

Installation Instructions for Hot Fix A8R001

The hot fix, A8R001x6.zip, contains the updated files that are required to address the documented issues.

Note: Do NOT extract the contents of the A8R001 ZIP file. The hot fix installation process extracts the contents as needed.

Important Notes

- Before you apply this hot fix, you must have SAS Firmwide Risk for Solvency II 3.1 installed on your system. Refer to [SN-35968](#) for instructions on how to determine which product releases you have installed.
- When the hot fix installs, the files delivered in this hot fix are automatically backed up. Regardless, it is recommended that you back up your system before you apply updates to software.
- You must have administrator privileges on your client and server machine.
- Before you apply this hot fix, you must terminate all active SAS sessions, daemons, spawners, and servers.
- When you install this hot fix, you must use the same user ID that was used to perform the initial software installation.
- **CONFIGURATION NOTE:** No automatic configuration is included in this hot fix.

If you have previously configured software installed, the SAS Deployment Manager might present a screen on which you see “Apply SAS Hot Fixes” and “Configure SAS Hot Fix” options. On this screen, ensure that the “Configure SAS Hot Fix” option is **not** selected. If this option is automatically selected, deselect it before you proceed to the next SAS Deployment Manager screen.

Caution: If you do not deselect the “Configure SAS Hot Fix” option, unintended consequences can occur when you apply this hot fix.

How to Install the Hot Fix

You use SAS Deployment Manager to install the A8R001 hot fix for SAS Firmwide Risk for Solvency II Server 3.1. By default, the SAS Deployment Manager searches in the `SASHome/InstallMisc/Fixes/New` directory for hot fixes to be applied. However, you will also be prompted for a location if you have downloaded hot fixes to a different directory.

After you download the A8R001a.zip file, follow the instructions for how to apply hot fixes in the [SAS Deployment Wizard and SAS Deployment Manager 9.4: User's Guide](#).

Caution: Make certain to review the CONFIGURATION NOTE above. You must ensure that the “Configure SAS Hot Fix” option in the SAS Deployment Wizard is deselected.

The content of this hot fix is listed in the [hot fix manifest](#).

Post-installation Instructions

This document provides instructions on how to run the `recreate_flow_instances.sas` script to restore an unsupported SAS Firmwide Risk for Solvency II 3.1 installation to the supported SAS Firmwide Risk for Solvency II 3.1 installation.

The target audience of this document is a system administrator who has access to all SAS Firmwide Risk for Solvency II user profiles.

Introduction

SAS Firmwide Risk for Solvency II 3.1 supports one federated area.

Only one federated area in a SAS Firmwide Risk for Solvency II 3.1 installation is supported because of how additional federated areas are handled. When more than one federated area is installed, existing job flow instances pick up input code, data, and flows, from the highest priority federated area. However, the internal data structure of SAS Infrastructure for Insurance Management continues to reflect the state of the system when the instances were created.

As a result, customers who have implemented SAS Infrastructure for Insurance Management 3.1 with multiple federated areas installed are not able to access instances that were created in different federated areas.

To correct this condition, the unsupported SAS Infrastructure for Insurance Management installation (with multiple federated areas) must be restored to the supported installation (with one federated area).

Complete the Prerequisite Setup

Before you begin the process to restore your SAS Infrastructure for Insurance Management installation to the supported installation, you must complete the following tasks:

1. Backup your system.

Ensure that you make a backup of the system (federated areas, persistent area, and PostgreSQL database) before you make any changes.

2. Manually consolidate all inputs from the multiple federated areas into the highest priority federated area.

All job flow instances that you want to re-create must be creatable from the highest priority federated area. This federated area can be a federated area that already exists or you can create a new one.

The inputs that must be consolidated from the multiple federated areas into the highest priority federated area include the following:

- Landing area data sets, templates, and so on, for all base dates and configurations
- Files located in the `config` folder (for example, `job_flow_definitions.csv`)
- Job flow definitions (`.bpmn` files) for all categories
- All configuration sets must be defined in the platform federated area, which is by default `fa.0.3.1`
- Node files, which are located in the source folder of a federated area

3. Ensure that the federated area or areas, including any new federated areas, are correctly defined in SAS Management Console.

If your deployment does not have any published instances, use SAS Management Console to delete every federated area except for the highest priority federated area from metadata.

