

SAS Visual Investigator 10.5

Hotfix 4

Release Notes

Document version: 1

The correct bibliographic citation for this document is as follows: SAS Institute Inc. 2019. *SAS Visual Investigator 10.5 Hotfix 4*. Cary, NC: SAS Institute Inc.

SAS Visual Investigator 10.5 Hotfix 4

Copyright © 2019, SAS Institute Inc., Cary, NC, USA

All rights reserved. Produced in the UK.

U.S. Government Restricted Rights Notice: Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions 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® Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing website at support.sas.com/publishing or call 1-800-727-3228.

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.

Contents

Release details	4
Updating SAS Visual Investigator Software	4
Overview	5
User Requirements for Performing the Update	5
Synchronize the Mirror Repository	5
(Optional) List the Packages That Are Available for Update.....	6
Generate a New Ansible Playbook.....	6
Post-Update Tasks.....	9
SAS Hadoop Plug-ins	9
Re-index and Resolve Data	9
Product changes	10
Product changes in SAS Visual Investigator 10.5 Hotfix 4.....	10
Product changes in SAS Visual Investigator 10.5 Hotfix 3.....	10
Product changes in SAS Visual Investigator 10.5 Hotfix 2.....	10
Issues addressed	10
Issues addressed in SAS Visual Investigator 10.5 Hotfix 4.....	10

Release details

This release		
Product Name	SAS Visual Investigator 10.5 Hotfix 4	
Supported SAS software		
Facet name	DU name	Version
svi-datahub	datahub	14.0.51
elasticsearch	elasticsrch	5.2.15
searchguard	searchguard	Not updated.
svi-ai	svi-ai	Not updated.
svi-alert	svi-alert	Not updated.
svi-audit	sviaudit	3.1.102
svi-configuration	sviconfig	Not updated.
svi-entity-resolution	svientity	Not updated.
svi-feature	svifeature	Not updated.
sasmobileinvestigator	svimobile	Not updated.
svi-network-analytics	svintwkanlyt	Not updated.
svi-sand	svi-sand	10.0.48
svi-vsd-service	sviscnrioadm	10.1.90
svi-transport	svi-transport-service	4.2.101
sasvisualinvestigator	svi-visual-investigator	7.4.54

Updating SAS Visual Investigator Software

The following instructions describe how to update your SAS Visual Investigator Software to the correct version.

Important

You can no longer use the steps in the *Update with Yum* and *Update with Zypper* sections in the “Updating Your SAS Viya Software” section of the SAS Viya 3.4 for Linux: Deployment Guide.

Overview

An update replaces some or all of your deployed software with the latest version of the software. You perform the update with the same command that was used to install SAS Viya, and use the same software order.

- To see what updates are available for your deployed software, go to the SAS Viya Hot Fix Availability web page at http://ftp.sas.com/techsup/download/hotfix/HF2/Viya_home.html.
- Use Ansible to update your installation.
- SAS might update components of the Ansible playbook that is used to deploy your SAS software. You will need to download the current version of the SAS Orchestration CLI to generate a new Ansible playbook for your deployment, and then run the new Ansible playbook.
- SAS Data Science is deployed as programming-only. However, to obtain the visual interface for the components of SAS Data Science, different products must be licensed. The resulting order should be deployed using the instructions in Adding SAS Viya Software. Contact your SAS account representative to determine the exact products that are required.

Here are other considerations when preparing for an update:

- The update process preserves any user-modified configuration values in the **vars.yml** file, but changes made to other files in the deployment might be lost. Therefore, SAS recommends that you make changes to **vars.yml** on all tenants when possible to avoid any loss of customizations that you made to other files.
Note: If you have added or changed any Consul keys, you should add those to the **vars.yml** or make the changes again after the upgrade.
- You will need the location of the directory on each machine where you stored deployment and maintenance files. For more information about this directory, see [Store the Playbook](#).
Note: If you have added any CAS servers to your initial deployment, you must update those CAS servers as well as each machine in your initial deployment.
- If you are using a PDF version of this guide, go to the Deployment Guides web page at <https://support.sas.com/en/documentation/install-center/sas-viya/deployment-guides.html> and verify that you have the latest version of the deployment documentation before you start the update process. The release date of each document is located in the bottom right corner of the front page.
- Updating SAS Viya software requires an outage period because some SAS Viya services are stopped and restarted automatically during the update process. The update process is the same regardless of whether the deployment is single-tenant or multi-tenant.

