



Sun Fire™ Midframe Server Best Practices Update for Firmware 5.13.x

*By James Hsieh, Customer Problem Resolution
Engineering, Americas*

Sun BluePrints™ OnLine—Oct 2002



<http://www.sun.com/blueprints>

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95045 U.S.A.
650 960-1300

Part No. 817-0080-10
Revision 06, 10/08/2002
Edition: October 2002

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California, 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. 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 United States and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Sun BluePrints, SunDocs, Sun Enterprise, Sun Fire, JumpStart, Sun StorEdge, SunSolve, OpenBoot, AnswerBook2, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the US and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. Netscape Navigator is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries.

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.

Sun Microsystems, Inc. has intellectual property rights relating to technology described in this document. In particular, and without limitation, these intellectual property rights may include one or more patents or pending patent applications in the U.S. or other countries.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

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 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054 Etats-Unis. Tous droits réservés.

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. 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 enregistrée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company Ltd.

Sun, Sun Microsystems, le logo Sun, Sun BluePrints, SunDocs, Sun Enterprise, Sun Fire, JumpStart, Sun StorEdge, SunSolve, OpenBoot, AnswerBook2, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, 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. Netscape Navigator est une marque de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPENDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Please
Recycle



Adobe PostScript

Sun Fire Midframe Servers Best Practices Update for Firmware 5.13.x

The Sun Fire™ 3800, 4800, 4810, and 6800 midframe servers provide new functionality to monitor, diagnose, and administer the system, which can increase overall system reliability, availability, and serviceability (RAS). Much of the new functionality is available through the Sun Fire System Controller (SC), which is a central component of the Sun Fire midframe server.

This article revisits existing best practices that were published in “Sun Fire™ Midframe Server Best Practices for Administration”, Sun BluePrints Online, October 2001. The October 2001 article has been used as the basis for this current article, so it is not a prerequisite of reading this article. However, you should familiarize yourself with the content in the previous best practices article so that you are familiar with the differences between the firmware updates.

This article introduces new best practices based on the enhancements that have been made to the SC in firmware update 5.13.x. If you are already familiar with the prior update, you can concentrate on the new enhancements, which include:

- SC failover (see “Configuring the Sun Fire SC Failover” on page 4)
- SNMP support (see “Configuring SNTP” on page 10)

This article also includes discussions on the following topics:

- Platform configuration
- Platform administration
- Platform security
- Error analysis and diagnosis
- SC maintenance procedures

While many recommendations made here apply to the majority of cases, not all recommendations may apply to every circumstance.

Platform Configuration

This section contains descriptions of how to configure the Sun Fire Midframe platform. The topics include:

- “Configuring the RS-232 Serial Port”
- “Configuring the Ethernet Port”
- “Configuring a Switched Private Network”
- “Configuring the Sun Fire SC Failover”
- “Configuring SNTP”
- “Setting the Date and Time on the Platform”
- “Changing POST Levels and Other Settings”
- “Configuring the Midframe Service Processor”
- “Configuring the Sun MC Software”
- “Preparing for Firmware Updates”
- “Configuring the Sun Explorer Software Utility”

Configuring the RS-232 Serial Port

You can access the SC through the built-in RS-232 serial port or through its 10/100BASE-T Ethernet port. Be sure that access to the serial port is available during the initial setup of the SC because it is the only connection on which the SC power-on self-test (SCPOST) output can be viewed. The port settings should be 9600 bps, 8-bits, no parity, and 1 stop bit (9600-8-N-1).

You can also access the serial port by using a network terminal server or by using the serial port on a Midframe Service Processor (MSP). For more information about the need for an MSP and on how to configure the MSP, refer to “Configuring the Midframe Service Processor” on page 11.

After you have set up the SC, serial port access should continue to be available on demand to provide an alternate access path to the SC in the event of a network problem, firmware updates, or SC reboots or resets. Serial port access is also required to monitor certain SC and platform related errors because it is where these errors will be displayed.

Configuring the Ethernet Port

The Ethernet port should be used as the primary connection path for the speed, multisection access, and logging capabilities it provides. Ethernet connections to the SC are accomplished by using a Telnet session. A 100BASE-T link is strongly recommended for the SC Ethernet connection and required for use with Sun™ Management Center (Sun MC) software. The Ethernet port should not be used *instead of* the RS-232 serial port connection, but should be used *in addition to* the RS-232 port.

Configuring a Switched Private Network

You should configure the SC on a switched private network. If you are configuring two SCs for the network, assign each SC a separate IP address so that they do not conflict with each other on the network. If the SC failover functionality introduced in firmware 5.13.x is used, a third IP address representing the *logical* hostname can be assigned.

FIGURE 1 illustrates a simplified network topology. The MSP is a workstation placed on the private ethernet network of the Sun Fire Platform to provide administrative support functions to the Sun Fire platform and the SCs.

