



Sun Fire™ V60x Compute Grid Rack System Release Notes

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 817-3074-11
October 2003, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

THIS PRODUCT CONTAINS CONFIDENTIAL INFORMATION AND TRADE SECRETS OF SUN MICROSYSTEMS, INC. USE, DISCLOSURE OR REPRODUCTION IS PROHIBITED WITHOUT THE PRIOR EXPRESS WRITTEN PERMISSION OF SUN MICROSYSTEMS, INC.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

This distribution may include materials developed by third parties. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Sun ONE, the Sun ONE logo, Sun Fire, AnswerBook2, docs.sun.com, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Products covered by and information contained in this manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical/ biological weapons, or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo, or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists, is strictly prohibited. Use of any spare or replacement CPUs is limited to repair or one-for-one replacement of CPUs in products exported in compliance with U.S. export laws. Use of CPUs as product upgrades, unless authorized by the U.S. Government, is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

CE PRODUIT CONTIENT DES INFORMATIONS CONFIDENTIELLES ET DES SECRETS COMMERCIAUX DE SUN MICROSYSTEMS, INC. SON UTILISATION, SA DIVULGATION ET SA REPRODUCTION SONT INTERDITES SANS AUTORISATION EXPRESSE, ECRITE ET PREALABLE DE SUN MICROSYSTEMS, INC.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Cette distribution peut comprendre des composants développés par des tierces parties. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Sun ONE, the Sun ONE logo, Sun Fire, AnswerBook2, docs.sun.com, Java, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

Ce produit est soumis à la législation américaine en matière de contrôle des exportations et peut être soumis à la réglementation en vigueur dans d'autres pays dans le domaine des exportations et importations. Les utilisations, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou reexportations vers les pays sous embargo américain, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exhaustive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites. L'utilisation de pièces détachées ou d'unités centrales de remplacement est limitée aux réparations ou à l'échange standard d'unités centrales pour les produits exportés, conformément à la législation américaine en matière d'exportation. Sauf autorisation par les autorités des Etats-Unis, l'utilisation d'unités centrales pour procéder à des mises à jour de produits est rigoureusement interdite.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Contents

| | |
|--|----|
| Sun Fire V60x Compute Grid Rack System Documentation Set Summary | 2 |
| Software Recovery Procedures | 3 |
| Recovering Red Hat Enterprise Linux 2.1 | 3 |
| Reinstalling the Operating System Software | 3 |
| Reconfiguring the Operating System Software | 10 |
| Cluster Grid Manager Software Recovery | 11 |
| Reinstalling Sun Control Station 2.0 Software | 11 |
| Reconfiguring the Java Plug-In Version For Mozilla | 12 |
| Installing the SCS Grid Engine Module | 14 |
| Installing Custom Scripts For Advanced Users | 15 |
| Using Scripts to Automate Installation Tasks | 16 |
| Using Scripts to Recreate a Lost <code>check.out</code> File | 16 |
| Using Scripts to Auto-Populate the AllStart Clients List | 17 |
| Using Scripts to Force All Nodes to Network Boot | 18 |
| Using Scripts to Add All Nodes as SCS Managed Hosts | 19 |
| Sun ONE Grid Engine Notes | 20 |
| AllStart Client Host Name Limitations | 20 |
| Grid Engine Settings | 20 |
| Grid Engine Configuration | 21 |

| | |
|--|----|
| Important Notes | 22 |
| Location of Kickstart Files For AllStart Clients | 22 |
| PXE Network Booting Conflict With LAN Management | 22 |
| Supported Browsers and Plug-Ins | 23 |

Sun Fire V60x Compute Grid Rack System Release Notes

These release notes supplement the information in the *Sun Fire V60x Compute Grid Rack System Installation Guide* (817-3072). The information is organized into the following sections:

- “Sun Fire V60x Compute Grid Rack System Documentation Set Summary” on page 2
- “Software Recovery Procedures” on page 3
- “Using Scripts to Automate Installation Tasks” on page 16
- “Sun ONE Grid Engine Notes” on page 20
- “Important Notes” on page 22
- “Supported Browsers and Plug-Ins” on page 23

Sun Fire V60x Compute Grid Rack System Documentation Set Summary

In addition to the documents created for the installation of the Sun Fire V60x Compute Grid system, many other documents are provided to supplement the information and to provide detailed information about system components after installation. This section provides a summary of the document set.

