



Sun™ Technical Compute Farm Product Overview and Software Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303
U.S.A. 650-960-1300

Part No. 806-7319-10
December 2000, [Revision A](#)

[Send comments about this document to: docfeedback@sun.com](mailto:docfeedback@sun.com)

Copyright 2000 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, Java, The Java Coffee Cup logo, docs.sun.com, Sun Enterprise, Sun Enterprise SyMON, Solstice, Solstice DiskSuite, Sun Management Center, Solaris, Solaris Resource Manager, Sun StorEdge, JumpStart, SunVTS, Solstice AdminSuite, Solaris Management Console, Solstice Backup, Netra, and Sun WorkShop 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 2000 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, the Sun logo, Java, The Java Coffee Cup logo, docs.sun.com, Sun Enterprise, Sun Enterprise SyMON, Solstice, Solstice DiskSuite, Sun Management Center, Solaris, Solaris Resource Manager, Sun StorEdge, JumpStart, SunVTS, Solstice AdminSuite, Solaris Management Console, Solstice Backup, Netra, et Sun WorkShop 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 REpondre 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 xi

1. Product Overview 1

Introducing the Sun Technical Compute Farm 1

 Typical Sun TCF System Applications 1

Understanding the Sun TCF Hardware 3

Understanding the Sun TCF Software 5

 Preinstalled Software 6

 Other Supported Software 9

2. Configuring the Sun TCF Software 11

Before You Start 12

 Solaris Version Information 14

Configuring the File Server 14

 Getting Started 15

 ▼ To Prepare to Configure the File Server 15

 ▼ To Configure the File Server 15

 ▼ To Set Up the Disk Arrays 18

Configuring the Compute Engines 23

 ▼ To Edit the Configuration Files 23

- ▼ To Run the Configuration Script 25
- Completing the Setup Procedure 27
- ▼ To Complete the Setup Procedure 27
- Overriding the Compute Engine Configuration 28
- ▼ To Override the Configuration Script 28
- Recovering/Restoring the Compute Engines 29
- ▼ To Restore a Compute Engine 30
- ▼ To Replace a Compute Engine 31
- Reinstalling Solaris Software on the File Server 31
 - Installing Solaris 8 Software 32
 - Reinstalling the Software Patches on the File Server 34

3. Using Solaris Management Console Software 37

- Configuring and Running Solaris Management Console Software 37
 - Before You Start 37
 - ▼ To Start Solaris Management Console 38
 - Working With Sun Management Console 39
 - ▼ To Add an Application to Sun Management Console 39
 - ▼ To Start an Application from Sun Management Console 40
 - ▼ To Manually Add a Machine to Solaris Management Console 41
 - ▼ To Use the Solaris Management Console Log File 42
- For More Information 42

4. Using Sun Management Center Software 43

- Sun Management Center Functions 43
- Configuring Sun Management Center 44
 - Configuring Email Notification of Alarms 44
 - ▼ To Add Email Addresses to the Sun Management Center Alarm Aliases 45

Incorporating Compute Engines Into the File Server's Domain	45
▼ To Perform a Discovery of the Compute Engines	45
Running Sun Management Center	52
Using the Browser	52
▼ To Use the File Scanning Module	53
▼ To View the Applications	53
Preconfigured Alarms	54
▼ To View the Alarms	55
▼ To View the Alarm Configuration	55
Monitoring the Performance of Your Sun TCF System	59
▼ To View Performance Statistics for a Particular Machine	59
Some Statistics to Monitor	62
5. Using Sun Grid Engine Software	63
Preconfigured Software Settings	63
Working With Sun Grid Engine	64
Setting Up Your Environment	64
Sun Grid Engine Functions	65
Sun Grid Engine Applications	65
▼ To Temporarily Stop Sun Grid Engine	67
▼ To Restart Sun Grid Engine	67
For More Information	68
6. Using Solstice DiskSuite Software	69
Solstice DiskSuite Functions	69
Starting the Solstice DiskSuite Software	70
▼ To Start Solstice DiskSuite From the Command Line	70
Understanding Solstice DiskSuite Metadevices	70
Root Mirroring	71

7. Using Solstice AdminSuite Software	73
Solstice AdminSuite Overview	73
Setting Up and Running Solstice AdminSuite	74
▼ To Start Solstice AdminSuite	74
8. Optimizing Your Sun TCF System	75
Optimizing File Server Performance	75
Editing the Cache Size Settings	75
Increasing NFS Threads	77
Optimizing I/O Performance	77
For More Information	78
9. Troubleshooting	79
Diagnostic Tools	79
Using SunVTS Software	79
▼ To Start SunVTS	80
▼ To Use the SunVTS Performance Meter	80
▼ To View the SunVTS Test Logs	81
Using SunVTS Test Tracing	81
Hardware Concerns	81
▼ To Boot From the Mirrored Disk	81

Tables

TABLE P-1	Sun TCF Related Documentation	xiii
TABLE P-2	Typographic Conventions	xiv
TABLE P-3	Shell Prompts	xv
TABLE 1-1	Sun TCF Software Components	6
TABLE 1-2	Other Supported Software	9
TABLE 2-1	Sun TCF Software Setup Worksheet	12
TABLE 2-2	Compute Engine Information Worksheet	13
TABLE 2-3	NFS Configurations for the Compute Engines	13
TABLE 2-4	Compute Engine Configuration Files	23
TABLE 2-5	Compute Engine Configuration Settings	27
TABLE 2-6	Preinstalled Software Interoperability/Compatibility Matrix	32
TABLE 2-7	Solaris 8 Version Compatibility Matrix	33
TABLE 2-8	Patches for Solaris 2.6 Software	34
TABLE 2-9	Patches for Solaris 7 Software	35
TABLE 2-10	Patches for Solaris 8 Software	35
TABLE 4-1	Preconfigured Alarms	54
TABLE 4-2	Performance Statistics	62
TABLE 5-1	Sun Grid Engine Settings	63
TABLE 6-1	Solstice DiskSuite Interoperability/Compatibility Matrix	69
TABLE 7-1	Solaris Easy Access Server Interoperability/Compatibility Matrix	73

Figures

FIGURE 1-1	Typical Sun Compute Farm Configuration	2
FIGURE 1-2	Sun TCF Hardware Components	4
FIGURE 1-3	Sun TCF Software Components	8
FIGURE 3-1	Solaris Management Console Main Screen	39
FIGURE 3-2	Add Application Dialog Box	40
FIGURE 3-3	Applications View Screen	41
FIGURE 4-1	Sun Management Center Login Screen	46
FIGURE 4-2	Sun Management Center Main Window	47
FIGURE 4-3	Discovery Request Dialog Box	48
FIGURE 4-4	New Discovery Request Dialog Box	49
FIGURE 4-5	Confirm Dialog Box	50
FIGURE 4-6	Discovery Window	51
FIGURE 4-7	Compute Engines in the Sun Management Center Domain	52
FIGURE 4-8	File Scanning Window	53
FIGURE 4-9	Applications View	54
FIGURE 4-10	Pop-Up Menu	56
FIGURE 4-11	Alarm Threshold Information	57
FIGURE 4-12	Actions Dialog Box	58
FIGURE 4-13	System Load Window	60
FIGURE 4-14	Pop-Up Menu	61

Preface

The *Sun Technical Compute Farm Product Overview and Software Guide* describes the Sun™ Technical Compute Farm (Sun TCF) product. This manual includes an overview of the software components, installation procedures, and administration references. This book is part of the Sun TCF software media kit shipped with the Sun TCF product for final assembly and use.

Sun Enterprise Services personnel install the Sun TCF hardware at your site. The Solaris™ 2.6 operating environment and the other software components are also preinstalled/.

Who Should Use This Book

This book is written for system administrators, service providers, and customers experienced with Sun workstations who need to install or administer the Sun TCF software system. Refer to the *Sun Technical Compute Farm Installation and Service Manual* and to the Product Notes for information on the Sun TCF hardware components as well as for instructions related to hardware assembly, configuration, and maintenance.

How This Book Is Organized

This book is divided into nine chapters, as follows:

Chapter 1 is an overview of Sun TCF and its associated components. This includes descriptions of how the Sun TCF components work together as a cohesive unit.

Chapter 2 describes how to configure the preinstalled software images on the file server and compute engines. This chapter also describes how to set up the software applications included in the preinstalled images.

Chapter 3 contains information on the Solaris Management Console™ software. This program serves as the main interface to the other Sun TCF applications.

Chapter 4 describes the Sun Management Center™ application. You use Sun Management Center to monitor the performance of your Sun TCF system.

Chapter 5 outlines some of the basic functions of the Sun Grid Engine software. Sun Grid Engine helps you manage queues, distribute jobs, and take advantage of unused or underused computing resources in the Sun TCF system.

Chapter 6 contains information on Solstice DiskSuite™. Use this application to manage the disk arrays in the Sun TCF system.

Chapter 7 discusses the Solstice AdminSuite™ software. Solstice AdminSuite allows you to create and maintain user and group accounts, remote hosts, and system resources.

Chapter 8 describes some methods you can use to optimize the performance of your Sun TCF system.

Chapter 9 provides reference information for diagnosing and resolving Sun TCF hardware and software problems.

Related Documentation

The following table describes other Sun TCF specific documentation included with your Sun TCF system.

TABLE P-1 Sun TCF Related Documentation

Title	Description	Part Number
<i>Sun Technical Compute Farm Installation and Service Manual</i>	This manual provides procedures for installing the hardware portion of your Sun TCF system.	806-7328-10
<i>Sun Technical Compute Farm Read Me First</i>	This document describes important last-minute information and other updates concerning your Sun TCF system.	806-7326-10 (English) 806-7329-10 (Japanese)
<i>Sun Technical Compute Farm Documentation Roadmap</i>	This document lists the manuals and other documentation included with the Sun TCF system.	806-7330-10

In addition, you should have the documentation for Solstice DiskSuite™, Sun Management Center™, and Solaris 2.6 and Solaris 8 software, as well as the documentation for your network hardware and for any optional software you choose to install. The Solstice DiskSuite and Sun Management Center documentation sets are included with the Sun TCF software media kits.

Refer to the *Sun Technical Compute Farm Documentation Roadmap* for a description of all of the manuals that come with the Sun TCF system. Some manuals are located only on CD-ROM or on `docs.sun.com`SM.

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/`

The Sun TCF product Web site at <http://www.sun.com/desktop/suntcf/docs.html> contains information on the Sun TCF product, including late-breaking news, specific equipment installation procedures, helpful information on software configuration, performance tips, and other useful topics. This Web site also links to the docs.sun.com Web site.

Typographic Conventions

TABLE P-2 Typographic Conventions

Typeface or Symbol	Meaning	Example
<i>AaBbCc123</i>	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. <code>machine_name%</code> You have mail.
AaBbCc123	What you type, contrasted with on-screen computer output	<code>machine_name%</code> su Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new words or terms, or words to be emphasized	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be root to do this.

Shell Prompts

TABLE P-3 Shell Prompts

Shell	Prompt
C shell prompt	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

Product Overview

This chapter contains the following topics:

- Introducing the Sun Technical Compute Farm on page 1
- Understanding the Sun TCF Hardware on page 3
- Understanding the Sun TCF Software on page 5

Introducing the Sun Technical Compute Farm

Sun Technical Compute Farm (Sun TCF) is a high-performance system that allows technical workgroups to more efficiently allocate and share computing resources. The Sun TCF hardware is a collection of independent, rack mounted technical compute engines connected to a file server, disk arrays, and networking hardware.

The Sun TCF software includes the Solaris operating environment, Sun system management software, and workload management software. These software layers work together to make the hardware appear to the end user as one computational entity.

Typical Sun TCF System Applications

The following list describes some typical uses for a Sun TCF system:

- Resource sharing for compute-intensive applications such as EDA (electronic design automation), engineering simulations, bioinformatics, software development, and financial analysis
- Maximizing CPU utilization

- Cost-effectively deploying more compute resources to end users
- Simultaneously executing many compute-intensive tasks
- Making development processes faster and more powerful by spawning multiple jobs simultaneously, thus reducing development cycle times

The Sun TCF system enables you to distribute compute-intensive jobs from individual users to the compute engines in the Sun TCF configuration, thus freeing your local user machines to handle user desktop tasks as well as relieving network traffic. Using Sun TCF with workload management software, such as Sun Grid Engine (formerly Codine), further enhances these capabilities.

Note – For information on workload management software, refer to the Sun TCF product Web site at <http://www.sun.com/desktop/suntcf>.

The following figure shows a typical compute farm arrangement using two Sun TCF systems to support three workgroups of users.

