

Initializing SAS[®] Environment Manager Service Architecture Framework for SAS 9.4M7

Last revised June 24, 2020



Copyright Notice

All rights reserved. Printed in the United States of America. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, by any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc. Limited permission is granted to store the copyrighted material in your system and display it on terminals, print only the number of copies required for use by those persons responsible for installing and supporting the SAS programming and licensed programs for which this material has been provided, and to modify the material to meet specific installation requirements. The SAS Institute copyright notice must appear on all printed versions of this material or extracts thereof and on the display medium when the material is displayed. Permission is not granted to reproduce or distribute the material except as stated above.

U.S. Government Restricted Rights Notice. Use, duplication, or disclosure of the software by the government is subject to restrictions as set forth in FAR 52.227-19 Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

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.

Table of Contents

Copyright Notice	ii
Pre-Initialization	1
Overview.....	1
Considerations	3
Terminology	3
Pre-Initialization Checklist	4
Initialization Process	4
Overview.....	4
Initialization Steps.....	5
SAS Environment Manager Service Architecture Initialization.....	5
Framework Validation Pre-Check	5
Framework Initialization.....	6
Final Validation and Operational Health Check.....	6
Enabling ACM ETL Process	7
Initializing and Enabling APM ETL	7
Troubleshooting Guide	10
Appendix A: Loading Data Mart Tables into Visual Analytics	14
Appendix B: Special Considerations for AIX	16
Appendix C: Methods of collecting logs in a multiple machine environment..	17
Sharing logs between machines using NFS mounts.....	17
Other methods of copying logs to the server.....	19
Appendix D: Upgrading the Service Architecture Framework.....	20
User-Created Content	20
User Modified Configuration Properties.....	21
User Modified Data Mart Location	21
Running the ugipInit script	21
Upgrading the Service Architecture Framework in a multi-machine deployment	22
Log Collection for Multi-machine Configurations	22
Appendix E: Modifying APM to work with Common Directories for Logging ..	24
Appendix F – Configuration modifications made when APM is initialized	26
Metadata changes	26
Logging configuration changes.....	26
Application Context directory changes.....	27
Appendix G – Multiple <levelroot> directories on the same host.....	28
Overview.....	28
Limitations	28
Implementation	28
Specifying an override file	29

Pre-Initialization

Overview

SAS Environment Manager offers an integrated, operationally focused, collection of administration and monitoring tools. This guide covers the initialization of the SAS Environment Manager Service Architecture framework. Designed to augment the basic capabilities of SAS Environment Manager, this package extends and automates many of the monitoring, auditing and user activity logging activities consistent with the operational requirements of a Service Oriented Architecture (SOA). However, this document is only the initialization guide for the Service Architecture framework. For detailed instructions and usage information on either the core functionality of SAS Environment Manager or the Service Architecture framework, refer to *SAS Environment Manager User's Guide*. For updates, enhancements, and a look to future offerings, visit our website at support.sas.com/rnd/emi.

The SAS Environment Manager Service Architecture framework is an operationally oriented collection of tools, settings, process automations, reports, and data mart. The following outline depicts the framework and its components:

- **SAS Environment Manager Service Architecture Framework**
 - *Extended Monitoring* – this component configures SAS Environment Manager to implement specific monitoring best practices that have been identified for SAS deployments and to allow SAS Environment Manager to exchange events with other applications.
 - **Best Practices**
 - *Alerts* – a number of alert definitions related to various SAS resources are included in the framework. This allows SAS Environment Manager to alert administrators to issues and potential problems as they are detected.
 - *Additional Resources* – the framework also defines additional resources that allow administrators to better monitor their SAS deployment. For example, a set of “HTTP check” resources are defined and regularly ping various SAS web applications to monitor their health and response time.
 - *Additional Resource Groups* – the framework also defines new resource groups that allow administrators to monitor logical collections of resources together.
 - *Metric Collection Tweaks* – the framework adjusts the metrics associated with various types of SAS resources for which data is collected; these tweaks can include enabling/disabling specific metrics or adjustments to how often a metric is collected.
 - **Event Integration**

- *Event Export* – the framework’s event export capabilities allows SAS Environment Manager to generate events to be consumed by third-party monitoring products.
- *Event Import* – the framework’s event import capabilities provide a command line interface for generating events to appear in SAS Environment Manager. This interface can be called from both internal SAS solutions and external applications.
- **Data Mart** – the data mart components of the framework provide a standardized data repository for data collected by the SAS Environment Manager agents and other monitoring tools.
 - **Agent-Collected Metrics (ACM)** – this data stream populates the data mart with the data collected by the SAS Environment Manager agent technology. This is the same data that drives the SAS Environment Manager web application. However, the data mart extends the length of time the data is stored and restructures the data to make it easier to work with for reporting and analysis.
 - **Audit, Performance Measurement (APM)** –this data stream provides the data mart with data extracted from the various SAS server log files and with any Solution Kits (see below) that have been installed and deployed. These log files provide information on resource consumption, usage of SAS artifacts (e.g. servers, programs, reports, etc.), SASJob logs, Visual Analytics and metadata activity.
- **Report Center** – this component extends the SAS Environment Manager web application to provide access to a rich set of out-of-the-box reports. These reports leverage the data stored in the data mart to provide both a set of usable reports and examples of the types of reports organizations may want to implement on their own using the preferred SAS reporting and analysis solutions.
- **Solution Kits** – the framework is designed to be extensible through solution kits which provide the mechanism that allows SAS solutions to leverage the other capabilities of SAS Environment Manager. These solution kits allow administrators to better manage and monitor the operational health of the SAS solutions. For example, a solution kit might identify additional resources to be monitored, adjust the metrics that are collected for a given resource or generate events to be surfaced in the SAS Environment Manager web application. Solution kits can also extend the data mart to support reporting and analysis of the operational health of a SAS solution.
- **Federated Data Mart** - For organizations with more than one SAS deployment, the Service Architecture includes a federated data mart capability. When enabled, this feature creates a separate, federated Data Mart, containing data imported from multiple instances of the SAS Environment Manager Data Mart. It is important to understand each imported data mart instance continues to exist and will continue to be updated. Administrators can choose to export data from individual data marts (either on an ad-hoc or regularly scheduled basis). This exported data can then be moved to the SAS deployment hosting the federated Data Mart and imported into the federated Data Mart. Consolidating data from multiple deployments into the federated Data Mart

enables organizations to compare usage information across multiple deployments. These multiple deployments could represent a traditional triplet of DEV/TEST/PROD environments or aggregate data from individual, departmentally focused deployments.

- *Visual Analytics Autofeed* - Ability to auto-feed Data Mart tables into the SAS Visual Analytics auto-loading capability.

Considerations

After the SAS Environment Manager Service Architecture framework installation and configuration process has been completed, an additional initialization process is required before the SAS Environment Manager Service Architecture is fully enabled. For most of the framework's capabilities, initialization is accomplished by executing a simple command.

