

High Frequency Electromagnetic Software

# SONNET<sup>®</sup> 12 Suites



*emCluster*  
Computing

*This page intentionally left blank.*

# **EMCLUSTER COMPUTING**

*Published: April 2009*

## **Release 12**

Sonnet Software, Inc.  
100 Elwood Davis Road  
North Syracuse, NY 13212  
Phone: (315) 453-3096  
Fax: (315) 451-1694

Technical Support: [support@sonnetsoftware.com](mailto:support@sonnetsoftware.com)

Sales Information: [sales@sonnetsoftware.com](mailto:sales@sonnetsoftware.com)

[www.sonnetsoftware.com](http://www.sonnetsoftware.com)

© Copyright 1989,1991,1993, 1995-2009 Sonnet Software, Inc. All Rights Reserved

Registration numbers: TX 2-723-907, TX 2-760-739

---

## Copyright Notice

---

Reproduction of this document in whole or in part, without the prior express written authorization of Sonnet Software, Inc. is prohibited. Documentation and all authorized copies of documentation must remain solely in the possession of the customer at all times, and must remain at the software designated site. The customer shall not, under any circumstances, provide the documentation to any third party without prior written approval from Sonnet Software, Inc. This publication is subject to change at any time and without notice. Any suggestions for improvements in this publication or in the software it describes are welcome.

---

## Trademarks

---

The program names *xgeom*, *emstatus*, *emvu*, *patvu*, *dxfgéo*, *ebridge*, *emgraph*, *gds*, *emserver*, *emclient*, *sonntcds*, and *sonntawr* (lower case bold italics), Co-calibrated, Lite, LitePlus, Level2 Basic, Level2 Silver, and Level3 Gold are trademarks of Sonnet Software, Inc. Sonnet®, *em*®, and *emCluster*® are registered trademarks of Sonnet Software, Inc. UNIX is a trademark of Unix Systems Labs. Windows NT, Windows2000, Windows ME, Windows XP and Windows Vista are trademarks of Microsoft, Inc. AutoCAD and Drawing Interchange file (DXF) are trademarks of Auto Desk, Inc. SPARCsystem Open Windows, SUN, SUN-4, SunOS, Solaris, SunView, and SPARCstation are trademarks of Sun Microsystems, Inc. HP, HP-UX, Hewlett-Packard are trademarks of Hewlett-Packard Company. ADS, Series IV, Touchstone, and Libra are trademarks of Agilent Technologies. Cadence and Virtuoso are trademarks of Cadence Design Systems, Inc. AWR and Microwave Office are registered trademarks and EM Socket is a trademark of Applied Wave Research, Inc. GDSII is a trademark of Calma Company. Acreso, FLEXlm, and FLEXnet are registered trademarks of Acreso Software. OSF/Motif is a trademark of the Open Software Foundation. IBM is a registered trademark of International Business Machines Corporation. Linux is a registered trademark of Linus Torvalds. Redhat is a registered trademark of Red Hat, Inc. SuSE is a trademark of Novell, Inc. Adobe® and Acrobat® are registered trademarks of Adobe, Inc. AWR and Microwave Office are trademarks of Applied Wave Research, Inc. Platform is a trademark and LSF® is a registered trademark of Platform Computing.

---



---

## TABLE OF CONTENTS

---



---

	TABLE OF CONTENTS . . . . .	5
1	SETTING UP SONNET CLUSTER . . . . .	7
	Introduction . . . . .	7
	Sonnet Cluster Overview . . . . .	9
	Sonnet Cluster Requirements . . . . .	11
	Floating Licenses . . . . .	11
	Cluster Node Networking . . . . .	12
	Firewalls. . . . .	12
	Sonnet Cluster Setup . . . . .	13
	Install Sonnet . . . . .	13
	Configuring the Master Computer . . . . .	13
	Starting <i>sonnetcluster</i> on the master computer. . . . .	17
	Adding a Slave to the Cluster . . . . .	18
	Configuring a Slave (Done on Each Slave) . . . . .	22
	Setting up a Client Computer . . . . .	30
	Testing the Sonnet Cluster . . . . .	37
	Configuring Slaves Using a Common Installation. . . . .	37
	Cluster Slave Versus Remote <i>em</i> Server . . . . .	40
2	ANALYZING USING SONNET CLUSTER . . . . .	41
	Running an Analysis on your Sonnet Cluster . . . . .	41
	An Analysis Example . . . . .	43
	Frequency Splitting . . . . .	48
	ABS Processing . . . . .	49
	Automatic Calculation of Discrete Data Points . . . . .	49
	User Defined Discrete Data Points . . . . .	50
	Overriding Cluster Options. . . . .	51
	Controlling Cluster Programs from a Command Line . . . . .	54
3	SETTING UP LSF CLUSTER . . . . .	57
	Requirements . . . . .	60
	LSF Setup . . . . .	60
	Sonnet Setup . . . . .	60

## Setting Up *emCluster*

---

Defining a Resource to Handle Sonnet Licensing . . . . .	68
Windows Firewall . . . . .	69
Troubleshooting LSF Cluster . . . . .	69
Testing the LSF Cluster . . . . .	69
Testing Sonnet from the Command Line . . . . .	70
Testing the <i>emclusterlsf</i> Script File . . . . .	71
<b>4 ANALYZING USING LSF CLUSTER . . . . .</b>	<b>73</b>
Running <i>emCluster</i> . . . . .	73
Frequency Splitting . . . . .	73
ABS Processing . . . . .	74
Automatic Calculation of Discrete Data Points . . . . .	75
Using Queues . . . . .	76
Limitations . . . . .	77
Running an Analysis on your LSF Cluster . . . . .	77
Overriding Cluster Computing Options . . . . .	78
Cleaning the <i>emCluster</i> Shared Directory . . . . .	81
<b>APPENDIX I THE <i>LSF Cluster</i> INITIALIZATION FILE . . . . .</b>	<b>83</b>
<b>APPENDIX II THE <i>Sonnet Cluster</i> INITIALIZATION FILE . . . . .</b>	<b>93</b>

# Chapter 1                      Setting up Sonnet Cluster

---

## Introduction

Sonnet's *emCluster* using Sonnet networking, also referred to as Sonnet Cluster, provides you with the ability to split your analysis project into multiple jobs which may then be processed in parallel on multiple computers to greatly reduce your processing time. This feature allows you to make the most efficient use of your processing resources and solver licenses, producing results in a fraction of the time it would take to run an analysis on a single computer.

This document assumes that you are familiar with using Sonnet. If you have not used Sonnet previously, we recommend going through the tutorials in the **Getting Started** manual before attempting to use Sonnet Cluster.

Sonnet Cluster provides a cluster computing capability for Sonnet that is based on the existing TCP/IP socket communication system already used in Sonnet's Remote EM Processing capability. Sonnet Cluster has the following features:

- The cluster computing queue may be shared by all users on your network, ensuring that your solver licenses are utilized as efficiently as possible.
- The status of each *emCluster* processing slave may be monitored at any time by any network user.
- All computing and hard disk utilization is local to each computing node; there are no shared file systems required for the *emCluster* computing nodes.

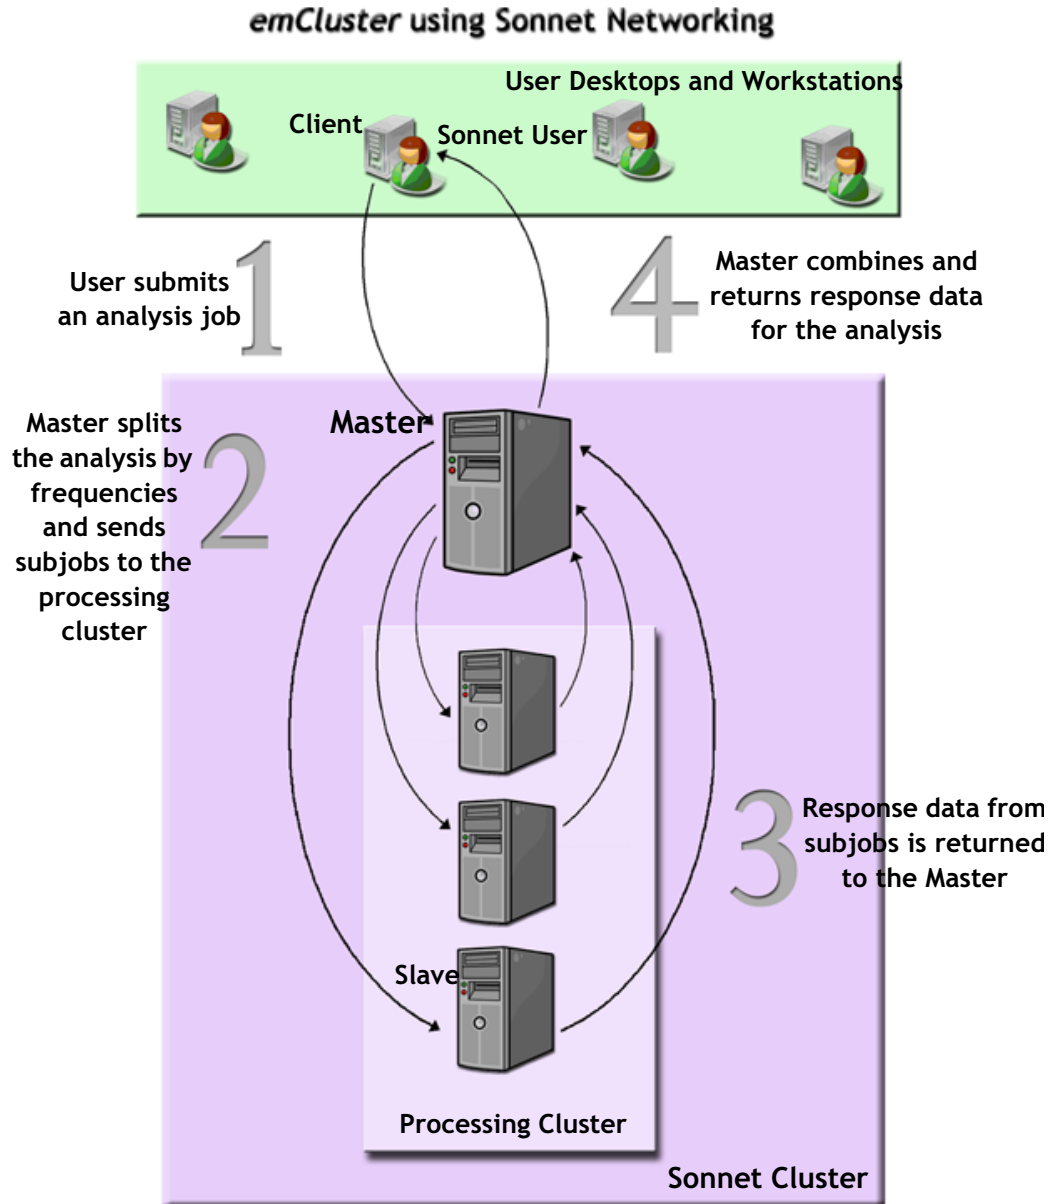
## Setting Up *emCluster*

---

- Minimal network bandwidth is required.
- Sonnet EM analysis jobs may be submitted either for parallel frequency point computation, or to compute entirely on one of the cluster processing nodes.

## Sonnet Cluster Overview

The Sonnet Cluster is made up of a master, clients and slaves. There is only one master computer, but you may have any number of client computers and slave computers. The setup of a Sonnet cluster is pictured below.



The different parts of the Sonnet cluster are defined below.

## Setting Up *emCluster*

---

- Client:**
- The computer from which a user submits an analysis job using Sonnet's analysis monitor (*emstatus*).
  - Can be a Windows, Linux or UNIX computer. Windows Vista is supported for clients.
  - Runs the program *emclient* to communicate with the master computer.
  - Receives and displays the analysis results from the master computer as each sub-frequency is completed.

- Master:**
- The computer that splits the analysis from the client into subjobs. These subjobs are sent to slaves for processing.
  - Must be a Linux or Windows XP system. Windows Vista is not supported at this time.
  - Runs the program *sonnetcluster* to communicate with slaves and clients.
  - Combines the analysis results from the slaves and returns those results to the user on the client from which the job was submitted.
  - Recommended that this computer be dedicated to the Sonnet Cluster.

- Slave**
- Performs electromagnetic analyses on subjobs sent from the master and returns the analysis results to the master.
  - Must be a Linux or Windows XP system. Windows Vista is not supported at this time.
  - May be 32-bit or 64-bit hardware and OS.
  - Runs the program *emserver* to communicate with the master and the program *em* to analyze the subjobs sent from the master.
  - Recommended that this computer be dedicated to the Sonnet Cluster.
  - Recommended that this computer be high performance hardware.

A user submits a job from a client which is sent to the master computer to be split into subjobs. Each subjob analyzes the project using a subset of the analysis frequencies; the default is one analysis frequency per subjob. The master computer sends each subjob to one of the slaves in the processing cluster. Each subjob is analyzed on a slave at only one of the analysis frequencies (default mode). If there are more subjobs than slaves in the cluster, multiple subjobs are submitted to the same slave. Once each slave finishes analyzing its subjob, the results are sent back to the master. The master combines all the data from all the subjobs run on the slaves and sends the response data back to the user on the client from which the job was submitted.

