Implementation Guide:
Guide to Installation—Part I: Sun™ Cluster Management Services

Chris Dotson, Enterprise Engineering
Steve Lopez, Enterprise Engineering

Sun BluePrints™ OnLine—April 2002

http://www.sun.com/blueprints

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 USA
650 960-1300 fax 650 969-9131
Part No.: 816-4684-10
Revision 1.0
Edition: April 2002

Please
Recycle

Module 2

Part I: Sun™ Cluster Management Services

The purpose of this module is to guide you through the tasks you must perform before you install the Sun™ Cluster 3.0 software. These tasks include setting up the administrative workstation and configuring the Sun Cluster 3.0 hardware components.

The exercises in this module explain how to install and configure a workstation to perform Sun Cluster 3.0 software administrative functions in a cluster environment. Additionally, we provide instructions for configuring the cluster, implementing best practices, and performing design verifications, as well as administering a two-node Sun Cluster 3.0 hardware cluster.

For information about managing a cluster, refer to the documents referenced in Appendix B.

Many of the steps in this guide refer to “manual (local) procedures,” which you should perform only if you have local (physical) access to the SunPlex™ platform. For example, resetting the terminal concentrator in order to activate specific settings.

Objectives

After completing this module, you will have successfully verified the installation and software configuration for each cluster component. These tasks must be performed before you can install the Sun Cluster 3.0 software on each cluster node. Additionally, you will have implemented the associated key practices during each task, including:

- Configuring the Solaris™ Operating Environment (Solaris OE). This task is site-specific.
- Verifying the administrative workstation setup on the management server.
- Verifying the terminal concentrator (TC) configuration.
- Installing and configuring the Cluster Console utility in the management server environment.
- Verifying each cluster node installs with the Solaris OE and patches.
- Configuring additional cluster management services, such as JumpStart™ software for each cluster node, plus additional platform management services. This task is performed only during manual installations.

Many of the steps for configuring the cluster require you to have physical access to the equipment. We have made special note of specific configuration procedures that require local (physical) access to the SunPlex platform. For example, forcing the TC into Monitor mode requires local, manual procedures. For this reason, these hands-on exercises specify steps that are only performed as part of local (or manual) installation procedures. These manual steps are not performed if you are accessing the Sun Cluster 3.0 hardware (SunPlex platform) remotely. Instead, these steps are included as documentation of the procedures required to complete the configuration of each cluster component (as referenced in Figure 1-1 of “Guide to Installation - Hardware Setup”). Additionally, it is assumed that it is not necessary to provide the detailed procedures for installing the Solaris OE (plus Solaris OE patches) on each node. Instead, we define the parameters required for installation.

Note – In this module, you must perform all of the associated instructions to verify that the correct configuration has been achieved as specified. It is very important that you do this prior to configuring the additional cluster components to ensure consistent, reliable cluster services and SunPlex platform operations.

Prerequisites

As stated in the introduction, this lab guide is intended for use by qualified network administrators.

For installation queries, shell usage questions, patches, and packages, refer to the Sun Educational Services manuals for the “Solaris System Administration 1” and “Solaris System Administration 2” courses.

Required Reading

This module builds on concepts presented in the Cluster Column published in January, 2002, “Guide to Installation—Hardware Setup.” Before you read this module, familiarize yourself with the content in that publication, paying special attention to the hardware configuration and components described in Figure 1-1, and Tables 1-1 through 1-5.

Introduction

The administrative workstation (`clustadm`) is required for setting up and configuring the cluster, and can be used for ongoing cluster management and operations.

The administrative workstation is a remote workstation, which can be used to administer each cluster node by connecting through the terminal concentrator, and can be used to administer and monitor the cluster through the cluster control panel, the command-line interface, or through the Sun™ Management Center software console (monitoring only).

Note – Beyond these administrative workstation functions, we recommend that you implement additional cluster management services. When configuring this type of management server, consider combining JumpStart software services, applications software (and /PATCH) repository, along with Sun Cluster 3.0 software administrative functions appropriate to your implementation.

In this module, we describe the procedures used to install the administrative workstation (`clustadm`). We explain how to confirm that all requirements are met, verify that all cluster components (including patches) are installed correctly, and ensure that the student-selected shell environment is configured.

Management Server Functions

Management servers can be used to perform the following functions:

- Management servers can be designated for distributing (saving) valuable cluster node (site-specific) configuration files.
- Consider, also, configuring the Solaris OE `syslog` facility to copy messages from each cluster node to a remote administrative workstation (or another designated host) for the purpose of notifying administrators of potential cluster disruptions (logs, which should be saved / reviewed, as necessary).
- A primary function demonstrated in these hands-on labs is the ability of the administrative workstation to run the cluster control panel (CCP), accessing the Cluster Console window of each cluster node in the SunPlex platform. As instructed throughout this guide, during the installation (and when monitoring cluster operation), commands will be entered simultaneously to each cluster node using a Cluster Console window (one window for each node), accessed from the CCP Hosts menu.
- Cluster management services should provide all standard administrative workstation functions and can provide additional platform management functions.

Note – These hands-on labs demonstrate remote cluster operations (for example, with CCP operations) including configuring and monitoring cluster and status, managing changes within the cluster (for example, shutting down individual cluster nodes on the SunPlex platform prior to performing upgrades), rebooting a cluster node, and ensuring that auto-cluster formation occurs.

- Time synchronization throughout the SunPlex platform.
- Solaris OE installation services using JumpStart software.
- Applications servers using VERITAS Volume Manager software (VxVM,) Solstice DiskSuite™ software, Sun Cluster software, and associated utilities).
- Online repository for applications (binaries) and patches (`/PATCH`).
- Online repository for site-specific data and configuration files.

For more information about cluster administration functions (for example, administering cluster interconnects and public networks, patching Sun Cluster software and firmware, and backing up and restoring a cluster), refer to the documentation referenced in Appendix B.

Optionally, the management server can be configured to act as a Sun Cluster 3.0 software-installation server, and can help simplify software installations for each cluster node on the SunPlex platform. For example, you can install the Sun Cluster 3.0 software packages on the management server and populate a custom JumpStart software directory.

Key Practice: Use the JumpStart technology to maintain consistency and fully automate the installation of the Solaris OE and additional software packages. The JumpStart software can minimize operator errors that occur during a manual installation process. Combining JumpStart software and Flash archives (for example, with the `flar` and `flarcreate` commands) can help enable quick and consistent disaster recovery operations.

Note – For information about implementing Flash archives, reference recent publications available from Sun BluePrints™ OnLine web site at `http://www.sun.com/blueprints/`, as well as the documentation referenced in Appendix B.

Not Required for Application Failover

Clusters do not require the use of the administrative workstation for normal applications processing to occur, or to achieve automatic failover to a standby node.

Section 2.1: Installing and Configuring the Management Server

This section describes how to install the Solaris 8 Operating Environment (Solaris 8 OE), plus patches, on the management server (`clustadm`).