To enable the APM (Audit Performance Measurement) capabilities of the framework, the initialization process will involve running some additional scripts. If your deployment spans multiple machines, you will need to decide on an approach for bringing log files from all of the machines to a single landing zone directory for the APM ETL to process. This is typically accomplished using the SAS Deployment Agent to copy files from remote machines to the central landing zone directory. Consult the SAS Environment Manager User's Guide for additional information.

The APM log discovery and collection processing is designed to work with the standard log directory hierarchy located under the *LevN/...* configuration directory tree. This automated processing will not work if you selected the custom deployment option *Create Common Directories* for a single common logging directory at deployment time. See *Appendix E: Modifying APM to work with Common Directories for Logging* to work around this limitation.

For multi-machine SAS deployments, the SAS Environment Manager Service Architecture package will be installed on one machine, in one SAS Application context, and is considered a singleton (only one framework and one Data Mart for a deployment). For the Service Architecture to continue to function, the selected Application context cannot be load-balanced across multiple server machines. Consult the SAS Environment Manager User's Guide for more information.

A final consideration is the location, protection, and backup practices associated with the service architecture log archive. The **Archive** directory, located in the **emi-framework** directory, is the archive of all logs gathered by the masterAPMETL process. Best practices would suggest the copying of the zip'ed logs to some form of shared network storage or tape media depending on your IT organization's requirements for protection of access and activity of historical records.

Terminology

Throughout this document **<levelroot>** will be used as a placeholder for the fully-qualified location of the *LevN* directory within your SAS Configuration directory (e.g. /opt/SAS/Config/Lev1). When used in commands in the following instructions, replace instances of **<levelroot>** with that full (absolute) path.

The machine on which the framework is deployed is called the SAS Environment Manager Enablement Kit Server. This is often the same machine to which the SAS Environment Manager Server is deployed

but that is not always the case. The framework is deployed in the **emi-framework** directory under **<levelroot>/Web/SASEnvironmentManager**, and only the SAS Environment Manager Enablement Kit Server machine should have a directory with that name.

Pre-Initialization Checklist

Note: The SAS Environment Manager Service Architecture Framework requires that SAS be fully configured and running on all participating hosts before initialization can proceed.

Before beginning the initialization process, ensure that you have access to the following:

1. Your SAS deployment configured with a functioning Stored Process Server (common for most installations) . If unsure, validate via SAS Management Console.
2. Your SAS installer ID and password (external host credentials). This account will be used for the Service Architecture ETL processing.
3. The path to your SAS Environment Manager Solution Architecture framework directory. The machine on which the framework is deployed is called the SAS Environment Manager Enablement Kit Server.
4. The name of your SAS Application Server context(s) on all hosts in your deployment
5. Your SAS administration ID and password (e.g. sasadm@saspw).
6. If you have configured your SAS deployment to use certificates, SSL and the HTTPS protocol rather than the HTTP protocol, confirm that you are able to access all of the SAS web applications using the HTTPS protocol before completing the steps described below. The SAS Environment Manager Service Architecture framework leverages the configuration changes you have made for the SAS Intelligence Platform. If you do not encounter any problems accessing the SAS Web applications via HTTPS, the framework should work as well.
7. AIX users will need to manually add **sasevs@saspw** account to the SAS Environment Managers Super User Role prior to initializing the Service Architecture, see Appendix C.

Initialization Process

Overview

Configuration of the package is broadly defined in three phases or stages. The main phases of configuration are:

1. Pre check, validation of the initial deployment of SAS and SAS Environment Manager.
2. Validation of the SAS Environment Manager Service Architecture framework and the initialization of the enhanced monitoring bundle.
3. Enabling either the ACM or APM ETLs, including an additional initialization step for the APM ETL. All ETL processes are optional and can be enabled at any time after the framework has been initialized. However, one or more ETL's are required to construct the Data Mart.

Administrators may choose to modify various facets of the package (e.g. – changing the default reports, data retention windows, etc.). SAS recommends that this customization be performed only after initial installation, configuration, and validation of the package using the supplied configuration

recommendations. This will ensure that the supplied package capabilities are operating properly and will allow a site to retain a base level of capability as a functional foundation.

Initialization Steps

IMPORTANT: If you are migrating from, or are performing an *Upgrade In Place* from an earlier release of the SAS Environment Manager Service Architecture framework, see **Appendix D: Upgrading the Service Architecture Framework**. Log collection from remote machines is handled differently now and may require changes if you presently use NFS mounts/SMB shares to deliver logs to the SAS Environment Manager Enablement Kit Server.

The following sections describe the steps necessary to initialize the various components of the package.

SAS recommends that all SAS configuration files and the SAS Metadata repository be backed up and archived before configuring the SAS Audit, Performance and Measurement environment. Refer to the SAS 9.4 Intelligence Platform: System Administration Guide for additional details.

1. Change directory (cd) to the bin subdirectory of the emi-framework. This will be `<levelroot>/Web/SASEnvironmentManager/emi-framework/bin` (UNIX), `C:\<levelroot>\Web\SASEnvironmentManager\emi-framework\bin` (Windows). All subsequent references to the execute commands will assume that your current working directory is the **bin** directory of the framework's installation unless instructed otherwise.

SAS Environment Manager Service Architecture Initialization

Framework Validation Pre-Check

2. **Validate framework:** Change directory (cd) to the bin subdirectory of emi-framework. Run the validate script.

On UNIX:

```
$ ./validate.sh --level 2
```

On Windows:

```
validate.bat --level 2
```

A successful execution of the validate script indicates that the Service Architecture framework has been properly installed.

Framework Initialization

3. **Initialize Framework:** Change directory (cd) to the bin subdirectory of emi-framework (AIX users will need to manually add sasevs@saspw account to the SAS Environment Managers Super User Role, see Appendix C).

Use **-h** option to show help

Invoke script as follows.

On UNIX:

```
$ ./emi_init.sh -i
```

On Windows:

```
emi_init.bat -i
```

Review log messages for successful execution.

4. Restart your SAS Object Spawner(s).

On UNIX:

```
$ <levelroot>/ObjectSpawner/ObjectSpawner.sh restart
```

On Windows:

```
C:<levelroot>\ObjectSpawner\ObjectSpawner.bat restart
```

Final Validation and Operational Health Check

5. **Validate framework:** Change directory (“cd”) into the **emi-framework/bin** (UNIX) **emi-framework\bin** (Windows) directory. Run Validate script with level 3 diagnostics.

On UNIX:

```
$ ./validate.sh --level 3
```

On Windows:

```
validate.bat --level 3
```

6. Review the messages for errors. Correct any reported errors including file modifications, ownership errors, server operational status and invalid file permissions. If you made any corrections, re-run

emi_init.sh with -f option to force replacement. Then restart Object Spawner and revalidate. Consult your admin guide for more details.