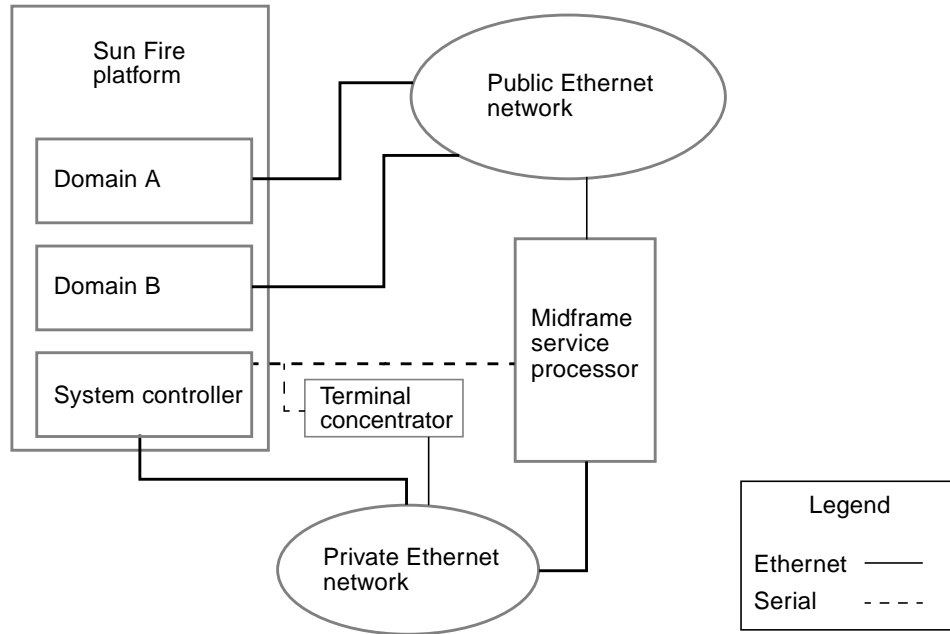


FIGURE 1 Simplified Network Topology

The serial connection in the illustration above can be replaced with a terminal concentrator if the same MSP is monitoring multiple Sun Fire platforms. If the terminal concentrator supports encrypted logins and sessions (for example, by using secure shell), the terminal concentrator may be connected to the public Ethernet network. A terminal concentrator is recommended to improve the ability to a single MSP to monitor multiple platforms.

Configuring the Sun Fire SC Failover

You should use two SCs in a Sun Fire system to provide failover of the SC functionality in case one of the SCs fail and to keep the domains in the system running. Prior to firmware version 5.13.0, if the main SC suffered a failure, administrative capabilities such as the ability to access domain consoles would be lost. With the introduction of firmware 5.13.0, full failover is available, so if the main SC fails, the spare SC can take over administrative functions, in addition to the system clock functionality.

When enabled, the two SCs communicate with each other by using an internal communications link. They also exchange health information and synchronize internal configuration information over the link. The SC that is acting as the *main* system controller also generates a heartbeat. If the heartbeat unexpectedly disappears, the *spare* SC takes over the *main* functionality.

Before enabling the SC failover feature, both SCs and all of the boards in a Sun Fire platform should be at the same firmware version. While it is possible to have mixed versions of firmware under certain circumstances, it is recommended that all of the boards and the SCs use the same version of firmware.

You can determine the firmware version by using the `showboards` command, as follows:

```
heslab-16-sc0:SC> showboards -p version

Component      Compatible Version
-----
SSC0           Reference  5.13.2
/N0/IB6        Yes        5.13.2
/N0/IB8        Yes        5.13.2
/N0/SB0        Yes        5.13.2
/N0/SB2        Yes        5.13.2

heslab-16-sc0:SC>
```

The above output does not include the version of firmware from the spare SC. To gather that information, you must connect to the spare SC and use the `showsc` command to determine the `ScApp` revision, as in the following:

```
heslab-16-sc1:sc> showsc

SC: SSC1
Spare System Controller
SC Failover: disabled

SC date: Sun Oct 06 14:06:58 PDT 2002
SC uptime: 2 days 6 hours 35 minutes 8 seconds

ScApp version: 5.13.2
RTOS version: 23

heslab-16-sc1:sc>
```

The SC failover software introduces a number of new commands and settings that need attention. You should use the `showfailover` command to check on failover status. The `-v` option gives the most information about the configuration.

```
heslab-16-sc0:SC> showfailover -v

SC: SSC0
Main System Controller
SC Failover: enabled and active.
Clock failover enabled.

heslab-16-sc0:SC>
```

The above information shows that the `showfailover -v` command was run on SSC0, that the SC is currently in the role of main, and that both SC administrative function failover and clock failover are enabled and functioning. You should run the `showfailover -v` command whenever you reboot a SC to ensure that the SC failover functionality has restarted and is functioning properly.

You can obtain an additional piece of information about the SC failover status by using the `showplatform -p sc` command, as in the following example.

```
heslab-16-sc0:SC> showplatform -p sc

SC POST diag Level: max
SC Failover: enabled and active.
Logical Hostname: heslab-16-sc

heslab-16-sc0:SC>
```

In the above example, the value for the logical hostname is displayed. Each SC continues to have a unique IP address assigned to it. The logical hostname is a third IP address that always points at whichever SC is currently functioning in the role of main. In the figure below, the *logical* SC is the logical hostname.

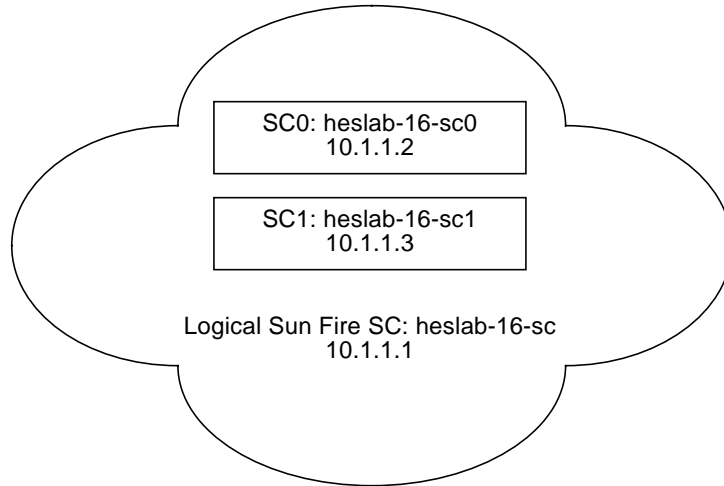


FIGURE 2 Sun Fire SC Logical Hostname

The following is an example of how to set up the Sun Fire SC failover functionality.

```
heslab-16-sc0:SC> setupplatform -p sc

SC
--
SC POST diag Level [max]: max
Enable SC Failover? [no]: yes
Logical Hostname or IP Address []: heslab-16-sc

heslab-16-sc0:SC>
```

To force the spare SC to assume the role of main, use the `setfailover force` command. This should not be necessary under normal operating conditions, but the functionality should be tested during a maintenance window after you enable the failover functionality to verify correct failover operations.

