SAS Visual Investigator 10.5.1 Hotfix 2

Release Notes

Contents

Release details

This release		
Product Name	SAS Visual Investigator 10.5.1 Hotfix 2	
Supported SAS software		
Facet name	**DU name**	**Version**
svi-datahub	datahub	14.1.165
Elasticsearch	elasticsrch	5.3.7
searchguard	searchguard	1.3.18
svi-ai	svi-ai	Not updated.
svi-alert	svi-alert	11.4.168
svi-audit	sviaudit	Not updated.
svi-configuration	sviconfig	Not updated.
svi-entity-resolution	svientity	Not updated.
svi-feature	svifeature	5.4.174
sasmobileinvestigator	svimobile	Not updated.
svi-network-analytics	svintwkanlyt	Not updated.
svi-sand	svi-sand	10.1.179
svi-vsd-service	sviscnrioadm	10.2.128
svi-transport	svi-transport-service	Not updated.
sasvisualinvestigator	svi-visual-investigator	7.5.131
sasadaptivelearning	adaptivelearning	Not updated.
adaptiveml	adaptiveml	Not updated.

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 <u>SAS Mirror Manager and the Mirror Repository</u>.

   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 <u>Store the Playbook</u>.

      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 <u>Store the Playbook</u>.

      • 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 (espvm) service by running the following command on Red Hat Enterprise Linux 6.x:

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

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

```
sudo systemctl stop sas-viya-espvm-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 -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
```

- 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.

Product changes

There are no product changes in the SAS Visual Investigator 10.5.1 Hotfix 2 release.

Issues addressed

This section gives details of issues that were addressed in the SAS Visual Investigator 10.5.1 Hotfix 2 release.

Issues addressed in SAS Visual Investigator 10.5.1 Hotfix 2

SAS Note	Details
___	Boolean operators are now correctly preserved when child object queries are used in conjunction with free text searches.
64630	An issue has been addressed which could cause linking to an existing case disposition to fail if the relationship had additional attributes.
___	Certificates are now correctly re-created in deployments with multiple Vault instances.
64629	The **logback.xml** file has been removed. As a result, the svi-vsd-service logs now contain all required information and use the same pattern as other service logs.
64625	After restarting, the SAS Visual Investigator feature service now generates a single refresh event when reloading all security rules.
___	The number of calls made to the Search Guard Search Action group has been reduced.
___	The number of calls made to return related objects and their labels has been reduced.
___	Scenarios containing multiple aggregations no longer produces incorrect results.
___	SAS Visual Investigator no longer exposes SAS software component version numbers to unauthenticated users.
___	The timeslider now correctly activates nodes on a network diagram when the timeslider passes the valid from date.
64621	An issue has been addressed which could cause saving changes to a flow in Scenario Administrator to fail.
___	The HtmlSanitizer whitelist regular expression no longer causes an object save to fail due to a StackOverflowError.
___	Bulk loading Roles and Role Mappings has been ported to Search Guard for Elasticsearch 5.6.12.
___	An issue has been addressed which could cause a horizontal scrollbar to appear at the bottom of the Home page when using a column layout.
___	An issue has been addressed which could see Search Guard fail on VerifyElasticsearchRunning if the system was set to a language other than English.
___	SAS Visual Investigator now uses the value of **rootLibPath** instead of the **tmp** directory for CAS global libraries. After upgrading to SAS Visual Investigator 10.5.1 Hotfix 2, you must delete the existing CAS Tenant libraries, for example **svi_<tenantId>**, before restarting Scenario Administrator and CAS.
64598	An issue has been addressed which could cause some tenants in a multi-tenant environment to be inaccessible after onboarding.
64597	SAS Visual Investigator users can correctly select options from the **Alert Activity** drop-down list using their mouse.
64571	The Alert scorecard now renders correctly when the scorecard contains more than one entity.
64570	An issue has been addressed which could cause constraint violations when running a flow.

SAS Note	Details
___	Using Advanced Search to create queries combining child entities no longer displays an error when a user runs the query.
64568	An issue has been addressed which could cause data loss when incrementally re-indexing relationships.
___	SAS Visual Investigator now correctly encodes and decodes invalid HTML tags in text fields.
64595	SCENOUT_ tables are no longer promoted, and as such, no longer use as much CAS system disk space.
___	A new REST API endpoint has been added to repair damaged data stores.
64594	SAS Visual Investigator now correctly encodes and decodes text input in angle brackets, for example, `<VEHICLE01>`. As a result, the text can now be returned in searches.