SAS® Viya® 3.4 Administration: Infrastructure Servers

Infrastructure Servers: Overview

Note: A programming-only deployment uses only one of the infrastructure servers—Apache HTTP Server.

SAS Viya contains these infrastructure servers:

- “SAS Configuration Server”
- “SAS Secrets Manager (Linux)”
- “SAS Infrastructure Data Server”
- “SAS Message Broker”
- “SAS Cache Locator and Cache Server”
- “Apache HTTP Server”

The following diagram identifies the infrastructure service components of a SAS Viya full deployment:

Figure 1 *SAS Viya Infrastructure Servers (Full Deployment)*

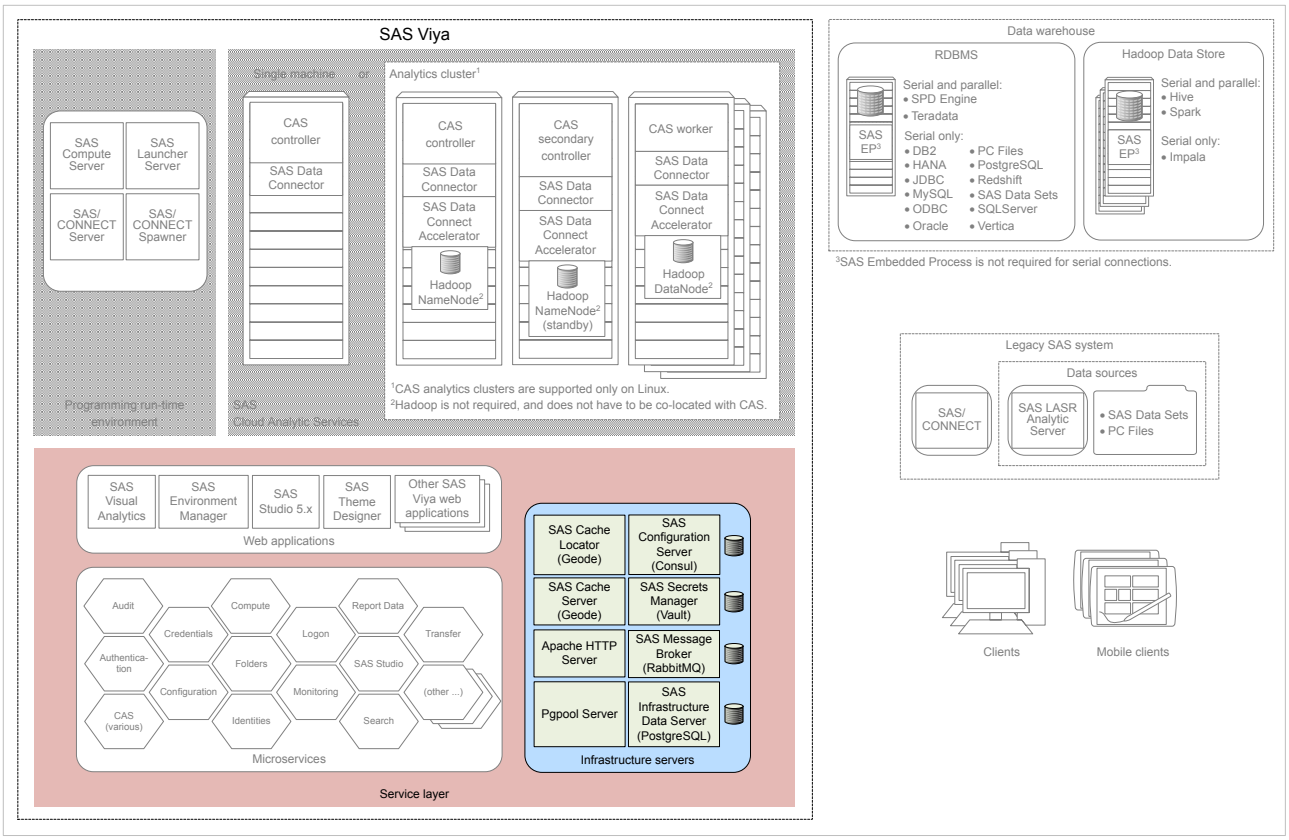

Note: SAS Secrets Manager does not run on Windows.

The following diagram shows that a programming-only deployment uses only the Apache HTTP Server from the SAS Viya server layer:

Figure 2 *SAS Viya Infrastructure Servers (Programming-only Deployment)*

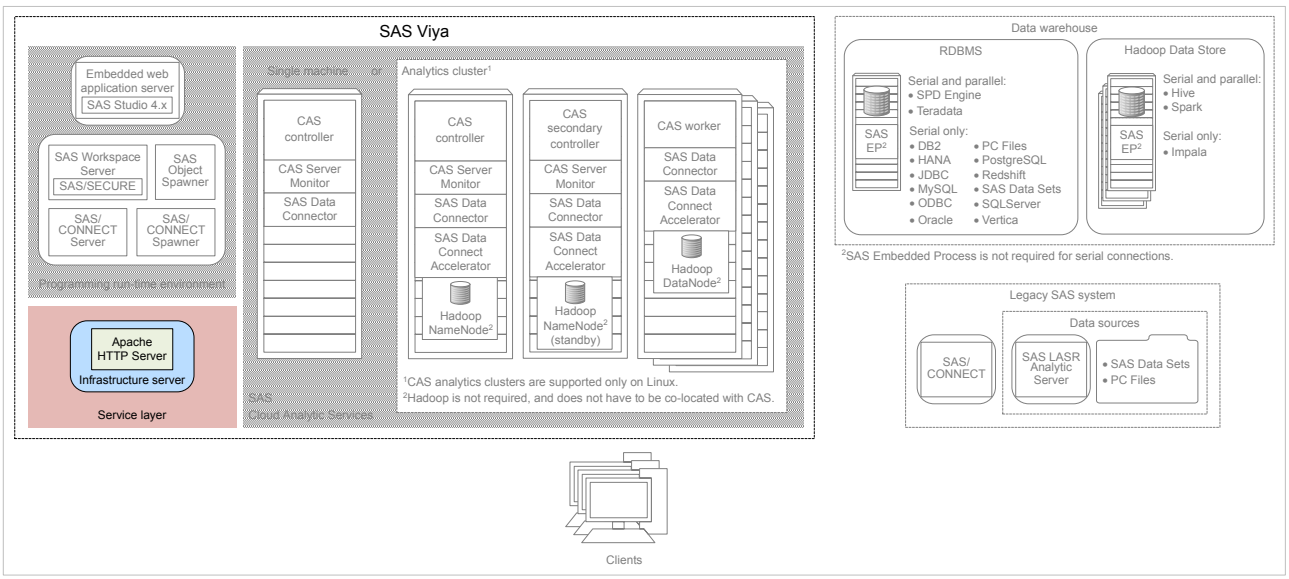

SAS Configuration Server

Overview

SAS Configuration Server is based on HashiCorp Consul. SAS Configuration Server acts as a service configuration registry that serves as a central repository for configuration data, service discovery, and health status.

Note: A programming-only deployment does not use SAS Configuration Server.

Operate (Linux)

SAS Viya provides a script in **/etc/init.d** that you use to stop, start, restart, and check the status of SAS Configuration Server. The script is named, `sas-viya-consul-default`.

Syntax

How you run `sas-viya-consul-default` depends on your operating system:

- Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

  sudo systemctl status | stop | start | restart sas-viya-consul-default

- Red Hat Enterprise Linux 6.x (or an equivalent distribution):

  sudo service sas-viya-consul-default status | stop | start | restart

Usage Notes and Tips

- You must be logged on to the machine where SAS Configuration Server resides. Also, you must have root-level privileges to run this script.
- For multi-machine deployments, run `sas-viya-consul-default` on every SAS Viya machine. Start SAS Configuration Server *first*. Stop SAS Configuration Server *last*.
- If there are multiple instances of SAS Configuration Server, start them in this sequence:
  - First, start `sas-viya-consul-default` on all machines in the `[consul]` host group.

    The `[consul]` host group is located in the Ansible playbook inventory.ini file and it defines which machines host the SAS Configuration Server instances.
  - Next, start `sas-viya-consul-default` on all other machines in the deployment, which launches the agent processes for SAS Configuration Server.
- If there are multiple instances of SAS Configuration Server, stop them in this sequence:
  - First, stop `sas-viya-consul-default` on all machines not in the `[consul]` host group.

    The `[consul]` host group is located in the Ansible playbook inventory.ini file.
  - Next, stop `sas-viya-consul-default` on machines in the `[consul]` host group.
- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.
- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-consul-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-consul-default` on RHEL 7.x with the `service` command, and later attempt to shut down SAS Configuration Server using the `systemctl` command, the configuration server stops responding and does not shut down.

Examples

- To check status of SAS Configuration Server on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status sas-viya-consul-default
```

- To stop SAS Configuration Server on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

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

- To start SAS Configuration Server on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl start sas-viya-consul-default
```

- To restart SAS Configuration Server on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-consul-default restart
```

Operate (Windows)

Using the Services snap-in in the Microsoft Management Console, you can start, stop, and restart SAS Configuration Server (Consul).

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.

Figure 3 *SAS Configuration Server in the Services snap-in*

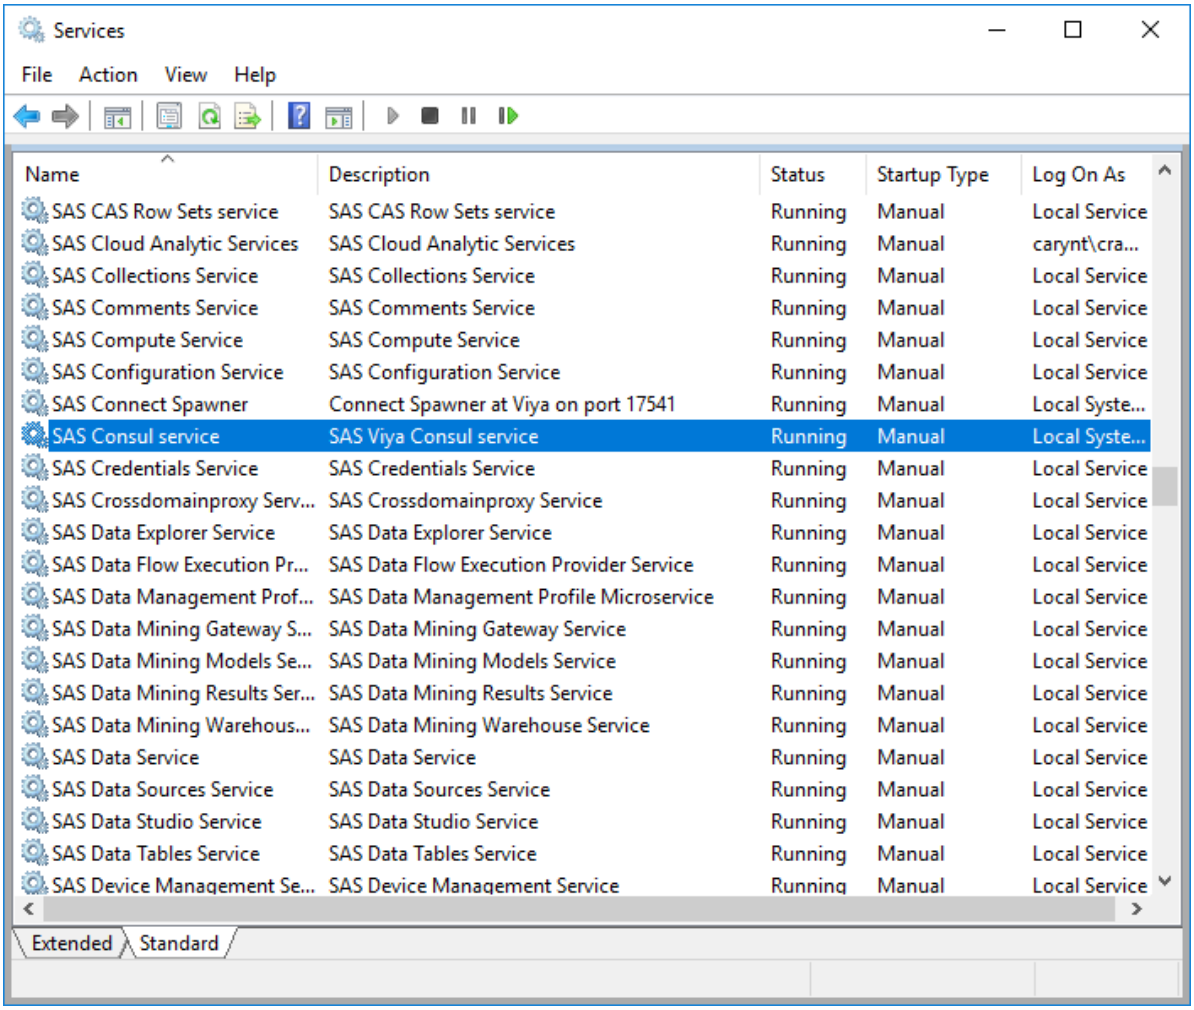

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.

Concepts

What Is SAS Configuration Server?

SAS Configuration Server is based on HashiCorp's Consul. Consul is a distributed, highly available registry that contains service configuration data and availability and overall performance (health) information.

Configuration data resides in SAS Configuration Server as key-value pairs. This data is used by SAS Viya microservices at start-up to load default values and to discover any service dependencies.

During run time, whenever a service's properties change, the service is notified, and it rereads its properties from SAS Configuration Server. (The exceptions are noted in "What Services Must Be Restarted?" in *SAS Viya Administration: Configuration Properties*.)

Each service registers its health checks when it starts. The Monitoring system periodically queries the status of the health checks.

How Does the SAS Configuration Service Work with SAS Configuration Server?

For information about how the SAS Configuration Service works with SAS Configuration Server, see "How SAS Viya Configuration Works" in *SAS Viya Administration: Configuration Properties*.

Log Files

SAS Configuration Server log files are located in **`/opt/sas/viya/config/var/log/consul/default`**.

SAS Secrets Manager (Linux)

Overview

SAS Secrets Manager is based on HashiCorp Vault. SAS Secrets Manager uses Vault to store and generate secrets such as Transport Layer Security (TLS) certificates.

Important: SAS Secrets Manager is not deployed on Windows.

Note: A programming-only deployment does not use SAS Secrets Manager. For more information, see "Deployment Types" in *SAS Viya Administration: Orientation*.

Operate

SAS Viya provides a script in **`/etc/init.d`** that you use to stop, start, restart, and check the status of SAS Secrets Manager. The script is named, `sas-viya-vault-default`.

Syntax

How you run `sas-viya-vault-default` depends on your operating system:

- Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

  `sudo systemctl status | stop | start | restart sas-viya-vault-default`

- Red Hat Enterprise Linux 6.x (or an equivalent distribution):

  `sudo service sas-viya-vault-default status | stop | start | restart`

Usage Notes and Tips

- You must be logged on to the machine where SAS Secrets Manager resides. Also, you must have root-level privileges to run this script.
- For multi-machine deployments, run `sas-viya-vault-default` on every SAS Viya machine that also contains SAS Configuration Server (Consul). SAS Secrets Manager (Vault) is always deployed on the same machine as the Configuration server. (Machines that contain Configuration agents do not have SAS Secrets Manager.)

  Start or restart SAS Secrets Manager immediately *after* you run SAS Configuration Server.

  Stop SAS Secrets Manager immediately *before* you stop SAS Configuration Server.

- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.
- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-vault-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-vault-default` on RHEL 7.x with the `service` command, and later attempt to shut down SAS Secrets Manager using the `systemctl` command, secrets manager stops responding and does not shut down.

Examples

- To check status of SAS Secrets Manager on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status sas-viya-vault-default
```

