SSas. | THE POWER TO KNOW.

Technical Paper

Configuring the SAS® Web Infrastructure Platform Data Server for High Availability on Windows

---

Table of Contents

Introduction

Overview

This document describes how to configure the SAS middle tier for high availability of the SAS SharedServices database on Windows. The SharedServices database is one of the databases installed on and managed by the SAS Web Infrastructure Platform (WIP) Data Server. If you want to configure the SAS Web Infrastructure Platform Data Server on UNIX, see "**Managing SAS® Web Infrastructure Platform Data Server High-Availability Clusters on UNIX"**.

This document describes the configuration steps for an out-of-the-box or new level configuration only. You cannot update an existing SAS Web Infrastructure Platform Data Server configuration. Configuring for high availability involves configuring two SAS Web Infrastructure Platform Data Servers in a master/slave relationship, both based on PostgreSQL 9.1.9. One is the primary data server, and one is a standby data server that contains a replica of the data in the SharedServices database on the primary data server.

The configuration enables the primary data server to replicate data to the standby data server that is configured to take over as primary server if a failure occurs. High availability is provided for the SAS Web Infrastructure Platform Data Server's SharedServices database. However, others in the middle tier (such as vdbadm and transportsvcs_db) can be configured for high availability in a similar fashion. The EVManager and Administration databases are hosted on a third stand-alone SAS Web Infrastructure Platform Data Server, which is configured by the SAS Deployment Wizard and is separate from the SAS Web Infrastructure Platform high availability servers.

WIP Data Server	Configuration	Hosted Databases
WIP Data Server 1	Set up for replication as primary.	SharedServices (primary read/write copy).
WIP Data Server 2	Set up for replication as standby (receives all data from primary; essentially, this is a mirror of the primary server).	SharedServices (a read-only copy that's replicated from server #1 in real time).
WIP Data Server 3	Stand-alone (not connected to the primary and standby WIP data servers).	EVManager, Administration (possibly others).

The databases on the independent data server (#3) are described in *SAS® 9.4 Intelligence Platform: Middle-Tier Administration Guide***.**

The configuration process involves these three high-level steps:

1. Configure the SAS Web Infrastructure Platform Data Servers for high availability, with a primary server replicating data to a standby server.

2. Configure the SAS High-Availability JDBC shim to allow the middle tier to switch SAS Web Infrastructure Platform Data Servers during failover.

3.   Configure SAS Environment Manager to monitor the primary data server and promote the standby server when the primary fails.

To maximize high availability, a three-machine setup is recommended with one SAS Web Infrastructure Platform Data Server per machine. All of the steps in this document are written for a three-machine setup. However, if your resources are constrained, you can set up the three servers on fewer machines. If you do so, you must change the port numbers to be distinct on each machine. For example, if SAS Web Infrastructure Platform Data Servers 1 and 3 are located on the same machine, then port numbers must be different for each server, for example, 9432 and 9433.

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 SharedServices database. Make regular backups of the SharedServices database 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 customized to be used with SAS.

This paper focuses on how to reconfigure the SAS Web Infrastructure Platform Data Server, which includes the SharedServices database. This default SAS Web Infrastructure Platform 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. The SAS Web Infrastructure Platform Data Server that hosts the SAS Environment Manager database and other SAS Web Infrastructure Platform databases is not shown.

The reconfigured middle tier provides high availability for the SharedServices database by using the native PostgreSQL JDBC driver and the SAS High-Availability JDBC shim, also shown.
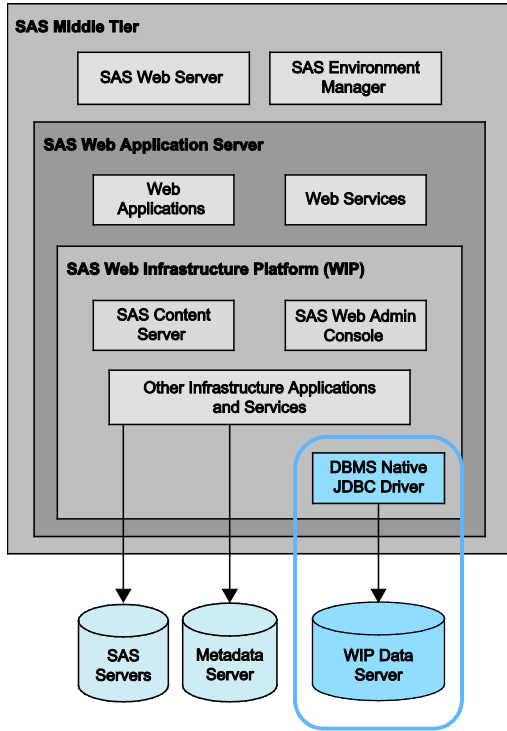
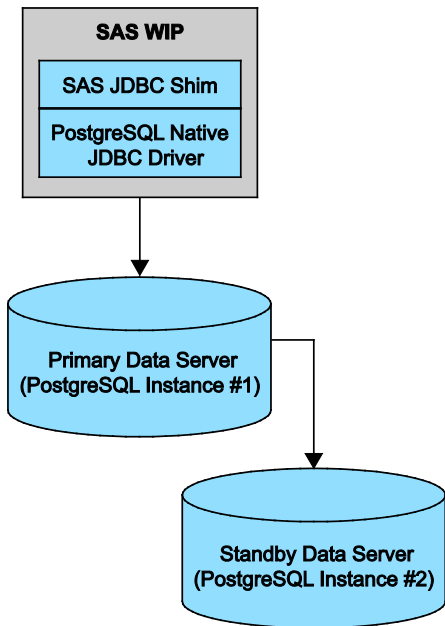*Figure 1. Default Configuration for SAS Web Infrastructure Platform Data Server*