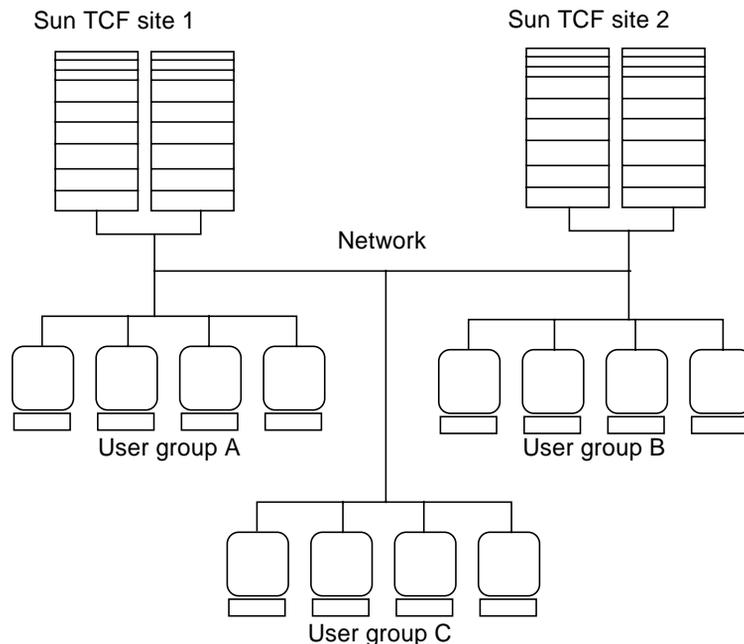


FIGURE 1-1 Typical Sun Compute Farm Configuration

Understanding the Sun TCF Hardware

The compute farm system hardware consists of a rack mounted file server and compute engines along with networking equipment, data storage assemblies, and any optional devices you ordered (such as additional racks, storage arrays, and so on). Except for the file server and the networking equipment, all of the Sun TCF hardware components can be added to or removed from the system while the compute farm is running.

For a complete and detailed description of the Sun TCF hardware, refer to the *Sun Technical Compute Farm Installation and Service Manual*.

Typically, your Sun TCF configuration is assembled and tested at the factory. The software images for the file server and compute engines are installed on the systems and configured to start the installation at power-up.

Sun Enterprise Services normally installs the Sun TCF hardware at your site for you, and then you proceed with configuring the software and installing any optional applications. If this is not the case with your Sun Sun TCF system, refer to the *Sun Technical Compute Farm Installation and Service Manual* before proceeding with the instructions in this document.

The following figure shows the components in the base rack of the compute farm.

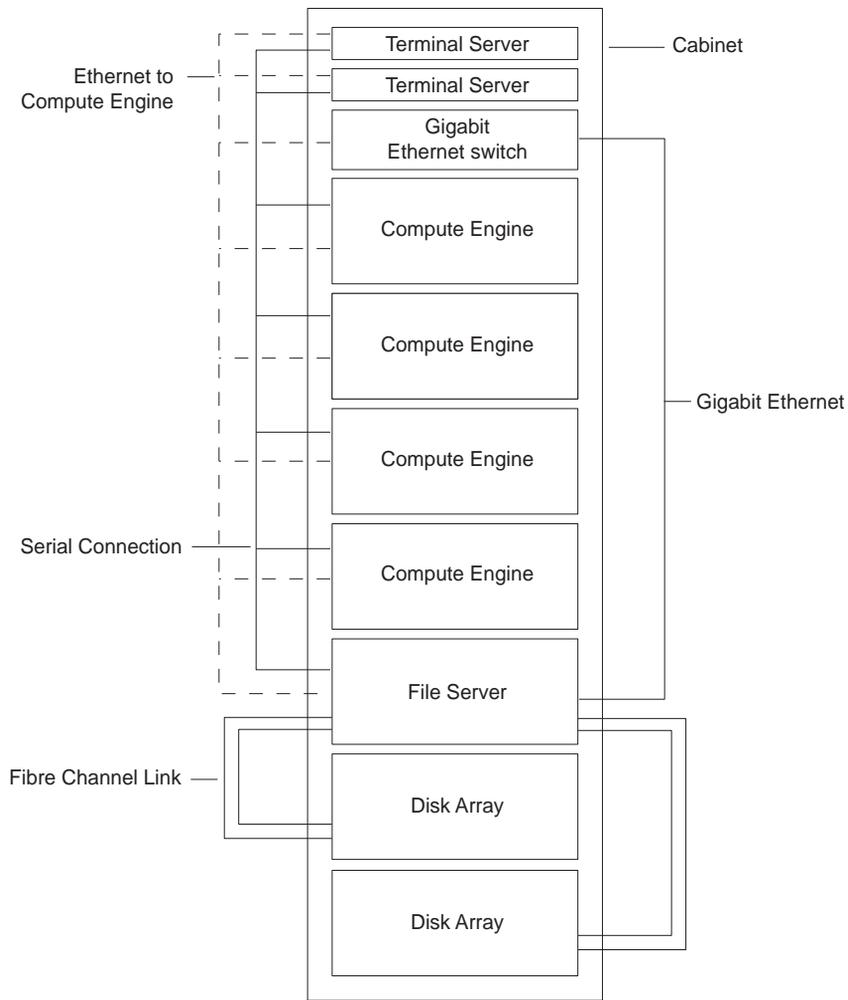


FIGURE 1-2 Sun TCF Hardware Components

The basic Sun TCF configuration ships with the following hardware components:

- File server

The Sun TCF system's file server is a Sun Enterprise™ 420R server with four CPUs. The file server also contains three cards: one that supports Gigabit Ethernet connection to the Ethernet switch through the Gigabit Ethernet adapter, and two cards for redundant Fibre Channel links to the disk arrays.

- **Compute engines**

The compute engines are Sun Enterprise 420R systems with four CPUs each. The base Sun TCF configuration supports five compute engines. Expansion racks contain nine additional compute engines each. For more information on the expansion racks and configurations, refer to the *Sun Technical Compute Farm Installation and Service Manual*.

- **Data storage**

The data storage device is a Sun StorEdge™ A5200 series disk array, designed for RAID-type data storage. The base Sun TCF system contains one disk array; a configuration that contains two disk arrays is also available.

- **Ethernet switch**

The Cisco 3524-XL Ethernet switch provides the 100BASE-T Ethernet connections between the compute engines and the file server, as well as supplying the GigabitEthernet connection to the file server.

- **Terminal servers**

The Cisco AS2511-RJ access servers (referred to in this manual as *terminal servers*) provide the RS-232 serial connections to the file server and compute engines for console access.

- **Rack**

The rack for the Sun TCF system is a Sun StorEdge 72-inch standard rack. The initial rack holds one complete Sun TCF system with one file server, one disk array, four compute engines, and the associated networking components. Expansion racks, which are 72-inch Sun StorEdge Expansion Cabinets, hold nine additional compute engines.

In addition to the Sun TCF system hardware, you need a router connection (also known as a *gateway*) to connect your Sun TCF system to your internal network.

Understanding the Sun TCF Software

The Sun TCF system comes with preinstalled software applications that are complete and ready to use. However, if you want to use other applications for some functions, the Sun TCF system does support certain other software packages.

The following table shows the key preinstalled software components of the Sun TCF system and the role that each application plays. In addition, the table shows some additional supported applications for certain functions.

TABLE 1-1 Sun TCF Software Components

Feature	Preinstalled Software	Other Optional Software Supported
Operating environment	Solaris	
Workload management	Sun Grid Engine	LSF
Monitoring	Sun Management Center	
Administration	Solstice AdminSuite/Solaris Management Console	
Storage software	Solstice DiskSuite	VERITAS Volume Manager
Configuration management	SCCS	Clearcase, Sun WorkShop™ TeamWare
Diagnostics	SunVTS™	
Backup	ufsdump and ufsrestore (included in the Solaris operating environment)	Solstice Backup™
Ease of use	Ready-to-use configuration scripts and recovery utilities	

“Other Supported Software” on page 9 describes some of other applications that are available from Sun and from third-party vendors.

Preinstalled Software

The basic compute farm ships with the following preinstalled software components:

- Solaris operating environment

The file server is preinstalled with the Solaris 2.6 operating environment. If desired, you can install Solaris 8 software (included on CD) or Solaris 7 with the Sun TCF system. For more information on the Solaris operating environment, refer to “Getting Started” on page 15.

- Sun Grid Engine Software

Sun Grid Engine software orchestrates all technical work that takes place on the Sun TCF system. Sun Grid Engine schedules and executes jobs, manages the utilization of the compute engines, and keeps accounting statistics on the overall performance of the compute farm. The software contains a fully-functional GUI

(graphical user interface) as well as a command-line interface. If desired, you can incorporate other systems into the Sun TCF cluster (such as desktop workstations), and you can execute jobs on them at appropriate times (such as nights and weekends). For more information on Sun Grid Engine, refer to Chapter 5 and to the Sun Grid Engine documentation.

- Sun Management Center Software

Use the Sun Management Center (formerly Sun Enterprise SyMON™) software to monitor all the systems in the Sun TCF system from a central, graphical Java™-based console. This software is a client-server application that can monitor performance and resource data such as CPU load, I/O transactions, system health, and disk usage. Sun Management Center can also automatically send alerts and perform actions when specified conditions occur. Sun Management Center also works with Sun Grid Engine to monitor workload processing and resource utilization, and to perform simple workload maintenance. If desired, you can incorporate other systems into the domain of Sun Management Center and monitor their status as well. For more information on Sun Management Center, refer to Chapter 4.

- Solaris Management Console Software

Use the Solaris Management Console software as a portal to all system management applications on the file server and the compute engines. From one console, you can access all essential applications on all systems, as well as any other applications you choose to incorporate. For more information on Solaris Management Console, refer to Chapter 3 and to the Solaris Management Console documentation.

- Solstice AdminSuite Software

Use Solstice AdminSuite™ software to administer user accounts, group accounts, networks, disks, and remote hosts. For more information on Solstice AdminSuite, refer to Chapter 7.

- SunVTS

Solaris SunVTS™ (Validation Test Suite) software serves as a diagnostic tool for the Sun TCF hardware. For more information on SunVTS, refer to Chapter 9.

- Solstice DiskSuite Software

Use the Solstice DiskSuite™ (SDS) software to configure and manage complex storage systems for better data throughput and redundancy. This application provides a GUI for viewing and manipulating storage structures, such as disk mirrors, stripes, RAIDs, and hot-spare pools.

The preinstalled software on the Sun TCF system also includes two disk utilities (`ufsdump` and `ufsrestore`) for the Sun StorEdge A5200 series disk array, utilities for the Gigabit Ethernet and Fibre Channel cards, and any required patches to the preinstalled software.

The following figure illustrates the relationships among the Sun TCF software applications.

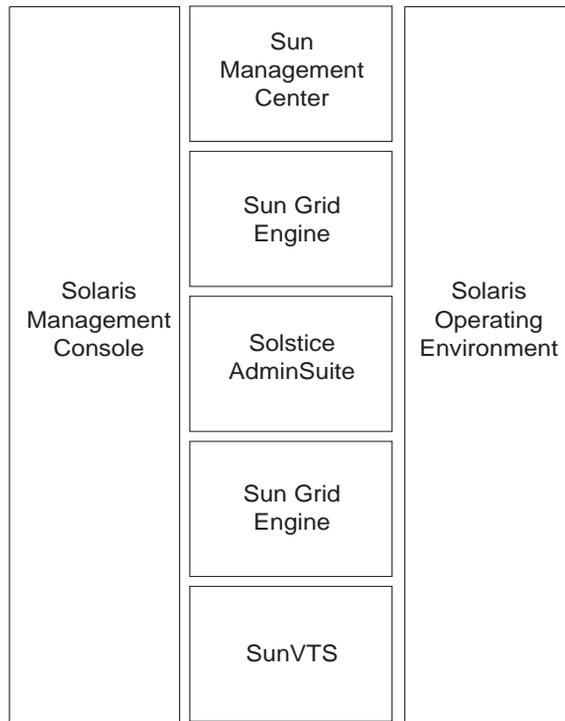


FIGURE 1-3 Sun TCF Software Components

Other Supported Software

Some Sun and third-party software packages can enhance the functionality and scalability of your Sun TCF system. The following table describes some of the functions you can add and the corresponding software products.

TABLE 1-2 Other Supported Software

Function	Sun Software	Third-Party Software
Workload management		LSF (Load Sharing Facility) from Platform Computing
Configuration management	Sun WorkShop TeamWare 2.1	ClearCase from Rational Software, RCS freeware
Volume management, Source control management	Source Code Control System (SCCS) (included in the Solaris software), Sun Enterprise Volume Manager™	VERITAS Volume Manager
Backup and Restore	Solstice Backup	
Compute engine configuration and recovery	JumpStart™	

Configuring the Sun TCF Software

This chapter contains the following topics:

- Before You Start on page 12
- Configuring the File Server on page 14
- Configuring the Compute Engines on page 23
- Completing the Setup Procedure on page 27
- Overriding the Compute Engine Configuration on page 28
- Recovering/Restoring the Compute Engines on page 29
- Reinstalling Solaris Software on the File Server on page 31

The following list outlines the steps you perform in this chapter.

1. Fill out the configuration worksheet and tables
2. Configure the file server
3. Set up the disk arrays
4. Edit the compute engine configuration files
5. Configure the compute engines
6. Complete the setup procedure

The Sun TCF system hardware should be installed before you perform the procedures in this chapter. Refer to the *Sun Technical Compute Farm Installation and Service Manual* before proceeding.

Note – You must install another workstation on the same network as the Sun TCF system to serve as a system administration workstation or system monitor. Without the additional workstation, you will not be able to access the Sun TCF system, since it does not come with a monitor.

Before You Start

The following tables contain the information you need to configure your Sun TCF file server. Have this information ready when you perform the procedures in the next section in this manual, “Configuring the File Server” on page 14.

