

Paper 463-2013

Monitoring 101: New Features in SAS 9.4 for Monitoring Your SAS Intelligence Platform

Amy Peters, Bob Bonham, Zhiyong Li, SAS Institute Inc., Cary, NC

ABSTRACT

Ever needed an alert on SASWORK storage usage at 80 percent? Or known when a SAS® user account has been locked out due to failed login attempts? Or understand the memory and swap usage of a computer hosting the SAS Stored Process Server? What if you could see a complete listing of process resource consumption for all physical machines hosting a given SAS deployment? Learn about options for answering these questions, including new tools in SAS 9.4 that auto-discover software resources (including SAS servers and web application servers) of your platform for SAS Business Analytics. Discover how to use agents to collect metrics that reflect availability, performance, utilization, and throughput, giving you a more proactive understanding of the operational state of your SAS deployment.

INTRODUCTION

Each SAS system administrator or IT operations specialist is faced with the challenging task of monitoring, managing and forecasting the needs of software, hardware and systems. So much so that even the language of discussing a problem, event, or analysis can become rather complex. Presented in Figure 1 is a diagram that depicts the monitoring “continuum” over time and the need for dynamic monitoring, which is typically not persisted, recent < 3 days review of system usage via SAS® Environment Manager and long term “forensics” type of usage and capacity planning offered by the SAS® Audit, Performance, and Measurement Package and the SAS® IT Intelligence solution.

Overview of SAS Operational Monitoring Continuum				
	Real-Time	Operational (Detailed)	Operational (Summary)	Capacity Planning/Forensics
Focus: Usage & Consumption	Real-time Process Monitoring (not persisted)		Audit Performance Measurement (APM Data Sets)	
Focus: OS Metrics & Events		SAS Environment Manager		SAS IT Intelligence (Performance Database)
Goals/Tasks/Uses	<ul style="list-style-type: none"> Dynamic visualization of real-time activity 	<ul style="list-style-type: none"> Monitor health of the environment. Alerting Review Logs 	<ul style="list-style-type: none"> Provide “Context” for Operational Activities Configuration Change Control 	<ul style="list-style-type: none"> Understand usage patterns of SAS content and data Audit Security Changes Capacity Planning Hardware Maintenance
Time Scale	< 1 minute	1 minute → 3 days	3 days → 10 days	> 10 days

Figure 1 – Overview of SAS Operational Monitoring Continuum

First available with SAS 9.4, SAS Environment Manager combines VMware’s vFabric™ Hyperic® application monitoring framework with a collection of SAS designed technologies providing a unique, SAS-aware, out-of-the-box monitoring solution.

For 9.4 initial release, SAS Environment Manager's focus is serving as the operational monitoring solution with plugins for the administration, management, and monitoring of SAS technologies including the administration of the SAS Web Application Server. Additional functionality includes agent based, auto-discovery of resources, collection of resource consumption, monitoring of log events, alerting, and many ways to present and view the collected information. Over the 9.4 life cycle, SAS Environment Manager will be extended to provide a more robust, centralized SAS administration and monitoring. In the meantime, SAS® Management Console will continue to serve as the primary interface for SAS content administration.

SAS Environment Manager consists of the following components:

- a web-based administrative interface
- an agent on each managed host with technology and solution aware plugins
- a central server to manage and deploy the plugins
- a centralized operational datastore
- a collection of reports and graphs helping characterize resource usage across the whole operational deployment.

HOW DOES IT WORK?

The basic architecture of SAS Environment Manager is an agent process running on each platform in a SAS deployment communicating to a central management server. Agents monitor detected resources and periodically report resource metrics back to the server. The server is responsible for presenting a GUI for interacting with those agents, managing the data collected by the agents, distributing plugins, providing the ability to create alerts and escalation procedures based on collected metrics and graphing of the metrics provided via the agent plug-ins.