Figure 2. Resulting Configuration with a Primary Data Server and a Second Standby Data Server

Install and Configure the Primary and Standby WIP Data Servers

Configuring the primary and standby data servers enables the SharedServices database that is used by the SAS Web Infrastructure Platform to be installed on these SAS Web Infrastructure Platform Data Server instances. As part of this process, you also deploy a separate, stand-alone SAS Web Infrastructure Platform Data Server to host and manage the SAS Environment Manager database (EVManager database). This deployment enables you to automate failover from the primary to standby SAS Web Infrastructure Platform Data Server via SAS Environment Manager. This EVManager database must be hosted on a completely separate SAS Web Infrastructure Platform Data Server instance.

CAUTION:
In order for automated failover to work, you must not install the EVManager database on the primary or standby SAS Web Infrastructure Platform Data Servers that you configure for replication. The EVManager database should reside on a separate data server that is stand-alone and not connected to the primary or standby servers.

These instructions assume that you have not yet deployed SAS and that you are preparing **SAS Web Infrastructure Platform Data Servers** that you will use to configure your main SAS deployment using the SAS Deployment Wizard. Using these instructions, you will use the SAS Deployment Wizard to install the SAS Web Infrastructure Platform Data Server, and then you will manually configure SAS Web Infrastructure Platform Data Servers needed for failover. You'll need access to the SAS Software Depot in order to run the installation.

There is a set of scripts available from SAS Technical Support that automate the configuration of the primary and standby servers for failover. These scripts can be run after steps 1–3 of the "Initialize and Configure the Primary Data Server" section are complete.

Initialize and Configure the Primary Data Server

To initialize and configure the primary data server, follow the steps below as the user that you will use to run the SAS deployment later in this process.

In these instructions, 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). Below are the example values used in the instructions. Substitute your site's values for these values.

Host value	Example Values Used in These Instructions
Host names	wipds1.sas.com (primary), wipds2.sas.com (standby)
IP Address	10.21.1.1 (primary), 10.21.1.2 (standby)
WIP data server port	9432
SASHome	C:\Program Files\SASHome
WIP data server data and log locations	C:\SAS\WIPDS

1. Start by identifying two machines to use—one for the primary SAS Web Infrastructure Platform Data Server instance and one for the standby data server instance.

2. Connect (via SSH, Remote Desktop, or whatever is appropriate for your environment) to the machine that will host the primary SAS Web Infrastructure Platform Data Server. In these instructions, this machine is named wipds1.sas.com.

3. Run the SAS Deployment Wizard.

   a. Choose your language.
   b. Choose **Install SAS Software**. Use the default path for SASHome. Or, choose another path, and click **Next**.
   c. Choose **Install Additional Software**.
   d. Scroll down in the products list and select **SAS Web Infrastructure Platform Data Server**.
   e. Click **Next** several times.
   f. Click **Start** to kick off the installation. Wait for the installation to finish.

4. Open a command prompt.

5. Add the path to the SAS Web Infrastructure Platform Data Server binaries to your PATH environment variable for this command prompt as follows, substituting your actual location for the example SASHome location:

   ```
   set PATH=C:\Program Files\SASHome\SASWebInfrastructurePlatformDataServer\9.4\bin;
   %PATH%
   ```

6. 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:

   ```
   C:\SAS\WIPDS
   ```

   You can use the `dir` directory command to see the amount of free space on the current drive.

   Change the directory (`CD`) into the directory you just created.

7. Create a log subdirectory within the parent directory you just created:

   ```
   mkdir log
   ```

8. Run the initialize database PostgreSQL command, `initdb`:

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

9. The `initdb` command creates a data directory where files containing the data within SAS Web Infrastructure Platform Data Server databases are stored. The `initdb` command also creates a DBMS administrative user. This name, which is `dbmsowner` by default, corresponds to the SAS Web Infrastructure Platform database administrative user that you set up during the SAS Deployment Wizard configuration. Specify the password (also set in the SAS Deployment Wizard for the administrative user) when prompted. Record the administrative user's user ID and password for future use.

10. Change the directory (`CD`) to the data directory.

11. Edit the remote access permissions configuration file pg_hba.conf to add the following lines. Specify the IP address in standard dotted decimal notation with a Classless Inter-Domain Routing (CIDR) mask length. Refer to your PostgreSQL documentation for information about setting the CIDR mask length.

IPv4:

```
host      replication   all    10.21.1.2/32 trust
host    all       all     all         md5
```

IPv6:

```
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 the example IP addresses. Substitute your actual IP values.

12. 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.

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

    ```
    copy postgresql.conf postgresql.orig
    ```

14. Edit the postgresql.conf configuration file to set the configuration values as shown below. Some settings contain recommended values for general SAS Web Infrastructure Platform database use. These values can be increased for deployments that include multiple web servers. Update the log_directory value to match the log subdirectory that you created in Step 8, as follows:

    ```
    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:
        #    'C:\\SAS\\WIPDS\\log'
    log_truncate_on_rotation = off
    log_rotation_age = 1d
    log_min_messages = info
    ```

    Note: Replace single backslashes in paths within the postgresql.conf file with two backslashes. Use double backslashes in the `log_directory` path as shown in the example above.

15. Register the new SAS Web Infrastructure Platform Data Server instance as a Windows service using the `pg_ctl` command:

    ```
    C:\> pg_ctl register -N "SAS HA WIP Data Server" -D "<path_to_data_dir>"
    -w -S a -o "-p <port>"
    ```

16. For example:

    ```
    C:\> pg_ctl register -N "SAS HA WIP Data Server" -D "C:\SAS\WIPDS\data" -
    w -S a -o "-p 9432"
    ```

17. Start the new Windows service you just created:

```
C:\> net start "SAS HA WIP Data Server"
```

18. Check the log file in the log directory to ensure that the server is running without problems. The contents of the file should be similar to this:

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

Initialize and Configure the Standby Data Server

To initialize the standby data server, follow these steps:

1. Connect to the machine that will run the standby SAS Web Infrastructure Platform Data Server, for example, via SSH or Remote Desktop.

2. Follow the same set of steps to install and configure a SAS Web Infrastructure Platform Data Server on the standby machine as performed on the primary machine. Once the standby SAS Web Infrastructure Platform Data Server is set up the same way as the primary via the instructions above, continue to the next step.

3. Stop the SAS Web Infrastructure Platform Data Server on the standby machine.

```
net stop "SAS HA WIP Data Server"
```

4. Copy the data directory from the primary server to the standby server.

   a. On the primary machine (with host name wipds1 in this example), specify the administrative user and put the primary SAS Web Infrastructure Platform Data Server into base backup mode:

   ```
   $psql -c "SELECT pg_start_backup('stage backup', true)" -p 9432 -U
   dbmsowner -d postgres
   ```

   b. Copy the primary server's data directory to the standby machine, host name wipds2.
   You can use Windows Explorer to copy the data directory from the primary to the standby machine. For example, follow these steps on the machine that hosts the standby server:

      i. Using Windows Explorer, open \\wipds1\c$\SAS\WIPDS.
      ii. Right-click the data directory and click **Copy**.
      iii. Navigate to the corresponding directory on the standby server, `C:\SAS\WIPDS`.
      iv. Right-click within the main Explorer pane and click **Paste**.

   c. Take the primary SAS Web Infrastructure Platform Data Server out of base backup mode. On the primary machine, specify the administrative user and run:

   ```
   $ psql -c "SELECT pg_stop_backup()" -p 9432 -U dbmsowner -d postgres
   ```

5. Edit the pg_hba.conf file in the data directory. Change the line that allows replication to use the IP address of the primary SAS Web Infrastructure Platform Data Server.

6. This will not be used initially because no SAS Web Infrastructure Platform Data Server will be replicating data from the standby server. However, it makes it easier to recover if the primary data server fails in the future.

7. Create a recovery.conf file in the data directory on the standby server.

8. This file causes the standby server to connect to the primary for streaming replication. That is, the standby server streams data from the primary in near-real time as it's added or updated. Add the following content to the recovery.conf file. Specify the administrative user. Be sure to update the IP address to the actual IP address of the machine that hosts your primary SAS Web Infrastructure Platform Data Server.

9. **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=9432 user=dbmsowner'
    # Specifies a trigger file whose presence should cause streaming replication
    to
    # end (i.e., failover).
    trigger_file = 'C:\\path\\to\\trigger.tmp'
    ```

10. **Note:** The backslashes (\) in the path to the trigger file must be escaped, which requires using double backslashes in the path.

11. Start the standby SAS Web Infrastructure Platform Data Server on host wipds2.sas.com.

12. If configured correctly, it starts streaming replication. Use the same command to start the SAS Web Infrastructure Platform Data Server that you used when you first installed it. (See the section above on configuring the primary SAS Web Infrastructure Platform Data Server for a reminder.)

13. On the standby server, check the SAS Web Infrastructure Platform Data Server log file to ensure that the server is operating correctly, that it hasn't raised errors, and that streaming replication is working.

14. At this point, the SAS Web Infrastructure Platform Data Server is now configured for replication. Look in the log directory for the most recent file. The log file content should look something like this:

    ```
    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
    ```

Prepare for and Execute the SAS Deployment Wizard

Before deploying SAS and using the SAS Deployment Wizard to populate the SharedServices database with data, you need to manually create the database. Follow these steps:

1. On the primary SAS Web Infrastructure Platform Data Server, run psql in interactive mode:

    ```
    $ psql -p 9432 -U <admin_user> -d postgres
    ```

2. Enter the following commands.

3. The commands create the SharedServices database, switch the psql connection to the new SharedServices database, and add the required lo (Large Object) extension to the database:

```
create database "SharedServices";
\c SharedServices
create extension lo;
\q
```

4. Validate the creation of the SharedServices database and large object extension by specifying the administrative user and running:

```
$ psql –d SharedServices –p 9432 –U dbmsowner –c "select * from pg_extension where extname='lo'"
```

5. You should see the following output:

```
extname | extowner | extnamespace | extrelocatable | extversion | extconfig | extcondition
---------+----------+--------------+----------------+------------+-----------+--------------
lo      |       10 |         2200 | t              | 1.0        |           |
(1 row)
```

You can now use the SAS Deployment Wizard to install and configure your SAS deployment, including the SAS Web Infrastructure Platform Database and SAS Environment Manager.

Let the SAS Deployment Wizard populate the SharedServices database with the appropriate data. You are asked to specify the connection information for the database. Specify the information for the primary SAS Web Infrastructure Platform Data Server instance, wipds1. When the SAS Deployment Wizard populates the SharedServices database on the primary instance with content, the data and subsequent updates are replicated automatically from the primary to the standby instance. For full instructions, refer to *SAS® 9.4 Intelligence Platform: Installation and Configuration Guide.*

To populate the SharedServices database on your replicated SAS Web Infrastructure Platform Data Servers and install SAS Environment Manager so that it can be used to monitor the SAS Web Infrastructure Platform Data Servers during SAS deployment, follow these guidelines:

1. On your middle-tier machine, separate the deployment into install and configure steps.

2. That is, run the SAS Deployment Wizard once to install the software, and again to configure it. This is needed because the SAS Web Infrastructure Platform installs the SAS Web Infrastructure Platform Data Server JDBC driver, which is needed during configuration of SAS Web Infrastructure Platform. Installing and then configuring SAS Web Infrastructure Platform as separate steps ensures that the JDBC driver is available on the file system at the time it's needed.

3. Choose to install and configure SAS Environment Manager.

4. Deploy one SAS Environment Manager server (typically installed on your middle-tier machine). Also, deploy a SAS Environment Manager agent on each SAS Web Infrastructure Platform Data Server machine—one on the primary server and one on the standby server.

5. In the Select Configuration Prompting Level window, select **Custom**.

6. In the SAS Web Infrastructure Platform Database: Data Server window, deselect (uncheck) the **Use SAS Web Infrastructure Platform Data Server** check box.

7. Selecting the check box causes the SAS Deployment Wizard to configure SAS Web Infrastructure

Platform to use a stand-alone SAS Web Infrastructure Platform Data Server.

8. In the SAS Web Infrastructure Platform: Database Type window, choose PostgreSQL as the database type.

9. In the JDBC Properties: Database Connections Properties window, fill in the host name and port of your SAS Web Infrastructure Platform Data Server.

10. Also, provide the SAS Deployment Wizard with the path to a copy of the required SAS Web Infrastructure Platform Data Server JDBC JAR file. During the installation step, the SAS Deployment Wizard installed the required JDBC driver in `<SASHome>\SASWebInfrastructureDataBaseJDBCDrivers\9.4\Driver`. Provide this path to the SAS Deployment Wizard on this screen. Leave the **Use Catalog** and **Use Schema Pattern** boxes unchecked, which they are by default.

11. In the second SAS Web Infrastructure Platform Database: JDBC Properties window, enter the credentials for the SAS Web Infrastructure Platform Data Server administrator that you created when you set up the server.

12. These credentials are used to connect to SAS Web Infrastructure Platform Data Server.

13. Proceed with the rest of the deployment process as normal.

14. Install and configure a SAS Environment Manager Agent on both the primary and standby SAS Web Infrastructure Platform Data Server machines, if not already done.

During configuration, specify the connection information for the SAS Environment Manager server that you installed earlier.

CAUTION:
Never run two or more primary SAS Web Infrastructure Platform Data Servers simultaneously. Data can be lost or corrupted and can be difficult to retrieve without a full backup of the SharedServices database. As a best practice, make regular backups of the SharedServices database.

Configure the SAS Middle Tier for WIP Data Server Failover

You must configure the SAS middle tier to use the SAS 9.4 High-Availability JDBC shim in order for failover to work. The JDBC shim redirects queries and connections to the SharedServices database from a failed data server to a running data server.

In conjunction with the SAS Environment Manager plug-in for SAS Web Infrastructure High-Availability Server, the SAS High-Availability JDBC shim allows the SAS middle tier to support fully automated SAS Web Infrastructure Platform Data Server failover for the SharedServices database. Both databases must contain the content populated by the SAS Deployment Wizard during the configuration process.

Configuration

To configure a currently running SAS Web Infrastructure Platform and SAS Web Application Server to use the SAS High-Availability JDBC shim, perform the following steps:

1. Stop the SAS Web Application Server using the Services applet in the Control Panel.

2. Open a command prompt and copy the High-Availability JDBC shim files (hapostgresql.jar and ha-jdbc-properties.xml) from the **WEBINFDJDBC** install directory.

3. If you installed SAS in C:\Program Files\SASHome, copy the files in this directory:

4. C:\Program Files\SASHome\SASWebInfrastructureDataBaseJDBCDrivers\9.4\Driver to the SASServer servlet container library within your config dir,
   <config_dir>\<LevN>\Web\WebAppServer\SASServer1_1\lib (for example,
   C:\SAS\Config\Lev1\Web\WebAppServer\SASServer1_1\lib).

5. Change directories to the lib directory into which you just copied the two files above.

6. Modify the High-Availability JDBC shim configuration file ha-jdbc-properties.xml to define the Active Servers.

7. **Note:** Make sure to modify the copy of the file within your SAS config directory, not the source copy in SASHome.

8. Be sure to set the ActiveServerCount property to 2 to enable high availability. A zero value for the ActiveServerCount or a missing configuration file forces the High-Availability JDBC shim into pass-through mode to the native PostgreSQL JDBC driver.

9. Change directories to ../conf -- that is, up one level, then into the conf subdirectory.

10. Copy server.xml to server.xml.orig to back up the file before editing it.

11. Open the server.xml file in your preferred editor.

CAUTION:
In the following step, modify only the lines for SharedServices resource files. You must not make changes to any other databases such as the Administration database.

1. Search for "SharedServices" and modify the SharedServices resource to reference the SAS High-Availability JDBC shim driver class and JDBC type. This requires two changes in the resource for SharedServices:

   a. For the driverClassName attribute, change org.postgresql.Driver to com.sas.postgres.ha.Driver.
   b. For the URL attribute, change `postgresql` to `hapostgresql`.

2. See the code highlighted in bold italic.