The default is to create one job per frequency, but you can override this for a given analysis by using options in the analysis monitor. For details on assigning subjobs to the processing cluster, see “Automatic Calculation of Discrete Data Points” on page 49.

Any computer capable of running Sonnet can be used as a client. Slaves or masters must use either a Linux operating system supported by Sonnet or a Windows XP operating system. Multi-core or multi-CPU computers may be used, with each core or processor being used as a server. We recommend that the master computer be a dedicated core which does not also serve as a client or slave in the Sonnet Cluster. The same core may be used as both a slave and a client in the Sonnet cluster although typically slaves and clients should be on different cores.

## Sonnet Cluster Requirements

In order to use Sonnet Cluster, you must have the following:

- A *emCluster* license. The Sonnet Cluster option is available only with the Sonnet Professional suite and must be purchased separately. This license will be used by the master computer. Please see your system administrator if you are unsure of the availability of this program.
- A Sonnet *em* license available for each slave in the network. Note that if you use node-locked licenses, then you will need one license for each slave in your network. If you have a floating license, then each slave must be able to simultaneously check out an *em* license.
- Each client will need access to an *xgeom* license in order to edit and submit jobs to the network.

Sonnet Cluster is available on all Linux operating systems supported by this release and on Windows XP. Windows Vista computers may be used as clients, but not as master or slaves in the cluster. UNIX computers may be used as clients, but not as master or slaves in the cluster. For information about Sonnet requirements, please refer to the appropriate installation manual or go online:

<http://www.sonnetsoftware.com/requirements>

## Floating Licenses

We do not recommend using a floating license for both general network use and for cluster slave use. If a non-cluster computer checks out an *em* license, and there are not enough *em* licenses left for all the slaves in the cluster, one or more of the


## Setting Up *emCluster*

---

slaves will wait until a license is available in order to analyze its subjob. Therefore, the analysis job submitted to the cluster will be delayed for the length of time that the non-cluster computer uses the *em* license.

### Cluster Node Networking

We recommend that the Master and slave nodes use a common local network to keep cluster message traffic out of the rest of your computer network. Network traffic between Master and slave nodes may occasionally be rather high (depending on the number of slaves and the analysis job size). Traffic between clients and the Master is comparatively much lower.



### Firewalls

The master and clients, and master and slaves need to communicate with each other in order for Sonnet cluster to function correctly. When setting up the computers in the cluster, you specify what ports are used on each type of computer to communicate with other computers in the cluster. If you have software or hardware firewalls in your network, you must identify the ports being used for the cluster to the firewall so that the computers in the cluster may communicate. If you have set up the Master and slaves in a local network as suggested above, we also recommend disabling any firewalls between the master and slave computers. Please see your system administrator for more information on firewalls on your computer and network.

## Sonnet Cluster Setup

The table below shows a summary of the actions that need to be completed to set up a Sonnet Cluster. You may use this as checklist to ensure that you have completed all the necessary steps.

### Overview of Sonnet Cluster Setup

Task	Directions
Install Sonnet on each slave, client and the master in your cluster.	See page 13
Configure the master Computer	See page 13
Adding Slaves to the Cluster	See page 18
Configuring a Slave (This must be done for each slave in the cluster.)	See page 22
Setting Up a Client computer (This must be done for each client in the Sonnet Cluster.)	See page 30

The next sections detail how to accomplish the tasks in the table above.

### Install Sonnet

#### 1 Install Sonnet on the master computer, all slave computers and all client computers.

Sonnet should be installed on all the computers in your network that will be used in the Sonnet Cluster including setting up licensing to enable all the software. Refer to the appropriate installation manual for the operating system on each computer. Once Sonnet has been installed and verified, continue on the next step. Note that all client computers and slave computers must have network access to the master.

### Configuring the Master Computer

The first step in setting up your Sonnet Cluster is defining and configuring your master computer. The master computer runs the program *sonnetcluster*. Client computers in the Sonnet cluster submit analysis jobs to the master computer and

## Setting Up *emCluster*

---

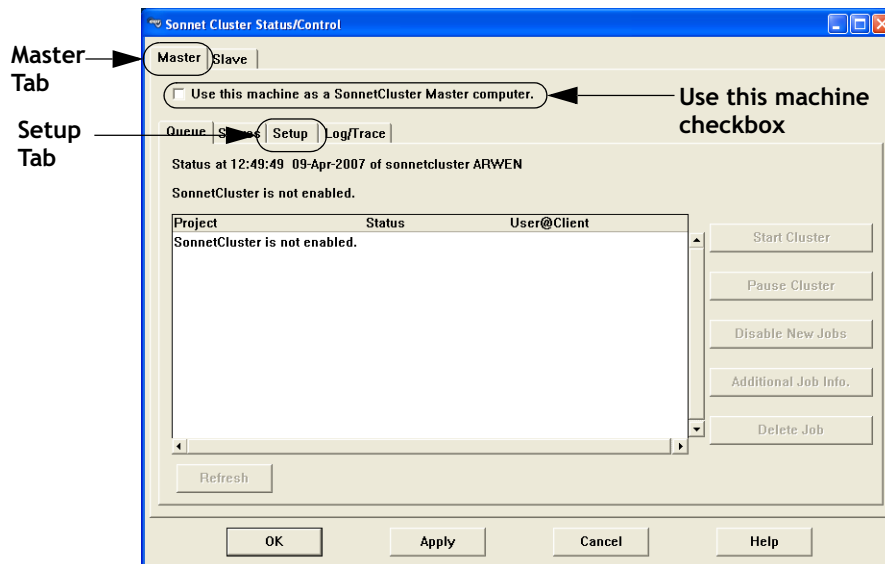
receive analysis results back from the master computer. Slave computers receive subjobs from the master computer and send back analysis results. The master computer combines these results before returning them to the submitting client.

**2 Open the Sonnet task bar on the master computer.**

For instructions on opening the Sonnet task bar, please refer to “Invoking Sonnet” on page 30 of the **Getting Started** manual.

**3 Select *Admin* ⇒ *SonnetCluster Status/Control* from the main menu of the Sonnet task bar.**

The Sonnet Cluster Status/Control window appears on your display as shown below.



**4 If it is not already selected, click on the Master tab in the window.**

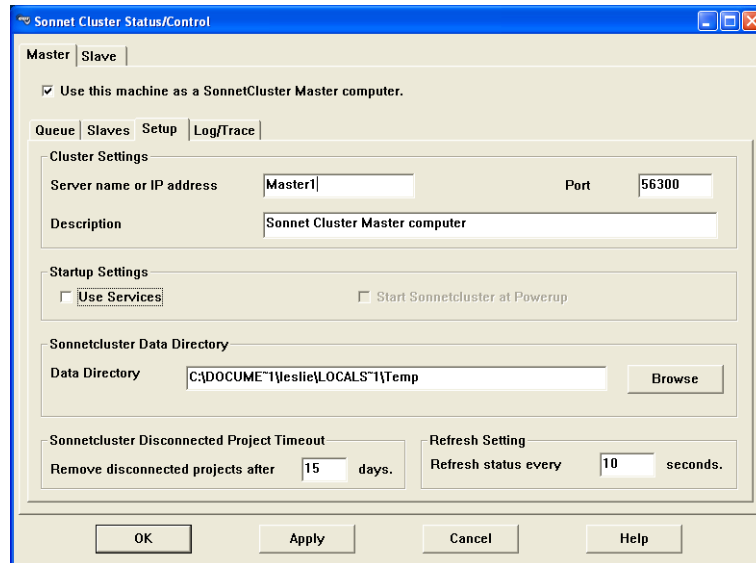
**5 Click on the “Use this machine as a SonnetCluster Master computer” check box.**

Selecting this checkbox designates this computer as the master computer in the Sonnet network.

### Cluster Settings

- 6 Click on the Setup tab in the SonnetCluster Status/Control window.

The appearance of the window below is updated.



- 7 Enter the Server name or IP address of this computer in the Server name or IP address text entry box in the Cluster Settings section of the window.

If you enter a server name, the name must be recognized by all slaves and clients in the Sonnet Cluster. This name is the same as the computer would use in a domain or on a Windows network.

- 8 If you do not wish to use the default port number of 56300, enter the desired port number in the Port text entry box.

This is the port on the Master used by the slave and client computers to communicate with the master computer. This is the same port number which should be entered when you add the cluster as a server to your client in Step 9 on page 34. Keep a record of this port number to use when you later set up your slaves.

## Setting Up *emCluster*

---

- 9 If you wish, you may enter a brief description of the master computer in the Description text entry box.

This description appears when the Sonnet Cluster identifies itself on the network.

---

**NOTE:**

The default values used for the data directory, project timeout, and refresh status settings do not usually need to be changed. For more information on these settings, please see online help for this dialog box.

---

### Starting *sonnetcluster* at Powerup for Windows

If you wish to have the program *sonnetcluster* start automatically at the startup of the master computer, you must install *sonnetcluster* as a service. You will need administrator privileges for this procedure.

- 10 To install *sonnetcluster* as a service, select the Use Services checkbox in the Startup Settings section of the window.

This enables the Start Sonnetcluster at Powerup checkbox.

- 11 Select the Start Sonnetcluster at Powerup checkbox.

Whenever the master computer is restarted, the program *sonnetcluster* will automatically be started. Continue below at Step 14.

### Starting *sonnetcluster* at Powerup for Linux

To change your system startup files for a Linux system, you must run the program *soninstallservice*.

---

**NOTE:**

The *soninstallservice* command will require root permissions to properly install the desired services. Failing to login as root before executing this command may result in unsuccessful initialization of the Sonnet Services.

---

For instructions on how to use this program,

- 12 Open a terminal on the master computer.

- 13 Type the following:

```
cd <Sonnet Directory>/bin
soninstallservice -h
```

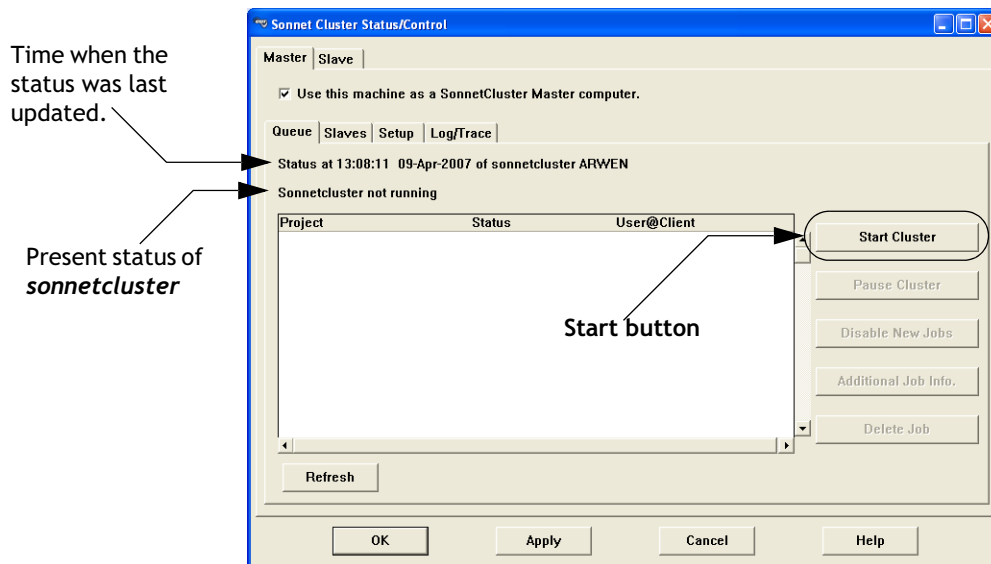
where <Sonnet Directory> is the directory in which the Sonnet Software was installed. Instructions for using *soninstallservice* appear on your display; follow these to install *sonnetcluster* as a service which starts at powerup.

### Starting *sonnetcluster* on the master computer.

You need to start the *sonnetcluster* program in order to test the connections as you set up the slaves and clients.


**14 Click on the Queue tab of the Sonnet Cluster Status/Control window.**

The appearance of the window changes. The present status of *sonnetcluster* is displayed just below the tab as shown below.



**15 Click on the Start Cluster button.**

Since you have made changes to your master computer setup, a query appears on your display.