TABLE 2-1 Sun TCF Software Setup Worksheet

Item	Value
Communication Information	
Language and locale	
Geographic location (us, uk, and so on)	
Time zone	
Terminal type for installation console	
File Server Information	
File server host name	
File server IP address	
Compute Engine Information	
Compute engine identity	Refer to TABLE 2-2.
Compute engine NFS mount points	Refer to TABLE 2-3.
Network Information	
Networked or standalone?	This should always be <i>networked</i> .
Naming service used (for example, NIS/yp, DNS)	
Domain name	
Naming service (NIS/yp or DNS) server host name	
Naming service server (NIS/yp or DNS) IP address	
Network subnet mask	

TABLE 2-1 Sun TCF Software Setup Worksheet

Item	Value
Solstice DiskSuite Information	
Root filesystem mirroring (yes/no)	
Configure the disk arrays (yes/no)	
• If yes, then: Number of volumes	
• If yes, then: Name of mount point for each volume	
• If yes, then: Type of each volume (RAID0, RAID1, RAID0+1)	
• If yes, then: Size of each volume	

Use the following table to record the configuration and identities of the compute engines. The host name and IP address should have been assigned already. If you are using NIS, make sure the compute engines are included in the NIS maps.

TABLE 2-2 Compute Engine Information Worksheet

Host Name	IP Address	Ethernet Address

Use the following table to record the configuration of the NFS mount points on the compute engines.

Note – Be sure to include an entry for the volumes on the file server, even though the volumes have not yet been configured.

TABLE 2-3 NFS Configurations for the Compute Engines

Remote Server Host Name	Remote Directory	Local Mount Point

Note – After you have finished entering the information in the tables, make a copy of the tables and attach it to your file server.

Solaris Version Information

The Sun TCF file server comes with a preinstalled software image that includes the Solaris 2.6 operating environment. The Solaris 8 operating environment is included in the Solaris 8 Media Kit, but is not preinstalled. If you want to run Solaris 8 software on your file server, you can install it from the CD during this initial configuration process, or you can install Solaris 8 later on.



Caution – If you install Solaris 8, the preinstalled software functions (such as compute engine recovery) will be permanently eliminated from the Sun TCF software. You will need to manually set up those functions in Solaris 8.

If you decide to install Solaris 8, you will also need to upgrade the applications in the preinstalled software so the program versions are compatible with Solaris 8. The upgraded versions of the preinstalled applications are included with the Solaris 8 Media Kit. Refer to the documentation for each of the software packages and follow the instructions. For information on the software versions included with Solaris 8, refer to “Installing Solaris 8 Software” on page 32.

The Sun TCF software supports Solaris 7. If you prefer, you may install it instead of Solaris 8. Solaris 7 is not included with the Sun TCF system, so you need to purchase the software separately. Contact your Sun representative for more information, or purchase the software online at <http://www.sun.com>.

Configuring the File Server

The setup script for the file server guides you through the following tasks:

- Installing an Asian locale (optional)
- Enabling root mirroring on the file server
- Configuring the disk array(s)

Once you have finished performing these steps, you configure the compute engines, as described in “Configuring the Compute Engines” on page 23.

Getting Started

Before you perform the procedures in this section, you should plan and configure the networking setup for your Sun TCF system. This includes deciding whether to connect the Sun TCF system to your internal network, and so on. Refer to the *Sun Technical Compute Farm Installation and Service Manual* for more information.

▼ To Prepare to Configure the File Server

1. **Ensure that the hardware installation is complete.**
2. **Make sure that the power is off to all of the units in the system; this includes the file server and all of the compute engines.**
3. **Power on the rack and configure the terminal server.**

Refer to Chapter 7 in the *Sun Technical Compute Farm Installation and Service Manual*.

▼ To Configure the File Server

1. **To access the serial port of the file server, type the following command at the system prompt:**

```
% telnet servername port
```

servername is the host name of the terminal server, and *port* is the port number attached to the file server (preconfigured as 2001).

2. **Power on the file server.**
3. **Follow the screen prompts to answer the system identification questions.**

Refer to the information you recorded in TABLE 2-1. Once you have finished entering the identification information, the file server automatically reboots.

4. When the file server reboots, watch the screen messages.

The software asks whether you want to install an Asian locale on the file server. The screen shows messages similar to the following:

```
*****
Asian locale setup
*****
Do you want to install one of the following locales?
0)<none>
1) Simplified Chinese EUC
2) Traditional Chinese EUC
3) Japanese EUC
4) Korean EUC
Please choose:
```

5. Type the number that corresponds to your choice.

The software shows the following screen:

```
=====
Root Mirroring
=====

IMPORANT NOTE: It is STRONGLY RECOMMENDED
that you set up root mirroring on this machine.

You have the following choices:

1) Set up root mirroring now, using
Solstice DiskSuite. After automatically
setting up mirroring, the system reboots.

2) Set up root mirroring later, using
Solstice DiskSuite. Run a single
command later to automatically set up mirroring.

3) Set up root mirroring on your own, using
a third-party product such as VERITAS.

Please select 1, 2 or 3:
```

Note – Root mirroring on the file server is strongly recommended so that the file server can recover from a single internal disk failure. If you do not want to use Solstice DiskSuite to configure root mirroring, use VERITAS Volume Manager or a similar third-party product to configure root mirroring after you have finished running the setup script.

- If you choose option 1, the system reboots and eventually displays the prompt shown below. Proceed from Step 6.
- If you choose option 2, you may invoke the root mirroring procedure at a later time using the `mirror_root` command. Note that this command causes the system to reboot. the system
- If you choose option 3, you should still configure root mirroring, using the documentation included with your choice of volume management software.

The system asks whether you want to configure the disk array(s) using Solstice DiskSuite software. The screen shows the following messages:

```
=====
Sun StorEdge Disk Array Configuration
=====

You have the following choices:

1) Run the Sun StorEdge configuration tool now.
This tool uses Solstice DiskSuite to set up
the Sun StorEdge disk array.

2) Run the Sun StorEdge configuration tool later.

3) Set up the Sun StorEdge disk array on your own,
using either Solstice DiskSuite or
a third-party product such as VERITAS.

Please select 1, 2, or 3:
```

6. Type the number that corresponds to your choice.

- If you choose option 1, the configuration script immediately proceeds with the configuration process. Refer to “To Set Up the Disk Arrays” on page 18.
- If you choose option 2, you may invoke the disk array setup procedure at a later time using the `storedge_config` command.
- If you choose option 3, you should still configure your disk arrays, using the documentation included with your volume management software.

At this point, the file server has restarted and is available for login. If you have not already configured the compute engines, do so before turning them on. Refer to “Configuring the Compute Engines” on page 23 for more information.

▼ To Set Up the Disk Arrays

You may perform this procedure as part of the setup script as described in Step 6 of the previous section, or you can set up the disk arrays from the system prompt by typing the `storedge_config` command. This section describes the setup process.

Note – Comments shown in **bold** type are examples of responses to the screen prompts. They do not automatically appear on the screens.

For more information on Solstice DiskSuite and how it works, refer to Chapter 6.

When you invoke the setup script on the file server, the screen displays the following message:

```
*****
          Configure Sun StorEdge A5200 Array
*****

Do you want to use Solstice DiskSuite to configure the Sun StorEdge
A5200 Array now?
(yes/no) yes
```

Note – The actual screen displays you see in this procedure differ slightly from the ones shown in this example. Your Sun TCF system has either one or two disk arrays; a two-disk system has 22 or 44 disks available. This example shows a system with only one disk array (22 disks).

1. To set up the disk arrays, type `yes` and press Return.

The following messages appear:

```
Scanning disks...

        Detected one A5x00 StorEdge A5x00 array.
        The following disks were found in your Sun StorEdge system.

c1t0d0  c1t1d0  c1t2d0  c1t3d0  c1t4d0  c1t5d0  c1t6d0
c1t7d0  c1t8d0  c1t9d0  c1t10d0
c2t16d0 c2t17d0 c2t18d0 c2t19d0 c2t20d0 c2t21d0 c2t22d0
c2t23d0 c1t24d0 c1t25d0 c1t26d0

        (Total Number of Disks=22)

        How many volumes do you want to create ?
? 3
```

2. Type the number of disk volumes that you want to create and press Return.

The following screen appears:

```
Please type the name of this volume's first mount point.
(e.g. /disk1) /data1
```

3. Type the name of the first mount point (for example, `/data1`) and press Return.

The following screen appears:

```
How many disks will be used in this volume?
(Exclude disks to be used for hot spares)
(You have 22 disks now.)
```

4. Type the number of disks you are using (for example, `12`) and press Return.

The following screen appears:

```
Please select one of these three choices for this volume:
1) Mirroring/Striping+Mirroring 2) Pure Concatenation 3) Striping
Note: If you choose mirroring the usable space is half the total
disk space.
        Your choice (1/2/3)?
```

5. Type the number that corresponds to the type of formatting that you plan to use with your disk arrays (for example, 1), and then press Return.

Note – The sample messages in the rest of this section show the responses that appear when you choose to mirror the disks (choice 1). The other formatting methods show similar messages.

The system responds with the following screen:

```
Your volumes are striped and then mirrored.

How many hot spare disks does this mirror volume contain?
(default 1) 2
```

6. Choose the number of disks that you want to set up as hot spares for your main system disks (for example, 2) and press Return.

The following message is displayed:

```
Please type the name of this volume's second mount point.
(e.g. /disk1)? /data2
```

7. Type the location for the second mount point (that is, the mount point for the hot spare disk).

The software displays the following message:

```
How many disks will be used in this volume?
(Exclude disks to be used for hot spares)
(You have 8 disks now.) 4
```

Note – Always use an even number of disks for a mirror partition. If you enter an odd number at this step, the installation script asks you for the formatting type, displays an error message saying Please use even number of disks for Mirror partition, and returns you to this step.

- 8. Type the number of disks you want to use for the hot spare (for example, 2) and press Return.**

The software displays the following screen:

```
Please select one of these three choices for this volume:
1) Mirroring/Striping+Mirroring 2) Pure Concatenation 3) Striping
Note: If you choose mirroring the usable space is half the total
disk space.
Your choice (1/2/3)?
```

- 9. Choose a formatting method for the second volume and press Return.**

The software displays the following screen:

```
How many hot spare disks does this mirror volume contain?
(default 1)
```

- 10. Type the number of hot spare disks for this volume and press Return.**

The software displays the following screen:

```
Please type the name of this volume's third mount point.
(e.g. /disk1)? /data3
```

- 11. Type the name of the third mount point and press Enter.**

The system displays the following message:

```
How many disks will be used in this volume?
(Exclude disks to be used for hot spares)
(You have 4 disks now.)
```

12. Type the number of disks you want to use for the volume and press Return.

The following messages are displayed.

```
Please select one of these three choices for this volume:
1) Mirroring/Striping+Mirroring 2) Pure Concatenation 3) Striping
Note: If you choose mirroring the usable space is half the total
disk space.
Your choice (1/2/3)? 3

Partitioning disks...

Please wait....
Done

The following metadevices have been created:

/dev/md/dsk/d1 :- Mirrored
/dev/md/dsk/d2 :- Concatenated
/dev/md/dsk/d3 :- Striped

Do you want to continue with this configuration?
(yes/no) yes
```

13. To continue the configuration process, type yes and press Return.

The following prompt appears:

```
Do you want to create file systems on these metadevices now?
(This might take several minutes.)
(yes/no) yes
```

14. To create the Unix File System (UFS), type yes and press Return.

The process may take from several minutes to about two hours. When it has completed, the following message appears:

```
Sun Storedge Array configuration completed
```

Configuring the Compute Engines

Note – If you are configuring a new expansion rack of compute engines to work with a previously installed file server, refer to “Overriding the Compute Engine Configuration” on page 28.

The file server stores two files that contain the configuration information for the compute engines. The following table describes these files.

TABLE 2-4 Compute Engine Configuration Files

Description	File Name	Purpose
Identity file	<code>/etc/suntcf/ hosts.conf</code>	Stores the identity information for the compute engines
NFS mounts file	<code>/etc/suntcf/nfs.conf</code>	Stores the NFS mount configuration

▼ To Edit the Configuration Files

Before you run the configuration script on the compute engines, perform the steps in this procedure.

1. Log into the file server as root.

If you do not know the root password, see your network administrator.

2. Open each of the files in TABLE 2-4 in the text editor of your choice.

3. Using the sample syntax in the files as an example, fill in the information for the compute engines that you recorded in TABLE 2-2 and TABLE 2-3.

The sample file for the `hosts.conf` file looks like this:

```
#
# Compute Engine autoconfiguration
# identification list
#
# format:
# <IP address> <hostname> <ethernet address>
#
# example:
#
# 10.192.186.21 farm02 8:0:20:ff:ee:dd:cc
#
# enter your information below
#
```

The sample file for the `nfs.conf` file looks like this:

```
#
# Compute Engine autoconfiguration
# NFS mount list
#
# format:
# <remote host> <remote dir> <local mount pt>
#
# example
#
# farm01 /export/tools /share/tools
#
# enter your information below
#
```

Note – If you enter information for NFS volumes that you have not yet created (for example, if you have not configured the disk array), you will need to remount the volumes on each compute engine after you create them. However, entering the information for those volumes in the `nfs.conf` file saves the trouble of having to manually enter the mounts in the `/etc/vfstab` file later on.

▼ To Run the Configuration Script

Note – Remember to leave the compute engines powered off until you have finished running the configuration script.