The use of SNTP, as discussed in “Configuring SNTP” on page 10, is strongly recommended with SC failover. If SNTP is not enabled, the time on the two SCs needs to be checked to ensure that they are the same. Because the domains on a Sun Fire platform derive their time based on the time set on the SCs, the time on the running domains could be affected after an SC failover if the time on the SCs is not synchronized.

Even though SC Failover copies configuration information from the main SC to the spare SC, it is not a replacement for backing up the SC. Users should still perform a `dumpconfig` of the SC after enabling failover and on a regular basis afterwards.

Setting the Date and Time on the Platform

When a Sun Fire platform is installed, the platform time should be set from the platform shell and in each individual domain using the `setdate` command. Each domain shell can have a separate time setting, so setting each one individually is required.

The following shows an example of how to set the date and time.

```
heslab-16-sc0:SC> setdate 09081228  
Sun Sep 08 12:28:00 PDT 2002  
heslab-16-sc0:SC>
```

You can set the time and date on the domains in a similar manner. Use the `setdate -h` command for additional help and options for setting the time.

The following is an example of the output from the `setdate -t list` command, which is helpful in determining the correct time zone to use for your locale:

```
heslab-16-sc0:SC> setdate -t list
list: is not a valid timezone, valid time zones are:

ACT  GMT+9.5  Central Standard Time (Northern Territory)
AET  GMT+10    Eastern Standard Time (New South Wales)
AGT  GMT-3     Argentine Time
ART  GMT+2     Eastern European Time
AST  GMT-9     Alaska Standard Time
BET  GMT-3     Brazil Time
BST  GMT+6     Bangladesh Time
CAT  GMT+2     Central African Time
CNT  GMT-3.5  Newfoundland Standard Time
CST  GMT-6     Central Standard Time
CTT  GMT+8     China Standard Time
EAT  GMT+3     Eastern African Time
ECT  GMT+1     Central European Time
EET  GMT+2     Eastern European Time
EST  GMT-5     Eastern Standard Time
GMT  GMT+0     Greenwich Mean Time
HST  GMT-10   Hawaii Standard Time
IET  GMT-5     Eastern Standard Time
IST  GMT+5.5  India Standard Time
JST  GMT+9     Japan Standard Time
MET  GMT+3.5  Iran Time
MIT  GMT-11   West Samoa Time
MST  GMT-7     Mountain Standard Time
NET  GMT+4     Armenia Time
NST  GMT+12   New Zealand Standard Time
PLT  GMT+5     Pakistan Time
PNT  GMT-7     Mountain Standard Time
PRT  GMT-4     Atlantic Standard Time
PST  GMT-8     Pacific Standard Time
SST  GMT+11   Solomon Is. Time
UTC  GMT+0     Coordinated Universal Time
VST  GMT+7     Indochina Time

heslab-16-sc0:SC>
```

Configuring SNTP

With SC firmware versions 5.13.0 and higher, the SC is capable of synchronizing its time-of-day clock with a network time server using SNTP. SNTP usage is encouraged to keep the SCs at an accurate time.

The following shows an example of how to enable SNTP.

```
heslab-16:A> setupdomain -p sntp

SNTP
----
SNTP server []: 10.1.63.251

heslab-16:A>
```

Changing POST Levels and Other Settings

To provide thorough testing of all components, the power-on self-test (POST) level for both the SC and domains should be set to maximum. Maximum is the default level for all domains. If you cannot always run the maximum POST level, you should at least use it during system installation. You should also use the maximum level of POST in other circumstances, such as if hardware is being replaced or moved after an unexpected system or power failure or when the hardware is suspected of causing system problems.

The following shows an example of the platform SCPOST level you should use:

```
heslab-16-sc0:SC> setupplatform -p sc

SC
--
SC POST diag Level [max]: max
Enable SC Failover? [no]: yes
Logical Hostname or IP Address []: heslab-16-sc

heslab-16-sc0:SC>
```

Note – For SCs running firmware versions lower than 5.13.0, the parameters for controlling SC Failover will not be visible.

The following show an example of the domain boot parameters we recommend:

```
heslab-12:B> showdomain -p bootparams

diag-level = max
verbosity-level = off
error-level = max
interleave-scope = within-board
interleave-mode = optimal
reboot-on-error = true
error-policy = diagnose
OBP.use-nvramrc? = true
OBP.auto-boot? = true
OBP.error-reset-recovery = sync

heslab-12:B>
```

Note – A value of `default` is equal to `max` in the case of the domain `diag-level` parameter that controls the domain POST level.

Of special note is the addition of the `error-policy` value in firmware versions 5.13.x, or higher. This value affects the behavior of the system if the SC detects a hardware error condition by determining whether the error message only is displayed or whether the error message is displayed with a diagnostic recommendation. It is recommended that `error-policy` be set to `diagnose`.

Configuring the Midframe Service Processor

An external system to help with the administration of Sun Fire servers is helpful because of the need to access and monitor the SC on a regular basis (console output or SC platform messages) and because the SC attempts to log messages to an external host (by using SNMP or `syslog`).

This section contains descriptions of how to configure the Midframe Service Processor (MSP):

- Configuring log messages on the MSP
- Configuring the Sun MC software
- Configuring the Sun Explorer software utility

The MSP provides a centralized and secured access point for logging these messages, and it provides support services that the SC cannot provide. While the Sun Fire platform is theoretically self-contained, for ease of problem diagnosis,

accessibility to platform information, and updating system firmware and software, an MSP is strongly recommended to provide a centralized location for these functions.