User Requirements for Performing the Update

To perform the update process, you must have administrator privileges for the machine. In addition, your account must have superuser (sudo) access. To verify sudo user privileges, run the following command: `sudo -v` or `sudo -l`.

Synchronize the Mirror Repository

If you are using a mirror repository, perform these steps.

1. (Optional) To list the packages that are available for the update process, run the following command on the machine where the mirror repository is located:

Note: Enter the command on a single line. Multiple lines are used here to improve readability.

```
mirrormgr mirror diff --deployment-data path-to-deployment-zip-file-from-SOE
--path path-to-mirror-destination --latest
```

Note: The `mirrormgr mirror diff` command returns the available files for all supported platforms of the products in the deployment. To filter out unwanted content from the output of the command, use a pipe and the `grep` command. For example, add the following at the end of the preceding command to filter out all file names containing `suse`:

```
| grep -v "suse"
```

2. Before performing an update, you must synchronize the mirror repository with SAS. Use the same options to update the mirror repository that you used to create the mirror repository. For more information, see [SAS Mirror Manager and the Mirror Repository](#).

To synchronize, run the following command on the machine where the connected mirror repository is located:

Note: Enter the command on a single line. Multiple lines are used here to improve readability.

```
mirrormgr mirror --deployment-data path-to-deployment-zip-file-from-SOE
--path path-to-mirror-destination --latest
```

3. If you are in a deployment without internet access, move the files from the machine where the connected mirror repository is located to the machine where the unconnected mirror repository is located.

(Optional) List the Packages That Are Available for Update

To list the packages that are available for the update process, run the following command:

on Red Hat Enterprise Linux:

```
sudo yum check-update "sas-*
```

Generate a New Ansible Playbook

You will need the location of the directory on each machine where you stored deployment and maintenance files.

To generate and apply a new Ansible playbook for your deployment:

1. (Optional) Record the existing list of installed software before you begin.
 - a. On each machine in your deployment, create a file that lists the names and versions of the RPM packages of the SAS Viya software that are installed. Create this file in the directory on each machine where you stored deployment and maintenance files. For more information about this directory, see [Store the Playbook](#).

Run the following command to create a text file that lists all the RPM packages:

```
sudo rpm -qg SAS > /sas/install/viya_rpms.txt
```

- b. On each machine in your deployment, create a file that lists the yum groups or packages that are installed. Create this file in the directory on each machine where you stored deployment and maintenance files. For more information about this directory, see [Store the Playbook](#).

- Run the following command to create a text file that lists the yum groups on Red Hat Enterprise Linux:

```
sudo yum grouplist "SAS*" > /sas/install/viya_yumgroups.txt
```

- Run the following command to create a text file that lists the RPM packages on SUSE Linux:

```
sudo rpm -qa | grep "sas-" > /sas/install/viya_packages.txt
```

Note: If you receive a message such as the following, it can be ignored.

```
Repository i is listed more than once in the configuration
```

2. Use the Software Order Email (SOE) for your original deployment to download the current version of the SAS Orchestration CLI.

3. Using the SAS Orchestration CLI that you downloaded, create a new playbook using the instructions on the SAS Orchestration Command Line Interface (CLI) download site. For more information, see [Create a Playbook](#).

Note: To modify the type of deployment from programming-only to full, use the correct `--deployment-type` command line option.