7. Details associated with the framework initialization can be located in the log **emi-framework/Logs/emiInit.<datetime>.log** (UNIX) **emi-framework\Logs\emiInit.<datetime>.log** (Windows).

Enabling ACM ETL Process

1. Run the **emi_init.sh** (UNIX) **emi_init.bat** (Windows) script using the **-enable** option and specify **ACM**. Note: The other valid values, **APM** and **ALL**, require initialization prior to enabling, see Initializing and Enabling APM section below. Any of ETL processes can also be disabled by using the **-disable** <[ACM,APM,ALL]>

Use -h option to show help

Invoke script as follows with the enable ACM option.

On UNIX:

```
$ ./emi_init.sh --enable ACM
```

On Windows:

```
emi_init.bat --enable ACM
```

Review log messages for successful execution.

2. Repeat the step **Final Validation and Operational Health Check** completed earlier.

Initializing and Enabling APM ETL

The APM subsystem modifies properties and configuration values in the SAS deployment to enable additional logging for Application Response Measurement (ARM) metrics and for audit records. For a description of the properties and configuration items that are modified, see *Appendix F – Configuration modifications made when APM is initialized*.

1. Verify that your SAS Metadata Server is running, then invoke the **./apm_init.sh** (UNIX) or **apm_init.bat** (Windows) script.

On UNIX:

```
$ ./apm_init.sh
```

On Windows:

```
apm_init.bat
```

Logs will be generated in the **Logs** subdirectory under the emi-framework root directory

(**apmInit_[datetime].log**). Check this log file for ERROR messages. The `apm_init` command issues a series of SAS Environment Manager *control actions* to each deployed SAS Environment Manager Agent, telling the agent to run the local **logDiscovery/ConfigureAPM** process. If the `apm_init` command indicates a failure of one of these control actions, check the most recent generated logs on that particular host under the SASEnvironmentManager **emi-client/Logs** directory (**logDiscovery_[datetime].log**).

As a fall-back, in case the SAS EV `apm_init/control` actions approach doesn't work, you can log into each monitored machine in the environment as the **SAS Install** user and issue the following commands:

```
$ cd <levroot>/Web/SASEnvironmentManager/emi-client/bin
$ ./log_discovery.sh --apmInit
```

Errors will be reported to your console and in the most recent `logDiscovery.log` in

```
<levroot>/Web/SASEnviromentManager/emi-client/Logs/
```

NOTE: By default, logs from each remote host are copied to the landing zone on the `emi-framework` host through the use of SASEnvironment Manager Agent and SAS Deployment Agent technologies. Alternative techniques are discussed in [Appendix C](#).

2. **Restart SAS Application Servers:** For each machine in your deployment, restart the SAS Metadata Server, Object Spawner, OLAP Server, and Connect Spawner (on each machine, if deployed) to load the new logging configurations. **DO NOT RESTART** the SAS Environment Manager Server (by default, `sas.servers restart` will restart the SAS Environment Manager Server). Ensure that no other users are working in this SAS Server environment, then invoke:

On UNIX:

```
$ <levelroot>/SASMeta/MetadataServer/MetadataServer.sh restart
$ <levelroot>/ObjectSpawner/ObjectSpawner.sh restart
$ <levelroot>/ConnectSpawner/ConnectSpawner.sh restart
$ <levelroot>/SASApp/OLAPServer/OLAPServer.sh restart << Repeat for each App
Context >>
```

On Windows:

For Windows, the SAS EBI servers are typically deployed as Windows Services. As a Windows Administrator User, use Windows Start -> Control Panel -> Administrative Tools -> Services to bring up a services control panel. Scroll through the window to the entries that start with the pattern 'SAS [Config-'. Find the Metadata Server and click Restart. This will automatically restart all dependent SAS Foundation Servers on this machine, and restart the Metadata Server. On a multiple-machine deployment, you may have other servers on other machines, and will need to restart those as well using the services control panel on those machines.

```
SAS [Config-Levl] SASMeta - Metadata Server
```

SAS Management Console: Validate that your SAS deployment is working properly. Use the Server Manager Plugin in the SAS Management Console, then validate each supported server (right-mouse click on each server name and select “Validate” if available). The validation step will cause the servers to generate some initial records in each server’s associated Logs or PerfLogs directory, for example, `<levelroot>/SASApp/WorkspaceServer/Logs` or `<levelroot>/SASApp/WorkspaceServer/PerfLogs`.

3. Enable APM ETL

Run the `emi_init.sh` (UNIX) `emi_init.bat` (Windows) script using the `-enable` option selecting **APM** or **ALL**

(Note: Any, or all, ETL processes can later be disabled by using the `-disable <[ACM,APM,ALL]>` option.)

Use `-h` option to show help

Invoke script as follows.

On UNIX:

```
$ ./emi_init.sh --enable APM
```

On Windows:

```
emi_init.bat --enable APM
```

4. Repeat the step “**Final Validation and Operational Health Check**” completed earlier.

This completes the initialization of your SAS Environment Manager Service Architecture framework, enhanced monitoring bundle and selected ETL processes. Please consult the *SAS Environment Manager 2.5: User’s Guide* for additional information.

SAS Environment Manager, its reports, and supporting data are provided to assist in the troubleshooting and performance tuning of SAS environments. Although audit, access, and connectivity data and reports are provided, SAS does not provide a guarantee as to their completeness.

If you require more extensive and comprehensive audit, access, and connectivity data, please contact SAS Professional Services for assistance.

Troubleshooting Guide

Below is a listing of the most common problems with SAS Environment Manager Service Architecture including a brief explanation and description of their possible resolution. Consult the *SAS Environment Manager 2.5: User's Guide* for additional details or contact our Technical Support Department <http://support.sas.com> or consult our website <http://support.sas.com/rnd/emi>. Most service architecture framework commands can be re-executed with a -f (force) option. However, some specific commands will re-initialize the Data Mart, wiping the data.

Issue/Error Message	What it means/Resolution
Unable to connect to the SAS Environment Manager Server	Confirm SAS Environment Manager server is running. Connect to the SAS Environment Manager web application; it should be available via a URL like this: http://myserver:7080 .
[program] failed with return code [rc]	All service architecture framework logs are stored in the emi-framework/Logs directory. Review specific log for detailed associated with the reported error.
SAS plugins were not found in SAS Environment Manager	Confirm sas-servers plugin has been installed via SAS Environment Manager Plug-in Manager (not common) . Contact Technical Support for additional instructions.
SAS EMI Framework plugin is not loaded	Confirm emi-framework plugin has been installed via SAS Environment Manager Plug-in Manager. (not common). Contact Technical Support for additional instructions.
[user] is not a member of the Super Users Role	Confirm [user] is a member of the SAS Environment Manager Super Users Role, Logon to SAS Environment Manager, select Manage -> List Roles. If it is not a member, follow Appendix C instructions for adding sasevs@saspw to the Super User Role as described for AIX users.
Local SAS Environment Agent is not properly initialized	Confirm that on the local and host machines, SAS Environment Manager agent is running properly. Restart if necessary. Logon SAS Environment Manager via port 7080 or access <code><levelroot>/Web/SASEnvironmentManager/agent-5.8.8-EE/bin. hq-agent.(sh/bat)</code> restart.