Key Practice: Ensure that all firmware is installed with the recent versions for all systems and subsystems (for example, for servers, disk arrays, and controllers). For example, on each system (for example, on each node), ensure the system `EEPROM` contains the most current OpenBoot™ PROM version. For the examples presented in this guide, we used `OpenBoot 3.15`. Additionally, ensure that all subsystems and controllers are configured using the latest versions, as appropriate.

For each node, you can find the latest version of the OpenBoot PROM at `http://sunsolve.sun.com`.

Step 2.1.1—Installing the Solaris OE on the Administrative Workstation

For local (manual) installations, install the Solaris OE on the administrative workstation. In subsequent steps, this workstation will be configured to act as a management server (sometimes referred to as "`clustadm`").

Note – The same version of the Solaris OE runs on both the administrative workstation (`clustadm`) and on each cluster node on the SunPlex platform. When configuring the SunPLEX platform, refer to the *actual* requirements to determine the appropriate version of Sun Cluster software, plus applications, and to determine the appropriate Solaris OE version and distribution packages to be installed. Determine which, if any, real dependencies exist. For example, selecting "Entire Distribution Plus OEM" may be required if there are any actual dependencies upon third-party software.

Key Practice: Configure a flexible, consistent disk partitioning scheme. For example, configure a consistent partitioning scheme that is flexible enough to allow for the use of either Solstice DiskSuite software or VxVM software.

Implement the following standard partitioning recommendations for boot disks:

- Partition each disk spindle identically to save time and provide additional configuration flexibility by ensuring consistency across nodes.
- Reserve cylinders 0–5 (the first six cylinders) on each disk. Both volume managers require these cylinders be left unused and available. For example, Solstice DiskSuite software assigns these cylinders to slice 7 on each disk (for replica metatstate databases). If a volume manager is not configured, there should not be any significant penalties for leaving these cylinders unused.

The following partitioning guidelines can be implemented for the management server (`clustadm`) boot disk, which reserves approximately 10MB for a volume manager and allocates 1GB (slice 1) for swap space, and assigns all remaining space to "`/`" (`root`, on slice 0).

Configure boot disk slices using the following guidelines:

Slice 0 = cylinders 500 - 7461 assigned to "`/`" (*all unallocated space*; approximately 10GB)

Slice 1 = cylinders 7 - 499 assigned to "`swap`" (*approximately 1GB;*)

Slice 2 = cylinders 0 - 7505 assigned as "`backup`"(*full extent of the disk*)

Slice 7 = cylinders 1 - 6 "`unassigned`" (*reserve 10MB for use by a volume manager*)

{Reserve cylinders 1 - 6: Solstice DiskSuite software would require slice 7; VxXM requires slices 3 and 4.}

Caution – Our example assumes an 18GB disk drive with 7506 cylinders. For each configuration, always ensure the slicing information matches the actual disk geometry.

Note – The actual swap space should be sized based on the actual Solaris OE version and the requirements of the applications to be hosted. During the procedures presented in this book, we will not configure a volume manager on the administrative workstation (`clustadm`).

Step 2.1.2—Verifying the Solaris OE Configuration

To verify that your Solaris OE configuration is correct, log in as the root user on the management server (`clustadm`) and ensure that the Solaris OE has been installed, as specified in this section, and confirm that the primary boot disk (`c0t0`) is configured, per the guidelines in this section.

Note – For the examples presented in this guide, the root password is set to `abc`.

Table 2-1 Solaris OE Configuration

Hostname:	`clustadm`
IP address:	129.153.xx.xxx
Name service:	None (local files)
Set subnet:	Yes
Subnet mask:	255.255.255.0
Default gateway:	None

Note – The values quoted in the preceding table are sample values. In a live installation, substitute the appropriate site-specific values.

Key Practice: To verify that the Solaris OE installation was successful and that no errors were reported:

1. Review the `/var/sadm/README` file to determine the location of the most recent installation logs (for example, to determine the location of `/var/sadm/system/logs`).
2. Examine the current Solaris OE installation log files for potential errors (for example, examine `begin.log`, `sysidtool.log`, or `install.log`).
3. Confirm the cause of any installation error messages that have occurred, resolving failures before proceeding further.

Step 2.1.3—Tracking Installation Errors

On the administrative workstation, examine the installation logs, ensuring that any Solaris OE installation errors do not go undetected or unresolved. To do this, enter the following commands:

```
clustadm# cd /var/sadm/system/logs
clustadm# pwd
/var/sadm/system/logs

clustadm# ls
begin.log             finish.log             install.log
begin.log_2000_04_13  finish.log_2000_04_13  sysidtool.log
{{sample dates only}}
clustadm# more *
It is important to look for installation error messages in
the logs.

Example:

pkgadd: ERROR: postinstall script did not complete
successfully
```

Step 2.1.4—Obtaining the Latest Solaris OE Patches

For local (manual) installations only, obtain the latest required Solaris OE patches from either the SunSolve℠ CD-ROM, or from the SunSolve Online℠ web site at `http://sunsolve.Sun.com` (click the Patches option on the left column).

Note – The SunSolve program is a contract service from Sun Enterprise Services. We recommend that you subscribe to this service, especially if you are running a production server.

Outside North America, the method you use for obtaining the most recent patches available may deviate from the following procedure, which obtains patches through the `http://sunsolve.sun.com` web site. Ask your local Sun service provider for the best method for getting the required patch clusters for your current operating environment.

Key Practice: Follow these recommendations when obtaining and using the latest available required Solaris OE patches:

- Before you install new patches, create a /PATCHES directory on a dedicated server to store the patches. This enables centralized patch management. For example, the hands-on lab hardware has been configured with a “master” JumpStart technology server that will serve all software binaries, patches, and act as the repository.
- Refer to the individual patch README files to review any installation prerequisites before installing patches. Using this practice could prevent conflicts with other patches, software, bootprom variables, or other unknowns.
- Always install the latest Solaris OE recommended patches from SunSolve. Maintaining the latest recommended patches assures your system provides the highest reliability.

Step 2.1.5—Installing Patches on the Administrative Workstation

For local (manual) installations only, install the latest Solaris OE recommended patches on the administrative workstation. Ensure patches are successfully installed and applied.

Key Practice: Follow these recommendations when installing patches on the administrative workstation.

- Review the /var/sadm/README file to identify important log files to be examined, including the /var/sadm/install data/Solaris 2.8 Recommended log file.
- Confirm the cause of any patch installation error messages which may occur.

Step 2.1.6—Rebooting the Management Server

For local (manual) installations, reboot the management server after all patches have been successfully installed and applied.

Reboot the system after patches have been installed and applied.

Note – It is often a good idea to reboot the system after changes are made to system software and configuration. For example, at this time, reboot the management server after the patches have been installed to ensure that changes are applied and that a consistent state has been achieved.

Step 2.1.7—Verifying Shell Environment Variables

On the `clustadm` workstation, log in as the `root` user. Verify the shell environment variables are established as listed in the following table. Make these settings permanent by editing either the `/.profile` or `/.login` (for C shell users) file. Generic examples of each of these files can be found in Appendix A.

Note the following:

- Export all variables.
- Use the `/.profile` or `/.login` file to make the root user settings.
- Edit the `/etc/passwd` file to change the shell—the default shell is `/sbin/sh`.