4. You must extract the new playbook to a location that is different from that of your original playbook. For example, if you extracted your original playbook to `/sas/install/`, you might extract the new playbook to `/sas/upgrade/` instead. You must extract the new playbook to a location that is different from the one that you used for your deployment for these reasons:

- a. To preserve the original `vars.yml` file and the inventory file.
- b. To ensure that the playbook directory correctly reflects what is delivered. If a new playbook is accidentally extracted over an existing playbook, files that were removed in the newer playbook would still be available and could negatively affect the process for researching and resolving deployment issues.

To extract the new playbook, use a command that is like the following:

```
tar xf SAS_Viya_playbook.tgz -C /sas/upgrade/
```

5. Merge the `vars.yml` file and the inventory file from the previous deployment into the new playbook. If the previous inventory file contains any spaces that are used to indent machine names, do not include the extra spaces.

- a. Compare the two `vars.yml` files and compare the two inventory files since there could be additions or changes in the newer set of files.

```
diff /sas/install/sas_viya_playbook/vars.yml
    /sas/upgrade/sas_viya_playbook/vars.yml
```

```
diff /sas/install/sas_viya_playbook/inventory-file
    /sas/upgrade/sas_viya_playbook/inventory.ini
```

- b. If the new files contain new content, then merge your customized edits from the two original files into the two new files. If a key/value pair in the original file is not included in the new file, you do not need to add the key/value pair to the new file. If you have any questions, contact SAS Technical Support.
- c. If the original `vars.yml` file from the deployment that is being upgraded contains a value for the `casenv_tenant` variable, it must be removed. Run the following commands to remove the registered CAS service.

Note: Enter each command on a single line. Multiple lines are used here to improve readability.

```
cd /opt/sas/viya/home/bin
./sas-bootstrap-config \
--token-file
/opt/sas/viya/config/etc/SASSecurityCertificateFramework/tokens/consul/default/c
lient.token \
agent service deregister \
"cas-{casenv_tenant}-default-http"
./sas-bootstrap-config \
--token-file
/opt/sas/viya/config/etc/SASSecurityCertificateFramework/tokens/consul/default/c
lient.token \
agent service deregister \
"cas-{casenv_tenant}-default"
```

- d. If you have questions about whether to add a key/value pair from an original file to the new file, contact SAS Technical Support.

6. If you have deployed SAS Event Stream Processing or SAS Event Stream Manager, perform the following steps:

- a. Stop the SAS Event Stream Processing Studio (esvvm) service by running the following command on Red Hat Enterprise Linux 6.x:

```
sudo service sas-viya-esvvm-default stop
```

on Red Hat Enterprise Linux 7.x or SUSE Linux:

```
sudo systemctl stop sas-viya-esvvm-default
```

- b. If you installed Streamviewer, stop the Streamviewer process:

```
$DFESP_HOME/bin/dfesp_xml_client -url "http://host-name:http-port/exit"
```

Replace *host-name* with the host name of the machine where Streamviewer is running.

Replace *http-port* with the port number that you provided when you started Streamviewer with the start-up script.

For more information, see [Starting Streamviewer](#).

- c. Stop the Metering Server:

```
dfesp_xml_client -url "http://host-name:http-port/SASESP/exit"
```

Replace *host-name* with the host name of the machine where the Metering Server is running.

Replace *http-port* with the port number for the Metering Server. By default, it uses port 31001.

7. To apply the new Ansible playbook, change to the directory where the new playbook is located:

```
cd /sas/upgrade/
```

Run the following command:

```
ansible-playbook site.yml
```

8. If you removed the CAS service that is associated with a `casenv_tenant` variable (described in Step 3), update any bookmarked URLs to remove that value and use `cas-shared-default-http` instead. For example, if your original deployment contained a `casenv_tenant` value of `viya32`, you should change it from `http://host.company.com/cas-viya32-default-http` to `http://host.company.com/cas-shared-default-http`.

Note: Do not include `casenv_tenant` in your new `vars.yml`. This variable is no longer used.

9. (Optional) Record the new list of installed software.