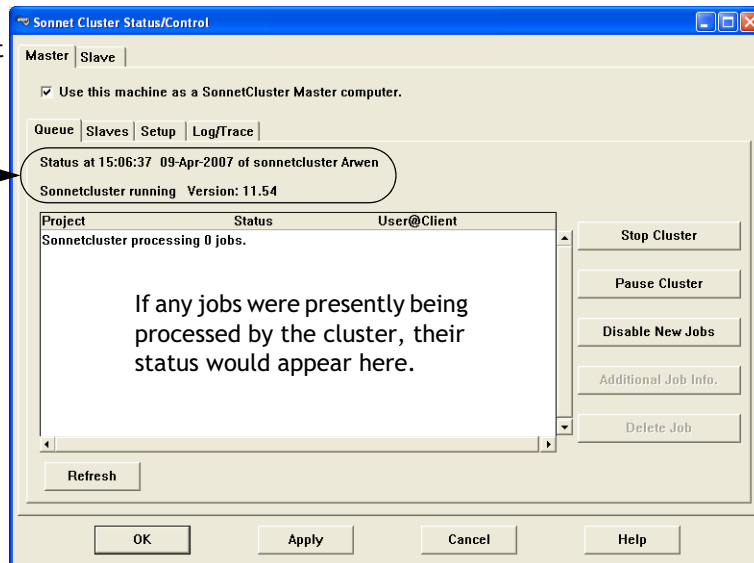
## Setting Up *emCluster*

---

### 16 Click on the Yes button in the query box.

The Status/Control window is updated to show that *sonnetcluster* is running as shown below. On Windows XP, an icon appears in your system tray to indicate that *sonnetcluster* is running on the master computer.

Status shows that *sonnetcluster* is up and running.



### TIP

You may also start *sonnetcluster* by using a command line. Be aware, however, that you do not receive feedback to the display as to the success of startup; using the GUI provides more direct error reporting. Please see “Controlling Cluster Programs from a Command Line” on page 54 on instructions on how to do so.

## Adding a Slave to the Cluster

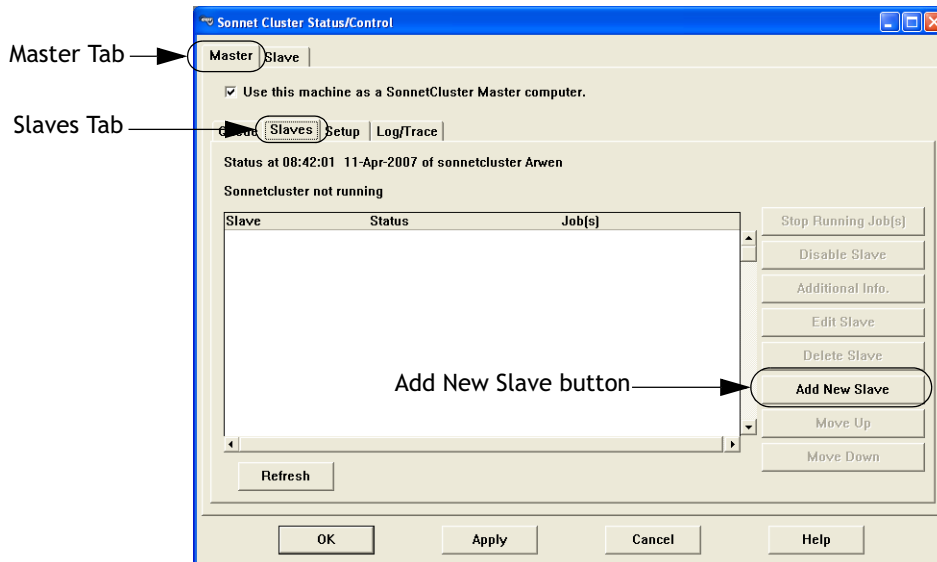
To add a slave definition to your cluster, follow the instructions below. If you are continuing from the previous section, the Sonnet Cluster Status/Control window is already open. If this is so, continue at Step 3 on page 19. If the window is not presently open, do the following:

- 1 **On the master computer, open the Sonnet task bar.**
- 2 **Select *Admin* ⇒ *SonnetCluster Status/Control* from the main menu of the Sonnet task bar.**

The SonnetCluster Status/Control window appears on your display.

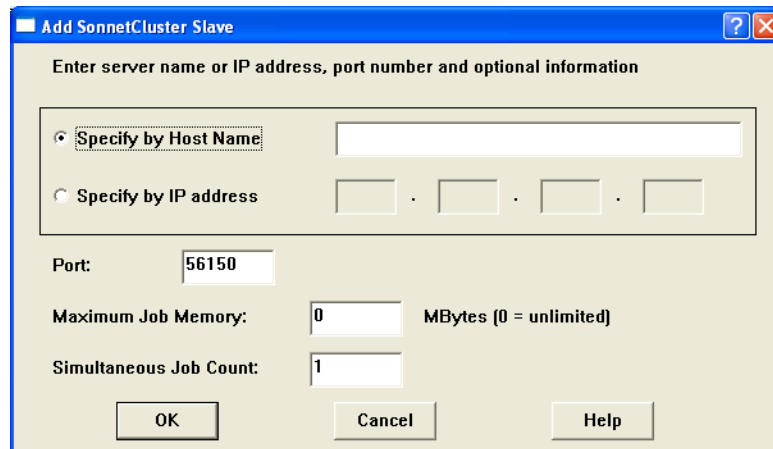
- 3 If it is not already selected, click on the Master tab at the top of the window.
- 4 Click on the Slaves tab in the Status/Control window.

The appearance of the window is updated and should appear similar to the view shown below.



- 5 To add a new slave to the cluster, click on the Add New Slave button in the Status/Control window.

The Add SonnetCluster Slave dialog box appears on your display.



## Setting Up *emCluster*

---

**6 Enter the host name or IP address of the slave computer you wish to add to your Sonnet Cluster.**

Click on the appropriate radio button, then enter the host name or IP address. If you enter a host name, it must be defined on the master computer. This name is the same as the computer would use in a domain or on a Windows network.

**7 Enter the Port number that will be used to communicate with the slave computer.**

This should match the port number entered when setting up the slave computer in Step 11 on page 26.

**8 Enter the maximum amount of memory which may be used for an analysis on the slave.**

The default value of zero allows an unlimited amount of memory on the slave computer to be used so that an analysis is only limited by the amount of physical memory on the slave computer. This setting ensures that the slave analyzes any size problem. If you wish to limit the size of the analysis subjobs that can be performed on this slave, enter the amount of memory to which you wish to limit the analysis in the Maximum Job Memory text entry box. The amount of memory should be entered in MBytes. Note that if you have a multi-core computer, you must assume that the total amount of memory allocated is the memory limit times the number of processing cores.

**9 Enter the simultaneous job count allowed on this slave.**

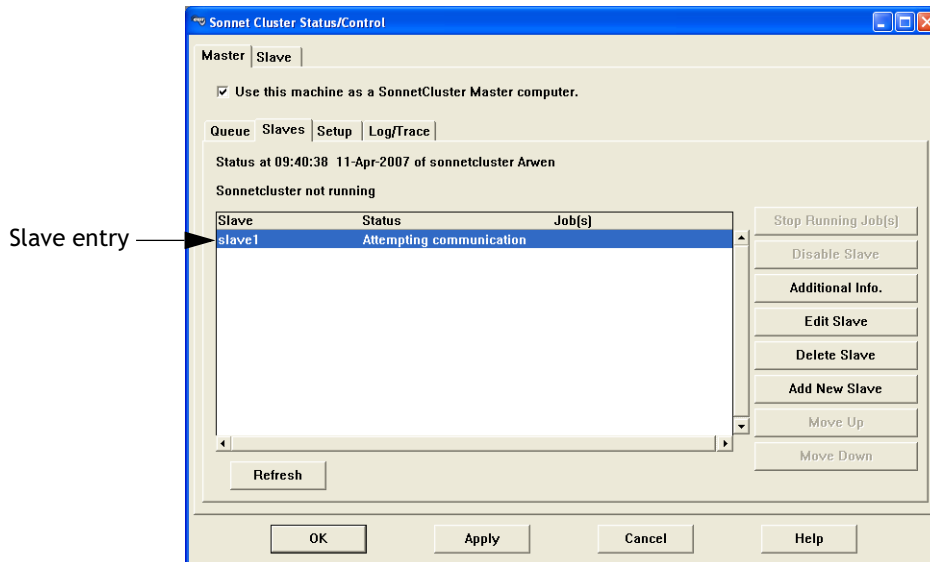
Enter the number of jobs that you wish to allow this slave to process at any given time. The default value is one. The value you enter is the number of jobs the master is permitted to send to this slave at the same time.

**10 Click on the OK button to close the dialog box and apply the changes.**


A confirmation message is displayed.

### 11 Click on Apply Now.

The slave computer just added now appears in the status window.



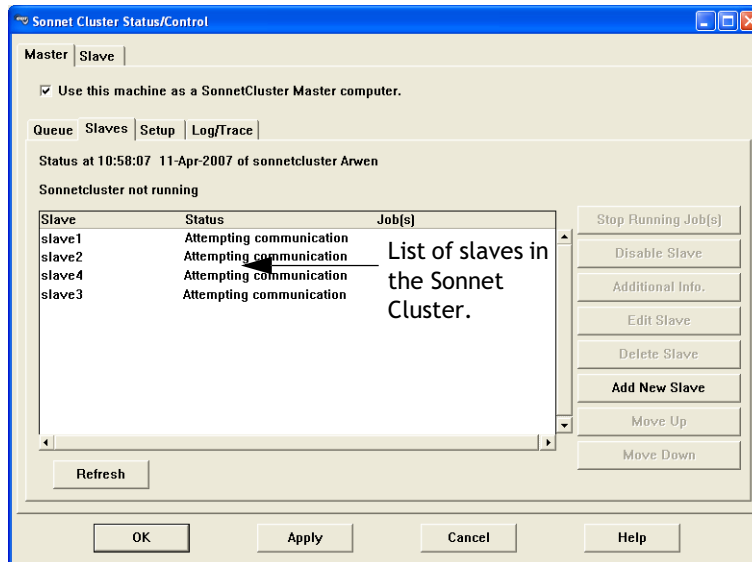
If the host name is not recognized, you receive the error message shown below. In this case, you should make sure that the host name is defined on your master computer.



## Setting Up *emCluster*

---

You should repeat this procedure until all the slaves are added to your Sonnet Cluster. The Sonnet Cluster Status/Control window is shown below listing four slave computers in the cluster. Note that jobs are submitted to the slaves in the



order in which the slaves appear in this list. The only exception is that if the master computer also functions as a slave, this is always the last slave to be used. You can override this setting in the *sonnetcluster* initialization file; please see the `UseThisServerAsSlaveLast` entry in Appendix II, “The Sonnet Cluster Initialization File” on page 94.

### Configuring a Slave (Done on Each Slave)

Once the master computer is defined and running, the next step in setting up your Sonnet Cluster is configuring the slaves that will do the analysis processing. The slave computers run the programs *emserver* and *em*. Slave computers receive subjobs from the master computer and send back analysis results. The master

computer combines these results before returning them to the user at the client from which the job was submitted. You will have to perform the procedure below on each slave computer in your Sonnet Cluster network.

**NOTE:**

If you are using a cluster comprised of computers all of which are using the same installation of the software (this would typically be computers using Linux operating systems), you can simplify setting up your slaves by using the “Apply slave configuration to ALL SonnetCluster Slave computers.” checkbox. If this is the case, please see “Configuring Slaves Using a Common Installation” on page 37 for instructions on setting up your slaves instead of using the next section.

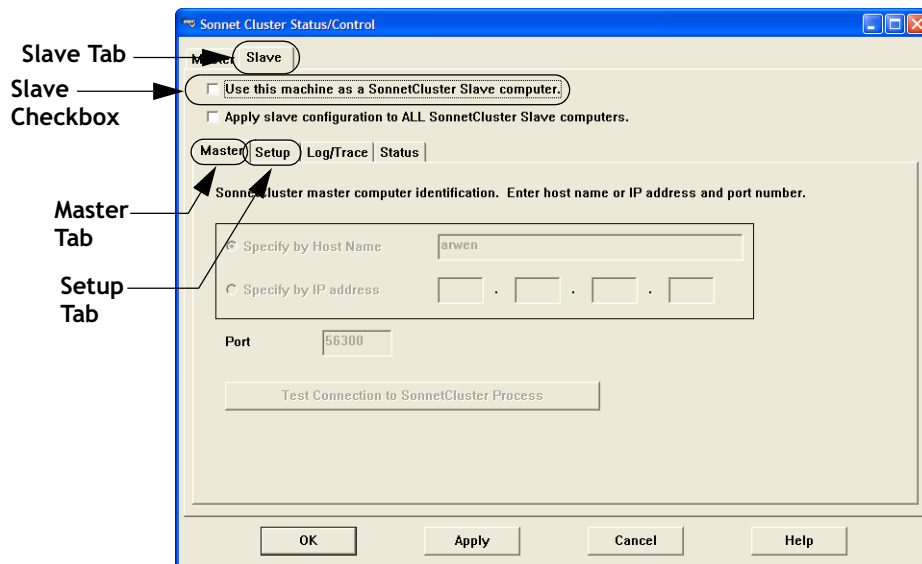
### Selecting a Computer as a Slave

- 1 Open the Sonnet task bar on the computer which you wish to setup as a slave in the cluster.
- 2 Select *Admin* ⇒ *SonnetCluster Status/Control* from the task bar main menu.

The Sonnet Cluster Status/Control window appears on your display.

- 3 Click on the Slave tab at the top of the window.

