

## Best Practices for Implementing High Availability for SAS® 9.4

Cheryl Doninger, SAS; Zhiyong Li, SAS; Bryan Wolfe, SAS

### ABSTRACT

There are many components that make up the mid-tier and server-tier of a SAS 9.4 deployment. There is also a variety of technologies that can be used to provide high availability of these components. This paper will focus on a small set of SAS recommended best practices for a consistent high availability strategy across the entire SAS 9.4 platform. We will focus on two technologies: clustering as well as the high availability features of SAS Grid Manager. For the clustering, we will detail newly introduced clustering capabilities in SAS 9.4 such as the middle tier SAS Web Application Server and the server tier SAS metadata clusters. We will also introduce the small, medium and large deployment scenarios or profiles which make use of each of these technologies. These different deployment scenarios will reflect the typical customer's environment and address their high availability, performance and scalability requirements.

### INTRODUCTION

What is a highly available SAS environment and why is it important? SAS analytics has become a mission critical component of most organizations. It is therefore critical that the SAS environment be highly available to analysts and business users at all times.

*High availability is a system design protocol and associated implementation that ensures a certain absolute degree of operational continuity during a given measurement period. Availability refers to the ability of the user community to access the system, whether to submit new work, update or alter existing work, or collect the results of previous work. If a user cannot access the system, it is said to be unavailable. (Wikipedia 2011)*

A planned SAS deployment installs and configures various components across the following tiers:

- SAS Metadata Server Tier
- SAS Application Server Tier
- SAS Web or Middle Tier
- SAS Client Tier

In order to ensure consistent availability of the SAS environment, you must eliminate any single points of failure across the first three tiers listed above. Client software is not generally part of a high availability design and therefore will not be discussed in this paper. SAS recommends the high availability features of SAS Grid Manager in addition to clustering as the best approach for a consistent and complete high availability solution for SAS 9.4. This eliminates the need to purchase additional third-party software or hardware, such as a load balancer. This paper will focus on the SAS Intelligence Platform components and will not cover other services specific to other products or solutions that run on top of the platform. The following SAS Analytic platform components must be deployed and configured for high availability:

- SAS® Metadata Server (Metadata Server)
- SAS Web Server (Web Server)
- SAS Web Application Server (Web Application Server)
- SAS Web Infrastructure Platform Data Server (WIP Data Server)
- SAS JMS Broker (JMS Broker)
- SAS Cache Locator (Cache Locator)
- SAS Object Spawner (Object Spawner)
- SAS OLAP Server (OLAP Server)
- SAS Environment Manager Server (EV Server)
- SAS Environment Manager Agent (EV Agent)
- SAS Deployment Agent (Deployment Agent)

The above list includes the full name of each component followed by an abbreviated name that will be used in the diagrams in this paper. The placement of some of the components in the diagrams was chosen to match the sample plans that are typically used for a SAS deployment. In some cases, multiple configuration directories must be created on the same node. It is recommended that you specify the same configuration level across all configuration

directories in the deployment.

## Terminology

The following terms and associated definitions will be used throughout this paper:

### **Primary Node**

The node running either the only active instance of a service or the master instance of a service in a master/slave scenario.

### **Cold Standby Node**

A node, other than the node where a service was previously running, where a new instance of a service is started as needed.

### **Warm Standby Node**

A node where a backup instance of a service is pre-started to take over as needed.

### **Restart**

The act of starting a failed service on the same node where it was previously running.

### **Failover**

The act of starting a service on a cold standby or promoting a service to master status on a warm standby.

### **Cluster**

Two (or more) active instances of a service to ensure that if one instance fails, there are one (or more) active instances available for continued operation. A cluster may also provide some level of load balancing or built-in failover. The services in a cluster can coexist on the same physical node (vertical clustering), on a group of nodes (horizontal clustering). Horizontal clustering should be used to achieve high availability because it requires two (or more) nodes and eliminates a single point of failure from a hardware perspective.

### **Active-Passive Mode**

A service will execute on only one node at any given time while cold standby nodes are available for executing the service in the event of a failure. While there are multiple nodes *capable* of executing the service, only one of those nodes is *actually* executing the service at a given point in time. The service will be started only on a cold standby node if the node on which the service was originally running becomes unavailable.

### **Active-Active Mode**

A service will execute on multiple nodes concurrently with each service instance providing full capabilities to consumers. It is often desirable to ensure that a minimum number of service instances are always functioning.

### **Master-Slave Mode**

A service will execute on multiple nodes concurrently. However, one instance will be designated the master and the instance to provide full capabilities to consumers. An additional instance will be designated the slaves and will take over the role of master in the event the master fails. The details of this transition are specific to individual technologies.

