

Installation Instructions for Platform Suite for SAS® Version 10.1 for Windows



The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2016. *Installation Instructions for Platform Suite for SAS® Version 10.1 for Windows*. Cary, NC: SAS Institute Inc.

Installation Instructions for Platform Suite for SAS® Version 10.1 for Windows

Copyright © 2016, SAS Institute Inc., Cary, NC, USA

All rights reserved. Produced in the United States of America.

For a hard-copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a Web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

U.S. Government License Rights; Restricted Rights: The Software and its documentation is commercial computer software developed at private expense and is provided with RESTRICTED RIGHTS to the United States Government. Use, duplication or disclosure of the Software by the United States Government is subject to the license terms of this Agreement pursuant to, as applicable, FAR 12.212, DFAR 227.7202-1(a), DFAR 227.7202-3(a) and DFAR 227.7202-4 and, to the extent required under U.S. federal law, the minimum restricted rights as set out in FAR 52.227-19 (DEC 2007). If FAR 52.227-19 is applicable, this provision serves as notice under clause (c) thereof and no other notice is required to be affixed to the Software or documentation. The Government's rights in Software and documentation shall be only those set forth in this Agreement.

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513-2414.

December 2016

SAS provides a complete selection of books and electronic products to help customers use SAS® software to its fullest potential. For more information about our offerings, visit support.sas.com/bookstore or call 1-800-727-3228.

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.

Table of Contents

Chapter 1 - Introduction	1
Platform RTM for SAS	2
Architecture	2
Installation Directories.....	4
Pre-Installation Requirements	5
Post-Installation Tasks.....	6
Chapter 2 - Installing Process Manager and LSF.....	7
Testing the Installation	27
Chapter 3 - Installing Grid Management Service (GMS).....	29
Testing the Installation	33
Chapter 4 - Installing LSF on Grid Nodes or SAS Foundation Grid Clients	34
Testing the Installation	47
Adding Nodes or SAS Foundation Clients to the Grid	48
Converting a Grid Node Machine to a Grid Client	48
Chapter 5 - LSF Quick Reference	49
Check Cluster Configuration (lsadmin)	50
Find Out Cluster Status (lsid and lsload)	51
Check LSF Batch Configuration (badmin)	52
Find Out LSF Batch System Status (bhosts and bqueues).....	53
For More Information	54
Appendix – Multiple Host Installation	55
Pre-Installation Requirements	55
Installation Procedure.....	55
Testing the Installation	55

Chapter 1 - Introduction

This document describes how to install the Platform Suite for SAS, version 10.1, on Windows hosts for use with SAS products and solutions. The Platform Suite for SAS can be an individual addition to several SAS products and solutions to provide enterprise-level scheduling capabilities on a single server environment. The Platform Suite for SAS is also included as part of the SAS Grid Manager product to enable:

- distributed enterprise scheduling
- workload balancing
- parallelized workload balancing

The Platform Suite for SAS includes the following components:

- **Process Manager** - the interface used by the SAS scheduling framework to control the submission of scheduled jobs to (LSF) Load Sharing Facility which manages any dependencies between the jobs. The Flow Manager and Calendar Editor clients are included with Process Manager and may be optionally installed. These clients are not required by SAS; however, they do provide additional functionality.
 - **Flow Manager** - provides a visual representation of flows that have been created for a Process Manager Server. These include flows that were created and scheduled in SAS Management Console's Schedule Manager, as well as reports that have been scheduled through SAS Web Report Studio. Platform Flow Manager provides information about each flow's status and associated dependencies. You can view or update the status of jobs within a flow, and you can run or rerun a single job regardless of whether the job failed or completed successfully.
 - **Calendar Editor** - a scheduling client for a Process Manager Server. This client enables you to create new calendar entries for time dependencies for jobs that are scheduled to run on the server. You can use it to create custom versions of the calendars that are used to create time dependencies for jobs.
- **LSF** – dispatches all jobs submitted to it, either by Process Manager or directly by SAS, and returns the status of each job. LSF also manages any resource requirements and performs load balancing across machines in a grid environment.
- **Grid Management Services** – provides the run-time information about jobs, hosts and queues for display in the SAS Grid Manager Plug-in for SAS Management Console.

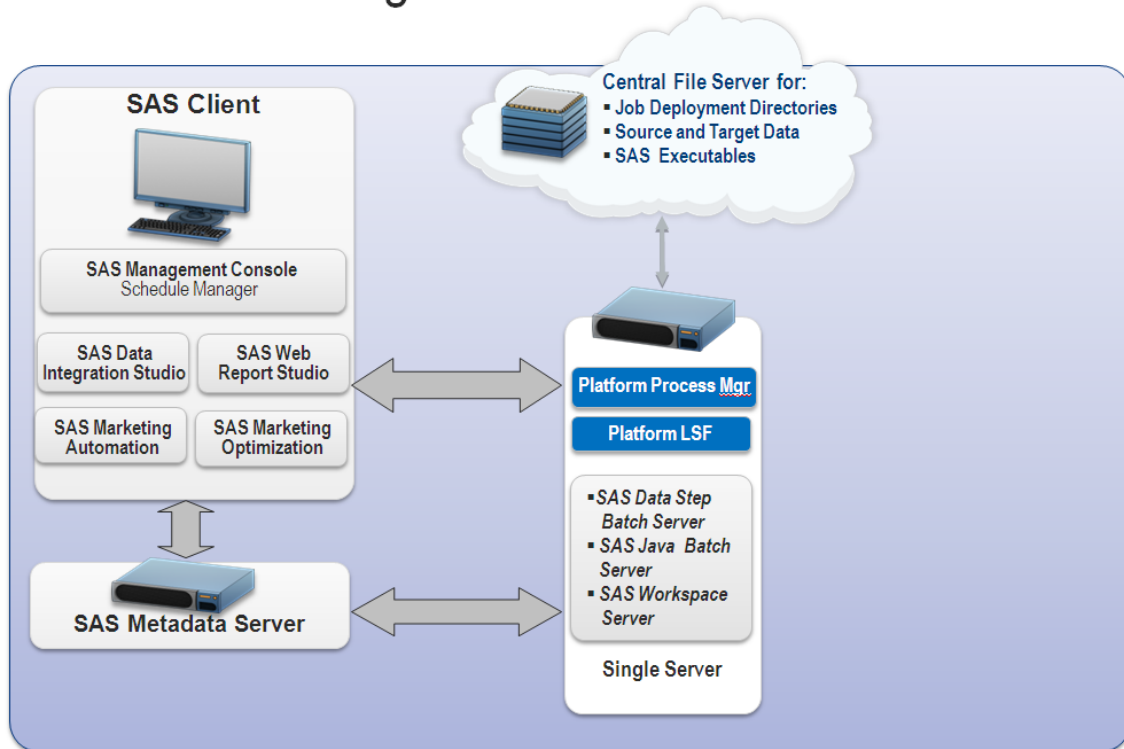
Platform RTM for SAS

Platform RTM (Report, Track & Monitor) for SAS is a Web-based tool that provides IT administrators a way to graphically view the status of devices and services within their SAS grid environment as well as manage the policies and configuration of their grid. It is a visual tool used to quickly track and diagnose issues before they affect service levels. Platform RTM for SAS includes drill-down capabilities to view details of hosts, jobs, queues, and user activities while instant alerts on job performance and grid efficiency allow administrators to optimize usage and workloads. It includes customizable graphs to visually analyze resource usage, workload trends, and job behavior. It also includes GUI interfaces that allow administrators to update the policies and rules in the grid configuration. Administrators can also set up high availability for any of the services that are critical to the operation of the grid as well as the applications executing in the grid. Platform RTM for SAS helps system administrators improve decision-making, reduce costs, and increase service levels for SAS grid deployments.

You can download Platform RTM for SAS from SAS Demos and Downloads site at http://www.sas.com/apps/demosdownloads/platformRTM_PROD__sysdep.jsp?packageID=000669. System requirements and installation instructions for Platform RTM for SAS are provided on the download page and are therefore not covered in this document.

Architecture

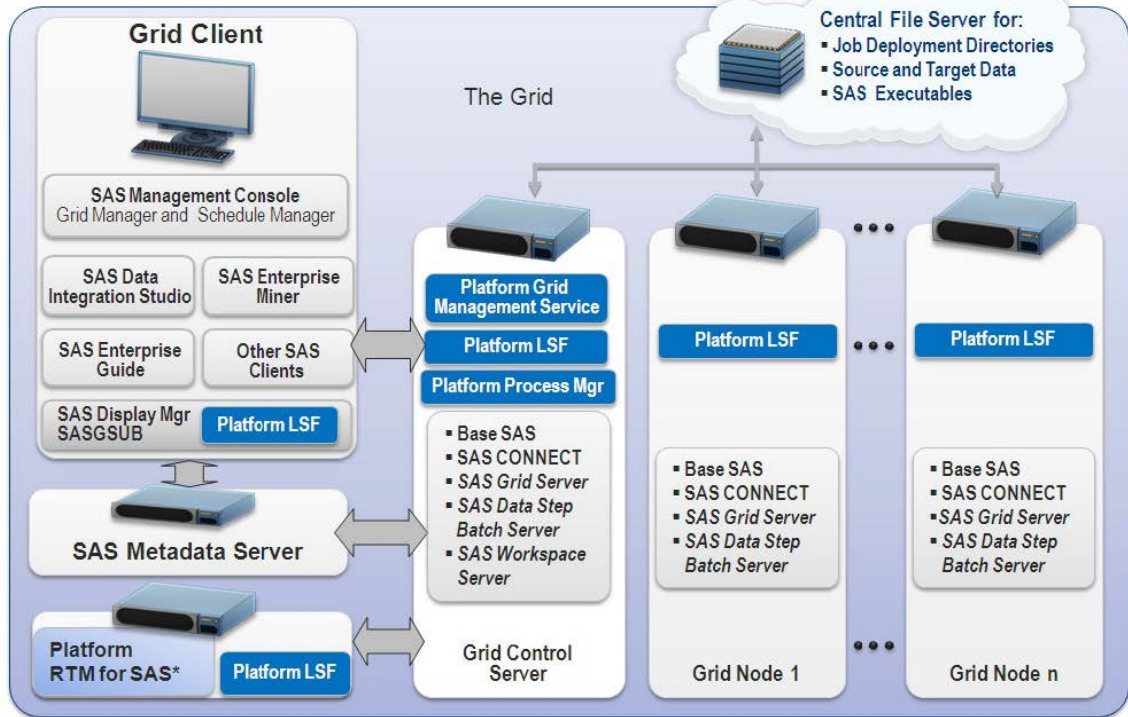
SAS® Scheduling Architecture



Copyright © 2010, SAS Institute Inc. All rights reserved.

Job scheduling on a single machine requires that you install Platform Process Manager 10.1. During the Process Manager install, you will also install Platform LSF 10.1. See “Chapter 2 - Installing Process Manager and LSF” for the instructions on installing on a single server.

SAS® Grid Computing Architecture



Copyright © 2010, SAS Institute Inc. All rights reserved.

* Linux and Windows only; see system requirements

Several types of machines make up a SAS grid environment. These machines have been defined to clarify the software components that must be installed on each one as well as the SAS metadata that must be configured. The SAS Metadata Server is shown on a separate machine in this sample architecture. It is common to dedicate a machine to running the SAS Metadata Server, but you may choose to run the metadata server on the grid control server. The three machine types specific to a grid installation are defined as follows:

- **grid client** - a grid client submits work to the grid but is not part of the grid resources available to execute work. Examples of a grid client include:
 - a SAS Data Integration Studio client (Platform LSF not installed on this client machine).
 - a SAS Enterprise Miner client (Platform LSF not installed on this client machine).
 - a SAS Management Console client using the Schedule Manager plug-in or any other applications scheduling SAS workflows. (Platform LSF not installed on this client machine).
 - a SAS Foundation install (minimum Base SAS, SAS/CONNECT, and Platform LSF) used to run a program that submits work—both whole programs or programs separated into parallel segments which are programs separated into parallel segments to the grid. Installation of the Platform LSF component is required in this case in order for SAS/CONNECT to submit the work to the grid.
 - Platform RTM for SAS (requires Platform LSF).

grid control server - any machine in the grid can be designated as the grid control server. More software is installed on the grid control server and more SAS metadata configuration takes place on this machine. It is highly recommended that you start the installation of the Platform Suite for SAS on this machine. In a SAS Data Integration Studio and SAS Enterprise Miner scenario the grid control server runs a workspace server that executes programs that utilize SAS/CONNECT to distribute work to the grid nodes. The grid control server can be configured as a grid resource capable of receiving work to execute or not, depending on the needs of your environment.

- **grid node** - a grid node is a grid computing resource capable of receiving the work that is being distributed. Each grid node must be running a minimum of Base SAS, SAS/CONNECT, and Platform LSF.

Installation of Platform Suite for SAS is performed first on the grid control server and is followed by installation on all of the grid node machines. Installation of Platform LSF on the grid control server can be installed as a part of the Process Manager installation or it can be installed by itself. This document will only show LSF being installed as part of the Process Manager installation.

Machines that do processing for the grid as well as machines that submit jobs to run on the grid must have Platform LSF installed. Grid clients such as SAS Data Integration Studio or SAS Enterprise Miner do not submit jobs directly but rather work with a SAS Workspace Server or a stored process server that does the job submission. Since those grid clients do not submit jobs, they do not need Platform LSF installed, but the machine where the workspace server or stored process server would need it installed. If you are writing your own grid-enabled SAS program in SAS Foundation and want to run the program, that grid client workstation must have Platform LSF installed since it will be doing the actual submission of jobs to the grid.

This document will assist you in installing Platform Suite for SAS to create a computer cluster and enable the cluster to work with the SAS Business Intelligence Platform. Please refer to the *Grid Computing* document, located at <http://support.sas.com/documentation/onlinedoc/gridmgr/index.html>.

SAS Grid Manager Control Server requires Platform Process Manager 10.1 and Platform Grid Management Service 8.01. Platform LSF 10.1 will be installed during the Platform Process Manager 10.1 installation. See “Chapter 2 - Installing Process Manager and LSF” for the instructions on installing Process Manager and LSF. See “Chapter 3 - Installing Grid Management Service (GMS)” for instructions on installing Grid Management Services. SAS Grid Manager Node and SAS Grid Manager Client require only Platform LSF 10.1. See “Chapter 4 - Installing LSF on Grid Nodes or SAS Foundation Grid Clients” for the instructions on installing LSF.

Installation Directories

Caution: Do not install Process Manager and LSF to the same directory.

The Platform Suite for SAS installation produces the following directory structure:

Process Manager Server and Client files are by default installed in C:\Program Files(x86)\Platform Computing\IBM Spectrum LSF Platform Process Manager\.

LSF uses a shared directory for configuration and management of jobs. The share directory should not be the same as the installation directory. All machines in the grid need to be able to access this shared directory. While the directory can be a share on any machine in the grid, it is recommended that the share not be on a machine in the grid for increased high availability.

The Grid Management Service (GMS) files are installed in the LSF directory structure under their own directory: gms

Pre-Installation Requirements

1. Read the Platform Web Services deployment documentation.
2. Create a domain Primary LSF administrator account. While the example screenshots in this document show a `<domain>\username`, it is recommended that you refer to the relevant section on User Accounts in the *SAS Intelligence Platform: Application Server Administration Guide* located at <http://support.sas.com/documentation/cdl/en/biasag/63854/HTML/default/viewer.htm#titlepage.htm> for guidance on the user ID to specify for your installation. The user account `<domain>\sas` will be used throughout this document as an example where a specific user account name is necessary for clarity. It is recommended that the same operating system account be used to deploy both SAS and Platform LSF. Separate accounts *can* be used but doing so means that the SAS install account must be given read/write access to the LSF configuration files. See the Platform Web Services deployment documentation for details. The Primary LSF administrator must have the following privileges on **grid control server**:

- Act as part of the operating system
- Adjust memory quotas for a process (increase quotas)
- Back up files and directories
- Bypass traverse checking
- Debug programs
- Log on as a service
- Replace a process level token
- Restore files and directories

The Primary LSF administrator must have the following privileges on EACH **grid node**:

- Act as part of the operating system
- Adjust memory quotas for a process (increase quotas)
- Debug programs
- Log on as a service
- Replace a process level token

Note: *The Primary LSF administrator account should belong to the Local Administrators group on each host.*

3. Locate the `SAS9*_*.txt` and `LSF10*_*.txt` file located in the `sid_files` directory in your SAS Software Depot. The `LSF10*_*.txt` file will license all components of Platform Suite for SAS in a scheduling capabilities on a single server environment. The `SAS9*_*.txt` will license all the components of Platform Suite for SAS as part of the SAS Grid Manager.
4. Create a shared directory which is not the same as the installation directory for LSF. Ensure the shared directory is accessible with the same path name from all machines in the grid. The LSF administrator needs to be able to write to this directory from all the machines in the grid. For example, if the host machine is named `fileserver` and the shared directory is named `lsf10` then you would enter `\\fileserver\lsf10` when the install process asks for the shared directory name.