For a full list of the documents shipped with the system, see “Related Documentation” in the Preface of the *Sun Fire V60x Compute Grid Rack System Installation Guide* (817-3072).

- Sun Fire V60x Compute Grid Rack System documents:
Use these documents for initial system installation, power-on, and software setup.
These two documents are shipped in hard copy with your system.
 - *Sun Fire V60x Compute Grid Rack System Installation Guide* (817-3072)
 - *Sun Fire V60x Compute Grid Rack System Release Notes* (817-3074)
- Sun Rack 900 documents:
Use these documents for more detailed information about the system rack. Tasks covered include unpacking, installation, and service of replaceable parts.
The Sun Rack documents are shipped in hard-copy with the system.
- Sun Fire V60x documents:
Use these documents for more detailed information about the Sun Fire V60x servers that are used as the system nodes. Tasks covered include removal and installation from the slide-rails in the rack, and detailed diagnostics and service of replaceable parts.
The Sun Fire V60x documents are shipped in a set that is a combination of hard-copy documents and PDF documents that are on the Sun Fire V60x and Sun Fire V65x Servers Documentation CD.
- Sun Control Station and Sun ONE Grid Engine, Enterprise Edition documents:
Use these documents for detailed information about setting up and using the software after the initial installation. Tasks covered include advanced settings for using the software to monitor and run the system as a grid.
The Sun Control Station and Sun ONE Grid Engine, Enterprise Edition documents are include as PDF documents that are integrated into the Help system of the Sun Control Station software. They are also shipped as PDF files on their respective CDs.
- The documents for the third-party hardware components are also included and vary, based on the supplier that is used.
Documents for replacing and using the network switches, terminal server, and keyboard unit are shipped in hard-copy and as PDF files on included CDs.

Software Recovery Procedures

Your Sun Fire V60x Compute Grid is shipped with the Red Hat Enterprise Linux 2.1 operating system and the Cluster Grid Manager software suite preinstalled on the Cluster Grid Manager (CGM) node. This section contains the procedures for recovering or reinstalling the Red Hat Enterprise Linux 2.1 operating system software and the Cluster Grid Manager software suite to the CGM node in the case that you have to replace a CGM node or reinstall the software for any reason.

- “Recovering Red Hat Enterprise Linux 2.1” on page 3
- “Cluster Grid Manager Software Recovery” on page 11

Recovering Red Hat Enterprise Linux 2.1

Use this procedure if you need to reinstall the Red Hat Enterprise Linux 2.1 distribution that was preinstalled on your CGM node. This section is divided into two procedures:

- “Reinstalling the Operating System Software” on page 3
- “Reconfiguring the Operating System Software” on page 10

Reinstalling the Operating System Software

Use this procedure to reinstall the Linux operating system software.

CDs Required For This procedure:

- Red Hat Enterprise Linux 2.1 distribution CDs
- Sun Fire V60x and Sun Fire V65x Servers Resource CD (705-0601)

- 1. Insert the Red Hat Enterprise Linux 2.1 CD 1 into the CGM node and wait for the first Red Hat installation screen to appear, then press Enter.**
- 2. At the Language Selection screen, select the language for your location, then click Next.**

The default setting is English.

3. **At the Keyboard Configuration screen, accept the default settings shown below, then click Next.**

The default settings are:

- Model: Generic 105-key international PC
- Layout: US English
- Dead keys: Enable dead keys

4. **At the Mouse Configuration screen, select Generic 3-button mouse (PS/2), then click Next.**

The default setting is generic 3-button mouse (PS/2).

5. **At the Welcome to Red Hat Linux screen, click Next.**

6. **At the Installation Type screen, choose Custom Installation Type, then click Next.**

The Disk Partitioning Setup screen appears.

7. **Create five RAID 1 partitions on each of the two hard drives, as follows:**

- a. **At the Disk Partitioning Setup screen, select Manually Partition With Disk Druid, then click Next.**

The Disk Setup screen appears.

- b. **At the Disk Setup screen, click New to begin creating a new partition.**

A New Partition dialog box appears.