1. **To start the script, log in to the file server as superuser (if you have not already done so).**
2. **Type the following command at the prompt:**

```
# ce_config
```

The configuration script displays the following screen:

```
=====
#       Compute Engine       #
#   Autoconfiguration Setup   #
=====

Please provide the information necessary
for this tool to complete the autoconfiguration.
Edit the following two files, which contain
examples of the proper syntax.

/etc/suntcf/hosts.conf
    Contains the host name, IP address,
    and Ethernet address for each compute engine.

/etc/suntcf/nfs.conf
    Contains the remote hostname, remote
    directory, and local mount point for each
    NFS volume.
```

3. **Review the information in these files to ensure that it is correct before proceeding.** Use the text editor of your choice to open these files. Make sure that all Ethernet addresses are accurate, that every compute engine has been included in the identity file, and that all of the NFS mounts have been included in the NFS mounts file.

4. Once you have determined that the information is correct, type *y* to continue configuring the compute engines.

The script executes and shows messages similar to the following:

```
Using the following information:

Compute Engine identification
-----
comp-eng1, IP is 000.000.000.000, Ethernet address is
0:0:0:0:0:0:0

NFS configuration
-----
serv1:/data1 will be mounted on /share/data1
serv2:/data2 will be mounted on /share/data2

Continue? (y/n)
```

5. Type *y* to continue.

The script continues to execute. When it has finished, it shows the following messages:

```
*****
*      Compute Engine      *
*  Autoconfiguration Setup  *
*      Complete           *
*****

Power on the compute engines.
They automatically configure
themselves within ten minutes.

Once the compute engines are configured,
run the "smc_setup" command to add the compute
engines to the Solaris Management Console(TM).
You also need to manually add each compute engine
to the Sun Management Center(TM) domain. Please
consult the Sun TCF Product Overview and
Software Guide for more information.
```

The script then exits to the superuser prompt.

6. After the script has exited, turn on power to the compute engines.

After approximately 10 minutes, the compute engines are ready for access.

The configuration script defines the following settings for the compute engines:

TABLE 2-5 Compute Engine Configuration Settings

Feature	Setting
Name service	Same as for the file server. Refer to TABLE 2-1.
Root password	Same as for the file server.
Sun Management Center	Use the file server as the Sun Management Center server.
NFS mount points	Set by the configuration script.
Sun Grid Engine	The compute engine is an execution host for the Sun Grid Engine domain. The file server is the Sun Grid Engine master host.

When you log in to each compute engine, the configuration information is displayed on your terminal.

Completing the Setup Procedure

After you configure and power on the file server and all compute engines, perform the following steps.

▼ To Complete the Setup Procedure

1. Configure root mirroring on the file server, if you did not already do so.

This measure ensures that the file server can be recovered from the failure of a single disk.

2. Perform a backup of the file server.

It is strongly recommended that you have the root filesystem of the file server stored on backup media. If both internal drives ever fail at the same time, this backup copy will help you restore your root filesystem.

3. Configure the storage array, if it has not already been done.

Refer to “To Set Up the Disk Arrays” on page 18.

4. Set up the Solaris Management Console software

Before you can use Solaris Management Console on the TCF, make the compute engines accessible from the file server. For information on how to do this, refer to Chapter 3.

5. Configure the Sun Management Center software.

Make sure you perform the following steps to complete the configuration process. These procedures are described in detail in Chapter 4.

a. On the file server, edit the file `/etc/aliases`. For the two aliases `sunmc-alert` and `sunmc-critical`, add the email addresses of the people responsible for the maintenance of the TCF.

b. Invoke the Sun Management Center console, and perform a “discovery” to incorporate the compute engines into the domain of the file server.

When you complete these steps, you will have a fully-operational compute farm. You can begin to submit jobs to the TCF and monitor the jobs’ execution.

Overriding the Compute Engine Configuration

You can override the automatic configuration script for the compute engines. You need to override the configuration script if you are using an expansion rack with an older-model base rack, since the that base rack does not support automatic configuration.

▼ To Override the Configuration Script

1. Access the serial port of the compute engine by typing the following command from any system on the network:

```
% telnet servername port
```

where *servername* is the host name of the compute engine, and *port* is the port number for that compute engine.

2. Press Ctrl-Z.

The `telnet` prompt is displayed:

```
telnet>
```

3. Type `send brk` and press Return several times until the `ok` prompt appears on the screen.

If the `ok` prompt does not appear, repeat Steps 2 and 3 until it does.

4. At the `ok` prompt, type the following command and press Return:

```
ok> boot disk0:a
```

The compute engine then reboots and skips the automatic configuration process.

5. If you want to install one of the Asian locales on your compute engine, select the number that corresponds to the locale as shown in the following screen. The default value is 5 (no Asian locales).

The installation script automatically supplies the responses to the questions displayed on the screen. When the installation process is complete, the console login is displayed.

6. If you want to use Sun Management Center, provide the name of the Sun Management Center server (typically, this is the same as the file server). Otherwise, you can choose not to set up Sun Management Center at this time.

The compute engine is now ready to run. If you need to set up additional compute engines, repeat the procedure in this section for all of the compute engines in your expansion rack.

Recovering/Restoring the Compute Engines

The Sun TCF preinstalled software includes a recovery function that enables you to restore the operating system and software on each compute engine to its initial configured state (that is, to the state that the compute engine was in after it was first automatically configured). This is useful if the disk has been corrupted, or if the compute engine hardware needs to be replaced for some reason.

Note – Before a compute engine can be recovered, you need make sure that the compute engine configuration script has been run on this compute engine, and that its information is included in the identity file. Refer to “To Run the Configuration Script” on page 25 for more information.

If the root disk (`disk0`) on the compute engine is corrupted, replace it with a blank disk. Make sure that your blank disk is an approved replacement unit. Contact your Sun representative for information.

Note – The second disk on the compute engine (`disk1`) will be erased during this procedure.

▼ To Restore a Compute Engine

To restore a compute engine to its original settings, power on the compute engine to be recovered and perform the following steps:

1. **Access the serial port of the compute engine by typing the following from any system on the network:**

```
% telnet servername port
```

where *servername* is the host name of the compute engine and *port* is the port number for that compute engine.

2. **Press Return a few times, until you see the `ok` prompt on the screen.**
3. **Type the following command at the `ok` prompt:**

```
ok> boot net - install
```

The script reinstalls the operating system and bundled software, and then configures the compute engine according to the configuration files that were set when the SunTCF system was initially installed. The whole process takes about an hour.

If a compute engine can still boot, but you still want to erase both internal drives and start with a fresh installation, follow these steps:

1. **Log in to the compute engine as root.**

2. Type the following commands at the shell prompt and press return after each command:

```
# rm /var/sadm/system/admin/INST_RELEASE
# reboot -- "net - install"
```

The script reinstalls the operating system and preinstalled software, and configures the compute engine according to the configuration files set during the initial installation of the SunTCF system. The whole process takes about an hour.

▼ To Replace a Compute Engine

If you need to replace an entire compute engine (rather than just one disk), perform the following steps:

1. Use your favorite text editor to edit the compute engine identity file `/etc/suntcf/hosts.conf` and modify the Ethernet address setting for that host name to reflect that of the new system.

You do not need to modify any other entries in that file.

2. Rerun the `ce_config` script. Refer to “To Run the Configuration Script” on page 25.

Make sure you run the `ce_config` script before powering on the new compute engine. Once this has been done, follow the procedure in the previous section, “To Restore a Compute Engine.”

Reinstalling Solaris Software on the File Server

The preinstalled software image that you configured earlier in this chapter contains additional system patches and other applications not included in the media kits.

For information on reinstalling the Solaris operating environment on your file server, refer to the documentation for your version of Solaris software. After you have reinstalled the Solaris operating environment, you need to restore the software and files you backed up before reinstalling the Solaris software. Refer to the Solaris documentation for your software packages.

Note – If you are reinstalling the Solaris 2.6 operating environment and preinstalled software, you also need to download and install certain software patches on the file server. Refer to “Reinstalling the Software Patches on the File Server” on page 34 for more information.

The following software is preinstalled on the file server along with the Solaris 2.6 operating environment:

- Solaris Easy Access Server 3.0, which includes Solstice DiskSuite 4.2.1, Solstice AdminSuite 3.0, Netscape Communicator 4.51, and Solaris Management Console 1.0.1
- Sun Management Center 2.1.1 server software
- SunVTS 2.1.3
- Various patches and updates that support the Solaris operating environment, the disk arrays, and the Fibre Channel card

The following table shows which version of management tools to use with the Solaris 2.6, Solaris 7, and Solaris 8 operating environments.

TABLE 2-6 Preinstalled Software Interoperability/Compatibility Matrix

Application	Solaris 2.6	Solaris 7	Solaris 8
Sun Management Center	Version 2.1.1	Version 2.1.1	Version 2.1.1
Solaris Easy Access Server	Version 3.0	Version 3.0	Version 3.0 Update
Solstice DiskSuite	Version 4.2	Version 4.2	Version 4.2.1

Note – Although the file server and compute engines support Solaris 7, it is not included with the Sun TCF system. If you do decide to install Solaris 7 on Sun TCF, it is recommended that you purchase a separate support contract for Solaris 7.

Installing Solaris 8 Software



Caution – If you install Solaris 8, the preinstalled software functions (such as compute engine recovery) will not work. You will need to manually set up those functions in Solaris 8.

If you want to install Solaris 8 on your file server, you need to install the operating environment from the Solaris 8 CD. Refer to the Solaris 8 documentation (included in the Solaris 8 Media Kit) for the procedures.

The following table describes the versions of the preinstalled software included in the Solaris 8 Media Kit, as well as the location of each application.

TABLE 2-7 Solaris 8 Version Compatibility Matrix

Application	Location
Solstice DiskSuite 4.2.1	CD 2 of 2, in the directory Solaris_8/EA/products
Solaris Management Console 2.0	CD 2 of 2, in the directory Solaris_8/EA/products
SunVTS 4.1	Solaris supplement CD
Sun GigabitEthernet 3.0 driver (required on the file server)	Solaris supplement CD
Solstice AdminSuite 3.0.1	Download this version from http://www.sun.com/bigadmin/ content/adminPack

Sun Grid Engine 5.2 and Sun Management Center 2.1.1 software come in separate kits with the Sun TCF software and documentation.

Note – With the Solaris 8 operating environment, the functionality in Solaris Easy Access Server 3.0 is included in the Solaris 8 Admin Pack. You can download the Solaris 8 Admin Pack from the following Web site:

<http://www.sun.com/bigadmin/content/adminPack>

Note – Before you start the process of installing Solaris 8 software, be sure to read the upgrade information in the Solstice DiskSuite documentation. You will need to perform some tasks in Solstice DiskSuite before you install the operating environment.

Reinstalling the Software Patches on the File Server

Your Sun TCF system comes fully configured, including all the necessary patch updates for the Solaris 2.6 operating environment and the preinstalled applications. If you ever need to reinstall the Solaris 2.6 operating environment and applications on the file server from the Sun TCF Media Kit, you need to download and install the patches. The patches are not included in the Sun TCF Media Kit.

You can download the patch from the following location:

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

The following tables list the patches required by the Sun TCF hardware and the preinstalled software. These patches are installed along with the preinstalled software; however, the Solaris 2.6 operating environment may require that you download additional patches for your other hardware and software applications.

TABLE 2-8 Patches for Solaris 2.6 Software

Patch Number	Product	Description
105181-23	Solaris 2.6	SunOS 5.6 kernel update patch
106764-07	Solaris 2.6	Sun Gigabit Ethernet 2.0 driver patch
105356-16	Fibre Channel card driver	/kernel/drv/ssd and /kernel/drv/sd patch
105357-04	Fibre Channel card driver	/kernel/drv/ses patch
107280-07	Fibre Channel card driver	/kernel/drv/ifp patch
105375-25	Fibre Channel card driver	sf and socal driver patch
106219-03	luxadm	man page patch
106387-04	SunVTS 2.1.3	Security, Netra™ ft 1800 support, pmem panic, and test fixes
106810-01	SunVTS 2.1.3	Update afbtest and m64test
105742-01	SunVTS 2.1.3	cpudatatest update
109590-01 (Solaris 2.6)	Sun Management Center	A5x00 agent support
106627-06	Solstice DiskSuite 4.2	Product patch

TABLE 2-9 Patches for Solaris 7 Software

Patch Number	Product	Description
105356-16	Fibre Channel card driver	/kernel/drv/ssd and /kernel/drv/sd patch
105357-04	Fibre Channel card driver	/kernel/drv/ses patch
107280-07	Fibre Channel card driver	/kernel/drv/ifp patch
105375-25	Fibre Channel card driver	sf and socal driver patch
106387-04	SunVTS 2.1.3	Security, Netra ft 1800 support, pmem panic, and test fixes
106810-01	SunVTS 2.1.3	Update afbtest and m64test
109589-01 (Solaris 7)	Sun Management Center	A5x00 agent support
106627-06	Solstice DiskSuite 4.2	Product patch

TABLE 2-10 Patches for Solaris 8 Software

Patch Number	Product	Description
109590-01 (Solaris 8)	Sun Management Center	A5x00 agent support
106627-06	Solstice DiskSuite 4.2	Product patch

These patches are installed on the file server. Only the SunVTS patches are also installed on the compute engines.

Note – The Solaris 2.6 Media Kit includes a CD(0) disc in case you need to reinstall Solaris 2.6 software. If the preinstalled software image setup routine fails for any reason, you can use this disc to install Solaris 2.6 software on the Sun TCF system.