- To stop SAS Secrets Manager on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

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

- To start SAS Secrets Manager on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl start sas-viya-vault-default
```

- To restart SAS Secrets Manager on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-vault-default restart
```

Concepts

What Is SAS Secrets Manager?

SAS Secrets Manager is based on HashiCorp Vault. Vault is a distributed, highly available server used to manage secrets. A *secret* is information that you want to secure, such as keys, passwords, certificates, and so on. Vault provides a secure interface to secrets, in addition to access control, and audit logging.

Here are some features of secrets manager and examples of how SAS Viya uses them:

- On-demand generation of secrets

  Secrets manager generates TLS certificates for SAS Viya servers at start-up.
- Secure storage for secrets

  Microservices use secure storage so that multiple microservice instances running on the same machine do not request multiple TLS certificates.
- Encrypt and decrypt data without storing it

  SAS Compute Server uses this feature when it sends a password to child processes.
- Revocation of secrets

  SAS Viya services use this feature when rotating security artifacts. (For example, services use vault tokens to request TLS certificates).

For more information, see "Concepts" in *Encryption in SAS Viya: Data in Motion*.

Dependency on SAS Configuration Server (Consul)

SAS Secrets Manager is installed on the same machines where SAS Configuration Server (Consul) resides. SAS Configuration Server contains a namespace where secrets manager stores secrets in encrypted form, which enables all instances of secrets manager access to consistent data. Also, secrets manager relies on the configuration server for locking and leader election. Therefore, in order for SAS Secrets Manager to be operational, the configuration server must be running.

In multiple-machine, fault tolerant deployments, SAS Secrets Manager has a primary (leader) and one or more standbys (hot standbys). For information about SAS Secrets Manager topology, see "Fault Tolerance in SAS Viya (Linux)" in *SAS Viya Administration: General Servers and Services*.

TTL Precedence Rules

The SAS Secrets Manager (Vault) time to live (TTL) properties have certain rules that you must follow. Failure to follow these rules can cause secrets manager not to start.

Figure 4 *Time to Live Properties Precedence Rules*

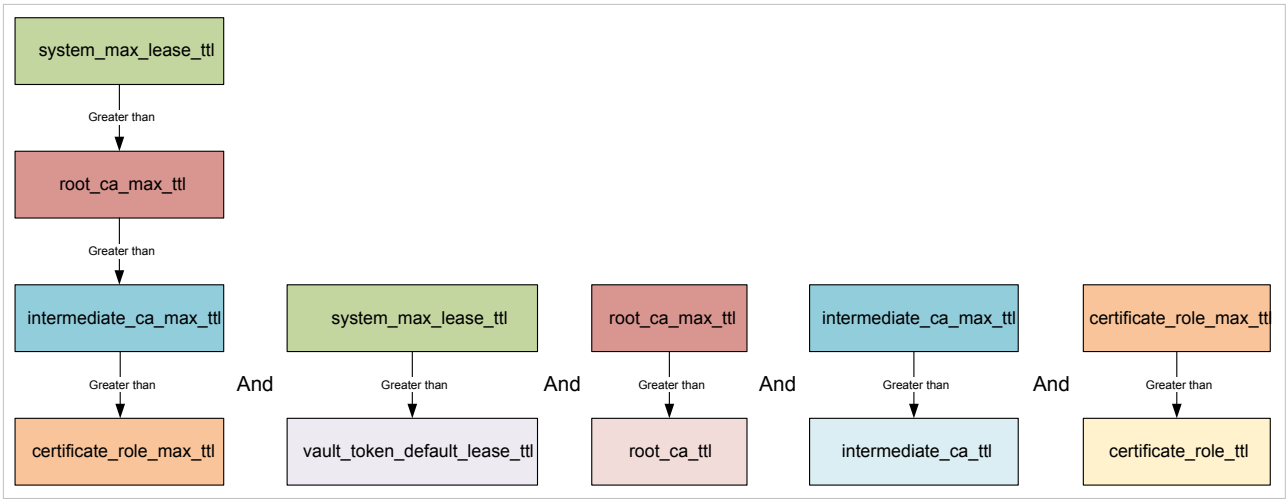

Log Files

SAS Secrets Manager log files are located in `/opt/sas/viya/config/var/log/vault/default`.

SAS Infrastructure Data Server

Overview

SAS Infrastructure Data Server is based on PostgreSQL version 9 and is configured specifically to support SAS software. SAS Infrastructure Data Server stores user content, such as reports, custom groups, comments, authorization rules, selected source definitions, attachments, audit records, and user preferences.

Note: A programming-only deployment does not use SAS Infrastructure Data Server.

How To (Cluster)

Operate a Cluster (Linux)

SAS Viya provides a script in **/etc/init.d** that you use to stop, start, restart, and check the status of a SAS Infrastructure Data Server cluster. The script is named, sas-viya-sasdatasvrc-postgres.

Syntax

How you run sas-viya-sasdatasvrc-postgres depends on your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl status | stop | start | restart sas-viya-sasdatasvrc-postgres

  Important: When using the status argument, sas-viya-sasdatasvrc-postgres does not display the data server node list. For this reason, we recommend that, on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, you make a direct call to sas-viya-sasdatasvrc-postgres when querying the data server status (for example, **cd /etc/init.d**, then **sudo ./sas-viya-sasdatasvrc-postgres status**.

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-postgres status | stop | start | restart

Usage Notes and Tips

- You must be logged on to the machine where the Pgpool server resides. Also, you must have root-level privileges to run this script.
- On multi-tenant SAS Viya systems, the script is named sas-tenant-ID-sas-viya-sasdatasvrc-postgres.
- In the list of running processes you see data server-related processes with **ct** in their names (for example, sas-viya-sasdatasvrc-postgres-node0-**ct**-pg_hba). These are processes that keep the data server configuration files in-sync with SAS Configuration Server (Consul). For more information, see "Synchronizing Configuration Files with SAS Configuration Server" on page 41.
- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.
- On Linux systems that support systemd, use the systemctl command when running sas-viya-sasdatasvrc-postgres. The systemctl command maintains a record of service status that the service command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, do not mix System V init and systemd commands. Mixing the System V init (service command) with the systemd (systemctl command) causes several issues. The systemctl command knows nothing about a SAS Viya service started with the service command. If you start sas-viya-sasdatasvrc-postgres on RHEL 7.*x* with the service command, and later attempt to shut down the data server cluster using the systemctl command, the data server stops responding and does not shut down.

Examples

- To check status of the data server cluster (and see the nodes list) on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
cd /etc/init.d
```

```
sudo ./sas-viya-sasdatasvrc-postgres status
```

- To check status of the data server cluster (and not see the nodes list) on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl status sas-viya-sasdatasvrc-postgres
```

- To stop the data server cluster on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres stop
```

- To start the data server cluster on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-sasdatasvrc-postgres
```

- To restart the data server cluster on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres restart
```

Operate a Cluster (Windows)

Using the Services snap-in in the Microsoft Management Console, you can start, stop, and restart SAS Infrastructure Data Server.

Figure 5 *SAS Infrastructure Data Server in the Services snap-in*

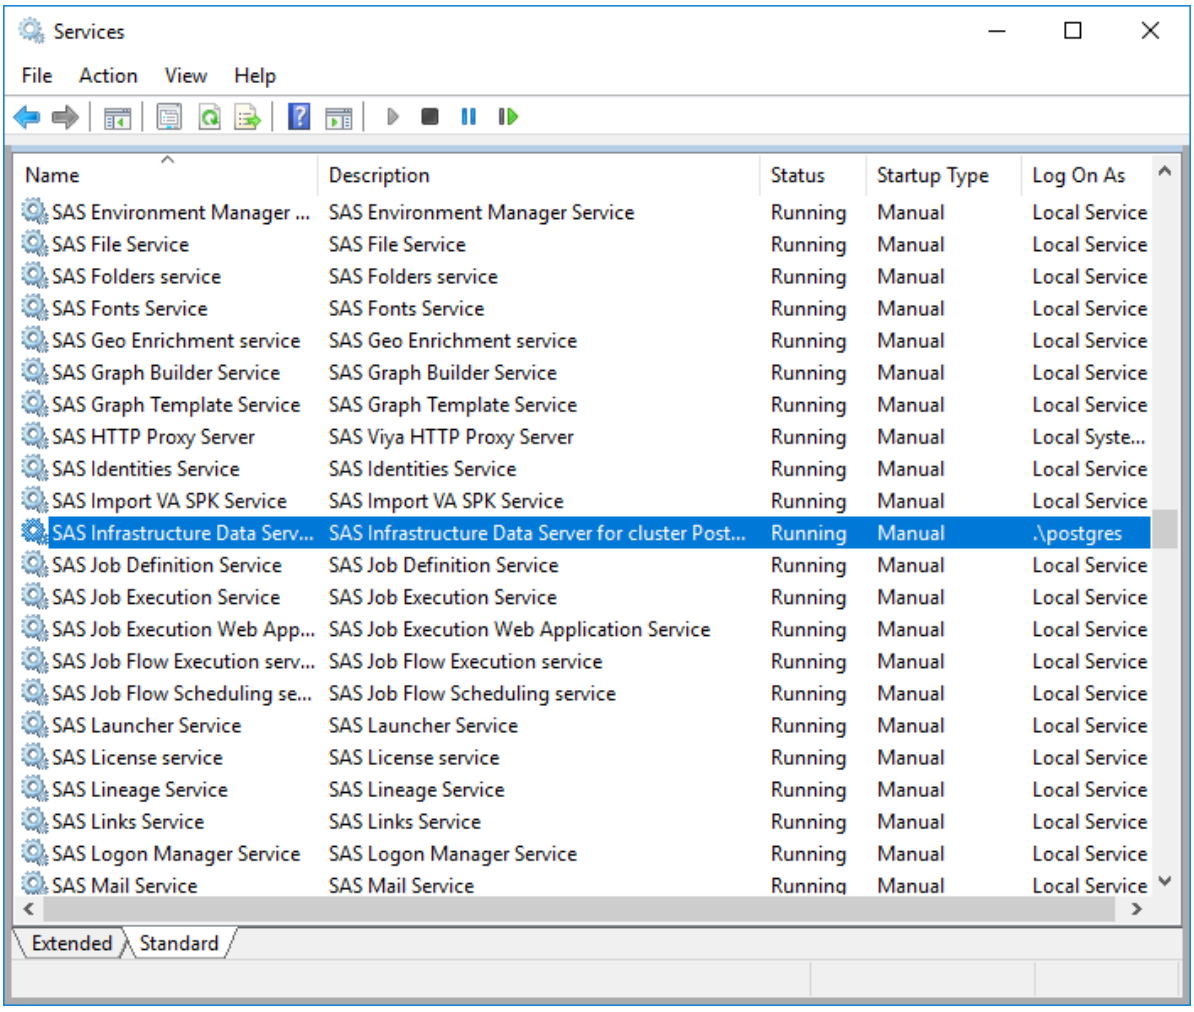

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other

SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.

Recover a Failed Single Node PostgreSQL Cluster

The SAS Infrastructure Data Server cluster is considered to be failed under these conditions: when it fails to start, and when all its data nodes are marked as `unhealthy` in the cluster definition file **`/opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/pool.cdf`**.

Common causes for cluster failure include a power failure, network connectivity issues, or a machine reboot. Another cause of cluster failure can be from lack of system resources. Examples are disk space, memory, number of processes, number of open files, ports, semaphores, and shared memory. In such cases, the data server logs contain failure information.

The cluster will not start until the problem that caused the cluster failure has been fixed, and the server nodes are marked as `healthy` in pool.cdf.

On the pgpool server machine, you can manually update pool.cdf, or you can run the repair_postgres_nodes.sh script. After the script updates pool.cdf, it attempts to start the cluster.

1. Before attempting any cluster repair procedures, do the following:
   - Examine the integrity of the data on the data server.

     One or more failovers might have occurred. Therefore, examine the server with the most current data.
   - Back up the data server data directories.
2. Fix the problem that caused the cluster to fail.
3. Verify the log files on both the Postgres and pgpool nodes for any errors.
   - Postgres log file:

     `/opt/sas/viya/config/var/log/sasdatasvrc/postgres/noden`, where *n* is the node number.
   - pgpool log file:

     `/opt/sas/viya/config/var/log/sasdatasvrc/postgres/pgpool0`
4. Stop the SAS Infrastructure Data Server cluster using the command that is appropriate for your operating system:
   - Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

     ```
     sudo systemctl stop sas-viya-sasdatasvrc-postgres
     ```
   - Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

     ```
     sudo service sas-viya-sasdatasvrc-postgres stop
     ```
5. Determine whether there are any leftover processes on one or more Postgres and pgpool nodes. The command lists only those leftover processes that belong to Postgres and pgpool. For command syntax details, see the Linux ps and egrep manual pages.

   ```
   ps -ef | egrep -i 'pgpool:|postgres:|all|bin/pgpool|bin/postgres'
   ```
6. Remove one or more leftover processes on the Postgres and pgpool nodes using the PID, as appropriate.

   ```
   kill PID
   ```

Note: Make sure that you remove only the pgpool parent process. When the parent process is removed, the child processes are removed by default.

Do not use the kill -9 command. This command removes only the parent process and leaves behind the child processes.

7 As the SAS install user (sas) or the user who has root-level privileges, sign in to the pgpool server machine.

8 Choose one of the following methods to update pool.cdf:

Note: Each method attempts to restart the cluster after pool.cdf has been modified.

- Run the repair_postgres_nodes.sh script to mark all cluster nodes as healthy in pool.cdf:

  sudo /opt/sas/viya/home/libexec/sasdatasvrc/script/maintenance/repair_postgres_nodes.sh

- Using a text editor, open the cluster definition file and mark all of the data server nodes as healthy:

