

SAS Visual Investigator 10.8 Hotfix 8.1 Release Notes



The correct bibliographic citation for this document is as follows: SAS Institute Inc. 2024. *SAS Visual Investigator 10.8 Hotfix 8.1*. Cary, NC: SAS Institute Inc.

SAS Visual Investigator 10.8 Hotfix 8.1

Copyright © 2024, 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 <https://support.sas.com/en/books.html> 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	5
Overview	5
User Requirements for Performing the Update	6
Synchronize the Mirror Repository	6
(Optional) List the Packages That Are Available for Update.....	7
Update with Ansible.....	7
Post-Update Tasks.....	9
SAS Hadoop Plug-ins	9
Re-index and Resolve Data	9
Update the Type of Deployment with Ansible	10
Overview	10
Update the Type of Deployment.....	10
Generate a New Ansible Playbook.....	10
Product changes.....	13
Issues addressed.....	13
Issues addressed in SAS Visual Investigator 10.8 Hotfix 8.1.....	14

Release details

This release		
Product Name	SAS Visual Investigator 10.8 Hotfix 8.1	
Supported SAS software		
Facet name	DU name	Version
svi-datahub	datahub	16.8.46
elasticsearch-secure	elasticsearchsecure	8.5.21
searchguard	searchguard	N/A
svi-ai	svi-ai	6.8.15
svi-alert	svi-alert	14.7.29
svi-audit	sviaudit	5.8.13
svi-configuration	sviconfig	5.3.0
svi-data-export	svi-data-export	2.7.11
svi-entity-resolution	svientity	5.7.10
svi-feature	svifeature	7.7.12
svi-indexer	svi-indexer	1.8.42
sasmobileinvestigator	svimobile	4.8.35
svi-network-analytics	svintwkanlyt	6.7.11
svi-sand	svi-sand	12.7.38
svi-vsd-service	sviscnrioadm	12.7.50
svi-scorecard	svi-scorecard	2.7.11
svi-transport	svi-transport-service	6.8.9
sasvisualinvestigator	svi-visual-investigator	9.11.138

Important

The Data Hub and Indexer services will always use "Etc/UTC" as the JVM Default Time Zone. Any attempts to change the JVM time zone for Data Hub and Indexer will be ignored.

Updating SAS Visual Investigator Software

Services delivered in the SAS Visual Investigator 10.8 Hotfix 8.1 release will fail to start if they are being updated from a SAS Visual Investigator 19w47 software order that does not contain SAS Visual Analytics. To ensure the update process completes successfully, you will require a new software order and upgrade-in-place deployment.

The following procedure can be used to determine if such actions are required:

1. Locate the YUM repository definitions on the Linux host containing the deployed version of SAS Visual Investigator.
2. Check if SAS Visual Analytics is included in the deployment. If included, a new software order is **not** required.
 - `grep -q '/va/' /etc/yum.repos.d/sas* && echo "VA is included"`
3. Check the repository version for SAS Viya common components. If version 126 is accessible, a new software order is **not** required.
 - `grep -q '/vcmnfnd/126/' /etc/yum.repos.d/sas* && echo "Order newer than 19w47"`

Overview

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 3.5 for Linux: Deployment Guide.

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 the same tool (Ansible, yum, or Zypper) to update that you used to install. For example, if you used an Ansible playbook for your initial installation, update with Ansible.
- If you want to update your deployment from programming-only to full, you can use only Ansible.
- 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 to a Deployment and Upgrading Products in SAS Viya 3.5](#). Contact your SAS account representative to determine the exact products that are required.
- SAS might update components of the Ansible playbook that are used to deploy your SAS Viya 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.

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** when possible in order to avoid any loss of customizations that you made to other files.
- 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/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 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. Synchronize the deployment's mirror repository with SAS's mirror repository. 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-"
```

Update with Ansible

To update a SAS Viya deployment using Ansible:

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 -qq 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
```

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
```

2. Review the *_deployment.* files (for example, **casconfig_deployment.lua**) in the existing deployment for any user-modified changes. If there are any user-modified changes to the *_deployment.* files, back up the files and update the **vars.yml** file with the changes before you perform the update. If you have questions, contact SAS Technical Support.

Note

SAS recommends that you add your customizations to the **vars.yml** file rather than to a *_deployment.* file in order to preserve your customizations. Otherwise, your customizations would be lost during the update process.

By default, the update process backs up the following files:

For CAS:

```
/opt/sas/viya/config/etc/cas/default/cas_usermods.settings
```

```
/opt/sas/viya/config/etc/cas/default/casconfig.lua
```

```
/opt/sas/viya/config/etc/cas/default/cas.hosts
```

For SAS Object Spawner:

```
/opt/sas/viya/config/etc/spawner/default/spawner.cfg
```

For SAS/CONNECT:

```
/opt/sas/viya/config/etc/sysconfig/connect/default/sas-connect
```

3. 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 (espm) service by running the following command:

- For Red Hat Enterprise Linux 6.x:

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

- 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 service:

```
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.

