from an interface perspective for most products remains very similar in the context of SAS Grid. For Enterprise Guide, Data Integration Studio or Enterprise Miner the change in workflow is one of some initial configuration during product setup after which the use of the product is very similar. SAS Grid then is a way to allow many users access to a collection of compute resources and SAS applications while providing workload management, high availability and performance benefits. SAS Grid also provides a highly valued ease of expansion due to its scalable architecture. This feature allows businesses to easily expand their investment in the value of the analytics produced by SAS code in an affordable and metered way. This single benefit is well received by system administrators and CFOs alike. Expanding a SAS Grid doesn’t require reinstallation of SAS but rather an incremental investment in additional hardware and software resources.

### WHAT IS HIGH PERFORMANCE ANALYTICS (HPA)?

SAS Grid is a part of strategic initiative at SAS called HPA which stands for High Performance Analytics. HPA contains three SAS products that are commonly called the “pillars” of the HPA platform. These products include SAS Grid, SAS In-Database and SAS In-Memory solutions. SAS Grid is considered the first pillar in the three part HPA platform. SAS Grid being the oldest and most mature product is also the most logical and common entry point for most all users of SAS software across the board when considering a SAS enterprise level architecture. This is because it enables parallelization, high availability, workload balancing and general performance improvements for most SAS code from Base SAS through a whole host of more robust SAS offerings.

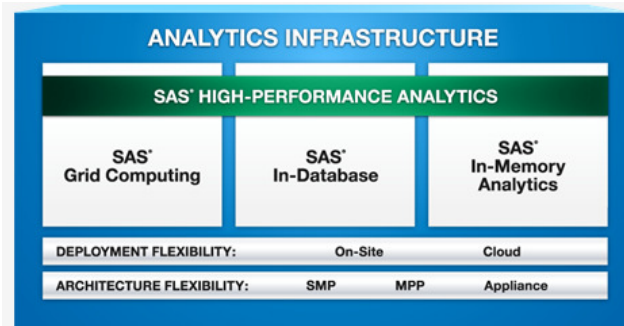


Figure 1: Graphical Representation of HPA

## WHAT IS PLATFORM SUITE FOR SAS?

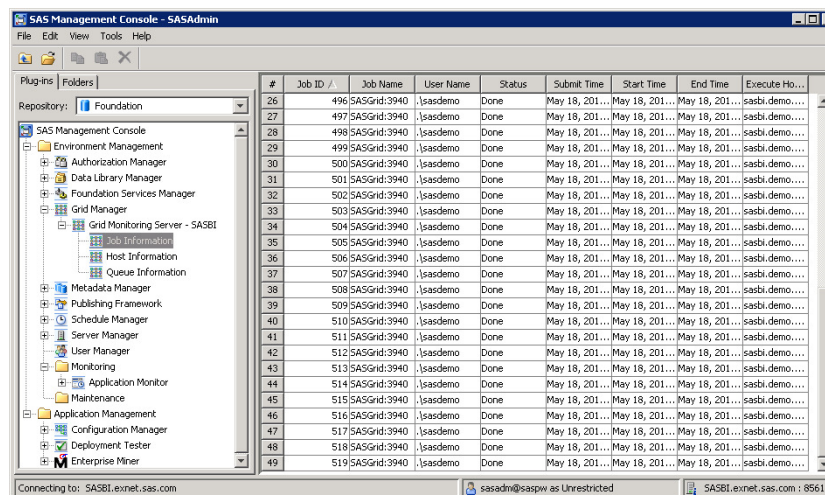
Platform Suite for SAS (aka, PSS) is a collection of software products from Platform Computing used as the core scheduling and process management tools that power the larger SAS Grid offering. PSS is the result of a long time partnership between SAS and Platform extending back over 10 years. PSS is current at version 6.1 with SAS 9.3. The products in Platform Suite for SAS include Platform LSF, Platform Process Manager, Platform MPI and Grid Management Service.

Platform LSF provides a robust scheduling and workload management capability across the compute nodes in the cluster. Platform Process Manager provides the ability to schedule workflows across the cluster. Many times Process manager is considered “the scheduler” in that it logically provides the ability to submit a workflow to the grid. In the strictest sense Process Manager is used to create a workflow and define the dependencies around the steps in that workflow (including items like launching time, workflow logic and failure exception handling). Platform LSF technically is eventually passed the defined workflow from Platform Process manager for execution on the nodes. In this scenario Platform Process Manager is a workflow engine while Platform LSF is actually “the scheduler”. This is simply a case of overloaded terminology but is a common point of confusion.

