



ANSYS Electromagnetics Suite 2021 R1 Unix/Linux Installation Guide



ANSYS, Inc.
Southpointe
2600 ANSYS Drive
Canonsburg, PA 15317
ansysinfo@ansys.com
<https://www.ansys.com>
(T) 724-746-3304
(F) 724-514-9494

Release 2021 R1
January 2021

ANSYS, Inc. and
ANSYS Europe,
Ltd. are UL
registered ISO
9001:2015
companies.

Copyright and Trademark Information

© 1986-2021 ANSYS, Inc. Unauthorized use, distribution or duplication is prohibited.

ANSYS, ANSYS Workbench, AUTODYN, CFX, FLUENT and any and all ANSYS, Inc. brand, product, service and feature names, logos and slogans are registered trademarks or trademarks of ANSYS, Inc. or its subsidiaries located in the United States or other countries. ICEM CFD is a trademark used by ANSYS, Inc. under license. All other brand, product, service and feature names or trademarks are the property of their respective owners. FLEXlm and FLEXnet are trademarks of Flexera Software LLC.

Disclaimer Notice

THIS ANSYS SOFTWARE PRODUCT AND PROGRAM DOCUMENTATION INCLUDE TRADE SECRETS AND ARE CONFIDENTIAL AND PROPRIETARY PRODUCTS OF ANSYS, INC., ITS SUBSIDIARIES, OR LICENSORS. The software products and documentation are furnished by ANSYS, Inc., its subsidiaries, or affiliates under a software license agreement that contains provisions concerning non-disclosure, copying, length and nature of use, compliance with exporting laws, warranties, disclaimers, limitations of liability, and remedies, and other provisions. The software products and documentation may be used, disclosed, transferred, or copied only in accordance with the terms and conditions of that software license agreement.

ANSYS, Inc. and ANSYS Europe, Ltd. are UL registered ISO 9001: 2015 companies.

U.S. Government Rights

For U.S. Government users, except as specifically granted by the ANSYS, Inc. software license agreement, the use, duplication, or disclosure by the United States Government is subject to restrictions stated in the ANSYS, Inc. software license agreement and FAR 12.212 (for non-DOD licenses).

Third-Party Software

See the legal information in the product help files for the complete Legal Notice for ANSYS proprietary software and third-party software. If you are unable to access the Legal Notice, please contact ANSYS, Inc.

Table of Contents

Table of Contents	Contents-1
1 - Getting Started	1-1
Installation Prerequisites	1-1
ANSYS Electromagnetics Suite Software	1-1
License File	1-1
License Server	1-1
Hardware and Software Requirements	1-2
Operating System Specific Prerequisites	1-2
Red Hat Enterprise Linux	1-2
License Server	1-3
Required Libraries to Install License Server	1-3
Required Libraries to Run License Server	1-3
Additional Prerequisites	1-4
Additional Resources	1-5
2 - Installing ANSYS Electromagnetics Software	2-1
Obtaining a License File	2-1
Verifying that Enough Disk Space is Available	2-1
Installing the Software	2-1
Integrating ANSYS Electromagnetics Products with ANSYS Work- bench	2-5
3 - High Performance Computing	3-1
Options for Distributed Analysis	3-1
Semaphores on Linux	3-2
Recommended settings	3-2
Displaying semaphore settings	3-3
Changing semaphore settings	3-3

Getting Information on Semaphores	3-3
Removing Semaphore Arrays	3-4
Semaphore Errors	3-4
Remote Simulation Manager (RSM)	3-5
When do you need RSM?	3-7
Installation/Configuration Overview	3-8
Installing RSM	3-8
Configuring RSM	3-9
Windows to Linux Job Submission	3-11
Prerequisites for Job Submission	3-11
Directory Shared between Windows and Linux	3-11
Network Access from Windows Host to Linux Job Management Host	3-11
Ansoft RSM Service Running on Job Management Host	3-11
Prerequisites for Job Monitoring	3-12
Network Access from Windows Host to Linux Cluster Hosts ...	3-12
Supported Schedulers	3-12
Select Scheduler Dialog	3-13
Submit Job Dialog	3-13
User Passwords are Encrypted	3-17
Configuration of AnsoftRSMService for Job Submission	3-18
IBM Platform LSF	3-18
PBS (PBSPro, Torque, and Moab/Torque)	3-18
SGE (including Oracle Grid Engine, Open Grid Scheduler and Univa)	3-19
Registering ANSYS Electromagnetics Product Engines for use with RSM	3-19
Supported Job Schedulers	3-20
Custom Scheduler Integration	3-20

Running HPC Test Jobs on a Cluster	3-21
Non-distributed job	3-22
Non-MPI job distributed to multiple hosts	3-23
MPI job distributed to multiple hosts	3-23
Submitting and Running the Tests	3-24
Submitting and running a non-distributed job	3-26
Submitting and running a non-MPI job distributed to multiple hosts	3-28
Submitting and running an MPI job distributed to multiple hosts	3-30
Scheduler Specific Configuration and Testing	3-33
IBM Platform LSF	3-33
Cluster Configuration Guidelines	3-33
Running Test Jobs	3-34
Non-distributed job	3-34
Non-MPI job distributed to multiple hosts	3-34
MPI job distributed to multiple hosts	3-35
Grid Engine (SGE, UGE, OGE)	3-36
Cluster Configuration Guidelines	3-36
qssh Command	3-36
Parallel Environments	3-37
Submitting Exclusive Jobs	3-38
Configuring an exclusive resource	3-39
Submitting an exclusive job on the command line	3-40
Submitting an exclusive job using the job submission GUI	3-40
Consumable Memory Limits	3-41
Running Test Jobs	3-43
Non distributed job	3-43

Non-MPI job distributed to multiple hosts	3-44
MPI job distributed to multiple hosts	3-44
PBSPro	3-45
Cluster Configuration Guidelines	3-45
Running Test Jobs	3-46
Non-distributed job	3-47
Non-MPI job distributed to multiple hosts	3-47
MPI job distributed to multiple hosts	3-48
PBS/Torque	3-48
Cluster Configuration Guidelines	3-48
Running Test Jobs	3-49
Non-distributed job	3-49
Non-MPI job distributed to multiple hosts	3-50
MPI job distributed to multiple hosts	3-50
4 - Updating Software and Licenses	4-1
Updating ANSYS Electromagnetics Software	4-1
Updating a Software License	4-2
Specifying License Server Machines via the ansyslmd.ini file or Environment Variables.	4-3
5 - Message Passing Interface (MPI) Installation	5-1
Linux	5-1
General Notes	5-3
Setting up passwordless ssh	5-3
Configuring passwordless ssh for the first time	5-4
Troubleshooting Section	5-5
A - Silent Installation	A-1
Running the Silent Installation	A-1
Response File	A-2

B - License hostid	B-1
Network Interface Card	B-1
Windows	B-2
Linux	B-2
Hardware Key	B-2
Installing a Hardware Key on a USB Port	B-2
Machine hostid	B-3
C - Setting Up a Printer on Linux	C-1

1 - Getting Started

Installation Prerequisites

Before installing ANSYS Electromagnetics Suite software, you need the following:

- The ANSYS Electromagnetics product CD or product download (see ["Additional Resources "](#) on page 1-5 for download information).
- A license file based on the ansoftd vendor daemon.
- The ANSYS License Manager software download.
- A computer with TCP/IP installed, for counted evaluation and purchased licenses.
- A computer that meets the minimum or recommended system requirements, which depends on the specific ANSYS Electromagnetics products you plan to run.

ANSYS Electromagnetics Suite Software

The ANSYS Electromagnetics product or equivalent download contains:

- The ANSYS Electromagnetics Suite software and project examples.
- The help system and PDFs.

License File

Most ANSYS Electromagnetics products require a license, which is sent as an e-mail attachment. After you receive a license file, save it to a temporary location.

The license file requires a license hostid, which is a hardware device used to validate the license. A list of supported hostids, and a general explanation of hostids, is available in ["License hostid "](#) on page B-1

If you do not have a license file for the ANSYS Electromagnetics software you are about to install, please contact your ANSYS sales representative.

License Server

Once you have purchased the software, you must designate a machine on your network to act as a license server. The ANSYS License Manager must be installed in one of the following two locations:

- On the computer where the software will be installed.
- On another machine accessible over the network.

Note You can install the ANSYS License Manager before or after you install the ANSYS Electromagnetics software.

See the ANSYS License Manager instructions available as part of the ANSYS License Manager download.

Hardware and Software Requirements

For supported platforms and system requirements, go to:

<http://www.ansys.com/Solutions/Solutions-by-Role/IT-Professionals/Platform-Support>

Operating System Specific Prerequisites

Red Hat Enterprise Linux

Red Hat Enterprise Linux is by default installed with only 64-bit binaries. ANSYS Electromagnetics products require some 32-bit libraries to run properly. When the application is started, it will identify any missing runtime libraries.

Additional Graphics Libraries

ANSYS Electromagnetics applications make use of OpenGL libraries. If you experience sluggishness when interacting with the UI, or if you see any of the following errors, you will need to install OpenGL libraries.

- Unable to resolve function `glXQueryExtension`
- Failed to initialize an OpenGL view

ANSYS recommends the use of OpenGL libraries provided by the manufacturer for the graphics card in your machine. Alternately, you may consider installing the Mesa OpenGL RPM package **mesa-libGL.i686** and its dependencies.

Required 32-bit Libraries to Run Applications

The following minimal set of 32-bit libraries and their dependencies are required to run applications. Note that these are the minimum versions required; you may use newer versions.

```
glibc.i686
glib2.i686
bzip2-libs.i686
```

compat-expat1.i686

libpng.el6.i686

libtiff.i686: Also includes required dependencies:

- libjpeg.i686
- zlib.i686

libXft.i686: Also includes required dependencies:

- fontconfig.i686
- freetype.i686
- libXrender.i686

libXxf86vm.i686

Required 32-bit Libraries to Install Applications

The following minimal set of 32-bit libraries and their dependencies are required to install applications. Note that these are the minimum versions required; you may use newer versions.

- glibc.i686

License Server

Once you have purchased the software, you must designate a machine on your network to act as a license server. The ANSYS License Manager must be installed in one of the following two locations:

- On the computer where the software will be installed.
- On another machine accessible over the network.

Note You can install the ANSYS License Manager before or after you install the ANSYS Electromagnetics software.

See the ANSYS License Manager instructions available as part of the ANSYS License Manager download.

Required Libraries to Install License Server

The following minimal set of 32-bit libraries and their dependencies are required to install applications. Note that these are the minimum versions required; you may use newer versions.

glibc.i686

compat-libstdc++-33.i686

Required Libraries to Run License Server

The following minimal set of 32-bit libraries and their dependencies are required to run applications. Note that these are the minimum versions required; you may use newer versions.

```
glibc.i686
glib2.i686
bzip2-libs.i686
compat-expat1.i686
compat-libstdc++-33.i686
libpng.el6.i686
libtiff.i686: Also includes required dependencies:
• libjpeg.i686
• zlib.i686
libXft.i686: Also includes required dependencies:
• fontconfig.i686
• freetype.i686
• libXrender.i686
libXxf86vm.i686
```

Additional Prerequisites

Red Hat Enterprise Linux 6.1 introduces **biosdevname**, an optional convention for naming network interfaces. **biosdevname** assigns names to network interfaces based on their physical location. Note, however that **biosdevname** is disabled by default, except for a limited set of Dell systems. The FLEX license manager requires a network device that is named **ethX**, where **X** is the lowest numbered eth device (e.g.: eth0, or eth1). Licensing does not work when the device has any other name. To resolve this, you will need to change the name of the device to match the Linux standard convention. There are four options you can use to enable the standard ethernet names:

1. If you do not need to use the Consistent Network Device Naming package, you can remove the biosdevname package from your installation by running the following command in a terminal as root:

```
yum remove biosdevname
```

This removes the package and on reboot will restore the traditional naming scheme.

2. You can update the name in the /etc/sysconfig/network-scripts/ifcfg-* file:

There should be a file in /etc/sysconfig/network-scripts that is named "ifcfg-" and the name of your network device. To restore the name, you can rename this file to "ifcfg-eth0" and rename the device name from the current name to eth0 in the

contents of the file. Once updated, restarting networking services or rebooting should enable the change.

- You can write rules in `/etc/udev/rules.d/70-persistent-net.rules` to change the device names. This will take precedence over the physical location naming scheme. Such rules may look like:

```
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*",
ATTR{address}=="00:11:22:33:44:55", ATTR {type}=="1",
KERNEL=="eth*", NAME="public"
```

- You can disable `biosdevname` in the kernel command line by rebooting the computer and bringing up the kernel command line in the boot menu. From the kernel command line, you can run:

```
biosdevname=0
```

You can find more information about this change:

<https://access.redhat.com>

<https://fedoraproject.org/wiki/Features/ConsistentNetworkDeviceNaming>

<http://rhn.redhat.com/errata/RHEA-2011-0644.html>

http://linux.dell.com/files/whitepapers/consistent_network_device_naming_in_linux.pdf

Additional Resources

General:	The ANSYS web site: http://www.ansys.com .
Downloads:	Software downloads (including FLEXlm for Ansoft) and installation-related documentation are available at the ANSYS customer portal, https://support.ansys.com/portal/site/AnsysCustomerPortal
Online Technical Support:	Available through the ANSYS customer portal: https://support.ansys.com/portal/site/AnsysCustomerPortal Under Online Support , choose ANSYS Electromagnetics Customer Support . The ANSYS Electromagnetics technical support site provides access to frequently asked questions, application notes, presentations, and example scripts.

Product Documentation:	The help system is accessible from the Help menu. Information about the release can be found in the <i>What's New</i> documents. In addition, PDF versions of the help system are accessible from within the help system.
-------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

2 - Installing ANSYS Electromagnetics Software

To install ANSYS Electromagnetics software on a Linux-based machine:

- Obtain the prerequisites as listed in chapter 1.
- Obtain a license file.
- Verify disk space.
- Install the ANSYS Electromagnetics software.
- Set up a license server if prompted to do so.

Obtaining a License File

ANSYS Electromagnetics requires a license for all software. If you do not have a license file for the ANSYS Electromagnetics software you are about to install, contact your ANSYS sales representative. See "[License File](#)" on page 1-1 for more information.

Verifying that Enough Disk Space is Available

The amount of disk space required by the ANSYS Electromagnetics software depends on which packages you are installing. Refer to the hardware requirements for each package you have purchased for more details. During the setup of a particular installation, you are informed of the disk space requirements for that software.

You need sufficient disk space in the following directories:

- The directory where the software is installed.
- The directory where the software stores temporary files while solving a project.

Note: On every machine where you are going to run the software, you also need enough swap space to accommodate at least the amount of RAM on the machine. Contact your system administrator if you have questions about swap space.

Installing the Software

To install ANSYS Electromagnetics software:

1. Run the installation program.
An initial **Installation Shell** screen automatically appears.
2. Click **Install <Product>** for the product you want to install.

A **Welcome** screen appears.

Note If you receive an error that InstallShield did not find enough temporary space to run, set your ISHIELD_TMP environment variable to a different directory.

1. Exit the installation.
2. Go to a prompt window, and type the following:
setenv ISHIELD_TMP <newdir> (csh)
export ISHIELD_TMP=<newdir> (bash)
3. Return to step 1 on "[Installing the Software](#)" on the [previous page](#).

3. Click **Next**.

The **License Agreement** screen appears.

4. Read the license agreement carefully. Click **I accept the terms of the license agreement** if you agree to the terms of the license agreement.

5. Click **Next**.

The **Choose Destination Location** screen appears.

6. Under **Directory Name**, do one of the following to specify a path:

- Accept the default installation directory as it appears under **Directory Name**.
- To select a different existing directory, either type the path directly, or browse to the location by clicking **Browse**. When you click **Browse**, the **Select a Directory** window appears. Select a directory, and click **Open**. *If you have installed other ANSYS Electromagnetics software, we recommend that you install new software in the same directory.*
- To create a new directory, either type the path directly, or browse to the location for the new directory by clicking **Browse**. When you click **Browse**, the **Select a Directory** window appears. Enter the new directory name in the **Enter file name** field, and click **Open**.

Note Multiple machines/nodes can run from the same installation tree. You do not have to install the application on each unix/linux machine/node. ANSYS Electromagnetics recommends installing to a network path.

7. Click **Next**.

Depending on the software and version being installed, one or both of the following screens may appear before the **Select**

License Information screen: the **Choose Temporary Directory Location Select Library Option** screen.

8. If the **Temporary Directory Location** screen appears, select a default location for all simulations (including local) to use as a temporary work space.
 - If desired, change the temporary directory assignment. You can edit the field directly or use the Browse button to display file browser to let you navigate the directories on your machine and network.
 - If desired, check the option to add read/write permissions for all users. If the installation will be used for remote simulation by multiple users, all users must have permissions on the temporary directory.
9. If the **Choose Library Option** screen appears, do one of the following to specify the location of ANSYS Electromagnetics libraries:

Note ANSYS Electromagnetics libraries are large databases of electronic components, developed by ANSYS and various manufacturers.

 - Click **This local installation** to copy and install the default ANSYS Electromagnetics component libraries, and then click **Next**.
 - Click **The following common location** to use libraries that are already installed (locally or on the network). Click **Browse**, select the appropriate path, and click **OK**. Click **Next**.
10. Click **Next**.

The **License Information** screen appears.

11. Select one of the following three license options.

I have a new license file

Select this if you recently received a license file from ANSYS (either evaluation or purchased).

a. Click **Next**.