## IMPLEMENTING HIGH AVAILABILITY FOR SAS 9.4

### SAS Grid Manager

SAS Grid Manager has the ability to monitor services on a node and automatically restart them on the same or different node in the event of a failure. This feature of SAS Grid Manager will be leveraged in the scenarios described in this paper to provide high availability of all critical services within a SAS deployment. SAS Grid Manager can be used to implement an active-passive mode of operation where a service is failed over to a cold standby node. An example of this type of failover would be running a non-clustered SAS Metadata Server on a single node in the deployment. If that node were to fail, SAS Grid Manager would automatically failover the SAS Metadata Server to a cold standby node in the deployment. It can also be used to implement an active-active mode of operation by monitoring all active instances of a service and restarting the service on the same node or failing over to a cold

standby node if a previous node became unavailable. For example, assume you have ten nodes in your environment and have licensed SAS® Integration Technologies to run on three of the nodes. In this case, you want to ensure that you always have three Object Spawners running at all times on any three of the ten nodes. SAS Grid Manager will monitor the three running instances and if one fails either restart on the same node or if the node fails, failover to another node in the environment thereby guaranteeing three available Object Spawners at all times. In addition, SAS Grid Manager can maintain the dependencies between the services across the deployment to ensure that services are restarted in the required order. The details of how to configure SAS Grid Manager to provide high availability are beyond the scope of this paper. Refer to the “Recommended Reading” section at the end of this paper for a link to “High Availability with SAS Grid Manager”.

A shared file system is required in a SAS Grid Manager deployment to enable all nodes in the Grid to be able to read and write to a common file system. This shared file system also plays an important role in providing high availability by allowing the use of shared configuration directories where applicable and a shared location for the SAS Metadata Repository in an active-passive configuration. It is assumed that high availability is a feature of the shared file system itself and configured as directed by the file system vendor.

**Note on licensing:** Use of SAS Grid Manager to provide high availability of components running on nodes outside of the server tier impacts calculation of the total core count for the SAS Grid Manager license. If SAS Grid Manager is used to provide active-passive high availability of a single instance of the SAS Metadata Server that is running on a (typically dedicated) node outside of the server tier, then the number of cores on that node must be included in the total core count for the SAS Grid Manager license. If SAS Grid Manager is used to provide active-active high availability of a SAS Metadata Server cluster that is running on nodes outside of the server tier, then the number of cores on *just one* of the nodes in the cluster must be included in the total core count for the SAS Grid Manager license.

### **SAS Metadata Server Cluster**

SAS 9.4 provides the option of implementing a SAS Metadata Server cluster. This feature provides redundancy and high availability of the SAS Metadata Server, which is a core component of the SAS infrastructure. Clustering ensures that the service will continue to operate in the event of a node failure.

A SAS Metadata Server cluster is a group of three or more nodes that have been configured as identical metadata servers. Each node runs its own server process and has its own server configuration information, journal file, and copy of the repository data sets. In addition, each node maintains a complete in-memory copy of the metadata repository. The nodes work together as if they were a single metadata server.

Client applications and users interact with the cluster in the same way that they would interact with a metadata server that is not clustered. A load-balancing process automatically distributes work among the nodes. If a node ceases to operate, the server continues to be available via the remaining nodes.

### **Other SAS Application Servers**

You can create load-balancing clusters of SAS Workspace Servers, SAS Stored Process Servers, SAS Pooled Workspace Servers, and SAS OLAP Servers. For all clusters except OLAP, load balancing is handled by the Object Spawners that are associated with each server. The SAS Object Spawner is aware of any failures and is responsible for restarting any failed servers based on each specific configuration. Therefore, the SAS Object Spawner is a critical component that must be configured for high availability; the individual servers that it spawns do not.

SAS OLAP Servers do not use Object Spawners. They handle load balancing directly and, therefore, must be configured for high availability.

### **SAS Web Application Server and SAS Cache Locator Cluster**

The SAS Web Application Server also supports clustering. In a cluster configuration, multiple SAS Web Application Server instances (managed servers) participate in a load-balancing scheme to handle client requests. The managed servers in a cluster run on two or more nodes in order to achieve high availability. The benefit of clustering the SAS Web Application Servers is that the web applications running within those servers are much more likely to continue to be available to users.

Even though the SAS Web Application Server supports clustering, it should be noted that not all SAS Web Applications support clustering due to technical limitations. A list of the SAS Web Applications that do not support clustering is available in *SAS 9.4 Intelligence Platform System Administration Guide* included in the Recommended Reading section at the end of this paper. Applications that do not support clustering are deployed into a single node and therefore will not be highly available via clustering. In order to make these exception applications highly available, the SAS Web Application Servers that host them will need to be configured to use SAS Grid Manager for

failover. The configuration details for that failover are beyond the scope of this paper. SAS continues to work on resolving these limitations so that all applications will support clustering in future releases.