Note – Prior to using a text editor (*vi*) to view and modify files, verify and set your terminal environment variable, as appropriate (for example, set it to **TERM=vt220**) for proper video display.

Table 2-2 Example Environment Variables

Variable	Label
TERM	`vt220`
stty	`istrip`

Prompt	<hostname># {{e.g, clustadm#}}
Add the following to the PATH variable	PATH=$PATH:/usr/bin:/usr/ucb:/etc:/sbin:/usr/sbin:/opt/SUNWcluster/bin:
Add the following to the MANPATH variable	MANPATH=$MANPATH:/usr/dt/man:/usr/man:/usr/openwin/share/man:/opt/SUNWcluster/man:

Summary of Key Practices

Use the JumpStart software to maintain consistency and to fully automate the software installations and configurations of each cluster node.

Ensure all firmware are installed with the most recent required versions for all systems and subsystems (for example, for servers, disk arrays, and controllers). For example, ensure the EEPROM contains the most current OpenBoot PROM version on each cluster node, and that disk subsystems are configured using the latest revisions, as appropriate.

Configure a flexible, consistent disk partitioning scheme.

As part of the initial OS installation preparation, and prior to configuring the system EEPROM, reset the EEPROM to the factory defaults.

Verify that the Solaris OE installation was successful and that all required packages are installed.

Create a /PATCHES directory on a dedicated server to store all patches to enable centralized patch management.

Ensure that no conflicts exist. Refer to the individual patch README files to review any installation prerequisites before installing patches.

Install the latest required Solaris OE patches.

Validate the installation of all patches, reviewing installation logs.

Reboot the system after all patches have been installed and applied.

End of Section 2.1

This completes this section. The Solaris OE has been manually installed and all patches have been applied and verified.

Section 2.2: Configuring the Terminal Concentrator

In this section, we configure the terminal concentrator (TC) for use in the cluster.

Note – Some steps for configuring the TC require you to gain physical access to the terminal concentrator and equipment (for example, forcing the TC into Monitor mode). Follow the steps in this section only when performing these steps locally, as during a manual installation.

Step 2.2.1—Preparing to Connect the Serial Cable

For local (manual) installations, prior to connecting the serial cable: Should this be an ordered list?

1. Ensure the TC power is off.
2. Connect the cable, noting the serial port (ttyb, is the default) you will be using on the administrative workstation.
3. As described in Figure 1-1 and Tables 1-1 through 1-5 of the “Guide to Installation - Hardware Setup,” a serial cable must be connected from a serial (ttya or b) port on the administrative workstation to port 1 on the TC. Port 1 (configuration port) of the TC is required for performing all local (manual) steps.

Note – The next step is not implemented for these hands-on labs. As Figure 1-1 and Tables 1-1 through 1-5 of the “Guide to Installation - Hardware Setup,” indicate, serial port “A” of the administrative workstation is connected to the TC (port 1) instead of (default) ttyb.

Step 2.2.2—Linking to the TC

For local (manual) installations, use the UNIX™ `tip` command to communicate with the TC during configuration.

Note – Before the `tip` command will work, you must ensure that the `/etc/remote` file includes the following lines appended to the end of the file.

```
annexterm:\
        :dv=/dev/term/n:br#9600:el=^C^S^Q^U^D:ie=%$:oe=^
D

{{In the above line, substitute the serial port letter
you are using for “n” - for example, if using ttyb,
replace “n” with “b”, as: “...../dev/term/b.....”}}
```

- An easy way create this entry is to simply copy the lines from the `hardwire` command, then change the entry from *hardwire* to *annexterm*—ensure the port letter is correct.
- You can use the `tip`(1) command to connect the cluster administration console I/O with the TC I/O using the parameters specified in this file.

Step 2.2.3—Connecting to the TC

For local (manual) installations, on the administrative workstation, connect to the TC by entering the following command.

```
clustadm# tip annexterm
connected
```

Note – For local (manual) installations, you will not receive a prompt from the TC until you physically power off the TC, power it back on, and within 30 seconds, press and hold the TEST button until the power LED starts flashing, as described in the next step.

Step 2.2.4—Configuring the TC and Verifying Settings

For local (manual) installations, configure the TC device in the following prescribed sequence:

1. Power on the TC while viewing the front panel and status LED indicators.
2. After powering on the unit, and *within 30 seconds*, press and hold the TEST button until the power LED starts flashing.

3. When the power LED starts to flash, release the TEST button for *at least* six seconds, then briefly push the TEST button again to initiate diagnostics.

Before proceeding to the next step, ensure that the TC successfully completes (the diagnostic tests take approximately 60 seconds to complete).

Key Practice: Verify that the TC settings are correct. TC firmware options and settings can vary between different TC revision levels; actual options will likely differ from those specified in this hands-on lab. When configuring the TC, refer to the manufacturers documentation to ensure settings are established correctly. Specifically, ensure settings for the TC Internet address, subnet mask, and broadcast address are as indicated in the following steps.

Step 2.2.5—Verifying the TC Settings

For local (manual) installations, when the diagnostic tests are completed, ensure that the Tip window of the administrative workstation appears as follows.

```
System Reset - Entering Monitor Mode
monitor:
```

Note – For the next few steps to configure the TC, the settings should be configurable “as-listed;” however, the TC firmware settings vary from unit-revision to unit-revision. Your actual options may differ. When configuring the TC, refer to the manufacturer’s documentation to ensure the settings are established correctly. Specifically, ensure the settings for the TC Internet address, subnet mask, and broadcast address, as indicated in the following steps.

Step 2.2.6—Setting TC Values

For local (manual) installations, use the addr command to set the Internet address, subnet mask, and broadcast address for the TC. To do this, enter the following commands.

```
System Reset - Entering Monitor Mode
monitor:addr
Enter Internet address [192.40.85.60]:: 192.9.200.4
Enter Subnet mask [255.255.255.0]:: <CR>
Enter Preferred load host Internet address [<any host>]::
<CR>
Enter Broadcast address [129.153.49.255]:: 192.9.200.255
Enter Preferred dump address [0.0.0.0]:: <CR>
Select type of IP packet encapsulation(ieee802/ethernet)
[<ethernet>]:: <CR>
Type of IP packet encapsulation: <Ethernet>
Load Broadcast Y/N:: [N] <CR>
monitor:: sequence

At this point you need to enter a list of 1 to 4
interfaces to attempt to use for downloading or upline
dumping. Enter them in the order they should be tried,
separated by commas or spaces.

Possible interfaces are:
Ethernet: net
SELF: self

Enter interface sequence [net]:: self
Interface sequence: self
monitor:: ~.     {{this command ends the tip session}}
```

The addr command displays and sets several annex TC operating parameters, IP address, subnet mask, preferred load host IP address, load/dump gateway IP addresses, broadcast address, and IP encapsulation type.

The sequence command edits the load/dump interface list. This list determines the order of the network interfaces and determines whether the local area network (LAN) or the SLIP interface will be used by the Annex TC for loading and dumping. The default Ethernet selection is "net," which uses the LAN interface.

The "self" selection specifies that the TC is self-boot configured.