  ```
  vi /opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/pool.cdf
  ```

  Change `node0=unhealthy` to `node0=healthy`.

9 Check the value of `network.databaseTraffic.enabled` is `true` by running the following commands:

Note: Specify each of these commands on a single line. Multiple lines are used here to improve readability

```
source /opt/sas/viya/config/consul.conf

export CONSUL_HTTP_TOKEN=$(sudo cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
tokens/consul/default/client.token)

/opt/sas/viya/home/bin/sas-bootstrap-config kv read --recurse config| grep network.databaseTraffic.enabled
```

10 If the value is true, verify that the SAS Configuration Server and SAS Secret Manager (vault) are alive and accessible:

```
/opt/sas/viya/home/bin/consul members
/opt/sas/viya/home/bin/vault status
```

Here is an example:

```
/opt/sas/viya/home/bin/consul members
```

```
node_id     | Address     | status | Type   |  Build | Protocol | DC    | Segment
------------+-------------+--------+--------+--------+----------+-------+----------
machine1 | ip-address  | alive  | server | 1.0.6  | 2        | viya | <all>
machine2 | ip-address  | alive  | client | 1.0.6  | 2        | viya | <default>
machine3 | ip-address  | alive  | client | 1.0.6  | 2        | viya | <default>
machine4 | ip-address  | alive  | client | 1.0.6  | 2        | viya | <default>
```

```
/opt/sas/viya/home/bin/vault status
```

```
Seal Type: <seal-type>
    Sealed: false
    Key Shares: 1
    Key Threshold: 1
    Unseal Progress: 0
    Unseal Nonce:
    Version: <version-number>
    Cluster Name: <vault-cluster-name>
    Cluster ID: <cluster-id>

    High-Availability Enabled: true
             Mode: active
            Leader Cluster Address: <cluster-address>
```

11 If the value is false, verify that at least the SAS Configuration Server is alive and accessible so that the cluster can be started.

```
/opt/sas/viya/home/bin/consul members
```

12 Start the cluster using the command that is appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  `sudo systemctl start sas-viya-sasdatasvrc-postgres`

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  `sudo service sas-viya-sasdatasvrc-postgres start`

13 Check the status of the data server cluster.

A **status** of `up` in the cluster status list indicates that the node is connected and is an active part of the cluster. There should be only one primary server, with no standby servers or one or more standby servers, as appropriate.

Failback a Cluster

Failback refers to the restoration of the high availability (HA) SAS Infrastructure Data Server cluster to its original configuration before the failover.

A PostgreSQL *original configuration* is indicated when the cluster status list displays the following:

- node0 has the **role** of `primary`
- all other nodes have a **role** of `standby`
- all nodes have a **status** of `up`

To failback a SAS Infrastructure Data Server cluster to its original configuration before the failover, follow these steps:

1 Make sure that the following servers are running and are accessible:

- SAS Configuration Server (Consul)
- Pgpool server
- SAS Infrastructure Data Server (PostgreSQL) cluster

2 As the SAS install user (sas) or the user who has root-level privileges, sign in to the pgpool server machine.

3 To ensure that the cluster is in a failover condition, check its status by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
cd /etc/init.d
sudo ./sas-viya-sasdatasvrc-service-name status
```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-service-name status
```

Verify that node0 has a **status** of `down` and a **role** of `standby`.

Here is an example of running the command on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution) where the *service-name* is named **postgres2**:

```
sudo service sas-viya-sasdatasvrc-postgres2 status
```

Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | down   | 0.250000  | standby | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | primary | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

4 In the cluster status list, if node0 has a **status** of `up` and a **role** of `standby`, you can go directly to Step 7.

5 To recover (start) node0, start the node using the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-sasdatasvrc-service-name-node0
```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-service-name-node0 start
```

Here is an example Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres2-node0 start
```

Here is typical output:

```
Starting sas-viya-sasdatasvrc-postgres2-node0 service...

                                                        [  OK  ]
```

6 Check the cluster status to verify that the node has started using the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
cd /etc/init.d
sudo ./sas-viya-sasdatasvrc-service-name status
```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-service-name status
```

Verify that node0 has a **status** of `up`.

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

sudo service sas-viya-sasdatasvrc-postgres2 status

Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | standby | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | primary | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

7 Now, stop the current primary node to failback to the original primary node by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl stop sas-viya-sasdatasvrc-*service-name-node-name*

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name-node-name* stop

Here is an example where *node-name* is named **node1** on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

sudo service sas-viya-sasdatasvrc-postgres2-node1 stop

Here is typical output:

```
Stopping sas-viya-sasdatasvrc-postgres2-node1 service...
                                                          [  OK  ]
```

8 Check the status of the cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  cd /etc/init.d

  sudo ./sas-viya-sasdatasvrc-*service-name* status

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name* status

Verify that **node0** has a **status** of `up` and a **role** of `primary`.

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution) where *service-name* is named **postgres2**:

sudo service sas-viya-sasdatasvrc-postgres2 status

Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | primary | 1          | false             | 0
 1       | machine2 | 5452 | down   | 0.250000  | standby | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

Note: The remaining running nodes of the cluster initially show a **status** of `down` while replication is established for the new primary node. Continue to monitor the cluster status until all running nodes have a **status** of `up`.

9 Now, recover (start) the stopped node by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl start sas-viya-sasdatasvrc-*service-name-node-name*

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name-node-name* start

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution) where the previous primary node is named **node1**:

sudo service sas-viya-sasdatasvrc-postgres2-node1 start

Here is typical output:

```
Starting sas-viya-sasdatasvrc-postgres2-node1 service...

                                                        [  OK  ]
```

10 Re-check the status of the cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  cd /etc/init.d

  sudo ./sas-viya-sasdatasvrc-*service-name* status

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name* status

Verify that the cluster has returned to its initial configuration:

- **node0** has the **role** of `primary`
- all other nodes have a **role** of `standby`
- all nodes have a **status** of `up`

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution) where the data server service is named **postgres2**:

sudo service sas-viya-sasdatasvrc-postgres2 status

Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | primary | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | standby | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

Add a Cluster (Ansible)

1 Sign on your Ansible controller with administrator privileges, and locate the file, **`/playbook/vars.yml`**.

2 Using a text editor, open vars.yml and locate the `INVOCATION_VARIABLES` section.

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

3 Copy and paste an existing cluster definition.

In this example, the new cluster is being added to Machine2:

```
INVOCATION_VARIABLES:
   Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres

   Machine2:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
```

```
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

4 Configure the new cluster definition for the pgpool server (pgpoolc) and the data server nodes (sasdatasvrc):

- Pgpool server definition parameters:
  - PCP_PORT

    Specifies the pcp port for the pgpool instance.
  - PGPOOL_PORT

    Specifies the pgpool port. This is the primary port through which all databases connect.
  - SANMOUNT

    Specifies the location of the data files. This path is typically the same value as the other data nodes.
  - SERVICE_NAME

    Specifies the unique service name for the data server cluster. `SERVICE_NAME` should be the same for the pgpool server and all nodes in the cluster.
- Data server node definition parameters:
  - NODE_NUMBER

    Specifies the sequential node identifier. The primary node is 0. Standby nodes start at 1 and are incremented sequentially.
  - NODE_TYPE

    Specifies the type of node that you are adding. The primary node should have a value of `P`. Standby nodes should have a value of `S`.
  - PG_PORT

    Specifies the Postgres database port. The pgpool server communicates with the database on this port. Clients use the `PGPOOL_PORT`. The port must be available for use on the deploy target.
  - SANMOUNT

    Specifies the location of the data files. This path is typically the same value as the other data nodes.
  - SERVICE_NAME

    Specifies the unique service name for the data server cluster. `SERVICE_NAME` should be the same for the pgpool server and all the nodes in the cluster.

Here is an example:

```
INVOCATION_VARIABLES:
  Machine2:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres2
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres2
    - NODE_NUMBER: '1'
```

```
    NODE_TYPE: S
    PG_PORT: '5432'
    SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
    SERVICE_NAME: postgres2
```

5 If the cluster machines that you are adding are already a part of your SAS Viya system, and are already in `[pgpoolc]` and `[sasdatasvrc]` host groups, then go to Step 6.

   Otherwise, add the machines to your Ansible inventory.ini file at the top of the file in the target list, and in the `[pgpoolc]` and `[sasdatasvrc]` host groups, respectively.

6 Run your Ansible playbook using the site.yml file.

   Here is an example:

   `ansible-playbook site.yml`

   For a complete list of playbook commands, see "Deploy the Software" in *SAS Viya for Linux: Deployment Guide*.

Delete a Node or a Cluster

CAUTION! Do not delete a primary node unless you plan to delete the entire cluster. Deleting a primary node would increase chances of introducing data corruption. To delete the primary node, failover the node to a standby node, and wait for all remaining nodes to indicate that they are available. When the nodes are available, it is safe to delete the former primary node. Do not delete a pgpool node without first moving the pgpool node of the cluster. Failure to do so will make the cluster unusable. If you choose to delete the node data using the `-d` option, its data files are deleted. Use caution when deciding to use the `-d` option.

1 As root or with an account that has root-level privileges, sign in to the machine where the node that you want to remove resides.

2 Change the directory to **`/opt/sas/viya/home/libexec/sasdatasvrc/script`**.

3 Run the sds_delete_node.sh script with the following options:

   Note: When the sds_delete_node.sh script runs, it stops the cluster.

   - -s *service-name*
   - -n *node-name*
   - -d y | n

   CAUTION! A yes (y) value specifies that the script deletes the node or the cluster data files.

   - -c *absolute-path*/sds_env_var.sh

   Here is an example:

```
sudo ./sds_delete_node.sh -s postgres -n node1 -d y
-c /opt/sas/viya/config/etc/sasdatasvrc/postgres/node1/sds_env_var.sh
```

   Every time the script runs, it generates a new log file in **`/tmp/sds_uninstall_log`**.

4 After the script runs, be sure to delete the node or the cluster definition in the `INVOCATION_VARIABLES` section of vars.yml. For more information, see "Add a Node (Ansible)" on page 28.

Recover a Cluster When All Nodes Are Stopped in a Multi-Node High Availability Cluster

Consider this example: in a three-node cluster where the primary is on node0, someone does the following:

1 Stops node0, and waits for fail-over to node1 to complete.

2 Stops node1, and waits for fail-over to node2 to complete.

3 Stop node2, and there are no more nodes to fail over to.

As a result of these actions, there is no primary node. The cluster is unable to be stopped or restarted.

To recover such a cluster, follow these steps:

1 As the root user or a user privileges, sign in to the SAS Configuration Server (Consul) machine.

2 Verify that the SAS Configuration Server and SAS Secret Manager (Vault) are running, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl status sas-viya-consul-default
  sudo systemctl status sas-viya-vault-default
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-consul-default status
  sudo service sas-viya-vault-default status
  ```

3 If you have made the SAS Viya service layer highly available, do the following:

- Verify that all the configuration servers are alive and accessible:

  Note: Specify each of these commands on a single line. Multiple lines are used here to improve readability

  ```
  source /opt/sas/viya/config/consul.conf

  /opt/sas/viya/home/bin/consul members

  export CONSUL_HTTP_TOKEN=$(sudo cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
  tokens/consul/default/client.token)

  /opt/sas/viya/home/bin/sas-bootstrap-config kv read config/postgres/sas.dataserver.conf/
  common/max_connections
  ```

- Verify that the secret manager leader and its hot standbys are alive and accessible:

  ```
  /opt/sas/viya/home/bin/vault status
  ```

Important: Bring up any nodes that are down, before you continue.

