Paper SAS4240-2016

Creating a Strong Business Case for SAS® Grid Manager: Translating Grid Computing Benefits to Business Benefits

Marlos A Bosso, SAS Institute Inc.; Sao Paulo, Brazil

ABSTRACT

SAS® Grid Manager, as well as other grid computing technologies, have a set of great capabilities that we, IT professionals, love to have in our systems. This technology increases high availability, allows parallel processing, facilitates increasing demand by scale out, and offers other features that make life better for those managing and using these environments. However, even when business users take advantage of these features, they are more concerned about the business part of the problem. Most of the time business groups hold the budgets and are key stakeholders for any SAS Grid Manager project. Therefore, it is crucial to demonstrate to business users how they will benefit from the new technologies, how the features will improve their daily operations, help them be more efficient and productive, and help them achieve better results. This paper guides you through a process to create a strong and persuasive business plan that translates the technology features from SAS Grid Manager to business benefits.

INTRODUCTION

Grid computing and all its benefits are not a new concept to IT professionals. You can find lots of software and systems that uses grid computing as a resource to achieve high performance, to provide high availability, and to reduce costs.

SAS Grid Manager is the SAS implementation of grid computing for SAS software. SAS Grid Manager adds many features to SAS systems, such as the following:

- load balancing
- high-availability
- parallel processing
- significant reduction of the total cost of ownership (TCO)
- ease of scaling out
- prioritization of tasks
- real-time monitoring
- better governance with single point for management and administration.

Some of these features (such as reduction of TCO) can directly be monetized to help provide a return of investment (ROI). However, there are many more business benefits that are below the surface, and could make a great difference in creating a business case for SAS Grid Manager. In this paper, we will go deeper in examining four benefits that have proven to be the most persuasive.

1. Decreasing the time taken from the start of development of a new model to deploying this model in production. We will call this time-to-market.
2. Increasing a models accuracy, also known as lift.
3. Improving productivity for SAS users.
4. Decreasing losses caused by unplanned outages of SAS.

SAS GRID MANAGER FEATURES

To ensure a common level of understanding, this section briefly introduces grid computing, then details what you need to know when creating a business case for SAS Grid Manager. Because the goal is not to go deep into technical aspects, this paper does not explain how SAS Grid Manager works and it does not

provide technical details. This section is written this way so that no previous technical background is required. Therefore, both IT and business professionals can follow the guidelines provided to create a strong business case.

The order of the features in this paper is not related to their importance. A feature's importance is different in each situation, and should be determined when constructing the business case.

WHAT IS GRID COMPUTING

I´m taking advantage of a well-known and publicly available knowledge source to define grid computing - Wikipedia. In this way we can be sure we are using a well-known definition.

"Grid computing is the collection of computer resources from multiple locations to reach a common goal. The grid can be thought of as a distributed system with non-interactive workloads that involve a large number of files."

"Grids are a form of distributed computing whereby a "super virtual computer" is composed of many networked loosely coupled computers acting together to perform large tasks. For certain applications, "distributed" or "grid" computing, can be seen as a special type of parallel computing that relies on complete computers (with onboard CPUs, storage, power supplies, network interfaces, etc.) connected to a computer network (private or public) by a conventional network interface, such as Ethernet"

In very simple terms, this means that to make a grid we need several smaller computers, instead of a single large one. In other words, we can say that a grid is a group of several computers working together in a very coordinated way.

PARALLEL PROCESSING

Parallel processing is distributing tasks to different resources. These resources work on these tasks at same time; processing in parallel and not in sequence.

Because we have several computers in a grid, it makes sense to send different tasks so that different computers execute them. This enables us to avoid high concurrence and achieve better performance. SAS Grid Manager, however, goes beyond that. It allows a single SAS program to split itself into several subtasks and to send these subtasks to different computers. The result is a very significant gain in performance.

A public case from Bank of America states that some processes have had execution time improved by 90%. However, we will be more conservative when estimating performance gain while constructing a business case. Based on previous experience, I suggest that an average of 30% performance improvement is a good estimate for ROI calculations.

LOAD BALANCING

Load balancing is the process that receives tasks from requestors and dispatches them to resources that will then execute these tasks. There are several methods to perform load balancing. The simplest method assigns one task to each resource, without considering the complexity of each task or the load that each resource is handling. The most sophisticated methods consider many details when selecting which resource is best suited to handle each task.

