

Sun™ HPC 2.0 SCI Guide



THE NETWORK IS THE COMPUTER™

Sun Microsystems Computer Company

A Sun Microsystems, Inc. Business
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
650 960-1300 fax 650 969-9131

Part No.: 805-1561-10
Revision A, November 1997

Copyright 1997 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303 USA. 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 U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook, SunDocs, Solaris, OpenWindows, Sun HPC Software, Ultra HPC, Ultra HPC Cluster, UltraSPARC, Sun Performance WorkShop Fortran, and Sun Performance Library are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

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 1997 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303 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 déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook, SunDocs, Solaris, OpenWindows, Sun HPC Software, Ultra HPC, Ultra HPC Cluster, UltraSPARC, Sun Performance WorkShop Fortran, et Sun Performance Library 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.

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 REPRENDRE 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.



Contents

Preface v

1. Preparing for Installation 1-1

- 1.1 Required Documentation 1-1
- 1.2 Hardware Overview 1-2
 - 1.2.1 Cluster Topologies 1-2
 - 1.2.2 2-Node SCI Link 1-3
 - 1.2.3 SCI Switch Link 1-4
 - 1.2.4 SCI SBus Card Scrubber Jumpers 1-5
 - 1.2.5 Limits on Cluster Size 1-6
- 1.3 FRU Part Numbers 1-7
- 1.4 Upgrade Kits 1-7
- 1.5 Software Requirements 1-8
- 1.6 Software Driver Overview 1-8
- 1.7 `sm_config` Configuration Utility 1-8
 - 1.7.1 Caution for NIS+ Users 1-9
- 1.8 Addresses for the Private Network 1-10

2. Task Maps 2-1

- 2.1 Standalone Sun HPC Server 2-2
- 2.2 2-Node Switchless Cluster 2-3

- 2.3 SCI Switch Cluster 2-4
- 2.4 Hardware and Software Upgrade Kit 2-5
- 2.5 Software Upgrade Kit 2-6

3. Installation Procedures 3-1

- 3.1 Installing a Standalone Sun HPC Server 3-1
- 3.2 Installing a Switchless Two-Node Cluster 3-2
- 3.3 Installing an SCI Switch Cluster 3-3
- 3.4 Installing a Sun HPC Hardware and Software Upgrade Kit 3-4
- 3.5 Installing a Sun HPC Software Upgrade Kit 3-6

4. SCI Interface Troubleshooting 4-1

- 4.1 General Verification Tests 4-1
- 4.2 SCI Switch 4-1
 - 4.2.1 General Hardware Inspection 4-1
 - 4.2.2 SCI Switch Status LED Locations 4-2
 - 4.2.3 Port Status LEDs 4-2
 - 4.2.4 General Switch Status LED 4-3
 - 4.2.5 The `get_ci_status` Command 4-3
 - 4.2.6 Client Net Failure 4-4
- 4.3 Incorrect Software Configuration 4-5
- 4.4 Incorrect Firmware 4-5

A. Man Pages A-1

- A.1 `sm_config(1m)` A-1
- A.2 `get_ci_status(1m)` A-6

B. SMA Messages B-1

- B.1 Message Directory B-1
- B.2 SMA Messages B-2

Index Index-1

Preface

The *Sun™ HPC 2.0 SCI Guide* is intended for experienced system administrators.

Using UNIX Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris 2.x Handbook for SMCC Peripherals*
- AnswerBook™ online documentation for the Solaris™ 2.x software environment
- Other software documentation that you received with your system

Typographic Conventions

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output.	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output.	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Command-line variable; replace with a real name or value.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be root to do this. To delete a file, type <code>rm filename</code> .

Shell Prompts

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

TABLE P-3 Related Documentation

Application	Title	Part Number
All	<i>Sun HPC Software 2.0 Release Notes</i>	805-2191-10
Installation	<i>Sun HPC Software 2.0 Installation Guide</i>	805-1562-10
All	<i>Sun HPC Software 2.0 System Administrator's Guide</i>	805-1554-10
LSF	<i>LSF User's Guide</i> <i>LSF User's Quick Reference</i> <i>LSF Administrator's Guide</i> <i>LSF Administrator's Quick Reference</i> <i>LSF Programmer's Guide</i>	No Sun part numbers are associated with LSF documentation.

Ordering Sun Documents

SunDocsSM is a distribution program for Sun Microsystems technical documentation. Contact SunExpress for easy ordering and quick delivery. You can find a listing of available Sun documentation on the World Wide Web.

TABLE P-4 SunExpress Contact Information

Country	Telephone	Fax
Belgium	02-720-09-09	02-725-88-50
Canada	1-800-873-7869	1-800-944-0661
France	0800-90-61-57	0800-90-61-58
Germany	01-30-81-61-91	01-30-81-61-92
Holland	06-022-34-45	06-022-34-46
Japan	0120-33-9096	0120-33-9097
Luxembourg	32-2-720-09-09	32-2-725-88-50
Sweden	020-79-57-26	020-79-57-27
Switzerland	0800-55-19-26	0800-55-19-27

TABLE P-4 SunExpress Contact Information (*Continued*)

United Kingdom	0800-89-88-88	0800-89-88-87
United States	1-800-873-7869	1-800-944-0661
World Wide Web: http://www.sun.com/sunexpress/		

Sun Documentation on the Web

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

Sun Welcomes Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. You can email or fax your comments to us. Please include the part number of your document in the subject line of your email or fax message.

- Email: `smcc-docs@sun.com`
- Fax: SMCC Document Feedback
1-650-786-6443

Preparing for Installation

This manual is a supplement to the standard hardware installation and service manuals for your system. It also supplements the *Sun HPC Software 2.0 Installation Guide*, Part Number 805-1562.

1.1 Required Documentation

Locate the standard installation and service manuals for your server before beginning the installation or upgrade. TABLE 1-1 lists the pertinent installation and service manuals by part number. Your system documentation set may include other manuals that are also pertinent to the installation.

TABLE 1-1 Manuals Required for Installation

Server	Installation Manual	Service Manual
Sun Ultra HPC 2	802-5934 (US version) or localized equivalent	802-2561
Sun Ultra HPC 3000	802-6050	802-6051
Sun Ultra HPC 4000	802-3844	802-3845
Sun Ultra HPC 5000	802-3844	802-3845
Sun Ultra HPC 6000	802-3844	802-3845
Sun Ultra HPC 10000	805-0309	805-0311

1.2 Hardware Overview

A Sun™ Ultra™ HPC server can be used alone or it can be configured as part of a cluster. Sun Ultra HPC Clusters use Scalable Coherent Interface (SCI) interconnects for high-speed communications within the cluster.

Note – Sun Ultra HPC Clusters use only a single SCI SBus interface card to support SCI interconnections. This is unlike SCI implementation in Ultra Enterprise clusters, which use two SCI SBus interface cards per node. The information in this manual applies to Ultra HPC Clusters only.

1.2.1 Cluster Topologies

The nodes in Sun Ultra HPC Clusters are standard Sun servers but have additional high-speed private interfaces that allow the nodes to share information with each other and process data in parallel.

The nodes are normally also connected to standard networks to receive and queue job requests.

Sun HPC Software supports two cluster topologies:

- *Topology A* – two-node configuration (see Section Section 1.2.2, “2-Node SCI Link”)
- *Topology B* – switched configuration for up to four nodes (see Section Section 1.2.3, “SCI Switch Link”)

1.2.2 2-Node SCI Link

A two-node HPC cluster can be linked either directly (topology A) or through an SCI switch (topology B).

In topology A, each node must have an SCI SBus card. The SCI SBus cards are connected directly (FIGURE 1-1) by an 80-line SCI station cable.




FIGURE 1-1 Two-Node SCI-to-SCI Link (Topology A)

Note – One card in a direct node-to-node link must have the scrubber jumper set on. The other card must have its scrubber jumper in the off position. See Section 1.2.4, “SCI SBus Card Scrubber Jumpers,” for additional details.

1.2.3 SCI Switch Link

In topology B, an SCI switch (FIGURE 1-2) can connect up to four nodes. Each node requires an SCI SBus card and an SCI station cable.