This tab allows you to designate this computer as a slave and identify the master computer in the Sonnet Cluster. The appearance of the window changes and should appear similar to the view shown below.



## Setting Up *emCluster*

---

**4 Select the Use this machine as a SonnetCluster Slave computer checkbox.**

This selects the computer as a slave in the cluster. All the controls in the window are enabled when this checkbox is selected.

**5 Click on the Master tab below the Use this machine as a SonnetCluster Slave computer checkbox.**

**6 Specify the master computer in the Sonnet Cluster.**

You may identify the master computer by a host name or IP address. This name is the same as the computer would use in a domain or on a Windows network. Select the type of identification by clicking on the appropriate radio button, then entering the host name or IP address in the corresponding entry field.

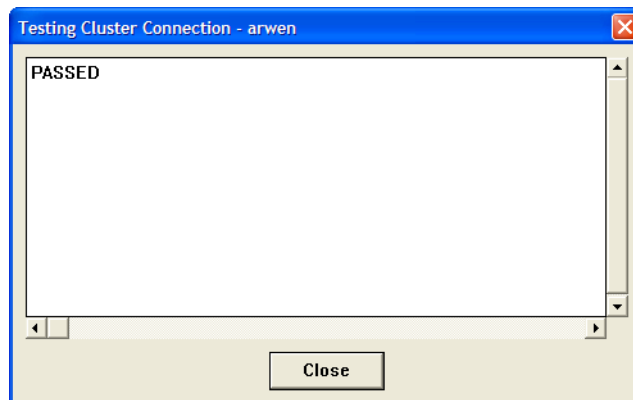
**7 Enter the port number by which this slave computer communicates with the master computer.**

This should be the same value entered as the port number when setting up the master computer. The default value is 56300.

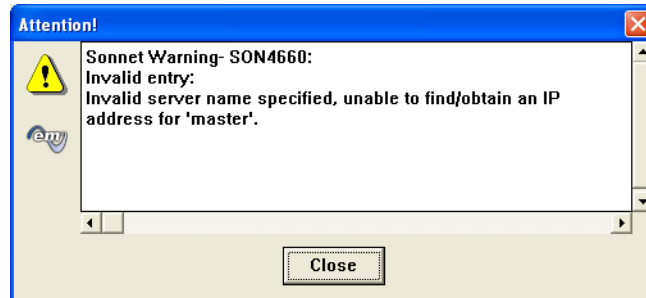
### Testing the Master/Slave Connection

**8 Click on the Test Connection to SonnetCluster Process button in the window.**

If the slave is able to communicate with the master computer across your network, then you will receive a PASSED message as shown below.



If the Host name or IP address entered is not valid, you will receive an error message as shown below.

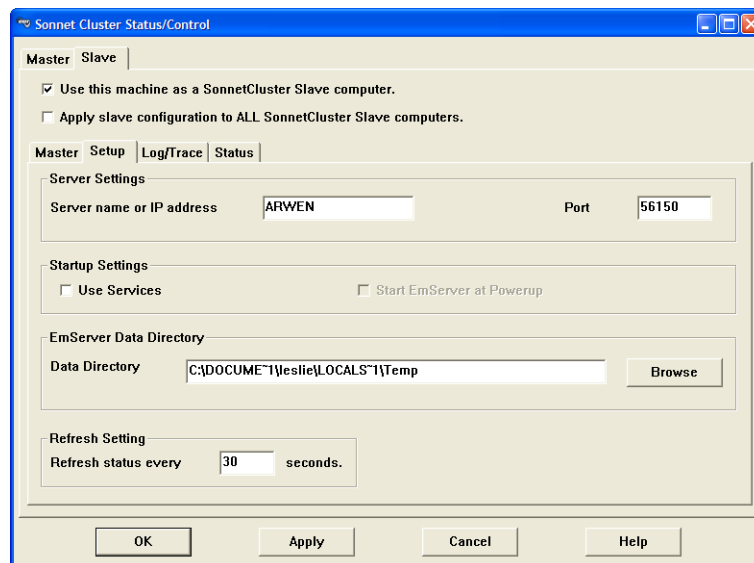


If you receive an error message, check to ensure that you are using the correct Host name or IP address for the master computer. If you are using a host name, you should check to see that the correct IP address is being used for the Host name. If the host name and IP address are correct, check your firewall configuration to make sure that the host name is correct. You should not continue the setup until you have established communications between the slave and master computer.

### Slave Settings

- 9 Click on the Setup tab in the Sonnet Cluster Setup/Control window.

The appearance of the window is updated as shown below.



## Setting Up *emCluster*

---

- 10 Enter the Server name or IP address of this computer in the Server name or IP address text entry box in the Server Settings section of the window.**

The server name is the host name of this slave computer. If you enter a server name, the name must be recognized by the master computer. This name is the same as the computer would use in a domain or on a Windows network.

- 11 If you do not wish to use the default port number of 56150, enter the desired port number in the Port text entry box.**

This is the port used by this slave computer to communicate with the master computer of the Sonnet Cluster. This is the port number used when defining a slave on the master computer, see Step 7 on page 20. If you change the port number for the slave, retain a record of the port number used in order to complete the master configuration.

---

**NOTE:** The default values used for the data directory and refresh status settings do not usually need to be changed. For more information on these settings, please see online help for this dialog box.

---

### Starting *emserver* at Powerup Windows

If you wish to have the program *emserver* start automatically at the startup of this slave computer, you must install *emserver* as a service. You will need administrator privileges to do so.

- 12 To install *emserver* as a service, select the Use Services checkbox in the Startup Settings section of the window.**

This enables the Start EmServer at Powerup checkbox.

- 13 Select the Start EmServer at Powerup checkbox.**

Whenever this slave computer is restarted, the program *emserver* will automatically be started. Continue below at Step 16.

### Starting *emserver* at Powerup for Linux

To change your system startup files for a Linux system, you must run the program *soninstallservice*. For instructions on how to use this program,

---

**NOTE:** The *soninstallservice* command will require root permissions to properly install the desired services. Failing to login as root before executing this command may result in unsuccessful initialization of the Sonnet Services.

---

14 Open a terminal on the slave computer.

15 Type the following:

```
cd <Sonnet Directory>/bin
soninstallservice -h
```

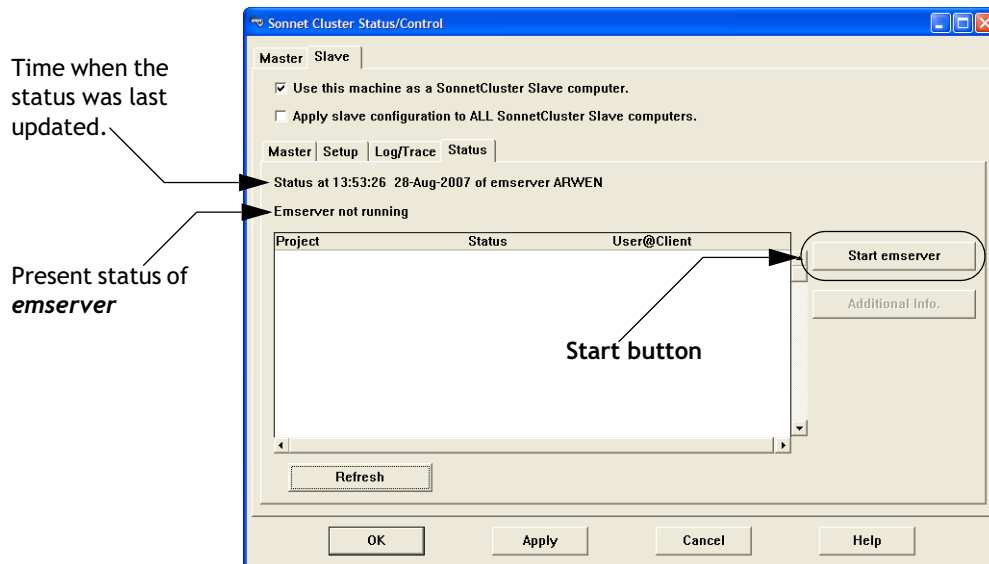
where <Sonnet Directory> is the directory in which the Sonnet Software was installed. Instructions for using soninstallservice appear on your display; follow these to install *emserver* (“install the remote *em* service”) as a service which starts at powerup.

### Starting *emserver*

Once you have entered all the setup values, you may start the *emserver* program on this slave as directed below.

16 Click on the Status tab of the Sonnet Cluster Status/Control window.

The appearance of the window changes. The present status of *emserver* is displayed just below the tab as shown below.




## Setting Up *emCluster*

---

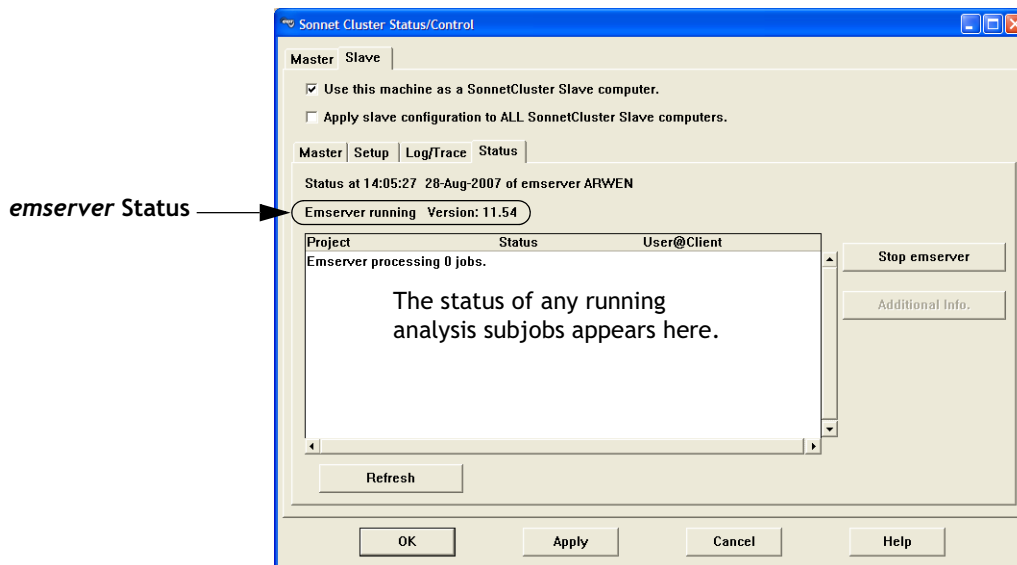
**17 Click on the Start *emserver* button.**

Since you have made changes to your slave setup, a query appears on your display.



**18 Click on the Yes button in the query box.**

The window is updated to indicate that *emserver* is running as shown below.



**19 Click on the OK button to close the dialog box and apply the changes.**

This completes setting up the slave computer.

---

**NOTE:** You must repeat this process for each slave you have added to the cluster.

---



### TIP

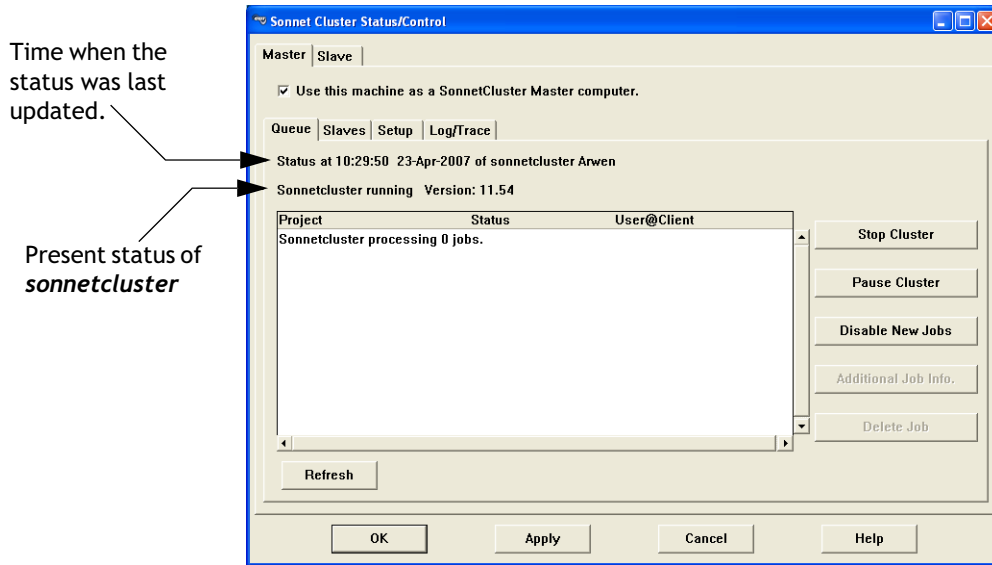
You may also start *emserver* by using a command line. Be aware, however, that you do not receive feedback to the display as to the success of startup; using the GUI provides more direct error reporting. Please see “Controlling Cluster Programs from a Command Line” on page 54 on instructions on how to do so.

---

### Checking the Slaves

Once you have completed adding all the slave computers to your Sonnet Cluster, you should check to ensure that all the slaves are enabled.