If your deployment has instances that have been published, all federated areas from which published instances were created must be retained in SAS Management Console in order to enable access to these instances.

Note: If you created a new federated area in which to consolidate the inputs from the multiple federated areas that exist, you must restart the SAS Firmwide Risk for Solvency II 3.1 server.

4. Create a `lib_mapping` data set. Store it in the persistent area.

For the `recreate_flow_instances.sas` script to re-create instances with modified input data sets, you must create a library mapping table named `lib_map.LIB_MAPPING` and store this table in the SAS Firmwide Risk for Solvency II persistent area.

Create the Library Mapping Data Set (LIB_MAPPING)

The LIB_MAPPING table contains mappings from the relative landing area folder path to its LIBRARY, which is used in SAS Firmwide Risk for Solvency II tasks.

The LIB_MAPPING table consists of the following two columns:

Column	Description
LIBRARY	Contains the library name as specified in the libnames.txt file.
LA_PATH	Contains the relative landing area path in the libname.txt file, but with all variables (such as %cs and %bd) resolved. Note: On Windows, backslashes (\) are the expected separators. On UNIX, forward slashes (/) are the expected separators.

For example, for GLOBAL library, the LA_PATH is /base/global.

When you create the LIB_MAPPING table, note the following:

- Create this mapping table based on the landing area structures for the highest priority federated area.
- Create the mapping table entries based on the libnames.txt file in the consolidated federated area.

You can use the following example code to create the LIB_MAPPING data set and the required data set structure:

```
libname lib_map "/local/install/cfgsas1/SAS94/PersistentArea";
data lib_map.LIB_MAPPING ;
attrib
    LIBRARY    length= $32
    LA_PATH    length= $128
;

infile cards dsd delimiter=',';
input
    LIBRARY    :$char.
    LA_PATH    :$char.
;
datalines4;
"GLOBAL", "/base/global"
"LA_STAGE", "/03312015"
"LA_STAGE", "/03312014"
"LA_STGM1", "/03312013"
"RD_ENV", "/rd_env"
"LA_STAT", "/base/eiopa_full_sii_2.0.1/static"
```

```
"LA_STAT", "/base/eiopa_full_sii_2.0.0/static"  
"LA_MAP", "/base/eiopa_full_sii_2.0.1/mapping"  
"LA_MAP", "/base/eiopa_full_sii_2.0.0/mapping"  
"ENV_DEF", "/base/env_def"  
;;;;  
run;
```

Run the Re-create Job Flow Instances Script

The following sections provide process information and instructions on how to run the `recreate_flow_instances.sas` script.

Note: Check the status of each step in the process before you proceed to the next step.

Before you proceed with the steps in the section, you must have completed the prerequisite setup documented in the [Complete the Prerequisite Setup](#) section.

What the Script Does

The algorithm of the `recreate_flow_instances.sas` script is as follows:

1. For each existing non-published job flow instance that created successfully (with no errors) for a specified user, the script gathers the following information, which is needed to re-create the instance:
 - Name
 - Job flow definition category and name
 - Job flow definition file
 - Entity and entity role
 - Configuration set
 - Base date
 - Description
 - Uploaded files, if any
2. Produces an HTML file (named `existing_instances_username.html`) that contains a summary of all existing instances for the specified user.
3. Copies uploaded files to a safe location.
4. Deletes the existing job flow instance.
5. Re-creates a new job flow instance using the information collected in Step 1.
6. Uploads any modified inputs and triggers the instance to execute. This step is repeated for all modified inputs for the instance.
7. Shares the instance if the original instance was shared.

8. Repeats steps until there are no remaining existing non-published and completed successfully job flow instances.
9. After the last existing instance is run, produces an HTML file (named `latest_instances_username.html`). The file contains a summary of all of the instances that were re-created and retained for the specified user.

How to Run the Script

After you have completed the prerequisite setup, complete the following steps to re-create error-free and non-published job flow instances:

1. Ensure that all SAS servers have been started and that the SAS Infrastructure for Insurance Management server is running.
2. Ensure that there are no users logged in.
3. Edit and save a copy of the `recreate_job_flow_instances.sas` script for each user whose instances you plan to re-create. (See [Edit the Script.](#))
4. To re-create instances for each user, run the user-modified version of the `recreate_job_flow_instances.sas` script. (See [Run the Script.](#))

Note: You must repeat this process every time a new federated area is installed.