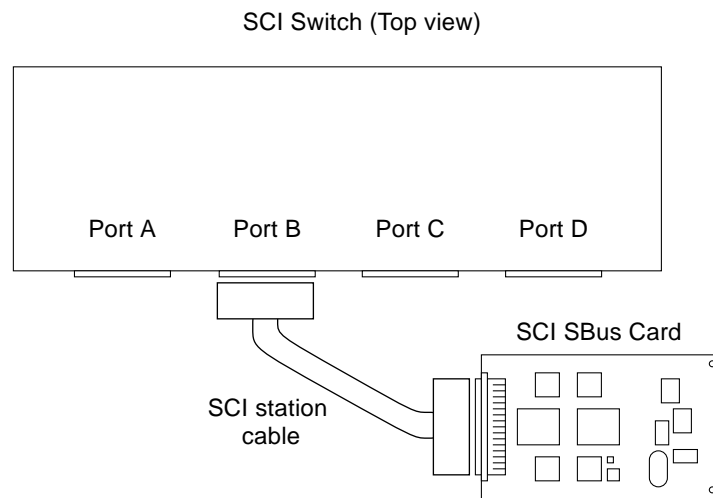



FIGURE 1-2 Basic SCI Switch Link (Topology B)

Note – In a switched HPC cluster, every SCI SBus card must have its scrubber jumper set on. See Section 1.2.4, “SCI SBus Card Scrubber Jumpers,” for additional details.

1.2.4 SCI SBus Card Scrubber Jumpers

Each SCI SBus card has a jumper called a scrubber jumper (FIGURE 1-3). The scrubber circuit performs link maintenance functions.



SCI SBus Card (Top view)

FIGURE 1-3 Location of the Scrubber Jumper



Caution – Examine the jumper setting on each SBus interface card before installing or replacing it. These cards are shipped without a default setting. If scrubber jumpers are set incorrectly when they are installed, communication between nodes can experience intermittent faults. See TABLE 1-2 for details.

TABLE 1-2 Scrubber Jumper Settings

Topology	SCI SBus Card Jumper Settings
Topology A, node-to-node	Two SBus cards: one jumper On, one jumper Off
Topology B, switch	All SBus cards: the jumper is set On

1.2.5 Limits on Cluster Size

The cable lengths determine the maximum distances between nodes in a cluster. The maximum length for an SCI interface cable is 10 meters. Shorter lengths are also available (TABLE 1-3).

TABLE 1-3 Part Numbers for SCI Cables

Length	Part Number
2 meters	F530-2360
5 meters	F530-2361
10 meters	F530-2362

To conserve space, desktop HPC nodes can be grouped on rack shelves or a tabletop, in which case shorter cables are more practical.

An SCI switch supports a cluster of two, three, or four nodes. See TABLE 1-4 for a list of the server types:

TABLE 1-4 Maximum Cluster Configurations

Server Type	Maximum Number of Servers
Sun Ultra HPC 2	4
Sun Ultra HPC 450	4
Sun Ultra HPC 3000	4
Sun Ultra HPC 4000	4
Sun Ultra HPC 5000	4
Sun Ultra HPC 6000	4
Sun Ultra HPC 10000	4

1.3 FRU Part Numbers

TABLE 1-5 lists the field-replaceable units (FRUs) that are common to HPC clusters. For other FRUs, refer to the service manual for the specific server model.

TABLE 1-5 Hardware Part Numbers

SCI SBus adapter card	F370-2868
SCI 4-port switch for switch-based configuration	F370-3146
SCI interface cable, 2-meter length	F530-2360
SCI interface cable, 5-meter length	F530-2361
SCI interface cable, 10-meter length	F530-2362

1.4 Upgrade Kits

Kits are available for upgrading Sun servers to HPC operation. Kits containing software only or software and hardware are available.

The server to be upgraded must meet the requirements outlined in TABLE 1-6.

TABLE 1-6 Upgrade Requirements

Category	Requirement
Server type	Sun Ultra2 Sun Ultra Enterprise 3000 Sun Ultra Enterprise 4000 Sun Ultra Enterprise 5000 Sun Ultra Enterprise 6000
Hardware	Must have an open SBus slot for an SCI card if server is to be part of a cluster
Software	Solaris 2.5.1 or 2.6.

1.5 Software Requirements

To support SCI operations, all the nodes in a Sun HPC Cluster must have

- Solaris 2.5.1 or 2.6
- The SUNWsci package
- The SUNWscid package
- The SUNWsma package

1.6 Software Driver Overview

The following packages comprise the software component of an SCI-based cluster interconnect:

- SUNWsci – the SCI hardware driver, provides an API to client drivers to handle management of shared-memory and message transfers between cluster nodes.
- SUNWscid – the SCI DLPI driver, provides TCP/IP functionality over the SCI interconnect. The SCI-based TCP/IP network interfaces behave similarly to Ethernet-based interfaces.
- SUNWsma – switch management software provides management of the SCI switch (if used), including SCI session management and various link and switch states.

1.7 sm_config Configuration Utility

`sm_config` is the SCI adapter configuration utility. It acts as a client of `sm_configd(1M)` daemon. `sm_config` contacts the `sm_configd(1M)` daemon on all nodes and works in a distributed fashion to retrieve the adapter inventory and configure the adapter cards on these nodes.

The IP addresses for all the nodes in the cluster are defined by `sm_config` using an equation based on the number of adapter cards and other topology information contained in a configuration template file that has been edited to represent the particular SCI topology that has been implemented.

Three SCI configuration templates are provided to simplify the editing task. These templates, called `sma2.hpc`, `sma3.hpc`, and `sma4.hpc`, are partially edited to represent two-node, three-node, and four-node SCI configurations—that is, the three SCI configurations supported by Sun HPC Software 2.0. These templates are described in Chapter 3 of the *Sun HPC Software 2.0 Installation Guide*.

Upon successful completion, a configuration file named `/etc/sma.config` is installed on all the nodes in the cluster. This file contains a snapshot view of the cluster members, switches and adapters. It also installs a file called `/etc/sma.ip`, which contains the IP addresses of all the SCI interfaces in the cluster.

The procedure for running `sm_config` is described in Chapter 7 of the *Sun HPC Software 2.0 Installation Guide*.

1.7.1 Caution for NIS+ Users

`sm_config` may have trouble contacting other node in the cluster in an NIS+ environment. By default, the NIS+ version of `/etc/nsswitch.conf` specifies the `services` entry as: `nisplus [NOTFOUND=return] files`. Since the `/etc/services` file is modified and used by `SUNWsma` and other packages, the `/etc/nsswitch.conf` entry should be as follows:

```
services:    files nisplus
```

Note – Place the term `files` first before other entries.

1.8 Addresses for the Private Network

The SCI connections from node to node, or node to SCI switch, form a private network. The private network addresses consist of 32 bits, of which the first three octets (24 bits) are user-selectable. The last octet is reserved to denote the SCI SBus card.

The private network addresses are based on the network IP address (shown in the configuration template as 204.68.34). The last eight bits of the private network addresses are added to the 24-bit network IP address.

For example, for topology A (FIGURE 1-4), the cards are numbered 1 and 2, and the corresponding private net addresses are 204.68.34.1 and 204.68.34.2.




FIGURE 1-4 Two-Node Configuration (Topology A)

If there are four cards in the system (the maximum configuration for topology B), the cards are numbered 1 through 4, and the corresponding private net addresses are 204.68.34.1, 204.68.34.2, 204.68.34.3, and 204.68.34.4.




FIGURE 1-5 Maximum Switch Configuration (Topology B)

Task Maps

In this guide, the hardware installation procedures vary according to the different Sun HPC Cluster configurations.

The following task maps describe individual situations and list specific sources of information that you will need to install:

- Standalone Sun HPC Server
- Two-node switchless cluster
- SCI switch cluster
- Upgrade kit containing Sun HPC hardware and software
- Upgrade kit containing Sun HPC Software only

2.1 Standalone Sun HPC Server

This task map is for a single server that is not part of a Sun HPC Cluster. Refer to the specific installation manual for the server (TABLE 1-1). This manual does not contain any modifications to the normal installation procedures.