- c. **In the New Partition dialog box, select hard drive `sda` from the list of Allowable Drives to create partitions on that drive first.**

- d. **In the New Partition dialog box, select Software RAID from the Filesystem Type pull-down menu.**

- e. **In the New Partition dialog box, define one of the five Software RAID partitions listed in TABLE 1.**

Note – Make the `/boot` partition your primary partition by selecting the box labeled, “Force to be primary partition.”

Note – You cannot enter the mount point for a partition until after you create the RAID 1 device in a later step.

TABLE 1 RAID 1 Partition Settings For System Recovery

| Mount Point | File System Type | RAID level | Partition Size (Mb) |
|-------------|------------------|------------|---------------------|
| / | ext3 | RAID 1 | 10000 |
| swap | swap | RAID 1 | 2000 |
| /boot | ext3 | RAID 1 | 64 |
| /var | ext3 | RAID 1 | 2000 |
| /scs | ext3 | RAID 1 | 20000 |

f. After you have defined the partition, click OK.

You are returned to the Disk Setup screen, where your new partition is listed.

g. Repeat Step b through Step f until you have created all five partitions in TABLE 1 on hard drive `sda`, then continue with Step h.

h. Create the same five partitions on hard drive `sdb` so that it will mirror hard drive `sda`.

Repeat Step b through Step f until you have defined the five partitions in TABLE 1 on hard drive `sdb`, then continue with Step i.

You are returned to the Disk Setup screen, where the 10 partitions you created are listed (5 partitions on hard drive `sda` and 5 partitions on hard drive `sdb`).

i. At the Disk Setup screen, click Make RAID.

A dialog box appears where you can select available partitions to make RAID.

j. In the dialog box, select a partition and edit the settings for the selected partition as follows, then click OK:

- Define the mount point for the partition. Refer to TABLE 1 and ensure that the mount point corresponds to the partition size that you defined earlier.

Note – There is no mount point for the `swap` partition.

- Select RAID 1 for the RAID Level.
- Select the Format Partition box.

k. Repeat Step j until you have defined the mount point and RAID Level for all 10 of the partitions.

8. After you have defined all of your partition settings, click Next to close the Disk Druid Disk Setup screen.

The Bootloader Configuration screen appears.

9. At the Bootloader Configuration screen, select LILO as the bootloader, then click Next.
10. At the Firewall Configuration screen, select No Firewall, then click Next.
11. At the Additional Language Support Selection screen, click Next.
12. At the Time Zone Selection screen, select the correct time zone for your locale, then click Next.
13. At the Account Configuration screen, type the root password, then click Next.
14. At the Authentication Configuration screen, click Next.
15. At the Package Group Selection screen, select the following group options, then click Next:
 - Printing Support
 - X Window System
 - GNOME
 - Network Support
 - Messaging and Web Tools
 - NFS File Server
 - SQL Database Server
 - Web Server
 - Network Managed Workstation
 - Emacs
 - Software Development
 - Kernel Development
 - Everything

The Video Card Configuration screen appears.

16. At the Graphical Interface (X) Configuration screen, make the following selections, then click Next.
 - ATI Mach64
 - Video Card RAM: 16 Mb

17. When the prompt that says About to Install appears, click Next.

The installation takes several minutes as the packages are installed and the partitions are formatted.

18. When you are prompted for the next CD in the Linux distribution, remove the current CD and replace it with the next CD.

When the installation is complete, the Boot Disk Creation screen appears.

19. At the Boot Disk Creation screen, select Skip Boot Disk Creation, then click Next.
20. At the Monitor Configuration screen, accept the default, then click Next.

If you are using a different monitor than the one in the KVM unit, select your monitor type rather than accepting the default.
21. At the Custom Graphics Configuration screen, make the following selections, then click Next.
 - Color depth: High-color 16-bit
 - Screen resolution: 1024x768
 - Desktop environment: Gnome
 - Login type: Graphical
22. At the screen that says, “Congratulations, Your installation is now complete,” click Exit.