The load balancing performed by SAS Grid Manager uses one of the more sophisticated methods. Before dispatching each task, SAS Grid Manager queries all resources to measure actual load and selects the best one to perform the task.

Using SAS Grid Manager to perform load balancing provides several advantages over using a third-party technology. I will highlight two of these advantages. First, because SAS Grid Manager understands SAS code and SAS programs, it is capable of splitting the tasks into subtasks (when it is appropriate). Second, the most common load balancing technologies are implemented by hardware (dedicated hardware that built to perform load balancing, commonly referred as a load balance appliance) and not by software. In this case, a benefit from the software industry plays on our side. Both hardware and software have big development costs, but only hardware technologies have high marginal production costs. Marginal production costs for software are so low that we can even consider them to be zero. Marginal production

costs are the costs associated with the production of one unit of the product after it has been fully developed. Thanks to the almost nonexistent marginal cost of software, in most cases it is much cheaper to implement the same features in software than in hardware.

Load balancing features also ensure the optimum utilization of your hardware. A possible alternative approach would be to have several servers to handle different user groups. For example, one server could be dedicated to users from the financial department, another server dedicated to marketing department, and so on. However, in this scenario there will be several moments when one server will be idle and the other server will have very high utilization. This situation means that we are not using the servers in the smartest way. Because each server needs to accommodate the peak demands from its user group by itself, we must increase the size of each server to accommodate these peaks. When you are using SAS Grid Manager, rather than dedicated servers, other servers will help in peak moments and don´t need to increase the size of the grid servers.

HIGH AVAILABILITY

SAS Grid Manager has two great resources to ensure that the SAS environment will be highly available. The first is a failover capability that constantly monitors critical services and automatically starts another instance if one of these services become unavailable. This process is known as failover, and it prevents SAS systems from going offline to users.

In addition to providing failover fors critical services, SAS Grid Manager has a rerun capability for running SAS tasks. Suppose that a very important task is running on a specific server and for some reason this server stops working. SAS Grid Manager identifies that the server has stopped, and it restarts the task in another server. This feature has the additional advantage that all work performed is not lost, because the task is restarted from the last checkpoint. This feature is very important when you have critical tasks that take long or run overnight.

Because this feature is implemented by software, it also has the benefit of having lower implementation costs compared to an implementation in hardware.

REDUCTION IN TOTAL COST OF OWNERSHIP

The first driver in TCO reduction is based on using software to implement several features that are most commonly implemented in hardware. Usually, SAS implementations use high-end hardware to ensure it is highly available and is load balanced. We benefit from a software industry advantage.

In addition, we can split the SAS environment among several smaller computers instead of a few large computers. The price of a server does not increase linearly with the number of cores in the server has. For servers with not many cores per server, the prices are very attractive. These servers are called commodity servers. However, if you need a bigger server, the price goes up exponentially. Being able to split computer capacity among several smaller servers is a big advantage. When this paper was written, we could easily find inexpensive server with up to 24 cores (2 processors with 12 cores per processor). However, the price increases very quickly for servers with more than 24 cores, making the price per core more expensive.

Because SAS Grid Manager has critical features implemented by software, we don´t need to use high-end hardware. Because the price for commodity hardware is much more attractive, that makes for a perfect choice. Using commodity hardware decreases the total cost, not just the purchasing costs, very significantly.

PRIORITIZATION OF TASKS

In all organizations, SAS systems perform tasks that have different requirements. Some critical tasks are crucial to business continuity, and must meet deadlines (for example, tasks to comply with regulatory requirements). There are also tasks that are not critical to business continuity. For those tasks, a delay would not cause a big problem

In traditional SAS systems there is no way to prioritize tasks. With SAS Grid Manager, you can set different priorities for tasks. You can even temporarily suspend less important tasks in order to run a critical task faster. This is a very valuable feature to help you meet important deadlines.

GOVERNANCE

SAS Grid Manager allows us to have a single SAS environment and to keep control of the amount of resources that each user group can access. This enables us to have a single environment to maintain, administer, configure, back up, and update.