<p>The script [scriptname] failed to compile</p>	<p>[scriptname] has been edited/corrupted locally. Contact Technical Support for a replacement/remove edits.</p>
<p>SAS Metadata server [servername] failed gyByName query.</p>	<p>Metadata server did not response to a “ping” request to its IP address. Run sas.servers start or run MetadataServers.(sh/bat).</p>
<p>Error: rc = [] running master_acm_etl.(sh/bat) or master_apm_etl.(sh/bat). Check log and return initialization with -f</p>	<p>Error detected in one of the etl processes. Review log message in the emi-framework/Logs directory for the etl process reporting the error. Correct error. Each etl can be executed separately to confirm status. Possibly causes could be inadequate storage, permissions, or operational status of SAS Environment Manager server or SAS Metadata server.</p>
<p>Request to enable APM with APM Initialized = false</p>	<p>APM ETL must be initialized prior to enabling. Consult documentation, then run apm_init.(sh/bat).</p>
<p>APM is not initialized/There was a problem validating external APM files</p>	<p>Issue with the APM ARM enablement. Rerun bin/apm_init.(sh/bat) with -f option.</p>
<p>APM has not been fully initialized. The file {...}/emi-framework/Conf/log_definitions.json needs to be updated (line <i>nn</i>).</p>	<p>For multi-machine deployments, edit the emi_framework/Conf/log_definitions.json file to specify local paths to mount points for the SAS <levelroot> directories on remote machines. This is described in 5. For multi-machine deployments, in section <i>Initializing and Enabling APM ETL</i> of this document.</p>
<p>INFO: (logarmprocesslog) Partial arm log. Source [(logfile)] does not appear to have been processed, ignoring log file</p>	<p>This info message will occur if you present to the log collection/log processing a partial log from a different SAS Deployment (missing header info) or if emi_init -resetDB -f was executed without stopping and restarting the SAS Application Servers. Restarting the servers will force the recreation of rolling logs including the required header information.</p>
<p>Messages similar to the following appear in the UpdateInstructions.html file generated during the upgrade:</p> <p style="text-align: center;">SAS Environment Manager Service Architecture</p> <p>A problem was encountered while updating the SAS Environment Manager Service</p>	<p>If the Service Architecture framework has <i>never</i> been initialized or enabled in your deployment, there is no need to perform an upgrade of the framework. In this situation, you can safely ignore these messages.</p>

<p>Architecture. Review the log <i>/config/Lev1/Logs/Configure/evmkitsvr_updateContent_2015-04-01-08.03.47.log</i> for more information on the problem. After addressing any issues, re-run the upgrade script from the command line. Refer to the appendix entitled Upgrading the Service Architecture in the Service Architecture Quickstart Guide for instructions on re-running this script. Please note that the appendix also outlines additional manual steps that may be needed to ensure your deployment of the Service Architecture remains fully functional.</p>	<p>If the Service Architecture has been initialized and your SAS deployment includes multiple machines, for example, one serving as the server tier and another for the middle-tier, the automated upgrade-in-place of the Service Architecture is likely to fail and generate these messages. In these cases, the solution is to ignore the error or warning messages and complete the upgrade process on all of the machines in the deployment. Once the rest of your SAS deployment has been upgraded and brought back on-line, you can upgrade the Service Architecture. To do so, run the <i>ugipInit</i> script as described in Appendix D “Upgrading the Service Architecture Framework”.</p> <p>If these messages are encountered and neither of the two previous situations apply, you should review the log file identified in the message. In addition, two other log files, with the name <i>ugipInit_{TimeStamp}.log</i> and <i>ugipInit_{TimeStamp}.sas.log</i> may be available in the Logs directory of the Service Architecture root directory. Review them for additional information on the problems encountered.</p>
<p>Duplicate listing of the reports displayed in the Report Center window and have performed a migration from SAS 9.4M2.</p>	<p>The duplicate reports can be ignored or can be deleted one-by-one via the SAS Management Console or the entire SAS Environment Manager folder tree under the Products folder can be deleted and reloaded by running <i>emi_init.sh -f -i</i></p> <p>Deleting the folder tree must be done with care. The deletion is not reversible. Best practices recommends backing up your metadata server prior to deleting the folder via SAS Management Console.</p>
<p>After updating host name of the SAS Stored Process machine with SAS Deployment Manager, the Report Center will not open.</p>	<p>SAS Deployment Manager currently does not automatically update <i>sasev.properties</i> path to the stored process server</p> <p>To resolve, updated the <i>sasev.properties</i> file in the <i>emi-framework</i> directory by replacing the current hostname in the <i>stp.url</i> with the new hostname of the server running the stored process server</p>

	<p>as follows: stp.url=http://<new hostname>:7980/SASStoredProcess by replacing the example string <new hostname> with new hostname excluding the angled brackets.</p>
--	--

Appendix A: Loading Data Mart Tables into Visual Analytics

In addition to the reporting tools and reports available in the Report Center, the data in the SAS Environment Manager Data Mart can be used with other SAS reporting and analysis tools. This includes analyzing and reporting on the data with SAS Visual Analytics. The **emi_init** script can be used to enable an automatic feed of the data mart data to SAS Visual Analytics. When enabled, this feed will copy the data mart data to a specific location, the VA autoloading dropzone, after the regular ETL processing has completed. Refer to the *SAS Visual Analytics: Administration Guide* for more information about how to enable its autoloading capabilities to make the data that has been copied to the drop zone available for analysis.

IMPORTANT: Please be sure the Visual Analytics autoloading drop zone location is accessible from the machine on which SAS Environment Manager Enablement Kit Server is running and on which the data mart ETL processes run. This is especially important if your deployment includes multiple machines providing compute servers.

For example, if your deployment includes both SAS Visual Analytics and another SAS solution, the SAS Environment Manager Enablement Kit Server *may not* be located on the same machine as SAS Visual Analytics. Since the Visual Analytics autoloading drop zone *will be* located on the SAS Visual Analytics server machine, it is important to ensure it is also available to the SAS Environment Manager Enablement Kit machine as well.

If that is not the same machine, the drop zone location could be made available via NFS mounts or Windows Shares. This is similar to the requirements for making log files available across machines discussed earlier in the section on enabling APM, and outlined in [Appendix C](#).