4 Log on to the primary SAS Infrastructure Data Server node machine as the SAS Installer user (sas), and set up the necessary environment variables and the Consul token:

```
source /opt/sas/viya/config/consul.conf

export CONSUL_HTTP_TOKEN=$(sudo cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
tokens/consul/default/client.token)
```

5 Verify that there is one (and only one) primary node:

```
/opt/sas/viya/home/bin/sas-bootstrap-config catalog service postgres
```

CAUTION! Failure to locate the correct primary node might cause data loss.

6 If two nodes are listed as primary, then locate the correct primary node. The node directory that contains the file, recovery.done, is the correct primary node. The size of the database for the primary node is usually

larger than the database for the other nodes. Check all the node directories (**/opt/sas/viya/config/data/sasdatasvrc/postgres/node0...*n***).

Here is an example of using the du (disk usage) command to find out the size of the database for node2:

```
du -sh /opt/sas/viya/config/data/sasdatasvrc/postgres/node2
```

7 After you have determined which node is the correct primary node, de-register the Postgres service for the correct primary node.

In this example, the correct primary node is node2:

```
/opt/sas/viya/home/bin/sas-bootstrap-config agent service deregister postgres-datanode2
```

8 Create a temporary JSON file to re-register node2 as a primary node using the tag **primary**:

a Start by making a copy of the file, service_node_registration.json.

In this example, the correct primary node is node2:

```
cd /opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/
```

```
cp /opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/service_node_registration.json
/opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/service_node_registration.json.primary.tmp
```

b Using a text editor, open the copied JSON file (service_node_registration.json.primary.tmp) and add **primary** to the tag list:

```
{
    "id": "postgres-datanode2",
    "name": "postgres",
    "tags": [
       "postgres", "primary"
    ],
    "address": "machine3.example.com",
    "port": 5432,
    "checks": [
      {
        "script": "/opt/sas/viya/home/libexec/sasdatasvrc/script/sds_consul_health_check.sh -config_path
              /opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/sds_env_var.sh",
        "interval": "30s",
        "timeout": "5s"
      }
    ]
  }
```

9 Re-register the correct node as a primary node.

In this example, the correct node is node2:

```
/opt/sas/viya/home/bin/sas-bootstrap-config agent service register --json
/opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/service_node_registration.json.primary.tmp
```

10 As you did in Step 5, check the PostgreSQL service for nodes and their tags:

```
/opt/sas/viya/home/bin/sas-bootstrap-config catalog service postgres
```

11 Repeat Step 8 – Step 10 to correct the wrong primary node, if any.

12 Stop the SAS Infrastructure Data Server cluster using the command appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl stop sas-viya-sasdatasvrc-postgres

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-postgres stop

13 Using a text editor, open **/opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/pool.cdf**, and specify the correct primary node as **healthy** and all standby nodes as **unhealthy**.

   In this example, node2 is the correct primary node:

```
cluster_definition_file_version=10.0.1
  pgpool_server_version=3.6.6
  postgresql_server_version=9.4.17
  pgpool0=healthy
  node0=unhealthy
  node1=unhealthy
  node2=healthy
  pgpoolConfigured
```

14 Using a text editor, open the PostgreSQL and Pgpool configuration files and verify that all parameters have proper values and that there are no missing values.

- Pgpool configuration file:

  /opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pgpool.conf

- PostgreSQL configuration files:

  /opt/sas/viya/config/data/sasdatasvrc/postgres/noden/postgresql.conf

> **TIP** Pgpool does not recognize a primary node if back-end parameters are missing in the pgpool.conf file: `backend_data_directoryn`, `backend_hostnamen`, `backend_portn`, `backend_weightn`, where *n* is the node number.

15 If you find any missing values in the configuration files that you verified in Step 14, then refresh those configuration files by stopping and restarting the corresponding consul-template services.

   Here is an example for node0 on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres-node0-ct-postgresql stop

sudo service sas-viya-sasdatasvrc-postgres-node0-ct-postgresql start

sudo service sas-viya-sasdatasvrc-postgres-node0-ct-pg_hba stop

sudo service sas-viya-sasdatasvrc-postgres-node0-ct-pg_hba start
```

   Here is an example for Pgpool on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pgpool stop

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pgpool start

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pool_hba stop

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pool_hba start

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pcp stop

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pcp start
```

16 After you have verified that the PostgreSQL and Pgpool configuration files are correct, start the cluster using the command appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl start sas-viya-sasdatasvrc-postgres

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

sudo service sas-viya-sasdatasvrc-postgres start

17 Check the status of the cluster using the command appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  cd /etc/init.d

  sudo ./sas-viya-sasdatasvrc-postgres status

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-postgres status

In this example, the two nodes are down:

```
PGPool is running with PID=21226
Checking Postgresql nodes status...
 node_id |       hostname       | port | status | lb_weight |  role   | select_cnt | load_balance_node |
replication_delay
---------+----------------------+------+--------+-----------+---------+------------+-------------------
+-------------------
 0       | machine1.example.com | 5432 | down   | 0.333333  | standby | 0          | false             | 0
 1       | machine2.example.com | 5432 | down   | 0.333333  | standby | 0          | false             | 0
 2       | machine3.example.com | 5432 | up     | 0.333333  | primary | 159        | true              | 0
(3 rows)
```

18 Recover the stopped nodes. Starting the node automatically recovers it and attaches it back to the cluster.

In this example, node0 and node1 are the stopped nodes that must be recovered. The commands used here are for Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres-node0 start

sudo service sas-viya-sasdatasvrc-postgres-node1 start
```

For Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-sasdatasvrc-postgres-node0

sudo systemctl start sas-viya-sasdatasvrc-postgres-node1
```

Note: If the database size is greater than 150 GB, use the script instead of the systemctl command to recover the node:

```
sudo /etc/init.d/sas-viya-sasdatasvrc-postgres-node0 start
```

If the database size is greater than 150 GB, using the systemctl command results in the following error message: **timeout was exceeded**. The systemctl command exits after the specified time has elapsed even though the recovery process is still running in the background. Therefore, the user will not know whether the recovery process has completed.

To resolve this issue, use the script **sas-viya-sasdatasvrc-postgres-node0**. The script runs until the recovery process has completed, and then it issues a message that indicates that the node was recovered successfully. Here is an example:

```
cd /etc/init.d
sudo ./sas-viya-sasdatasvrc-postgres-node0 start

Starting sas-viya-sasdatasvrc-postgres-node0 service...


Success
```

19 Re-check the status of the cluster.

(For the necessary command, see Step 17.)

In this example, the two nodes that were down are now up:

```
PGPool is running with PID=21226
Checking Postgresql nodes status...
 node_id |       hostname       | port | status | lb_weight |  role   | select_cnt | load_balance_node |
replication_delay
---------+----------------------+------+--------+-----------+---------+------------+-------------------
+-------------------
 0       | machine1.example.com | 5432 | up     | 0.333333  | standby | 0          | false             | 0
 1       | machine2.example.com | 5432 | up     | 0.333333  | standby | 0          | false             | 0
 2       | machine3.example.com | 5432 | up     | 0.333333  | primary | 159        | true              | 0
(3 rows)
```

How To (Nodes)

Check the Status of a Node (Linux)

SAS Viya uses the operating system's default init system or systemd command to launch a script that can check the status of a SAS Infrastructure Data Server node. This script, `sas-viya-sasdatasvrc-postgres-noden`, resides in **`/etc/init.d`**.

Note: As the SAS install user (sas) or the user who has root-level privileges, you must be signed in to the machine where the node resides.

Note: Each node script is numbered, starting at zero (0).

Run the script, using the command appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  `sudo systemctl status sas-viya-sasdatasvrc-postgres-noden`

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  `sudo service sas-viya-sasdatasvrc-postgres-noden status`

Note: On Red Hat Enterprise Linux 7.*x*, the first time you run the `sas-viya-sasdatasvrc-postgres-noden` command, you must run the command twice. Red Hat Enterprise Linux 7.*x* has backward compatibility for the system V service, and it does not have initial systemd unit files. The first time the `sas-viya-sasdatasvrc-postgres-noden` command runs, it builds the service unit file from the system V init file. Therefore, until the unit file is built, the `sas-viya-sasdatasvrc-postgres-noden` commands do not function properly.

Stop a Node (Linux)

SAS Viya uses the operating system's default init system or systemd command to launch a script that can stop a SAS Infrastructure Data Server node. This script, `sas-viya-sasdatasvrc-postgres-noden`, resides in **`/etc/init.d`**.

CAUTION! The act of stopping individual nodes changes the cluster state. Stopping the primary node causes a failover to occur in a cluster of two or more nodes. In addition, stopping a standby node removes the node from the active cluster. Stopped nodes must be recovered in order for them to added back to the cluster. The recovery of stopped nodes occurs automatically during node start-up. During failover of the primary node (0), other healthy standby nodes (2 and 3) go through a process of "following" the new primary node (1). During failover, nodes 2 and 3 briefly detach from the cluster and display a status of 3 (unhealthy). Wait several minutes and then recheck the cluster status. Eventually, nodes 2 and 3 should re-attach to the cluster and display a **status** of `up` (healthy).

Note: As the SAS install user (sas) or the user who has root-level privileges, you must be signed in to the machine where the node resides.

Note: Each node script is numbered, starting at zero (0).

Note: On Red Hat Enterprise Linux 7.*x*, the first time you run the `sas-viya-sasdatasvrc-postgres-noden` command, you must run the command twice. Red Hat Enterprise Linux 7.*x* has backward compatibility for the system V service, and it does not have initial systemd unit files. The first time the `sas-viya-sasdatasvrc-postgres-noden` command runs, it builds the service unit file from the system V init file. Therefore, until the unit file is built, the `sas-viya-sasdatasvrc-postgres-noden` commands do not function properly.

Run the script, using the command appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl stop sas-viya-sasdatasvrc-postgres-noden
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-postgres-noden stop
  ```

Start a Node (Recover a Node) (Linux)

A SAS Infrastructure Data Server data node is considered to be “unhealthy” when it has a **status** of `down` in the cluster status list. If a PostgreSQL data node has been stopped or has been taken offline, the pgpool server removes this node from the cluster.

When you restart an unhealthy node, pgpool server automatically initiates the node recovery process. To recover an unhealthy data node, follow these steps:

1. Make sure that the following servers are running and accessible:
   - SAS Configuration Server (Consul)
   - pgpool server
   - SAS Infrastructure Data Server (PostgreSQL) cluster
2. As the SAS install user (sas) or the user who has root-level privileges, sign in to the pgpool server machine, and run the following command, appropriate for your operating system:
   - Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

     ```
     cd /etc/init.d
     sudo ./sas-viya-sasdatasvrc-service-name status
     ```

   - Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

     ```
     sudo service sas-viya-sasdatasvrc-service-name status
     ```

   Verify that the unhealthy data node has a **status** of `down` and a **role** of `standby`.

   Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution), where the data server service is named **`postgres2`**:

   ```
   sudo service sas-viya-sasdatasvrc-postgres2 status
   ```

   Here is typical output:

   Note: In this example, the unhealthy node is **`node0`**.

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | down   | 0.250000  | standby | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | primary | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

3 As the SAS install user (sas) or the user who has root-level privileges, sign in to the machine that contains the unhealthy data node.

4 Make sure that the unhealthy node is stopped by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl stop sas-viya-sasdatasvrc-*service-name-node-name*

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name-node-name* stop

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution), where the unhealthy node is named **node0**.

sudo service sas-viya-sasdatasvrc-postgres2-node0 stop

Here is typical output:

```
Service sas-viya-sasdatasvrc-postgres2-node0 is not running.
                                                            [  OK  ]
```

5 Recover the node as a standby server by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl start sas-viya-sasdatasvrc-*service-name*-node0

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name*-node0 start

The pgpool server automatically starts the unhealthy node.

A node **status** of `up` indicates that the node is connected and is an active part of the cluster. There should be only one server with a **role** of `primary`, with zero or more servers with a **role** of `standby`.

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

sudo service sas-viya-sasdatasvrc-postgres2-node0 start

Here is typical output:

```
Starting sas-viya-sasdatasvrc-postgres2-node0 service...