Imagine your company has two user groups. One group has 30 heavy users from the finance department. The other group has users from marketing department. This group has only five users and they have a much lighter use of the SAS system. It is obvious that the first group needs more hardware resources than the second group. If you put all of the users on the same server, there is no way to split the amount of resources available to each group, and splitting costs between the groups would be an issue. The budget for SAS usage for the second group is probably not compatible with the budget from first group. It will be very challenging to meet the requirements for both groups in a single system. The most common solution is to create two different SAS environments, which doubles the administration efforts, two environments to maintain, configure, back up, update, and so on. Besides, there is a significant risk that the two environment's configurations will become different over time, making them even more of a challenge to maintain.

With SAS Grid Manager we can control the hardware utilization. We can set the maximum number of simultaneous tasks that each group can submit, while using the same hardware. So we avoid duplicating administration efforts as well the risk of different configurations due to human error, while ensuring that computational power is split as needed.

REAL-TIME MONITORING

SAS Grid Manager also add features to monitor server utilization in real time. We can see not only resource utilization, but also who is consuming the resource. Using a friendly web interface, we can easily check how long each task has been running and how many resources (such as CPU, memory, and I/O) it is using.

The utilization history is also available, so it is possible to determine how many resources each individual user and each user group has used. This allows us to create a billing process based on users' utilization. Charging users by utilization helps to reduce total costs, and it disciplines users to become more efficient and intelligent in their use of SAS.

SCALE OUT

In traditional SAS systems the only way to increase capacity is to increase the size of the servers where SAS is installed. SAS Grid Manager provides the advantage of allow scale out. Because SAS Grid Manager is a distributed system, when demand rises we only need to add servers to the grid. We do not need to increase the size of our servers.

The new servers do not have to be identical to the existing servers. They do not even need to have equal resources. We can add servers with different numbers of cores and different amounts of memory, and they will nicely be part of the same grid and will take on different amounts of tasks based on their utilization.

This capability greatly decreases the initial costs of a SAS system, because we can start small and add resources to meet demand as it increases. Because we are only adding more servers to the grid, there are no associated migration projects, which means environment expansion becomes much cheaper, easier, and smoother.

BUSINESS BENEFITS

Some features convert to business benefits and allow us to calculate a ROI in a very direct way (for example, the TCO cost reduction). However, some other features might be hidden from the eyes of someone who is working in a SAS Grid Manager business case for the first time.

In this section we are going to explore both of these features.

DIRECT BUSINESS BENEFIT

Cost reduction is a direct business benefit and does not require any translation from the technical details. TCO reduction and load balancing do not need much elaboration in the business case, just a way to estimate its value.

All companies with heterogeneous and diverse hardware have estimated costs for different hardware. When we position SAS Grid Manager we suggest a move to commodity hardware. To calculate the savings generated by this change we can talk to IT infrastructure teams and ask them to estimate annual costs for two scenarios. The first scenario is high-end hardware without SAS Grid Manager, and the second scenario is commodity hardware with SAS Grid Manager. It is important to make sure IT includes all dependencies, servers, storage, network, support, datacenter, and so on. It is very likely that the SAS Grid Manager scenario will have a significantly lower TCO.

Load balancing allows SAS Grid Manager to optimize hardware utilization, because tasks are intelligently distributed across all grid nodes. Because of this distribution, the size of the environment with SAS Grid Manager is smaller than the environment without it, which makes TCO even lower.

The annual difference in costs is one of the components of the ROI we want to calculate. See Figure 1 for an illustration.