The SAS Cache Locator can also be run in a cluster on two or more nodes in order to achieve high availability.

### **Other SAS Middle Tier Components**

The SAS Web Infrastructure Platform Data Server, SAS Environment Manager Server and the SAS JMS Broker each support a master-slave mode of clustering where there are two running instances with one instance running as the master and the second instance running as a slave or warm standby. If the master instance fails, steps must be completed for the warm standby to assume the role of master. The specific steps required for the warm standby to assume the role of master are specific to each of the three technologies. Further details can be found in *SAS 9.4 Intelligence Platform: Middle Tier Administration Guide, Second Edition* listed in the Recommended Reading section.

## **MINIMUM CONFIGURATION FOR HIGH AVAILABILITY**

This initial scenario uses the smallest number of nodes possible to deploy the SAS Metadata Server Tier, SAS Application Server Tier and SAS Middle Tier in a highly available environment. In order to eliminate the single point of failure in a single node environment, you must have at least two physical nodes in order to achieve high availability. The two nodes must be sized adequately with respect to number of cores, memory, disks, etc. to meet the needs of the workload that will be run. The goal of this paper is to start with the minimum configuration needed to provide high availability and then build on this configuration through a series of scenarios. Each subsequent scenario will add additional nodes to scale out for an increasing number of users, larger workloads and maximum performance.

### *SAS Metadata Server Tier*

In this scenario the SAS Metadata Server is running in an active-passive mode and is *not* clustered since a minimum of three nodes is required to implement a SAS Metadata Server cluster. SAS Grid Manager provides automatic restart or failover of the SAS Metadata Server as needed in the event of a failure. This requires the SAS Metadata Server and repository to be deployed onto a shared file system so that if the SAS Metadata Server must failover to a cold standby node, it will still have access to the executables, configuration and repository.

### *SAS Application Server Tier*

The SAS Object Spawner and SAS OLAP Server are actively running on the second node and SAS Grid Manager will provide restart on the second node or failover to the first node in the event that the second node fails. Deploying a shared configuration of the SAS Object Spawner and the SAS OLAP Server facilitates initial deployment as well as failover.

The SAS Environment Manager Agent and SAS Deployment Agent are actively running on each node in the deployment. If either fails, SAS Grid Manager will attempt to restart these services on the same node. Failover to another node does not apply in this case because there cannot be more than one instance of these services running on a node. Therefore the configuration of these components must be done on a per-node basis.

### *SAS Middle Tier*

The SAS Web Server is running in active-passive mode and is deployed onto a shared file system accessible by all, in this case two, nodes participating in the active-passive configuration. SAS Grid Manager will provide restart on the first node or failover to the second node in the event that the first node fails.

The SAS Web Infrastructure Platform Data Server, SAS Environment Manager Server and the SAS JMS Broker are all running in master-slave mode with a master instance on one of the nodes and a warm standby on the other node. If the master instance fails or the node it is running on becomes unavailable, the warm standby will assume the role of master.

The SAS Web Application Server and SAS Cache Locator are clustered across both nodes and SAS Grid Manager is used to restart any failed instance on the same node in an effort to constantly maintain two running instances of each.

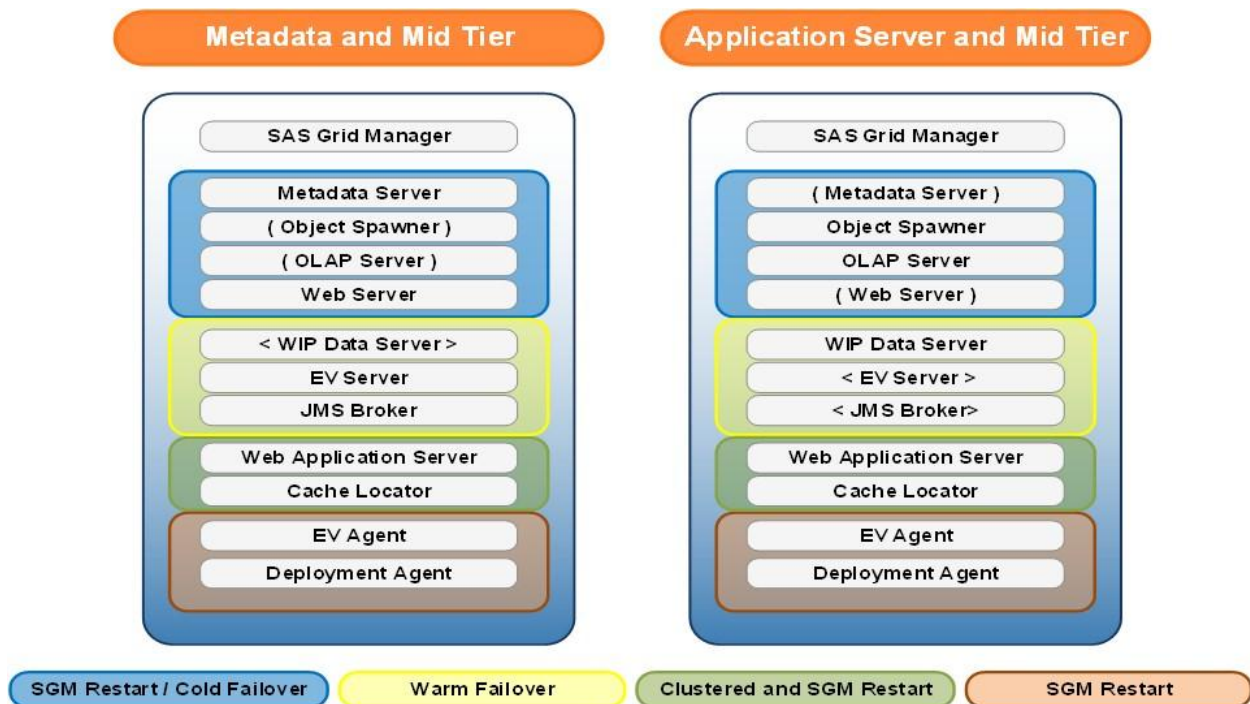
A shared configuration is recommended for the following components:

- SAS Metadata Server
- SAS Object Spawner
- SAS OLAP Server
- SAS Web Server

An independent or per-node configuration is recommended for the following components:

- SAS Environment Manager Agent
- SAS Web Infrastructure Platform Data Server
- SAS Environment Manager Server
- SAS JMS Broker
- SAS Web Application Server
- SAS Cache Locator

In **Figure 1**, the primary function of the node on the left is to run the SAS Metadata Server as well as the SAS Middle Tier. The node on the right provides the SAS Application Server tier with the SAS Web Infrastructure Platform Data Server as well as additional active instances of clustered SAS Middle Tier components.



**Figure 1. Minimum Configuration for High Availability**

**Notes:**

- Parentheses ( ) denote cold stand-by for failover.
- Chevrons < > denote warm stand-by, the slave node of a master/slave cluster.
- The SAS Web Server can be configured in active/active clustered mode following the instructions in *SAS 9.4 Intelligence Platform: Middle-Tier Administration Guide* in the Recommended Reading section.

**INCORPORATING A CLUSTERED SAS METADATA SERVER**

This scenario builds on the previous scenario by adding an additional node to provide two things:

- a clustered SAS Metadata Server
- additional compute resources for SAS servers

### *SAS Metadata Server Tier*

In this case, there is an active instance of the SAS Metadata Server running on each of the three nodes. SAS Grid Manager provides restart of each instance on the same node in the event of a failure of one of the SAS Metadata Server processes

In a SAS Metadata Server cluster, each node runs its own server process and has its own server configuration information, journal file, and copy of the repository data sets. The SAS Metadata Server instances are not using a shared configuration in this scenario.

### *SAS Application Server Tier*

The SAS Object Spawner and SAS OLAP Server are active on the third node with the majority of the compute resources of this node available to these SAS servers. SAS Grid Manager will provide restart on the third node or failover to the second node if the third node becomes unavailable. Deploying a shared configuration of the SAS Object Spawner and the SAS OLAP Server facilitates initial deployment as well as failover.

The SAS Environment Manager Agent and SAS Deployment Agent are actively running on each node in the deployment. If either fails, SAS Grid Manager will attempt to restart these services on the same node. Failover to another node does not apply in this case because there cannot be more than one instance of these services running on a node. Therefore the configuration of these components must be done on a per-node basis.

### *SAS Middle Tier*

The SAS Web Server is running in active-passive mode and is deployed onto a shared file system accessible by all, in this case two, nodes participating in the active-passive configuration. SAS Grid Manager will provide restart on the first node or failover to the second node in the event that the first node fails.

The SAS Web Infrastructure Platform Data Server, SAS Environment Manager Server and the SAS JMS Broker are all running in master-slave mode with a master instance on one of the nodes and a warm standby on the other node. If the master instance fails or the node it is running on becomes unavailable, the warm standby will assume the role of master.