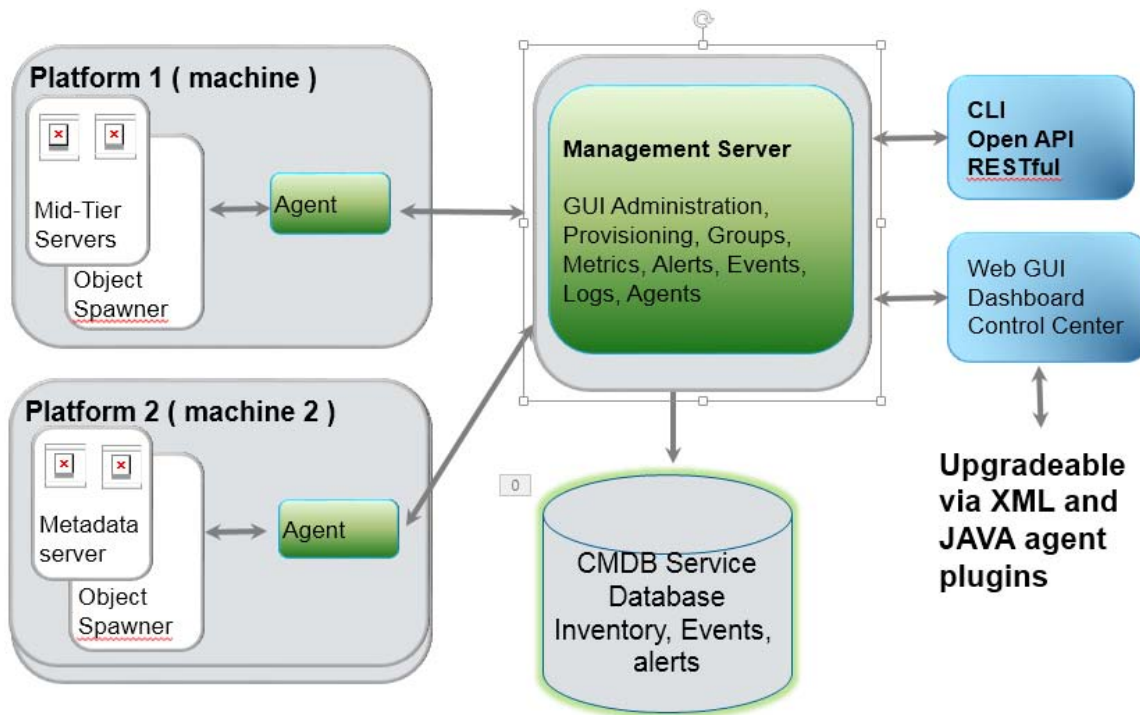


Figure 2 - Monitoring and Management Architecture

SAS Environment Manager can automatically perform simple system administration tasks to prevent and resolve a broad range of issues, interface directly with your network operations center (NOC), and provide notification of system and configuration file changes, escalation notifications, and alert workflow.

A SAS Environment Manager agent is a software process that runs on and interacts with a particular machine (platform). Each managed physical and virtual machine has its own agent. Each agent performs the following functions:

- auto-discovers the software components that run on its platform
- periodically re-scans its platform for configuration changes
- gathers resource metrics and confirms availability
- sends the inventory and performance data that it collects to the management server
- performs log scraping and monitors configuration changes
- enables you to perform resource control actions or script execution

The SAS Environment Manager database, also known as a configuration management database (CMDB), is a repository of information related to all resources known by SAS Environment Manager. You can monitor and manage only those resources discovered by the agents plug-in and accepted into the management environment. SAS Environment Manager stores its inventory information in a Postgres database and performs the following functions:

- provides facilities for managing your software
- displays resources in your deployment
- implements its own access control model
- detects when alerts are triggered
- performs alert notifications or escalation processes as defined.

DISCOVERING YOUR ENVIRONMENT

The SAS Environment Manager server is typically installed on a mid-tier platform. An agent is installed on each mid-tier/server machine in your deployment. The SAS Environment Manager server is typically configured to listen on port 7080. The agent auto-reports back to their server, go through a discovery phase of their installation, and provide operational and log details back to the server.

After login, the GUI presents an option to select a role, which has an associated dashboard. The dashboard is a landing page providing an at-a-glance operational view of the systems in the deployment. Select the Resources tab and you see the active, known, and monitored resources that have been auto-discovered and added into the inventory of your deployment. These resources can be

- platforms (physical or virtual machines, or logical, proxied representations)
- servers running on those platforms (database servers, web servers, web application servers, SAS servers, messaging servers, and guest operating systems)
- services running in those servers (examples include queues in Activemq, web applications in SAS Web Application Server, or database tables and indexes)
- logical groupings, application groupings or mixed groupings of platforms, servers, and / or services.

You can drill into the different resources to get status and answer questions such as these:

- What is the CPU usage trend for a given platform?
- What is the memory and swap usage of the SAS Stored Process Server?
- What are my IO rates for a given file system or disk?
- What is the thread count for a JVM?
- How many clients are currently connected to an Object Spawner?