Note – Solaris 7 software is not included with the Sun TCF Media Kit.

Using Solaris Management Console Software

Solaris Management Console serves as the main user interface for the other applications in the Sun TCF system. This program provides a single location in which to find all the system management applications you need without having to log in to each individual system in the TCF. Use Solaris Management Console as a portal for all the system management GUIs, such as AdminSuite, `qmon` (the Sun Grid Engine graphical user interface), Solstice DiskSuite, and so on.

This chapter contains the following topics:

- Configuring and Running Solaris Management Console Software on page 37
- For More Information on page 42

Configuring and Running Solaris Management Console Software

Before You Start

- Before invoking Solaris Management Console for the first time, add the compute engines to the console by typing the following command at the file server system prompt:

```
% smc_setup
```

Note – This command incorporates all of the compute engines that were included in the identity file during the compute engine configuration procedure. Any system that was not automatically configured must be added manually. The next section describes how to manually add compute engines. For more information on the configuration script, refer to Chapter 2.

Always launch Solaris Management Console from the file server as the first application. You can launch the other applications from within the Solaris Management Console user interface.

▼ To Start Solaris Management Console

- **To start Solaris Management Console, type the following command at the system prompt:**

```
% smc &
```

Note – In some cases, it might not be possible to log into Solaris Management Console, due to the server process not starting properly. If this is the case, enter the following commands at the system prompt:

```
% /etc/init.d/ehttpd stop  
% /etc/init.d/ehttpd start
```

You should then be able to log into Solaris Management Console.

FIGURE 3-1 shows the main screen of the Solaris Management Console program.



FIGURE 3-1 Solaris Management Console Main Screen

Working With Sun Management Console

The Solaris Management Console opening screen lists tools and utilities by server. The left side of the window displays all the management applications, grouped by functionality type. The right side displays the applications in one specific group.

▼ To Add an Application to Sun Management Console

1. **To find an application installed on a server, double-click on the folder icons.**

FIGURE 3-2 shows the Add Application dialog box.

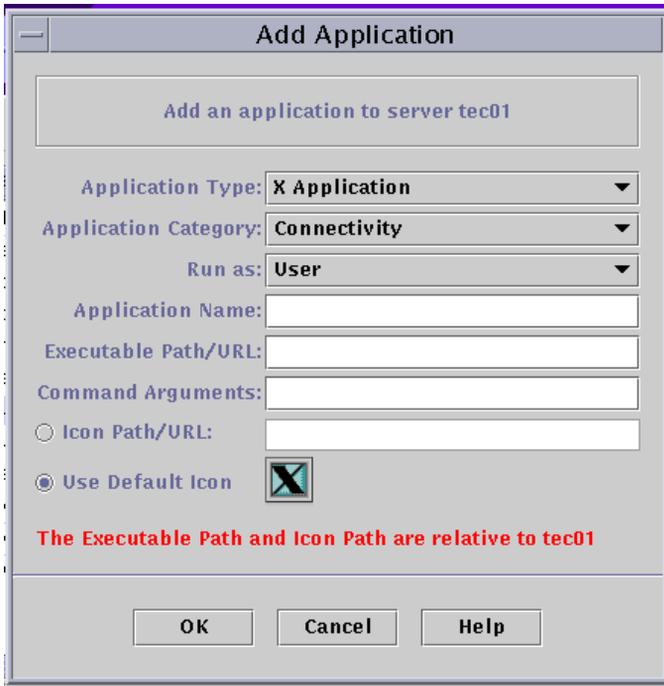


FIGURE 3-2 Add Application Dialog Box

2. Click on the drop-down menus to locate the application.
 3. When you have finished, click OK.
- ▼ To Start an Application from Sun Management Console
- To start the application, double-click the application icon on the right side.
- FIGURE 3-3 shows the Applications View in Sun Management Console.



FIGURE 3-3 Applications View Screen

▼ To Manually Add a Machine to Solaris Management Console

Solaris Management Console allows you to manage multiple machines from one console. In the Sun TCF system, you can use Solaris Management Console to monitor your file server, your terminal server, and the compute engines.

It is recommended that you operate Solaris Management Console from the file server and add the compute engines to it using the Add SMC option. Configuring Solaris Management Console in this fashion allows you to manage the entire compute farm from one location.

To add a file server or compute engine to Solaris Management Console, perform the following step:

- **From the Console menu, select the Add Server option.**

▼ To Use the Solaris Management Console Log File

- **Click the Show Log button to view a log of all Solaris Management Console activities.**

For More Information

Solaris Management Console does not include printed documentation. For more information on Solaris Management Console, refer to the application's online help or refer to the `docs.sun.com` Web site.

Using Sun Management Center Software

Sun Management Center software is a client/server application for monitoring a group of systems from one central system. The full Sun Management Center software (server, agent, console, and Help) is installed on the Sun TCF file server. The Sun Management Center agents (clients) are installed on the compute engines.

This chapter discusses the following topics:

- Sun Management Center Functions on page 43
- Configuring Sun Management Center on page 44
- Running Sun Management Center on page 52
- Preconfigured Alarms on page 54
- Monitoring the Performance of Your Sun TCF System on page 59

Sun Management Center Functions

You can use Sun Management Center to perform the following administrative functions for your Sun TCF system:

- Monitor the performance of functions and peripherals.
- Track information such as CPU load average, I/O transactions, and disk capacity.
- Automatically collect data and display a history as a series of graphs in order to view system trends.
- Create and maintain a physical inventory of client systems.
- Monitor software or hardware alarms or problems that might occur in the Sun TCF system.
- Perform actions automatically triggered by specific conditions.

- Monitor the functioning of workload management software, such as Sun Grid Engine or LSF from Platform Computing.

For more information on configuring and using Sun Management Center, refer to the *Sun Management Center User's Guide*.

Note – It is recommended that you install the Sun Management Center console on another workstation on your network. Since the file server does not have a monitor, you should use the additional workstation to monitor the compute farm with Sun Management Center. Installing the console on a second workstation improves the system's performance and eliminates the need for `xhost` connections. You do not need to install server/agent support on the second workstation.

To take full advantage of the Sun Management Center monitoring facilities, make sure the program's agents are running on each of the compute engines in your compute farm. If not, the Sun Management Center software can obtain only general information about each of the compute engines. Sun Management Center is installed automatically on each of the compute engines as part of the preinstalled software image.

If you have installed third-party workload management software (such as LSF) on your Sun TCF system, contact your workload management software vendor to obtain the modules that integrate the WMS with Sun Management Center.

Use Sun Management Center 2.1.1 with the Solaris 2.6, Solaris 7, and Solaris 8 operating environments.

Configuring Sun Management Center

After you turn on and configure the file server and all compute engines, you need to perform some final steps to complete the setup of Sun Management Center on the TCF.

Configuring Email Notification of Alarms

Your Sun TCF system is configured with preset alarms on the compute engines. These alarms send email to two aliases on the file server, serving as central points of notification. By default, these email messages are sent to `root` on the file server, but it is advisable to add the email addresses of the TCF administrators so that the

administrators are directly notified of alarms. The subject line and body text of the email message contains the hostname of the affected system and a short description of the condition.

Two kinds of alarms have been configured: *alert* and *critical*. The alert alarm is for noncritical situations that require further investigation. For example, if the system load on a compute engine were to exceed 6, the compute engine would still be functional, but its performance might be compromised.

The critical alarm applies to situations that warrant immediate attention. For example, if the root file system on a Sun TCF system exceeds 95% capacity, you need to take immediate action before it fills up completely.

▼ To Add Email Addresses to the Sun Management Center Alarm Aliases

1. **Log into the file server and open the file `/etc/aliases` in the text editor of your choice.**
2. **At the bottom of the file, look for the lines containing `smc-critical` and `smc-alert`. After each colon, add a list of email addresses to which you would like those emails sent. Separate each email address with a comma.**
3. **Save your changes and exit the text editor.**
4. **Type `newalias` at the command prompt to update and save the email alias.**

Incorporating Compute Engines Into the File Server's Domain

The file server also acts as the Sun Management Center server in your Sun TCF system. Before monitoring can take place, the Sun Management Center server needs to be aware of all the machines that it will be monitoring. The group of machines monitored by a Sun Management Center server is referred to as a *domain*. The compute engines are included in the file server's domain by performing a process called a *discovery*.

▼ To Perform a Discovery of the Compute Engines

1. **Invoke the Sun Management Center console by double-clicking the Sun Management Center console icon within the Solaris Management Console.**



FIGURE 4-1 Sun Management Center Login Screen

2. Log in using the following values:

User name: root

Password: your chosen root password

Server: the host name of your file server

The main window of the Sun Management Center console appears.

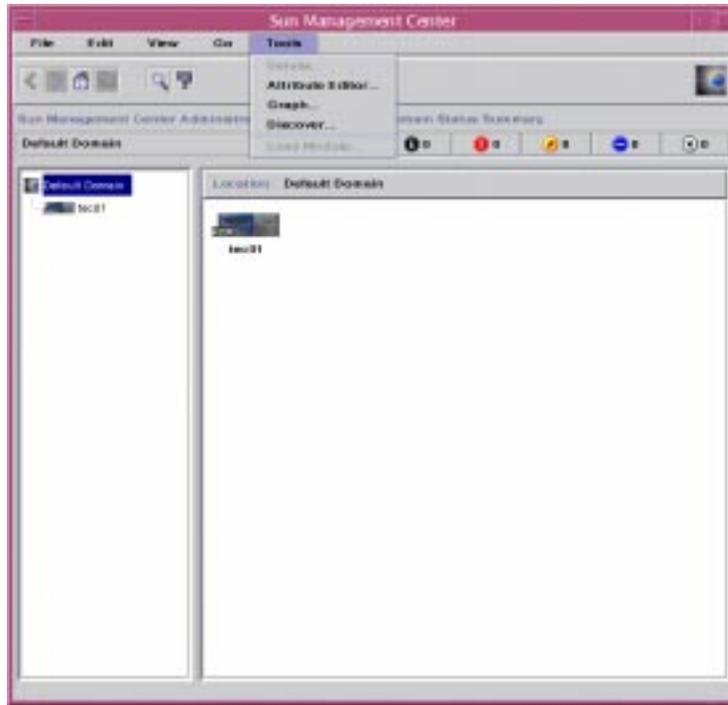


FIGURE 4-2 Sun Management Center Main Window

3. Under the Tools menu item, select Discovery.

The following screen is displayed:

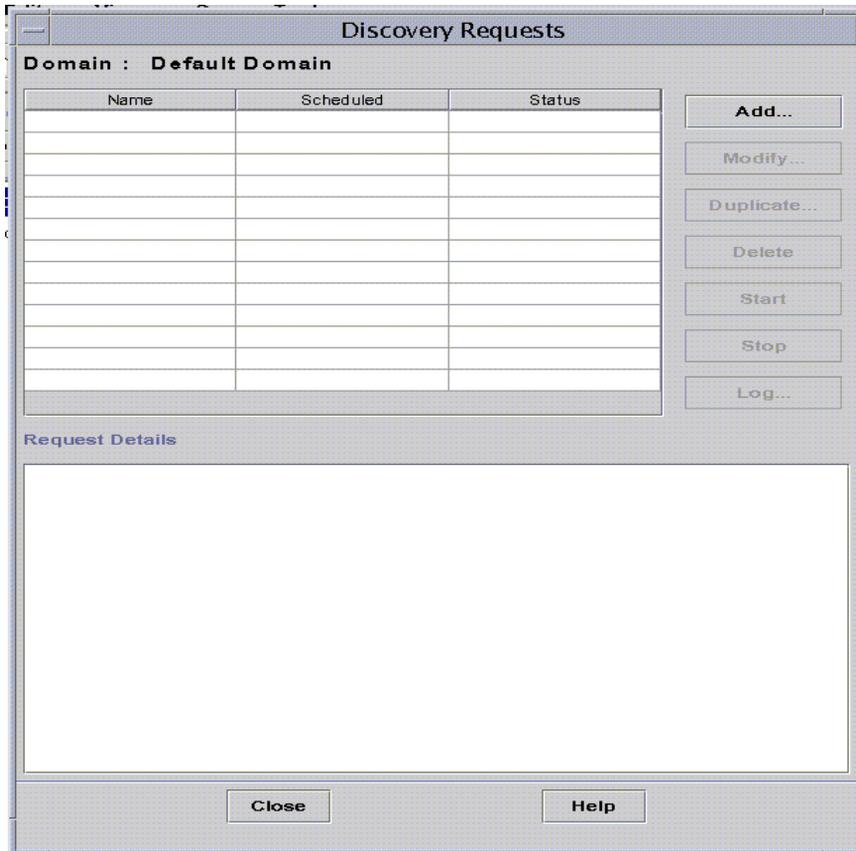


FIGURE 4-3 Discovery Request Dialog Box

4. Click on Add.

The following screen is displayed.

