

# Log entries, Events, Performance Measures, and SLAs: Understanding and Managing your SAS® Deployment by Leveraging the SAS® Environment Manager Data Mart

Robert Bonham, Gregory A. Smith, SAS Institute Inc., Cary NC

## ABSTRACT

SAS Environment Manager is included within the release of SAS 9.4. This exciting new product allows administrators to monitor the performance and operation of their SAS deployments. What very few people are aware of is that the data collected by SAS Environment Manager is stored in a centralized data mart that is designed to help administrators to better understand the behavior and performance of the components of their SAS solution stack. This data mart can also be used to help organizations meet their IT Infrastructure Library (ITIL) reporting and measurement requirements. In addition to the information about alerts, events and performance metrics collected by the SAS Environment Manager agent technology, this data mart includes the metadata audit and content usage data previously only available from the SAS® Audit and Performance Measurement (APM) package.

## INTRODUCTION

Introduced with the SAS 9.4 Intelligence Platform release, SAS Environment Manager provides a framework for SAS administrators and IT operations staff to monitor the performance, health, and operation of their SAS deployment. Designed by SAS for SAS, SAS Environment Manager provides an integrated set of agent and server-based capabilities that are aware of SAS processes and are designed to monitor the status of a SAS deployment's servers, services, and web applications. These capabilities can help IT organizations in establishing the deployment's operational norms and service levels. As a documented part of the ITIL best practices, the establishment of customer service levels and resource consumption baselines are critical to maintaining the operational health and stability of a software deployment. By providing essential information about the usage of their IT resources, SAS Environment Manager can assist IT departments in providing a stable, yet scalable SAS deployment for both the current and future needs of their organization.

The SAS Environment Manager framework and kit enablement package is a valuable enhancement to the SAS Environment Manager. It extends the capabilities of SAS Environment Manager beyond the near real-time monitoring capabilities surfaced in the SAS Environment Manager web application user interface. At the center of these expanded capabilities is the SAS Environment Manager Data Mart. This is a specialized data store designed specifically to drive the reporting and analysis of the data that SAS Environment Manager collects. To facilitate this work, the data mart stores detailed data for a longer period of time than the SAS Environment Manager web application. In addition, it provides a fuller context by combining the standard SAS Environment Manager agent-collected metric information with information extracted from the log files generated by the various servers, solutions and web applications of a SAS deployment. This data is made available in a set of out-of-the-box reports and a set of reporting tools. However, organizations are not limited to this set of tools. The data mart enables the full range of analysis capabilities provided by the SAS Intelligence Platform to be applied to the collected data.

## THE IT INFRASTRUCTURE LIBRARY

The IT Infrastructure Library is a set of well-established practices for IT service management that focuses on aligning IT services with the needs of the business. At their core, ITIL's five guidelines map the IT service life-cycle, beginning with the identification of customer needs, through design and implementation of the service into operational best practices and finally, to the monitoring of customer focused services based on mutually established service level agreements.

At the heart of the ITIL services management framework are three key management objectives:

- align IT services with the current and future needs of the business and customers
- improve the quality of IT services
- reduce the long-term cost of IT services

The ITIL service design guidelines encompass process and procedures to assist with delivery and implementation of software, hardware, and services. Rather than focus on a specific area of technology, ITIL service design addresses the processes necessary to successfully stage, test, implement, and support a given hardware or solution.

## **THE BENEFITS OF USING THE DATA MART**

SAS Environment Manager, including the framework enhancements, can be used to address a variety of systems management challenges within an organization. It provides functionality that does the following:

- Creates a performance data mart (the SAS Environment Manager Data Mart), containing resource consumption and log records associated with the SAS deployment for an extended period of time.
- Collects data that is specific to a site's SAS deployment health and reporting requirements.
- Enforces consistent references to common measurements across disparate data sources.
- Leverages the domain knowledge of the SAS Research and Development staff directly through the bundling of key performance reports, system health measures, and automated alerting when the deployments health has been impacted.
- Provides a variety of out-of-the-box reports to accommodate the needs of the data administrators, performance analysts, and high-level information consumers focused on ensuring the operational integrity of the SAS Intelligence Platform.

