

# Managing SAS® Web Infrastructure Platform Data Server High-Availability Clusters

Ken Young, SAS Institute Inc.

## ABSTRACT

The SAS Web Application Server is a lightweight server that provides enterprise-class features for running SAS middle-tier web applications. This server can be configured to use the SAS Web Infrastructure Platform Data Server for a transactional storage database. You can meet the high-availability data requirement in your business plan by implementing a SAS Web Infrastructure Platform Data Server cluster. This paper focuses on how the SAS Web Infrastructure Platform Data Server on the SAS middle tier can be configured for failover and data replication involving multiple nodes. SAS Environment Manager and pgpool-II are used to enable these high-availability strategies, to monitor the server status, and to initiate failover as needed.

## INTRODUCTION

### OVERVIEW

Why is high availability important in the SAS Web Application Server? The SAS Web Infrastructure Platform is a collection of services and applications that provide common infrastructure features for use by SAS web applications. By default, the SAS Web Infrastructure Platform uses the SAS Web Infrastructure Platform Data Server to store alerts, comments, workflows, and SAS Content Server data. If the SAS Web Infrastructure Platform Data Server becomes unavailable, then the SAS Web Infrastructure Platform services will likewise become unavailable. This effect will cause the entire middle-tier node to effectively go offline, since none of the web server applications will be able to log on successfully. For SAS 9.4, the SAS Web Infrastructure Platform Data Server is version 9.1.9 of the PostgreSQL server.

The configuration enables the primary data server to replicate data to one or more standby data servers. These standby servers are configured to take over as primary server if a failure occurs. High availability is provided for all databases in the SAS Web Infrastructure Platform Data Server that are used by the SAS Web Application Server, including the SharedServices, EVManager, and Administration databases. These databases are described in the [SAS 9.4 Intelligence Platform: Middle-Tier Administration Guide](#).

This configuration is summarized in the following table:

SAS Web Infrastructure Platform Data Server	Configuration	Hosted Databases
SAS Web Infrastructure Platform Data Server 1	Set up for replication as primary. This server sends all data to one or more standbys.	SharedServices, EVManager, Administration, vdbadm, transportsvc_db (and others).
SAS Web Infrastructure Platform Data Server 2+	Set up for replication as standby. This server receives all data from the primary server, but remains in a read-only state. There are one or more standby servers.	Read-only copies of the databases, which are replicated from SAS Web Infrastructure Platform Data Server 1 in real time.

The installation and configuration process involves these seven high-level steps:

1. Install and configure the SAS Metadata Server and the first SAS Web Infrastructure Platform Data Server using the SAS Deployment Wizard.
2. Configure the first SAS Web Infrastructure Platform Data Server for high availability.
3. Install and configure a standby high-availability SAS Web Infrastructure Platform Data Server.

4. Configure pgpool-II with each SAS Web Infrastructure Platform Data Server.
5. Install and configure the pgpool-II Administration Tool (pgpoolAdmin).
6. Configure the middle tier to use the SAS Web Infrastructure Platform Data Server high-availability cluster.
7. Maintain the SAS Web Infrastructure Platform Data Server high-availability cluster using pgpoolAdmin.

**Caution:** There is a potential for data loss or corruption if the failover components are configured improperly. Never configure more than one primary SAS Web Infrastructure Platform Data Server for use simultaneously. Doing so can result in data loss or corruption that can be very difficult to recover from without full backups of the databases. Make regular backups of the databases as a best practice. The high-availability configuration is not a substitute for making regular backups.

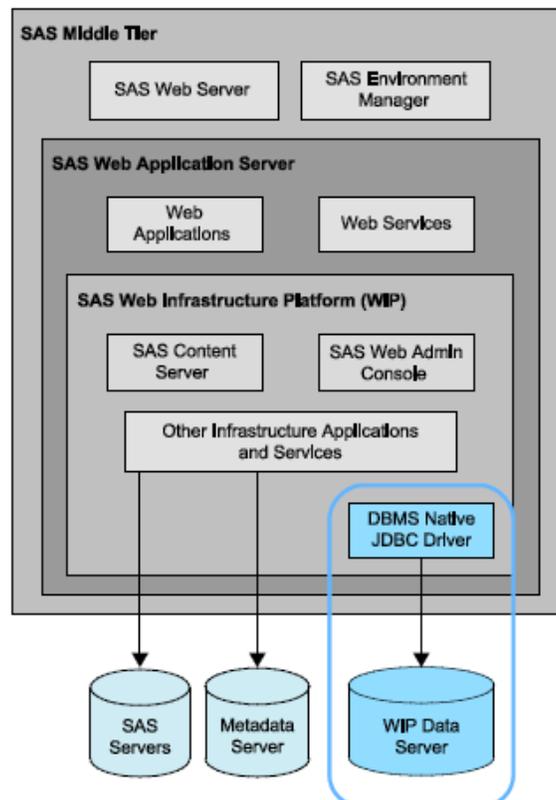
## BACKGROUND

The SAS Web Infrastructure Platform is a collection of middle-tier services and applications. This collection provides basic integration services and infrastructure that support SAS applications and solutions. The collection is delivered as part of the SAS Integration Technologies package. As such, all business intelligence applications, data integration applications, and SAS solutions have access to the SAS Web Infrastructure Platform as part of their standard product bundling.

The SAS Web Infrastructure Platform Data Server is a SAS version of the PostgreSQL DBMS that is configured to be used with SAS.

This paper focuses on how to reconfigure the SAS Web Infrastructure Platform Data Server on a middle-tier node. The default data server configuration is outlined in blue in Figure 1. The resulting configuration, shown in Figure 2, includes a primary data server and a second standby data server. Additional standby data servers can be added, if necessary.