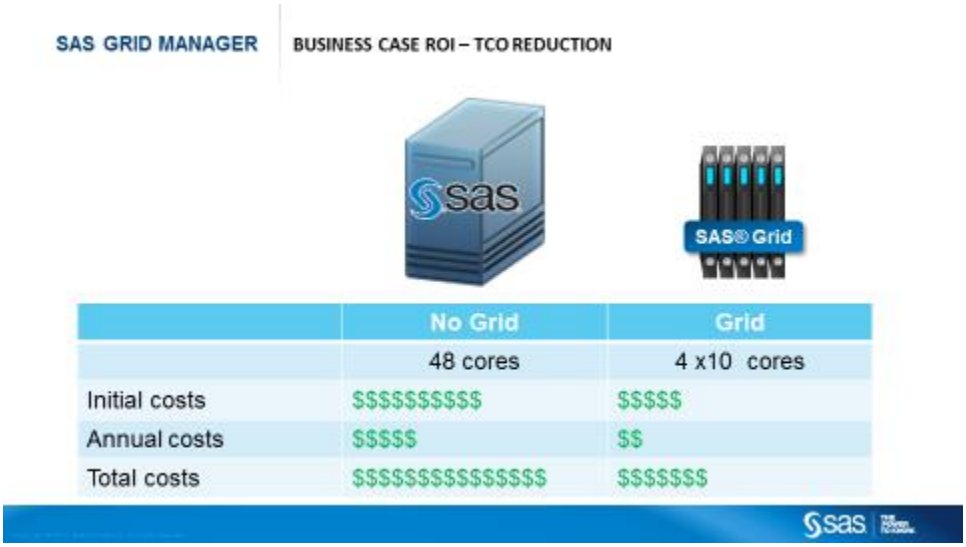

Figure 1 - Representation of ROI Due to TCO Reduction

DECREASE TIME-TO-MARKET

Time-to-market is the amount of time spent between determining a business problem that needs to be solved and starting to collect benefits from the model developed to address this problem. This period includes all of the steps of the analytical cycle: identify the business problem, prepare date, explore data, transform and select data, create the model, test the model, and deploy the model.

Models are created to solve business problems, regardless of whether they it will reduce costs, increase revenue, segment the customer base, optimize investments, and so on. Models are always supposed to produce gains. Therefore, if you can start collecting gains earlier, you will gain more.

As explained before, we can estimate that using SAS Grid Manager will improve performance by 30%. With better performance, all tasks in SAS will be faster, decreasing the amount of time spent in computer tasks during the analytical cycle. Since time spent in computer tasks is lower, the total time is also lower.

To calculate a ROI based on the improved performance, we need a few indicators. First, we can estimate time reduced in processing. Here again, we will be very conservative and estimate a 10% reduction in time to market. Second, we need to work with business analysts to estimate an average gain per model

and the number of models deployed to production per year. With both of these indicators, we will be able to calculate an annual gain due to faster deployment of models.

For example, suppose that before SAS Grid Manager, time-to-market was 90 days. Business analysts told us that about 20 models are deployed per year. We estimate a 10% reduction in the time to market, which corresponds to 81 days after implementing SAS Grid Manager. Because there are 20 deployments per year and SAS Grid Manager can provide a gain of nine days for each deployment, we will have a total of 180 more days of gains. Using this estimate of gain produced for each new model per day, we can arrive at values. See representation of this in Figure 2.

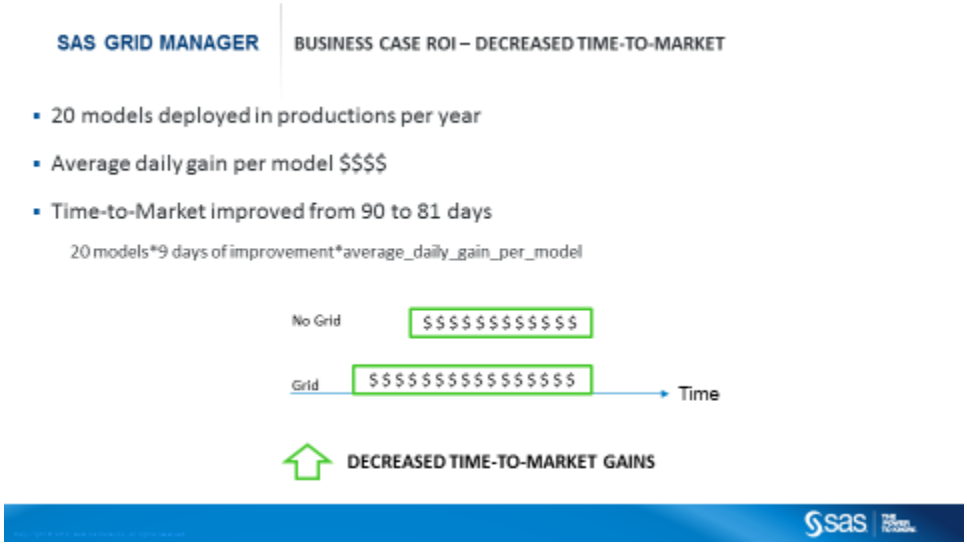

Figure 2 - Representation of Decreased Time-to-Market Gains

MODELS LIFT IMPROVEMENT

In the real world we live in, everyone has deadlines to meet. All analysts that are working in a new model have a limited amount of time to complete their jobs. Creation of statistical and analytical models is complicated task. There is still a portion of the job of creating new models that relies on running lots of tests. We know that good information, process, and methodology can decrease the dependency on running tests, but it difficult to eliminate it entirely. It is also difficult to determine the amount of time required to improve a model's lift, or accuracy. However, it is easy to see that with more time for testing, better results will be achieved.