This article does not recommend any particular type of MSP because each site's needs (for instance, the number of systems to monitor and the requirement for the Sun MC software) are generally different. In addition, the requirements of individual sites might conflict. For example, `syslog(3)` does not require as many system resources to monitor hosts as does the Sun MC software. However, because of the limited number of `syslog(3)` logging facilities available per host, it might not be possible to monitor as many systems as a single, larger Sun MC server can, without generating large unmanageable log files.

▼ To Configure Log Messages on the MSP

To be able to log messages sent to the MSP with `syslog(3)` from a SC, you need to make additions to the default `/etc/syslog.conf` file on the `syslog(3)` host, as described in the following steps. The additions should correspond to the settings made on the Sun Fire platform.

1. Add the following entries to the `syslog.conf` file:

```
local0.notice    /var/log/messages.platform
local1.notice    /var/log/messages.domain-A
local2.notice    /var/log/messages.domain-B
local3.notice    /var/log/messages.domain-C
local4.notice    /var/log/messages.domain-D
```

The above entries should be separated by tabs; otherwise, `syslogd` will fail.

2. Create the log files before restarting `syslog(3)` by entering the following commands.

You must ensure that the files have the appropriate permissions.

```
nerm# touch /var/log/messages.platform
nerm# touch /var/log/messages.domain-A
nerm# touch /var/log/messages.domain-B
nerm# touch /var/log/messages.domain-C
nerm# touch /var/log/messages.domain-D
```

3. Restart the `syslogd(1M)` daemon, or reboot the MSP.

4. Verify that `syslog(3)` is working correctly by using the `logger(1)` command:

```
nerm# logger -p local0.notice "Platform test message"
nerm# logger -p local1.notice "Domain A test message"
nerm# logger -p local2.notice "Domain B test message"
nerm# logger -p local3.notice "Domain C test message"
nerm# logger -p local4.notice "Domain D test message"
```

5. Verify that the test message is logged in the appropriate log file.
6. Verify that the SC is logging properly by entering the `setkeyswitch` command:

```
heslab-12:B> setkeyswitch off
heslab-12:B> setkeyswitch on
```

7. Verify that the POST results are sent to the log files.

Periodically, you should rotate the log files to prevent them from growing too large. You can do this by setting up additional scripts such as `/usr/lib/newsyslog` to run on a regular basis. The `newsyslog` script modifies the contents of the additional scripts to rotate the specified log files. Rotate the files on a monthly basis, and keep archived copies of the information for at least two months.

- a. For the Solaris 8 OE, and earlier releases, add five entries to the existing `/usr/lib/newsyslog` file to rotate the five log files referenced above.

The following code contains an example of an entry:

```
#
LOGDIR=/var/log
LOG=messages.platform
if test -d $LOGDIR
then
    cd $LOGDIR
    if test -s $LOG
    then
test -f $LOG.6 && mv $LOG.6 $LOG.7
        test -f $LOG.5 && mv $LOG.5 $LOG.6
        test -f $LOG.4 && mv $LOG.4 $LOG.5
        test -f $LOG.3 && mv $LOG.3 $LOG.4
        test -f $LOG.2 && mv $LOG.2 $LOG.3
        test -f $LOG.1 && mv $LOG.1 $LOG.2
        test -f $LOG.0 && mv $LOG.0 $LOG.1
        mv $LOG $LOG.0
        cp /dev/null $LOG
        chmod 644 $LOG
        sleep 40
    fi
fi
```

Be sure to create five unique entries for each of the five log files by changing the name of the file in the LOG value. Also, be sure that the last line in the file remains as follows:

```
kill -HUP `cat /etc/syslog.pid`
```

- b. For the Solaris 9 OE, use the `logadm` command to set up the rotation of the log files.

The `logadm` command replaces the functionality of the `/usr/lib/newsyslog` file.

```
nerm# logadm -w /var/log/messages.platform -C 8 -a 'kill -HUP `cat
/var/run/syslog.pid`'
```

This command must be entered on one line.

Use the following command for the domains:

```
nerm# logadm -w /var/log/messages.domain-A -C 8 -a 'kill -HUP `cat  
/var/run/syslog.pid`'
```

Be sure to change the message file name for all of the domains on the system.

As mentioned, `syslog(3)` facilities are limited, so plan ahead and organize the limited number of resources to effectively enable an administrator to quickly locate data. In addition, you should establish a mechanism to parse and sort the incoming information on a regular basis and to send the administrator an email message of the changes. Further information on the configuration of `syslog` can be found in the Solaris OE system administration guides.

Configuring the Sun MC Software

A Sun MC software server normally requires a higher level of system resources, such as a correctly configured dual processor system capable of having 1 gigabyte of RAM or more. However, a Sun MC software server also has a greater capability to monitor and administer a large number of systems. Whether or not the Sun MC software proxy agent is running on the same host as the server agent might influence the Sun MC software server configuration.

The Sun MC software should be implemented with two systems. One small system should act as a proxy agent for one or more Sun Fire platforms, and the second system should be a larger Sun MC software server that is tasked with monitoring the entire network. This configuration provides additional monitoring capabilities in case the system containing the Sun MC software server becomes unavailable. It also provides flexibility in the MSP and security configuration.

To be able to monitor SNMP traps generated by the SC, you must install the Sun MC 3.0 Platform Update 1, or higher. This version of Sun MC is available with the Solaris 8 OE 04/01 release. Currently, the Sun MC software is the only package that can understand the SNMP traps generated by the SC. No MIBS are publicly available. Refer to the *Sun Management Center 3.0 Supplement for Sun Fire 6800, 4810, 4800, and 3800 Systems* for additional installation and setup information.