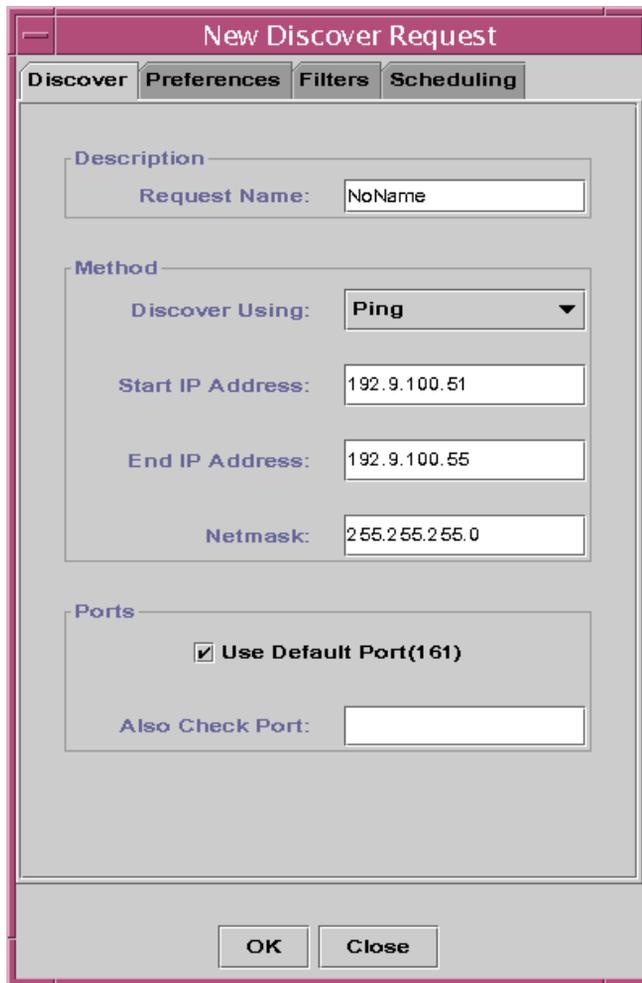


FIGURE 4-4 New Discovery Request Dialog Box

5. Fill in the range with the range of IP addresses used for the compute engines and click OK.
6. When the system asks Do you want to start the request upon creation?, click Yes.



FIGURE 4-5 Confirm Dialog Box

The discovery screen shows the progress of the just-created discovery.

- 7. When the discovery screen returns the message `Complete`, close this window.**

The screen appears similar to the following:



FIGURE 4-6 Discovery Window

8. Double-click on the icon indicating the subnet.

You should see a screen with each compute engine indicated by a representation of a Sun Enterprise 420R server.

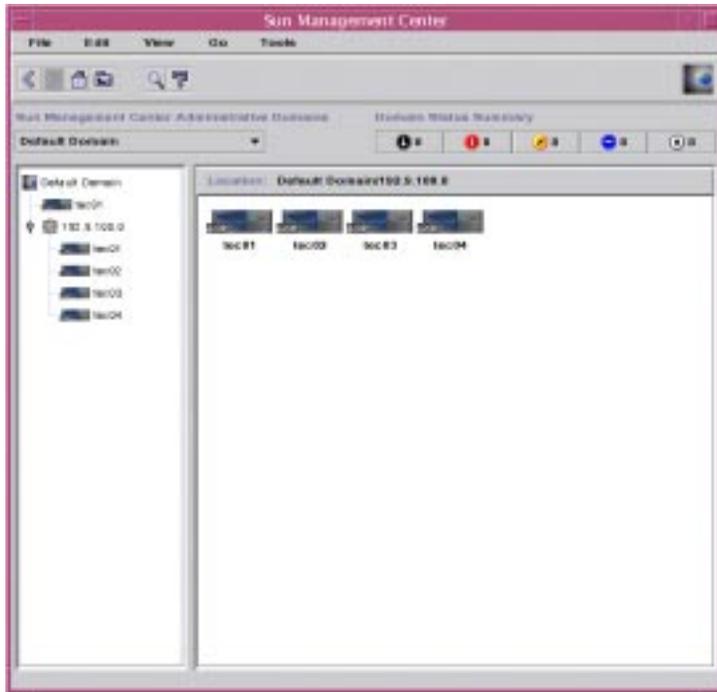


FIGURE 4-7 Compute Engines in the Sun Management Center Domain

At this point, the setup of Sun Management Center is complete. To view the state of any of the machines in the Sun TCF system, double-click on the corresponding server icon. Use the browser to explore any conditions reported by each module.

Running Sun Management Center

To start the Sun Management Center application, use the Solaris Management Console user interface. Refer to Chapter 3 for more information.

Using the Browser

The Sun Management Center Browser enables you to gather and examine information about your Sun TCF system. This section contains some brief procedures that show you how to use the browser.

▼ To Use the File Scanning Module

1. Double-click on the icon representing the compute engine you want to scan.
2. Click on the **Browser** tab at the top of the window.
3. Click on **Local Applications** in the left side of the window, and then click on **File Scanning**.
4. Double-click the **File Scan** folder icon.

The following information appears in the browser window:

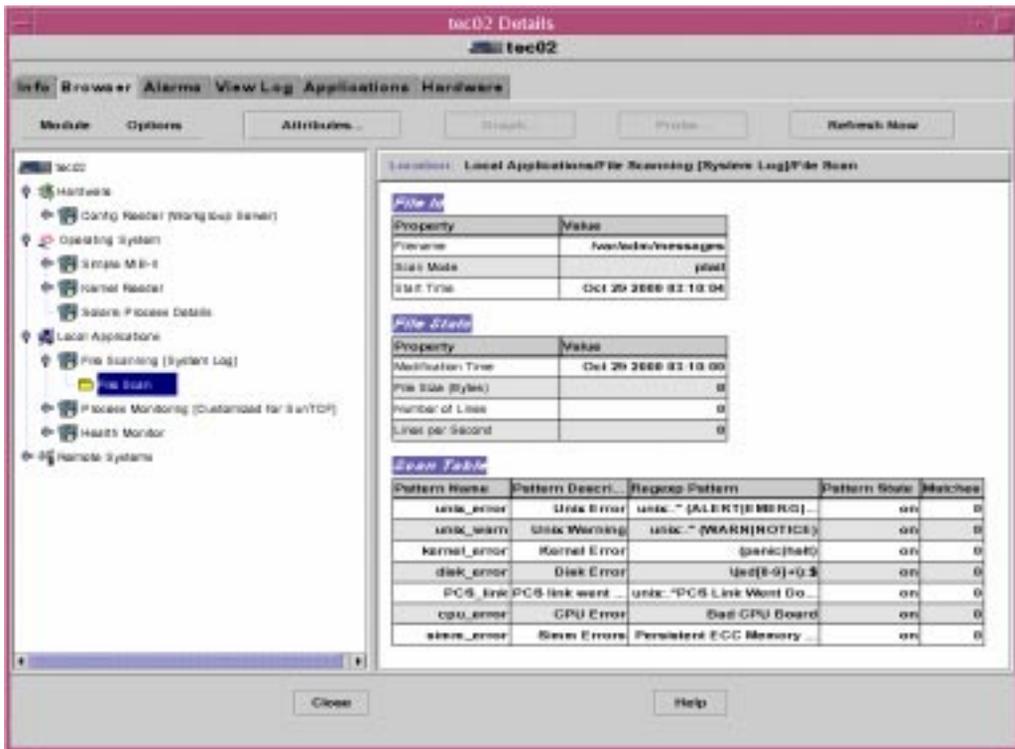


FIGURE 4-8 File Scanning Window

▼ To View the Applications

1. Double-click on the icon representing the compute engine you want to scan.
2. Click on the **Applications** tab at the top of the window.

The following window is displayed:

TABLE 4-1 Preconfigured Alarms

Module	Entry	Criteria	Alarm Type
File Scanning	unix_error, kernel_error, disk_error	Occurs = critical	Critical alarm
	unix_warn	Occurs = alert	Alert alarm
	PCS_link_down	Occurs = alert	Alert alarm
	Bad CPU Board	Occurs = critical	Critical alarm
	Persistent ECC Memory Error	Occurs = critical	Critical alarm
Health Monitor	All standard	All standard	Not set, since alarm conditions are transitory/advisory
Solaris Process Detail	Standard	Standard	None

You may modify these settings at any point, as well as add any additional modules and alarms. You should consider incorporating the Sun Grid Engine module information into the Sun Management Center domain. Refer to Chapter 5 for more information.

▼ To View the Alarms

- 1. Double-click on the icon representing the compute engine you want to scan.**
The Details window displays.
- 2. Click on the Alarms tab at the top of the Details window.**
The list of all current alarms appears in the window.

▼ To View the Alarm Configuration

- 1. Double-click on the icon representing the compute engine you want to scan.**
- 2. Click on the Browser tab at the top of the window.**
- 3. Click on Operating System in the left side of the window, and then click on Kernel Reader.**

4. Double-click on the System Load icon.

The System Load window appears.

5. Click the right mouse button on the System Load window.

A pop-up menu appears, similar to the following:

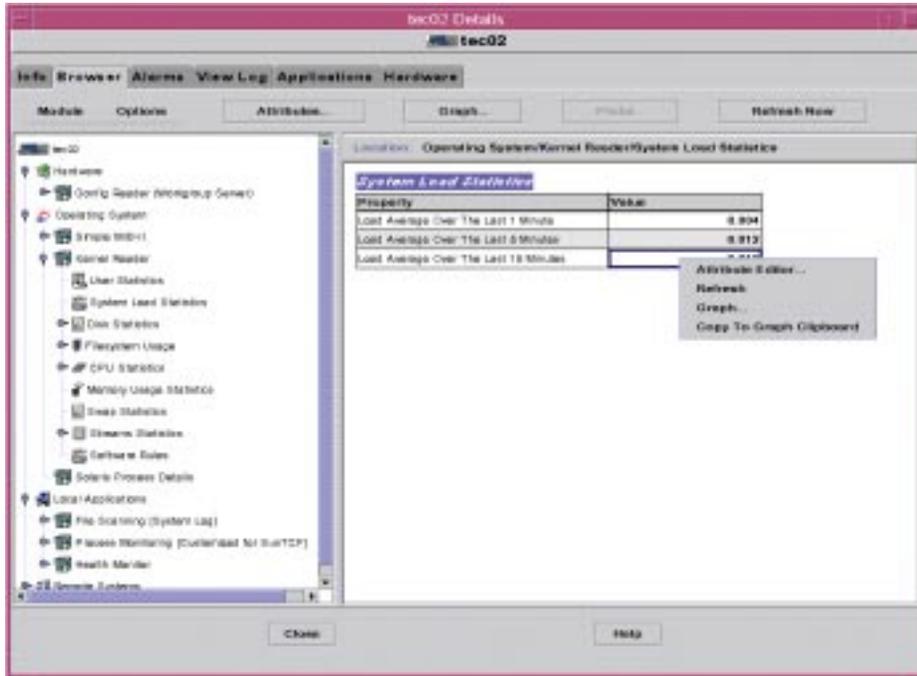


FIGURE 4-10 Pop-Up Menu

6. Select Attributes.

7. Click on the Alarms tab. The following information appears:



FIGURE 4-11 Alarm Threshold Information

You can add to or modify the alarm thresholds from this screen.

8. When you have finished editing the threshold information, click **Apply**.

9. To specify what actions to take when alarm conditions exceed the thresholds, click on the **Actions** tab.

The following screen is displayed:

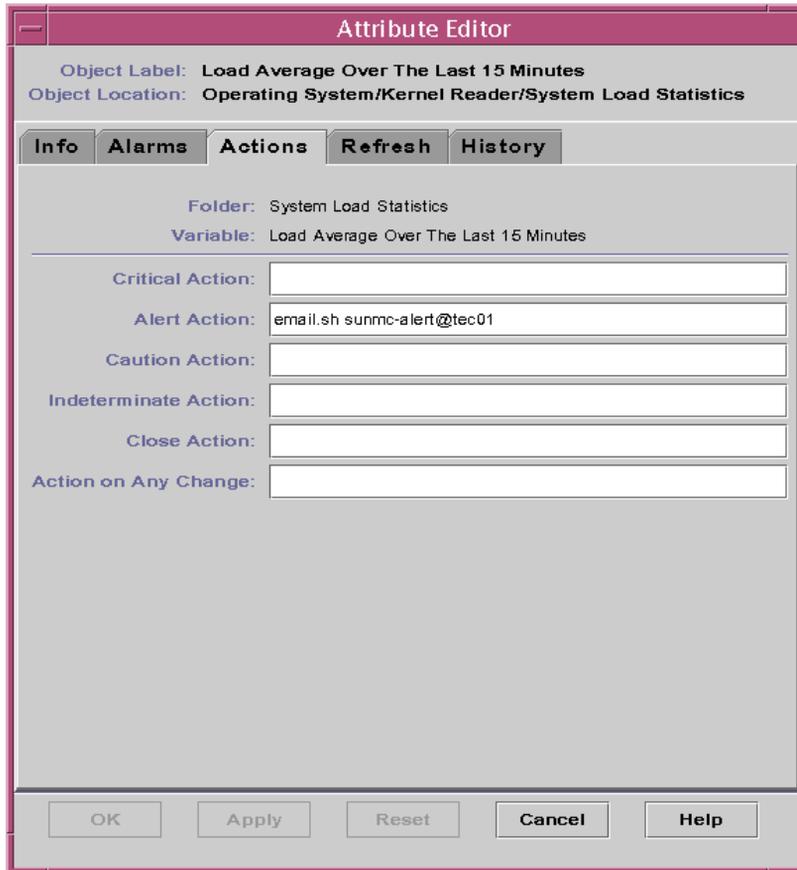


FIGURE 4-12 Actions Dialog Box

10. When you have finished entering your alarm action information, click **Apply**, and then click **OK** to close the dialog box.