```
Old values for driverClassName= and url=:
<Resource auth="Container" driverClassName="org.postgresql.Driver"
factory="com.sas.vfabrictcsvr.atomikos.BeanFactory" maxPoolSize="100"
minPoolSize="10" name="sas/jdbc/SharedServices"
password="${pw.sas.jdbc.SharedServices}" testQuery="select 1"
type="com.atomikos.jdbc.nonxa.AtomikosNonXADataSourceBean"
uniqueResourceName="sas/jdbc/SharedServices"
```

```
url="jdbc:postgresql://server.example.com:9432/SharedServices"
user="postgres"/>
```

3. New values for driverClassName= and url=:

```
<Resource auth="Container" driverClassName="com.sas.postgres.ha.Driver"
factory="com.sas.vfabrictcsvr.atomikos.BeanFactory" maxPoolSize="100"
minPoolSize="10" name="sas/jdbc/SharedServices"
password="${pw.sas.jdbc.SharedServices}" testQuery="select 1"
type="com.atomikos.jdbc.nonxa.AtomikosNonXADataSourceBean"
uniqueResourceName="sas/jdbc/SharedServices"
url="jdbc:hapostgresql://server.example.com:9432/SharedServices"
user="postgres"/>
```

4. Ensure that you changed only the line for the SharedServices database and not for another database.

5. Save and close the server.xml file. Open the log4j.xml file, which is in the same directory into which you copied the driver (SASServer1_1/lib).

6. Add the class logging statements for the High-Availability JDBC shim (below) to the log4j.xml file. Some sections of log4j.xml are commented out. Be sure to add the statements to a part of the file that isn't commented out, so it'll take effect. Add this to log4j.xml:

```
<category additivity="false" name="com.sas.postgres.ha">
    <priority value="DEBUG"/>
    <appender-ref ref="CONSOLE"/>
    <appender-ref ref="FILE"/>
  </category>
```

7. Start the SAS Web Application Server and wait for this process to complete. The SAS Web Application Server has completely finished starting when it writes a statement like this to its server.log file:

```
2013-07-15 10:50:46,766 INFO (main) [org.apache.catalina.startup.Catalina]
Server startup in 621182 ms
```

Verification

To verify that the SAS Web Application Server started without error, check the servlet container log file (server.log in the `<Config_dir>\<LevN>\Web\WebAppServer\SASServer1_1\logs` directory).

It might take some time (perhaps 30 minutes) for the SAS Web Application Server to complete its start up. You see High-Availability JDBC shim log messages in this file. You should see log statements similar to these:

```
INFO [main] - Driver#connect SAS 9.4 High Availability Driver for PostgreSQL JDBC4
(build 0001)
DEBUG [main] - Driver#connect High availability properties were loaded from 'ha-
jdbc-properties.xml'
DEBUG [main] - Driver#findActiveServer Attempting to set active server #1 ...
DEBUG [main] - Driver#findActiveServer Client URL:
jdbc:postgresql://wipds1.sas.com:9432/dbitest?user=dbitest&password=xxxxxx
DEBUG [main] - Driver#findActiveServer Active Server URL:
jdbc:postgresql://wipds1.sas.com:9432/dbitest?user=dbitest&password=xxxxxx
DEBUG [main] - Driver#findActiveServer Active Server: Active server #1
(wipds1.sas.com) was set.
DEBUG [main] - Driver#connect Connected in 45 milliseconds,
url=jdbc:postgresql://wipds1.sas.com:9432/dbitest?user=dbitest&password=xxxxxx
```

CAUTION:

Never run two or more primary SAS Web Infrastructure Platform data servers simultaneously. Data can be lost or corrupted. It can be difficult to retrieve without a full backup of the SharedServices database. As a best practice, make regular backups of the SharedServices database.

Configure SAS Environment Manager to Manage WIP Data Server

The SAS Environment Manager plug-in for Web Infrastructure High-Availability Server (HA plug-in) must be configured, which involves these high-level steps:

1. Create a PostgreSQL password file.

2. Configure the plug-in to monitor both the primary SAS Web Infrastructure Platform Data Server and the standby SAS Web Infrastructure Platform Data Server.

3. Configure the plug-in to promote the standby server if the primary server fails.

4. These steps are described in detail in the following instructions.

Create the PostgreSQL Password File

You must create a PostgreSQL password file that provides the credentials needed by the SAS Environment Manager agent in order to log on to a SAS Web Infrastructure Platform Data Server. This password file is automatically used by PostgreSQL client components, including the psql command-line tool, which is used by the SAS Environment Manager plug-in for High Availability to monitor SAS Web Infrastructure Platform Data Servers.

Since the EV agent service on Windows is typically run as the LocalSystem account, the password file must be created in a special path contained within the Windows System32 directory.

To create the password file, follow the steps below. Start by performing these steps on the primary SAS Web Infrastructure Platform Data Server machine, and then follow the same procedure for the standby SAS Web Infrastructure Platform Data Server machine.

1. Before using credentials from the password file, the PostgreSQL client tools check for the environment variables PGPASSWORD and PGPASSFILE. Unless you plan to use alternative methods of password storage, ensure that these variables are not set.

2. Create the PostgreSQL password file in the appropriate place for your operating system.

3. If the EV agent service is configured to run as the default LocalSystem account, create the file pgpass.conf in this directory:

   `%SYSTEMROOT%\system32\config\systemprofile\AppData\Roaming\postgresql`

4. If this directory doesn't exist, create it before you create the pgpass.conf file.

5. The format of the information in the password file is:

   `hostname:port:database:username:password`

6. An asterisk (*) in any of the first three fields specifies that the PostgreSQL client tools are to allow the specified user and password credentials to be used for connections regardless of the field's value. For example, an asterisk in the port field specifies for the PostgreSQL client tools to use those credentials when connecting to SAS Web Infrastructure Platform Data Server on any port.

7. Choose a SAS Web Infrastructure Platform Data Server user that is appropriate for the EV agent to use to access the SAS Web Infrastructure Platform Data Server. An example would be the administrative user that you created earlier when you set up WIP data server. No special permissions are needed within the SAS Web Infrastructure Platform Data Server other than the ability to connect to the Postgres system database.

8. Add the credentials for the user to the password file. For example, the file might look like the following:

    ```
    *:*:*:dbmsowner:mypassword
    ```

Perform the above steps on the standby SAS Web Infrastructure Platform Data Server machine as well.

Configure the SAS Environment Manager Plug-In for Web Infrastructure High Availability

The following sections provide steps to configure the SAS Environment Manager plug-in for Web Infrastructure High Availability.

Prepare to Configure the Plug-In

1. Verify that you are logged on to the SAS Environment Manager as sasadm@saspw using the password you set during deployment in order to configure the plug-in.

2. Select **Start ▶ All Programs ▶ Windows PowerShell version ▶ Windows PowerShell.**

3. On the PowerShell command line, query the execution policy by submitting `Get-ExecutionPolicy.`

4. If the policy is not set to `Unrestricted`, then submit `Set-ExecutionPolicy Unrestricted.`