## **THE SAS ENVIRONMENT MANAGER COMPONENTS**

SAS Environment Manager consists of a number of components that work together to monitor the health of a SAS deployment, collect performance and log data, and maintain the data mart. The key components include

- SAS Environment Manager server
- SAS Environment Manager agents
- SAS Environment Manager plugins
- SAS Environment Manager web application

### **SAS ENVIRONMENT MANAGER SERVER**

The SAS Environment Manager server runs as part of the SAS middle tier. It coordinates all of the other components and serves as the central collection point for all of the data collected from the agents. It is responsible for interacting with the SAS Web Infrastructure Data Server, which is a Postgres-based database server that serves as the back-end data store for the SAS Environment Manager web application.

### **SAS ENVIRONMENT MANAGER AGENTS**

The SAS Environment Manager agents are deployed onto every monitored server in a SAS deployment. These components are responsible for the activity that SAS Environment Manager performs on the servers. This activity includes collecting metrics, monitoring events, and executing control actions initiated by the users through the web application.

The SAS Environment Manager agent's function is to "wake up" periodically, collect performance metrics on the various resources in the environment, and transmit that information to the SAS Environment Manager server. The metrics that are collected will vary, depending on the resources the agent has determined to be running ("discovered") in the environment. Monitored resources range from high-level resources such as the operating environment to very specific resources like the garbage collection associated with an instance of the SAS Web Application Server.

### **SAS ENVIRONMENT MANAGER PLUGINS**

The SAS Environment Manager plugin technology provides a framework for SAS Environment Manager to monitor and interact with a wide variety of resources. Individual plugins provide the information that SAS Environment Manager needs when working with a specific type of resource. The plugins define how the SAS Environment Manager agents can recognize that a particular resource is present. Plugins also define both the list of possible metrics that can be collected for a resource and which of those metrics are enabled by default. SAS Environment Manager includes plugins that allow the SAS Environment Manager application to monitor resources specific to SAS, such as the SAS Metadata Server and the SAS Stored Process Server. Other plugins allow SAS Environment Manager to monitor other components that can be found in a SAS deployments, such as a Hadoop deployment or web servers.

In addition to defining the performance measurements that the agent should collect for a given type of resource, the SAS Environment Manager plugins can also define the types of events that should be recorded for those resources.

These events are surfaced through the SAS Environment Manager Event Center and Alert Center areas in the web application interface. There are four types of events that can be recorded:

- messages written to a monitored log file
- changes made to a monitored configuration file
- control actions initiated by administrators through the SAS Environment Manager web application
- alerts that have been triggered

The SAS Environment Manager framework and kit enablement package provides out-of-the-box definitions for many essential event alerts for excessive resource usage or errors. Furthermore, additional alerts can be defined to meet specific site requirements. Alerts can be triggered based on a specific measurement falling outside of some expected boundary or a specific message being detected in a monitored log file. The data mart includes tables with information about the events that were recorded and the alerts that have been triggered. These events allow administrators to look for unusual activity and correlate the events with the other resource usage metrics or log messages.

### **SAS ENVIRONMENT MANAGER WEB APPLICATION**

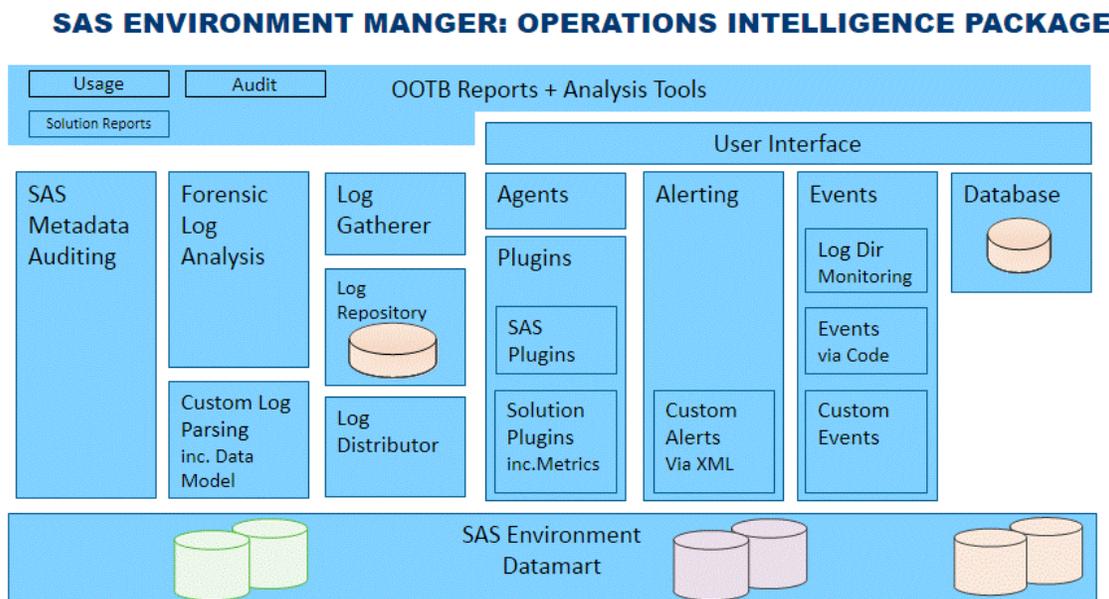
The SAS Environment Manager web application provides the primary user interface to SAS Environment Manager. This web application provides administrators a central place from which they can monitor and interact with their SAS deployment. From within the application, administrators can perform actions such as:

- configure custom dashboards
- define alerts
- monitor the health of resources, including seeing current metric values within the context of the last few days
- review key events and triggered alerts
- initiate control actions

For many resources, the application also provides views that surface additional resource-specific capabilities such as access to related reports or collecting additional information for SAS Technical Support.

## THE SAS ENVIRONMENT MANAGER FRAMEWORK ENCHANEMENTS

The SAS Environment Manager framework extends the capabilities of SAS Environment Manager beyond the valuable but somewhat limited capabilities surfaced through the SAS Environment Manager web application. These extended capabilities are designed to expand monitoring to include more than just what can be detected and collected by the SAS Environment Manager agent technology. The package makes the valuable information locked away in a deployment's various log files available, and integrates that information with the resource measurements collected by the agents into a report-ready data mart.



**Figure 1: Conceptual Overview of SAS Environment Manager Framework Package**

The key components of the SAS Environment Manager package include:

- Log Collector and Log Distributor web services
- log parsers
- data mart ETL programs and scripts
- a pre-configured data mart
- tuned alerting, log, and grouping capabilities
- a report center
- a solution kit framework
- event capability integration

## LOG COLLECTOR AND LOG DISTRIBUTOR WEB SERVICES

SAS Environment Manager collects important information through the parsing and analysis of various log files (such as server logs, solution logs, and web application logs) generated by both SAS and third-party resources in a SAS deployment. One challenge when working with log files is that they are often scattered across all of the servers in a deployment and are stored in a variety of directory locations. Periodically, the Log Collector (a component deployed with the SAS Environment Manager agent) identifies new log files (or the new messages in an open log file) and

copies this new log information to a central repository. Batch jobs execute each night and use the Log Distributor, which is a REST-based web service, to pull the new log information from this central repository for further processing.

## **LOG PARSERS**

The package's log parser technologies enable the SAS Environment Manager to leverage the log files generated by many of the components of a SAS deployment. Similar to the approach taken with SAS Environment Manager plugins, the SAS Environment Manager log parsing architecture is designed to be extensible and modular. Rather than one massive program that needs to handle every possible type of log file, individual parsers are designed to handle specific types of log files. For example, one parser might be responsible for parsing the log files generated by the SAS stored process server logs, while another handles web server logs or the log files associated with a specific SAS solution. The parsers allow SAS Environment Manager to understand the log file format and to extract the valuable bits of information from the log file. Regardless of which parser module processed the log file, the extracted important information is always standardized in the same way before being added to the SAS Environment Manager Data Mart.

This technology is derived from the SAS APM package, which has been available as an experimental download from the SAS support website for many years. Like its predecessor technology, the current technology extracts two types of information:

- metadata audit information, which is information about changes made to the objects within the SAS Metadata Server repository
- performance measurement information that captures information about how various SAS artifacts (such as content and servers) are being used and measurements of the resulting resource consumption.

## **DATA MART ETL PROGRAMS AND SCRIPTS**

The package includes the scripts and programs necessary to populate the SAS Environment Manager Data Mart. This code pulls the agent collected data from the back-end database that supports the SAS Environment Manager, and, like the parser technology, performs the transformations necessary to standardize the data before loading the data into the data mart.