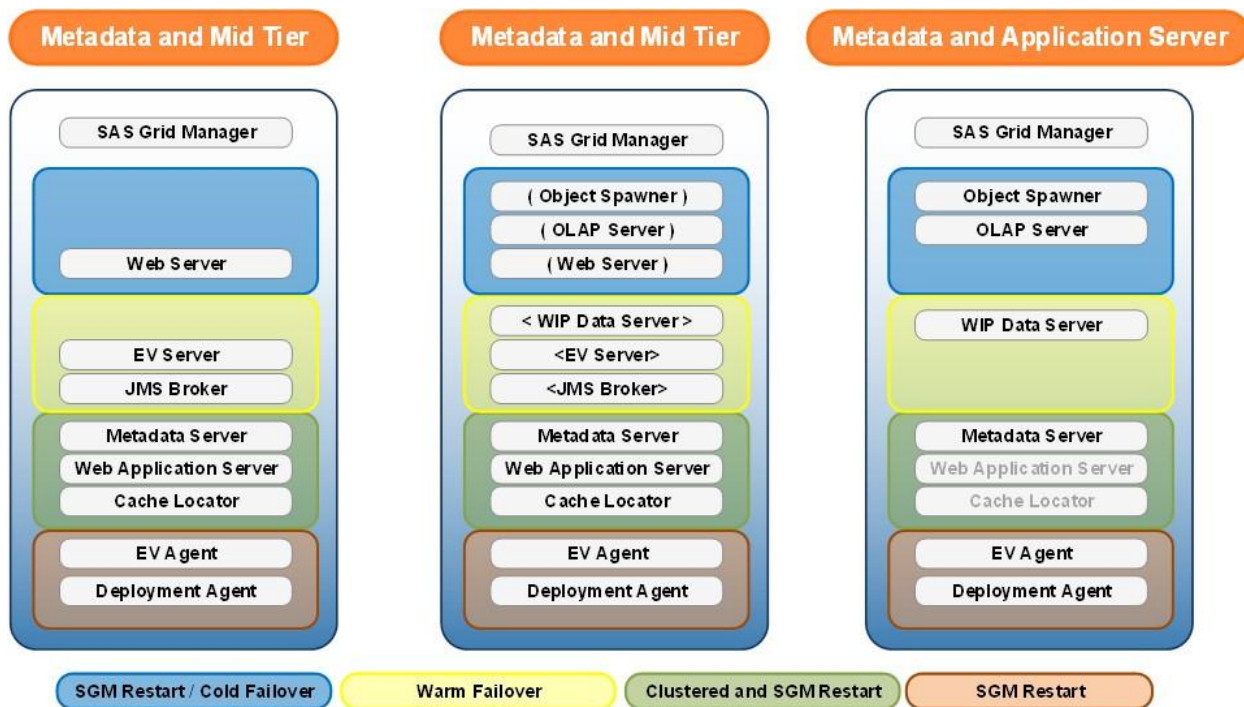
The SAS Web Application Server and SAS Cache Locator are clustered across the first and second nodes and SAS Grid Manager is used to restart any failed instance on the same node in an effort to constantly maintain two running instances of each. As an option, these services could include the third node in the cluster (grey text in **Figure 2**) at no additional licensing cost.

A shared configuration is recommended for the following components:

- SAS Object Spawner
- SAS OLAP Server
- SAS Web Server

An independent or per-node configuration is recommended for the following components:

- SAS Metadata Server
- SAS Environment Manager Agent
- SAS Web Infrastructure Platform Data Server
- SAS Environment Manager Server
- SAS JMS Broker
- SAS Web Application Server
- SAS Cache Locator



**Figure 2. Clustered Metadata Server**

**Notes:**

- Parentheses ( ) denote cold stand-by for failover.
- Chevrons < > denote warm stand-by, the slave node of a master/slave cluster.
- The SAS Web Server can be configured in active/active clustered mode following the instructions in *SAS 9.4 Intelligence Platform: Middle-Tier Administration Guide* in the Recommended Reading section.
- The mid-tier services can be clustered on the first two nodes only or also across the third node (grey text) provided adequate sizing has been performed.

## SCALING OUT THE SAS APPLICATION SERVER TIER

The next step is to scale out to two nodes dedicated to the SAS Application Server Tier and two nodes dedicated to the SAS Middle tier while maintaining the SAS Metadata Server cluster.

### *SAS Metadata Server Tier*

In this scenario the SAS Metadata Server is running clustered in active-active mode. SAS Grid Manager provides restart of an instance on the same node in the event of a failure of one of the SAS Metadata Server processes. However, because each node in a clustered metadata server has its own server configuration information, journal file, and copy of the repository data sets, the concept of failing over to a cold standby node does not apply in this case. The SAS Metadata Server instances are not using a shared configuration in this scenario.

### *SAS Application Server Tier*

In this scenario, the SAS Object Spawner and SAS OLAP Server are running in active-passive mode on the third node with SAS Grid Manager providing restart of both of these services or failover to the fourth node in the event that the third node fails. As an option, both the SAS Object Spawner and SAS OLAP Server can be run in an active-active mode (grey text **Figure 3**) on both nodes three and four. However, this would require additional SAS Integrations Technology and SAS OLAP Server licenses on a per-node basis which would require additional license fees. In either case, deploying a shared configuration of the SAS Object Spawner and the SAS OLAP Server facilitates initial deployment as well as failover.