Figure 1



## ENHANCEMENTS TO THE CURRENT SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER HIGH AVAILABILITY TECHNICAL PAPER

The process described in this paper corrects most of the limitations of the process described in the previously published SAS Technical Paper “[Configuring the Web Infrastructure Platform Data Server for High Availability](#).” Specifically, the following features are now supported:

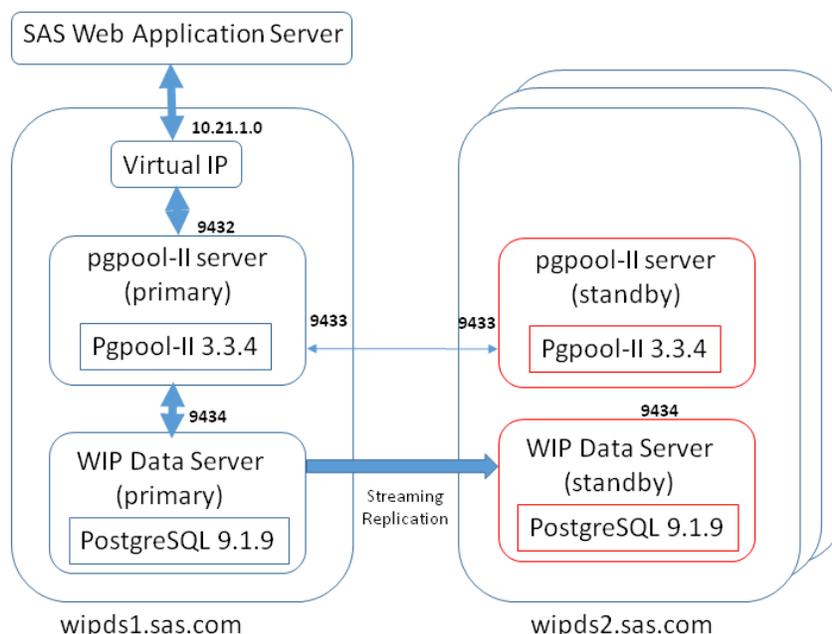
- Multiple standby high-availability servers can be configured.
- High availability is configured on the server side only. There are no longer any client-side high-availability components.
- All SAS Web Infrastructure Platform databases can be configured for high availability, including the EVManager and Administration databases.
- All SAS Web Infrastructure Platform Data Server interfaces (for example, JDBC, ODBC, and pg\_ctl) are included.
- The SAS JDBC Driver for High Availability is no longer used.
- There is no single point of failure, since SAS Environment Manager no longer controls failover.
- An existing SAS Web Infrastructure Platform Data Server can be updated without SAS Deployment Wizard reconfiguration.

## GATHER INFORMATION ABOUT THE SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER HIGH-AVAILABILITY CLUSTER

These instructions assume that you want to deploy one primary data server and one standby data server. However, you can deploy additional standby data servers by repeating the steps in the “Initialize and Configure the Standby Data Server” section.

Before deploying the SAS Web Infrastructure Platform Data Servers, record the IP address, host name, and port number for each server. In our example, the primary data server host is wipds1.sas.com (IP address 10.21.1.1) and the standby data server host is wipds2.sas.com (IP address 10.21.1.2). The 10.21.1.0 IP address is used by pgpool-II as the virtual address for the SAS Web Infrastructure Platform Data Server. Figure 2 shows how these elements relate to one another.

Figure 2



The following table shows the example values that are used in the instructions. Substitute your site's values for these values.

<b>Example Values Used in These Instructions</b>	
Hostnames	wipds1.sas.com (primary), wipds2.sas.com (standby)
IP Addresses	10.21.1.0 (virtual) 10.21.1.1 (primary) 10.21.1.2 (standby)
pgpool-II data port	9432 <sup>1</sup>
pgpool-II administration port	9433
SAS Web Infrastructure Platform Data Server port	9434
SASHome	/install/SAS/SASHome
WIPCONFIG <sup>2</sup> (SAS Web Infrastructure Platform Data Server data and log locations)	/install/SAS/config/Lev1/WebInfrastructurePlatformDataServer
WIPINSTALL <sup>2</sup> (data server installation directory)	/install/SASHome/SASWebInfrastructurePlatformDataServer/9.4
SAS Web Infrastructure Platform Data Server administrator account (internal)	dbmsowner
Apache HTTP server and pgpool-II account (external)	apache (this account needs root or SUDO privileges)
SAS installer account (external)	cfgsas1

*Notes:*

1. Clients of the SAS Web Infrastructure Platform Data Server connect to the pgpool-II data port (9432), not to the actual SAS Web Infrastructure Platform Data Server port (9434).
2. WIPCONFIG and WIPINSTALL are UNIX environment variables. Set these to the appropriate values used in your deployment.

## **STEP 1: INSTALL AND CONFIGURE THE SAS METADATA SERVER AND THE FIRST SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER**

The first step is to install and configure the primary and standby SAS Web Infrastructure Platform Data Servers. The databases that are used by the SAS Web Infrastructure Platform can then be installed on those data server instances.

These instructions assume that you have not yet deployed SAS and that you are using the SAS Deployment Wizard to prepare the SAS Web Infrastructure Platform Data Servers for your main SAS deployment. You will use the SAS Deployment Wizard to install two or more SAS Web Infrastructure Platform Data Servers, and then you will manually configure the data servers in a master/slave arrangement for failover. You will need access to the SAS Software Depot in order to run the installation.

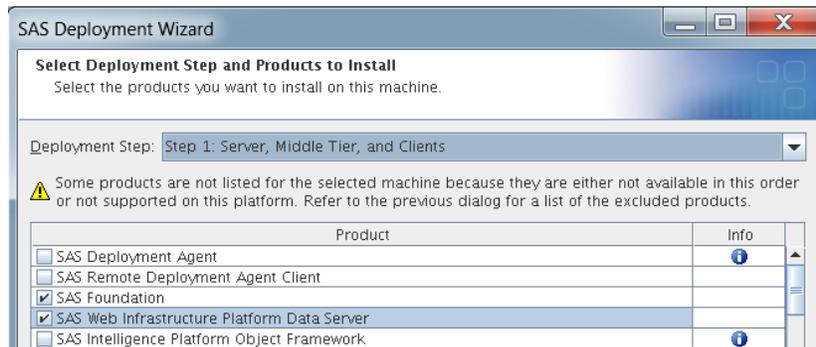
The SAS Web Infrastructure Platform Data Server has a dependency on the SAS Metadata server. In these instructions, the metadata server is deployed in conjunction with the data server. Alternatively, the metadata server can be deployed on a separate node prior to deploying the data server. Our example deploys the metadata server on wipds1.sas.com, which will host the primary data server.

Prior to deploying the SAS Web Infrastructure Platform Data Server, the virtual IP address (10.21.1.0) should be setup on wipds1.sas.com. The virtual address should be specified in the SAS Deployment Wizard as the host for the data server.

