



SAS[®] Viya[™] 3.2 Administration: SAS[®] Infrastructure Data Server

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.

SAS Infrastructure Data Server: How To

Operate a Cluster

SAS Viya uses the operating system's default init system or systemd command to launch a script that can stop, start, restart, and check the status of the SAS Infrastructure Data Server cluster. (A data server cluster consists of all of the PostgreSQL nodes and pgpool-II.) This script, `sas-viya-sasdatasvrc-postgres`, resides in `/etc/init.d`.

Note: You must be signed in to the pgpool-II machine as the SAS install user (`sas`) or with `sudo` privileges to run this script.

To operate a data server cluster, run the following command:

```
sas-viya-sasdatasvrc-postgres status | stop | start | restart
```

Note: There is a script with which you can manage and view the running state of all SAS Viya servers and services. For more information, see [“All Servers and Services”](#) in [SAS Viya Administration: General Servers and Services](#).

For your convenience, here are a few examples:

- checking status of the data server cluster using a direct call:

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

- stopping the data server cluster using the **Red Hat Linux version 6** init system command:

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

2

- starting the data server cluster using the **Red Hat Linux version 7** systemd command:

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

- restarting the data server cluster using a direct call:

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

Check the Status of a Node

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: You must be signed in to the node machine as the SAS install user (`sas`) or with `sudo` privileges.

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

- Red Hat Linux version 6:

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

- Red Hat Linux version 7:

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

Stop a Node

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! 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 it from the active cluster. Stopped nodes must be recovered in order for them to be added back to the cluster. (This is done 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). When that happens, wait for a couple of minutes and then recheck the cluster status. Eventually, nodes 2 and 3 should re-attach to the cluster and display a status of 2 (healthy).

Note: You must be signed in to the node machine as the SAS install user (`sas`) or with `sudo` privileges.

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

- Red Hat Linux version 6:

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

- Red Hat Linux version 7:

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

Start a Node (Recover a Node)

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

When you start a node, PostgreSQL automatically goes through the node recovery process.

Note:

SAS Configuration Server, which is based on Consul, must be running before you can recover a stopped SAS Infrastructure Data Server node.

1 Sign in to the ppgool-II machine as the SAS install user (sas) or with sudo privileges.

2 Run the following command:

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

Verify that the node to be started (recovered) is a **standby** server and is disconnected from the cluster (status of 3).

```
node_id | hostname | port | status | lb_weight | role
-----+-----+-----+-----+-----+-----
0       | machine1 | 5432 | 2      | 0.250000 | primary
1       | machine2 | 5432 | 3      | 0.250000 | standby
(2 rows)
```

3 Sign in to the stopped node machine as the SAS install user (sas) or with sudo privileges.

4 Make sure that the node is indeed stopped, by running the appropriate command:

Note: In this document, the stopped node is node 1.

■ Red Hat Linux version 6:

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

■ Red Hat Linux version 7:

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

5 Start (recover) the node as a standby server, by running the appropriate command:

■ Red Hat Linux version 6:

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

■ Red Hat Linux version 7:

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

Note: When you start a node, PostgreSQL automatically starts the node as a standby server. If you want to make a standby node a primary server, see [Failback an HA Cluster](#).

6 Sign in to the ppgool-II machine as the SAS install user (sas) or with sudo privileges.

7 Verify that the node has been successfully added to the cluster, by running the node status command:

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

A ppgool status of two indicates that the node is connected and an active part of the cluster. There should be only one primary server, with zero or more standby servers:

```
node_id | hostname | port | status | lb_weight | role
-----+-----+-----+-----+-----+-----
0       | machine1 | 5432 | 2      | 0.250000 | primary
1       | machine2 | 5432 | 2      | 0.250000 | standby
(2 rows)
```

Note: If starting (recovering) a node fails, refer to [the Troubleshooting section on page 9](#).

Get Current Passwords

1 Sign on to the machine that runs the SAS Configuration Server, which is based on Consul, as the SAS install user (sas) or with sudo privileges.