Monitoring the Performance of Your Sun TCF System

To get the most from the Sun TCF system, it is important to monitor its performance in the actual environment in which it is being used. You can easily monitor all nodes through the Sun Management Center software.

Different applications have different requirements. By observing how the nodes (file server and compute engines) perform, you can make optimizations that are specific to the installation.

Note – Often, it is more effective to tune the workload management software (such as Sun Grid Engine), not the operating environment. Tuning the workload management software is a complex task, and it is highly dependent upon your specific environment. Refer to the Sun Grid Engine documentation for ideas on how to tune it for your needs.

▼ To View Performance Statistics for a Particular Machine

1. **Double-click on the icon representing the compute engine you want to monitor.**
2. **Click on the Browser tab at the top of the window.**
3. **Click on Operating System in the left side of the window, and then click on Kernel Reader.**
4. **Double-click on the System Load icon. Information similar to the following is displayed for your compute engine:**

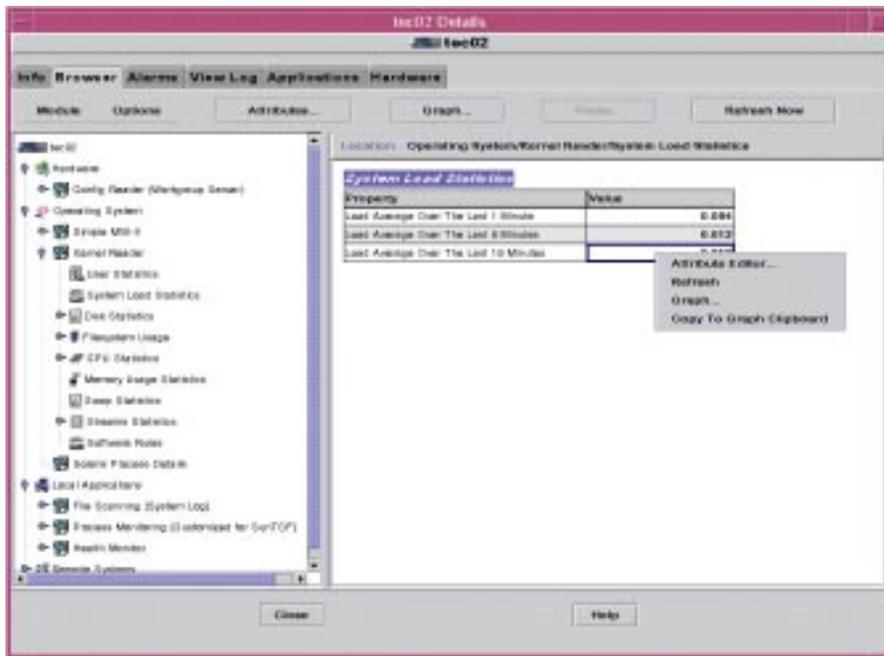


FIGURE 4-14 Pop-Up Menu

- To graph the system load performance over time, select Graph from the pop-up menu.

For more information, refer to the Sun Management Center documentation.

Some Statistics to Monitor

The following is a brief list of performance statistics that are worth monitoring, and what they indicate.

TABLE 4-2 Performance Statistics

Condition	Explanation	Remedy
High % Idle Time on a compute engine (low % CPU Utilization value)	The compute engine is not being used effectively.	Modify the configuration of your workload management software to better allocate resources for compute jobs. For example, if the jobs are single threaded and individually consume less than 1GB of RAM, then you should allow four jobs to run on one compute engine.
High % CPU Wait Time value on the file server or on the compute engines	Possible performance problems with the storage array, network switch, or network interface(s)	Try using <code>cacheFs</code> to help with the storage array performance. Refer to “Optimizing I/O Performance” on page 77 for more information.
High Disk Service Time value on the file server	Possible performance problems with the storage array. If the Disk Busy percentage is consistently high, or if the Average Service Time is showing large spikes relative to the average, it could indicate an overloaded disk array.	Some solutions for this situation include the following: <ul style="list-style-type: none">• Move some data to another file server• Modify the parameters used in setting up the metadvice for the disk array(s)• Use <code>cacheFs</code> for data that is usually not written (only read).

Using Sun Grid Engine Software

This chapter discusses the following topics:

- Preconfigured Software Settings on page 63
 - Working With Sun Grid Engine on page 64
 - Sun Grid Engine Functions on page 65
 - For More Information on page 68
-

Preconfigured Software Settings

Your SunTCF system comes with Sun Grid Engine preinstalled and preconfigured. The following table summarizes the settings.

TABLE 5-1 Sun Grid Engine Settings

Sun Grid Engine Information	Preconfigured Setting
Installation directory	<code>/share/gridware</code>
Administrator user account	<code>grdadmin</code>
TCP communication port	536
Cluster name	default
Queues	One queue per compute engine, each with 4 job slots. Queues are named as <code>hostname.q</code> , where <i>hostname</i> is the name of the compute engine.



Caution – The password for the administrator user (`grdadmin`) is not set by default. You should log in as root and set the password for this account as soon as possible. If you do not set the password for `grdadmin`, your system could be vulnerable to a breach in security.

Working With Sun Grid Engine

Sun Grid Engine allows you to submit and monitor jobs to the TCF from any system on the network. The only requirement is that you must mount the shared directory that contains the Sun Grid Engine installation on the system (in this example, the shared directory is `/share/gridware` on the file server).

Setting Up Your Environment

To work with Sun Grid Engine, you must first set up your environment. The way this is done depends on your current shell environment.

- If you are operating under the C shell or TC shell, type the following command at the system prompt:

```
% source /share/gridware/default/common/settings.csh
```

- If you are using the Korn, Bash, or Bourne shell, type this command:

```
$ . /share/gridware/default/common/settings.sh
```

Once you set up your environment, you can use the commands described in this section.

Note – None of the commands in this section work unless you perform this setup step.

Sun Grid Engine Functions

Sun Grid Engine 5.2 (formerly known as Codine) uses these main commands for job submission:

- `qmon`
- `qsub`
- `qrsh`
- `qtcsh`

Each of these commands has a related `man` (manual) page. You should also refer to the `man` pages for `submit(1)` and `qtcsh(1)` for more details on how those commands work with the submission commands.

The Sun Grid Engine Web site contains more information about the product:

`http://www.sun.com/gridware`

Sun Grid Engine Applications

The following sections provide brief explanations for how the submission commands work. For more information on these commands, refer to the Sun Grid Engine documentation.

`qmon`

The `qmon` application serves as the GUI (Graphical User Interface) to Sun Grid Engine. It allows you to submit jobs, as well as monitor their progress and change settings on your TCF system.

Note – Many of the administrative features of `qmon` are only available if you run this tool as the `grdadmin` user (as opposed to running as `root` or `superuser`).

- **To invoke the GUI, type `qmon` at the system prompt.**

You can also start `qmon` from within the Solaris Management Console.

qsub

Use `qsub` for traditional batch submit operations (where I/O is directed to a file). Note that `qsub` only accepts shell scripts, and not executable files. Sun Grid Engine contains an application script, `qs`, that allows `qsub` to accept executable files directly.

qrsh

The `qrsh` command is similar to the Solaris `rsh` command, except that a host name is not needed. Instead, `qrsh` runs a shell script or an executable file, potentially on any node in the cluster. I/O is directed back to the submitter's terminal window.

By default, if the job cannot be run immediately, `qrsh` will not queue the job. To allow jobs to queue, use the `-now no` flag with `qrsh`. Note that I/O can be redirected with the shell redirect operators. For example, to run the `uname -a` command under `qrsh`, type the following command at the prompt:

```
% qrsh uname -a
```

The scheduler selects a machine in the cluster and then displays the `uname` information for that machine to the submitting terminal. To redirect the output to a file, type the following command:

```
% qrsh uname -a > /tmp/myfile
```

The output from the `uname` command is then written to `/tmp/myfile` on the submitting host. To allow the `uname` command to queue, you would then type the following:

```
% qrsh -now no uname -a
```

If a suitable host is not immediately available, the `qrsh` command waits to execute the `uname` command until a suitable host becomes available. At that time, the command output will be displayed on the submitting terminal.

qtcsh

The `qtcsh` command is a modified version of `tcsh`. This command automatically submits jobs that are listed in a task file to the cluster. See the `qtcsh(1)` man page for more information on how to use this command.

▼ To Temporarily Stop Sun Grid Engine

If you want to temporarily stop Sun Grid Engine (for example, if you want to use another workload management program), perform these steps:

1. **Log in to the file server as the Sun Grid Engine administrator user (`grdadmin`).**
2. **Type the following commands at the prompt:**

```
% qconf -ke
% qconf -ks
% qconf -km
```

These commands cause Sun Grid Engine to cease normal operations, such as scheduling and executing jobs.

Note – If the file server or any of the compute engines is rebooted, Sun Grid Engine will start up again automatically. This means that the commands in Step 2 would have to be re-executed.

If you want to disable Sun Grid Engine, log in as `grdadmin` to each individual machine (the file server and all of the compute engines) in the Sun TCF system and follow this procedure:

1. **Type the following command at the prompt and press Return:**

```
% mv /etc/rc2.d/S95codine5 /etc/rc2.d/s95codine5
```

2. **Reboot the system.**

▼ To Restart Sun Grid Engine

- **To restart the Sun Grid Engine software without rebooting a system, log in as root to the file server and each compute engine, and type the following command at the prompt:**

```
# /share/gridware/default/common/codine5
```

If you disabled Sun Grid Engine as described in the previous section, log in as root to each individual machine (file server and all of the compute engines) and perform these steps to reenble the program:

1. Type the following command at the prompt and press Return:

```
# mv /etc/rc2.d/s95codine5 /etc/rc2.d/s95codine5
```

2. Type this command to restart Sun Grid Engine:

```
# /share/gridware/default/common/codine5
```

For More Information

For more information on Sun Grid Engine, refer to the Codine 5.2 documentation included on the Sun Grid Engine CD.

Using Solstice DiskSuite Software

This chapter discusses the following topics:

- Solstice DiskSuite Functions on page 69
 - Starting the Solstice DiskSuite Software on page 70
 - Understanding Solstice DiskSuite Metadevices on page 70
-

Solstice DiskSuite Functions

Solstice DiskSuite is used to create and manage complex storage structures such as disk mirrors, stripes, RAIDs, and hot spare pools. This application is installed only on the file server.

In addition, Solstice DiskSuite contains data storage software applications that enable you to perform the following tasks with your Sun TCF disk arrays:

- Configuring devices
- Volume/disk management

The following table shows which version of Solstice DiskSuite to use with the Solaris 2.6, Solaris 7, and Solaris 8 operating environments.

TABLE 6-1 Solstice DiskSuite Interoperability/Compatibility Matrix

Application	Solaris 2.6	Solaris 7	Solaris 8
Solstice DiskSuite	Version 4.2	Version 4.2	Version 4.2.1

Solstice DiskSuite is included in the preinstalled software image, as a part of the Solaris Easy Access Server 3.0 application suite. In addition to Solstice DiskSuite 4.2, this suite includes Solstice AdminSuite 3.0, and Solaris Management Console 1.0.1. Updates to the software are available for downloading at <http://www.sun.com>.

For more information on Solstice DiskSuite, refer to the Solstice DiskSuite documentation.

Starting the Solstice DiskSuite Software

You can start Solstice DiskSuite from within the Sun Management Center interface, or you can run it from the command line.

▼ To Start Solstice DiskSuite From the Command Line

- Type the following command:

```
% /usr/opt/SUNWmd/sbin/metatool &
```

The Solstice DiskSuite main screen is displayed.

In addition to the graphical interface, Solstice DiskSuite provides a set of commands that you can use to manipulate the physical or virtual storage structures on your network (referred to in Solstice DiskSuite as *metadevices*). The *Solstice DiskSuite User Guide* contains complete information on both the usage and the syntax of these commands.

Understanding Solstice DiskSuite Metadevices

Solstice DiskSuite metadevices can be physical devices, or they can be virtual storage structures on your Sun TCF system. These devices can include the following:

- Mirrors

Disk mirrors are two or more disks that contain identical information at all times so that if one fails, the other can take over instantly. Even root disks can be mirrored.

- Stripe and concat devices

A stripe or concat (for concatenation) device is used to create large logical partitions from several smaller disks. Stripes and concats form the building blocks for the other types of metadevices.

- Hot spare pool

Solstice DiskSuite allows you to create a pool of hot spare disks. In the event of a failure, Solstice DiskSuite pulls a functional disk from the hot spare pool and replaces the failed disk.

- Trans devices

A trans device is used for UFS logging. A trans device reduces time that a system spends on the `fsck` routine during boot. Usually, a trans device contains both a master device and a logging device.

- RAID5 metadevices

A RAID5 metadevice uses striping plus additional parity data to provide redundancy. This can allow the loss of one or more disks without permanent data loss.

Note – Trans devices and RAID 5 metadevices are not recommended for the Sun TCF system, even though the Solstice DiskSuite software contains these capabilities.

Root Mirroring

Root mirroring is an essential part of the Sun TCF architecture. It helps provide high availability on the file server, and protects against permanent data loss in the event of a disk failure.

In a compute farm environment, high availability is not needed on compute engines, since there are multiple of these; in addition, data loss on a compute engine is not critical, since they do not contain any permanent data. However, it is critical that the file server remain available at all times, and any settings you have made on the file server need to be preserved in the event of a disk failure. Root mirroring provides a solution to both of these issues.