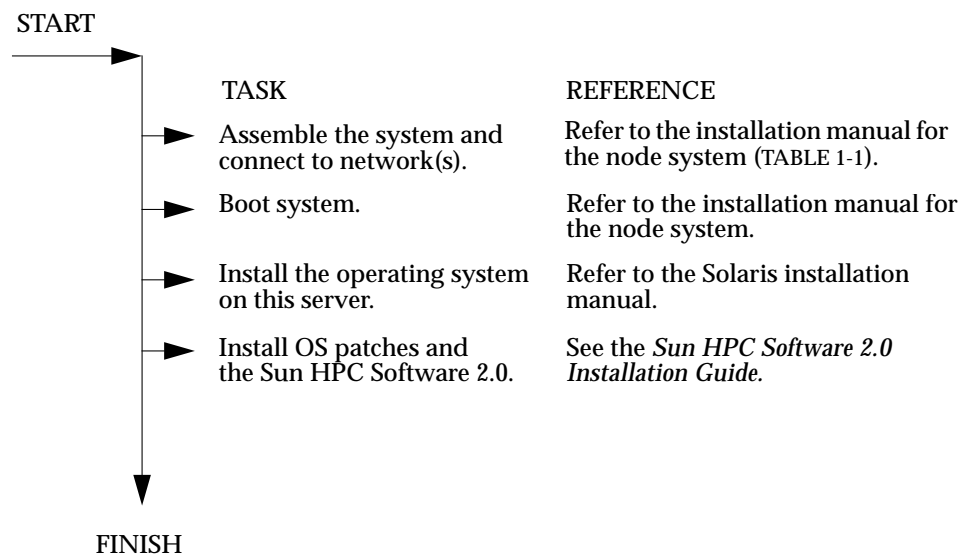


FIGURE 2-1 Task Map for a Standalone Sun HPC Server

2.2 2-Node Switchless Cluster

This task map is for a cluster in which one node connects directly to the second node (topology A in FIGURE 1-1). For instructions, see Section 3.2 “Installing a Switchless Two-Node Cluster” on page 3-2.

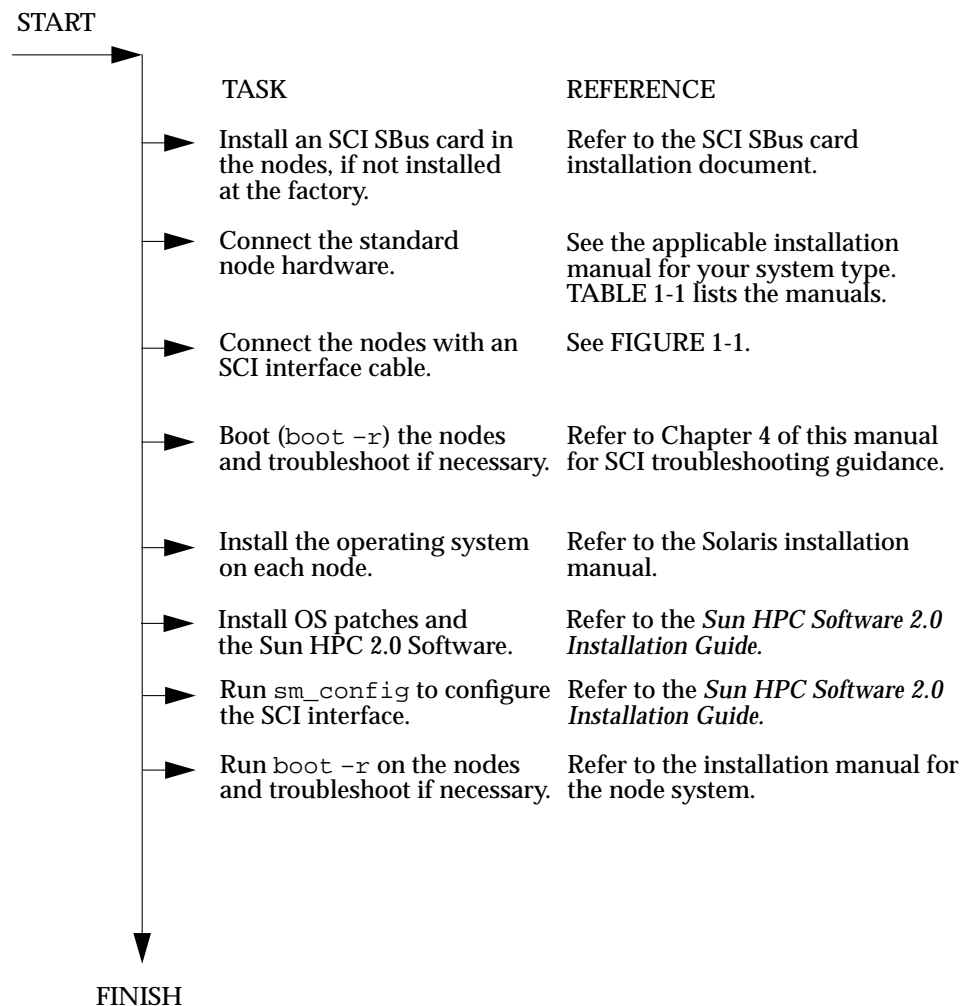


FIGURE 2-2 Task Map for a 2-Node Cluster

2.3 SCI Switch Cluster

This task map is for a cluster in which the nodes connect directly to an SCI switch (topology B in FIGURE 1-2). For instructions, see Section 3.3 “Installing an SCI Switch Cluster” on page 3-3.

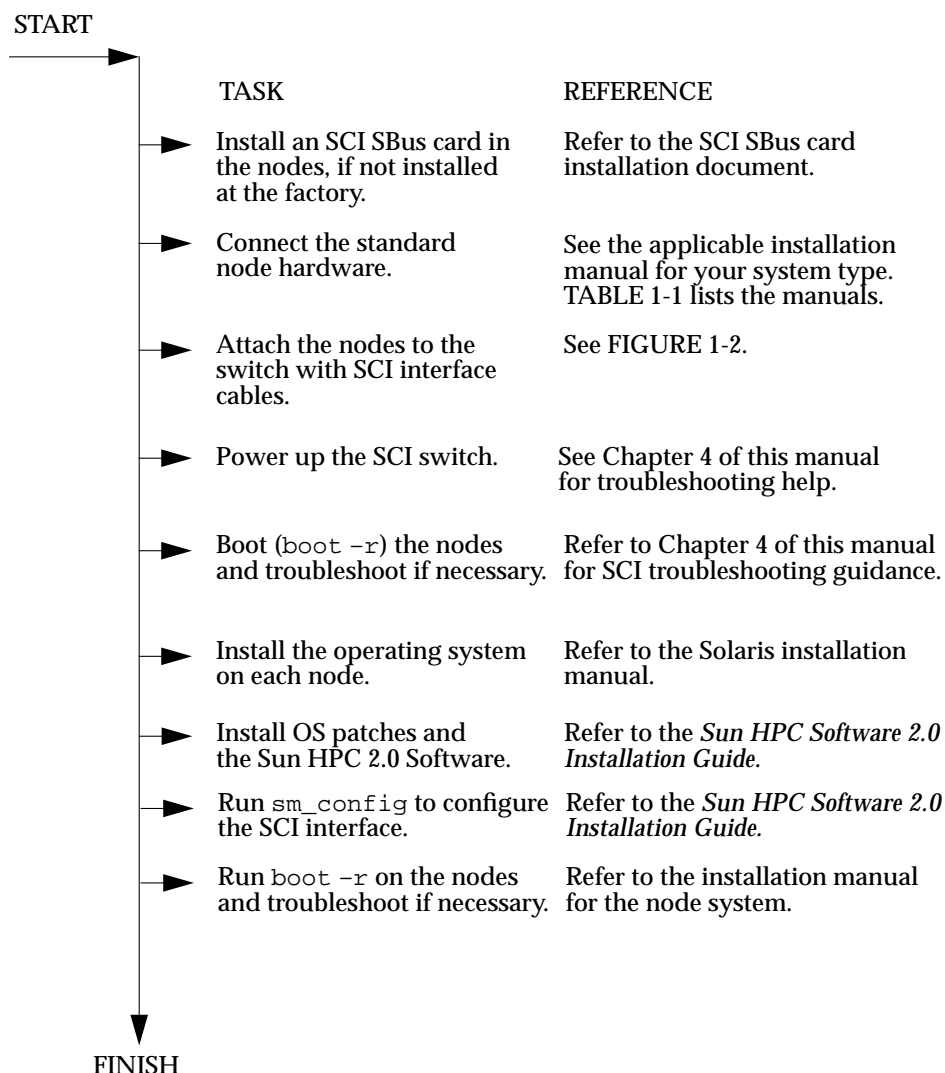


FIGURE 2-3 Task Map for a Switch Cluster

2.4 Hardware and Software Upgrade Kit

This task map is for upgrading a non-HPC server and connecting it to an SCI switch cluster (topology B in FIGURE 1-2.) For instructions, see Section 3.4 “Installing a Sun HPC Hardware and Software Upgrade Kit” on page 3-4.