SAS Grid Manager cannot postpone the deadline from analysts, which would be ideal. However, it can run SAS tasks faster, which produces the same result. Because SAS tasks run faster, analysts can run more tests in the same amount of time. Running more tests increases the average lift of the models.

It is challenging to determine which model will be better and how much better it will be. Nevertheless, when considering a portfolio that contains lots of models, the increase in accuracy is measurable. Estimating the amount of increase in lift is still challenging. We can take two approaches to estimate the increase. The first method is to estimate a very low increase, such as something around 3%, which is a value that will not require much justification because it is so low. The second approach is to partner with business analysts so that they can estimate a value that makes sense in their reality.

After we get an estimate of the increase in lift, we can calculate the ROI by multiplying the increase by the total number of models in production and by the average gain per model. See Figure 3 for a representation.

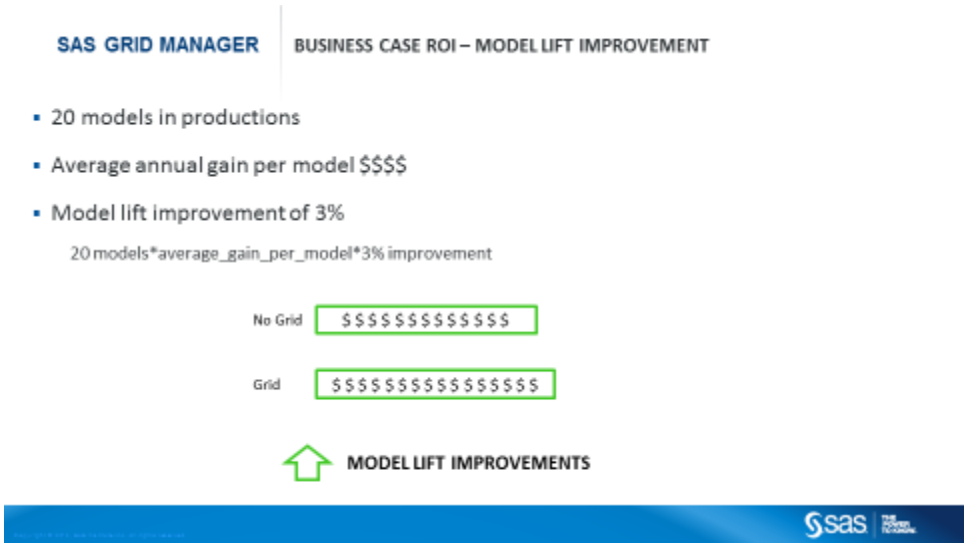

Figure 3 - Gains Due to Models Lift Improvement

USERS' PRODUCTIVITY IMPROVEMENT

Groups of SAS users are an important part of SAS systems, so they must be considered as part of the business case. Some users spend almost their full day working with SAS, while others spend just a few hours per week. Regardless of the level of usage, all users' activity will benefit from better performance. The gain in productivity is very simple to calculate.

Once again, we will be conservative in our estimates. Let´s assume that on average each SAS users spends 2 hours per day working with SAS. Let´s also assume that SAS Grid Manager will improve performance by 25% (this figure is even more conservative than our previous benefits calculations). It means that each SAS user will be able to perform the same tasks that would normally take two hours per day in just 1.5 hours per day.

We can estimate costs for employee per hour. Every company has this indicator, so it should be easy to find this information. Using this cost and the total number of SAS users, we can calculate the productivity gain by multiplying the cost per hour per employee by the number of SAS users and multiplying that by half an hour (the time saved). Finally, we multiply this figure by 365 to determine the gain per year. See Figure 4 for an illustration.

In some situations, this metric could sound insensitive. In these cases, we can transform this productivity gain to an ability to perform more tasks per year and use that as the value.