Another License

Information window appears, asking you to specify a directory and license file.

b. Click **Browse**.

The **Select License File** window appears.

c. Find and select the license file you received from ANSYS Electromagnetics. Click **Open**, and then click **Next**.

d. If prompted, enter the name of the license server. You may also select or clear the **Specify server TCP/IP port number (advanced)** option. Specifying a TCP/IP port means that the software only looks at that port for licenses. The default port for the ANSYS License Manager is 1055.

e. Click **Next**.

I want to specify a license server

Select this if you are using a counted license and already know on which machine the ANSYS License Manager is/will be installed.

a. Click **Next**.

Another License

Information window appears, asking if you want to use a single server or a redundant server.

b. Select or clear the **Specify redundant servers** check box, and enter the server name(s) in the field(s).

The redundant-server setup requires three license servers, each serving an identical license file. The order in which the servers are entered must match the license files on each server.

c. Select or clear the **Specify server TCP/IP port number (advanced)** option. Specifying a TCP/IP port means that the software only looks at that port for licenses. The default port for the ANSYS License Manager is 1055.

I want to keep the current license configuration

Select this if you already have a license file for previous ANSYS software and want to keep the same configuration.

• Click **Next**.

For additional information about the three license options, click **Help me decide which of the following options to choose**.

Note If you want to change the license server after installation, modify this file: <install_dir>/shared_files/li-

censing/ansyslmd.ini. For details, see [Updating a Software License](#).

12. In the **Summary Information** screen, review the information displayed under **Current Settings**, and then do one of the following:
 - If the settings are correct, click **Next** to start copying files.
 - If the settings are incorrect, click **Back** as many times as necessary, and change your selections. Verify each screen again before clicking **Next** to move on to the next screen. When all information is correct, return to the **Summary Information** screen, and click **Next**.

The **Setup Status** screen appears, displaying a progress bar, and installation begins.

Once the files have all been copied, a message appears telling you to verify that the latest version of ANSYS License Manager is installed on the machine you specified as the license server.

Note If you receive a message about read-only files, you typically want to overwrite them with the newer versions.

13. Click **OK** in the message.

When the installation is complete, a **Setup Complete** window appears for the product you installed.

14. Click **Finish**.

Integrating ANSYS Electromagnetics Products with ANSYS Workbench

After the installation is complete, you can integrate ANSYS Electromagnetics with ANSYS Workbench by running the following command:

```
<install_dir>/An-  
sysEM21.1/Linux64/scripts/IntegrateWithANSYS19.5.pl.
```

and following the prompts for integration.

3 - High Performance Computing

ANSYS Electromagnetics products have a variety of ways to use multiple threads, multiple cores or multiple hosts to improve productivity. These additional computational resources may be used to solve larger problems, to solve a given problem faster, or to improve the quality of a solution (using more stringent convergence criteria, for example) in the same amount of time. This section describes the configuration and requirements for using additional resources by ANSYS Electromagnetics products.

To make effective use of the resources allocated to the analysis, the job settings must match the job requirements. In addition, the cluster configuration and the resources must be allocated to the job in a way that the products are able to make effective use of the resources.

Note For additional information about high performance computing not in this guide, see the *ANSYS Electromagnetics HPC Administrator's Guide* (HPC_Admin.pdf) in <installdir>/AnsysEM<release_number>/Linux64/Help.

Options for Distributed Analysis

There are several options for configuring ANSYS Electromagnetics products to take advantage of additional computational resources. The simplest approach is to use the Remote Simulation Manager (RSM) for ANSYS Electromagnetics products. RSM may be installed, configured, and started on one or more hosts. It runs as a daemon. An analysis may be distributed to any of the hosts where the Remote Simulation Manager is running. RSM is bundled with ANSYS Electromagnetics products.

There are a number of job schedulers available from third parties, such as IBM Platform LSF, Univa Grid Engine, or PBSPro. A job scheduler may also be called a distributed resource management system, a cluster manager, a batch system or a distributed resource manager. These systems provide queuing of jobs, load balancing, assigning resources to jobs, and many other management and monitoring features. ANSYS Electromagnetics products integrate with several of the most widely used job schedulers. The user is able to submit jobs to the scheduler. When the job runs, the product uses the resources allocated to the job to accomplish the analysis.

Users may be able to use job schedulers which are not included in the list of supported job schedulers via "custom scheduler integration." The end user, working with ANSYS Electromagnetics products support engineers, will provide an IronPython script or code library to act as an interface between the scheduler and the ANSYS Electromagnetics products.

The RSM for ANSYS Electromagnetics products may be used for running an analysis on a remote host or for running an analysis distributed to multiple hosts. However, the user is required to manage the jobs and resources used by these jobs; no ability to queue jobs until needed resources are available is provided. In general, a job scheduler is more appropriate if there are a number of hosts shared between a number of users, because a scheduler is better able to manage the jobs and the resources. It can queue a job and delay the start the job until appropriate resources become available. Custom scheduler integration is needed to allow jobs to be run on a cluster that is not managed by one of the supported job schedulers.

Semaphores on Linux

The Linux kernel provides a number of features to facilitate synchronization between threads and/or processes. Semaphores are one such mechanism. Linux supports two styles of semaphores, System V IPC semaphores and POSIX semaphores. On Linux, ANSYS EM software uses MainWin, which uses System V IPC semaphores. If there are not enough System V IPC semaphores available, the software may fail to start or terminate due to errors.

System V IPC semaphores are allocated as semaphore arrays. Each array may contain a few or many semaphores. The Linux kernel limits both how many semaphore arrays can be allocated, as well as the total number of semaphores that can be created.

Recommended settings

The following kernel semaphore settings are recommended:

- SEMMSL (maximum number of semaphores per semaphore set) = 256
- SEMMNS (system-wide limit on the number of semaphores in all semaphore sets) = 40000
- SEMOPM (maximum number of operations that may be specified in a semop() call) = 32
- SEMMNI (system-wide limit on the number of semaphore identifiers) = 32000

Displaying semaphore settings

The following Linux command will display the current semaphore settings:

```
cat /proc/sys/kernel/sem
```

The semaphore settings are displayed or set in the order: SEMMSL SEMMNS SEMOPM SEMMNI

Sample output:

```
% cat /proc/sys/kernel/sem
250 32000 32 128
```

Changing semaphore settings

The following Linux command may be used to change the semaphore settings to the recommended values:

```
echo 256 40000 32 32000 > /proc/sys/kernel/sem
```

This command must be run as a privileged user, e.g., root. This command immediately changes the kernel limits. Rebooting the machine resets the semaphore limits to the default values.

Contact a system administrator to configure the semaphore settings at boot time.

Getting Information on Semaphores

Use the following Linux command to show information about the currently allocated System V semaphore arrays:

```
/usr/bin/ipcs -s
```

Sample output:

```
%/usr/bin/ipcs -s
----- Semaphore Arrays -----
key          semid      owner     perms    nsems
0x00000000   23887872  user1    664      2
0x00000000   19202049  user1    664      2
0x00000000   22413314  user1    664      2
0x00000000   3735555   user1    664      2
0x00000000   22446084  user2    664      2
0x00000000   22478853  user2    664      2
```

Use the following Linux command to show additional information about the semaphore array with semid id:

```
/usr/bin/ipcs -s -i id
```

Sample output:

```
%/usr/bin/ipcs -s -i 23887872
Semaphore Array semid=23887872
uid=789 gid=123 cuid=456 cgid=100
mode=0664, access_perms=0664
nsems = 2
otime = Wed Mar 11 18:48:16 2020
ctime = Wed Mar 12 18:48:16 2020

semnum  value  ncount  zcount  pid
0        0       0        0       0
1        1       0        0      31667
```

Removing Semaphore Arrays

System V IPC semaphore arrays are not automatically removed when a process terminates. Normally, processes that create semaphore arrays remove them on shutdown. If a process that creates some semaphore arrays terminates in an abnormal manner, such as due to a crash or a kill command, the semaphore arrays created by the process may not be removed. These semaphore arrays will remain allocated, reducing the number of semaphore arrays available for other processes. On a reboot, all semaphore arrays are removed.

You can also remove semaphore arrays using the following Linux command, where *semid* is the id of the semaphore array to remove:

```
/usr/bin/ipcrm -s semid
```

Semaphore arrays should only be removed if they are no longer being used by any processes.

Semaphore Errors

Error messages due to running out of semaphores may look like the following:

*sem_lock->semop->op_op: Invalid argument

*sem_unlock->semop->op_op: Invalid argument

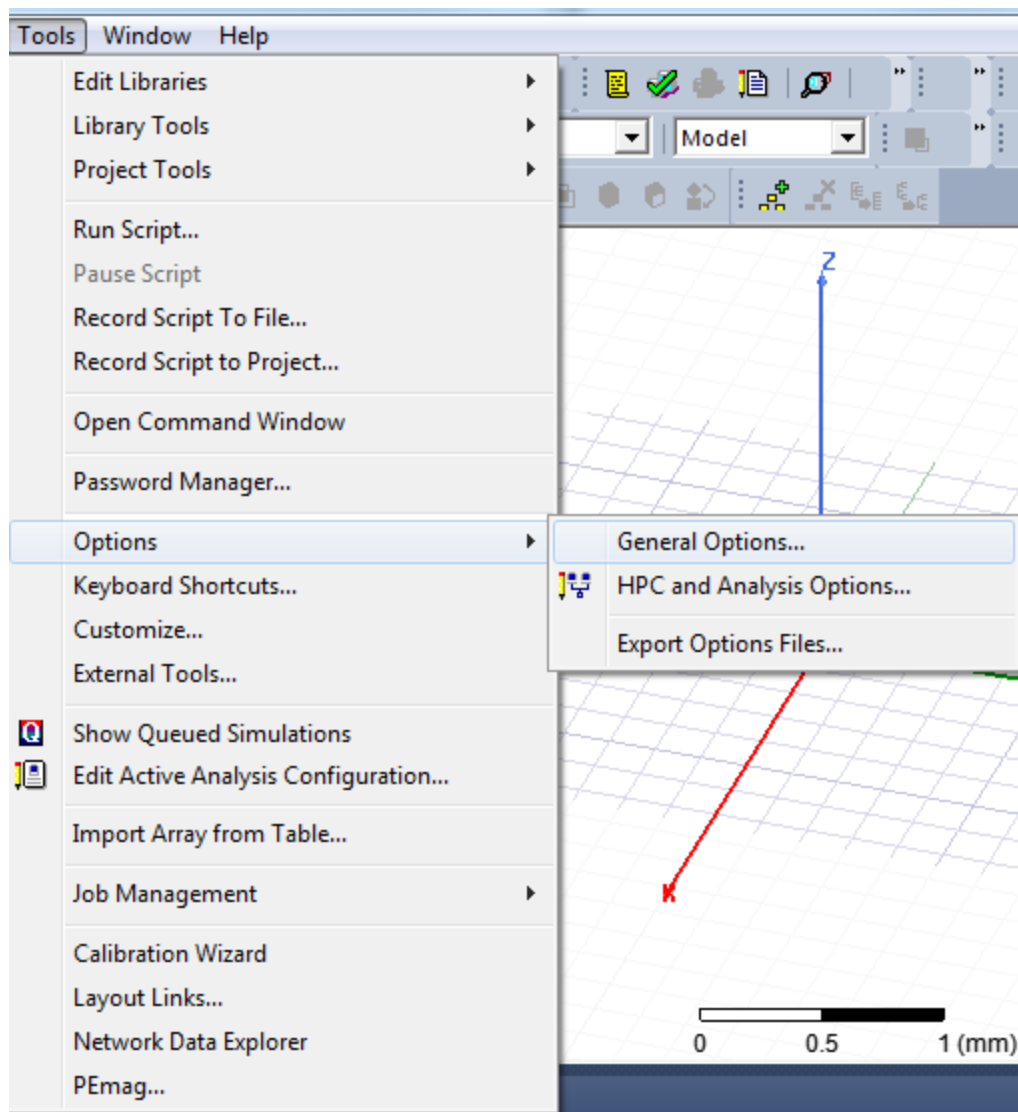
*sem_unlock->semctl: Invalid argument

If any errors similar to the above are seen, use the linux command `cat /proc/sys/kernel/sem` to display the semaphore limits, and use the `/usr/bin/ipcs -s` command to display the semaphore sets currently in use. To make more semaphore arrays available, use the commands shown above to change the semaphore limits, or to remove semaphore sets that are no longer in use.

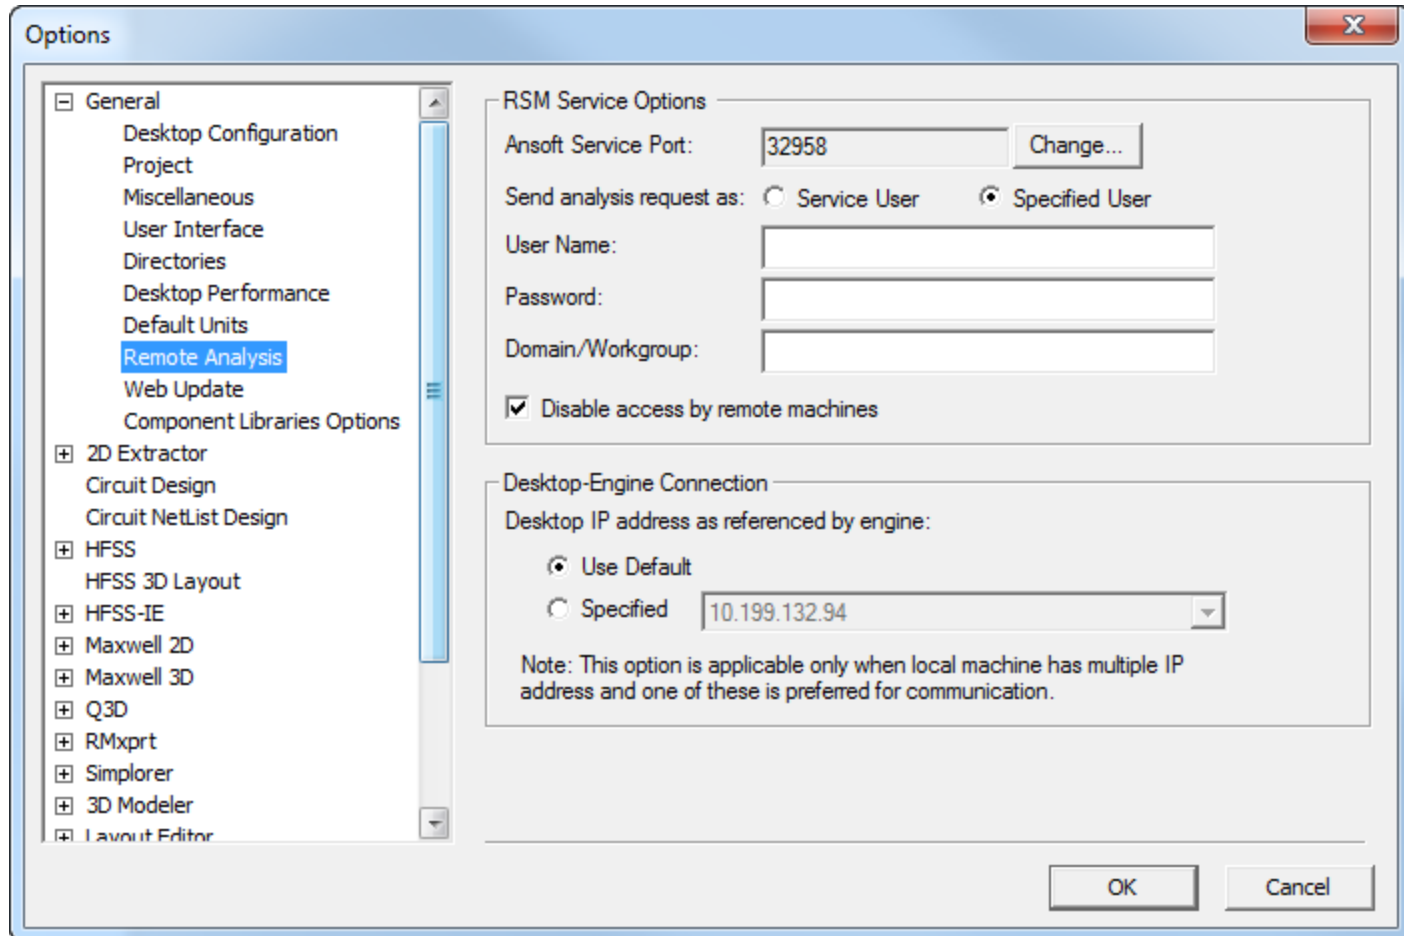
Remote Simulation Manager (RSM)

The Remote Simulation Manager (RSM) for ANSYS Electromagnetics products may be used for running an analysis on a remote host or for running an analysis distributed to multiple hosts. When using RSM, the users are required to manage the jobs and resources used by these jobs. No load balancing or queueing of jobs is provided. The Remote Simulation Manager is not used if the analysis is run as a job on one of the supported job schedulers, such as LSF, PBS, or Windows HPC. See ["Supported Job Schedulers" on page 3-20](#) for a list of these jobs schedulers. RSM is also not used if the analysis is run as a job using custom scheduler integration with a scheduler that is not supported by default. See ["Custom Scheduler Integration" on page 3-20](#) for more information.