The steps are as follows:

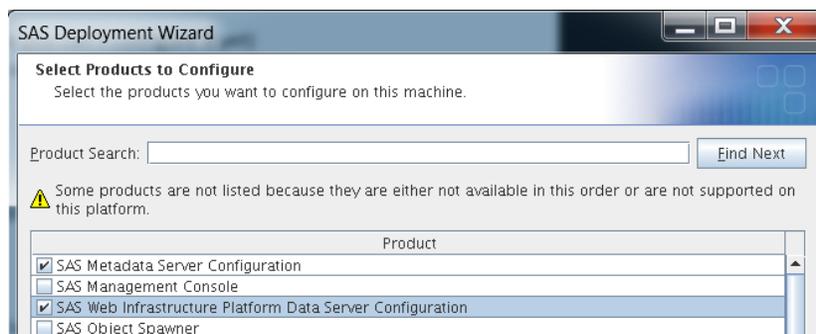
1. Start the SAS Deployment Wizard.
2. Choose your language.
3. Choose **Install SAS Software**. Use the default path for SASHome, or choose another path. Then click **Next**.
4. Choose **Perform a Planned Deployment**. Check both **Install SAS Software** and **Configure SAS Software**. Then click **Next**.
5. Choose a plan file that contains both the SAS Metadata Server and SAS Web Infrastructure Platform Data Server configurations. Then click **Next**.
6. On the Deployment Step and Products to Install wizard page, select a **Server Tier** deployment step. Then check **SAS Foundation** and the **SAS Web Infrastructure Platform Data Server** product to install, as show in Figure 3. Then click **Next**.

Figure 3



7. On the Products to Configure wizard page, check **SAS Metadata Server Configuration** and **SAS Web Infrastructure Platform Data Server Configuration**, as shown in Figure 4. Then click **Next**.

Figure 4



8. On the Local Machine Name wizard page, specify the long (**wipds1.sas.com**) and short (**wipds1**) local host names of the current machine.
9. On the SAS Metadata Server wizard page, specify the host name and port of the SAS Metadata Server.
10. Click **Next** on the subsequent wizard pages.
11. Click **Start** to initiate the installation and configuration. Wait for the deployment to finish.

## STEP 2: CONFIGURE THE FIRST SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER FOR HIGH AVAILABILITY

On `wipds1.sas.com`, manually configure the SAS Web Infrastructure Platform Data Server to be the primary high-availability server, as follows:

1. Open a command shell.
2. Stop the SAS Web Infrastructure Platform Data Server by executing the `webinfdsvrc.sh` command in the configuration directory. For example:

```
$ $WIPCONFIG/webinfdsvrc.sh stop
```

3. Edit the SAS Web Infrastructure Platform Data Server control file, `$WIPCONFIG/webinfdsvrc.sh`. In the file, change the port value for the server from the value that was specified in the SAS Deployment Wizard (9432 is the default) to the value that you recorded in the table on page 4 (for example, 9434).

For example, suppose the control file contains the following line:

```
nohup $COMMAND start -D "/install/SAS/config/Levl/WebInfrastructurePlatformDataServer/data" -  
o "-i -p 9432" > $LOGSDIR/webinfdsvrc_console.log 2>&1 &
```

You would change the line to the following:

```
nohup $COMMAND start -D "/install/SAS/config/Levl/WebInfrastructurePlatformDataServer/data" -  
o "-i -p 9434" > $LOGSDIR/webinfdsvrc_console.log 2>&1 &
```

As stated earlier, the SAS Web Application Server will connect to the `pgpool-II` data port and not directly to the SAS Web Infrastructure Platform Data Server.

4. Edit the remote access permissions configuration file, `$WIPCONFIG/data/pg_hba.conf`. Add the following lines:

```
IPv4:  
#TYPE      DATABASE      USER      IP-ADDRESS      IP-MASK      METHOD  
host      replication  all       10.21.1.2       255.255.255.255  trust  
host      all          all       all              all           md5
```

```
IPv6:  
#TYPE      DATABASE      USER      IP-ADDRESS/MASK      METHOD  
host      replication  all       fe80::90a7:5556:b102:3fc3%11/128  trust  
host      all          all       ::1/128           md5
```

**Note:** 10.21.1.2 and fe80::90a7:5556:b102:3fc3%11 are example IP addresses. Substitute your actual IP values. This allows a process from any remote host to connect with MD5 authentication, and it allows the standby SAS Web Infrastructure Platform Data Server (that is installed later) to connect to this server to replicate its data.

5. Make a backup copy of the `postgresql.conf` configuration file:

```
$ cp postgresql.conf postgresql.orig
```

6. Edit the `postgresql.conf` configuration file to set the configuration values, as shown in this example:

```
listen_addresses = '*'  
wal_level = hot_standby  
max_wal_senders=10  
wal_keep_segments=200  
hot_standby=on  
checkpoint_segments=30  
log_truncate_on_rotation = off  
log_rotation_age = 1d  
log_min_messages = info
```

7. Start the primary SAS Web Infrastructure Platform Data Server by executing the `webinfdsvrc.sh` command in the configuration directory, as follows.

```
$ $WIPCONFIG/webinfdsvrc.sh start
```

You can check the status of SAS Web Infrastructure Platform Data Server with this command:

```
$ $WIPCONFIG/webinfdsvrc.sh status
```

8. Check the log file in the log directory (\$WIPCONFIG/Logs) to make sure that the server is running without problems. The contents of the file should be similar to the following:

```
LOG: database system is ready to accept connections
```

```
LOG: autovacuum launcher started
```

### STEP 3: INSTALL AND CONFIGURE A STANDBY HIGH-AVAILABILITY SERVER

#### INSTALL A STANDBY SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER

Using the SAS Deployment Wizard, install a second SAS Web Infrastructure Platform Data Server on wipds2.sas.com. The server will be manually configured as the standby data server later on. Follow these steps:

1. Start the SAS Deployment Wizard.
2. Choose your language.
3. Choose **Install SAS Software**. Use the default path for SASHome, or choose another path. Then click **Next**.
4. Choose **Install Additional Software**.
5. Scroll down in the products list and select **SAS Web Infrastructure Platform Data Server**.
6. Click **Next** on the subsequent wizard pages.
7. Click **Start** to begin the installation. Wait for the installation to finish.

#### CONFIGURE AND INITIALIZE THE STANDBY SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER

Log on as the user that you will use to run the SAS deployment later in this process. Then follow these steps to initialize and configure a standby data server. In our example, the steps are executed on wipds2.sas.com.