- a. On each machine in your deployment, create a file that lists the names and versions of the RPM packages of the SAS Viya software that are installed. Create this file in the directory on each machine where you stored deployment and maintenance files. For more information about this directory, see [Store the Playbook](#).

- Run the following command to create a text file that lists the RPM packages:

```
sudo rpm -qq SAS > /sas/install/new_viya_rpms.txt
```

- b. On each machine in your deployment, create a file that lists the SAS yum groups or packages that are installed. Create this file in the directory on each machine where you stored deployment and maintenance files. For more information about this directory, see [Store the Playbook](#).

- Run the following command to create a text file that lists the yum groups on Red Hat Enterprise Linux:

```
sudo yum grouplist "SAS*" > /sas/install/new_viya_yumgroups.txt
```

- Run the following command to create a text file that lists the RPM packages on SUSE Linux:

```
sudo rpm -qa | grep "sas-" > /sas/install/new_viya_packages.txt
```

Note: If you receive a message such as the following, it can be ignored.

```
Repository repository-name is listed more than once in the configuration
```

You can see the differences between the previous and current deployments by comparing the lists of installed software that precedes the update (Step 1) and that follows the update.

Post-Update Tasks

After you complete an update, you must perform additional tasks. This section lists those updates, how to determine whether the update has been applied, and the manual steps required to finish the installation.

SAS Hadoop Plug-ins

Note: If you are not using SASHDAT on HDFS, skip this section.

If you receive an updated RPM for SAS Hadoop Plug-ins, for each Hadoop environment, you must determine whether to install it. Use the newer version of the RPM in each Hadoop environment. Install the updated RPM only if its version is newer than the one that is currently installed.

1. Install the updated RPM for SAS Hadoop Plug-ins on the CAS controller.
2. To find out the version number of the RPM that was just installed:

```
rpm --queryformat "%{VERSION}" -q sas-hdatplugins
```

The version number will be displayed in a format of three two-digit numbers separated by periods. Examples are 03.13.00 or 03.04.02. The variable *VERSION* is used to represent the actual version number in your environment.

3. On each Hadoop name node or data node, use the appropriate method to find the currently installed version:

- For SAS Hadoop Plug-ins that were installed on the Hadoop cluster with the `sashdat-install.sh` command:

```
cat /opt/sas/HDATHome/SAS_VERSION
```

The output should include a version number in the *VERSION* format that is shown in the output of Step 2.


- For SAS Hadoop Plug-ins that were installed on the Hadoop cluster with Cloudera Manager, the parcel name should be `SASHDAT-VERSION.p0.1`.
 - For SAS Hadoop Plug-ins that were installed on the Hadoop cluster with Ambari, the stack name should be `SASHDAT-VERSION.s01`.
4. If the installed version that is displayed by Step 3 is equal to or greater than the version number of the RPM displayed by the command in Step 2, then skip the rest of this task.
 5. If the installed version that is displayed by Step 3 is less than the version number of the RPM that is displayed by the command in Step 2, follow the steps described in [Deploying SAS Plug-ins for Hadoop](#) to update your Hadoop environment to the newer version of the SAS Hadoop Plug-ins.

Re-index and Resolve Data

As a result of upgrading to Elasticsearch 5.6.12 and the changes to the index structure for child entity searching, after upgrading to SAS Visual Investigator 10.5 Hotfix 4 you must re-index and resolve all your data.

Important

You do not need to re-index and resolve data if you are upgrading from SAS Visual Investigator 10.5 Hotfix 3 to SAS Visual Investigator 10.5 Hotfix 4.

1. In the administrative application, click **Data Objects**.
2. On the **Data Objects** toolbar, click **Entities**.
3. On the toolbar, click  and select **Re-index and resolve entities** from the menu.

You are warned that re-indexing and resolving all entities might take a long time.

4. Click **Yes** to continue.

A pop-up message notifies you that the re-indexing and resolving process has started.