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:

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

- Install accounts with sudo privileges (that are not the install user):

```
export CONSUL_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

The script, sds_change_user_pw.sh, changes either the [dbmsowner](#) account password or the [sas account](#) password for SAS Infrastructure Data Server. When run, sds_change_user_pw.sh also synchronizes the new password with configuration files and the SAS Configuration Server, which is based on Consul.

CAUTION! To avoid data loss, change the sas user account password only during a scheduled maintenance window, when users are not accessing the SAS Viya system. 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](#).

1 Log on to the SAS Infrastructure Data Server machine as the SAS install user (sas).

2 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`.

3 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 10](#).

4 Using the location of sds_env_var.sh noted in [Step 2](#), invoke the script using the following command:

```
/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.

- 5 Enter the information that you collected in [Step 3](#) as the script prompts you for it.

After you provide values for 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 to 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.

- 6 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 dbmsowner postgres
```

- 7 When prompted, enter the new password for dbmsowner.
- 8 Restart all SAS Viya services.

For more information, see [“All Servers and Services” in SAS Viya Administration: General Servers and Services](#).

Recover a Failed HA Cluster

If your primary node on the SAS Infrastructure Data Server has failed to start, has crashed, or becomes unavailable, then you must promote a standby server to primary status. If you do not have any more standby servers available for promotion, then your cluster has failed and needs to be rebuilt. The cluster does not start until you complete the recovery process.

- 1 Fix the problem that caused the cluster to fail.
- 2 Sign in to the pgpool-II machine as the SAS install user (`sas`) or with `sudo` privileges.
- 3 Stop the data server cluster by running the appropriate command:

- Red Hat Linux version 6:

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

- Red Hat Linux version 7:

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

- 4 Using a text editor, open the cluster definition file and mark all of the data server nodes as healthy. (Performing this step causes the data server start-up script to start the previously failed node.)

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

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

- 5 Restart the data server cluster by running the appropriate command:

- Red Hat Linux version 6:

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

- Red Hat Linux version 7:

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

- 6 Check the status of the HA data server cluster by running the following command:

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

A pgpool status of 2 indicates that the node is connected and an active part of the cluster. There should be only one primary server, with zero or more standby servers.

Failback an HA Cluster

Failback means reconstituting the high availability (HA) SAS Infrastructure Data Server cluster to its original configuration before the failover.

Note: SAS Configuration Server, which is based on Consul, must be running before you can failback a high availability SAS Infrastructure Data Server cluster.

Note: This document assumes that the primary node in your HA cluster is node 0.

- 1 Sign in to the pgpool-II machine as the SAS install user (sas) or with sudo privileges.
- 2 Check the status of the cluster to determine that the current primary is healthy, by running the following command:

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

In the following output, node 0 has a status of 3, indicating that it is unhealthy. The failover primary, node1, is healthy, because it has a status of 2:

node_id	hostname	port	status	lb_weight	role
0	machine1	5432	3	0.250000	standby
1	machine2	5432	2	0.250000	primary
2	machine3	5432	2	0.250000	standby
3	machine4	5432	2	0.250000	standby

(4 rows)

- 3 Sign in to the node 0 cluster machine as the SAS install user (sas) or with sudo privileges.
- 4 Perform the steps described in [Start a Node \(Recover a Node\)](#).
- 5 Sign in to the pgpool-II machine as the SAS install user (sas) or with sudo privileges.
- 6 Check the status of the cluster to determine that node 0 is now a healthy standby, by running the following command:

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

In the following output, node 0 has a status of 2, indicating that it is now healthy:

node_id	hostname	port	status	lb_weight	role
0	machine1	5432	2	0.250000	standby
1	machine2	5432	2	0.250000	primary
2	machine3	5433	2	0.250000	standby
3	machine4	5433	2	0.250000	standby

(4 rows)

- 7 Sign in to the current primary node machine as the SAS install user (sas) or with sudo privileges.
- 8 Restart the current primary node, by running the appropriate command:

Note: In this document, the current primary is node 1.

- Red Hat Linux version 6:

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

- Red Hat Linux version 7:

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

Restarting the current primary node triggers a failover. The pgpool-II server automatically selects node 0 as the new primary node, if it is healthy.

- 9 Sign in to the pgpool-II machine as the SAS install user (sas) or with sudo privileges.

- 10 Check the status of the cluster.

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

Determine that node 0 is now the primary, and the previous primary node is now a healthy standby:

node_id	hostname	port	status	lb_weight	role
0	machine1	5432	2	0.250000	primary
1	machine2	5432	2	0.250000	standby
2	machine3	5433	2	0.250000	standby
3	machine4	5433	2	0.250000	standby

(4 rows)

Clean up after a Hardware Failure

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

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

- 1 Sign in to the pgpool-II machine as the SAS install user (sas) or with sudo privileges.

- 2 Stop the HA data server cluster by running the appropriate command:

- Red Hat Linux version 6:

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

- Red Hat Linux version 7:

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

- 3 Delete any socket-lock (".s.PGSQL.xxxx") or PID file (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.lock
- /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 appropriate command:

- Red Hat Linux version 6:


```
sudo service sas-viya-sasdatasvrc-postgres start
```
- Red Hat Linux version 7:


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

SAS Infrastructure Data Server: Concepts

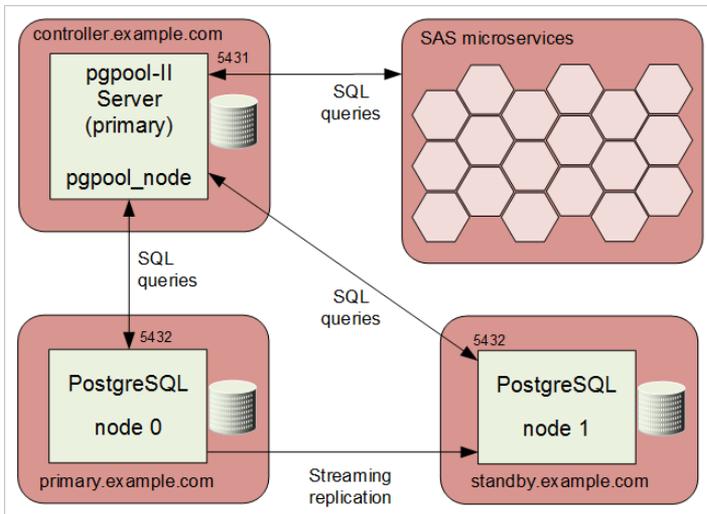
What Is the SAS Infrastructure Data Server?

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.

The server is automatically configured during installation and configuration. By default, the SAS install account (sas) 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 Recovery](#).

SAS Infrastructure Data Server Architecture



pgpool-II

SAS provides pgpool-II (version 3) open-source software to enable you to manage PostgreSQL clusters. 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.

pgpool-II is capable of providing the following:

- high availability (failover management)
- load balancing (read-only scaling with SELECT queries, but not writes)

- connection pooling

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 [“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.

If a standby node is shut down, SAS Infrastructure Data Server fails to restart this node the next time the cluster is started.

Explanation:

SAS Infrastructure Data Server does not follow the proper recovery process for standby nodes that are shut down.

The `pgpool-II` server runs `failover.sh` on the standby node and marks the node “unhealthy.” The next time the cluster starts, this unhealthy node remains down.

Resolution:

The workaround is to start (recover) the standby node following the steps listed in [Start a Node \(Recover a Node\) on page 2](#). SAS has fixed this issue in the next version of SAS Viya.

SAS Infrastructure Data Server: 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, then 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 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 be at least six characters long.
- The password can contain a mix of alphanumeric characters.
- The following special characters are allowed:
 - hashtag (#)
 - asterisk (*)
 - plus sign (+)
 - question mark (?)
- There are no restrictions for including numbers or mixed-case characters.

Environment Parameters

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

- /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