Step 2.2.7—Initializing Changes to the Configuration

For local (manual) installations, the TC must be power cycled for the previous (manual) configuration changes to take effect. To put the changes into effect:

1. Power off the TC.
2. Power on the TC and wait approximately 90 seconds for it to configure.

Key Practice: Because port 1 of the TC is the configuration port, minimize security vulnerability by disconnecting the administrative workstation from port 1 of the TC after the configuration. This will prevent unauthorized access to the TC's configuration port.

Default Router Configuration

Though this is not required for these hands-on exercises, if the TC requires access from an adjacent network, the `defaultrouter` configuration must be performed on each cluster node. This would be performed later, after the Solaris OE installation has completed on each cluster node. At that time, configure the default router information on the cluster nodes by creating the file `/etc/defaultrouter` and inserting the IP address of your gateway.

The following example shows how to configure the gateway address when the TC requires access from an adjacent network.

```
192.9.200.254 {{sample gateway address}}
```

Note – In this example, we use the gateway IP address instead of the hostname (see following best practice).

Key Practice: For each node being managed, maintain all IP addresses in one location. Edit the file `/etc/defaultrouter` and add the hostname of the gateway system. Next, edit the `/etc/inet/hosts` file and reference the gateway (hostname, IP address, plus an appropriate comment to identify the entry—for example, default gateway). In this manner, you can easily record (and maintain) a complete list of all nodes on the SunPlex platform.

If you want to enable a default route without rebooting and want to verify your new router setting, enter the following commands.

```
clustadm# route add default 192.9.200.254 {{sample gateway
address}}

clustadm# netstat -rn {{look for 192.9.200.254 entry with
“ug” flag}}
```

Step 2.2.8—Completing the Configuration

For local (manual) installations, complete the configuration of the TC by entering the following commands. Enter the data as shown in the following example, where prompted.

```
clustadm# telnet tc
Trying 192.2.200.4...
Connected to tc.
Escape character is '^]'.
cli
Annex Command Line Interpreter * Copyright 1991 Xylogics,
Inc.
annex: su
Password: 192.9.200.4 {{the password defaults to the assigned IP address, but
does not echo to the screen}}
annex# edit config.annex
Ctrl-W: save and exit Ctrl-X: exit Ctrl-F: page down Ctrl-
B: page up
       # The following are definitions of the gateway
entries
       #
       %gateway
       #
       # The following are definitions of the macro
entries
       #
       %macros
       %include macros
       #
       # The following are definitions of the rotary
entries
       #
       %rotary
       %include rotaries
       <Cntrl>+W #to exit and save
annex# admin
admin: port all
admin: set port mode slave
You may need to reset the appropriate port, Annex
subsystem or reboot the Annex for changes to take effect.
admin: reset
reset default port set [y] ? y
admin: quit
annex# hangup {{ends the login session}}

clustadm#
```

If the TC requires access from an adjacent network, the default router should be included in the `config.annex` file of the TC.

Key Practice: To avoid security vulnerability, change the default password for the TC. After changing the password, maintain your telnet session with the TC and use another window to telnet to the TC to verify your new password works. A typographical error when entering a password change will render the TC inaccessible. See the SunSolve Online article SRDB ID 24989 for resetting a lost root password.

Step 2.2.9—Changing the TC Password

For local (manual) installations, and to avoid security vulnerability, change the TC password, which is similar to changing a UNIX system password, as follows:

1. Telnet to the TC, enter the `cli` command at the port prompt (to enable the command interpreter).
2. Enter `su`, then the default password (the TC IP address).
3. Enter the `passwd` command at the prompt and change the default password. Record any changes, and the new password setting.

Note – While changing the password is important in the field, for the purpose of these manual procedures, we use the default password. In case the default password is lost, you must enter the `erase` command to clear the EEPROM contents and then re-enter all data. See the SunSolve Online article SRDB ID 24989.

Summary of Key Practices

Verify the TC settings are correct. TC firmware options/settings can vary between different TC units.
Disconnect the administrative workstation (serial port) from the TC (port 1) to avoid security vulnerability.
Maintain all IP addresses in a central location for each node being managed.
Change the default password of the TC to avoid security vulnerability.

End of Section 2.2

This completes this section. The TC configuration has been verified.

Section 2.3: Configuring the Solaris JumpStart™ Server

Configure the management server (administrative workstation) to act as the Solaris JumpStart server, for each cluster node.

Note – This section is “optional” and should only be performed during a local (manual) installation, when configuring the administrative workstation (`clustadm`) as a JumpStart server.

Use JumpStart software to maintain consistency and to fully automate the installation of the Solaris Operating Environment (Solaris OE) and additional software packages. The JumpStart software can minimize operator errors that occur during a manual installation process. Combining JumpStart software and Flash archives (for example, `flar` and `flarcreate`), enables quick (consistent) disaster / recovery operations.

Note – For information about implementing Flash archives (`flar` and `flarcreate`), refer to recent publications available from Sun BluePrints web site, at `http://www.sun.com/blueprints/`, and the references in Appendix B.

Step 2.3.1—Inserting the Solaris OE CD

For local (manual) installations, insert the Solaris 8 OE CD into the CD-ROM drive of the management server (administrative workstation). (The `vold(1M)` daemon will automatically mount in the `/cdrom` directory).

Step 2.3.2—Creating the JumpStart Sotware OS Directory

For local (manual) installations, create the JumpStart OS directory by entering the following command in the management server (administrative workstation).

```
clustadm# mkdir -p /JumpStart/OS/Solaris_2.8_07.01
```

Key Practice: Where possible, use a directory name that reflects the contents of the directory. For example, the name "`/JumpStart/OS/Solaris_2.8_07.01`" clearly identifies the release date and version of the operating system.

Caution – In the following step, ensure that the directory you create has enough space to hold the contents of the entire operating system CD (approximately 1.3 Gbyte). Use the `df -k` command to confirm that adequate disk space is available.

Step 2.3.3—Setting up the JumpStart Server

For local (manual) installations, continue setting up the JumpStart server by entering the following command into the administrative workstation (this step takes approximately one hour to complete).

```
clustadm# cd /cdrom/cdrom0/s0/Solaris_8/Tools

clustadm# ./setup_install_server /JumpStart/OS/Solaris_2.8_07.01

verifying target directory...
Calculating the required disk space for the Solaris_8 product
Copying the CD image to disk...
Install Server setup complete
```

Step 2.3.4—Accessing Directories Over a Network With NFS

For local (manual) installations, to ensure the operating system image can be installed on the nodes from a network, edit the `/etc/dfs/dfstab` file to include the following line, making the directories accessible over a network using NFS.

```
 share -F nfs -o ro,anon=0 /JumpStart

{{Note: Add this line. The anon=0 option allows the
install clients access as root to the specified file
system}}
```

Step 2.3.5—Sharing Directories Without Rebooting

For local (manual) installations, ensure the NFS daemon (nfsd) is running, then share the directories by entering the following command at the administrative workstation.

```
clustadm# /etc/init.d/nfs.server stop
clustadm# /etc/init.d/nfs.server start
{{Note: This makes the directories available without
rebooting}}
```