## **PRE-CONFIGURED DATA MART**

Once enabled, the SAS Environment Manager Data Mart is populated with a core set of reporting data sets containing information about the resources most commonly of interest. This includes the various server platforms, file mounts, network interfaces, and SAS servers. These reporting data sets allow administrators to monitor the behavior of these critical resources immediately, right of the box.

## **TUNED ALERTING, LOG, AND GROUPING CAPABILITIES**

SAS Environment Manager provides event collection and message log monitoring of SAS IOM servers. Expanding upon this capability, the package includes extensive tuning of the metric templates (the default set of metrics enabled for specific types of resources) and event and alert definitions. It also enables auto-generation of resource groups designed to simplify monitoring and data collection, and highlighting of critical resource usage and consumption. Although inherently flexible and extendable, the default metric collection process has been refined, based on the expertise within SAS, to provide a "production smart" collection of logs and metrics necessary to provide the monitoring essentials of a successful deployment.

## **REPORT CENTER**

The Report Center includes tools to allow administrators to leverage the data in the SAS Environment Manager Data Mart. In addition, the package includes a set of pre-built out-of-the-box reports that use the data captured in the core default set of reporting data sets. The reports are useful both as reports that administrators can use to do their job, and as models of how the reporting tools can be used. These samples enable administrators to develop their own reports to meet their unique needs.

## **SOLUTION KIT FRAMEWORK**

The SAS Environment Manager Solution Kit is a framework for extending SAS Environment Manager to support SAS solutions and technologies. This framework provides a flexible mechanism for solution developers to integrate their solution into the SAS Environment Manager monitoring environment. A solution kit contains the tools needed to support a specific solution, including identification of resources, key metrics and measurements, definitions of events and alerts, and log parsers for the relevant log files.

## EVENT CAPABILITIES INTEGRATION

The SAS Environment Manager framework package provides components to facilitate integration of its event capabilities with other solutions. These components can be leveraged by the solution kits to support other SAS solutions. They can even be leveraged by organizations to surface events from their own, in-house, developed SAS solutions. In addition, the package provides the ability to surface SAS Environment Manager events to an organization's third-party monitoring solutions.

## DATA PROCESSING FLOW

Although it's not critical for data mart consumers to have a detailed understanding of the processing flows that the data follows, a high-level understanding can be helpful. As described previously, much of the data is collected by the SAS Environment Manager agents and passed to the SAS Environment Manager server. The SAS Environment Manager server stores the data in a database within the SAS Web Infrastructure Data Server. This database was designed specifically to drive the SAS Environment Manager web application and to ensure that this interactive application is responsive. Toward that end, the database maintains only a couple of days' worth of detailed data, and regularly summarizes it to the hourly, 6-hour, and daily level before discarding old data. This allows administrators to put the current data into some longer-term context.

The data mart, on the other hand, maintains detailed data for a much longer period of time and does not do any type of automatic summarization. It is designed to allow administrators to work with a larger set of data and permits more detailed analysis and reporting. The data mart ETL process is responsible for moving the detailed data from the database before it is rolled-off and auto-summarized. After extracting the data from the database, the ETL process stores the detailed data in a staging area. From there, additional ETL processing is responsible for transforming the raw data into a more usable form and ultimately making it available through the data mart. Like any data mart creation process, the goal is to transform the data from a format that is optimized for storage efficiency to one that is focused on facilitating analyst efficiency.

For the data mart data that is derived from forensic log analysis, the processing flow is slightly different. As outlined earlier, the Log Collector and Log Distributor services are responsible for ensuring that all important log data is available for processing. Periodically, this information is processed and the log files are parsed by specialized parsers that are designed to process specific types of log files. Once the important information has been extracted, it goes through a similar process of standardization and transformation as it is loaded into the data mart.