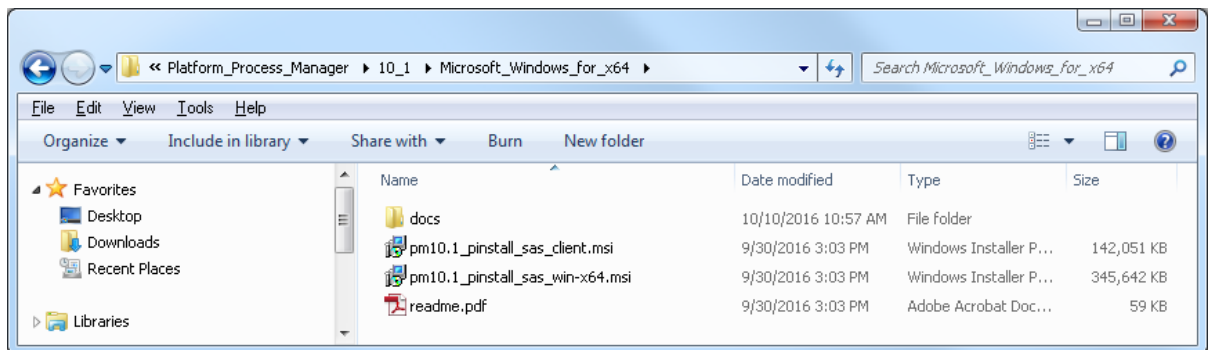
Post-Installation Tasks

Once you have installed Platform Suite for SAS, you still must enable the cluster to work with SAS Software. First, consult *Grid Computing in SAS*, located at <http://support.sas.com/documentation/onlinedoc/gridmgr/index.html>. After you complete the tasks in that document, you should then refer to the *Scheduling in SAS* document, located at <http://support.sas.com/documentation/onlinedoc/sasmc/index.html>. This must be done after SAS has been installed and configured.

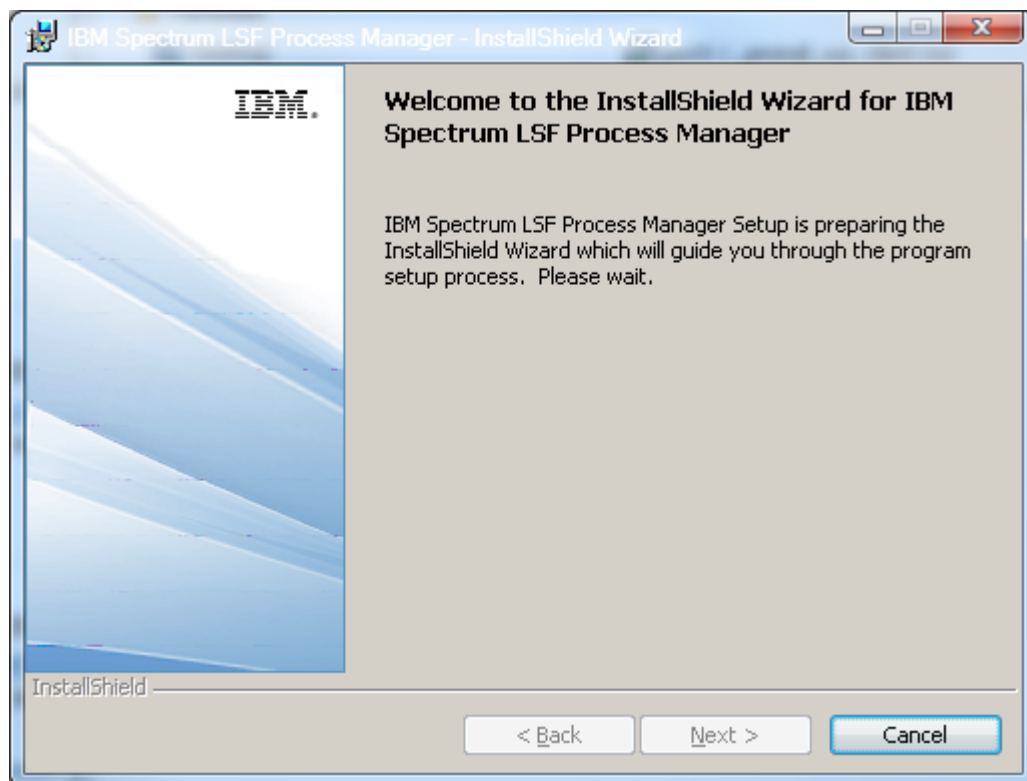
Chapter 2 - Installing Process Manager and LSF

Caution: Do not install Process Manager and LSF to the same directory.

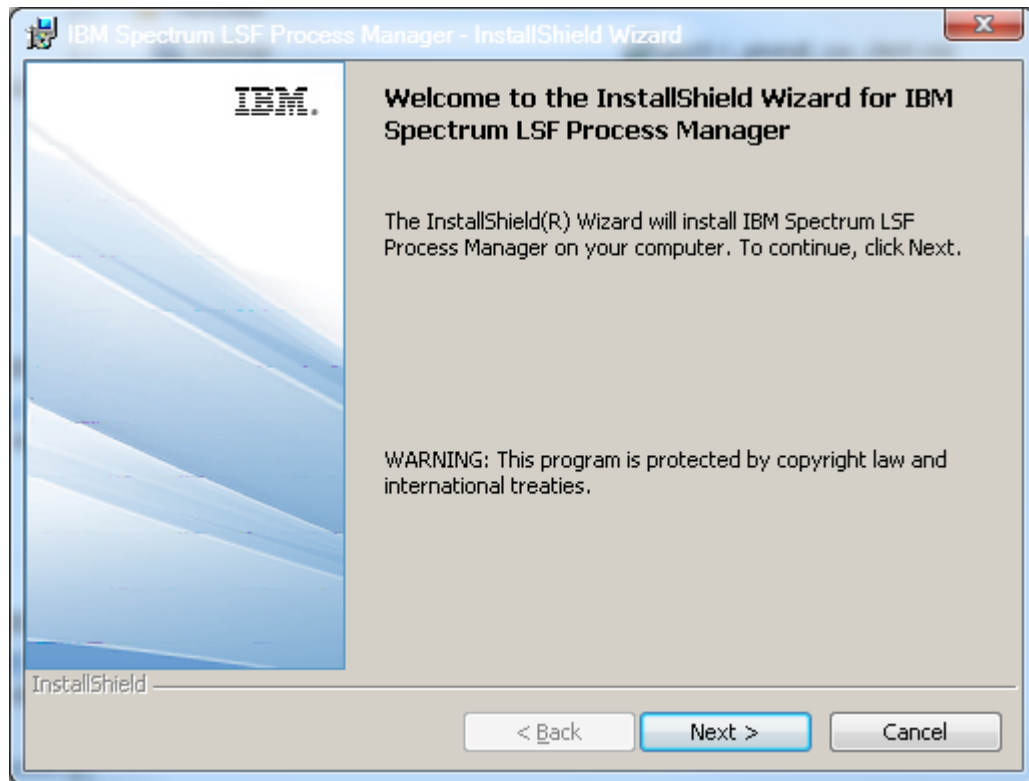
1. Run the executable located in the SAS Software Depot. The location is in the `third_party/Platform_Process_Manager/10_1` directory under the sub-directory named for the operating system you are installing on.



Next, the Platform Process Manager – InstallShield Wizard dialog displays showing the extraction progress.

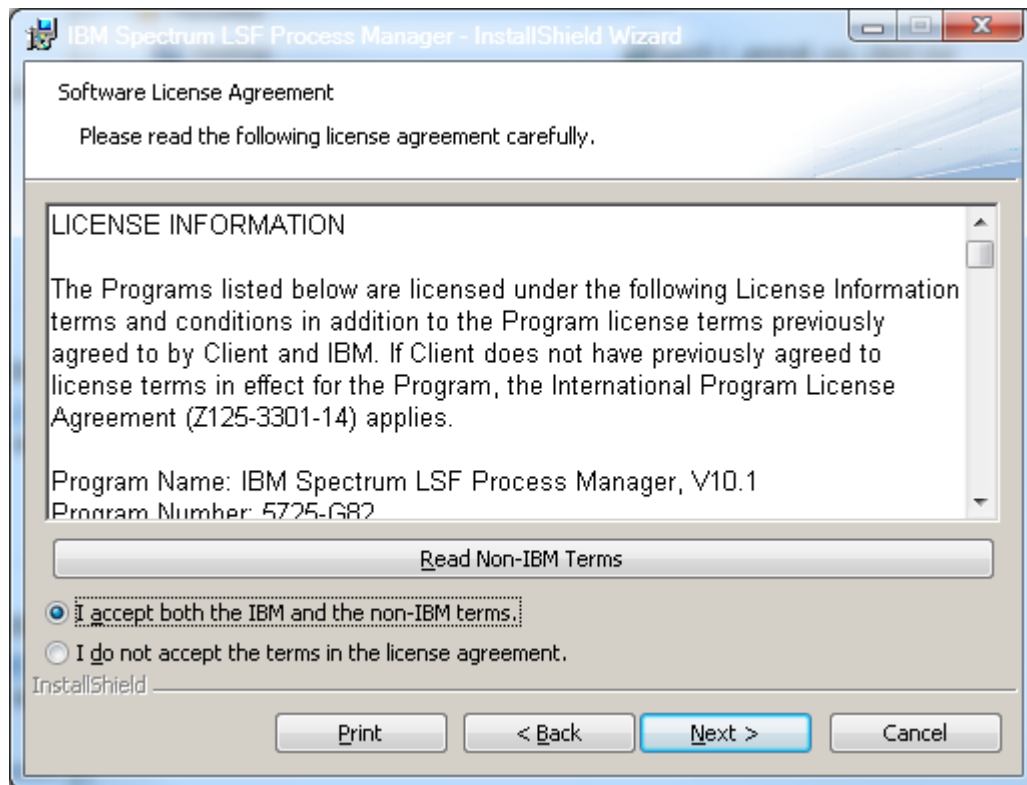


When the extraction process is complete, the Platform Process Manager – InstallShield Wizard dialog confirms installation on your computer enabling the **Next** button.



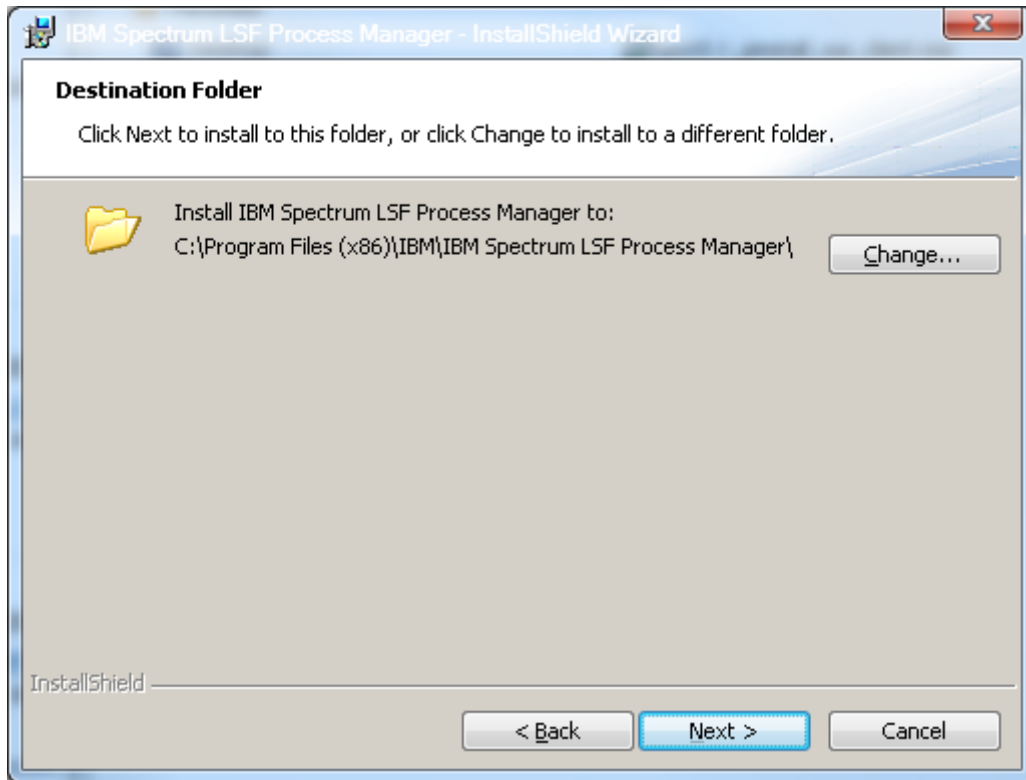
2. Click **Next** to continue.

The Platform Process Manager – InstallShield Wizard dialog End User Software License Agreement is displayed.



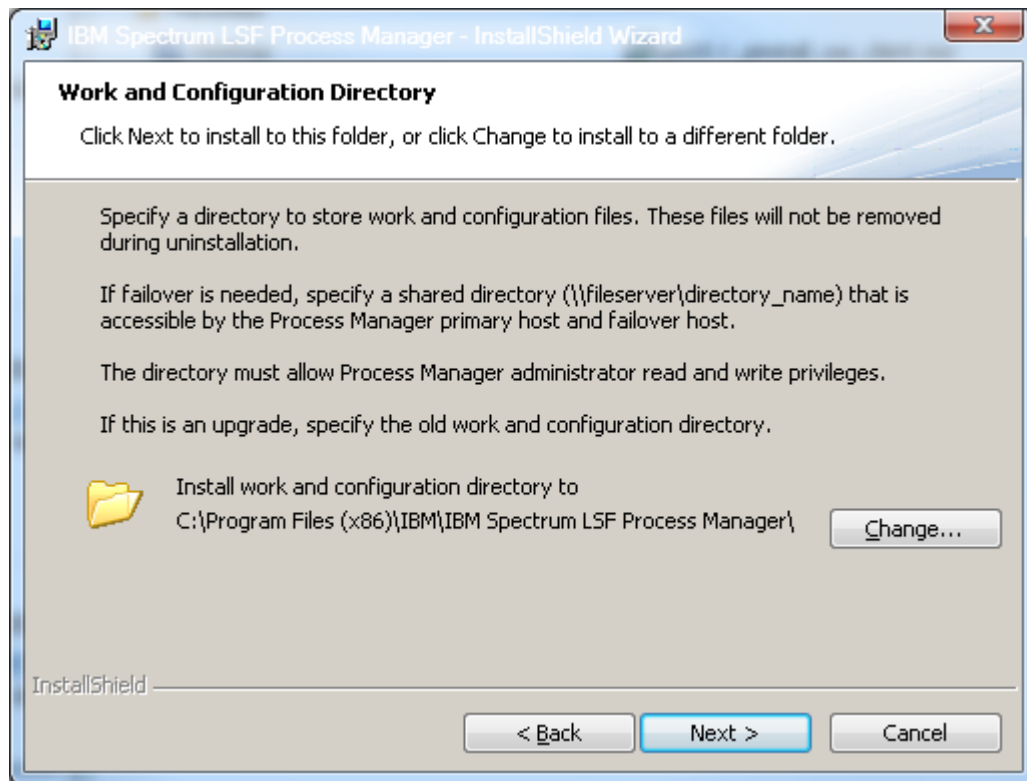
3. Select **Accept** after reading the agreement (the default is **Do not accept**). Click **Next** to continue.

The Platform Process Manager - InstallShield Wizard/Destination Folder dialog asks for the location to install the Platform Process Manager files.



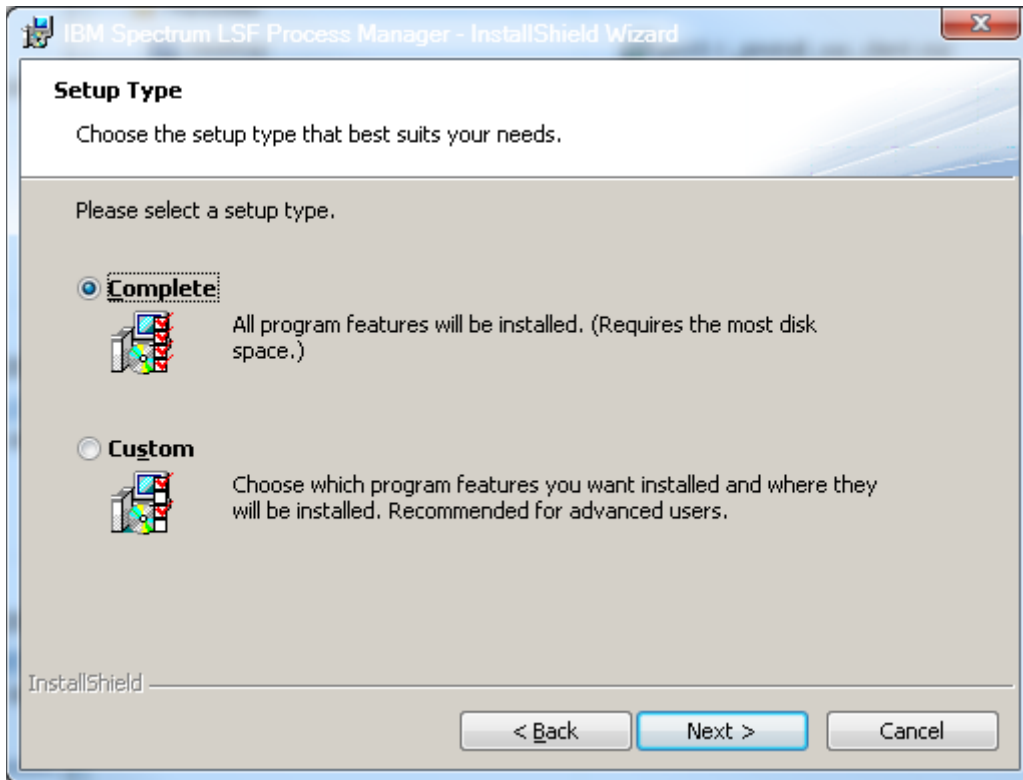
4. Click **Next** to continue.