5. Verify that the command did not produce any errors.

6. Type `Exit` at the PowerShell command prompt to end the session.

Verify That the Correct Plug-In Is Installed

1. Log in to SAS Environment Manager and click **Manage** in the toolbar.

2. Click **Plugin Manager** under Server Settings.

3. Verify that the emi-services plug-in is listed on the following window:

Configure the Monitoring Process

The following steps configure the monitoring process for both the primary server and the standby server.

1. Click Resources in the top toolbar.

2. Click the name of the machine that is running the primary instance of the SAS Web Infrastructure Platform Data Server. Then, pull down the **Tools** menu and select **New Platform Service**.

3. Give the new service a name, for example, `Postgres HA`. Optionally, give the service a description.

4. Set the new service type as **PGHA** and click **OK**.

5. In the next window, click **Configuration Properties.**

6. On the next window, you will enter the new SAS Web Infrastructure Platform Data Server information that is unique to your installation. Use the descriptions in the table below to complete steps 7 and 8.

Host Value	Description
pg_path	Specify the path to the WIP data server directory that contains the bin directory. Do not put bin on the end. For example: `C:\Program Files\SASHome\SASWebInfrastructurePlatformDataServer\9.4`
pg_user	Specify the WIP data server user ID that you want to use to monitor the WIP data server. This is case sensitive. This should be the same user that you added to the PostgreSQL password file earlier in this section.
g_port	Specify the port that the WIP data server is listening on, for example, 9432.
pg_log	Specify a location for the plug-in to write log files, or leave blank to take the default location.
pg_trig	Specify the exact path to the WIP data server trigger file. This path must match the path that was set as the trigger_file value in recovery.conf. If you are configuring the standby server, double-check the trigger file path that you specify. If the path is incorrect, the standby server is not promoted during failover, and the standby server does not take over from the failed primary. If you are configuring the primary server, enter the exact path to the trigger file you would use if the two servers swapped roles such that the standby became the new primary and this primary were recovered as a standby. This value is initially not used but can make recovering from a failure easier in the future.

7. After specifying the information for the primary server on the window shown below, click **OK**.

8. Repeat steps 2–7, specifying information specific to the standby server and click **OK**.

Check the Availability of Primary and Standby Servers

To check the availability of the primary server, follow these steps:

1. Click **Monitor** in the Resource Monitor window toolbar.
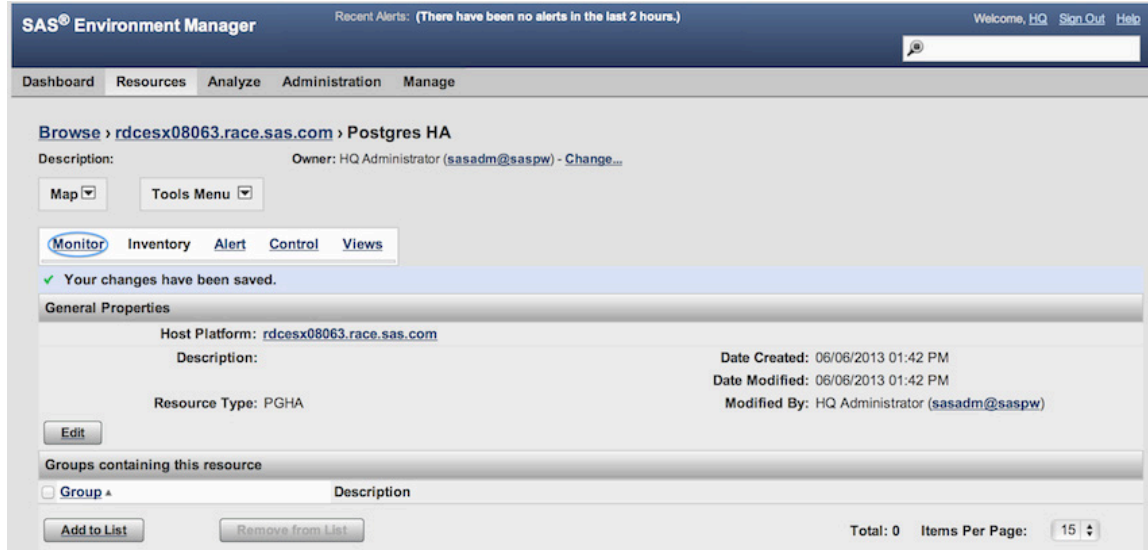
2. Check the availability bar to ensure that it is at least partially green, for example, as shown in the below. Green indicates that the primary SAS Web Infrastructure Platform Data Server instance is up and available. Repeat these steps to verify that the standby server is also available.
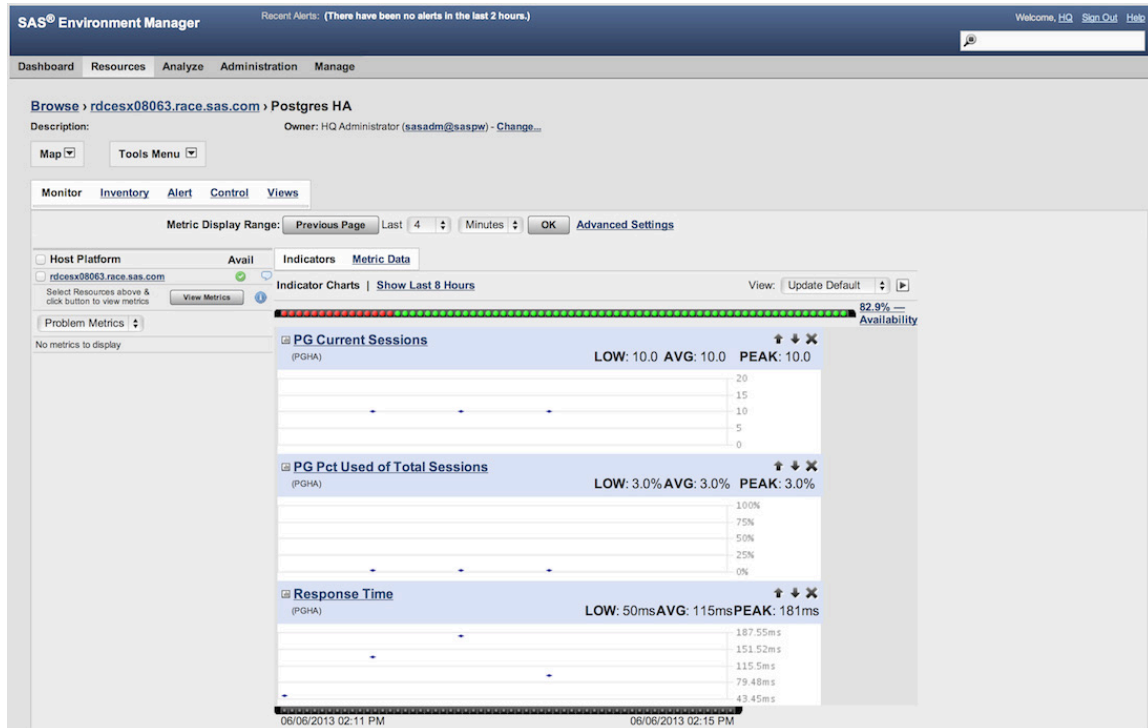
Configure Failover Alert

To configure SAS Environment Manager to generate an alert when the SAS Web Infrastructure Platform Data Server is unavailable and automatically promote the standby SAS Web Infrastructure Platform Data Server to primary, continue in the Resources Monitor window and follow these steps:

1. Click the name of the SAS Web Infrastructure Platform Data Server in the top browse list.

2. Click the PGHA service that you created earlier.

3. Click **Alert** in the toolbar.
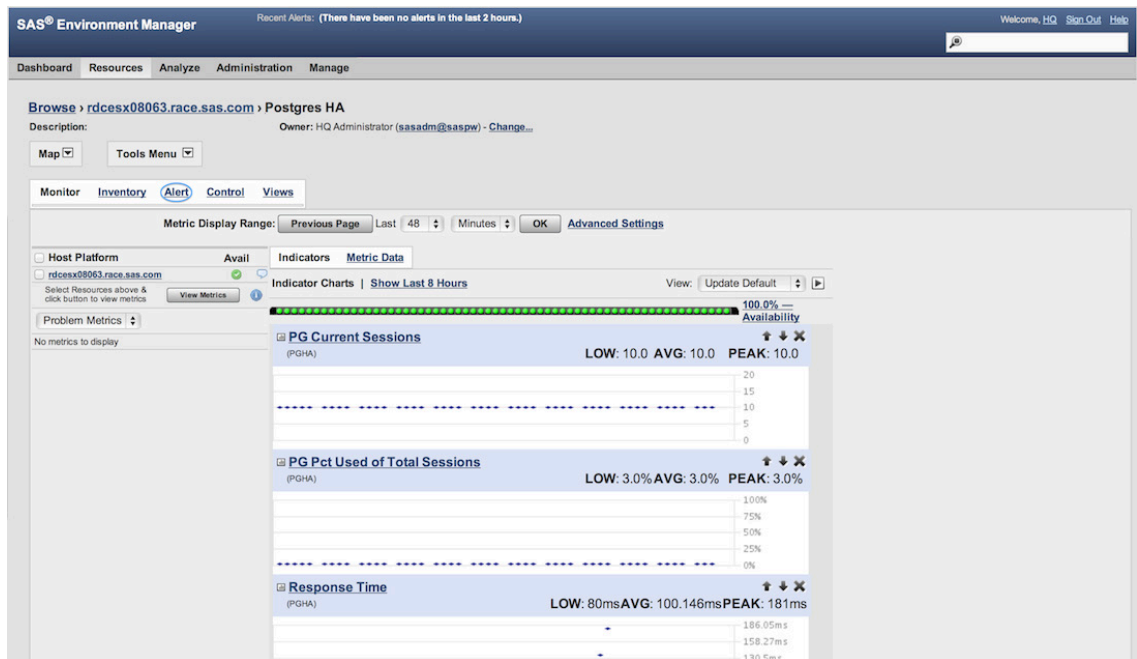
4. On the Alert Configuration window (not shown), click **Configure**.

5. Click the **New** button to bring up the next window.

6. On the Alert Definition window (shown below), configure your new alert.

   Note: Carefully follow the configuration in the window below. You are setting an alert for when the primary server's availability drops below 100%. You can use the percent sign (%) in the value field for the condition, as shown in the window.

   Optionally, you can enable the action filter Generate one alert and then disable alert definition until fixed.

   If this action filter is enabled when the primary server fails, the SAS Environment Manager raises one alert and makes one attempt to promote the standby server. Then, SAS Environment Manager stops until the problem is marked fixed. By default, SAS Environment Manager raises alerts every minute for as long as the primary server is down.

7. Click **OK** to bring up the next window.

8. In the Configure Actions window, click **Control Action**.

9. On the next window, click **Edit**, near the bottom of the screen to add a control action definition.

10. Provide **Control Action** details for your site. Under Resource Name, be sure to select the name of your *standby* SAS Web Infrastructure Platform Data Server. The control action promotes the standby server to primary, so this action must be taken on the standby server.