1. Connect (via SSH, Remote Desktop, or whatever is appropriate for your environment) to the machine that will host the standby data server. In our example, this machine is named wipds2.sas.com.
2. Open a command shell.
3. Add the path to the SAS Web Infrastructure Platform Data Server binaries to your PATH environment variable for this command shell process. Set the path as shown below, replacing the example SASHome location with your actual location. Also, set the LD\_LIBRARY\_PATH environment variable. The example is for a Linux system. Use the appropriate environment variable for your system (for example, \$LIBPATH or \$SHLIB\_PATH).

```
$ export PATH=$WIPINSTALL/bin:$PATH
```

```
$ export LD_LIBRARY_PATH=$WIPINSTALL/lib:$LD_LIBRARY_PATH
```

4. Create a parent directory to store the SAS Web Infrastructure Platform Data Server data and log directories. This should be outside of SASHome. For example, create the same directory structure that was used with the primary SAS Web Infrastructure Platform Data Server, including the “data” and “Logs” directories:

```
$ mkdir $WIPCONFIG/data
```

```
$ mkdir $WIPCONFIG/Logs
```

You can use the df -hl command to find the list of local storage devices and the free space on each.

5. Run the initialize database command, initdb, as in the following example:

```
$ initdb -D $WIPCONFIG/data -U dbmsowner -W -A md5 -E UTF8
```

The initdb command creates a data directory for the files that contain the data for the SAS Web Infrastructure Platform Data Server databases. The initdb command also creates a DBMS administrative user. You can choose the name of the administrative user; for example, you could use the name dbmsowner. Specify a password for the administrative user when prompted. Record the administrative user's user ID and password for future reference. This user should have the same name and password as the administrative user on the primary data server.

6. Edit the remote access permissions configuration file, \$WIPCONFIG/data/pg\_hba.conf. Add the following lines:

```
IPv4:
#TYPE  DATABASE      USER      IP-ADDRESS      IP-MASK          METHOD
host   replication  all      10.21.1.2      255.255.255.255  trust
host   all           all      all            all              md5

IPv6:
#TYPE  DATABASE      USER      IP-ADDRESS/MASK          METHOD
host   replication  all      fe80::90a7:5556:b102:3fc3%11/128  trust
host   all           all      all                       md5
```

**Note:** 10.21.1.2 and fe80::90a7:5556:b102:3fc3%11 are example IP addresses. Substitute your actual IP values. This allows a process from any remote host to connect with MD5 authentication.

7. Make a backup copy of the postgresql.conf configuration file, using a command similar to the following:

```
$ cp $WIPCONFIG/data/postgresql.conf $WIPCONFIG/data/postgresql.orig
```

8. Edit the postgresql.conf configuration file to set the appropriate configuration values, as shown in the following example. Update the log\_directory value to match the log subdirectory that you created in step 4.

```
listen_addresses = '*'
max_connections = 256 # Recommended setting for a single web server.
shared_buffers = 256MB # Recommended setting for a single web server.
max_prepared_transactions = 256 # Recommended setting for a single web server.
work_mem = 2MB # Recommended setting for a single web server.
wal_level = hot_standby
wal_buffers = 16MB # Recommended setting for a single web server.
checkpoint_segments=30
checkpoint_completion_target = 0.9 # Recommended setting for a single web server.
max_wal_senders=10
wal_keep_segments=200
hot_standby=on
effective_cache_size = 2GB # Recommended setting for a single web server.
log_destination = 'stderr'
logging_collector = on
log_directory = '<path_to_log_dir>' # For example:
# '/install/SAS/config/Levl/WebInfrastructurePlatformDataServer/Logs'
log_truncate_on_rotation = off
log_rotation_age = 1d
log_min_messages = info
```

9. Use the following command to verify that the standby SAS Web Infrastructure Platform Data Server will start.

```
$ pg_ctl -D $WIPCONFIG/data -o "-p 9434" start
```

You can check the status of the data server with this command:

```
$ pg_ctl status -D $WIPCONFIG/data
```

10. Check the log file in the log directory to make sure that the server is running without problems. The contents of the log should be similar to the following:

```
LOG: database system is ready to accept connections
LOG: autovacuum launcher started
```

## CONFIGURE THE STANDBY SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER FOR HIGH AVAILABILITY

1. Stop the SAS Web Infrastructure Platform Data Server on the standby machine, using the following command:
2. Copy the data directory from the primary data server directory to the standby data server directory, as follows:

```
$ pg_ctl -D $WIPCONFIG/data stop
```

- a) Use the following command to put the primary data server machine (with host name wipds1 in this example) into base backup mode:

```
$ psql -c "SELECT pg_start_backup('stage backup', true)" -h wipds1.sas.com -p 9434 -U dbmsowner -d postgres
```

- b) Copy the primary server's data directory to the standby data server's machine (for example, wipds2). You can use the command rsync or the command scp to perform the copy. For example, on the standby machine, you would run the following command:

```
$ rsync -a wipds1.sas.com:$WIPCONFIG/data $WIPCONFIG/data --exclude postmaster.pid
```

- c) Use the following command to take the primary data server out of base backup mode:

```
$ psql -c "SELECT pg_stop_backup()" -h wipds1.sas.com -p 9434 -U dbmsowner -d postgres
```

3. Edit the pg\_hba.conf file in the data directory. In the line that enables replication, specify the IP address of the primary data server. This line will not be used initially, because no SAS Web Infrastructure Platform Data Server will initially be replicating data from the standby server. However, the line will make recovery easier if the primary data server fails in the future.
4. Create a recovery.conf file in the data directory on the standby data server. This file causes the standby server to connect to the primary server for streaming replication. The standby server streams data from the primary server in near-real time as it is added or updated.
5. Add the following content to the recovery.conf file. Be sure to replace the IP address with the actual IP address of the machine that hosts your primary SAS Web Infrastructure Platform Data Server.