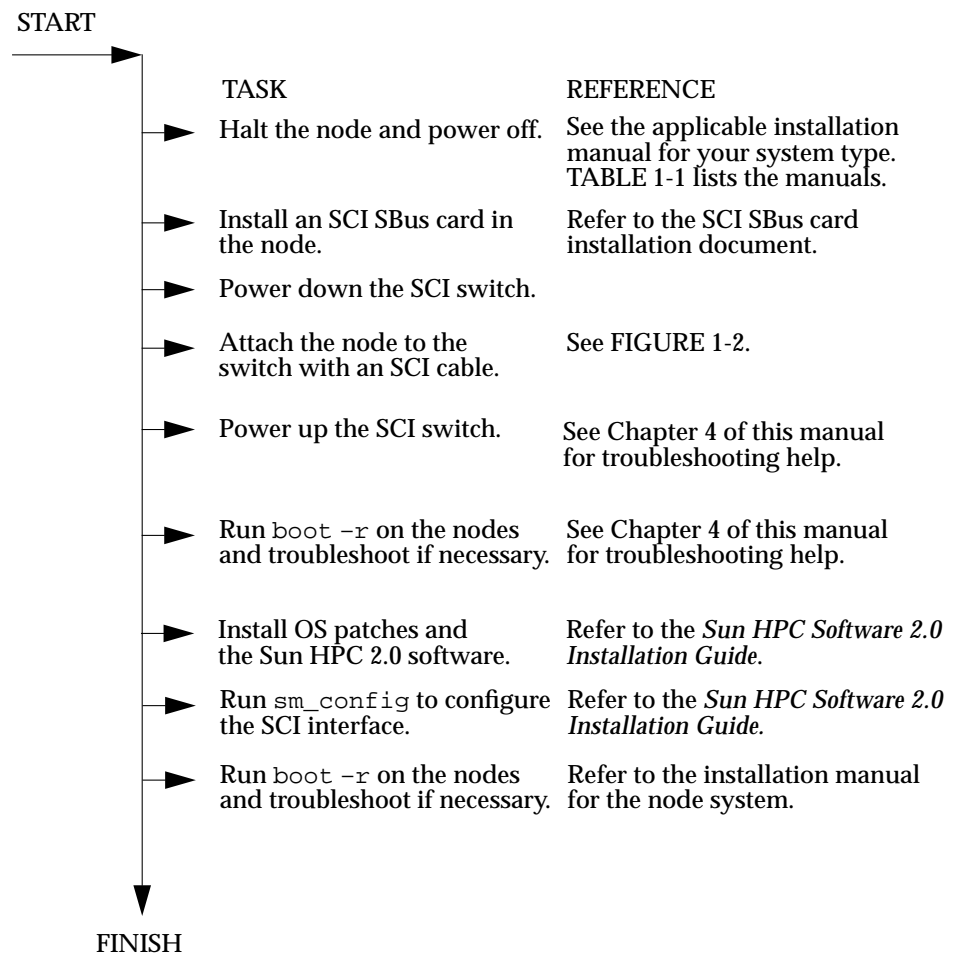


FIGURE 2-4 Task Map for a Hardware and Software Upgrade Kit

2.5 Software Upgrade Kit

This task map is for upgrading a non-HPC server that is not to be part of an HPC cluster. Refer to the documents listed in the task map for detailed instructions.

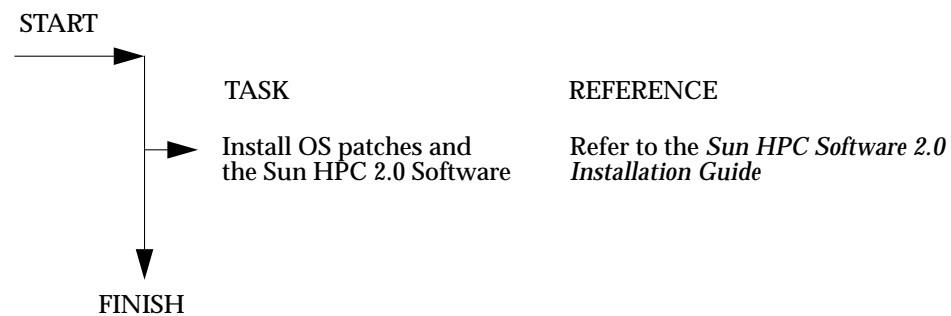


FIGURE 2-5 Task Map for a Sun HPC Software Upgrade Kit

Installation Procedures

3.1 Installing a Standalone Sun HPC Server

There are no supplementary hardware installation instructions in this guide for a standalone Sun Ultra HPC server.

This is a summary of the installation steps:

1. Install the system hardware.

Refer to the system installation guide (TABLE 1-1).

2. Install the operating system.

Refer to the Solaris installation guide.

Once these steps are complete, install the Sun HPC Software as described in the *Sun HPC Software 2.0 Installation Guide*.

3.2 Installing a Switchless Two-Node Cluster




FIGURE 3-1 2-Node Configuration

1. Position the nodes in the desired configuration.

The distance between nodes is limited by the length of the SCI station cable. See TABLE 1-3 for details.

2. Install SCI SBus interface cards in any node that does not already have one.

Remember, each node must have one, and only one SCI SBus interface card. See Section 3.4, “Installing a Sun HPC Hardware and Software Upgrade Kit,” for additional information.

3. Refer to the applicable system hardware installation manual(s) listed in TABLE 1-1 to complete the hardware installation of each node.

4. Power on and boot the nodes.

Refer to Chapter 4 of this manual for SCI troubleshooting guidance.

5. Install the general operating system on each node.

Refer to the Solaris document set for instructions.

6. Install any necessary operating system patches.

Refer to the *Sun HPC Software 2.0 Installation Guide* and release notes for HPC-specific patch requirements.

7. Install the HPC software.

Refer to the *Sun HPC Software 2.0 Installation Guide* for instructions.

3.3 Installing an SCI Switch Cluster




FIGURE 3-2 Maximum Switch Configuration

1. Position the nodes in the desired configuration.

The distance between nodes is limited by the length of the SCI station cables. See TABLE 1-3 for details.

2. Connect each node to an SCI switch port:

a. For each node, connect an SCI cable to the SCI SBus card.

b. Connect the SCI cable to a port on the SCI switch.

It is suggested that you start with port A.

c. Verify that the AC power switch on the SCI switch is in the OFF position, then connect the power cord of the SCI switch to a 120 VAC outlet.

Note that if fewer than four nodes are connected to the SCI switch, they can be connected to any ports—that is, there is no required port population sequence. By starting with port A and populating in consecutive order, however, you take full advantage of the SCI configuration templates that have been partially edited.

3. Refer to the applicable system hardware installation manual(s) listed in TABLE 1-1 to complete the hardware installation of each node.

4. Power on the SCI switch.

5. Power on and boot the nodes.

Refer to Chapter 4 of this manual for SCI troubleshooting guidance.

6. Refer to the Solaris document set to install the general operating system on each node.

7. Install the Sun HPC Software.

Refer to the *Sun HPC Software 2.0 Installation Guide* for instructions.

3.4 Installing a Sun HPC Hardware and Software Upgrade Kit

1. Prepare the SCI SBus card.

Set the SCI SBus interface card's scrubber jumper to the required position (FIGURE 3-3):

- If the interface card will be in a switchless two-node configuration, on interface card must have its scrubber jumper in the **ON** position and one must be in the **OFF** position.
- If the interface card will be connected to an SCI switch, its scrubber jumper must be in the **ON** position.