The Platform Process Manager - InstallShield Wizard/Work and Configuration Directory dialog asks for the location to install the work and configuration files.



5. Click **Next** to continue.

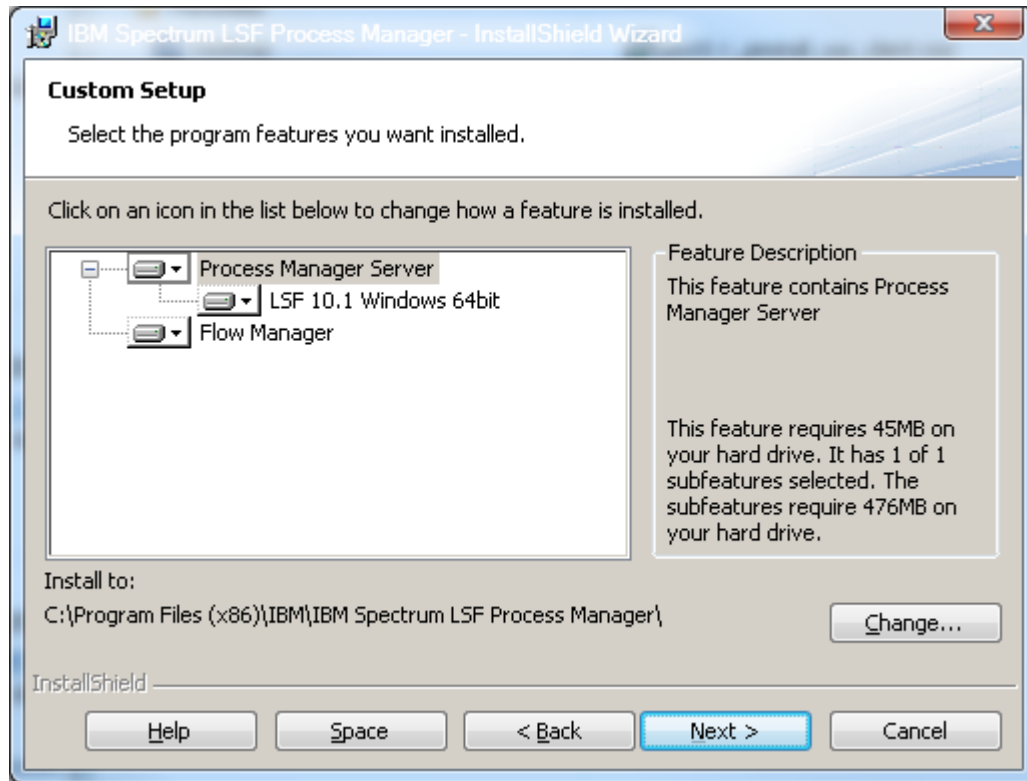
The Platform Process Manager - InstallShield Wizard/Setup Type dialog offers two types of setups: Complete or Custom.



Since both Process Manager and LSF are being installed, select the **Complete** option.

6. Click **Next** to continue.

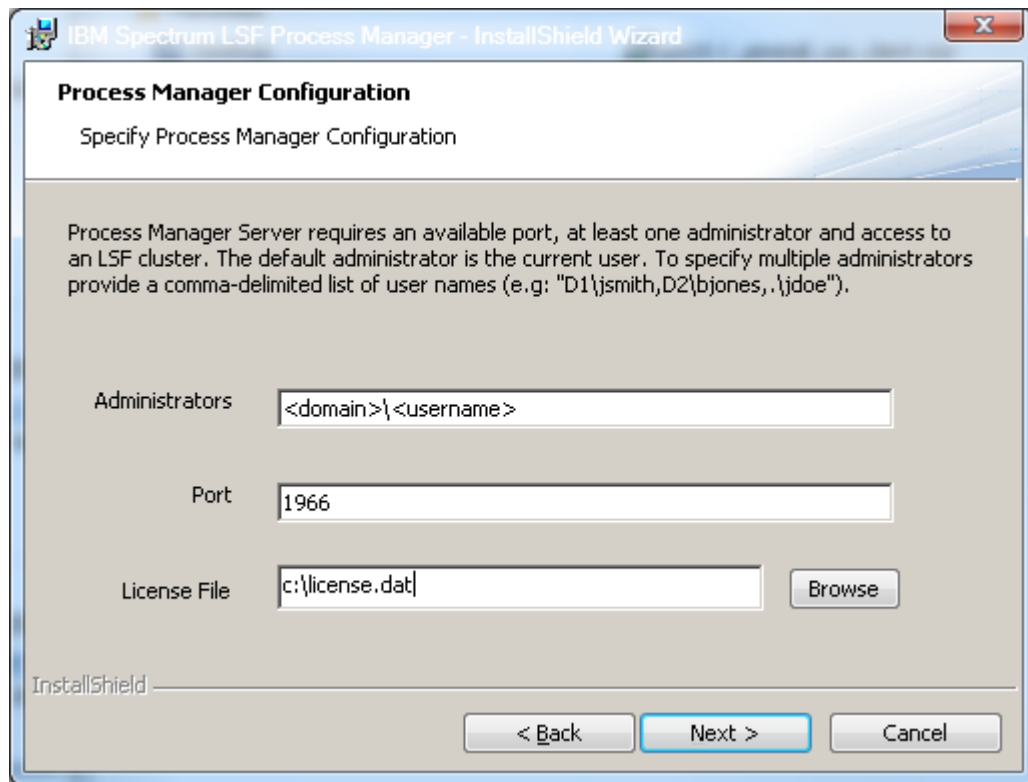
(Optional) In the Platform Process Manager - InstallShield Wizard/Custom Setup dialog, if the **Custom Setup** is selected, then the installation prompts for the components you want to install.



Since you are installing both Process Manager and LSF, make sure everything is selected. If LSF is already installed, remove the checkmark next to the **Platform LSF** option.

7. Click **Next** to continue (the default is the **Back** button).

The Platform Process Manager - InstallShield Wizard/Specify Process Manager Configuration dialog asks for the name of the cluster Administrators field and the path of the License File.



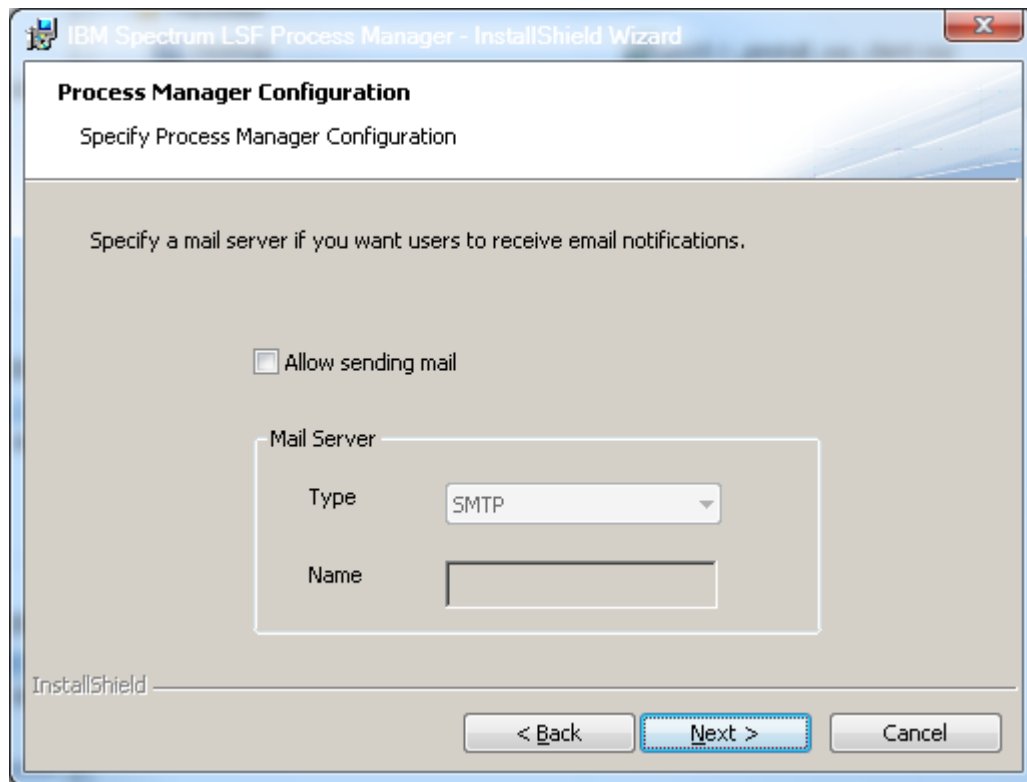
8. Enter the domain LSF administrator you created in the pre-installation requirements (usually `<domain>\sas` where `<domain>` is your Windows domain name). Find the license file provided by SAS, including its path, mentioned in the pre-installation requirements. Leave the port number at 1966.
9. Click **Next** to continue.

Next, the Platform Process Manager - InstallShield Wizard/Process Manager Configuration dialog asks for the user name to run the Process Manager and LSF service programs.

The image shows a Windows-style dialog box titled "IBM Spectrum LSF Process Manager - InstallShield Wizard". The main heading is "Process Manager Configuration" with the subtitle "Specify Process Manager Configuration". The text inside explains that the user specified will be the login for the Process Manager Server and will be added to the Administrators group. It provides instructions on the format for the username: `<DOMAIN>\<USERNAME>`. Below this, there are two input fields: "Username" with the placeholder text `<domain>\<username>` and "Password" with a masked input (dots). At the bottom, there are three buttons: "< Back", "Next >" (highlighted in blue), and "Cancel". The "InstallShield" logo is visible in the bottom left corner of the dialog area.

10. Enter the domain Primary LSF administrator you created in the SAS pre-installation checklist (usually `<domain>\sas` where `<domain>` is your Windows domain name).
11. Click **Next** to continue.

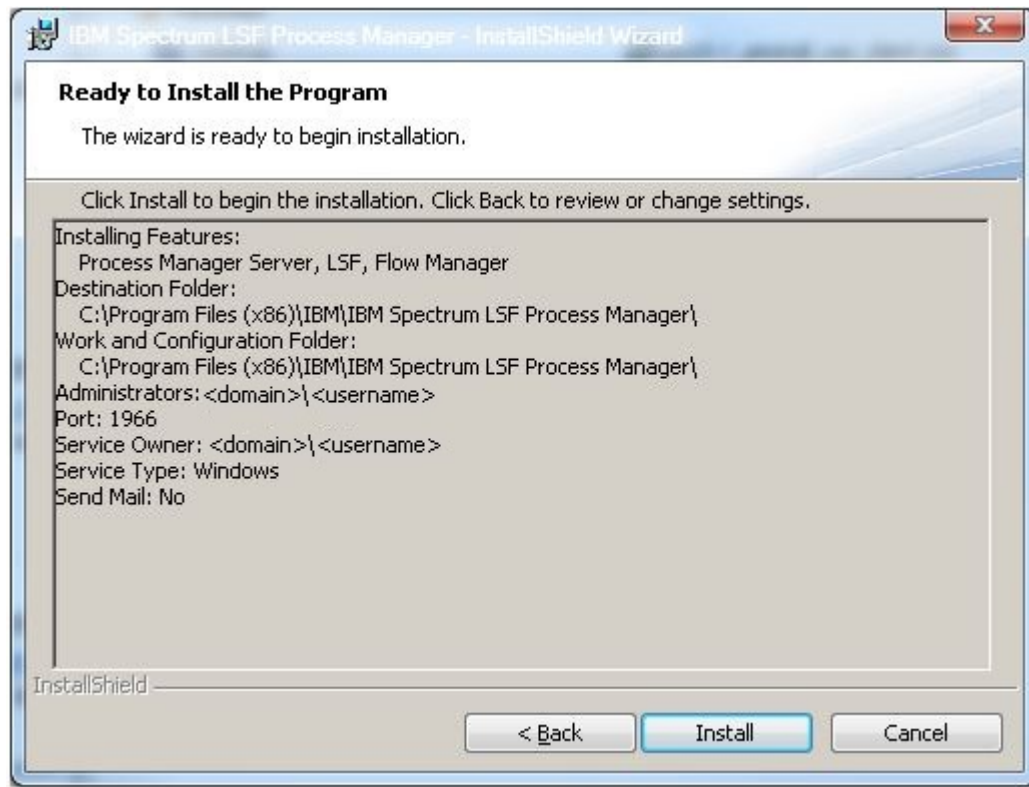
Next, Platform Process Manager - InstallShield Wizard/Process Manager Configuration dialog offers the option of receiving an email alert whenever a job is completed.



The email configuration is optional. If you don't want Process Manager or LSF to send an email, then, leave the checkbox unchecked.


12. Click **Next** to continue.

The Platform Process Manager - InstallShield Wizard/Ready to Install the Program dialog has gathered all the information required and displays a summary before starting the installation.



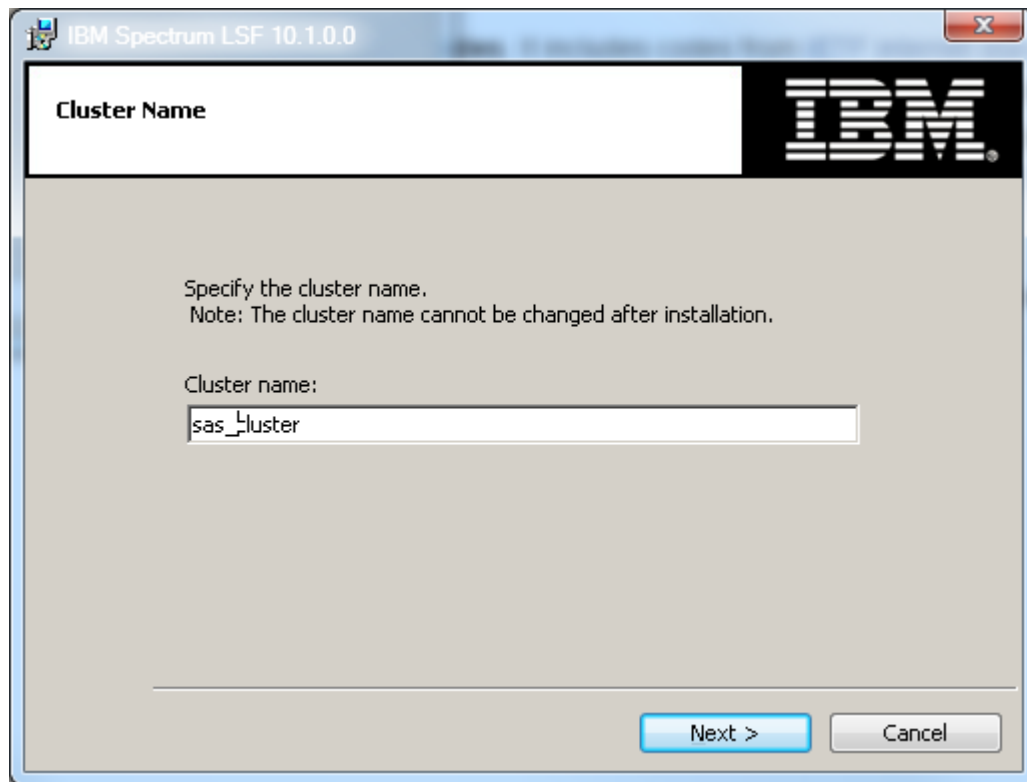
13. Click **Install** to continue.

The Platform Process Manager - InstallShield Wizard/Installing IBM Platform Process Manager dialog displays the progress of program features being installed.



14. Click **Next** to continue.

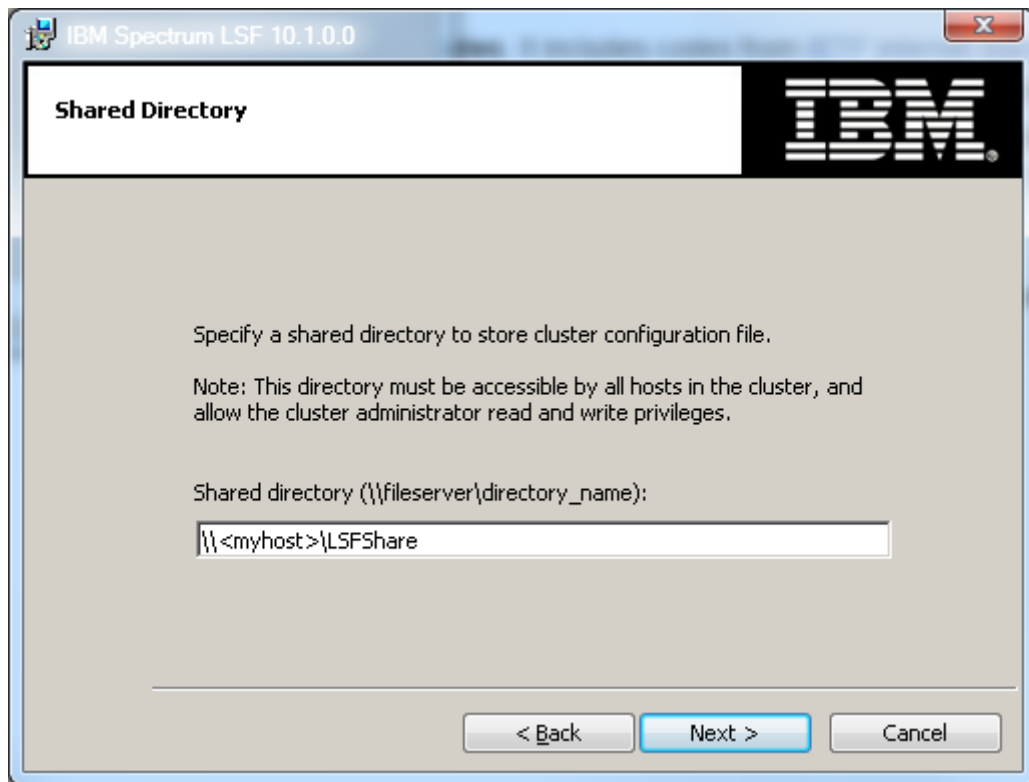
The IBM Platform LSF dialog requests the name of the cluster. This name is used to identify LSF services and configuration files.