If you need to determine the path to the Visual Analytics autoloading drop zone, do the following:

1. Launch the *SAS Management Console* and sign in.
2. Navigate to the *Data Library Manager* plugin and expand the folder structure to see the list of libraries
3. Select the library labeled “*Environment Manager Data Mart LASR*”, right-mouse click on it and select Properties.
4. Click on the tab labeled “*Extended Attributes*” in the “*Environment Manager Data Mart LASR Properties*” dialog.
5. The Visual Analytics autoloading drop zone location is shown as the value associated with the field named “*VA.AutoLoad.Location*”.
6. This is the path that the data mart ETL processing will use as the destination when copying over data mart data. The copy operation is executed as a file copy. If this destination is not accessible from the SAS Environment Manager Enablement Kit Server machine, the copy will fail and data will not be available for use with SAS Visual Analytics. If this destination is not accessible from the SAS Visual Analytics server machine, then the autoloading task will fail to import the data.
7. Specify a path that is accessible from both the SAS Environment Manager Enablement Kit Server machine and the SAS Visual Analytics server machine, then click the OK button to close the dialog.
8. If a new path was specified, be sure that path and the following subfolders exist: Append, Formats, Logs, Unload

To enable the data feed for SAS Visual Analytics, follow these steps:

1. If your deployment includes multiple SAS solutions or multiple compute server machines, confirm the drop-zone is accessible as described above. Even if your deployment does not include multiple compute server machines, you may want to confirm the drop zone is accessible.
2. Change directory into the **emi-framework/bin** (UNIX) **emi-framework\bin** (Windows) directory.
3. Run the **emi_init.sh** (UNIX) **emi_init.bat** (Windows) script using the **-vafeed** option and specifying a value of **ON**.
4. Note: To later disable this feature, run the same command and but specify **OFF** for the **-vafeed** option.

Appendix B: Special Considerations for AIX

SAS Environment Manager Service Account

On AIX systems, the SAS Environment Manager Service Account (sasevs@saspw) needs to be granted the special permission of the SAS Environment Manager Super User Role. To do this, complete the following steps:

- Logon to SAS Environment Manager, select tab **Manage**
- select **List Roles**
- select **Super User Role**
- In the Assigned Users section, middle of the page, select **Add to List**
- select choice box next to the sasevs@saspw account, select **+ arrow button**,
- select **OK**

Appendix C: Methods of collecting logs in a multiple machine environment

SAS Environment Manager can be configured to collect and report on audit and performance data extracted from the log files generated by the various SAS and other servers in your SAS deployment. These logs are collected (i.e. copied) to a central “Landing Zone” under the **emi-framework** directory.

If your SAS environment is deployed across multiple hosts, the SAS Environment Manager must be able to collect log data from those remote hosts.

By default, logs from these remote hosts are copied to the landing zone on the emi-framework host through the use of SAS Environment Manager Agent and SAS Deployment Agent technologies. Using an automated process, logs are collected from the SAS servers into a local landing zone on each host. The SAS Deployment Agent then copies these logs from each remote landing zone into the centralized landing zone on the SAS Environment Manager Enablement Kit Server host. These logs are then parsed to extract information which is loaded into the SAS Environment Manager Data Mart.

At some sites, there may be site restrictions that prevent or discourage this mechanism. For those sites, there are the following alternative techniques:

Sharing logs between machines using NFS mounts

An alternative method of collecting logs from remote machines is to set up NFS mount points for shared access to the host landing zones. This technique allows the SAS Environment Manager Log Collector, running on the remote machines, to write logs to an NFS mount point that refers to the physical landing zone on the SAS Environment Manager Enablement Kit Server host’s file system. To accomplish this, follow these steps:

NOTE: In the following examples, the emi-framework is installed on *node1.domain.com* and the emi-framework directory is **/disk1/Config/Lev1/Web/SASEnvironmentManager/emi-framework**. The remote hosts include *node2.domain.com* and *node3.domain.com*.

1. Log in as the original installer account. Some actions below must be performed as the superuser; all other actions should be performed as the original installer.
2. The daemon processes ‘nfsd’ and ‘mountd’ should be running on the SAS Environment Manager Enablement Kit Server host and on any other host that you would like to collect data from, such as the other hosts in the multi-machine SAS environment. On Linux, you can confirm that these daemon processes are running by issuing this command:

```
/sbin/service nfs status
```

3. On the SAS Environment Manager Enablement Kit Server host, change directory to the **emi-framework** directory. For example,

```
cd /disk1/Config/Lev1/Web/SASEnvironmentManager/emi-framework
```

4. Export the **LandingZone** directory so that it can be NFS-mounted on other machines, and processes running on those other hosts can write to this directory. NOTE: You must be superuser to accomplish this. For example,

```
exportfs -o "rw,sync"  
node2.domain.com:/disk1/Config/Lev1/Web/SASEnvironmentManager/emi-  
framework/LandingZone
```

NOTE: You can repeat the exportfs command for each remote machine, or you may specify an asterisk (*) as the remote machine, to export to all hosts.

5. On the remote hosts from which log files will be collected, NFS-mount this location in place of the original **LandingZone** directory (located under the **emi-client** directory). First, rename the existing **LandingZone** so that any current files are kept intact. For example,

```
cd /disk1/Config/Lev1/Web/SASEnvironmentManager/emi-client  
mv LandingZone LandingZone.original
```

Then create a new **LandingZone** directory and mount the remote file location here. This location must be set to be writeable.

For example,

```
mount -t nfs node1.domain.com:/disk1/Config/Lev1/Web/SASEnvironmentManager/emi-  
framework/LandingZone ./LandingZone
```

NOTE: You must be superuser to issue the mount command.

6. Logs on the SAS Environment Manager Enablement Kit Server host need to mimic the mounted-directory configuration as above by linking the **LandingZone** directory in the local **emi-client** directory to the central **LandingZone** in the **emi-framework** directory. First, rename the existing **LandingZone** so that any current files are kept intact, then link the local **LandingZone** directory to the central **LandingZone** directory. For example,

For UNIX/Linux:

```
cd /disk1/Config/Lev1/Web/SASEnvironmentManager/emi-client  
mv LandingZone LandingZone.original  
ln -s /disk1/Config/Lev1/Web/SASEnvironmentManager/emi-framework/LandingZone  
./LandingZone
```

For Windows, the equivalent commands would be similar to:

```
cd C:\SAS\Config\Lev1\Web\SASEnvironmentManager\emi-client  
rename LandingZone LandingZone_Original  
mklink/d LandingZone C:\SAS\Config\Lev1\Web\SASEnvironmentManager\emi-  
framework\LandingZone
```

(note: the arguments are reversed compared to the UNIX/Linux "ln" command)

7. On the SAS Environment Manager Enablement Kit Server host, edit the **emi-framework/sasev.properties** file and add this line:

```
logCentralize.enabled=false
```

That tells the log processing that remote logs are delivered via shared directories, so it does not need to explicitly collect log bundles from the remote hosts.

8. The SAS Environment Manager Log Collector, running on the remote hosts, will now collect log files and write them into the landing zone, and the SAS Environment Manager Server will be able to read those logs on the SAS Environment Manager Enablement Kit Server host. Note that as the Service Architecture Framework processes these logs, successfully-processed log files will be deleted from the shared mount point location, effectively deleting them from the remote host also.