Step 2.3.6—Validating Directory Sharing

Validate that the directories have been successfully shared by entering the following command at the administrative workstation.

```
clustadm# share
-               /JumpStart   ro,anon=0  ""
clustadm#
```

Step 2.3.7—Creating the sysidcfg File

For local (manual) installations, create the /JumpStart/OS/JumpStart/sysidcfg file and include the following lines: (the system locale and time zone entries should correspond to your JumpStart area and time).

```
system_locale=en_US #sample
timezone=US/Central #sample
terminal=sun
name_service=NONE
timeserver=JumpStarthost #should be hostname of admin workstation
network_interface=hme0{netmask=255.255.255.0}
```

Step 2.3.8—Making sysidcfg File Readable

For local (manual) installations, change the permissions of the `sysidcfg` file to be readable by all by entering the following command at the administrative workstation (management server).

```
clustadm# chmod 644 /JumpStart/sysidcfg
```

Step 2.3.9—Configuring the hosts File

To ensure the `/etc/inet/hosts` file is configured to support the SunPlex platform, add the IP address for each Solaris OE install client (cluster node) to be serviced by the JumpStart server. Also add the additional hostname entries required to support the cluster.

- At this time, ensure that the corresponding entries are created in support of the SunPlex platform, including entries for each cluster node to be managed, for the `clustadm` workstation, for the terminal concentrator (`tc`), and for any additional site-specific entries, such as the two logical host entries (as required for our two data services).

```
clustadm# more /etc/inet/hosts

. . . {output omitted}

xxx.xxx.xx.xxx  clustadm loghost
xxx.xxx.xx.xxx  clustnode1
xxx.xxx.xx.xxx  clustnode2
xxx.xxx.xx.xxx  lh-hanfs
xxx.xxx.xx.xxx  lh-apache
xxx.xxx.xx.xxx  tc

. . . {output omitted}
```

Step 2.3.10—Copying Sample Files

For local (manual) installations, copy the following files from the `/JumpStart/OS/Solaris_2.8_07.01/Solaris_8/Misc/jumpstart_sample` directory to the `/JumpStart` directory.

- `rules`
- `check`
- `set_root_pw`

Verify these files were copied correctly.

Note – The `set_root_pw` file is not needed when using the `sysidcfg` file.

Step 2.3.11—Enabling Solaris JumpStart Software

For local (manual) installations, enable Solaris JumpStart software for both cluster nodes (`clustnode1` and `clustnode2`) by adding each as Solaris 8 OE install clients. Recall that your Ethernet (MAC) address will be different from the example shown. Ensure that these site-specific entries are correct for your implementation.

Example, adding `clustnode1` as a Solaris JumpStart software client.

```
clustadm# cd /JumpStart/OS/Solaris_2.8_07.01/Solaris_8/Tools

clustadm# ./add_install_client -e 8:0:20:cf:27:62 \
-i 192.9.200.2 \
-s clustadm:/JumpStart/OS/Solaris_2.8_07.01 \
-p clustadm:/JumpStart/sysidcfg
-c clustadm:/JumpStart \
clustnode1 sun4u
```

Where:
`/JumpStart/OS/Solaris_2.8_07.01/Solaris_8/Tools` equals the location of the `add_install_client` utility.

Where:
`-e 8:0:20:cf:27:62` equals the Ethernet address should correspond with the data gathered for this node during Module 1.

Note – Do not enter the Ethernet address used in these examples. Use the Ethernet address specific to your machine.

`-i 192.9.200.2` equals the IP address assignment for `clustnode1`.

`-s clustadm:/JumpStart/OS/Solaris_2.8 07.01` equals the location of the CD-ROM image.

`-p clustadm:/JumpStart/sysidcfg` equals the location of the `sysidcfg` file previously verified.

`-c clustadm:/JumpStart` equals the location of the `rules.ok` file that was previously verified.

`clustnode1 sun4u` specifies that `clustnode1` is a Sun Enterprise™ 220R server (UltraSPARC™ II architecture).

The following example shows how to add `clustnode2` as a Solaris JumpStart software client.

```
clustadm# cd /JumpStart/OS/Solaris_2.8_07.01/Solaris_8/Tools

clustadm# ./add_install_client -e 8:0:20:cf:23:a3 \
-i 192.9.200.3 \
-s clustadm:/JumpStart/OS/Solaris_2.8_07.01 \
-p clustadm:/JumpStart/sysidcfg
-c clustadm:/JumpStart \
clustnode2 sun4u
```

Where:
`/JumpStart/OS/Solaris_2.8_07.01/Solaris_8/Tools` equals the location of the `add_install_client` utility.

Where:
`-e 8:0:20:cf:23:a3` equals the Ethernet address should correspond with the data gathered for this node, during Module 1.

Note – Do not enter the Ethernet address used in these examples. Use the Ethernet address specific to your machine.

`-i 192.9.200.3` equals the IP address assignment for `clustnode2`.

`-s clustadm:/JumpStart/OS/Solaris 2.8 07.01` equals the location of the CD-ROM image.

`-p clustadm:/JumpStart/sysidcfg` equals the location of the `sysidcfg` file previously verified.

`-c clustadm:/JumpStart` equals the location of the `rules.ok` file that was previously verified.

`clustnode2 sun4u` specifies that `clustnode2` is a Sun Enterprise 220R server (UltraSPARC-II architecture).

Step 2.3.12—Verifying Successful Completion

Verify the `add_install_client` commands completed successfully.

Note – In this exercise the nodes to be configured for use with JumpStart software are named `clustnode1` and `clustnode2`.

Confirm that the add_install_client command was executed by viewing the /etc/bootparams, and /etc/ethers files.

```
clustadm# more /etc/bootparams
clustnode1
root=clustadm:/JumpStart/OS/Solaris_2.8_07.01/Solaris_8/Tools/Boot
install=clustadm:/JumpStart/OS/Solaris_2.8_07.01/ boottype=:in
sysid_config=clustadm:/JumpStart install_config=clustadm:/
JumpStart/sysidcfg  rootopts=:rsize=32768

clustnode2
root=clustadm:/JumpStart/OS/Solaris_2.8_07.01/Solaris_8/Tools/Boot
install=clustadm:/JumpStart/OS/Solaris_2.8_07.01/ boottype=:in
sysid_config=clustadm:/JumpStart install_config=clustadm:/
JumpStart/sysidcfg  rootopts=:rsize=32768

clustadm# more /etc/ethers
8:0:20:cf:27:62 clustnode1 {{your ethernet address will differ}}
8:0:20:cf:23:a3 clustnode2 {{your ethernet address will differ}}
```

Note – Refer to the Appendix B references about configuring a JumpStart server. Note, also, that if the add_install_client command fails to add these entries correctly, or an older version of the file(s) had previously existed, it may be necessary to remove a failed entry (for example, from the /etc/bootparams file).

To remove a suspected entry for a specific failing install client, enter the following command.

```
clustadm# cd /JumpStart/OS/Solaris_2.8_07.01/Solaris_8
clustadm# ./rm_install_client client_hostname
```

Note – After removing a failing node, re-run the add_install_client commands, as shown in the previous step where this was performed. This may be necessary if, somehow, the previous attempts failed and entries were not correctly removed.