4. To verify the health of the SAS Infrastructure Data Server before running the playbook, perform the task in [Verify the Health of SAS Infrastructure Data Server](#).
5. When you initiate the update, use the same command and options that you used when you performed the initial deployment. For more information, see [Deploy the Software](#).

If you added additional CAS servers to your deployment, run the command for each CAS server. When you added the CAS servers to your deployment, you created **inventory.ini** and **vars.yml** files for each additional CAS server. Those same files must be specified in the command to initiate the update.

Note

Before running the playbook, ensure that the following services are running:

- Consul
- Vault
- Local Consul
- HA Postgres Consul Template

To initiate the update:

```
ansible-playbook -i CAS-server-inventory-file-name site.yml -e "@CAS-server-vars-file-name"
```

Be sure to add the appropriate options that you used for the initial deployment of the CAS servers. Repeat this command for each additional CAS server.

6. (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 -qg SAS > /sas/install/new_viya_rpms.txt
```
- b. On each machine in your deployment, create a file that lists the yum groups or 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 yum groups on Red Hat Enterprise Linux:

```
sudo yum grouplist "SAS*" > /sas/install/new_viya_yumgroups.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 list of installed software that precedes the update ([Step 1](#)) and the list that follows the update.

- c. To verify that a specific update was applied, compare the contents of the text file created in [Step 6a](#) to the packages listed for the specific update. The package list for a specific update is available in the Manifest View for the update on the SAS Viya Hot Fix Availability web page at [http://ftp.sas.com/techsup/download/hotfix/HF2/Viya\\_home.html](http://ftp.sas.com/techsup/download/hotfix/HF2/Viya_home.html).

## 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

If you are updating a single tenant environment from SAS Visual Investigator 10.6 to SAS Visual Investigator 10.8 Hotfix 8.1 where your relationships use a dedicated index, you must re-index all relationships.

**Note**

Multi-tenant installations are unaffected.

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**.

# Update the Type of Deployment with Ansible

## Overview

Using Ansible, you can modify your deployment from programming-only to full.

**Note**

If you make changes to the **proxy.conf** file and then rerun the playbook, those changes are overwritten. A copy of the **proxy.conf** should be created in the `/etc/httpd/conf.d` directory when you rerun the playbook. Use this copy, along with the instructions in Post-installation Tasks, to make changes to the updated **proxy.conf** file.

## Update the Type of Deployment

To update a SAS Viya deployment from programming-only to full:

1. Complete the task in [Generate a New Ansible Playbook](#).

**Note**

Do not select the programming-only option in the SAS Orchestration CLI when generating the new playbook. The default playbook generated will be for a full deployment.

2. To change from a programming-only deployment to a full deployment, you must complete the configuration for the visual attributes of the deployment. For more information, see [Post-installation Tasks](#).

**Note**

When the programming-only interface is deployed, SAS Studio is accessible on port 7080. However, when updating from a programming-only deployment to a full deployment, SAS Studio uses a dynamically assigned port deployment.

## Generate a New Ansible Playbook

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](#).

If updates are needed in the Ansible playbook, 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 -qq 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
```

#### 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
```

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:

- To preserve the original `vars.yml` file and the inventory file.
- 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 similar to 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 inventory files to check for additions or changes in the newer set of files. Be sure to evaluate the comments to determine whether the requirements for host groups changed between releases of the software.

```
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 of the three commands 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/client.token

agent service deregister "cas-{casenv_tenant}-default-http"

./sas-bootstrap-config --token-file

/opt/sas/viya/config/etc/SASSecurityCertificateFramework/tokens/consul/default/client.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. Run the following command:

For Red Hat Enterprise Linux 6.X:

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

- 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 5), 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`, 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 -qg 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
```

**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 list of installed software that precedes the update (Step 1) and the list of installed software that follows the update.

- c. To verify that a specific update was applied, compare the contents of the text file created in Step 9a to the packages listed for the specific update. The package list for a specific update is available in the Manifest View for the update on the SAS Viya Hot Fix Availability web page at [http://ftp.sas.com/techsup/download/hotfix/HF2/Viya\\_home.html](http://ftp.sas.com/techsup/download/hotfix/HF2/Viya_home.html).

## Product changes

SAS Visual Investigator is now deployed with OpenSearch. OpenSearch is the core driver for entity indexing and searching.

After upgrading to SAS Visual Investigator 10.8 Hotfix 8.1, you must re-index all entities and all relationships.

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 all 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**.

## Issues addressed

This section gives details of issues that were addressed in the SAS Visual Investigator 10.8 Hotfix 8.1 release.

## Issues addressed in SAS Visual Investigator 10.8 Hotfix 8.1

| SAS Note | Details                                                                                                                                             |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| —        | An issue has been addressed that caused scenario-fired events from previous Alert versions to display on new alerting events.                       |
| —        | When importing a flow, validation messages have been improved in situations where two relationships exist for the same tables and join keys.        |
| —        | An issue has been addressed that enabled users without the Export from tables and grids capability to export data from grids using the column menu. |