Figure 3 presents a view of SAS Environment Manager's monitoring window plotting a platform's free memory, consumption, and load average for the past 8 hours. Note its green status bar showing availability.

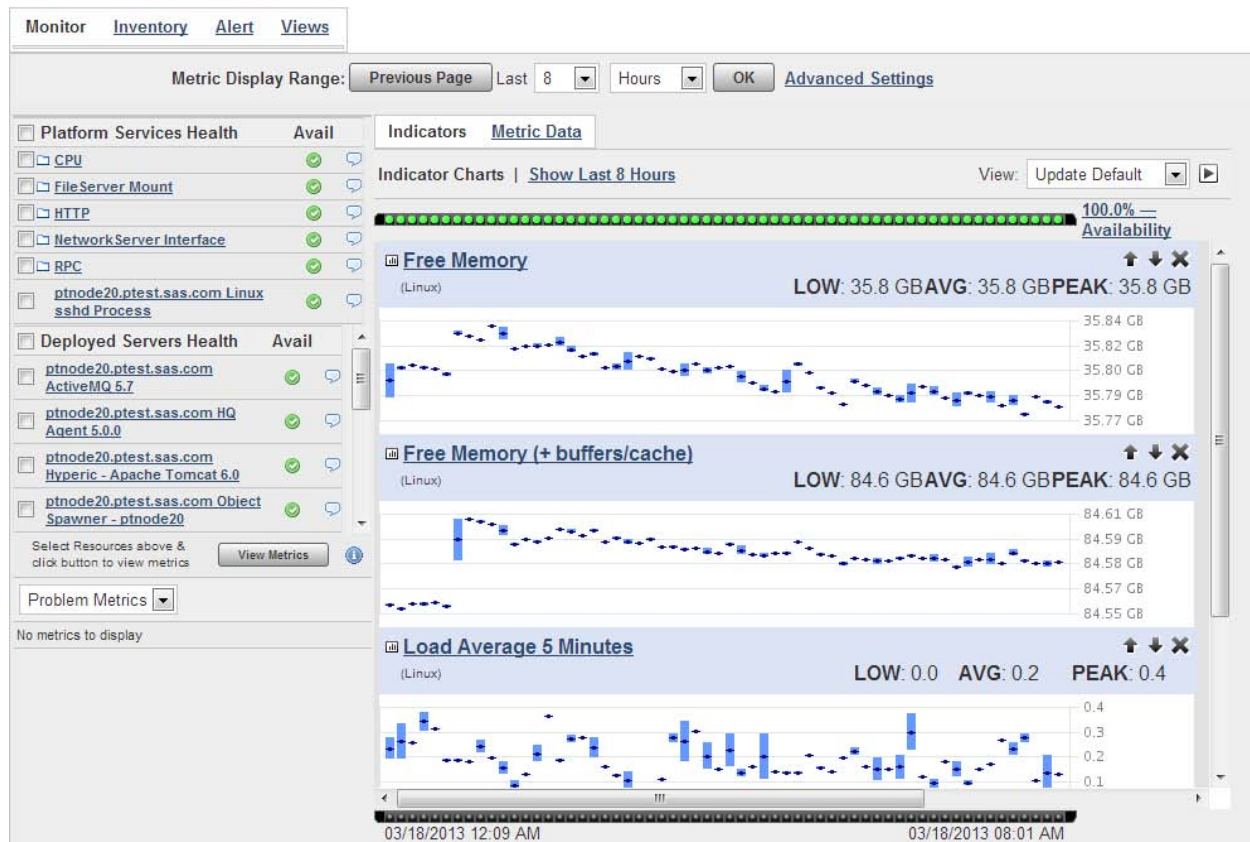


Figure 3 – SAS Environment Manager Monitoring Sample Showing System Free Memory & Load Average Over Time

MANAGING YOUR ENVIRONMENT

Plug-ins are extendable, agent-side, monitoring code or scripts that are designed to provide deep and broad-based insight into a SAS solutions deployment. They are the interface between SAS® Environment Manager and the platforms, servers, and services in your deployment. SAS Environment Manager is delivered with a full complement of plug-ins capable of recognizing most parts of a SAS deployment including plug-ins to monitor the system resources, middle tier, SAS Foundation Servers, and SAS Application Servers. The plug-ins are distributed automatically to the agents as they are installed and connect to the server. In addition, SAS Environment Manager provides a plug-in manager interface allowing you to add, update, remove, and synchronize plugins across remote agents, ensuring that your monitoring capabilities are consistent across your entire deployment infrastructure.