NOTE: Although the example syntax above is given for Linux systems, NFS is available on Windows as well, as the “Server for NFS” feature. See the Microsoft documentation for your specific Windows operating system for details.

Other methods of copying logs to the server

If neither the SAS Deployment Agent method nor the NFS mount method are suitable at your site, there are a number of other methods that you can utilize to get logs from the remote machines to the appropriate server location.

- You can have an automated FTP/SFTP/SCP script to transfer the log files from each remote LandingZone into the EMI Framework Landing Zone.
- You can use an alternative form of network-mounted disk access (SMB, etc.). Any method that results in the log files being transferred or written to the expected location on the SAS Environment Manager EMI Framework host will work, as long as the log file arrives intact and in readable form.
- You can even “sneaker net” the files from the remote hosts to the EMI Framework Landing Zone (i.e. - copy the files onto removable storage such as a thumb-drive, tape, or other media, and physically transporting the data to the EMI Framework host). The parsing logic in the EMI Framework can process a series of multiple logs; so you could allow logs to accumulate on the remote hosts for several days and physically transfer them to the EMI Framework host on a periodic basis, such as weekly.

Appendix D: Upgrading the Service Architecture Framework

If the Service Architecture framework has never been initialized or enabled in your deployment, there is no need to perform an upgrade of the framework. In this situation, you can safely ignore any WARNING messages about the SAS Environment Manager Service Architecture during an upgrade of your SAS deployment.

The upgrade-in-place process automatically upgrades a deployment of the SAS Environment Manager Service Architecture. This involves updating source code, delivering new capabilities and changes to the data mart schema. For the majority of customers, those who have not customized or altered, their deployment, this automatic upgrade process should be able to take care of everything. However, for sites that have extended or modified the Service Architecture, there may be some additional steps needed. This appendix identifies modifications or extensions sites may have made and provides guidance on how to ensure the updated deployment works as desired. It is expected that very few customers will need to perform any of the steps discussed below and it is unlikely that any customer would need to perform all of the steps below.

User-Created Content

If your site has created any additional stored processes, reports or other content that depends on the SAS Environment Manager Data Mart, you may need to perform additional steps to continue to use that content. In most cases, the existing Data Mart data will be available after an upgrade-in-place. While there have been some changes to the Data Mart tables, the changes have been minor (e.g. changes in variable lengths, assignment of formats, etc.) and should not require significant changes to your user-created content. The library and tables names remain the same.

The metadata folders containing the Report Center reports and stored processes were re-created during the upgrade-in-place. Virtually all of the existing reports which were stored as “stored process report” objects were replaced with new reports in the form of “stored process” objects. In most cases, the new reports have the same names as the previous reports and convey the same information. However, the report output is now generated using ODS graphics technology and the appearance of the reports may have changed.

It was not expected that customers would store their own content in the metadata folders used by the Report Center. However, if you did store content in those folders, that content was deleted when the folders were re-created. To provide some measure of protection in such cases, the contents of the metadata folders were exported, as a SAS package file, before the folders were deleted. This SAS package file will have a name of the form **productsFolderBackup. {TIMESTAMP}. spk**, and should be in the **SASPackages** sub-directory of the upgraded Service Architecture root directory, for example, **<levelroot>/Web/SASEnvironmentManager/emi-framework/SASPackages/productsFolderBackup.20161026_105926 .spk**. The SAS Management Console includes a metadata import wizard that can be used to recover content from this file.

If your content included any stored processes and you chose to store the source code for these stored processes in the Service Architecture directory structure, that source code should be available in the upgrade-in-place back-up directory. In fact, this back-up directory should contain any files that were present in the Service Architecture directory tree when the upgrade-in-place was started, except for the

Data Mart data which was moved to the new Service Architecture root directory. The name of this backup directory incorporates the timestamp of when the upgrade-in-place process was started, for example, `<levelroot>/Web/SASEnvironmentManager/emi-framework_2016-10-26_09.49.30/SASPACKAGES/emiContent.spk`.

User Modified Configuration Properties

In some cases, changes made to the Service Architecture configuration properties via the Configuration Manager plug-in to the SAS Management Console may be over-written during an upgrade-in-place. If your site has made such modifications, you should use the plug-in to review the current settings to confirm you are comfortable with the current values or update the values to reflect your preferences.

User Modified Data Mart Location

Most sites have used the default location within the Datamart directory under the Service Architecture root directory for the Data Mart data libraries, i.e. `<levelroot>/Web/SASEnvironmentManager/emi-framework/Datamart`. The upgrade-in-place process should handle that situation without any issues. The process uses a series of host system move or copy commands to make the existing data available after an upgrade-in-place. However, if your site has used soft (or symbolic) links, NFS mounts, or some other mechanism to redirect the Data Mart data libraries at the operating system level, you will need to review the state of things after the upgrade-in-place to determine if the data was moved successfully. If the data redirect is no longer working, you will need to fix it and then re-run some of the upgrade-in-place processing to ensure the Data Mart tables are structured appropriately for this version of the Service Architecture. Information on re-running the appropriate portion of the upgrade-in-place process are included in the section titled “**Running the ugiplnit script**”

If your site modified the LIBNAME statements in any of the various Service Architecture autoexec files, those changes were not carried forward. These files have names that start with EVDMAUTOEXEC and are located in the Datamart directory. As described above, the previous versions of these files can be found in the Service Architecture upgrade-in-place backup directory. However, do NOT simply replace the new autoexec files with the previous ones. Doing so will render the Service Architecture unusable. Instead, you should review the previous versions to identify what you changed and make corresponding changes to the new versions. Note that two of the previous autoexec files, `EVDMAUTOEXEC_KITS.sas` and `EVDMAUTOEXEC_APM.sas`, are no longer needed. In general, modifications to these files are strongly discouraged. Although there is a `sasev.datamartDir` property defined in the `sasev.properties` file located in the Service Architecture root directory, modifying this property is *not sufficient* to specify an alternate location for the SASEV Datamart.

Running the ugiplnit script

After you have completed any of the changes described above that apply to your situation, you will need to run a script to finish the upgrade-in-place process. The script can be run without specifying any arguments or parameters. Messages will be generated as the script executes updating you on what it is doing. At the final message should indicate if the script completed successfully or not. If not, the messages will identify a log file for you to review for more information on what caused the script to fail. Correct the issues identified and re-run the script.

1. Change directory (cd) to the bin subdirectory of the Service Architecture root directory. This will be `<levelroot>/Web/SASEnvironmentManager/emi-framework/bin` (UNIX), `C:<levelroot>\Web\SASEnvironmentManager\emi-framework\bin` (Windows). Where `<levelroot>` represents the appropriate level directory under your SAS Configuration directory (e.g. `/install/SAS/Config/Lev1`).

2. Run the script:

On UNIX:

```
$ ./ugip_init.sh
```

On Windows:

```
ugip_init.bat
```

A successful execution of the upgrade script indicates that the Service Architecture framework has been updated to the current version.