Platform MPI is an implementation of MPI or Message Passing Interface originally created by a now acquired company formerly called Scali. This technology was eventually purchased and improved by Platform. Platform MPI is a core technology used in a number of alternate technology offerings not used by SAS at this time. While Platform MPI is installed along with PSS, it is technically not used at this time but may be leveraged in the future to enable parallelization of underlying technology in SAS Grid or use by advanced users who may be creating their own SAS PROCs including calls to C level code in SAS programming for distributed memory computing (aka grids).

Finally Grid Management Service is a monitoring daemon that allows SAS users access to the information provided by Platform LSF including job status and load information within a SAS context. This includes the passing of information to the SAS Management Console to the Grid Manager plug-in installed with SAS Grid.

SAS Grid then includes PSS but also other components including the licensing and additional software necessary to connect a user to grid itself. This includes the SAS Management console plug-in called the Grid Manager plug-in, SASGSUB a command line submission client as well as the code included in individual SAS clients such as SAS Enterprise Miner, SAS Data Integration Studio and SAS Enterprise Guide that include both configuration options and the ability to submit jobs to the grid via SAS MetaData Server. This work is the culmination of the combined working experience with SAS and compute grids by both SAS R&D as well Platform Computing for last decade.



The screenshot shows the SAS Management Console - SASAdmin interface. On the left, there is a tree view of the repository structure, including folders like 'SAS Management Console', 'Environment Management', 'Authorization Manager', 'Data Library Manager', 'Foundation Services Manager', 'Grid Manager', 'Grid Monitoring Server - SASBI', 'Metadata Manager', 'Publishing Framework', 'Schedule Manager', 'Server Manager', 'User Manager', 'Monitoring', 'Application Monitor', 'Maintenance', 'Application Management', 'Configuration Manager', 'Deployment Tester', and 'Enterprise Miner'. The main area displays a table of job execution details.

#	Job ID	Job Name	User Name	Status	Submit Time	Start Time	End Time	Execute Ho...
26	496	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
27	497	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
28	498	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
29	499	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
30	500	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
31	501	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
32	502	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
33	503	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
34	504	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
35	505	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
36	506	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
37	507	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
38	508	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
39	509	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
40	510	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
41	511	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
42	512	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
43	513	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
44	514	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
45	515	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
46	516	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
47	517	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
48	518	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...
49	519	SASGrid:3940	..sasdemo	Done	May 18, 201...	May 18, 201...	May 18, 201...	..sasbi.demo...

Figure 2 SAS Management Console Grid Manager Plug-in

## YOU MIGHT NEED SAS GRID IF:

These reasons generally fall into three basic categories including lack of performance of existing infrastructure, hardware refresh/update and lack of robustness of infrastructure including single points of failure. The following is a list of commonly referenced pain points and reasons for considering SAS Grid along with their corresponding discussion points.

- **Critical workloads queuing**

Without a workload management system many sites simply use the called “sneaker net” meaning they simply communicate over the cube wall regarding who is next in line. SAS users organize amongst themselves the best way to utilize valuable compute cycles for mission critical applications. In larger groups this can be confusing at best and disastrous at worst.

- **Slow SAS Infrastructure response time (per job)**

As a SAS user, sometime response times from a shared compute infrastructure leaves something to be desired. In many cases classic bottlenecks seen in single server environments are easily handled with SAS Grid.

- **No ability to control workloads systematically or centrally**

Not having the ability apply a set of policies systematically using a workload management system over the compute resources means a lack of control over one of the most important business assets in today’s technology powered business. The lack of use of a workload management system also means that there is little or no ability to be efficient with existing resources or to understand their use over time. A workload management system provides all these benefits as well as flexible, efficient and configurable use of the investment in technology.

- **Aging infrastructure (over 3 to 5 years old)**

On average technology buyers refresh their investment in their compute infrastructure every three to five years. This phenomenon is probably most directly attributed to some combination of Moore’s Law and the time it takes various businesses to end of life and replace critical systems. Not doing so means not leveraging the latest and greatest in not only hardware but software as well. Ultimately this is a competitive issue.