SYSTEM PLUG-INS

SAS Environment Manager is shipped with plugins to interface with and gather consumption metrics associated with the operating system (OS) usage of its hardware and a monitored processes usage of operation system resources. Examples from the OS include CPU processor usage, memory usage, network interface IO, disk IO, and file system activity. Processor usage includes each process's virtual memory size, resident memory size, CPU time consumed, threads, and open file descriptors. Together, these plugins bring together a wealth of resource usage consumption which, when viewed over time, helps administrators profile the resource usage patterns and needs of their software deployment.

Plug-ins are also shipped to support the data tier and mid-tier. For example, the SAS Web Application Server plug-in includes the following capabilities:

- control actions including starting, stopping, and restarting of a SAS web application.

- configure the operational control parameters and JVM run-time options such as heap size and debugging, default servers values for JSP and static content, and various SAS Web Application Server connectors.
- define the configuration of a group of SAS Web Application Server instances using command line utilities.
- deploy a SAS web application to a single or predefined group of servers.

SAS FOUNDATION SERVERS PLUG-IN

The SAS Foundation Servers Plug-in supports a variety of SAS servers including the SAS Metadata Server, SAS OLAP Server, SAS Object Spawner, SAS Stored Process Server, SAS Workspace Server, SAS Connect Spawner, and the SAS Pooled Workspace Server. SAS servers are automatically detected as they are deployed and started. Some servers are detected by querying the system process table. Running servers are matched up to their definitions in the SAS metadata repository and additional properties are filled in (install path, context name, metadata ID, version, and so on).

Servers spawned by the SAS Object Spawner are detected slightly differently. First, the SAS Object Spawner itself is detected. The list of servers that can be spawned by this object spawner is retrieved from the SAS Metadata Server. The logical and physical servers that the SAS Object Spawner is configured to launch become child resources of the SAS Object Spawner and have their own availability status and their own performance metrics.

Every SAS server resource supported by the plug-in implements two important metrics:

- Availability - A quick "sniff test" of a server to test whether it is running.
- Health - A deeper check that indicates whether the server can actually service requests.

Key process consumption metrics include:

- process virtual memory size
- process resident memory
- process page faults
- process CPU system time.

Metrics specific to each server are provided in addition to these. For example, SAS Metadata Server metrics also include current clients, journal space available, and journal transactions processed. The SAS Stored Process Server, however, can report on the number of connected, running, and waiting clients.

Where possible, the SAS Foundation Servers Plug-in implements server control actions. Servers can be started, stopped, restarted, paused, and resumed directly through the SAS Environment Manager interface.

Special server resources called SAS Home Directory and SAS Config Directory are also instantiated in the SAS Foundation Servers Plug-in. These services represent the presence of a SAS installation (regardless of which servers are configured or running) and facilitates functions such as monitoring disk usage for the deployment's SASWORK and changes to the deployment's configuration files. You can monitor configuration changes by using a wildcard regular expression to indicate what files should be included (*.cfg files, as an example), then if changes are made to any file meeting the condition, an event is created. The event includes the differences between the original file and the modified file.

SAS APPLICATION SERVER TIER PLUG-IN

The most significant difference between the SAS Application Server Tier plug-in and the SAS Foundation Servers Plug-in is that the SAS Application Server Tier Plug-in defines a new top-level resource (a 'platform') to represent the entire SAS environment across one or more machines. The SAS Foundation Servers Plug-in provides a per-machine view of the SAS environment. The 'parent' resource for the SAS servers is a specific platform (in this case machine). For the SAS Application Server Tier Plug-in, the top-level parent resource is a representation of the entire logical SAS Application Server tier of the SAS Foundation servers. Child resources are logical entities (contexts and SAS servers that are not associated with a context such as operating system services and object spawners).

KEEPING UP WITH WHAT'S GOING ON

Event management is the process that monitors events occurring in the SAS deployment. Events are defined as any detectable, discernible, or signature occurrence of a condition of significance. This definition is extremely broad and helps to highlight another significant feature of SAS Environment Manager – your ability to create the following alerts, to automate and manage problem detection and problem response:

- trigger and report an alert when a specified condition (or set of conditions) occurs.
- notify designated personnel when an alert is triggered.
- run designated resource control operations when an alert is triggered.
- track the resolution status of problems that are revealed by alerts.
- review alert and alert action history.
- establish an escalation process for an alert. An alert that has an associated escalation process has an acknowledged status that indicates whether designated responsible or interested parties are aware of the problem.