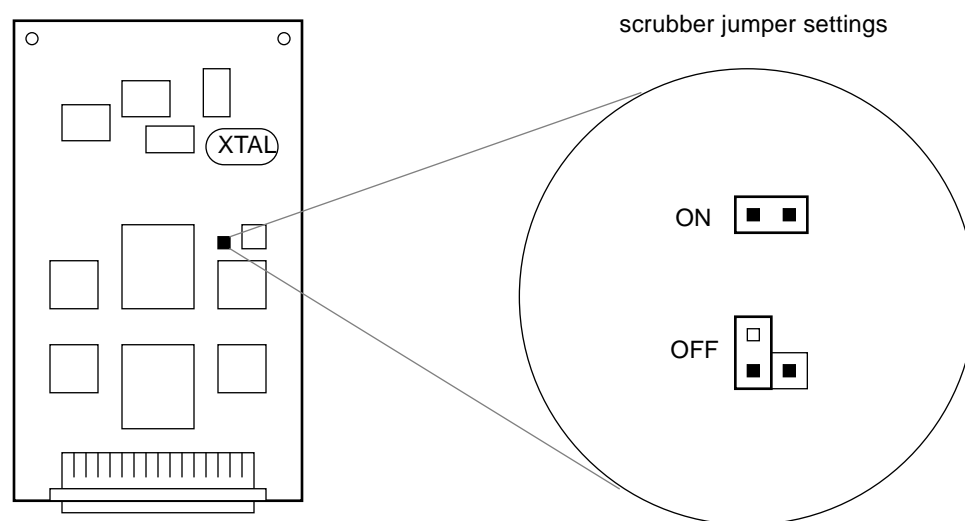


FIGURE 3-3 Scrubber Jumper on the SCI SBus Card

2. If the server to be upgraded is operating, halt and power down the server.

3. Position the node(s) to be upgraded in the desired configuration.

The distance between nodes is limited by the length of the SCI station cables. See TABLE 1-3 for details.

4. Install the SCI SBus interface card in an open SBus slot.

5. Connect the SCI switch to the server:

- a. Connect an SCI cable to the SCI SBus card.

b. Connect the SCI cable to a port on the SCI switch.

It is suggested that you start with the lowest-numbered port.

c. Apply power to the SCI switch.

Note that if fewer than four nodes are connected to the SCI switch, they can be connected to any ports—that is, there is no required port population sequence. By starting with port A and populating in consecutive order, however, you take full advantage of the SCI configuration templates that have been partially edited.

6. If the upgrade involves installing a new system, refer to the applicable system hardware installation manual(s) listed in TABLE 1-1 to complete the hardware installation of the server.

7. Refer to the Solaris document set to install the general operating system on the upgraded server.

8. Install the Sun HPC Software.

Refer to the *Sun HPC Software 2.0 Installation Guide* for instructions.

3.5 Installing a Sun HPC Software Upgrade Kit

1. **If necessary, upgrade the operating system first.**

The operating system must be Solaris 2.5.1 or 2.6.

2. **Install the Sun HPC Software.**

Refer to the *Sun HPC Software 2.0 Installation Guide* for instructions.

SCI Interface Troubleshooting

4.1 General Verification Tests

1. Use `get_ci_status` on all cluster nodes to verify interconnectivity. See Section 4.2.5, “The `get_ci_status` Command” and Appendix A for information about this utility.
2. Use `ifconfig -a` to verify that all SCI interfaces are up and have the correct IP addresses.
3. Use `ping` to verify connectivity between all nodes and along paths.

4.2 SCI Switch

4.2.1 General Hardware Inspection

Perform the following checks to determine the physical state of various SCI subsystem components. Verify that:

- All SCI scrubber jumpers are properly set, depending on the cluster topology.
- All SCI cables are properly seated.
- All SCI switches have power applied.
- No SCI status LEDs are red—see TABLE 4-1 and TABLE 4-2.

4.2.2 SCI Switch Status LED Locations

Clusters with two, three, or four nodes can be connected through an SCI switch. The switch status LEDs provide information that can be used to troubleshoot SCI switch failures (FIGURE 4-1). Guidelines for interpreting these LEDs are provided in Section 4.2.3, “Port Status LEDs,” and Section 4.2.4, “General Switch Status LED.”




FIGURE 4-1 SCI Status LED Locations

4.2.3 Port Status LEDs

The four port status LEDs located on the switch front panel can be used to troubleshoot individual port failures (TABLE 4-1).

Note – A switch port synchronization error can result from a cable being removed.

TABLE 4-1 SCI Switch Port Status LEDs

Situation	Port LED Status
No power	All four LEDs not lit.
Fatal switch errors: fatal hardware error, temperature too high, fan(s) not operative, power supply problem	All four LEDs red.
Port errors: SCI cable out, sync error	Associated port LED is red.
Port operative, no transactions	Associated port LED is green.
Port operative, with transactions	Associated port LED is blinking green.

4.2.4 General Switch Status LED

The switch status LED located on the rear panel indicates overall switch failures (TABLE 4-2).

TABLE 4-2 SCI Switch Rear Panel LED

Situation	LED Status
Fatal switch errors: fatal hardware error, temperature too high, fan(s) not operative, power supply problem	Red
Switch operational	Green

4.2.5 The get_ci_status Command

You can use the results of the `get_ci_status` command to troubleshoot clusters that have SCI switches. For example, for the configuration in FIGURE 4-2, if the `get_ci_status` command is used on `interconn1`, a typical output would be:

```
# /opt/SUNWsma/bin/get_ci_status
sma: sci #0: sbus_slot# 1; adapter_id 8 (0x08); ip_address 1; switch_id# 0;
port_id# 0; Adapter Status - UP; Link Status - UP
sma: Switch_id# 0
sma: port_id# 1: host_name = interconn2; adapter_id = 72; active | operational
sma: port_id# 2: host_name = interconn3; adapter_id = 136; active | operational
sma: port_id# 3: host_name = interconn4; adapter_id = 200;inactive|inoperational
#
```

In this example, the line

```
sma: port_id# 3: host_name = interconn4; adapter_id = 200;inactive|inoperational
```

indicates that the path between SCI switch 0, port 3 and `interconn4` is inactive and not operational.




FIGURE 4-2 Typical Four-Node Configuration with an SCI Switch

In this instance, if the `get_ci_status` command were run on all four nodes, and if the same path was inactive and inoperative between SCI switch 0, port 3 and interconn4, it is more than likely that either the SCI switch 0, port 3, the cable, or the interconn4 host adapter is faulty.

However, if the `get_ci_status` command indicates that the same path is inactive and inoperative for one node only, such as in the instance of interconn1, then it is more than likely that either the interconn 1 host adapter, the cable, or SCI switch 0, port 0 is faulty.

Note that some aspects of the `get_ci_status` command output, such as host names, will vary according to your configuration.

4.2.6 Client Net Failure

System console messages will identify the specific port that has failed. Otherwise, for information on test commands as well as additional troubleshooting, refer to the documentation that came with your client network interface card.

4.3 Incorrect Software Configuration

Make sure that:

- The working copy of the `sm_config` template file correctly matches the hardware configuration and cluster topology.
- `sm_config` ran successfully on only one of the cluster nodes.
- All nodes were rebooted after `sm_config` was executed.

4.4 Incorrect Firmware

If an SCI SBus card is loaded with the wrong firmware, the SCI cards will not be detected upon system power-on or reboot/reset.

Improper loading of the firmware can happen two ways:

- Old firmware programmed into new SBus2b cards
- New firmware programmed into old SBus2 cards

If proper firmware is loaded, a banner (containing the word `FCode`) will be printed from each SCI card *twice* during power-on or reboot or reset. No banner will be printed at all for a card loaded with improper firmware.

The following are sample console messages (which are not saved in the message file):

1. One SCI card is working in the node:

```
rebooting...
Resetting ...

DOLPHIN SBus-to-SCI (SBus2b) Adapter - 9029, Serial #5017
FCode 9029 $Revision: 2.3 $ - d9029_52 $Date: 1996/10/30 07:47:53
$

Executing SCI adapter selftest.      Adapter OK.
screen not found.
Can't open input device.
Keyboard not present.  Using ttya for input and output.

DOLPHIN SBus-to-SCI (SBus2b) Adapter - 9029, Serial #5017
FCode 9029 $Revision: 2.3 $ - d9029_52 $Date: 1996/10/30 07:47:53
$

Executing SCI adapter selftest.      Adapter OK.

Sun Ultra 1 SBus (UltraSPARC 167MHz), No Keyboard
```

2. No SCI cards are working in the node:

```
rebooting...
Resetting ...

screen not found.
Can't open input device.
Keyboard not present.  Using ttya for input and output.

Sun Ultra 1 SBus (UltraSPARC 167MHz), No Keyboard
```

Note – If SCI cards do not show up during boot time, check the physical installation of the cards. If reseating the cards does not correct the problem, the SCI cards may be damaged and should be returned. Contact your local SunServiceSM representative or your local SunService authorized service provider for assistance.