Upgrading the Service Architecture Framework in a multi-machine deployment

If your SAS deployment includes multiple machines, for example, one serving as the server tier and another for the middle-tier, the automated upgrade-in-place of the Service Architecture is likely to fail. An error or warning messages similar to the following will appear in the `UpdateInstructions.html` file generated during the upgrade:

SAS Environment Manager Service Architecture

A problem was encountered while updating the SAS Environment Manager Service Architecture. Review the log `/config/Lev1/Logs/Configure/evmkitsvr_updateContent_2015-04-01-08.03.47.log` for more information on the problem. After addressing any issues, re-run the upgrade script from the command line. Refer to the appendix entitled `Upgrading the Service Architecture` in the `Service Architecture Quickstart Guide` for instructions on re-running this script. Please note that the appendix also outlines additional manual steps that may be needed to ensure your deployment of the Service Architecture remains fully functional.

These issues are caused by portions of the upgrade process needing to communicate across the various tiers and encountering problems since some of the tiers may not be up or may not have been upgraded yet. In these cases, the solution is to ignore the error or warning messages and proceed with the rest of the upgrade process on all of the machines in the deployment. Once the rest of your SAS deployment has been upgraded and brought back on-line, you can upgrade the Service Architecture. To do so, run the `ugipInit` script as described above. Note that the script only needs to be run once and only on the machine on which the Service Architecture root directory.

Log Collection for Multi-machine Configurations

As part of the upgrades to the Service Architecture Framework beginning with SAS 9.4M3, a number of configuration steps were simplified, and log collection was automated using the SAS Deployment Agent and log discovery, collection, and centralization tasks.

In the SAS 9.4M2 release of the Service Architecture Framework, to allow the APM ETL to collect logs from remote machines, the administrator had to configure those machines to share their SAS <levroot> directories to the SAS Environment Manager Enablement Kit Server host, typically through the use of NFS mount points or SMB shares. The administrator then updated a **log_definitions.json** file to define where to find those mount points in the local host's directory tree (which could be at any arbitrary points in the tree).

The SAS 9.4M3 release of the Service Architecture Framework still supports using network-shared directories, but requires that those exported paths be mounted to a specific place in the SAS Environment Manager Enablement Kit Server host directory tree under the **emi-framework** directory. Rather than requiring access to the complete <levroot> subtree on the remote host, the administrator only needs to export a landing zone for the SAS Environment Manager Enablement Kit Server host and mount that path into the local landing zone on each remote host. See **Sharing logs between machines using NFS mounts** in **Appendix C: Methods of collecting logs in a multiple machine environment**.

The default log collection method avoids NFS mounts/SMB shares altogether, using the SAS Deployment Agent to collect compressed archives of logs from the remote machines to the SAS Environment Manager Enablement Kit Server host where they are expanded, parsed, and archived. This process happens automatically at scheduled intervals. SAS recommends that the administrator use this method for centralizing logs going forward.

Appendix E: Modifying APM to work with Common Directories for Logging

The APM log discovery, collection, and ETL processing is designed to work with the standard log directory hierarchy located under the LevN/... configuration directory tree. This automated processing will not work if you selected the custom deployment option *Create Common Directories* for a single common logging directory at deployment time.

If you have enabled Common Directories, you can manually add symbolic links in the deployed LevN directory hierarchy for the traditional log directories to allow APM to collect and process logs from the common log directory (typically located at LevN/Logs).

On each SAS host that has Common Directories enabled for logging, create symbolic links for each supported server's log directory(ies) in each application context in the LevN hierarchy. For example, for the SAS Metadata Server in a standard deployment there are two log directories: Logs and AuditLogs, but only AuditLogs is used by APM.

Create symbolic links from those directory names to the common logging directory. Assuming SAS was deployed into the normal default <levelroot> directory:

Linux/Unix, deployed to:

```
/usr/local/sas/Config/Lev1
```

with a common logging directory located at

```
/usr/local/sas/Config/Lev1/Logs
```

As the SAS install user, create a symbolic link as follows, this example is for MetadataServer/AuditLogs:

```
$ ln -s /opt/sas/config/Lev1/Logs /opt/sas/config/Lev1/SASMeta/MetadataServer/AuditLogs
```

Windows, deployed to:

```
C:\SAS\Config\Lev1
```

with a common logging directory located at

```
c:\SAS\Config\Lev1\Logs
```

As the SAS install user, create a symbolic link as follows, this example is for MetadataServer/AuditLogs:

```
sas> mklink/d c:\SAS\Config\Lev1\SASMeta\MetadataServer\AuditLogs c:\SAS\Config\Lev1\Logs
```

Note: the Unix `ln -s` command and the Windows `mklink/d` command have different target argument ordering.

The APM log discovery and collection process looks for the following log directories. Some of these directories may not exist in a given deployment's configuration, depending on options chosen for that deployment and/or host:

At the main <levelroot> directory:

```
<levelroot>/ConnectSpawner/Logs  
<levelroot>/ObjectSpawner/AuditLogs  
<levelroot>/Web/Logs
```

For the MetadataServer context (default is SASMeta):

```
<levelroot>/SASMeta/MetadataServer/AuditLogs
```

For *each* application context (default is SASApp):

```
<levelroot>/SASApp/BatchServer/PerfLogs  
<levelroot>/SASApp/ConnectServer/Logs  
<levelroot>/SASApp/ConnectServer/PerfLogs  
<levelroot>/SASApp/GridServer/Logs  
<levelroot>/SASApp/GridServer/PerfLogs  
<levelroot>/SASApp/OLAPServer/Logs/AuditLogs  
<levelroot>/SASApp/OLAPServer/Logs/Logs  
<levelroot>/SASApp/OLAPServer/Logs/PerfLogs  
<levelroot>/SASApp/PooledWorkspaceServer/Logs  
<levelroot>/SASApp/PooledWorkspaceServer/PerfLogs  
<levelroot>/SASApp/StoredProcessServer/Logs  
<levelroot>/SASApp/StoredProcessServer/PerfLogs  
<levelroot>/SASApp/WorkspaceServer/Logs  
<levelroot>/SASApp/WorkspaceServer/PerfLogs
```

Make links from each of those directories to your common logging directory:

Unix/Linux (typical): /usr/local/sas/Config/Lev1/Logs

Windows (typical): c:\SAS\Config\Lev1\Logs

After creating the required log directory links for each host in the SAS deployment, and after enabling the APM log collection (see section *Initializing and Enabling APM ETL*), check the APM logs in

```
Lev1/Web/SASEnvironmentManager/emi-client/Logs/logCollect_<datetime>.log
```

to verify that logs are being collected through the symbolic links. Collection processing may be delayed up to four hours before the next automatic discovery/collection cycle is started.