The SAS Environment Manager Agent and SAS Deployment Agent are actively running on each node in the deployment. If either fails, SAS Grid Manager will attempt to restart these services on the same node. Failover to another node does not apply in this case because there cannot be more than one instance of these services running on a node. Therefore the configuration of these components must be done on a per-node basis.

### SAS Middle Tier

The SAS Web Server is running in active-passive mode and is deployed onto a shared file system accessible by all, in this case two, nodes participating in the active-passive configuration. SAS Grid Manager will provide restart on the first node or failover to the second node in the event that the first node fails.

The SAS Web Infrastructure Platform Data Server, SAS Environment Manager Server and the SAS JMS Broker are all running in master-slave mode with a master instance on one of the nodes and a warm standby on the other node for failover. If the master instance fails or the node it is running on becomes unavailable, the warm standby will assume the role of master.

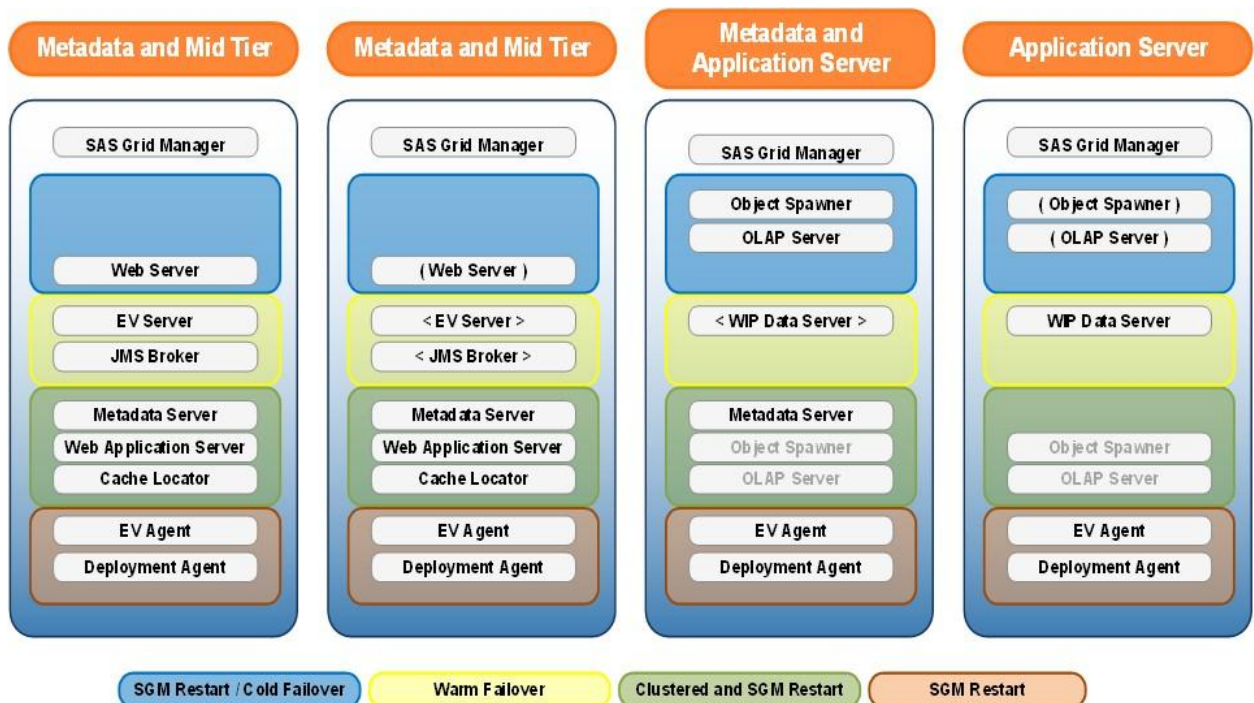
The SAS Web Application Server and SAS Cache Locator are clustered across the first and second nodes and SAS Grid Manager is used to restart any failed instance on the same node in an effort to constantly maintain two running instances of each.

A shared configuration is recommended for the following components:

- SAS Object Spawner
- SAS OLAP Server
- SAS Web Server

An independent or per-node configuration is recommended for the following components:

- SAS Metadata Server
- SAS Environment Manager Agent
- SAS Web Infrastructure Platform Data Server
- SAS Environment Manager Server
- SAS JMS Broker
- SAS Web Application Server
- SAS Cache Locator





### Figure 3. Scaling Out SAS Application Server Tier

#### Notes:

- Parentheses ( ) denote cold stand-by for failover.
- Chevrons < > denote warm stand-by, the slave node of a master/slave cluster.
- The SAS Web Server can be configured in active/active clustered mode following the instructions in SAS 9.4 *Intelligence Platform: Middle-Tier Administration Guide* in the Recommended Reading section.

