

SAS® Grid Manager – Troubleshooting

25Jul2007

INTRODUCTION

This document contains suggestions to help debug problems that may occur when running SAS Grid Manager using the Platform Suite for SAS. This document assumes you have an understanding of TCP/IP networking and an understanding of Platform LSF.

BACKGROUND

A SAS grid consists of grid clients working with SAS/CONNECT client sessions to create grid jobs that start remote SAS/CONNECT sessions on grid nodes. For Foundation SAS, the interactive or batch Foundation SAS process is both the grid client and the SAS/CONNECT client. For SAS applications such as DI Studio, the grid client is the DI Studio application and the SAS/CONNECT client is either the workspace server or stored process server that processes the SAS code created by the application.

COMMUNICATION

In order to process SAS code on remote SAS sessions, communication channels must be created between the remote SAS/CONNECT session and the client SAS/CONNECT session. To create the channel to the client SAS/CONNECT session, the remote SAS/CONNECT session uses the short hostname and port of the SAS/CONNECT client provided by the grid command script (sasgrid.bat). Once that channel has been made, the remote SAS/CONNECT session uses the channel to send its IP address and port to the client SAS/CONNECT session to allow it to create the channel back to the remote SAS/CONNECT session. Since both host names and IP addresses are passed back and forth, hostname to IP address mapping needs to be correct at both ends. For machines that have more than one network adapter, it is imperative that the default IP address returned for the hostname is accessible by the other machines in the grid.

CREDENTIALS

The remote SAS/CONNECT session is started using the user credentials of the user running the SAS/CONNECT client session. For Foundation SAS grid client, the user is the account which started the interactive or batch Foundation SAS process. For a workspace server, the user is determined by the grid client application such as DI Studio. For stored processes, the user is the account used to start the stored process server by the object spawner.

#1 - CHECK NETWORK SETUP

HOST ADDRESSIBILITY

A common problem at many customer sites is that machine names are mapped to 127.0.0.1 in the /etc/hosts file on one or more grid nodes. This will cause the signon connection to the grid node to either fail or hang because the SAS session being invoked on the grid node will not be able to determine the correct IP address of the machine on which it is running. A correct IP address must be returned to the client session in order to complete the connection. Check the /etc/hosts file on each of the grid nodes to ensure that it does not map the machine's name to the 127.0.0.1 address. For example, if the following line exists in /etc/hosts

```
127.0.0.1 myserver localhost.localdomain localhost
```

you would need to delete the name "myserver".

HOST CONNECTIVITY

Many grid problems are a direct result of improper network setup and the ability of one machine to know the correct address of all other machines in the grid. Here are some suggestions to test the network setup:

- Make sure each machine in the grid (including the grid control machine, grid nodes, and SAS/CONNECT clients) can 'ping' every other grid machine since SAS/CONNECT clients submit jobs to the grid that can run on any grid node.
- Run 'hostname' on each SAS/CONNECT client machine (i.e. machines that will be used to submit jobs). Make sure that when you ping the SAS/CONNECT client from all grid nodes, you use that host name without the domain suffix. Also, when you run ping from the SAS/CONNECT client and try to ping itself using the value returned from 'hostname', make sure the IP address returned is the same IP address seen from the grid nodes. On machines with multiple network adapters, this may not necessarily be the case.
- Repeat step #2 on each grid node and run ping from the SAS/CONNECT client. This ensures the IP address returned by the remote SAS session is accessible from the SAS/CONNECT client.

What do you do if a problem occurs? Either the DNS server needs to be corrected or entries need to be made in each machine's 'hosts' files. Contact your network administrator for the best way to fix the problem. Also, Platform LSF assumes that each host in the grid has a single name, can resolve the IP address from the name, and can resolve the official name from the IP address. If this is not the case, LSF will need its own 'hosts' file located in its configuration subdirectory (<LSF_TOP>/conf/hosts).

PORTS

Another area that needs to be checked is whether the ports SAS and LSF use for communication are accessible from other machines. The ports may not be accessible if a firewall is running on one or more machines. If firewalls are running ports need to be opened so that communication will work between the LSF daemons and the instances of SAS. Issuing the command `'telnet <host> <port>'` will return an error if *port* is not open on *host*.