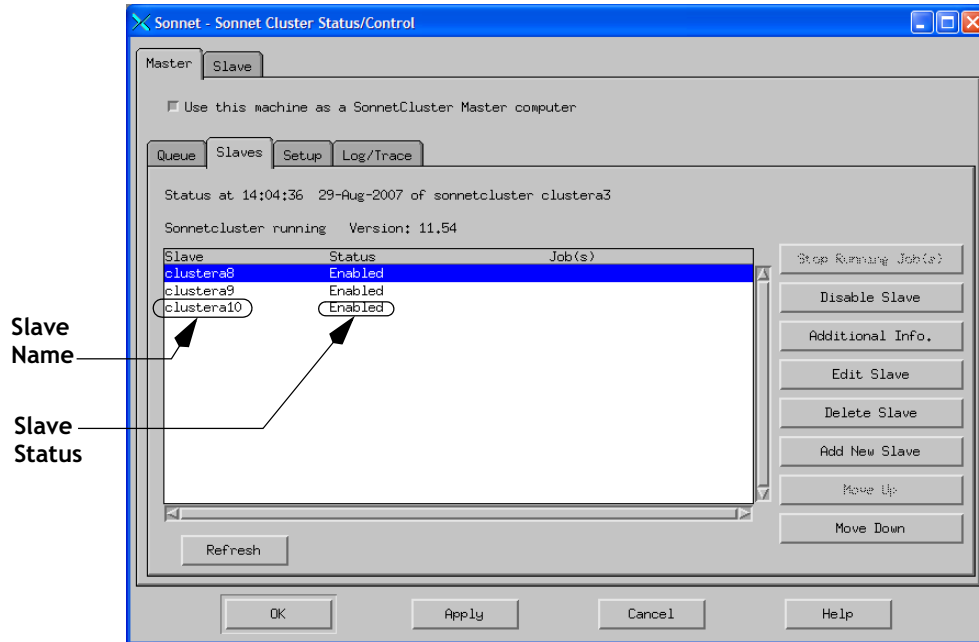
- 1 **Open the Sonnet task bar on the master computer and select *Admin* ⇒ *SonnetCluster Status/Control* from the main menu of the Sonnet task bar.**
- 1 The Sonnet Cluster Status/Control window appears on your display as shown below:



## Setting Up *emCluster*

### 2 Click on the Slaves tab.

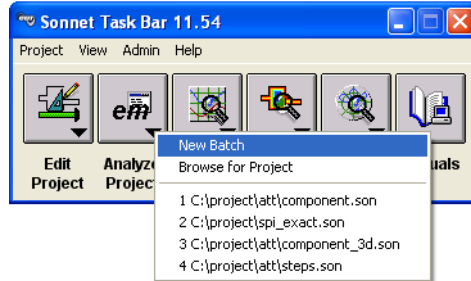
The appearance of the window changes to display the status of the slaves as shown below. The word “Enabled” should appear after each slave in the list.



## Setting up a Client Computer

Once you have configured the master computer and slave computers for the cluster, you will need to set up the clients. Client computers are the computers from which a user will submit an analysis job to the master computer to be sent to the processing cluster for analysis. The program *emclient* runs on a client computer so that the client can communicate with the master computer in the cluster and submit jobs. Setting up the client is done in the analysis monitor. Detailed instructions for setting up a client computer are below. You should perform this setup on each client computer you wish to use in your Sonnet Cluster.

- 1 **Open the Sonnet task bar on the computer you wish to setup as a client.**

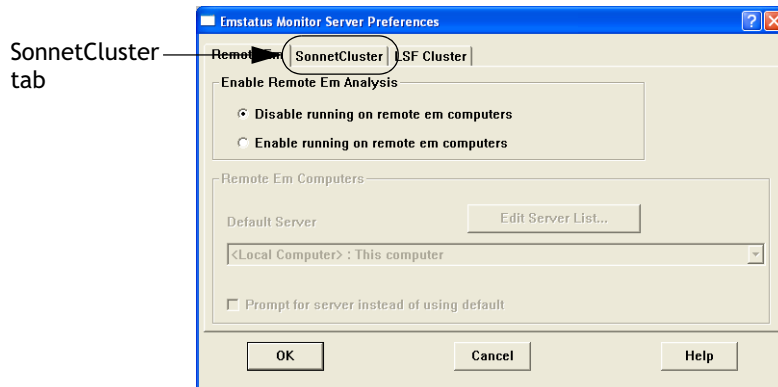


- 2 **Click on the Analyze Project button, then select New Batch from the pop-up menu.**

If the New Batch Creation window appears, select <Local Computer> from the list of servers. The analysis monitor appears on your display.

- 3 **Select File ⇒ Preferences from the main menu of the analysis monitor.**

The Emstatus Monitor Server Preferences dialog box appears on your display.



- 4 **Click on the SonnetCluster tab in the dialog box.**

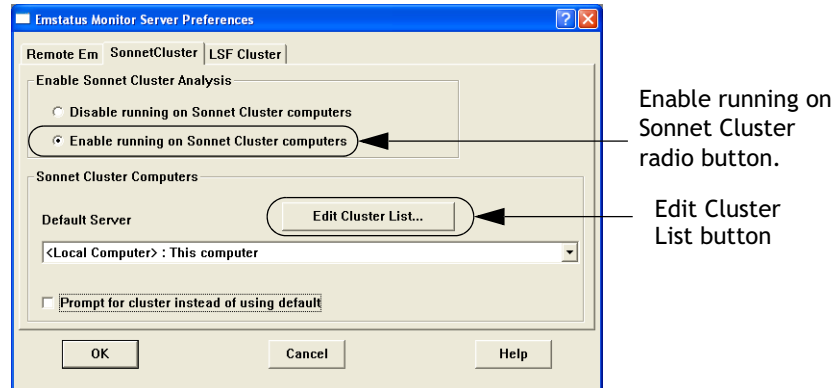
The appearance of the dialog box is updated.

## Setting Up *emCluster*

---

**5 Click on the Enable running on Sonnet Cluster computers radio button.**

This enables the controls in the window.



**6 Click on the Edit Cluster List button.**

This opens the Edit SonnetCluster Server List dialog box which allows you to define servers for this client. Servers are computers, or computer clusters that will perform processing submitted by this computer. You may add multiple servers at one time, but for our example, only one cluster is added to the server list.

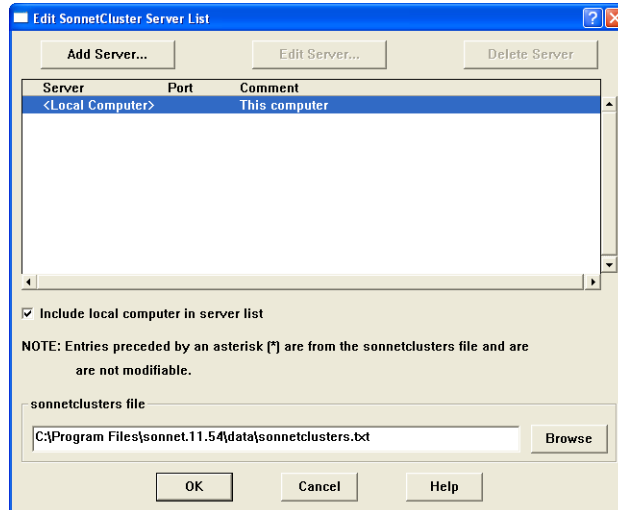
By default, the client computer is always included in the list as a server. If you do not wish the local computer to appear in the list, then clear the Include local computer in server list checkbox.

You must add the Sonnet Cluster you have defined to the server list. You identify the cluster by the name of the master computer. You may also directly edit the server list by editing the text file specified in the sonnetclusters file text entry box. Changes made to the specified text file will be reflected in this dialog box.



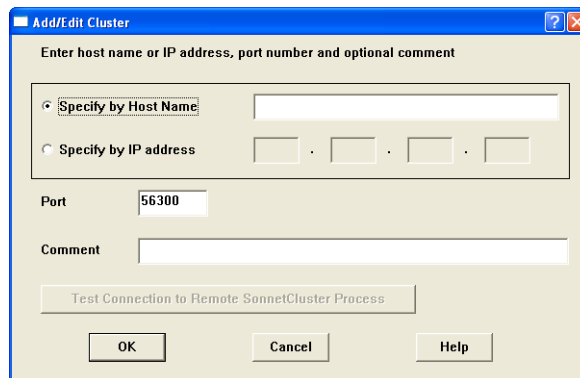
### TIP

It is possible to keep the sonnetcluster file on a network drive and allow multiple clients access, so that changes made in the server list affect all clients instead of having to make a change on each client.



- 7 Click on the Add Server button to add your Sonnet Cluster as a server.

The Add/Edit Cluster dialog box appears on your display.



- 8 Specify the master computer of the cluster you wish to add by entering the host name or IP address.

If you use a host name, the client computer must recognize the host name. This name is the same as the computer would use in a domain or on a Windows network.

## Setting Up *emCluster*


---

- 9 **If you do not wish to use the default port number of 56300, enter the desired port number in the Port text entry box.**

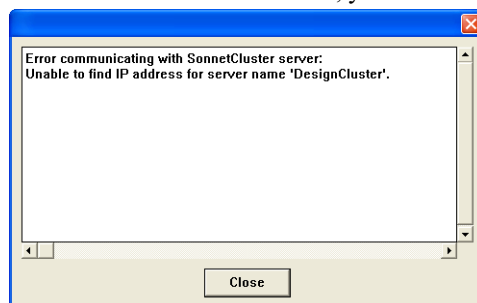
This port number is the port that the client computer uses to communicate with the master computer. This should be the same as the port number entered when configuring the master in Step 8 on page 15. Optionally, you may enter a comment in the Comment text entry box which will appear in the server list and help identify this cluster.

- 10 **Once you have defined the master computer, click on the Test Connection to Remote SonnetCluster Process.**

This will test communications between this client computer and the master computer. If the test is successful, you will receive a PASSED message as shown below.



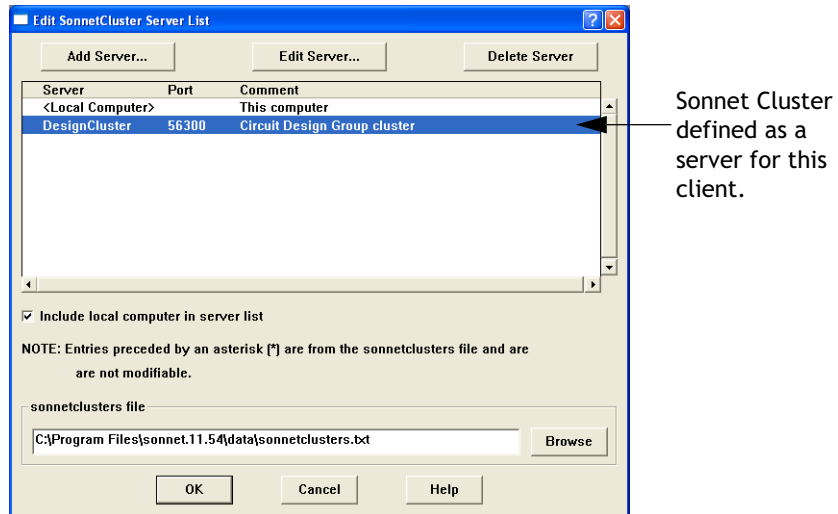
If the communications test fails, you will receive the error message shown below.



If you receive an error message, check to ensure that you have the correct IP address for the master computer or that the host name is correctly identified and points to the correct IP address. If the host name and IP address are correct, check your firewall configuration to make sure that the host name is correct.

- 11 Click on the OK button to close the dialog box and apply the changes.

The Sonnet Cluster you defined as a server now appears in the server list.



**NOTE:**

You may define multiple Sonnet Clusters as servers for any given client.

---

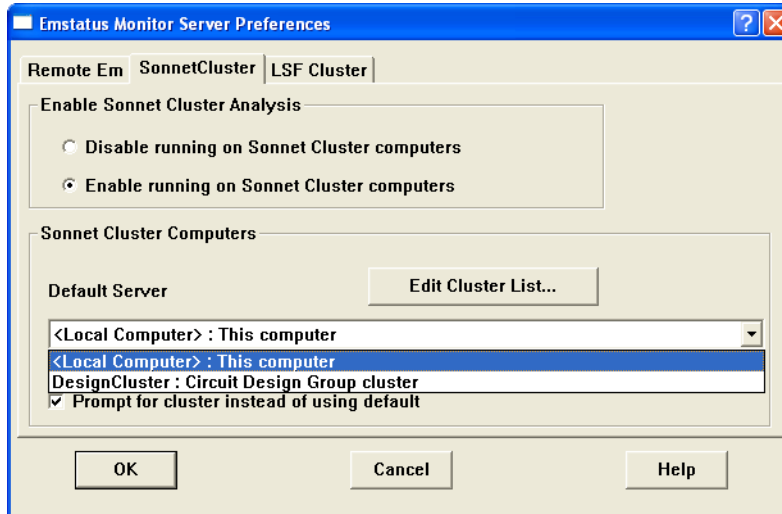
## Setting Up *emCluster*

---

12 Click on the OK button to close the dialog box.

13 Select the desired default server from the drop list.


The Sonnet Cluster that you just added as a server appears in this list. Analysis jobs will default to being run on the server selected here. If you wish to choose



which server to use each time you submit a job from this client, select the “Prompt for cluster instead of using default” checkbox. When you run an analysis from this client, a window appears which allows you to choose on which server you wish to run your analysis. For more details, see “Running an Analysis on your Sonnet Cluster” on page 41.