If you suspect that an SCI SBus interface card is loaded with the wrong firmware, perform the following steps to investigate:

- 1. With the system powered off, note the serial numbers of the adapter cards that are physically installed.**
- 2. Power on the system.**
- 3. Run `/opt/SUNWsci/bin/sciadm` and enter the `identify` command.**

This command displays the firmware version, fcode version, and serial number of each adapter board found.
- 4. Compare the number of cards found by `sciadm` against the number of adapters physically installed.**

Two cards should be displayed in the output. If not, there is at least one bad card in the system.
- 5. Compare the adapter board serial numbers from the output of the `identify` command, to the serial number on each adapter card physically installed.**

Note which serial number(s) are displayed. Cards that do not have their serial numbers displayed are bad and need replacement.

APPENDIX **A**

Man Pages

This appendix contains man pages for:

- `sm_config`
- `get_ci_status`

A.1 `sm_config(1m)`

```
sm_config(1M)           Maintenance Commands           sm_config(1M)

NAME
    sm_config - SCI adapter configuration utility for clusters

SYNOPSIS
    sm_config [-t] -f filename

AVAILABILITY
    SUNWsma

INTERFACE CLASSIFICATION
    Sun Private

DESCRIPTION
    sm_config is the SCI adapter configuration utility. It acts
    as a client of sm_configd(1M) daemon. sm_config contacts
    the sm_configd(1M) daemon on all the hosts and works in a
    distributed fashion to retrieve the adapter inventory and
    configure the adapter cards on these hosts. The configura-
```

tion process involves programming -

- (a) the adapter Node-Ids into the adapter's flash memory and
- (b) IP addresses into the cards.

Upon successful completion, a configuration file named `/etc/sma.config` is installed on all the hosts in the cluster. This file contains a snapshot view of the cluster members, switches, adapters etc. It also installs a file called `/etc/sma.ip` which contains the IP addresses of all the SCI interfaces in the cluster.

OPTIONS

`-t` starts `sm_config` in debug mode.

`-f filename` takes the filename as an input file. The input file template is available in `/opt/SUNWsmab/bin` directory as (a) `template.pdb` (for PDB clusters) and (b) `template.hpc` (for HPC clusters). These template files provide detailed information about the type of information required by `sm_config`.

This input file template contains 8 sections -

1. Cluster configuration section - specifies the type of cluster being configured (PDB or HPC). A sample template for this section -

Cluster is configured as = PDB

2. Host names section - requires the names of all the hosts in the cluster. If the hosts in the cluster do not have full public-net connectivity then the name of the host without connectivity must be preceded by `"_%"`. This indicates to `sm_config` not to contact this host via the public-net.

For example, consider a case where `host2` in a cluster of `host1`, `host2`, `host3` and `host4` lacks public-net connectivity. When `sm_config` is started with the following `template_1`, it will contact `host1`, `host3` and `host4` over the net and configure their SCI interfaces. However, it is now the user's responsibility to run `sm_config` on `host2` in the stand-alone mode using `template_2` below.


```
template_1 - used on host1, host3 and host4 :-
```

```
HOST 0 = host1
HOST 1 = _%host2
HOST 2 = host3
HOST 3 = host4
```

```
template_2 - used on host2 :-
```

```
HOST 0 = _%host1
HOST 1 = host2
HOST 2 = _%host3
HOST 3 = _%host4
```

A caveat to keep in mind when running `sm_config` in stand-alone mode is that, `sm_config` cannot guarantee the coherency of the `/etc/sma.config` generated during the different invocations (for eg. in the above case - `/etc/sma.config` on host2 versus the ones on host1, host3 and host4) if the user were to supply inconsistent input data for the two cases.

3. Number of Switches section - Accepts input for the total no. of switches in the cluster.

However, if the cluster being configured has some unused SCI adapters meant for use in the future, then the cluster should be configured as it would look in the future, when all the adapters are fully connected. For instance, a 1-switch cluster containing 4 hosts with 2 adapters on each (second set of adapters idle), which will later evolve into a 2-switch cluster should be configured as a 2-switch cluster.

This ensures that when the cluster evolves to its final form in future, new communication channels (SMA sessions) will be created on the new links (say, through a new switch) on the fly. This eliminates having to run `sm_config` later and rebooting the machine. A detailed example of this is given in the input template file.

A sample template for this section -

```
Number of Switches in cluster = 2
```

4. Number of Direct Links section - Accepts input for the total no. of direct SCI links in the cluster. A sample template for this section -

```
Number of Direct Links in cluster = 2
```

5. Allow Rings section - Whether the cluster supports configurations with multiple hosts connected to the same port of the same switch. A sample template for this section -

Allow Rings in cluster (Y/N)? = N

6. Adapter information section - Accepts detailed information for each adapter on each host. A sample template for this section is -

host 0 :: adp 0 is connected to = switch 0 :: port 0
host 0 :: adp 1 is connected to = link 1 :: endpt 0

7. Network IP address section - Accepts the first 3 octets (network) of the IP address for a particular switch or link. A sample template for this section -

Network IP address for Switch 0 = 204.152.65

8. Netmask section - Accepts the netmask for the private SCI sub-nets. For example, a cluster with less than 15 hosts per switch should select a netmask of 0xf0 while a cluster with 15 hosts or more but less than 31 hosts would choose 0xe0. A sample template for this section -

Netmask = f0

USAGE

The root user can start `sm_config` from the command line using the `-f` option to provide an input file to it.

A cluster can have 3 topologies -

(i) Switched - All hosts are connected to each other via SCI switches. Can have 2 or more hosts.

(ii) Non-switched - Two hosts connected directly via SCI cables (direct links). Has exactly 2 hosts.

(iii) Hybrid - Contains switches and direct links. Can have 2 or more hosts.

NOTE - At this point, PDB does not support more than two switches in a cluster (see (i) above), nor does it support case (iii) from above.

NOTE - `sm_config` can be run on any host in the cluster, but it should not be run on multiple hosts simultaneously (eg. via `cconsole`). If this occurs, the results are unpredictable - in the worst case, the adapter flash memory might get

programmed with corrupt data.

NOTE - After running `sm_config`, the system should be rebooted.

FILES

`/opt/SUNWcluster/bin/sm_config`
`/etc/sma.config`
`/etc/sma.ip`

SEE ALSO

`sm_configd(1M)`

DIAGNOSTICS

`sm_config` prints error and warning messages to `stderr`. If a fatal error occurs on any host or locally where `sm_config` is running, then the process is aborted and no `/etc/sma.config` is generated till the error is rectified. Do not reboot the machine till a successful run of `sm_config` has been completed.

RELEASE NOTES -

If `nis+` is being used as the name service then the default behaviour is to look up the global `nis+` map but if that doesn't exist, the local `/etc/services` file is not searched. This behaviour is different from the default `nis` behaviour. In this scenario `inetd` will be unable to start the `sm_configd` daemon.

SunOS 5.5.1

Last change: 30 March 1997

A.2 `get_ci_status(1m)`

`get_ci_status(1m)` Maintenance Commands `get_ci_status(1m)`

NAME

`get_ci_status` - Displays the Cluster configuration, the SCI adapter status and the SMA session status.

SYNOPSIS

`get_ci_status [-l]`

AVAILABILITY

SUNWsma

INTERFACE CLASSIFICATION

Sun Private

DESCRIPTION

`get_ci_status` displays the cluster configuration, the SCI adapter status and the SMA session status. It queries the SCI driver for information about the local SCI adapters and tests the connectivity to SCI adapters on other hosts, either via a switch or a direct link.

For each adapter in the cluster, `get_ci_status` displays the host it is on, the port on a switch it is connected to, its adapter-id and whether the local adapters can communicate with the adapters on other hosts.

In addition, for each local adapter `get_ci_status` displays the SBus slot# it is attached to, the host part of its IP address and whether the adapter is functional.

OPTIONS

`-l` Displays the local SCI adapter status only.

no option

Displays the local SCI adapter status and global cluster status. When displaying the global status, `get_ci_status` reports whether the remote adapter can be

reached at the hardware level (via SCI_PROBES) and/or at the software session level (via SMA sessions).