By default, LSF uses ports 6878, 6879, 6881, and 6882. GMS uses port 1976 and PM uses 1966. If the port numbers need to be changed, the LSF ports are defined in `lsf/conf/lsf.conf`, GMS ports are defined in `gms/conf/ga.conf`, and PM ports are defined in `pm/conf/js.conf`. Changing the GMS or the PM port will require changing the metadata for the Grid Monitoring Server and Job Scheduler Server respectively. Also, a port may be in use by another program. With the LSF daemons stopped, you can issue the `'netstat -an | <search-tool> <port#>'` command and look for the specific LSF ports in the output. The search tool on UNIX is `grep`. The search tool on Windows is `findstr`. If the port is in use, reassign the port or stop the program using the port.

SAS will assign random ports to use for connections, but you can restrict the range of ports SAS will use using the `'-tcpportfirst <first_port>'` and the `'-tcpportlast <last_port>'` options. These can be specified in the SAS configuration file or on the SAS command line. For remote SAS/CONNECT sessions these would need to be specified in the grid command script (`sasgrid.bat`) that is run on the grid nodes. As an example, adding the following parameters to the SAS command line in `sasgrid.bat` will restrict the ports the remote SAS/CONNECT session will use to between 5000 and 5005:

-tcpportfirst 5000 -tcpportlast 5005

#2 – MAKE SURE LSF IS RUNNING

When the install has completed, you should make sure all the LSF daemons are running on each machine. To do this, go to each machine and run the `'lsid'` command. If that command completes correctly, the daemons are running on that machine. The daemons create log files that will help debug problems. They are either located in the machine's `<LSF_TOP>\logs` subdirectory (on Windows) or the shared `<LSF_TOP>/log` subdirectory (on UNIX). If the daemon does not have access to the share on UNIX, it will put the log files in the `/tmp` subdirectory.

Some things that may cause the command to fail include:

- On Windows machines, you need to have the `LSF_ENVDIR` and `LSF_BINDIR` environment variables set to the proper locations along with adding `%LSF_BINDIR%` to the `PATH` environment variable. `LSF_ENVDIR` is usually set to `\\<master_machine>\<LSF_Share>\etc`. `LSF_BINDIR` is usually set to `<LSF_Install_Directory>\bin`. The installation does change the registry to enable this, but the values will not be available to grid jobs unless the machine is rebooted.
- On UNIX machines, you may have to source the `<LSF_TOP>/conf/profile.lsf` file to setup the LSF environment first).
- A machine cannot access the configuration files. Make sure the machine has access to the shared directory that contains the binary and configuration files. If the master host that is sharing the drive comes up after the grid node machine that is trying to access the shared drive, the daemons on the node may not start. You can add the `LSF_GETCONF_TIMES` environment variable to the system environment and set it to the number of times you want the daemon to try accessing the share in 5 second intervals before the daemon quits. Setting it to a value of 600 will result in the node trying for 50 minutes $((600*5)/60)$ before quitting.
- The license file is invalid or missing. If LSF cannot find a license file, some daemons may not start or work correctly. Make sure the license file exists and is accessible by the daemons.
- Not all daemons are running. Restart the daemons on every machine in the grid using the `'lsfrestart'` command.
- Cannot connect to cluster master machine (the SAS grid control machine). This is the first machine listed in the `lsf.cluster.<cluster_name>` file. Make sure the master host has its daemons running and both machines can communicate with each other.

If changes need to be made that require the daemons to be restarted, you can run the `'lsfrestart'` command to stop and start the LSF daemons. If that does not work, under UNIX you can run the `'/etc/init.d/lsf restart'` command. Under Windows, you can go to the Services administration tool and stop the SBD, RES, and LIM services in that order and then start the LIM, RES, and SBD services in that order.

#3 – CHECK LSF SETUP

Some grid problems are caused by the host name being specified incorrectly in the `conf/lsf.cluster.<cluster_name>` file or the resource is not specified in `lsf.shared` file. Here are some suggestions to make sure the configuration is correct.