*Note:* The recovery.conf file must be in the DATA directory.

```
standby_mode = 'on'
# Specifies a connection string which is used for the standby server to
# connect with the primary.
primary_conninfo = 'host=10.21.1.1 port=9434 user=dbmsowner'
# Specifies a trigger file whose presence should cause streaming
# replication to end (i.e., failover). Use double-backslashes on Windows.
trigger_file = '/path/to/trigger.tmp'
```

6. Start the standby SAS Web Infrastructure Platform Data Server on host wipds2.sas.com. If configured correctly, it starts streaming replication. Use the same command to start the server that you used when you first installed it. (See the preceding section on configuring the primary SAS Web Infrastructure Platform Data Server for a reminder.)
7. On the standby server, check the SAS Web Infrastructure Platform Data Server log file to ensure that the server is operating correctly, that it has not raised errors, and that streaming replication is working. At this point, the data server is configured for replication. Look in the log directory for the most recent file. The log file contents should look similar to the following:

```
LOG: entering standby mode
LOG: consistent recovery state reached at 0/2EF9B8F0
LOG: redo starts at 0/2EF0E4D0
```

```
LOG: database system is ready to accept read only connections
LOG: streaming replication successfully connected to primary
```

## STEP 4: CONFIGURE PGPOOL-II WITH EACH SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER

Manually install and configure pgpool-II on each SAS Web Infrastructure Platform Data Server host (wipds1.sas.com and wipds2.sas.com). The pgpool-II tool monitors the status of the primary data server status and triggers failover to a standby data server when the primary server becomes unavailable. This tool also dynamically controls IP address assignment on each machine using a watchdog process. When the installation and configuration is complete, you will have a running pgpool-II server on each host. Each host will have a populated pgpool-II configuration directory. These files are available as a separate download:

```
$WIPCONFIG/pgpool/etc
    pgpool.conf
    failover.sh
    follow_master.sh
    recovery.sh
    pgpool_start.sh
    pcp.conf
    pgpool_passwd
```

Follow these steps to install pgpool-II and its extensions on each SAS Web Infrastructure Platform Data Server:

1. Retrieve the pgpool-II build for your platform from the appropriate repository, and then use the following command to install pgpool-II:

```
$ sudo yum localinstall pgpool-II-91-3.3.2-1.server.rpm
```

2. Using the following command, copy the pgpool-II shared objects needed by the SAS Web Infrastructure Platform Data Server from the pgpool-II installation directory:

```
$ cp /usr/lib64/pgsql/*.so $WIPINSTALL/lib
```

3. Use the following command to install the pgpool\_regclass function. This is needed by pgpool-II to handle duplicate table names in different schemas.

```
$ psql -p 9432 -U dbmsowner -f /usr/share/pgsql/contrib/pgpool-
regclass.sql template1
```

4. Use the following command to create the pgpool\_catalog.insert\_lock table. This table is used by pgpool-II to coordinate table locking when insert\_lock is set in replication mode.

```
$ psql -p 9432 -U dbmsowner -f /usr/share/pgsql/contrib/insert_lock.sql
template1
```

5. Use the following command to enable the pgpool-II online recovery functions. These functions enable pgpoolAdmin to stop, restart, and reload the SAS Web Infrastructure Platform Data Servers.

```
$ psql -p 9432 -U dbmsowner -f /usr/share/pgsql/contrib/pgpool-
recovery.sql template1
```

6. Configure the pgpool-II control interface file, pcp.conf. This file is used by pgpoolAdmin to collect pgpool-II status and terminate pgpool-II processes remotely. The file must be readable by the Apache HTTP Server. Use the following command to encrypt and store the password (“mypassword” in the example below) of pgpool-II for PHP:

```
$ /usr/local/bin/pg_md5 mypassword
```

The MD5 encrypted password will be displayed. Use the following command to save the “apache” user account and the encrypted password to the pcp.conf file:

```
$ sudo echo "apache:34819d7beeabb9260a5c854bc85b3e44" > /usr/local/etc/pcp.conf
```

7. Configure the pgpool-II settings file, \$WIPCONFIG/pgpool/pgpool.conf. Each back-end SAS Web Infrastructure Platform Data Server instance must be defined, beginning with the primary data server. The key parameters to change are:

```
memqcache_oiddir =
'/install/SAS/config/Levl/WebInfrastructurePlatformDataServer/pgpool/log/pgpool/oiddir'
failover_command =
'/install/SAS/config/Levl/WebInfrastructurePlatformDataServer/pgpool/failover.sh %d %h %p %D
%m %M %H %P %r %R'
port = 9433 # pgpool-II port
load_balance_mode = on # Activate load balancing mode
pcp_port = 9434 # pgpool-II admin port
pid_file_name = '/install/SAS/pgpool/run/pgpool.pid'
logdir = '/install/SAS/config/Levl/WebInfrastructurePlatformDataServer/pgpool/log'
follow_master_command =
'/install/SAS/config/Levl/WebInfrastructurePlatformDataServer/pgpool/follow_master.sh %d %h
%p %D %m %M %H %P %r %R'
# Watchdog configuration
use_watchdog = on # Activate watchdog
delegate_IP = '10.21.1.0' # Virtual IP address
wd_hostname = '10.21.1.1' # IP address of this watchdog
wd_port = 9000 # port number for watchdog service
ifconfig_path = '/home/apache/sbin' # ifconfig command path
arping_path = '/home/apache/sbin' # arping command path
wd_interval = 3 # lifecheck interval in seconds
other_pgpool_hostname0 = '10.21.1.2' # IP address other pgpool server
other_pgpool_port0 = 9433 # Port number for other pgpool
other_wd_port0 = 9434 # Port number for other watchdog
# Primary WIP DS
backend_hostname0 = 'wipds1'
backend_port0 = 9432
backend_weight0 = 1
backend_data_directory0 = '/install/SAS/config/Levl/WebInfrastructurePlatformDataServer/data'
# Standby WIP DS
backend_hostname1 = 'wipds2'
backend_port1 = 9432
backend_weight1 = 1
backend_data_directory1 = '/install/SAS/config/Levl/WebInfrastructurePlatformDataServer/data'
```

8. Update the script \$WIPCONFIG/failover.sh with connection information for pgpool-II and the SAS Web Infrastructure Platform Data Server, as shown in the following example. This script is used by pgpool-II to promote a standby data server to become the primary server. The key parameters to change are:

```
myconfigdir=/install/SAS/config/Levl/WebInfrastructurePlatformDataServer
myinstalldir=/install/SAS/SASHome/SASWebInfrastructurePlatformDataServer/9.4
log=$myconfigdir/pgpool/log/failover.log
pg_ctl=$myhomedir/bin/pg_ctl
cluster=$myconfigdir/data
```

9. Update the script \$WIPCONFIG/follow\_master.sh with configuration information for pgpool-II and the SAS Web Infrastructure Platform Data Server, as shown in the following example. After a failover occurs, this script is used by pgpool-II to degenerate all remaining standby servers and have them follow the new primary server. The key parameters to change are:

```
myconfigdir=/install/SAS/config/Levl/WebInfrastructurePlatformDataServer
myinstalldir=/install/SAS/SASHome/SASWebInfrastructurePlatformDataServer/9.4
log=$myconfigdir/pgpool/log/failover.log
pg_ctl=$myhomedir/bin/pg_ctl
cluster=$myconfigdir/data
PCP_PORT=9433
pcp_super_user=apache
```