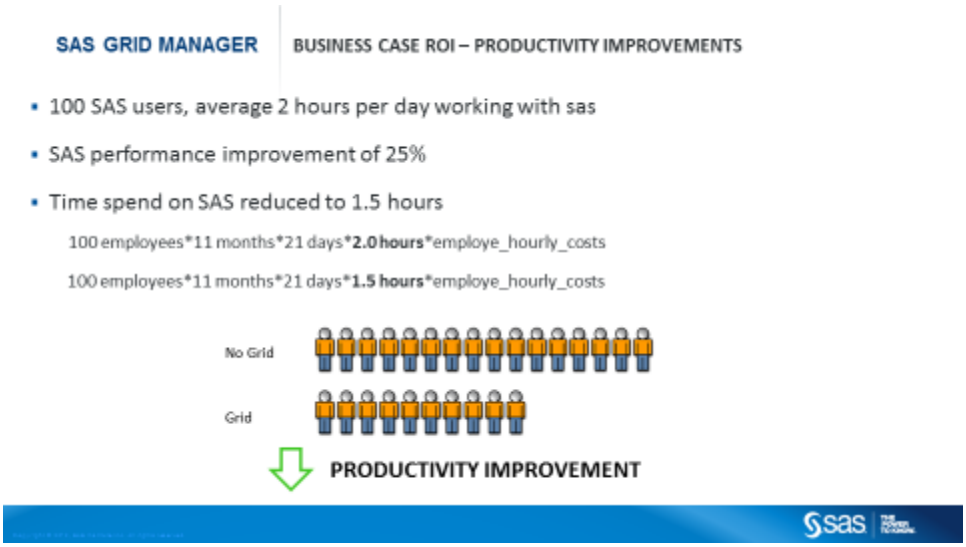

Figure 4 – Users' Productivity Improvement

REDUCED LOSSES DUE TO UNPLANNED OUTAGES

It is very common to have critical business process that rely on SAS. This means that any unplanned outage of SAS environments will result in business problems. Most of these problems can easily be associated with losses.

Because SAS Grid Manager implements high availability, we can calculate a ROI that is due to reduction of unplanned outages. As we have done throughout this paper, we will be conservative and estimate a reduction in outage hours per year, rather than a total elimination.

The key here is to partner with business analysts to understand how many critical processes rely on SASChose one of these processes and estimate the losses per hour of outage. We can then use this estimate to estimate the total losses per outage. The strategy to be conservative helps here once again, because we are estimating only one process loss. If we used this figure as an average, someone might disagree with using that figure as an acceptable value. Therefore, the idea is to underestimate the total loss so that everyone will agree with the total loss.

For example, suppose a company has 50 critical business process that rely on SAS. We choose one of this processes and work with business analysts to estimate an hourly loss that would be caused by this process not running because of a SAS outage.

Suppose the customer has a process to calculate how much of a discount they should offer to a customer who asks to cancel their services. Based on customer data, offering a calculated discount is an intelligent decision, retaining a customer who is good for the business while avoiding having to offer higher discounts than needed. If there is a SAS outage, the company cannot determine the best discount to offer, so the customer will not be retained, and losses will occur.

Business analysts can help identify the losses for a single process. We can then multiply this value by the number of critical processes to estimate total losses for all 50 critical process. Because we always want to be conservative, we can multiply the loss per process by 20 instead of 50.

Finally, to determine a ROI, we just need to estimate the improvement. We find the total unplanned unavailability for the last year and reduce this figure by 40%. For example, if there were 40 hours of unplanned availability last year, we estimate that this figure will be reduced to 24 hours this year. Calculate the savings and add to the ROI. See Figure 5 for an illustration.