11. Click **OK**.

Post-Setup Notes

Following the steps above, you performed the following tasks:

- Configure the primary and standby SAS Web Infrastructure Platform Data Servers.

- Enable the SAS High-Availability JDBC shim.

- Configure monitoring and promotion for both the primary and standby servers.

The system is fully set up to handle the failure of the primary SAS Web Infrastructure Platform Data Server and to automatically fail over to the standby if this occurs.

With this high-availability configuration, your standby server maintains a constant duplicate of the data on the primary server. The standby server is read-only and is not usable by the SAS middle tier as a host for a live SharedServices database. However, when a failure of the primary server is detected, the SAS Environment Manager plug-in for Web

Infrastructure High-Availability Server creates a trigger file in the path indicated by the trigger_file setting that you configured in the recovery.conf file. When the standby server detects that this trigger file has been created, it promotes itself to become a primary server, which is read/write.

In many cases, administration of this new setup is the same as with a single node instance. For example, use the Services applet to stop or restart either the primary or standby SAS Web Infrastructure Platform Data Server. (Be careful stopping the primary server once high availability has been configured, because SAS Environment Manager will detect that it is down and automatically promote the standby server.) You need to follow different steps when promotion of the standby occurs or the nodes fall out of sync. See the next section.

CAUTION:
Never run two or more primary SAS Web Infrastructure Platform Data Servers simultaneously. Data can be lost or corrupted, and the data can be difficult to retrieve without a full backup of the SharedServices database. As a best practice, make regular backups of the SharedServices database.

Repairing Streaming Replication after a Failover

If the primary server becomes unavailable and a failover occurs, the system should continue running normally after a brief period of possible errors as the failover occurs. After a failover, you no longer have a high-availability system. That is, you have only a single SAS Web Infrastructure Platform Data Server—the original standby. This original data server is now the new primary SAS Web Infrastructure Platform Data Server. Now, you should bring up a second server, so the system can tolerate failures on the new primary server.

The easiest thing to do is to bring up the original, failed primary server as a standby server. That is, the easiest course of action after a failover is to let the new primary server continue to act as the primary and bring up the original primary as the new standby. This effectively swaps the original roles of the two servers. To do this, follow the same steps as when you originally configured the standby server: Put the primary into base backup mode, copy its data directory to the standby, take the primary out of base backup mode, and then start the standby server.

Follow the steps from the "Initialize and Configure the Standby Server" section of this document to bring up your original primary as the new standby and synchronize it with the new primary server. Make sure that the new standby machine has a recovery.conf file and that the trigger file is specified correctly, both in recovery.conf and in the configuration options for the PGHA plug-in.

In SAS Environment Manager, be sure to first verify that the trigger file path is specified correctly in the configuration options for the PGHA plug-in. Then, follow the steps in the in the "Configure Failover Alert" section.

CAUTION:
There is a potential for data loss or corruption if this feature is used improperly. If two or more primary SAS Web Infrastructure Platform Data Server are used simultaneously, data loss or corruption can occur. This problem can be very difficult to recover from without full backups of the SharedServices database. Be very careful to not allow two active primary SAS Web Infrastructure Platform Data Servers to be used. And, be sure to make regular backups of the SharedServices database. This feature is not a substitute for making regular backups.

Recovering a Failed Standby Server

If the standby server fails (and the primary is still functioning normally), use the same set of steps in the "Initialize and Configure the Standby Server" section of this document to resynchronize the standby server with the primary:

1. Put the primary server into base backup mode.

2. Copy the data directory to the standby server.

3. Take the primary out of base backup mode.

4. Start the standby server.

When you start the standby server, it begins streaming replication as normal. If you want to check the log files to make sure, see Step 7 in the "Initialize and Configure the Standby Server" section.