## SCALING OUT TO MULTIPLE NODES FOR EACH TIER

The final scenario deploys each SAS tier onto multiple nodes to provide maximum performance and scalability for a large number of simultaneous users, workload and data.

#### *SAS Metadata Server Tier*

The SAS Metadata Server tier has a minimum of three nodes in order to deploy a clustered SAS Metadata Server. In this scenario the SAS Metadata Server is running clustered in active-active mode. SAS Grid Manager provides restart of an instance on the same node in the event of a failure of one of the SAS Metadata Server processes. However, because each node in a clustered metadata server has its own server configuration information, journal file, and copy of the repository data sets, the concept of failing over to a cold standby node does not apply in this case. The SAS Metadata Server instances are not using a shared configuration in this scenario.

#### *SAS Application Server Tier*

The SAS Application Server tier is a typical SAS Grid Manager compute tier with two or more nodes. The SAS Object Spawner and SAS OLAP server can be licensed and run on a single node (as shown in **Figure 4**) in the SAS Application Server tier making use of SAS Grid Manager for restart or failover in the case of a process or node failure. As an option, the SAS Object Spawner and SAS OLAP Server can be run on two or more of the nodes (grey text **Figure 4**) to provide scalability and performance as required by the workload. However, this would require additional SAS Integrations Technology and SAS OLAP Server licenses on a per-node basis resulting in additional license fees. In either case, deploying a shared configuration of the SAS Object Spawner and the SAS OLAP Server facilitates initial deployment as well as failover.

The SAS Environment Manager Agent and SAS Deployment Agent are actively running on each node of each tier in the deployment. If either fails, SAS Grid Manager will attempt to restart these services on the same node. Failover to another node does not apply in this case because there cannot be more than one instance of these services running on a node. Therefore the configuration of the SAS Environment Manager Agent must be done on a per-node basis. The SAS Deployment Agent does not require any configuration.

#### *SAS Middle Tier*

The SAS Middle Tier makes use of a minimum of two nodes. The SAS Web Server is running in active-passive mode and is deployed onto a shared file system accessible by all, in this case two, nodes participating in the active-passive configuration. SAS Grid Manager will provide restart on the first node or failover to the second node in the event that the first node fails.

The SAS Web Infrastructure Platform Data Server, SAS Environment Manager Server and the SAS JMS Broker are all running in master-slave mode with a master instance on one of the nodes and a warm standby on the other node for failover. If the master instance fails or the node it is running on becomes unavailable, the warm standby will assume the role of master.

The SAS Web Application Server and SAS Cache Locator are clustered across the first and second nodes and SAS Grid Manager is used to restart any failed instance on the same node in an effort to constantly maintain two running instances of each.

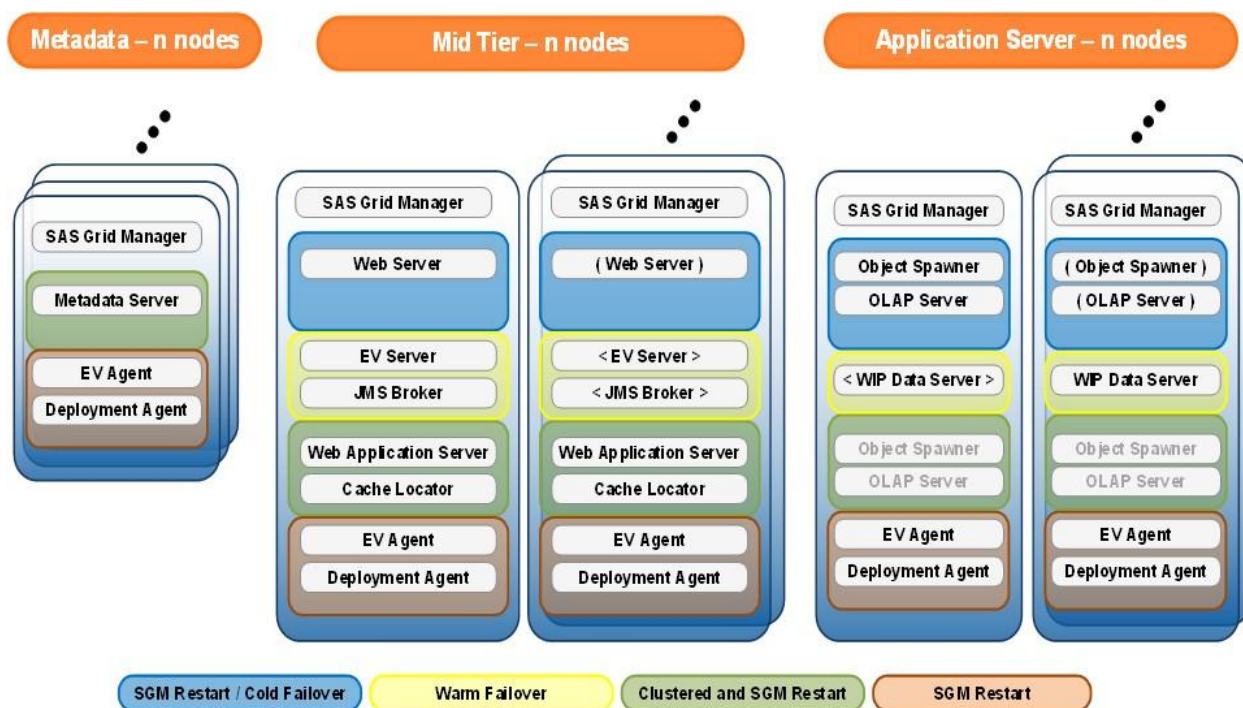
Nodes identical to the second node in this tier can be added to provide additional resources and increased performance as needed. Additional licenses are not required to scale out this tier and, therefore, it can be done at no additional cost from a licensing perspective.