1. Log in as an LSF administrator (i.e., one of the usernames listed in the `lsf.cluster.<cluster_name>` under the **Administrators=username1 username2 ... usernameN** line) on one of the machines in the grid, preferably the grid control server machine.
2. Run the command **lsadmin ckconfig -v** to check the LSF configuration files for errors.
3. Run the command **badadmin ckconfig -v** to check the batch configuration files for errors.
4. Run the command **lshosts** to list all the hosts in LSF and make sure all the hosts are listed with the proper resources.
5. Run the command **bhosts** to list all the hosts in LSF's batch system. Make sure all hosts are listed. Make sure the Status for all hosts is set to **ok** and the MAX column has the correct number of jobs slots defined for each host (i.e. the maximum number of jobs the host can process at the same time).

If any problems occur, correct the LSF configuration file and issue the commands **lsadmin reconfig** and **badadmin reconfig** so that the daemons will use the updated configuration files. If you have added or removed hosts from the grid, you must restart the master batch daemon by issuing the command **badadmin mbdrestart**. To restart everything, you can issue the **lsfrestart** command.

#4 – CHECK LSF JOB EXECUTION

Some problems only come about when you actually try to run jobs on the grid. To minimize what could go wrong, debug jobs can be submitted to specific machines to isolate the problems.

To do this, run the command **bsub -I -m <host_name> set** from the grid client machine to each grid node so that you can be sure a job runs on that machine. The command should dump out the environment to the screen for a job running on the remote machine.

If this job fails, you can see what happened by executing the **bhist -I <job_id>** command using the job id output to the terminal. This output will list a number of things including the username of the person submitting the job, the command, and all the problems LSF had executing the job. Some messages in the **bhist** output from common problems include

- **"Failed to logon user with password"** – the password in the Windows' `passwd.lsfuser` file is invalid. Update the password using the **lspasswd** command.
- **"Unable to determine user account for execution"** – the user specified does not have an account on the destination machine. This may occur when going from a Windows grid client to a UNIX grid node because the Windows user has a domain prefixed to the user name. This can be eliminated by adding the line **LSF_USER_DOMAIN=<domain>** to the Windows' `lsf.conf` file to strip the domain from the username.

#5 – CHECK SAS GRID METADATA

SAS needs to retrieve information about the grid from a SAS Metadata Server in order to operate properly. Bring up the SAS Management Console and make sure the following is setup correctly:

LOGICAL GRID SERVER

- In the Server Manager under the application server context, there should be a logical grid server defined.
- The properties for the logical grid server should include the path to the script file that will be executed on the grid node. Make sure the path exists on every node in the grid and that the command is valid on every node in the grid.

GRID MONITORING SERVER

- In the Server Manager, a server should have been defined for monitoring the grid. The server's connection properties should have the name/address of the machine that is running the Grid Monitoring Server daemon (typically the SAS grid control machine) and the port should be the port value defined (1976 by default) in the GMS configuration file.

If the values are not correct in the metadata, the grid control server's configuration needs to be repaired by running the SAS Configuration Wizard again on the grid control server.

If the grid command does not exist on a grid node, the grid node's configuration needs to be repaired by running the SAS configuration wizard again on the grid node.

#6 – CHECK GRID MONITORING

The Grid Manager Plug-in for SAS Management Console displays information about the grid's jobs, hosts, and queues. Once the Grid Monitoring Server metadata has been setup and the Grid Management Service is running on the control server, grid information should be able to be displayed in SAS Management Console.

COMMON PROBLEMS

- Message “Connection timed out” or “Connection refused”
 - The Grid Management Service is not running. Go to the grid control machine and make sure the service has been started.
- Message “Your userid or password is invalid. Please try again or contact your systems administrator”.
 - Often this is the case when the user's credentials stored in the metadata has no password defined for the login associated with the authorization domain used by the Grid Monitoring Server connection. For example, “Grid 1 Monitoring Server” is defined in the metadata to use the “DefaultAuth” authorization domain. The user “User1” has a login defined in the User Manager for the “DefaultAuth” domain, but the login only has the user ID specified and the password is blank.