Edit the Script

Before you run the `recreate_flow_instances.sas` script to re-create job flow instances, you must edit the script to specify the user, site, and job-specific parameters. After you edit the script, save the script with the associated user name appended to the end (for example, `recreate_flow_instances_username.sas`).

The section of the file that contains the parameters that you must edit is preceded by the following comment:

```
REQUIRED PARAMETERS WHOSE VALUES MUST BE SET ON A PER SITE BASIS
```

To edit the `recreate_flow_instances.sas` script, complete the following steps:

1. Navigate to the script, which is located in the following directories:

UNIX:

```
/SASHome/SASFoundation/9.4/rmifirmmva/sasmisc/recreateInstanceJob
```

Windows:

```
\SASHome\SASFoundation\9.4\rmifirmmva\sasmisc\recreateInstanceJob
```

2. (Required) Open the file and locate the section that contains the site-specific parameters and enter values for the following:

- **persistent_area_loc** – The complete path to the SAS Firmwide Risk for Solvency II persistent area.

For example:

```
/local/install/IRM/SAS94/PersistentArea
```

- **highest_fa_id** – The ID of the highest priority federated area in the system. This federated area contains all of the information needed to re-create the job flow instances.
- **srchostname01** – The host name.

For example:

```
irmserver.abc.xyz.com
```

- **srcimport01** – The server port.

For example:

```
irmserver.abc.xyz.com:7980
```

- **srcusername01** – The user name of the user that created the job flow instances. The job flow instances will be deleted and then re-created for this specific user.

For example:

```
irmuser1
```

Note: If the user name specified for the **srcusername01** parameter is a qualified user ID, use the @ domain delimiter when you specify the user name.

For example:

```
user@domain
```

- **srcpassword01** – The password for the user specified for the **srcusername01** parameter. You can enter this value as plain text or SAS-encoded. The password is masked when you execute the script.

3. (Optional) The following parameters have default values, but you can change the values as desired:
 - **debug** – Sets the amount of information logged when the script is executed. Possible values are **TRUE** (log debugging information) and **FALSE** (do not log debugging information). The default is **TRUE**.
 - **logOptions** – Specifies the amount of macro execution information in the log. Valid values are **mprint**, **symbolgen**, **mlogic**, and other SAS supported options. The default is blank.
 - **pollInt** – The interval, in seconds, at which the status of job flow instance creation is checked. The default is **10**.
 - **maxWait** – The maximum amount of time, in seconds, the script waits for the creation of a job flow instance to complete. The default is **3600**.
4. Save the file with the name of the user added to the filename.

For example:

```
recreate_flow_instances_username.sas
```

5. For each user whose job flow instances you want to re-create, repeat steps 1 through 4.

After you edit and save the `recreate_flow_instances.sas` script for each user, you can then run the script to re-create the instances in the designated federated area.

Run the Script

To run the script, complete the following steps:

1. View the `recreate_flow_instances_username.sas` file and verify that the parameters are correct and that the correct user is specified.
2. From a command line, enter the following commands to run the script.

UNIX:

```
cd /SASHome/SASFoundation/9.4/rmifirmmva/sasmisc/  
recreateInstanceJob
```

```
/SASHome/SASFoundation/9.4/sas.exe -sysin recreate_flow_  
_instances_username.sas -log recreate_flow_instances_username.log
```

Windows:

```
cd \SASHome\SASFoundation\9.4\rmifirmmva\sasmisc\  
recreateInstanceJob
```



```
\SASHome\SASFoundation\9.4\sas.exe -sysin recreate_flow_instances_username.sas -log recreate_flow_instances_username.log
```

where *username* is the name of the user for which you edited the file.

3. When the script has completed, scan the log for any errors that might have occurred during the re-creation process. If there are errors, resolve the issues, and repeat the process to run the script. For information about how to troubleshoot error messages, see the [Troubleshooting Error Conditions](#) section.
4. When the script has completed, view the HTML output files that are generated by the script. The files are located in the same directory as the script and are named `existing_instances_username.html` and `latest_instances_username.html`.

These output files contain a list of all of the job flow instances that were re-created.

5. Repeat steps 1 through 4 with each script that you created for a user.
6. From the SAS Firmwide Risk for Solvency II 3.1 web application, verify that all of the job flow instances that you re-created are accessible and defined correctly.