If the the Remote Simulation Manager is running on a host, by default, jobs started on remote hosts may distribute jobs to the local host. A user may disable this access or re-enable this access using the Options dialog. The **Tools>Options>General Options** menu item pops up the Options dialog:



The **Disable access by remote machines** check box, in the **RSM Service Options** group box, on the **General>Remote Analysis** page controls access. If this check box is not checked, the default setting, then remote hosts can distribute jobs to the local host. If this check box is checked, then remote hosts cannot distribute jobs to the local host.



When do you need RSM?

RSM is needed only to run an analysis on a remote host or an analysis distributed to multiple hosts. RSM is also required if some of the solvers are running on a hosts that run under a different OS than the OS host that is running ANSYS Electronics Desktop. If RSM is required, then the RSM daemon or service must be configured and running on all remote hosts for the analysis.

Please note that RSM is **not** needed:

- To run an analysis only on the local host.
- To run an analysis using integration with a supported job scheduler or using custom scheduler integration.

You should not install RSM on hosts that are managed by a scheduler.

Installation/Configuration Overview

Both RSM and ANSYS Electromagnetics applications can be installed in one central location. You do not have to install the application on each machine/node. We recommend using the network-based directions below.

Network Installation: Overview

1. Install all ANSYS Electromagnetics products in a network location accessible from all machines. See ["Installing the Software" on page 2-1](#) for details on installing an ANSYS Electromagnetics product.
2. Install RSM in a network location accessible from all machines.
3. On every machine, configure and start RSM.
4. On one machine, register ANSYS Electromagnetics product installations with RSM.

Non-Network Installation: Overview

1. Install ANSYS Electromagnetics products on every machine. See ["Installing the Software" on page 2-1](#) for details on installing an ANSYS Electromagnetics product.
2. Install RSM on every machine.
3. On every machine, configure and start RSM.
4. On every machine, register ANSYS Electromagnetics product installations with RSM.

Installing RSM

To install ANSYS Electromagnetics RSM software:

1. If you have received a CD or DVD for the ANSYS Electromagnetics product you wish to install, insert it into the CD-ROM drive and run `install` from the top-level directory of the CD-ROM.

If you downloaded the ANSYS Electromagnetics product from ANSYS Electromagnetics software downloads website (see ["Additional Resources " on page 1-5](#)), unzip the download, and run the **install** file in the newly created directory.

An initial **Installation Shell** screen automatically appears.

2. Click **Install RSM** for the product you want to install.
A **Welcome** screen appears.

Note If you receive an error that InstallShield did not find enough temporary space to run, set your ISHIELD_TMP environment variable to a different directory.

1. Exit the installation.
2. Go to a prompt window, and type the following:

```
setenv ISHIELD_TMP <newdir> (csh)
export ISHIELD_TMP=<newdir> (bash)
```
3. Return to step 1.
3. To install the ANSYS Electromagnetics Remote Simulation Manager, click **Next**.
The **License Agreement** screen appears.
4. Read the license agreement carefully. Click **I accept the terms of the license agreement** if you accept the terms of the agreement.
5. Click **Next**.
The **Choose Destination Location** screen appears.
6. **Under Directory Name, do one of the following to** accept the default path or specify a new one:
 - Accept the default installation directory listed under **Directory Name**.
 - To select a different directory, click **Browse**. The **Select a directory** window appears. Select a directory, and click **OK**. *If you have other ANSYS Electromagnetics software, we recommend that you install new software in the same directory.*To create a new directory from the **Select a directory** window, navigate to the desired location, and click the Create a New Folder icon. Enter the new name in the **File name** field. Click **OK** to close the **Select a directory** window.
7. Click **Next** from the Remote Simulation Manager Setup.
A summary information window appears, listing the location where the ANSYS Electromagnetics Remote Simulation Manager will be installed, and the total size.
8. Click **Next**.
The install executes and installation summary appears.
9. Click **Finish** to close the install window.

Configuring RSM

Note: Configuring RSM on a Linux Job Management Host for submitting jobs to a cluster is separate from using RSM to launch solver processes. That is, a single RSM process cannot be used to do both functions. In addition, RSM may run as the root user when used to launch solver processes, but it should not be run as the root user when used for submission of jobs to a Linux cluster.

General RSM syntax

RSM's command-line interface uses an LSB-compliant syntax:

<path to RSM installation>/rsm/ansoftrmservice command, where *command* includes:

- **help** : Show full details of all RSM command options

- **start** : Start RSM
- **status** : Show status of RSM
- **stop** : Stop RSM

- **condrestart** : Restart RSM if it is already running
- **restart** : Restart RSM, or start it if it is not running
- **reload** : Restart RSM, or start it if it is not running (same as restart)

- **startonboot** : Configure RSM to start automatically on boot as root (command must be run as root)
- **startonboot -user <username>**: Configure RSM to start automatically on boot as a specific user (command must be run as root)

Determining Who Should Run RSM

The user who runs RSM determines who launches solver processes.

User running RSM	User who owns/launches solver processes
Non-root user	User running RSM
root	User who requests to solve, or if that user does not exist on this machine, the user who is running RSM

First-time RSM configuration

RSM should be configured to start automatically on boot on each node to which remote simulation requests will be sent. You must also decide which user will run RSM ("[Determining Who Should Run RSM](#)" on the [previous page](#)).

To configure RSM to start automatically as a specific user, run:

```
<path to RSM installation>/rsm/ansoftrsmervice startonboot -user  
<username>
```

This option will write the appropriate files in /etc/rc.d as well as start the RSM service as that user.

Windows to Linux Job Submission

Given a set of prerequisites, ANSYS Electronics Desktop can permit Windows to Linux job submission as part of HPC.

Prerequisites for Job Submission

Directory Shared between Windows and Linux

For all jobs submitted to a Linux cluster, the project file is required to be in a directory that is accessible from all execution hosts used by the job. For submission of jobs from a Windows host to a Linux cluster, the project file must also be accessible from the Windows host where the GUI runs. There must be a directory shared with both Windows and Linux hosts, and the project file may be in a subdirectory (at any level) of the shared directory.

Network Access from Windows Host to Linux Job Management Host

The job is submitted to the cluster from a Linux host configured for submission of jobs to the Linux cluster. We call this Linux host the "Job Management" host. The information about the job to be submitted is transmitted to the Job Management host over the network. As a result, the Windows host where the GUI runs must have network access to the Job Management host. If this communication is blocked, then job submission from a Windows host to the Linux cluster will not be possible. Communication could be blocked if there is a firewall or if the Linux cluster is only on a private network, for example.

Ansoft RSM Service Running on Job Management Host

The ansoftrsmervice must be running on the Linux Job Management host. Before the ansoftrsmervice is started, it must be configured for submission of jobs to the cluster. The SchedulerName and

ConfigString fields in the Scheduler block of the ansoftrmservice.cfg configuration file must be specified. The contents of these fields are described in the following table:

Field Name	Contents	Examples
SchedulerName	Identifier of Scheduler	IBM Platform LSF: 'lsf' PBSPro or Torque: 'pbs' Univa, SGE, etc.: 'sge'
ConfigString	Directory containing scheduler commands	IBM Platform LSF: '' (not required) PBSPro or Torque: '/opt/pbspro/PBSPro_13.0.0/default/bin' Unive, SGE, etc.: '/opt/univa/bin/lx-amd64'

The environment should be configured for job submission before starting the ansoftrmservice. The ansoftrmservice should be run as a non-privileged user; no special privileges are required. It should be run as a user without login privileges, so that only privileged users have access to this process.

Prerequisites for Job Monitoring

For job monitoring, all prerequisites for job submission are required. One additional requirement, described below, is also required for job monitoring.

Network Access from Windows Host to Linux Cluster Hosts

In order to obtain full monitoring information from a job, the Windows host needs access to some of the job processes. That is, the Windows monitoring host requires network access to the processes running on the Linux cluster execution hosts. If this communication is blocked, then only limited monitoring information is available.

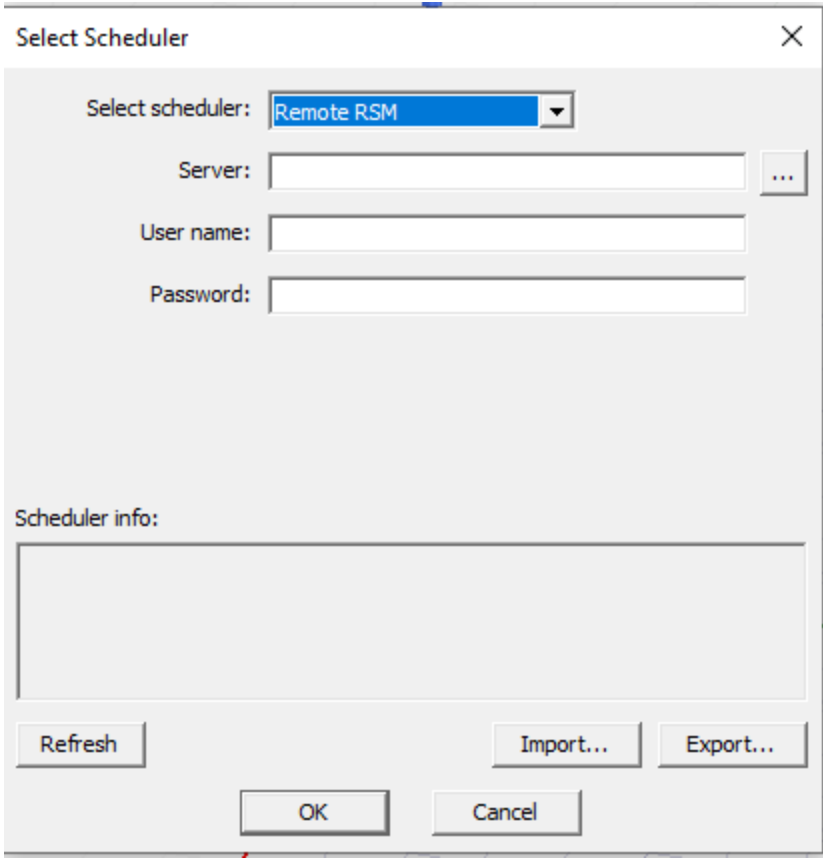
Supported Schedulers

This feature may be used with all Linux schedulers for which job submission from the GUI is supported:

- IBM Platform LSF
- Univa Grid Engine (formerly Sun Grid Engine, Oracle Grid Engine)
- PBSPro
- Torque

Select Scheduler Dialog

If you select **Remote RSM**, you can enter a username and password. This username and password are used when the job is submitted to the Linux scheduler. The AnsoftRSMService must be running on the remote host.

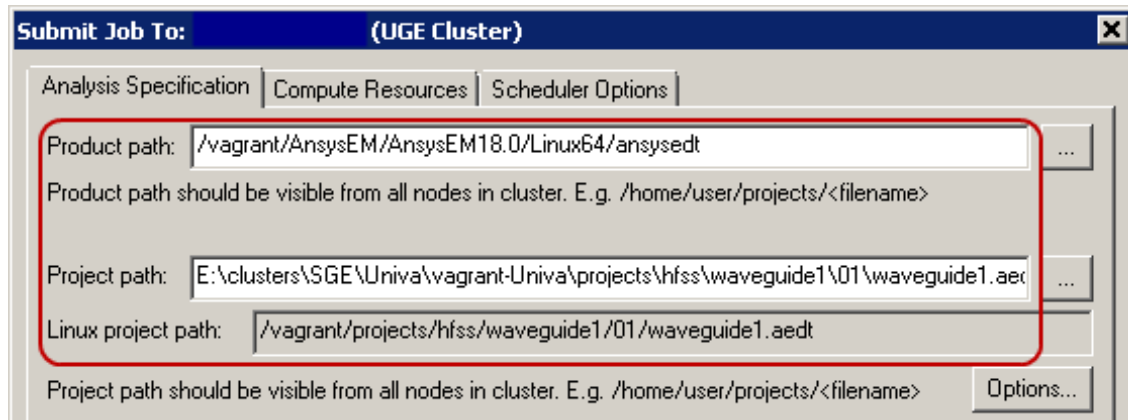


The screenshot shows a dialog box titled "Select Scheduler". At the top, there is a dropdown menu labeled "Select scheduler:" with "Remote RSM" selected. Below this are three text input fields: "Server:" (with a browse button "..."), "User name:", and "Password:". At the bottom of the dialog, there is a "Scheduler info:" label above a large empty text area. At the very bottom, there are five buttons: "Refresh", "Import...", "Export...", "OK", and "Cancel".

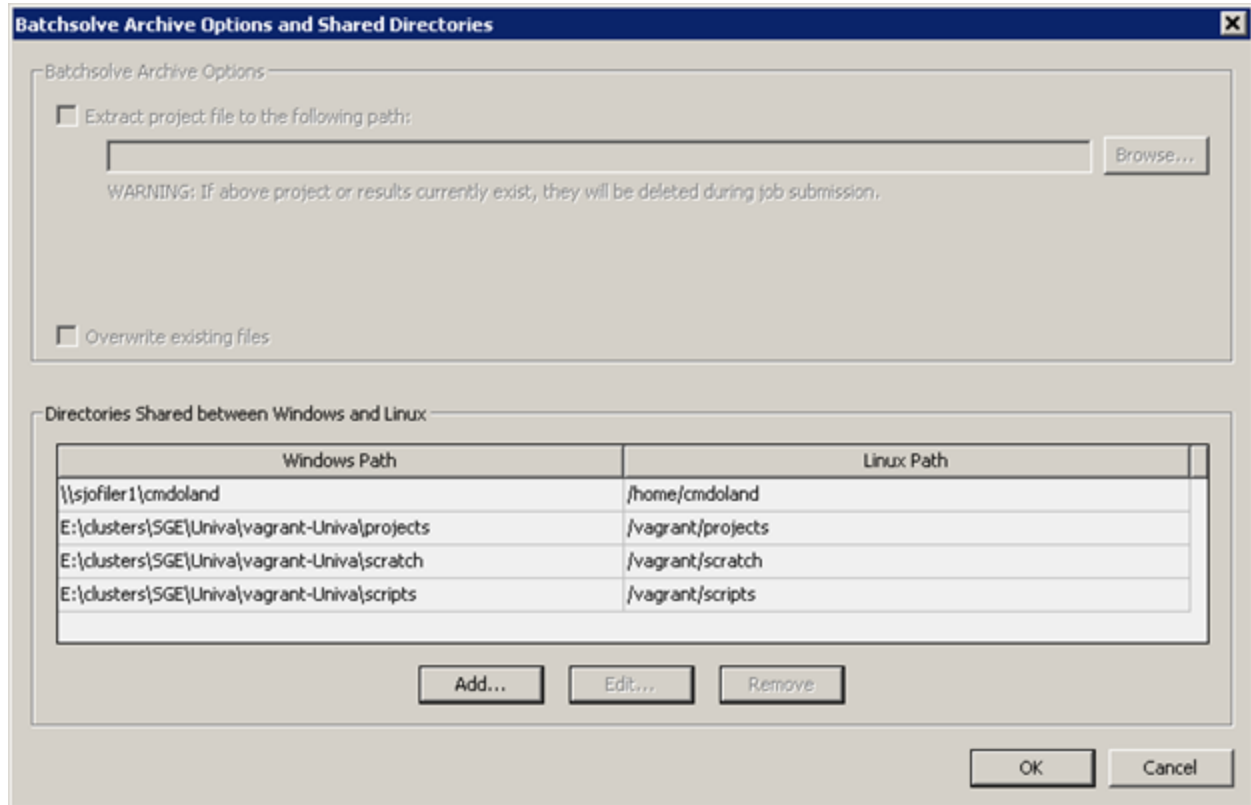
Submit Job Dialog

In the **Submit Job** dialog, you must enter the Linux path to the product in the Product path edit control. The browse button (labeled "...") may not be used to browse for the product. There is no requirement for the product installation directory to be accessible from the Windows GUI host.

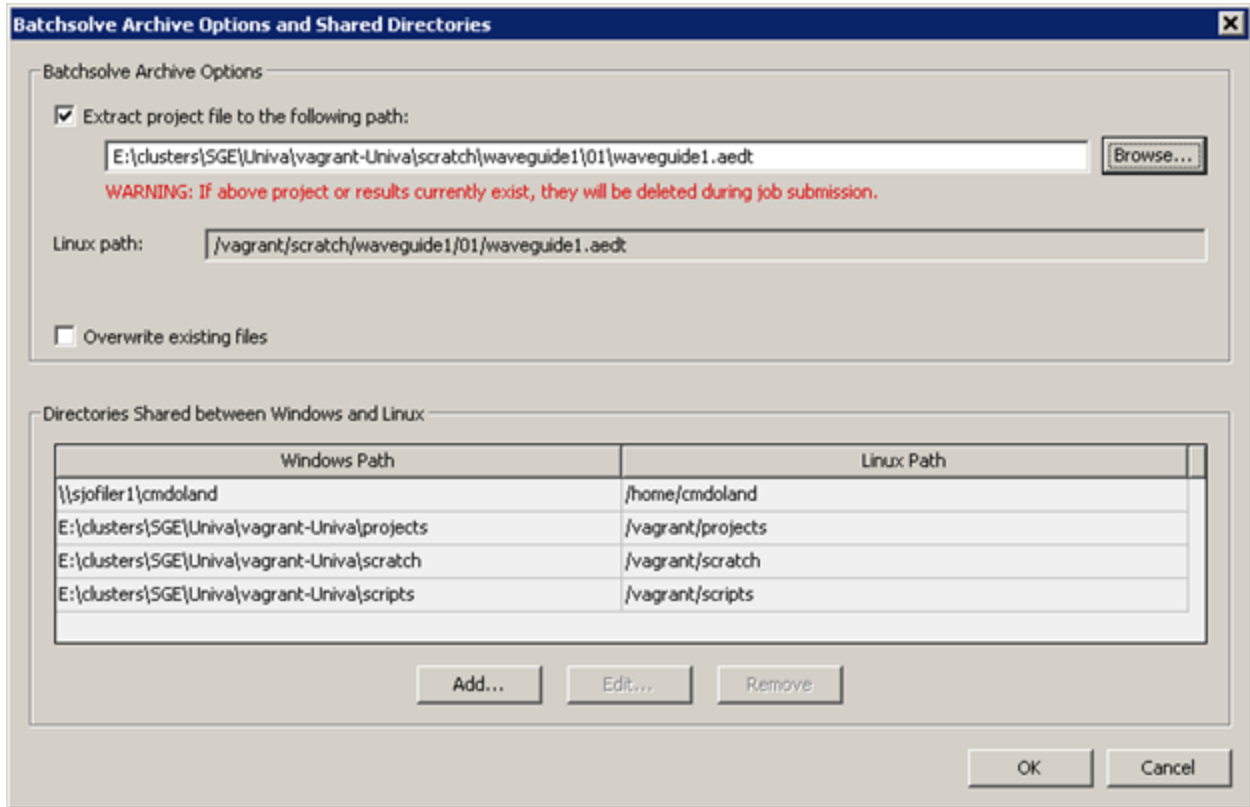
You must enter the Windows path of the project file in the Product path edit control or use the browse button (labeled "...") to select the Windows path of the project file. If the Linux path of the project file can be determined from the specified Windows path and the directories shared between Windows and Linux, then the Linux path of the project file is shown in the Linux project path edit control. This edit control cannot be edited directly.