The SCI Probe reachability is indicated by active or inactive keywords following the status for the remote adapter in question. The software SMA session reachability is indicated by operational or inoperational keywords. For example, an output of the following form -

```
sma: sci #0: sbus_slot# 1; adapter_id 8 (0x08); ip_address
1; switch_id# 0; port_id# 0; Adapter Status - UP; Link
Status - UP
sma: sci #1: sbus_slot# 2; adapter_id 12 (0x0c); ip_address
17; switch_id# 1; port_id# 0; Adapter Status - UP; Link
Status - UP
sma: Switch_id# 0
sma: port_id# 1: host_name = interconn2; adapter_id = 72;
active | operational
sma: port_id# 2: host_name = interconn3; adapter_id = 136;
active | operational
sma: port_id# 3: host_name = interconn4; adapter_id = 200;
active | operational
sma: Switch_id# 1
sma: port_id# 1: host_name = interconn2; adapter_id = 76;
active | inoperational
sma: port_id# 2: host_name = interconn3; adapter_id = 140;
inactive | operational
sma: port_id# 3: host_name = interconn4; adapter_id = 204;
inactive | inoperational
```

indicates that there are 2 local adapters (adapter_id 8 and 12), both of which are functioning OK (keyword UP) with respect to SCI Probes to themselves. In case, a local adapter is unable to complete a successful SCI Probe to itself, the status of that local adapter is shown as DOWN.

The global status is shown in the set of lines associated with a switch. The status of the communication channel from the local adapter (adapter_id 8) to remote adapters on interconn2, interconn3 and interconn4 via the first switch (Switch_id# 0) is - SCI Probe status OK (keyword active) and SMA sessions functional (keyword operational).

However, the status of the communication channels from the local adapter (adapter_id 12) to remote adapters via the

second switch (Switch_id #1) have the following problems -

1. Adapter_id 76 => SCI Probes - reachable (keyword active) and SMA session - not established (keyword inoperational)
2. Adapter_id 140 => SCI Probes - unreachable (keyword inactive) and SMA session - established (keyword inoperational). This is a brief transitional stage.
3. Adapter_id 204 => SCI Probes - unreachable (keyword inactive) and SMA session - not established (keyword inoperational)

USAGE

`get_ci_status` can be run from the command line by any user. However, it can only be run after the adapter cards have been initialized using `sm_config(1M)`. This ensures that all the adapter node-ids have been properly programmed and the configuration file `/etc/sma.config` exists.

FILES

`/opt/SUNWsma/bin/get_ci_status`
`/etc/sma.config`

SEE ALSO

`sm_config(1M)`,

DIAGNOSTICS

`get_ci_status` prints error and warning messages to `stderr`.

SunOS 5.5.1

Last change: 30 March 1997

SMA Messages

If the Sun HPC Cluster uses an SCI switch, the SMA software can display the messages listed in this appendix. The message codes are listed alphabetically in Section B.1, and output from the messages is shown in Section B.2, in numerical order by message code suffix.

B.1 Message Directory

TABLE B-1 lists SMA messages in alphabetical order.

TABLE B-1 SMA Messages

Message Code	Description
SUNWcluster.sma.smactl.4007	Cannot create logical adapter: None found
SUNWcluster.sma.smactl.4008	Cannot create logical adapter:no response
SUNWcluster.sma.smad.1030	\$clustername adapter \$adp selected
SUNWcluster.sma.smad.1101	smad(\$pid): entering stand-alone mode
SUNWcluster.sma.smad.1102	smad: Cluster
SUNWcluster.sma.smad.1103	smad: Cluster
SUNWcluster.sma.smad.1104	smad: Cluster
SUNWcluster.sma.smad.1105	smad: Cluster
SUNWcluster.sma.smad.3104	SCI Adapter \$adp (\$from_aid): Session to
SUNWcluster.sma.smad.4004	smad(\$pid): exiting by request

TABLE B-1 SMA Messages (Continued)

Message Code	Description
SUNWcluster.sma.smad.4030	\$clustername no adapter available to select
SUNWcluster.sma.smad.5010	\$clustername adapter \$adp de-selected
SUNWcluster.sma.smak.1001	SCI Adapter \$adp: Card operational
SUNWcluster.sma.smak.1051	SCI Adapter \$adp: Link operational
SUNWcluster.sma.smak.3001	SCI Adapter \$adp (\$from_aid): Session to
SUNWcluster.sma.smak.3002	SCI Adapter \$adp (\$from_aid): Session to
SUNWcluster.sma.smak.3003	SCI Adapter \$adp (\$from_aid): Session to
SUNWcluster.sma.smak.4001	SCI Adapter \$adp: Card not operational
SUNWcluster.sma.smak.4051	SCI Adapter \$adp: Link not operational
SUNWcluster.sma.watchdog.1002	smad(\$pid) started
SUNWcluster.sma.watchdog.2001	child (\$pid) \$exit died \$status
SUNWcluster.sma.watchdog.4002	smad (\$pid) exiting

B.2 SMA Messages

The following pages list SMA messages in numerical order.

```
#####
# SUNWcluster.sma.watchdog.1002 - smad($pid) started
#####
msgid "SUNWcluster.sma.watchdog.1002.message"
msgstr "SMAD is starting."
msgid "SUNWcluster.sma.watchdog.1002.error"
msgstr "The process-id is mentioned in brackets."
msgid "SUNWcluster.sma.watchdog.1002.fix"
msgstr "Not Applicable. "

#####
# SUNWcluster.sma.watchdog.2001 - child ($pid) $exit died $status
#####
```



```

msgid "SUNWcluster.sma.watchdog.2001.message"
msgstr "The SMAD child daemon is dead. If necessary, another SMAD child"
      "daemon will be started up automatically."
msgid "SUNWcluster.sma.watchdog.2001.error"
msgstr "The process-id exit code and exit status are mentioned in the"
      "message."
msgid "SUNWcluster.sma.watchdog.2001.fix"
msgstr "Not Applicable. "

#####
# SUNWcluster.sma.watchdog.4002 - smad ($pid) exiting
#####
msgid "SUNWcluster.sma.watchdog.4002.message"
msgstr "The SMAD watchdog parent daemon is also exiting, because of a"
      "request to exit, or a fatal error."
msgid "SUNWcluster.sma.watchdog.4002.error"
msgstr "The process-id of the parent daemon is in brackets."
msgid "SUNWcluster.sma.watchdog.4002.fix"
msgstr "Not Applicable. "

#####
# SUNWcluster.sma.smad.4004 - smad($pid): exiting by request
#####
msgid "SUNWcluster.sma.smad.4004.message"
msgstr "The SMAD child daemon is dying by request."
      "request to exit, or a fatal error."
msgid "SUNWcluster.sma.smad.4004.error"
msgstr "Probably a due to a shutdown or a pkgrm."
msgid "SUNWcluster.sma.smad.4004.fix"
msgstr "Not Applicable. "

#####
# SUNWcluster.sma.smak.1001 - SCI Adapter $adp: Card operational
#####
msgid "SUNWcluster.sma.smak.1001.message"
msgstr "The adapter is working."
msgid "SUNWcluster.sma.smak.1001.error"
msgstr "Not Applicable. "
msgid "SUNWcluster.sma.smak.1001.fix"
msgstr "Not Applicable. "

#####
# SUNWcluster.sma.smak.1051 - SCI Adapter $adp: Link operational
#####
msgid "SUNWcluster.sma.smak.1051.message"

```

```

msgstr "The link from the mentioned adapter is working."
msgid "SUNWcluster.sma.smak.1051.error"
msgstr "Not Applicable. "
msgid "SUNWcluster.sma.smak.1051.fix"
msgstr "Not Applicable. "

#####
# SUNWcluster.sma.smak.4001 - SCI Adapter $adp: Card not operational
#####
msgid "SUNWcluster.sma.smak.4001.message"
msgstr "The adapter is not working."
msgid "SUNWcluster.sma.smak.4001.error"
msgstr "Not Applicable. "
msgid "SUNWcluster.sma.smak.4001.fix"
msgstr "Not Applicable. "

#####
# SUNWcluster.sma.smak.4051 - SCI Adapter $adp: Link not operational
#####
msgid "SUNWcluster.sma.smak.4051.message"
msgstr "The link from the mentioned adapter is not working."
msgid "SUNWcluster.sma.smak.4051.error"
msgstr "Not Applicable. "
msgid "SUNWcluster.sma.smak.4051.fix"
msgstr "Not Applicable. "

#####
# SUNWcluster.sma.smad.1030 - $clustername adapter $adp selected
#####
msgid "SUNWcluster.sma.smad.1030.message"
msgstr "The mentioned adapter has been selected to act as the"
      "logical adapter."
msgid "SUNWcluster.sma.smad.1030.error"
msgstr "Not Applicable. "
msgid "SUNWcluster.sma.smad.1030.fix"
msgstr "Not Applicable. "

#####
# SUNWcluster.sma.smad.4030 - $clustername no adapter available to select
#####
msgid "SUNWcluster.sma.smad.4030.message"
msgstr "There is no adapter available which can act as the logical adapter."
msgid "SUNWcluster.sma.smad.4030.error"
msgstr "This indicates significantly broken connectivity."
msgid "SUNWcluster.sma.smad.4030.fix"