Troubleshoot Error Conditions

After you re-create job flow instances in the highest federated area, the re-created instances should be available in the SAS Firmwide Risk for Solvency II web application. If any of the re-created instances are not accessible from the web interface, an error occurred during the re-creation process.

The following table lists error conditions that might occur during the re-creation process and the possible causes for those error conditions:

Condition	Possible Causes
Errors in the <code>recreate_flow_instances</code> log	<p>If there are errors written to the log, examine the log on a case-by-case basis.</p> <p>Verify that the user-supplied parameters are set correctly.</p> <p>There are several reasons that re-creating instances with modified input might fail, including the following:</p> <ul style="list-style-type: none"> • Did the user provide the correct <code>LIB_MAPPING</code> table? <p>The script checks the delimiter for <code>LA_PATH</code> column inside the <code>LIB_MAPPING</code> table before it re-creates any instances. If the script finds any mismatched delimiters between systems, it generates a message similar to the following:</p>

ERROR: Library mapping table does not have the correct delimiter. Windows \ and Linux /.

The script also checks if a data file being uploaded can find the corresponding library in the associated instance. If the library cannot be found via the library mapping table for a data file, the script generates the following error:

```
ERROR: Some modified data from the instance
and instanceKey could not find the
corresponding library from the lib_mapping
table.
```

Note: If the script encounters the two conditions above, it terminates the re-create before any instances are re-created, regardless of whether it is a modified instance or not.

Recommended Action: Fix the LIB_MAPPING table and rerun the script.

- Is the data key found?

The script checks if the data file to be uploaded can find its data key in the re-created instance. If the data key is not found in the re-created instance, the script skips the upload and generates the following error message:


```
ERROR: data set "&Lib-nm" "&up_root_nm." Could
not find the datakey for the instance.
```


Recommended Actions:

1. Verify that the LIB_MAPPING table provides the correct file path to library mapping. If the libname.filename cannot be found in the instance, the data key cannot be resolved.
2. Ensure that the consolidated federated area is correctly merged. For example, ensure that the job flow definition in the consolidated FA has the libname.filename in it.
3. If necessary, fix existing conditions and rerun the script.

- Does the upload fail?

If the LIB_MAPPING table checks pass, the script begins to delete existing instances, re-create the instances in the highest federated area, and upload the modified input.

	<p>The script calls the <code>irm_rest_upload_file</code> macro to perform the upload. This macro sends an HTTP request to the server and the server returns the status code for the upload.</p> <p>The file upload might fail to execute the <code>irm_rest_upload_file</code> macro. There is no way to check if the macro executed unless the macro is executed. The HTTP return status code contains information about whether the upload is successful or not.</p> <p>The most common return error code is "500 Internal Server Error". When this error code is returned, the script terminates the process and generates the following error message in the log file:</p> <pre>ERROR: the upload failed for file &filename. in instance &topInstanceKey.</pre> <p>When this error occurs, the original instance has been deleted during the process but not re-created in the highest priority federated area. However, the system has backed up the Excel files for all of the modified input under <code>persistent_area/backup/&orig_instance_key</code> for all of the modified instance. This enables the files associated with the instance to be retrieved.</p> <p>Recommended Actions:</p> <ol style="list-style-type: none"> 1. Verify that the data model stays the same for the data set in the consolidated federated area as with the data set to be uploaded. If the data models do not match, the upload fails. 2. Verify that the LIB_MAPPING table provides correct mapping for file paths to libraries. 3. Verify that the data file to be uploaded resolves with the correct library name. 4. If necessary, fix any conditions that are causing the error and rerun the script.
Re-create instance fails	<p>If a job flow instance is not successfully re-created, examine the instance and view the <code>recreate_flow_instances</code> log to determine the cause of failure. Instances that are not successfully indicated are a  in the web application interface)</p>

	<p>If the failure was caused by an unassigned library, data sets that are not found, or similar type problems, this might indicate that all required inputs were not correctly merged into the highest priority federated area.</p>
<p>The upload of a modified data sets fails.</p>	<p>Instances that contain modified inputs are re-created by the script. The modified data sets are located in the persistent area of the re-created instance. In the web application, the instance is displayed with a  indicator.</p> <p>If there are no errors in the log associated with the failed upload, view the instance's log to identify the cause of the failure. Typically, this failure occurs when all of the required inputs were not correctly merged into the highest priority federated area.</p>