10. Start pgpool-II on each SAS Web Infrastructure Platform Data Server node, beginning with the

primary node. To verify correct operation, execute the following command:

```
$ ps -e -o pid,cmd | grep pgpool
```

Messages like the following should appear, indicating that the pgpool-II process is running and awaiting connections:

```
10119 /usr/local/bin/pgpool -D -n -f $WIPCONFIG/pgpool/pgpool.conf -F
$WIPCONFIG/pgpool/pcp.conf -a $WIPCONFIG/pgpool/pool_hba.conf
10120 pgpool: wait for connection request
10121 pgpool: wait for connection request
10122 pgpool: wait for connection request
10123 pgpool: wait for connection request
10124 pgpool: wait for connection request
10125 pgpool: wait for connection request
10126 pgpool: wait for connection request
10127 pgpool: wait for connection request
10128 pgpool: wait for connection request
10129 pgpool: wait for connection request
```

11. Repeat steps 1-9 on all remaining standby SAS Web Infrastructure Platform Data Server nodes.

## STEP 5: INSTALL AND CONFIGURE THE PGPOOL-II ADMINISTRATION TOOL

The pgpool-II Administration Tool (pgpoolAdmin) is a web-based management tool for pgpool-II. It is used to monitor, start, and stop pgpool-II. This tool also controls pgpool-II's settings. Either YUM (Yellowdog Updater Modifier) or the RPM package manager can be used to retrieve and install UNIX packages. Most of the installation packages need to be retrieved and installed with root system privileges. This is typically done using the "apache" user account. The Apache HTTP server and PHP must be installed on the primary SAS Web Infrastructure Platform Data Server high-availability node (wipds1.sas.com). Follow these steps:

1. Install the latest version of PHP, using the following command. This version must be 4.4.2 or higher.

```
$ sudo yum install php, php -ini
```

2. Install the PostgreSQL extension for PHP, using the following command:

```
$ sudo yum install php-pgsql
```

3. Install the latest version of Apache HTTP Server, using the following command:

```
$ sudo yum install httpd
```

4. Install pgpoolAdmin, using the following command:

```
$ sudo yum install pgpoolAdmin
```

To verify the correct installation of pgpoolAdmin, execute the following command:

```
$ rpm -qi pgpoolAdmin
```

Messages similar to the following should appear:

```
Name       : pgpoolAdmin           Relocations: (not relocatable)
Version    : 3.4.0                Vendor: (none)
Release    : 1.pgdg              Build Date: Fri 07 Nov 2014 01:12:45 AM EST
Install Date: Fri 05 Dec 2014 03:16:04 PM EST   Build Host: 93eb297ac0f2
Group      : Applications/Databases   Source RPM: pgpoolAdmin-3.4.0-1.pgdg.src.rpm
Size       : 2565595               License: BSD
Signature  : (none)
URL        : http://pgpool.net/
Summary    : PgpoolAdmin - web-based pgpool administration
```

5. Configure pgpoolAdmin to run with the Apache HTTP Server. File ownership of key pgpoolAdmin files need to be changed to the apache user.

```
$ sudo chmod 777 /var/www/html/pgpoolAdmin/templates_c
```

```

$ sudo chown apache:apache /usr/local/etc/pcp.conf
$ sudo chmod 644 /usr/local/etc/pcp.conf
$ sudo chown apache:apache /var/www/html/pgpoolAdmin/conf/pgmgt.conf.php
$ sudo chmod 644 /var/www/html/pgpoolAdmin/conf/pgmgt.conf.php

```

6. Start the Apache HTTP Server, using the following command:

```
$ sudo /etc/init.d/httpd start
```

To verify the correct installation of the Apache HTTP Server, execute the following command:

```
$ sudo /etc/init.d/httpd status
```

The following message should appear indicating that the server is running:

```
httpd (pid 9655) is running...
```

7. To verify the correct integration and operation of pgpoolAdmin, you should be able to access the PHP information page and the pgpoolAdmin login screen at the following URLs:

```

http://wipds1.sas.com/pgpoolAdmin/install/phpinfo.php
http://wipds1.sas.com/pgpoolAdmin/login.php

```

## STEP 6: CONFIGURE THE MIDDLE TIER TO USE THE SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER HIGH-AVAILABILITY CLUSTER

Using the SAS Deployment Wizard, configure the SAS Web Application Server to use the pgpool-II data port on the primary high-availability node (for example, port 9432 on wipds1.sas.com). Since the SAS Metadata Server has already been configured, do not configure it again during middle-tier configuration (see Figure 5). However, the SAS Web Infrastructure Platform Data Server configuration needs to be run again in order to configure the SAS Web Infrastructure Platform with the high-availability data server cluster (see Figure 6).