Preparing for Firmware Updates

For purposes of firmware updates to the SC, you must set up an FTP or HTTP service on the MSP. You can set up an anonymous FTP server by following the instructions in the `ftpd(1M)` man page, or you can use normal FTP by specifying a user and password in the FTP URL. If the MSP uses the Solaris 8 OE or higher, a

version of the Apache Web server is provided with the Solaris OE, which you can use to provide HTTP services. Because the HTTP service is more configurable than the FTP service and because it may be restricted to listen only on certain network interfaces, the HTTP can have less of a security impact than FTP. For additional information, refer to the “Securing the Sun Fire™ Midframe System Controller Updated for 5.13.x”, which is available from: <http://www.sun.com/blueprints>

You can install the operating system for Sun Fire domains either from an attached DVD-ROM drive or over the network from a Solaris JumpStart™ software server. The function of a JumpStart software server may also be well suited for an MSP. Detailed instructions for setting up a Solaris JumpStart software server can be found in the Solaris OE systems administration guides.

When choosing an appropriate MSP (or MSPs), some additional capabilities need to be considered, such as how to access the serial ports on multiple SCs and how many devices need to be monitored on the same system. For example, the Sun StorEdge™ T3 Array may also need to be monitored by the same MSP.

Configuring the Sun Explorer Software Utility

After completing the initial installation of a Sun Fire server, you should install the Sun Explorer software utility on both the server and the MSP, and you should set it up to periodically collect system configuration information and error messages. Check the following site regularly for updates to the Sun Explorer software:
<http://sunsolve.sun.com>

If possible, the output from the Sun Explorer software should be automatically sent to the Sun Explorer software database at the email address you specified when you set up the software. You should use version 3.5.3.1, or higher, because it has the capability to gather data from the SC. Version 3.5, or higher, can collect data from the SC; however, version 3.5.3.1, or higher, is recommended because it includes the ability to gather information about the components in the system, such as the field-replaceable unit ID (FRUID). The FRUID information is stored in each component in a Sun Fire system. It contains information about the parts, such as the revision level, serial number, and manufacturing information.

The following command gathers information from the SC. You should execute it on the MSP. In addition, use of the following command assumes that the Sun Explorer software has already been installed on the system in the default location:
`/opt/SUNWexplo`

```
nerm# /opt/SUNWexplo/bin/explorer -w fru,sextended,default
```

If you execute this command from an MSP, the Explorer software will collect data from the MSP and the SC. To collect data from the SC, the Explorer software uses a telnet connection; therefore, the MSP must be able to establish a telnet session on the SC.

If security considerations prevent the automatic sending of the Sun Explorer software results to the Sun Explorer software database, you should still install the Sun Explorer software utility so that it is available to collect information in the event that service is required on the system and information needs to be collected.

The initial installation is also a good time to record and check the system serial number, `hostid`, and MAC address provided with the system and to become familiar with how these values are reported by the SC `showplatform -v` command. Keep this information where it can be easily accessed in case a SC replacement is required.

Platform and Domain Administration

The SC has a administration scheme in which operations affecting the entire system are administered through a *platform* shell and operations affecting separate domains are administered through a *domain* shell.

This section contains general descriptions of administering the SC and detailed descriptions of the following topics:

- “Monitoring Through the Serial Port, `syslog`, and SNMP”
- “Setting Up Information Recording and Logging”

You can access multiple platform shells simultaneously. The platform shell can view the status of any component within the system and can also control its allocation. One example of how the platform shell is used to manage the entire system is through access control lists (ACLs). You can set up ACLs by using the `setupplatform -p acl` command. You can also use the ACLs to restrict domain access to resources.

While the platform shell manages and administers overall system resources, domain specific operations, such as the turning of the virtual keyswitch, are controlled by the domain shell. The domain shell can only access resources specified in the ACL set up for it by the platform shell, and only one shell per domain can be active at any time. The ACL restricts the domain shell so that it views only the resources that the domain is currently using, resources that are allocated to the domain, or any resources that are unassigned on the platform and are available to the domain according to the ACL.

One advantage to this setup is that you can restrict the access to the platform shell (and administration of the overall system resources) to a group of administrators, who are separate from a group of administrators for the domains. You can control access to platform and domain shells by using passwords that you can set and change by using the `password` command on the SC. From the platform shell, you can set or change the platform and domain shell passwords. From a domain shell, you can change only the password of the particular domain. Because the platform has additional privileges, its password should be different from those selected for the domains.

Monitoring Through the Serial Port, `syslog`, and SNMP

Administration of a Sun Fire server is designed to be performed primarily through the SC, which can be accessed in two ways: RS-232 serial port and 10/100BASE-T Ethernet. Both ports should remain available at all times, and you should monitor the platform console and all of the domain consoles.

After the initial platform setup, the serial port performs an important role in the administration of the Sun Fire server. The system in the event of a network outage and the location where the system outputs platform and SC error messages to the serial port. While output is buffered (up to 4 kilobytes) and can be directed to a `syslog` host if a network is configured, you should continuously monitor the serial port output either with a logging terminal server or by connecting the serial port to an MSP, which can capture the terminal output with a mechanism such as `script(1)` in the Solaris™ Operating Environment (Solaris OE). This mechanism provides additional error detection and information in the event of a failure. Finally, network connectivity to the SC will be lost during firmware updates because the SC is rebooted. Thus, serial port access is critical to monitor the progress of the update procedure and for diagnostics, should a problem arise.

For routine tasks, the preferred method of accessing the SC is the Ethernet interface. In addition to providing multiple high-speed shell connections, the Ethernet interface on the SC enables you to set up `syslog(3)` and SNMP messages to be sent to a designated administration platform. The Ethernet interface is required for performing firmware updates on the system and for saving and restoring the SC configuration information. You should enable and configure both `syslog` and SNMP facilities; however, you should record system consoles by using a mechanism such as `script(1M)` because not all of the messages can be logged with `syslog` or SNMP.