The node reboots automatically.
23. After the system returns to a Red Hat login screen, log in as the root user.
24. Install the required E1000 network drivers and configure the Ethernet device as follows:
 - a. Insert the Sun Fire V60x and Sun Fire V65x Server Resource CD (shipped with your system) into the CGM node and mount the CD by typing the following command.

```
# mount /dev/cdrom /mnt/cdrom
```
 - b. Copy the required network drivers from the Resource CD and install them to the CGM node by typing the following commands:

```
# cd /mnt/cdrom/drivers/src
# cp e1000-4.4.19.tar.gz /root
# cd /root
# tar -zxvf e1000-4.4.19.tar
# cd e1000-4.4.19/src
# make install
# insmod e1000
```
 - c. Remove the Resource CD from the system after you type the following command:

```
# umount /dev/cdrom
```
 - d. Reboot the system by typing the following command:

```
# reboot
```
 - e. After the system returns to a Red Hat login screen, log in as the root user.

- f. Verify that the e1000 network drivers were installed by typing the following commands and looking for the e1000 entry for eth1 in the /etc/modules.conf file.**

```
# cd /etc
# more /etc/modules.conf
```

Sample file contents are shown here:

```
alias parport_lowlevel parport_pc
alias scsi_hostadapter aic79xx
alias eth0 e1000_4412k1
alias usb-controller usb-uhci
alias eth1 e1000
```

g.

- h. From the Gnome desktop menu bar, select Program > System > Internet Configuration Wizard.**

- i. In the Add New Device Type dialog box, select your Ethernet connection, then use the wizard to configure the Ethernet device and to activate it.**

Consult with your system administrator to select settings that are compatible with your network. The factory-default IP address of the CGM node is 192.168.160.5.

- j. Restart the network service by typing the following the command:**

```
# service network restart
```

- k. Verify that the system sees the Ethernet device by typing the following command.**

```
# ifconfig -a
```

25. Download and install the required Adaptec SCSI driver as follows:

- a. On the CGM node, use a browser to go to the Sun Fire V60x download site:**

<http://www.sun.com/servers/entry/v60x/downloads.html>

- b. Navigate to the Device Drivers download links for Red Hat Enterprise Linux 2.1 software.**

- c. Download the following tar file to a /tmp directory on the CGM node:**

Adaptec SCSI Driver RPMs 1.3.10 for Red Hat Enterprise Linux 2.1
(as-aic79xx.tar.gz)

- d. Extract the contents of the tar file into the /tmp directory by typing the following commands:**

```
# cd /tmp
# tar -zxf /tmp/as-aic79xx.tar.gz
```

- e. Determine which kernel version is running on your system by typing the following commands:

```
# uname -a | awk '{print $3}'
```

The kernel version on your system is displayed similar to the following example:

```
2.4.9-e.12smp
```

- f. Locate the correct drivers for your kernel version in the `as-aic79xx` folder by typing the following commands:

```
# cd as-aic79xx/  
# ls *kernel-version*
```

Where *kernel-version* is the kernel version you determined in Step e. Using the example in the previous step, the command and response would look as follows:

```
# ls *e.12*  
aic79xx-1.3.10_2.4.9_e.12-rh21as_1.i686.rpm  
aic79xx-1.3.10_2.4.9_e.12-rh21as_1.src.rpm  
aic79xx-enterprise-1.3.10_2.4.9_e.12-rh21as_1.i686.rpm  
aic79xx-smp-1.3.10_2.4.9_e.12-rh21as_1.i686.rpm
```

- g. Install the required SCSI drivers by typing the following commands:

```
# rpm -ivh driver-version
```

Where *driver-version* is the driver that you determined in Step f. Using the example in the previous step, the commands would look as follows:

```
# rpm -ivh aic79xx-1.3.10_2.4.9_e.12-rh21as_1.i686.rpm  
# rpm -ivh aic79xx-smp-1.3.10_2.4.9_e.12-rh21as_1.i686.rpm
```

Note – The two required SCSI drivers are the smp/i686 driver and the uniprocessor i686 driver (non-enterprise), as shown in the previous example.

- h. Inform the boot loader where to find the new initial ramdisk (initrd) image by typing the following commands:

```
# lilo  
# reboot
```

26. Continue with “Reconfiguring the Operating System Software” on page 10.

Reconfiguring the Operating System Software

Use this procedure to reconfigure the Linux operating system after you reinstall it.