Using Solstice AdminSuite Software

This chapter discusses the following topics:

- Solstice AdminSuite Overview on page 73
- Setting Up and Running Solstice AdminSuite on page 74

Solstice AdminSuite Overview

Use Solstice AdminSuite software to add, remove, and modify user and group accounts, remote hosts, and system resources such as networks and disks. This application also allows you to modify the system resources, such as mailing lists and NFS administration.

Solstice AdminSuite is part of the Solaris Easy Access Server 3.0 product (included in the preinstalled software image). In addition to Solstice AdminSuite, SEAS includes Solstice DiskSuite 4.2 and Solaris Management Console 1.0.1. Updates to the SEAS software are available for downloading at <http://www.sun.com>.

For more information on Solaris Management Console, refer to Chapter 3. For more information on Solstice DiskSuite, refer to Chapter 6.

The following table shows which version of Solaris Easy Access Server to use with the Solaris 2.6, Solaris 7, and Solaris 8 operating environments.

TABLE 7-1 Solaris Easy Access Server Interoperability/Compatibility Matrix

Application	Solaris 2.6	Solaris 7	Solaris 8
Solaris Easy Access Server	Version 3.0	Version 3.0	Version 3.0 Update

Setting Up and Running Solstice AdminSuite

▼ To Start Solstice AdminSuite

You can start Solstice AdminSuite through the Solaris Management Console interface, or you can start it directly from the command prompt. This section describes how to start Solstice AdminSuite directly.

- **To start the main Solstice AdminSuite application, type the following command:**

```
% /opt/SUNWseam/3_0/bin/admapp &
```

For more information on how to use this application with your Sun TCF system, refer to the Solstice AdminSuite 3.0 online help.

Optimizing Your Sun TCF System

This chapter includes the following topics:

- Optimizing File Server Performance on page 75
- Optimizing I/O Performance on page 77
- For More Information on page 78

Optimizing File Server Performance

The Sun TCF file server, which is attached to the disk storage array(s), is normally used only for file serving and not for running any applications. You can change particular system settings on the file server nodes to optimize their performance.

Note – Log in as superuser before you try to edit files in the `/etc` subdirectory. Otherwise, you will not be able to save any of your changes.

Editing the Cache Size Settings

Use your favorite text editor to modify the settings for the parameters in this section.

Note – The Solaris 2.6 preinstallation script automatically configures these settings for the Sun TCF system. If you have to reinstall the operating system on your compute farm, you will also have to reset these values manually.

Setting the Maximum Buffer Cache Size

In `/etc/system`, the parameter `bufhwm` sets the maximum size of the buffer cache. The buffer cache stores certain UFS (Unix File System) disk I/O operations. Normally, the buffer cache is automatically set to 2% of available memory; however, the nodes in the SunTCF have 4GB of memory, and 2% of this amount would be excessive.

- **To limit the amount of memory used for cache, include the following line in the `/etc/system` file on the file server:**

```
set bufhwm=10240
```

This limits the maximum size of the buffer cache to 10 Mbytes, which should be more than adequate.

Setting the Size of the Lookup Cache

The parameter `ncsize` sets the size of the Directory Name Lookup Cache (DNLC). This cache holds the name of files that have been accessed. A larger cache can significantly reduce the load on an NFS server with many clients.

- **To set `ncsize` to the largest possible value of 34,906, include the following line in the `/etc/system` file:**

```
set ncsiz=34906
```

Setting the Size of the Inode Cache

The parameter `ufs_ninode` sets the size of the cache that holds the disk files' inodes (an inode is a pointer to the information on the disk). Set this parameter to the maximum tested value ensures that more inodes are accessed quickly from memory, rather than having to be read from disk.

- **To set the value of `ufs_ninode`, add the following line to your `/etc/system` file:**

```
set ufs_ninode=34906
```

Increasing NFS Threads

In the file `/etc/init.d/nfs.server`, you can increase the number of NFS threads from 16 to 64 or 128. Changing the number of NFS threads increases the number of threads available on the file server to service I/O requests from the compute engines. This could increase I/O response, since having a larger number of threads, even if they are unused, does not usually adversely affect performance.

To increase the NFS threads, perform this procedure:

1. Use your favorite text editor to open the file `/etc/init.d/nfs.server`.

The file contains a section of code that appears as follows:

```
if [ $startnfsd -ne 0 ]; then
    /usr/lib/nfs/mountd
    /usr/lib/nfs/nfsd -a 16
elif [ ! -n "$_INIT_RUN_LEVEL" ]; then
    echo "NFS service was not started because" \
        "/etc/dfs/dfstab has no entries."
fi
```

2. Change the number 16 to 64 or 128.
3. Save your changes and exit the text editor.

Optimizing I/O Performance

The `cachefs` file system and its associated utilities are included with the Solaris operating environment. `cachefs` can help optimize file I/O performance on a compute engine for file systems that are remotely mounted from the file server. This feature works best when the files on the remote file system change infrequently, as compared to the number of times they are read.

When `cachefs` is implemented, each compute engine retains a local cache of files from the specified remote file system. When a cached file is accessed, `cachefs` checks the remote file to determine whether any changes have taken place. If so, it refreshes the local copy. If not, it simply uses the cached version. As a result, `cachefs` can help reduce the demand and response time of the file system and network.

Refer to the system administration guide for your version of Solaris for instructions on how to implement `cachefs`.

For More Information

For more information on tuning the file server for optimum performance, refer to the *NFS Tuning Guide* in the Solaris 2.6 or Solaris 8 Media Kit (depending on which version of the operating environment you are using). You can also find general information on performance tuning in general in these two books:

1. *Solaris Internals* by Richard McDougall and Jim Mauro (Palo Alto, CA: Sun Microsystems Press, 2001)
2. *Sun Performance and Tuning: Java and the Internet* by Adrian Cockcroft (Englewood Cliffs, NJ: Prentice Hall, 1998)

For tips on optimizing your Solstice DiskSuite configuration, refer to the *Solstice DiskSuite Reference Guide* and the *Solstice DiskSuite User Guide* (included in the Solaris Media Kit).

Troubleshooting

This chapter describes the following:

- Diagnostic Tools on page 79
- Hardware Concerns on page 81

Diagnostic Tools

This section describes some of the software-based diagnostic tools included with the Sun TCF system. For more information on the diagnostics tools and verifying the Sun TCF system hardware, refer to the *Sun Technical Compute Farm Installation and Service Manual*.

Using SunVTS Software

SunVTS (Validation Test Suite) tests and validates Sun hardware by verifying the connectivity and functionality of most hardware controllers, devices, and platforms. SunVTS 2.1.3 is preinstalled on each Sun TCF node with the Solaris 2.6 operating environment.

Note – If you are using Solaris 8 with your Sun TCF system, you must upgrade your SunVTS version to 4.1. SunVTS 4.1 is included with the Solaris 8 software. Refer to your Sun Enterprise 420R documentation for instructions on how to install and operate the software.

Use SunVTS to determine the cause of a problem with a hardware object identified in Sun Management Center. For example, if Sun Management Center detects an error or failure in one of the compute engines, you can use SunVTS to troubleshoot the problem with the compute engine.

For more information on SunVTS and its functions, refer to the SunVTS documentation.

▼ To Start SunVTS

You can start SunVTS from within the Solaris Management Console interface, or you can start it directly from the system prompt. This section shows how to start SunVTS from the system prompt.

1. **Type the following command at the system prompt:**

```
% /opt/SUNWvts/bin/sunvts
```

The SunVTS screen displays.

2. **Select which object you want SunVTS to test (for example, CPU, network, or SCSI), and then run the tests.**

Some tests run for only a few seconds. Other tests run for longer periods of time.

3. **To end a test, select Stop.**

▼ To Use the SunVTS Performance Meter

The Meter button in the tool bar displays an itemized graph of many system functions (for example: CPU usage, network packets, paging, available swap, interrupts). The graph displays ongoing activity for the selected object and allows you to track object and system performance over time.

1. **From the SunVTS GUI, specify which host you wish to monitor.**
2. **Select the objects you wish to monitor in that host.**

Examples of objects include CPU usage and available swap space on the disk arrays. For more information on the types of objects you can monitor with SunVTS, refer to the SunVTS documentation.

▼ To View the SunVTS Test Logs

After the tests complete, or you have stopped an ongoing test, you can view the detailed results.

- **To view the test results, select the Log button.**

The Log Files window is displayed. Click Print to print the test results, Delete to clear them from the window, or Close to close the window.

Using SunVTS Test Tracing

Test tracing traces all system calls that a SunVTS test makes, and then saves the resulting information in the form of `truss` output. Use the Trace Test option in the Commands menu to pick a test to trace. When the trace is finished, all the (`truss`-style) information from that test appears in the Test Messages box in the bottom of the main window. You can also write the test trace data to a specified file. Refer to the SunVTS documentation for more information.

Note – If you want to test serial ports, be sure to connect loopback connectors before launching SunVTS.

For more information on SunVTS, refer to the SunVTS documentation.

Hardware Concerns

This section discusses hardware-related problems with the Sun TCF system.

▼ To Boot From the Mirrored Disk

If you are unable to boot your Sun TCF system from the default boot disk due to a hardware problem, you can boot the system from the mirrored disk.

1. **Press Stop-A (the Stop key plus the A key).**

The `ok` prompt is displayed.

2. Type the following command:

```
ok boot disk1:a
```

The Sun TCF system then boots from the mirrored disk.

You need to order a replacement for your default boot disk from Sun. After the replacement is installed, be sure to run Solstice DiskSuite software. Use the Solstice DiskSuite command menu to resync the two disks (your mirrored disk and the replacement).

Index

A

- Asian locales
 - setting up on the compute engine, 29

B

- booting from a mirrored disk, 81

C

- cache size, 75
- cachefs, 77
- compute engine
 - optimizing performance, 77
 - running Sun Management Center agents, 44
- concat devices, 71
- concat disks, 19

D

- diagnostic tools, 79
- disk arrays
 - configuring, 18
 - formatting, 19
- disk formatting, 19
- disk mirrors, 70
- disk stripes, 71
- documentation, xiii
 - online, xiii

F

- file server
 - configuring the preinstalled software image, 14
 - optimizing performance, 75
 - software worksheet, 12
 - Solaris 2.6 software patches, 34

G

- grdadmin, 63
 - setting the password, 64

H

- hardware description, 3
- hot spare pool, 71
- hot spares, 20, 69
 - defined, 71

L

- log file, 42
- LSF, 9
 - integrating with Sun Management Center, 44

M

- metadevices, 70
 - defined, 70

- metatool, 70
- mirror partitions, 20
- mirrored disk
 - booting, 81
- mirroring, 19
- mirrors
 - booting from a mirrored disk, 81
 - defined, 70
 - number of disks, 20

N

- networking
 - setup worksheet, 12

O

- online documentation, xiii

P

- performance
 - monitoring, 59
- performance meter
 - in SunVTS, 80
- performance statistics, 59
- preinstalled software, 69
 - applications on the compute engine, 32
 - Solaris 2.6 patches, 34
- preinstalled software image, 14

Q

- qmon, 65
- qrsh, 66
- qsub, 66
- qtchsh, 66

R

- reinstalling
 - Solaris 2.6 parameters, 75
 - Solaris 2.6 software patches, 34

- related documentation, xiii
- replacement disks
 - setup, 82

S

- SCCS, 6
- servers
 - adding in Solaris Management Console, 41
- software patches
 - for Solaris 2.6, 31, 34
- Solaris
 - description, 6
 - reinstalling, 31
- Solaris 2.6
 - preset parameters, 75
 - software patches, 34
- Solaris 7, 32
- Solaris 8
 - installing, 14
 - installing instead of Solaris 2.6, 14, 33
 - Solstice DiskSuite considerations, 33
 - SunVTS version, 79
- Solaris Easy Access Server, 32
 - components, 69, 73
- Solaris Management Console
 - adding servers, 41
 - description, 7
 - log file, 42
- Solaris SunVTS, 79
- Solstice AdminSuite, 7
 - description, 73
- Solstice DiskSuite
 - configuring the disk arrays, 18
 - description, 69
 - metadevices, 70
 - setting up a replacement disk, 82
- striping, 19, 71
- Sun Grid Engine, 9
 - functions, 65
 - monitoring system performance, 59
 - preconfigured settings, 63
 - qmon, 65
 - qrsh, 66
 - qsub, 66
 - qtchsh, 66

- restarting, 67
- setting the administrator password, 64
- setting up the environment, 64
- stopping, 67
- Sun Management Center, 29
 - description, 7, 43
 - improving performance, 44
 - installing on system administration workstation, 44
 - using SunVTS to detect failures, 80
- Sun StorEdge arrays, 18
- Sun TCF
 - description, 1
 - diagnostic tools, 79
 - hardware description, 3
 - monitoring performance, 59
 - planning the networking setup, 15
 - product Web site, xiv, 2
 - related documentation, xiii
 - software setup worksheet, 12
- SunVTS, 79
 - detecting failures shown in Sun Management Center, 80
 - performance meter, 80
 - test logs, 81
 - test tracing, 81

T

- truss output, 81

V

- VERITAS Volume Manager, 9

W

- WMS, 59
- workload management software, 9, 59
 - integration with Sun Management Center, 44
 - purpose, 2
- worksheet
 - for software setup, 12