```

```

msgstr "It is likely that the cluster will reconfigure itself. If the cluster"
      "does not reconfigure, administrator needs to take some corrective"
      "action, such as aborting a node or fixing the link problems."

#####
# SUNWcluster.sma.smad.5010 - $clustname adapter $adp de-selected
#####
msgid "SUNWcluster.sma.smad.5010.message"
msgstr "The mentioned adapter has been de-selected and is no longer the"
      "logical adapter."
msgid "SUNWcluster.sma.smad.5010.error"
msgstr "This adapter was acting as the logical adapter and it is no longer"
      "the logical adapter. If recovery happens, some other adapter will be"
      "chosen to act as the logical adapter."
msgid "SUNWcluster.sma.smad.5010.fix"
msgstr "Not Applicable. "
#####
# SUNWcluster.sma.smad.1101 - smad($pid): entering stand-alone mode
#####
msgid "SUNWcluster.sma.smad.1101.message"
msgstr "SMAD is starting up and things seem to be fine. No PDB software is "
      "running."
msgid "SUNWcluster.sma.smad.1101.error"
msgstr "Not Available."
msgid "SUNWcluster.sma.smad.1101.fix"
msgstr "Not Applicable."

#####
# SUNWcluster.sma.smad.1102 - smad: Cluster '$clustname' monitoring
#####
msgid "SUNWcluster.sma.smad.1102.message"
msgstr "CMM has asked SMAD to monitor the cluster status and it is now being "
      "monitored. It can also happen if the SMAD was already in 'monitor' "
      "mode and had died and restarted."
msgid "SUNWcluster.sma.smad.1102.error"
msgstr "Not Available."
msgid "SUNWcluster.sma.smad.1102.fix"
msgstr "Not Applicable."

#####
# SUNWcluster.sma.smad.1103 - smad: Cluster '$clustname' running
#####
msgid "SUNWcluster.sma.smad.1103.message"
msgstr "CMM has informed SMAD of the cluster membership. SMAD continues to "
      "monitor the entire cluster. It can also happen if the SMAD was "

```

```

"already in 'cluster' mode and had died and restarted."
msgid "SUNWcluster.sma.smad.1103.error"
msgstr "Not Available."
msgid "SUNWcluster.sma.smad.1103.fix"
msgstr "Not Applicable."

#####
# SUNWcluster.sma.smad.1104 - smad: Cluster '$clustname' returning
#####
msgid "SUNWcluster.sma.smad.1104.message"
msgstr "The cluster is entering the 'return' transition. SMAD does not "
"manage the logical adapter in this state, until 'step1' (running) "
"occurs. SMAD will continue to monitor the cluster for physical "
"failures. It can also happen if the SMAD was already in 'return' "
"mode and had died and restarted."
msgid "SUNWcluster.sma.smad.1104.error"
msgstr "Not Available."
msgid "SUNWcluster.sma.smad.1104.fix"
msgstr "Not Applicable."

#####
# SUNWcluster.sma.smad.1105 - smad: Cluster '$clustname' no longer running
#####
msgid "SUNWcluster.sma.smad.1105.message"
msgstr "This node has left the cluster. CMM has informed SMAD of this. The "
"cluster is no longer being monitored. It can also happen if the SMAD "
"had died and restarted."
msgid "SUNWcluster.sma.smad.1105.error"
msgstr "Not Available."
msgid "SUNWcluster.sma.smad.1105.fix"
msgstr "Not Applicable."

#####
# SUNWcluster.sma.smad.3104 - SCI Adapter $adp ($from_aid): Session to
# $to_aid not open
#####
msgid "SUNWcluster.sma.smad.3104.message"
msgstr "The session from adapter ($adp), whose SCI id is $from_aid, to the "
"SCI adapter with SCI id $to_aid could not be opened."
msgid "SUNWcluster.sma.smad.3104.error"
msgstr "It will be opened when possible, later."
msgid "SUNWcluster.sma.smad.3104.fix"
msgstr "Not Applicable."

#####

```

```

# SUNWcluster.sma.smak.3001 - SCI Adapter $adp ($from_aid): Session to
#                               $to_aid active
#####
msgid "SUNWcluster.sma.smak.3001.message"
msgstr "The session from adapter ($adp), whose SCI id is $from_aid, to the "
        "SCI adapter with SCI id $to_aid is active."
msgid "SUNWcluster.sma.smak.3001.error"
msgstr "Not Available."
msgid "SUNWcluster.sma.smak.3001.fix"
msgstr "Not Applicable."

#####
# SUNWcluster.sma.smak.3002 - SCI Adapter $adp ($from_aid): Session to
#                               $to_aid closed
#####
msgid "SUNWcluster.sma.smak.3002.message"
msgstr "The session from adapter ($adp), whose SCI id is $from_aid, to the "
        "SCI adapter with SCI id $to_aid has been closed."
msgid "SUNWcluster.sma.smak.3002.error"
msgstr "This is probably because a node is leaving the cluster, or "
        "rebooting."
msgid "SUNWcluster.sma.smak.3002.fix"
msgstr "Not Applicable."

#####
# SUNWcluster.sma.smak.3003 - SCI Adapter $adp ($from_aid): Session to
#                               $to_aid failed
#####
msgid "SUNWcluster.sma.smak.3003.message"
msgstr "The session from adapter ($adp), whose SCI id is $from_aid, to the "
        "SCI adapter with SCI id $to_aid has been closed."
msgid "SUNWcluster.sma.smak.3003.error"
msgstr "This is probably because of a failure, or a shutdown of the remote "
        "node."
msgid "SUNWcluster.sma.smak.3003.fix"
msgstr "Not Applicable."

#####
# SUNWcluster.sma.smactl.4007 - Cannot create logical adapter: None found
#####
msgid "SUNWcluster.sma.smactl.4007.message"
msgstr "The unavailability of a logical adapter has been detected at step1 "
        "of the reconfiguration process. This node will abort itself out of "
        "the cluster."
msgid "SUNWcluster.sma.smactl.4007.error"

```

```

msgstr "Not Available."
msgid "SUNWcluster.sma.smactl.4007.fix"
msgstr "Not Applicable."

#####
# SUNWcluster.sma.smactl.4008 - Cannot create logical adapter:no response
#                               from SMAD
#####
msgid "SUNWcluster.sma.smactl.4008.message"
msgstr "SMAD has not responded to a request to create a logical adapter.It "
      "is likely that SMAD has aborted / died in a fatal manner."
msgid "SUNWcluster.sma.smactl.4008.error"
msgstr "Not Available."
msgid "SUNWcluster.sma.smactl.4008.fix"
msgstr "Not Applicable."

```

Index

A

addresses, private network, 1-10

C

cable part numbers, 1-7

E

/etc/nsswitch.conf, 1-9
/etc/services file, 1-9
/etc/sma.config file, 1-9
/etc/sma.ip address file, 1-9

F

FRU part numbers, 1-7

H

hardware part numbers, 1-7

I

identify command, 4-7

J

jumper on SCI SBus card, 1-5

L

location of the scrubber jumper, 1-5

M

man pages
 get_ci_status, A-6
 smconfig, A-1
maximum cable length, 1-6

N

NIS+, 1-9
nisplus [NOTFOUND=return], 1-9

O

/opt/SUNWsci/bin/sciadm, 4-7

P

part numbers
 cables, 1-6
 documentation, 1-1
private network addressing, 1-10

R

requirements
 documentation, 1-1
 for upgrade, 1-7

S

SCI SBus card part number, 1-7
SCI switch part number, 1-7
scrubber jumper, 1-5
sm_config, 1-8
station cable part numbers, 1-7
SUNWsci hardware driver, 1-8
SUNWscid DLPI driver, 1-8
SUNWsma switch management software, 1-8

T

topologies, 1-2
troubleshooting, 4-1
 hardware
 network
 client net, 4-4