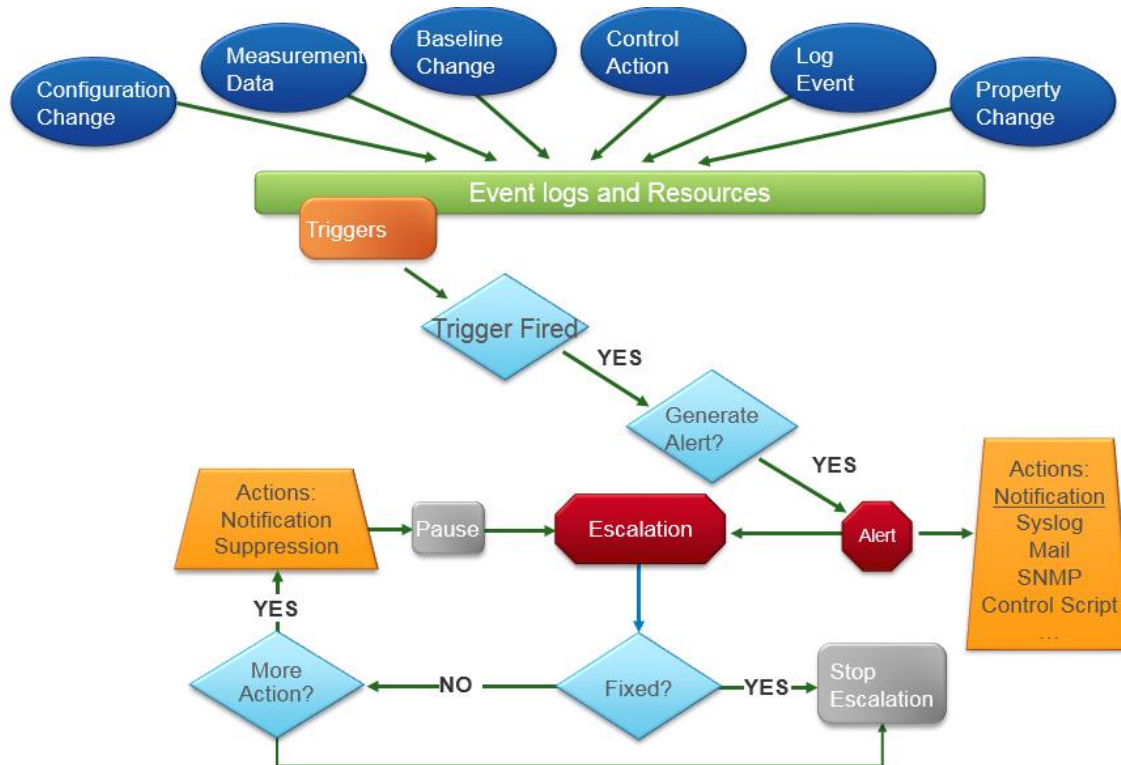


Figure 4 - Workflow of events, alert triggers, and possible escalation

SAS server log files are automatically monitored for new messages. By default, certain key warning messages in addition to all error messages will trigger an event in SAS Environment Manager and can be configured to perform an escalation alert. Log message pattern matching and alert definitions are configurable and can be customized by the operating administrator as-needed.

Let's look at an example of using the alerting capability in SAS Environment Manager to assist with tuning your SAS Web Application Server. Heap memory is one of the most important JVM parameters that you need to adjust. Setting it too high wastes the resource and increases the potential of having a high garbage collection time. Setting it too low can cause out-of-memory errors and increase the frequency of the garbage collection. Let's define two alerts to keep an eye on usage, so you can increase or decrease the max heap size based on that usage:

- Alert 1: Free heap memory low - if the free heap memory is less than 20 percent of the max heap size, an event will be generated that triggers an alert. You can set the priority of this alert as "high" and notify all super users when this happens.
- Alert 2: Free heap memory high - if the free heap memory is greater than 50 percent of the max heap size, an event will be generated that triggers an alert. You can set the priority of this alert as "medium".

If you get a lot of Alert 1 triggers, you need to increase the max heap memory. On the other hand, if you get a lot of Alert 2 triggers, you can consider decreasing the max heap memory to free up some resources.

Another good example of improving your service levels via alert triggers is the ability to define an event that monitors SASWORK storage consumption. As previously highlighted, the SAS Home Directory service contains the consumption metrics for SASWORK including both capacity (gigabytes used) and use percent. A common practice would be to trigger an alert if SASWORK usage exceeded 80% of available free space.