Appendix F – Configuration modifications made when APM is initialized

When the APM subsystem is initialized, it makes modifications to logging definitions, stored process invocation, and related SAS Metadata Server metadata properties. Each of these changes is logged in

Lev1/Web/SASEnvironmentManager/emi-framework/Logs/logDiscovery_<datetime>.log

The APM configuration step is automatically run as part of each log discovery cycle to make required changes to any application contexts added after the initial deployment.

Metadata changes

APM enables Application Response Measurement (ARM) metrics logging for jobs run on certain SAS servers such as the Logical Pooled Workspace Servers, the Logical Stored Process Servers, and the Batch Servers. As part of this, it makes modifications to metadata property definitions associated with those server types. Note: these modifications are *required* for APM to function correctly.

For each Logical Pooled Workspace Server and each Logical Stored Process Server found in metadata, the associated properties are modified as follows. In default SAS installations, these properties are not defined. Both the old and new values are logged:

Initialization program = SASEnvironment/SASMacro/apm_request_init.sas

Termination program = SASEnvironment/SASMacro/apm_request_term.sas

For each Batch Server found in metadata, the associated property is modified as follows:

BatchServer.DataStep.Property.LogExt.xmlKey.txt = _ (note: underscore)

Logging configuration changes

As part of enabling ARM metrics logging and audit logging, APM changes the initialization for the related servers to use the server's supplied *logconfig.apm.xml* logging definition file instead of the default *logconfig.xml* definition file. This change is implemented by appending a new

-logconfigloc <path>

entry to each server's *sasv9_usermods.cfg* file which will override the default log definition. A copy of the original *sasv9_usermods.cfg* file is saved as a backup file with an appended date-time, as follows:

sasv9_usermods.cfg_<yyyymmdd>_<hhmmss>

The new *sasv9_usermods.cfg* file has a comment wrapper that serves as a flag to the APM configuration process to indicate that this file has already been updated and does not need to be changed again.

The following server types receive modified logging definitions:

- MetadataServer
- BatchServer
- ConnectServer
- GridServer
- OLAPServer
- PooledWorkspaceServer
- StoredProcessServer
- WorkspaceServer

Application Context directory changes

For each application context (SASApp, etc.), the APM configuration step copies the required ARM macros into that context's *SASEnvironment/SASMacro* subdirectory.

For each application context, an invocation of the `%apm_armsetup` macro is appended to the context's *appserver_autoexec_usermods.sas* file to initialize the ARM environment. A copy of the original *appserver_autoexec_usermods.sas* file is saved as a backup file with an appended date-time, as follows:

```
appserver_autoexec_usermods.sas _<yyyymmdd>_<hhmmss>
```

The new *appserver_autoexec_usermods.sas* file has a comment wrapper that serves as a flag to the APM configuration process to indicate that this file has already been updated and does not need to be changed again.

Appendix G – Multiple <levelroot> directories on the same host

Overview

In some customized configurations, SAS may be deployed with multiple distinct server tiers on the *same* host machine, resulting in multiple LevN directories for that SAS deployment on a given host. Typically, this is done to allow an initial single host or smaller set of host machines to function as if it were a larger multi-machine deployment, with the goal of eventually moving the additional server tiers to their own host machines.

To support this, the 9.4M7 release of the SAS Environment Manager Service Architecture framework has added a feature which allows an administrator to direct the APM log discovery and collection tools to scan multiple LevN/ directory hierarchies on the local host.

Limitations

A SAS deployment may be configured on a single host or across multiple hosts. In either case, the Service Architecture framework only supports collecting APM log and metric data from a single SAS deployment – it does not support collecting APM data from multiple distinct SAS deployments, for example Dev/Test/Prod or Lev1/Lev2/Lev3/etc., into a single data mart.

Implementation

For a typical SAS deployment, one <levelroot> is configured on each host of the deployment. That <levelroot> may contain multiple combined server tiers, for example, a single `/opt/sas/Config/Lev1` could contain each of SAS Metadata Server, SAS Compute, SAS WebAppServer, SAS OLAP Server, etc.

At regular intervals, a SAS Environment Manager Agent in the <levelroot> directory hierarchy on each host of the deployment triggers APM tools to scan that local directory hierarchy, “discovering” the SAS server logs supported by the APM ETL and collecting them into a local landing zone directory. Those logs are eventually copied from the landing zone to the SAS Environment Manager Enablement Kit Server for processing. See: *SAS Environment Manager 2.5: User’s Guide – Chapter 8, Understanding SAS Environment Service Architecture*.

In some cases, SAS may be deployed with multiple distinct server tiers in different LevN directories on the *same* host machine. For example, a site might configure the SAS Metadata Server tier and the related SAS Compute tier on the same host, but would configure them under different base <levelroot> directories to simplify splitting them out to independent hosts at a later date:

```
/opt/sas/MetaConfig/Lev1/ ...
/opt/sas/ComputeConfig/Lev1/ ...
```

Only one running SAS Environment Manager Agent is allowed *per deployment on a given host machine*, so the Agents are either not configured or are disabled in all but one of those <levelroot> hierarchies.

Note: the above is a greatly simplified description of the actual configuration definition and deployment process for multiple independent server tier levelroots on the same host.

Each running Agent will trigger log discovery and collection only in its own *<levelroot>* directory hierarchy by default. Any additional local *<levelroot>* directory hierarchies will be ignored.

To support discovering and collecting logs from those additional *<levelroot>* directory hierarchies, APM now supports adding a *levroots_override.txt* file to the APM configuration on each host to specify an explicit list of *<levelroot>* directories to scan on the given host.

Specifying an override file

In the *<levelroot>* directory hierarchy that implements the *active running SAS Environment Agent* for that host, in the

<levelroot>/Web/SASEnvironmentManager/emi-client/Conf/

sub-directory, create a file called *levroots_override.txt* which contains fully-qualified paths to each *<levelroot>* that should be scanned on this host, one per line, including an entry for the current *<levelroot>*. For example:

/opt/sas/MetaConfig/Lev1

/opt/sas/ComputeConfig/Lev1

(or Windows)

C:\SAS\ConfigMeta\Lev1

C:\SAS\ConfigMidtier\Lev1

No blank lines, leading or trailing spaces, or other items (comments) are allowed. The file is assumed to be plain text. Be careful in that some utilities such as Windows Notepad may put extraneous characters or character-encodings in the file. On Windows platforms, the backslashes in paths should be escaped with additional backslashes. Spaces in path names should also be escaped.

The discovery scripts do a simple sanity-check to verify that the specified path(s) exist on the host and complain if they are not found. The discovery process will log INFO messages in the

.../emi-client/Logs/logDiscovery_<date>_<time>.log

for each override directory that it successfully adds from the *levroots_override.txt* file

<date> <time> INFO [main] [logDiscovery] - Adding override levroot <path>

The discovery tool will then scan through each of those hierarchies, adding relevant APM log location entries to a list used to drive the log collection process.

If one or more of the additional server tiers are migrated to other hosts, remove those *<levelroot>* entries from the *levroots_override.txt* file.