15. Enter the name of the grid (cluster), without spaces.

16. Click **Next** to continue.

Next, the IBM Platform LSF/Shared Directory dialog requests where to install the configuration files for the cluster.

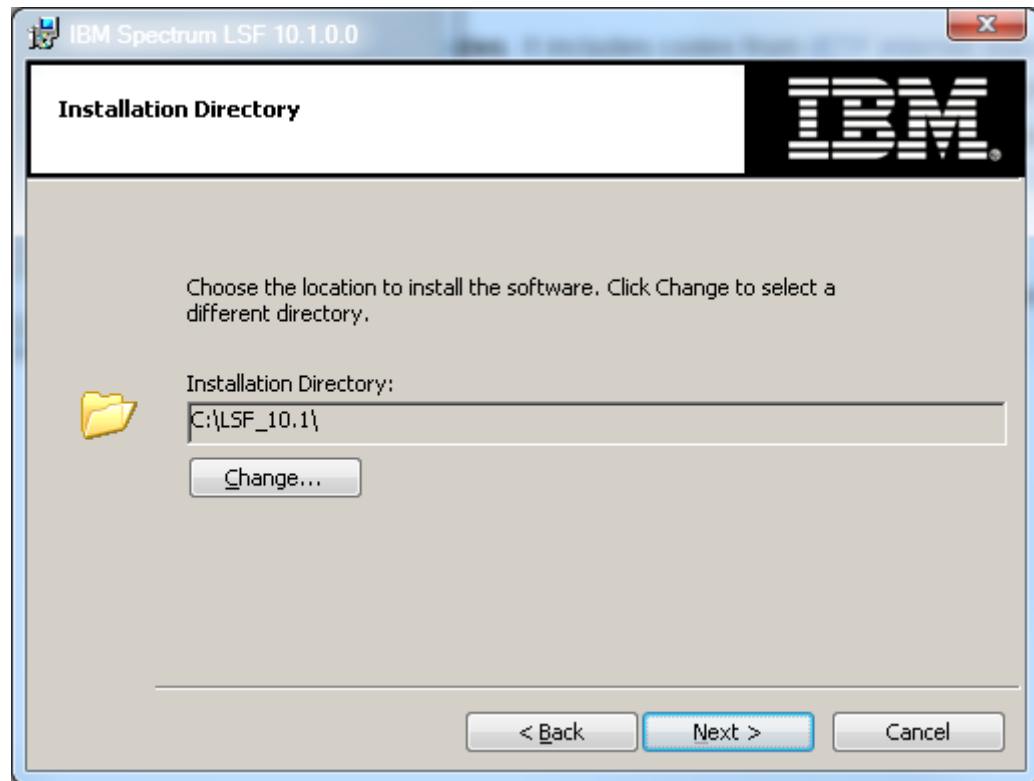


Note: Make sure the Primary LSF administrator account (such as <domain>\sas) has write permission to the share. Otherwise, you will receive Error 30035 about the failure to run the egoconfig mghost command.

17. Enter a new path or choose the default option.

18. Click **Next**. You must use a UNC path.

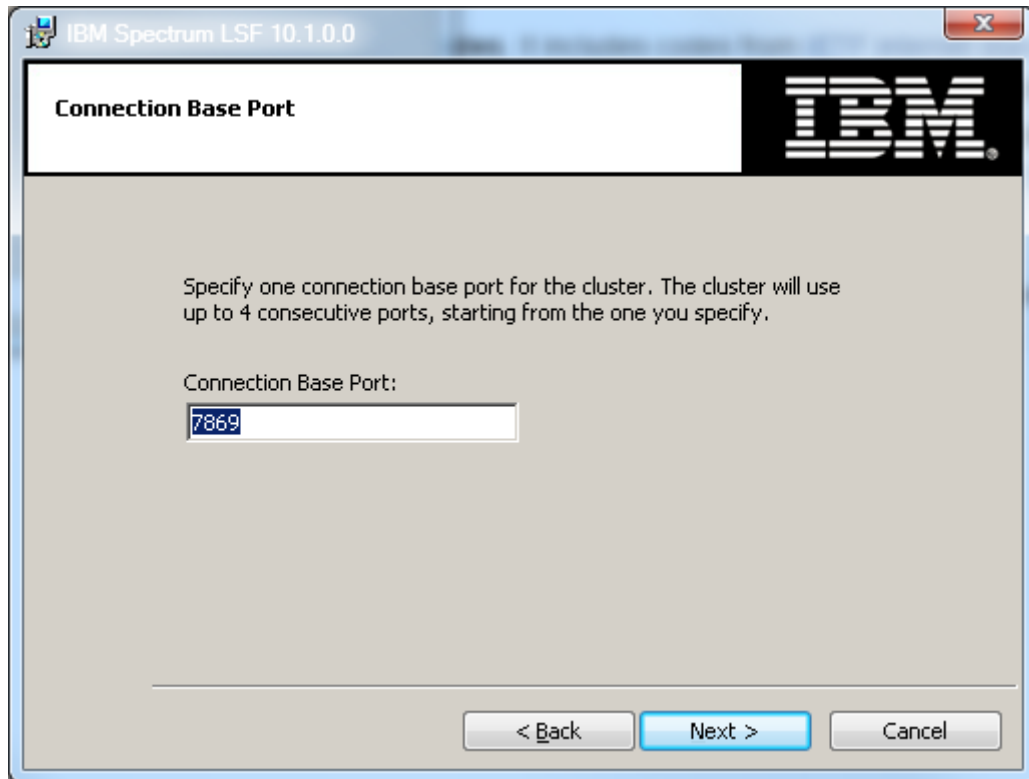
Next, the IBM Platform LSF/Installation Directory dialog requests where the files are to be installed.



Note: Make sure the installation directory is not the same as the shared directory in the previous step. Otherwise, you will receive error during the installation.

19. Click **Next** to continue.

Next, the IBM Platform LSF dialog requests the Connection Base Port for the grid (cluster) connection.

The image shows a Windows-style dialog box titled "IBM Spectrum LSF 10.1.0.0". The dialog has a header bar with the title and a close button. Below the header, there is a section titled "Connection Base Port" with the IBM logo to its right. The main area of the dialog contains the following text: "Specify one connection base port for the cluster. The cluster will use up to 4 consecutive ports, starting from the one you specify." Below this text is a label "Connection Base Port:" followed by a text input field containing the number "7869". At the bottom of the dialog, there are three buttons: "< Back", "Next >" (which is highlighted in blue), and "Cancel".

Unless you know there is a conflict, leave the port number at 7869. In this example, LSF uses ports 7869, 7870, 7871, and 7872.

20. Click **Next** to continue.

Next, the IBM Platform LSF dialog requests the LSF port for grid (cluster) services.



The image shows a Windows-style dialog box titled "IBM Spectrum LSF 10.1.0.0". The dialog has a tab labeled "LSF Ports" and the IBM logo in the top right corner. The main area contains the instruction "Specify the LSF ports for the cluster." followed by three input fields: "LSF RES Port:" with the value "6878", "LSF SBD Port:" with the value "6882", and "LSF MBD Port:" with the value "6881". At the bottom, there are three buttons: "< Back", "Next >" (which is highlighted in blue), and "Cancel".

Unless you know there is a conflict, leave the default values as is.

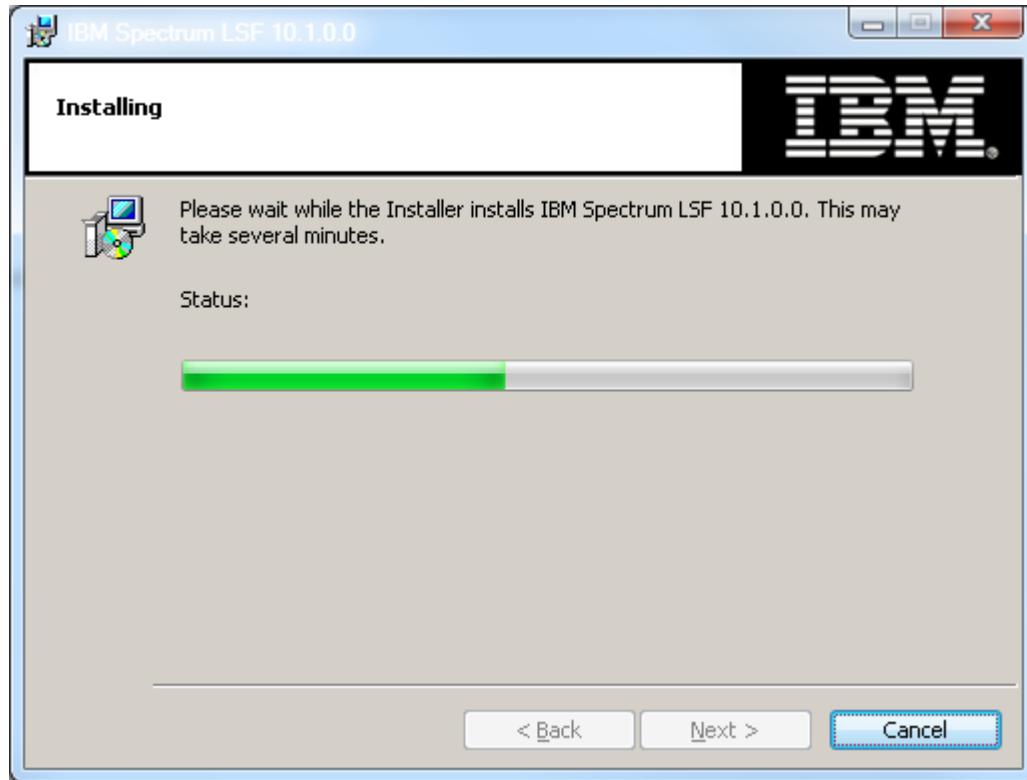
21. Click **Next** to continue.

When the install is complete, the IBM Platform LSF dialog displays an installation results summary.

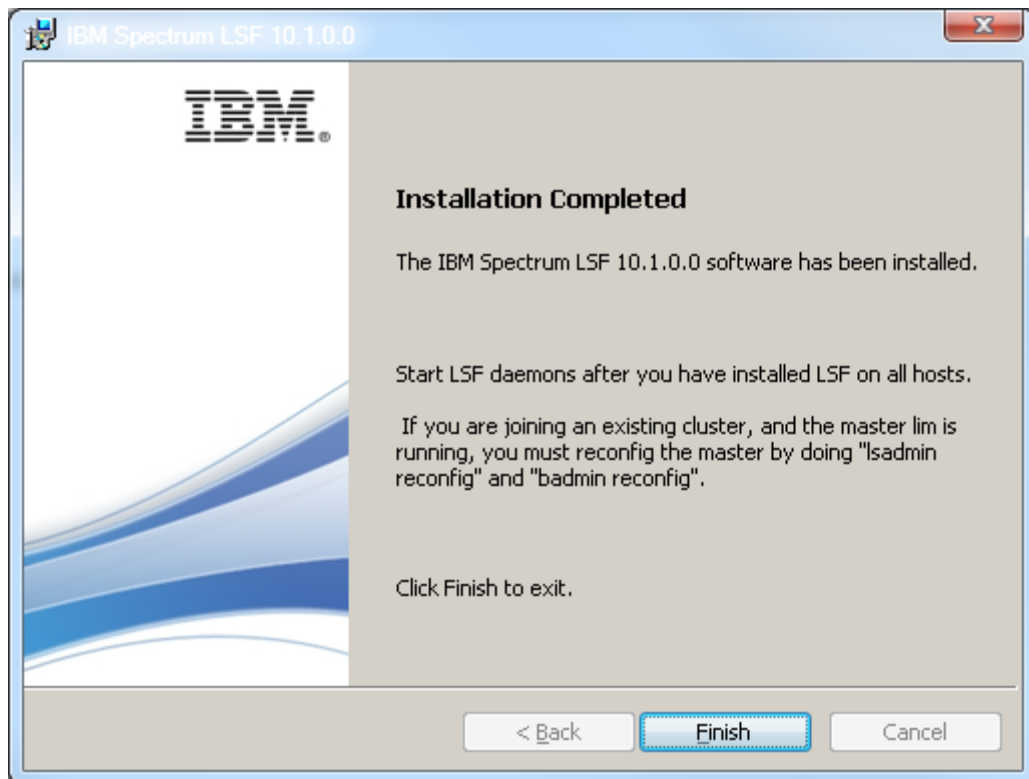


22. Click **Install** to continue. A status bar indicates the progress of the installation.

Make sure that the installation was successfully completed.

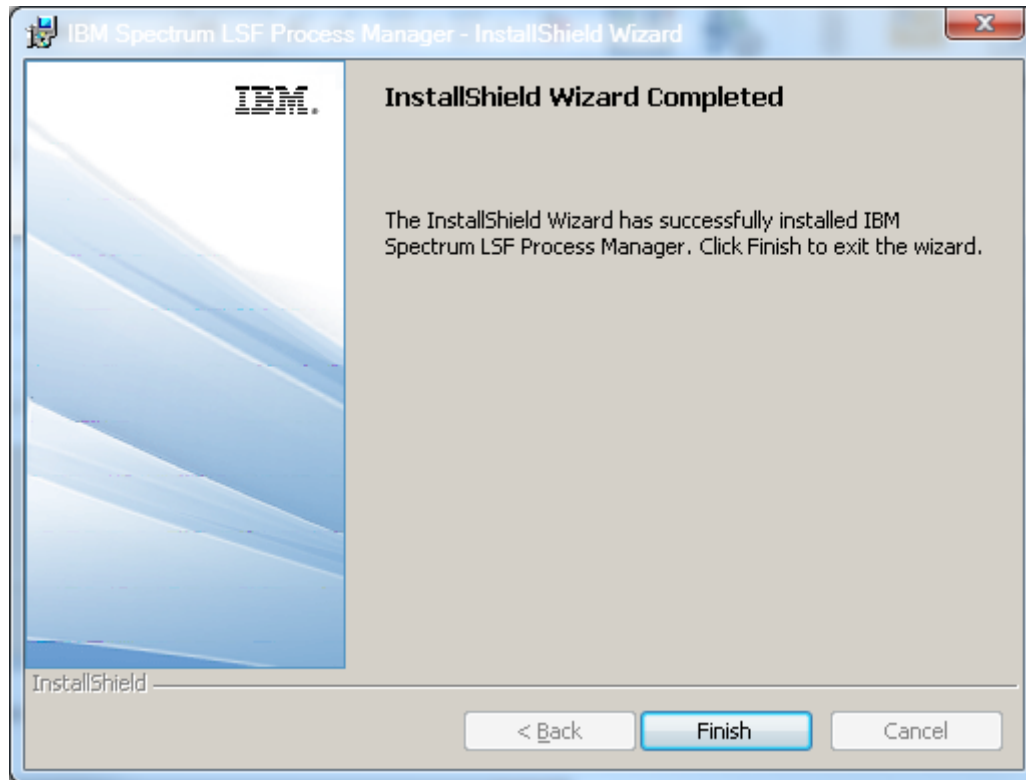


The IBM Platform LSF dialog states that the installation is complete and displays its final screen.



23. Click **Finish** to continue.

The IBM Platform Process Manager/InstallShield Wizard dialog has completed the installation and displays a final screen.



24. Click **Finish** to continue.

25. Reboot your system.

*Note: Windows users must register their Windows user account passwords with LSF by running the command **lspasswd**. If users change their passwords, they must use this command to update LSF. A Windows job command does not run if the password is not registered in LSF. If your site has a utility for updating passwords across your systems, then you may want to consider adding the **lspasswd** command to that utility.*

Testing the Installation

Once the system has rebooted, you can follow these steps to make sure LSF on the grid control server (LSF master machine) is operating properly.

1. Log onto the machine as one of the users you added to LSF using the **lspasswd** command. These are usually `<domain>\sas`, `<domain>\sasadm`, `<domain>\sastrust`, or `<domain>\sasdemo`.
2. Open a Command Prompt window by selecting **Start→Programs→Accessories→Command Prompt**.
3. Run the command **lsid**. This will display the cluster name and the grid control server (LSF master machine) name.
4. Run the command **lshosts**. This will display static information about the grid control server (LSF master machine).