1. Enable serial redirection on the CGM node as follows:

- a. **Modify the CGM nodes's /etc/lilo.conf file to add the following line after the lines that read, read-only:**

```
append="console=tty0 console=ttyS1,9600"
```

This change enables serial redirection of the output from the LILO boot loader and the early boot process.

- b. **Modify the CGM nodes's /etc/inittab file to add the following line after the line that reads, 6:2345:respawn:/sbin/mingetty tty6:**

```
7:2345:respawn:/sbin/mingetty ttyS1
```

- c. **Modify the CGM nodes's /etc/securetty file to add the following line at the end of the file:**

```
ttyS1
```

- d. **Reboot the server to enable the serial redirection settings.**

2. Configure the X windows environment on the CGM node as follows:

- a. **At a Linux command line, log in as the root user.**
- b. **Start the Red Hat Linux configuration utility by typing the following command:**

```
# setup
```

- c. **Select X Configuration from the menu of setup selections.**
- d. **Accept all default X configuration options, except for the following changes you must make:**
 - Video card memory = 8MB
 - Color resolution = 24-bit, 1024x768

After you make these configuration changes, you can start the X windows environment by typing the **startx** command at a Linux command line.

Note – You might not be able to resize the X windows because of a Red Hat bug. You can work around this bug by performing the following steps:

- i) Click on MainMenu on the toolbar at the bottom of the screen.
 - ii) Select Programs > Setting > Sawfish Window Manager > Moving and Resizing
 - iii) Deselect the box labeled, "Show current dimensions of window while resizing."
 - iv) Click Apply.
 - v) Click OK.
-

3. Continue with "Cluster Grid Manager Software Recovery" on page 11.

Cluster Grid Manager Software Recovery

Use this procedure to reinstall the Cluster Grid Manager software suite that was preinstalled on your CGM node. This section is organized into the following procedures that should be performed in the order they are listed here:

- “Reinstalling Sun Control Station 2.0 Software” on page 11
- “Reconfiguring the Java Plug-In Version For Mozilla” on page 12
- “Installing the SCS Grid Engine Module” on page 14
- “Installing Custom Scripts For Advanced Users” on page 15

CD Required For These Procedures

- Cluster Grid Manager Software Recovery CD (798-4973)

Note – You must install the operating system before performing this procedure, as described in “Recovering Red Hat Enterprise Linux 2.1” on page 3.

Reinstalling Sun Control Station 2.0 Software

Use this procedure to reinstall the Sun Control Station (SCS) software.

- 1. Insert the Cluster Grid Manager Software Recovery CD into your CGM node.**

If the CD does not mount automatically, mount it by typing the following commands:

```
# mount /dev/cdrom /mnt/cdrom
# cd /mnt/cdrom
```

- 2. Copy the SCS tar file from the CD to the `/scs` directory on your CGM node by typing the following command:**

```
# cp scs-2.0-release.tgz /scs
```

This file is approximately 370 Mb, so the copying might take several minutes.

- 3. After the copy operation has finished, type the following commands to install the new SCS software:**

```
# cd /scs
# tar -zxvf scs-2.0-release.tgz
# cd scs-2.0/install
# ./install -factoryinstall
```

The installation might take several minutes.

4. Install the SCS patch as follows:

- a. Copy the SCS patch from the CD to the root directory on the CGM node by typing the following command:**

```
# cp scs-2.0p1.tgz /root
```

- b. After copy operation finishes, extract the tar file by typing the following commands:**

```
# cd /root
# tar -zxvf scs-2.0p1.tgz
```

- c. Install the SCS patch by typing the following commands:**

```
# cd scs-2.0p1
# ./install/install
```

- d. Delete the patch tar file after the installation finishes by typing the following commands:**

```
# cd ..
# cd /root
# rm scs-2.0p1.tgz
```

- e. Reboot the CGM node to initialize the SCS database by typing the following command:**

```
# reboot
```

5. Continue with “Reconfiguring the Java Plug-In Version For Mozilla” on page 12.

Reconfiguring the Java Plug-In Version For Mozilla

The Java™ plug-in for the Mozilla™ browser that is supplied with the RedHat Linux software is not compatible and it must be replaced by the Java plug-in supplied with the SCS software. Use the following procedure to reconfigure the Java plug-in version.