Summary of Key Practices

Where possible, use a directory name that reflects the contents of the directory. Confirm add_install_client was performed correctly.

End of Section 2.3

This completes this section. The Solaris JumpStart software has been verified, and each cluster node can be quickly (and consistently) installed with the Solaris OE version supported by the cluster.

Section 2.4: Installing SUNWccon Package on the Administrative Workstation (Management Server)

Install the Sun Cluster 3.0 software on the administrative workstation, for accessing the console on each cluster node.

Step 2.4.1—Installing the SUNWccon Package

Install the Sun Cluster 3.0 software SUNWccon package, which contains the administrative and console software, on the management server.

For example, change to the Sun Cluster 3.0 software “Packages” directory containing the SUNWccon package by entering the following commands.

```
clustadm# cd  /cdrom/suncluster_3_0/SunCluster_3.0/Packages
clustadm# pkgadd -d . SUNWccon
```

The installation of Sun Cluster 3.0 software is now complete. The cluster software provides the tools needed to build and manage the cluster. These tools will be used after configuring the cluster software.

Note – To access the `man` (1M) pages and binaries for the Sun Cluster 3.0 software packages, the `MANPATH` and `PATH` variables will need updating.

Step 2.4.2—Automatically Setting Environment Variables

To automatically set the environment variables each time a new shell is invoked, add the following environment variables to your startup files (`/.profile` or `/.cshrc`) as shown in the examples in Appendix A.

- Add `/opt/SUNWcluster/bin:` to the `PATH` variable.
- Add `/opt/SUNWcluster/man:/opt/SUNWconn/man:` to the `MANPATH` variable.

Caution – It is important that you accurately enter the names of each variable. Ensure that the variables are set correctly to avoid unpredictable results.

Step 2.4.3—Initializing Changes

For the changes made in the previous step to take effect, log out and log back in as the root user, or execute the `source` command as shown in the following example.

```
clustadm# source /.cshrc {{for csh users}}

{{or, for the bourne shell you may enter: }}
clustadm# . /.profile
```

End of Section 2.4

This completes this section. The SUNWccon software has been installed.

Section 2.5: Configuring the Management Server to Administer Cluster Nodes

Configure the management server to administer each of the cluster nodes as explained in the following sections.

Step 2.5.1—Creating the etc/clusters File

The first step is to create the /etc/clusters file with an entry for this cluster, as follows.

```
<Name of Cluster> <Node1 Name> <Node2 Name> ...
```

Where:

- <Name of Cluster> is nhl
- <Node1 Name> is clustnode1
- <Node2 Name> is clustnode2

On the management server, create the /etc/clusters file, configuring clustnode1 and clustnode2. Verify that the /etc/clusters file is correct as shown in the following example.

```
clustadm# cat /etc/clusters
nhl clustnode1 clustnode2
```

Step 2.5.2—Enabling the Connection to the Terminal Concentrator

Configure the /etc/serialports file on the management server (administrative workstation) enabling a connection to the terminal concentrator (tc) through the ttya (serial) port. Create the /etc/serialports file by using the following format.

At the management server (administrative workstation), verify that /etc/serialports is correct, as follows.

```
clustadm# cat /etc/serialports
clustnode1 tc 5002
clustnode2 tc 5003
```

Note – In this example, the TC has eight ports, 1–8. The `5002` entry refers to the TC physical port 2, and the `5003` entry refers to physical port 3. The `tc` entry must correspond to the hostname entry in `/etc/inet/hosts`.

Step 2.5.3—Disabling snmpdx Daemon

To avoid a potential conflict between the cluster control software and the Solaris OE `snmpdx` daemon, disable the `snmpdx` daemon by changing the file name that invokes the `snmpdx` daemon from `/etc/rc3.d/S76snmpdx` to `/etc/rc3.d/s76snmpdx` (note the lower case `s`). By changing the name of the file, this daemon will not start when the system initializes (init 3).

At the management server (administrative workstation), rename the file by entering the following command.

```
clustadm# mv /etc/rc3.d/S76snmpdx /etc/rc3.d/s76snmpdx
```

End of Section 2.5

This completes this section. You may now start the CCP.

Section 2.6: Configuring the Cluster Control Panel

During this section, you will start the cluster control panel (CCP) by entering the `ccp` command for the cluster named `nhl`. Please read this entire step before entering any commands. After starting the cluster control panel, you will double-click the Cluster Console icon.

Step 2.6.1—Enabling Remote Display

If you are accessing the administrative workstation (`clustadm`) from a remote system, first use the `xhost +` command on your local system to enable remote display from `clustadm` to your local system.

When accessing the administrative workstation remotely, you must also set the `DISPLAY` environment variable on the administrative workstation (`clustadm`) to point to your local system (for example, `setenv DISPLAY` *`yoursystem`*`:0.0`).

Note – This step can be performed when accessing the administrative workstation from a remote workstation. It is often useful to access the CCP remotely, as appropriate, or when configuring (administering) the Sun Cluster software.

If accessing `clustadm` remotely, you must set the `DISPLAY` variable before invoking the CCP.

1. First, on your local workstation type the following (example only).

```
yoursystem# /usr/openwin/bin/xhost +clustadm
```

2. Next, on `clustadm` type the following (example only).

```
clustadm# setenv DISPLAY yoursystem:0.0
```

Step 2.6.2—Displaying the CCP

To display the CCP, first verify that your shell environment is setup and that shell variables are set according to the environment variables listed in Appendix A. Enter the following commands in the shell window on the management server (administrative workstation).

```
clustadm# which ccp
/opt/SUNWcluster/bin/ccp

clustadm# ccp nhl &
```

When the ccp command is executed, it will display the Cluster Control Panel window; a menu bar and icon panel will display all available tools. See the following figure.

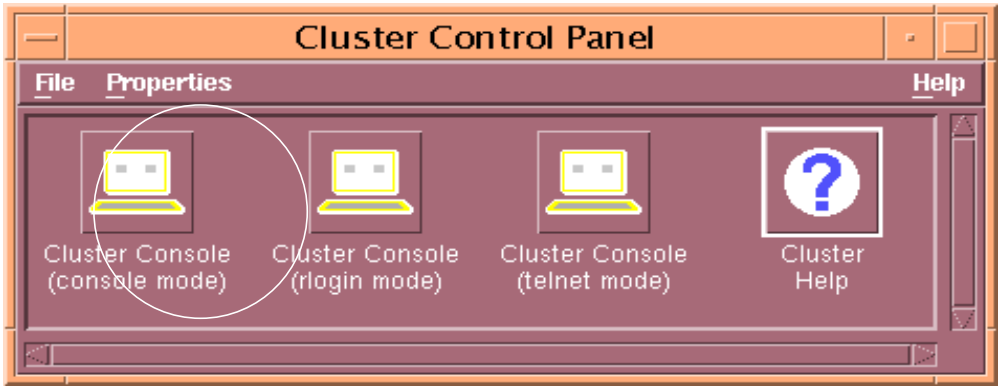

Figure 2-1 Cluster Control Panel Window

Step 2.6.3—Customizing the CCP Display

Refer to the preceding figure. Double-click the Cluster Console (console mode) icon (circled), to display the Cluster Console window, as shown in the following figure.