Figure 5

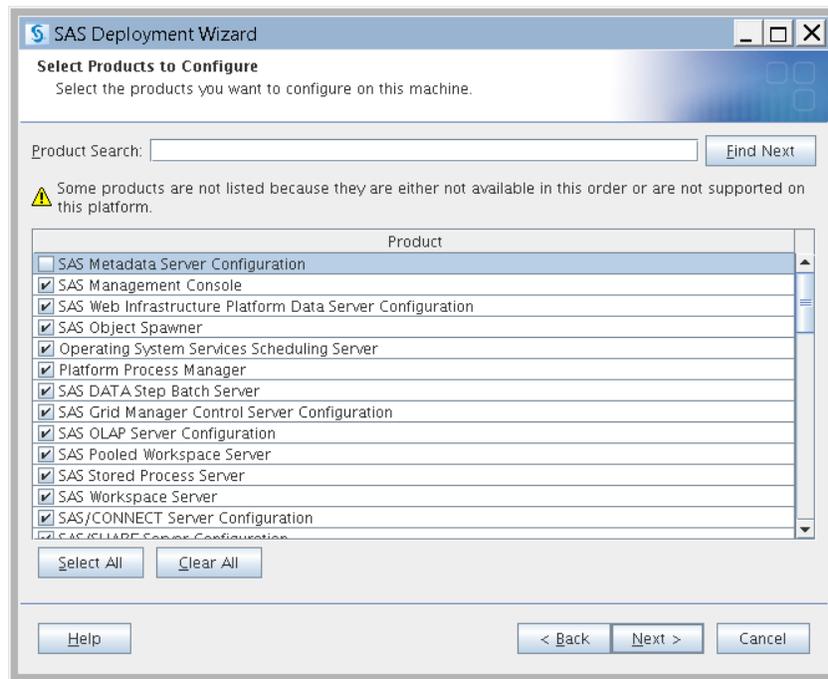
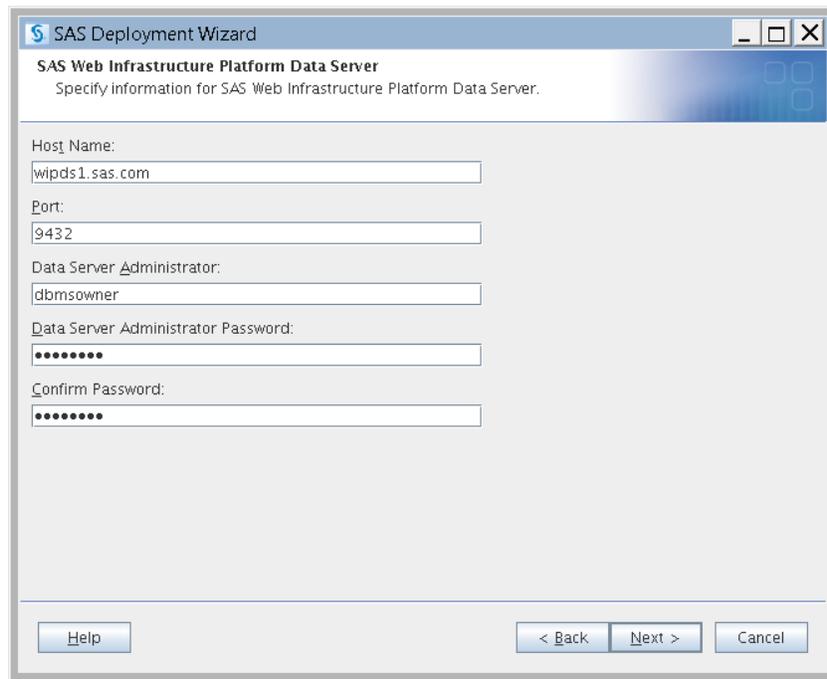


Figure 6

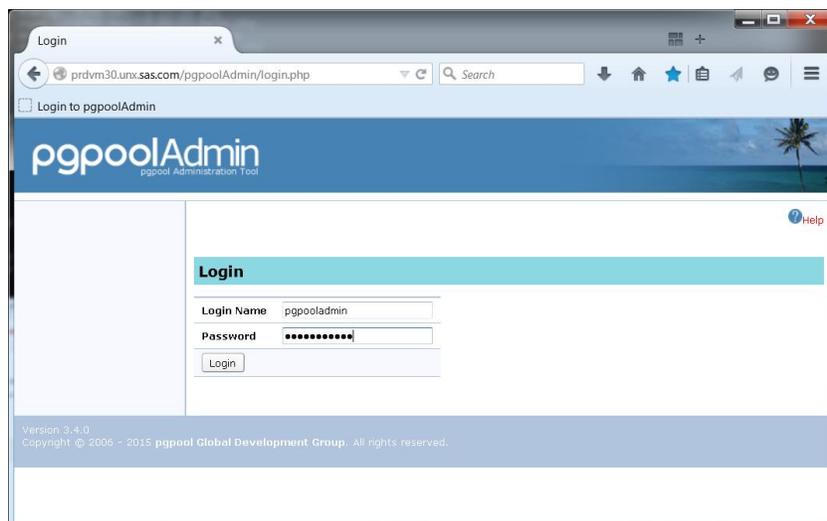


## STEP 7: MAINTAIN THE SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER HIGH-AVAILABILITY CLUSTER USING PGPOOLADMIN

### LOG IN TO PGPOOLADMIN

Use pgpoolAdmin (Figure 7) to manage the SAS Web Infrastructure Platform Data Server high-availability cluster.

Figure 7



Using this tool, you can do the following:

1. Manage WIP Data Server instance status (stop, restart, reload configuration, remove from cluster, and add new server), as shown in Figure 8.

Figure 8

**Backend info (PostgreSQL)**

Summary Process Info. **Node Info.**

**Node Info.**

	IP Address	Port	Status		Weight						
node 0	wipds1	9432	Up. Connected. Running as primary server	postgres: Up	0.500	Disconnect	Stop	Restart	Reload	Remove	
node 1	wipds2	9432	Up. Connected. Running as standby server	postgres: Up	0.500	Disconnect	Promote	Stop	Restart	Reload	Remove

Add a new backend node

2. Manage the pgpool-II instance status (stop, restart, and reload), as shown in Figure 9.

Figure 9

**pgpool**

Stop pgpool Restart pgpool Reload

3. Manually trigger failover and failback (Figure 10).

Figure 10

**Failover and Failback**

Parameter	Value
Command to execute when detaching a node failover_command (string)	%D %m %M %H %P %r %R
Command to execute when attaching a node failback_command (string)	/install/keyoun/pgpletc/failt
Fail over when socket communication error to backend occurs fail_over_on_backend_error	<input checked="" type="checkbox"/>

**Online Recovery**

Parameter	Value
PostgreSQL user name for online recovery recovery_user (string)	keyoun
PostgreSQL user password for online recovery recovery_password (string)	
Recovery command to be excuted on the first stage recovery_1st_stage_command (string)	basebackup.sh
Recovery command to be excuted on the second stage recovery_2nd_stage_command (string)	
Number of seconds to wait for online recovery to complete. 0 means no wait recovery_timeout (integer)	90