1. Configure the Java plug-in version by typing the following commands:

```
# cd /usr/lib/mozilla/plugins
# rm libjavaplugin_oji.so
# ln -s \
/usr/java/j2sdk1.4.1_02/jre/plugin/i386/ns610/libjavaplugin_oji.
so
```

2. Verify that the new Java plug-in version was configured as follows:

- a. Close all Mozilla applications.**
- b. Start a Mozilla browser.**

- c. At the top of the Mozilla window, click on Help > About Plug-ins.**
 - d. Verify that the following version of the Java plug-in is listed:**
Java™ Plug-in1.4.1_02-b06
- 3. Continue with “Installing the SCS Grid Engine Module” on page 14.**

Installing the SCS Grid Engine Module

After you install the SCS software, you must install the Grid Engine module separately.

1. **Start a browser and type the following URL.**

`http://n.n.n.n`

Where *n.n.n.n* is the IP address that you assigned to the CGM node.

2. **Read the Sun Control Station license agreement that appears and accept the license agreement if you agree with the terms.**

A Sun Control Station Welcome page appears.

3. **Go to the Sun Control Station login page for your CGM node by entering the URL in the format that is shown on the Welcome page:**

`https://n.n.n.n:8443/sdui`

Where *n.n.n.n* is the IP address that you assigned to the CGM node.

Note – The URL uses the `https` format.

4. **At the Sun Control Station login page, log in as the SCS administrator using the default entries shown below, then click the Login button.**

User Name: **admin**

Password: **admin**

5. **On the Cluster Grid Manager main page, click on Administration > Modules in the left-side panel.**

The Control Modules window appears.

6. **On the Control Modules window, click on Add Module.**

The Add Module window appears.

7. **Select Location as File and browse to the Grid Engine module file on the Cluster Grid Manager Software Recovery CD:**

`/mnt/cdrom/gridModule-1.0-14.mapp`

8. **Click on Install Now.**

Accept any security certificates or warnings that appear.

Note – You might have to log in to SCS again after you install the Grid Engine module to see the Grid Engine module selection in the menu.

9. **Continue with “Installing Custom Scripts For Advanced Users” on page 15.**

Installing Custom Scripts For Advanced Users

Several useful scripts are included on the Cluster Grid manager Software Recovery CD. Use the following procedure to install the scripts to your CGM node.

1. **Create a `/usr/mgmt/diag` directory on your CGM node by typing the following command:**

```
# mkdir /usr/mgmt/diag
```

2. **Copy and extract the scripts tar file from the Cluster Grid manager Software Recovery CD to your CGM node by typing the following commands:**

```
# cp /mnt/cdrom/customerdiag1.2.tar /usr/mgmt/diag
# cd /usr/mgmt/diag
# tar -xvf customerdiag1.2.tar
```

3. **Remove the Cluster Grid manager Software Recovery CD from your CGM node after you type the following command:**

```
# umount /dev/cdrom
```

Note – After you install the custom scripts, you can use them to automate several of the more redundant SCS AllStart software deployment activities. The procedures in the following sections describe how to use the scripts, but they are optional and recommended for advanced users.

Using Scripts to Automate Installation Tasks

This section contains procedures that describe how to use scripts that are included on the Cluster Grid Manager Software Recovery CD. These scripts can be used to automate some of the redundant tasks that are required when using the SCS AllStart module to deploy software to the compute nodes.

You should first review the basic AllStart module procedures in the *Sun Fire V60x Compute Grid Rack System Installation Guide* (817-3072) before you use these procedures.

Note – You must first install the scripts, as described in “Installing Custom Scripts For Advanced Users” on page 15.

Using Scripts to Recreate a Lost `check.out` File

When your system is manufactured, a file named `check.out` is created on the CGM node that lists the MAC addresses for all the nodes in your system. If this file is lost for any reason, you can use one of the custom scripts as described in this procedure to recreate the `check.out` file.

1. Type the following commands to run the script:

```
# cd /usr/mgmt/diag
# ./config -c n.n.n.n check TS-port-numbers
```

Where `n.n.n.n` is the IP address of the system's terminal server and *TS-port-numbers* is a range or list of terminal server ports to which compute nodes are connected. For example, 1-32 would denote the range for a fully configured, 32-node system. If your system is not fully configured, your *TS-port-numbers* value might look like 1,2,4,6-16.