CUSTOMIZING YOUR EXPERIENCE

You can control what information is presented to you on your dashboard. Portlets are sections of the dashboard that provide a quick view of information like resource health or recent alerts. Portlets are available to:

- list all resources that have problem metrics and provides details including availability status, number of alerts per resource, number of times the metric has been out of bounds, and the most recent time that the out-of-bounds metric was collected
- show new and changed resources that have been auto-discovered and enable you to add them to your inventory
- indicate the availability of resources that you deem important
- display your favorite charts.

In addition to your personal dashboard, you have access to a distinct dashboard for each operational role including guest, administrator and config operations.

There are many ways to view metrics for your resources and change how they are collected. For example, you can change the collection interval rate for the availability metric to more accurately monitor a platform.

Custom plugins can be added to your system to discover, collect data from, and control resources of your choosing.

USING OTHER INTERFACES TO TAKE SAS® ENVIRONMENT MANAGER ACTIONS

In addition to using the graphical user interface of SAS® Environment Manager, you can directly access core operational functions and actions. This is useful for streamlining common implementation and configuration activities. For example, you can use command line tools to perform bulk updates to inventory and configuration data. The API also allows you to implement interfaces between SAS® Environment Manager and other management systems by writing a utility or use pre-built functionality that calls the APIs to extract inventory data into data you could analyze using SAS software.

There are several ways to interface to SAS® Environment Managers JAVA classes including the ability to call the JAVA methods directly or in a groovy script, access the methods via fully exposed RESTful web services (RESTful web services are Representational State Transfer style of web service access), or via a Command Line Interface (CLI) API package.

With the API you can:

- create, update, and extract data about platforms, servers, services, and groups
- create, update, and extract metric collection settings for resource types and individual resources
- define alerts for resource types and individual resources
- create and update native users and roles.

INTEGRATION WITH IT SYSTEMS

You might need to interface with and support complimentary technologies to meet the holistic needs of your enterprise. Some systems management tools are not very extendable; others require you to build your own plug-ins to collect the metrics you need, which is often a non-trivial and time-consuming task. SAS® Environment Manager supports extendable plug-in technology which gives you instant access to ready-to-run resource plug-ins, providing deep and broad-based insight to your SAS deployment without adding to your development workload. In addition, as updates and extensions to plugins are sometimes necessary as monitoring needs evolve, SAS Environment Manager's Plug-in Manager allows you to keep all your plugins synchronized automatically across all of your SAS resources; ensuring monitoring capabilities are consistent across your entire SAS application infrastructure.

This paper highlights several integration options available to administrators to help customize and interface with other systems including the ability to leverage open-source plugins, design custom plugins, CLI API's and RESTful web services. Going forward, more options are planned to help SAS Environment Manager share its understanding and knowledge of SAS solutions and technologies with common enterprise class network and control products.

CONCLUSION

Organizations are focusing their efforts on delivering new services and technologies at higher service levels using the latest technologies such as virtualization and cloud computing. To accomplish this, IT Operations specialists and SAS Administrators need the tools and automation software like SAS Environment Manager to provide the monitoring, control, and orchestration necessary to achieve operational excellence and enterprise readiness. SAS Environment Manager is targeted to become your web-based window into a complete and seamless operations and administrative experience for SAS Solutions and technologies.

RECOMMENDED READING

This paper was written prior to the production release of SAS 9.4. Features and implementation details might change in the production release. Check the product documentation for more up-to-date information, including the following:

- *SAS Intelligence Platform: System Administration Guide* (available at a future date at <http://support.sas.com>)
- Help for SAS Environment Manager

CONTACT INFORMATION

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

Name: Amy Peters
Enterprise: SAS Institute, Inc.
Address: SAS Campus Drive
City, State ZIP: Cary, NC 27513
Work Phone: (919)531-7325
Fax: (919)677-4444
E-mail: Amy.Peters@sas.com
Web: www.sas.com

Name: Bob Bonham
Enterprise: SAS Institute, Inc.
Address: SAS Campus Drive
City, State ZIP: Cary, NC 27513
Work Phone: (919)531-6847
Fax: (919)677-4444
E-mail: Bob.Bonham@sas.com
Web: www.sas.com

Name: Zhiyong Li
Enterprise: SAS Institute, Inc.
Address: SAS Campus Drive
City, State ZIP: Cary, NC 27513
Work Phone: (919)531-9068
Fax: (919)677-4444
E-mail: Zhiyong.Li@sas.com
Web: 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.