In this example, three windows are displayed: one small Cluster Console window, and two larger *cconsole: host [name]* windows. Note that each of the larger windows is associated with a specific host, or cluster node.

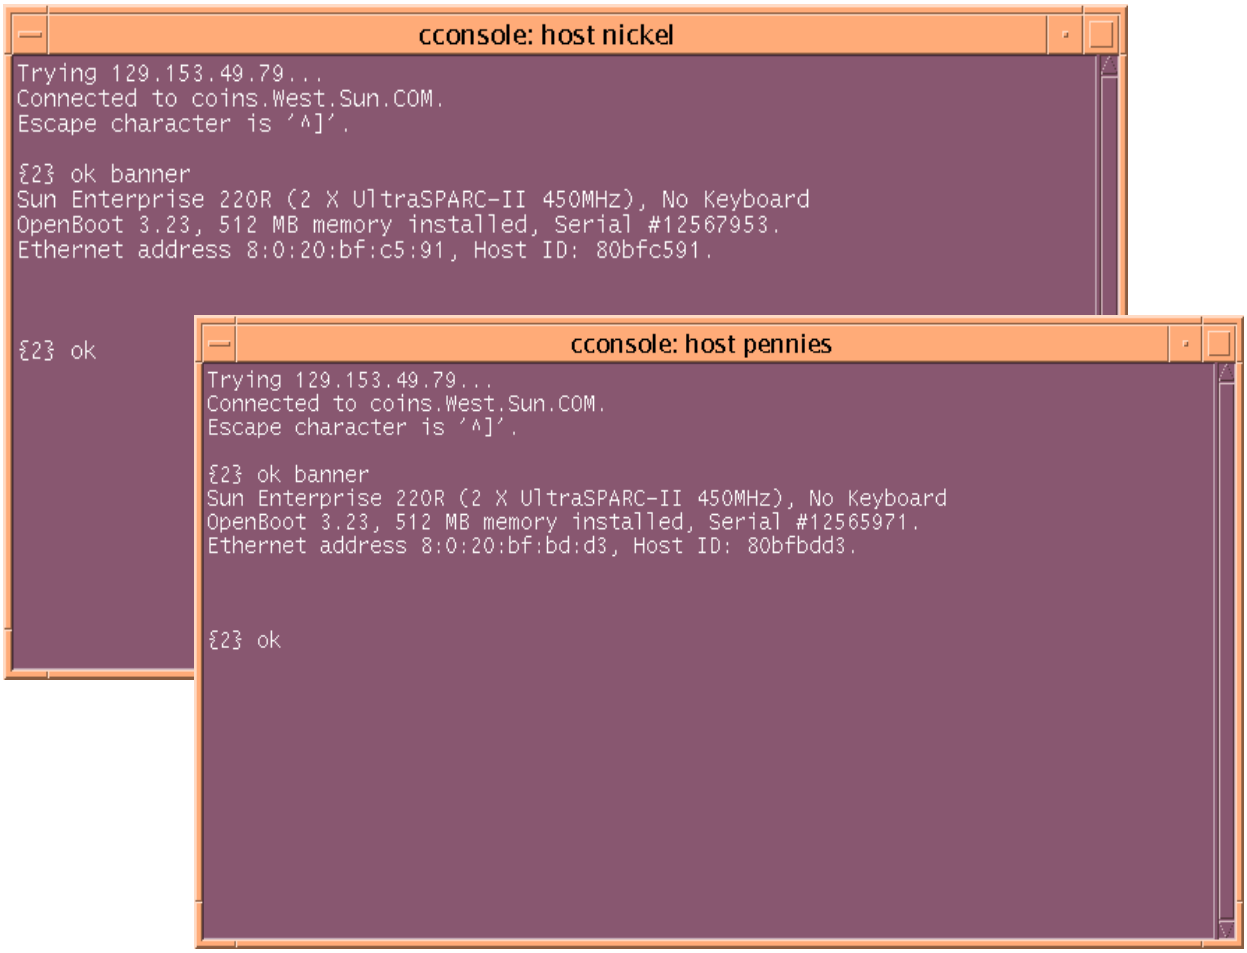

Figure 2-2 Figure: 2-2 Cluster Console (Console Mode) and cconsole Windows

Caution – The Cluster Console utility provides a method of entering commands into multiple cluster nodes either simultaneously or individually, as required. Always be aware of which window is active before entering commands.

If a Cluster Console window does not appear for a cluster node, follow these steps to display the missing window:

1. From the Cluster Console window (console mode), select Hosts, and click Select Hosts.

2. Verify or insert an entry for each cluster node (for example, there should be entries for `clustnode1`, `clustnode2`).

At this time, arrange each window to your personal viewing preferences. To ungroup the Cluster Console window from the *cconsole: host [name]* windows, select Options from the Cluster Console window and uncheck the Group Term Windows checkbox.

For example, arrange the *cconsole* windows to be able to see each window clearly (and at the same time) by moving the Cluster Console window away from the other cluster node windows. This is done to ensure commands are entered correctly into one, or both nodes, as required during these exercises (and to prevent entering commands into the wrong window).

It is not necessary to do so at this time, but when you wish to close the Cluster Console window, select Exit from the Host menu on the Cluster Console window.

Note – If you need to issue a `Stop-A` command to the cluster nodes to place them into the OpenBoot PROM mode, position the cursor over the Cluster Console window, left mouse click, then press the `^]` (Ctrl +]) keys. This will access the `telnet>` prompt. At the `telnet>` prompt, enter the `send brk` command into each hosts cconsole window to return to the OpenBoot PROM `ok` prompt.

Step 2.6.4—Configuring Each Cluster Node

Verify the CCP utility is functioning and begin the software configuration of each cluster node.

At this time, ensure that each cluster node is at the OpenBoot PROM `ok` prompt.

Note – For local (manual) installations, prior to installing Solaris Operating Environment (Solaris OE) on each cluster node, we will reconfigure the OpenBoot PROM settings for each cluster node. This is achieved by executing commands at the OpenBoot PROM `ok` prompt (the `ok` prompt should be viewable through the Cluster Control Panel window). We are able to do this because we have previously configured

(and connected) the serial (console) port from each cluster node to the terminal concentrator (refer to Figure 1-1 in the "Guide to Installation—Hardware Setup," for cabling / connections).

Key Practice: On each node on the SunPLEX platform, ensure the system EEPROM has the current (for example, latest or required) OpenBoot PROM version. This could prevent undesirable anomalies (bugs) from affecting your system reliability / performance. To obtain or verify that you have the most the most current version, refer to http://sunsolve.sun.com.

Step 2.6.5—Resetting EEPROM to Factory Defaults

Prior to configuring or customizing the system EEPROM, before installing any software, and to meet cluster requirements, reset the system EEPROM to its factory default. To do this:

1. Position the cursor into the Cluster Console window and left-click to make the window active.
2. Enter the set-defaults command. As you enter the command, you should see the word *set-defaults* echoed in the cconsole window, on each cluster node.