2. Reset each compute node by pressing the Reset button on each node's front panel.

As each node resets, it will provide output of MAC addresses to a file named `/usr/mgmt/diag/customer_check.out`.

Using Scripts to Auto-Populate the AllStart Clients List

Perform the following procedure to use the MAC addresses from your `customer_check.out` file to auto-populate the AllStart Clients list.

Note – Use this procedure after you have already used AllStart to create your Distribution, Payload, and Profile, as described in the Sun Fire V60x Compute Grid Rack System Installation Guide (817-3072).

1. Type the following commands to run the script:

```
# cd /usr/mgmt/diag
# ./as_mac.pl -i NODE_BASE_IP -f customer_check.out
```

Where `NODE_BASE_IP` is the base, or starting IP address for your node range. All nodes have their IP addresses incremented by one, following this lowest IP address.

The script uses the MAC addresses in the `customer_check.out` file to populate the AllStart Clients list. Allstart adds clients, starting with `NODE_BASE_IP`, for each MAC address in the `customer_check.out` file, up to, but not including, the CGM node.

2. Verify that the clients were added by looking at the AllStart Clients list. At the Cluster Grid Manager main window, click on AllStart > Clients.

All of the new clients should be listed, although they have no payload or profile associated with them yet.

3. Modify the clients that you just created to associate them with the AllStart deployment you are creating, as follows:

a. On the AllStart Clients window, click **Select All**.

b. Click **Modify**.

c. **Modify the settings for the clients as described in the Sun Fire V60x Compute Grid Rack System Installation Guide.**

When you finish making the settings, you are returned to the AllStart Clients window.

Note – Be sure to set up the clients settings so that they are associated with the settings for the distribution, payload, and profiles that you have already created for this AllStart deployment.

d. In the AllStart Clients window, click **Select All**.

e. Click Enable.

All client entries are enabled so that they are visible to the system. Enabled clients are indicated by a Y character under the Enabled heading on the AllStart Clients window.

4. Modify the DHCP configurations for the clients as follows:

a. On the Cluster Grid Manager main window, click AllStart > Service.

The AllStart Current Service Settings window appears.

b. Click Modify.

The Modify Service Settings window appears.

c. Verify that the DHCP Enabled box is selected.

d. Click Modify DHCP Info.

e. Select the DHCP subnet and click Edit.

f. Enter the router and DNS server IP addresses for your servers. Do not add anything to the Network/netmask or IP Range fields.

Using Scripts to Force All Nodes to Network Boot

Use the following procedure to force all nodes to network boot, as required when you are deploying software to compute nodes.

1. Type the following commands to run the script:

```
# cd /usr/mgmt/diag
# ./config -c n.n.n.n pxe TS-port-numbers
```

Note – You must ensure that none of the ports given in the *TS-port-numbers* node range are currently in use when you use the command to run this script. The script must have access to the serial ports of each node to take control of the nodes.

Where *n.n.n.n* is the IP address of the system's terminal server and *TS-port-numbers* is a range or list of terminal server ports to which compute nodes are connected. For example, 1-32 would denote the range for a fully configured, 32-node system. If your system is not fully configured, your *TS-port-numbers* value might look like 1,2,4,6-16.

2. Reset or power on the nodes by pressing the Reset or Power buttons on the front panel of each node.

The script causes each node to network boot and pull the software deployment from the CGM node.

Using Scripts to Add All Nodes as SCS Managed Hosts

Before you can deploy the Sun ONE Grid Engine, Enterprise Edition software to the system compute nodes so that they can be managed as a grid, you must first add the nodes as Sun Control Station managed hosts. Perform this procedure to use a script to add all nodes as SCS managed hosts.

1. Type the following command to create a file named **nodelist**, which contains the list of Allstart clients that will be added as SCS managed hosts:

```
# cd /usr/mgmt/diag
# ./createNodeList.pl > nodelist
```

2. Type the following command to run the script that adds the clients as SCS managed hosts and installs the SCS client daemons on them:

Note – Do not run this script in the background. Monitor the progress of the activity by watching the SCS Administration > Hosts window. Refresh the window until all the added hosts appear in the list.

```
# ./devMgrParallel.pl add file nodelist
```

Sun ONE Grid Engine Notes

This section contains information about the Sun ONE Grid Engine, Enterprise Edition (S1GEEE) v 5.3p4 software that is preinstalled on your system's CGM node.