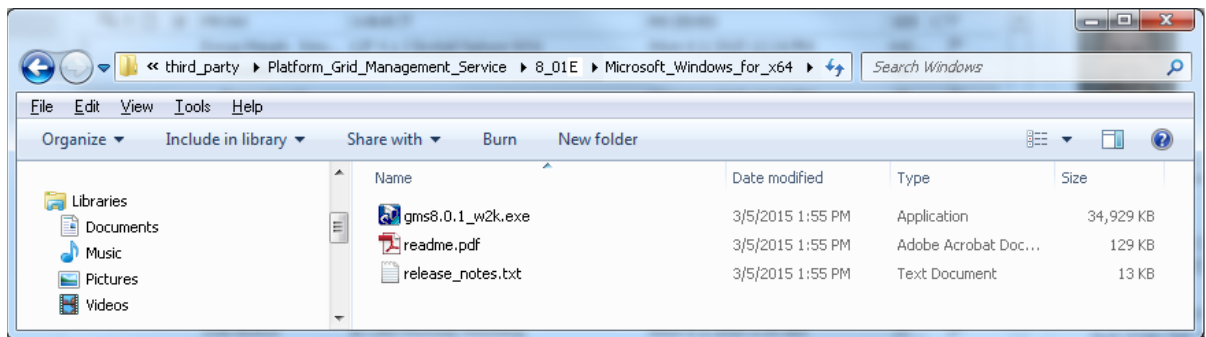
5. Run the command **lsload**. This will display dynamic information about the grid control server (LSF master machine).
6. Run the command **bsub sleep 100**. This will submit a job to the cluster.
7. Run the command **bjobs**. This will display the job information. As you repeat this command, you should see the job go from **PEND**, to **RUN**, to being removed from the queue.
8. Run the command **jid**. When prompted for username and password, provide the Primary LSF administrator credentials. This will display static information about the Process Manager Server.
9. Start Flow Manager by selecting **Start→Programs→IBM→Flow Manager** to run a client application to verify client communication to the Process Manager Server.

Note: *If you are installing Platform Suite for SAS for single machine scheduling, the task is complete at this stage and you can stop here. If you are installing Platform Suite for SAS for use with SAS Grid Manager, continue with the next chapter.*

Chapter 3 - Installing Grid Management Service (GMS)

Grid Management Service is a daemon that is used by the Grid Manager Plug-in for the SAS Management Console to display grid information. After installing this service, you can use the SAS Management Console to view grid information.

1. Log onto the grid control server as the Primary LSF administrator.
2. Run the installation program by double-clicking `gms8.0.1_w2k.exe`. The self-extracting executable is provided by SAS.

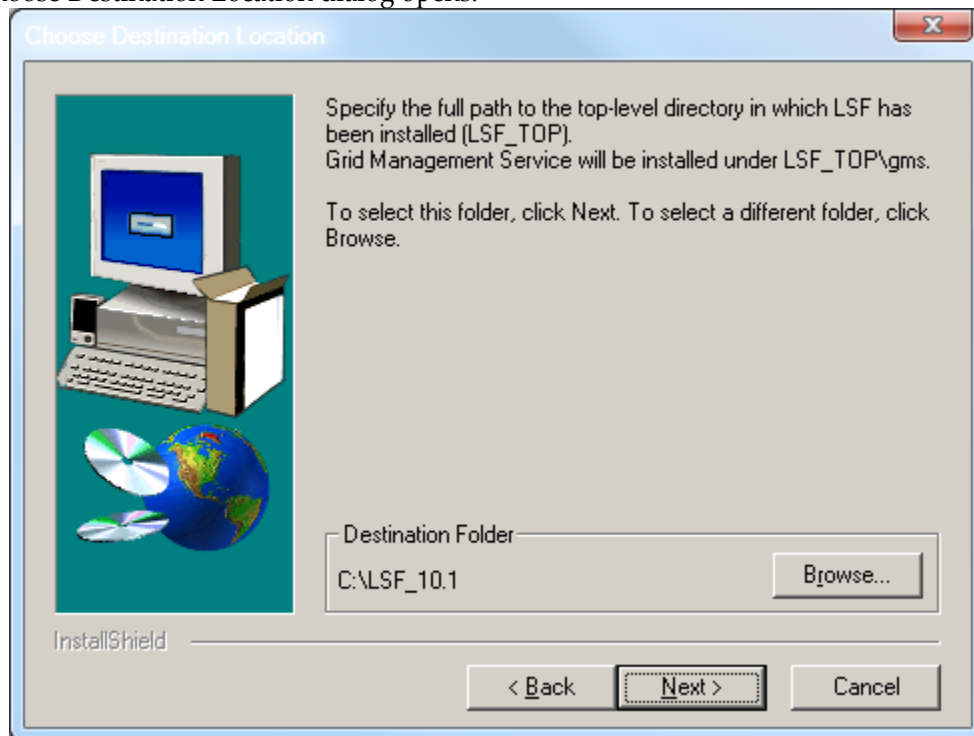


The Welcome to the Platform Grid Management Service installer dialog opens.



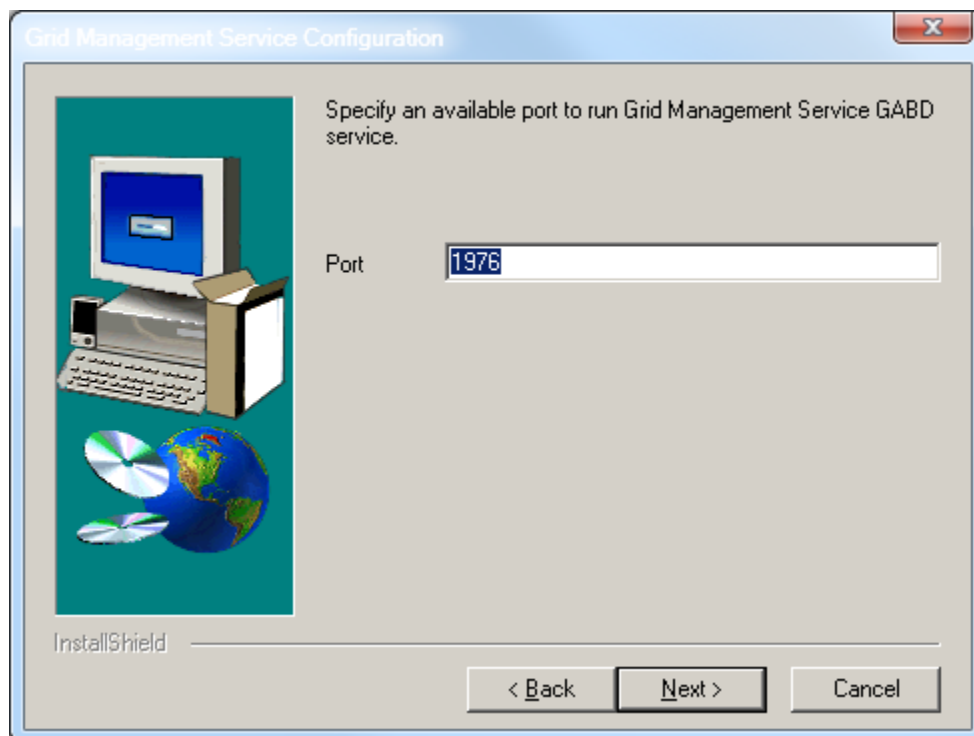
3. Click **Next** to continue.

The Choose Destination Location dialog opens.



4. Choose the preferred destination location.
5. Click **Next** to continue, unless you need to change the installation location.

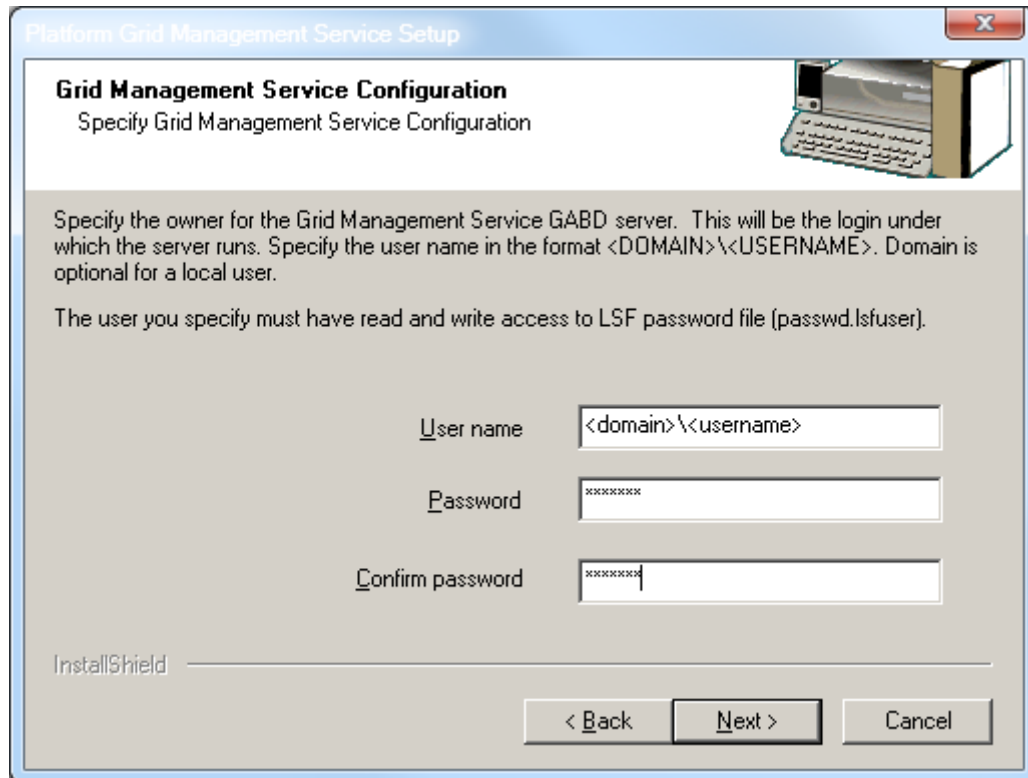
The Grid Management Service Configuration dialog opens.



Unless there is a conflict, leave the port number in the **Port** field at the default, 1976. Make sure your port is not already in use.

6. Click **Next** to continue.

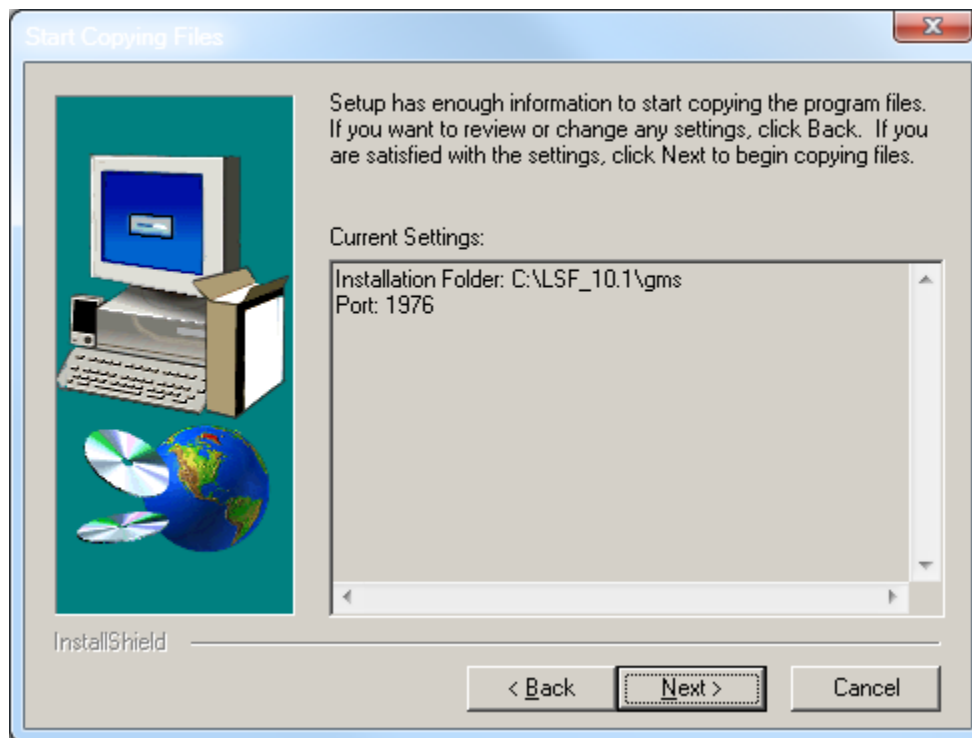
The Platform Grid Management Service Setup/Configuration dialog opens.



The screenshot shows a Windows-style dialog box titled "Platform Grid Management Service Setup". Inside, the main heading is "Grid Management Service Configuration" with the subtitle "Specify Grid Management Service Configuration". To the right of the text is a small icon of a computer monitor and keyboard. The main text area contains instructions: "Specify the owner for the Grid Management Service GABD server. This will be the login under which the server runs. Specify the user name in the format <DOMAIN>\<USERNAME>. Domain is optional for a local user." and "The user you specify must have read and write access to LSF password file (passwd.lsfuser).". Below this are three input fields: "User name" with the placeholder text "<domain>\<username>", "Password" with masked characters "xxxxxxxx", and "Confirm password" with masked characters "xxxxxxx". At the bottom left is the "InstallShield" logo. At the bottom right are three buttons: "< Back", "Next >", and "Cancel".

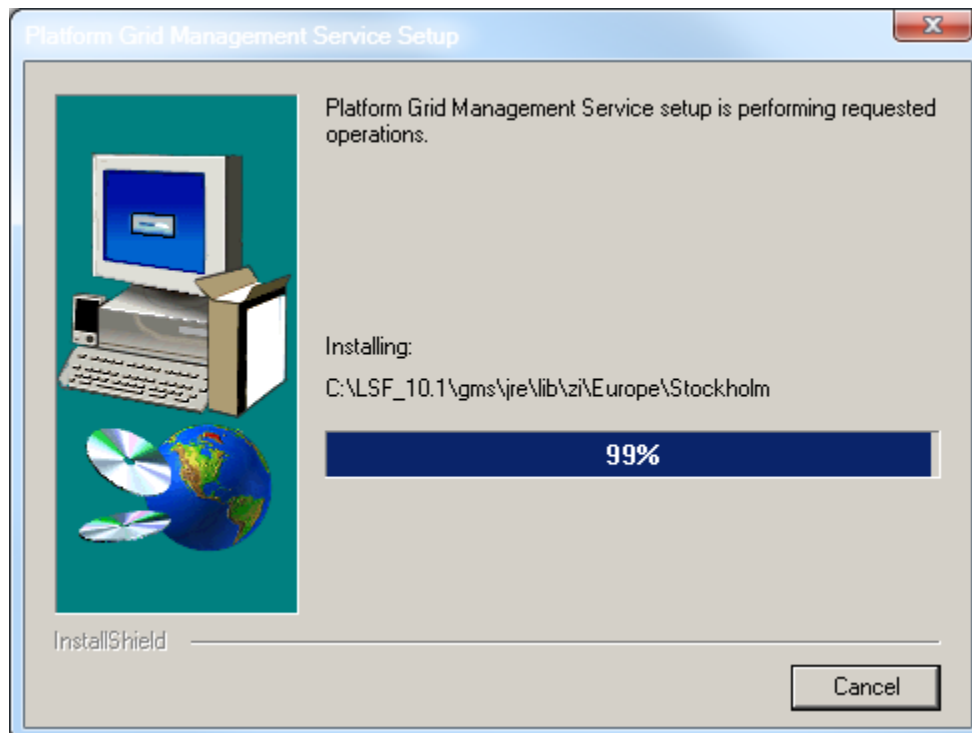
7. Specify the domain and user name under which the Grid Management Service server will run.
It is recommended that Grid Management Service use the same installation account as LSF. The account needs to be an LSF administrator. Specify the password for this user name. Confirm the password.
8. Click **Next** to continue.

The Start Copying Files dialog opens.

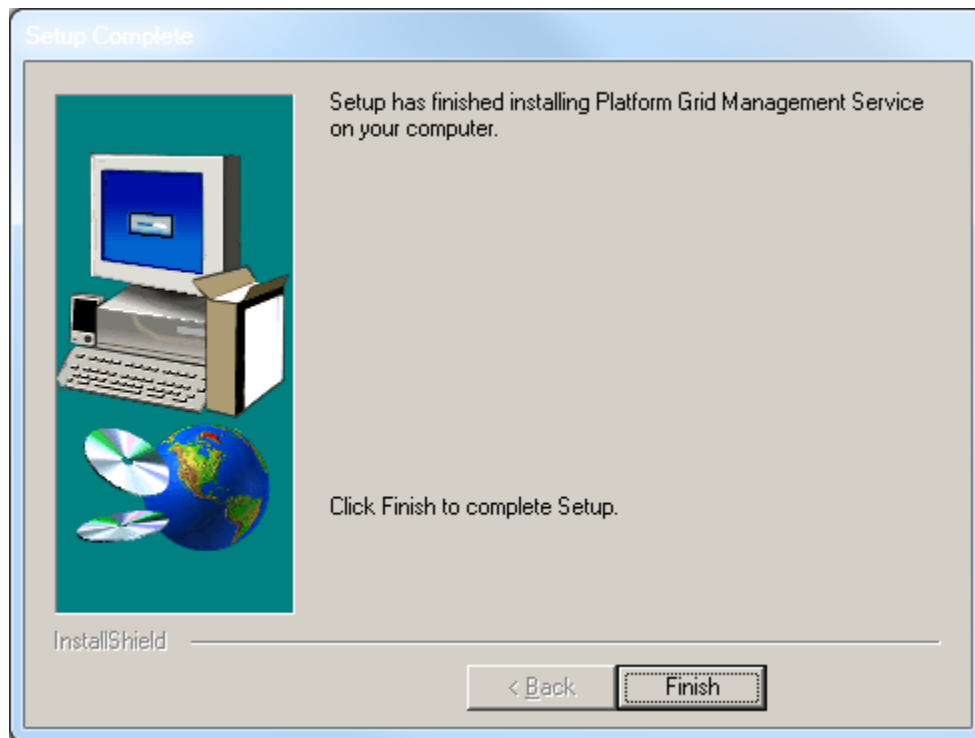


9. Click **Next** to start the installation.

In the Platform Grid Management Service Group dialog, the installation progress indicator displays.