Note – SNMP should not be enabled unless a Sun MC server has been configured to support the system.

Setting Up Information Recording and Logging

You should perform the setup of information recording and logging during the initial platform setup. Although you can also make these changes at other times, you should set up `syslog(3)` to record the information to a central location so that information can be quickly located in the event of a problem and to prevent the SC message buffer from being overwritten.

Because the SC message buffer is in volatile memory, messages can be lost if the SC loses power. The SC maintains a 4 kilobyte ring buffer for messages from each domain and the platform. It also maintains an additional buffer to record particular failure messages for retrieval by authorized service personnel.

When setting up the platform or domain using the `setupplatform` or `setupdomain`, respectively, you are prompted for a `syslog(3)` log host. You can supply a `syslog(3)` log host by using an IP address or hostname, as well as a facility level.

Note – For older firmware, you will not see the Log Facility prompt, as shown in the following example of how to set up log host. The syntax for the log hostname is `hostname:facility` (for example, `mysysloghost:local0`).

The following shows an example of how to set up the log host.

```
heslab-16-sc0:SC> setupplatform -p loghost

Loghosts
-----
Loghost [10.1.63.251]: loghost_name
Log Facility [local0]:

heslab-16-sc0:SC>

heslab-16-sc0:B> setupdomain -p loghost

Loghosts
-----
Loghost [ ]: loghost_name
Log Facility [local0]: local1

heslab-16-sc0:B>
```

Corresponding changes need to be made to the `/etc/syslog.conf` file on the `syslog(3)` host or on the MSP. You can find further information on the configuration of `syslog(3)` in the Solaris OE system administration guides in the Solaris OE System Administrator AnswerBook2™ collection.

Based on the number of systems and `syslog(3)` devices that a single `syslog(3)` host will be monitoring, you should establish a convention to maximize use of the limited number of `syslog(3)` facilities. There are only eight `syslog(3)` facilities available for users in the Solaris OE, so it is likely that an administrator will quickly run out of unique `syslog(3)` facilities. Organization of message logging is important to enable administrators to quickly find the desired information. A good way to organize `syslog(3)` logging is to assign the `local0` facility to all platform messages and then assign `local1-4` to domains A through D, respectively. The `syslog(3)` facility in the Solaris 8 OE identifies each `syslog(3)` entry with the originating hostname and `syslog(3)` facility. This identification makes it easy to quickly separate messages coming from different hosts.

When setting up the platform, you can configure the SC to interface with the latest versions of the Sun MC software through SNMP. This improves the monitoring and administration capabilities of the platform. It is strongly recommended that the default community strings be changed during installation for security reasons. The following values for platform and domain public and private community strings are set by default.

```
Platform Public: P-public
Platform Private: P-private

Domain A Public: A-public
Domain A Private: A-private

Domain B Public: B-public
Domain B Private: B-private

Domain C Public: C-public
Domain C Private: C-private

Domain D Public: D-public
Domain D Private: D-private
```

SNMP must be enabled on the platform by using the `setupplatform` command before SNMP can be enabled on any of the domains. The following shows an example of the `setupplatform` command.

```
heslab-16:SC> setupplatform -p snmp

SNMP
----
Platform Description [Sun Fire 3800]:
Platform Contact [email_address]:
Platform Location [Lab]:
Enable SNMP Agent? [no]: yes
Trap Hosts [ ]: 10.1.221.11
Public Community String [ ]: P-public
Private Community String [ ]: P-private
```

For more information on how to set up SNMP, refer to the security articles by Alex Noordergraaf at the Sun BluePrints Online website at:

<http://www.sun.com/blueprints/>

The following shows an example of the `setupdomain` command.

```
heslab-16:A> setupdomain -p snmp

SNMP
----
Domain Description [test]:
Domain Contact [email_address]:
Trap Hosts [ ]: 10.1.221.11
Public Community String [ ]: A-public
Private Community String [ ]: A-private

heslab-16:A>
```

The port for the Trap Hosts value can be entered in the form of `hostname:port` in firmware 5.13.0, or higher. Do not change the port setting unless specifically instructed to do so for the installation of other software on the trap host. To find additional information on configuring the Sun MC software, refer to the Sun MC software documentation at the following site:

<http://docs.sun.com>

In addition to setting up `syslog(3)` and SNMP, you should monitor domain console sessions in a manner similar to that described for the platform and the serial port connection. While the SC has a buffer for each domain shell's messages, the SC will

not send domain console messages or error messages generated by the Solaris OE (such as panic strings) to an external log host. Therefore, if you do not constantly monitor the domain consoles, critical messages and valuable diagnostic information could be lost in the event of a failure. With multiple domains to monitor, you should access the domain shells through the Ethernet port because it allows multiple connections.

Platform Security

System security is important for any computing system, and the Sun Fire server is no exception. This section contains descriptions of the following basic platform security topics:

- “Recommendations for User Authorization”
- “Telnet Sessions”
- “Keyswitch Settings”

The following sections contain general concepts of platform security. For detailed information, refer to the security documents in “References” on page 27.

Because the Sun Fire domains run the same Solaris OE as other systems, basic security practices that apply to any Solaris OE system also apply to the Sun Fire servers. These practices include regular patch maintenance, stopping unnecessary network services, and choosing good passwords to prevent account abuse. However, with the SC, the architecture of the Sun Fire server causes additional security considerations. In addition, because the SC is key to the administration and operation of a Sun Fire server, the security capabilities of the SC are more limited.

Great care should be taken in the setup of the system to ensure that access is restricted only to authorized personnel. Failure to properly secure access to the SC can adversely affect the operation of the Sun Fire server. Refer to the Sun BluePrints OnLine article “Securing the Sun Fire™ Midframe System Controller (Update for 5.13.x)” for more information.