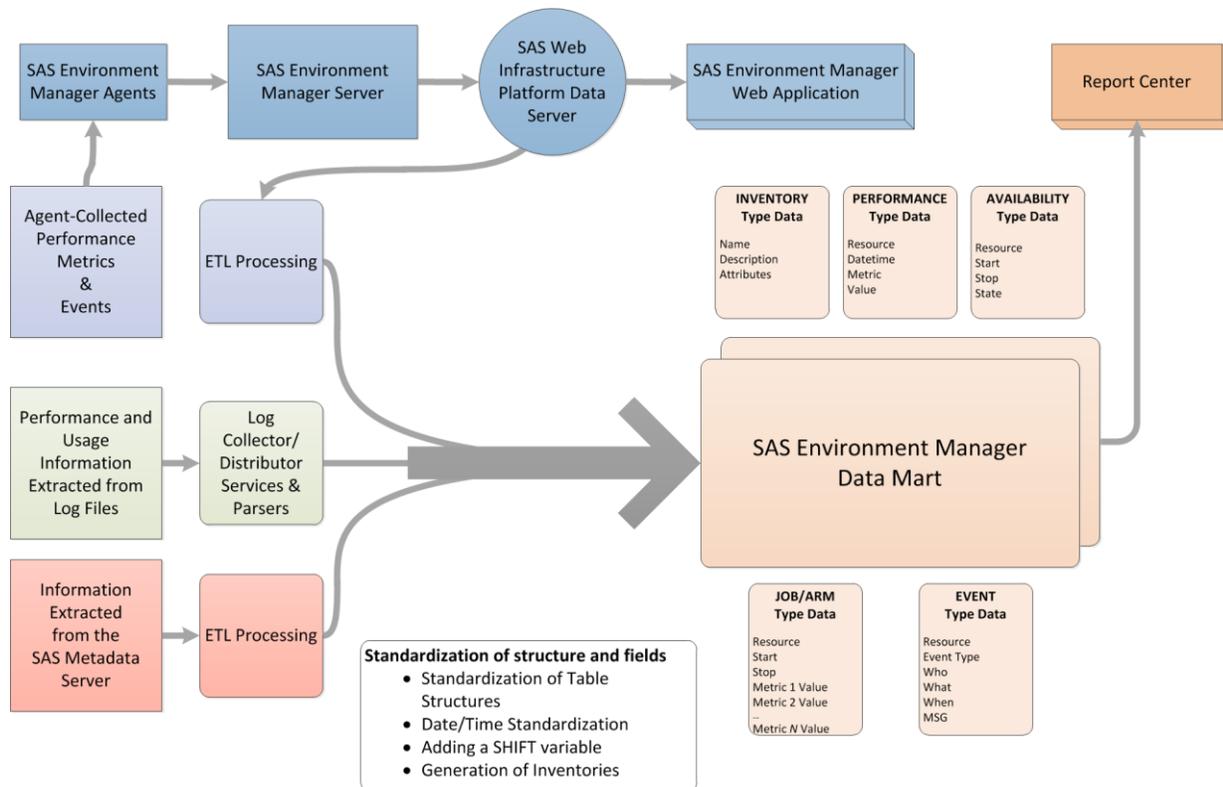


Figure 2: Data Processing Flow

## THE SAS ENVIRONMENT MANAGER DATA MART

Moving data from its raw form to the more usable and report- and analytics-friendly format of the data mart involves several types of transformations. These transformations include:

- standardization of key fields and table structures
- consistent handling of events and alerts
- maintaining consolidated inventories of resources and measurements containing both current and historical information
- retaining data for a fixed (but configurable) time-frame