The Setup Complete dialog opens.



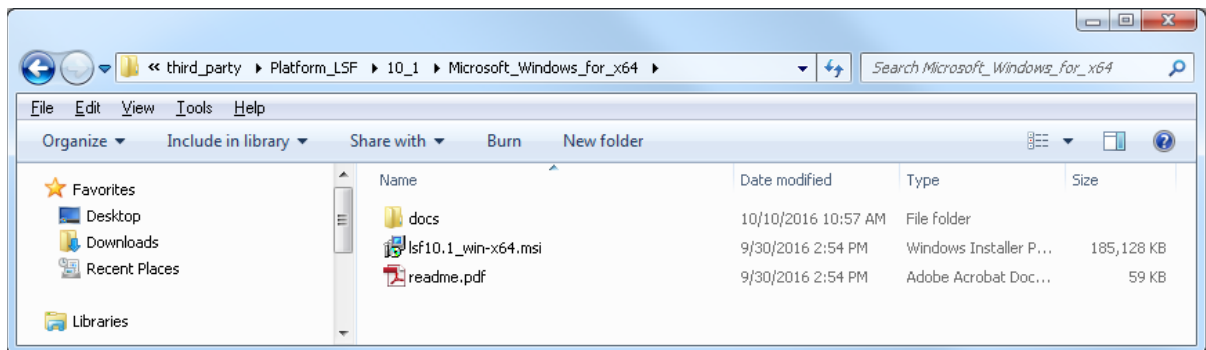
10. Click **Finish**.

Testing the Installation

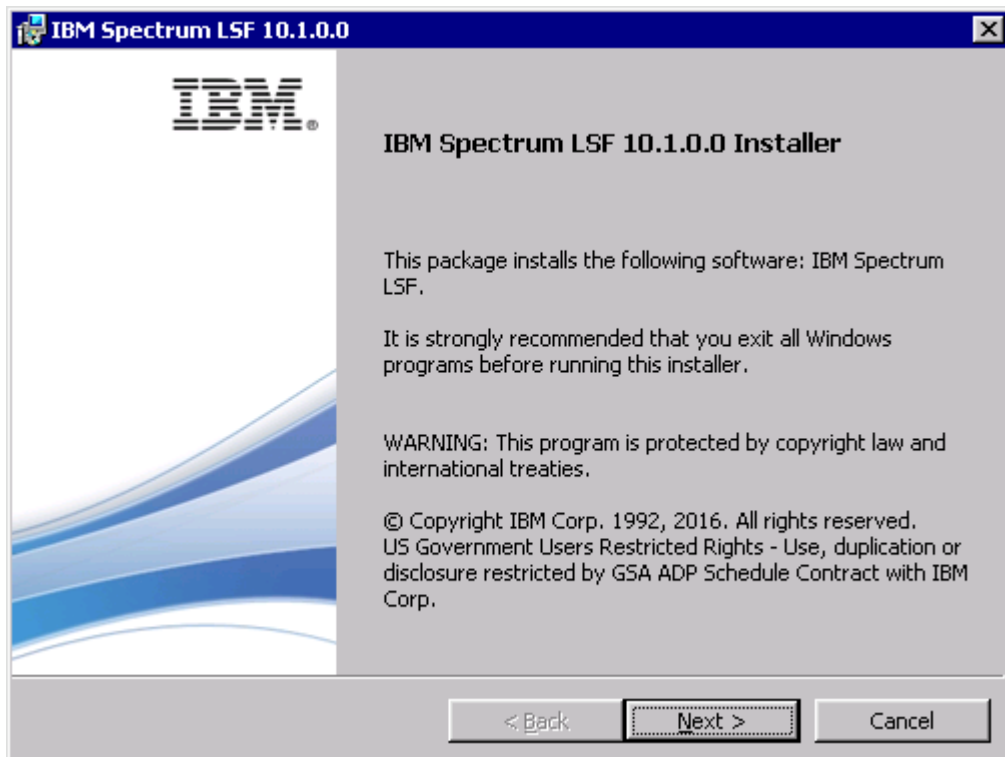
Once SAS has been installed and configured, the SAS Grid Manager Plug-in in the SAS Management Console can be used to test and confirm that the Grid Management Service is working.

Chapter 4 - Installing LSF on Grid Nodes or SAS Foundation Grid Clients

1. After logging on as an administrator on the machine that becomes a grid node or a grid client, run the Windows installer (lsf10.1*.msi file) located in the SAS Software Depot. For 9.4, the location is in the `third_party/Platform_LSF/10_1` directory under the sub-directory named for the operating system you are installing on.

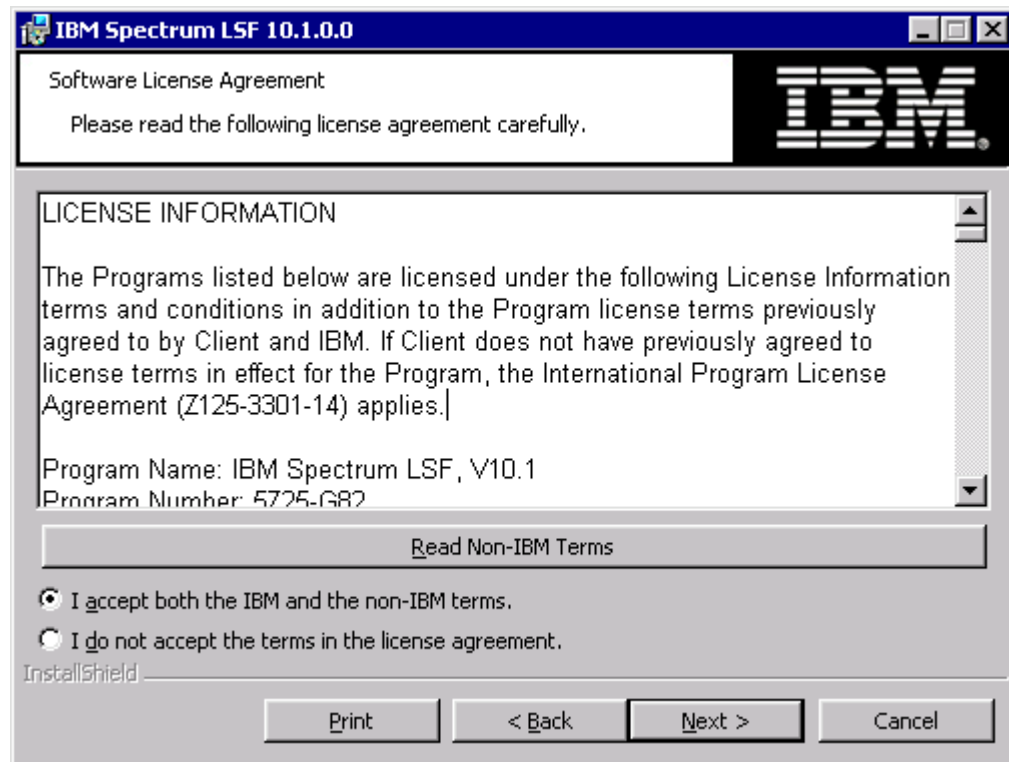


The Platform LSF Installer dialog displays.



2. Click **Next** to continue.

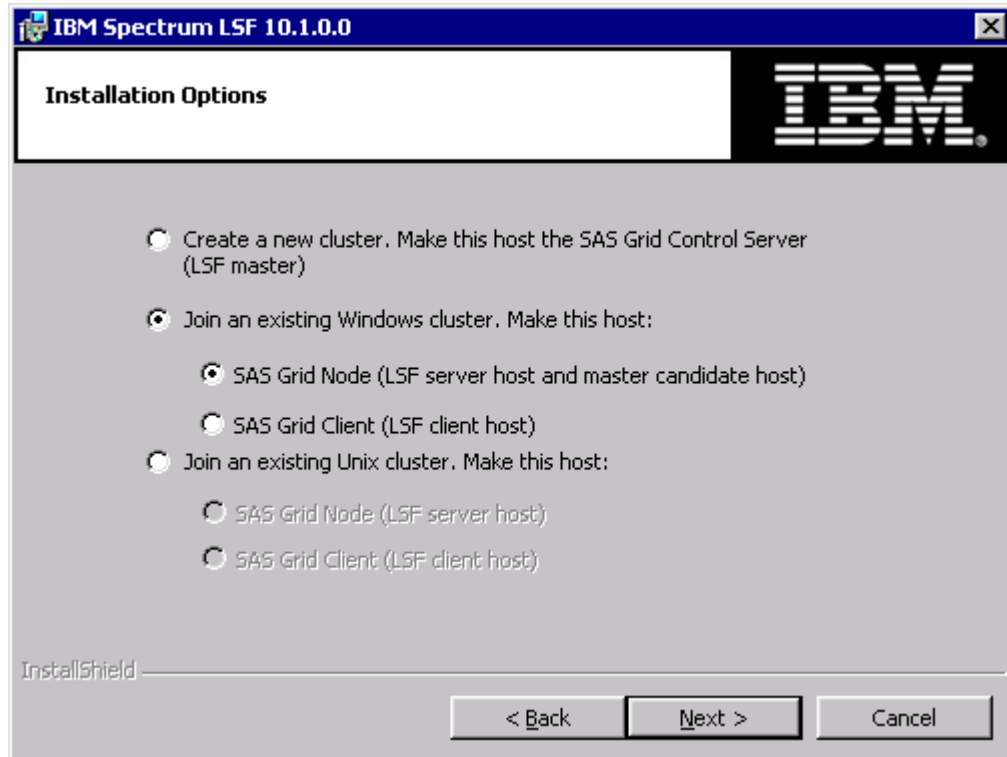
The Platform LSF End User Software License Agreement dialog displays.



3. Select **Accept** after reading the agreement (the default is **Do not accept**) and click **Next** to continue.

You will be asked if this is a new cluster or if you are adding a host to an existing cluster. Since you are installing on the grid node machine, we will be adding to an already existing cluster.

The Platform LSF/Installation Options dialog displays.



4. Select, **Join an existing Windows cluster.**

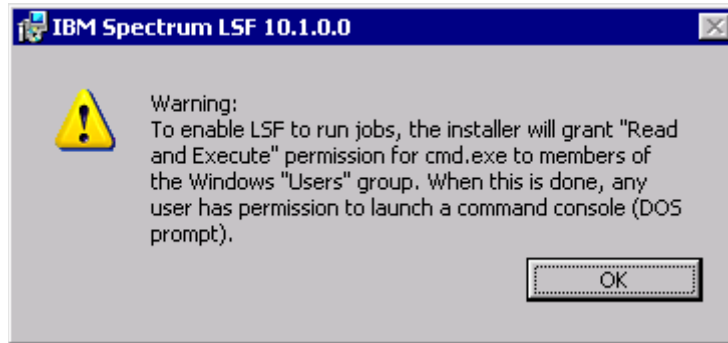
There are two types of hosts you can add. One is a **SAS Grid Node** which allows the machine to participate in the cluster for executing jobs. The other is a **SAS Grid Client** for the purpose of submitting jobs to the grid without the machine participating as a grid node. You can add the Windows machine to an existing cluster on Windows or a UNIX cluster. For a UNIX cluster, do not add a **SAS Grid Node**. The option is provided by the LSF installer but you should only add a **SAS Grid Client** to an existing UNIX cluster with a Windows installer. For **SAS Grid Client** you will need to set the environment variable LSF_FULL_VERSION to 10.1 as this is not automatically executed like it is for a **SAS Grid Node**.

If you are installing Platform LSF on a machine that you want to be a grid node, then select **SAS Grid Node** to an existing Windows cluster. Click **Next** to continue.

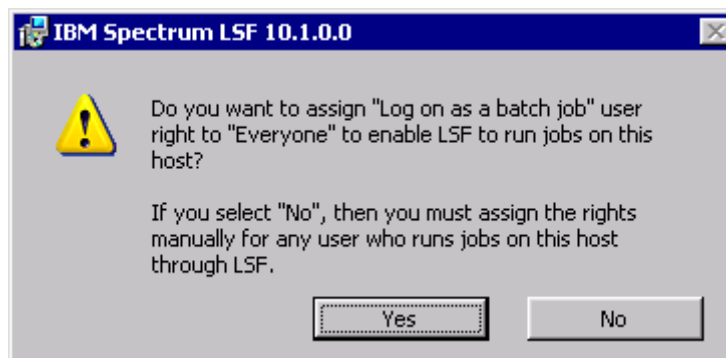
If you are installing Platform LSF on a machine solely for the purpose of being a SAS Foundation grid client, then select **SAS Grid Client**. Follow the step below if the existing cluster is on a Windows-based machine. Click **Next** to continue, then go to step 9.

5. If you are installing SAS Grid Client on an existing UNIX cluster. Click **Next** to continue, then go to step 11.

The Platform LSF dialog (which assigns user rights) displays.

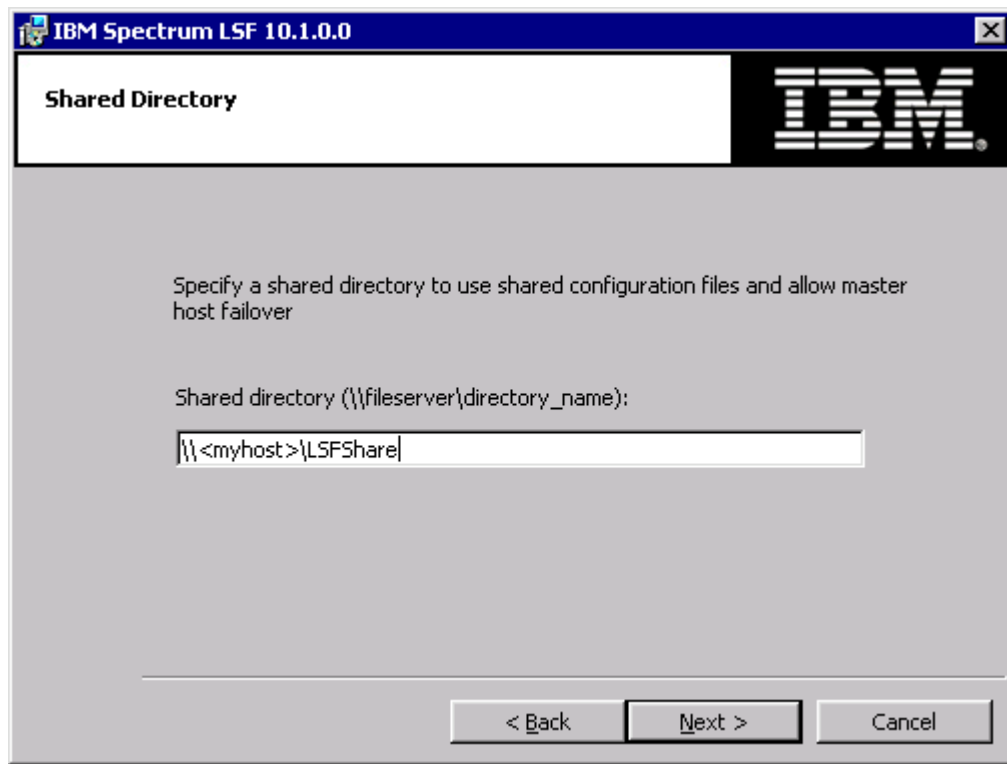


6. Select **OK**.
7. To run **SAS Grid Node** on an existing Windows cluster, you will be asked if you want to assign rights to allow LSF to run jobs on the host.



8. Select **Yes**.

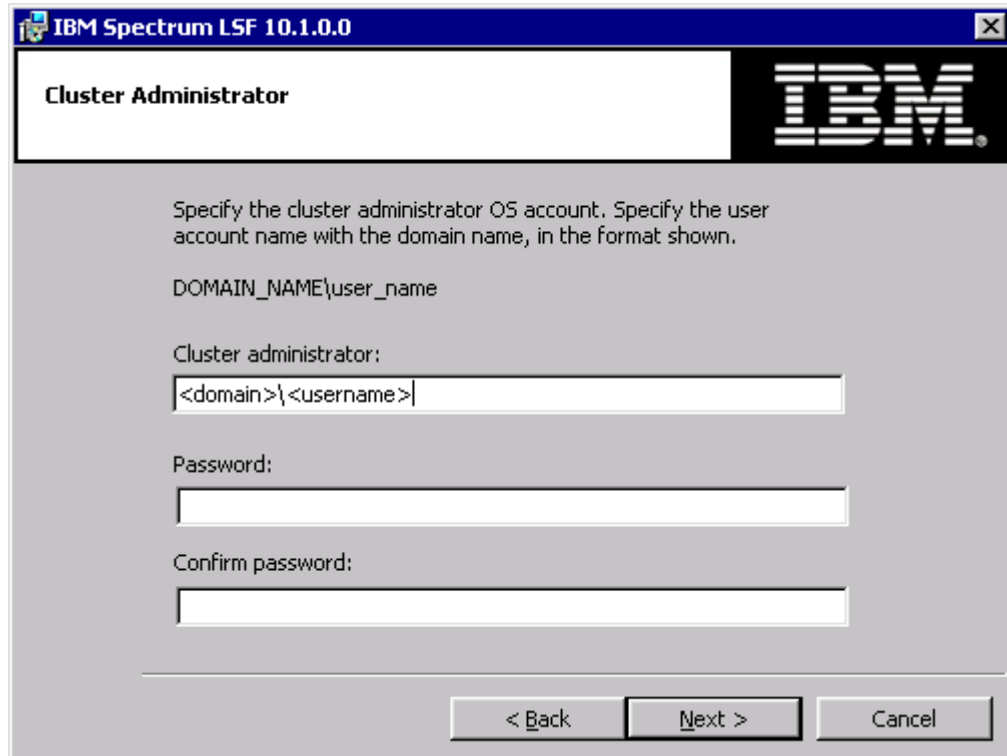
The Platform LSF/Shared Directory displays, you are asked for the shared directory location for configuration.