The **Options** button activates a dialog that you can use to specify archive options for a job. The lower portion of this dialog allows you to specify one or more directories shared between Windows and Linux. The mapping of directories between Windows and Linux is shown in a grid which displays the Windows path and the Linux path for each shared directory. There are also buttons to add a new shared directory, to edit an existing shared directory, or to delete one or more shared directories. The Windows path or the Linux path of any shared directory may be selected in the grid and directly edited, as well.



If you specify a project in an archive, the dialog activated by the **Options** button may be used to specify the Windows pathname of the project to be extracted from the archive. If this is done, the Linux pathname of the target project is determined from the directories shared between Windows and Linux and shown in the upper portion of this dialog.



If you specify a batchextract script, the Linux path of the batch-extract script is determined from the Windows path of the batch-extract script and the directories shared between Windows and Linux.

Submit Job To: 192.168.10.109 (UGE Cluster)

Analysis Specification | Compute Resources | Scheduler Options

Product path: ...
 Product path should be visible from all nodes in cluster. E.g. /home/user/projects/<filename>

Project path: ...
 Linux project path: ...
 Project path should be visible from all nodes in cluster. E.g. /home/user/projects/<filename> Options...

Analysis setups

All setups in project

All setups in design:

Single setup: Use large scale D80

Monitor job (This must be checked to allow monitoring from the user interface.)

Wait for license

Analysis options

Batchoptions:

Add... Remove Edit...

Environment: ...

Use batch extract

Script path: ...
 Script path should be visible from all nodes in cluster. E.g. /home/user/projects/<filename>

Save Settings As Default Import... Export... Import Configuration

Preview Submission Show advanced options Submit Job Cancel

User Passwords are Encrypted

Jobs are submitted to the Linux cluster using the user name and password entered in the **Select Scheduler** dialog. These settings are persistent; in general, these settings need to be entered only if they change. To ensure security, user passwords are stored in an encrypted format. When a job is submitted from a Windows host using the

ansoftrmservice running on a Linux submission host, the user credentials are sent over the network in an encrypted format.

Configuration of AnsoftRSMService for Job Submission

The use of AnsoftRSMService for job submission is supported only on Linux. Before starting the AnsoftRSMService, the configuration file, `ansoftrmservice.cfg` in the AnsoftRSMService installation directory should be modified as described below for the associated scheduler. The AnsoftRSMService daemon should be started as a non-privileged user (that is, as a user other than root). This user should have the login password disabled so that only privileged users are able to access the AnsoftRSMService process.

Generally, the AnsoftRSMService should be started with the environment configured for submission of jobs to the Linux cluster. Other configuration requirements are specified in the scheduler specific sections below.

IBM Platform LSF

In the 'AnsoftCOMDaemon' block of the `ansoftrmservice.cfg` configuration file, the `SchedulerName` should be set to `lsf`, and the `ConfigString` should be set to an empty string. The AnsoftRSMService must be started with the environment configured for submission of jobs to the Linux cluster.

Example `ansoftrmservice.cfg` configuration file:

```
$begin 'AnsoftCOMDaemon'
  $begin 'Managed COM Servers'
  $end 'Managed COM Servers'
  $begin 'Scheduler'
    'SchedulerName'='lsf'
    'ConfigString'=''
  $end 'Scheduler'
$end 'AnsoftCOMDaemon'
```

PBS (PBSPro, Torque, and Moab/Torque)

In the 'AnsoftCOMDaemon' block of the `ansoftrmservice.cfg` configuration file, the `SchedulerName` should be set to `pbs`, and the `ConfigString` should be set to the directory containing the PBS commands. The AnsoftRSMService must be started with the environment configured for submission of jobs to the Linux cluster.

Example `ansoftrmservice.cfg` configuration file:

```
$begin 'AnsoftCOMDaemon'  
  
    $begin 'Managed COM Servers'  
    $end 'Managed COM Servers'  
    $begin 'Scheduler'  
        'SchedulerName'='lsf'  
        'ConfigString'='/opt/pbspro/PBSPro_13.0.0/default/bin'  
    $end 'Scheduler'  
  
$end 'AnsoftCOMDaemon'
```

SGE (including Oracle Grid Engine, Open Grid Scheduler and Univa)

In the 'AnsoftCOMDaemon' block of the ansoftrmservice.cfg configuration file, the SchedulerName should be set to sge, and the ConfigString should be set to the directory containing the SGE commands. The AnsoftRSMService must be started with the environment configured for submission of jobs to the Linux cluster.

Example ansoftrmservice.cfg configuration file:

```
$begin 'AnsoftCOMDaemon'  
  
    $begin 'Managed COM Servers'  
    $end 'Managed COM Servers'  
    $begin 'Scheduler'  
        'SchedulerName'='sge'  
        'ConfigString'='/share/univa/bin/lx-amd64'  
    $end 'Scheduler'  
  
$end 'AnsoftCOMDaemon'
```

Registering ANSYS Electromagnetics Product Engines for use with RSM

You must make each installation of RSM aware of each ANSYS Electromagnetics product, so that remote simulation requests are properly routed to the product. For network installations, you need only to register each ANSYS Electromagnetics product once, regardless of the number of machines. Otherwise, you need to register each ANSYS Electromagnetics product on each machine running RSM.

Registration of ANSYS Electromagnetics products is handled through **RegisterEngineswithRSM.pl**, located in the product installation

directory (for example, /apps/ansoft/AnsysEM21.1/RegisterEnginesWithRSM.pl).

RegisterEnginesWithRSM.pl syntax

<product-specific installation directory>/RegisterEnginesWithRSM.pl
command, where *command* includes:

- **add** : Add this product installation's RSM engines to RSM
- **help** : Show full details of all RegisterEnginesWithRSM.pl command options
- **remove** : Remove this product installation's engines from RSM
- **status** : Show RSM registration status for each RSM engine

Supported Job Schedulers

The supported job schedulers of ANSYS Electromagnetics Suite 2021 R1 are listed in [Job Schedulers and Queuing Systems Support](#).

The versions earlier than the latest tested version indicated in the table may also work fine and might have been tested in previous versions of the product. Newer versions may also work if backward compatibility is maintained by the scheduler.

Custom Scheduler Integration

For the above job schedulers in a standard configuration, ANSYS Electromagnetics products provide an out-of-the-box integrated solution. For integration with other schedulers or non-standard scheduler configurations, user (or scheduler vendor) needs to provide a component called "custom scheduler proxy." Using this small component, AnsysEM mainly identifies the compute nodes allocated by a scheduler to the AnsysEM job, and launches solver processes on those nodes.

A custom scheduler proxy is usually an IronPython script file that has a class that implements a specified interface. Functionally, the methods of this class mediate the interactions between AnsysEM processes and the scheduler. Physically, such a proxy is a small text file, and does not require to be compiled/linked. This is often the preferred approach, as it is easier to install and test. The details of this interface and its methods are described elsewhere.

Installing a custom scheduler proxy usually means just adding a .py file under schedulers subdirectory of the AnsysEM installation directory. By providing a custom scheduler proxy written in IronPython,

you can even override the default scheduler proxy for a supported scheduler.

Note: IronPython is an implementation of the Python programming language targeting the .NET runtime. You do not need to install IronPython, as AnsysEM already comes with IronPython. AnsysEM supports it on both Windows and Linux.

Alternately, a custom scheduler proxy can be a dynamically linked library (on Microsoft Windows) that implements a set of functions. This library is loaded by the AnsysEM at runtime, and if it is running as part of a scheduler job, the AnsysEM interacts with the library to get information from the scheduler, and to start additional processes on specified hosts. Each such library implements a set of extern "C" functions needed to mediate the interactions between AnsysEM and the scheduler. The details of these functions are described elsewhere.

Limitations

Implementing your own custom scheduler proxy does not allow use of the job management GUI. It also does not provide tight-integration of scheduler for MPI-based jobs.

Running HPC Test Jobs on a Cluster

This section describes how to submit and run a small number of test jobs to verify that the product installation, the cluster configuration, and the MPI installation and configuration are satisfactory for running ANSYS Electromagnetics analysis jobs. These tests are not comprehensive, but successful completion of all of the jobs indicates that the basic requirements for running ANSYS Electromagnetics jobs have been met. Before running these tests, the following steps should be completed:

- The ANSYS Electromagnetics product or products should be installed
- The cluster should be configured for running ANSYS Electromagnetics analysis jobs, as described in "[Scheduler Specific Configuration and Testing](#) " on page 3-33
- MPI should be installed and configured as described in "[Message Passing Interface \(MPI\) Installation](#) " on page 5-1

Each of the test jobs should run to completion in a few minutes or less. They should be run in the order specified, proceeding to the

next test only if any problems found in the preceding test have been addressed. Three different test jobs are described.

- [Non-distributed job](#)
- [Non-MPI job distributed to multiple hosts](#)
- [MPI job distributed to multiple hosts](#)

Non-distributed job

This is a simple job that uses one core on one host. This test uses the Tee.aedt example project, which is located in the "Examples/HFSS/RF Microwave" subdirectory of the product installation directory. The files Tee.PDF and Tee.aedt are needed for this test. This test uses Setup1, but it does not use the frequency sweep, Sweep1, so the frequency sweep must be disabled.

Use the following steps to copy the project files and disable the frequency sweep, if needed.

1. Copy the files Tee.aedt and Tee.PDF from the Examples/HFSS/RF Microwave subdirectory of the product installation directory to another directory.
2. Run the product using the ansyedt script file for Linux.
3. Open the copy of the project.
4. In the Project Manager window, expand the tree item TeeModel (Driven Model)>Analysis>Setup1>Sweep1. If frequency sweep Sweep1 is enabled, right-click on TeeModel (Driven Model)>Analysis>Setup1> Sweep1, and select the item **Disable Sweep** in the popup menu to disable the frequency sweep.
5. Save and close the project.

Successful completion of this job indicates that the following requirements have been met:

- The job submission host is correctly configured for submitting jobs to the cluster.
- The job user account is authorized to run jobs on the cluster.
- The product installation is accessible from the cluster host allocated to the job.
- The project directory is accessible from the cluster host allocated to the job.
- Communication between different job processes on the cluster host allocated to the job is not blocked.

Non-MPI job distributed to multiple hosts

This job uses multiple cores distributed to multiple hosts. It does not use MPI. This job will start on one host, but portions of the analysis will be distributed to the other hosts allocated to the job. This test also uses the Tee.aedt example project, which is located in the Examples/HFSS/RF Microwave subdirectory of the product installation directory. The files Tee.PDF and Tee.aedt are needed for this test, the same as the previous test. This test uses Setup1, and it uses the frequency sweep, Sweep1, so the frequency sweep must be enabled.

Use the following steps to copy the project files and enable the frequency sweep, if needed.

1. Run the product using the ansysedt script file for Linux.
2. Open the copy of the project.
3. In the Project Manager window, expand the tree item TeeModel (Driven Model)>Analysis>Setup1>Sweep1. If frequency sweep Sweep1 is disabled, right click on TeeModel (Driven Model)>Analysis>Setup1>Sweep1, and select the item **Enable Sweep** in the popup menu to enable the frequency sweep.
4. Save and close the project.

Successful completion of this job indicates that the following additional requirements have been met:

- A job process on one cluster host allocated to the job is able to use the cluster to start additional job processes on other cluster hosts allocated to the job.
- Communications between job processes running on different cluster hosts allocated to the job are not blocked.

MPI job distributed to multiple hosts

This job uses MPI processes distributed to multiple hosts. This job will start on one host, but portions of the analysis will be distributed to MPI ranks running on the first execution host and ranks running on other hosts allocated to the job. This test uses the differential_stripline.aedt example project, which is located in the Examples/HFSS/Transmission Lines subdirectory of the product installation directory. The files differential_stripline.pdf and differential_stripline.aedt are needed for this test. This test uses Setup1, but it does not use the frequency sweep, DC_10GHz, so the frequency sweep must be disabled.

Use the following steps to copy the project files and disable the frequency sweep, if needed.

1. Copy the files `differential_stripline.aedt` and `differential_stripline.pdf` from the "Examples/HFSS/Transmission Lines" sub-directory of the product installation directory to another directory.
2. Run the product using the `ansyedt` script file for Linux.
3. Open the copy of the project. In the Project Manager window, expand the tree item `differential_stripline` -> `Differential Stripline (Driven Terminal)>Analysis>GHz>DC_10GHz`. If frequency sweep `DC_10GHz` is enabled, right-click on `differential_stripline>Differential Stripline (Driven Terminal)>Analysis>GHz>DC_10GHz`, and select the item **Disable Sweep** in the popup menu to disable the frequency sweep.
4. Save and close the project.

Successful completion of this job indicates that the following additional requirements have been met:

- If needed, MPI is installed and configured on the cluster hosts.
- If needed, user passwords are registered with MPI on the cluster hosts.
- If needed, the MPI prerequisites are installed on the cluster hosts.

Submitting and Running the Tests

This section includes general instructions for submitting and running the jobs. For additional notes on running any of the test jobs on a specific cluster, see the information for the specific scheduler in "[Scheduler Specific Configuration and Testing](#)" on page 3-33

The test jobs may be submitted to a cluster using the Ansys Electromagnetics job submission GUI. To use this GUI, start the product, using the `ansyedt` script file for Linux, on a host which is configured as a submission host for the cluster. The product should be started in an environment configured for submitting jobs to the cluster.

Before submitting the first job, the scheduler must be selected. Use menu item **Tools>Job Management>Select Scheduler** to pop up the **Select Scheduler** dialog.

For a Windows HPC cluster, select **Windows HPC** in the **Choose scheduler** drop-down list. Enter the name of the head node in the **Head node** edit control. The "..." button may be used to access a dialog to browse computers accessible via a network. For other clusters, select the scheduler type (lsf, pbs, or sge) in the **Choose scheduler** drop-down list. Select lsf for IBM Platform LSF. Select pbs for PBS Professional or PBS/Torque. Select SGE for Sun Grid Engine, Oracle Grid Engine, Univa Grid Engine, or Open Grid Scheduler/Grid Engine.

Pressing the **Refresh** button causes the connection to the cluster to be verified, and basic cluster information to be displayed in the **Scheduler info** text control. When finished, click **OK**. The selected scheduler (and head node, if applicable) will be used for subsequent jobs submission using the Submit Job GUI. To change the scheduler or head node, pop up the **Select Scheduler** dialog again, and select a different scheduler or head node.

To submit a job, use menu item **Tools>Job Management>Submit Job** to pop up the **Submit Job To:** dialog. The controls applicable to options that are used less frequently are only shown if the **Show advanced options** check box is checked. Three tabs are shown on this dialog:

- **Analysis Specification:** Basic analysis parameters, including the product pathname, the project pathname, and the setup or setups to analyze. Several additional options may be specified using this tab.
- **Compute Resources:** The user selects the resources to be allocated to the job and options for job distribution. Scheduler specific resources may be specified using the "Resource selection parameters."
- **Scheduler Options:** Options for the job to be submitted to the scheduler, such as the job name and job priority.

Use the **Preview Submission** button to display information about the job to be submitted to the cluster. The dialog shows the job submission command and the contents of the job script for the job. If a problem is detected with the selected job parameters, a message box containing an error message may be shown, instead of the job submission parameters.

Use the **Submit Job** button to submit the job to the selected cluster. If there is an error submitting the job, then an error message will be shown. If the job is submitted successfully, then a dialog box showing the job ID is displayed to the user. If the **Begin monitoring this job now** check box is checked when this dialog is dismissed,

then the **Monitor Job** dialog pops up to monitor the job. The **Monitor Job** dialog may be dismissed at any time, and the job will continue to run. To resume monitoring the same job, select menu item **Tools>Job Management>Monitor Jobs** to access the **Monitor Job** dialog. Selecting a project file using the **Recent Jobs** drop-down list or the file browser popped up by the "..." button will resume monitoring of the most recent job submitted for the selected project. Monitoring will not be possible from the **Monitor Job** dialog if there is no network connection between the submission host and the cluster execution host where the Desktop process runs, or if this connection is blocked.