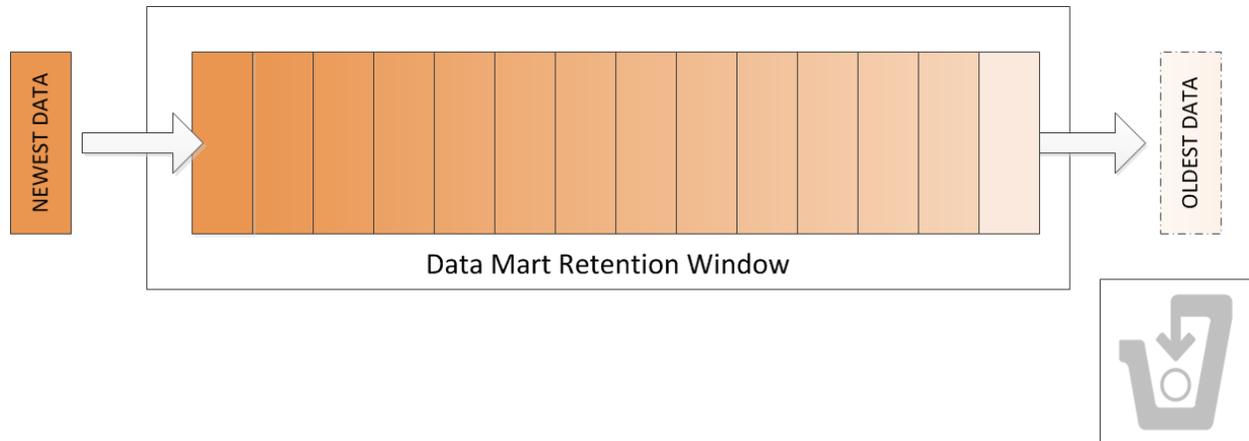


Figure 3: The SAS Environment Manager Data Mart Maintains a "Rolling" Window of Data

### STANDARDIZATION

Regardless of the original source of the data, the process of loading the data into the data mart enforces standardization. This standardization is designed to make it easier for administrators and analysts to work with data. Specific attributes that are standardized include:

#### Date and Times

**UCT** The records within the data mart are time-stamped to identify when the measurement was taken or the log message was recorded. This date time information is always stored in UTC (Universal Coordinated Time) to allow data from disparate sources (possibly collected from servers in different time zones) to be correlated.

**Time Zone Offset** Although storing data in a standard time zone makes it easier to work with data from multiple time zones, administrators will often need to be able to determine the date and time in its original time zone. To facilitate this, another field is added that represents the difference between "local" time and UTC. Obtaining the original time becomes simple task of adding the offset to the UTC time.

**Shift** One place where local time is important is in calculating work shifts. For example, many organizations divide the work week into three shifts: PRIMETIME, OFFHOURS, and WEEKEND. This shift information is often an important part of the analysis. As part of the standardization process, a field identifying the shift is added to the data mart tables. The mapping of date and time to shift is defined by the shift mapping scheme, which is stored in one of the data mart auxiliary tables. Organizations can customize the shift mapping scheme as needed to match their policies and processes.

#### Reporting Data Set Structure

**Standard Resource Identifier Information** Although SAS Environment Manager monitors a wide variety of resources and content, the reporting data sets in the data mart are structured in a consistent way. In addition to the time and time zone fields described above, the structure includes a standard set of fields to identify a resource regardless of where the resource exists in the resource hierarchy. This makes it easier to use the same tools and code for any resource.

**Measurement Metadata** Reporting data sets differ in the resource measurements they contain. However, standardization also makes it easier for analysts to work with these data sets. For example, the variable

label for every measurement field provides additional information about the measurement, including the units used and how the measurement is expected to change over time. Information about how the measurement can change could include whether the measurement trends upward over time, as a measurement like system uptime might, or if the measurement is dynamic, meaning it will bounce up and down, such as the measurement of CPU load might.

## RESOURCE AND MEASUREMENT INVENTORIES

The data mart includes a set of tables that contain information about each of the resources that are currently being monitored or have previously been monitored, and the measurements that are available for each resource. These tables make it easy to generate inventory reports and exception reports that identify discrepancies in measurements that are enabled across a particular type of resource.

## EVENTS AND ALERTS

Different types of events can generate different bits of information. For example, when a disk space usage alert is triggered, the important information might include the actual disk space being used. On the other hand, if an event was recorded because an ERROR message was written to a monitored log file, different information is important (in this case, the name of the log file and the message text). To make reporting easier, the data mart uses a standard set of fields for every event and alert, regardless of type. This includes fields to identify the following:

- resource involved
- affected component (such as the name of the log file or the name of the alert triggered)
- user involved
- time when the event occurred
- any additional information or messages that were generated

## OUT OF THE BOX REPORTING DATA SETS

A standard set of reporting data sets are configured to be available in the data mart after the data mart is initially deployed. These reporting data sets provide useful information about the resources that are most likely to be of interest to organizations. However, the data mart is designed to be very flexible and can be customized to store the data of most interest to a specific organization.

The standard, out-of-the-box reporting data sets are described below.