- **Inability to meet business SLAs**

When technology managers, systems administrators and analyst cannot provide the information they glean from SAS in a timely manner to their management this ultimately not only puts a company’s future but individual careers at risk. Meeting the needs of fast paced information based companies, whether codified in a true Service Level Agreement or not, can be addressed by the use of SAS Grid.

- **Lack of Highly available infrastructure (single points of failure)**

Many companies work very happily along for many years without giving much thought to single points of failure in their infrastructure. The problem is that eventually with scale concerns and overall IT disaster planning eventually coming to light this problem needs to be addresses. This is done in two ways. Firstly Platform LSF itself is highly available with an ability to fail over to some configurable set of compute nodes to keep the grid itself running. Additional components of the grid and for that matter other software of a critical nature can also take advantage of a built in feature of LSF called “EGO” which is a highly flexible software for keeping software services running.

- **No ability to do any form of job parallelism**

SAS users in situations simply run their workload in series never aware that part of all of their workload could be broken logically into independent units and run in a fraction of the time across multiple nodes.

- **No ability to understand overall efficiency**

This is again a major benefit of having an integrated workload management system as a part of the grid infrastructure. Without the ability to track what occurs on a daily basis the simple act of problem identification is difficult if not impossible. Also many times a user’s perception of overall system efficiency doesn’t match reality or they are simply not aware of the entire load on the system. In either case a workload manager provides the ability to analyze and explain what is happening on systems over time including the efficiency of license use, total job throughput as well as who and what is using the most compute cycles, memory and the all important IO.

- **Lack of correct number of SAS licenses to support total workload**

SAS Grid provides a flexible infrastructure for maximizing the investment made in software licensing. The ability to provide configurable access to the compute infrastructure based upon business role translates to job placement on the grid according to a policy not based upon persistence or gaming of the resources by the submitting user. This means the projects and programs most important to business success are being addressed first while allowing flexible access to resources during non peak periods.

- **No clear path to a scalable easily expandable infra to support long term growth**

Due to the nature of the way that SAS Grid is installed adding compute nodes to expand analytical capabilities does not require architecting a completely new solution or even reinstallation of the existing software. Expansion is as simple as the addition of compute nodes and licenses to meet the need.

## CONCLUSION

The decision to include SAS Grid during an upgrade or consolidation is a multifaceted one. There are many benefits of SAS Grid to SAS users, IT and entire business. The value of a robust and configurable enterprise level solution for SAS applications and users is invaluable to many individual careers and businesses alike. There are obviously many considerations well beyond technology not covered here including licensing and hardware costs as well as the need for training not only users but also IT staff in the effective use of all the new technology involved. All that said SAS Grid has been successfully deployed at many of largest SAS using institutions around the world for many years now and should be considered a stable and well supported technology. A smaller SAS Grid can be a great entry point for rapidly expanding new businesses or a game changing life saver for long established consumers of SAS who really simply need a way to efficiently control resources and costs as well.

## REFERENCES

- <http://www.sas.com/high-performance-analytics/>
- <http://www.sas.com/grid>
- <http://support.sas.com/rnd/scalability/platform/index.html>
- [http://support.sas.com/rnd/scalability/grid/PSS6.1\\_Unix\\_Install.pdf](http://support.sas.com/rnd/scalability/grid/PSS6.1_Unix_Install.pdf)
- [http://www.hpcwire.com/hpcwire/2007-11-02/platform\\_acquires\\_scali\\_manage\\_business.html](http://www.hpcwire.com/hpcwire/2007-11-02/platform_acquires_scali_manage_business.html)
- <http://www-03.ibm.com/press/us/en/pressrelease/36372.wss>

## RECOMMENDED READING

- <http://support.sas.com/rnd/scalability/grid/gridpapers.html>
- <http://support.sas.com/rnd/scalability/grid/griddocs.html>

## CONTACT INFORMATION

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

Adam H. Diaz  
IBM Platform Computing  
500 SAS Campus Drive R1465  
Cary, NC 27513  
919-531-1644  
[adam.diaz@sas.com](mailto:adam.diaz@sas.com) or [adiaz@us.ibm.com](mailto:adiaz@us.ibm.com)  
<http://www.sas.com/grid>

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.