"Submitting and running a non-distributed job " below

"Submitting and running a non-MPI job distributed to multiple hosts" on page 3-28

"Submitting and running an MPI job distributed to multiple hosts" on page 3-30

Submitting and running a non-distributed job

1. Copy the files Tee.aedt and Tee.PDF from the Examples/HFSS/RF Microwave subdirectory of the product installation directory to a directory which is accessible from the cluster hosts using the same path. This directory, the project directory, must be readable and writeable by the user running the job. Disable the frequency sweep as described above.

The scheduler (and head node, if applicable) should be set as described above before submitting the job to the cluster.

2. Select menu item **Tools>Job Management>Submit Job** to activate the **Submit Job To:** dialog.
3. In the **Analysis Specification** tab, ensure that the **Product path** edit control contains a pathname of the product that is valid for the cluster execution nodes.
 - a. Specify the path to the project file in the **Project** edit control.
 - b. Specify the setup "TeeModel:Nominal:Setup1" in the **Analysis Setups** group box.
4. Switch to the **Compute Resources** tab, and make sure that the **Use automatic settings** check box is not checked.
 - a. For the Method, select **Number of Tasks and Cores** in the drop-down list.
 - b. Set the **Total number of tasks** to 1.

- c. Set the **Cores per distributed task** to 1.
 - d. Ensure that the **Limit number of tasks per node to:** check box is not checked. The other options on all tabs may be set to any values.
5. Press the **Preview Submission** button. If no errors are found, then the **Preview Submit Job Results** dialog should display information on how the job will be submitted to the cluster. If an error is displayed instead of the **Preview Submit Job Results** dialog, then address the error before continuing. This dialog should indicate that only one host is needed for the job. If this is not the case, then review the job resources specified on the **Compute Resources** tab, and retry. If the job and task parameters are correct, then close the **Preview Submit Job Results** dialog and continue.
 6. Submit the job using the **Submit Job** button. If an error occurs on job submission, address the error before submitting the job again.

If the job fails to complete successfully, check the following items:

- The user submitting the job must be authorized to run jobs on the cluster.
- The pathname of the product must be accessible on the cluster execution node(s) where the job runs, and the job user must have permission to run the product executables.
- The pathname of the project must be accessible on the cluster execution node(s) where the job runs. The job user must have permission to read from and write to the directory containing the project file, and there must be adequate disk space for the job results.
- The temp directory used by the job must exist and it must have adequate disk space for the job temporary files.
- Check for errors or warnings in the **Monitor Job** dialog, and address these issues.
- Check the cluster for errors or warnings related to the job and address these issues. The job ID is displayed in a dialog when the job is successfully submitted. The job ID of the last job run on the specified project is also shown in the **Monitor Job** dialog.

Submitting and running a non-MPI job distributed to multiple hosts

If the cluster type is not SGE, the cluster must have at least two hosts with the same number of cores per host. Let N be the smallest number of cores per host, such that the cluster contains at least two hosts with N cores per host. The job will be submitted to run on two hosts with N cores per host, requesting a total of $2 * N$ cores. If the cluster type is SGE (which includes SGE, OGE, UGE, or OGS/GE), then select an integer N such that there is a parallel environment with an integral allocation rule of N , and there are at least two hosts with at least N cores that may be used by the job. For many clusters, selecting $N = 4$ may be a good choice. The parallel environment with an integral allocation rule of N should be specified when the job is submitted.

1. Copy the files Tee.aedt and Tee.PDF from the "Examples/HFSS/RF Microwave" subdirectory of the product installation directory to a directory which is accessible from the cluster hosts using the same path. This directory, the project directory, must be readable and writable by the user running the job. Enable the frequency sweep as described above.

The scheduler (and head node, if applicable) should be set as described above before submitting the job to the cluster.

2. Select menu item **Tools>Job Management>Submit Job** to activate the **Submit Job To:** dialog.
3. In the **Analysis Specification** tab, ensure that the **Product path** edit control contains a pathname of the product that is valid for the cluster execution nodes.
 - a. Specify the path to the project file in the **Project** edit control.
 - b. Specify the setup "TeeModel:Nominal:Setup1" in the **Analysis Setups** group box.
4. Switch to the **Compute Resources** tab, and make sure that the **Use automatic settings** check box is not checked.
 - a. For the Method, select **Number of Tasks and Cores** in the drop-down list.
 - b. Set the **Total number of tasks** to $2 * N$.
 - c. Set the **Cores per distributed task** to 1.
 - d. Ensure that the **Limit number of tasks per node to:** check box is not checked. The other options on all tabs may be set to any values.

- e. In the **Job distribution** group box, press the **Modify** button to access the **Job Distribution** dialog.
 - f. In the **Enable Distribution Types** section, uncheck the **Use defaults** check box, and check only the **Frequencies** check box.
 - g. In the **Distribution levels** section, select the **Single level only** radio button.
 - h. Click **OK** to close the **Job Distribution** dialog.
5. For the schedulers which support setting the minimum and maximum cores per node for a job (LSF), set both of these values to the value N, determined earlier, as follows:
- a. Click on the "... " button to activate the **Compute Resource Selection Parameters** dialog.
 - b. For the MinCoresPerNode parameter, check the **Specify Parameter** check box, and set the value to N.
 - c. For the MaxCoresPerNode parameter, check the **Specify Parameter** check box, and set the value to N.
 - d. Click **OK** to close the **Compute Resource Selection Parameters** dialog.

For SGE schedulers (including SGE, OGE, UGE, and OGS/GE), specify a parallel environment with an integral allocation rule equal to the value N, selected earlier.

- a. Determine which parallel environments have an allocation rule equal to the value N. Then, click on the "... " button to activate the **Compute Resource Selection Parameters** dialog.
- b. For the ParallelEnvironment parameter, check the **Specify Parameter** check box, and select one of the parallel environments having an allocation rule equal to the value N from the drop-down list.
- c. Click **OK** to close the **Compute Resource Selection Parameters** dialog.

The other options on all tabs may be set to any values.

6. Press the **Preview Submission** button. If no errors are found, then the **Preview Submit Job Results** dialog should display information on how the job will be submitted to the cluster. If an error is displayed instead of the **Preview Submit Job Results** dialog, then address the error before continuing. This dialog should indicate that a total of $2 * N$ cores are required for the job, and that N cores will be needed on each host allocated to the job. If this is not the case, then review the job

resources specified on the **Compute Resources** tab, and retry. If the job and task parameters are correct, then close the **Preview Submit Job Results** dialog and continue.

7. Submit the job using the **Submit Job** button. If an error occurs on job submission, address the error before submitting the job again.

If the job fails to complete successfully, check the following items, in addition to the items mentioned for the first test:

- The Desktop process will use the cluster to launch processes on the other hosts allocated to the job. Verify that the cluster configuration allows a job to launch processes on the other hosts allocated to the job.
- Job processes on different hosts allocated to the job need to communicate. Verify that communication between these job processes is not blocked.

Submitting and running an MPI job distributed to multiple hosts

Use the Intel MPI Vendor will be used for MPI jobs. This test should use the same MPI Vendor that will be used by cluster users. If no MPI Vendor is specified by the user, then the default MPI Vendor (Intel) is used. Intel MPI is denoted by the string "Intel."

If the cluster type is not SGE, then the cluster must have at least two hosts with the same number of cores per host. Let N be the smallest number of cores per host, such that the cluster contains at least two hosts with N cores per host. The job will be submitted to run on two hosts with N cores per host, requesting a total of $2 * N$ cores. If the cluster type is SGE (which includes SGE, OGE, UGE, or OGS/GE), then select an integer N such that there is a parallel environment with an integral allocation rule of N, and there are at least two hosts with at least N cores that may be used by the job. For many clusters, selecting $N = 4$ may be a good choice. The parallel environment with an integral allocation rule of N should be specified when the job is submitted.

1. Copy the files `differential_stripline.aedt` and `differential_stripline.pdf` from the `Examples/HFSS/Transmission Lines` sub-directory of the product installation directory to a directory which is accessible from the cluster hosts using the same path. This directory, the project directory, must be readable and writable by the user running the job. Disable the frequency sweep as described above.

- The scheduler (and head node, if applicable) should be set as described above before submitting the job to the cluster.
2. Select menu item **Tools>Job Management>Submit Job** to activate the **Submit Job To:** dialog.
 3. In the **Analysis Specification** tab, ensure that the **Product path:** edit control contains a pathname of the product that is valid for the cluster execution nodes.
 - a. Specify the path to the project file in the **Project:** edit control.
 - b. Specify the setup Differential Stripline:Nominal:3GHz in the **Analysis Setups** group box.
 - c. In the **Analysis options** group box, set the **HFSS/MPIVendor** batchoption to the desired MPI Vendor, "Intel". If this batchoption is present, but set to a different value, then you can use the **Edit** button to change the value; edit the value in the **Edit Batchoption** dialog and press the **Update** button to register the new value before closing the dialog with the **Done** button. If it is not set, then use the **Add** button to add it; in the **Add Batchoption** dialog, select the registry key "HFSS/MPIVendor" in the grid, and then edit the text in the **Value** control, if needed, and press **Add** to register the new value, before closing the dialog with the **Done** button.
 4. Switch to the **Compute Resources** tab, and make sure that the **Use automatic settings** check box is not checked.
 - a. For the Method, select **Number of Tasks and Cores** in the drop-down list.
 - b. Set the **Total number of tasks** to $2 * N$.
 - c. Set the **Cores per distributed task** to 1.
 - d. Ensure that the **Limit number of tasks per node to:** check box is not checked.
 - e. In the **Job distribution** group box, press the **Modify** button to access the **Job Distribution** dialog.
 - f. In the **Enable Distribution Types** section, uncheck the **Use defaults** check box, and check only the **Direct Solver Memory** check box.
 - g. In the **Distribution levels** section, select the **Single level only** radio button.
 - h. Click **OK** to close the **Job Distribution** dialog.
 5. For the schedulers which support setting the minimum and maximum cores per node for a job (LSF), set both of these values

to the value N, determined earlier, as follows:

- a. Click on the "..." button to activate the **Compute Resource Selection Parameters** dialog.
- b. For the MinCoresPerNode parameter, check the **Specify Parameter** check box, and set the value to N.
- c. For the MaxCoresPerNode parameter, check the **Specify Parameter** check box, and set the value to N.
- d. Click **OK** to close the Compute Resource Selection Parameters dialog.

For SGE schedulers (including SGE, OGE, UGE, and OGS/GE), specify a parallel environment with an integral allocation rule equal to the value N, selected earlier.

- a. Determine which parallel environments have an allocation rule equal to the value N. Then, click on the "..." button to activate the **Compute Resource Selection Parameters** dialog.
- b. For the ParallelEnvironment parameter, check the **Specify Parameter** check box, and select one of the parallel environments having an allocation rule equal to the value N from the drop-down list.
- c. Click **OK** to close the **Compute Resource Selection Parameters** dialog.

The other options on all tabs may be set to any values.

6. Press the **Preview Submission** button. If no errors are found, then the **Preview Submit Job Results** dialog should display information on how the job will be submitted to the cluster. If an error is displayed instead of the **Preview Submit Job Results** dialog, then address the error before continuing. This dialog should indicate that a total of $2 * N$ cores are required for the job, and that N cores will be needed on each host allocated to the job. If this is not the case, then review the job resources specified on the **Compute Resources** tab, and retry. If the job and task parameters are correct, then close the **Preview Submit Job Results** dialog and continue.
7. Submit the job using the **Submit Job** button. If an error occurs on job submission, address the error before submitting the job again.

If the job fails to complete successfully, in addition to the items mentioned for the first and second tests, MPI configuration issues may cause job failures. Check the following MPI configuration issues:

- Verify that the Intel MPI Vendor is specified using a batch-option setting in the **Submit Job To:** dialog. The controls for specifying batchoptions are located in the **Analysis Options** group box on the **Analysis Specification** tab of the **Submit Job To:** dialog.
- Verify that the specified MPI Vendor is installed (if needed) and configured correctly for the user submitting the job, as described in ["Message Passing Interface \(MPI\) Installation "](#) on [page 5-1](#) Pay close attention to the registration of user passwords because the lack of a cached password when one is required, or a cached password that is incorrect or out of date (even if no cached password is required) may cause authentication errors. Note that passwords are cached separately on each cluster host, so a job may have authentication errors on some hosts but not on other hosts.

Scheduler Specific Configuration and Testing

The sections below describe configuration requirements for each supported scheduler. Methods for running tests to verify the configuration are included.

["IBM Platform LSF" below](#)

["Grid Engine \(SGE, UGE, OGE\)" on page 3-36](#)

["PBSPro" on page 3-45](#)

["PBS/Torque" on page 3-48](#)

IBM Platform LSF

Cluster Configuration Guidelines

To allow users to submit exclusive jobs, the users must have access to queues which are configured with the EXCLUSIVE property set to Y. Appropriate queues should be configured with the EXCLUSIVE property set if users need to submit exclusive jobs.

Using the ANSYS Electromagnetics job submission GUI, a user may submit jobs where the user specifies tasks automatically and specifies the "Number of Cores and (Optional) RAM." However, the RAM per core may only be specified if the LSF cluster is configured with parameter RESOURCE_RESERVE_PER_SLOT set. This LSF cluster configuration parameter should be set to allow user to specify the RAM per core for the job.

Running Test Jobs

If there is a problem connecting to the scheduler when the job submission node is the local node (where the GUI is running) and the lsf scheduler is specified, check the following items:

- The cluster is up and running normally.
- The local node is configured as a job submission node for the cluster.
- The user running the GUI has permission to submit jobs to the cluster.
- The environment was configured for submission of jobs to the cluster before starting the product, `ansysedt`.

Non-distributed job

- If the job fails to complete successfully, check the following items:
- The user submitting the job must be authorized to run jobs on the cluster.
- The pathname of the product must be accessible by the user running the job from the cluster execution node(s) where the job runs.
- The pathname of the project must be accessible by the user running the job from the cluster execution node(s) where the job runs. The directory containing the project file must be writable by the user running the job from the cluster execution node(s) where the job runs, and there must be adequate disk space for the job results.
- The temp directory specified by `batchoptions` or using the installation default setting on the cluster execution node(s) where the job runs must have adequate disk space for the job temporary files.
- Check for errors or warnings in the **Monitor Job** dialog, and address these issues.

Non-MPI job distributed to multiple hosts

This job is submitted to run on two hosts with N cores per host, requesting a total of $2 * N$ cores, where N is the smallest number of cores per host, such that the cluster contains at least two hosts with N cores per host.

For this test, the product, `ansysedt`, runs on one of the hosts, and the LSF `blaunch` command is used to start an `ansoftrsm` process on the

other host. The `ansoftrsm` process is used to start any additional processes needed by the job on the other host. In some cases, multiple `ansoftrsm` processes are started on the other host, but only one such process should be running on a given host at any time.

If the job fails to complete successfully, check the following items, in addition to the items mentioned for the first test:

- Verify that the LSF `blaunch` command is able to start the `ansoftrsm` process on the other host.

MPI job distributed to multiple hosts

This job is submitted to run on two hosts with N cores per host, requesting a total of $2 * N$ cores, where N is the smallest number of cores per host, such that the cluster contains at least two hosts with N cores per host.

Normally, passwordless `ssh` is used to start MPI tasks on remote hosts. The `RemoteSpawnCommand` batchoption setting may be used to specify whether to use passwordless `ssh` or another option. Here are the allowed values of the `RemoteSpawnCommand` batchoption setting, along with their meanings:

- **SSH** (the default): Use passwordless `ssh` for launching remote MPI processes
- **RSH**: Use the `rsh` (remote shell) command for launching remote MPI processes
- **Scheduler**: Use the scheduler for launching remote MPI processes

If the **SSH** option is specified, then cluster users are required to configure passwordless `ssh` on the cluster hosts. If the **RSH** option is specified, then the `rsh` (remote shell) command must be enabled on the cluster hosts. The "Scheduler" batchoption value is only allowed for jobs running under an LSF or SGE cluster and only if the MPI Vendor is Intel; no additional configuration is needed in this case. See "[Message Passing Interface \(MPI\) Installation](#)" on page 5-1 for more information.

- Verify that the correct MPI Vendor ("Intel") is specified via a batchoption setting when submitting the job; otherwise the default, "Intel," is used for the job.
- Verify that the specified MPI Vendor is installed and configured correctly for the user submitting the job, as described in "[Message Passing Interface \(MPI\) Installation](#)" on page 5-1. Note that for Intel MPI, MPI password registration is done

separately for each cluster host. If authentication errors occur for some cluster hosts but not for other cluster hosts, then the errors may be due to an incorrect password or no password registered on some hosts.

Grid Engine (SGE, UGE, OGE)

This section covers the various versions of Grid Engine. The recent versions include Sun Grid Engine (SGE). When Sun was acquired by Oracle, this product was renamed to Oracle Grid Engine (OGE). Later, Univa took over support for this product. The Univa version is called Univa Grid Engine (UGE). UGE includes some features not available in the earlier versions. A supported open source version based on a previous (open source) version of Sun Grid Engine is known as Open Grid Scheduler/Grid Engine. The comments in this section are generally applicable to all of these versions, although there are some differences between them. This is especially true for Univa Grid Engine, which has some features not available on the other versions.