TABLE NAME	Description of Included Resource Types	Sample Measures Collected
HOSTPLATFORMS	Operating system instances that generally map to a single physical (or virtual) server	<ul style="list-style-type: none"> <li>• CPU TIME (IDLE, SYSTEM, USER)</li> <li>• MEMORY (FREE, SWAP, USED)</li> <li>• LOAD AVERAGES (1 Min, 5 Min, 15 Min)</li> <li>• NUMBER OF PROCESSES (RUNNING, SLEEPING, STOPPED, TOTAL, ZOMBIE)</li> </ul>
FILEMOUNTS	File systems mounted in an environment	<ul style="list-style-type: none"> <li>• CAPACITY (FREE, USED, TOTAL)</li> <li>• DISK ACCESSES (READS, WRITES)</li> </ul>
NETWORKINTERFACE	Network interfaces configured on every server in the environment	<ul style="list-style-type: none"> <li>• TRAFFIC IN/OUT (PACKETS/BYTES/BITS)</li> <li>• TRANSMIT ERRORS</li> </ul>

METADATASVRS	SAS Metadata Servers configured in the environment	<ul style="list-style-type: none"> <li>• NUMBER OF CLIENTS</li> <li>• MEMORY USAGE (CURRENT, HIGHEST)</li> <li>• JOURNAL TRANSACTIONS</li> </ul>
IOMSERVERS	SAS IOM servers that are spawned via the SAS Object Spawner, including Workspace Server and stored process server sessions	<ul style="list-style-type: none"> <li>• CURRENT CONNECTED CLIENTS</li> <li>• MEMORY SIZE</li> <li>• NUMBER OF PROCESSES</li> </ul>
WEBAPPSERVERERS	SAS web application server instances configured in the environment	<ul style="list-style-type: none"> <li>• MEMORY USAGE (HEAP, PERMGEN)</li> <li>• GARBAGE COLLECTION</li> </ul>

**Table 1: Reporting Data Sets**

## THE REPORT CENTER

ITIL describes the concept of a service desk to align IT services and resources. The goal of a service desk is to create documented and concise application and solution Service Level Agreements (SLAs). However, the successful instrumentation and monitoring of a production solution can be challenging, and will vary from site to site or even from deployment to deployment.

At the heart of a successful ITIL framed service desk is the establishment of service norms and the monitoring of a large number of resource consumption measures and log events, along with specific domain expertise. Providing the expertise of knowing what to examine is the key design principle of the SAS Environment Manager framework and kit enablement package. SAS Environment Manager couples this expertise with the performance metrics and log parsers, a data mart, and the analytic capabilities of SAS.

The SAS Environment Manager Report Center surfaces the integrated monitoring capabilities of SAS Environment Manager. It can help IT staff in creating an ITIL service desk representation of a SAS deployment. Even if an organization has not fully adopted ITIL processes, the capabilities of SAS Environment Manager can be immensely valuable to understanding the state of their SAS deployment.

SAS Environment Manager Report Center provides graphical and text reports to accommodate the needs of the data administrators, performance analysts, and high-level information consumers. It includes reports designed to present:

- forecast and project resource usage
- consolidation and evaluation of data from a large number of resources
- charts and drill down information to perform detailed analysis via plotting or tabular reports
- charting techniques for cumulative analysis
- charting and tabular reports for events analysis, used to derive the consumption based on event frequency
- simple, linear event correlation analysis
- forensics and exception analysis of events, log messages, and resource consumption

The SAS Environment Manager Report Center offers a collection of out-of-the-box resource consumption and error reports that provide both a health check and a historical record of usage. The Report Center leverages the SAS platform's stored process and stored process report technologies. The Report Center offers a single source of performance and troubleshooting forensics and resource usage of the SAS deployment's infrastructure. SAS Environment Manager Report Center offers a unique assessment of the resource usage and operational health of a deployment from the perspective of those who build the software.