                                                            [  OK  ]
```

6 Make sure that the node has been successfully added to the cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  cd /etc/init.d

  sudo ./sas-viya-sasdatasvrc-*service-name* status

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name* status

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

sudo service sas-viya-sasdatasvrc-postgres2 status

Here is typical output:

Note: In this example, the previously unhealthy node (**node0**) has a **status** of up.

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | standby | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | primary | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

Note: If starting (recovering) a node fails, refer to the Troubleshooting section on page 41.

Add a Node (Ansible)

Adding a data node to your SAS Infrastructure Data Server cluster consists of modifying the vars.yml file and running your Ansible playbook.

1. With administrator privileges, sign in to your Ansible controller , and locate the file, **/playbook/vars.yml**.

2. Using a text editor, open vars.yml and locate the INVOCATION_VARIABLES section.

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

3. Copy an existing node definition and place it under the deploy target on which the node will be configured.

   Here is an example:

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
```

```
        PGPOOL_PORT: '5431'
        SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
        SERVICE_NAME: postgres
      sasdatasvrc:
      - NODE_NUMBER: '0'
        NODE_TYPE: P
        PG_PORT: '5432'
        SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
        SERVICE_NAME: postgres
      - NODE_NUMBER: '0'
        NODE_TYPE: P
        PG_PORT: '5432'
        SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
        SERVICE_NAME: postgres
```

4 Configure the node definition in order to meet the requirements of the cluster:

- NODE_NUMBER

  Specifies the sequential node identifier. Standby nodes start at 1 and are incremented sequentially.

  For example, if you have only a primary node, the node that you are adding should have a `NODE_NUMBER` of 1. If the last standby node in your cluster has the value of 1, the node that you are adding should have a `NODE_NUMBER` of 2.

- NODE_TYPE

  Specifies the type of node that you are adding. The only acceptable value is `S` (standby). After initial deployment, you cannot add a primary node.

- PG_PORT

  Specifies the Postgres database port. The pgpool server communicates with the database on this port. Clients use the `PGPOOL_PORT`. The port must be available for use on the deploy target.

- SANMOUNT

  Specifies the location of the data files. This path is typically the same value as the other data nodes.

- SERVICE_NAME

  Specifies the service name for the data server cluster. It must be an exact match of the name of the cluster to which you are adding a data node.

Here is an example:

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    - NODE_NUMBER: '1'
      NODE_TYPE: S
      PG_PORT: '5433'
```

```
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

5 Run your Ansible playbook using the sitedefault.yml file.

Here is an example:

`ansible-playbook site.yml`

For a complete list of playbook commands, see "Deploy the Software" in *SAS Viya for Linux: Deployment Guide*.

Move a Node (Ansible)

Moving a data node to your SAS Infrastructure Data Server cluster consists of modifying the vars.yml file and running your Ansible playbook.

1 With administrator privileges, sign in to your Ansible controller , and locate the file **`/playbook/vars.yml`**.

2 Using a text editor, open vars.yml and locate the `INVOCATION_VARIABLES` section.

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    - NODE_NUMBER: '1'
      NODE_TYPE: S
      PG_PORT: '5433'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

3 Copy the existing node definition and place it under the deploy target to which you want to move the node.

In this example, the deploy target is Machine2:

```
INVOCATION_VARIABLES:
  Machine2:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres2
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres2
    - NODE_NUMBER: '1'
      NODE_TYPE: S
      PG_PORT: '5433'
```

```
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres2
```

4 Configure the node definition to meet the requirements of the cluster:

- NODE_NUMBER

  Specifies the sequential node identifier. This number should change to fit with the target cluster. For example, if the last standby node in the cluster is 1, the node that you are moving should have a `NODE_NUMBER` of 2. If there is only a primary node in the target cluster, the node that you are moving should have a `NODE_NUMBER` of 1.

- NODE_TYPE

  Specifies the type of node that you are moving. The only acceptable value is `s` (standby). After initial deployment, you cannot move a primary node.

- PG_PORT

  Specifies the Postgres database port. The pgpool server communicates with the database on this port. Clients use the `PGPOOL_PORT`. The port must be available for use on the deploy target.

- SANMOUNT

  Specifies the location of the data files. This path is typically the same value as other data nodes.

- SERVICE_NAME

  Specifies the service name for the data server cluster.

  Note: Do not change this value.

5 Run your Ansible playbook using the sitedefault.yml file.

Here is an example:

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

For a complete list of playbook commands, see "Deploy the Software" in *SAS Viya for Linux: Deployment Guide*.

Change the Port Number or the Data Directory for a Node (Ansible)

CAUTION! To avoid data corruption, do not change the port number or the data directory on a primary node.

1 As the SAS install user (sas) or the user who has root-level privileges, sign in to the pgpool machine.

2 To stop the node whose port or directory you want to change, run the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl stop sas-viya-sasdatasvrc-service-name-node-name
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name-node-name stop
  ```

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres-node3 stop
```

3 To check the status of the node, run the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  cd /etc/init.d
  sudo ./sas-viya-sasdatasvrc-service-name status
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name* status

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres status
```

Verify that failover has successfully occurred. In the cluster status list, the **status** of the node should be `down`.

4 With administrator privileges, sign in to your Ansible controller, and locate the file**/playbook/vars.yml**.

5 Using a text editor, open vars.yml and locate the `INVOCATION_VARIABLES` section.

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    - NODE_NUMBER: '1'
      NODE_TYPE: S
      PG_PORT: '5433'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

6 Make the necessary changes to the port number or the data directory in the definition for the node:

- NODE_NUMBER

  Specifies the sequential node identifier.

  Note: Do not change this value.

- NODE_TYPE

  Specifies the type of node: `P` (primary) or `S` (standby).

  Note: Do not change this value.

- PG_PORT

  Specifies the Postgres database port. The pgpool server communicates with the database on this port. Clients use the `PGPOOL_PORT`. The port must be available for use on the deploy target.

- SANMOUNT

  Specifies the location of the data files. This path is typically the same value as other data nodes.

- SERVICE_NAME

  Specifies the service name for the data server cluster.

  Note: Do not change this value.

7 Run your Ansible playbook using the sitedefault.yml file.

Here is an example:

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

For a complete list of playbook commands, see "Deploy the Software" in *SAS Viya for Linux: Deployment Guide*.

8 Check the status of the node. In the cluster status list, the **status** of the node should be `up`.

How To (General)

Get Current Passwords (Linux)

1 As the SAS install user (sas) or the user who has root-level privileges, sign in to any SAS Infrastructure Data Server machine.

2 Obtain the security token from the configuration server, and set it as an environment variable, using the appropriate command:

- Install user or root accounts:

```
source /opt/sas/viya/config/consul.conf
```

```
export CONSUL_HTTP_TOKEN=$(cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
tokens/consul/default/client.token)
```

- With root-level privileges (but not as the install user), install accounts:

```
source /opt/sas/viya/config/consul.conf
```

```
export CONSUL_HTTP_TOKEN=$(sudo cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
tokens/consul/default/client.token)
```

3 Run the sas-bootstrap-config script for the data server user ID whose password you want to obtain:

- sas

```
/opt/sas/viya/home/bin/sas-bootstrap-config kv read config/postgres/sas.dataserver.pooluser/
common/sr_check_password
```

- dbmsowner

```
/opt/sas/viya/home/bin/sas-bootstrap-config kv read config/application/sas/database/postgres/password
```

Change User Passwords (Linux)

The script, sds_change_user_pw.sh, changes SAS Infrastructure Data Server passwords and synchronizes them with SAS Configuration Server (Consul) and configuration files.

CAUTION! To avoid data loss, change the sas user account password only during a scheduled maintenance when users are not accessing SAS Viya. The data server must be running when you change the sas user's password Changing the password for the database user, sas, causes all nodes on the database cluster to restart.

Note: To change the password, you must know the current password. For more information, see "Get Current Passwords (Linux)".

1 As the SAS install user (sas), sign in to the SAS Infrastructure Data Server Pgpool machine.

Note: The change user password script requires sudo execution privileges.

2 You can determine the status of your cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

cd /etc/init.d

sudo ./sas-viya-sasdatasvrc-*service-name* status

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name* status

Before you run the change password script, verify that the cluster is in its initial configuration state (running and healthy):

- node0 has the **role** of `primary`
- all other nodes have a **role** of `standby`
- all nodes have a **status** of `up`

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution), where the data server service is named **postgres2**:

sudo service sas-viya-sasdatasvrc-postgres2 status

Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | primary | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | standby | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

3 Locate the data server environment variables file, sds_env_var.sh, and record its location.

  By default, sds_env_var.sh resides in **/opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0**.

4 The script prompts for the following information. Have this information ready when you run the script in a later step:

  - database user name
  - current database password
  - new database password

  Note: Your password must conform to the data server password policy on page 43.

5 Using the location of sds_env_var.sh noted in Step 3, run the script using the following command:

```
sudo -Hu sas /opt/sas/viya/home/libexec/sasdatasvrc/script/sds_change_user_pw.sh
-config_path
/opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/sds_env_var.sh
```

  TIP If you run the script from the directory where it resides, you might see several `cannot open [No such file or directory]` messages. This is a known issue, and you can safely ignore these messages.

6 Enter the information that you collected in Step 4 as the script prompts you for it.

After you provide the values in response to the prompts, the script connects to SAS Configuration Server and updates all instances of the database user password that it finds. Changes made in the configuration server are synchronized with the proper SAS Infrastructure Data Server configuration files. Finally, the script issues the necessary SQL commands in the data server to update the permissions for the database user.

7 To validate that your password has successfully changed, connect to the data server first database, **postgres**, using the PostgreSQL interactive terminal, psql:

```
/opt/sas/viya/home/bin/psql -h data-server-machine-name -U user-IDservice-name
```

8 When prompted, enter the new password for dbmsowner.

9 Restart all SAS Viya services.

For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.

Clean Up after a Hardware Failure (Linux)

If the machine on which the high availability (HA) SAS Infrastructure Data Server cluster runs was stopped unexpectedly, you might need to perform some cleanup steps after you restart the machine.

These steps involve removing any socket-lock files and any PID files that might have become orphaned after the PostgreSQL and pgpool servers were improperly shut down.

1 As the SAS install user (sas) or with root-level privileges, sign in to the pgpool machine.

2 Stop the HA data server cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl stop sas-viya-sasdatasvrc-service-name
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name stop
  ```

3 Delete any socket-lock file (in the form .s.PGSQL.**xxxx**) or any PID file (in the form *server*.pid) that corresponds to your HA data server cluster ports.

For the default HA data server instance with one data node, remove the following files:

- `/tmp/.s.PGSQL.5430`
- `/tmp/.s.PGSQL.5431`
- `/tmp/.s.PGSQL.5432`
- `/tmp/.s.PGSQL.5432.lck`
- `/opt/sas/viya/config/data/sasdatasvrc/postgres/node0/postmaster.pid`
- `/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/run/pgpool.pid`

4 Restart the HA data server cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl start sas-viya-sasdatasvrc-service-name
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name start
  ```

Remove a Persistent Lock on a Database Table (Linux)

Persistent locks on a SAS Infrastructure Data Server database table are caused by an uncommitted transaction or a long running query. To fix this problem, you must identify the process IDs of the client connections that are locking the table and terminate these connections.

1. As the SAS install user (sas) or with root-level privileges, sign in to the pgpool server machine.
2. If you know the PostgreSQL dbmsowner (superuser) password, go to Step 3. Otherwise, follow the steps in "Get Current Passwords (Linux)".
3. If you have not already done so, install pgAdmin on any machine (including Microsoft Windows) that has access to the machine that is running the pgpool server.
4. In pgAdmin, perform the following steps:
   a. Create a **New Server Registration** object and specify the following information:
      - **Host**: ***`machine-name`***

        The name of the machine on which the pgpool server resides.
      - **Port**: ***`pgpool-client-connection-port`***

        The default port is `5431`.
      - **Maintenance DB**: **`SharedServices`**

        Do not use the default, `postgres`.
      - **Username**: *superuser*

        The database superuser. The default is `dbmsowner`.
      - **Password**: string

        The superuser password.
   b. Connect to the pgpool server.
   c. Highlight the server name, and choose **Tools** ⇨ **Server Status**.

      The status panel shows all the client connections in the top panel. The second panel shows the persistent locks.
   d. Choose **Actions** ⇨ **Refresh** multiple times in order to determine whether the listed locks are transient or persistent. A transient lock disappears, and a persistent lock remains throughout refreshes.
   e. (Optional) You can cross-reference the process identifiers (PIDs) for the locked tables with the connection listing in order to identify the client (the application name) that has locked the table.

      Note: If you choose this option, open a SAS Technical Support track about this issue. Include the **PID**, **Application Name**, **Connection State**, and **Query** (if applicable) from the top connections section. Also include the PID and the persistent locked table names from the **Lock** section.
   f. To clear the locks, run the **pg_terminate_backend()** command on each PID that has a persistent lock.

      To do this, go back to the main pgAdmin panel. Highlight the **SharedServices** database name in the server **Object Browser** and choose **Tools** ⇨ **Query Tool** to open an SQL query execution window.
   g. Execute the **pg_terminate_backend(__PID__)** command to close each connection that is associated with a table that has a persistent lock.

      Here is an example:

```
SELECT pg_terminate_backend(14826);
SELECT pg_terminate_backend(16697);
SELECT pg_terminate_backend(22246);
```

h Select **Tools** ⇨ **Server Status** and refresh the panel (**Actions** ⇨ **Refresh** from the menu).

If all the persistent locks have been removed from the second **Locks** section, the persistent locks are successfully removed.

i Exit pgAdmin.

Routine Maintenance Tasks

Overview

Routine maintenance for a SAS Infrastructure Data Server consists of the following tasks:

- Adhering to a rigid schedule of performing database backups.
- Performing a re-index, vacuuming, and analyzing each table in the database during a maintenance cycle.
- Inspecting the data server logs periodically to make sure that there is no data corruption.
- Removing large orphaned data objects from the database periodically to free disk space.
- Track PostgreSQL software patches, and apply the patches that contain critical fixes, such as the CVE-2018-10915 security patch.

Re-Index, Vacuum, and Analyze Database Tables (Linux)

Follow these steps to re-index, vacuum, and analyze each table in the SAS Infrastructure Data Server databases. SAS recommends that you perform these steps during a maintenance cycle in order to reduce the chance of a PostgreSQL database command hanging because of a long-term lock on a table. If you encounter a hang condition, to remove the lock, you might need to restart the SAS Infrastructure Data Server.

1 As the SAS install user (sas) or with root-level privileges, sign in to the pgpool server machine.

2 If you know the PostgreSQL dbmsowner (superuser) password, go to Step 3. Otherwise, follow the steps in "Get Current Passwords (Linux)".

3 Run the following commands to set up the PostgreSQL command-line environment:

```
export PATH=/opt/sas/viya/home/bin:$PATH
```

```
export LD_LIBRARY_PATH=/opt/sas/viya/home/lib:/opt/sas/viya/home/
lib64:$LD_LIBRARY_PATH
```

4 (Optional) Stop all SAS Viya services, and run only SAS Infrastructure Data Server.

5 Run the following commands:

Note: For illustration, `5431` is used as the client connection port, **`/opt/sas/viya/home`** is used as the installation directory, and `dbmsowner` is used as the database superuser. Substitute the values that are appropriate for your site.

> **TIP** To prevent having to enter the superuser password multiple times, you can create a **`~/.pgpass`** file.

- Re-index all databases:

```
./reindexdb -a -p 5431 -h localhost -U dbmsowner
```

- Perform a full vacuum and analyze all databases:

```
./vacuumdb -p 5431 -h localhost -U dbmsowner -f -v -z -a
```

Note: If you encounter a hang condition, you might need to restart the SAS Infrastructure Data Server to remove the lock.

If there were no errors in the previous step, then you are done.

If you stopped the all of the SAS Viya services, then you can now stop the SAS Infrastructure Data Server and restart all the SAS Viya services.

Remove Large Orphaned Data Objects (Linux)

Large objects in the SAS Infrastructure Data Server database are stored separately from the tables that reference them. When a particular row is updated or deleted, these objects can become orphaned (unattached) from a table. Periodically, these orphaned large objects must be manually removed to free disk space.

1 Create an SQL command file named lo-cleanup.sql with the following content:

```
DROP FUNCTION IF EXISTS sas_lob_cleanup();
CREATE FUNCTION sas_lob_cleanup() RETURNS VOID AS $function$
DECLARE
   possible_lob_col record;
   possible_oid_row record;
   possible_oid_val bigint;
BEGIN
   DROP TABLE IF EXISTS sas_possible_lobs;

   CREATE TABLE sas_possible_lobs (table_schema TEXT NOT NULL, table_name
   TEXT NOT NULL, column_name TEXT NOT NULL, column_value BIGINT);

   FOR possible_lob_col IN SELECT * FROM information_schema.columns WHERE
   udt_name IN ('int4', 'int8', 'numeric', 'oid', 'text', 'varchar',
   'char', 'lo') AND NOT (table_schema = 'pg_catalog' AND table_name =
   'pg_shdepend') AND NOT (table_schema = 'pg_catalog' AND table_name =
   'pg_largeobject') AND table_name != 'sas_possible_lobs'

   LOOP

      BEGIN

          FOR possible_oid_val IN EXECUTE 'SELECT CAST(' ||
      possible_lob_col.column_name || ' AS BIGINT) FROM ' ||
      possible_lob_col.table_schema || '.' || possible_lob_col.table_name
      || ' WHERE ' || possible_lob_col.column_name || ' IS NOT NULL AND
      CAST(' || possible_lob_col.column_name || ' AS BIGINT) < ((2 ^ 32) -
      1) AND CAST(' || possible_lob_col.column_name || ' AS BIGINT) > 0'

          LOOP

          --raise notice 'successfully cast % % %',
          possible_lob_col.table_schema,
          possible_lob_col.table_name,
          possible_lob_col.column_name;

              INSERT INTO sas_possible_lobs (table_schema, table_name,
          column_name, column_value) VALUES (possible_lob_col.table_schema,
          possible_lob_col.table_name, possible_lob_col.column_name,
          possible_oid_val);
```

```
            END LOOP;

            EXCEPTION

                WHEN cannot_coerce THEN

                   --RAISE NOTICE 'error coercing % % %',
          possible_lob_col.table_schema, possible_lob_col.table_name,
          possible_lob_col.column_name;

                WHEN invalid_text_representation THEN

                   --RAISE NOTICE 'error casting % % %',
          possible_lob_col.table_schema, possible_lob_col.table_name,
          possible_lob_col.column_name;

                WHEN others THEN

                   RAISE NOTICE 'unexpected failure';

            END;

      END LOOP;

      SELECT LO_UNLINK(lo.loid) FROM pg_catalog.pg_largeobject lo GROUP BY loid
      HAVING (NOT EXISTS (SELECT 1 FROM public.sas_possible_lobs pl WHERE lo.loid
      = pl.column_value));

      DROP TABLE IF EXISTS sas_possible_lobs;

   END;

   $function$ LANGUAGE PLPGSQL;

   SELECT sas_lob_cleanup();
```

2 As the SAS install user (sas) or with root-level privileges, sign in to the pgpool server machine

3 Run the following command for each database:

Note: For illustration, `5431` is used as the client connection port, **`/opt/sas/viya/home`** is used as the installation directory, and `dbmsowner` is used as the database superuser. Substitute the values that are appropriate for your site.

```
psql -p 5431 -h localhost -U dbmsowner -d postgres -a -f lo-cleanup.sql
psql -p 5431 -h localhost -U dbmsowner -d SharedServices -a -f lo-cleanup.sql
```

Apply the CVE-2018-10915 Security Patch

A new security patch, CVE-2018-10915, fixes a vulnerability was found in libpq, the default PostgreSQL client library.

Sites that deploy SAS Viya 3.4 on Windows have a newer version of PostgreSQL (version 9.4.19) that contains the fix for this security issue. However, sites running SAS Viya 3.4 or earlier on Red Hat Enterprise Linux and SUSE Linux Enterprise Server, must manually apply patch SAS Infrastructure Data Server with patch CVE-2018-10915.

1 As the SAS install user (sas) or with root-level privileges, sign in to the pgpool server machine.

2 Stop all SAS Viya services.

3 Perform an update using `yum`, or run Ansible with the appropriate playbook.

4 Restart all the SAS Viya services.

Concepts

What is the SAS Infrastructure Data Server?

SAS Infrastructure Data Server is used for transactional storage by SAS middle-tier software. It is also used by some SAS solutions software for user content such as reports, custom groups, comments, authorization rules, selected source definitions, attachments, and user preferences. The server is configured specifically to support SAS software, and is based on PostgreSQL version 9.

By default, the SAS installer account is used to start the server.

The databases that are managed by the server are backed up and restored with the Backup and Recovery Deployment Tool. For more information, see *SAS Viya Administration: Backup and Restore*.

Figure 6 *SAS Infrastructure Data Server Architecture*

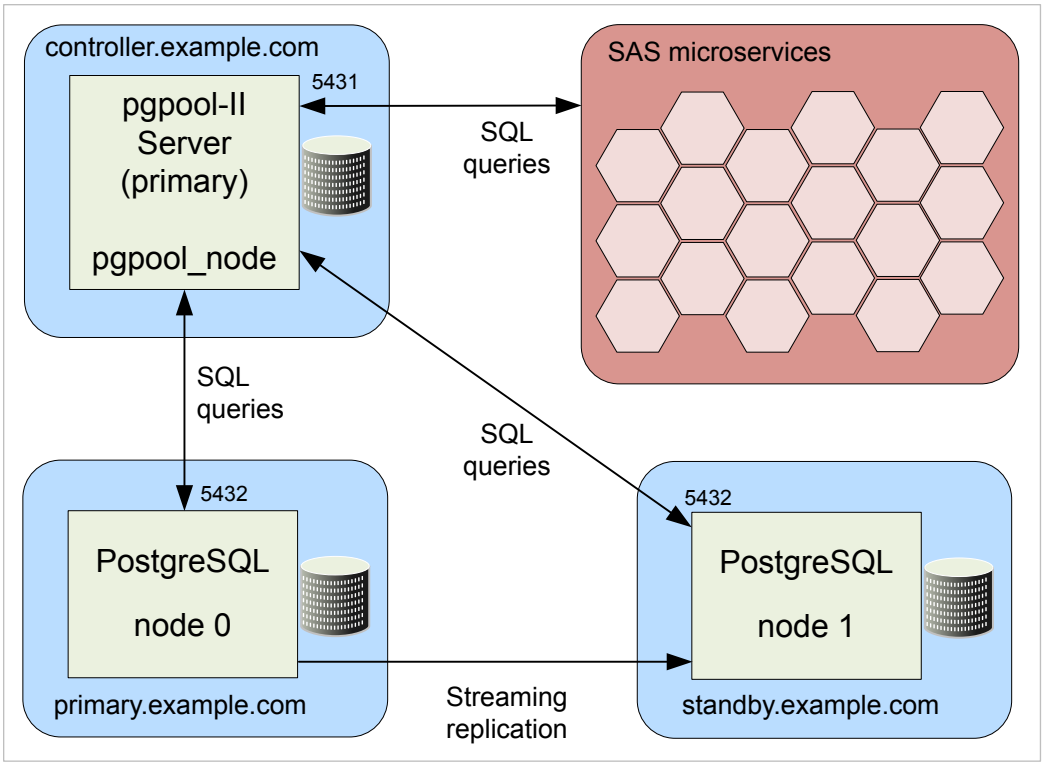

Pgpool-II

SAS provides Pgpool-II (version 3) open-source software to enable you to manage PostgreSQL clusters for high availability (failover management). The Pgpool-II software resides and operates between SAS Infrastructure Data servers and clients. All data connections and database requests are routed through the pgpool service.

Synchronizing Configuration Files with SAS Configuration Server

The SAS Infrastructure Data Server uses a daemon called consul-template to keep configuration files synchronized with configuration properties that are modified in SAS Environment Manager and stored as key-value pairs in SAS Configuration Server (Consul). For more information, see "Non-Spring-Based Servers" in *SAS Viya Administration: Configuration Properties*.

Each configuration file for each Pgpool node and PostgreSQL data node of the cluster has an init script to manage the synchronization. The consul-template init scripts for the SAS Infrastructure Data Server cluster are the following:

- For each Pgpool
  - sas-viya-sasdatasvrc-postgres-pool0-ct-pgpool
  - sas-viya-sasdatasvrc-postgres-pool0-ct-pool_hba
  - sas-viya-sasdatasvrc-postgres-pool0-ct-pcp
- For each data server node, where *n* is the node number, such as node0:
  - sas-viya-sasdatasvrc-postgres-node*n*-pg_hba
  - sas-viya-sasdatasvrc-postgres-node*n*-postgresql

Troubleshooting

psql: server closed the connection unexpectedly. This probably means the server terminated abnormally before or while processing the request.

Explanation:

The SAS Viya environment was shut down abnormally.

Resolution:

Restart the SAS Viya environment using the sas-viya-all-services start command. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.

/opt/sas/viya/config/etc/sasdatasvrc/../node.cdf was already marked with 'recoveryInProgress=y'. Exiting from auto-recovery.

Explanation:

PostgreSQL was in the process of recovering the SAS Infrastructure Data Server node when it encountered an error, and stopped the recovery process. Whenever it restarts a data server node, PostgreSQL always inserts the line, `recoveryInProgress=y`, in the node.cdf file to avoid a simultaneous recovery.

Resolution:

1. Review the recovery log to determine what the problem is.

   (The recovery log is located here: **`/opt/sas/viya/config/var/log/sasdatasvrc/cluster/nodex/sds_auto_recovery_node.log`**.)

2. Fix the problem.

3. Remove the following line from the node's node.cdf file:

   `recoveryInProgress=y`

4. Restart (recover) the node.

EDTERROR: missing chunk number 0 for toast value 9558737 in pg_toast_2619
EDTCONTEXT: automatic analyze of table "SharedServices.public.sas_audit"
EDTERROR: could not read block 3062 in file "base/18797/19703": read only 0 of 8192 bytes
***yyyy-mm-dd* EDTERROR: unexpected data beyond EOF in block 0 of relation base/16715/107679**

Explanation:

There is a high probability that your SAS Infrastructure Data Server database is corrupted.

Resolution:

After you correct the cause of the data corruption, and recover the database using a restored backup.

ERROR: Cluster stop failed. Please review the log file. /opt/sas/viya/config/var/log/sasdatasvrc/postgres/pgpool0/sas-viya-sasdatasvrc-postgres-service_YYYYMMDD_####.log [FAILED]
Unexpected response code: 500 ERROR: Unable to read a key Unexpected response code: 500 (rpc error: failed to get conn: dial tcp someotherhost.com:8300: getsockopt: connection refused) ERROR: Unable to list the nodes that provide the service 'postgres'

Explanation:

SAS Infrastructure Data Server fails to stop because it cannot connect to SAS Configuration Server (Consul) even though Consul is running.

This problem occurs in a multi-machine deployment where the primary data server node is not on the same system as the primary Consul server (the server designated in the inventory.ini file). If the primary Consul server has already been shut down, the data server fails to stop, even if a local Consul service is running.

Resolution:

Always stop SAS Infrastructure Data Server before the primary Consul server, regardless of which machine it is located on. For more information, see "Read This First: Start and Stop Servers and Services" in *SAS Viya Administration: General Servers and Services*.

Reference

Database

> **TIP** All PostgreSQL data servers have a *first database* named **postgres**. For more information, see Creating a Database in PostgreSQL documentation.

In a SAS Viya deployment, SAS Infrastructure Data Server is configured to manage the SharedServices database. SAS Viya microservices create database schemas within SharedServices.

If your deployment includes SAS solutions software that supports SAS Infrastructure Data Server, more databases might be configured on the server.

Default Users

dbmsowner
: The PostgreSQL database owner and the SAS database administrator user.

sas
: The SAS Viya install user and the account used for SAS Infrastructure Data Server cluster management.

Network Access

SAS Infrastructure Data Server is configured to accept connections on all network interfaces, and it requires password authentication. By default, SAS configures the server to use network port number 5431.

PostgreSQL instances are configured with JDBC data sources that reference the SharedServices database.

Password Policy

The user name and password for the SAS Infrastructure Data Server administrator are specified during deployment. The password can be updated. Passwords for SAS Infrastructure Data Server are subject to the following guidelines:

- The password must not contain any non-alphanumeric characters.

  Examples are underscores (_), hyphens (-), and periods (.).
- The password must be at least six characters long.
- The password can contain alphanumeric characters.
- There are no restrictions for including numbers or mixed-case characters.

Environment Parameters (Linux)

Export the following path in order to execute PostgreSQL and pgpool commands:

```
export LD_LIBRARY_PATH=/opt/sas/viya/home/lib:/opt/sas/viya/home/lib64
```

Configuration Files (Linux)

- `/opt/sas/viya/config/etc/sasdatasvrc/postgres/node0/node.cdf`
- `/opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/pool.cdf`
- `/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pgpool.conf`
- `/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pcp.conf`
- `/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pool_hba.conf`
- `/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pool_passwd`
- `/opt/sas/viya/config/data/sasdatasvrc/postgres/node0/postgresql.conf`
- `/opt/sas/viya/config/data/sasdatasvrc/postgres/node0/pg_hba.conf`

Configuration File (Windows)

The SAS Infrastructure Data Server configuration file on Windows is `\ProgramData\SAS\Viya\etc\sasdatasvrc\postgres\node0\node.cdf`.

Log Files

SAS Infrastructure Data Server log files are located in the following path, depending on your operating system:

- Linux:

  `/opt/sas/viya/config/var/log/sasdatasvrc`
- Windows:

  `\ProgramData\SAS\Viya\var\log\sasdatasvrc\postgres\node0`

SAS Message Broker

Overview

SAS uses a set of event APIs that are dependent on Spring Integration and Spring AMQP to interact with the message broker. The AMQP-compliant message broker that SAS uses is Pivotal's RabbitMQ, version 3. RabbitMQ includes the Erlang platform, version 20.

Note: A programming-only deployment does not use SAS Message Broker.

Operate (Linux)

SAS Viya provides a script in **/etc/init.d** that you use to stop, start, restart, and check the status of SAS Message Broker. The script is named, `sas-viya-rabbitmq-server-default`.

Syntax

How you run `sas-viya-rabbitmq-server-default` depends on your operating system:

- Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

  sudo systemctl status | stop | start | restart sas-viya-rabbitmq-server-default

- Red Hat Enterprise Linux 6.x (or an equivalent distribution):

  sudo service sas-viya-rabbitmq-server-default status | stop | start | restart

Usage Notes and Tips

- You must be logged on to the machine where SAS Message Broker resides. Also, you must have root-level privileges to run this script.
- If there are multiple instances of SAS Message Broker:
  - start them in the reverse sequence in which you stopped them.
  - stop them in the reverse sequence in which you started them.
- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.
- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-rabbitmq-server-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-rabbitmq-server-default` on RHEL 7.x with the `service` command, and later attempt to shut down SAS Message Broker using the `systemctl` command, the configuration server stops responding and does not shut down.

Examples

- To check status of SAS Message Broker on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status sas-viya-rabbitmq-server-default
```

- To stop SAS Message Broker on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-rabbitmq-server-default stop
```

- To start SAS Message Broker on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl start sas-viya-rabbitmq-server-default
```

- To restart SAS Message Broker on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-rabbitmq-server-default restart
```

Operate (Windows)

Using the Microsoft Management Console (MMC) Services snap-in, you can start, stop, and restart SAS Message Broker.

Figure 7 *SAS Message Broker in the Services Snap-In*

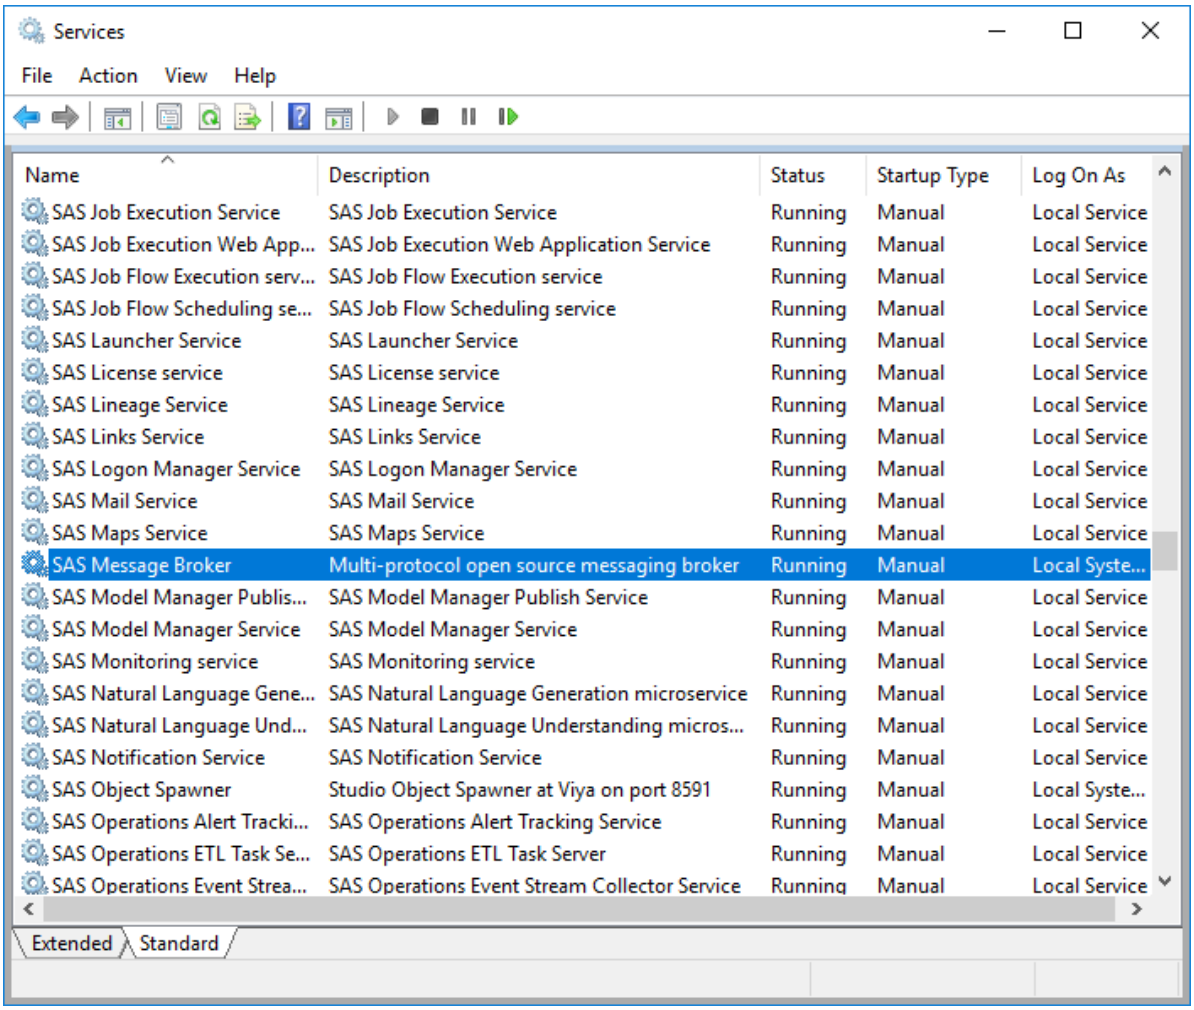

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other

SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.

Concepts

What is SAS Message Broker?

SAS Message Broker is an integral part of the event-driven architecture in which SAS Viya services participate. SAS uses a set of event APIs that are dependent on Spring Integration and Spring AMQP for interacting with the message broker. The AMQP-compliant message broker that SAS uses is Pivotal's RabbitMQ. The SAS event APIs provide a layer of abstraction between the message broker and its clients. The SAS event APIs also prevent code from breaking, which could result if SAS changed its third-party message broker from RabbitMQ to another third-party message broker in the future.

How Does Message Broker Work?

SAS Message Broker accepts messages in a standard format and routes them through exchanges and queues, which provide transaction acknowledgment, message persistence, and redundancy. Message broker exchanges accept messages from publishers and route them to queues, as appropriate. The exchange type controls whether messages are sent to a specific queue, to all associated queues, or only to queues that accept a particular message routing key or that match a key pattern.

SAS Message Broker Reference

Exchanges

SAS Message Broker uses the following exchanges:

- sas.application
- sas.application.backup
- sas.backup.topic
- sas.ledger
- sas.log
- sas.metric
- sas.notification
- sas.search.schema.topic

Configuration Files

SAS Message Broker configuration files are located in `/opt/sas/viya/config/etc/rabbitmq-server/`.

Note: Change these configuration files only when instructed to do so by SAS Technical Support.

Log Files

SAS Message Broker log files are located in **/opt/sas/viya/config/var/log/rabbitmq-server/default**.

SAS Cache Locator and Cache Server

Overview

SAS Cache Locator and SAS Cache Server provide a distributed cache technology to microservices in SAS Viya.

See Also

SAS Viya Administration: Tuning

Operate (Linux)

SAS Viya provides scripts in **/etc/init.d** that you use to stop, start, restart, and check the status of SAS Cache Locator and SAS Cache Server. The scripts are named `sas-viya-cachelocator-default` and `sas-viya-cacheserver-default`, respectively.

Syntax

How you run `sas-viya-cachelocator-default` and `sas-viya-cacheserver-default` depends on your operating system:

- Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

  sudo systemctl status | stop | start | restart sas-viya-cachelocator-default

  sudo systemctl status | stop | start | restart sas-viya-cacheserver-default

- Red Hat Enterprise Linux 6.x (or an equivalent distribution):

  sudo service sas-viya-cachelocator-default status | stop | start | restart

  sudo service sas-viya-cacheserver-default status | stop | start | restart

Usage Notes and Tips

- You must be logged on to the machine where the SAS Cache Locator and SAS Cache Server services reside. Also, you must have root-level privileges to run these scripts.
- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.
- On Linux systems that support systemd, use the `systemctl` command when running the individual service and server scripts. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-cachelocator-default` on RHEL 7.x with the `service` command, and later attempt to shut down SAS Cache Locator using the `systemctl` command, the cache locator stops responding and does not shut down.

Examples

- To check status of SAS Cache Locator on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status sas-viya-cachelocator-default
```

- To stop SAS Cache Server on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

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

- To start SAS Cache Locator on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl start sas-viya-cachelocator-default
```

- To restart CAS Cache Server on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-cacheserver-default restart
```

Operate (Windows)

Using the Microsoft Management Console (MCC) Services snap-in, you can start, stop, and restart SAS Cache Locator and the SAS Cache Server.