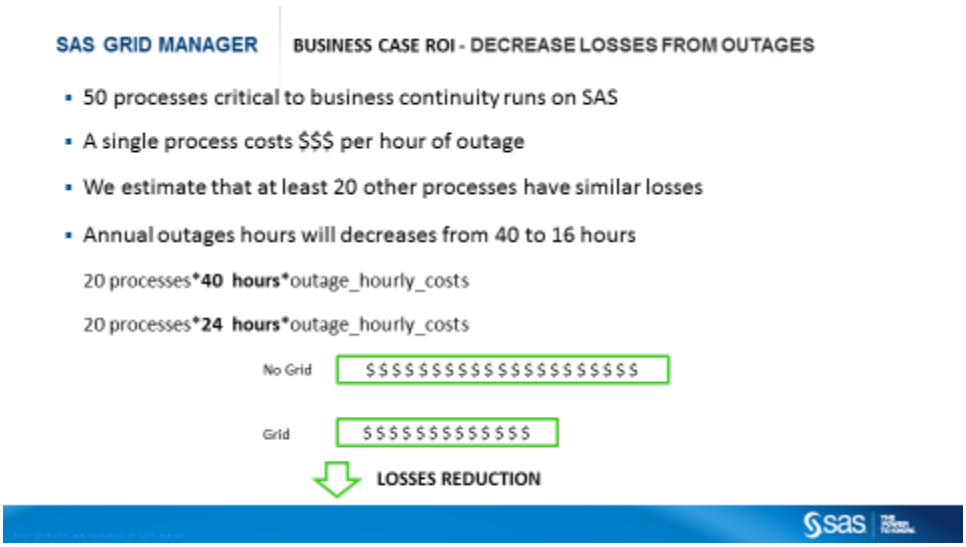

Figure 5 - Representation of Losses Due to Unplanned Outages in SAS System

CONCLUSION

SAS Grid Manager adds many features to SAS systems. IT professionals really like these features, because they make their lives better. However, most of the benefits generated by SAS Grid Manager are technical and do not translate directly to a business benefit. This paper has guided you through a process to translate these features into business benefits, which leads to a strong business case with an attractive ROI for a SAS Grid Manager implementation project.

We have talked about four indirect benefits for businesses (decreasing models' time-to-market, improving models' lift, improving users' productivity and decreasing losses from unplanned outages) plus one direct benefit (TCO reduction). In most cases you will not be able to use all these benefits, so we recommend that you choose two that are the most compelling. Elaborate your choices more fully, but present the two individually (because each one can be presented as a strong argument on its own). Present the two remaining benefits together, along with other benefits from the project (if any), and a few not as strong arguments. Using all of these arguments together results in a strong overall argument.

One key factor to make a successful business plan strong is the partnership between the SAS system users and the IT professionals who will work to implement and maintain the plan. It is challenging for only one of these groups to obtain all of the information required and construct a full business case without help. Because SAS Grid Manager has many benefits, it is much easier to create a collaborative group and have both groups work together.

The last suggestion in this paper is to always underestimate when using estimated values in the business case. The reason for this is that that people will test a business case during a business case analysis to make sure the case is is consistent. When your case is tested, you will need to explain how you determined the estimated values, and you must be ready to prove that the values are achievable. This task becomes much easier when you are using underestimated values. Although we have seen that some companies have experienced a 90% performance gain with SAS Grid Manager, we use only a 25-30% gain in our business case. It could be difficult to prove a 90% improvement, but using one-third of this value is easily accepted, even in situations where 90% was achieved.

REFERENCES

"Grid computing" Wikipedia. Accessed February 18th, 2016. Available https://en.wikipedia.org/wiki/Grid_computing

SAS Institute. “Bank of America avoids gridlock in credit risk scoring, forecasting” Available http://www.sas.com/en_us/customers/bank-of-america-credit-risk.html

ACKNOWLEDGMENTS

The process of elaborating business cases for SAS Grid Manager that is described in this paper was developed and improved through time with help of several customers. These customers consistently provided feedback and shared crucial business information that led to successful business plans. I would like to acknowledge the collaboration of all customers who I have worked with over the years. These customers helped to create and develop the guidelines explained in this paper.

I also need to acknowledge the great collaboration from my peers at SAS Brazil, all colleagues from the Pre-Sales and Sales teams with which I have engaged in SAS Grid Manager opportunities, and especially Gulherme Reis and Rodrigo Africani.

Finally, I especially acknowledge and offer thanks to my technical reviewers and good friends Adriana Silva and Sergio Zaccarelli.

RECOMMENDED READING

- *Grid Computing in SAS 9.4*
- *Platform RTM for SAS Administrator Guide*
- *Scheduling in SAS 9.4*

CONTACT INFORMATION

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

Marlos Bosso
SAS Instituto Brasil
marlos.bosso@sas.com
https://br.linkedin.com/in/marlosbosso