AllStart Client Host Name Limitations

When you use the SCS AllStart module to create the client nodes to which you will deploy the software payloads, you are required to enter network interface information for those clients. In the AllStart Clients > Enter Network Interface Information window, you must enter the host name for the client node that you are creating. (See the *Sun Fire V60x Compute Grid Rack System Installation Guide* for the full procedure.)

When entering the host name, you *cannot* use the full host name format, which would include the domain name. Instead, you must use a short host name format. For example:

Use this host name format: `host1`

Do not use this format: `host1.mydomain.com`

If you use the full host name format, the S1GEEE software cannot resolve the host name properly and the host (client node) is not able to join the grid or act as the grid master host.

Grid Engine Settings

When the grid engine is deployed, the following settings are automatically used:

- SGE admin user name: `sgeadmin`
- SGE admin user ID: 268
- SGE commd TCP port: 536
- SGE root directory: `/gridware/sge`
- SGE cell name: default
- NFS sharing: the `/gridware/sge/default/common` directory is NFS-shared from the designated master host to all other hosts. All other files and directories, such as binaries and spool directories, are installed locally on all grid hosts.

Grid Engine Configuration

When you configure a compute host, one default queue is created for it. The queue settings are the same as that for a regular (standalone) S1GEEE deployment, with the exception of the following:

In the Sun Fire V60x Compute Grid environment, the `rerunnable` parameter is set to "y". In other words, jobs running in the queue can be restarted on other compute hosts of the system in certain circumstances; for example, when a compute host is being removed from the grid.

After you have deployed the grid engine, you can modify the configuration parameters on the queues that were automatically set up to anything you want, or even delete the queues entirely.

For details on grid engine settings, refer to the *Sun ONE Grid Engine, Enterprise Edition 5.3 Administration and User's Guide*. This document is accessible through the Help interface of the SCS software, or at the following URL:

<http://www.sun.com/products-n-solutions/hardware/docs/Software/S1GEEE/index.html>

Important Notes

This section contains information about known issues and considerations regarding the system and its operation.

Location of Kickstart Files For AllStart Clients

You can verify that your AllStart clients have been correctly configured by checking for their listing in the `/scs/allstart/ksconfig/` directory on your CGM node.

Each compute node that has been configured as an AllStart client is identifiable by its MAC address, as listed in the following Kickstart file format:

```
/scs/allstart/ksconfig/ks.MAC-address.cfg
```

PXE Network Booting Conflict With LAN Management

When you use the Sun Control Station AllStart module to deploy software from the CGM node to the compute nodes, you force the target compute node to network (PXE) boot and pull the software from the CGM node.

The PXE boot process involves UDP network transactions. If the DHCP/PXE server tries to assign an IP address that already belongs to another locally networked node that has LAN management enabled, the PXE boot process might fail. Note that even if the bootloader appears successful, the LAN-managed node might still have corrupted the transaction.

If you encounter this problem, there are several solutions:

- Deactivate LAN management software on the local network before you perform the PXE booting process.
- Assign different IP addresses to the LAN-managed nodes that are in conflict with the DHCP/PXE server's assignments.
- Do not use the client node's LAN-managed Ethernet interface; use the other Ethernet interface. Note that this would require you to reconfigure any Kickstart information to use the new interface.

Supported Browsers and Plug-Ins

For viewing Sun Control Station 2.0 software, the following browsers and plug-ins have been tested and are officially supported on the indicated operating system platforms at this time.

- Linux software (tested on Red hat Linux 7.3 and Red Hat Enterprise Linux 2.1):
 - Browser: Mozilla 1.2.1
 - Java™ Plug-in: Java Plug-in 1.4.1
- Solaris 8 and Solaris 9 operating system software:
 - Browser: Netscape 7.0
 - Java Plug-in: Java Plug-in 1.4.1
- Windows NT 4.0:
 - Browser: Internet Explorer 6
 - Java Plug-in: Java Plug-in 1.4.1