14 Click on the OK button to close the Preferences dialog box.

A notice appears on your display. In order to use the cluster you have just set up, you will need to exit from this session of the analysis monitor (batch) and open a new batch.



15 Click on the OK button to close the notice.

This completes the setup of Sonnet Cluster.

## Testing the Sonnet Cluster

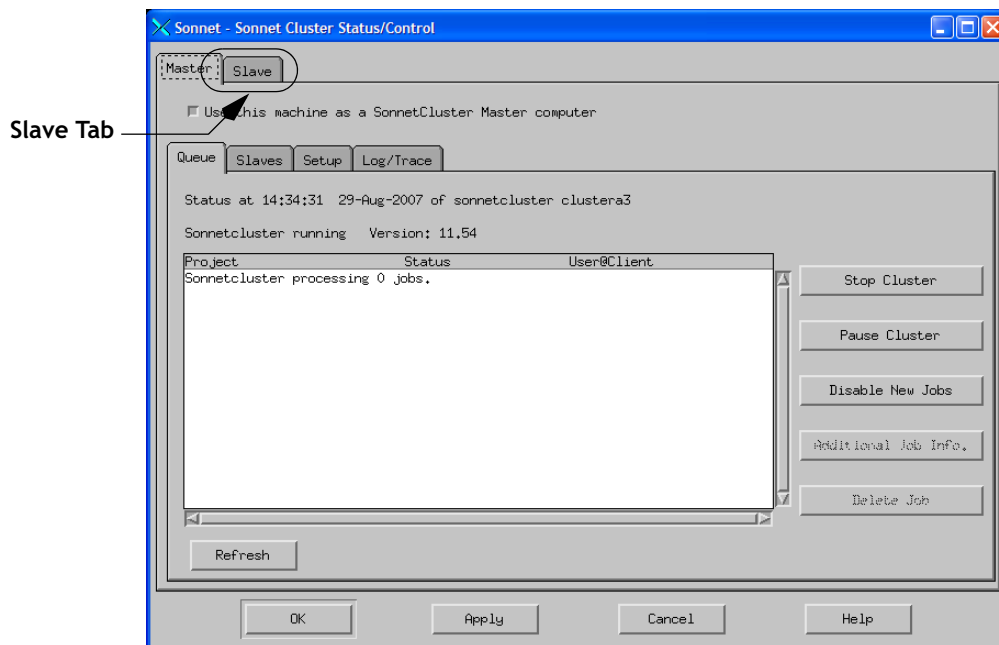
Once you have completed setting up your Sonnet Cluster, we suggest running an analysis on a simple example file, such as [Steps](#), to ensure that the cluster is functioning correctly. For directions on running an analysis on the cluster, please see “Running an Analysis on your Sonnet Cluster” on page 41.

## Configuring Slaves Using a Common Installation

If you are using a cluster comprised of computers all of which are using the same installation of the software (this would typically be computers using Linux operating systems), you can simplify setting up your slaves by using the “Apply slave configuration to ALL SonnetCluster Slave computers.” You first configure your master computer as a slave using the settings you wish to propagate. To do so, follow the directions below. Note that your master computer should be configured and running the program *sonnetcluster* BEFORE performing this procedure.

- 1 **On the Master computer, select *Admin* ⇒ *SonnetCluster Status/Control* from the main menu of the Sonnet task bar.**

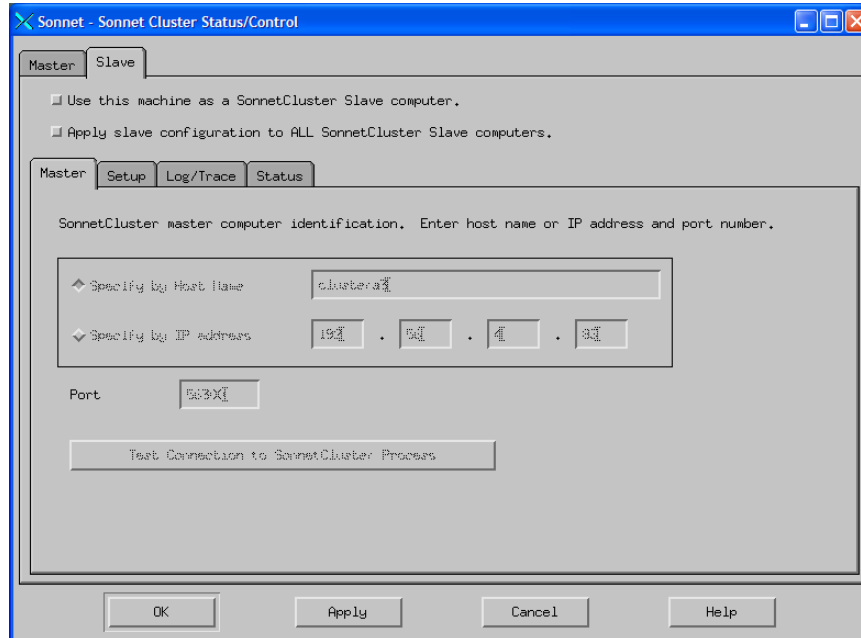
The SonnetCluster Status/Control window appears on your display.



## Setting Up emCluster

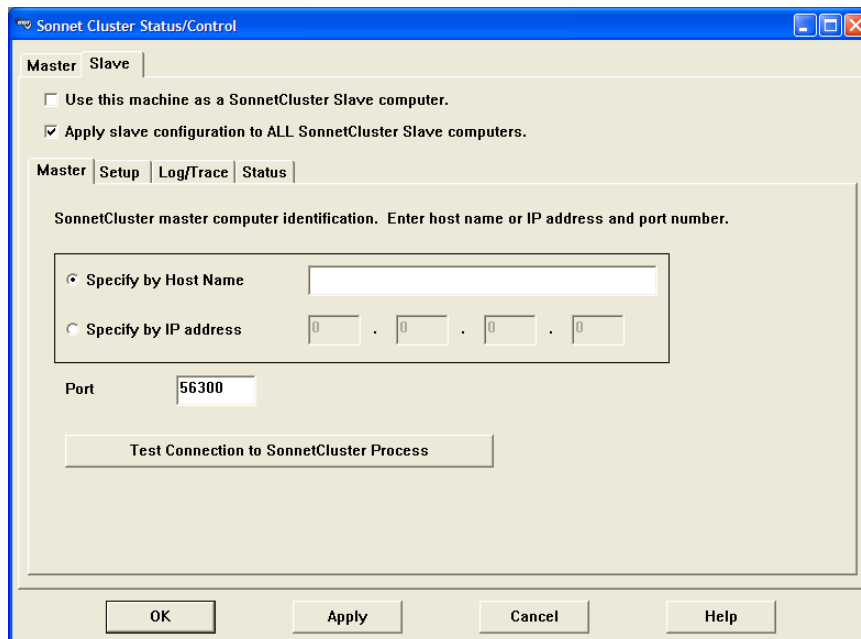
- 2 Click on the Slave tab in the Sonnet Status/Control window.

The appearance of the window is updated.



- 3 Select the “Apply slave configuration to All Sonnet Cluster Slave computers.” checkbox.

This enables the controls on this tab.



- 4 **Enter the Server name or IP address of the Master computer in the Server name or IP address text entry box.**

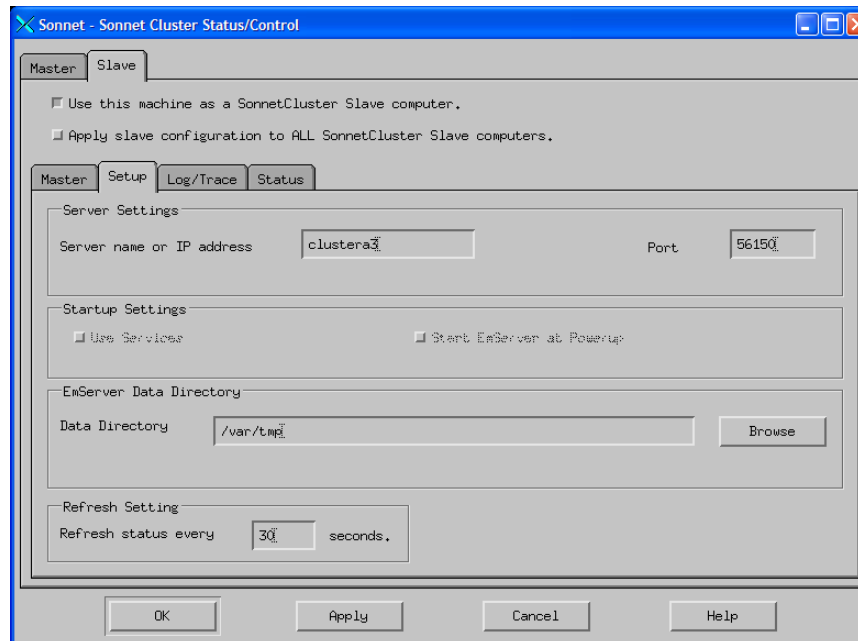
If you enter a server name, the name must be recognized by all slaves and clients in the Sonnet Cluster. This name is the same as the computer would use in a domain or on a Windows network.

- 5 **If you do not wish to use the default port number of 56300, enter the desired port number in the Port text entry box.**

This is the port on the Master used by the slave and client computers to communicate with the master computer. This is the same port number which should be entered when you add the cluster as a server to your client in Step 9 on page 34.

- 6 **Click on the Setup tab in the Sonnet Cluster Setup/Control window.**

The appearance of the window is updated as shown below.



- 7 **Enter the Server name or IP address of this computer in the Specify by Host Name or IP address text entry box in the Server Settings section of the window.**

The server name is the host name of this slave computer. If you enter a server name, the name must be recognized by the master computer. This name is the same as the computer would use in a domain or on a Windows network.

## Setting Up *emCluster*

---

- 8 If you do not wish to use the default port number of 56150, enter the desired port number in the Port text entry box.**

This is the port used by this slave computer to communicate with the master computer of the Sonnet Cluster. This is the port number used when defining a slave on the master computer, see Step 7 on page 20. If you change the port number for the slave, retain a record of the port number used in order to complete the master configuration.

**NOTE:**

---

**The default values used for the data directory and refresh status settings do not usually need to be changed. For more information on these settings, please see online help for this dialog box.**

---

- 9 Once you have completed entering the settings, click on the “Apply slave configuration to ALL SonnetCluster Slave computers” checkbox.**

When you click on the OK or Apply button in the window, all of the settings here are applied to all of the slaves in the cluster.

You will still need to go to each slave and start the program *emserver*. For instructions on how to do this, see “Starting *emserver*” on page 27. Also, if you wish to have *emserver* start up automatically after reboot, you need to go to each slave and install *emserver* as a service. For instructions on how to do this, see “Starting *emserver* at Powerup for Linux” on page 26.

Once you have configured all your slaves, you should continue your setup at “Checking the Slaves” on page 29.

## Cluster Slave Versus Remote *em* Server

You should not define the same computer as both a slave in a Sonnet Cluster and as a remote *em* server. A slave in a Sonnet Cluster may not run simultaneous jobs. Therefore, if a remote *em* job is submitted by a remote user ahead of a subjob from the cluster master, the subjob will be queued behind the remote *em* job and is not processed until the remote *em* job is completed. Since the master computer must wait until data from all the subjobs is returned from the slaves before combining the analysis data and sending it back to the submitting client, the analysis on the cluster is delayed for the amount of time it takes to complete the remote *em* job on that one slave. This reduces the efficiency of your Sonnet cluster.

## Chapter 2 Analyzing using Sonnet Cluster

---

Using a Sonnet Cluster allows you to take advantage of your computer cluster to improve the efficiency of your Sonnet *em* analysis. The discussion below provides some suggestions on how to use a Sonnet Cluster for analysis followed by instructions on running an analysis and overriding the default setting of one analysis frequency per subjob.

You should be aware, however, that there is an overhead cost to splitting the job to run on the cluster and assembling the analysis results. Using a Sonnet cluster for analysis of small jobs requiring a low number of frequency samples may show smaller efficiency gains due to cluster overhead. Larger analyses that are dominated by solve time show better gains. Therefore, we recommend small jobs be processed directly on the user's local computer and larger jobs be analyzed on the Sonnet cluster.

### Running an Analysis on your Sonnet Cluster

In order to run an analysis on the cluster, Sonnet must be installed and configured on all the cluster computers: master, slaves and clients. The program *sonnetcluster* must be running on the master and the program *emserver* must be running on all the slaves. For directions on doing this, refer to [“Sonnet Setup” on page 60](#).