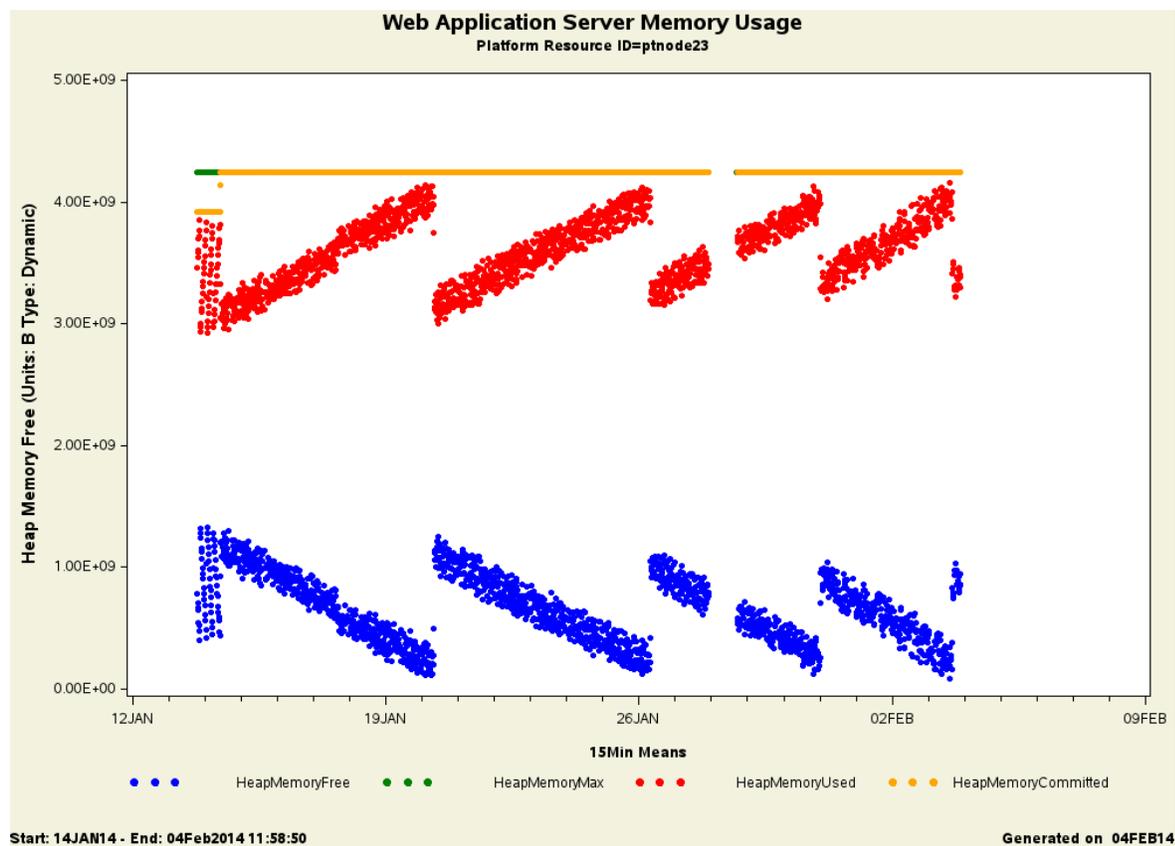


Figure 4: A Sample Report from the SAS Environment Manager Report Center

## CONCLUSION

SAS Environment Manager provides IT staff with a valuable tool for understanding the health and performance of their SAS deployments. The components of the SAS Environment Manager framework package build upon this foundation to provide a more comprehensive view of a SAS deployment. By combining the performance data collected by the SAS Environment Manager agent technology with the information extracted from a deployment's various log files, the SAS Environment Manager Data Mart gives IT staff an unprecedented view into behavior of their SAS deployment. This allows IT staff to diagnose complex resource interactions, discern long term usage trends, and perform exception analysis across several layers of the software stack. These capabilities allow IT staff to reduce downtime and increase reliability of their deployment. The SAS Environment Manager technologies matrix also provides a framework for SAS Research and Development domain experts to share their knowledge with customers. By identifying key metrics and providing definitions for event and alert conditions, these experts enable organizations to benefit from their experience. The solution kit framework allows SAS Environment Manager to monitor SAS solutions as organizations expand their use of SAS. At the same time, the components of the package also allow SAS to fit into an organization's IT ecosystem by allowing the specialized knowledge of SAS Environment Manager to be integrated into the organization's third-party monitoring environment. Regardless of whether an organization has adopted the ITIL approach, SAS Environment Manager provides them with a mechanism to effectively measure performance, establish baselines, and monitor the health of their SAS deployment.

## CONTACT INFORMATION

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

Robert Bonham, Gregory Smith  
100 SAS Campus Drive  
Cary, NC 27513  
SAS Institute Inc.  
[Bob.Bonham@sas.com](mailto:Bob.Bonham@sas.com), [Greg.Smith@sas.com](mailto:Greg.Smith@sas.com)  
<http://www.sas.com>

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

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