Figure 8 *SAS Cache Locator in the Services Snap-In*

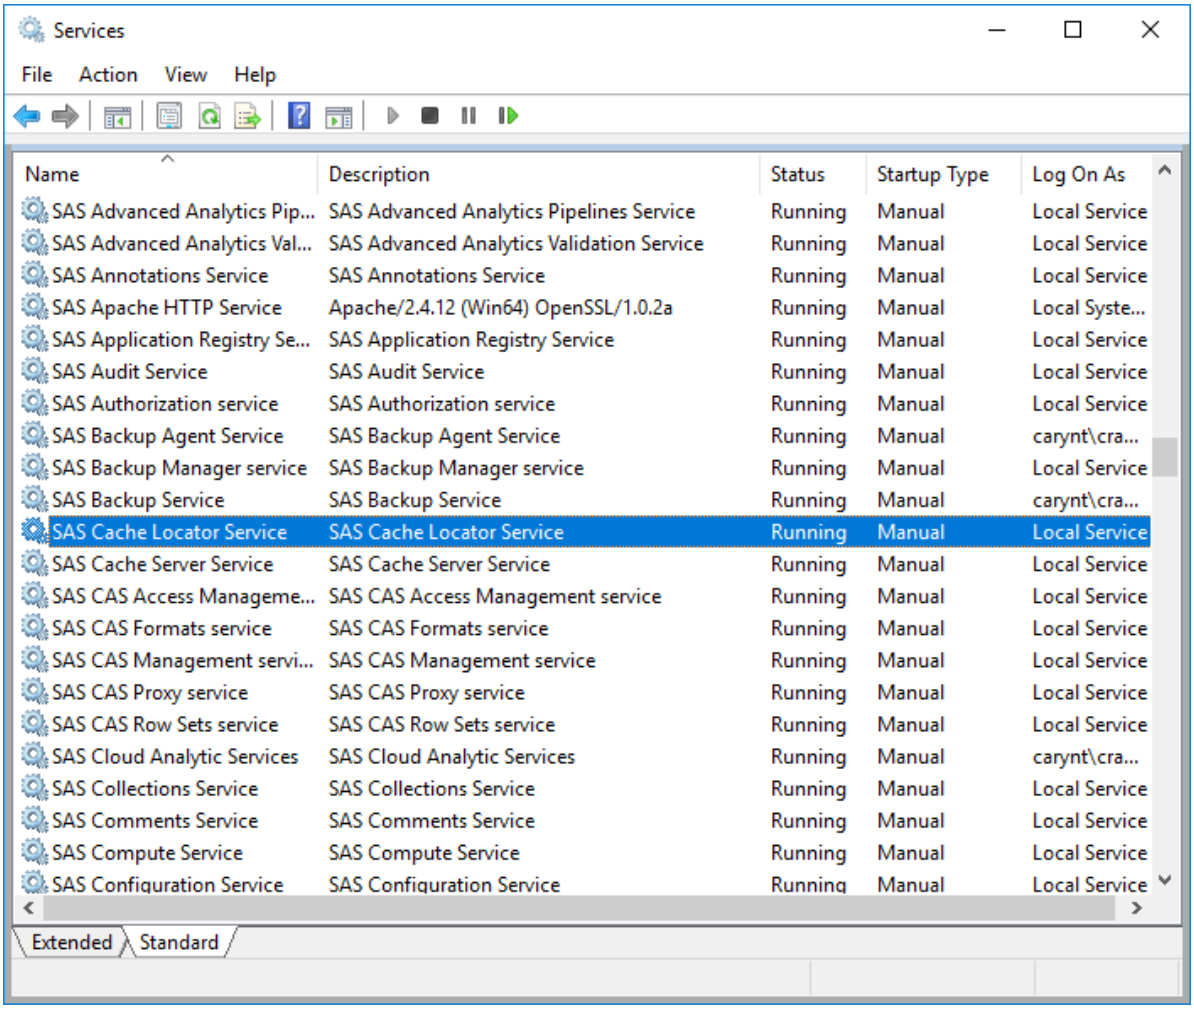

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.

Concepts

SAS Cache Locator

SAS Cache Locator is a server that provides discovery information to SAS Viya microservices for the purpose of forming a distributed data cache. SAS Cache Locator is based on the open-source Apache Geode project.

SAS Cache Server

SAS Cache Server hosts long-lived data regions (a cache) and serves the contents to SAS Viya microservices. Like SAS Cache Locator, SAS Cache Server is based on the open-source Apache Geode project.

Configuration

SAS Cache Locator and SAS Cache Server embed the Apache Geode API within their respective SAS Viya microservices, cachelocator and cacheserver.

The cachelocator and cacheserver microservices enable the cache locator and the cache server to gain access to SAS Configuration Server (Consul) in order to dynamically register and to retrieve properties with the SAS Viya Configuration service. For more information, see "Non-Spring-Based Servers" in *SAS Viya Administration: Configuration Properties*.

When configuration changes are made to cachelocator and cacheserver, you must restart SAS Cache Locator and SAS Cache Server in order for their changes to take effect. For information about how to modify the configuration for cachelocator and cacheserver, see "Edit Configuration Instances" in *SAS Viya Administration: Configuration Properties*.

Log Files

Log files for SAS Cache Locator and SAS Cache Server are located in `/opt/sas/viya/config/var/log/cachelocator/default` and `/opt/sas/viya/config/var/log/cacheserver/default`.

Apache HTTP Server

Overview

SAS Viya uses Apache HTTP Server to serve static HTML content and to proxy client connections. A high-availability proxy environment is not installed by default, but is a supported configuration.

For information about software requirements, see "Apache httpd" in *SAS Viya for Linux: Deployment Guide*.

How To

Operate (Linux)

Note: You must be logged on to the machine where SAS HTTP Proxy Server resides, and you must have root-level privileges to run this script.

Note: For complete information about httpd arguments, see https://httpd.apache.org/docs/2.0/programs/httpd.html.

To run httpd and sas-viya-httpproxy-default scripts, use the commands that are appropriate to your operating system:

- On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution):

  ```
  sudo systemctl status | stop | start | restart httpd
  sudo systemctl status | stop | start | restart sas-viya-httpproxy-default
  ```

- On SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl status | stop | start | restart apache2
  sudo systemctl status | stop | start | restart sas-viya-httpproxy-default
  ```

- On Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service httpd status | stop | start | restart
  sudo service sas-viya-httpproxy-default status | stop | start | restart
  ```

Operate (Windows)

Using the Microsoft Management Console (MMC) Services snap-in, you can start, stop, and restart SAS HTTP Proxy Server.

Figure 9 *SAS Apache HTTP Server in the Services Snap-In*

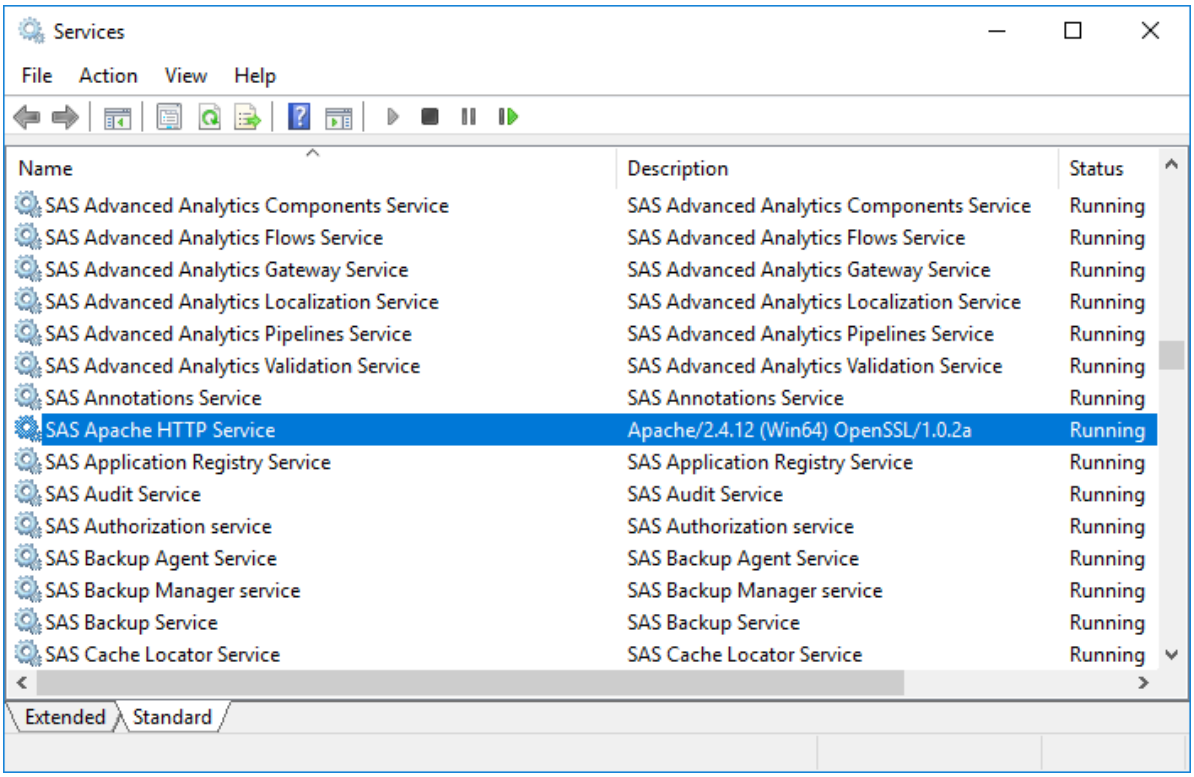

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services" in *SAS Viya Administration: General Servers and Services*.

Configure External Reverse Proxy

To configure SAS Viya for an external reverse proxy, including load balancers that act as a reverse proxy:

1 Add properties to ensure that applications that generate links to SAS Viya objects (such as SAS Visual Analytics reports) are aware of the external reverse proxy.

   From the machine where the SAS Viya internal Apache HTTP Server resides, run the following commands to add the appropriate property values. Here is an example:

   Note: Specify each of these commands on a single line. Multiple lines are used here to improve readability

```
source /opt/sas/viya/config/consul.conf

/opt/sas/viya/home/bin/sas-bootstrap-config kv write
config/viya/sas.httpproxy.external.hostname reverseproxy_hostname

/opt/sas/viya/home/bin/sas-bootstrap-config kv write
```

```
config/viya/sas.httpproxy.external.port reverseproxy_port
```

2 On the external reverse proxy, set the X-Forwarded-Proto and X-Forwarded-Port headers to the protocol and port that the client is using to connect to the external reverse proxy.

3 From the SAS Viya internal Apache HTTP server machine, comment out the X-Forwarded-Proto and X-Forwarded-Port lines in the petrichor.conf file as follows:

   a Locate the petrichor.conf file according to the platform used:

UNIX	**/etc/http/conf.d**
SUSE Linux	**/etc/apache2/conf.d**
Windows	**C:\ProgramData\SAS\Viya\etc\httpd\conf.d**

   b Make a backup copy of the petrichor.conf file.

   c Locate the following lines in the petrichor.conf file:

```
# Default the X-Forwarded-Proto header to http
RequestHeader set X-Forwarded-Proto http

# Set the X-Forwarded-Proto header if request is over HTTPS
RewriteCond "%{HTTPS}" "on"
RewriteRule ^.*$ - [ENV=HTTPS:true]
RequestHeader set X-Forwarded-Proto https env=HTTPS
RequestHeader set X-Forwarded-Port 443 env=HTTPS
```

   d Comment out these lines so that they appear as follows:

```
# Default the X-Forwarded-Proto header to http
#RequestHeader set X-Forwarded-Proto http

# Set the X-Forwarded-Proto header if request is over HTTPS
#RewriteCond "%{HTTPS}" "on"
#RewriteRule ^.*$ - [ENV=HTTPS:true]
#RequestHeader set X-Forwarded-Proto https env=HTTPS
#RequestHeader set X-Forwarded-Port 443 env=HTTPS
```

4 Start (or restart) the SAS Viya internal Apache HTTP Server.

Note:

- The external reverse proxy must use HTTPS.
- The external reverse proxy must forward requests through the SAS Viya internal Apache HTTP Server without changing the URL path.
- In environments with multiple SAS Viya internal Apache HTTP servers, you should configure the external reverse proxy to route requests to an active HTTPD instance. Round-robin routing or load-balanced routing is recommended.

HTTP Request Methods

SAS Viya supports the following HTTP request methods. These methods are used by REST APIs for creating, reading, updating, and deleting resources and for performing other domain-specific actions.

- DELETE
- GET
- HEAD
- OPTIONS
- PATCH
- POST
- PUT

For more details, refer to HTTP Method Registry.

Concepts

SAS Viya uses Apache HTTP Server as a web server. Apache HTTP Server serves static HTML content and proxies client communication.

A third-party load balancer is required in order to provide high availability for Apache HTTP Server. You can also install your own web server on a separate machine in order to proxy connections from the internet to Apache HTTP Server. For more information about making HTTP Server highly available, see "Apache httpd" in *SAS Viya for Linux: Deployment Guide*.

Log Files

Note: You must be logged on with root-level privileges to the machine where the service resides in order to view log files.

By default, Apache HTTP Server log files are located in the following directory, as appropriate for your operating system:

- On Red Hat Enterprise Linux and equivalent distributions:

  `/var/log/httpd`

- On SUSE Linux Enterprise Server 12.*x*:

  `/var/log/apache2`