To fix this you can provide the password, remove the login for the authorization domain, or use a different authorization domain for the grid monitoring server connection. If you provide the correct password, the user will not be asked for a username and password. If you remove the login for that authorization domain or change the Grid Monitoring Server connection to use a different authorization domain, the user will be prompted for their user ID and password to connect to the machine where the Grid Monitoring Server is running.
- Message “Your username or password is invalid. Try again, or contact your systems administrator”.
 - This occurs when the user ID or password are invalid on the grid control machine. Make sure you provide a user ID and password that allows the user to log onto the machine where the Grid Management Service is running.

#7 – CHECK SAS JOB EXECUTION

SAS has a grid test program available on the SAS support website that should be run from a grid client to test connectivity to all nodes in the grid. It is located at <http://support.sas.com/rnd/scalability/grid/gridfunc.html#testprog>.

1. Copy and paste the grid test program into a Foundation SAS Display Manager Session.
2. If the application server associated with your logical grid server in your metadata is not 'SASMain', change all occurrences of 'SASMain' in the test program to the name of the application server associated with your logical grid server. For example, some SAS installations have the application server named 'SASApp', so all occurrences of SASMain should be replaced with SASApp.
3. Submit the code.

This will try to start one remote SAS session for every job slot available in the grid. This means it may start more than one job on multi-processor machines since LSF assigns one job slot for each CPU by default.

COMMON PROBLEMS:

- You get error messages
 - “Grid Manager not licensed” - Make sure your SID contains a license for SAS Grid Manager, aka slot 447.
 - “Grid Manager cannot be loaded” – Make sure Platform Suite for SAS has been installed and the LSF and PATH environment variables are defined properly.
 - “Invalid resource requested” – The application server name or workload value has not been defined in the lsf.shared file. Also make sure you associate the value with the hosts you want to run SAS programs in the lsf.cluster.<cluster_name> file.
- Number of grid nodes is 0. This can occur for a number of reasons:
 - The application server name has not been defined as a resource name in the lsf.shared file.
 - The application server name has not been associated with any grid nodes in the lsf.cluster.<cluster_name> file.
 - Grid client where job is submitted cannot communicate with the entire grid.
- Number of grid nodes is not the number of grid node machines.
 - As shipped, the number of grid nodes equals the number of CPUs in all grid nodes. If Hotfix **E9BA05** has been applied, the number of grid nodes equals the number of job slots in the grid. By default, the number of job slots is equal to the number of CPUs, but the number of job slots for a grid node can be changed.
 - The application server name has not been associated with all the grid nodes in the lsf.cluster.<cluster_name> file.

- Jobs fail to start. This can occur for a number of reasons:
 - Grid command defined in the logical grid server metadata is not valid on grid nodes or does not bring up SAS on the grid node when run. Logging onto a grid node and running the command defined in the logical grid server definition should attempt to bring up SAS on a grid node. SAS will try to start but fail since it is missing some grid parameters. If SAS tries to come up, the command is working properly. Platform will return a return code of 127 if the command to be executed is not found and a 128 return code if the command is found, but there is a problem executing the command.
 - Incorrect version of SAS installed on grid nodes. SAS version 9.1.3 SP3 is a minimum. Often a return code of 231 is associated with this problem.
 - Unable to communicate between grid client and grid node. Make sure the network is setup properly as indicated in **CHECK NETWORK SETUP** above.
- Jobs run on machines that are only supposed to be grid clients.
 - By default, all machines listed in the `lsf.cluster.<cluster_name>` file are part of the grid and can process jobs. If you want a machine to be able to submit jobs to the grid (i.e., a grid client), but not be one of the machines that can process the job (i.e., a grid node), you must set its maximum job slots to 0. See the **"Converting a Grid Node to a Grid Client"** section of the [Step-by-Step Installation of Platform Suite for SAS](#) document.

FOR MORE INFORMATION

For more information about SAS and grid computing, visit the following websites:

SAS Scalability and Performance Community

<http://support.sas.com/rnd/scalability/grid>

SAS Grid website

<http://www.sas.com/grid>

SAS/CONNECT Grid Syntax

<http://support.sas.com/rnd/scalability/grid/gridfunc.html>

Example wrapper code for load balancing multiple users SAS jobs

<http://support.sas.com/rnd/scalability/grid/gridfunc.html>

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.