Enter the share path created when the grid control server was installed (see step 14 of "Chapter 2 - Installing Process Manager and LSF"). This will allow LSF to access information from the configuration files for the cluster.

9. Click **Next**.

The Platform LSF/Cluster Administrator dialog displays.



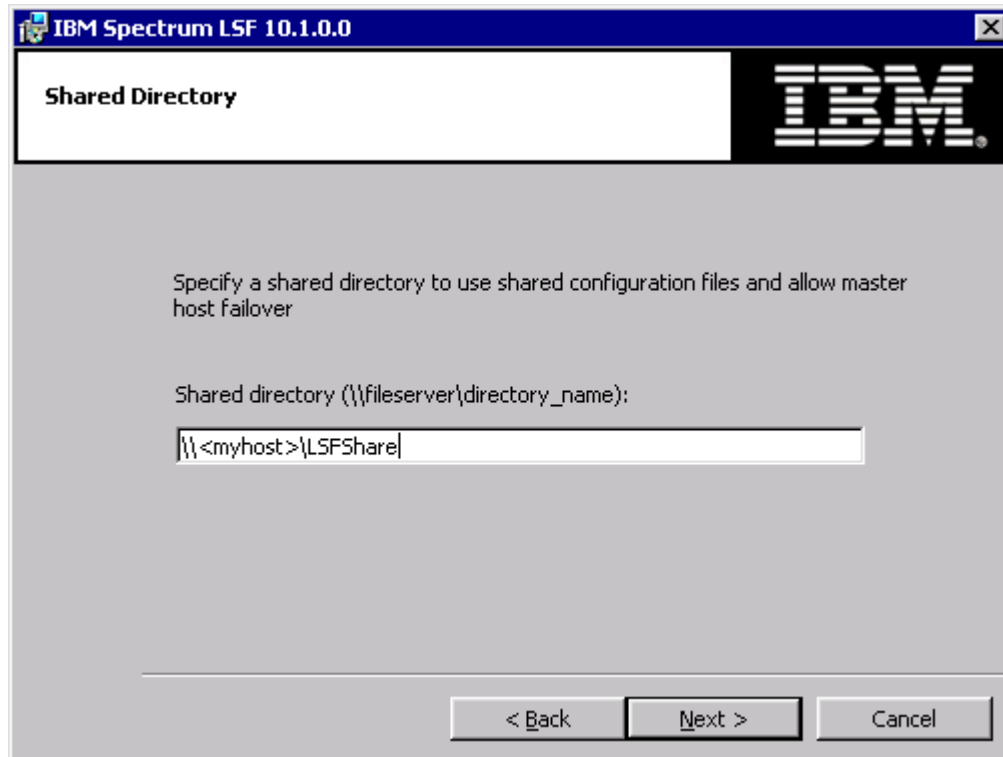
The image shows a Windows-style dialog box titled "IBM Spectrum LSF 10.1.0.0". The dialog has a header bar with the title and a close button. Below the header, the title "Cluster Administrator" is displayed on the left, and the IBM logo is on the right. The main area of the dialog contains the following text: "Specify the cluster administrator OS account. Specify the user account name with the domain name, in the format shown." Below this is the text "DOMAIN_NAME\user_name". Then, there is a label "Cluster administrator:" followed by a text input field containing the placeholder text "<domain>\<username>". Below that is a label "Password:" followed by a password input field. Then, there is a label "Confirm password:" followed by another password input field. At the bottom of the dialog, there are three buttons: "< Back", "Next >", and "Cancel".

If the domain Primary LSF Administrator is not automatically populated, enter the domain Primary LSF administrator you created in the SAS pre-installation checklist (usually `<domain>\sas` where `<domain>` is your Windows domain name).

10. Click **Next** to continue, then go to step 17.

The Platform LSF/Shared Directory dialog displays.

To run SAS Grid Client on an existing Windows cluster, you are asked for the shared directory location for configuration.



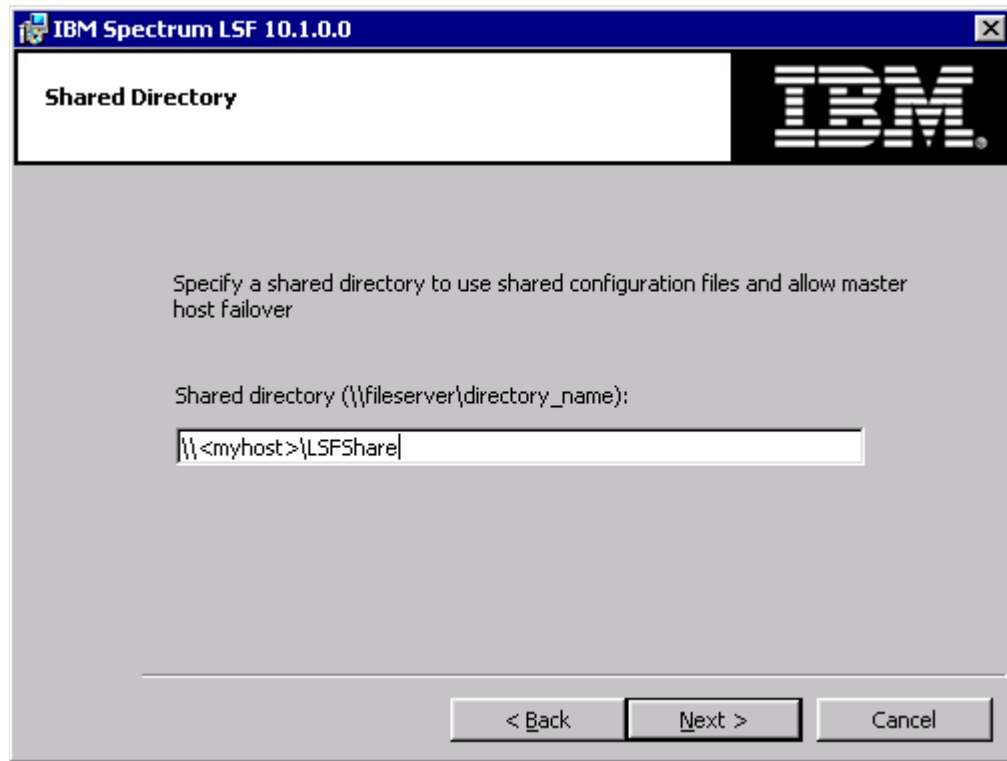
Enter the share path created when the grid control server was installed (see "Chapter 2 - Installing Process Manager and LSF"). This will allow LSF to access information from the configuration files for the cluster.

11. Click **Next** to continue, then go to step 17.

The Platform LSF Shared Directory dialog displays.

For **SAS Grid Client** in an existing UNIX cluster, you are asked for the shared directory location for configuration.


Note: This is not the UNIX grid configuration directory.



12. Click **Next**.

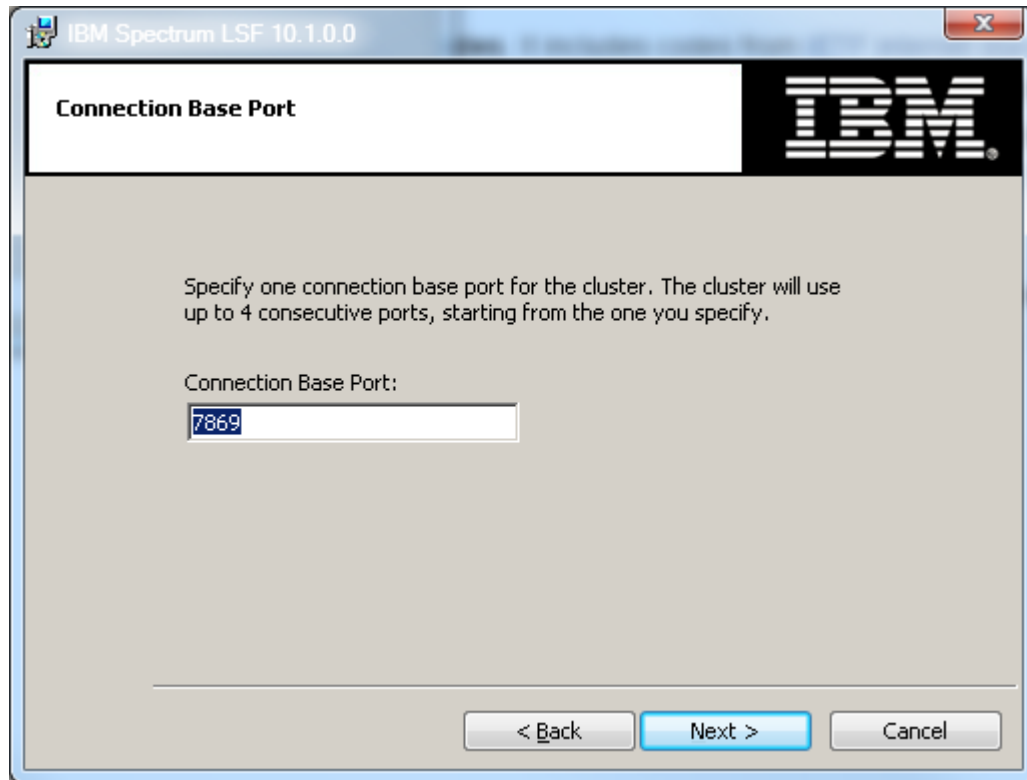
The Platform/LSF Server Hosts displays by default.

Note: *Duplicating the LSF_SERVER_HOSTS value from the UNIX lsf.conf file as the value in this case is usually sufficient.*



13. Click **Next**.


The Platform LSF Connection Base Port dialog displays. Provide the base port in the UNIX cluster field.




14. Click **Next** to continue.

The Platform LSF/Warning dialog displays.

It will warn you that the configuration does not have the LSF_USER_DOMAIN set if it hasn't been listed in the UNIX cluster file. The LSF_USER_DOMAIN will need to be set to the Windows domain used on the Windows client.



If the host is not in the cluster file as a client or you are not using FLOATING_CLIENTS, then you may see the following warning:



15. Select **Cancel** to cancel the operation and correct the problem.
16. Select **Continue** to proceed with the installation.

The Platform LSF/Installation Directory dialog displays. The three selections made in step 3 culminate to bring you to this dialog.

17. Enter your preferred path for installation location or accept the default path offered.



18. Click **Next** to continue.

This is the same path that was specified earlier in “Chapter 2 - Installing Process Manager and LSF”.

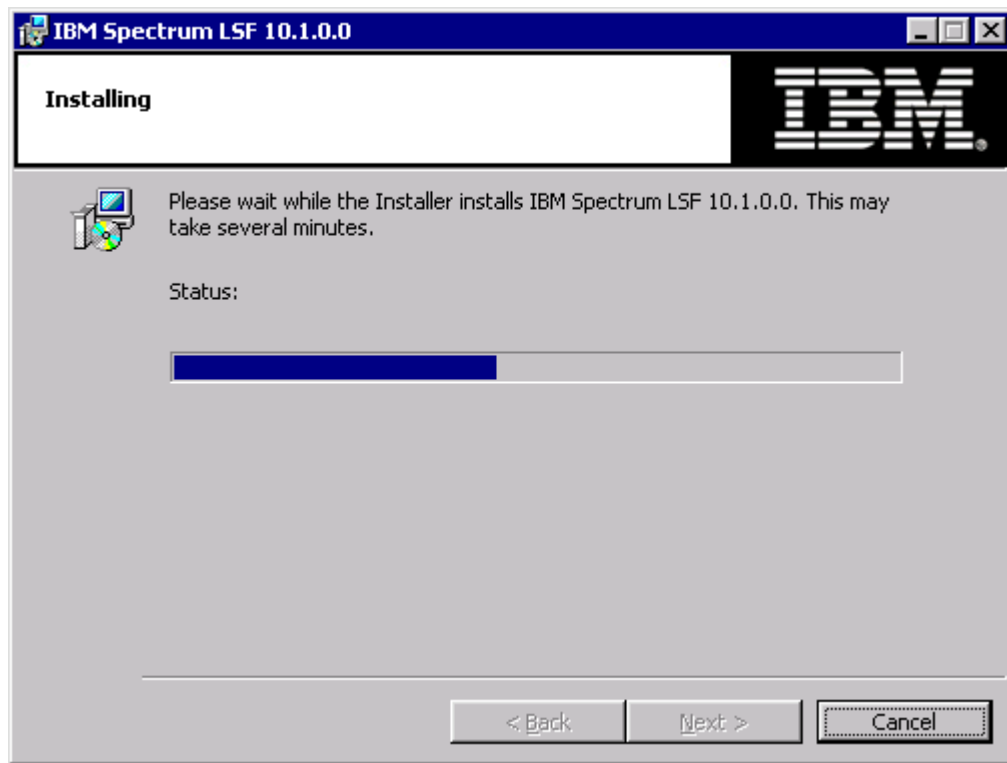
Note: Make sure the installation directory is not the same as the shared directory in the previous step. Otherwise, you will receive an error during the installation.

LSF has completed gathering all the information it needs and displays a summary before starting the install as shown in the Platform LSF/Summary dialog below.

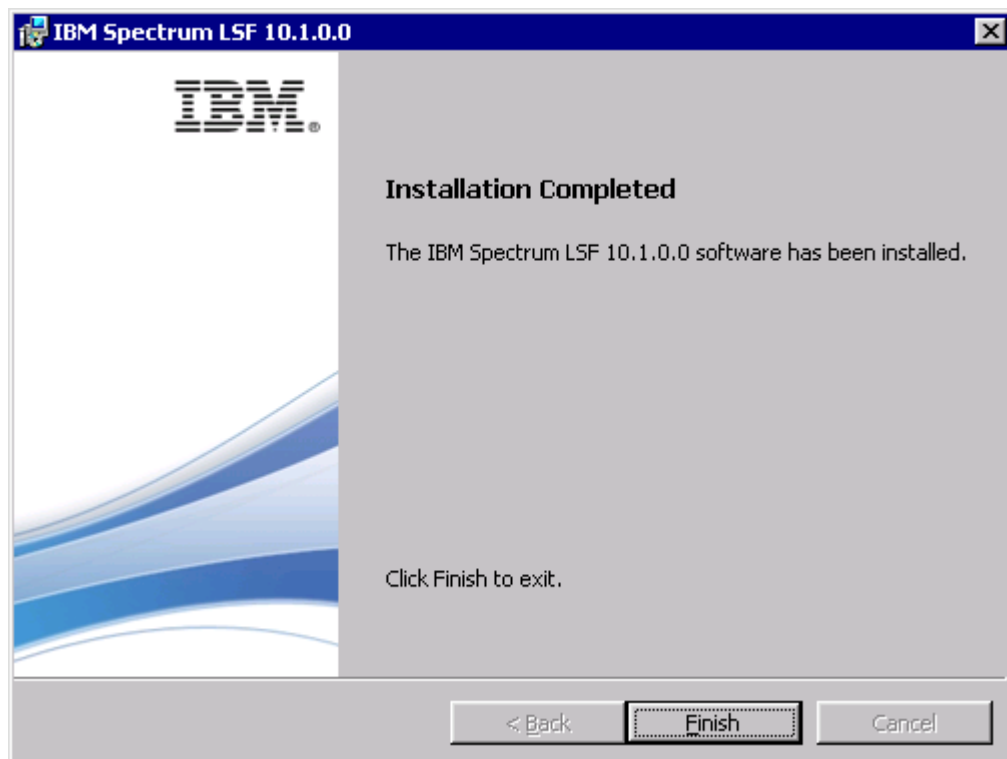


19. Click **Install** to start the installation.

The Platform LSF/Installing dialog displays.



The Platform LSF/Installation Completed dialog displays.



20. Click **Finish**.

21. **Reboot the new grid node machine** and run the **lsfrestart** command on the grid control server so that the cluster is updated about the new grid node machine. To run the command, start a Command Prompt window using **Start→Programs→Accessories→Command Prompt** and execute the command.

Testing the Installation

Once the system has rebooted, you can follow these steps to make sure LSF on the cluster is operating properly.

1. Log onto the grid control server as an LSF administrator or user.
2. Open a command prompt to execute the command in the following steps.
3. Run the command **lshosts**. This displays static information about the grid control server and all grid nodes.
4. Run the command **lsload**. This displays dynamic information about the grid control server and all grid nodes.
5. Run the command **bsub sleep 1000**. This submits a job to the cluster. Repeat this command once for each node in the cluster.
6. Run the command **bjobs**. This displays the job information. As you repeat this command, you see the job go from **PEND**, to **RUN**, to being removed from the queue.