Jobs may be submitted via the command line or via the job submission GUI. When using the GUI, the software will create the ANSYS Electromagnetics product command line as well as the job submission command line (the bsub command line) from the parameters specified by the user in the GUI. The options of the ANSYS Electromagnetics product command line will be selected to effectively use the resources allocated to the job by the scheduler. The user may specify a specific queue and/or parallel environment for the job using the GUI. If no parallel environment is specified by the user, then a parallel environment will be automatically selected by the job submission software. If no queue is specified by the user, then the job submission command will not contain a queue specification, and the scheduler will select the queue for the job. The state of the **Use automatic settings** check box and other job submission GUI controls will affect the job submission command options.

Cluster Configuration Guidelines

qrsh Command

ANSYS EM parallel batch jobs use the SGE qrsh command to launch engine processes on remote hosts. If the qrsh command is not working correctly, then the parallel job is unable to launch engine processes on remote hosts. If this problem occurs, the batch log for the job typically includes one or more error messages indicating that a COM engine was unable to be started on a remote host. If this

occurs, the user or cluster administrator should verify that the SGE `qrsh` command is working correctly, and correct the problem if the SGE `qrsh` command is not working correctly.

The `qrsh` command may be tested by running a simple command on a specified host, such as `qrsh -l hostname=host1 hostname` or `qrsh -l hostname=host1 ls /tmp`, where `host1` is the remote host name. The first test should echo back the hostname of the remote machine. The second test should list the contents of the `/tmp` directory on the remote machine.

The failures of the SGE `qrsh` command are associated with the following global sge configuration parameters, listed below with values that may cause the failures:

- `qrsh_command /usr/bin/ssh -t`
- `rsh_command /usr/bin/ssh -t`
- `rlogin_command /usr/bin/ssh -t`

If these parameter settings are removed or changed to builtin, then the SGE built-in mechanisms are used for `qrsh`, `rsh`, and `rlogin`. No problems with the built-in versions have been reported. The SGE `qconf -sconf` global command may be used to view these parameter settings. The SGE `qconf -mconf` global command may be used to modify these parameter settings.

Parallel Environments

For these clusters, no parallel environment is needed to submit a job that is allocated only one core on one host. A parallel environment is required to submit a job that requires two or more cores, whether the cores are on the same host or on different hosts. We expect that most users will want to submit distributed jobs at least occasionally, either to run larger jobs or to complete the job more quickly. Appropriate parallel environments should be configured for the cluster to run the types of parallel jobs that the users would like to run.

All parallel environments for ANSYS Electromagnetics jobs must be configured with the `control_slaves` parameter set to `TRUE`.

If the **Use automatic settings** check box is not checked, then you must also specify the total number of tasks, `T`, for the job, and the number of cores per task, `C`. For this case, the allocation rule of the parallel environment must be an integer. The software will allocate the same number of tasks, `N`, on each node, where `N` is determined

by the nodes available for the job and the user-specified queue and/or parallel environment, if any. A minimum of $N * C$ cores is required per node, so the allocation rule must be an integer that is equal to or greater than $N * C$. Note that if the allocation rule is an integer greater than $N * C$, then the extra cores allocated to the job are not used by the ANSYS Electromagnetics product, leading to less than ideal resource usage. We recommend configuring multiple parallel environments with different integral allocation rules to allow a good match of resources needed by the job to the resources allocated to the job. For example, if all nodes have 8 cores, configuring parallel environments with allocation rules from 1 to 8 will allow a good resource match for a variety of jobs with different numbers of cores per task specified by the users.

If the **Use automatic settings** check box is checked, then you can specify the **Number of Nodes and Cores** for the job or the **Number of Cores and (Optional) RAM** for the job. For the first case, the **Number of Nodes and Cores** specified for the job, the job is allocated the same number of cores on each node, and the parallel environment must be an integer, as for the case when the **Use automatic settings** check box is not checked. The same parallel environment configurations are recommended for this case. For the second case, the **Number of Cores and (Optional) RAM** specified for the job, a parallel environment with any allocation rule may be used. We recommend that at least one parallel environment with an allocation rule of `$fill_up` be provided, as this allows the job to be mostly run on hosts where no other job is running. An allocation rule of `$round_robin` may also be useful to allow such jobs to be distributed across all available hosts.

In summary, we recommend that parallel environments with the following configurations be provided for ANSYS Electromagnetics parallel jobs. At least one parallel environment with an allocation rule of `$fill_up`. At least one parallel environment with an allocation rule of `$round_robin`. A selection of parallel environments with integral allocation rules from 1 up to the maximum number of cores per host for the hosts associated with the parallel environment. All of the parallel environments for ANSYS Electromagnetics jobs must be configured with the `control_slaves` parameter set to `TRUE`.

Submitting Exclusive Jobs

Some users require or desire the ability to submit jobs to a cluster that have exclusive access to each host allocated to the job. That

is, any host running the job is not available for use by any other jobs. This may be needed to limit contention for resources that are not tracked by the cluster, such as network bandwidth. By default, an SGE cluster may not be configured for exclusive jobs. If the cluster configuration includes a complex of type `BOOL`, with `relop EXCL`, that is both consumable and requestable, then this complex may be used to submit exclusive jobs. If the cluster configuration does not contain such a complex, then step 1 below indicates how to add an exclusive complex to the cluster configuration. The value of this complex must be set to `TRUE` for each execution host, as shown in step 2 below. Step 2 also indicates how to check this requirement if an exclusive resource is already configured for the cluster.

Configuring an exclusive resource

1. If there is no complex configured for the cluster of type `BOOL`, with `relop EXCL`, that is both consumable and requestable, then this step shows how to configure such a complex. Use the command `qconf -mc` to add a new complex to the table of complexes. Recommended attributes are:
 - `name : exclusive`
 - `shortcut : excl`
 - `type : BOOL`
 - `relop : EXCL`
 - `requestable : YES`
 - `consumable : YES`
 - `default : 0`
 - `urgency : 0`
2. After the complex is configured, the value of this complex for each execution host must be set. The SGE SGE command `qconf -me hostname`, where `hostname` is the name of the host, may be used to display the host configuration settings for a specific host. The `"complex_values"` line should include the setting `"exclusive_name=TRUE"`, where `exclusive_name` is the name of the exclusive complex. A backslash (`\`) at the end of a line indicates that the line is continued on the next line. Set the value of the exclusive complex to `TRUE` for each execution host using the SGE command `qconf -me hostname`. Add the setting `"exclusive_name=TRUE"` to the `complex_values` line or lines, where `exclusive_name` is the name of the exclusive complex. Different settings on the `complex_values` line are separated by commas.

Submitting an exclusive job on the command line

When submitting a job using the `qsub` command, the job will have exclusive access to each host allocated to the job if the exclusive resource name ("exclusive") or exclusive resource shortcut ("excl") is included in the resource list specified by the `qsub -l` option. If the resource list does not include the exclusive resource name or the exclusive resource shortcut, then the job will not have exclusive access to each host allocated to the job; other jobs may run on the same host or hosts at the same time as this job.

Example `qsub` command line for exclusive serial job:

```
qsub -b y -l excl /opt/AnsysEM/AnsysEM21.1/ansysedt -ng -  
BatchSolve -machinelist num=1 ~/Ansoft/OptimTee.aedt
```

Although serial jobs use only one slot, no other jobs will run on the host where this job is running, even if additional slots are present.

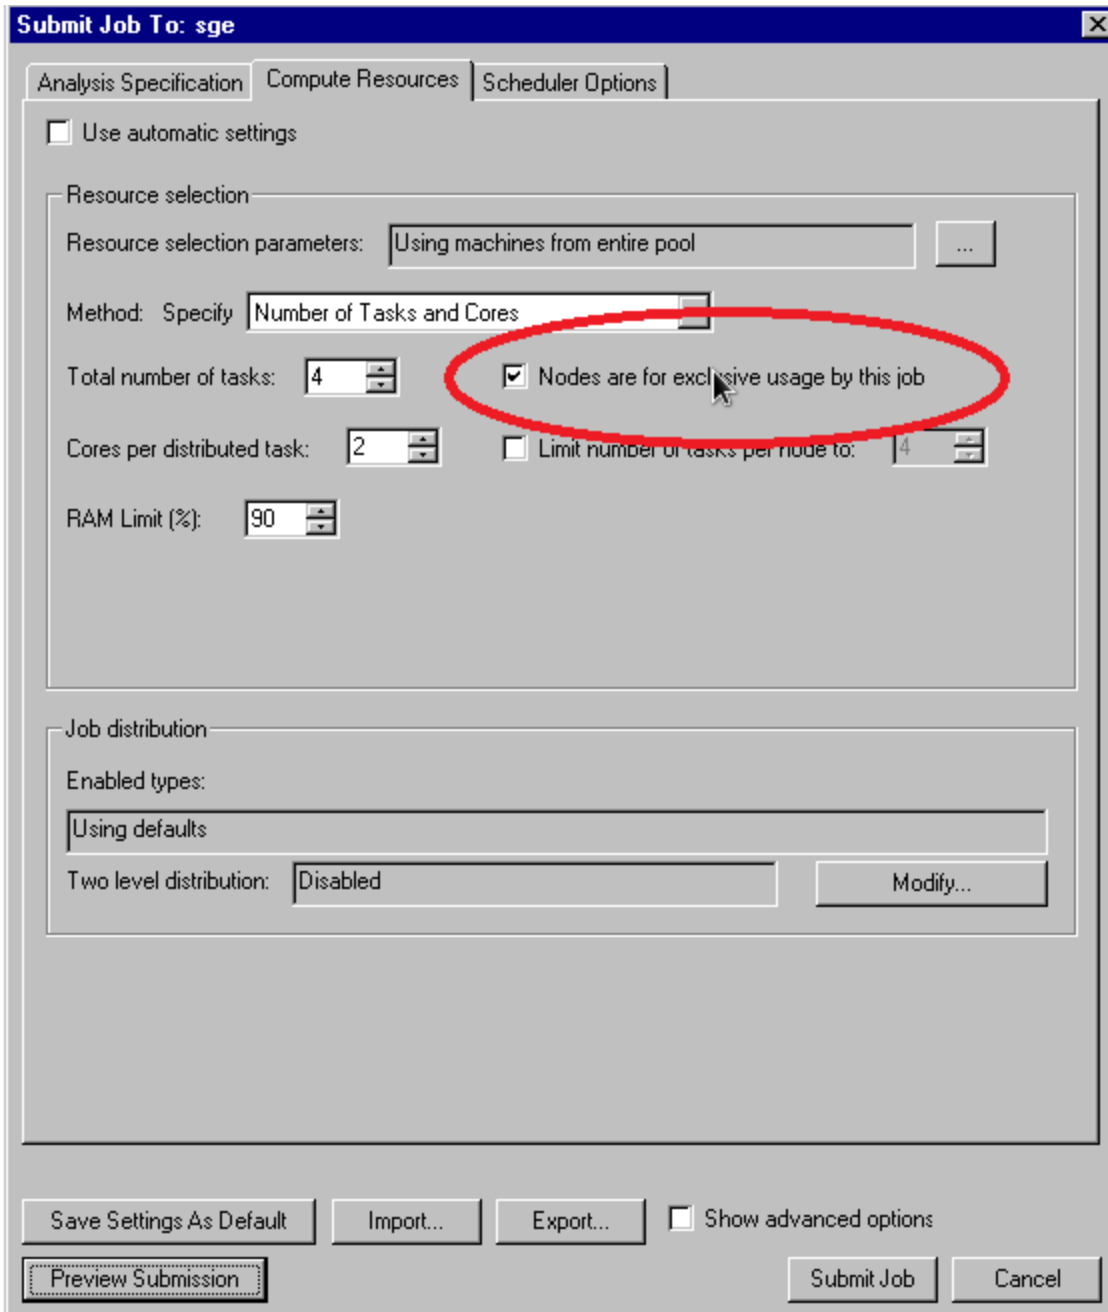
Example `qsub` command line for exclusive parallel job using eight engines, each using a single thread of execution:

```
qsub -b y -l excl -pe pe1 8 /opt/AnsysEM/AnsysEM21.1/ansysedt -  
ng -BatchSolve -Distributed -machinelist num=8 ~/An-  
soft/OptimTee.aedt
```

None of the hosts allocated to this job will be allowed to run other jobs while this job is running.

Submitting an exclusive job using the job submission GUI

For jobs submitted to an SGE cluster using the job submission dialog, the state of the **Nodes are for exclusive usage by this job** check box controls whether the job is submitted as an exclusive job. This check box is located on the **Compute Resources** tab of the job submission dialog as shown below. If this check box is checked, then the job submission command line will include the exclusive resource option. The resource name does not need to be "exclusive"; any complex of type `BOOL`, with `relop EXCL`, that is both consumable and requestable may be used.



Consumable Memory Limits

When submitting memory intensive jobs on an SGE cluster, it may be useful to specify a consumable memory limit for the job. If the resource request for the job specifies a consumable limit, then, while the job is running, the specified resource is consumed, that is, not available for other jobs. Jobs requiring the same resource will only be scheduled on a host if the total quantity of the

resource on the host exceeds the sum of the consumable resource needs of the jobs. For non-consumable resources, the resource requirement for each job is independently compared to the quantity of the resource available on the host. Consumable resources may be configured as per-slot resources or as per-job resources.

Consumable memory limits are important for jobs submitted to an SGE cluster using the job submission dialog, with the **Use automatic settings** check box checked, and the user specifying the **Number of Cores and (Optional) RAM** for the job. If the RAM per core is specified by the user, then the job submission command includes a consumable memory resource requirement. If a complex of type MEMORY is configured to be requestable and consumable (per-slot, rather than per-job), then this complex will be used for the consumable memory resource requirement to meet the RAM per core requirement specified by the user. If no such complex is configured, then the RAM per core may not be specified by the user.

Some versions of Grid Engine come with "m_mem_free" and "mem_free" complexes already configured. If either one or both are configured as a requestable and consumable per-slot resource, there is no more configuration required. For other grid engine versions, it may be necessary to configure a complex to use as a memory resource requirement to allow users to use automatic settings and specify the Number of Cores and RAM for the job.

To configure this resource, add a complex with the following attributes:

- name: m_mem_free or mem_free or another name, if desired
- shortcut: mfree or mf or another alias, if desired
- type: MEMORY
- relop: <=
- requestable: YES
- consumable: YES
- default: 0
- urgency: 0

The **SGE sconf -sc** command may be used to examine the complexes and their configurations.

The **SGE sconf -mc** command may be used to modify the complexes or their configurations.

In addition to adding the complex to the cluster for the requestable and consumable memory resource, each host must be configured to specify the limit of this resource. The SGE **sconf -sel** command may be used to list the execution hosts of the cluster. The SGE **sconf -se hostname** command may be used to examine the configuration of the host with name **hostname**. The SGE **sconf -me hostname** command may be used to modify the configuration of the host with name **hostname**. The memory limit should be added to the `complex_values` for the host. If there is no `complex_values`, then the `complex_values` entry should be added: the value of the `m_mem_free` complex must be the same as the amount of memory on the host. For example, if the complex name is `m_mem_free`, and the host has 32 GB of memory, then the `complex_values` of the host configuration should include the specification "`m_mem_free=32G`". Different complex values are separated by commas in the host configuration.

Running Test Jobs

If there is a problem connecting to the scheduler when the job submission node is the local node (where the GUI is running) and the sge scheduler is specified, check the following items:

- The cluster is up and running normally.
- The local node is configured as a job submission node for the cluster.
- The user running the GUI has permission to submit jobs to the cluster.
- The environment was configured for submission of jobs to the cluster before starting the product, `ansysedt`.

["Non distributed job" below](#)

["Non-MPI job distributed to multiple hosts" on the next page](#)

["MPI job distributed to multiple hosts" on the next page](#)

Non distributed job

If the job fails to complete successfully, check the following items:

- The user submitting the job must be authorized to run jobs on the cluster.
- The pathname of the product must be accessible by the user running the job from the cluster execution node(s) where the job runs.

- The pathname of the project must be accessible by the user running the job from the cluster execution node(s) where the job runs. The directory containing the project file must be writable by the user running the job from the cluster execution node(s) where the job runs, and there must be adequate disk space for the job results.
- The temp directory specified by batchoptions or using the installation default setting on the cluster execution node(s) where the job runs must have adequate disk space for the job temporary files.
- Check for errors or warnings in the **Monitor Job** dialog, and address these issues.

Non-MPI job distributed to multiple hosts

This job will be submitted to run on two hosts using N cores per host, requesting a total of $2 * N$ cores. Here, N is an integer such that there is a parallel environment with an integral allocation rule of N. The ParallelEnvironment parameter should be specified to be one of the parallel environments having an allocation rule equal to the value N when the job is submitted.

For this test, the product, ansysedt, will run on one of the hosts, and the SGE qsh command will be used to start an ansoftrsm process on the other host. The ansoftrsm process will be used to start any additional processes needed by the job on the other host. In some cases, multiple ansoftrsm processes will be started on the other host, but only one such process should be running on a given host at any time.

If the job fails to complete successfully, check the following items, in addition to the items mentioned for the first test:

- Verify that the SGE qsh command is able to start the ansoftrsm process on the other host.

MPI job distributed to multiple hosts

This job will be submitted to run on two hosts using N cores per host, requesting a total of $2 * N$ cores. Here, N is an integer such that there is a parallel environment with an integral allocation rule of N. The ParallelEnvironment parameter should be specified to be one of the parallel environments having an allocation rule equal to the value N when the job is submitted.