A shared configuration is recommended for the following components:

- SAS Object Spawner
- SAS OLAP Server
- SAS Web Server

An independent or per-node configuration is recommended for the following components:

- SAS Metadata Server
- SAS Environment Manager Agent
- SAS Web Infrastructure Platform Data Server
- SAS Environment Manager Server
- SAS JMS Broker
- SAS Web Application Server
- SAS Cache Locator



**Figure 4. Multiple Machines Each Tier**

**Notes:**

- Parentheses ( ) denote cold stand-by for failover.
- Chevrons < > denote warm stand-by, the slave node of a master/slave cluster.
- The SAS Web Server can be configured in active/active clustered mode following the instructions in SAS 9.4 *Intelligence Platform: Middle-Tier Administration Guide* in the Recommended Reading section.

**SUMMARY**

You use SAS Analytics to gain insight into increasing amounts of data, to make mission critical decisions and to gain competitive advantage. You cannot afford downtime that interrupts your ability to access any part of your SAS

infrastructure. This paper provides best practices for leveraging the SAS Grid Manager failover and restart capabilities in conjunction with the clustering technologies supported by the SAS Metadata Server and SAS Middle Tier in order to maintain a highly available SAS 9.4 deployment. Of course you may choose to implement high availability for only some of the tiers or components in your deployment. You may also choose to use either SAS Grid Manager or clustering but not both. If you use SAS Grid Manager without clustering, you lose the load balancing capabilities provided by multiple instances of a particular service. If you use clustering without SAS Grid Manager you lose the automatic restart and failover capability if a service instance dies or the machine running it crashes. Also, there are other technologies available that provide high availability, such as hardware load balancers, that may be standard within your IT organization but which do not provide restart and failover capabilities. Therefore, the SAS best practice recommendation is to use SAS Grid Manager and clustering as the minimal set of technologies to provide a consistent high availability solution for all components within a SAS 9.4 deployment.

## ACKNOWLEDGMENTS

The authors would like to acknowledge Scott Parrish and Doug Haigh for their feedback and review of this paper.

## RECOMMENDED READING

- SAS Institute Inc. 2011. SAS Institute white paper. "High Availability with SAS Grid Manager." Available at <http://support.sas.com/rnd/scalability/grid/HA/GridMgrHAServices.pdf>
- SAS Institute Inc. 2013. *SAS 9.4 Intelligence Platform: Middle-Tier Administration Guide, Second Edition*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/documentation>.
- SAS Institute Inc. 2013. "Using Metadata Server Clustering." In *SAS 9.4 Intelligence Platform: System Administration Guide, Second Edition*. Cary, NC: SAS Institute Inc. Available at <http://support.sas.com/documentation>.
- SAS Institute Inc. 2013. "Understanding Server Load Balancing." In *SAS 9.4 Intelligence Platform: Application Server Administration Guide*. Cary NC: SAS Institute Inc. Available at <http://support.sas.com/documentation>.
- Li, Zhiyong, and Alec Fernandez. 2014. "Migrating SAS Java EE Applications from WebLogic, WebSphere, and JBoss to Pivotal tc Server." When they are available, see the *Proceedings of the SAS Global Forum 2014 Conference*. Cary, NC: SAS Institute Inc.

## CONTACT INFORMATION

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

Name: Cheryl Doninger  
Organization: SAS  
Work Phone: (919) 531-7941  
Email: [cheryl.doninger@sas.com](mailto:cheryl.doninger@sas.com)

Name: Zhiyong Li  
Organization: SAS  
Work Phone: (919) 531-9068  
Email: [zhiyong.li@sas.com](mailto:zhiyong.li@sas.com)

Name: Bryan Wolfe  
Organization: SAS  
Work Phone: (919) 531-5706  
Email: [bryan.wolfe@sas.com](mailto:bryan.wolfe@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.