For local (manual) installations, at this time (at each cluster node, prior to any subsequent EEPROM settings), enter the following OpenBoot PROM command (ok prompt).

```
ok set-defaults
```

Using the set-defaults command at this step establishes a consistent, known (default) state of all OpenBoot PROM variables (prior to customizing the OpenBoot PROM environment).

Caution – Resetting the system EEPROM should only be performed at this time, during the initial preparation for the Solaris OE installation. This command resets all EEPROM (OpenBoot PROM) variables to their known, factory default values. All subsequent steps assume the EEPROM has been reset at this point in the exercise. During the next few steps, the EEPROM will be modified or customized.

Ensure a consistent state on each cluster node before proceeding to configure site-specific (custom) OpenBoot PROM settings. Prior to implementing any configuration changes, and as part of *initial* Solaris OE installation preparations, reset the EEPROM to the factory defaults. This is done only once, and at this point in the procedure and will easily (quickly) ensure that a known and consistent state is achieved before further customization occurs.

Step 2.6.6—Documenting System Information

Execute the banner command, as follows, and document the system information, such as the system model number, OpenBoot PROM version, Ethernet address, hostid, and serial number.

```
ok banner
```

Each node will respond with configuration information.

Summary of Key Practices

On each node on the SunPLEX platform, ensure the system EEPROM has the current/latest OpenBoot PROM version.

Prior to customization, reset the system EEPROM to a consistent, known state (factory default).

Execute the OpenBoot PROM banner command, and verify (document) important system information.

End of Section 2.6

This completes this section. The CCP has been configured and verified.

Appendix A: System Configuration Files

This appendix presents the following information:

- Sample `/.profile` and `/.login` startup scripts
- Sample `/etc/inet/hosts`

Output From `/.profile` and `/.login` Files

The following example configuration files are referenced in this module.

Sample `/.profile` File

```
#!/bin/ksh
# root ksh environment settings
TERM=vt220
stty=istrip
Prompt='/bin/hostname'
PATH=/usr/bin:/usr/ucb:/etc:/sbin:/usr/sbin:\
/usr/cluster/bin:/opt/SUNWcluster/bin:.
MANPATH=/usr/dt/man:/usr/man:/usr/openwin/share/man:\
/usr/cluster/man:/opt/SUNWcluster/man:/opt/SUNWconn/man:
export TERM stty Prompt PATH MANPATH
umask 022
```

Sample `/.login` File

```
# The root csh environment settings
setenv TERM vt220
setenv stty istrip
set prompt = 'bin/hostname'%
setenv PATH /usr/bin:/usr/ucb:/etc:/sbin:/usr/sbin:\
/usr/cluster/bin:/opt/SUNWcluster/bin:.
setenv MANPATH /usr/dt/man:/usr/man:/usr/openwin/share/man:\
/usr/cluster/man:/opt/SUNWcluster/man:/opt/SUNWconn/man
umask 022
```

Note – Discussing shells and their usage can be highly controversial. The above scripts have been written as generic as possible. You are welcome to add to the scripts or not use them at all, however, keep in mind that you will have to troubleshoot any errors the shell scripts create.

Sample /etc/inet/hosts File

```
# more /etc/inet/hosts

# Internet host table

127.0.0.1        localhost

129.153.xx.xxx clustadm loghost

129.153.xx.xxx clustnode1

129.153.xx.xxx clustnode2

129.153.xx.xxx lh-hanfs

129.153.xx.xxx lh-apache

129.153.xx.xxx tc
```

Note – Example only. Shows standard hostname entries for the SunPLEX platform. Note the output includes our two logical host entries ('lh-xxx'), plus the terminal concentrator (tc). Verify your configuration matches actual site-specific configuration requirements.

Appendix B: References

This appendix provides a list of URLs and references to online documents and Sun BluePrints OnLine articles to review when implementing Sun Cluster platform solutions.

On-Line References to Sun Cluster 3.0 Documents and Sun BluePrints Publications

Cluster Platform 220/100 Architecture - A Product from SunTone™ Platforms Portfolio by Enrique Vargas can be found at:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Sun Web Start Flash information can be found by browsing
`http://www.sun.com/solaris/webstartflash/`

Sun BluePrints OnLine articles and valuable writings are available by browsing
`http://www.sun.com/blueprints`

Availability - What It Means, Why It's Important, and How to Improve It, by Richard McDougal
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Disaster Recovery Requirements, by Stan Stringfellow can be found by browsing
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

High Availability Best Practices, by Enrique Vargas can be found by browsing
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

HA: Boot/Root/Swap, by Jeannie Johnstone Kobert can be found by browsing
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

The Art of Production Environment Engineering by Bill Walker, can be found by browsing
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Datacenter Naming Scheme, by Mark Garner, can be found by browsing
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Robust Clustering: A Comparison of SunCluster 3.0 vs. Sun Cluster 2.2
`http://www.sun.com/blueprints/0901/sc30vs22.html`

Architecting a Service Provider Infrastructure for Maximum Growth, by Stan Stringfellow can be found by browsing
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Building Secure N-Tier Environments, by Alex Noordergraaf can be found by browsing:
`http://www.sun.com/software/solutions/blueprints/browsesubject.html`

Planning for Large Configurations of Netra t1 Servers, by Stan Stringfellow can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Planning to Fail, by John S. Howard:
http://www.sun.com/software/solutions/blueprints/browsesubject.html

High Availability Fundamentals, by Enrique Vargas:
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Tales from the Trenches: The Case of the RAM Starved Cluster, by Richard McDougal can be found by browsing:
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Building a JumpStart Infrastructure, by Alex Noordergraaf can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Using NTP to control and Synchronize System Clocks (Parts 1 - 3) can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Solaris 8 Additions to "sysidcfg," by Rob Snevely can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html

An Overview of Methodology, by Adrian Cockcroft can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html

JumpStart™ Mechanics: Using JumpStart Application for Hands-Free Installation of Unbundled Software Parts 1 and 2, by John S. Howard
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Setting Up A Solaris Operating Environment Install Server and the Solaris JumpStart Feature, by Rob Snevely
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Sun GigaBit Ethernet vs. Alteon WebSystems, Inc. : Performance (throughout) Benchmark Analysis and Datasheet
http://www.sun.ca/products-n-solutions/hw/networking/connectivity/sungigabitethernet/gem.html#analysis

Scrubbing Disks using Solaris OE "format" Program, can be found by browsing at
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Fast Oracle Exports, by Stan Stringfellow can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html

IP Network Multipathing, by Mark Garner can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Sun / Oracle Best Practices can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html

MR System for Rapid Recovery, by John S. Howard can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html

VxVM Private Regions: Mechanics and Internals of the VxVM Configuration Database, by Gene Trantham
http://www.sun.com/software/solutions/blueprints/browsesubject.html

Veritas VXVM Storage Management Software, by Gene Trantham
http://www.sun.com/software/solutions/blueprints/browsesubject.html

SSP Best Practices, by John S. Howard can be found by browsing
http://www.sun.com/software/solutions/blueprints/browsesubject.html