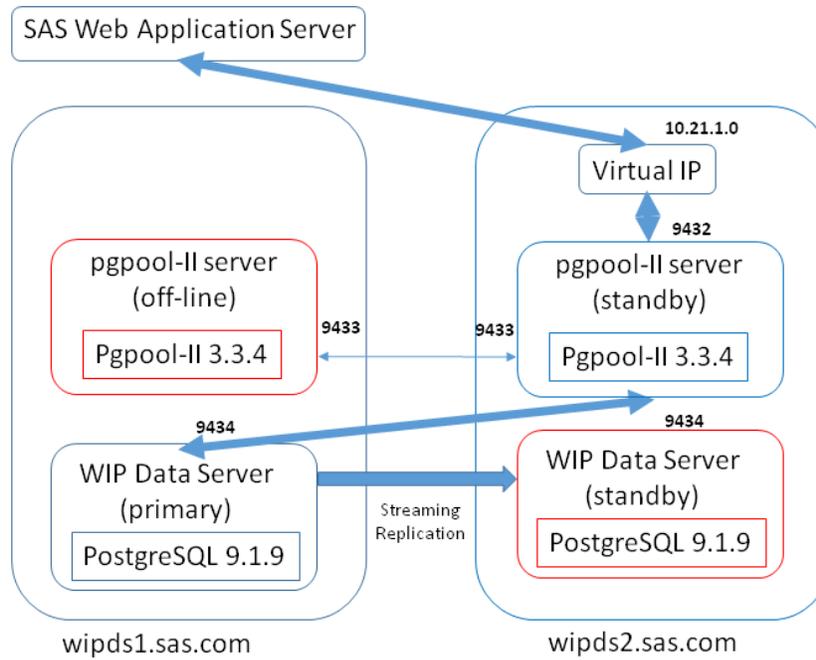
4. Change pgpool-II settings.

## PGPOOL-II SERVER FAILOVER SCENARIO

If the primary pgpool-II server fails (Figure 8), the following occurs:

1. The standby pgpool-II server is promoted to primary status.
2. The new primary pgpool-II server connects to the primary SAS Web Infrastructure Platform Data Server, which is still on wipds1.sas.com.
3. The virtual IP address (10.21.1.0) is reestablished on wipds2.sas.com.

Figure 8

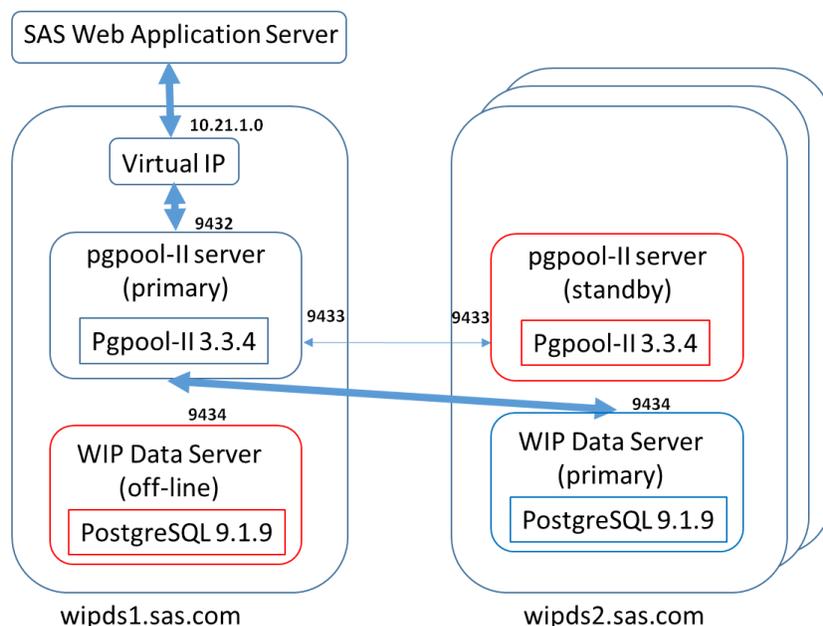


## SAS WEB INFRASTRUCTURE PLATFORM DATA SERVER FAILOVER SCENARIO

If the primary SAS Web Infrastructure Platform Data Server fails (Figure 9), the following occurs:

1. The standby SAS Web Infrastructure Platform Data Server is promoted to primary status.
2. pgpool-II connects to the new primary SAS Web Infrastructure Platform Data Server on wipds2.sas.com.
3. Any other standby SAS Web Infrastructure Platform Data Server is reset to receive replicated data from the new primary server.

Figure 9



## CONCLUSION

A robust high-availability solution for the SAS Web Infrastructure Platform Data Server can ensure that the middle tier's SAS Web Infrastructure Platform data will always be accessible. This is critical in maintaining a smoothly running SAS Web Application Server and SAS Web Server for your business. This software-based solution should be coupled with hardware-based solutions, such as ECC RAM and redundant disk storage, to achieve a superior high-availability SAS Web Server experience.

## REFERENCES

SAS Institute Inc. (2014). "Configuring the Web Infrastructure Platform Data Server for High Availability." Available at <http://support.sas.com/resources/papers/configure-wip-4-availability.pdf>.

PgPool Global Development Group. 2014. "pgpool-II." Accessed February 9, 2015. <http://www.pgpool.net/docs/latest/pgpool-en.html/>.

PgPool Global Development Group. "pgpool-II Tutorial [watchdog in master-slave mode]." Accessed February 9, 2015. [http://www.pgpool.net/pgpool-web/contrib\\_docs/watchdog\\_master\\_slave\\_3.3/en.html](http://www.pgpool.net/pgpool-web/contrib_docs/watchdog_master_slave_3.3/en.html).

SAS® 9.4 Intelligence Platform Administration Documentation available at <http://support.sas.com/documentation/onlinedoc/intellplatform/tabs/admin94.html>, in particular these references:

SAS Institute Inc. 2014. *SAS® 9.4 Intelligence Platform: Middle-Tier Administration Guide, Second Edition*. Cary, NC: SAS Institute Inc.

SAS Institute Inc. 2014. *SAS® 9.4 Intelligence Platform: Web Application Administration Guide, Third Edition*. Cary, NC: SAS Institute Inc.

SAS® 9.4 Intelligence Platform Installation, Configuration, and Migration Documentation available at <http://support.sas.com/documentation/onlinedoc/intellplatform/tabs/install94.html>, in particular this reference:

SAS Institute Inc. 2014. *SAS® 9.4 Intelligence Platform: Installation and Configuration Guide*. Cary, NC: SAS Institute Inc.

## ACKNOWLEDGMENTS

Sincere thanks are given to Salman Maher and Mark Kirk for reviewing this paper.

## RECOMMENDED READING

A recommended SAS Global Forum paper on the following topic can be found at <http://support.sas.com/resources/papers/proceedings14/>:

*Best Practices for Implementing High Availability for SAS® 9.4*

## CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Ken Young  
100 SAS Campus Drive  
Cary, NC 27513  
SAS Institute Inc.  
(919) 531-9319  
[Ken.Young@sas.com](mailto:Ken.Young@sas.com)  
[www.sas.com](http://www.sas.com)

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.