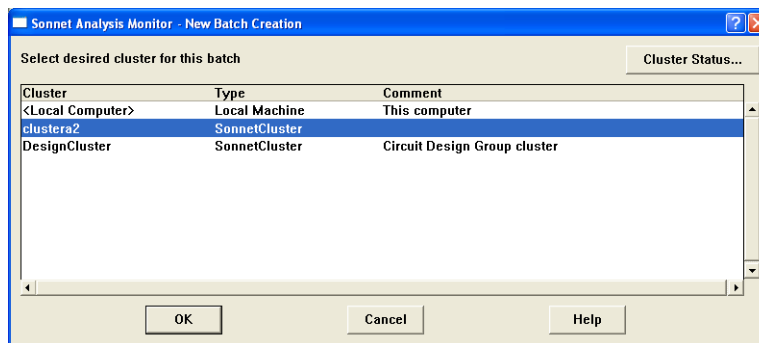
## Setting Up *emCluster*

---

To run an analysis on the Sonnet cluster, do the following:

- 1 On a cluster client, open a Sonnet project that is ready for analysis in the project editor and click on the Analyze button on the project editor's tool bar.**
- 2 The New Batch Creation dialog box appears on your display.**

Note that this dialog box only appears if you have selected the Prompt for server checkbox in the Preferences dialog box of the analysis monitor when you setup your cluster computing. This dialog box contains a list of Sonnet clusters available to which you may submit jobs.



If you are not prompted, then the Sonnet cluster is your default processing server, and you may continue at Step 5.

- 3 Click on the Sonnet Cluster to select it for the analysis.**

The selected server is highlighted. In the picture above, the cluster “clustera2” is selected.

- 4 Click on the OK button in the New Batch Creation dialog box.**

This closes the dialog box. The Sonnet cluster is now associated with the analysis monitor window which will subsequently open. All jobs added to the batch list of the analysis monitor will be analyzed on the Sonnet cluster.


- 5 When the New Batch Creation dialog box is closed, the analysis monitor appears on your display and your analysis job starts running.**

Information is returned to the analysis monitor on the client as it is processed by the master computer. The types of information which appear depend on the splitting, if any, of your project. If any errors occur with the cluster, a status message is posted in the status section of the analysis monitor just under the progress bar.

If there is a loss of communication between the master computer and the submitting client while an analysis is running, it is possible to reconnect to an ongoing analysis or recover the results of a completed analysis. For details on reconnecting or recovering disconnected jobs, please see online Help for the analysis monitor.

### An Analysis Example

An example of running an analysis on a cluster with examples of output received during the analysis is shown below. This example shows an ABS analysis, from 0.2 GHz to 2 GHz, of the example project `spi_exact.son` analyzed on the “cluster2” Sonnet cluster which consists of seven slaves: `cluster3`, `cluster5`, `cluster6`, `cluster7`, `cluster8`, `cluster9`, and `cluster10`. The circuit is pictured below.



Note: This circuit is using a much smaller cell size than needed to accurately model the response in order to make the processing time significant per frequency.

Once the job is submitted to the cluster, the master computer subsections the circuit and calculates the seven discrete frequencies at which to analyze the circuit in order to create the adaptive response data for the ABS analysis. Seven discrete

## Setting Up emCluster

frequencies is the default number for an ABS analysis run on a cluster; for more details, please refer to “ABS Processing” on page 49. Seven subjobs are then sent to the slaves in the cluster. The analysis monitor is shown below.

Indicates that the project spi\_exact.son is being analyzed on the Sonnet Cluster “clustera2”

The project was subsectioned on the master clustera2 before being sent to the slaves.

Seven subjobs are running on the cluster; none are pending. There were enough slaves on the cluster to allow all the subjobs to be analyzed simultaneously. If there had been less slaves or more subjobs, some jobs would be listed as pending

You may observe the status of the cluster while the analysis is running by selecting the command *View* ⇒ *Remote Status* from the main menu of the analysis monitor. The SonnetCluster status window appears on your display. To see the status of your submitted analysis job, click on the Queue tab as shown below.

Name of Sonnet Cluster

User “leslie” submitted this job from the client “Arwen.”

Project Name

The job is presently running on the cluster.

## Chapter 2 Analyzing using Sonnet Cluster

If you click on the “Slaves” tab in the status window, the status of the subjobs on the slaves is shown as pictured below.

The screenshot shows a window titled "SonnetCluster status - clustera2". It has two tabs: "Queue" and "Slaves". The "Slaves" tab is selected. The window displays the following information:

SonnetCluster running. Version: 11.54

Slave	Status	Job[s]
clustera3:56150	Running	spi_exact.son - leslie@ARWEN
clustera5:56150	Running	spi_exact.son - leslie@ARWEN
clustera6:56150	Running	spi_exact.son - leslie@ARWEN
clustera7:56150	Running	spi_exact.son - leslie@ARWEN
clustera8:56150	Running	spi_exact.son - leslie@ARWEN
clustera9:56150	Running	spi_exact.son - leslie@ARWEN
clustera10:56150	Running	spi_exact.son - leslie@ARWEN

Annotations in the image:

- An arrow points to the "Slaves" tab with the text: "Version of the software being run on the Sonnet cluster."
- An arrow points to the "SonnetCluster running. Version: 11.54" text with the text: "Version of the software being run on the Sonnet cluster."
- An arrow points to the "Slave" column header with the text: "List of slaves in the cluster. The name of each slave is followed by the port through which the slave is communicating with the master. All of these slaves are using the default port 56150."
- An arrow points to the "Job[s]" column header with the text: "This entry line indicates that a subjob of the analysis job spi\_exact.son, submitted to the cluster by user 'leslie' from client 'Arwen,' is running on the slave 'clustera10.'"

Buttons at the bottom: Close, Refresh, Additional Info., Help.

This entry line indicates that a subjob of the analysis job spi\_exact.son, submitted to the cluster by user “leslie” from client “Arwen,” is running on the slave “clustera10.”

As each slave finishes a subjob, the results are returned to the master computer and the response data is output to the analysis monitor window. An example is shown below.

```

Job 6:
Mon Apr 30 17:09:51 2007. Frequency Sweep Combinations. ← Start Time
Em version 12.52 on clustera3 (Lnx-86) cluster ← Sonnet Version and Slave ID

Frequency sweep 1.

Frequency: 1.7 GHz ← Analysis Frequency
Frequency completed Mon Apr 30 17:16:27 2007. → Response Data
De-embedded S-Parameters. 50.0 Ohm Port Terminations.
Magnitude/Angle. Touchstone Format. (S11 S21 S12 S22).
1.70000000 0.798105 -55.05 0.570045 -153.5 0.570045 -153.5 0.797264 -70.31
Job 6 completed Mon Apr 30 17:16:27 2007. ← Completion Time of Subjob
    
```

## Setting Up emCluster

As analysis results are delivered back to the master from the slaves, the progress bar in the analysis monitor is updated and response data is displayed in the output window as shown below.

The screenshot shows the 'spi\_exact.son' analysis monitor window. The top status bar indicates 'Freqs: 6 Complete' and 'Latest Freq: 7 min. 9 sec.'. A progress bar shows 75% completion. Below the progress bar, it states 'The jobs are running...' and '0 pending; 1 running; 6 completed'. The left sidebar contains buttons for 'Response Data', 'Errors/Warnings', 'Timing Info', and 'Batch List'. The main output window displays the following text:

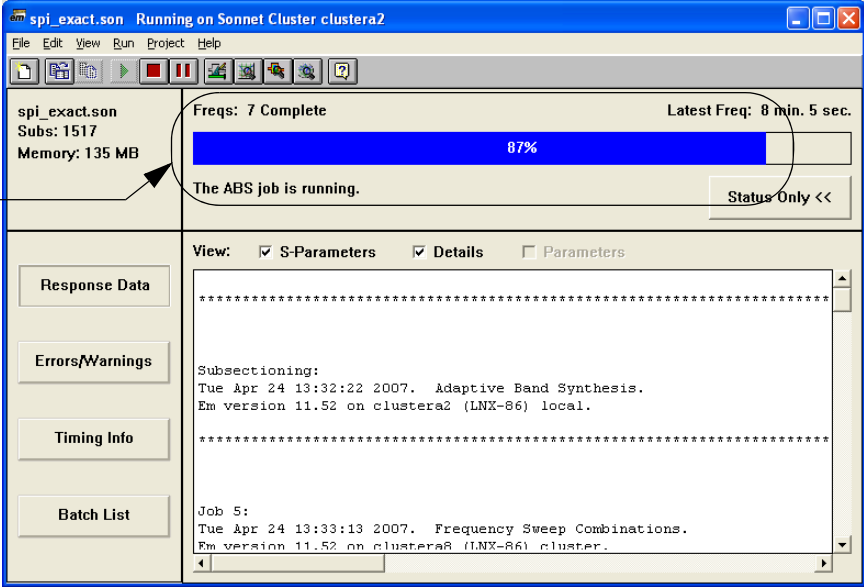
```
Subsectioning:
Tue Apr 24 13:32:22 2007. Adaptive Band Synthesis.
Em version 11.52 on clustera2 (LNX-86) local.
*****
Job 5:
Tue Apr 24 13:33:13 2007. Frequency Sweep Combinations.
Em version 11.52 on clustera8 (LNX-86) cluster.
Frequency sweep 1.
Frequency: 1.4 GHz
```

Six of the seven subjobs are completed and one is still running. Since each subjob is analyzing at one frequency, six frequencies are shown as complete, and the job is 75% complete.

Subjob 5 is being run on the slave "clustera8" and is running the analysis at 1.4 GHz.

When running an ABS analysis, as in this example, once all the subjobs have been returned to the master, the master combines the analysis results from the slaves and sends a single subjob with the combined discrete data to one of the slaves to fit the polynomial and create the adaptive data for the ABS analysis. This is shown in the analysis monitor pictured below.

The seven discrete frequencies are complete, and the ABS job-fitting the polynomial and creating the adaptive data-is being run. The analysis is 87% complete.



The screenshot shows a software window titled "spi\_exact.son Running on Sonnet Cluster clustera2". The window has a menu bar (File, Edit, View, Run, Project, Help) and a toolbar. The main area is divided into several sections. On the left, there is a sidebar with buttons for "Response Data", "Errors/Warnings", "Timing Info", and "Batch List". The top right of the main area displays "Freqs: 7 Complete" and "Latest Freq: 8 min. 5 sec.". Below this is a progress bar showing "87%" completion. A status message reads "The ABS job is running." with a "Status Only <<" button. The bottom section is a text area with a "View:" label and checkboxes for "S-Parameters" (checked), "Details" (checked), and "Parameters" (unchecked). The text area contains the following output:

```
*****  
Subsectioning:  
Tue Apr 24 13:32:22 2007. Adaptive Band Synthesis.  
Em version 11.52 on clustera2 (LNX-86) local.  
*****  
  
Job 5:  
Tue Apr 24 13:33:13 2007. Frequency Sweep Combinations.  
Em version 11.52 on clustera8 (LNX-86) cluster.  
*****
```

If the seven discrete data points are not enough to produce a model, then *em* runs an analysis at additional discrete data points until there is sufficient data to create an accurate model. In this case, the additional frequencies are analyzed sequentially on the same slave until the model is achieved.

## Setting Up emCluster

Once the adaptive data has been calculated and sent back to the master, the master combines the discrete data and adaptive data and sends the analysis results for the job back to the submitting client.

The screenshot shows a window titled "spi\_exact.son Finished on Sonnet Cluster clustera2". The window has a menu bar (File, Edit, View, Run, Project, Help) and a toolbar. The main area is divided into several sections:

- Top Left:** Project information: "spi\_exact.son", "Subs: 1517", "Memory: 135 MB".
- Top Right:** Frequency information: "Freqs: 7 Discrete [361 Adaptive]", "Latest Freq: 8 min. 5 sec.". Below this is a blue progress bar at "100%".
- Middle:** A status bar indicating "Analysis completed Tue Apr 24 13:42:41 2007." and a "Status Only <<" button.
- Bottom Left:** A sidebar with buttons for "Response Data", "Errors/Warnings", "Timing Info", and "Batch List".
- Bottom Right:** A text area showing the final analysis results. It includes the text: "Final ABS analysis: Tue Apr 24 13:42:04 2007. Adaptive Band Synthesis. Em version 11.52 on clustera3 (LNX-86) cluster." and "Final ABS analysis completed Tue Apr 24 13:42:31 2007." Below this, a summary of frequencies and timing is shown: "Discrete frequencies analyzed : 7", "Adaptive frequencies analyzed : 361", "Time submitted to sonnetcluster : Tue Apr 24 13:32:16 2007", "Time completed on sonnetcluster : Tue Apr 24 13:42:42 2007", and "Total time : 10 minutes 26 seconds".

Annotations with arrows point to the 100% progress bar and the final analysis text in the bottom right.

Analysis is complete and the status bar is at 100%.

The final ABS analysis was processed on the slave "cluster3."

The final output shows how many discrete frequencies were analyzed, how many adaptive frequencies were calculated, and the total running time on the cluster.