Normally, passwordless ssh is used to start MPI tasks on remote hosts. The RemoteSpawnCommand batchoption setting may be used to

specify whether to use passwordless ssh or another option. Here are the allowed values of the RemoteSpawnCommand batchoption setting, along with their meanings:

- SSH (the default): Use passwordless ssh for launching remote MPI processes
- RSH: Use the rsh (remote shell) command for launching remote MPI processes
- Scheduler: Use the scheduler for launching remote MPI processes

If the SSH option is specified, then cluster users are required to configure passwordless ssh on the cluster hosts. If the RSH option is specified, then the rsh (remote shell) command must be enabled on the cluster hosts. The "Scheduler" batchoption value is only allowed for jobs running under an LSF or SGE cluster and only if the MPI Vendor is Intel; no additional configuration is needed in this case. See ["Message Passing Interface \(MPI\) Installation "](#) on page 5-1 for more information.

If the job fails to complete successfully, in addition to the items mentioned for the first and second tests, MPI configuration issues may cause job failures. Check the following MPI configuration issues:

- Verify that the Intel MPI Vendor is specified via a batchoption setting when submitting the job.
- Verify that the RemoteSpawnCommand is specified via a batchoption setting when submitting the job; otherwise the default, SSH, is used for the job.
- Verify that the specified MPI Vendor is installed and configured correctly for the user submitting the job, as described in ["Message Passing Interface \(MPI\) Installation "](#) on page 5-1.

PBSPro

Cluster Configuration Guidelines

If the PBSPro configuration file pathname is different from the default, /etc/pbs.conf, then the PBS_CONF environment variable should be set to the pathname of this file. Cluster users need to ensure that the PBS_CONF environment variable is set appropriately when submitting jobs to the cluster and when running jobs on the cluster.

If the PBS_SERVER parameter in the PBSPro configuration file is not set to the default PBSPro server name, then the PBS_DEFAULT

environment variable should be set to the default server name. Cluster users need to ensure that the PBS_DEFAULT environment variable is set appropriately when submitting jobs to the cluster and when running jobs on the cluster.

Users who submit parallel jobs for ANSYS Electromagnetic products to the cluster using the job submission GUI may select between several different methods for specifying the resources for the job. If a user specifies the "Number of Nodes and Cores" for the job, with the number of tasks for the job determined automatically, then the job submission command resource specification will select the nodes with the greatest number of cores per node that meet the other job requirements. If the cluster is not homogeneous (that is, the cluster execution nodes have different configurations), this may not make best use of available resources.

Nodes with many cores per node may be allocated to the job, even if the job could have been distributed across more nodes, with each node having fewer cores per node. In some cases, a job may be allocated one or more nodes with many cores per node, even if the job could run on the same number of nodes with fewer cores per node. This issue may be mitigated by creating appropriate queues, and specifying the queue to use when submitting the job with the job submission GUI. For example, separate queues may be created for small jobs, for medium sized jobs, and for large jobs, where "small" nodes are available to the queue for small jobs, "medium sized" nodes are available to the queue for medium sized jobs, and "large" nodes are available to the queue for large jobs.

Running Test Jobs

If there is a problem connecting to the scheduler when the job submission node is the local node (where the GUI is running) and the PBSPro scheduler is specified, check the following items:

- The cluster is up and running normally.
- The local node is configured as a job submission node for the cluster.
- The user running the GUI has permission to submit jobs to the cluster.
- The environment was configured for submission of jobs to the cluster before starting the product, ansysedt.

["Non-distributed job" on the facing page](#)

["Non-MPI job distributed to multiple hosts" on the facing page](#)

["MPI job distributed to multiple hosts" on the next page](#)

Non-distributed job

If the job fails to complete successfully, check the following items:

- The user submitting the job must be authorized to run jobs on the cluster.
- The pathname of the product must be accessible by the user running the job from the cluster execution node(s) where the job runs.
- The pathname of the project must be accessible by the user running the job from the cluster execution node(s) where the job runs. The directory containing the project file must be writable by the user running the job from the cluster execution node(s) where the job runs, and there must be adequate disk space for the job results.
- The temp directory specified by batchoptions or using the installation default setting on the cluster execution node(s) where the job runs must have adequate disk space for the job temporary files.
- Check for errors or warnings in the **Monitor Job** dialog, and address these issues.

Non-MPI job distributed to multiple hosts

This job will be submitted to run on two hosts with N cores per host, requesting a total of $2 * N$ cores, where N is the smallest number of cores per host, such that the cluster contains at least two hosts with N cores per host.

For this test, the product, `ansyedt`, will run on one of the hosts, and the PBS `pbsdsh` command will be used to start an `ansoftrsm` process on the other host. The `ansoftrsm` process will be used to start any additional processes needed by the job on the other host. In some cases, multiple `ansoftrsm` processes will be started on the other host, but only one such process should be running on a given host at any time.

If the job fails to complete successfully, check the following items, in addition to the items mentioned for the first test:

- Verify that the PBS `pbsdsh` command is able to start the `ansoftrsm` process on the other host.

MPI job distributed to multiple hosts

This job will be submitted to run on two hosts with N cores per host, requesting a total of $2 * N$ cores, where N is the smallest number of cores per host, such that the cluster contains at least two hosts with N cores per host.

Normally, passwordless ssh is used to start MPI tasks on remote hosts. The RemoteSpawnCommand batchoption setting may be used to specify whether to use passwordless ssh or another option. Here are the allowed values of the RemoteSpawnCommand batchoption setting, along with their meanings:

- SSH (the default): Use passwordless ssh for launching remote MPI processes
- RSH: Use the rsh (remote shell) command for launching remote MPI processes
- Scheduler: Use the scheduler for launching remote MPI processes

If the SSH option is specified, then cluster users are required to configure passwordless ssh on the cluster hosts. If the RSH option is specified, then the rsh (remote shell) command must be enabled on the cluster hosts. The "Scheduler" batchoption value is only allowed for jobs running under an LSF or SGE cluster. See ["Message Passing Interface \(MPI\) Installation " on page 5-1](#) for more information.

If the job fails to complete successfully, in addition to the items mentioned for the first and second tests, MPI configuration issues may cause job failures. Check the following MPI configuration issues:

- Verify that the Intel MPI Vendor is specified via a batchoption setting when submitting the job.
- Verify that the RemoteSpawnCommand is specified via a batchoption setting when submitting the job; otherwise the default, SSH, will be used for the job.
- Verify that the specified MPI Vendor is installed and configured correctly for the user submitting the job, as described in ["Message Passing Interface \(MPI\) Installation " on page 5-1](#).

PBS/Torque

Cluster Configuration Guidelines

If the Torque home directory, TORQUEHOME, is different from the default directory, /var/spool/torque, then the environment variable

ANSOFT_TORQUEHOME should be set to the pathname of the Torque home directory. The environment of each ANSYS Electromagnetics job should include this environment variable setting.

Running Test Jobs

If there is a problem connecting to the scheduler when the job submission node is the local node (where the GUI is running) and the PBS/Torque scheduler is specified, check the following items:

- The cluster is up and running normally.
- The local node is configured as a job submission node for the cluster.
- The user running the GUI has permission to submit jobs to the cluster.
- The environment was configured for submission of jobs to the cluster before starting the product, `ansyedt`.

["Non-distributed job" below](#)

["Non-MPI job distributed to multiple hosts" on the next page](#)

["MPI job distributed to multiple hosts" on the next page](#)

Non-distributed job

If the job fails to complete successfully, check the following items:

- The user submitting the job must be authorized to run jobs on the cluster.
- The pathname of the product must be accessible by the user running the job from the cluster execution node(s) where the job runs.
- The pathname of the project must be accessible by the user running the job from the cluster execution node(s) where the job runs. The directory containing the project file must be writable by the user running the job from the cluster execution node(s) where the job runs, and there must be adequate disk space for the job results.
- The temp directory specified by `batchoptions` or using the installation default setting on the cluster execution node(s) where the job runs must have adequate disk space for the job temporary files.
- Check for errors or warnings in the **Monitor Job** dialog, and address these issues.

Non-MPI job distributed to multiple hosts

This job will be submitted to run on two hosts with N cores per host, requesting a total of $2 * N$ cores, where N is the smallest number of cores per host, such that the cluster contains at least two hosts with N cores per host.

For this test, the product, `ansysedt`, will run on one of the hosts, and the PBS `pbsdsh` command will be used to start an `ansoftrsm` process on the other host. The `ansoftrsm` process will be used to start any additional processes needed by the job on the other host. In some cases, multiple `ansoftrsm` processes will be started on the other host, but only one such process should be running on a given host at any time.

If the job fails to complete successfully, check the following items, in addition to the items mentioned for the first test:

- Verify that the PBS `pbsdsh` command is able to start the `ansoftrsm` process on the other host.

MPI job distributed to multiple hosts

This job will be submitted to run on two hosts with N cores per host, requesting a total of $2 * N$ cores, where N is the smallest number of cores per host, such that the cluster contains at least two hosts with N cores per host.

Normally, passwordless `ssh` is used to start MPI tasks on remote hosts. The `RemoteSpawnCommand` batchoption setting may be used to specify whether to use passwordless `ssh` or another option. Here are the allowed values of the `RemoteSpawnCommand` batchoption setting, along with their meanings:

- `SSH` (the default): Use passwordless `ssh` for launching remote MPI processes
- `RSH`: Use the `rsh` (remote shell) command for launching remote MPI processes
- `Scheduler`: Use the scheduler for launching remote MPI processes

If the `SSH` option is specified, then cluster users are required to configure passwordless `ssh` on the cluster hosts. If the `RSH` option is specified, then the `rsh` (remote shell) command must be enabled on the cluster hosts. The "Scheduler" batchoption value is only allowed for jobs running under an LSF or SGE cluster. See ["Message Passing Interface \(MPI\) Installation"](#) on page 5-1 for more information.

If the job fails to complete successfully, in addition to the items mentioned for the first and second tests, MPI configuration issues may cause job failures. Check the following MPI configuration issues:

- Verify that the Intel MPI Vendor is specified via a batchoption setting when submitting the job.
- Verify that the RemoteSpawnCommand is specified via a batch-option setting when submitting the job; otherwise the default, SSH, will be used for the job.
- Verify that the specified MPI Vendor is installed and configured correctly for the user submitting the job, as described in ["Message Passing Interface \(MPI\) Installation "](#) on page 5-1.

4 - Updating Software and Licenses

After installation, you may need to reinstall the software, update a license, or add a software component, such as a library.

Updating ANSYS Electromagnetics Software

To update ANSYS Electromagnetics software:

1. If you have received a CD for the ANSYS Electromagnetics product you wish to install, insert it into the CD-ROM drive, and run install from the top-level directory of the CD-ROM.

If you downloaded the ANSYS Electromagnetics product from the ANSYS Electromagnetics software downloads website (see ["Additional Resources "](#) on page 1-5), unzip the download, and run the install file in the newly created directory.

2. Run install, which is located in the top-level directory where the CD-ROM is mounted (for example, /cdrom/cdrom0/install).

The initial **Installation Shell** screen automatically appears.

3. Click **Install <Product>** for the product you want to reinstall. An **Installation Shell** window appears.

4. Follow the instructions as you would for a regular software installation. Make sure to specify the same directory as the previous installation.

See ["Installing the Software"](#) on page 2-1 for more details.

5. On the **License Information** screen, select the new license file.

Note Select **I have a new license file** to change the license file (if you received a new one from ANSYS Electromagnetics, either evaluation or purchased).

Select **I want to specify a license server** if you need to change the name of the machine where the ANSYS License Manager is installed.

See ["Installing the Software"](#) on page 2-1 for more details on available license options.

6. When you reach the **Summary Information** screen, review the information displayed under **Current Settings**, and then do one of the following:

- If the settings are correct, click **Next** to start copying files.

- If the settings are incorrect, click **Back** as many times as necessary, and change your selections. Verify each screen again before clicking **Next** to move on to the next screen. When all information is correct, return to the **Summary Information** screen, and click **Next**.

The **Setup Status** screen appears, displaying a progress bar, and the update begins.

When the software update is complete, a **Setup Complete** window appears for the product you installed.

7. Click **Finish**.

Updating a Software License

There may be times when you need to update a software license. This may occur for a number of reasons, such as:

- When you change servers, do the following:
 1. Install ANSYS License Manager on the new server machine, and add the new license file.
 2. Run each ANSYS Electromagnetics software installation in **Maintenance mode**, and specify the names of the new servers.
- When you receive an updated purchased license file. You only need to update the license file using the ANSYS License Manager.

If you are changing the name of the license server used by the software, you need to modify the client installation:

1. Click **Start>ANSYS EM Suite x**.
2. Click the desired ANSYS product program folder, and then click **<ANSYS Electromagnetics Product> Maintenance**.

The **Maintenance Mode** screen appears for the desired ANSYS Electromagnetics software.

3. Click **Modify**.
4. Click **Next** in each window.
5. On the **License Information** screen, select the new license file.

Note Select **I have a new license file** to change the license file (if you received a new one from ANSYS Electromagnetics, either evaluation or purchased).

Select **I want to specify a license server** if you need to change the name of the machine where the ANSYS License Manager is installed.

See "Select one of the following three license options." on page 2-4 for more details on available license options.

Specifying License Server Machines via the `ansyslmd.ini` file or Environment Variables.

The Specify the License Server Machine option during installation creates or updates the `ansyslmd.ini` file that is located in the licensing directory. Entries in the `ansyslmd.ini` file tell ANSYS, Inc. products which license server(s) to query to find a valid license. Using this option allows all users at your site to use this setting without having to individually set the **ANSYSLMD_LICENSE_FILE** or **ANSYSLI_SERVERS** environment variables to specify the license server machine(s). It also eliminates the need to have a copy of the license file on every system at your site.

The order that the SERVER lines are listed in the `ansyslmd.ini` file dictates the order in which the license request is granted. To reorder the entries in the list, highlight a server and use the **Move up** or **Move down** buttons.

You are not limited to designating one set of license server machines for your network. You can have multiple single-server or three-server (redundant triad) licensing systems on your network. In this situation, you would have certain licenses connected to a set of server machines (one or three) on the network, and other licenses connected to a different set of server machines (one or three) on the network.

The format of the `ansyslmd.ini` file created by **ANSLIC_ADMIN** follows. Each server's specification entry in this file will typically begin with `ANSYSLI_SERVERS=` and `SERVER=` to specify the Licensing Interconnect and FlexNet port numbers, respectively.

On a single server:

```
ANSYSLI_SERVERS=<ansysliport>@<host>  
SERVER=<flexnetport>@<host>
```

For multiple single servers, each server should have its own `ANSYSLI_SERVERS=` and `SERVER=` lines.

On redundant (triad) servers on Linux platforms. **Note** that Linux uses colons as separators:

```
ANSYSLI_  
SERVERS=<  
  ansysliport>@<host1>:<ansysliport>@<host2>:<ansysliport>@<host3>  
SERVER=<  
  flexnetport>@<host1>:<flexnetport>@<host2>:<flexnetport>@<host3>
```

Windows platforms use semicolons (;) instead of colons:

```
ANSYSLI_  
SERVERS=<  
  ansysliport>@<host1>;<ansysliport>@<host2>;<ansysliport>@<host3>  
SERVER=<  
  flexnetport>@<host1>;<flexnetport>@<host2>;<flexnetport>@<host3>
```

Do not use commas as separators; use colons (:) on Linux machines and semicolons (;) on Windows machines.

The ANSYS Licensing Interconnect default port number is 2325 and the ANSYS FlexNet default port number is 1055. The host is the license server hostname. For example, if the license server name is alpha:

```
ANSYSLI_SERVERS=2325@alpha  
SERVER=1055@alpha
```

The order of the ANSYSLI_SERVERS and the SERVER lines in the ansyslmd.ini file specifies the order in which the requested license will be granted.

You must use the port@host format; you cannot enter a path or a filename in place of the hostname.

Overriding the ansyslmd.ini File – If you want to override the server specification settings in the ansyslmd.ini file, you can do so by setting the **ANSYSLMD_LICENSE_FILE** and **ANSYSLI_SERVERS** environment variables on individual machines. These environment variables are useful if you want to temporarily point to a different license server machine without disrupting the machine's configuration. Use the **ANSYSLI_SERVERS** environment variable to specify the Licensing Interconnect port number. Use the **ANSYSLMD_LICENSE_FILE** to specify the FlexNet port number.

If the **ANSYSLMD_LICENSE_FILE** environment variable is set but the **ANSYSLI_SERVERS** environment variable is not set, the same server machines will be used to specify the Licensing Interconnect but the port number will be replaced by the Licensing Interconnect default port of 2325. When both variables are set, **ANSYSLMD_LICENSE_FILE**

explicitly defines the FlexNet servers while **ANSYSLI_SERVERS** explicitly defines the Licensing Interconnect servers.

If you set the **ANSYSLMD_LICENSE_FILE** or **ANSYSLI_SERVERS** environment variables on a three-server (redundant) system, specify all three systems in the same order as the SERVER lines are listed in the license file. If you specify only the master and it is down, you could see a "License Server Down" or "No License Found" message and the search for a license could fail. Join redundant or multiple single server systems by separating the system names with colons on Linux systems and semicolons on Windows systems.

To have your license server preference known each time you log in, set the environment variables **ANSYSLMD_LICENSE_FILE** and **ANSYSLI_SERVERS** (both Windows and UNIX systems). On UNIX systems, place these environment variables in your login startup file (e.g., .cshrc file).