Recommendations for User Authorization

To help deter unauthorized access, passwords should be set on the SC platform and domain shells. You can set these passwords by using the `password` command with the `-d domain` option from either the SC platform or domain shells. The `password` command issued from a domain shell, can only be used to change the password for that particular domain.

The SC does not enforce any password standards, and it maintains no records of failed login attempts or the source of the login attempts. Given the importance of these passwords, especially in terms of restricting access to critical system resources, chose passwords that cannot be easily guessed or discovered using a brute-force attack. Passwords for the SC can and should be longer than eight characters. It is strongly recommended that passwords for platform access and superuser (`root`) access on the domains be different.

Serial Port Access

It is extremely important to carefully control access to the SC serial port. The serial port is the lowest level of access to the SC, so an unprotected serial port could have serious consequences to the operation of the Sun Fire system because access to the serial port can compromise the application that runs on the SC. Because that application controls the entire Sun Fire system, improper access could result in undesired changes to critical settings or in system outages. Attach the serial port connection of the SC to a password-controlled terminal server or directly to the MSP where access can be monitored and logged.

Telnet Sessions

There is no current capability for session encryption between a host and the SC through the Ethernet interface. Because telnet sessions are used to make connections to the SC, maintaining a secure network is of the utmost importance. It is strongly recommended that the MSP and the SCs be placed on a private, switched, nonrouted network. The MSP should be the only way to access the SCs, and access to the MSP should be carefully secured, monitored, and encrypted if possible. This includes using a terminal server that supports encryption (for example, Secure Shell) if possible. If the terminal server does not support encryption, be sure to place the terminal server on the private network, not on the public network.

In cases where corporate security policy dictates the use of encrypted sessions on all networks, you can connect the Ethernet interface on the SC directly to a secured workstation with an Ethernet cross-over cable. Then, you can require encrypted sessions between other hosts and the secured workstation. Due diligence should also be exercised to keep the MSP and SC up-to-date in terms of patches.

To further enhance the security of the MSP, you should use a tool such as the Solaris™ Security Toolkit (formerly known as *JASS*) to install and improve the security profile of the MSP. You can obtain the Solaris Security Toolkit from:

<http://www.sun.com/security/jass>

While it is advantageous to set up an `httpd` web server and/or an anonymous FTP server on the MSP to facilitate firmware updates to the SC, both of these network services have traditionally been common sources for security issues. Because a compromised MSP could also compromise the SC and the entire Sun Fire platform, it is important to pay close attention to the setup of these services, to the security patches applied for these services, and to the access to the MSP, which must be carefully configured.

With SC firmware 5.13.0, and higher, you can control the behavior of Telnet access to the SC. The following shows an example of how to control Telnet access.

```
heslab-16-sc0:SC> setupplatform -p security

Security Options
-----
Enable telnet servers? [yes]: yes
Idle connection timeout (in minutes; 0 means no timeout) [0]: 5

heslab-16-sc0:SC>
```

It is preferable to access the SCs through the Ethernet port using a Telnet connection. Thus, Telnet services should remain enabled unless there is a requirement to turn it off. Inactive connections should be disconnected after some period of time because they potentially represent a security problem.

Keyswitch Settings

During normal operations, it is recommended that the virtual domain keyswitch be set to *secure* by using the following command:

```
heslab-16-sc0:A> setkeyswitch secure
```

Setting the keyswitch to *secure* prevents firmware updates to I/O and system boards in the domain. It also prevents an operator from sending a break command to the running domain and accidentally terminating the Solaris OE. The keyswitch needs to be changed to the *on* position before sending a break command or updating firmware. Only domain administrators with access to the domain shell can set the keyswitch to the *secure* position.

Error Analysis and Diagnosis

Sun Fire servers provide significantly enhanced diagnostics capabilities. In the event of a system fault, the system should provide data for both software and hardware failures that you can use to help determine the source of the fault. Errors can be generated and logged to several places, depending on the type of error. Use a utility such as Explorer to gather data from the system so that all error messages can be collected in a central location.

After the appropriate error messages have been located, use the flow charts in the *Sun Fire 6800/4810/4800/3800 Systems Troubleshooting Manual* to isolate the source of the error, as far as possible. Based on the results, attempt to verify the failure using component blacklisting, segmenting, or other reconfiguration before attempting to remove or replace components in the case of a suspected hardware problem.

Maintenance Functions

There are maintenance functions that you must perform on a regular basis. The following functions are described in this section:

- “Restoring the Sun Fire SC Configuration”
- “Updating the Firmware and Real Time Operating System”
- “Removing the SC from Platform Use”

Restoring the Sun Fire SC Configuration

If an SC fails, you might need to manually restore the SC configuration information. After the configuration of the platform has been completed, including setting up domains and segments, create a backup of your SC configuration so that a quick restoration will be possible.

The following shows an example of how to create a backup of the Sun Fire SC configuration.

```
heslab-12:sc> dumpconfig -f ftp://me:passw0rd@heslab-05/dumps
```

The following shows an example of how to restore a Sun Fire SC configuration.

```
heslab-12:sc> restoreconfig -f ftp://me:passwd@heslab-05/dumps
```

You should create a backup of the SC configuration on a routine basis to ensure that the dump file is up-to-date. To help with a quick recovery if a primary SC fails, make sure that the secondary SC has the same configuration information as the primary after the domains have been configured.

Updating the Firmware and Real Time Operating System

Periodically, updates to the SC firmware and RTOS will be made available. These updates often contain critical bug fixes and functionality enhancements to the SC and should be applied as part of a regular patch maintenance routine.

Before applying a firmware update using the `flashupdate` command, carefully read the release notes and `Install.info` files in the patch package before proceeding with the update to familiarize yourself with the procedures. Backing up the SC configuration before updating is also recommended.