## Frequency Splitting

One way to increase the efficiency of your processing using the Sonnet Cluster is frequency splitting as described above. Multiple duplicate projects are created, each analyzing at a single frequency point and simultaneously submitted to the Sonnet cluster so that you are performing your *em* processing in parallel. Therefore, if you are analyzing at 10 discrete frequencies, 10 *em* jobs, each containing the same circuit, are submitted to the Sonnet cluster, one for each frequency. This is the default mode for a Sonnet Cluster.

You may override the one job per frequency setting for a particular batch window by using the *Run* ⇒ *Override Cluster Options* command in the analysis monitor window. This window allows you to run the analysis interactively on only one slave computer, or to limit the number of subjobs the analysis can be split into. In this case, each subjob is analyzed at some subset of the analysis frequencies. For example, if you are analyzing a circuit from 1 to 8 GHz with 1.0 GHz steps and limiting the analysis to two subjobs, each subjob would analyze the circuit at four frequencies.

For more details on using the overrides, please see “Automatic Calculation of Discrete Data Points” on page 49.

---

**NOTE:** You may not use frequency splitting when running an optimization or parameter sweep; the analysis is automatically run interactively on one slave.

---

## ABS Processing

The Adaptive Band Synthesis (ABS) technique provides a fine resolution response for a frequency band requiring only a small number of analysis points. *Em* performs a full analysis at a few points and uses the resulting internal, or cache, data to synthesize a fine resolution band. The output data consists of the discrete data points, frequencies at which the analysis engine, *em*, performs an electromagnetic analysis, and the adaptive data, which is data calculated using a rational polynomial fit. This is often the most efficient process in Sonnet by which you can obtain response data over a band. For more information about ABS, please refer to [Chapter 9, “Adaptive Band Synthesis \(ABS\)”](#) in the **Sonnet User’s Guide**.

By using your cluster to perform the analyses at the discrete data points in parallel, you can further increase the efficiency of an ABS sweep. There are two ways of accomplishing an ABS sweep when using a computer cluster. You may submit an ABS sweep and allow *sonnetcluster* to automatically calculate at which discrete data points to run a full analysis (automatic) or you may define linear sweeps or single frequency points in which you determine, in advance, at which discrete data points you run an analysis before attempting the ABS sweep (user-defined). You must use the Frequency Sweep combinations to combine the different types of sweeps when doing the user-defined ABS. The ABS sweep should be the last sweep specified.

### Automatic Calculation of Discrete Data Points

The automatic calculation of discrete data points is used when you submit an Adaptive sweep to the Sonnet cluster for analysis. In this case, *em* defaults to analyzing at seven discrete data points. These discrete data points always include the starting and ending frequencies of the band and the rest of the points are evenly distributed across the band. These discrete data points are then split up into one job per frequency point and submitted for analysis to the Sonnet cluster. Once all the results from the discrete data points are complete, *em* solves for the rational polynomial and produces the adaptive data. The default of seven data points is

## Setting Up *emCluster*

---

conservative to ensure that there are enough data points to produce an accurate model. An adaptive sweep run on your local computer would solve one discrete data point at a time and attempt to “fit” a polynomial after solving at each data point. Setting up seven discrete data points beforehand allows you to take advantage of the frequency splitting capabilities of cluster computing.

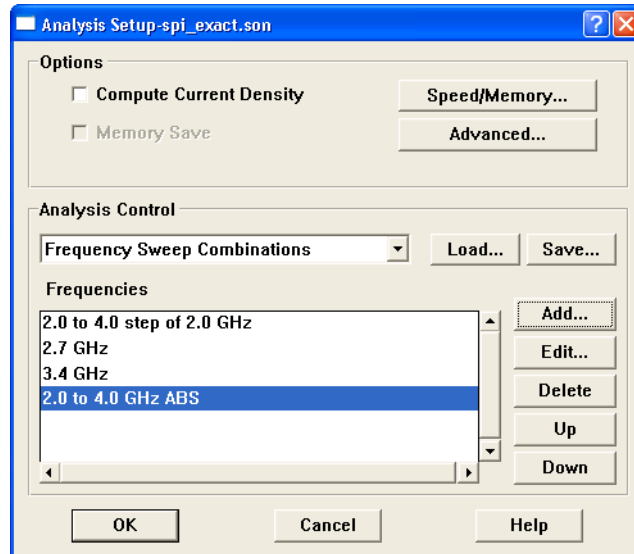
If the seven discrete data points are not enough to produce a model, then *em* runs an analysis at additional discrete data points until there is sufficient data to create an accurate model. In this case, the additional frequencies are done sequentially on the same server until the model is achieved.

### User Defined Discrete Data Points

For some ABS sweeps, you may have knowledge about your circuit such that you know the number of discrete data points needed or you have specific data points at which you desire a full analysis. In these cases, it may be more efficient to set up your analysis with a linear sweep first which analyzes the circuit at the desired frequency points before performing the ABS sweep. The linear sweep and the ABS sweep should be set up together using the Frequency Sweep Combinations option for Analysis Control.

For example, you are going to do an ABS analysis of a circuit from 2.0 to 4.0 GHz and are aware from previous analyses that analyzing the discrete data points 2.0 GHz, 4.0 GHz, 2.7 GHz, and 3.4 GHz provide enough discrete data for a polynomial fit. Therefore, there is no need to analyze at the default seven discrete

data points used for an ABS analysis on a cluster. You only need to analyze at four discrete points saving the processing time of analyzing at an additional three frequencies. In this case, your analysis setup would be as pictured below:



Note that the first entry in the Frequency Sweep Combinations analyzes the circuit at 2.0 GHz and 4.0 GHz, and the next two entries analyze at 2.7 GHz and 3.4 GHz respectively, followed by the ABS sweep. The ABS sweep will use the data obtained from the previous sweeps to attempt a polynomial fit before performing an analysis at any additional discrete data points.

This method also allows you to produce data for a specific frequency that is critical in your design.

## Overriding Cluster Options

When running an analysis on the cluster, you may override the default settings. To do so, perform the following:

- 1 **Open the Sonnet task bar and click on the Analyze Project button.**

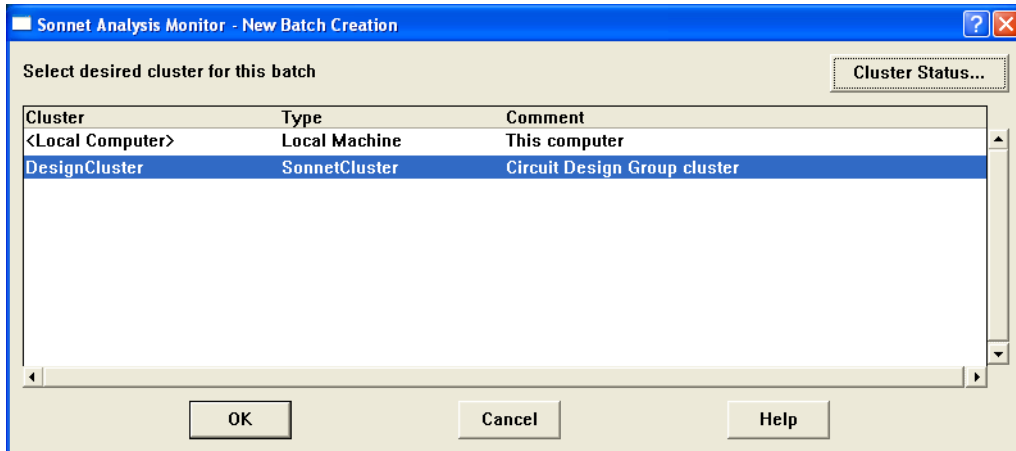
A pop-up menu appears on your display.

## Setting Up *emCluster*

---

**2 Select “New Batch” from the pop-up window.**

If you have the “Prompt for Server” option enabled in your Preferences, then the New Batch Creation window will appear on your display. If the New Batch Creation window does not appear and Sonnet cluster is selected as your default processing server, continue at Step 5.



**3 Select the desired cluster from the Server/Cluster list in the New Batch Creation window.**

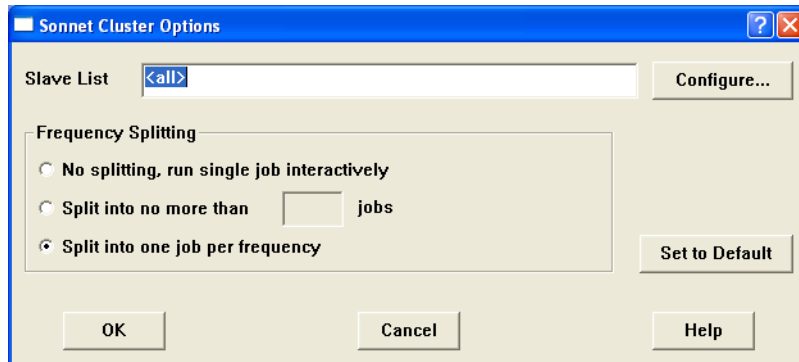
In the above illustration, the Sonnet cluster “DesignCluster” is selected to perform your processing. If you wish to check the present status of the cluster, click on the Cluster Status button in the upper right hand corner of the window.

**4 Click OK to close the window and execute the changes.**

The analysis monitor window appears in batch list mode.

- 5 Select **Run** ⇒ **Override Cluster Options** from the analysis monitor main menu.

The Sonnet Cluster Options dialog box appears on your display. The values which appear in this dialog box are the defaults.



- 6 Enter the desired slaves in the Slave List.

You may directly edit the slave list or click on the Configure button to the right which opens the Sonnet Cluster Host List Selection dialog box which allows you to select the desired slave computers. All of the slaves defined in your Sonnet cluster appear here.

If you do not enter any slaves in the slave list, then *sonnetcluster* uses the default, which is all the slaves available in the cluster. If this is true, then <all> appears in the list.



### WARNING

---

The cluster must be up and running in order to change the slave list. If you attempt to enter anything besides “ALL” when the cluster is not running, you will receive an error message.

---

- 7 Select the desired radio button to control the Frequency Splitting for the job.

#### NOTE:

---

You may not use frequency splitting when running an optimization or parameter sweep; the analysis is automatically run interactively on one slave regardless of the settings in this dialog box.

---

If you select “No splitting,” then the job is submitted as one analysis job to the cluster and one slave analyzes your circuit at all the specified analysis frequencies. When this option is selected, then response data and status information, such as you would see with a local analysis, is sent to the analysis monitor. This is defined as running interactively. This is similar to Remote *em* processing, except that the computing resource selection is controlled by *sonnetcluster*.

The Split into no more than <n> jobs selection limits the number of jobs your analysis may be split into. The project is duplicated the number of times you enter in the adjacent text entry box and each duplicate project analyzes the circuit for a subset of the requested analysis frequencies. For example, if your analysis is for the frequencies 2, 4, 6, 8, 10, 12, 14, and 16 GHz and you have entered 4 as the maximum number of jobs then the first job analyzes at 2 and 4 GHz, the second project at 6 and 8 GHz and so on.

The “Split into one job per frequency” option splits the project into one job per requested analysis frequency; this is the default mode. If more jobs are created than there are server hosts in your cluster, then multiple jobs will be submitted to some server hosts. If you have a sufficient number of server hosts in your cluster, this option results in the greatest efficiency for your processing.

- 8 If you want to reset the dialog box to the original values, then click on the Set to Defaults button.**

You will lose any changes made up until then in the dialog box.

- 9 Once all your entries are complete, click on the OK button to close the dialog box and apply the changes.**

This batch window will now run any jobs submitted to the cluster using the operating parameters input to the Override Cluster Options dialog box. These changes will only apply to this batch window. If a new batch window is opened, it will use the default values unless you once again override those options.

- 10 Add the projects you wish to use the cluster to analyze to the batch list.**

You may add projects by selecting *File* ⇒ *Add Project(s)* from the main menu of the analysis monitor window or by clicking on the Add Projects button on the analysis monitor tool bar.

- 11 Click on the Run button on the analysis monitor tool bar to analyze the projects in the batch list.**

## Controlling Cluster Programs from a Command Line

It is possible to start and stop both the *sonnetcluster* program, which runs on the Master, and the *emserver* program, which runs on all slaves, from a command line. Be aware, however, that you do not receive feedback to the display as to the success of the startup; using the GUI provides more direct error reporting. For instructions on starting each program, see the instructions below.

- 1 On UNIX and Linux systems, open a terminal. On a Windows system, open a Command prompt window.
- 2 Change the directory to “<Sonnet Directory>/bin” where <Sonnet Directory> is the directory in which Sonnet is installed.
- 3 Enter the command for the desired action below.

**Starting *sonnetcluster*:** `sonnetcluster -DaemonProcess`

**Stopping *sonnetcluster*:** `emclient -Mode SonnetCluster -SHUTDOWN`

**Starting *emserver*:** `emserver -DaemonProcess`

**Stopping *emserver*:** `emclient -Mode RemoteEm -SHUTDOWN`

- 4 When you are done executing the desired command, close the terminal or command prompt window.