5 - Message Passing Interface (MPI) Installation

For Linux, the supported Message Passing Interface (MPI) option in ANSYS Electromagnetics Suite is Intel. MPI is required only when using multiple nodes. MPI is never required when solving on a single node.

This chapter provides information about the Linux installations of MPI:

Linux

Follow these steps for a Linux installation:

1. Install ANSYS Electromagnetics Suite.
 - The recommended installation is in a single location on a networked drive that is available to all machines.
 - If you choose to install on each machine individually, make sure to use the same location on each machine (for example, /share/install/).
2. Install and configure the Remote Simulation Manager (RSM), if needed.
 - The local machine is the machine on which you are running the desktop. If you set up your simulations so that the first machine in the distributed machine list is the local machine, you do not need to install or run the RSM.
 - If you are running through a scheduler, the scheduler will launch the simulation, and RSM is not needed.
 - If the first machine in the distributed machine list is not the local machine and you are not running through a scheduler, then RSM must be used. RSM needs to be installed and running on the remote machine that will host the simulation, and the ANSYS Electromagnetics Suite installation must be registered with RSM on that machine (see the Remote Analysis topic in the online help for details).
3. Configure MPI.

Note: It is not necessary to install MPI on any machine: Intel MPI is provided in the ANSYS Electromagnetics Suite installation.

Set up passwordless ssh to enable communication between machines that will be used in the simulation. At a minimum, this must be

set up from the simulation host machine (the first machine in the distributed machine list) to all of the remaining machines. It is also possible to use RSH, but this is not recommended.

4. Define the temp directory on all machines, assigning Full Control to the designated user.
5. If a firewall is installed, it should be off or configured as specified in the Firewall Configuration topic in the online help.
6. Select Intel to use in the **MPI Vendor** box on the **Options** tab under **Tools > Options > HPC and Analysis Options**. Also select whether to use SSH or RSH for the Remote Spawn Command (SSH is recommended). If the product is running as an interactive scheduler job on an LSF or SGE cluster, and the MPI Vendor is Intel, then there is a third option for the Remote Spawn Command, Scheduler. Selecting the Scheduler option means that the LSF or SGE scheduler will be used to launch remote commands, rather than RSH or SSH.
7. If you choose a remote machine to host the simulation, then RSM needs to be installed and running on that remote machine. There are two RSM settings that can affect the simulation. The first is the user who is running RSM. The second setting is in the Desktop in **Tools > Options > General Options** on the **Remote Analysis** section. Choose to send the analysis request as a **Service User** or a **Specified User**. Here is additional information about these settings:

RSM User	Remote Simulation Options	Simulation runs on remote machine as	Notes
root	Service User	The user who is running the ANSYS Electromagnetics desktop	t RSM will try to run as the user, but may fail based on local privileges
root	Specified User = user_A	user_A	Runs as user_A
User_A	Service User	user_A	Runs as user_A
User_A	Specified User = user_B	Will fail due to privileges	Not recommended

General Notes

If using a scheduler (for example, SGE or LSF) with a network installation, the installation, setup and execution are the same as for previous releases. If there are any questions, refer to the online help for more information.

If you are interested in controlling which MPI Vendor to use with Batchoptions, the options for setting this are:

```
'HFSS/MPIVendor'='Intel'  
'HFSS-IE/MPIVendor'='Intel'  
'Q3D Extractor/MPIVendor'='Intel'
```

Similarly, the options for setting the Remote Spawn Command with Batchoptions are:

```
'HFSS/RemoteSpawnCommand'='RSH'  
'HFSS-IE/RemoteSpawnCommand'='RSH'  
'Q3D Extractor/RemoteSpawnCommand'='RSH'
```

or

```
'HFSS/RemoteSpawnCommand'='SSH'  
'HFSS-IE/RemoteSpawnCommand'='SSH'  
'Q3D Extractor/RemoteSpawnCommand'='SSH'
```

or

```
'HFSS/RemoteSpawnCommand'='Scheduler'  
'HFSS-IE/RemoteSpawnCommand'='Scheduler'  
'Q3D Extractor/RemoteSpawnCommand'='Scheduler'
```

The Scheduler option can only be used for jobs running under an LSF or SGE scheduler and only if the MPI Vendor is Intel.

Setting up passwordless ssh

1. First verify that passwordless ssh is not already configured by running the following command on one node to access another:

```
ssh <remote node> echo hi
```

If you receive a response of hi, then passwordless ssh is already configured between this node and the remote node. Otherwise, you will need to configure passwordless ssh. If you have already configured passwordless ssh and it is not working, you should run the following command to obtain information for diagnosing the problem:

```
ssh -v -v -v <remote node> echo hi
```

Configuring passwordless ssh for the first time

1. Generate ssh key by entering the following:

```
ssh-keygen -t rsa
```

2. Save the file to the default location (\$HOME/.ssh/id_rsa).
3. Press **Enter** for no passphrase.
4. Press **Enter** again to confirm no passphrase.
5. Ensure that permissions for the following directories are correct:

- `chmod go-w ${HOME}`
- `chmod 700 ${HOME}/.ssh`
- `chmod 600 ${HOME}/.ssh/id_rsa`

6. Append your personal public key file to your personal authorized_keys list

```
cat ${HOME}/.ssh/id_rsa.pub >> ${HOME}/.ssh/authorized_keys  
(if required) cat ${HOME}/.ssh/id_rsa.pub >> ${HOME}/.ssh/authorized_keys2
```

7. Test that local passwordless ssh is working.

```
ssh <local node> echo hi
```

If you are prompted for a password, there is probably an access issue. Run `ssh -v -v -v <local node>` to obtain information for diagnosing the problem.

8. The next steps depend on whether you have a user-specific HOME directory that is shared across machines.

If so, in most cases it is sufficient to test that remote passwordless ssh is working to one node:

```
ssh <remote node> echo hi
```

- If the command returns hi, you probably have a shared HOME directory, and the configuration is complete. If it fails, then you need to diagnose the failure by running `ssh -v -v -v <remote node>` to obtain information for diagnosing the problem.

- If you know that your HOME directory is not shared, you need to copy the contents of your local id_rsa.pub file to the authorized_keys file on the remote machine:

```
cat ${HOME}/.ssh/id_rsa.pub | ssh <remote node> 'cat >>  
${HOME}/.ssh/authorized_keys'; ssh <remote node> 'chmod 600  
${HOME}/.ssh/authorized_keys'
```

Run `ssh <remote node> echo hi` to verify that you can log into the remote node from the local node.

Troubleshooting Section

- **I am installing and have a dual monitor setup, but the installation wizard isn't displaying correctly.**

If there is a conflict with the RedHat gnome-shell window manager during installation, replace gnome-shell via the terminal emulator command shell with "metacity -replace &" and continue with the installation. After the installation is complete, change "metacity - replace &" back to "gnome-shell - replace &".

- **I am getting an error message that mentions AnsoftRSMService. Is there a problem with my installation?**

[error] Unable to locate or start COM engine on 'your_machine_name_here' : Unable to reach AnsoftRSMService. Check if the service is running and if the firewall allows communication...

Do not immediately assume there is a problem with RSM, especially if your_machine_name_here is your local machine. The message is stating that it cannot start the Com engine. This can be the result of a misspelled machine name or an issue with the firewall. Check these first.

If the host machine is a remote machine (your_machine_name_here is a remote computer), there might be an issue with the RSM service on the remote computer. Check the RSM installation and registration on that machine.

- **Is there an easy way to test the computers in my Distributed Machine Configurations?**

1. Select **HPC and Analysis Options**.
2. Click **Edit**, and select one or more machine names in the list on the **Machines for Distributed Analysis** dialog.
3. Click the **Test Machines** button. A set of tests will be run from the current machine to the selected machines that includes pinging the computer and checking the RSM availability.

- **When I attempt to simulate, I get the following error message:**

[error] The attempted launch of solvers via MPI failed while connecting to communication pipes. The probable cause is failure to install the vendor MPI on one or more machines or password authentication failure for MPI during the launch attempt.

Followed by

Unable to create child process: mpirun.

Possible causes for the messages are:

- MPI is not installed on one or more machines.
- The wrong version of MPI is installed on one or more machines.
- The password registration is incorrect on one or more machines.
- ANSYS Electromagnetics Suite is not installed in the same location on all machines.
- There is an issue with the firewall.
- **When I attempt to simulate, I get the following error message:**

[error] Distributing matrix solutions for domains, process mpirun error: Simulation directory creation failed for one or more MPI solvers. Check access privileges and disk space for the following:...

This message means you do not have the correct privilege for the temp directory on one or more of the remote machines. In many cases the software uses the installation temp directory that is defined during the ANSYS Electromagnetics Suite installation on the given machine. It is recommended that this directory be set for Everyone to have full Read/Write privileges on all machines.

- **My setup was working last week, but today I am getting a password authentication failure error.**

This error may occur if anything changed in your network configuration. The most common cause is your login password was changed as part of a periodic update for network security. If the passwords were registered for MPI, they need to be updated on all machines every time the login passwords are changed.

- **When I attempt to simulate with a solver utilizing MPI on a local network, the job fails during MPI initialization.**

Error messages include but are not limited to: *Unable to create child process:* and *Could not start memory inquiry.*

This often happens if there is an issue with the host key based authentication. When this happens check the terminal (console) where you launched the software for a warning:

The authenticity of host '<hostname> (<ip address>)' can't be established.

Manually initiate an ssh connection from your current machine to this host using the exact hostname as listed in the warning and answer yes at the prompt to save the host key.

A - Silent Installation

When you use the ANSYS Electromagnetics silent installation feature, the software is installed on your computer with no graphical interface appearing during the installation process. InstallShield suppresses all dialog boxes, including error messages and informational message boxes.

To install ANSYS Electromagnetics software on a Linux machine using a silent installation:

- Obtain a license file for the software you are installing.
- Install the ANSYS Electromagnetics software from a command prompt, as described in this chapter.

Running the Silent Installation

To install ANSYS Electromagnetics software on a Unix machine using a silent installation, you need to record a response file, and then use that file to install the software on other similar computers.

To record a response file:

1. Go to the command prompt.
2. At the command line, use the `cd` command to change to the product subdirectory (which contains the file **setup.exe**).
3. Do the following to run Setup in Initial mode to record the response file:

```
setup.exe -options-record "<response file name>"
```

See ["Response File " on the next page](#) for an example of a typical response file.

To install the software on other similar computers (i.e., replay the response file):

1. Copy the response file so that you can use it on the other computers.
2. On other computers where you want to install the software, go to the command prompt, and change to the directory containing **setup.exe**.
3. Type the following, and press ENTER:

```
setup.exe -options "<response file name>" -silent
```

Note Although each value is commented, there is no way to display feedback if the information is invalid. Therefore, be careful

with the syntax and spacing of the commands.

In silent mode, no confirmation windows of any kind are displayed. In this mode, InstallShield suppresses all dialog boxes, including error messages and informational message boxes.

Any line in the response file beginning with the "#" symbol is a comment. Blank lines are ignored.

You can change values only to the right of the "=" sign.

Note If a failure occurs in silent mode, Setup aborts just as if you clicked the **Abort** button in a wizard dialog box during a normal installation. In this case, an error is written to the log file.

The installation typically fails for one of the following two reasons:

- An invalid path was specified on the command line.
- There were differences in the dialog box sequences between the recorded response file and the current silent installation--for example, an additional message appearing during the playback due to differences in the two computer systems. Since most computers would normally be configured identically by a system administrator, this problem should not occur often.

Response File

The response file contains values you specified when recording it. It can be used to configure the silent installation with the options you prefer.

Following is an example of a typical response file:

```
#####-
#
# InstallShield Options File
#
# Wizard name: Install
# Wizard source: setup.jar
# Created on: Wed Sep 17 14:32:59 EDT 2003
# Created by: InstallShield Options File Generator
```

```
#
# This file contains values that were specified during a recent execution of
# Install. It can be used to configure Install with the options specified below
# when the wizard is run with the "-options" command line option.
Read each
# setting's documentation for information on how to change its value.
#
# A common use of an options file is to run the wizard in silent mode. This lets
# the options file author specify wizard settings without having to run the
# wizard in graphical or console mode. To use this options file for silent mode
# execution, use the following command line arguments when running the wizard:
#
# -options "settings.txt" -silent
#
#####
#####
#
# License Agreement State
#
# The initial state of the License Agreement panel. Legal values are:
#
# 0 - Nothing will be selected
```

```
# 1 - "I &accept the terms of the license agreement." will be selected
# 2 - "I &do not accept the terms of the license agreement." will be selected
#
# For example, to configure the panel to initially display "I &do not accept the
# terms of the license agreement.", use
#
# -W Agree.selection=2
#
-W Agree.selection=1
#####
#
# User Input Field - groupID
#
#
-W GroupPanel.groupID=""
#####
#
# Install Location
#
# The install location of the product. Specify a valid directory into which the
# product should be installed. If the directory contains spaces, enclose it in
# double-quotes. For example, to install the product to C:\Program Files\My
# Product, use
#
```



```
# -P installLocation="C:\Program Files\My Product"  
#  
-P installLocation="/usr/ansoft"  
  
...  
  
...  
  
<remainder of file deleted>
```


B - License hostid

Software licenses for each module of ANSYS products you purchased are contained in one or more license files provided by ANSYS. Each license file requires one or more license hostids, which are physical devices that contain a specific identification. License hostids are always required for the license server, and they may also be required for each machine running ANSYS software.

Each supported operating system supports certain hostids:

Operating System	Permitted hostids
Windows	<ul style="list-style-type: none"> • The ethernet address of a network interface card (MAC address). • Flexera FLEXid USB port hardware key. • Disk serial number. <p>To obtain the disk serial number, open a command prompt.</p> <p>Type vol <drive letter>:, for example vol C:</p>
Linux	<ul style="list-style-type: none"> • The ethernet address of a network interface card (MAC address). • Flexera FLEXid USB port hardware key.
Solaris	Machine hostid

This utility will provide the disk id by default:

<http://www.ansys.com/Support/Licensing/Capture+License+Server+Information>

Network Interface Card

A network interface card is a physical device used to connect to a network (also known as an ethernet card). Each network interface card includes a unique ID that can be used as a license hostid.

You can obtain an ethernet address to use as a valid license hostid. The method for doing so differs depending on whether you are running Windows or Red Hat Enterprise Linux.

Windows

To obtain an ethernet address as a valid license hostid when running Microsoft Windows:

1. If you are using a laptop, disconnect the laptop from the docking station if docked.
2. Go to the command prompt.
3. Type `ipconfig /all`, and press **Enter**.
A listing of network interface cards is shown.
4. For a valid network interface card, look for the **Physical Address** line. The ethernet address for that device is the value of the physical address, without dashes.

Note Other types of adapters, including Point-to-Point Protocol (PPTP) adapters and the Microsoft Loopback Adapter, may be listed here. These are not valid for license hostids.

Linux

To obtain an ethernet address as a valid license hostid when running Red Hat Enterprise Linux:

1. If you are using a laptop, disconnect the laptop from the docking station if docked.
2. Start a terminal session.
3. Type `ifconfig` (normally located in the `/sbin` directory) to see a listing of all active network devices.
4. Devices with an ethernet address show a **Link encap** value of **Ethernet**. The ethernet address for each device is the value in the **HWaddr** field, without colons.

Hardware Key

Installing a Hardware Key on a USB Port

Only a single device can be installed on each USB port.



Flexera FLEXid hardware key for a USB Port

To install the hardware key on a USB port:

1. Attach the hardware key to the USB port.
2. If a prompt appears asking you to install or specify a driver, choose **Cancel**.

See the ANSYS License Manager documentation for instructions on using USB dongles.

Machine hostid

A machine hostid is present on most Unix machines, including Sun machines. To determine the hostid for a supported Unix machine, enter the following command:

Sun Solaris hostid

C - Setting Up a Printer on Linux

To print from ANSYS Electromagnetics software on Linux, you must first configure a printer. To do this, launch the MainWin control panel.

1. Run `./mwcontrol &` in the installation subdirectory.
The **MainWin Control Panel** appears.
2. Double-click on the "Printers" icon to start the **MainWin Printers** panel.
3. Then double-click on the "Add New Printer" icon.
This starts the **Add Printer Wizard**.
4. Select the **Let the wizard search for printers** radio button and click **Next**.
5. In the **Identify your Unix Printer** dialog do one of the following:
 - If your printer is listed, select it.
 - If your printer is not listed, you will need to cancel and get someone with root permission to setup a printer queue on your machine (and then you will need to come back and run this wizard later).

Note On Solaris you set up a new print queue by running "lpadmin" (as root). On Red Hat Linux, you can run 'System Settings/Printing' to launch printconf-gui (as root).
6. Click **Next**.
The **Print Command** dialog appears.
7. Change the Print Command only if instructed to do so by your user administrator.
8. Click **Next**.
The Choose PPD File dialog appears.
9. Select your printer manufacturer and model from the list or use the **Choose File** button to browse to a PPD file provided by your printer manufacturer. Click **Next**.
10. The **Printer Name** dialog appears. Enter a **Name** to identify the printer. Click **Next**.
11. Choose whether this printer should be the default and click **Next**.
12. Choose whether you would like to print a test page and click **Next**.

13. In the **Finish Adding New Printer** dialog, verify the printer setup information. If the information is incorrect, use the **Back** button to return to the appropriate dialog and correct the entry. If the information is correct, click **Finish** to complete the setup of your printer.

With a print queue setup, and the printer added, you should then see the printer when running ANSYS Electromagnetics products.