Also be sure to perform the firmware updates regularly, and for Sun Fire systems which have two SCs, remember to update the firmware on both SCs.



Caution – You must follow the instructions in the `Install.info` file, included with each patch release, to ensure that both ScApp and RTOS are updated together. ScApp should only be run with the accompanying version of RTOS. Upgrading from some versions of the firmware may require that the upgrade to the SCs be done in a specific order. Be sure to read and follow the instructions carefully.

You can retrieve updates from the SunSolveSM program website:

<http://sunsolve.sun.com>

You should also have copies of important SC parameters that are displayed by the `showplatform` and `showboards` commands, as well as those displayed by the OpenBootTM PROM commands `printenv` and `devalias`. See “Appendix: SC Parameters” on page 28 for examples of the output.

Refer to the *Sun Fire 6800/4810/4800/3800 System Controller Command Reference Manual* for more information on the commands discussed here.

Removing the SC from Platform Use

If an SC needs to be removed for maintenance purposes, you must follow the instructions for SC replacement that are specific for the version of firmware on the SCs. For specific instructions, refer to the *Sun Fire 6800/4810/4800/3800 Systems Platform Administration Manual* for your version of firmware.

In general, an SC should never be removed from a system unless the SC can be powered off, either by using the `poweroff SSCx` command or by removing the power to the entire platform.

References

The following sources were referenced in this article:

- *Sun Fire 6800/4810/4800/3800 System Controller Command Reference Manual*
- *Sun Fire 6800/4810/4800/3800 Systems Service Manual*
- *Sun Management Center 3.0 Software Installation Guide*
- *Sun Management Center 3.0 Supplement for Sun Fire 6800, 4810, 4800, and 3800 Systems*
- “Securing the Sun Fire Midframe System Controller (Updated for SC Firmware 5.13.0)”, Sun BluePrints Online, September 2001
- *Sun Field Engineer Handbook*
- Global Knowledge Engineering Explorer Tool at:
<http://sunsolve.sun.com>
- Solaris Security Toolkit (formerly known as *jass*) at:
<http://www.sun.com/jass/>

About the Author

James Hsieh is a member of the Sun CPR Engineering group and is part of the team responsible for providing engineering support of Sun's midrange server line. Prior to his current role with CPR Engineering, James worked for Sun Enterprise Services supporting mission critical customers. Prior to Sun, James worked for over eleven years with UNIX and Sun systems as a software engineer and as a systems administrator for large groups of UNIX systems.

Ordering Sun Documents

The SunDocsSM program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals through this program.

Accessing Sun Documentation Online

The `docs.sun.com` web site enables you to access Sun technical documentation online. You can browse the `docs.sun.com` archive or search for a specific book title or subject. The URL is: <http://docs.sun.com/>

To reference Sun BluePrints OnLine articles, visit the Sun BluePrints OnLine Web site at: <http://www.sun.com/blueprints/online.html>

Appendix: SC Parameters

You should also have copies of important SC parameters that are displayed by the `showplatform` and `showboards` commands, as well as the output from the OpenBootTM PROM commands `printenv` and `devalias`.

Here is the output of the printenv command

```
heslab-16-sc0:SC> console a

Connected to Domain A

{0} ok printenv
Variable                Value                (Default Value)
-----
nvramrc
secondary-diag?         false                (false)
auto-boot-timeout       512                  (512)
scsi-initiator-id       7                    (7)
powerfail-time          0                    (0)
fcode-debug?           true                 (true)
output-device           /sgcn                (/sgcn)
input-device            /sgcn                (/sgcn)
load-base               0x4000               (0x4000)
boot-command            boot                 (boot)
auto-boot?              true                 (true)
error-reset-recovery    sync                 (sync)
diag-file               ()                   ()
boot-file               ()                   ()
diag-device             disk diskifp diskglm diskc net (disk diskifp diskglm diskc
net)
boot-device             disk diskifp diskglm diskc net (disk diskifp diskglm diskc
net)
local-mac-address?     false                (false)
ansi-terminal?         true                 (true)
screen-#columns         80                   (80)
screen-#rows            34                   (34)
silent-mode?           false                (false)
use-nvramrc?           false                (false)
oem-logo                ()                   ()
oem-logo?              false                (false)
oem-banner              ()                   ()
oem-banner?            false                (false)
security-mode           ()                   ()
security-password       *****
security-#badlogins     ()                   ()
hardware-revision       ()                   ()
last-hardware-update    ()                   ()
#power-cycles           ()                   ()
diag-switch?           false                (false)
{0} ok
```

Here is the output of the devalias command:

```
{0} ok devalias

Alias                Device Path
-----
xdisk                /ssm@0,0/pci@d,600000/pci@1/SUNW,qlc@5/fp@0,0/disk@0,0
xnet                 /ssm@0,0/pci@d,700000/pci@1/SUNW,qfe@3,1
diskc                /ssm@0,0/pci@lc,700000/pci@1/SUNW,isptwo@4/sd@0,0
netc                 /ssm@0,0/pci@lc,700000/pci@1/SUNW,hme
cdrom                /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo@4/sd@6,0:f
disk                 /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo@4/sd@0,0
disk0                /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo@4/sd@0,0
disk1                /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo@4/sd@1,0
disk2                /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo@4/sd@2,0
disk3                /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo@4/sd@3,0
disk4                /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo@4/sd@4,0
disk5                /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo@4/sd@5,0
def-cn               /sgcn
diskifp              /ssm@0,0/pci@19,700000/SUNW,ifp@1/ssd@0,0
diskglm              /ssm@0,0/pci@18,700000/scsi@2/disk@0,0
scsi                 /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo
net                  /ssm@0,0/pci@18,700000/pci@1/SUNW,hme
{0} ok
```