You can view the progress of the re-index, and whether it completes successfully by clicking **Jobs**.

Product changes

Product changes in SAS Visual Investigator 10.5 Hotfix 4

There are no product changes in this release of SAS Visual Investigator 10.5 Hotfix 4.

Product changes in SAS Visual Investigator 10.5 Hotfix 3

For information about product changes and issues addressed in SAS Visual Investigator 10.5 Hotfix 3, see [SAS Visual Investigator 10.5 Hotfix 3 Release Notes](#)

Product changes in SAS Visual Investigator 10.5 Hotfix 2

For information about product changes and issues addressed in SAS Visual Investigator 10.5 Hotfix 2, see [SAS Visual Investigator 10.5 Hotfix 2 Release Notes](#).

Issues addressed

This section gives details of issues that were addressed in the SAS Visual Investigator 10.5 Hotfix 4.

Issues addressed in SAS Visual Investigator 10.5 Hotfix 4

SAS Note	Details
—	<p>When viewing SAS Visual Investigator audits, you can now choose to hide the Session ID value.</p> <p>To hide the Session ID on a single tenant, set the following Consul key to <code>false</code>:</p> <pre>config/application/tenants/{tenant}/svi-audit/displaySessionID</pre> <p>To hide the Session ID on multiple tenants, set the following Consul key to <code>false</code>:</p> <pre>config/svi-audit/svi-audit.displaySessionID</pre> <p>If this key is undefined or set to <code>true</code>, the Session ID value will show in SAS Visual Investigator audits.</p>
—	<p>An issue has been addressed where accessing pages might fail and display an authentication error.</p>
—	<p>Fields which have the Contains HTML option selected are now validated against an HTML sanitizer. Administrators can create a whitelist of accepted HTML tags and attributes. Items that are not in the whitelist are HTML encoded to mitigate potential security issues. For example:</p> <p>Entering <code><p><ABC123></p></code> when the Contains HTML setting is selected, stores the string as <code><p>&lt;ABC123&gt;</p></code>.</p> <p>The same string is stored as <code><p><ABC123></p></code> if the Contains HTML setting is not selected.</p>
—	<p>Labels added to an object on the Dates and Locations tab in the administrator application now correctly display when the object is viewed on a map or on the timeline.</p>
—	<p>Creating an external entity with a string field populated with HTML containing the <code>onerror</code> attribute no longer causes an error to display in a user's browser.</p>

SAS Note	Details
—	Users can now clear a selection from a drop-down list if the field and control are not mandatory.
—	Adding a Category when creating or updating a data object no longer fails when the object is saved.
—	Tenant initialization performance has been improved, in particular for deployments containing large numbers of entities and security rules.
—	The <code>-XX:+DisableExplicitGC</code> option has been removed from BareOS deployment scripts as it could lead to an Elasticsearch out of memory error.
—	When connecting to MySQL and SQL Server data stores, the password is no longer written to the log in plain text.
—	<p>Entering text contained inside angle brackets (for example, <code><ABC123></code>) into a rich text field is now searchable. That is, after saving the contents of a rich text field, users can search for and return results for <code>ABC123</code> by doing one of the following:</p> <ul style="list-style-type: none"> • Enter the search text in quotes. For example: <code>"<ABC123>"</code> • Remove the brackets. For example: <code>ABC123</code> • Use the query builder with the Phrase operator. For example: <code>+_type:Customers +description:"<ABC123>"</code> • Use the Home page form search. For example: <code>Customers AND description: "<ABC123>"</code>
—	SAS Visual Investigator now validates the contents of the Name and Description fields when adding an attachment to an object.
—	Sand service requests no longer fail when an Elasticsearch node becomes unavailable. The request is now sent to another Elasticsearch node.
—	When you perform a network expansion, you can no longer change view until the expansion completes, as indicated by a loading indicator. This means that objects returned remain consistent across different views.
—	Exporting table data to Microsoft Excel now correctly includes the Last Modified Date column.