The following screen capture is sample output of a heterogeneous cluster where the grid control server (where LSF is the master machine) is **D15003.testgrid.com** running the Windows operating system and the grid nodes are **grid1.testgrid.com**, **grid2.testgrid.com**, and **grid3.testgrid.com**, all running Linux. For a homogeneous cluster, all machines will be running the same operating system.

```

[lsfadmin@grid3 ~]$ . /grid/conf/profile.lsf
[lsfadmin@grid3 ~]$ lshosts
HOST_NAME      type      model  cpuf  ncpus  maxmem  maxswp  server  RESOURCES
d15003.test    NTX86     PC450   13.2    1  2038M  3926M   Yes (SASMain nt)
grid1.testg    LINUX86   Opteron8 60.0    1   250M   511M   Yes (SASMain linux)
grid2.testg    LINUX86   Opteron8 60.0    1   250M   511M   Yes (SASMain linux)
grid3.testg    LINUX86   Opteron8 60.0    1   250M   511M   Yes (SASMain linux)
[lsfadmin@grid3 ~]$ lsload
HOST_NAME      status  r15s  r1m  r15m  ut    pg    ls    it    tmp    swp    mem
grid2.testgrid.  ok     0.1   0.0   0.5   2%    9.3   0     8 9888M 512M 211M
grid1.testgrid.  ok     0.5   0.0   0.4   2%    8.2   0     7 9736M 512M 210M
grid3.testgrid.  ok     0.6   0.7   0.8  17%  464.5 1     0 9888M 510M 168M
d15003.testgrid  ok     1.0   1.3   0.7  55%  16.4 1     0 8906M 2741M 877M
[lsfadmin@grid3 ~]$ bhosts
HOST_NAME      STATUS    JL/U    MAX  NJOBS    RUN  SSUSP  USUSP  RSV
d15003.testgrid.co ok        -      1     0     0     0     0     0
grid1.testgrid.com ok        -      1     0     0     0     0     0
grid2.testgrid.com ok        -      1     0     0     0     0     0
grid3.testgrid.com ok        -      1     0     0     0     0     0
[lsfadmin@grid3 ~]$

```

Adding Nodes or SAS Foundation Clients to the Grid

Adding new grid nodes or SAS Foundation clients to an existing cluster is identical to adding them to a new cluster. Follow the steps presented earlier in this chapter.

Converting a Grid Node Machine to a Grid Client

When you run SAS Foundation solely for the purposes of submitting jobs to the grid without allowing that machine to participate as a grid node and you installed Platform LSF on a machine as an “LSF Server” host type, then the following steps will prevent jobs from running on the machine. Essentially, this makes it an “LSF Client” machine. This is accomplished by changing the state of a machine to ‘closed’. To change a machine’s state to ‘closed’, do the following:

1. Log on as the LSF Administrator.
2. Run the command **badmin hclose** *<host_name>*.

When you run the **bhosts** command, the host should display a status of ‘closed’.

Chapter 5 - LSF Quick Reference

Command	Description										
lsid	Displays version number, cluster name, and the grid control server (master host) name. Useful to see if the grid daemons are running and if running in SAS mode.										
lshosts	Displays information about the hosts recognized by LSF along with their static resource information.										
lsload	Displays the dynamic resource information for the hosts in the grid (cluster).										
bhosts	Displays batch information about all hosts in the grid (cluster).										
bjobs	Displays information about current user's LSF jobs.										
lsfstartup	Starts the LIM, RES, sbatchd, and mbatchd daemons on all hosts in the cluster. Must be run as root and all hosts must be running rsh or ssh daemons.										
lsfrestart	Restarts the LIM, RES, sbatchd, and mbatchd daemons on all hosts in the cluster. Must be run as root and all hosts must be running rsh or ssh daemons.										
lsfshutdown	Shuts down the LIM, RES, sbatchd, and mbatchd daemons on all hosts in the cluster. Must be run as root and all hosts must be running rsh or ssh daemons.										
lsadmin	<p>Administrative tool for LSF available to LSF administrators. Useful subcommands are:</p> <table> <tr> <td>reconfig</td><td>Restarts all LIMs in the cluster to read any changes in the configuration files.</td></tr> <tr> <td>limstartup</td><td>Starts LIM on the local host</td></tr> <tr> <td>limrestart</td><td>Restarts LIM on the local host</td></tr> <tr> <td>resstartup</td><td>Starts RES on local host</td></tr> <tr> <td>resrestart</td><td>Restarts RES on local host</td></tr> </table>	reconfig	Restarts all LIMs in the cluster to read any changes in the configuration files.	limstartup	Starts LIM on the local host	limrestart	Restarts LIM on the local host	resstartup	Starts RES on local host	resrestart	Restarts RES on local host
reconfig	Restarts all LIMs in the cluster to read any changes in the configuration files.										
limstartup	Starts LIM on the local host										
limrestart	Restarts LIM on the local host										
resstartup	Starts RES on local host										
resrestart	Restarts RES on local host										
bhist	<p>Displays historical information about jobs. Useful parameters are:</p> <table> <tr> <td>-p -r -d -a</td><td>Displays information about specific jobs (<u>p</u>ending, <u>r</u>unning, <u>d</u>one, or <u>a</u>ll).</td></tr> <tr> <td>-l</td><td>Display in long format.</td></tr> <tr> <td>-u <user> all</td><td>Displays job for specified or all users.</td></tr> <tr> <td><job ID></td><td>Displays only specified job information.</td></tr> </table>	-p -r -d -a	Displays information about specific jobs (<u>p</u> ending, <u>r</u> unning, <u>d</u> one, or <u>a</u> ll).	-l	Display in long format.	-u <user> all	Displays job for specified or all users.	<job ID>	Displays only specified job information.		
-p -r -d -a	Displays information about specific jobs (<u>p</u> ending, <u>r</u> unning, <u>d</u> one, or <u>a</u> ll).										
-l	Display in long format.										
-u <user> all	Displays job for specified or all users.										
<job ID>	Displays only specified job information.										

badmin	Administrative tool for LSF's batch processing facility available to LSF administrators. Useful subcommands are:	
	reconfig	Reconfigures the batch facility without restarting sbatchd or mbatchd to read any changes in the configuration files.
	hstartup	Starts sbatchd on the local host
	hrestart	Restarts sbatchd on the local host
	mbdrestart	Restarts mbatchd. Needs to be done when new hosts are added to the grid (cluster).
	hclose <host>	Closes a host preventing it from running jobs.
	hopen <host>	Opens a host to allow it to run jobs.
bsub	Submit a job to the grid. Useful parameters are:	
	-I	Interactive. Remote output displayed locally.
	-m	Submit to a specific host.
	-R "res_req"	Submit with specified resource

The LSF commands shown in this section include examples of typical output. The output you see differs according to your local configuration.

The commands are described briefly so that you can easily use them to check your LSF installation. See the *LSF Reference* for complete usage and command options. You can use these commands on any LSF host. If you get proper output from these commands, your cluster is ready to use. If your output from the commands discussed in this section has errors, see the *LSF Reference* for help.

Check Cluster Configuration (lsadmin)

```
lsadmin ckconfig -v
```

The **lsadmin** command controls the operation of an LSF cluster and LSF configuration files. The -v flag displays detailed information about the LSF configuration:

```
lsadmin ckconfig -v
```

```
Checking configuration files ...
```

```
Checking configuration files ...
```

```
EGO 3.4.0 build 238688, Sep 13 2016
Copyright International Business Machines Corp. 1992, 2016.
US Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corp.
```

```
    binary type: win-x64
Reading configuration from
\\myhost\LSFShare\conf\ego\sas_cluster\kernel\ego.conf
Oct 26 10:24:52 2016 12560:9624 5 3.4.0 EGO 3.4.0 build 238688, Sep 13 2016
Copyright International Business Machines Corp. 1992, 2016.
```

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

binary type: win-x64

```
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 Lim starting...
Oct 26 10:24:52 2016 12560:9624 5 3.4.0 main: The specified EGO_LOG_MASK
value <(null)> is not valid. The system will use the defau
lt value <LOG_WARNING>.
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 LIM is running in advanced workload
execution mode.
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 Master LIM is not running in
EGO_DISABLE_UNRESOLVABLE_HOST mode.
Oct 26 10:24:52 2016 12560:9624 5 3.4.0 C:\LSF_10.1\10.1\etc\lim.exe -C
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 initEntitlement: EGO_AUDIT_MAX_SIZE
was not set. Default value <100> will be used.
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 initEntitlement:
EGO_AUDIT_MAX_ROTATE was not set. Default value <20> will be used.
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 LIM is running as IBM Spectrum LSF
Standard Edition.
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 LIM is running as IBM Spectrum
Conductor Edition.
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 LIM is running as EGO Edition.
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 reCheckClass: numhosts 1 so reset
exchIntvl to 15.00
Oct 26 10:24:52 2016 12560:9624 6 3.4.0 Checking Done.
```

No errors found.

The messages shown are typical of normal output from **lsadmin ckconfig -v**. Other messages may indicate problems with your LSF configuration. See the *LSF Reference* for help with some common configuration errors.

Find Out Cluster Status (lsid and lsload)

lsid

Notifies you if your LSF environment is set up properly. **lsid** displays the current LSF version number, cluster name, and host name of the current grid control server (LSF master host) for your cluster. The grid control server (LSF master host) name displayed by **lsid** may vary, but it is usually the first host configured in the `Hosts` section of `LSF_CONFDIR/lsf.cluster.cluster_name`.

lsid

```
IBM Spectrum LSF Standard 10.1.0.0, Sep 13 2016
Copyright International Business Machines Corp. 1992, 2016.
US Government Users Restricted Rights - Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM Corp.
```

```
My cluster name is sas_cluster
My master name is myhost
Cluster in ISV mode : SAS
```

lsload

Displays the current load levels of the cluster. The output contains one line for each host in the cluster. The status should be acceptable for all hosts in your cluster. For example:

```
lsload
```

```
HOST_NAME status r15s r1m r15m ut pg ls it tmp swp mem
hosta ok 0.0 0.0 0.0 6% 0.2 2 1365 97M 65M 29M
hostb -ok 0.0 0.0 0.0 9% 0.0 4 1 130M 319M 12M
hostc ok 2.5 2.2 1.9 64% 56.7 50 0 929M 931M 4000M
hostd ok 0.2 0.2 0.2 1% 0.0 0 367 93M 86M 50M
hoste busy *6.0 2.2 1.9 64% 56.7 50 0 929M 931M 4000M
hostf unavail
```

A `busy` status is shown for hosts with any load index beyond its configured thresholds. An asterisk (*) marks load indices that are beyond their thresholds, causing the host status to be busy. A minus sign (-) in front of the value `ok` means that RES is not running on that host.

If you see the message

```
LIM is down
```

or

```
LIM is not responding
```

after starting or reconfiguring LSF, wait a few seconds and try `lsload` again to give the LIMs time to initialize. `lsload` also shows if LSF is licensed for the host. If you see the message

```
Host does not have a software license
```

you must install a valid LSF license or make sure that the license server is running properly.

There are also a couple of other useful commands:

- The `lshosts` command displays configuration information for LSF hosts and their static resource information.
- The `lsinfo` command displays cluster configuration information about resources, host types, and host models.

Check LSF Batch Configuration (badmin)

```
badmin ckconfig -v
```

The `badmin` command controls and monitors the operation of the LSF Batch system. Use the `badmin ckconfig` command to check the LSF Batch configuration files. The `-v` flag displays detailed information about the configuration:

```
badmin ckconfig -v
Checking configuration files ...
-----
No errors found.
```

The messages shown above are the normal output from `badmin ckconfig -v`. Other messages may indicate problems with the Platform LSF Batch configuration. See the *LSF Reference* for help with some common configuration errors.

Find Out LSF Batch System Status (*bhosts* and *bqueues*)

bhosts

The `bhosts` command tells you if LSF Batch is running properly. `Bhosts` displays the status and other details about the grid nodes (LSF Batch server hosts) in the cluster:

- maximum number of job slots allowed by a single user
- total number of jobs in the system, jobs running, jobs suspended by users, and jobs suspended by the system
- total number of reserved job slots

The status should be `ok` acceptable for all grid nodes (hosts) in your cluster. For example:

```
bhosts

HOST_NAME STATUS JL/U MAX NJOBS RUN SSUSP USUSP RSV
hosta ok - - 0 0 0 0 0
hostb ok - - 0 0 0 0 0
hostc ok - - 0 0 0 0 0
hostd ok - - 0 0 0 0 0
```

If you see the message

```
lsbatch daemons not responding
```

after starting or reconfiguring LSF, wait a few seconds and try `bhosts` again to give the SBDs time to initialize.

bqueues

LSF Batch queues organized jobs with different priorities and different scheduling policies. The `bqueues` command displays available queues and their configuration parameters. For a queue to accept and dispatch jobs, the status should be `Open:Active`.

```
bqueues

QUEUE_NAME PRIO STATUS MAX JL/U JL/P JL/H NJOBS PEND RUN SUSP
owners 43 Open:Active - 6 - - 0 0 0 0
priority 43 Open:Active - - - - 0 0 0 0
night 40 Open:Active - - - - 0 0 0 0
chkpnt_rerun_qu 40 Open:Active - - - - 0 0 0 0
short 35 Open:Active - - - - 0 0 0 0
license 33 Open:Active - - - - 0 0 0 0
normal 30 Open:Active - - - - 0 0 0 0
idle 20 Open:Active - - - - 0 0 0 0
```

The queue information displayed by `bqueues` is configured in `lsb.queues`. Eight queues are defined by default in `lsb.queues`. Modify this file to add, delete, or change queues.

bqueues -l

To see more detailed queue information, use `bqueues -l`:

```
bqueues -l normal
```

```
QUEUE: normal
-- For normal low priority jobs, running only if hosts are lightly
loaded. This is the default queue.
PARAMETERS/STATISTICS
PRIO NICE STATUS MAX JL/U JL/P JL/H NJOBS PEND RUN SSUSP USUSP RSV
30 20 Open:Active - - - - 8 8 0 0 0 0
STACKLIMIT MEMLIMIT
2048 K 5000 K
SCHEDULING PARAMETERS
r15s r1m r15m ut pg io ls it tmp swp mem
loadSched - - - - - - - - - -
loadStop - - - - - - - - - -
USERS: all users
HOSTS: all hosts used by the LSF Batch system
```

`bqueues -l` shows the following kinds of information about the queue:

- what kinds of jobs are meant to run on the queue
- resource usage limits
- hosts and users are able to use the queue scheduling threshold values:
 - `loadSched` is the threshold for LSF to dispatch a job automatically
 - `loadStop` is the threshold for LSF to suspend a job automatically

Other useful commands include:

- The `bparams` command displays information about the LSF Batch configuration parameters.
- The `bhist` command displays historical information about jobs.

For More Information

See the *LSF Administrator's Guide* for more information about seeing the status of your cluster.

See the *LSF Reference* for detailed information about the commands described in this section.

See *Administering Process Manager* for detailed information about Process Manager configuration and maintenance.

These documents are also available at

<http://support.sas.com/rnd/scalability/platform/index.html>.

Appendix – Multiple Host Installation

Pre-Installation Requirements

1. Microsoft Installer (MSI) version 2.0 or later is required.
2. Check the permissions of the Primary LSF administrator (owns all the configuration and log files; for example, *<domain>\sas*). **Important:** The account is supposed to belong to the Local Administrators group on each host and have the following privileges for this account on each LSF host:
 - Act as part of the operating system
 - Adjust memory quotas for a process (increase quotas)
 - Debug programs
 - Log on as a service
 - Replace a process level token
3. Download the free PsExec from Microsoft. You can download the PsExec onto the master host, and then simply copy `PsExec.exe` onto the execution path (for example, `C:\WINDOWS\System32`) of the host. You will install all compute hosts from this host.
4. Make sure the Remote Registry Service is started on every compute host. This service should have been started automatically by default.

Installation Procedure

Follow these steps to install LSF10.1/SAS on multiple hosts.

1. Install the grid control server. Log on as the Primary LSF administrator on the grid control server and run the installer (for example, `lsf10.1_win-x64.msi`). Make a note of the base port, `sbatchd`, `res`, and `mbatchd` port for the following steps.
2. Edit the `install.bat` file located in LSF top install folder (for example, `C:\LSF_10.1`) as directed in the comments in the `install.bat` file.
3. Execute the `install.bat` command to inconspicuously install on all grid nodes.
4. Run **lsfstartup** to start the LSF cluster.

Testing the Installation

When the system has rebooted, follow these steps to ensure LSF on the grid control server is operating properly.

1. Log onto the machine as one of the users you added to LSF using the `lspasswd` command. These are usually *<domain>\sas*, *<domain>\sasadm*, *<domain>\sastrust*, or *<domain>\sasdmo*.
2. Open a Command Prompt window by selecting **Start→Programs→Accessories→Command Prompt**.
3. Run the command **lsid**. This displays the cluster name and the grid control server (LSF master machine) name.

4. Run the command **lshosts**. This displays static information about the grid control server (LSF master machine).
5. Run the command **lsload**. This displays dynamic information about the grid control server (LSF master machine).
6. Run the command **bsub sleep 100**. This submits a job to the cluster.
7. Run the command **bjobs**. This displays the job information. As you repeat this command, you should see the job go from PEND, to RUN, to being removed from the queue.



SAS is the leader in [business analytics](#) software and services, and the largest independent vendor in the business intelligence market. Through innovative solutions, SAS helps customers at more than 70,000 sites improve performance and deliver value by making better decisions faster. Since 1976 SAS has been giving customers around the world THE POWER TO KNOW®.